

WhereScape Red Teradata

User Guide

8.1.1.0

WhereScape RED Teradata User Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Copyright Notice

Copyright © 2002-2018 WhereScape Software Limited. All rights reserved. This document may be redistributed in its entirety and in this electronic or printed form only without permission; all other uses of this document and the information it contains require the explicit written permission of WhereScape Software limited.

Due to continued product development, this information may change without notice. WhereScape Software Limited does not warrant that this document is error-free.

Trademarks

WhereScape and WhereScape RED are trademarks or registered trademarks of WhereScape Software Limited. Other brands or product names are trademarks or registered trademarks of their respective companies.

WhereScape USA, Inc

1915 NW AmberGlen Parkway
Suite 400, Beaverton
Oregon 97006
United States
T: 503-466-3979
F: 503-466-3978

WhereScape Limited

P.O.Box 56569, Auckland 1446
12-16 Tapora Street
Quay Park
Auckland 1010, New Zealand
T: +64-9-358-5678
F: +64-9-358-5679

WhereScape Europe

450 Brook Drive
Green Park
Reading RG2 6UU
United Kingdom
T: +44-118-914-4509
F: +44-118-914-4508

WhereScape Asia Pte. Ltd

300 Tampines Avenue 5
#09-02 Singapore 529653
T: +65-6679-5728

CONTENTS

Overview	1
Overview of WhereScape RED	2
How to use this Guide	5
Design	7
Objects and Windows	9
Object Types	10
Working with Objects	15
Connections	15
Load Tables	17
Dimension Tables	22
Data Store Tables	25
EDW 3NF Tables	27
Stage and Aggregate Tables	28
OLAP Cubes	29
OLAP Dimensions	32
Procedures	35
Scripts	37
Indexes	38
Views	39
Join Indexes	40
Exports	41
Source Mapping	42
Object Check-Outs and Check-Ins	44
Re-Create Dialog	49
Organizing Objects	51
Adding Objects to Projects	54
Removing Objects from Projects	56
Using Project/Object Maintenance	57
Adding Projects to Groups	58
Removing Projects from Groups	59
Moving Projects within Groups	59
List Projects Memberships for an Object(s)	59
Windows and Panes	60
Builder Window	61
Scheduler Window	66
Procedure Editor Window	69
Export Middle Pane Output	71
Find Function	73

Tutorials	75
<hr/>	
Default Settings	77
<hr/>	
Settings - Options	78
Settings - Repository Identification	79
Settings - Repository Privacy Settings	80
Settings - Object Types	83
Object Type Availability	83
Object Type Names	84
Object Type Ordering	85
Object Type End User Setting	86
Object Type Icon	87
Object Type Color	90
Object Sub Types	91
Settings - Global Naming Conventions	93
Case Conversion	93
Global Naming of Tables	94
Global Naming of Source Mappings	95
Global Naming of Indexes	96
Global Naming of Key Columns	97
Global Naming of Procedures	99
Settings - DSS Tables and Columns	100
DSS Tables	102
DSS Columns	103
DSS Columns for Custom Targets	105
Settings - Check-Out and Check-In	107
Code Generation	108
General	108
Default Update Procedure Options	110
Settings - Storage	112
Target Usage	112
Target Location	113
Table Storage	116
Default Optional CREATE Clause	117
Index Type	118
Settings - Metadata Versioning	119
Settings - Documentation	120
Settings - Available Load Types	121
Settings - Other	122
Settings - User Preferences	123
Settings - Common	124
Look and Feel	124
Local Naming Conventions	130
Local Paths	135
Outputs	136
Other	138
Settings - Current Repository	139
Look and Feel	139

Settings - Language Options	140
Parameters	143
<hr/>	
Connections	145
<hr/>	
Connection Types	147
Database - Data Warehouse/Metadata Repository	148
Database	155
ODBC	159
Windows	164
Other	166
UNIX	167
Hadoop	172
Microsoft Analysis Server 2005+	178
Microsoft Analysis Server 2005+ - OLAP Cubes	179
Microsoft Analysis Server 2005+ - Tabular Mode	181
Browsing a Connection	184
Connection Browse Properties	186
Changing a Connection's Properties	188
Reset Meta Database Connections	189
Connection Settings for BDA	190
Configuring the BDA Server	191
Big Data Adapter Settings	192
Configuring your database for use by BDA	192
Big Data Adapter Settings	193
Connection Extended Properties	195
<hr/>	
Table Properties	197
<hr/>	
Properties	198
Rebuilding Update Procedures	200
Storage	202
Table Storage Screen - Teradata	203
Location	203
Storage	204
Other	205
Table Storage Screen - Tabular	205
Location	206
Other	206
Processing	206
Bulk Table Storage Change	206
Override Create DDL	209
Source	210
Documentation Fields	210
Documentation Fields Screen	211
Table Extended Properties	212

Notes	213
Loading Data	215
Choosing the Best Load Method	217
Load Drag and Drop	218
Data Type Mappings	220
Database Link Load	221
Database Link Load - Properties	221
Database Link Load - Source Screen	223
ODBC Based Load	225
Native ODBC Based Load	225
Native ODBC Based Load - Source Screen	226
File Actions	228
Native Loads using UNIX and LINUX	231
TPT Load	233
TPT Load - Source Screen	235
Cleanup after TPT Load Failure	238
TPT UNIX/Linux Script Load	240
TPT UNIX Script Load - Properties	241
SSIS Loader	244
Loading Data into RED Load Tables using SSIS	245
Flat File Loads	253
Loading Data from Flat Files using SSIS	259
Flat File Load - Source Screen	268
Script based loads	275
XML File Load	277
External Load	281
Apache Sqoop Load	281
Handling Missing Source Columns	288
Load Table Transformations	290
Post-Load Procedures	290
Changing Load Connection and Schema	291
Dimensions	293
Dimensions Overview	294
Building a Dimension	295
Generating the Dimension Update Procedure	303
Source tab	308
Dimension Artificial Keys	314
Dimension Column Properties	316
Dimension Column Transformations	325
Dimension Hierarchies	326
Adding a Dimension Hierarchy	327
Using a Maintained Hierarchy	329
Snowflake	330

Creating a Snowflake	331
Dimension Language Mapping	332
Staging	333
<hr/>	
Building the Stage Table	334
Generating the Staging Update Procedure	337
Generating a Procedure	337
Procedure type	337
Locking Request Modifier	339
Source Table Mapping	339
Parameter selection	342
Dimension Joins	343
Dimension history information	344
Building and Compiling the Procedure	345
Stage Table Custom Procedure	345
Stage Table Column Properties	345
Stage Table Column Transformations	351
Permanent Stage Tables	352
Generating the Permanent Staging Update Procedure	352
Set Merge Procedure	358
Data Store Objects	364
<hr/>	
Data Store Objects Overview	365
Building a Data Store Object	367
Drag and Drop	367
Data Store Object Properties	367
Create and Load	368
Deleting and Changing columns	369
Adding additional columns	369
Create the table	370
Generating the Data Store Update Procedure	371
Generating a Procedure	371
Processing tab	371
Source tab	375
Building and Compiling the Procedure	376
Data Store Artificial Keys	378
Data Store Column Properties	382
Data Store Column Transformations	387
EDW 3NF Tables	388
<hr/>	
EDW 3NF Tables Overview	389
Building EDW 3NF Table	391
Generating the EDW 3NF Update Procedure	395
Processing tab	395

Source tab	399
Indexes	400
Converting an existing EDW 3NF Table to a EDW 3NF History Table	401
EDW 3NF Table Artificial Keys	403
EDW 3NF Table Column Properties	406
EDW 3NF Table Column Transformations	411
Data Vaults	413
<hr/>	
Data Vault Functions and Features	414
Load Table Meta Data Columns	414
Data Vault Stage Table	416
Hash Key Generation Wizard	417
Hub, Link and Satellite Creation Wizard	418
Data Vault Templates	419
Data Vault Settings	420
Object Types settings:	420
Global Naming Conventions settings:	421
DSS Tables and Columns settings:	422
Table Column Properties	423
Maintain Hash Key Columns	425
Building Data Vault Objects	426
Creating Load Tables	426
Creating Data Vault Stage Tables	428
Generating Update Procedures for the Data Vault Stage Table	437
Creating the Hub, Link and Satellite Tables	440
Creating the Hub table	440
Creating the Link table	444
Creating the Satellite table	448
Generating Update Procedures for Hub, Link and Satellite Tables	452
Hub table	452
Link and Satellite Tables	454
Changing the Data Vault Hash Key Function in WhereScape RED 6.9.1.0 and above	455
Custom Objects	459
<hr/>	
Model Tables	461
<hr/>	
Model Table Overview	462
Building a Model Table	464
Drag and Drop	464
Model Table Properties	465
Create and Load	466
Deleting and Changing columns	467
Adding additional columns	467
Create the table	468

Generating the Model Table Update Procedure	469
Generating a Procedure	469
Business Key definition	469
Locking Request Modifier	471
Source Table Mapping	472
Building and Compiling the Procedure	474
Model Table Artificial Keys	475
Model Table Custom Procedure	475
Model History Tables	476
Generating History Table Update Procedures	478
Model Table Column Properties	481
Model Table Column Transformations	486
Fact Tables	487
<hr/>	
Detail Fact Tables	488
Creating Detail Fact Tables	488
Generating the Detail Fact Update Procedure	490
Generating a Procedure	490
Define Fact Procedure Type and Options	490
Template	491
Define Fact Business Key Columns	491
Source Tab	494
Fact Table Column Properties	495
Fact Table Column Transformations	500
Fact Table Language Mapping	501
Aggregation	503
<hr/>	
Creating an Aggregate Table	504
Creating an Aggregate Summary Table	504
Aggregate Table Column Properties	505
Aggregate Table Column Transformations	510
Join Indexes	511
<hr/>	
Creating a Join Index	512
Views	515
<hr/>	
One to One Views	516
Dimension Views for Aliasing	518
Compound Views, Facts and Dimensions	520
Dimension View Hierarchies	524
Adding a Dimension View Hierarchy	524
Creating a Custom View	526

View Aliases	528
Analysis Services OLAP Cubes	531
OLAP Overview	532
OLAP Defining the Data Source for the OLAP Cube	533
OLAP Defining an OLAP Cube	535
Building a New OLAP Cube	535
Setting Cube Properties	538
OLAP Inspecting and Modifying Advanced Cube Properties	539
OLAP Creating an OLAP Cube on the Analysis Services Server	540
OLAP Cube Objects	541
OLAP Cube Properties	541
OLAP Cube Measure Groups	546
OLAP Cube Measure Group Processing/Partitions	548
OLAP Cube Measure Group Partitions	552
OLAP Cube Measures	554
OLAP Cube Calculations	557
OLAP Cube Key Performance Indicators	560
OLAP Cube Actions	563
OLAP Cube Dimensions	567
OLAP Cube Measure Group Dimensions	570
OLAP Dimension Objects	573
OLAP Dimension Overview	573
OLAP Dimension Attributes	577
OLAP Dimension Attribute Relationships	580
OLAP Dimension Hierarchies	582
OLAP Dimension User Defined Hierarchy Levels	584
OLAP Changing OLAP Cubes	586
OLAP Retrofitting an OLAP Object	588
Transformations	593
Column Transformations	594
Column Transformation Properties	595
Load Table Column Transformations	597
Database Link During Load Transformations	598
File During Load Transformations	600
After Load Transformations	601
Teradata User Defined Functions	602
Teradata UDF Example	603
Re-usable Transformations	605
Creating a New Re-usable Transformation	605
Specify the Name of the Transformation	606
Enter Re-usable Transformation Metadata	608
Define the Transformation Model	609
Completed Re-usable Transformation	610

Changing a Re-usable Transformation	611
Applying Changes to Dependant Transformations	612
Using Re-usable Transformations	613
Exporting Data	614
<hr/>	
Building an Export Object	615
File Attributes	619
File Attributes - SSIS Exports	622
Export File Definition	623
SQL Server Integration Services (SSIS)	623
Export Column Properties	624
Script based Exports	627
Procedures and Scripts	629
<hr/>	
Procedure Generation	631
Procedure Editing	637
Procedure Loading and Saving	640
Procedure Comparisons	642
Procedure Compilation	643
Procedure Running	643
Procedure Syntax	644
Procedure Properties	645
Macros	647
BTEQ Scripts	647
Script Generation	647
Script Generation (Windows/Teradata)	648
Windows PowerShell Scripts	654
Script Editing	658
Script Testing	659
Script Syntax	659
Script Environment Variables	661
Calling a Batch File from a Script	667
Scheduling Scripts	669
Manually created scripts	671
Templates	673
<hr/>	
Template Properties	675
Template Editor	677
Evaluating an API Outline Template	678
Template Usage	682
Windows PowerShell Templates	683
Scheduler	685
<hr/>	

Scheduler Options	686
Auto	689
Tools	690
Select Job Report Fields	690
Scheduler States	691
Scheduling a Job	693
Working with Jobs	699
Creating a Job	701
Editing a Job	711
Editing Tasks in a Job	715
Editing Task Dependencies	722
Show Dependencies Diagram	727
Inserting a Copy of a Job	729
Deleting a Job	731
Deleting Job Logs	732
Starting a Job	734
Halting a Job	735
Aborting a Job	736
Restarting a Job	737
Creating an Application from a Job	742
Stand Alone Scheduler Maintenance	745
SQL to return Scheduler Status	748
Reset Columns in Job and Task View	749
Stopping a Linux/UNIX Scheduler from within RED	750
Indexes	753
<hr/>	
Index Definition	754
Documentation and Diagrams	759
<hr/>	
Creating Documentation	760
Batch Documentation Creation	763
Reading the Documentation	764
Diagrams	765
Types of Diagrams	765
Schema Diagram	767
Source Diagram	769
Joins Diagram	773
Links Diagram	774
Impact Diagram	775
Dependency Diagram	776
Working with Diagrams	778
Creating a Job from a Diagram	781
Creating an Application from a Diagram	783
Creating a Project from a Diagram	785

Reports	787
Dimension-Fact Matrix	788
OLAP Dimension-Cube Matrix	789
Dimension Views for a Specified Dimension	790
Column Reports	791
Columns without Comments	791
All Column Transformations	793
Re-Usable Column Transformations	794
Column Track-Back	795
Column Track-Forward	797
Table Reports	799
Tables without Comments	799
Load Tables by Connection	801
Export Objects by Connection	802
Records that failed a Dimension Join	803
External Source Tables/files	805
Procedure Reports	806
Modified Procedures	806
Custom Procedures	807
Object Reports	808
Objects-Projects Matrix	808
Modified Objects (excluding indexes)	809
Objects Checked-out	811
Loaded or Imported Objects	812
Objects with Extended Properties	813
Job Reports	816
Object-Job Matrix	816
Jobs with an Object	818
Tasks of a Job	819
Operational Reports	821
Object Performance History	821
Job Performance History	822
Task Performance History	823
Validate	825
Validate Meta-data	826
Validate Workflow Data	826
Validate Table Create Status	826
Validate Load Table Status	827
Validate Procedure Status	827
List Meta-data Tables not in the Database	827
List Database Tables not in the Meta-data	828
List Tables with no related Procedures or Scripts	829
List Procedures not related to a Table	831

Teradata: View of Model Validate	832
Query Data Warehouse Objects	833
Promoting Between Environments	835
Applications	836
Application Creation	837
Creating an Application	837
Define an Application distribution	837
Output Directory	838
Objects to Add/Replace	840
Objects to Delete	841
Application Loading	842
Creating and Loading Applications from the Command Line	843
Importing Object Metadata	843
Importing Language Files	845
Data Warehouse Testing	846
Backing Up and Restoring Metadata	851
Backup using DB Routines	852
Unloading Metadata	853
Loading an Unload	856
Altering Metadata	861
Validating Tables	862
Validating Source (Load) Tables	864
Validating Procedures	865
Altering Tables	866
Validating Indexes	867
Recompiling Procedures	867
Callable Routines	869
Introduction to Callable Routines	870
Callable Routines API	870
Callable Routines per RDBMS	872
Callable Routines Names Qualifier	873
Callable Routines Common Input	873
Callable Routines Invocation	874
Alternative Invocation Methods	875
Ws_Api_Glossary	877
Ws_Connect_Replace	879
Ws_Job_Abort	881
Ws_Job_Clear_Archive	882
Ws_Job_Clear_Logs	885

Ws_Job_Clear_Logs_By_Date	887
Ws_Job_Create	889
Ws_Job_CreateWait	893
Ws_Job_Dependency	897
Ws_Job_Release	900
Ws_Job_Restart	902
Ws_Job_Schedule	905
Ws_Job_Status	907
Ws_Load_Change	913
Ws_Maintain_Indexes	915
Ws_Version_Clear	918
WsParameterRead	921
WsParameterReadF	922
WsParameterReadG	923
WsParameterWrite	925
WsWrkAudit	926
WsWrkAuditBulk	929
WsWrkError	932
WsWrkErrorBulk	935
WsWrkTask	938
Ws_admin_v Views	941
Ws_admin_v_audit	942
Ws_admin_v_error	943
Ws_admin_v_sched	944
Ws_admin_v_task	944
Retrofitting	947
Migrating the Data Warehouse Database Platform	948
Importing a Data Model	957
Re-Targeting Source Tables	964
Retro Column Properties	966
Retro Column Properties Screen	970
Retro Column Transformations	971
Integrating WhereScape RED into an Existing Warehouse	973
Rebuilding	974
Integrating	975
Integrating, Host Scripts	976
Integrating, Selecting a Table Type	977
Integrating, Questions	978
Integrating, Procedures	981
Integrating, Views	982

Integrating, WhereScape Tables	982
Relationship Maintenance	983
Add Relationship	984
List Relationships	985
Generate Relationships	986
Upgrading RED	987
Login Checks	989
Data Type Mappings	991
Using Data Type Mapping Sets	992
Maintaining Data Type Mapping Sets	994
Creating a New Data Type Mapping Set	996
Copying a Data Type Mapping Set	1000
Editing a Data Type Mapping Set	1003
Deleting a Data Type Mapping Set	1006
Loading Custom Data Type Mapping Sets	1008
Exporting Custom Data Type Mapping Sets	1010
Custom Data Type Mapping Examples	1012
Column Context Menu	1019
Properties	1020
Change Column(s)	1023
Add Column	1025
Duplicate Column	1026
Delete Column	1027
Re-space Order Number	1028
Impact	1029
Sync Column order with database	1031
Send Columns to Another Object	1032
Database Functions	1035
Using Database Function Sets	1036
Column Transformation Properties Dialog	1036
Transformation Definition screen	1037
Connection Properties screen	1038
Maintaining Database Function Sets	1039
Creating a New Database Function Set	1041
Copying a Database Function Set	1043
Editing a Database Function Set	1045

Deleting a Database Function Set	1056
Loading Database Function Sets	1058
Exporting Database Function Sets	1061
Gather Statistics	1063
<hr/>	
Define Statistics	1063
Extended Properties	1065
<hr/>	
Defining Extended Properties	1066
Creating an Extended Property Definition	1067
Variable Settings	1068
Scope Settings	1069
Maintaining Extended Property Definitions	1070
Extended Properties Value Assignment	1072
Setting Up Extended Property Values for a Connection	1072
Setting Up Extended Property Values for an Object	1073
Extended Properties Lookup	1073
Extended Properties Data Migration Between Repositories	1074
Exporting Extended Properties	1074
Loading Extended Properties	1076
Table and Column Comments	1079
<hr/>	
Defining Table Comments	1080
Exporting Table Comments	1081
Loading Table Comments	1084
Defining Column Comments	1088
Exporting Column Comments	1089
Loading Column Comments	1093
Importing Comments from an External Source	1097
Viewing the Import/Export Logs	1099
Multi Source Processing	1101
<hr/>	
Multi Source Functions and Features	1102
Source Mapping Object	1102
Source Mapping Tool	1103
Global Naming of Source Mappings	1104
Independent Execution of Update Procedures	1105
Table Column Properties	1106
Adding Source Mapping to Objects	1107
Drag and Drop	1109
Maintaining Source Column Mappings	1112
Generating Update Procedures for Source Mapping Objects	1116
Executing Update Procedures via Scheduler	1119

CHAPTER 1

OVERVIEW

WhereScape RED User Guide for Teradata

The WhereScape RED **User Guide** for Teradata is available either as online help, as a PDF, or in a printed manual. The User Guide provides information on how to use WhereScape RED to build a data warehouse.

This online help version of WhereScape RED's documentation is for RED 8.1.1.0, last updated Wednesday, April 18, 2018.

IN THIS CHAPTER

Overview of WhereScape RED	2
How to use this Guide	5

OVERVIEW OF WHERESCAPE RED

Traditionally data warehouses take too long to build and are too hard to change. WhereScape RED is an Integrated Development Environment to support the building and managing of data warehouses. It has the flexibility to enable you to build an entire enterprise data warehouse or just the user facing views, aggregates and summaries.

In all cases, the core values of WhereScape RED are twofold: its rapid capabilities that enable better data warehouses to be built, faster, and its integrated environment that simplifies management.

As a data warehouse specific tool, WhereScape RED embodies a simple, pragmatic approach to building data warehouses. With WhereScape RED you specify what you want to achieve by dragging and dropping objects to create a meta view, and then let WhereScape RED do the heavy lifting of creating the necessary tables, procedures, etc. Data warehouse wizards prompt for additional information at critical points to provide the maximum value from the generated objects.

Different data warehousing approaches including agile, prototyping and waterfall are supported by WhereScape RED. Agile developers will find specific functionality has been included to support common agile practices. For developers who are new to data warehousing, or are looking for a fast, pragmatic approach, WhereScape's Pragmatic Data Warehousing Methodology can be used.

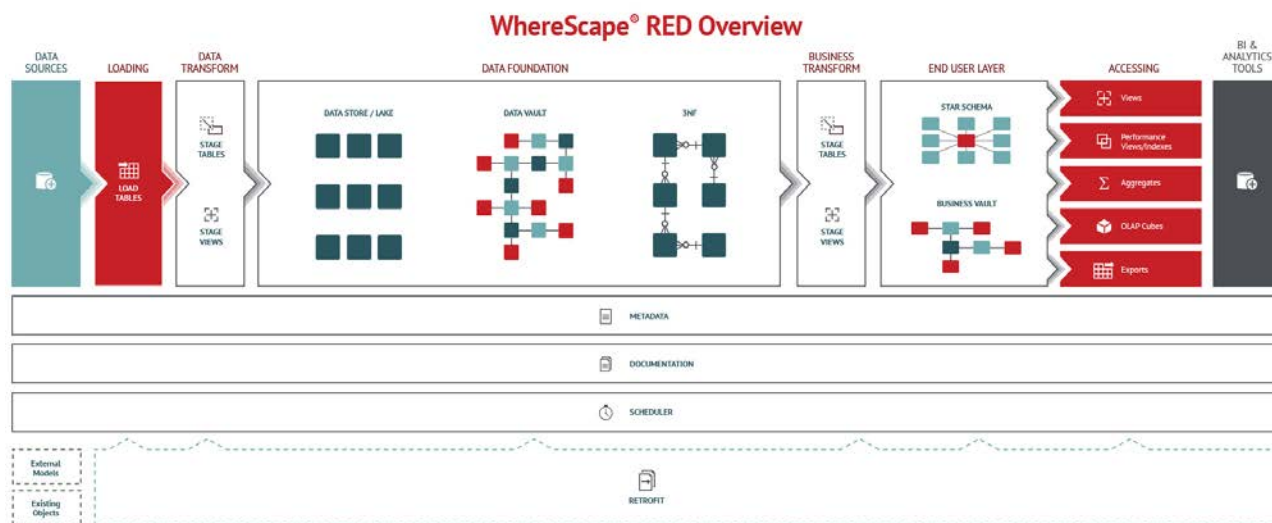
The basic concepts behind WhereScape's Pragmatic Data Warehousing Methodology are:

- minimize the impact on the source systems
- centralize management within the data warehouse
- store transactional data at the lowest practical grain within the data warehouse
- snapshot, combine and rollup transactional tables to provide additional value
- utilize views or cubes for end user access
- allow for incremental loads from day one
- use an iterative approach
- minimize cleansing and transformations to ease source system reconciliation
- design the data warehouse independently from the end user tool layer

WhereScape RED supports these concepts to facilitate very rapid building of data warehouses.

WhereScape RED controls the flow of data from the source systems through transforming and modeling layers to analysis areas. Different styles of data warehousing (EDW 3NF, dimensional, etc.) are supported and utilize different objects, but all follow the same basic flow.

The diagram below shows the objects and the information flow:



Data Flow

Data is moved from source tables to load tables via scripts, database links and ODBC links. These load tables are created by dragging and dropping from a connection object. Load tables are generally based on source system tables. Their main purpose is to be a destination for moving data as simply and quickly as possible from the source system. Load tables will generally hold a single unit of data (e.g. last night or last month), and will be truncated at the start of each extract. Transformations can be performed on the columns during the load process if required.

Load tables feed stage tables, which in turn feed data store, dimension tables. Data from multiple load tables can be combined at this level.

First tier transactional tables (fact or dimension tables) are created and updated from stage tables. Second tier tables (summary rollup, aggregate, join indexes, etc.) are created and updated from lower level tables.

Feeds for downstream systems (exports) can be created by dragging and dropping from a dimension, view or aggregate object. Fast export or parallel transporter are used to generate file exports.

Cubes can be created from transactional tables or views.

Procedural code

WhereScape RED for Teradata generates procedural code using Teradata database procedures at each stage in the data warehouse build process. The generated code is, in nearly all cases, sufficient to create a rapid prototype of the data warehouse.

While the generation of code is often seen as a key benefit of WhereScape RED, the ability to control and manage custom code is also critical to the long-term management of the data warehouse environment.

In most cases 85-100% of the generated code will be taken through to production with no customization required.

Scheduler

The flow of data from the source systems to data warehouse tables is controlled and managed by the WhereScape RED scheduler. All generated code includes audit and error logging logic that is used by the scheduler.

The scheduler provides a single point of control for the warehouse. From the scheduler, the state of all jobs can be ascertained. Any warning or error messages can be investigated and if a problem occurs, the scheduler controls the restart of the job from the point of failure.

Documentation

Documenting the warehouse is often a task left until last, and in many cases done once (if at all!) and not kept up to date. WhereScape RED generates user and technical documentation, including diagrams, in HTML format.

Technical documentation includes copies of all current procedures.

User documentation includes a glossary of business terms available independently of any end user tool.

Where additional specific information needs to be included in the documentation, WhereScape RED supports the inclusion of custom HTML pages in the generated output. This means in many cases the entire documentation requirements can be managed from one location, and regenerated as changes occur.

WhereScape RED and ETL Tools

WhereScape RED's core strength is in the rapid building of data warehouse structures. Organizations that have already purchased traditional ETL tools can use WhereScape RED as a pureplay data warehouse toolset. WhereScape RED can be used to iteratively build data marts or presentation layer objects that need to be constantly updated to keep relevant for end users. In most cases, customers will find that WhereScape RED has enough ETL capabilities to build the entire data warehouse, using the database rather than a proprietary engine to perform ETL processing.

The cross over in functionality between ETL tools and WhereScape RED is not large. WhereScape RED is tightly integrated into the data warehouse database and has an embedded data warehouse building approach. For WhereScape data movement is the start of the process—from source system to load tables. The key benefits of the product: development productivity and an integrated environment to manage and maintain your warehouse, come after the data movement stage. Where a traditional ETL tool is already in use, the output of the ETL process is a WhereScape RED Load, Stage or Dimension table from which WhereScape RED builds more advanced data warehouse structures.

HOW TO USE THIS GUIDE

The WhereScape RED **User Guide**, the online help, and the WhereScape RED **Installation and Administration Guide** assume that you are proficient in the use of the Windows operating system.

WhereScape RED often provides multiple ways to accomplish a task. In some cases, you can use the main menu, the right-click menu, or a toolbar, or a key combination (e.g. Alt/M and Ctrl/M to raise menus). Instructions in this documentation generally include only the most convenient method of accomplishing a task.

The following sources of information are available with WhereScape RED:

WhereScape RED Installation and Administration Guide

The **Installation and Administration Guide** is available either as online help, as a PDF, or in printed format. The Installation and Administration Guide provides the information needed to:

- Install the WhereScape RED software
- Validate the various software components required by WhereScape RED
- Install the WhereScape RED metadata
- Install a scheduler
- Optionally install third party data warehouse applications
- Upgrade the WhereScape RED software
- Create and load language files.

WhereScape RED User Guide for Teradata

The WhereScape RED **User Guide** for Teradata is available either as online help, as a PDF, or in a printed manual. The User Guide provides information on how to use WhereScape RED to build a data warehouse.

Sql Admin User Guide

The **Sql Admin User Guide** is available either as online help, or as a PDF. It provides documentation in the use of the stand-alone SQL query tool shipped as part of the WhereScape RED product.

WhereScape Forum

A **web forum** is available at <http://www.wherescape.com>. This forum contains information on the latest version, and bug reports that may be relevant for installation. In addition the WhereScape Blog is available at <http://www.wherescape.com> which may provide additional information.

CHAPTER 2

DESIGN

Design Introduction

WhereScape RED for Teradata can be used to build data warehouses based on any number of design philosophies from EDW 3NF enterprise data warehouses with consumer data marts through to federated or conformed star schema based warehouses. In the absence of another approach, the following methodology can be used for the design of data warehouses.

Note: This section can be skipped if you already have data warehouse design experience or a methodology you wish to utilize. It is meant to provide the novice designer with some tips for designing a data warehouse.

Design Approach

The concepts behind the WhereScape Pragmatic Data Warehouse Methodology are as follows:

- 1 Building an enterprise-wide data warehouse is a process—an evolution rather than a big bang. Start small and grow the warehouse in manageable chunks until all the pieces are in place. Once you reach that stage, changes and new source systems will continue the process.
- 2 You need to understand the big picture, but not get lost in it. Talk to all the various departments, business units and companies within the organization. Do so at a relatively high level and try to understand how the information from each area impacts or affects the others. Identify commonalities and areas where the same information is handled in different ways. This process should take days or weeks not months.
- 3 Identify the high value, high return and possibly easiest areas of the business. Drill down in these areas and break down the workload into small manageable chunks of work, for example, one to two analysis areas. Agree on the first component of the data warehouse and do that.
- 4 Get an understanding of the source system for this first component or analysis area. If possible, get an entity relationship diagram and talk to the people who built or support the application. Identify the tables that contain the key information you will need. The goal is a quick and initial view, a detailed specification is not required.
- 5 Design the first component. This design should be a first draft, and can be written rather than using a design tool. Remember at this stage what the end users want is not really known, so don't set the design in concrete, or spend a large amount of time in this area.

Note: Experienced users of Wherescape RED will often dispense with a design and go straight to building a prototype.

- 6 Build a prototype. In most cases this should not take more than one or two weeks - experienced WhereScape RED developers can expect to build prototypes in hours or days. Concentrate on the detailed and descriptive data, unless you have a clear picture of the summarized requirements. Do as much as possible in terms of validating the data back to the

source system. If dealing with a large or complex source system then only deliver a segment in this prototype, e.g. one branch, one store, one product group, etc. Keep It Simple.

- 7** Demonstrate the prototype to a group of the key users. Then drill down to a subset of key users (we recommend no more than three) who will help you go forward with the design. If possible give these users access to the prototype and get them using the data. Stress that data accuracy is not the issue at this stage, rather the look and feel.
- 8** Enhance the prototype with the feedback provided by the users. Again a quick process. If complicated requirements evolve then create a plan to implement, doing the highest value parts first. The goal is to get quick buy in and support from the two or three key users.
- 9** Provide key users access to the reworked prototype and get them using the data. Have them define the business names for all the measures and attributes, and to define any pre-calculated measures that they frequently use. Get them to define the hierarchies in the data. Ascertain the commonly utilized queries and reports, and see if there would be a better way of presenting these.
- 10** From the user feedback look at the need or possibility of using higher level fact tables, such as summaries, aggregates, snapshot or composite rollup tables.

The concepts and methodologies for designing and building a data warehouse are beyond the scope of this manual. It is assumed that the reader understands the basic concepts of a data warehouse, and is familiar with modeling, EDW 3NF, star and snowflake schemas, dimensions, fact tables, etc.

Refer to the WhereScape web site for a basic overview of data warehouse design if required.

CHAPTER 3

OBJECTS AND WINDOWS

WhereScape RED makes use of an object concept when dealing with the different components that make up a data warehouse solution.

The main object types are: Connections, Load Tables, Dimensions, Stage Tables, Fact Tables, OLAP Cubes, Aggregates, Procedures, Host Scripts, Indexes, Retros and Exports.

This chapter explains and provides an overview of each of these object types and how they can be managed and organized. The full functionality of each object is covered in the following chapters.

The various Windows, Panes and Views that form the WhereScape RED tool are also explained.

IN THIS CHAPTER


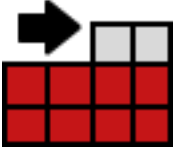
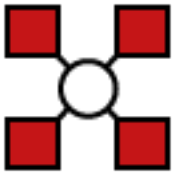

Object Types	10
Working with Objects	15
Organizing Objects	51
Windows and Panes	60
Export Middle Pane Output	70
Find Function.....	73


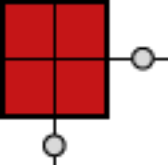

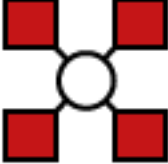
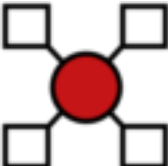

OBJECT TYPES

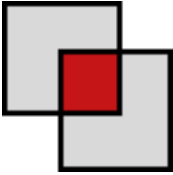




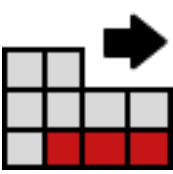

WhereScape RED has a concept of objects which are combined to build the data warehouse. Each WhereScape RED object has properties that allow the data warehouse developer to change how the object is used.







Note: Some objects may not be available for certain types of WhereScape RED licenses.


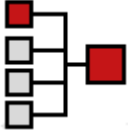

WhereScape RED objects include:

Object Type	Purpose
Connection 	<p>Connections define the path to external objects such as source data. Examples of connection object types are databases, analysis services cubes, operating systems or ODBC sources. Connections isolate environments simplifying, for example, the promotion of code between development and production.</p>
Load Table 	<p>Load tables are the first entry point of data into the data repository, and typically hold the latest set of change data. These objects contain their definition. Load tables can be defined as external, file, script or XML. Based on their definition they will, for example, run a predefined script or create a load script at run time. Pre-load actions (e.g. truncate) or post load procedures can be defined as part of a load object. In addition, transformations (either during or after the load) can be defined against columns in a load table.</p>
Dimension 	<p>Dimension tables are the constraining elements in the star schema design, and are defined by this object type. WhereScape RED will automatically generate procedural code for the three standard types of slowly changing dimensions, as well as date ranged dimensions (where the current version is defined by an external system). WhereScape RED also ships with a standard time dimension which can of course be extended. Dimensions can also be defined as mapping or work tables which do not appear in the generated user documentation.</p>
Dimension View 	<p>Note: WhereScape RED version 8.1.1.0 and above no longer supports the Dimension View object type when creating new objects. However, users that have existing Dimension View objects in their data warehouse can retain and continue to use them.</p> <p>A dimension view is a database view of a dimension table. It may be a full or partial view. A common usage is to create views where multiple date dimensions exist for one fact table. Other types of views supported by WhereScape RED include fact views, other table views, work views and user defined views.</p>

Object Type	Purpose
Stage Table 	<p>Stage tables are used in the transformation of raw data into model or star schema format. They typically hold only the latest set of change data. As well as custom procedures, WhereScape RED can generate different types of procedural code based on the complexity and speed of the data set. A stage table can also be defined as a work table, which has the same properties as a stage table but does not appear in the generated user documentation.</p>
EDW 3NF Table 	<p>An EDW 3NF table is a data warehouse object used to build third normal form enterprise data warehouses. In WhereScape RED, EDW 3NF objects have many of the code generating attributes of stage, dimension and fact tables. Third normal form enterprise data warehouses can be thought of as a source system for star schema data marts. Alternatively, they may be reported off directly by users and reporting tools.</p>
Data Store Table 	<p>A Data Store Table is a data warehouse object used to store any type of data for later processing. In WhereScape RED, Data Store objects have many of the code generating attributes of stage, dimension and fact tables. Data objects can be thought of as a source system for the data warehouse. Alternatively they may be reported off directly by users and reporting tools. Data Store Objects can be considered either reference or transactional in nature.</p>
Model Table 	<p>Note: WhereScape RED version 8.1.1.0 and above no longer supports the Model Table object type when creating new objects. However, users that have existing Model table objects in their data warehouse can retain and continue to use them.</p> <p>Model objects are used to create EDW 3NF models in an enterprise data warehouse. They may contain surrogate keys to other model tables.</p>
Fact Table 	<p>Fact tables are the central table in a star schema design. This object type allows the definition of fact tables. They support transactional, rollup, snapshot or partitioned (detail, rollup or exchange) fact tables. Changing a fact table's properties to partitioned will start a partitioning wizard that prompts for the required information.</p>
Aggregate 	<p>The aggregate object type provides a means to speed up access by summarizing data to a higher grain.</p>

Object Type	Purpose
Join index 	Join index is a Teradata specific object used for performance across multiple tables.
View 	View objects are usually created as end user objects from any table in the data warehouse. The data or columns may be restricted or extra descriptions may be added for use by the end user or reporting tools.
OLAP Cube 	The OLAP Cube object type provides a means to develop and manage Microsoft Analysis Services cubes. Cubes would normally be built from fact tables and provide summarized data.
OLAP Dimension 	An OLAP Dimension is built by WhereScape RED for every dimension table associated with the fact (or aggregate) table the OLAP Cube is derived from. OLAP Dimensions are shared across one or more OLAP Cubes. In analysis services, a dimension is a group of attributes that represent an area of interest related to the measures in the cube and which are used to analyze the measures in the cube.
Index 	This object type defines database indexes used to improve the access times on any of the table object types (i.e. Load, Stage, Dimension and Aggregate).
Export 	Exports are used to manage exports from the data repository.
Retro 	Retros are used to load predefined data models from modeling tools and to retrofit existing tables into the WhereScape RED metadata.

Object Type	Purpose
Retro Copy 	<p>Retros can be used to copy data from an existing data warehouse into WhereScape RED metadata. Retros can be set as Retro Copy objects to enable data transfer from the existing data warehouse to the new data warehouse.</p>
Procedure 	<p>The procedure object type is used to define and hold database stored procedures. As such it may contain functions, procedures and packages that are generated, modified or custom developed.</p>
Host script 	<p>Host script objects are either Windows or UNIX scripts. These scripts are maintained within the WhereScape RED environment and can be scheduled to run in their host environments.</p>
Template 	<p>Template objects are used to generate DDL, update procedures and host scripts. Once a template has been created it can be associated with a table and an operation on that table. The template is then used to generate the script used for the associated operation.</p> <p>Each template is assigned a type and a target database, these properties are used to assist with filtering when associating table operations to templates. Note that not all operations support template script generation on all target databases.</p> <p>Utility type templates can contain common code for use by other templates.</p>
Hub 	<p>A Hub is a table of unique business keys, they usually contain a hash key, business key(s), load date and record source. Hubs should normally have at least one Satellite.</p>
Link 	<p>Links are many-to-many tables representing current and past relationships between two or more Hub entities and are used to describe associations, transactions, hierarchies and redefinitions of Hub entities in a Data Vault. Links have their own hash key and the hash keys for the Hubs that are linked as well as a Load Date and Record Source. The attributes describing the context of a link are stored in Satellite Tables (see below).</p>

Object Type	Purpose
<p>Satellite</p> 	<p>Satellites are Data Vault objects which contain metadata that provides context for Hub and Link entities at a given time or over a period of time. Each Satellite entity can contain information on one Hub or Link.</p> <p>Satellite tables contain a hash key for the parent Hub or Link, a timestamp for the date of change and relevant descriptive fields. Satellites are usually created once per source system. However, descriptive attributes can change at different rates, so Satellites can also be created based on rate of change.</p>
<p>Source Mapping</p> 	<p>Source mapping objects can be defined for a number of table object types, for example stage tables, hubs, links and satellites. They allow these target tables to be sourced from more than one set of source tables.</p> <p>Source mapping objects don't result in data warehouse tables but contain source table and column information metadata, including any transformations and join criteria. This metadata is used to generate a procedure or script to populate the target table. Source mappings support more than one insert/update routine from different sources to be defined for the target table, and either executed collectively or independently.</p> <p>Refer to the section Multi Source Processing (on page 1101) for details.</p>
<p>Custom1/Custom2</p> 	<p>Custom1 and Custom2 objects are user defined objects. These Object Types can be renamed in the Tools>Options>Object Types>Object Names menu.</p>

Connections are normally the first objects created. These connections are then used in the creation of load tables through the drag and drop functionality. Subsequent objects can also be created, using drag and drop.

It should be noted that although the object types have names that correspond with their primary usage, they can be used for other purposes. For example, the fact object type could be used to create persistent stage tables if required.

Some objects are not supported by all databases, and some advanced properties are specific to the different databases.

WORKING WITH OBJECTS

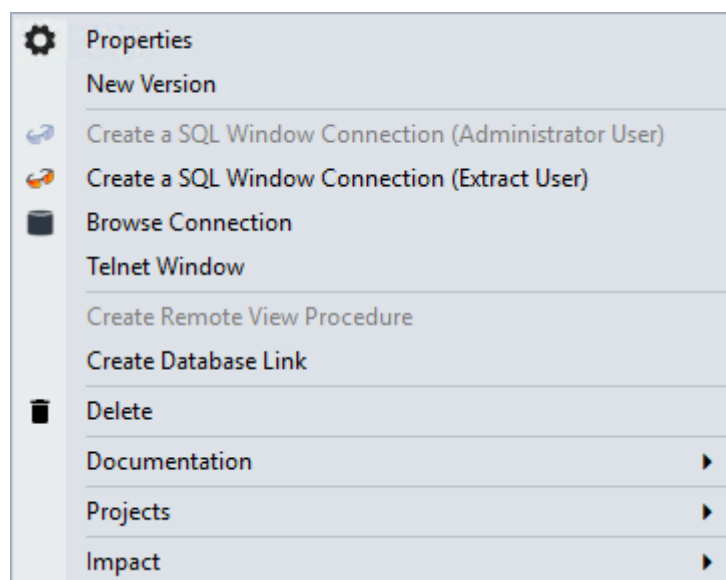
Most object types perform some form of action in the data warehouse. For example stage, dimension and aggregate table based objects are 'Updated' in the data warehouse via the defined update procedure. Procedures can be executed in the database.

When positioned on an Object in the left pane of the RED Builder window, the right-click pop-up menu provides a number of options for manipulating the object. Further options may be available through the menus provided in the various windows.

The operations of each of the objects is discussed in the following chapters. A brief overview of some of the more common operations follows:


CONNECTIONS

Connections, once defined are typically browsed and used as a source for drag and drop operations. For database connections, a database link is normally required. This link can be created via the right-click menu associated with a connection. See sample menu below:




Other operations available through the menu are editing the properties of the connection, creating a version of the connection, creating a telnet window for UNIX connections and creating a remote view creation procedure where required for database connections and loads.

The **Documentation** menu option can be used to generate (or read if already generated) the WhereScape RED HTML documentation for the selected connection. Two options are available: **Display** and **Create**:




The **Projects** menu option can be used to remove the connection from the current project or to add the connection to the current project. The **List Projects** option displays a list of projects which contain the current object, results are shown in the bottom pane. Multiple objects can be selected by double-clicking the **Connection** icon in the left pane and Ctrl + clicking multiple connections in the middle pane.

If there aren't any projects in the repository, these options are unavailable.

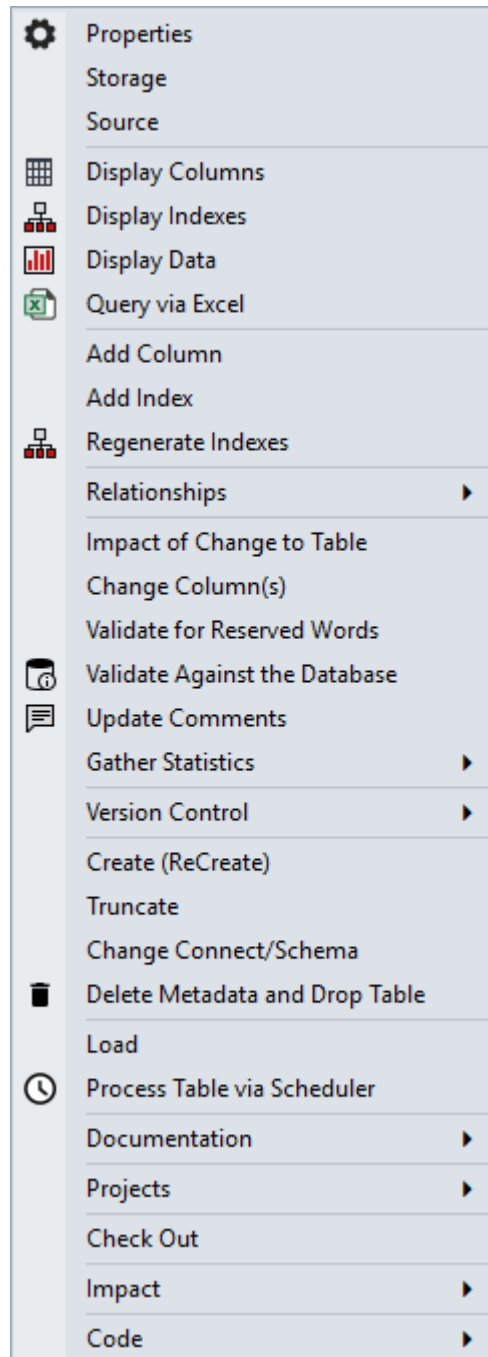


The **Impact** menu enables you to run reports on load and export objects associated with the connection.



LOAD TABLES

Load tables once defined would normally be created and loaded, unless these actions were performed as part of the drag and drop operation. The menu below shows the operations that can be performed on Load tables.



