

Oracle® Retail Advanced Science Cloud Services
Implementation Guide
Release 16.0.206
E93438-01

February 2018

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Primary Author: Judith Meskill

Primary Author: Mike Besse

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	xvii
Preface	xix
1 Introduction	
Oracle Retail Cloud Services and Business Agility	1-1
Oracle Retail Advanced Clustering Cloud Service	1-1
Reporting and Analysis	1-2
Oracle Retail Assortment and Space Optimization Cloud Service	1-2
Dynamic Creation of Space Clusters	1-3
Conduct Micro-Space Optimization What-if Analysis	1-3
Preview Results Leveraging Shelf Preview Capabilities	1-3
Oracle Retail CDT Science Cloud Service and DT Science Cloud Service	1-3
Customer Decision Trees	1-3
Demand Transference Science	1-3
Oracle Retail Attribute Extraction Cloud Service	1-4
Oracle Retail Market Basket Insight Cloud Service	1-4
Common Workflow	1-4
Interacting with Oracle Retail Advanced Science	1-4
Hardware and Software Requirements	1-6
2 Implementation Overview	
Implementation Process	2-1
Implementation Steps	2-1
Configure the Application Roles and Users	2-1
Data Load Overview	2-1
Edit and Load Common Oracle Retail Advanced Science Seed Data	2-2
Perform Attribute Preprocessing for CDT and DT, as Appropriate	2-3
Mandatory Configuration Parameters	2-3
3 Customer Decision Trees	
Input Data	3-1
Overview	3-1
Transactions Data Requirements	3-2
Attribute Data Requirements	3-2

Attributes with a Large Number of Values.....	3-3
Grouped vs. Raw Attribute Values	3-4
Attribute Splitting.....	3-4
Functional-Fit Attributes.....	3-4
Customer Segments	3-5
Location Hierarchy	3-5
Setting Up Categories.....	3-5
Calculating Customer Decision Trees.....	3-6
Setting the Top Attribute	3-6
Excluding Attributes from the Calculation.....	3-7
Handling of the Brand Attribute	3-7
Limitations of the CDT Calculation	3-7
Choosing the Time Interval	3-7
Understanding the Filter Settings.....	3-8
Segments vs. Location	3-8
Setting the Escalation Path.....	3-9
How the CDT Score is Calculated	3-9
Understanding CDT Pruning.....	3-9
Overriding the CDT Calculation.....	3-9
Using the Calculation Stage	3-10
Setup Stage.....	3-10
Data Filtering Stage.....	3-10
Calculation Stage.....	3-10
Advanced Use.....	3-12

4 Demand Transference

DT and CDT	4-1
Demand Transference Model	4-1
An Example.....	4-2
Historical Similarity Data.....	4-2
Historical Sales Data.....	4-3
The Role of Attributes in Calculating Similarities.....	4-4
Attribute Data Requirements	4-4
Guidelines on Number of Attributes and Attribute Values	4-5
The Effect on Similarity Values.....	4-6
Avoiding Attributes with Many Values	4-6
Functional-Fit Attributes.....	4-6
Customer Segments	4-7
Location Hierarchy	4-7
Setting Up Categories	4-8
Using Demand Transference	4-8
Seasonality in Historical Sales Data	4-8
Assortment Elasticity and the Cannibalization Factor	4-9
Calculating Assortment Elasticity	4-10
The Substitutable Demand Percentage	4-10
No Need for Time Interval.....	4-11
Segments vs. Locations	4-12

Setting the Escalation Path.....	4-12
Automatic Updating.....	4-12
5 Using Demand Transference	
Seasonality in Historical Sales Data.....	5-1
Assortment Elasticity.....	5-1
The Importance of Assortment Changes in Historical Data.....	5-2
The Meaning of the Possible Values of Assortment Elasticity.....	5-3
The Substitutable Demand Percentage.....	5-3
No Requirement for a Time Interval.....	5-4
Segments vs. Locations.....	5-4
Setting the Escalation Path.....	5-4
Automatic Updating.....	5-5
Avoiding Categories with Small Assortments.....	5-5
Implementing DT for Fashion Categories.....	5-5
Proper Level for Fashion Categories.....	5-6
Seasonality (Life Cycle) Considerations.....	5-6
6 Assortment Recommender	
Prerequisites.....	6-1
Producing Better Assortments.....	6-1
Run Groups and Run Frequency.....	6-2
The Run Group Parameters.....	6-2
Data Used by the Assortment Recommender.....	6-3
Halo Effects.....	6-4
Troubleshooting.....	6-5
7 Advanced Clustering	
Overview.....	7-1
Data Requirements.....	7-2
Multiple Hierarchies and Level Support.....	7-2
Clustering Criteria Supported in Store Clustering.....	7-3
Attributes in Store Clustering.....	7-4
Configuration Process.....	7-5
Multiple Clustering Approach.....	7-6
New Stores or Stores with Poor History.....	7-8
Outliers.....	7-8
Export to Excel.....	7-9
8 Assortment and Space Optimization	
Overview.....	8-1
ASO Data Input Requirements.....	8-3
Assortment Data.....	8-3
Planogram Data.....	8-4
Assortment-to-Planogram Mapping.....	8-5

Assortment to POG Mapping Process	8-5
Input Data	8-5
Automated Process	8-6
Mapping Errors	8-6
Replenishment Data.....	8-7
Optimization Science	8-7
Optimization Algorithm Overview	8-7
Sales and Inventory Model	8-8
Sales and Inventory Modeling Considers All Possibilities	8-8
Inventory Levels After Replenishment.....	8-9
Calculating the Facing Capacity for a Product/Fixture Combination.....	8-9
Maximum Capacity of a Product.....	8-10
Replenishment Parameters.....	8-10
Objective Function	8-11
Constraints	8-12
Product Family Group Constraints	8-12
Diagnosis of Visual Guidelines	8-13
Product Groups	8-17
Validation Tool (Sanity Checker).....	8-18
Diagnosis of Dropped Products.....	8-18
Checklist for Optimization Results Diagnosis	8-19
FAQs.....	8-20

9 Customer Segmentation

Overview	9-1
Data Requirements	9-1
Multiple Hierarchies and Support.....	9-2
Supported Segment Criteria.....	9-3
Customer Segmentation Attributes	9-3
Configuration Process	9-4
Attribute Preprocessing	9-5
Segmenting Approach.....	9-6
Customer Segment Store Profile Generation	9-6
Preprocessing	9-6

10 Configuration

User Interface Authentication and Authorization	10-1
User Management Configuration: Configuring Users and Roles	10-1
User Roles.....	10-2
Configuration	10-2
Generic Configuration.....	10-3
Advanced Clustering Configurations	10-4
Assortment and Space Optimization Configurations.....	10-6
Customer Decision Tree Configurations	10-12
Demand Transference Configurations.....	10-15
Returns Logistics Configurations	10-18
Market Basket Insight Configurations	10-19

Internationalization	10-21
Configuration Updates	10-21
Configuration Tables	10-22
Email Notification Configuration	10-24
11 Attribute Processing	
Attribute Preprocessing	11-1
Process Overview	11-2
Product Attribute Loading	11-3
Introduce New Attribute	11-4
Determine Attribute Source and Define in Oracle Retail Advanced Science Tables	11-4
W_PRODUCT_D or W_PRODUCT_ATTR_D	11-4
W_RTL_ITEM_GRP1_D or W_RTL_ITEM_GRP2_D	11-4
Populate RSE_PROD_ATTR_GRP_VALUE_STG Interface	11-5
Populate RSE_PROD_ATTR_VALUE_XREF_STG Interface	11-6
12 Market Basket Insights	
Overview	12-1
Data Requirements	12-1
MBA_ARM_SRVC_LOC_STG	12-2
Science Algorithms/Services	12-2
ARM_PH	12-2
ARM_PH_PROMO	12-2
ARM_PH_CS	12-3
Configurations	12-3
Data Output	12-4
13 Returns Logistics	
Overview	13-1
Data Inputs	13-1
Data Outputs	13-2
14 Data Interfaces	
Inbound Interface Requirements by Application	14-1
Advanced Clustering.....	14-1
Assortment and Space Optimization	14-3
Customer Decision Tree	14-4
Demand Transference.....	14-6
Customer Segmentation.....	14-7
Returns Logistics	14-9
Market Basket Insights	14-9
RASE Web Services.....	14-11
Authentication and Authorization	14-12
Summary of Web Services	14-12
Outbound Interfaces by Application	14-19

Advanced Clustering.....	14-19
Customer Decision Tree.....	14-19
Demand Transference.....	14-19
Assortment and Space Optimization.....	14-20
Customer Segmentation.....	14-20
Market Basket Insights.....	14-20
Interface Details by Logical Group.....	14-20
Retail Insights.....	14-21
W_DOMAIN_MEMBER_DS_TL.dat.....	14-21
W_EXCH_RATE_GS.dat.....	14-21
W_INT_ORG_DHS.dat.....	14-21
W_INT_ORG_DS.dat.....	14-22
W_INT_ORG_DS_TL.dat.....	14-22
W_MCAL_PERIOD_DS.dat.....	14-22
W_PARTY_PER_DS.dat.....	14-23
W_PRODUCT_ATTR_DS.dat.....	14-23
W_PRODUCT_DS.dat.....	14-23
W_PRODUCT_DS_TL.dat.....	14-24
W_PROD_CAT_DHS.dat.....	14-24
W_RTL_CMG_PRODUCT_MTX_DS.dat.....	14-24
W_RTL_CO_HEAD_DS.dat.....	14-25
W_RTL_CO_LINE_DS.dat.....	14-25
W_RTL_CUSTSEG_DS.dat.....	14-25
W_RTL_CUST_CUSTSEG_DS.dat.....	14-26
W_RTL_ITEM_DEL_TMP.dat.....	14-26
W_RTL_ITEM_GRP1_DS.dat.....	14-26
W_RTL_IT_LC_DEL_TMP.dat.....	14-26
W_RTL_PROMO_COMP_TYPE_DS.dat.....	14-27
W_RTL_PROMO_DS.dat.....	14-27
W_RTL_RECLASS_DP_GP_TMP.dat.....	14-27
W_RTL_RECLASS_IT_SC_CL_TMP.dat.....	14-27
W_RTL_SLS_TRX_IT_LC_DY_FS.dat.....	14-28
W_RTL_SLSFC_IT_LC_WK_F.dat.....	14-29
W_RTL_TRADE_AREA_DS.dat.....	14-29
W_RTL_TRADE_ATRA_LOC_MTX_DS.dat.....	14-29
RI Market Interfaces.....	14-29
Oracle Retail Advanced Science Common.....	14-30
rse_fake_cust_stg.txt.....	14-30
rse_like_loc_stg.txt.....	14-30
rse_prod_attr_grp_value_stg.txt.....	14-31
rse_prod_attr_value_xref_stg.txt.....	14-32
rse_sls_pr_lc_cs_wk_stg.txt.....	14-33
rse_sls_pr_lc_wk_stg.txt.....	14-34
Customer Decision Tree.....	14-35
attr.csv.dat (cdt_attribute_exp_vw).....	14-35
drtyattrvaltx.csv.ovr (cdt_attr_val_prod_xref_exp_vw).....	14-36
cdt_export.tar.gz.....	14-36

Demand Transference.....	14-37
drtyassrtelasv.csv.ovr (dt_ae_exp_vw).....	14-37
dt_assort_mult.csv (dt_assort_mult_exp_vw).....	14-37
ipdmdtfrpcti.csv (dt_assort_trans_exp_vw).....	14-38
drtyattrwgtv.csv.ovr (dt_attr_wgt_exp_vw).....	14-38
dt_new_items.csv (dt_new_items_exp_vw).....	14-39
dt_new_item_ros.csv (dt_new_item_ros_exp_vw).....	14-39
drtysiminv.csv.ovr (dt_sim_exp_vw).....	14-39
dt_loc_wk_excl_stg.txt.....	14-40
dt_mdl_prod_exp_stg.txt.....	14-40
dt_prod_loc_excl_stg.txt.....	14-41
Assortment and Space Optimization – Assortment Files.....	14-41
so_assortment_finalized_stg.txt.....	14-41
so_assortment_stg.txt.....	14-42
so_assort_cluster_member_stg.txt.....	14-43
so_assort_cluster_stg.txt.....	14-44
so_assort_phprod_attr_stg.txt.....	14-45
so_assort_phprod_finalized_stg.txt.....	14-45
so_assort_phprod_like_prod_stg.txt.....	14-46
so_assort_product_strcltr_stg.txt.....	14-47
so_assort_proloc_fcst_stg.txt.....	14-48
so_assort_proloc_priccost_stg.txt.....	14-49
Assortment and Space Optimization – POG Files.....	14-49
so_bay_fixture_shelf_stg.txt.....	14-49
so_bay_fixture_stg.txt.....	14-50
so_display_style_fixture_stg.txt.....	14-51
so_display_style_stg.txt.....	14-51
so_disp_style_orientation_stg.txt.....	14-52
so_fixture_disp_config_stg.txt.....	14-52
so_fixture_stg.txt.....	14-53
so_pegboard_disp_config_stg.txt.....	14-54
so_pog_bay_stg.txt.....	14-55
so_pog_display_style_stg.txt.....	14-55
so_pog_stg.txt.....	14-56
so_pog_store_cda_stg.txt.....	14-57
so_pog_store_stg.txt.....	14-57
so_prod_display_style_stg.txt.....	14-58
so_shelf_stg.txt.....	14-58
Assortment and Space Optimization – Miscellaneous Files.....	14-59
so_pog_assort_mapping_stg.txt.....	14-59
so_pog_assort_seas_mapping_stg.txt.....	14-60
so_prod_loc_repl_param_stg.txt.....	14-61
so_prod_stack_height_limit_stg.txt.....	14-62
Advanced Clustering.....	14-63
rsestrclst.csv (cis_cluster_set_exp_vw).....	14-63
cis_cluster_tmpl_stg.txt.....	14-63
cis_cluster_tmpl_prod_xref_stg.txt.....	14-66

cis_cluster_tmpl_lvl_attr_stg.txt	14-67
Returns Logistics	14-67
rl_price_elasticity_stg.txt	14-67
rl_price_ladder_stg.txt	14-68
rl_shipping_cost_stg.txt	14-69
Market Basket Insights	14-69
mba_arm_srvc_loc_stg.txt	14-69
mba_arm_run_exp.txt	14-70
mba_arm_result_exp.txt	14-71
Security	14-73
Web Services	14-73
SOAP	14-73
REST	14-74

15 Batch Processing

Overview	15-1
Custom Batch Requests	15-1
Managing Custom Batch Requests	15-1
Handling Data Files	15-2
Supported PROCESS_QUEUE Trigger Values	15-2
ORASE Incremental Exports	15-3
Batch Process Flow	15-3

16 Social Analytics

Data Inputs	16-1
--------------------------	------

17 Science Innovation Workbench

Components	17-1
Retailer Workspace Schema	17-2
Oracle APEX	17-2
Workspace	17-2
Users and Roles	17-2
Oracle Advanced Analytics	17-6
Oracle Data Mining	17-6
How to Invoke Oracle Data Mining	17-6
Oracle R	17-10
ORASE Restful Service	17-16
Troubleshooting	17-19
DBMS Scheduler	17-20
ORASE Schema Objects	17-20

Glossary of Acronyms

List of Figures

1-1	Interacting with Oracle Retail Advanced Science	1-5
7-1	Simple Clustering.....	7-7
7-2	Nested Attribute Clustering.....	7-7
7-3	Mixed Attribute Clustering	7-8
8-1	ASO Assortment Flow.....	8-2
8-2	ASO Workflow Overview.....	8-3
8-3	Blocking, Product Family Group, Same shelf, Sorting Illustrated.....	8-13
8-4	Primary as Vertical	8-14
8-5	Primary as Horizontal.....	8-15
8-6	Add Visual Guidelines.....	8-15
8-7	Export to Excel.....	8-16
8-8	Product Counts for Blocking Strategies.....	8-17
8-9	Dropped Products.....	8-19
8-10	Validation Errors.....	8-19
14-1	Manage Credential Stores.....	14-15
15-1	Batch Process Flow	15-4
17-1	Science Innovation workbench Key Components	17-2
17-2	Manage Users and Groups	17-3
17-3	Create User.....	17-3
17-4	Assign Group.....	17-3
17-5	Require Change of Password	17-4
17-6	SQL Workshop Object Browser: Reading Database Objects	17-4
17-7	SQL Workshop Create Object	17-4
17-8	SQL Workshop SQL Command: Executing Ad Hoc SQLs.....	17-5
17-9	SQL Workshop SQL Scripts: Uploading, Executing, and Running.....	17-5
17-10	Application Builder	17-5
17-11	Machine Learning Algorithms in Database	17-6
17-12	Oracle R	17-10
17-13	Identify an Application	17-12
17-14	Add Page.....	17-12
17-15	Add Report	17-13
17-16	Edit Page.....	17-13
17-17	Shared Components	17-14
17-18	Authentication Scheme	17-14
17-19	HTTP Header.....	17-14
17-20	Pre-Configured Scheme	17-15
17-21	Scheme Values.....	17-15
17-22	Application Window	17-16
17-23	RESTful Service	17-16
17-24	Create RESTful Service.....	17-17
17-25	Create Module	17-17
17-26	Resource Template.....	17-17
17-27	Resource Handler.....	17-18
17-28	Test	17-18
17-29	RESTful Refresh	17-18
17-30	Retailer Workspace.....	17-19
17-31	SQL Workspace.....	17-19
17-32	SQL Command.....	17-20

List of Tables

1-1	Interacting with Oracle Retail Advanced Science	1-6
2-1	Mandatory Common RSE Database Configuration Parameters.....	2-2
2-2	Mandatory CDT Configuration Parameters	2-3
2-3	Mandatory DT Configuration Parameters	2-4
2-4	Mandatory AC Configuration Parameters.....	2-5
2-5	Mandatory ASO Configuration Parameters	2-5
7-1	Data Requirements	7-2
7-2	Manual Overrides	7-5
7-3	Enable or Disable Cluster By	7-6
8-1	Vertical, Horizontal Strategy	8-14
8-2	Horizontal, Vertical Strategy	8-14
8-3	Product Group Constraints	8-17
9-1	Data Requirements	9-1
9-2	Manual Overrides	9-5
9-3	Enable or Disable Segment By	9-5
9-4	Attribute Configuration	9-6
9-5	Filter Step Conditions.....	9-7
9-6	Sample Step Conditions	9-7
10-1	ORASE Default Enterprise Roles.....	10-2
10-2	Oracle Retail Advanced Science Generic Configurations.....	10-3
10-3	AC Configurations.....	10-4
10-4	ASO Configurations	10-7
10-5	CDT Configurations	10-13
10-6	DT Configurations	10-15
10-7	Returns Logistics Configurations	10-18
10-8	MBI Configurations	10-19
10-9	Customer Decision Tree Configuration Tables	10-22
10-10	Advanced Clustering and Customer Segmentation Configuration Tables.....	10-22
10-11	Demand Transference Configuration Tables.....	10-23
10-12	ORASE Generic Configuration Tables.....	10-23
10-13	Assortment and Space Optimization Configuration Tables.....	10-24
11-1	RSE_PROD_ATTR_GRP_VALUE_STG	11-5
11-2	Flavor Attribute Values.....	11-6
11-3	RSE_PROD_ATTR_VALUE_XREF_STG.....	11-6
11-4	Adding a New Flavor Attribute	11-7
12-1	Data Elements.....	12-1
12-2	Implementation Suffixes	12-3
13-1	Data Inputs.....	13-1
13-2	Data Outputs	13-2
14-1	Advanced Clustering Inbound	14-1
14-2	Assortment and Space Optimization Inbound	14-3
14-3	Customer Decision Tree Inbound.....	14-5
14-4	Demand Transference Inbound	14-6
14-5	Customer Segmentation Inbound	14-7
14-6	Return Logistics Inbound	14-9
14-7	MBI Inbound Interface	14-9
14-8	Mapped Administrator Roles	14-12
14-9	Advanced Clustering Export Parameters.....	14-14
14-10	Customer Segment Export Parameters.....	14-14
14-11	Manage Credential Stores Information.....	14-15
14-12	Configuration	14-15
14-13	AIP/Replenishment Results Parameters.....	14-16
14-14	Output Aggregated Across Approved Runs Parameters	14-16

14-15	Assortment Results Details Parameters.....	14-16
14-16	Cross Reference Between Planograms and Finalized Assortments Parameters	14-17
14-17	POG Header Parameters.....	14-17
14-18	POG Equipment Components Parameters	14-17
14-19	POG/Stores Cross Reference Parameters	14-18
14-20	Finalized Assortment Product Hierarchies Parameters.....	14-18
14-21	Finalized Assortment Products Parameters.....	14-18
14-22	Product Display Style Information Parameters.....	14-18
14-23	Advanced Clustering Outbound	14-19
14-24	Customer Decision Tree Outbound	14-19
14-25	Demand Transference Outbound.....	14-19
14-26	Assortment and Space Optimization Outbound.....	14-20
14-27	Customer Segmentation Outbound	14-20
14-28	Market Basket Insights Outbound.....	14-20
14-29	W_DOMAIN_MEMBER_DS_TL.dat	14-21
14-30	W_EXCH_RATE_GS.dat.....	14-21
14-31	W_INT_ORG_DHS.dat	14-21
14-32	W_INT_ORG_DS.dat.....	14-22
14-33	W_INT_ORG_DS_TL.dat.....	14-22
14-34	W_MCAL_PERIOD_DS.dat	14-22
14-35	W_PARTY_PER_DS.dat.....	14-23
14-36	W_PRODUCT_ATTR_DS.dat	14-23
14-37	W_PRODUCT_DS.dat.....	14-23
14-38	W_PRODUCT_DS_TL.dat.....	14-24
14-39	W_PROD_CAT_DHS.dat	14-24
14-40	W_RTL_CMG_PRODUCT_MTX_DS.dat.....	14-24
14-41	W_RTL_CO_HEAD_DS.dat	14-25
14-42	W_RTL_CO_LINE_DS.dat	14-25
14-43	W_RTL_CUSTSEG_DS.dat.....	14-25
14-44	W_RTL_CUST_CUSTSEG_DS.dat	14-26
14-45	W_RTL_ITEM_DEL_TMP.dat	14-26
14-46	W_RTL_ITEM_GRP1_DS.dat.....	14-26
14-47	W_RTL_IT_LC_DEL_TMP.dat	14-26
14-48	W_RTL_PROMO_COMP_TYPE_DS.dat.....	14-27
14-49	W_RTL_PROMO_DS.dat.....	14-27
14-50	W_RTL_RECLASS_DP_GP_TMP.dat.....	14-27
14-51	W_RTL_RECLASS_IT_SC_CL_TMP.dat.....	14-27
14-52	W_RTL_SLS_TRX_IT_LC_DY_FS.dat.....	14-28
14-53	W_RTL_SLSFC_IT_LC_WK_F.dat	14-29
14-54	W_RTL_TRADE_AREA_DS.dat.....	14-29
14-55	W_RTL_TRADE_ATRA_LOC_MTX_DS.dat.....	14-29
14-56	RI Market Interfaces	14-29
14-57	rse_fake_cust_stg.txt.....	14-30
14-58	rse_like_loc_stg.txt.....	14-31
14-59	rse_prod_attr_grp_value_stg.txt.....	14-31
14-60	rse_prod_attr_value_xref_stg.txt.....	14-33
14-61	rse_sls_pr_lc_cs_wk_stg.txt.....	14-33
14-62	rse_sls_pr_lc_wk_stg.txt	14-34
14-63	attr.csv.dat (cdt_attribute_exp_vw)	14-35
14-64	drtyattrvaltx.csv.ovr (cdt_attr_val_prod_xref_exp_vw).....	14-36
14-65	drtyassrtelasv.csv.ovr (dt_ae_exp_vw).....	14-37
14-66	dt_assort_mult.csv (dt_assort_mult_exp_vw).....	14-37
14-67	ipdmdtfrpcti.csv (dt_assort_trans_exp_vw)	14-38
14-68	drtyattrwgtv.csv.ovr (dt_attr_wgt_exp_vw)	14-38
14-69	dt_new_items.csv (dt_new_items_exp_vw)	14-39

14-70	dt_new_item_ros.csv (dt_new_item_ros_exp_vw).....	14-39
14-71	drtysiminv.csv.ovr (dt_sim_exp_vw)	14-39
14-72	dt_loc_wk_excl_stg.txt	14-40
14-73	dt_md1_prod_exp_stg.txt.....	14-41
14-74	dt_prod_loc_excl_stg.txt.....	14-41
14-75	so_assortment_finalized_stg.txt.....	14-41
14-76	so_assortment_stg.txt	14-42
14-77	so_assort_cluster_member_stg.txt.....	14-44
14-78	so_assort_cluster_stg.txt	14-44
14-79	so_assort_phprod_attr_stg.txt.....	14-45
14-80	so_assort_phprod_finalized_stg.txt	14-46
14-81	so_assort_phprod_like_prod_stg.txt.....	14-47
14-82	so_assort_product_strcltr_stg.txt.....	14-47
14-83	so_assort_proloc_fcst_stg.txt.....	14-48
14-84	so_assort_proloc_priccost_stg.txt	14-49
14-85	so_bay_fixture_shelf_stg.txt	14-50
14-86	so_bay_fixture_stg.txt	14-50
14-87	so_display_style_fixture_stg.txt.....	14-51
14-88	so_display_style_stg.txt	14-51
14-89	so_disp_style_orientation_stg.txt	14-52
14-90	so_fixture_disp_config_stg.txt	14-53
14-91	so_fixture_stg.txt.....	14-53
14-92	so_pegboard_disp_config_stg.txt.....	14-54
14-93	so_pog_bay_stg.txt	14-55
14-94	so_pog_display_style_stg.txt	14-55
14-95	so_pog_stg.txt.....	14-56
14-96	so_pog_store_cda_stg.txt.....	14-57
14-97	so_pog_store_stg.txt	14-58
14-98	so_prod_display_style_stg.txt.....	14-58
14-99	so_shelf_stg.txt	14-59
14-100	so_pog_assort_mapping_stg.txt	14-59
14-101	so_pog_assort_seas_mapping_stg.txt.....	14-60
14-102	so_prod_loc_repl_param_stg.txt	14-61
14-103	so_prod_stack_height_limit_stg.txt	14-62
14-104	rsestrclst.csv (cis_cluster_set_exp_vw).....	14-63
14-105	cis_cluster_tmpl_stg.txt.....	14-64
14-106	cis_cluster_tmpl_prod_xref_stg.txt	14-66
14-107	cis_cluster_tmpl_lvl_attr_stg.txt.....	14-67
14-108	rl_price_elasticity_stg.txt	14-67
14-109	rl_price_ladder_stg.txt	14-68
14-110	rl_shipping_cost_stg.txt	14-69
14-111	mba_arm_srvc_loc_stg.txt	14-69
14-112	mba_arm_run_exp.txt	14-70
14-113	mba_arm_result_exp.txt	14-71
15-1	Trigger Values	15-2
15-2	Supported Zip Files	15-4
15-3	Handling Data Files	15-6
16-1	Data Input	16-1
17-1	ORASE Schema Objects	17-20

Send Us Your Comments

Oracle Retail Advanced Science Cloud Services Implementation Guide, Release 16.0.206

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

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

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

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

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

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

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

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

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

Preface

Oracle Retail Advanced Science Cloud Services Implementation Guide provides detailed information useful for implementing and configuring the application. It helps you to understand the behind-the-scenes processing of the application.

Audience

This document is for users and administrators. This includes merchandisers, buyers, business analysts, implementation team partners, the ORC, and administrative personnel.

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 Advanced Science Cloud Services documentation set:

- *Oracle Retail Advanced Science Cloud Services Administration Guide*
- *Oracle Retail Advanced Science Cloud Services User Guide*
- *Oracle Retail Advanced Science Cloud Services Release Notes*
- *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced 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.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Introduction

The Oracle Retail Advanced Science Cloud Services is comprised of the following Cloud Services:

- Oracle Retail Advanced Clustering Cloud Service
- Oracle Retail Assortment and Space Optimization Cloud Service
- Oracle Retail Customer Decision Tree Science and Demand Transference Science Cloud Service
- Oracle Retail Attribute Extraction Cloud Service
- Oracle Retail Market Basket Insight Cloud Service

These Cloud Services support the retail business processes of store cluster creation, optimization of item facings to available space, and insights about customer behavior patterns and product preferences. When incorporated within the end-to-end Assortment Planning and Optimization process, retailers can move beyond traditional planning processes and create customer centric and targeted assortments, improving customer satisfaction and overall business profitability.

Oracle Retail Cloud Services and Business Agility

The Oracle Retail Advanced Science Cloud Services are hosted in the Oracle Cloud with the security features inherent to Oracle technology and a robust data center classification, providing significant uptime. The Oracle Cloud team is responsible for installing, monitoring, patching, and upgrading retail software. Included in the service are continuous technical support, access to software feature enhancements, hardware upgrades, and disaster recovery. The Cloud Service model helps to free customer IT resources from the need to perform these tasks, giving retailers greater business agility to respond to changing technologies and to perform more value-added tasks focused on business processes and innovation.

Oracle Retail Software Cloud Service is acquired exclusively through a subscription service (SaaS) model. This shifts funding from a capital investment in software to an operational expense. Subscription-based pricing for retail applications offers flexibility and cost effectiveness.

Oracle Retail Advanced Clustering Cloud Service

Oracle Retail Advanced Clustering Cloud Service is an enterprise-specific clustering solution that leverages data mining capabilities to create store groupings at various product levels using multiple inputs. These inputs include performance data, product attributes, store attributes, third-party data such as demographic data as well as

consumer segments. Using embedded science and automation capabilities, retailers are able to identify patterns within available data to create the necessary customer-centric and targeted clusters to be used by downstream assortment planning, allocation/replenishment, pricing, and promotions planning processes.

The store clustering process enables the creation, review, and approval of store clusters for downstream solution use, while also providing the ability to define and use clustering templates that can be specific to given product/location combinations.

The Oracle Retail Advanced Clustering Cloud Service provides retailers with multiple clustering generation approaches and methods. These include the creation of simple, nested, and mixed attribute clusters using multiple methods, including those that support discrete and non-discrete attributes.

The types of clusters include the following:

- Performance-based clusters (Sales Revenue, Sales Units, Gross Profit%, and so on)
- Product attribute-based clusters (Brand, Color Family, Price Band, and so on)
- Location attribute-based clusters (Store Size, Climate, Population Size, and so on)
- Consumer profile-based clusters (Consumer Segment Profiles)

In addition to the above, users have the ability to create multiple clustering scenarios within a single cluster run. This enables the ability to leverage embedded rankings, scoring logic, as well as solution recommendations to define and approve the most appropriate clusters for use in intended planning or execution processes.

Reporting and Analysis

Users can access and review the following reporting information to drive decisions related to the clustering process.

Users can perform the following:

- Determine what categories or merchandise classifications benefit most from clustering; determine the level of product or location hierarchy at which to cluster; and determine what attributes should be leveraged.
- Analyze details related to the available cluster recommendations, assessing areas such as cluster composition, performance, and attributes, as well as store level scores (in relation to total clusters).
- Review cluster scenario comparison features, visually assessing differences between the respective store cluster details.

Oracle Retail Assortment and Space Optimization Cloud Service

The Oracle Retail Assortment and Space Optimization Cloud Service can help maximize return on space, sales, revenue, and profits while improving customer satisfaction by optimizing assortments and facings to available space.

Leveraging key inputs such as optimization goals, demand transference science, and visual guidelines as well as inventory and replenishment factors, retailers are presented with a recommended shelf/fixture layout that can be leveraged in downstream execution processes.

Dynamic Creation of Space Clusters

Leveraging available fixture data, the Oracle Retail Assortment and Space Optimization Cloud Service dynamically groups stores (known as space clusters) with common fixture dimensions, enabling retailers to optimize and refine their assortments at the planogram or store level.

Conduct Micro-Space Optimization What-if Analysis

The Oracle Retail Assortment and Space Optimization Cloud Service provides retailers with the ability to conduct 'what-if' analysis by adjusting fixture lengths during an optimization run. The solution allows for a visual review, comparison, and validation of the results. This provides the ability to dynamically manage and assess the impacts of adding or removing fixture space from a particular store (or store group). The solution can help plan for and conduct store projects by recommending the re-allocation of space to planograms with an optimal return on space.

Preview Results Leveraging Shelf Preview Capabilities

Prior to approving optimization results for downstream execution, retailers are able to review shelf previews, assessing variation from current or historical planograms as well as confirming that recommended results align with expectations. Updates to the respective shelf preview may be made in near real-time, with forecasted results updated in a real-time manner.

Oracle Retail CDT Science Cloud Service and DT Science Cloud Service

Customer Decision Trees

The Oracle Retail Customer Decision Tree Science and Demand Transference Science Cloud Service enables retailers to create customer segment-specific decision trees using available transaction level data. These customer decision trees are specific to their customer segments and the respective geographies they operate within, and retailers are provided a better understanding of their most important products and product attributes. Using this detailed information, the retailer is able to effectively analyze assortment coverage and identify the duplication of item types as well as prevent the removal of core items that would cause a loss of customers.

Demand Transference Science

Using the Oracle Retail Customer Decision Tree and Demand Transference Science Cloud Service, retailers can analyze a significant number of households (for example, in the thousands) to identify and rank which products are truly unique and whose sales are incremental, as opposed to those that can be discontinued because they are repetitive in nature and can be substituted with other products.

Understanding the incremental and substitutable sales associated with each item within an assortment, category managers can optimize the breadth of their assortments, as experienced by their customer's purchase preferences, with the optimal number of SKUs, given space constraints or financial goals.

Oracle Retail Attribute Extraction Cloud Service

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.

Oracle Retail Market Basket Insight Cloud Service

The Oracle Retail Market Basket Insight Cloud Service lets retailers review the analysis about their customer market baskets. The system calculates association rules from the provided sales transaction data, which provides insight into customer shopping patterns. The process examines sales transaction data and identifies associations between different types of products. Such information can help a retailer understand that promoting one product is sufficient to help drive sales of another product, given the sales associations they exhibit.

Common Workflow

The Oracle Retail Advanced Science solutions have a similar workflow and user interface (UI). The workflow lets users implement new science modules using similar techniques. For example, a retailer who uses the Demand Transference Science Cloud Service and the Customer Decision Tree Science Cloud Service may then be able to more easily learn and use Advanced Clustering and other aspects of demand modeling. This approach lowers the future total cost of implementing various science modules.

The *Oracle Retail Advanced Science Cloud Services User Guide* provides details about using these applications.

Interacting with Oracle Retail Advanced Science

Two connection channels are used for interaction with the Oracle Retail Advanced Science application:

Browser-Based

The application is accessed through a URL. The user is authenticated in order to gain access to the application. Access rights are controlled by the customer administrator through a Web application (Oracle Access Manager). Note that a role-based security policy is used. This allows the administrator to specify which applications and the tasks associated with those applications are accessible to which users.

Bulk Data Movement

A scheduled ETL extraction process must be used to extract the required data on the customer side and send it to the application through SFTP. Similarly, a schedule-based set of processes must be set up to process data coming in the opposite direction: from the application to on-premise. Note that this connection is still initiated from on-premise. Data is made available by the application for the download to on-premise location and processed further. All the necessary processes and credentials are set up during implementation.

Web Services

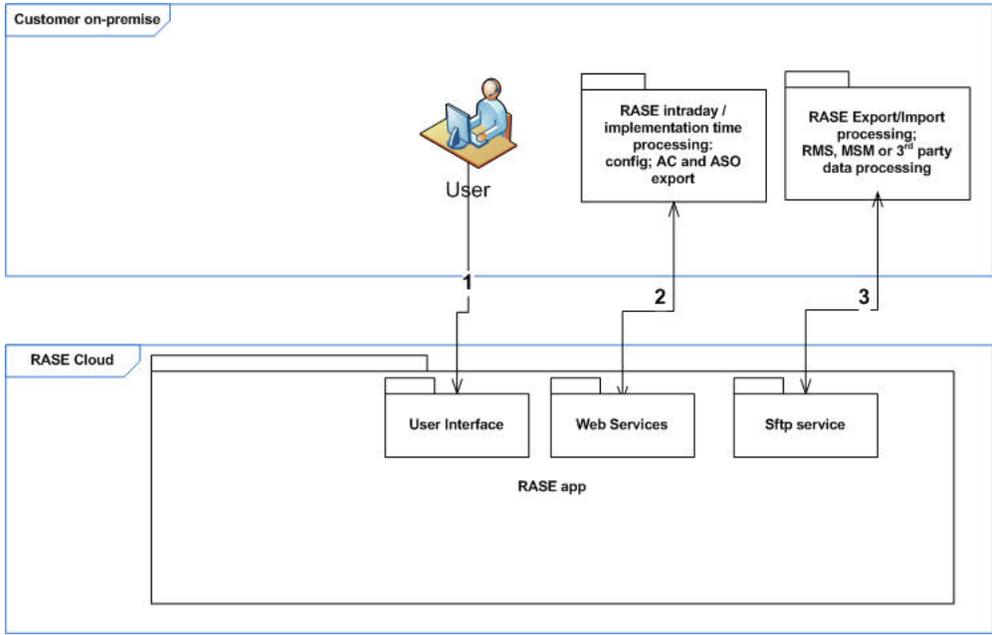
RASE web services are REST-based. RASE web services provide access to some of RASE application data and functionality but do not fully mirror the user interface or the export and import features of the backend. They are not a replacement for bulk data export, which must still be done at a scheduled time as part of batch processing. However, access to the configuration can be used during implementation and upgrade time, while AC and ASO export web services can serve as the means of obtaining incremental update data from a specified point in time (driven by a query parameter) as means of intra-day processing.

In the Cloud

The application processes are hooked up to a cloud scheduler to work in concert with what is sent (uploaded) from on-premise and what must be published to the outgoing SFTP directory for on-premise download.

The interactions with the application are illustrated in [Figure 1-1](#).

Figure 1-1 Interacting with Oracle Retail Advanced Science



Here are the major steps.

Table 1–1 Interacting with Oracle Retail Advanced Science

Step #	Protocol	Direction	External/ Internal	Description	Type of Data Sent
1	https/443	Inbound	External (Internet)	Used by customer to communicate with the application UI, OAM (login), OIM.	Cluster, DT, CDT, space optimized results parameters
2	sftp/2222	Inbound	External (Internet)	Data synchronization content (from RMS); Import/export files uploaded or downloaded by customer scripts (ODI for RMS-sourced data).	Includes hierarchies, calendar, sales transactions, and so on. Clusters generated, like entities, CDT xml, similarity files, space optimized results.
3	nfs/2049	Copy to/from local NFS mount	Internal	Copy files uploaded via SFTP to application server for processing with the data processing jobs. Copy exported result files to SFTP server for customer download.	Clusters generated, like entities, CDT xml, similarity files, and so on.
4	nfs/2049	Copy from local NFS mount	Internal	Pull files uploaded to SFTP server to application server for processing by Oracle Data Integrator (ODI) data load jobs.	Hierarchies, calendar, sales transactions, and so on.
5	SQLnet/1521	App->DB	Internal	DB create, update and delete operations from the application.	Cluster parameters, query parameters, and so on.

Hardware and Software Requirements

Oracle Retail Advanced Science has the following requirements:

Supported Browsers

- Mozilla Firefox Enterprise Version 45+
- Internet Explorer 11.0 (32-bit)
- Chrome (latest version)

Desktop Requirements

- Windows 7 or 8.1
- MS Excel 2003 or higher

Other Requirements

The user's source IP address must be communicated to the application cloud administration team for security purposes.

The SFTP client used for uploading and downloading data must be compatible with the SFTP protocol used by the application. Examples include:

- Putty Command line client
- Win SCP

- WS_FTP Pro Version 9

Note that all file exchange must be carried out in binary format.

Implementation Overview

This chapter provides an overview of the implementation of the Oracle Retail Advanced Science applications.

Implementation Process

This section provides details about the implementation process. It assumes that the underlying platforms have been properly implemented. This includes servers, Oracle database, RADM, and WebLogic application servers.

Implementation Steps

The order of steps provided here is designed to simplify the process. The advanced user may be able to change the process order or skip some steps; however, that is not recommended and not documented here.

Note: See [Chapter 10, "Configuration"](#) for details about application configurations that can be modified as part of a deployment.

Configure the Application Roles and Users

Before any user can log into any Oracle Retail Advanced Science application, you must set up application roles, add users, and assign users to the correct roles. To do this, complete the steps described in [Chapter 10](#).

Data Load Overview

Note: Prior to running any installed .ksh scripts, you must source the RSE Environment Setup file located here: \$RSE_HOME/common/scripts/lib/rse.env. To source this file, use the command

```
.$RSE_HOME/common/scripts/lib/rse.env
```

During an implementation of any modules, several steps are required. This section provides some details about this process.

The rse_config.ksh and the rse_master.ksh script are located in the \$RSE_HOME/common/scripts/bin directory. In addition, similar scripts are located within each of the application directories, for example, \$RSE_HOME/cdm/cis/scripts/bin has a cis_config.ksh and a cis_master.ksh script. All of the *config.ksh and *master.ksh

are similar in nature, so this section focuses on the `rse_config.ksh` and `rse_master.ksh` as examples. However, the concepts apply equally to the application-specific `*config.ksh` and `*master.ksh` scripts.

Edit and Load Common Oracle Retail Advanced Science Seed Data

All the applications share a set of configurable parameters that must be loaded into the RSE_CONFIG table. All have default values and are configurable and editable by the administrator. This section explains how to load and, if desired, edit these parameters.

The .ctl files for common configuration data must be edited and loaded into the staging tables. This data is common to all the applications. The application-specific .ctl files are located in their own application seed_data folders (for example, `orase\installer\orase16\so\db\seed_data`).

Review the .ctl files in that directory and adjust any configurations needed for the environment. Some configurations cannot be changed after the application has been used; therefore, you must carefully consider the parameters listed in [Table 2-1](#). The remainder are optional and default to reasonable valuables.

The following configuration parameters must be initialized during setup. The values for hierarchy level and type are recommended for any installation that integrates with the CMPO installation and must match for all installed applications.

Table 2-1 Mandatory Common RSE Database Configuration Parameters

Application	Parameter	Description	Value
RSE	CAL_PERIOD_LEVEL	This is the calendar hierarchy level that is used to drive RSE processing.	4
RSE	CHAIN_LEVEL_DESC	The description to use for any top level hierarchy element, when one must be manually created.	CHAIN
RSE	CMGRP_HIER_TYPE	The hierarchy ID to use for the CMPO. Recommendation is for a normal installation with CMPO.	1
RSE	CMGRP_LEVEL_ID	The hierarchy level ID that contains the level of the product hierarchy where the CMPO level exists (installation configuration). Recommendation is for a normal installation with CMPO.	5
RSE	MT_TZ	The time zone that is used by application server(s), that is, by the middle-tier. It must match <code>SELECT tzone FROM V\$TIMEZONE_NAMES</code> .	America/New_York
RSE	PRIMARY_LANGUAGE_CODE	The name of the language code to use for all RSE data sourced from RI.	EN
RSE	RA_FISCAL_CAL_ID	The ID of the calendar to use from RI (since RI supports multiple calendars).	1240
RSE	TRADE_AREA_HIER_TYPE	The hierarchy ID to use for the Trade Area (installation configuration).	6
RSE	UI_TZ	The time zone for display. It must match <code>SELECT tzone FROM V\$TIMEZONE_NAMES</code> .	America/New_York

Perform Attribute Preprocessing for CDT and DT, as Appropriate

Product attributes are required by CDT and DT and are stored in the RADM. Attribute preprocessing is independent of the Oracle Retail Advanced Science database and happens in RI or flat files generated by the user. Once these tables and files are correct, you can load the resulting attributes in the Oracle Retail Advanced Science schema as part of the data load process.

Here are the basic attribute pre-processing steps:

1. Populate RADM with raw attribute data.
2. Optionally, perform translation.
3. Parse.
4. Cleanse and standardize.
5. Categorize and label.
6. Define the attributes.
7. Bin and group.

For details on these steps, see [Chapter 11, "Attribute Processing"](#).

Mandatory Configuration Parameters

[Table 2–2](#) contains the mandatory configuration parameters for CDT.

Table 2–2 Mandatory CDT Configuration Parameters

Application	Parameter	Description	Value
CDT	CDT_CAL_HIER_TYPE	The hierarchy ID to use for the fiscal calendar (installation configuration).	11
CDT	CDT_CAL_LEVEL_ID	The hierarchy level ID that contains the level of the calendar hierarchy that CDT operates on. (This should equate to the Week - Installation configuration).	4
CDT	CDT_CMGRP_LEVEL_ID	The hierarchy level ID that contains the level of the product hierarchy that CDTs are created for (installation configuration).	5
CDT	CDT_CUSTSEG_HIER_TYPE	The hierarchy ID to use for customer segment (installation configuration).	4
CDT	CDT_CUSTSEG_LEVEL_ID	The hierarchy level ID that contains the level of the customer segment hierarchy that CDTs are created for (installation configuration).	2
CDT	CDT_LOC_HIER_TYPE	The hierarchy ID to use for location (installation configuration).	2

Table 2–2 (Cont.) Mandatory CDT Configuration Parameters

Application	Parameter	Description	Value
CDT	CDT_LOC_LEVEL_ID	The hierarchy level ID that contains the level of the location hierarchy that CDTs are created for (installation configuration).	Best equivalent for trade area level
CDT	CDT_PROD_HIER_TYPE	The hierarchy ID to use for the CMPO (installation configuration). The recommendation is for a normal installation with CMPO.	1
CDT	DEF_NUM_WEEKS_FOR_SIMILARITY	The default number of weeks of sales transaction data to be used by the similarity process. This is used when the user does not specify time intervals.	15

Table 2–3 contains the mandatory configuration parameters for DT.

Table 2–3 Mandatory DT Configuration Parameters

Application	Parameter	Description	Value
DT	AE_CALC_INT_LENGTH	The number of weeks to group together for in an interval for the AE calculation.	4
DT	CDT_SIMILARITY_AVAILABLE	Whether the CDT similarity has been made available to DT.	Y
DT	DT_CAL_HIER_TYPE	The hierarchy ID to use for the fiscal calendar.	11
DT	DT_CAL_LEVEL_ID	The hierarchy level ID that contains the level of the calendar hierarchy that DT operates on. (It should equate to week.)	4
DT	DT_CMGRP_LEVEL_ID	The hierarchy level ID that contains the level of the product hierarchy that DTs are created for.	5
DT	DT_LOC_HIER_TYPE	The hierarchy ID to use for location.	5
DT	DT_LOC_LEVEL_ID	The hierarchy level ID that contains the level of the location hierarchy that DTs are created for.	Best equivalent for trade area level
DT	DT_PROD_HIER_TYPE	The hierarchy ID to use for the CMPO.	1
DT	PR_LOC_STATUS_LAST_COMPLETED_WK	The last completed week for the SKU/Store ranging data copying.	Week ID from RSE_CAL_HIER
DT	WGT_CALC_INTERVAL_LENGTH	The number of weeks to group into an interval that is then used to perform weight calculations.	4

Table 2–4 contains the mandatory configuration parameters for AC.

Table 2–4 Mandatory AC Configuration Parameters

Application	Parameter	Description	Value
CIS	CIS_DFLT_CALENDAR_HIER_TYPE_ID	the default calendar hierarchy for clustering.	11
CIS	CIS_DFLT_LOCATION_HIER_TYPE_ID	The default location hierarchy for clustering.	2
CIS	CIS_DFLT_PRODUCT_HIER_TYPE_ID	The default product hierarchy for clustering.	1
CIS	PERF_CIS_APPROACH	The approach to use for performance based clustering. The available options are CDT and DT.	CDT

Note: There are no mandatory configuration parameters for MBA.

Table 2–5 Mandatory ASO Configuration Parameters

Application	Parameter	Description	Value
SO	SO_CAL_HIER_TYPE	The hierarchy ID to use for the calendar (installation configuration).	10
SO	SO_FISCAL_CAL_HIER_TYPE	The hierarchy ID to use for the fiscal calendar (installation configuration).	11
SO	SO_LOC_HIER_TYPE	The hierarchy ID to use for location (installation configuration).	2
SO	SO_PROD_HIER_LEVEL_FOR_LEAF_NODE	The product hierarchy level number for leaf node.	7
SO	SO_PROD_HIER_TYPE	The hierarchy ID to use for the product (installation configuration).	1

Customer Decision Trees

This chapter provides details about the use of the Customer Decision Tree Science Cloud Service application.

Input Data

This section describes setting up the data that the CDT application uses to calculate CDTs.

Overview

The calculation of CDTs is based on a retailer's historical data, especially customer-linked transactions data that includes subsets of transactions from the same customer. The CDT calculation does not require any data about the customer; it does require that the transactions are flagged to indicate that they came from the same customer.

The CDT calculation uses this customer-linked transactions data to determine, for a particular category at a particular store, the switching behavior of the customer among the items in the category at that store. By seeing what fraction of all historical customers of the category consider two specific items substitutable, CDT generates a similarity for the two items, which is a number between 0 and 1 that indicates how substitutable those two items are.

It is important to have data from a large numbers of customers shopping in the category in order to be more certain of the similarity values. In general, it is not recommended to perform CDT calculation for categories where customer-linked transactions data is available only for a few hundred customers.

The CDT calculation also relies on attributes, since attributes are at the nodes of the CDT. The CDT calculation applies the similarity calculation to attribute values as well as to items in order to find the similarities between attribute values. The CDT is then generated by examining the relationships between the attribute-value similarities and the item-level similarities. So good attribute information is also important.

Notice that the CDT calculation is all within a particular category, and thus the CDT models the customer's choice process only within a category. The CDT calculation generates separate CDTs, using separate calculations, for each category that the user chooses.

The CDT calculation does not filter out the effects of promotions or price changes, because these effects can cause customers to switch to a different item. This is valuable since the basis of the CDT calculation is examining switching behavior among the customers. Generally, more switching behavior in the historical data leads to a better CDT.

Transactions Data Requirements

The historical transactions data for the CDT calculation must meet the following requirements:

Customer Linked

Since the calculation involves examining switching behavior by customers, it is necessary to identify which transactions came from the same customer. This can be done using a loyalty card or a generated customer ID. No actual information about a specific customer is required; all that is needed is a way to identify which transactions come from the same customer. Note that it is possible to have customer IDs from a retailer where the customer ID is not that of an actual customer but rather a cashier loyalty card that was used for many different customers. These customer IDs, and their associated history, cannot be used for the CDT calculation, since the data comes from a large number of different customers. The data load for CDT automatically filters out such "fake customers."

Repeat Purchase

The category used for the calculation must be one where the typical customer performs several transactions per year. Examples include grocery items such as milk or yogurt, which are typically purchased weekly, and batteries, which are typically purchased several times per year. Item such as electronics are not appropriate, as such items may only be purchased every few years. Note that it can be possible to trade off purchase frequency and history length. It is also possible to trade off purchase frequency with the number of customers who shop in the category.

Attribute Data Requirements

The attribute values for the CDT calculation must meet the following requirements:

Set of Attributes

Each category is characterized by a unique set of attributes. These attributes differ from category to category. For example, for yogurt, the attributes might be size, flavor, brand, fat percentage, and pack size. For chocolate, the attributes might be size, brand, milk/dark, nut type, and package type. Two categories can both have brand, but that the brand attribute will have different values for each of the categories. So brand is actually a different attribute for each category.

Mapping

Each item in the category must be mapped to its set of attribute values. This information must be obtained from the retailer. Null values are acceptable as long as they are not too numerous. The CDT application can still run even if some attribute values are listed as null for some items in a category, but too many null values decrease the reliability of the generated CDTs.

Significance

The attributes for a category must be the ones that the customers actually pay attention to when shopping in the category. They are attributes that actually affect the customers' purchasing decisions.

The process of obtaining attributes for a category and performing a mapping of items in the category to attribute values is likely to require a significant amount of time and labor, even if the retailer has the information available, since this has to be done for every category.

Attributes with a Large Number of Values

A raw attribute is one that has a large number of attribute values. For example, the brand attribute for yogurt may be a list of 50 different brands at a large grocer. Using the raw attributes directly for the CDT calculation presents a problem, because each node of a CDT expands into a set of branches whose number is equal to the number of attribute values of the attribute at the node. An expansion into 50 different branches, one for each brand, is not desirable because the CDT would become too large to examine or interpret.

Such raw attributes must be turned into grouped attributes. This involves grouping the attribute values into a small number of bins. This grouping should be done in consultation with the retailer, who may have specific requirements. For example, a retailer may want to group soft-drink brands into Brand A, Brand B, and a third group that includes all other brands.

Another approach is to divide the values into two attributes (known as attribute splitting). For example, if the color attribute has many values, the single color attribute can be divided into two attributes, with one attribute representing the primary color and the second attribute representing a modifier. The CDT application's schema directly supports attribute grouping.

Attributes with a large number of values (for example, in the hundreds) can cause the CDT Calculation Stage to require a lot of time. Here are some approaches for handling categories that have attributes with a large number of values. The retailer should help determine which approach is appropriate.

Position of the Attribute Within the Tree

Typically, an attribute with a large number of values must not be at the top or near the top of the CDT. With such a large number of values, it is unlikely customers are first selecting from among such a large number of values and then selecting from other attributes with a smaller number of values. In addition, if such an attribute were at the top of the tree, the tree would be extremely wide and shallow. It would be extremely wide because the tree would then split into as many branches as there are attribute values. If there are 100 brands, and Brand was at the top of the tree, the tree would split into 100 different branches. As a result, the CDT would not be useful to the retailer. The tree would also be quite shallow; with 100 different branches, each branch would probably have very few SKUs, and so the branch could not be expanded much further.

In such a category, customers generally first use another attribute with a smaller number of values, and then choose an attribute with a large number of values. The following example demonstrates what happens when an attribute with a large number of values is lower in the tree.

In this example, the top attribute in the tree indicates the Market Segment, so that the SKUs in the category are split into various sub-categories. The Brand has a large number of attribute values. Because the Brand is below Market Segment, the branch for each segment only has a small subset of the Brands. Although Brand has many attribute values along each branch, only a subset apply, because each Brand only applies to one or two market segments. As a result, CDT never branches by all of the values in Brand, and only branch by a small subset of Brand.

It is possible for the CDT calculation to move Brand lower in the tree by itself, but in order to improve the performance of the CDT calculation in such a case, it can be helpful to direct the calculation to move the attribute with a large number of values lower in the tree.

There is no direct way in the CDT application to force an attribute to be lower, but here are some indirect strategies to use:

- In the Calculation Stage, set the Market Segment attribute to be a Top Attribute. This forces the Market Segment to be at the top of the tree, and so Brand will not be at the top of the tree. This can improve the performance of the Calculation Stage of the CDT application because it has fewer options to consider when expanding the tree.
- Set multiple attributes as the Top Attribute. It is possible that multiple attributes in combination delineate the market segment, for example. There could be a main segment and a sub-segment, as two separate attributes. In such a case, set both of them as Top Attribute. The CDT Calculation Stage will still determine the ordering of the Top Attributes.

For more information, see [Functional-Fit Attributes](#).

Grouped vs. Raw Attribute Values

CDT supports grouped attributes, which turns the raw attribute into one with many fewer values by grouping the attribute values into a small number of bins (for example, 3 to 5). CDT does not have an automated way of performing this grouping, so it is best if the grouping is done in consultation with the retailer, who may have specific requirements.

For example, a retailer may want the following grouping of soft-drink brands: {Soda A, Soda B, everything else}, because they are most interested in their two main brands and are willing to bin all other smaller brands together. The grouped-attribute approach is primarily useful for retailers who are not concerned with every specific value of an attribute that may have a large number of values.

A CDT with a grouped attribute will have branches only for the groups, and not for the original raw values within the groups. In the example above, the branches for Brand would consist only of Soda A, Soda B, Everything Else. This approach can provide additional insights into shopping behavior that would not be available with the raw-attribute approach. For example, if the Brands were grouped into three groups, each representing a particular price tier (say High, MainStream, and Budget), then the CDT can show the behavior of Mainstream customers vs. the behavior of Budget customers. The portion of the CDT underneath Mainstream would look different from the portion of the CDT underneath Budget. In essence, the portions are CDTs that show how each type of customer shops. This insight would not be available using only the raw Brand attribute.

Attribute Splitting

Another approach to handling attributes that have a large number of attribute values is to break them into two attributes (known as attribute splitting). For example, if the Color attribute has many values representing combinations of colors, it may be best to break the single color attribute into two attributes, with one attribute representing the primary color and the second attribute representing a modifier. However, this is an advanced technique, and grouping is recommended over attribute splitting.

Functional-Fit Attributes

A functional fit attribute is one where there is no substitution across the attribute's values. For example, batteries of different sizes cannot be substituted for one another. Any category where size determines the functional suitability of the item will have size as a functional-fit attribute. Information about which attributes are functional fit ones must be loaded into the CDT application.

Designating an attribute as functional fit can also be useful any time the attribute is unlikely to have substitution across it (for example, caffeinated vs. decaffeinated coffee). This is not exactly functional fit; however, substitution is unlikely, so it is better to mark the attribute as functional fit.

Functional-fit attributes always appear at the top of the CDT. The order of the functional-fit attributes will be some arbitrary order, but the order is not meaningful since there is no sense in which one functional fit attribute is more important than another. What the functional-fit attributes do is partition the set of items in a category into their own separate groups, each of which then has its own CDT.

This same effect can be achieved via the UI in the Calculation Stage, by using the Top Attribute functionality of the pop-up called Category Attributes Setup. Using the Top Attribute check box in this pop-up indicates to the Calculation Stage that the attribute should be put at the top of the tree.

Customer Segments

The CDT application can calculate CDTs by customer segment. Customer segments are set of groupings of the customer IDs that are used to identify the transactions. The retailer must provide the customer segment information as a data feed.

Customer segments are useful for retailers who believe that different customer segments shop differently. They may want CDTs by segment only for particular categories. The Calculation stage provides check boxes that control whether or not the run generates CDTs by segment. You can use these check boxes to calculate CDTs by segment for particular categories only.

Frequently, the groups will overlap, since a customer can belong to more than one segment. The CDT application functions normally in this case, since a separate CDT calculation is done per segment.

Location Hierarchy

The CDT application supports calculating CDTs by location hierarchy. The lowest level of the hierarchy should be above store; in general, calculating CDTs per store is not recommended. Per-store CDTs may have too little data to be reliable, and the calculation time involved can be quite long.

Some retailers may have stores that vastly differ in size and assortments. For example, a grocery chain may have both convenience stores and supermarkets. It may be necessary to arrange a separate calculation of CDTs for convenience stores vs. supermarkets, because people may shop differently at the two types of stores and the assortments may be different at the two types of stores.

One approach to this is to arrange a separate calculation by creating separate store clusters for convenience stores vs. supermarkets. The CDT application has the capability of calculating CDTs for each element of the location hierarchy, so it can calculate CDTs for the separate store clusters and thus produce separate CDTs for convenience stores vs. supermarkets.

Setting Up Categories

In general, a category is a set of items that are substitutable with each other (if there are no functional-fit attributes). The categories at a retailer can all be derived by choosing the correct level of the merchandise hierarchy at the retailer. The CDT application configuration supports choosing which level of the merchandise hierarchy is to be used as the category level.

A retailer may want categories that consist of unions of nodes of its merchandise hierarchy, because no level of its merchandise hierarchy suffices as the category level. The CDT application does support this, in that it allows defining an alternate merchandise hierarchy, where the categories can consist of arbitrary collections of items. However, before investing time in setting up an alternate hierarchy, make sure that it is necessary for meaningful CDT calculations.

Consider a category that consists of related though not substitutable products. For example, a single category of hair products can include both shampoo, conditioner, and items that are both shampoo and conditioner. There may be other hair-care related products in the category as well. A reasonable approach is to create an attribute called "Item Type" or "Market Segment" to indicate why the customer is buying it. The Market Segment attribute will segment the category into several sub-categories, and in the Market Segment, attribute should be set as a Top Attribute (see [Setting the Top Attribute](#)).

Calculating Customer Decision Trees

This section suggests ways to using the stages of the CDT application effectively.

Setting the Top Attribute

The Category Attribute Setup pop-up in the Calculation Stage contains check boxes that force particular attributes to the top of the tree. This is useful in several cases:

- The category has an attribute that has a large number of values (50 or more). See [Position of the Attribute Within the Tree](#).
- The category has a functional fit attribute. See [Functional-Fit Attributes](#).
- The category has an attribute that distinguishes market segments or product use. See [Setting Up Categories](#).
- You want to use the same top attribute that is present in a CDT from another source in order to make comparisons with that CDT easier.

It is possible to set more than one top attribute by checking multiple check boxes. In this case, all of the attributes will be at the top, but the Calculation Stage will determine the ordering of the attributes along each branch. This is useful if, for example, there are several attributes that together determine market segment.

In the case of using the top attribute as a market-segmenting attribute, it is possible to experiment with not using this attribute as the top attribute and letting the Calculation Stage determine the attribute ordering. This is useful if the market-segmenting attribute is not truly a functional fit attribute; that is, consumers can substitute across some of the market segments. For example, in the Cookie category, most likely customers can substitute across most of the market segments, because almost all cookies are desserts. In such a case, the Calculation Stage can give additional insight, by showing, for example, that customers actually shop by brand, so that even when they substitute across market segments they tend to stay with the same brand. This can be valuable information. However, if the retailer is interested only in substitutions within market segments, then it is proper to designate the market segment attribute as the Top Attribute.

However, in the case where the category has a very large number of items (greater than 1000), or the category has an attribute with a large number of values (50 or more), it is unwise to try such experiments, because the Calculation Stage may run too long. For such categories, setting the market-segmenting attribute as the top attribute is the best approach.

Excluding Attributes from the Calculation

The Category Attribute Setup pop-up in the Calculation Stage allows excluding particular attributes from the calculation. Use this to avoid meaningless attributes in the tree and also to decrease the calculation time of the Calculation Stage. Include only attributes in the tree that actually affect customers' purchasing. For example, Country of Origin may or may not be a relevant attribute, depending on whether it is visible on the package and plays a role in customers' decisions. Excluding such attributes will not only create a more useful CDT, but will also cut down on the execution time of the Calculation Stage

Handling of the Brand Attribute

Almost all categories will have a Brand attribute. The power of brands is well-known in retail, and in most categories, customers tend to stick with the same brand. Because of this, the Brand attribute will tend to show up near the top of the CDT. This is the correct scientific result, but not necessarily a useful one, for two reasons:

- It is known that customers shop by brand.
- Brand may have many attribute values, and the resulting tree will be shallow if Brand is high in the tree (see [Position of the Attribute Within the Tree](#)).

