

Oracle® Fusion Middleware

Developing Applications with Data Service Integrator

12c (12.1.3)

E47942-01

May 2014

Oracle Fusion Middleware Developing Applications with Data Service Integrator, 12c (12.1.3).

E47942-01

Copyright © 2014, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

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

Contents

1 Introduction to Data Services

1.1	Concepts	1-1
1.1.1	Data in the 21st Century	1-1
1.1.2	Data Access Integration Architecture	1-2
1.1.3	Oracle Data Service Integrator: Roles and Responsibilities.....	1-3
1.2	How-to.....	1-3
1.2.1	How to Configure the Retail Dataspace Sample Application for Oracle Data Service Integrator 1-4	
1.2.1.1	Prerequisites	1-4
1.2.1.2	About WorkSpace Studio, Data Services Studio, and Eclipse	1-4
1.2.1.3	Start WorkSpace Studio	1-4
1.2.1.4	Start the Server	1-6
1.2.1.5	Deploy Your Projects	1-7
1.2.1.6	Create the Retail Dataspace Sample Web Application	1-7
1.2.1.7	See Also	1-9
1.2.2	How to Configure the Retail Dataspace Sample Application for OSDI Studio	1-9
1.2.2.1	Prerequisites	1-10
1.2.2.2	About WorkSpace Studio, Data Services Studio, and Eclipse	1-10
1.2.2.3	Start Studio	1-10
1.2.2.4	Start the Server	1-12
1.2.2.5	Deploy Your Projects	1-12
1.2.2.6	See Also	1-12
1.3	Example: How to Create Your First Data Services	1-13
1.3.1	Goal of the Tutorial.....	1-13
1.3.1.1	Requirements	1-13
1.3.1.2	Before You Begin	1-13
1.3.2	Creating a Dataspace Project.....	1-15
1.3.2.1	Set Up a Folder for Physical Data Services	1-15
1.3.3	Creating Physical Data Services	1-16
1.3.3.1	Select a Data Source.....	1-17
1.3.3.2	Schemas Directory	1-19
1.3.3.3	Publish Your Projects	1-20
1.3.4	Creating a Logical Data Service.....	1-21
1.3.4.1	Attempt To Publish Your Dataspace Project	1-22
1.3.4.2	Bottom Up or Top Down.....	1-23
1.3.4.3	Add an Operation to CUST_ORDERS_ITEMS	1-23

1.3.4.4	Building Your Query.....	1-26
1.3.4.5	Building Your FLWR Statement Graphically.....	1-27
1.3.4.6	Add a Parameter.....	1-28
1.3.4.7	Map Elements to the Return Type.....	1-30
1.3.4.8	Populating the Return Clause.....	1-30
1.3.4.9	Set Statement Scoping.....	1-33
1.3.4.10	Creating Joins - the Where Clauses.....	1-36
1.3.4.11	Associate a Parameter with a For Node.....	1-37
1.3.5	Creating, Saving, and Associating the XML Type.....	1-39
1.3.5.1	Modifying the XML Type.....	1-40
1.3.6	Testing Your Data Service Function.....	1-42
1.3.6.1	View Test Run Results.....	1-43
1.3.7	Adding Create-Update-Delete Functions to Your Data Service.....	1-45
1.3.8	Updating Your Results.....	1-47
1.3.9	Reviewing the Query Plan.....	1-48
1.3.10	Reviewing the Update Map.....	1-49
1.3.11	Archiving Your Project.....	1-50
1.3.11.1	Saving Project to a ZIP File.....	1-50
1.3.12	Summary.....	1-51
1.4	Reference.....	1-52
1.4.1	Oracle Data Service Integrator Start Menu.....	1-52
1.4.1.1	Oracle Data Service Integrator Start Menu.....	1-52
1.4.2	Data Service Types and Functions.....	1-52
1.4.3	Data Service Characteristics.....	1-53
1.4.4	Operational Characteristics.....	1-53
1.5	Related Topics.....	1-55
1.5.1	Getting the Most from the WebLogic Eclipse Plugin Framework.....	1-55
1.5.1.1	Data Services Eclipse for WebLogic.....	1-56
1.5.2	Create a Data Service with a Flat Return Type.....	1-56
1.5.2.1	Overview.....	1-57
1.5.2.2	Create a Dataspace Project.....	1-57
1.5.2.3	Create the Return Type.....	1-58
1.5.2.4	Create Physical Data Services.....	1-59
1.5.2.5	Create a Logical Data Service.....	1-59
1.5.2.6	Create the Query Map.....	1-60
1.5.2.7	See Also.....	1-62

2 Developing and Managing Dataspace Projects

2.1	Data Service File Validation During Deployment.....	2-1
2.2	How-to.....	2-1
2.2.1	How To Create, Build, Clean, and Delete Dataspace Projects.....	2-2
2.2.1.1	Creating a Dataspace Project.....	2-2
2.2.1.2	Building a Dataspace Project.....	2-2
2.2.1.3	Cleaning a Dataspace Project.....	2-2
2.2.1.4	Deleting a Dataspace Project.....	2-3
2.2.2	How to Publish, Configure, and Remove Dataspace Projects.....	2-3
2.2.2.1	Publishing Server Projects.....	2-3

2.2.2.2	Configuring Server Projects	2-4
2.2.2.3	Managing Configured Projects Through the Servers Window	2-5
2.2.2.4	Removing Dataspace Projects from a Server	2-6
2.2.3	Exporting Dataspace Projects or Project Folders	2-6
2.2.4	Exporting Dataspace Project Artifacts Using Oracle Data Service Integrator Export Wizards 2-7	
2.2.4.1	Exporting Dataspace Artifacts	2-7
2.2.4.2	Generating a Data Service Definitions and Artifacts JAR	2-8
2.2.4.3	Generating a Mediator Client JAR File.....	2-10
2.2.4.4	Generating a JAR File Containing Data Service-to-Web Service Maps.....	2-11
2.2.5	Importing a Dataspace Project.....	2-13
2.2.5.1	Importing a File-based Project.....	2-13
2.2.6	How To Handle Error Conditions in a Dataspace Project.....	2-13
2.2.7	How To Validate, Build, Export, and Package Dataspace Projects from the Command Line 2-15	
2.2.7.1	Data Service File Validation During Deployment.....	2-15
2.2.7.2	Dataspace Packaging from the Command-line	2-15
2.2.7.3	Syntax Summary.....	2-15
2.2.7.4	Command-Line Ant Build Targets	2-16
2.2.7.5	Command-line Examples using Ant and Java	2-19
2.3	Reference	2-20
2.3.1	Dataspace Projects Cheatsheet.....	2-20
2.3.2	Setting Eclipse for WebLogic Initialization Parameters (.ini)	2-21
2.4	Related Topics	2-22

3 Creating and Updating Physical Data Services

3.1	Concepts.....	3-1
3.1.1	Creating Physical Data Services by Importing Source Metadata	3-1
3.1.1.1	Source View	3-2
3.1.2	Physical Data Services from Java Functions Overview	3-2
3.1.2.1	Simple Java Types and Their XQuery Counterparts.....	3-3
3.1.2.2	Physical Data Service from a Java Function - Example Code.....	3-4
3.2	How to Create Physical Data Services	3-6
3.2.1	How To Create Physical Data Services from Relational Tables and Views.....	3-6
3.2.1.1	Setting Up the Physical Data Service Creation Wizard	3-6
3.2.1.2	Setting Up the Import Wizard for Relational Objects	3-8
3.2.1.3	Selecting SQL Table and View Objects for Import	3-9
3.2.1.4	Setting Properties for New Data Service Operations	3-11
3.2.1.5	Verifying Data Service Composition	3-12
3.2.1.6	Database-specific Catalog and Schema Considerations	3-12
3.2.1.7	XML Name Conversion Considerations	3-13
3.2.2	How To Create Physical Data Services from Stored Procedures.....	3-13
3.2.2.1	Importing Stored Procedure Metadata Using the Physical Data Service Creation Wizard 3-14	
3.2.2.2	Setting Up the Physical Data Service Creation Wizard	3-14
3.2.2.3	Setting Up the Import Wizard for Relational Objects	3-15
3.2.2.4	Selecting Stored Procedure Objects for Import	3-16

3.2.2.5	Configuring Selected Stored Procedures	3-17
3.2.2.6	Stored Procedure Configuration Reference	3-20
3.2.2.7	Setting Properties for New Data Service Operations	3-21
3.2.2.8	Verifying Data Service Composition	3-23
3.2.2.9	Adding Operations to an Existing Data Service	3-23
3.2.2.10	Support for Stored Procedures in Popular Databases.....	3-24
3.2.3	How To Create Physical Data Services Based on SQL Statements.....	3-27
3.2.3.1	Setting Up the Physical Data Service Creation Wizard	3-27
3.2.3.2	Setting Up the Import Wizard for Relational Objects	3-29
3.2.3.3	Entering a SQL Statement	3-30
3.2.3.4	Setting Properties for New Library Functions	3-31
3.2.3.5	Verifying Data Service Composition	3-32
3.2.4	How To Create Physical Data Services Based on Database Functions	3-32
3.2.4.1	Setting Up the Physical Data Service Creation Wizard	3-33
3.2.4.2	Setting Up the Import Wizard for Relational Objects	3-34
3.2.4.3	Providing Database Function Details	3-35
3.2.4.4	Verifying Data Service Composition	3-36
3.2.5	How To Create a Physical Data Service from a Web Service.....	3-37
3.2.5.1	Setting Up the Physical Data Service Creation Wizard	3-37
3.2.5.2	Accessing a Web Service.....	3-39
3.2.5.3	Selecting Web Service Operations to Import.....	3-44
3.2.5.4	Setting Characteristics of Imported Web Service Operations.....	3-45
3.2.5.5	Setting the Data Service Name	3-46
3.2.6	Preparing to Create Physical Data Services From Java Functions.....	3-48
3.2.7	How To Create a Physical Data Service from a Java Function	3-48
3.2.7.1	Setting Up the Physical Data Service Creation Wizard	3-48
3.2.7.2	Accessing Java Functions	3-50
3.2.7.3	Selecting Java Functions to Import	3-52
3.2.7.4	Setting Characteristics of Imported Java Functions	3-52
3.2.7.5	Setting the Physical Data Service Name.....	3-53
3.2.7.6	See Also	3-54
3.2.8	How To Create a Physical Data Service from XML Data	3-54
3.2.8.1	Setting Up the Physical Data Service Creation Wizard	3-54
3.2.8.2	Specifying XML Data Schema and File	3-56
3.2.8.3	Setting Properties for New Library Functions	3-57
3.2.8.4	Verifying Data Service Composition	3-58
3.2.8.5	XML File Import Sample	3-58
3.2.9	How To Create a Physical Data Service from a Delimited File.....	3-59
3.2.9.1	Setting Up the Physical Data Service Creation Wizard	3-60
3.2.9.2	Specifying Delimited File Information	3-61
3.2.9.3	Setting Properties for New Library Functions	3-63
3.2.9.4	Verifying Data Service Composition	3-64
3.3	How to	3-64
3.3.1	How To Enable Optimistic Locking.....	3-64
3.3.1.1	Set the Locking Policy	3-65
3.3.1.2	Select the Locking Fields	3-65
3.3.1.3	See Also	3-66

3.3.2	How To Update Physical Data Service Metadata	3-66
3.3.2.1	Topics	3-67
3.3.2.2	Scope of Metadata Update	3-68
3.3.2.3	Important Considerations When Updating Source Metadata	3-69
3.3.2.4	Using the Update Source Metadata Wizard	3-69
3.3.2.5	Inspecting and Reverting Changes Using Local History.....	3-72
3.3.3	Creating SOAP Handlers for Imported WSDLs.....	3-72
3.3.3.1	Create a Java Class Implementing the Generic Handler Interface.....	3-72
3.3.3.2	Compile your intercept handler into a JAR file.	3-73
3.3.4	Creating XMLBean Support for Java Functions.....	3-74
3.3.4.1	Supported XMLBean Standards.....	3-74
3.3.4.2	Creating XMLBean Classes for Java Functions	3-75
3.3.4.3	See Also	3-79
3.3.5	How To Browse and Select a Schema Type	3-79
3.3.5.1	Browsing and Selecting Schema Types	3-80
3.3.5.2	Browsing Schema Types.....	3-80
3.3.5.3	Selecting a Schema Type.....	3-81
3.3.6	Physical Data Service from a Java Function - Example Code	3-84
3.3.6.1	Using a Function Returning an Array of Java Primitives.....	3-84
3.3.6.2	Processing complex types represented via XMLBeans.....	3-84
3.4	Example: XMLBeans Example Using a Metadata-rich Java Class.....	3-85
3.4.1	Java Source.....	3-85
3.4.2	Schema Definition.....	3-86
3.4.3	Data Service Function	3-87
3.5	Reference	3-88
3.5.1	Stored Procedure Configuration Reference	3-88
3.5.1.1	In Mode, Out Mode, Inout Mode	3-88
3.5.1.2	Procedure Profile	3-88
3.5.1.3	Supporting Stored Procedures with Nullable Input Parameter(s).....	3-89
3.5.2	Simple Java Types and Their XQuery Counterparts	3-90
3.6	Related Topics	3-91
3.6.1	How To Add an External Function to an Existing Physical Data Service.....	3-91
3.6.1.1	Additional Constraints	3-92

4 Designing Logical Data Services

4.1	Concepts	4-1
4.1.1	Building Logical Entity Data Services	4-1
4.1.1.1	The Benefits of Logical Services	4-1
4.1.1.2	Design View	4-2
4.1.1.3	Query Map View	4-6
4.1.1.4	Update Map View	4-7
4.1.1.5	Test View	4-9
4.1.1.6	See Also	4-10
4.1.2	Data Service Keys	4-10
4.1.2.1	Overview	4-10

Key for the CUSTOMER Table 11

4.1.2.2	Parts of a Key.....	4-11
4.1.2.3	Composite Keys	4-14
4.1.2.4	See Also	4-15
4.1.3	XML Types and Return Types	4-15
4.1.3.1	Where XML Types are Used	4-15
4.1.3.2	Where Return Types are Used.....	4-15
4.2	How-to.....	4-16
4.2.1	How To Add a Read Function	4-16
4.2.1.1	Overview	4-16
4.2.1.2	Create the Function in Eclipse for WebLogic	4-17
4.2.1.3	See Also	4-18
4.2.2	How To Add a Library Function or Procedure.....	4-18
4.2.2.1	Overview	4-18
4.2.2.2	Add the Function or Procedure.....	4-18
4.2.2.3	Test in Eclipse for WebLogic.....	4-20
4.2.3	How To Create Logical Data Service Keys	4-20
4.2.3.1	Generate a Key	4-21
4.2.3.2	Select Elements for a Key	4-22
4.2.3.3	Select a Key Schema File.....	4-24
4.2.3.4	View and Map a Key.....	4-25
4.2.3.5	See Also	4-27
4.2.4	How To Declare a Security Resource in Eclipse for WebLogic.....	4-27
4.2.4.1	Choose a Technique	4-27
4.2.4.2	Create the Security Resource	4-27
4.2.4.3	Use the Security Resource in XQuery.....	4-28
4.2.4.4	Assign Security Resources.....	4-31
4.2.4.5	Test Security	4-31
4.2.4.6	See Also	4-32
4.3	Examples	4-32
4.3.1	How to Create a Logical Data Service with a Group By Clause.....	4-32
4.3.1.1	Overview	4-33
4.3.1.2	Design the Return Type Schema	4-33
4.3.1.3	Create the Logical Data Service	4-34
4.3.1.4	Create the Group By Node.....	4-34
4.3.1.5	Create the For Node	4-36
4.3.1.6	Add an Aggregate Function.....	4-37
4.3.1.7	Test the Service	4-37
4.3.1.8	See Also	4-38
4.3.2	How To Create a Data Service with a Flat Return Type	4-38
4.3.2.1	Overview	4-39
4.3.2.2	Create a Dataspace Project	4-39
4.3.2.3	Create the Return Type.....	4-40
4.3.2.4	Create Physical Data Services.....	4-41
4.3.2.5	Create a Logical Data Service	4-42
4.3.2.6	Create the Query Map.....	4-44
4.3.2.7	See Also	4-45
4.4	Reference	4-45

4.4.1	XQuery Source of a Logical Entity Service	4-45
4.4.1.1	Source Code.....	4-45
4.4.1.2	See Also.....	4-48
4.5	Related Topics	4-48

5 Modeling Data Services Relationships

5.1	Relationship Between Data Services and Models	5-1
5.2	How to.....	5-2
5.2.1	Create Your First Data Services Model.....	5-2
5.2.1.1	Introduction	5-2
5.2.1.2	Building a Simple Data Service Relationship Model.....	5-4
5.2.1.3	Setting Relationship Properties	5-5
5.2.1.4	Configuring Navigation Functions.....	5-7
5.2.2	Work with Large Models.....	5-8
5.2.2.1	Search	5-9
5.2.2.2	Outline Mode	5-9
5.2.3	Generate a Relationship Modeler Report.....	5-9
5.2.3.1	Model Report Format.....	5-10
5.3	Reference.....	5-10
5.3.1	Relationship Modeler Options.....	5-11
5.3.1.1	Model Right-click Menu Options.....	5-11
5.3.2	Model Diagram Rules	5-12
5.3.3	Notable Relationship Modeler Properties.....	5-12
5.3.4	Relationship Models in Source View	5-13

6 Building XQueries

6.1	How To.....	6-1
6.1.1	Create a Return Type.....	6-1
6.1.1.1	Choose a Technique	6-1
6.1.1.2	Write a Return Type Schema	6-3
6.1.1.3	Generate a Schema File	6-5
6.1.1.4	See Also	6-6
6.1.2	Add a Complex Child Element to a Return Type.....	6-6
6.1.2.1	Add the Child Element Visually	6-6
6.1.2.2	Edit the XML Source	6-9
6.1.2.3	See Also	6-10
6.1.3	Check Namespaces in Return Types.....	6-11
6.1.3.1	Check Prefix Bindings.....	6-11
6.1.3.2	Edit the Namespace.....	6-12
6.1.3.3	See Also	6-13
6.1.4	Create Conditional Elements in Return Types.....	6-13
6.1.4.1	Add the Condition.....	6-13
6.1.4.2	Create the Expression	6-14
6.1.4.3	See Also	6-16
6.1.5	Add a Where Clause to a Query.....	6-16
6.1.5.1	Define the Condition.....	6-16

6.1.5.2	Join Tables with a Where Clause.....	6-17
6.1.5.3	Use an XQuery Function in a Where Clause	6-18
6.1.5.4	See Also	6-19
6.1.6	Use the XQuery Expression Editor	6-19
6.1.6.1	Overview	6-19
6.1.6.2	The fn-bea:value Function.....	6-20
6.1.6.3	See Also	6-21
6.1.7	Use the Source Editor.....	6-21
6.1.7.1	What is the Source Editor?	6-21
6.1.7.2	Searching Source.....	6-22
6.1.7.3	Navigating to Specific Functions.....	6-22
6.1.7.4	Color Coding	6-23
6.1.7.5	Code Completion.....	6-23
6.2	Reference	6-25
6.2.1	XQuery Language Version Support.....	6-25
6.2.2	Built-in XQuery Functions.....	6-25
6.3	Related Topics	6-26

7 Testing Update Procedures Using SDO Data Graphs

7.1	Key Points	7-1
7.2	Updates in Test View	7-1
A Data Graph with Old and New Data 2		
7.3	Optimistic Locking	7-2

8 Understanding Query Plans

8.1	Using Query Plan View	8-1
8.1.1	Query Plan Information and Warnings.....	8-3
8.1.2	Printing or Saving Your Query Plan.....	8-3
8.1.3	Loading a Previously Saved Query Plan.....	8-4
8.2	Analyzing a Sample Query.....	8-4
8.3	Working with a Query Plan	8-4
8.3.1	Identifying Problematic Conditions Through the Query Plan	8-4

9 Managing Update Maps

9.1	Understanding Update Maps.....	9-1
9.1.1	The Target Box	9-3
9.1.1.1	The Input Type.....	9-3
9.1.1.2	Procedure Icons.....	9-3
9.1.1.3	For Each Blocks	9-4
9.1.1.4	Update Blocks	9-4
9.1.1.5	The Return Key Block	9-7
9.1.1.6	Customization.....	9-8
9.2	Changing a Mapping.....	9-9
9.2.1	Example.....	9-10
9.3	Removing a Mapping.....	9-12
9.4	Reverting Customizations	9-13

9.5	Adding a Condition to an Update Block	9-13
9.5.1	Example.....	9-14
9.6	Editing XQuery Expressions	9-15
9.6.1	Overview.....	9-15
9.6.2	The fn-bea:value Function	9-16
9.7	Adding an Update Map Procedure.....	9-16
9.7.1	Overview.....	9-17
9.7.2	Generating Default Procedures	9-17
9.7.3	Designing Custom Procedures	9-19
9.8	Determining the Scope of a Variable.....	9-20
9.8.1	Variable Types and Scoping Rules.....	9-20
9.8.1.1	Example: an update map for a customer-orders data service.....	9-21
9.8.1.2	Example: an outside mapping to a key value	9-22
9.8.2	Updating Foreign Key Values.....	9-23
9.8.2.1	Example: coalesce-equal	9-24
9.9	Reference	9-25
9.9.1	Update Map Functions	9-25
9.10	How To	9-26
9.10.1	How To Recognize When Something is Wrong.....	9-26
9.10.1.1	Understand the Symptoms	9-26
9.10.1.2	Check the Problems Tab	9-27
9.10.1.3	Resolve Errors and Warnings	9-28
9.10.2	How To Understand Mappings with Different Data Types	9-29
9.10.2.1	Overview	9-30
9.10.2.2	Built-In Cast Functions	9-30
9.10.2.3	Custom Cast Functions.....	9-31
9.10.3	How to Cast Using a Built-in XQuery Function.....	9-31
9.10.3.1	Example.....	9-31
9.10.4	How To Cast Using a Custom XQuery Function.....	9-32
9.10.4.1	Example.....	9-32
9.10.5	How To Test an Update Map Cast.....	9-33
9.10.5.1	Example.....	9-33
9.10.6	How To Handle Disabled Procedures in Underlying Data Sources.....	9-35
9.10.6.1	Check the Data Sources	9-35
9.10.6.2	Resolve the Disabled Procedures	9-36
9.10.6.3	Add or Enable Procedures in the Underlying Data Source	9-37
9.10.6.4	Change the XML Return Type.....	9-37
9.10.7	How To Handle Non-Unique Joins.....	9-38
9.10.7.1	Understand the Join	9-38
9.10.7.2	Correct the Block Scope	9-41
9.10.7.3	Correct the Table Join.....	9-41
9.10.7.4	Enable Update Blocks and Procedures.....	9-42
9.10.7.5	Test a Non-Unique Join	9-43
9.10.8	How To Handle Non-Unique Values	9-43
9.10.8.1	Example.....	9-43
9.10.9	How To Handle Unmapped Required Values.....	9-44
9.10.9.1	Overview	9-45

9.10.9.2	Draw the Mapping	9-46
9.10.9.3	Cast a Constant	9-46
9.11	Testing Update Maps	9-47
9.11.1	Configure Audit Properties.....	9-47
9.11.2	Capture the Data Graph.....	9-47
9.11.3	Submit the Update.....	9-48
9.12	How To Test an Update Procedure.....	9-49
9.12.1	Configure Audit Properties.....	9-50
9.12.2	Capture the Data Graph.....	9-50
9.12.3	Submit the Update.....	9-50

10 Preparing Services for Clients

10.1	Generating a Mediator Client JAR File	10-1
10.1.1	Using the IDE	10-1
10.1.2	Using the Command-Line Tool	10-2
10.2	Generating a Web Services Mediator Client JAR File	10-3
10.2.1	Overview	10-3
10.2.2	Using Eclipse for WebLogic	10-3
10.2.3	Using the Command-Line Tool	10-4
10.3	Generating a Web Service Map and WSDL from a Data Service	10-5
10.3.1	Creating a Map File	10-5
10.3.2	Generating a WSDL File from a Map File	10-7
10.3.3	Examining the Generated WSDL	10-7
10.3.4	Testing the Generated WSDL.....	10-7
10.3.5	Modifying the Map File	10-7
10.3.5.1	Adding Data Services and Operations	10-8
10.3.5.2	Deleting Data Services and Operations from a Map File	10-8
10.3.5.3	Renaming Mapped Operations	10-8
10.4	Configuring Security for Web Services Applications.....	10-8
10.4.1	Configuring Basic Authentication.....	10-8
10.4.2	Configuring Transport Level Security (HTTPS)	10-8
10.4.3	Configuring Web Services Security (WSS).....	10-9
10.4.4	Specifying Policies	10-9
10.4.4.1	Specifying Global Policies	10-9

Sample Map File 10

10.4.4.2	Specifying Policies for a Function	10-10
10.5	Web Services Map File Reference	10-11
10.5.1	Map File-Level Properties	10-11
10.5.2	Operation Level Properties	10-12
10.5.3	Map File XML Schema Definition	10-13

Web Services Map File Schema Definition 13

10.5.4	Mapping of Data Service Type to WSDL Message Type	10-15
10.5.4.1	Two Schema Elements Per Function.....	10-15

Operation Element and Return Element 16

10.5.4.2	Mapping of Update Functions with DataGraphs	10-16
----------	---------------------------------------------------	-------

10.5.4.3	Overloading Data Service Functions	10-17
Overloaded Functions 17		
10.5.5	Examining the Generated WSDL	10-17
10.5.6	Testing the Generated WSDL.....	10-18
10.5.7	Copying and Saving a WSDL Generated from a Map	10-18
10.6	Understanding SQL Maps	10-19
10.6.1	Overview	10-19
10.6.2	Publishable Operations.....	10-20
10.6.3	General Conditions.....	10-21
10.7	Map Functions and Procedures to SQL Objects	10-21
10.7.1	Creating an SQL Map.....	10-21
10.7.2	Removing an SQL Map.....	10-22
10.8	SQL Object Mapping Rules	10-22
10.9	Constraints on Publishing Data Service Objects to SQL	10-23
10.9.1	Non-Tabular Element Types Affect Ability to Publish Functions as SQL Objects	10-24

11 Data Service Annotations

11.1	Overview	11-1
11.2	XDS Annotations.....	11-1
11.2.1	General Properties	11-3
11.2.1.1	Standard Document Properties	11-3
11.2.1.2	User-Defined Properties	11-3
11.2.2	Data Access Properties.....	11-3
11.2.2.1	Relational Data Service Annotations	11-4
11.2.2.2	Source Binding Provider	11-4
11.2.2.3	Web Service Data Service Annotations	11-5
11.2.2.4	Java Function Data Service Annotations.....	11-5
11.2.2.5	Delimited Content Data Service Annotations	11-6
11.2.2.6	XML Content Data Service Annotations	11-6
11.2.2.7	User Defined View XDS Annotations	11-6
11.2.3	Target Type Properties.....	11-7
11.2.3.1	Native Type Properties	11-7
11.2.3.2	Update-related Type Properties	11-8
11.2.4	Key Properties	11-8
11.2.5	Relationship Properties.....	11-9
11.2.6	Update Properties.....	11-9
11.2.6.1	Optimistic Locking Fields	11-10
11.2.6.2	Security Properties	11-10
11.3	Function Annotations.....	11-10
11.3.1	General Properties	11-13
11.3.2	UI Properties.....	11-13
11.3.3	Cache Properties	11-14
11.3.4	Transaction Properties	11-14
11.3.5	Behavioral Properties	11-14
11.3.5.1	Inverse Functions.....	11-14
11.3.5.2	Equivalent Transforms	11-15

11.3.6	Polymorphic Functions	11-15
11.3.7	Signature Properties	11-16
11.3.8	Native Properties	11-16
11.3.8.1	SQL Query Properties	11-16
11.3.8.2	SOAP Handler Properties	11-17
11.3.9	Implementation Properties.....	11-17
11.4	XFL Annotations	11-17
11.4.1	General Properties	11-18
11.4.2	Data Access Properties.....	11-18
11.4.3	Security Properties.....	11-19
11.5	Data Service Annotations Schema.....	11-19

12 Best Practices When Building Data Services

12.1	Overview	12-1
12.1.1	Understanding the Oracle Data Service Integrator Server	12-1
12.1.2	Understanding the Oracle Data Service Integrator Client.....	12-2
12.2	Oracle Data Service Integrator Development Best Practices	12-2
12.2.1	Organizing Data Services Using Projects	12-2
12.2.2	Building Data Services in Layers.....	12-3
12.2.2.1	Physical Layer	12-3
12.2.2.2	Logical Layer	12-4
12.2.2.3	Integration Layer	12-5
12.2.2.4	Application Layer	12-6
12.3	Performance and Optimization Best Practices	12-7
12.3.1	Database Access	12-7
12.3.1.1	Retrieving Only the Necessary Data.....	12-7
12.3.1.2	Designing Functions Which can be Pushed to the Database	12-8
12.3.1.3	Verifying that Joins are Implemented as Left-Outer or Inner/Natural Joins...	12-8
12.3.1.4	Pushing Left Outer Joins to the Database	12-8
12.3.1.5	Using a ppci-impl Join	12-9
12.3.1.6	Avoiding Casting.....	12-10
12.3.1.7	Optimizing Ad Hoc Queries	12-10
12.3.1.8	Writing Your Own SQL for Oracle Data Service Integrator to Use	12-11
12.3.2	Exercising Care when Using Fail-over, Fail-over-retry, and Timeout.....	12-11
12.3.3	Using Inverse Functions	12-12
12.3.4	Using Caching and Auditing	12-13
12.3.5	Query Plans	12-13
12.3.5.1	Evaluating Performance Before Running the Query.....	12-13
12.3.5.2	Precompiling Query Plans	12-13
12.3.5.3	Evaluating the Performance by Running the Query	12-14
12.3.6	Monitoring Operational Performance and Service Level Agreements.....	12-14
12.4	How To Get More Help	12-14
12.5	Related Topics	12-14

Preface

This document describes how to develop applications for the Oracle Data Service Integrator.

Audience

This document is intended for application developers.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documents

For more information, see the following documents in the Oracle Data Service Integrator documentation set:

- *Oracle Fusion Middleware Using Data Service Integrator XQuery Engine*
- *Oracle Fusion Middleware Administering Data Service Integrator*
- *Oracle Fusion Middleware Installing Data Service Integrator*
- *Oracle Fusion Middleware Developing Applications with Data Service Integrator*
- *Oracle Fusion Middleware Data Services Java API for Oracle Data Integrator*

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 to Data Services

This chapter provides information, examples, and tutorials on data services basics.

This chapter includes the following sections:

- [Section 1.1, "Concepts"](#)
- [Section 1.2, "How-to"](#)
- [Section 1.3, "Example: How to Create Your First Data Services"](#)
- [Section 1.4, "Reference"](#)
- [Section 1.5, "Related Topics"](#)

1.1 Concepts

This section presents the following topics:

- [Section 1.1.1, "Data in the 21st Century"](#)
- [Section 1.1.2, "Data Access Integration Architecture"](#)
- [Section 1.1.3, "Oracle Data Service Integrator: Roles and Responsibilities"](#)

1.1.1 Data in the 21st Century

In modern enterprises data is generally readily available. While this has reduced the need to move physical data into data warehouses, data marts, data mines, or other costly replications of existing data structures, the problems of dynamic data integration, immediate secured access and update, data transformation, and data synchronization remain some of the most vexing challenges facing the IT world.

Oracle Data Service Integrator provides a comprehensive approach to this challenge by:

- Providing a unified means of importing metadata representing the structure of any data source using its Metadata Import wizard.
- Allowing for the creation of hierarchical data structures from traditional column-row data.
- Providing a query-driven interface to extend the physical model so data specialists can create powerful transformations of existing data and queries.
- Automatically creating data models that introspect physical data structures (and their contents) in situ, normalizes representation of diverse data, and allow the representation of the relationship of physical and logical data.

- Maintaining the accuracy of metadata through automated updates from the data source.

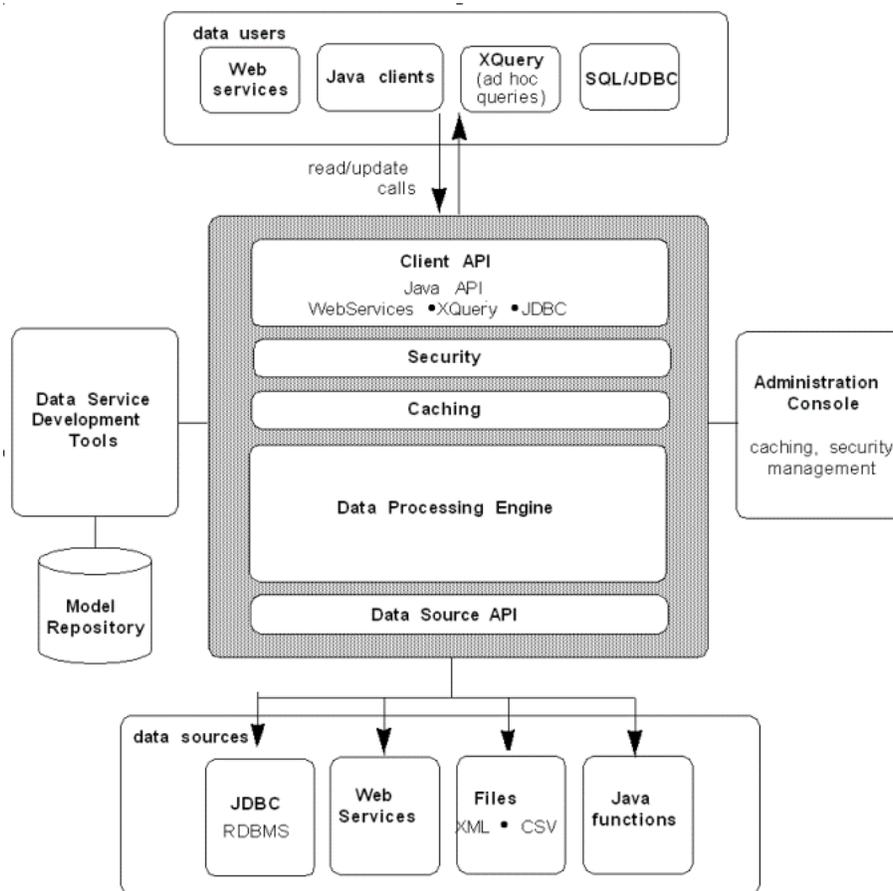
Oracle Data Service Integrator can be used to create, refine, and validate logical data structures through a process of importing data sources, creating physical and logical models, and designing queries for use by applications in an infrastructure that provides for easy maintenance, while enhancing security and performance.

Through standardized Service Data Objects (SDO) technology, web-based applications can automatically read and update relational data. Through simple Java programs Oracle Data Service Integrator update capabilities can be extended to support any logical data source.

1.1.2 Data Access Integration Architecture

In contemporary enterprise computing, data typically passes through multiple processing and storage layers. While enterprise data can easily be accessed, turning that data into useful information economically and efficiently, particularly updateable information, remains a difficult and high-maintenance task.

Figure 1-1 Component Architecture



Data users access data through the client API and data sources gain access through the data source API.

Oracle Data Service Integrator approaches the problem of creating integration architectures by building logical data services around physical data sources and then allowing business logic to be added as part of easily maintained, graphically designed XML query functions (also called XQueries).

Any Eclipse for WebLogic application can include Oracle Data Service Integrator-based projects. And any application can access Oracle Data Service Integrator queries — including update functions — through a mediator API or an Oracle Data Service Integrator Control. In the case of relational data, updates can be performed automatically through Service Data Objects (SDO) (For details see "Programming with Service Data Objects" in the *Oracle Data Service Integrator Client Application Developer's Guide*.)

Oracle Data Service Integrator provides for the development of integrated queries within any Eclipse for WebLogic application. Each application can contain multiple Oracle Data Service Integrator-based projects, as well as any other types of projects offered by Eclipse for WebLogic.

1.1.3 Oracle Data Service Integrator: Roles and Responsibilities

The following summarizes typical roles and responsibilities related to creating and maintaining data services.

- **Physical Data Service Development.** Any team member can quickly create a set of physical data services from enterprise data sources.
- **Entity Data Service Development.** A data architect with knowledge of the relationships between enterprise data sources can then create data services based on physical and previously developed *logical data services*.
- **Query Development.** Once data services are created, an IT team member can create reusable query functions using the graphical XQuery Editor. The editor is directly tied to a Source View that facilitates code-based modifications to automatically-generated designs.
- **Deployment.** Once data services are developed, they can be deployed from the IDE or by an administrator through the Oracle Data Service Integrator Administration Console.
- **Application Development.** Application designers can use data service query functions in their Oracle WebLogic applications. Through Service Data Objects (SDO) and the Mediator API or an Oracle Data Service Integrator Eclipse Control, applications can retrieve and update data, yet remaining insulated from the complexities of managing the underlying data interaction.
- **Metadata Management.** Administrators, architects, and designers can use the Service Explorer for real-time introspection of disparate data source metadata that has been developed through Oracle Data Service Integrator.

1.2 How-to

This section presents the following sections:

- [Section 1.2.1, "How to Configure the Retail Dataspace Sample Application for Oracle Data Service Integrator"](#)
- [Section 1.2.2, "How to Configure the Retail Dataspace Sample Application for OSDI Studio"](#)

1.2.1 How to Configure the Retail Dataspace Sample Application for Oracle Data Service Integrator

This section describes how to set up the Retail Dataspace Sample Application after completing the installation of Oracle Data Service Integrator.

- [Section 1.2.1.1, "Prerequisites"](#)
- [Section 1.2.1.2, "About WorkSpace Studio, Data Services Studio, and Eclipse"](#)
- [Section 1.2.1.3, "Start WorkSpace Studio"](#)
- [Section 1.2.1.4, "Start the Server"](#)
- [Section 1.2.1.5, "Deploy Your Projects"](#)
- [Section 1.2.1.6, "Create the Retail Dataspace Sample Web Application"](#)
- [Section 1.2.1.7, "See Also"](#)

1.2.1.1 Prerequisites

A prerequisite to configuring the retail dataspace sample application is to have the Oracle Data Service Integrator installed on a supported platform.

Choose the guide for the version you are running:

- Installation Guide for OSDI

1.2.1.2 About WorkSpace Studio, Data Services Studio, and Eclipse

Note: WorkSpace Studio was formerly named Data Services Studio.

This tutorial uses the version of Eclipse that is installed with Oracle Data Service Integrator.

The Eclipse framework often provides multiple ways of achieving same result. In many cases there is no "correct" or "better" way. In other words, there are often many paths to the same results.

1.2.1.3 Start WorkSpace Studio

Open OSDI from Eclipse by selecting:

Window > Open Perspective > Oracle Data Service Integrator

1.2.1.3.1 Select a Workspace OSDI Studio projects are called *dataspace* projects. These Projects in turn are located in a workspace folder.

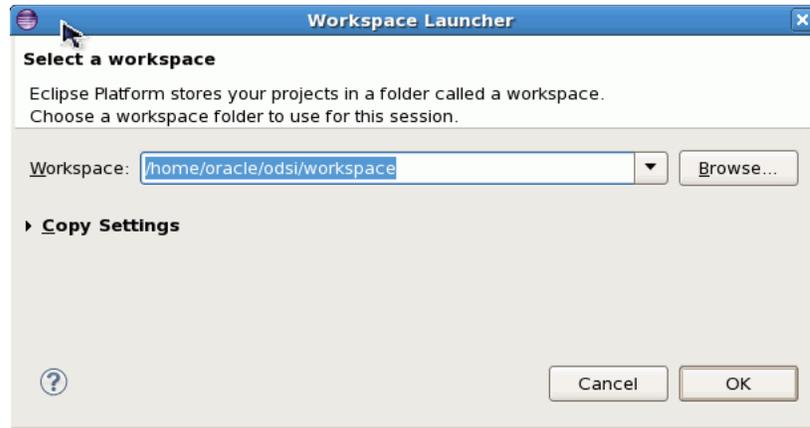
The first step in creating a dataspace is to select a workspace.

1. Use the default location:

Install Retail Dataspace

2. Click **OK**.

Figure 1–2 Oracle Data Service Integrator Selecting a Workspace

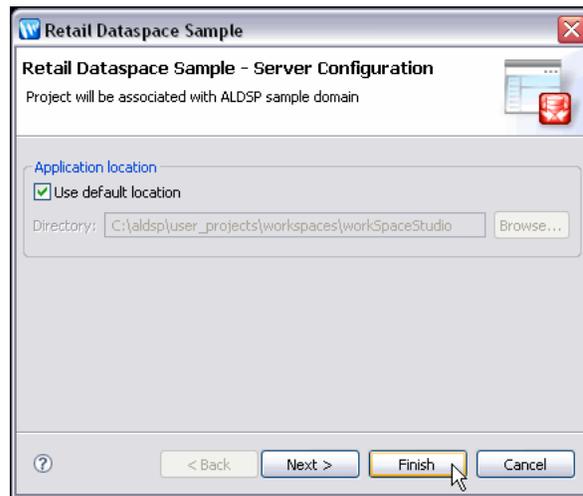


If this is the first time you have opened Studio, the WorkSpace Studio screen displays. Browse to choose a workspace folder to use for the current session.

In the Samples section click on:

Install Retail Dataspace Sample

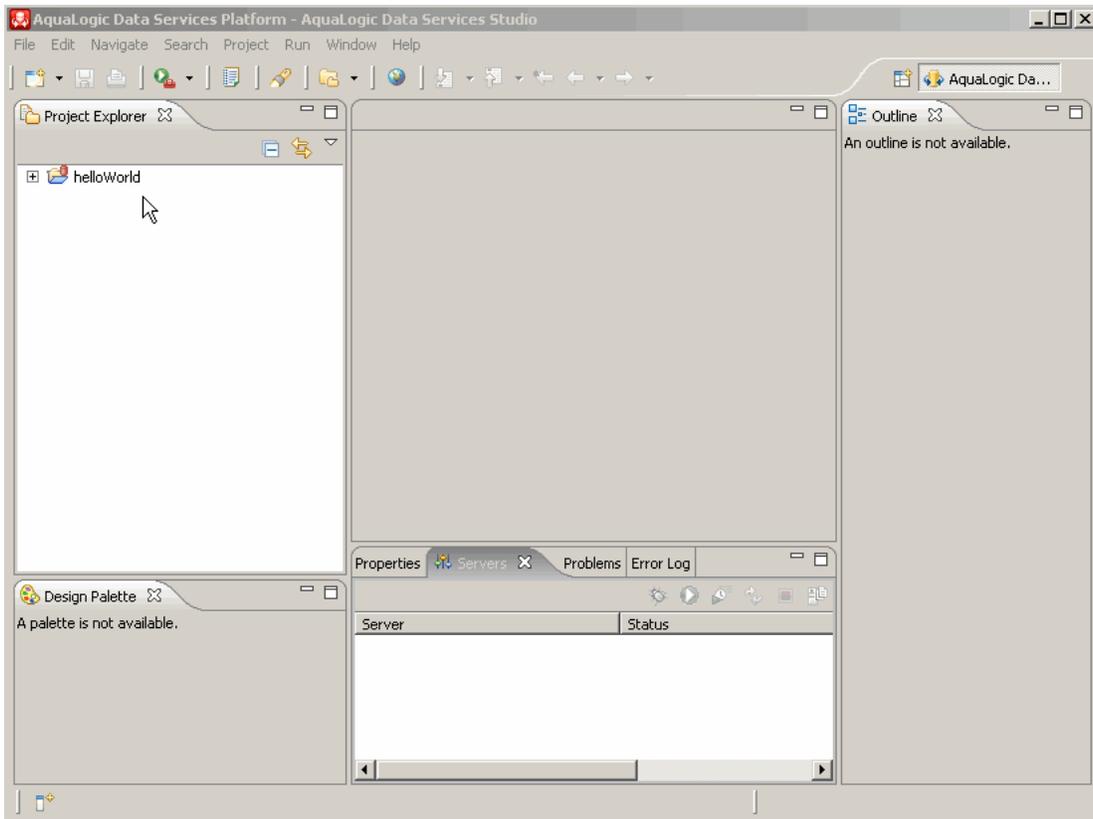
Figure 1–3 Retail Dataspace Server Configuration Dialog



The Server Configuration dialog displays. You can choose the default location or browse for an application location.

1. Click **Finish**. This imports the RetailDataspace project and finds or creates a server for the Oracle Data Service Integrator sample domain and associates it with the project.
2. Answer **Yes** to the question about associating your project with the Oracle Data Service Integrator perspective.

Figure 1–4 Initial Oracle Data Service Integrator Perspective



The initial Data Services Studio screen displays. There are Project Explorer, Design Palette, Properties, and Outline tabs.

1.2.1.3.2 Start the Server An Oracle Data Service Integrator-enabled server is a version of WebLogic Server with additional functionality to support Oracle Data Service Integrator deployment and runtime. The Oracle Data Service Integrator server must be running in order to access sample data and to deploy your project.

To start your server from Studio:

1. Locate the Servers window. If it isn't visible, use the following option command:

Window > Show View > Servers

1.2.1.4 Start the Server

An Oracle Data Service Integrator-enabled server is a version of WebLogic Server with additional functionality to support Oracle Data Service Integrator deployment and runtime. The Oracle Data Service Integrator server must be running in order to access sample data and to deploy your project.

To start your server from Studio:

1. Locate the Servers window. If it isn't visible, use the following option command:

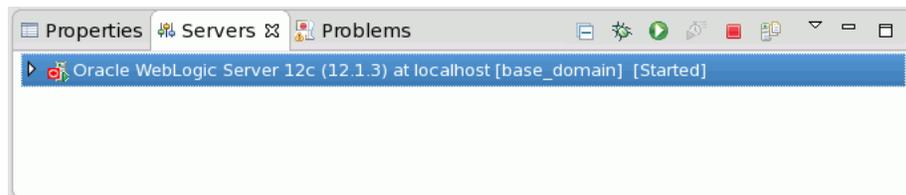
Window > Show View > Servers

2. In the Servers window locate the **Oracle Data Service Integrator Samples Server** (this may be the only server listed). Notice that its status is **Stopped**.

- Right-click on the server name and select **Start**. (The start-up operation can take several minutes.) Notice the running log of server startup actions in the Console window.

You can also start the server by selecting the server and clicking the **Start** icon.

Figure 1–5 Server Window



The Server window shows the Oracle Data Service Integrator Samples Server with the status Started and the State Synchronized.

1.2.1.5 Deploy Your Projects

Each project should be deployed to validate the installation.

- Right-click on the server.
- Choose **Publish** from the menu. A message should appear indicating successful deployment.
- Click **OK**.

Also deploy the **RetailDataspaces** project.

1.2.1.6 Create the Retail Dataspaces Sample Web Application

If Eclipse for WebLogic Platform 10.2 is installed into the same `BEA_HOME` which contains `ALDSP_HOME`, you can create the a web-based sample application.

Caution: Platform 10.0 MP1 contains version 10.0 of Eclipse. However, the sample application required features available in Eclipse 10.2. The workaround is to install the 10.2 version of Eclipse for WebLogic Platform. This version is available from the download site. Current that link is:

<http://commerce.bea.com/showproduct.jsp?family=WLV&major=10.2&minor=0>

Follow these steps to access the web-based sample application:

- To make sure the necessary dataspace are deployed, redeploy (right-click > **Deploy Project**):
 - ElectronicsWS** and **RetailDataspaces** projects
- From the WorkSpace 1.1 menu select: **Window > Show View > Servers**
- Locate the option:

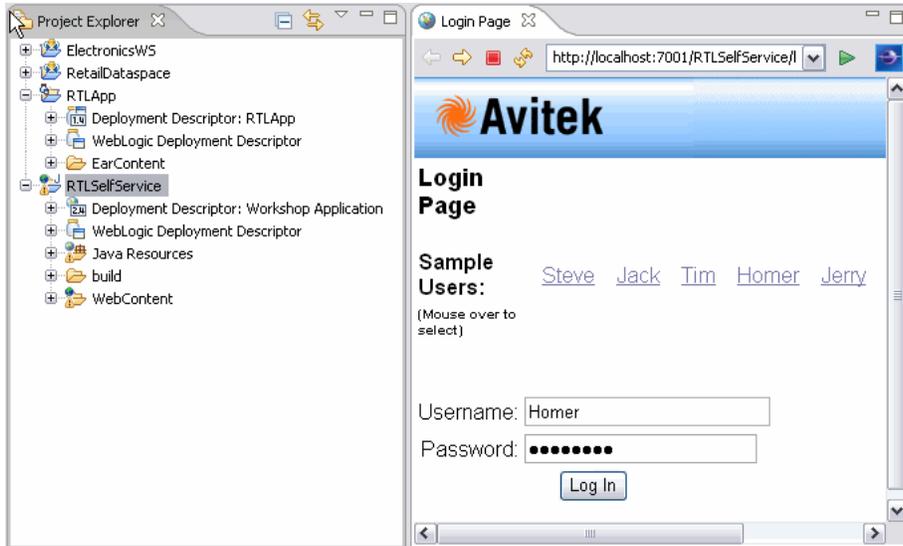
Retail Dataspace Sample Web Application (WebLogic Eclipse only)

4. Click **Next**, then **Finish**.
5. If asked if you want to open the OSDI Studio Perspective, click **Yes**.
6. In the Project Explorer view, right-click on the **RTLSelfService** project, and choose:

Run As > Run on Server

This will initially deploy your projects and then open the sample Avitek login page.

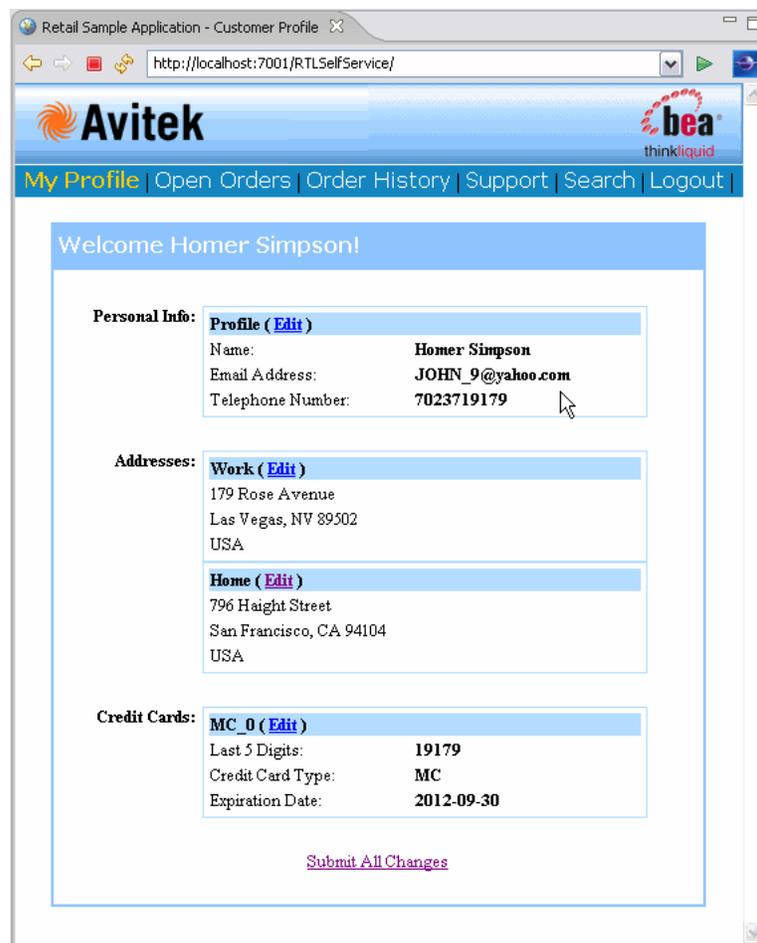
Figure 1–6 Avitek Login Page



The Avitek login page shows the Project Explorer and a place to enter your username and password.

7. Mouse over one of the names and log in. After a few moments information about the fictitious customer will appear.

Figure 1–7 Avitek Welcome Page



The Avitek Welcome Page has six tabs: Profile, Open Orders, Order History, Support, Search, and Logout. The Profile page displays Personal Info, Addresses, and Credit Cards. You can edit this data.

1.2.1.7 See Also

- [Example: How to Create Your First Data Services](#)
- *Retail Database Sample Application Guide*

1.2.2 How to Configure the Retail Dataspace Sample Application for OSDI Studio

This section describes how to set up the Retail Dataspace Sample Application after completing the installation of OSDI Studio.

- [Section 1.2.2.1, "Prerequisites"](#)
- [Section 1.2.2.2, "About Workspace Studio, Data Services Studio, and Eclipse"](#)
- [Section 1.2.2.3, "Start Studio"](#)
- [Section 1.2.2.4, "Start the Server"](#)
- [Section 1.2.2.5, "Deploy Your Projects"](#)

- [Section 1.2.2.6, "See Also"](#)

1.2.2.1 Prerequisites

A prerequisite to configuring the retail datasource sample application is to have the ODSI Data Services Studio installed on a supported platform.

1.2.2.2 About WorkSpace Studio, Data Services Studio, and Eclipse

Note: Data Services Studio became part of WorkSpace Studio with Eclipse. Where possible, the generic term *Studio* is used.

This tutorial uses the version of Eclipse that is installed with ODSI.

The Eclipse framework often provides multiple ways of achieving same result. In many cases there is no "correct" or "better" way. In other words, there are often many paths to the same results.

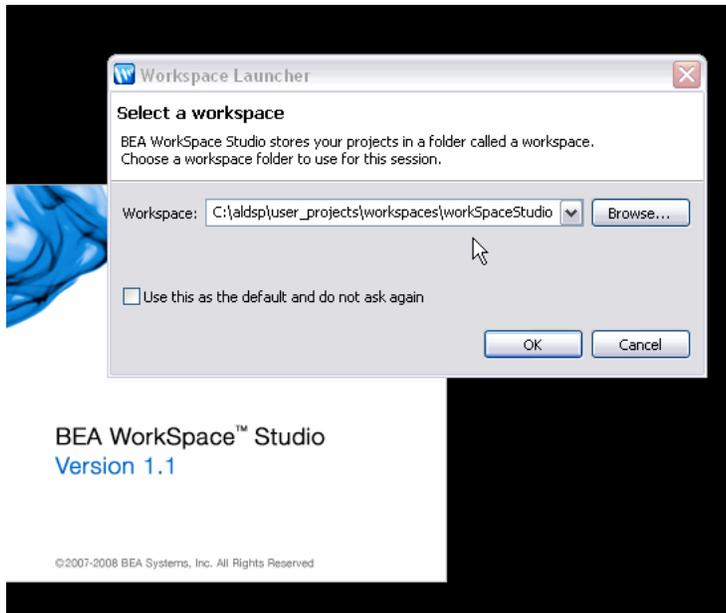
1.2.2.3 Start Studio

Open Studio using the following Windows Start menu command:

Start > All Programs > BEA Products > BEA AquaLogic Data Services Platform 3.0 > Data Services Studio

1.2.2.3.1 Select a Workspace

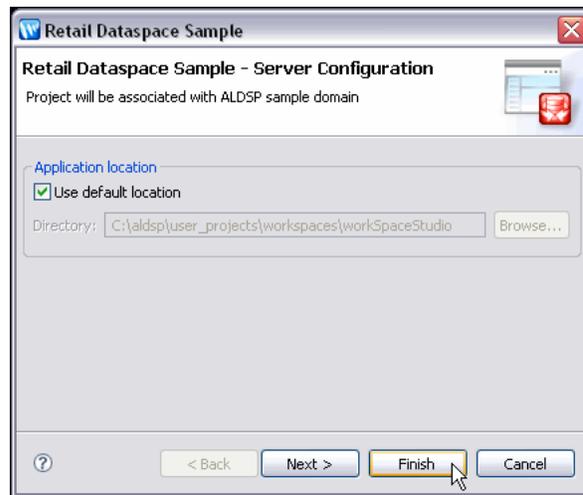
Figure 1–8 Selecting a Workspace



If this is the first time you have opened Studio, the WorkSpace Studio screen displays. Browse to choose a workspace folder to use in the current session.

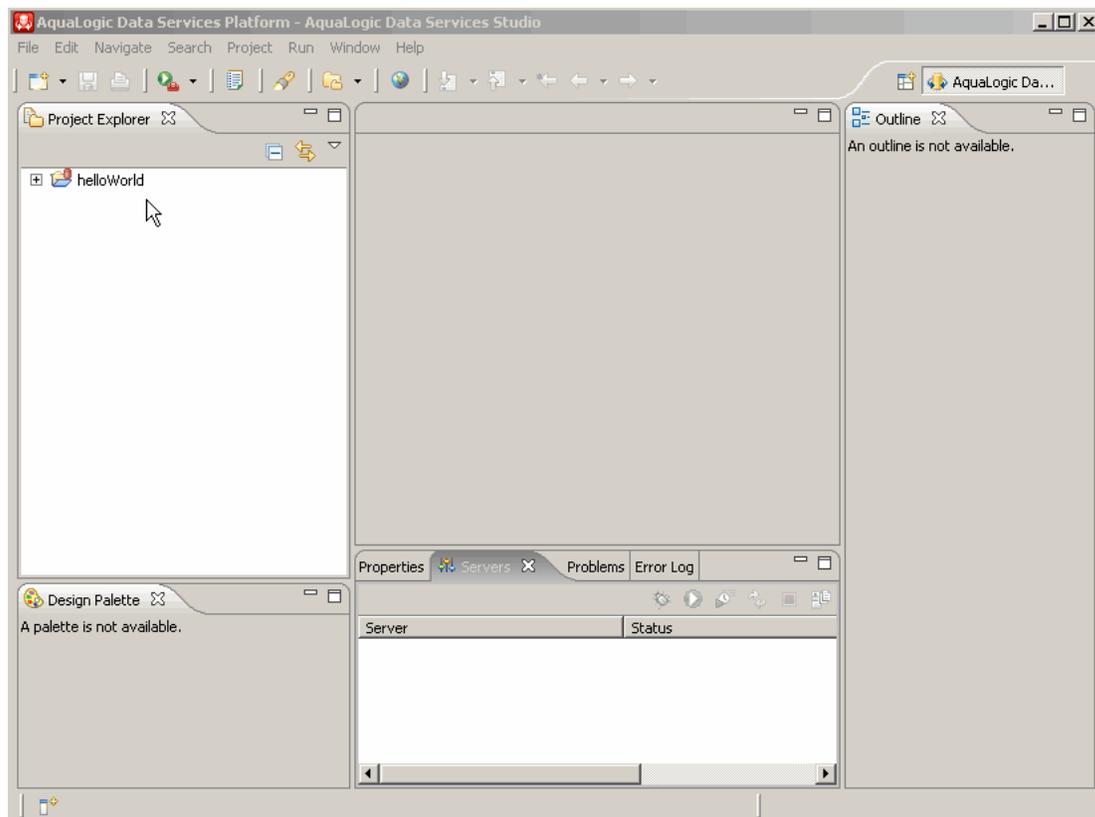
In the Install Sample Applications section click on:

Retail Dataspace Sample

Figure 1–9 Retail Dataspace Server Configuration Dialog

The Server Configuration dialog displays. You can choose the default location or browse for another application location.

1. Click **Finish**. This will import the RetailDataspace project and find or create a server for the OSDI Studio sample domain and associate it with the project.
2. Click **Yes** if you are asked about associating your project with the OSDI Studio perspective.

Figure 1–10 Initial OSDI Studio Perspective

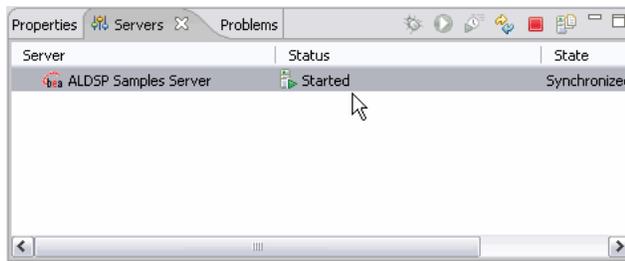
The initial Services Studio screen displays. There are Project Explorer, Design Palette, Properties, and Outline tabs.

1.2.2.4 Start the Server

An OSDI Studio-enabled server is a version of WebLogic Server with additional functionality to support OSDI Studio deployment and runtime. The OSDI Studio server must be running in order to access sample data and to deploy your project.

1. Locate the Servers window. If it isn't visible, use the following option command:
Window > Show View > Servers
2. In the Servers window locate the OSDI Studio Samples Server (this may be the only server listed). Notice that its status is Stopped.
3. Right-click on the server name and select **Start**. (The start-up operation can take several minutes.) Notice the running log of server startup actions in the Console window.

Figure 1–11 Server Window



The Server window shows the OSDI Studio Samples Server with the status Started and the State Synchronized.

1.2.2.5 Deploy Your Projects

Each project should be deployed to validate the installation.

1. Right-click on the server.
2. Choose **Publish** from the menu. A message should appear indicating successful deployment.
3. Click **OK**.

Also deploy the **RetailDataspaces** project.

1.2.2.6 See Also

- Install the OSDI Studio Sample Retail Application
- [Example: How to Create Your First Data Services](#)
- *Retail Dataspaces Sample Application Guide*

1.3 Example: How to Create Your First Data Services

Creating a data service from scratch — as you will if you follow this tutorial — is a good way to get the feel of working with Eclipse for WebLogic, as well as other aspects of data services. In the process of creating a logical data service you also automatically create several physical data services. Physical data services represent physical data sources.

This tutorial contains the following sections:

- [Section 1.3.1, "Goal of the Tutorial"](#)
- [Section 1.3.2, "Creating a Dataspace Project"](#)
- [Section 1.3.3, "Creating Physical Data Services"](#)
- [Section 1.3.4, "Creating a Logical Data Service"](#)
- [Section 1.3.5, "Creating, Saving, and Associating the XML Type"](#)
- [Section 1.3.6, "Testing Your Data Service Function"](#)
- [Section 1.3.7, "Adding Create-Update-Delete Functions to Your Data Service"](#)
- [Section 1.3.8, "Updating Your Results"](#)
- [Section 1.3.9, "Reviewing the Query Plan"](#)
- [Section 1.3.10, "Reviewing the Update Map"](#)
- [Section 1.3.11, "Archiving Your Project"](#)
- [Section 1.3.12, "Summary"](#)

1.3.1 Goal of the Tutorial

The goal of this tutorial is to illustrate an approach to creating a logical data service, including creating an XML Type (schema), using Eclipse for WebLogic. Along the way you will use many of the facilities:

- Drag-and-drop Query Map
- Source Editor
- Test Editor
- Query Plan
- Update Map

This example uses data provided with the Retail Dataspace Sample Application (RTLApp).

1.3.1.1 Requirements

The requirement for the demonstration project are to develop a logical data service from several physical data services. When run by a client, the data service will return a consolidated view of a particular customer's orders, as well as all the items in each order.

1.3.1.2 Before You Begin

Before you can begin the tutorial make sure you:

- Properly install Oracle Data Service Integrator.

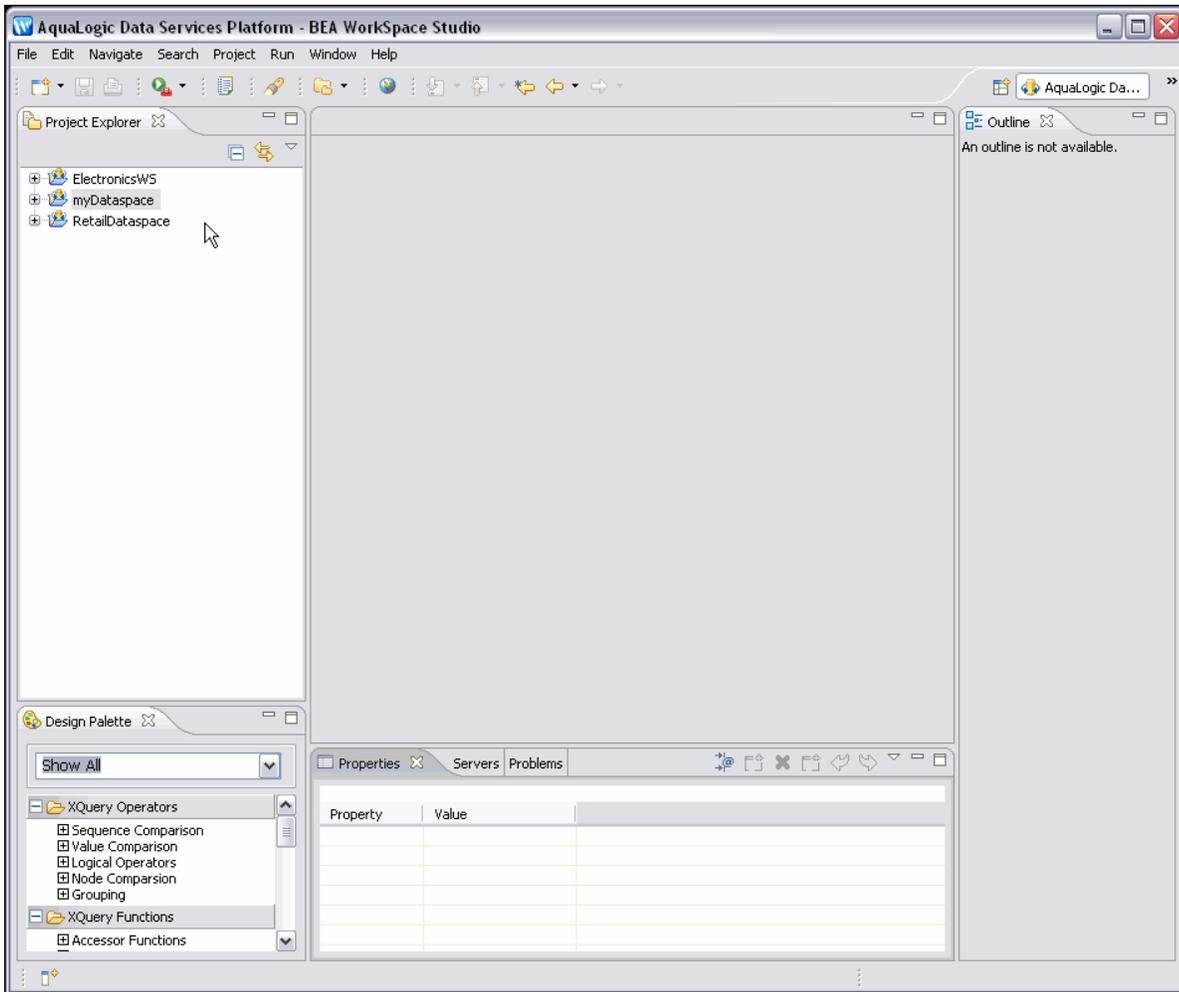
Reference:

Oracle Data Service Integrator Installation Guide

- Configure the Retail Dataspace Sample Application.
 - Configure the Retail Dataspace Sample Application
 - Configure the Retail Dataspace Sample Application for ODSI
- Have the application open in Eclipse for WebLogic and the Oracle Data Service Integrator-enabled Oracle WebLogic 10.3 server running.

Also describe in "Configure the Retail Dataspace Sample Application."

Figure 1–12 Oracle Data Service Integrator Default Perspective After Adding myDataspace



In the WorkSpace Studio, the Project Explorer tab shows project directories and the Design Palette tab shows XQuery operators and functions.

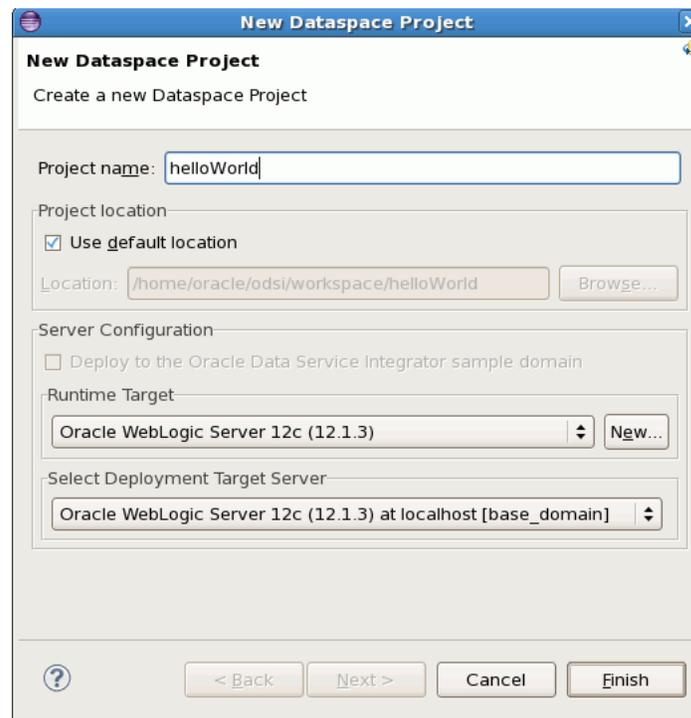
Note: Click on image to view it enlarged in a separate window.

1.3.2 Creating a Dataspace Project

Data services are created within Eclipse for WebLogic as Eclipse projects, called dataspace projects. With the Oracle Data Service Integrator-enabled server running, the first step is to create a new dataspace project.

1. From the menu select:
File > New > Dataspace Project
2. Give your project a name such as:
myDataspace
3. Click the **Finish** button.

Figure 1–13 *Creating a New Dataspace Project*



The New Dataspace Project screen provides a space for entering your project name, choosing a directory for project contents (or selecting the default location), choosing a runtime target, and selecting a deployment target server.

1.3.2.1 Set Up a Folder for Physical Data Services

Data services are typically created inside project folders. The recommended first step in creating one or several data services is to create containers (folders).

In this tutorial two folders will be created:

- One for physical data services.
 - One for logical data services.
1. In the Project Explorer window right-click on **myDataspace**, choose:
New > Folder

2. Name your folder:
logical
3. Click the **Finish** button.
4. Create another folder under **myDataspac** named:
physical
5. Click the **Finish** button.

Physical data services represent physical data such as tables in relational databases or web services. Logical data services are build upon existing physical or logical data services.

- [Chapter 3, "Creating and Updating Physical Data Services"](#)
- [Chapter 4, "Designing Logical Data Services"](#)

Figure 1–14 *Creating a New Folder*



The Folder dialog lets you enter the parent folder for your folder, or select the parent folder from a list. There is field for specifying the folder name. An Advanced button offers additional choices.

6. Right-click on your new physical folder and choose:
New > Physical Data Service

1.3.3 Creating Physical Data Services

Physical data services are based on existing data sources.

Whenever you create physical data services, you must first identify the data source. Available options include:

- Relational
- Web Service
- Java Function
- Delimited Data
- XML Data

To take advantage of data provided with the sample application, a relational data source is used.

The sample databases **RTLAPPLOMS** and **RTLCUSTOMER** provided with the Retail Sample Application contain five tables. In this section you will create physical data services corresponding to those tables.

Table 1–1 Data Sources and Data Services

Data Source	Name	Table	Data Service
RTL Appliance Order Management System	RTLAPPOMS	■ CUSTOMER_ORDER	■ CUSTOMER_ORDER.DS
		■ CUSTOMER_ORDER_LINE_ITEM	■ CUSTOMER_ORDER_LINE_ITEM.DS
RTL Customer Data	RTLCUSTOMER	■ ADDRESS	■ ADDRESS.DS
		■ CUSTOMER	■ CUSTOMER.DS

1.3.3.1 Select a Data Source

The select a data source dialog initially allows you to select a data source type (such as relational or web service). Once that selection is made, additional options appear. The following table lists the actions required to select the relational data sources that will be used throughout this tutorial.

Format similar to that shown in the table below is used to describes the steps needed to work through multi-page wizards.

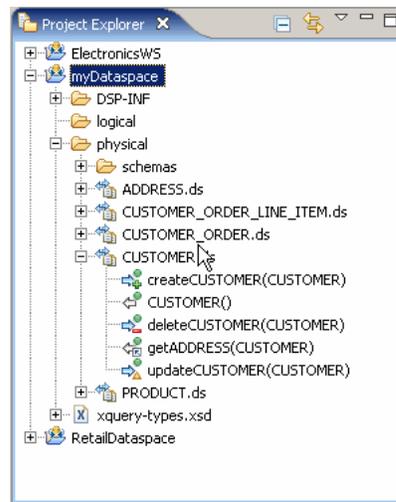
Table 1–2 Setting Up Sources for Data Services

Step	Dialog	Field/Column	Action	Comment
	Select Data Source	Save in:		Use default (/myDataspaces/physical)
1.		Data source type:	Select Relational	From dropdown list.
2.		Data source:	Select dspSamplesDataSource	
3.			Click the Next button.	
4.	Select SQL Sources	Select SQL objects:	<ul style="list-style-type: none"> ■ Checkbox next to RTLAPPLOMS. ■ Checkbox next to RTLCUSTOMER. 	Expand (+ symbol to left of data source name) to see tables in the data sources.

Table 1–2 (Cont.) Setting Up Sources for Data Services

Step	Dialog	Field/Column	Action	Comment
5.			Click the Next button.	The information retrieved through introspection of relational data sources is represented as the potential creation of the five primary Read operations, as well as their containing data services.
6.	Review New/Updated Data Service Operation(s)	Public	Mark all five operations Public by clicking the checkbox in the Public column.	Public operations are available to any authorized calling application. Note: The Primary option only applies to create, update, and delete functions.
7.		Select Common XML Type Namespace... button	Click the button.	Because you are building up an XML Type for your logical data service from several physical data services that each have an underlying XML type, it is necessary for each type to share a namespace.
8.	XML Type Namespace	Select XML Type Namespace:	Enter custOrdersItems .	
9.		Click the OK button.		Notice that the target namespace column now shows the new namespace for your operations.
10.	Review New/Updated Data Service Operation(s)		Click the Next button.	
11.	Review New Data Service(s)		Click the Finish button.	It is necessary to modify names when: <ul style="list-style-type: none"> ■ A data service of the same name already exists in the specified folder. ■ You are attempting to import two data sources with the same name. In this example, however, there are no name conflicts and no changes are needed.
12.	Open Data Service Files	Option to open each new physical data service in Eclipse for WebLogic	Select No .	

Your new data services appear in your physical folder in the Project Explorer.

Figure 1–15 Newly Created Data Services

The Project Explorer shows the data services you have just created.

If you expand your new data services you will see that each physical data service has been created with functions corresponding to standard relational operations. For example the `CUSTOMER.ds` data service contains the following operations:

- `createCUSTOMER(CUSTOMER)`
- `CUSTOMER()`
- `deleteCUSTOMER(CUSTOMER)`
- `getADDRESS(CUSTOMER)`
- `updateCUSTOMER(CUSTOMER)`

Some relationship operations (such as `getADDRESS(CUSTOMER)`) have been created automatically. This operation returns an `ADDRESS` type when it is passed a `CUSTOMER` type as a parameter. The operation can be inferred during the data service creation process because `ADDRESS` contains a foreign key that is a unique `custID` in the `CUSTOMER` data service (and underlying source). Relationship functions are described in detail in the [Modeling Data Services Relationships](#) section.

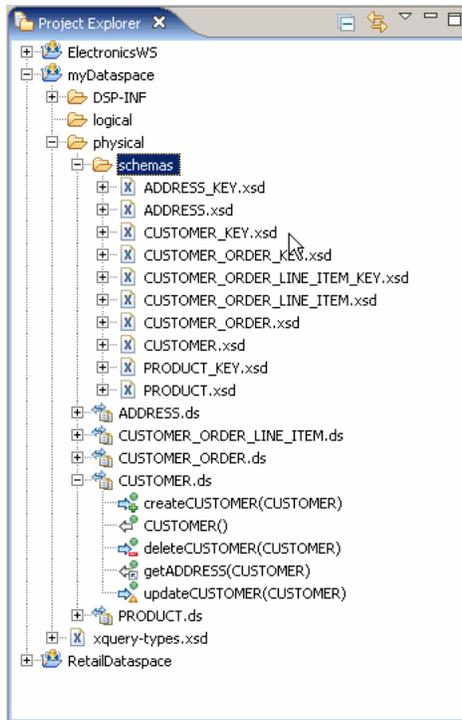
1.3.3.2 Schemas Directory

You should find a `schemas` folder adjacent to the newly created data services. This folder contains schema files created during the metadata import process. For relational sources, schemas are created for both the data source (table or view) and the primary keys found during the introspection of the relational source. For example:

- `CUSTOMER.xsd`
- `CUSTOMER_KEY.xsd`

If you look in the `schemas` directory you will see that for each physical data service created, two schemas were created. One representing the physical data service and the other to describe the primary keys in the data source.

Figure 1–16 Expanded View of Project Explorer



In the expanded view of the schemas directory, the data source/primary key pairs are shown: ADDRESS.xsd and ADDRESS_KEY.xsd; CUSTOMER.xsd and CUSTOMER_KEY.xsd; CUSTOMER_ORDER.xsd and CUSTOMER_ORDER_KEY.xsd; CUSTOMER_ORDER_LINE_ITEM.xsd and CUSTOMER_ORDER_LINE_ITEM_KEY.xsd; and PRODUCT.xsd and PRODUCT_KEY.xsd.

When a logical entity data service is created, it is either:

- Associated with an existing schema or
- A return type associated with a function becomes the basis of a generated XML type that is then associated with the data service.

1.3.3.3 Publish Your Projects

Using Eclipse for WebLogic, you can publish your dataspace projects to a server when it is ready for testing and debugging. Publishing is also useful during the project development phase because in its default configuration, when you publish a project in Eclipse for WebLogic, it is automatically built and validated. The validation process identifies error conditions, if any.

Note: When publishing a project to a server, the project is validated and only valid projects are successfully published.

A dialog displays the progress and, upon successful completion, the status of the server changes to **Synchronized**.

1.3.4 Creating a Logical Data Service

This section describes the following topics:

- [Section 1.3.4.1, "Attempt To Publish Your Dataspace Project"](#)
- [Section 1.3.4.2, "Bottom Up or Top Down"](#)
- [Section 1.3.4.3, "Add an Operation to CUST_ORDERS_ITEMS"](#)
- [Section 1.3.4.4, "Building Your Query"](#)
- [Section 1.3.4.5, "Building Your FLWR Statement Graphically"](#)
- [Section 1.3.4.6, "Add a Parameter"](#)
- [Section 1.3.4.7, "Map Elements to the Return Type"](#)
- [Section 1.3.4.8, "Populating the Return Clause"](#)
- [Section 1.3.4.9, "Set Statement Scoping"](#)
- [Section 1.3.4.10, "Creating Joins - the Where Clauses"](#)
- [Section 1.3.4.11, "Associate a Parameter with a For Node"](#)

A logical data service can be thought of as a "virtual" data source. Logical data services are built upon existing physical or logical data services.

Note: The Oracle Data Service Integrator Retail Sample Application is a good source for best practices associated with creating layered data services.

To create a logical data service:

1. Right-click on the folder named **logical** that you previously created.
 - New > Logical Data Service**
2. Set the data service name to:
 - CUST_ORDERS_ITEMS**
3. Click **Finish**.

After making these selections, your new entity data service appears in **Overview** mode.

Since no functions have yet been added to your data service, the work area of the data service is empty.

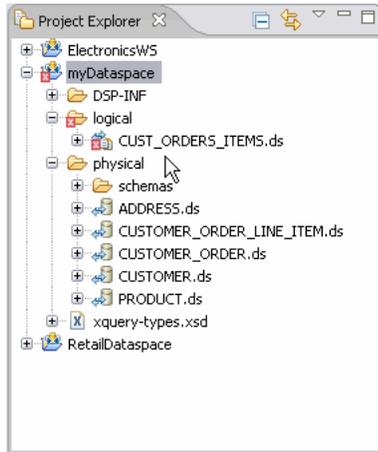
Options available for creating and testing your new data service appear at the bottom of the workspace. In addition to **Overview**, you will see the following tabs:

- **Query Map**
- **Update Map**
- **Plan**
- **Test**
- **Source**

1.3.4.1 Attempt To Publish Your Dataspace Project

There are times when attempts to publish your data service under development will not be successful. This is expected since as you create your query in the Query Map, source is created simultaneously. (When a data service is in such a state, you will notice a red x on its associated icon in Project Explorer.)

Figure 1–17 Project After Unsuccessful Publish Effort



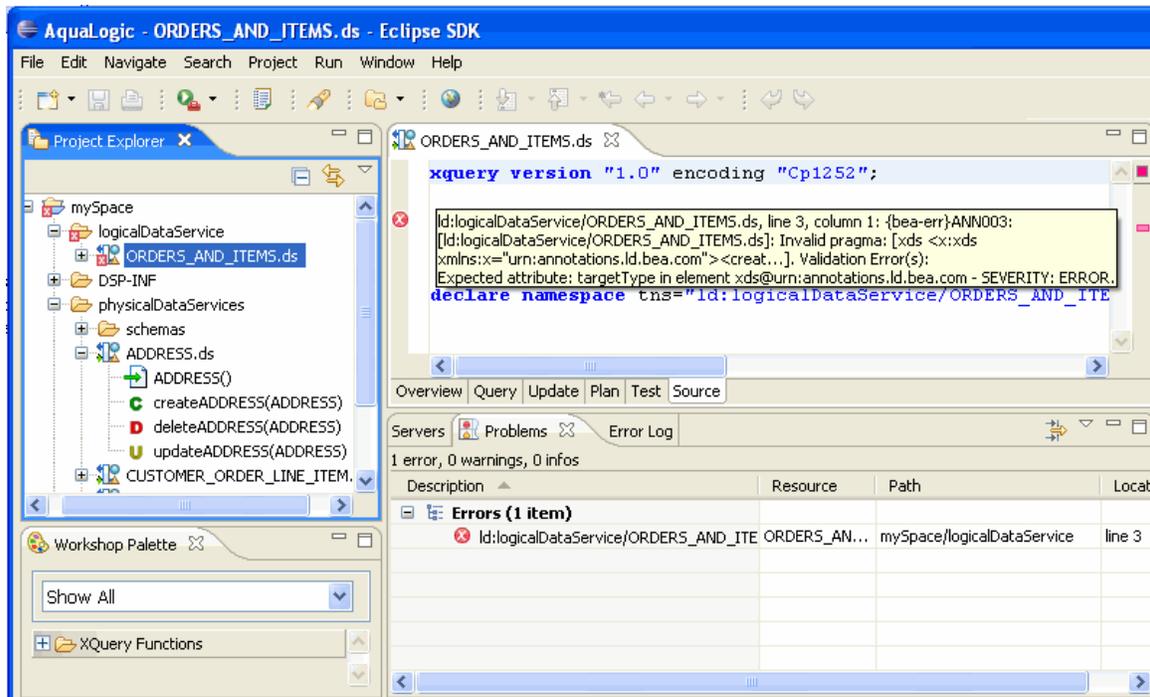
A red x on the icon beside `CUST_ORDERS_ITEMS.ds` indicates that it is invalid.

Unlike the previously successful publish operation, you will now get a message indicating that your project contains build errors and cannot be published.

In this case your newly created `CUST_ORDERS_ITEMS` data service is invalid. You can verify this several ways after clicking **OK**.

- Inspect your code by clicking on the **Source** tab.
- Double-click on the error reported in the **Problems** window.
- Inspect the contents of the **Error log** window.

Figure 1–18 Incomplete Logical Data Service Validation Error



The Problems tab shows an error message. It shows the description (including Id), resource, path, and location of the error in the code.

Although an error condition exists, you can continue creating on your data service.

1.3.4.2 Bottom Up or Top Down

Data services can be designed from the top-down or bottom-up. The following table compares these two approaches.

Table 1–3 Data Service Design Models

Data Service Design Model	Description
Top-down	The new data service is based on an existing XML Type (schema) that is either drawn from an existing data service or developed externally.
Bottom-up	The new data service is created by: <ul style="list-style-type: none"> ■ Identifying one or more data sources. ■ Building up a Return type in the Query Map. ■ Saving your data service and associating it with the schema created from the newly designed Return type.

This tutorial uses a bottom-up design.

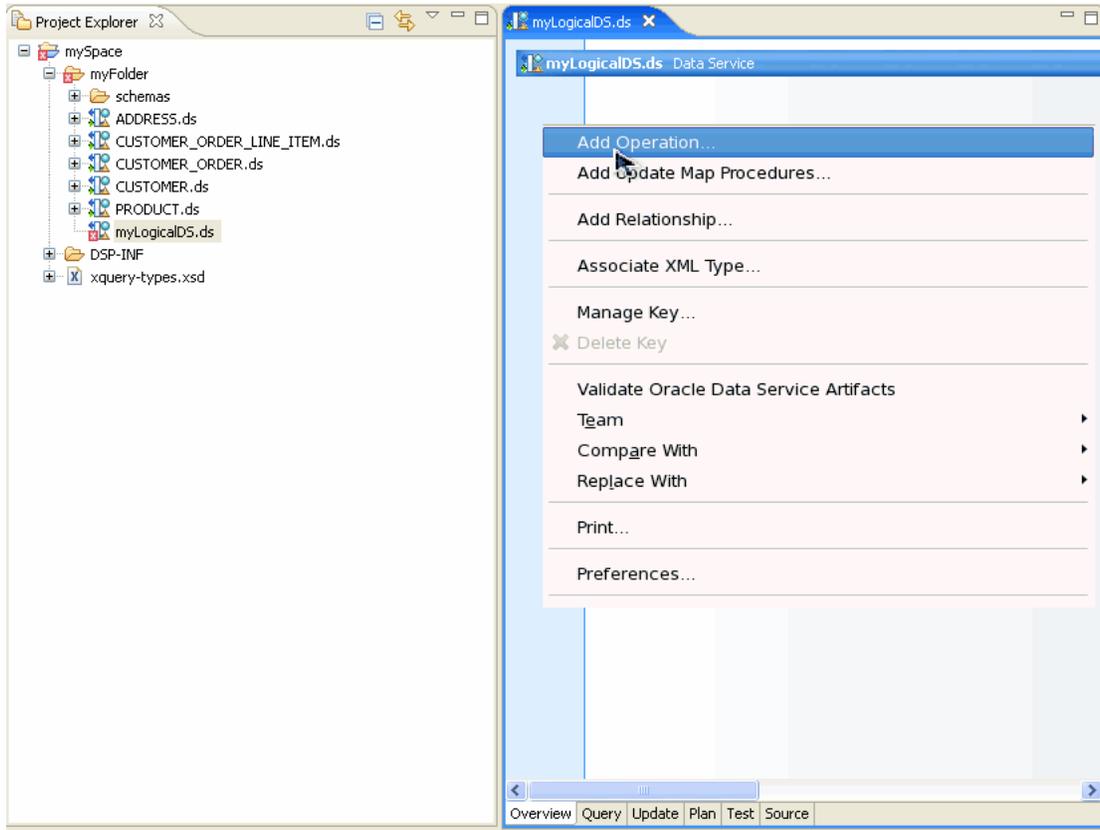
1.3.4.3 Add an Operation to CUST_ORDERS_ITEMS

The next step is to add a read function to your new data service that will return a document containing all the orders placed by a particular customer, and all the items in each order.

To add your new function:

1. Select the **Overview** tab.
2. Right-click in the **CUST_ORDERS_ITEMS** data service's work area.
3. Choose **Add Operation...** from available options.

Figure 1–19 *Creating a New Operation*



In the myLogical DS data service workspace, Add Function is selected from the context menu.

This figure has nothing to do with the previous text.

The next steps will create a publicly available Read function for your new data service.

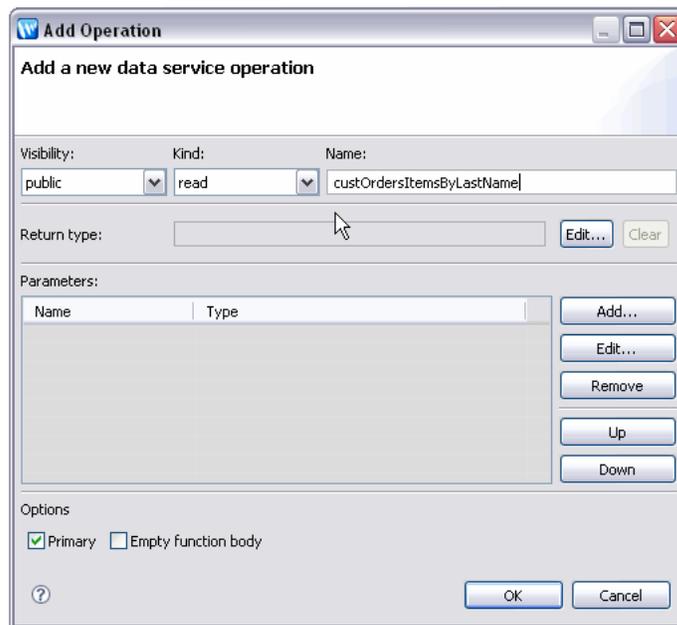
Table 1–4 *Add Operation Dialog Options*

Step	Option	Action	Comment/Reference
	Visibility		Options are private (internal to data service), protected (from public), and public. Default setting is public.
	Kind		All operations are functions other than library procedures. The Read function simply retrieves information from your data source. Default operation is read.
1.	Name	custOrdersItemsByLastName	Any valid XML name can be entered; spaces are not allowed.

Table 1–4 (Cont.) Add Operation Dialog Options

Step	Option	Action	Comment/Reference
	Return Type:		Bottom-up designs of a data service create the Return type in the Query Map.
	Parameters:		Can be added here or in the Query Map. Leave unselected.
	Options: Primary		Defines function as the Primary Read function in the entity data service. Default is selected.
	Options: Empty Function Body		Default is not selected.
2.		Click OK .	

Figure 1–20 Add Operation Dialog

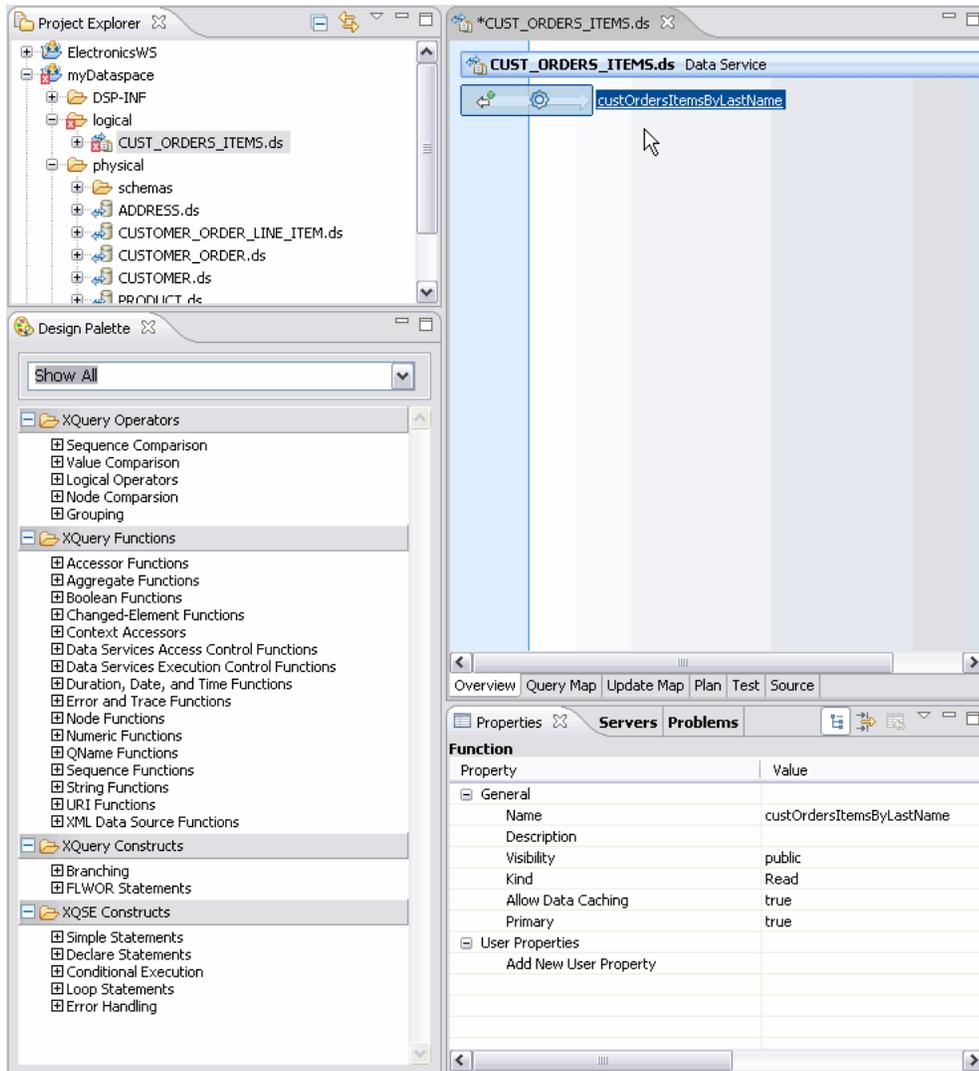


The Add Operation dialog adds a new data service operation. Here, Visibility is set to public, Kind to read, Name to custOrdersItemsByLastName. No parameters are set. The Options are set to Primary and Empty function body is deselected.

Every artifact and artifact element in Overview has properties. In some cases these properties — such as name and type — are either directly editable or adjustable through dropdown list boxes. The Properties window is, by default, visible in the Eclipse for WebLogic perspective. If the Properties window is not visible you can retrieve it using the command:

Window > Reset Perspective

Figure 1–21 New Data Service Operation and Properties

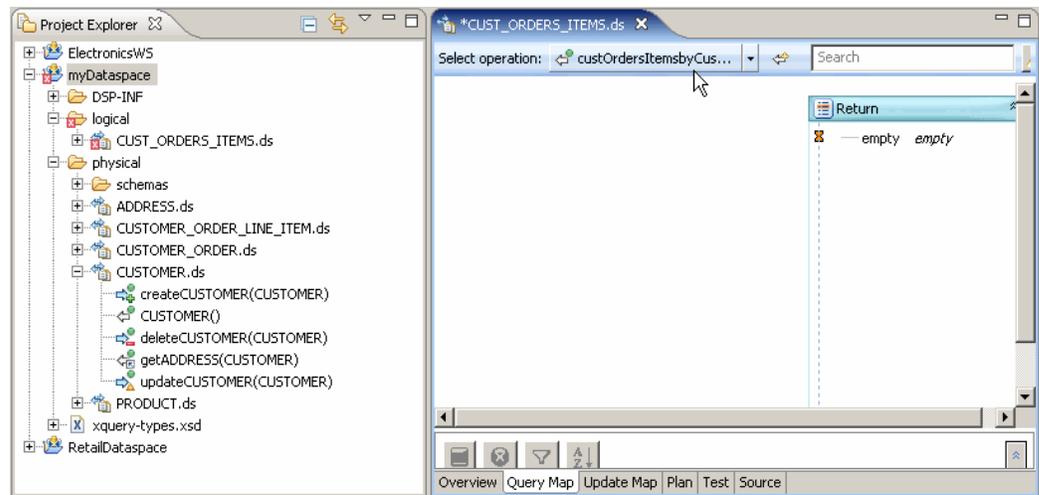


The Project Explorer, Design Palette, Properties, and CUST_ORDERS_ITEMS.ds panels are shown. In the Project Explorer, CUST_ORDERS_ITEMS.ds is selected. The properties set at the previous step are shown. In the CUST_ORDERS_ITEMS.ds Data Service work area, the custOrdersItemsByLastName function is shown.

1.3.4.4 Building Your Query

Click on the `custOrdersItemsByLastName` function name in the work area to enter Query Map mode

Figure 1–22 Initial Query View



In the CUST_ORDERS_ITEMS.ds workspace, the Return panel in Query Map mode shows an **empty empty** placeholder element.

Changes made in the Query Map editor are immediately reflected in source and vice-versa. When there is an error in source, the Query map may not be available. You can typically correct such a condition using the **Undo** menu option or **Ctrl-Z**. Alternatively, click the Source tab and edit as needed.

1.3.4.5 Building Your FLWR Statement Graphically

XQueries are often described as being build upon "FLWR" statements:

- For/Let
- Where
- Return

Changes made in source are immediately rendered graphically in the query map.

1.3.4.5.1 Adding Data Sources to Query View - the For/Let Statements It is through the Query Map that you can bring together representations of existing data sources and associate their elements with the Return type of a new data service.

In the current example your new data service is to provide a consolidated view drawn from the CUSTOMER, CUSTOMER_ORDER, and CUSTOMER_ORDER_LINE_ITEM data services. The Read functions from these physical data services therefore need to be represented in the work area of the new data service.

Follow these steps to add these representations to your Query map:

1. In the physical folder expand the following data services:
 - CUSTOMER.ds
 - CUSTOMER_ORDER.ds
 - CUSTOMER_ORDER_LINE_ITEM.ds
2. Drag and drop the **Read** operations of the following data services **CUSTOMER**, **CUSTOMER_ORDER**, and **CUSTOMER_ORDER_LINE_ITEM** into the query

work area. Read operations are identified by the a white-arrow-with-green-ball icon as shown below.

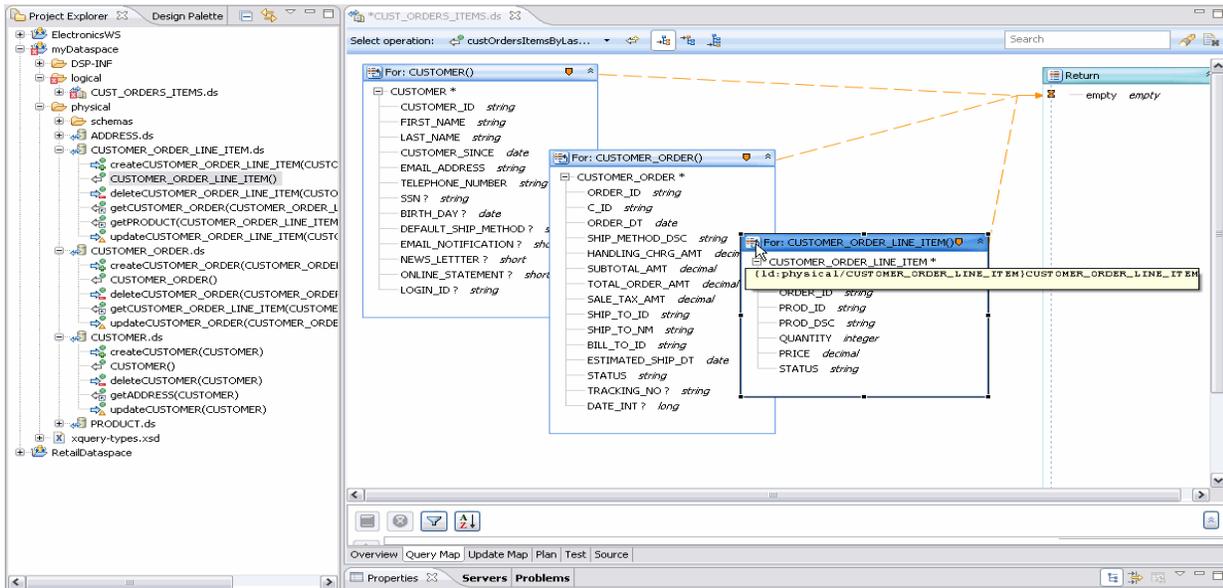
Figure 1–23 Read Function Icon



Read function icon, white arrow with green ball, next to CUSTOMER_ORDER().

Each of these operational building blocks will become **for** statements in the XQuery description of your new data service.

Figure 1–24 Data Source Representations in Work Area



The Data Source Representations in Work Area graphic shows the artifacts useful in tailoring your query:

- Data sources are represented in three XQuery For: statements.
- The 'empty empty' element in the Return type is a placeholder for the elements and their type that will eventually be projected.
- The lines from the three statements to the empty global element in the Return type represents current scopings. By adjusting these lines when a Return type is populated you can alter the arrangement of information returned by your query. (Described below.)

1.3.4.6 Add a Parameter

Parameters can be added when your operation is created or in the Query Map. Parameters can be of simple (primitive) type or complex, such as the XML type from another data service.

In this case you create a single `xs:string` parameter that will allow retrieval of one or more records by a customer's last name.

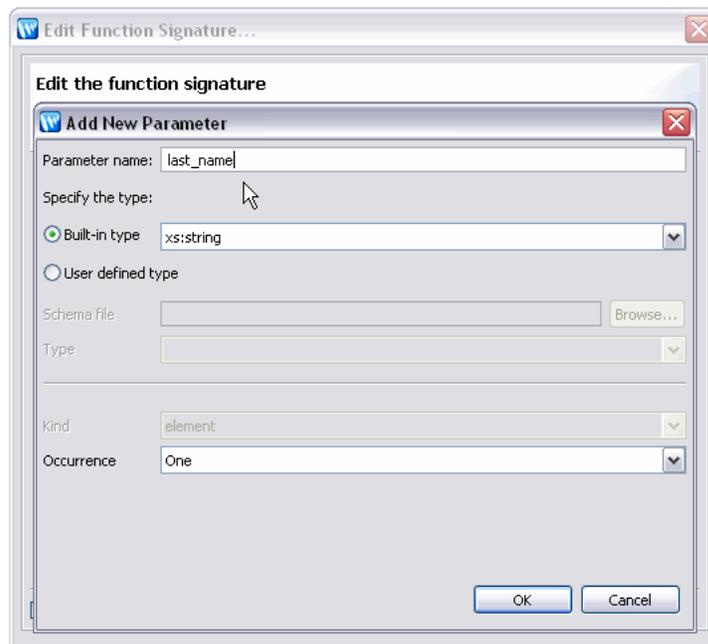
To add a parameter:

1. In the Query Map work area right-click in a blank area and select:
Edit Signature...
2. If asked to save modified resources click **OK**.
3. If asked to save modified resources click
4. Complete the **Edit Function Signature...** dialog.

Table 1–5 Edit Function Signature Dialog Options

Step	Field	Action	Comment/Reference
1.	Parameter name	last_name	
	Parameter type		xs:string is the default primitive type.
	Occurrence		Default is One.
2.		Click OK	
3.		Click OK	In the Edit Function Signatures dialog.

Figure 1–25 Add New Parameter Dialog



In the Add New Parameter dialog, the Parameter name is last_name, the Built-in type is xs.string, and the Occurrence is One.

The last_name parameter appears in the work area.

Example 1–1 Partial Source of CUST_ORDERS_ITEMS After Addition of Read Functions and last_name Parameter

```
xquery version "1.0" encoding "UTF-8";

(:: pragma ... ::)
```

```

declare namespace cus2= "ld:physical/CUSTOMER";
declare namespace cus1= "ld:physical/CUSTOMER_ORDER";
declare namespace ust= "custOrdersItems";
declare namespace cus= "ld:physical/CUSTOMER_ORDER_LINE_ITEM";
declare namespace tns="ld:logical/CUST_ORDERS_ITEMS";

(:: pragma ... ::)

declare function tns:custOrdersItemsByLastName(){
for $CUSTOMER in cus:CUSTOMER()
for $CUSTOMER_ORDER in cus1:CUSTOMER_ORDER()
for $CUSTOMER_ORDER_LINE_ITEM in cus2:CUSTOMER_ORDER_LINE_ITEM()
return
    ()
};

```

1.3.4.7 Map Elements to the Return Type

Three icons associated with projecting elements to the Return type appear above the Query Map work area. (You may need to widen your window to see all three icons.)

Table 1–6 Mapping Mode Icons

Icon	Mapping Mode	Keyboard Equivalent	Description
	Value	None	Maps simple or complex elements to identical values in the Return type. For example, a simple element can be projected to a comparable simple element in the Return type.
	Overwrite	Ctrl-Drag object	Overwrites simple or complex element in the Return type with the selected simple or complex element.
	Append	Ctrl-Shift-Drag object	Maps simple or complex object as a child to the Return type element it is associated with.

You will use these options to map representations of source data to the Return type of your new data service.

1.3.4.8 Populating the Return Clause

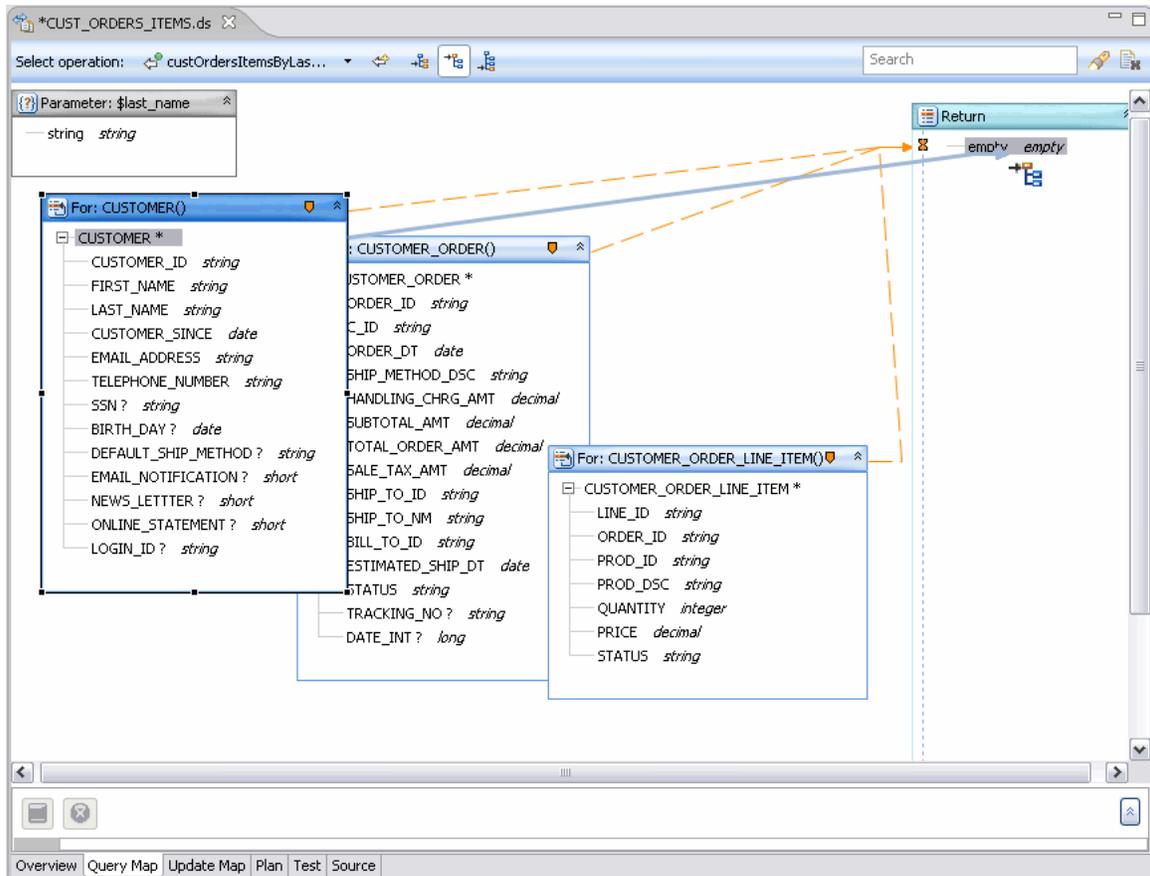
From the three mapping icons in the Select operation line at the top of the query map, select the second of the three icons, Overwrite mapping:

Drag the **CUSTOMER** complex element

CUSTOMER*

over the global element placeholder labeled "**empty**" in the Return type.

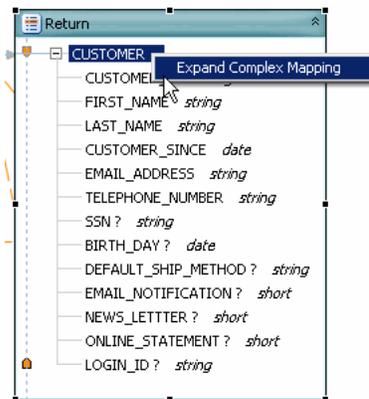
Figure 1–26 Mapping Complex Element to Return Type



This figure shows mapping complex element types to Return types by dragging.

Right-click on the new **CUSTOMER** element in the **Return** type and select:
Expand Complex Mapping

Figure 1–27 Expanding Complex Mapping



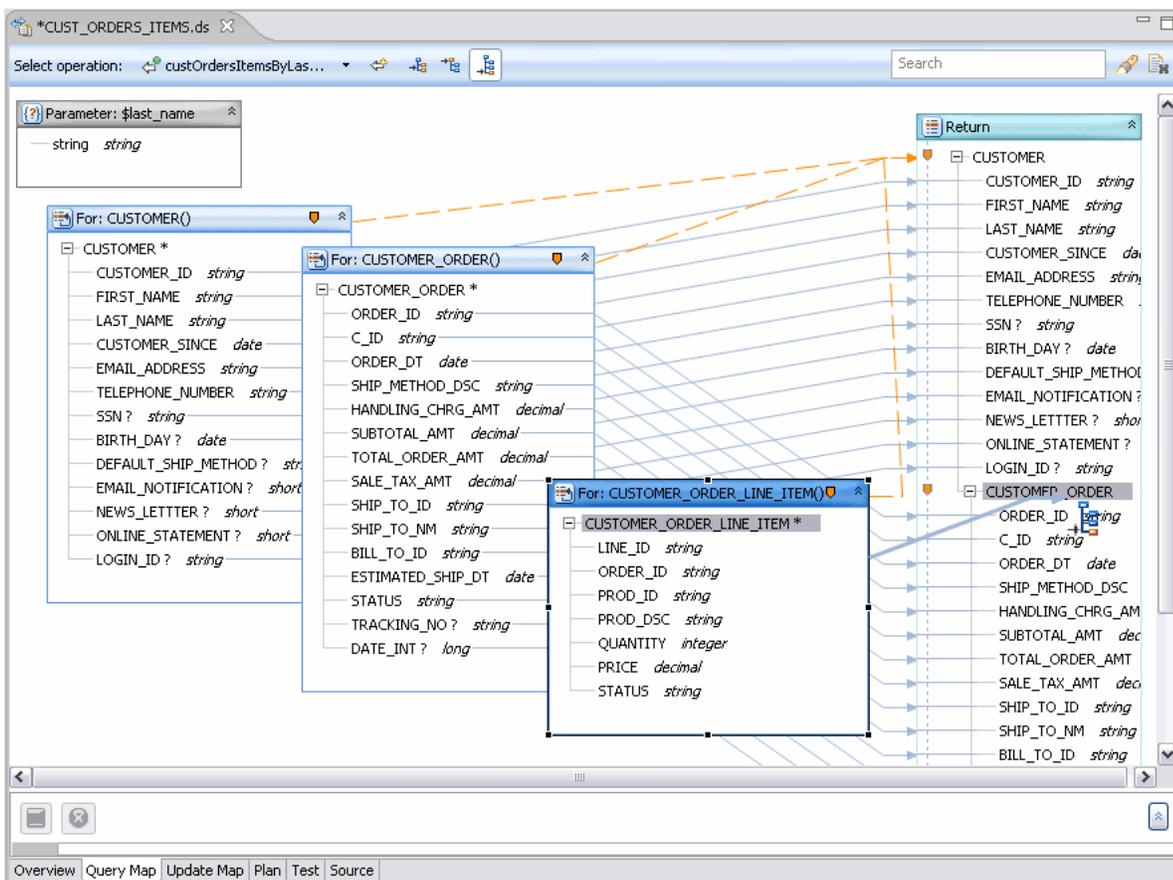
Expanding the complex mapping reveals contents of the **CUSTOMER** element in the **Return** type. The contents include **FIRST_NAME**, **LAST_NAME**, **CUSTOMER_SINCE**, **EMAIL_ADDRESS**, **TELEPHONE_NUMBER**, **SSN**, **BIRTH_DAY**, **DEFAULT_SHIP_**

METHOD, EMAIL_NOTIFICATION, NEWS_LETTER, ONLINE_STATEMENT, LOGIN_ID.

This gesture is a shortcut for drawing lines from each element for the statement into the **Return** type. This gesture is also necessary if you want to add a complex child element to the type. Notice that individual mapping lines now connect each element in the For: node with an element in the **Return** type. Individual mappings can be added or deleted using drag-and-drop or the **Delete** key, respectively. The next steps will add elements from the **CUSTOMER_ORDER** data service to your Return type.

1. Select **Append Mapping** mode.
2. Drag the **CUSTOMER_ORDER** complex element **CUSTOMER_ORDER*** over the **CUSTOMER** element in the Return type. Notice that the **CUSTOMER_ORDER** global element and the names of its children now appear after the **CUSTOMER** elements.
3. Expand complex mapping for the **CUSTOMER_ORDER** global element.
4. From the work area drag the **CUSTOMER_ORDER_LINE_ITEM** complex element over the **CUSTOMER_ORDER** element in the Return type.
5. Expand complex mapping for these elements.

Figure 1–28 Adding Child Elements to Return Type



An arrow between the expanded **CUSTOMER_ORDER_LINE_ITEM** complex element and the expanded **CUSTOMER_ORDER** element in the Return type show mapping for

these elements. Similar mappings are shown for CUSTOMER() and CUSTOMER_ORDER().

1.3.4.9 Set Statement Scoping

Click the **Source** tab to inspect your generated code. Notice that the **Return** type contains all three For: statements.