Properties, **Storage** and **Source** all launch the properties window for the Load table, albeit focused on different tabs within this window.

The columns and indexes of the load table can be displayed using **Display Columns** and **Display Indexes**. Any data in the load table can be displayed using **Display Data**. If the data is displayed, only the first 100 rows are returned from the table. Either the Sql Admin tool (accessible via the WhereScape start menu option), or the Excel query must be used if more detailed data analysis is required. Query the columns in Excel using **Query via Excel**.




TIP: When a column list has been displayed in the central pane, it is sorted based on the **order** field associated with each column. Clicking the column label **Col name** will sort the columns into alphabetic order. A subsequent click will re-sort based on the **order** field.

New columns and indexes can be manually added through this menu using **Add Column** and **Add Index**. Normally columns are added via drag and drop and most common indexes are created during the procedure generation phase.

The **Regenerate Indexes** menu option is used to add missing standard indexes. Selecting this menu item displays a dialog box with options to regenerate missing indexes in the metadata and recreate them or to just regenerate the missing indexes in the metadata.

The **Relationships** menu options enable the management of enhanced relationships. The **Add Relationship** option opens the Add Relationships dialog, the **List Relationships** option displays a list of enhanced relationships in the Drop Target Pane for the selected object and **Generate Relationships** generates relationships which have not yet been defined in metadata.



The **Impact of Change to Table** menu option produces a list of objects that will be potentially impacted by a change to the load table structure.

The **Change Column(s)** menu option to apply changes to a selected number of columns.

The **Validate for Reserved Words** menu option produces a list of table or column names where reserved words have been used; enabled for supported ODBC Drivers.

The metadata for the load table can be compared with the physical table resident in the database, using **Validate Against the Database** and where required, the table altered to match the metadata.

The **Update Comments** context menu option refreshes the table and column comments from the metadata, e.g. the description fields in the table and column Properties screen are updated.

Notes: You can manage the table and column comments outside the data warehouse environment via the **Tools>Table and Column Comments** menu option. Refer to the **Table and Column Comments** (on page 1079) chapter for more information.

Use **Gather Statistics** to gather statistics on a table. This action enables the underlying database to optimize each query, based on the statistics collected about the data that is being accessed. Users can either chose to **Define Statistics** or to **Refresh Full Statistics**.

Gathering statistics can be performed on any table by selecting this option from a table's right click menu, or to automate this process, by adding a statistics task to a job being processed by the scheduler (Stats, Quick Stats, Analyze or Quick Analyze). For more information about adding statistics tasks to jobs, see *Editing Tasks in a job* (on page 715).

The statistics process from the object context menu for Teradata tables is described below:


Teradata tables	
Define Statistics	Using the DlgDefineStats dialog, we save the DDL and execute it. This is object specific.
Refresh Full Statistics	COLLECT STATISTICS ON database.table (this requires statistics to have been executed on the object).



Tip: To Define Statistics on a Table Object see also *Gather Statistics* (on page 1063).

A version of a Load table is a copy of the metadata definition of the table at the time of the versioning. This version information can be used to create a new Load table, or can simply be left as a backup and reference point. Use **Version Control > New Version** to version a Load table. The **Build Application** menu option enables you to build an application file for the Load table and the **Duplicate Object** menu option enables you to create a new Load table as a duplicate of this table.

Note: The extended property values set in the original Load table is copied to the duplicate Load table. Refer to the *Extended Properties* (on page 1065) chapter for more information.



The **Create (ReCreate)** menu option creates the table in the database based on the definition stored in the metadata. To alter a table, select the Validate against database option (see the section on table validation).

The **Truncate** menu option truncates the table.

The **Change Connect/Schema** menu option enables the rapid changing of the connection information associated with the Load table. This information can be changed en-bulk for a number of Load tables. See *Changing Load Connection and Schema* (on page 291) for details.

The **Delete Metadata and Drop Table** menu option deletes the metadata definition for the table. It also gives you the option to drop the table in the database (dropping the table in the database is the default option). This is a permanent delete and no recovery is provided, so please use with caution. A version of the object's metadata is normally auto created (depends on the settings in **Tools>Options**).

The **Load** menu option performs an interactive load of the data. The method of loading depends on the type of connection. This menu option is intended for use with small data volumes as in a prototype environment. Large data volumes would normally be scheduled. The Load locks the WhereScape RED screen until completed.

Note: The Load option does not drop or create any indexes. Use the Process option, if indexes need to be maintained.

The **Process Table via Scheduler** menu option sends a request to the scheduler to immediately process the load table. This process will drop any indexes marked as `pre_drop`, load the data and rebuild any required indexes. Control is immediately returned to the user and the loading will occur via the scheduler.

The **Reset Failed TPT Load** menu option is used when a TPT load has failed which results in a residual database lock on the Load table and a TPT checkpoint file being created.

If the cause of the failure was temporary and has been resolved, then restarting the failed load is the recommended action. This is normally done by restarting the failed job in the scheduler. Refer to *Cleanup after TPT Load Failure* (on page 238) for details.


However, if the cause of the failed load persists (for example, an outage of the source database) then the **Reset Failed TPT Load** function from the Load table's context menu can be used. This option uses TPT functionality to clear the checkpoint file and release the database lock on the Load table.

Note:


- If the Load table is not populated via a TPT load or is loaded via TPT load but no failure was detected, then a message displays that no reset is required when **Reset Failed TPT Load** is selected.
 - For Unix/Linux TPT and Windows TPT Script loads, you need to ensure that the **TPT Job Name** (see "*TPT Load - Source Screen*" on page 235) is a constant value so that the Failed and Reset jobs are using the same TPT job name.
-

The **Check Out** menu option can be used to check out objects for editing to prevent any other users from being able to modify, update or delete any of their associated objects while you are making changes to them. Refer to *Object Check-Outs and Check-Ins* for details.

The **Documentation** menu option can be used to generate (or read if already generated) the WhereScape RED HTML documentation for the selected load table. Two options are available: **Display** and **Create**:




The **Projects** menu option can be used to remove the Load table from the current project or to add the Load table to the current project. If there aren't any projects in the repository, these two options are unavailable.



The **Impact** menu option enables you to produce a number of reports and diagrams:

- **Track Back, Track Forward** or **Dependent Jobs** report
- **Track Back** and **Track Forward** diagrams



The **Code** menu option can be used to view a procedure associated to a table or to regenerate the table's update procedure.

Hover over this option to display an additional menu containing a list of procedures associated with the table:



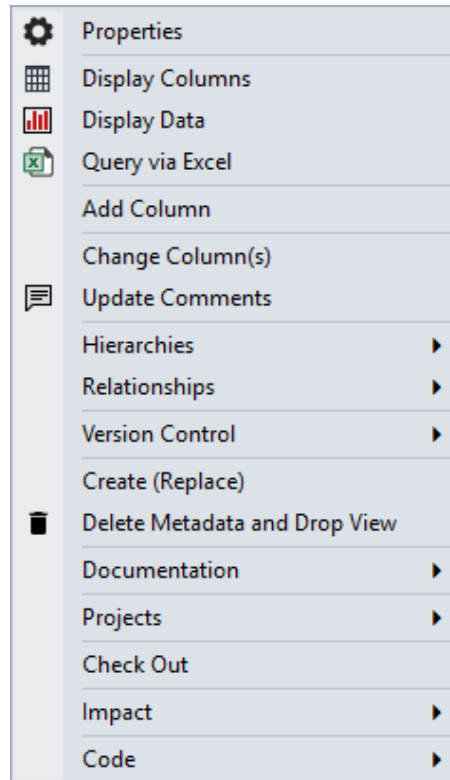
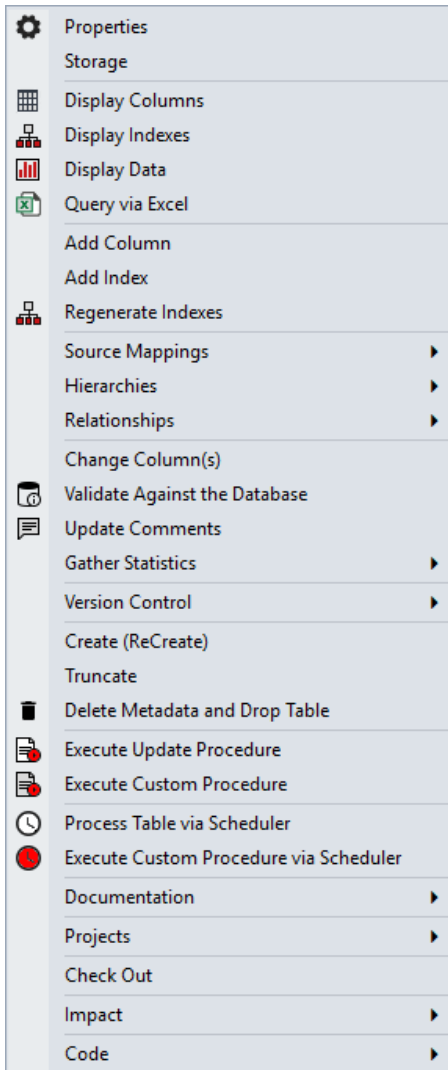
Choose a procedure from the list to open in the procedure editor in view mode or choose to regenerate the update procedure.

Notes:

1. Only Load tables with one or more defined procedures have the **Code view** option.
 2. If the Load table is Script-based load type, either option is displayed in the **Code** sub menu:
 - Generate Load Script (if a script is not yet generated).
 - Regenerate <script name> (if a script has already been generated).
 Refer to **Script based loads** for details.
-

DIMENSION TABLES


The standard pop-up menus for **dimensions** follow (dimension tables on the left, dimension views on the right):



Note: WhereScape RED version 8.1.1.0 and above no longer supports the Dimension View object type when creating new objects. However, users that have existing Dimension View objects in their data warehouse can retain and continue to use them.

The bulk of these menu options are the same as for Load tables and are described under the Load table section above. The differences are:

The **Source Mappings** sub-menu contains two options:




Source mapping objects are created when there is a requirement to have more than one set of source tables to populate the target table and provides for the generation of a update and custom procedure per source mapping object.

Source Mappings can be added using **Add Source Mapping** and listed using **List Source Mappings**.

Refer to the section **Multi Source Processing** (on page 1101) for details.

The **Hierarchies** sub-menu contains the following options:



Hierarchies can be added using **Hierarchies>Add Hierarchy** and listed using **Hierarchies>List Hierarchy**. Hierarchy elements can be listed using **Hierarchies>List Hierarchy Elements**. The **Hierarchies>Copy Hierarchies from Source** feature copies all hierarchies from the source table to the destination table. Source hierarchies are copied to the destination table automatically during table creation, but this feature is useful if the source table has been updated since destination table was created.

The **Relationships** menu options allow the management of enhanced relationships. The **Add Relationship** option opens the Add Relationships dialog, the **List Relationships** option displays a list of enhanced relationships in the Drop Target Pane for the selected object and **Generate Relationships** generates relationships which have not yet been defined in metadata.

The **Execute Update Procedure** menu option executes the procedure defined as the 'update procedure' for the table. The procedure is executed interactively and locks the screen until completed. This menu option is only intended for use when working with small/prototype data volumes, and no index handling is performed.


The **Execute Custom Procedure** menu option executes the procedure defined as the 'custom procedure' for the table. As with Update, the procedure is executed interactively.

The **Execute Custom Procedure via Scheduler** menu option executes the procedure defined as the 'custom procedure' for the table, via the Scheduler.

The **Code** menu option can be used to view a procedure attached to a table or to rebuild the table's update procedure; and to view or rebuild the Get Key function on the Dimension/Dimension View.

Note: The **Get Key** function is not available for dimension objects that are stored in custom database targets.

Hover over this option to display an additional menu containing available options:




Choose a procedure from the list to open in the procedure editor in view mode or choose to rebuild the update procedure.

Note: Only tables with one or more defined procedures have the **Code** option.


DATA STORE TABLES

The standard pop-up menus for **Data Store** tables is:



The menu options available for Data Store Tables is a subset of the options for Dimension Tables, except for an additional option under the code menu.

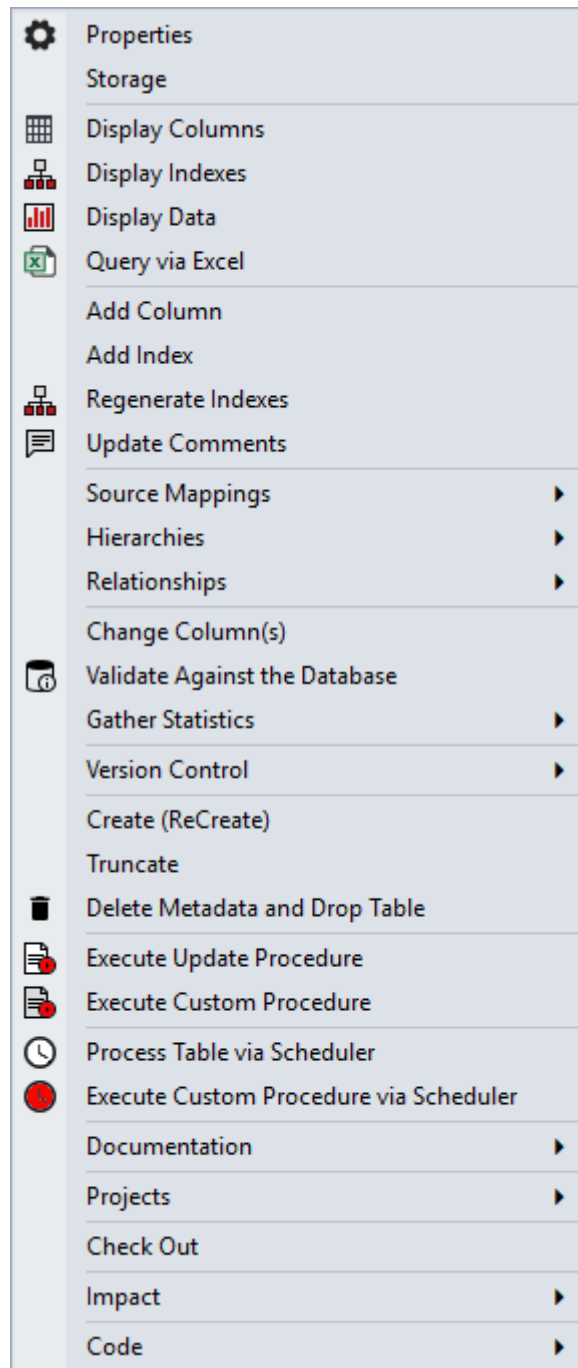
The **Code** menu option can be used to view a procedure attached to a table or to rebuild or regenerate the table's update procedure. Hover over this option to display an additional menu, containing available options:



Choose a procedure from the list to open in the procedure editor in view mode or choose to rebuild or regenerate the update procedure.

EDW 3NF TABLES

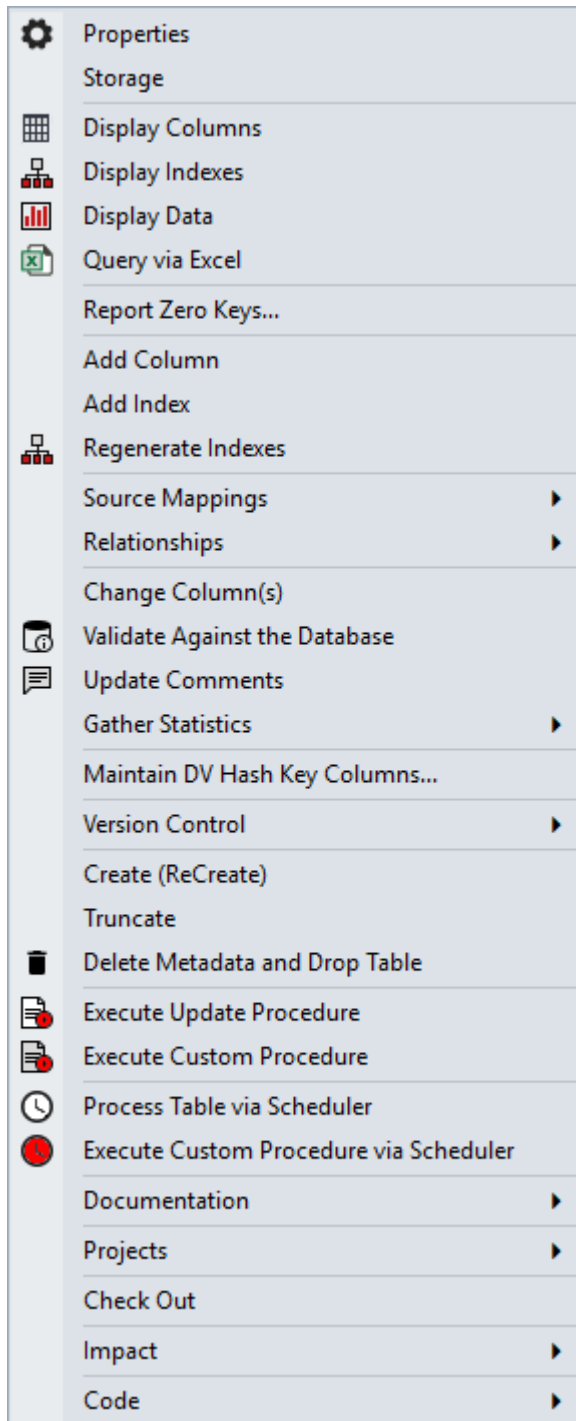
The standard pop-up menu for **EDW 3NF tables** is:



The menu options available for EDW 3NF Tables is a subset of the options for Data Store Tables.

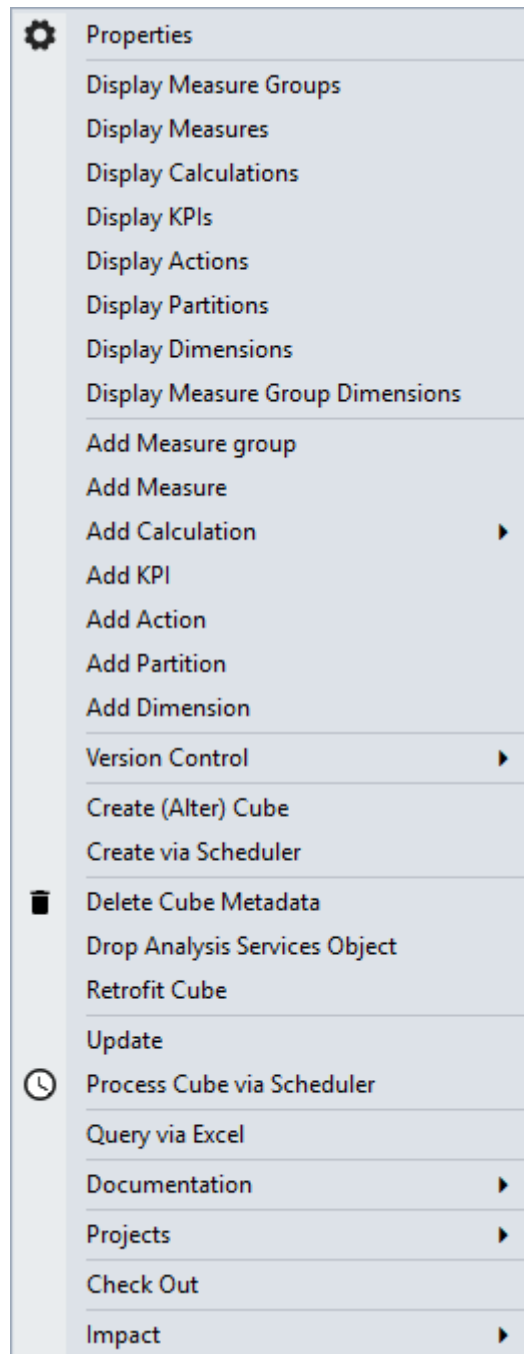
STAGE AND AGGREGATE TABLES

All of these table types have a similar pop-up menu. The standard **stage table** menu is as follows:



OLAP CUBES

The standard pop-up menu for **OLAP Cubes** is:



The **Properties** menu option opens the properties dialog that defines cube creation options and access to documentation tabs.

The **Display Measure Groups** menu option shows the details of the measure groups associated with the cube in the middle pane.

The **Display Measures** menu option lists the measures associated with the measure groups in the cube. This is the default view in the middle pane when a cube is selected in the left pane with a single click.

The **Display Calculations** menu option lists all the calculated members defined in the cube.

The **Display KPIs** menu option lists all the Key Performance Indicators defined in the cube.

The **Display Actions** menu option lists all the actions defined in the cube.

The **Display Partitions** menu option lists all the partitions defined against the related measure groups within the cube.

The **Display Dimensions** menu option lists all of the dimensions defined in the cube.

The **Display Measure Group Dimensions** menu option displays a cross tab report in the middle pane showing cube dimensions relating cube measure groups.

The **Add Measure group** menu option allows a new measure group to be added to the cube.

The **Add Measure** menu option adds another measure to the cube.

The **Add Calculation** menu option adds a new calculated member to the cube.

The **Add KPI** menu option adds a new KPI to the cube.


The **Add Action** menu option adds a new action to the cube.

The **Add Partition** menu option adds a new partition to a measure group in the cube.

The **Add Dimension** menu option adds an existing OLAP Dimension to the cube.

A version of an OLAP cube is a copy of the metadata definition of the cube at the time of the versioning. This version information can be used to create a new OLAP cube, or can simply be left as a backup and reference point. Use **Version Control, New Version** to version an OLAP cube.

The **Build Application** option allows you to build an application file for the OLAP cube and the **Duplicate Object** option allows you to create a new OLAP cube as a duplicate of this OLAP cube.



The **Create (Alter) Cube** menu option creates the cube and supporting objects in Analysis Services (including cube database, data source view (DSV) and dimensions) based on the definition in WhereScape RED.

The **Create via Scheduler** menu option submits a create of the OLAP cube to the scheduler.

The **Delete Cube Metadata** menu option deletes the cube definition from WhereScape RED.

The **Drop Analysis Services Object** menu option drops the selected object in Analysis Services.

The **Retrofit Cube** menu option retrofits the OLAP cube from Analysis Services.


The **Update** menu option processes the cube in Analysis Services interactively from the WhereScape RED interface.

The **Process Cube via Scheduler** menu option generates a WhereScape RED scheduler job to process the cube in Analysis Services.

The **Query via Excel** menu option opens up an .oqy file in Microsoft Excel.


Note: Due to a shortcoming in the Microsoft Office installation, it may be necessary to associate the .oqy file extension with Microsoft Excel before this option will succeed.

The **Documentation** menu option can be used to generate (or read if already generated) the WhereScape RED HTML documentation for the selected OLAP cube. Two options are available: **Display** and **Create**:




The **Projects** menu option can be used to remove the OLAP cube from the current project or to add the OLAP cube to the current project. The **List Projects** option displays a list of projects which contain the current object, results are shown in the bottom pane. Multiple objects can be selected by double-clicking the **OLAP Cube** icon in the left pane and Ctrl + clicking multiple connections in the middle pane.

If there aren't any projects in the repository, these options are unavailable.

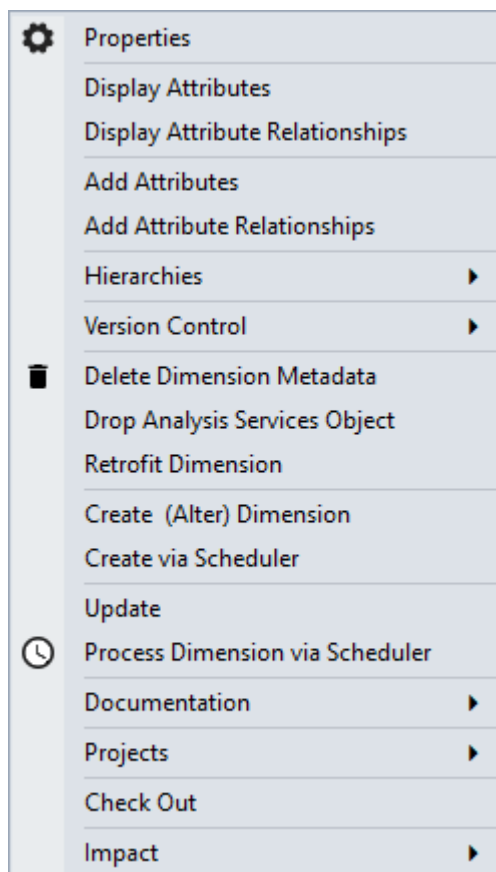


The **Impact** menu option enables you to produce a **Track Back** diagram on the OLAP cube.



OLAP DIMENSIONS

The standard pop-up menu for an **OLAP Dimension** is:



The **Properties** menu option displays the OLAP dimension properties dialog which includes documentation tabs.

The **Display attributes** menu option lists the attributes for the selected dimension. This is the default view when an OLAP Dimension is selected in the left pane.

The **Display Attribute relationships** menu option shows the relationships between the dimensional attributes.

The **Display hierarchies** menu option lists the hierarchies associated with the selected dimension.

The **Display hierarchy levels** menu option lists all the levels for all hierarchies for that dimension.


The **Add Attributes** menu option adds a new attribute to the dimension.

The **Add Attribute relationships** menu option adds an attribute relationship for the selected dimension.

The **Add hierarchy** menu option adds a new hierarchy to the dimension.

The **Add hierarchy level** menu option adds a level to a hierarchy.

A version of an OLAP Dimension is a copy of the metadata definition of the OLAP Dimension at the time of the versioning. This version information can be used to create a new OLAP Dimension, or can simply be left as a backup and reference point. Use **Version Control, New Version** to version an OLAP Dimension. The **Build Application** option allows you to build an application file for the OLAP Dimension and the **Duplicate Object** option allows you to create a new OLAP Dimension as a duplicate of this OLAP Dimension.



The **Delete Dimension Metadata** menu option deletes the cube definition from WhereScape RED metadata.

The **Drop Analysis Services object** menu option provides the ability to drop the selected object from Analysis Services.

The **Retrofit Dimension** menu option retrofits the OLAP Dimension from Analysis Services.


The **Create (Alter) Dimension** menu option creates the OLAP Dimension and supporting objects in Analysis Services (including cube database and dsv) based on the definition in WhereScape RED. This option requires connection and cube database information populated in the OLAP Dimension properties.

The **Create via Scheduler** menu option submits a create of the OLAP Dimension to the scheduler.

The **Update** menu option processes the OLAP Dimension in Analysis Services interactively from the WhereScape RED interface. This option requires connection and cube database information populated in the OLAP Dimension properties.


The **Process Dimension via Scheduler** menu option generates a WhereScape RED scheduler job to process the OLAP Dimension in Analysis Services. This option requires connection and cube database information populated in the OLAP Dimension properties.

The **Documentation** menu option can be used to generate (or read if already generated) the WhereScape RED HTML documentation for the selected OLAP dimension. Two options are available: **Display** and **Create**:

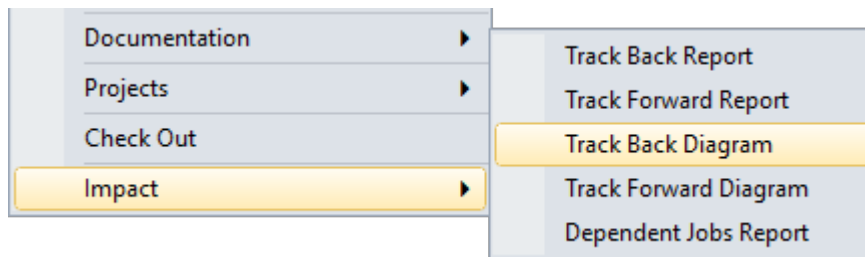


The **Projects** menu option can be used to remove the OLAP dimension from the current project or to add the OLAP dimension to the current project. The **List Projects** option displays a list of projects which contain the current object, results are shown in the bottom pane. Multiple objects can be selected by double-clicking the **OLAP Dimension** icon in the left pane and Ctrl + clicking multiple connections in the middle pane.

If there aren't any projects in the repository, these options are unavailable.

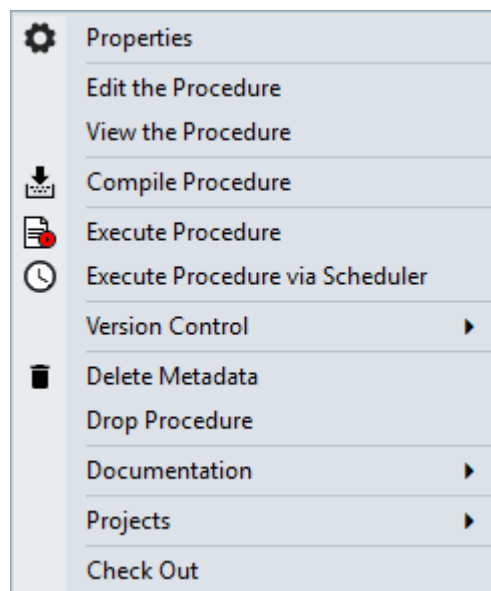


The **Impact** menu option enables you to produce a **Track Back** diagram on the OLAP dimension.



PROCEDURES

Procedures are commonly auto built through the Properties screen of one of the table types. They can also be created manually. Once created they can be edited, compiled, etc. The menu displayed when right-clicking a procedure name is as follows:



If a procedure has been locked as the result of the WhereScape RED utility being killed or failing or a database failure, then it can be unlocked via the properties screen associated with the procedure.

The **Edit the Procedure** menu option invokes the procedure editor and loads the procedure. A procedure can be compiled and executed within the procedure editor.


The **View the Procedure** menu option displays a read only copy of the procedure. If the procedure is locked by another user, then viewing the procedure is the only option available.

The **Compile Procedure** menu option compiles the procedure from the metadata.

The **Execute Procedure** menu option executes the procedure and displays the results in the results window.

The **Execute Procedure via Scheduler** menu option sets up a job to execute the procedure via the scheduler.

A version of a Procedure can be created at any time via the **Version Control, New Version** menu option. The various versions of the procedure can be viewed from within procedure editor, or a new procedure can be created from the version. The **Build Application** option allows you to build an application file for the Procedure and the **Duplicate Object** option allows you to create a Procedure as a duplicate of this Procedure.




Selecting the **Delete Metadata** menu option deletes the procedure from the meta data. It then asks if the procedure should also be dropped from the database.

Selecting the **Drop Procedure** menu option drops the procedure from the meta data and from the database as well.


Note: Procedures cannot be dropped or deleted if they are protected by, for example, an **Edit Lock**.

The **Documentation** menu option can be used to generate (or view if already generated) the WhereScape RED HTML documentation for the selected procedure. Two options are available: **Display** and **Create**:



The **Projects** menu option can be used to remove the procedure from the current project or to add the procedure to the current project. The **List Projects** option displays a list of projects which contain the current object, results are shown in the bottom pane. Multiple objects can be selected by double-clicking the **Procedure** icon in the left pane and Ctrl + clicking multiple connections in the middle pane.

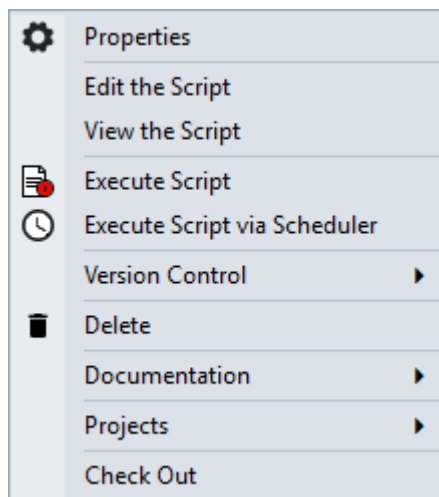
If there aren't any projects in the repository, these options are unavailable.



SCRIPTS

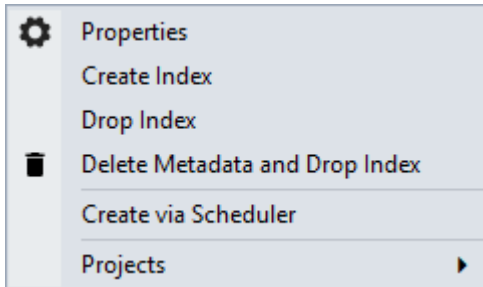
Scripts are very similar in their operations to procedures. Some key differences are as follows:

- If a script is deleted, it is only removed from the metadata. However, because WhereScape RED never stores the script on the host system, this removes the script permanently.
- Unix scripts, Windows CMD scripts and Windows **PowerShell scripts** (see "**Windows PowerShell Scripts**" on page 653) are available on Load objects.
- Windows CMD scripts and Windows PowerShell scripts are available on all table object types, where the table is stored in a custom database target and when using a Windows scheduler.
- The table object needs to specify the connection appropriate for the script (typically Windows or Unix).



INDEXES

Indexes are always associated with a table. To define a new index, the menu option associated with the table that is to have the index must be used. Once defined, the following operations can be performed.



The **Properties** menu option displays the **Properties** screen, which contains the entire definition of the index, including the columns in use, storage parameters and index type, etc. The way the scheduler handles the index is also defined in the Properties. An index can be set so that it will be dropped by the scheduler prior to a table update and then rebuilt by the scheduler once the update has been completed. It can also be defined for rebuild on certain days.

The **Storage** menu option displays the storage properties screen, containing storage options and parameters.

The **Create Index** menu option creates the index in the database. This may take some time for large indexes, and in such cases, it would be better to schedule a create of the index. See the chapter on the Scheduler, if such an activity is required. This menu option is intended for use when working with prototype data volumes.


The **Drop Index** menu option drops the index in the database.

The **Delete Metadata and Drop index** removes the metadata definition of the index and drops it from the database. No recovery is possible once this option is actioned.

The **Create via Scheduler** menu option submits a create of the index to the scheduler.

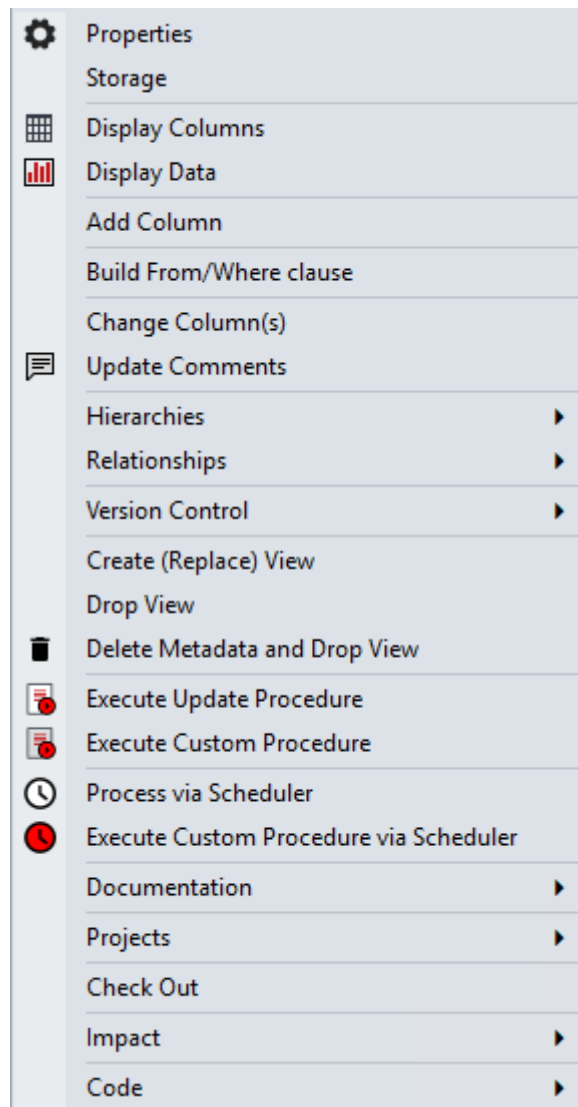
The **Projects** menu option can be used to remove the index from the current project or to add the index to the current project. The **List Projects** option displays a list of projects which contain the current object, results are shown in the bottom pane. Multiple objects can be selected by double-clicking the **Index** icon in the left pane and Ctrl + clicking multiple connections in the middle pane.

If there aren't any projects in the repository, these options are unavailable.



VIEWS

Views are primarily built to alias column names or add locking clauses to table objects (Load, Stage, Dimension and Aggregate Tables). The menu below shows the operations that can be performed on views.



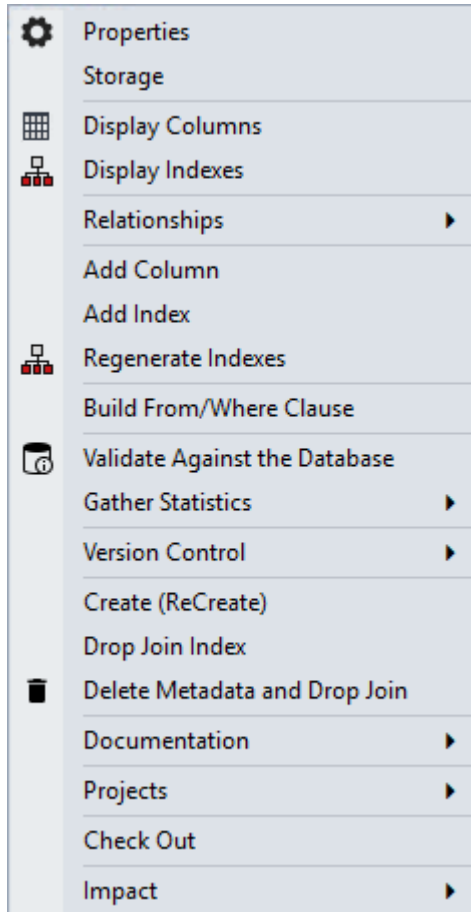
The menu options available for views are almost a subset of those for stage tables, described under the stage section above.

The additional menu option is **Build From/Where clause**. This option provides a way to invoke a wizard to define a join between two or more tables.

Display Indexes, Query table via Excel, Add Index, Regenerate Indexes are not available for views.

JOIN INDEXES

Join indexes are an indexing structure containing columns from one or more base tables. They are used to summarize data, pre-perform table joins and re-hash Teradata storage, by providing alternate primary indexing. The standard join index menu is as follows.



The bulk of these menu options are the same as for views and are described under the view section above. The differences are:

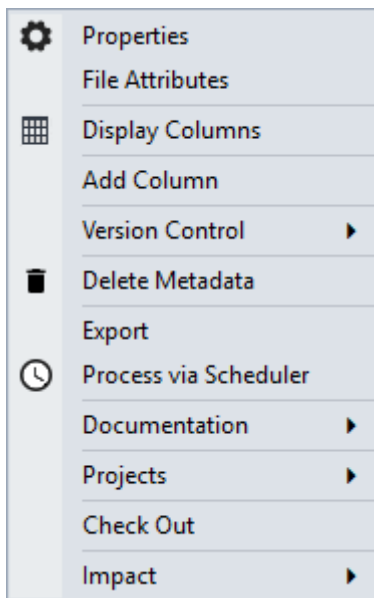
A join index may have an index on it, so the **Add Index** option is provided.

Join indexes are automatically maintained by Teradata, so do not have the **Execute Update Procedure**, **Execute Custom Procedure** or **Process via Scheduler** options.

EXPORTS

Export objects do not exist as an object in their own right in the Teradata database. They are created from a single model, view or aggregate object and are used to generate file exports for a downstream system. Export files are built using Teradata tools, such as FastExport and Teradata Parallel Transporter.

The pop-up menu displayed when a right-click is used on an export is as follows:



File Attributes launches the properties window for the export table, focusing on the File Attributes tab within this window.

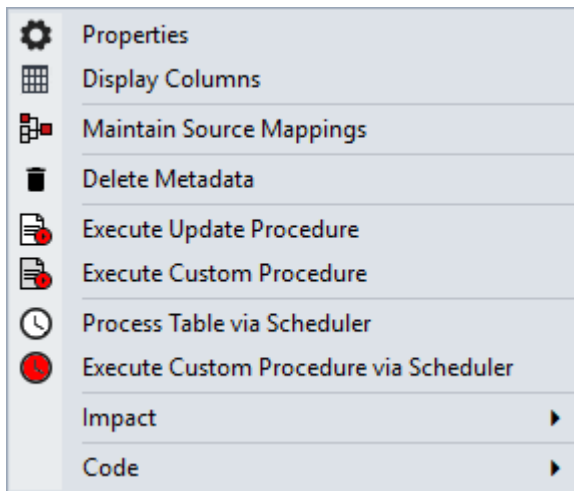
All the other options for export objects are described above under other object types.

SOURCE MAPPING

Source Mapping objects are used in WhereScape RED to define source column mappings when there is a requirement for the target table to be populated by more than one set of source tables. These objects support the generation of an update procedure for the target table, sourced from each set of the source tables.

Refer to the section **Multi Source Processing** (on page 1101) for details.

The context menu displayed when you right-click a source mapping object is as follows:




Maintain Source Mappings launches the Source Mapping tool which enables you to manage the column mappings of the selected target table.

The **Execute Update Procedure via Scheduler** menu option sets up a WhereScape RED scheduler job to execute the source mapping object's update procedure.

The **Execute Custom Procedure via Scheduler** menu option executes the procedure defined as the 'custom procedure' for the table via the the WhereScape RED scheduler.

The **Impact** sub-menu option enables you to produce the following:

- **Track Back Report** – Lists the source tables of the selected source mapping object.
- **Track Forward Report** – Lists the parent table object of the selected source mapping object.
- **Track Back Diagram** – Displays the source tables of the selected source mapping object.
- **Track Forward Diagram** – Displays the parent table object of the selected source mapping object.
- **Dependent Jobs Report** – Lists jobs that includes the currently selected source mapping object or the parent table object.



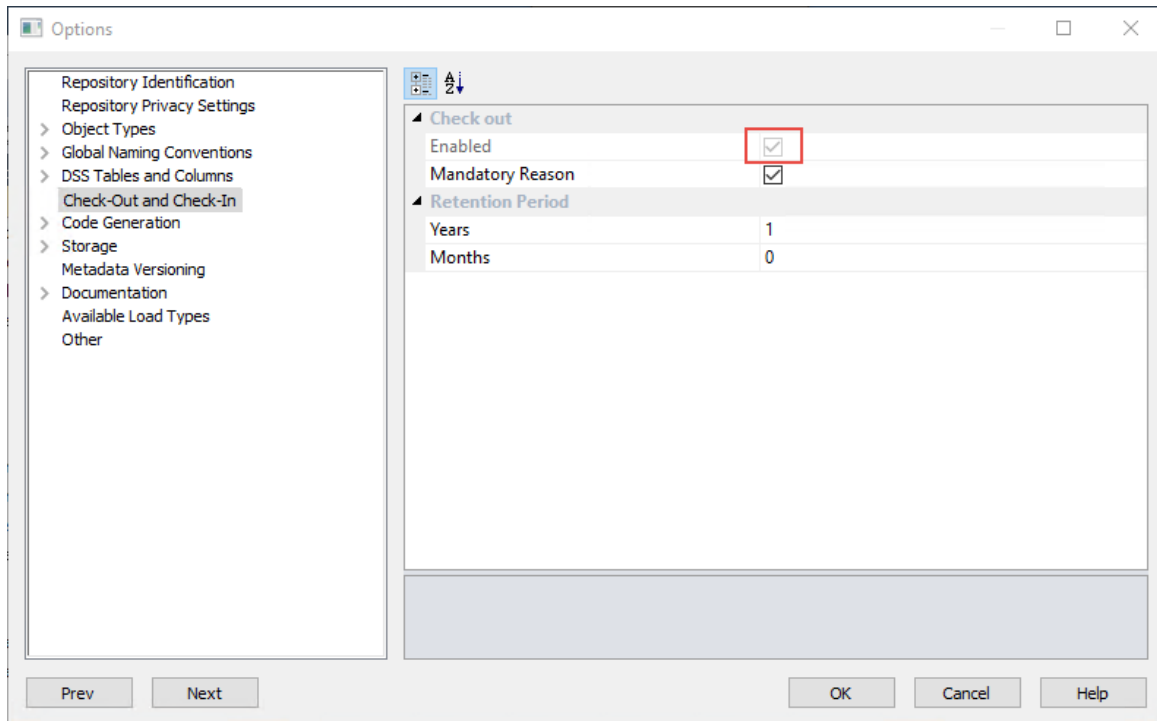
All the other options for source mapping objects are described above, under other object types.

OBJECT CHECK-OUTS AND CHECK-INS

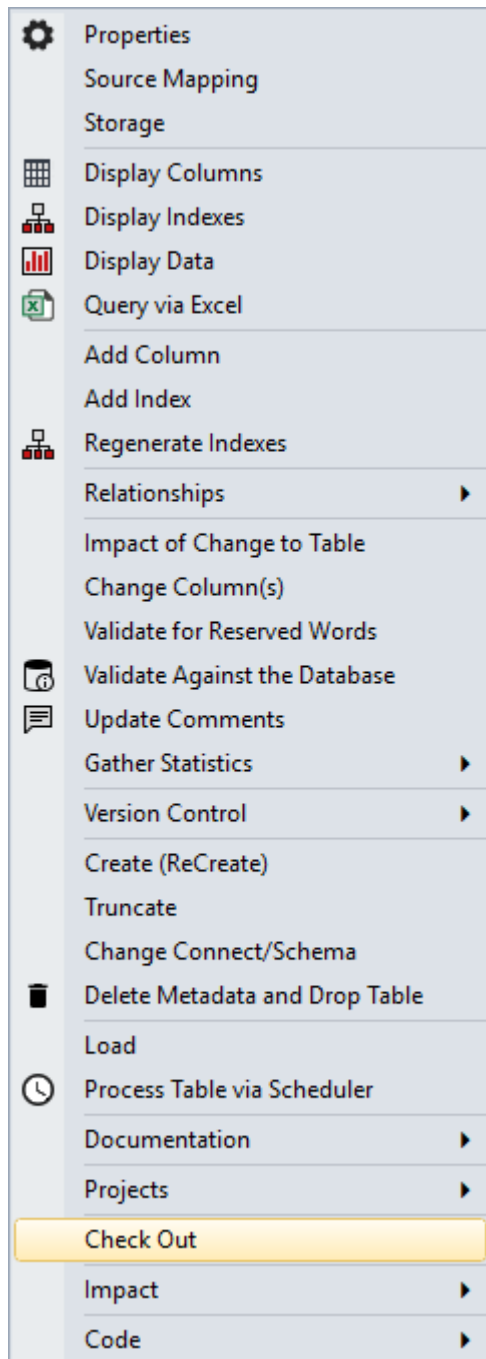
Objects in RED can be checked out for editing to prevent any other users from being able to modify, update or delete any of their associated objects while you are making changes to them.