Here are some strategies for getting around these effects:

- Exclude Brand from the tree (see [Excluding Attributes from the Calculation](#)). The resulting tree will describe customer behavior in the other attributes. This indicates customer behavior, assuming that they shop by Brand. Given that they shop by brand, what are the effects of the other attributes on their purchasing behavior? CDT answers that question.
- Use the Top Attribute functionality to move Brand lower in the tree. See [Setting the Top Attribute](#).
- Group the brands, so that Brand becomes a grouped attribute. See [Grouped vs. Raw Attribute Values](#). This is a reasonable approach if taken in conjunction with the retailer, and can offer additional insight into shopping behavior not available without grouping. However, this approach is best taken as a phase 2 task, rather than immediately.

Limitations of the CDT Calculation

Because the CDT calculation uses historical data, the resulting CDT depends on the historical assortment represented in the data. If a particular attribute value does not have any representation in historical assortments among a particular group of stores, then the CDT for those stores will not have this attribute value in it. Similarly, if the assortments carried many more items of a particular attribute value compared to another attribute value, which limits the customer's choices, this can affect the CDT.

It is important to select historical data that reflects the retailer's current assortment, if the retailer has changed assortments in the last few years.

Choosing the Time Interval

The data used to calculate CDTs can be restricted to specific time intervals in the Data Setup stage. Thus, it is not necessary to use all of the available historical data to calculate CDTs. Some possible reasons for restricting the data to specific time intervals are:

The retailer may decide that particular time intervals, such as the two months before Christmas, represent periods where the buying behavior of its customers is

significantly different for certain categories. In this case, you can run the CDT application for just for these categories. Choose these categories in the UI, and then also choose the particular time intervals where the retailer believes shopping behavior is different.

If the retailer has changed business practices for certain categories, it is better to exclude the historical data from before the change, so that the CDTs reflect the retailer's current business practices and assortments, not the past ones.

One caution about selecting time intervals: there is always the danger of selecting too narrow a time interval, so that the amount of historical data in the interval is too little. See [Transactions Data Requirements](#).

In general, it is better not to restrict the data too much.

Understanding the Filter Settings

The Data Filtering stage of the CDT UI can be used to filter the data in order to remove outlier data that may result in incorrect CDTs. The user can adjust the values for the filters in order to control the extent of the filtering. The Data Filtering stage has five filters.

The three absolute filters have values that represent absolute limits that the data in question must pass in order not to be filtered out. For example, the maximum on missing attribute values states an absolute maximum that items must meet in order to be used in the CDT calculation. Items that have more than the maximum allowable missing attribute values will not be used in the CDT calculation.

The two relative filters have field values that are relative to the median of each category. The filters use median instead of the more-common average because the median is less likely to be affected by extreme outliers in the data. The average can be brought up (or down) by a single extreme outlier; this cannot happen with the median.

For example, the Transaction History Minimum is a percentage of the median transaction history length for a particular category. It is possible that the transaction history length varies by category. In generating a CDT for a particular category, the goal is to filter out customers who have transaction histories that are too short.

The default value of the filter for Transaction History Minimum is set to 1%, which filters out the customers that are truly outliers for the category because their history length is much smaller than median.

Segments vs. Location

Calculating CDTs by both segment and location hierarchy is not recommended. This calculation generates a large number of CDTs, since it will generate one CDT for each combination of location and segment, which takes a large amount of computation time. The large number of CDTs generated are not considered practically useful. You should either generate CDTs by segment, using Location Chain, or generate CDTs by location, using Segment Chain.

For a first calculation of CDTs, it is best to calculate them at Segment Chain/Location Chain. This quickly generates one CDT per category. It is a good way to check that everything has been done correctly and that the CDTs being produced are not unreasonable.

Setting the Escalation Path

The last stage in the CDT application involves setting the escalation path. If you are using only the segment hierarchy or only the location hierarchy, the escalation path is simply the hierarchy that you are using, and you set the escalation path according to the hierarchy. If you are using both a location hierarchy and a segment hierarchy, then usually you should set the escalation path to go up the segment hierarchy first, and then the location hierarchy. It is better to use only one of the hierarchies.

When using both hierarchies, the escalation path is necessary in order to tell the application which parent it should go to when moving up from a given segment/location node. With both hierarchies in play, every segment/location node has multiple higher-level nodes that do not lie along a single path. The escalation path is necessary to tell the application in what order the higher-level nodes should be considered. When only one hierarchy is used, the higher-level nodes form a single path.

How the CDT Score is Calculated

The terminal nodes of a CDT are the lowest-level nodes in the tree. The terminal nodes of the tree partition the items in the category. The items within each terminal node should be quite similar to each other, and less similar to the items in the other terminal nodes. The terminal nodes provide a clustering of the items in the category. A numerical score for the clustering given by the terminal nodes can be calculated.

Unconstrained clustering using any of the standard clustering algorithms using the similarities as the distance measure can also be created. This clustering can be compared with the clustering score for the clustering by terminal nodes. The terminal-node cluster score will be lower than the score for the unconstrained clustering because the unconstrained clustering had no constraints when performing the clustering. The closer the terminal-node clustering score is to the unconstrained score, the better the CDT. The CDT score in the CDT application is represented as a percentage of the unconstrained clustering score.

Typically, you should eliminate any CDT that has a score of below 60 percent, using the Pruning stage of the application.

Understanding CDT Pruning

The Evaluation Stage of the CDT application performs an operation called pruning, in which entire CDTs are removed. In the Evaluation Stage, the CDT as a whole is deemed reliable or not. An unreliable CDT is removed in its entirety; there is no automatic mechanism for making small adjustments to a CDT. The only mechanism the CDT application has for making small adjustments to a CDT is the manual editing of a CDT allowed in the CDT editor.

Overriding the CDT Calculation

It may be necessary, because of prior knowledge concerning the business of the retailer, or knowledge about the historical transactions at the retailer, to override portions of the calculation performed in the Calculation Stage of the CDT application. The override mechanism there allows you to specify what the topmost attributes of the CDT should be. For example, from an understanding of the retailer's business, it may be clear that in a particular category, brand should be at the top level of the tree. The override mechanism allows you to specify brand as the top level of the tree. The override mechanism is also flexible enough to allow specifying only the top level of the tree, while the rest of the tree is filled in by the usual calculation.

While it is possible to obtain the same effect by manual editing of the CDT, manual editing is much slower, especially if you have generated multiple CDTs for each category.

Using the Calculation Stage

This section provides step-by-step instructions for setting up the Calculation Stage, with a few comments on using the other stages. The focus here is mainly on the Calculation Stage, because the settings in this stage can directly affect how the CDTs look and because the Calculation Stage generally takes up most of the execution time.

If you are just beginning to use the CDT application, experiment with smaller categories (fewer than 1,000 items) initially. Smaller categories are easier to work with because they take less execution time in the Calculation Stage than larger ones, so it is easier to do multiple runs and examine results.

Setup Stage

When first starting to use the CDT application, it is best to set up only one category at a time in the Setup Stage. In this way, each run is for one single category. It requires some experience to include multiple categories in the same run, and it is not recommended as a starting point. The instructions assume only one category has been set up in this run.

Before selecting a category for a CDT run, review the data requirements in [Transactions Data Requirements](#) to be sure that your desired category meets the data requirements.

Data Filtering Stage

This stage is usually straightforward, in that the default values of the fields are usually suitable. However, it is important to check the Data Filtering Summary at the bottom of the screen *after* the stage has completed running. You must click **Refresh** in the summary table in order to see the results related to the latest run. Check each filter in the summary to see how much data it filtered out. If too much data was filter out, then determine whether the data may have a problem, or whether you need to adjust the filter so that it is less stringent.

If this is the first time you have run the Data Filtering Stage on a particular category, then you should run only the Setup Stage and the Data Filtering stage on the category, without running the Calculation Stage. This allows you to check the Data Filtering Summary before spending time running the Calculation Stage. Once you have run the Data Filtering Stage on a category and have checked the Data Filtering Summary, then you can re-run the Data Filtering Stage on the same category without checking the Summary, unless you have loaded new data for the category.

Calculation Stage

The steps here are simplified to help you get started in properly using the Calculation Stage. Once you become familiar with using this process, you can alter and expand them to use more of the capabilities of the CDT application. The process presented here represents the minimal set of steps needed to produce CDTs and to get you going in the right direction.

Take care during this process so that the Calculation Stage can complete within 1 or 2 hours for categories that have more than 1,000 items or that have an attribute with more than 50 values. The steps detail any additional consideration needed for large

categories. After performing a run with these steps, if the time for the Calculation Stage to run turns out to be acceptable, then these restrictions can be relaxed on subsequent runs of the category.

Each step may reference sections of this chapter that can provide further details.

1. Check both top level check boxes (one for Segments and one for Location). With these settings, the Calculation Stage will generate only one CDT, representing the CDT for all customers and all locations. This is the recommended way to start using the Calculation Stage. In particular, these settings are recommended for very large categories, where the calculation time for multiple CDTs may be quite prohibitive and not worth the investment until you have generated one CDT.
2. Exclude any unnecessary attributes (see [Excluding Attributes from the Calculation](#)). It is good practice to exclude unnecessary attributes, but it is even more so when working with large categories in order to avoid unnecessary computation. For large categories, consider excluding attributes that you know are less important to the retailer, even if they may have an effect on customers' purchasing in the category.
3. Handle Brand properly. Brand frequently has many attribute values, and handling it properly is especially important when the category also has a large number of items (1,000 or more). You can skip this step if the category has fewer than 1,000 items. See [Handling of the Brand Attribute](#).
4. Set Top Attributes properly. In particular, if the category has some type of market segment, product type, or product usage attributes, then force these attributes to be at the top of the tree. If the category has a large number of items, then it is likely to require some of these attributes, because with that many items, the items will likely have different segments or types. It is unlikely that the entire set of items is completely interchangeable in the customers' mind, and so it is proper to put segmenting attributes at the top of the tree. In addition to being scientifically proper, this decreases the execution time of the Calculation Stage because there are fewer combinations for the stage to consider and because other attributes such as Brand with a very large number of values will be moved down the tree, where fewer items are involved in the calculation. For more information, see [Setting the Top Attribute](#).
5. For large categories, consider setting the SKU Percentage of the termination condition to a lower value (possibly 0%). A value of X in this field in the Calculation Stage UI specifies that a branch will end when a node on the branch contains fewer than X% of the items in the category. If $X = 5\%$, which is the default, and the category contains 2,000 items, then the branch ends when a node on the branch contains fewer than 100 items. A threshold of 100 is probably too high, and if left at 5%, various branches may not be expanded to their full extent. If the value is expanded until a node on the branch contains fewer than 10 items, the SKU percentage field must be set to $10 / 2,000 = 0.5\%$. However, the field only accepts integer percentages, and so it must be set to 0%, which will let the Calculation Stage use other internal criteria to end a branch. This field can also be used in a reverse manner; that is, by setting a higher value the tree will become shallower and the calculation time will be reduced. For trial runs, you may wish to leave it at 5%, and see how far the branches are expanded before trying a run with a setting of 0%. For smaller categories, with fewer than 1,000 items, using the default of 5% is likely to be reasonable and no adjustment is needed. The use of percentage in this field can also be handy if you are performing runs that have more than one category. If the categories are related, so that you want trees of roughly the same depth for them, the percentage nature of the field will help produce this result.

Advanced Use

The process described in [Calculation Stage](#) is intended as a starting point, and is the shortest path to getting one CDT per category. Once that has been achieved, it is possible to consider some more advanced uses of the CDT application. Here are some suggestions for these more advanced uses. These are not steps to be performed but individual suggestions.

- Generate CDTs that are specific to location or segment. The Calculation Stage can generate CDTs that are segment specific or location specific, by unchecking the appropriate Top level check boxes in the stage. This makes it possible to see whether purchasing behavior differs by segment or by location. Note that on large categories with more than 1,000 items, such a calculation can take many hours, because the stage must calculate one CDT per segment or per location. Note also that it is *not* advisable to uncheck both check boxes, as that will produce a CDT for each segment/location combination. This is a large number of CDTs that will take a very long time to run.
- Set up grouped attributes. See [Grouped vs. Raw Attribute Values](#).
- Experiment with setting different attributes as the top attribute, or with not setting a top attribute at all. See [Setting the Top Attribute](#). Different settings here can produce different insights. However, keep in mind the points raised in [Handling of the Brand Attribute](#).

Demand Transference

This chapter provides details about the use of the Demand Transference Science Cloud Service application.

DT and CDT

The DT and CDT applications differ in significant ways. The CDT application has more stringent requirements for data than the DT application. CDT requires customer-linked, frequent transactions. Many retailers in various areas of retail do not have this type of data readily available. DT only requires SKU-store-week sales-units aggregates.

Demand Transference Model

A mathematical model of how the transference happens is required in order to calculate the transfer of demand in response to assortment changes. It is essential to understand the model at a basic level in order to best use DT. DT generates parameters that go into the model, so an understanding of the model can help when using the DT parameters.

The model is known as a cannibalization model. In this type of model, each item in an assortment has an associated value called its "full demand," which is the demand the item would have if it were the only item in the assortment. The full demand of an item is then multiplied by a factor, called the "cannibalization factor," which has a value of 1 if there are no other items in the assortment, but becomes progressively less than 1 as more and more items are added to the assortment. As the assortment becomes larger, the demand for each of the items decreases from its full demand because of cannibalization. The reverse is also true. If items are removed from the assortment, then the cannibalization factors increase, representing demand transferred from the removed item to the items remaining in the assortment. The cannibalization factors decrease from a value of 1 when the assortment becomes larger, and increase (up to a limit of 1) when the assortment becomes smaller.

The degree of change in an item's cannibalization factor indicates how similar the added items are. Item A's cannibalization factor will decrease more for added items that are very similar to A. The similarity of items is a key input to the Demand Transference model.

The cannibalization factor of an item accounts for similarities and also for a quantity called "assortment elasticity." The assortment elasticity determines how much of a decrease in the cannibalization factor occurs due to the addition of items of a particular similarity. The assortment elasticity is a number that depends on the particular category for which demand transference is being calculated. In one category,

adding item B to the assortment may cause item A's cannibalization factor to go from 0.7 to 0.6, whereas in another category, adding an item Y may cause item X's cannibalization factor to go from 0.7 to 0.5, even though the similarity of X and Y is the same as the similarity of A and B. In other words, similarities alone are not enough to calculate cannibalization factors. The assortment elasticity is necessary to tell us, for each category, how much change in cannibalization factors will occur for items of a given similarity.

The two components of the cannibalization factor, the similarities and the assortment elasticity, are calculated by DT from historical data. (The similarities can also be imported instead of calculated.) DT then exports the similarities and the assortment elasticities to any applications that want to calculate demand transference. It is up to the consuming application to properly use the cannibalization model in conjunction with the exported similarities and assortment elasticities to calculate transfers of demand when assortments change.

Note that demand transference only occurs within a category. All calculations are based on items cannibalizing each other, and there are no complimentary (halo) effects. DT calculates assortment elasticities at a level always higher than item. A single category/segment/location combination receives just one assortment elasticity.

An Example

This simple example explains how applications such as Category Management Planning and Optimization (CMPO) use the demand transference model to generate forecasts after assortment changes. In an assortment in the Cookies subcategory, one cookie SKU is removed from the assortment. The cannibalization factors of the rest of the items increase, because each is now cannibalized less after the removal of the one SKU. Because the cannibalization factors increase, the model predicts an increase in sales in accordance with the increase in the cannibalization factors. The removal of the SKU caused these increases, and some of the SKU's demand has transferred to the other SKUs.

Historical Similarity Data

DT has two different options for obtaining similarities. It can calculate them by itself or it can import them from CDT.

If the retailer has not implemented CDT for a category, then of course only the first option is possible.

The second option is recommended for a category only if the retailer has implemented CDT and run it for the category, since only in that case are similarities for the category available from CDT.

If CDT similarities are available, the recommendation is that you use them, instead of having the Similarity Calculation stage calculate its own similarities. The similarities from CDT are generally preferable to the attribute-based similarities that DT can calculate on its own because CDT similarities do not rely on attributes. They are extracted purely from historical transactions data.

The transactions data held in the RADM schema is used to feed both CDT and to generate the SKU-store-week aggregates for DT, so in option 2, consistency between the similarities and the SKU-store-week aggregates is automatic.

The similarities obtained from CDT may not cover all of the SKUs that are currently in the historical data loaded for DT. For example, it is possible that since the CDT similarities were calculated, the retailer has added some new SKUs to some

assortments. This situation requires no special handling, because DT can extend the CDT similarities to cover the added SKUs. This extension does require attribute values for the new SKUs.

Note that the CDT similarities for a category may be only at Segment-Chain/Location-Chain. In this case, there is only one set of similarities for the category, or they may exist at various levels of the location hierarchy or the segment hierarchy, depending on what options the user selected in the CDT application's Calculation stage.

Historical Sales Data

DT requires SKU-store-week sales-units aggregates. The data loader for DT automatically produces the needed SKU-store-week sales-units aggregates from transactions data that is held in the RADM schema, so it is not necessary to implement a separate loader for SKU-store-week aggregates.

Typically, the data cannot be aggregated to a higher location level than store because different stores usually have different assortments. Some atypical cases can occur in which aggregation across some stores is legitimate because the assortments are the same or nearly so, but this is generally not the case.

In addition to the SKU-store-week sales-units aggregates, DT also requires promotion data. A flag indicates which SKU-store-weeks contain major promotions (ones that caused a very large increase in sales units, such as three times normal). DT uses the promotion data to flatten the promotional spikes in the SKU-store-week sales-units aggregates. DT uses the flattened data called the baseline to calculate assortment elasticity.

DT calculates the assortment elasticity by examining the historical assortment changes and seeing their effect on base sales rates of the items remaining in the assortment. Promotional spikes can affect this calculation by obscuring the true effect on base sales rates. These promotional spikes are removed to decrease the sales rates back to their base rates.

For example, suppose the historical data for a store S indicates that the Cookies assortment has one fewer SKU in week 10 compared to week 1. That is, a cookie SKU was removed. To see where the demand from this SKU transferred to in week 10, the sales units of the remaining SKUs between week 1 and week 10 are compared. This comparison is made across many pairs of weeks (though not all possible pairs of weeks). A promotion in week 10 of particular items can interfere with the analysis of the changes in demand that were due to the assortment changes.

Note that, in CDT, the effect of promotions is left in because it is an external influence that helps cause switching behavior in customers. However, in DT, promotions can affect the calculation of demand transference in the case where one item in the category is promoted and another is not, which is why the promotions are flattened.

An alternative to flattening the promotions is to ignore the SKU-store-weeks where promotions occurred. However, to implement this, it is necessary to ignore all SKU-store-weeks in any week where a SKU-store was promoted, because it is not clear what the effect is on the other items when one SKU is promoted and the others are not. Removing that many SKU-store-weeks can leave little data remaining, especially since many retailers promote quite frequently. For this reason, it is better to flatten promotions and keep more data instead.

Similarly to the way in which promotions are handled, the calculation of the baseline also involves removing short-term downward spikes that are due to very short-term stock outs. Note that inventory information is not an input to the DT system, and so

the algorithm finds large but short-lasting dips in weekly sales units and fills those in. This handling of out-of-stock is not related to long-term out-of-stock conditions, discussed in a later section.

The Role of Attributes in Calculating Similarities

Without customer-linked transactions data, DT must use the attribute values of the SKUs to calculate similarities. The similarity of two SKUs is based in part on how many attribute values they have in common (the more in common, the higher the similarity of the two SKUs). The attributes used in the calculation are the raw attributes, not the grouped attributes that CDT uses. So it is not necessary to group the attribute values for DT.

Because the attributes play such an important role in calculating similarities, attribute quality is important when DT performs the calculation.

Attributes are also used in performing any necessary extensions of the CDT similarities to cover new SKUs.

Note that the similarities calculated by DT are only at Segment-Chain/Location-Chain. In contrast, the CDT similarities can be at multiple levels.

Attribute Data Requirements

The attribute values for the DT calculation must meet the following requirements:

Set of Attributes

Each category is characterized by a unique set of attributes. These attributes differ from category to category. For example, for yogurt, the attributes might be size, flavor, brand, fat percentage, and pack size. For chocolate, the attributes might be size, brand, milk/dark, nut type, and package type. Two categories can both have brand, but the brand attribute will have different values for each of the categories. So brand is actually a different attribute for each category.

Mapping

Each item in the category must be mapped to its set of attribute values. This information must be obtained from the retailer. Null values are acceptable as long as they are not too numerous. DT can still run even if some attribute values are listed as null for some items in a category, but too many null values decrease the reliability of the generated DTs. In particular, too many SKU pairs may come out as less similar than they should be, which would decrease demand transference between those pairs (which leads to an underestimate of demand transference in applications such as CMPO).

Null values have a particular use in accommodating categories that are actually a union of more than one category. See "[Setting Up Categories](#)" for more information.

Significance

The attributes for a category must be the ones that the customers actually pay attention to when shopping in the category. They are attributes that actually affect the customers' purchasing decisions.

Note that the similarity calculation will still complete even with attributes that do not affect customer behavior, but the similarities produced will be less distinguishing. For example, a category has a Supplier attribute, which indicates for a given product which supplier shipped an item to the grocer. This attribute may be important to the

grocer for accurate bookkeeping, but it has no effect on the customer's purchasing behavior because it is not reflected in the item itself nor is it something that the customer is concerned about. However, if it is included when setting up attributes, then the effect would be to increase the similarity of items that were from the same supplier. This is a false similarity, since it does not reflect how the customer actually views these items. In particular, if the supplier is a duplicate of the Brand attribute, then the similarity of products within the same Brand would be unintentionally increased.

The process of obtaining attributes for a category and performing a mapping of items in the category to attribute values is likely to require a significant amount of time and labor, even if the retailer has the information available, since this must be done for every category.

Guidelines on Number of Attributes and Attribute Values

The number of attributes and attribute values must be enough to distinguish the SKUs within a category. That is, for a given set of attribute values, the number of SKUs in the category all having those values must be a small number. A maximum of seven SKUs is recommended. For example, the Cookie category at a grocer has only three attributes, Brand, Package Size, and Organic. If Brand has seven values, Package Size has three values, and Organic has two values (either Yes or No), then the total number of combinations of attribute values is $7 \times 3 \times 2 = 42$. For 600 different cookie SKUs, the average for each combination of attribute values will represent $600 / 42 =$ approximately 14 different SKUs. The distribution of SKUs among the 42 different sets of values will not be an even 14, as some sets of values will have much more than 14, while others will have less. The three attributes alone are not enough to provide enough distinguishing power among the cookie SKUs. If Flavor is an important determinant of customer purchases, it should be added to the Cookie category. The guideline of a maximum of seven indicates that additional attributes are necessary. It is worth examining those sets of attribute values that have the largest number of SKUs associated with them in order to see what attributes can be added to reduce the number of SKUs.

It is not just the number of attributes that is important, but how many values each attribute has. For example, if Brand had 100 values instead of seven values, then the total number of attribute-value combinations is $100 \times 3 \times 2 = 600$. It might seem that an easy way to achieve the maximum of seven is to expand the number of values in the attributes. However, this results in each SKU being similar to only a small number of other SKUs. For a single attribute for cookies with 600 different values, it might then be possible to assign one value to each cookie SKU, separating all of the 600 cookies SKUs with a single attribute. However, this would make each cookie SKU completely dissimilar (similarity of 0) from all other cookie SKUs, and the result would be no transference between the SKUs. Putting all 600 SKUs each into a separate Brand causes a complete loss of any similarity information among the SKUs, and no transference will result. For an opposite example, consider 11 attributes, each with only two values. There are a total of $2^{11} = 2048$ combinations of values, so that may be enough to encode 2,000 SKUs, even though there are only $2 \times 11 = 22$ distinct attribute values over the 11 attributes. In general, having more attributes is better, and it is better to increase the number of attributes rather than increase the number of attribute values of a single attribute. However, this is not always possible, and it is better to have the attribute with many values than not have the attribute at all. Flavor, for example, can have many values, as can Color. See "[Avoiding Attributes with Many Values](#)" regarding attributes that have many values.

The more SKUs in the category, the more attributes and attribute values will be needed to achieve the maximum of 7.

The Effect on Similarity Values

Suppose the set A of SKUs consists of 22 SKUs, all with the same attribute values, and the set B of SKUs consists of 25 SKUs, all with the same attribute values (but a different set of attribute values from set A). If the set A consists of cookie SKUs all with a package size of Small, and set B consists of cookie SKUs with the same attribute values as A except the size is Medium, then every SKU in A has a similarity of 1 to every other SKU in A, and every SKUs in A is similar to at least 22 other SKUs. Every SKU in A is similar to every SKU in B, since they only differ in one attribute value (namely size). So a SKU in set A is similar to at least $21 + 25 = 46$ SKUs, which means that if a SKU in A were deleted from an assortment, its demand would have significant transference to about 46 other SKUs, assuming all 46 remained in the assortment. It is possible that a SKU in A being similar to 46 other SKUs in fact represents reality, but if it does not, then using additional attributes that distinguish the SKUs in A and in B will reduce the number of similar SKUs.

Avoiding Attributes with Many Values

Attributes with a large number of values occur frequently. For example, a color attribute in any clothing category might have several shades of each color. Midnight blue, sea blue, and sky blue may all be separate attribute values of the Color attribute; the problem is that in the similarity calculation, a midnight blue item and a sea blue item would be considered completely dissimilar colors, because the two color attribute values are different; in reality, because they are both shades of blue, they should be somewhat similar. One solution is to split the color attribute into two separate attributes, a primary color attribute and a modifier. In this example, the primary color would be blue and the modifiers midnight, sea, and sky.

Functional-Fit Attributes

A functional fit attribute is one where there is no substitution across the attribute's values. For example, batteries of different sizes cannot be substituted for one another. Any category where size determines the functional suitability of the item will have size as a functional-fit attribute.

Information about which attributes are functional fit ones must be loaded into DT. The information is used to perform the similarity extension process of CDT similarities and to correctly calculate attribute-based similarities.

In either case, the functional-fit attributes are used to set the similarity of two SKUs to be 0 if the SKUs differ in any functional-fit attribute. Without the functional-fit information, the two SKUs may have non-zero similarity, and there would be erroneous demand transference between the two SKUs, such as batteries of different sizes.

Designating an attribute as functional fit can also be useful any time the attribute is unlikely to have substitution across it (for example, caffeinated vs. decaffeinated coffee). This is not exactly functional fit; however, substitution is unlikely, so it is better to mark the attribute as functional fit.

One approach to avoiding having to define large numbers of attributes and attribute values is to use functional-fit attributes. This approach does not help achieve the maximum of 7, but it can help decrease the number of SKUs that are similar to a given SKU. For example, with the sets A and B of cookie SKUs, if size were designated as functional fit, then the similarity between SKUs in A and SKUs in B would become 0. However, that designating size as functional fit does nothing about the 22 SKUs in A that all have a similarity of 1 to each other, since their attribute values are all the same. (Similar comments apply to set B.)

If the attributes and attribute values are insufficient to reach the maximum of 7 SKUs per set of attribute values, functional-fit attributes can be used to decrease the number of SKUs to which transference occurs. This is a second-best approach, and it is better to design a proper set of attributes and attribute values, in order to:

- Achieve the maximum of 7 SKUs
- Provide transference between SKUs that should have transference. Using functional-fit attributes reduces transference, but it may reduce it too much and remove transference from pairs of SKUs that should have transference. For example, in the sets A and B, the similarity between a SKUs in A and a SKU in B becomes 0, which does not reflect reality since the SKUs in A and in B share common attribute values except for size.
- Keep the second-best approach as a last resort, in case time is insufficient for designing a good set of attributes for a category.

Customer Segments

DT can calculate assortment elasticities by customer segment. This involves dividing the customer IDs into groups (the groups do not have to be disjoint). Retailers who want to use segments must, as with CDT, create the necessary groupings of customer IDs. DT uses the segments to produce segment-SKU-store-week aggregates of sales units, instead of just SKU-store-week aggregates. The segment-SKU-store-week aggregates are produced by aggregating transactions data, just as with the SKU-store-week aggregates. The difference is that the aggregation is by segment.

There is always a Segment-Chain for the segment hierarchy, and so there is always a segment that contains all customers. The Segment-Chain level of segment-SKU-store-week aggregates is not necessarily the sum of the lower-level segment-SKU-store-week aggregates, because it is possible that the segments are not disjoint (meaning a customer can belong to more than one segment). The Segment-Chain-level aggregates are produced by a separate aggregation of transactions data instead of by aggregating lower-level aggregates.

Using segments allows DT to calculate separate assortment elasticities for each segment. This means that demand transference can differ by segment.

Note that when using customer segments, references in this document to "SKU-store-week" data should be read as "segment-SKU-store-week" data. For example, the SKU-store-week sales-units aggregates mentioned above become segment-SKU-store-week sales-units aggregates.

Location Hierarchy

DT supports calculating assortment elasticities by location hierarchy. The lowest level of the hierarchy should be above store; in general, assortment elasticities should not be calculated per store. Per-store assortment elasticities may have too little data to be reliable. The calculation time involved can be quite large to handle all stores individually. The calculation of assortment elasticities depends on having assortment changes in the historical data, and the store level may contain too few assortment changes to produce reliable assortment elasticities.

Some retailers may have stores that differ in size and assortments. For example, a grocery chain may have both convenience stores and supermarkets. It may be necessary to arrange a separate calculation of CDTs for convenience stores vs. supermarkets, because people may shop differently at the two types of stores and the assortments may be different at the two types of stores.

One approach to this is to arrange a separate calculation by creating separate store clusters for convenience stores vs. supermarkets. DT has the capability of calculating CDTs for each element of the location hierarchy, so it can calculate CDTs for the separate store clusters and thus produce separate CDTs for convenience stores vs. supermarkets.

Setting Up Categories

In general, a category is a set of items that are substitutable with each other (if there are no functional-fit attributes). The categories at a retailer can all be derived by choosing the correct level of the merchandise hierarchy at the retailer. The DT configuration supports choosing which level of the merchandise hierarchy is to be used as the category level.

Demand transference can only occur within the category, since the categories define the sets of items that cannibalize each other.

A retailer may want categories that consist of unions of nodes of its merchandise hierarchy because no level of its merchandise hierarchy suffices as the category level. DT does support this, in that it allows defining an alternate merchandise hierarchy, where the categories can consist of arbitrary collections of items. However, before investing time in setting up an alternate hierarchy, make sure that it is necessary for meaningful DT calculations.

For example, it is possible that the set of all yogurt SKUs at a retailer is not at any level of the merchandise hierarchy. The retailer may have the category Dairy Products, which is too large because it contains yogurt and milk, and the retailer might have the category Store-brand Yogurt, which is too small because it leaves out the yogurt SKUs that are not store brand. In such a case, it may be necessary to set up an alternate hierarchy so that all the yogurts can be put together in their own set. On the other hand, if a level of the existing merchandise hierarchy contains most of the yogurt SKUs, but not quite all of them, an alternate hierarchy may not be worth the effort.

Frequently, retailers will have categories that are actually unions of categories. For example, a retailer might have a Hair Care category that contains shampoo, conditioner, and hair oil. The retailer may not want to separate out this category into three separate categories of Shampoo, Conditioner, and Hair Oil, if, for example, a single person in the organization is responsible for all three. The problem is that these three types of products do not share a common set of attributes. The attributes describing Hair Oil are not the same ones needed for describing Shampoo or Conditioner. The types of products may share common attributes, such as Scent, but each type of product also needs its own set of attributes. The solution is to define, in addition to the common attributes, a set of attributes for each product type. If an attribute applies only to Shampoo, and not to Conditioner or Hair Oil, then Conditioner SKUs and Hair Oil SKUs should have Null for the value of that attribute. This is a common use of null attribute values, and makes it possible to handle the case of a category that is really the union of smaller sub-categories.

Using Demand Transference

This section suggests ways to using the stages of DT effectively.

Seasonality in Historical Sales Data

DT assumes that within a category, all of the items at a store have a common seasonality. This assumption is correct for categories in which each item does not have a predetermined point of obsolescence or in which the point of obsolescence is years

after the item was first introduced. Most grocery categories or basic clothing items meet this assumption. Electronics items frequently have defined life cycles that are measured in years. The situation where care would be needed is a category where, within the same store, the items have differing life cycles and the life cycles are short, so that at a given moment, in this category, the store may have items that are at various points in their life cycles. This is the situation where the common seasonality assumption is invalid. This situation commonly occurs with fashion merchandise.

Assortment Elasticity and the Cannibalization Factor

An assortment elasticity of 0 turns all cannibalization factors into constant 1, meaning the assortment has no cannibalization. This is unlikely. However, it does show that a small-magnitude value for assortment elasticity indicates a category where cannibalization is small. Similarly, a high magnitude of assortment elasticity indicates a category where cannibalization is large. It is possible for the magnitude to be too large.

It is also possible for the Calculation stage of DT to produce assortment elasticities that are positive. Such positive values for assortment elasticity are an indication that there is some unidentified problem with the data, because a positive assortment elasticity means cannibalization factors increase with increasing assortment size, which in turn means each item in the assortment sells more the larger the assortment gets. In the Evaluation stage of DT, such positive assortment elasticities are removed and replaced by assortment elasticities from escalation.

A simple example for understanding cannibalization factors involves adding identical, or nearly identical, items to an assortment. (In practice a retailer would never do this, but it is useful as an example.) With only one of these items in the assortment, it takes the entire market share for the item. If another item that is so similar as to be almost identical is added, the two items split the market share evenly between them, half to each item. The cannibalization factors are now half for both items. If a third such item is added, a 3-way even split is created, one-third for each and the cannibalization factors are all one-third. This pattern continues as more items are added; the cannibalization factors all slowly approach 0 (but never reach 0). As an aside, this example also happens to show how adding items to an assortment does not necessarily produce more market share overall for the assortment, since the new item may siphon off sales of existing items.

In this example, the cannibalization factors were all equal, but in a real example they likely would all be different.

The cannibalization factor is actually a power-law, meaning the assortment elasticity enters into the cannibalization factor as an exponent. The cannibalization factor consists of a positive value, called the Total Assortment Effect (TAE), raised to the assortment elasticity. Each item in an assortment has its own TAE; the TAE increases as items are added to the assortment. Therefore, the assortment elasticity is a negative number in order for the cannibalization factor to decrease as TAE increases. (In the above example, the TAE could be the count of the number of items added so far, and the assortment elasticity would then be -1, thus producing one-half, one-third, and so on.)

Note the similarity to the more-conventional idea of power-law price elasticity, which involves a price raised to a negative power (the negative power being the price elasticity). In the cannibalization model, the TAE plays the role of price.

The cannibalization factor also accounts for the similarity of the items being added to the assortment, so that similar items cannibalize each other more than non-similar

ones. The similarity values are used to calculate the TAE; higher similarities produce a larger TAE, providing a larger decrease in cannibalization factors.

The cannibalization factor depends on both the similarity values and the assortment elasticity. It may seem that similarity alone determines cannibalization (as a similarity of 0.5 between items A and B means that A takes half of B's share if A is added), but it is not that simple. By separating out the concepts of TAE and assortment elasticity, the model is more robust; if all of the similarity values are biased lower or higher for some reason, the bias can be accounted for by adjusting the magnitude of the assortment elasticity so that the cannibalization factors are still correct.

Calculating Assortment Elasticity

In order to calculate assortment elasticity, DT requires historical data that contains assortment changes, because DT uses historical data to determine how much cannibalization occurred when historical TAEs changed. From the relationship between changes in historical TAEs and changes in cannibalization, DT then calculates the assortment elasticity. This is similar to more conventional calculations of price elasticity. In order to determine price elasticity from historical data, it is necessary to have price changes in the historical data, and the more changes the better.

For example, suppose the historical data for a store S shows that the Cookies assortment has one fewer SKU in week 10 compared to week 1. That is, a cookie SKU has been removed. The TAEs for the remaining cookie SKUs will all decrease between week 1 and week 10 because of the removal of the one SKU. DT then examines the changes in historical sales units of the SKUs in the cookie assortment at S between week 1 and week 10. By relating the changes in the sales units to the changes in TAEs, DT can calculate the assortment elasticity. The calculation will produce an elasticity of large magnitude if the changes in TAE caused a large increase in sales units; a small-magnitude elasticity will result if the increases are moderate.

In practice, such historical comparisons are always more complex than in this example. It is rare to find a pair of weeks where the only assortment change was the removal of a single SKU. Typically, in each pair of weeks, there are many assortment changes, involving both additions and removals, and the changes in TAE are a result of all of those changes. In the end, though, the relationship between the changes in TAE and the changes in sales units is summarized in a single number, the assortment elasticity, across all pairs of weeks. Because this single number summarizes the vast number of pairs of weeks and SKUs where TAEs changed, it is an average over all the pairs of weeks and SKUs in the historical data and is not tuned to any particular SKU.

Consider if CMPO is used to remove a single SKU from an assortment. It is likely that no pair of weeks in sales history exists where exactly this SKU was removed and only this SKU was removed. For forecasting the results of this removal, CMPO is making an extrapolation from the historical analysis described above, and using the assortment elasticity that is not tuned to this particular situation of removing only this one particular SKU.

The Substitutable Demand Percentage

The substitutable demand percentage of an item in an assortment is the fraction of its demand that is retained by the assortment if the item is removed from the assortment. It is a measure of how substitutable the item is. For example, if the substitutable percentage is 100 percent, then removing the item will not decrease the total sales units of the assortment, since all of the demand for the item will transfer to the other items that remain in the assortment. On the other hand, if it is 50 percent, then removal of the item from the assortment means 50 percent of its demand is lost, and 50 percent is

retained. The total assortment sales units will decrease if this item were to be removed from the assortment.

The magnitude of the assortment elasticity has an influence on the substitutable percentage. Increasing the magnitude of the assortment elasticity increases the substitutable percentage. DT only calculates assortment elasticity for the entire category (not per item), so changing the value of assortment elasticity changes the substitutable percentage for all items in the category all at once.

It is possible for the magnitude of assortment elasticity to be too large, and the indication of this is that the substitutable percentage for several of the items in the assortment is over 100 percent. It is acceptable for a few items to have substitutable percentages over 100 percent because those are probably outliers. If the assortment is large, having a few such outliers is almost a certainty. However, having 10 percent of items in the assortment over 100 percent requires attention.

DT provides a tool for examining the substitutable percentage, and also for decreasing the assortment elasticity if too many items have a substitutable percentage over 100 percent. The following are some guidelines on how to use this tool.

Selecting Time Interval

Select a time interval that is likely to contain assortments that are representative of the retailer's current assortments. Since the retailer is going to be using the assortment elasticity to forecast what happens when modifying current assortments, it makes sense to test the assortment elasticity against assortments that are as similar as possible to the current ones.

Adjusting Assortment Elasticity

Adjust the assortment elasticity by setting a maximum substitution percentage. DT then calculates an assortment elasticity that results in substitution percentages that do not exceed the set maximum. When using this feature, you may want to set the maximum to something higher than 100 percent if there are some outlier items that have high substitution percentages. Forcing these outliers down to 100 percent may result in a very small-magnitude assortment elasticity, which may mean unacceptably small substitution percentages for all except the outlier items. So you may want to select a maximum that is higher than 100 percent but that still brings most items down to 100 percent, leaving a few outliers above 100 percent.

Setting Maximum Percentage

Set the maximum percentage even if all substitution percentages are already below 100 percent. You may know that a particular category should exhibit a substitution percentage of at most 70 percent. In such a case, the tool can be used to bring the substitution percentages down to 70 percent.

No Need for Time Interval

The Data Setup stage in CDT can be used to set up a time interval for the CDT calculation. The Data Setup stage of DT has no equivalent.

The cannibalization factor is actually a power-law, meaning that the assortment elasticity enters into the cannibalization factor as an exponent. The cannibalization factor consists of a positive value, called the Total Assortment Effect (TAE), raised to the assortment elasticity. Each item in an assortment has its own TAE; the TAE increases as items are added to the assortment. Therefore, the assortment elasticity is a negative number, in order for the cannibalization factor to decrease as TAE increases.

The cannibalization factor already directly incorporates information about the assortment through the TAE, and thus the cannibalization model can handle fairly large assortment changes. This makes it less necessary to use a time interval for DT, compared to CDT, because historical assortment changes can be directly accounted for in the model as changes in TAE.

Segments vs. Locations

In the Calculation stage for DT, as with CDT, it is possible to set up the calculation so that it is performed at all combinations of levels of the segment hierarchy and the location hierarchy. However, the recommendation is to use only one of the two hierarchies in the Calculation stage. Set either the segment hierarchy or the location hierarchy (or both) to be Chain. Because the calculation of assortment elasticity requires assortment changes in history, generating assortment elasticities at all levels may mean that at lower levels, the data does not contain enough assortment changes in history. If the assortment changes are infrequent, you should only calculate a Segment- Chain/Location Chain assortment elasticity.

Setting the Escalation Path

The last stage in DT involves setting the escalation path. If you are using only the segment hierarchy or only the location hierarchy, the escalation path is simply the hierarchy that you are using, and you set the escalation path according to the hierarchy. If you are using both a location hierarchy and a segment hierarchy, then usually you should set the escalation path to go up the segment hierarchy first, and then the location hierarchy. It is better to use only one of the hierarchies.

When using both hierarchies, the escalation path is necessary in order to tell the application to which parent it should go to when moving up from a given segment/location node. With both hierarchies in play, every segment/location node has multiple higher-level nodes that do not lie along a single path. The escalation path is necessary to tell the application in what order the higher-level nodes should be considered. When only one hierarchy is used, the higher-level nodes form a single path.

Automatic Updating

DT can automatically and periodically update the assortment elasticities as new sales history comes in. New assortment elasticities can be loaded into the consuming applications and immediately used. When new historical transactions enter the RADM schema, DT automatically aggregates them and produces new SKU-store-week sales-units aggregates. These new aggregates are then appended to the older SKU-store-week aggregates, and the resulting data set is then used in a new calculation of assortment elasticities. Note the following about the calculation:

- It does not re-run all of the stages. It just calculates assortment elasticity.
- It only updates assortment elasticities, not the similarities from the Similarity Calculation stage.
- It uses a mix of old data and more recent data. As a result, the values of the assortment elasticities change slowly over time as the data set becomes more tilted towards newer data.
- The assortment elasticities that were overridden using the Substitutable Percentage tool stay overridden, and are not updated.

Using Demand Transference

This chapter provides details about using Demand Transference.

Seasonality in Historical Sales Data

The DT application assumes that, within a category, all of the items at a store have a common seasonality. This assumption is generally correct for long life cycle categories, where each item does not have a predetermined point of obsolescence or where the point of obsolescence is years from the point of introduction of the item. Examples include most grocery categories and basic clothing items. Electronics items frequently have defined life cycles, generally measured in years.

It is important to address the situation in which different items in the same store have different life cycles and those life cycles are short. In this situation, the store may have items that are at various points in their life cycles and there is no common seasonality. This frequently occurs with fashion merchandise (see "[Implementing DT for Fashion Categories](#)").

Assortment Elasticity

Here is an example that explains the cannibalization model. In this example, all the cannibalization factors are equal; in a real example, the factors would all be different.

Identical, or nearly identical, items are being added to an assortment. If only one of these items is added to the assortment, it takes the entire market share for the item. When another item that is extremely similar is added, the two items split the market share evenly between them. The cannibalization factors are now half for both items. The addition of a third such item creates a three-way even split, one-third for each and the cannibalization factors are each one-third. As more such items are added, the cannibalization factors approach but never reach zero. (Parenthetically, this example also shows how adding items to an assortment does not necessarily produce more market share overall for the assortment, since the new item may simply siphon off sales of existing items.)

The cannibalization factor is actually a power-law, that is, the assortment elasticity enters into the cannibalization factor as an exponent. The cannibalization factor consists of a positive value, the Total Assortment Effect (TAE), raised to the assortment elasticity. Each item in an assortment has its own TAE; the TAE increases as items are added to the assortment. Therefore, the assortment elasticity is a negative number, in order for the cannibalization factor to decrease as TAE increases. (In the example, the TAE could simply be the count of the number of items added so far, and the assortment elasticity would then be -1, thus producing $1/2, 1/3, \dots$)

Note the similarity to the more conventional idea of power-law price elasticity, which involves a price raised to a negative power (the negative power being the price elasticity). In the cannibalization model, the TAE plays the role of price.

The cannibalization factor also accounts for the similarity of the items being added to the assortment, so that similar items cannibalize each other more than non-similar ones. The similarity values are used to calculate the TAE; higher similarities produce a larger TAE, providing a larger decrease in cannibalization factors.

The cannibalization factor depends on both the similarity values and the assortment elasticity. It might seem that similarity alone determines cannibalization, as a similarity of 0.5 between items A and B means that A takes half of B's share if A is added, but that is not the case. In particular, by separating the concepts of TAE and assortment elasticity, the model is more robust; if all of the similarity values are biased lower or higher for some reason, the bias can be accounted for by adjusting the magnitude of the assortment elasticity so that the cannibalization factors are still correct.

The Importance of Assortment Changes in Historical Data

In order to calculate assortment elasticity, DT requires historical data that contains assortment changes, because DT examines historical data to determine how much cannibalization occurred when historical TAEs changed. From the relationship between changes in historical TAEs and changes in cannibalization, DT then calculates the assortment elasticity. This is similar to calculations of price elasticity. In order to determine price elasticity from historical data, it is necessary to have price changes in the historical data, and the more changes the better.

Using the example from "[Assortment Elasticity](#)", suppose, in the historical data for a particular store *S*, that the cookies assortment has one fewer SKU in week 10 compared to week 1 (that is, some cookie SKU was removed). The TAEs for the other remaining cookie SKUs will all decrease between week 1 and week 10 because of the removal of the one SKU. DT then examines the changes in historical sales units of the SKUs in the cookie assortment at *S* between week 1 and week 10. By relating the changes in the sales units to the changes in TAEs, DT can calculate the assortment elasticity. A larger magnitude elasticity will result if the changes in TAE caused a large increase in sales units; a smaller magnitude elasticity will result if the increases are moderate.

In reality, the comparisons in this historical analysis that DT does are always more complex than in this simple example. It is rare to find a pair of weeks where the assortment change was just removal of a single SKU. Typically, in each pair of weeks, there are many assortment changes, involving both additions and removals, and the changes in TAE are a result of all of those changes. In the end, though, the relationship between the changes in TAE and the changes in sales units are summarized in a single number, the assortment elasticity, across all pairs of weeks. Because this single number summarizes the number of pairs of weeks and SKUs where TAEs changed, it is necessary an average over all the pairs of weeks and SKUs in the historical data and is not tuned to any particular SKU.

If Category Management Planning and Optimization (CMPO) is used to remove a single SKU from an assortment, it is likely that no pair of weeks in sales history exists in which exactly this SKU was removed and only this SKU was removed. For forecasting the results of this removal, CMPO makes an extrapolation from the historical analysis described above and uses the assortment elasticity that is not tuned to this particular situation of removing only this one particular SKU.

The Meaning of the Possible Values of Assortment Elasticity

An assortment elasticity of 0 turns all cannibalization factors into constant 1, meaning the assortment has no cannibalization. This is highly unlikely. However, it does show that small-magnitude assortment elasticity indicates a category where cannibalization is small. Likewise, a high magnitude of assortment elasticity indicates a category where cannibalization is large. It is possible for the magnitude to be too large (see "[The Substitutable Demand Percentage](#)").

It is also possible for the Calculation Stage of DT to produce assortment elasticities that are positive. Such positive values for assortment elasticity are an indication that there is some unidentified problem with the data, because a positive assortment elasticity means cannibalization factors increase with increasing assortment size, which in turn means each item in the assortment sells more the larger the assortment becomes. This is presumed to be a nonsensical result, and, in the Evaluation Stage of DT, such positive assortment elasticities are removed and replaced by assortment elasticities from escalation (that is, the elasticities are replaced with higher-level ones).

It is possible, with sufficient data analysis, to figure out what problem with the historical data caused the positive assortment elasticity. However, such analysis is difficult to automate, and escalation is used instead.

The Substitutable Demand Percentage

The substitutable demand percentage, or just substitutable percentage, of an item in an assortment is the fraction of its demand that is retained by the assortment if the item is removed from the assortment. It is a measure of how substitutable the item is. For example, if the substitutable percentage is 100 percent, then removing the item will not decrease the total sales units of the assortment, since all of the demand for the item will transfer to the other items that remain in the assortment. If, on the other hand, it is 50 percent, then the removal of the item from the assortment means that 50 percent of its demand is lost, and 50 percent is retained. The total assortment sales units will decrease if this item were to be removed from the assortment.

The magnitude of the assortment elasticity has an influence on the substitutable percentage. Increasing the magnitude of the assortment elasticity increases the substitutable percentage. DT only calculates the assortment elasticity for the entire category (not per item), so changing the value of the assortment elasticity changes the substitutable percentage for all items in the category all at once.

It is possible for the magnitude of the assortment elasticity to be too large. This is indicated by a substitutable percentage for several of the items in the assortment that is over 100 percent. A few items can have substitutable percentages over 100 percent, because those are probably outliers. If the assortment is large, it is likely that a few such outliers exist. If 10 percent of items in the assortment are over 100 percent, then the results should be examined.

DT provides a tool for examining the substitutable percentage and for decreasing the assortment elasticity if too many items have a substitutable percentage over 100 percent. Here are some guidelines for using this tool.

- When selecting the time interval for the tool, select one that is likely to contain assortments that are representative of the retailer's current assortments. Since the retailer is going to be using the assortment elasticity in forecasts of what happens when current assortments are modified, it makes sense to test the assortment elasticity against assortments that are as similar as possible to the current ones.
- It is possible to use the tool to dial down the assortment elasticity. Using Setting maximum substitution percentage, DT calculates an assortment elasticity that

results in substitution percentages that do not exceed the set maximum. When using this feature, you may want to set the maximum to a value higher than 100 percent if there are some outlier items that have high substitution percentages. Forcing these outliers down to 100 percent may result in a small-magnitude assortment elasticity, which may mean unacceptably small substitution percentages for all except the outlier items. You may want to select a maximum that is higher than 100 percent but that still brings most items down to 100 percent, leaving a few outliers above 100 percent.

- It is possible to use this tool to set the maximum percentage even if all substitution percentages are already below 100 percent. You may have business knowledge, or a directive from the retailer, and know that a particular category must exhibit a substitution percentage of at most 70 percent. In this case, this tool can be used to bring the substitution percentages down to 70 percent. This can make the difference between acceptance and rejection by the client.

No Requirement for a Time Interval

A time interval for the CDT calculation can be set in the CDT Data Setup stage. No equivalent exists in the Data Setup stage of DT.

The cannibalization factor directly incorporates information about the assortment through the TAE, and so the cannibalization model can handle large assortment changes. This makes it less necessary to use a time interval for DT, compared to CDT, because historical assortment changes can be directly accounted for in the model as changes in TAE.

Segments vs. Locations

In the Calculation Stage for both DT and CDT, it is possible to set up the calculation so that it is performed at all combinations of levels of the segment hierarchy and the location hierarchy. This is a more practical possibility for assortment elasticity than for the CDT calculation, because the assortment elasticity is not examined directly by people (unlike the CDTs), and producing thousands of values will not cause an issue. However, it is recommended to use only one of the two hierarchies in the Calculation Stage. Set either the segment hierarchy or the location hierarchy (or both) to be Chain. Because the calculation of assortment elasticity requires assortment changes in history, generating assortment elasticities at all levels may mean that, at lower levels, the data does not contain enough assortment changes in history. You may want to use your business knowledge of the particular retailer or particular category here, since you may know for the retailer or for the category whether assortment changes are frequent or not in the historical data you have. If the assortment changes are infrequent, you may be better off calculating a Segment- Chain/Location Chain assortment elasticity only.

Setting the Escalation Path

The last stage in DT involves setting the escalation path. If you are using only the segment hierarchy or only the location hierarchy, the escalation path is simply the hierarchy that you are using, and you set the escalation path according to the hierarchy. If you are using both a location hierarchy and a segment hierarchy, then usually you should set the escalation path to go up the segment hierarchy first, and then the location hierarchy. It is better to use only one of the hierarchies.

When using both hierarchies, the escalation path is necessary in order to tell the application which parent it should go to when moving up from a given

segment/location node. With both hierarchies in play, every segment/location node has multiple higher-level nodes that do not lie along a single path. The escalation path is necessary to tell the application in what order the higher-level nodes should be considered. When only one hierarchy is used, the higher-level nodes form a single path.

Automatic Updating

DT can automatically and periodically update the assortment elasticities as new sales history is available. This feature is unique to DT; CDT does not perform automatic updating because it makes less sense to automatically produce new CDTs. New assortment elasticities can be loaded into the consuming applications and thus immediately used; however, the value of new CDTs is less clear.

When new historical transactions enter the RA schema, DT will automatically aggregate them and produce new SKU-store-week sales-units aggregates. These new aggregates are then appended to the older SKU-store-week aggregates, and the resulting data set is then used in a new calculation of assortment elasticities.

Note the following about this calculation:

- It does not in any way run the full DT application, that is, re-run all of the stages. The calculation is more targeted and just calculates assortment elasticity.
- It only updates assortment elasticities, not the similarities from the Similarity Calculation Stage.
- Because it uses a mix of old data and more recent data, the values of the assortment elasticities will change slowly over time as the data set becomes more tilted towards newer data. This is by design. It is not desirable to have sudden changes in assortment elasticity, since that would result in sudden changes in cannibalization and demand transference.
- Any assortment elasticities that were overridden using the Substitutable Percentage tool (see "[The Substitutable Demand Percentage](#)"), stay overridden, and are not updated.

Avoiding Categories with Small Assortments

It is possible for a retailer to have categories where the assortments are very small, that is, 20 or fewer items in the assortment. Such categories can pose a problem for DT because of the small amount of sales data for just 20 items, and also the number of assortment changes may be quite few.

It is better if the assortment is small but items from a much larger set have been added or removed frequently from the assortment. That is, the category has a much larger set of items, but only 20 of them are in an assortment at a given time. It is possible that the assortment changes were frequent enough that more than 20 items have sales history, and in this case DT results may be reliable even though the assortment is small.

Implementing DT for Fashion Categories

For various reasons, fashion categories require some special consideration. This section describes what is different about them and how to handle the differences.

Proper Level for Fashion Categories

The lowest level of data must not be the SKU level, that is, the Size level of the merchandise hierarchy. Because size is a functional-fit attribute, or nearly so, the level of calculation must be at least one level above size (Style-Color). The historical sales-units data must be aggregated at least up to Style-Color. This also helps avoid problems with low sales rates and noisy sales rates at the SKU level, both common problems for fashion categories. It also helps decrease the number of SKUs within a category, since a multitude of sizes are possible for each Style-Color.

It is worth considering whether color is necessary. Aggregating to the Style level means that transference among colors cannot be calculated. However, it is not clear how useful calculating transference among colors would be, since the colors change for every selling season, and calculating historical transfereces among colors may not be particularly useful. A possible halfway approach here might be to aggregate to Style-Primary-Color, where there are only a few primary colors. The primary colors chosen can be the ones that are stable season after season, so that historical transfereces among them might be useful in future selling seasons. The primary colors can be chosen to be groupings of the actual colors (so for example, midnight blue and sky blue would become blue). In general, in fashion, the number of colors can be large, and it is unlikely that calculating transfereces among such colors would be useful. Aggregating to Style-Primary-Color or even to Style can help avoid low sales rates and noisy sales rates.

It is possible to employ different approaches in grouping the colors. One is to use the primary color. It is also possible to group the colors based on the type of customer the color is designed to attract. For example, the colors can be grouped in "trendy colors" vs. "basic colors." The grouping should be decided in consultation with the retailer, to determine how the retailer uses colors in the category. The retailer may already have a grouping of colors that it uses, and the simplest approach may be to use this grouping.

Here, it is assumed that the aggregation of the historical sales-units data to be either Style-Primary-Color or Style. The term "item" should be understood to mean Style-Primary-Color or Style, depending on the chosen aggregation.

Typically, for fashion, the number of colors is large because of all the color variants. If it is necessary to retain all of the colors instead of following the recommendations above, then it will be necessary to split the color attribute into at least two attributes, a primary color and a secondary color. For more information, see "[The Role of Attributes in Calculating Similarities](#)".

Seasonality (Life Cycle) Considerations

DT makes the assumption that the items within a category at a store all have a common seasonality (see "[Seasonality in Historical Sales Data](#)"). Because of the short (tens of weeks) life cycles of fashion items, and because items within a category may have different introduction times within the same store, the assumption of common seasonality across the items in a category is probably not valid for fashion categories. It is possible within the category to have items that are at various points in their life cycles within the same store. At the same time, some items may be in the uptrend part of their life cycle, while other items are in the downtrend part of their life cycle.

There are some ways to deal with this by properly setting up the input data for DT.

One simple approach is to approximate the life cycle of an item by using the SKU-store-ranging described in [Demand Transference](#). In this approximation, the range for an item is set to start at x percent sell-through of the item and end at y percent sell-through. Choose x to be 5, and y to be 70. X must not be 0 (that is, the point of introduction of the item), since it takes sometime after the point of

introduction of the item for customers to start buying it in quantity and for the item to start having any kind of cannibalization effect on the other items. \mathbf{Y} must be set to a point where significant numbers of customers have started to either lose interest in the item or where the item no longer has sufficient numbers of sizes available. In either case, customers are now transferring their demand to the other items in the assortment, so it is as if the item were no longer in the assortment. This is an approximation as the item is still in the assortment and is still selling, just at a significantly lower rate than its peak sales rate.

Assortment Recommender

This chapter provides details about the ORASE Assortment Recommender, a batch-run system that provides recommendations for changing the current assortments in each store and each category in order to increase revenue or gross profit from that category at that store.

Prerequisites

The Assortment Recommender relies on output from the DT module of ORASE. In particular, every category/store combination that is to receive assortment recommendations must have DT results in an approved DT version. Some of the concepts used in the Assortment Recommender, such as substitutable and incremental demand, are concepts that the DT application also uses.

Producing Better Assortments

The Assortment Recommender starts from the current assortments for each category in each store and finds assortment changes that will increase either the total revenue or total gross profit of the assortment. Each category-store combination is optimized separately.

The total revenue or total gross profit calculation accounts for

- Cannibalization effects, by using demand transference, and
- Halo effects, meaning additional revenue a SKU in category brings by encouraging the purchase of complementary SKUs in other categories.

The assortment changes recommended are those that improve revenue or gross profit after accounting for the above effects. For example, the Assortment Recommender may recommend dropping an item that is very similar to the other items in the assortment and replacing it with an item that is less similar to the other items, because dropping the similar item does not decrease total revenue by much, and adding the dissimilar item brings in additional revenue. Similarly, the Assortment Recommender may recommend dropping a SKU that brings little halo revenue in favor of a SKU that brings in more halo revenue (or gross profit).

The calculation of revenue or gross profit does not include any revenue that an assortment change in other categories may bring though halo effects. Each category at a store is optimized separately, so when optimizing a category B, it is not possible to include revenue that another category A may bring to B through the halo effect of a SKU in A on B.

The Assortment Recommender lets you choose which of the following to maximize:

- total assortment sales units
- total assortment revenue
- total assortment gross profit

These quantities include halo sales units, halo revenue, or halo gross profit, respectively

The Assortment Recommender runs as a batch process and on a set schedule produces recommended assortment changes for each category/store combination using whatever is the current assortment at the time it runs. See the following sections for how the schedule can be configured.

Run Groups and Run Frequency

It is unnecessary to obtain frequent assortment-change recommendations for every single category/store combination. For many category/store combinations, recommendations may only be necessary at infrequent intervals, or not at all. The Assortment Recommender provides interfaces to allow external control of the run frequency of category/store combinations in the following way.

The interfaces allow for defining a run group, which consists of a set of categories and a set of stores. When a run group executes, every combination of category/store from the set of categories and set of stores receives new recommendations. The run group as a whole is associated with a run frequency, which specifies how often the category/store combinations in it receive new recommendations. In this way, the run group is a mechanism to control:

- The category/store combinations that receive recommendations. For example, if, for category A, only stores 1, 2, and 3 should receive recommendations, then a run group can be set up with category A, and stores 1, 2, and 3.
- The frequency of recommendations. A retailer may be particularly interested in certain key categories, so it makes sense to schedule recommendations only for those categories. However, not every category may need frequent recommendations.

Keep in mind that when the run group executes, every combination of category/store in the run group receives recommendations. If the run group has 10 categories and 500 stores, this is $10 \times 500 = 5,000$ sets of recommendations and 5,000 separate calculations to produce those recommendations.

The Run Group Parameters

Various parameters control how the recommendation calculation proceeds when a run group executes. These parameters must be present in the run-group tables in the database. Within a run group, each category has its own set of these parameters because the parameters depend on the category. The parameters are:

- The set of Must-Keep SKUs. Each category in a run group can have a list of such SKUs, which indicate to the Assortment Recommender that these SKUs must not ever be removed from the assortment. For example, key items for the retailer in the category must be on the list. The list can be empty, in which case the assortment recommender is free to swap out any SKU currently in the assortment.
- The Assortment-Size Change. This is an integer that can be negative, positive, or 0. It indicates the change the Assortment Recommender must make to the number of SKUs in the assortment. Suppose the value is C . Then the Assortment Recommender will change the assortment size by C . If C is 0, then the assortment

size stays the same. If negative, the assortment size decreases by $-C$, and if positive the assortment size increases by C . There is one value of C for each category in a run group. For a given category, C applies to all of the stores in the run group. It may not be possible for the Assortment Recommender to achieve a change of C for a particular category in a particular store, but it will attempt to get as close as possible.

- The Min-Keep Percent. This is a non-negative percentage. Suppose the value is M . Then the Assortment Recommender keeps at least M percent of the SKUs in the assortment to be the same. Suppose M were 50 percent. Then at least half the assortment will consist of SKUs that are already in the assortment, but the Assortment Recommender is free to choose which 50 percent to change. This parameter, along with the Must-Keep SKUs, is useful for ensuring that the Assortment Recommender does not make too many changes to the assortment.
- The assortment metric to maximize. The choices here are: sales units, revenue, or gross profit. This is a total over the entire assortment, and the Assortment Recommender recommends assortment changes that increase the chosen metric from what it is for the current assortment. The choice of assortment metric is per category within a run group, and it applies to all of the locations in the run group. For example, if a certain category is a loss leader but drives the customer to make other purchases, then the choice of metric for this category might be sales units.

Data Used by the Assortment Recommender

The Assortment Recommender requires the following data to generate assortment recommendations. Each data element is derived automatically from other data in the ORASE schema and fed into the Assortment Recommender. The following list describes the data and how it is derived.

Required Data

- The current assortment for each category at each store. As discussed above, for each category/store combination in a run group, the Assortment Recommender starts with the current assortment and makes changes to it to increase the chosen assortment metric. For each category/store combination, the system takes the SKUs that were selling in the store during the last available week of historical data.

The weekly sales-units rate of each SKU in the current assortment. This is calculated through an average over the most recent four weeks of historical data.