Example 1–2 Function cust_orders_items_byLastName(string) in Source View

```
declare function tns:custOrdersItemsByLastName($last_name as xs:string) {
for $CUSTOMER in cus:CUSTOMER()
for $CUSTOMER_ORDER in cus1:CUSTOMER_ORDER()
for $CUSTOMER_ORDER_LINE_ITEM in cus2:CUSTOMER_ORDER_LINE_ITEM()
return
  <ust:CUSTOMER>
    <CUSTOMER_ID>{fn:data($CUSTOMER/CUSTOMER_ID)}</CUSTOMER_ID>
    <FIRST_NAME>{fn:data($CUSTOMER/FIRST_NAME)}</FIRST_NAME>
    <LAST_NAME>{fn:data($CUSTOMER/LAST_NAME)}</LAST_NAME>
    <CUSTOMER_SINCE>{fn:data($CUSTOMER/CUSTOMER_SINCE)}</CUSTOMER_SINCE>
    <EMAIL_ADDRESS>{fn:data($CUSTOMER/EMAIL_ADDRESS)}</EMAIL_ADDRESS>
    <TELEPHONE_NUMBER>{fn:data($CUSTOMER/TELEPHONE_NUMBER)}</TELEPHONE_NUMBER>
    <SSN?>{fn:data($CUSTOMER/SSN)}</SSN>
    <BIRTH_DAY?>{fn:data($CUSTOMER/BIRTH_DAY)}</BIRTH_DAY>
    <DEFAULT_SHIP_METHOD?>{fn:data($CUSTOMER/DEFAULT_SHIP_METHOD)}</DEFAULT_SHIP_METHOD>
    <EMAIL_NOTIFICATION?>{fn:data($CUSTOMER/EMAIL_NOTIFICATION)}</EMAIL_NOTIFICATION>
    <NEWS_LETTTER?>{fn:data($CUSTOMER/NEWS_LETTTER)}</NEWS_LETTTER>
    <ONLINE_STATEMENT?>{fn:data($CUSTOMER/ONLINE_STATEMENT)}</ONLINE_STATEMENT>
    <LOGIN_ID?>{fn:data($CUSTOMER/LOGIN_ID)}</LOGIN_ID>
  {
    <ust:CUSTOMER_ORDER>
      <ORDER_ID>{fn:data($CUSTOMER_ORDER/ORDER_ID)}</ORDER_ID>
      <C_ID>{fn:data($CUSTOMER_ORDER/C_ID)}</C_ID>
      <ORDER_DT>{fn:data($CUSTOMER_ORDER/ORDER_DT)}</ORDER_DT>
      <SHIP_METHOD_DSC>{fn:data($CUSTOMER_ORDER/SHIP_METHOD_DSC)}</SHIP_METHOD_DSC>
      <HANDLING_CHRG_AMT>{fn:data($CUSTOMER_ORDER/HANDLING_CHRG_AMT)}
      </HANDLING_CHRG_AMT>
      <SUBTOTAL_AMT>{fn:data($CUSTOMER_ORDER/SUBTOTAL_AMT)}</SUBTOTAL_AMT>
      <TOTAL_ORDER_AMT>{fn:data($CUSTOMER_ORDER/TOTAL_ORDER_AMT)}</TOTAL_ORDER_AMT>
      <SALE_TAX_AMT>{fn:data($CUSTOMER_ORDER/SALE_TAX_AMT)}</SALE_TAX_AMT>
      <SHIP_TO_ID>{fn:data($CUSTOMER_ORDER/SHIP_TO_ID)}</SHIP_TO_ID>
      <SHIP_TO_NM>{fn:data($CUSTOMER_ORDER/SHIP_TO_NM)}</SHIP_TO_NM>
      <BILL_TO_ID>{fn:data($CUSTOMER_ORDER/BILL_TO_ID)}</BILL_TO_ID>
      <ESTIMATED_SHIP_DT>{fn:data($CUSTOMER_ORDER/ESTIMATED_SHIP_DT)}
      </ESTIMATED_SHIP_DT>
      <STATUS>{fn:data($CUSTOMER_ORDER/STATUS)}</STATUS>
      <TRACKING_NO?>{fn:data($CUSTOMER_ORDER/TRACKING_NO)}</TRACKING_NO>
      <DATE_INT?>{fn:data($CUSTOMER_ORDER/DATE_INT)}</DATE_INT>
    {
      <ust:CUSTOMER_ORDER_LINE_ITEM>
        <LINE_ID>{fn:data($CUSTOMER_ORDER_LINE_ITEM/LINE_ID)}</LINE_ID>
        <ORDER_ID>{fn:data($CUSTOMER_ORDER_LINE_ITEM/ORDER_ID)}</ORDER_ID>
        <PROD_ID>{fn:data($CUSTOMER_ORDER_LINE_ITEM/PROD_ID)}</PROD_ID>
        <PROD_DSC>{fn:data($CUSTOMER_ORDER_LINE_ITEM/PROD_DSC)}</PROD_DSC>
        <QUANTITY>{fn:data($CUSTOMER_ORDER_LINE_ITEM/QUANTITY)}</QUANTITY>
        <PRICE>{fn:data($CUSTOMER_ORDER_LINE_ITEM/PRICE)}</PRICE>
        <STATUS>{fn:data($CUSTOMER_ORDER_LINE_ITEM/STATUS)}</STATUS>
      </ust:CUSTOMER_ORDER_LINE_ITEM>
    }
  }
}
```

```

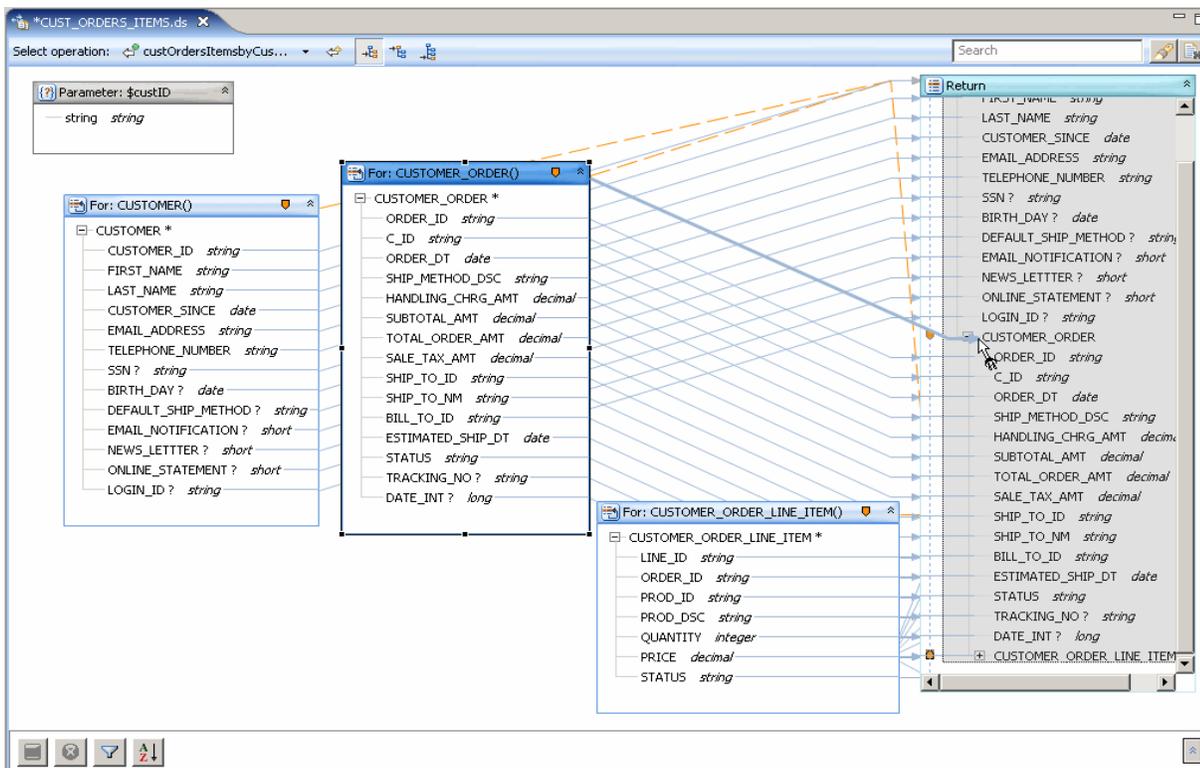
    }
  }
</ust:CUSTOMER>
};

```

Note: The current query is — in relational terminology — a cross-product or a *Cartesian join*. Such queries when run are very CPU intensive. In the case of this example, scoping and joining should occur before the query is run.

Using the **Query Map** you can adjust this quite easily by changing the scoping of the subordinate data services in the **Return** type, as shown in the following steps.

Figure 1–29 Adjusting Scoping Rules in the Return Type



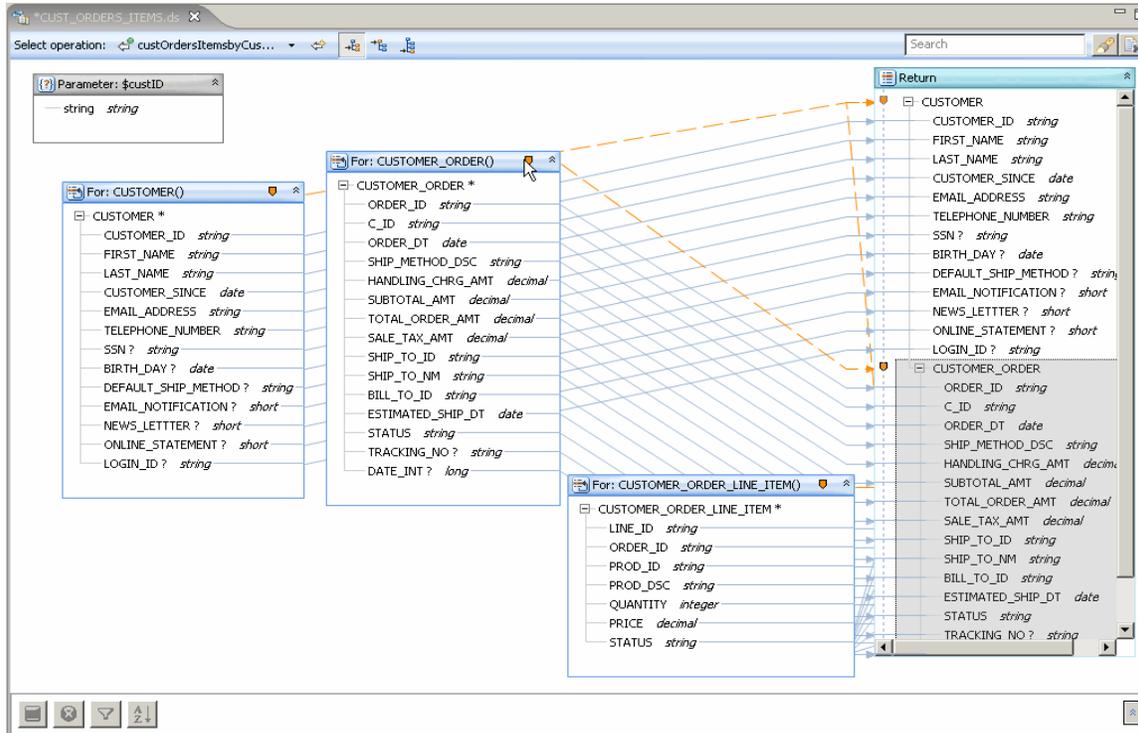
In the Query Map mode, there is a zone icon on the For: \$CUSTOMER_ORDER() node. There is an arrow between the \$CUSTOMER_ORDER() node and the CUSTOMER_ORDER element in the Return type. After the icon is dragged to the corresponding Return element, the icon appears next to the element name. Similarly, a zone icon appears on the For:CUSTOMER() and For:CUSTOMER_ORDER_LINE_ITEM() nodes.

1. Return to **Query Map** mode.
2. With your mouse select the zone icon in the node:
For: \$CUSTOMER_ORDER()
3. Drag the zone icon over the corresponding **CUSTOMER_ORDER** element in the **Return** type.

Notice that the zone line from the **CUSTOMER_ORDER** node moves to the subordinate complex type (**CUSTOMER_ORDER**).

4. Drag the zone icon of **CUSTOMER_ORDER_LINE_ITEM** to its corresponding element in the **Return** type.

Figure 1–30 Nested Zoning in the Return Type



The zone icon appears on the For:\$CUSTOMER_ORDER_LINE_ITEM node. Drag the icon to its corresponding Return element, which is not shown here.

Switch to **Source** view to verify that the for statements are nested in the **Return** clause. Now, when a parameter is passed with the operation, all the customers with a particular last name will be returned which contains orders and order line items associated with that customer.

Example 1–3 Source View of Return Type with Nested Return Types

```

declare function tns:custOrdersItemsByLastName($last_name as xs:string) {
for $CUSTOMER in cus:CUSTOMER()

return
  <ust:CUSTOMER>
    <CUSTOMER_ID>{fn:data($CUSTOMER/CUSTOMER_ID)}</CUSTOMER_ID>
    <FIRST_NAME>{fn:data($CUSTOMER/FIRST_NAME)}</FIRST_NAME>
    <LAST_NAME>{fn:data($CUSTOMER/LAST_NAME)}</LAST_NAME>
    <CUSTOMER_SINCE>{fn:data($CUSTOMER/CUSTOMER_SINCE)}</CUSTOMER_SINCE>
    <EMAIL_ADDRESS>{fn:data($CUSTOMER/EMAIL_ADDRESS)}</EMAIL_ADDRESS>
    <TELEPHONE_NUMBER>{fn:data($CUSTOMER/TELEPHONE_NUMBER)}</TELEPHONE_NUMBER>
    <SSN?>{fn:data($CUSTOMER/SSN)}</SSN>
    <BIRTH_DAY?>{fn:data($CUSTOMER/BIRTH_DAY)}</BIRTH_DAY>
    <DEFAULT_SHIP_METHOD?>{fn:data($CUSTOMER/DEFAULT_SHIP_METHOD)}</DEFAULT_SHIP_METHOD>

```

```

<EMAIL_NOTIFICATION?>{fn:data($CUSTOMER/EMAIL_NOTIFICATION)}</EMAIL_NOTIFICATION>
<NEWS_LETTTER?>{fn:data($CUSTOMER/NEWS_LETTTER)}</NEWS_LETTTER>
<ONLINE_STATEMENT?>{fn:data($CUSTOMER/ONLINE_STATEMENT)}</ONLINE_STATEMENT>
<LOGIN_ID?>{fn:data($CUSTOMER/LOGIN_ID)}</LOGIN_ID>
{
  for $CUSTOMER_ORDER in cus1:CUSTOMER_ORDER()
  return
  <ust:CUSTOMER_ORDER>
    <ORDER_ID>{fn:data($CUSTOMER_ORDER/ORDER_ID)}</ORDER_ID>
    <C_ID>{fn:data($CUSTOMER_ORDER/C_ID)}</C_ID>
    <ORDER_DT>{fn:data($CUSTOMER_ORDER/ORDER_DT)}</ORDER_DT>
    <SHIP_METHOD_DSC>{fn:data($CUSTOMER_ORDER/SHIP_METHOD_DSC)}</SHIP_METHOD_DSC>
    <HANDLING_CHRG_AMT>{fn:data($CUSTOMER_ORDER/HANDLING_CHRG_AMT)}
      </HANDLING_CHRG_AMT>
    <SUBTOTAL_AMT>{fn:data($CUSTOMER_ORDER/SUBTOTAL_AMT)}</SUBTOTAL_AMT>
    <TOTAL_ORDER_AMT>{fn:data($CUSTOMER_ORDER/TOTAL_ORDER_AMT)}</TOTAL_ORDER_AMT>
    <SALE_TAX_AMT>{fn:data($CUSTOMER_ORDER/SALE_TAX_AMT)}</SALE_TAX_AMT>
    <SHIP_TO_ID>{fn:data($CUSTOMER_ORDER/SHIP_TO_ID)}</SHIP_TO_ID>
    <SHIP_TO_NM>{fn:data($CUSTOMER_ORDER/SHIP_TO_NM)}</SHIP_TO_NM>
    <BILL_TO_ID>{fn:data($CUSTOMER_ORDER/BILL_TO_ID)}</BILL_TO_ID>
    <ESTIMATED_SHIP_DT>{fn:data($CUSTOMER_ORDER/ESTIMATED_SHIP_DT)}
      </ESTIMATED_SHIP_DT>
    <STATUS>{fn:data($CUSTOMER_ORDER/STATUS)}</STATUS>
    <TRACKING_NO?>{fn:data($CUSTOMER_ORDER/TRACKING_NO)}</TRACKING_NO>
    <DATE_INT?>{fn:data($CUSTOMER_ORDER/DATE_INT)}</DATE_INT>
    {
      for $CUSTOMER_ORDER_LINE_ITEM in cus2:CUSTOMER_ORDER_LINE_ITEM()
      return
      <ust:CUSTOMER_ORDER_LINE_ITEM>
        <LINE_ID>{fn:data($CUSTOMER_ORDER_LINE_ITEM/LINE_ID)}</LINE_ID>
        <ORDER_ID>{fn:data($CUSTOMER_ORDER_LINE_ITEM/ORDER_ID)}</ORDER_ID>
        <PROD_ID>{fn:data($CUSTOMER_ORDER_LINE_ITEM/PROD_ID)}</PROD_ID>
        <PROD_DSC>{fn:data($CUSTOMER_ORDER_LINE_ITEM/PROD_DSC)}</PROD_DSC>
        <QUANTITY>{fn:data($CUSTOMER_ORDER_LINE_ITEM/QUANTITY)}</QUANTITY>
        <PRICE>{fn:data($CUSTOMER_ORDER_LINE_ITEM/PRICE)}</PRICE>
        <STATUS>{fn:data($CUSTOMER_ORDER_LINE_ITEM/STATUS)}</STATUS>
      </ust:CUSTOMER_ORDER_LINE_ITEM>
    }
  </ust:CUSTOMER_ORDER>
}
</ust:CUSTOMER>
};

```

1.3.4.10 Creating Joins - the Where Clauses

Where clauses satisfy either specific conditions (such as where \$i=5) or join conditions such as:

```
where $CUSTOMER_ORDER/ORDER_ID eq $CUSTOMER_ORDER_LINE_ITEM/ORDER_ID
```

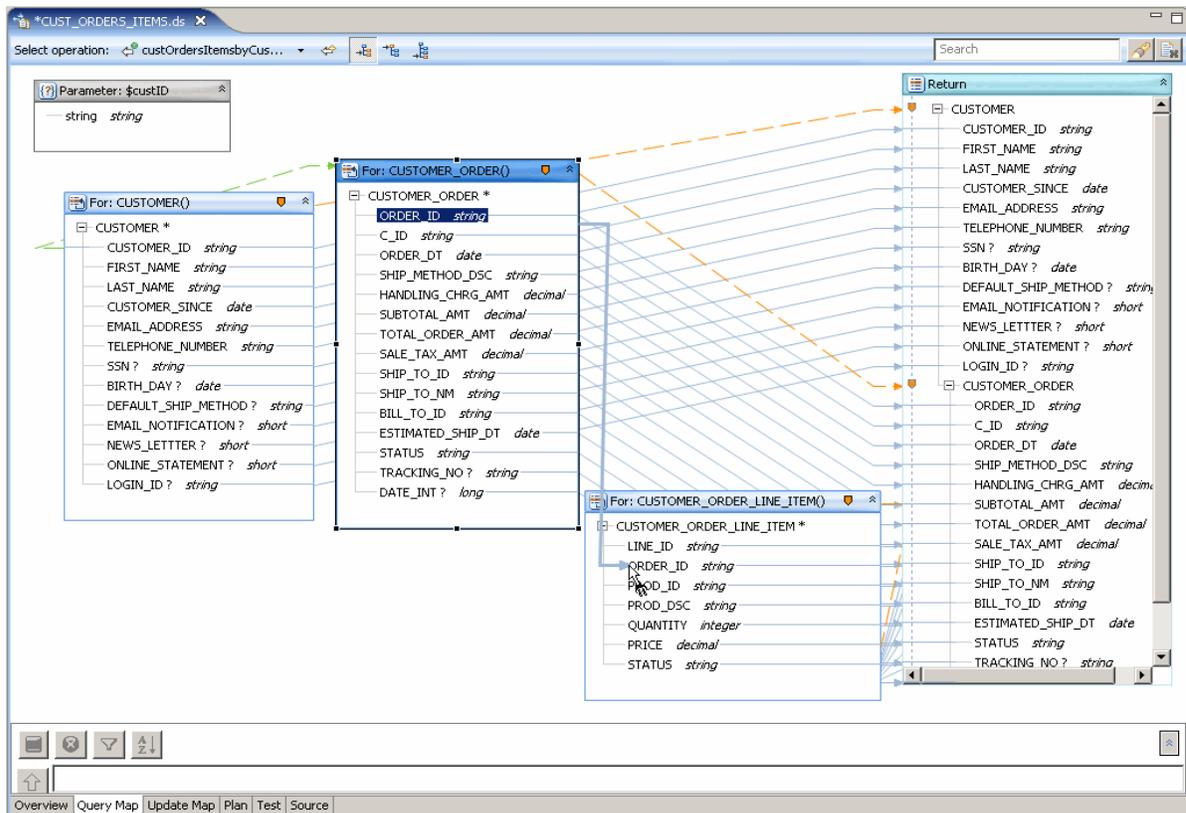
1. Return to **Query Map** mode.
2. To establish join conditions among your data sources, drag the specified element in one **For :** statement to the specified element in the target **For** statement:

Source and element	Target and element
\$CUSTOMER/CUSTOMER_ID	\$CUSTOMER_ORDER/C_ID

Source and element	Target and element
\$CUSTOMER_ORDER/ORDER_ID	\$CUSTOMER_ORDER_LINE_ITEM/ORDER_ID

Note: You may need to move the `For` nodes around in the work area to expose the elements.

Figure 1–31 Setting Up a Join Condition



In Map Query mode, there is an arrow pointing from source `CUSTOMER_ORDER/ORDER_ID` to target `CUSTOMER_ORDER_LINE_ITEM/ORDER_ID`, showing that they are joined.

You can verify your first join clause by clicking on target (`CUSTOMER_ORDER`) object. Alternatively, you can look in Source view to verify that the new where clause is modifying the `CUSTOMER_ORDER_LINE_ITEM` type.

1.3.4.11 Associate a Parameter with a For Node

An additional necessity where the condition that directs the query results to a particular customer can be created by adding a parameter to an element in a node. Parameters can be simple or complex.

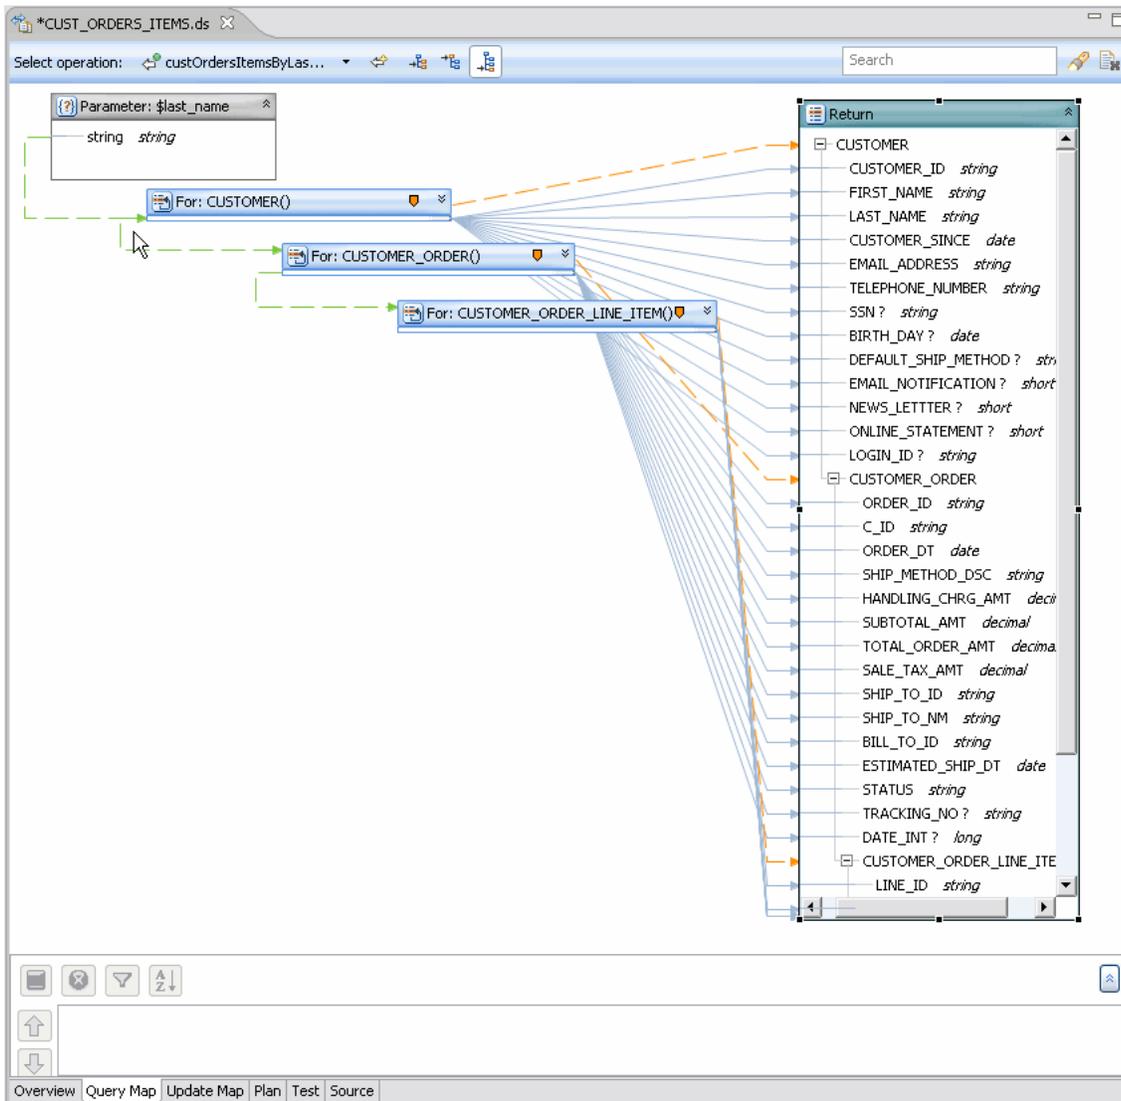
This project requires use of a single parameter: `last_name`.

In the Query Map drag the element:

string string

in the `$last_name` parameter over the `LAST_NAME` element in the `CUSTOMER` node.

Figure 1–32 Mapped Parameter and Where Clause



In Query Map mode, a line connecting the parameter to the node appears. This will also be reflected in the Query Map Expression editor when you click on the `CUSTOMER` `For:` node.

The results of this operation can also be viewed in the **Source** tab.

```
declare function tns:custOrdersItemsByLastName($last_name as xs:string) /
as element(ust1:CUST_ORDERS_ITEMS)* {
for $CUSTOMER in cus:CUSTOMER() where $last_name eq $CUSTOMER/LAST_NAME
return...
```

In **Source** you will also notice that the `for` statements now contain `where` clauses based on your graphical gestures.

```
for $CUSTOMER in cus:CUSTOMER()
where $last_name eq $CUSTOMER/LAST_NAME
```

```

return
...
for $CUSTOMER_ORDER in cus1:CUSTOMER_ORDER()
where $CUSTOMER/CUSTOMER_ID eq $CUSTOMER_ORDER/C_ID
return
...
for $CUSTOMER_ORDER_LINE_ITEM in cus2:CUSTOMER_ORDER_LINE_ITEM()
where $CUSTOMER_ORDER/ORDER_ID eq $CUSTOMER_ORDER_LINE_ITEM/ORDER_ID
return
...

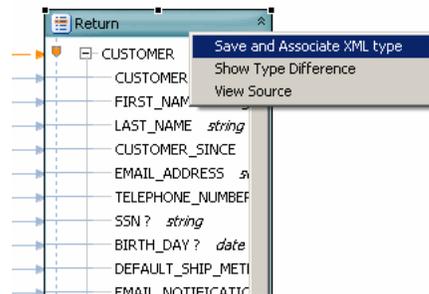
```

1.3.5 Creating, Saving, and Associating the XML Type

Since this entity data service is being created "bottom up", it is not yet associated with an XML Type (schema).

Now that you have a Return type, however, you create a valid XML Type by saving your Return type and associating it with a namespace that is unique to the project.

1. Go to **Query Map**.
2. Right-click on the **Return** type's title bar.
3. Select **Save and Associate XML type**.



Save and Associate XML Type

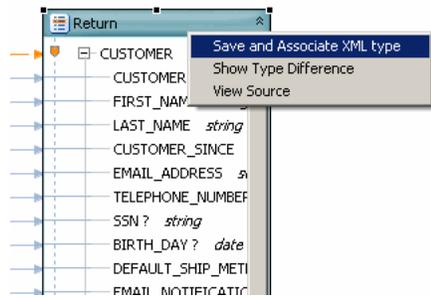
4. When asked if you want to save modified resources, choose **OK**.
5. In the **Save and Associate XML Type** dialog you will notice that the current name and namespace setting of the Return type conflicts with that of an existing type in the `CUSTOMER.xsd` file. Change the **Name of the Return type** global element from:

CUSTOMER

to:

CUST_ORDERS_ITEMS
6. Leave the **Update references** option selected. (This option — which is by default selected — means that XML Type references in source will be updated to reflect the changes you are making.)

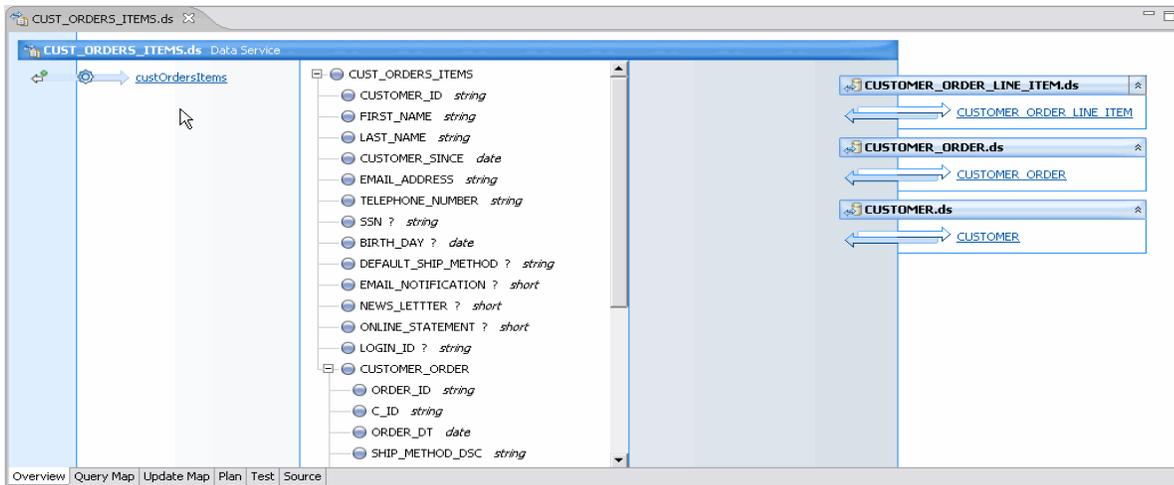
Figure 1–33 Save and Associate XML Type



The Save and Associate XML Type dialog contains Location, Namespace, and Name fields. The Update references box is checked.

7. Click **Preview**. This mode shows what changes will be performed by the name change (refactoring) operation. In this case a new schema file will be created and the target type will be renamed to **CUST_ORDERS_ITEMS**.
8. Click **OK**.
9. Notice that the target type (root element) in your Return type has been renamed.
10. Click **Overview**; you will see that your entity data service is now associated with an XML type.

Figure 1–34 Newly Associated XML Type



The root element in the Return type has been renamed CUST_ORDERS_ITEMS. The entity service is now associated with an XML type.

11. Publish your project. This operation should be successful.

1.3.5.1 Modifying the XML Type

When an XML Type is generated, complex elements by default return a single instance of their type (for example, one CUSTOMER_ORDER will be returned even if there are many).

In order to return all customer orders and all of each orders' line items minor changes to the data service's XML type are needed. The XML markup for this is:

```
maxOccurs="unbounded"
```

In other words, the element returns **n**, any number of document fragments that meet the criteria.

To modify your new CUST_ORDERS_ITEMS XML Type:

1. Click on the **Overview** tab, if it is not already selected.
2. Right-click on the topmost element in the XML type: **CUST_ORDER_ITEMS**.
3. Select **Edit Schema**. The Eclipse schema editor opens.
4. Click the schema editor's Source tab (below the editor's work area).
5. Locate the first qualified element: **CUSTOMER_ORDER**.
6. Place your cursor where you want to add the statement (just between the double-quote and the closing angle bracket (>) at the end of the line)
7. Enter a space.
8. Activate the code assistant with the combination:

Ctrl + spacebar

You will get a code completion dialog.

9. Perform the **Ctrl+space** operation twice, once for the `max_occurs`, and again to add the unbounded statement. The line now appears as:

```
<xs:element form="qualified" name="CUSTOMER_ORDER" maxOccurs="unbounded">
```

10. Follow Steps 5-9 for the second qualified element, **CUSTOMER_ORDER_LINE_ITEM**.
11. Save the **CUST_ORDERS_ITEMS.xsd** file.

File > Close

Example 1-4 CUST_ORDERS_ITEMS Schema (XSD File)

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema targetNamespace="custOrdersItems" xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CUST_ORDERS_ITEMS">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="CUSTOMER_ID" type="xs:string"/>
        <xs:element name="FIRST_NAME" type="xs:string"/>
        <xs:element name="LAST_NAME" type="xs:string"/>
        <xs:element name="CUSTOMER_SINCE" type="xs:date"/>
        <xs:element name="EMAIL_ADDRESS" type="xs:string"/>
        <xs:element name="TELEPHONE_NUMBER" type="xs:string"/>
        <xs:element name="SSN" maxOccurs="1" minOccurs="0" type="xs:string"/>
        <xs:element name="BIRTH_DAY" maxOccurs="1" minOccurs="0" type="xs:date"/>
        <xs:element name="DEFAULT_SHIP_METHOD" maxOccurs="1" minOccurs="0" type="xs:string"/>
        <xs:element name="EMAIL_NOTIFICATION" maxOccurs="1" minOccurs="0" type="xs:short"/>
        <xs:element name="NEWS_LETTER" maxOccurs="1" minOccurs="0" type="xs:short"/>
        <xs:element name="ONLINE_STATEMENT" maxOccurs="1" minOccurs="0" type="xs:short"/>
        <xs:element name="LOGIN_ID" maxOccurs="1" minOccurs="0" type="xs:string"/>
        <xs:element form="qualified" name="CUSTOMER_ORDER" maxOccurs="unbounded">
          <xs:complexType>
            <xs:sequence>
              <xs:element name="ORDER_ID" type="xs:string"/>
              <xs:element name="C_ID" type="xs:string"/>
              <xs:element name="ORDER_DT" type="xs:date"/>
            </xs:sequence>
          </xs:complexType>
        </xs:element>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

```

<xs:element name="SHIP_METHOD_DSC" type="xs:string"/>
<xs:element name="HANDLING_CHRG_AMT" type="xs:decimal"/>
<xs:element name="SUBTOTAL_AMT" type="xs:decimal"/>
<xs:element name="TOTAL_ORDER_AMT" type="xs:decimal"/>
<xs:element name="SALE_TAX_AMT" type="xs:decimal"/>
<xs:element name="SHIP_TO_ID" type="xs:string"/>
<xs:element name="SHIP_TO_NM" type="xs:string"/>
<xs:element name="BILL_TO_ID" type="xs:string"/>
<xs:element name="ESTIMATED_SHIP_DT" type="xs:date"/>
<xs:element name="STATUS" type="xs:string"/>
<xs:element name="TRACKING_NO" maxOccurs="1" minOccurs="0" type="xs:string"/>
<xs:element name="DATE_INT" maxOccurs="1" minOccurs="0" type="xs:long"/>
<xs:element form="qualified" name="CUSTOMER_ORDER_LINE_ITEM" maxOccurs="unbounded">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="LINE_ID" type="xs:string"/>
      <xs:element name="ORDER_ID" type="xs:string"/>
      <xs:element name="PROD_ID" type="xs:string"/>
      <xs:element name="PROD_DSC" type="xs:string"/>
      <xs:element name="QUANTITY" type="xs:integer"/>
      <xs:element name="PRICE" type="xs:decimal"/>
      <xs:element name="STATUS" type="xs:string"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
</xs:sequence>
</xs:complexType>
</xs:element>
</xs:schema>

```

1.3.6 Testing Your Data Service Function

Having created a parameterized read function for your logical data service, you can now test it.

1. Having created a parameterized read function for your logical data service, you can now test it.

2. Using the dropdown in the **Select** operation field, choose the function:

```
custOrdersItemsByLastName(string)
```

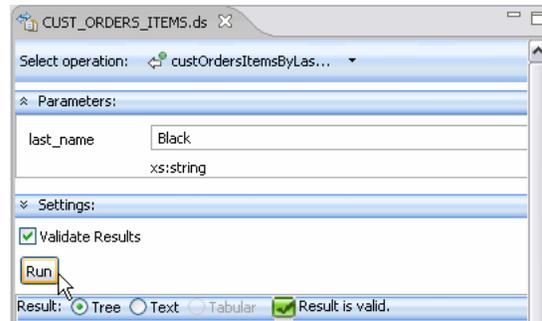
3. Enter:

Black

as the last name parameter.

Note: Entries are case-sensitive.

4. Click **Run**. Your project should be republished successfully and your data then appear.
5. Click the **+** to the left of **CUST_ORDERS_ITEMS** to view your data in Tree format. Notice that all the customer's orders are listed under customer information. If you open **CUSTOMER_ORDER** you will see that items for each order are also listed.

Figure 1–35 Testing a Parameterized Query

The dialog shows the selected operation as `custOrderItemsByLastName`. In the Parameters section, the last name is set to `Black`. In the Settings section, the `Validate Results` box is checked. Click on the `Run` button to test.

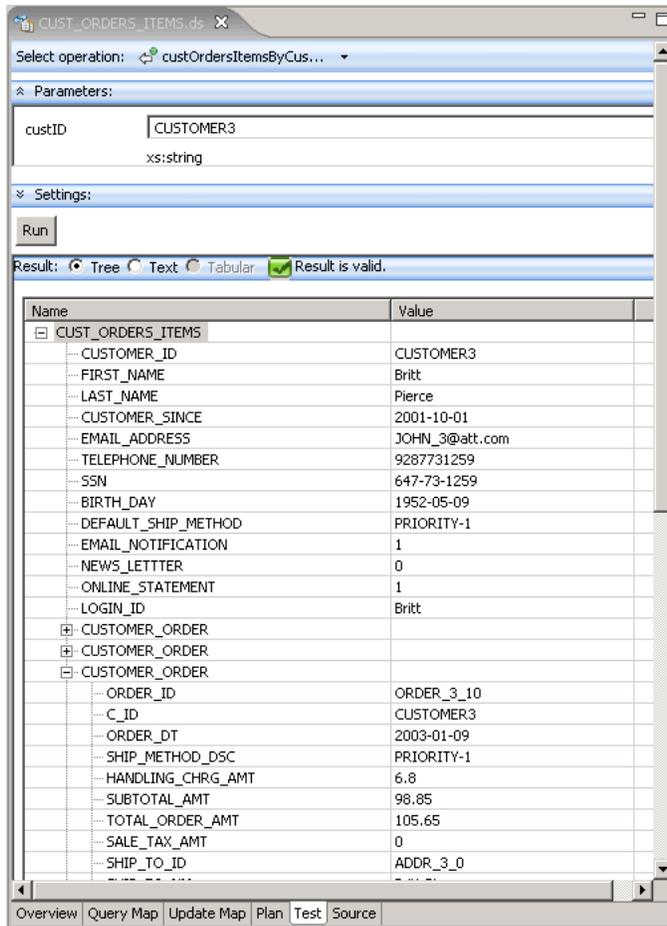
1.3.6.1 View Test Run Results

Test results from this function can be viewed in two ways:

- Tree
- Text

Note: The **Tabular** option is only available for flat (non-nested) results.

Figure 1–36 Test Run Results in Tree Style Format



Results are shown organized into a tree style format.

1.3.6.1.1 Query Statistics in the Console Window The Console window will always contain information on a successfully executed query. Access the Console with:

Window > Show View > Other... > General > Console

Click **OK**.

Sample console output is shown below.

Figure 1–37 Query Details in the Console Window



Query details are displayed in the console window. Statistics include query compilation time, query evaluation time, and operation duration. Audit event statistics are shown.

1.3.7 Adding Create-Update-Delete Functions to Your Data Service

You can also edit results in the Test area. In other words, you can update your data.

To do this an update procedure based on your data service must exist. Until then, the **Edit**, **Submit** and **Cancel** buttons at the bottom of the Test mode work area



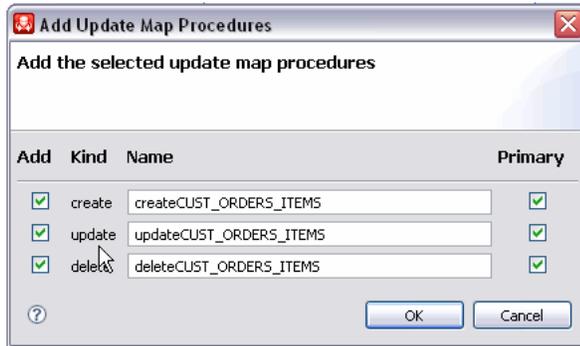
will be grayed out.

The easiest way to create an update procedure for your logical data service is to generate a default update map procedure. When you do this you will also be given the option of creating delete and insert procedures.

To add the new procedures:

1. In the Overview tab, right-click in the work area choose **Add Update Map Procedures...**

Figure 1–38 Add Update Map Procedures Dialog

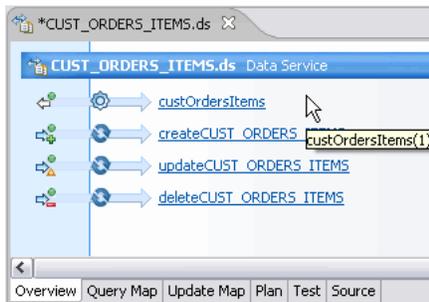


The Add Update Map Procedures dialog enables you to add the selected update map procedures. Check boxes let you select whether you want to add the procedures. There are three procedure kinds shown: create, update, and delete. In this example, for the create kind, the name is createCUST_ORDERS_ITEMS; for the update kind, the name is updateCUST_ORDERS_ITEMS; and for the delete kind, the name is deleteCUST_ORDERS_ITEMS. Check boxes let you indicate whether the selected procedure is primary.

2. Leave the default **Add and Primary** checkbox options selected for each function and click **OK**.

Notice that the procedures are added to your data service.

Figure 1–39 Update Map Procedures



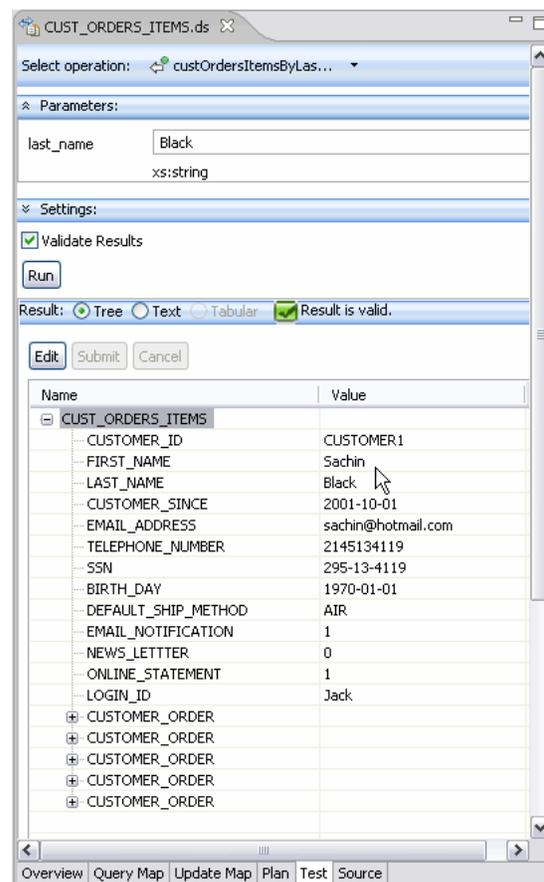
On the CUST_ORDERS_ITEMS overview panel, createCUST_ORDERS_ITEMS, updateCUST_ORDERS_ITEMS, and deleteCUST_ORDERS_ITEMS have been added to the data service.

1.3.8 Updating Your Results

Now that you have an `updateCUST_ORDERS_ITEMS` procedure, you can update data -- either through the **Test** tab or through authorized client applications. Here are the steps:

1. Click on the **Test** tab and scroll to the top of the window.
2. From the **Select operation** dropdown select the `createCUST_ORDERS_ITEMS(CUST_ORDERS_ITEMS())` operation to review the generated type.
3. From the **Select operation** dropdown select the read function `custOrdersItemsByLastName()`.
4. Run the function using **Black** as the `last_name` value.
5. Your project may need to be saved.
6. Click **Edit**.
7. Expand the top element in the **CUST_ORDERS_ITEMS** tree.
8. Change the customer's first name from **Jack** to **Sachin** using the built-in line editor. Optionally change the email address as well.
9. Click the **Submit** button at the bottom of the work area. A message indicating that your data has been successfully submitted appears.

Figure 1–40 Changing an Element in Test View



The Test tab shows how you can edit values such as first name and email address. The FIRST_NAME has been changed to Sachin and the email address has been changed to sachin@hotmail.com.

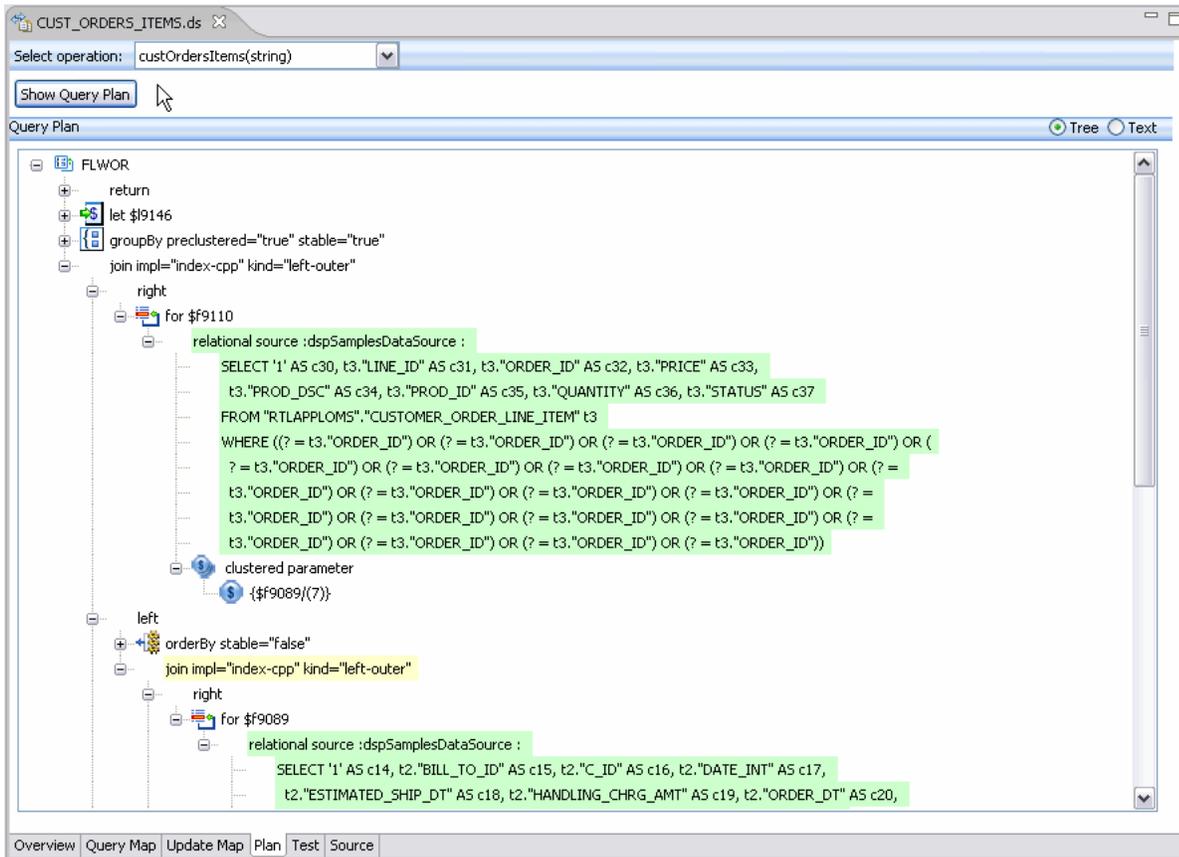
10. Re-run your function to see that the first name field reflects the changes you made.

1.3.9 Reviewing the Query Plan

After a data service has been successfully published, the query plan for the service's read functions can be examined through the Plan tab. The plan can be display in tree or text mode.

1. Click the **Plan** tab.
2. Choose the `custOrdersItemsByLastName(string)` function from the **Select operation** dropdown.
3. Click **Show Query Plan**.

Figure 1–41 Tree View of Query Plan



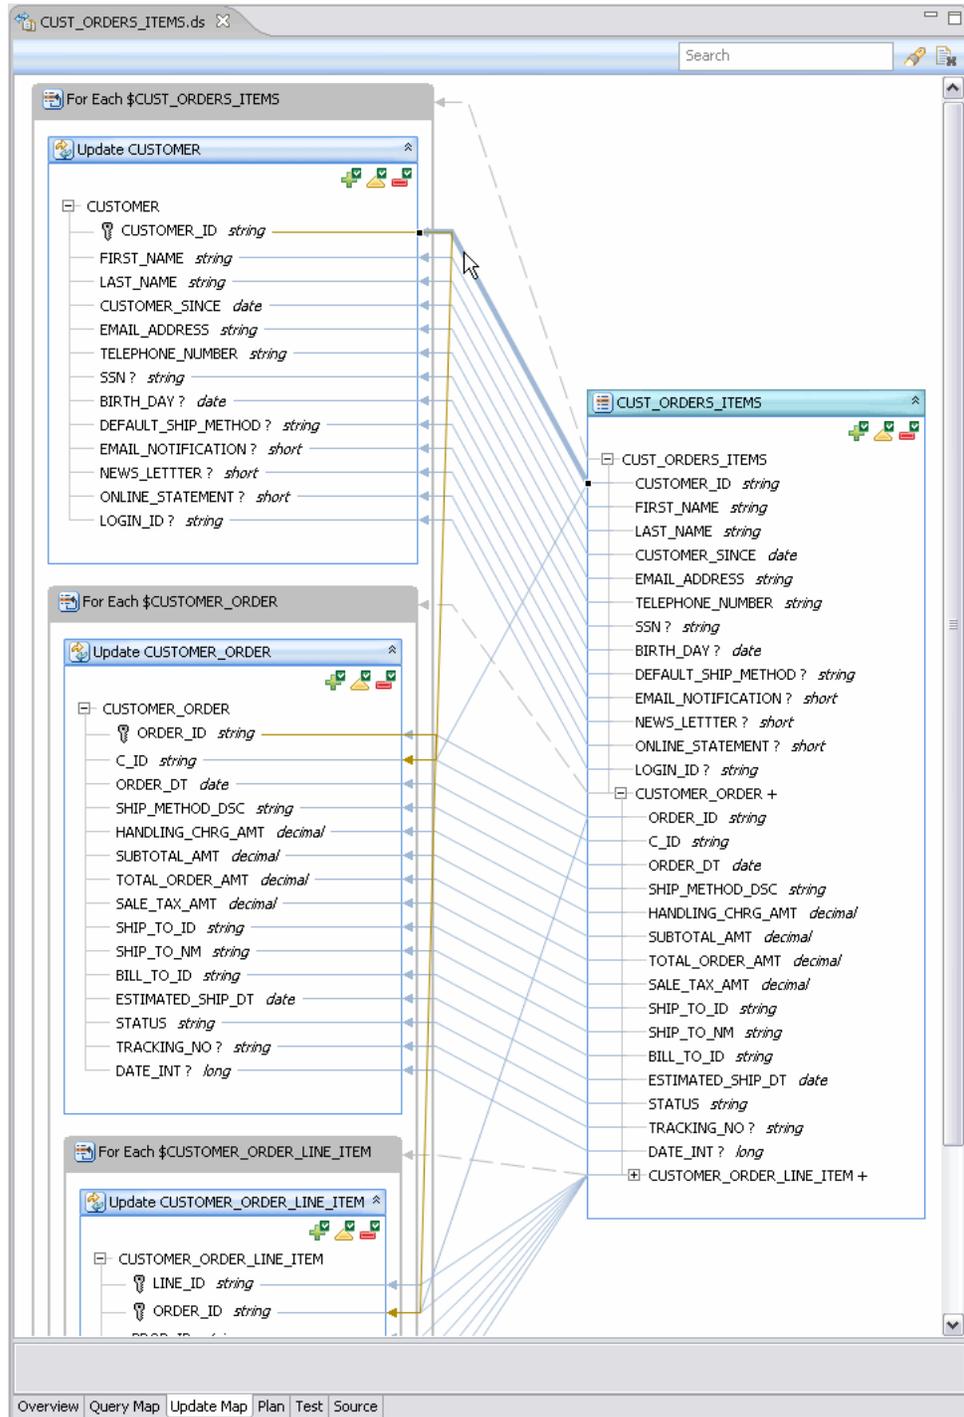
On the Plan tab, the tree version of the query plan for the service's read functions is shown. The `join impl="index-cpp" kind="left-outer"` function is expanded.

1.3.10 Reviewing the Update Map

After an entity data service is successfully published and contains an update function, its update map can be inspected and, as necessary, edited.

- Click the **Update Map** tab.

Figure 1-42 CUST_ORDERS_ITEMS Update Map



The CUST_ORDERS_ITEMS Update Map shows the mapping between CUST_ORDERS_ITEMS, Update CUSTOMER, and Update CUSTOMER_ORDER.

For more information, see "[Understanding Update Maps](#)" on page 9-1.

1.3.11 Archiving Your Project

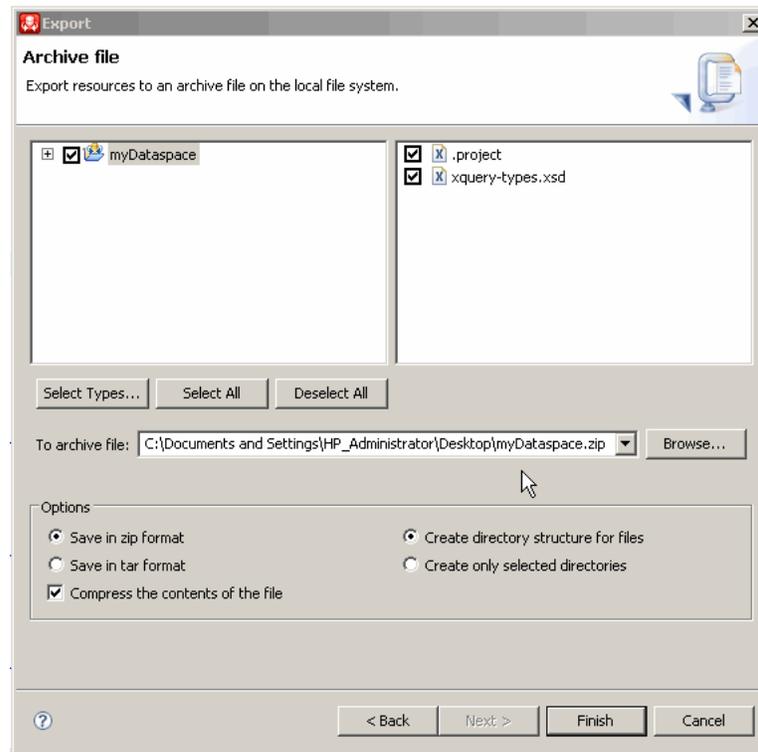
You can save your entire project to a ZIP file. Then, when you need to load it again, you can do so with a simple Import operation.

Other examples in the Oracle Data Service Integrator documentation use this or similar examples, so having this project available will be make it easier to experiment with other Oracle Data Service Integrator faculties.

1. In Project Explorer, right-click on the **myDataspaces Project**.
2. Choose **Export**.
3. In the Export dialog choose:
General > Archive File
4. Click **Next**.

1.3.11.1 Saving Project to a ZIP File

1. In the Archive file dialog the **myDataspaces** project is pre-selected. Browse to the location where you want to put your archive file.
2. Name your file:
myDataspaces
Leave all other options unchanged.
3. Click **Save**.

Figure 1–43 Creating the Archive File

The Archive file dialog exports resources to an archive file on the local file system. The myDataspace project is selected. In the To archive file: field, myDataspace.zip on the local drive has been selected. Options Save in zip format, Compress the contents of the file, and Create directory structure for files are selected. The Back, Finish, and Cancel buttons are active.

4. Click **Finish**.

A file myDataspace.zip will be created in the directory you specified.

1.3.12 Summary

Congratulations! In just a few minutes you have:

- Started Oracle Data Service Integrator.
- Created several physical data services based on existing data.
- Created a logical data service based on elements from three physical sources.
- Build a function to retrieve based on information on a particular customer, the customer's orders, and each item in each order.
- Created an XML Type based on the Return type of your function.
- Modified the XML Type to better support a master-detail arrangement of information.
- Tested your results.
- Edited your results.
- Viewed the query plan and the updated map.

- Create an archive file of your dataspace.
About 150 lines of XQuery have been generated.

1.4 Reference

This section describes the following sections:

- [Section 1.4.1, "Oracle Data Service Integrator Start Menu"](#)
- [Section 1.4.2, "Data Service Types and Functions"](#)
- [Section 1.4.3, "Data Service Characteristics"](#)
- [Section 1.4.4, "Operational Characteristics"](#)

1.4.1 Oracle Data Service Integrator Start Menu

The Oracle Data Service Integrator Start menu provides easy access to components used to develop Oracle Data Service Integrator data services. Access is from the Windows Start menu:

Start > All Programs > Oracle WebLogic

1.4.1.1 Oracle Data Service Integrator Start Menu

The following table describes the menu options available from the main Oracle WebLogic menu.

Table 1–7 Start Menu Options

Option	Usage
Oracle Data Service Integrator	Provides access to the Oracle Data Service Integrator online documentation.
Online Documentation	The Oracle WebLogic documentation home page.
QuickStart	Provides links to help get started with installed Oracle products
SmartUpdate	Used in conjunction with your Oracle Support ID to download any applicable patches and maintenance packs.
Uninstall Oracle Data Service Integrator	Uninstalls Oracle Data Service Integrator.
Eclipse for WebLogic	Oracle Data Service Integrator Eclipse-based IDE
User Projects	User-created domains.

1.4.2 Data Service Types and Functions

Oracle Data Service Integrator functions can have a number of attributes. This section describes those attributes and the conditions under which they are applicable.

Figure 1–44 Oracle Data Service Integrator Data Service Types and Attributes



The Oracle Data Service Integrator Data Service Types and Attributes figure shows information about functions and data services. For access, the types are public,

protected, and private. For primacy, the types are primary and non-primary. For kind, the types are create, update, delete, read, navigate, and library. For operation, the types are function and procedure. For implementation, the implementation types are XQuery, XQSE, template-based, and external. For data services, there are two types: logical and physical. For DS shape, the types are library and entity.

1.4.3 Data Service Characteristics

The following table describes the characteristics of Oracle Data Service Integrator data services. Data service characteristics are defined in the XQuery source pragma.

Table 1–8 Data Service Characteristics

Characteristic	Description
Type	<p>There are two types of data services:</p> <ul style="list-style-type: none"> ■ Physical. The data service is directly based on metadata imported from underlying data sources. Physical data services are created during the metadata import process. ■ Logical. The data service is based wholly or partially on data derived from other data services. Logical data services are created either through the Query Map Editor or in source.
Shape	<p>The shape of a data service is determined by its XML type, or underlying schema, if any. Shapes are:</p> <ul style="list-style-type: none"> ■ Entity. An entity data service is associated with an XML type. For example, physical data services based on relational tables are entity data services. For any given entity data service, all read functions return information in the shape of the primary XML type. ■ Library. A library data service is not associated with an XML type. Library data services contain routines that can be used by other library or entity data services.

For more information, see [Chapter 6, "Building XQueries."](#)

1.4.4 Operational Characteristics

The following table describes the characteristics that can be used to describe functional routines in Oracle Data Service Integrator. These characteristic descriptions are also part of the function's signature, visible in data service Source editor.

Table 1–9 Oracle Data Service Integrator Operations Characteristics

Characteristic	Description
Access	<p>Access or visibility to a functional routine can be set as:</p> <ul style="list-style-type: none"> ■ Public. A public operation can be called from any operation in the same data space and from an Oracle Data Service Integrator client API. Public operations are the only ones that can be called from client APIs such as Web services or the Java Mediator API. ■ Protected. An operation with protected visibility can be called from any operation in the same data space. Protected operations cannot be accessed from Oracle Data Service Integrator client APIs. An operation in the data space can access the function. Functions in physical data services are, by default, protected. ■ Private. The function can only be accessed by other functions in its data service. Operations with private visibility are also off-limits to client APIs.
Primacy	<p>Every logical entity data service identifies a single primary function for each kind of function. For example, if there are several read functions, one will be set as primary.</p> <p>In the case of read functions, the data service relies on the primary read function in the data service to determine the shape of the Return type.</p> <p>For create, update, and delete functions, the primacy setting is used by update templates of component data services.</p> <p>In an entity data service, a function can be set as primary. Other functions of a similar type are automatically considered non-primary.</p> <p>Note that library functions have no Return type and are not categorized as primary or non-primary.</p>
Kind	<p>Oracle Data Service Integrator has several kinds of functions. For physical data services, the kind of function is inferred during the data service creation process, when metadata is imported.</p> <p>Four of the functions are actually CRUD (create-read-update-delete) procedures, which operate on the underlying data.</p> <ul style="list-style-type: none"> ■ Read. Returns data from an underlying data source. ■ Create. Creates one or several records. ■ Update. Updates one or several records. ■ Delete. Deletes one or several records. <p>Other kinds of functions include:</p> <ul style="list-style-type: none"> ■ Navigate. Navigate function have the current data service Return type as one of the input parameters; it typically returns a sequence of the return schema element from the related data Service. Example: Return type Order instead of Return type Customer. ■ Library. Functions, which are independent of the data service XML type. Library functions can appear in either data services and library data services

Table 1–9 (Cont.) Oracle Data Service Integrator Operations Characteristics

Characteristic	Description
Operation	<p>There are two types of operations:</p> <ul style="list-style-type: none"> ■ Functional. General-purpose data service functions are designed to retrieve data for clients. Functions cannot have side-effects. Functions can be defined through XQuery or XQSE. If XQSE is used, the fact that the routine is identified as a function means that it does not have side effects. ■ Procedural. The purpose of a procedural function (also called a procedure or side-effecting procedure) is to affect external processes. A classic example of a side-effecting procedure is an RDBMS stored procedure that modifies underlying data. When a stored procedure is invoked, it operates on the data in the RDBMS without necessarily returning anything to the caller. Similarly, in Oracle Data Service Integrator, a procedural function will primarily invoke an external process. Create-update-delete operations are, by definition, procedural. <p>Note: There is an important distinction between functions and procedures from the perspective of the data service optimizing engine. Procedures are always considered to have side-effects and are therefore never optimized by the XQuery engine in such a way that they do not independently execute. While a delete() function might not be executed (i.e., "optimized away"), a delete() procedure will always be called.</p>
Implementation	<p>Functions can be implemented in the following ways:</p> <ul style="list-style-type: none"> ■ XQuery. The most common means of implementing an Oracle Data Service Integrator function is through XQuery. Of course the data service itself is implemented in XQuery. ■ XQSE. The XQuery Scripting Extension provides a procedural language to extend XQuery to support certain kinds of operations. ■ Template-based. An update template defines the data flow and order for update operations for a logical data service. The update engine in the Oracle Data Service Integrator server executes a procedure based on a template; is typically a Java routine used to manage updates of non-relational data. The same template is used by create, update, and delete routines. ■ External. External functions are based on physical sources such as Java, web services, XML, flat files, or relational sources. External functions can be created in entity or library data services.

1.5 Related Topics

This section describes the following topics:

- [Section 1.5.1, "Getting the Most from the WebLogic Eclipse Plugin Framework"](#)
- [Section 1.5.2, "Create a Data Service with a Flat Return Type"](#)

1.5.1 Getting the Most from the WebLogic Eclipse Plugin Framework

Oracle Data Service Integrator dataspace are initially created as projects in the WebLogic Eclipse plugin framework. The Eclipse IDE is a rich, open development environment.

While some aspects of Eclipse are described in this section, no attempt is made to replicate the large body of documentation available for Eclipse developers.

References

- Eclipse home site
- Eclipse Help documentation
- Eclipse user guides

1.5.1.1 Data Services Eclipse for WebLogic

Eclipse for WebLogic run inside the Eclipse framework.

Table 1–10 Workshop for WebLogic Artifacts in the Oracle Data Service Integrator Perspective

Artifact	Purpose
Project Explorer	Contains project artifacts including data services and their functions.
Properties editor	Contains read/write and read only properties associated with the selected artifact. For example a function may be set to public, protected, or private through its Properties editor.
Outline manager	Provides a scrollable view of your model, query, or update mapper. This is particularly useful in large projects since the work area may not be large enough to show all the artifacts.
Console	The console appears whenever the server is accessed.
Servers tab	Display the status of the Oracle Data Service Integrator server which in turn provides clients with access to data services and their underlying data sources.
Problems tab	Displays problems encountered by the project.
Error Log tab	Displays errors associated with the project.

The **Windows > Show View** menu option can be used to add additional windows to the perspective.

In addition, several Oracle Data Service Integrator Perspective menu options are provided under:

File > New

These allow you to, first, create an Oracle Data Service Integrator dataspace project and then to add various types of data services, models and web service mapper.

1.5.2 Create a Data Service with a Flat Return Type

This topic shows you how to create an update map from a logical data service with a flat, non-nested return type, using the sample database that ships with Oracle Data Service Integrator.

- [Section 1.5.2.1, "Overview"](#)
- [Section 1.5.2.3, "Create the Return Type"](#)
- [Section 1.5.2.3, "Create the Return Type"](#)
- [Section 1.5.2.5, "Create a Logical Data Service"](#)
- [Section 1.5.2.5, "Create a Logical Data Service"](#)

- [Section 1.5.2.6, "Create the Query Map"](#)
- [Section 1.5.2.7, "See Also"](#)

1.5.2.1 Overview

A return type can be non-nested, or flat, even if it joins two relational tables, where one table has a one-to-many relationship with the other table. An example is one customer in a CUSTOMER table with many Orders in an ORDERS table. One approach to the return type is to nest an Orders element of multiple cardinality beneath the Customer element.

Figure 1–45 A Nested Customer-and-Orders Schema



A nested schema is shown. CustomerOrder is expanded to CUSTOMER and CUSTOMER_ORDER. CUSTOMER is expanded to CUSTOMER_ID, FIRST_NAME, LAST_NAME, and SSN. CUSTOMER_ORDER is expanded to ORDER_ID, C_ID, TOTAL_ORDER_AMT, and STATUS.

Because you can design a logical data service with any structure, regardless of the underlying data sources, it is just as valid to define a flat return type to model the relationship between Customers and Orders.

Figure 1–46 A Flat Customer-and-Orders Schema



A flat schema is shown. Under CUSTOMERS_AND_ORDERS, the following appear: CUSTOMER_ID, FIRST_NAME, LAST_NAME, EMAIL_ADDRESS, ORDER_ID, ORDER_DT, TOTAL_ORDER_AMT.

1.5.2.2 Create a Dataspace Project

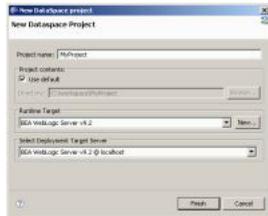
First, create a new dataspace project to contain your physical and logical data services:

1. In Eclipse for WebLogic, choose **File > New > Dataspace Project**.
2. Enter a project name such as FlatReturnType, then click **Finish**.
3. Right-click the new dataspace project name, and choose **New > Folder**.

4. Create folders named physical and logical. Within logical, create a folder named schemas.

Using separate folders for physical and logical services helps separate the physical and logical integration layers.

Figure 1–47 Adding a New Dataspace Project



The New Dataspace Project dialog contains fields for Project Name, Project Contents, Runtime Target, and Select Deployment Target Server.

1.5.2.3 Create the Return Type

The return type the logical data service uses combines data from the CUSTOMER table and the ORDERS table. It has a non-nested XML structure, even though the data shows that customers and orders have a one-to-many relationship.

You can define the return type by creating an XML schema (XSD) file. In an XML editor, create a schema file like this one:

Example 1–5 XML schema (XSD) File

```
<?xml version="1.0" encoding="UTF-8" ?>
<xs:schema targetNamespace="ld:logical/FlatReturnType" xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CUSTOMERS_AND_ORDERS">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="CUSTOMER_ID" type="xs:string"/>
        <xs:element name="FIRST_NAME" type="xs:string"/>
        <xs:element name="LAST_NAME" type="xs:string"/>
        <xs:element name="EMAIL_ADDRESS" type="xs:string"/>
        <xs:element name="ORDER_ID" type="xs:string"/>
        <xs:element name="ORDER_DT" type="xs:date"/>
        <xs:element name="TOTAL_ORDER_AMT" type="xs:decimal"/>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

Be sure to:

- Define targetNamespace to make sense for your dataspace project.
Make sure you have only one top-level element of the name you choose (here, CUSTOMERORDER) in your target namespace. You can give the targetNamespace the same name as the dataspace project, but you are not required to.
- Save the schema file in the logical/schemas folder within your dataspace project.

Note that the cardinality of all elements uses the default values, minOccurs="1" and maxOccurs="1". Each customer has many orders, but there is only one combination of

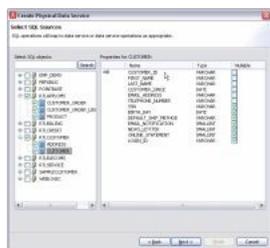
customer and order, so the cardinality of the order elements (ORDER_ID, ORDER_DT, and TOTAL_ORDER_AMT) is still 1.

1.5.2.4 Create Physical Data Services

Now, create physical data services based on the sample database or your own physical data sources.

1. In Project Explorer, right-click the physical folder in your dataspace project.
2. Choose **New > Physical Data Service**.
3. Choose **Relational** for Data source type and *dspSamplesDataSource* for Data source, then click **Next**.
4. Expand *RTLCUSTOMER* and select *CUSTOMER*.
5. Select **Public** for both *CUSTOMER* and *CUSTOMER_ORDER*, then click **Next**.
6. Click **Finish**.
7. When asked if you want to open the new data services, click **No**.

Figure 1–48 Adding Physical Data Services



The Create Physical Data Service dialog lets you choose the data source, database type, and relational object for the new data service. There are field for Container, Data source type, Data source, and Database type. Under Relational database object, choose Tables and views, Stored procedures, SQL statement, or Database function.

This graphic does not match the previous text.

1.5.2.5 Create a Logical Data Service

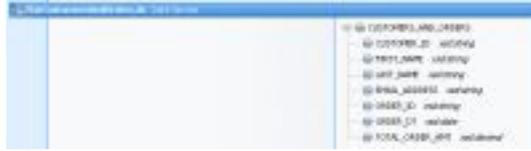
Now that you have physical data services and a schema for the return type, you can create the logical data service.

1. Right-click the logical folder, then choose **New > Logical Data Service**.
2. Enter a name for the service, such as *FlatCustomersAndOrders*.
3. Make sure **Entity Data Service** is selected, then click **Finish**.

Now associate a return type with the service:

1. Right-click in the **Overview** tab and choose **Associate XML Type**.
2. Select the schema and click **OK**.

Figure 1–49 A New Logical Data Service with a Return Type

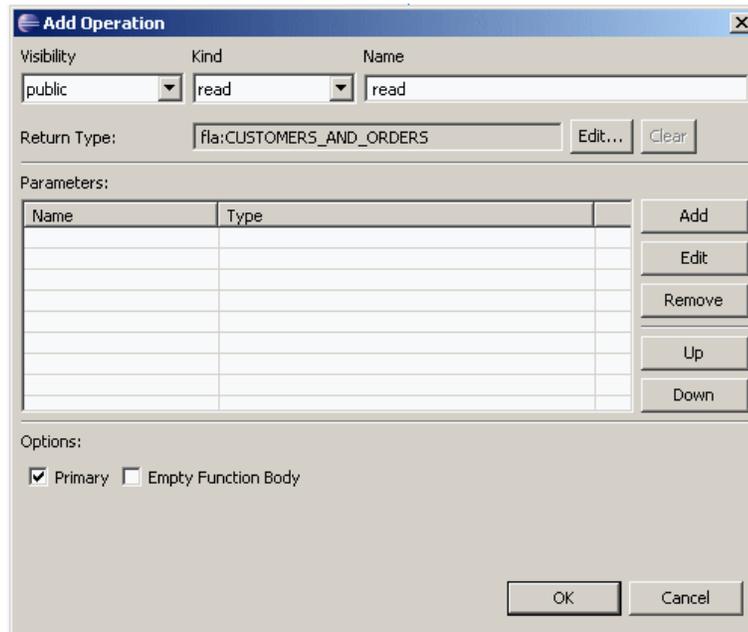


A flat Customers and Orders data services shows a CUSTOMERS_AND_ORDERS return type.

You also need to define a primary Read function, in order to create both the query map and update map.

1. Right-click in the service name bar at the top, and choose **Add Operation**.
2. Make sure **Kind** is set to **read**, then enter a function name, such as `read`.
3. Make sure **Primary** is selected, then click **OK**.

Figure 1–50 Creating a Primary Read Function



On the Add Operation dialog, Visibility is set to public, Kind is set to read, and name is set to read. The Return Type is fla:CUSTOMERS_AND_ORDERS. The Primary option is selected.

1.5.2.6 Create the Query Map

Now you need to create the query map visually in Eclipse for WebLogic, which in turn generates an update map.

1. Click the **Query Map** tab.
2. In Project Explorer, expand the physical data services **CUSTOMER.ds** and **CUSTOMER_ORDER.ds**.
3. Drag the **Read** function from each physical service to the mapping area.
 Notice that you cannot scope the **CUSTOMER_ORDER** block to a subtype in the return type, because the return type has no subtypes.
4. Drag mappings from the **CUSTOMER** block on the left to the return type for **CUSTOMER_ID**, **FIRST_NAME**, **LAST_NAME**, and **EMAIL_ADDRESS**.
5. Drag mappings from the **CUSTOMER_ORDER** block on the left to the return type for **ORDER_ID**, **ORDER_DT**, and **TOTAL_ORDER_AMT**.
6. In the **For** blocks, drag from **CUSTOMER/CUSTOMER_ID** to **CUSTOMER_ORDER/CUSTOMER_ID**.

This creates a join between the two data sources.

At this point, the query map looks like this.

Figure 1–51 A Query Map with Mappings and a Join



Mappings to the return type are shown, as well as the join (the dotted line) between **CUSTOMER** and **CUSTOMER_ORDER**.

If you click the **Source** tab and expand the **Read** function, you see XQuery code like this:

Example 1–6 XQuery Code for read function

```

declare function tns:read() as element(flac:CUSTOMERS_AND_ORDERS) *{
  for $CUSTOMER_ORDER in cus1:CUSTOMER_ORDER()
  for $CUSTOMER in cus:CUSTOMER()
  where $CUSTOMER/CUSTOMER_ID eq $CUSTOMER_ORDER/C_ID
  return
    <flac:CUSTOMERS_AND_ORDERS>
      <CUSTOMER_ID>{fn:data($CUSTOMER/CUSTOMER_ID)}</CUSTOMER_ID>
      <FIRST_NAME>{fn:data($CUSTOMER/FIRST_NAME)}</FIRST_NAME>
      <LAST_NAME>{fn:data($CUSTOMER/LAST_NAME)}</LAST_NAME>
      <EMAIL_ADDRESS>{fn:data($CUSTOMER/EMAIL_ADDRESS)}</EMAIL_ADDRESS>
      <ORDER_ID>{fn:data($CUSTOMER_ORDER/ORDER_ID)}</ORDER_ID>
      <ORDER_DT>{fn:data($CUSTOMER_ORDER/ORDER_DT)}</ORDER_DT>
      <TOTAL_ORDER_AMT>{fn:data($CUSTOMER_ORDER/TOTAL_ORDER_AMT)}</TOTAL_ORDER_AMT>
    </flac:CUSTOMERS_AND_ORDERS>
};
  
```

Notice that the XQuery code has a for statement nested directly within another for statement. This creates an inner join between the two tables in SQL. To confirm the SQL that is created:

1. Click the **Test** tab.
2. At Select operation, make sure the primary Read function is selected.
3. Click **Run** (saving your data service as necessary).

You should see an XQuery FLWOR statement node. If you expand it, you should see a SQL query like this, showing an inner join:

```
SELECT t1."ORDER_DT" AS c1, t1."ORDER_ID" AS c2, t1."TOTAL_ORDER_AMT" AS c3,  
       t2."CUSTOMER_ID" AS c4, t2."EMAIL_ADDRESS" AS c5, t2."FIRST_NAME" AS c6, t2."LAST_NAME" AS c7  
FROM "RTLAPPLOMS"."CUSTOMER_ORDER" t1  
JOIN "RTLCUSTOMER"."CUSTOMER" t2  
ON (t2."CUSTOMER_ID" = t1."C_ID"
```

The inner join is created because the logical data service has a flat return type. When you mouse over the SQL query, you see this message:

Generated SQL query does not have a WHERE clause. This may cause the query to take longer to finish and use excessive memory resources.

1.5.2.7 See Also

[Example: How to Create Your First Data Services](#)

Developing and Managing Dataspace Projects

This chapter describes the concepts of data service validation during deployment, and how to create, build, clean, configure, remove, export, import, and build dataspace projects.

This chapter contains the following topics:

- [Section 2.1, "Data Service File Validation During Deployment"](#)
- [Section 2.2, "How-to"](#)
- [Section 2.3, "Reference"](#)
- [Section 2.4, "Related Topics"](#)

2.1 Data Service File Validation During Deployment

In the Eclipse IDE a dataspace project's data service (.ds) files are validated automatically according to the following deployment model:

- The publish to server action validates the dataspace projects.
- All the project's artifacts are collected.
- The collected artifacts are deployed to the server.

2.2 How-to

This section describes the following topics:

- [Section 2.2.1, "How To Create, Build, Clean, and Delete Dataspace Projects"](#)
- [Section 2.2.2, "How to Publish, Configure, and Remove Dataspace Projects"](#)
- [Section 2.2.3, "Exporting Dataspace Projects or Project Folders"](#)
- [Section 2.2.4, "Exporting Dataspace Project Artifacts Using Oracle Data Service Integrator Export Wizards"](#)
- [Section 2.2.5, "Importing a Dataspace Project"](#)
- [Section 2.2.6, "How To Handle Error Conditions in a Dataspace Project"](#)
- [Section 2.2.7, "How To Validate, Build, Export, and Package Dataspace Projects from the Command Line"](#)

2.2.1 How To Create, Build, Clean, and Delete Dataspace Projects

This section describes the following topics:

- [Section 2.2.1.1, "Creating a Dataspace Project"](#)
- [Section 2.2.1.2, "Building a Dataspace Project"](#)
- [Section 2.2.1.3, "Cleaning a Dataspace Project"](#)
- [Section 2.2.1.4, "Deleting a Dataspace Project"](#)

A dataspace project is developed in Eclipse for WebLogic and deployed to a local server. While the development process typically is an iterative cycle of modification and deployment, it is important to keep in mind that the existence of a project in Eclipse for WebLogic is only loosely coupled with its deployed status. This loose coupling has implications for several types of operations:

- Development and deployment to the server
- Publishing to the server
- Configuring projects on the server
- Removal of the project
- Removal of the project from the server

2.2.1.1 Creating a Dataspace Project

You can create a new dataspace projects using the Eclipse for WebLogic File menu.

File > New > Dataspace Project

For more information, see ["Creating a Dataspace Project"](#) on page 1-15.

2.2.1.2 Building a Dataspace Project

In Eclipse for WebLogic it is often a good practice to set your project to be built automatically every time you modify a file in your project. You can establish this setting through the Eclipse for WebLogic Project menu:

Project > Build Automatically

A checkbox appears when this option is selected.

2.2.1.3 Cleaning a Dataspace Project

Applying a "clean" to a project clears out any existing build problems and build states. If your build runs into error conditions or other problems, cleaning and redeploying your project is a recommended first step.

Project > Clean...

WARNING: A dataspace project can only be deployed when no other process has an editing lock on the Oracle Data Service Integrator configuration that contains your dataspace. The Oracle Data Service Integrator configuration can be locked through the Oracle Data Service Integrator Administration Console (Lock and Edit), by a client process (MBean API or WLST script), or during deployment from Eclipse/WebLogic for Eclipse.

2.2.1.4 Deleting a Dataspace Project

Dataspace projects are both created and deleted through the Project Explorer.

To delete a project:

1. Right-click on the project's name in the Project Explorer.
2. Select **Delete**.

You will be given two options:

- **Delete content from the file system?** If you choose:
 - Do not delete content
you will be able to import the project at a later time.
- **Delete the dataspace on the server?** The deployed dataspace will be removed from the server. If this option is not selected, the dataspace will remain in one of two states, depending on selected options:
 - Available to be configured on the server
 - Configured on the server

Note: Data services can also be removed from their server through the Oracle Data Service Integrator Administration Console.

2.2.2 How to Publish, Configure, and Remove Dataspace Projects

Dataspace projects are created in the Eclipse for WebLogic Eclipse plugin framework. A project that builds successfully is ready to be made available from a local supported version of Oracle WebLogic server.

Several terms can be used to describe the process of managing a server's dataspace projects.

- **Publish.** All projects associated with a server can be deployed at once.
- **Available.** These are projects that have been published and are available to be configured on the server. A project must both present on the server and configured before it can be access by client applications. A project with an Available status can be thought of as staged.
- **Configured.** A configured project is available to authorized calling applications. A project with Configured status on the server can be thought of as released.

Tip: Use the Add and Remove Projects dialog to move projects between **Available** and **Configured** state.

This section describes the following topics:

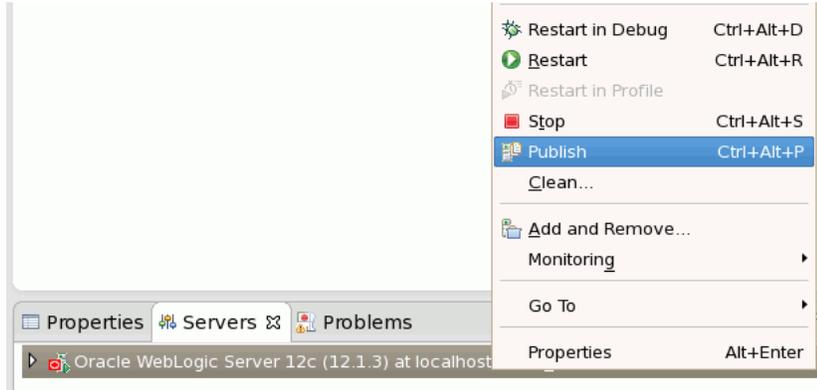
- [Section 2.2.2.1, "Publishing Server Projects"](#)
- [Section 2.2.2.2, "Configuring Server Projects"](#)
- [Section 2.2.2.3, "Managing Configured Projects Through the Servers Window"](#)
- [Section 2.2.2.4, "Removing Dataspace Projects from a Server"](#)

2.2.2.1 Publishing Server Projects

You can publish all the dataspace projects associated with a workspace.

Right-click on the name of your server in the Servers window and choose Publish.

Figure 2–1 Publishing Server Projects



Publishing Server Projects

The state in the Servers window will be changed to **Republish**.

Publishing or republishing a set of projects does not affect the configuration status of each project on the server. You can modify the configuration status through the **Add and Remove Projects...** dialog.

2.2.2.2 Configuring Server Projects

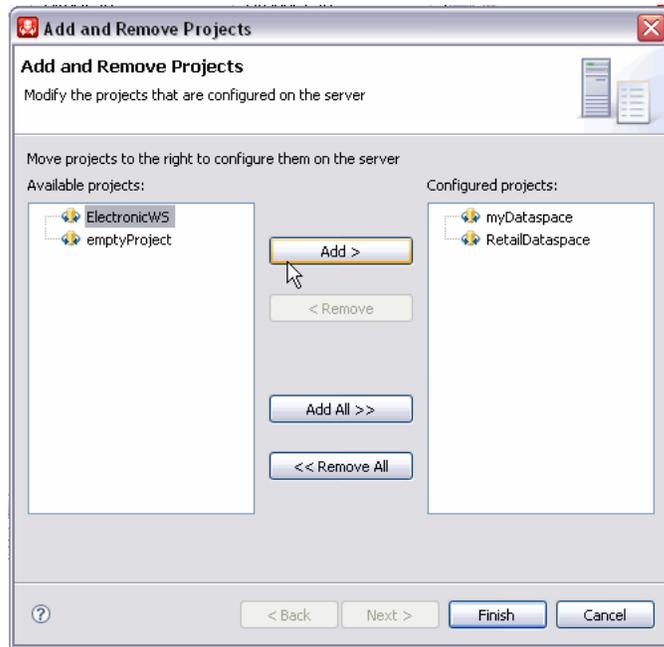
Projects on a server are considered either configured on the server or available to be configured on the server. Configuration status is managed either through the Add and Remove Projects... dialog or directly from the Servers window.

Note: Only configured projects are available to client applications.

2.2.2.2.1 Managing Configured Projects Through Dialog To access the dialog:

1. Right-click on the name of the server in the Servers window.
2. Select **Add and Remove Projects...**

Figure 2–2 Add and Remove Projects Dialog

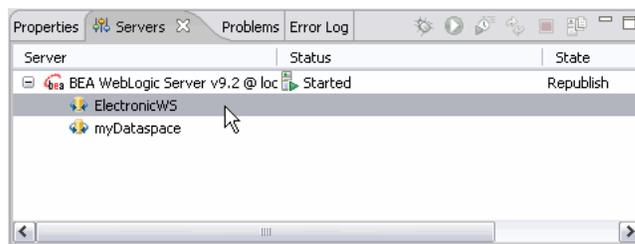


The Add and Remove Projects dialog modifies projects that are configured on the servers. There are two lists: **Available projects** and **Configured projects**. Select a project on the Available projects list and click the **Add** button to move it to the Configured projects list. To remove a project, click on the project in the Configured projects list, click **Remove**, and the project moves to the Available projects list. Similarly, there are **Add all** and **Remove all** buttons.

2.2.2.3 Managing Configured Projects Through the Servers Window

You can also change a project's configuration status through the Servers window.

Figure 2–3 Server and Projects Dialog



Server and Projects Dialog

1. Click on the + symbol next to the server name.
2. Right-click on the project you wish to unconfigure.
3. Select **Remove**.

Alternatively, just select your project and click the **Delete** key.

You can use the Add and Remove Projects Dialog to change the configuration status of your project.

2.2.2.4 Removing Dataspace Projects from a Server

You can permanently remove a project from the server through the right-click menu **Delete** option in the Project Explorer.

For more information, see ["Deleting a Dataspace Project"](#) on page 2-3.

2.2.3 Exporting Dataspace Projects or Project Folders

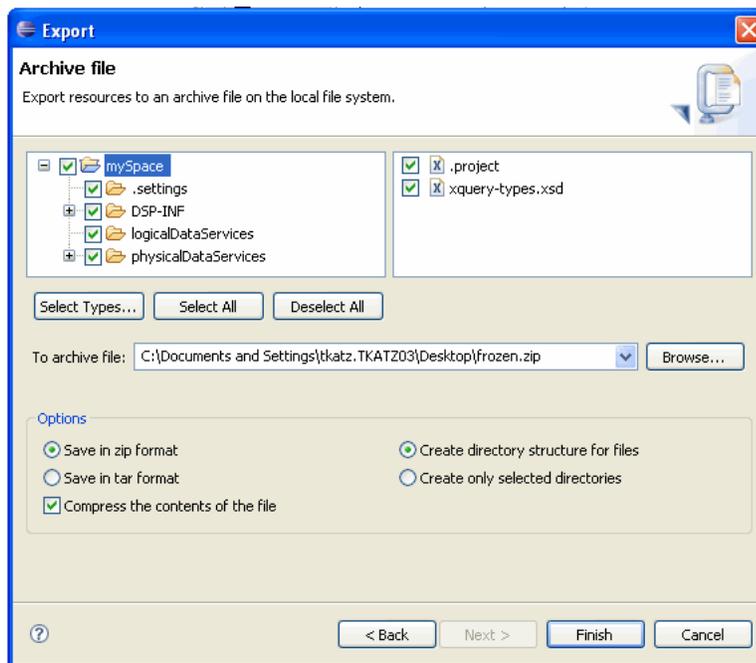
Oracle Data Service Integrator dataspace projects or their component folders can be exported in EAR (archive) format using standard Eclipse mechanisms. The export target is the local file system.

If an entire project is exported as an EAR, it constitutes a back-up of the project which can then be re-imported into an Eclipse-compatible IDE.

To create an archive file of a dataspace project:

1. In Project Explorer right-click on your project (or folder).
2. Navigate to:
Export > General > Archive File
3. In the Archive File wizard select the entire project or one or several folders.
4. Select from available export options.
5. Click **Finish**.

Figure 2–4 *Creating an EAR File for a Dataspace Project*



The Export dialog lets you export resources to an archive file on the local file system. There are two panels: on the left, the mySpace folder is open and selected, and the subfolders are selected. On the right, the project and xquery-types.xsd file are selected.

There are three buttons: Select Types, Select All, and Deselect All. Browse to specify the To archive file location. The Save in zip format and Create directory structure for files options are selected, as well as the Compress the contents of the file checkbox. The Save in tar format and Create only selected directories options are unselected.

See Eclipse Documentation (<http://www.eclipse.org/documentation/>) for general information about the Export operation.

2.2.4 Exporting Dataspace Project Artifacts Using Oracle Data Service Integrator Export Wizards

This section describes the various types of export operations available for Oracle Data Service Integrator dataspace projects.

- [Section 2.2.4.1, "Exporting Dataspace Artifacts"](#)
- [Section 2.2.4.2, "Generating a Data Service Definitions and Artifacts JAR"](#)
- [Section 2.2.4.3, "Generating a Mediator Client JAR File"](#)
- [Section 2.2.4.4, "Generating a JAR File Containing Data Service-to-Web Service Maps"](#)

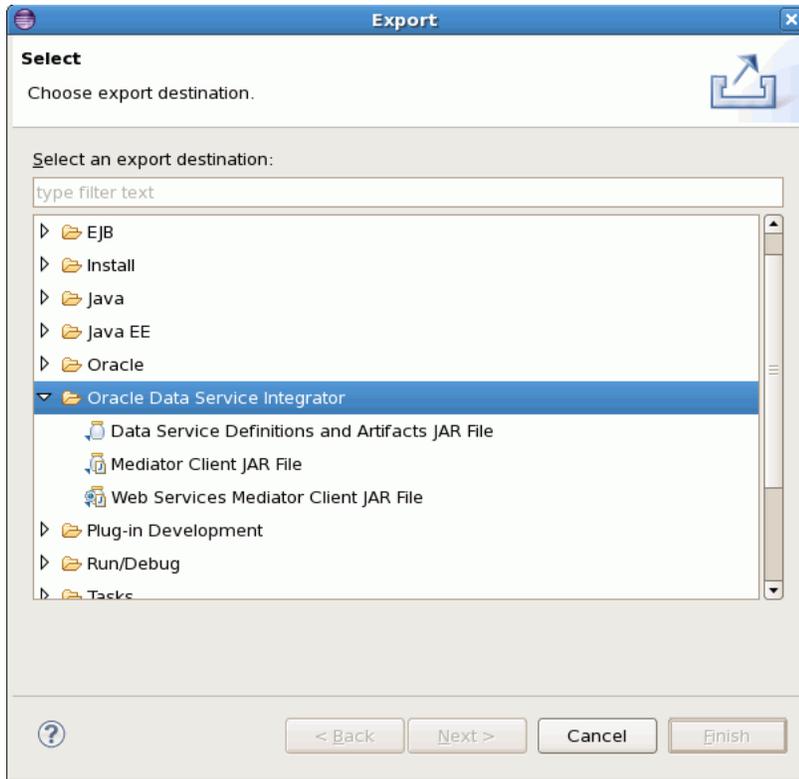
2.2.4.1 Exporting Dataspace Artifacts

The Oracle Data Service Integrator export wizards can be accessed using the **File > Export** menu, then expand on the Oracle Data Service Integrator export wizard category.

File > Export... > Oracle Data Service Integrator

Note: Artifacts can only be exported from deployable dataspace projects; if your project is not deployable, the export operation will not succeed.

Figure 2–5 Exporting a Dataspace Project



The Export dialog lets you choose an export destination. There is a field to typing filter text. A directory tree is shown and the AquaLogic Data Services Platform folder is selected. There are Back, Next, Finish, and Cancel buttons.

In the Oracle Data Service Integrator category there are three wizards. Each generates a specific type of JAR file.

Table 2–1 Types of Oracle Data Service Integrator JAR File Export Operations

Export Type	Effect
Data service definitions and artifacts	All Oracle Data Service Integrator deployable artifacts are bundled into a JAR file.
Mediator client	A Java interface for accessing data services is created.
Web services mediator client	A web services interface for accessing data services is created.

Any dataspace projects pre-selected in the Project Explorer will automatically be selected in the export wizard.

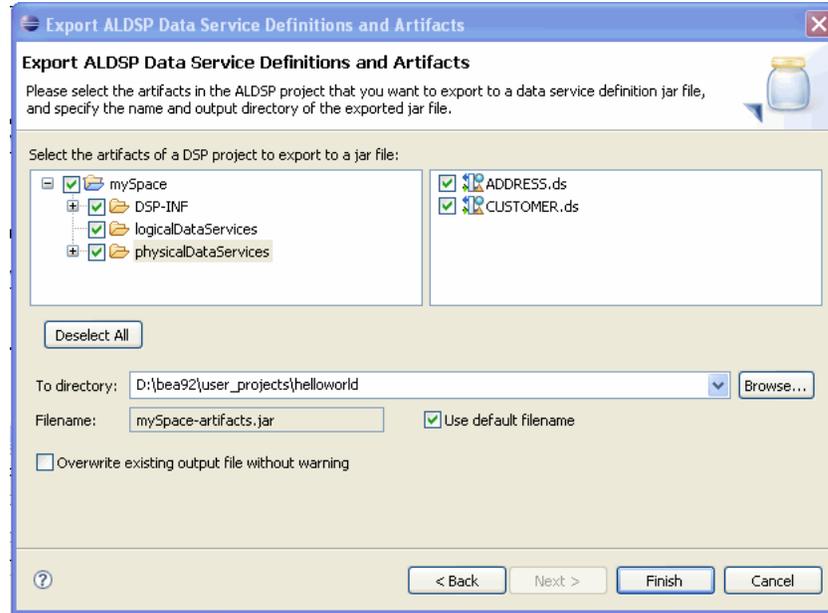
2.2.4.2 Generating a Data Service Definitions and Artifacts JAR

An exported JAR file containing a single project's server-deployable definitions and artifacts. Such a file can be imported into another Oracle Data Service Integrator-enabled version of Eclipse. In addition, the definitions and artifacts JAR can be useful:

- As a means of transporting a dataspace from one application to another.
- In conjunction with certain refactoring operations.

- For deployment on multiple servers (clusters) at a later time.
- For debugging purposes.

Figure 2–6 Export Data Service Definitions and Artifacts



The Export OSDI Studio Data Service Definitions and Artifacts dialog shows two panels containing files in a flat structure. All folders are selected. There is a Deselect All button. Browse to specify the directory. You can specify the filename or select Use default filename. The Overwrite existing output file without warning option is deselected. There are Back, Next, Finish, and Cancel buttons.

In the wizard, the contents you identify using their adjacent checkbox will be exported. For example if you check the box next to the project name, all of that projects server-deployable components will be selected.

You can fine-tune your selection by clicking on a folder. The folder's contents will appear in the right-hand column where you can use a checkbox to control which artifact will be exported.

Table 2–2 Actions Associated with Generating a Data Service Definitions and Artifacts JAR

	Item	Recommended Setting or Action	Details – Comments
1	Export Data Service Definitions and Artifacts Page		
2	Check folders or their contents that you want to export		
3	Deselect All		Convenience if more than one project is selected.
4	To Directory:	/user_projects/helloworld	JAR file can be exported anywhere on the system.
5	Filename:	-artifacts.jar	Example: mySpace-dsp-client.jar

Table 2–2 (Cont.) Actions Associated with Generating a Data Service Definitions and Artifacts JAR

	Item	Recommended Setting or Action	Details – Comments
6	Use default filename option	Selected	When selected editing of the generated filename is not allowed.
7	Overwrite existing file without warning option	Unselected	If unselected you will be asked if you want to overwrite any existing file of the same name.
8	Success message	Finish	Identifies project and target name.

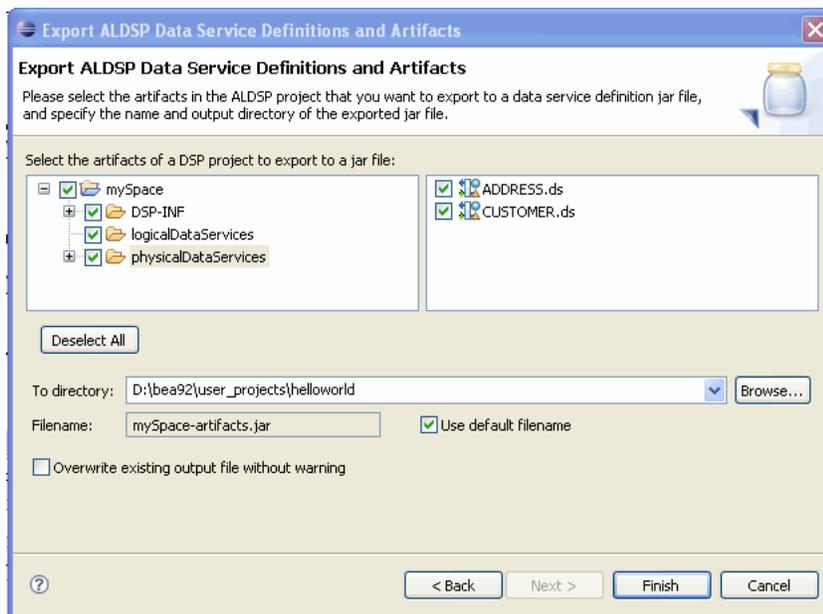
2.2.4.3 Generating a Mediator Client JAR File

This wizard presents a list of open Oracle Data Service Integrator projects to select from, and it generates an Oracle Data Service Integrator mediator client JAR file from the currently selected project. Projects are exported one at a time.

Java programs access data services through the Oracle Data Service Integrator Mediator API. This API is generated from the Oracle Data Service Integrator Eclipse platform and is based on a project that can be successfully built and deployed.

For more information, see *Accessing Data Services from Java Clients*.

Figure 2–7 Export Mediator Client JAR Wizard



The Export OSDI Studio Mediator Client Jar dialog lets you select the OSDI Studio project from which you want to generate a mediator client jar file, and specify the name and output directory for the jar file. There is a pane for selecting the DSP project for generating a mediator client jar file. The mySpace folder is selected. There is a Deselect All button. You can browse to the directory you want. There is a space to specify the filename. The Use default filename option is checked. The Overwrite existing output file without warning option is unchecked. There are Back, Next, Finish, and Cancel buttons.

Table 2–3 Steps Associated with Generating a Mediator Client API JAR File

	Item	Recommended Setting or Action	Details – Comments
1	Select the Mediator Client JAR File export wizard		
2	Type:	Mediator Client JAR File	
3	Select a Dataspace Project Page	Next	
4	Pick a dataspace project	mySpace	Only one project at a time can be exported.
5	Deselect All		Convenience if more than one project is selected.
6	To Directory:	/user_projects/helloworld	JAR file can be exported anywhere on the system.
7	Filename:	-dsp-client.jar	Example: mySpace-dsp-client.jar
8	Use default filename option	Selected	When selected editing of the generated filename is not allowed.
9	Overwrite existing file without warning option	Unselected	If unselected you will be asked if you want to overwrite any existing file of the same name.
10	Success message	Finish	Identifies project and target name.

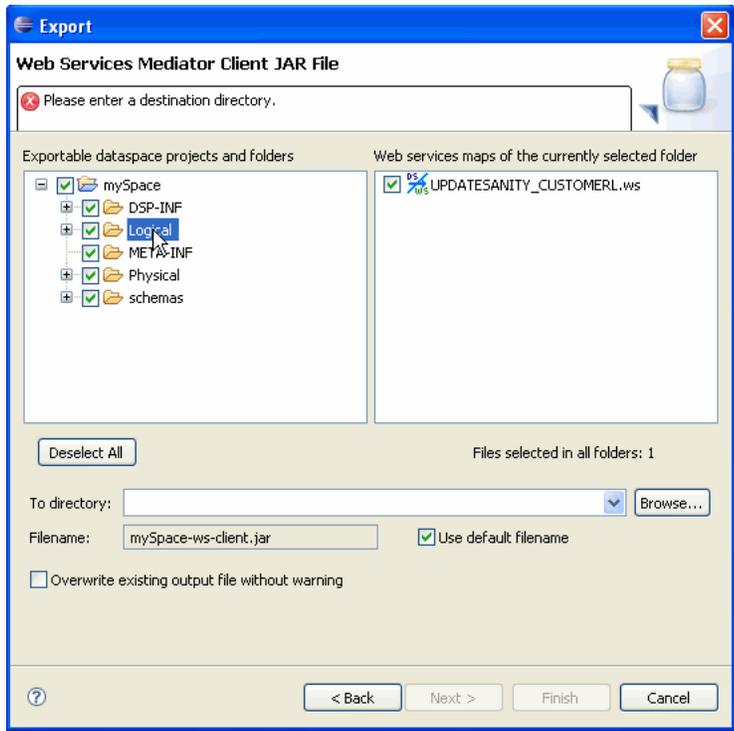
2.2.4.4 Generating a JAR File Containing Data Service-to-Web Service Maps

After you have created a web service map of one or more data services you can create an exported JAR file containing these maps.

Then the created JAR file can be used as your applications web service interface to available data services.

Note: Only publicly available operations can be turned into web service operations. You can adjust access level to a data service function through the Properties window.

Figure 2–8 Exporting a Web Services Mediator Client JAR File



In the Web Services Mediator Client JAR File dialog, there are two panels: one for exportable dataspace projects and folders, and another for Web services maps of the currently selected folder. All items on both panels are selected. There is a Deselect All button. Browse to specify a destination directory. Specify a filename. The Use default filename option is selected. The overwrite existing output file without warning option is deselected.

Table 2–4 Actions Associated with Generating a Web Services Mediator Client JAR File

	Item	Recommended Setting or Action	Details – Comments
1	Export Web Services Mediator Client JAR File		
2	Check folders or their contents that contain web service maps	Select folder	Click on the exportable dataspace folder to located selected WS file.
3	Deselect All		Convenience if more than one project is selected.
4	To Directory:	/user_projects/helloworld	JAR file can be exported anywhere on the system.
5	Filename:	-ws-client.jar	Example: mySpace-ws-client.jar

Table 2–4 (Cont.) Actions Associated with Generating a Web Services Mediator Client JAR File

	Item	Recommended Setting or Action	Details – Comments
6	Use default filename option	Selected	When selected editing of the generated filename is not allowed.
7	Overwrite existing file without warning option	Unselected	If unselected you will be asked if you want to overwrite any existing file of the same name.
8	Success message	Finish	Identifies project and target name.

2.2.5 Importing a Dataspace Project

A JAR file containing an Oracle Data Service Integrator project can be imported into your workspace.

1. Create a dataspace project, naming it appropriately. (Alternatively you may be able to use an existing project if there are no naming conflicts.)

File > New > Dataspace Project

2. Click on your new project.
3. Choose:
File > Import > General > Archive File
4. **Next.**
5. Browse to the directory location of your JAR file.
6. **Open.**
7. Answer **Yes to All** to the question regarding overwriting `xquery-types.xsd`.
8. **Finish.**
9. Deploy your project to verify a successful build and deployment.

2.2.5.1 Importing a File-based Project

A project in an accessible file system can be imported into Eclipse for WebLogic. You can import one or several projects at the same time.

1. Choose:
File > Import... > General > Existing Projects into Workspace
2. Browse to your dataspace project directory.
3. Select the project or projects you wish to import.

2.2.6 How To Handle Error Conditions in a Dataspace Project

During the course of creating your project there are times when the project will be in an error condition. There are many reasons for this. Generally speaking the way to handle such conditions is to either:

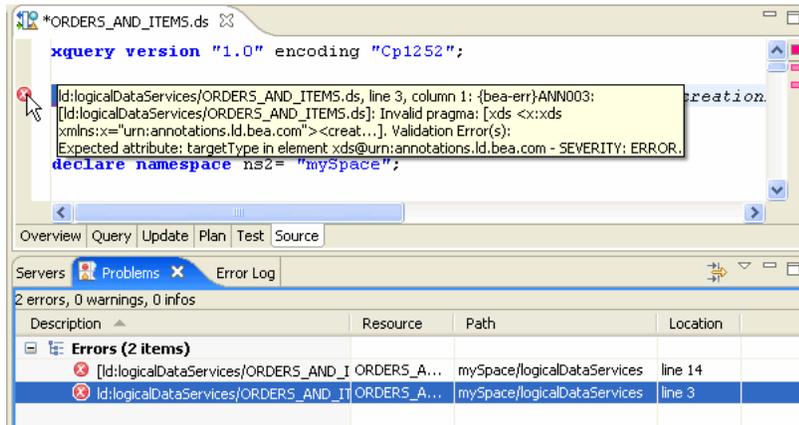
- Go forward because you understand why the condition has occurred.
- Revert using **Undo**.

There may be some cases, however, when the error condition comes as a surprise and/or there is no easy way to revert. Information about such conditions can be found in two places.

Table 2–5 Problem Reporting in Dataspace Projects

Tabular window	Purpose
Problems (Window > Show View > Problems)	Collects and displays errors in the data service source file.
Error Log (Window > Show View > Error Log)	Collects and displays project-related error conditions.

Figure 2–9 Problem and Error Log Tabs in a Dataspace



The Problems tab displays a description, resource, path, and location for each error. In this case, two errors are listed under the Problems tab. A detailed error report appears in the source code in the upper half of the dialog.

The error log contains several types of messages; icons are used to differentiate their type.

Table 2–6 Error Log Icons and Their Meaning

Icon	Meaning
	Error.
	Error with log or stack trace.
	Warning.
	Informational.
	Process icon.

Double-clicking on each line in the Error Log window will open a separate dialog that will allow you to see more information. Examples:

- Double-clicking on an error might open a dialog that will contain the related stack trace.
- Double-clicking on each line in the Problems view will, if possible, open the file having errors and highlight the error.

2.2.7 How To Validate, Build, Export, and Package Dataspace Projects from the Command Line

This section describes how to validate, build, export and package Oracle Data Service Integrator dataspace projects from the command line.

- [Section 2.2.7.1, "Data Service File Validation During Deployment"](#)
- [Section 2.2.7.2, "Dataspace Packaging from the Command-line"](#)
- [Section 2.2.7.3, "Syntax Summary"](#)
- [Section 2.2.7.4, "Command-Line Ant Build Targets"](#)
- [Section 2.2.7.5, "Command-line Examples using Ant and Java"](#)

2.2.7.1 Data Service File Validation During Deployment

In the Eclipse IDE a dataspace project's data service (.ds) files are validated automatically according to the following deployment model:

- The publish to server action validates the dataspace projects.
- All the project's artifacts are collected.
- The collected artifacts are deployed to the server.

The Oracle Data Service Integrator Export mechanism allows for a dataspace project's artifacts to be packaged in a JAR the contents of which are identical to what would be generated from the IDE for deployment to an Oracle WebLogic server.

2.2.7.2 Dataspace Packaging from the Command-line

There is also an occasional need for operations such as validate, build, export, and package to be available in a scripting environment. This section describes an Ant script file, `cmdline_build.xml`, provided in the "bin" directory under the Oracle Data Service Integrator installation that can be invoked by a user to:

- Validate a dataspace project
- Generate a deployment JAR file of a dataspace project

Note: For those not wishing or able to use Ant, Java equivalent command-line options are also described.

2.2.7.3 Syntax Summary

This section describes commands and syntax.

Command	Syntax
help	help [{cmd} all]
validate-project	validate-project {project}
validate-dataspace	validate-dataspace {dataspace-path}
export-mediator-client	export-mediator-client {project} {jardir} [jarname.jar]
export-ws-client	export-ws-client {project} {jardir} [jarname.jar] [ws_locator,...]
export-artifacts	export-artifacts {project} {jardir} [jarname.jar]

2.2.7.4 Command-Line Ant Build Targets

This section describes available Oracle Data Service Integrator ant build targets. It describes the following topics:

- [Section 2.2.7.4.1, "Build XML File"](#)
- [Section 2.2.7.4.2, ""help""](#)
- [Section 2.2.7.4.3, "Build Invocation Syntax via Java"](#)
- [Section 2.2.7.4.4, "Build Invocation Syntax via Ant"](#)
- [Section 2.2.7.4.5, ""validate-project""](#)
- [Section 2.2.7.4.6, ""validate-dataspace""](#)
- [Section 2.2.7.4.7, ""export-mediator-client""](#)
- [Section 2.2.7.4.8, ""export-ws-client""](#)
- [Section 2.2.7.4.9, ""export-artifacts""](#)

2.2.7.4.1 Build XML File

The build XML file:

```
cmdline_build.xml
```

will be provided in the directory:

```
<odsi_home>/bin
```

To see a list of build targets with short descriptions in the Ant build XML file, invoke the command below at the prompt window:

```
ant -f <bea_home>/odsi_10.3/bin/cmdline_build.xml -projecthelp
```

Note:

- It is assumed that Ant is available on your computer and is on your path. Some targets require:
 - ECLIPSE_HOME environment variable points to the Eclipse installation directory.
 - javac be available on the PATH variable.
- Commands other than "help" involving an Eclipse project requires specification of the Eclipse workspace directory that contains the project.
 - For Java commands the directory is specified via the "-data" option.
 - For Ant command, it is specified as the "-Dworkspace" property.

2.2.7.4.2 "help"

The "help" target is the default build target. It shows a list of available Oracle Data Service Integrator commands and the syntax needed to invoke the command in Java.

Command	Syntax
help	help [{cmd} "all"]

2.2.7.4.3 Build Invocation Syntax via Java

Note that the syntax shows the portion starting with "odsi_command" below. However, the full syntax to be entered for Java at the prompt window is:

```
java -cp <eclipse_home_dir>/startup.jar org.eclipse.core.launcher.Main -data <workspace_dir>
-application com.bea.dsp.ide.app.runCmdline <odsi_command> <cmd_param> ...
```

2.2.7.4.4 Build Invocation Syntax via Ant

If invoked via Ant, Oracle Data Service Integrator command parameters should be specified as Ant properties. For example, to get help about the "export-artifacts" command, enter:

```
ant -f <odsi_install_dir>/bin/cmdline_build.xml help -Dcmd=export-artifacts
```

To get help on all Oracle Data Service Integrator commands, specify the following property:

```
-Dcmd=all
```

or omit the optional -Dcmd property completely:

```
ant -f <odsi_install_dir>/bin/cmdline_build.xml help
```

2.2.7.4.5 "validate-project"

The "validate-project" target validates the data service (.ds) files in the specified dataspace project. Data service error messages that would appear in the Eclipse IDE's Problems view are sent to stdout when this target is invoked. A "fail" status is returned by this target if any error exists in a .ds data service file in the project.

```
ant -f <odsi_install_dir>/bin/cmdline_build.xml -Dworkspace=/bea/projects/myworkspace
-Dproject=MyODSIProject validate-project
```

Command	Syntax
validate-project	validate-project {project}

2.2.7.4.6 "validate-dataspace"

The validate-dataspace target validates the dataspace project at the specified path.

```
ant -f <odsi_install_dir>/bin/cmdline_build.xml -Ddataspace-path=dataspace-path validate-dataspace
```

Command	Syntax
validate-dataspace	validate-dataspace {dataspace-path}

2.2.7.4.7 "export-mediator-client"

The export-mediator-client target is for generating an Oracle Data Service Integrator mediator client JAR file of a dataspace project.

```
ant -f <odsi_install_dir>/bin/cmdline_build.xml -Dworkspace=/bea/projects/myworkspace
-Dproject=MyODSIProject -Djardir=/temp -Djardirname=myoutput.jar export-mediator-client
```

The default value for the output JAR file name is:

```
<project>-dsp-client.jar
```

Command	Syntax
export-mediator-client	export-mediator-client {project} {jardir} [jardirname.[jar]]

2.2.7.4.8 "export-ws-client"

The `export-ws-client` target generates a web services mediator client JAR file from the specified comma-separated list of wsmmap file locators in a dataspace project.

```
ant -f <odsi_install_dir>/bin/cmdline_build.xml -Dworkspace=/bea/projects/myworkspace
-Dproject=MyODSIProject -Djardir=/temp -Djarname=myoutput.jar
-Dws_locators=ld:customer.ws,ld:order.ws export-ws-client
```

The default value of the `ws_locators` is all wsmmap file locators in the project.

```
<project>-ws-client.jar
```

An example of a wsmmap file locator is:

```
ld:logical/wsmaps/CUSTOMER.ws
```

Command	Syntax
<code>export-ws-client</code>	<code>export-ws-client {project} {jardir} [jarname[.jar]] [ws_locator,...]</code>

2.2.7.4.9 "export-artifacts"

The "export-artifacts" target creates a JAR file containing the definitions and artifacts of the dataspace project.