This functionality is configured through the **Tools>Options>Check-Out and Check-In** setting.

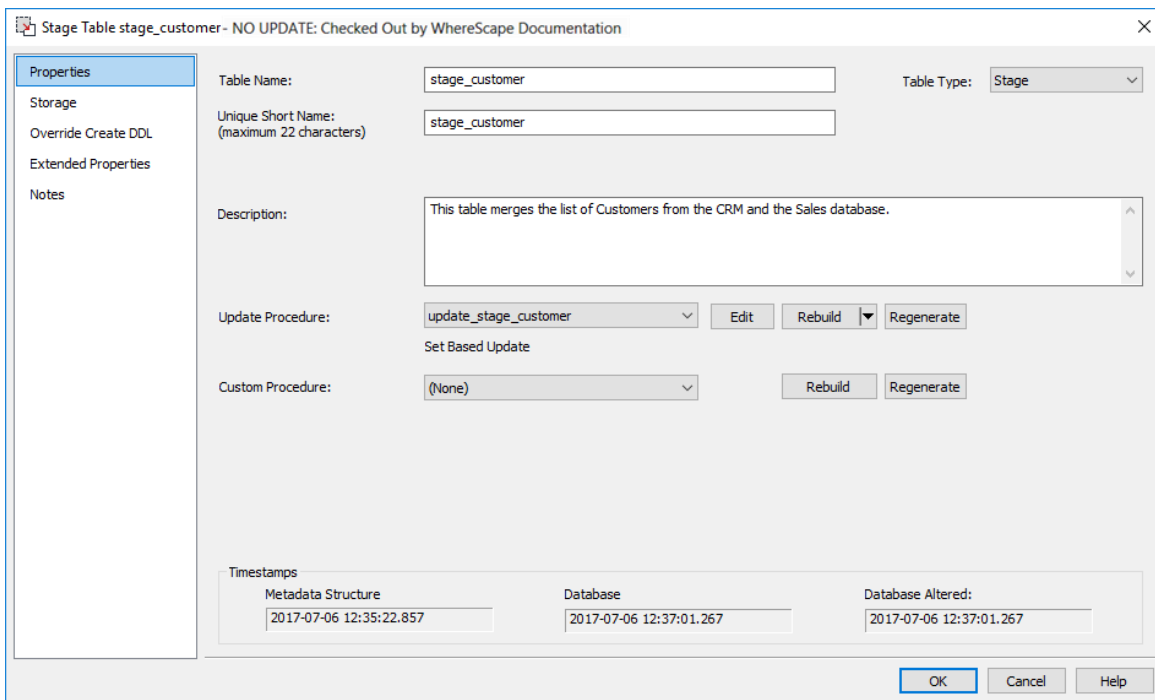
The **Check-Out** option is enabled by default.



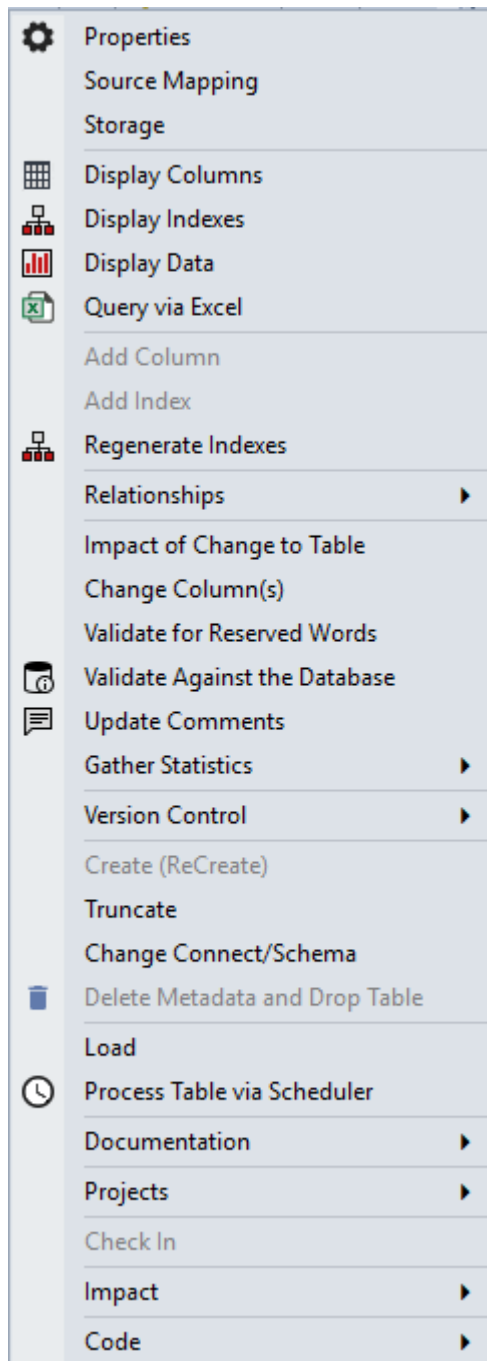
This setting enables users to check-out/check-in objects in WhereScape RED through the object's right-click context menus.



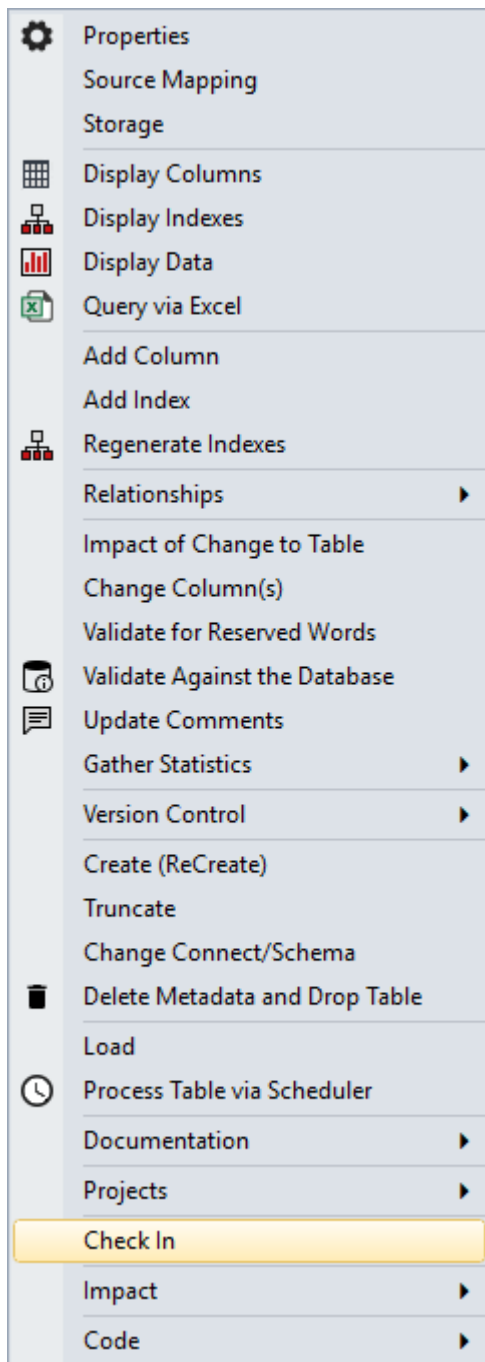
Checked out objects and their associated procedures and scripts cannot be modified, updated or deleted by other users. If a different user tries to access the properties screen or procedure windows for those objects, their fields become unavailable and headed NO UPDATE: Checked Out by (user name).



Other users logging in to the same metadata repository will have some of the checked-out object's context menu options unavailable, as shown below.

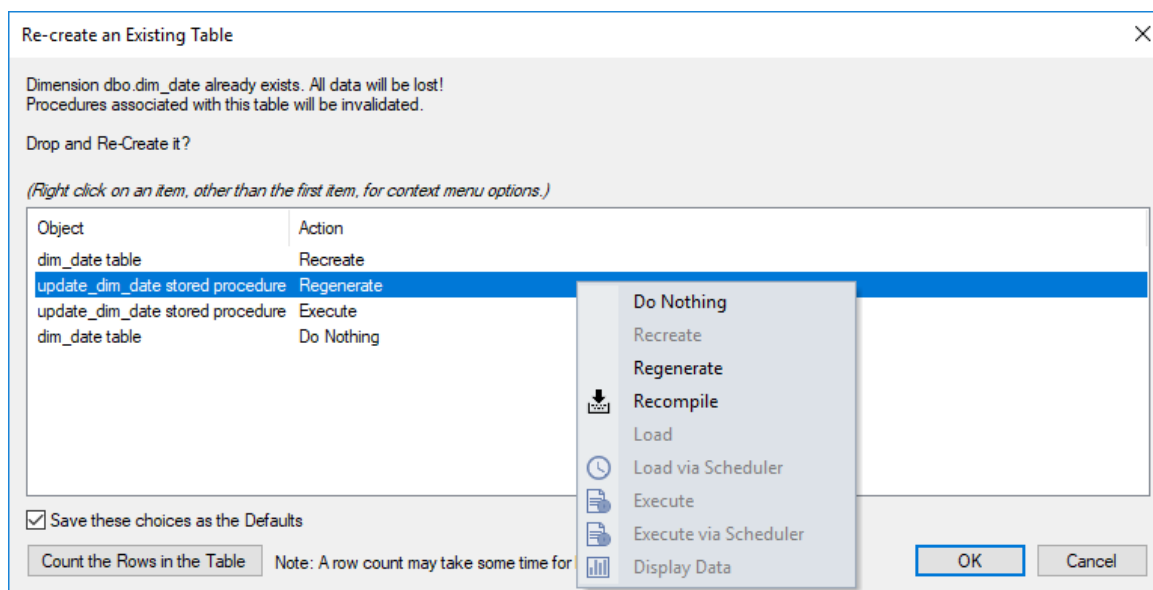


To check-in a checked-out object, simply right-click it and select **Check In** from the context menu.



RE-CREATE DIALOG

The Re-create window provides available options for dropping and recreating an existing table and associated procedures. Objects and actions available depend on the table being recreated.



Actions in the table are completed from top to bottom; the action performed on each object can be set as desired:

Object(s)	Action	Description
Table	Recreate	Recreate the table structure only. Do not make changes to associated procedures. This action is mandatory.
Artificial key sequence (Oracle only)	Do Nothing	Do not recreate the key sequence. Skip to the next action.
	Recreate (Default)	Recreate the artificial key sequence. Do not make changes to associated procedures.
Update procedure	Do Nothing	Do not regenerate the update procedure, skip to the next action.
	Regenerate (Default)	Recreate the table structure, regenerate the procedure definition in metadata based on the current table and column properties, and recreate the procedure in the database. This is the option to use if there have been changes to column or table properties, e.g. column transformations.
	Recompile	Recreate the table structure and recreate the update procedure based on the existing procedure definition in the metadata.

Object(s)	Action	Description
Load tables	Do Nothing (Default)	Do nothing for this object, skip to the next action.
	Load	Perform an interactive load of the data into the table. WhereScape recommends performing this via the Scheduler for large tables.
	Load via Scheduler	Add a data load task to the scheduler. Useful for large tables where processing may take some time.
Load script	Regenerate	Regenerate the table's load script. This option is only available for Load tables with a load type of script and with the source connection setting defined. If the Load table does not have an existing script then a new script is created.
Update procedure	Do Nothing (Default)	Do nothing for this object, skip to the next action.
	Execute	Execute the update procedure. WhereScape recommends performing this via the Scheduler for large tables.
	Execute via Scheduler	Execute the update procedure via the scheduler. Useful for large tables where processing may take some time.
Tables	Do Nothing	Do not display the data.
	Display Data	Display the table data after recreating the table and update procedure (if applicable). This is the default setting for Views and Dimension Views. This is the default setting for tables when an update procedure is present and has been executed in the previous action. Unavailable when an update procedure is present but has not been executed in the Recreate window.

Save these choices as the Defaults

This option can be used to store the current actions as the default choices for all object types. The Recreate window is then automatically populated with the saved actions each time it appears.

Count the Rows in the Table

Displays the total number of rows in the table at the bottom of the Recreate window. Knowing the number of rows in a table can help estimate how long Load and Execute actions may take.

ORGANIZING OBJECTS

As mentioned in the previous section there are many types of object in WhereScape RED.

The objects in the metadata repository are displayed in the left pane of the Builder window. They are displayed in a tree structure which can be expanded and closed as required. The tree can be refreshed by using the F5 or Ctrl/R key.

Object Groups

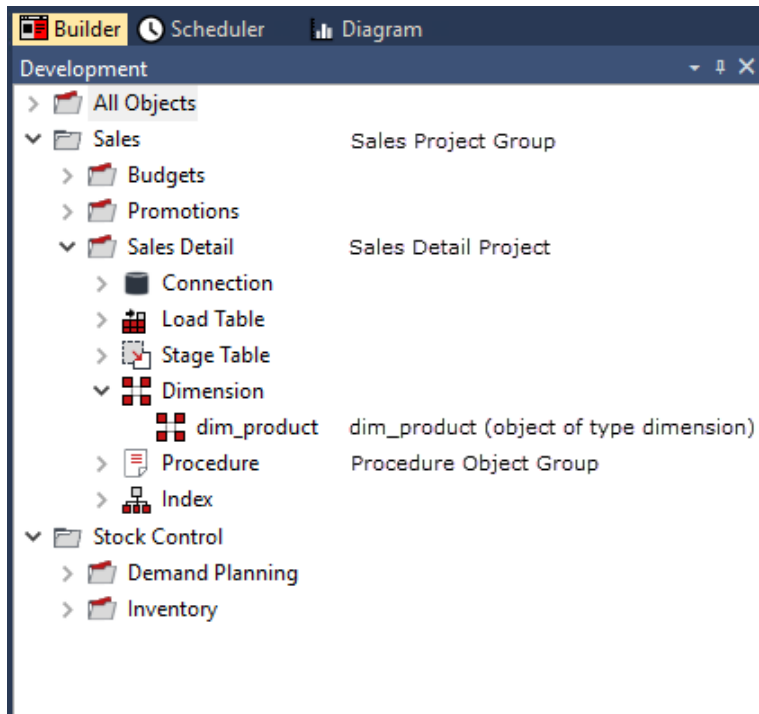
The objects created in WhereScape RED are grouped together into object groups based on object type. For example, we store all the dimension objects in the Dimension **object group**. Optionally, we can choose to display dimension tables as the **Dimension** object group and dimension views as the **Dimension View** object group.

Projects

These object groups are in turn stored within **projects**. When WhereScape RED is first started, the special project called **All Objects** is the only project. This project will always contain all the objects that exist in the metadata repository. Additional projects can be created if desired. These additional projects can hold some or all of the objects as seen in **All Objects**. An object as such, only exists once in the metadata. Therefore, if we have a dimension object called dim_product, there is and can only be one copy of the object dim_product within the metadata.

Projects are used to hold a group of objects that relate back to a similar module or analysis area of the data warehouse. In the example below, we have additional projects called Budgets, Promotions, Sales, Demand Planning and Inventory. In this way projects enable us to restrict the amount of information (objects) we need to deal with to just those relevant to the area being worked on.

The example below shows a meta repository, using two project groups (Sales and Stock Control) and five additional projects (Budgets, Promotions, Sales, Demand Planning and Inventory).



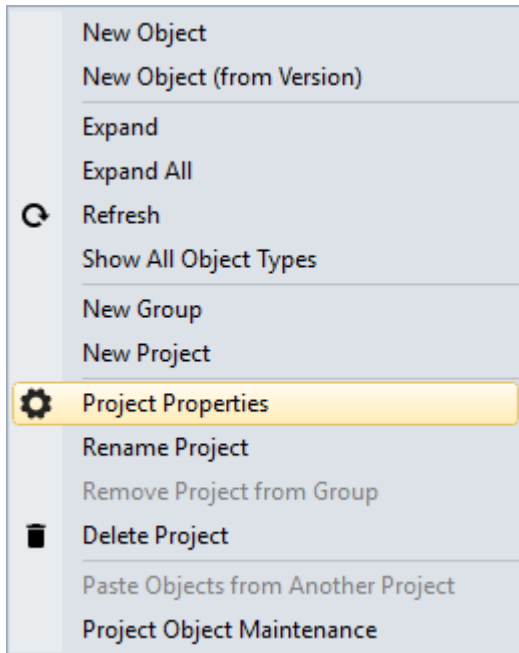
Note: If you delete an object using the right-click menu, then the object will be deleted from the metadata and will be removed from all Projects. Remove the object from the project instead of deleting it.

It is important to understand that these projects are only a means of visualizing the objects. Even though an object may appear in many projects, it only exists *once* in the metadata.


To create a project, right-click in the left pane below the last project and select **New Project**.

The **File>New Project** menu option may also be used. Projects can be renamed, removed from the project group or deleted by using the right-click menu when positioned on the project name. Deleting a project does not delete the objects from the metadata; it simply removes their reference in the project being deleted.

To view or make changes to the Project's Properties; right-click the project name and select **Project Properties**.



The following screen shows four check boxes.



Include in User Documentation - When selected, the project is included in the User Documentation. The default state is selected.

Include in Tech Documentation - When selected, the project is included in the Technical Documentation. The default state is checked.

Local Project - When selected as local, the project is not included in the Application files. The default state is unchecked.

Show Unused Object Types in Object Browser - When selected, all object types are displayed in the Object Browser Pane. When unchecked, only object types currently in use are displayed. The default state is unchecked. Refreshing the Object Browser Pane is required after changing this setting.

Project Groups

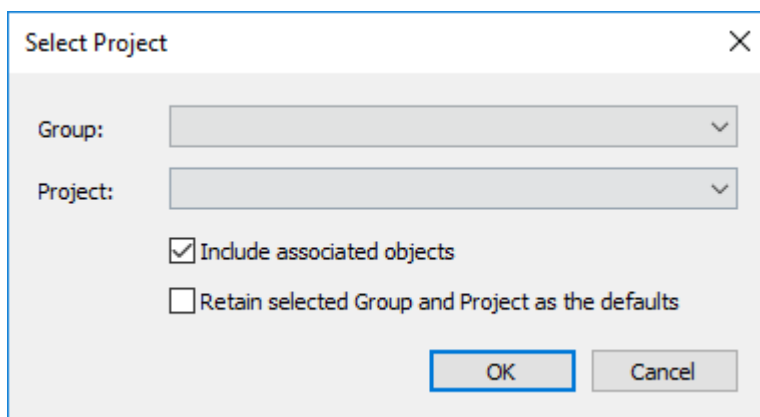
Projects can in turn be grouped together into Project Groups. A project must only appear in one Project Group. To create a Project Group, use the **File>New Group** menu option. Project Groups can be removed by using the right-click menu when positioned on a group name.

ADDING OBJECTS TO PROJECTS

There are several different ways to add objects to a project:

- Click an object in the left pane, press **Ctrl+C**, click the target project (either the project folder or any object group or object within it) and press **Ctrl+V**
- **Drag** the object to the required project
- Right-click on an object in the left pane and select **Projects>Add to Project**
- Highlight a number of objects in the middle pane, right-click and select **Projects>Add to Project**
- *Use the Project Object Maintenance Facility* (see "*Using Project/Object Maintenance*" on page 57)

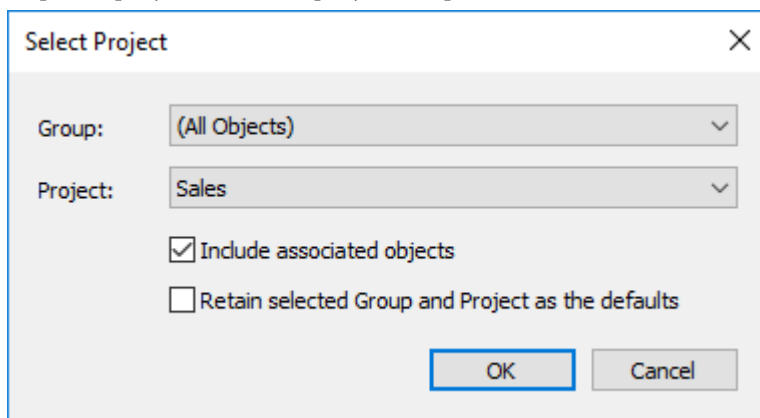
Both options above that use the right-click **Projects>Add to Project** menu, displays the following dialog box:



The dialog box is titled "Select Project" and features a close button (X) in the top right corner. It contains two dropdown menus: "Group:" and "Project:". Below these are two checkboxes: "Include associated objects" (checked) and "Retain selected Group and Project as the defaults" (unchecked). At the bottom, there are "OK" and "Cancel" buttons.

To add an object to an independent project (a project that is not in a group)

- 1 Right-click the object from the **Objects** list pane and then select **Projects > Add to Project**.
- 2 Select the required project from the **Project** drop-down list and then click **OK**. Choose the required project from the project drop-down list and click **OK**.



Select Project

Group: (All Objects) ▾

Project: Sales ▾

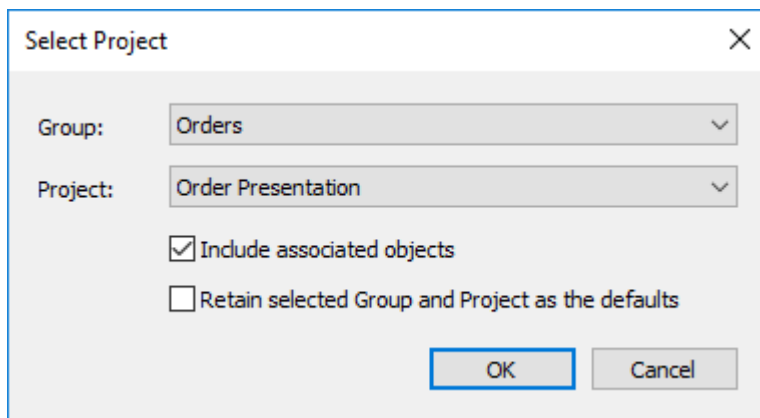
Include associated objects

Retain selected Group and Project as the defaults

OK Cancel

To add an object to a dependent project

- 1 Select the required group from the **Group** drop-down list.
- 2 Select the required project from the **Project** drop-down list and then click **OK**.



Select Project

Group: Orders ▾

Project: Order Presentation ▾

Include associated objects

Retain selected Group and Project as the defaults

OK Cancel

Notes:

- Selecting the **Include associated objects** check box, also adds any indexes, procedures and scripts to the selected project.
- Selecting the **Retain selected Group and Project as the defaults** check box enables you to retain the last values defined in the **Group** and **Project** fields and use them by default when an object is added to a Project or Group.

REMOVING OBJECTS FROM PROJECTS

To remove objects from a project

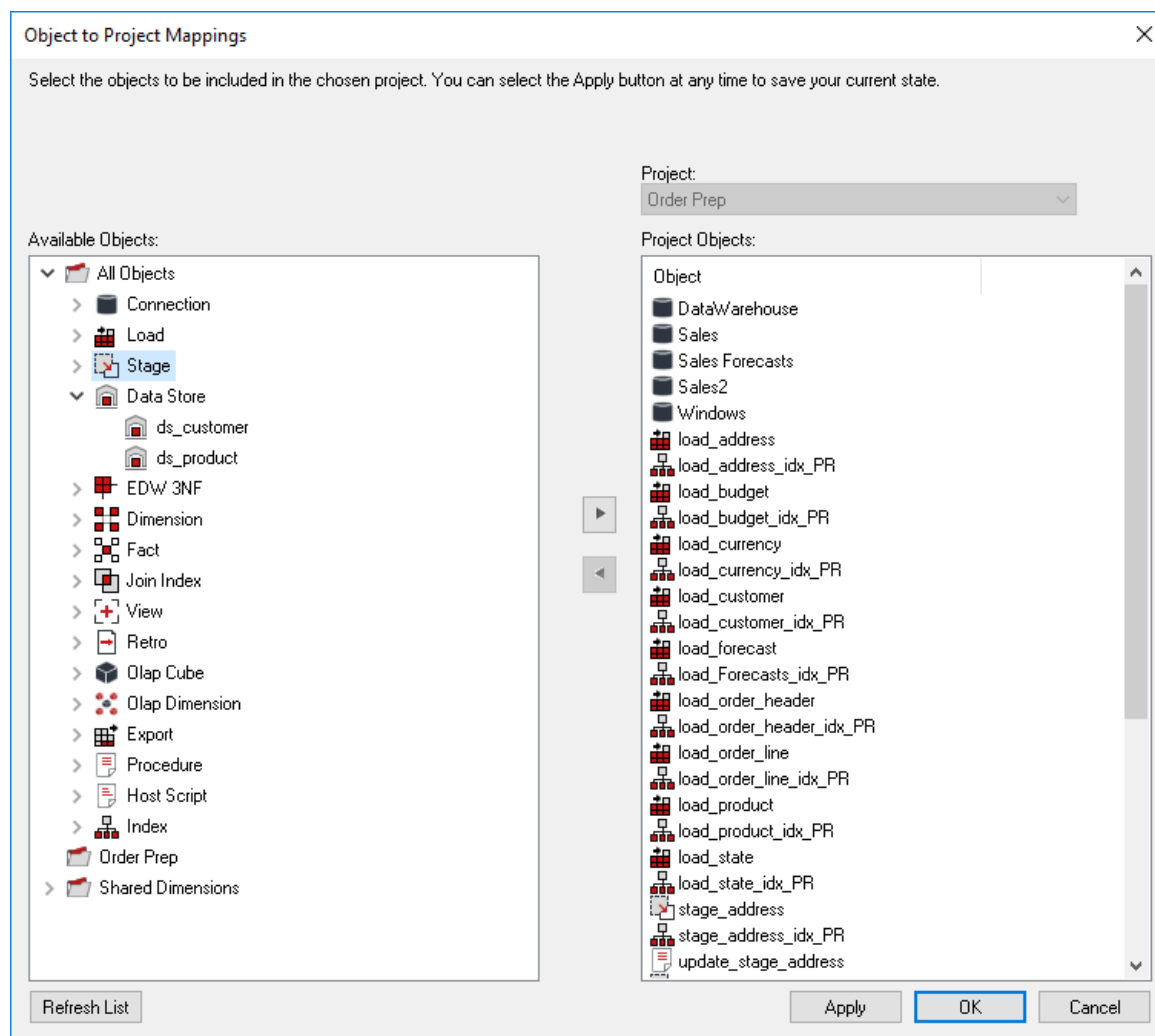
There are different ways to remove objects from a project:

- **Drag** the object to the blank area at the bottom of the pane
- **Drag** the object into the middle pane
- **Drag** the object to the **All Objects** project
- Right-click on an object in the left pane and select **Projects>Remove from Project**
- Highlight a number of objects in the middle pane, right-click and select **Projects>Remove from Project**
- *Use the Project Object Maintenance Facility* (see "*Using Project/Object Maintenance*" on page 57)

USING PROJECT/OBJECT MAINTENANCE

The Project Object Maintenance Facility is invoked by right-clicking a project and choosing **Project Object Maintenance**.

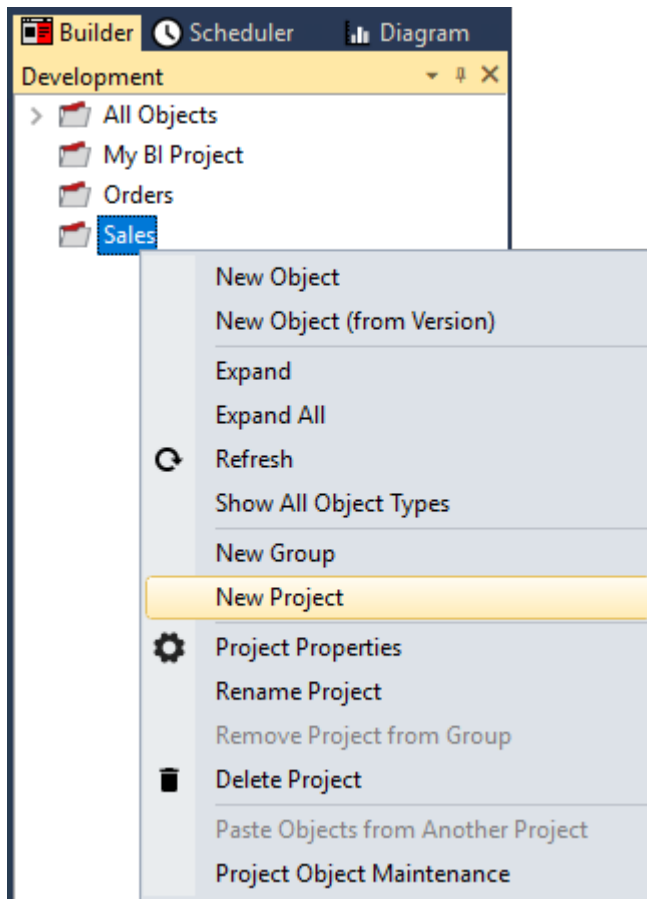
The following window is displayed:



- To **add** objects to a project, move objects from the left to the right using the > button. By default, associated object (procedures, scripts and indexes) are also moved to the left. These can be manually removed if not required.
- To **remove** objects from a project, move objects from the right to the left using the < button.
- When done, click **OK**.
- Clicking **Apply** updates object/project relationships without having to exit the Project Object Maintenance Facility.
- Clicking **Refresh** refreshes the object tree in the left pane.

ADDING PROJECTS TO GROUPS

The best way to move a project into a group is to first create the project in the group. This is done by right-clicking on the group and selecting **New Project**:



The other option is to move an existing project from another group into this group. See **Moving Projects within Groups** (on page 59)

REMOVING PROJECTS FROM GROUPS

A project can be removed from a group by:

- **Dragging** the project to the blank area at the bottom of the pane
- **Dragging** the project into the middle pane
- **Dragging** the project to the **All Objects** project
- **Removing** the project using right-click **Remove Project from Group**
- **Deleting** the project using right-click **Delete Project**

Note: The first four methods move the project out from the group to become an independent project. The last option removes the project from the metadata repository.

MOVING PROJECTS WITHIN GROUPS

Note: An object can be in any number of projects, but a project can only be in one Group.


Performing a drag from one group to another group will simply create an additional project->group mapping.

To move a project from one group to another group:

- 1 **'Copy'** the project by dragging the project from the one group to the other group.
- 2 **Remove** the project from the original group by right-clicking and selecting **Remove Project from Group**.

LIST PROJECTS MEMBERSHIPS FOR AN OBJECT(S)

To list which projects contain a specific object, right click on the object in the left pane and select **Project>List Projects**. Results are shown in the bottom pane.



Multiple objects of one type can be selected by double-clicking the **Object Type** icon in the left pane (e.g. **Connection**, **Load Table**, etc.) and Ctrl + clicking multiple connections in the middle pane.

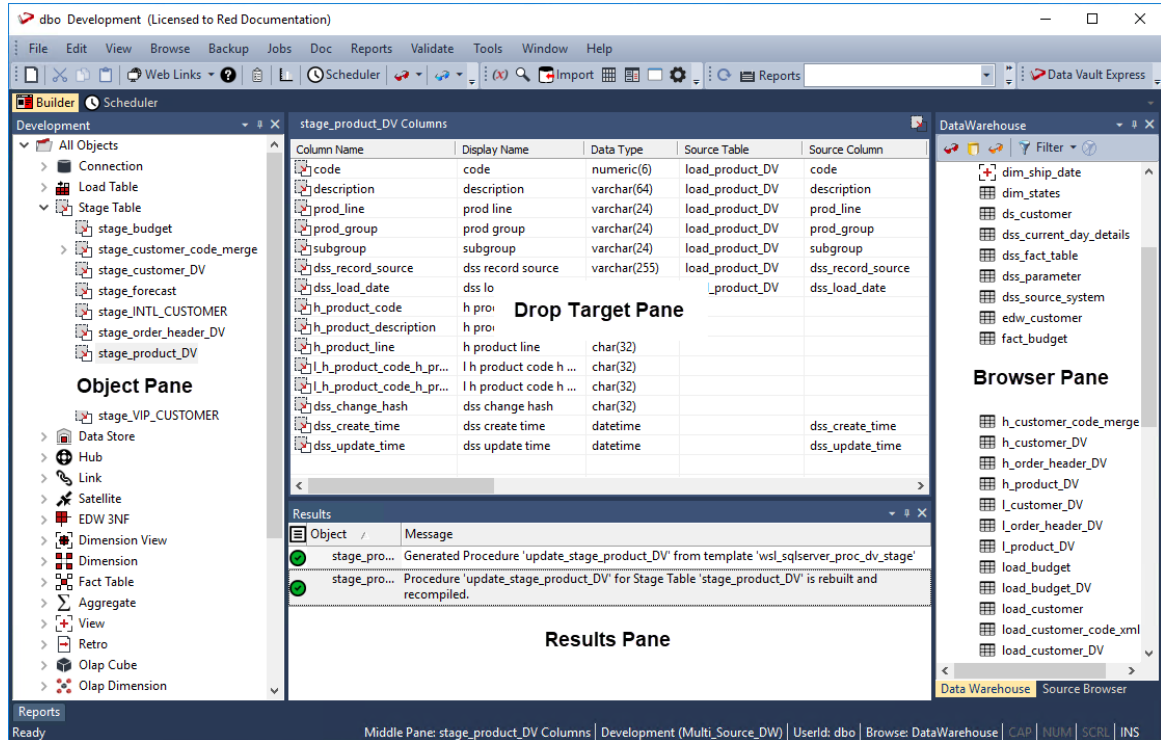
WINDOWS AND PANES

WhereScape RED has a number of different windows that are utilized in the building and maintenance of a data warehouse. Each window may in some cases be broken into panes. There are four main windows that are used extensively in the building of a data warehouse:

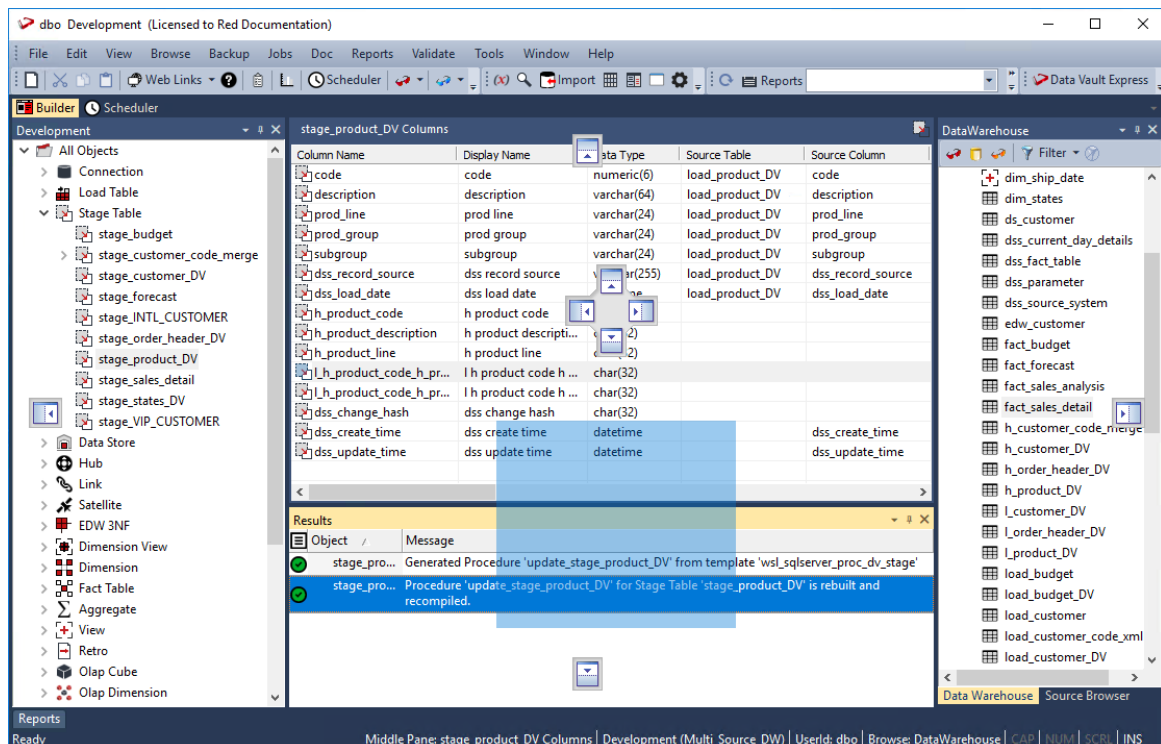
- The **Builder** window
- The **Scheduler** window
- The **Diagram** window
- The **Procedure Editor** window

BUILDER WINDOW

The **Builder Window** has four panes.



The left, bottom and right panes can be dragged out of its docking place and docked elsewhere. Docking handles appear when a pane is dragged.

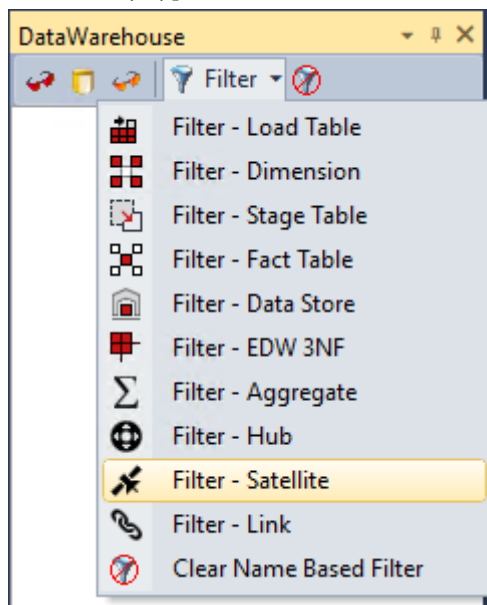


The **left pane** contains all the **objects** within the metadata repository. These objects are stored in object groups (e.g. Model). The object groups are in turn optionally stored in Projects, and the Projects are optionally stored within Project Groups.

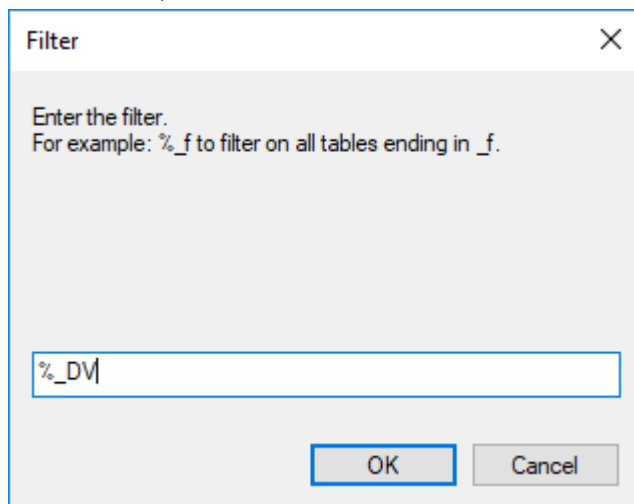
The **middle pane** is used to show the **results** of various **queries** on both the metadata and the underlying source and database tables. The middle pane is also used as the drop target in drag and drop operations. The status line at the bottom of the screen displays the current contents of the middle pane. To send the middle pane output to a file or to clipboard, see **Export Middle Pane Output** (on page 70).

The **right pane** is the **Browser Pane**, it shows both source and data warehouse systems. There are two browser panes available at any one time. The source may be the data warehouse itself. Typically, this pane is used as the source of information in the drag and drop operations. The Browser Pane can be filtered:

- To filter by type click the down arrow next to the **Filter** button.



- Or to filter by name click the **Filter** button and enter the filter criteria.



Name based filters can be cleared by clicking the down arrow next to the **Filter** button and selecting **Clear Name Based Filter**, seen in the image above.


The **bottom** pane is the **results** pane and it shows the results of any command executed on an object. Multiple messages can be displayed. Expand the '+' sign next to an object to see a complete list of messages relating to an object. When a report is run from the main menu, the results are displayed in a separate tab in the bottom pane.

Pop-up menus are typically available in all four panes.

F5 or Ctrl+R key can be used to refresh the left and right panes, when the cursor is positioned within these panes.

Toolbar

The **Builder** toolbar is shown below:

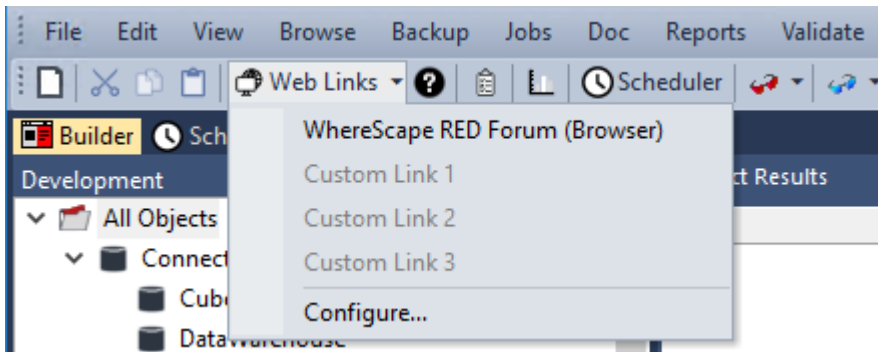


The **Diagram** button switches to the diagram view window which provides a graphical representation of the star schema and track back diagrams.

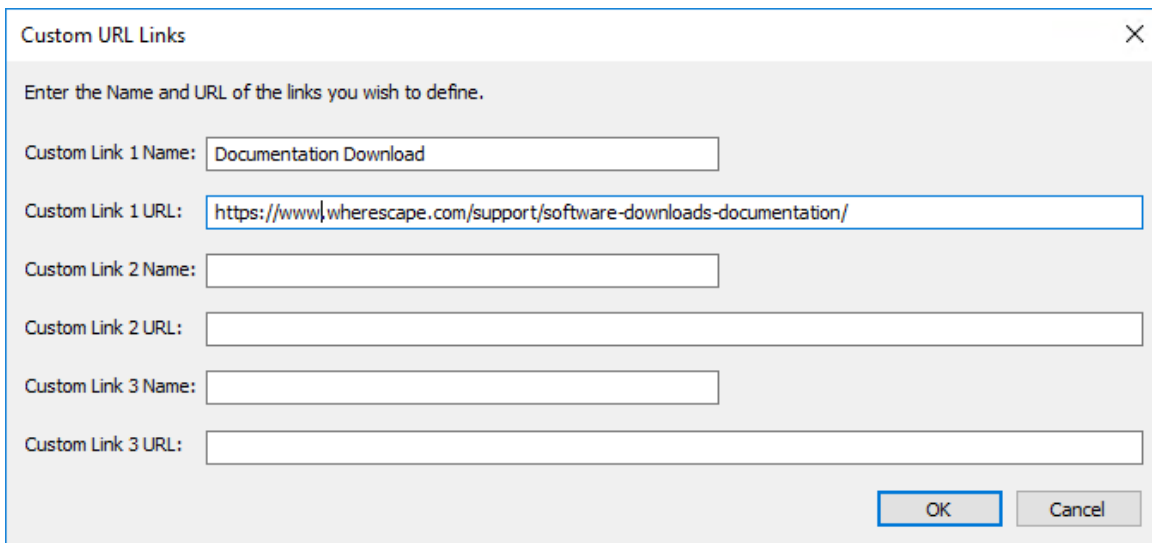
The **Scheduler** button switches to the scheduler control window.

The **two source browse** buttons (orange and blue) provides a quick method of invoking the source browser which populates the 'Browser pane'. Each of the two browse buttons, when clicked browses to the connection last used for that button. To change the connection being browsed, click the down arrow beside the glasses icon.

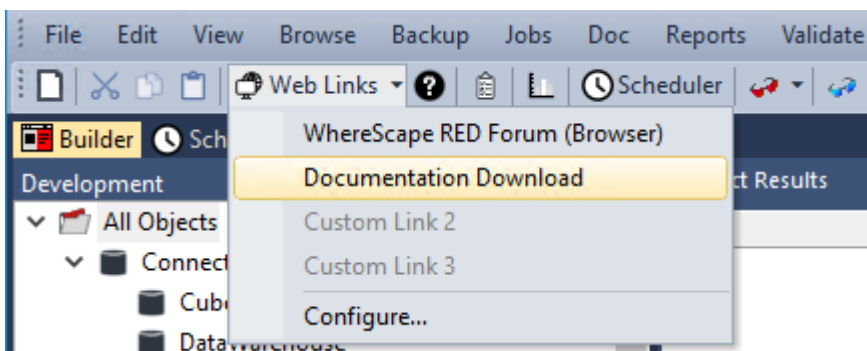
The **Web Links** button brings up the online WhereScape forum in a new tab. To select or enter other web links, click the down arrow beside the Web Links button.