- The price of each SKU in the current assortment. Historical price data is not used for this, but instead the total revenue over the most recent four weeks of historical data is divided by the total sales units over the same four weeks. This provides an average historical price, based on the last four weeks of historical data.
- The gross profit of each SKU in the current assortment. Historical price or cost data is not used; instead, an average is taken similarly to the price calculation. The total gross profit is taken over the most recent four weeks of historical data and divided by the total sales units over the same four weeks.

The sales-units rate, price, and gross profit are required in order to support the possible assortment metrics.

Notice that in addition to data about the current assortment, the Assortment Recommender requires data about possible SKUs that it can swap into the assortment, since otherwise, in the case of keeping the assortment sizes the same or expanding the assortment, no recommendations would be possible. (The case of decreasing

assortment sizes is discussed separately below.) For each category/store combination in a run group, the system calculates the following for the possible SKUs to be swapped into an assortment:

SKU Data

- For a category/store combination, the set of new SKUs is the SKUs that are in the current assortments of other stores but not in the current assortment of this store. (See above for how the current assortment of a store is determined.)
- The price of the new SKU is the average of the prices in the stores in which it is part of the current assortment. The price at each store in which it is selling is determined as described in ["Required Data"](#).
- The gross profit of the new SKU is the average of the gross profits in the stores in which it is part of the current assortment. The gross profit at each store in which it is selling is determined as described in ["Required Data"](#).
- Assigning the sales units rate of the new SKU is the trickiest to handle, since the sales-units rate of the new SKU must be forecast based on the assumption that it is selling at a store where it may not have sold before. Here, the Assortment Recommender identifies like items of the new SKU among SKUs that are currently in the assortment at the store, and from the like items it makes a forecast of the new SKU's sales-units rate at this store. Identifying the like items is done through the use of similarities. For more information about similarities, see the section ["The Role of Attributes in Calculating Similarities"](#).

If the Assortment-Size Change parameter is set to a negative value, then it is possible to decrease the size of the assortment without having data about possible SKUs to swap into the assortment, as in this situation the Assortment Recommender would simply be finding SKUs to delete from the assortment while still maximizing the chosen assortment metric.

Halo Effects

As mentioned above, the Assortment Recommender accounts for halo effects when calculating the selected assortment metric. This calculation uses the output of the Market Basket Insights. The MBI determines halo effects at the sub-class level. That is, sub-class A of one category has a halo effect on sub-class B of another category, meaning some significant fraction of the people who purchase in a SKU in A also purchase a SKU in B. Suppose the Assortment Recommender is running for a particular category C. The Assortment Recommender, when it considers putting a SKU C into the assortment, adjusts upward the amount of the assortment metric that C brings in order to include the halo effect. For example, suppose C is a SKU in sub-class A, and sub-class A brings a halo lift of 10 percent to sub-class B. If the metric the Assortment Recommender is maximizing is sales units, then to the sales units U of C itself, the Assortment Recommender adds sales units of $0.1U$ to represent the sales units of B bought by purchasers of C. Similarly, if the chosen metric is revenue, then to the revenue brought by C itself, $0.1U$ times the average price of SKUs in sub-class B is added. The average price of SKUs in sub-class B is calculated by a weighted average of prices, with the weights being the weekly sales-units rates.

The handling of gross profit is similar to the handling of revenue, except that a weighted average of gross profits of B is used instead of the weighted average of prices.

The above discussion involves adding SKU C to the assortment, but the same discussion holds if the Assortment Recommender is removing SKU C from the assortment.

Troubleshooting

Several conditions can prevent the Assortment Recommender from producing recommendations for specific category/store combinations. When any of these occur, the Assortment Recommender will not produce an error but will simply not produce recommendations for the particular category/store combination.

- The DT application was not run for a particular category, or the category does not have an approved DT version associated with it. This means the category does not have results from the DT application, and without those results, it is not possible for the assortment recommender to run since it cannot account for demand-transference effects. In this case, the category will not receive any assortment recommendations regardless of store.
- The Assortment Recommender was not able to find any new SKUs for the particular category/store combination. In the above description about finding the set of new SKUs, it is possible that the procedure does not find any new SKUs at all, perhaps because all of the stores are assorted identically at that point in time for this particular category. This may happen with categories that are less important to the retailer, so that the retailer does not see any benefit in tailoring the assortment within each store.
- The Assortment Recommender was not able to find any assortment changes that result in an increased assortment metric. This can happen if:
 - There were no new SKUs available (see previous item).
 - The number of new SKUs available was very small.
 - The run-group parameters for the category are too restrictive. For example, too many SKUs are listed as must-keep SKUs or the min-keep percentage was set too high (greater than 80 percent).
- The Assortment Recommender may run, but without using halo effects if the halo effects are not available. For example, the MBI may not have run or may not have produced halo effects for the category in question.

Advanced Clustering

This chapter describes the Advanced Clustering Cloud Service. It provides details on configuration and implementation.

Overview

Advanced Clustering (AC) lets users create store clusters based on common features such as customer demographics in order to manage merchandise assortments and pricing strategies in a targeted way. Clusters can help retailers understand who shops in their stores and what their preferences are. Clusters can be used to inform decisions about assortment, pricing, promotion, forecasting, allocation, and supply chain processes based on selling patterns in stores. An understanding of the characteristics of the customers who shop in a store and what they buy can help a retailer target specific customers

The application 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 the business objectives defined when the clusters are designed. Users can define a hierarchy cluster of stores based on a store attribute such as format and then cluster further using performance attributes in order to determine which stores have high, medium, and low sales. What-if scenarios and ranking can be used to compare how cohesive and well separated clusters are in each scenario as the number of cluster centers is increased. The application uses scoring to indicate which clusters fall below defined thresholds and may require manual intervention. Business Intelligence graphics illustrate the patterns in the data and the attributes that are important in each cluster.

The key features available in AC are:

- Dynamic nested clustering, in which a user can cluster on a criteria, analyze the results, and then decide whether or not to further sub-cluster.
- Mixed attribute clustering. A cluster can be created on continuous attributes such as performance (sales, revenue, and gross profit) as well as discrete attributes such as store size, demographics, and seasonality, all at the same time.
- Configurable clustering criteria such as customer profiles, product attributes, performance, and store attributes.
- Recommendations are made for the optimal number of clusters and the scores for each cluster. These are based on the quality of the clusters: how cohesive and well separated the clusters are.

Data Requirements

Advanced Clustering relies on following data and it uses ETL to load the data.

Table 7–1 Data Requirements

Objects	Granularity	Required/Optional
Hierarchies	Product, Location, and Fiscal	Required
Location Attribute	Store	Required
Product Attribute	SKU	Required
Aggregate Sales Data	Week/SKU/Store	Required
Customer Segment Profiles	Store/Customer Segment/(Category or All Merchandise)	Optional
Alternate Hierarchy	CM Group or Trade Area	Optional
Like Locations	Store	Optional
Product Attribute Group and Value	Category or Subcategory	Optional
Aggregate Forecast Sales Data	Week/SKU/Store	Optional

Multiple Hierarchies and Level Support

Store clusters can be generated for the following combinations of hierarchies.

Product Hierarchies

This includes:

- Core merchandise hierarchy
- Alternate hierarchy

Clusters can be defined for either of the hierarchies and for different hierarchy levels. For example, clusters can be defined for Chain, Department, or Category along the product hierarchy.

Note that store clustering can only be defined for a product hierarchy level higher than Item. Item level store clustering is not supported. However, store clusters can be generated for item groups by defining an item group as a level in the alternate hierarchy.

Location Hierarchies

This includes:

- Core location hierarchy
- Alternate hierarchy (optional)

Clusters can be defined for either of the hierarchies and for different hierarchy levels. For example, clusters can be defined for Chain, Trade Area, or Region along the location hierarchy. Store clusters can be generated for channel, if channels are configured as a level in the location hierarchy.

Calendar Hierarchies

This includes:

- Core fiscal calendar hierarchy (week, month/period, quarter, half, year)

- Gregorian calendar (week, month, quarter, half, year). Leverages a start and stop date (day level date range)
- Planning period. Leverages alternate hierarchies, including planning period, buy periods, and defined holiday time periods such as back to school and fourth of July. This is optional.

Clusters can be defined for any of these three calendar hierarchies (the cluster effective period). Note that the source time period for historical data only uses the core fiscal calendar hierarchy, supporting hierarchy levels that include week, month/period, quarter, half, and year.

Clustering Criteria Supported in Store Clustering

In store clustering, the cluster criteria are a set of attributes that define store clusters. These attributes can be either discrete or continuous. A group of these clusters is called "Cluster by." For example, demographic data such as income and store properties such as store format can be grouped into a store attribute Cluster by.

These default Cluster by are supported.

Customer Profile

Stores are clustered based on the similarity in the mix of customer profiles shopping in the stores and trading areas. These clusters form the basis for further analysis to understand which customers shop in which stores and how they shop. Retailers obtain market data from market research firms such as the Nielsen Corporation and use the data to create customer profiles for their stores. An ORASE feed can be used to provide this information to AC at the category or all merchandise level.

Location Attribute

Stores are clustered based on how shopping behavior varies by store attribute. This provides information about who is shopping in a store or trading area as well as demographic data such as ethnicity, income levels, education, household size, and family status. Retailers can analyze the cluster composition and related business intelligence in order to better understand the shopping behavior of their customers. This can help retailers make assortment and pricing decisions.

Product Attributes

In this type of clustering, stores are grouped together that have similar sales shares for one or more product attributes (for example, coffee brands such as premium, standard, and niche). The percentage of each store's contribution is calculated using the Sales Retail \$ for each product attribute value to the total retail sales of the category or subcategory of each location.

Product attributes can be configured only at category or sub-category levels.

Performance Criteria

Stores are clustered based on historic sales by considering performance at different merchandise levels while performing store clustering and analyzing how the shopping behavior varies by category. This can help to identify high, medium, and low volume stores.

Forecast Criteria

Stores are clustered based on forecast sales by considering future sales at different merchandise levels while performing store clustering and analyzing how the shopping

behavior varies by category. This can help to identify high, medium, and low volume stores based on the predicted sales.

Mixed Criteria

Discrete and continuous attributes are combined together. Retailers can cluster stores using attributes from all the above defined criteria at the same time.

Attributes in Store Clustering

Cluster by uses a collection of attributes, including consumer profile attributes, sales metric attributes, location attributes, and product attributes.

Sales Metrics

Store clustering uses a fixed set of sales metrics. These attributes cannot be extended. Supported attributes include Sales Retail \$, Sales Unit, Sales AUR, Gross Margin R, and Gross Margin %.

Forecast Sales Metrics

Store clustering uses a fixed set of predicted sales metrics. These attributes cannot be extended. Supported attributes include Forecast Sales Retail \$, Forecast Sales Unit, Forecast Sales AUR, Forecast Gross Margin R, and Forecast Gross Margin %.

Location Attributes

Store clustering relies on location attributes that can loaded into ORASE as either core attributes or as user-defined attributes. These attributes are defined for stores. They define the store properties, including demographic and geographic details. During installation, only attributes that have 15 distinct values are configured. Attributes with higher discrete values are not considered for store clustering.

Product Attributes

Product attributes based on store clustering use two types of attributes, raw attributes and grouped attributes. The former are product attributes, identified as important attribute values. The latter are fed into Oracle Retail Advanced Science and are available to the store clustering process only when the CDT module is enabled. The clustering process groups together stores that have similar attributes values for a product category. A store share is calculated using Sales Retail \$, which is the ratio of the sales retail of each product attribute value to the total sales retail of the category or subcategory of the specific location.

Each category or subcategory must have raw attributes and grouped attributes.

Grouped attributes (the default) classify the items in the class or subclass. This set of attributes differs from class to class. For example, for yogurt, the attributes are: size, flavor, brand, fat percentage, and pack size. For chocolate, the attributes are: size, brand, milk/dark, nut type, and package type. The two classes can both have the attribute of brand, but the brand attribute will have different values for each of the categories. Group attributes have a mapping of each item in the category to its set of attribute values. This information is provided as a data feed to the application. If grouped attributes are not available then raw attributes are used for the store clustering of product attributes.

Raw attributes typically have a large number of attribute values. For example, the brand attribute for yogurt may list 50 different brands at a large grocer. Using raw attributes, the system runs a preprocess to identify n (default = 3) attributes values that are most frequently sold for each attribute in a category or subcategory.

Configuration Process

Default configuration occurs during the installation and upgrade. The configuration process is responsible for installing or upgrading any new attributes in the application. This process ensures that any existing manual overrides introduced by the retailers are not overridden and any new additions are brought into the clustering process. The default configuration includes the following:

- All attributes are enabled by default and weights are normalized among all the configured attributes.
- Any discrete location attribute that has more than $n=15$ attributes values is not configured by default. Note that the value of n is a configuration and can be modified at the time of deployment.
- The UI formatting of each attribute is identified based on the data type of the attributes.
- Nesting is enabled by default for all types of Cluster by (except mixed, which is an alternative approach to clustering).
- The deployment of clusters at multiple hierarchies or levels is enabled.

The following configurations may require manual overrides if the default configuration is not acceptable or data is not available.

Table 7–2 Manual Overrides

Name	Description
Enable or disable Cluster by	Disable Cluster by. For example, for the consumer profiles for each store for a retailer, the Cluster by for the consumer segment must be turned off.
Enable or disable nesting	Allow multiple nesting levels under an existing Cluster by. For example, cluster first by product performance, with nested clustering by store attributes.
Enable or disable attributes	Enable an attribute to be considered for clustering or contextual BI. For example, label the population density attribute as a BI attribute instead of a clustering attribute, as very few stores have data for population density, and it is not significant enough for store clustering.
Change UI formatting	Change formatting associated with the attributes, such as label, decimals, percent, and currency. These are configurations for each attribute and do not rely on XLIF entries.
Cluster deployment	Enable or disable hierarchy at which clusters can be deployed. For example, if CMPO store clusters are only defined for categories, then only clusters at the category level will be approved. The other levels can be enabled if needed.
Outlier Rule	Change default outlier rules for a Cluster by. By default, the distance from centroid rule is enabled. See section below for other supported outlier rules.
New Store Rule	Change default store rules for a Cluster by. By default, the like location rule is enabled. See section below for other supported new store rules.

Table 7-3 Enable or Disable Cluster By

Cluster By	Description	Example	Enable
Customer segment profile	Cluster store using consumer segment distribution	20% soccer mom, 30% empty nesters	Enable if consumer segment profiles are available for each store.
Store attribute	Cluster store using location attributes	Income, climate, size, store format	Enable if location attributes are available for each store
Performance	Cluster store using sales metrics	Sales revenue, sales unit, gross margin \$	Enable if retailer wants to cluster store using performance metrics
Product attribute	Cluster store using product attribute sales shares	Brand, color, seasonality, size/fit	Enable if retailer wants to cluster store using product attributes sales share
Mixed attributes	Cluster store using mixed attributes, combining attributes across all the cluster criteria	Income, climate, size, store format, sales revenue, sales unit, gross margin	Enable if retailer wants to cluster store using combination of attributes
Product forecast	Cluster store using predicted sales metric	Forecast sales revenue \$, forecast sales unit, forecast gross margin \$	Enable if retailer wants to cluster store using future sales metrics

Multiple Clustering Approach

Store clustering supports three types of clustering: simple, nested, and mixed. By default, all three approaches are enabled. These approaches are applicable to all Cluster by. Store clustering functionality supports dynamic nesting capabilities. For example, the user can cluster on a criteria, analyze the results, and then decide to further sub-cluster.

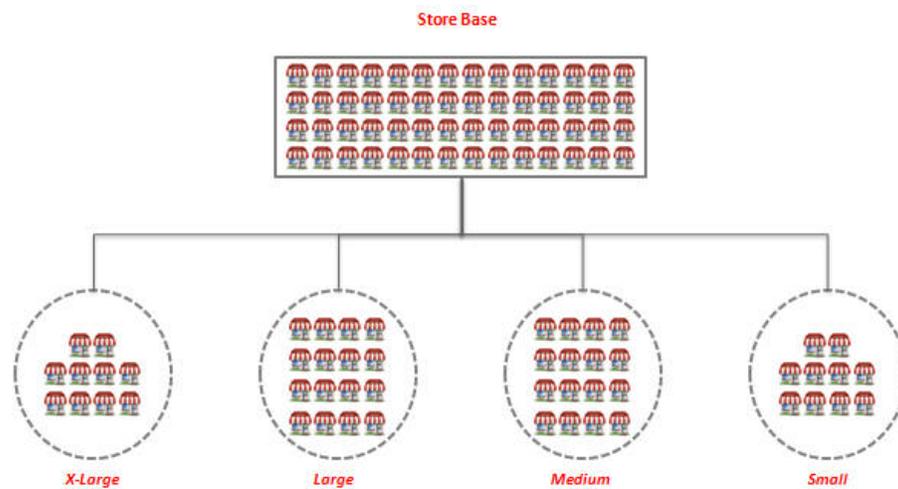
Mixed attribute clustering is also supported. For example, it is possible to cluster on continuous attributes such as performance (sales, revenue, and gross profit) as well as discrete attributes (store size and demographics) at the same time.

Simple

Users can select attributes from a Cluster by. For example, users can select location attribute Cluster by and generate clusters using location attributes such as store size.

Users select the most important store attribute (based on category) and group the stores accordingly. Clusters may or may not cross trade areas, regions, and districts. This depends on the approach as well as the responsibility of the planning team.

Figure 7-1 Simple Clustering

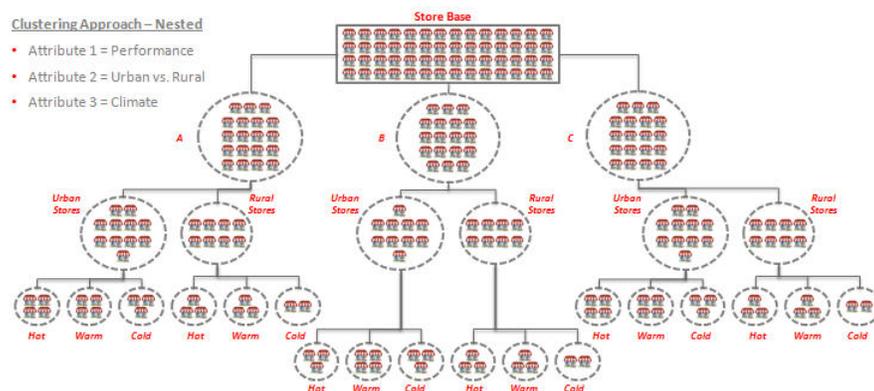


Nested

All Cluster by, except Mixed Attributes, can use a nested hierarchy by default. For example, performance clusters can be further clustered using location attributes. This approach allows a dynamic hierarchy of clusters. Nesting can be enabled or disabled in the AC configurations.

Users can select the most important store attributes (based on the category and the group of stores selected during the wizard process) and group the stores accordingly to ensure that a more refined assortment can be created by category or location selection. Clusters may or may not cross trade areas, regions, and districts. This depends on the approach as well as the responsibilities of the planning team. Once the initial clusters are created, users can further cluster using attributes to define nested clusters. This can be done within a trade area as well as at the total company level. The number of clusters is granular, as compared to mixed attributes.

Figure 7-2 Nested Attribute Clustering

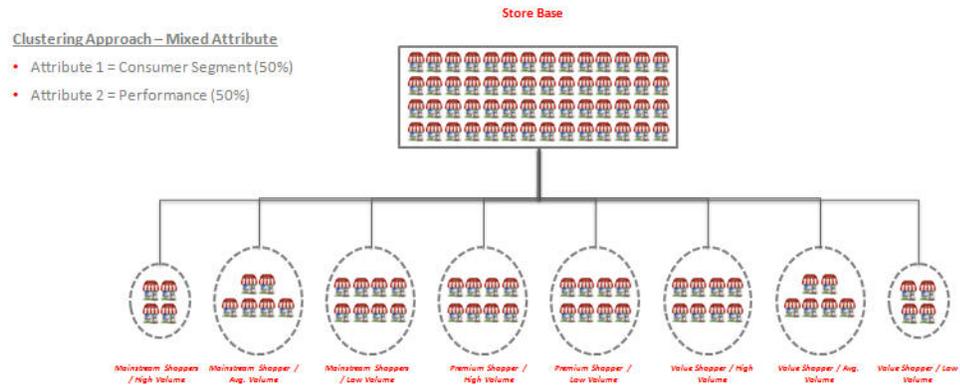


Mixed

Mixed attributes, including Cluster By such as consumer segment, location attributes, performance, and product attributes, are supported by default. Users can combine attributes from different Cluster by. For example, users can combine attributes from consumer segment and performance Cluster by and generate clusters using sales revenue and consumer segment distributions.

Users can use the most important store attributes (based on the category and the group of stores selected during the wizard process) and group the stores accordingly. Clusters may or may not cross trade areas, regions, and districts. This depends on the approach as well as the responsibilities of the planning team. These clusters are the final ones used in assortment process. The number of clusters in this case is confined, compared to the nested clustering approach.

Figure 7–3 Mixed Attribute Clustering



New Stores or Stores with Poor History

Store clustering supports three rules for allocating new stores to a cluster. These rules can be configured for each Cluster by. A rule is applied after the clusters are generated.

These rules require a feed into Oracle Retail Advanced Science that defines a mapping between a location and like locations. This mapping can be configured by merchandise. One location can be mapped to multiple locations with different weights.

Like Stores (Default Rule)

Stores with new history or poor history are allocated to the same cluster in which the like locations are allocated. For example, a new store or a store whose poor history has been corrected can be allocated to a valid performance cluster.

Largest Clusters

New stores or stores with a poor history can be allocated to the largest cluster identified by the clustering analytics. For example, a new store that has not yet formed a customer base can be allocated to a larger cluster until significant customer profiles have been collected.

Cohesive Clusters

New stores or stores with poor history are allocated to the most compact cluster identified by the clustering analytics. For example, stores can be assigned to a cluster that has not been not affected by outliers.

Outliers

Store clustering supports two rules that identify stores as outliers in a cluster. These rules can be configured for each Cluster by. The rule is applied after the clusters are generated.

Distance from Centroid (Default Rule)

The distance from a store to the centroid is identified. If the distance is beyond a defined limit of the configured threshold from the centroid, then the cluster is identified as having outliers. The user must investigate such clusters.

Cluster Size

The percentage of stores that are allocated to certain clusters is identified. If they fall beyond a certain limit in comparison to total number of stores, the cluster is identified as having outliers. The user must investigate such clusters.

Export to Excel

To use any reporting tool with an Excel file exported from Advanced Clustering, you may need to adjust the format. Here are some examples of possible formatting adjustments.

- The Text column should remain as is.
- If the Percentage column uses a percent symbol, then that symbol must be removed.
- If the Currency column uses either commas or currency symbols, then those must be removed.
- If the Number column uses commas, then they must be removed.

Assortment and Space Optimization

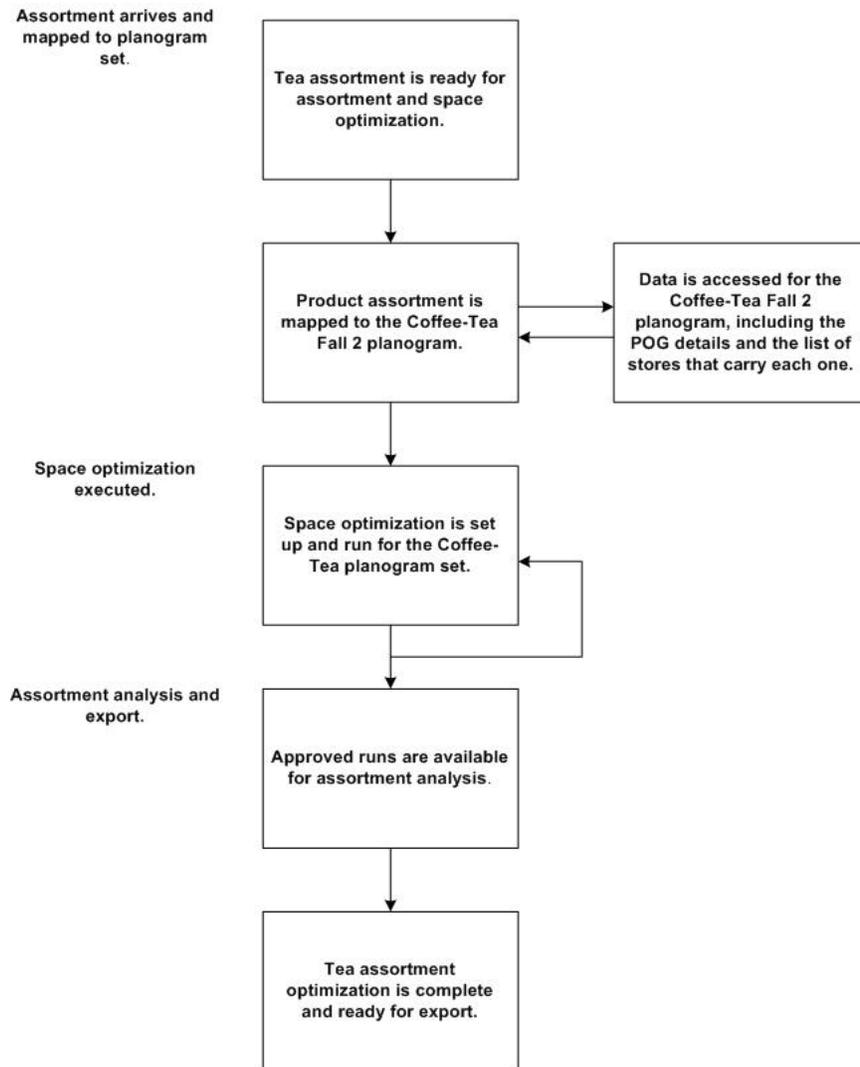
This chapter describes the Assortment and Space Optimization Cloud Service.

Overview

Assortment and Space Optimization (ASO) is used to determine the optimal selection and arrangement of products within stores by optimizing the product assortment and product placement on a planogram. ASO uses information about available space in stores, product dimensions, fixture configurations, expected demand and demand transference, replenishment schedules, target service levels, merchandising rules, visual guidelines, and category goals to create VPOGs that optimize total performance. In short, ASO generates space-informed optimal product assortments in the form of a virtual planogram. (A VPOG is ASO's assortment recommendations in product facings, depth, stacking, and SKU order in form of a planogram.)

The products that are selected for the VPOGs are the ones that ASO recommends for the finalized assortment. The recommended product level data is used inside CMPO, while the virtual planograms created in ASO are available for use in space planning applications. [Figure 8-1](#) shows the flow of an assortment through ASO.

Figure 8–1 ASO Assortment Flow



ASO supports store level and space cluster level optimizations runs. 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 planogram length, depth, height, and fixture type, if additional planogram attributes are selected). Space clusters are a level between assortment cluster and store. ASO creates space clusters by partitioning stores from an assortment cluster into smaller groups of stores that have the same product list and planogram attributes. In a store level optimization run, each store is optimized independently using the store's specific data to generate one optimal planogram for each store. On the other hand, in a space cluster-level optimization run, every space cluster is optimized independently, using aggregated store data to produce one optimal planogram for each space cluster.

ASO supports shelves, pegboards, freezer chests, and a combination of shelf and peg-board planogram fixture types. The smart start process in ASO creates fixture details using a combination of default values and user selection and assigns shelves to the base fixture of the fixture profile.

At a high level, ASO starts with an assortment that is ready for optimization. The assortment is mapped to one or more planograms and one or more optimization runs.

Approved runs are then available for assortment analysis and can then be finalized and exported.

Figure 8–2 shows an overview of ASO workflow, which is described here:

- **Receive Assortment:** ASO receives a preliminary assortment list from CMPO that is to be space optimized. The assortment is mapped to a set of available planograms through assortment-to-planogram mapping files.
- **Set up Optimization:** Allows users to specify the location level for the optimization, select planogram and optimization locations, select or update available fixture configurations, view or modify product merchandising options, and demand and replenishment data. Note that the majority of this information is pre-loaded into the solution.

Figure 8–2 ASO Workflow Overview



- **Assortment and Merchandising Rules:** Allows users to specify visual blocking and sorting rules, product and planogram constraints, assortment rules like the list of products that must be kept together, pick at least three from a list of five products, and so on. Note that mandatory items from CMPO are automatically available here.
- **Analyze Optimized Space and Assortment:** Allows users to view and analyze optimized assortment and associated facings recommendations as well as a visualization of the optimized POG. The user can also review KPIs such as optimized service levels, expected sales, and profit.
- **Finalize Cluster and Store Virtual POG:** Allows users to interact with VPOG and override recommendations before finalizing and approving optimized assortments.

ASO Data Input Requirements

This section provides information about setting up the data that the ASO application uses to generate optimal product assortment and placement on virtual planograms. See *Oracle Retail Advanced Science Cloud Services Implementation Guide* for detailed file formats and definitions.

Assortment Data

ASO requires product assortment data that contains information about assortments, assortment clusters and products within an assortment cluster, placeholder product information, like-items for placeholder products, price, cost, and forecast data for products in the assortment.

Note that this section does not discuss the assortment files related to placeholder attributes, finalized assortment placeholder products, or assortment finalization data. Information about these files can be found in *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface*.

Here is a list of the assortment data files.

- **Assortment staging file:** This file specifies the assortment header and general information about the assortment, including assortment goal.

- Assortment cluster file: This file provides information on assortment clusters.
- Assortment cluster membership file: This file contains information on stores to assortment cluster assignments.
- Assortment product store/cluster file: This file defines the list of products (including placeholder products) for a given assortment and store/assortment-cluster combination. It also contains the IPI and product priority indices.
- Assortment product location forecast file: This file contains the weekly forecast for a list of products (including placeholder products) for a given assortment and store/assortment-cluster combination.
- Assortment product location price and cost file: This file provides the regular retail price and the cost for the list of products (including placeholder products) for a given assortment and store/assortment-cluster combination.
- Assortment placeholder product like product file: This file contains the list of placeholder products and a like product for each placeholder product in the assortment. ASO uses the like product information available in the file to extract product size data, product attributes, and replenishment data.

Planogram Data

ASO requires planogram (POG) data that is used to define POG dimensions, categories, seasonal information, and product display geometry. CMPO provides ASO with the files it requires. Thus, CMPO customers can export POG data to files that ASO can import readily. Customers who use another POG or space planning solution must provide POG data to meet the ASO interface definitions.

The following POG files are needed for ASO.

- POG definition file: This file defines the major characteristics of a POG, including name, category, status, season, and dimensions.
- POG store file: This file maps a POG to a particular store or a set of stores.
- POG display style file: This file lists the display styles used in certain planograms.
- Product display style orientation file: This file is a cross reference between display styles and orientation. This lists the valid orientations for each display style. It is required that each display style must be mapped to at least one orientation.
- POG bay configuration file: This file provides a list of bays used by the POG.
- Fixture definition file: This file provides a list of the fixtures that define the POG. Fixture attributes specific to pegboards and freezer chests are also defined in this file.
- Fixture configuration file: This file describes the fixture layout in a bay. A fixture can be a shelf, pegboard, or freezer chest.
- Display style compatibility file: The file cross references fixture types and display styles. It lists the fixtures for which the display style is valid.
- Shelf definition file: This file is required for shelf fixture planograms. It provides the details for each individual shelf in the fixture.
- Shelf configuration file: This file describes the shelf layout in a fixture.
- Product display style file: This file specifies product to display style mappings. It provides a list of display styles that are available for a specific product.

- Display style definition file: This file provides the display style product settings and dimensions.
- Shelf product configuration file: This file describes the product layout on the shelf fixture.
- Pegboard/freezer product configuration file: This file describes the product layout on the pegboard/freezer fixture.
- Store custom attributes: This file provides user-defined POG attributes for every store/POG.

It is recommended that, at implementation time, planogram data is imported in bulk rather than on ad-hoc basis. Additional planogram data can be loaded incrementally.

Assortment-to-Planogram Mapping

A partially automated process in ASO attempts to map assortments to planograms. The process matches the seasons associated with planograms and assortments and considers demand spread factors for products in assortments that are assigned to multiple planograms at one time. The user can achieve the desired mappings by creating or editing a pair of mapping files. The assortment-to-planogram mapping files must be created before either of the two components can be used in ASO. Introduction of new assortments or planograms mandates the update and load of new mapping files before using the new assortments or planograms in ASO.

Here is the list of mapping files.

- POG to assortment mapping file: This file contains the POG hierarchy to assortment product mapping information, and it is used to identify which POG should be used for each product in the assortment.
- POG season-to-assortment mapping file: This file contains the POG season-to-assortment date mapping. Once the mapping from product to POG has been performed, a second pass examines this table to identify the specific season for the POG to use based on the assortment start date.

Assortment to POG Mapping Process

The assortment to POG mapping process has two main components. An automated mapping process assigns assortments from APO to POG sets. Afterwards the mapping, the user can access a UI to review and potentially override the mapping results.

Input Data

The following data is input for the mapping process.

- APO assortment, with associated data (SO_ASSORTMENT_STG)
 - Product list from APO assortment, at either the store level or the cluster level (SO_ASSORT_PRODUCT_STRCLTR_STG)
 - Stores within the assortment clusters, including effective assortment dates (SO_ASSORT_CLUSTER_STG and SO_ASSORT_CLUSTER_MEMBER_STG)
- Historical POG data, including POG hierarchy information (that is, POG department, POG category, POG sub-category) and POG seasonal attributes (SO_POG_STG)
 - Store to POG data. List of stores that have each POG, along with the effective dates for the POG at the store level (SO_POG_STORE_STG)

- Mapping tables
 - Product to POG node mapping table (SO_POG_ASSORT_MAPPING_STG)
 - Seasonal attribute mapping table (SO_POG_ASSORT_SEAS_MAPPING_STG)
- Merchandise hierarchy (Core data provided separately)

Automated Process

The automated process consists of the following steps:

1. Find one or more POG set(s) for each product.
 - The process iterates through the list of products associated with each cluster or store (depending on the assortment level) and performs a lookup to identify a POG node and demand spread factor for each product.
 - ASO searches the mapping tables for the lowest level merchandise hierarchy node that matches that product, starting at the product level, then, if the mapping is not there, looks for the product's parent, the product's grandparent, and so on. The highest level for escalation is the assortment's product category.
 - Exception case: If no matching row for the product (or ancestors) is found in the mapping table, the product will be considered unmapped and will be handled manually via the UI. Nothing else is done for that product by the automated process.
2. Find the seasonal attribute.
 - For each location within the assortment cluster and mapped POG node found for the products above, ASO performs a year-independent lookup for the seasonal attribute using the location's start date against the season's provided in the mapping tables. If no match is found, the process will use the current season. Current season is defined as the one that corresponds to the POG with the latest (most recent) effective_start date for any loaded POG within the POG node.
3. Find a POG for each one of the stores in the assortment cluster associated with the product. Stores within the same assortment clusters will very likely have different POGs assigned to them. This step finds those POGs
 - If a season is found, then for every store, the process looks for a specific POG within the POG node and seasonal attribute identified above.
 - If a season is not specified, then the application uses the Current Season to find specific POGs for each store.
 - If a season is specified in the mapping table but there is no such season in the historical POGs, then the process treats this as if the season was not specified and uses the Current Season.
 - Exception case: If none of the above techniques find a match for some stores, then the store will be unmapped and that information will be made available in the UI. The user can attempt to fix this by selecting some other POG node/season. Stores that remain unmapped after the POG mapping/manual overrides, will fall out and do not go through SO.

Mapping Errors

The following mapping errors can occur:

- Demand Spread Factor (DSF) out of range (1-100). The product is mapped to more than one POG set and the total DSF across these POG sets for the product is less than 1 or greater than 100.
- Unmapped Store. After matching the products from the assortment and looking for POGs within the identified POG sets, none of the stores in the cluster have a matching POG. This means that the assortment was mapped to a POG set for a product for which ASO cannot optimize any store.
- Unmapped Product. This is raised because mapping data either does not exist for the product because it was not provided or it was provided but the POG Set does not exist.

Replenishment Data

ASO consumes replenishment parameters at the product/location level and simulates a (s, S) inventory model to estimate service levels as a function of number of facings, which is one of the key inputs to the optimization engine. Lost sales are estimated as a part of service level calculations.

In addition to mapping and replenishment feeds, ASO also had another feed related to the Assortment and Space Optimization Product Stack Height Limit file. For details, see *Oracle Retail Advanced Science Cloud Services Implementation Guide*.

Optimization Science

This section describes the science behind the optimization algorithms used in the ASO application. It does not provide all the details of the algorithm. It does provide some guidance so that the user can troubleshoot and resolve issues quickly during an implementation.

Optimization Algorithm Overview

The optimization algorithm analyzes trade-offs and selects the best set of items from a given assortment. It then informs the user where to place these items, using how many facings, in a given planogram. The algorithm uses sophisticated mathematical modeling techniques to analyze all possible solutions to generate the best possible planogram. The optimization algorithm is provided an objective (for example, maximize total profits over all items in the assortment), business rules or restrictions, and the sales for a particular item i with k number of facings and at a particular elevation in the planogram (see [Sales and Inventory Model](#)). The algorithm analyzes the trade-offs between all possible solutions and picks the solution that gives the best value for the objective. All the restrictions imposed by the user are treated as required; that is, all possible solutions must satisfy that particular criterion.

To provide a sense of the analysis, here are some of the trade-offs that are analyzed.

- Is there sufficient space to pick all items and assign the maximum number of facings? Or instead, because space is limited, which items must be dropped?
- Should a particular item be included in the ultimate solution? Is there sufficient space? Or, since space is limited, does giving space to this item result in the dropping of a more profitable item?
- Does picking an item cause a conflict with the imposed business rules? For example, a user might determine that the item must be placed at eye-level, but there are other items that are equally profitable that do not have any elevation restrictions.

- Does dropping an item have a significant impact on the demand lost and thus revenues? Or instead, does dropping an item causes no significant loss in demand since some of the demand is transferred to other items that are selected?
- For an item that is picked, is it profitable to add another facing? Or instead, is it profitable to select another item? Does adding another facing for an item cause the item to be dropped since it cannot be fit anywhere in the planogram?
- Does adding another facing help in meeting the service level requirement? Or instead, can the minimum service level requirement never be met and that item must be dropped?

Sales and Inventory Model

At a high level, this model generates sales for a particular item i with k number of facings and at a particular elevation in the planogram. This model consumes the replenishment parameters and historical sales and computes the expected sales for a particular item i with k number of facings and at a particular elevation in the planogram.

The optimization that ASO performs makes decisions about which products to put on the shelves and how many facings to give each product. The basis for these decisions is demand information for each product and how much of the demand can be satisfied with a given amount of inventory for the item. Items with high demand may need more shelf space to hold the necessary inventory, and the ASO optimization balances that against other products that may have lower demand and also lower shelf space requirements.

The term "sales" refers to that portion of the product's demand that can be satisfied by a given amount of inventory. If inventory is sufficiently high, then all of the demand can be satisfied, and then demand is equal to sales. However, for lower levels of inventory, sales may be less than demand. For each product, optimization requires the relationship between the product's demand, sales, and inventory, since ultimately sales are what matters.

The feature of ASO called Sales and Inventory modeling performs the calculations of how much of a product's demand will turn into sales for a given level of inventory. The results of these calculations are fed to the optimizer. The sales units of a product are dependent on the following:

- Demand for the product
- Replenishment of the product, which determines how much inventory is available to meet demand

Replenishment of the product in turn depends on:

- Various replenishment parameters, such as the frequency of replenishment. These replenishment parameters are for each product-store combination.
- Inventory levels immediately after replenishment. This level is a separate calculation on its own, and is not a single parameter, which is why it is discussed separately.

Sales and Inventory Modeling Considers All Possibilities

Because the Sales and Inventory modeling occurs before optimization, the modeling does not know on which shelf or fixture the product will ultimately end up or how many facings the products will have. Thus, the modeling performs separate calculations for each fixture that can possibly hold the product and for each possible facing value. This is a type of what-if calculation, which gives the optimizer the

information about what the sales of a product would be if it were to be placed on a particular fixture with a particular number of facings. The optimizer requires this information for all possible fixtures and facings for each product so that it can make the best choices.

Inventory Levels After Replenishment

When replenishment of a product occurs, the portion of the shelf belonging to the product is re-stocked. The level to which the product is restocked (the order-up-to level) is partly determined by how much product will fit in the shelf space that the product will occupy. The calculation of how many units of the product will fit depends on the number of facings and how many units will fit in each facing. The latter is the units per facing capacity of the product, and depends on the product and the fixture the product is being placed on. For simplicity, units per facing capacity is called just "facing capacity".

Calculating the Facing Capacity for a Product/Fixture Combination

The calculation of facing capacity consists of several steps:

1. **Orientation.** The orientation of the product is how the product sits on the shelf. For simplicity, in the remaining text, the terms "height, width, and depth" refer to the dimensions of the product *after* the product has been oriented correctly on the shelf. The orientation can certainly affect the number of units of the product that can fit on the shelf, and thus it is essential for each product to have its correct orientation for the shelf-space capacity calculation to produce correct results.
2. **Stacking portion of Facing Capacity.** This is the portion of Facing Capacity where the products are stacked one unit atop another. The number of units that can be stacked depends on the height of the shelf, the height of the product, and the product's Max Stack Quantity, which takes precedence over the number of units allowed by the shelf height. This is then multiplied by the number of units of the product that will fit one behind the other within the depth of the fixture. This calculation also depends on the Above Gap and Behind Gap of the product.
 - Elements in this calculation: shelf height, product height, product's Max Units High, fixture depth, product depth, Above Gap, Behind Gap.
3. **Cap portion of Facing Capacity.** This is the number of units of the product that can be put on top of the stacking portion with the product in a different orientation. The Facing Capacity is then the sum of the stacking portion and the cap portion. The cap space has dimensions shelf depth x product width x remaining height, where remaining height is the height left over after the product has been stacked. The cap-units calculation simply determines how many units of product will fit into this cap space, using the product's Cap Height, Cap Depth, and Cap Width. The product's Max Cap Quantity is a maximum of how many units of product can be stacked in the cap space.
 - Elements in this calculation: product's dimensions, Cap Height, Cap Depth, Cap Width, Max Cap Quantity.

If the results from sales-inventory modeling show that a product has very low service levels, meaning the product is not receiving the inventory that it needs to serve its demand, check whether the facing capacity is sufficient to meet the demand of the product. This may involve checking the elements identified above to see if they are correct.

If the facing capacity is insufficient for the demand of the product, then service levels will be low, regardless of the values for the other inventory-related settings for the

product, since there is not enough room on the shelf for sufficient quantities of the product.

Maximum Capacity of a Product

In addition to the units per facing capacity, the other determinant of the order up-to level of a product is the product's Maximum Capacity. This number can be interpreted in three different ways, depending on the product's settings:

- The number of units of the product. This is the most straightforward, as the Maximum Capacity is itself just a count of units of the product.
- The number of case packs of the product. In this case, the case pack size of the product (given in the replenishment parameters for the product) is multiplied by the Maximum Capacity to convert case packs to units. (The case pack size is given in units of product.)
- The number of days of supply of the product. Here, the weekly demand for the product (from the replenishment parameters of the product) is divided by the number of days in a week, and multiplied by the Maximum Capacity to convert the Maximum Capacity into units.

Elements in this calculation: Maximum Capacity of the product, casepack size, weekly demand of the product.

The Maximum Capacity is a setting for the product and does not depend on the number of facings that the product is given.

For a given number of facings, the sales and inventory modeling determines the order up-to level by comparing two unit quantities:

- The product of the number of facings and the units per facing capacity
- The Maximum Capacity, after it is converted to units.

The smaller of these two values is the order up-to level used by the Sales and Inventory modeling for this number of facings. Thus, for example, if the fixture simply does not have enough space to hold the necessary inventory for a product, it will not help to increase the Maximum Capacity. If you attempt to increase the Maximum Capacity, and still the inventory for the product is not enough to meet demand, then it is time to check the Facing Capacity.

In general, if the results of the sales and inventory modeling show that a product is not receiving sufficient inventory, in addition to the elements affecting the facing capacity, also check the elements affecting the Maximum Capacity calculation.

Replenishment Parameters

The replenishment frequency can also affect whether the product receives sufficient inventory. The product must receive enough inventory to meet demand at least until the next review point. Thus, for example, if replenishment occurs only once per week for a product, the product's order-up-to level must be high enough to sustain at least seven days of demand for the product. If the fixture is too small to hold that much product (see the Facing Capacity section above), then increasing the replenishment frequency may help.

- Element to check: replenishment frequency

The transit time and trigger type are related replenishment parameters. The recommended trigger type is Demand Based, which means the order point, which is the inventory level that will trigger replenishment, is calculated by accounting for the demand and the transit time. With a trigger type Demand based, the order point is

high enough to leave enough units to meet demand until the replenishment arrives, and thus a larger transit time will mean a higher order point.

It is possible for the transit time to be so large that the order point is too high to be contained on the fixture. For example, suppose the Facing Capacity is lower than the order point. In this case, the product will continually run out of inventory, and it is necessary to either give the product more room on the fixture or decrease the transit time for the product.

- Element to check: transit time

Objective Function

The objective function specified by the user plays a major role in determining which solutions are considered best. For example, if the user specifies the objective function as maximize profit margin, then the algorithm will look for solutions that are superior on profit margin and not necessarily on the other KPIs such as revenue and sales volume. Sometimes, understanding why an item is included or dropped and why it is given so many facings might be as trivial as looking at the objective function contribution of that particular item to the objective function.

Note that the objective function not only considers contributions from products that are included in the final planogram solution but also from the dropped products, by using the demand transference module. The key idea here is that, for a dropped product, no sales are lost; some sales may be transferred back to other substitutable products. The Demand Transference module generates the demand transference values that may not be consumed *as is*. ASO provides flexibility for the ASO users to dampen the demand transference values if they are deemed too high.

There are a few choices for specifying an objective.

- 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. This objective tells the optimization to fill the planogram with items that results in best possible revenue but at the same minimize the inventory units carried. As one can imagine, carrying too much inventory will result in higher revenues but at a higher cost of excess inventory. This metric lets the user strike a balance between these two metrics.
- 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 Category Manager.
- 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 Category Manager.
- 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 Category Manager.

- 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 Category Manager.

Weighted metrics use IPI values provided by the Category Manager. It is generally expected that IPI values are positive and hence negative or 0 values for products will be flagged as warnings since it can result in products getting dropped.

Constraints

If the objective function focuses on the best possible solution, then constraints work in the opposite direction, by restricting the set of possible solutions. For example, if the objective function says to select the most profitable item A and assign the maximum number of facings (for example, 15), a constraint on item A may say it is not possible to have more than three facings for item A in order to provide a minimum number of facings to the items of the other brands.

Optimization enforces all constraints as required except for Product Family Group constraints; that is, it finds all the solutions that satisfy all the constraints except for Product Family Group constraints that are specified by the user.

Sometimes, inadvertently, the user might specify conflicting constraints that can result in no solution or unexpected solutions. Often, a resolution can be found by just understanding the implications of individual constraints. Some examples of how to analyze the constraints:

- Why was item B dropped? Is the Average Weekly Demand too low?
- Why did item C receive so many facings? Is this item part of Match Facings group? Or what is the minimum number of facings rule?
- Why is an item placed at this elevation? What are the elevation restrictions on this item?
- Why is an item in second bay and not in the first bay? Is there a blocking rule in place?

More often than not, the user must analyze the interplay between two or more constraints to see why an item is included or dropped and why it is assigned so many facings, and why it is placed in a particular bay and at a particular elevation. A Sanity Checker is provided in the ASO Actions menu that can be used to identify some of the logical conflicts that have arisen due to the definition of the constraints.

Product Family Group Constraints

Among all the constraints specified by the user, Product Family Group constraints are enforced softly. To elaborate on the PFG constraints: Product Family Group constraints are defined using the sort attributes (up to three attributes) mentioned in the Visual Guidelines screen. Note that these Sort Attributes are also used in laying out items within a shelf. These constraints tell the optimization that products belonging to a particular combination of sort attribute1, sort attribute 2 and sort attribute 3 are to be kept together or "close enough". An example of a Product Family Group is Deodorants belonging to a Brand, Size and Scent to be placed together.

It should be noted that Product Family Group constraints are different from Same Shelf constraints or Visual Guidelines as follows:

- Product Family Group constraints are a relaxation of same shelf constraints - products do not have to be kept in same shelf but in nearby shelves.

- Product Family Group constraints are different from Visual Guidelines as they do not enforce any order between the groups.

Some recommendations on using Product Family Group Constraints:

- ASO provides three tools to help in achieving the intended layout: Same Shelf, Visual Guidelines, and Product Family Group. Note that Sorting Attributes are dual purpose as they also help in laying out items within a shelf. [Figure 8-3](#) explains the three concepts.

Figure 8-3 Blocking, Product Family Group, Same shelf, Sorting Illustrated



- The user should make use of different tools provided to achieve the intended layout. For example, the user can define the order using the Visual Guidelines and further keep products close enough within a block, and the user can define Product Family Group constraints. PFG constraints provide the ability to keep the products together within a block (for example, vertical block + horizontal block).
- An example of a redundant PFG strategy would be to use same attribute to define visual guidelines (for example, use brand to define primary vertical blocking) and for product family group (for example, use Brand for defining Product Family Group). In such a case, Visual Guidelines will keep products for each Brand together, and create blocks using the order specified. Product Family Groups are redundant here.

Diagnosis of Visual Guidelines

Visual guidelines are imposed as a set of constraints that restrict the placement of an item. Specification of the visual guidelines contains the following:

- Blocking strategy
- Primary blocking attributes (up to two attributes)
- Secondary blocking attributes (up to two attributes)

A blocking strategy can have a design in which primary blocking is vertical and secondary blocking is horizontal or primary blocking is horizontal and secondary blocking is vertical. When the primary blocking is vertical, the application allows the user to specify horizontal blocking attributes by each vertical block. Similarly, when primary blocking is horizontal, the user can specify the vertical blocks by each horizontal block. Note that both strategies allow the user to specify additional horizontal blocks called top or bottom, which allows the user to put items in the top shelves or bottom shelves.

The following illustration clarifies these strategies. Assume that there are four brands, A, B, C, and D. Users also have the ability to create a merged block, as shown in the following example. Brand C and Brand D are merged to create a combined vertical block, as the user believes that these are premium brands that do not have as many SKUs as Brand A and Brand B. In the first scenario, the user wants the items placed in the following order. From left to right, the user wants to see Brand A, then Brand B,

and then Brand C and D. Further, from top to bottom, the user wants to distinguish Brand A by size (since the user believes the size plays a major part in customer purchasing decision for Brand A). For Brand B, the user wants the flavors arranged from top to bottom. The (Vertical, Horizontal) strategy is shown in [Table 8–1](#). Alternatively, the user can decide that size is the most important attribute across all brands and specify that horizontal rows of shelves should be of same size. The (Horizontal, Vertical) strategy is shown in [Table 8–2](#).

Table 8–1 Vertical, Horizontal Strategy

Brand A	Brand B	Brand C
12 oz	Vanilla	
24 oz	Chocolate	
36 oz	Strawberry	Brand D
	Multipack	

Table 8–2 Horizontal, Vertical Strategy

Brand A	Brand B	Brand C and D
12 oz	12 oz	12 oz
24 oz + vanilla	24 oz + vanilla	24 oz + vanilla
24 oz + others	24 oz + others	24 oz + others
36 oz	36 oz	36 oz

These two strategies provide the flexibility for the user to define which takes precedence, vertical or horizontal, using only attributes. Note that the user does not actually indicate which shelves or elevation a product should be placed at. Instead, the optimization decides the shelves or elevations, based on the order of items specified by the user, using attributes. When the primary attribute is vertical, the optimization tries to follow the order specified for vertical blocks and then tries to place products by the horizontal blocks defined. In contrast, when the primary attribute is horizontal, the optimization tries to generate solutions that adhere to the horizontal blocking order and vertical blocking order that is common to all horizontal blocks. In the example above, for the former, Brand A has three horizontal blocks, Brand B has four horizontal blocks, and the merged block has two horizontal blocks. For latter, there are four horizontal blocks and three vertical blocks for all horizontal blocks.

An example of solution to the primary as vertical and secondary as horizontal strategy stated in Table 6-1 is shown in [Figure 8–4](#).

Figure 8–4 Primary as Vertical

Bay 1	Bay 2		Bay 3	Bay 4	
Brand A + 12 oz	Brand A + 12 oz	Brand B + Vanilla	Brand B + Vanilla	Brand B + Vanilla	Brand C
Brand A + 12 oz	empty space	Brand B + Vanilla	Brand B + Vanilla		
Brand A + 12 oz	Brand A + 24 oz	Brand B + Chocolate	Brand B + Vanilla	Brand B + Vanilla	Brand C
Brand A + 24 oz		Brand B + Chocolate	Brand B + Chocolate	empty space	Brand C
Brand A + 24 oz	Brand A + 24 oz	Brand B + Chocolate	Brand B + Chocolate	Brand B + Chocolate	Brand D
Brand A + 36 oz	empty space	Brand B + Chocolate	Brand B + Chocolate	Brand B + Chocolate	Brand D
Brand A + 36 oz	Brand A + 36 oz	Brand B + Strawberry	Brand B + Multipack	empty space	Brand D
Brand A + 36 oz	Brand A + 36 oz	Brand B + Multipack		empty space	Brand D

An example of solution to the primary as horizontal and secondary as vertical strategy stated in Table 6-2 is shown in Figure 8-5.

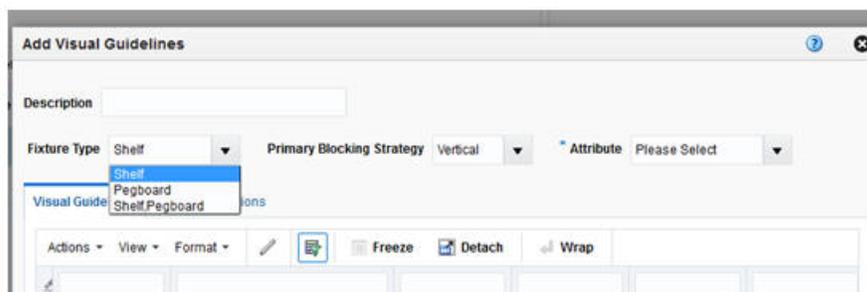
Figure 8-5 Primary as Horizontal

Bay 1	Bay 2		Bay 3	Bay 4	
Brand A+12 oz	Brand A +12 oz	Brand B+12 oz	Brand B +12 oz	Brand B +12 oz	Brand C&D +12 oz
Brand A+12 oz	empty space	Brand B+12 oz	Brand B +12 oz		
Brand A+12 oz	Brand A + 24 oz +Vanilla	Brand B + 24 oz + Vanilla	Brand B + 24 oz + Vanilla	Brand B + 24 oz + Vanilla	Brand C&D + 24 oz + Vanilla
Brand A + 24 oz +Vanilla			Brand B + 24 oz +Vanilla		
Brand A + 24 oz +Others	Brand A + 24 oz +Others	Brand B + 24 oz + Others	Brand B + 24 oz +Others	empty space	Brand C&D + 24 oz +Others
Brand A + 24 oz +Others	Brand A + 24 oz +Others	Brand B + 24 oz + Others	Brand B + 24 oz +Others	Brand B + 24 oz + Others	Brand C&D + 24 oz +Others
Brand A + 36 oz	Brand A + 36 oz	Brand B +36 oz	Brand B + 36 oz	empty space	Brand C&D + 36 oz
Brand A + 36 oz	Brand A + 36 oz	Brand B +36 oz		Brand B + 36 oz	Brand C&D + 36 oz

Visual Guidelines provide flexibility to model various layouts, and yet improper specification of visual guidelines can result in empty spaces or sparsely populated optimal planograms. The next few paragraphs provide some recommendations on how to use Visual Guidelines:

1. The user can specify visual guidelines by fixture type or common to all fixture types as shown in Figure 8-6. The user must define separate guidelines for different fixture types since products or product attributes are not common to both fixture types.

Figure 8-6 Add Visual Guidelines



2. In general, adding visual guidelines restricts the optimization, so the user should create a run without any visual guidelines. Later, the user can create runs with a variety of visual guidelines. Comparing the no-VG and VG runs provides the impact of imposing visual guidelines, in terms of revenue, sales units, and so on.
3. It is preferred that when adding visual guidelines, the user should bear in mind to start with a simple blocking strategy (for example, primary only) and to check the results and increase complexity in terms of additional and secondary strategies.
4. As stated above, there are two block strategies that are supported by ASO: (a) Primary is Vertical, Secondary is Horizontal (b) Primary is Horizontal, Secondary is Vertical. In case of shelf fixtures only, each horizontal block will need at least one shelf to satisfy and thus it is a restrictive strategy than the former. Further, the second strategy, in case of shelf fixtures, needs at least one shelf for each horizontal block and thus, the number of horizontal blocks is limited to number of shelves in the bay. For example, it is pointless to define ten horizontal blocks for strategy defined in Table 8-1 as each bay has at most eight shelves

5. It is essential that the user does a preliminary analysis on the product counts and KPIs like Sales Units to see how well the blocks are defined. Very few products will result in thin blocks or empty spaces; low KPIs typically result in sacrificing the block for another block with better products. ASO provides Export to Excel on Product Constraint tab so that the user can check product counts. Here is an example of the process:
 - a. The user can select Custom Attributes (up to three attributes at a time) and click the **Show Attributes** button.
 - b. The user can then click **Export to Excel** and download the products with attributes information into Excel as shown in [Figure 8-7](#).

Figure 8-7 Export to Excel



- c. An example of the analysis is shown below: the user would like to define blocks using two attributes: Segment and Pack-size.

Exporting data to Excel and performing a pivot on the attributes gives us the Table 6-3. Attribute values highlighted in green color are Segment and others are Pack-size values.

Notice that the Mouthwash segment seems the largest in terms of product counts whereas all other segments have few products. Given the disproportionate distribution of product counts, it might be better to start with only two blocks, one for Mouthwash and rest all in others. This gives even distribution of products in all the blocks.

Next, the user needs to check the value of each block. It is possible that some blocks are not valuable and optimization can trade that block for another block with higher value. ASO Visual Guidelines screen provides the ability to see the KPIs like Sales Units, and Revenue. This gives the user the ability to understand the value of each block.

Figure 8–8 Product Counts for Blocking Strategies

Block #	Attribute 1 + Additional Attribute	No. of Products	No. of Products in block		Block #	Attribute 1 + Additional Attribute	No. of Products	No. of Products in block				
Block 1	Value 1		22		Block 1	Value 1		44				
	Vanilla	9				Vanilla	9					
	Strawberry	5				Strawberry	5					
	Chocolate	1				Chocolate	1					
	Value 2					Value 2						
Vanilla	1	Vanilla	1									
Chocolate	6	Chocolate	6									
Block 2	Value 3		7			Value 3	Pack		3	Value 4	Pack	3
	Pack	3					Single		1			
	Single	1					Pack		3			
Block 3	Value 5		8	Value 5	Pack	4	Value 6	Cinnamon	1			
	Pack	4			Single	4						
Block 4	Value 6		3	Cinnamon	Mint	2	Value 8	All Pack Sizes	4			
	Cinnamon	1			Mint	2						
Block 5	Value 7		58	Value 7	Vanilla	11	Block 2	Vanilla	11			
	Vanilla	11			Strawberry	24		Strawberry	24			
	Strawberry	24			Chocolate	18		Chocolate	18			
	Chocolate	18			Others	5		Others	5			
Block 6	Value 8		4	All Pack Sizes	All Pack Sizes	4						
	All Pack Sizes	4										
Poor Blocking Strategy					Better Blocking Strategy							

Finally, the user needs to decide whether to define primary as vertical or primary as horizontal. In general, the rule of thumb is to define primary as vertical since it is less restrictive than primary as horizontal.

After optimization, the user can quickly check how the blocks are formed. Optimization tries its best to satisfy all constraints and provide optimal amount of space for each product. This can result sometimes in empty spaces within each block. In such a case, the user should revisit blocking strategy; perhaps few smaller blocks can be merged into the bigger block.

Product Groups

Product groups provide a set of constraints for the optimization that specify the relation between any pair of items. For example, retailer has to pair a high-margin product with low-margin product or retailer has to match number of facings for shampoos and conditioners. ASO provides a few variants on these constraints.

Table 8–3 Product Group Constraints

Constraint	Description
At Least	At least m items must be selected in the final assortment.
Exactly	Exactly m items must be selected in the final assortment.
At Most	At most m items must be selected in the final assortment.
All or Nothing	If one item from this group is selected, then all other items in this group must be selected in the final assortment.
Match Facings	Whatever items are selected, all the items selected must be given same number of facings.

Table 8–3 (Cont.) Product Group Constraints

Constraint	Description
Same Shelves	Whatever items are picked, they must be placed on the same shelf in the final planogram (only applicable for shelf fixture type).

Validation Tool (Sanity Checker)

This feature helps the user to identify logical conflicts that may occur because of the constraints imposed and provides guidance on how to resolve the issues. The validation feature generates two types of alerts: Error and Warning. Error is generated when the set of constraints result in no solution. Warning is informative in nature and does not necessarily result in no solution. The user should try to understand why an error or warning is generated and examine the resolution provided.

Diagnosis of Dropped Products

The Dropped Products tab provides information on which products are dropped and the high-level reason why they are dropped. Coupled with Validation Errors/Warnings, the user can get an idea on why a product is dropped. The user should try to understand why an error or warning is generated and examine the resolution provided.

Further, the user should use Export to Excel to understand the high-level issues for dropped products. They are categorized primarily as follows:

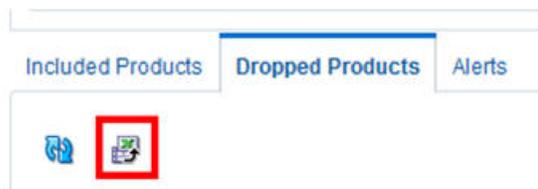
- Do Not Include.
- No Solution. When the optimizer cannot satisfy the constraints specified, then it does not return any solution. For example, if all items need minimum of ten facings and are mandatory. This is not physically possible on the POG with such constraints and hence optimizer will return no solution.
- User Constraints. When a product cannot fit into any shelf due to its geometry or display orientation. To identify which products, the user should refer to Validation Errors/Warnings panel.
- Missing Sales and/or Forecast Data. As stated, a product does not get included because there is missing data. To identify which data elements are missing, the user should refer to Validation Errors/Warnings panel.
- Invalid Sales and/or Forecast Data. As stated, a product does not get included because there is invalid data (for example, price = 0). To identify which data elements are invalid, the user should refer to Validation Errors/Warnings panel.
- Minimum Service Level. When a product cannot satisfy the minimum service level (for example, 90%) then the product is dropped from the POG. Refer to the Sales and Inventory Model section of the Implementation Guide on how it can be improved.
- Solver Choice. This can be due to the objective defined and contribution of the item compared to the space needed for that item. Some examples of why Solver (or Optimization) drops the items:
 - Item has very low demand; the user should check why the demand is low.
 - Item got dropped due to low service levels since it needs many facings to meet minimum service level requirement.

- Item got dropped because optimization compared the item's contribution to the objective function to the space needed and decided that it is best if that space is given to another product.