```
ant -f <odsi_install_dir>/bin/cmdline_build.xml -Dworkspace=/bea/projects/myworkspace
-Dproject=MyODSIProject -Djardir=/temp -Djarname=myoutput.jar export-artifacts
```

The content would be identical to the artifact JAR file created in the IDE. By default, the name of the artifact JAR file is:

```
<project>-artifacts.jar
```

Command	Syntax
<code>export-artifacts</code>	<code>export-artifacts {project} {jardir} [jarname[.jar]]</code>

Notes:

Referenced Java Projects

Since a dataspace project may reference other Java projects in the same Eclipse workspace, you should make certain that:

- Referenced projects in your build script are also built.
- The resulting JAR files and dependent JAR files are copied to the dataspace project's `DSP-INF/lib` directory.

This needs to be done prior to exporting a deployable JAR file using the `export-artifacts` command in order for all referenced/required JAR files to be included in the artifact JAR file.

Invoking Build Commands Without Ant

The Ant targets described in the previous sections are actually implemented in Java. So the actual implementation can be invoked at the prompt window using Java directly -- or any script process -- instead of Ant.

2.2.7.5 Command-line Examples using Ant and Java

This section contains several examples of invoking the Oracle Data Service Integrator command using Ant and Java. It describes the following topics:

- [Section 2.2.7.5.1, "Getting the help text of all the commands using Ant and Java at the prompt window"](#)
- [Section 2.2.7.5.2, "Getting the help text of a specific command using Ant and Java at the prompt window"](#)
- [Section 2.2.7.5.3, "Exporting the artifacts of a dataspace project"](#)

2.2.7.5.1 Getting the help text of all the commands using Ant and Java at the prompt window

Ant:

```
ant -f <bea_home>\odsi_10.3\bin\cmdline_build.xml help -Dcmd=all
```

Java:

```
java -cp <eclipse_home>/eclipse/startup.jar org.eclipse.core.launcher.Main
-application com.bea.dsp.ide.app.runCmdline help all
```

2.2.7.5.2 Getting the help text of a specific command using Ant and Java at the prompt window

Ant:

```
ant -f <bea_home>\odsi_10.3\bin\cmdline_build.xml help -Dcmd=export-artifacts
```

Java:

```
java -cp <eclipse_home>/startup.jar org.eclipse.core.launcher.Main
-application com.bea.dsp.ide.app.runCmdline help export-artifacts
```

2.2.7.5.3 Exporting the artifacts of a dataspace project

This example exports the project:

DspProj

in workspace:

/MyWorkspace

to:

/temp

directory using the default JAR file name:

<project>-artifacts.jar

Ant:

```
ant -f <bea_home>\odsi_10.3\bin\cmdline_build.xml -Dworkspace=/MyWorkspace
-Dproject=DspProj -Djardir=/temp export-artifacts
```

Java:

```
java -cp <eclipse_home>/startup.jar org.eclipse.core.launcher.Main -data /MyWorkspace
/MyWorkspace -application com.bea.dsp.ide.app.runCmdline
-application com.bea.dsp.ide.app.runCmdline export-artifacts DspProj /temp
```

2.3 Reference

This section describes the following topics:

- [Section 2.3.1, "Dataspaces Projects Cheatsheet"](#)
- [Section 2.3.2, "Setting Eclipse for WebLogic Initialization Parameters \(.ini\)"](#)

2.3.1 Dataspaces Projects Cheatsheet

The following table presents the tasks involved with managing dataspaces projects.

Table 2–7 *Dataspaces Projects Cheatsheet*

Task	Location	Action	Comments
Build a project with every save	Project menu	Build Automatically	
Clean selected project	Project menu	Clean	
Close a project	Project Explorer	Right-click > Close Project	
Close a project	Project menu	Project > Close Project	
Close unrelated projects	Project Explorer	Right-click > Close Unrelated Projects	
Copy-Paste a project	Project Explorer	Right-click on project name > Copy-Paste	
Create a new project	File menu	File > New > Dataspaces Project	
Delete selected project from a workspace	Project Explorer	Right-click > Delete	
Export selected project	File menu	Export	
Import selected project into a workspace	File menu	Import	
Open selected project	Project Explorer	Open	
Project properties	Project menu	Right-click > Properties	
Project properties	File menu	Properties	
Publish all the projects in a workspace	Server window	Right-click > Publish	
Refactor Rename-Remove	Project Explorer	Right-click > Refactor > Rename-Remove	Refactor provides for "safe" renaming or deleting of projects or project components.
Refresh display	Project Explorer	Right-click > Refresh File > Refresh	
Remove a previously deployed (configured) project from the server	Server window	Right-click > Delete > Yes > Yes to 'Delete the dataspaces on the server?'	
Rename a project	File menu	Rename	See Safely rename...

Table 2–7 (Cont.) Dataspace Projects Cheatsheet

Task	Location	Action	Comments
Rename a project safely	Project Explorer	Right-click > Refactor > Rename...	
Safely delete selected project from a workspace	Project Explorer	Right-click > Refactor > Delete...	See Safely delete...
Search for a project in a workspace	Search menu	<ul style="list-style-type: none"> ■ Search... ■ Ctrl-h 	
See dataspace currently deployed on the server	Server window	Click the + next to the name of the running server.	
Show project explorer	Window menu	Show View > Project Explorer	
Show project in Service Assembly Modeler	Project Explorer	Right-click > Show in Service Assembly Modeler	
Submit a project to the Oracle Enterprise Repository	Project Explorer	Right-click > Show in Service Assembly Modeler	
Undeploy a project	Server window	Right-click > Delete > Yes > Yes to 'Delete the dataspace on the server?'	
Update a project's metadata	Project Explorer	Right-click > Update metadata	
Validate project artifacts	Project Explorer	Right-click > Validate...	

2.3.2 Setting Eclipse for WebLogic Initialization Parameters (.ini)

Eclipse for WebLogic is an Eclipse plugin. As such, it uses an initialization file that is similar to the eclipse.ini file available in basic installations of Eclipse. This file in turn provides directives to the Java Virtual Machine (JVM) in which the application runs.

The default parameters provided with Oracle Data Service Integrator are located in the following file:

```
{BEA_HOME}/workshop_10.3/workshop.ini
```

Note: When running `workshop.exe` the settings in this file supersede any configuration settings in the `eclipse.ini` file.

By default the `workshop.ini` memory settings are:

```
-vmargs
-Xms384m
-Xmx768m
```

If you encounter out-of-memory errors associated with running an Oracle Data Service Integrator project, try increasing the maximum memory setting. For example:

```
-Xmx1024m
```

2.4 Related Topics

For more information, see ["Example: How to Create Your First Data Services"](#) on page 1-13.

Creating and Updating Physical Data Services

This chapter includes how-to information about a wide range of tasks pertaining to sources such as relation tables and views, stored procedures, SQL statements, and Java functions.

This chapter has the following sections:

- [Section 3.1, "Concepts"](#)
- [Section 3.2, "How to Create Physical Data Services"](#)
- [Section 3.3, "How to"](#)
- [Section 3.4, "Example: XMLBeans Example Using a Metadata-rich Java Class"](#)
- [Section 3.5, "Reference"](#)
- [Section 3.6, "Related Topics"](#)

3.1 Concepts

This section describes the following topics:

- [Section 3.1.1, "Creating Physical Data Services by Importing Source Metadata"](#)
- [Section 3.1.2, "Physical Data Services from Java Functions Overview"](#)

3.1.1 Creating Physical Data Services by Importing Source Metadata

In Oracle Data Service Integrator metadata around a particular data source is developed during the process of creating a physical data service. For example, a list of the tables and columns in a relational database is metadata. A list of operations in a Web service is metadata.

In Oracle Data Service Integrator, a physical data service is typically primarily based on metadata describing the structure of those physical data sources.

Physical data services are the building blocks for the creation of logical data services.

Table 3–1 Data Source Support for Creating Physical Data Services

Source Type	Venue
Relational (including tables, views, stored procedures, and SQL)	JDBC
Web services (WSDL files)	URI, UDDI, WSDL

Table 3–1 (Cont.) Data Source Support for Creating Physical Data Services

Source Type	Venue
Delimited (CSV files)	File-based data, such as spreadsheets.
Java functions (.java)	Programmatic
XML (XML files)	File- or data stream-based XML

When information about physical data is developed during the creation of physical data services, two things happen:

- A physical data service (extension .ds) is created in your Oracle Data Service Integrator-based project.
- A companion schema of the same name (extension.xsd), is created. This schema describes quite exactly the XML type of the data service. Such schemas are placed in a directory named schemas which is a sub-directory of your newly created data service.

3.1.1.1 Source View

The introspection process is done through the Physical Data Service Creation wizard. This wizard introspects available data sources and identifies data objects that can be rendered as operations for either entity or library data services. Once created, physical data services become the building-blocks for queries and logical data services through a series of programs created in the query source.

For example, the following source resulted from importing a Web service operation:

```
(::pragma function <f:function xmlns:f="urn:annotations.ld.oracle.com" kind="read"
nativeName="getCustomerOrderByOrderID"nativeLevel1Container="ElecDBTest"
nativeLevel2Container="ElecDBTestSoap" style="document"/>:

declare function f1:getCustomerOrderByOrderID($x1 as element(t1:getCustomerOrderByOrderID))
as schema-element(t1:getCustomerOrderByOrderIDResponse) external;
```

Notice that the imported Web service is described as a "**read**" function in the pragma. "**External**" refers to the fact that the schema is in a separate file.

For some data sources such as web services imported metadata represents functions which typically return void (in other words, these functions do something other than return data). Such routines are sometimes called side-effecting functions or procedures.

The following source resulted from importing Web service metadata that includes an operation that has been identified as a side-effecting procedure:

```
(::pragma function <f:function xmlns:f="urn:annotations.ld.oracle.com"
kind="hasSideEffects" nativeName="setCustomerOrder" style="document"/>:

declare function f1:setCustomerOrder($x1 as element(t3:setCustomerOrder))
as schema-element(t3:setCustomerOrderResponse) external;
```

In the above pragma the function is identified as "hasSideEffects".

3.1.2 Physical Data Services from Java Functions Overview

In Oracle Data Service Integrator, you can create physical data services based on user-defined functions implemented as Java classes. Oracle Data Service Integrator supports Java functions returning the following types:

- Java primitive types and single-dimension arrays
- Global elements and global element arrays through XMLBean classes
- Global elements and global element arrays through SDO DataObjects

Oracle Data Service Integrator packages operations marked as create, update, or delete functions in an Entity data service. Otherwise, the resulting data service is of type Library. Functions determined to return void are automatically marked as library procedures. When creating a new physical data service, you can change the nominated function type.

The Java method name, when used in an XQuery, becomes the XQuery function name qualified with a namespace.

Note: The following restrictions apply to Java functions:

- Java functions intended for import into a data service must be declared as static
 - Function overloading is based on the number of arguments, not the parameter types
 - Array support is restricted to single-dimension arrays only
 - In functions returning complex types, the return element needs to be extracted from a valid XML document
-
-

3.1.2.1 Simple Java Types and Their XQuery Counterparts

The following outlines the mapping between simple Java types and the corresponding XQuery or schema types:

Table 3–2 Java Types and Corresponding XQuery or Schema Types

Java Simple or Defined Type	XQuery/Schema Type
boolean	xs:boolean
byte	xs:byte
char	xs:char
double	xs:double
float	xs:float
int	xs:int
long	xs:long
short	xs:short
string	xd:string
java.lang.Date	xs:datetime
java.lang.Boolean	xs:boolean
java.math.BigInteger	xs:integer
java.math.BigDecimal	xs:decimal
java.lang.Byte	xs:byte
java.lang.Char	xs:char
java.lang.Double	xs:double

Table 3–2 (Cont.) Java Types and Corresponding XQuery or Schema Types

Java Simple or Defined Type	XQuery/Schema Type
java.lang.Float	xs:float
java.lang.Integer	xs:integer
java.lang.Long	xs:long
java.lang.Short	xs:short
java.sql.Date	xs:date
java.sql.Time	xs:time
java.sql.Timestamp	xs:datetime
java.util.Calendar	xs:datetime

Java functions can consume parameters and return values of the following types:

- Java primitives and types listed in the previous table
- Apache XMLBeans
- Oracle XMLBeans
- SDO DataObject (typed or untyped)

Note: The elements or types referred to in the schema should be global elements.

3.1.2.2 Physical Data Service from a Java Function - Example Code

This topic provides examples showing the use of imported Java functions in an XQuery and the processing of complex types.

3.1.2.2.1 Using a Function Returning an Array of Java Primitives As an example, the Java function `getRunningTotal` can be defined as follows:

```
public static float[] getRunningTotal(float[] list) {
    if (null == list || 1 >= list.length)
        return list;
    for (int i = 1; i < list.length; i++) {
        list[i] = list[i-1] + list[i];
    }
    return list;
}
```

The corresponding XQuery for executing the above function is as follows:

```
Declare namespace f1="ld:javaFunc/float"
Let $y := (2.0, 4.0, 6.0, 8.0, 10.0)
Let $x := f1:getRunningTotal($y)
Return $x
```

The results of the query is as follows:

```
2.0, 6.0, 12.0, 20.0, 30.0
```

3.1.2.2.2 Processing complex types represented via XMLBeans Consider a schema called `Customer` (`customer.xsd`), as shown in the following:

```
<?xml version="1.0" encoding="UTF-8" ?>
  <xs:schema targetNamespace="ld:xml/cust:/BEA_BB10000"
    xmlns:xs="http://www.w3.org/2001/XMLSchema">
    <xs:element name="CUSTOMER">
      <xs:complexType>
        <xs:sequence>
          <xs:element name="FIRST_NAME" type="xs:string" minOccurs="1"/>
          <xs:element name="LAST_NAME" type="xs:string" minOccurs="1"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:schema>
```

You could compile the schema using XMLBeans to generate a Java class corresponding to the types in the schema.

```
xml.cust.beaBB10000.CUSTOMERDocument.CUSTOMER
```

For more information, see <http://xmlbeans.apache.org>.

Following this, you can use the **CUSTOMER** element as shown in the following:

```
public static xml.cust.beaBB10000.CUSTOMERDocument.CUSTOMER[]
  getCustomerListGivenCustomerList(xml.cust.beaBB10000.CUSTOMERDocument.CUSTOMER[]
    ipListOfCust)
  throws XmlException {
  xml.cust.beaBB10000.CUSTOMERDocument.CUSTOMER [] mylocalver = pListOfCust;
  return mylocalver;
}
```

The resulting metadata information produced by the New Physical Data Service wizard will be:

```
(::pragma function <f:function xmlns:f="urn:annotations.ld.oracle.com" kind="datasource"
  access="public">
  <params>
  <param nativeType=" [Lxml.cust.beaBB10000.CUSTOMERDocument$CUSTOMER; "/>
  </params>
  </f:function::)

declare function f1:getCustomerListGivenCustomerList($x1 as element(t1:CUSTOMER)*) as
  element(t1:CUSTOMER)* external;
```

The corresponding XQuery for executing the above function is:

```
declare namespace f1 = "ld:javaFunc/CUSTOMER";
let $z := (
  validate(<n:CUSTOMER xmlns:n="ld:xml/cust:/BEA_BB10000"><FIRST_NAME>John2</FIRST_NAME>
  <LAST_NAME>Smith2</LAST_NAME>
  </n:CUSTOMER>),

  validate(<n:CUSTOMER xmlns:n="ld:xml/cust:/BEA_BB10000"><FIRST_NAME>John2</FIRST_NAME>
  <LAST_NAME>Smith2</LAST_NAME>
  </n:CUSTOMER>),

  validate(<n:CUSTOMER xmlns:n="ld:xml/cust:/BEA_BB10000"><FIRST_NAME>John2</FIRST_NAME>
  <LAST_NAME>Smith2</LAST_NAME>
  </n:CUSTOMER>),

  for $zz in $z
  return
```

3.2 How to Create Physical Data Services

This section describes the following topics:

- [Section 3.2.1, "How To Create Physical Data Services from Relational Tables and Views"](#)
- [Section 3.2.2, "How To Create Physical Data Services from Stored Procedures"](#)
- [Section 3.2.3, "How To Create Physical Data Services Based on SQL Statements"](#)
- [Section 3.2.4, "How To Create Physical Data Services Based on Database Functions"](#)
- [Section 3.2.5, "How To Create a Physical Data Service from a Web Service"](#)
- [Section 3.2.6, "Preparing to Create Physical Data Services From Java Functions"](#)
- [Section 3.2.7, "How To Create a Physical Data Service from a Java Function"](#)
- [Section 3.2.8, "How To Create a Physical Data Service from XML Data"](#)
- [Section 3.2.9, "How To Create a Physical Data Service from a Delimited File"](#)

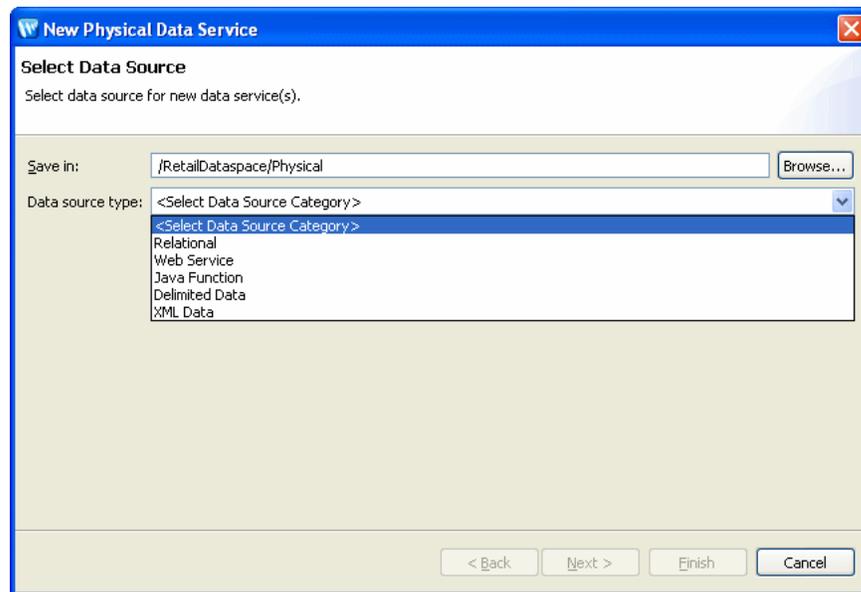
3.2.1 How To Create Physical Data Services from Relational Tables and Views

The following topics describe how to create physical data services from relational tables and views:

- [Section 3.2.1.1, "Setting Up the Physical Data Service Creation Wizard"](#)
- [Section 3.2.1.2, "Setting Up the Import Wizard for Relational Objects"](#)
- [Section 3.2.1.3, "Selecting SQL Table and View Objects for Import"](#)
- [Section 3.2.1.4, "Setting Properties for New Data Service Operations"](#)
- [Section 3.2.1.5, "Verifying Data Service Composition"](#)
- [Section 3.2.1.6, "Database-specific Catalog and Schema Considerations"](#)
- [Section 3.2.1.7, "XML Name Conversion Considerations"](#)

3.2.1.1 Setting Up the Physical Data Service Creation Wizard

Physical data services are created using a wizard.

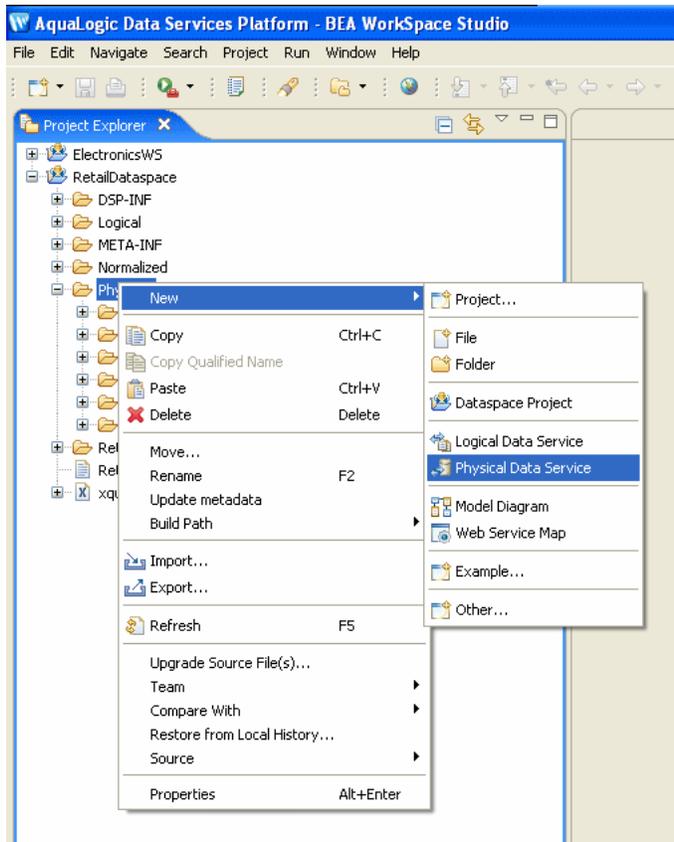
Figure 3–1 Physical Data Service Creation Wizard

Browse to specify a directory where you want to save the data service. The possible data source types are Relational, Web Service, Java Function, Delimited Data, and XML data.

3.2.1.1.1 Starting the Wizard To start the physical data service creation wizard:

- Right-click on your dataspace project or any folder in your project.
- Choose **New > Physical Data Service**

Figure 3–2 Creating a New Physical Data Service



In the Project Explorer, a project is selected and the context window is displayed. New is highlighted and Physical Data Service is selected.

3.2.1.2 Setting Up the Import Wizard for Relational Objects

When importing a relational object available options include the ability to:

1. Set a location for your new data service to be saved within your project.
2. Select a data source from the dropdown listbox.
3. Select the database type for the selected source (**PointBase** for the sample RDBMS) from the dropdown listbox.
4. Select among the relational source types listed in the following table.

Table 3–3 Types of available relational data sources

Relational Type	Description
Tables and Views	Displays all public tables and views in the selected data source.
Stored Procedures	Displays all public stored procedures in the selected data source.
SQL Statement	Allows creation of a SQL statement for extracting relational data from the data source.
Database Function	Allows creation of an XQuery function in a library data service based on build-in or custom database functions.

3.2.1.3 Selecting SQL Table and View Objects for Import

To create a physical data service based on a relational table or view:

1. Select the **Tables and Views** option
2. Click **Next**.

A list of available database table and view SQL objects appears.

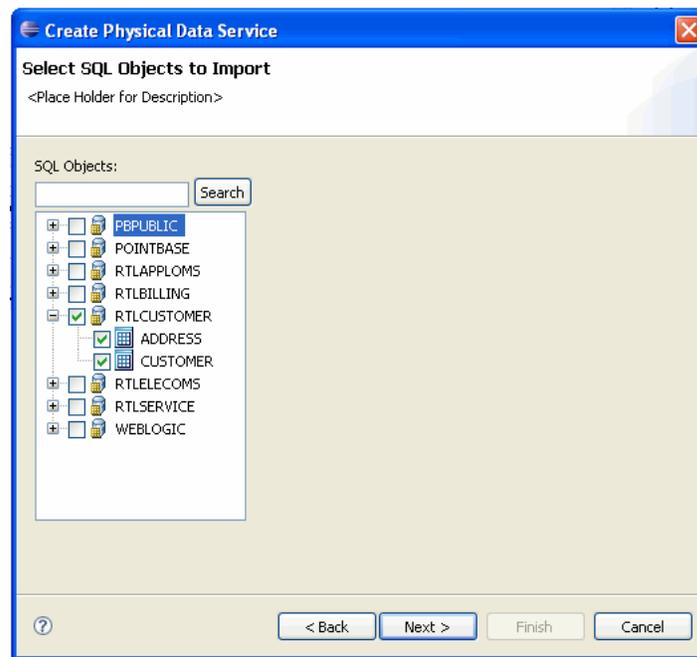
Objects are grouped based on the relational data sources catalog and/or schema.

In the example of an **RTLCUSTOMER** catalog, the **ADDRESS** and **CUSTOMER** tables both become physical data services.

For more information, see "[Database-specific Catalog and Schema Considerations](#)" on page 3-12.

Simply check the desired objects or their container, which will select all enclosed tables or views.

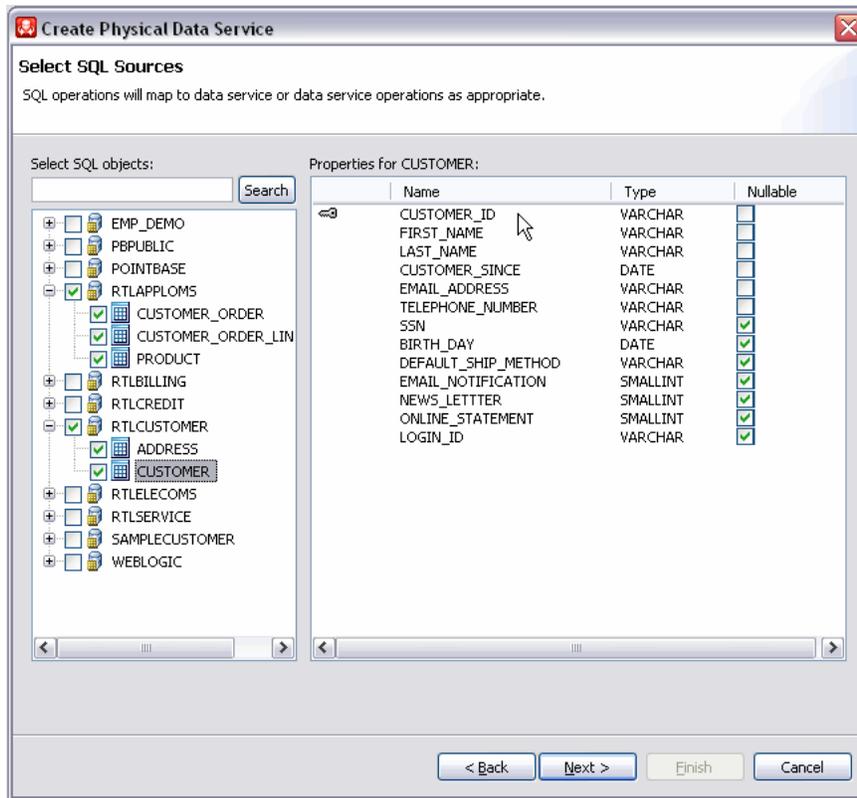
Figure 3-3 Table and View Objects Selected for Import



In the Select Objects to Import dialog, the **RTLCUSTOMER**, **ADDRESS**, and **CUSTOMER** objects are selected.

If you click on an individual object such as **ADDRESS** or **CUSTOMER**, information describing the database's primary key(s), column name, type and nullability appears. For example the **CUSTOMER** table contains a **CUSTOMER_ID** field of type **VARCHAR**. That column is not nullable, meaning that it must be supplied with any updates.

Figure 3–4 Physical Data Service Properties



The Select SQL Sources dialog shows an object tree on the left, with CUSTOMER selected, and a panel with four columns on the right for CUSTOMER properties: key, name, type, and nullable.

3.2.1.3.1 Filtering SQL Objects Using Search The Search option available when creating a physical data service can be especially useful when:

- You know specific names of the data source objects you want to turn into data services.
- Your data source may be so large that a filter is needed.
- You may be looking for objects with specific naming characteristics such as:

`%audit2003%`

The above search command retrieves all objects that contain the enclosed string.

3.2.1.3.2 Using JDBC Syntax to Search SQL Objects You can search through available SQL objects using standard JDBC wildcard syntax.

- An underscore (`_`) creates a wildcard for an individual character.
- A percentage sign (`%`) indicates a wildcard for a string. Entries are case-sensitive.

Another example:

`CUST%, PAY%`

entered in the Tables/Views field the above search string returns all tables and views starting with either CUST or PAY.

3.2.1.3.3 Special Considerations When Searching Stored Procedures If no items are entered for a particular field, all matching items are retrieved. For example, if no filtering entry is made for the Procedure field, all stored procedures in the data object will be retrieved.

3.2.1.4 Setting Properties for New Data Service Operations

Each new entity data service is created with a Read function that contains all the metadata elements identified during data service creation. It can be thought of as comparable to the following construct in the relational world:

```
select * from <table>
```

Use the Properties dialog to:

- Optionally modify the operation name.
- Set the **Public** option (check if you want your function to be available to client applications).
- Set the kind of operation (in some cases only **Read** will be available).
- Set the **Primary** option (check if you want your function to be the primary of its type).

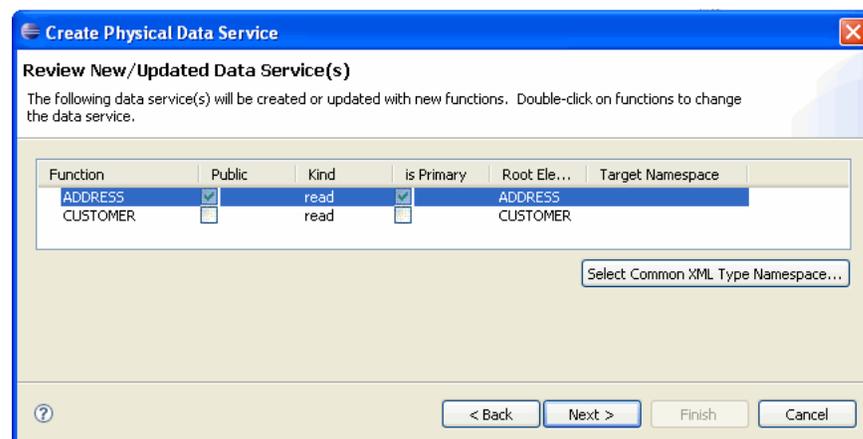
Note: In some cases this option may not be available.

- Select a common XML namespace for the entire data service or individual target namespaces for specific operations.
- Set the target namespace.

The root element, which is read-only, is also displayed.

Note: Initially the root element name matches the name of the data service.

Figure 3–5 Setting Properties for New Data Service Functions



There are six columns on the Review New/Updated Data Service(s) dialog: Function, Public, Kind, is Primary, Root Element, and Target Namespace. For the ADDRESS function, the Public checkbox is selected, the read Kind is displayed, the is Primary checkbox is selected, the Root Element is ADDRESS, and the Target Namespace column is blank.

3.2.1.4.1 Default Naming Conventions There are several default naming conventions associated with new data services:

- When a table, view, or other data source object is the source for a data service, the nominated name is wherever possible the same as the source object name. In some cases, however, names are adjusted to conform with XML naming conventions.
- Initially the root element name matches the name of the data service.

For more information, see [XML Name Conversion Considerations](#).

3.2.1.5 Verifying Data Service Composition

On the Review New Data Service(s) page you can set, confirm or, optionally, change suggested data service names depending on the type of physical data service you are creating.

3.2.1.5.1 Default Physical Data Service Names The nominated name for a new data service is, wherever possible, the same as the source object name. In some cases, however, names are adjusted to conform with XML naming conventions.

3.2.1.5.2 About Automatic Data Service Name Changes Name conflicts occur when there is a data service of the same name present in the target directory. Name conflicts are highlighted in red.

There are several situations where you will need to change the name of your data service:

- There already is a data service of the same name in your application.
- You are trying to create multiple data services with the same name.

Data services always have the file extension:

.ds

3.2.1.6 Database-specific Catalog and Schema Considerations

Database vendors variously support database catalogs and schemas.

Table 3–4 Vendor Support for Catalog and Schema Objects

Vendor	Catalog	Schema
Oracle	Does not support catalogs. When specifying database objects, the catalog field should be left blank.	Typically the name of an Oracle user ID.
DB2	If specifying database objects, the catalog field should be left blank.	Schema name corresponds to the catalog owner of the database, such as db2admin.
Sybase	Catalog name is the database name.	Schema name corresponds to the database owner.

Table 3–4 (Cont.) Vendor Support for Catalog and Schema Objects

Vendor	Catalog	Schema
Microsoft SQL Server	Catalog name is the database name.	Schema name corresponds to the catalog owner, such as dbo. The schema name must match the catalog or database owner for the database to which you are connected.
Informix	Does not support catalogs. If specifying database objects, the catalog field should be left blank.	Not needed.
PointBase	PointBase database systems do not support catalogs. If specifying database objects, the catalog field should be left blank.	Schema name corresponds to a database name.

3.2.1.7 XML Name Conversion Considerations

When a source name is encountered that does not fit within XML naming conventions, default generated names are converted according to rules described by the SQLX standard. Generally speaking, an invalid XML name character is replaced by its hexadecimal escape sequence (having the form xUUUU).

For additional details see section 9.1 of the W3C draft version of this standard:

<http://www.sqlx.org/SQL-XML-documents/5WD-14-XML-2003-12.pdf>

3.2.2 How To Create Physical Data Services from Stored Procedures

Stored procedures are database objects that group an executable set of SQL and native database programming language statements together to perform a specific task locally. Advanced DBMS systems utilize stored procedures to improve query performance, manage and schedule data operations, enhance security, and so forth.

In Oracle Data Service Integrator you can, for specifically supported databases, create physical data services based on stored procedures.

It is often convenient to leverage independent routines as part of managing enterprise information through a data service. An obvious example would be to leverage standalone update or security functions through data services. Such functions have no XML type; in fact, they typically return nothing (or void).

Stored procedures are very often side-effecting from the perspective of the data service, since they perform internal operations on data. In such cases all you need to do is identify the stored procedures as a data service procedure when your physical data service is created.

After you have identified the stored procedures that you want to add to your data service, you also have an opportunity to identify which of these should be identified as data service procedures.

Each stored procedure that is imported becomes a separate data service. In other words, if you have five stored procedures, you will create five data services.

The following topics describe how to create a physical data service from a stored procedure:

- [Section 3.2.2.1, "Importing Stored Procedure Metadata Using the Physical Data Service Creation Wizard"](#)
- [Section 3.2.2.2, "Setting Up the Physical Data Service Creation Wizard"](#)
- [Section 3.2.2.3, "Setting Up the Import Wizard for Relational Objects"](#)

- [Section 3.2.2.4, "Selecting Stored Procedure Objects for Import"](#)
- [Section 3.2.2.5, "Configuring Selected Stored Procedures"](#)
- [Section 3.2.2.6, "Stored Procedure Configuration Reference"](#)
- [Section 3.2.2.7, "Setting Properties for New Data Service Operations"](#)
- [Section 3.2.2.8, "Verifying Data Service Composition"](#)
- [Section 3.2.2.9, "Adding Operations to an Existing Data Service"](#)
- [Section 3.2.2.10, "Support for Stored Procedures in Popular Databases"](#)

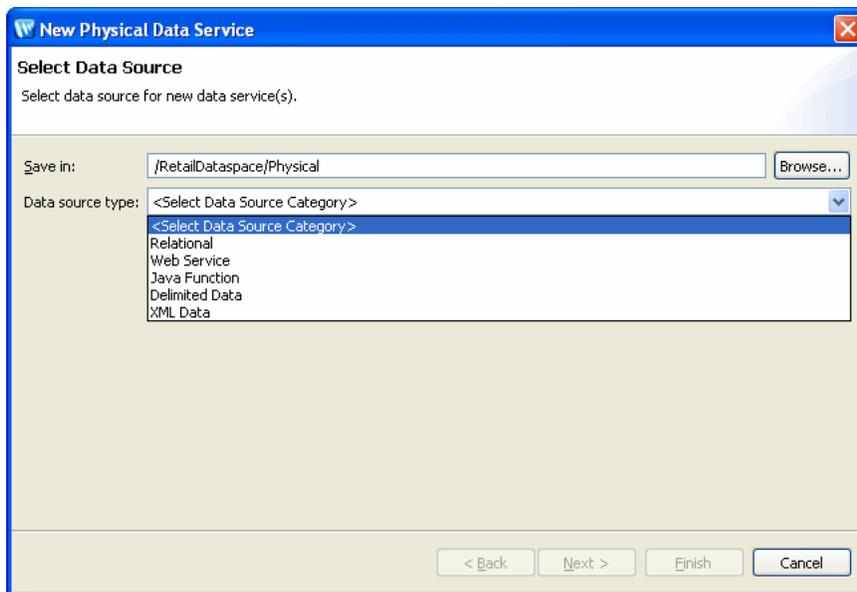
3.2.2.1 Importing Stored Procedure Metadata Using the Physical Data Service Creation Wizard

The following topics cover the actions necessary to create physical data services from relational stored procedures.

3.2.2.2 Setting Up the Physical Data Service Creation Wizard

Physical data services are created using a wizard.

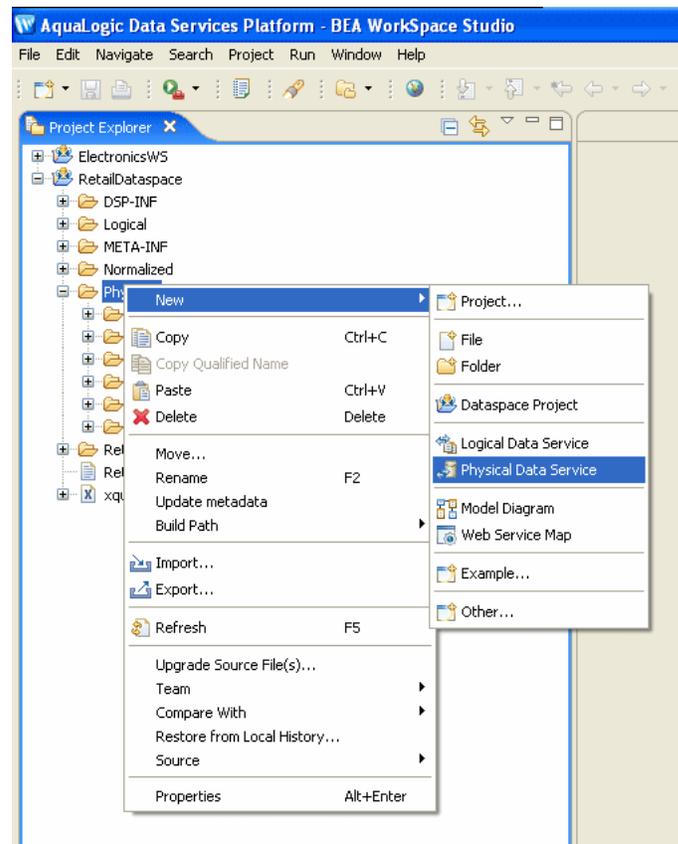
Figure 3–6 Physical Data Service Creation Wizard



Browse to specify a directory where you want to save the date service. The possible data source types are Relational, Web Service, Java Function, Delimited Data, and XML data.

3.2.2.2.1 Starting the Wizard

- To start the physical data service creation wizard:
1. Right-click on your dataspace project or any folder in your project.
 2. Choose New > Physical Data Service

Figure 3–7 Creating a New Physical Data Service

In the Project Explorer, a project is selected and the context window is displayed. New is highlighted and Physical Data Service is selected.

3.2.2.3 Setting Up the Import Wizard for Relational Objects

When importing a relational object available options include the ability to:

1. Set a location for your new data service to be saved within your project.
2. Select a data source from the dropdown listbox.
3. Select the database type for the selected source (PointBase for the sample RDBMS) from the dropdown listbox.
4. Select among the relational source types listed in the following table.

Table 3–5 Types of available relational data sources

Relational Type	Description
Tables and Views	Displays all public tables and views in the selected data source.
Stored Procedures	Displays all public stored procedures in the selected data source.
SQL Statement	Allows creation of a SQL statement for extracting relational data from the data source.
Database Function	Allows creation of an XQuery function in a library data service based on build-in or custom database functions.

3.2.2.4 Selecting Stored Procedure Objects for Import

To create physical data services based on stored procedures:

1. Select the **Stored Procedures** option.
2. Click **Next**.

A list of available stored procedures appears.

Objects are grouped based on the relational data sources catalog and/or schema.

You can use wildcards to support importing metadata on internal stored procedures. For example, entering the following string as a stored procedure filter:

```
%TRIM%
```

retrieves metadata on the system stored procedure:

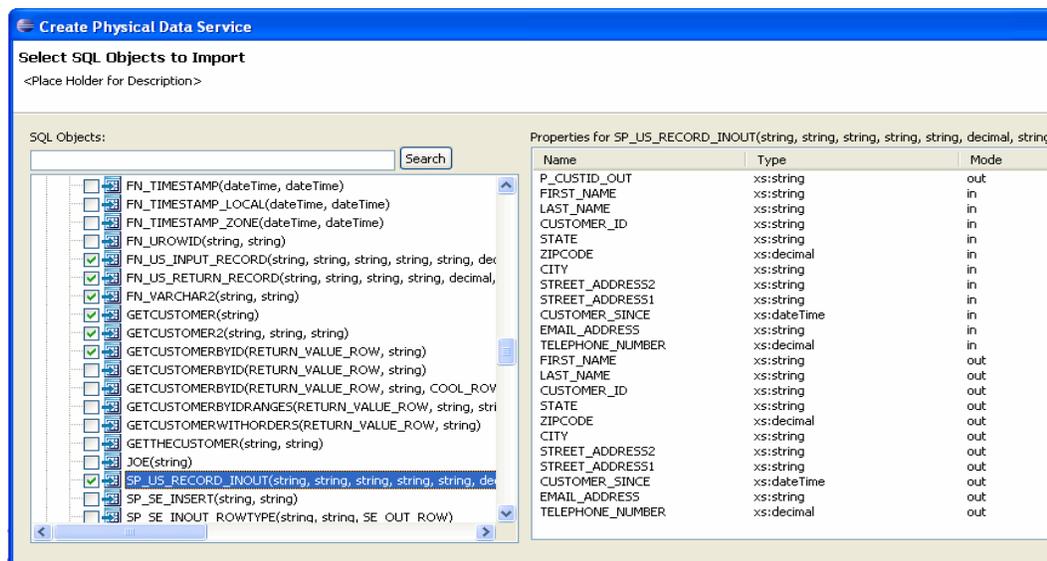
```
STANDARD.TRIM
```

In such a situation you may want to make a "nonsense" entry in the **Table/View** field in order to avoid retrieving all tables and views in the database.

For more information, see ["Database-specific Catalog and Schema Considerations"](#) on page 3-12.

Simply check the desired objects or their container, which will select all enclosed stored procedures.

Figure 3–8 Stored Procedure Objects Selected for Import



On the Select SQL Objects to Import dialog, there are two columns. On the left there is a list of SQL objects. On the right is a list of properties. There are checkboxes next to several SQL objects. The SP_US_RECORD_INOUT object is checked and selected, and its properties appear in the right column.

3.2.2.4.1 Filtering SQL Objects Using Search The Search option available when creating a physical data service can be especially useful when:

- You know specific names of the data source objects you want to turn into data services.
- Your data source may be so large that a filter is needed.
- You may be looking for objects with specific naming characteristics such as:

```
%audit2003%
```

3.2.2.4.2 Using JDBC Syntax to Search SQL Objects You can search through available SQL objects using standard JDBC wildcard syntax.

- An underscore (_) creates a wildcard for an individual character.
- A percentage sign (%) indicates a wildcard for a string. Entries are case-sensitive.

Another example:

```
CUST%, PAY%
```

entered in the **Tables/Views** field the above search string returns all tables and views starting with either **CUST** or **PAY**.

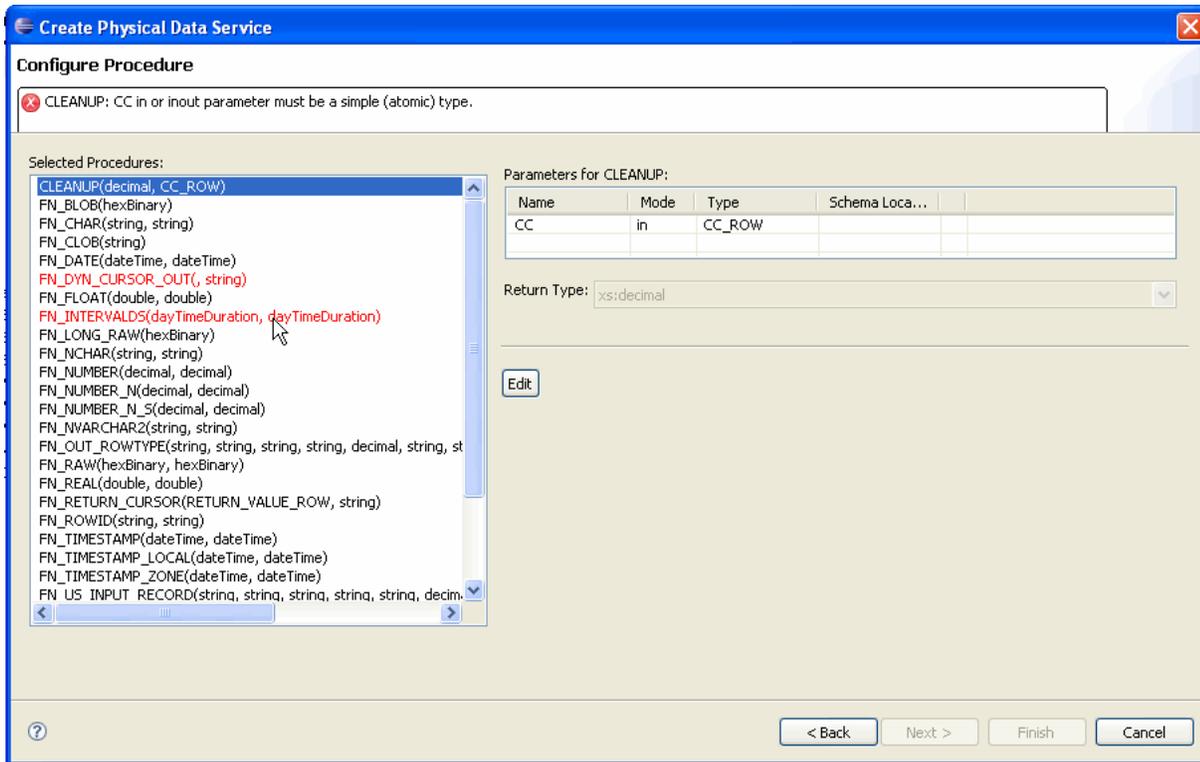
3.2.2.4.3 Special Considerations When Searching Stored Procedures If no items are entered for a particular field, all matching items are retrieved. For example, if no filtering entry is made for the Procedure field, all stored procedures in the data object will be retrieved.

3.2.2.5 Configuring Selected Stored Procedures

When Oracle Data Service Integrator introspects a stored procedure, the process may not be complete. For example, a required item of information such as a schema file or type cannot be determined. When such introspection problems occur, the stored procedure in question is highlighted in red. This setting means that additional information about the procedure must be provided by the user before the data service can be created.

Your goal in correcting an "<unknown>" condition associated with a stored procedure is to bring the metadata obtained by the import wizard into conformance with the actual metadata of the stored procedure. In some cases this will be by correcting the location of the return type. In others you will need to adjust the type associated with an element of the procedure or add elements that were not found during the initial introspection of the stored procedure.

Figure 3–9 Configure Stored Procedure Dialog



On the Configure Stored Procedure dialog, error messages are displayed for selected procedures. The problem procedures are highlighted in red in the column on the left. A table containing the name, mode, type, and schema location is on the right. An error message for a selected procedure is displayed at the top of the dialog. Use the Edit button to correct configuration settings.

When several stored procedures are selected at the same time for physical data service creation, all the selected procedures must be adequately configured before any data services based on the procedures can be created.

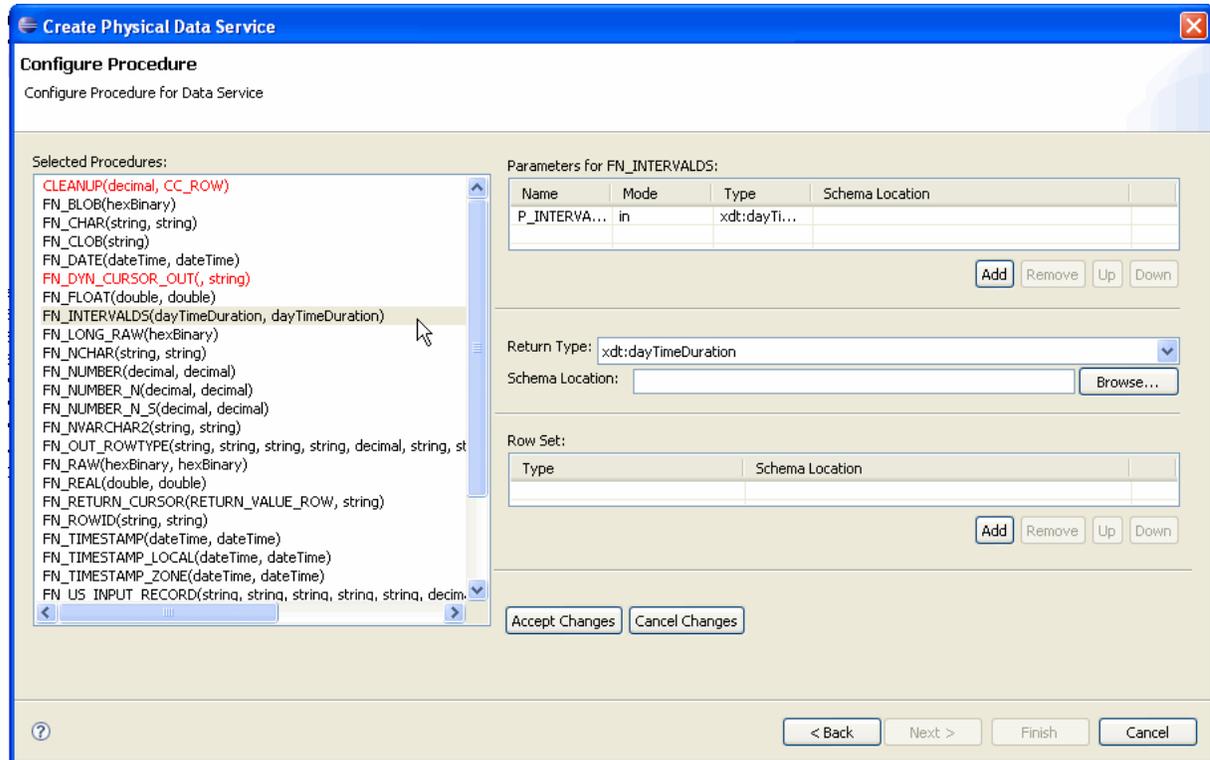
Note: An alternative to configuring an incomplete stored procedure before proceeding is to use the wizard Back button to de-select the procedure in question.

Here are the steps involved in editing a set of stored procedures that will be imported as data services:

1. Scroll through the list of selected procedures.
2. For each procedure in red type, use the **Edit** button to correct the configuration settings.
3. Make any other changes. (In some cases the data architect may know of requirements that are not identified during the introspection process.)
4. Click **Next** when all the procedures in the selected set are valid.

3.2.2.5.1 Editing Stored Procedure Configurations Stored procedure configuration can be complicated. An understanding of the characteristics of the stored procedure in your database is an essential prerequisite. This section describes stored procedure options in detail.

Figure 3–10 Stored Procedure Metadata Editing Options



The Configure Procedure dialog lets you edit parameters for selected procedures. For a selected procedure and parameter, return type and schema location fields are displayed. Browse to specify the schema location. You can set a row by specifying the type and schema location of the row. Buttons let you accept changes or cancel changes.

Once in stored procedure configuration edit mode, options are available in three general areas:

- **Parameters.** Stored procedures requiring complex parameters can only be turned into data services once a schema has been identified. In addition, retrieved information on parameters required by a stored procedure may be incorrect. For example, additional parameters may be needed.
- **Return type.** Stored procedures returning complex data require a local schema to handle data returned from the call. In addition, retrieved information on stored procedure return types may be incorrect or it may be the case that no returned data is wanted.
- **Row set.** A row set identifies a schema and its associated library data service to hold information returned by a stored procedure. In some cases multiple row sets may need to be specified.

Table 3–6 Stored Procedure Editing Options

Category	Options	Settings	Discussion
Parameters	Name	Parameter name	Editable
	Mode	on/out/inout	
	Type	XQuery type	May be derived from the stored procedure. Primitive XQuery type settings are also available.
	Schema location	XSD file	Schema file must be in the project.
Return type	Type	XQuery type or global type from selected schema	
	Schema location	XSD file	Schema file must be in the project.
Row set	Type	Data service	Derived from selected schema.
	Schema location	XSD file	Schema file must be in the project.

3.2.2.6 Stored Procedure Configuration Reference

The following topics provide detailed information regarding various configuration options associated with creating data services based on stored procedures.

3.2.2.6.1 In Mode, Out Mode, Inout Mode In, Out, and Inout mode settings determine how a parameter passed to a stored procedure is handled.

Parameter Mode	Effect
In	Parameter is passed by reference or value.
Inout	Parameter is passed by reference.
Out	Parameter is passed by reference. However the parameter being passed is first initialized to a default value. If your stored procedure has an OUT parameter requiring a complex element, you may need to provide a schema.

3.2.2.6.2 Procedure Profile Each element in a stored procedure is associated with a type. If the item is a simple type, you can simply choose from the pop-up list of types. If the type is complex, you may need to supply an appropriate schema. Click on the schema location button and either enter a schema pathname or browse to a schema. The schema must reside in your application.

After selecting a schema, both the path to the schema file and the URI appear.

3.2.2.6.3 Complex Parameter Types Complex parameter types are supported under only three conditions:

- As the output parameter
- As the Return type
- As a rowset

3.2.2.6.4 About Rowsets A rowset type is a complex type.

In some cases the wizard can automatically detect the structure of a rowset and create an element structure. However, if the structure is unknown, you will need to provide it.

Note: All rowset-type definitions must conform to this structure.

The name of the rowset type can be:

- The parameter name (in case of a input/output or output only parameter).
- An assigned name.
- The referenced element name (result rowsets) in a user-specified schema.

Not all databases support rowsets. In addition, JDBC does not report information related to defined rowsets.

3.2.2.6.5 Using Rowset Information In order to create data services from stored procedures that use rowset information, you need to supply the correct ordinal (matching number) and a schema. If the schema has multiple global elements, select the one you want from the Type column. Otherwise the type used match the first global element in your schema file.

The order of rowset information is significant; it must match the order in your data source. Use the **Move Up** / **Move Down** commands to adjust the ordinal number assigned to the rowset.

Note: XML types in data services generated from stored procedures do not display native types. However, you can view the native type in the Source editor; it is located in the pragma section.

3.2.2.6.6 Stored Procedure Version Support Only the most recent version of a particular stored procedure can be imported into Oracle Data Service Integrator. For this reason you cannot identify a stored procedure version number when creating a physical data service based on a stored procedure. Similarly, adding a version number for your stored procedure in the Source editor will result in a query exception.

3.2.2.6.7 Supporting Stored Procedures with Nullable Input Parameter(s) If you know that an input parameter of a stored procedure is nullable (can accept null values), you can change the signature of the function in Source View to make such parameters optional by adding a question mark at end of the parameter.

For example (question-mark (?) shown in bold):

```
function myProc($arg1 as xs:string) ...
```

would become:

```
function myProc($arg1 as xs:string?) ...
```

3.2.2.7 Setting Properties for New Data Service Operations

Each new entity data service is created with a Read function that contains all the metadata elements identified during data service creation. It can be thought of as comparable to the following construct in the relational world:

```
select * from <table>
```

Use the Properties dialog to:

- Optionally modify the operation name.
- Set the **Public** option (check if you want your function to be available to client applications).
- Set the kind of operation (in some cases only **Read** will be available).
- Set the **Primary** option (check if you want your function to be the primary of its type).

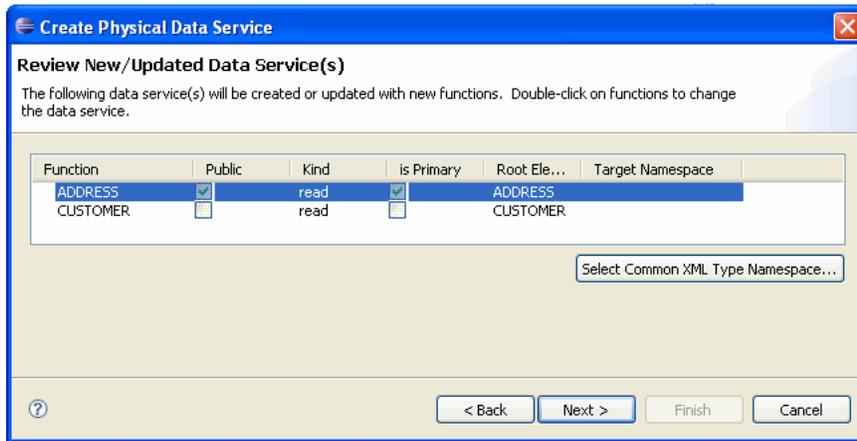
Note: In some cases this option may not be available.

- Select a common XML namespace for the entire data service or individual target namespaces for specific operations.
- Set the target namespace.

The root element, which is read-only, is also displayed.

Note: Initially the root element name matches the name of the data service.

Figure 3–11 Setting Properties for New Data Service Functions



On the Review New/Updated Service(s) dialog, selected data services will be updated with new functions. In the table, function ADDRESS is selected. Public is checked, Kind is read, is Primary is checked, and Root Element is ADDRESS. There is a Selected Common XML Type Namespace button. There are Back, Next, Finish, and Cancel buttons.

3.2.2.7.1 Default Naming Conventions There are several default naming conventions associated with new data services:

- When a table, view, or other data source object is the source for a data service, the nominated name is wherever possible the same as the source object name. In some cases, however, names are adjusted to conform with XML naming conventions.

- Initially the root element name matches the name of the data service.

For more information, see XML Name Conversion Considerations.

3.2.2.8 Verifying Data Service Composition

On the Review New Data Service(s) page you can set, confirm or, optionally, change suggested data service names depending on the type of physical data service you are creating.

3.2.2.8.1 Default Physical Data Service Names The nominated name for a new data service is, wherever possible, the same as the source object name. In some cases, however, names are adjusted to conform with XML naming conventions.

For more information, see XML Name Conversion Considerations.

3.2.2.8.2 About Automatic Data Service Name Changes Name conflicts occur when there is a data service of the same name present in the target directory. Name conflicts are highlighted in red.

There are several situations where you will need to change the name of your data service:

- There already is a data service of the same name in your application.
- You are trying to create multiple data services with the same name.

Data services always have the file extension:

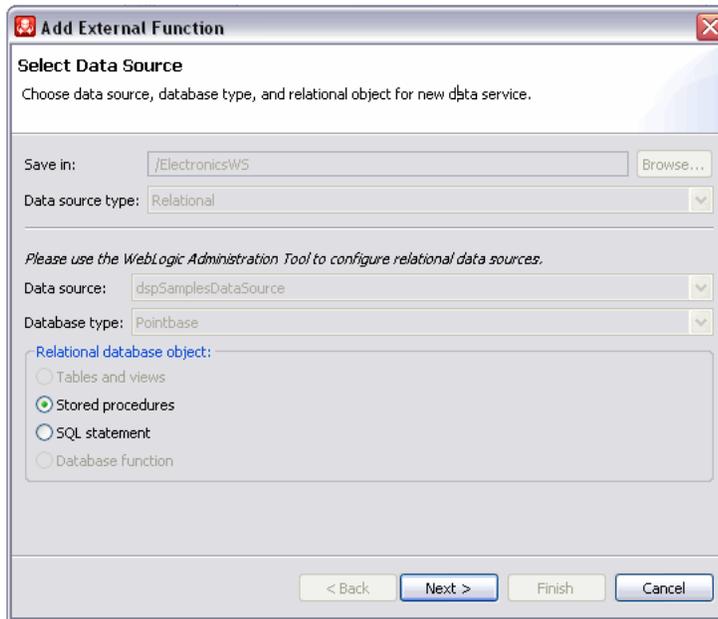
.ds

3.2.2.9 Adding Operations to an Existing Data Service

You can add SQL statement or stored procedure operations based on the same data source to an existing physical data service based a stored procedure.

For more information, see ["How To Add an External Function to an Existing Physical Data Service"](#) on page 3-91.

Figure 3–12 Adding a Stored Procedure or SQL Statement to a Data Service



In the Select Data Source dialog, you can choose the data source, datasource type, and relational object for a new data service. In the Relational database object section of the dialog, the Stored Procedures option is selected. All other options are disabled except for the SQL statement option.

3.2.2.10 Support for Stored Procedures in Popular Databases

Each database vendor approaches stored procedures differently. Oracle Data Service Integrator support limitations generally reflect JDBC driver limitations.

3.2.2.10.1 General Restrictions There are several restrictions that apply to stored procedures generally:

- Oracle Data Service Integrator does not support rowset as an input parameter.
- Only data types supported by Oracle Data Service Integrator can be imported as part of stored procedures.

Note: For a list of database types supported by Oracle Data Service Integrator XQuery-SQL Mapping Reference.

3.2.2.10.2 Oracle Stored Procedure Support The following table describes data service creation support for Oracle stored procedures.

Term	Usage
Procedure types	<ul style="list-style-type: none"> ■ Procedures ■ Functions ■ Packages

Term	Usage
Parameter modes	<ul style="list-style-type: none"> ■ Input only ■ Output only ■ Input/Output ■ None
Parameter data types	<p>Any Oracle PL/SQL data type except:</p> <ul style="list-style-type: none"> ■ ROWID ■ UROWID <p>Note: When defining function signatures, note that the Oracle %TYPE and %ROWTYPE types must be translated to XQuery types that match the true types underlying the stored procedure's %TYPE and %ROWTYPE declarations. %TYPE declarations map to simple types; %ROWTYPE declarations map to rowset types.</p>
Data returned from a function	Oracle supports returning PL/SQL data types such as NUMBER, VARCHAR, %TYPE, and %ROWTYPE as parameters.
Comments	<p>The following identifies limitations associated with importing Oracle database procedure metadata.</p> <ul style="list-style-type: none"> ■ The data service creation process can only detect the data structure for cursors that have a binding PL/SQL record. For a dynamic cursor you need to manually specify the cursor schema. ■ Data from a PL/SQL record structure cannot be retrieved due to an Oracle JDBC driver limitations. ■ The Oracle JDBC driver supports rowset output parameters only if they are defined as reference cursors in a package. ■ The Oracle JDBC driver does not support NATURALN and POSITIVEN as output only parameters.

3.2.2.10.3 Sybase Stored Procedure Support The following table describes data service creation support for Sybase stored procedures.

Term	Usage
Procedure types	<ul style="list-style-type: none"> ■ Procedures ■ Grouped procedures ■ Functions are categorized as a scalar or inline table-valued and multi-statement table-valued function. Inline table-valued and multi-statement table-valued functions return rowsets.
Parameter modes	<ul style="list-style-type: none"> ■ Input only ■ Output only
Parameter data types	For a list of database types supported by Oracle Data Service Integrator, see the XQuery-SQL Mapping Reference.
Data returned from a function	<p>Sybase functions supports returning a single value or a table. Procedures return data in the following ways:</p> <ul style="list-style-type: none"> ■ As output parameters, which can return either data (such as an integer or character value). ■ As return codes, which are always an integer value. ■ As a rowset for each SELECT statement contained in the stored procedure or any other stored procedures called by that stored procedure. ■ As a global cursor that can be referenced outside the stored procedure supports, returning single value or multiple values.

Term	Usage
Comments	<p>The following identifies limitations associated with importing Sybase database procedure metadata:</p> <ul style="list-style-type: none"> ▪ The Sybase JDBC driver does not support input/output or output only parameters that are rowsets (including cursor variables). ▪ The Jconnect driver and some versions of the Oracle Sybase driver cannot detect the parameter mode of the procedure. In such a case, the return mode will be UNKNOWN, preventing importation of the metadata. To proceed, you need to set the correct mode.

3.2.2.10.4 IBM DB2 Stored Procedure Support The following table describes data service creation support for IBM DB2 stored procedures.

Term	Usage
Procedure types	<ul style="list-style-type: none"> ▪ Procedures ▪ Functions ▪ Packages where each function is also categorized as a scalar, column, row, or table function. ▪ A scalar function returns a single-valued answer each time it is called. ▪ A column function is one which conceptually is passed a set of like values (a column) and returns a single-valued answer (AVG()). ▪ A row function is a function that returns one row of values. ▪ A table function is a function that returns a table to the SQL statement that referenced it.
Parameter modes	<ul style="list-style-type: none"> ▪ Input only ▪ Output only ▪ Input/output
Parameter data types	For a list of database types supported by Oracle Data Service Integrator see the XQuery-SQL Mapping Reference. For a list of database types supported by Oracle Data Service Integrator, see the XQuery-SQL Mapping Reference.
Data returned from a function	DB2 supports returning a single value, a row of values, or a table.
Comments	<p>The following identifies limitations associated with creating physical data services based on DB2 stored procedures:</p> <ul style="list-style-type: none"> ▪ Column type functions are not supported. ▪ Rowsets as output parameters are not supported. ▪ The DB2 JDBC driver supports float, double, and decimal input only and output only parameters. Float, double, and decimal data types are not supported as input/output parameters.

3.2.2.10.5 Microsoft SQL Server Stored Procedure Support The following table describes data service creation support for Microsoft stored procedures.

Term	Usage
Procedure types	SQL Server supports procedures, grouped procedures, and functions. Each function is also categorized as a scalar or inline table-valued and multi-statement table-valued function. Inline table-valued and multi-statement table-valued functions return rowsets.
Parameter modes	SQL Server supports input only and output only parameters.

Term	Usage
Parameter data types	SQL Server procedures/functions support any SQL Server data type as a parameter. For a list of database types supported by Oracle Data Service Integrator, see the XQuery-SQL Mapping Reference.
Data returned from a function	<p>SQL Server functions supports returning a single value or a table. Data can be returned in the following ways:</p> <ul style="list-style-type: none"> ■ As output parameters, which can return either data (such as an integer or character value) or a cursor variable (cursors are rowsets that can be retrieved one row at a time). ■ As return codes, which are always an integer value. ■ As a rowset for each SELECT statement contained in the stored procedure or any other stored procedures called by that stored procedure.
Comments	<p>The following identifies limitations associated with importing SQL Server procedure metadata.</p> <ul style="list-style-type: none"> ■ Result sets returned from SQL server (as well as those returned from Sybase) are not detected automatically. Instead you will need to manually add parameters as a result. ■ The Microsoft SQL Server JDBC driver does not support rowset input/output or output only parameters (including cursor variables).

3.2.3 How To Create Physical Data Services Based on SQL Statements

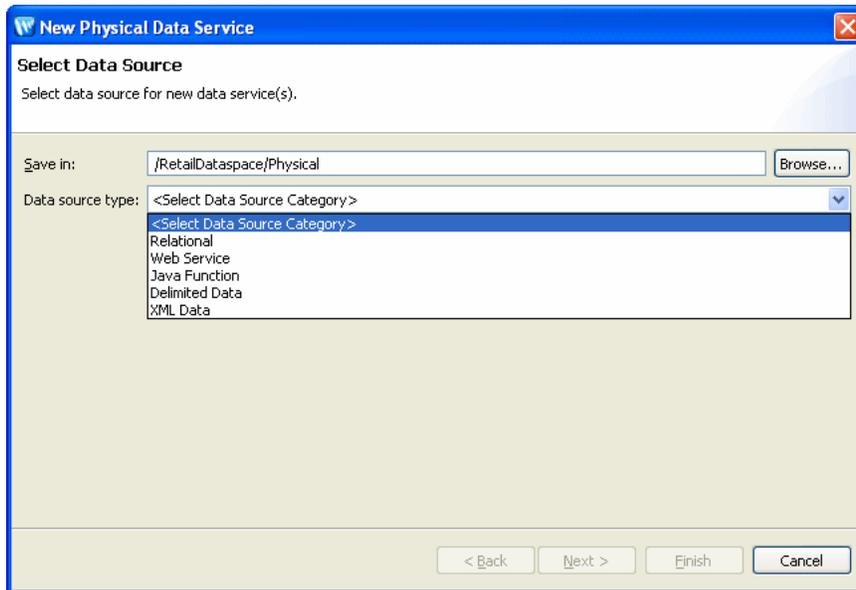
The following topics cover the actions necessary to create physical data services from SQL statements:

- [Section 3.2.3.1, "Setting Up the Physical Data Service Creation Wizard"](#)
- [Section 3.2.3.2, "Setting Up the Import Wizard for Relational Objects"](#)
- [Section 3.2.3.3, "Entering a SQL Statement"](#)
- [Section 3.2.3.4, "Setting Properties for New Library Functions"](#)
- [Section 3.2.3.5, "Verifying Data Service Composition"](#)

3.2.3.1 Setting Up the Physical Data Service Creation Wizard

Physical data services are created using a wizard.

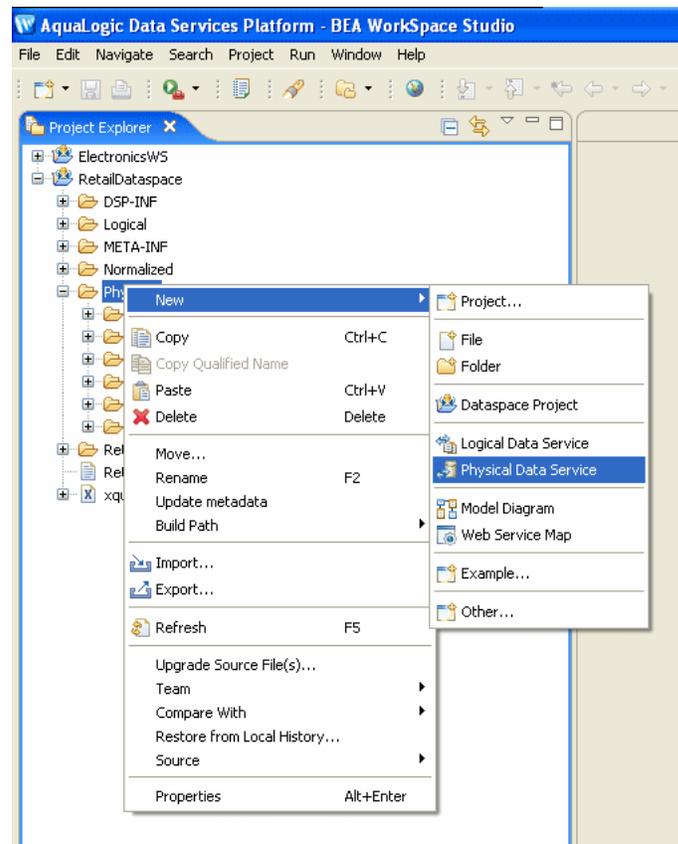
Figure 3–13 Physical Data Service Creation Wizard



Browse to specify a directory where you want to save the data service. The possible data source types are Relational, Web Service, Java Function, Delimited Data, and XML data.

3.2.3.1.1 Starting the Wizard To start the physical data service creation wizard:

1. Right-click on your dataspace project or any folder in your project.
2. Choose **New > Physical Data Service**.

Figure 3–14 Creating a New Physical Data Service

In the Project Explorer, a project is selected and the context window is displayed. New is highlighted and Physical Data Service is selected.

3.2.3.2 Setting Up the Import Wizard for Relational Objects

When importing a relational object available options include the ability to:

1. Set a location for your new data service to be saved within your project.
2. Select a data source from the dropdown listbox.
3. Select the database type for the selected source (PointBase for the sample RDBMS) from the dropdown listbox.
4. Select among the relational source types listed in the following table.

Table 3–7 Types of available relational data sources

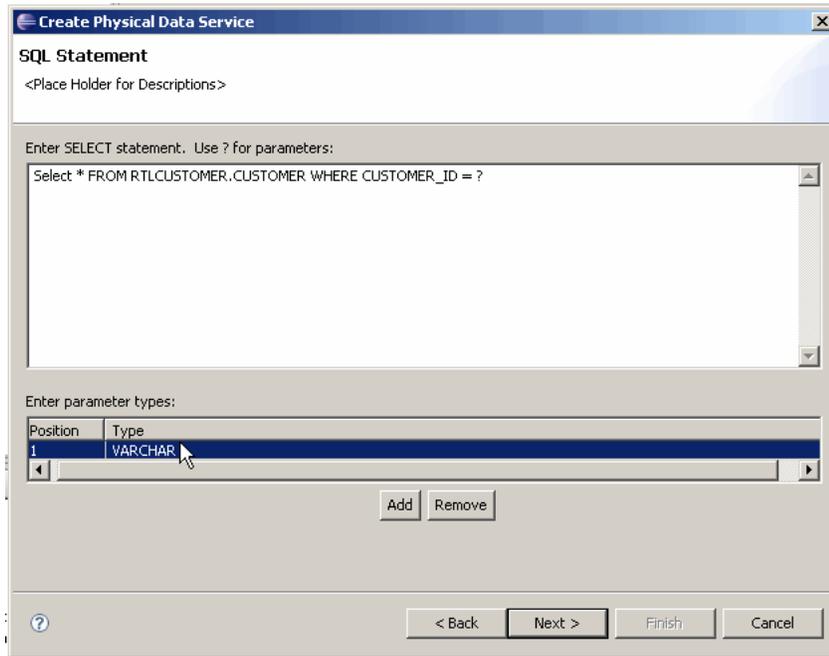
Relational Type	Description
Tables and Views	Displays all public tables and views in the selected data source.
Stored Procedures	Displays all public stored procedures in the selected data source.
SQL Statement	Allows creation of a SQL statement for extracting relational data from the data source.
Database Function	Allows creation of an XQuery function in a library data service based on build-in or custom database functions.

3.2.3.3 Entering a SQL Statement

You can build library data service functions based on SQL statements. The XQuery engine uses the statement to retrieve metadata which is, in turn, formulated into a function that can be used by other data services or made public.

After selecting the SQL Statement option the next page of the wizard allows you to enter a `SELECT` statement and any necessary parameters.

Figure 3–15 SQL Statement Entry Dialog



On the SQL Statement dialog, there is a space for entering `SELECT` statements. You can use a question mark for parameters. There is also a space for entering parameter types. A table shows a position and type.

You can type or paste your `SELECT` statement into the `SELECT` statement box, indicating parameters with a question-mark symbol.

?

Using one of the Oracle Data Service Integrator samples, the following `SELECT` statement can be used:

```
SELECT * FROM RTLCUSTOMER.CUSTOMER WHERE CUSTOMER_ID = ?
```

For the parameter field, you would need to select a data type. In this case, `CHAR` or `VARCHAR`.

1. Click **Add** to insert a new row into the parameter table, which indicates a parameter for the SQL statement.
2. Select **Parameter Type** from the drop-down combo box.

Notes:

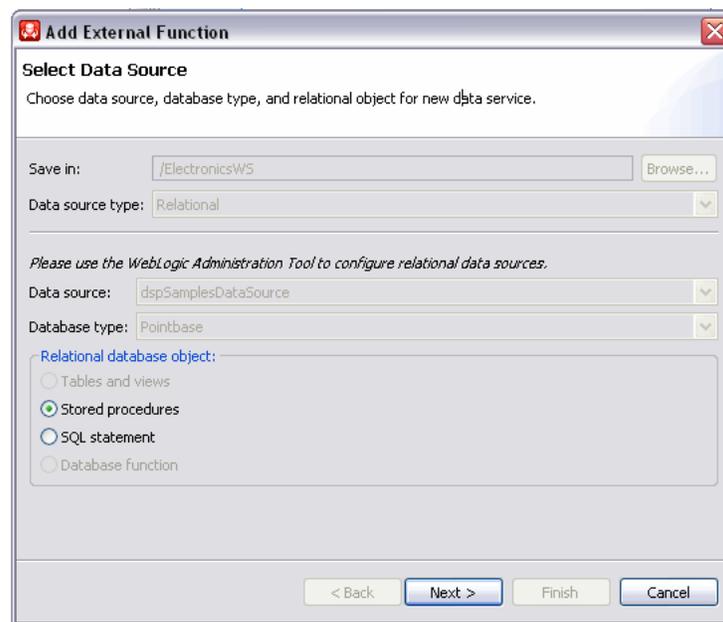
- When you run your query under Test View, you will need to supply the parameter in order for the query to run successfully.

- Oracle Data Service Integrator needs to be able to refer to the columns of the result of your SQL statement by name. To ensure that this is possible, you should use aliases as needed to ensure that computed columns indeed have usable names.
 - The position of the parameter is significant.
3. In Test view run your query, supplying a parameter such as CUSTOMER3.

3.2.3.3.1 Adding Operations to an Existing Data Service You can add SQL statement or stored procedure operations based on the same data source to an existing physical data service based a SQL statement.

For more information, see ["How To Add an External Function to an Existing Physical Data Service"](#) on page 3-91.

Figure 3–16 Adding a Stored Procedure or SQL Statement to a Data Service



In the Select Data Source dialog, you can choose the data source, datasource type, and relational object for a new data service. In the Relational database object section of the dialog, the Stored Procedures option is selected. All other options are disabled except for the SQL statement option.

3.2.3.4 Setting Properties for New Library Functions

This general topic applies to setting properties for all types of library data service functions.

Use the Review New Data Service Operations page to:

- Change the function name.
- Set the **Public** option (check if you want your function to be available to client applications).
- Set the kind of function (in some cases only one option will be available).

- Set the **Primary** option (check if you want your function to be the primary of its type).

Note: In some cases this option may not be available.

- Select a common XML namespace for the entire data service.
- Set the target namespace.

3.2.3.5 Verifying Data Service Composition

On the Review New Data Service(s) page you can set, confirm or, optionally, change suggested data service names depending on the type of physical data service you are creating.

3.2.3.5.1 Default Physical Data Service Names The nominated name for a new data service is, wherever possible, the same as the source object name. In some cases, however, names are adjusted to conform with XML naming conventions.

For more information, see "[XML Name Conversion Considerations](#)" on page 3-13.

3.2.3.5.2 About Automatic Data Service Name Changes Name conflicts occur when there is a data service of the same name present in the target directory. Name conflicts are highlighted in red.

There are several situations where you will need to change the name of your data service:

- There already is a data service of the same name in your application.
- You are trying to create multiple data services with the same name.

Data services always have the file extension:

.ds

3.2.4 How To Create Physical Data Services Based on Database Functions

You can create library physical data services based on two types of database functions:

- Functions that are provided with your database.
- *Custom functions* that you have created and stored in your database.

Note: A library data service created based on database functions is restricted to that type of function. For example, a library function based on a stored procedure cannot be added to a library data service that contains database functions.

A library data service created based on database functions is restricted to that type of function. For example, a library function based on a stored procedure cannot be added to a library data service that contains database functions.

- [Section 3.2.4.1, "Setting Up the Physical Data Service Creation Wizard"](#)
- [Section 3.2.4.2, "Setting Up the Import Wizard for Relational Objects"](#)
- [Section 3.2.4.3, "Providing Database Function Details"](#)
- [Section 3.2.4.4, "Verifying Data Service Composition"](#)

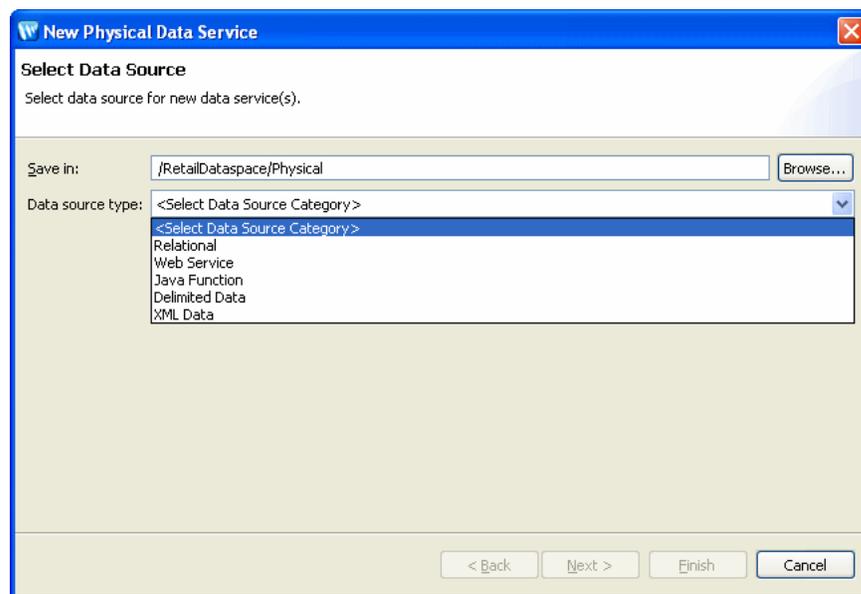
You can use the physical data service creation wizard to:

- Select relational as the Data Source type.
- Select a data source from available relational sources.
- Choose a database type. Database types listed would be drawn from the list of available database providers for your data source. By default Generic, the base platform provider, and Pointbase are provided.
- Select the Database function option.

3.2.4.1 Setting Up the Physical Data Service Creation Wizard

Physical data services are created using a wizard.

Figure 3–17 Physical Data Service Creation Wizard

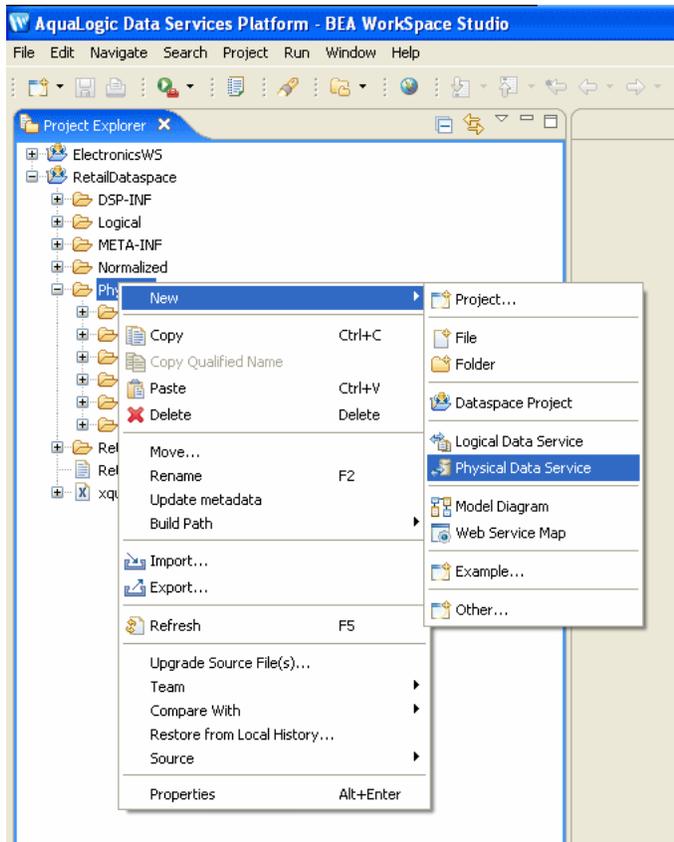


Browse to specify a directory where you want to save the data service. The possible data source types are Relational, Web Service, Java Function, Delimited Data, and XML data.

To start the physical data service creation wizard:

1. Right-click on your dataspaces project or any folder in your project.
2. Choose **New > Physical Data Service**

Figure 3–18 Creating a New Physical Data Service



In the Project Explorer, a project is selected and the context window is displayed. New is highlighted and Physical Data Service is selected.

3.2.4.2 Setting Up the Import Wizard for Relational Objects

When importing a relational object available options include the ability to:

1. Set a location for your new data service to be saved within your project.
2. Select a data source from the dropdown listbox.
3. Select the database type for the selected source (PointBase for the sample RDBMS) from the dropdown listbox.
4. Select among the relational source types listed in the following table.

Table 3–8 Types of available relational data sources

Relational Type	Relational Type
Tables and Views	Displays all public tables and views in the selected data source.
Stored Procedures	Displays all public stored procedures in the selected data source.
SQL Statement	Allows creation of a SQL statement for extracting relational data from the data source.
Database Function	Allows creation of an XQuery function in a library data service based on build-in or custom database functions.

1. In the Select a Data Source dialog choose **Database function**.
2. Click Next.

Figure 3–19 Importing Database Function Metadata

The screenshot shows a dialog box titled "Create Physical Data Service" with a sub-header "Select Data Source". The instructions inside the dialog are: "Choose data source, database type, and relational object for new data service." The form contains the following fields and options:

- Container:** A text box containing "/Imports/physicalDataServices" and a "Browse..." button.
- Data source type:** A dropdown menu set to "Relational".
- Instructions:** "Please use the WebLogic Administration Tool to configure relational data sources."
- Data source:** A dropdown menu set to "dspSamplesDataSource".
- Database type:** A dropdown menu.
- Relational database object:** A group box containing five radio buttons:
 - Tables and views
 - Stored procedures
 - SQL statement
 - Database function** (selected)

At the bottom of the dialog are four buttons: "< Back", "Next >", "Finish", and "Cancel".

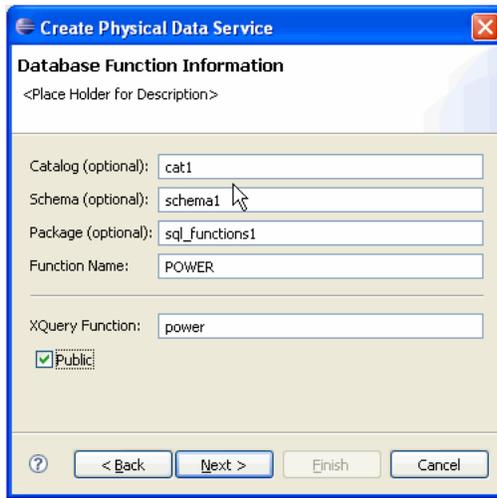
The Select Data Source dialog lets you choose the data source, database type, and relational object for the new data service. Browse to specify a container. Choose a data source type from the drop-down menu. In the next section of the form, there is a note: "Please use the WebLogic Administration Tool to configure relational data sources." Select a data source from the drop-down menu. Select a database type from the drop-down menu. Choose one out of five relational database objects: Tables and views, Stored procedures, SQL statement, or Database function.

3.2.4.3 Providing Database Function Details

Here are instructions for how to provide database function information.

1. Select a data source from the dropdown list of data sources available to your server. You should identify a data source that contains the built-in or user-defined database functions you want to access through your data services.
2. Enter the information necessary to identify your database function.
3. Complete the function definition including identifying parameters in Source view.

Figure 3–20 Entering Database Function Information



On the Database Function Information dialog, there are five fields: Catalog (optional), Schema (optional), Package (optional), Function Name, and XQuery Function. There is a check mark in the Public check box.

Table 3–9 Database Function Information Dialog Options

Option	Action	Comment/Reference
Catalog:	Enter catalog name, if needed by your RDBMS	
Schema:	Enter schema name, if needed by your RDBMS	
Package:	Enter package name, if needed by your RDBMS	
Function name:	Database function name	Required.
XQuery function	XQuery function name	Required; will invoke the database function.
Public	Select, if you want to make your operation public	Default for created XQuery functions is protected.
	Click Next	
Review	Enter library data service name	If the name of an existing library data service is provided.

3.2.4.4 Verifying Data Service Composition

On the Review New Data Service(s) page you can set, confirm or, optionally, change suggested data service names depending on the type of physical data service you are creating.

3.2.4.4.1 Default Physical Data Service Names

For more information, see "[XML Name Conversion Considerations](#)" on page 3-13.

3.2.4.4.2 About Automatic Data Service Name Changes Name conflicts occur when there is a data service of the same name present in the target directory. Name conflicts are highlighted in red.

There are several situations where you will need to change the name of your data service:

- There already is a data service of the same name in your application.
- You are trying to create multiple data services with the same name.

Data services always have the file extension:

.ds

3.2.5 How To Create a Physical Data Service from a Web Service

In Oracle Data Service Integrator, three top-level provider types are identified:

- WSDL File
- URI
- ALSB Proxy

A Web service is a self-contained, platform-independent unit of business logic that is accessible through application adaptors, as well as standards-based Internet protocols such as HTTP or SOAP.

Web services greatly facilitate application-to-application communication. As such they are increasingly central to enterprise data resources. A familiar example of an externalized Web service is a frequent-update weather portlet or stock quotes portlet that can easily be integrated into a Web application.

Similarly, a Web service can be effectively used to track a drop shipment order from a seller to a manufacturer.

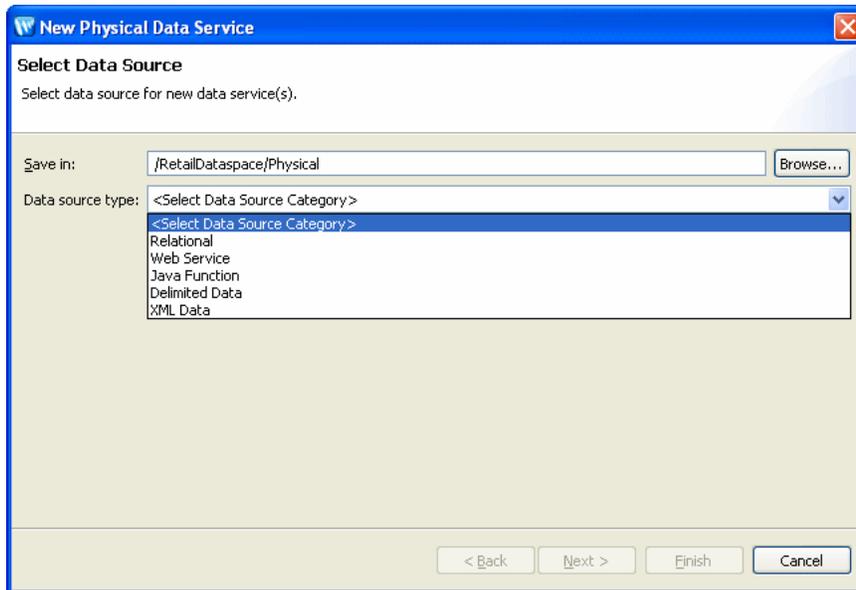
This section describes the following topics:

- [Section 3.2.5.1, "Setting Up the Physical Data Service Creation Wizard"](#)
- [Section 3.2.5.2, "Accessing a Web Service"](#)
- [Section 3.2.5.3, "Selecting Web Service Operations to Import"](#)
- [Section 3.2.5.4, "Setting Characteristics of Imported Web Service Operations"](#)
- [Section 3.2.5.5, "Setting the Data Service Name"](#)

3.2.5.1 Setting Up the Physical Data Service Creation Wizard

Physical data services are created using a wizard.

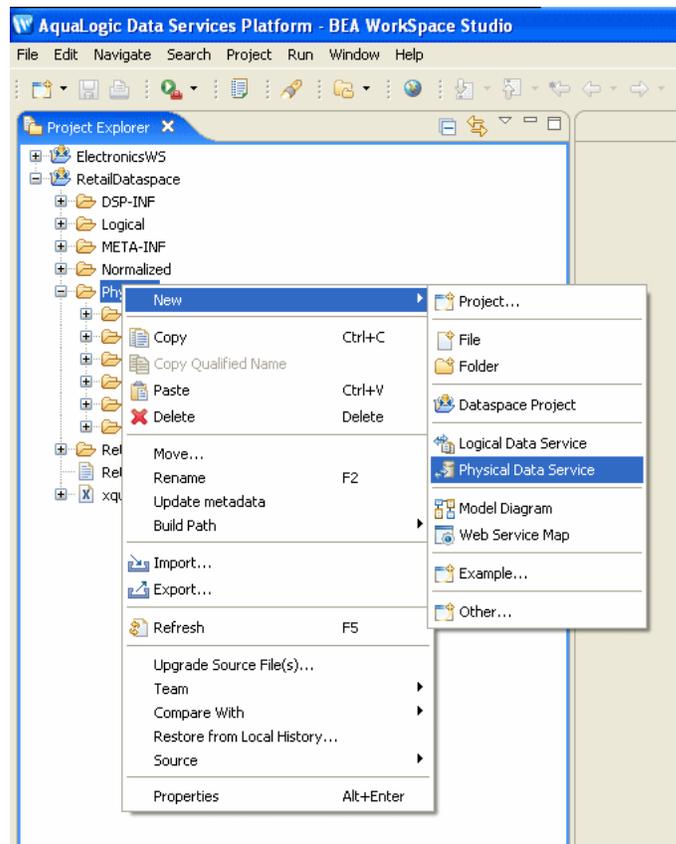
Figure 3–21 Physical Data Service Creation Wizard



Browse to specify a directory where you want to save the data service. The possible data source types are Relational, Web Service, Java Function, Delimited Data, and XML data.

3.2.5.1.1 Starting the Wizard To start the physical data service creation wizard:

1. Right-click on your dataspace project or any folder in your project.
2. Choose **New > Physical Data Service**

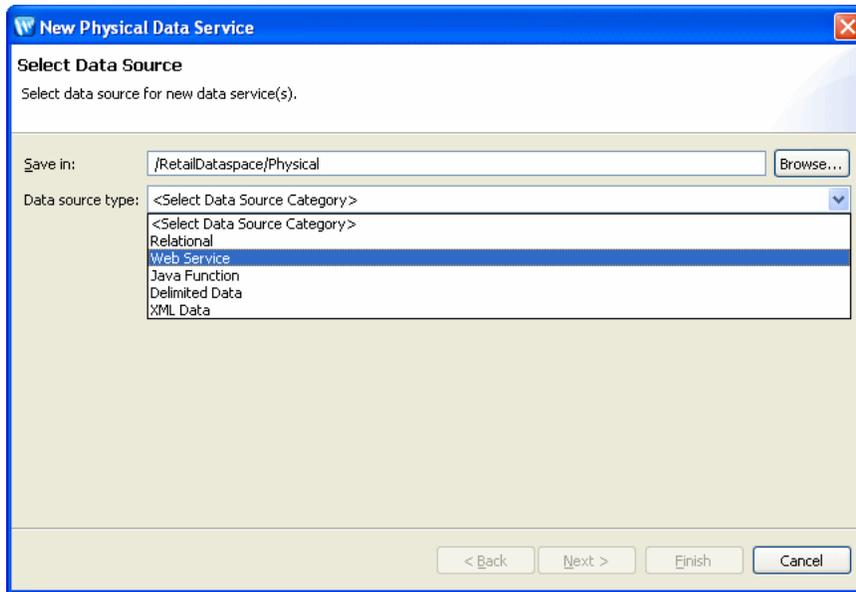
Figure 3–22 Creating a New Physical Data Service

In the Project Explorer, a project is selected and the context window is displayed. New is highlighted and Physical Data Service is selected.

3.2.5.2 Accessing a Web Service

After you select web service as your data source, you are given the option of specifying a WSDL file, URI, or ALSB proxy service.

Figure 3–23 Selecting Web Service as a Data Source



In the Select Data Source dialog, a Save In field lets you enter a directory. Under Data Source Type, Web Service is selected.

There are several ways to access a specific web service:

- From a Web Service Description Language (WSDL) file located in your current OSDI Studio project or through the AquaLogic Service Consumption Framework (SCF).
- From a WSDL accessible via a URL.
- Through an ALSB proxy service.

3.2.5.2.1 Locating a WSDL File You can select a WSDL file in two ways:

- From your project using the Browse button or
- By downloading a WSDL to your project through the consumption framework

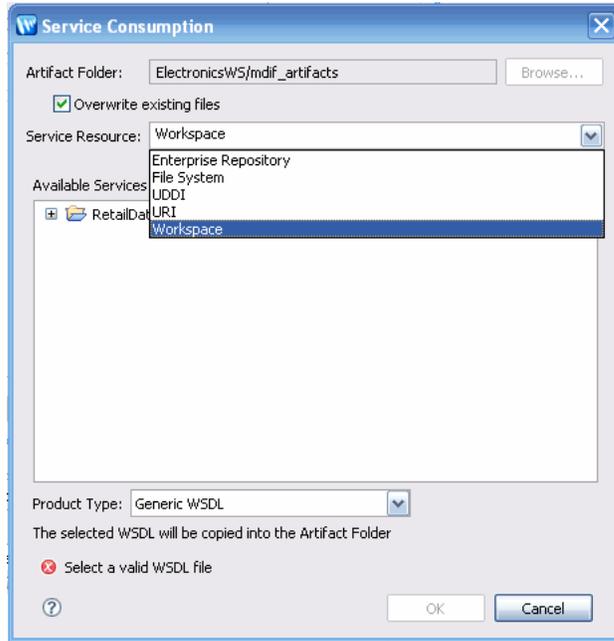
3.2.5.2.2 Browsing to a WSDL Click Browse to navigate to a WSDL in your current dataspace project.

3.2.5.2.3 Downloading a WSDL File via the Service Consumption Framework To download a WSDL file using the Service Consumption Framework:

1. Click on **Download WSDL...**
2. Select from available Service Resource Options.
3. Click **OK**. A SCF message indicating success or failure will appear.
4. Click **Next**.
5. Choose from available operations.

Note: Downloading a WSDL is similar to downloading a text file. Once it is in your project, you can treat it as local and make changes. Changes made to such a WSDL will not be reflected in the source WSDL. Similarly, any changes to the source WSDL (new operations or changes to the signature) would only be reflected if the WSDL was again downloaded and re-imported.

Figure 3–24 Using the Service Consumption Framework to Access a WSDL



The Service Consumption dialog lets you browse to select an artifact folder. There is a checkbox that lets you specify whether you want to overwrite existing files. In the Service Resource pull-down menu, there are five choices: Enterprise Repository, File System, UDDI, URI, and Workspace (which is selected). There is also a space that shows available services.

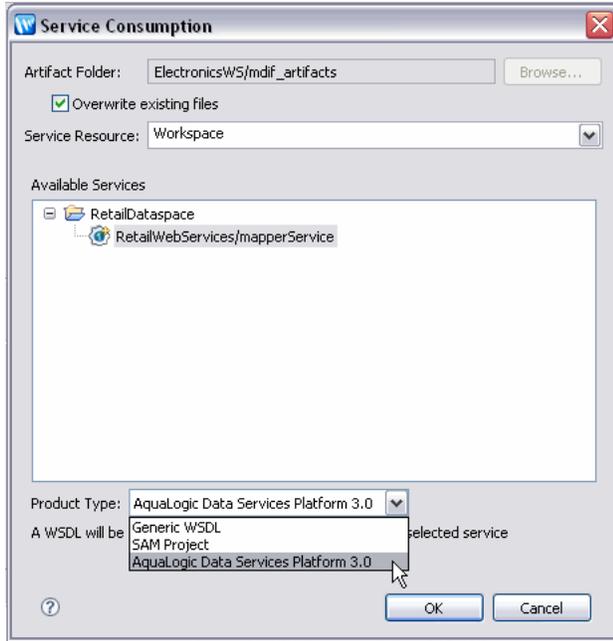
The following table briefly describe available service resources:

Resource	Description
Enterprise Repository	WSDLs in the associated artifacts folder of the Enterprise Repository
UDDI	UDDI resources registered with the UDDI registry.
URI	Any valid WSDL accessible via URI (alternatively, use the URI option from the primary dialog)
File System	A WSDL retrievable from the local file system
Workspace	WSDLs or services in other projects in the Workspace (alternatively, use the Browse button for WSDLs in the current project on the primary dialog).

3.2.5.2.4 Selecting the Product Type For a Workspace service resource, services can be consumed from several types of products. By default, three product types are available:

- Generic WSDL
- SAM (Service Assembly Modeler)
- AquaLogic Data Services Platform

Figure 3–25 Selecting a WSDL from an OSDI Studio Workspace



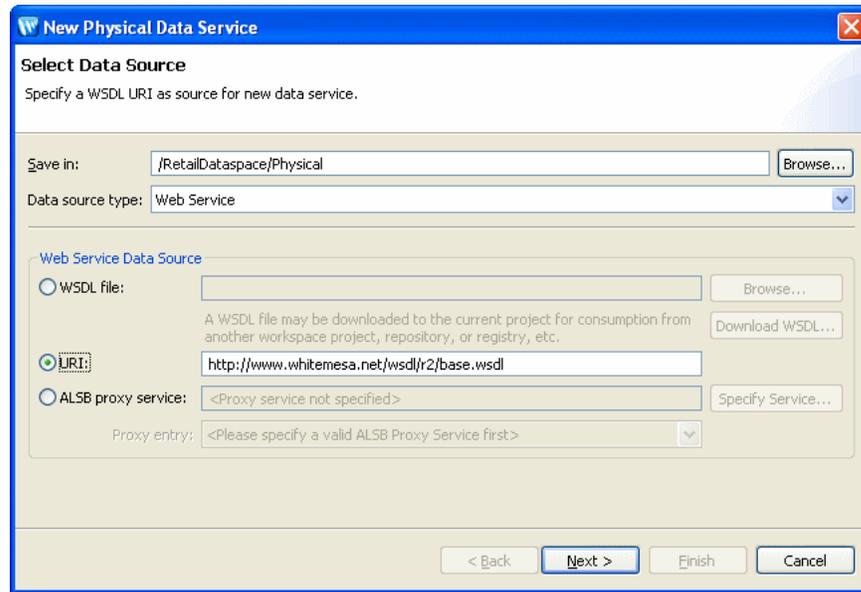
On the Service Consumption dialog, you can browse to select an artifact folder. There is a checkbox that lets you specify whether you want to overwrite existing files. For Service Resource, Workspace is selected. Under Available Services, in the RetailDataspace folder, RetailWebServices/mapperService is selected.

If you download a WSDL via SCF through either the Enterprise Repository or Workspace option and you are using the OSDI Studio or ALSB product type, you can locate and view the originated service of the WSDL using the Navigate to External Service right-click menu option on the WSDL file in the Project Explorer.

3.2.5.2.5 Specifying a WSDL URI Specifying a WSDL URI

You can test the ability to create a physical data service based on a web service using the following WSDL (available as of this writing):

`http://www.whitemesa.net/wsdl/r2/base.wsdl`

Figure 3–26 Importing Metadata from a WSDL

On the Select Data Source dialog, browse to find a directory where you want to save. From the Data source type pull-down, Web Service is selected. For Web Service Data Source, there are three choices: WSDL file, URI, and ALSB proxy serviced. URI is selected and a website is specified. There is a note: A WSDL file may be downloaded to the current project for consumption from another workspace project, repository, or registry, etc.

3.2.5.2.6 The ALSB Proxy Service Option To access web services through AquaLogic Service Bus (ALSB) you need to:

- Provide access and credential information to AquaLogic Service Bus.
- Select a proxy service (if there is more than one).

AquaLogic Service Bus access requires providing the following:

- Server name
- Port number
- User name
- Password

This information should be available from your AquaLogic Service Bus administrator.

Note: ■ You must configure the ALSB proxy service to use the sb transport protocol to enable access through OSDI Studio.

- The Select Proxy list only shows the WSDL-based AquaLogic Service Bus transport proxies in the ALSB server you are connected to.
-

After the required information is provided, the WSDL will become available using the name of the selected proxy service.

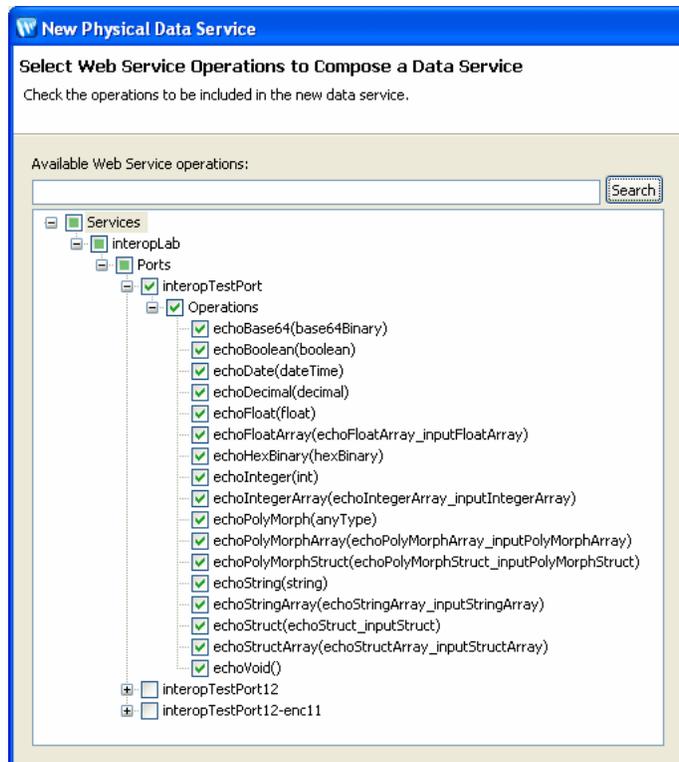
3.2.5.2.7 Steps in Importing a Web Service

1. Specify a Web service URL, local WSDL file, or ALSB proxy service.
2. Click Next.

3.2.5.3 Selecting Web Service Operations to Import

From the list of available webservice operations grouped by serviceName and portname, choose the operation that you want to turn into data service operation.

Figure 3–27 Selecting Web Service Operations



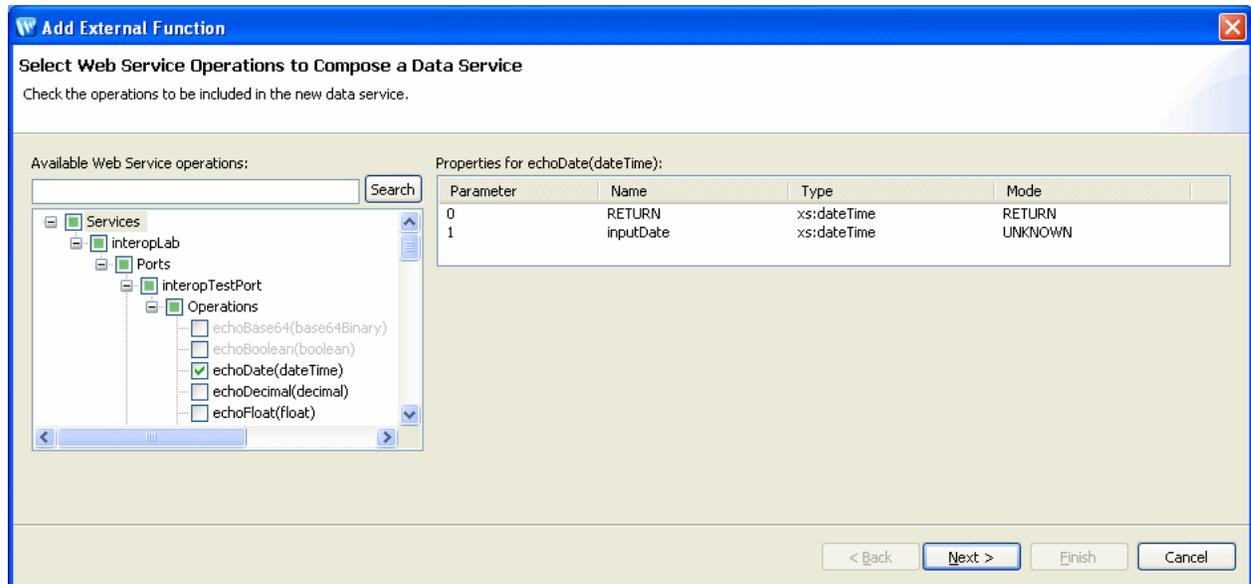
The Select Web Services Operations to Compose a Data Service dialog, a web services tree includes checkboxes next to operations you want to add. Check the operations to be included in the new data service. There is Search button that lets you navigate to available Web Services operations. On the dialog, in the web services tree, all operations are selected.

Note: During the import process you will be choosing the operations you want to import, setting names and other characteristics. These choices will determine whether a Library or Entity data service will be created. Thus a familiarity with the operations of your Web service is needed.

3.2.5.3.1 Adding Operations to an Existing Data Service You can add operations to an existing physical data service based a web service by adding an external function from the same WSDL.

For more information, see ["How To Add an External Function to an Existing Physical Data Service"](#) on page 3-91.

Figure 3–28 Adding an External Operation to a Data Service



The Select Web Services Operations to Compose a Data Service dialog, a web services tree includes checkboxes next to operations you want to add. Check the operations to be included in the existing data service. There is Search button that lets you navigate to available Web Services operations. On the dialog, in the web services tree, one operation is selected. A Properties table displays on the right side of the dialog with Parameter, Name, Type, and Mode columns.

3.2.5.3.2 Steps Involved in Selecting Web Service Operations

1. Select the operations you want to turn into data services or library data service functions.
2. Click Next.

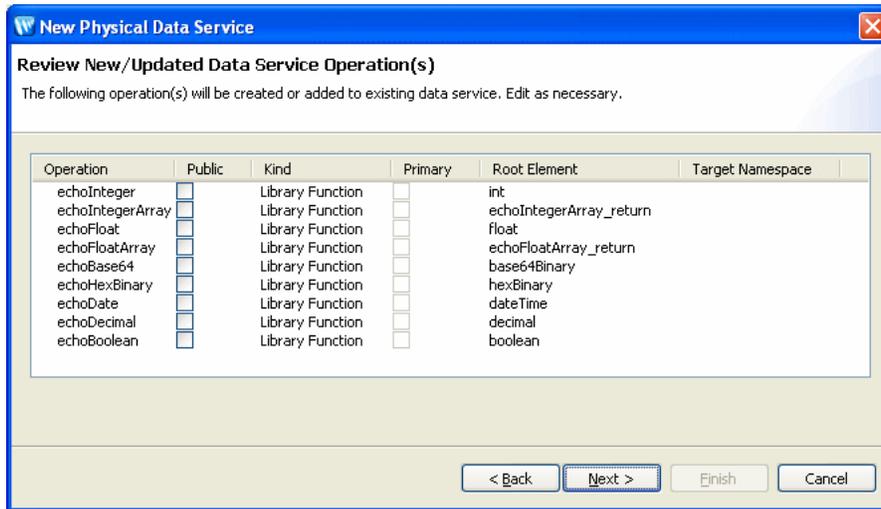
3.2.5.4 Setting Characteristics of Imported Web Service Operations

The following table describes available options for each operation you have selected to import.

Table 3–10 Options Available for Imported Web Service Operations

Characteristic	Options	Comment
Operation	adjust as needed	You can change the nominated name to any legal XML name using the built-in line editor.
Public	Boolean	By default Web service-derived operations are protected. A checkbox allows you to mark any function or procedure as public. (Once in a data service, operations can be marked private as needed.)
Kind	<ul style="list-style-type: none"> ■ Read ■ Create ■ Update ■ Delete ■ Library function ■ Library procedure 	<p>Operations determined to return void are automatically marked as library procedures.</p> <p>You can change the nominated function type. The wizard attempts to correctly set the function type being imported.</p> <p>Note: Operations marked as create, update, or delete functions will be packaged in an Entity data service. Otherwise, the resulting data service will be of type Library.</p>
Primary	Boolean	Not applicable for web service operations.
Root Element	Root element of the operation	For complex data types the topmost element is listed. In case of RPC-style web services the top-most generated element is listed.
Target Namespace	imported value	This represents the target namespace of the generated data service.

Figure 3–29 Setting Characteristics of Imported Web Service Operations



The Review New/Updated Data Service Operation(s) dialog shows the operation(s) that will be created or added to the existing data service. The dialog lets you edit information about the operations. There are six columns: Operation, Public, Kind, Primary, Root Element, and Target Namespace.

3.2.5.5 Setting the Data Service Name

You can change the name of your data service to any legal name that does not conflict with another name in the current data space.

In addition, if there already is a data service in your project based on the same WSDL an option to add the new operation to the existing data service appears.

Note: When importing a web service operation that itself has one or more dependent (or referenced) schemas, the wizard creates second-level schemas according to internal naming conventions. If several operations reference the same secondary schemas, the generated name for the secondary schema may change if you re-import or synchronize with the Web service.

3.2.5.5.1 Implementation Notes This section contains implementation notes.

3.2.5.5.2 Special Considerations when Creating a Data Service Based on a RPC-Style Web Service In case of RPC-style web services, results are return as qualified or unqualified based on the setting of the schema attribute:

`elementFormDefault`

In the general case of web services, `elementFormDefault` can be overridden by setting the form attribute for any child element. However, such individual settings are ignored for RPC-style web services since only the global setting (qualified or unqualified) is taken into account.

For example:

```
<s:schema elementFormDefault="qualified"
  targetNamespace="http://temp.openuri.org/SampleApp/CustomerOrder.xsd"
  xmlns:s0="http://temp.openuri.org/SampleApp/CustomerOrder.xsd"
  xmlns:s="http://www.w3.org/2001/XMLSchema">
  <s:complexType name="ORDER">
    <s:sequence>
      <s:element minOccurs="0" maxOccurs="1" form="unqualified" name="ORDER_ID" type="s:string"/>
      <s:element minOccurs="0" maxOccurs="1" form="unqualified" name="CUSTOMER_ID"
        type="s:string"/>
    </s:sequence>
  </s:complexType>
</s:schema>
```

In the above code the global element is qualified but a child element (ORDER_ID) is unqualified.

In the standard case, the special setting of "unqualified" for ORDER_ID will be honored. In the case of RPC-style web services, however, the runtime will generate "qualified" attributes for all the elements, including ORDER_ID.

Note: RPC-style web services such as those generated by ADO.NET may contain child elements with "form" attributes which do not match the schema's `elementFormDefault` declaration. In order for such web services to be turned into executable data service operations, make sure that all form element attributes and the `elementFormDefault` attribute are in agreement (either "qualified" or "unqualified").

3.2.5.5.3 Multi-dimensional Arrays in RPC Mode Multi-dimensional arrays in RPC mode are not supported.

3.2.5.5.4 See Also

For more information, see [How to Create SOAP Handlers for Imported WSDLs](#).

3.2.6 Preparing to Create Physical Data Services From Java Functions

This topic provides an overview of how to create a new physical data service from Java functions.

Before you can create physical data services based on custom Java functions, you need to create a Java class containing both the schema and function information. The entire process involves the following:

1. Using Apache XMLBeans, Oracle XMLBeans, or SDO DataObjects, create a schema of the data that is being used as parameters and return values by the Java functions.
2. Create the XMLBean classes or SDO DataObject classes and package them in a JAR file.

Tip: For more information, see [Creating XMLBean Support for Java Functions](#).

3. Place the JAR file in the `DSP-INF/lib` folder of the project in which you want to create the new Physical Data Service.
4. Create the new Physical Data Service based on your custom Java functions by importing the corresponding `.class` file.

3.2.7 How To Create a Physical Data Service from a Java Function

The following sections describe how to create physical data services based on custom Java functions that return both simple and complex types.

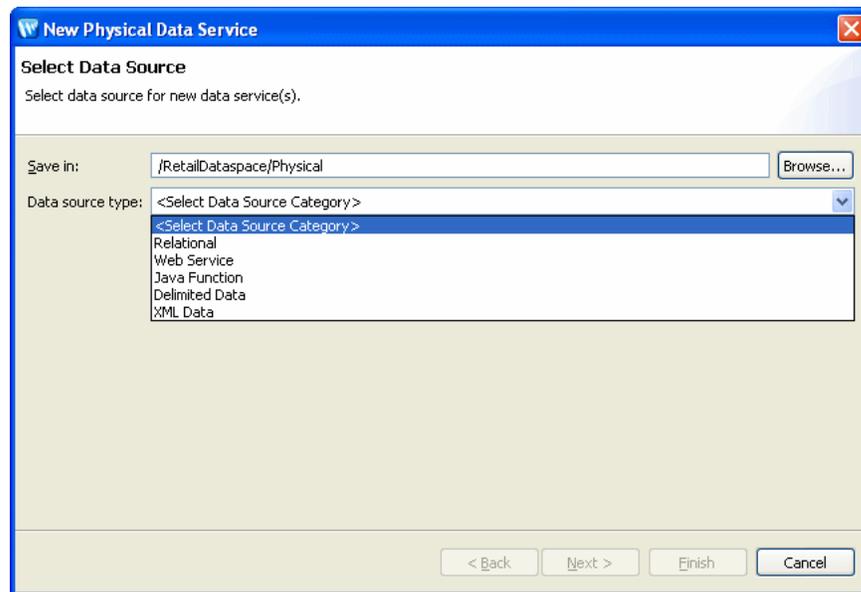
Note: Before you can create physical data services based on custom Java functions, you must create a Java class containing both the schema and function information. For more information, see ["Preparing to Create Physical Data Services From Java Functions"](#) on page 3-48.

For more information about supported Java types and the corresponding generated data services, see ["Physical Data Services from Java Functions Overview"](#) on page 3-2.

- [Section 3.2.7.1, "Setting Up the Physical Data Service Creation Wizard"](#)
- [Section 3.2.7.2, "Accessing Java Functions"](#)
- [Section 3.2.7.1, "Setting Up the Physical Data Service Creation Wizard"](#)
- [Section 3.2.7.4, "Setting Characteristics of Imported Java Functions"](#)
- [Section 3.2.7.5, "Setting the Physical Data Service Name"](#)
- [Section 3.2.7.6, "See Also"](#)

3.2.7.1 Setting Up the Physical Data Service Creation Wizard

Physical data services are created using a wizard.

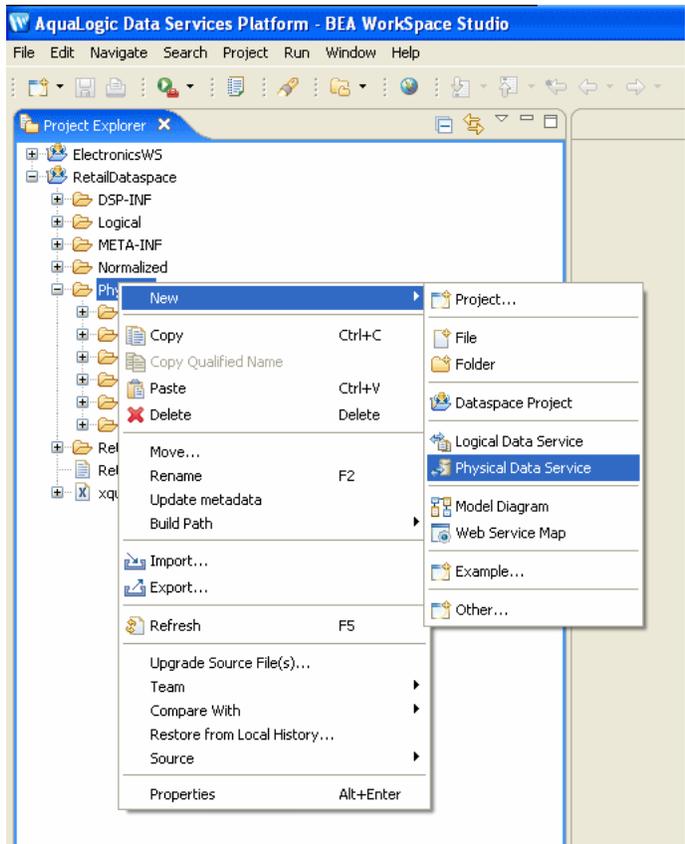
Figure 3–30 Physical Data Service Creation Wizard

Browse to specify a directory where you want to save the data service. The possible data source types are Relational, Web Service, Java Function, Delimited Data, and XML data.

3.2.7.1.1 Starting the Wizard To start the physical data service creation wizard:

1. Right-click on your dataspace project or any folder in your project.
2. Choose **New > Physical Data Service**.

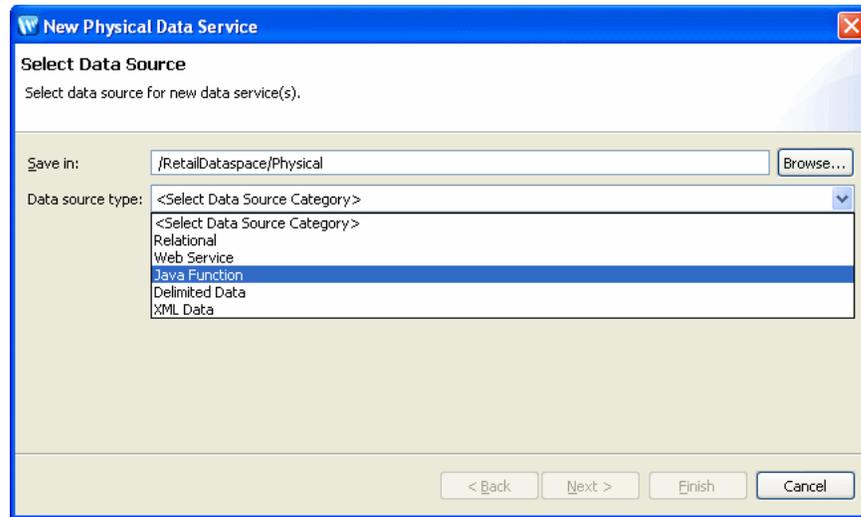
Figure 3–31 Creating a New Physical Data Service



In the Project Explorer, a project is selected and the context window is displayed. New is highlighted and Physical Data Service is selected.

3.2.7.2 Accessing Java Functions

After you choose Java Function as your data source, you need to specify a class name containing the Java functions.

Figure 3–32 Choosing Java Function as a Data Source

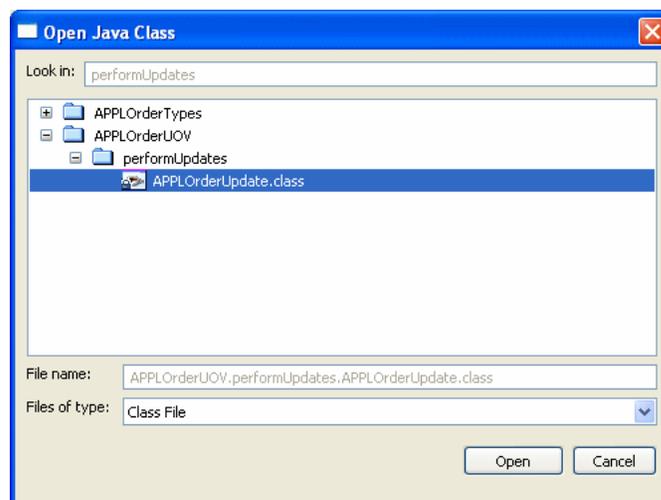
On the Select Data Source dialog, Java Function is selected as the data source type.

3.2.7.2.1 Choosing the Java class: To choose the Java class containing the Java functions:

1. Choose Java Function from the Data source type drop-down list.
2. Click **Browse**. The Open Java Class dialog appears.
3. Select the Java `.class` file and click **Open**.

The `.class` file must reside in the same dataspaces into which you are importing the Java functions.

4. Click **Next**.

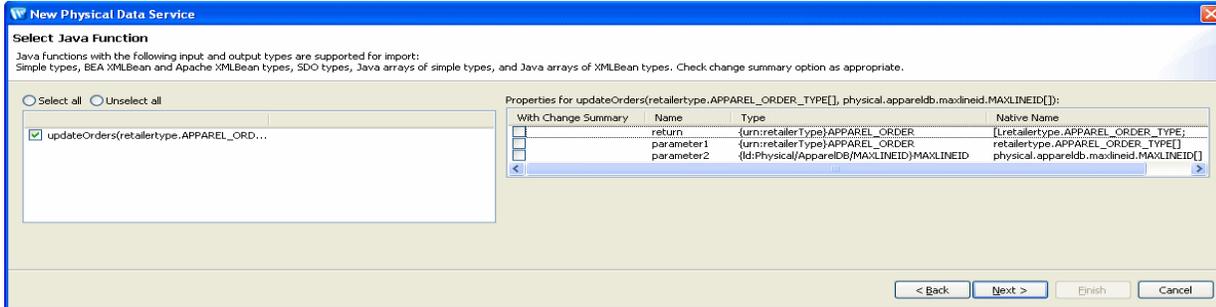
Figure 3–33 Open Java Class Dialog

On the Open Java Class dialog, the `.class` file is selected.

3.2.7.3 Selecting Java Functions to Import

After you select Java Function as your data source, you need to select the Java functions to import.

Figure 3–34 Selecting Java Function Dialog



On the Select Java Function dialog, there is one panel containing function names on the left. The function name checkbox is checked. The Select all and Unselect all options are not selected. On the right is a reflecting properties for the selected function. There are four columns: With Change Summary, Name, Type, and Native Name. There are checkboxes in the With Change Summary column next to each property.

To select the Java functions to import:

1. Select the Java functions you want to import by checking the corresponding box.
 Select **With Change Summary** to have Oracle Data Service Integrator declare the parameter or return value as changed-element enabling you to use it with update operations. This option is only available for SDO DataObject-generated classes.
2. Click Next.

3.2.7.4 Setting Characteristics of Imported Java Functions

After choosing the Java functions to import, you can optionally set the characteristics of the functions.

Figure 3–35 Setting Characteristics of Imported Java Functions



On the Review New/Updated Data Service Operation(s) dialog, there is a table with six columns: Operation, Public, Kind, Primary, Root Element, Target Namespace. Characteristic for the updateOrders operation are provided in the table.

The following table describes the available options for each function you have selected to import.

Table 3–11 Options Available for Imported Java Functions

Characteristic	Options	Comment
Operation name	Adjust as needed	You can change the nominated name to any legal XML name using the built-in line editor.
Public	Boolean	By default Java function-derived operations are protected. A checkbox allows you to mark any function or procedure as public. (Once in a data service, operations can be marked private as needed.)
Kind	<ul style="list-style-type: none"> ■ Read ■ Create ■ Update ■ Delete ■ Library function ■ Library procedure 	<p>Functions determined to return void are automatically marked as library procedures.</p> <p>You can change the nominated function type. The wizard attempts to correctly set the function type being imported.</p> <p>Note: Operations marked as create, update, or delete functions will be packaged in an Entity data service. Otherwise, the resulting data service will be of type Library.</p>
is Primary	Boolean	Not applicable for Java functions.
Root Element	Root element of the operation	For complex data types the topmost element is listed.
Target Namespace	Imported value	This represents the target namespace of the generated data service.

To set the characteristics of imported Java functions:

1. Optionally edit the details of each operation:
2. Click **Next**.

3.2.7.5 Setting the Physical Data Service Name

You can set the name of your data service to any legal name that does not conflict with another name in the current dataspace.

To complete the wizard:

1. Type the name of the data service in the Data service name field.
2. Click **Finish**.

Oracle Data Service Integrator creates a pragma (visible in Source view) that defines the function signature and relevant schema type for complex types such as schema elements or SDO types.

If there are existing data services in your project, you have the option of adding functions and procedures to that library or creating a new library for them. All the Java file functions are located in the same data service.

Note: When importing a Java function that itself has one or more dependent (or referenced) schemas, the wizard creates second-level schemas according to internal naming conventions. If several operations reference the same secondary schemas, the generated name for the secondary schema may change if you re-import or synchronize with the Java class.

3.2.7.6 See Also

Concepts

- [Section 3.1.2, "Physical Data Services from Java Functions Overview"](#)

How Tos

- [Section 3.2.6, "Preparing to Create Physical Data Services From Java Functions"](#)
- [Section 3.3.4.2, "Creating XMLBean Classes for Java Functions"](#)

3.2.8 How To Create a Physical Data Service from XML Data

XML files are a convenient means of handling hierarchical data. XML files and associated schemas are easily turned into library data service functions.

The following topics cover the actions necessary to create physical data services from XML data:

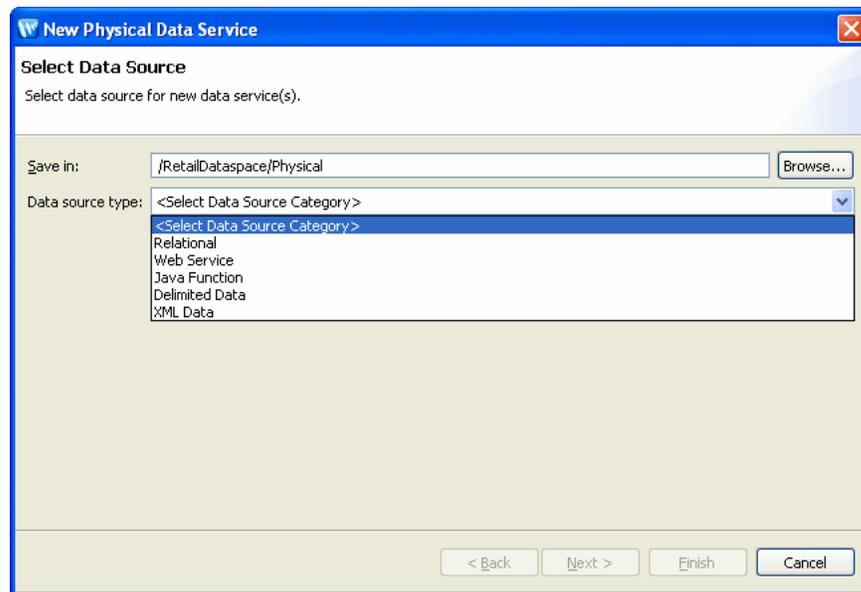
- [Section 3.2.8.1, "Setting Up the Physical Data Service Creation Wizard"](#)
- [Section 3.2.8.2, "Specifying XML Data Schema and File"](#)
- [Section 3.2.8.3, "Setting Properties for New Library Functions"](#)
- [Section 3.2.8.4, "Verifying Data Service Composition"](#)
- [Section 3.2.8.5, "XML File Import Sample"](#)

You can use the physical data service creation wizard to:

- Select XML Data as the Data Source type.
- Select a schema file and option data file.
- Create a Library data service based on the XML data.

3.2.8.1 Setting Up the Physical Data Service Creation Wizard

Physical data services are created using a wizard.

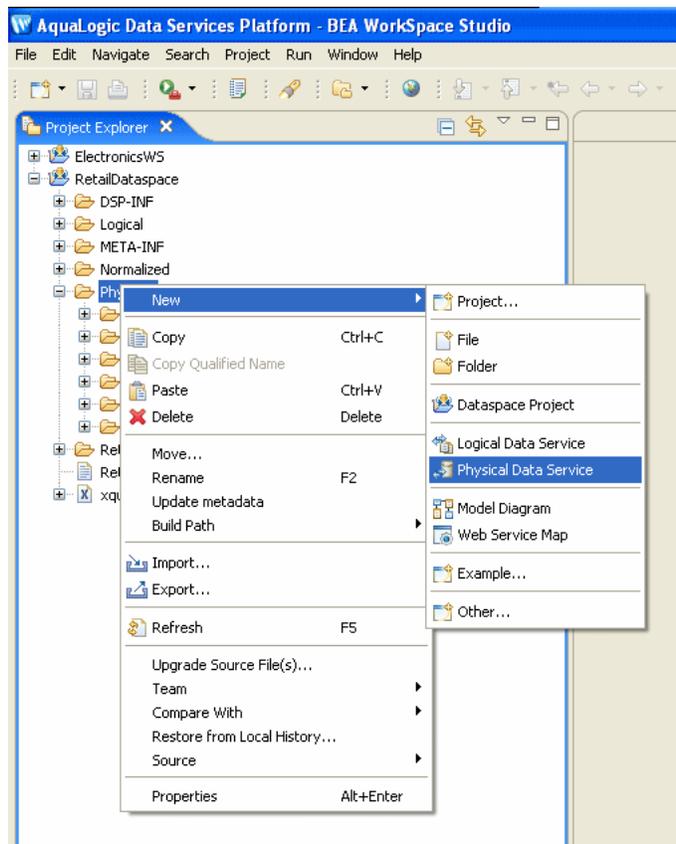
Figure 3–36 Physical Data Service Creation Wizard

Browse to specify a directory where you want to save the data service. The possible data source types are Relational, Web Service, Java Function, Delimited Data, and XML data. Select XML data.

3.2.8.1.1 Starting the Wizard To start the physical data service creation wizard:

1. Right-click on your dataspace project or any folder in your project.
2. Choose **New > Physical Data Service**

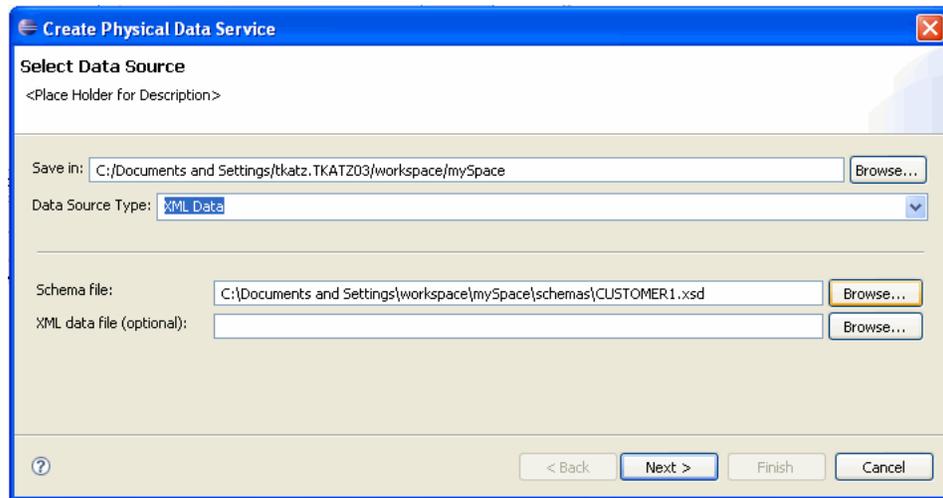
Figure 3–37 Creating a New Physical Data Service



In the Project Explorer, a project is selected and the context window is displayed. New is highlighted and Physical Data Service is selected.

3.2.8.2 Specifying XML Data Schema and File

A physical data service based on XML data requires identification of a valid XML schema and, optionally, a data source.

Figure 3–38 Import XML Data Wizard

On the Select Physical Data Service dialog, browse to specify the directory where you want to save the XML schema. From a pull-down menu, data source type XML Data. Browse to find the schema file. Browse to find the optional XML data file.

The scheme must be available in your dataspace.

The data source can be:

- File-based
- URI-based

In most cases the XML data will be available at runtime, through a URI.

However, in cases where the XML data is also in your project you can specify an absolute location for the file. You can also import data from any XML file on your system using an absolute path prepended with the following:

```
file:///
```

For example, on Windows systems you can access an XML file such as Orders.xml from the root C: directory using the following URI:

```
file:///c:/Orders.xml
```

On a UNIX system, you would access such a file with the following URI:

```
file:///home/Orders.xml
```

3.2.8.3 Setting Properties for New Library Functions

This general topic applies to setting properties for all types of library data service functions.

Use the Review New Data Service Operations page to:

- Change the function name.
- Set the **Public** option (check if you want your function to be available to client applications).
- Set the kind of function (in some cases only one option will be available).

- Set the **Primary** option (check if you want your function to be the primary of its type).

Note: In some cases this option may not be available.

- Select a common XML namespace for the entire data service.
- Select a common XML namespace for the entire data service.
- Set the target namespace.

The root element, which is read only, is also displayed.

3.2.8.4 Verifying Data Service Composition

On the Review New Data Service(s) page you can set, confirm or, optionally, change suggested data service names depending on the type of physical data service you are creating.

3.2.8.4.1 Default Physical Data Service Names The nominated name for a new data service is, wherever possible, the same as the source object name. In some cases, however, names are adjusted to conform with XML naming conventions.

When a source name is encountered that does not fit within XML naming conventions, default generated names are converted according to rules described by the SQLX standard.

Generally speaking, an invalid XML name character is replaced by its hexadecimal escape sequence (having the form `xUUUU`).

3.2.8.4.2 About Automatic Data Service Name Changes Name conflicts occur when there is a data service of the same name present in the target directory. Name conflicts are highlighted in red.

There are several situations where you will need to change the name of your data service:

- There already is a data service of the same name in your application.
- You are trying to create multiple data services with the same name.

Data services always have the file extension:

`.ds`

3.2.8.5 XML File Import Sample

An XML file import sample can be found in the sample RTLApp directory:

`DataServices/Demo`

3.2.8.5.1 Testing the Import Wizard with an XML Data Source When you create metadata for an XML data source but do not supply a data source name, you will need to identify the URI of your data source as a parameter when you execute the data service's read function.

The identification takes the form of:

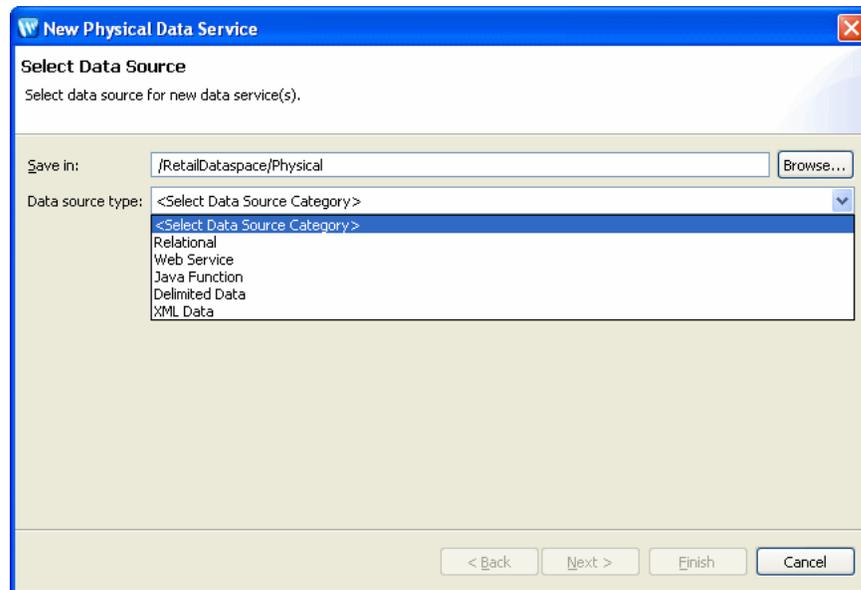
`<uri>/path/filename.xml`

where *uri* is representative of a path or path alias, *path* represents the directory and *filename.xml* represents the filename. The `.xml` extension is required.

3.2.9 How To Create a Physical Data Service from a Delimited File

Spreadsheets offer a highly adaptable means of storing and manipulating information, especially information which needs to be changed quickly. You can easily turn such spreadsheet data into a data services.

Figure 3–39 Physical Data Service Creation Wizard



Browse to specify a directory where you want to save the data service. The possible data source types are Relational, Web Service, Java Function, Delimited Data, and XML data. Select XML data.

You can use the physical data service creation wizard to:

- Select a delimited file as the Data Source type.
- Select either a schema file or a file with delimited data.
- Specify whether the information has a header or not.
- Specify delimiter.
- Specify a fixed width value for each column.

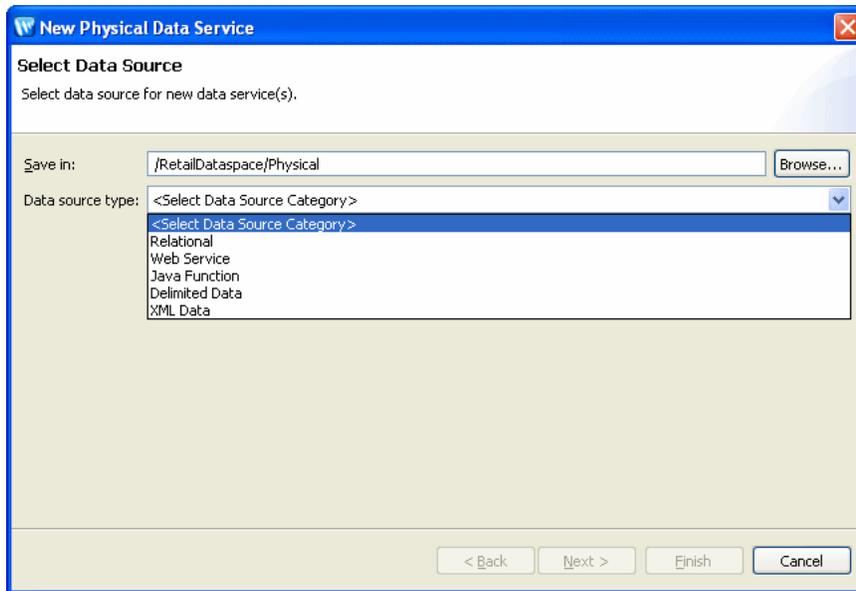
The following topics cover the actions necessary to create physical data services from delimited files:

- [Section 3.2.9.1, "Setting Up the Physical Data Service Creation Wizard"](#)
- [Section 3.2.9.2, "Specifying Delimited File Information"](#)
- [Section 3.2.9.3, "Setting Properties for New Library Functions"](#)
- [Section 3.2.9.4, "Verifying Data Service Composition"](#)

3.2.9.1 Setting Up the Physical Data Service Creation Wizard

Physical data services are created using a wizard.

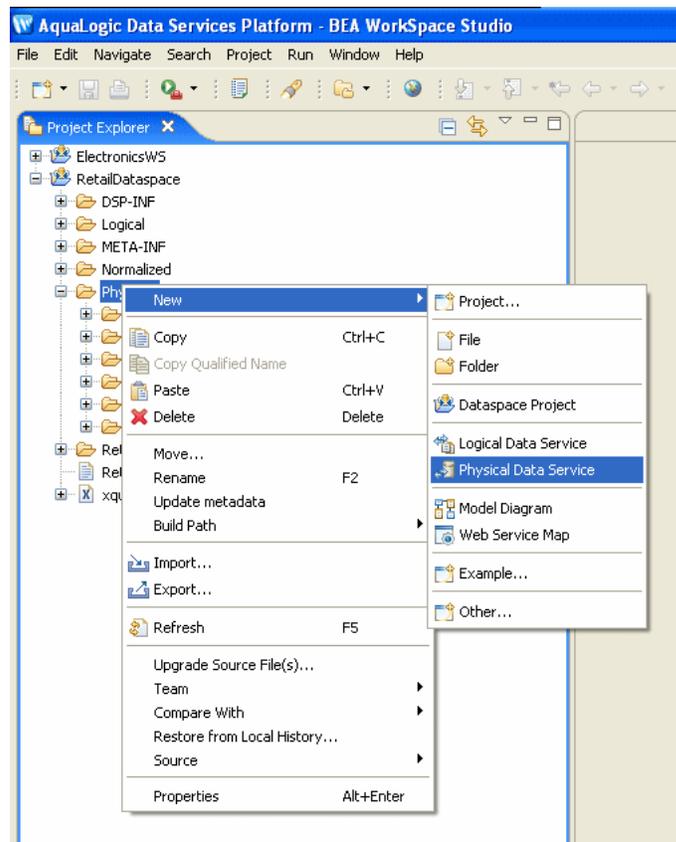
Figure 3–40 Physical Data Service Creation Wizard



Browse to specify a directory where you want to save the data service. The possible data source types are Relational, Web Service, Java Function, Delimited Data, and XML data.

3.2.9.1.1 Starting the Wizard To start the physical data service creation wizard:

1. Right-click on your dataspace project or any folder in your project.
2. Choose **New > Physical Data Service**.

Figure 3–41 Creating a New Physical Data Service

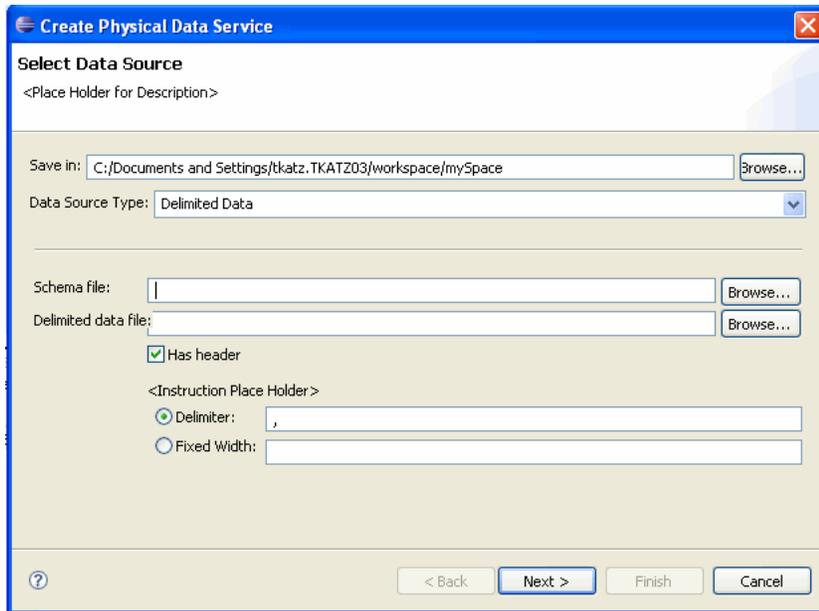
In the Project Explorer, a project is selected and the context window is displayed. New is highlighted and Physical Data Service is selected.

3.2.9.2 Specifying Delimited File Information

A Library data service based on delimited data requires:

1. Schema in your project and/or a
2. Location of the delimited data file

Figure 3–42 Import Delimited File Data Wizard



On the Import Delimited File Data Wizard, browse to a directory where you want to save. Delimited Data is shown as the Data Source Type. You can browse for or enter a Schema file name and Delimited data file name. A checkbox indicates Has header. In the Instruction Place Holder section, Delimiter is selected and a comma is specified. As an alternative, you can select Fixed Width.

The schema and data file must be available in your dataspace.

3.2.9.2.1 Providing a Document Name, a Schema Name, or Both There are several approaches to developing metadata around delimited information, depending on your needs and the nature of the source.

- Provide a delimited document name only. If you supply the import wizard with the name of a valid CSV file, the wizard will automatically create a schema based on the columns in the document. All the columns will be of type string, although you can later modify the generated schema with more accurate type information. The generated schema will have the same name as the source file.
- Providing a schema name only. This option is typically used when the source file is dynamic; for example, when data is streamed.
- Providing both a schema and a document name. Providing a schema with a CSV file gives you the ability to more accurately type information in the columns of a delimited document.

3.2.9.2.2 Locating the CSV File Using the import wizard you can browse to any file in your project. You can also import data from any CSV file on your system using an absolute path prepended with:

file:///

For example, on Windows systems you can access an XML file such as `Orders.xml` from the root `C:` directory using the following URI:

file:///<c:/home>/Orders.csv

On a UNIX system, you would access such a file with the URI:

```
file:///<home>/Orders.csv
```

3.2.9.2.3 Import Delimited Data Options

There are two options:

- **Header.** Indicates whether the delimited file contains header data. Header data is located in the first row of the spreadsheet. If you check this option, the first row will not be treated as imported data.
- **Delimited or Fixed Width.** Data in your file is either separated by a specific character (such as a comma) or is of a fixed width (such as 10 spaces). If the data is delimited, you also need to provide the delimited character. By default the character is a comma.

3.2.9.2.4 Supported Datatypes

The following datatypes are supported for delimited file metadata import operations:

```
XMLSchemaType.BASE64BINARY
XMLSchemaType.BOOLEAN
XMLSchemaType.DATE
XMLSchemaType.DATETIME
XMLSchemaType.DECIMAL
XMLSchemaType.DOUBLE
XMLSchemaType.FLOAT
XMLSchemaType.INT
XMLSchemaType.INTEGER
XMLSchemaType.LONG
XMLSchemaType.STRING
XMLSchemaType.SHORT
```

3.2.9.2.5 Additional Considerations

Consider the following:

- The number of delimiters in each row must match the number of header columns in your source minus one (# of columns-1). If subsequent rows contain more than the maximum number of delimiters (fields), subsequent use of the data service will not be successful.
- If the delimited file has rows with a variable number of delimiters (fields), you can supply a schema that contains optional elements for the trailing set of extra elements.
- Not all characters are handled the same way. Some characters may need special escape sequences before spreadsheet data can be accessed at runtime.

3.2.9.3 Setting Properties for New Library Functions

This general topic applies to setting properties for all types of library data service functions.

Use the Review New Data Service Operations page to:

- Change the function name.
- Set the **Public** option (check if you want your function to be available to client applications).
- Set the kind of function (in some cases only one option will be available).

Note: In some cases this option may not be available.

- Select a common XML namespace for the entire data service.
- Set the target namespace.

The root element, which is read only, is also displayed.

3.2.9.4 Verifying Data Service Composition

On the Review New Data Service(s) page you can set, confirm or, optionally, change suggested data service names depending on the type of physical data service you are creating.

3.2.9.4.1 Default Physical Data Service Names The nominated name for a new data service is, wherever possible, the same as the source object name. In some cases, however, names are adjusted to conform with XML naming conventions.

For more information, see XML Name Conversion Considerations.

3.2.9.4.2 About Automatic Data Service Name Changes Name conflicts occur when there is a data service of the same name present in the target directory. Name conflicts are highlighted in red.

There are several situations where you will need to change the name of your data service:

- There already is a data service of the same name in your application.
- You are trying to create multiple data services with the same name.

Data services always have the file extension:

.ds

3.3 How to

These sections describe procedures to create and update physical data services:

- [Section 3.3.1, "How To Enable Optimistic Locking"](#)
- [Section 3.3.2, "How To Update Physical Data Service Metadata"](#)
- [Section 3.3.3, "Creating SOAP Handlers for Imported WSDLs"](#)
- [Section 3.3.4, "Creating XMLBean Support for Java Functions"](#)
- [Section 3.3.5, "How To Browse and Select a Schema Type"](#)
- [Section 3.3.6, "Physical Data Service from a Java Function - Example Code"](#)

3.3.1 How To Enable Optimistic Locking

These sections describe how to enable optimistic locking in order to update a physical relational data source.

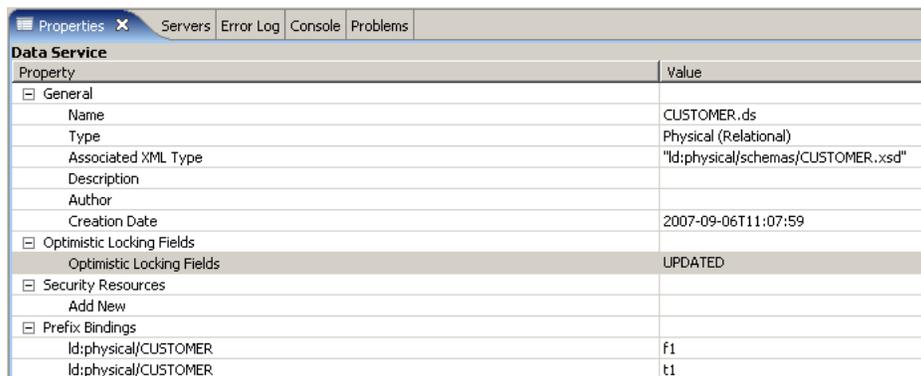
- [Section 3.3.1.1, "Set the Locking Policy"](#)
- [Section 3.3.1.2, "Select the Locking Fields"](#)
- [Section 3.3.1.3, "See Also"](#)

3.3.1.1 Set the Locking Policy

Define the optimistic locking policy on the physical data sources that support your logical data service before you attempt to test an update in Test view or use an update map. Optimistic locking is used with physical data sources that are relational.

The current value of optimistic locking is defined in the **Optimistic Locking Fields** property. You can see this property in the **Properties** tab in Overview mode.

Figure 3–43 Checking the Optimistic Locking Policy



On the Properties tab, under Data Service, properties and values are listed. Optimistic Locking Fields is selected with a value of UPDATED.

Updates to relational data sources use a special XML structure called a data graph. The root element of data graph is <sdo:datagraph>, and the data graph also has a <changeSummary> element.

You can use any of these values for **Optimistic Locking Fields**. They describe how the elements in the data graph compare to fields in the relational data source.

Value of Optimistic Locking Fields	Effect
PROJECTED	All elements in the data graph are mapped to the data source to verify whether it can be updated. Default value.
UPDATED	Only elements that have changed in your data graph are used to verify whether the data source has changed.
SELECTED FIELDS	Selected elements are used to verify whether the data source has changed. The elements must be non-key elements.

To set the locking policy:

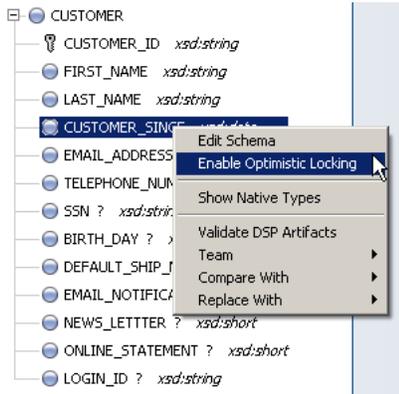
1. Open a physical data service in OSDI Studio.
2. Click the **Overview** tab, then below it, the **Properties** tab.
3. At **Optimistic Locking Fields**, click in the **Value** column, then choose a value.

3.3.1.2 Select the Locking Fields

If you choose **SELECTED FIELDS**, you must also select the fields used to verify changes in the data source. You can select any number of non-key fields. The key fields are used to identify the data records to be updated. If you select a complex element, its child elements also become selected elements.

You can also disable a field once it is selected.

Figure 3–44 Choosing Fields for Optimistic Locking



The context menu displays when you right click on a field. Enable Optimistic Locking is selected.

To select the fields used for optimistic locking:

1. Click the **Overview** tab.
2. Right-click a non-key element in the return type.
Key elements are marked with a key symbol.
3. Choose **Enable Optimistic Locking**.

When you enable optimistic locking for a field, its icon (in the return type in the Overview tab) changes to . You can also see the optimistic locking fields in the pragma statement at the top of the service's Source tab:

```
(::pragma xds <x:xds targetType="t:CREDITRATING" xmlns:x="urn:annotations.ld.oracle.com"
xmlns:t="ld:physical/CREDITRATING"> ... <optimisticLockingFields> <field name="RATING"/>
</optimisticLockingFields>
```

3.3.1.3 See Also

For more information, see the following resources:

3.3.1.3.1 How To

- ["How To Test an Update Procedure"](#) on page 9-49

3.3.1.3.2 Concepts

- Brief Overview of Service Data Objects (for Eclipse for WebLogic)
- Data Programming Model and Update Framework (in depth, for client applications).

3.3.2 How To Update Physical Data Service Metadata

When you first create a physical data service its underlying metadata is, by definition, consistent with its data source. Over time, however, your metadata may become "out of sync" for several reasons:

- The structure of underlying data sources may have changed, in which case it is important to be able to identify those changes so that you can determine when and if you need to update your metadata.
- You have modified schemas or added relationships to your data service.

In some cases relationships between data services will be preserved during metadata update. See ["Using the Update Source Metadata Wizard"](#) on page 3-69 for details.

This section describes the following topics:

- [Section 3.3.2.1, "Topics"](#)
- [Section 3.3.2.2, "Scope of Metadata Update"](#)
- [Section 3.3.2.3, "Important Considerations When Updating Source Metadata"](#)
- [Section 3.3.2.4, "Using the Update Source Metadata Wizard"](#)
- [Section 3.3.2.5, "Inspecting and Reverting Changes Using Local History"](#)

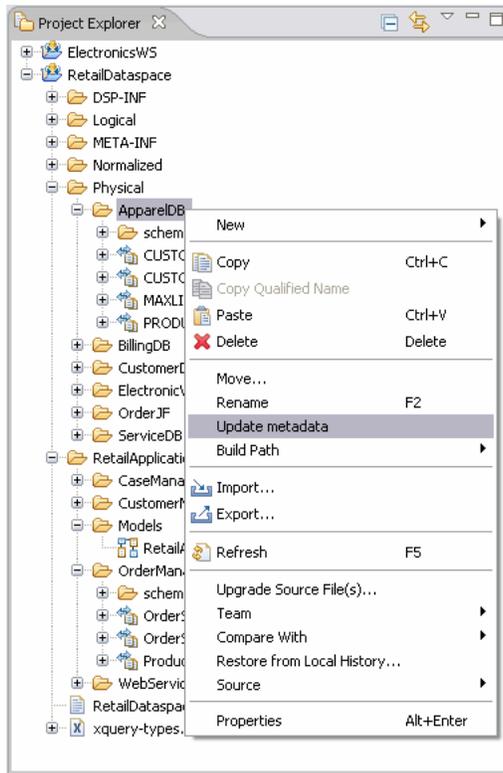
3.3.2.1 Topics

This section describes the following topics:

- Scope of Metadata Update
- Important Considerations When Updating Source Metadata
- Using the Update Source Metadata Wizard
- Inspecting and Reverting Changes Using Local History

In Project Explorer you can use the right-click menu option Update metadata to see if there are any differences between your source metadata files and the underlying source.

Figure 3–45 Update Metadata Option in Project Explorer



In the Project Explorer, a metadata source is highlighted and Update metadata is selected on the right-click menu.

The Update metadata option can be used with:

- Relational table and view associated with changes to the relational database including providerID, the sourceBindingProviderClassName, columns, and optimistic locking fields.
- Web services
- Java functions
- Delimited files

Metadata update cannot be applied to data services based on:

- Relational stored procedures
- XML files

3.3.2.2 Scope of Metadata Update

When you run the Metadata update option, differences between your physical data service and the underlying data source are categorized according to the following scheme:

Category	Meaning
Objects added	The data source contains one or more objects that are not currently represented in the physical data service. From the perspective of the data source, information from the existing data service is added back after the metadata update. Another way to look at this is from the perspective of the data service. In this view, certain artifacts are retained. A typical example is a relationship with another data service. Existing relationships are identified and retained, the metadata is updated to reflect the current data source, and the relationships are added back to the data service.
Objects deleted	One or more objects in the physical data service is not found in the underlying data source. A typical artifact that will be marked for deletion would be a schema that is referenced by an operation (such as a relationship function) in the data service. Objects marked for delete generally appear together. Note: You should carefully inspect the update wizard for items marked for deletion. In the case of schemas, in particular, a prudent course of action would be to retain the schema (uncheck the delete option) unless you are certain that it is not needed by an operation in your data service. Deleting a needed schema will make your data service invalid and undeployable.
Objects changed	One or more objects in the physical data service and the underlying data source do not match and an adjustment will be made. An example of an artifact that will be marked as changed would be if the relational providerID underlying the data source has changed or is unavailable.
Source unavailable	The data source underlying the physical data service could not be accessed.

3.3.2.3 Important Considerations When Updating Source Metadata

The update metadata operation can have both direct and indirect consequences.

Source metadata should be updated with care by someone who is quite familiar with the underlying data source. For example, if you have added a relationship between two physical data services, updating your source metadata may remove the relationship from both data services. If the relationship appears in a model diagram, the relationship line will appear in red, indicating that the relationship is no longer described by the respective data services.

3.3.2.3.1 Direct and Indirect Effects Direct effects apply to physical data services. Indirect effects occur to logical data services, since such services are themselves based — at least indirectly — on physical data services.

For example, if you have created a new relationship between a physical and a logical data service (not a recommended practice), updating the physical data service can invalidate the relationship. In the case of the physical data service, there will be no relationship reference. The logical data service will retain the code describing the relationship but it will be invalid if the opposite relationship notations is no longer be present.

3.3.2.4 Using the Update Source Metadata Wizard

The Update metadata wizard allows you to update your source metadata.

Note: Before attempting to update source metadata you should make sure that your build project has no errors.

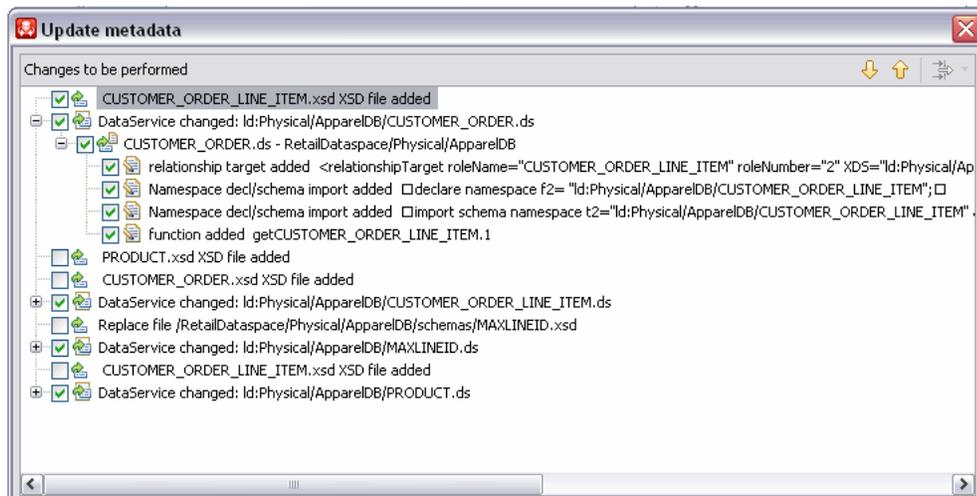
You can perform a metadata update on your entire dataspace project, folders from the project, or any qualified data service. Generally speaking, metadata updates should be performed as specifically as possible.

Note: Use Shift-click or Ctrl-click to select multiple data services or folders in a single dataspace.

After you select your target(s), the wizard identifies the metadata that will be verified and any differences between your metadata and the underlying source.

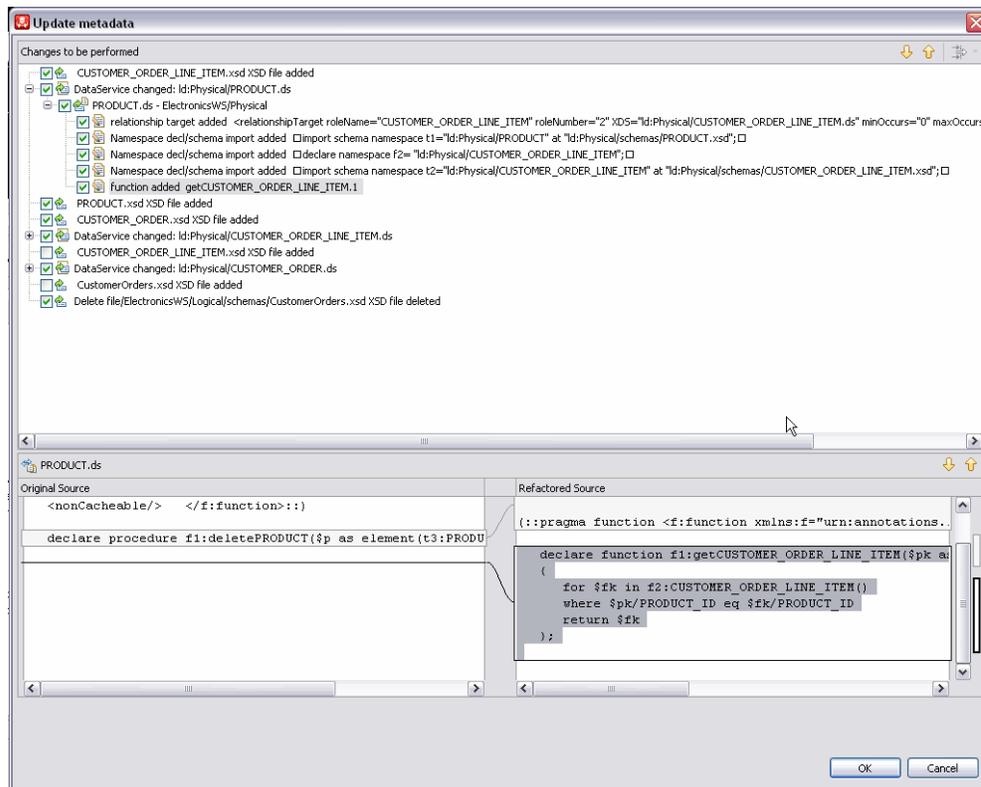
You can select/deselect any data service listed in the dialog using the checkbox to the left of the name. You can also choose to select/deselect specific changes for the data service using the checkbox to the left of the change description.

Figure 3–46 Update Metadata Command



Shows an Update Metadata command for ElectronicsWS project in the sample application.

Figure 3–47 Original Source and Refactored Source Details



The Update Methods screen, there are three sections: Changes to be performed, Product - Original Source and Refracted Source.

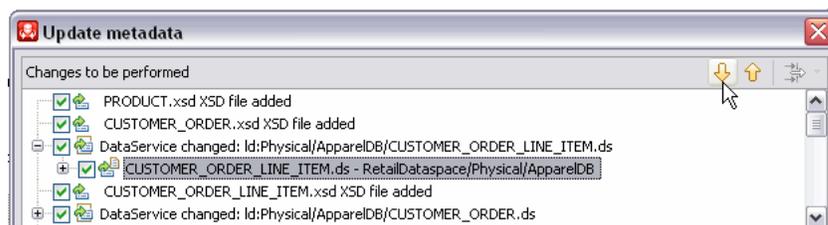
The upper portion of the Update metadata plan shows the changes to be performed. In some cases items are presented and selected (checked). In other cases items are presented but unchecked.

In the details view, the left-hand side shows the current source (called Original Source). The right-hand side shows what the result will be after metadata update (called Refactored Source).

Your only options in the dialog are to select or deselect specific changes using the adjacent checkboxes.

Up/down arrows are available on the Update Metadata titlebar to move through the possible changes. (The Filter Changes option icon next to the arrows is not applicable to metadata update and is not active).

Figure 3–48 Update Metadata Wizard Navigation Arrows



Up/down arrows are available on the Update Metadata titlebar to move through the possible changes. (The Filter Changes option icon next to the arrows is not applicable to metadata update and is not active.)

3.3.2.5 Inspecting and Reverting Changes Using Local History

You can use the Local History option provided with Eclipse to review changes that have been made through the Metadata Update Wizard.

Here are the steps involved:

1. In the Project Explorer right-click on your data service.
2. Select:

Compare With > Local History...

The Compare With Local History window will open. If there have been several changes made, each will be identified through a timestamp.

It is also often possible to revert a metadata update using a similar mechanism:

1. In the Project Explorer right-click on your data service.
2. Select:

Replace With > Local History...

The Replace With Local History window will open. If there have been several changes made, each will be identified through a timestamp.

Note: If you just want to revert to the immediate previous change, use the right-click option:

Replace With > Previous from Local History...

3.3.3 Creating SOAP Handlers for Imported WSDLs

When you import metadata from web services for Oracle Data Service Integrator, you can create SOAP handler for intercepting SOAP requests and responses. The handler will be invoked when a web service method is called. You can chain handlers that are invoked one after another in a specific sequence by defining the sequence in a configuration file.

To create and chain handlers, the following steps are involved:

1. Create a Java Class Implementing the Generic Handler Interface.
2. Compile your intercept handler into a JAR file.
3. Define a Configuration File.
4. Define the Interceptor Configuration.
5. Concluding Actions.

3.3.3.1 Create a Java Class Implementing the Generic Handler Interface

The GenericHandler interface is:

```
javax.xml.rpc.handler.GenericHandler
```

The following code illustrates an example of implementing a generic handler.

```

package WShandler;

import java.util.Iterator;
import javax.xml.rpc.handler.MessageContext;
import javax.xml.rpc.handler.MessageContext;
import javax.xml.soap.SOAPElement;
import javax.xml.rpc.handler.HandlerInfo;
import javax.xml.rpc.handler.GenericHandler;
import javax.xml.namespace.QName;

/**
 * Purpose: Log all messages to the Server console
 */
public class WShandler extends GenericHandler
{
    HandlerInfo hinfo = null;

    public void init (HandlerInfo hinfo) {
        this.hinfo = hinfo;
        System.out.println("*****");
        System.out.println("ConsoleLoggingHandler r: init");
        System.out.println(
            "ConsoleLoggingHandler : init HandlerInfo" + hinfo.toString());
        System.out.println("*****");
    }
}

```

For more information, see *Creating and Using Client-Side SOAP Message Handlers* in Oracle WebLogic documentation.

3.3.3.2 Compile your intercept handler into a JAR file.

The steps are to compile your intercept handler and JAR the class file.

3.3.3.2.1 Define a Configuration File The configuration file specifies the handler chain and the order in which the handlers will be invoked.

The following is an example of the handler chain configuration. The handler-class attribute specifies the fully-qualified name of the handler.

Example 3-1 Code Sample: Handler Chain Configuration

```

<weblogic-wsee-clientHandlerChain
xmlns="http://www.oracle.com/ns/weblogic/90"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:j2ee="http://java.sun.com/xml/ns/j2ee">
<handler>
    <j2ee:handler-name>sampleHandler</j2ee:handler-name>
    <j2ee:handler-class>WShandler.WShandler</j2ee:handler-class>
    <j2ee:init-param>
        <j2ee:param-name>ClientParam1</j2ee:param-name>
        <j2ee:param-value>value1</j2ee:param-value>
    </j2ee:init-param>
</handler>
</weblogic-wsee-clientHandlerChain>

```

3.3.3.2.2 Define the Interceptor Configuration In your Oracle Data Service Integrator application, define the interceptor configuration for the method in the data service to which you want to attach the handler.

Example 3–2 Code Sample: Intercept Configuration

```
xquery version "1.0" encoding "WINDOWS-1252";

(::pragma xds <x:xds xmlns:x="urn:annotations.ld.bea.com"
  targetType="t:echoStringArray_return"
  xmlns:t="ld:SampleWS/echoStringArray_return">
<creationDate>2005-05-24T12:56:38</creationDate>
<webService targetNamespace=
"http://soapinterop.org/WSDLInteropTestRpcEnc"
wsdl="http://webservice.bea.com:7001/rpc/WSDLInteropTestRpcEncService?WSDL"/></x:xds>::)

declare namespace f1 = "ld:SampleWS/echoStringArray_return";

import schema namespace t1 = "ld:AnilExplainsWS/echoStringArray_return" at
  "ld:SampleWS/schemas/echoStringArray_param0.xsd";

(::pragma function <f:function xmlns:f="urn:annotations.ld.bea.com" kind="read"
  nativeName="echoStringArray" nativeLevel1Container="WSDLInteropTestRpcEncService"
  nativeLevel2Container="WSDLInteropTestRpcEncPort" style="rpc">
<params>
  <param nativeType="null"/>

<interceptorConfiguration aliasName="LoggingHandler"
  fileName="ld:SampleWS/handlerConfiguration.xml" />
  </f:function>::)

declare function f1:echoStringArray($x1 as element(t1:echoStringArray_param0)) as
  schema-element(t1:echoStringArray_return) external;
<interceptorConfiguration aliasName="LoggingHandler"
  fileName="ld:testHandlerWS/handlerConfiguration.xml">
```

In the file the `aliasName` attribute specifies the name of the handler chain to be invoked and the `fileName` attribute specifies the location of the configuration file.

3.3.3.2.3 Concluding Actions

- Place the JAR file that was based on the intercept handler (created above) in your project's `dsp-inf/lib` folder.
- Compile and run your application. Your handlers will be invoked in the order specified in the configuration file.

3.3.4 Creating XMLBean Support for Java Functions

Before you can create a Physical Data Service from Java functions, you need to create a `.class` file that contains XMLBean classes based on global elements and compiled versions of your Java functions. This topic describes how to create XMLBean classes based on a schema of your data.

- [Section 3.3.4.1, "Supported XMLBean Standards"](#)
- [Section 3.3.4.2, "Creating XMLBean Classes for Java Functions"](#)
- [Section 3.3.4.3, "See Also"](#)

3.3.4.1 Supported XMLBean Standards

Imported Java functions containing complex types must have a schema that conforms to one of the following XMLBean standards:

Version	URL
Apache	org.apache.xmlbeans
Oracle	com.bea.xml

If your Java routines were compiled under previous versions, they will need to be recompiled before they can be imported.

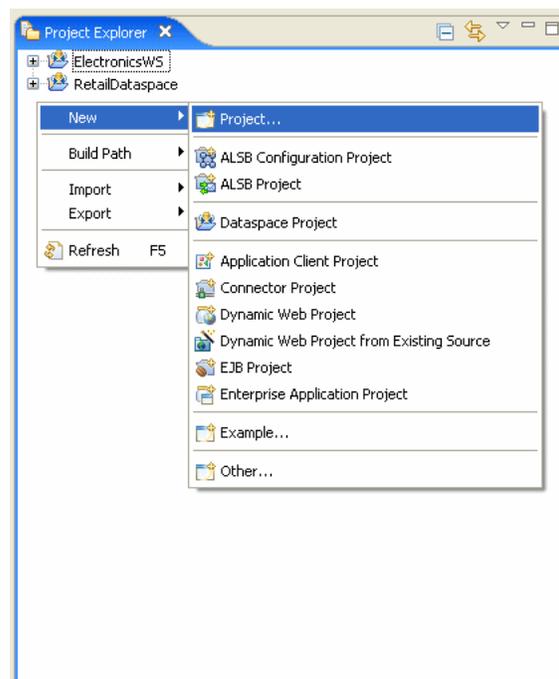
Note: The New Physical Data Service wizard requires that all the complex parameter or return types used by the functions correspond to XMLBean global element types whose content model is an anonymous type. Thus only functions referring to a top level element are imported.

3.3.4.2 Creating XMLBean Classes for Java Functions

This topic describes how to create XMLBean classes based on a schema of your data.

3.3.4.2.1 Creating a New Project You need to create a new project to build the XMLBean classes.

Figure 3–49 *New Project*



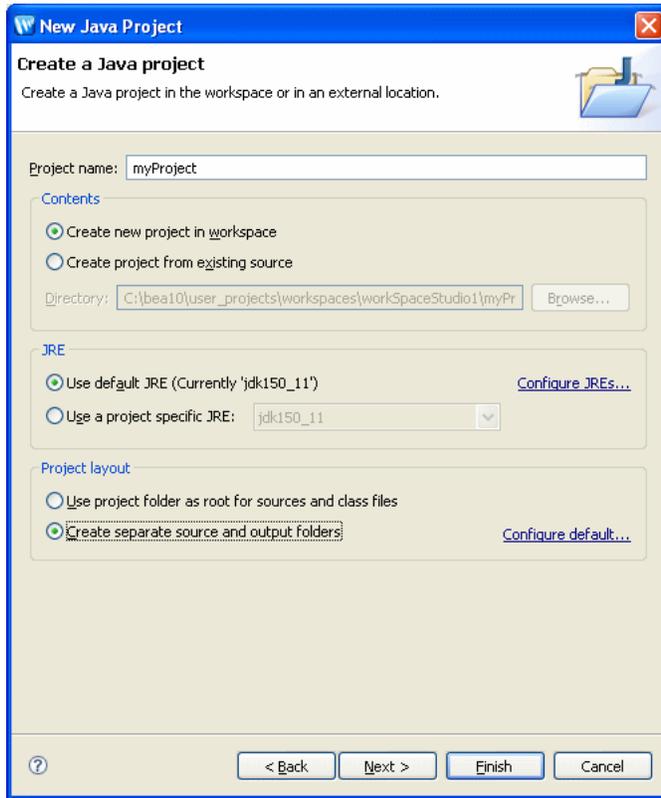
On the right-click menus in the Project Explorer, **New** is highlighted and **Project** is selected.

To create a new project:

1. Using Eclipse for WebLogic, create a new project by right-clicking in the Project Explorer and choosing **New > Project** in the menu. The New Project wizard is launched.

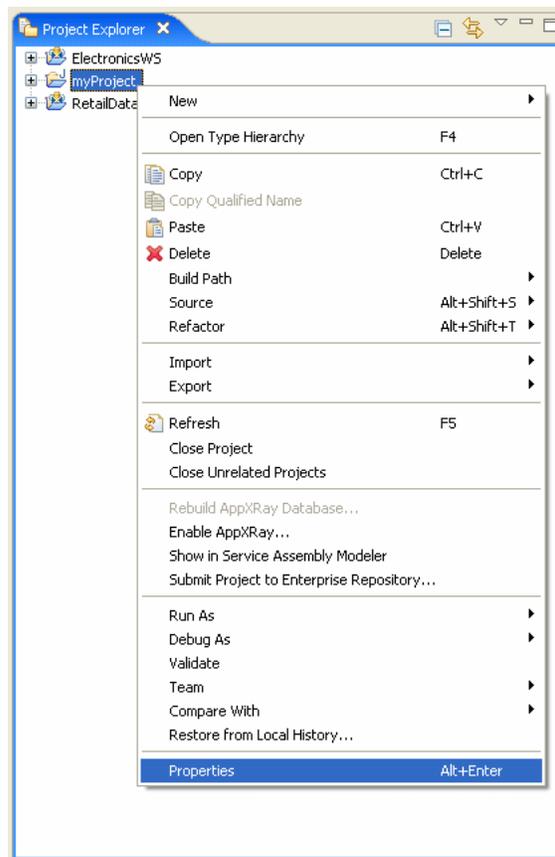
2. Choose **Java > Java Project** and click **Next**.
3. Type a name for the project and select **Create separate source and output folders** in the Project layout area.
4. Click **Finish**. Eclipse for WebLogic creates a new project in the Project Explorer.

Figure 3–50 New Java Project Wizard



The Create a Java project wizard enables you to create a new project. There is a field for you to specify the project name, in this case, myProject. There are two choices for Contents: Create new project in workspace (selected) or Create project from existing source. If Create project from existing source is selected, you can browse to specify a directory. For JRE, there are two choices: Use default JRE and Create separate source and output folders. You can click to configure the JRE. For Project layout, there are two choices: Use project folder as root for sources and class files and Create separate source and output folders. You can click to configure the default.

3.3.4.2.2 Enabling XMLBeans Builder You need to enable the XMLBeans Builder in the project to allow it to create classes based on the Java source and XML schema files.

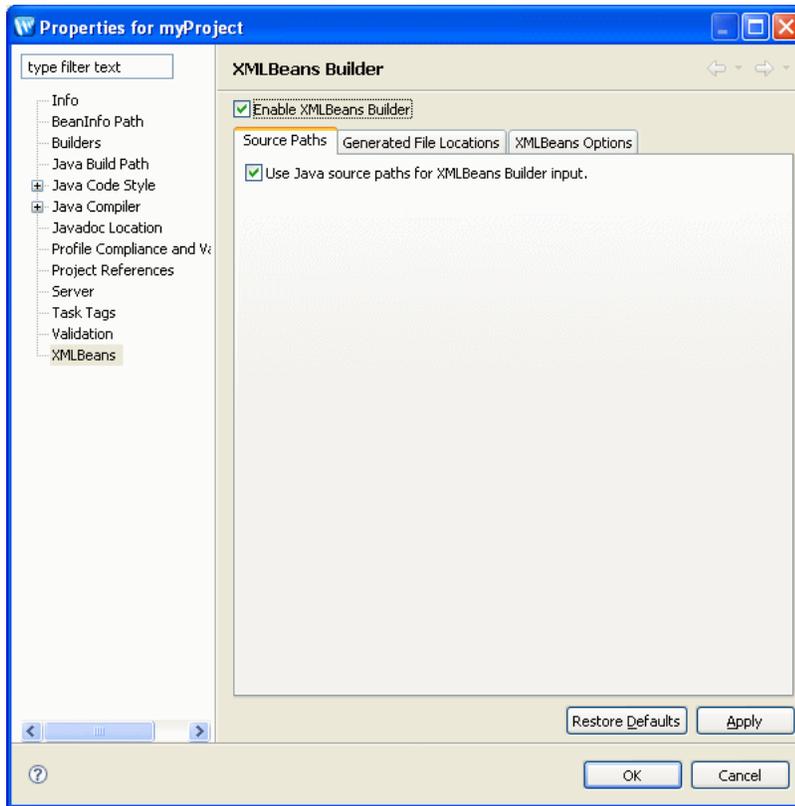
Figure 3–51 Project Properties

In the Project Explorer, the new project is selected and Properties is selected on the right-click menu.

To enable XMLBeans Builder:

1. Right-click the new project, and choose **Properties** from the menu. The Properties dialog appears.
2. Click **XMLBeans**, select the **Enable XMLBeans Builder** checkbox, and click **OK**.

Figure 3–52 Enabling XMLBeans Builder



The XML Beans Builder dialog enables you to create classes based on the Java source and XML schema files. There are two panels. On the left is a field that includes "type filter text." XMLBeans is shown at the bottom of the tree displayed. On the right, the Enable XMLBeans Builder checkbox is selected. On the Source Paths tab, the Use Java source paths for XMLBeans Builder input checkbox is selected. There are Restore Defaults and Apply buttons as well as OK and Cancel.

3.3.4.2.3 Importing Schema and Java Source Files You need to import the schema files and Java source files into the project.

To import the schema and Java source files:

- Copy the schema files representing the data used by the Java functions along with the Java source files into the src folder.

3.3.4.2.4 Creating a Project Reference The final step involves creating a reference to your XMLBeans-based project from the dataspace in which you want to use the Java functions.

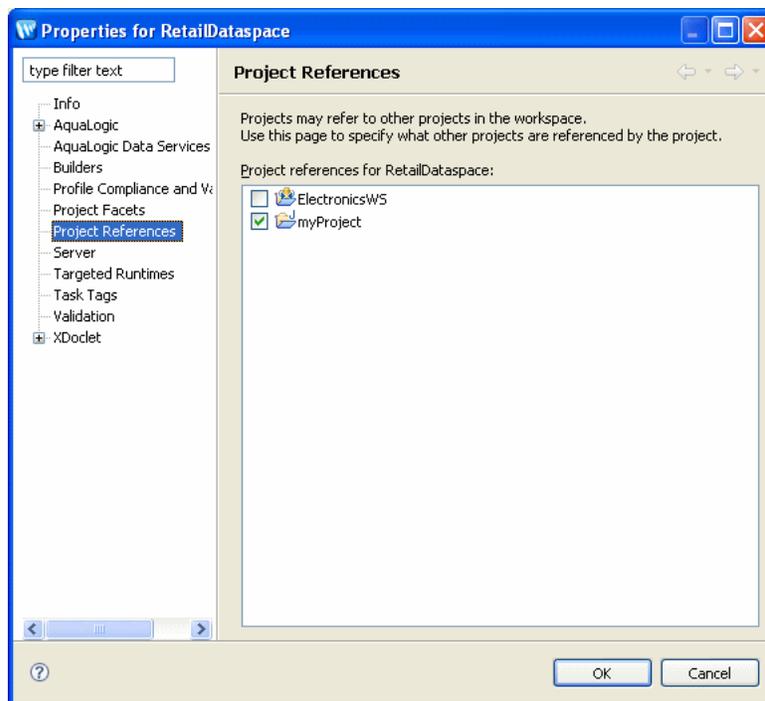
To create the project reference:

1. Right-click the dataspace project in which you want to use the Java functions and choose **Properties** from the menu.
2. Select **Project References** in the Properties dialog.
3. Select your XMLBeans-based project, and click **OK**.

When your project is deployed, Oracle Data Service Integrator does the following:

- Rebuilds your XMLBeans-based project, if required, and generates a JAR file
- Copies the JAR file to the DSP-INF/lib folder in the dataspace project

Figure 3–53 Project References



The Project References dialog lets you specify the projects referenced by the current project. Project References is selected on the right. The project references for the RetailDataspace project are shown on the right with checkboxes. The myProject checkbox is selected.

3.3.4.3 See Also

How Tos

- [Section 3.2.7, "How To Create a Physical Data Service from a Java Function"](#)

Reference

- [Section 3.1.2, "Physical Data Services from Java Functions Overview"](#)
- [Section 3.4, "Example: XMLBeans Example Using a Metadata-rich Java Class"](#)

3.3.5 How To Browse and Select a Schema Type

This topic describes how to use the Schema Type Browser to locate and associate schema types within your dataspace projects.

- [Section 3.3.5.1, "Browsing and Selecting Schema Types"](#)
- [Section 3.3.5.2, "Browsing Schema Types"](#)
- [Section 3.3.5.3, "Selecting a Schema Type"](#)

3.3.5.1 Browsing and Selecting Schema Types

When developing Oracle Data Service Integrator dataspace projects, it is typically the case that the associated schema documents are distributed throughout multiple directories within the project. You can use the Schema Type Browser in Eclipse for WebLogic to quickly locate schema types and associate a schema type with a data service without having to remember the specific schema document that contains the type.

3.3.5.2 Browsing Schema Types

You can browse the schema types within your dataspace projects using the Schema Type Browser. Browsing for schema types enables you to search for global types declared in schema files in the selected project or in all available dataspace projects.

To browse the schema types:

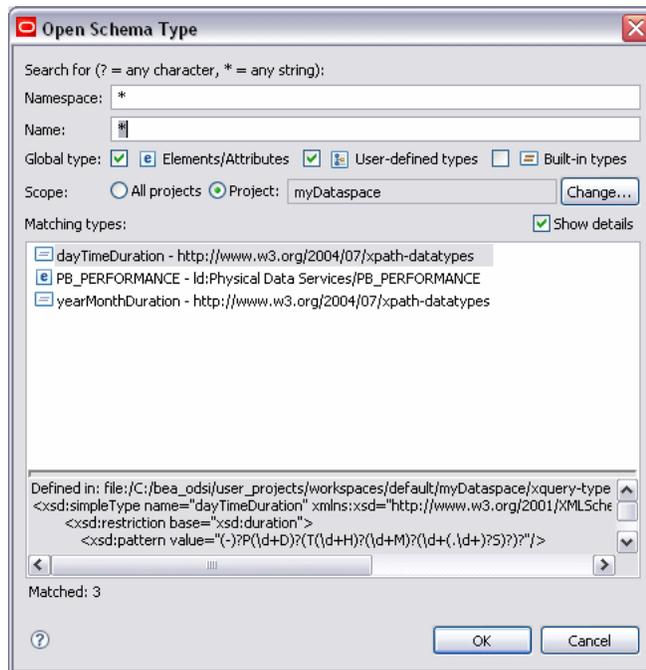
1. Select a dataspace project, a folder in the project, or an .xsd file using the Project Explorer.

This serves as the initial scope within the Schema Type Browser.

2. Choose **Navigate > Open Schema Type** from the main menu.

Alternatively, you can right-click the entity in the Project Explorer and choose **Open Schema Type** from the context-sensitive menu. The **Open Schema Type** dialog appears.

Figure 3–54 *Open Schema Type Dialog*



The **Open Schema Type** dialog has fields where you can type a namespace and name. You can use a question mark (?) to search any character or an asterisk (*) to search any string. There are checkboxes for three global types: **Elements/Attributes**, **User-defined types**, and **Built-in types**. There are two options for **Scope**: **All Projects** and **Project**. A **Change** button lets you change projects. Matching types are displayed.

3. Type a search string for the namespace and the name, as appropriate, in the respective fields.

You can use the wildcard character "?" to represent any character and "*" to represent any string.

4. Select the global types, including elements and attributes, user-defined types, and built-in types, to include in the search.
5. Select the scope of the search.

By default, the scope of the search is the project or folder selected when you opened the Schema Type Browser. You can select all projects or choose a specific folder by clicking the **Change** button. The schema types matching the search criteria appear in the Matching types list, along with the namespace and associated icon. The number of matches appears below the list. The results are updated as you specify new search criteria.

6. Select a type in the list to display the location of the schema file in which the type is declared.

Click **Show Details** to enable this additional information to be displayed about the selected type.

7. Double-click a type (or select the type and click **OK**) to display the type using the default XML schema editor.

3.3.5.3 Selecting a Schema Type

You can select schema types within your Oracle Data Service Integrator dataspace projects using the Schema Type Browser when specifying the signature for an operation or when associating an XML type with a data service.

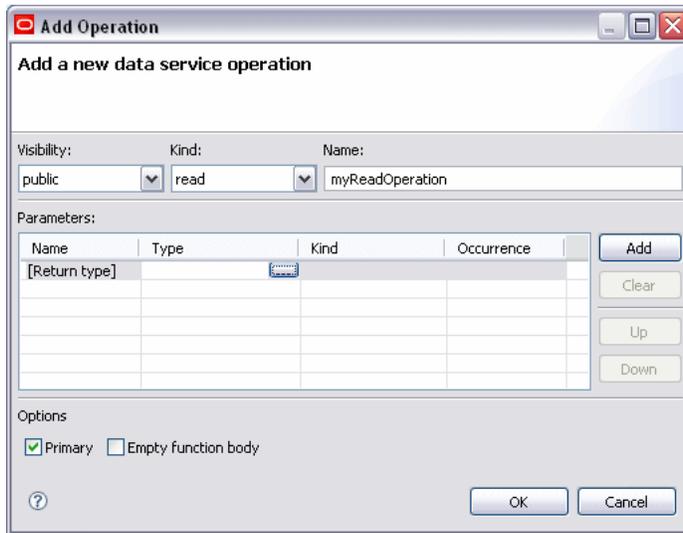
3.3.5.3.1 Editing a Signature You can select a schema type directly from an input field in cases, for instance, when you are adding an operation or editing an operation signature and need to associate a type directly with a parameter. When selecting a schema type in this manner, the schema scope is restricted to the specific project and the matches return only element types.

To select a schema type when editing a signature or adding an operation:

1. Click the browse button that appears in a schema type field.

For example, you could click the browse button in the **Type** field when adding a new data service operation.

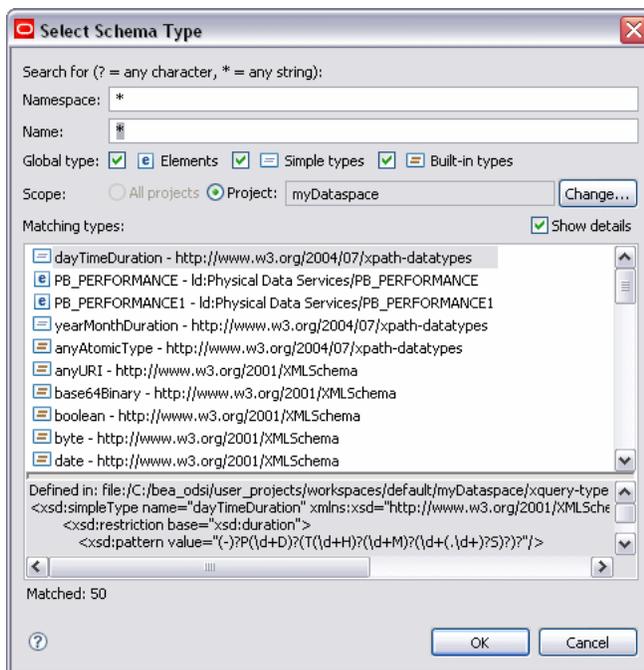
Figure 3–55 Add Operation Dialog



The Add Operation dialog enables you to add a new data service operation. There are two choices from pull-down menus: Visibility (public is selected) and Kind (read is selected). There is a field for you to specify the name of the operation. In the Parameters section, there are four columns: Name, Type, Kind, and Occurrence. There is a browse button in the Type column. Four buttons appear on the right side of the Parameters matrix: Add, Clear, Up, and Down. There are two Options: Primary (selected) and Empty function body.

The Select Schema Type dialog appears.

Figure 3–56 Select Schema Type Dialog



On the Select Schema Type dialog, there are fields that let you specify the Namespace and Name of the schema you want to find. You can use a question mark (?) to search for any character and an asterisk (*) to search for any string. For Global type, there are three choices: Elements (selected), Simple types (selected), and Built-in types (selected). For Scope, there are two choices: All projects and Project (which displays the name of the project). A Change button lets you change the name of the project. The Matching Types section displays the types matched by your search.

2. Type a search string for the namespace and the name, as appropriate, in the respective fields.

You can use the wildcard character "?" to represent any character and "*" to represent any string.

3. Select the global types, including elements, simple types, and built-in types, to include in the search.
4. Select the scope of the search.

By default, the scope of the search is the current dataspace project. You can choose a specific folder by clicking the **Change** button. The schema types matching the search criteria appear in the **Matching types** list, along with the namespace and associated icon. The number of matches appears below the list. The results are updated as you specify new search criteria.

5. Select a type in the list to display the location of the schema file in which the type is declared.

Click **Show Details** to enable this additional information to be displayed about the selected type.

6. Double-click a type (or select the type and click **OK**) to populate the field with the selected type.

3.3.5.3.2 Associating a Type with a Data Service You can associate an XML type with a data service using the Select Schema Type dialog.

To associate an XML type:

1. Open a data service using the Project Explorer.
2. Right-click in the data service and choose **Associate XML Type** from the context-sensitive menu.

The Select Schema Type dialog appears.

3. Select the scope of the search.

By default, the scope of the search is the project selected when you opened the Schema Type Browser. You can choose a specific folder by clicking the **Change** button. The schema types matching the search criteria appear in the Matching types list, along with the namespace and associated icon. The number of matches appears below the list. The results are updated as you specify new search criteria.

4. Select a type in the list to display the location of the schema file in which the type is declared.

Click **Show Details** to enable this additional information to be displayed about the selected type.

5. Double-click a type (or select the type and click **OK**) to associate the XML type with the data service.

3.3.6 Physical Data Service from a Java Function - Example Code

This topic provides examples showing the use of imported Java functions in an XQuery and the processing of complex types.

- [Using a Function Returning an Array of Java Primitives](#)
- [Processing complex types represented via XMLBeans](#)

3.3.6.1 Using a Function Returning an Array of Java Primitives

As an example, the Java function `getRunningTotal` can be defined as follows:

```
public static float[] getRunningTotal(float[] list) {
    if (null == list || 1 >= list.length)
        return list;
    for (int i = 1; i < list.length; i++) {
        list[i] = list[i-1] + list[i];
    }
    return list;
}
```

The corresponding XQuery for executing the above function is as follows:

```
Declare namespace fl="ld:javaFunc/float"
Let $y := (2.0, 4.0, 6.0, 8.0, 10.0)
Let $x := fl:getRunningTotal($y)
Return $x
```

The results of the query is as follows:

```
2.0, 6.0, 12.0, 20.0, 30.0
```

3.3.6.2 Processing complex types represented via XMLBeans

Consider a schema called `Customer` (`customer.xsd`), as shown in the following:

```
<?xml version="1.0" encoding="UTF-8" ?>
<xs:schema targetNamespace="ld:xml/cust:/BEA_BB10000"
xmlns:xs="http://www.w3.org/2001/XMLSchema">
<xs:element name="CUSTOMER">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="FIRST_NAME" type="xs:string" minOccurs="1"/>
      <xs:element name="LAST_NAME" type="xs:string" minOccurs="1"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
</xs:schema>
```

You could compile the schema using XMLBeans to generate a Java class corresponding to the types in the schema.

```
xml.cust.beaBB10000.CUSTOMERDocument.CUSTOMER
```

For more information, see <http://xmlbeans.apache.org>.

Following this, you can use the `CUSTOMER` element as shown in the following:

```
public static xml.cust.beaBB10000.CUSTOMERDocument.CUSTOMER[]
getCustomerListGivenCustomerList(xml.cust.beaBB10000.CUSTOMERDocument.CUSTOMER[]
ipListOfCust) throws XmlException {
    xml.cust.beaBB10000.CUSTOMERDocument.CUSTOMER [] mylocalver = pListOfCust;
```

```

return mylocalver;
}

```

The resulting metadata information produced by the New Physical Data Service wizard will be:

```

(::pragma function <f:function xmlns:f="urn:annotations.ld.oracle.com"
kind="datasource" access="public">
  xml.cust.beaBB10000.CUSTOMERDocument.CUSTOMER [] mylocalver = pListOfCust;
  return mylocalver;
}

```

The corresponding XQuery for executing the above function is:

```

declare namespace fl = "ld:javaFunc/CUSTOMER";
let $z := (
validate(<n:CUSTOMER xmlns:n="ld:xml/cust:/BEA_BB10000"><FIRST_NAME>John2
</FIRST_NAME><LAST_NAME>Smith2</LAST_NAME>
</n:CUSTOMER>),
validate(<n:CUSTOMER xmlns:n="ld:xml/cust:/BEA_BB10000"><FIRST_NAME>John2
</FIRST_NAME><LAST_NAME>Smith2</LAST_NAME>
</n:CUSTOMER>),
validate(<n:CUSTOMER xmlns:n="ld:xml/cust:/BEA_BB10000"><FIRST_NAME>John2
</FIRST_NAME><LAST_NAME>Smith2</LAST_NAME>
</n:CUSTOMER>),
validate(<n:CUSTOMER xmlns:n="ld:xml/cust:/BEA_BB10000"><FIRST_NAME>John2
</FIRST_NAME><LAST_NAME>Smith2</LAST_NAME>
</n:CUSTOMER>))
for $zz in $z
return

```

3.4 Example: XMLBeans Example Using a Metadata-rich Java Class

This topic shows an example of an XMLBeans-based Java function that can be imported into a dataspace project. Importing the Java function into a physical data service results in data service functions corresponding to the source Java functions.

The topic provides the following source listings for the example:

- [Section 3.4.1, "Java Source"](#)
- [Section 3.4.2, "Schema Definition"](#)
- [Section 3.4.3, "Data Service Function"](#)

3.4.1 Java Source

The Java code in the following listing calculates the price of all the items in ORDER_LINE_ITEM to determine the total order amount.

Example 3-3 Java Source Listing

```

package performUpdates;

import java.math.BigDecimal;
import java.util.List;

```

```

import java.util.Iterator;

import org.openuri.temp.sampleapp.customerorder.ELEC_ORDER;
import org.openuri.temp.sampleapp.customerorder.ELEC_ORDER;
import org.openuri.temp.sampleapp.customerorder.ELEC_ORDER.ITEMS.ORDER_LINE_ITEM;

public class ELECOrderUpdate {

    // Change the TotalOrderAmount based on the updated LINE ITEMS.
    // Input and return are the Data objects

    public static ELEC_ORDER[] updateOrders(ELEC_ORDER[] custOrders) {
        System.out.println("\n\n>>> ELECOrderUpdate updateTotalAmount started.");

        if (custOrders == null)
            return custOrders;

        int size = custOrders.length;
        for (int i=0; i<size; i++) {
            ELEC_ORDER order = (ELEC_ORDER)custOrders[i];
            BigDecimal subTotal = new BigDecimal(0);
            BigDecimal totalOrderAmount = new BigDecimal(0);
            BigDecimal saleTax = new BigDecimal(0);

            ITEMS items = order.getItems();
            List itemlist = items.getOrder_Line_Item();
            Iterator item = itemlist.iterator();

            String maxID = "0";

            // Calculate the subTotal and totalOrderAmount
            while (item.hasNext()) {
                ORDER_LINE_ITEM lineitem = (ORDER_LINE_ITEM)item.next();
                BigDecimal quantity = new BigDecimal(Integer.toString(lineitem.getQUANTITY()));
                subTotal = subTotal.add(quantity.multiply(lineitem.getPRICE()));

                lineitem.setLINE_ID(maxID);
                maxID = Integer.toString(new Integer(maxID).intValue() + 1);
            }

            System.out.println(">>> ELECOrderUpdate updateTotalAmount completed.\n\n");
            return custOrders;
        }
    }
}

```

3.4.2 Schema Definition

The schema used to create XMLBeans is shown below. It simply models the structure of the complex element; it could have been obtained by first introspecting the data directly.

Example 3-4 Schema Definition for ELEC_ORDER

```

<_p_r_e_:schema targetNamespace="http://temp.openuri.org/SampleApp/CustomerOrder.xsd"
elementFormDefault="qualified" attributeFormDefault="unqualified"
xmlns="http://temp.openuri.org/SampleApp/CustomerOrder.xsd" xmlns:_p_r_e_
="http://www.w3.org/2001/XMLSchema" xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:conv="http://www.openuri.org/2002/04/soap/conversation/" xmlns:s0="http://www.openuri.org/"
xmlns:cw="http://www.openuri.org/2002/04/wsd1/conversation/"

```

```

xmlns:http="http://schemas.xmlsoap.org/wsdl/http/" xmlns:s0="http://www.openuri.org/"
xmlns:jms="http://www.openuri.org/2002/04/wsdl/jms/"
xmlns:mime="http://schemas.xmlsoap.org/wsdl/mime/" xmlns:s="http://www.w3.org/2001/XMLSchema"
xmlns:s0="http://www.openuri.org/" xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
xmlns:s0="http://www.openuri.org/" xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
xmlns:soapenc="http://schemas.xmlsoap.org/soap/encoding/">
  <s:element name="ELEC_ORDER">
    <s:complexType>
      <s:sequence>
        <s:element name="ORDER_ID" type="xsd:string" minOccurs="1"/>
        <s:element name="CUSTOMER_ID" type="xsd:string" minOccurs="1"/>
        <s:element name="ORDER_DATE" type="xsd:date" minOccurs="1"/>
        <s:element name="SHIPMENT_METHOD" type="xsd:string" minOccurs="1"/>
        <s:element name="HANDLING_CHARGE" type="xsd:decimal" minOccurs="1"/>
        <s:element name="SUBTOTAL" type="xsd:decimal" minOccurs="1"/>
        <s:element name="TOTAL_ORDER_AMOUNT" type="xsd:decimal" minOccurs="1"/>
        <s:element name="SALE_TAX" type="xsd:decimal" minOccurs="1"/>
        <s:element name="SHIP_TO" type="xsd:string" minOccurs="1"/>
        <s:element name="SHIP_TO_NAME" type="xsd:string" minOccurs="1"/>
        <s:element name="BILL_TO" type="xsd:string" minOccurs="1"/>
        <s:element name="ESTIMATED_SHIP_DATE" type="xsd:date" minOccurs="1"/>
        <s:element name="STATUS" type="xsd:string" minOccurs="1"/>
        <s:element name="TRACKING_NUMBER" type="xsd:string" minOccurs="0" nillable="true"/>
        <s:element name="ITEMS">
          <s:complexType>
            <s:sequence>
              <s:element name="LINE_ID" type="xsd:string" minOccurs="1"/>
              <s:element name="ORDER_ID" type="xsd:string" minOccurs="1"/>
              <s:element name="PROD_ID" type="xsd:string" minOccurs="1"/>
              <s:element name="PROD_DESC" type="xsd:string" minOccurs="1"/>
              <s:element name="QUANTITY" type="xsd:int" minOccurs="1"/>
              <s:element name="PRICE" type="xsd:decimal" minOccurs="1"/>
              <s:element name="STATUS" type="xsd:string" minOccurs="1"/>
            </s:sequence>
          </s:complexType>
        </s:element>
      </s:sequence>
    </s:complexType>
  </s:element>
</_p_r_e_:schema>

```

3.4.3 Data Service Function

The following listing shows the generated data service function based on the imported Java code.

Example 3–5 *Generated Data Service Function*

```

xquery version "1.0" encoding "UTF-8";

(::pragma xfl <x:xfl xmlns:x="urn:annotations.ld.bea.com">
<creationDate>2008-04-23T10:40:54</creationDate>
<creationDate>2008-04-23T10:40:54</creationDate>
<javaFunction class="performUpdates.ELECOOrderUpdate"/>
</x:xfl>::)

```

```

declare namespace fl = "ld:UpdateOrder";

import schema namespace t1 = "http://temp.openuri.org/SampleApp/CustomerOrder.xsd" at
"ld:schemas/CustomerOrders.xsd";

(::pragma function <f:function xmlns:f="urn:annotations.ld.bea.com" visibility="protected"
kind="library" isPrimary="false" nativeName="updateOrders">

<params>
  <param nativeType="[Lorg.openuri.temp.sampleapp.customerorder.ELEC_ORDER;"/>
</params>
</f:function>::)

declare function fl:updateOrders($parameter1 as element(t1:ELEC_ORDER)*) as
schema-element(t1:ELEC_ORDER)* external;

```

3.5 Reference

The following are reference sections for creating and updating physical data services:

- [Section 3.5.1, "Stored Procedure Configuration Reference"](#)
- [Section 3.5.2, "Simple Java Types and Their XQuery Counterparts"](#)

3.5.1 Stored Procedure Configuration Reference

The following topics provide detailed information regarding various configuration options associated with creating data services based on stored procedures.

- [Section 3.5.1.1, "In Mode, Out Mode, Inout Mode"](#)
- [Section 3.5.1.2, "Procedure Profile"](#)
- [Section 3.5.1.3, "Supporting Stored Procedures with Nullable Input Parameter\(s\)"](#)

3.5.1.1 In Mode, Out Mode, Inout Mode

In, **Out**, and **Inout** mode settings determine how a parameter passed to a stored procedure is handled.

Parameter Mode	Effect
In	Parameter is passed by reference or value.
Inout	Parameter is passed by reference.
Out	Parameter is passed by reference. However the parameter being passed is first initialized to a default value. If your stored procedure has an OUT parameter requiring a complex element, you may need to provide a schema.

3.5.1.2 Procedure Profile

Each element in a stored procedure is associated with a type. If the item is a simple type, you can simply choose from the pop-up list of types. If the type is complex, you may need to supply an appropriate schema. Click on the schema location button and either enter a schema pathname or browse to a schema. The schema must reside in your application.

After selecting a schema, both the path to the schema file and the URI appear.

3.5.1.2.1 Complex Parameter Types Complex parameter types are supported under only three conditions:

- As the output parameter
- As the Return type
- As a rowset

3.5.1.2.2 About Rowsets A rowset type is a complex type.

The rowset type contains a sequence of a repeatable elements (for example called CUSTOMER) with the fields of the rowset.

In some cases the wizard can automatically detect the structure of a rowset and create an element structure. However, if the structure is unknown, you will need to provide it.

Note: All rowset-type definitions must conform to this structure.

The name of the rowset type can be:

- The parameter name (in case of a input/output or output only parameter).
- An assigned name.
- The referenced element name (result rowsets) in a user-specified schema.

Not all databases support rowsets. In addition, JDBC does not report information related to defined rowsets.

3.5.1.2.3 Using Rowset Information In order to create data services from stored procedures that use rowset information, you need to supply the correct ordinal (matching number) and a schema. If the schema has multiple global elements, select the one you want from the Type column. Otherwise the type used match the first global element in your schema file.

The order of rowset information is significant; it must match the order in your data source. Use the Move Up / Move Down commands to adjust the ordinal number assigned to the rowset.

Note: XML types in data services generated from stored procedures do not display native types. However, you can view the native type in the Source editor; it is located in the pragma section.

3.5.1.2.4 Stored Procedure Version Support Only the most recent version of a particular stored procedure can be imported into Oracle Data Service Integrator. For this reason you cannot identify a stored procedure version number when creating a physical data service based on a stored procedure. Similarly, adding a version number for your stored procedure in the Source editor will result in a query exception.

3.5.1.3 Supporting Stored Procedures with Nullable Input Parameter(s)

If you know that an input parameter of a stored procedure is nullable (can accept null values), you can change the signature of the function in Source View to make such parameters optional by adding a question mark at end of the parameter.

For example (question-mark (?)):

```
function myProc($arg1 as xs:string) ...
```

would become:

```
function myProc($arg1 as xs:string?) ...
```

3.5.2 Simple Java Types and Their XQuery Counterparts

The following outlines the mapping between simple Java types and the corresponding XQuery or schema types:

Java Simple or Defined Type	XQuery/Schema Type
boolean	xs:boolean
byte	xs:byte
char	xs:char
double	xs:double
float	xs:float
int	xs:int
long	xs:long
short	xs:short
string	xd:string
java.lang.Date	xs:datetime
java.lang.Boolean	xs:boolean
java.math.BigInteger	xs:integer
java.math.BigDecimal	xs:decimal
java.lang.Byte	xs:byte
java.lang.Char	xs:char
java.lang.Double	xs:double
java.lang.Float	xs:float
java.lang.Integer	xs:integer
java.lang.Long	xs:long
java.lang.Short	xs:short
java.sql.Date	xs:date
java.sql.Time	xs:time
java.sql.Timestamp	xs:datetime
java.util.Calendar	xs:datetime

Java functions can consume parameters and return values of the following types:

- Java primitives and types listed in the previous table
- Apache XMLBeans
- Oracle XMLBeans
- SDO DataObject (typed or untyped)

Note: The elements or types referred to in the schema should be global elements.

3.6 Related Topics

The following section describes how to add an external function to an existing physical data service.

3.6.1 How To Add an External Function to an Existing Physical Data Service

You can add qualified external operations (functions and procedures) from the same data source to existing physical data services based on:

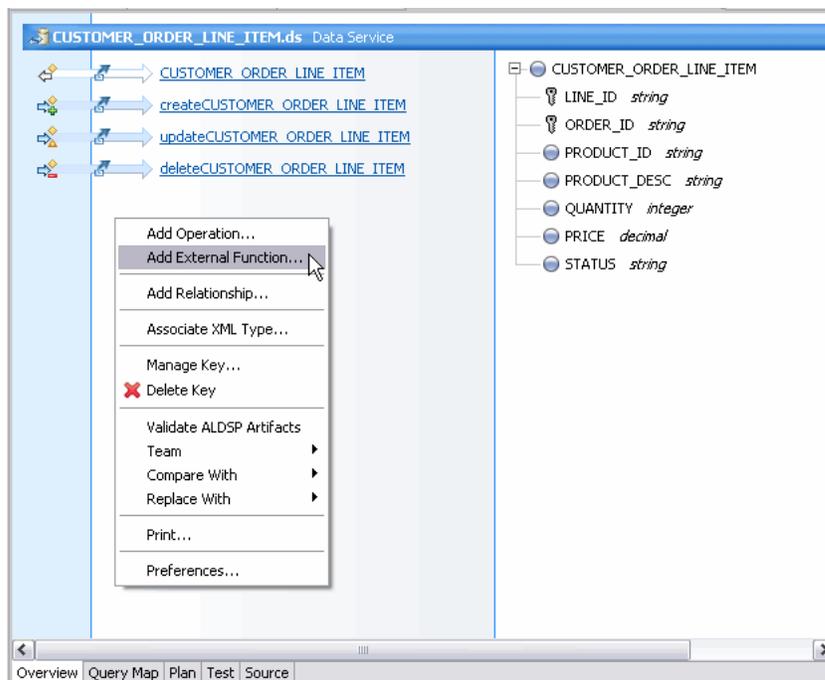
- Relational
- Web service
- Java functions

This is a very convenient way of enhancing a data service based on changes in underlying data or other business needs.

The steps involved in adding an external function to a qualified data service are:

1. Open the data service in Project Explorer.
2. Select **Overview mode**.
3. Right click > **Add External Function...**
4. A wizard appropriate to the data service type will appear. Complete the steps as you would when creating a data service. For example, select from the set of currently unselected operations in the WSDL that underlies a web service-based data service.

Figure 3–57 Adding an External Function to a Data Service



The CUSTOMER_ORDER_LINE_ITEM.ds Data Service panel is shown in the Project Explorer. The right-click menu is displayed and Add External Function is selected.

External operations cannot be added to physical data services based on:

- XML data
- Delimited data

Table 3–12 Qualified Operation and Physical Data Service Type Matrix

Artifact	Physical Data Service Type	Comment
Operation	Web Service	Only visible operations will be from the WSDL that underlies the physical data service.
Function	Java	Only functions from the Java class defined by the underlying data service will be visible.
Stored Procedure	Relational	Only stored procedures from the same data source defined by the underlying data service will be visible.
SQL Statement	Relational	Query must be to the same database as that underlying the data service.

3.6.1.1 Additional Constraints

- Table-based functions cannot be added to data services.
- Only library operations can be added to library data services.
- Read or primary create-update-delete functions can be added to entity data services as long as the entity data services constraints are not violated.

Designing Logical Data Services

Logical dataservices let you create a new loosely coupled architecture by piecing together the data assets you already have. This means combining data from relational data sources, web services, XML files, other files, or Java functions.

This chapter describes the following topics:

- [Concepts](#)
- [How-to](#)
- [Examples](#)
- [Reference](#)
- [Related Topics](#)

4.1 Concepts

This section describes the following topics:

- [Building Logical Entity Data Services](#)
- [Data Service Keys](#)
- [XML Types and Return Types](#)

4.1.1 Building Logical Entity Data Services

This topic introduces you to logical entity data services.

- [The Benefits of Logical Services](#)
- [Design View](#)
- [Query Map View](#)
- [Update Map View](#)
- [Test View](#)
- [See Also](#)

4.1.1.1 The Benefits of Logical Services

The benefit of data services is the ability to combine multiple data sources of different types into service-oriented architectures. Enterprise data is often stored in relational databases, non-relational databases, packaged applications (such as SAP, PeopleSoft, Siebel, and others), custom applications, or files of various types. You might also be accessing data from web services.

The goal is to create a new loosely coupled architecture by piecing together the data assets you already have. In a practical sense, this means combining data from relational data sources, web services, XML files, other files, or Java functions. Logical data services are of two types, entity and library.

Logical entity services allow you to design, model, and create a data view from many underlying data sources. Logical library services are simply a collection of related functions and procedures within a data service container. This topic introduces logical entity services.

On a tangible level, a logical entity service is an XQuery source file with functions and procedures that act on data. A logical entity service has:

- Exactly one XML schema that represents the data the service returns (its return type).
- Any number of create, update, or delete procedures, where up to one of each type is primary.
- Any number of library functions and procedures.
- Any number of relationships with other entity services.

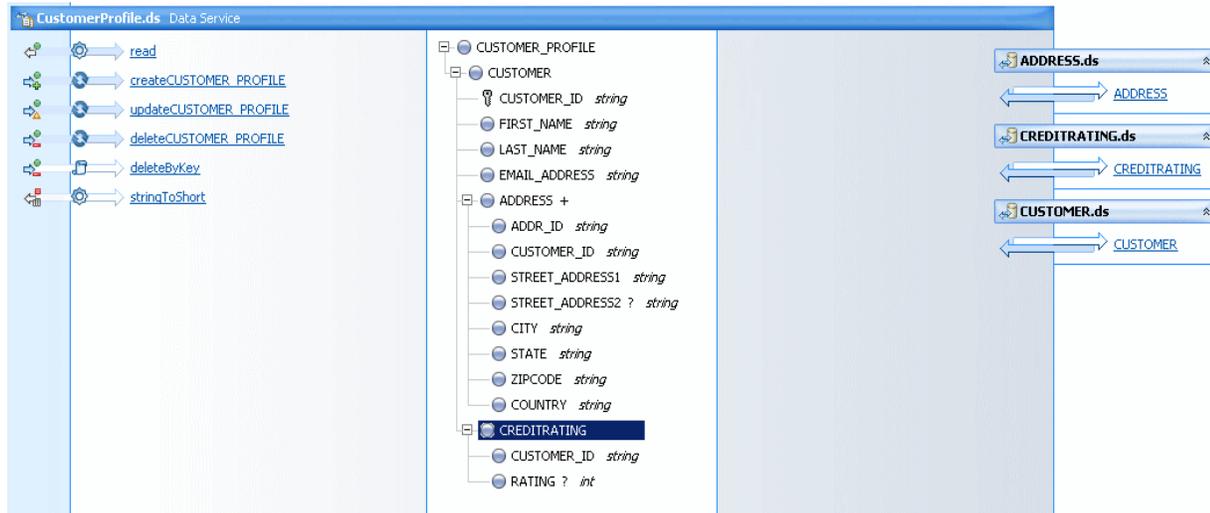
In addition, a logical entity service must have a primary read function if you want the service to have an update map.

4.1.1.2 Design View

Logical data services have their foundation in XML web services. The backbone of a logical data service is its return type, which is a combination of data you design expressed as an XML schema.

You can see the return type in the **Overview** tab in Eclipse for WebLogic.

Figure 4-1 Design View of a Logical Data Service



On the **Overview** tab in the Eclipse for WebLogic, the logical service's return type is shown in the center. You can right-click it to see the XML schema source. On the left, you see the definitions for the service. On the right, you see other data services that underlie the logical data service.

The underlying data services can be physical or logical.

The beauty of a logical data service is that a return type is a model. Logical models capture the complexity of data integration once, and allow you to write clients that remain the same even when underlying physical data sources change.

The structure of a return type does not need to match the structure of the underlying data sources. Here, the CUSTOMER element has a 1-to-many relationship with its child element ADDRESS, and a 1-to-1 relationship with its other child element, CREDITRATING. Each complex element represents a separate physical data source.

Example 4–1 The Return Type Schema

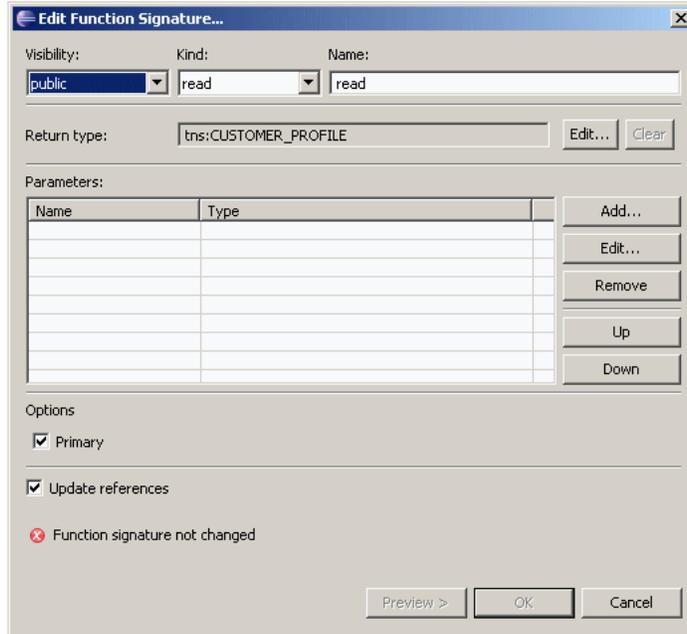
```
<?xml version="1.0" encoding="UTF-8" ?>
<xs:schema targetNamespace="ld:logical/CustomerProfile"
xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CUSTOMER_PROFILE">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="CUSTOMER">
          <xs:complexType>
            <xs:sequence>
              <xs:element name="CUSTOMER_ID" type="xs:string"/>
              <xs:element name="FIRST_NAME" type="xs:string"/>
              <xs:element name="LAST_NAME" type="xs:string"/>
              <xs:element name="EMAIL_ADDRESS" type="xs:string"/>
              <xs:element name="ADDRESS" maxOccurs="unbounded">
                <xs:complexType>
                  <xs:sequence>
                    <xs:element name="ADDR_ID" type="xs:string"/>
                    <xs:element name="CUSTOMER_ID" type="xs:string"/>
                    <xs:element name="STREET_ADDRESS1" type="xs:string"/>
                    <xs:element name="STREET_ADDRESS2" type="xs:string" minOccurs="0"/>
                    <xs:element name="CITY" type="xs:string"/>
                    <xs:element name="STATE" type="xs:string"/>
                    <xs:element name="ZIPCODE" type="xs:string"/>
                    <xs:element name="COUNTRY" type="xs:string"/>
                  </xs:sequence>
                </xs:complexType>
              </xs:element>
            </xs:sequence>
          </xs:complexType>
        </xs:element>
        <xs:element name="CREDITRATING" maxOccurs="1">
          <xs:complexType>
            <xs:sequence>
              <xs:element name="CUSTOMER_ID" type="xs:string"/>
              <xs:element name="RATING" type="xs:int" minOccurs="0"/>
            </xs:sequence>
          </xs:complexType>
        </xs:element>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

However, this structure is only by design. You could also have designed the return type with fewer elements, or in a flat structure, depending on how you want the service to return data.

4.1.1.2.1 The Primary Read Function The functions and procedures in a logical entity service are implemented in XQuery, which queries XML data much as SQL queries relational data. You can get information about any function or procedure by right-clicking it in the Overview tab.

A read function, for example, often takes no parameters and returns an instance of the return type.

Figure 4–2 Viewing the Signature of a Read Function



The Edit Function Signature dialog lets you view and edit the signature of a selected function. There are three fields at the top of the dialog: View (drop-down menu), Kind (drop-down menu), and Name (in this case, read). The Return Type is displayed. An Edit button lets you edit the return type. A Clear button lets you clear your edits. There is a Parameters matrix with two columns: Name and Type. There are five buttons: Add, Edit, Remove, Up, and Down. The Primary and Update references options are selected. Function signature not changed is displayed.

In a logical entity service, you can designate one read function as primary. A primary read function captures the main data integration logic in the service. Oracle Data Service Integrator generates the create, update, and delete procedures and the update map from the primary read function.

You can see the source code of the primary read function in the Source tab.

Example 4–2 Checking the Primary Read Function Source

```
declare function tns:read() as element(tns:CUSTOMER_PROFILE) *{
for $CUSTOMER in cus1:CUSTOMER()
return
  <tns:CUSTOMER_PROFILE>
    <CUSTOMER>
      <CUSTOMER_ID>{fn:data($CUSTOMER/CUSTOMER_ID)}</CUSTOMER_ID>
      <FIRST_NAME>{fn:data($CUSTOMER/FIRST_NAME)}</FIRST_NAME>
      <LAST_NAME>{fn:data($CUSTOMER/LAST_NAME)}</LAST_NAME>
```

```

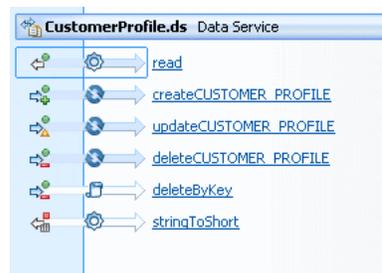
<EMAIL_ADDRESS>{fn:data($CUSTOMER/EMAIL_ADDRESS)}</EMAIL_ADDRESS>
{
  for $ADDRESS in add:ADDRESS()
  where $CUSTOMER/CUSTOMER_ID eq $ADDRESS/CUSTOMER_ID
  return
  <ADDRESS>
    <ADDR_ID>{fn:data($ADDRESS/ADDR_ID)}</ADDR_ID>
    <CUSTOMER_ID>{fn:data($ADDRESS/CUSTOMER_ID)}</CUSTOMER_ID>
    <STREET_ADDRESS1>{fn:data($ADDRESS/STREET_ADDRESS1)}</STREET_ADDRESS1>
    <STREET_ADDRESS2?>{fn:data($ADDRESS/STREET_ADDRESS2)}</STREET_ADDRESS2>
    <CITY>{fn:data($ADDRESS/CITY)}</CITY>
    <STATE>{fn:data($ADDRESS/STATE)}</STATE>
    <ZIPCODE>{fn:data($ADDRESS/ZIPCODE)}</ZIPCODE>
    <COUNTRY>{fn:data($ADDRESS/COUNTRY)}</COUNTRY>
  </ADDRESS>
}
{
  for $CREDITRATING in cre:CREDITRATING()
  where $CUSTOMER/CUSTOMER_ID eq $CREDITRATING/CUSTOMER_ID
  return
  <CREDITRATING>
    <CUSTOMER_ID>{fn:data($CREDITRATING/CUSTOMER_ID)}</CUSTOMER_ID>
    <RATING?>{fn:data($CREDITRATING/RATING)}</RATING>
  </CREDITRATING>
}
</CUSTOMER>
</tns:CUSTOMER_PROFILE>
};

```

This read function returns a **CUSTOMER_PROFILE** element with a nested **CUSTOMER** element. Each **CUSTOMER** element has some number of **ADDRESS** elements and some number of **CREDITRATING** elements, where the **CUSTOMER_ID** in **ADDRESS** or **CREDITRATING** matches the **CUSTOMER_ID** in **CUSTOMER**. (The XQuery where clauses create table joins; see "Add a Where Clause to a Query" on page 6-16).

4.1.1.2.2 Create, Update, and Delete Procedures A logical entity service also typically has create, update, and delete procedures that act on underlying data sources. (The difference between a function and a procedure is that a procedure can have side effects, while a function cannot; see "Data Service Types and Functions" on page 1-52).

Figure 4-3 Viewing Functions and Procedures

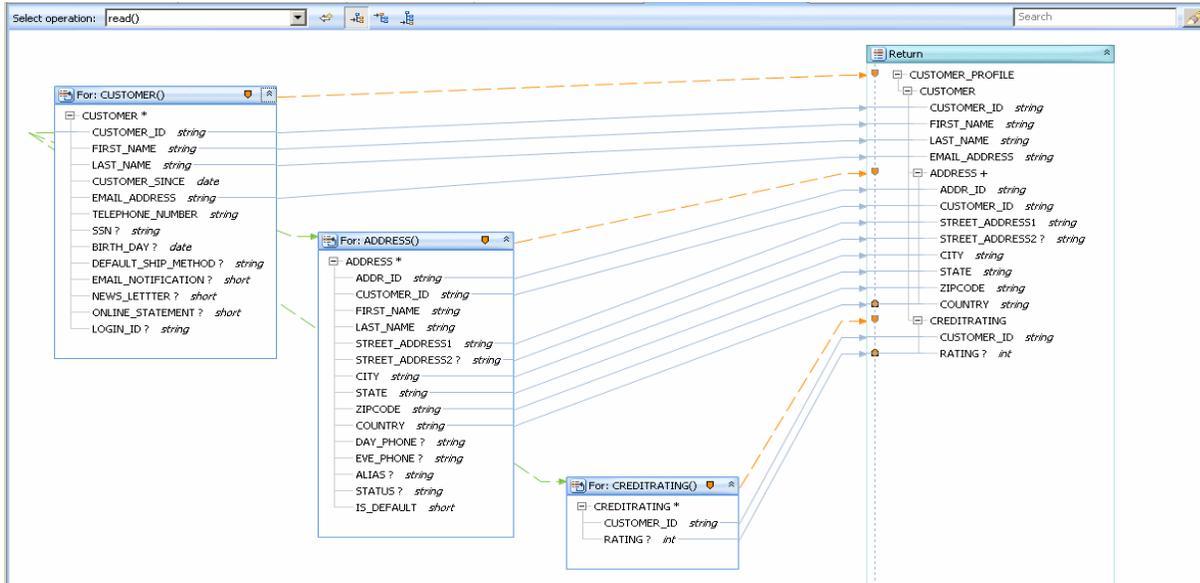


The Customer Profile service has one create procedure, one update procedure, and two delete procedures. It also has a library procedure named `stringToShort`, which casts two data types.

4.1.1.3 Query Map View

The Query Map view maps elements in data sources to the return type.

Figure 4–4 Mapping Data Sources to the Return Type



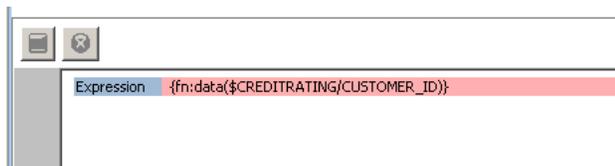
The Query Map view maps elements in the data sources to the return type. Data sources are on the left and the return type is on the right. The blue lines map elements from the data sources to elements in the return.

The green dashed lines between the data source blocks create joins, which become where clauses in the XQuery source, for example:

```
for $ADDRESS in add:ADDRESS()
where $CUSTOMER/CUSTOMER_ID eq $ADDRESS/CUSTOMER_ID
return
```

If you click a data element (not a container element) in the return type, you see its XQuery expression in the expression editor.

Figure 4–5 Mapping Data in an XQuery Expression

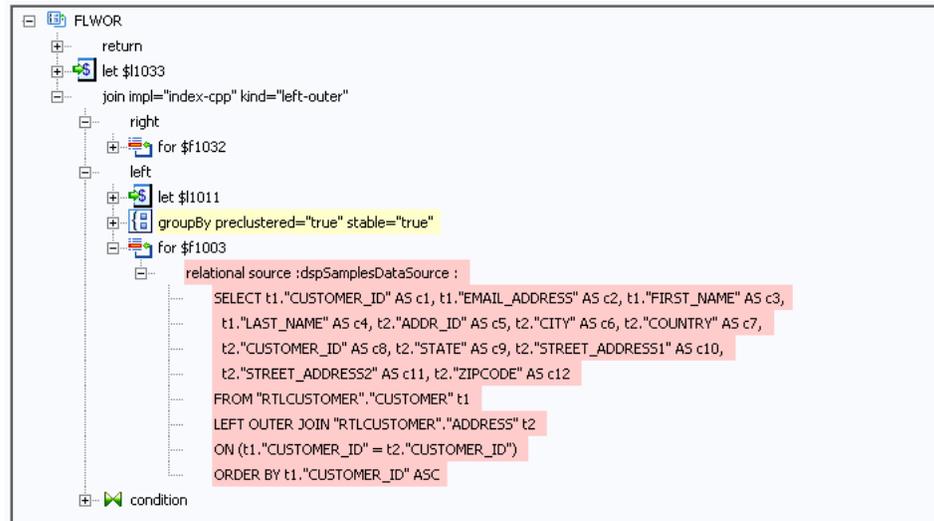


In the Expression editor, an XQuery expression is shown.

Notice that the mapping expressions use the built-in XQuery function `fn:data`, which extracts the data value from an XML element.

As you map elements visually in the Query Map, Oracle Data Service Integrator creates XQuery source (for example, the `read` function shown above). The XQuery source is later converted to SQL queries, which you can see in Plan view.