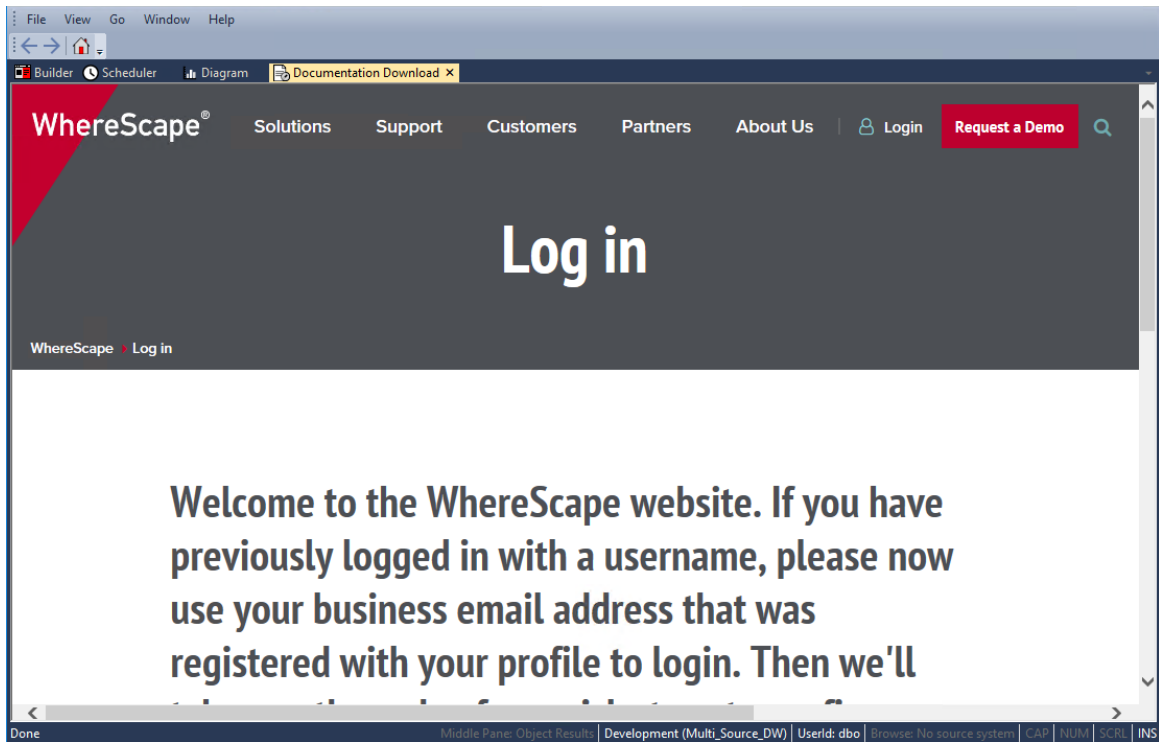
For example, let us add a Web Link to the Documentation Download on WhereScape's web page. Select **Configure** to open the following dialog box, allowing you to enter three custom URLs. Enter the name and the URL for the first custom link. Click **OK**.



To select the newly entered web link, click the down arrow beside the **Web Links** button. Select the newly entered Documentation Download option.



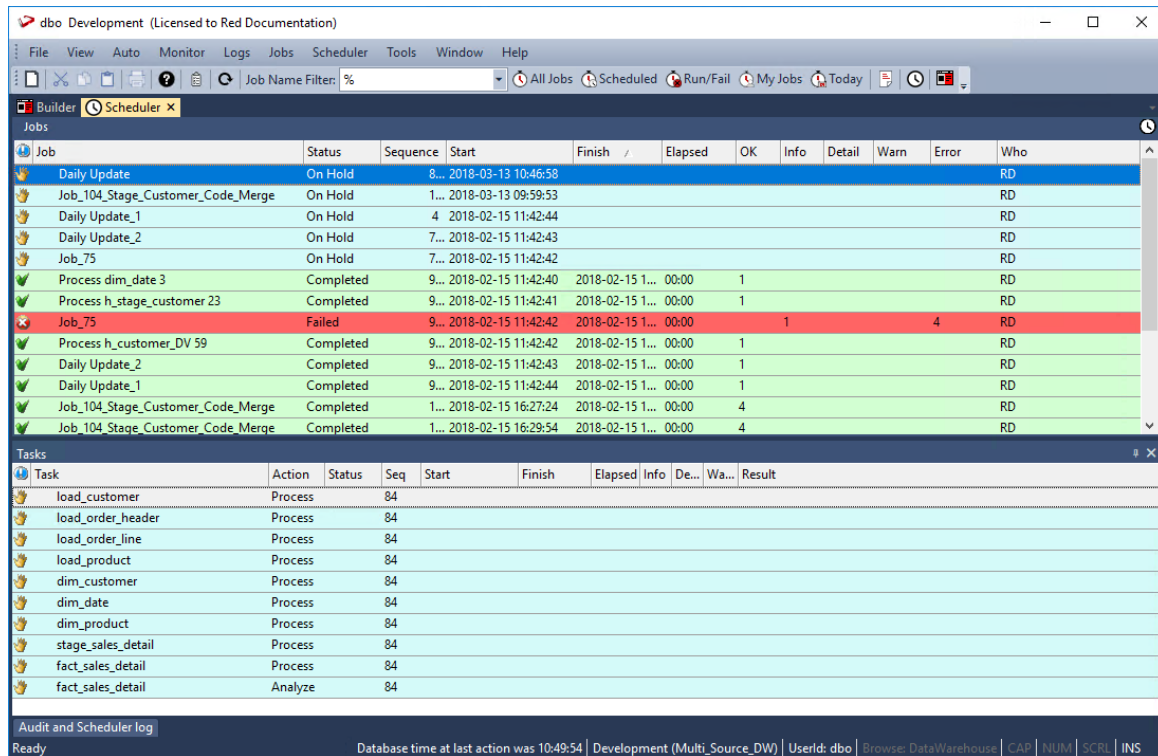
The relevant web page is displayed in a new tab.



Quick access buttons on the Builder Toolbar also include **versioning**, **building application**, **reports** and **document creation**.

SCHEDULER WINDOW

The **Scheduler Window** is used as the main interface to the scheduler.




Jobs can be scheduled, monitored, edited and deleted through this window. The window consists of two panes. The toolbar provides a quick way to display various job selections in the **top pane**. Double-click on the job in the top pane to see the tasks of the selected job in the **bottom pane**. Double-click on the task in the bottom pane to see the audit trail displayed in a separate tab in the bottom pane.

See the *Scheduler* (on page 685) chapter for more details.

Toolbar

The scheduler toolbar is shown below:



The **New Job** button invokes the dialog to create a new scheduled job.

The **Builder Window** button switches back to the main builder window.

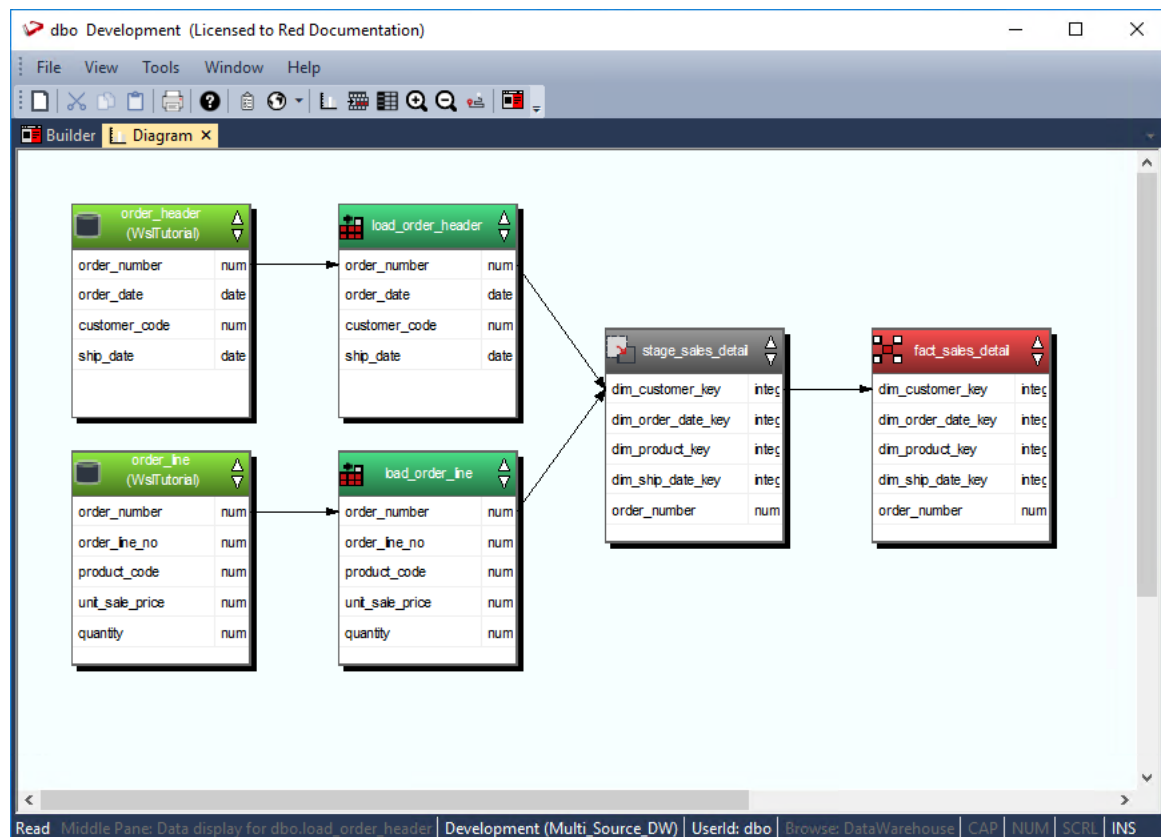
The **Auto Refresh** button, when depressed, will result in a refresh of the current (right pane) display every 10 seconds. Click the button again to stop the auto refresh.

The refresh interval can be adjusted through the menu option **Auto/Refresh** interval.

Quick access to different categories of jobs are also available via the toolbar.

DIAGRAM WINDOW

The **Diagram Window** is used to display the tables of the data warehouse in diagrammatic form, showing the various sources or targets of the selected object.




The **Diagram Selection** is as follows:

- Schema Diagram
- Source Diagram
- Joins Diagram
- Links Diagram
- Impact Diagram
- Dependency Diagram

See the **Diagrams** (on page 765) chapter for more details.

Toolbar

The **Diagram** toolbar is shown below:



The **New Diagram** button provides a dialog to allow the selection of the diagram type and table.

The **Overview** button provides a diagram showing the various objects in the WhereScape metadata and the standard flow of data through these objects. Repeated clicking of the **Overview** button will step through each stage of the data flow.

The **Toggle** button switches between display only diagrams and a printable variant. When the printable variation of a diagram is displayed, the **Grid** button will toggle the display of grid lines.

The **Builder** button switches back to the main builder window.

PROCEDURE EDITOR WINDOW

The **Procedure Editor Window** provides a means of viewing, editing, compiling, comparing and running procedures.

```

-----
-- DBMS Name      :   Sql Server 2000
-- Script Name    :   update_dim_customer
-- Description    :   Update the dimension dim_customer
-- Generated by   :   WhereScape RED
-- Generated for  :   WhereScape Limited
-- Author        :   WhereScape QuickStart
-----
-- Notes / History
--
CREATE PROCEDURE update_dim_customer
    @p_sequence      integer
    , @p_job_name     varchar(256)
    , @p_task_name    varchar(256)
    , @p_job_id       integer
    , @p_task_id      integer
    , @p_return_msg   varchar(256) OUTPUT
    , @p_status       integer      OUTPUT
AS
SET XACT_ABORT OFF  -- Turn off auto abort on errors
SET NOCOUNT ON    -- Turn off row count messages

-----
-- Control variables used in most programs
-----
DECLARE
    @v_msgtext      varchar(255)  -- Text for audit_trail
    , @v_sql         nvarchar(255) -- Text for sql statements
    , @v_set         integer       -- commit set
    , @v_analyze_flag integer     -- analyze flag
    , @v_step        integer       -- return code
    , @v_update_count integer     -- no. of records updated

```


Multiple such windows can be open at any one time, each processing a different procedure.

Comments (identified by a leading double dash --) are displayed in green in this window and the procedural code in black. The font is a fixed pitch font (by default) to make the indentation and alignment of code easier to view. The font, colors and indent size can all be changed if desired.

See the ***Procedures and Scripts*** (on page 629) chapter for more details.

Toolbar

The procedure editor toolbar is shown below:



The **Save** button will write the procedure to the WhereScape metadata repository in the database.

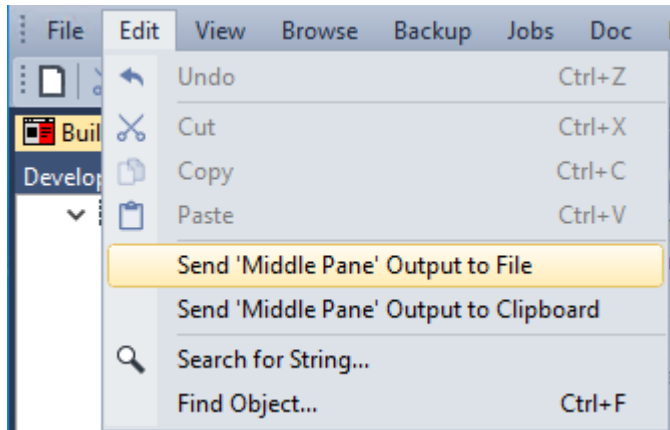
The **View Other Procedures** button allows the concurrent viewing of older versions of the current procedure, other procedures in the metadata, compiled procedures in the database and templates.

The **Compile** button will attempt to compile the procedure. Once compiled the procedure is stored within the database as well as in the metadata.

The **Execute** button will run a procedure that conforms to the WhereScape parameter syntax. See the chapter on procedures for more details.

EXPORT MIDDLE PANE OUTPUT

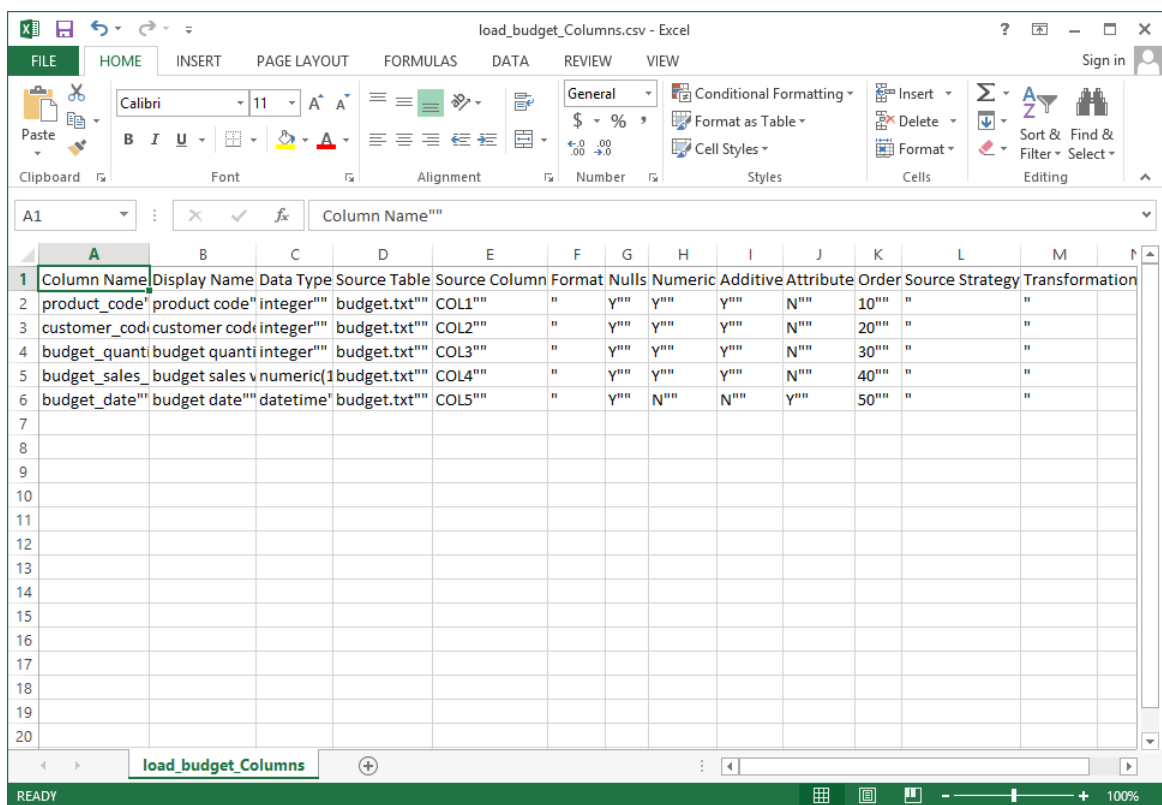
To send the middle pane output to a file, go to the **Edit** menu and select **Send 'Middle Pane' Output to File**.



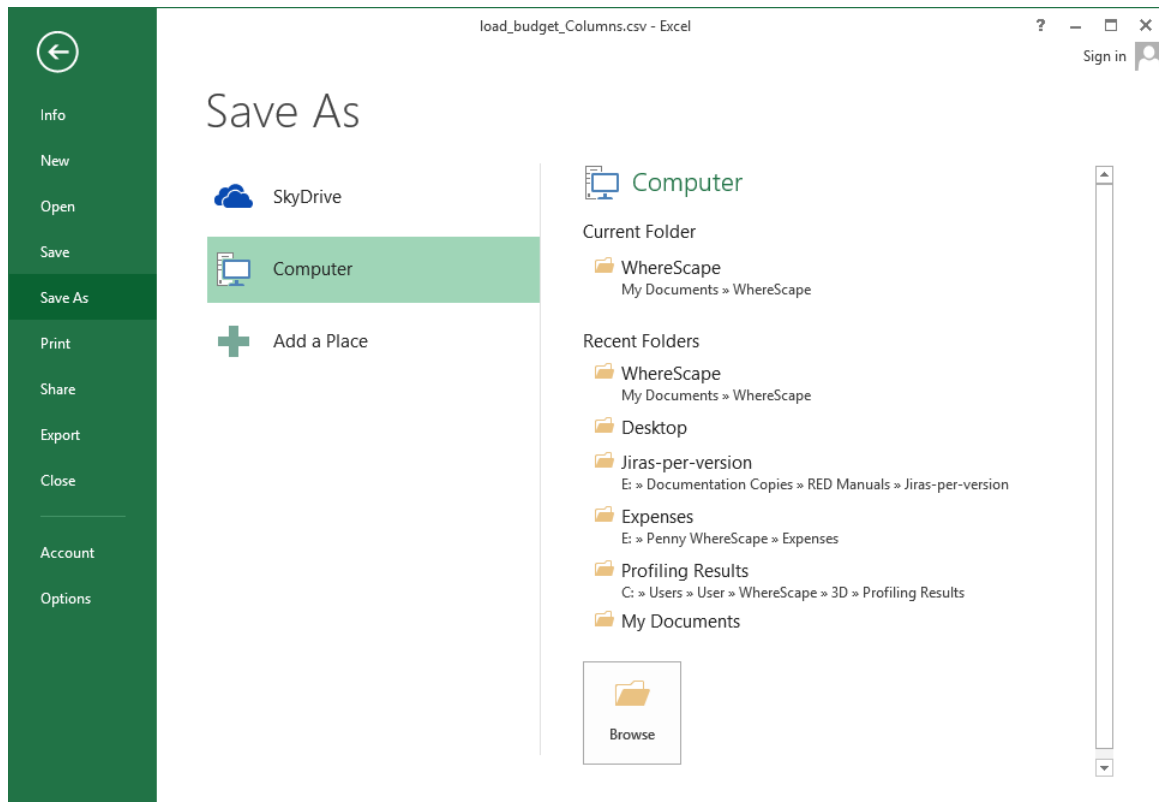
A file is created in the directory chosen in **Tools>User Preferences>Outputs>Output File Directory**.

If the Output File Extension is set to .csv in **Tools>User Preferences>Outputs>Output File Extension**, then an Excel file is created.

If the file is set to auto-open in **Tools>User Preferences>Outputs>Output File Auto Open**, then the file opens automatically.

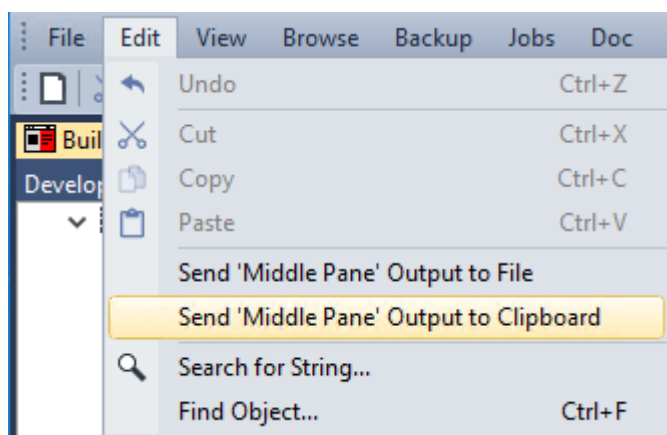


To save to a File, click on the **File tab** and then select either **Save** or **Save As** to save to a file.



To view the settings for **Middle Pane File Output**, see *Export Middle Pane Output Settings* (see "**Outputs**" on page 136).

To send the middle pane output to clipboard, go to the **Edit** menu and select **Send 'Middle Pane' Output to Clipboard**.

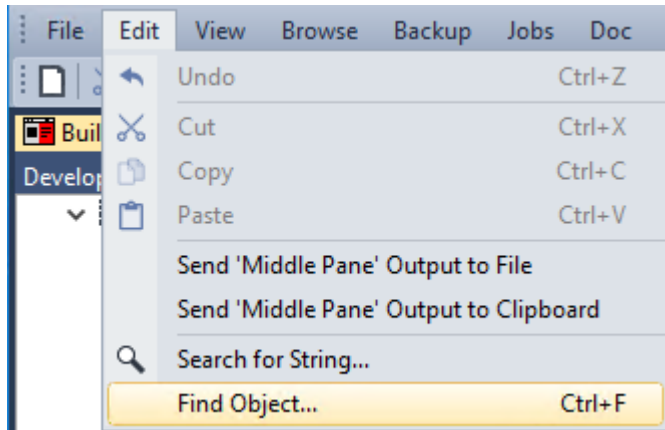


To view the settings for **Middle Pane Clipboard Output**, see *Export Middle Pane Output Settings* (see "**Outputs**" on page 136).

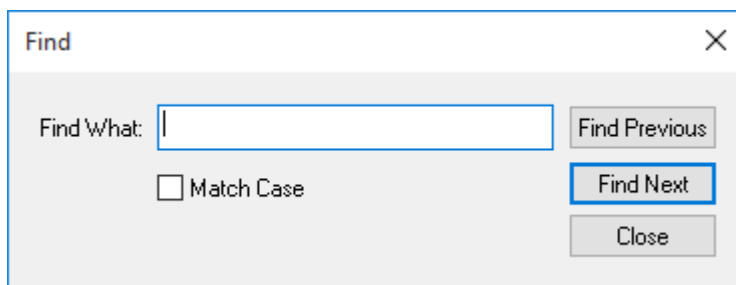
FIND FUNCTION

The Find function can help users quickly find a table when the list of tables has grown to a large size. The Find function can be accessed two ways within WhereScape RED:

- Click anywhere in the Object Pane or Browser Pane and press **Ctrl + F**, or
- select **Edit > Find Object** (this option searches the Object Pane only).



Both methods open the following dialog:



CHAPTER 4

TUTORIALS

Note: The WhereScape RED Tutorial Help has been removed from Release 8.0.2.0 and above.

CHAPTER 5

DEFAULT SETTINGS

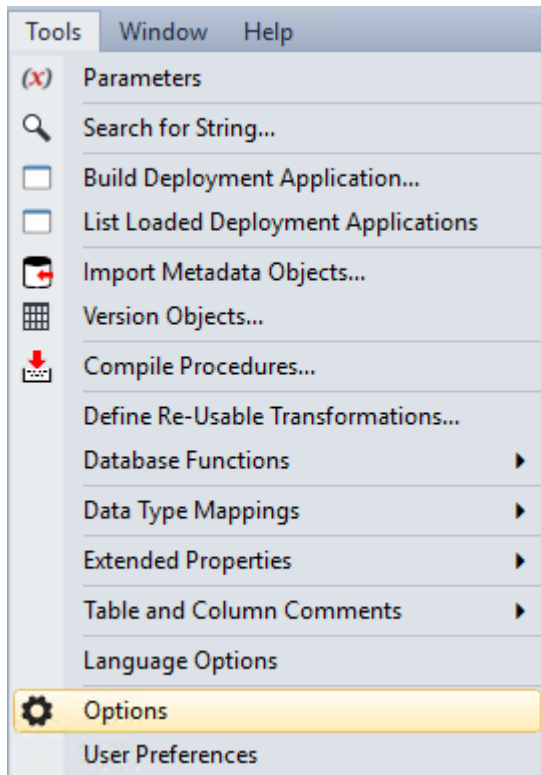
This chapter describes the settings and default values that can be set for a metadata repository. To access these settings, select the **Tools** menu; and then select either **Options** or **User Preferences**.

IN THIS CHAPTER

Settings - Options	78
Settings - User Preferences.....	123
Settings - Language Options	140

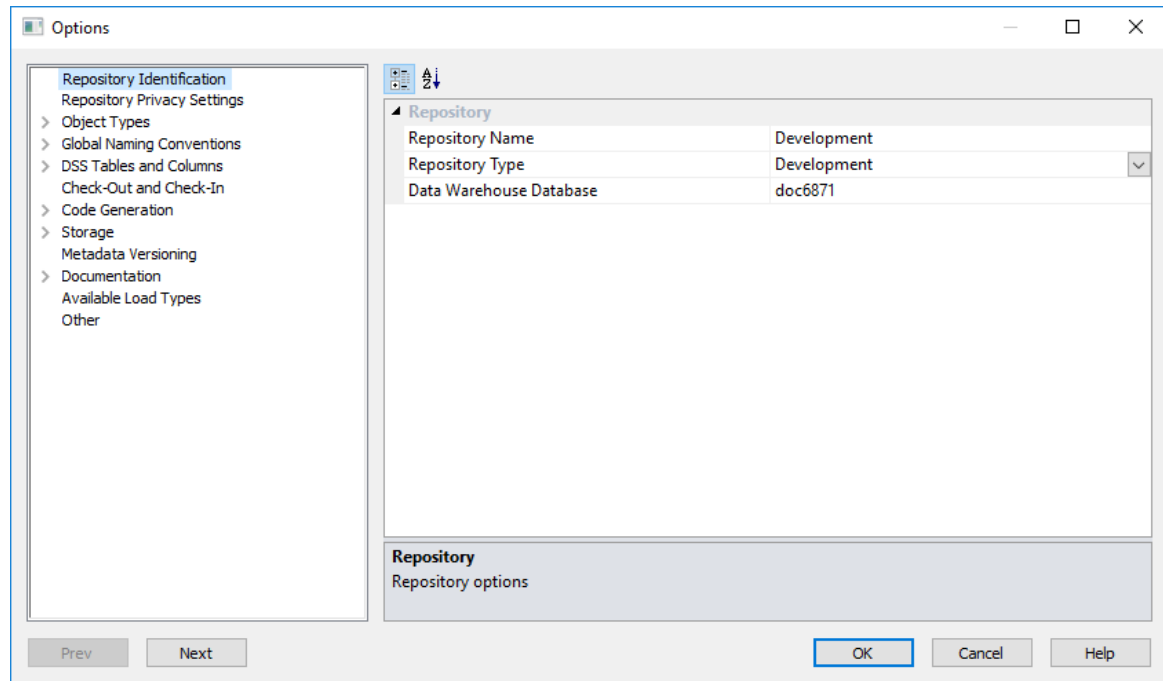
SETTINGS - OPTIONS

Select **Options** from the **Tools** menu.



SETTINGS - REPOSITORY IDENTIFICATION

This option enables users to set the **Repository Identification** settings.



Repository Name

Set the **name** for the repository. This name appears in the top left corner of the title bar in WhereScape RED. Restart WhereScape RED for repository name changes to take effect.

Repository Type

Set the **type** for the repository. The repository type must reflect the environment. For example, a 'Production' type must be chosen for the production environment.

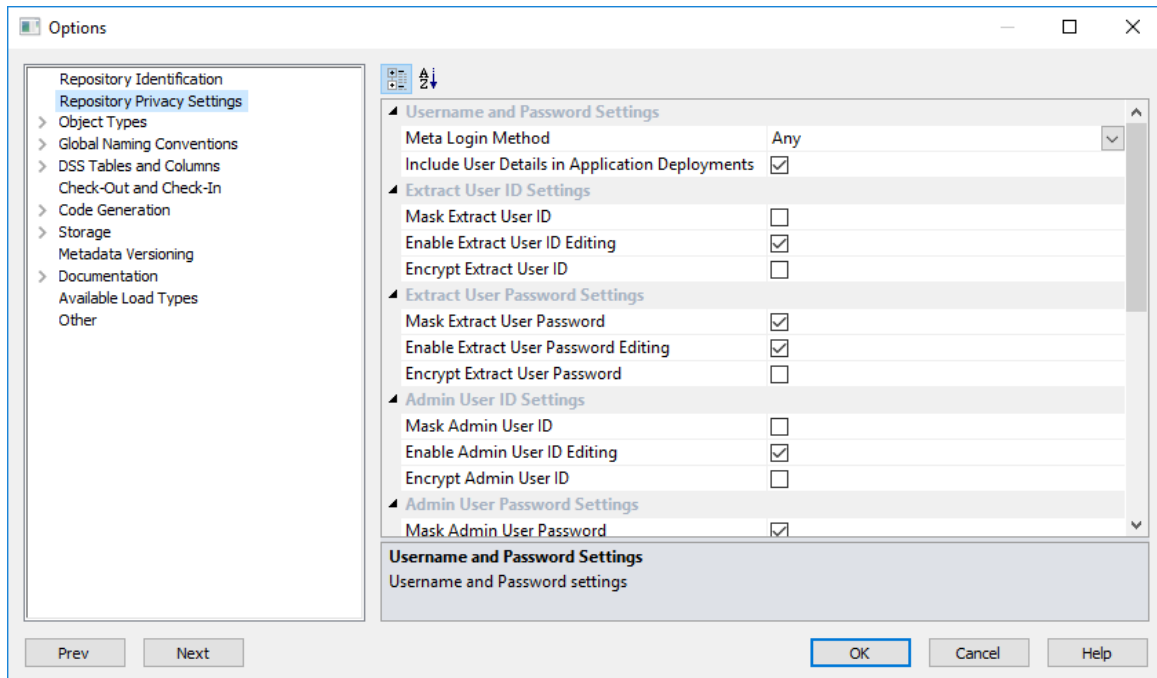
Data Warehouse Database

The database of the data warehouse that is used by the WhereScape RED scheduler. You should not normally need to change this value.

SETTINGS - REPOSITORY PRIVACY SETTINGS

This option allows users to set the **Repository Privacy Settings**.

WARNING: For UNIX/Linux scheduler processing, the **Encrypt** User and Password options cannot be used. Encrypt options are only supported when using a Windows scheduler.



Changing Repository Settings

Since the repository privacy settings can be configured from the Tools (Options) menu, for an environment to be secured, **a database administrator must change the permissions on table ws_meta_admin table to read-only after the appropriate repository privacy change settings in WhereScape RED have been made.**

Note: Changing this set of permissions to read-only is something which occurs outside of WhereScape RED and will be dependent on the specific meta data database.

Username and Password Settings

- **Meta Login Method** - This option can be set to restrict users to using a particular login method for the meta repository
- **Include User Details in Application Deployments** - Includes or excludes User Details in Application Deployment packages

Extract User ID Settings

- **Mask Extract User ID** - Masks the input of the "Extract/Unix/Windows User ID" on the connection properties
- **Enable Extract User ID Editing** - Enables editing the "Extract/Unix/Windows User ID" via the connection properties
- **Encrypt Extract User ID** - Encrypts "Extract/Unix/Windows User ID" in the meta repository using WhereScape encryption

Extract User ID Settings

- **Mask Extract User Password** - Masks the input of the "Extract/Unix/Windows User Password" on the connection properties
- **Enable Extract User Password Editing** - Enables editing "Extract/Unix/Windows User Password" via the connection properties
- **Encrypt Extract User Password** - Encrypts "Extract/Unix/Windows User Password" in the meta repository using WhereScape encryption

Admin User ID Settings

- **Mask Admin User ID** - Masks the input of the "Admin/DSS User ID" on the connection properties
- **Enable Admin User ID Editing** - Enables editing the "Admin/DSS User ID" via the connection properties
- **Encrypt Admin User ID** - Encrypts "Admin/DSS User ID" in the meta repository using WhereScape encryption

Admin User Password Settings

- **Mask Admin User ID** - Masks the input of the "Admin/DSS User ID" on the connection properties
- **Enable Admin User ID Editing** - Enables editing the "Admin/DSS User ID" via the connection properties
- **Encrypt Admin User ID** - Encrypts "Admin/DSS User ID" in the meta repository using WhereScape encryption

Teradata Wallet User ID Settings

- **Mask Teradata Wallet User ID** - Masks the input of the "Teradata Wallet User ID" on the connection properties
- **Enable Teradata Wallet User ID Editing** - Enables editing the "Teradata Wallet User ID" via the connection properties
- **Encrypt Teradata Wallet User ID** - Encrypts the "Teradata Wallet User ID" in the meta repository using WhereScape encryption

Teradata Wallet String Settings

- **Mask Teradata Wallet String** - Masks the input of the "Teradata Wallet String" on the connection properties
- **Enable Teradata Wallet String Editing** - Enables editing the "Teradata Wallet String" via the connection properties
- **Encrypt Teradata Wallet String** - Encrypts the "Teradata Wallet String" in the meta repository using WhereScape encryption

JDBC User ID Settings

- **Mask JDBC User ID** - Masks the input of the "JDBC User ID" on the connection properties
- **Enable JDBC User ID Editing** - Allows editing the "JDBC User ID" via the connection properties

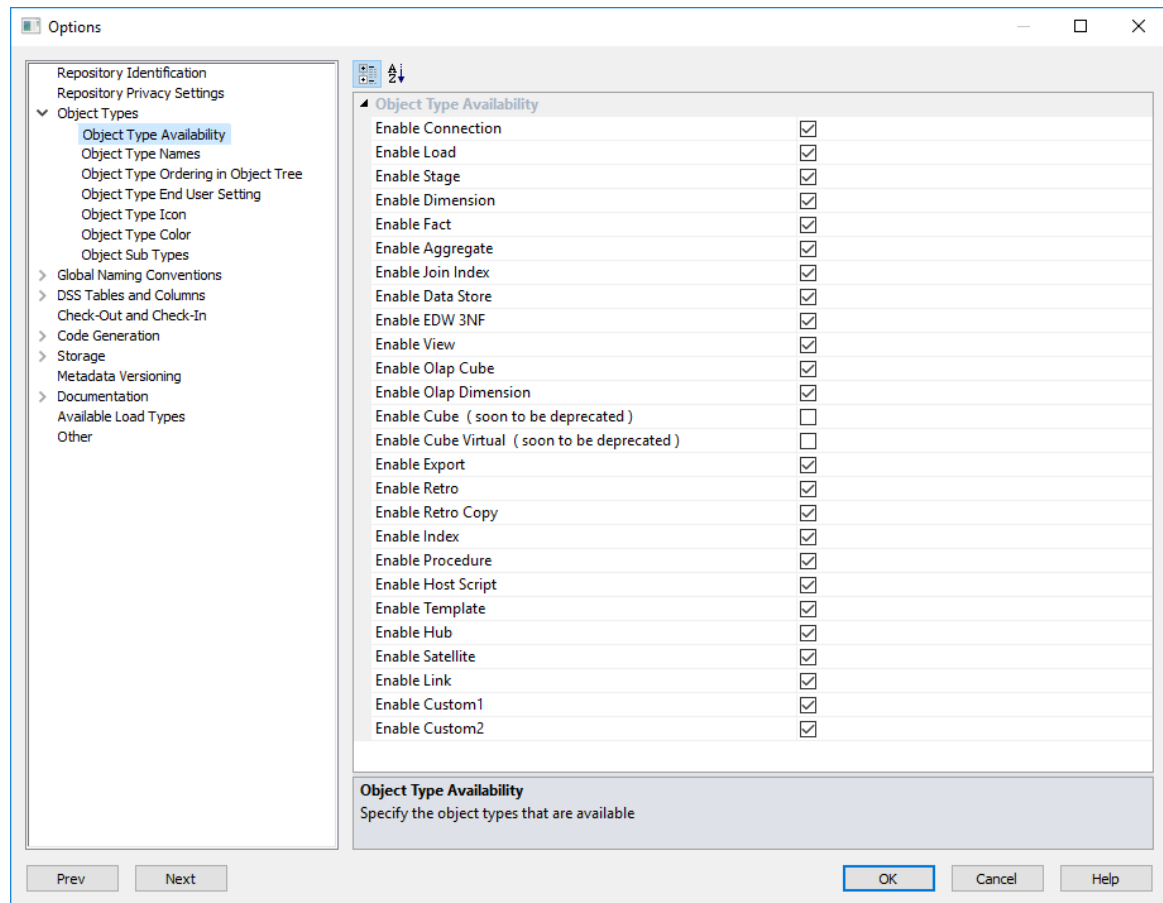
JDBC User Password Settings

- **Mask JDBC User Password** - Masks the input of the "JDBC User Password" on the connection properties
- **Enable JDBC User Password Editing** - Allows editing the "JDBC User Password" via the connection properties

SETTINGS - OBJECT TYPES

OBJECT TYPE AVAILABILITY

This option enables you to activate or deactivate the various object types within the data warehouse repository.

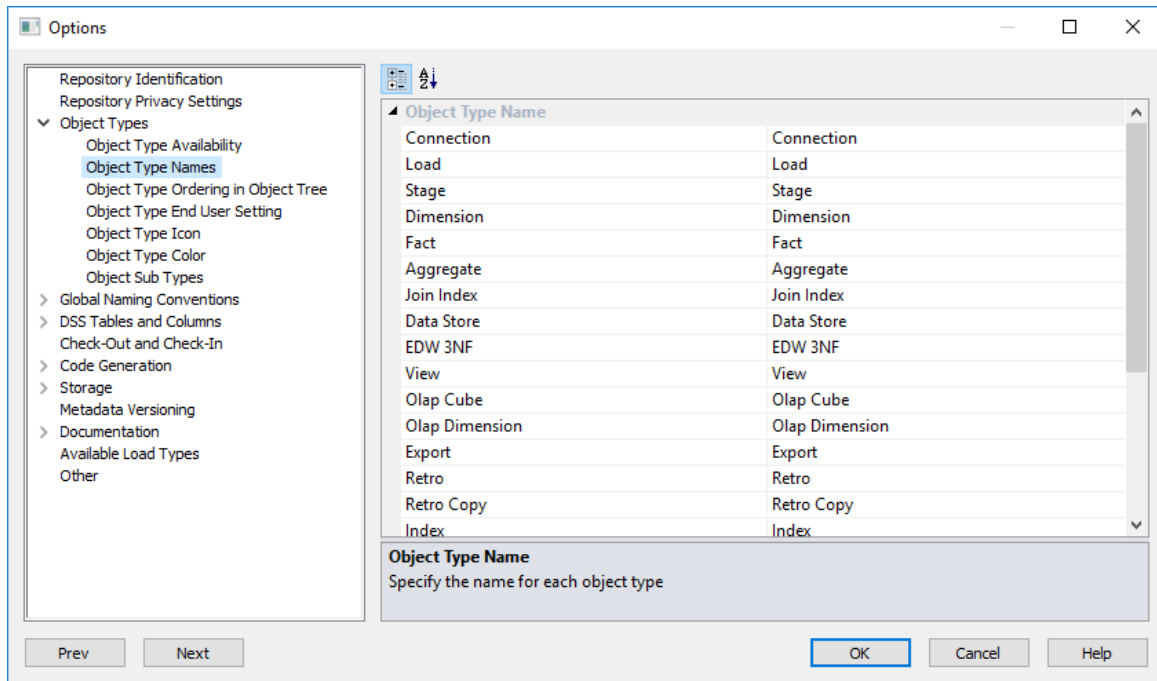


Object Type Availability

Enable/Disable object types in the data warehouse by selecting/clearing the availability checkboxes for each object type. All settings are enabled by default, except for the **Enable Cube** and **Enable Virtual Cube** options which are soon to be deprecated.

OBJECT TYPE NAMES

This option enables you to set the **names** for the various object types in RED.



Object Type Name

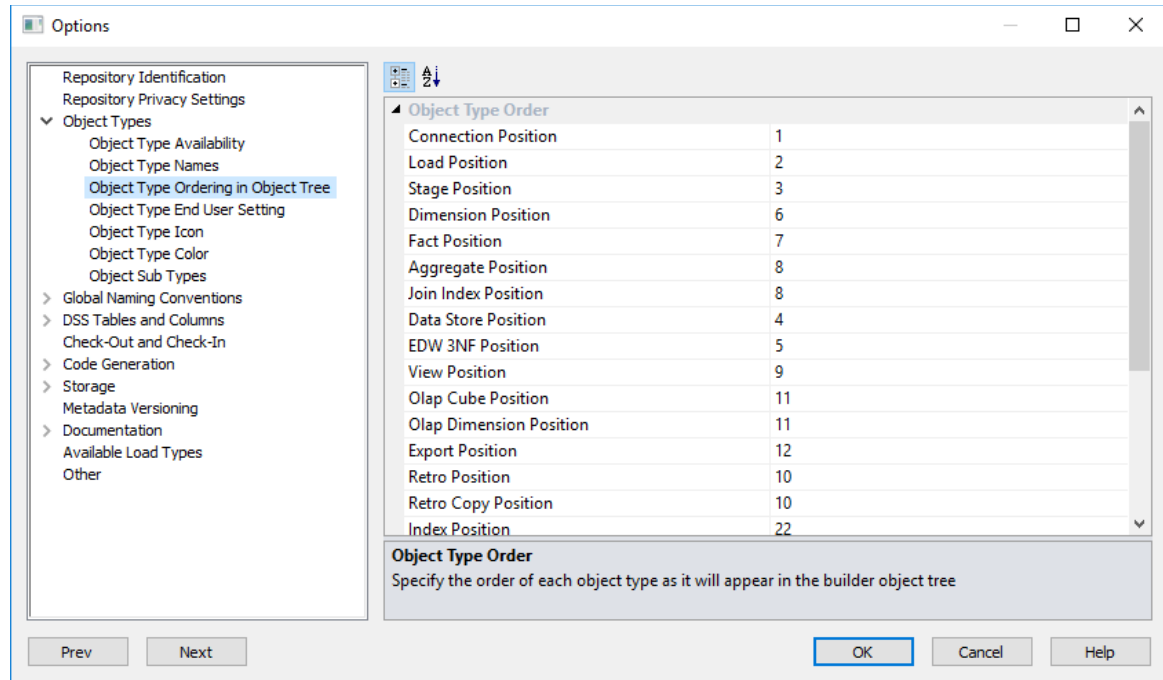
Set the desired **name** for each object type.

NOTE - Data Vault Repository Types: Users with Data Vault model type licenses that chose a **Data Vault repository type** while creating the RED metadata repository will have appropriate Data Vault repository default settings, such as Object Type Names, Global Naming of Tables, Indexes, Key Columns and Procedures/Scripts, as well as other repository settings/user preferences.

These repository types will have their default Object Type Names of Normalized and Data Store objects set to Hub/Link and Satellite.

OBJECT TYPE ORDERING

This option enables you to set the **ordering** in which the object types appear in the object tree on the left pane.



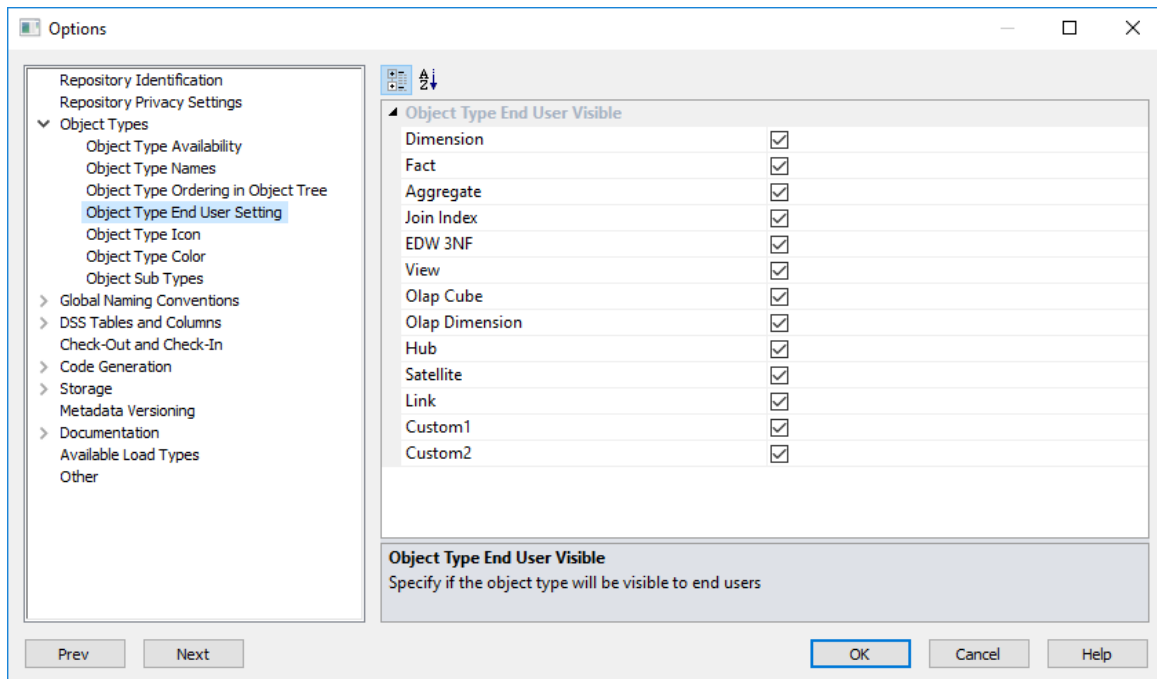
Object Type Order

Set the **ordering** of the object types as displayed in the object tree pane.

OBJECT TYPE END USER SETTING

This option enables you to set the Object types as **end user objects**.

All the options in this screen are selected by default.