In general, the user can conduct the diagnosis of the errors/dropped products as follows:

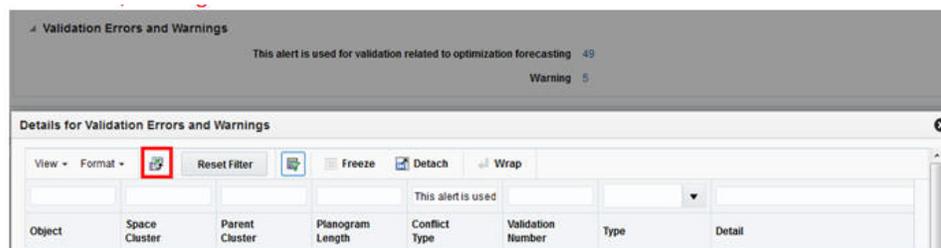
1. Go to Dropped Product panel and use the Export to Excel functionality. Looking at the Reason column, the user can determine why a product was dropped.

Figure 8–9 Dropped Products



2. To identify further specifics about the errors, the user can click **Validation/Errors** and click the hyperlink and use the Export to Excel functionality. This gives specific product-cluster errors/warnings.

Figure 8–10 Validation Errors



Checklist for Optimization Results Diagnosis

When an SR is submitted, provide the following:

- Issue description and expected behavior. As the below details are gathered, be sure the selections in the UI match the expectations and that the constraints were not specified for the wrong cluster.
- Validation Errors/Warnings Excel (using Export to Excel feature)
- Dropped Products Excel
- Visual Guidelines Excel (using Export to Excel feature in Visual Guidelines UI)
- Analysis of Visual Guidelines Excel (for example, Product Counts by Attributes from the Visual Guidelines Primary Block Details section)
- VPOG Screenshot (Using All Visual Guidelines display option)
- Replenishment Data
- Run Number, Space Cluster ID

FAQs

Here are some answers to typical questions that can arise.

An item cannot meet the service level with the defined facings. Should the solver give more facings to meet the minimum service level restriction?

Not necessarily. This depends on the feed from the sales and inventory model, which indicates what the possible service level is for a given number of facings. If it indicates that adding any number of facings will not necessarily help, then the solver drops that item if it is not mandatory or returns no solution if it is mandatory.

Why are empty spaces present in the planogram?

The goal of the optimization is never to fill up the entire space. Empty spaces can arise for a variety of reasons. This can occur if more than sufficient space is available for all the products. In this case, the solver places the products in more than one elevation and location. Alternatively, many constraints may be in play at the same time and one of them can cause some products to be dropped.

The solution generated does not seem optimal, since I can generate a better solution. Why?

Such a situation does not arise normally. It can sometimes happen because the optimization searches for the best possible solution within an allocated time. The allocated time can vary, depending on how many jobs are running and how much processing power is available for this job. To negate such uncertainty, user can increase the time allocated for the optimization from the default value of 180 seconds.

Why are items not sorted across the shelves?

The sorting functionality sorts items within a particular shelf. It does not sort across shelves as this negates the optimal placement determined by the optimization and can sometimes result in infeasibility. For these reasons, sorting is not allowed across shelves. The user is advised to examine the visual guidelines to obtain the desired structure.

Why are the horizontal blocks arranged in a zigzag pattern instead of being truly horizontal?

This can arise for the shelf fixture type. If the bays are actually identical, then the optimization can generate horizontal blocks. However, most often, the shelves are not aligned perfectly across bays and so the optimization is asked to look for almost horizontal blocks. In the case of the pegboard fixture type, the optimization can generate actual horizontal blocks, as it is not limited by shelf elevations.

Why is an assortment not ready for optimization?

An assortment may stop at the pre-optimization stage and not be ready for optimization because errors were found either in the assortment data during data loading, during the checking of integration with other data sets (such as missing replenishment, display style, and so on), or during the mapping process and requiring user review. Mapping errors can be examined in the Assortment Mapping UI. In addition, assortments that are ready for SO may not be available for optimization if the POG set to which they are mapped has multiple assortments mapped to it and at least one of those assortments is not ready for optimization due to mapping errors. In such cases, the POG set is meant to optimize multiple assortments; having errors in at least one of them prevents the whole POG set from becoming available for optimization.

Can I copy an old run from an assortment for the same category rather than a new assortment and use it against the new assortment (most of the products/stores should be the same)?

No, a run's scope is limited to the POG set and specific assortments used in that run. Once the assortments from that run are exported and finalized, there is nothing more to do with it. The user can still copy the run and perform some test scenarios, but the data used in the run corresponds to the old assortment and there is no relationship with newest data from more recent assortments.

I just sent some POG and mapping data to fix some of the errors in an assortment. However, the assortment still shows that mapping results require review. Why wasn't the newest data used to fix the mapping errors?

The assortment to POG mapping process runs every time an assortment is loaded into ASO and every time the direct assortment data is updated. In this case, the data that was delivered is not part of the assortment, so there is not an automatic mapping process for the assortment. An authorized user must use the Assortment Mapping UI and manually request the Re-map action from the drop-down menu. The mapping process will then run again. It will use the latest data available within the application, and, if all the data elements provided fix the old mapping errors, the assortment will map correctly and will be upgraded to the ready for optimization status.

Customer Segmentation

This chapter provides details about the use of the Customer Segmentation application.

Overview

Customer Segmentation (CS) lets users create segments of customers based on common attributes, such as customer demographics, in order to help a retailer manage merchandise and sales strategies in a targeted way. Segments can help retailers understand the types of customers who shop in their stores and gain insight into their typical shopping patterns. This understanding can help retailers target specific customers.

The application optimizes segments in order to determine the minimum number of segments that best describes the data used in the analysis and that best meets the business objectives defined when the segments are designed. What-if scenarios and ranking can be used to compare how cohesive and well separated the segments are in each scenario as the number of segments is increased. The application uses scoring to indicate which segments fall below defined thresholds and may require manual intervention. Business Intelligence graphics illustrate the patterns in the data and the attributes that are important in each segment.

The key features available in CS are:

- Recommendations are provided for the important attributes to use in creating segments.
- Segments can be created on continuous attributes such as sales performance as well as discrete attributes such as customer gender.
- Configurable segmenting criteria such as demographics and RFM (Recency and Frequency Measures) are provided.
- Recommendations are made for the optimal number of segments and the scores for each segment. These are based on the quality of the segments: how cohesive and well separated the segments are.

Data Requirements

Customer Segmentation relies on following data, and it uses ETL to load the data.

Table 9–1 *Data Requirements*

Object	Granularity	Required/Optional
Hierarchy	Product, Location, and Fiscal	Required

Table 9–1 (Cont.) Data Requirements

Object	Granularity	Required/Optional
Customer Attribute	Demographic or User Defined	Required for demographic segmentation
Sales Transaction	Customer-identified, Product, Date, Transaction ID	Required for RFM segmentation
Alternate Hierarchy	CM Group or Trade Area	Optional

Multiple Hierarchies and Support

Customer segments can be generated for the following combinations of hierarchies.

Product Hierarchy

One of the following hierarchies can be used:

- Core merchandise hierarchy
- Alternate hierarchy

The level at which segments should be created can be configured. The user interface is used to create the specific node identified. This allows the creation of different segments of customers for different geographic regions. For example, different segments can be defined for Canada vs. France.

Location Hierarchy

One of the following hierarchies can be used:

- Core location hierarchy
- Alternate hierarchy (optional)

The level at which segments should be created is configurable. The user interface is then used to create the specific node identified. This allows the creation of different segments of customers for different geographic regions. For example, different segments can be defined for Canada vs. France.

Calendar Hierarchy

This includes:

- Core fiscal calendar hierarchy (week, month/period, quarter, half, year)
- Gregorian calendar (week, month, quarter, half, year). Leverages a start and stop date (day level date range)
- Planning period. Leverages alternate hierarchies, including planning periods, buy periods, and defined holiday time periods such as back to school and Fourth of July. This is optional.

Segments can be defined for any of these three calendar hierarchies (the segment effective period). Note that the source time period for historical data only uses the core fiscal calendar hierarchy. A configuration permits data aggregation at either the fiscal period or fiscal quarter, so the user can select any level that is at that level or above in the user interface.

Supported Segment Criteria

In customer segmentation, the segment criteria consist of a set of attributes that define customer segments. These attributes can be either discrete or continuous. A group of these segments is called "Segment by." For example, demographic data, such as income and gender, and store properties, such as store and formal, can be grouped into a Customer Demographic Segment by.

Here is the list of the default Segment bys that are supported.

Customer Demographics

Customers are segmented based on the similarity in the values of the various customer attributes. Examples include gender, income, educational background, age, and range. Additionally, it is possible to use several user-defined numeric or discrete attributes.

RFM and Customer Behavior

Customers are segmented based on attributes that are comprised of aggregate metrics regarding their purchase behavior. Examples include the number of purchases, amount of sales, average basket size, and share of products purchased while being promoted. Retailers can analyze the segment composition and related business intelligence in order to better understand the customer shopping behavior associated with the segments.

Category Purchase Behavior

Customers are segmented in a similar manner to RFM and Customer Behavior, with the difference that these attributes are calculated for a selection of the most important product categories. The category level used here is configurable. It helps enable segmentation by product category purchase behavior.

Customer Segmentation Attributes

Segment by uses a collection of attributes, including demographics, purchase behavior, product purchase behavior, product profiles, and user defined. Each quarter, Oracle Retail Advanced Science batch processing creates new versions for each location at the configured level of the location hierarchy. During this process, attributes are summarized and their data is analyzed for usefulness for segmentation. An attribute can have a different level of usefulness for each of the different versions. For example, if the majority of customers in the Canada location provide a gender attribute, and the majority of the customers in the France location do not provide a gender attribute, then gender can be used in Canada, but not in France.

Furthermore, within the same location, it is possible for an attribute to be considered useful for the most recent quarter, while in the prior quarter it was not useful because there were insufficient values available during that time. The aggregate statistics about the attributes for a version can be seen using the Explore Data screen after selecting the segment criteria. Segments created in the previous quarter have different statistics than those that are created during a different quarter.

Demographics

Demographic segmentation relies on customer attributes that are loaded into RI. The set of attributes used from RI's customer dimension is fixed. If alternative attributes are needed, see [User Defined](#). Once loaded into RI, they can be used by Customer Segmentation. They define details about each customer that can be used for creating customer segments using those values. During the time when a new version is created,

only attributes that have 15 or fewer discrete values are used. Attributes with a higher number of discrete values are not considered for customer segmentation.

Purchase Behavior

Customer segmentation uses a fixed set of sales transaction metrics, which are obtained from sales transactions identified by a customer ID. The values include Sales Quantity, Sales Retail \$, Gross Margin \$, Promotional Sales Quantity, Promotional Sales Retail \$, Promotional Gross Margin \$, Number of Transactions, Average Transaction Count per configured period, SKU count, Transaction Basket Size, and Promoted Sales Share of Total Sales.

Product Purchase

Customer segmentation uses a selection of the top categories to portray the shopping patterns for each customer for some product categories. Each time a version is created, the categories with the highest amount of sales are picked as the categories for which Product Purchase based assessments are done. These attributes are similar in concept to the Purchase Behavior segment; however, these are specific to the top categories for the location associated with the version. The attributes include Sales Quantity, Sales Retail \$, Gross Margin \$, Promotional Sales Quantity, Promotional Sales Retail \$, Promotional Gross Margin \$, Number of Transactions, Average Number of Transactions, SKU Count, Transaction Basket Size, Promoted Sales Share of Total Sales, Average Basket Sales \$, and Average Basket Gross Margin \$.

Product Profile

Customer Segmentation uses a fixed set of category-based sales profile values. For the same set of top categories described in [Product Purchase](#), the share of the customer's total purchases for the category is calculated. The share of the Promotional Sales Retail \$ of the customers total sales is also calculated.

User Defined

For any attribute that is available for the customer, but is not accounted for in the default set of attributes, there are provisions for loading a set of customer or user defined attributes into RI. These attributes can be either numeric values or discrete values. If the attribute value is numeric (such as a zip code), but must be treated as discrete rather than a ranged numeric value, then the attribute must be loaded to an appropriate text attribute column in RI. Any attribute that has more than 15 distinct values will not be used by the segmentation process.

Once an attribute is defined, it is possible to adjust the configuration data in the database to assign a more context-suitable name for the attribute. This enables the user interface to identify the attribute as a specific attribute, and not just as a generic Custom Text Attr or Custom Number Attr.

Configuration Process

Default configuration occurs during the installation and upgrade. The configuration process is responsible for enabling or disabling any attributes in the application. This ensures that the desired attributes are available for use during the segmentation process.

- All attributes are enabled by default.
- Any discrete attribute that has more than $n=15$ attributes values is not configured by default. Note that the value of n is a configuration and can be modified at the time of deployment.

- The UI formatting of each attribute is identified based on the data type of the attributes and by the name of the attribute.

The following configurations may require manual overrides if the default configuration is not acceptable or data is not available.

Table 9–2 Manual Overrides

Name	Description
Enable or disable Segment by	Disable Segment by. For example, if there are no customer demographic attributes, then the Customer Demographics segment by can be disabled.
Enable or disable attributes	Enable an attribute to be considered for segmentation. For example, if there is no intention on loading user defined attributes, they can all be disabled so that when new versions are created, no processing time is spent analyzing those attributes.
Change UI formatting	Change formatting associated with the attributes such as label, decimals, percent, and currency. These are configurations for each attribute, and do not rely on XLIF entries.
Location hierarchy level	If customer segments are desired for each location at a given location level, then the configuration can be adjusted so that all processing is done at that level. This also requires an adjustment of the approval level so that it allows the segments to be approved.
Outlier rule	Change default outlier rules for a Segment by. By default, the distance from centroid rule is enabled. See section below for other supported outlier rules.

Table 9–3 Enable or Disable Segment By

Segment By	Description	Example	Enable
Customer Demographics	Segment customers using demographic values	Age range, income, gender	Enable if consumer demographic attributes are available
RFM and Customer Behavior	Segment stores using location attributes	Income, climate, size, store format	Enable if location attributes are available for each store
Category Purchase Behavior	Segment stores using sales metrics	Sales revenue, sales unit, gross margin	Enable if retailer wants to segment stores using performance metrics

Attribute Preprocessing

Before customer segments are created, the available customer data must be preprocessed in order to identify the sample of customers to use for segmentation and to determine which attributes are the most beneficial for use while creating segments. A few configurations can be manipulated to help improve this process. The following table defines some configurations that can be adjusted to control how attributes are used by the system. These values can be manipulated in the table CIS_TCRITERIA_ATTR.

Table 9–4 Attribute Configuration

Column	Description
DELETE_FLG	Set value to Y to prevent an attribute from being used by any processing.
SAMPLE_FLG	Set value to Y so that the attribute to be used has a stratified sample of customers. This should help ensure that an appropriate selection of customers are represented in the sample. Up to three attributes can be set to Y. If no attributes are configured with a Y value, then a random sample of customers will be used.
DISPLAY_FORMAT_ID	This can be adjusted to use different display formats, as defined in RSE_DISPLAY_FORMAT.
NUM_BUCKETS	Allows a different number of buckets for use when showing summary metrics for a numeric attribute.
ATTR_IMP_FLG	When set to a Y value, the importance of this attribute is analyzed for usefulness during segmentation. If an attribute will never be used for segmentation, setting the value to N will exclude it from the attribute importance calculations.

Another configuration that can be adjusted is the CIS_BUS_OBJ_TCRITERIA_ATT_XREF.VALIDATING_ATTR_FLG. For any attribute that has a Y value for this column, the attribute is shown in the Insights portion of the UI. If it is determined that an attribute should not be displayed, then this value can be changed to a value N so that it is excluded from display in the Insights results.

Segmenting Approach

Customer segmentation uses Oracle Data Mining for the creation of its segments. The k-means approach that is used results in the creation of segments in a hierarchical manner. The process automatically determines which attribute is the best attribute to split into an additional segment. This process is continued until the desired number of clusters has been achieved.

Customer Segment Store Profile Generation

Customer Segmentation calculates the sales share of customer segments for each store. These store profiles can be generated by the user for the approved customer segments from the user interface. They can then be consumed by RI to generate business reports and Store Clustering to generate customer centric store clusters.

Preprocessing

As part of preprocessing when a version is created, Customer Segmentation first filters the customer and then samples the customer.

Filter Customer

The filtering step allows the implementer to set the following conditions in order to remove outlier customers during different phases of Customer Segmentation,

Table 9–5 Filter Step Conditions

Filter Rule	Operation
Fake Customers	This rule discards customers if the number of transaction per day exceeds the value x (the max daily transaction threshold to identify a fake customer - default 10)
Discarding Customers	Keep customers where customer sales fall between min (sales or # of transaction) and max (sales or # of transaction). Minimum/Maximum percentile for amount of sales transactions. (default .001) Minimum/Maximum percentile for number of sales transactions. (default .001)
New Customer Rules	Include new customers if there is enough existing sales history for the customer. Use configurations that control the min/max values for average transaction count and average transaction amounts used to include new customers: Minimum/Maximum percentile for amount of sales transactions. (default .001) Minimum/Maximum percentile for number of sales transactions. (default .001)
Top Categories	A configuration that can limit the number of categories (n) that a user picks. The system picks it based on the default - Sales Revenue of the category to be used to calculate the top categories. Allowable values are limited to SLS_AMT, SLS_QTY, PROFIT_AMT, and SHARE variations of the same (i.e., SLS_AMT_SHARE).

Sample Customer

The sampling step allows the implementer to enable and adjust sampling customers.

Table 9–6 Sample Step Conditions

Sample Rule	Operation
Target Sample	Set the value to Y so that the attribute to be used has a stratified sample of customers. This should help ensure that an appropriate selection of customers is represented in the sample. Up to three attributes can be set to Y . If no attributes are configured with a Y value, then a random sample of customers will be used. This can be adjusted by changing CIS_TYPE_CRITERIA in Data Management UI.
Sample Size	The sampling percentage for Customer Segmentation. This can be adjusted by changing RSE_CONFIG in the Data Management UI.

This chapter describes the major configuration points in Oracle Retail Advanced Science, including:

- [User Interface Authentication and Authorization](#)
- [User Management Configuration: Configuring Users and Roles](#)
- [Configuration](#)
- [Internationalization](#)
- [Market Basket Insight Configurations](#)

Note: Since MBI is distinct from the other applications, much of what is described here is not applicable for MBI. For clarity, MBI implementation, configuration, operations and data model are described separately in [Chapter 12, "Market Basket Insights."](#)

User Interface Authentication and Authorization

Note: For more information, see *Oracle Retail Advanced Science Cloud Services Administration Guide*.

For authorization, Oracle Retail Advanced Science modules have been built with role-based access. Access to application user interface components is done by assigning application roles. Application roles are defined as part of the application and deployed as part of the installation process. Application roles are mapped to enterprise roles during initial environment provisioning. Enterprise roles exist as LDAP groups in OID. Refer to the *Oracle Retail Advanced Science Cloud Services User Guide* for the definition of standard user roles.

User Management Configuration: Configuring Users and Roles

This section provides detailed instructions on setting up enterprise-level user management using Oracle WebLogic 12c with Enterprise Manager. The user management configuration is handled using the WLS Console and the WLS Enterprise Manager (EM).

User Roles

Oracle Retail Advanced Science supports the roles listed in [Table 10–1](#). For more information, see *Oracle Retail Advanced Science Cloud Services Administration Guide*.

Table 10–1 ORASE Default Enterprise Roles

Application Module	Default Enterprise Roles	Corresponding Application Roles
CDT	ANALYTIC_EXPERT_JOB	Customer Decision Tree Duty
DT	ANALYTIC_EXPERT_JOB	Demand Transference Duty
ASO	CATEGORY_MANAGER_JOB	Category Manager Duty
	SPACE_PLANNER_JOB	Micro Space Optimization Analyst Duty
	MERCHANDISING_ANALYST_JOB	
	SPACE_ADMINISTRATOR_JOB	ASO Administrator Duty
	FORECAST_MANAGER_JOB	Analytic Super User Duty
CS	ASSORTMENT_PLANNER_JOB	Customer Segment Business Duty
	CUSTOMER_ANALYST_JOB	
	MARKET_ANALYST_JOB	
	MERCHANDISER_JOB	
	CUSTOMER_SEGMENT_ADMINISTRATOR_JOB	Customer Segment Advanced Duty
AC	ASSORTMENT_PLANNER_JOB	Advanced Clustering Business Duty
	MERCHANDISER_JOB	
	CLUSTERING_ADMINISTRATOR_JOB	Advanced Clustering Advanced Duty
AE	ATTRIBUTE_EXTRACTION_JOB	Attribute Extraction Duty
Admin	CONFIG_ADMINISTRATOR_JOB	Configuration Administrator Duty
	INTEGRATION_ADMINISTRATOR_JOB	Integration Administrator Duty

Configuration

This section provides details about application configurations that can be modified as part of a deployment. The list of configurations is limited to those settings that most likely need to be reviewed and adjusted before the applications are used. Some configuration points cannot be adjusted after initial setup, so a careful review of the configurations should be done. For a complete list of configurable values, see the complete list of configuration values in RSE_CONFIG.

Any installation configuration must be set before the application is initialized with data and used.

Note: Before beginning any advanced customization, you must consult with development.

Generic Configuration

The following is a list of configurations that can be adjusted. This list does not include internal configuration values. All of these entries are adjusted by manipulating the corresponding rows in RSE_CONFIG where APPL_CODE='RSE'.

Table 10–2 Oracle Retail Advanced Science Generic Configurations

Parameter Name	Example Value	Description	Notes
CAL_PERIOD_LEVEL	4	This is the calendar hierarchy level that is used to drive the processing (installation configuration).	Usually this value should be set to the Fiscal Week level of the Fiscal Calendar. The number 4 refers to the fourth level in RSE_HIER_LEVEL table. This cannot be adjusted after initial setup.
CMGRP_HIER_TYPE	1	The hierarchy ID to use for the CM Group (installation configuration).	If an alternate hierarchy is to be used, then this configuration must specify that hierarchy type. Otherwise it must reflect the primary product hierarchy type. The value represented here relates to the ID in RSE_HIER_TYPE. This cannot be adjusted after initial setup.
CMGRP_LEVEL_ID	5	The hierarchy level ID that contains the level of the product hierarchy where the CM Group level exists (installation configuration).	Once the CMGRP_HIER_TYPE is configured, this level must be set to indicate the level of the hierarchy (RSE_HIER_LEVEL) that defines the categories.
PRIMARY_LANGUAGE_CODE	EN	The name of the language code to use for all RSE data sourced from RI (installation configuration).	Data values stored in the database are not multi-language capable, and are not affected by the UI language settings, like the UI labels are. This setting must select the language code for which data should be shown in the UI.
RA_FISCAL_CAL_ID	1240	ID of the calendar to use from RI since RI supports multiple calendars (installation configuration).	This value must match the ID of the desired fiscal calendar loaded into W_MCAL_PERIOD_DS.MCAL_CAL_ID.
CHAIN_LEVEL_DESC	CHAIN	The description to use for any top level hierarchy element when one must be manually created.	This description must be adjusted if a different description for the top level of the hierarchy is desired.
DEFAULT_LOCALE	en_US	The default locale to use for rendering elements that cannot support multiple locales.	Adjust this value so that it contains the correct LOCALE setting for most users of the application.

Table 10–2 (Cont.) Oracle Retail Advanced Science Generic Configurations

Parameter Name	Example Value	Description	Notes
DISPLAY_DATE_FORMAT	Mon dd, yyyy	The default date mask to be used by UI.	Adjust this so that the format of dates can be displayed as desired.
FAKE_CUST_DAY_TXN_THRESHOLD	10	The maximum daily transaction threshold to identify a fake customer.	This setting must be adjusted to the maximum number of daily transactions that a normal customer usually has. This allows all other customers with higher transactions counts to be excluded from processing.
UI_TZ	America/New_York	Timezone for display. Must match SELECT tzname FROM V\$TIMEZONE_NAMES.	This setting must be adjusted so that it contains a proper time zone setting indicating where most users run the UI.

Advanced Clustering Configurations

The following is a list of configurations that can be adjusted for the Advanced Clustering Cloud Service. This list does not include internal configuration values. All of these entries are adjusted by manipulating the corresponding rows in RSE_CONFIG where APPL_CODE='CIS'.

Table 10–3 AC Configurations

Parameter Name	Example Value	Description	Notes
PERF_CIS_APPROACH	CDT	The approach to use for performance-based clustering. Available options are CDT and DT.	If the CDT application is in use, then this configuration must be set to CDT, which will result in the use of attribute groups for product attribute clustering. Otherwise, it must be set to DT, which will result in the use of the top raw attribute values. This cannot be adjusted after initial setup.
ATTR_NAME_SEPARATOR	-	The separator character(s) to use to separate the different components of the attribute names in CIS_TCRITERIA_ATTR.	Adjust this value if a different separator is desired in the UI for attribute names built from multiple values. This cannot be adjusted after initial setup.
CIS_CONTR_SLS_SRC_COLUMN	SLS_AMT	Source column for contribution BI chart sales values. SLS_AMT, SLS_QTY, PROFIT_AMT are the allowable values.	Adjust this value as needed so that the BI shows contributions based on the desired sales column.

Table 10-3 (Cont.) AC Configurations

Parameter Name	Example Value	Description	Notes
CIS_CONTR_X_SRC_COLUMN	SLS_AMT	Source column for the x axis of the contribution BI chart. SLS_AMT, SLS_QTY, PROFIT_AMT are the allowable values.	Adjust this value as needed so that the BI shows contributions based on the desired sales column.
CIS_CONTR_Y_SRC_COLUMN	SLS_QTY	Source column for the y axis of the contribution BI chart. SLS_AMT, SLS_QTY, PROFIT_AMT are the allowable values.	Adjust this value as needed so that the BI shows contributions based on the desired sales column.
CIS_DFT_PIVOT_LVL	6	Default pivot level to show in explore data.	This configuration indicates the lowest level of the organization hierarchy (see <code>rse_hier_level</code>) that should be shown in the Explore Data pivot table. A value of 6 allows store locations to be visible, but this can be adjusted to a higher level if this level of detail is not desired.
CIS_IDX_AVG_SRC_COLUMN	SLS_AMT	Name of the column to use for index-to-average BI calculations. SLS_AMT, SLS_QTY, PROFIT_AMT are the allowable values.	Adjust this value as needed so that the BI shows index to averages based on the desired sales column.
CIS_NUMERIC_DFT	0	Default attribute value for numeric.	This can be adjusted so that attributes without a value are displayed with the desired value. This cannot be adjusted after initial setup.
CIS_STRING_DFT	UNKNOWN	Default attribute value for string.	This can be adjusted so that attributes without a value are displayed with the desired value. This cannot be adjusted after initial setup.
CIS_VARIABILITY_IDX_SRC_COLUMN	SLS_AMT	Name of the column to use for variability index BI calculations. SLS_AMT, SLS_QTY, PROFIT_AMT are the allowable values.	Adjust this value as needed so that the BI shows variability indexes based on the desired sales column.
DEFAULT_NUM_ATTR_VALUE	15	Constant for number of discrete values allowed for store attribute clustering.	Any attribute that has more than this number of distinct categorical values is not available for use as an attribute.
DEFAULT_STR_CATEGORICAL_ATTR	UNKNOWN	Default string description for row added in <code>cis_criteria_attr_type_value</code> table for unmatched grouping.	This can be adjusted so that attributes without a value are displayed with the desired value. This cannot be adjusted after initial setup.

Table 10–3 (Cont.) AC Configurations

Parameter Name	Example Value	Description	Notes
INSIGHT_BI_SALES_AMT_MARGIN	N	Flag to show sales BI based on sales amount and margin.	Either INSIGHT_BI_SALES_AMT_MARGIN or INSIGHT_BI_SALES_AMT_UNIT can be assigned a Y so that the BI displays the BI using the desired columns.
INSIGHT_BI_SALES_AMT_UNIT	Y	Flag to show sales BI based on sales amount and units	
MAX_HIST_WEEK_CNT	104	The maximum number of weeks that must be selectable by the UI when processing historic data.	
MAX_ITEMS_IN_GRAPH_CLUSTER_DETAIL	-1	Maximum number of clusters to be displayed in cluster details graph.	A value of -1 results in an unlimited number of values. Adjust if necessary.
MAX_ITEMS_IN_GRAPH_CLUSTER_LIST	-1	Maximum number of clusters to be displayed in cluster list graph.	A value of -1 results in an unlimited number of values. Adjust if necessary.
MNG_RUN_NO_WKS	26	Display run for past n weeks.	
PERF_ATTR_TOPN_COUNT	3	The number of attribute values to be used per product category for performance-based clustering.	If PERF_CIS_APPROACH configuration is DT, then this configuration will limit the number of attribute values to this number of values with the greatest sales.
PERF_NUM_WEEKS_FOR_SLS_SHARE	16	The number of weeks to be used while calculating the sales share for the product attributes.	
PERF_NUM_WEEKS_FOR_TOPN_CALC	16	The number of weeks to be used while identifying the top <i>n</i> attributes	
SELECT_ALL_MERCH_NODES	N	Flag to identify if all (or only first) merchandise node(s) to be selected by default.	
SUMM_CAL_ALL_LVL	N	Flag to identify whether performance summarization allowed at all available calendar levels.	
SUMM_MERCH_ALL_LVL	N	Flag to identify whether performance summarization allowed at all available merchandise levels.	

Assortment and Space Optimization Configurations

The following is a list of configurations that can be adjusted for the Assortment and Space Optimization Cloud Service. This list does not include internal configuration values. All of these entries are adjusted by manipulating the corresponding rows in RSE_CONFIG where APPL_CODE='SO'.

Table 10–4 ASO Configurations

Parameter Name	Example Value	Description	Notes
SO_LOC_HIER_TYPE	2	The hierarchy ID to use for location (installation configuration).	To use an alternate location hierarchy for ASO, adjust this configuration accordingly.
SO_PROD_HIER_TYPE	1	The hierarchy ID to use for the product (Installation configuration).	To use an alternate product hierarchy for ASO, adjust this configuration accordingly.
ALRT_LESS_THAN_PCT_USED_SPACE	0.8	An alert will be triggered if the run optimization results use less space than the value specified by this global parameter.	
ALRT_LESS_THAN_SERVICE_LEVEL_AMT	0.80	An alert will be triggered if the run optimization results have a sales service level lower than the value specified by this global parameter.	
ALRT_LESS_THAN_SERVICE_LEVEL_QTY	0.80	An alert will be triggered if the run optimization results have a quantity service level lower than the value specified by this global parameter.	
ALRT_MORE_THAN_CNT_PRODUCT_DROPPED	10	An alert will be triggered if the run optimization results dropped more products than the value specified by this global parameter.	
ALRT_MORE_THAN_PCT_PRODUCT_DROPPED	0.2	An alert will be triggered if the run optimization results dropped a percentage of product higher than the value specified by this global parameter.	
ALRT_NO_FEASIBLE_SOLUTION	0	An alert will be triggered if the run optimization results have no results.	
ALWAYS_REVIEW_MAPPING_RES_FLG	N	Default value is N. A Y flag indicates that a user mapping review is always required (regardless of results or errors). A N flag triggers a review based on other flags and conditions.	
CAPACITY_RANGE_UNITS	25	Capacity range units used by SO Solver. This parameter value maps to a CRU row with this value ID within so_prod_constr_range_values table.	
DEFAULT_BAY_MERGE_CONSTR_FLG	N	Default indicator for the use of merging bays constraint.	

Table 10–4 (Cont.) ASO Configurations

Parameter Name	Example Value	Description	Notes
DEFAULT_BLOCKING_CONSTR_FLG	Y	Default indicator for the use of blocking constraint.	
DEFAULT_SPACING_CONSTR_FLG	Y	Default indicator for use of spacing constraint.	
DEFAULT_USABLE_SPACE_CONSTR_FLG	N	Default flag indicating if space constraint must be used.	
DEFAULT_USABLE_SPACE_CONSTR_PCT	1	Default usable space constraint percentage.	
DEMAND_DISTRIBUTION	Normal	Demand distribution used by SO Solver.	
DFLT_HORIZONTAL_BLOCKING_FLG	N	A Y value in this flag indicates the analytics that combining adjacent attribute blocks should be done (when possible).	
DFLT_REPL_CASEPACK	1	Default replenishment parameter for casepack.	
DFLT_REPL_FACINGS_LIFT	0	Default facing lift.	
DFLT_REPL_SHELF_PARAM	0	Default shelf replenishment parameter.	
DFLT_REPL_SHELF_TT	2	Default replenishment type.	
DFLT_REPL_TYPE	2	Default replenishment type.	
DFLT_SHELF_THICKNESS	1	This is the default shelf thickness that is used by the POG-shelf interface to create the initial-bottom shelf for empty shelf fixtures.	
GV_DAYS_TO_VALIDATE_WO_CHANGES	21	Number of days without direct changes the validation process considers data objects for validation.	
GV_RESULT_DETAIL_LEVEL	SUMMARY	Level of detail for each validation that is used to produce the results (DETAIL:rows for every failure or SUMMARY: a row at the data object level).	
GV_VALIDATION_SECTIONS_TO_RUN	ASSORTMENT_POG_MAPPING_DS	Global validations are executed for the selected data objects. ASSORTMENT, POG, MAPPING and DS (Display Style).	
MAX_CAPACITY_RANGE	80	Maximum capacity range used by SO Solver.	
MAX_HEIGHT_RANGE	72	Maximum height range used by SO Solver.	
MAX_NUMBER_OF_FACINGS	5	Maximum number of facings used by SO Solver.	

Table 10–4 (Cont.) ASO Configurations

Parameter Name	Example Value	Description	Notes
MAX_NUM_OPT_LOC_BLOCK	10	Maximum number of blocks per optimization location.	
MAX_SHELF_THICKNESS	2.5	This is the maximum shelf thickness that can be used for shelf fixture edits.	
MAX_STACK_DFLT_LIMIT	5	The default maximum stack limit that is used for display styles that have a missing value.	
MIN_CAPACITY_RANGE	0	Minimum capacity range used by SO Solver.	
MIN_HEIGHT_RANGE	0	Minimum height range used by SO Solver.	
MIN_NUMBER_OF_FACINGS	1	Minimum number of facings used by SO Solver.	
MIN_SHELF_DEPTH	2	This is the minimum shelf depth that can be used for shelf fixture edits. The maximum shelf depth is defined by the fixture depth.	
MIN_SHELF_THICKNESS	0.5	This is the minimum shelf thickness that can be used for shelf fixture edits.	
MIN_SHELF_VERTICAL_GAP	2.5	This is the specific smallest allowable vertical offset (SAVO) value. Ensure that any edit action leaves at least this much space between shelves.	
MNG_ASSORT_NO_WKS	52	Display assortments for past <i>n</i> weeks.	
MNG_RUN_NO_WKS	52	Display run for past <i>n</i> weeks.	
NUMBER_OF_SIMULATED_DAYS	1000	Number of simulated days used by SO Solver.	
OPT_LOC_LVL1_NAME_STR	All Locations	This value is used entirely or as a prefix to generate the pogset location and optimization location top level names.	
OPT_LOC_LVL2_NAME_STR	PC_	This value is used as a prefix to generate the pogset location and optimization location mid level names.	
OPT_LOC_LVL3_NAME_STR	SC_	This value is used as a prefix to generate the pogset location and optimization location bottom level names.	
PC_SUM_CAPRANGE	Set Capacity Range	Capacity range label for product constraint summary.	

Table 10–4 (Cont.) ASO Configurations

Parameter Name	Example Value	Description	Notes
PC_SUM_ELEVATION	Elevation	Elevation label for product constraint summary.	
PC_SUM_ELEVRANGE	Set Elevation Range	Elevation range label for product constraint summary.	
PC_SUM_FACERANGE	Set Facing Range	Facing range label for product constraint summary.	
PC_SUM_FACINGS	Facings	Facings label for product constraint summary.	
PC_SUM_INCLUSION	Inclusion	Inclusion label for Product constraint summary	
PENALTY_PFG_MAX	10	Maximum product family group penalty.	
PENALTY_PFG_NORM	0.15	Normalized value that affects how close products of the same family are placed together.	
POGC_SUM_MERGEBAYS	Merge Adjacent Bays	Label to display in the UI for POG Constraint - Merge Adjacent Bays.	
POGC_SUM_PRODSPACE	Use Product Spacing	Label to display in the UI for POG Constraint - Use Product Spacing.	
POGC_SUM_SERVICELEVEL	Set Minimum Service Level	Label to display in the UI for POG Constraint - Minimum Service Level.	
POGC_SUM_USABLESPACE	Set Usable Space	Label to display in the UI for POG Constraint - Set Usable Space.	
POG_SET_LVL1_NAME_STR	All Planograms	This value is used to generate the name for the top level node on planogram list.	
PRODUCT_INCLUSION	2	Product Inclusion rule used by SO Solver. This parameter value maps to a IN row with this value ID within so_prod_constr_range_values table.	
PRODUCT_STACKING_HEIGHT_LIMIT	24	Product stacking height limit that applies as a global setting to all top products.	
PROD_ATTR_NAME_DELIMITER	-	This value is used as a delimiter between the product name/descr and the attribute name/descr when setting up POG attributes. A NULL value here results in no concatenations.	

Table 10–4 (Cont.) ASO Configurations

Parameter Name	Example Value	Description	Notes
REVIEW_DSFE_ERROR_FLG	Y	A Y flag indicates a user review is required for DSF errors. N lets the process move forward to the next stage using the DSF available.	
REVIEW_UNMAPPED_PROD_FLG	Y	A Y flag indicate a user review is required for unmapped products. N lets the process move forward to next stage eliminating unmapped products. This is not desired for products.	
REVIEW_UNMAPPED_STORE_FLG	Y	A Y flag indicates a user review is required for unmapped stores. N lets the process move forward to the next stage eliminating unmapped stores.	
SO_MIN_SERVICE_LEVEL	0.8	Minimum target service level for SO optimization process.	
SO_PROD_HIER_LEVEL_FOR_LEAF_NODE	7	Product hierarchy level number for leaf node.	
STD_ADJUSTMENT_COEFFICIENT_1	0.05	Analytical parameter. Demand standard deviation adjustment parameter 1.	
STD_ADJUSTMENT_COEFFICIENT_2	0.19	Analytical parameter. Demand standard deviation adjustment parameter 2.	
TOP_SHELF_STACKING_HEIGHT_LIMIT	18	Top shelf stacking height limit that applies as a global setting to all top shelves.	
UI_CONFIG_PC_RENDERED_COL_7	N	UI configuration for product constraints render column 7. Default Y means column will be rendered.	
UI_CONFIG_PC_RENDERED_COL_8	N	UI configuration for product constraints render column 8. Default Y means column will be rendered.	
UI_CONFIG_PC_RENDERED_COL_9	N	UI configuration for product constraints render column 9. Default Y means column will be rendered.	
UI_CONFIG_PC_VISIBLE_COL_1	N	UI configuration for product constraints visible column 1. Default Y means column will be visible.	
UI_CONFIG_PC_VISIBLE_COL_2	N	UI configuration for product constraints visible column 2. Default Y means column will be visible.	

Table 10–4 (Cont.) ASO Configurations

Parameter Name	Example Value	Description	Notes
UI_CONFIG_PC_VISIBLE_COL_3	N	UI configuration for product constraints visible column 3. Default Y means column will be visible.	
UI_CONFIG_PC_VISIBLE_COL_4	Y	UI configuration for product constraints visible column 4. Default Y means column will be visible.	
UI_MAX_POG_CONFIG_LENGTH	600	UI configuration for maximum length bound for the Create Lengths pop-up.	
UI_MAX_POG_CONFIG_NO_OF_BAYS	10	UI configuration for maximum number of bays bound for the Create Lengths pop-up.	
UI_MIN_POG_CONFIG_NO_OF_BAYS	1	UI configuration for minimum number of bays bound for the Create Lengths pop-up.	
UI_THRESHOLD_SL	Y	UI configuration for Thresholds Configurable for Service Level Formatting.	
UI_THRESHOLD_SL_MAX	0.85	UI configuration for Thresholds MAX after which color green would be shown.	
UI_THRESHOLD_SL_MIN	0.75	UI configuration for Thresholds MIN below which color red would be shown.	
USE_OPT_DT	N	SO global indicator for applying DT.	If the DT application is not in use, then this configuration should be set to an N to disable this feature.
USE_SERVICE_LEVEL_CONSTRAINT	Y	SO global indicator for applying service level constraints.	

Customer Decision Tree Configurations

The following is a list of configurations that can be adjusted the Customer Decision Tree Science Cloud Service. This list does not include internal configuration values. All of these entries are adjusted by manipulating the corresponding rows in RSE_CONFIG where APPL_CODE='CDT'. See [Chapter 3](#) for more details.

Table 10–5 CDT Configurations

Parameter Name	Example Value	Description	Notes
CDT_CAL_LEVEL_ID	4	The hierarchy level ID that contains the level of the calendar hierarchy that CDT operates on (should equate to Week - Installation configuration).	Normally, this configuration must be the same value as the common CAL_PERIOD_LEVEL configuration.
CDT_CMGRP_LEVEL_ID	5	The hierarchy level ID that contains the level of the product hierarchy that CDTs are created for (installation configuration).	Normally, this must be the same setting as the common CMGRP_LEVEL_ID configuration.
CDT_LOC_HIER_TYPE	2	The hierarchy ID to use for location (installation configuration).	This setting must be set to use either the trade area hierarchy type or the organization hierarchy type.
CDT_LOC_LEVEL_ID	4	The hierarchy level ID that contains the level of the location hierarchy that CDTs are created for (installation configuration).	If CDT_LOC_HIER_TYPE is set to trade area, this must be set to a value of 2. Otherwise, it must be set to the level of the organization hierarchy for which CDT output is desired.
CDT_PROD_HIER_TYPE	1	The hierarchy ID to use for the CM Group (installation configuration).	Normally, this setting must be the same value as the common CMGRP_HIER_TYPE configuration. It can either be the primary product hierarchy or an alternate product hierarchy.
CDT_CALC_RAW_ATTR_SIM	Y	Determines whether or not to execute Raw Attribute Value Similarities routines.	If the RI application is in use, then this setting must be set to a Y to enable calculation of data indented to be used by RI.
CDT_UI_DEF_CALC_PARENT_SEGMENT_FLG	Y	UI default for the calculate only parent consumer segments flag.	
CDT_UI_DEF_CALC_PARENT_TRADE_AREA_FLG	N	UI default for calculate only parent trade areas flag.	
CDT_UI_DEF_CDT_SCORE_HIST_CNT	20	UI default for the number of histogram buckets for the CDT scores histogram.	
CDT_UI_DEF_DATA_FILTER_HIST_CNT	20	UI default for the number of histogram buckets for the data filtering histograms.	
CDT_UI_DEF_EXCLUDE_CUST_CNT	1000	UI default for minimum require customer counts for pruning process.	
CDT_UI_DEF_EXCLUDE_MIN_SCORE	0.25	UI default for minimum CDT score required for the pruning process.	
CDT_UI_DEF_EXCLUDE_SKU_CNT	10	UI default for minimum number of SKUs for the pruning process.	

Table 10–5 (Cont.) CDT Configurations

Parameter Name	Example Value	Description	Notes
CDT_UI_DEF_EXCLUDE_TREE_LEVEL_CNT	2	UI default for minimum number of levels of the tree for the pruning process.	
CDT_UI_DEF_LOWEST_EXPANSION_LEVEL	15	UI default for lowest number of levels allowed for a tree.	
CDT_UI_DEF_MAX_CUST_AVG_DY_TXN	100	UI default for maximum number of times more than average a customers daily transaction count can be.	
CDT_UI_DEF_MAX_MISS_ATTR_CNT	3	UI default for maximum number of missing attributes a SKU can have	
CDT_UI_DEF_MIN_ATTR_SKU_CNT	5	UI default for minimum number of SKUs assigned to an attribute, to be used by the process.	
CDT_UI_DEF_MIN_ATTR_VALUE_SKU_CNT	5	UI default for minimum number of SKUs assigned to an attribute value, to be used by the process.	
CDT_UI_DEF_MIN_CUST_TXN_CNT	0.01	UI default for minimum number of transactions required for a customer, as a percent of the average number.	
CDT_UI_DEF_MIN_NODE_ITEM_CNT_PCT	0.05	UI default for the minimum percent of SKUs required for a node of the tree before it is considered a terminal node.	
CDT_UI_DEF_MIN_SKU_TXN_CNT	0.01	UI default for minimum number of transactions required for a SKU, as a percent of the average number.	
CDT_UI_DEF_PRUNING_HIST_CNT	20	UI default for the number of histogram buckets for the pruning histograms.	
CDT_XML_PRECISION	4	Default precision of weight field in CDT XML.	Adjust this to control the amount of precision in the generated CDT XML files.

Table 10–5 (Cont.) CDT Configurations

Parameter Name	Example Value	Description	Notes
HISTOGRAM_DEFAULT_BIN_APPROACH	C	The default histogram bin approach (C = Custom, W = Width)	
HISTOGRAM_DEFAULT_NUM_BINS	7	The default number of bins to display for CDT histograms	
MAX_NUM_WEEKS_FOR_SIMILARITY	104	The maximum number of weeks of sales transaction data to be used by the similarity process. This prevents the process from using too much data.	Adjust this to reduce the use of too much input data for the process. Enabling a high number of weeks will result in slower performance; however, it may be necessary if suitable data is not available in a smaller number of weeks.

Demand Transference Configurations

The following is a list of configurations that can be adjusted for the Demand Transference Science Cloud Service. This list does not include internal configuration values. All of these entries are adjusted by manipulating the corresponding rows in RSE_CONFIG where APPL_CODE='DT'.

Table 10–6 DT Configurations

Parameter Name	Example Value	Description	Notes
AE_CALC_INT_LENGTH	8	The number of weeks to group together for in an interval for the AE calculation.	This setting must be adjusted according to the data quality and quantity available for a customer. Higher values improve performance; however, it limit the amount of data available to be processed.
ATTRIBUTE_LIST_SEPARATOR	-	A separator to be used to display a list of attributes in Similarity Calculation screen.	
CDT_SIMILARITY_AVAILABLE	Y	Whether CDT similarity has been made available to DT.	If the CDT application is being used, this setting enables DT to use the similarities that CDT may have calculated.
DT_CAL_LEVEL_ID	4	The hierarchy level ID that contains the level of the calendar hierarchy that DT operates on (should equate to Week).	Normally, this configuration must be the same value as the common CAL_PERIOD_LEVEL configuration.
DT_CMGRP_LEVEL_ID	5	The hierarchy level ID that contains the level of the product hierarchy that DTs are created for.	Normally, this must be the same setting as the common CMGRP_LEVEL_ID configuration.

Table 10–6 (Cont.) DT Configurations

Parameter Name	Example Value	Description	Notes
DT_LOC_HIER_TYPE	2	The hierarchy ID to use for location.	This must be adjusted to trade area hierarchy or be left at the default organization hierarchy.
DT_LOC_LEVEL_ID	4	The hierarchy level ID that contains the level of the location hierarchy that DTs are created for.	If DT_LOC_HIER_TYPE is set to trade area, this must be set to a value of 2. Otherwise, it must be set to the level of the organization hierarchy for which DT output is desired.
DT_MDL_AP_EXP_WKS_BACK_END	1	The number of weeks back from the last date that range data has been loaded for (PR_LOC_STATUS_LAST_COMPLETED_WK) to end using for model apply export.	Adjust this and MDL_AP_EXP_WKS_BACK_START to control which weeks should be used during data export.
DT_PROD_HIER_TYPE	1	The hierarchy ID to use for the CM Group.	Normally, this setting must be the same value as the common CMGRP_HIER_TYPE configuration. It can either be the primary product hierarchy or an alternate product hierarchy.
DT_REMOVE_REDUNDANCY	N	If set to Y, then remove redundancy while calculating attribute-based similarities.	
DT_SIM_DISPLAY_ROWNUM	9999999	The number of distinct similarity values to show in the UI pop-up. Setting to a high number effectively eliminates this limit.	
GENERIC_SEPARATOR		A separator to be used to display a list of items, for example. SKU prod_ext_code name.	This value is used in the UI to separate lists of items. For example, when a list of attributes is shown in the UI, they will be delimited by this value.
HISTOGRAM_DEFAULT_BIN_APPROACH	W	The default histogram bin approach (C = Custom, W = Width).	
HISTOGRAM_DEFAULT_NUM_BINS	7	The default number of buckets in the contextual BIs.	
MAX_NUM_WEEKS_FOR_ATTR_WGT	104	The maximum number of weeks of input data to use for calculating attribute weights.	Setting this value to a higher value allows the use of more weeks of data, which will slow performance; however, enables better results for categories with infrequent sales.

Table 10-6 (Cont.) DT Configurations

Parameter Name	Example Value	Description	Notes
MAX_NUM_WEEKS_FOR_AVG_SLS	104	The maximum number of weeks of input data to use for calculating the average sales.	Setting this value to a higher value allows the use of more weeks of data, which will slow performance; however, enables better results for categories with infrequent sales.
MAX_NUM_WEEKS_FOR_FILTERING	104	The maximum number of weeks of input data to use for data filtering. Setting this value lower than the other MAX_NUM_WEEKS_FOR* configurations overrides those other configurations.	Setting this value to a higher value allows the use of more weeks of data, which will slow performance; however, enables better results for categories with infrequent sales.
MAX_NUM_WEEKS_FOR_MDL_CALC	104	The maximum number of weeks that should be used during model build calculation.	Setting this value to a higher value allows the use of more weeks of data, which will slow performance; however, enables better results for categories with infrequent sales.
MAX_NUM_WEEKS_FOR_MDL_UPDT	104	The maximum number of weeks that should be used during model build update calculation.	Setting this value to a higher value allows the use of more weeks of data, which will slow performance; however, enables better results for categories with infrequent sales.
MAX_NUM_WEEKS_FOR_SIMILARITY	104	The maximum number of weeks of input data to use for calculating similarity.	Setting this value to a higher value allows the use of more weeks of data, which will slow performance; however, enables better results for categories with infrequent sales.
MDL_AP_EXP_WKS_BACK_START	4	The number of weeks back from the last date that range data has been loaded for (PR_LOC_STATUS_LAST_COMPLETED_WK) to start using for model apply export.	Adjust this and DT_MDL_AP_EXP_WKS_BACK_END to control which weeks should be used during data export.
PRUNED_CATEGORIES_SEPARATOR	,	A separator to be used to display a list of pruned attributes in the Calculation screen.	
UI_DEF_CALC_PARENT_CS_ONLY_FLG	N	The UI default for calculate only parent customer segments flag.	
UI_DEF_CALC_PARENT_TA_ONLY_FLG	N	The UI default for calculate only parent trade areas flag.	
UI_DEF_MAX_MISS_ATTR_CNT	3	The maximum number of missing attributes for a SKU, before requiring it to be filtered from use.	

Table 10–6 (Cont.) DT Configurations

Parameter Name	Example Value	Description	Notes
UI_DEF_MIN_SKU_CNT	10	The UI default for minimum number of SKUs required for a segment/store.	
UI_DEF_MIN_SKU_TXN_LEN_PCT	0.01	The UI default for minimum SKU transaction length as a percentage of the CM Group average.	
UI_DEF_MIN_TOT_SLS_UNIT_PCT	0.01	The UI default for minimum total sales units as a percentage of the CM group average.	
WGT_CALC_INTERVAL_LENGTH	4	The number of weeks to group into an interval that is then used to perform weight calculations with.	This setting can be adjusted according to the data quality and quantity available for a customer. Higher values improve performance; however, it limits the amount of data available to be processed.

Returns Logistics Configurations

The following is a list of configurations that can be adjusted for the Returns Logistics process. All of these entries are adjusted by manipulating the corresponding rows in RSE_CONFIG where APPL_CODE='RL'.

Table 10–7 Returns Logistics Configurations

Parameter Name	Example Value	Description	Notes
RL_NUM_TOP_CATEGORIES	10	The number of top selling categories for Returns Logistics to process.	
RL_NO_OF_SIMULATION_RUNS	100	The number of simulated forecast runs to perform.	
RL_IN_SEASON_DEMAND_WEIGHT	0.99	The percentage of the in-season demand and returns rate to use. This is used to calculate the adjusted demand and returns rate.	
RL_HISTORICAL_DEMAND_WEIGHT	.01	The percentage of the historical demand and returns rate to use. This is used to calculate the adjusted demand and returns rate.	
RL_GUR_TIMELIMIT	60	GUROBI parameter. Time limit.	
RL_GUR_THREADS	0	GUROBI parameter. Number of allowed threads.	
RL_PROCESSING_THREADS	0	The number of threads to create for data processing. A value of 0 creates one thread for each machine processor.	

Market Basket Insight Configurations

The following is a list of configurations that can be adjusted for the Market Basket Insight Cloud Service. This list does not include internal configuration values. All of these entries are adjusted by manipulating the corresponding rows in RSE_CONFIG where APPL_CODE='MBA'. Some of the configurations note that overrides are possible. When this is required, see the RSE_CONFIG_CODE configuration table and use the appropriate Hierarchy Level value as the PARAM_CODE to override the value for only that hierarchy level.

Table 10–8 MBI Configurations

Parameter Name	Example Value	Description	Notes
ARM_CS_HIER_LEVEL	SBC	The highest level of the hierarchy to run the process for. Valid values are: CLS, SBC.	
ARM_CS_MAX_LIFT	100	The maximum lift required for an association rule.	
ARM_CS_MAX_RULE_COUNT	9999	The maximum number of rules desired for association rules.	Affects the count of rules per set size.
ARM_CS_MAX_SET_SIZE	2	The maximum number of hierarchy members to include in an association rule.	Maximum allowed is 4, although a setting this high can negatively affect performance with some datasets.
ARM_CS_MIN_CONFIDENCE	.05	The minimum confidence required for an association rule.	
ARM_CS_MIN_LIFT	.05	The minimum lift required for an association rule.	
ARM_CS_MIN_REV_CONFIDENCE	.05	The minimum reverse confidence required for an association rule.	
ARM_CS_MIN_SUPPORT	.001	The minimum percentage of transactions require for an association rule.	
ARM_CS_MIN_SUPPORT_TXN_CNT	1000	The minimum number of sales transactions required for creating association rules.	
ARM_CS_WEEK_CNT	1	The number of weeks that should be processed while calculating the association rules.	Only set this to more than 1 if there is a need to reprocess weeks that would have been processed in the prior batch window.
ARM_PH_MAX_LIFT	100	The maximum lift required for an association rule. Override with PARAM_CODE of the hierarchy level name.	
ARM_PH_MAX_RULE_COUNT	9999	The maximum number of rules desired for association rules.	Affects the count of rules per set size.
ARM_PH_MAX_SET_SIZE	2	The maximum number of hierarchy members to include in an association rule. Override with PARAM_CODE of the hierarchy level name.	Maximum allowed is 4, although a setting this high can negatively affect performance with some datasets.

Table 10–8 (Cont.) MBI Configurations

Parameter Name	Example Value	Description	Notes
ARM_PH_MIN_CONFIDENCE	.05	The minimum confidence required for an association rule. Override with PARAM_CODE of the hierarchy level name.	
ARM_PH_MIN_LIFT	.05	The minimum lift required for an association rule. Override with PARAM_CODE of the hierarchy level name.	
ARM_PH_MIN_REV_CONFIDENCE	.05	The minimum reverse confidence required for an association rule. Override with PARAM_CODE of the hierarchy level name.	
ARM_PH_MIN_SUPPORT	.001	The minimum percentage of transactions require for an association rule. Override with PARAM_CODE of the hierarchy level name.	
ARM_PH_MIN_SUPPORT_TXN_CNT	1000	The minimum number of sales transactions required for creating association rules.	
ARM_PH_TOP_LEVEL	CLS	The highest level of the hierarchy to run the process for. Valid values are: DEPT, CLS, SBC.	
ARM_PH_WEEK_CNT	1	The number of weeks that should be processed while calculating the association rules.	Only set this to more than 1 if there is a need to reprocess weeks that would have been processed in the prior batch window.
ARM_PROMO_HIER_LEVEL	SBC	The highest level of the hierarchy to run the process for. Valid values are: CLS, SBC.	
ARM_PROMO_MAX_LIFT	100	The maximum lift required for an association rule. Override with PARAM_CODE of the hierarchy level name.	
ARM_PROMO_MAX_RULE_COUNT	9999	The maximum number of rules desired for association rules.	Affects the count of rules per set size.
ARM_PROMO_MAX_SET_SIZE	2	The maximum number of hierarchy members to include in an association rule. Override with PARAM_CODE of the hierarchy level name.	Maximum allowed is 4, although a setting this high can negatively affect performance with some datasets.
ARM_PROMO_MIN_CONFIDENCE	.05	The minimum confidence required for an association rule. Override with PARAM_CODE of the hierarchy level name.	
ARM_PROMO_MIN_LIFT	.05	The minimum lift required for an association rule. Override with PARAM_CODE of the hierarchy level name.	

Table 10–8 (Cont.) MBI Configurations

Parameter Name	Example Value	Description	Notes
ARM_PROMO_MIN_REV_CONFIDENCE	.05	The minimum reverse confidence required for an association rule. Override with PARAM_CODE of the hierarchy level name.	
ARM_PROMO_MIN_SUPPORT	.001	The minimum percentage of transactions require for an association rule. Override with PARAM_CODE of the hierarchy level name.	
ARM_PROMO_MIN_SUPPORT_TXN_CNT	1000	The minimum number of sales transactions required for creating association rules.	
ARM_PROMO_WEEK_CNT	1	The number of weeks that should be processed while calculating the association rules.	Only set this to more than 1 if there is a need to reprocess weeks that would have been processed in the prior batch window.

Internationalization

The user interface supports multiple languages in a single instance, but the underlying database only supports a single language in an instance.

The database default language is selected at installation. Once set, there is no support for switching the database language.

The application user interfaces adhere to the language setting for each user's browser. For example, to change the language for the Firefox browser:

1. Select Tools from the menu bar.
2. Select Options.
3. Select Choose.
4. Select the language to add.

The following language are supported: English, German, Greek, Spanish, French, Croatian, Hungarian, Italian, Japanese, Korean, Dutch, Polish, Brazilian Portuguese, Russian, Swedish, Turkish, Simplified Chinese, and Traditional Chinese.

Configuration Updates

Use the Manage Configuration user interface screen in order to edit the values in the configuration tables. In order to access this screen, you must be assigned the role of CONFIG_ADMINISTRATOR_JOB. The tables you can edit within Manage Configuration depends upon which application roles you have been assigned, as described in [Table 10–1](#).

Once you have accessed Manage Configuration, you select, from the configuration tables available to you, the table that you want to review and edit. Note that each configuration table allows specific actions. You may or may not be able to view the table, edit some or all of the columns, or add and delete rows. (Example needed again)

Certain of the tables in Manage Configuration are not intended for configuration but instead are provided to display data that may be useful during implementation. For

each Oracle Retail Advanced Science interface file, a table is available that contains any records that have failed the data validation rules. Such tables have an extension of BAD. For example, for the `rse_prod_attr_grp_value_stg.txt` interface, the data is loaded into a table named `RSE_PROD_ATTR_GRP_VALUE_STG`. The error table is named `RSE_PROD_ATTR_GRP_VALUE_BAD`.

You can see all the error tables for all the interfaces via Manage Configuration. When an error occurs with an interface, you can review the BAD table to see what records failed validation and why. Each BAD table includes the following columns: `ERROR_ROWID`, `ERROR_ID`, `ERROR_DESCR`, and `ERROR_DT`. The `ERROR_ID` is useful as a reference. The `ERROR_DESCR` provides a message to explain the validation error.

Configuration Tables

The following tables can be viewed in the Manage Configuration screen.

Table 10–9 Customer Decision Tree Configuration Tables

Table Name	Description
<code>CDT_EXCLUDE</code>	Defines the various types of pruning filters that can be used to prevent a CDT from being used during the escalation phase of the CDT workflow.
<code>CDT_FILTER</code>	Defines the various types of data filters that can be used to filter sales transaction data used for the CDT calculation.
<code>CDT_VERSION</code>	Defines a version to collectively group together a batch of CDTs that have been created for a particular purpose.

Table 10–10 Advanced Clustering and Customer Segmentation Configuration Tables

Table Name	Description
<code>CIS_ALGORITHM_ATTR</code>	Defines the possible attributes for any algorithm.
<code>CIS_BUSINESS_OBJECT</code>	Hosts the list of applications that are registered and configured to use the clustering feature.
<code>CIS_BUSSOBJ_OBJ_ALG_XREF</code>	This cross reference is provided so that you can use the same algorithm to generate different cluster objectives. The same algorithm can be used to generate customer clusters as well as store clusters. At the same time it is possible to list multiple algorithms that can be used to achieve a similar objective.
<code>CIS_BUSSOBJ_TCRIT_HIER_XREF</code>	Specifies the possible hierarchy levels for each hierarchy type (merchandise and location) that are permitted for the combination of objective ID, business objective ID, and type criteria ID.
<code>CIS_BUS_OBJ_HIER_DEPLOY_XREF</code>	Identifies the business object, objectives, product hierarch types, and levels that can be selected for deployment. This ensures that only authorized types of clusters are exported to external systems (CMPO).
<code>CIS_BUS_OBJ_NESTED_TCRITERIA</code>	Determines the possible child cluster types for a parent cluster.
<code>CIS_BUS_OBJ_TCRITERIA_XREF</code>	Specifies the possible cluster types that are permitted for the combination of objective ID and business objective ID.
<code>CIS_BUS_OBJ_TCRIT_ALGO_ATTR</code>	Defines possible attributes for any algorithm, business objective, objective, and criteria.
<code>CIS_CLUSTER_OUTLIER_RULE</code>	Specifies the possible outlier rules for a type criteria.

Table 10–10 (Cont.) Advanced Clustering and Customer Segmentation Configuration

Table Name	Description
CIS_EFFECTIVE_PERIOD	Specifies the planning period information.
CIS_ODM_SETTINGS	Temporary storage for the data mining settings provided to ODM. You can store multiple runs of settings by using different mining functions and run IDs. To provide this data to ODM, you must write a view that provides the appropriate rows of data.
CIS_SALES_TYPE	Specifies the type of sales information - historical, forecasted, or planned.
CIS_SRC_ENTITY	Defines the different database views available for use in the various clustering implementations. These settings help control how the attributes are used throughout the system.
CIS_TCRITERIA_SRC_XREF	Cross reference to SRC_ENTITY_NAME and for the settings for Partitioning, Informations for the SRC_ENTITY_NAME attributes.
CIS_TCRIT_SRC_TYPE_XREF	Describes the xref of type_criteria and sales.
CIS_TYPE_CRITERIA	List of different clustering types and criteria that can be used to generate clusters.

Table 10–11 Demand Transference Configuration Tables

Table Name	Description
DT_EXCLUDE	Defines the different types of pruning filters available to prevent a DT result from being used during the escalation phase of the DT workflow.
DT_FILTER	Defines the different types of data filters used during the DT data filtering process.