Figure 4–6 Viewing a SQL Query in Plan View



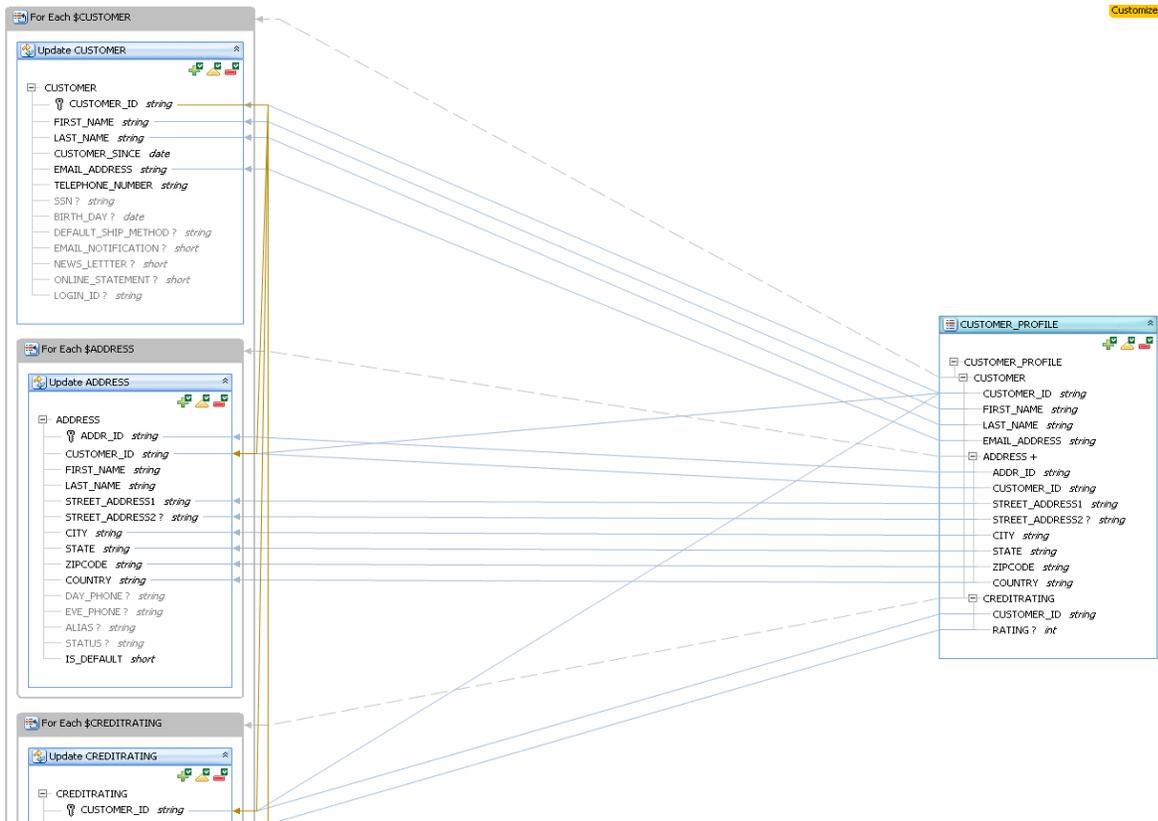
In the query plan, the left outer join between the CUSTOMER and ADDRESS relational tables is shown. This was created by the green dashed line drawn between the Customer and Address blocks in Query Map view.

When you build XQuery functions and procedures visually in Query Map view or by editing in Source view, you can test and run them on an Oracle Data Service Integrator server. During server runtime, the functions and procedures are compiled into an executable query plan. Examine the query plan before you finalize the queries. Query Plan view gives you a peek into a query's execution logic and flags potential performance and memory problems. Building XQuery functions is an iterative process of test, view plan, and edit.

4.1.1.4 Update Map View

While Query Map view shows how a service reads from data sources, Update Map view shows how the service writes data to them.

Figure 4–7 Checking Update Map View

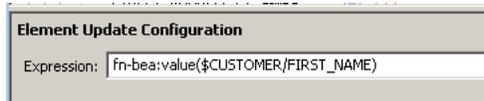


In the map view, the data sources are on the left, with updates coming from the return type on the right.

The return type is available to client applications, where users update data.

The blocks on the left are update blocks. Each mapped element in an update block has an XQuery expression that defines how the element is updated. You can see the expression in the expression editor below the mapping area.

Figure 4–8 Viewing an XQuery Expression



The Element Update Configuration dialog shows an update expression.

Oracle Data Service Integrator generates the update map for you when you create a logical data service under these conditions:

- Your service has a primary read function
- Your service has a primary read function

(If you are using other data source types, you must edit the update template.)

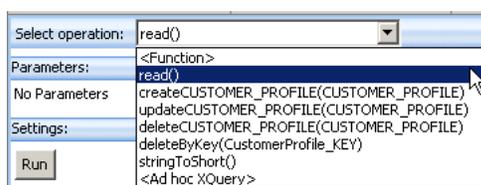
You can then customize the update map and test it in Test view, without programming.

An application client uses the Service Data Objects programming model to update data sources. SDO is an application framework that allows you to update data sources while disconnected from them, using a flexible, optimistic concurrency model. You use only one API, the SDO API, to update multiple data sources -- relational, web service, XML files, and so on.

4.1.1.5 Test View

The Test view available in Eclipse for WebLogic works like a built-in client where you can easily test any function or procedure in the data service, before you build a custom client.

Figure 4–9 Selecting an Operation



The Test view lets you select an operation to test. Select operation is a drop-down menu that displays the operation, in this case, read(). Under Parameter, No Parameters is displayed. Under Settings, there is a Run button.

Testing a read function, for example, returns data as the service would to a client, in the shape of the return type.

Figure 4–10 Reading Customer Profile Data

Name	Value
[-] CUSTOMER_PROFILE	
[-] CUSTOMER	
CUSTOMER_ID	CUSTOMER0
FIRST_NAME	Kevin
LAST_NAME	Smith
EMAIL_ADDRESS	Kevin@aol.com
[-] ADDRESS	
[-] ADDRESS	
[-] CREDITRATING	
[-] CUSTOMER_PROFILE	

Customer profile data is shown. For the CUSTOMER element, values are shown for CUSTOMER_ID, FIRST_NAME, LAST_NAME and EMAIL_ADDRESS.

To test a simple update, click the Edit button, edit some data in the result, then click Submit. When you test the read function again, the results show the change.

You can also test an SDO update by submitting a datagraph with a change summary (see Test an Update Procedure).

4.1.1.6 See Also

For more information, see the following sources:

Concepts:

- [Understanding Update Maps](#)

How tos:

- [Create a Return Type](#)
- [Example: How to Create Your First Data Services](#)
- [Add a Where Clause to a Query](#)

Reference:

- [XQuery Source of a Logical Entity Service](#)

Other Source:

- [Introduction to Service Data Objects \(ibm.com\)](#)

4.1.2 Data Service Keys

This topic describes what data service keys are and how they are used.

- [Overview](#)
- [Parts of a Key](#)
- [Composite Keys](#)
- [See Also](#)

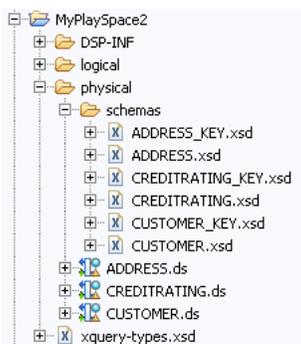
4.1.2.1 Overview

You are probably familiar with the concept of keys from relational databases, where a key is a set of one or more columns whose combined values are unique among all occurrences in a table.

When you create a physical data service, Oracle Data Service Integrator computes keys by introspecting the physical data sources. A physical data service key can have one or more fields, which are elements taken from the service's return type. Tangibly, a key is defined as an XML schema in an XSD file.

You can see the physical data service keys in your dataspace project in Eclipse for WebLogic. They appear in schema files with names such as:

```
datasource_KEY.xsd
```

Figure 4–11 Physical Data Service Keys in Eclipse for WebLogic

In the Eclipse for WebLogic, six keys are shown in the schemas folder. The key extension is .xsd.

In the generated XSD file, a key for a physical data service looks something like this.

Example 4–3 Key for the CUSTOMER Table

```
<?xml version="1.0" encoding="UTF-8" ?>
<xs:schema targetNamespace="ld:physical/CUSTOMER"
xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CUSTOMER_KEY">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="CUSTOMER_ID" type="xs:string"/>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

In this case, CUSTOMER_ID is the primary key in a relational table named CUSTOMER.

In a logical data service, a key also uniquely defines a data record. However, the data in the record can originate from multiple data sources of different types and can have a structure unlike the underlying physical data sources.

For a logical entity service, you must create the key. You can choose one of these options:

- Have Oracle Data Service Integrator generate the key based on the service's primary read function. Oracle Data Service Integrator generates a minimal key.
- Select the fields that make up the key. The elements that comprise the key must have a cardinality of 0 or 1 in the service's return type (with `maxOccurs="1"` or `maxOccurs="0"`, but not `maxOccurs="unbounded"`).

4.1.2.2 Parts of a Key

Suppose a logical service has a nested return type where a parent element with single cardinality can have multiple child elements, say one CUSTOMER element with many CUSTOMER_ORDER child elements.

Example 4–4 A Nested Return Type with a One-to-Many Relationship

```
<?xml version="1.0" encoding="UTF-8" ?>
<xs:schema targetNamespace="ld:logical/CustomersAndOrders"
```

```

xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CustomersAndOrders">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="CUSTOMER">
          <xs:complexType>
            <xs:sequence>
              <xs:element name="CUSTOMER_ID" type="xs:string"/>
              <xs:element name="FIRST_NAME" type="xs:string"/>
              <xs:element name="LAST_NAME" type="xs:string"/>
              <xs:element name="SSN" type="xs:string" minOccurs="0"/>
              <xs:element name="CUSTOMER_ORDER" maxOccurs="unbounded">
                <xs:complexType>
                  <xs:sequence>
                    <xs:element name="ORDER_ID" type="xs:string"/>
                    <xs:element name="C_ID" type="xs:string"/>
                  </xs:sequence>
                </xs:complexType>
              </xs:element>
            </xs:sequence>
          </xs:complexType>
        </xs:element>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>

```

This is the key that Oracle Data Service Integrator auto-generates from this return type, from the unique CUSTOMER_ID field:

Example 4-5 An Auto-Generated Simple Key

```

<?xml version="1.0" encoding="UTF-8"?>
<xs:schema targetNamespace="ld:logical/CustomerOrder" xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CustomersAndOrders_KEY">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="CUSTOMER_ID" type="xs:string"/>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>

```

If you choose to select the key fields, you need to use a unique field or fields with single cardinality. You can choose CUSTOMER_ID or SSN, or both. You cannot define the key on ORDER_ID or C_ID, because they belong to the CUSTOMER_ORDER element, which has multiple cardinality.

If you choose SSN, the key schema file looks like this.

Example 4-6 A Manually Selected Key

```

<?xml version="1.0" encoding="UTF-8"?>
<xs:schema targetNamespace="ld:logical/CustomerOrder" xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CustomersAndOrders_KEY">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="SSN" maxOccurs="1" minOccurs="0" type="xs:string"/>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>

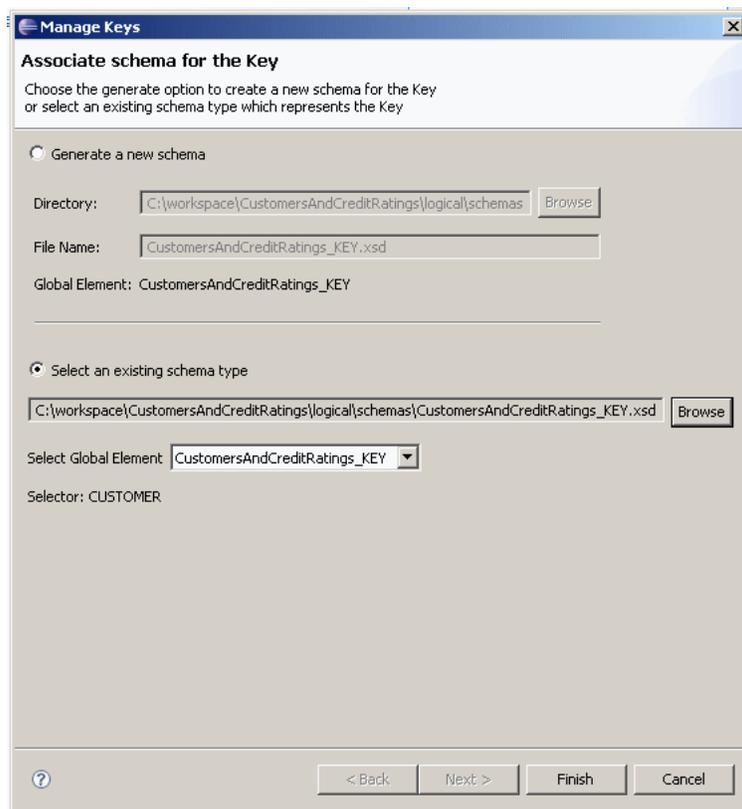
```

```
</xs:element>
</xs:schema>
```

A data service key has distinct parts:

- A selector. A key selector identifies a collection of data records. A key's selector is the element that contains the key field in the service's return type. You can see a key's selector in the Manage Key dialog when you create the key (below, it's the CUSTOMER element):

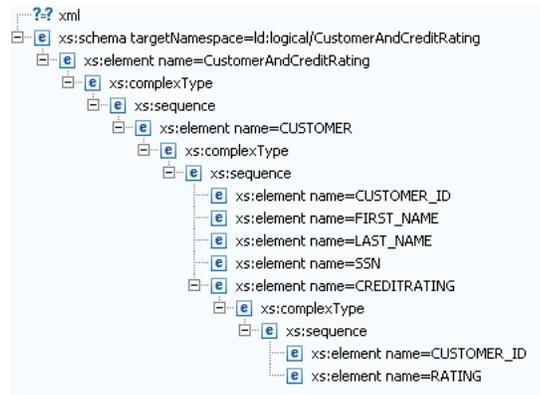
Figure 4-12 Associate Schema for the Key



The Manage Keys dialog enables you to choose the generate option to create a new schema for the key, or select an existing schema type which represents the key. There are two options: Generate a new schema and Select an existing schema. Under the Generate a new schema option, browse to select a directory and provide a file name. The Global Element is already selected. Under the Select an existing schema type option, browse to select the schema type. There is a pulldown menu that lets you select the global element. The Selector is already selected.

You can see that the CUSTOMER element is the root element of the return type:

Figure 4–13 Return Type



A return type tree is shown.

- The key fields. The fields that make up the key uniquely identify an element in the collection. For example, one customer identified by a CUSTOMER_ID value. Within Oracle Data Service Integrator, a key field is stored as a path which must not contain any repeating elements. Therefore, you cannot use elements with multiple cardinality in keys.

4.1.2.3 Composite Keys

With a logical service, a key can also be a composite key of multiple elements, as long as the elements have single cardinality in the return type. This is especially easy with a flat return type.

Example 4–7 A Flat, Non-Nested Return Type

```
<?xml version="1.0" encoding="UTF-8" ?>
<xs:schema targetNamespace="ld:logical/MyFlatOne" xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CUSTOMERORDER">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="CUSTOMER_ID" type="xs:string"/>
        <xs:element name="FIRST_NAME" type="xs:string"/>
        <xs:element name="LAST_NAME" type="xs:string"/>
        <xs:element name="EMAIL_ADDRESS" type="xs:string"/>
        <xs:element name="ORDER_ID" type="xs:string"/>
        <xs:element name="ORDER_DT" type="xs:date"/>
        <xs:element name="TOTAL_ORDER_AMT" type="xs:decimal"/>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

Oracle Data Service Integrator auto-generates a composite key using the key fields from the underlying physical data sources (in this example, CUSTOMER_ID and ORDER_ID). The composite key generated from this return type is shown below.

Example 4–8 An Auto-Generated Composite Key

```
<?xml version="1.0" encoding="UTF-8" ?>
<xs:schema targetNamespace="ld:logical/MyFlatOne" xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="MyFlatOne_KEY">
    <xs:complexType>
```

```

<xs:sequence>
  <xs:element name="CUSTOMER_ID" type="xs:string"/>
  <xs:element name="ORDER_ID" type="xs:string"/>
</xs:sequence>
</xs:complexType>
</xs:element>
</xs:schema>

```

This key allows you to identify a unique combination of **Customer** and **Order**, that is, one order for one customer.

4.1.2.4 See Also

For more information, refer to the following sources:

4.1.2.4.1 How Tos

- Create Logical Data Service Keys

4.1.3 XML Types and Return Types

In entity data services there are two types of types:

- Return types
- XML types

XML types and return types are very closely related. In data service operations involving entity data services, XML types define the shape of the data service.

Physically XML Types are represented a global elements in XML schemas (XSD files.) In other words, the XML types represents in hierarchical form the shape of the data service.

A way to think of these two artifacts is to first consider the class and the instance of the class in such languages as Java.

XML types can be thought of as a class from which objects in the form of functions are created. In many cases the information needed by these functions is either:

- A subset of the overall XML types -- for example, a function that returns last name and address but not first name or social security number.
- In need of further specification -- for example, adjusting a query to list all orders inside each customer rather than to repeat customer information each time.

Note: Return and XML types can be see in action in the following example:

Creating Your First Data Services

4.1.3.1 Where XML Types are Used

Oracle Data Service Integrator uses XML types in its model diagrams, entity data services, query editor, update mapper, and metadata browser.

4.1.3.2 Where Return Types are Used

Return types are sometimes called target schemas.

Return types can be thought of as the backbone of both data services and data models. Programmatically, return types are the "r" in FLWR (for-let-where-return) queries.

Return types have the following main purposes:

- Provide a template for the mapping of data from a variety of data sources and, in the case of updates, back to those data sources.
- Help determine the arrangement of the XML document generated by the XQuery.

Return types describes the structure or shape of data that a query produces when it is run.

Note: In order to maintain the integrity of Oracle Data Service Integrator queries used by your application, it is important that the query return type match the XML type in the containing data service. Thus if you make changes in the return type, you should use the XQuery Editor's **Save and associate schema** command to make the data service's XML type consistent with query-level changes. Alternatively, create a new data service based on your return type. For details see [Creating a Simple Data Service Function](#).

4.2 How-to

This section describes the following topics:

- [How To Add a Read Function](#)
- [How To Add a Library Function or Procedure](#)
- [How To Create Logical Data Service Keys](#)
- [How To Declare a Security Resource in Eclipse for WebLogic](#)

4.2.1 How To Add a Read Function

This topic describes how to add a read function to a logical entity service.

- [Overview](#)
- [Create the Function in Eclipse for WebLogic](#)
- [See Also](#)

4.2.1.1 Overview

A read function in a logical entity service retrieves data from underlying data sources, either physical or logical, and returns XML elements in the shape of the service's return type. You can build a logical service without a read function. However, the service must have at least one read function, marked primary, to have an update map. Only one read function in a service can be primary.

A read function is associated with exactly one XML schema, which is the service's return type. The read function must return the return type, but cannot take any other actions or have any side effects.

When you create a primary read function visually in Eclipse for WebLogic, Oracle Data Service Integrator generates a pragma annotation and XQuery source. The pragma looks something like this:

```
(::pragma function <f:function kind="read" visibility="public" isPrimary="true"
xmlns:f="urn:annotations.ld.bea.com"/>::)
```

The initial XQuery source, before you map data types in Query Map view, shows that the read function returns an instance of the service's return type:

```
declare function tns:read() as element(tns:CustomerAndAddress)*{
  <tns:CustomerAndAddress>
    <CUSTOMER>
      <CUSTOMER_ID></CUSTOMER_ID>
      <FIRST_NAME></FIRST_NAME>
      <LAST_NAME></LAST_NAME>
      <SSN?></SSN>
      {
        <ADDRESS>
          <ADDR_ID></ADDR_ID>
          <FIRST_NAME></FIRST_NAME>
          <ZIPCODE></ZIPCODE>
          <COUNTRY></COUNTRY>
        </ADDRESS>
      }
    </CUSTOMER>
  </tns:CustomerAndAddress>
};
```

At this point, the return type has no values. The values are added after you map data sources to the return type in Query Map view:

```
declare function tns:read() as element(tns:CustomerAndAddress)*{
  for $CUSTOMER in cus1:CUSTOMER()
  return
    <tns:CustomerAndAddress>
      <CUSTOMER>
        <CUSTOMER_ID>{fn:data($CUSTOMER/CUSTOMER_ID)}</CUSTOMER_ID>
        <FIRST_NAME>{fn:data($CUSTOMER/FIRST_NAME)}</FIRST_NAME>
        <LAST_NAME>{fn:data($CUSTOMER/LAST_NAME)}</LAST_NAME>
        <SSN?>{fn:data($CUSTOMER/SSN)}</SSN>
        {
          for $ADDRESS in add:ADDRESS()
          return
            <ADDRESS>
              <ADDR_ID>{fn:data($ADDRESS/ADDR_ID)}</ADDR_ID>
              <CUSTOMER_ID>{fn:data($ADDRESS/CUSTOMER_ID)}</CUSTOMER_ID>
              <FIRST_NAME>{fn:data($ADDRESS/FIRST_NAME)}</FIRST_NAME>
              <ZIPCODE>{fn:data($ADDRESS/ZIPCODE)}</ZIPCODE>
              <COUNTRY>{fn:data($ADDRESS/COUNTRY)}</COUNTRY>
            </ADDRESS>
        }
      </CUSTOMER>
    </tns:CustomerAndAddress>
};
```

4.2.1.2 Create the Function in Eclipse for WebLogic

Follow these steps to create the function:

1. Create a logical entity service. See [Section 1.3, "Example: How to Create Your First Data Services"](#)
2. In the Overview tab, right-click at the left, right, or top, and choose **Add Operation**.
3. At **Visibility**, choose an access level.

Public means the procedure can be called from the same dataspace and from client APIs; protected, only from the same dataspace; private, only from the same data service.

4. At **Kind**, choose read.
5. Enter a name for the function.
6. At **Return Type**, click **Edit**.
7. Click **Complex Type**, and choose a schema file.
8. At **Kind**, choose element.
9. At **Occurrence**, choose **Zero** or **More**.
10. Select **Primary**, and click **OK**.

4.2.1.3 See Also

How Tos

- ["Create a Return Type"](#) on page 6-1
- [Test a Read Function and Simple Update](#)

Concepts

- ["Data Service Types and Functions"](#) on page 1-52

4.2.2 How To Add a Library Function or Procedure

This topic describes how to add a library function or procedure to a data service.

- [Overview](#)
- [Add the Function or Procedure](#)
- [Test in Eclipse for WebLogic](#)

4.2.2.1 Overview

Library functions and procedures are utility operations that you can add to any service, physical, logical, or library. Library functions and procedures:

- Have a kind of library
- Are not marked as primary or non-primary
- Have a visibility of public, protected or private

4.2.2.2 Add the Function or Procedure

The example in this section is a library function that casts a value from xs:integer to xs:string.

1. Open the service and click the **Overview** tab.
2. Right-click at the left, right, or top, and choose **Add Operation**.
3. Select a value at **Visibility** (public = call from anywhere; protected = from the same dataspace; private = from the same data service).
4. At **Kind**, choose `libraryFunction` or `libraryProcedure`.

Figure 4–14 Add Operation - Library Function

On the Add Operation dialog, there are three fields at the top: Visibility, Kind, and Name. Select visibility with the drop-down menu. Select Library Function in the Kind with the drop-down menu. Specify the name. The Return Type is pre-selected. An Edit button enables you to edit the return type. A Clear button lets you clear your edits. There is a Parameters table with two columns: Name and Type. There are five buttons: Add, Edit, Remove, Up, and Down. There are two options: Primary and Empty Function Body. Primary is selected.

5. Give your function or procedure a name.
6. At **Return Type**, click **Edit** and choose a simple or complex return type. Click **OK**.
7. At Parameters, click Edit. Enter a parameter name, and choose a simple or complex return type. Click **OK**.
8. Click **Empty Function Body**, then **OK**.
9. Click the **Source** tab.

Oracle Data Service Integrator has generated a pragma statement and an empty function or procedure body, like this:

```
(::pragma function <f:function kind="library" visibility="public"
isPrimary="false" xmlns:f="urn:annotations.ld.bea.com"/>::)
```

```
declare function cus2:integerToString($theInt as xs:positiveInteger) as
xs:string* {
    $var-bea:tbd
};
```

10. In the function body, delete `$var-bea:tbd` and add your own XQuery code, for example:

```
declare function cus2:integerToString($theInt as xs:positiveInteger) as
xs:string* {
};
```

4.2.2.3 Test in Eclipse for WebLogic

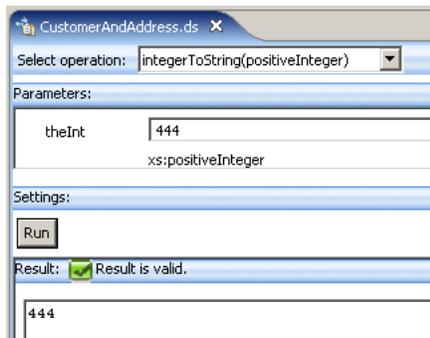
You can test the library function or procedure directly in Eclipse for WebLogic, before you use it from a client application.

1. Open the service, and click the **Test** tab.
2. At **Select Operation**, choose the library function or procedure you want to test.
3. Enter a value in the **Parameters** box.
4. (Optional) Expand **Settings** and enter new values for results, transactions, and authentication.
5. Click **Run**.

If the function or procedure works, you see valid results.

If not, you see an exception message that provides details, so that you can correct the error.

Figure 4–15 Result Validation in Eclipse for WebLogic



In Eclipse for WebLogic, the CustomerAndAddress data service is open. There is a Select operation drop-down menu; `integerToString(positiveInteger)` is selected. In the Parameters section, 444 is entered and `xs:positiveInteger` is displayed. In the Settings section, there is a Run button. In the Result section, the checkbox is checked and green, and Result is valid is displayed. The number 444 is displayed.

4.2.2.3.1 How Tos

- [How To Add a Read Function](#)
- [Add Update Map Procedures](#)

4.2.2.3.2 Reference

- [Data Service Types and Functions](#)

4.2.3 How To Create Logical Data Service Keys

This topic describes how to create a key for a logical data service.

- [Generate a Key](#)
- [Select Elements for a Key](#)
- [Select a Key Schema File](#)
- [View and Map a Key](#)

- [See Also](#)

A logical data service key uniquely identifies a data record the logical service defines. Because a logical service combines data from various physical and logical services, its key can combine or be different from the keys defined on underlying data sources.

For example, you might have a logical data service with a flat return type that combines data from two relational tables, **CUSTOMER** and **ORDER**. These tables have keys **CUSTOMER_ID** and **ORDER_ID**, respectively. In your logical data service, each data record is a unique combination of **Customer** and **Order**, so you create a composite key that combines **CUSTOMER_ID** and **ORDER_ID**.

Create procedures return a key to identify the data record that was inserted. Update and Delete procedures act on the data record the key identifies. A logical data service can have one key, although you can have multiple key schema files from which you select the key. You can have Oracle Data Service Integrator auto-generate the key, choose the elements you want in the key, or select an available schema (XSD) file to use for the key. The key definition requires specific knowledge of your data and the update map the service uses.

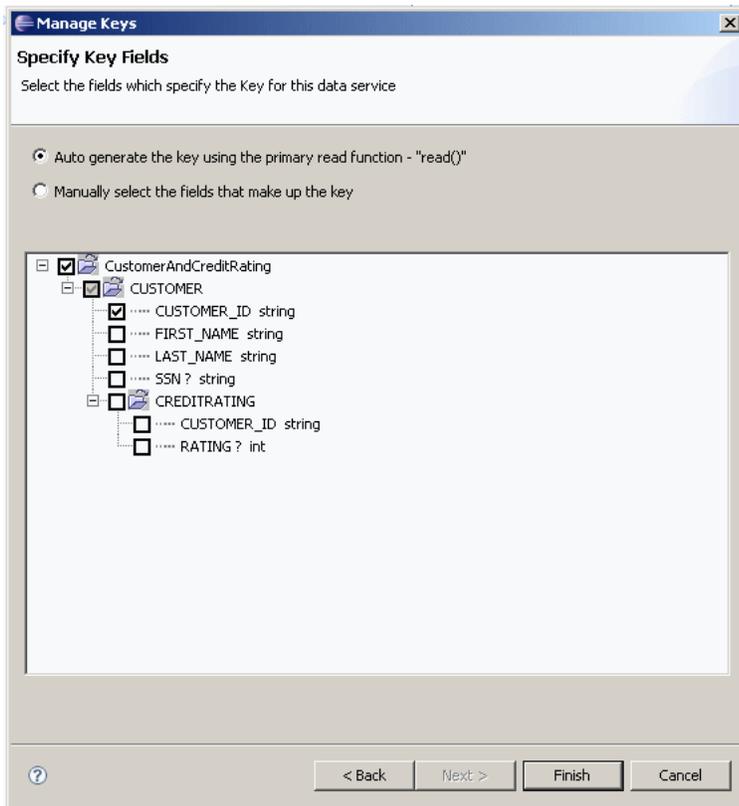
You can create a key for any logical data service that has a primary Read function. Once you create the key, you can view it in an update map and test it.

4.2.3.1 Generate a Key

To auto-generate the key:

1. Be sure the logical data service has a primary Read function.
2. Open the logical data service in Eclipse for WebLogic, and click the **Overview** tab.
3. Right-click in the service name bar, or at the left or right of the screen, and choose **Manage Key**.
4. Select **Generate a New Schema**.
5. Accept the default key name, or give your key a name ending in `.xsd`.
6. Click **Next**.
7. Select **Auto Generate the Key**, then click **Finish**.

Figure 4–16 Auto-Generating a Logical Data Service Key



The Manage Keys dialog lets you automatically select the fields that specify the Key for the data service. There are two choices: Auto generate the key using the primary function - "read()" (selected) and Manually select the fields that make up the key (deselected). There is a tree that displays the elements you could select manually if you choose manual selection.

You can now use the key as an argument or return type to an update map procedure, such as a Create, Update, or Delete procedure.

If you create a key, then delete it and create another one, you need to edit the signature of your Create procedure to return the new key:

Overview tab > right-click > Edit Signature

4.2.3.2 Select Elements for a Key

When you select elements for a key, you can add any element with single (1..1) or zero (0..1) cardinality, whether or not it is a key element in the underlying data source. An element with zero cardinality is optional and might contain null values, but you can use it as a key element. This allows you to create a wider variety of keys.

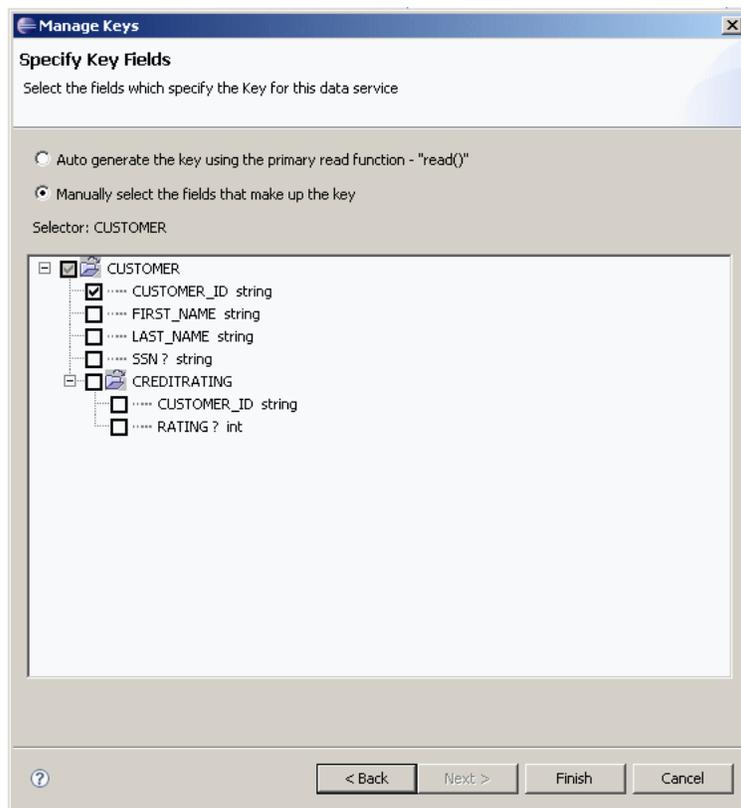
For example, you might have two data sources, one using a Social Security Number to identify records, and the other, a tax identification number. Your logical data service might have a return type that joins the two sources, so that a data record has either a social security number or a tax ID number. In the return type, both the social security number and the tax ID number are optional. The key can use either element to identify the record.

Note: You cannot select an element that has multiple (0..m or 1..m) cardinality to be part of a key.

To create a key with elements that you select:

1. Be sure the logical data service has a primary Read function.
2. Open the logical data service in Eclipse for WebLogic.
3. Click the **Overview** tab.
4. Right-click in the service name bar, or the left or right of the screen, and choose **Manage Keys**.
5. Click **Generate a New Schema**.
6. Give your key schema a name ending in `.xsd`.
7. Click **Manually select the fields that make up the key**.
8. Select the key fields you want, then click **Finish**.

Figure 4-17 *Selecting Elements for a Key*



The Manage Keys dialog lets you select the fields that specify the Key for the data service. There are two choices: Auto generate the key using the primary function - "read()" (deselected) and Manually select the fields that make up the key (selected). There is a tree that displays the elements you can choose.

4.2.3.3 Select a Key Schema File

You can also select an existing schema (XSD) file to use as the key:

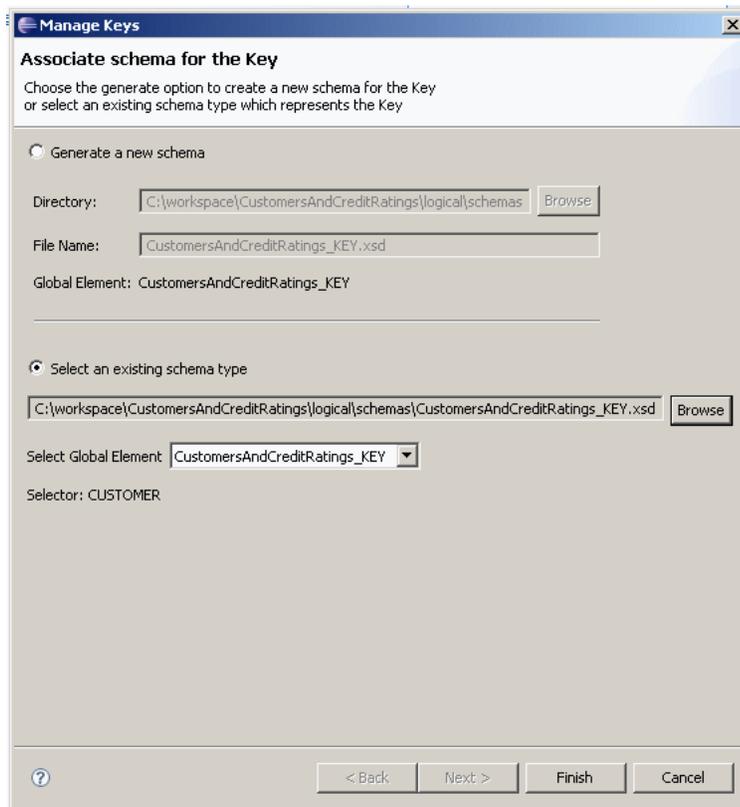
1. Be sure the logical data service has a primary Read function.
2. Open the logical data service in Eclipse for WebLogic.
3. Click the **Overview** tab.
4. Right-click in the service name bar, or the left or right of the screen, and choose **Manage Key**.
5. Click **Select an existing schema type**, then **Browse**.

The Manage Keys dialog shows you the key schema's global element and selector element.

6. Click **Finish**.

The schema in the Overview tab now displays a key icon next to the current key element or elements.

Figure 4–18 *Selecting the Key Schema*

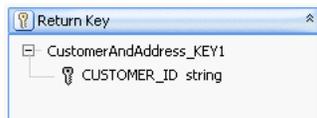


The Associate Schema for Key dialog enables you to choose the generate option to create a new schema for the key, or select an existing schema type which represents the key. There are two options: Generate a new schema (deselected) and Select an existing schema (selected). Under the Generate a new schema option, browse to select a directory and provide a file name. The Global Element is already selected. Under the Select an existing schema type option, browse to select the schema type. There is a pulldown menu that lets you select the global element. The Selector is already selected.

4.2.3.4 View and Map a Key

Once you create the key (whether by auto-generating, identifying key fields, or selecting a key schema file), you can see the key elements in the service's update map, at the lower left.

Figure 4–19 Viewing the Key in the Update Map

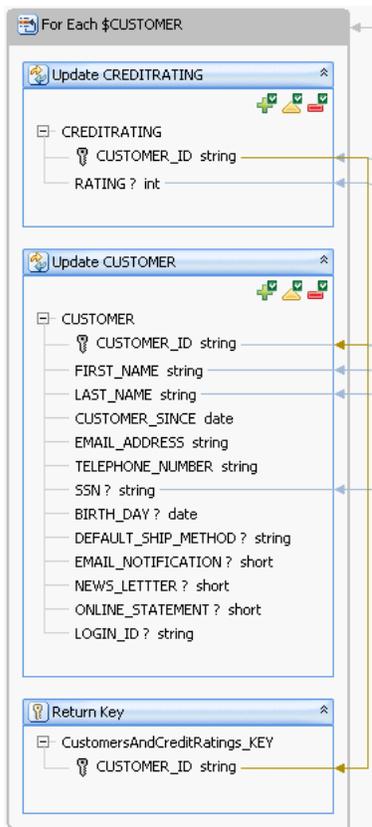


In the update map, the CUSTOMER_ID string is shown for CustomerAndAddress_KEY1.

The Return Key block represents the key elements a Create procedure returns when a new data record is added. In most cases, the key fields are automatically mapped to elements in the data sources on the left. If they are not mapped, you can add a mapping.

1. Locate the **Update** block on the left that contains the key element.
2. Drag from the key element in the **Update** block to the key element in the **Return Key** block.

Figure 4–20 Mapping a Key Element from an Update Block to the Return Key



The update map shows how the CUSTOMER_ID return key is mapped in the Update CUSTOMER block and the Update CREDITRATING block.

Note: Map the key element from an update block on the left, not from the return type on the right. If you map the key from the return type on the right, you allow the key value to be updated from data a user enters.

Once the key element is mapped, you can test it (preferably using sample data):

1. Click the **Test** tab.
2. At **Select operation**, choose one of the service's **Create** procedures.
3. Enter data in the XML template in the **Parameters** box.
4. Click **Run**.

The key value is returned in the Result box:

Figure 4–21 Key Value Return



In the Result box, for key CUSTOMER_ID, value CUSTOMER44 is returned.

You can also view the key schema file by locating the key in the Project Explorer, right-clicking, and choosing an XML editor to open the file. A key schema looks like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema targetNamespace="ld:logical/CustomerOrder" xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CustomersAndOrders_KEY">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="CUSTOMER_ID" type="xs:string"/>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

In the key schema, all elements must be in the same namespace as the root element. In the previous example, the namespace of the root element is:

ld:logical/CustomerOrder

A key schema cannot contain elements in different namespaces.

Note: If you have key schema files from a previous version of Oracle Data Service Integrator that you want to reuse, be sure that all elements within the schema are in the same namespace.

4.2.3.5 See Also

For more information, see the following resource:

- ["Data Service Keys"](#) on page 4-10

4.2.4 How To Declare a Security Resource in Eclipse for WebLogic

This topic describes how to add a security resource to a data service, so that the service returns data only if the caller has proper access.

- [Choose a Technique](#)
- [Create the Security Resource](#)
- [Use the Security Resource in XQuery](#)
- [Assign Security Resources](#)
- [Test Security](#)
- [See Also](#)

4.2.4.1 Choose a Technique

You can add a security resource to a data service in two ways:

- The first way is to use the Oracle Data Service Integrator Console to set elements and attributes that should be secured based on a security policy set by an administrator. This technique works in most cases for which you want to add a security policy.
- The other way, described here, is to create a custom security resource for an entity or library data service in Eclipse for WebLogic. The custom security resource is used directly in an XQuery expression to secure all or part of the service's return type. You can use the same custom security resource more than once in a single data service.

You can add a security resource to any data service, physical or logical, entity or library.

4.2.4.2 Create the Security Resource

You add a security resource to a logical entity service in Eclipse for WebLogic and then activate it using the Oracle Data Service Integrator Console.

Note: You can follow these steps on a physical or logical entity service. Be sure the service has a query map and a primary read function.

To create a security resource:

1. Open the service in Eclipse for WebLogic.
2. Make sure the **Properties** tab is displayed:
 Window > Show View > Properties
3. Click **Overview**, then **Properties**.
4. Expand the schema in the center. Locate the element you want to add the security resource to.

5. In the **Properties** tab, locate **Security Resources**.
6. Click the **Add New** field below it, then click the plus sign.
7. In the **Value** column, enter the name of the element you want to secure.

Figure 4–22 Creating a Security Resource

Property	Value
Creation Date	2007-10-22T13:36:48
[-] Security Resources	
Add New	
Delete Security Resources	
Security Resource (1)	CUSTOMER

Use the Properties tab to add the name of the element you want to secure. There are two columns: Property and Value. The Security Resource has the value of CUSTOMER.

Use just an element name (CUSTOMER), not a pathname (CUSTOMER_PROFILE/CUSTOMER) or a variable (\$CUSTOMER). You can use a simple element, a complex element, or the root element of the return type.

8. If needed, add more security resources and elements.
9. Click the **Source** tab.

The pragma statement at the top of the XQuery source file shows the new security resource:

```
(::pragma xds <x:xds targetType="cus:CustomerOrder"
  xmlns:x="urn:annotations.ld.oracle.com" xmlns:cus="ld:logical/CustomerOrder">
<creationDate>2007-10-22T13:36:48</creationDate>
<userDefinedView/>
<key name="DefaultKey" inferred="true" inferredSchema="true"
type="cus:CustomersAndOrders_KEY">
  <selector xpath="CUSTOMER"/>
</key>
<secureResources>
  <secureResource>CUSTOMER</secureResource>
</secureResources>
</x:xds>::-)
```

4.2.4.3 Use the Security Resource in XQuery

The next step is to add a condition to the return type so that it is returned only if the caller has access. To do this, make changes visually in the Query Map. You want to add a conditional statement to the service's primary read function, something like this:

```
declare function tns:read() as element(cus:CustomerOrder)*{
for $CUSTOMER in cus1:CUSTOMER()
return
  <cus:CustomerOrder>
    {
      if (add-authentication-expression-here) then
        <CUSTOMER>
          return type here ..
    }
  </CUSTOMER>
  else
```

```

    <CUSTOMER>{return nothing here}</CUSTOMER>
  }
</cus:CustomerOrder>

```

The following example shows how to create a security resource on an element in the return type, using the primary read function.

4.2.4.3.1 Create the If Condition

1. Click the **Query Map** tab.
2. At **Select Operation**, choose the primary read function.

Figure 4–23 Selecting the Read Operation

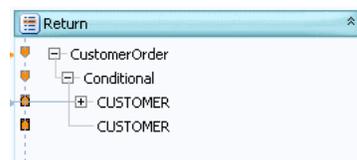


On the Query Map tab, you can select the read operation. Selection is shown with an arrow and a green circle.

3. In the return type, right-click the element for which you created a security resource in the **Properties** tab. Choose **Make Conditional**.

A node named **Conditional** is added to the return type.

Figure 4–24 Adding a Conditional Return Type



A conditional return type is added. Conditional is under Customer Order, and CUSTOMER is under Conditional.

4. Click the **Conditional** node.

You see the default conditional expression, (true), in the expression editor.

Figure 4–25 Expression Editor



The default conditional expression, (true), is shown in the expression editor.

5. Make sure the Design Palette is displayed (**Window > Show View > Design Palette**), then click it.
6. Expand:

XQuery Functions > Data Services Access Control Functions

7. In the mapping area, click the double arrow icon to open the expression editor.
8. Click the expression label in the editor.
9. Double-click (true), then delete it.

10. Drag the function `fn-bea:is-access-allowed` from the Design Palette to the editor.

```
fn-bea:is-access-allowed($label, $data_service)
```

11. For the `$label` argument, enter the name of your security resource as a string within quotes.

Use the same name you used in the Properties tab.

12. For the `$data_service` argument, enter the namespace-qualified name of your data service as a string within quotes:

```
fn-bea:is-access-allowed("CustomerOrder/CUSTOMER",
    "ld:logical/CustomersAndOrders.ds")
```

13. Click the Source tab, and check the read function. Make sure it has no errors.

Notice that the new expression is added to the if expression in the read function:

```
declare function tns:read() as element(cus:CustomerOrder)*{
  for $CUSTOMER in cus1:CUSTOMER()
  return
    <cus:CustomerOrder>
    {
      if (fn-bea:is-access-allowed("CUSTOMER",
        "ld:logical/CustomersAndOrders.ds")) then
        <CUSTOMER>
        ...
      </CUSTOMER>
    else
      <CUSTOMER>
      ...
    </CUSTOMER>
  }
}
```

14. Click **Save**.

You now need to define what is returned in the else clause.

4.2.4.3.2 Create the Else Condition

1. Click the Query Map tab.
2. In the return type, click the second conditional element.

Figure 4–26 Return Type



The second conditional element (CUSTOMER) is shown in the return type.

3. In the expression editor, enter "NA", and click **Save**.
4. Click the **Source** tab.

The read function now shows the return value for the else clause as the string "NA".

```
declare function tns:read() as element(cus:CustomerOrder)*{
  for $CUSTOMER in cus1:CUSTOMER()
  return
    <cus:CustomerOrder>
      {
        if (fn-bea:is-access-allowed("CUSTOMER",
          "ld:logical/CustomersAndOrders.ds")) then
          <CUSTOMER>
            <CUSTOMER_ID>{fn:data($CUSTOMER/CUSTOMER_ID)}</CUSTOMER_ID>
            <FIRST_NAME>{fn:data($CUSTOMER/FIRST_NAME)}</FIRST_NAME>
            <LAST_NAME>{fn:data($CUSTOMER/LAST_NAME)}</LAST_NAME>
            <SSN?>{fn:data($CUSTOMER/SSN)}</SSN>
            ...
          </CUSTOMER>
        else
          <CUSTOMER>{"NA"}</CUSTOMER>
        }
      }
    </cus:CustomerOrder>
}
```

4.2.4.4 Assign Security Resources

The next step is to use the Oracle Data Service Integrator console to create a security policy.

For more information, see [Securing Oracle Data Service Integrator Platform Resources](#)

All you need to do in the Oracle Data Service Integrator console is create a security policy. You have already created a custom security resource and added it to an XQuery function or procedure.

4.2.4.5 Test Security

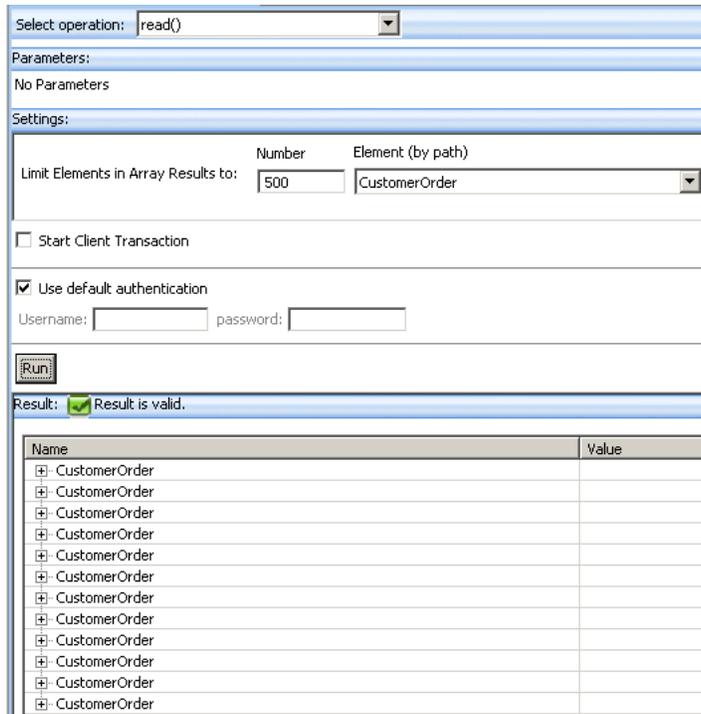
Once you establish security resources, you should test security in Test view.

To test a security resource:

1. Open the service in Eclipse for WebLogic.
2. Click the **Test** tab.
3. At **Select Operation**, choose the function you want to test.
4. Enter any parameters the function requires.
5. Expand **Settings** and enter the authentication credentials you want to use.
6. Click **Run**.

Check that the function returns either valid results if the authentication credential passes the security policy, or the string **NA** if it is not.

Figure 4–27 Testing the Read Function in Test View



Test view lets you select and test your operation. In the Test View, in the Select Operation pull-down menu, select your operation (read). In the Parameter section, No Parameters is displayed. In the Settings section, Limit Elements in Array Results is set to number 500, and Element (by path) is set to Customer Order, selected from the pull-down menu. Start Client Transaction is deselected. Use default authentication is selected. There is a Run button. Under Result, there is a green check mark and the words Result is Valid. Below, CustomerOrder appears many times.

4.2.4.6 See Also

- ["How To Add a Read Function"](#) on page 4-16
- [Securing Oracle Data Service Integrator Resources](#)

4.3 Examples

This section describes the following topics:

- [How to Create a Logical Data Service with a Group By Clause](#)
- [How To Create a Data Service with a Flat Return Type](#)

4.3.1 How to Create a Logical Data Service with a Group By Clause

This topic shows how to add a group by clause to a logical data service, using the Oracle extensions to XQuery.

- [Overview](#)
- [Design the Return Type Schema](#)
- [Create the Logical Data Service](#)

- [Create the Group By Node](#)
- [Create the For Node](#)
- [Add an Aggregate Function](#)
- [Test the Service](#)
- [See Also](#)

4.3.1.1 Overview

In relational data sources, a SQL `GROUP BY` statement is used with aggregate functions to group retrieved data by one or more columns. If you want to retrieve a list of distinct customers and the total amount of all orders each customer has placed from a relational data source, you might use a SQL statement like this:

```
SELECT CUSTOMER_ID, SUM(TOTAL_ORDER_AMOUNT) FROM ORDERS
      GROUP BY CUSTOMER_ID
```

The output produced groups all orders by customer and then totals the order amounts for each:

CUSTOMER_ID	TOTAL_OF_ALL_ORDERS
Customer0	9155.10
Customer1	5336.5
Customer2	11245.05
Customer3	1419.95

Oracle Data Service Integrator logical data services use XQuery 1.0 to query data. XQuery, as defined by the W3C standard, does not support group by clauses. However, Oracle Data Service Integrator has extended XQuery to allow a group by clause in an XQuery FLWOR statement:

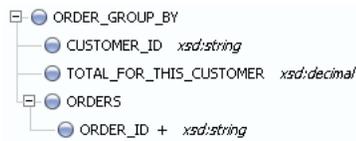
```
declare function tns:read() as element(ord1:ORDER_GROUP_BY)*{
for $CUSTOMER_ORDER in cus:CUSTOMER_ORDER()
group $CUSTOMER_ORDER as $CUSTOMER_ORDER_group by $CUSTOMER_ORDER/CUSTOMER_ID as $CUSTOMER_ID_group
return ...
```

You can add the XQuery group by statement to a logical data service visually in Eclipse for WebLogic. You should first make sure the service has a return type that supports the group by.

Suppose that after you retrieve all customer orders, group them by customer, and find the total amount of all orders each customer has placed, you also want a list of order IDs for each customer. You can design a logical data service to do this, doing part of the work in the mapping editor (in Eclipse for WebLogic) and part in the XQuery source.

4.3.1.2 Design the Return Type Schema

The return type schema needs an element to group by, such as a customer ID, and an element to hold an aggregate value, such as a sum or an average. The return type can also have a complex element that contains additional elements that provide information. This example provides the list of order IDs that are totalled for each customer, as one element with multiple cardinality within a complex element.

Figure 4–28 Return Type Schema for a Group By

The return type schema for a Group By provides a list of order IDs that are totalled for each customer, as one element with multiple cardinality within a complex element. Under `ORDER_GROUP_BY`, there are three elements: `CUSTOMER_ID`, `TOTAL_FOR_THIS_CUSTOMER`, and `ORDERS`. Under `ORDERS`, there is `ORDER_ID + xsd:string`.

If you want to design the schema top down using an XML editor, you can start with code like this and refactor it for your use case:

```

<?xml version="1.0" encoding="UTF-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema" targetNamespace="ld:logical/OrderGroupBy">
  <xs:element name="ORDER_GROUP_BY">
    <xs:complexType>
      <xs:element name="CUSTOMER_ID" type="xs:string" />
      <xs:element name="TOTAL_FOR_THIS_CUSTOMER" type="xs:decimal" />
      <xs:element name="ORDERS">
        <xs:complexType>
          <xs:sequence>
            <xs:element name="ORDER_ID" type="xs:string" maxOccurs="unbounded"
              form="unqualified" />
          </xs:sequence>
        </xs:complexType>
      </xs:element>
    </xs:complexType>
  </xs:element>
</xs:schema>

```

You can also create the return type bottom up, as you design the query map (see [Example: How to Create Your First Data Services](#)).

4.3.1.3 Create the Logical Data Service

Once you have defined the return type, create the logical data service and add the group by statement visually, using the mapping editor.

1. Create a new data space and import physical data sources (see ["Example: How to Create Your First Data Services"](#) on page 1-13).
2. Create a new logical data service.
3. Click **Overview**, right-click the name bar, choose Associate XML Type, and select the schema file for the return type.
4. Create a primary **Read** function.
5. Click **Query Map**. Drag the primary **Read** function from the relevant physical data source.

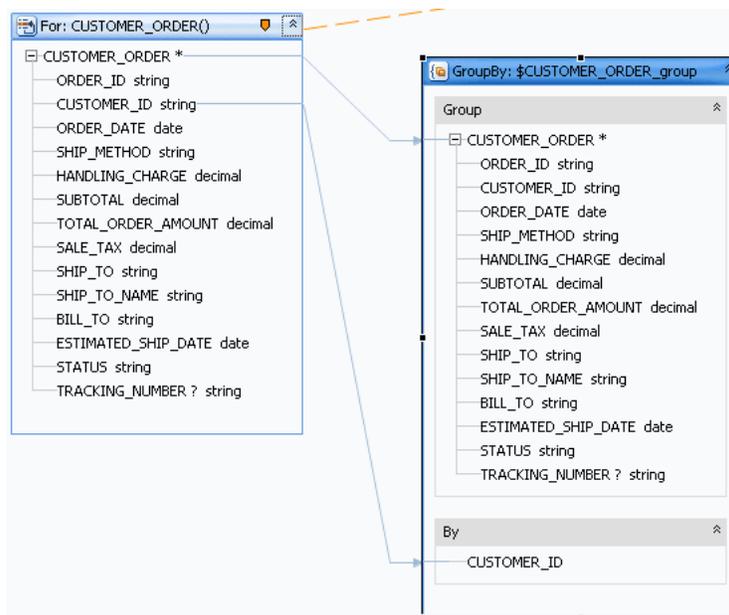
4.3.1.4 Create the Group By Node

Now create the group by node visually:

1. Right-click the element in the **For** block that you want to use as a grouping element, and select **Create Group By**.

A **Group By** node is created, and mappings are automatically drawn to it. The lower section of the **Group By** block shows the grouping element.

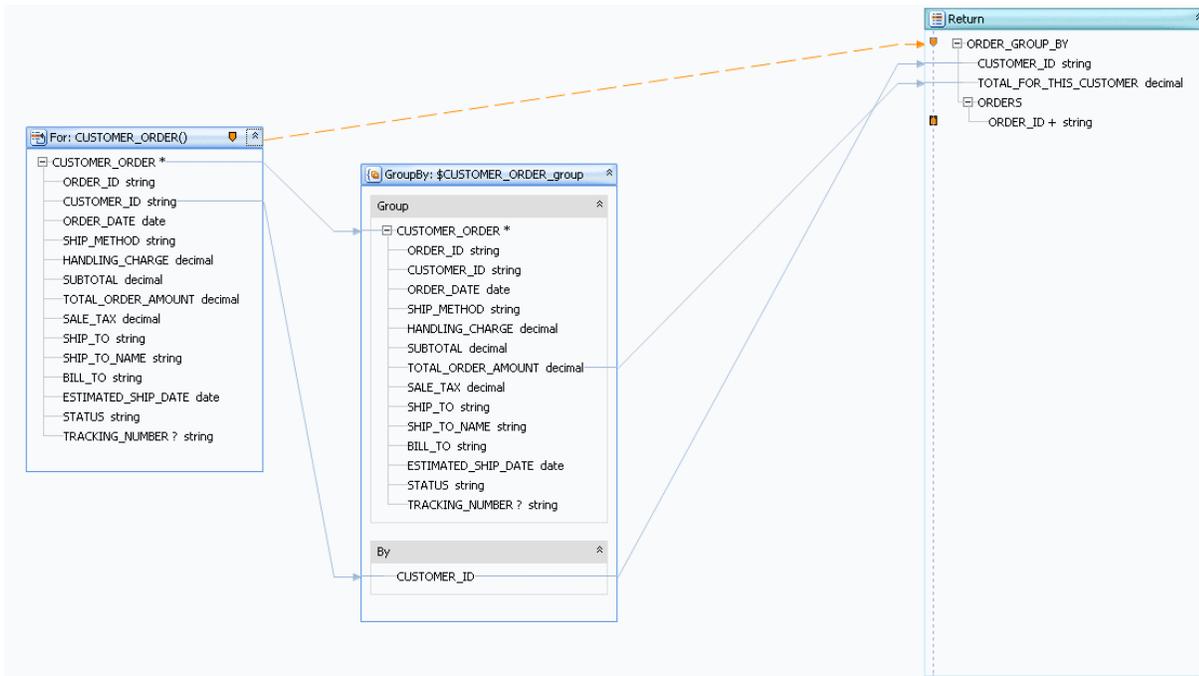
Figure 4–29 Add Group by Node



Mappings are automatically drawn when a Group By node is created. The lower section of the Group By block shows the grouping element: `CUSTOMER_ID`.

2. Drag a mapping from the grouping element in the **By** section of the Group By node to the grouping element in the return type (here, from `GroupBy CUSTOMER_ID` to `Return CUSTOMER_ID`).
3. Drag a mapping from the appropriate element in the top section of the Group By node to the aggregate element in the return type (here, from `Group By TOTAL_ORDER_AMOUNT` to `Return TOTAL_FOR_THIS_CUSTOMER`).

Figure 4–30 Mapping from the Group By Node



In the Group by block, TOTAL_ORDER_AMOUNT is mapped to TOTAL_FOR_THIS_CUSTOMER, and CUSTOMER_ID is mapped to CUSTOMER_ID in the Return type. Blue lines show the mapping.

4.3.1.5 Create the For Node

To map the information element, edit the XQuery code in the Source tab.

1. In the Source tab, add an XQuery for clause to the correct node in the primary Read function (here, the **ORDERS** node):

```

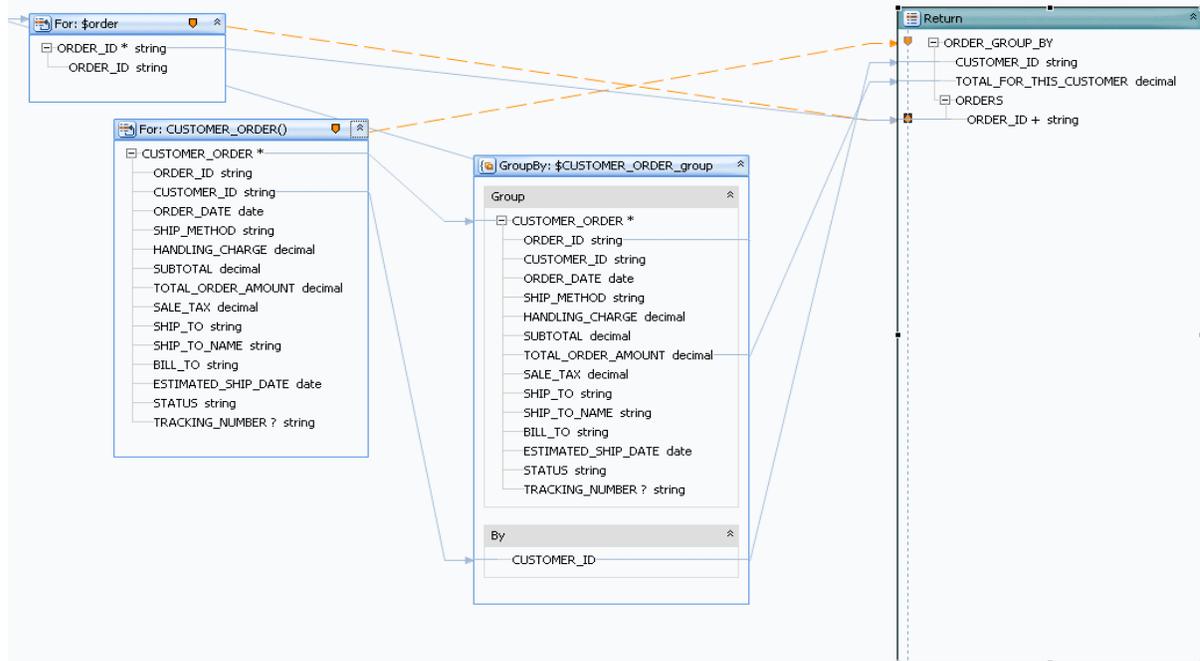
declare function tns:read() as element(ord1:ORDER_GROUP_BY) * {
  for $CUSTOMER_ORDER in cus:CUSTOMER_ORDER()
  group $CUSTOMER_ORDER as $CUSTOMER_ORDER_group by $CUSTOMER_ORDER/CUSTOMER_ID
  as $CUSTOMER_ID_group
  return
    <ord1:ORDER_GROUP_BY>
      <CUSTOMER_ID>{fn:data($CUSTOMER_ID_group)}</CUSTOMER_ID>
      <TOTAL_FOR_THIS_CUSTOMER>{fn:data($CUSTOMER_ORDER_group/
      TOTAL_ORDER_AMOUNT)}</TOTAL_FOR_THIS_CUSTOMER>
      <ORDERS> {
        for $order in $CUSTOMER_ORDER_group/ORDER_ID
        return
          <ORDER_ID>{fn:data($order)}</ORDER_ID>
      }
    </ORDERS>
  </ord1:ORDER_GROUP_BY>
};

```

The for statement declares a variable (here \$order) and then looks for an element (\$CUSTOMER_ORDER_group/ORDER_ID) in the first group the group by statement declares (CUSTOMER_ORDER_group). The for clause then returns the value of the element using the fn:data function.

- Click **Query Map**. Notice that a **For** node has been added.

Figure 4–31 Adding a For Node to a Group By



A For node for variable \$order has been created.

4.3.1.6 Add an Aggregate Function

Last, add an aggregate function to the aggregate element in the return type (here, `TOTAL_FOR_THIS_CUSTOMER`).

- In Query Map, click the aggregate element in the return type.

Notice that it uses the `fn:data` function, for example:

```
{fn:data($CUSTOMER_ORDER_group/TOTAL_ORDER_AMOUNT) }
```

- Click in the expression. Make sure the **Save** and **Cancel** icons are enabled.
- Click the Design Palette (**Window > Show View > Design Palette**).
- Expand **XQuery Functions**, then **Aggregate Functions**.
- Choose a function (here, the `fn:sum` function with one argument) and drag it to the expression editor. Leave the existing expression there.
- Edit the expression to use the existing expression as an argument to the aggregate function, for example:

```
{fn:sum( fn:data($CUSTOMER_ORDER_group/TOTAL_ORDER_AMOUNT) ) }
```

- Click **Save**.

4.3.1.7 Test the Service

The only way to test a logical data service with a group clause is to run the primary Read function in the **Test** tab. This type of data service does not have an update map,

so you cannot edit data and submit it or test an Update procedure. Likewise, you cannot test a Create procedure.

1. Click **Test**.
2. At **Select Operation**, choose the primary **Read** function.
3. Click **Run**.

You should see data grouped by the grouping element, with a result for the aggregate element, and containing a number of information elements.

Figure 4–32 Results of a Group By Statement

Name	Value
ORDER_GROUP_BY	
CUSTOMER_ID	CUSTOMER0
TOTAL_FOR_THIS_CUSTOMER	9155.1
ORDERS	
ORDER_ID	ORDER_10_0
ORDER_ID	ORDER_10_1
ORDER_ID	ORDER_10_2
ORDER_ID	ORDER_10_3
ORDER_ID	ORDER_10_4
ORDER_ID	ORDER_10_5
ORDER_ID	ORDER_10_6
ORDER_ID	ORDER_10_7
ORDER_ID	ORDER_10_8
ORDER_ID	ORDER_10_9
ORDER_ID	ORDER_10_10
ORDER_ID	ORDER_10_11
ORDER_ID	ORDER_10_12
ORDER_ID	ORDER_10_13

Data is grouped by grouping element. There is a result for the aggregate element, containing a number of information elements. Under ORDER_GROUP_BY, there are three elements: CUSTOMER_ID, TOTAL_FOR_THIS_CUSTOMER, and ORDERS. Under ORDERS, there are many ORDER_IDs with their corresponding values.

4.3.1.8 See Also

Examples

- [Example: How to Create Your First Data Services](#)

Other Resources

- [W3C XQuery Language Specification](#)
- [Extending XQuery for Grouping, Duplicate Elimination, and Outer Joins](#)

4.3.2 How To Create a Data Service with a Flat Return Type

This topic shows you how to create an update map from a logical data service with a flat, non-nested return type, using the sample database that ships with Oracle Data Service Integrator.

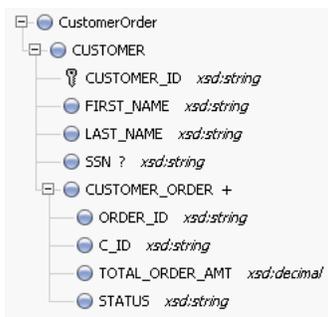
- [Overview](#)
- [Create a Dataspace Project](#)
- [Create the Return Type](#)
- [Create Physical Data Services](#)
- [Create a Logical Data Service](#)

- [Create the Query Map](#)
- [See Also](#)

4.3.2.1 Overview

A return type can be non-nested, or flat, even if it joins two relational tables, where one table has a one-to-many relationship with the other table. An example is one customer in a CUSTOMER table with many **Orders** in an ORDERS table. One approach to the return type is to nest an **Orders** element of multiple cardinality beneath the **Customer** element.

Figure 4–33 A Nested Customer-and-Orders Schema



A nested customer and orders schema is shown. Under CustomerOrder, there are two entries: CUSTOMER and CUSTOMER_ORDER. Under CUSTOMER, there are four entries: CUSTOMER_ID (key), FIRST_NAME, LAST_NAME, and SSN. Under CUSTOMER_ORDER, there are four entries: ORDER_ID, C_ID, TOTAL_ORDER_AMT, and STATUS.

Because you can design a logical data service with any structure, regardless of the underlying data sources, it is just as valid to define a flat return type to model the relationship between **Customers** and **Orders**.

Figure 4–34 A Flat Customer-and-Orders Schema



A flat customer and orders schema is shown. Under CUSTOMERS_AND_ORDERS, there are seven entries: CUSTOMER_ID, FIRST_NAME, LAST_NAME, EMAIL_ADDRESS, ORDER_ID (key), ORDER_DT, and TOTAL_ORDER_AMT.

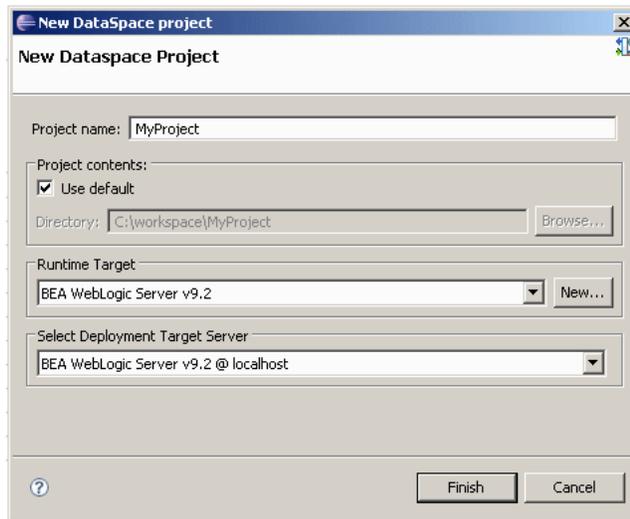
4.3.2.2 Create a Dataspace Project

First, create a new dataspace project to contain your physical and logical data services:

1. In Eclipse for WebLogic, choose **File > New > Dataspace Project**.

2. Enter a project name such as `FlatReturnType`, then click **Finish**.
3. Right-click the new dataspace project name, and choose **New > Folder**.
4. Create folders named `physical` and `logical`. Within `logical`, create a folder named `schemas`.
Using separate folders for physical and logical services helps separate the physical and logical integration layers.

Figure 4–35 New Dataspace Project



The New Dataspace Project dialog lets you create a new project. There is a field for a project name. Under Project Contents, there is a Use Default checkbox, which in this case is checked. The Directory field is inactive. There is a drop down menu for Runtime Target, as well as a New button. There is a drop-down menu for Select Deployment Target Service.

4.3.2.3 Create the Return Type

The return type the logical data service uses combines data from the `CUSTOMER` table and the `ORDERS` table. It has a non-nested XML structure, even though the data shows that customers and orders have a one-to-many relationship.

You can define the return type by creating an XML schema (XSD) file. In an XML editor, create a schema file like this one:

```
<?xml version="1.0" encoding="UTF-8" ?>
<xs:schema targetNamespace="ld:logical/FlatReturnType" xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CUSTOMERS_AND_ORDERS">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="CUSTOMER_ID" type="xs:string"/>
        <xs:element name="FIRST_NAME" type="xs:string"/>
        <xs:element name="LAST_NAME" type="xs:string"/>
        <xs:element name="EMAIL_ADDRESS" type="xs:string"/>
        <xs:element name="ORDER_ID" type="xs:string"/>
        <xs:element name="ORDER_DT" type="xs:date"/>
        <xs:element name="TOTAL_ORDER_AMT" type="xs:decimal"/>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

```
</xs:element>  
</xs:schema>
```

Be sure to:

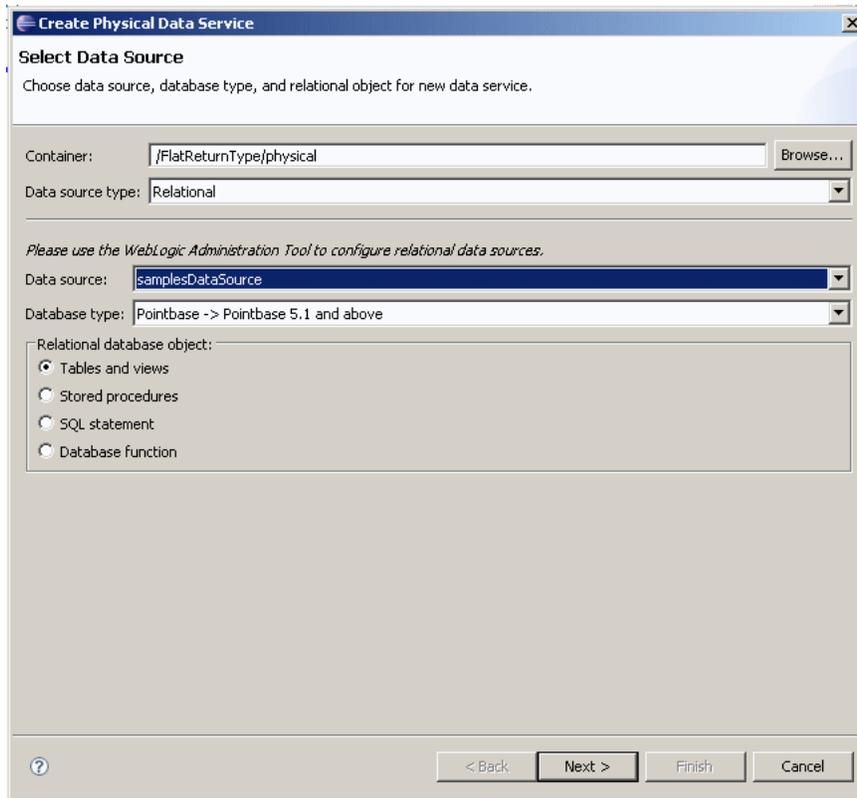
1. Define `targetNamespace` to make sense for your dataspace project.
Make sure you have only one top-level element of the name you choose (here, `CUSTOMERORDER`) in your target namespace. You can give the `targetNamespace` the same name as the dataspace project, but you are not required to.
2. Save the schema file in the `logical/schemas` folder within your dataspace project.

Note that the cardinality of all elements uses the default values, `minOccurs="1"` and `maxOccurs="1"`. Each customer has many orders, but there is only one combination of customer and order, so the cardinality of the order elements (`ORDER_ID`, `ORDER_DT`, and `TOTAL_ORDER_AMT`) is still 1.

4.3.2.4 Create Physical Data Services

Now, create physical data services based on the sample database or your own physical data sources.

1. In Project Explorer, right-click the physical folder in your dataspace project.
2. Choose **New > Physical Data Service**.
3. Choose **Relational** for Data source type and **dspSamplesDataSource** for Data source, then click **Next**.
4. Expand **RTLCUSTOMER** and select **CUSTOMER**.
5. Expand **RTLAPPLOMS** and select **CUSTOMER_ORDER**, then click **Next**.
6. Select **Public** for both **CUSTOMER** and **CUSTOMER_ORDER**, then click **Next**.
7. Click **Finish**.
8. When asked if you want to open the new data services, click **No**.

Figure 4–36 Adding Physical Data Services

The Select Data Service dialog lets you choose a data source, database type, and relational object for the new data service.

4.3.2.5 Create a Logical Data Service

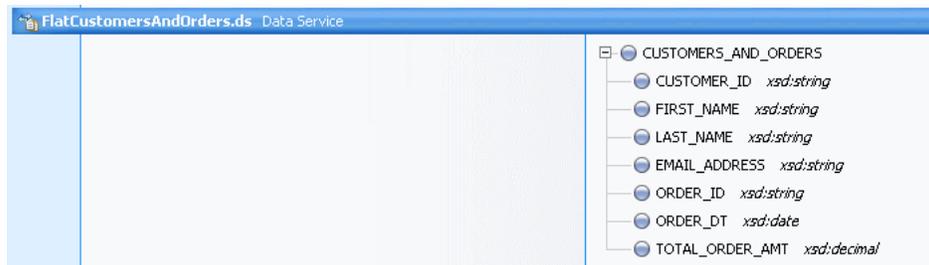
Now that you have physical data services and a schema for the return type, you can create the logical data service.

1. Right-click the logical folder, then choose **New > Logical Data Service**.
2. Enter a name for the service, such as `FlatCustomersAndOrders`.
3. Make sure **Entity Data Service** is selected, then click **Finish**.

Now associate a return type with the service:

1. Right-click in the **Overview** tab and choose **Associate XML Type**.
2. Select the schema and click **OK**.

Figure 4–37 A New Logical Data Service with a Return Type

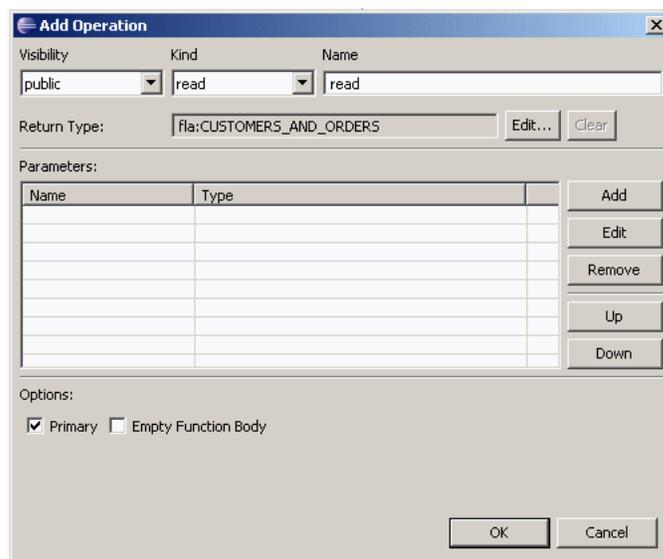


In the FlatCustomersAndOrders.ds Data Service display, a CUSTOMER_AND_ORDERS return type is shown. Below CUSTOMER_AND_ORDERS are CUSTOMER_ID, FIRST_NAME, LAST_NAME, EMAIL_ADDRESS, ORDER_ID, ORDER_DT, and TOTAL_ORDER_AMT.

You also need to define a primary Read function, in order to create both the query map and update map.

1. Right-click in the service name bar at the top, and choose **Add Operation**.
2. Make sure **Kind** is set to read, then enter a function name, such as read.
3. Make sure **Primary** is selected, then click **OK**.

Figure 4–38 Creating a Primary Read Function



The Add Operation dialog lets you create a primary read function. There are three fields at the top of the dialog. Visibility is a drop-down menu set to public. Kind is a drop-down menu set to read. Name is set to read. The Return Type is displayed. There is an Edit button that lets you edit the return type. A Clear button lets you clear your edits. There is a Parameters table with two columns: Name and Type. There are five buttons: Add, Edit, Remove, Up, and Down. The Primary option is selected and the Empty Function Body option is deselected.

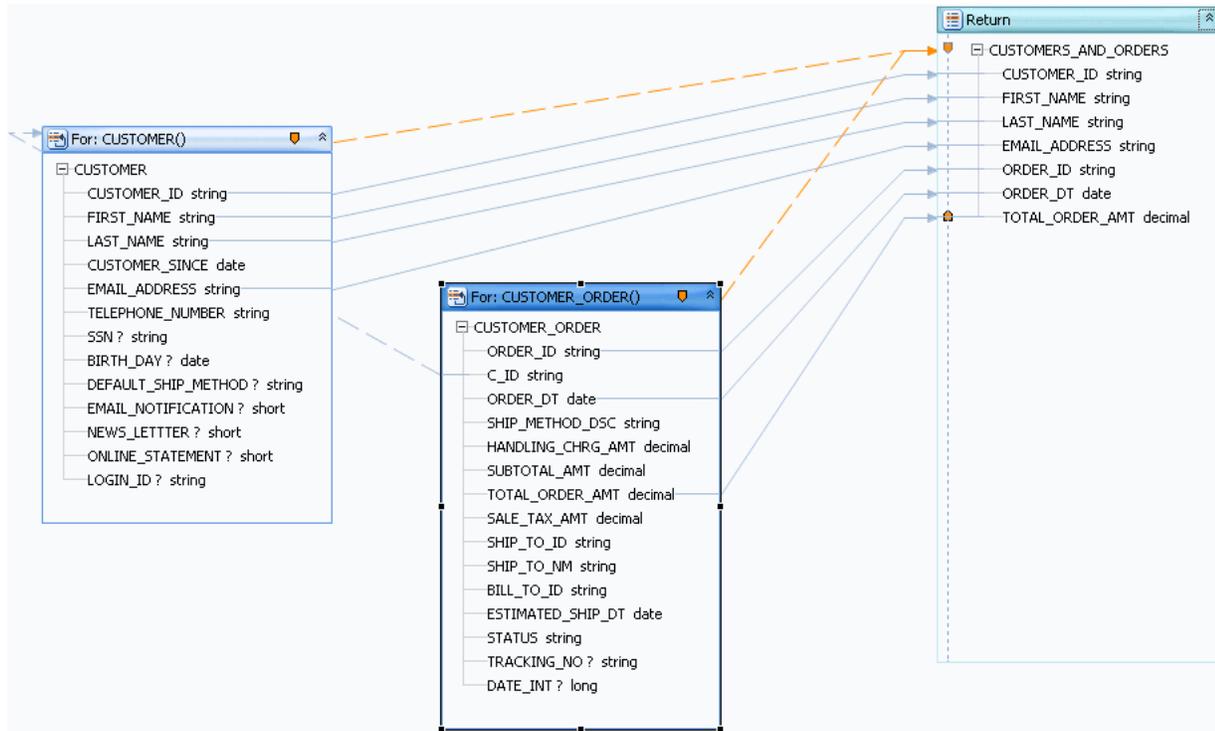
4.3.2.6 Create the Query Map

Now you need to create the query map visually in Eclipse for WebLogic, which in turn generates an update map.

1. Click the **Query Map** tab.
2. In **Project Explorer**, expand the physical data services `CUSTOMER.ds` and `CUSTOMER_ORDER.ds`.
3. Drag the **Read** function from each physical service to the mapping area.
 Notice that you cannot scope the `CUSTOMER_ORDER` block to a subtype in the return type, because the return type has no subtypes.
4. Drag mappings from the `CUSTOMER` block on the left to the return type for `CUSTOMER_ID`, `FIRST_NAME`, `LAST_NAME`, and `EMAIL_ADDRESS`.
5. Drag mappings from the `CUSTOMER_ORDER` block on the left to the return type for `ORDER_ID`, `ORDER_DT`, and `TOTAL_ORDER_AMT`.
6. In the For blocks, drag from `CUSTOMER/CUSTOMER_ID` to `CUSTOMER_ORDER/CUSTOMER_ID`.

This creates a join between the two data sources.

Figure 4–39 A Query Map with Mappings and a Join



A query map is shown with mappings and a join. Blue lines connect `CUSTOMER` and `CUSTOMER_ORDER` with the return type. The join between `CUSTOMER` and `CUSTOMER_ORDER` is shown with dotted orange lines.

If you click the Source tab and expand the Read function, you see XQuery code like this:

```

declare function tns:read() as element(fla:CUSTOMERS_AND_ORDERS)*{
for $CUSTOMER_ORDER in cus1:CUSTOMER_ORDER()
for $CUSTOMER in cus:CUSTOMER()
where $CUSTOMER/CUSTOMER_ID eq $CUSTOMER_ORDER/C_ID
return
  <fla:CUSTOMERS_AND_ORDERS>
    <CUSTOMER_ID>{fn:data($CUSTOMER/CUSTOMER_ID)}</CUSTOMER_ID>
    <FIRST_NAME>{fn:data($CUSTOMER/FIRST_NAME)}</FIRST_NAME>
    <LAST_NAME>{fn:data($CUSTOMER/LAST_NAME)}</LAST_NAME>
    <EMAIL_ADDRESS>{fn:data($CUSTOMER/EMAIL_ADDRESS)}</EMAIL_ADDRESS>
    <ORDER_ID>{fn:data($CUSTOMER_ORDER/ORDER_ID)}</ORDER_ID>
    <ORDER_DT>{fn:data($CUSTOMER_ORDER/ORDER_DT)}</ORDER_DT>
    <TOTAL_ORDER_AMT>{fn:data($CUSTOMER_ORDER/TOTAL_ORDER_AMT)}</TOTAL_ORDER_AMT>
  </fla:CUSTOMERS_AND_ORDERS>
};

```

Notice that the XQuery code has a for statement nested directly within another for statement. This creates an inner join between the two tables in SQL. To confirm the SQL that is created:

1. Click the **Test** tab.
2. At **Select** operation, make sure the primary **Read** function is selected.
3. Click **Run** (saving your data service as necessary).

You should see an XQuery FLWOR statement node. If you expand it, you should see a SQL query like this, showing an inner join:

```

SELECT t1."ORDER_DT" AS c1, t1."ORDER_ID" AS c2, t1."TOTAL_ORDER_AMT" AS c3,
       t2."CUSTOMER_ID" AS c4, t2."EMAIL_ADDRESS" AS c5, t2."FIRST_NAME" AS c6, t2."LAST_NAME" AS c7
FROM "RTLAPPLOMS"."CUSTOMER_ORDER" t1
JOIN "RTLCUSTOMER"."CUSTOMER" t2
ON (t2."CUSTOMER_ID" = t1."C_ID"

```

The inner join is created because the logical data service has a flat return type. When you mouse over the SQL query, you see this message:

Generated SQL query does not have a WHERE clause. This may cause the query to take longer to finish and use excessive memory resources.

4.3.2.7 See Also

- [Example: How to Create Your First Data Services](#)

4.4 Reference

The following example is for your reference.

4.4.1 XQuery Source of a Logical Entity Service

This topic shows sample XQuery source code for a logical entity data service.

- [Source Code](#)
- [See Also](#)

4.4.1.1 Source Code

```

xquery version "1.0" encoding "UTF-8";
(::pragma xds <x:xds targetType="cus:CUSTOMER_PROFILE"

```

```

xmlns:x="urn:annotations.ld.oracle.com" xmlns:cus="ld:logical/CustomerProfile">
  <creationDate>2007-10-05T10:29:01</creationDate>
  <userDefinedView/>
  <key name="DefaultKey" inferred="true" inferredSchema="true" type="cus:CustomerProfile_KEY">
    <selector xpath="CUSTOMER"/>
  </key>
</x:xds:::)

import schema namespace cus="ld:logical/CustomerProfile" at
  "ld:logical/schemas/CustomerProfile.xsd";

declare namespace cus1= "ld:physical/CUSTOMER";

declare namespace add= "ld:physical/ADDRESS";

declare namespace cre= "ld:physical/CREDITRATING";

import schema namespace cus2="ld:logical/CustomerProfile" at
  "ld:logical/schemas/CustomerProfile_KEY.xsd";

declare namespace tns="ld:logical/CustomerProfile";

declare function tns:stringToShort($theString) as xs:short {
  xs:short($theString)
};

(:::pragma function <f:function kind="read" visibility="public" isPrimary="true"
xmlns:f="urn:annotations.ld.oracle.com">
  <uiProperties>
    <component identifier="returnNode" minimized="false" x="842" y="11" w="244" h="601">
      <treeInfo id="0">
        <collapsedNodes>
          <collapsedNode>CUSTOMER_PROFILE\CUSTOMER</collapsedNode>
          <collapsedNode>CUSTOMER_PROFILE\CUSTOMER\ADDRESS</collapsedNode>
          <collapsedNode>CUSTOMER_PROFILE\CUSTOMER\CREDITRATING</collapsedNode>
        </collapsedNodes>
      </treeInfo>
    </component>
    <component identifier="CUSTOMER" x="44" y="56" h="300" w="219" minimized="false"/>
    <component identifier="ADDRESS" x="303" y="216" h="336" w="193" minimized="false"/>
    <component identifier="CREDITRATING" x="547" y="485" h="102" w="170" minimized="false"/>
  </uiProperties>
</f:function:::)

declare function tns:read() as element(tns:CUSTOMER_PROFILE)*{
for $CUSTOMER in cus1:CUSTOMER()
return
  <tns:CUSTOMER_PROFILE>
    <CUSTOMER>
      <CUSTOMER_ID>{fn:data($CUSTOMER/CUSTOMER_ID)}</CUSTOMER_ID>
      <FIRST_NAME>{fn:data($CUSTOMER/FIRST_NAME)}</FIRST_NAME>
      <LAST_NAME>{fn:data($CUSTOMER/LAST_NAME)}</LAST_NAME>
      <EMAIL_ADDRESS>{fn:data($CUSTOMER/EMAIL_ADDRESS)}</EMAIL_ADDRESS>
      {
        for $ADDRESS in add:ADDRESS()
        where $CUSTOMER/CUSTOMER_ID eq $ADDRESS/CUSTOMER_ID
        return
          <ADDRESS>
            <ADDR_ID>{fn:data($ADDRESS/ADDR_ID)}</ADDR_ID>
            <CUSTOMER_ID>{fn:data($ADDRESS/CUSTOMER_ID)}</CUSTOMER_ID>
      }
    </CUSTOMER>
  </tns:CUSTOMER_PROFILE>
}

```

```

        <STREET_ADDRESS1>{fn:data($ADDRESS/STREET_ADDRESS1)}</STREET_ADDRESS1>
        <CITY>{fn:data($ADDRESS/CITY)}</CITY>
        <STATE>{fn:data($ADDRESS/STATE)}</STATE>
        <ZIPCODE>{fn:data($ADDRESS/ZIPCODE)}</ZIPCODE>
        <COUNTRY>{fn:data($ADDRESS/COUNTRY)}</COUNTRY>
    </ADDRESS>
  }
  {
    for $CREDITRATING in cre:CREDITRATING()
    where $CUSTOMER/CUSTOMER_ID eq $CREDITRATING/CUSTOMER_ID
    return
    <CREDITRATING>
      <CUSTOMER_ID>{fn:data($CREDITRATING/CUSTOMER_ID)}</CUSTOMER_ID>
      <RATING?>{fn:data($CREDITRATING/RATING)}</RATING>
    </CREDITRATING>
  }
</CUSTOMER>
</tns:CUSTOMER_PROFILE>
};

(::pragma function <f:function kind="delete" visibility="public" isPrimary="true"
xmlns:f="urn:annotations.ld.oracle.com">
  <nonCacheable/>
  <implementation>
    <updateTemplate/>
  </implementation>
</f:function>::)

declare procedure tns:createCUSTOMER_PROFILE($arg as element(tns:CUSTOMER_PROFILE)* as
element(tns:CustomerProfile_KEY)* external;

(::pragma function <f:function kind="create" visibility="public" isPrimary="true"
xmlns:f="urn:annotations.ld.oracle.com">
  <nonCacheable/>
  <implementation>
    <updateTemplate/>
  </implementation>
</f:function>::)

(::pragma function <f:function kind="update" visibility="public" isPrimary="true"
xmlns:f="urn:annotations.ld.oracle.com">
  <nonCacheable/>
  <implementation>
    <updateTemplate/>
  </implementation>
</f:function>::)

declare procedure tns:updateCUSTOMER_PROFILE($arg as changed-element(tns:CUSTOMER_PROFILE)* as
empty() external;

(::pragma function <f:function kind="delete" visibility="public" isPrimary="false"
xmlns:f="urn:annotations.ld.oracle.com"/>::)

declare procedure tns:deleteByKey($arg0 as element(tns:CustomerProfile_KEY)){
  do return ();
};

```

4.4.1.2 See Also

Concepts

- [Building Logical Entity Data Services](#)

How Tos

- [Example: How to Create Your First Data Services](#)

4.5 Related Topics

For more information, refer to the following sources.

How Tos:

- [Example: How to Create Your First Data Services](#)
- [Create a Return Type](#)
- [Add a Complex Child Element to a Return Type](#)
- [Check Namespaces in Return Types](#)
- [Create Conditional Elements in Return Types](#)
- [Test a Read Function and Simple Update](#)
- [Test a Create or Delete Procedure](#)

Concepts:

- [Data Service Types and Functions](#)

Modeling Data Services Relationships

This chapter describes modeling data services relationships.

This chapter contains the following sections:

- [Section 5.1, "Relationship Between Data Services and Models"](#)
- [Section 5.2, "How to..."](#)
- [Section 5.3, "Reference"](#)

5.1 Relationship Between Data Services and Models

In large enterprises modeling is — or at least should be — an early task in developing a data services layer. By starting with a graphical representation of data resources it is easier to view data resources globally, leveraging existing information in interesting and useful ways. It is also easy to see opportunities for creating additional business logic in the form of logical services.

Model diagrams are quite flexible; they can be based on existing data services (and corresponding underlying data sources), planned data services, or a combination. You can also create and modify data services and data service XML types directly from the model.

Relationships can be surfaced through the **Relationship Modeler** in several ways:

- **Automatically.** By dragging two or more relational-based data services into a model diagram simultaneously. In such cases primary/foreign key relationships -- already available in the respective data service -- appear.
- **Graphically.** Through gestures you make in your model diagram or through the right-click menu*.*
- **Programmatically.** Through a data service **Source** editor.

Relationship functions allows data associated with one data service (such as **Customer**) to serve as a complex parameter for a related data service (such as **Orders**). Models can represent any combination of logical and physical data services.

A visual representation of a relationship between two data services can convey a considerable amount of information:

- **Cardinality.** Is the relationship one-to-zero (customers and promotional offers), one-to-one (customer and primary email), one-to-many (customers and orders), or many-to-many (customer orders and ordered items)?
- **Direction.** Arrows indicate possible navigation paths. Is there an originating entity associated with a subordinate entity (such as orders and order items) or is the relationship bidirectional (such as customers and orders)?

- **Roles.** A name matching the name of the adjacent data services navigation function (see below). Does the assigned relationship name capture the purpose of the navigation function it represents?

Many data service-related operations can be performed from the relationship modeler including:

- Modeling a high-level, visual view of data resources
- Viewing and adding to the relationships between data services
- Accessing or creating a data service
- Add operations to a data service
- Change the XML type (schema) associated with a data service

Navigation functions are visible as properties of each data service in the binary relationship. They can be fully inspected in the Source editor for each data service. Navigation functions also appear as mouse-over text over each endpoint of the relationship line.

By default, types shown in model diagrams are XML schema types, but you can change this to display native data source types in the case of physical data services.

Tip: For more information on data service modeling concepts see *Modeling and a Service-Oriented Architecture in the ODSI Concepts Guide*.

5.2 How to...

This section provides procedures for modeling in Oracle Data Service Integrator.

This section describes the following sections:

- [Section 5.2.1, "Create Your First Data Services Model"](#)
- [Section 5.2.2, "Work with Large Models"](#)
- [Section 5.2.3, "Generate a Relationship Modeler Report"](#)

5.2.1 Create Your First Data Services Model

Modeling Data Services

This section provides a basic overview of modeling in Oracle Data Service Integrator and a tutorial.

This section describes the following sections:

- [Section 5.2.1.1, "Introduction"](#)
- [Section 5.2.1.2, "Building a Simple Data Service Relationship Model"](#)
- [Section 5.2.1.3, "Setting Relationship Properties"](#)
- [Section 5.2.1.4, "Configuring Navigation Functions"](#)

5.2.1.1 Introduction

Using Oracle Data Service Integrator, you can create and maintain models of your enterprise data services. A model diagram is a graphical representation of a data model supported by Oracle Data Service Integrator.

Through data models you can:

- Model a high-level, visual view of data resources.
- View and extend relationships between data services.
- Access and create a data service.
- Add operations to a data service.
- Change the XML type (schema) associated with a data service.

In model diagrams, a relationship can be created by the gesture of drawing a line from one data service to another. In some cases (such as relational data services) relationships and the lines representing the relationship can be automatically inferred. In other cases, you need to create the relationship.

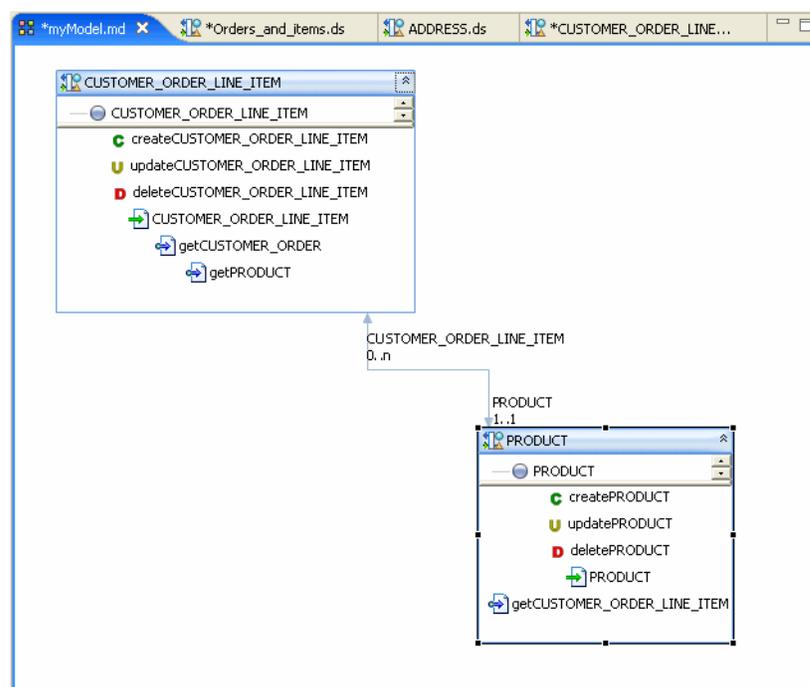
A visual representation of a relationship between two data services conveys a considerable amount of information:

- **Cardinality.** Is the relationship one-to-zero (customers and promotional offers), one-to-one (customer and primary email), one-to-many (customers and orders), or many-to-many (customer orders and ordered items)?
- **Direction.** Arrows indicate possible navigation paths. Is there an originating entity associated with a subordinate entity (such as orders and order items) or is the relationship bidirectional (such as customers and orders)?
- **Roles.** A name matching the name of the adjacent data services navigation function (see below). Does the assigned relationship name capture the purpose of the navigation function it represents?

Navigation functions are visible as properties of each data service in a binary relationship. Navigation functions also appear as mouse-over text over each endpoint of the relationship line.

Types shown in model diagrams are XML schema types.

Figure 5-1 Model Diagram of Physical Data Services



This figure shows a model diagram of physical data services

5.2.1.2 Building a Simple Data Service Relationship Model

You can create a sample data service relationship model by selecting a dataspace project and choosing:

File > New > Relationship Modeler

You can locate your model diagram anywhere in your project. Any legal filename can be used.

5.2.1.2.1 About the Data Services This example assumes that you are using the Oracle Data Service Integrator RTLApp as a data source.

The physical data services used in this sample are:

- CUSTOMER
- CUSTOMER_ORDER
- CUSTOMER_ORDER_LINE_ITEM

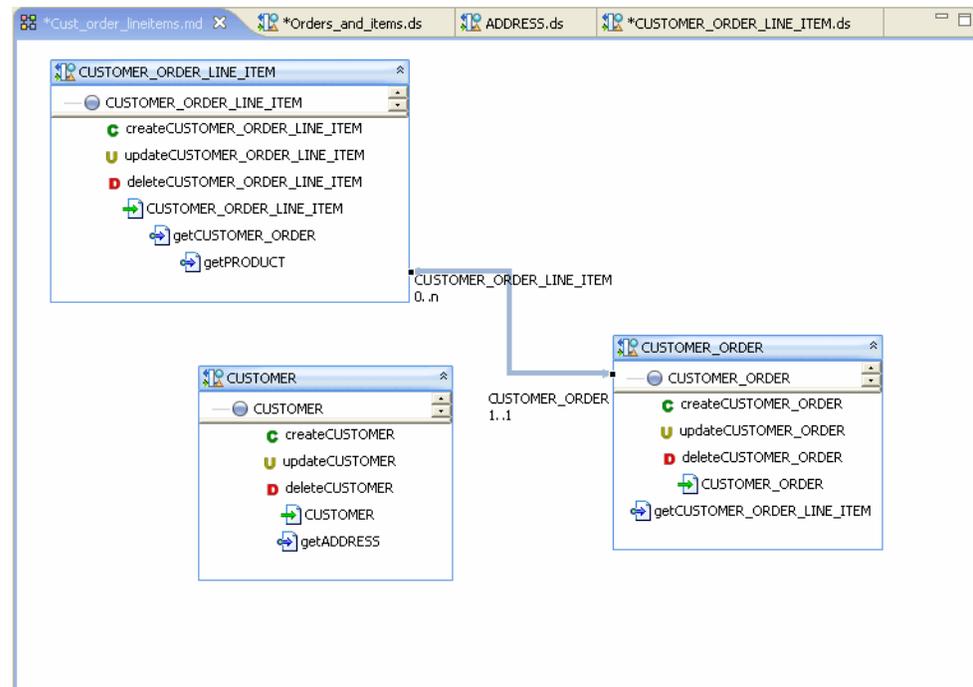
Note: See [Chapter 3, "Creating and Updating Physical Data Services"](#).

5.2.1.2.2 Adding Data Services to the Modeler You can add data services to your model using simple drag-and-drop from the **Project Explorer**. In the **Project Explorer** you can multi-select data services using either:

- Shift-click (contiguous services) or
- Control-click (individual services)

If you drag a set of data services into a model diagram, existing relationships to other data services in the model will be shown.

Figure 5-2 Populating the Relationship Modeler



Shows how to populate the relationship modeler.

Since the data services in this example are representations of relational sources, a several bidirectional relationships between **CUSTOMER_ORDER** and **CUSTOMER_ORDER_LINE_ITEMS** were inferred:

- The role named **CUSTOMER_ORDER** has a 1-to-1 relationship with **CUSTOMER_ORDER_LINE_ITEM**, meaning that a line item can only belong to one order.
- The role named **CUSTOMER_ORDER_LINE_ITEM** has a 0-to-n relationship with **CUSTOMER_ORDER**, meaning that there can be many line items associated with an order.

5.2.1.2.3 Creating an Additional Relationship A next step could be to create a relationship between **CUSTOMER** and **CUSTOMER_ORDER**.

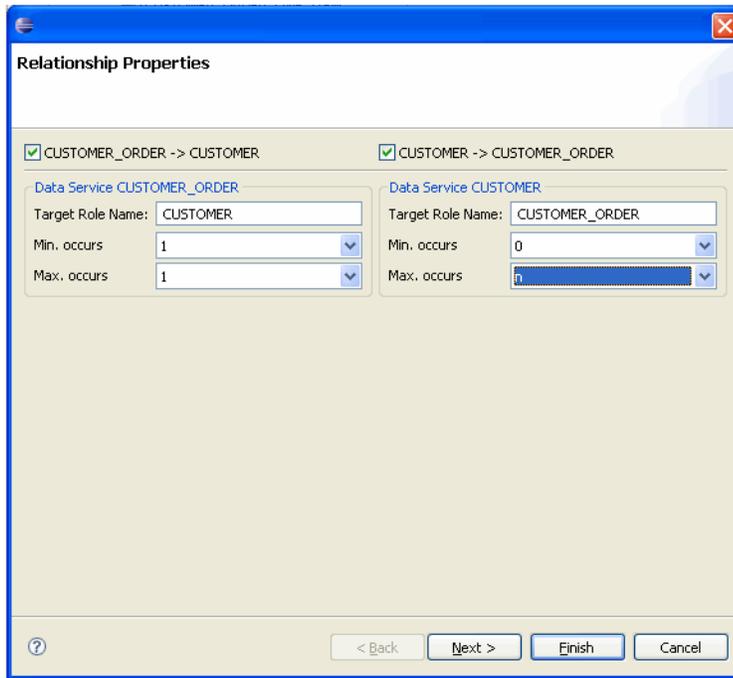
1. Right-click on **CUSTOMER_ORDER** node.
2. Select **Create Relationship to another data service**.
3. Select **CUSTOMER** as the target data service.
4. Click **OK**.

The **Relationship Properties** wizard appears.

5.2.1.3 Setting Relationship Properties

Relationship properties can be uni- or bi-directional.

Figure 5–3 Setting Relationship Options



Shows how to set properties in the Relationship Properties dialog.

Table 5–1 Relationship Properties Dialog Options

Option	Action	Comment/Reference
Set directionality	Select the directions to be supported in the relationship. The example is bidirectional so the default checked condition for the following relationships need not be changed: <ul style="list-style-type: none"> ■ CUSTOMER_ORDER -> CUSTOMER ■ CUSTOMER -> CUSTOMER_ORDER 	Creating relationships in a model automatically creates relationship functions between data services. Bi-directional settings mean that “get” functions for the related data service will be created on both sides of the relationship. By default, relationships are bidirectional.
Target Role Name	Enter the name of the role function. In the example, default names can be used: <ul style="list-style-type: none"> ■ CUSTOMER ■ CUSTOMER_ORDER 	By default the name will be based on the name of the related data service. It can be changed to any unique and legal name in your dataspace project.
Set maximum and minimum occurrences	Enter cardinality settings for the respective function. For the example the following settings are used: <ul style="list-style-type: none"> ■ CUSTOMER_ORDER -> CUSTOMER: 1-to-1 ■ CUSTOMER -> CUSTOMER_ORDER 0-to-n 	The minimum and maximum occurrence settings definite the nature of the relationship between the two services.

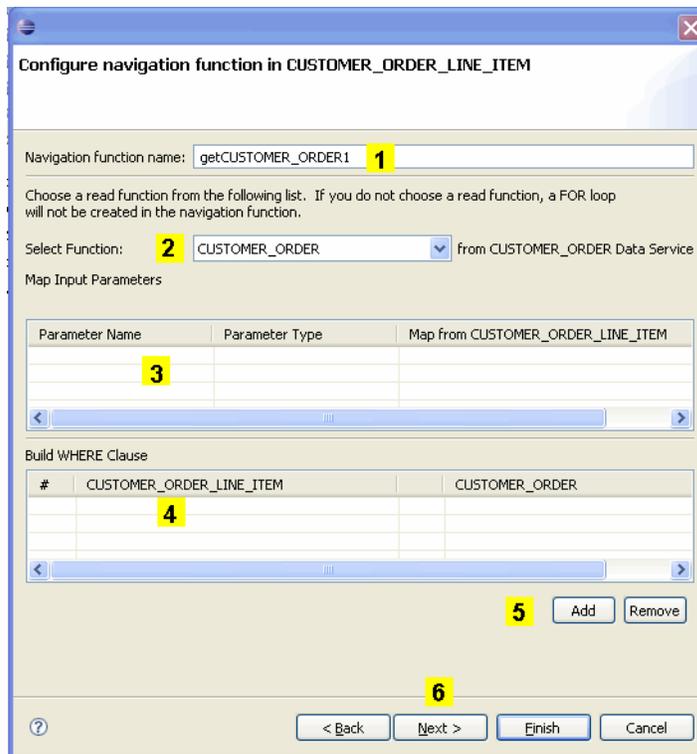
Click Next.

5.2.1.4 Configuring Navigation Functions

Each navigation function (one or two) being created also needs to be configured. Configuration includes:

- Setting a name for the navigation function
- Selecting a function from the newly related data service.
- Mapping input parameters
- Building a **WHERE** clause

Figure 5–4 *Configuring a Navigation Function*



Shows how to configure a navigation function in CUSTOMER_ORDER_L?INE_ITEM.

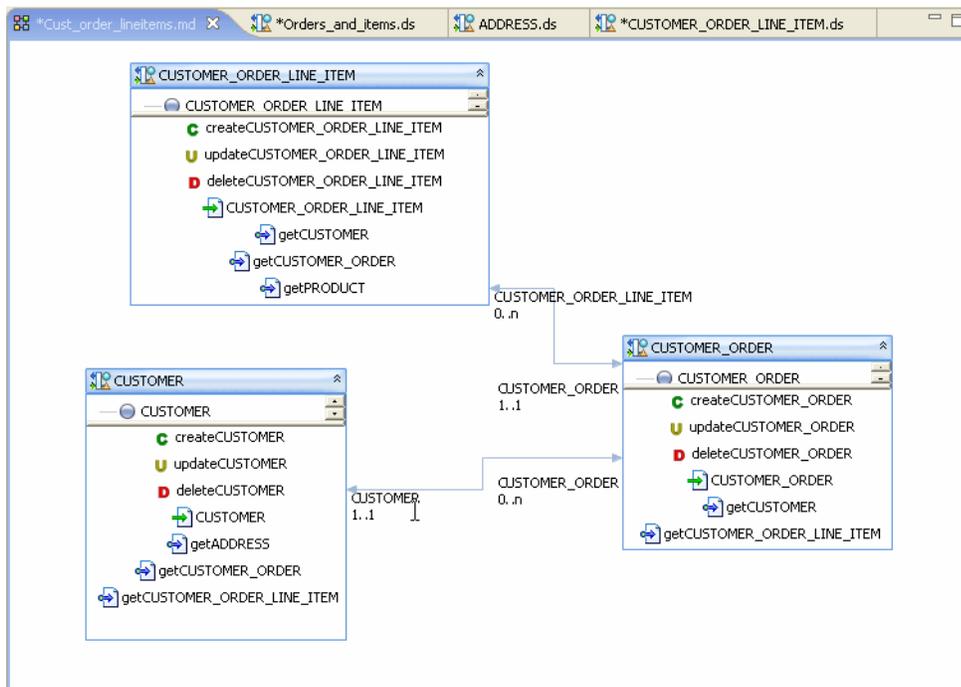
Table 5–2 *Specifying Relationship Wizard Function Name, Parameters, and Where Clauses*

Element	Purpose
Navigation function name	By default, the navigation function name is the name of the target data service with “get” prepended, as in “getCustomer.” If a function of that name exists, numbers will be appended to the function name as in getCustomer1. Note: When you invoke the Relationship wizard through a model diagram the opposite data service is determined by the gesture of drawing a line from one data service to another. In such cases the option of selecting a navigation function name is not present.
Related data service function	By default, the root function in the target data service is selected. However, you can select any available read function in the target data service.

Table 5–2 (Cont.) Specifying Relationship Wizard Function Name, Parameters, and Where Clauses

Element	Purpose
Map input parameters	If the related function has input parameters, the name and type of the available parameters appear. You can then use a pull-down menu to select an element from the target data service to map as the input parameter.
Build WHERE clause	Where clauses can be added to the function using pull-down menus that allow you to select join elements from each side of the relationship.
Add or Remove	Allows you to add additional where clauses or delete an identified where clause.
Next	When the relationship between data services is bidirectional clicking Next changes the focus to the second data service, where you can identify a navigation function name, parameters, and add where clauses for the second side of the relationship.

Figure 5–5 Customer-Order-Item Model



Shows the customer-order-item model.

5.2.2 Work with Large Models

How To Work with Large Model Diagrams

Model diagrams can hold any number of data services. The only limitation is that each data service must reside in the same dataspace.

Some tools are available in cases where very large models have been created.

5.2.2.1 Search

You can locate any data services in your model diagram.

1. Right-click in the white space (not on a data service representation) and select **Find Data Service**.
2. Type in the name of your data service using standard search options available in the dialog.

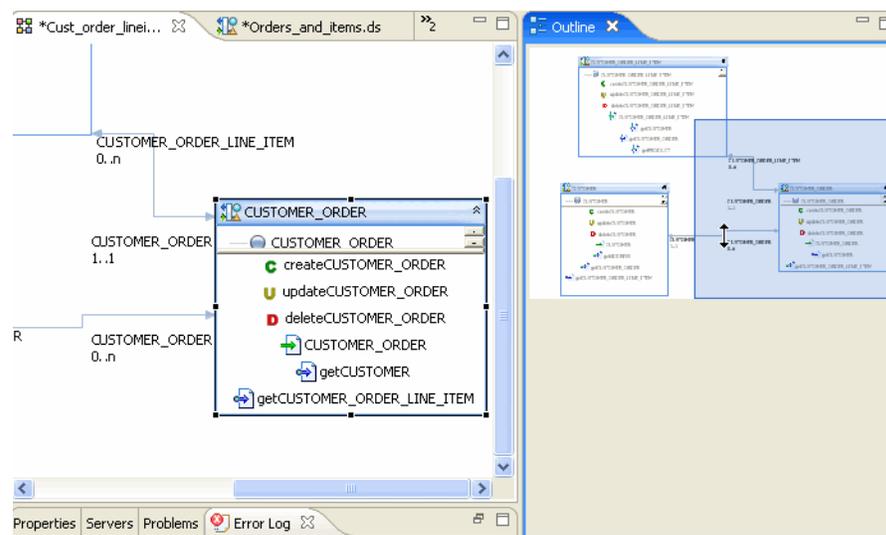
A dialog will appear containing a dropdown list of matching data services. The data service you select will be appear.

5.2.2.2 Outline Mode

For larger models you can use **Outline** view which will allow you to scroll through your model.

Window > Show View > Outline

Figure 5–6 Model Diagram Outline View



Shows the model diagram outline view.

5.2.3 Generate a Relationship Modeler Report

Generating Reports on Your Models

Both summary and detailed reports are available from the currently selected relationship model.

To create a report:

1. Right-click in any blank space in your model diagram.
2. Select:
 - Generated Report > Detailed or Generate Report > Summary**
3. Select a filename and location.
4. Click **Finish**.

Table 5–3 Summary and Detailed Report Categories Compared

Type	Description
Summary Report	<p>Provides general information related the model including:</p> <ul style="list-style-type: none"> ■ Location of each data service in the model ■ Type: logical or physical ■ Allows updates: true/false ■ Data source type ■ Data source name ■ Owner (if any) ■ Comment (if any) ■ Date created ■ Date last modified
Detailed Report	<p>A detailed model report contains all summary information listed above and, for each relationship between data services, the following additional information:</p> <ul style="list-style-type: none"> ■ Return type fully qualified name (the qname) ■ Details on each Read function including Return type, description, and comments ■ Details on the data service relationships including role name, target data service, minimum and maximum occurrences, opposite role name, navigation functions including Return type, description, comment and user-defined properties ■ Dependencies — a list of all dependent data services

When you choose the **Create a Model Report** right-click option you are asked to select a name for the HTML document that is generated. By default, the name of the summary report is:

```
<model_name>_md_summary.html
```

and the name of the detail report is:

```
<model_name>_md_detail.html
```

You can save the report to any location in your application, including to a new folder.

5.2.3.1 Model Report Format

The model report is in HTML format.

Note: Print your report from any browser or application that supports HTML printing.

5.3 Reference

This section describes the following sections:

- [Section 5.3.1, "Relationship Modeler Options"](#)
- [Section 5.3.2, "Model Diagram Rules"](#)
- [Section 5.3.3, "Notable Relationship Modeler Properties"](#)
- [Section 5.3.4, "Relationship Models in Source View"](#)

5.3.1 Relationship Modeler Options

This section describes some of the common operations you will use when working with the relationship modeler.

5.3.1.1 Model Right-click Menu Options

You can edit your model using a combination of right-click menu options and the model **Property Editor**. [Table 5-4](#) and [Table 5-5](#) describe right-click options based on the functional area of the model diagram that is in scope.

Table 5-4 Data Model: Notable Data Model Options

Command	Meaning
New Data Service	Allows you to create a new data service. After selecting a name and physical location for the data service, the service is created and placed on the diagram.
Find Data Service	Locates a data service within your model.
Select Router Type	Adjusts visual presentation of relationship lines based on the Manhattan model or shortest-path model.
Generate Report	Creates either a Summary or Detail report in an Eclipse HTML-based page. The report describes data services in your model, their bilateral relationships, and a description of each data service.

Table 5-5 Data Service: Notable Data Model Options

Command	Meaning
Open	Dialog allows you to select from a list of data services in the model diagram. As with drawing a line between two data services, this option brings up the Relationship wizard. (See Using the Relationship Wizard to Create Navigation Functions).
Create Relationship to Another Data Service	The Add Related command is available when one or several data services are selected in the model. Add Related lists data services that contain navigation functions referencing your currently selected data source. Click on the service you want to add and then repeat the process to add other available related services, if any.
Add Related Data Service	Adjusts visual presentation of relationship lines based on the Manhattan model or shortest-path model.
Remove from Diagram	Removes the selected data service from the model diagram. Alternatively, use the Delete key. Note: This operation does not affect the underlying data service.
Refactor	Provides for either safe delete or renaming of the currently selected data service. This is comparable to operations available for a data service from the Overview tab.
Associate XML Type	Provides a dialog where a different schema (XSD) file can be selected from the current project. Note: Changing a schema type for a data service can affect its functions as well as its relationships to other data services.
Manage Key...	Opens the Manage Key dialog box, allowing for modification of the key associated with the current data service.
Delete Key...	Deletes any key associated with the current data service.

Table 5–5 (Cont.) Data Service: Notable Data Model Options

Command	Meaning
Add Operation	Adds an operation (function or procedure) to the currently selected data service.
Show/Hide Native XML Types	Optionally displays/hides native types for elements representing physical objects associated with simple data types. Example: <code>VARCHAR(25)</code> .
Show Function Signatures	Displays/hides full read function signatures such as: <code>getAddress () as element (Address)</code>

Tip: Relationship lines connecting data services can be deleted by first selecting the line, then pressing **Delete**.

5.3.2 Model Diagram Rules

Rules Governing Model Diagrams

Model diagrams follow a set of rules:

- Each entity in the model has a title which is the local name of the data service (the fully-qualified name is visible as a mouse-over).
- Associated **Read** functions can be displayed, with or without signatures.
- Model diagrams do not “own” data services, but simply reference them. Multiple models can, without limit, contain representations of the same data service or relationships between data services.
- Models are not nested. That is, a model diagram cannot reference another.
- Multiple models can be defined and located anywhere in your project.
- Changes made to a model diagram can be reversed using the **Edit Undo** command. However it is important to keep in mind that changes to any underlying files such as schemas (XML types) or data services made through the model will not be undone. Instead, edit the data service directly or close and reopen your application before saving your changes.

Note: Changes to a model diagram that affect data services such as when a new relationship is created are only made permanent in Eclipse for WebLogic after you do a **File > Save All**.

5.3.3 Notable Relationship Modeler Properties

Properties both reflect and define relationships created in the model diagram.

Table 5–6 Notable Data Modeling Properties

Scope	Property	Settings	Comments
Data Service			Properties described in Managing Your Data Service.
Relationships	Nodes	Read only	Shows names of the related data services and their respective roles. Roles are assigned as source data service and target data service, but these assignments are arbitrary in the case of bidirectional relationships.
	Source and Target Cardinality	Drop down	Value can be 0-to-1, 0-to-many, 1-to-1, 1-to-many, and many-to-many.
Operations			Properties described in Managing Your Data Service.

5.3.4 Relationship Models in Source View

Generated Relationship Declarations in Source View

An example of a navigation function in the underlying source is:

```
(::pragma function <f:function xmlns:f="urn:annotations.ld.oracle.com"
kind="navigate" roleName="ADDRESS"/>:::-)
```

This specifies a relationship to the **Address** data service from the **Customer** data service.

Data services also contain declarations describing the nature of the relationship; this information is the source for the role names and cardinality values that appear in your model diagram.

For example, the data service **Address** contains the following relationship declarations:

```
<relationshipTarget roleName="CUSTOMER" roleNumber="1"
XDS="ld:DataServices/CustomerDB/CUSTOMER.ds" opposite="ADDRESS"/>
```

For each data service, a relationship is created which identifies its role name, cardinality, opposite data service, and a unique (to the data service) role number.

In the above example, a navigation function is automatically created that retrieves customer information based on the **customerID**.

In the case of the relationship between **Customer** and **Address**, the relationship is 0-to-n for the **Address** role (it can make an appearance any number of times or not at all) based on **CustomerID** being a foreign key in **Address** and a primary key in the **Customer** data service (and the underlying relational data sources respectively).

Since the relationships are bilateral, **Customer's** opposite is **Address** while **Address's** opposite is **Customer**.

If your data model is composed of both physical and logical data services, you should keep in mind that a metadata update on any underlying physical data services will remove any relationships you have created involving those data services.

Building XQueries

This chapter describes how to build XQueries.

This chapter contains the following sections:

- [Section 6.1, "How To"](#)
- [Section 6.2, "Reference"](#)
- [Section 6.3, "Related Topics"](#)

6.1 How To

These sections describe how to build XQueries:

- [Section 6.1.1, "Create a Return Type"](#)
- [Section 6.1.2, "Add a Complex Child Element to a Return Type"](#)
- [Section 6.1.3, "Check Namespaces in Return Types"](#)
- [Section 6.1.4, "Create Conditional Elements in Return Types"](#)
- [Section 6.1.5, "Add a Where Clause to a Query"](#)
- [Section 6.1.6, "Use the XQuery Expression Editor"](#)
- [Section 6.1.7, "Use the Source Editor"](#)

6.1.1 Create a Return Type

How To Create a Return Type

This topic describes the basics of creating return types for logical entity data services in the Query Mapper and directly in XML.

This section describes the following topics:

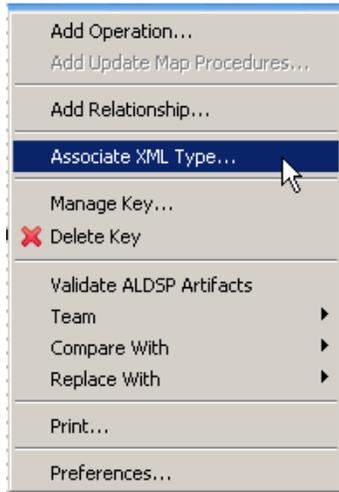
- [Section 6.1.1.1, "Choose a Technique"](#)
- [Section 6.1.1.2, "Write a Return Type Schema"](#)
- [Section 6.1.1.3, "Generate a Schema File"](#)
- [Section 6.1.1.4, "See Also"](#)

6.1.1.1 Choose a Technique

Data services use both XML types and return types.

XML types represent the shape of a logical data service, in the form of an XML schema. They are templates from which return types are created, comparable to a Java class. You use an XML type when you first create a logical entity service and add an XML schema to define its shape.

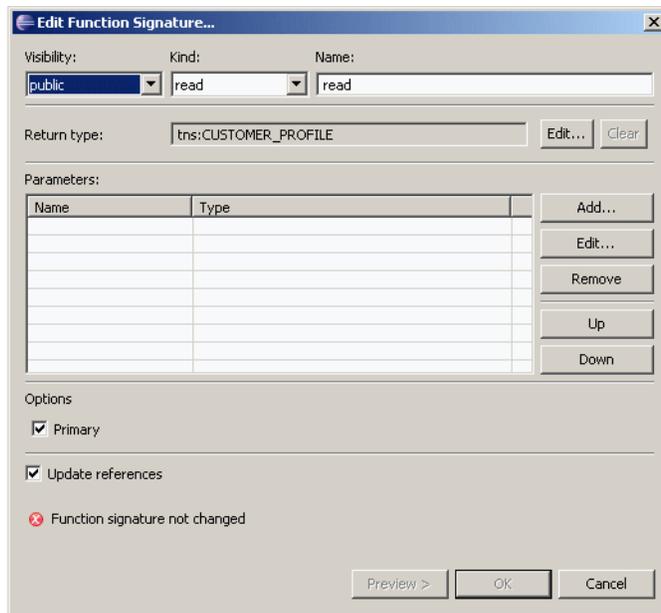
Figure 6–1 Adding an XML Type to a Service



This graphic shows the Associate XML Type option.

Return types represent the shape of data that a query produces when it is run. They are specific instances of an XML type, comparable to a Java object. Return types are the **R** in an **XQuery FLWOR** clause. For example, a service's primary read function returns a return type.

Figure 6–2 Checking the Return Type of a Read Function



This graphic shows the Edit Function Signature dialog.

An XML type is the backbone of a logical data service, because it defines the data the service returns. The XML schema that represents the XML type can combine any elements from any data sources the logical data service uses, including relational sources, web services, XML files, text files, and Java methods.

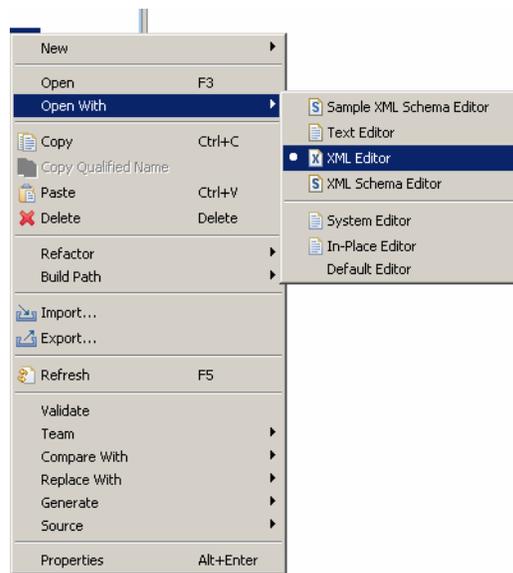
The schema for the logical data service is designed as a separate layer of the dataspace project, regardless of the actual structure of the underlying physical data sources. The schema is not required to use all elements in, or the same structure as, the physical data sources.

You can create a return type schema, an XSD file, in two ways:

- Top down, in an XML editor, either the one built into Eclipse for WebLogic or a standalone editor.
- Bottom up, by building the service visually in **Query Map** view and then using the **Save and Associate XML Type** command.

You should create the XSD file in the logical layer of your dataspace project, as it belongs to the logical data service. Eclipse for WebLogic provides several XML editors, which you can see if you right-click an XSD file in the **Project Explorer** and choose **Open With**.

Figure 6–3 *Choosing an XML Editor in Eclipse for WebLogic*



This graphic shows the Open with > XML Editor option.

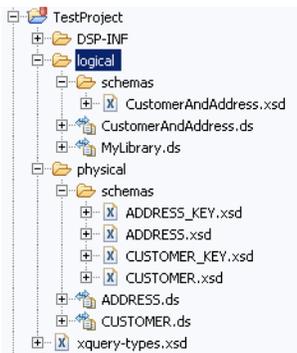
6.1.1.2 Write a Return Type Schema

To create the schema in an XML editor in Eclipse for WebLogic:

1. Choose a location for logical data service schemas in your dataspace project.

You may want to create a folder for schemas in the logical layer of your project (for example, `MyDataSpace/logical/schemas`) separate from the schemas folder that Oracle Data Service Integrator auto-generates for physical data services.

Figure 6–4 Project structure



This option shows the tree structure of a project.

2. Choose **File > New > Other**.
3. Choose **XML > XML Schema**, and click **Next**.
4. Choose a folder, enter a file name that ends in `.xsd`, and click **Finish**.

The generated schema looks something like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<schema xmlns="http://www.w3.org/2001/XMLSchema"
        targetNamespace="http://www.example.org/MySchema"
        xmlns:tns="http://www.example.org/MySchema" elementFormDefault="qualified">
</schema>
```

5. In the XML editor, change the URL of `targetNamespace` to one within your dataspace project:

```
targetNamespace="ld:logical/MyLibrary"
```

The `targetNamespace` URL should start with the prefix `ld:`, and `logical` indicates that the schema resides in the folder named `logical` in your dataspace project. The identifier that follows (here, **CustomerAndAddress**) defines the namespace.