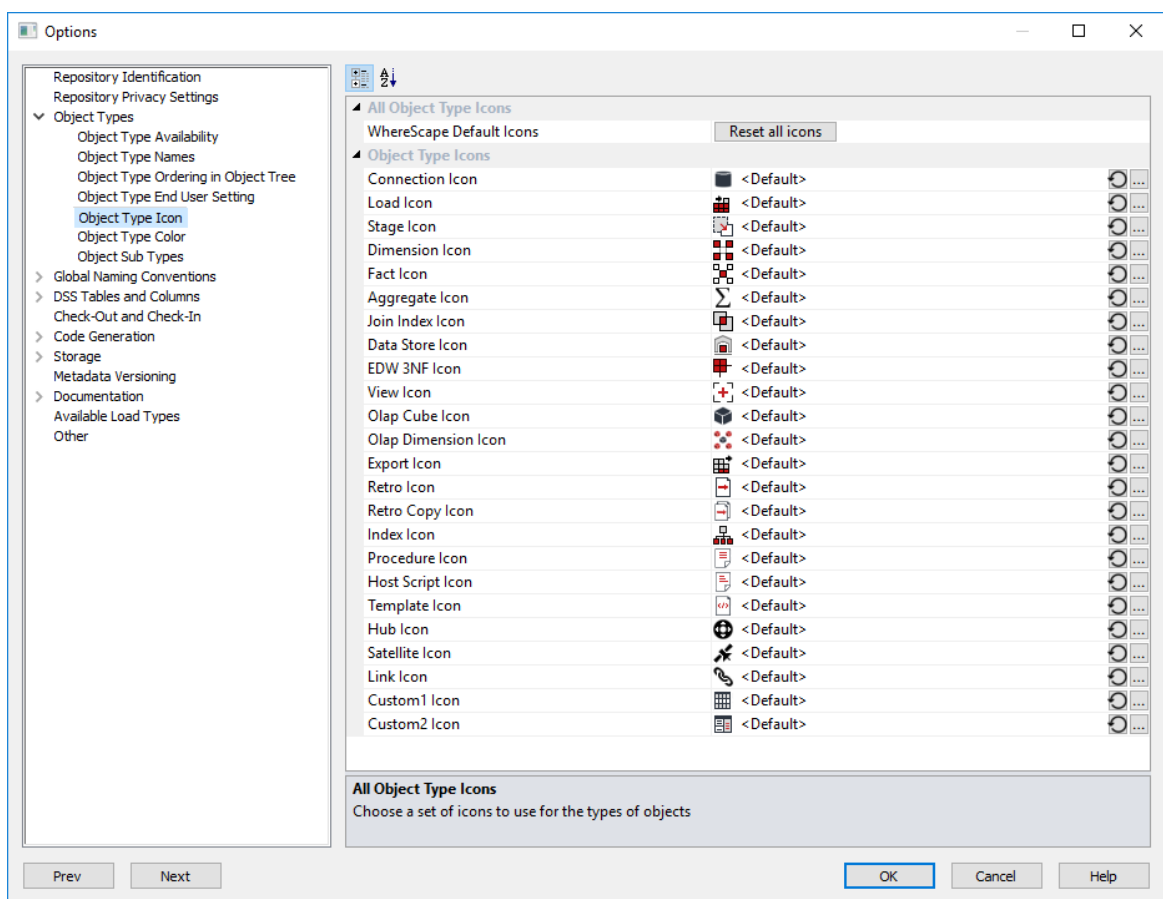
Object Type End User Visible

Clear an object type check box to remove it as an end user object, else leave selected.

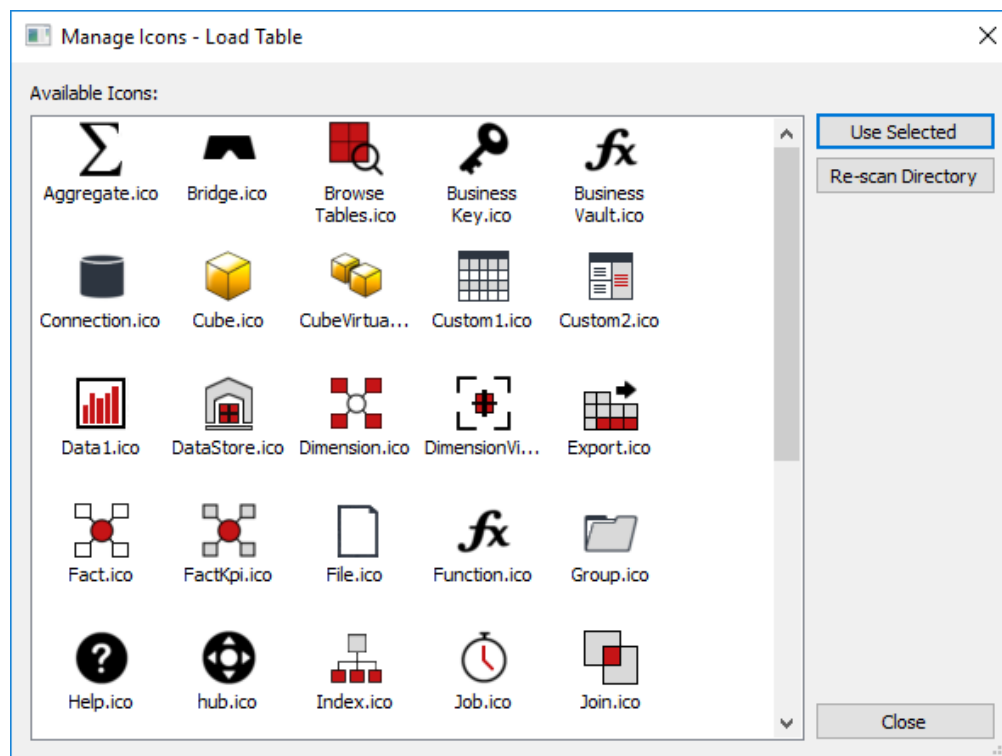
OBJECT TYPE ICON

This option enables you to configure the **Icons** for all Object Types. To configure custom Object Type icons:

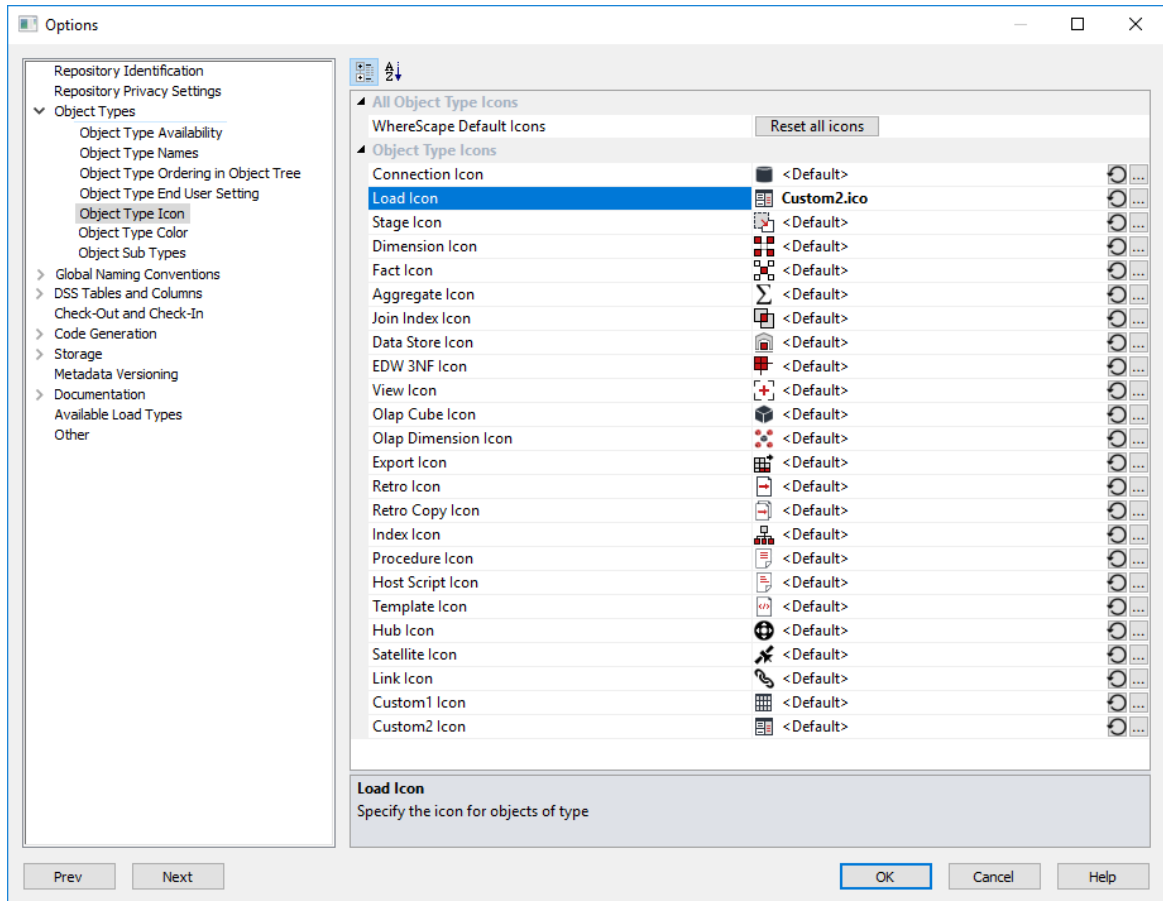
- 1 Create an 'Icons' folder in the WhereScape RED install directory, if it doesn't exist. For example:
C:\Program Files (x86)\WhereScape\Icons
- 2 Place custom '.ico' files in the Icons folder.
- 3 In RED, select **Options>Object Types>Object Type Icon**.



- 4 Click the ellipses button '...' next to each Object Type and select the desired icon.



- 5 All configured icons are displayed as file names in the screen. To save changes click **OK**.



To Reset An Icon

To reset an icon to default, click the reset button (⌛) next to the icon.

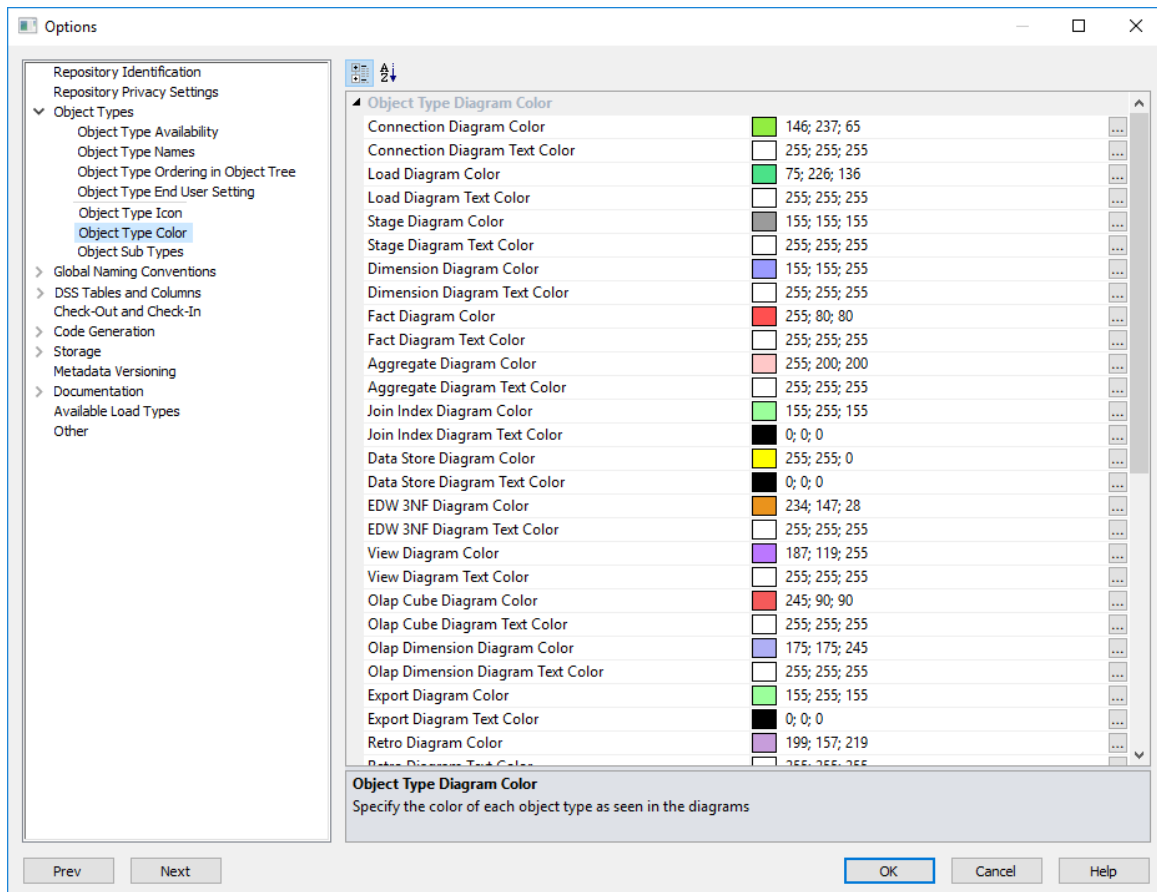
Reset All Icons

To reset all icons to default, click the **Reset All Icons** button at the top of the dialog.

Note: All installations must have a copy of the icon directory.

OBJECT TYPE COLOR

This option enables you to set the **diagram colors** for each object type in RED.

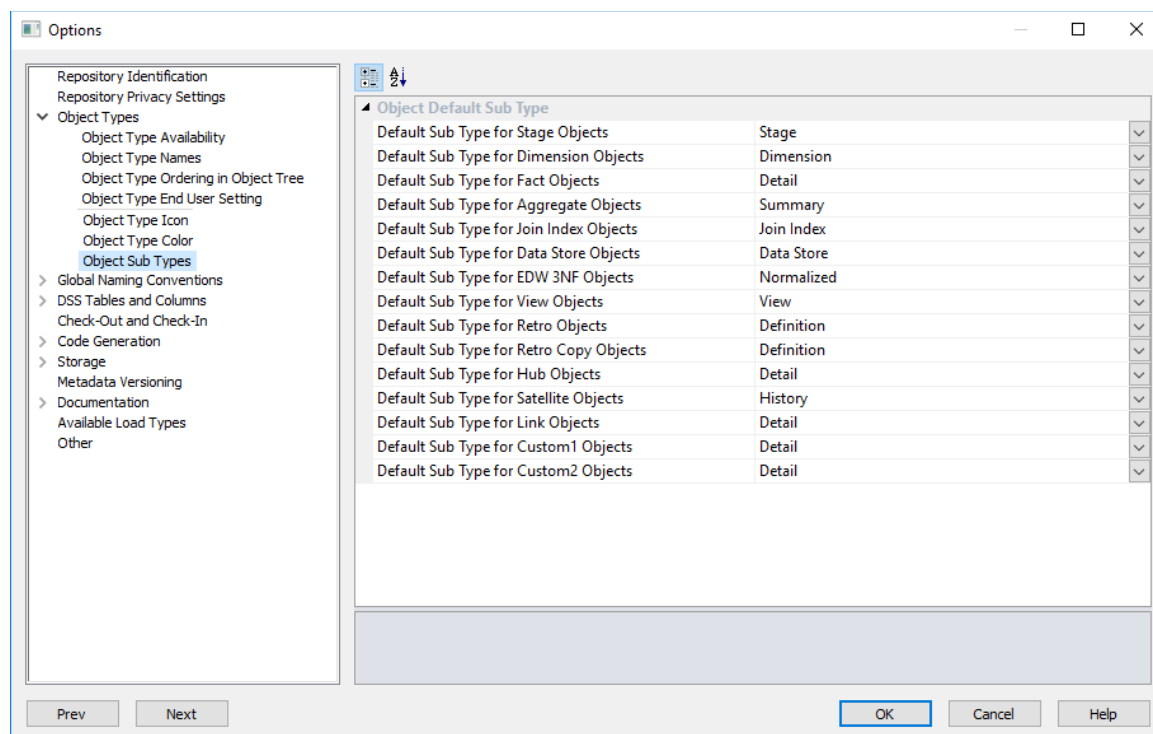


Object Type Diagram Color

Set the **Diagram color** for each object type.

OBJECT SUB TYPES

This option enables you to set the **default sub type** for activated objects in RED.



Object Default Sub Type

This option enables specifying the **Default Sub Type** for enabled object types. Select the desired default sub-types from the Object's drop-down lists.

For example, to have **Dimension** objects created in RED as **Changing Dimensions** at the time of drag and drop, select the **Changing Dimension** option in the **Default Sub Type for Dimension Objects**.

After the table is dragged and dropped, users can simply hit enter to proceed on the Dimension Type where the **Slowly Changing** type is already defaulting to the sub type option previously selected in **Tool>Options** menu.

Dimension Type ✕

Four methods are provided for managing dimensions. Please select the desired method. 🔗

1. Normal. The dimension is updated based on a business key, with new records being added if required. All columns except the business key can change.
Normal
2. Slowly changing. Changes in the values of selected columns result in new dimensional records being created. In all other respects the same as Type 1.
Slowly Changing
3. Previous data retained. The previous values of selected columns are stored in additional columns. In all other respects the same as Type 1.
Previous values
4. Date Ranged. The source system provides a date ranged business key. Similar to Type 2 except that we deal with the record as a whole and the dates are provided.
Date Ranged

The Dimension Properties screen reflects the selected table sub type on the **Table Type** drop-down list.

Dimension dim_customer_changing
✕

- Properties
- Storage
- Override Create DDL
- Language Mapping
- Extended Properties
- Purpose
- Concept
- Grain
- Examples
- Usage
- Notes

Table Name:
Table Type: Changing Dimension

Unique Short Name: (maximum 22 characters)

Business Display Name (EUL):

Description:

Update Procedure: (None)
Rebuild ▼

Custom Procedure: (None)
Rebuild Regenerate

Mnemonic (EUL):

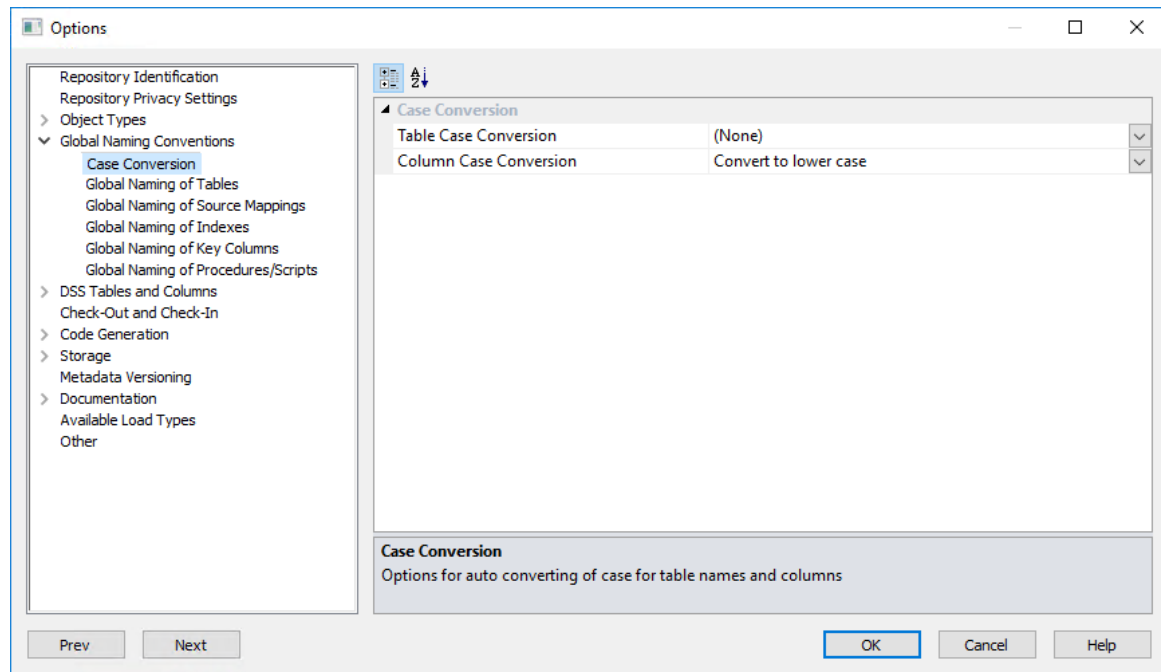
Timestamps
Metadata Structure
Database
Database Altered:

OK
Cancel
Help

SETTINGS - GLOBAL NAMING CONVENTIONS

CASE CONVERSION

This option enables you to set the **case conversion** methods for the Tables and Columns in RED.

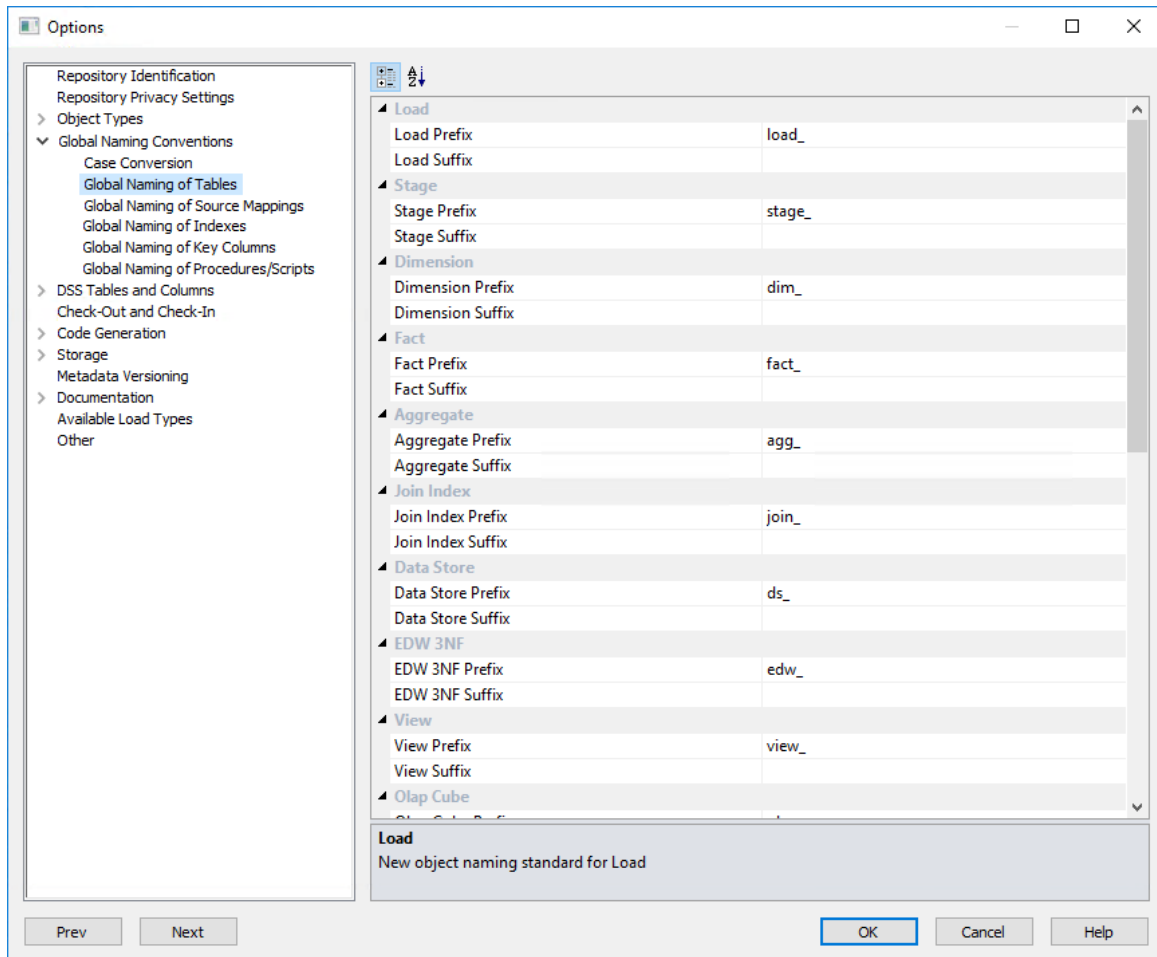


Case Conversion

Set the **Table Case Conversion** method and the **Column Case Conversion** method from the drop-down lists.

GLOBAL NAMING OF TABLES

This option enables you to set the **Global Naming of Tables** options.



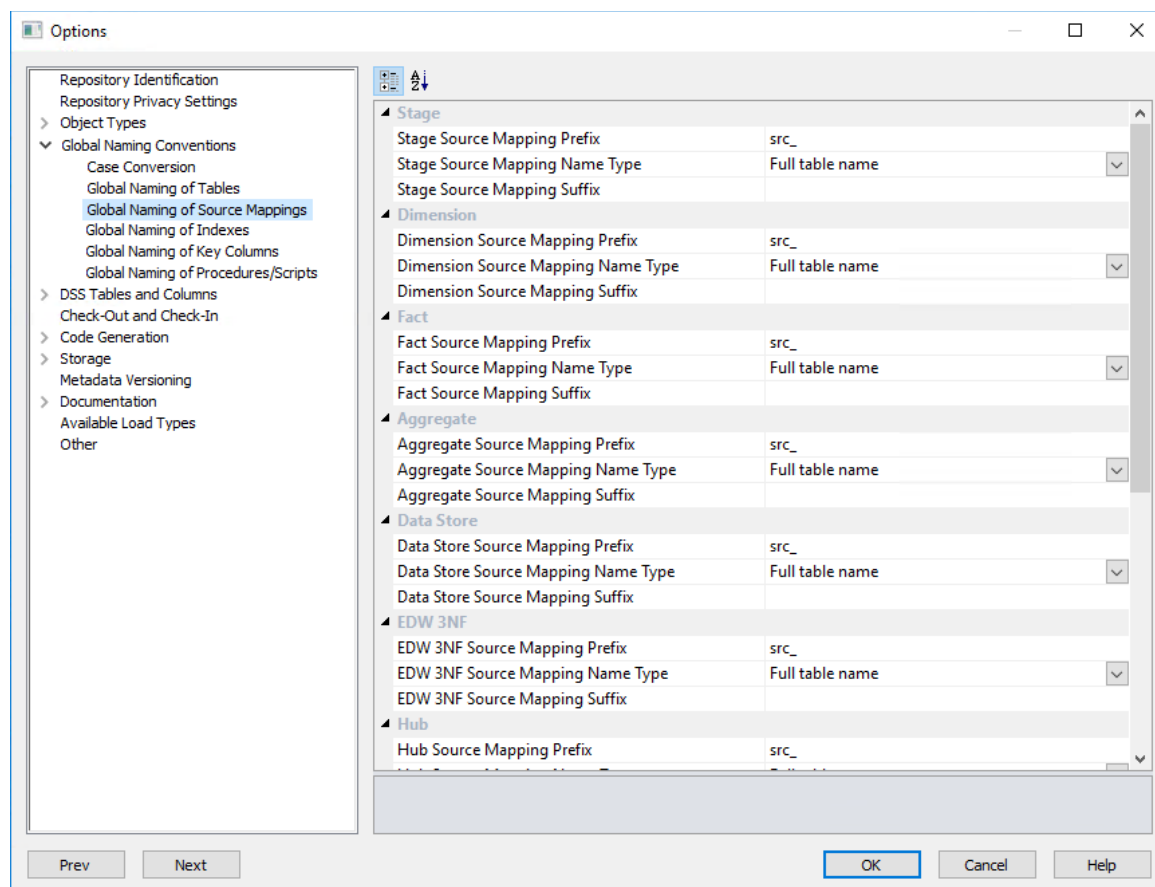
A prefix and/or a suffix string can be applied to an object name. Within Oracle and IBM DB2, a table name may be a maximum of 30 characters long, so these pre and post fix strings should not be more than eight characters long (WhereScape RED short names are a maximum of 22 long in Oracle and SQL Server and 12 long in DB2).

From the example screen above, if a source table called **customer** (with a short name of 'customer') was dragged into a Load table drop target then the default name would be **load_customer**.

The object name defaults shown above are the values that are installed with the base metadata. They can be changed at any stage, however, the change does not affect any existing objects. Therefore, if a new naming regime is chosen any existing objects will need to be renamed through the Properties screen of the object.

GLOBAL NAMING OF SOURCE MAPPINGS

This option enables users to set the **Global Naming of Source Mappings** options.



Similar to table objects created in RED, a prefix and/or a suffix string can also be applied to source mapping child objects.

Within Oracle and IBM DB2, a table name may be a maximum of 30 characters long, so these pre and post fix strings should not be more than eight characters long (WhereScape RED short names are a maximum of 22 long in Oracle and SQL Server and 12 long in DB2).

From the example screen above, if a source table was dragged into a table drop target called **stage_customer** (with a full table name of 'customer') then the default name for the first source mapping object created would be **src_stage_customer_1**.

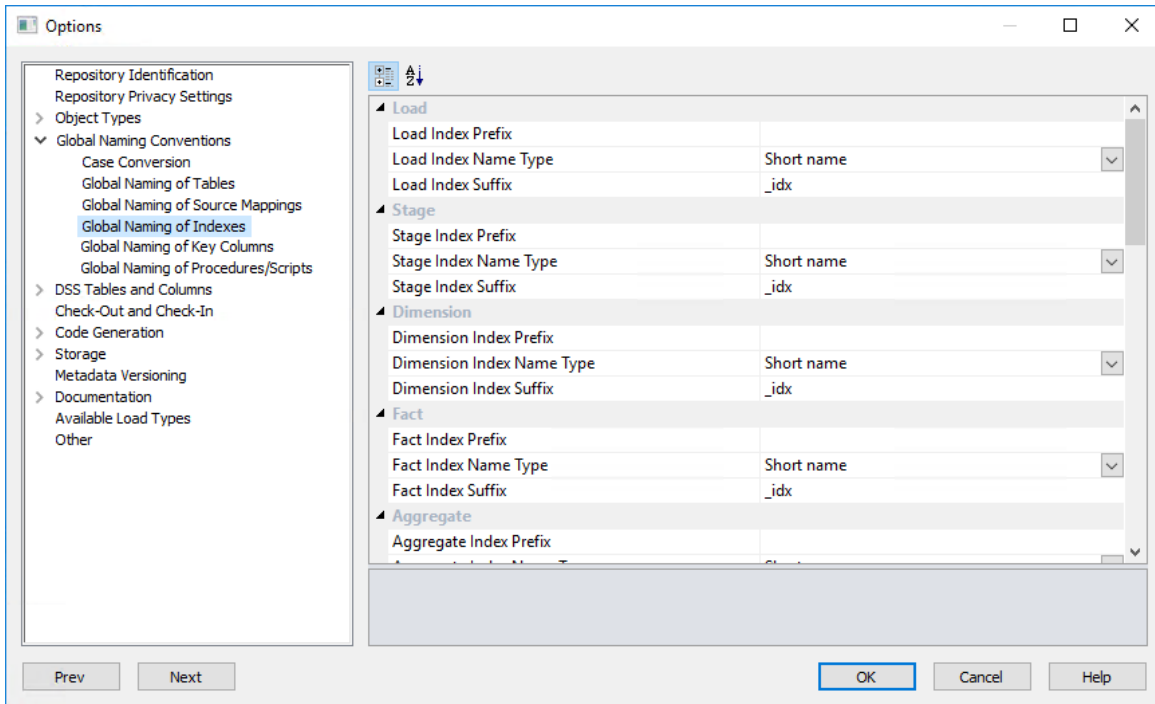
The same applies for the second source mapping object created—the object would be named **src_stage_customer_2**.

The object name defaults shown above are the values that are installed with the base metadata.

They can be changed at any stage, however, the change does not affect any existing objects. Therefore, if a new naming regime is chosen any existing objects will need to be renamed through the Properties screen of the object.

GLOBAL NAMING OF INDEXES

This option enables users to set the **Global Naming of Indexes** options.



Whenever a new procedure is defined, WhereScape RED builds or rebuilds a standard set of indexes for the table. These indexes will be created using the standard defined. As with the key naming, we can set either a pre-fix or suffix value, or in fact both, as well as choosing the use of either the table name or the short name associated with the table.

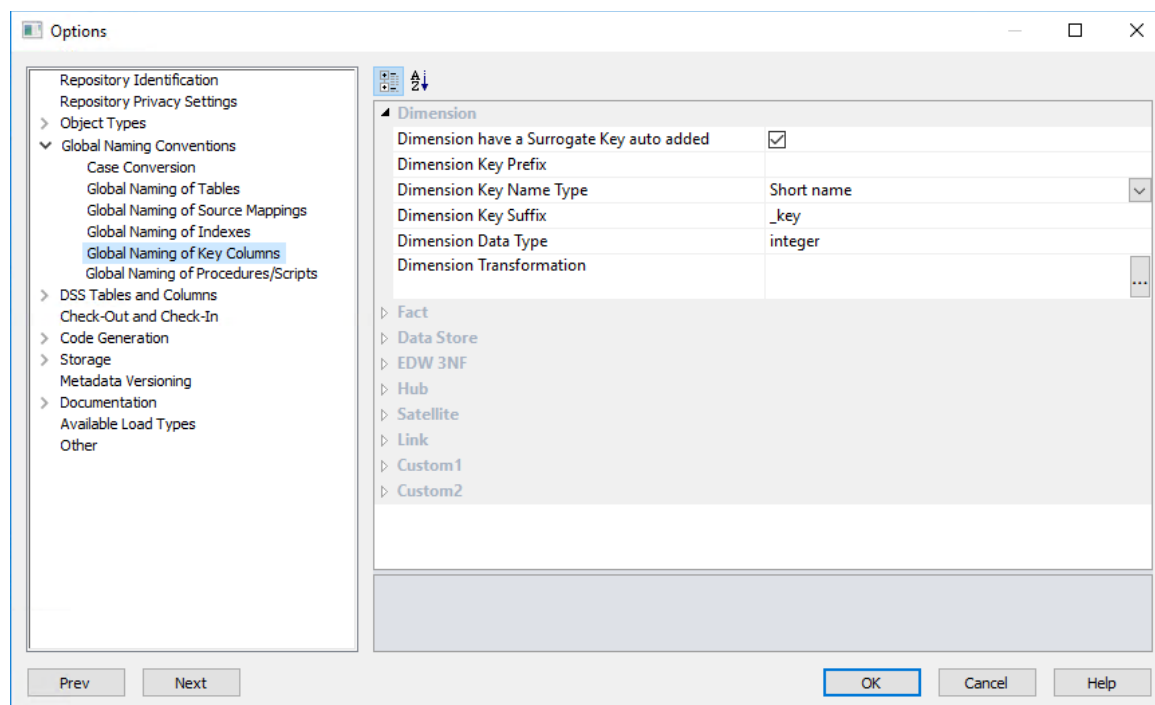
In addition to the naming specified above, WhereScape RED will add up to a further 3 characters to the end of the index name. These additional values will be "_0" through "_99", or "_A" through "_Z", or "_PR". When a new index is manually added, it will have the additional value of "_x" by default. This should be changed. The WhereScape RED naming standard for indexes is described below, but any valid name may be used.

From the example screen above, a dimension table would have indexes generated using the short name and with a suffix of "_idx". Therefore, a dim_sales dimension table would have indexes such as dim_sales_idx_x.

Ultimate suffix	meaning
_0	artificial key
_A	primary business key
_B through _Z	secondary business keys
_PR	primary index

GLOBAL NAMING OF KEY COLUMNS

This option enables you to set the **Global Naming of Key Columns**.



During the drag and drop generation of new tables, WhereScape RED builds an artificial (surrogate) key for the table if surrogate keys are enabled.

- The naming convention for the surrogate key can be set through the same menu option as above.
- **Prefix** and **suffix** values can also be added.
- There is a choice between the inclusion of the **full table name**, **short name** or **base name** assigned to each table.
- In the example screen above, which is the default, a dimension table key would use the table **short name** and have a suffix of "**_key**": for example your load_customer table would generate a key called **dim_customer_key** if it was dragged into a dimension drop target.

To have a table with non identity columns as surrogate keys, you can set the table's **Data Type** to **integer**. During the procedure generation this will create a logic that associates a sequential number to the artificial key of the dimension when a new row is inserted into the table.

- The example above displays the defaults for **Dimension** options but to set these fields on **Fact**, **Data Store** and **EDW 3NF** tables, expand the fields below Dimension to view and set your required options.
- For more on **Artificial Keys** see *Dimension Artificial Keys*, *Data Store Artificial Keys* and *EDW 3NF Artificial Keys*.

Dimension have a Surrogate Key auto added

Set this field if a Surrogate key column is to be added automatically to a table. Default for Dimension is set. Default for Fact, Data Store and EDW 3NF is not set.

Dimension Key Prefix

Key prefix that can be added to a new Dimension Key.

Dimension Key Name Type

Key name type for new Dimension keys. Select between Short name, Full table name and Base name.

Dimension Key Suffix

Key suffix that can be added to a new Dimension key.

Dimension Data Type

Default data type for the Dimension surrogate column definition. Set to Integer if you want a non identity column to be used as the surrogate key.

Dimension Transformation

Transformation where a database compliant SQL statement can be used for the surrogate key on new Dimension entries.

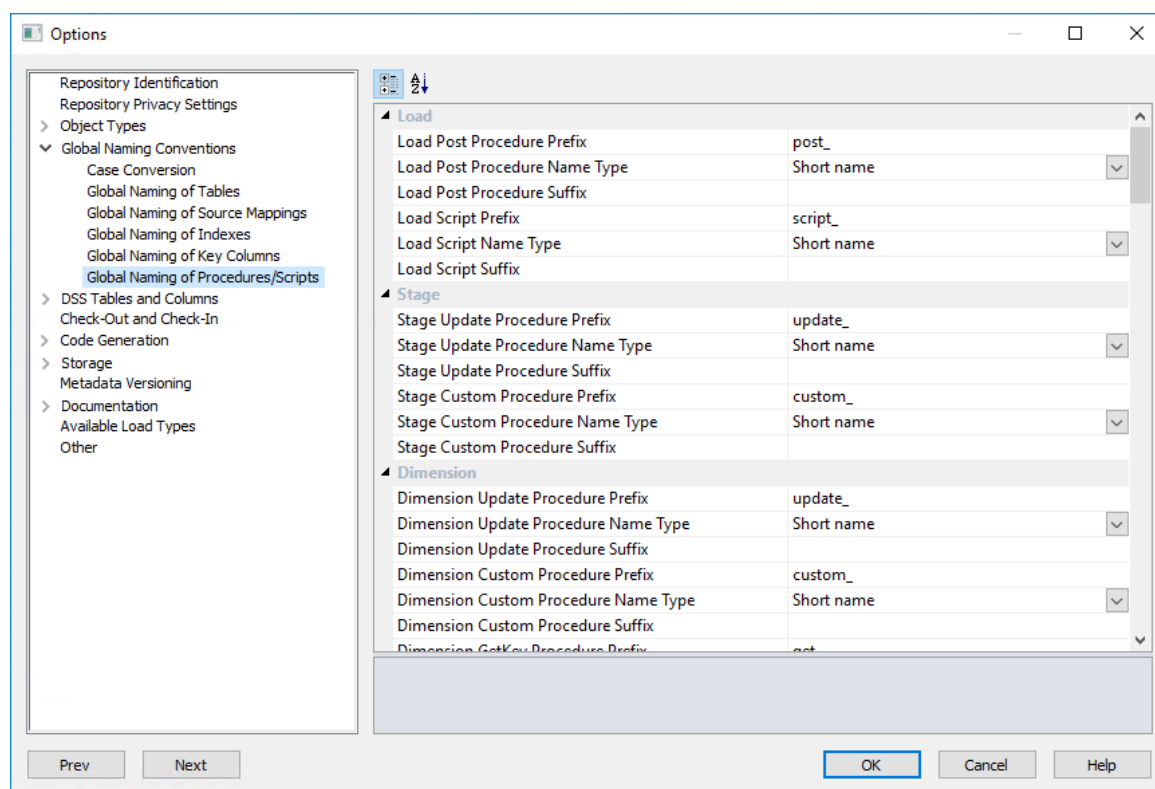
GLOBAL NAMING OF PROCEDURES

This option enables you to set the default naming conventions for generated procedures in RED.

A procedure is generated by selecting the **(Build Procedure...)** option from the **Update Procedure** drop-down list in the table's Properties screen.

Procedure name defaults

The values shown below are the default settings, which can be changed to meet the customer requirements. The only restriction is on the size of the resultant name, which is database dependent.



The contents of the **prefix** and **suffix** fields must contain characters that are valid in a database stored procedure name and should preferably not contain spaces.

The **Name Type** may be either the full table name or the unique short name assigned to each table. In the case of smaller table names, the short name is usually the same as the table name.

For example, if we have a stage table called `stage_product`, then from the example screen above, the two possible generated procedures would be called `update_stage_product` and `custom_stage_product`.

SETTINGS - DSS TABLES AND COLUMNS

When building the data warehouse, WhereScape RED makes use of a number of special tables and columns. Two tables are used, which are called **dss_source_system** and **dss_fact_table** by default.

The special columns used are defined in the table below.

Column name	Description
dss_batch	Not used at this stage.
dss_source_system_key	Added to support dimension tables that cannot be fully conformed, and the inclusion of subsequent source systems. See the section below for more details.
dss_fact_table_key	Used in composite rollup fact tables to identify the source fact table that contributed the particular row.
dss_create_time	Indicates when a record was created.
dss_update_time	Indicates when the record was last updated in the data warehouse. Used in the updating of rollup fact tables and aggregate tables.
dss_start_date	Used for dimension history tables as the start date for a particular version of a row.
dss_end_date	Used for dimension history tables as the end date for a particular version of a row.
dss_count	Applied to fact tables. Provides a simple row count variable that can be used by end user tools.
dss_current_flag	Used for dimension history tables. This flag identifies the current record where multiple versions exist.
dss_version	Used for model history tables. This column contains the version number of a history record. Numbered from 1 upwards with the highest number being the latest or current version. It forms part of the unique constraint for the business key of a dimension history table.
dss_file_name	Identifies a table holding files loaded into load tables.
dss_change_hash	Used to identify for a Satellite table. This column identifies the differences in the descriptive columns of a Satellite table which is used for generating the change hash key for creating a Satellite object.

All of these special columns can be renamed through the **Tools>Options>DSS Tables and Columns** menu option. All columns in the sample screen below (except for `dss_source_system`) can be renamed. The two tables however, require valid table names that meet certain criteria. See the appropriate sections below.

Note: When using table names other than the defaults for `dss_source_system`, it is worth considering the fact that by default the metadata backups will include any table that begins with "dss_". Therefore, if a table is used it is recommended that it have a name starting with "dss_". The advantages are that a working meta repository will be established through a backup and restore, if these tables are included in the backup set.

dss_source_system

This pseudo dimension table is designed to identify a data source for a dimension row. Its purpose is to handle changes in source systems. If its use is not desired (default) then leave this field blank.

For example:

An organization has a number of factories. These factories are referenced by all of the operational systems. The production system has its own code for each factory and this is the unique means of identifying the factory. The distribution system has a factory short name which it uses for the unique identifier. The raw materials system simply uses the factory name. It is probably not practical or even desirable to force these source systems to utilize a standard factory identification method, so instead we allow the dimension table to be non conformed. We do however, insist on a standard factory naming convention, so that our reports and queries will join information when the factory name is used.

In such an example, the `dss_source_system_key` is used to identify the source of the data for the dimension table row. It also adds to the unique business key, so that two source systems can utilize the same code to refer to different entities. This key also provides a degree of future proofing in the data warehouse, to assist in the possible changing of an underlying source system.

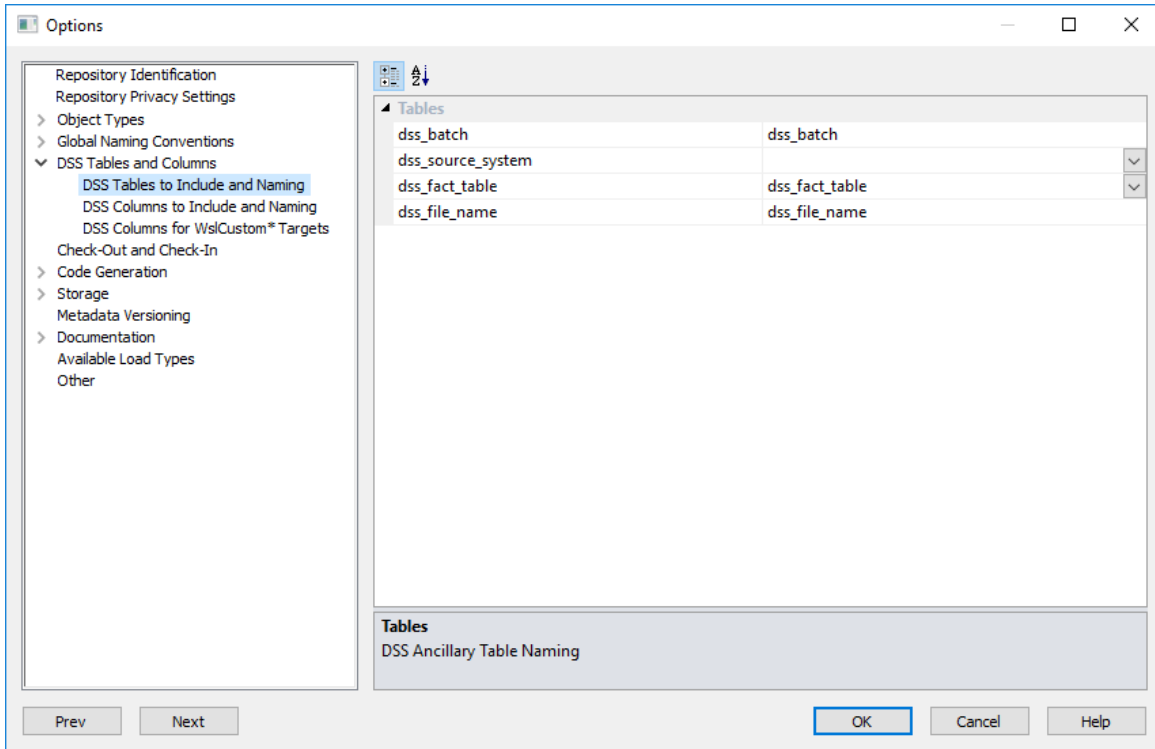
The generated procedure code will always set the key value of this table to 1. Therefore, manual code changes is required to make use of the functionality that this table offers.

If this table is to be given a different name then it and all its columns can be renamed or the following steps can be taken:

- 1 Create a new table by dragging the column `dss_source_system_name` from `dss_source_system` into a dimension target.
- 2 Rename the `dss_source_system_name` column to match the new table name.
- 3 Delete the last two columns.
- 4 Under the table Properties, change the table type to **Mapping table**. This prevents the table from being seen as a dimension table in the documentation.
- 5 Change the `dss_source_system` table name shown in the screen below, via the **Tools>Options>DSS Tables to Include and Naming** tab.

DSS TABLES

This option enables users to set the **DSS Tables**.