Table 10–12 ORASE Generic Configuration Tables

Table Name	Description
RSE_AGGR_SRVC_CONFIG_LEVELS	Defines the different hierarchy types and levels that must have aggregated data created as part of the hierarchy configuration.
RSE_BUSINESS_OBJECT_ATTR_MD	Defines the attributes for business objects and also provides relevant details about where from external table sources to obtain the data for this attribute.
RSE_BUSINESS_OBJECT_DB_SRC	Defines the source database objects for the attributes.
RSE_CONFIG	Provides configuration names and their values for the settings that can be changed and affect the operation of the software.
RSE_CONFIG_CODE	Provides configuration values for those configurations that can have different values, depending on how other values are configured. For example, if a configuration is required for a default error tolerance, but department 1 requires a different value, then a row here with a PARAM_CODE of 1 will enable a different value than the base configuration in RSE_CONFIG for just that department.
RSE_EMAIL_CFG_DISTR	Intersection table between the email configuration table and the email distribution list table used to resolve the many-to-many relationships.
RSE_EMAIL_DISTR_LIST	Provides the email distribution lists.

Table 10–12 (Cont.) ORASE Generic Configuration Tables

Table Name	Description
RSE_EMAIL_NOTIF_CFG	Defines the email messaging configuration.
RSE_EXP_GRP	Hosts the list of applications that are registered and configured to use the clustering application.
RSE_HIER_LEVEL	Defines the various levels for all the hierarchies.
RSE_HIER_TYPE	Defines the hierarchies that are available for use by the RSE applications.
RSE_LOAD_SRVC_CFG	Defines a data loader that can be executed through the data loading framework.
RSE_LOAD_VALDT_RULES_CFG	Defines the validation rules that a data loader performs, along with some configurable options that impact rows that fail the validation.
RSE_PROC_TASK_TMPL	Defines the templates for processing tasks used by the applications.
RSE_PROC_TMPL	Defines the processing templates for asynchronous or synchronous invocable from Java applications
RSE_SRVC_CONFIG	Defines all the database service routines that can be invoked through the database service framework in RSE.

Table 10–13 Assortment and Space Optimization Configuration Tables

Table Name	Description
SO_BI_ELEMENT	Contains configuration metadata for BI elements.
SO_BI_ELEMENT_CHART	Contains the metadata to configure BI element charts.
SO_INT_TRANSFORMATION_KEY	Used to help perform interface data transformation. The values in this table are used to align data from external sources with the data expected by ASO. It helps to isolate translation conversion issues.
SO_POG_FIXTCONF_ALG_PARAM	Stores the list of different algorithm parameters that a user can customize while running the fixture smart start process.
SO_POG_FIXT_CONFIG_ALGORITHM	Stores the list of available algorithms to perform the shelf fixture smart start process to create shelves for empty fixtures.
SO_PROD_CONSTR_RANGE_VALUES	Contains the list of product constraint values supported by the application.
SO_REPL_PARAM_DESCRIPTION	Stores the list of replenishment parameters that a user can change. These parameters have a defined list of valid values that are stored in this table so that they can be used by the UI for display.
SO_RUN_OBJECTIVE_FUNC	Contains the list of possible run objective functions that are supported by the application.
SO_STACK_CAP_STYLE	

Email Notification Configuration

The RSE_EMAIL_CFG_DISTR, RSE_EMAIL_DISTR_LIST, and RSE_EMAIL_NOTIF_CFG tables listed in the tables in [Configuration Tables](#) can be used to enable email notifications for certain batch processes. The RSE_EMAIL_NOTIF_CFG table contains a list of processes that can generate an email notification. Configurations are available

for loading interface files (PROCESS_TYPE=LOAD), as well as configurations for running the automated batch processes (PROCESS_TYPE=BATCH).

The RSE_EMAIL_DISTR_LIST table is used to define email distribution lists that should receive email notifications. If all email notifications are to be sent to a the same email address, then a row with a DEFAULT_FLG=Y can be used so that the distribution list is used for all email notices.

If some processes require different email recipients, then additional distribution lists can be created for those in RSE_EMAIL_DISTR_LIST. Once created there, a record can be created in RSE_EMAIL_CFG_DISTR that links the distribution list with the email notification. Any email configuration without a specific entry in RSE_EMAIL_CFG_DISTR is sent to the default distribution list.

By default, all LOAD failures trigger an email notification with details about the load that failed. Additionally, all ORASE batch processes send an email notification when the entire batch process completes. This email indicates whether the batch process was successful or if a failure occurred.

Attribute Processing

This chapter addresses attribute preprocessing. It contains the following sections:

- [Attribute Preprocessing](#)
- [Product Attribute Loading](#)

Attribute Preprocessing

Attributes provide context about products and enhance the accuracy of DT and CDT models. Attributes are stored within RI and are derived from product descriptions and merchandise hierarchy.

RADM may or may not contain product attributes. Any attributes found in RADM may have been created for BI reporting or other purposes and may need mining or preprocessing to make them suitable for Oracle Retail Advanced Science.

Some steps in attribute preprocessing require manipulating attribute data. Oracle Enterprise Product Data Quality is a licensed software package that facilitates some of the preprocessing data manipulation steps required to make attributes suitable for CDT and DT modeling.

Here is an example of product information for yogurt.

- Product description: Brand A non-fat organic 6 oz.
- Class description: Dairy product.
- Sub-class description: Yogurt.

SKU/Store attributes determined by preprocessing:

- Brand
- Price
- Size

Note that CDT modeling works optimally when there are five or fewer possible values for any given SKU-store attribute. For example, many price points are available for yogurt. For CDT, it is better to define between three and five price bins (that is, budget, regular, premium, and elite).

For ASO, the application itself does not have any specific requirements; the business requirements for the attribute values are what matters. ASO supports the use of attribute value groups to control the layout of products. If the business requirement states that products should be organized by many different attribute values, then, for ASO, the attribute value groups must have as many values as needed for the

organization specified. Note that care must be taken, as these two requirements can contradict one another.

Process Overview

The basic steps for attribute preprocessing are as follows:

- Populating RADM with attribute data
- Translating (optional)
- Parsing
- Cleansing and standardizing
- Categorizing and labeling
- Defining attributes
- Binning and grouping

Populating RADM with Attribute Data

A few steps are required to make RADM attributes suitable for Oracle Retail Advanced Science so that the applications can use this data.

The first requirement is to ensure that the attribute values are populated in RADM. This is the source for the attribute data and must be loaded there in order to be available to Oracle Retail Advanced Science.

Regarding RADM attributes: In RADM, an attribute can be defined in multiple ways. Flex attributes are typically stored in a column of the `W_PRODUCT_ATTR_D` table. RADM has a table `W_RTL_METADATA_G` that contains a list of defined attribute locations. Consult this list to see if there is already a defined place to store a particular attribute value.

RI also offers the ability to store Item Differentiators for products. These are essentially User Defined Attributes (UDAs), which consist of a lookup code for the attribute and the attribute value. These lookup codes are then defined in RADM's standard translation table (`W_DOMAIN_MEMBER_LKP_TL` with domain codes of `ITEM_UDA_HEAD` and `ITEM_UDA`). The actual association of an item to one of the UDAs is performed in the `W_RTL_ITEM_GRP_1_D` table.

Once attributes are available in RADM, it is necessary to define these attributes in the `RSE_BUSINESS_OBJECT_ATTR_MD` table. This requires engagement with OCI in order to configure the data correctly for retrieval from Oracle Retail Advanced Science. This table must be set up with appropriate metadata to define the source of the attributes from RADM. The sample `seed_data` file for this table contains some standard attributes that can be defined in RADM, but the table must be adjusted to contain the complete list of attributes that should be available for the modules to use. This must include Flex Attributes as well as User Defined Attributes.

Once attributes are defined in the `RSE_BUSINESS_OBJECT_ATTR_MD`, the next step is to provide custom lists of attributes that should be used per product category. This can be done through the `RSE_PROD_ATTR_GRP_VALUE_STG` and `RSE_PROD_ATTR_VALUE_XREF_STG` interfaces. The first interface is used to define the output of the binning and grouping of attributes. For example, if Coffee needs a Brand Tier attribute, and it should have values of Premium, Value, and Mainstream, then this interface would define this Coffee Brand Tier attribute, along with the values of Premium, Value, and Mainstream, and it should specify what source attribute is to be used for this (the source is in `RSE_BUSINESS_OBJECT_ATTR_MD`). The second table of the interface (`RSE_PROD_ATTR_VALUE_XREF_STG`), enables the association of

specific Brand attributes to the binned/grouped attribute values from the first interface (RSE_PROD_ATTR_GRP_VALUE_STG).

One concept to consider for these attributes and attribute values, is that they must be unique across all product categories. This offers the ability to classify one Brand as Premium for one product category, while it could be Mainstream for another product category. Additionally, it enables a different selection of attribute values for each product category. For example, another product category might not have a Premium Brand Tier, and therefore the interface would not include this value in this interface for that product category.

Translating

This step is only needed when the product data is in a different language than the customer's primary language.

Parsing

This step identifies and extracts target key words, such as "Brand A," "small," "blue," and "non-fat." from the source data (such as product description). It is done through semantic recognition, usually by software such as Oracle Enterprise Product Data Quality.

Cleansing and Standardization

This step edits the text and corrects spelling and grammar. For example, "Addr." will be recognized and converted into "Address" and "St." into "Street." EPDQ can facilitate this step.

Categorizing and Labeling

This step classifies targeted key words into the pre-defined categories, such as "Brand A" for "Brand," "small" for "Size" and "blue" for "Color." The product record can thus be labeled by the category values. EPDQ can facilitate this step.

Defining Attributes

With the extracted categories from the product data, attributes are defined. They can be some or all of the categories identified, based on contextual business knowledge and how populated the categories are.

Binning and Grouping

Binning and grouping are used to consolidate and reduce the number of possible values for an attribute into a manageable number.

- Binning divides numerical attributes, such as price, discounts, and mileage, into discrete sets of ranges, such as <=\$10, \$10~\$25, and >\$25.
- Grouping combines textual attributes that are too granular into a smaller set of attribute values. For example, tea weight can have dozens of values; grouping merges the values into coarser ranges (like small or large) and reduces the number of possible attribute values.

Product Attribute Loading

This section provides an example of adding an attribute for use by Oracle Retail Advanced Science into all the relevant tables. In this example, a new attribute is added to represent Flavor within the Coffee product category.

The process flow for this involves:

1. Identify the need to add a new product attribute for a product category.
2. Determine where the attribute data is found within RADM.
3. Coordinate with OCI to add the attribute definition in the tables, if it not already present.
4. Coordinate with OCI to run the batch process to load attribute data from RADM.
5. Determine if the attribute data requires any special grouping or binning.
6. Populate the RSE_PROD_ATTR_GRP_VALUE_STG staging table with attribute definition and values.
7. Populate the RSE_PROD_ATTR_VALUE_XREF_STG staging table with data to associate raw RADM attribute values to the Attribute Groups defined above.
8. Coordinate with OCI to run the batch process that processes the interface staging tables.
9. Coordinate with OCI to update the CIS attribute data to reflect the new attribute (product attributes).
10. Coordinate with OCI to update the CIS attribute data to reflect the new attributes (non-product attributes).

Introduce New Attribute

The first step in the process is the catalyst that triggers the remaining steps. The catalyst is the new attribute that has been introduced and must be made available within Oracle Retail Advanced Science.

Determine Attribute Source and Define in Oracle Retail Advanced Science Tables

The new attribute is loaded from RADM for each of the products that require this attribute. RADM has multiple ways of loading attributes, so the approach used varies, depending on where and how the data is stored in RADM. The process involves defining the source table and then defining the column (or column filter values) used to identify the attribute. Once the source is determined, the appropriate values are loaded into RSE_BUSINESS_OBJECT_ATTR_MD and possibly RSE_BUSINESS_OBJECT_DB_SRC.

W_PRODUCT_D or W_PRODUCT_ATTR_D

RADM's W_PRODUCT_D table and W_PRODUCT_ATTR_D table can provide attributes from any of the available columns in these tables. The W_PRODUCT_D table contains named columns with data of a specific logical value, while the W_PRODUCT_ATTR_D table contains a more flexible set of Number, Text, and Date columns that can contain varying values, depending on the implementation. From an attribute point of view, these tables are effectively the same and require the same type of handling.

W_RTL_ITEM_GRP1_D or W_RTL_ITEM_GRP2_D

The W_RTL_ITEM_GRP1_D and W_RTL_ITEM_GRP2_D tables in RADM are different than the other product attribute sources, in that these tables can have attributes implemented as unique rows and specific columns. These tables contain a PROD_GRP_TYPE column, which defines the type of data in the table. Values of ITEMUDA are used for User Defined Attributes. Rows in which the PROD_GRP_TYPE corresponds to the BRAND, COLOR, FLAVOR, SCENT, FABRIC, and STYLE WID columns (ex. BRAND_WID) are also possible.

Populate RSE_PROD_ATTR_GRP_VALUE_STG Interface

Once the attribute data has been reviewed and groups have been defined, it is necessary to define the attribute groups and process them into the database. The output of the prior step must be loaded into the staging table for Attribute Value Groups (RSE_PROD_ATTR_GRP_VALUE_STG). This interface defines two sets of data and is used to load two different tables.

Table 11-1 RSE_PROD_ATTR_GRP_VALUE_STG

Column	Example	Description
PROD_HIER_TYPE_NAME	Product Hierarchy	Must match the NAME from RSE_HIER_TYPE that has the ID equal to the RSE_CONFIG for CMGRP_HIER_TYPE.
PROD_EXT_KEY	CLS~1000~10000	The external key used to identify the product category (for example, Coffee Class). This value is the same as in RADM's INTEGRATION_ID of the W_PROD_CAT_DH, and also the PROD_EXT_KEY of the RSE_PROD_SRC_XREF table.
ATTR_SHORT_DB_NAME	FLAVOR	This must match the SHORT_DB_NAME of the RSE_BUSINESS_OBJECT_ATTR_MD table for the newly added attribute.
PROD_ATTR_GRP_EXT_KEY	CLS~1000~10000~flavor_yn CLS~1000~10000~flavor_type	This must be a unique value to describe the attribute to be used by the modules. Since the source Flavor attribute is being defined as two different attributes, two example values are shown here.
PROD_ATTR_GRP_NAME	FlavorYN FlavorType	A name to be displayed in the UI for the new attribute. Two example values are shown here.
PROD_ATTR_GRP_DESCR	Flavor Y/N Identifier Flavor Type	An optional/additional descriptive value that can be displayed in the UI for the new attribute.
PROD_ATTR_VALUE_KEY	(See additional table below)	A unique/external identifier for the new attribute values.
PROD_ATTR_VALUE_NAME	(See additional table below)	A name displayed in the UI for the attribute value.
PROD_ATTR_VALUE_DESCR	(See additional table below)	An optional/additional descriptive value that could be shown in the UI for the new attribute value.
FUNC_ATTR_FLG	N	This is a Y/N flag to indicate whether this attribute is considered to be an attribute associated with a specific function or role (Y) or not (N). For example, a customer cannot choose a product with a different value for the auto wiper blade size because each car model has a specific size requirements.

Here is a table showing the different values for adding the example Flavor Attribute Values.

Table 11–2 Flavor Attribute Values

PROD_ATTR_GRP_NAME	PROD_ATTR_VALUE_KEY	PROD_ATTR_VALUE_NAME	PROD_ATTR_VALUE_DESCR
FlavorYN	CLS~1000~10000~flavor_yn~y	Y	Yes
FlavorYN	CLS~1000~10000~flavor_yn~n	N	No
FlavorType	CLS~1000~10000~flavor_type~non	Non Flavored	Non Flavored
FlavorType	CLS~1000~10000~flavor_type~fruit	Fruit Flavored	Fruit Flavored
FlavorType	CLS~1000~10000~flavor_type~mild	Mild Flavored	Mild Flavored
FlavorType	CLS~1000~10000~flavor_type~special	Specialty	Specialty

Populate RSE_PROD_ATTR_VALUE_XREF_STG Interface

Once the RSE_PROD_ATTR_GRP_VALUE_STG interface has been loaded, it is possible to load the RSE_PROD_ATTR_VALUE_XREF_STG interface with a mapping of actual product attribute values (otherwise known as base attributes) to the attribute groups that were loaded via RSE_PROD_ATTR_GRP_VALUE_STG. The format of data to be loaded here depends on the format of the base attributes. Only one set of attribute value columns should be populated for this interface. These sets are MIN_ATTR_NUM_VALUE and MAX_ATTR_NUM_VALUE (for numeric attributes), ATTR_STRING_VALUE (for text attributes), MIN_ATTR_DATE_VALUE and MAX_ATTR_DATE_VALUE (for date attributes), ATTR_VALUE_EXT_CODE (for dimension based attributes). The sets are mutually exclusive of each other for this interface.

Table 11–3 RSE_PROD_ATTR_VALUE_XREF_STG

Column	Example	Description
PROD_ATTR_VALUE_KEY	CLS~1000~10000~flavor_yn~y	Must match a PROD_ATTR_VALUE_KEY that was loaded via the RSE_PROD_ATTR_GRP_VALUE_STG interface.
MIN_ATTR_NUM_VALUE	0	Minimum numeric value to associate with this attribute group value. Only applicable if this attribute uses the ATTR_NUM_VALUE column to store the base attribute value.
MAX_ATTR_NUM_VALUE	7	The maximum numeric value to associate with this range. Only applicable in conjunction with MIN_ATTR_NUM_VALUE.
ATTR_STRING_VALUE	Y	A string value to associate with this attribute group value. Only applicable if this attribute uses the ATTR_STRING_VALUE column to store the base attribute value.

Table 11-3 (Cont.) RSE_PROD_ATTR_VALUE_XREF_STG

Column	Example	Description
MIN_ATTR_DATE_VALUE	2010-01-01	The minimum date value to associate with this attribute group value. Default date format for provided control file is YYYY-MM-DD. Only applicable if this attribute uses the ATTR_DATE_VALUE column to store the base attribute value.
MAX_ATTR_DATE_VALUE	2010-01-31	The maximum date value to associate with this attribute group value. Default date format for provided control file is YYYY-MM-DD. Only applicable in conjunction with MIN_ATTR_DATE_VALUE.
ATTR_VALUE_EXT_CODE	32	For base attributes that are sourced from W_RTL_ITEM_GRP1_D, this column can be used to specify the key from the appropriate source column. This is applicable if this attribute uses ATTR_VALUE_EXT_CODE to store the attribute value.

Here is a table containing some examples for adding a new flavor attribute, using string-based attributes.

Table 11-4 Adding a New Flavor Attribute

PROD_ATTR_VALUE_KEY	ATTR_STRING_VALUE
CLS~1000~10000~flavor_yn~y	BLUEBERRY
CLS~1000~10000~flavor_yn~y	RASPBERRY
CLS~1000~10000~flavor_yn~y	VANILLA
S~1000~10000~flavor_yn~y	CARAMEL
CLS~1000~10000~flavor_yn~y	CINNAMON
CLS~1000~10000~flavor_yn~y	HAZELNUT
CLS~1000~10000~flavor_yn~n	PLAIN
CLS~1000~10000~flavor_type~non	PLAIN
CLS~1000~10000~flavor_type~fruit	BLUEBERRY
CLS~1000~10000~flavor_type~fruit	RASPBERRY
CLS~1000~10000~flavor_type~mild	HAZELNUT
CLS~1000~10000~flavor_type~mild	VANILLA
CLS~1000~10000~flavor_type~special	CINNAMON
CLS~1000~10000~flavor_type~special	CARAMEL

Market Basket Insights

This chapter describes Market Basket Insights Cloud Service.

Overview

Market Basket Insights (MBI) is used to gain insights into customer shopping patterns. A key component of MBI is the process of Association Rule Mining (ARM). This process examines sales transaction data and identifies associations between types of products. Such information can help a retailer understand that promoting one product is sufficient to help drive sales of another product, given the sales associations they exhibit.

The processing of these algorithms occurs each week as part of the weekly batch execution, and a set of output files are provided to expose the association rules that have been discovered by the process.

Data Requirements

Market Basket Insights relies on the following data elements. These must be provided via text files, which are then loaded.

Table 12–1 Data Elements

Object	Notes	Required/Optional
Product Hierarchy	The ARM processing mainly operates at Sub Class, but it can be configured to different levels.	Required
Location Hierarchy		Required
Fiscal Calendar		Required
Sales Transactions	Must contain transaction IDs as part of the data. If the transactions include Customer ID, then customer segment results are possible.	Required
Customer Segments	Customer IDs and their association to a segment allows customer segment-specific results.	Optional
MBA_ARM_SRVC_LOC_STG	Can be used to limit the scope of locations processed, or to specify a set of locations to exclude from processing.	Optional

In order to calculate association rules, it is necessary to receive sales transaction data that include a transaction ID. This is used to identify which products were purchased

by a customer as part of a single transaction. If the customer transactions also include a customer ID to identify the customer who purchased the transaction, and a customer segment dimension is provided that links customer IDs to customer segments, then it is possible to provide some results for each customer segment.

MBA_ARM_SRVC_LOC_STG

When specifying which locations to process or which locations to not process, the MBA_ARM_SRVC_LOC_STG interface can be used to limit the scope of locations to be processed. The data in this interface can be at any level of the location hierarchy. A customer may want to limit the scope of locations for the following reasons.

- Improve performance by only sampling some locations.
- Exclude locations that contain many wholesale transactions, where the transactions contain data for more than a single customer.
- Exclude locations that are experiencing a significant interruption to their normal sales pattern (for example, when undergoing a large scale renovation).
- Exclude locations that normally do not include customer-linked transactions from the ARM_PH_CS implementation, since suitable data to include for processing will not be available.

The SRVC_NAME column of this interface allows the specification of the service that must be filtered. If, however, all executions must have the same set of locations, then this column can be provided as a NULL value. The effect will be to use the same dataset for all the services. If, however, it is necessary to have some services use a different set of locations, then it is possible to provide the data specific to the different services. If data is provided for a SRVC_NAME, then the data must be provided with a SRVC_NAME specified. The valid SRVC_NAME values that can be provided are: ARM_PH (Product Hierarchy results), ARM_PH_PROMO (Product Hierarchy with Promotions results), and ARM_PH_CS (Product Hierarchy and Customer Segment, with Promotions results).

Science Algorithms/Services

This section describes the science algorithms and services.

ARM_PH

This implementation calculates association rules for a configurable set of product hierarchy levels. It supports the creation of association rules for Sub Classes, Classes, and Department, which can be controlled by a system configuration. All system configurations that affect this algorithm exist in the RSE_CONFIG configuration table, and are named with "ARM_PH_" as the prefix. Because this implementation supports being run for multiple hierarchy levels, if there is a need to set a configuration uniquely for a specific hierarchy level, this can be accomplished via the RSE_CONFIG_CODE table using the hierarchy level name as the PARAM_CODE value. If no such row exists in RSE_CONFIG_CODE, then the configuration will be taken from the corresponding RSE_CONFIG row.

ARM_PH_PROMO

This implementation calculates association rules at the Sub Class level of the hierarchy and is restricted to only rules where the IF side of the rule is promoted and the THEN side of the rule is not promoted. In order to be able to execute this and have results for

this implementation, it is necessary to provide promotion details with the sales transaction data. All system configurations that affect this algorithm exist in the RSE_CONFIG configuration table and are named with "ARM_PROMO_" as the prefix.

This implementation is used to focus on how products are associated when that promotion is in effect. This data can help a retailer understand the sales patterns that exist when promotions are involved, which can help the retailer avoid promoting too many items in an effort to help improve profit.

ARM_PH_CS

This implementation calculates association rules at the Sub Class level of the hierarchy and is restricted to rules where the IF side of the rule is promoted and the THEN side of the rule is not promoted. In order to be able to execute this and have results for this implementation, it is necessary to provide promotion details with the sales transaction data. All system configurations that affect this algorithm exist in the RSE_CONFIG configuration table and are named with "ARM_CS_" as a prefix.

This implementation provides the same type of information as the ARM_PH_PROMO implementation; however, it provides results that are specific to a customer segment. Therefore, this implementation requires the receipt of transactions that include the customer ID of the customer who purchased the transaction and the customer segment dimension, along with the association of the customers to each customer segment.

Configurations

There is a consistent pattern in the naming of the configurations for the MBI implementation. As described above, each implementation has a specific naming prefix. The suffixes are also similar across the implementations. These suffixes are described in more detail in [Table 12-2](#).

Table 12-2 Implementation Suffixes

Suffix	Example	Description
HIER_LEVEL	SBC	Indicates the name of the hierarchy level that the process is to be executed for. The values here are the same values as provided as LEVEL_NAME values in the W_PROD_CAT_DH interface. Not applicable to ARM_PH.
TOP_LEVEL	SBC	Indicates the highest level of the product hierarchy for which processing should be executed. Can contain SBC, CLS, or DEPT. Only applicable to ARM_PH.
MIN_SUPPORT	.001	Expresses the minimum percentage of transactions that are required to have the set of items in the same transaction.
MIN_SUPPORT_TXN_CNT	1000	In the event that sales volume is low, this is another way to express the minimum number of sales transactions that are required for the set of items to be sold together. The implementation uses the greater of the two values.
MIN_CONFIDENCE	.05	The minimum confidence value as calculated by the rule mining algorithm for an association rule.

Table 12–2 (Cont.) Implementation Suffixes

Suffix	Example	Description
MIN_REV_CONFIDENCE	.05	The minimum confidence as calculated by reversing the placement of the numbers in the calculation. Setting this value higher can help prevent redundancy in the rule expressions where the IF and THEN items are transposed.
MIN_LIFT	.05	The minimum lift as calculated by the rule mining algorithm for an association rule.
MAX_LIFT	100	The maximum lift as calculated by the rule mining algorithm for an association rule.
MAX_SET_SIZE	2	The maximum number of hierarchy members to include in the resulting rules. The set size includes the count of both the IF and THEN components. The maximum allowed is four, although it can be an expensive to calculate that many components.
MAX_RULE_COUNT	9999	The maximum number of rules that are retained per execution of the algorithm, per set size. This allows for the reduction of results to eliminate less important results.
WEEK_CNT	1	The number of weeks that are processed when the execution runs. Care should be taken when changing this to more than one week, as this can negatively affect performance.

Data Output

The results of the association rule mining can be obtained from two export interface files. One export interface contains summary information (`mba_arm_run_exp`) about the execution, along with various metrics that explain what the results are for. The second export interface file (`mba_arm_result_exp`) contains the details for each execution of the process. It is possible that a run may not contain any results to be exported. The data between the two interfaces can be joined to each other by the first column in each interface file (the `RUN_ID`).

In addition to metrics that quantify the rule (its frequency, its confidence, and its lift), the results also include sales values for the different components of the association rule. These sales values can help quantify the involved sales volume that is involved in the association rule.

Even if the weekly process that runs is executed for a single week each time, it is still possible to estimate the effects of the rule across multiple weeks by aggregating data across the weeks. The process for doing this requires locating the same product set across the different weeks within the same execution type. This means to join data in the `mba_arm_result_exp` interface by `if_prod_ext_key1`, `if_prod_ext_key2`, `if_prod_ext_key3`, `if_promo_flg1`, `if_promo_flg2`, `if_promo_flg3`, `then_prod_ext_key`, `then_promo_flg`, and the data in the `mba_arm_run_exp` interface by `run_type`, `if_hier_level`, `then_hier_level`, `loc_ext_key`, `custseg_ext_key`. The data between the `mba_arm_run_exp` and `mba_arm_result_exp` files are joined by the `run_id`.

Once this appropriate data has been gathered, the various sales metrics can be aggregated as needed. In order to calculate a new set of Frequency, Confidence, Lift, or Reverse Confidence values for a rule, it is possible to recalculate the values, as shown below. Note that in the these calculations, the following abbreviations are used: `run = mba_arm_run_ext`, `result = mba_arm_result_exp`.

Frequency = $\text{SUM}(\text{result.rule_txn_count}) / \text{SUM}(\text{run.tot_txn_cnt})$

Confidence = $\text{SUM}(\text{result.rule_txn_cnt}) / \text{SUM}(\text{result.if_tot_txn_count})$

Reverse Confidence = $\text{SUM}(\text{result.rule_txn_cnt}) / \text{SUM}(\text{result.then_tot_txn_count})$

Lift = $\text{SUM}(\text{result.rule_txn_cnt}) * \text{SUM}(\text{run.tot_txn_cnt}) / \text{SUM}(\text{result.if_tot_txn_count}) / \text{SUM}(\text{result.then_tot_txn_count})$

Returns Logistics

This chapter provides an overview of the Returns Logistics or Returns Optimization capabilities available for Oracle Retail Advanced Science.

Overview

The Returns Logistics process takes data inputs such as price, inventory, shipping costs, price ladders, sales, and returns information and passes them into a scientific algorithm that makes recommendations about products that have been returned. These recommendations include where to ship them from, where to ship them to, and in what quantities.

The bulk of the processing occurs outside of the database; however, the correct data must be in place in order for the algorithm to produce meaningful output. Internal database processes reside in Oracle Retail Advanced Science in order to serve the data to the external algorithm.

The Returns Logistics process runs weekly as a batch process for products in the top x selling categories, where x is a database configuration parameter stored in the RSE_CONFIG table.

Data Inputs

The Returns Logistics process uses data from both RI and ORASE databases. Price, inventory, sales, and returns data are all pulled directly from the RI database before being served to the external algorithm. The configuration parameters for the process are stored in Oracle Retail Advanced Science. For example, the process can be executed on the value for the parameter Top Selling Categories. The data for the top selling categories (which in turn determine the subset list of products), resides solely within the application. Price ladders, shipping costs, and price elasticity are all data sources provided directly within the application.

Table 13–1 Data Inputs

Data Input	Source	
	Database	Source Table
Current Price	RI	W_RTL_PRICE_IT_LC_G
Regular Price	RI	W_RTL_PRICE_IT_DY_F
Inventory	RI	W_RTL_INV_IT_LC_G
Sales	RI	W_RTL_SLS_IT_LC_WK_A
Returns	RI	W_RTL_SLS_IT_LC_WK_A

Table 13–1 (Cont.) Data Inputs

Data Input	Source Database	Source Table
Season	RI	W_RTL_SEASON_D
Season Products	RI	W_RTL_SEASON_IT_D
Product	RI	W_PRODUCT_D
Calendar	RI	W_MCAL_DAY_D, W_MCAL_WEEK_D
Shipping Costs	ORASE	RL_SHIPPING_COST
Price Ladder	ORASE	RL_PRICE_LADDER
Price Elasticity	ORASE	RL_PRICE_ELASTICITY
Top Selling Categories	ORASE	RSE_SLS_PH_LC_WK_A
Calendar Hierarchy	ORASE	RSE_CAL_SRC_XREF
Product Hierarchy	ORASE	RSE_PROD_HIER_TC, RSE_PROD_SRC_XREF
Location Hierarchy	ORASE	RSE_LOC_SRC_XREF

Data Outputs

The external algorithm of the Returns Logistics process writes recommendations back into the application database. Database tables are written to that provide shipping recommendations (for a quantity of product from location to location) and also the optimal demand and revenue of a product at a location.

These recommendations answer the question "What should be done with the products that have been returned?" Answers include shipping them from one store to another where demand may be higher.

Table 13–2 Data Outputs

Data Output	Database	Table
Shipping Recommendation	ORASE	RL_SHIPPING_RECOMMENDATION
Optimal Demand	ORASE	RL_OPT_DEM_REV

In addition to the tables, two corresponding database views are provided that can be used as the basis for an export process for example, to RI) as required.

Data Interfaces

Oracle Retail Advanced Science requires several interfaces in order to support the various applications it provides. This chapter does not duplicate the interface documentation, but instead focuses on identifying the tables and columns that it requires data in. For additional details about any interfaces, see Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface.

Inbound Interface Requirements by Application

Advanced Clustering

Table 14–1 Advanced Clustering Inbound

Logical Group	Inbound File	Notes
RI	W_MCAL_PERIOD_DS.dat	
RI	W_RTL_CMG_PRODUCT_MTX_DS.dat	Optional. Only required if the product hierarchy is not a suitable representation for a category.
RI	W_RTL_CUSTSEG_DS.dat	This file is not required if Customer Segmentation is used or if W_RTL_SLS_TRX_IT_LC_DY_FS is not used.
RI	W_RTL_CUST_CUSTSEG_DS.dat	This file is not required if Customer Segmentation is used or if W_RTL_SLS_TRX_IT_LC_DY_FS is not used.
RI	W_PARTY_ORG_DS.dat	An empty file is sufficient.
RI	W_PARTY_PER_DS.dat	Optional. Only required if loading data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_INT_ORG_DS.dat	
RI	W_INT_ORG_DHS.dat	
RI	W_INT_ORG_DS_TL.dat	
RI	W_PRODUCT_DS.dat	
RI	W_PRODUCT_ATTR_DS.dat	
RI	W_RTL_ITEM_GRP1_DS.dat	
RI	W_PROD_CAT_DHS.dat	

Table 14–1 (Cont.) Advanced Clustering Inbound

Logical Group	Inbound File	Notes
RI	W_PRODUCT_DS_TL.dat	
RI	W_DOMAIN_MEMBER_DS_TL.dat	
RI	W_RTL_PROMO_DS.dat	Required due to dependencies from data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_RTL_PROMO_DS_TL.dat	Required due to dependencies from data in W_RTL_PROMO_DS.
RI	W_RTL_PROMO_COMP_TYPE_DS.dat	Required due to dependencies from data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_EXCH_RATE_GS.dat	Optional. Required if W_RTL_SLS_TRX_IT_LC_DY_FS is used. This can be a one-time load. It needs at minimum to contain one row, even if only one currency is used.
RI	W_RTL_SLS_TRX_IT_LC_DY_FS.dat	Optional. If sales transaction details are not desired to be loaded, then RSE_SLS_PR_LC_WK_STG and RSE_SLS_PR_LC_CS_WK_STG interfaces can be used instead.
RI	W_EMPLOYEE_DS.dat	Empty file is sufficient unless W_RTL_SLS_TRX_IT_LC_DY_FS contains employee identifiers.
RI	W_REASON_DS.dat	Empty file is sufficient unless customer orders are loaded via the W_RTL_CO_HEAD_DS interface.
RI	W_RTL_CO_LINE_DS.dat	Empty file is W_RTL_SLS_TRX_IT_LC_DY_FS does not include customer orders.
RI	W_RTL_CO_HEAD_DS.dat	Empty file is sufficient if W_RTL_SLS_TRX_IT_LC_DY_FS does not include customer orders.
RI	W_RTL_CO_HEAD_STATUS_FS.dat	Empty file is sufficient if no customer orders are being loaded into W_RTL_CO_HEAD_DS.
RI	W_RTL_TRADE_AREA_DS.dat	Optional. Only required if the organization hierarchy is not suitable for grouping similar locations.
RI	W_RTL_TRADE_ATRA_LOC_MTX_DS.dat	Optional. Only required if the W_RTL_TRADE_AREA_DS is used.
RI	W_RTL_ITEM_DEL_TMP.dat	
RI	W_RTL_IT_LC_DEL_TMP.dat	
RI	W_RTL_RECLASS_DP_GP_TMP.dat	
RI	W_RTL_RECLASS_IT_SC_CL_TMP.dat	
RI	W_RTL_CUSTSEG_ALLOC_DS.dat	
RI	W_RTL_CONSUMERSEG_DS.dat	

Table 14–1 (Cont.) Advanced Clustering Inbound

Logical Group	Inbound File	Notes
RI	W_RTL_SLSFC_IT_LC_WK_F.dat	An empty file is sufficient unless store clustering is expected to use forecast data as an input.
RI	W_RTL_CHANNEL_DS.dat	This is a one-time file load.
ORASE	rse_prod_attr_grp_value_stg.txt	
ORASE	rse_prod_attr_value_xref_stg.txt	
ORASE	rse_sls_pr_lc_wk_stg.txt	Optional. Only required if data is not loaded via W_RTL_SLS_TRX_IT_LC_DY_FS.
ORASE	rse_sls_pr_lc_cs_wk_stg.txt	Optional. Only required if data is not loaded via W_RTL_SLS_TRX_IT_LC_DY_FS.
AC	rse_like_loc_stg.txt	

Assortment and Space Optimization

Table 14–2 Assortment and Space Optimization Inbound

Logical Group	Inbound File	Notes
RI	W_MCAL_PERIOD_DS.dat	
RI	W_RTL_CMG_PRODUCT_MTX_DS.dat	Optional. Only required if the product hierarchy is not a suitable representation for a category.
RI	W_PARTY_PER_DS.dat	
RI	W_INT_ORG_DS.dat	
RI	W_INT_ORG_DHS.dat	
RI	W_INT_ORG_DS_TL.dat	
RI	W_PRODUCT_DS.dat	
RI	W_PRODUCT_ATTR_DS.dat	
RI	W_RTL_PRODUCT_IMAGE_DS.dat	
RI	W_RTL_ITEM_GRP1_DS.dat	
RI	W_PROD_CAT_DHS.dat	
RI	W_PRODUCT_DS_TL.dat	
RI	W_DOMAIN_MEMBER_DS_TL.dat	
RI	W_RTL_TRADE_AREA_DS.dat	Optional. Only required if the organization hierarchy is not suitable for grouping similar locations.
RI	W_RTL_TRADE_ATRA_LOC_MTX_DS.dat	Optional. Only required if the W_RTL_TRADE_AREA_DS is used.
RI	W_RTL_ITEM_DEL_TMP.dat	
RI	W_RTL_IT_LC_DEL_TMP.dat	
RI	W_RTL_RECLASS_DP_GP_TMP.dat	

Table 14–2 (Cont.) Assortment and Space Optimization Inbound

Logical Group	Inbound File	Notes
RI	W_RTL_RECLASS_IT_SC_CL_TMP.dat	
ORASE	rse_prod_attr_grp_value_stg.txt	
ASO	so_assortment_finalized_stg.txt	
ASO	so_assort_phprod_finalized_stg.txt	
ASO	so_assortment_stg.txt	
ASO	so_assort_cluster_stg.txt	
ASO	so_assort_cluster_member_stg.txt	
ASO	so_assort_phprod_like_prod_stg.txt	
ASO	so_assort_product_strcltr_stg.txt	
ASO	so_assort_proloc_pricecost_stg.txt	
ASO	so_assort_proloc_fcst_stg.txt	
ASO	so_assort_phprod_attr_stg.txt	
ASO	so_pog_stg.txt	
ASO	so_pog_store_stg.txt	
ASO	so_pog_store_cda_stg.txt	
ASO	so_pog_bay_stg.txt	
ASO	so_prod_display_style_stg.txt	
ASO	so_display_style_stg.txt	
ASO	so_pog_display_style_stg.txt	
ASO	so_fixture_stg.txt	
ASO	so_bay_fixture_stg.txt	
ASO	so_shelf_stg.txt	
ASO	so_bay_fixture_shelf_stg.txt	
ASO	so_disp_style_orientation_stg.txt	
ASO	so_display_style_fixture_stg.txt	
ASO	so_fixture_disp_config_stg.txt	
ASO	so_pegboard_disp_config_stg.txt	
ASO	so_prod_loc_repl_param_stg.txt	
ASO	so_prod_stack_height_limit_stg.txt	
ASO	so_pog_assort_mapping_stg.txt	
ASO	so_pog_assort_seas_mapping_stg.txt	

Customer Decision Tree

Table 14–3 Customer Decision Tree Inbound

Logical Group	Inbound File	Notes
RI	W_MCAL_PERIOD_DS.dat	
RI	W_RTL_CMG_PRODUCT_MTX_DS.dat	Optional. Only required if the product hierarchy is not a suitable representation for a category.
RI	W_RTL_CUSTSEG_DS.dat	If Customer Segmentation is used, this file does not need to be provided externally.
RI	W_RTL_CUST_CUSTSEG_DS.dat	If Customer Segmentation is used, this file does not need to be provided externally.
RI	W_PARTY_PER_DS.dat	
RI	W_INT_ORG_DS.dat	
RI	W_INT_ORG_DHS.dat	
RI	W_INT_ORG_DS_TL.dat	
RI	W_PRODUCT_DS.dat	
RI	W_PRODUCT_ATTR_DS.dat	
RI	W_RTL_ITEM_GRP1_DS.dat	
RI	W_PROD_CAT_DHS.dat	
RI	W_PRODUCT_DS_TL.dat	
RI	W_DOMAIN_MEMBER_DS_TL.dat	
RI	W_RTL_PROMO_DS.dat	Required due to dependencies from data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_RTL_PROMO_COMP_TYPE_DS.dat	Required due to dependencies from data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_EXCH_RATE_GS.dat	
RI	W_RTL_SLS_TRX_IT_LC_DY_FS.dat	
RI	W_RTL_CO_LINE_DS.dat	Only required if rows in W_RTL_SLS_TRX_IT_LC_DY_FS include customer orders.
RI	W_RTL_CO_HEAD_DS.dat	Only required if rows in W_RTL_SLS_TRX_IT_LC_DY_FS include customer orders.
RI	W_RTL_TRADE_AREA_DS.dat	Optional. Only required if the organization hierarchy is not suitable for grouping similar locations.
RI	W_RTL_TRADE_ATRA_LOC_MTX_DS.dat	Optional. Only required if the W_RTL_TRADE_AREA_DS is used.
RI	W_RTL_ITEM_DEL_TMP.dat	
RI	W_RTL_IT_LC_DEL_TMP.dat	
RI	W_RTL_RECLASS_DP_GP_TMP.dat	
RI	W_RTL_RECLASS_IT_SC_CL_TMP.dat	

Table 14–3 (Cont.) Customer Decision Tree Inbound

Logical Group	Inbound File	Notes
ORASE	rse_prod_attr_grp_value_stg.txt	
ORASE	rse_prod_attr_value_xref_stg.txt	
ORASE	rse_fake_cust_stg.txt	Optional interface to designate customer IDs as either fake or not.
CDT	cdt_import.tar.gz	Optional interface to import CDTs created outside the application.

Demand Transference

Table 14–4 Demand Transference Inbound

Logical Group	Inbound File	Notes
RI	W_MCAL_PERIOD_DS.dat	
RI	W_RTL_CMG_PRODUCT_MTX_DS.dat	Optional. Only required if the product hierarchy is not a suitable representation for a category.
RI	W_RTL_CUSTSEG_DS.dat	If Customer Segmentation is used, this file does not need to be provided externally.
RI	W_RTL_CUST_CUSTSEG_DS.dat	If Customer Segmentation is used, this file does not need to be provided externally.
RI	W_PARTY_PER_DS.dat	Optional. Only required if loading data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_INT_ORG_DS.dat	
RI	W_INT_ORG_DHS.dat	
RI	W_INT_ORG_DS_TL.dat	
RI	W_PRODUCT_DS.dat	
RI	W_PRODUCT_ATTR_DS.dat	
RI	W_RTL_ITEM_GRP1_DS.dat	
RI	W_PROD_CAT_DHS.dat	
RI	W_PRODUCT_DS_TL.dat	
RI	W_DOMAIN_MEMBER_DS_TL.dat	
RI	W_RTL_PROMO_DS.dat	Required due to dependencies from data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_RTL_PROMO_DS_TL.dat	Required due to dependencies from data in W_RTL_PROMO_DS.
RI	W_RTL_PROMO_COMP_TYPE_DS.dat	Required due to dependencies from data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_EXCH_RATE_GS.dat	Optional. Required if W_RTL_SLS_TRX_IT_LC_DY_FS is used.

Table 14–4 (Cont.) Demand Transference Inbound

Logical Group	Inbound File	Notes
RI	W_RTL_SLS_TRX_IT_LC_DY_FS.dat	Optional. If sales transaction details are not desired to be loaded, then RSE_SLS_PR_LC_WK_STG and RSE_SLS_PR_LC_CS_WK_STG interfaces can be used instead.
RI	W_RTL_CO_LINE_DS.dat	Only required if rows in W_RTL_SLS_TRX_IT_LC_DY_FS include customer orders.
RI	W_RTL_CO_HEAD_DS.dat	Only required if rows in W_RTL_SLS_TRX_IT_LC_DY_FS include customer orders.
RI	W_RTL_CO_HEAD_STATUS_FS.dat	Empty file is sufficient if there are no customer orders being loaded into W_RTL_CO_HEAD_DS.
RI	W_RTL_TRADE_AREA_DS.dat	Optional. Only required if the organization hierarchy is not suitable for grouping similar locations.
RI	W_RTL_TRADE_ATRA_LOC_MTX_DS.dat	Optional. Only required if the W_RTL_TRADE_AREA_DS is used.
RI	W_RTL_IT_LC_DS.dat	
RI	W_RTL_ITEM_DEL_TMP.dat	
RI	W_RTL_IT_LC_DEL_TMP.dat	
RI	W_RTL_RECLASS_DP_GP_TMP.dat	
RI	W_RTL_RECLASS_IT_SC_CL_TMP.dat	
ORASE	rse_prod_attr_grp_value_stg.txt	
ORASE	rse_prod_attr_value_xref_stg.txt	
ORASE	rse_sls_pr_lc_wk_stg.txt	Optional. Only required if data is not loaded via W_RTL_SLS_TRX_IT_LC_DY_FS.
ORASE	rse_sls_pr_lc_cs_wk_stg.txt	Optional. Only required if data is not loaded via W_RTL_SLS_TRX_IT_LC_DY_FS.
DT	dt_loc_wk_excl_stg.txt	
DT	dt_prod_loc_excl_stg.txt	
DT	dt mdl_prod_exp_stg.txt	Quarterly interface.

Customer Segmentation

Table 14–5 Customer Segmentation Inbound

Logical Group	Inbound File	Notes
RI	W_MCAL_PERIOD_DS.dat	
RI	W_RTL_CMG_PRODUCT_MTX_DS.dat	Optional. Only required if the product hierarchy is not a suitable representation for a category.

Table 14–5 (Cont.) Customer Segmentation Inbound

Logical Group	Inbound File	Notes
RI	W_PARTY_ORG_DS.dat	An empty file is sufficient.
RI	W_PARTY_PER_DS.dat	
RI	W_INT_ORG_DS.dat	
RI	W_INT_ORG_DHS.dat	
RI	W_INT_ORG_DS_TL.dat	
RI	W_PRODUCT_DS.dat	
RI	W_RTL_ITEM_GRP1_DS.dat	
	W_PROD_CAT_DHS.dat	
RI	W_PRODUCT_DS_TL.dat	
RI	W_DOMAIN_MEMBER_DS_TL.dat	
RI	W_RTL_PROMO_DS.dat	Required due to dependencies from data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_RTL_PROMO_DS_TL.dat	Required due to dependencies from data in W_RTL_PROMO_DS.
RI	W_RTL_PROMO_COMP_TYPE_DS.dat	Required due to dependencies from data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_EXCH_RATE_GS.dat	This can be a one-time load, and needs to contain at minimum one row, even if only one currency is used.
RI	W_RTL_SLS_TRX_IT_LC_DY_FS.dat	
RI	W_EMPLOYEE_DS.dat	Empty file is sufficient unless W_RTL_SLS_TRX_IT_LC_DY_FS contains employee identifiers.
RI	W_REASON_DS.dat	Empty file is sufficient unless customer orders are loaded via the W_RTL_CO_HEAD_DS interface.
RI	W_RTL_CO_LINE_DS.dat	Empty file is sufficient if W_RTL_SLS_TRX_IT_LC_DY_FS does not include customer orders.
RI	W_RTL_CO_HEAD_DS.dat	Empty file is sufficient if W_RTL_SLS_TRX_IT_LC_DY_FS does not include customer orders.
RI	W_RTL_CO_HEAD_STATUS_FS.dat	Empty file is sufficient if no customer orders are being loaded into W_RTL_CO_HEAD_DS.
RI	W_RTL_TRADE_AREA_DS.dat	Optional. Only required if the organization hierarchy is not suitable for grouping similar locations.
RI	W_RTL_TRADE_ATRA_LOC_MTX_DS.dat	Optional. Only required if the W_RTL_TRADE_AREA_DS is used.
RI	W_RTL_ITEM_DEL_TMP.dat	
RI	W_RTL_IT_LC_DEL_TMP.dat	
RI	W_RTL_RECLASS_DP_GP_TMP.dat	

Table 14–5 (Cont.) Customer Segmentation Inbound

Logical Group	Inbound File	Notes
RI	W_RTL_RECLASS_IT_SC_CL_TMP.dat	
RI	W_RTL_CHANNEL_DS.dat	This is a one-time file load.

Returns Logistics

Table 14–6 Return Logistics Inbound

Logical Group	Inbound File	Notes
RL	rl_price_elasticity_stg.txt	Required. The price elasticity for each product-location. If none available, recommendation is to load default value of 2.
RL	rl_price_ladder_stg.txt	Required. Price ladders are all the valid price points for a product. No location information is considered, only product.
RL	rl_shipping_cost_stg.txt	Required. These are the shipping costs per unit of product from one location to another. Individual product detail is not considered.

Market Basket Insights

Table 14–7 MBI Inbound Interface

Logical Group	Inbound File	Notes
RI	W_MCAL_PERIOD_DS.dat	
RI	W_PRODUCT_DS.dat	
RI	W_PRODUCT_ATTR_DS.dat	
RI	W_PRODUCT_DS_TL.dat	
RI	W_PARTY_PER_DS.dat	
RI	W_RTL_PROD_HIER_IMAGE_DS.dat	An empty file is sufficient.
RI	W_PROD_CAT_DHS.dat	
RI	W_RTL_PROD_HIER_ATTR_LKP_DHS.dat	An empty file is sufficient.
RI	W_RTL_ITEM_GRP1_DS.dat	
RI	W_INT_ORG_DHS.dat	
RI	W_RTL_LOC_TRAITS_DS_TL.dat	An empty file is sufficient.
RI	W_RTL_PROMO_DS.dat	Required due to dependencies from data in W_RTL_SLS_TRX_IT_LC_DY_FS.

Table 14–7 (Cont.) MBI Inbound Interface

Logical Group	Inbound File	Notes
RI	W_RTL_PROMO_DS_TL.dat	Required due to dependencies from data in W_RTL_SLS_TRX_IT_LC_DY_FS.
RI	W_RTL_CHANNEL_DS.dat	This is a one-time file load.
RI	W_EMPLOYEE_DS.dat	Empty file is sufficient unless W_RTL_SLS_TRX_IT_LC_DY_FS contains employee identifiers.
RI	W_EXCH_RATE_GS.dat	This can be a one-time load. It needs at minimum to contain one row, even if only one currency is used.
RI	W_INT_ORG_DS.dat	
RI	W_INT_ORG_ATTR_DS.dat	
RI	W_INT_ORG_DS_TL.dat	
RI	W_PARTY_ORG_DS.dat	An empty file is sufficient.
RI	W_DOMAIN_MEMBER_DS_TL.dat	
RI	W_REASON_DS.dat	Empty file is sufficient unless customer orders are loaded via the W_RTL_CO_HEAD_DS interface.
RI	W_RTL_PRODUCT_COLOR_DS.dat	An empty file is sufficient.
RI	W_RTL_PRODUCT_ATTR_DS.dat	An empty file is sufficient.
RI	W_RTL_PRODUCT_ATTR_DS_TL.dat	An empty file is sufficient.
RI	W_RTL_PRODUCT_BRAND_DS.dat	An empty file is sufficient.
RI	W_RTL_PRODUCT_BRAND_DS_TL.dat	An empty file is sufficient
RI	W_RTL_CMG_PRODUCT_MTX_DS.dat	Empty file is sufficient.
RI	W_RTL_CO_HEAD_DS.dat	Empty file is sufficient if W_RTL_SLS_TRX_IT_LC_DY_FS does not include customer orders.
RI	W_RTL_CO_HEAD_STATUS_FS.dat	Empty file is sufficient if W_RTL_SLS_TRX_IT_LC_DY_FS does not include customer orders.
RI	W_RTL_CO_LINE_DS.dat	Empty file is sufficient if W_RTL_SLS_TRX_IT_LC_DY_FS does not include customer orders.
RI	W_RTL_CO_LINE_STATUS_FS.dat	Empty file is sufficient if W_RTL_SLS_TRX_IT_LC_DY_FS does not include customer orders.
RI	W_RTL_CO_SHIP_TYPE_DS.dat	Empty file is sufficient if W_RTL_SLS_TRX_IT_LC_DY_FS does not include customer orders.
RI	W_RTL_TRADE_AREA_DS.dat	Optional.
RI	W_RTL_TRADE_AREA_LOC_MTX_DS.dat	Optional.

Table 14–7 (Cont.) MBI Inbound Interface

Logical Group	Inbound File	Notes
RI	W_RTL_RECLASS_IT_SC_CL_TMP.dat	
RI	W_RTL_IT_LC_DEL_TMP.dat	
RI	W_RTL_RECLASS_DP_GP_TMP.dat	
RI	W_RTL_ITEM_DEL_TMP.dat	
RI	W_RTL_SLS_TRX_IT_LC_DY_FS.dat	
RI	W_RTL_CUSTSEG_DS.dat	This file would not be required if Customer Segmentation Cloud Service is used to create customer segments.
RI	W_RTL_CUST_CUSTSEG_DS.dat	This file would not be required if Customer Segmentation Cloud Service is used to create customer segments.
ORASE	rse_prod_attr_grp_value_stg.txt	Optional.
ORASE	rse_prod_attr_value_xref_stg.txt	Optional.
ORASE	mba_arm_srvc_loc_stg.txt	Optional. An empty file is sufficient, unless it is necessary to restrict locations to be processed.

RASE Web Services

RASE web services provide access to DB-bound configurations when a direct connection to the database is not desirable or possible; intra-day and ad-hoc access to certain application outputs is also provided. RASE web services are REST-based; it is assumed that you are familiar with basic REST principles (such as the usage of HTTP verbs). RASE web services provide access to a subset of application and output data, but do not fully mirror the user interface or export and import features of the backend. They are not a replacement for bulk data export, which must be done on a schedule as part of batch processing. However, access to the configuration can be used during implementation and upgrade periods, and AC and ASO export web services can serve as a means of obtaining incremental update data from a specified point in time (driven by a query parameter) as a means of intra-day processing.

All services support the query parameter `contentType` and the HTTP header `Content-Type`, with supported values `application/json` and `application/xml`. The query parameter takes precedence; if no content type is supplied, then `application/json` serves as the default.

The AC and ASO export services have a `dateMask` query parameter that must adhere to the Java `java.text.SimpleDateFormat` rules (for example, `dateMask=yyyyMMdd&exportDate=20151012`). All date parameters must be sent in this format, and output dates and timestamps are returned according to this format.

The Json/XML structure follows the corresponding DB table and view while being converted to Java standard. That is, underscores are removed, camel case is used, and first letter is lowercase. For example, the `RSE_CONFIG APPL_CODE` column is returned as `applCode`. All data is returned as type string.

Authentication and Authorization

Basic authentication is used, so you may use any client software that supports it. Authorization is done for ADF-LDAP (OID) mapped roles, and only administrator roles are used (that is, the calling user must be in a duty that is mapped to the roles in [Table 14–8](#)).

Table 14–8 Mapped Administrator Roles

Service	Role	Mapped Role
DT	DemandTransferenceRole	ANALYTIC_EXPERT_JOB
CDT	CustomerDecisionTreeRole	ANALYTIC_EXPERT_JOB
AC	StoreClusterAdvancedRole	CLUSTERING_ADMINISTRATOR_JOB
ASO	Administrator	SPACE_ADMINISTRATOR_JOB
CS	CustomerSegmentAdvancedRole	CUSTOMER_SEGMENT_ADMINISTRATOR_JOB

Summary of Web Services

This section provides a summary of web services.

Access to RSE_CONFIG Table

Fetch Config Data

GET on `/rase/resources/rse/parameters` returns all RSE_CONFIG entries this user has access to (that is, all applCode RSE entries plus corresponding module entries). For example, if a user is in the ASO Administrator role, then the applCode SO will be returned as well. Here is a list of the fields:

- applCode
- paramName
- paramValue
- configurableFlg.
- descr

Fetch One Entry

GET on `applCode,paramName`. For example, `/rase/resources/rse/parameters/RSE,PRIMARY_LANGUAGE_CODE` returns a particular entry, if the user has access to it. See [Fetch Config Data](#).

Create an Entry

POST to `/rase/resources/rse/parameters` with a form having a single field called "content" and a value having the appropriate new parameter data in content type as specified by query parameter and header, if the user has access to it. See [Fetch Config Data](#). The structure of the content must match what is returned by GET.

Update an Entry

PUT to `applCode,paramName`. For example, `/rase/resources/rse/parameters/RSE,LOC_HIER_TYPE` with the form having fields called `paramName`, `paramValue`, `descr`, and `configurableFlg`, along with appropriate values, if the user has access to it. See [Fetch Config Data](#).

Batch Update Entries

PUT to /rase/resources/rse/parameters with a form having a single field called "content" and a value having appropriate parameter data in the content type as specified by query parameter and header. Note that if the user does not have access to a particular entry, it will be skipped. See [Fetch Config Data](#). The structure of the content must match what is returned by GET.

Access to RSE_CONFIG_CODE Table

Fetch Data for an Entry

GET on applCode,paramName, paramCode. For example, /rase/resources/rse/parameters/DT,SIM_DISPLAY_CODE_PCT,2 returns data, if user has access to it. See [Fetch Config Data](#).

- applCode
- paramName
- paramCode
- paramValue
- configurableFlg
- descr

Create an Entry

POST to applCode,paramName,paramCode. For example, /rase/resources/rse/parameters/DT,SIM_DISPLAY_CODE_PCT,2} with a form having a single field named "content" and a value having appropriate new parameter data in the content type as specified by query parameter and header, if the user has access to it. See [Fetch Config Data](#). The structure of the content must match what is returned by GET.

Update an Entry

PUT to applCode,paramName, paramCode, For example, /rase/resources/rse/parameters/DT,SIM_DISPLAY_CODE_PCT,2 with form having a single field named "content" and value having appropriate parameter data in content type as specified by query parameter / header, if user has access to it. See [Fetch Config Data](#). The structure of the content must match what is returned by GET.

Delete an Entry

DELETE to applCode,paramName, paramCode. For example, /rase/resources/rse/parameters/DT,SIM_DISPLAY_CODE_PCT,2, if the user has access to it. See [Fetch Config Data](#).

Advanced Clustering Export

Get Clusters

This service is based on rsestrclst.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/cis/export/cluster.

Parameters (in addition to dateMask and contentType) are all joined by logical AND:

Table 14–9 Advanced Clustering Export Parameters

Name	Type	Required	Query Logic for Corresponding Column if Parameter is Provided
exportedDt	Date	Required	Greater than or equal to this value
effStartDt	Date	Required	Greater than or equal to this value
effEndDt	Date	Required	Less than or equal to this value
prodHierTypeExtKey	String	Required	Equal to this value []
prodExtKey	String	Required	Equal to this value
locExtKey	String	Required	Equal to this value
locHierTypeExtKey	String	Required	Equal to this value

Customer Segment Export

Get Segments

This service is based on rsestrclst.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/cis/export/segment.

Parameters (in addition to dateMask and contentType) are all joined by logical AND:

Table 14–10 Customer Segment Export Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
exportedDt	Date	Greater than or equal to this value
effStartDt	Date	Greater than or equal to this value
effEndDt	Date	Less than or equal to this value
prodExtKey	String	Equal to this value

Oracle Retail Advanced Science Integration with ORCE (Customer Engagement)

This integration enables the application to send approved customer segments into Oracle Retail Customer Engagement (ORCE). Whenever customer segments are approved for merchandise, location in the application, web service call is made to ORCE in order to save segments, along with its members and relevant attributes.

Credential Store

You can use the Manage Credential Stores screen, shown in [Figure 14–1](#), to maintain ORCE credentials in the application.

Figure 14–1 Manage Credential Stores

Enter the following information into Manage Credential Stores:

Table 14–11 Manage Credential Stores Information

Field	Description
Username	Credential username created by ORCE to enable integration
Password	Credential password (at least one character) created by ORCE to enable integration
Confirm Password	Confirm password
Description	Credential store used for Customer Segment integration with Oracle Customer Engagement

Configuration

The following configuration must be updated for integration. By default, ORCE integration is disabled.

Table 14–12 Configuration

Configuration	Description
CUST_SEG_WS_RELATE_FLG	Flag to identify whether to publish customer segment to customer engagement. (Y/N)
CUST_SEG_WS_RELATE_ORGID	ORGID is the 3-letter ID provided by Oracle Customer Engagement to create and implement an authentication key. (REL)
CUST_SEG_ATTR_WS_RELATE_FLG	Flag to identify whether to publish customer segment attribute to customer engagement. (Y/N)
CUST_SEG_RELATE_HOSTNAME	URL host name to publish customer segment to customer engagement. This should be same as CN name in the certificate.
CUST_SEG_RELATE_PORT	Port number to publish customer segment to customer engagement.
CUST_SEG_REL_SEGSERV_VERSION	Segment service version to use for publishing customer segment to customer engagement. (format vX_0 - v3_0)
CUST_SEG_REL_PROXY_HOSTNAME	Proxy host name to publish customer segment to customer engagement.

Table 14–12 (Cont.) Configuration

Configuration	Description
CUST_SEG_REL_HTTP_PROXY_PORT	HTTP proxy port number to publish customer segment to customer engagement.
CUST_SEG_REL_HTTPS_PROXY_PORT	HTTPS proxy port number to publish customer segment to customer engagement.
CUST_SEG_REL_BATCH_SIZE	Number of customers to add in batches after customer segment to customer engagement is published.

ASO Exports**AIP/Replenishment Results**

This service is based on so_assort_aiprepl_int.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/so/export/soAssortAipreplInt

Parameters (in addition to dateMask and contentType) are all joined by logical OR:

Table 14–13 AIP/Replenishment Results Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
assortmentSetId	String	Equal to this value
exportedDt	Date	Greater than or equal to this value

Output Aggregated Across Approved Runs

This service is based on so_assort_cm_int.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/so/export/soAssortCmInt.

Parameters (in addition to dateMask and contentType) are all joined by logical OR:

Table 14–14 Output Aggregated Across Approved Runs Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
assortmentSetId	String	Equal to this value
exportedDt	Date	Greater than or equal to this value

Assortment Result Details

This service is based on so_assort_int.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/so/export/soAssortInt.

Parameters (in addition to dateMask and contentType) are all joined by logical OR:

Table 14–15 Assortment Results Details Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
assortmentSetId	String	Equal to this value

Table 14–15 (Cont.) Assortment Results Details Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
exportedDt	Date	Greater than or equal to this value

Cross Reference Between Planograms and Finalized Assortments

This service is based on planogram_assortment.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/so/export/planogramAssortment.

Parameters (in addition to dateMask and contentType) all joined by logical OR:

Table 14–16 Cross Reference Between Planograms and Finalized Assortments Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
pogKey	String	Equal to this value
exportDate	Date	Greater than or equal to this value

POG Header

This service is based on planogram.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/so/export/planogram.