6. Delete the namespace definition for `xmlns:tns`, if your service binds `tns` to a different namespace. You can check this by clicking the **Overview** tab, then the **Properties** tab.

Figure 6–5 Data Service Properties

Data Service	
Property	Value
[-] General	
Name	MyLibrary.ds
Type	Logical (Library)
Description	
Author	
Creation Date	2007-10-15T13:31:36
[-] Security Resources	
Add New	
[-] Prefix Bindings	
ld:logical/MyLibrary	cus
ld:logical/MyLibrary	tns
[-] User Properties	
Add new property:	

This graphic shows the properties tab.

At this point, your schema file should like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<schema xmlns="http://www.w3.org/2001/XMLSchema"
        targetNamespace="ld:logical/CustomersAndOrders"
        elementFormDefault="qualified">
</schema>
```

7. Continue adding complex types, elements, and attributes using the XML editor.
8. Save the file, then right-click anywhere in it and choose **Validate**.

You can also create the return type schema using an XML editor outside Eclipse for WebLogic and then move the XSD file to your Eclipse for WebLogic dataspace project.

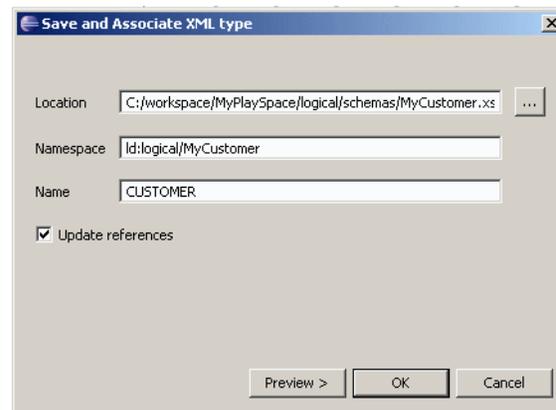
6.1.1.3 Generate a Schema File

You can also have Oracle Data Service Integrator generate the return type schema after you build the query map visually.

To generate the schema in Eclipse for WebLogic, follow these instructions (or see [Section 1.3, "Example: How to Create Your First Data Services"](#) for detailed instructions):

1. Create a dataspace.
2. Create physical data services in the dataspace.
3. Also in the dataspace, create a logical data service (**File > New > Logical Data Service**).
4. Create a Read function in the logical data service (**Overview** tab, right-click, **Add Operation**).
5. Drag the **Read** functions of the physical services you want to use to the **Query Map** tab.
6. Click **Overwrite** , and drag the root element in the **For** box to the root element in the **Return** type.
7. Right-click on the complex element in the Return type, and choose **Expand Complex Mapping**.
8. Right-click the return type box, and choose **Save and Associate XML Type**.

Figure 6–6 Save and Associate XML Type dialog



This graphic shows the Save and Associate XML Type dialog.

For **Location**, select the correct folder for logical schemas. In **Namespace**, enter a namespace that starts with `ld:logical`, such as `ld:logical/MyCustomer`. Be sure that the name of the root element (here, `CUSTOMER`) is unique within the namespace. (The `ld` namespace refers to the original name of Oracle Data Service Integrator, Liquid Data).

9. Click **OK**.
10. Save the file, then right-click anywhere in it and choose **Validate**.

6.1.1.4 See Also

How Tos

- [Section 6.1.2, "Add a Complex Child Element to a Return Type"](#)
- [Section 1.3, "Example: How to Create Your First Data Services"](#)

Other Resources

- XML Schema Tutorial (W3Schools)
<http://www.w3schools.com/schema/default.asp>
- XML Schema Part 1: Structures (W3C)
<http://www.w3.org/TR/xmlschema-1/>
- XML Schema Part 2: Datatypes (W3C)
<http://www.w3.org/TR/xmlschema-2/>

6.1.2 Add a Complex Child Element to a Return Type

Add a Complex Child Element to a Return Type

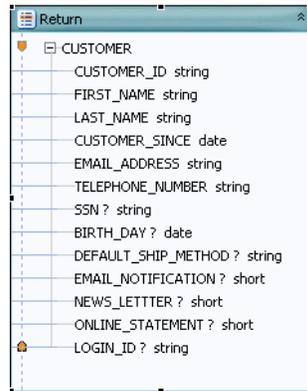
This topic describes how to add a complex child element to a return type, in Eclipse for WebLogic or in the XML source of the return type.

This section describes the following topics:

- [Section 6.1.2.1, "Add the Child Element Visually"](#)
- [Section 6.1.2.2, "Edit the XML Source"](#)
- [Section 6.1.2.3, "See Also"](#)

6.1.2.1 Add the Child Element Visually

Once you create a return type, you can add a complex type as a child of any element, in **Query Map** view. The complex child element must represent a physical data service. The parent element can have a one-to-many or one-to-one relationship with the child, depending on how you want the result data returned.

Figure 6–7 A Simple Return Type Before Adding a Child Element

This graphic shows a sample project tree.

To add a complex child element to a return type visually:

1. Open the logical data service in Eclipse for WebLogic.
2. Check **Project Explorer**. Be sure that your dataspace project has a physical data service for the complex child element you want to add. If it does not, add one
File > New > Physical Data Service
3. Click the **Query Map** tab.
4. In the return type, right-click the new parent element, and choose **Add Complex Child Element**.

Figure 6–8 Add Complex Child Element dialog

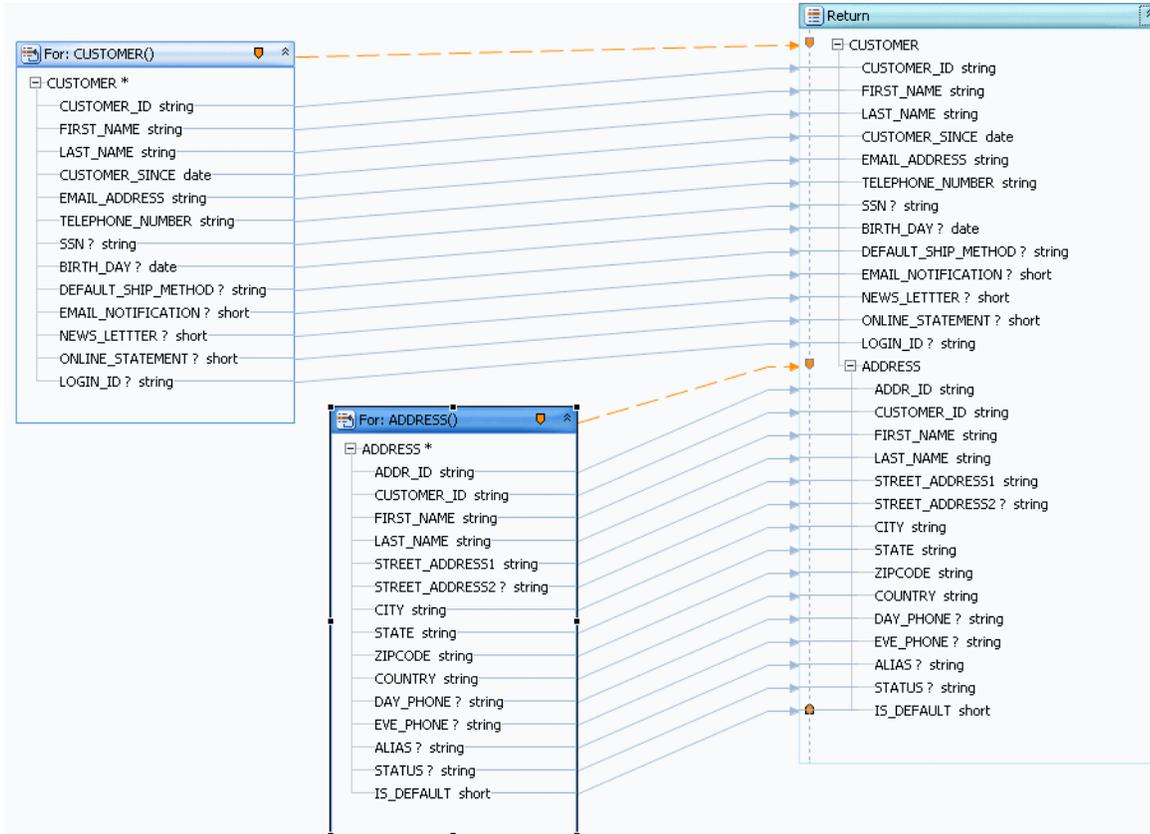
This graphic shows the Add Complex Child Element dialog.

5. For the **Schema File** field, browse (...) to the schema of the physical data service that represents the complex child element.
6. For **Type**, choose a complex type from the schema, then click **OK**.
7. From **Project Explorer**, drag the primary read function of the physical data service to the **Query Map**.
8. Starting from the child element's **For** block, drag the zone icon to the child element in the return type.

- Starting from the child element's For block, drag the parent type of the complex element to the return type.

This step maps all of the elements in the complex child to the return type.

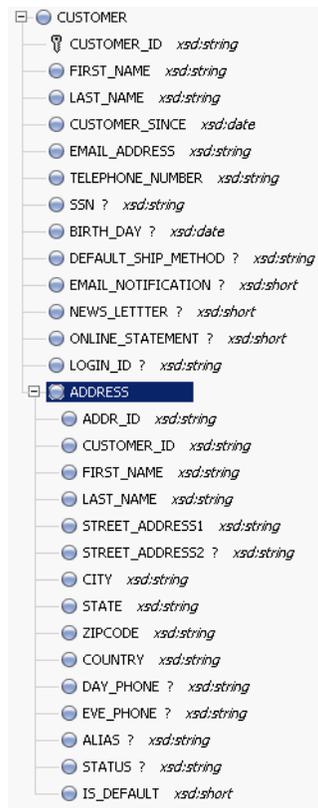
Figure 6–9 Mapping Elements



This graphic shows the mappings.

- Right-click the title bar of the return type, and choose **Save and Associate XML Type**.
- Click the **Overview** tab, and expand the schema to view the complex child in the return type.

You can also right-click the schema and choose **Edit Schema** to view the XML source.

Figure 6–10 Mapped Element

This graphic shows the mapped ADDRESS table.

6.1.2.2 Edit the XML Source

Adding the complex child element to the return type in the XML source accomplishes the same thing as adding it visually.

To add a complex child element to a return type in XML source:

1. Open the logical data service in Eclipse for WebLogic.
2. Check **Project Explorer**. Be sure that your dataspace project has a physical data service for the complex child element you want to add. If it does not, add it:

File > New > Physical Data Service

3. Click the **Overview** tab.
4. Right-click the return type schema in the center, and choose **Edit Schema**.

You see the schema for the logical data service, without the child element:

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema targetNamespace="ld:logical/MyCustomer"
xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="CUSTOMER">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="CUSTOMER_ID" type="xs:string"/>
        <xs:element name="FIRST_NAME" type="xs:string"/>
        <xs:element name="LAST_NAME" type="xs:string"/>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

```

        <xs:element name="CUSTOMER_SINCE" type="xs:date" />
    </xs:sequence>
</xs:complexType>
</xs:element>
</xs:schema>

```

5. In **Project Explorer**, right-click the schema file of the physical data service that represents the child element, and choose **Open With**.

You see the schema of the child element.

6. Copy the complex type from the physical data service schema to the logical data service schema. Take only the complex type:

```

<xs:complexType>
  <xs:sequence>
    <xs:element name="ADDR_ID" type="xs:string" />
    <xs:element name="CUSTOMER_ID" type="xs:string" />
    <xs:element name="FIRST_NAME" type="xs:string" />
    <xs:element name="LAST_NAME" type="xs:string" />
    <xs:element name="STREET_ADDRESS1" type="xs:string" />
    <xs:element name="STREET_ADDRESS2" type="xs:string" minOccurs="0" />
    <xs:element name="CITY" type="xs:string" />
    <xs:element name="STATE" type="xs:string" />
    <xs:element name="ZIPCODE" type="xs:string" />
    <xs:element name="COUNTRY" type="xs:string" />
    <xs:element name="DAY_PHONE" type="xs:string" minOccurs="0" />
    <xs:element name="EVE_PHONE" type="xs:string" minOccurs="0" />
    <xs:element name="ALIAS" type="xs:string" minOccurs="0" />
    <xs:element name="STATUS" type="xs:string" minOccurs="0" />
    <xs:element name="IS_DEFAULT" type="xs:boolean" />
  </xs:sequence>
</xs:complexType>

```

7. Right-click in the schema, and choose **Validate**.

6.1.2.3 See Also

How Tos

- [Section 6.1.1, "Create a Return Type"](#)
- [Section 6.1.3, "Check Namespaces in Return Types"](#)

Other Resources

- XML Schema Tutorial (W3Schools)
<http://www.w3schools.com/schema/default.asp>
- XML Schema Part 1: Structures (W3C)
<http://www.w3.org/TR/xmlschema-1/>
- XML Schema Part 2: Datatypes (W3C)
<http://www.w3.org/TR/xmlschema-2/>

6.1.3 Check Namespaces in Return Types

How To Check Namespaces in Return Types

This topic shows you how to make sure the namespaces used in your return type are correct.

This section describes the following topics:

- [Section 6.1.3.1, "Check Prefix Bindings"](#)
- [Section 6.1.3.2, "Edit the Namespace"](#)
- [Section 6.1.3.3, "See Also"](#)

6.1.3.1 Check Prefix Bindings

In the return type, a child element must be in the same namespace as its parent. If a return type uses elements in different namespaces, you cannot deploy the logical data service to the server or test it from Eclipse for WebLogic.

The exception to this rule is when the parent and child are in different namespaces, but both namespaces have the same prefix binding. Check prefix bindings first, and then edit the namespace, if needed.

To check prefix bindings in the Overview tab:

1. Click the **Overview** tab.
2. Click the **Properties** tab (if it's not visible, choose **Window > Show View > Properties**).

Figure 6–11 Properties tab

Property	Value
Type	Logical (Entity)
Associated XML Type	"ld:logical/schemas/MyCustomer.xsd"
Description	
Author	
Creation Date	2007-10-17T15:48:32
<input type="checkbox"/> Security Resources	
Add New	
<input type="checkbox"/> Prefix Bindings	
ld:logical/MyCustomer	myc
ld:logical/MyCustomer	myc1
ld:physical/ADDRESS	add
ld:physical/ADDRESS	phy
ld:physical/CUSTOMER	cus

This graphic shows the Properties tab.

To check prefix bindings in Source:

1. Click the **Source** tab.
2. Look for the XQuery namespace statements:

```
import schema namespace myc="ld:logical/MyCustomer" at
"ld:logical/schemas/MyCustomer.xsd";
declare namespace cus= "ld:physical/CUSTOMER";
import schema namespace myc1="ld:logical/MyCustomer" at
"ld:logical/schemas/MyCustomer_KEY.xsd";
```

In both these examples, the `myc` and `myc1` namespaces have the same prefix binding. You can have a parent element in one and a child element in another. But if you have a parent element in `myc` and a child in `cus`, you need to change one namespace in the return type.

6.1.3.2 Edit the Namespace

Once you check the prefix binding, you can check a namespace used in a return type and change it in the Query Map or Source view.

To edit a namespace in Query Map view:

1. Click the **Query Map** tab.
2. Select the parent element in the return type, then click it.

Be sure to select and then click; do not double-click.

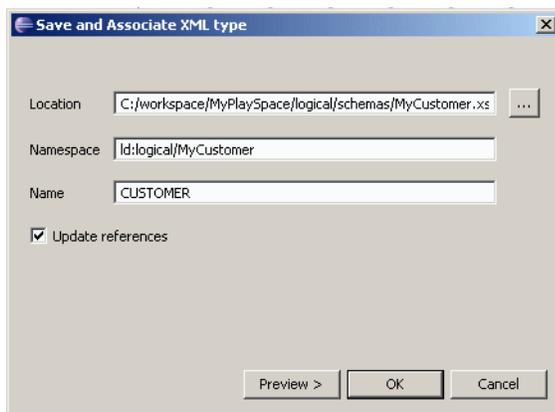
Figure 6–12 Parent Element



This graphic shows the selected parent element.

3. Select the child element in the return type, then click it.
4. If the child element is in a different namespace, change it to the namespace of the parent.
5. Right-click the title bar of the return type, and choose **Save and Associate XML Type**.

Figure 6–13 Save and Associate XML Type dialog



This graphic shows the Save and Associate XML Type dialog

6. Enter the correct location, namespace, and root element name for the return type. Click **OK**.

To edit a namespace in Source view:

1. Click the **Source** tab.

- Expand the primary Read function:

Figure 6–14 Primary Read Function

```
Ⓢ declare function myc:read() as element(myc:CUSTOMER)*{
```

- Locate the namespace of the child element and change it to the namespace of the parent, both in the start and end elements:

```
declare function myc:read() as element(myc:CUSTOMER)*{
  for $CUSTOMER in cus:CUSTOMER()
  return
    <myc:CUSTOMER>
      ...
      {
        for $ADDRESS in add:ADDRESS()
        where $CUSTOMER/CUSTOMER_ID eq $ADDRESS/CUSTOMER_ID
        return
          <myc:ADDRESS >
            ...
          </myc:ADDRESS>
      }
}
```

- Save the changes.

6.1.3.3 See Also

How Tos

- [Section 6.1.1, "Create a Return Type"](#)
- [Section 6.1.2, "Add a Complex Child Element to a Return Type"](#)

Other Resources

- XML Schema Tutorial (W3Schools)
<http://www.w3schools.com/schema/default.asp>
- XML Schema Part 1: Structures (W3C)
<http://www.w3.org/TR/xmlschema-1/>
- XML Schema Part 2: Datatypes (W3C)
<http://www.w3.org/TR/xmlschema-2/>

6.1.4 Create Conditional Elements in Return Types

These sections describe how to add a condition to a return type and determine the elements that are returned when the condition is true or false:

- [Section 6.1.4.1, "Add the Condition"](#)
- [Section 6.1.4.2, "Create the Expression"](#)
- [Section 6.1.4.3, "See Also"](#)

6.1.4.1 Add the Condition

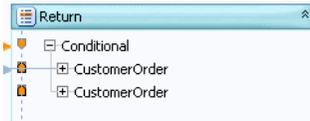
A condition in a return type defines two groups of elements: those returned when an expression is true, and those returned when an expression is false. When you add a condition to a return type, you see two groups of return type elements.

To add a condition to the return type:

1. Click the **Query Map** tab.
2. Right-click an element in the return type, and choose **Make Conditional**.

The conditional element is now duplicated.

Figure 6–15 Duplicated Element



This graphic shows the duplicated conditional element.

6.1.4.2 Create the Expression

You must add the conditional expression, that determines which element is returned, in the XQuery source. You cannot add a conditional expression in the expression editor.

1. Click the **Source** tab.

The primary Read function now has an if...else clause:

```
declare function tns:read() as element(cus:CustomerOrder)*{
  for $CUSTOMER in cus1:CUSTOMER()
  return
    if (true()) then
      <cus:CustomerOrder>
      ...
    </cus:CustomerOrder>
    else
      <cus:CustomerOrder>
      ...
    </cus:CustomerOrder>
};
```

The expression after the if statement is evaluated, and the service returns either the first or second set of elements. The XQuery true() function simply returns the Boolean value true.

2. In the XQuery source, replace true() with another XQuery expression, for example:

```
if ( fn:data( $CUSTOMER/LAST_NAME ) = "Black" ) then
```

You can use any XQuery expression that returns a value of true or false. In this example, if a customer has the last name **Black**, the first element group is returned. If not, the second element group is returned.

To add the value of an element in a For block, use the XQuery fn:data function, which takes the value of an element:

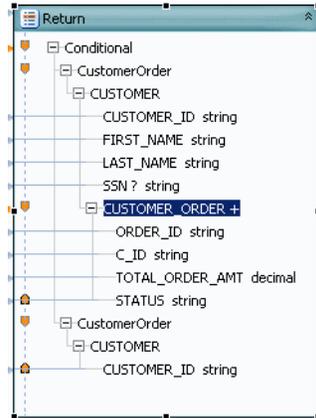
```
<LAST_NAME>Black</LAST_NAME>
```

3. Click the **Query Map** tab.

- In the return type, add or delete elements in either group to create the return groups you want.

Remember that the first group is returned if the expression is true, and the second group if the expression is false.

Figure 6–16 Add or Delete Elements



This graphic shows adding or deleting elements.

- Click the **Test** tab. Choose the **Read** function, and click **Run**. Check that the results are what you intend.

In this example, the full group of elements is only returned for customers with the last name **Black**. For other customers, only the **CUSTOMER_ID** is returned.

Figure 6–17 Results from Read Function

Select operation: read()

Parameters:
No Parameters

Settings:

Result: Result is valid.

Name	Value
[-] CUSTOMER	
[-] CUSTOMER_ID	CUSTOMER1
[-] FIRST_NAME	Jack
[-] LAST_NAME	Black
[-] SSN	295-13-4119
[+] CUSTOMER_ORDER	
[-] CustomerOrder	
[-] CUSTOMER	
[-] CUSTOMER_ID	CUSTOMER2
[-] FIRST_NAME	
[-] LAST_NAME	
[-] SSN	
[-] CustomerOrder	
[-] CUSTOMER	
[-] CUSTOMER_ID	CUSTOMER3
[-] FIRST_NAME	
[-] LAST_NAME	
[-] SSN	

This graphic shows the results from the Read function.

6.1.4.3 See Also

How Tos

- [Section 6.1.1, "Create a Return Type"](#)
- [Section 6.1.2, "Add a Complex Child Element to a Return Type"](#)
- [Test a Read Function and Simple Update](#)

Other Resources

- [Introduction to XQuery](#)
<http://www.devx.com/xml/Article/8046/0/page/3>
- [XQuery Tutorial](#)
<http://www.w3schools.com/xquery/default.asp>
- [XQuery 1.0 Specification](#)
<http://www.w3.org/TR/xquery/>

6.1.5 Add a Where Clause to a Query

Add a Where Clause to a Query

These sections describe several ways of adding XQuery where clauses to queries to join relational data sources:

- [Section 6.1.5.1, "Define the Condition"](#)
- [Section 6.1.5.2, "Join Tables with a Where Clause"](#)
- [Section 6.1.5.3, "Use an XQuery Function in a Where Clause"](#)
- [Section 6.1.5.4, "See Also"](#)

6.1.5.1 Define the Condition

A where clause in XQuery specifies criteria defining some return data. This is a simple XQuery where clause:

```
where $CUSTOMER/CUSTOMER_ID = "1111"
```

A where clause is usually part of an XQuery FLWOR (for-let-where-order by-return) expression. The where clause can be any XQuery expression, including another FLWOR expression. A common use of a where clause is to join two relational data sources, for example:

```
for $CUSTOMER_ORDER in cus1:CUSTOMER_ORDER()
where $CUSTOMER/CUSTOMER_ID eq $CUSTOMER_ORDER/C_ID
return
... xml elements here ...
```

The where clause here specifies a condition that defines a subset of results to return. The SQL statement Oracle Data Service Integrator generates from this XQuery expression creates a left outer join between two tables:

```
SELECT t1."CUSTOMER_ID" AS c1, t1."FIRST_NAME" AS c2, t1."LAST_NAME" AS c3,
t1."SSN" AS c4,t2."C_ID" AS c5, t2."ORDER_ID" AS c6, t2."STATUS" AS c7, t2.
```

```

"TOTAL_ORDER_AMT" AS c8
FROM "RTLCUSTOMER"."CUSTOMER" t1
LEFT OUTER JOIN "RTLAPPLOMS"."CUSTOMER_ORDER" t2
ON (t1."CUSTOMER_ID" = t2."C_ID")
ORDER BY t1."CUSTOMER_ID" ASC

```

Before you add a where clause to a logical data service, think about how to structure it. If you want to join two data sources, you can only do so on a key field that appears in both. In this example, the **CUSTOMER** table has a primary key named **CUSTOMER_ID** joined to a **CUSTOMER_ORDER** table with a foreign key named **C_ID**.

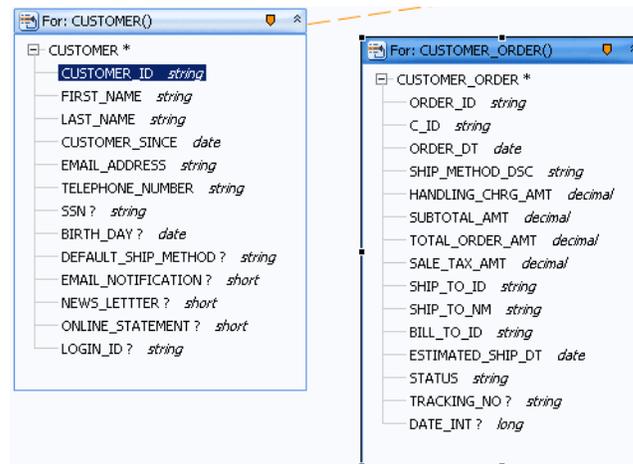
6.1.5.2 Join Tables with a Where Clause

The simplest way to create a where clause between two relational data sources is to map it in Query Map view.

To map the where clause:

1. Open a logical data service in Eclipse for WebLogic.
2. Click **Query Map**.
3. Drag the read functions  of at least two physical data sources from **Project Explorer** to the **Query Map** view.

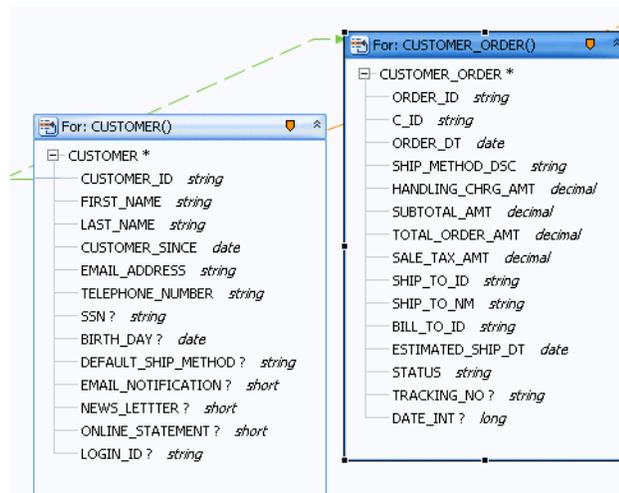
Figure 6–18 Read Functions



This graphic shows the read function from the data source.

4. In **Query Map** view, drag from a key element in the first data source to the corresponding key element in the second.

Figure 6–19 Key Elements



This graphic shows the association of the key element between the data sources.

If you click the second data source, you see the XQuery where clause in the expression editor:

Where \$CUSTOMER/CUSTOMER_ID eq \$CUSTOMER_ORDER/C_ID

6.1.5.3 Use an XQuery Function in a Where Clause

A where clause can also contain an XQuery function, including any built-in or Oracle-defined functions available from the **Design Palette**. The where clause is defined on an element within a For node.

To create a where clause with an XQuery function:

1. Click **Query Map**.
2. Click the **For** title bar of the node that contains the element.
3. Click **Add Where Clause**  to insert the where clause.
4. Open the **Design Palette (Window > Show View > Design Palette)**.
5. Expand XQuery Functions, then choose a function (for example: **Duration, Date, and Time Functions > fn:year-from-date**).
6. Drag the function to the expression editor.
7. Delete \$arg in the function, then click the element in the For node that you want to add.
8. Add an operator and a value to complete the expression.

```
fn:year-from-date($CUSTOMER/CUSTOMER_SINCE) < 2000
```

You can use any of the XQuery operators available in **Design Palette > XQuery Operators**.

9. Click **Save** .

In **Source** view, the where clause in the read function looks like this:

```
declare function tns:read() as element(tns:CUSTOMER_PROFILE) *{
```

```

for $CUSTOMER in cus1:CUSTOMER()
where fn:year-from-date($CUSTOMER/CUSTOMER_SINCE) < 2000
return
... xml elements here ...

```

10. Test the query in **Test** view, preferably on sample data, to make sure the results are what you expect.

6.1.5.4 See Also

How Tos

- Test a Read Function and Simple Update
- Test a Create or Delete Procedure

6.1.6 Use the XQuery Expression Editor

How To Edit XQuery Expressions

This topic describes how to edit XQuery expressions in the expression editor in Eclipse for WebLogic.

This section describes the following topics:

- [Section 6.1.6.1, "Overview"](#)
- [Section 6.1.6.2, "The fn-bea:value Function"](#)
- [Section 6.1.6.3, "See Also"](#)

6.1.6.1 Overview

You can edit the generated XQuery expressions in an update map using the expression editor.

Figure 6–20 The Expression Editor in an Update Map



This graphic shows the expression editor.

The update map expression language is a subset of XQuery syntax. In an update map, you can use any of the following XQuery constructs.

Table 6–1 XQuery Constructs

Type	Description	Example
Variable	A variable already defined in a For Each or Update block in the update map. \$\$root is a special predefined variable that refers to the root of the service's XML type.	\$ORDER_WITH_LINE_ITEM \$CUSTOMER
Constant	A numeric, string, or other constant.	"a" "12345"
Constant Cast	A constant cast to another XSD data type using the parentheses operator.	xsd:date("2007-01-01")
Function	A call to any XQuery function. You can see the built-in and Oracle-provided functions in the Design Palette. You can use a variable, path, or constant as an argument to a function.	fn-bea:value(\$CUSTOMER/FIRST_NAME)
Path	An expression that locates an XML element in a tree using variables, elements, and attributes. The syntax is: \$VARIABLE_NAME /elementName @attributeName	\$ORDER_WITH_LINE_ITEM/CUSTOMER_ORDER/ORDER_ID

Namespace prefixes are declared in the data service's XQuery source, which you can see in the **Source** tab. If a namespace is only used in the update map, and not in the logical data service, you must declare it. If a namespace cannot be resolved, it is shown with the prefix `ns?`.

The most common ways you use the expression editor are to:

- Add a constant to an unmapped element
- Cast a constant to an XSD data type, especially to resolve update block elements with no mappings
- Use an XQuery function available in the **Design Palette** to cast a value
- Use a custom XQuery cast function you have written

6.1.6.2 The fn-bea:value Function

A mapping between an element in a return type and an element in an update block uses the `fn-bea:value` function with a path name, for example:

```
fn-bea:value($CUSTOMER/CUSTOMER_ID)
```

An update mapping should always use `fn-bea:value`, whether Oracle Data Service Integrator auto-generates the mapping or you draw it. If you remove the `fn-bea:value` function from the expression and simply use an XQuery path expression (`$CUSTOMER/CUSTOMER_ID`), the element becomes disabled in the update map and you see this error message:

```
The expression does not match the expected type for this element
The expression assigned to this element is not valid
Hint: did you forget to use the value function?
```

The `fn-bea:value` function is required, because an update map updates a Service Data Object (SDO) and requires a special XML structure called a *datagraph* that includes a

change summary showing both the old and new values. The `fn-bea:value` function handles the update to the SDO correctly.

If you do not use `fn-bea:value`, Oracle Data Service Integrator throws an exception when you attempt to update the value.

6.1.6.3 See Also

Concepts

- [Understanding Update Maps](#)

How Tos

- [How To Handle Unmapped Required Values](#) (includes Cast a Constant)
- [How to Cast Using a Built-in XQuery Function](#)
- [How To Cast Using a Custom XQuery Function](#)

Other Sources

- [W3Schools XQuery Tutorial](#)

6.1.7 Use the Source Editor

This section describes the Oracle Data Service Integrator Source editor and highlights its editing features. The following topics are included:

- [Section 6.1.7.1, "What is the Source Editor?"](#)
- [Section 6.1.7.2, "Searching Source"](#)
- [Section 6.1.7.3, "Navigating to Specific Functions"](#)
- [Section 6.1.7.4, "Color Coding"](#)
- [Section 6.1.7.5, "Code Completion"](#)

6.1.7.1 What is the Source Editor?

The Source editor is available from a tab in the Oracle Data Service Integrator Eclipse perspective. As you build up your data service, the underlying source is always available from this editor.

Data service source typically:

- References a schema as the data service's XML type (for Entity data services).
- Defines functions in the data service.
- Declares namespaces for referenced data services.
- Contains `pragma` directives to the query engine.

In addition, data services created from physical data sources contain physical source metadata. For example, data services based on relational data describe the XML type (such as `xs:string`), the XPath, native size, native type, null-ability setting and so forth.

In developing data services there are many occasions when it is necessary or convenient to view and/or modify source.

The Source editor allows you to directly edit data service source code, as well as schemas. Changes to source are immediately reflected in other data service modes

such as the Query view editor; similarly, source is immediately updated when changes are made through the Query editor or in Overview mode.

Tip: When a data service is created the root level of your dataspace has "ld:" as its namespace. ld referred to the original name of Oracle Data Service Integrator, Liquid Data.

```
declare namespace ns4= "ld:Update/PhysicalDSs/SDO_WLCO_SET";
```

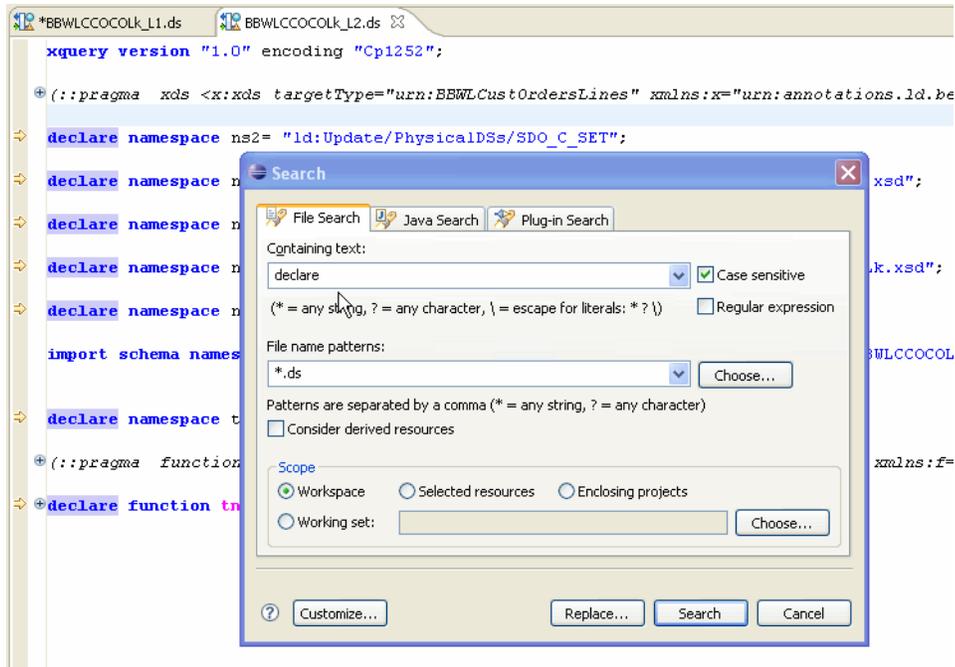
Note: Data Service Annotations

6.1.7.2 Searching Source

Eclipse offers several types of search.

- You can find all occurrences of a string in Source view using **Eclipse Search** menu. Each instance of the term in your project will be highlighted.
- You can use page search (**Ctrl-F**). search to find the next occurrence of a term. Standard search/replace functionality is available.

Figure 6–21 File Search in Eclipse



This graphic shows the File Search function in Eclipse.

6.1.7.3 Navigating to Specific Functions

To open **Source View** to a particular query function in **Overview** mode, first select the function, then click the **Source** tab.

6.1.7.4 Color Coding

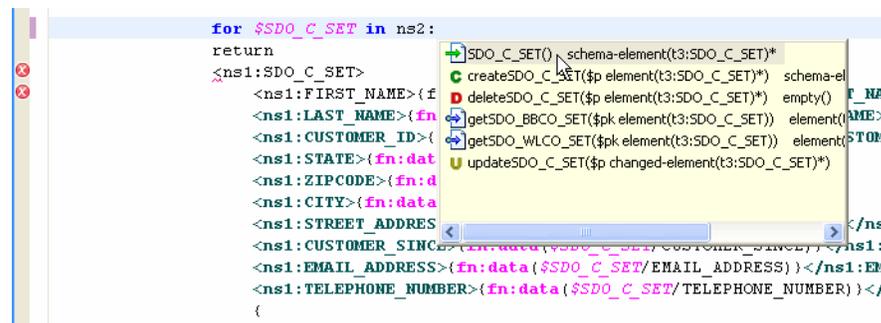
XQuery documents in **Source View** are color-coded according to the following scheme:

Color	Meaning
Blue	Keywords
Dark gray	Comments
Magenta	Variable
Dark green	XML markup
Red	Error conditions

6.1.7.5 Code Completion

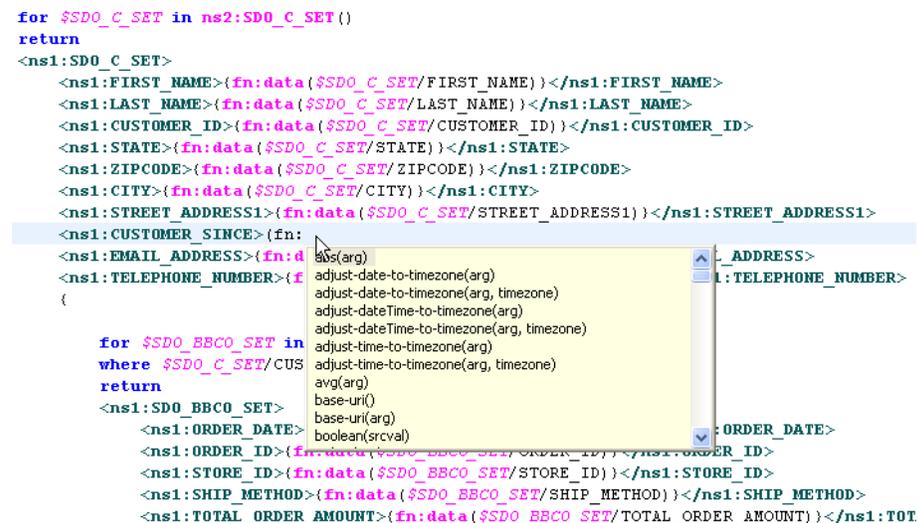
Code completion is available for XPath built-in and user-defined functions. Similarly, function completion is invoked when you type a namespace prefix followed by a colon.

Figure 6–22 Function Completion from Namespace



This graphic shows the code completion for a function.

Figure 6–23 XPath Completion in Source View



This graphic shows the code completion for an XPath.

Figure 6–24 Selecting from Available XQuery Functions

```

for $SDO_C_SET in ns2:SDO_C_SET()
return
<ns1:SDO_C_SET>
  <ns1:FIRST_NAME>{fn:data($SDO_C_SET/FIRST_NAME)}</ns1:FIRST_NAME>
  <ns1:LAST_NAME>{fn:data($SDO_C_SET/LAST_NAME)}</ns1:LAST_NAME>
  <ns1:CUSTOMER_ID>{fn:data($SDO_C_SET/CUSTOMER_ID)}</ns1:CUSTOMER_ID>
  <ns1:STATE>{fn:data($SDO_C_SET/STATE)}</ns1:STATE>
  <ns1:ZIPCODE>{fn:data($SDO_C_SET/ZIPCODE)}</ns1:ZIPCODE>
  <ns1:CITY>{fn:data($SDO_C_SET/CITY)}</ns1:CITY>
  <ns1:STREET_ADDRESS1>{fn:data($SDO_C_SET/STREET_ADDRESS1)}</ns1:STREET_ADDRESS1>
  <ns1:CUSTOMER_SINCE>{fn:data($SDO_C_SET/
  <ns1:EMAIL_ADDRESS>{fn:data($SDO_C_SET/E
  <ns1:TELEPHONE_NUMBER>{fn:data($SDO_C_SE
  (
    for $SDO_BBCO_SET in ns3:SDO_BBCO_SE
    where $SDO_C_SET/CUSTOMER_ID eq $SDO
    return
    <ns1:SDO_BBCO_SET>
      <ns1:ORDER_DATE>{fn:data($SDO_BB
      <ns1:ORDER_ID>{fn:data($SDO_BBCO
      <ns1:STORE_ID>{fn:data($SDO_BBCO_SET/STORE_ID)}</ns1:STORE_ID>
      <ns1:SHIP_METHOD>{fn:data($SDO_BBCO_SET/SHIP_METHOD)}</ns1:SHIP_METHOD>
      <ns1:TOTAL_ORDER_AMOUNT>{fn:data($SDO_BBCO_SET/TOTAL_ORDER_AMOUNT)}</ns1:TOTAL
  (

```

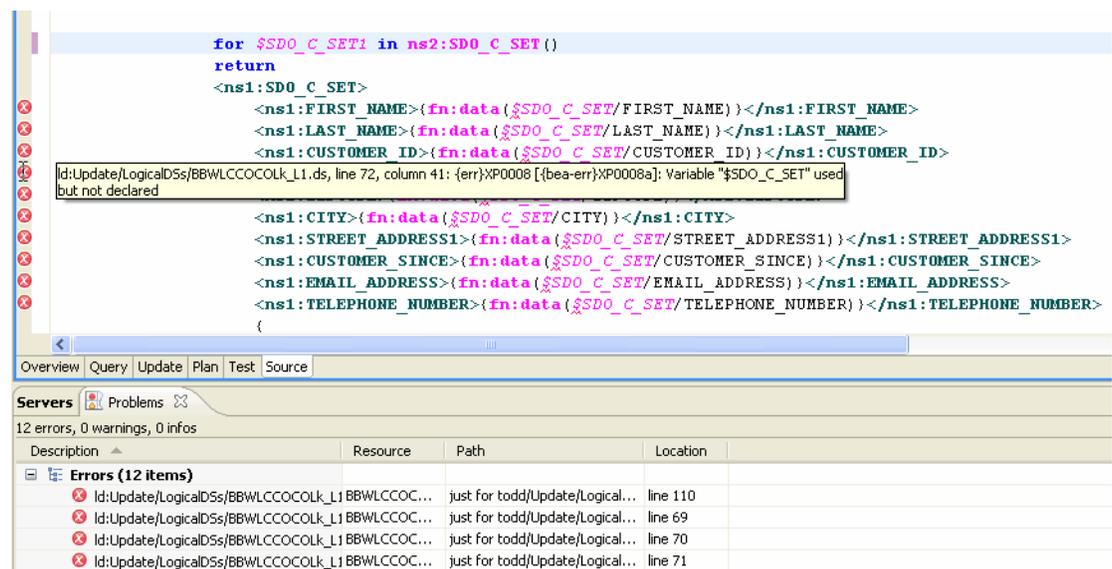
This graphic shows the code completion for an XQuery function.

6.1.7.5.1 Error Identification

Syntax errors that occur in source either as a result of editing or as a result of changes made in the XQuery Editor appear in the **Problems** tab.

Windows > Show View > Problems

Figure 6–25 Induced Error Condition in Source View



This graphic shows an error condition.

- Tip:** ■ Click on the error condition in the **Problems** tab, your cursor will be placed on the relevant line of code.
- Mouse over the error indicator in the **Source** editor, the complete error condition will appear.
 - Right-click on the left margin of **Source** view several options appear including the option to make line numbering active.
 - Right-click anywhere in Source view to access **Source** editor Preferences including permanently displaying line numbering.

6.2 Reference

These sections provide reference information for building XQueries:

- [Section 6.2.1, "XQuery Language Version Support"](#)
- [Section 6.2.2, "Built-in XQuery Functions"](#)

6.2.1 XQuery Language Version Support

Oracle Data Service Integrator supports the XQuery language as specified in XQuery 1.0: An XML Query Language, W3C Working Draft of July, 23, 2004. You can use any feature of the language described by the specification.

Oracle Data Service Integrator supplements the base XQuery syntax with a set of elements and directives that appear in Source View as pragmas. Pragmas are a standard XQuery feature that give implementors and vendors a way to include custom elements and directives within XQuery code.

The Oracle implementation of XQuery also contains some extensions to the language and additional functions. Oracle extensions to XQuery and links to W3C documentation are described in the *XQuery and XQSE Developer's Guide* (http://download.oracle.com/docs/cd/E13162_01/odsi/docs10gr3/xquery/index.html).

6.2.2 Built-in XQuery Functions

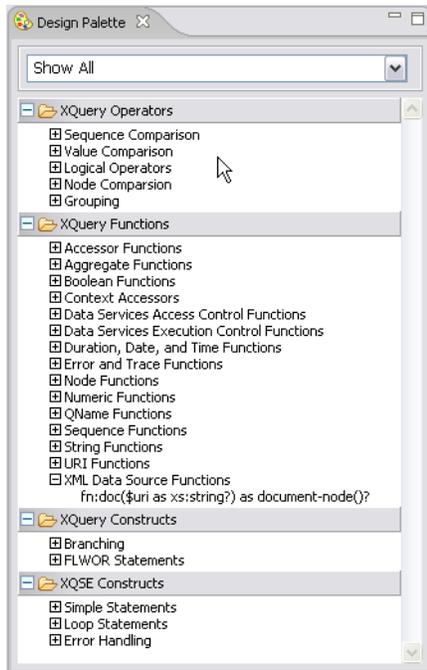
XQuery Functions in Eclipse for WebLogic

Eclipse for WebLogic provides numerous XQuery functions in the Design Palette. If it is not visible you can access it with:

Window > Show View > Design Palette

XQuery functions can be utilized in both **Query** and **Update Map** views.

Figure 6–26 XQuery Functions in the Design Palette



This graphic shows the XQuery functions.

Tip:

- For information on fn-bea XQuery functions, see the *XQuery and XQSE Developer's Guide*.
- For information on standard XQuery functions, see the W3C XQuery 1.0 and XPath 2.0 Functions and Operators specification.

6.3 Related Topics

How-to...

- [Section 1.3, "Example: How to Create Your First Data Services"](#)
- Test a Read Function and Simple Update
- Test a Create or Delete Procedure
- How To Develop Good XQSEs

Concepts...

- Understanding Data Service Annotations

Reference...

- XQuery Scripting Extensions

Testing Update Procedures Using SDO Data Graphs

This chapter is a brief overview of Service Data Objects and data graphs, which you use to test update procedures in Eclipse for WebLogic. This chapter contains the following sections:

- [Section 7.1, "Key Points"](#)
- [Section 7.2, "Updates in Test View"](#)
- [Section 7.3, "Optimistic Locking"](#)

7.1 Key Points

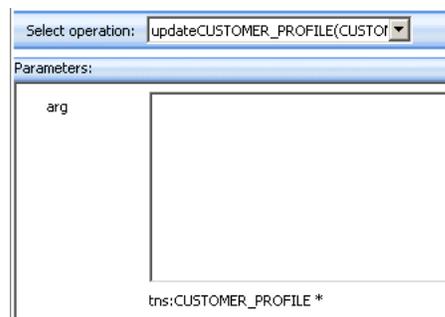
Remember these points:

- To test an Update procedure in Test view, you must submit a data graph as an argument.
- A data graph is an XML structure that contains the data you are changing, as well as the original data.
- When you update a relational data source, Oracle Data Service Integrator uses optimistic locking. The data source is locked at update, not when the data is initially retrieved.

7.2 Updates in Test View

When you test an Update procedure in Test view, you are actually updating a Service Data Object (SDO) from within Eclipse for WebLogic.

Figure 7-1 Update in Test View



Selecting an Update Procedure in Test View

SDO is a programming model for Java platforms that unifies data programming across many types of data sources. SDO is based on data objects, which are simply object instances that contain data. You can update the data objects using either static or dynamic data APIs. With a static API, the shape of the data is defined in advance. However, with a dynamic data API, you can update properties at run time that are not known at development time.

The SDO model is based on data graphs, which are collections of tree-structured data, usually XML. A client retrieves a data graph from a data source, modifies it, and applies the data graph back to the data source.

A data graph contains a `<changeSummary>` element with the original data you are updating. It also contains an XML element with the new data. When both the old and new data are passed back to the data object, the object can be updated.

Example 7-1 A Data Graph with Old and New Data

```
<sdo:datagraph xmlns:sdo="commonj.sdo">
  <changeSummary>
    <sim:SIMPLE_CUSTOMER sdo:ref="#/sdo:datagraph/sim:SIMPLE_CUSTOMER"
xmlns:sim="ld:logical/SimpleCustomer">
      <CUSTOMER_SINCE>1999-01-01T00:00:00</CUSTOMER_SINCE>
    </sim:SIMPLE_CUSTOMER>
  </changeSummary>
  <sim:SIMPLE_CUSTOMER xmlns:sim="ld:logical/SimpleCustomer">
    <CUSTOMER_ID>CUSTOMER7</CUSTOMER_ID>
    <CUSTOMER_SINCE>2007-11-11T00:00:00</CUSTOMER_SINCE>
  </sim:SIMPLE_CUSTOMER>
</sdo:datagraph>
```

7.3 Optimistic Locking

When an SDO updates a relational source, it uses optimistic locking to avoid change conflicts. With optimistic locking, the data source is not locked after the client acquires the data. Later, when an update is needed, the data in the source is compared to a copy of the data taken when it was acquired. If any of the underlying data was changed before the client applies the changes, the update is rejected, and the client must recover.

The optimistic locking policy is set for each relational data source.

Understanding Query Plans

This chapter describes how to obtain and use query plans. This chapter contains the following sections:

- [Section 8.1, "Using Query Plan View"](#)
- [Section 8.2, "Analyzing a Sample Query"](#)
- [Section 8.3, "Working with a Query Plan"](#)

8.1 Using Query Plan View

To obtain a query plan for any function in your data service, select the Query Plan tab and select a function, just as you would in Test View. In addition, as a convenience, you can obtain an ad hoc query plan for XQuery or SQL.

The interface for Query Plan View is quite similar to that used for testing your query functions. You select a function or procedure from a drop down list and then click the **Show Query Plan** button.

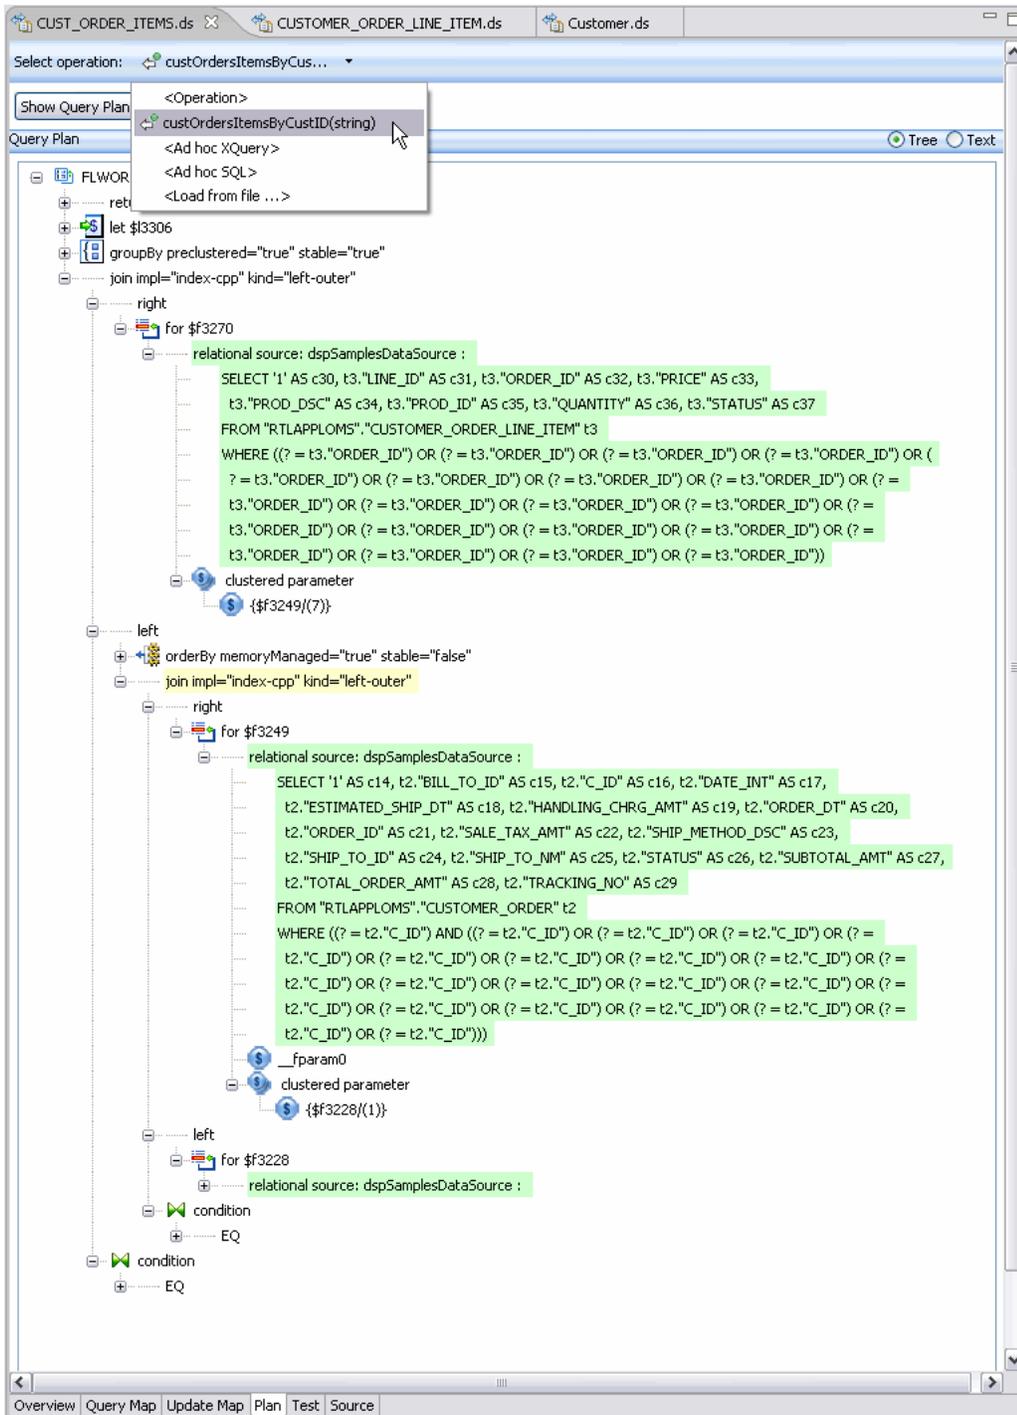
A query plan identifies the following query components:

- Joins
- Outer join
- Select statements
- Data sources
- Custom function calls
- Order-bys
- Remove duplicates

There are several ways that a query plan can be viewed:

- Tree view. A collapsible graphical presentation of the query plan.
- Text view. Presents the information as text.

Figure 8–1 Customer Order Items Query Plan



Customer order items query plan

8.1.1 Query Plan Information and Warnings

The query plan shows both informational and warning messages. When a section of the plan is flagged with a warning, the plan segment is highlighted in red. If you mouse over the segment, the warning message appears.

Informational messages also can appear with plan segments. Such segments are highlighted in yellow.

Table 8–1 Informational and Warning Messages Associated With Query Plans

Warning Message Type	Informational Message Type
XQuery compiler: Typematch. Typematch issues will be resolved by the compiler (may affect performance)	Audit. Auditing has been set for this particular function (will affect performance).
XQuery compiler: No where clause. There is no predicate associated with the query function (will affect performance).	Cache. Function is cached (may enhance performance).
XQuery compiler: Untyped data. Possible untyped atomic data found in the node constructor.	SQL pushdown generation details.
XQuery compiler: No such element. The element (name provided) is not found in in-scope schemas.	NA
SQL generation: missing key. Underlying table/view does not have a key.	NA
SQL generation: cannot generate subquery. isSubquery property is set to false on the data service. (See the "Function Annotations" section of the Understanding Data Services Annotations section of the XQuery Developer's Guide.	NA
SQL generation: cannot generate SQL for join expression. Unable to translate join condition.	NA
SQL generation: cannot generate SQL for aggregate expression (named). Function does not operate on a sequence.	NA
SQL generation: fn:string() function encountered. Use xs:string() instead since xs:string() can be pushed down to the database for processing.	NA

8.1.2 Printing or Saving Your Query Plan

There are two right-click options associated with query plans:

- Prints the plan
- Saves the plan

The default file name for the saved file will appear in the form:

```
<dataServiceName_qp>
```

If you right-click on the root element of the plan, Plan A right-mouse option on the root element in the plan allows you to print a query plan to a printer or a file. Right-click on any node in the plan and select either the print or print to a file option.

If you print to a file the filename will be of type XML. The name of the file will be the function name followed by the letters `_qp`, as in: `getCustomerView_qp.xml`

The file can be saved anywhere in your application.

8.1.3 Loading a Previously Saved Query Plan

You can load a previously saved query plan using the following steps:

1. Select
Load from file...
from the plan drop down box.
2. In the Browse File dialog locate an existing query plan in the current project.
3. Click **Open**.
The selected query plan will be appear.

8.2 Analyzing a Sample Query

Assume a query returns data related to order details after it is passed an order ID and a customer ID.

The following is a "pseudocode" description of the query:

```
for electronic orders matching CustomerID and OrderID
  return order information and ship-to information
  for credit card information matching an AddressID
    return credit information and bill-to address information
  for electronic line item information matching the line item in the order
    return line item information
```

The statements represent mappings or projections in the data service. This can be useful when trying to trace performance issues.

The join conditions are identified in the plan as a left-outer join driven by a complex parameter. By definition, joins have left and right sides, each of which can contain additional joins. One of the best uses of the query plan is to see how the query logic works up the various data threads to return results.

8.3 Working with a Query Plan

Two options are available in Query Plan:

- **Expand All.** This right-click menu option expands the currently selected element and any children. If applied to the top-most element in the plan, all elements are expanded.
- **Match highlighting.** When you click on a variable name any elements (open or closed) containing a match for that variable are highlighted. This feature helps you trace variables in the query plan.

8.3.1 Identifying Problematic Conditions Through the Query Plan

When you show a query plan for a particular function, you may notice red or yellow highlighting of particular routines. These correspond to warnings or informational messages from the plan interpreter. For example, if a for statement is missing a where clause (potentially leading to slow performance or retrieval of a massive amount of data) a red warning will appear adjacent to the statement. Simply mouse-over the highlighted section of the plan to view the information or warning.

Managing Update Maps

This chapter describes how to manage update maps. This chapter contains the following sections:

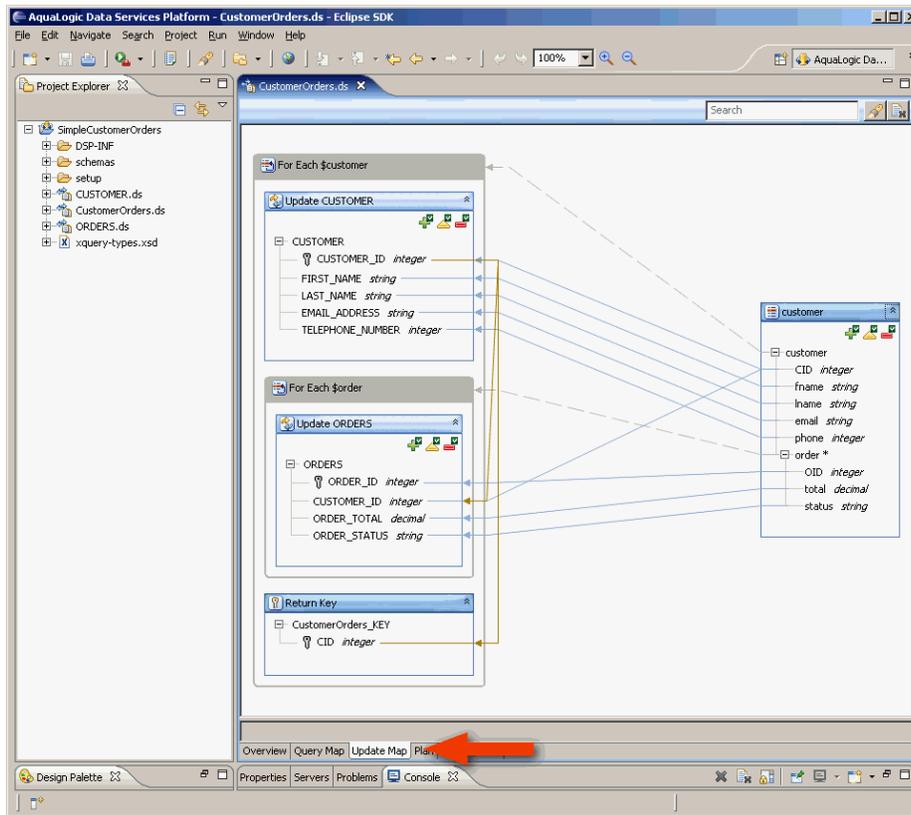
- [Section 9.1, "Understanding Update Maps"](#)
- [Section 9.2, "Changing a Mapping"](#)
- [Section 9.3, "Removing a Mapping"](#)
- [Section 9.4, "Reverting Customizations"](#)
- [Section 9.5, "Adding a Condition to an Update Block"](#)
- [Section 9.7, "Adding an Update Map Procedure"](#)
- [Section 9.8, "Determining the Scope of a Variable"](#)
- [Section 9.9, "Reference"](#)
- [Section 9.10, "How To"](#)
- [Section 9.11, "Testing Update Maps"](#)
- [Section 9.12, "How To Test an Update Procedure"](#)

9.1 Understanding Update Maps

An update map allows you to easily update your logical entity data service without having to write Java or XQSE code. This overview provides a foundation for understanding what an update map is and how you can use one.

Oracle Data Service Integrator generates a default update map automatically when you create a logical entity data service with a primary read function. You can see the update map associated with a data service by clicking the "**Update Map**" tab at the bottom of the screen (see the example that follows).

Figure 9–1 Custom Order Update Map



Custom Order Map

In this overview, as a running example we use an update map for a data service that joins together customers and orders.

The image to the left shows the update map for the data service (CustomerOrders.ds). The orange arrow identifies the location of the "Update Map" tab.

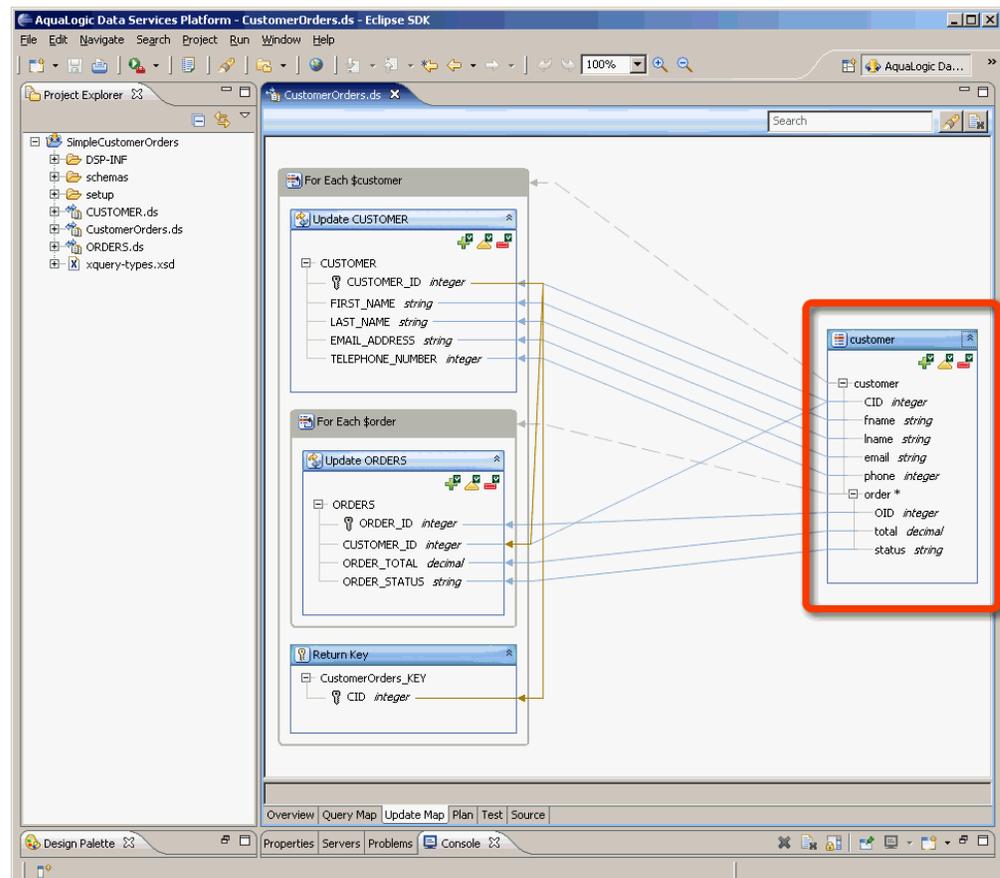
An *update map procedure* is a create, update, or delete procedure that is implemented by an update map. The update map maps values from the input to the update map procedure to the inputs of the procedures in the underlying data services. These underlying data services that the logical entity data service is composed of are referred to as the source data services. In the previous example, the input is mapped to the two source data services **CUSTOMER** and **ORDERS**. The blue arrows in the update map show how the values are mapped.

A logical entity data service has a target type that describes the entity that the data service is about. All read functions in the data service must return instances of the target type and all update map procedures must accept instances of the target type as input. For example, say that we have an entity data service about customers. The read functions of this data service must return customers and update map procedures must take customers as input.

9.1.1 The Target Box

The target box displays the data type of the input to the update map procedures, and the procedure icons. There is exactly one target block in an update map and it is displayed on the right.

Figure 9–2 Target Box



Shows the Target Box which displays the data type of the input to the update map procedures, and the procedure icons.

9.1.1.1 The Input Type

The input type (or, target type) is the type of the data that is passed to the update map procedures. Elements and attributes from the input type are mapped to the update blocks on the left.

9.1.1.2 Procedure Icons

The **Create**, **Update**, and **Delete** procedure icons indicate the status of the corresponding update map procedures. They appear in the upper-right corner of the target box. Each icon may have a green check, a yellow exclamation or a red 'X'. A green check indicates that the update map is fully capable of implementing the procedure. A yellow exclamation indicates that you can invoke the procedure, but there may be problems at runtime. A red 'X' indicates there is a serious problem that needs to be addressed. Any time that there is a red 'X' or a yellow exclamation on the

icon, you can hover the mouse pointer over the icon to get a tool tip providing more information (see Figure 9-3).

Figure 9-3 Procedure icons

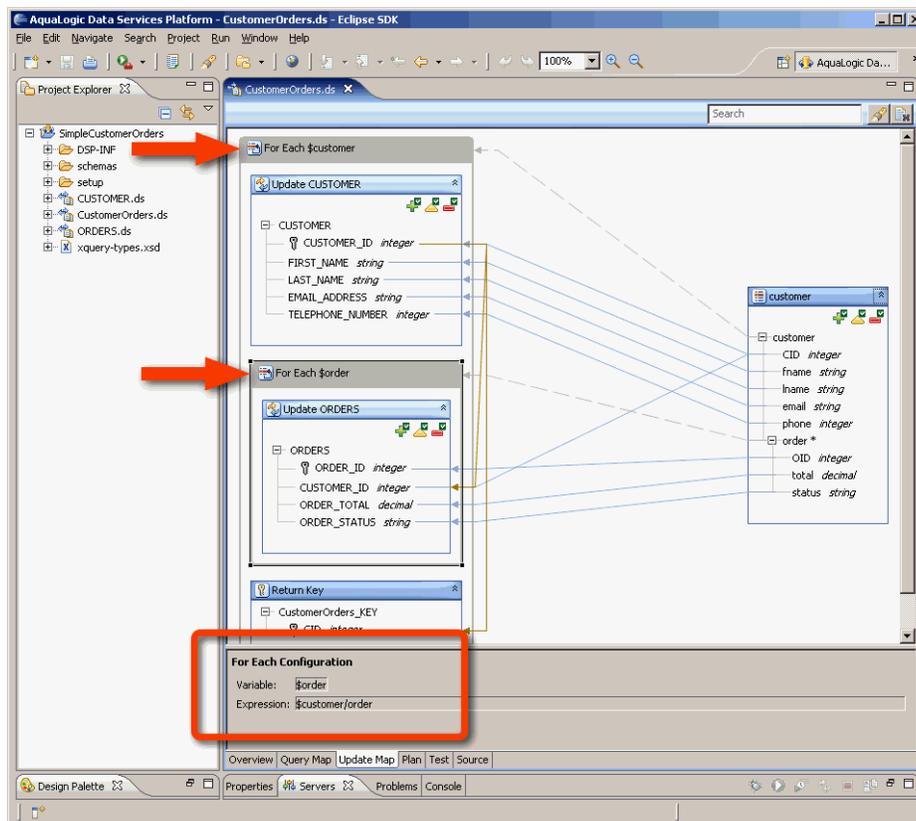


Shows procedure icons indicating the status of update map procedures.

9.1.1.3 For Each Blocks

A for each block loops over elements in the input to the update map procedure. A for each block is associated with a variable and a path expression. The path expression defines the sequence to iterate over and the variable binds to elements in the sequence. The variable may be referenced by expressions inside the for each block.

Figure 9-4 For Each Blocks



For Each Blocks

9.1.1.4 Update Blocks

An update block invokes the primary create, update, or delete procedure of a source data service. It will invoke a procedure every iteration of the for each block that

contains it. The contents of the update block represent the type of the input given to the procedure. Each element and attribute in the update block is assigned a mapping expression that determines what its value will be when the procedure is invoked. You can select an element or attribute to view or change the expression that determines what value it receives when the procedure is invoked (see the example below).

9.1.1.4.1 Procedure icons Like the target box, an update block also has a procedure status icon. Here the icons indicate the ability of the update block to propagate creates, updates, and deletes to the underlying data service. Otherwise, the meaning of the icons is the same as it is for the target box.

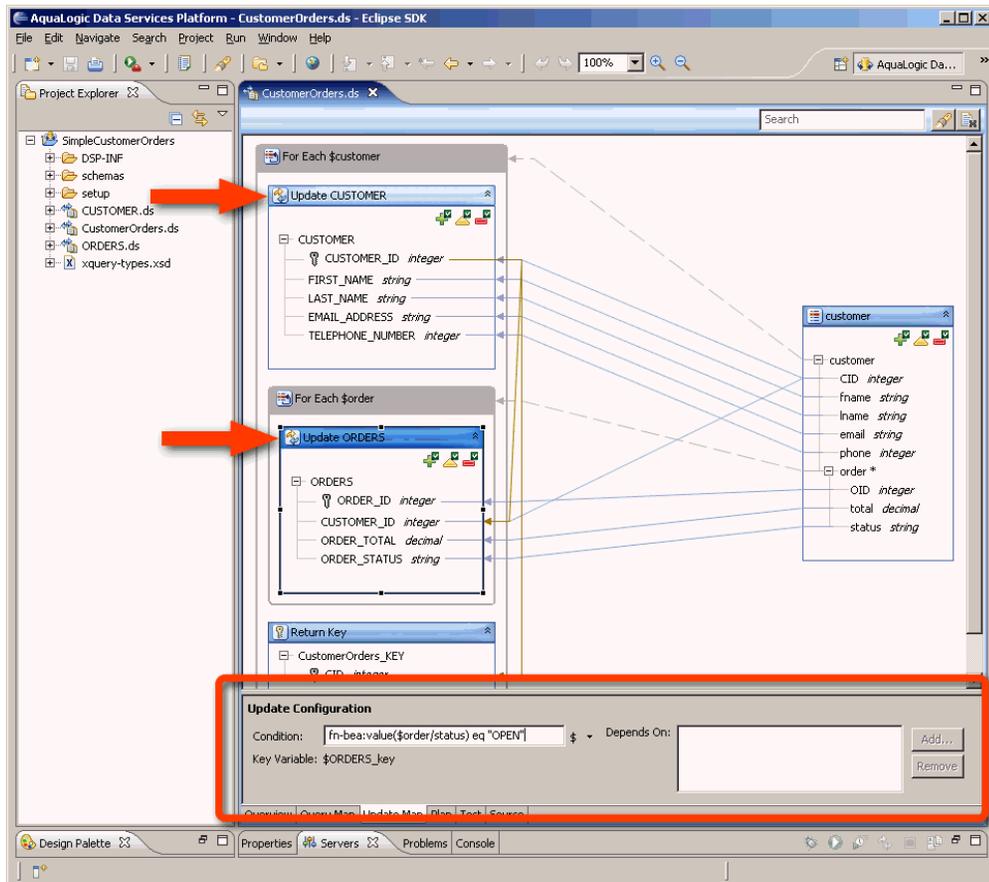
9.1.1.4.2 Output variable A primary create procedure may return a key. If the update block invokes a primary create procedure, it will bind the returned key to the output variable (also referred to as the key variable). The purpose of having the output variable available is for cases when the key value is generated automatically by an external source but is not part of the input. For example, your source data service is a wrapper for a customers relational database table. Say that the key of this table is an attribute **CUSTOMER_ID** which is an auto-generated number. If you are inserting a customer and some orders at the same time, you may need the auto-generated value for **CUSTOMER_ID** to pass to the input of the create procedure for **ORDERS**. When an update block results in the underlying update or delete procedure being invoked, the output variable will bind to the empty sequence.

9.1.1.4.3 Condition An update block can optionally have a condition. The condition is a Boolean expression that determines if the update block should be invoked or not. If there is no condition, then the update block will always be invoked (see the example below).

9.1.1.4.4 Dependencies When two or more update blocks appear as siblings within the same for-each block, it may be desirable to specify dependencies between them (e.g., due to referential constraints), so an update block can also include a list of dependencies. If update block A depends on update block B, update block B will execute before update block A in the case of a create or update operation (and in the opposite order in the case of a delete). Dependencies between update blocks that are not within the same for each block are not necessary, as the execution of an update map is implicitly outside-in.

9.1.1.4.5 Disabling an update block An update block can be disabled so that it will never be invoked at runtime. You can disable an update block by right clicking on it and selecting "Disable" The update block should then appear yellow instead of white to indicate that it has been disabled. Disabling an update block is effectively the same as adding a condition that is always false.

Figure 9-5 Update Blocks

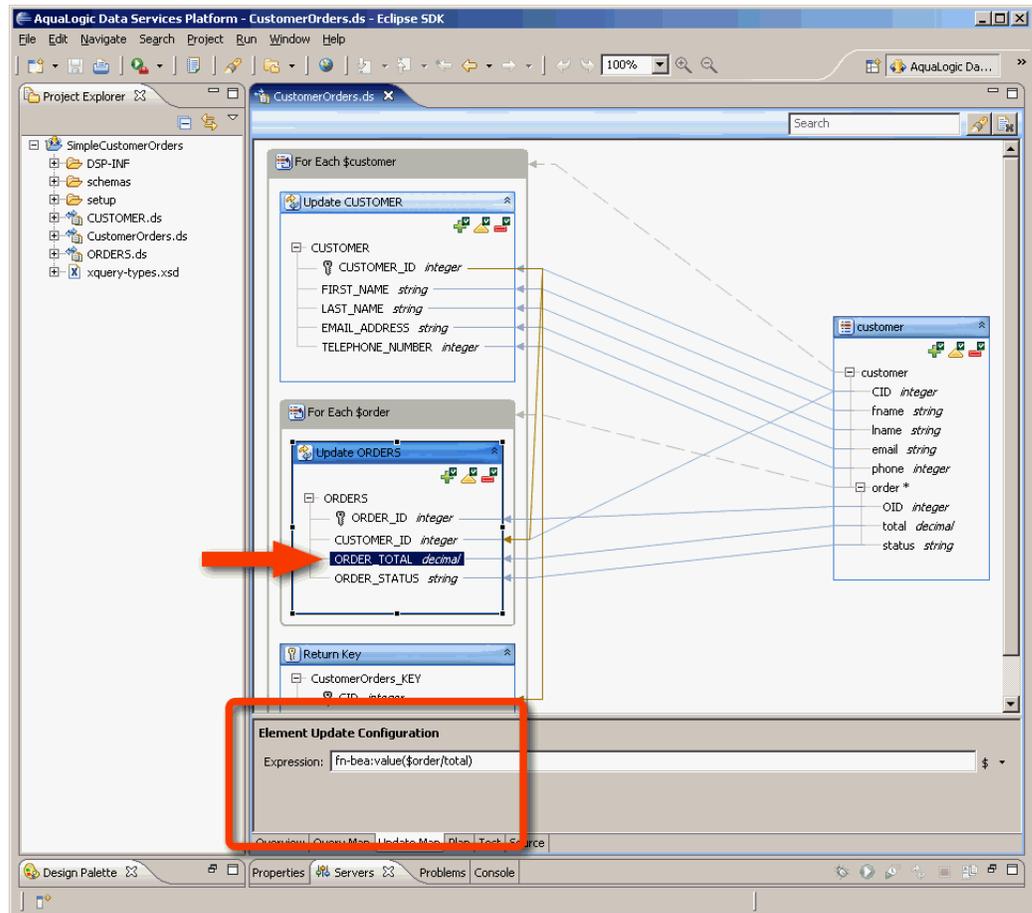


Update Blocks

The images to the left show the update map for the data service CustomerOrders.ds. In the first image, the two orange arrows identify the two update blocks in the map. One update map is for the source data service **CUSTOMER** and the other is for **ORDERS**.

In this case, the ORDERS update block is selected and its details are identified by the orange rectangle (select an update block by clicking on it). We can see that the output variable for this update operation is \$ORDERS_key. The condition is set to fn-bea:value(\$order/status) eq "OPEN" which means that this update block will only be executed when the input element status has the value "OPEN". \$order is a variable that is defined by the for each block containing the update block ("For Each \$order").

Figure 9-6 Update Block Element



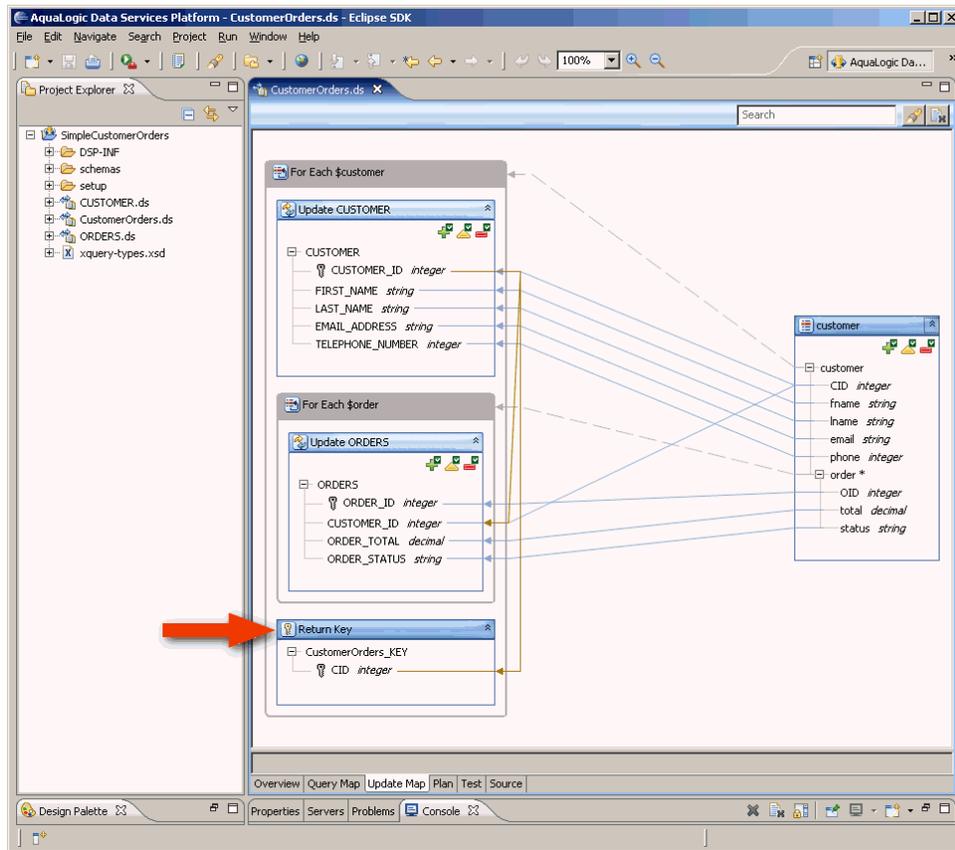
Update Block Element

In the second image, the orange arrow identifies the **ORDER_TOTAL** element of the **ORDERS** update block. The orange rectangle identifies the mapping expression (`$fn-bea:value($order/total)`) for **ORDER_TOTAL** which is displayed because **ORDER_TOTAL** is currently selected. The **ORDER_TOTAL** element will receive the value of the total element when the source data service procedure is invoked.

9.1.1.5 The Return Key Block

The key block describes what will be returned by the update map create procedure. If the data service does not have a key specified, then there will not be a key block and there will never be more than one key block for an update map.

Figure 9–7 The Return Key Block



The Return Key Block

The image to the left shows the update map for the data service `CustomerOrders.ds`. The orange arrow identifies the key block in the update map. The key specified for the **CustomerOrders** data service is the element `CID` so the key block constructs the `CID` element to be returned and uses the output variable of the **CUSTOMER** update block to get the value.

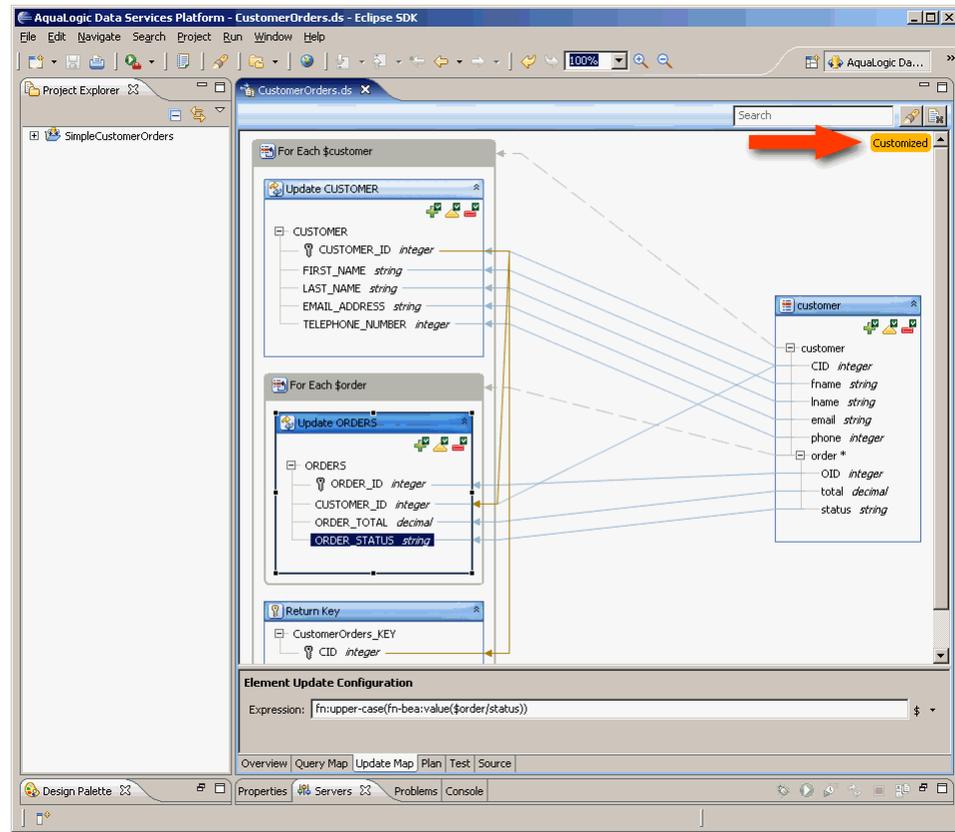
9.1.1.6 Customization

Oracle Data Service Integrator generates a default update map automatically when you create a logical entity data service with a primary read function. This default update map is generated based on the primary read function of the data service. As you change the primary read function, the update map will be regenerated automatically.

There are several ways to customize an update map. See the following topics for more information:

- [Changing a Mapping](#)
- [Removing a Mapping](#)
- [Reverting Customizations](#)
- [Editing XQuery Expressions](#)
- [Adding a Condition to an Update Block](#)

Figure 9–8 Customize



Customize

The image to the left shows the update map for the data service `CustomerOrders.ds`. The orange arrow identifies the "customized" symbol that appears after something in the update map has been changed. In this case, it is the mapping expression for `ORDER_STATUS` that has been modified.

Clicking the **Customized** icon and choosing **View Customization** from the menu displays a dialog showing the current customizations to the update map, including the enabling of update blocks, any conditions that have been set, and any changes to the mappings. This can help you to identify potential problems with the update map that might occur after you make a change to the primary read function, for example. If a problem has been identified, clicking OK accepts the correction and generates a new update map.

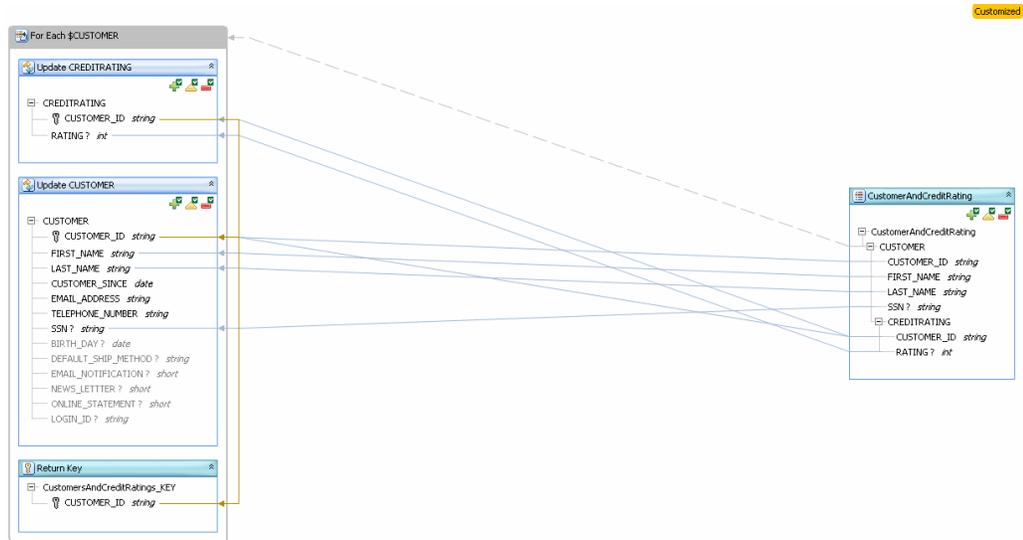
Note that in previous versions of Oracle Data Service Integrator, customizing the update map and then changing the primary read function resulted in the update map no longer being automatically regenerated. In the current version of Oracle Data Service Integrator, the update map is updated incrementally after customizations, as required.

9.2 Changing a Mapping

This section describes how to change a mapping in a default update map generated in Eclipse for WebLogic.

Once you have generated an update map, you can customize it by adding or removing mappings, changing an XQuery expression, adding dependencies, or changing the return type--all in Eclipse for WebLogic.

Figure 9–9 Sample Update Map



This graphic displays a sample update map.

Initially, an update map is generated from the primary read function of a logical data service and changes with the read function.

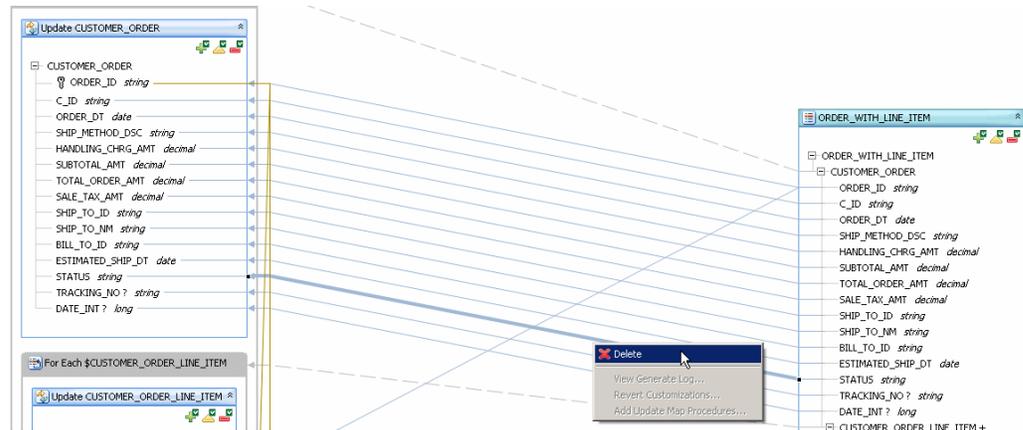
Once you customize an update map, it is no longer linked to the primary read function. If you change the primary read function after customizing the update map, either in a dialog box or in the Source tab, the update map does not change as a result. To re-link the update map to the primary read function, you must revert customizations.

9.2.1 Example

To change a mapping:

1. Click the **Update Map** tab.
2. Right-click an existing mapping line, and choose **Delete**.

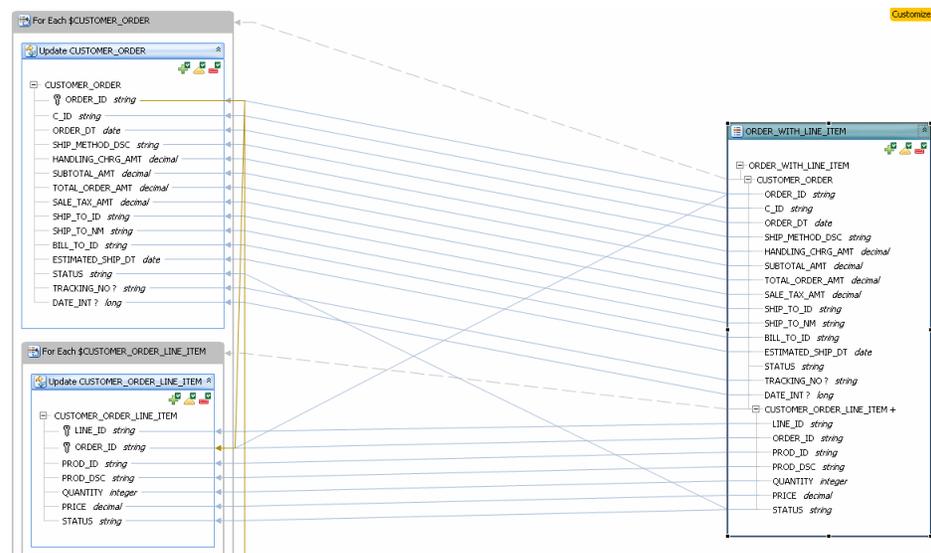
Figure 9–10 Deleting a Status Mapping



This graphic displays the deletion of a status mapping.

3. Drag from an element in the return type on the right to a new element in a data source on the left.

Figure 9–11 Creating a Status Mapping



This graphic displays the creation of a status mapping.

4. Make sure that the **Create**, **Update**, and **Delete** procedure icons (on both the right and left sides) are still enabled and not disabled.
5. Test the new mapping in the Test tab.

The CustomerOrderLineItem Service

In this service, you can draw a new mapping between elements of the same type.

1. Click the **Update Map** tab.

2. Right-click the mapping line between CUSTOMER_ORDER/STATUS in the return type and CUSTOMER_ORDER/STATUS in the update block, and choose **Delete**.
3. Drag a new mapping from CUSTOMER_ORDER_LINE_ITEM/STATUS, the child element in the return type, to CUSTOMER_ORDER/STATUS in the update block. These elements have the same data type.
4. Make sure that the procedure icons are enabled.
5. Click CUSTOMER_ORDER/STATUS on the left, and check the new mapping in the expression editor.
6. Click the **Test** tab.
7. Run a read function, then click **Edit**.
8. Choose a CUSTOMER_ORDER element, then change the value of the first CUSTOMER_ORDER_LINE_ITEM/STATUS child element.
9. Click **Submit**.
10. Run the read function again, then check that the value of CUSTOMER_ORDER/STATUS has changed.

In this example, the child element (CUSTOMER_ORDER_LINE_ITEM) has a multiple cardinality, while the parent element (CUSTOMER_ORDER/STATUS) has a single cardinality. You can see this by checking the XML return type in the Overview tab. By default, the first child element value is read to update the data source. You can override this behavior by adding a dependency or writing a custom update function.

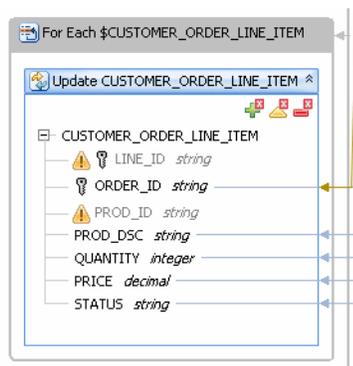
When you map one element to another, be sure that the elements have the same or compatible data types. To be compatible, data types must be in the same type hierarchy in the XML Schema DataTypes specification, such as xs:integer and xs:decimal. These types are cast automatically. If you draw a mapping between two elements of different types and hierarchies, you must cast one data type to the other, using a built-in cast function or a custom cast function.

9.3 Removing a Mapping

This section describes how to remove a mapping from an update map.

An update map shows mappings for required, optional, and key elements. In an update map, optional elements are displayed with a question mark, and key elements with a key symbol. A key element is usually required. If you remove a mapping from a key element, it becomes disabled with a warning icon.

Figure 9–12 Mappings to LINE_ID and PROD_ID Deleted



This graphic displays mappings to LINE_ID and PROD_ID deleted.

Removing a mapping might also cause create, update, or delete procedures to become disabled. However, you can correct either of these conditions, by handling unmapped required values.

If you need to remove a mapping, you can do so in either the update map or query map.

To remove a mapping in the update map:

1. Click the **Update Map** tab.
2. Right-click the mapping line, then choose **Delete**.
3. If the element becomes disabled in the update block on the left, resolve it.

To remove a mapping in the query map:

1. In the Query Map tab, right-click the mapping, then choose **Delete**.
2. Handle any required unmapped values in the update map.

9.4 Reverting Customizations

This section describes how to remove anything you have changed in an update map, regenerating the update map from the primary read function.

You can undo all changes you have made to an update map. Undoing changes creates a new update map, generating it from the primary read function. When you choose Revert Customizations, all changes you have made to the update map are lost, even changes that you have previously saved.

If the update map had errors or warnings that your changes corrected, the errors or warnings will reappear.

To undo changes and generate a new update map:

1. Click the **Update Map** tab.
2. Right-click and choose **Revert Customizations**.
3. Correct any warnings, errors, or disabled procedure icons that appear.

9.5 Adding a Condition to an Update Block

This section describes how to add a condition to an update block in an update map.

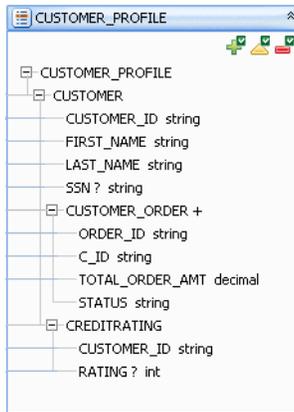
In the update map, you can override an Update block by defining conditions in the expression editor that determine when the block is updated.

A condition is a Boolean expression based on XQuery functions and values defined in the update map, for example:

```
fn-bea:value($CUSTOMER/CUSTOMER_ORDER/TOTAL_ORDER_AMT) > 1000
```

For example, you might have a logical data service with a return type that combines Customer, Order, and CreditRating data. Each customer can have multiple orders and one credit rating.

Figure 9–13 Return Type with Customer, Order, and CreditRating Data



This graphic displays a logical data service with a return type that combines Customer, Order, and CreditRating data.

9.5.1 Example

In the update map, you may want to set a condition that a customer's credit rating can only be updated if the customer places an order with an amount greater than 1000.00.

To set an update map condition:

1. Click the **Update Map** tab.
2. Click the update block on the left that contains the element for which you want to set the condition (for example, the CREDITRATING box for the CREDITRATING/RATING element).

You can now enter a condition in the expression editor.

3. Enter a condition in the Condition box, for example:

```
fn-bea: value ($CUSTOMER/CUSTOMER_ORDER/TOTAL_ORDER_AMT) > 1000.00
```

4. Save the data service.
5. Click the **Test** tab.

The logical data service returns the data in [Figure 9–14](#):

Figure 9–14 Data Returned from Logical Data Service

Name	Value
[-] CUSTOMER_PROFILE	
[-] CUSTOMER_PROFILE	
[-] CUSTOMER	
-- CUSTOMER_ID	CUSTOMER1
-- FIRST_NAME	Jack
-- LAST_NAME	Black
-- SSN	295-13-4119
[-] CUSTOMER_ORDER	
-- ORDER_ID	ORDER_1_4
-- C_ID	CUSTOMER1
-- TOTAL_ORDER_AMT	1283.65
-- STATUS	
[+] CUSTOMER_ORDER	
[-] CREDITRATING	
-- CUSTOMER_ID	CUSTOMER1
-- RATING	600

This graphic displays the data returned by the logical data service.

6. Run a read function, then click **Edit** and attempt to submit a value for the element that has the condition.

When you test the update map, you can only update the **CREDITRATING** data source if **TOTAL_ORDER_AMT** for any of the customer's orders is greater than **1000.00**.

9.6 Editing XQuery Expressions

The following sections describe how to edit XQuery expressions in the expression editor in Eclipse for WebLogic:

- [Section 9.6.1, "Overview"](#)
- [Section 9.6.2, "The fn-bea:value Function"](#)

9.6.1 Overview

You can edit the generated XQuery expressions in an update map using the expression editor.

The update map expression language is a subset of XQuery syntax. In an update map, you can use any of the following XQuery constructs.

Table 9–1 Editing XQuery Expressions

Type	Description	Example
Variable	A variable already defined in a For Each or Update block in the update map. \$\$root is a special predefined variable that refers to the root of the service's XML type.	\$ORDER_WITH_LINE_ITEM \$CUSTOMER
Constant	A numeric, string, or other constant.	"a" "12345"

Table 9–1 (Cont.) Editing XQuery Expressions

Type	Description	Example
Constant Cast	A constant cast to another XSD data type using the parentheses operator.	xsd:date("2007-01-01")
Function	A call to any XQuery function. You can see the built-in and Oracle-provided functions in the Design Palette. You can use a variable, path, or constant as an argument to a function.	fn-bea:value(\$CUSTOMER/FIRST_NAME)
Path	An expression that locates an XML element in a tree using variables, elements, and attributes. The syntax is: \$VARIABLE_NAME /elementName @attributeName	\$ORDER_WITH_LINE_ITEM/CUSTOMER_ORDER/ORDER_ID

Namespace prefixes are declared in the data service's XQuery source, which you can see in the Source tab. If a namespace is only used in the update map, and not in the logical data service, you must declare it. If a namespace cannot be resolved, it is shown with the prefix ns?.

The most common ways you use the expression editor are to:

- Add a constant to an unmapped element
- Cast a constant to an XSD data type, especially to resolve update block elements with no mappings
- Use an XQuery function available in the Design Palette to cast a value
- Use a custom XQuery cast function you have written

9.6.2 The fn-bea:value Function

A mapping between an element in a return type and an element in an update block uses the fn-bea:value function with a path name, for example:

```
fn-bea:value($CUSTOMER/CUSTOMER_ID)
```

An update mapping should always use fn-bea:value, whether Oracle Data Service Integrator auto-generates the mapping or you draw it. If you remove the fn-bea:value function from the expression and simply use an XQuery path expression (\$CUSTOMER/CUSTOMER_ID), the element becomes disabled in the update map and you see this error message:

The expression does not match the expected type for this element

The expression assigned to this element is not valid

Hint: did you forget to use the value function?

The fn-bea:value function is required, because an update map updates a Service Data Object (SDO) and requires a special XML structure called a datagraph that includes a change summary showing both the old and new values. The fn-bea:value function handles the update to the SDO correctly.

If you do not use fn-bea:value, Oracle Data Service Integrator throws an exception when you attempt to update the value.

9.7 Adding an Update Map Procedure

This section describes how to add a create, update, or delete procedure to a logical entity service.

9.7.1 Overview

In a logical entity service, you can add create, update, and delete procedures (called update map procedures) that act on underlying data sources. A procedure is an operation that can have side effects, for example, a create procedure that adds a new record to a database table and returns a key value.

You can create update map procedures visually in Eclipse for WebLogic and have the framework generate XQuery pragma statements and source code, or you can write the source code directly in XQuery or XQSE.

The XQuery pragma statement looks something like this:

```
(::pragma function <f:function kind="create" visibility="public" isPrimary="true"
xmlns:f="urn:annotations.ld.bea.com">
```

This statement defines a create procedure, with public visibility, that is primary. Even though the pragma statement uses the keyword function, the operation you define is a procedure, as you can see from the declaration:

```
declare procedure cus:createCustomerAndAddress($arg as
element(cus:CustomerAndAddress)* ) as element(cus:CustomerAndAddress_KEY)\*
external;
```

This line declares the procedure with the name `createCustomerAndAddress`, defines one argument with the service's return type, and specifies a key as a return value.

9.7.2 Generating Default Procedures

When you generate default update map procedures, they have these parameters and return values:

Table 9–2 Parameters and Return Values

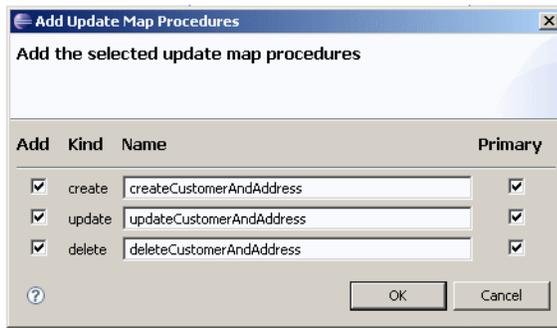
Type	Parameters	Return Value
Create	The service's Return type	The current key, empty if no key is defined
Update	The service's Return type using a changed-element kind	Empty
Delete	The service's Return type	Empty

Before you create update map procedures, especially create procedures, add a key to your service. A primary create procedure must return a key. Primary update and delete procedures require the Return type as an argument; their non-primary equivalents can be written to accept a key instead.

To generate a default update map procedure:

1. Create a key for your service.
2. In the Overview tab, right-click at the left, right, or top, and choose **Add Update Map Procedures**.

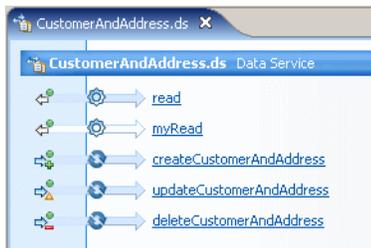
Figure 9–15 Add Update Map Procedures



This graphic displays the Add Update Map Procedures dialog.

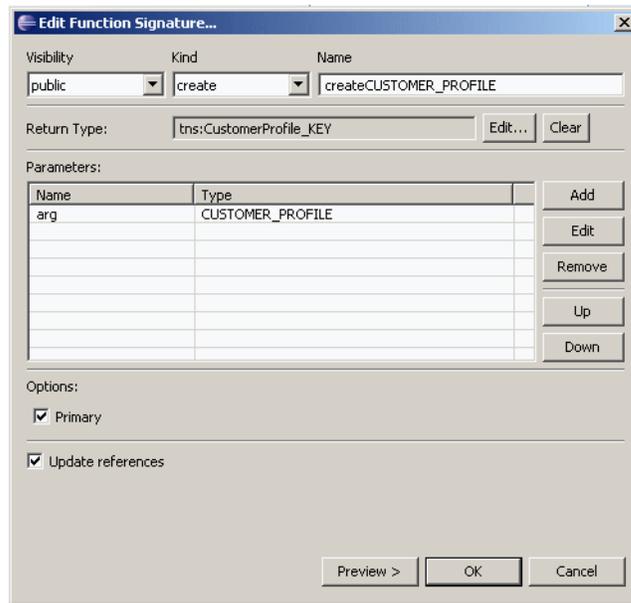
3. Select **Add** to indicate which procedures to add.
4. Add names in the **Name** fields.
5. Mark **Primary** to indicate if each procedure should be primary.
6. Click **OK**.

Figure 9–16 Added Procedures



This graphic displays the procedures added to the data service.

7. In the Overview tab, right-click a procedure name and choose **Edit Signature**.

Figure 9–17 Creating a Procedure Signature

This graphic displays the Edit Function Signature dialog.

8. Make any necessary changes to the procedure signature in the dialog box.

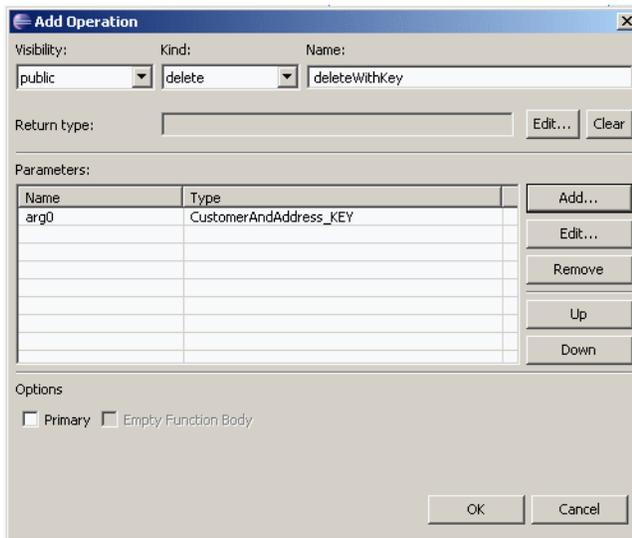
9.7.3 Designing Custom Procedures

You can also create procedures with the arguments and return types you choose. This is useful for procedures in addition to the primary create, update, and delete procedures.

To design custom procedures:

1. Click **Overview**.
2. Right-click at the top, left, or right, and choose **Add Operation**.

Figure 9–18 Add Operation



This graphic displays the Add Operation dialog box.

3. Choose a value for **Visibility**.
4. In the **Kind** field, choose create, update, or delete.
5. In the **Name** field, enter a procedure name.
6. (Optional) At Return Type, click **Edit**. Choose a primitive or complex type, then click **OK**.
7. Click **Add** in the Parameters pane.
8. Choose a primitive or complex type from an XML or XSD file, then click **OK**.
9. In the **Kind** field, choose a value.
 Choose *element* to use the exact XML element you selected as a parameter; *changed-element*, if values in the element must be updated; *schema-element*, if the element must be validated according to an XML schema.
10. Choose a value for **Occurrence**.
11. Click **OK** in both dialog boxes.

9.8 Determining the Scope of a Variable

This section describes how variables may be used when customizing an update map expression. If you are new to update maps, it is recommended that you first read Understanding Update Maps.

9.8.1 Variable Types and Scoping Rules

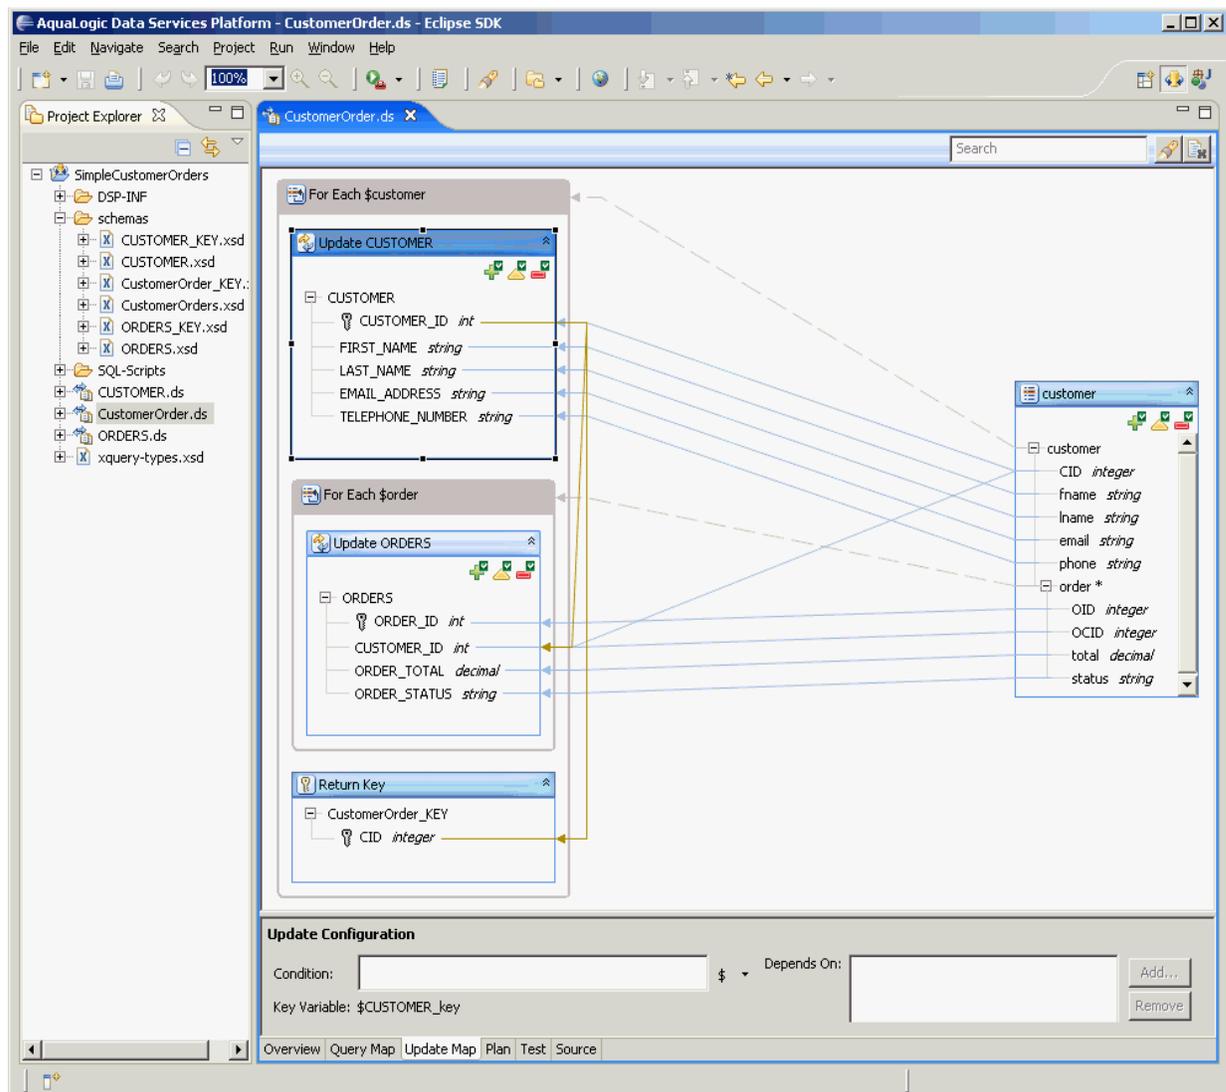
Variables may be defined by a for-each block or by an update block (as an output variable). An output variable may be used in an expression if the expression is contained within an update block that depends on the update block that defines the variable. A for each block variable may be used by an expression if the expression is immediately inside the for-each block that defines the variable. However, if the

defining for-each block contains any other for each blocks which also contain the expression, the variable may not be used in the expression.

9.8.1.1 Example: an update map for a customer-orders data service

The image to the left shows an update map for a logical data service about customers and orders (see Understanding Update Maps for more info). Notice that the update block for customers is selected and the name of its output variable `$CUSTOMER_key` is shown at the bottom of the screen (the output variable is also referred to as the key variable). In this case, `$CUSTOMER_key` can be referenced anywhere from within the update block for orders (e.g. in a mapping expression or in the condition for the update block).

Figure 9–19 Using Variables for Customizing Maps



Using Variables for Customizing

Also notice that we have two for-each blocks that define the variables `$customer` and `$order`. Within the update block for customers, only the `$customer` variable is visible

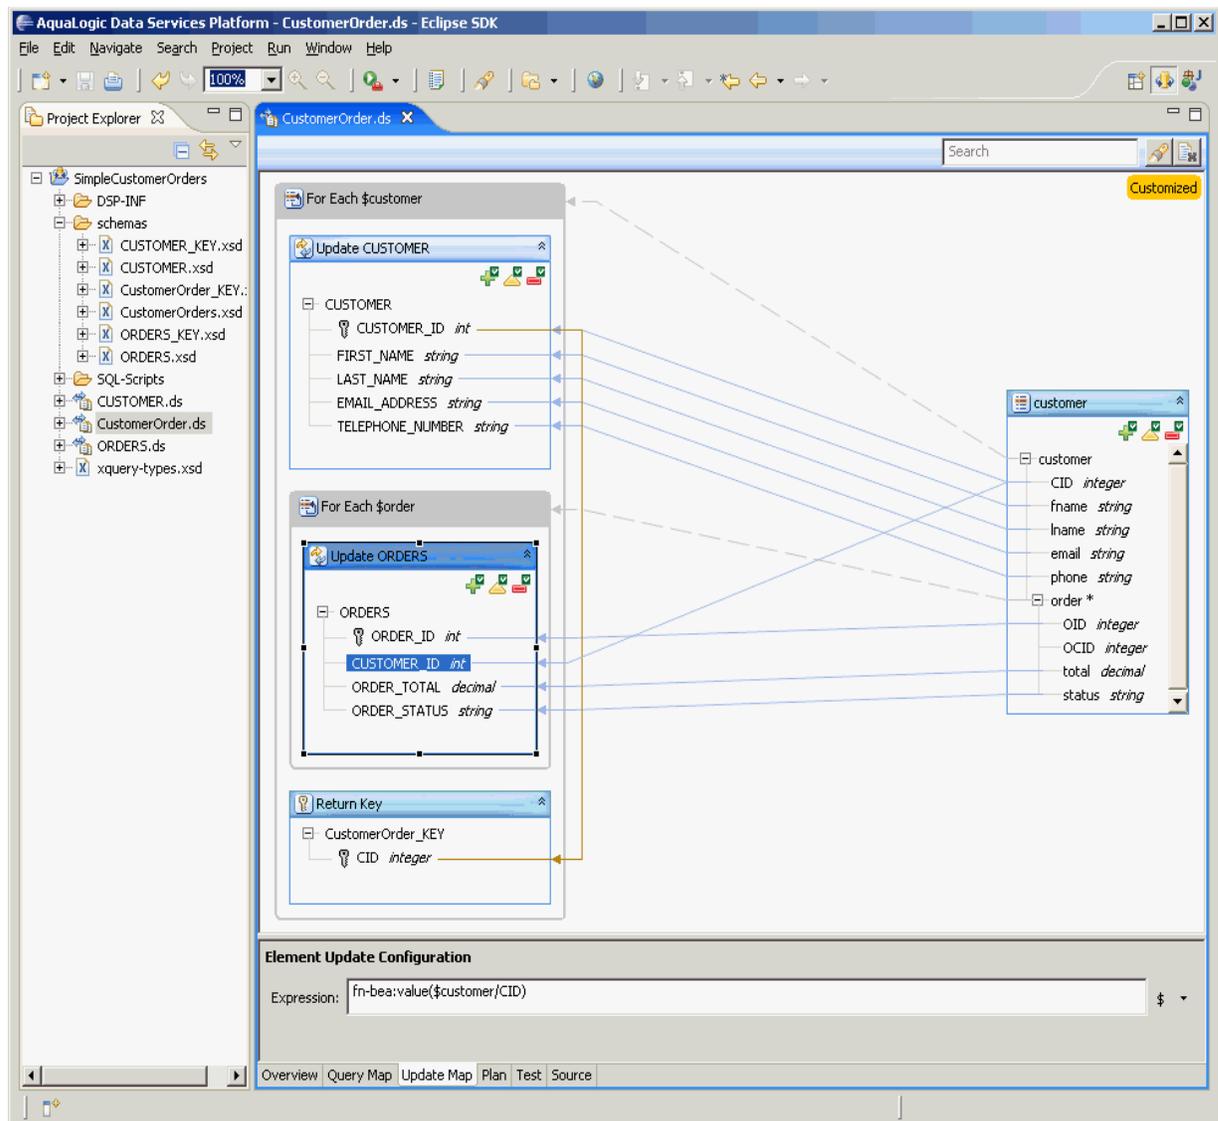
and within the update block for orders only the `$order` variable is visible (with one exception which will be discussed next).

These restrictions are in place to prevent unintuitive or complicated behavior by the update map at runtime. If these rules are too restrictive for your application, you may want to consider using XQSE. However, there is one exception to the for each block variable usage rule. If the variable is used as part of a path expression that references a key value, then usage is valid as long as the expression is within the defining block. It is important to note that updates will be effectively disabled for such "outside" mappings (creates and deletes will still work).

9.8.1.2 Example: an outside mapping to a key value

In the image on the left, the `CUSTOMER_ID` element in the update block for orders is selected and its mapping expression (`fn-bea:value($customer/CID)`) is shown at the bottom of the screen. Normally, the variable `$customer` could not be used within the orders update block since it is not directly contained within the corresponding for-each block. However, since it is being used to reference `CID` it is allowed because `CID` is equivalent to the key value of the customer data service. Updates for this outside mapping will be disabled. That is, if `customer/CID` is modified in the input to the update map procedure, the `CUSTOMER_ID` element in the orders update block will not have the modified value.

Figure 9–20 Outside Mapping



Outside Mapping

9.8.2 Updating Foreign Key Values

The function `fn-bea:coalesce` takes 1 or more arguments and simply returns the first argument that is not empty. The function `fn-bea:coalesce-equal` works the same way except that it additionally checks that all non empty arguments are equal. If it finds that any two non empty arguments are not equal, it will throw an exception at runtime. The automatically generated update map may use `fn-bea:coalesce-equal` for the mapping expression for foreign key values if it can be inferred from the target data service that the values should always be equal.

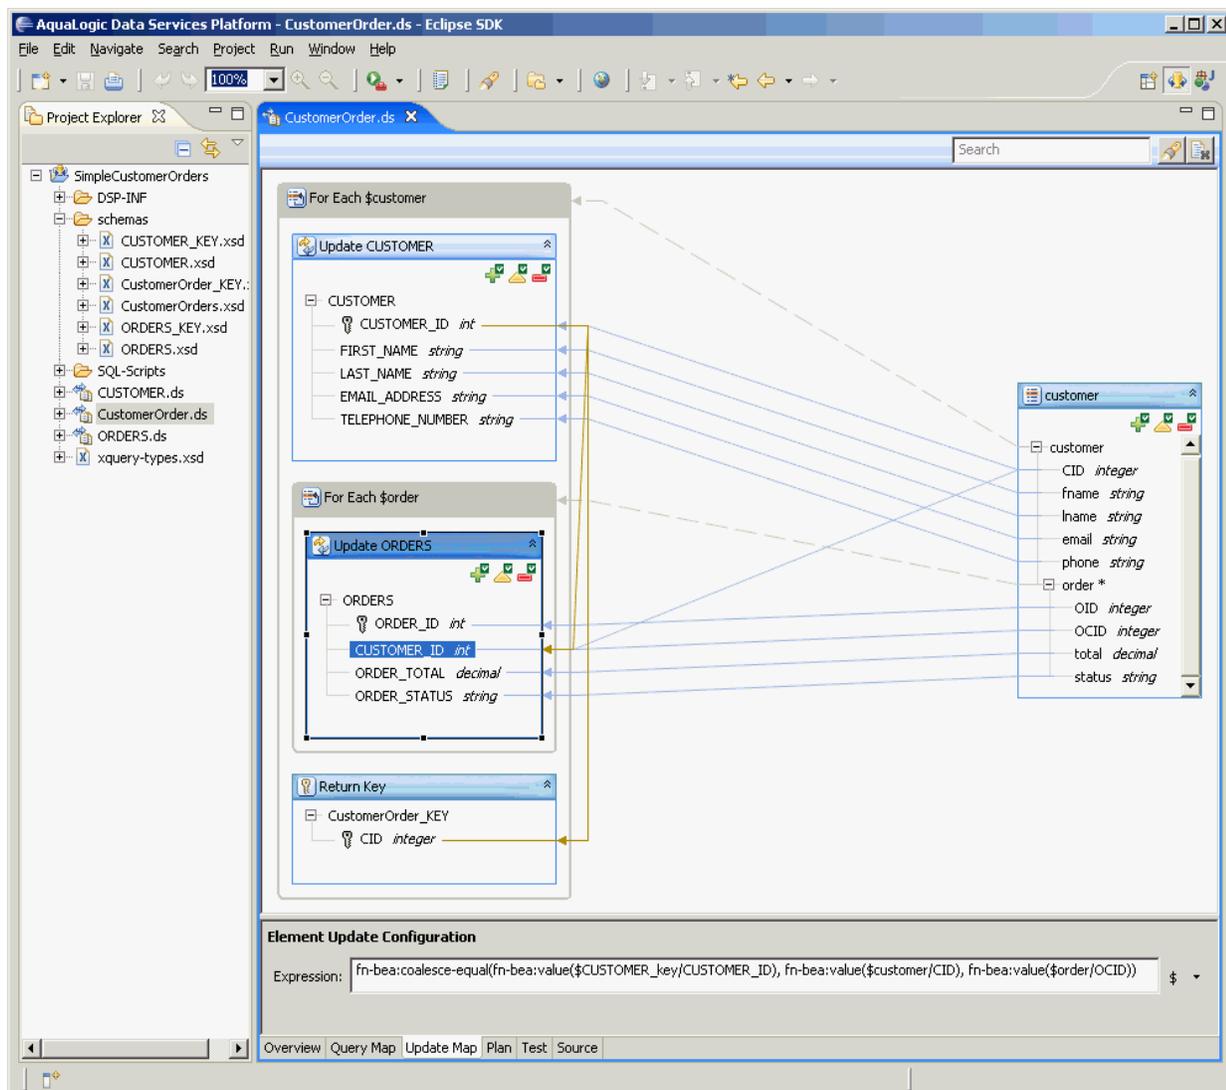
If an argument to `fn-bea:coalesce-equal` contains an a path expression that falls under the exception to the for each block variable rule mentioned above, then updates will be disabled for the entire expression containing `fn-bea:coalesce-equal`. If your automatically generated update map is in this situation and you wish to be able to

update the foreign key value, you can simply remove the argument that contains the offending mapping. (See the example that follows)

9.8.2.1 Example: coalesce-equal

In the first image on the left, the CUSTOMER_ID element in the update block for orders is selected and its mapping expression (fn-bea:coalesce-equal(...)) is shown at the bottom of the screen. When the update map is used to create a customer with orders, the value for CID may not be known as it may be auto generated by an underlying relational database. This means that \$customer/CID and \$order/OCID may be empty. In this case the generated key value will be returned and the orders will get the value for CUSTOMER_ID via fn-bea:value(\$CUSTOMER_key/CUSTOMER_ID). If the update map is used to update or delete customers and orders, \$CUSTOMER_key/CUSTOMER_ID will be empty but \$customer/CID and \$order/OCID should not be.

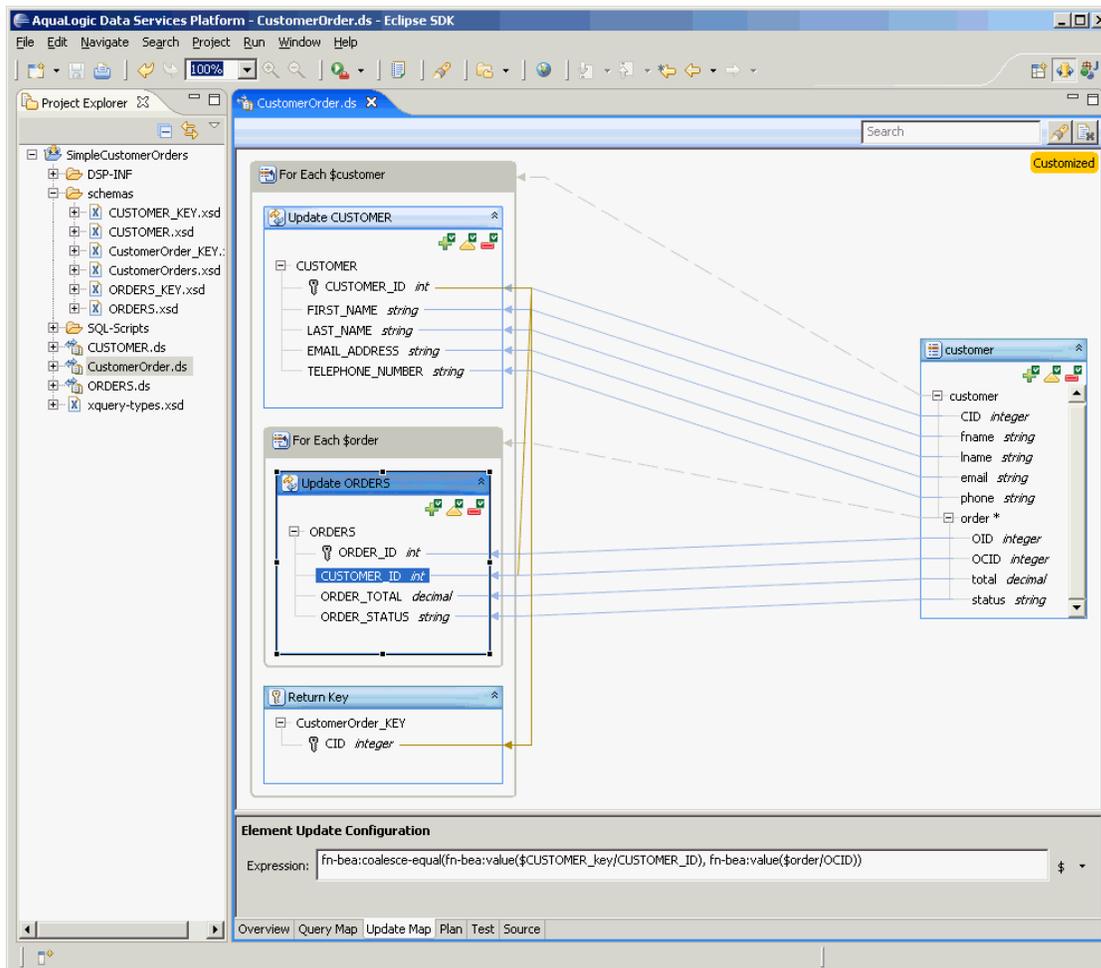
Figure 9–21 Updating Foreign Key Values: Customer ID



Updating Foreign Key Values: Customer ID

This mapping contains an outside reference to the variable `$customer` so updates will be disabled for `customer/order/OCID` in the input to the update map procedure. To enable updates to `OCID`, we can remove the outside mapping. The second image on the left shows what the update map looks like after this modification.

Figure 9–22 Outside Reference Mapping



Outside reference mapping

9.9 Reference

This section provides reference information for managing update maps.

9.9.1 Update Map Functions

The following functions are useful in update map expressions (e.g. update block conditions and mapping expressions).

```
fn-bea:coalesce($arg ... as xdt:anyAtomicType) as xdt:anyAtomicType?
```

The function `fn-bea:coalesce` takes 1 or more arguments and returns the first that is not empty.

```
fn-bea:coalesce-equal($arg ... as xdt:anyAtomicType) as xdt:anyAtomicType?
```

The function `fn-bea:coalesce-equal` takes 1 or more arguments and returns the first that is not empty. If any of its non empty arguments are not equal then it will throw an exception at runtime. (see also [How to update a foreign key values mapped using fn-bea:coalesce-equal](#))

```
fn-bea:value($arg as item()?) as xdt:anyAtomicType?
```

The function `fn-bea:value` is essentially the same as `fn:data` except that it additionally indicates to the Oracle Data Service Integrator runtime that the variable `$arg` may bind to a `changed-element()` depending on the context (like in the case of an update procedure driven by the update map). As a general rule, `fn-bea:value` should always be used in place of `fn:data` in update map expressions since they should be written to support all three flavors of update procedures (create, update, and delete).

```
fn-bea:ambiguous($arg ... as xdt:anyAtomicType) as xdt:anyAtomicType?
```

The function `fn-bea:ambiguous` may appear in the default update map when there are multiple target values (on the right) that map to the same source value (on the left). For example, if the same value is projected (that is, returned) more than once in the primary read function, this may result in `fn-bea:ambiguous` being used in the corresponding update map by default. It is expected that the user will manually remove the call to `fn-bea:ambiguous` when resolving the ambiguous mapping. For example, the user may choose to pick one of the arguments in the `fn-bea:ambiguous` call to be the new mapping expression and disregard the others.

9.10 How To

This section describes procedures for managing update maps.

9.10.1 How To Recognize When Something is Wrong

The following sections describe why an update map might appear disabled and point you to solutions:

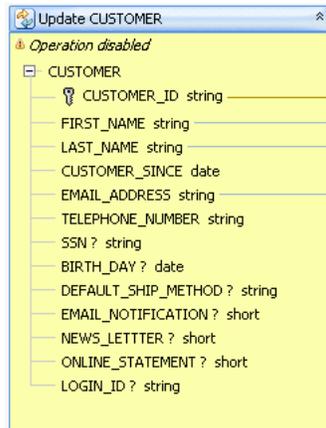
- [Section 9.10.1.1, "Understand the Symptoms"](#)
- [Section 9.10.1.2, "Check the Problems Tab"](#)
- [Section 9.10.1.3, "Resolve Errors and Warnings"](#)

9.10.1.1 Understand the Symptoms

The signs of a disabled update map appear on the update map itself, in the **Generate Log**, and in the **Problems** tab.

In the update map, you may see disabled (or yellow) update blocks. When an update block is completely or partially disabled, updates do not occur in the data source the block maps to.

Figure 9–23 Disabled Update Block



A totally disabled update block

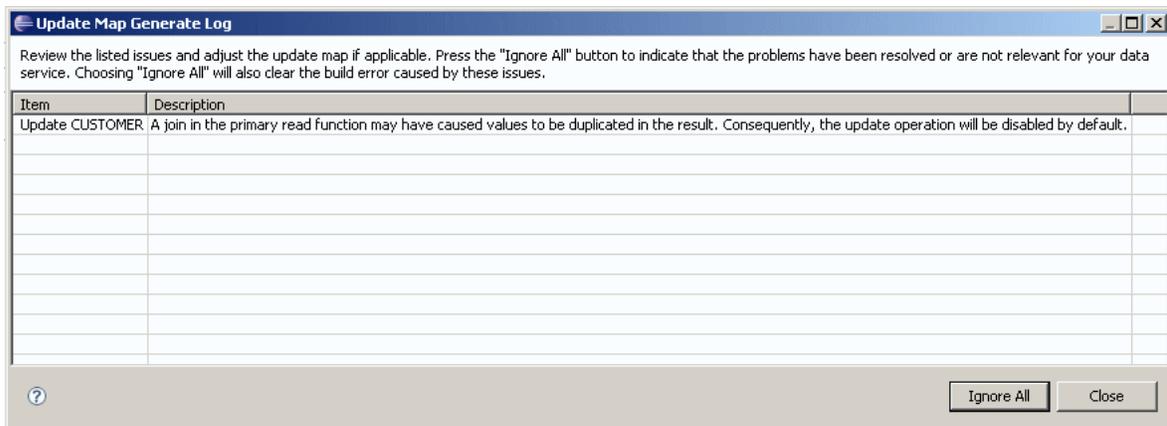
9.10.1.1.1 A Disabled Update Block An update map procedure that is disabled has a yellow or red status indicator at the upper right.

9.10.1.1.2 Disabled Procedure Icons You might also see a message with a link to view the Generate Log.

9.10.1.1.3 A View Generate Log Message Clicking the link displays the Update Map Generate Log window.

9.10.1.1.4 Update Map Generate Log Here is an example of an Update Map Generate log.

Figure 9–24 Update Map Generate Log



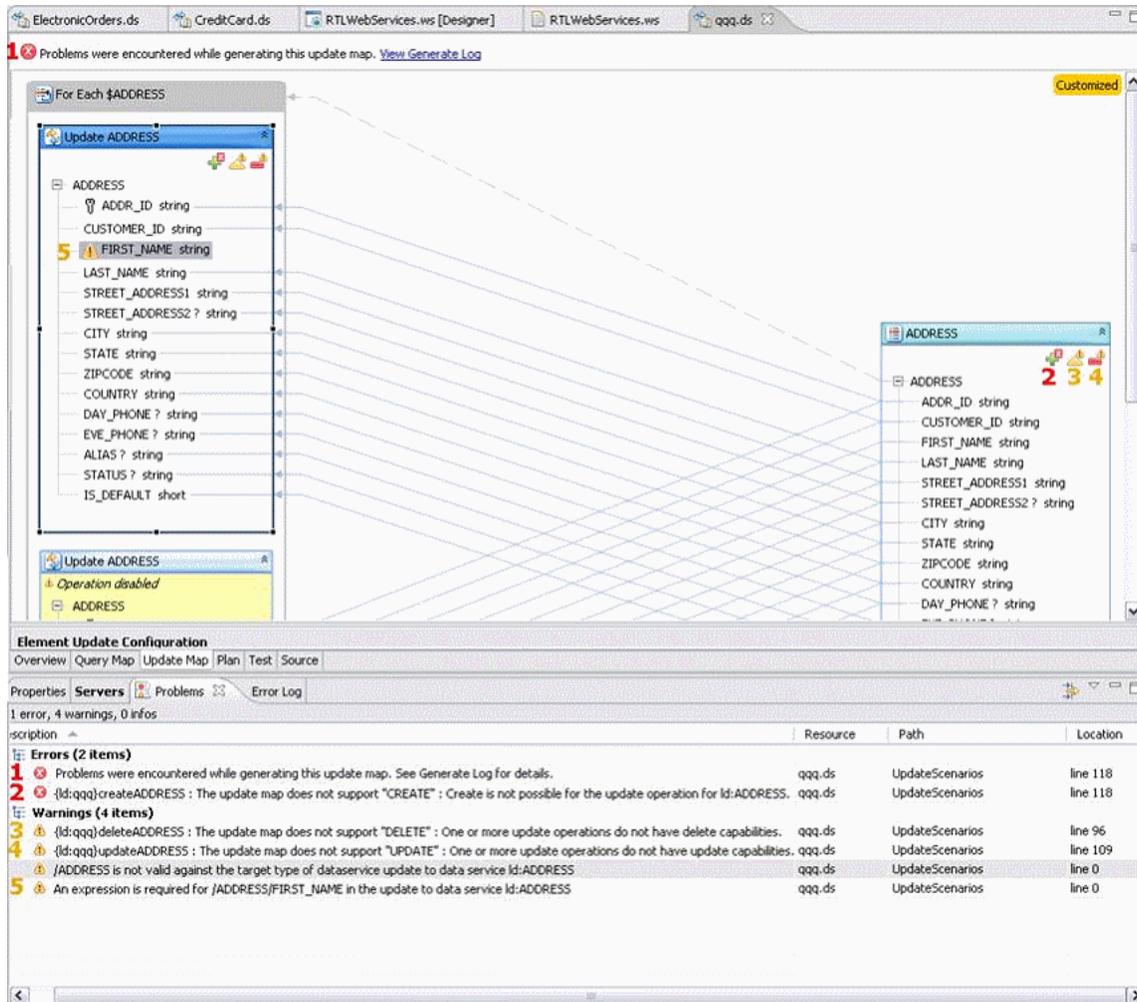
Example of an Update Map Generate log

9.10.1.2 Check the Problems Tab

If you see disabled procedure icons or other symptoms, you should also check the Problems tab for detailed Error and Warning messages. The Problems tab shows errors

and warnings that the View Generate Log message does not. For example, this update map shows two errors and three warnings.

Figure 9–25 Check the Problems Tab for a Disabled Update Map



Checking for problems.

Errors prevent you from deploying the update map to the Oracle Data Service Integrator server and testing it. Warnings tell you that something is not supported in the update map, but the update will proceed.

To sort the Problems tab, as shown above:

1. Click the **Problems** tab.
2. Click the triangle icon at the upper right, and choose **Sorting**.
3. Sort first by **Resource**, then by **Severity** and **Description**.

9.10.1.3 Resolve Errors and Warnings

You may have a valid reason to use a certain logical data service design that initially generates an update map with constraints. This is fine. You can find workarounds and resolve most disabled update map conditions.

9.10.1.3.1 Disabled Update Blocks When you encounter a disabled, or yellow, update block, you can right-click it and choose Enable. The most likely reasons an update block is disabled are shown below.

Once you enable the update block, you will likely see:

- Elements that have warnings or are completely disabled
- Disabled Create-Update-Delete procedure icons

9.10.1.3.2 Disabled Procedure Icons When you see disabled procedure icons, check the update blocks on the left. The procedure icons in the return type on the right naturally result from those on the left.

In general, the status indicators for update map procedures are:

- Green if the update map will work at run time as you have designed it, even if parts of it are disabled
- Yellow if some parts of the update map will work at run time, but you might see run-time errors on other parts
- Red if the update map will not work at run time

If you want to correct an update map before run time, a red or yellow status indicator on the left can have any of the following meanings.

Table 9–3 Status Indicators

Status	Type of Procedure	Meaning	Solution
Red	Create, Update, Delete	The data service does not have a primary procedure of that type.	Create a primary procedure (Overview tab, right-click, Add Update Map Procedures, select Primary)
Red	Update, Delete	The data service does not have a key.	Create Logical Data Service Keys
Red	Create	The update block has missing mappings or mappings of the wrong data type	Understand Mappings with Different Data Types Handle Unmapped Required Values
Red	Create	The return type contains non-element or non-attribute XML items that are not allowed.	Handle an Unsupported Node Constructor Error
Red	Create, Update, Delete	The update block references a variable from another disabled update block.	Right-click the disabled block, and choose Enable.
Red	Update, Delete	The data service has a key, but one or more key fields have missing mappings, mappings of the wrong data type, or mappings to invalid items in the return type.	Handle Unmapped Required Values
Yellow	Update, Delete	The update block has missing mappings, mappings of the wrong data type, or mappings to invalid items in the return type.	Understand Mappings with Different Data Types Handle Unmapped Required Values

9.10.2 How To Understand Mappings with Different Data Types

The following sections describe casts between elements of different data types in an update map:

- [Section 9.10.2.1, "Overview"](#)

- [Section 9.10.2.2, "Built-In Cast Functions"](#)
- [Section 9.10.2.3, "Custom Cast Functions"](#)

9.10.2.1 Overview

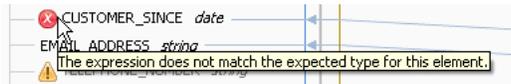
In an update map, you may need to map elements of different data types between a return type and an underlying data source.

For example, a return type might contain an `xsd:dateTime` element that maps to an `xs:date` element in the data source. When data types differ, you need to cast between them in order to enable the update map. Type differences occur because a logical data service design can differ from actual physical data sources or because data types used by an underlying data source are unknown at design time.

When the update map is first generated, the element in the data source has no mapping and a warning icon.

If you draw a mapping line in Update Map view, from the `xsd:dateTime` value in the return type to the `xsd:date` value in the update block, the element becomes disabled.

Figure 9–26 An Error Due to Data Type Mismatch



Element Data Type Mismatch

You can fix this type of error by using different techniques to cast, according to the data types you are casting.

First, review the built-in datatypes chart in the XML Schema Datatypes specification to understand the hierarchies of data types used in XML Schema. The type `xs:string` and its subtypes belong to one type hierarchy, and the type `xs:decimal` and its subtypes belong to another.

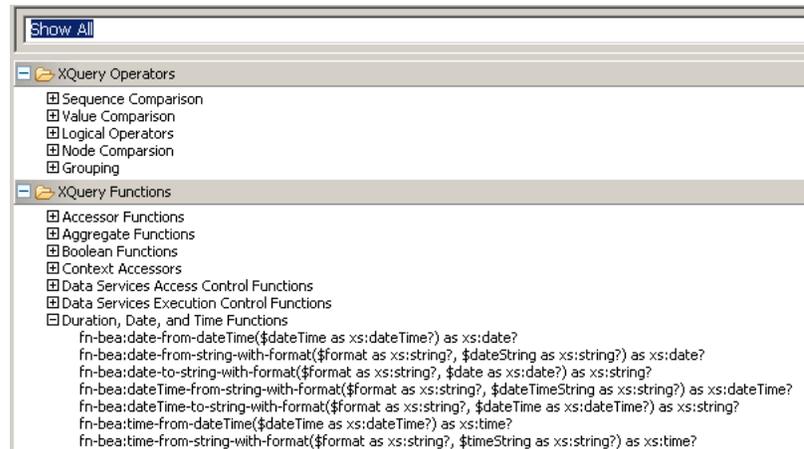
Casts between elements of different types are handled in one of three ways:

1. **Type promotion.** If both data types are in the same type hierarchy and the cast moves up the hierarchy, Oracle Data Service Integrator casts them implicitly. This is known as type promotion. For example, `xs:token` is promoted to `xs:string` and `xs:integer` is promoted to `xs:decimal`. Implicit casts are implemented in Oracle Data Service Integrator according to the XQuery 1.0 specification.
2. **Built-in cast function.** If the types do not use a cast up the same type hierarchy and type promotion does not occur, you can use a built-in XQuery function available from the Design Palette.
3. **Custom cast function.** If a built-in XQuery function is not available, you can write your own custom cast function in the Source tab of your primary logical data service or in a specialized library data service that performs casting.

9.10.2.2 Built-In Cast Functions

If a built-in function provides the cast you need, you can simply drag it from the Design Palette to the expression editor and enter argument values.

Figure 9–27 Built-in XQuery Casting Functions



Built-in XQuery Casting Functions

9.10.2.3 Custom Cast Functions

Before you write a custom XQuery cast function, make sure that XQuery allows the cast you want to perform. Check the casting section in the XQuery 1.0 specification to understand the rules for casting between types in XQuery, especially the chart that describes casting between primitive types.

Remember these general guidelines:

- The primitive type chart shows which casts can be performed between primitive types. For example, an integer (such as 44) can always be cast to a string ("44"). However, a string can only be cast to an integer in some cases. The string "55" can be cast to the integer 55, but the string "hello" cannot be cast to an integer.
- If both the source and target types are derived from the same primitive type, you can cast between them.
- If the source and target types are derived from different primitive types, you are casting across the type hierarchy. In general, you need to cast the source type up the hierarchy to its primitive type; then, cast from the primitive type of the source to the primitive type of the target; and last, cast from the primitive type of the target to the target type (see the rules in the XQuery 1.0 specification).

Once you write the cast function, you can test it in Eclipse for WebLogic, before you run it with a client application.

9.10.3 How to Cast Using a Built-in XQuery Function

The following sections describe how to use a built-in XQuery function to cast values of different data types in an update map.

9.10.3.1 Example

You can cast an element from one data type to another using a built-in XQuery cast function when:

- Type promotion does not occur.
- The data comes from a variable or an other source that is not a constant

- A built-in function that performs the cast you want is available in the Design Palette.

To cast using a built-in XQuery function:

1. Click the **Update Map** tab.
2. Click the disabled element in an update block on the left.

In the expression editor, you see an expression that uses `fn-bea:value()` to map from the return type on the right, for example:

```
fn-bea:value($CUSTOMER/CUSTOMER_SINCE)
```

This expression represents a `dateTime` value coming from the return type.

3. Open the Design Palette
Window > Show View > Design Palette
4. Expand XQuery Functions, then a category (for example, Duration, Date, and Time Functions).
5. Drag the function you want to the expression editor (for example, `fn-bea:date-from-dateTime`), leaving the existing expression there.
6. If feasible, use the existing expression as an argument to the function, for example:

```
fn-bea:date-from-dateTime( fn-bea:value($CUSTOMER/CUSTOMER_SINCE) )
```

Here the original value is used as the `$dateTime` argument to `fn-bea:date-from-dateTime()`.

7. Test the update map cast to make sure it works as you expect.

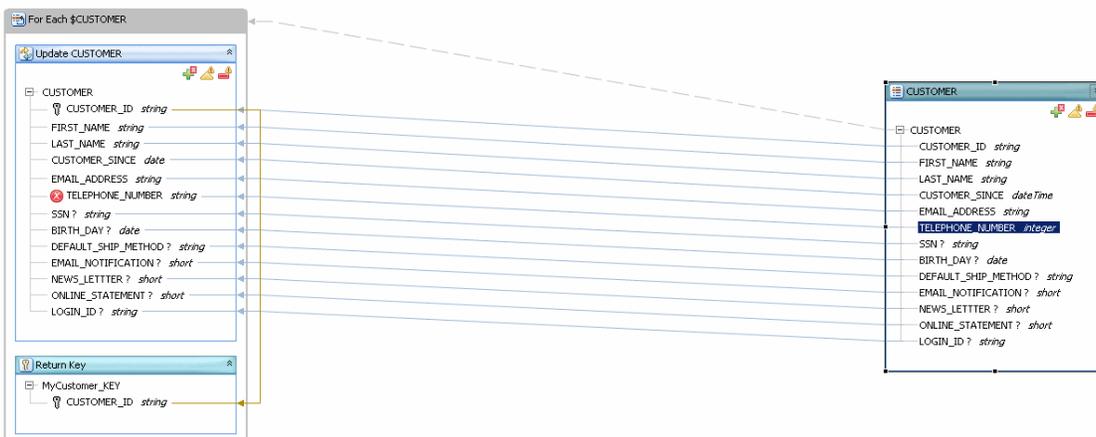
9.10.4 How To Cast Using a Custom XQuery Function

This section describes how to write a custom XQuery function to cast between elements of different data types in an update map.

9.10.4.1 Example

An example of a custom XQuery cast function is one that casts from integer to string. Suppose the logical data service's return type uses `xsd:integer` for the **TELEPHONE_NUMBER** element, while the underlying data source uses `xsd:string`.

Figure 9–28 Mapping from Integer to String



Mapping from Integer to String

The mapping between the two `TELEPHONE_NUMBER` elements is initially disabled. The value from the return type is something like 4155551212, which can easily be converted between `xsd:integer` and `xsd:string`. Check the type casting chart in the XQuery 1.0 specification to make sure the cast you want to perform is allowed.

Caution: When you test the cast function, you also need to perform the opposite cast (in this case, `xsd:string` to `xsd:integer`).

To write a custom XQuery cast function:

1. Click the Source tab.
2. Write an XQuery function that takes an argument of the data type you are casting from and returns a value of the data type you are casting to, for example:

```
declare function tns:intToString($theint as xs:integer) as xs:string {
  xs:string($theint)
};
```

Assign your function to an XML namespace your logical data service uses. Be sure both the parameter and return type are valid XML Schema data types. Then, write a statement that performs the cast.

3. In the Update Map tab, click the element in the data source on the left.

At this point, the element is disabled: . Its value is taken from the return type, so its XQuery expression looks something like this:

```
fn-bea:value($CUSTOMER/TELEPHONE_NUMBER)
```

Remember that the value from the return type is an `xs:integer`.

4. Add your new cast function, using the existing expression as its argument, for example:

```
tns:intToString(fn-bea:value($CUSTOMER/TELEPHONE_NUMBER))
```

At this point, the update map should be completely enabled.

5. If the disabled icon on the element does not disappear immediately, click another element in the update map.
6. Test the update map cast to make sure it works as you expect.

9.10.5 How To Test an Update Map Cast

This section describes how to test a cast between elements of different data types in an update map.

9.10.5.1 Example

The easiest way to test an update map cast function is to use Read-Edit-Submit from the Test tab in Eclipse for WebLogic.

Suppose you are casting from `xs:integer` to `xs:string`. To test the cast function, you need to retrieve data from the data source as `xs:string` and display it in the Test tab as `xs:integer`, so you also need to cast in the reverse direction. The primitive types casting chart in the XQuery 1.0 specification shows that you can always cast from `xs:integer` to `xs:string`, but you can only cast from `xs:string` to `xs:integer` in some cases.

To test an update map cast using Read-Edit-Submit, you first edit the source code of the primary Read function to do a comparable cast when the data is read from the data source. For example, suppose you want to cast from `dateTime` to `date` during an update. To test, you must first cast the date value to `dateTime` when you read it from the data source.

Before you use this test method, check the casting chart in the XQuery specification to make sure the XQuery cast you want to perform works in both directions. In the example given here, the cast is from `xs:dateTime` to `xs:date` in the update map and from `xs:date` to `xs:dateTime` in the primary Read function. Both casts must be valid in XQuery.

1. Click the **Source** tab.
2. Locate the primary **Read** function, which looks something like this:

```
declare function tns:read() as element(cus:CUSTOMER)*{
  for $CUSTOMER in cus1:CUSTOMER()
  return
    <cus:CUSTOMER>
      <CUSTOMER_ID>{fn:data($CUSTOMER/CUSTOMER_ID)}</CUSTOMER_ID>
      <FIRST_NAME>{fn:data($CUSTOMER/FIRST_NAME)}</FIRST_NAME>
      <LAST_NAME>{fn:data($CUSTOMER/LAST_NAME)}</LAST_NAME>
      <CUSTOMER_SINCE>{fn:data($CUSTOMER/CUSTOMER_SINCE)}</CUSTOMER_SINCE>
      <EMAIL_ADDRESS>{fn:data($CUSTOMER/EMAIL_ADDRESS)}</EMAIL_ADDRESS>
      <TELEPHONE_NUMBER>{fn:data($CUSTOMER/TELEPHONE_NUMBER)}
    </TELEPHONE_NUMBER>
      <SSN?>{fn:data($CUSTOMER/SSN)}</SSN>
      <BIRTH_DAY?>{fn:data($CUSTOMER/BIRTH_DAY)}</BIRTH_DAY>
      <DEFAULT_SHIP_METHOD?>{fn:data($CUSTOMER/DEFAULT_SHIP_
METHOD)}</DEFAULT_SHIP_METHOD>
      <EMAIL_NOTIFICATION?>{fn:data($CUSTOMER/
EMAIL_NOTIFICATION)}</EMAIL_NOTIFICATION>
      <NEWS_LETTTER?>{fn:data($CUSTOMER/NEWS_LETTTER)}</NEWS_LETTTER>
      <ONLINE_STATEMENT?>{fn:data($CUSTOMER/ONLINE_STATEMENT)}
    </ONLINE_STATEMENT>
      <LOGIN_ID?>{fn:data($CUSTOMER/LOGIN_ID)}</LOGIN_ID>
    </cus:CUSTOMER>
  };
```

3. Locate the element you want to cast and add a XQuery cast expression to it. For example, this casts an `xs:date` to an `xs:dateTime` in the **CUSTOMER_SINCE** element:

```
<CUSTOMER_SINCE>{ xs:dateTime(fn:data($CUSTOMER/CUSTOMER_SINCE)) }
</CUSTOMER_SINCE>
```

To cast an `xs:string` to an `xs:integer` in **TELEPHONE_NUMBER**, enter this:

```
<TELEPHONE_NUMBER>{xs:integer( fn:data($CUSTOMER/TELEPHONE_NUMBER))
}</TELEPHONE_NUMBER>
```

4. Click the **Test** tab.
5. At **Select Operation**, choose the service's primary **Read** function and click **Run**.
In the Result pane, you might see that the values have been cast, if the new type looks different.
6. Click a customer record, then **Edit**.
7. Change one of the values you have just cast.

If you are working with `xs:date` and `xs:dateTime`, change the date portion of the value, rather than the time. The time is truncated when you store the value in the data source as an `xs:date`. When you read it back as an `xs:dateTime`, it looks like 00:00:00.

8. Click **Submit.**

You should see this message:

Data has been submitted

9. Click **Run again to verify the change.**

9.10.6 How To Handle Disabled Procedures in Underlying Data Sources

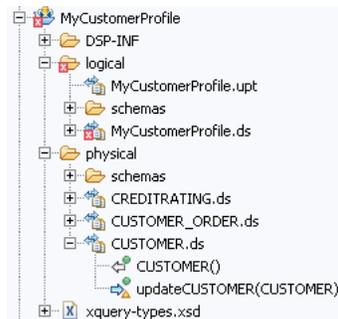
This section explains how to enable an update map for a logical data service when an underlying data source has disabled procedures.

9.10.6.1 Check the Data Sources

If a Create, Update, or Delete procedure is disabled in a data source that your logical data service uses, part of the update map is disabled as well. Specifically, the update block that maps to the data source is disabled.

For example, you might have a physical data service that is missing a Create, Update, or Delete procedure.

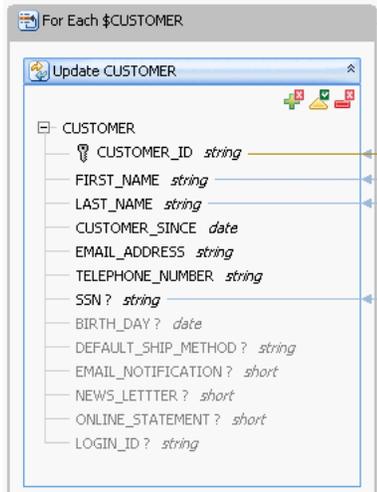
Figure 9–29 Physical Data Service with No Create or Delete Procedure



Physical Data Service with No Create or Delete Procedure

As a result, the update block that maps to this data source has its Create and Delete procedures disabled.

Figure 9–30 Update Block with Disabled Create and Delete Procedures



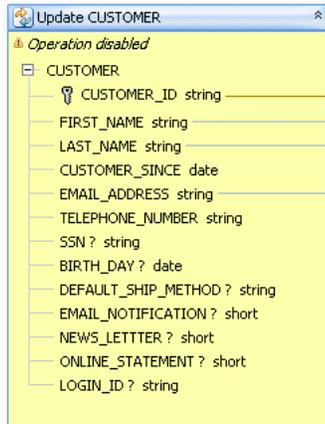
Update Block with Disabled Create and Delete Procedures

9.10.6.2 Resolve the Disabled Procedures

If you do not need to use the procedures that are disabled in the underlying data source, you can disable the entire update block:

1. Click **Update Map**.
2. Right-click the update block, and choose **Disable**.

Figure 9–31 Disable the Update Block



Disable the Update Block

Disabling the block might also disable procedures or key elements in other blocks.

3. Resolve any mappings that become disabled.

9.10.6.3 Add or Enable Procedures in the Underlying Data Source

You can also enable procedures in or add them to the underlying data source. For example, to add a procedure to a physical data service:

1. Open the physical data service, and click the **Overview** tab.
2. Right-click near the top, and choose **Add Operation**.

Figure 9–32 Add Operation

Add Operation

3. Choose the **Visibility** and **Kind** of the procedure, then enter a name.
4. Click **Add** to add a parameter. Enter a **Parameter Name**, then choose a **Type**, **Kind**, and **Occurrence**. Click **OK**.
5. Select **Primary** if you want the procedure to be primary for its type.
6. Click the **Update Map** tab.
7. Right-click in the update map, then choose **Revert Customizations**.

Be sure that the procedures in the update block that maps to the underlying data source is enabled.

9.10.6.4 Change the XML Return Type

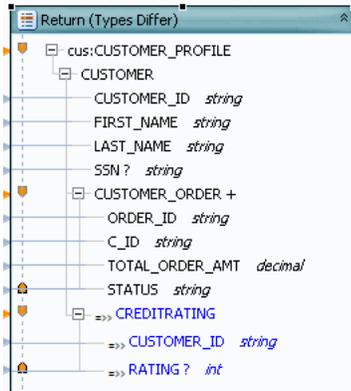
You can also change the XML schema the logical data service uses for its return type. For example, you might remove the element that attempts to update the disabled data source. You can even do this dynamically from Eclipse for WebLogic.

To change the return type from Eclipse for WebLogic:

1. Open the logical data service, and click the **Overview** tab.
2. Right-click the schema, then choose **Edit Schema**.
3. Remove the entire element, between the `<xs:element>` and `</xs:element>` tags.
4. Click the **Query Map** tab.

5. Right-click the return type, then choose **Show Type Difference**.
You should see the removed elements in blue.

Figure 9–33 Show Type Difference



Show Type Difference

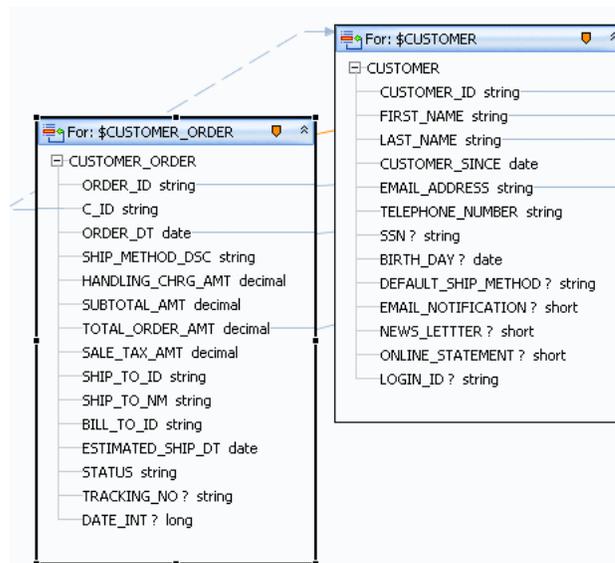
6. Right-click the removed element, and choose **Remove Element**.
7. Click the **Update Map** tab.
8. Resolve any disabled elements or procedures.

9.10.7 How To Handle Non-Unique Joins

This section shows how to enable an update map when a logical data service uses a non-unique join between relational data sources.

9.10.7.1 Understand the Join

In a logical data service, you can join tables visually in the Query Map by dragging from a key element in one data source to a corresponding key element in another data source.

Figure 9–34 Joining tables in the Query Map

Joining tables in the Query Map

You can also create a join by adding an XQuery WHERE statement in the expression editor or the Source tab:

```
where $CUSTOMER/CUSTOMER_ID eq $CREDIT_CARD/CUSTOMER_ID
```

If both tables are in the same database, the XML return type is nested, and you are joining on a unique key, Oracle Data Service Integrator creates a left outer join. You can see the SQL in the query plan for the service (click the Plan tab, then Show Query Plan):

```
SELECT ...
FROM "RTLCUSTOMER"."CUSTOMER" t1
LEFT OUTER JOIN "RTLCUSTOMER"."ADDRESS" t2
ON (t1."CUSTOMER_ID" = t2."CUSTOMER_ID")
```

If the XML return type is flat, Oracle Data Service Integrator creates an inner join, and the SQL looks like this:

```
SELECT ...
FROM "RTLAPPLOMS"."CUSTOMER_ORDER" t1
JOIN "RTLCUSTOMER"."CUSTOMER" t2
ON (t2."CUSTOMER_ID" = t1."C_ID")
```

A left outer join returns rows from the left (meaning, the first) table, even if they do not match any rows in the right (second) table.

An inner join requires that a value in the left table match a value in the right table in order for the left values to be included in the result. For example, you might match one customer to many orders, creating a joined table like this:

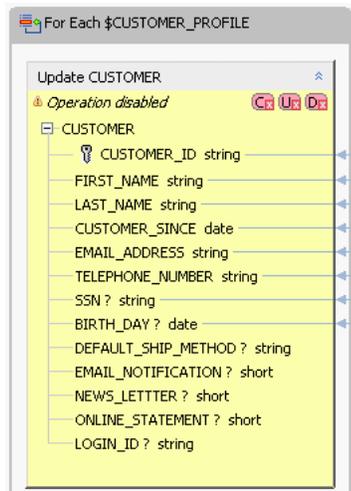
Table 9–4 Inner Join Value Requirements

CUSTOMER_ID	FIRST_NAME	LAST_NAME	EMAIL_ADDRESS	ORDER_ID	ORDER_DT	TOTAL_ORDER_AMT
CUSTOMER1	Jack	Black	jack@yahoo.com	ORDER_1_0	2001-10-01	156.39
CUSTOMER1	Jack	Black	jack@yahoo.com	ORDER_1_1	2002-02-17	596.65
CUSTOMER1	Jack	Black	jack@yahoo.com	ORDER_1_2	2002-07-07	656.65

Here, CUSTOMER_ID is a unique key and has one row in the relational source. However, in the joined table, CUSTOMER1 has three orders and three rows. If you update information for CUSTOMER1 such as FIRST_NAME in the joined table, where each customer has multiple rows, the value to use to update the underlying data source is ambiguous.

With a non-unique join, all or part of the update map is temporarily disabled and looks like this:

Figure 9–35 A Disabled Update Block



A Disabled Update Block

When you click **View Generate Log** in the update map, you see a message like this one:

The primary read function has a non-unique join involving this data source.

In your function or procedure code, in the Source tab, you might see for statements directly nested within each other, without an intervening WHERE clause:

```
for $CUSTOMER in ns1:CUSTOMER()
for $CREDIT_CARD in ns2:CREDIT_CARD()
return
```

Or, you might see XML elements directly nested within each other without intervening SQL statements:

```
<ns7:CUSTOMER_PROFILE>
  <CUSTOMER>
    ...
    {
      <CREDIT_CARD>
```

```

    ...
    </CREDIT_CARD>
  }
</CUSTOMER>
</ns7:CUSTOMER_PROFILE>

```

These are all symptoms of a non-unique join. You need to enable the update map so that you can deploy the service, test it, and make it available to client applications.

In an update map, the most common causes of a non-unique join are:

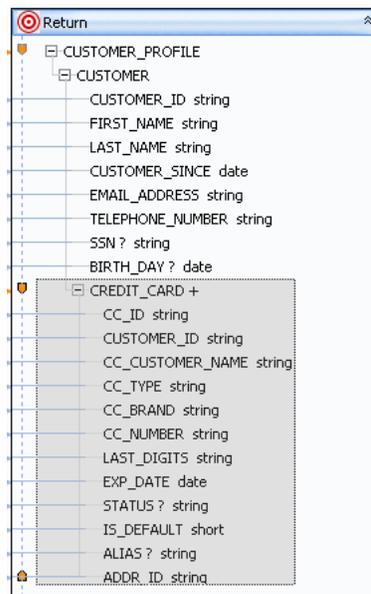
- A logical data service with a flat (non-nested) return type.
- An incorrect block scope in the query map.
- An incorrect table join, or no table join, in the query map.
- An attempt to join on a field other than a key field.

9.10.7.2 Correct the Block Scope

If your logical data service has a nested XML return type, scope the data sources to XML blocks within the return type.

1. In Query Map, click the zone icon of a data source.
2. Drag the zone icon from the data source to the nested element in the return type.
3. Mouse over the zone icon in the data source. Verify that only the nested element is highlighted in the return type.

Figure 9–36 *Checking the Scope in the Return Type*



Checking the Scope in the Return Type

9.10.7.3 Correct the Table Join

You might also get a non-unique join if the data sources are not joined correctly. You can join the tables either visually in the Query Map or by entering a WHERE clause in

the expression editor or the Source tab. Be sure to join tables on a key element, marked like this:

```
ORDER_ID string-
```

To join tables visually:

1. Click the Query Map tab.
2. Drag from a key element in one data source to the same key element in another data source (for example, \$CUSTOMER/CUSTOMER_ID to \$ADDRESS/CUSTOMER_ID).
3. Click the Source tab and expand the read function to check the location of the WHERE clause. For example, if your XML return type is nested, the XQuery code should also be nested:

```
for $CUSTOMER in ns1:CUSTOMER()
return
  ...
  for $CREDIT_CARD in ns2:CREDIT_CARD()
  where $CUSTOMER/CUSTOMER_ID eq $CREDIT_CARD/CUSTOMER_ID
  return
  ...
```

To use the expression editor:

1. Click the **Query Map** tab.
2. Click the **For** block of the data source you are joining to.
3. In the expression editor, click **Add Where Clause**.
4. After the Where keyword, add the elements to be joined (for example, \$CUSTOMER/CUSTOMER_ID eq \$CREDIT_CARD/CUSTOMER_ID).
5. Click **Save**.
6. Check the WHERE clause in the **Source** tab, as described above.

Remember that Oracle Data Service Integrator creates a left outer join if both tables are in the same database and the XML return type is nested. If the XML return type is flat, Oracle Data Service Integrator creates an inner join.

9.10.7.4 Enable Update Blocks and Procedures

If your service has a return type with a flat structure, you may get a non-unique join, even if the join is correct in the Query Map and the Source tab.

If this happens, or if all or part of the update map is disabled for any reason, you can enable an update block or the Create-Update-Delete procedures within the block.

To enable a disabled (yellow) update block:

1. Right-click in the block, and choose **Enable**.
The update block should now have a white (enabled) background. The Create, Update, or Delete procedure icons might still appear red or yellow, if they are disabled. However, you should be able to test the primary read function.
2. Click the **Test** tab.
3. At **Select Operation**, choose the primary read function, and click **Run**.

To enable an update map procedure:

1. If an element is marked with a Warning icon indicating that a mapping is required, select it.

2. In the expression editor, give the element a value with the correct data type.
3. Continue for all disabled elements.
4. In the Test tab, test an update procedure to ensure that the value overrides you have entered do what you want.

9.10.7.5 Test a Non-Unique Join

Let's go back to the sample joined table data (which we can see in the Test tab, by choosing the primary read function and clicking Run):

Table 9–5 Testing a Non-Unique Join

CUSTOMER_ID	FIRST_NAME	LAST_NAME	EMAIL_ADDRESS	ORDER_ID	ORDER_DT	TOTAL_ORDER_AMT
CUSTOMER1	Jack	Black	jack@yahoo.com	ORDER_1_0	2001-10-01	156.39
CUSTOMER1	Jack	Black	jack@yahoo.com	ORDER_1_1	2001-02-17	596.65
CUSTOMER1	Jack	Black	jack@yahoo.com	ORDER_1_2	2001-07-07	656.65

In this case, the XML return type is flat, and Oracle Data Service Integrator has created an inner join between the CUSTOMER and CUSTOMER_ORDER tables in underlying relational data sources. In the joined table view, one customer has many orders. The CUSTOMER_ID can appear multiple times, but the ORDER_ID is unique.

Once the update map is enabled, you can update data in either the CUSTOMER or CUSTOMER_ORDER table in the data sources:

1. Click a row in the joined table data, then click Edit.
2. Locate the correct node in the XML tree data, and expand it.
3. Click the value you want to change, then edit it.
4. Click Submit.

If you update TOTAL_ORDER_AMT, from the CUSTOMER_ORDER table, the amount changes in one row of the joined table view.

However, if you update EMAIL_ADDRESS, the email address changes in one row of the data source table and in all rows for that customer in the joined table view.

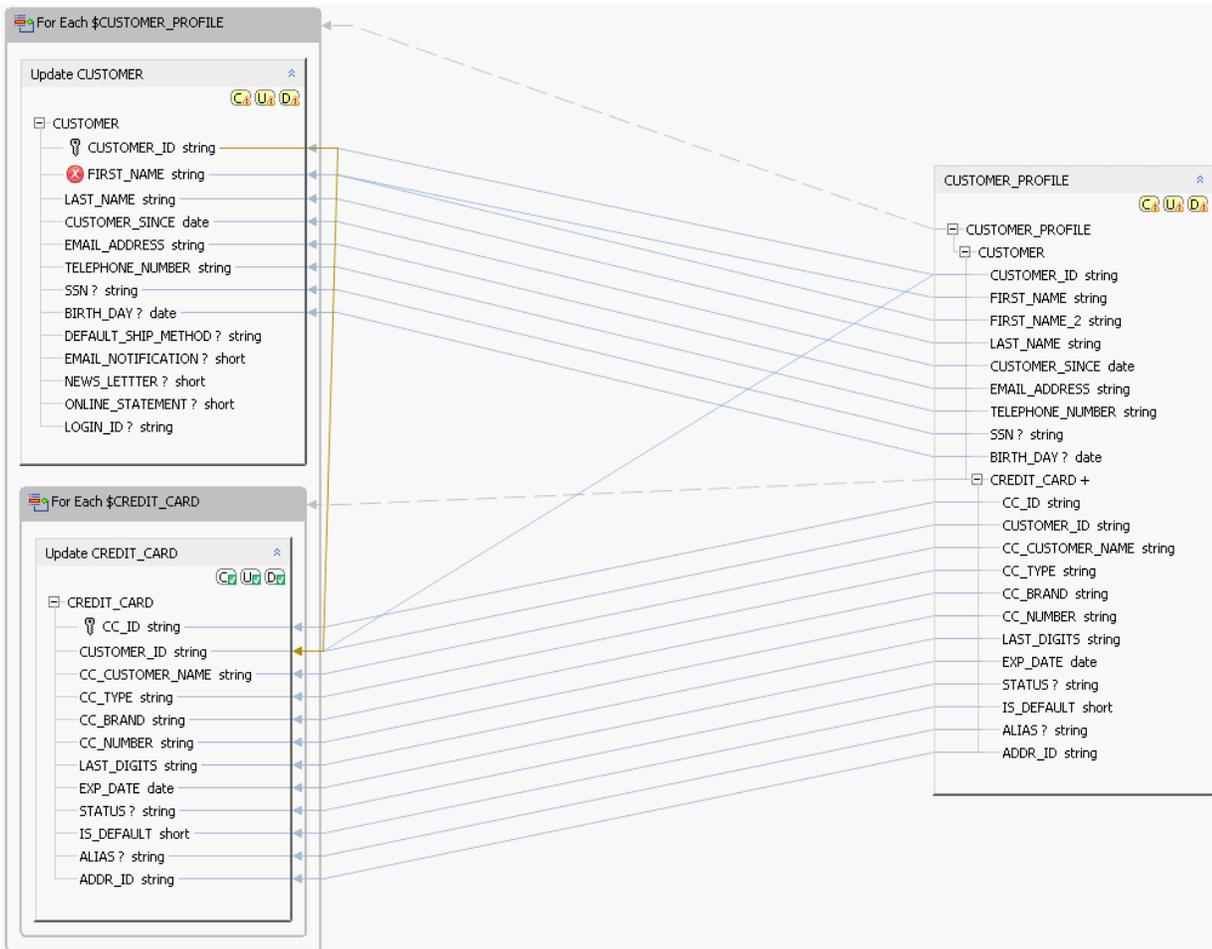
9.10.8 How To Handle Non-Unique Values

This section describes how to handle an update map that is disabled because two values in a return type map to one value in a data source.

9.10.8.1 Example

In a query map, you might attempt to map one value in a data source to two values in an XML return type. When the update map is generated and the flow is reversed, two values map from the return type to one in the data source, which creates an update error.

Figure 9–37 An Error from a Non-Unique Value



An Error from a Non-Unique Value

The cause of the error is that two values are attempting to update one in the data source. This creates a build error in the logical data service, and you cannot deploy or test it. You cannot right-click and enable the update block either. The update doesn't work unless you write a custom update function in XQSE.

The best solution is to disable the multiple mapping in the Query Map tab:

1. Click Query Map.
2. Delete the mapping line from the data source to the second, duplicate element in the return type. This should reverse the error.
3. Save the data service and click Update Map to check the change.
4. If the error still exists, right-click and choose Revert Customizations.

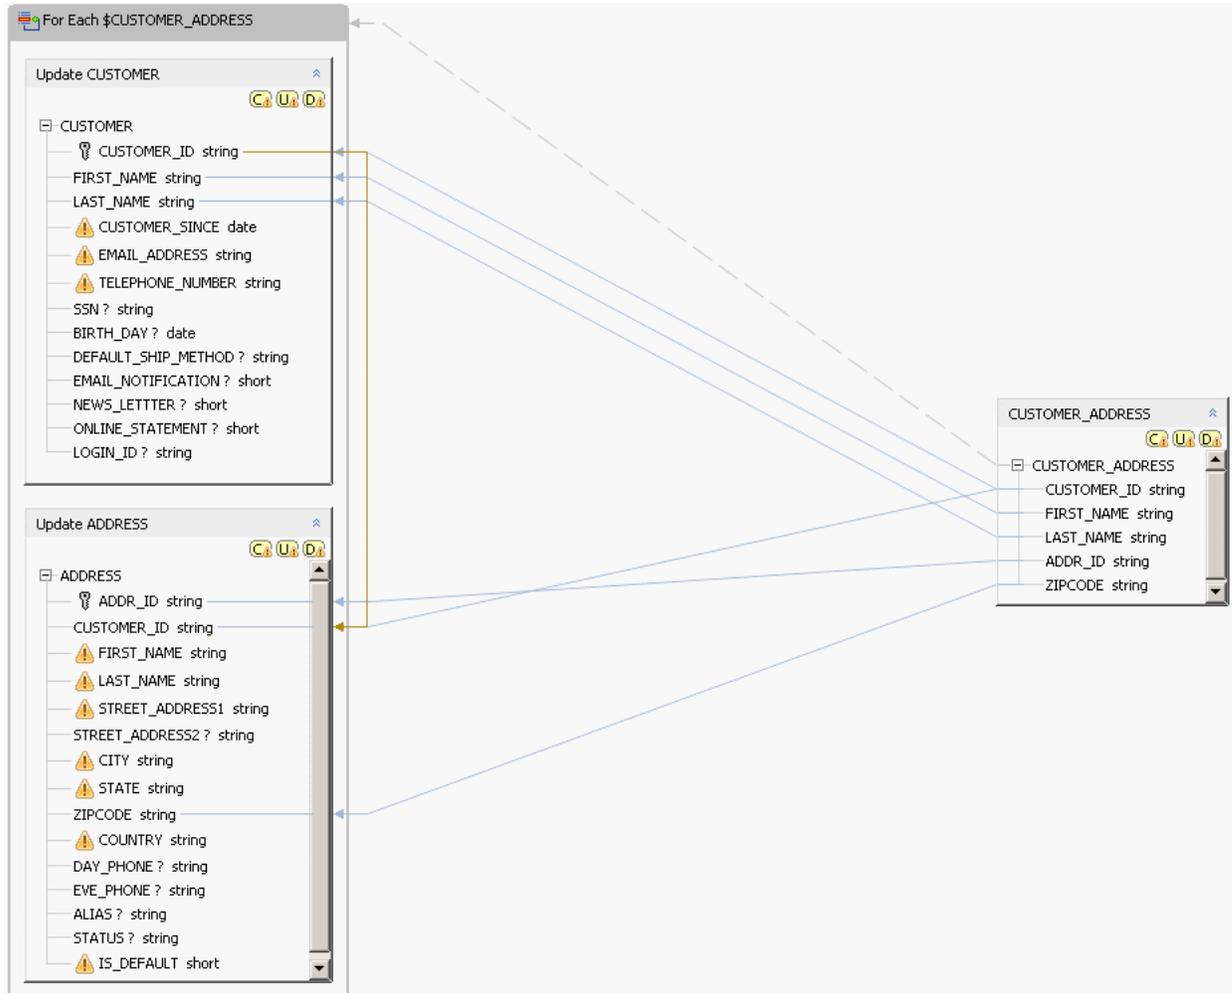
9.10.9 How To Handle Unmapped Required Values

This section describes how to enable an update map when the data sources on the left have required elements that are not mapped from the return type on the right.

9.10.9.1 Overview

When required mappings are missing, the Create-Update-Delete procedures for the update block are disabled. That means you cannot create, update, or delete the underlying data sources. In Eclipse for WebLogic, the update map looks like this.

Figure 9–38 Required Mappings Are Missing



Required Mappings Are Missing

- A mapping that was deleted from or did not exist in the query map.
- An XML return type that does not contain all required elements. This can be valid, especially if you do not want to expose all elements in your data sources to a client application.

If an element is required but does not have a value, it is marked with a Warning icon.

In either case, the Create, Update, or Delete procedures do not work, so you need to resolve the error. You can do either of these:

- Draw the mapping in Query Map view.
- Enter an override value (either an expression or a constant) in the expression editor.

9.10.9.2 Draw the Mapping

To draw the mapping in the Query Map tab:

1. Click **Query Map**.
2. Drag from an element in a data source on the left to the matching element in a return type on the right.

Make sure the elements have the same data types or similar data types that are cast implicitly.

9.10.9.3 Cast a Constant

If you enter a constant to override the missing mapping, it is only used with Create procedures, to insert data into the data source. Update procedures ignore the override values you enter and leave the data source unchanged. (Of course, Delete procedures delete a record from the data source, so override values are not relevant to them.)

When you enter an override value, make sure the value you enter has the data type the element in the physical data source requires. You can enter a constant like "44" or "2007-01-01" and cast it to an XML Schema data type such as xs:integer or xs:date, using either of these:

- A built-in XQuery cast function
- The parentheses cast operator, as in xs:date("string"), to invoke an XML Schema type constructor function

The parentheses cast operator uses any XML Schema data type outside the parentheses and a string that is appropriate for the data type you are casting to within the parentheses. For example, you can perform these casts:

```
xs:date("2007-01-01")
xs:dateTime("2007-01-01T16:44:44")
xs:integer("44")
```

But you cannot perform these:

```
xs:date("2007-01-01T16:44:44")
xs:dateTime("date")
xs:integer("text")
```

To cast a constant in the expression editor:

1. Click the Update Map tab.
2. Click an unmapped element in a data source on the left.
3. In the expression editor, enter a constant that has the data type the element requires. For example, for an element of type xs:string, you might enter:

```
"Bob"
```

If the element has another data type, enter a string within a cast expression, for example:

```
xs:integer("44")
xs:dateTime("2007-07-17T09:00:00")
```

4. Continue for all disabled elements.
5. In the Test tab, test an update using Run - Edit - Submit to make sure the value overrides work as you expect.

9.11 Testing Update Maps

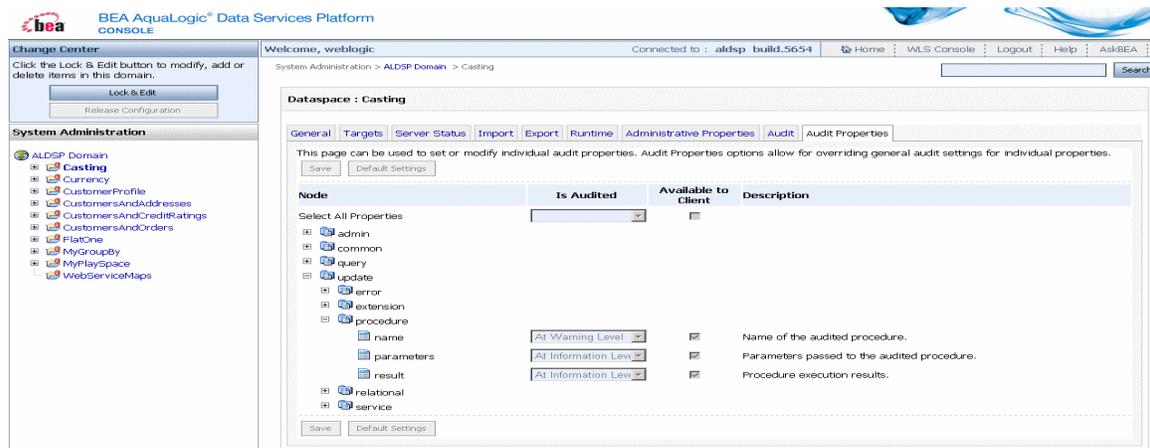
This section describes how to test an Update procedure in Test view in Eclipse for WebLogic.

9.11.1 Configure Audit Properties

To test an Update procedure in Eclipse for WebLogic, you must submit a data graph in the Parameters box in Test view. A data graph is an XML structure with a root element of <sdo:datagraph> and a <changesummary> element. The easiest way to submit a data graph is to capture one from an audit.

First, configure audit properties in the Oracle Data Service Integrator Console.

Figure 9–39 Configuring Audit Properties in the Oracle Data Service Integrator Console



Configuring Audit Properties in the Oracle Data Service Integrator Console

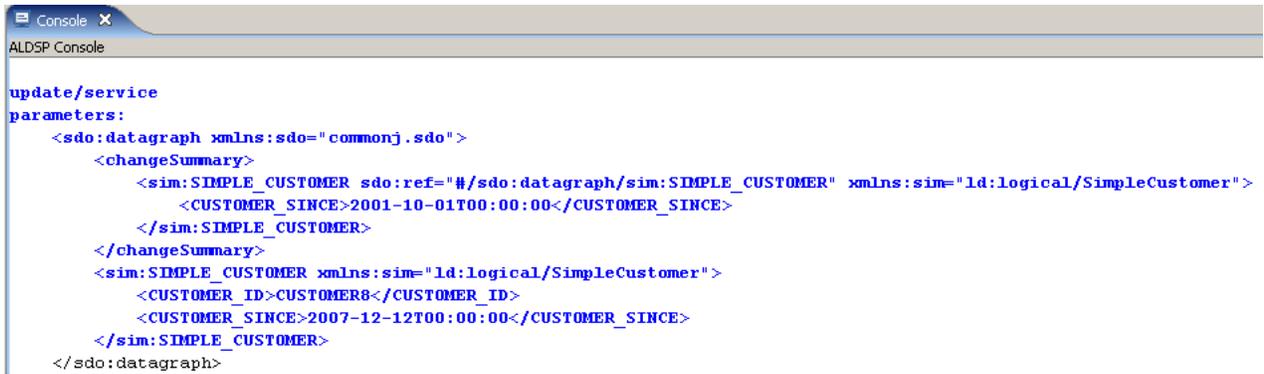
To configure audit properties so that Oracle Data Service Integrator generates data graphs:

1. Open the Oracle Data Service Integrator Console and log in.
2. Click the name of a data space project.
3. Click the Audit Properties tab.
4. Click **Lock & Edit** in the upper left pane.
5. Navigate to the Update > Service node (be careful not to move to Update > Error > Procedure).
6. For Name, Parameters, and Result, choose Always from the Is Audited menu.
7. Click **Save**.
8. Click **Activate Changes** in the upper left pane.

9.11.2 Capture the Data Graph

You can then capture a data graph from the audit messages displayed in the Eclipse for WebLogic Console tab, and edit the data graph to submit to the Update procedure in Test view.

Figure 9–40 Viewing a Data Graph in the Eclipse for WebLogic Console Tab



Viewing a Data Graph in the Eclipse for WebLogic Console Tab

To capture the data graph:

1. Open a logical data service in Eclipse for WebLogic.
2. Click the Test tab.
3. Choose the service's primary Read function, then click **Run**.
4. Click **Edit**, edit a value, then click **Submit**.
5. (Optional) Check the Eclipse for WebLogic Console tab.

If you see the Oracle WebLogic Server console data, not the Oracle Data Service Integrator console data, click the drop-down arrow next to the console icon, and choose Oracle Data Service Integrator Console.

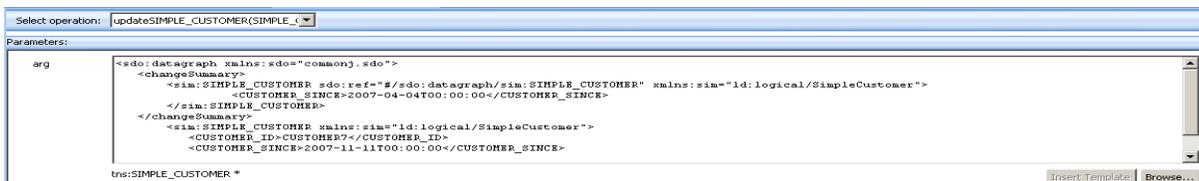
6. Scroll up in the Eclipse for WebLogic Console tab until you locate the data graph, right-click, and copy it.

9.11.3 Submit the Update

When you update relational sources, the SDO update mechanism uses optimistic locking to avoid change conflicts. With optimistic locking, the data source is not locked when the SDO client acquires the data. Later, when the client wants to update, the data in the source is compared to a copy of the data at a time when it was acquired. If there are discrepancies, the update is not committed. Before you submit the data graph to the Update procedure, be sure that optimistic locking is enabled in the underlying data source you are updating.

You can then submit the data graph to the Update procedure. However, you may need to edit it, as the data graph you captured from the Eclipse for WebLogic Console tab reflected the last change you made, not the change you are presently submitting to the Update procedure.

Figure 9–41 Submitting the Data Graph to the Update Procedure



Submitting the Data Graph to the Update Procedure

The data graph you submit to the Update procedure takes the place of the return type as an argument, even if you are updating only some of the elements in the return type.

To submit the data graph to an Update procedure:

1. Enable optimistic locking on any physical relational data sources the data graph is updating.
2. Open a data service in Eclipse for WebLogic, and click the Test tab.
3. At Select Operation, choose an Update procedure.
4. Copy a data graph you have captured from the Eclipse for WebLogic Console tab to the Parameters box.
5. Edit the data graph for the change you want to make.

The data graph you captured applies to a change made in the visual interface. Update the change summary to the values the object presently has, and the remaining elements to the new values you want to set. For example, this is a change summary captured from the Eclipse for WebLogic Console tab:

```
<sdo:datagraph xmlns:sdo="commonj.sdo">
  <changeSummary>
    <sim:SIMPLE_CUSTOMER sdo:ref="#/sdo:datagraph/sim:SIMPLE_CUSTOMER"
xmlns:sim="ld:logical/SimpleCustomer">
      <CUSTOMER_SINCE>1999-01-01T00:00:00</CUSTOMER_SINCE>
    </sim:SIMPLE_CUSTOMER>
  </changeSummary>
  <sim:SIMPLE_CUSTOMER xmlns:sim="ld:logical/SimpleCustomer">
    <CUSTOMER_ID>CUSTOMER7</CUSTOMER_ID>
    <CUSTOMER_SINCE>2007-11-11T00:00:00</CUSTOMER_SINCE>
  </sim:SIMPLE_CUSTOMER>
</sdo:datagraph>
```

This version has been updated in the Parameters box (note the difference in the CUSTOMER_SINCE dates):

```
<sdo:datagraph xmlns:sdo="commonj.sdo">
  <changeSummary>
    <sim:SIMPLE_CUSTOMER sdo:ref="#/sdo:datagraph/sim:SIMPLE_CUSTOMER"
xmlns:sim="ld:logical/SimpleCustomer">
      <CUSTOMER_SINCE>2007-11-11T00:00:00</CUSTOMER_SINCE>
    </sim:SIMPLE_CUSTOMER>
  </changeSummary>
  <sim:SIMPLE_CUSTOMER xmlns:sim="ld:logical/SimpleCustomer">
    <CUSTOMER_ID>CUSTOMER7</CUSTOMER_ID>
    <CUSTOMER_SINCE>2008-04-04T00:00:00</CUSTOMER_SINCE>
  </sim:SIMPLE_CUSTOMER>
</sdo:datagraph>
```

6. Click **Run**. You should see this message in Test view:

```
Operation was successful.
```

9.12 How To Test an Update Procedure

This section describes how to test an Update procedure in Test view in Eclipse for WebLogic.

9.12.1 Configure Audit Properties

To test an Update procedure in Eclipse for WebLogic, you must submit a data graph in the Parameters box in Test view. A data graph is an XML structure with a root element of `<sdo:datagraph>` and a `<changesummary>` element. The easiest way to submit a data graph is to capture one from an audit.

First, configure audit properties in the Oracle Data Service Integrator Console.

To configure audit properties so that Oracle Data Service Integrator generates data graphs:

1. Open the Oracle Data Service Integrator Console and log in.
2. Click the name of a data space project.
3. Click the Audit Properties tab.
4. Click **Lock & Edit** in the upper left pane.
5. Navigate to the Update > Service node (be careful not to move to Update > Error > Procedure).
6. For Name, Parameters, and Result, choose Always from the Is Audited menu.
7. Click **Save**.
8. Click **Activate Changes** in the upper left pane.

9.12.2 Capture the Data Graph

You can then capture a data graph from the audit messages displayed in the Eclipse for WebLogic Console tab, and edit the data graph to submit to the Update procedure in Test view.

To capture a data graph:

1. Open a logical data service in Eclipse for WebLogic.
2. Click the Test tab.
3. Choose the service's primary Read function, then click **Run**.
4. Click **Edit**, edit a value, then click **Submit**.
5. (Optional) Check the Eclipse for WebLogic Console tab.

If you see the Oracle WebLogic Server console data, not the Oracle Data Service Integrator console data, click the drop-down arrow next to the console icon, and choose Oracle Data Service Integrator Console.

6. Scroll up in the Eclipse for WebLogic Console tab until you locate the data graph, right-click, and copy it.

9.12.3 Submit the Update

When you update relational sources, the SDO update mechanism uses optimistic locking to avoid change conflicts. With optimistic locking, the data source is not locked when the SDO client acquires the data. Later, when the client wants to update, the data in the source is compared to a copy of the data at a time when it was acquired. If there are discrepancies, the update is not committed. Before you submit the data graph to the Update procedure, be sure that optimistic locking is enabled in the underlying data source you are updating.

You can then submit the data graph to the Update procedure. However, you may need to edit it, as the data graph you captured from the Eclipse for WebLogic Console tab reflected the last change you made, not the change you are presently submitting to the Update procedure.

The data graph you submit to the Update procedure takes the place of the return type as an argument, even if you are updating only some of the elements in the return type.

To submit the data graph to an Update procedure:

1. Enable optimistic locking on any physical relational data sources the data graph is updating.
2. Open a data service in Eclipse for WebLogic, and click the **Test** tab.
3. At **Select Operation**, choose an **Update** procedure.
4. Copy a data graph you have captured from the Eclipse for WebLogic Console tab to the **Parameters** box.
5. Edit the data graph for the change you want to make.

The data graph you captured applies to a change made in the visual interface. Update the change summary to the values the object presently has, and the remaining elements to the new values you want to set. For example, this is a change summary captured from the Eclipse for WebLogic Console tab:

```
<sdo:datagraph xmlns:sdo="commonj.sdo">
  <changeSummary>
    <sim:SIMPLE_CUSTOMER sdo:ref="#/sdo:datagraph/sim:SIMPLE_CUSTOMER"
xmlns:sim="ld:logical/SimpleCustomer">
      <CUSTOMER_SINCE>1999-01-01T00:00:00</CUSTOMER_SINCE>
    </sim:SIMPLE_CUSTOMER>
  </changeSummary>
  <sim:SIMPLE_CUSTOMER xmlns:sim="ld:logical/SimpleCustomer">
    <CUSTOMER_ID>CUSTOMER7</CUSTOMER_ID>
    <CUSTOMER_SINCE>2007-11-11T00:00:00</CUSTOMER_SINCE>
  </sim:SIMPLE_CUSTOMER>
</sdo:datagraph>
```

This version has been updated in the Parameters box (note the difference in the **CUSTOMER_SINCE** dates):

```
<sdo:datagraph xmlns:sdo="commonj.sdo">
  <changeSummary>
    <sim:SIMPLE_CUSTOMER sdo:ref="#/sdo:datagraph/sim:SIMPLE_CUSTOMER"
xmlns:sim="ld:logical/SimpleCustomer">
      <CUSTOMER_SINCE>2007-11-11T00:00:00</CUSTOMER_SINCE>
    </sim:SIMPLE_CUSTOMER>
  </changeSummary>
  <sim:SIMPLE_CUSTOMER xmlns:sim="ld:logical/SimpleCustomer">
    <CUSTOMER_ID>CUSTOMER7</CUSTOMER_ID>
    <CUSTOMER_SINCE>2008-04-04T00:00:00</CUSTOMER_SINCE>
  </sim:SIMPLE_CUSTOMER>
</sdo:datagraph>
```

6. Click **Run**. You should see this message in Test view:

Operation was successful.

Preparing Services for Clients

This chapter describes how to prepare client JAR files and web service maps for client applications. This chapter includes the following sections:

- [Section 10.1, "Generating a Mediator Client JAR File"](#)
- [Section 10.2, "Generating a Web Services Mediator Client JAR File"](#)
- [Section 10.3, "Generating a Web Service Map and WSDL from a Data Service"](#)
- [Section 10.4, "Configuring Security for Web Services Applications"](#)
- [Section 10.5, "Web Services Map File Reference"](#)
- [Section 10.6, "Understanding SQL Maps"](#)
- [Section 10.7, "Map Functions and Procedures to SQL Objects"](#)
- [Section 10.8, "SQL Object Mapping Rules"](#)
- [Section 10.9, "Constraints on Publishing Data Service Objects to SQL"](#)

10.1 Generating a Mediator Client JAR File

To use the Static Mediator API in a client application, you must generate a Mediator Client JAR file. This JAR file contains the Static Mediator API interfaces, plus all the necessary SDO-compiled schemas for a dataspace.

One Java method is generated for each mapped data service operation. Method names match the mapped data service operation names. Client developers access data service operations by calling these methods.

This section explains how to generate a Mediator Client JAR by these two methods:

- [Section 10.1.1, "Using the IDE"](#)
- [Section 10.1.2, "Using the Command-Line Tool"](#)

Note: You can also generate a Mediator Client JAR using the Administration Console. See the *Oracle Data Service Integrator Administration Guide* for details.

10.1.1 Using the IDE

To generate a mediator client JAR file using the IDE:

1. Select **File > Export**.

2. In the Export dialog, select **Oracle Data Service Integrator > Mediator Client JAR File**.
3. Click **Next**.
4. Complete the Mediator Client JAR File dialog as follows:
 1. Select a Dataspace project to export. You can only select one Dataspace project at a time.
 2. Specify a directory in which to place the exported JAR file. You can use the drop down list to select a recently specified directory or use the **Browse** button to locate one.
 3. Unselect the **Use default name** checkbox if you want to enter a name for the JAR file.
 4. Click **Finish** to create the JAR file.

The Oracle Data Service Integrator Console view displays the export task status and any errors that may have occurred. You can click the **Cancel** button to cancel the export task before it has completed.

10.1.2 Using the Command-Line Tool

This section explains how to generate a Mediator Client JAR file using the command-line tool. Before using the command-line tool, be sure you have the following:

- Oracle WebLogic Server 10.3 installed with Oracle Data Service Integrator installed in the default location `BEA_HOME/odsi_10.3`.
- A Dataspace project on your local file system that contains data service (`.ds`) and schema (`.xsd`) files. Miscellaneous IDE files within the project folder are allowed and will not affect the export.
- Ant installed and in your path.

To generate the client JAR, use this Ant command:

```
ant -Dapproot=PROJECT_HOME -f BEA_HOME/odsi/bin/sdo_dspclientgen.xml
```

where `PROJECT_HOME` is the full path to the Data Space project's root folder, and `BEA_HOME` is the root path for your Oracle WebLogic installation.

For example (all on one line):

```
ant -Dapproot=/home/myprojects/myapp -f /home/bea/odsi/bin/sdo_dspclientgen.xml
```

This Ant script produces a file named `PROJECTNAME-dsp-client.jar` in `PROJECT_HOME`, where `PROJECTNAME` is the name of the directory `PROJECT_HOME` (as opposed to the full path to that directory). For example, the above script produces the Mediator Client JAR file:

```
/home/myprojects/myapp/myapp-dsp-client.jar.
```

Optional command-line features include:

- Your environment must contain a `WL_HOME` environment variable, pointing to the WLS 9.2 installation. If it does not, you can provide an alternate by adding `-Dwl.home=/path` to specify the WLS root directory.

- Your Oracle Data Service Integrator installation must be in the default directory `BEA_HOME/odsi_10.3`. If it is not, you can provide an alternate by adding `-Ddsp.home=/path` to specify the directory.
- To specify a full directory path for the output, add `-Doutdir=/dirpath` to the Ant command. You must provide an absolute path; a relative path, including ".", will not work, as it is assumed to be relative to `PROJECT_HOME`.
- To specify a different name for the JAR file, add `-Dsdojarname=name.jar`.

10.2 Generating a Web Services Mediator Client JAR File

This section explains how to generate a Web Services Mediator Client JAR file. This JAR is required by developers writing Java clients that access data services through web services using the Static Mediator API.

This section includes these sections:

- [Section 10.2.1, "Overview"](#)
- [Section 10.2.2, "Using Eclipse for WebLogic"](#)
- [Section 10.2.3, "Using the Command-Line Tool"](#)

10.2.1 Overview

To use the Static Mediator API in a web services-enabled client application, you must generate a Web Services Mediator Client JAR file. This JAR file contains the Static Mediator API interfaces, plus all the necessary SDO-compiled schemas for a dataspace.

One Java method is generated for each data service function that is mapped to a WSDL operation. Method names match the mapped WSDL operation name. Client developers access data service functions through the web service by calling these methods. If the web service requires message-level security, you can add a credential provider and trust manager through initial context properties. For more information on security, see *Configure Security for Web Services Applications*.

Note: You can also generate a Mediator Client JAR using the Administration Console. See the *Oracle Data Service Integrator Administration Guide* for details.

10.2.2 Using Eclipse for WebLogic

To generate a Web Services Mediator Client JAR file using Eclipse for WebLogic:

1. Select **File > Export**.
2. In the Export dialog, select **Oracle Data Service Integrator > Web Services Mediator Client JAR File**.
3. Click **Next**.
4. Complete the Web Services Mediator Client JAR File dialog as follows:
 1. In the left panel, select the Dataspace project that contains the .ws file(s) to export. You can only export .ws files in one Dataspace project at a time. Checking or unchecking the checkbox next to a project or a folder automatically checks or unchecks all the sub-folders and .ws files under that project/folder.

2. In the right panel, select the Web Service Map file to export. You can select one or more `.ws` files. To see and selectively check the `.ws` files in a sub-folder, expand and click on the folder on the left panel. The message under the right panel shows the total number of `.ws` files currently checked for export.
3. Specify a directory in which to place the exported JAR file. You can select any location on your system. You can use the dropdown list to select a recently specified directory or use the Browse button to locate one. By default, the exported JAR will be named: `<data_space_name>-ws-client.jar`.
4. Unselect the **Use default name** checkbox if you want to enter a name for the JAR file.
5. Click **Finish**.

The Oracle Data Service Integrator Console view displays the export task status and any errors that may have occurred. You can click the **Cancel** button to cancel the export task before it has completed.

10.2.3 Using the Command-Line Tool

This section explains how to generate the Web Services Mediator Client JAR file using Ant and presents example Ant commands. Before using the command-line tool, be sure you have the following:

- Oracle WebLogic Server 10.3 installed with Oracle Data Service Integrator installed in the default location `BEA_HOME/odsi_10.3`.
- A Dataspace project on your local filesystem that contains data service (`.ds`) and schema (`.xsd`) files. Miscellaneous IDE files within the project folder are allowed and will not affect the export.
- Ant installed and in your path.

To generate the client JAR, use this Ant command:

```
ant -Dapproot=PROJECT_HOME -Dwslocator=locator -f BEA_HOME/odsi_10.3/bin/sdo_
dspclientgen.xml
```

Where:

- `PROJECT_HOME` is the path to the Dataspace project. You must specify a full path for the values of `BEA_HOME` and `PROJECT_HOME`.
- The `locator` option takes one of these values:
 - `d:URI` - Specifies a URI (or a semicolon-separated or space-separated list of URIs) to a `.ws` file in the Dataspace project from which to generate the JAR file. For example:


```
ld:MediatorTestDataServices/CustomerWeb.ws
```
 - `ALL` - Generates the JAR for all `.ws` files in the dataspace.

The result of executing this Ant script is a file named `PROJECT-ld-client.jar` in `PROJECT_HOME`, where `PROJECT` is the name of the directory `PROJECT_HOME` (as opposed to the full path to that directory).

Optional command-line features include:

- Your environment must contain a `WL_HOME` environment variable, pointing to the WLS 9.2 installation. If it does not, you can provide an alternate by adding `-Dwl.home=/path` to specify the WLS root directory.

- Your Oracle Data Service Integrator installation must be in the default directory `BEA_HOME/odsi_10.3`. If it is not, you can provide an alternate by adding `-Ddsp.home=/path` to specify the directory.
- To specify a full directory path for the output, add `-Doutdir=/dirpath` to the Ant command. You must provide an absolute path; a relative path, including ".", will not work, as it is assumed to be relative to `PROJECT_HOME`.
- To specify a different name for the JAR file, add `-Dsdojarname=name.jar`.

This example specifies multiple `.ws` files. The command must be entered on one line.

```
ant -Daproot=/home/myprojects/myapp
-Dwslocator='ld:MediatorTestDataServices/CustomWeb.ws;
ld:MediatorTestDataServices/OtherCustomWeb.ws' -f /home/bea/odsi_10.3/bin/
sdo_dspclientgen.xml
```

This example generates a JAR that includes all of the `.ws` files in the dataspace. The command must be entered on one line.