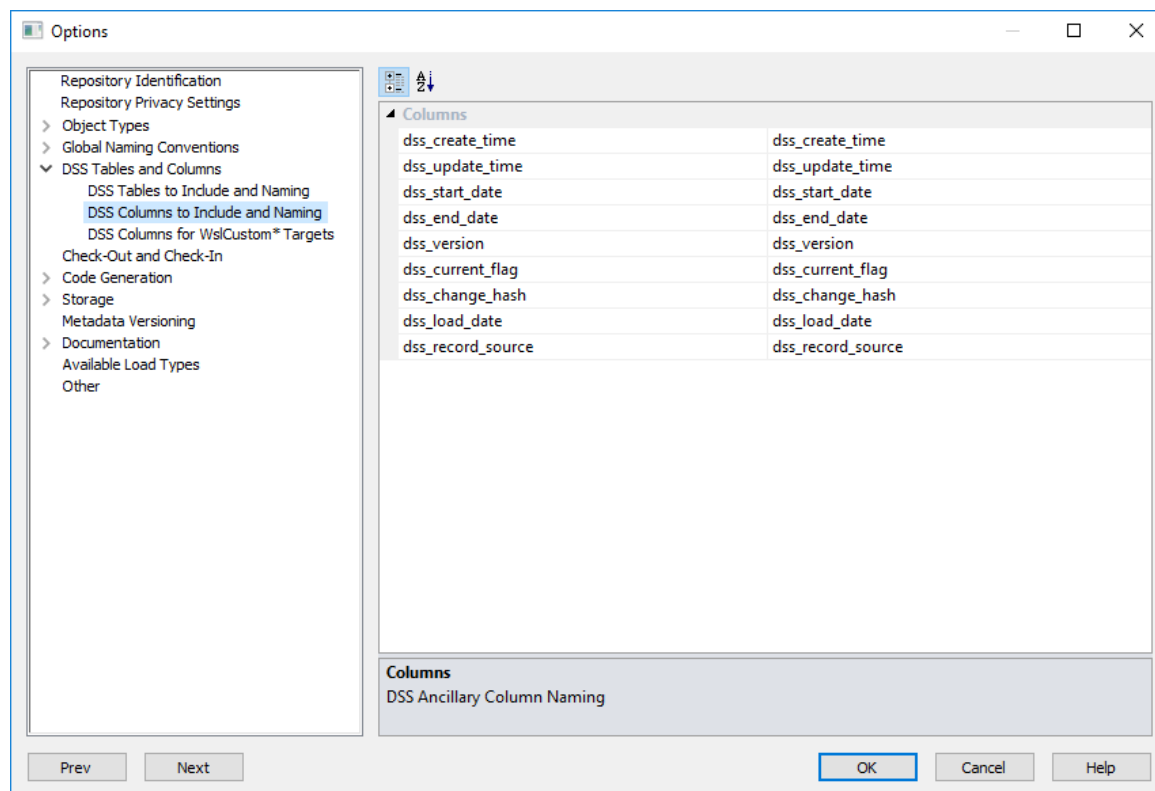


Tables

Set the **DSS Tables**.

DSS COLUMNS

This option enables you to set the **DSS Columns**.



Columns

Set the **DSS Columns**.

dss_create_time

Column added to all stage, ODS, EDW 3NF, dimension, fact and aggregate tables for information only. Leave the field blank to disable or add a name for the dss_create_time column, e.g. dss_create_time.

dss_update_time

Column added to all dimension and stage tables. It is required if the generated code for fact and aggregate tables is to be used.

dss_start_date

Column used for dimension history tables. It is used to identify when a dimension table row was replaced. This is a required field.

dss_end_date

Column used for dimension history tables. It is used to identify when a dimension table row was replaced. This is a required field.

dss_version

Column used for dimension history tables. Is it used to store the version of a dimension table row. This is required for unique constraints.

dss_current_flag

Column used for dimension history tables. It is used to identify the current dimension table row. This is a required field.

dss_change_hash

Column used to identify the differences in the descriptive columns of a Satellite table which is used for generating the change hash key for a Satellite object. Refer to ***Creating Data Vault Stage Tables*** (on page 427) for more details.

dss_load_date

Column used to store the date when the value in this row was loaded into the metadata repository. This column and the **dss_record_source** column below is added to new Load tables that have the option **Add meta data columns to table** selected. Refer to the ***Data Vaults chapter*** (see "***Data Vaults***" on page 413) for more details.

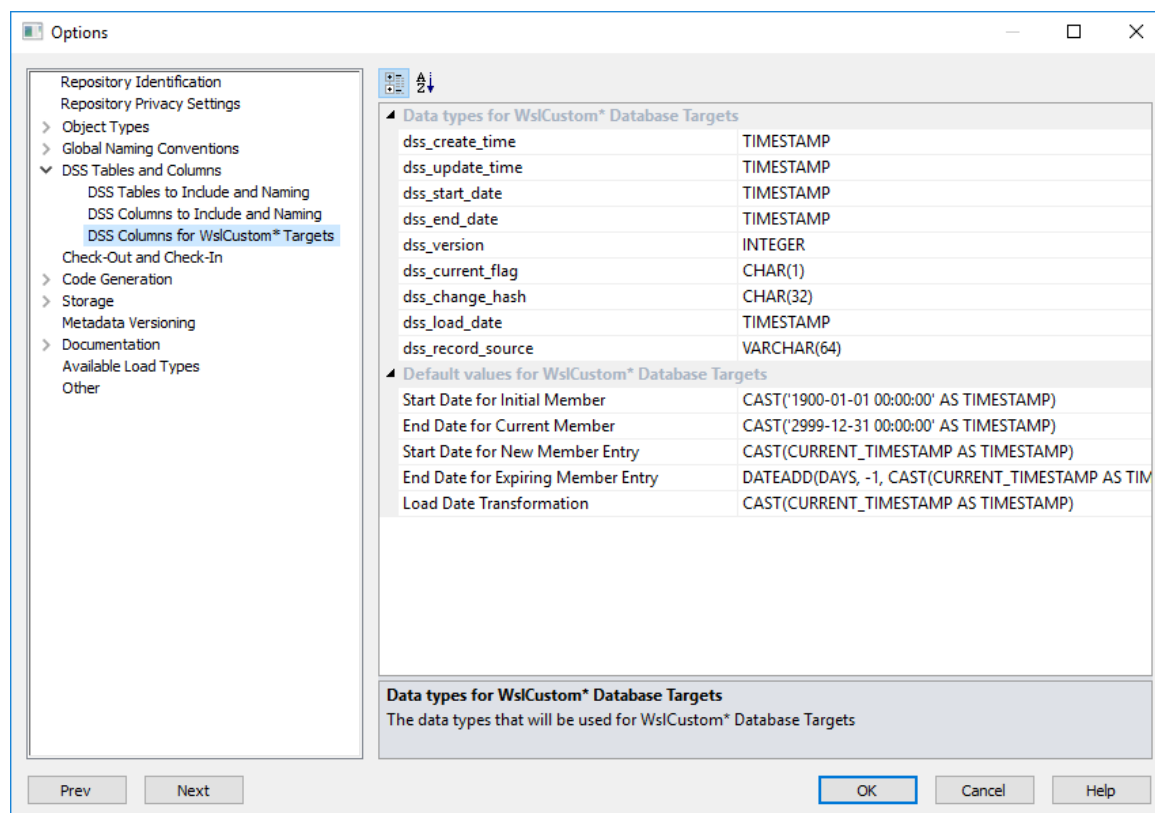
dss_record_source

Column used to store a descriptive term to identify the source of this record. This column and the **dss_load_date** column above is added to new load tables that have the option **Add meta data columns to table** selected. Refer to the ***Data Vaults chapter*** (see "***Data Vaults***" on page 413) for more details.

DSS COLUMNS FOR CUSTOM TARGETS

Note: These settings are only available/displayed if your WhereScape RED license includes support for custom target databases.

This option enables you to set the **DSS Columns** for custom database targets.



Data types for Custom Database Targets

Enables you to set the data types that are used in the DSS columns added to new table objects.

Notes:

1. The **Hash Key Generation wizard** (on page 416) creates new hub and link hash keys in a Data Vault Stage table with the same data type defined for the **dss_change_hash** column, if the table is in a Custom database target.
2. If the table is *not* in a Custom database target, the data type CHAR(32) is used.
3. Changing this data type has no impact on any existing hash key columns.

Default values for Custom Database Targets

Enables you to set the default values that are used when generating update procedures for Slowly Changing Dimension tables and Load tables.

Start Date for Initial Member

Default date function to use for the start date of rows added to a slowly changing dimension table.

End Date for Current Member

Default date function to use for the end date of current rows in a slowly changing dimension table.

Start Date for New Member Entry

Default date function to use for the start date of current rows in a slowly changing dimension table.

End Date for Expiring Member Entry

Default date function to use for the end date of expiring rows in a slowly changing dimension table.

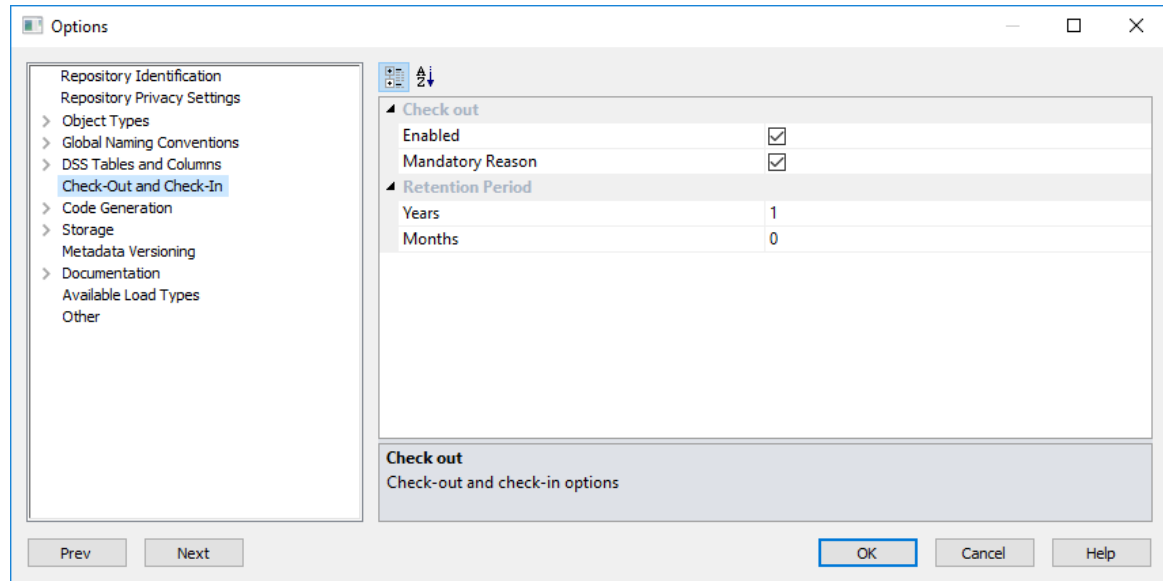
Note: Refer to the Change Detection section of *Generating the Dimension Update Procedure* (on page 303) for more information.

Load Date Transformation

Default date function to use for populating the new records inserted in load tables.

SETTINGS - CHECK-OUT AND CHECK-IN

This option enables users to set up for the **Check-out** or **Check-In** of Procedures.



Check out

Enabled: This option is selected by default to enable procedures to be checked-in or checked-out.

Mandatory Reason: Select this option if a reason is mandatory for checking out or checking in procedures, else deselect.

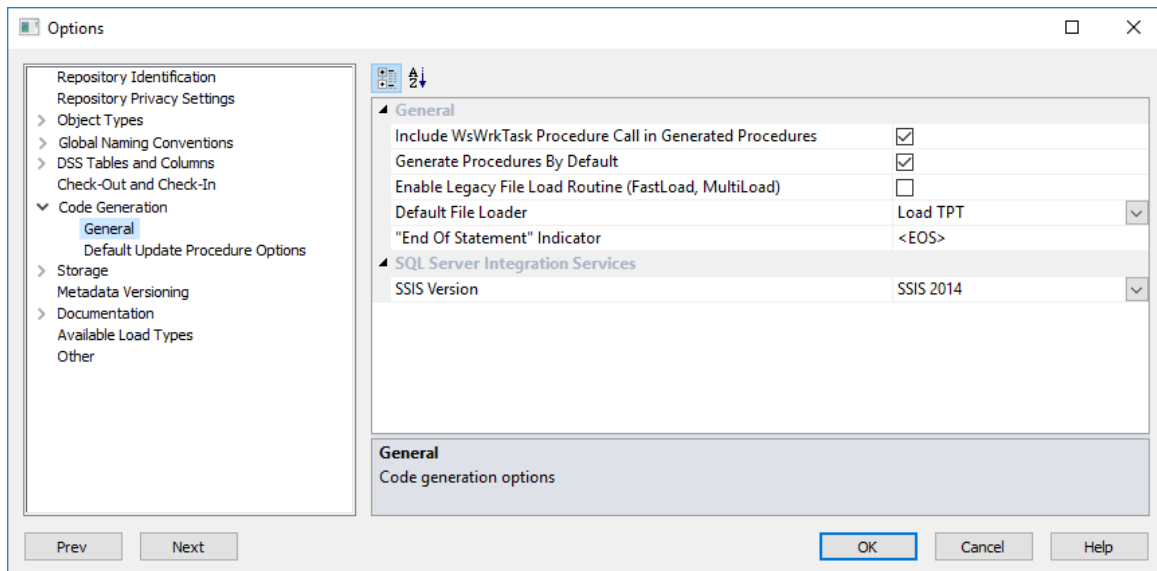
Retention Period

Set the length of time; **Years** and **Months**, for which procedures may be checked-out.

CODE GENERATION

GENERAL

This option enables users to set some general **Code Generation** settings.

**General****Include WsWrkTask Procedure**

This option is selected by default and results in a call to the WsWrkTask function being placed at the end of most of the generated update procedures. These calls to WsWrkTask result in counters being set in the meta table `ws_wrk_task_log`. These counters can be viewed via a query on the view `ws_admin_v_task`.

Generate Procedures By Default

Set this option to generate Procedures by default.

Enable Legacy File Load Routine (FastLoad, MultitLoad)

Set this option to enable FastLoad and MultiLoad legacy file load routines in RED. This option is not selected by default, e.g. FastLoad and Multitload options are not available in the **Default File Loader** drop-down field.

Note: This setting also controls the options available in the **File Load Routine** drop-down list in the **Source** tab of the Load table Properties screen.

Default File Loader

Options for the default file loader are:

- Load TPT (default)
- Update TPT
- Stream TPT
- No Load

Notes:

-When importing a Model from 3D to RED, select **Load TPT** instead of Fastload as the Default File Loader method. FastLoad is not a valid option for loading Linux files to Teradata.

- The **Enable Legacy File Load Routine** option is not enabled by default. The **Default File Loader** is set to **Load TPT** by default.

"End of Statement" Indicator

Set the indicator to separate multiple SQL statements in a SQL block. If left blank the default value of <EOS> is used.

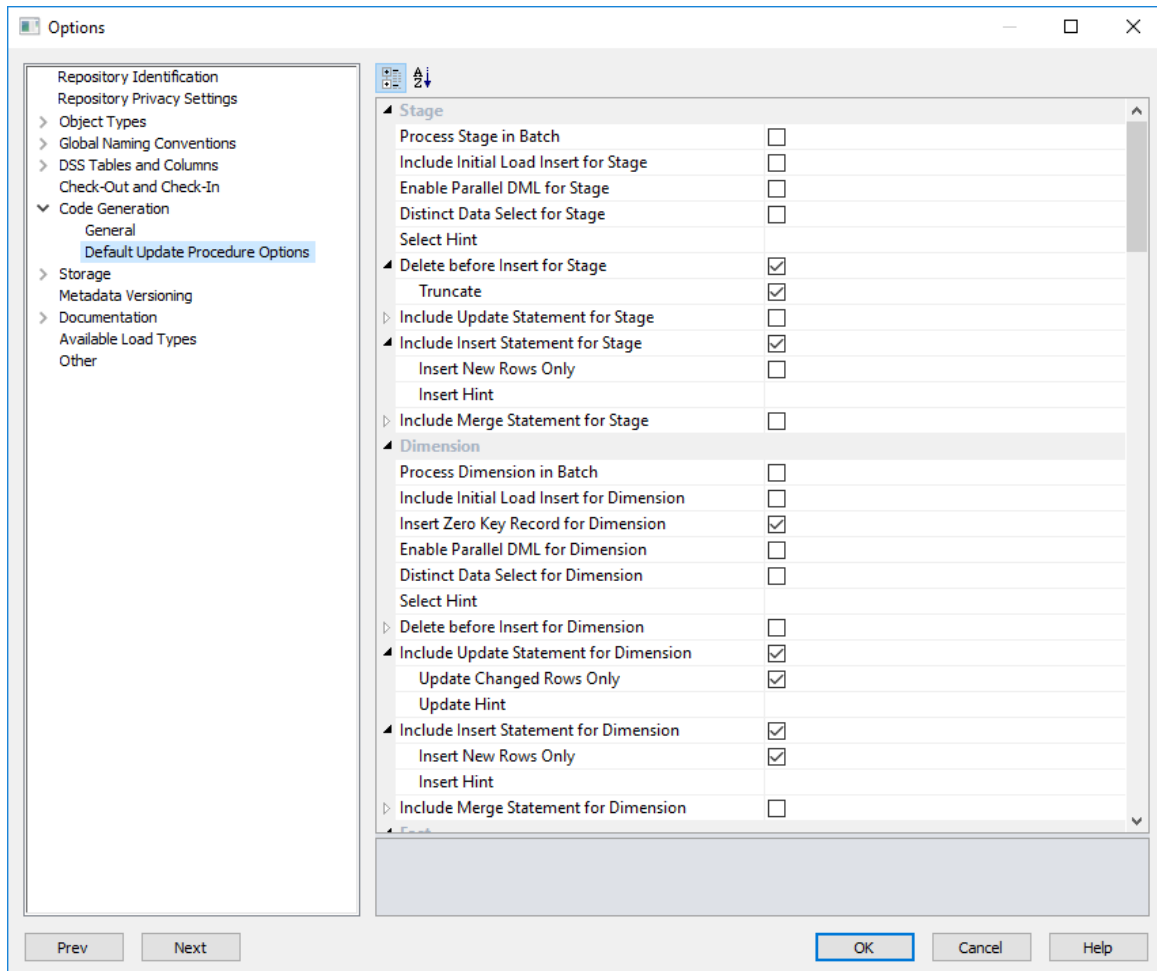
SQL Server Integration Services

SSIS Version

Available version of SQL Server Integration Services. SSIS is not enabled by default for Teradata. To use SSIS to load data, the relevant version of SSIS needs to be selected on this drop-down list.

DEFAULT UPDATE PROCEDURE OPTIONS

This option enables you to set **default update procedure** settings for specific object types.



Process in Batch - enables users to select a column to drive data processing in a loop based on the distinct ordered values of the selected column.

Include Initial Load Insert - enables users to include an additional insert statement to the update procedure. If the target table is empty, the new insert statement is run in place of the standard generated code.

Insert Zero Key Record - enables users to add an insert statement for an unknown record with an artificial key of zero. Only applicable to tables with an artificial key.

Parallel DML - enables users to add all code required to the update procedure for enabling Oracle parallel inserts. **Note:** Oracle only.

Distinct Data Select - enables users to ensure duplicate rows are not added to the table.

Select Hint - enable users to enter a database-compliant hint to be used in the SELECT statement. Parameters \$TABLE\$ and \$INDEX\$ is automatically replaced at procedure generation time.

Delete before Insert - enables users to add a delete statement to the update procedure before any update or insert statement.

Truncate - performs a DDL operation to delete all records from a table, rather than a predicate based DML operation to delete individual rows. This option is automatically enabled when **Delete before Insert** is enabled, but can be deactivated separately.

Include Update Statement - enables users to include an update statement in the procedure to update changing rows in the table.

Update Changed Rows Only - enables users to use change detection to work out what rows require updating.

Update Hint - enables users to enter a database hint to be used in the UPDATE statement.

Include Insert Statement - enables users to include an insert statement in the procedure to insert new rows in the table.

Insert New Rows Only - enables users to use change detection to work out what rows require inserting.

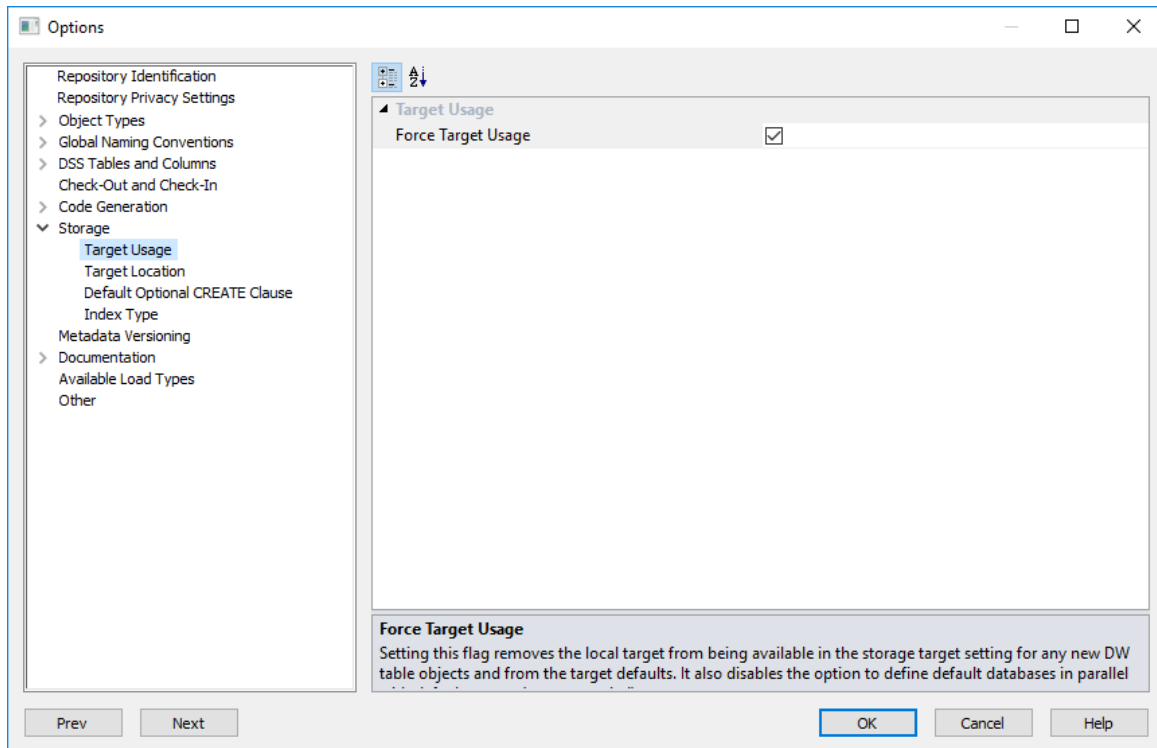
Insert Hint - enables users to enter a database hint to be used in the INSERT statement.

Include Merge Statement - enables users to include a merge statement in the procedure to merge new/changed rows in the table.

SETTINGS - STORAGE

TARGET USAGE

This option enables you to force use table target locations defined in the Connection Properties screen, and prevents the use of legacy local targets when creating new table objects in RED. This option is enabled by default for new repositories—it is recommended to define targets for the data warehouse objects that is separate from the metadata.

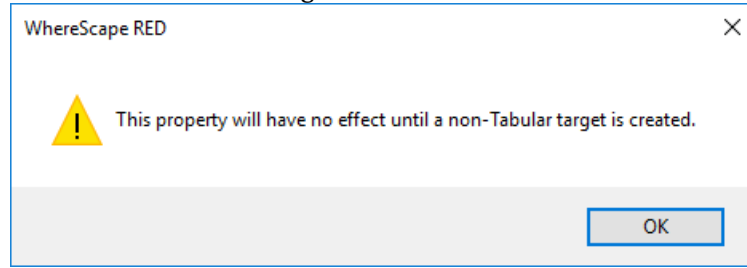


The **Force Target Usage** setting removes the option to use local targets in the **Default Target** drop-down list of the **Target Location** (on page 113) option and also hides the **Table Storage** (on page 116) option, which is used to set the storage locations for each table object type created in RED.

Notes:

- Existing RED objects or objects loaded into RED with local set as their target retains the local target setting.
- If the target location is changed from local to a specific target, then there is no option to set it back to local without deactivating the Force Target Usage setting.
- If table target(s) has not been defined, a message is displayed to warn users that this setting will

have no effect until a target is created.



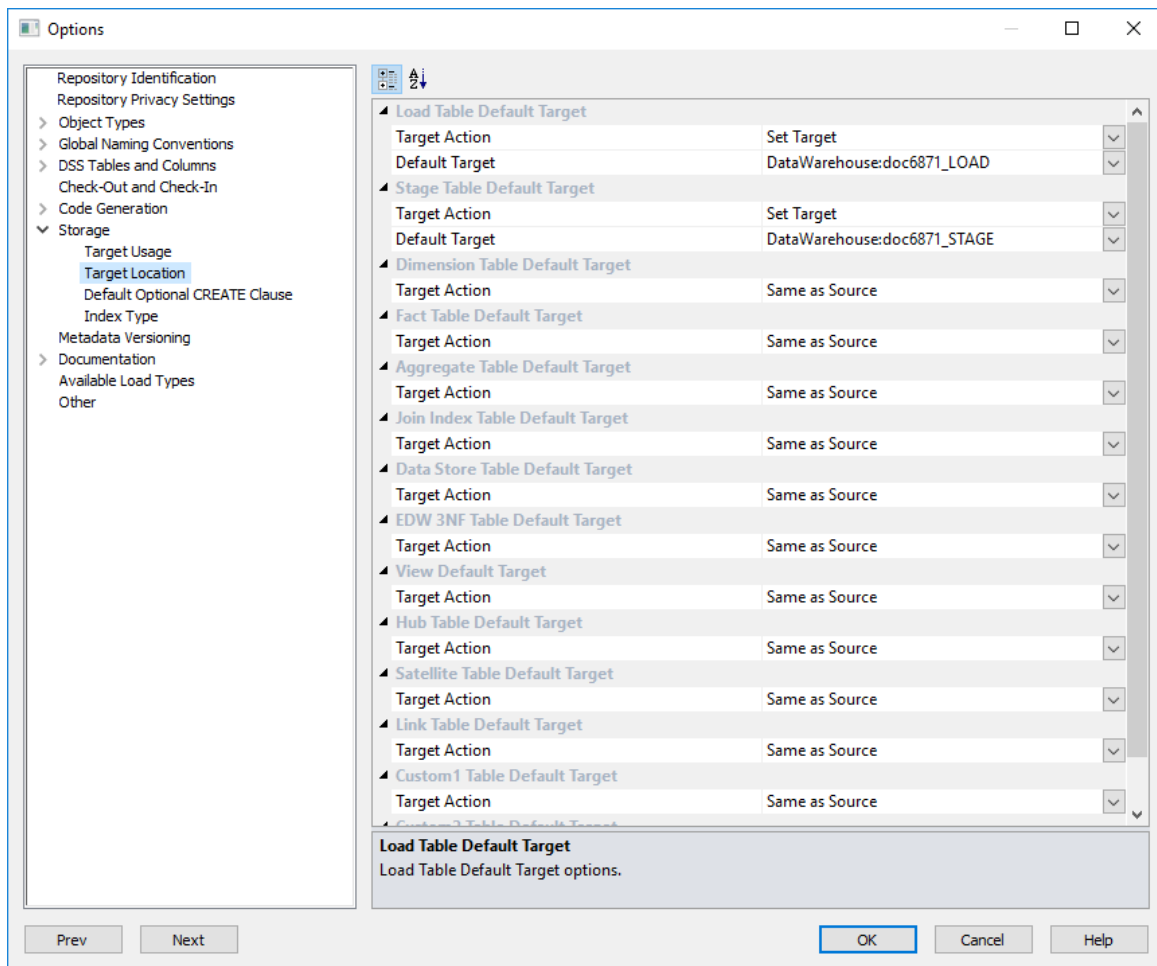
To see more about creating target locations see *Connection to the Data Warehouse* (see "Database - Data Warehouse/Metadata Repository" on page 147).

TARGET LOCATION

Target Location options enables users that are placing objects across multiple databases to set default target locations for new tables.

Default table target locations can be set for the following objects:

- **Load**
- **Stage**
- **Dimension**
- **Kpi Fact**
- **Fact**
- **Aggregate**
- **Join Index**
- **Data Store**
- **EDW 3NF**
- **View**
- **Hub Table**
- **Satellite**
- **Link**
- **Custom**



Target Action

Set Target

This option enables you to set a default target location for new tables to be created. It enables the **Default Target** drop-down list where a specific target location for new tables can be defined.

Note: If the **Force Target Usage** setting is enabled in the **Target Usage** (on page 112) option, selecting **Set Target** automatically sets the **Default Target** field with the first target location value available from drop-down list. The option to use local target is also removed.

Same as Source

This option is selected if the table's default storage must be the same as the original source, where the table is coming from. This option cannot be selected for Load tables.

Default Target

A default target location can only be entered if the **Set Target** action has been selected in the **Target Action** drop-down list.

With this option, users can choose between setting a table's default location:

- **(local)** if the **Force Target Usage** setting is **not** enabled in the **Target Usage** (on page 112) option, or
- to any other **target locations** that have been defined in the relevant connection Properties.

To set a default target location on a table by table basis:

- 1 Select the **Set Target>Same as Source** option from the **Target Action** drop-down list.
- 2 To have tables located in a specific target location, select a **default target** where the new object should be placed as the object is dragged and dropped to the middle work pane.

To see more about creating target locations see **Connection to the Data Warehouse** (see "**Database - Data Warehouse/Metadata Repository**" on page 147).

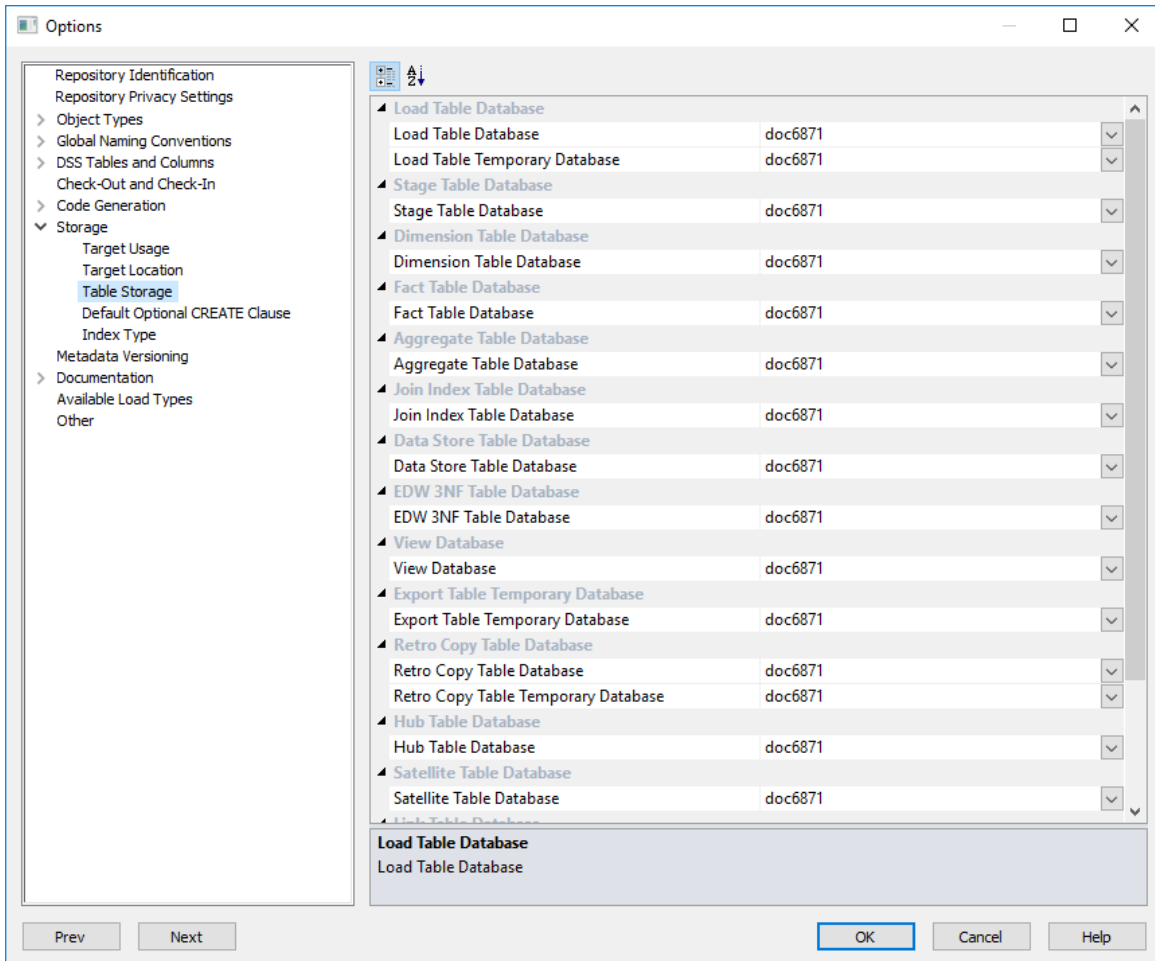
Even though the default target location can be set in the **Target Location** options, this setting can also be changed after the table has been created via the **Storage** tab of each table's Properties screen.

To see more information about changing the target location after a table has been created, see **Storage** (on page 202).

TABLE STORAGE

Note: The **Table Storage** option is only available/visible if the *Target Usage > Force Target Usage* (see "*Target Usage*" on page 112) setting is **not** enabled.

This option enables you to set the **Storage** locations for each table object type created in RED.



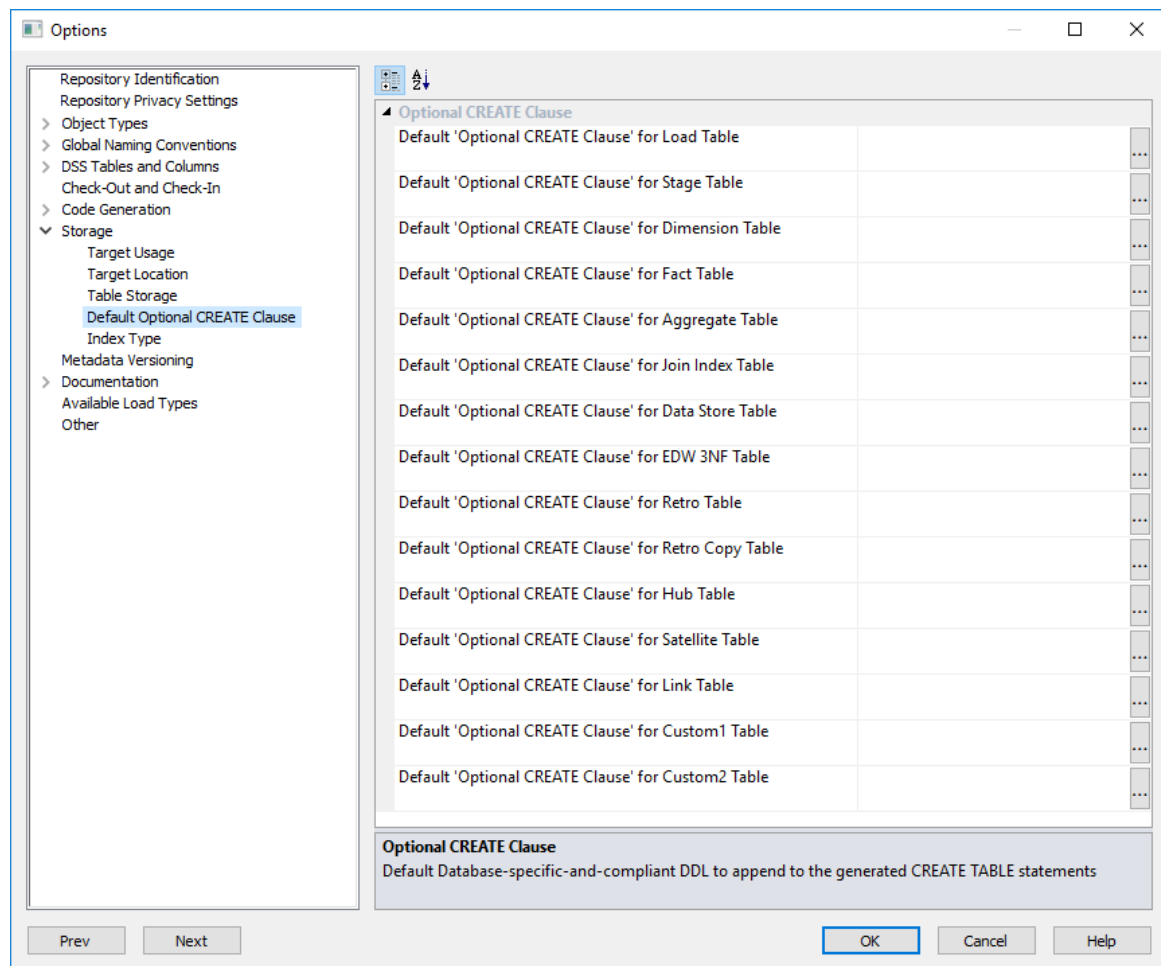
Set the **Storage** locations for each table type.

These defaults are applied when a table is created. They can be changed by selecting the **Storage** tab on the Properties screen of a table.

DEFAULT OPTIONAL CREATE CLAUSE

This option enables you to define a default value for the "Optional CREATE Clause" property of each object type, which is populated when the object is first created.

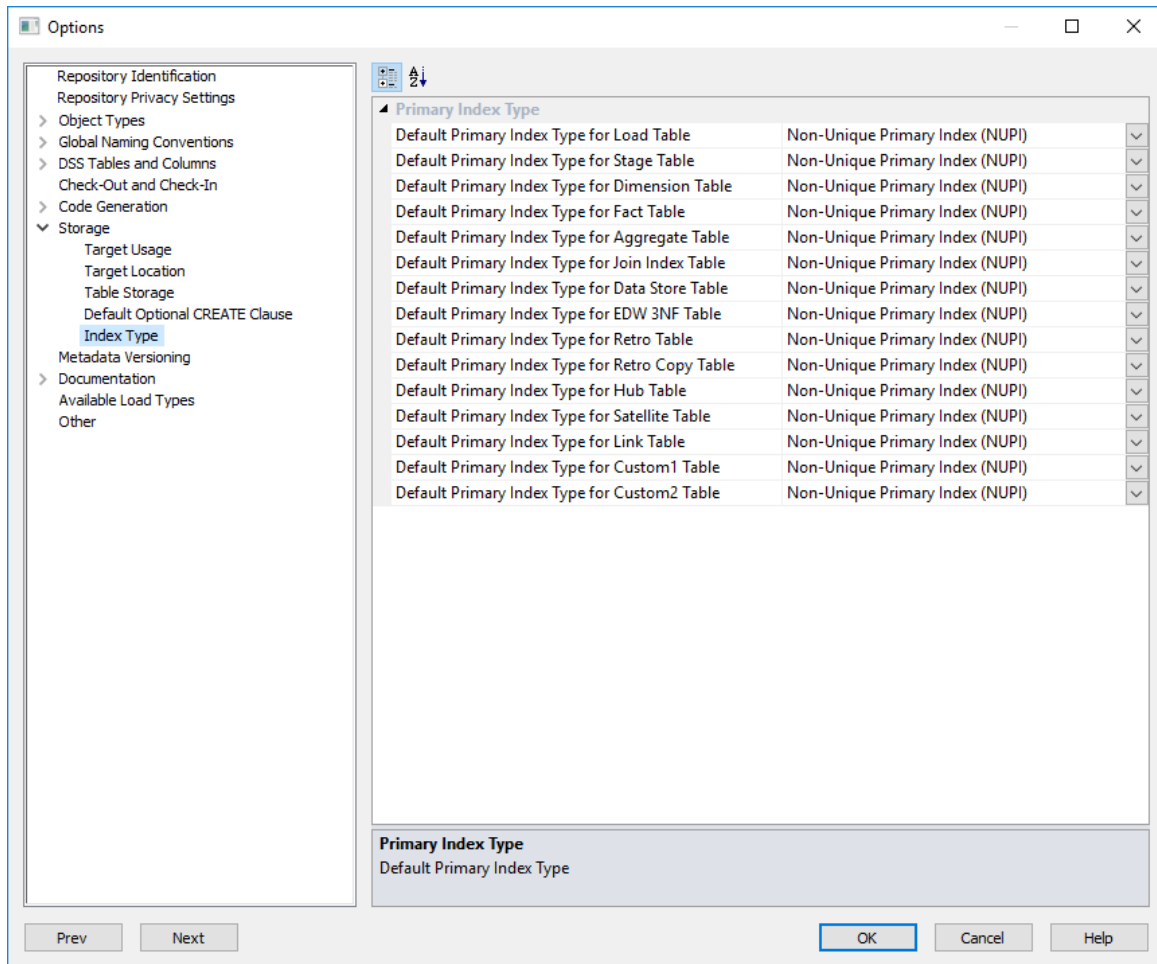
The Optional CREATE Clause text is appended to the DDL CREATE statement when the table is generated.



TIP: This option is only used to set the default optional create clause for new objects. To edit the Optional CREATE Clause of an **existing** object or edit the clause on a table by table basis, go to the object's **Properties** screen, click the **Storage** tab and edit the **Optional CREATE Clause** field.

INDEX TYPE

This option enables you to set the default type of **primary index type** for each table type.



Set the **default primary index type** for each table type.

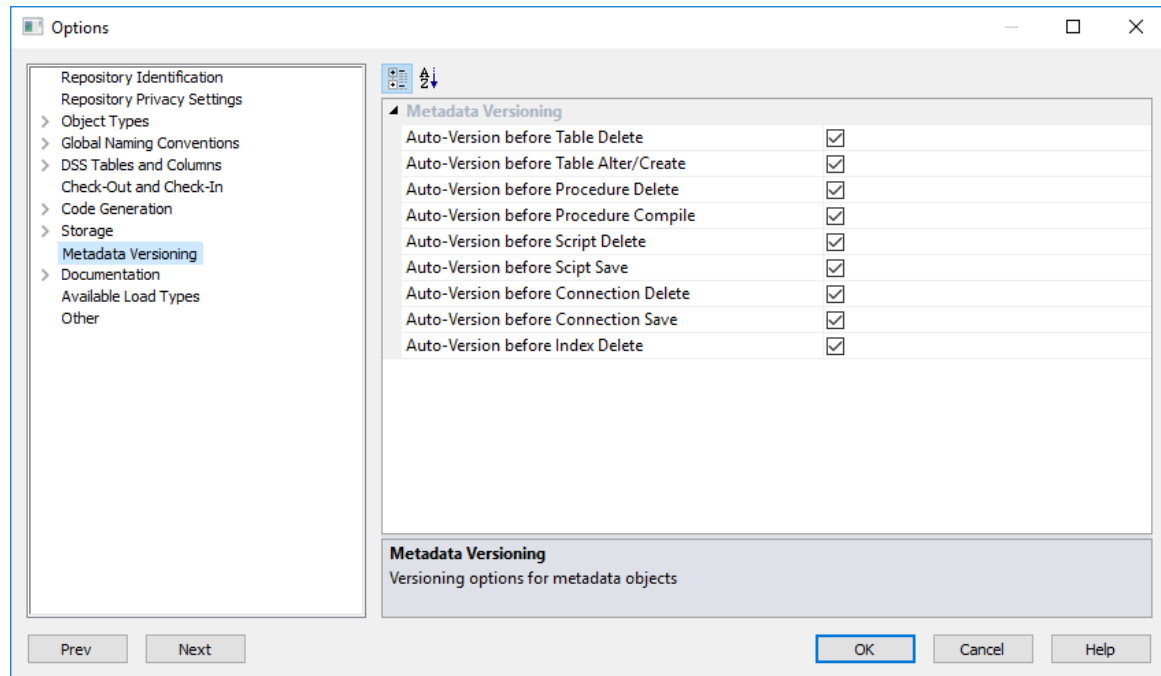
The options are:

- Non-Unique Primary Index (NUPI)
- Unique Primary Index (UPI)
- No Primary Index (NOPI)

These defaults are applied when an index definition is created. They can be changed by selecting the **Storage** tab on the Properties screen of an index.

SETTINGS - METADATA VERSIONING

This option enables you to alter the **Metadata Versioning** settings.