Parameters (in addition to dateMask and contentType) all joined by logical OR:

Table 14–17 POG Header Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
pogKey	String	Equal to this value
exportDate	Date	Greater than or equal to this value

POG Equipment Components

This service is based on equipment.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/so/export/equipment.

Parameters (in addition to dateMask and contentType) are all joined by logical OR:

Table 14–18 POG Equipment Components Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
pogKey	String	Equal to this value
exportDate	Date	Greater than or equal to this value

POG/Stores Cross Reference

This service is based on planogram_store.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/so/export/planogramStore.

Parameters (in addition to dateMask and contentType) are all joined by logical OR:

Table 14–19 POG/Stores Cross Reference Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
pogKey	String	Equal to this value
storeKey	String	Equal to this value
exportDate	Date	Greater than or equal to this value

Finalized Assortment Product Hierarchies

This service is based on product_hierarchy.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/so/export/productHierarchy.

Parameters (in addition to dateMask and contentType) are all joined by logical OR:

Table 14–20 Finalized Assortment Product Hierarchies Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
exportDate	Date	Greater than or equal to this value

Finalized Assortment Products

This service is based on product_position.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/so/export/productPosition.

Parameters (in addition to dateMask and contentType) are all joined by logical OR:

Table 14–21 Finalized Assortment Products Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
pogKey	String	Equal to this value
skuKey	String	Equal to this value
exportDate	Date	Greater than or equal to this value

Product Display Style Information

This service is based on sku_details.csv. See *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details.

GET on /rase/resources/so/export/skuDetails.

Parameters (in addition to dateMask and contentType) are all joined by logical OR:

Table 14–22 Product Display Style Information Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
skuKey	String	Equal to this value
skuName	String	Equal to this value
effectiveDate	Date	Greater than or equal to this value

Table 14–22 (Cont.) Product Display Style Information Parameters

Name	Type	Query Logic for Corresponding Column if Parameter is Provided
expiryDate	Date	Greater than or equal to this value
exportDate	Date	Greater than or equal to this value

Outbound Interfaces by Application

Advanced Clustering

Table 14–23 Advanced Clustering Outbound

Logical Group	Outbound File	Notes
AC	rsestrclst.csv	
AC	cis_custseg_attr_exp.csv	
AC	cis_custseg_cat_attr_exp.csv	
AC	cis_custseg_cust_export.csv	
AC	cis_custseg_exp.csv	
AC	cis_store_cluster_attr_exp.csv	
AC	cis_store_cluster_exp.csv	
AC	cis_store_cluster_mem_exp.csv	
AC	cis_store_cluster_prop_exp.csv	

Customer Decision Tree

Table 14–24 Customer Decision Tree Outbound

Logical Group	Outbound File	Notes
CDT	attr.csv.dat	
CDT	drtyattrvaltx.csv.ovr	
CDT	cdt_export.tar.gz	

Demand Transference

Table 14–25 Demand Transference Outbound

Logical Group	Outbound File	Notes
DT	drtyassrtelasv.csv.ovr	
DT	dt_assort_mult.csv	
DT	drtyattrwgtv.csv.ovr	
DT	dt_new_items.csv	
DT	dt_new_item_ros.csv	
DT	ipdmdtfrpcti.csv	

Table 14–25 (Cont.) Demand Transference Outbound

Logical Group	Outbound File	Notes
DT	drtysiminv.csv.ovr	

Assortment and Space Optimization

Table 14–26 Assortment and Space Optimization Outbound

Logical Group	Outbound File	Notes
ASO	equipment.csv	
ASO	planogram_assortment.csv	
ASO	planogram.csv	
ASO	planogram_store.csv	
ASO	product_hierarchy.csv	
ASO	product_position.csv	
ASO	sku_details.csv	
ASO	so_assort_aiprepl_int.txt	
ASO	so_assort_cm_int.txt	
ASO	so_assort_int.txt	

Customer Segmentation

Table 14–27 Customer Segmentation Outbound

Logical Group	Outbound File	Notes
ORCS	cis_custseg_exp.csv	
ORCS	cis_custseg_attr_exp.csv	
ORCS	cis_custseg_cust_export.csv	
ORCS	cis_custseg_cat_attr_exp.csv	

Market Basket Insights

Table 14–28 Market Basket Insights Outbound

Logical Group	Outbound File	Notes
MBI	mba_arm_run_exp.txt	
MBI	mba_arm_result_exp.txt	

Interface Details by Logical Group

Retail Insights

For all of the interfaces listed in this section, refer to the Retail Insights documentation to learn about other data requirements or more details about the interface. This section is present to identify the minimally required content in the interfaces that is necessary for the applications to function. If additional content is provided for other columns, this may result in the need to provide other interfaces in order to maintain relational integrity. By providing more data, it would enable the use of additional RI reporting capabilities other than what is minimally required for the application. It is assumed that all of these interfaces are provided daily.

Note that the order of columns presented in this table is not necessarily the order in which the columns should be provided in the interface.

W_DOMAIN_MEMBER_DS_TL.dat

Table 14–29 W_DOMAIN_MEMBER_DS_TL.dat

Column Name	Notes
DOMAIN_CODE	
DOMAIN_TYPE_CODE	
DOMAIN_MEMBER_CODE	File Sourced
DOMAIN_MEMBER_NAME	Number of weeks in the period (either 4 or 5)
LANGUAGE_CODE	
SRC_LANGUAGE_CODE	
DATASOURCE_NUM_ID	

W_EXCH_RATE_GS.dat

Table 14–30 W_EXCH_RATE_GS.dat

Column Name	Notes
START_DT	
EXCH_RATE	
W_FROM_CURCY_CODE	
W_TO_CURCY_CODE	
DATASOURCE_NUM_ID	
INTEGRATION_ID	W_FROM_CURCY_CODE~W_TO_CURCY_CODE~ YYYYMMDD

W_INT_ORG_DHS.dat

The number of ORG_HIER_##_NUM columns provided depends on the number of levels defined in the organization hierarchy.

Table 14–31 W_INT_ORG_DHS.dat

Column Name	Notes
ORG_TOP_NUM	
ORG_HIER13_NUM	
ORG_HIER12_NUM	

Table 14-31 (Cont.) W_INT_ORG_DHS.dat

Column Name	Notes
ORG_HIER11_NUM	
ORG_HIER10_NUM	
ORG_HIER9_NUM	
LEVEL_NAME	
DATASOURCE_NUM_ID	
INTEGRATION_ID	LEVEL_NAME~ORG_HIER##_NUM

W_INT_ORG_DS.dat**Table 14-32 W_INT_ORG_DS.dat**

Column Name	Notes
ORG_NUM	
DATASOURCE_NUM_ID	
INTEGRATION_ID	ORG_NUM

W_INT_ORG_DS_TL.dat**Table 14-33 W_INT_ORG_DS_TL.dat**

Column Name	Notes
ORG_NAME	
ORG_DESCR	
LANGUAGE_CODE	
SRC_LANGUAGE_CODE	
DATASOURCE_NUM_ID	
INTEGRATION_ID	ORG_NUM

W_MCAL_PERIOD_DS.dat**Table 14-34 W_MCAL_PERIOD_DS.dat**

Column Name	Notes
MCAL_CAL_ID	
MCAL_CAL_NAME	
MCAL_CAL_CLASS	File Sourced
MCAL_PERIOD_TYPE	Number of weeks in the period (either 4 or 5)
MCAL_PERIOD_NAME	
MCAL_PERIOD	
MCAL_PERIOD_ST_DT	
MCAL_PERIOD_END_DT	
MCAL_QTR	

Table 14–34 (Cont.) W_MCAL_PERIOD_DS.dat

Column Name	Notes
MCAL_YEAR	
MCAL_QTR_START_DT	
MCAL_QTR_END_DT	
MCAL_YEAR_START_DT	
MCAL_YEAR_END_DT	
INTEGRATION_ID	MCAL_CAL_ID~YEAR~PERIOD
DATASOURCE_NUM_ID	1
W_INSERT_DT	
W_UPDATE_DT	

W_PARTY_PER_DS.dat**Table 14–35 W_PARTY_PER_DS.dat**

Column Name	Notes
CUSTOMER_NUM	
DATASOURCE_NUM_ID	
INTEGRATION_ID	CUSTOMER_NUM

W_PRODUCT_ATTR_DS.dat

Oracle Retail Advanced Science can use data in any of the PRODUCT_ATTR* columns that are populated with data. Although no specific column is required, it is recommended that many columns be provided.

Table 14–36 W_PRODUCT_ATTR_DS.dat

Column Name	Notes
PROD_NUM	
W_CATEGORY	
DATASOURCE_NUM_ID	
INTEGRATION_ID	PROD_NUM

W_PRODUCT_DS.dat**Table 14–37 W_PRODUCT_DS.dat**

Column Name	Notes
PROD_NUM	
PROD_CAT5	
DATASOURCE_NUM_ID	
INTEGRATION_ID	PROD_NUM

W_PRODUCT_DS_TL.dat**Table 14–38** *W_PRODUCT_DS_TL.dat*

Column Name	Notes
PRODUCT_NAME	
PRODUCT_DESCR	
LANGUAGE_CODE	
SRC_LANGUAGE_CODE	
DATASOURCE_NUM_ID	
INTEGRATION_ID	PROD_NUM

W_PROD_CAT_DHS.dat

The number of LVL#ANC_PROD_CAT_ID can vary, depending on the number of levels in the hierarchy.

Table 14–39 *W_PROD_CAT_DHS.dat*

Column Name	Note
TOP_LVL_PROD_CAT_ID	
LVL8ANC_PROD_CAT_ID	
LVL7ANC_PROD_CAT_ID	
LVL6ANC_PROD_CAT_ID	
LVL5ANC_PROD_CAT_ID	
LVL4ANC_PROD_CAT_ID	
LEVEL_NAME	CMP, DIV, GRP, DEPT, CLS, SBC.
DATASOURCE_NUM_ID	
INTEGRATION_ID	<p>LEVEL_NAME~[LVL#ANC TOP_LVL]_PROD_CAT_ID</p> <p>For Subclass data:</p> <p>SBC~[LVL6ANC_PROD_CAT_ID]~[LVL5ANC_PROD_CAT_ID]~[LVL4ANC_PROD_CAT_ID]</p> <p>Example:</p> <p>SBC~1234~5678~9012</p> <p>For Class Data:</p> <p>CLS~[LVL6ANC_PROD_CAT_ID]~[LVL5ANC_PROD_CAT_ID]</p> <p>Example:</p> <p>CLS~1234~5678</p>

W_RTL_CMG_PRODUCT_MTX_DS.dat**Table 14–40** *W_RTL_CMG_PRODUCT_MTX_DS.dat*

Column Name	Notes
CAT_MGNT_LEVEL	

Table 14–40 (Cont.) W_RTL_CMG_PRODUCT_MTX_DS.dat

Column Name	Notes
CAT_MGNT_NUM	
CAT_MGMT_DESC	
PROD_IT_NUM	
PROD_SC_NUM	
PROD_CL_NUM	
PROD_DP_NUM	
PROD_GP_NUM	
PROD_DV_NUM	
LEVEL_NAME	
DATASOURCE_NUM_ID	
INTEGRATION_ID	

W_RTL_CO_HEAD_DS.dat**Table 14–41 W_RTL_CO_HEAD_DS.dat**

Column Name	Notes
CO_HEAD_ID	
DATASOURCE_NUM_ID	
INTEGRATION_ID	CO_HEAD_ID
ETL_PROC_WID	

W_RTL_CO_LINE_DS.dat**Table 14–42 W_RTL_CO_LINE_DS.dat**

Column Name	Notes
CO_LINE_ID	
CO_HEAD_ID	
DATASOURCE_NUM_ID	
INTEGRATION_ID	CO_HEAD_ID~CO_LINE_ID
ETL_PROC_WID	

W_RTL_CUSTSEG_DS.dat**Table 14–43 W_RTL_CUSTSEG_DS.dat**

Column Name	Notes
CUSTSEG_ID	
CUSTSEG_NAME	
DATASOURCE_NUM_ID	
INTEGRATION_ID	CUSTSEG_ID

W_RTL_CUST_CUSTSEG_DS.dat**Table 14-44** *W_RTL_CUST_CUSTSEG_DS.dat*

Column Name	Notes
CUSTSEG_ID	
CUST_ID	
DATASOURCE_NUM_ID	
INTEGRATION_ID	CUSTSEG_ID~CUST_ID

W_RTL_ITEM_DEL_TMP.dat**Table 14-45** *W_RTL_ITEM_DEL_TMP.dat*

Column Name	Notes
PROD_NUM	
CREATED_ON_DATE	
INTEGRATION_ID	PROD_NUM
DATASOURCE_NUM_ID	

W_RTL_ITEM_GRP1_DS.dat**Table 14-46** *W_RTL_ITEM_GRP1_DS.dat*

Column Name	Notes
PROD_NUM	
PROD_GRP_TYPE	
FLEX_ATTRIB_1_CHAR	
FLEX_ATTRIB_3_CHAR	For ITEM_UDA PROD_GRP_TYPE records, this columns is expected to be populated with data.
DATASOURCE_NUM_ID	
INTEGRATION_ID	ITEM_DIFF: DIFF_TYPE~DIFF_ID~PROD_NUM~ITEM_DIFF ITEM_UDA: PROD_NUM~ITEMUDA~UDA_HEAD~UDA_DTL

W_RTL_IT_LC_DEL_TMP.dat

This interface should be provided when an item is to be considered deleted from use by the system.

Table 14-47 *W_RTL_IT_LC_DEL_TMP.dat*

Column Name	Notes
PROD_NUM	
CREATED_ON_DATE	

Table 14-47 (Cont.) W_RTL_IT_LC_DEL_TMP.dat

Column Name	Notes
INTEGRATION_ID	PROD_NUM
DATASOURCE_NUM_ID	

W_RTL_PROMO_COMP_TYPE_DS.dat**Table 14-48 W_RTL_PROMO_COMP_TYPE_DS.dat**

Column Name	Notes
PROMO_COMPONENT_TYPE_ID	
PROMO_COMPONENT_TYPE	
INTEGRATION_ID	PROMO_COMPONENT_TYPE_ID
DATASOURCE_NUM_ID	

W_RTL_PROMO_DS.dat**Table 14-49 W_RTL_PROMO_DS.dat**

Column Name	Notes
PROMO_EVENT_ID	
PROMO_PARENT_ID	
PROMO_DETAIL_ID	
PROMO_LEVEL	
PROMO_COMPONENT_TYPE	
PROMO_COMPONENT_ID	
DATASOURCE_NUM_ID	
INTEGRATION_ID	PROMO_EVENT_ID~PROMO_PARENT_ID~ PROMO_COMPONENT_ID~PROMO_ DETAIL_ID

W_RTL_RECLASS_DP_GP_TMP.dat**Table 14-50 W_RTL_RECLASS_DP_GP_TMP.dat**

Column Name	Notes
IDNT	
TABLE_NAME	
DATASOURCE_NUM_ID	
INTEGRATION_ID	

W_RTL_RECLASS_IT_SC_CL_TMP.dat**Table 14-51 W_RTL_RECLASS_IT_SC_CL_TMP.dat**

Column Name	Notes
ITEM	

Table 14-51 (Cont.) W_RTL_RECLASS_IT_SC_CL_TMP.dat

Column Name	Notes
NEW_DEPT	
NEW_CLASS	
NEW_SUBCLASS	
OLD_DEPT	
OLD_CLASS	
OLD_SUBCLASS	
DATASOURCE_NUM_ID	
INTEGRATION_ID	

W_RTL_SLS_TRX_IT_LC_DY_FS.dat**Table 14-52 W_RTL_SLS_TRX_IT_LC_DY_FS.dat**

Column Name	Notes
SLS_TRX_ID	
PROD_IT_NUM	
ORG_NUM	
CO_HEAD_ID	-1 If actual order header ids are provided, then data must be provided in W_RTL_CO_HEAD_DS.
CO_LINE_ID	-1. If actual order line ids are provided, then data must be provided in W_RTL_CO_LINE_DS.
DAY_DT	
VOUCHER_ID	-1
RTL_TYPE_CODE	
MIN_NUM	
EMPLOYEE_NUM	-1
SLS_QTY	
SLS_AMT_LCL	
SLS_PROFIT_AMT_LCL	
LIA_QTY	Required if values other than -1 are present in CO_HEAD_ID and CO_LINE_ID.
LIA_AMT_LCL	Required if values other than -1 are present in CO_HEAD_ID and CO_LINE_ID.
DATASOURCE_NUM_ID	
INTEGRATION_ID	SLS_TRX_ID~PROD_IT_NUM~VOUCHER_ID~ DAY_DT
PROMO_COMP_ID	
CUST_REF_TYPE	
CUST_REF_NUMBER	

W_RTL_SLSFC_IT_LC_WK_F.dat**Table 14-53** *W_RTL_SLSFC_IT_LC_WK_F.dat*

Column Name	Notes
PROD_IT_NUM	
ORG_NUM	
SLSFC_FOR_EOW_DT	
SLSFC_ON_DAY_DT	
SLSFC_QTY	This is the sales unit that has been forecasted for the given time frame period.
SLSFC_ID	
DATASOURCE_NUM_ID	
INTEGRATION_ID	

W_RTL_TRADE_AREA_DS.dat**Table 14-54** *W_RTL_TRADE_AREA_DS.dat*

Column Name	Notes
TRADE_AREA_NUM	
TRADE_AREA_TYPE	
DATASOURCE_NUM_ID	
INTEGRATION_ID	TRADE_AREA_NUM

W_RTL_TRADE_ATRA_LOC_MTX_DS.dat**Table 14-55** *W_RTL_TRADE_ATRA_LOC_MTX_DS.dat*

Column Name	Notes
TRADE_AREA_NUM	
ORG_NUM	
DATASOURCE_NUM_ID	
INTEGRATION_ID	TRADE_AREA_NUM~ORG_NUM

RI Market Interfaces

See the *Oracle Retail Insights Cloud Service Suite/Oracle Retail Advanced Science Cloud Services Data Interface* for details about the following RI Market interfaces:

Table 14-56 *RI Market Interfaces*

Table Name	Description
W_RTL_MKTSLS_TA_CH_CNG_WK_FS	This table contains the market sales data at market product/trade area/retail type/consumer group/channel/week level.
W_RTL_MARKET_PRODUCT_DS	This table contains the market product data.
W_RTL_MARKET_PRODUCT_MTX_DS	This table contains the market product matrix data.

Table 14–56 (Cont.) RI Market Interfaces

Table Name	Description
W_RTL_MARKET_PRODUCT_DS_TL	This is the staging table for W_RTL_MARKET_PRODUCT_D_TL.
W_RTL_CONSUMER_GRP_DS	This is the consumer group dimension staging table. Consumers represent a group of unknown or target people with certain characteristics. This is expected to be a full load.
MARKET_PROD_LVL_NAME_DHS	This is the staging table for W_RTL_MARKET_PROD_DH.

Oracle Retail Advanced Science Common

rse_fake_cust_stg.txt

This is the staging table for specifying customers who are considered as fake customers. A fake customer is a customer who purchases too many transactions to be considered a single customer. Examples are generic store cards.

Notes

This interface allows a user to manually define the fake customers. This interface can be used instead of (or in addition to) the automated routine that is provided to automatically detect fake customers. A fake customer is a customer who purchases an unusual number of transactions, and therefore cannot be an actual person shopping for specific needs. The importance of removing fake customers from processing is so that the unusual buying patterns they demonstrate do not interfere with the analysis done to try to understand customer purchase decisions.

The CUSTOMER_NUM should be related to the W_PARTY_PER_D.CUSTOMER_NUM column.

This interface can overwrite the FAKE_CUST_FLG value for existing rows, while the automated process cannot. Therefore, if a customer is detected as a fake customer by the automated detection routine, this loader can be used to signify that the customer is not fake, and then on subsequent executions of the automated process, the customer will not be identified as a fake customer.

Table 14–57 rse_fake_cust_stg.txt

Position	Column Name	Data Type	Description
1	CUSTOMER_NUM	VARCHAR2(80)	The customer ID to be updated.
2	FAKE_CUST_FLG	VARCHAR2(1)	A flag to indicate whether the customer should be identified as fake (Y) or not (N).

rse_like_loc_stg.txt

This is the staging table used to load the like stores for CM Group or Category.

Table 14–58 *rse_like_loc_stg.txt*

Position	Column Name	Data Type	Description
1	LOC_EXT_KEY	VARCHAR2(80)	External ID for the store location.
2	LIKE_LOC_EXT_KEY	VARCHAR2(80)	External ID for the like store location.
3	PROD_HIER_TYPE_NAME	VARCHAR2(255)	The name of the product hierarchy type associated with this store.
4	PROD_EXT_KEY	VARCHAR2(80)	The external key to identify the product hierarchy this like store relates to.
5	WEIGHT	NUMBER(9,4)	Weight of the like store, associated with the store.
6	EFFECTIVE_START_DT	DATE	The date of the start of the effective period. (Day)
7	EFFECTIVE_END_DT	DATE	The date of the end of the effective period. (Day)
8	ACTIVE	VARCHAR2(1)	This is a Y/N flag to indicate whether this like store mapping is valid or not.
9	NEW_FLG	VARCHAR2(1)	This is a Y/N flag to indicate whether this store is new or existing store [poor history].

rse_prod_attr_grp_value_stg.txt

This is the staging table used to load the associations of CM Groups to product attributes.

Notes

This table defines the set of attributes and attribute values for those attributes. The only NULLABLE columns for this interface is the PROD_ATTR_GRP_DESCR, PROD_ATTR_VALUE_DESCR columns.

The data should be joinable to RSE_PROD_SRC_XREF via the PROD_EXT_KEY, where the LEAF_FLG = Y. The PROD_HIER_TYPE_NAME should be joinable to the RSE_HIER_TYPE table on the Name column.

The values in the PROD_ATTR_GRP_EXT_KEY must be uniquely assigned to a PROD_EXT_KEY.

Table 14–59 *rse_prod_attr_grp_value_stg.txt*

Position	Column Name	Data Type	Description
1	PROD_HIER_TYPE_NAME	VARCHAR2(255)	The name of the product hierarchy type associated with this Attribute Group value.

Table 14–59 (Cont.) rse_prod_attr_grp_value_stg.txt

Position	Column Name	Data Type	Description
2	PROD_EXT_KEY	VARCHAR2(80)	The external key to identify the product hierarchy this product attribute group value relates to.
3	ATTR_SHORT_DB_NAME	VARCHAR2(30)	The short name for the attribute that this product attribute group is related to.
4	PROD_ATTR_GRP_EXT_KEY	VARCHAR2(80)	The external key to uniquely identify the product attribute group.
5	PROD_ATTR_GRP_NAME	VARCHAR2(255)	The name for the product attribute group.
6	PROD_ATTR_GRP_DESCR	VARCHAR2(255)	The description for the product attribute group.
7	PROD_ATTR_VALUE_KEY	VARCHAR2(255)	The external key to uniquely identify the product attribute group value.
8	PROD_ATTR_VALUE_NAME	VARCHAR2(255)	The name for the product attribute group value.
9	PROD_ATTR_VALUE_DESCR	VARCHAR2(255)	The description for the product attribute group value
10	FUNC_ATTR_FLG	VARCHAR2(1)	This is a Y/N flag to indicate whether this attribute is considered to be an attribute associated with a specific function or role (Y) or not (N).

rse_prod_attr_value_xref_stg.txt

This table contains a cross reference of product attribute values to the CM Group Attribute Value Groups.

Notes

This table must be joinable to the RSE_PROD_ATTR_GRP_VALUE_STG table. This interface must be unique across all columns of this table. The PROD_ATTR_VALUE_KEY must be joinable to data that was provided by the related RSE_PROD_ATTR_GRP_VALUE_STG table. The MIN_ATTR_NUM_VALUE/MAX_ATTR_NUM_VALUE columns must be provided as a set, when one is provided. The MIN_ATTR_DATE_VALUE/MAX_ATTR_DATE_VALUE columns must be provided as a set, when one is provided. Every PROD_ATTR_VALUE_KEY that was provided by the RSE_PROD_ATTR_GRP_VALUE_STG should have some rows provided in this interface, so that attribute values can be found and associated with this attribute value. The ATTE_VALUE_EXT_CODE is expected to be joined with RI's W_RTL_ITEM_GRP1_D. Only one set of value columns should be provided per row. (that is, MIN_ATTR_NUM_VALUE & MAX_ATTR_NUM_VALUE but not at the same time as specifying a value for ATTR_STRING_VALUE.

Table 14–60 *rse_prod_attr_value_xref_stg.txt*

Position	Column Name	Data Type	Description
1	PROD_ATTR_VALUE_KEY	VARCHAR2(255)	External key to identify the product attribute group value this xref is for.
2	MIN_ATTR_NUM_VALUE	NUMBER(22,5)	Minimum number value for this xref. Inclusive of this value.
3	MAX_ATTR_NUM_VALUE	NUMBER(22,5)	Maximum number value for this xref. This value is not inclusive in this xref.
4	ATTR_STRING_VALUE	VARCHAR2(255)	An attribute string value to associate with this attribute group value.
5	MIN_ATTR_DATE_VALUE	DATE	Minimum date value to associate with this attribute group value. This value is inclusive.
6	MAX_ATTR_DATE_VALUE	DATE	Maximum attribute value to include for this attribute group value. This value is not inclusive in this range.
7	ATTR_VALUE_EXT_CODE	VARCHAR2(255)	An external attribute code to associate with this attribute group value.

rse_sls_pr_lc_cs_wk_stg.txt

This is the staging table to load aggregate sales data for a product, location, customer segment, and week.

Notes

The data should be unique for the WK_END_DT, PROD_EXT_KEY, LOC_EXT_KEY, and CUSTSEG_EXT_KEY columns. The WK_END_DT should be related to historical date that is within the fiscal calendar hierarchy. This interface expects the data to be provided in a weekly aggregate form, according to the definition of the week in the fiscal calendar hierarchy. The PROD_EXT_KEY should be related to a LEAF_NODE_FLG=Y row in the RSE_PROD_SRC_XREF for the primary product hierarchy. The LOC_EXT_KEY should be related to a LEAF_NODE_FLG=Y row in the RSE_LOC_SRC_XREF table for the primary location hierarchy. The CUSTSEG_EXT_KEY should be related to a LEAF_NODE_FLG=Y row in the RSE_CUSTSEG_SRC_XREF table. The SLS_QTY, SLS_AMT, and PROFIT_AMT columns should represent the sales of the product during the week period, which are not attributed to a promotion. The SLS_PR_QTY, SLS_PR_AMT, and SLS_PR_PROFIT_AMT columns should represent the sales of the product during the week period, which are attributed to a promotion.

Table 14–61 *rse_sls_pr_lc_cs_wk_stg.txt*

Position	Column Name	Data Type	Description
1	WK_END_DT	DATE	The date of the end of the fiscal week

Table 14–61 (Cont.) rse_sls_pr_lc_cs_wk_stg.txt

Position	Column Name	Data Type	Description
2	PROD_EXT_KEY	VARCHAR2(80)	External ID for the Product
3	LOC_EXT_KEY	VARCHAR2(80)	External ID for the store location.
4	CUSTSEG_EXT_KEY	VARCHAR2(80)	External ID for the customer segment.
5	SLS_QTY	NUMBER(38,20)	Quantity of units sold for this entity, while not on promotion.
6	SLS_AMT	NUMBER(38,20)	Global sales amount sold for this entity while not on promotion.
7	PROFIT_AMT	NUMBER(38,20)	Amount of profit for this entity while not on promotion.
8	SLS_PR_QTY	NUMBER(38,20)	The number of units sold that was associated with a promotion.
9	SLS_PR_AMT	NUMBER(38,20)	The global sales currency amount that was associated with a promotion.
10	SLS_PR_PROFIT_AMT	NUMBER(38,20)	Amount of global current profit amount for this entity that was associated with a promotion.

rse_sls_pr_lc_wk_stg.txt

This is the staging table to load aggregate sales data for a product, location and week.

Notes

The data should be unique for the WK_END_DT, PROD_EXT_KEY, LOC_EXT_KEY columns. The WK_END_DT should be related to historical date that is within the fiscal calendar hierarchy. This interface expects the data to be provided in a weekly aggregate form, according to the definition of the week in the fiscal calendar hierarchy. The PROD_EXT_KEY should be related to a LEAF_NODE_FLG=Y row in the RSE_PROD_SRC_XREF for the primary product hierarchy. The LOC_EXT_KEY should be related to a LEAF_NODE_FLG=Y row in the RSE_LOC_SRC_XREF table for the primary location hierarchy. The SLS_QTY, SLS_AMT, and PROFIT_AMT columns should represent the sales of the product during the week period, which are not attributed to a promotion. The SLS_PR_QTY, SLS_PR_AMT, and SLS_PR_PROFIT_AMT columns should represent the sales of the product during the week period, which are attributed to a promotion.

Table 14–62 rse_sls_pr_lc_wk_stg.txt

Position	Column Name	Data Type	Description
1	WK_END_DT	DATE	The date of the end of the fiscal week
2	PROD_EXT_KEY	VARCHAR2(80)	External ID for the Product

Table 14–62 (Cont.) rse_sls_pr_lc_wk_stg.txt

Position	Column Name	Data Type	Description
3	LOC_EXT_KEY	VARCHAR2(80)	External ID for the store location.
4	SLS_QTY	NUMBER(38,20)	Quantity of units sold for this entity, while not on promotion.
5	SLS_AMT	NUMBER(38,20)	Global sales amount sold for this entity while not on promotion.
6	PROFIT_AMT	NUMBER(38,20)	Amount of profit for this entity while not on promotion.
7	SLS_PR_QTY	NUMBER(38,20)	The number of units sold that was associated with a promotion.
8	SLS_PR_AMT	NUMBER(38,20)	The global sales currency amount that was associated with a promotion.
9	SLS_PR_PROFIT_AMT	NUMBER(38,20)	Amount of global current profit amount for this entity that was associated with a promotion.

Customer Decision Tree

attr.csv.dat (cdt_attribute_exp_vw)

This view provides the complete set of category specific attributes, and their attribute values for an export to the CMPO application.

Notes

This view provides the complete set of category specific attributes, and their attribute values for an export to the CMPO application. The data in this table should be unique by the ATTRIBUTE_VALUE_ID column. The CDT_ATTR_VAL_PROD_XREF_EXT_VW provides additional details that are related to this export's data. These two exports are joinable by the ATTRIBUTE_VALUE_ID column. The source data for this export is the RSE_PROD_ATTR_GRP_VALUE and RSE_PROD_ATTR_GRP tables.

Table 14–63 attr.csv.dat (cdt_attribute_exp_vw)

Position	Column Name	Data Type	Description
1	ATTRIBUTE_VALUE_ID	VARCHAR2(80)	The external identifier for the attribute value.
2	ATTRIBUTE_VALUE_NAME	VARCHAR2(255)	A descriptive name for the attribute value.
3	ATTRIBUTE_EXTERNAL_ID	VARCHAR2(80)	The external identifier for the category specific attribute.

Table 14–63 (Cont.) attr.csv.dat (cdt_attribute_exp_vw)

Position	Column Name	Data Type	Description
4	ATTRIBUTE_NAME	VARCHAR2(255)	A descriptive name for the category specific attribute.

drtyattrvaltx.csv.ovr (cdt_attr_val_prod_xref_exp_vw)

This view provides the data to provide to CMPO application for an export of Products and their Product Attribute Values.

Notes

This view provides a complete set of exportable attribute values and the products that have these attribute values. This view can be used to provide data to the CMPO application, so that they be aware of the product/attribute value assignments, which CDT uses. The MERCHANDISE_ID is the external key that relates to the RSE_PROD_SRC_XREF table's PROD_EXT_KEY. This value should be limited to only leaf node records in that table, and should be related to the hierarchy type as configured in RSE_CONFIG's CDT_PROD_HIER_TYPE configuration. The ATTRIBUTE_EXTERNAL_ID is the external identifier as obtained from the RSE_PROD_ATTR_GRP's PROD_ATTR_GRP_EXT_KEY column. This should be unique across a given CM Group. This value should also be provided in the CDT_ATTRIBUTE_EXP_VW as the ATTRIBUTE_EXTERNAL_ID column. The ATTRIBUTE_VALUE_ID is the unique value for a given attribute value. This value should be associated with the RSE_PROD_ATTR_GRP_VALUE's PROD_ATTR_EXT_KEY column.

Table 14–64 drtyattrvaltx.csv.ovr (cdt_attr_val_prod_xref_exp_vw)

Position	Column Name	Data Type	Description
1	MERCHANDISE_ID	VARCHAR2(80)	The external identifier for a product/SKU.
2	ATTRIBUTE_EXTERNAL_ID	VARCHAR2(80)	The external identifier for the Category specific attribute.
3	ATTRIBUTE_VALUE_ID	VARCHAR2(80)	The external identifier for the product attribute value.

cdt_export.tar.gz

This tar file contains multiple XML files that represent a CDT for a given location and customer segment. The format of the XML is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
elementFormDefault="qualified" targetNamespace="http://oracle/rgbu/cdt/2.0"
xmlns:cdt="http://oracle/rgbu/cdt/2.0">
  <xs:element name="CDT">
    <xs:complexType>
      <xs:sequence>
        <xs:element ref="cdt:attribute"/>
      </xs:sequence>
      <xs:attribute name="segment" use="required"/>
      <xs:attribute name="category" use="required" type="xs:NCName"/>
      <xs:attribute name="tradingarea" use="required" type="xs:NCName"/>
    </xs:complexType>
  </xs:element>
```