```
ant -Daproot=/home/myprojects/myapp -Dwslocator=ALL -f /home/bea/
odsi_10.3/bin/sdo_dspclientgen.xml
```

10.3 Generating a Web Service Map and WSDL from a Data Service

If you intend to access a data service through web services using the Data Services Mediator API, you must generate a web service map file first. A web service map file maps data service functions to web service operations. The map file is also used for setting and configuring security policies for web services applications.

This section describes the following topics:

- [Section 10.3.1, "Creating a Map File"](#)
- [Section 10.3.2, "Generating a WSDL File from a Map File"](#)
- [Section 10.3.3, "Examining the Generated WSDL"](#)
- [Section 10.3.4, "Testing the Generated WSDL"](#)
- [Section 10.3.5, "Modifying the Map File"](#)

10.3.1 Creating a Map File

This section describes the basic steps that are required to create a map file. You can accomplish all of these tasks using the Oracle Data Service Integrator Eclipse IDE. The procedure assumes that you have created or have access to the data service (`.ds`) file from which you want to create a data service.

There are two ways to create a web service map file (`.ws` file):

Method 1

1. Obtain access to the data space project containing the data service you wish to make accessible from a web service.
2. Right-click on the data service name in Project Explorer and select **Create Web Service Map**. The map file (`.ws` file) is created with the same name as the data service file, and the map file is opened in the editor.

Method 2

1. Obtain access to the data space project containing the data service you wish to make accessible from a web service.
2. Right-click on the data service name in Project Explorer and select **New > Web Service Map**.
3. Use the dialog to create an empty web service map file with a name of your choosing (example: OrderService).
4. Click **Finish**. The empty map file opens in the editor.
5. Drag either an entire data service file onto the map file or drag individual data service operations.

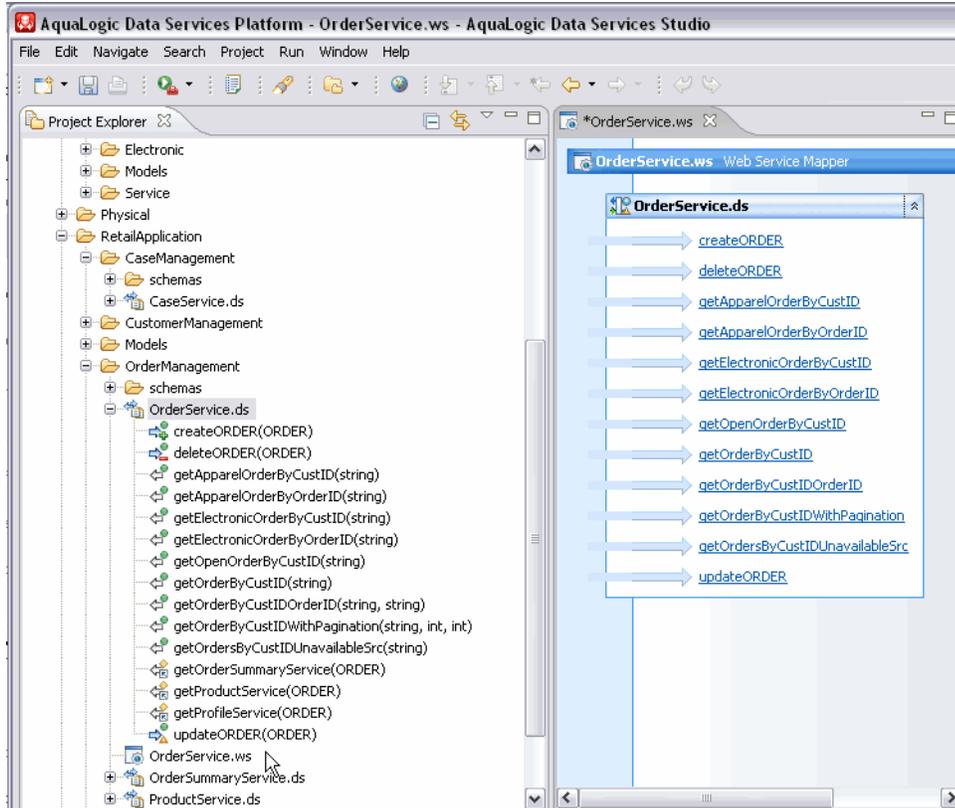
Example: RetailApplication > OrderManagement > OrderService.ds

6. Click **OK**. A file named OrderService.ws is created in the same folder as the source data service.

Note: Only the data service functions that are mapped in the map file are available to clients. Only public data service operations can be mapped.

Figure 10–1 shows a data service file called OrderService.ds as the source for a map file called OrderService.ws (created using Method 2).

Figure 10–1 Adding a Data Service to a Web Services Map File



This graphic displays a data service file as the source for a map file.

Note: You can add additional public operations from other data services to the same web service map file.

10.3.2 Generating a WSDL File from a Map File

To generate a WSDL file from a .WS file:

1. Right-click on the .WS file.
2. Choose **Save WSDL As**.
3. Specify a name for the WSDL file.
4. Click **Save**. Your new WSDL file should appear in the directory identified in the Save as dialog.
5. Double-click on the WSDL file to verify the SOAP service for the WSDL.

Note: A WSDL that has more than one schema section pointing to the same target namespace results in validation errors with the Eclipse WTP default WSDL validator. The WSDL generated in the above example is valid; however the project indicates a validation error condition.

The error condition does not interfere with your ability to build and deploy the project. Also, you can use the following Eclipse option settings to prevent the validation error report from displaying:

1. Select **Project > Properties > Validation**.
2. Select **Override** validation preferences.
3. Uncheck the **Build** option associated with the WSDL Validator.
4. Click **Apply**, then **OK**.
5. Select **Project > Clean**.
6. Click **OK**.

The validation error warnings should disappear.

10.3.3 Examining the Generated WSDL

You can examine the generated WSDL file. See [Section 10.5, "Web Services Map File Reference"](#) for details.

10.3.4 Testing the Generated WSDL

You can test the generated WSDL file. See [Section 10.5, "Web Services Map File Reference"](#) for details.

10.3.5 Modifying the Map File

This section describes additional ways to add data services and operation to a map, and how to delete operations from an existing map.

10.3.5.1 Adding Data Services and Operations

You can drag and drop either an entire data service or individual data service operations from the Project Explorer onto an existing map file in the map file editor.

You can right-click in the map editor and select **Add Data Services/Operations to Map**. Use the Select Resources to Add to Map dialog to add data service resources to the map.

10.3.5.2 Deleting Data Services and Operations from a Map File

To delete one or more operations, select the operations and right-click on the selected operations, and then select **Delete**.

To delete all operations that are related to a data service, right-click on the `.ds` dataservice box and select **Delete**.

10.3.5.3 Renaming Mapped Operations

To rename a mapped operation, select the operation, right-click and select **Rename Operation**. Then enter a new name for the mapped operation.

10.4 Configuring Security for Web Services Applications

This section describes the following topics:

- [Section 10.4.1, "Configuring Basic Authentication"](#)
- [Section 10.4.2, "Configuring Transport Level Security \(HTTPS\)"](#)
- [Section 10.4.3, "Configuring Web Services Security \(WSS\)"](#)
- [Section 10.4.4, "Specifying Policies"](#)[Section 10.4.3, "Configuring Web Services Security \(WSS\)"](#)

Oracle Data Service Integrator Native Web Services supports the following security features:

- Basic authentication (Web Application Security)
- Transport level security (HTTPS)
- Message level security (Web Services Security)

Note: When configuring a BPEL process with `SOAPReference`, you must set `oracle.soa.ws.outbound.omitWSA` to true by updating the `composite.xml` file or using the **Properties > Composite** option. Then save the process and re-run the project.

10.4.1 Configuring Basic Authentication

To use basic authentication, set the **Basic Auth Required** property of the web services map file to **true**. For more information, see [Section 10.5, "Web Services Map File Reference"](#).

10.4.2 Configuring Transport Level Security (HTTPS)

Use the web service map file property editor to change the **Transport Type** to **HTTPS**. HTTP is the default. For more information, see [Section 10.5, "Web Services Map File Reference"](#).

For HTTPS, you can configure either 1-way or 2-way SSL. For detailed information on transport level security, see the Oracle WebLogic Server document *Configuring Security: Configuring Transport-Level Security*.

10.4.3 Configuring Web Services Security (WSS)

WSS provides message level security. For WSS, Oracle Data Service Integrator Native Web Services supports the same standards that are supported by Oracle WebLogic Server. For detailed information on WSS, see the Oracle WebLogic Server document, *Configuring Security: Updating a Client Application to Invoke a Message-Secured Web Service*.

The supported standards include:

- SOAP Message Security
- Username Token Profile
- X.509 Certificate Token Profile
- SAML Token Profile

To use Web Services Security with an Oracle Data Service Integrator web services application:

1. Choose the type of web services security you want to use with your Oracle Data Service Integrator application.
2. Configure security policies through the appropriate policy file(s). See the Oracle WebLogic Server document *Configuring Security: Overview of Web Services Security* for detailed information on configuring policy files for each type of web services security.
3. Edit the web services mapping file to include your policy file(s). You can associate policies with an entire mapping file or for specific operations within the file. See [Section 10.4.4, "Specifying Policies"](#) for details.

10.4.4 Specifying Policies

You can specify policies for a map file or for individual operations in a map file.

10.4.4.1 Specifying Global Policies

To specify a policy for web services security for a map file:

1. Create the policy file. See the Oracle WebLogic Server document *WebLogic Web Services: Security* for detailed information on configuring policy files for each type of web services security.
2. Import the policy file into your Oracle Data Service Integrator project. The easiest way to do this is to use the IDE to import the file as a resource. The policy file must reside in the DSP-INF/policies directory.
3. Configure the web services map file to include the policy.

The following listing shows an example `.ws` file that includes the optional, top-level policies element. Each policy element describes one policy file. The policies element can contain one or more policy elements. The locator attribute contains either an Oracle Data Service Integrator locator for the policy file or a fixed URI that describes the location of the standard WLS policy file.

Oracle Data Service Integrator supports three security policy types. Their URIs are: `policy: Auth`, `policy: Encrypt`, and `policy: Sign`. These are abstract policy files

provided by Oracle WebLogic Server that describe authentication, encryption, and digital signature policies. These policy files do not have to physically reside in DSP project repository.

The policy element contains a required attribute `Direction`. This attribute represents at which direction the security policy will apply. The policy direction can be: `REQUEST`, `RESPONSE`, or `REQUEST_RESPONSE`.

- **REQUEST** - The policy applies only to the inbound request message.
- **RESPONSE** - The policy applies only to the response message.
- **REQUEST_RESPONSE** - The policy applies to both inbound request and the response message.

Refer to the schema definition for detailed information on the structure of the map file (see [Section 10.5, "Web Services Map File Reference"](#)).

Example 10–1 Sample Map File

```
<?xml version="1.0" encoding="UTF-8"?>
<web:WebServicesMap targetNamespace="ld:myMapper.ws" soapVersion="SOAP_1.1"
transportType="HTTP" ADODotNETEnabled="false" basicAuthRequired="false"
xmlns:web="http://www.oracle.com/odsi/management/configuration/webservices">
  <web:policies>
    <web:policy locator="ld:mypolicy.xml">
      <web:policy direction="REQUEST_RESPONSE">
    </web:policies>
  <web:dataServices>
    <web:dataService locator="ld:CUSTOMER.ds">
      <web:function name="deleteCUSTOMER" arity="1"
operation="deleteCUSTOMER" returnInHeader="false">
        <web:parameterMapping>
          <web:parameter name="p" wsdlMapping="SOAP_BODY"/>
        </web:parameterMapping>
      </web:function>
      <web:function name="updateCUSTOMER" arity="1"
operation="updateCUSTOMER" returnInHeader="false">
        <web:parameterMapping>
          <web:parameter name="p" wsdlMapping="SOAP_BODY"/>
        </web:parameterMapping>
      </web:function>
    </web:dataService>
  </web:dataServices>
</web:WebServicesMap>
```

10.4.4.2 Specifying Policies for a Function

To specify policies for a function in a map file:

1. Follow the same basic instructions for specifying a policy for a web service map file, described previously.
2. In the `.ws` file, add the `policies` element to the function element. The `policies` element contains one or more policy element. A policy element represents the security policy that applies to the WSDL operation. The optional child element `ParameterMapping` for the function element contains a list of parameters that are mapped to the SOAP header.

10.5 Web Services Map File Reference

The web services map file is an XML file that provides an explicit mapping between Oracle Data Service Integrator data service functions and web service operations. The map file is the basis for generating the WSDL that describes the web services interface for a data service. This section discusses the configurable parts of the map file in detail. For details on creating a map file, see [Section 10.3, "Generating a Web Service Map and WSDL from a Data Service."](#)

This section describes the following topics:

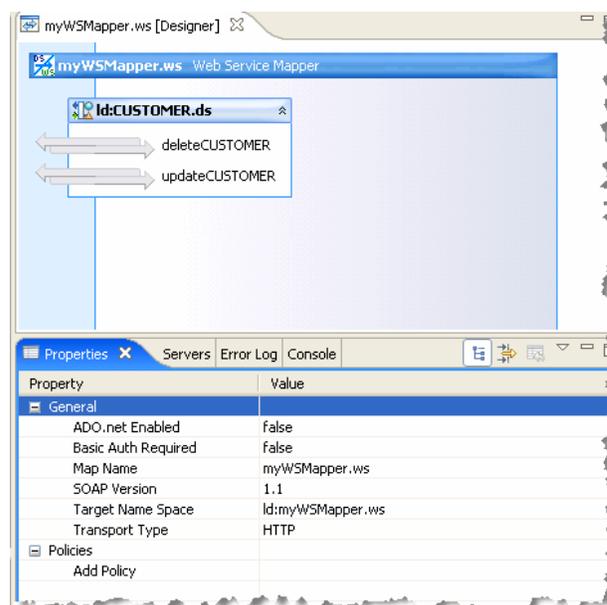
- [Section 10.5.1, "Map File-Level Properties"](#)
- [Section 10.5.2, "Operation Level Properties"](#)
- [Section 10.5.3, "Map File XML Schema Definition"](#)
- [Section 10.5.4, "Mapping of Data Service Type to WSDL Message Type"](#)
- [Section 10.5.5, "Examining the Generated WSDL"](#)
- [Section 10.5.6, "Testing the Generated WSDL"](#)
- [Section 10.5.7, "Copying and Saving a WSDL Generated from a Map"](#)

10.5.1 Map File-Level Properties

The Oracle Data Service Integrator Eclipse IDE lets you create the web services map file (as explained in [Section 10.3, "Generating a Web Service Map and WSDL from a Data Service"](#)) and configure the map file. The New Web Service Map wizard creates a .ws file in a specified location within the Dataspace project. This section describes the configurable map file properties. To configure these properties, use the Properties editor in the IDE.

[Figure 10–2](#) shows a sample map file that maps functions from a data service called **Customer.ds**. To view properties for a map file, select the map file in the IDE and select **Window > Show View > Properties**.

Figure 10–2 *Map File-Level Properties*



This graphic displays a sample map file.

Table 10–1 describes each of the map file properties.

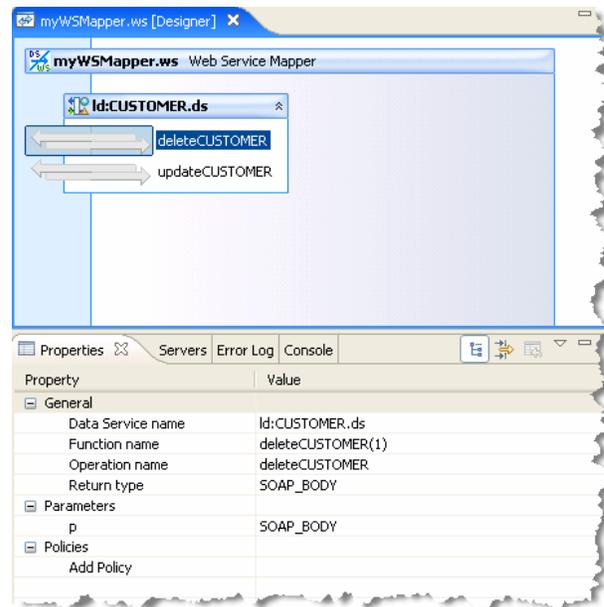
Table 10–1 Map File Properties

Property	Description
ADO.net Enabled	If enabled, a .NET style WSDL is generated. This WSDL includes .NET datasets in the WSDL construct. Disabled by default. For more information on ADO.NET, see the <i>Client Application Developer's Guide</i> .
Basic Auth Required	If true, basic authentication is required to access the WSDL operations.
Map Name	(Read-only) The name of the map file.
SOAP Version	SOAP 1.1 and 1.2 are supported. The version is used by Oracle Data Service Integrator to decide which kind of SOAP binding to create during WSDL generation. The default is 1.1. SOAP 1.2 encoding is not supported. Encoding is an optional feature defined by the SOAP 1.2 specification.
Target Name Space	The default value is generated from the web service based on the location of the map file and the file name.
Transport Type	HTTP and HTTPS are the only supported types. Default is HTTP.
Policies	Lets you specify security policies that apply to all the functions in the map. For information on policies, see Section 10.4 , "Configuring Security for Web Services Applications."

10.5.2 Operation Level Properties

This section describes the operation-level properties that you can modify in the IDE. Operations match up with data service functions. Each data service function maps to a WSDL operation. Operation-level properties apply to the specific operation only.

Figure 10–3 shows the properties displayed for a selected data service function. To view properties for a data service operation, select the operation in the IDE and select **Window > Show View > Properties**.

Figure 10–3 Operation-Level Properties

This graphic displays properties for a selected data service function.

Table 10–2 describes each of the operation properties.

Table 10–2 Operation Properties

Property	Description
Data Service Name	Read-only.
Function Name	Read-only.
Operation Name	The WSDL operation name that is used to generate a WSDL. This name has to be unique within the map file.
Return Type	Maps the WSDL operation return type to either a SOAP header or body.
Parameters	Lists all parameters for the operation and lets you map each parameter to either a SOAP header or body.
Policies	Specifies security policies that apply to the operation. For information on policies, see Section 10.4, "Configuring Security for Web Services Applications."

10.5.3 Map File XML Schema Definition

Example 10-1 is the schema file for the map (.ws file) definition.

Example 10–2 Web Services Map File Schema Definition

```
<xs:schema
targetNamespace="http://www.oracle.com/dsp/management/configuration/webservices"
xmlns:tns="http://www.oracle.com/dsp/management/configuration/webservices"
xmlns="http://www.oracle.com/dsp/management/configuration/webservices"
xmlns:xs="http://www.w3.org/2001/XMLSchema"
elementFormDefault="qualified">
```

```
<xs:element name="WebServicesMap">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="policies" type="PoliciesType" minOccurs="0"/>
      <xs:element name="dataServices" type="DataServicesType"/>
    </xs:sequence>
    <xs:attribute name="targetNamespace" type="xs:anyURI" use="required"/>
    <xs:attribute name="soapVersion" type="SoapVersionType" default="SOAP_1.1"/>
    <xs:attribute name="transportType" type="TransportTypeType" default="HTTP"/>
    <xs:attribute name="ADODotNETEnabled" type="xs:boolean" default="false"/>
  </xs:complexType>
</xs:element>

<xs:simpleType name="SoapVersionType">
  <xs:restriction base="xs:string">
    <xs:enumeration value="SOAP_1.1"/>
    <xs:enumeration value="SOAP_1.2"/>
  </xs:restriction>
</xs:simpleType>

<xs:simpleType name="TransportTypeType">
  <xs:restriction base="xs:string">
    <xs:enumeration value="HTTP"/>
    <xs:enumeration value="HTTPS"/>
  </xs:restriction>
</xs:simpleType>

<xs:complexType name="PoliciesType">
  <xs:sequence>
    <xs:element name="policy" type="PolicyType" minOccurs="1"
maxOccurs="unbounded"/>

    </xs:sequence>
  </xs:complexType>

  <xs:complexType name="PolicyType">
    <xs:attribute name="locator" type="xs:string" use="required"/>
    <xs:attribute name="direction" type="PolicyDirectionType" default="REQUEST_
RESPONSE"/>

  </xs:complexType>

  <xs:simpleType name="PolicyDirectionType">
    <xs:restriction base="xs:string">
      <xs:enumeration value="REQUEST"/>
      <xs:enumeration value="RESPONSE"/>
      <xs:enumeration value="REQUEST_RESPONSE"/>
    </xs:restriction>
  </xs:simpleType>

  <xs:complexType name="DataServicesType">
    <xs:sequence>
      <xs:element name="dataService" type="DataServiceType"
minOccurs="1" maxOccurs="unbounded"/>
    </xs:sequence>
  </xs:complexType>

  <xs:complexType name="DataServiceType">
    <xs:sequence>
      <xs:element name="function" type="FunctionType" minOccurs="1"
```

```

        maxOccurs="unbounded" />
    </xs:sequence>
    <xs:attribute name="locator" type="xs:string" use="required" />
</xs:complexType>

<xs:complexType name="FunctionType">
  <xs:sequence>
    <xs:element name="policies" type="PoliciesType" minOccurs="0" />
    <xs:element name="parameterMapping" type="ParameterMappingType"
minOccurs="0" />
  </xs:sequence>
  <xs:attribute name="name" type="xs:string" use="required" />
  <xs:attribute name="arity" type="xs:integer" use="required" />
<xs:attribute name="operation" type="xs:string" use="required" />
<xs:attribute name="returnInHeader" type="xs:boolean" default="false" />
</xs:complexType>

<xs:complexType name="ParameterMappingType">
  <xs:sequence>
    <xs:element name="parameter" type="ParameterType" minOccurs="1"
maxOccurs="unbounded" />
  </xs:sequence>
</xs:complexType>

<xs:complexType name="ParameterType">
  <xs:attribute name="name" type="xs:string" use="required" />
<xs:attribute name="wsdlMapping" type="WSDLMappingType" use="required" />
</xs:complexType>

<xs:simpleType name="WSDLMappingType">
  <xs:restriction base="xs:string">
    <xs:enumeration value="SOAP_HEADER" />
    <xs:enumeration value="SOAP_BODY" />
  </xs:restriction>
</xs:simpleType>
</xs:schema>

```

10.5.4 Mapping of Data Service Type to WSDL Message Type

This section explains how data service types are mapped to WSDL message types when you map a data service function to a WSDL operation.

- [Two Schema Elements Per Function](#)
- [Mapping of Update Functions with DataGraphs](#)
- [Overloading Data Service Functions](#)

10.5.4.1 Two Schema Elements Per Function

For each data service function, two WSDL schema elements are generated. The first element is the name of the request message, and it is the same as the data service function name that is mapped to the WSDL message. The second represents the response message. The response message name is the same as the function name with "Response" appended to it. The following listing shows an example schema where `getCustomer` is the request name and `getCustomerResponse` is the response name. The response element contains the return type of the data service function, which can be complex or simple.

Example 10–3 Operation Element and Return Element

```

<types>
  <xsd:schema targetNamespace="ld:DataServices/RTLServices/Customer.ws"
             xmlns:dsns0="urn:retailerType">
    <xsd:import namespace="urn:retailerType" />

    <xsd:element name="getCustomer">
      <xsd:complexType>
        <xsd:sequence/>
      </xsd:complexType>
    </xsd:element>

    <xsd:element name="getCustomerResponse">
      <xsd:complexType>
        <xsd:sequence>
          <xsd:element ref="dsns0:CUSTOMER_PROFILE" minOccurs="0"
                      maxOccurs="unbounded" />
        </xsd:sequence>
      </xsd:complexType>
    </xsd:element>
  </types>

```

10.5.4.2 Mapping of Update Functions with DataGraphs

This section explains how a data service update operation's parameters and return type are mapped to a WSDL schema definition.

Consider the following data service definition for an operation called updateADDRESS:

```

(::pragma function <f:function xmlns:f="urn:annotations.ld.oracle.com"
visibility="public" kind="update" isPrimary="true" nativeName="ADDRESS"
nativeLevel2Container="RTLCUSTOMER" style="table">

```

```

<nonCacheable/> </f:function::)

```

```

  declare procedure f1:updateADDRESS($p as changed-element(t1:ADDRESS)*)
as empty() external;

```

Note that the operation's parameter type is changed-element(UserType). In this case the element is ADDRESS. The changed-element type is translated to a DataGraph in the WSDL schema. The WSDL schema must also include a schema definition for the DataGraph. The following listing shows the translated updateADDRESS operation and the schema definition for the DataGraph.

Example 10–4

```

<xs:element name="updateADDRESS">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="p">
        <xs:complexType>
          <xs:sequence>
            <xs:element ref="dsns0:ADDRESSDataGraph" minOccurs="0" maxOccurs="unbounded" />
          </xs:sequence>
        </xs:complexType>
      </xs:element>
    </xs:sequence>
  </xs:element>

```

```

        </xs:complexType>
    </xs:element>
<xs:element name="updateADDRESSResponse">
<xs:complexType>
    <xs:sequence />
</xs:complexType>
</xs:element>
...
<xs:schema targetNamespace="ld:ADDRESS" xmlns:dsns0="ld:ADDRESS"
    xmlns:sdo="commonj.sdo">
    <xs:import namespace="commonj.sdo"
        schemaLocation="http://www.osoa.org/sdo/2.1/schemas/datagraph.xsd" />
    <xs:element name="ADDRESSDataGraph" type="dsns0:ADDRESSDataGraphType" />
<xs:complexType name="ADDRESSDataGraphType">
<xs:complexContent>
<xs:extension base="sdo:BaseDataGraphType">
<xs:sequence>
    <xs:element ref="dsns0:ADDRESS" />
</xs:sequence>
</xs:extension>
</xs:complexContent>
</xs:complexType>
</xs:schema>

```

10.5.4.3 Overloading Data Service Functions

Data service functions can be overloaded, meaning that two functions in the same data service have the same name but a different number of parameters. For example, in the following listing two `getCustomer()` functions are declared, each with a different parameter set. To support WSDL generation for overloaded data service functions, the web services map requires the overloaded function to be mapped to a different WSDL operation. In other words, if you drag two functions with the same name from a data service onto a web service map file, Oracle Data Service Integrator generates different WSDL operation names for the two functions. You can accept the default names or change them.

Example 10–5 Overloaded Functions

```

declare function ns9:getCustomer() as element(ns2:CUSTOMER_PROFILE)*

declare function ns9:getCustomer($customerID as xs:string) as
element(ns2:CUSTOMER_PROFILE)*

```

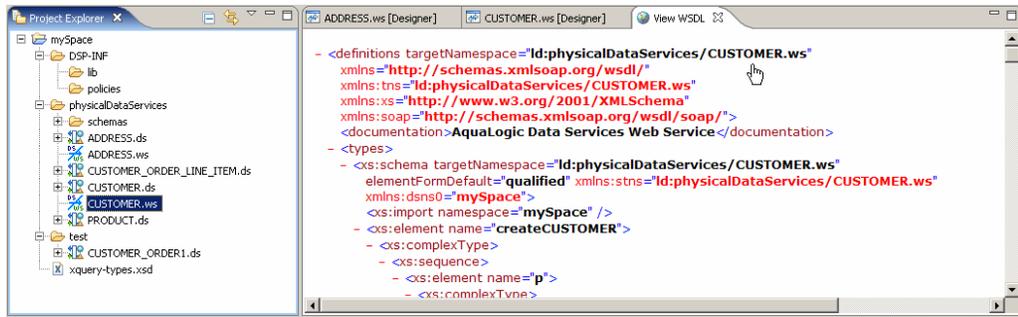
10.5.5 Examining the Generated WSDL

You can examine the generated WSDL file. The dataspace project's associated Oracle WebLogic server must be started and the dataspace project be deployed to the server to view the WSDL or test the Web Service.

1. Right-click on the web service file name (example: CUSTOMER.ws).
2. Choose *View WSDL*.

The WSDL appears in its own window in the work area.

Figure 10–4 View of Generated WSDL



This graphic displays a view of the generated WSDL.

You can also request the WSDL for a deployed project by entering the following URL:

http://host:port/dataSpaceProjectName/folderName/.../mapFileName.ws?WSDL

For example:

http://localhost:7001/myDataSpace/myWMapper.ws?WSDL

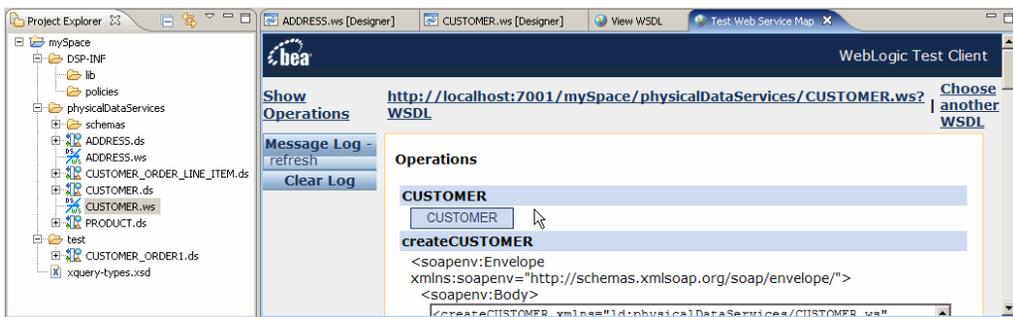
10.5.6 Testing the Generated WSDL

You can test the generated WSDL file using these steps. The dataspace project's associated Oracle WebLogic server must be started and the dataspace project be deployed to the server to view the WSDL or test the Web Service.

1. Right-click on the web service file name (example: **CUSTOMER.ws**).
2. Choose **Test Web Service**.

The WSDL appears in its own window.

Figure 10–5 View of Tested Web Service



This graphic displays a view of the tested web service.

10.5.7 Copying and Saving a WSDL Generated from a Map

You can copy or save a WSDL by right-clicking the map file and selecting **Copy WSDL URL** or **Save WSDL As**.

10.6 Understanding SQL Maps

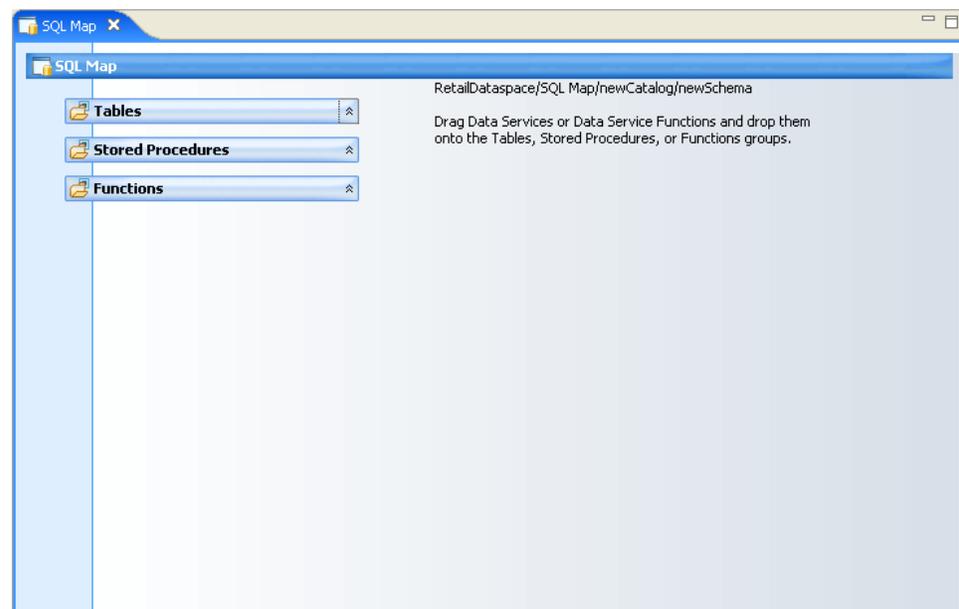
This section describes mapping functions in data services to SQL objects. It describes the following topics:

- [Section 10.6.1, "Overview"](#)
- [Section 10.6.2, "Publishable Operations"](#)
- [Section 10.6.3, "General Conditions"](#)

10.6.1 Overview

A SQL Map enables you to publish data service functions as SQL objects (which are created when you specify the mapping). Using SQL Maps, you can expose data services modeled in Oracle Data Service Integrator as relational data sources. This enables you to use reporting tools (such as Crystal Reports and Microsoft Access, among others), Java applications, and development tools (such as Data Tools Platform or SQL Explorer) to access information from data services using SQL queries (through a JDBC client).

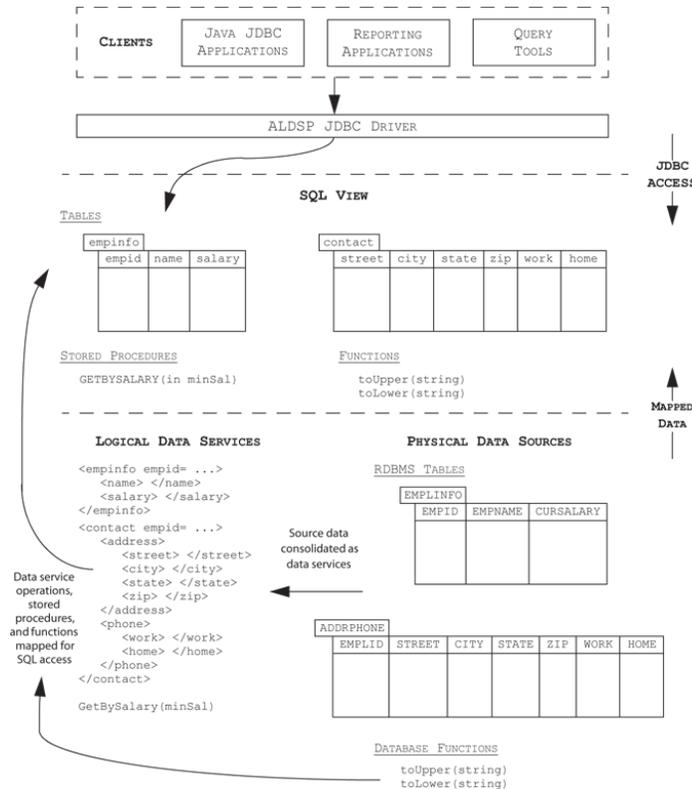
Figure 10–6 Sample SQL Map



This graphic displays a sample SQL map.

As [Figure 10–7](#) shows, source data can be consolidated, integrated, and transformed using Oracle Data Service Integrator data services. The source data itself can come from disparate sources throughout the enterprise, including relational databases and web services, among others. Using SQL Maps you can, in turn, expose the data service operations as a relational data source accessible using SQL queries. This enables JDBC clients to access data consolidated through Oracle Data Service Integrator.

Figure 10–7 SQL Mapping Overview



This graphic displays an overview of SQL mapping.

You can publish the following types of data service artifacts for SQL access:

- Data service functions, either with or without parameters.
- External database functions. These are database-specific functions which are either built into a particular commercial database or which were custom-designed on the database side and then declared to Oracle Data Service Integrator as external database functions.

10.6.2 Publishable Operations

SQL mappable data service functions can be thought of as relationally-compatible XQuery functions. Depending on their signature, you can publish such functions for use as SQL tables or stored procedures. The association between the function and the SQL object is defined at design time when creating a SQL map.

The following summarizes the types of data service functions you can publish as SQL tables or stored procedures:

- You can map non-parameterized data service functions as SQL tables, and parameterized data service functions as stored procedures.
- You cannot map private or protected functions as part of a SQL Map.
- You cannot map procedures as part of a SQL Map.

You can map library data service database functions to functions, but not to SQL tables and stored procedures.

Note: See [Section 10.8, "SQL Object Mapping Rules"](#) for details about permitted mappings.

10.6.3 General Conditions

The following general conditions apply when exposing data service operations as relational data sources:

- The exposed data service XQuery function signatures must only involve types that are supported by the relational (JDBC) type system.
- The structure of the underlying schema (the Return type) must have a relationally-compatible data shape, which means that the data service type cannot include repeating data (data elements with cardinality greater than 1). This is because SQL provides a traditional, two-dimensional approach to data access, as opposed to the multi-level, non-normalized approach defined by XML.

Note: For more information about creating data services with flat schemas, see [Section 1.5.2, "Create a Data Service with a Flat Return Type."](#)

- You cannot map data service operations that return scalar (primitive) values as SQL tables or stored procedures.

10.7 Map Functions and Procedures to SQL Objects

This section describes how to create SQL objects (tables, stored procedures, and database functions) from dataspace project operations including conforming physical and logical functions and procedures. Once created, the objects will be available to client applications through JDBC.

- [Creating an SQL Map](#)
- [Removing an SQL Map](#)

10.7.1 Creating an SQL Map

To create an SQL map:

1. In Eclipse for WebLogic, right-click a dataspace project folder in the Project Explorer and choose **Add SQL Map**. Eclipse for WebLogic creates an SQL Map with a default catalog and schema.

Note: You can define only a single SQL Map for a dataspace. You can, however, add multiple catalogs and schemas to an SQL Map.

You can rename the default catalog and schema by expanding the SQL Map in the Project Explorer, right-clicking the catalog or schema, choosing **Rename**, and entering the new name.

2. Select the schema to which you want to map the data service functions and procedures. Eclipse for WebLogic displays folder tabs for Tables, Stored Procedures, and Functions.

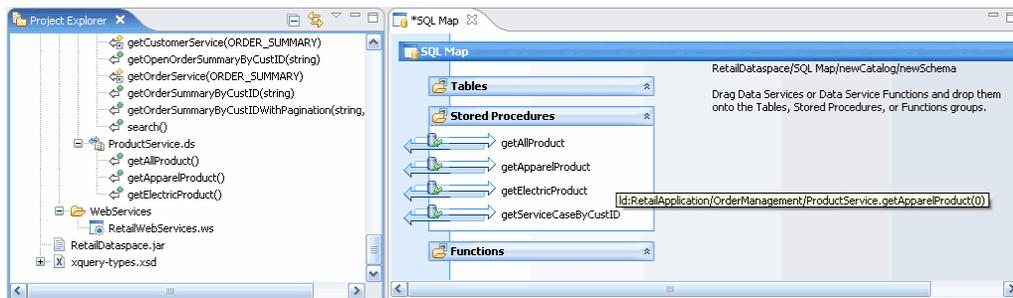
3. Drag-and-drop the data service functions and procedures from the Project Explorer to the corresponding folder tab in the SQL Map.

If you drag-and-drop an entire data service to a folder tab in the SQL Map, Eclipse for WebLogic attempts to map all functions contained in the data service to the corresponding SQL object type (Table, Stored Procedure, or Function).

Note: If you attempt to map a dataspace object which does not meet SQL map criteria, a dialog will appear, explaining the problem.

You can map the same data service function or procedure to multiple schemas. You can also map a function or procedure to multiple SQL object types, however, Eclipse for WebLogic displays an alert dialog (see Map Data Service Functions for SQL Use Alert Dialog) in case of a naming conflict and suggests a new SQL name. You can edit this new name, as required.

Figure 10–8 Populated SQL Map



This graphic displays a populated SQL map.

10.7.2 Removing an SQL Map

To remove an SQL map, right-click the dataspace project folder and choose **Remove SQL Map**.

10.8 SQL Object Mapping Rules

The function and procedure types used in data services map to various types of SQL objects. The general mapping rules are described in [Table 10–3](#).

Table 10–3 Mapping Rules

Type	Element Type	Optional	Functions
Read functions	Yes	Yes	No
Navigation functions	Yes	Yes	No
Private functions	No	No	No
Protected functions	No	No	No
Procedures	No	No	No
Library data service functions (non-database)	Yes	Yes	No
Library data service database functions	No	No	Yes

10.9 Constraints on Publishing Data Service Objects to SQL

There are semantic and structural constraints to publishing data service objects to SQL.

Semantic constraints include some general types of objects as private functions.

Table 10–4 outlines the structural constraints on publishing data service artifacts to SQL.

Table 10–4 Structural Constraints

Limitation	Discussion
Limitation affecting all SQL objects	Limitations in this section affect publication to any type of SQL object.
Functions referring to types that are neither simple nor elements	Examples of such types include item, node, and attribute.
Functions with anonymous element types	Functions containing elements where the name is not defined are not mapable. For example: <pre>declare function f() as element()</pre>
Functions declarations using recursive XML types	For example, a function declaration with a complex type (PersonType) containing an element that is also of type PersonType is not mapable: <pre><element name="PERSON" type="tns:PersonType" /> <complexType name="PersonType"> <sequence> <element name="first_name" type="string" /> <element name="last_name" type="string" /> <element name="contact" type="tns:PersonType" /> </sequence> </complexType></pre>
XML types with content models containing wildcards	XML wildcards include: <ul style="list-style-type: none"> • xs:any • xs:anyAttribute
XML types with mixed content	<pre><a> <child/> this is simply text <child/> </pre>
Limitations affecting publishing as a SQL Table	Limitations in this category affect publishing as SQL tables.
Functions with parameters	Functions with parameters can be mapped as stored procedures.
Functions containing simple return types	Functions containing simple return types can be mapped as SQL functions.

Table 10–4 (Cont.) Structural Constraints

Limitation	Discussion
Functions containing any non-tabular element type	See Section 10.9.1, "Non-Tabular Element Types Affect Ability to Publish Functions as SQL Objects" . Also applies to stored procedures.
Functions with any <code>AtomicType</code> types	Also applies to stored procedures.
Limitations affecting publishing as a stored procedure	Limitations in this category affect publishing as a stored procedure.
Functions accepting element parameter types	These functions cannot be published as stored procedures.
Functions containing a sequence of simple return types, such as <code>xs:string*</code>	The function declaration is not eligible. For example: <pre>declare function f(\$p as xs:string*) as xs:int</pre>
Functions with any <code>AtomicType</code> types	Also applies to tables.
Functions with any non-tabular element types	See Section 10.9.1, "Non-Tabular Element Types Affect Ability to Publish Functions as SQL Objects" . Also applies to tables.
Limitations affecting publishing as a SQL Function	Limitations in this category affect publishing as a SQL functions.
Function with a sequence parameter type and an arity greater than 1.	An example shows <code>xs:int*</code> as the sequence parameter type: <pre>declare function f(\$p as xs:int*, \$q as xs:string) as xs:int</pre>
Functions with element types	 <pre>declare function f (\$p as element(e)) as xs:int</pre>

10.9.1 Non-Tabular Element Types Affect Ability to Publish Functions as SQL Objects

The structure of a data service function determines whether it can be mapped to an SQL object or not. For example, a parameterized function cannot be published as an SQL table since by definition SQL tables do not take parameters. Some structural constraints are practically self-evident; others are less obvious.

Note: A quick way to determine if a particular function can be published to a particular type of SQL object is to drag the function to a SQL object table, stored procedure, or functions folder. Even if the function is grayed out — meaning that it cannot be published to any type of SQL object — an alert dialog will appear explaining why the selected object cannot be published.

For example, functions with non-tabular element types cannot be published as tables or stored procedures because XML output structure cannot be mapped to a normalized SQL table.

Underlying each data service is an XML type, or schema. Some XML types are readily mapped for JDBC use because they are — like SQL tables — two dimensional.

```

<CUSTOMER>
  <FIRST_NAME>
  <LAST_NAME>
  <CUSTOMER_ID>
</CUSTOMER>

```

When published as SQL, the table structure corresponds to the following:

FIRST_NAME	LAST_NAME	CUSTOMER_ID
Jack	Black	CUSTOMER1

As long as the object mapper can reduce the structure of the XML document to rank-one, the mapping can occur. For example:

```

<CUSTOMER>
  <FIRST_NAME>
  <LAST_NAME>
  <CUSTOMER_ID>
  <CUSTOMER_ORDER>
    <ORDER_ID>
    <C_ID>
    <ORDER_DT>
  </CUSTOMER_ORDER>
</CUSTOMER>

```

This is publishable as a table in the following form as long as there is *one or fewer* customer orders associated with the customer:

FIRST_NAME	LAST_NAME	CUSTOMER_ID	ORDER_ID	C_ID	ORDER_DT
Jack	Black	CUSTOMER1	ORDER_1_0	CUSTOMER1	2001-10-01

If, however, the CUSTOMER_ORDER type is unbounded, meaning that it can represent more than one order associated with a single customer, the structure no longer corresponds to a well-formed relational table and the mapping is not allowed.

Data Service Annotations

This chapter describes the syntax and semantics of annotations in data service documents developed within Eclipse for WebLogic. This chapter contains the following sections:

- [Section 11.1, "Overview"](#)
- [Section 11.2, "XDS Annotations"](#)
- [Section 11.3, "Function Annotations"](#)
- [Section 11.4, "XFL Annotations"](#)
- [Section 11.5, "Data Service Annotations Schema"](#)

11.1 Overview

Data service documents define collections of XQuery functions and/or XQSE functions or procedures. Annotations are XML fragments comprising the character content of XQuery pragmas.

There are two types of annotations:

- Global annotations: these pertain to the entire entity or library data service document. Global annotations are also referred to as XDS or XFL annotations respectively.
- Local annotations: these pertain to a particular function. Local annotations are also referred to as function annotations.

11.2 XDS Annotations

There is a single XDS ("XQuery Data Service") annotation per entity data service document, which appears before all function annotations. The identifier for the pragma carrying the XDS annotation is `xds`. The qualified name of the top level element of the XML fragment corresponding to an XDS annotation has the local name `xds` and the namespace URI:

```
urn:annotations.ld.oracle.com
```

Each entity data service is associated with a unique target type. The prime type of the return type of every read function must match its target type. The target type of an entity data service is an element type whose qualified name is specified by the `targetType` attribute of the `xds` element. It is defined in a schema file associated with the entity data service.

The contents of the top-level xds element is a sequence of the following blocks of properties:

- [Section 11.2.1, "General Properties"](#)
- [Section 11.2.2, "Data Access Properties"](#)
- [Section 11.2.3, "Target Type Properties"](#)
- [Section 11.2.4, "Key Properties"](#)
- [Section 11.2.5, "Relationship Properties"](#)
- [Section 11.2.6, "Update Properties"](#)

The following is an example of an XDS annotation. In this case, the target type `t:CUSTOMER` associates the entity data service with a `t:CUSTOMER` type in a schema file.

```
(::pragma xds <x:xds xmlns:x="urn:annotations.ld.oracle.com"
targetType="t:CUSTOMER" xmlns:t="ld:oracleDS/CUSTOMER">

<author>Joe Public</author>
<relationalDB name="OracleDS" />

<field type="xs:string" xpath="FIRST_NAME">
  <extension nativeFractionalDigits="0" nativeSize="64"
    nativeTypeCode="12" nativeType="VARCHAR2"
    nativeXPath="FIRST_NAME" />
  <properties nullable="false" />
</field>

<field type="xs:string" xpath="LAST_NAME">
  <extension nativeFractionalDigits="0" nativeSize="64"
    nativeTypeCode="12" nativeType="VARCHAR2"
    nativeXPath="LAST_NAME" />
  <properties nullable="false" />
</field>

<field type="xs:string" xpath="CUSTOMER_ID">
  <extension nativeFractionalDigits="0" nativeSize="64"
    nativeTypeCode="12" nativeType="VARCHAR2"
    nativeXPath="CUSTOMER_ID" />
  <properties nullable="false" nativeKey="true" />
</field>

<field type="xs:dateTime" xpath="CUSTOMER_SINCE">
  <extension nativeFractionalDigits="0" nativeSize="7"
    nativeTypeCode="93" nativeType="DATE"
    nativeXPath="CUSTOMER_SINCE" />
  <properties nullable="false" />
</field>

<field type="xs:string" xpath="EMAIL_ADDRESS">
  <extension nativeFractionalDigits="0" nativeSize="32"
    nativeTypeCode="12" nativeType="VARCHAR2"
    nativeXPath="EMAIL_ADDRESS" />
  <properties nullable="false" />
</field>

<key name="CUSTOMER_ID" />

<relationshipTarget roleName="CUSTOMER_ORDER" roleNumber="2"
  XDS="ld:oracleDS/CUSTOMER_ORDER.xds" minOccurs="0"
```

```

    maxOccurs="unbounded" opposite="CUSTOMER"/>
  </x:xds>:::-)

```

11.2.1 General Properties

There are two types of general XDS properties:

- [Section 11.2.1.1, "Standard Document Properties"](#)
- [Section 11.2.1.2, "User-Defined Properties"](#)

11.2.1.1 Standard Document Properties

You can specify a set of standard document properties consisting of optional XML elements containing information pertaining to the author, creation date, or version of the document. You can also use the optional element named "documentation" to specify related documentation. The names and types of the elements in the standard document properties block, as well as examples of their use, are shown in [Table 11–1](#).

Table 11–1 Standard Document Properties

Element Name	Element Type	Optional	Example Instance
author	xs:string	Yes	<author>J. Public</author>
creationDate	xs:date	Yes	<creationDate>2004-05-31</creationDate>
version	xs:decimal	Yes	<version>2.2</version>
documentation	xs:string	Yes	<documentation> Models an online Customer </documentation>

11.2.1.2 User-Defined Properties

In addition to the standard properties, you can specify custom properties pertaining to the entire data service document using a sequence of zero (0) or more "property" elements. Each property element must be named using its "name" attribute and may contain any string content. For example:

```
<property name="data-refresh-rate">week</property>
```

11.2.2 Data Access Properties

A data service may be used to model access to an external data source or to model a transformation on top of one or more data sources or other transformations. Data services modeling external data sources are referred to as *physical*. Transformation data services not representing a particular data source are referred to as *logical*.

The block of data access properties allows each data service to define whether it is physical or not. When a data service is physical, the data access annotation describes the type of the external source being accessed by its external functions (there may be a single external source per data service) and its connection properties. When a data service is logical, the data service is designated as a user-defined view, and no connection information is required.

The following types of physical data services are supported:

- Relational
- Web service
- Java function

- Delimited content
- XML content

The following sections describe the data access annotation for the physical data service types, as well as for data services that are designated as user-defined views. You can specify only one of these annotations in each data service. If no annotation is provided, the data service is considered a user-defined view.

- [Section 11.2.2.1, "Relational Data Service Annotations"](#)
- [Section 11.2.2.2, "Source Binding Provider"](#)
- [Section 11.2.2.3, "Web Service Data Service Annotations"](#)
- [Section 11.2.2.4, "Java Function Data Service Annotations"](#)
- [Section 11.2.2.5, "Delimited Content Data Service Annotations"](#)
- [Section 11.2.2.6, "XML Content Data Service Annotations"](#)
- [Section 11.2.2.7, "User Defined View XDS Annotations"](#)

11.2.2.1 Relational Data Service Annotations

The data access annotation for a relational data service consists of the element `relationalDB` with two required attributes, described in [Table 11-2](#):

Table 11-2 Required Attributes for the `relationalDB` Element

Attribute	Description
<code>name</code>	The JNDI name by which the external relational data source has been registered with the application server.
<code>providerId</code>	The identifier of the Oracle Data Service Integrator relational provider in use for the specified relational data source.

```
<relationalDB name="OracleDS" providerId="Oracle-9" />
```

In addition, the `relationalDB` element can contain the following optional parts:

- An optional element, `properties`, that exposes relational provider-specific attributes, such as the values of specific settings of the Relational Database Management System (RDBMS) represented by the relational source.
- An optional attribute, `sourceBindingProviderClassName`, that specifies the transformation used to determine the relational source to be used at system runtime in the place of the statically defined source.

11.2.2.2 Source Binding Provider

The value of the optional `sourceBindingProviderClassName` attribute should be bound to the fully-qualified name of a user-defined Java class implementing the interface:

```
com.bea.ld.bindings.SourceBindingProvider
```

defined by the following:

```
package com.bea.ld.bindings;
public interface SourceBindingProvider
{
    public String getBinding(String genericLocator, boolean isUpdate);
}
```

The user-defined implementation should provide the transformation that, given the statically configured relational source name (parameter `genericLocator`) and a Boolean flag indicating whether the relational source is accessed in query or update mode (parameter `isUpdate`), determines the name of the relational source name used by the system at runtime.

You can use this transformation mechanism to perform credential mapping. In this case, a single set of query or update operations to be performed in the name of two distinct users *U1* and *U2* against the same statically-configured relational source *R0*, is executed against two distinct relational sources *R1* and *R2* respectively (where all sources *R0*, *R1*, *R2* represent the same RDBMS and the security policies applied to the connection credentials used for *R1* and *R2* correspond to the security policies applied to the application credentials of user *U1* and *U2*, respectively).

Note: Set the source binding provider name uniformly across all relational data services sharing the same relational source JNDI name. Although this restriction is not enforced, its violation could result in unpredictable behavior at runtime.

11.2.2.3 Web Service Data Service Annotations

The data access annotation for a data service based on a Web service consists of the empty element `webService` with two required attributes, described in [Table 11-3](#):

Table 11-3 Required Attributes for the `webService` Element

Attribute	Description
<code>wSDL</code>	A valid <code>http:</code> or <code>ld:</code> URI pointing to the location of the WSDL file containing the definition of the external Web service source. Note: You must configure the Eclipse proxy information in order to access the external web service. Refer to the Eclipse documentation for details.
<code>targetNamespace</code>	A valid URI that is identical to the <code>targetNamespace</code> URI of the WSDL.

For example:

```
<webService targetNamespace="urn:GoogleSearch"
  wSDL="ld:google/GoogleSearch.wSDL" />
```

In addition, if the physical data service models an Oracle Service Bus proxy service, the `webService` element can carry the optional attributes described in [Table 11-4](#):

Table 11-4 Optional Attributes for the `webService` Element

Attribute	Description
<code>sbProxyServiceName</code>	The name of the Oracle Service Bus proxy service.
<code>sbTransportProtocol</code>	The name of the protocol used by the proxy service. Valid values are: <code>t3</code> , <code>iiop</code> , <code>http</code> , <code>t3s</code> , <code>iiops</code> or <code>https</code> .

11.2.2.4 Java Function Data Service Annotations

The data access annotation for a Java function data service consists of the empty element `javaFunction` with a single required attribute named `class`, whose value you set to the fully qualified name of the Java class serving as the external source.

For example:

```
<javaFunction class="com.example.Test" />
```

11.2.2.5 Delimited Content Data Service Annotations

The data access annotation for a delimited content data service is the empty element `delimitedFile`, accepting the optional attributes described in [Table 11-5](#):

Table 11-5 *Optional Attributes for the delimitedFile Element*

Attribute	Description
file	A valid URI pointing to the location of the delimited file.
schema	A valid URI pointing to the location of the XML schema file defining the type (structure) of the delimited contents. If absent, the schema is derived based on the contents.
inferredSchema	Specifies whether the schema was inferred or provided by the user. The default value is false.
delimiter	The string used as the delimiter. If absent, the <code>fixedLength</code> attribute should be present.
fixedLength	The fixed length of the tokens contained in fixed length content. If absent, the <code>delimiter</code> attribute should be present.
hasHeader	A Boolean flag indicating whether the first line of the content should be interpreted as a header. The default value is false.

For example:

```
<delimitedFile schema="ld:df/schemas/ALL_TYPES.xsd" hasHeader="true"
  delimiter="," file="ld:df/ALL_TYPES.csv" />
```

11.2.2.6 XML Content Data Service Annotations

The data access annotation for an XML content data service is the empty element `xmlFile` accepting the attributes described in [Table 11-6](#).

Table 11-6 *Attributes for the xmlFile Element*

Attribute	Description
file	(Optional) A valid URI pointing to the location of the XML file.
schema	A valid URI pointing to the location of the XML schema file defining the type (structure) of the XML contents.

For example:

```
<xmlFile schema="ld:xml/somewhere/CUSTOMER.xsd"
  file="ld:xml/CUSTOMER_NESTED.xml" />
```

11.2.2.7 User Defined View XDS Annotations

The data access annotation for a user-defined view data service is also known as a logical data service. It consists of the single empty element:

```
userDefinedView
```

For example:

```
<userDefinedView/>
```

11.2.3 Target Type Properties

The optional block of target type properties enables you to annotate simple valued fields in the target type of the entity data service with native type information pertaining to the following:

- The type of the corresponding field in the underlying external source (applicable only to data source data services)
- Information about the field's properties with respect to its update behavior. Each annotated field is represented by the element named "field" with two required attributes, described in [Table 11-7](#):

Table 11-7 Required Attributes for the field Element

Attribute	Description
xpath	An XPath value pointing to the field.
type	The qualified name of the field's simple XML schema or XQuery type.

The following is an example of a field element definition:

```
<field type="xs:string" xpath="FIRST_NAME">
  <extension nativeSize="64" nativeTypeCode="12" nativeType="VARCHAR2"
    nativeXPath="FIRST_NAME" />
  <properties nullable="false" />
</field>
```

11.2.3.1 Native Type Properties

Each "field" element can contain an optional "extension" element that accepts the optional attributes described in [Table 11-8](#):

Table 11-8 Optional Attributes for the extension Element

Attribute	Description
nativeXPath	A native XPath value pointing to the corresponding native field in the external source.
nativeType	The native name of the native type of the corresponding native field, as it is known to the external source.
nativeTypeCode	The native type code of the native type of the corresponding native field, as it is known to the external source. In the case of relational sources, this is the type code as reported by JDBC.
nativeSize	The native size of the native type of the corresponding native field, as it is known to the external source. In the case of relational sources, this is the size as reported by JDBC.
nativeFractionalDigits	The native scale of the native type of the corresponding native field, as it is known to the external source. In the case of relational sources, this is the scale as reported by JDBC.
nativeKey	A Boolean value indicating whether the field participates in the native record's key. The default is false.

11.2.3.2 Update-related Type Properties

Each "field" element can also contain an optional "properties" element that accepts the optional attributes described in [Table 11–9](#):

Table 11–9 *Optional Attributes for the properties Element*

Attribute	Description
immutable	A Boolean value specifying whether the field is immutable (read-only) or not. The default value is false.
nullable	A Boolean value specifying whether the field accepts null values or not. The default value is false.

11.2.4 Key Properties

The optional block of key properties enables you to specify an identity constraint (key) on the entity data service target type. An identity constraint for an entity data service is represented by the element "key" along with an XML schema specifying the key type.

The "key" element accepts a required attribute "type", whose value should be bound to the qualified name of the element type defining the locations of the data fields comprising the key. The key type should in turn be specified by an XML schema imported by the data service.

The "key" element may also carry the optional attributes in [Table 11–10](#):

Table 11–10 *Optional Attributes for the key Element*

Attribute	Description
name	Serves as the key alias. Might be used as a user-friendly description of the semantic constraints expressed by the key.
inferred	A Boolean value specifying whether the key was auto-derived or user-defined. The default is true.
inferredSchema	A Boolean value specifying whether the key schema was auto-derived or user-defined. The default is true.

In most cases, the identity constraint refers to the collection of data bindings returned by the entity data service's read functions, with each binding's type being the data service target type. In the case that a data service returns an XML document, the collection on which the identity constraint may be specified is normally defined by some element nested within the document element. In such a case, the "key" element contains an optional "selector" element that is used to specify the collection. The "selector" element carries a required "xpath" attribute, whose value is an XPath value pointing to the nested element defining the collection root. The XPath forms accepted by this attribute are simplified XPaths, using only the element or attribute axes and no predicates.

The following is an example of a "key" element definition:

```
<key name="CUSTOMER_ID" />
  <selector xpath="CUSTOMER" />
</key>
```

11.2.5 Relationship Properties

The optional block of relationship properties enables you to specify a set of relationship targets. A relationship target of an entity data service is an entity data service with which first service maintains a unidirectional or bidirectional relationship. Unidirectional relationships are realized through one or more navigate functions in the first data service that returns one or more instances of objects of the second service target type. Bidirectional relationships require that reciprocal functions are present in the second data service as well.

A relationship target is represented by the element `relationshipTarget` that accepts the attributes described in [Table 11–11](#):

Table 11–11 *Attributes for the relationshipTarget Element*

Attribute	Description
<code>roleName</code>	A string that uniquely identifies the relationship target inside the data service.
<code>roleNumber</code>	(Optional) Either 1 or 2 (default is 1). The <code>roleNumber</code> specifies the index of the relationship target within the relationship.
<code>XDS</code>	The Oracle Data Service Integrator URI of the data service serving as the relationship target.
<code>minOccurs</code>	(Optional) The minimum cardinality of relationship target instances participating in this relationship. Possible values are all non-negative integers and the empty string. The default value is the empty string.
<code>maxOccurs</code>	(Optional) The maximum cardinality of relationship target instances participating in this relationship. Possible values are all positive integers, the string unbounded, and the empty string. The default is the empty string.
<code>opposite</code>	(Optional) String attribute that indicates the reciprocal relationship target in the case of bidirectional relationships. The value of this attribute is the identifier used to identify this data service as a relationship target in the data service identified by the value of the <code>XDS</code> attribute.

Additionally, the `relationshipTarget` element can itself contain the element "relationship" which in turn contains the nested element "description" that contains a human readable description about the relationship.

The following is an example of a `relationshipTarget` element definition:

```
<relationshipTarget roleName="CUSTOMER_ORDER" roleNumber="2"
  XDS="ld:oracleDS/CUSTOMER_ORDER.xds" minOccurs="0"
  maxOccurs="unbounded" opposite="CUSTOMER"/>
```

11.2.6 Update Properties

The optional block of update properties enables you to specify a set of properties that establish certain policies about updating an entity data service's underlying sources. In particular, you can specify the following policies:

- [Section 11.2.6.1, "Optimistic Locking Fields"](#)
- [Section 11.2.6.2, "Security Properties"](#)

11.2.6.1 Optimistic Locking Fields

SDO update assumes optimistic locking transactional semantics. The data service being updated can specify the fields that should be checked for updates during the interim using the empty element `optimisticLockingFields` that accepts one of the following as its content:

- An empty element, named `updated`, to specify only updated fields.
- An empty element, named `projected`, to specify all projected fields.
- One or more elements, named "field", that accept a required string-valued attribute named `name` to specify user-specified fields.

The following is an example of a `functionForDecomposition` element definition:

```
<optimisticLockingFields>
  <updated/>
</optimisticLockingFields>
```

11.2.6.2 Security Properties

You can use a data service to define one or more user-defined, logical protected resources.

The element `secureResources`, containing one or more string-valued elements named `secureResource`, can be used for this purpose.

For example:

```
<secureResources>
  <secureResource>MyResource</secureResource/>
  <secureResource>MyOtherResource</secureResource/>
</secureResources>
```

You can link a logical resource defined using this syntax to a user-provided security policy using the Oracle Data Service Integrator Console. Query content can inquire about a user's ability to access a logical resource using the built-in function `isAccessAllowed()`.

11.3 Function Annotations

There is a single function annotation per data service function or procedure, which appears before the function or procedure declaration in the document. The identifier for the pragma carrying the function annotation is "function". The qualified name of the top level element of the XML fragment corresponding to a function annotation has the local name "function" and the namespace URI `urn:annotations.ld.oracle.com`.

Modeling Kind

Each entity data service function or procedure is classified using one of the following categories:

- Create procedure
- Read function
- Update procedure
- Delete procedure
- Navigate function

- Library function or procedure

The classification of a data service method is determined by the value of the optional attribute "kind" in the function element, which accepts the values create, read, update, delete, navigate, or library to denote the corresponding categories. The default value is library.

Each library data service function or procedure is always of kind library.

The prime type of the return type of a read function must match the target type of the entity data service. In addition, the function element for a navigate function must carry a string-valued attribute returns whose value must match the role name of a relationship target defined in the data service. Moreover, the prime type of the return type of a navigate function must match the target type of the data service serving as the relationship target.

An operation designated as a procedure has in the general case side-effects. In other words, its invocation entails modifications of the state of the affected data sources. Therefore, a procedure may not be referenced by Oracle Data Service Integrator functions.

A library function residing in a relational database function library data service file is always external. It may not be invoked directly by clients. Instead, it should be referenced by other data service functions or ad-hoc queries.

Visibility

Functions or procedures may also be classified based on their visibility using one of the following categories:

- Public
- Protected
- Private

The classification of a data service method is determined by the value of the optional attribute "visibility" in the function element, which accepts the values public, protected, or private to denote the corresponding categories. The default value is protected.

Public methods are accessible by Oracle Data Service Integrator dataspace clients as well as other data services within the dataspace.

Protected methods are not accessible by Oracle Data Service Integrator dataspace clients but can be accessed by other data services within the dataspace.

Private methods may be accessed only by other methods within the data service in which they are defined.

Primary

The optional boolean attribute `isPrimary` may also be used to classify entity data service methods as primary or non-primary. The default value is false.

This property is applicable only to create, update and delete procedures or read functions.

In the case of a procedure, when this property is set to true, it denotes that the procedure should be the one to be automatically used by the update maps of logical data services directly depending on the data service defining the procedure, in order to perform the corresponding update operation (that is, create, update or delete).

In the case of a read function, when this property is set to true, it denotes that the read function should be the one to be used to infer the data service update map.

There may exist at most one primary method of each kind specified within an entity data service.

URI

Finally, the namespace URIs of the qualified names of all the functions and/or procedures in a data service must specify the location of the data service document in the Oracle Data Service Integrator repository. For example:

```
ld:{directory path to data service folder}/{data service file name without extension}
```

The function element accepts the additional optional attributes described in [Table 11–12](#).

Table 11–12 *Optional Attributes for the function Element*

Attribute	Description
nativeName	Applicable to data source functions or procedures, nativeName is the name of the function or procedure as it is known to the external source. In the case of relational sources, for example, it corresponds to the table name.
nativeLevel1Container	Applicable to data source functions or procedures that represent external sources employing hierarchical containment schemes; nativeLevel1Container is the name of the top-level native container, as it is known to the external source. In the case of relational sources, for example, it corresponds to the catalog name, whereas, in the case of Web service sources, it corresponds to the service name.
nativeLevel2Container	Applicable to data source functions or procedures that represent external sources employing hierarchical containment schemes; nativeLevel2Container is the name of the second-level native container, as it is known to the external source. In the case of relational sources, for example, it corresponds to the schema name. In the case of Web service sources, it corresponds to the port name.
nativeLevel3Container	Applicable to data source functions or procedures that represent external sources employing hierarchical containment schemes; nativeLevel3Container is the name of the top-level native container, as it is known to the external source. In the case of relational sources, for example, it corresponds to the stored procedure package name.
style	Applicable to data source functions or procedures, style is a native qualifier by which the function is known to the external source (e.g. table, view, storedProcedure, or sqlQuery for relational sources; rpc or document for Web services).
roleName	Applicable to navigate functions, roleName should match the value of the roleName attribute of the relationshipTarget implemented by the function.

The content of the top-level function element is a sequence of the following blocks of properties:

- [Section 11.3.1, "General Properties"](#)
- [Section 11.3.2, "UI Properties"](#)

- [Section 11.3.3, "Cache Properties"](#)
- [Section 11.3.4, "Transaction Properties"](#)
- [Section 11.3.5, "Behavioral Properties"](#)
- [Section 11.3.6, "Polymorphic Functions"](#)
- [Section 11.3.7, "Signature Properties"](#)
- [Section 11.3.8, "Native Properties"](#)
- [Section 11.3.9, "Implementation Properties"](#)

The following is an example of a function annotation:

```
( ::pragma function
<f:function xmlns:f="urn:annotations.ld.oracle.com" kind="read"
nativeName="CUSTOMER" nativeLevel2Container="RTL" style="table">
<nonCacheable/>
</f:function>:::-)
```

11.3.1 General Properties

All standard document properties and user-defined properties defined in [Section 11.2.1.1, "Standard Document Properties"](#) and [Section 11.2.1.2, "User-Defined Properties"](#) are applicable to function annotations.

11.3.2 UI Properties

A set of user interface properties may be introduced by the XQuery Editor to persist location information about the graphical components representing the expression in the function body. UI properties are represented by the element `uiProperties` which accepts a sequence of one or more elements, named `component`, as its content. Each "component" element accepts the attributes described in [Table 11–13](#).

Table 11–13 *Attributes for the component Element*

Attribute	Description
identifier	An identifier for the UI component.
minimized	A Boolean flag indicating whether the UI component has been minimized or not.
x	The x-coordinate for the UI component.
y	The y-coordinate for the UI component.
w	The width of the UI component.
h	The height of the UI component.
viewPosX	The x-coordinate of the scrollbar position of the component.
viewPosY	The y-coordinate of the scrollbar position of the component.

In addition, each "component" element may optionally contain one or more `treeInfo` elements containing information about the tree representation of the types pertaining to the component. In the absence of the above property, the query editor uses the default layout.

11.3.3 Cache Properties

You can use the optional block of cache properties to specify whether a function can be cached or not. You should specify a function whose results for the same set of arguments are intrinsically highly volatile as non-cached. On the other hand, you should specify a function whose results for the same set of arguments are either fixed or remain unchanged for a period of time as cacheable.

This property of a function is represented by the empty element `nonCacheable`. In the absence of the `nonCacheable` element, a function is considered to be potentially cacheable. The following is an example:

```
<nonCacheable/>
```

11.3.4 Transaction Properties

You can use the optional block of transaction properties to specify whether a procedure can participate in a transaction or not. This property is applicable only to physical procedures bound to external data sources of type Java or Oracle Service Bus proxy service. A transactional procedure should rollback its effects if the overall transaction, in which it participates, fails.

This property is represented by the empty element `nonTransactional`. In the absence of the `nonTransactional` element, a procedure is considered to be transactional. The following is an example:

```
<nonTransactional/>
```

11.3.5 Behavioral Properties

The optional block of behavioral properties allows you to provide information related to known associations between a function's input and its output, or across two or more functions. In particular, you may specify the following:

- [Section 11.3.5.1, "Inverse Functions"](#)
- [Section 11.3.5.2, "Equivalent Transforms"](#)

11.3.5.1 Inverse Functions

Given an XQuery function f , the optional block of inverse functions may be used in order to denote a function g , defined over the range of f , that, when composed with f (that is, $g(f)$), renders one of the parameters of f . If f has multiple parameters, an inverse function may be defined for each one of its parameters.

The inverse functions block is represented by an optional element, named `inverseFunctions`, which accepts as its content a sequence of empty elements, named `inverseFunction`. Each `inverseFunction` element accepts the following attributes:

- `parameterIndex`: optional attribute denoting the index of the parameter for which the inverse function is defined. The index of the first parameter is assumed to be 1. It may be omitted if the function being annotated has a single parameter.
- `name`: required attribute denoting the fully-qualified name of the inverse function.

Note: Both the annotated and the inverse function must be either built-in or external XQuery functions.

The following is an example of an `inverseFunctions` element definition:

```
<inverseFunctions>
  <inverseFunction parameterIndex="2" name="p:MyInverse" xmlns:p="urn:test"/>
</inverseFunctions>
```

11.3.5.2 Equivalent Transforms

Given an XQuery function: f , the optional block of equivalent transforms may be used in order to denote a pair of functions C_1 and C_2 with identical signatures and equivalent semantics, that accept f as one of their parameters. In simple terms, the equivalence is perceived to mean that each occurrence of $C_1(\dots f, \dots)$ may be safely substituted with: $C_2(\dots f, \dots)$.

The equivalent transforms block is represented by an optional element, named `equivalentTransforms`, which accepts as its content a sequence of empty elements, named `pair`. Each `pair` element accepts the following required attributes:

- `source`: denotes the fully qualified name of the source transform (that is, C_1).
- `target`: denotes the fully qualified name of the target transform (that is, C_2).
- `arity`: denotes the (common) arity of the source and target transforms.

Note: The source transform may be either a built-in or external function. Both source and target transforms must not be defined as invertible functions.

The following is an example of an `equivalentTransforms` element definition:

```
<equivalentTransforms>
  <pair source="p:sourceFunction_1" target="p:targetFunction_1" arity="1"
  xmlns:p="urn:test1"/>
  <pair source="q:sourceFunction_2" target="q:targetFunction_2" arity="3"
  xmlns:q="urn:test2"/>
</equivalentTransforms>
```

11.3.6 Polymorphic Functions

A library function residing in a relational database function library data service may be designated as polymorphic if its actual return type can be determined from the actual type of one of its parameters. A polymorphic function is annotated by an optional element, `isPolymorphic`, which accepts as its content an empty element, named `parameter`. The `parameter` element accepts the following optional attribute:

- `index`: denotes the index of the parameter whose actual type determines the function's actual return type. The index of the first parameter is assumed to be 1. It may be omitted if the function being annotated has a single parameter.

The following is an example of an `isPolymorphic` element definition:

```
<isPolymorphic>
  <parameter index = "2"/>
</isPolymorphic>
```

11.3.7 Signature Properties

You can use the optional block of signature properties to annotate the parameters of a data service function or procedure with additional information to that provided by the function signature. These properties are applicable to physical data service functions or procedures.

The signature properties block is represented by the element `params` which accepts a sequence of one or more elements, named `param`, as its content. Each `param` element is an empty element that accepts the optional attributes described in [Table 11-14](#):

Table 11-14 *Optional Attributes for the param Element*

Attribute	Description
<code>name</code>	The name of the parameter, as it is known to the external source.
<code>nativeType</code>	The native type of the parameter, as it is known to the external source.
<code>nativeTypeCode</code>	The native type code of the parameter, as it is known to the external source.
<code>xqueryType</code>	The qualified name of the XML Schema or XQuery type used for the parameter.
<code>kind</code>	One of the following values: unknown , in , inout , out , return or result (applicable to stored procedures).

The following is an example of a `params` element definition:

```
<params>
  <param nativeType="java.lang.String"/>
  <param nativeType="java.lang.int"/>
</params>
```

11.3.8 Native Properties

You can use native properties to further annotate a data source function or procedure based on the type of the external source that it represents. There are two types of native properties pertaining to relational and Web service sources respectively:

- [Section 11.3.8.1, "SQL Query Properties"](#)
- [Section 11.3.8.2, "SOAP Handler Properties"](#)

11.3.8.1 SQL Query Properties

The function annotation element of a function that represents a user-defined SQL query has its style attribute set to `sqlQuery` and accepts a nested element, named `sql`. The `sql` element accepts string content that corresponds to the statement of the (possibly parameterized) SQL query that the function represents.

If required, the statement can be escaped inside a CDATA section to account for reserved XML characters (for example: `<`, `>`, `&`). The `sql` element also accepts the optional attribute `isSubquery` whose boolean value indicates whether the SQL statement may be used as a nested SQL sub-query. If the attribute is absent, its value defaults to `true`.

The following is an example of a `sqlQuery` element definition:

```
<sql isSubquery="true">
  SELECT t.FIRST_NAME FROM RTLALL.dbo.CUSTOMER t</sql>
```

11.3.8.2 SOAP Handler Properties

The "function" annotation element of a function or procedure that represents a Web service call accepts a nested element, `interceptorConfiguration`. The `interceptorConfiguration` element accepts two required attributes:

Table 11–15 Required Attributes for the `interceptorConfiguration` Element

Attribute	Description
<code>fileName</code>	The location of the file containing the configuration of the SOAP handler chains that are applicable to the Web service.
<code>aliasName</code>	The alias name by which the SOAP handler chain has been configured.

11.3.9 Implementation Properties

You can use implementation properties to specify that an external create, update or delete procedure is implemented by the update map of the data service in which it is defined.

The optional element `implementation` accepts the required empty element `updateTemplate` as its content.

For example:

```
<implementation>
  <updateTemplate/>
</implementation>
```

11.4 XFL Annotations

There is a single XFL ("XQuery Function Library") annotation per library data service document, which appears before any function annotation in the document. The identifier for the pragma carrying the XFL annotation is `xfl`. The qualified name of the top level element of the XML fragment corresponding to an XFL annotation has the local name:

`xfl`

and the namespace URI:

`urn:annotations.ld.oracle.com`

The contents of the top-level `xfl` element is a sequence of the following blocks of properties.

- [Section 11.4.1, "General Properties"](#)
- [Section 11.4.2, "Data Access Properties"](#)
- [Section 11.4.3, "Security Properties"](#)

The following sections provide detailed descriptions of each block of properties, while the following excerpt provides an example of a XFL annotation, which may serve as a reference.

```
(::pragma xfl <x:xfl xmlns:x="urn:annotations.ld.oracle.com">
<creationDate>2005-03-09T17:48:58</creationDate>
```

```
<webService targetNamespace="urn:GoogleSearch"
  wsdl="ld:google/GoogleSearch.wsdl" />
</x:xf1>:::-)
```

11.4.1 General Properties

The general properties applicable to an library data service document are identical to the general properties for an entity data service document, as described in [Section 11.2.1, "General Properties"](#).

11.4.2 Data Access Properties

Each library data service document defines one or more XQuery functions and/or XQSE functions or procedures that serve as library operations that can be used either inside other entity or library data service documents.

Since library data service documents do not have a target type, the return types of the library functions found inside these document may differ from each other. In particular, a function inside a library data service document may return a value having a simple type (or any other type). Library data service functions can be external data source functions or user-defined.

The following types of library data service documents are supported:

- Relational (physical)
- Web service (physical)
- Java function (physical)
- Relational database function (physical)
- User-defined view (logical)

You can specify only one of the annotations in each library data service. If no annotation is provided, the library data service is considered a user-defined view.

The data access properties for Relational, Web service, Java function, and user-defined view library data service documents are the same as the corresponding properties for entity data service documents, as described above.

A relational database function library data service contains native functions, either database vendor-provided or user-defined in the database, from one or more relational data sources, modeled as external XQuery functions.

The data access annotation for a relational database function library data service comprises an element named `customNativeFunctions` with a single child element, named `relational`, whose content is a sequence of one or more elements named `dataSource`. Each `dataSource` element contains a single text value, which should be set to the JNDI name by which the external relational source has been registered with the application server.

For example:

```
<customNativeFunctions>
<relational>
  <dataSource>oracleDS1</dataSource>
  <dataSource>oracleDS2</dataSource>
</relational>
</customNativeFunctions>
```

11.4.3 Security Properties

The same as in entity data services.

11.5 Data Service Annotations Schema

This section provides the schema for data service annotations.

```
<?xml version="1.0"?>
<xs:schema targetNamespace="urn:annotations.ld.oracle.com"
xmlns:tns="urn:annotations.ld.oracle.com"
xmlns:xs="http://www.w3.org/2001/XMLSchema" elementFormDefault="unqualified"
attributeFormDefault="unqualified">
  <!--=====-->
  <!-- XDS annotation -->
  <!--=====-->
  <xs:element name="xds">
    <xs:complexType>
      <xs:sequence>
        <!-- document properties -->
        <xs:element name="author" type="xs:string" minOccurs="0"/>
        <xs:element name="comment" type="xs:string" minOccurs="0"/>
        <xs:element name="creationDate" type="xs:dateTime" minOccurs="0"/>
        <xs:element name="documentation" type="xs:string" minOccurs="0"/>
        <xs:element name="version" type="xs:decimal" minOccurs="0"/>
        <!-- user defined properties -->
        <xs:sequence minOccurs="0" maxOccurs="unbounded">
          <xs:element name="property">
            <xs:complexType>
              <xs:simpleContent>
                <xs:extension base="xs:string">
                  <xs:attribute name="name" type="xs:string"/>
                </xs:extension>
              </xs:simpleContent>
            </xs:complexType>
          </xs:element>
        </xs:sequence>
        <!-- data access properties -->
        <xs:choice>
          <!-- choice 1: java functions -->
          <xs:element name="javaFunction">
            <xs:complexType>
              <xs:attribute name="class" type="xs:string" use="required"/>
            </xs:complexType>
          </xs:element>
          <!-- choice 2: web services -->
          <xs:element name="webService">
            <xs:complexType>
              <xs:attribute name="wsdl" type="xs:anyURI" use="required"/>
              <xs:attribute name="targetNamespace" type="xs:anyURI" use="required"/>
              <xs:attribute name="sbProxyServiceName" type="xs:string"/>
              <xs:attribute name="sbTransportProtocol" type="tns:SBTransportProtocolType"/>
            </xs:complexType>
          </xs:element>
          <!-- choice 3: relational sources -->
          <xs:element name="relationalDB">
            <xs:complexType>
              <xs:sequence>
                <xs:element name="properties" minOccurs="0">
                  <xs:complexType>
```

```

        <xs:anyAttribute processContents="lax" />
    </xs:complexType>
</xs:element>
</xs:sequence>
<xs:attribute name="name" type="xs:string" use="required"/>
<xs:attribute name="providerId" type="xs:string" />
<xs:attribute name="dbType" type="xs:string"/>
<xs:attribute name="dbVersion" type="xs:string"/>
<xs:attribute name="driver" type="xs:string"/>
<xs:attribute name="uri" type="xs:string"/>
<xs:attribute name="username" type="xs:string"/>
<xs:attribute name="password" type="xs:string"/>
<xs:attribute name="SID" type="xs:string"/>
<xs:attribute name="sourceBindingProviderClassName" type="xs:string"/>
</xs:complexType>
</xs:element>
<!-- choice 4: delimited files -->
<xs:element name="delimitedFile">
    <xs:complexType>
        <xs:attribute name="file" type="xs:anyURI"/>
        <xs:attribute name="schema" type="xs:anyURI"/>
        <xs:attribute name="inferredSchema" type="xs:boolean" default="false"/>
        <xs:attribute name="delimiter" type="xs:string"/>
        <xs:attribute name="fixedLength" type="xs:positiveInteger"/>
        <xs:attribute name="hasHeader" type="xs:boolean" default="false"/>
    </xs:complexType>
</xs:element>
<!-- choice 5: XML files -->
<xs:element name="xmlFile">
    <xs:complexType>
        <xs:attribute name="file" type="xs:anyURI"/>
        <xs:attribute name="schema" type="xs:anyURI" use="required"/>
    </xs:complexType>
</xs:element>
<!-- choice 6: user defined view -->
<xs:element name="userDefinedView" minOccurs="0"/>
<!-- choice 7: nothing, defaults to userDefinedView -->
<xs:sequence/>
</xs:choice>
<!-- field annotations -->
<xs:sequence minOccurs="0" maxOccurs="unbounded">
    <xs:element name="field">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="extension" minOccurs="0">
                    <xs:complexType>
                        <xs:sequence minOccurs="0">
                            <xs:element name="autoNumber">
                                <xs:complexType>
                                    <xs:attribute name="type" type="tns:autoNumberType" use="required"/>
                                    <xs:attribute name="sequenceObjectName" type="xs:string"/>
                                </xs:complexType>
                            </xs:element>
                        </xs:sequence>
                    </xs:complexType>
                </xs:element>
                <xs:attribute name="nativeXpath" type="xs:string"/>
                <xs:attribute name="nativeType" type="xs:string"/>
                <xs:attribute name="nativeTypeCode" type="xs:int"/>
                <xs:attribute name="nativeSize" type="xs:int"/>
                <xs:attribute name="nativeFractionalDigits" type="tns:scaleType"/>
                <xs:attribute name="nativeKey" type="xs:boolean" default="false"/>
            </xs:sequence>
        </xs:complexType>
    </xs:element>
</xs:sequence>

```

```

        <!-- relational: autoNumber -->
        <!-- relational: native column names and types -->
    </xs:complexType>
</xs:element>
<xs:element name="properties">
    <xs:complexType>
        <xs:attribute name="immutable" type="xs:boolean" default="false"/>
        <xs:attribute name="nullable" type="xs:boolean" default="false"/>
        <xs:attribute name="transient" type="xs:boolean" default="false"/>
    </xs:complexType>
</xs:element>
</xs:sequence>
<xs:attribute name="xpath" type="xs:string" use="required"/>
<xs:attribute name="type" type="xs:string" use="required"/>
</xs:complexType>
</xs:element>
</xs:sequence>
<!-- keys -->
<xs:sequence minOccurs="0" maxOccurs="unbounded">
    <xs:element name="key">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="selector" minOccurs="0">
<!-- defaults to . -->
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="extension" minOccurs="0">
                                <xs:complexType>
                                    <xs:attribute name="nativeXPath" type="xs:string" use="required"/>
                                </xs:complexType>
                            </xs:element>
                        </xs:sequence>
                        <xs:attribute name="xpath" type="xs:string" use="required"/>
                    </xs:complexType>
                </xs:element>
            </xs:sequence>
            <xs:attribute name="name" type="xs:string"/>
            <xs:attribute name="type" type="xs:QName"/>
            <xs:attribute name="inferred" type="xs:boolean" default="true"/>
            <xs:attribute name="inferredSchema" type="xs:boolean" default="true"/>
        </xs:complexType>
    </xs:element>
</xs:sequence>
<!-- relationships -->
<xs:sequence minOccurs="0" maxOccurs="unbounded">
    <xs:element name="relationshipTarget">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="relationship" minOccurs="0">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="description" type="xs:string" minOccurs="0"/>
                        </xs:sequence>
                    </xs:complexType>
                </xs:element>
            </xs:sequence>
            <xs:attribute name="roleName" type="xs:string" use="required"/>
            <xs:attribute name="roleNumber" type="tns:roleType" default="1"/>
            <xs:attribute name="XDS" type="xs:string" use="required"/>
            <xs:attribute name="minOccurs" type="tns:allNNI" default="1"/>

```

```

        <xs:attribute name="maxOccurs" type="tns:allNNI" default="1"/>
        <xs:attribute name="opposite" type="xs:string"/>
    </xs:complexType>
</xs:element>
</xs:sequence>
<!-- SDO elements -->
<xs:element name="functionForDecomposition" minOccurs="0">
    <xs:complexType>
        <xs:attribute name="name" type="xs:QName" use="required"/>
        <xs:attribute name="arity" type="xs:int" use="required"/>
    </xs:complexType>
</xs:element>
<xs:element name="javaUpdateExit" minOccurs="0">
    <xs:complexType>
        <xs:attribute name="className" type="xs:string" use="required"/>
        <xs:attribute name="classFile" type="xs:string"/>
    </xs:complexType>
</xs:element>
<xs:element name="optimisticLockingFields" minOccurs="0">
    <xs:complexType>
        <xs:choice>
            <xs:element name="updated">
                <xs:complexType/>
            </xs:element>
            <xs:element name="projected">
                <xs:complexType/>
            </xs:element>
            <xs:element name="field" minOccurs="0" maxOccurs="unbounded">
                <xs:complexType>
                    <xs:attribute name="name" type="xs:string" use="required"/>
                </xs:complexType>
            </xs:element>
        </xs:choice>
    </xs:complexType>
</xs:element>
<!-- security -->
<xs:element name="secureResources" minOccurs="0">
    <xs:complexType>
        <xs:sequence>
            <xs:element name="secureResource" type="xs:NCName" minOccurs="0"
maxOccurs="unbounded"/>
        </xs:sequence>
    </xs:complexType>
</xs:element>
<xs:element name="readOnly" minOccurs="0">
    <xs:complexType/>
</xs:element>
</xs:sequence>
<xs:attribute name="targetType" type="xs:QName" use="required"/>
</xs:complexType>
</xs:element>
<!--=====-->
<!-- XFL annotation -->
<!--=====-->
<xs:element name="xfl">
    <xs:complexType>
        <xs:sequence>
            <!-- document properties -->
            <xs:element name="author" type="xs:string" minOccurs="0"/>
            <xs:element name="comment" type="xs:string" minOccurs="0"/>

```

```

<xs:element name="creationDate" type="xs:dateTime" minOccurs="0"/>
<xs:element name="documentation" type="xs:string" minOccurs="0"/>
<xs:element name="version" type="xs:decimal" minOccurs="0"/>
<!-- user defined properties -->
<xs:sequence minOccurs="0" maxOccurs="unbounded">
  <xs:element name="property">
    <xs:complexType>
      <xs:simpleContent>
        <xs:extension base="xs:string">
          <xs:attribute name="name" type="xs:string"/>
        </xs:extension>
      </xs:simpleContent>
    </xs:complexType>
  </xs:element>
</xs:sequence>
<!-- data access properties -->
<xs:choice>
  <!-- choice 1: java functions -->
  <xs:element name="javaFunction">
    <xs:complexType>
      <xs:attribute name="class" type="xs:string" use="required"/>
    </xs:complexType>
  </xs:element>
  <!-- choice 2: web services -->
  <xs:element name="webService">
    <xs:complexType>
      <xs:attribute name="wsdl" type="xs:anyURI" use="required"/>
      <xs:attribute name="targetNamespace" type="xs:anyURI" use="required"/>
      <xs:attribute name="sbProxyServiceName" type="xs:string"/>
      <xs:attribute name="sbTransportProtocol" type="tns:SBTransportProtocol"/>
    </xs:complexType>
  </xs:element>
  <!-- choice 3: relational sources -->
  <xs:element name="relationalDB">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="properties" minOccurs="0">
          <xs:complexType>
            <xs:anyAttribute processContents="lax" />
          </xs:complexType>
        </xs:element>
      </xs:sequence>
      <xs:attribute name="name" type="xs:string" use="required"/>
      <xs:attribute name="providerId" type="xs:string" />
      <xs:attribute name="dbType" type="xs:string"/>
      <xs:attribute name="dbVersion" type="xs:string"/>
      <xs:attribute name="driver" type="xs:string"/>
      <xs:attribute name="uri" type="xs:string"/>
      <xs:attribute name="username" type="xs:string"/>
      <xs:attribute name="password" type="xs:string"/>
      <xs:attribute name="SID" type="xs:string"/>
      <xs:attribute name="sourceBindingProviderClassName" type="xs:string"/>
    </xs:complexType>
  </xs:element>
  <!-- choice 6: user defined view -->
  <xs:element name="userDefinedView" minOccurs="0"/>
  <!-- choice 7: nothing, defaults to userDefinedView -->
  <xs:sequence/>
  <!-- choice 8: custom native functions -->
  <xs:element name="customNativeFunctions">

```

```

    <xs:complexType>
      <xs:choice>
        <xs:element name="relational">
          <xs:complexType>
            <xs:sequence>
              <xs:element name="dataSource" type="xs:string" maxOccurs="unbounded" />
            </xs:sequence>
          </xs:complexType>
        </xs:element>
      </xs:choice>
    </xs:complexType>
  </xs:element>
</xs:choice>
<!-- field annotations -->
<xs:sequence minOccurs="0" maxOccurs="unbounded">
  <xs:element name="field">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="extension" minOccurs="0">
          <xs:complexType>
            <xs:sequence minOccurs="0">
              <xs:element name="autoNumber">
                <xs:complexType>
                  <xs:attribute name="type" type="tns:autoNumberType" use="required" />
                  <xs:attribute name="sequenceObjectName" type="xs:string" />
                </xs:complexType>
              </xs:element>
            </xs:sequence>
            <xs:attribute name="nativeXPath" type="xs:string" />
            <xs:attribute name="nativeType" type="xs:string" />
            <xs:attribute name="nativeTypeCode" type="xs:int" />
            <xs:attribute name="nativeSize" type="xs:int" />
            <xs:attribute name="nativeFractionalDigits" type="tns:scaleType" />
            <!-- relational: autoNumber -->
            <!-- relational: native column names and types -->
          </xs:complexType>
        </xs:element>
        <xs:element name="properties">
          <xs:complexType>
            <xs:attribute name="immutable" type="xs:boolean" default="false" />
            <xs:attribute name="nullable" type="xs:boolean" default="false" />
            <xs:attribute name="transient" type="xs:boolean" default="false" />
          </xs:complexType>
        </xs:element>
      </xs:sequence>
      <xs:attribute name="xpath" type="xs:string" use="required" />
      <xs:attribute name="type" type="xs:string" use="required" />
    </xs:complexType>
  </xs:element>
</xs:sequence>

<xs:element name="secureResources" minOccurs="0">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="secureResource" type="xs:NCName" minOccurs="0"
maxOccurs="unbounded" />
    </xs:sequence>
  </xs:complexType>
</xs:element>
</xs:sequence>

```

```

    </xs:complexType>
  </xs:element>
  <!--=====-->
  <!-- function annotation -->
  <!--=====-->
  <xs:element name="function">
    <xs:complexType>
      <xs:sequence>

        <!-- standard properties -->
        <xs:element name="author" type="xs:string" minOccurs="0"/>
        <xs:element name="comment" type="xs:string" minOccurs="0"/>
        <xs:element name="version" type="xs:decimal" minOccurs="0"/>
        <xs:element name="documentation" type="xs:string" minOccurs="0"/>

        <!-- user defined properties -->
        <xs:sequence minOccurs="0" maxOccurs="unbounded">
          <xs:element name="property">
            <xs:complexType>
              <xs:simpleContent>
                <xs:extension base="xs:string">
                  <xs:attribute name="name" type="xs:string"/>
                </xs:extension>
              </xs:simpleContent>
            </xs:complexType>
          </xs:element>
        </xs:sequence>

        <!-- UI properties -->
        <xs:element name="uiProperties" minOccurs="0">
          <xs:complexType>
            <xs:sequence maxOccurs="unbounded">
              <xs:element name="component">
                <xs:complexType>
                  <xs:sequence>
                    <xs:element name="treeInfo" minOccurs="0" maxOccurs="unbounded">
                      <xs:complexType>
                        <xs:sequence>
                          <xs:element name="collapsedNodes" minOccurs="0">
                            <xs:complexType>
                              <xs:sequence>
                                <xs:element name="collapsedNode" type="xs:string" minOccurs="0"
maxOccurs="unbounded"/>
                              </xs:sequence>
                            </xs:complexType>
                          </xs:element>
                        </xs:sequence>
                      </xs:complexType>
                    </xs:element>
                  </xs:sequence>
                </xs:complexType>
              </xs:element>
            </xs:sequence>
            <xs:attribute name="id" type="xs:string"/>
          </xs:complexType>
        </xs:element>
      </xs:sequence>
      <xs:attribute name="identifier" type="xs:string"/>
      <xs:attribute name="minimized" type="xs:boolean" default="false"/>
      <xs:attribute name="x" type="xs:int"/>
      <xs:attribute name="y" type="xs:int"/>
      <xs:attribute name="w" type="xs:int"/>
      <xs:attribute name="h" type="xs:int"/>
      <xs:attribute name="viewPosX" type="xs:int"/>
      <xs:attribute name="viewPosY" type="xs:int"/>
    </xs:complexType>
  </xs:element>

```

```
        </xs:element>
    </xs:sequence>
</xs:complexType>
</xs:element>

<!-- sql statement -->
<xs:element name="sql" minOccurs="0">
    <xs:complexType>
        <xs:simpleContent>
            <xs:extension base="xs:string">
                <xs:attribute name="isSubquery" type="xs:boolean" default="true"/>
            </xs:extension>
        </xs:simpleContent>
    </xs:complexType>
</xs:element>

<!-- cache -->
<xs:element name="nonCacheable" minOccurs="0">
    <xs:complexType/>
</xs:element>

<!-- transactions -->
<xs:element name="nonTransactional" minOccurs="0">
    <xs:complexType/>
</xs:element>

<!-- optimization -->
<xs:element name="outputIsOrderedBy" minOccurs="0">
    <xs:complexType>
        <xs:sequence>
            <!-- absent for parameters whose order in the function signature
                 coincides with their order in the order by list -->
            <xs:element name="parameter" minOccurs="0" maxOccurs="unbounded">
                <xs:complexType>
                    <!-- 1, 2, ... -->
                    <xs:attribute name="index" type="xs:int" use="required"/>
                    <!-- overrides default -->
                    <xs:attribute name="mode" type="tns:orderingModeType"/>
                </xs:complexType>
            </xs:element>
        </xs:sequence>
        <xs:attribute name="mode" type="tns:orderingModeType" use="required"/>
    </xs:complexType>
</xs:element>

<xs:element name="inverseFunctions" minOccurs="0">
    <xs:complexType>
        <xs:sequence>
            <xs:element name="inverseFunction" minOccurs="1" maxOccurs="unbounded">
                <xs:complexType>
                    <!-- 1, 2, ... -->
                    <xs:attribute name="parameterIndex" type="xs:int"/>
                    <xs:attribute name="name" type="xs:QName" use="required"/>
                </xs:complexType>
            </xs:element>
        </xs:sequence>
    </xs:complexType>
</xs:element>

<xs:element name="equivalentTransforms" minOccurs="0">
    <xs:complexType>
```

```

    <xs:sequence>
      <xs:element name="pair" minOccurs="1" maxOccurs="unbounded">
        <xs:complexType>
          <xs:attribute name="source" type="xs:QName" use="required"/>
          <xs:attribute name="target" type="xs:QName" use="required"/>
          <xs:attribute name="arity" type="xs:int" use="required"/>
        </xs:complexType>
      </xs:element>
    </xs:sequence>
  </xs:complexType>
</xs:element>

<!-- polymorphism -->
<xs:element name="isPolymorphic" minOccurs="0">
  <xs:complexType>
    <xs:choice>
      <xs:element name="parameter">
        <xs:complexType>
          <xs:sequence/>
          <!-- optional: defaults to 1 -->
          <xs:attribute name="index" type="xs:nonNegativeInteger"/>
        </xs:complexType>
      </xs:element>
    </xs:choice>
  </xs:complexType>
</xs:element>

<!-- signature: used by java functions and stored procedures -->
<xs:element name="params" minOccurs="0">
  <xs:complexType>
    <xs:sequence maxOccurs="unbounded">
      <xs:element name="param">
        <xs:complexType>
          <xs:attribute name="name" type="xs:string"/>
          <xs:attribute name="nativeType" type="xs:string"/>
          <xs:attribute name="nativeTypeCode" type="xs:int"/>
          <xs:attribute name="xqueryType" type="xs:QName"/>
          <xs:attribute name="kind" type="tns:paramKindType"/>
        </xs:complexType>
      </xs:element>
    </xs:sequence>
  </xs:complexType>
</xs:element>

<!-- interceptor configuration: used by webservice SOAP interceptors -->
<xs:element name="interceptorConfiguration" minOccurs="0">
  <xs:complexType>
    <xs:attribute name="aliasName" type="xs:string" use="required"/>
    <xs:attribute name="fileName" type="xs:string" use="required"/>
  </xs:complexType>
</xs:element>

<!-- implementation -->
<xs:element name="implementation" minOccurs="0">
  <xs:complexType>
    <xs:choice>
      <xs:element name="updateTemplate">
        <xs:complexType/>
      </xs:element>
    </xs:choice>
  </xs:complexType>

```

```

    </xs:element>
</xs:sequence>

<xs:attribute name="visibility" type="tns:functionVisibilityType" default="protected"/>
<xs:attribute name="kind" type="tns:functionKindType" default="library"/>
<xs:attribute name="isPrimary" type="xs:boolean" default="false"/>
<xs:attribute name="roleName" type="xs:string"/>
<xs:attribute name="nativeName" type="xs:string"/>
<xs:attribute name="nativeLevel1Container" type="xs:string"/>
<xs:attribute name="nativeLevel2Container" type="xs:string"/>
<xs:attribute name="nativeLevel3Container" type="xs:string"/>
<xs:attribute name="style" type="tns:functionStyleType"/>
</xs:complexType>
</xs:element>
<!--=====-->
<!-- common types -->
<!--=====-->
<xs:simpleType name="functionVisibilityType">
  <xs:restriction base="xs:string">
    <xs:enumeration value="public"/>
    <xs:enumeration value="protected"/>
    <xs:enumeration value="private"/>
  </xs:restriction>
</xs:simpleType>
<xs:simpleType name="functionKindType">
  <xs:restriction base="xs:string">
    <xs:enumeration value="read"/>
    <xs:enumeration value="navigate"/>
    <xs:enumeration value="create"/>
    <xs:enumeration value="update"/>
    <xs:enumeration value="delete"/>
    <xs:enumeration value="library"/>
  </xs:restriction>
</xs:simpleType>
<xs:simpleType name="functionStyleType">
  <xs:restriction base="xs:string">
    <xs:enumeration value="table"/>
    <xs:enumeration value="view"/>
    <xs:enumeration value="storedProcedure"/>
    <xs:enumeration value="sqlQuery"/>
    <xs:enumeration value="document"/>
    <xs:enumeration value="rpc"/>
  </xs:restriction>
</xs:simpleType>
<!-- used by stored procedures -->
<xs:simpleType name="paramKindType">
  <xs:restriction base="xs:string">
    <xs:enumeration value="unknown"/>
    <xs:enumeration value="in"/>
    <xs:enumeration value="inout"/>
    <xs:enumeration value="out"/>
    <xs:enumeration value="return"/>
    <xs:enumeration value="result"/>
  </xs:restriction>
</xs:simpleType>
<!-- used by maxOccurs in relationship -->
<xs:simpleType name="allNNI">
  <xs:union memberTypes="xs:nonNegativeInteger">
    <xs:simpleType>
      <xs:restriction base="xs:string">

```

```

        <xs:enumeration value="unbounded"/>
        <xs:enumeration value=""/>
    </xs:restriction>
</xs:simpleType>
</xs:union>
</xs:simpleType>
<!-- used by relationships -->
<xs:simpleType name="roleType">
    <xs:restriction base="xs:nonNegativeInteger">
        <xs:enumeration value="1"/>
        <xs:enumeration value="2"/>
    </xs:restriction>
</xs:simpleType>
<xs:simpleType name="autoNumberType">
    <xs:restriction base="xs:string">
        <xs:enumeration value="identity"/>
        <xs:enumeration value="sequence"/>
        <xs:enumeration value="userComputed"/>
    </xs:restriction>
</xs:simpleType>
<xs:simpleType name="nullSortOrderType">
    <xs:restriction base="xs:string">
        <xs:enumeration value="high"/>
        <xs:enumeration value="low"/>
        <xs:enumeration value="unknown"/>
    </xs:restriction>
</xs:simpleType>
<xs:simpleType name="scaleType">
    <xs:union memberTypes="xs:int">
        <xs:simpleType>
            <xs:restriction base="xs:string">
                <xs:enumeration value="null"/>
            </xs:restriction>
        </xs:simpleType>
    </xs:union>
</xs:simpleType>
<xs:simpleType name="orderingModeType">
    <xs:restriction base="xs:string">
        <xs:enumeration value="ascending"/>
        <xs:enumeration value="descending"/>
    </xs:restriction>
</xs:simpleType>
<xs:simpleType name="stringListType">
    <xs:list itemType="xs:string"/>
</xs:simpleType>
<xs:simpleType name="dataSourcesType">
    <xs:restriction base="tns:stringListType">
        <xs:minLength value="1"/>
    </xs:restriction>
</xs:simpleType>
<xs:simpleType name="SBTransportProtocolType">
    <xs:restriction base="xs:string">
        <xs:enumeration value="t3"/>
        <xs:enumeration value="iiop"/>
        <xs:enumeration value="http"/>
        <xs:enumeration value="t3s"/>
        <xs:enumeration value="iiops"/>
        <xs:enumeration value="https"/>
    </xs:restriction>
</xs:simpleType>

```

</xs:schema>

Best Practices When Building Data Services

This topic introduces a series of best practices you can consider when building data services using Oracle Data Service Integrator. This chapter contains the following sections:

- [Section 12.1, "Overview"](#)
- [Section 12.2, "Oracle Data Service Integrator Development Best Practices"](#)
- [Section 12.3, "Performance and Optimization Best Practices"](#)
- [Section 12.4, "How To Get More Help"](#)
- [Section 12.5, "Related Topics"](#)

12.1 Overview

The Oracle Data Service Integrator runtime is hosted inside a WebLogic server container and can co-exist with other server platforms such as AquaLogic Service Bus 10gR3 (Oracle Data Service Integrator 10gR3 and higher), WebLogic Integration (WLI) or WebLogic Portal. Clients access the Oracle Data Service Integrator data services through the Data Service Mediator API, Web Services API, the Oracle Data Service Integrator JDBC driver, and the Service Bus Oracle Data Service Integrator Transport. Oracle Workshop for WebLogic-based applications can access Oracle Data Service Integrator data services through an Oracle Data Service Integrator Control.

Oracle Data Service Integrator was one of the first products to introduce the concept of dataspace. A dataspace is a unit of deployment, administration, and security policy control. A single Oracle Data Service Integrator runtime environment can host one or more dataspace. A dataspace contains a set of related data services. Data services hosted in a dataspace run in the same context and can be reused from other data services.

12.1.1 Understanding the Oracle Data Service Integrator Server

When the Oracle Data Service Integrator server receives a request, it does the following:

1. The Oracle Data Service Integrator server looks up the data service by namespace, and the operation by function name and number of arguments.
2. The server looks in the query plan cache for a compiled query plan. If the server does not find one, it compiles the query and caches the resulting query plan. (The query plan cache uses a "most-recently-used" algorithm, with a default size of 100 plans.)
3. After the server has a query plan, it binds the operation's arguments to the plan.

4. The server runs the query and streams the data from the back-end systems through the Oracle Data Service Integrator server.

WARNING: There are cases, however, that require full materialization of the data inside the Oracle Data Service Integrator server, such as during an in-memory sort operation run in the Oracle Data Service Integrator server middle tier. For these types of operations, the server can optionally use secondary storage (disk) as needed so as not to exhaust the virtual memory of the JVM. Here are some specifics:

- If the operation is initiated using the data service mediator API or the Oracle Data Service Integrator JDBC driver, the query result can be streamed all the way to the client as the client iterates through the results.
 - If the operation is initiated using the Web Service client API or DSP transport, the results are streamed through the server but fully materialized inside the server before sending the results to the client.
-

12.1.2 Understanding the Oracle Data Service Integrator Client

For detailed information about Oracle Data Service Integrator clients, refer to *Introducing Data Services for Client Applications in the Client Application Developer's Guide*.

12.2 Oracle Data Service Integrator Development Best Practices

Oracle Data Service Integrator, similar to other advanced development systems, offers a powerful and flexible environment for creating high performance and feature-rich programs. This section provides a series of best practice guidelines for organizing and structuring your projects.

Best Practice	Overview
Organize using dataspace projects	You can reduce build and deployment times by organizing your data services into projects.
Build data services in layers	Building projects using a layered approach can simplify development and increase reliability

12.2.1 Organizing Data Services Using Projects

Oracle Data Service Integrator treats all artifacts within a dataspace project as being within the same context. This means that you are free to organize projects in any way that makes logical sense. For example, you could choose to create a single project with 100 data services, or 10 projects with 10 data services each.

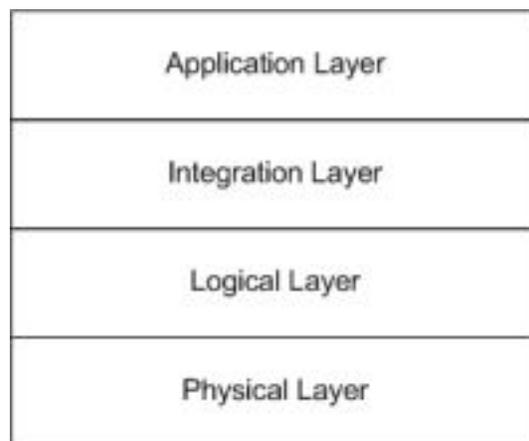
In practical terms, however, there are advantages to organizing in specific ways. For example, Studio builds a project at a time. Therefore Studio can typically build a project containing fewer data services faster than one containing a larger number of data services.

Tip: In general, you should consider organizing data services into projects if the data services (across the projects) do not need to participate in building views. This means, for example, grouping logically-related data services in one project and non-related services in another project, and so on.

12.2.2 Building Data Services in Layers

Oracle Data Service Integrator enables you to design and build your data services as a series of layers extending from the physical data sources to your business and application logic. This section provides an overview of the most common layers used in Oracle Data Service Integrator and offers suggestions on the types of operations that you should include in each layer.

Figure 12–1 Oracle Data Service Integrator Layers



The following table provides an overview of the layers in a typical Oracle Data Service Integrator dataspace project:

Table 12–1 Oracle Data Service Integrator Layers

Layer	Overview
Physical	Contains imported operations related to resources such as databases, Web services, XML files, and Java functions
Logical	Contains operations matching the physical layer operations that perform simple transformations to the data.
Integration	Contains the coarse-grained business objects that are manipulated by the dataspace project
Application	Can add a further layer of abstraction and specialization to a data service project.

12.2.2.1 Physical Layer

The physical layer contains imported operations related to resources such as databases, Web services, XML files, and Java functions.

To create the physical layer, do the following:

1. At the top-level of the dataspace project, create a folder named Physical.

2. Create a subdirectory for each physical resource in the Physical folder.
3. Import the operations for each resource into the respective folders.

Note: You typically should not modify these physical data services. This makes it easier to re-import changes to the physical resources, if necessary.

12.2.2.2 Logical Layer

The logical layer contains operations matching the physical layer operations that perform simple transformations to the data.

To create the logical layer, do the following:

1. At the top-level of the dataspace project, create a folder named Logical.
2. Mirror the directory structure from the Physical folder in the Logical folder. Create a corresponding data service in the Logical folder for each data service in the Physical folder.
3. Create a schema for each data service.
4. For each data service in the Logical folder, create a new operation that calls the corresponding physical operation.

Assign meaningful names to elements. For example, **CUST_ID** might become `CustomerId`. You can also add computation and other logic to operations in the logical layer, for example, **LoanDuration** = `$LOAN/END_DATE - $LOAN/START_DATE`.

Note: You can also use the logical layer to define inverse functions, as required. Ideally, this is the only layer that should have expanded mappings, although the application layer can also use expanded element mappings to retrieve required elements.

The logical layer typically does not include significant joins or selections. The layer instead focuses on data access and element construction. Therefore the resulting query plans will simply be calls to the physical source along with element construction.

This means that there is typically little opportunity to improve query efficiency in the logical layer, so it normally does not help to examine the query plans at this layer.

The following shows a sample operation in the logical layer:

```
function getCustomer() {  
  
  for $c in CUSTOMER()  
  return  
    <Customer>  
      <CustomerId>{$c/CUSTOMER_ID}</CustomerId>  
      <LastName>{$c/LAST_NAME}</ LastName >  
      <FirstName>{$c/FIRST_NAME}</FirstName>  
      <Email>{$c/EMAIL}</Email>  
      <CustomerSince>{$c/CUSTOMER_SINCE}</CustomerSince>  
      <SupportLevel>{$c/SUPPORT_LEVEL}</SupportLevel>  
    </Customer>  
}
```

12.2.2.3 Integration Layer

The integration layer contains the coarse-grained business objects that are manipulated by the dataspace project.

To create the integration layer, do the following:

1. At the top-level of the dataspace project, create a folder named Integration.
2. Create a folder in the Integration directory for each group of coarse-grained business objects you plan to create.

You might, for example, have folders named Ordering, Customer, and Product.

Note: The integration layer is where you should perform data integration, joining, for instance Customer, Order, and LineItem to create the coarse-grained XML structure CustomerOrderLineItem.

3. Create a data service for each coarse-grained business object.

If possible, also create a 'get all' base function in the data service (such as CustomerOrderLineItem) that constructs an element by joining elements such as Customer >Order > LineItem.

Note: Do not put selection operations, such as 'by CustomerId' for example, in this 'get all' function. Also, operations in the integration layer should only call operations in the logical layer.

If possible, design data services in the integration layer to only use compound mappings, since expanded mappings are expensive in terms of time and memory. The example below shows how to construct nested output using compound mappings only. Note the introduction of envelope elements. The **LineItemEnvelope** element is not required since it is at the bottom; it is included for consistency.

```
function getCustomers() {
  for $c in Customer()
  return
    <CustomerEnvelope>
      {$c}
      {
        for $o in Order()
          where $o/CustomerId eq $c/CustomerId
          return
            <OrderEnvelope>
              {$o}
              {
                for $l in LineItem()
                  where $l/OrderId eq $o/OrderId
                  return
                    <LineItemEnvelope>
                      {$l}
                    </LineItemEnvelope>
              }
            </OrderEnvelope>
          }
        </CustomerEnvelope>
      }
}
```

After you have created the base function, you can write specific selection functions using the base function. For example:

```
function getCustomerById( $customerId as xs:string) {
  for $c in getCustomers()
    where $c/CustomerId eq $customerId
    return
      $c
}

function getCustomersByLastName( $lastName as xs:string) {
  for $c in getCustomers()
    where $c/LastName eq $lastName
    return
      $c
}

function getCustomersByOrderAmountExceeding( $orderAmount as xs:decimal) {
  for $c in getCustomers()
    where some $o in $c/OrderEnvelope/Order
      satisfies $o/OrderAmount > $orderAmount
    return
      $c
}
```

Since the integration layer contains joins and possibly selections, there is the possibility of inefficient query plans. As you create each function, examine the resulting query plan and evaluate its efficiency.

In some cases, it may even be useful to examine the query plan as you develop each function. For instance, when creating the base operation, `getCustomers()`, consider examining the query plan after you add the join on `Orders`, but before you add the join on `LineItems`.

If adding the join `Customer > Order` produces an inefficient query plan, there is little point in continuing to complicate the query by also joining to `LineItem` before correcting the join on `Customer > Order`.

12.2.2.4 Application Layer

For simple dataspace projects, three layers will likely be sufficient. However, if your dataspace has more demanding requirements, you can add an application layer.

The application layer adds a further layer of abstraction and specialization to a data service project. For instance, you could add row or column-level security in the application layer, or only expose required columns to reduce unnecessary data retrieval. Similarly, you could use the application layer to create a highly-specialized data service to fulfill a specific need.

The structure of the application layer resembles the integration layer, and makes calls directly to the integration layer.

The benefit of these layers is that once they are created, you will have constructed a virtual database around all your disparate data. You will be able to create and modify data services that are only loosely coupled with the underlying physical data. Changes in the underlying data will be much less likely to affect your data service developers both in term of data service development and maintenance.

And your application developers will be able to write to layer that is highly insulated and far more powerful as compared to writing to separate data sources and processing the results locally (perhaps in yet another database).

12.3 Performance and Optimization Best Practices

Oracle Data Service Integrator attempts to optimize the operation of your dataspace project as well as interactions between your dataspace and the underlying data sources. This section describes a series of best practices that you can use to potentially increase the levels of optimization and boost the performance of your dataspace projects.

Category	Best Practice
Database Access	<ul style="list-style-type: none"> ■ Retrieve only necessary data ■ Design functions that can be pushed to the database ■ Verify that joins are implemented as left-outer or inner/natural joins ■ Push left-outer joins to the database ■ Use a ppci-impl join ■ Avoid casting ■ Optimize ad hoc queries ■ Write your own SQL in some cases
Fail-Over	Exercise care when using fail-over, fail-over-retry, and timeout
Inverse Functions	Use inverse functions
Caching/Auditing	Take advantage of caching
Query Plans	<ul style="list-style-type: none"> ■ Evaluate performance before running the query ■ Precompile query plans ■ Evaluate performance by running the query
Performance Monitoring	Evaluate and monitor operational performance

12.3.1 Database Access

This section describes performance and optimization best practices related to data access.

12.3.1.1 Retrieving Only the Necessary Data

Examine the query plan and verify that the SQL statements retrieve only the necessary data. For example, suppose you have the following selection in your XQuery:

```
where $c/POSTAL_CODE eq $postalCode
```

You should then find that condition in the SQL statements as:

```
WHERE POSTAL_CODE = ?
```

If this clause does not appear in the SQL, you need to determine why this is the case to enable the code generated by Oracle Data Service Integrator to be pushed to the database.

12.3.1.2 Designing Functions Which can be Pushed to the Database

To have a function pushed to the database, and thereby increase efficiency and performance, the function must have an equivalent database operation or an inverse operation. For example, suppose you have the following selection in your XQuery:

```
where match($c/LAST_NAME, $regexpr)
```

Oracle Data Service Integrator cannot push this selection to the database because there is no equivalent to the `match()` function in SQL, and there is no inverse function for `match()` defined in Oracle Data Service Integrator.

Note: In most cases, your XQuery functions will test for simple equality and inequality. Oracle Data Service Integrator can automatically push these types of operators to the database.

12.3.1.3 Verifying that Joins are Implemented as Left-Outer or Inner/Natural Joins

You need to consider how joins are implemented in your XQuery. For example, suppose you have a join in your XQuery similar to the following:

```
where $c/CUSTOMER_ID eq $o/CUSTOMER_ID
```

If both tables are in the same database and the returned XML is nested, the query plan should contain SQL with a left-outer-join of these tables. If the returned XML is flat, the query plan should include SQL with a inner/natural join of the two tables.

If the join is not implemented as either a left-outer or inner/natural join in the SQL, you need to determine why. Similarly, you need to examine options for ensuring that the joins are pushed to the database.

12.3.1.4 Pushing Left Outer Joins to the Database

The following conditions need to be satisfied for a join producing nested output to be pushed to the database as a left-outer join:

- Both tables need to be in the same database.
- The driving (left) table must have a unique key.

The second condition is necessary because the result set from SQL will contain duplicate rows, and Oracle Data Service Integrator needs to detect these rows so that the right-hand side can be grouped by the left-hand side key. If the left table does not have a unique key, Oracle Data Service Integrator is forced to get the records from the left table only, generate a key, retrieve corresponding rows from the right table, and join them with those from the left. If this occurs, you will find two SQL statements, one for each table, joined using an `index-cpp` join and a `generate-key()` operator above the left table.

If the conditions are met, there is no limit to the number of levels of nested output that can be joined in a single SQL statement. However, if within the nesting, there are sibling child elements, the second child cannot be in the join. For example, `Customers > Orders > LineItems > Products > Pricings` can be pushed as a single SQL statement with multiple left-outer-joins. But in `Customers > (Orders, Addresses)`, `Customers >`

Orders can be pushed as a left-outer join, but Addresses will be implemented as a clustered, parameter-passing join (cpp).

The reason that Address is not joined as another left-outer join in the same SQL is because adding it would produce a cross-product between Orders and Addresses as they are not dependent on each other. So, if a Customer had 1,000 orders and 1,000 addresses, the resulting SQL would return 1,000,000 rows and Oracle Data Service Integrator would be left to remove the duplicates.

By splitting this into two separate SQL statements, the first would retrieve 1,000 rows and the second would retrieve another 1,000 rows (total 2,000). In addition, the reason that Address is not retrieved in the same SQL statement through a union is that this would disrupt the streamability of the query, and Oracle Data Service Integrator would need to read all of the Customer > Order rows before it got to the first Customer > Address row.

Although this would not be a problem for small results, Oracle Data Service Integrator is unaware of the result set size and, whenever possible, prepares a query plan that can handle results of any size.

12.3.1.5 Using a ppci-impl Join

When Oracle Data Service Integrator cannot push the join of two tables to the database, the system implements the join as a parameter-passing, clustered-indexed join. Specifically, consider a JDBC datasource project that needs to retrieve the list of customers and their associated orders. You could use code similar to the following:

```
// will need stmt2 later
stmt2=conn.prepareStatement("select * from ORDER where CUSTOMER_ID = ?");

stmt1= conn.prepareStatement("select * from CUSTOMER where ZIPCODE=?");
stmt1.setString(1, zipcode);
rs1 = stmt1.executeQuery();
while (rs1.next() ){
    System.out.println("customer :"+rs1.getString("CUSTOMER_ID"));
    stmt2.setString(1, rs1.getString("CUSTOMER_ID")); // set parameter
    rs2=stmt2.executeQuery();
    while( rs2.next() ) {
        System.out.println("    order :"+rs2.getString("ORDER_ID"));
    }
}
}=
```

This uses a parameter-passing join (the parameter is CUSTOMER_ID). If, however, you have a lengthy list of customers from the first SQL statement, you will need to call the second SQL statement quite a few times. One thing you could do to reduce the number of calls would be to fetch orders from several customers at a time. Once you have the orders for several customers, you can manually join the rows with the customers.

This is exactly what Oracle Data Service Integrator does automatically for you! So, if you have a join that can not be pushed as a single statement, Oracle Data Service Integrator performs a batched (clustered) parameter-passing join. The first SQL statement would therefore be as follows:

```
select ... from CUSTOMER ...
```

Similarly, the second statement would be as follows:

```
select ... from ORDER where CUSTOMER_ID = ? or CUSTOMER_ID = ? ....
```

You may wonder why there are 20 arguments. Why not use the exact number of customers, for example? The answer is that there are limits to the number of arguments that you can pass to an SQL statement.

But, more importantly, database performance quickly degrades as more arguments are passed (at some point, a full-table scan would be a good idea if there are many lookups to perform). That's the first reason why there are 20 arguments.

Second, if you were to specify the exact number of arguments, you would need to create a separate SQL statement (and `PreparedStatement`) in the worst case, for up to 20 arguments resulting in 20 times the number of `PreparedStatements` that you really need.

In addition, since the database needs to compile `PreparedStatements` not encountered earlier, this would require 20 times the number of compilations. In some cases, this could even cause statements to be discarded from the cache, resulting in even more compiling.

Note: When a JDBC connection pool maintains a cached `PreparedStatement`, it uses resources (cursors) in the database. This is a further motivation to keep the number of `PreparedStatements` to a minimum.

The final point to consider is what happens when the number of customers is not an exact multiple of 20? In this case, the final argument is repeated to fill in the remaining places since the logical operation `CUSTOMER='ZZZ' OR CUSTOMER='ZZZ'` is the same as `CUSTOMER='ZZZ'`. This is a very simple optimization for a database and testing has shown that it does not significantly impact performance.

Note that `ppci-join` seems like a lot of effort simply to eliminate a few database round trips. If the join is on an indexed column, there will be 20 indexed lookups independent of whether you make one call with 20 parameters, or make 20 calls with one parameter.

While this is the case, if the join is on an unindexed column, one call with 20 parameters takes one full table scan and 20 calls with one parameter takes 20 full table scans. Obviously you may need a noticeable performance difference between these two approaches.

12.3.1.6 Avoiding Casting

You should generally avoid casting data, especially casting that cannot be performed by the database.

The reason for this is as follows: If there is a cast from one type to another for the purpose of a join or a selection, and that cast cannot be handled by the database, Oracle Data Service Integrator is forced to retrieve all the rows and perform the cast itself, followed by the join or selection.

You can determine if the casting cannot be performed by the database by examining the cast operators in the query plan.

12.3.1.7 Optimizing Ad Hoc Queries

Clients can submit ad hoc queries, which are essentially data service functions that exist only while the queries are run. From the engine's perspective, ad hoc queries go through the same life-cycle as regular queries.

Oracle Data Service Integrator caches the query plan from an ad hoc query, just as it does with a query plan for a data service function. If the same ad hoc query is run again, the cached query plan is used.

If you are relatively certain that the same ad hoc query will not be requested a second time, you can have the client application instruct Oracle Data Service Integrator not to cache the ad hoc query plan using the following request config attribute:

```
DO_NOT_CACHE_QUERY_PLAN
```

However, use of this directive is generally not needed. (Note that using the Oracle Data Service Integrator Filter API generates an ad-hoc query based on the filter.)

You can use the Oracle Data Service Integrator Query Plan view to display the query plan and present hints regarding operators that may not be optimal. Oracle Data Service Integrator also supplies audit information that shows additional details of the query run. Both of these tools are helpful for developing Oracle Data Service Integrator dataspace projects.

12.3.1.8 Writing Your Own SQL for Oracle Data Service Integrator to Use

You should only consider writing your own SQL statements for Oracle Data Service Integrator to use when it's easier than letting Oracle Data Service Integrator generate the SQL statements.

However, keep the following in mind if you choose to write your own SQL:

- It gets more difficult as the queries become more complex.
- Oracle Data Service Integrator is limited in the optimizations that it can perform with SQL that you have written.
- Oracle Data Service Integrator cannot use your SQL statements with a ppci-impl join.
- Oracle Data Service Integrator cannot use your SQL statements containing a left-outer join to implement nested elements.

Note that Oracle Data Service Integrator attempts to push all access to the same database into a single SQL statement, with the following exceptions:

- Oracle Data Service Integrator does not push unions to the database because this does not offer a significant improvement and forces the system to group the retrieved rows in the query engine. It can also break streamability.
- Oracle Data Service Integrator does not push cross-products to the database because a cross product with 1,000 rows on the left and 1,000 rows on the right results in 1,000,000 rows (instead of 1,000 + 1,000).
- Oracle Data Service Integrator does not push a left-outer join to the database in cases when the left-hand side does not have a unique key. For example, Oracle Data Service Integrator does not push the following:

```
A left-outer-join B on A.x=B.x left-outer-join C on A.y = C.y
```

This results in a cross-product as B and C are not related.

12.3.2 Exercising Care when Using Fail-over, Fail-over-retry, and Timeout

The BEA XQuery fail-over function is an XQuery extension that enables you to catch unexpected, unavoidable exceptions, such as a database or a Web service being unavailable.

To appreciate the advantage of using the fail-over function, consider the example of a Web service that is overloaded to the point that calling it blocks for five minutes and then fails. If an XQuery calls this Web service 100 times, it could take up to 500 minutes to complete the query.

The fail-over mechanism provides a solution in that once the primary expression fails, the assumption is that it will continue to fail for some time, so there is no need to re-evaluate within the same query execution.

Notwithstanding this feature, however, you should exercise care in using the fail-over, fail-over-retry, and timeout functions. This is because when you specify that a certain portion of the XQuery is to fail-over, the compiler and optimizer does exactly that and cannot combine expressions outside the fail-over with expressions inside the fail-over.

For example, suppose that you are retrieving a database table using the CUSTOMER() function, but want it to fail-over if the database is not available. You could use the following expression:

```
for $CUSTOMER in fail-over( CUSTOMER_in_db_1(), CUSTOMER_in_db_2())
```

Consider further that in a higher-level function, you select from that result using the following clause:

```
where $CUSTOMER/CUSTOMER_ID = $CustId
```

Because of the fail-over, Oracle Data Service Integrator is forced to read the entire CUSTOMER table and then, in memory, select the one customer that you need.

Note that use of Java functions put similar restrictions on the Oracle Data Service Integrator compiler and optimizer, since it cannot examine the Java functions (which are opaque to Oracle Data Service Integrator). This can also thwart optimizations.

One solution is to move Java function logic into Oracle Data Service Integrator operations. If this is not possible, and you find that the Java functions are blocking optimizations, you might explore using inverse functions.

12.3.3 Using Inverse Functions

Oracle Data Service Integrator enables data service developers to register inverse functions with the system, enabling you to define and use general user-defined data transformations without sacrificing query pushdown and updateability. Using this information, Oracle Data Service Integrator is able to perform a reverse transformation of the data when analyzing query predicates or attempting to decompose updates into underlying data source updates.

Consider the following example: You have a column called dateTimes in a database, defined as the number of milliseconds since January 1, 1970. You then declare a function called millisToDateTime() which converts the millisecond values to dateTime enabling you to query on that column and retrieve a dateTime that you can display as part of the query output.

Now suppose that you have a dateTime value in XQuery, either from a query argument or from another data source, and you want to retrieve all the rows that match the dateTime from the database table. The only obvious solution is to retrieve every row from the table, convert the milliseconds to dateTime, and then compare it to the dateTime value in the XQuery (discarding all rows that do not match).

This is very inefficient. It would be better to convert the dateTime to milliseconds, and push that selection to the database. You could write a dateTimeToMillis() function and use the output to select directly on the milliseconds column. But there are a couple problems with this approach.

First, you would need to expose the milliseconds column to the outside world. And second, you would need to know information about the milliseconds column to use the `dateTimeToMillis()` function.

A much more elegant solution is to define an inverse function for the `millisToDateTime()` function that, given the output of `millisToDateTime()`, returns the input argument. Once defined, Oracle Data Service Integrator can run the `dateTime` value through the inverse function to obtain the correct value in milliseconds to push to the database.

Since Oracle Data Service Integrator can determine that the milliseconds value from the database was sent through `millisToDateTime()`, it knows that applying the inverse function to a `dateTime` will return a value suitable for the database.

12.3.4 Using Caching and Auditing

Caching or auditing a specific function requires that the function call remain in the query plan. This means that if you cache a function such as `getOrders($customer)`, the optimizer cannot combine the function `getCustomers()` with `getOrders($customer)` to produce the following SQL statement:

```
select ... from CUSTOMER C, ORDER O where C.CID eq O.CID
```

Instead, the optimizer has to leave the `getOrders($customer)` function as the following:

```
select ... from ORDER O where O.CID = ?
```

Another consideration may be less obvious. Suppose you have a function `getCustomerCreditScore($cid)` that retrieves the credit score from the `getAllCustomerInfo($cid)` function. If the `getAllCustomerInfo($cid)` function is not cached, then the resulting query plan needs to access only the database tables necessary to get the credit score.

However, if the `getAllCustomerInfo($cid)` function is cached, then the optimizer cannot determine exactly the data from `getAllCustomerInfo($cid)` that may be required by a different function that calls `getAllCustomerInfo($cid)`.

This means that Oracle Data Service Integrator will have no choice but to retrieve and cache everything. In the case that the `getAllCustomerInfo($cid)` function retrieves the customer profile, orders, customer support cases, credit information, and so on, and all you need is the credit score, you would be forced to retrieve (and cache) an excessive amount of data just to get the credit score.

12.3.5 Query Plans

This section describes performance and optimization best practices related to query plans.

12.3.5.1 Evaluating Performance Before Running the Query

You should consider examining the query plan for an operation and evaluating its expected performance before running the query. This is useful because, in some cases, the operation may run endlessly as a result of an infinite loop.

12.3.5.2 Precompiling Query Plans

Precompiling query plans can increase performance, especially since the initial compile time can sometimes be significant. However, this is not always the case.

Instead of immediately precompiling your query plans, examine the plan to see if you can reduce the compile time in other ways. For example, using compound mappings can often help.

Tip: Post to the Oracle Data Integration forum if you would like to be forwarded a startup class that precompiles query plans.

12.3.5.3 Evaluating the Performance by Running the Query

Ultimately, the best way to evaluate the performance of an operation is to run the query and examine the results.

To evaluate performance, do the following:

1. Turn on auditing.

Note: Do not enable auditing of individual functions; this changes the query plan.

2. Run the query and examine the audit information.
3. If the query does not finish running, check the query plan. Look at the individual SQL execution times. If you see simple SQL statements that take excessively long to run, check that you have the appropriate indexes defined.
4. Run the SQL through 'explain plan' and check whether the SQL is retrieving more rows than needed.
5. Finally, verify that the selections and joins are getting pushed to the database.

12.3.6 Monitoring Operational Performance and Service Level Agreements

For information about monitoring operational performance and Service Level Agreements (SLA), refer to the audit and profiler samples installed with Oracle Data Service Integrator.

```
<ALDSP_HOME>/samples/Audit  
<ALDSP_HOME>/samples/Profiler
```

The listener sample writes audit records to a database and produces the necessary reports.

12.4 How To Get More Help

For additional help and further suggestions for dataspaces optimization, post to Post to the Oracle Data Integration forum. Customers can also open cases with Oracle Customer Support, as required.

12.5 Related Topics

XQuery Reference Information

W3C XML Query (XQuery)

XQuery 1.0 and XPath 2.0 Data Model (XDM)

XQuery 1.0 and XPath 2.0 Functions and Operators