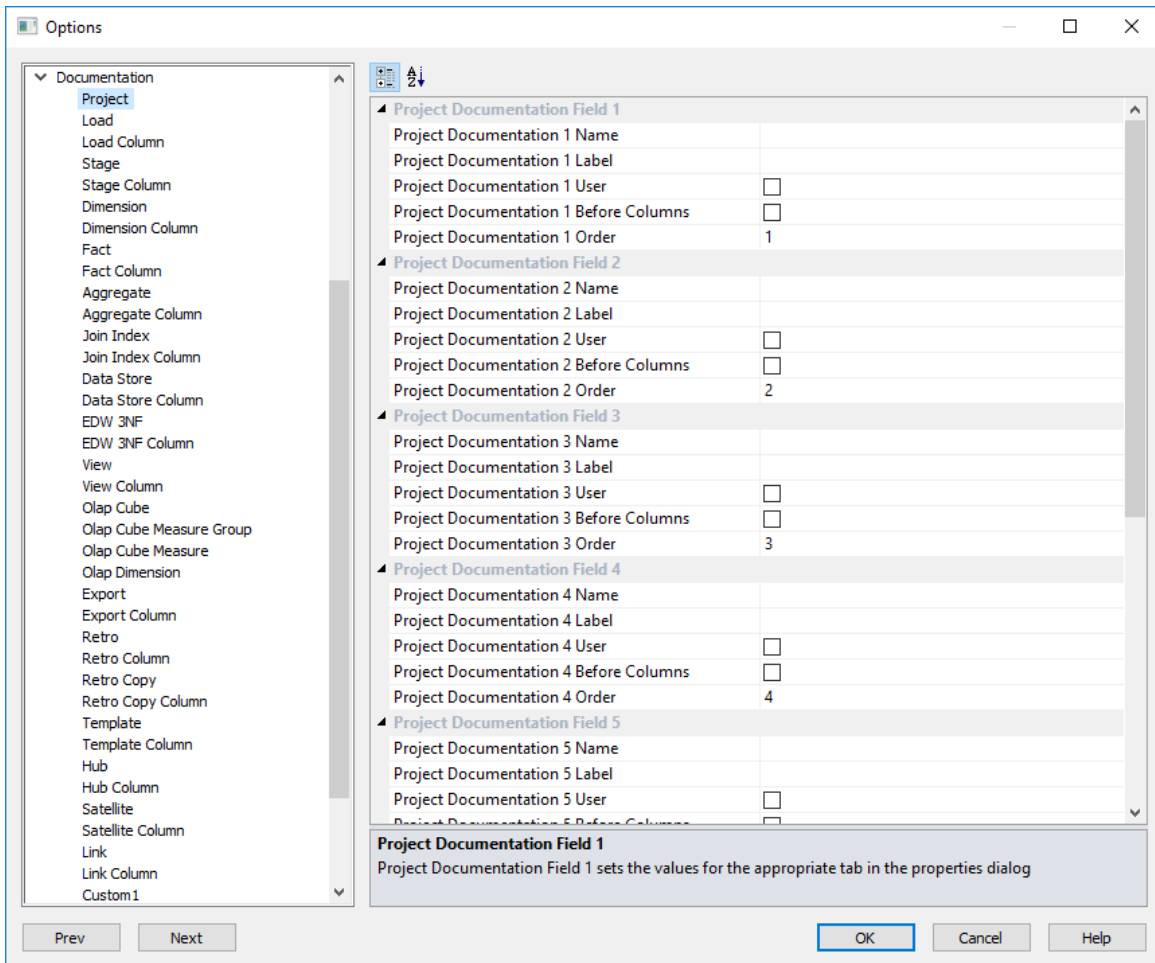


Metadata Versioning

Select or clear the corresponding check box to set when to auto-version the metadata. All settings are selected by default.

SETTINGS - DOCUMENTATION

This options enables you to alter the **documentation** settings.



The **Documentation Name** sets the name of the appropriate tab in the properties dialog.

The **Documentation Label** sets the the label or description of the appropriate documentation tab.

The **Documentation User** defines if the documentation information is visible to end users and included in end user documentation.

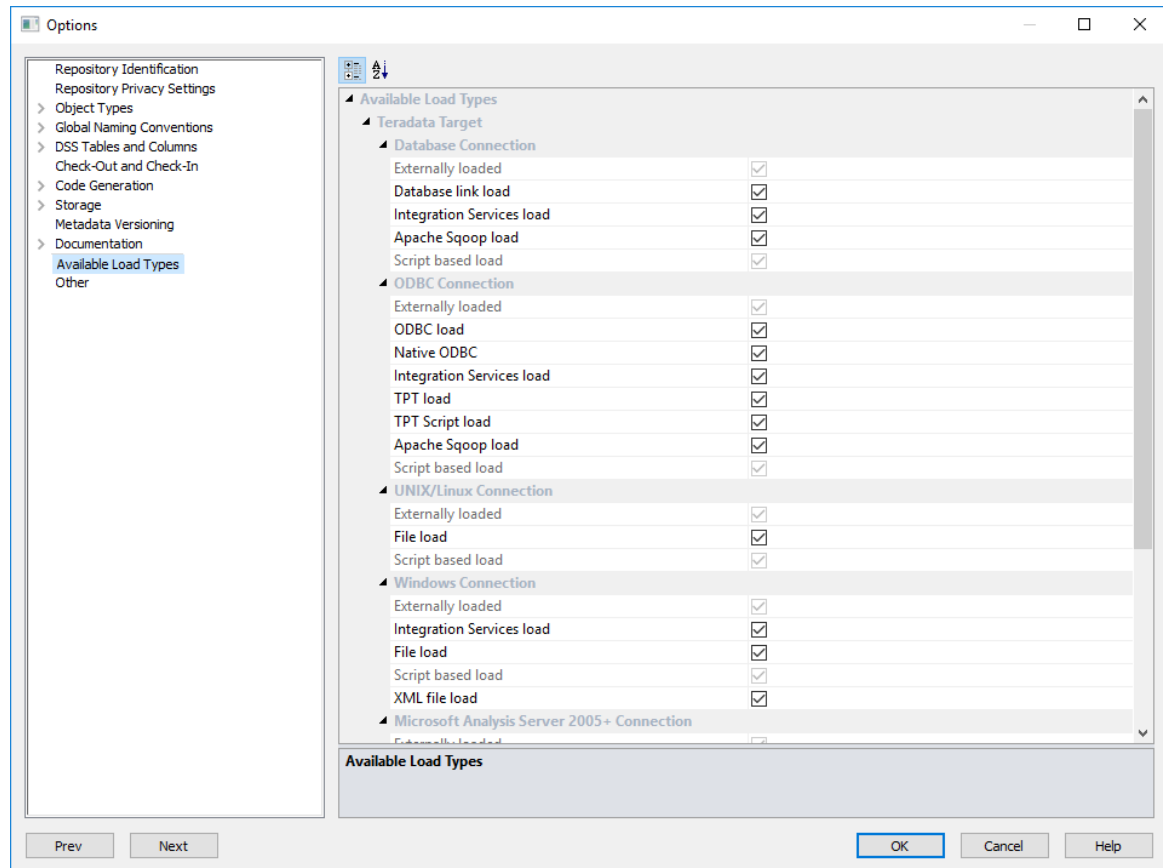
The **Documentation Before Columns** defines if the documentation tab information is shown in the documentation before or after the column information.

The **Documentation Order** defines the field order on the Properties screen tabs.

SETTINGS - AVAILABLE LOAD TYPES

This option enables you to activate or deactivate the various load types available in RED, based on Target and Source database.

Select from the combination of load types supported by target and source databases, to control the default load type options that can be set by users from the Connection and the Load table Properties screens.

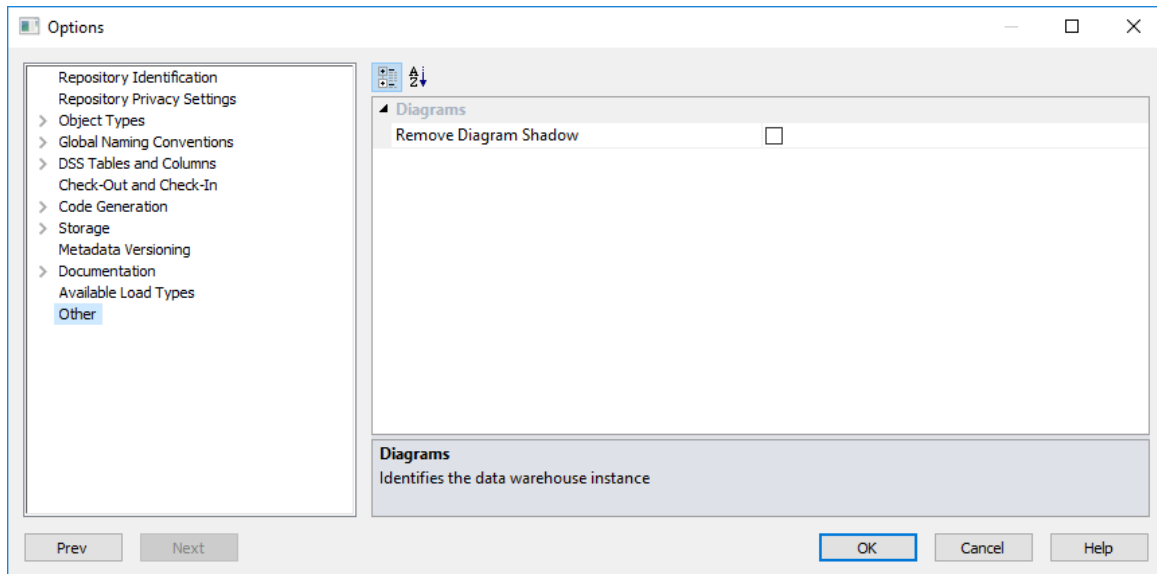


Note: The following load types are **not** enabled by default for newly created Teradata repositories:

- Integration Services (SSIS) load in Database, ODBC and Windows connections.
- Native ODBC load in an ODBC connection.
- XML file load in a Windows connection.

SETTINGS - OTHER

This option enables you to add or remove shadows in the diagrams.

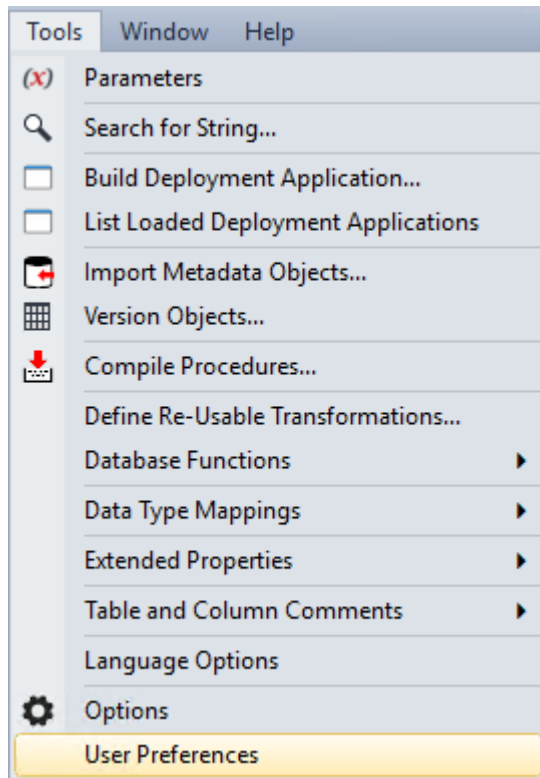


Remove Diagram Shadow

Set to prevent a shadow appearing on all printable diagrams produced in the diagrammatic window.

SETTINGS - USER PREFERENCES

Select **User Preferences** from the **Tools** menu.

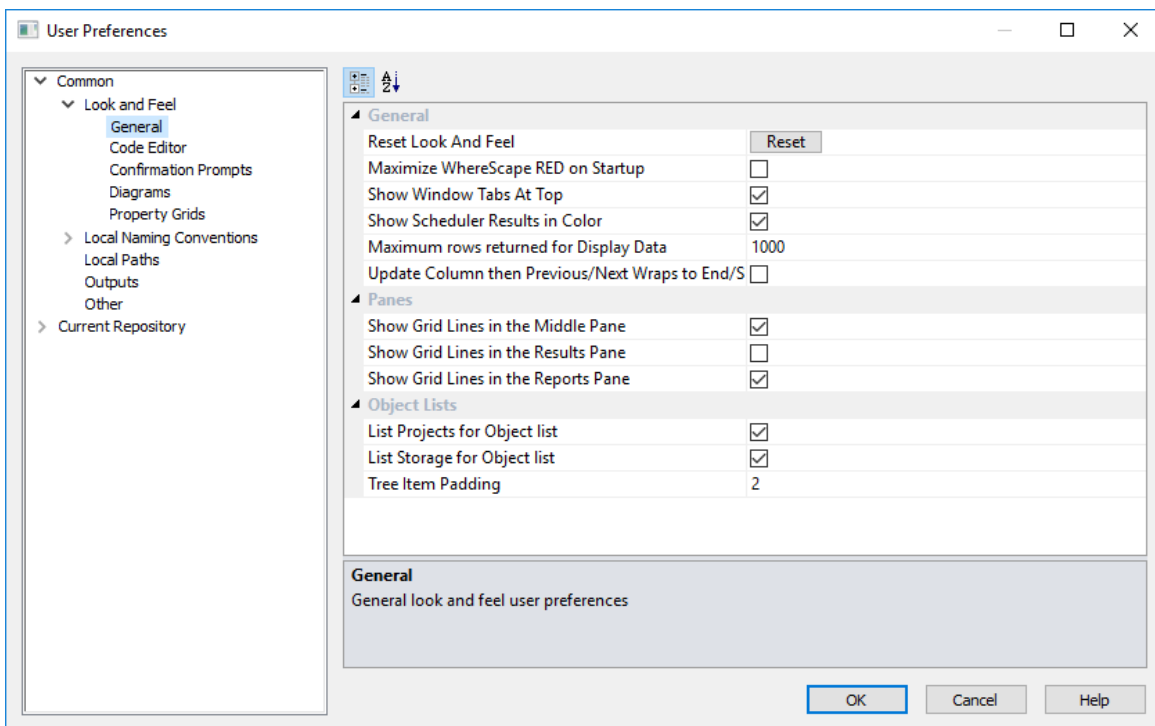


SETTINGS - COMMON

LOOK AND FEEL

General

This option enables you to set the **look and feel in general**.



General

The **Reset Look And Feel** option enables you to reset all window tab positions for Builder and Scheduler panes. Reset scheduler and report headings.

The **Maximize WhereScape RED on Startup** option starts WhereScape RED in full screen mode.

The **Show Window Tabs At Top** option displays the window tabs at the top of the screen.

The **Scheduler Results in Color** option turns on job status color coding in the scheduler.

The **Maximum rows returned for Display Data** option sets the maximum number of rows that are returned when displaying data.

The **Update Column then Previous/Next Wraps to End/Start** option controls the behavior of the directional **Update** buttons on the Column Properties dialogs. When enabled, the '<-Update' button will wrap to the last column when it moves beyond the first column; and the 'Update ->' button will wrap to the first column when it moves beyond the last column. When disabled (default), the dialog closes after an attempt to navigate before the first column or after the last column.

Panes

The **Show Grid Lines in the Middle Pane** option when selected, shows grid lines in the main work area.

The **Show Grid Lines in the Results Pane** option when selected, shows grid lines in the results area.

The **Show Grid Lines in the Reports Pane** option when selected, shows grid lines in the reports area.

Object Lists

The **List Projects for Object list** option when selected, shows the projects for each object in the middle pane object list.

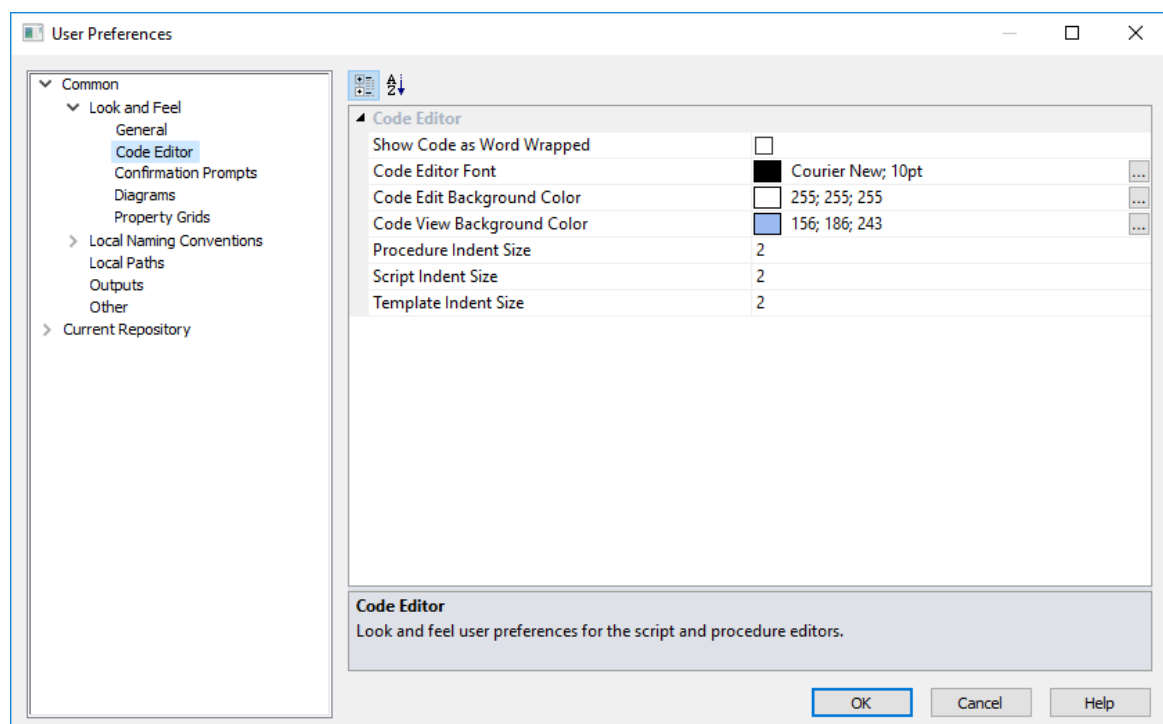
The **List Storage for Object list** option when selected, shows the storage for each object in the middle pane object list.

The **Tree Item Padding** option enables you to select the number of pixels used to pad tree items when the tree items represent an Object; for example, in the Object Pane and the Browser Pane. Padding is added to the top and bottom of each tree item. Padding can be set from 0-10 pixels, default value is 2.

Warning: Both these options impact on the speed lists are generated. Since they are enabled by default, both options can be deactivated to speed up the process or if considered irrelevant according to user's preferences.

Code Editor

This option enables you to set the **look and feel in code editor**.



Code Editor

The **Show Code as Word Wrapped** option enables you to have word wrapping applied to code by default.

The **Code Editor Font** option enables you to select the font used in code editors.

The **Code Editor Background Color** option enables you to select the background color when editing code.

The **Code View Background Color** option enables you to select the background color when viewing code.

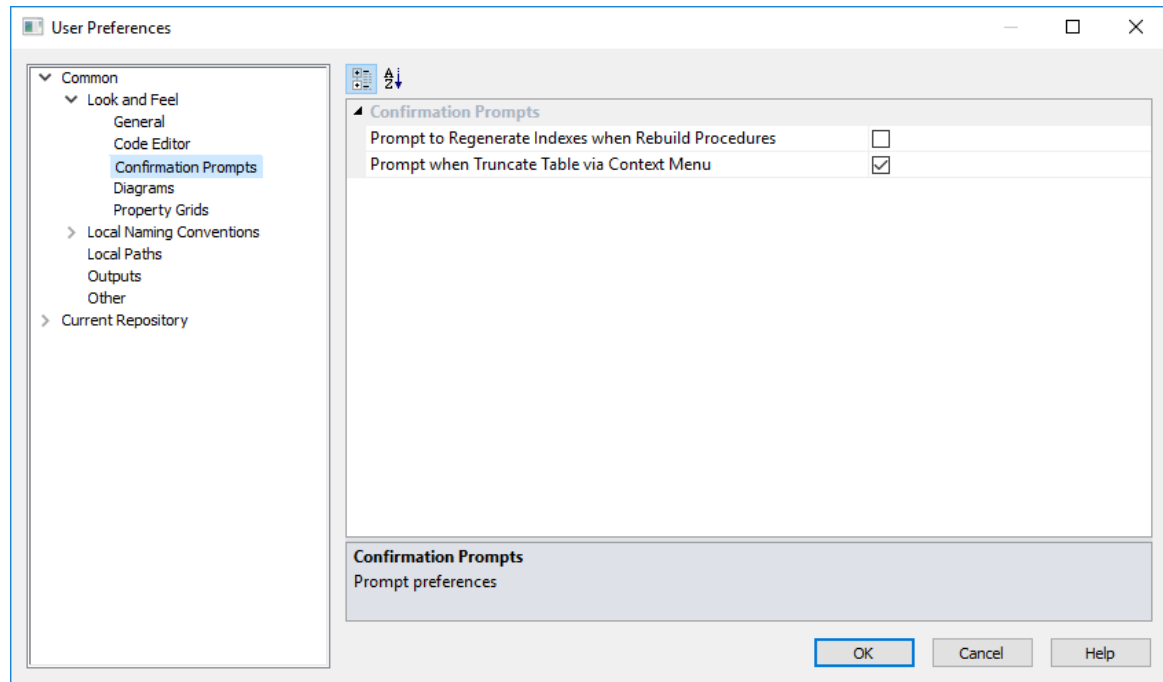
The **Procedure Indent Size** option enables you to specify the number of spaces that are generated when a TAB character is used within the Procedure editor. Permitted range is 2 through 10.

The **Script Indent Size** option enables you to specify the number of spaces that are generated when a TAB character is used within the Script editor. Permitted range is 2 through 10.

The **Template Indent Size** option enables you to specify the number of spaces that are generated when a TAB character is used within the Template editor. Permitted range is 2 through 10.

Confirmation Prompts

This option enables you to set the **look and feel** in confirmation prompts.

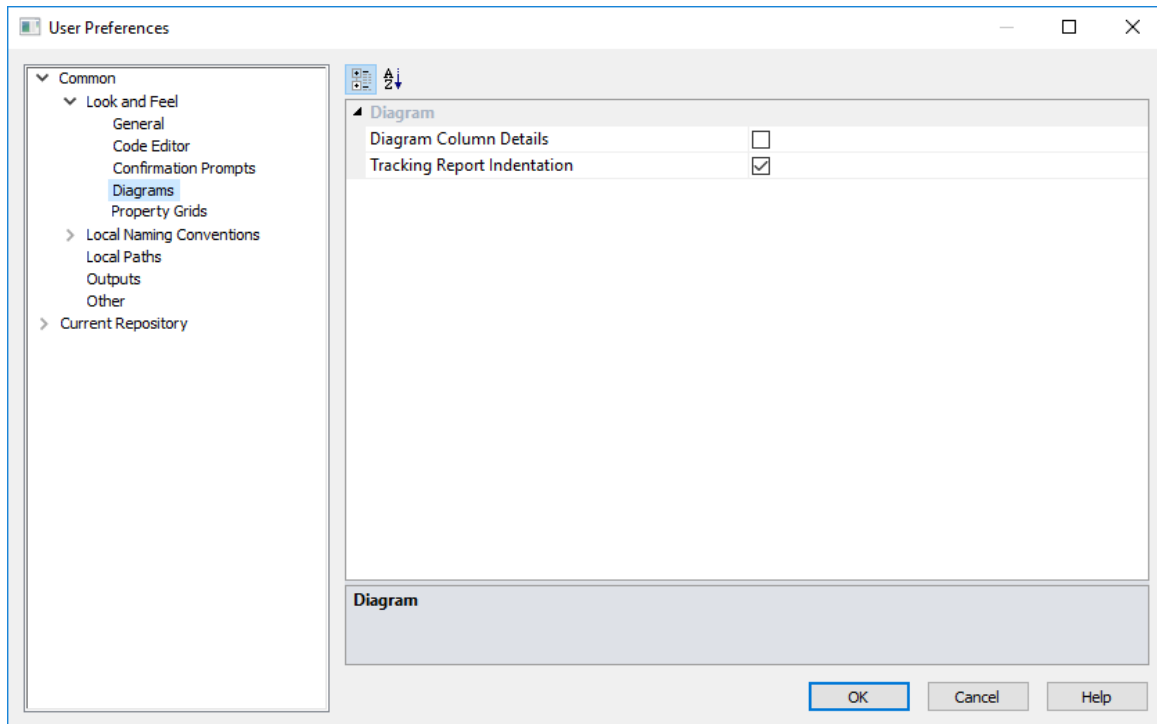


Confirmation Prompts

- **Prompt to Regenerate Indexes when Rebuild Procedures** - If set, always prompts for index regeneration whenever an update procedure is rebuilt.
- **Prompt when Truncate Table via Context Menu** - If set, always pops up a confirmation message before the truncate command is executed.

Diagrams

This option enables you to set the **look and feel** in the **diagrams**.



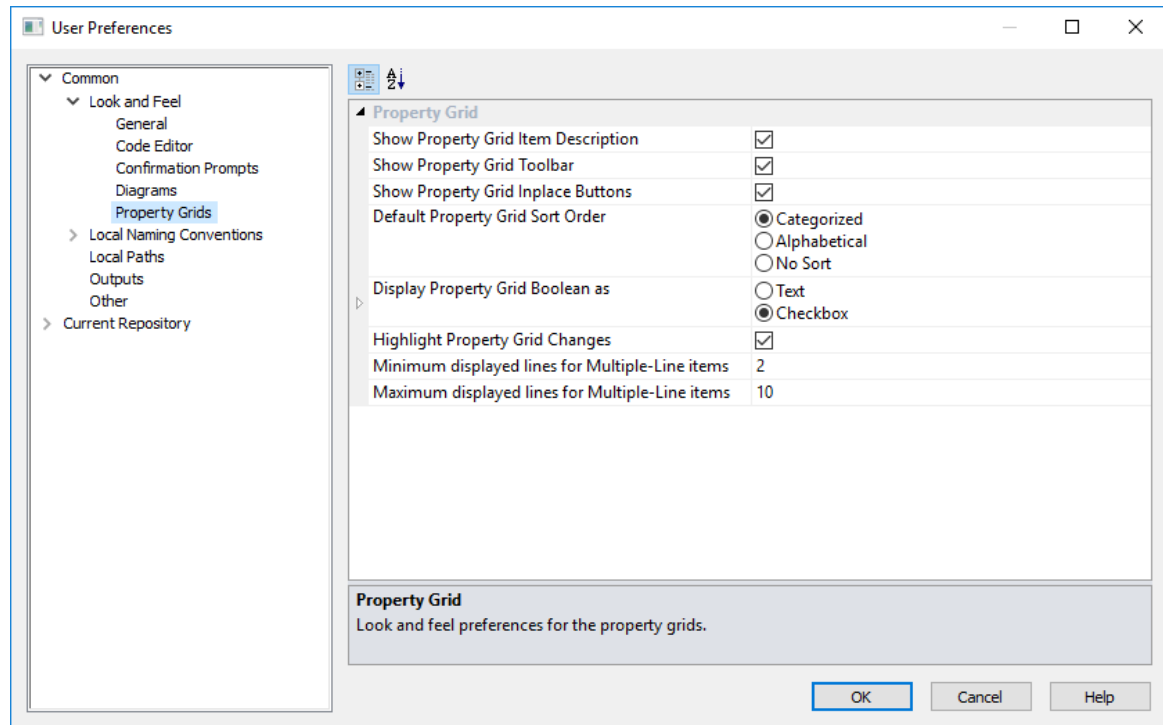
Diagram

The **Diagram Column Details** option shows the columns as the initial diagram.

When set, the **Tracking Report Indentation** output includes tabs to show dependency level.

Property Grids

This option enables you to set the **look and feel** in the **property grids**.



Property Grid

The **Show Property Grid Item Description** option shows the property grid item description. The default is selected.

The **Show Property Grid Toolbar** option shows the property grid toolbar. The default is selected.

The **Show Property Grid Inplace Buttons** option shows the property grid buttons for all items. The default is selected.

The **Default Property Grid Sort Order** option enables you to select the default property grid sort order for items. The options are Categorized, Alphabetical and No Sort and the default is Categorized.

The **Display Property Grid Boolean as** option enables you to select how boolean items are to be displayed. The options are Text and Check box and the default is Text.

The **Text for Boolean True** option enables you to enter the text for the boolean value True.

The **Text for Boolean False** option enables you to enter the text for the boolean value False.

The **Highlight Property Grid Changes** option enables you to highlight changed items in the property grid. The default is True.

The **Minimum displayed lines for Multiple-Line items** option sets the minimum display lines for multi-line inputs.

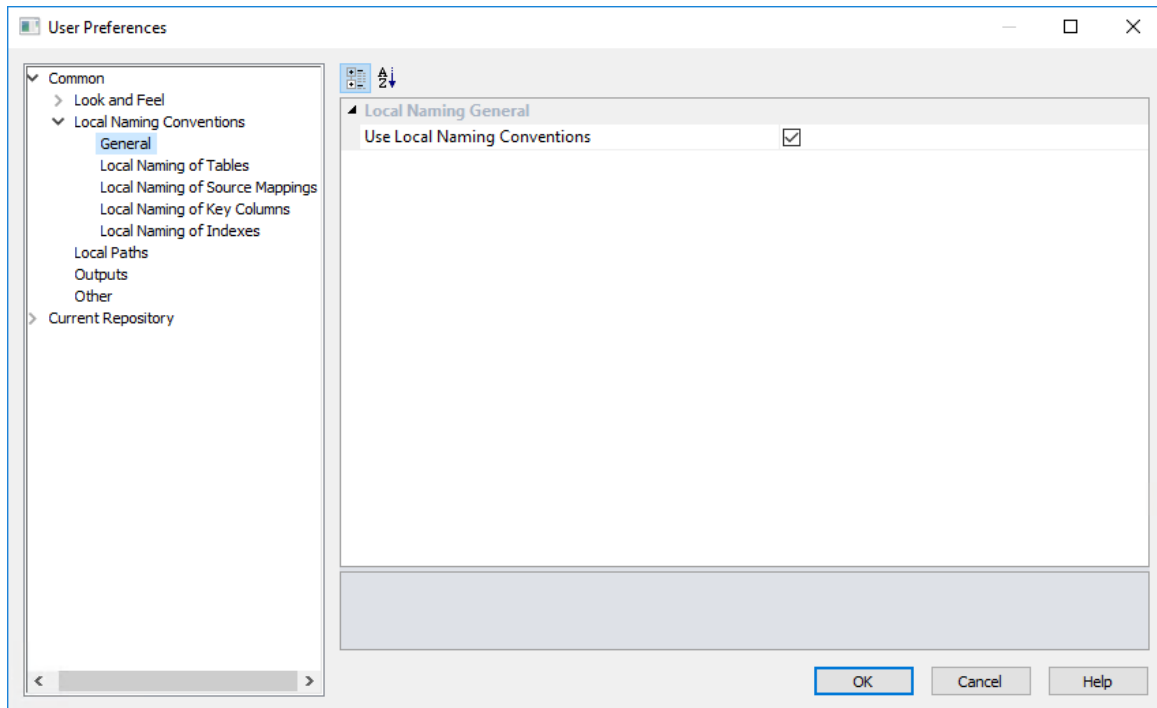
The **Maximum displayed lines for Multiple-Line items** option sets the maximum display lines for multi-line inputs.

LOCAL NAMING CONVENTIONS

The various options are described below.

General

This option enables you to set the **local naming conventions**.



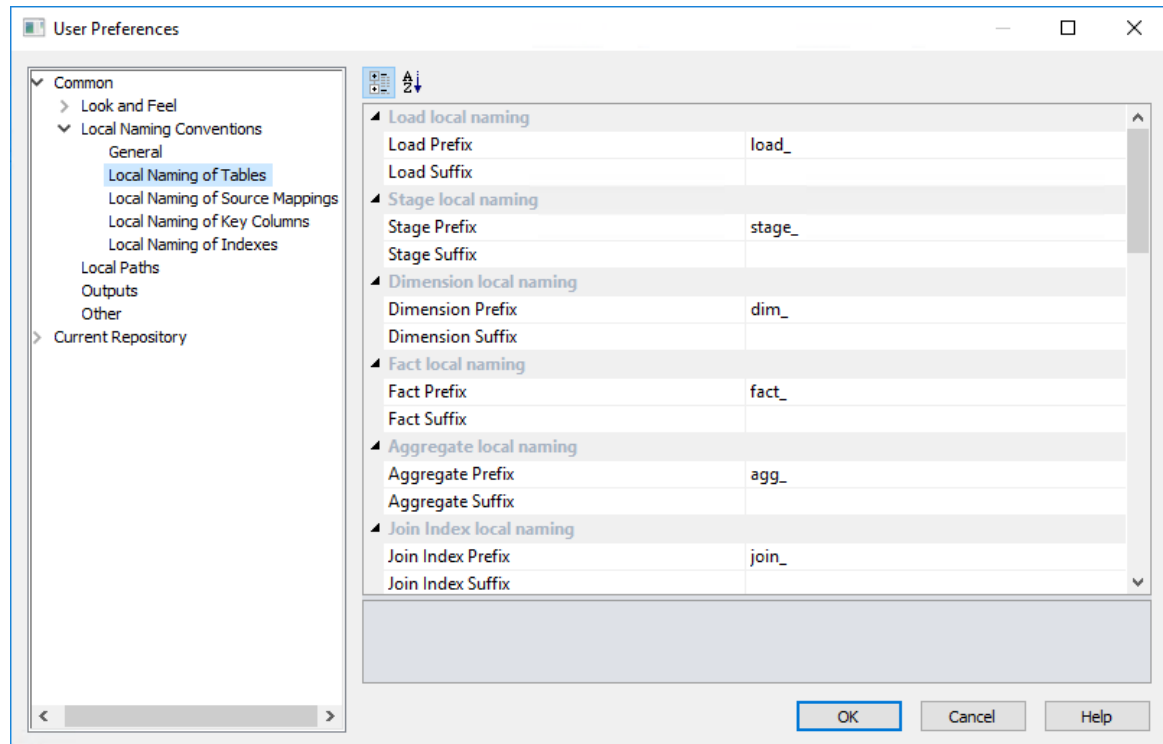
Local Naming General

Set this option if you want to **Use Local Naming Conventions**. If this option is set, the **Local Naming of Tables**, **Key Columns** and **Indexes** options is enabled in the object tree.

Note: If this option is set, it can overwrite short names and object prefixes.

Local Naming of Tables

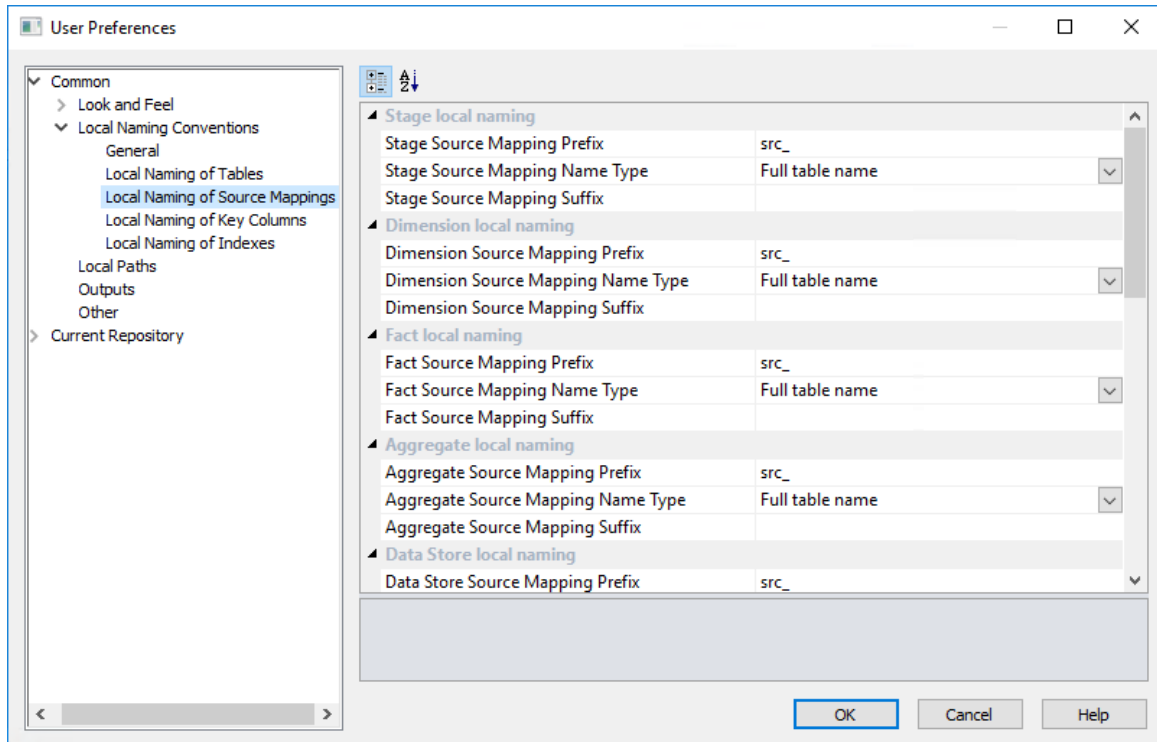
This option enables you to set the **Local Naming of Tables**.



Define the **prefix** and **suffix** used in the default naming convention for each table type.

Local Naming of Source Mappings

This option enables users to set the **Local Naming of Source Mappings** options.



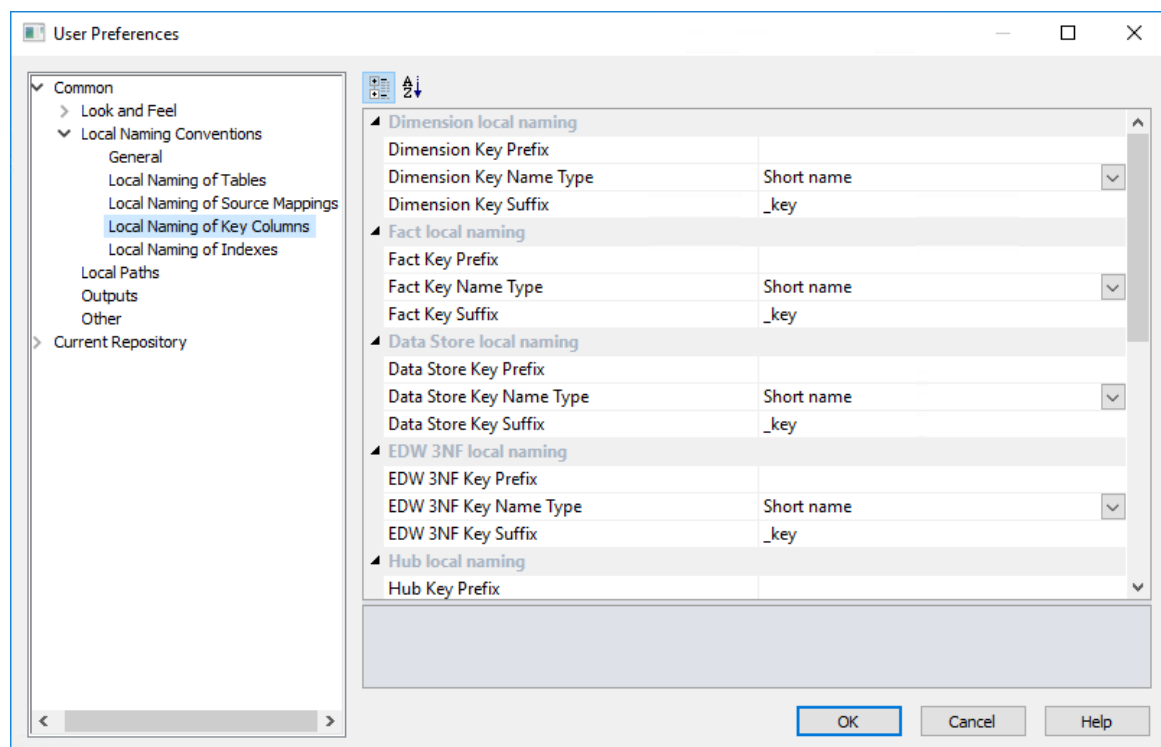
The **Source Mapping Prefix** option enables you to set the prefix used in the default source mapping naming convention.

The **Source Mapping Name Type** enables you to set the basis for the source mapping naming.

The **Source Mapping Suffix** enables you to set the suffix used in the default source mapping naming convention.

Local Naming of Key Columns

This option enables you to set the **Local Naming of Key Columns**.



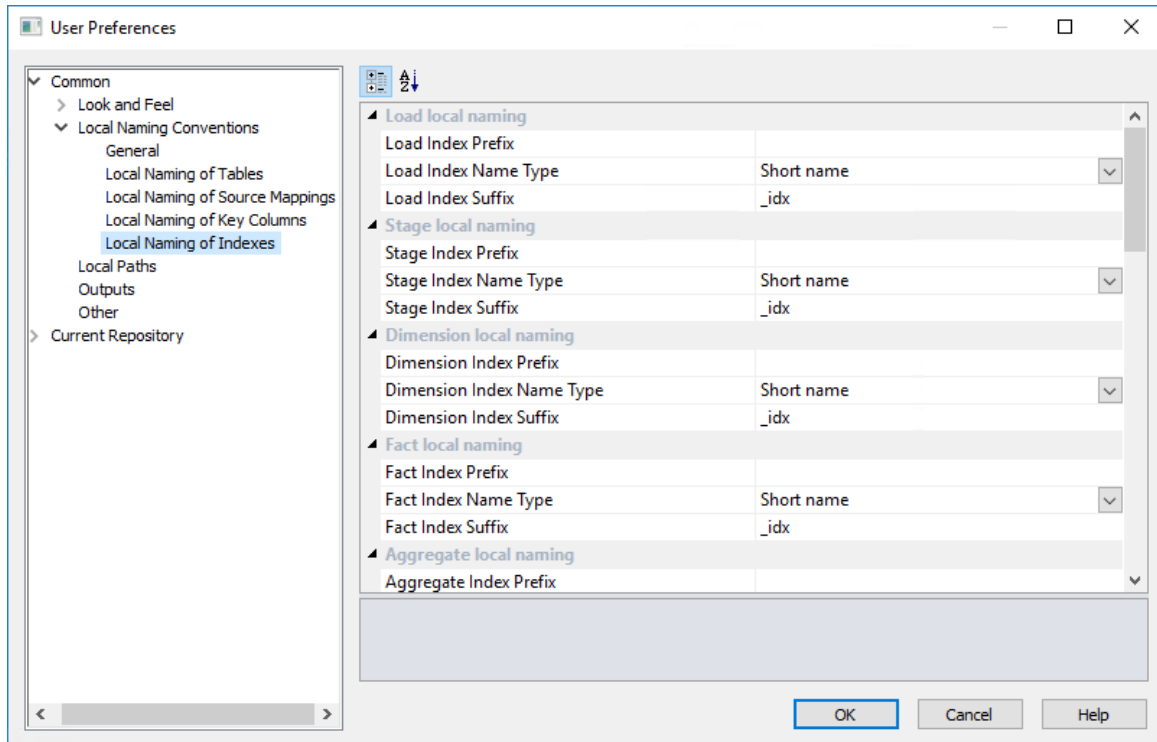
The **Key Prefix** option sets the prefix used in the default key naming convention.

The **Key Name Type** option sets the basis for the key naming.

The **Key Suffix** option sets the suffix used in the default naming convention.

Local Naming of Indexes

This option enables you to set the **Local Naming of Indexes**.



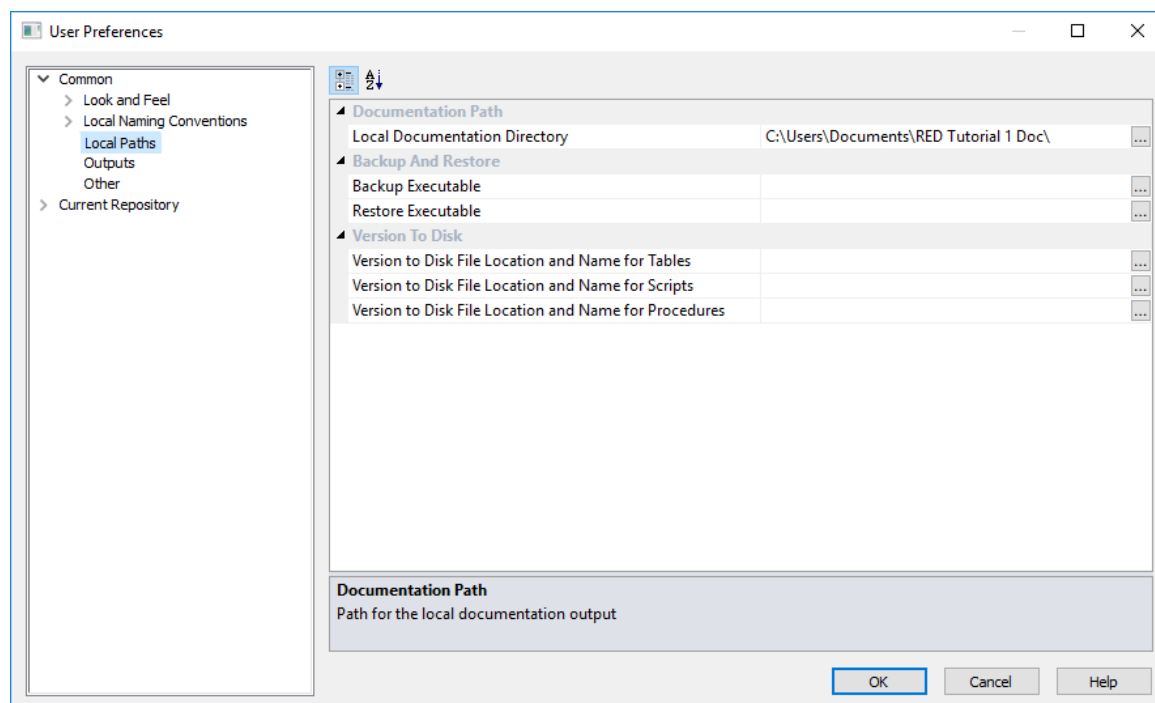
The **Index Prefix** option enables you to set the prefix used in the default index naming convention.

The **Index Name Type** enables you to set the basis for the index naming.

The **Index Suffix** enables you to set the suffix used in the default index naming convention.

LOCAL PATHS

This option enables you to set the **local paths** for documentation, backup and restore and for versioning to disk.



Documentation Path

Sets the **local documentation directory**.

Backup And Restore

The **Backup Executable** option sets the override for backup executable. By default, WhereScape RED tries to find the path of the backup executable. This is bcp.exe for SQL Server, exp.exe for Oracle and db2cmd.exe for IBM DB2. This edit box provides the ability to specify the exact location and name of the executable. This is useful when WhereScape RED cannot find the program or if there are multiple versions of the program on the PC.