```

<xs:element name="attribute">
  <xs:complexType>
    <xs:sequence>
      <xs:element minOccurs="0" maxOccurs="unbounded" ref="cdt:attribute"/>
    </xs:sequence>
    <xs:attribute name="name" use="required" type="xs:NCName"/>
    <xs:attribute name="value" type="xs:NCName"/>
  </xs:complexType>
</xs:element>
</xs:schema>

```

Demand Transference

drtyassrtelasv.csv.ovr (dt_ae_exp_vw)

This view provides all the exportable data elements needed to provide the Assortment Elasticity metric that the DT application has calculated. This view only provides results for Active results.

Table 14–65 drtyassrtelasv.csv.ovr (dt_ae_exp_vw)

Position	Column Name	Data Type	Description
1	PROD_EXT_KEY	VARCHAR2(80)	The external identifier for the category this data belongs to.
2	CUSTSEG_EXT_KEY	VARCHAR2(80)	The external identifier for the customer segment associated with this data.
3	LOC_EXT_KEY	VARCHAR2(80)	The external identifier for the location.
4	AE	NUMBER(22,7)	The assortment elasticity that DT has calculated.
5	EFFECTIVE_DT_FROM	DATE	The date when this data was activated.
6	EFFECTIVE_DT_TO	DATE	Not used.

dt_assort_mult.csv (dt_assort_mult_exp_vw)

This view provides an exportable list of assortment multipliers to RDF so that the impact of assortment changes can influence RDF results.

Table 14–66 dt_assort_mult.csv (dt_assort_mult_exp_vw)

Position	Column Name	Data Type	Description
1	EFF_START_DT	DATE	The date that this assortment multiplier is effective for use.
2	PROD_EXT_KEY	VARCHAR2(80)	The external ID for the product.
3	LOC_EXT_KEY	VARCHAR2(80)	The external ID for the store location.

Table 14–66 (Cont.) dt_assort_mult.csv (dt_assort_mult_exp_vw)

Position	Column Name	Data Type	Description
	ASSORT_MULT	NUMBER(38,20)	The assortment multiplier associated with the product and location, as a result of changes to the assortment at the location.

ipdmdtfrpcti.csv (dt_assort_trans_exp_vw)

This view provides an exportable list of assortment multipliers to AIP so that the impact of assortment changes can influence AIP results.

Table 14–67 ipdmdtfrpcti.csv (dt_assort_trans_exp_vw)

Position	Column Name	Data Type	Description
1	LOC_EXT_KEY	VARCHAR2(80)	The external ID for the store location.
2	PROD_EXT_KEY	VARCHAR2(80)	The external ID for the product.
3	REPL_PROD_EXT_KEY	VARCHAR2(80)	The external ID for the replacement product.
4	TRANSFER_PCT	NUMBER	The percentage of demand that could transfer from the first product to the second (replacement) product.

drtyattrwgtv.csv.ovr (dt_attr_wgt_exp_vw)

This view provides the export data that is to be provided to the CMPO system for attribute weights used by the DT application.

Table 14–68 drtyattrwgtv.csv.ovr (dt_attr_wgt_exp_vw)

Position	Column Name	Data Type	Description
1	PROD_EXT_KEY	VARCHAR2(80)	The external ID for the category
2	CUSTSEG_EXT_KEY	VARCHAR2(80)	The external ID for the customer segment
3	LOC_EXT_KEY	VARCHAR2(80)	The external ID for the location hierarchy/trade area
4	ATTRIBUTE_EXTERNAL_ID	VARCHAR2(80)	The external ID for the attribute
5	ATTR_WGT	NUMBER(22,7)	The decimal weight that this attribute represents within the category/customer segment/trade area

Table 14–68 (Cont.) drtyattrwgtv.csv.ovr (dt_attr_wgt_exp_vw)

Position	Column Name	Data Type	Description
6	FUNC_ATTR_IND	NUMBER	An indicator to express with the attribute is a functional attribute (1) or not (0). A functional attribute is one that fits a specific purpose and cannot be substituted by other products with other values for this attribute.

dt_new_items.csv (dt_new_items_exp_vw)

This view provides a list of items and locations, for which the item is newly added to the assortment at the location in an exportable view for use by RDF.

Table 14–69 dt_new_items.csv (dt_new_items_exp_vw)

Position	Column Name	Data Type	Description
1	PROD_EXT_KEY	VARCHAR2(80)	The external ID for the product.
2	LOC_EXT_KEY	VARCHAR2(80)	The external ID for the store location.
3	EFF_START_DT	DATE	The date that this product is considered to be added to the assortment.

dt_new_item_ros.csv (dt_new_item_ros_exp_vw)

This view provides an exportable view of new items and their forecasted rate of sale to RDF.

Table 14–70 dt_new_item_ros.csv (dt_new_item_ros_exp_vw)

Position	Column Name	Data Type	Description
1	PROD_EXT_KEY	VARCHAR2(80)	The external ID for the product.
2	LOC_EXT_KEY	VARCHAR2(80)	The external ID for the store location.
3	FCST_ROS	NUMBER(38,20)	The forecasted rate of sale for this product at this location.

drtysiminv.csv.ovr (dt_sim_exp_vw)

This view provides an export of product similarities calculated within the DT application for export to CMPO.

Table 14–71 drtysiminv.csv.ovr (dt_sim_exp_vw)

Position	Column Name	Data Type	Description
1	PROD_EXT_KEY1	VARCHAR2(80)	The external identifier for one half of the product pair.

Table 14–71 (Cont.) drtysiminv.csv.ovr (dt_sim_exp_vw)

Position	Column Name	Data Type	Description
2	CUSTSEG_EXT_KEY	VARCHAR2(80)	The external identifier for the customer segment this data relates to.
3	LOC_EXT_KEY	VARCHAR2(80)	The external identifier for the store location this data relates to.
4	PROD_EXT_KEY2	VARCHAR2(80)	The external identifier for the other half of the product pair.
5	PROD_SIM	NUMBER(22,7)	The measurement of how similar the two products are to each other. The values range between values of 0 (completely dissimilar) to values of 1 (completely similar).
6	EFFECTIVE_DT_FROM	DATE	The date this similarity value became effective.
7	EFFECTIVE_DT_TO	DATE	The date this similarity value is effective until.

dt_loc_wk_excl_stg.txt

This is a staging table that loads a list of locations and dates that should be excluded from Demand Transference processing. This can be useful to exclude abnormal or corrupted data points.

Table 14–72 dt_loc_wk_excl_stg.txt

Position	Column Name	Data Type	Description
1	LOC_EXT_KEY	VARCHAR2(80)	External ID for the location.
2	WK_END_DT	DATE	The week end date of a week to be excluded.
3	UPDT_CODE	VARCHAR2(1)	A code to indicate how to update the target table. U = Update/Create, D = Delete existing record.

dt_mdl_prod_exp_stg.txt

This table is the staging table that provides a list of products that are eligible for processing Model Apply in order to receive product to product demand transferences.

The values loaded here should be of the same product hierarchy that DT is associated to work with.

Table 14–73 dt_mdl_prod_exp_stg.txt

Position	Column Name	Data Type	Description
1	PROD_EXT_KEY	VARCHAR2(80)	The external product identifier for the hierarchy to be included in later model apply processing.

dt_prod_loc_excl_stg.txt

The staging table for products and locations that should be excluded from DT processing as if they were out of the assortment.

Table 14–74 dt_prod_loc_excl_stg.txt

Position	Column Name	Data Type	Description
1	PROD_EXT_KEY	VARCHAR2(80)	External product key
2	LOC_EXT_KEY	VARCHAR2(80)	External ID for the location.
3	WK_FROM_DT	DATE	Week date to start the exclusion at.
4	WK_TO_DT	DATE	Week date to end the exclusion at.

Assortment and Space Optimization – Assortment Files

so_assortment_finalized_stg.txt

The interface table to accept assortment finalization general details. This data is used to summarize the multiples assortments in the set.

Notes

This staging table is used to receive assortment data from external sources. Each assortment provided must have a unique assortment_id. This data is mandatory.

Table 14–75 so_assortment_finalized_stg.txt

Position	Column Name	Data Type	Description
1	ASSORTMENT_SET_ID	VARCHAR2(80)	Unique assortment set ID. This value is used to group together multiple assortments (user requests).
2	PRODUCT_CATEGORY_KEY	VARCHAR2(80)	This value has to match a node in merchandise hierarchy. This is the external ID that is known and shared across applications.

Table 14–75 (Cont.) so_assortment_finalized_stg.txt

Position	Column Name	Data Type	Description
3	ASSORT_LABEL	VARCHAR2(80)	This is a user entered field with default value. This value is presented within the UI as the familiar label/name recognizable by the user. It can be NULL.
4	TRADE_AREA_LABEL	VARCHAR2(80)	CMPO trading area used to extract the assortment.
5	ASSORT_ROLE	VARCHAR2(50)	This is a CMPO data element that should be passed to SO. SO shows this value within BI modules. This is relevant for the user when they pick the optimization objective function.
6	ASSORT_TACTIC	VARCHAR2(100)	This is a CMPO data element that should be passed to SO. SO shows this value within BI modules. This is relevant for the user when they pick the optimization objective function.
7	ASSORT_GOAL	VARCHAR2(50)	This is a CMPO data element that should be passed to SO. SO shows this value within BI modules. This is relevant for the user when they pick the optimization objective function.

so_assortment_stg.txt

The interface table to accept assortment header/general information.

Notes

This staging table is used to receive assortment data from external sources. Each assortment provided must have a unique assortment_id. Mandatory data.

Table 14–76 so_assortment_stg.txt

Position	Column Name	Data Type	Description
1	ASSORTMENT_SET_ID	VARCHAR2(80)	Unique assortment set ID. This value is used to group together multiple assortments (user requests).
2	ID	VARCHAR2(80)	System generated Primary Key. internal assortment identifier.

Table 14–76 (Cont.) so_assortment_stg.txt

Position	Column Name	Data Type	Description
3	PRODUCT_CATEGORY_KEY	VARCHAR2(80)	This value has to match a node in merchandise hierarchy. This is the external ID that is known and shared across applications.
4	ASSORT_LABEL	VARCHAR2(80)	This is a user entered field with default value. This value is presented within the UI as the familiar label/name recognizable by the user. It can be NULL.
5	TRADE_AREA_LABEL	VARCHAR2(80)	CMPO trading area used to extract the assortment.
6	REQUEST_TYPE	NUMBER(2)	This field can accept two values. 1= Optimization Request, 2= Finalized Assortment Reporting.
7	ASSORT_LOC_TYPE	NUMBER(2)	This field can accept two values. 1= Cluster Assortment, 2 = Store Assortment. This field indicates the level at that the assortment is delivered.
8	ASSORT_ROLE	VARCHAR2(50)	This is a CMPO data element that should be passed to SO. SO shows this value within BI modules. This is relevant for the user when they pick the optimization objective function.
9	ASSORT_TACTIC	VARCHAR2(100)	This is a CMPO data element that should be passed to SO. SO shows this value within BI modules. This is relevant for the user when they pick the optimization objective function.
10	ASSORT_GOAL	VARCHAR2(50)	This is a CMPO data element that should be passed to SO. SO shows this value within BI modules. This is relevant for the user when they pick the optimization objective function.

so_assort_cluster_member_stg.txt

This staging table receives the relationship of stores assigned to a specific cluster for the given assortment.

Notes

This is the store list for the stores delivered within the assortment interface, grouped within clusters. SO is expecting to always receive stores grouped within clusters. In a case where stores need to be sent individually, a cluster should be created for than single store. Product list will be linked directly to a store whenever assortment type = 2 (Store). Start and end date will only be included for this table for assortments delivered at the store level. Mandatory data.

Table 14–77 *so_assort_cluster_member_stg.txt*

Position	Column Name	Data Type	Description
1	ASSORTMENT_ID	VARCHAR2(80)	ID that identified the assortment. It must match an assortment key within the assortment file.
2	LOCATION_KEY	VARCHAR2(80)	This is the external store ID, known and shared across applications.
3	CLUSTER_KEY	VARCHAR2(80)	Internal CMPO cluster key. This key must match one of the cluster keys provided within the assortment cluster file.
4	START_DT	DATE	Start date range to be used for the store to retrieve forecast data. The format should be YYYY-MM-DD.
5	END_DT	DATE	End date range to be used for the store to retrieve forecast data. The format should be YYYY-MM-DD.

so_assort_cluster_stg.txt

This is the staging table used to receive assortment placeholder products included within the assortment.

Notes

This is the Cluster List for the clusters delivered within the assortment interface. SO is expecting to always receive stores grouped within clusters. In a case where stores need to be sent individually, a cluster should be created for than single store. Product list will be linked directly to a cluster whenever assortment type = 1 (Cluster). Start and end date will only be included for assortments delivered at the cluster level. Mandatory data.

Table 14–78 *so_assort_cluster_stg.txt*

Position	Column Name	Data Type	Description
1	ASSORTMENT_ID	VARCHAR2(80)	ID that identified the assortment. It must match an assortment key within the assortment file.
2	CLUSTER_KEY	VARCHAR2(80)	CMPO cluster key, a value that can be returned from SO to CMPO to uniquely identify the cluster_id.

Table 14–78 (Cont.) so_assort_cluster_stg.txt

Position	Column Name	Data Type	Description
3	CLUSTER_NAME	VARCHAR2(80)	Name associated to the cluster, end user should recognize this name as the cluster name seen or entered within CMPO.
4	START_DT	DATE	Start date range to be used for the stores within the cluster to retrieve forecast data. The format should be YYYY-MM-DD.
5	END_DT	DATE	End date range to be used for the stores within the cluster to retrieve forecast data. The format should be YYYY-MM-DD.

so_assort_phprod_attr_stg.txt

This is the staging table used to receive attribute data for assortment placeholder products.

Notes

This table includes only placeholder product attributes. Attribute names must match existing attributes already available within SO and shared with the other products. Optional data.

Table 14–79 so_assort_phprod_attr_stg.txt

Position	Column Name	Data Type	Description
1	ASSORTMENT_ID	VARCHAR2(80)	ID that identified the assortment. It must match an assortment key within the assortment file.
2	PLACEHOLDER_PRODUCT_KEY	VARCHAR2(80)	CMPO product key for placeholder product specific to the assortment, must be different from other formalized products
3	ATTR_NAME	VARCHAR2(50)	Name of the product attribute that is being passed. Must match a known product attribute.
4	ATTR_VALUE	VARCHAR2(50)	Specific value that should be used for the placeholder product/attribute combination.

so_assort_phprod_finalized_stg.txt

This staging table is used to receive finalized assortment placeholder products included within the assortment. This data is used to transform the placeholder name and ID.

Notes

Optional data; if not, placeholder products are included. List of placeholder products included in the assortment. Each placeholder item must be paired with an existing product. SO uses the next pieces of data from the existing product (like item).

- MSM-type data like sizes, other product merchandising info
- Merchandise hierarchy info (where this product sits in the hierarchy)
- Product attributes (like what would be used for DT calls)
- SO-only data (sku/store repl parameters used in SO, other SO inputs)

The Product Key for placeholder items must always be different than the one for any known product.

Table 14–80 *so_assort_phprod_finalized_stg.txt*

Position	Column Name	Data Type	Description
1	ASSORTMENT_SET_ID	VARCHAR2(80)	Unique assortment set ID. This value is used to group together multiple assortments (user requests).
2	PLACEHOLDER_PRODUCT_KEY	VARCHAR2(80)	CMPO product key for placeholder product specific to the assortment, must be different from other formalized products.
3	FINALIZED_PRODUCT_NAME	VARCHAR2(80)	Tag that describes the finalized placeholder item.
4	FINALIZED_PRODUCT_KEY	VARCHAR2(80)	This field must match a Product Key Definition in RSE Core. This is the external ID that is known and shared across applications. The like product key must be one of the known products also included within the assortment.

so_assort_phprod_like_prod_stg.txt

The staging table is used to receive assortment placeholder products included within the assortment.

Notes

Optional data; if not, placeholder products are included. List of placeholder products included in the assortment. Each placeholder item must be paired with an existing product. SO uses the next pieces of data from the existing product (like item).

- MSM-type data like sizes, other product merchandising info
- Merchandise hierarchy info (where this product sits in the hierarchy)
- Product attributes (like what would be used for DT calls)
- SO-only data (sku/store repl parameters used in SO, other SO inputs)

The Product Key for placeholder items must always be different than the one for any known product.

Table 14–81 *so_assort_phprod_like_prod_stg.txt*

Position	Column Name	Data Type	Description
1	ASSORTMENT_ID	VARCHAR2(80)	ID that identified the assortment. It must match an assortment key within the assortment file.
2	PLACEHOLDER_PRODUCT_KEY	VARCHAR2(80)	CMPO product key for placeholder product specific to the assortment, must be different from other formalized products.
3	PLACEHOLDER_PRODUCT_NAME	VARCHAR2(80)	Tag that describes the placeholder item, it is used by the UI when looking at product level data.
4	LIKE_PRODUCT_KEY	VARCHAR2(80)	This field must match a Product Key Definition in RSE Core. This is the external ID that is known and shared across applications. The like product key must be one of the known products also included within the assortment.

so_assort_product_strcltr_stg.txt

Staging table that receives the list of products within the assortment. This data is delivered at the Product/ (store or cluster) level.

Notes

Placeholder products must also be included within this table. An assortment can be delivered either at the store or cluster level, but not both at the same time. Product key for placeholder products must always be different than the one used for any known product. Mandatory data.

Table 14–82 *so_assort_product_strcltr_stg.txt*

Position	Column Name	Data Type	Description
1	ASSORTMENT_ID	VARCHAR2(80)	ID that identified the assortment. It must match an assortment key within the assortment file.

Table 14–82 (Cont.) so_assort_product_strcltr_stg.txt

Position	Column Name	Data Type	Description
2	CLUSTER_STORE_KEY	VARCHAR2(80)	This field can be either an assortment cluster key or a location key. The actual value is determined by the assortment type (If Assortment_type = 1 (ClusterAssortment key) and If Assortment_type = 2 (Store Assortment key). The external store ID must be the one known and shared across applications.
3	PRODUCT_KEY	VARCHAR2(80)	This is the external product ID that is known and shared across applications. In case of placeholder products this field will contain CMPO placeholder product key that must be different than any known product.
4	IPI_VALUE	NUMBER(18,4)	This value can be NULL if not available.
5	PRIORITY	NUMBER(2)	This field can take 4 different values, 1 = mandatory, 2 = core, 3 = optional, and -1=dropped. Records with -1 may be filtered out.

so_assort_proloc_fcst_stg.txt

This is the staging table used to receive assortment forecast data.

Notes

This table receives forecast data for all the products within the assortment, including placeholder products. The forecast must cover the range of dates specified for the cluster or stores. Mandatory data.

Table 14–83 so_assort_proloc_fcst_stg.txt

Position	Column Name	Data Type	Description
1	ASSORTMENT_ID	VARCHAR2(80)	ID that identified the assortment. It must match an assortment key within the assortment file.
2	PRODUCT_KEY	VARCHAR2(80)	This is the external ID that is known and shared across applications. In case of placeholder products this field contains CMPO placeholder product key.

Table 14–83 (Cont.) so_assort_proloc_fcst_stg.txt

Position	Column Name	Data Type	Description
3	LOCATION_KEY	VARCHAR2(80)	This is the external store ID, known and shared across applications.
4	WEEKLY_PERIOD	DATE	Week start date for which the forecast is provided.
5	DEMAND	NUMBER(18,4)	Forecast demand for the week.
6	ERROR_TERM	NUMBER(18,4)	Not currently in use. Default to NULL.

so_assort_proloc_pricecost_stg.txt

This staging table is used to receive price and cost data for assortment products.

Notes

The data on this table must be delivered at the product/location level. This table must include the corresponding price and cost for placeholder products (if any is included within the assortment). Mandatory data.

Table 14–84 so_assort_proloc_pricecost_stg.txt

Position	Column Name	Data Type	Description
1	ASSORTMENT_ID	VARCHAR2(80)	ID that identified the assortment. It must match an assortment key within the assortment file.
2	PRODUCT_KEY	VARCHAR2(80)	This is the external ID that is known and shared across applications. In case of placeholder products this field contains CMPO placeholder product key.
3	LOCATION_KEY	VARCHAR2(80)	This is the external store ID, known and shared across applications.
4	PRICE	NUMBER(18,4)	Product price, single currency unit determined at configuration level. No multi-currency is allowed.
5	COST	NUMBER(18,4)	Product cost, single currency unit determined at configuration level. No multi-currency is allowed.

Assortment and Space Optimization – POG Files**so_bay_fixture_shelf_stg.txt**

This table receives the information that describe the Shelf layout in a Fixture. (It is used for Shelf Fixture only.)

Notes

This table has the information about the specific shelves that are included within a given shelf fixture. This table is only populated for planograms that include shelf fixtures; data is not available within this table for Pegboard and freezer chest fixtures.

Table 14–85 *so_bay_fixture_shelf_stg.txt*

Position	Column Name	Data Type	Description
1	BAY_KEY	VARCHAR2(80)	Bay external ID. Bay is a direct dependent of the planogram.
2	FIXTURE_KEY	VARCHAR2(80)	Fixture external ID. Fixture is a direct dependent of the Bay.
3	SHELF_KEY	VARCHAR2(80)	Shelf external ID. Shelf is a direct dependent of shelf fixture.
4	POS_X	NUMBER(18,4)	Position of the shelf on the X axis. Origin point: bottom, left, back (within the fixture).
5	POS_Y	NUMBER(18,4)	Position of the shelf on the Y axis. Origin point: bottom, left, back (within the fixture).
6	POS_Z	NUMBER(18,4)	Position of the shelf on the Z axis. Origin point: bottom, left, back (within the fixture).

so_bay_fixture_stg.txt

This table receives the fixture layout within a Bay. Fixture can be Shelf, Pegboard or Freezer.

Notes

This table has the information about the different fixtures that are assigned to every planogram bay. There can be multiple fixtures within a bay; each fixture within a bay keeps the same layout from left to right.

Table 14–86 *so_bay_fixture_stg.txt*

Position	Column Name	Data Type	Description
1	BAY_KEY	VARCHAR2(80)	Bay external ID. Bay is a direct dependent of the planogram.
2	FIXTURE_KEY	VARCHAR2(80)	Fixture external ID. Fixture is a direct dependent of the Bay.
3	POS_X	NUMBER(18,4)	Position of the fixture on the X axis. Origin point: bottom, left, back (within the bay)

Table 14–86 (Cont.) so_bay_fixture_stg.txt

Position	Column Name	Data Type	Description
4	POS_Y	NUMBER(18,4)	Position of the fixture on the Y axis. Origin point: bottom, left, back (within the bay)
5	POS_Z	NUMBER(18,4)	Position of the fixture on the Z axis. Origin point: bottom, left, back (within the bay).

so_display_style_fixture_stg.txt

This table contains the cross reference between display style and fixture types.

Notes

This is a compatibility object that defines that display styles can be used for the distinct fixture types.

Table 14–87 so_display_style_fixture_stg.txt

Position	Column Name	Data Type	Description
1	DISPLAY_STYLE_KEY	VARCHAR2(80)	External display style ID
2	FIXTURE_TYPE	VARCHAR2(80)	Fixture type that supports the display style (Shelf, Pegboard or Freezer chest)

so_display_style_stg.txt

This table contains the list of available display styles for products.

Notes

A display style defines the product physical dimensions as well as the different options that can be used to place the product within a planogram.

Table 14–88 so_display_style_stg.txt

Position	Column Name	Data Type	Description
1	KEY	VARCHAR2(80)	External ID that identified the display style
2	NAME	VARCHAR2(80)	Display style name
3	DESCR	VARCHAR2(80)	Display style description
4	DEPTH	NUMBER(18,4)	Product depth relevant to default from t-0 position
5	HEIGHT	NUMBER(18,4)	Product height relevant to default from t-0 position
6	WIDTH	NUMBER(18,4)	Product width relevant to default from t-0 position
7	FINGER_SPACE_ABOVE	NUMBER(18,4)	Product gap above between same product
8	FINGER_SPACE_BESIDE	NUMBER(18,4)	Product gap beside between same product

Table 14–88 (Cont.) so_display_style_stg.txt

Position	Column Name	Data Type	Description
9	FINGER_SPACE_BEHIND	NUMBER(18,4)	Product gap behind between same product
10	INTER_PRODUCT_GAP	NUMBER(18,4)	Gap between products. This field captures gap beside between different products.
11	MAX_STACK	NUMBER(10)	Number of items that could be stacked together; equal to 1 if not stackable
12	NESTING_HEIGHT	NUMBER(18,4)	Product nesting height. The product will not allow nesting if all nesting dimensions are 0
13	NESTING_WIDTH	NUMBER(18,4)	Product nesting width. The product will not allow nesting if all nesting dimensions are 0
14	NESTING_DEPTH	NUMBER(18,4)	Product nesting depth. The product will not allow nesting if all nesting dimensions are 0
15	COLOR	VARCHAR2(30)	Product color. It can be NULL.
16	DISPLAY_UNITS	NUMBER(3)	For unit display style it is 1, else it is something >1
17	TYPE	VARCHAR2(80)	Display style type. CASE, UNIT, TRAY or other display style type.

so_disp_style_orientation_stg.txt

This is the cross reference between display style and product valid orientations.

Notes

This table contains the list of valid orientations for a given display style. The table includes a default orientation that should be used for each display style.

Table 14–89 so_disp_style_orientation_stg.txt

Position	Column Name	Data Type	Description
1	DISPLAY_STYLE_KEY	VARCHAR2(80)	External display style ID
2	ORIENTATION_KEY	VARCHAR2(80)	External orientation ID
3	DEFAULT_FLG	VARCHAR2(30)	Y indicates the orientation should be consider as the default for the display style. N indicates the orientation is valid for the display style but not a default.

so_fixture_disp_config_stg.txt

This table receives the historical planogram product data for shelf fixtures.

Notes

This table receives the list of products and distribution of them across historical planograms. It includes the position and orientation of the products within the planogram as well as the number of facings per products. This data will be later used to estimate facing lift.

Table 14–90 *so_fixture_disp_config_stg.txt*

Position	Column Name	Data Type	Description
1	DISPLAY_STYLE_KEY	VARCHAR2(80)	External ID that identifies a display style associated to a single product.
2	BAY_KEY	VARCHAR2(80)	Bay external ID. Bay is a direct dependent of the planogram.
3	FIXTURE_KEY	VARCHAR2(80)	Fixture external ID. Fixture is a direct dependent of the Bay.
4	SHELF_KEY	VARCHAR2(80)	Shelf external ID. Shelf is a direct dependent of shelf fixture.
5	ORIENTATION_KEY	VARCHAR2(80)	External ID that identifies orientation used.
6	POS_X	NUMBER(18,4)	Position of the product on the X axis. Origin point: bottom, left, back
7	POS_Y	NUMBER(18,4)	Position of the product on the Y axis. Origin point: bottom, left, back
8	POS_Z	NUMBER(18,4)	Position of the product on the Y axis. Origin point: bottom, left, back
9	FACING_QUANTITY	NUMBER(5)	Number of facings of the product

so_fixture_stg.txt

This table contains the list of fixtures within a historical planogram.

Notes

This table has the list of fixtures that define the historical planogram layout. Fixtures are planogram components within a Bay; each bay can include one or more fixtures.

Table 14–91 *so_fixture_stg.txt*

Position	Column Name	Data Type	Description
1	KEY	VARCHAR2(80)	External fixture identifier
2	FIXTURE_TYPE	VARCHAR2(80)	Fixture type values can one of the following: Shelf, Pegboard or Freezer Chest
3	DEPTH	NUMBER(18,4)	Fixture depth
4	HEIGHT	NUMBER(18,4)	Fixture height

Table 14–91 (Cont.) so_fixture_stg.txt

Position	Column Name	Data Type	Description
5	WIDTH	NUMBER(18,4)	Fixture width
6	VERTICAL_SPACING	NUMBER(18,4)	Vertical spacing. This value is only relevant for pegboard fixtures.
7	HORIZONTAL_SPACING	NUMBER(18,4)	Horizontal spacing. This value is only relevant for pegboard fixtures.
8	MAX_LENGTH	NUMBER(18,4)	Max length. This value is only relevant for pegboard fixtures.
9	CAPACITY_X	NUMBER(18,4)	Freezer Chest Capacity X (length). This value is only relevant for freezer chest fixtures.
10	CAPACITY_Y	NUMBER(18,4)	Freezer Chest Capacity Y (depth). This value is only relevant for freezer chest fixtures.
11	CAPACITY_Z	NUMBER(18,4)	Freezer Chest Capacity Z (height). This value is only relevant for freezer chest fixtures.

so_pegboard_disp_config_stg.txt

This table receives the historical planogram product data for pegboard and freezer chest fixtures.

Notes

This table receives the list of products and distribution of them across historical planograms. It includes the position and orientation of the products within the planogram as well as the number of facings per products. This data will be later used to estimate facing lift.

Table 14–92 so_pegboard_disp_config_stg.txt

Position	Column Name	Data Type	Description
1	DISPLAY_STYLE_KEY	VARCHAR2(80)	External ID that identifies a display style associated to a single product.
2	BAY_KEY	VARCHAR2(80)	Bay external ID. Bay is a direct dependent of the planogram.
3	FIXTURE_KEY	VARCHAR2(80)	Fixture external ID. Fixture is a direct dependent of the Bay.
4	ORIENTATION_KEY	VARCHAR2(80)	External ID that identifies orientation used.
5	POS_X	NUMBER(18,4)	Position of the product on the X axis. Origin point: bottom, left, back

Table 14–92 (Cont.) so_pegboard_disp_config_stg.txt

Position	Column Name	Data Type	Description
6	POS_Y	NUMBER(18,4)	Position of the product on the Y axis. Origin point: bottom, left, back
7	POS_Z	NUMBER(18,4)	Position of the product on the Y axis. Origin point: bottom, left, back
8	FACING_QUANTITY	NUMBER(5)	Number of facings of the product

so_pog_bay_stg.txt

This table receives the list of bays that are used to build a planogram.

Notes

A Bay is the level under the planogram that is used to position fixtures to build the final planogram layout. The bay is directly linked to a unique planogram; fixtures can then be linked to the bay.

Table 14–93 so_pog_bay_stg.txt

Position	Column Name	Data Type	Description
1	BAY_KEY	VARCHAR2(80)	Bay external ID. Bay is a direct dependent of the planogram.
2	POG_KEY	VARCHAR2(80)	External planogram identifier
3	BAY_SEQUENCE	NUMBER(3)	Sequence from left to right in which the bay appear within the planogram

so_pog_display_style_stg.txt

This table is the cross reference between historical planograms and product display styles.

Notes

Product to display style mapping. It provides a list of display styles available to choose for certain products.

Table 14–94 so_pog_display_style_stg.txt

Position	Column Name	Data Type	Description
1	POG_KEY	VARCHAR2(80)	External planogram identifier
2	DISPLAY_STYLE_KEY	VARCHAR2(80)	External display style identifier. This identifier links a historical planogram with a specific product.

so_pog_stg.txt

This table is used to receive planogram level details. This is the planogram header data.

Notes

The data in this table is used internally to generate the distinct POG Set (SO concept). The content of this table is transformed into planograms and POG Sets. The rows within this table correspond to Historical Planograms layouts that are received from external sources.

Table 14–95 so_pog_stg.txt

Position	Column Name	Data Type	Description
1	KEY	VARCHAR2(80)	External planogram identifier
2	NAME	VARCHAR2(80)	Planogram name
3	DESCR	VARCHAR2(80)	planogram description
4	SEASON_CODE	VARCHAR2(30)	Seasonal code used by the historical planogram.
5	SEASONAL_ATTRIBUTE	VARCHAR2(30)	Seasonal attribute used for the historical planogram
6	EFFECTIVE_START_DT	DATE	Planogram start date
7	EFFECTIVE_END_DT	DATE	Planogram end date
8	STATUS	VARCHAR2(30)	Identifier that describes the planogram status
9	CATEGORY_KEY	VARCHAR2(80)	POG category key. The second lowest level of POG category hierarchy
10	CATEGORY_NAME	VARCHAR2(80)	POG category name.
11	SUB_CATEGORY_KEY	VARCHAR2(80)	POG sub-category key. The lowest level of POG category hierarchy
12	SUB_CATEGORY_NAME	VARCHAR2(80)	POG sub-category name
13	DEPT_KEY	VARCHAR2(80)	POG department key
14	DEPT_NAME	VARCHAR2(80)	POG department name
15	LENGTH	NUMBER(18,4)	The total length of a planogram. It must be equal to the sum of the length for all the bays within the planogram.
16	DEPTH	NUMBER(18,4)	The total depth of a planogram. Should be equal to the greatest depth within all the fixtures in the planogram
17	HEIGHT	NUMBER(18,4)	The total height of a planogram. It should be equal to the highest fixture within the planogram.

so_pog_store_cda_stg.txt

This is an staging table to load customer defined attributes for POG/store combinations. These attributes are static values that are used as informational attributes within the UI.

Table 14–96 so_pog_store_cda_stg.txt

Position	Column Name	Data Type	Description
1	POG_KEY	VARCHAR2(80)	
2	STORE_KEY	VARCHAR2(80)	
3	ATTR_NUM_VALUE_1	NUMBER(18,4)	
4	ATTR_NUM_VALUE_2	NUMBER(18,4)	
5	ATTR_NUM_VALUE_3	NUMBER(18,4)	
6	ATTR_NUM_VALUE_4	NUMBER(18,4)	
7	ATTR_NUM_VALUE_5	NUMBER(18,4)	
8	ATTR_NUM_VALUE_6	NUMBER(18,4)	
9	ATTR_NUM_VALUE_7	NUMBER(18,4)	
10	ATTR_NUM_VALUE_8	NUMBER(18,4)	
11	ATTR_NUM_VALUE_9	NUMBER(18,4)	
12	ATTR_NUM_VALUE_10	NUMBER(18,4)	
13	ATTR_DATE_VALUE_1	DATE	
14	ATTR_DATE_VALUE_2	DATE	
15	ATTR_DATE_VALUE_3	DATE	
16	ATTR_DATE_VALUE_4	DATE	
17	ATTR_DATE_VALUE_5	DATE	
18	ATTR_STRING_VALUE_1	VARCHAR2(80)	
19	ATTR_STRING_VALUE_2	VARCHAR2(80)	
20	ATTR_STRING_VALUE_3	VARCHAR2(80)	
21	ATTR_STRING_VALUE_4	VARCHAR2(80)	
22	ATTR_STRING_VALUE_5	VARCHAR2(80)	
23	ATTR_PCT_VALUE_1	NUMBER(5,4)	
24	ATTR_PCT_VALUE_2	NUMBER(5,4)	
25	ATTR_PCT_VALUE_3	NUMBER(5,4)	
26	ATTR_PCT_VALUE_4	NUMBER(5,4)	
27	ATTR_PCT_VALUE_5	NUMBER(5,4)	

so_pog_store_stg.txt

This table includes the list of stores that used the historical planogram.

Notes

This is a cross reference between historical planograms and stores for which the planogram is valid (depending on dates).

Table 14–97 *so_pog_store_stg.txt*

Position	Column Name	Data Type	Description
1	POG_KEY	VARCHAR2(80)	External planogram identifier
2	STORE_KEY	VARCHAR2(80)	This is the external store ID, known and shared across applications.
3	EFFECTIVE_START_DT	DATE	Start date for which the historical planogram is valid for the given store
4	EFFECTIVE_END_DT	DATE	End date for which the historical planogram is valid for the given store

so_prod_display_style_stg.txt

This table is the product to display style cross reference.

Notes

Product to display style mapping. It provides a list of display styles available to choose for certain product.

Table 14–98 *so_prod_display_style_stg.txt*

Position	Column Name	Data Type	Description
1	PRODUCT_KEY	VARCHAR2(80)	This is the external ID that is known and shared across applications. This cannot be a placeholder product.
2	DISPLAY_STYLE_KEY	VARCHAR2(80)	External ID that identifies the display style
3	DEFAULT_FLG	VARCHAR2(1)	Y - Indicates the default display style for a given product. N - Indicates the combination should not be considered as a default. Each product should have one default display style.

so_shelf_stg.txt

This table contains the list of shelves within a historical planogram that uses shelf fixtures.

Notes

This table has the list of shelves that define a shelf fixture within the historical planogram layout. Shelf are planogram components within a Shelf fixture; each shelf fixture can contain one or more shelves.

Table 14–99 *so_shelf_stg.txt*

Position	Column Name	Data Type	Description
1	KEY	VARCHAR2(80)	External shelf identifier
2	DEPTH	NUMBER(18,4)	Shelf depth
3	HEIGHT	NUMBER(18,4)	Shelf Height. This is the physical shelf height/thickness
4	WIDTH	NUMBER(18,4)	Shelf width

Assortment and Space Optimization – Miscellaneous Files

so_pog_assort_mapping_stg.txt

This is the staging table used to receive the cross reference data to perform POG to assortment mapping.

Notes

This file contains the POG hierarchy to assortment product mapping information. This data is used to identify that POG should be used for each product within an assortment.

Table 14–100 *so_pog_assort_mapping_stg.txt*

Position	Column Name	Data Type	Description
1	POG_DEPT_KEY	VARCHAR2(80)	This is the POG dept key. This is a POG hierarchy external key known to the external source. This is a mandatory value.
2	POG_CATEGORY_KEY	VARCHAR2(80)	This is the POG category key. This is a POG hierarchy external key known to the external source. This is a mandatory value.
3	POG_SUB_CATEGORY_KEY	VARCHAR2(80)	This is the POG subcategory key. This is a POG hierarchy external key known to the external source. This is a mandatory value.
4	ASSORT_PRODUCT_LEVEL	VARCHAR2(80)	This is an identifier to the product level within the product hierarchy. This value must match with the product hierarchy available within SO.

Table 14–100 (Cont.) so_pog_assort_mapping_stg.txt

Position	Column Name	Data Type	Description
5	ASSORT_PRODUCT_KEY	VARCHAR2(80)	This is an identifier to a node within the merchandise hierarchy. It could be a specific product or any other node not higher than the assortment product category level within the merchandise hierarchy.
6	DEMAND_SPREAD_FACTOR	NUMBER(6,3)	This is the demand spread factor. This value is normally null, meaning a 100% demand is assigned to the POG node, in specific cases where the product is placed on multiple POG nodes a demand spread factor could be used to split the demand across those multiple POGs.

so_pog_assort_seas_mapping_stg.txt

This is the staging table used to receive the cross reference data to perform the assortment to POG season mapping.

Notes

Once the mapping from product to POG has been performed, a second pass looks into this data to identify the specific POG season to use, based on the assortment start date.

Table 14–101 so_pog_assort_seas_mapping_stg.txt

Position	Column Name	Data Type	Description
1	POG_DEPT_KEY	VARCHAR2(80)	This is the POG dept key. This is a POG hierarchy external key known to the external source. This is a mandatory value.
2	POG_CATEGORY_KEY	VARCHAR2(80)	This is the POG category key. This is a POG hierarchy external key known to the external source. This is a mandatory value.
3	POG_SUB_CATEGORY_KEY	VARCHAR2(80)	This is the POG subcategory key. This is a POG hierarchy external key known to the external source. This is a mandatory value.

Table 14–101 (Cont.) so_pog_assort_seas_mapping_stg.txt

Position	Column Name	Data Type	Description
4	SEASONAL_ATTRIBUTE	VARCHAR2(30)	This field refers to a specific year independent time period (season) for a CMPO assortment and a POG set. Examples include Spring, holiday, back to school, year-round, Fall, and Winter.
5	MIN_ASSORT_START_DT	DATE	The year component is irrelevant; the year component should be delivered as 0000. This is a year independent time period. The assortment start date is matched within the date range specified by this min assort start date and the max assort start date.
6	MAX_ASSORT_START_DT	DATE	The year component is irrelevant; the year component should be delivered as 0000. This is a year independent time period. The assortment start date is matched within the date range specified by the min assort start date and this max assort start date.

so_prod_loc_repl_param_stg.txt

This is the staging table used to receive the replenishment parameters at the product/location level.

Notes

Replenishment parameters are not directly linked to any assortment. This is generic data; however, these parameters must exist for all the known product/store combinations provided within an assortment.

Table 14–102 so_prod_loc_repl_param_stg.txt

Position	Column Name	Data Type	Description
1	PRODUCT_KEY	VARCHAR2(80)	This is the external ID that is known and shared across applications. In case of placeholder products this field contains CMPO placeholder product key.
2	LOCATION_KEY	VARCHAR2(80)	This is the external store ID, known and shared across applications.
3	CASEPACK	NUMBER(18,4)	Product casepack for the given store

Table 14–102 (Cont.) so_prod_loc_repl_param_stg.txt

Position	Column Name	Data Type	Description
4	REPLENISHMENT_ FREQ	NUMBER(18,4)	Replenishment frequency (RF) = number of replenishments to the shelf per week
5	REPLENISHMENT_TYPE	NUMBER(10)	Replenishment source/type - two options: 1 = from DC/vendor, 2 = from back room
6	TRANSIT_TIME	NUMBER(10,2)	Transit time (TT) = number of days it takes an order to go from source (DC or back room) to shelf
7	SHELF_ REPLENISHMENT_TT	NUMBER(10)	Shelf replenishment trigger type. 3 options: 1 = cover demand over repl period + transit time, 2 = repl when inventory gets to a target percentage of capacity and 3 = repl when a casepack can fit.
8	SHELF_ REPLENISHMENT_ PARAM	NUMBER(18,4)	Shelf replenishment parameter (currently only applies for option 2)
9	STDEV_BOOSTER	NUMBER(10,6)	Standard deviation booster (number greater than or equal to 0, makes sense to limit to 1)
10	DAYS_OF_SALES_PER_ WK	NUMBER(3,2)	Days of sales per week (number between 1 and 7)
11	FACINGS_LIFT	NUMBER(5,4)	Facing lift parameter

so_prod_stack_height_limit_stg.txt

This table is used to accept an optional client feed that provides product-specific stacking height limits.

Table 14–103 so_prod_stack_height_limit_stg.txt

Position	Column Name	Data Type	Description
1	PRODUCT_KEY	VARCHAR2(80)	This is the external ID that is known and shared across applications.
2	STACK_HEIGHT_LIMIT	NUMBER(18,4)	This is the stacking height limit for the specific product. The value here must be provided using the same units of measure used for all other product dimensions.

Table 14–103 (Cont.) so_prod_stack_height_limit_stg.txt

Position	Column Name	Data Type	Description
3	ENABLED_FLG	VARCHAR2(1)	This flag indicates if the product stacking height limit should be used or not. Y means the value specified here is used, N means the value is ignored and the application global value is used instead for the product.

Advanced Clustering

rsestrclst.csv (cis_cluster_set_exp_vw)

This view provides an exportable set of clusters to send to CMPO.

Table 14–104 rsestrclst.csv (cis_cluster_set_exp_vw)

Position	Column Name	Data Type	Description
1	EFF_START_DT	VARCHAR2(11)	The starting date that the cluster is effective.
2	EFF_END_DT	VARCHAR2(11)	The ending date for which the cluster is effective.
3	PROD_EXT_KEY	VARCHAR2(80)	The external ID for the product hierarchy this cluster is applicable to.
4	LOC_EXT_KEY	VARCHAR2(80)	The external ID for the store location that belongs in this cluster.
5	CLUSTER_ID	NUMBER(10)	The identifier for the cluster.
6	CLUSTER_LABEL	VARCHAR2(50)	A descriptive name/label for the cluster.

cis_cluster_tmpl_stg.txt

This is a staging table used to populate the CIS_CLUSTER_TMPL and CIS_CLUSTER_TMPL_LEVEL tables with externally provided data.

The various hierarchy columns in this are considered optional. If however, a hierarchy level value is provided anywhere, then there must be a hierarchy type value for that dimension. A single column is used for the hierarchy types, since the value cannot fluctuate between the different levels.

At least one TYPE_CRITERIA_NAME# column must be provided with a value, unless the DELETE_FLG indicates to delete the TEMPLATE. The TYPE_CRITERIA_NAME# columns must be provided in sequence, with no internal gaps.

The values provided in this table must be legal combinations as defined in other metadata tables.

Notes

This interface file is not handled as part of routine weekly processing. It can be provided as a special interface to adjust the configured templates available for use in the UI.

Table 14–105 *cis_cluster_tmpl_stg.txt*

Position	Column Name	Data Type	Description
1	EXT_KEY	VARCHAR2(80)	This is the unique external ID for a template.
2	NAME	VARCHAR2(80)	The name of the template.
3	DESCR	VARCHAR2(255)	A column that offers space for more descriptive text for this template.
4	BUSOBJ_SHORT_NAME	VARCHAR2(10)	A short name of the business object that this template relates to. Related to CIS_BUSINESS_OBJECT.SHORT_NAME.
5	OBJECTIVE_NAME	VARCHAR2(50)	The name of the objective this template relates to. Relates to the CIS_OBJECTIVE.NAME column.
6	DELETE_FLG	VARCHAR2(1)	A Y/N flag to indicate whether the template should be considered deleted (Y) or not (N).
7	PROD_HIER_TYPE	VARCHAR2(255)	The NAME of the product hierarchy type that this template would use. This is an optional column.
8	LOC_HIER_TYPE	VARCHAR2(255)	The NAME of the location hierarchy type that this template would use. This is an optional column.
9	TYPE_CRITERIA_NAME1	VARCHAR2(80)	The TYPE_NAME of the CIS_TYPE_CRITERIA that is to be selected as the first cluster level.
10	PROD_HIER_LEVEL1	VARCHAR2(80)	The name of the product hierarchy level that should be selected for this level of the cluster hierarchy. This is an optional column.
11	LOC_HIER_LEVEL1	VARCHAR2(80)	The name of the location hierarchy level that should be selected for this level of the cluster hierarchy. This is an optional column.
12	TYPE_CRITERIA_NAME2	VARCHAR2(80)	The TYPE_NAME of the CIS_TYPE_CRITERIA that is to be selected as the second cluster level.

Table 14–105 (Cont.) cis_cluster_tmpl_stg.txt

Position	Column Name	Data Type	Description
13	PROD_HIER_LEVEL2	VARCHAR2(80)	The name of the product hierarchy level that should be selected for this level of the cluster hierarchy. This is an optional column.
14	LOC_HIER_LEVEL2	VARCHAR2(80)	The name of the location hierarchy level that should be selected for this level of the cluster hierarchy. This is an optional column.
15	TYPE_CRITERIA_NAME3	VARCHAR2(80)	The TYPE_NAME of the CIS_TYPE_CRITERIA that is to be selected as the third cluster level.
16	PROD_HIER_LEVEL3	VARCHAR2(80)	The name of the product hierarchy level that should be selected for this level of the cluster hierarchy. This is an optional column.
17	LOC_HIER_LEVEL3	VARCHAR2(80)	The name of the location hierarchy level that should be selected for this level of the cluster hierarchy. This is an optional column.
18	TYPE_CRITERIA_NAME4	VARCHAR2(80)	The TYPE_NAME of the CIS_TYPE_CRITERIA that is to be selected as the fourth cluster level.
19	PROD_HIER_LEVEL4	VARCHAR2(80)	The name of the product hierarchy level that should be selected for this level of the cluster hierarchy. This is an optional column.
20	LOC_HIER_LEVEL4	VARCHAR2(80)	The name of the location hierarchy level that should be selected for this level of the cluster hierarchy. This is an optional column.
21	TYPE_CRITERIA_NAME5	VARCHAR2(80)	The TYPE_NAME of the CIS_TYPE_CRITERIA that is to be selected as the fifth cluster level.
22	PROD_HIER_LEVEL5	VARCHAR2(80)	The name of the product hierarchy level that should be selected for this level of the cluster hierarchy. This is an optional column.

Table 14–105 (Cont.) cis_cluster_tmpl_stg.txt

Position	Column Name	Data Type	Description
23	LOC_HIER_LEVEL5	VARCHAR2(80)	The name of the location hierarchy level that should be selected for this level of the cluster hierarchy. This is an optional column.

cis_cluster_tmpl_prod_xref_stg.txt

This is a staging table that can be used to populate into the CIS_CLUSTER_TMPL_PROD_XREF table. This table provides the ability to associate a template to different product hierarchy members.

Although the product hierarchy related columns of this table are all optional, only one template can be missing the PROD_HIER_TYPE column, and only one template can be assigned to a PROD_HIER_TYPE/PROD_HIER_LEVEL combination where the PROD_EXT_KEY is NULL. There can be any number of products assigned to a given template.

When providing data for this interface, it is expected that all product assignments are provided for a template. Any product that was previously assigned to a template, but is not assigned in a subsequent load of data for that template, results in the removal of the prior product assignment.

Notes

This interface file is not handled as part of routine weekly processing. It can be provided as a special interface to adjust the configured templates available for use in the UI.

Table 14–106 cis_cluster_tmpl_prod_xref_stg.txt

Position	Column Name	Data Type	Description
1	CLUSTER_TMPL_EXT_KEY	VARCHAR2(80)	External identifier for the cluster template.
2	PROD_HIER_TYPE	VARCHAR2(255)	The product hierarchy type that this template is applicable to. If NULL, then this template is the default template for all products.
3	PROD_HIER_LEVEL	VARCHAR2(80)	The product hierarchy level name that this template is applicable to. This can only be NULL if the PROD_HIER_TYPE column is NULL.
4	PROD_EXT_KEY	VARCHAR2(80)	The external identifier for the product hierarchy member that is assigned to this template. If NULL, then this template is applicable to all products that match the provided PROD_HIER_TYPE and PROD_HIER_LEVEL.

cis_cluster_tmpl_lvl_attr_stg.txt

This is the staging table for optionally loading default weights for attributes to be used by a particular template level. Data is not required here; however, if it is provided, then when the template is used to set up a new cluster scenario, the weights and attributes defined here are used as the participating attributes for the cluster configuration. Otherwise, the default attributes and weights are used.

For attributes that are associated with product hierarchies, the PROD_EXT_KEY can be provided to limit scope to those products. This is an optional column for this interface.

Notes

This interface file is not handled as part of routine weekly processing. It can be provided as a special interface to adjust the configured templates available for use in the UI.

Table 14–107 *cis_cluster_tmpl_lvl_attr_stg.txt*

Position	Column Name	Data Type	Description
1	CLUSTER_TMPL_EXT_KEY	VARCHAR2(255)	The external key to identify the cluster template that these attributes are applicable to.
2	CLUSTER_LEVEL	NUMBER(10)	The level number of the cluster hierarchy that these attribute weights should be associated with.
3	ATTR_EXT_KEY	VARCHAR2(50)	The external key ATTR_EXT_KEY Field contains the reference to the loader to identify an attribute.
4	PROD_EXT_KEY	VARCHAR2(80)	For attributes that are product specific, this column allows for the specification of the products whose attributes should be used.
5	DFLT_ATTR_WEIGHT	NUMBER(18,4)	The default weight that should be used for this attribute, for this cluster level.

Returns Logistics**rl_price_elasticity_stg.txt**

This is the input data file for the staging table for specifying price elasticities. Price elasticity is defined at a product-location intersection.

Table 14–108 *rl_price_elasticity_stg.txt*

Position	Column Name	Data Type	Description
1	PROD_HIER_TYPE_NAME	VARCHAR2(255)	The name of the appropriate product hierarchy.

Table 14–108 (Cont.) *rl_price_elasticity_stg.txt*

Position	Column Name	Data Type	Description
2	PROD_EXT_KEY	VARCHAR2(80)	External product identifier.
3	LOC_HIER_TYPE_NAME	VARCHAR2(255)	The name of the appropriate location hierarchy.
4	LOC_EXT_KEY	VARCHAR2(80)	External location identifier.
5	PRICE_ELASTICITY	NUMBER	The name of appropriate product hierarchy.

rl_price_ladder_stg.txt

This is the input data file for the staging table that defines price ladders for each product. These are all of the valid price points for a product.

Table 14–109 *rl_price_ladder_stg.txt*

Position	Column Name	Data Type	Description
1	PROD_HIER_TYPE_NAME	VARCHAR2(255)	The name of the product hierarchy type associated with this store.
2	PROD_EXT_KEY	VARCHAR2(80)	External product identifier.
3	FULL_PRICE	NUMBER	Full or regular dollar price.
4	PRICE_POINT1	NUMBER	Price Point 1
5	PRICE_POINT2	NUMBER	Price Point 2
6	PRICE_POINT3	NUMBER	Price Point 3
7	PRICE_POINT4	NUMBER	Price Point 4
8	PRICE_POINT5	NUMBER	Price Point 5
9	PRICE_POINT6	NUMBER	Price Point 6
10	PRICE_POINT7	NUMBER	Price Point 7
11	PRICE_POINT8	NUMBER	Price Point 8
12	PRICE_POINT9	NUMBER	Price Point 9
13	PRICE_POINT10	NUMBER	Price Point 10
14	PRICE_POINT11	NUMBER	Price Point 11
15	PRICE_POINT12	NUMBER	Price Point 12
16	PRICE_POINT13	NUMBER	Price Point 13
17	PRICE_POINT14	NUMBER	Price Point 14
18	PRICE_POINT15	NUMBER	Price Point 15
19	PRICE_POINT16	NUMBER	Price Point 16
20	PRICE_POINT17	NUMBER	Price Point 17
21	PRICE_POINT18	NUMBER	Price Point 18
22	PRICE_POINT19	NUMBER	Price Point 19

Table 14–109 (Cont.) rl_price_ladder_stg.txt

Position	Column Name	Data Type	Description
23	PRICE_POINT20	NUMBER	Price Point 20
24	PRICE_POINT21	NUMBER	Price Point 21
25	PRICE_POINT22	NUMBER	Price Point 22
26	PRICE_POINT23	NUMBER	Price Point 23
27	PRICE_POINT24	NUMBER	Price Point 24
28	PRICE_POINT25	NUMBER	Price Point 25

rl_shipping_cost_stg.txt

This is the input data file for the staging table that defines shipping costs between store locations. Shipping costs are defined per unit of product.

Table 14–110 rl_shipping_cost_stg.txt

Position	Column Name	Data Type	Description
1	SHIP_FROM_LOC_ID	VARCHAR2(30)	Identifier for the location being shipped from.
2	SHIP_FROM_LOC_TYPE	VARCHAR2(30)	Store location or distribution center (S or D).
3	SHIP_FROM_LOC_NAME	VARCHAR2(30)	Store location or distribution center name.
4	SHIP_TO_LOC_ID	VARCHAR2(30)	Identifier for the location being shipped to.
5	SHIP_TO_LOC_NAME	VARCHAR2(30)	Store location or distribution center name.
6	SHIP_COST_PER_UNIT	NUMBER	Dollar cost of shipping one unit.

Market Basket Insights**mba_arm_srvc_loc_stg.txt**

This interface file allows for the limiting of locations for various Association Rules services.

Table 14–111 mba_arm_srvc_loc_stg.txt

Position	Column Name	Data Type	Description
1	SRVC_NAME	VARCHAR2(30)	The name of the Association Rule service. Valid values are ARM_PH, ARM_PH_PROMO, ARM_PH_CS, or a Null value.
2	LOC_EXT_CODE	VARCHAR2(80)	External identifier for the location hierarchy node

Table 14–111 (Cont.) mba_arm_srvc_loc_stg.txt

Position	Column Name	Data Type	Description
3	INCLUDE_FLG	VARCHAR2(1)	A Yes/No flag to indicate whether the location should be included (Y) or excluded (N) from processing.

mba_arm_run_exp.txt

This view provides an export of the results of the association rule calculations, which will include the components of the association, as well as several metrics to quantify it.

Table 14–112 mba_arm_run_exp.txt

Position	Column Name	Data Type	Description
1	RUN_ID	NUMBER(10)	Internal ID to identify a run of the association rule mining process.
2	RUN_TYPE	VARCHAR2(30 CHAR)	Implementation name for the particular run.
3	DATE_FROM	DATE	Starting date used to generate the association rules (YYYY-MM-DD)
4	DATE_TO	DATE	Ending date used to generate the association rules (YYYY-MM-DD)
5	IF_HIER_LEVEL	VARCHAR2(80 CHAR)	Hierarchy level which will be reported on the IF side of any association rule.
6	THEN_HIER_LEVEL	VARCHAR2(80 CHAR)	Hierarchy level which will be reported on the THEN side of any association rule.
7	LOC_EXT_KEY	CHAR	When applicable, the external identifier for the location hierarchy node which the process used during its processing.
8	CUSTSEG_EXT_KEY	CHAR	When applicable, the external identifier for the customer segment which the process used during its processing.
9	TOT_TXN_CNT	NUMBER(10)	The total number of sales transactions which were processed by this run.
10	TOT_SLS_AMT	NUMBER(20,4)	The total of sales amounts which were processed by this run.
11	TOT_SLS_PROFIT_AMT	NUMBER(20,4)	The total of sales profit amounts which were processed by this run.

Table 14–112 (Cont.) mba_arm_run_exp.txt

Position	Column Name	Data Type	Description
12	TOT_SLS_QTY	NUMBER(20,4)	The total of sales quantities which were processed by this run.
13	CREATE_DT	DATE	The date the process was executed. (YYYY-MM-DD)

mba_arm_result_exp.txt

This view provides an export of the results of the association rule calculations, which will include the components of the association, as well as several metrics to quantify it.

Table 14–113 mba_arm_result_exp.txt

Position	Column Name	Data Type	Description
1	RUN_ID	NUMBER(10)	This a foreign key to the MBA_ARM_RUN_EXP_VW.RUN_ID, to associate with the summary info it provides.
2	SET_SIZE	NUMBER(2)	The number of components of the rule, in both the IF and THEN components.
3	IF_PROD_EXT_KEY1	VARCHAR2(80 CHAR)	The external key for the first if product hierarchy member of the association rule.
4	IF_PROD_EXT_KEY2	VARCHAR2(80 CHAR)	The external key for the second if product hierarchy member of the association rule.
5	IF_PROD_EXT_KEY3	VARCHAR2(80 CHAR)	The external key for the third if product hierarchy member of the association rule.
6	IF_PROMO_FLG1	VARCHAR2(1)	When applicable, a Y/N flag to indicate if the corresponding product hierarchy member was promoted (Y) or not (N).
7	IF_PROMO_FLG2	VARCHAR2(1)	When applicable, a Y/N flag to indicate if the corresponding product hierarchy member was promoted (Y) or not (N).
8	IF_PROMO_FLG3	VARCHAR2(1)	When applicable, a Y/N flag to indicate if the corresponding product hierarchy member was promoted (Y) or not (N).

Table 14–113 (Cont.) mba_arm_result_exp.txt

Position	Column Name	Data Type	Description
9	THEN_PROD_EXT_KEY	VARCHAR2(80 CHAR)	The external key for the then product hierarchy member of the association rule.
10	THEN_PROMO_FLG	VARCHAR2(1)	When applicable, a Y/N flag to indicate if the corresponding product hierarchy member was promoted (Y) or not (N).
11	RULE_TXN_COUNT	NUMBER(10)	The number of transactions where this set of items were sold together.
12	RULE_FREQ	NUMBER(38,20)	The frequency expressed as a percent of transactions, where this set of items were sold together.
13	RULE_CONF	NUMBER(38,20)	The confidence that the association rule is not caused by random chance.
14	RULE_LIFT	NUMBER(38,20)	The lift of the association rule, which is another means for validating the association rule is not a random occurrence.
15	RULE_REV_CONF	NUMBER(38,20)	The reverse confidence that the association rule is not caused by random chance. This is similar to rule_conf, although the components of the calculation are reversed.
16	IF_TOT_TXN_COUNT	NUMBER(10)	The total number of transactions (not constrained by the associated then products) where the if components were sold together.
17	IF_SLS_AMT	NUMBER(20,4)	The total sales amount for the if products for transactions with this association rule.
18	IF_SLS_PROFIT_AMT	NUMBER(20,4)	The total sales profit amount for the if products for transactions with this association rule.
19	IF_SLS_QTY	NUMBER(20,4)	The total sales quantity for the if products for transactions with this association rule.

Table 14–113 (Cont.) mba_arm_result_exp.txt

Position	Column Name	Data Type	Description
20	THEN_TOT_TXN_COUNT	NUMBER(10)	The total number of transactions (not constrained by the associated if products) where the then products were sold.
21	THEN_SLS_AMT	NUMBER(20,4)	The total sales amount for the then products for transactions with this association rule.
22	THEN_SLS_PROFIT_AMT	NUMBER(20,4)	The total sales profit amount for the then products for transactions with this association rule.
23	THEN_SLS_QTY	NUMBER(20,4)	The total sales quantity for the then products for transactions with this association rule.

Security

Advanced Science Cloud Services uses web services to send information to Customer Engagement and to provide configuration and application incremental data and reports to customers.

Web Services

The web service is stateless, so state is not stored or managed. Pagination such as the batch size of data or parameters such as export data time, product, location, and so on are used to manage payload size and handle session time outs.

SOAP

An outbound interface is provided to send customer segments and its members to ORCE (Customer Engagement). This interface supports the following security features.

Message authentication is enabled in ORCE, and Advanced Science Cloud Services messages include authentication information in the HTTP header for the message. This authentication information is specific to ORCE and is stored in credential stores. Credential stores are created or updated from the Data Management task, enabled for an administrator. A Base64 encoding tool is used to encode the authorization key that is sent as part of the message HTTP header request. The credential store uses APIs that applications can use to create, read, update, and manage credentials securely and marks code as being privileged. This affects the determination of subsequent access.

Configuration is provided that can be used to set up proxy settings for both HTTP and HTTPS.

The XML that is sent as part of the message relies on marshalling and un-marshalling to or from Java objects generated using the WSDL/Schema exposed via ORCE. This enforced XML generated is well formed and valid. It is the responsibility of ORCE to convert XML, and Advanced Science Cloud Services does not complete any XML conversion. There are no concerns concerning XXE and XEE.

REST

An outbound interface exports data (GET request) and uses REST to expose data. These web services are REST-based; it is assumed that callers are familiar with basic REST principles (such as the use of HTTP verbs). AC and ASO export web services can serve as a means of obtaining incremental update data from a specified point in time. All services support the query parameter `contentType` and the HTTP header `content-type`, with supported values `application/json` and `application/xml`. The query parameter takes precedence; if no content type is supplied, then `application/json` serves as the default. Basic authentication is used, so you may use any client software that supports it. Authorization is done for ADF-LDAP (OID) mapped roles, and only administrator roles are used. That is, the calling user must be assigned a duty that is mapped to the defined administrator roles. JSON/XML parsing is done using standard JAXB request parameters that are validated before data is fetched.

Batch Processing

This chapter provides an overview of the batch processing capabilities available for Oracle Retail Advanced Science.

Overview

The implementation process for Cloud Services involves loading data files for dimensions and fact data into the database. For new implementations, the best practice is to test the interfaces in a logical sequence, in small test cycles, using a Custom Batch Request process.

Once all required data has been loaded and all interfaces have been tested, the scheduled batch cycles that perform different tasks can be used, depending on the frequency involved. Oracle Retail Advanced Science has INTRADAY processes that are used for ASO, as well as DAILY, WEEKLY, and QUARTERLY batch cycles, each of which performs different tasks, depending on which applications are being used.

Custom Batch Requests

A custom batch request provides some flexibility in the execution of batch routines during the application initialization and setup stage. This process should not be used once the application is running its normal scheduled batch cycles. During this stage of the application setup, it is generally necessary to test interfaces to make sure they follow the correct formats and contain the proper data. In this way, an implementer can perform tests in a self-sufficient manner.

Managing Custom Batch Requests

To initiate a custom batch, upload a PROCESS_QUEUE file that contains entries to trigger the execution of the processes that are associated with those identifiers. Since most processes are triggered based on the receipt of inbound files, or may be a request to trigger the execution of processes required to create an outbound file, the values that can be used inside the PROCESS_QUEUE file are generally the names of the data files. The values that can be used to trigger other batch steps are described in [Table 15-1](#).

Once the PROCESS_QUEUE has been uploaded to the inbound directory of the FTP server, a PROCESS_QUEUE.complete file can be uploaded and created. This triggers the execution of the batch steps. Once the batch process is complete, a verification email notification is sent, provided the Manage Configuration screen has been configured for such email notification.

If the `PROCESS_QUEUE` contains a list of any inbound data files, these files must be uploaded prior to the creation of the `PROCESS_QUEUE.complete` file.

After the batch process completes, a file named `PROCESS_QUEUE.log` is created in the `EXPORT` directory of the FTP server. This file contains any details that may be relevant for the implementer. It may include `SQL*Loader` log file contents if errors occurred during the processing. Log files for the programs that were executed may also be included. Such information can help in determining the cause of the error. When the batch process completes, if any outbound files to be created are placed in the `EXPORT` directory on the FTP server so that they can be retrieved.

Handling Data Files

For the process described in this section, it is assumed that the `PROCESS_QUEUE` file contains the value of `W_PRODUCT_DS.dat`, which can trigger the execution of the batch processing for loading that file.

The data to be processed can be provided as a text file (for example, `W_PRODUCT_DS.dat`) or as a compressed file (for example, `W_PRODUCT_DS.dat.gz`). For RI interfaces, a context file can also be provided that lists the columns in the interface either as a text file (for example, `W_PRODUCT_DS.dat.ctx`) or a compressed file (for example, `W_PRODUCT_DS.dat.ctx.gz`). The `PROCESS_QUEUE` file specifies the interface name of `W_PRODUCT_DS.dat`, and the process that collects the data files then retrieves any file of these filename patterns.

If the process request requires that multiple files be processed, these files can also be provided in a zip file. The file handler looks for a file named `ORASE_PROCESS_TRIGGER.zip`, unzips the contents, and uses any files listed in `PROCESS_QUEUE`. If a file that was previously included in the `ORASE_PROCESS_TRIGGER.zip` file must be adjusted, it is possible to send that file individually, so that the entire zip file does not need to be recreated and retransmitted.

Supported `PROCESS_QUEUE` Trigger Values

In addition to supporting any inbound or outbound data files, some additional values, described in [Table 15-1](#), can be used to trigger the execution of some specific batch processes.

Table 15-1 *Trigger Values*

Process Queue Trigger Text	Description
<code>SIL_INIT</code>	Initializes RI MCAL Current Date. This may be used as required to advance the business date.
<code>SO_POST_PROC</code>	Triggers the execution of a series of steps that perform the data processing required to operate after the successful staging and loading of individual SO data files.
<code>EXPORT_PREP_DAILY</code>	Many ORASE export files provide incremental data exports that have occurred since the most recent export process was run. This step resets the from/to date range for daily exports to include changes up through the time this process is executed. The from date is set to the date/time that was previously the to date value.

Table 15–1 (Cont.) Trigger Values

Process Queue Trigger Text	Description
EXPORT_PREP_WEEKLY	Many of the application export files provide incremental data exports for periods that begin with the date of the last time the export process was run. This step resets the from/to date range for weekly exports so that it includes changes up through the time this process is executed. The From date is set to the date and time that were previously used for the To values.
EXPORT_PREP_QUARTERLY	Many of the application export files provide incremental data exports for periods that begin with the date of the last time the export process was run. This step resets the from/to date range for quarterly exports so that it includes changes up through the time this process is executed. The From date is set to the date and time that were previously used for the To date values.
EXPORT_PREP_INTRADAY	Many of the application export files provide incremental data exports for periods that begin with the date of the last time the export process was run. This step resets the from/to date range for intraday exports so that it includes changes up through the time this process is executed. The From date is set to the date and time that were previously used for the To date values.

ORASE Incremental Exports

As described in [Table 15–1](#), all incremental export files for Oracle Retail Advanced Science are controlled by a set of dates that define the beginning and ending range of data to be exported. This data is stored in a configuration table called RSE_EXP_GRP, and can be seen in the Manage Configuration screen. Although the screen does not show the time associated with the dates, the export uses a time component of the date columns. Each incremental export has a date associated with the data to be exported. Only data that has a date value between the FROM_DT and TO_DT columns of the RSE_EXP_GRP that it is associated with, will be exported when the export file is created.

When testing an application, it is important to realize that if a test export of data is required, you must make sure that data is available to be exported and that the data is associated with a date that is in the range of the export group. If an export runs and does not produce any data in the file, you should check the values of the Export Group to ensure the dates were not set incorrectly.

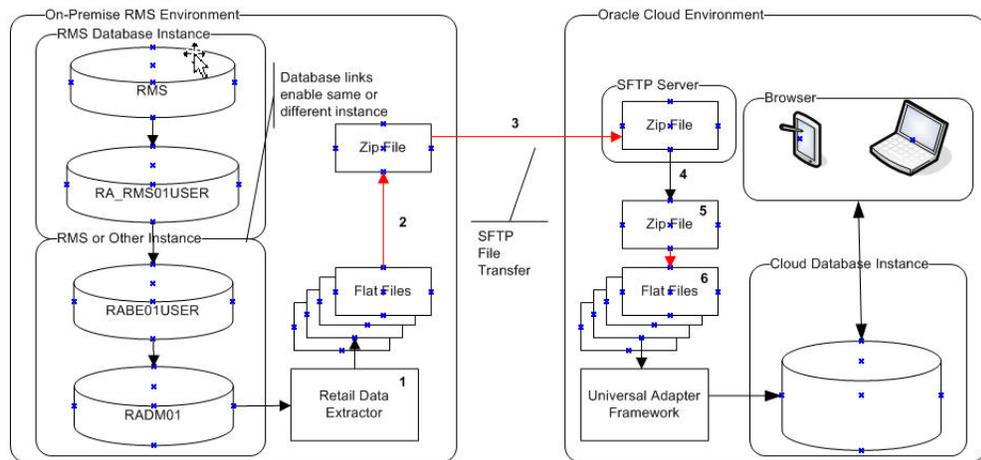
When you create or process data in the application user interface, want to test the export of that data, you must advance Export Group's date range by running the appropriate export preparation step as described in this chapter. This causes the date range to advance and enables the exporting of the data that is available for exporting. Note that if the Export Group date range is advanced too many times, the data that you want to export may no longer be in the current range for exporting.

You may encounter such issues when using this custom batch process to trigger the execution of exports; however, these issues will not occur once the application is running the batch routines in an automated manner, because the batch processes are only executed once per batch cycle.

Batch Process Flow

[Figure 15–1](#) illustrates the batch process flow.

Figure 15–1 Batch Process Flow



Here is the process.

1. The on-premise batch shell script extracts data to files initiated by the customer scheduler.
2. Merch batch script creates the zip file named RI_RMS_DATA.zip. Additionally, zip files named RI_CE_DATA.zip and RI_MFP_DATA.zip should be created.
3. You should sftp the three zip files. Then, create a file named "COMPLETE" in the sftp directory COMMAND.
4. After the COMPLETE file is found in the COMMAND directory, the file watcher initiates the processing of files and places them in the landing directory of the cloud server.
5. The presence of the COMPLETE file in the landing directory releases the batch load processing.
6. The batch load process begins with tasks that
 - a. Archive the files that have been received in a date/time stamped directory.
 - b. Perform the presence validation exercise that verifies that all expected files for the customer’s subscribed modules in the zipped files. This terminates if any expected files are missing.
 - c. Clear the previous day’s files from the \$MMHOME/data/staging directory.
 - d. Unzip the zip file into the \$MMHOME/data/staging directory.

Table 15–2 lists the zip files.

Table 15–2 Supported Zip Files

Zip File Name	Frequency	File Type	Notes
RI_RMS_DATA.zip	Daily	Inbound	All files which start with W_* can be placed in any combination of the RI*zip files.
RI_CE_DATA.zip	Daily	Inbound	All files which start with W_* can be placed in any combination of the RI*zip files.
RI_MFP_DATA.zip	Daily	Inbound	All files which start with W_* can be placed in any combination of the RI*zip files.

Table 15-2 (Cont.) Supported Zip Files

Zip File Name	Frequency	File Type	Notes
ORASE_WEEKLY.zip	Weekly	Inbound	Any inbound file that does not start with W_* and has a weekly frequency can be placed in here.
ORASE_INTRADAY.zip	Intraday	Inbound	Any inbound file that has an intraday frequency can be placed in here.
ORASE_WEEKLY_extract.zip	Weekly	Outbound	Any outbound file that has a weekly frequency will be placed in here.
ORASE_INTRADAY_extract.zip	Intraday	Outbound	Any outbound file that has an intraday frequency will be placed in here.

Table 15-3 column headings have been shortened because of space considerations. The complete headings are as follows: File Name, Frequency Daily, Frequency Weekly, Frequency Quarterly, Frequency Intraday, Frequency Customer Segments, Assortment and Space Optimization, Customer Decision Trees, Demand Transference, Offer Optimization, Innovation Workbench, Market Basket Insights, Inbound/Outbound.

Table 15-3 Handling Data Files

File Name	Daily	Wkly	Qtrly	Intraday	Clustr	CS	ASO	CDT	DT	OO	IW	MBI	I/O
RA_SRC_CURR_PARAM_G.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_MCAL_PERIOD_DS.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_PRODUCT_DS.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_PRODUCT_ATTR_DS.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_PRODUCT_DS_TL.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_PARTY_PER_DS.dat	Y	Y	Y	O	R	R	R	R	O	R	O	R	I
W_RTL_PROD_HIER_IMAGE_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_PROD_CAT_DHS.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_RTL_PROD_HIER_ATTR_LKP_DHS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_RTL_ITEM_GRP1_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_INT_ORG_DHS.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_RTL_LOC_TRAITS_DS_TL.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_RTL_PROMO_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_RTL_PROMO_DS_TL.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_RTL_CHANNEL_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I

Table 15-3 (Cont.) Handling Data Files

File Name	Daily	Wkly	Qrtly	Intrday	Clustr	CS	ASO	CDT	DT	OO	IW	MBI	I/O
W_EMPLOYEE_DS.dat	Y	Y	Y	O	R	R	R	O	O	O	O	R	I
W_EXCH_RATE_GS.dat	Y	Y	Y	O	R	R	R	O	O	O	O	R	I
W_INT_ORG_DS.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_INT_ORG_ATTR_DS.dat	Y	Y	Y	R	O	O	O	O	O	O	O	O	I
W_INT_ORG_DS_TL.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_PARTY_ORG_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_PARTY_ATTR_DS.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_DOMAIN_MEMBER_DS_TL.dat	Y	Y	Y	R	R	R	R	R	R	R	R	R	I
W_RTL_PRODUCT_IMAGE_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_REASON_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_RTL_PRODUCT_COLOR_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_RTL_PRODUCT_ATTR_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_RTL_PRODUCT_BRAND_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_RTL_PRODUCT_BRAND_DS_TL.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_RTL_CMG_PRODUCT_MTX_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I
W_RTL_CONSUMERSEG_DS.dat	Y	Y	Y	O	O	O	O	O	O	O	O	O	I

Table 15-3 (Cont.) Handling Data Files

File Name	Daily	Wkly	Qrtly	Intrday	Clustr	CS	ASO	CDT	DT	OO	IW	MBI	I/O
W_RTL_CONSUMER_GRP_DS.dat	Y	Y	Y			O							I
W_RTL_CO_HEAD_DS.dat	Y	Y	Y		O	O		O	O	O	O	O	I
W_RTL_CO_HEAD_STATUS_FS.dat	Y	Y	Y		O	O		O	O	O	O	O	I
W_RTL_CO_LINE_DS.dat	Y	Y	Y		O	O		O	O	O	O	O	I
W_RTL_CO_LINE_STATUS_FS.dat	Y	Y	Y		O	O		O	O	O	O	O	I
W_RTL_CO_SHIP_TYPE_DS.dat	Y	Y	Y		O	O		O	O	O	O	O	I
W_RTL_CUSTSEG_ALLOC_DS.dat	Y	Y	Y		R								I
W_RTL_MARKET_PRODUCT_DS.dat	Y	Y	Y			O							I
W_RTL_MARKET_PROD_DHS.dat	Y	Y	Y			O							I
W_RTL_MARKET_PRODUCT_DS_TL.dat	Y	Y	Y			O							I
W_RTL_MARKET_PRODUCT_MTX_DS.dat	Y	Y	Y			O							I
W_RTL_MARKET_PROD_ATTR_DS.dat	Y	Y	Y		O	O		O	O	O	O	O	I
W_RTL_MARKET_PROD_ATTR_MTX_DS.dat	Y	Y	Y		O	O		O	O	O	O	O	I
W_RTL_MARKET_PROD_BRAND_DS.dat	Y	Y	Y		O	O		O	O	O	O	O	I
W_RTL_MARKET_PROD_DH_MTX_DS.dat	Y	Y	Y		O	O		O	O	O	O	O	I

Table 15-3 (Cont.) Handling Data Files

File Name	Daily	Wkly	Qrtly	Intrday	Clustr	CS	ASO	CDT	DT	OO	IW	MBI	I/O
W_RTL_MKTSLT_TA_CHI_CNG_WK_FS.dat	Y	Y	Y			O							I
W_RTL_MKTSLT_TA_CHI_HG_WK_FS.dat	Y	Y	Y			O							I
W_RTL_TRADE_AREA_DS.dat	Y	Y	Y		O	O	O	O	O	O	O	O	I
W_RTL_TRADE_AREA_LOC_MTX_DS.dat	Y	Y	Y		O	O	O	O	O	O	O	O	I
W_RTL_RECLASS_IT_SC_CL_TMP.dat	Y	Y	Y		R	R	R	R	R	R	R	R	I
W_RTL_IT_LC_DEL_TMP.dat	Y	Y	Y		R	R	R	R	R	R	R	R	I
W_RTL_RECLASS_DP_GP_TMP.dat	Y	Y	Y		R	R	R	R	R	R	R	R	I
W_RTL_ITEM_DEL_TMP.dat	Y	Y	Y		R	R	R	R	R	R	R	R	I
W_RTL_SLS_TRX_IT_LC_DY_FS.dat	Y	Y	Y		O	R	O	R	O	R	O	R	I
W_RTL_SLSFC_IT_LC_WK_FS.dat	Y	Y	Y		O								I
W_RTL_CUSTSEG_DS.dat	Y	Y	Y		O	O	O	R	R	R	O	R	I
W_RTL_CUST_CUSTSEG_DS.dat	Y	Y	Y		O	O	O	R	R	R	O	R	I
rse_like_loc_stg.txt	Y	Y	Y		O								I
rse_prod_attr_grp_value_stg.txt	Y	Y	Y		R	R	R	R	R	R	R	R	I
rse_prod_attr_value_xref_stg.txt	Y	Y	Y		R	R	R	R	R	R	R	R	I
rse_md_cda_stg.txt	Y	Y	Y		O								I
rse_md_cda_values_stg.txt	Y	Y	Y		O								I

Table 15-3 (Cont.) Handling Data Files

File Name	Daily	Wkly	Qrtly	Intrday	Clustr	CS	ASO	CDT	DT	OO	IW	MBI	I/O
rse_pr_lc_cda_stg.txt	Y	Y	Y	O									I
rse_pr_lc_cal_cda_stg.txt	Y	Y	Y	O									I
cis_cluster_tmpl_stg.txt	Y	Y	Y	O	O								I
cis_cluster_tmpl_lv1_attr_stg.txt	Y	Y	Y	O	O								I
cis_cluster_tmpl_prod_xref_stg.txt	Y	Y	Y	O	O								I
rse_wkly_sls_stg.txt	Y	Y	Y	R				R	O				I
rse_wkly_sls_seg_stg.txt	Y	Y	Y					R	O				I
rse_fake_cust_stg.txt	Y	Y	Y	O	O								I
dt_loc_wk_excl_stg.txt	Y	Y	Y						O				I
dt_prod_loc_excl_stg.txt	Y	Y	Y						R				I
cdt_import.tar.gz	Y	Y	Y					R					I
rsestrclst.csv	Y	Y	Y	R	R								O
so_assortment_finalized_stg.txt			Y				R						I
so_assort_phprod_finalized_stg.txt			Y				R						I
so_assortment_stg.txt	Y	Y	Y				R						I
so_assort_cluster_stg.txt	Y	Y	Y				R						I
so_assort_cluster_member_stg.txt	Y	Y	Y				R						I
so_assort_phprod_like_prod_stg.txt	Y	Y	Y				R						I
so_assort_product_strcltr_stg.txt	Y	Y	Y				R						I

Table 15-3 (Cont.) Handling Data Files

File Name	Daily	Wkly	Qrtly	Intrday	Clustr	CS	ASO	CDT	DT	OO	IW	MBI	I/O
so_assort_proloc_pricecost_stg.txt			Y				R						I
so_assort_proloc_fcst_stg.txt			Y				R						I
so_assort_phprod_attr_stg.txt			Y				R						I
so_pog_stg.txt			Y				R						I
so_pog_store_stg.txt			Y				R						I
so_pog_store_cda_stg.txt			Y				R						I
so_pog_bay_stg.txt			Y				R						I
so_prod_display_style_stg.txt			Y				R						I
so_pog_display_style_stg.txt			Y				R						I
so_fixture_stg.txt			Y				R						I
so_fixture_stg.txt so_bay_bay_fixture_stg.txt			Y				R						I
so_shelf_stg.txt so_bay_fixture_shelf_stg.txt			Y				R						I
so_disp_style_orientation_stg.txt			Y				R						I
so_display_style_fixture_stg.txt			Y				R						I
so_fixture_disp_config_stg.txt			Y				R						I
so_pegboard_disp_config_stg.txt			Y				R						I
so_prod_loc_repl_param_stg.txt			Y				R						I
so_prod_stack_height_limit_stg.txt			Y				R						I

Table 15-3 (Cont.) Handling Data Files

File Name	Daily	Wkly	Qrtly	Intrday	Clustr	CS	ASO	CDT	DT	OO	IW	MBI	I/O
so_pog_assort_mapping_stg.txt			Y				R						I
so_pog_assort_mapping_stg.txt			Y				R						I
dt_mdl_prod_exp_stg.txt	Y	Y	Y					R					I
ipdmdfrpcti.csv			Y						R				O
attr.csv.dat	Y	Y	Y					R					O
drtyattrvalbx.csv.ovr	Y	Y	Y					R					O
drtyassrelasv.csv.ovr	Y	Y	Y					R					O
drtyattrwgtv.csv.ovr	Y	Y	Y					R					O
drtyminv.csv.ovr.gz	Y	Y	Y					R					O
dt_assort_mult.csv	Y	Y	Y					R					O
dt_new_items.csv	Y	Y	Y					R					O
dt_new_item_ros.csv	Y	Y	Y					R					O
cdt_export.tar.gz	Y	Y	Y					R					O
so_assort_aiprepl_int.txt			Y				R						O
so_assort_int.txt			Y				R						O
so_assort_cm_int.txt			Y				R						O
planogram_assortment.csv			Y				R						O
planogram.csv			Y				R						O
equipment.csv			Y				R						O
planogram_store.csv			Y				R						O
product_hierarchy.csv			Y				R						O
product_position.csv			Y				R						O
sku_details.csv			Y				R						O

Table 15-3 (Cont.) Handling Data Files

File Name	Daily	Wkly	Qrtly	Intrday	Clustr	CS	ASO	CDT	DT	OO	IW	MBI	I/O
so_global_val_smry.txt			Y				R						O
so_global_val_detail.txt			Y				R						O
cis_store_cluster_exp.csv	Y	Y			R								O
cis_store_cluster_attr_exp.csv	Y	Y			R								O
cis_store_cluster_mem_exp.csv	Y	Y			R								O
cis_store_cluster_prop_exp.csv	Y	Y			R								O
cis_custseg_exp.csv	Y	Y				R							O
cis_custseg_attr_exp.csv	Y	Y				R							O
cis_custseg_cust_export.csv	Y	Y				R							O
cis_custseg_cat_attr_exp.csv	Y	Y				R							O
cis_custseg_store_distr_exp.csv	Y	Y				R							O
rl_shipping_cost_stg.txt	Y	Y								O			I
rl_price_ladder_stg.txt	Y	Y								O			I
rl_price_elasticity_stg.txt	Y	Y								O			I
pro_baseline_stg.txt	Y	Y								R			I
pro_custseg_clv_stg.txt	Y	Y								O			I
pro_inventory_stg.txt.gz	Y	Y								R			I
pro_lifecycle_fatigue_stg.txt	Y	Y								O			I
pro_plan_promotion_stg.txt	Y	Y								R			I

Table 15-3 (Cont.) Handling Data Files

File Name	Daily	Wkly	Qrtly	Intrday	Clustr	CS	ASO	CDT	DT	OO	IW	MBI	I/O
pro_plan_promotion_lift_stg.txt	Y	Y	Y							O			I
pro_price_cost_stg.txt	Y	Y	Y							R			I
pro_price_elasticity_stg.txt	Y	Y	Y							R			I
pro_price_ladder_stg.txt	Y	Y	Y							R			I
pro_sales_return_stg.txt	Y	Y	Y							O			I
pro_season_stg.txt	Y	Y	Y							R			I
pro_season_product_stg.txt	Y	Y	Y							R			I
pro_season_period_stg.txt	Y	Y	Y							R			I
pro_season_curr_opt_metric_stg.txt.gz	Y	Y	Y							R			I
pro_season_prod_mkdn_edt_stg.txt	Y	Y	Y							R			I
pro_seasonality_stg.txt	Y	Y	Y							R			I
pro_to_marketing_redemrate_stg.txt	Y	Y	Y							R			I
pro_to_mechanic_redem_rate_stg.txt	Y	Y	Y							R			I
pro_to_redemption_rate_stg.txt	Y	Y	Y							R			I
pro_model_dates_stg.txt	Y	Y	Y							R			I
hos_loc_hier_item_stg.txt	Y	Y	Y										I
hos_order_type_stg.txt	Y	Y	Y										I
hos_tender_media_stg.txt	Y	Y	Y										I

Table 15-3 (Cont.) Handling Data Files

File Name	Daily	Wkly	Qrtly	Intrday	Clustr	CS	ASO	CDT	DT	OO	IW	MBI	I/O
hos_day_part_stg.txt	Y	Y	Y										I
mr_class_stg.txt	Y	Y	Y										I
mr_item_class_stg.txt	Y	Y	Y										I
hos_guest_check_hist_stg.txt.gz	Y	Y	Y										I
hos_guest_chk_line_it_hist_stg.txt.gz	Y	Y	Y										I
hos_family_group_stg.txt	Y	Y	Y										I
hos_family_group_master_stg.txt	Y	Y	Y										I
hos_menu_item_stg.txt	Y	Y	Y										I
hos_menu_item_master_stg.txt	Y	Y	Y										I
hos_menu_it_dy_part_ttl_stg.txt	Y	Y	Y										I
hos_menu_item_price_stg.txt	Y	Y	Y										I
mr_rhs_item_cat_stg.txt	Y	Y	Y										I
hos_discount_stg.txt	Y	Y	Y										I
hos_service_charge_stg.txt	Y	Y	Y										I
hos_revenue_center_stg.txt	Y	Y	Y										I
hos_major_group_master_stg.txt	Y	Y	Y										I
mba_arm_srvc_loc_stg.txt	Y	Y	Y									O	I
mba_arm_run_exp.txt	Y	Y	Y									R	O
mba_arm_result_exp.txt	Y	Y	Y									R	O

Social Analytics

The Social Analytics (SA) application takes data inputs from two sources, social data and sales data. Through a joint analysis of these two data inputs, SA delivers insights regarding social metrics and the potential links between the social metrics and sales trends.

The SA process runs weekly as a batch process for the most recent social data and sales data of the past eight weeks.

Data Inputs

The SA process uses social data and weekly sales data as well as retailer's hierarchy data (including product hierarchy, location hierarchy, and calendar hierarchy).

The social data is provided as a data feed by Oracle Social Relationship Management (SRM) for a specific retailer; the retailer must have a license for SRM in order to be able to implement and use SA.

The SRM platform is a web-based application that helps retailers analyze the content from different social outlets and monitor social metrics such as trending topics, trending sources, number of online references to a particular brand, color, and social sentiments.

The social data feed is provided based on the topics and filters that a user creates in the SRM platform. The feed includes topics, keywords (such as color, fabric, style, and so on), activities (that is, conversations together with the timestamp, location, and source) in different social outlets (for example, tweet messages), and sentiments. This data is stored in the ORASE schema and is updated as part of the weekly batch process.

To learn more about how to create topics and filters, see the Oracle Social Relationship Management documentation.

Social data, sales data, and hierarchy data are all pulled from the ORASE database.

Table 16–1 Data Input

Data Input	Source Database	Source Table
Topics	ORASE	SA_TOPIC_DEFINITION
Keywords	ORASE	SA_INDICATOR_KEYWORD_FILTER
Activities	ORASE	SA_SDF_MESSAGE
Sentiments	ORASE	SA_SDF_SNIPPET
Weekly Sales	ORASE	RSE_SLS_PH_LC_WK_A

Table 16–1 (Cont.) Data Input

Data Input	Source Database	Source Table
Product Hierarchy	ORASE	RSE_PROD_HIER_TC, RSE_PROD_SRC_XREF
Location Hierarchy	ORASE	RSE_LOC_SRC_XREF
Calendar Hierarchy	ORASE	RSE_CAL_SRC_XREF

Science Innovation Workbench

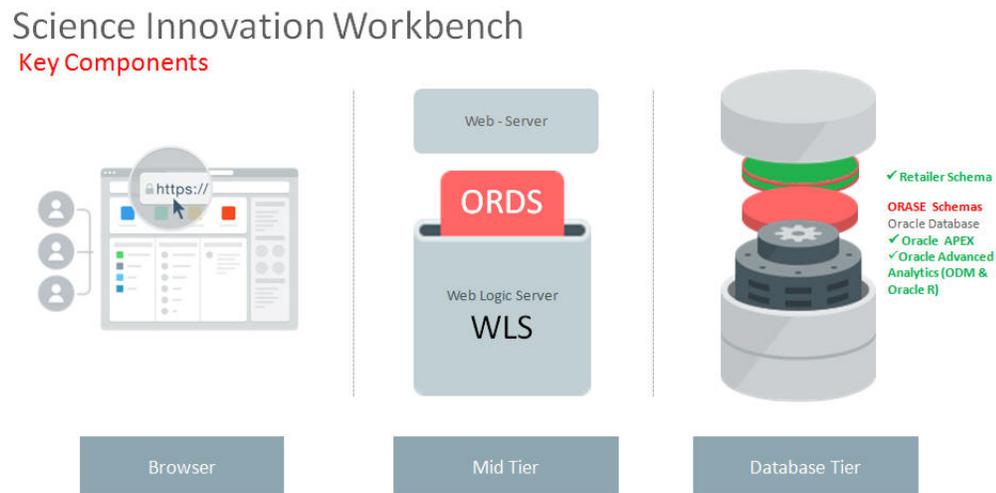
Science Innovation Workbench is a workspace that provides read-only access to ORASE data objects and clean data by using Oracle APEX. This extension is a workspace for advanced analytics users to add new implementations by using Oracle Advanced Analytic (Oracle R/ODM) algorithms that are implemented as SQL/PLSQL functions. This chapter provides features, examples, and implementation details that are available in Science Innovation Workbench in ORASE.

The key features available in Science Innovation Workbench are:

- **ORASE Schema as a Service**, in which a user can use ORASE cleansed, read-only data, upload and combine data to gain insights, and upload retail application data to further analyze and mine data.
- **Advanced Analytics**, allows a user to use Oracle Machine Learning and Advanced Analytic algorithm that are implemented as SQL/PLSQL functions.
- **Visualize Data**, enables a user to explore data, develop and visualize data using charts, and review reports using the browser.
- **RESTful Web Service**, allows a user to declaratively create and edit RESTful service definitions using SQL Query.
- **SQLWorkshop**, enables a user to view and manage database objects.

Components

The key components of Science Innovation Workbench are Retailer Workspace Schema, Oracle APEX, and Oracle Advanced Analytics (Oracle Data Mining and Oracle R Enterprise).

Figure 17–1 Science Innovation workbench Key Components

Retailer Workspace Schema

Science Innovation Workbench provides retailer with a logical work area (Retailer Workspace) that is associated with a predefined database schema. The schema(s) store the database objects and provide read access to an existing retailer's ORASE application data objects. The retailer workspace schema is an independent schema for a retailer to use and these data objects are owned by retailer.

Oracle APEX

This section describes Oracle APEX, a web-based software development environment.

Workspace

A workspace <RETAILER_WORKSPACE> is a predefined workspace for the retailer where workspace users can create database objects and applications. The workspace has privileges to the allocated <RETAILER_WORKSPACE_SCHEMA> database schema.

Users and Roles

Science Innovation Workbench has two types of users: application developers and workspace administrators, and they can be created using Oracle APEX.

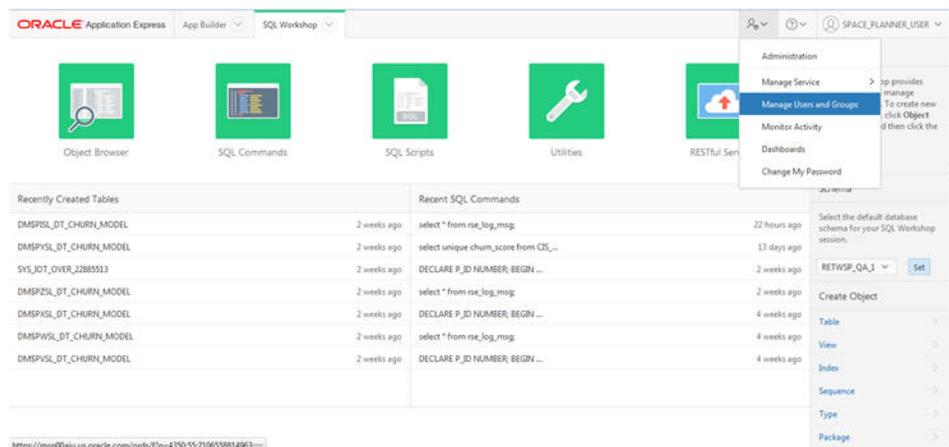
Science Innovation Workbench Administrators

The workspace administrator role is already created for the retailer; the administrator can create and edit developer accounts, manage groups, and manage development services.

Science Innovation Workbench Developer

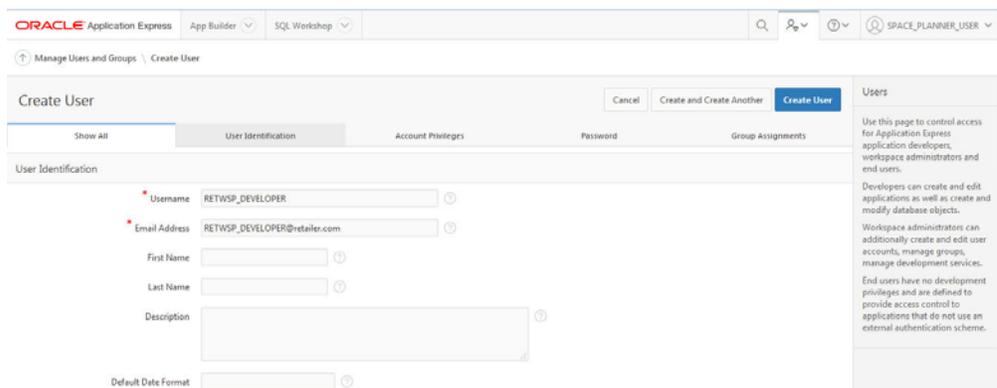
Workspace administrators can create workbench developers by selecting Manage Users and Groups. Developers can create and modify applications and browse database objects in an allocated workspace and schema. The retailer workspace schema has privileges required by Oracle Data Mining and Oracle R Enterprise for executing analytic models.

Figure 17-2 Manage Users and Groups



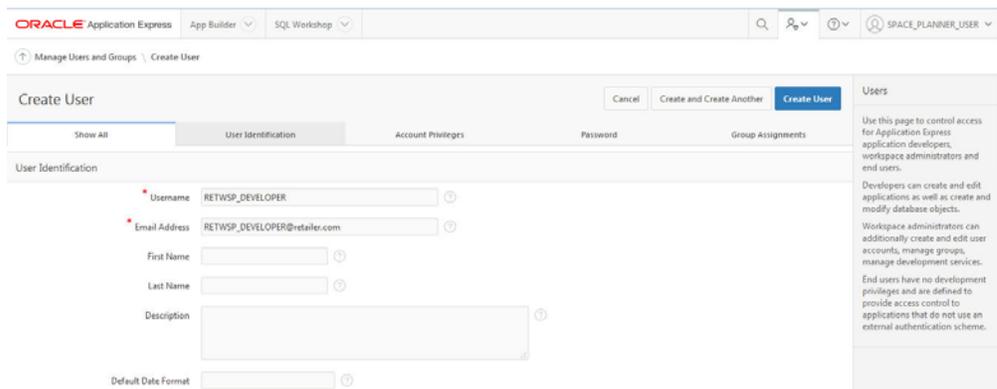
1. Create User. This area displays the Oracle APEX Create User screen that can be used to create a new Developer account and assign a RESTful Service group. Note that this user must also be created in Identity Management with the same username and password.

Figure 17-3 Create User



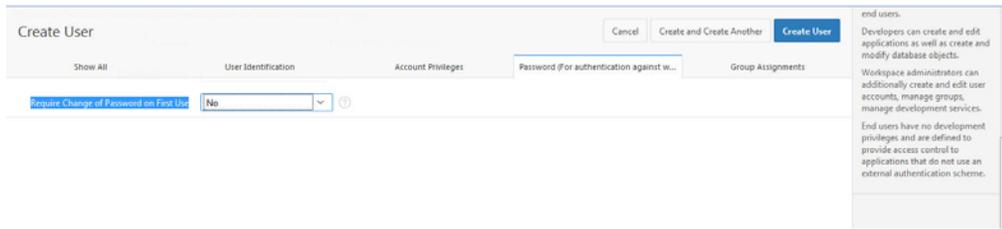
2. Assign Group RESTful Service.

Figure 17-4 Assign Group



3. Set Require Change of Password on First Use to No.

Figure 17–5 Require Change of Password



SQL Workshop

The SQL Workshop provides tools to view and manage database objects. To create new database objects, click **Object Browser** and then click **Create**.

Figure 17–6 SQL Workshop Object Browser: Reading Database Objects

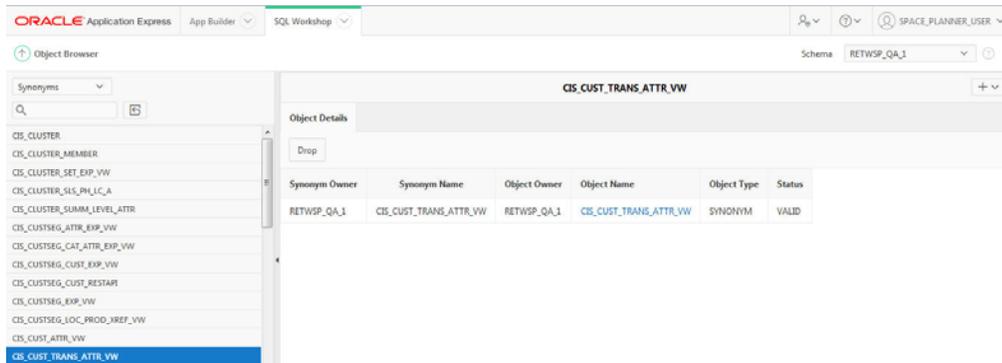


Figure 17–7 SQL Workshop Create Object

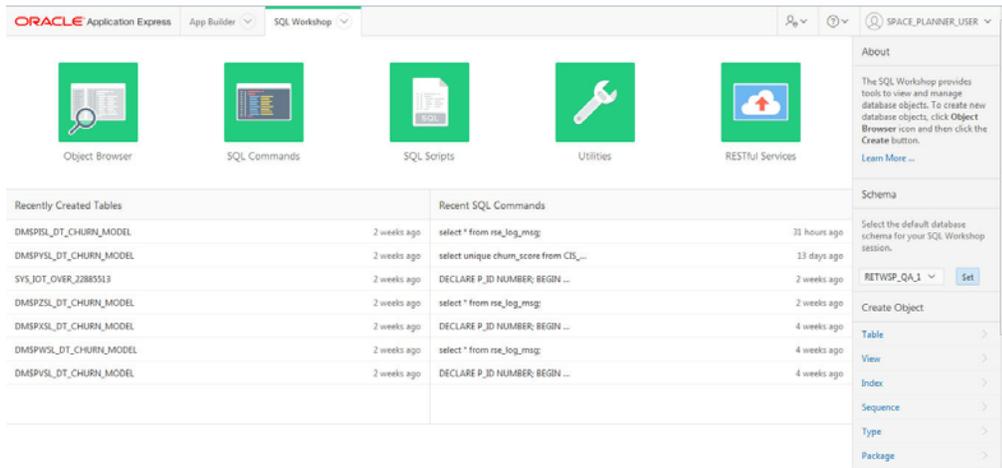


Figure 17–8 SQL Workshop SQL Command: Executing Ad Hoc SQLs

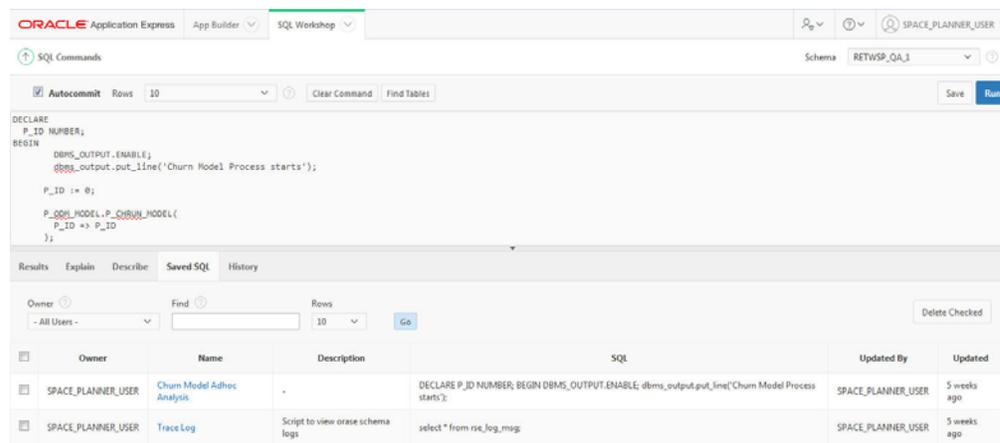
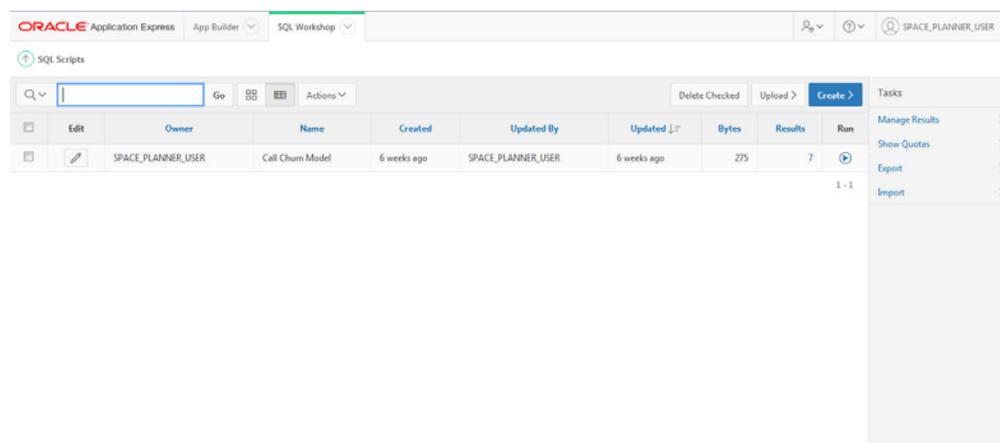


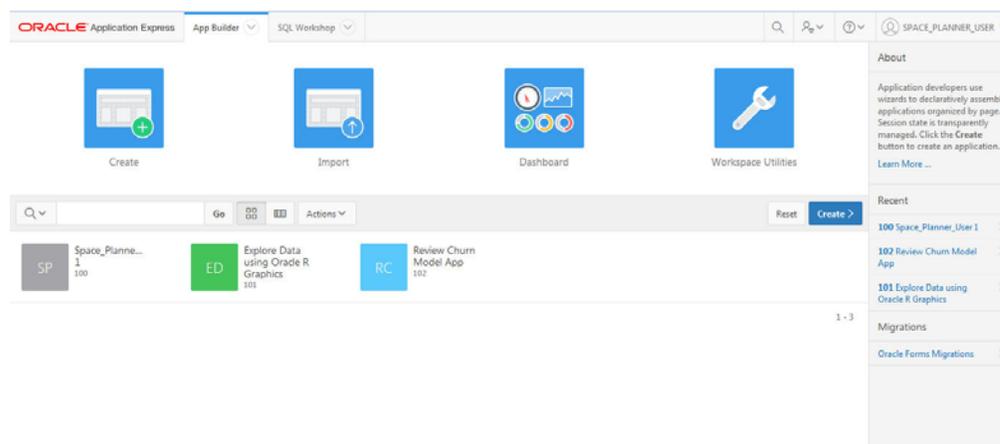
Figure 17–9 SQL Workshop SQL Scripts: Uploading, Executing, and Running



Application Builder

Application developers use wizards to declaratively assemble applications organized by page. The session state is transparently managed. Click **Create** to create an application.

Figure 17–10 Application Builder



Oracle Advanced Analytics

Oracle Advanced Analytics has two components in Oracle Database Enterprise Edition, Oracle Data Mining and Oracle R Enterprise. Science Innovation Workbench assigns privileges required by Oracle Data Mining and Oracle R Enterprise for execution.

Oracle Data Mining offers a comprehensive set of in-database algorithms for performing a variety of mining tasks, such as classification, regression, anomaly detection, feature extraction, clustering, and market basket analysis.

Oracle R Enterprise, integrates R, the open-source statistical environment, with Oracle Database.

Figure 17–11 Machine Learning Algorithms in Database

Classification <ul style="list-style-type: none"> • Decision Tree • Logistic Regression • Naïve Bayes • Support Vector Machine • RandomForest 	Clustering <ul style="list-style-type: none"> • Hierarchical k-Means • Orthogonal Partitioning Clustering 	Market Basket Analysis <ul style="list-style-type: none"> • Apriori – Association Rules
Regression <ul style="list-style-type: none"> • Linear Model • Generalized Linear Model • Multi-Layer Neural Networks • Stepwise Linear Regression • Support Vector Machine 	Attribute Importance <ul style="list-style-type: none"> • Minimum Description Length 	Feature Extraction <ul style="list-style-type: none"> • Nonnegative Matrix Factorization • Principal Component Analysis • Singular Value Decomposition
	Anomaly Detection <ul style="list-style-type: none"> • 1 Class Support Vector Machine 	Time Series <ul style="list-style-type: none"> • Single Exponential Smoothing • Double Exponential Smoothing

Oracle Data Mining

Oracle Data Mining can be used to build and deploy predictive and descriptive data mining applications, add intelligent capabilities to existing applications, and generate predictive queries for data exploration.

The Oracle Data Mining developers guide, samples, and tutorials are available at the following websites.

Oracle Data Mining Developer's Guide

http://www.oracle.com/pls/db121/vbook_subject?subject=dma

Data Mining Concepts

<https://docs.oracle.com/database/121/DMCON/toc.htm>

Oracle Data Mining Sample Programs

<http://www.oracle.com/technetwork/database/options/advanced-analytics/odm/odm-samples-194497.html>

Samples can be downloaded `odm12csampleprograms-2184025.7z`

How to Invoke Oracle Data Mining

The following scripts show how to invoke Oracle Data Mining Classification script that creates a classification model using the Decision Tree algorithm.

```
CREATE OR REPLACE PACKAGE BODY pkg_odm_model
```

```

AS
  -- DESCRIPTION - This script creates a classification model using the Decision
Tree algorithm.
PROCEDURE proc_churn_model(
  p_id NUMBER)
IS
BEGIN
  DECLARE ----- start drop RETWSP_CUST_CHURN_MODEL
    not_found EXCEPTION;
    PRAGMA EXCEPTION_INIT(not_found, -40203);
  BEGIN
    dbms_output.put_line('Start Drop RETWSP_CUST_CHURN_MODEL Tables');
    DBMS_DATA_MINING.DROP_MODEL('RETWSP_CUST_CHURN_MODEL');
    dbms_output.put_line('End Drop RETWSP_CUST_CHURN_MODEL Tables');
  EXCEPTION
  WHEN not_found THEN
    dbms_output.put_line('RETWSP_CUST_CHURN_MODEL not found');
  END; ----- end drop RETWSP_CUST_CHURN_MODEL
  -----
  -- CREATE A SETTINGS TABLE
  --
  -- The default classification algorithm is Naive Bayes. In order to override
  -- this, create and populate a settings table to be used as input for
  -- CREATE_MODEL.
  --
  DECLARE ----- start drop RETWSP_CUST_CHMDL_SETTINGS
    not_found EXCEPTION;
    PRAGMA EXCEPTION_INIT(not_found, -40203);
  BEGIN
    dbms_output.put_line('Start Drop RETWSP_CUST_CHMDL_SETTINGS Tables');
    EXECUTE IMMEDIATE 'DROP TABLE RETWSP_CUST_CHMDL_SETTINGS';
    dbms_output.put_line('End Drop RETWSP_CUST_CHMDL_SETTINGS Tables');
  EXCEPTION
  WHEN not_found THEN
    dbms_output.put_line('RETWSP_CUST_CHMDL_SETTINGS not found');
  END; ----- end drop RETWSP_CUST_CHMDL_SETTINGS
  DECLARE ----- start drop RETWSP_CUST_CHMDL_COST
    not_found EXCEPTION;
    PRAGMA EXCEPTION_INIT(not_found, -40203);
  BEGIN
    dbms_output.put_line('Start Drop RETWSP_CUST_CHMDL_COST Tables');
    EXECUTE IMMEDIATE 'DROP TABLE RETWSP_CUST_CHMDL_COST';
    dbms_output.put_line('End Drop RETWSP_CUST_CHMDL_COST Tables');
  EXCEPTION
  WHEN not_found THEN
    dbms_output.put_line('RETWSP_CUST_CHMDL_COST not found');
  END; ----- end drop RETWSP_CUST_CHMDL_COST
  DECLARE ----- start create table RETWSP_CUST_CHMDL_SETTINGS
    already_exists EXCEPTION;
    PRAGMA EXCEPTION_INIT(already_exists, -00955);
  BEGIN
    dbms_output.put_line('Start Create RETWSP_CUST_CHMDL_SETTINGS Tables');
    EXECUTE IMMEDIATE 'CREATE TABLE RETWSP_CUST_CHMDL_SETTINGS (
setting_name VARCHAR2(30),
setting_value VARCHAR2(4000))';
    dbms_output.put_line('End Create RETWSP_CUST_CHMDL_SETTINGS Tables');
  EXCEPTION
  WHEN already_exists THEN
    dbms_output.put_line('Exception not found');
  END; ----- end create table RETWSP_CUST_CHMDL_SETTINGS

```

```

DECLARE ----- Create RETWSP_CUST_CHMDL_COST Tables begins
    already_exists EXCEPTION;
    PRAGMA EXCEPTION_INIT(already_exists, -00955);
BEGIN
    dbms_output.put_line('Start Create RETWSP_CUST_CHMDL_COST Tables');
    EXECUTE IMMEDIATE 'CREATE TABLE RETWSP_CUST_CHMDL_COST (
actual_target_value          NUMBER,
predicted_target_value      NUMBER,
cost                        NUMBER)';
    dbms_output.put_line('End Create RETWSP_CUST_CHMDL_COST Tables');
EXCEPTION
WHEN already_exists THEN
    dbms_output.put_line('RETSWP_CUST_CHMDL_COST not found');
END; ----- Create RETWSP_CUST_CHMDL_COST Tables ends
-- CREATE AND POPULATE A COST MATRIX TABLE
--
-- A cost matrix is used to influence the weighting of misclassification
-- during model creation (and scoring).
-- See Oracle Data Mining Concepts Guide for more details.
--
dbms_output.put_line('Start Insert records into RETWSP_CUST_CHMDL_COST');
DECLARE ----- sub-block begins
    already_exists EXCEPTION;
    PRAGMA EXCEPTION_INIT(already_exists, -00955);
BEGIN
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (0,0,0)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (0,1,1)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (0,2,2)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (0,3,3)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (1,0,1)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (1,1,0)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (1,2,2)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (1,3,3)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (2,0,3)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (2,1,2)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (2,2,0)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (2,3,1)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (3,0,3)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (3,1,2)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (3,2,1)';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_COST VALUES (3,3,0)';
    dbms_output.put_line('End Insert Records');
EXCEPTION
WHEN already_exists THEN
    dbms_output.put_line('RETSWP_CUST_CHMDL_COST not found');
END; ----- sub-block ends
dbms_output.put_line('End Insert records into RETWSP_CUST_CHMDL_COST');
-- Populate settings table
DECLARE ----- sub-block begins
    already_exists EXCEPTION;
    PRAGMA EXCEPTION_INIT(already_exists, -00955);
    v_stmt          VARCHAR2(4000);
    v_algo_name     VARCHAR2(100);
    v_algo_decision_tree VARCHAR2(100);
BEGIN
    dbms_output.put_line('Start Populate settings table' || dbms_data_mining.algo_
name);
    dbms_output.put_line('Start Populate settings table' || dbms_data_mining.algo_
decision_tree);
    v_algo_name      := dbms_data_mining.algo_name;

```

```

    v_algo_decision_tree := dbms_data_mining.algo_decision_tree;
    v_stmt                := 'INSERT INTO RETWSP_CUST_CHMDL_SETTINGS (setting_name,
setting_value) VALUES ('' || v_algo_name || ''','' || v_algo_decision_tree ||
''')';
    dbms_output.put_line('Start Populate settings table v_stmt --' || v_stmt);
    EXECUTE IMMEDIATE v_stmt;
EXCEPTION
WHEN already_exists THEN
    dbms_output.put_line('Exception not found');
END; ----- sub-block ends
DECLARE ----- sub-block begins
    already_exists EXCEPTION;
    PRAGMA EXCEPTION_INIT(already_exists, -00955);
    v_table_name  VARCHAR2(100);
    v_matrix_cost VARCHAR2(100);
BEGIN
    v_table_name := dbms_data_mining.clas_cost_table_name;
    v_matrix_cost := 'RETWSP_CUST_CHMDL_COST';
    EXECUTE IMMEDIATE 'INSERT INTO RETWSP_CUST_CHMDL_SETTINGS (setting_name,
setting_value) VALUES' || '(' || v_table_name || ',' || v_matrix_cost ||
''')';
    dbms_output.put_line('End Populate settings table');
EXCEPTION
WHEN already_exists THEN
    dbms_output.put_line('Exception not found');
END; ----- sub-block ends
-----
-- CREATE A NEW MODEL
--
-- Build a DT model
dbms_output.put_line('Start Create Churn Model');
DBMS_DATA_MINING.CREATE_MODEL( model_name => 'RETWSP_CUST_CHURN_MODEL', mining_
function => dbms_data_mining.classification, data_table_name => 'cis_cust_attr_
vw', case_id_column_name => 'CUSTOMER_ID', target_column_name => 'CHURN_SCORE',
settings_table_name => 'RETWSP_CUST_CHMDL_SETTINGS');
dbms_output.put_line('End Create Churn Model');
-----
-- DISPLAY MODEL SIGNATURE
--
column attribute_name format a40
column attribute_type format a20
SELECT attribute_name,
       attribute_type
FROM user_mining_model_attributes
WHERE model_name = 'RETWSP_CUST_CHURN_MODEL'
ORDER BY attribute_name;
END p_chrun_model;
END pkg_odm_model;

```

Test ODM Model

```

DECLARE
    RUN_ID NUMBER;
BEGIN
    DBMS_OUTPUT.ENABLE;
    dbms_output.put_line('Churn Model Process starts');
    RUN_ID := 1001;
    pkg_odm_model.proc_churn_model( RUN_ID => RUN_ID );
    dbms_output.put_line('Churn Model Process ends');
END;

```

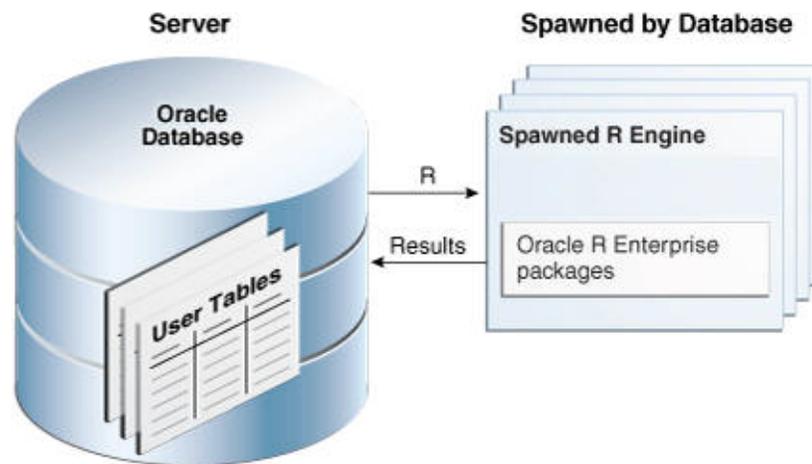
Oracle R

Science Innovation Workbench allows an advanced user to execute the embedded R engine in Oracle Database. It allows R users to off-load calculations that require either more resources such as those available to Oracle Database or database-driven data parallelism. It supports R scripts that can be embedded in SQL or PL/SQL programs.

Science Innovation Workbench via Oracle R Enterprise Database Engine enables scale for large datasets, access to tables, views in the database, the use of SQL query parallel execution, and the use of in-database statistical and data mining functionality.

Note: R Client is not available in Science Innovation Workbench and only Embedded R Executions using SQL Interface are supported.

Figure 17–12 Oracle R



Oracle R Enterprise user guide, samples, and tutorials are available at the following websites.

Oracle R Enterprise User Guide

https://docs.oracle.com/cd/E40980_01/doc.14/e39886/toc.htm

Oracle R Enterprise Sample

https://docs.oracle.com/cd/E11882_01/doc.112/e36763/appa_typicalinstall.htm#OREAD265

Tutorials on how to use Oracle R Enterprise

https://apexapps.oracle.com/pls/apex/f?p=44785:24:4786379229847::P24_CONTENT_ID,P24_PREV_PAGE,P24_PROD_SECTION_GRP_ID:8984,141,

SQL Interface for Embedded R Execution

https://docs.oracle.com/cd/E67822_01/OREUG/GUID-7F101F72-1C62-4961-BEA9-0F6E0B183F4E.htm#OREUG543

How to execute Oracle R Model

Here is an example that shows how to use R packages.

Execute R lm to fit linear models in Oracle Database. It can be used to carry out regression, analysis of variance and covariance

```
BEGIN
  sys.rqScriptDrop('FittingLinearModel');
  sys.rqScriptCreate('FittingLinearModel', 'function(dat,datastore_name) { mod <-
  lm(LOYALTY_SCORE ~ AVG_AGE + AVG_INCOME, dat) ore.save(mod,name=datastore_name,
  overwrite=TRUE) }');
END;
/

SELECT *
FROM TABLE(rqTableEval( CURSOR
  (select AVG_INCOME, AVG_AGE, LOYALTY_SCORE from cis_cust_attr_vw
  ), CURSOR
  (SELECT 1 AS "ore.connect", 'myDatastore' AS "datastore_name" FROM dual
  ), 'XML', 'FittingLinearModel' ));
```

How to Execute Oracle R Graphics

The following scripts show how to use the R Pairs function to create a correlation matrix plot between parameters in the data. The following example shows how to generate a scatter plot matrix using the R script RQG\$pairs in Oracle R.

```
drop table retwsp_cust_attr_corr;
create TABLE retwsp_cust_attr_corr as (SELECT * FROM TABLE
(
  rqTableEval(
  cursor( ((select AVG_INCOME, AVG_AGE, LOYALTY_SCORE from cis_cust_attr_vw)) )
  ), -- Input Cursor
  cursor(select 'Scatterplot Matrices' as MAIN from DUAL), -- Param Cursor
  'PNG', -- Output Definition
  'RQG$pairs' -- R Script
  )
  ));
```

How to View an Oracle R Graphics Image in APEX

The following steps show how to create an application to display graphic images generated using Oracle R graphics. The process includes creating an application to explore input data using scatter plot matrices.

1. Identify an application.

Figure 17–13 Identify an Application

Create an Application

● Name

User Interface **Desktop**

Schema **RETWSP_DEMO_1** ⓘ

Name ⓘ

Application **101** ⓘ

Theme ^ ⓘ

Theme Style v ⓘ

< Cancel Create Application **Next >**

2. Add a page to the application.

Figure 17–14 Add Page

Create an Application

● Pages

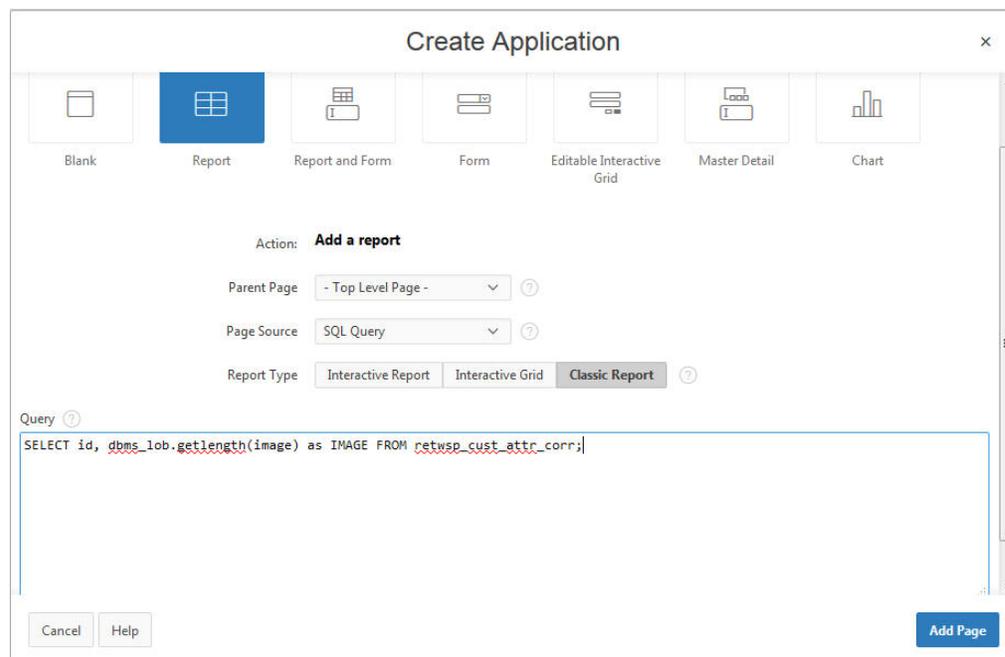
Page	Name	Type	Page Mode	Source Type	Source	Parent Page
	1	Home	Blank	Normal	-	-

[Add Page](#)

< Cancel Create Application **Next >**

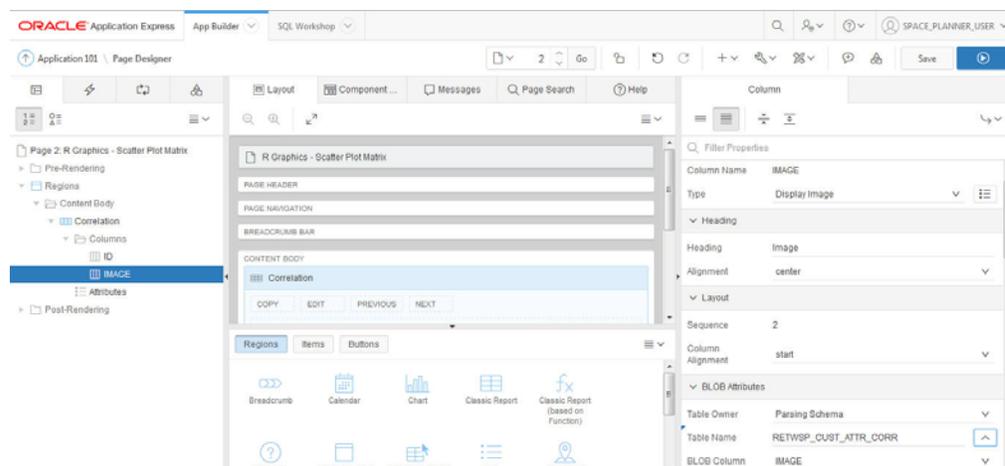
3. Select the report and page source as an SQL Query and report type as a Classic Report. Input Query is `SELECT id, dbms_lob.getlength(image) as IMAGE FROM retwsp_cust_attr_corr;`

Figure 17–15 Add Report



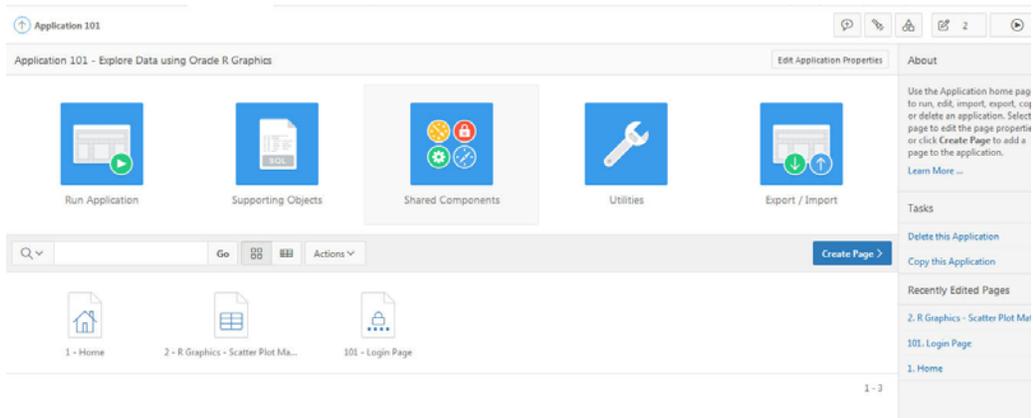
4. Edit Page in Page Designer. Select Content Body, Scatter Plot Matrix, Columns - IMAGE, and type as Display Image.

Figure 17–16 Edit Page



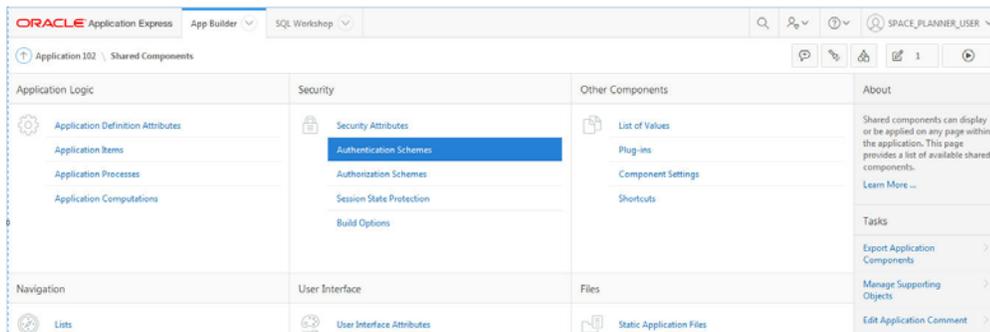
5. Set the following BLOB Attributes and click **Save**.
 - Set Table Name as retwsp_cust_attr_corr.
 - Set BLOB column as IMAGE.
 - Set Primary Key Column 1 as ID.
6. Select Shared Components.

Figure 17–17 Shared Components



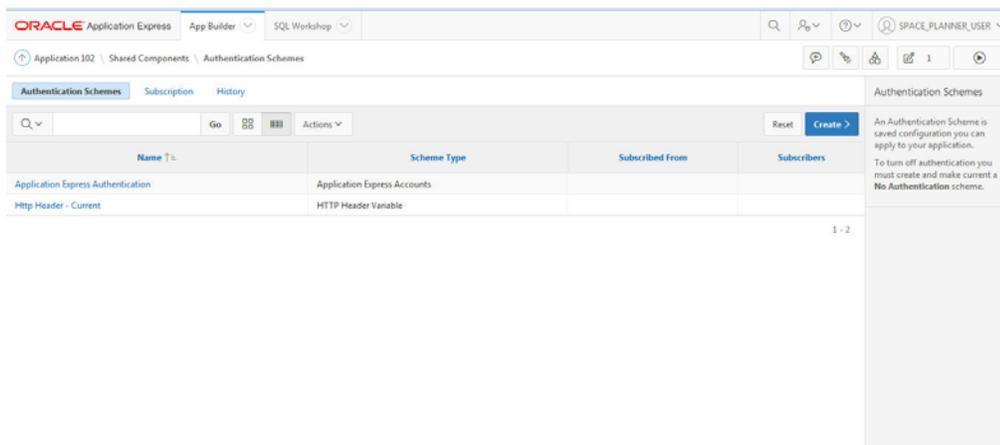
7. Select Authentication Scheme.

Figure 17–18 Authentication Scheme

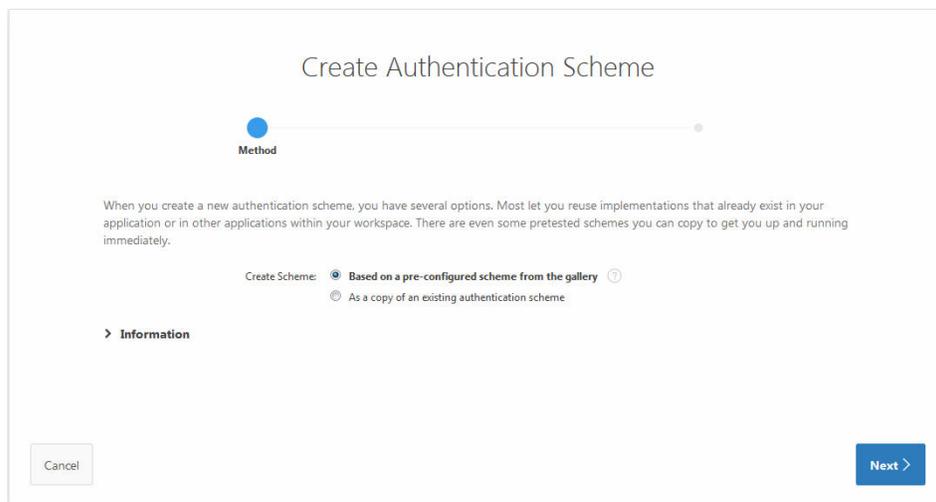


8. Ensure that HTTP Header is selected and is marked as Current. This will ensure that the Login screen is not displayed while the application is running.

Figure 17–19 HTTP Header

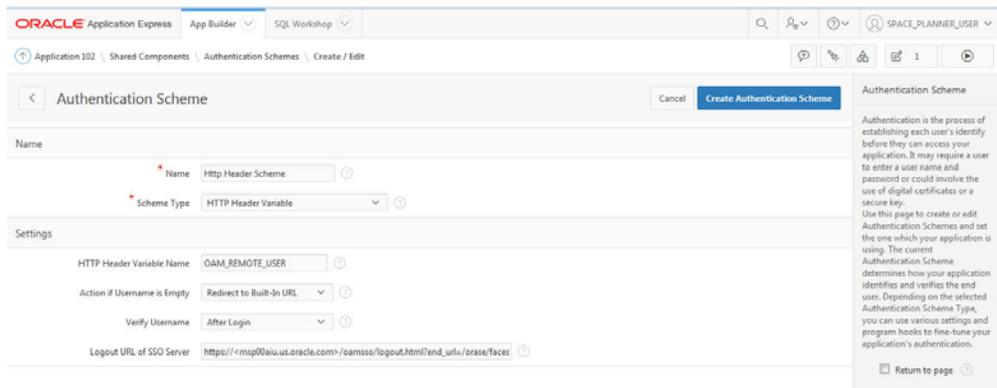


9. If only Application Express Authentication is present, select **Create a New Scheme** to enable Single Sign On. While creating the scheme, select **Based in a pre-configured scheme from the gallery**.

Figure 17–20 Pre-Configured Scheme

10. Set the following values:

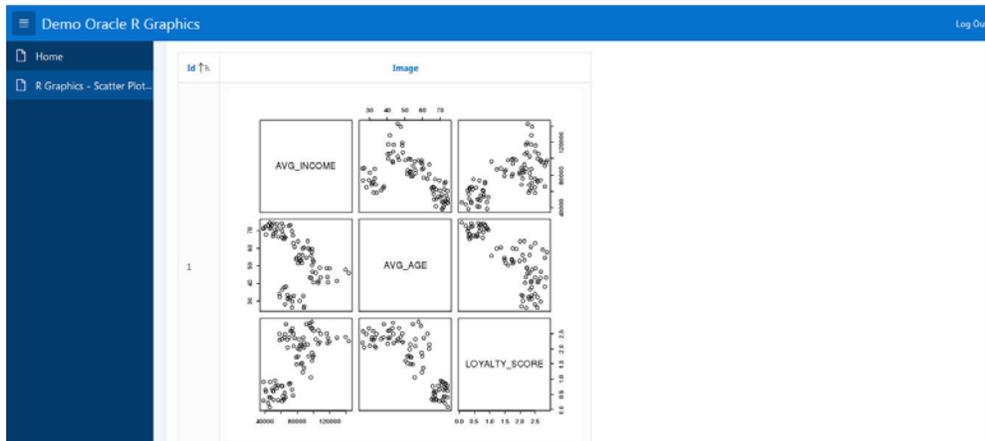
- Select Scheme Type - HTTP Header Variable.
- Set HTTP Header Variable Name - OAM_REMOTE_USER will be selected.
- Select action if username is Empty - Redirect to Built in URL
- Select verify Username - After login.
- Change Logout Url of SSO Server to `https://<server>/oamssso/logout.html?end_url=/orase/faces/home`, Where `<server>` is Host URL

Figure 17–21 Scheme Values

11. Click **Current Authentication Scheme**.

12. Run the application. This opens a new window with the application. If the Login window appears, then make sure the authentication scheme is set to the HTTP Header or contact administrator to ensure that the security scheme is handled appropriately.

Figure 17–22 Application Window



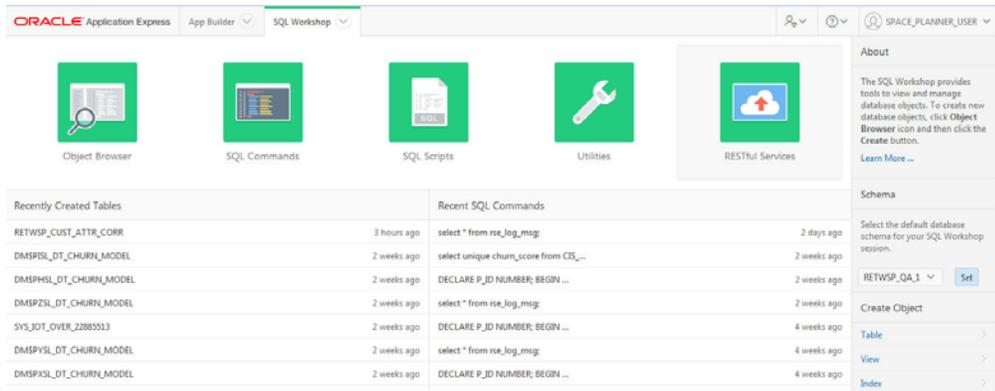
ORASE Restful Service

RESTful Services allow for the declarative specification of RESTful access to the database. They are created by configuring a set of Uniform Resource Identifiers (URIs) to a SQL query or anonymous PL/SQL block. The set of URIs is identified by a URI template.

To create a RESTful service:

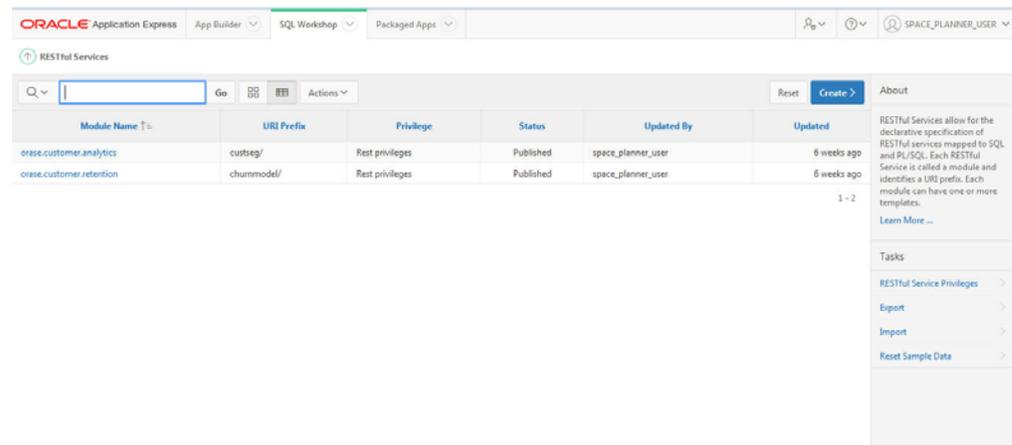
1. Select SQL Workshop ' RESTful Services.

Figure 17–23 RESTful Service



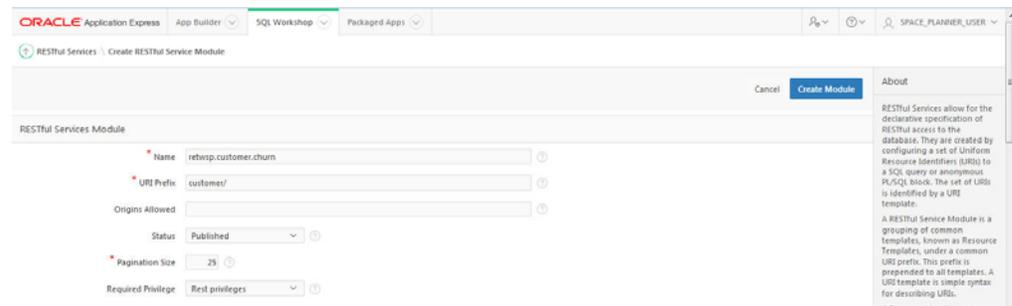
2. On selecting RESTful Services, you see the option to create a RESTful Service.

Figure 17–24 Create RESTful Service



- A RESTful Service Module is a grouping of common templates, known as Resource Templates, under a common URI prefix. This prefix is prepended to all templates.

Figure 17–25 Create Module



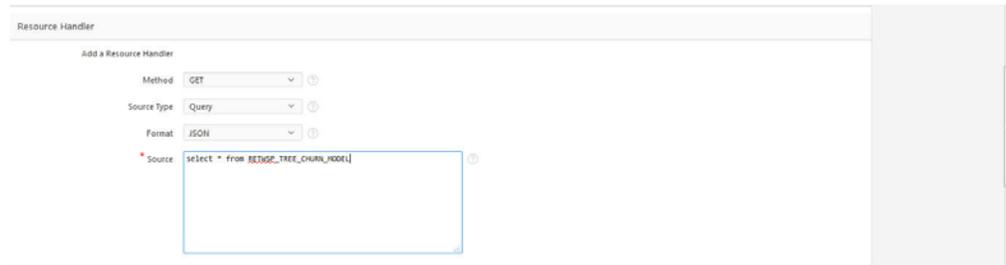
- A URI template is a simple syntax for describing URIs. You populate the required fields as shown in Figure 17–26 and make sure to set the required privileges as Rest Privileges.

Figure 17–26 Resource Template



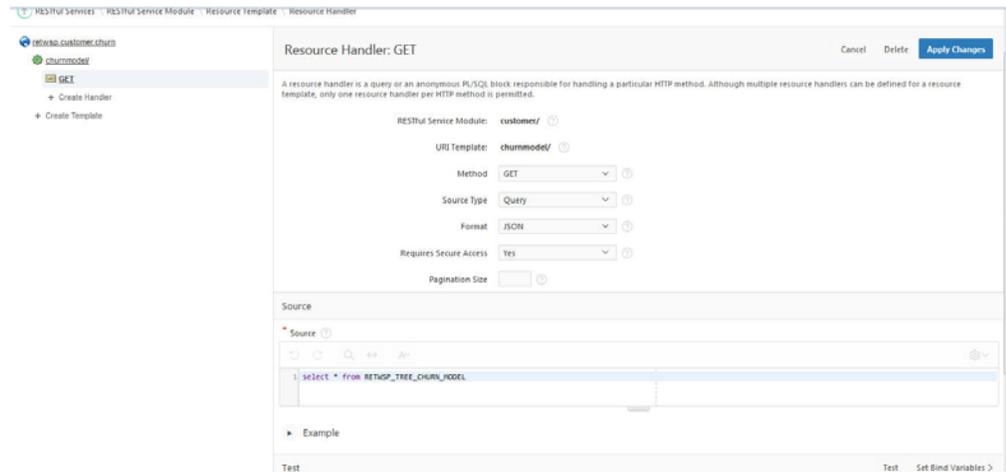
- A Resource Handler is a query or an anonymous PL/SQL block responsible for handling a particular HTTP method. Multiple handlers can be defined for a Resource Template; however, only one handler per HTTP method is permitted. You can select method, source type, format, and SQL Query to read data from the schema.

Figure 17–27 Resource Handler



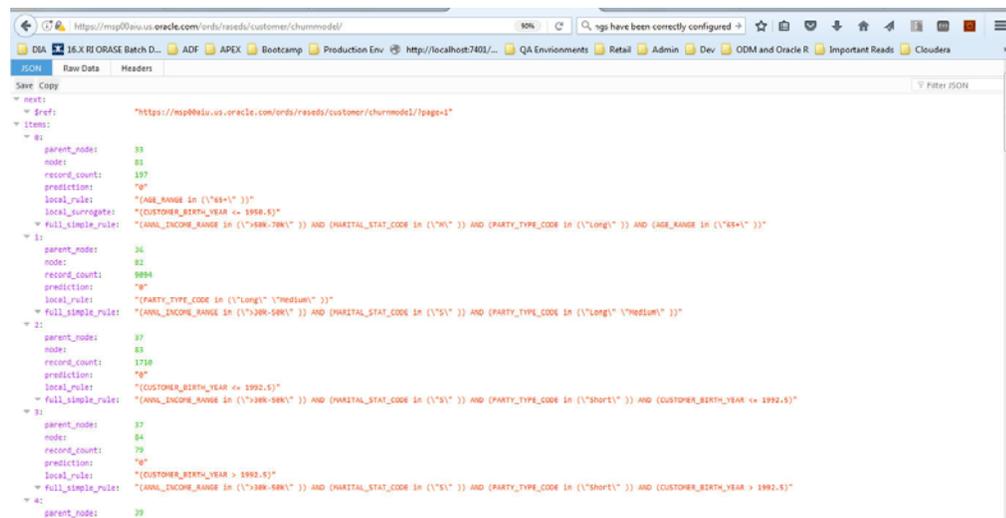
6. To test the rest API, click **Test Button** (see Test at the bottom right corner.)

Figure 17–28 Test



7. The page refreshes and displays data in JSON format.

Figure 17–29 RESTful Refresh

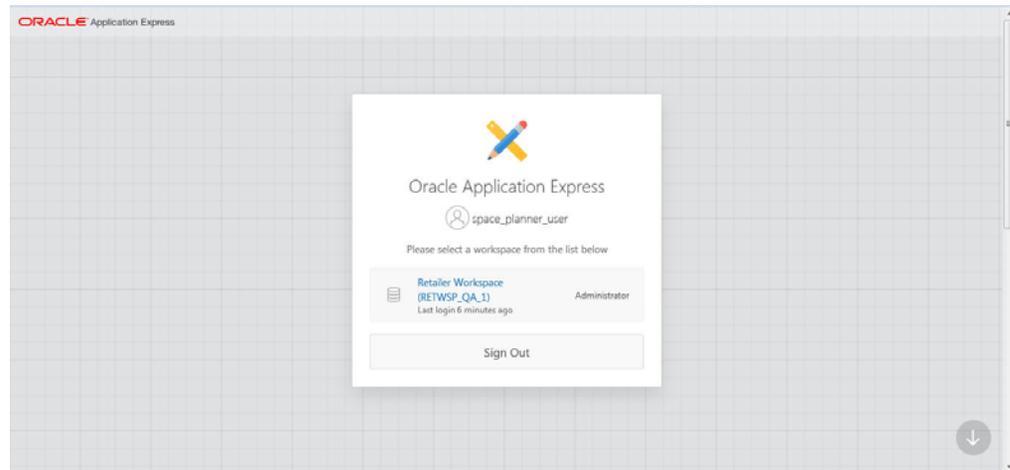


Troubleshooting

Science Innovation Workbench can also be used for monitoring ORASE logs to troubleshoot issues in an ORASE process due to an error in a batch or business process.

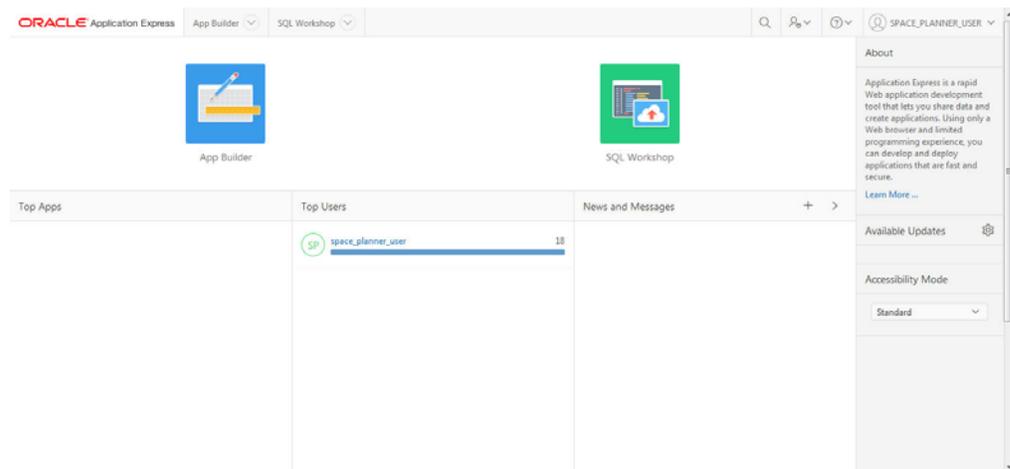
1. Click **Science Innovation Workbench**.
2. Select **Retailer Workspace**.

Figure 17–30 Retailer Workspace



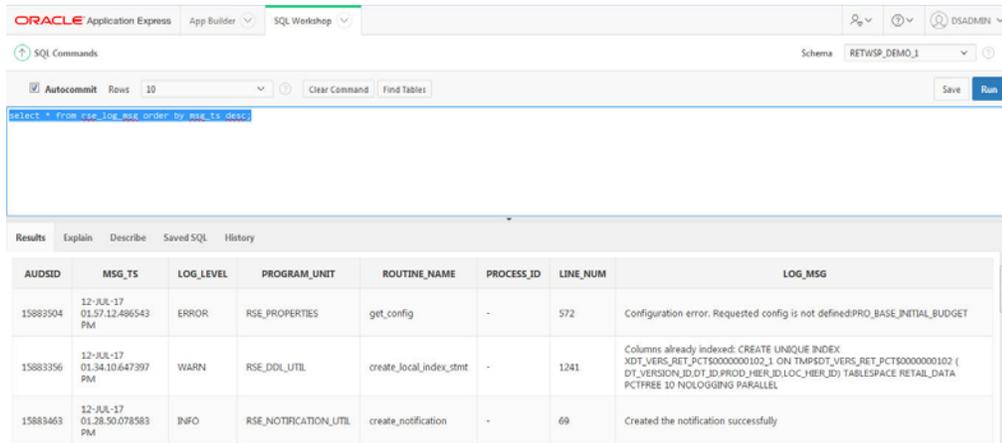
3. Select **SQL Workshop**.

Figure 17–31 SQL Workspace



4. Select **SQL Command**.
 - a. Enter SQL `select * from rse_log_msg order by msg_ts desc;`
 - b. Select run and view the log message in the bottom panel.
 - c. You can describe objects using `describe rse_log_msg;`

Figure 17–32 SQL Command



DBMS Scheduler

Science Innovation Workbench can be used for data mining tasks that are long running and have to rely on database scheduling functions and procedures that can be called from any PL/SQL program.

Any jobs created in the Retailer Workspace using Science Innovation Workbench must be created under Job Class RETAILER_WORKSPACE_JOBS. Jobs created in a default job class or any other job class other than RETAILER_WORKSPACE_JOBS can be disabled by an Oracle Administrator while managing resources.

```

BEGIN
DBMS_SCHEDULER.CREATE_JOB (
  job_name          => 'retwsp_churn_model',
  job_type          => 'STORED_PROCEDURE',
  job_action        => 'retwsp.pkg_customer_analytics.proc_churn_model',
  start_date        => '01-JAN-17 07.00.00 PM US/Pacific',
  repeat_interval   => 'FREQ=YEARLY; BYDATE=0331,0630,0930,1231; ',
  end_date          => '31-DEC-17 07.00.00 PM US/Pacific',
  job_class         => 'RETAILER_WORKSPACE_JOBS',
  comments          => 'Retailer workspace churn model job');
END;
/
    
```

ORASE Schema Objects

The following ORASE database objects are available for the advanced analyst to use.

Table 17–1 ORASE Schema Objects

Table Name	Description
rse_cal_hier	This table is used to hold all calendar hierarchies. Examples are the normal calendar hierarchy, and can also contain an alternate hierarchy for the fiscal calendar.
rse_prod_hier	This table is used to hold all product hierarchies. Examples are the normal product hierarchy, and can also contain an alternate category hierarchy.
rse_loc_hier	This table is used to hold all location hierarchies. Examples are the normal organizational hierarchy, and can also contain an alternate hierarchy for trade areas.

Table 17-1 (Cont.) ORASE Schema Objects

Table Name	Description
rse_prod_loc_status	Status of the item at this location for this time frame. A-Active; I-Inactive; C-Discontinued; D-Deleted
rse_ret_lc_wk_a	This table contains aggregate sales data for the dimensions of a location and a week.
rse_sls_lc_wk_a	This table contains aggregate sales data for the dimensions of a location and a week.
rse_sls_pr_lc_cs_wk	This table contains aggregate sales data for a Product, Location, Customer Segment and Week. The SLS_PR columns represent the metrics for that week that were on promotion, while the other metrics represent the sales metrics while the item was not on promotion.
rse_sls_pr_wk_a	This table contains aggregate sales data for the dimensions of a product and a week.
rse_sls_ph_lc_wk_a	This table contains aggregate sales transaction data for different product hierarchy/levels, at the store location/week dimension.
rse_sls_pr_lc_wk	This table contains aggregate sales data for a Product, Location, and Week. The SLS_PR columns represent the metrics for that week that were on promotion, while the other metrics represent the sales metrics while the item was not on promotion.
rse_sls_txn	This table contains sales transaction data.
rse_prod_attr	This is the table that holds product attributes.
rse_prod_attr_grp	This is the table used to load the associations of CM Groups to product attributes.
This is the table used to load the associations of CM Groups to product attributes and its values	rse_prod_attr_grp_value
rse_prod_attr_grp_value_map	This is the table used to load the associations of CM Groups to product attributes, group values and actual product attribute values
rse_like_loc	This is the table used to load the like stores for CMGroup or Category.
rse_hier_level	This table defines the various levels for all the hierarchies.
rse_hier_type	This table defines the available hierarchies for use within the RSE applications.
rse_fake_cust	Table for specifying customers who are considered as fake customers. A fake customer is a customer who purchases too many transactions to be considered a single customer. Examples are generic store cards.
rse_loc_src_xref	This table contains integration ID information that enables interaction with other systems, using IDs that other systems can accommodate for the Location Hierarchy.
rse_prod_src_xref	This table contains integration ID information that enables interaction with other systems, using IDs that other systems can accommodate for the Product Hierarchy.
rse_log_msg	This table contains messages logged while database, batch or business processing.
w_party_per_d	This table contains customer data and its attribute.

Table 17-1 (Cont.) ORASE Schema Objects

Table Name	Description
cis_cust_attr_vw	Customer Attributes - This view provides basic customer attributes.
cis_cust_trans_attr_vw	Customer Transaction Attributes - This view provides attributes for customer transactions.
cis_cust_trans_ph_attr_vw	Customer Transaction Attributes - This view provides product attributes for customer transactions.
cis_custseg_attr_exp_vw	This view provides an export of the attributes that define a segment.
cis_custseg_cat_attr_exp_vw	This view provides an export of the product attributes that define a segment.
cis_custseg_cust_exp_vw	This view provides the members for a an exportable set of segments.
cis_custseg_exp_vw	This view provides an exportable set of clusters for customer segmentation.
cis_sls_cust_cal_a	This table contains aggregate customer sales data for a configurable level of the calendar hierarchy. The table is to be partitioned by Calendar, and is also be suitable for sub partitioning by Customer using a Hash Partition strategy, so that subsequent uses can operate within the confines of a given Hash partition.
cis_sls_cust_ph_cal_a	This table contains aggregate customer sales data for a configurable level of the calendar hierarchy, for a selected set of product hierarchies. The table is to be partitioned by Calendar, and is also be suitable for sub partitioning by Customer using a Hash Partition strategy, so that subsequent uses can operate within the confines of a given Hash partition.
cis_sls_ph_a	This table contains aggregate sales data for all product hierarchy members of a configured hierarchy type and level. This can be used to identify the Top Categories for use by things like Customer Segmentation.
cis_cluster_set_exp_vw	This view provides an exportable set of clusters to send to Cat Man.
cis_store_cluster_exp_vw	This view provides an exportable set of clusters for stores.
cis_store_cluster_mem_exp_vw	This view provides the members for a an exportable set of segments.
cis_store_cluster_prop_exp_vw	This view provides an exportable set of clusters for stores.
cis_cluster_sls_ph_lc_a	This table contains calendar level aggregates for the various clusters. The table is to be partitioned by Calendar.
cis_cluster_summ_level_attr	These are metrics generated at cluster/attribute/business object level. There are metrics that are generated at that level such as centroid.
cis_prod_attr_loc_share	This table contains aggregate sales data for product attribute values as well as the share that these values are with respect to the total sales for the product hierarchy, calendar, location. The share value is a configurable value, which can either be based on sales units, sales amount or profit amount.

Glossary of Acronyms

AC

Advanced Clustering, also known as CIS.

ASO

Oracle Retail Assortment and Space Optimization.

BI

Business Intelligence.

CDT

Customer Decision Tree.

DB

Database.

DT

Demand Transference.

MBA

Market Basket Analysis.

MDS

Oracle MetaData Services.

ORASE

Oracle Retail Advanced Science Engine, also known as RSE. It contains CDT, DT, AC, ASO, and MBA.

POG Hierarchy

Defined by three levels: POG department, POG category, and POG subcategory. The POG hierarchy is used to organize POGs within POG sets. For example, a leaf to root path in the POG hierarchy: Grocery -> Snacks -> Crackers.

POG Node

A leaf node (POG subcategory) within a POG set.

POG Set

Historical POGs in the same POG subcategory and with the same seasonal attribute. An ASO term.

RI

Retail Insights, formerly known as Retail Analytics.

RADM

Retail Analytics Data Model, also known as RA Schema.

RCM

Oracle Retail Category Management.

RDF

Oracle Retail Demand Forecasting.

Seasonal Attribute

Refers to a specific year independent time period for an APO assortment and a POG set. Examples include Spring, holiday, back to school, year-round. (Also, Season Attribute.)