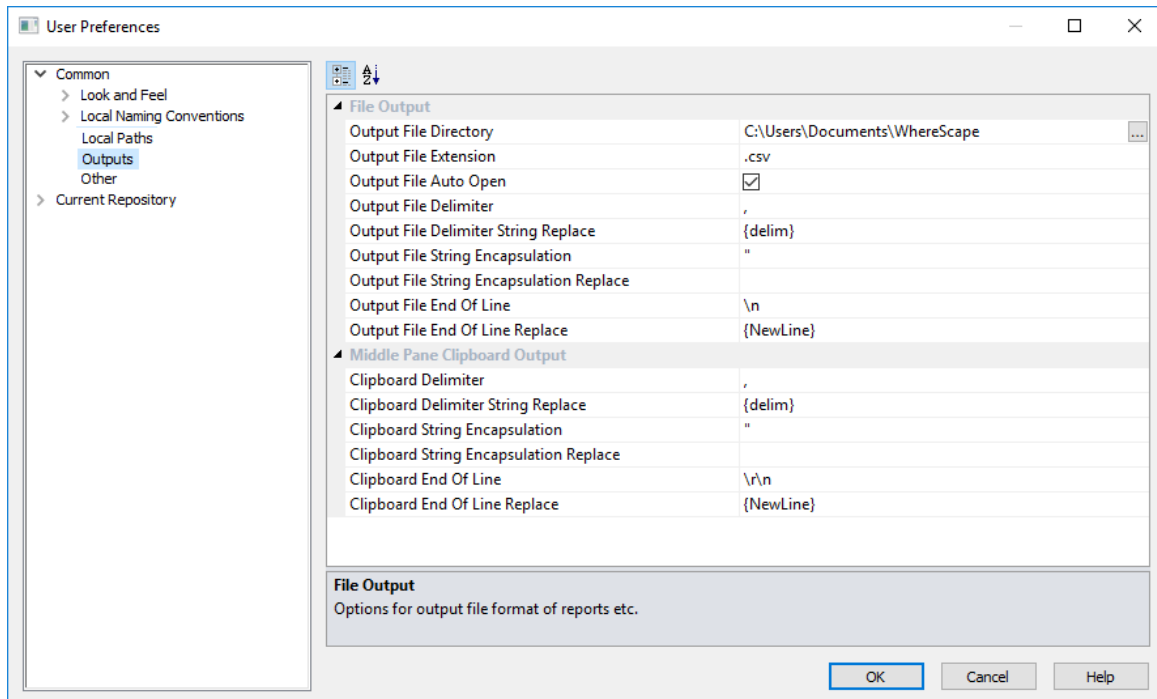
The **Restore Executable** option sets the override for restore executable. By default, WhereScape RED tries to find the path of the restore executable. This is bcp.exe for SQL Server, imp.exe for Oracle and db2cmd.exe for IBM DB2. This edit box provides the ability to specify the exact location and name of the executable. This is useful when WhereScape RED cannot find the program or if there are multiple versions of the program on the PC.

Version to Disk

Set the locations and names for **Versions to disk**. If any of the three version to disk paths are set, WhereScape RED will automatically create ascii files containing the applicable ddl or code each time an automated version occurs in the entered directory.

OUTPUTS

This option enables you to set the **Output** user preferences.



File Output

The **Output File Directory** option enables you to set the path for output files created from the middle pane.

The **Output File Extension** option enables you to set the file extension for output files created from the middle pane. This value determines the program that will auto open files.

The **Output File Auto Open** option enables you to automatically open the results in files created from the middle pane.

The **Output File Delimiter** option enables you to set the characters that separate each field within each record of output files created from the middle pane. Common values are , and |.

The **Output File Delimiter String Replace** option enables you to set the characters that will replace the delimiter character if it occurs inside a field.

The **Output File String Encapsulation** option enables you to set the characters that are used to enclose string values of files created from the middle pane. Common values are " and '.

The **Output File String Encapsulation Replace** option enables you to set the characters that will replace the encapsulation string if it occurs inside a field.

The **Output File End Of Line** option enables you to set the characters saved at the end of each record of files created from the middle pane. Common values are \n, \r and \t.

The **Output File End Of Line Replace** option enables you to set the characters that will replace the end of line string if it occurs inside a field.

Middle Pane Clipboard Output

The **Clipboard Delimiter** option enables you to set the characters that separate each field within each record of clipboard output created from the middle pane. Common values are , and |.

The **Clipboard Delimiter String Replace** option enables you to set the characters that will replace the delimiter character if it occurs inside a field.

The **Clipboard String Encapsulation** option enables you to set the characters that are used to enclose string values of clipboard output created from the middle pane. Common values are " and '.

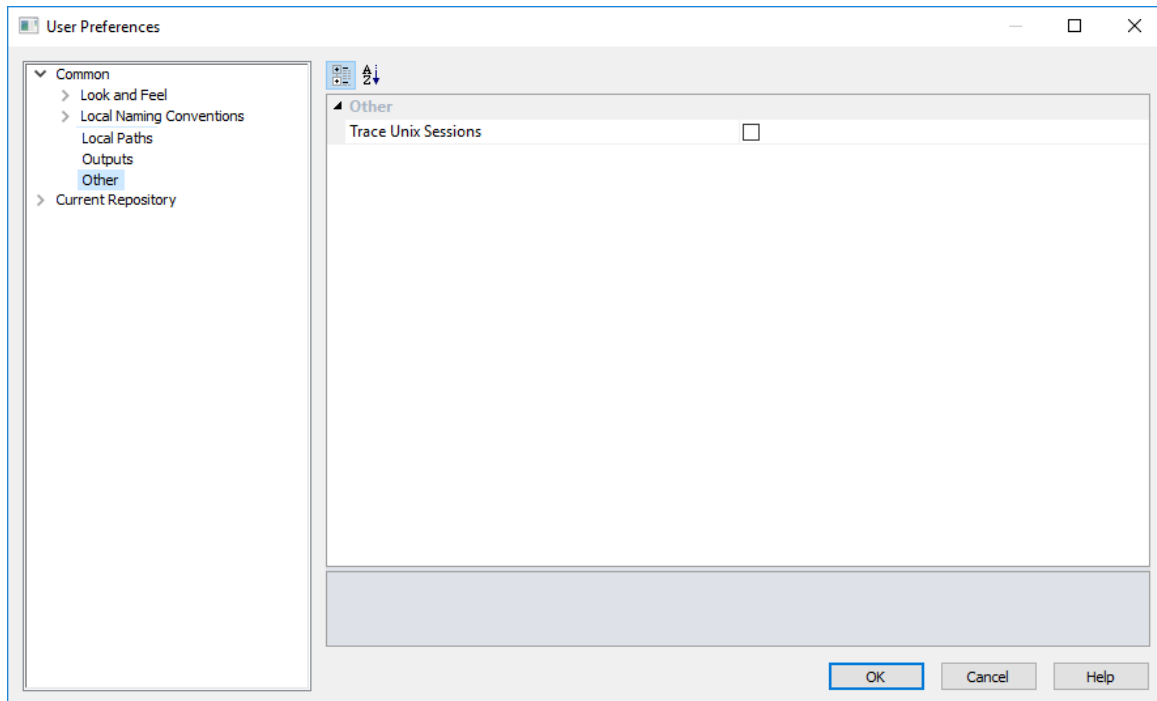
The **Clipboard String Encapsulation Replace** option enables you to set the characters that will replace the encapsulation string if it occurs inside a field.

The **Clipboard End of Line** option enables you to set the characters saved at the end of each record of clipboard output created from the middle pane. Common values are \n, \r and \t.

The **Clipboard End of Line Replace** option enables you to set the characters that will replace the end of line string if it occurs inside a field.

OTHER

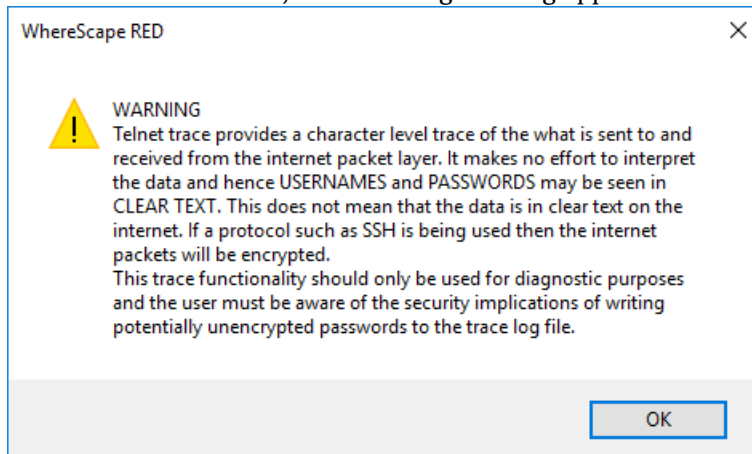
This option enables you to set the **Other** user preferences.



Other

The **Trace Unix Sessions** option traces all Unix activity undertaken by WhereScape RED until it is terminated. The file `WslMedTelnet.txt` is created in the program directory for WhereScape RED. This option is intended for debugging of specific Unix problems and the setting of this switch would normally be done at the request of WhereScape when attempting to solve a Telnet issue. This setting is only relevant for the PC on which the setting is made (e.g. it is not a global setting for the repository).

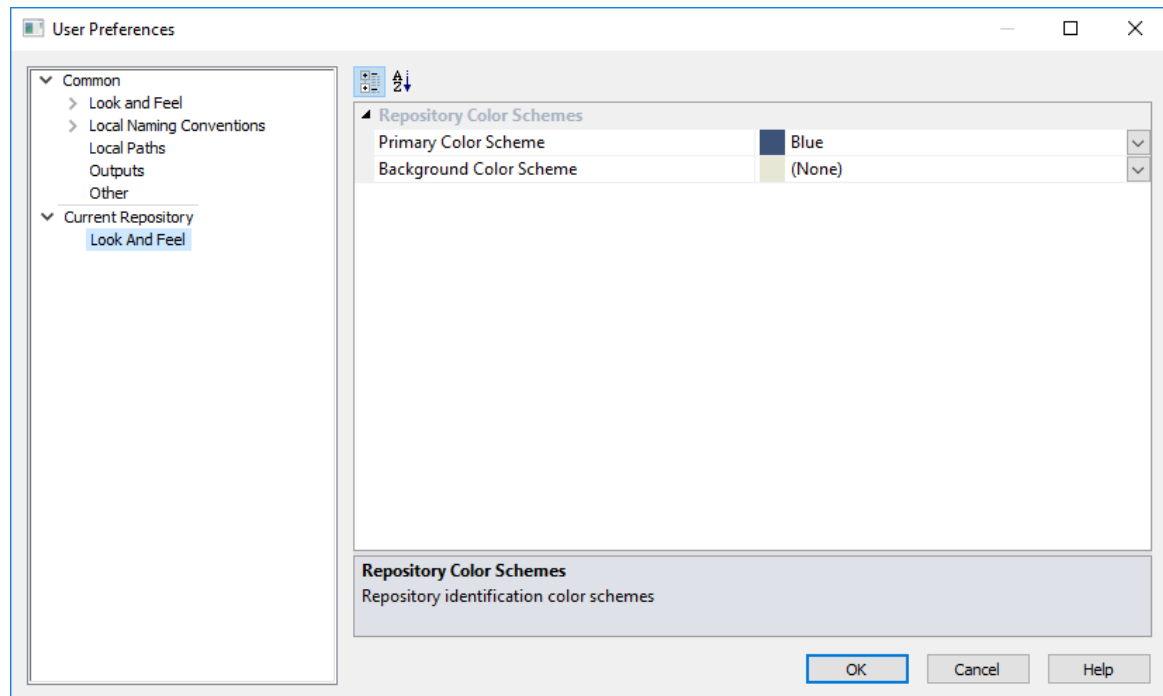
Note: When set to true, the following warning appears:



SETTINGS - CURRENT REPOSITORY

LOOK AND FEEL

This option enables you to set the **look and feel** for the **current repository**.

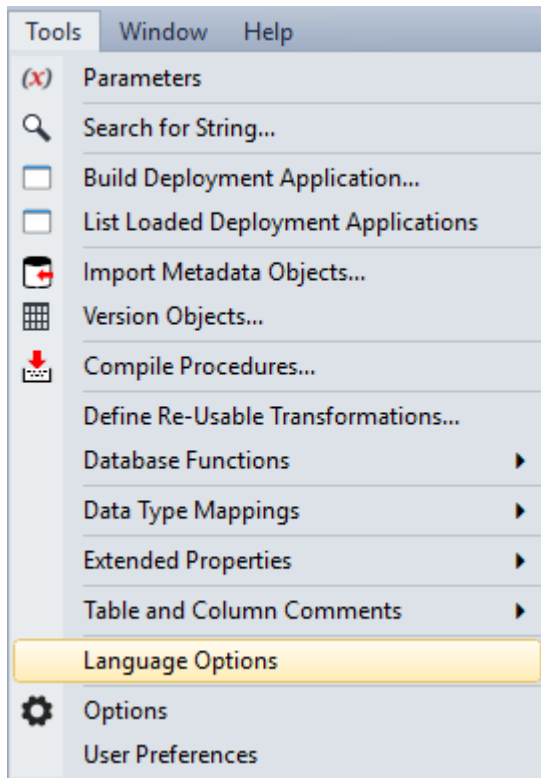


Repository Color Schemes

Set the primary and background **color schemes** for the current repository.

SETTINGS - LANGUAGE OPTIONS

Select **Language Options** from the **Tools** menu.



Languages can be defined via the **Tools>Language Options** menu. This option is only available for SQL Server databases and applies only to dimension, fact and OLAP objects. A blank entry means that no languages have been defined and thus no translations can be saved.

To add a language

Click the **Add Language** button; enter the new language ID and click **OK**.

To delete a language

First select the language from the drop-down list and then click the **Delete Language** button.

Note: All translations for the selected language will also be deleted.

Language Options

Language: French

Add Language

Delete Language

Language Description: This is the master Customer List for the Company.

Analysis Services Language: French (France)

OK Cancel

Language

The language Reference/ID.

Language Description

The language Description.

Analysis Services Language

Used to identify the language ID as used in Analysis Services.


CHAPTER 6

PARAMETERS

Parameters are a means of passing information between two or more procedures and between the RED environment and procedures. They can be edited within the RED environment by selecting the **Tools>Parameters** menu option. A list of parameters is displayed as per the example below:

Parameter listing		
Parameter	Value	Comments
(x) BUSINESS_DAYS_IN_MONTH	22	updated via daily_time_roll
(x) BUSINESS_DAYS_MTD	9	updated via daily_time_roll
(x) CURRENT_CAL_MONTH	201808	updated via daily_time_roll
(x) CURRENT_CAL_MONTH_NO	8	updated via daily_time_roll
(x) CURRENT_CAL_YEAR	2018	updated via daily_time_roll
(x) CURRENT_DATE	09-AUG-18	updated via daily_time_roll
(x) CURRENT_FIN_MONTH	201808	updated via daily_time_roll
(x) CURRENT_FIN_MONTH_NO	8	updated via daily_time_roll
(x) CURRENT_FIN_YEAR	2018	updated via daily_time_roll
(x) PREVIOUS_CAL_MONTH	201807	updated via daily_time_roll
(x) PREVIOUS_CAL_MONTH_NO	7	updated via daily_time_roll
(x) PREVIOUS_CAL_YEAR	2017	updated via daily_time_roll
(x) PREVIOUS_FIN_MONTH	201807	updated via daily_time_roll
(x) PREVIOUS_FIN_MONTH_NO	7	updated via daily_time_roll
(x) PREVIOUS_FIN_YEAR	2017	updated via daily_time_roll
(x) load_budget_0	18	Rowcount from native load at 08/08
(x) load_customer_0	6	Rowcount from native load at 18/10
(x) load_forecast_0	18	Rowcount from native load at 18/10
(x) load_order_header_0	9	Rowcount from native load at 18/10
(x) load_order_line_0	21	Rowcount from native load at 18/10

A parameter can be added, edited, copied or deleted by using the right-click menu in the Parameter column:



Typical parameter usage may be the global definition of how many days should be looked back for change data, a month or processing period etc.

Parameters can be used in **load tables** to place limits in a 'Where' clause, etc. See *Database Link Load - Source Mapping* (see "*Database Link Load - Source Screen*" on page 223) for more information.

They are also used by **stage table procedures** as variables. See *Generating the Staging Update Procedure* (on page 337) for more information.

Two procedures are provided to enables procedures to read and write parameters. These procedures are *WsParameterRead* (on page 921) and *WsParameterWrite* (on page 925). Using these procedures, a procedure can load and use the contents of a parameter, modify an existing parameter, or add a new parameter.

CHAPTER 7

CONNECTIONS

Connection objects serve several purposes in WhereScape RED:

- 1 They are used to browse potential source data in source systems and to acquire metadata. Potential source data includes database tables and flat files.

For database tables, WhereScape RED:

- Uses the **ODBC Source** set on each connection to browse the source system.
- Acquires the metadata for new **load tables** built from the source system using drag and drop.

For files, WhereScape RED:

- Connects directly to Windows, UNIX/Linux or Hadoop to analyze the source file for the new load table and acquire its metadata.
- Prompts for user input for any metadata not available in the source file.

NOTE1: ODBC connections must be either **User DSN** or **System DSN**. **File DSN** connections are **not** supported.

NOTE2: *Windows* and *UNIX* connections do not have an **ODBC Source** property. *UNIX* connections are use for UNIX and Linux systems.

- 2 Load tables with a connection of **Connection type ODBC** extract data from source systems using ODBC. The **ODBC Source** of the connection is the ODBC DSN used for the extract.

NOTE: If a **Teradata TPT compliant ODBC DSN** is defined in the **Connection Properties** menu, the **TPT DSN** is used for TPT ODBC Loads.

- 3 Each data warehouse metadata repository must have a Data Warehouse connection to use drag and drop to create new objects (other than load tables) in the data warehouse.

WhereScape RED:

- Uses the **ODBC Source** set on the Data Warehouse connection to browse the Data Warehouse database.
- Acquires the metadata for any **tables** built from existing data warehouse tables.

NOTE: This connection always has a **Connection type of Database**.

- 4 Cube objects require a connection to define the Analysis Services server used to create and load cubes. This is a connection with a **Connection type** *Microsoft Analysis Server 2005+*.
- 5 Export objects require a connection to define the target environment where exported data is written. This is a connection with a **Connection type** of *UNIX* or *Windows*.

IN THIS CHAPTER

Connection Types	147
Browsing a Connection	184
Changing a Connection's Properties.....	188
Reset Meta Database Connections.....	188
Connection Settings for BDA.....	189
Connection Extended Properties	195

CONNECTION TYPES

Connections can be set up via the following methods:

- **Connections to the Data Warehouse/Metadata Repository** (see "**Database - Data Warehouse/Metadata Repository**" on page 147)
- **Connections to Another Database** (see "**Database**" on page 155)
- **ODBC Based Connections** (see "**ODBC**" on page 159)
- **Connections to Windows** (see "**Windows**" on page 163)
- **Connections to UNIX/Linux**
- **Connections to Hadoop**
- **Connections to Microsoft Analysis Servers** (see "**Microsoft Analysis Server 2005+**" on page 178)

DATABASE - DATA WAREHOUSE/METADATA REPOSITORY

This section describes the connection to the Data Warehouse. It describes in greater detail the connection Properties as they apply to **Database** type connections and specifically to a Data Warehouse or Metadata repository connection.

The Data Warehouse connection is the connection that stores the metadata repository and it is the connection that is used in the drag and drop functionality to create the dimension, stage, fact and aggregate tables. This connection is also used to create cubes.

Connection types also impact the available load methods.



TIP: The Data Warehouse connection must exist, if you wish to use drag and drop to create dimensions, stage tables, fact tables, aggregates and cubes.

Data Warehouse connection example

- A User ID and Password must be specified or
- As below a **Teradata Wallet User ID** and **TD Wallet String**.

The screenshot shows the 'Connection DataWarehouse' dialog box with the following configuration details:

Property	Value
Connection Name	DataWarehouse
Connection Type	Database
Database Type	Teradata
ODBC Data Source Name (DSN)	RED_REPOSITORY
WhereScape RED Metadata Connection Indicator	<input checked="" type="checkbox"/>
Source System	
Database ID	TD_15_00
Database Link Name	
Provider Name	
Big Data Adapter Settings	
JDBC Connection String (JDBC URL)	jdbc:teradata://192.168.60.226/DATABASE=WSL_REPOSIT...
JDBC Driver Class Name	com.teradata.jdbc.TeraDriver
Omit Sqoop Driver Option	<input type="checkbox"/>
Sqoop Connection Manager Class	
Include Database/Schema Name in Sqoop Table Option	<input checked="" type="checkbox"/>
Include Sqoop Columns Option	<input checked="" type="checkbox"/>
Database Credentials	
Extract User ID	dssdemo
Extract User Password	*****
Administrator User ID	
Administrator User Password	
Teradata Wallet User ID	
Teradata Wallet String	
JDBC User ID	dssdemo
JDBC Password	*****
ODBC User Default	Extract User...

Buttons: OK, Cancel, Help

General

Connection Name

Name used to label the connection within WhereScape RED. Typically this is **Data Warehouse**

Connection Type

Indicates the connection source type or the connection method such as Database, ODBC, Windows, Unix. Here the connection type is **Database**.

Database Type

Type of database such as DB2, Greenplum, Hive, Netezza, Oracle, SQL Server, Teradata. Default is **(local)**.

ODBC Data Source Name (DSN)

ODBC Data Source Name (DSN) as defined in the Windows 32-bit **ODBC Data Source Administrator**.

Note: The ODBC Source Name defined in RED must be the same on all machines that use the corresponding connection.

WhereScape RED Metadata Connection Indicator

Distinguishes the special connection that identifies the WhereScape RED data warehouse/metadata repository. This option must be enabled for **Data Warehouse/Metadata Repository** type connections.

Note: There should only be one data warehouse/metadata connection in a WhereScape RED repository.

Source System

Database ID

Database Identifier (e.g. Oracle SID or TNS Name, Teradata TDPID) or Database Name (e.g. as in DB2 or SQL Server).

Database Link Name

Optional name of a Database Link that is used to access the database. Only required field for Database type connections where the database is not on the same server as the data warehouse. If the server is on the same database the link doesn't need to be defined and the field can be left blank.

For connections to databases located in different servers where a database link load is required:

- if the database link already exists to this source database then enter that link name in this field
- if the link doesn't exist then enter a link name

Big Data Adapter Settings

JDBC Connection String (JDBC URL)

Connection string used by the WhereScape Big Data Adapter to access this database. This is required for Apache Sqoop loads involving this connection.

The token `$OBJECT_DATABASE$` will be replaced by the name of the database containing the object (e.g. load table) being operated on.

Users loading into Teradata from a Hive or Hadoop connection using the Teradata connection manager for Sqoop, who need to load into more than one database, will need to add `DATABASE=$OBJECT_DATABASE$` into their JDBC URL on the DataWarehouse connection (e.g. `jdbc:teradata://192.168.60.226/DATABASE=$OBJECT_DATABASE$`).

BDA will replace `$OBJECT_DATABASE$` with the database containing the load table when loading into this connection, and with the source schema defined on the load table when loading from this connection.

JDBC Driver Class Name

JDBC driver class to be used by the WhereScape Big Data Adapter. This field must be set if the JDBC URL is set.

Select the appropriate JDBC Driver class name from the drop-down list. If this is left empty this will not be specified in generated commands.

Omit Sqoop Driver Option

If set, the `--driver` option to Sqoop will be omitted. This is required for certain connection types such as Oracle connections.

If you select the **Omit Sqoop Driver Option** check-box, the driver parameter will not be used in sqoop command line. This is a requirement for Oracle at the moment, as suggested by Sqoop documentation for 1.4.5.

Sqoop Connection Manager Class

Custom Sqoop connection manager class. Corresponds to the `--connection-manager` command line argument. Leave blank if this is not required.

Include Database/Schema Name in Sqoop Table Option

If set, the `--table` option to Sqoop will include the database/schema name of the destination table when performing Apache Sqoop loads into this connection. This is incompatible with some connection managers, for example the Cloudera Connector Powered by Teradata. If this is not set, users must ensure that the database/schema is otherwise communicated to Sqoop, for example by using the `$OBJECT_DATABASE$` token in the JDBC Connection String.

Include Sqoop Columns Option

If set, the `--columns` option to Sqoop will be included when performing Apache Sqoop loads into this connection. This is incompatible with some connection managers, for example the Cloudera Connector Powered by Teradata. If this is not set, users must ensure that the order of columns in the loads match the order in the metadata.

Database Credentials

Extract User ID

Database User that has access to SELECT from the source system tables to extract data.

Extract User Password

The password of the data warehouse user. For SQL Server, this field can be left blank if using a trusted login, or the server login password.

Administrator User ID

Left blank.

Administrator User Password

Left blank.

Teradata Wallet User ID

Database User ID that has access to SELECT from the source system tables to extract data.

Teradata Wallet String

The Teradata Wallet String is the string replacing the username and password for your connection. Teradata TD Wallet is a Teradata product part of the TTU (Tools and Utilities). Refer to Teradata documentation if you don't have a TD Wallet created already.

JDBC User ID

User ID to login as using JDBC via WhereScape Big Data Adapter (optional).

JDBC Password

Password to login with when using JDBC via WhereScapeBig Data Adapter (optional).

ODBC User Default

Select either Extract User ID or Teradata Wallet from the drop-down menu as the default log on method.

Other

Default Schema for Browsing

Optional comma-delimited list of schemas for the browser pane filter. Enter the schema(s) you want the connection to browse by default on the right browser pane.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection.

Note: The available options in this drop-down list is configured from Tools>Options>**Available Load Types** (see "**Settings - Available Load Types**" on page 120).

New Table Default Load Script Connection

The default Script Connection used for new Load Tables, created using this connection.

New Table Default Load Script Template

The default Template used for generating Scripts for new Load Tables created using this connection.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination. The SSIS Connection String is a required field for SSIS based loads.

For more details on how to create a SSIS Connection String and load data via an Integration Services Load package, see section **SSIS Loader** or **Loading Data from Flat Files using SSIS**.

Note: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited.

SSIS Use Column Names

Determines whether to use column names or column titles for SSIS loads. Teradata OLE DB driver by default returns titles. Leave this option disabled.

Data Type Mapping Set

Mapping Set to use when converting from a source database data type to a destination database data type.

Default Transform Function Set

Function Set that is selected by default in the Transformation dialogs.

When Connection is an OLAP Data Source

This section of fields is only relevant and will only be visible if the **Datawarehouse** field is enabled. These fields are required so that the data warehouse can be used as a source for the Analysis Services cubes.

MSAS Connection String

Connection string to be used by Microsoft Analysis Services (MSAS) to connect to the data warehouse.

Note: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited.

Connection Provider/Driver

Name of the Connection Provider/Driver to use to connect to the data warehouse database when it is used as the data source for OLAP cubes. Set to **TDOLEDB**.

Data Warehouse Server

Data Warehouse Server Name, which is used when the data warehouse is used as the data source for OLAP cubes. Set this to the Teradata TDPID.

Data Warehouse Database ID

Data Warehouse Database Identifier (e.g. Oracle SID or TNS Name, Teradata TDPID) or Database Name (e.g. as in DB2 or SQL Server), which is used when the data warehouse is used as the data source for OLAP cubes.

Target Connection Settings

Default Table Create DDL Template

The default Template used for generating DDL for new Tables created on targets of this connection.

Default View Create DDL Template

The default Template used for generating DDL for new Views created on targets of this connection.

Enable Automatic Creation of Indexes

This setting enables you to turn ON/OFF the automatic creation of metadata for indexes when creating new objects. Neither the RED metadata nor the physical index is created, if this setting is not enabled for objects.

This option is not selected by default. You can turn ON this option for databases that have indexes, and if you require RED to automatically create metadata indexes for objects.

Notes: This setting only applies when creating metadata indexes for new objects and does not affect existing metadata index records.

Turning this option OFF does not delete an object's existing metadata records or the physical index.

Default Update Script Connection

The default Script Connection for new Update Scripts in tables created using this connection.

This setting is only displayed in connections which have a target database type for which Update Script functionality is enabled, e.g. Custom database type.

Target Table Location [For target enabled licenses]

Add new Target Location

This option enables you to add new target database locations for objects in this connection.

- 1 Click the **Add** button to add the required target location for this connection.
- 2 Give the new target location a name and then enter the **Target Database** and/or **Temp Database**.
- 3 Default target database location(s) for **New Tables** can also be set from the **Tools>Options** menu – See **Settings - Storage - Target Location** (see "**Target Location**" on page 113).
- 4 For more details on setting specific target database locations on a table by table basis see **Storage** (on page 202).

Temp Databases specify the location for temporarily created tables used in Load and Export processes.



TIP: Once the connection has been set up, you can right-click the connection in the middle pane or double click the connection name from the left pane to view or edit the connection's Properties.

DATABASE

This section describes connections to another database source inside the same Teradata server, but not in the WhereScape RED meta repository.

Example

Sample database connection object properties screen:

The screenshot shows the 'Connection Tutorial' Properties dialog box. The 'General' section is expanded, showing the following fields:

Field	Value
Connection Name	Tutorial (OLTP)
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Tutorial_Teradata
WhereScape RED Metadata Connection Indicator	<input type="checkbox"/>

The 'Source System' section is also expanded, showing:

Field	Value
Database ID	TD_14_00
Database Link Name	

The 'Big Data Adapter Settings' section is expanded, showing:

Field	Value
JDBC Connection String (JDBC URL)	
JDBC Driver Class Name	
Omit Sqoop Driver Option	<input type="checkbox"/>
Sqoop Connection Manager Class	

The 'Database Credentials' section is expanded, showing:

Field	Value
Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	
Teradata Wallet User ID	tutorial
Teradata Wallet String	Your TD Wallet String
JDBC User ID	
JDBC Password	
ODBC User Default	Extract User

The 'Other' section is expanded, showing:

Field	Value
Default Schema for Browsing	tutorial
New Table Default Load Type	Database link load

At the bottom of the dialog, there is a 'Connection Name' field with the value 'tutorial' and a description: 'Name used to label the connection within WhereScape RED.' The dialog also has 'OK', 'Cancel', and 'Help' buttons.

The connection object Properties window has the following fields:

General

Connection Name

Name used to label the connection within WhereScape RED.

Connection Type

Indicates the connection source type or the connection method such as Database, ODBC, Windows, Unix. Set to **Database**.

Database Type

Type of database such as DB2, Greenplum, Hive, Netezza, Oracle, SQL Server, Teradata. Default is **(local)**.

ODBC Data Source Name (DSN)

ODBC Data Source Name (DSN) as defined in the Windows 32-bit **ODBC Data Source Administrator**.

Note: The ODBC Source Name defined in RED must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Only required for the Data Warehouse/metadata repository connection. Leave this check box unselected.

Source System

Database ID

Database Identifier (Teradata TDPID).

Database Link Name

This field is always blank for Teradata.

Database Credentials

Extract User ID

Database User that has access to SELECT from the source system tables to extract data.

Extract User Password

The password of the data warehouse user.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

Teradata Wallet User ID and Teradata Wallet String

Enter the relevant credentials when using the Teradata Wallet log on method instead of Extract User ID and Password for connecting to another database.

ODBC User Default

Select either Extract User ID or Teradata Wallet from the drop-down menu as the default log on method.

Other

Default Schema for Browsing

Optional comma-delimited list of schema(s) for the browser pane filter. Enter the schema(s) you want the connection to browse by default on the right browser pane.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection.

Note: The available options in this drop-down list is configured from Tools>Options>**Available Load Types** (see "**Settings - Available Load Types**" on page 120).

New Table Default Load Script Connection

The default Script Connection used for new Load Tables, created using this connection.

New Table Default Load Script Template

The default Template used for generating Scripts for new Load Tables created using this connection.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination. The SSIS Connection String is a required field for SSIS based loads.

For more details on how to create a SSIS Connection String and load data via an Integration Services Load package, see section **SSIS Loader**.

Note: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited.

Data Type Mapping Set

Mapping Set to use when converting from a source database data type to a destination database data type. Setting this field to **(Default)** makes RED automatically select the relevant mapping set, otherwise you can choose one of the standard mapping sets from the drop-down list or create a new set.

Target Connection Settings

Default Table Create DDL Template

The default Template used for generating DDL for new Tables created on targets of this connection.

Default View Create DDL Template

The default Template used for generating DDL for new Views created on targets of this connection.

Enable Automatic Creation of Indexes

This setting enables you to turn ON/OFF the automatic creation of metadata for indexes when creating new objects. Neither the RED metadata nor the physical index is created, if this setting is not enabled for objects.

This option is not selected by default. You can turn ON this option for databases that have indexes, and if you require RED to automatically create metadata indexes for objects.

Notes: This setting only applies when creating metadata indexes for new objects and does not affect existing metadata index records.

Turning this option OFF does not delete an object's existing metadata records or the physical index.

Default Update Script Connection

The default Script Connection for new Update Scripts in tables created using this connection.

This setting is only displayed in connections which have a target database type for which Update Script functionality is enabled, e.g. Custom database type.

Target Table Location [For target enabled licenses]**Add new Target Location**

This option enables you to add new target database locations for objects in this connection.

- 1 Click the **Add** button to add the required target location for this connection.
- 2 Give the new target location a name and then enter the **Target Database** and/or **Temp Database**.
- 3 Default target database location(s) for **New Tables** can also be set from the **Tools>Options** menu – See **Settings - Storage - Target Location** (see "**Target Location**" on page 113).
For more details on setting specific target database locations on a table by table basis see **Storage** (on page 202).

Notes:

- The database and schema names for **Custom** database connections are used as follows:

<database> . <schema> . objectname

Leave database name blank if not required. Leave schema name blank to use the default schema.

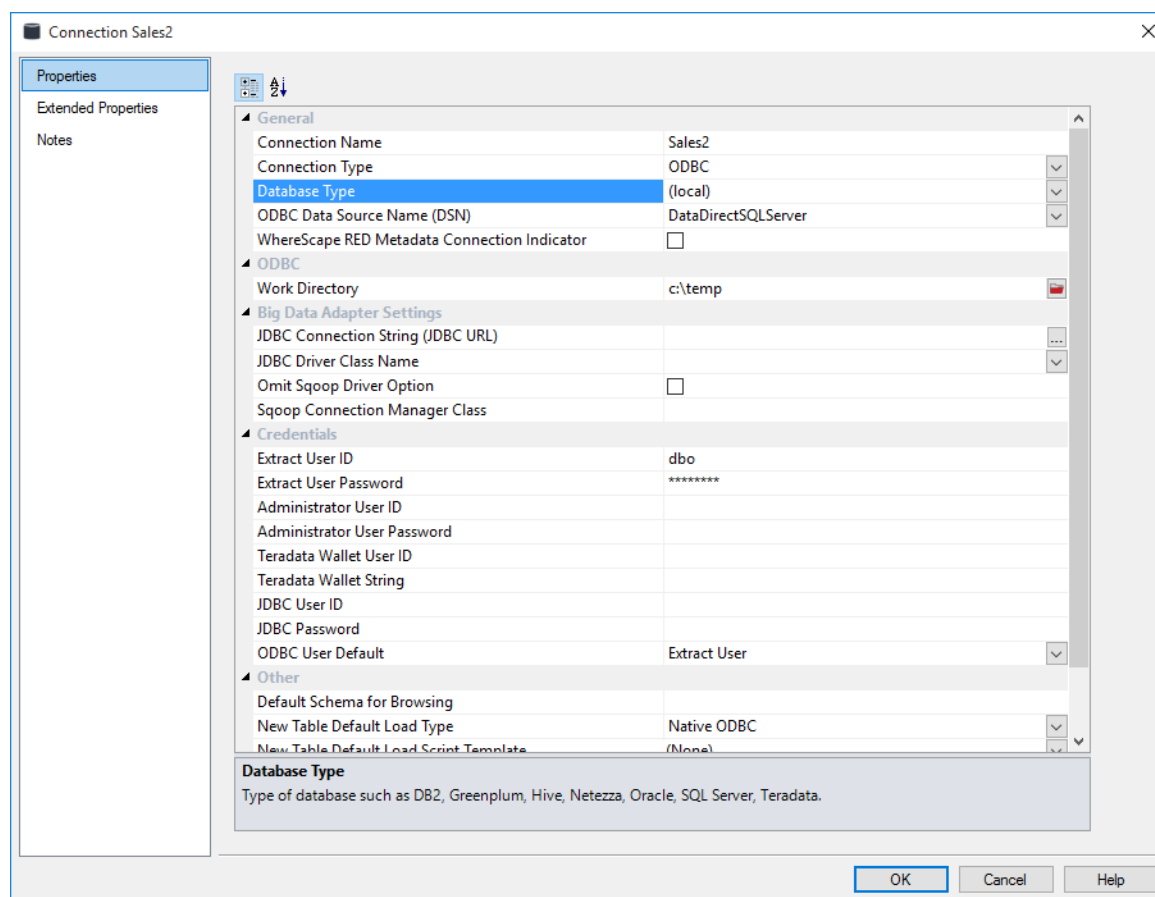
- Once the connection has been set up, you can right-click the connection in the middle pane or double click on the connection name from the left pane to view or edit the connection's Properties.

ODBC

This connection is via an ODBC link. All data movement is performed using the ODBC connection.

Native ODBC Load Example

- A User ID and Password must be specified or
- As below a **Teradata Wallet User ID** and **TD Wallet String**.



General

Connection Name

Name used to label the connection within WhereScape RED.

Connection Type

Indicates the connection source type or the connection method. Set to **ODBC**.

Database Type

Type of database such as DB2, Greenplum, Hive, Netezza, Oracle, Sql Server, Teradata, etc.

ODBC Data Source Name (DSN)

ODBC Data Source Name (DSN) as defined in the Windows 32-bit **ODBC Data Source Administrator**.

Note: The ODBC Source Name defined in RED must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Only required for the Data Warehouse/metadata repository connection. Leave this check box unselected.

Note: There must only be one metadata connection in a WhereScape RED metadata repository.

ODBC

Work Directory

Windows directory used by WhereScape RED to create temporary files for minimal logged extracts. The directory must exist and allow write access. There must be a different work directory for each WhereScape RED Scheduler running on the same machine to avoid file conflicts. Typically C:\Temp or a sub-directory of C:\Temp is used. See *Native ODBC Based Load* (on page 225).

Big Data Adapter Settings

JDBC Connection String (JDBC URL)

Connection string used by the WhereScape Big Data Adapter to access this database.

JDBC Driver Class Name

JDBC driver class to be used by the WhereScape Big Data Adapter. This field must be set if the JDBC URL is set.

Select the appropriate JDBC Driver class name from the drop-down list. If this is left empty this will not be specified in generated commands.

Omit Sqoop Driver Option

If set, the --driver option to Sqoop will be omitted. This is required for certain connection types such as Oracle connections.

If you select the **Omit Sqoop Driver Option** check-box, the driver parameter will not be used in sqoop command line. This is a requirement for Oracle at the moment, as suggested by Sqoop documentation for 1.4.5.

Sqoop Connection Manager Class

Custom Sqoop connection manager class. Corresponds to the --connection-manager command line argument. Leave blank if this is not required.

Database Credentials

Extract User ID

Database User that has access to SELECT from the source system tables to extract data.

Extract Password

Database Password to use with the Extract User ID to login to the source system to extract data.

Administrator User ID

Leave blank.

Administrator Password

Leave blank.

Teradata Wallet User ID and Teradata Wallet String

Enter the relevant Teradata Wallet credentials instead of the Extract user and password for the ODBC connection.

JDBC User ID

User ID to login as using JDBC via WhereScape Big Data Adapter (optional).

JDBC Password

Password to login with when using JDBC via WhereScape Big Data Adapter (optional).

ODBC User Default

Select either Extract User ID or Teradata Wallet from the drop-down menu as the default log on method.

Other

Default Schema for Browsing

Optional comma-delimited list of schema(s) for the browser pane filter. Enter the schema(s) you want the connection to browse by default on the right browser pane.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection. Set to the desired type of **ODBC load**—Native, ODBC, TPT, TPT Script load, Integration Services load or Externally Loaded.

Note: The available options in this drop-down list is configured from Tools>Options>**Available Load Types** (see "**Settings - Available Load Types**" on page 120).

New Table Default Load Script Connection

The default Script Connection used for new Load Tables, created using this connection.

New Table Default Load Script Template

The default Template used for generating Scripts for new Load Tables created using this connection.

SSIS Connection String (OLEDB)

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination. The SSIS Connection String is a required field for SSIS based loads.

For more details on how to create a SSIS Connection String and load data via an Integration Services Load package, see section **SSIS Loader**.

Note: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited.

TPT ODBC Data Source Name (DSN)

The Teradata TPT compliant ODBC Data Source Name (DSN) that is to be used in the TPT Read Operator for TPT ODBC Loads. If not populated, this defaults to the DSN defined above for this Connection.

Data Type Mapping Set

Mapping Set to use when converting from a source database data type to a destination database data type. Setting this field to **(Default)** will cause RED to automatically select the relevant mapping set.

Target Connection Settings

Default Table Create DDL Template

The default Template used for generating DDL for new Tables created on targets of this connection.

Default View Create DDL Template

The default Template used for generating DDL for new Views created on targets of this connection.

Enable Automatic Creation of Indexes

This setting enables you to turn ON/OFF the automatic creation of metadata for indexes when creating new objects. Neither the RED metadata nor the physical index is created, if this setting is not enabled for objects.

This option is selected by default. You can turn OFF this option for databases that do not have indexes, and if you do not require RED to automatically create metadata indexes for objects.

Notes: This setting only applies when creating metadata indexes for new objects and does not affect existing metadata index records.

Turning this option OFF does not delete an object's existing metadata records or the physical index.

Default Update Script Connection

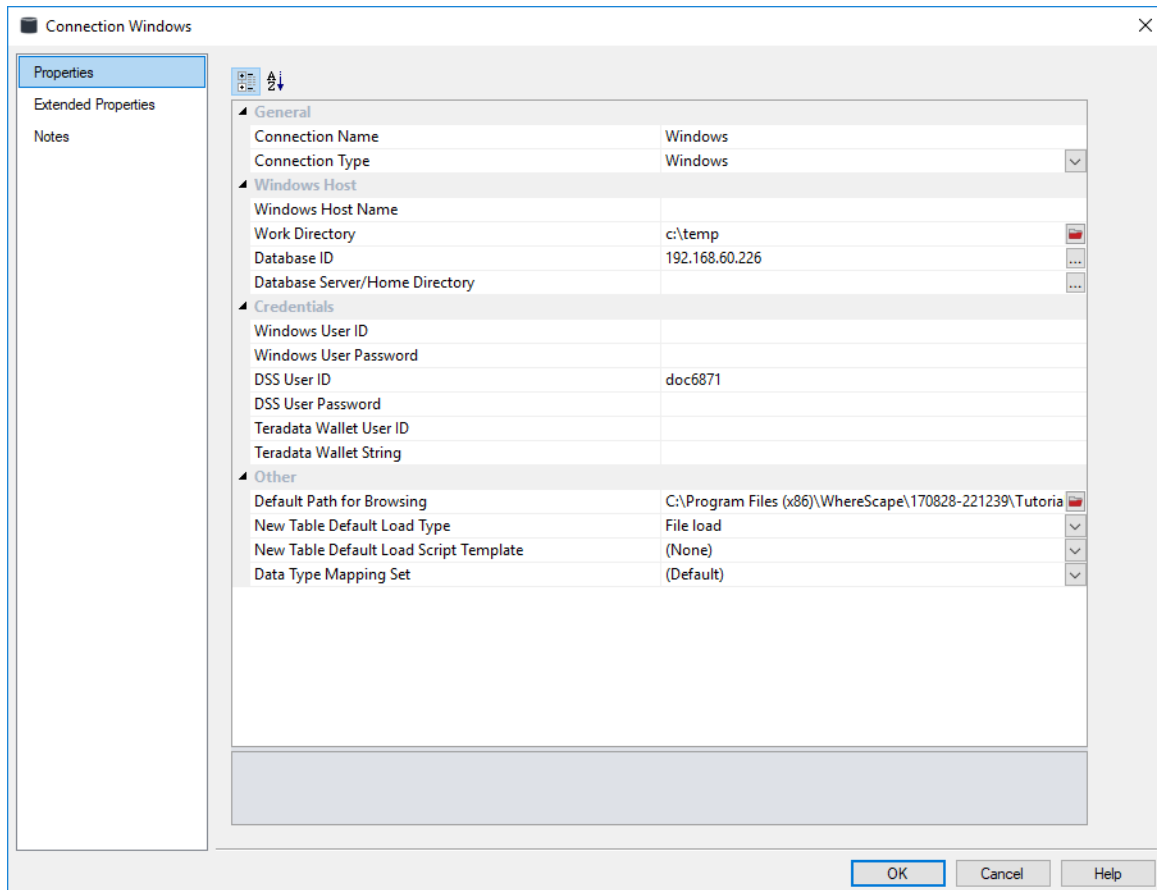
The default Script Connection for new Update Scripts in tables created using this connection.

This setting is only displayed in connections which have a target database type for which Update Script functionality is enabled, e.g. Custom database type.

Once the connection has been set up, you can right-click on the connection in the middle pane or double click on the connection name from the left pane to view or edit the connection's Properties.

WINDOWS

This connection is back to the PC that you are working on, or to a host Windows PC.



General

Connection Name

Name used to label the connection within WhereScape RED.

Connection Type

Indicates the connection source type or the connection method, such as Database, ODBC, Windows, Unix. Set to **Windows**.

Windows Host

Windows Host Name

IP address or host name that identifies the Windows machine. Leave blank to connect to the local machine.

Work Directory

Windows directory used by WhereScape RED to create temporary files for minimal logged extracts. The directory must exist and allow write access. There must be a different work directory for each WhereScape RED Scheduler running on the same machine to avoid file conflicts. Typically C:\Temp or a sub-directory of C:\Temp is used.

Database ID

Database Identifier (Teradata TDPID).

Database Server/Home Directory

Optional to specify the Database Home Directory if it is different from the standard home directory.

Credentials

Windows User ID and Password

Leave this blank if you are connecting to your own PC. Enter details if you are connecting remotely to another Windows system.

Dss User ID and Password/ Teradata Wallet User ID and Teradata Wallet String

Enter the relevant details for connecting to the **Data Warehouse**. Enter either DSS User ID and Password or Teradata Wallet credentials depending on the log on method selected.

OTHER

Default Path for Browsing

Optional default Path for browser pane filter. When a path has been selected in this field, it becomes the initial point for browsing and it is also expanded on open in the right hand browser pane.

New Table Default Load Type

The default Load type for new tables created using this connection.

Notes:

The available options in this drop-down list is configured from Tools>Options>*Available Load Types* (see "*Settings - Available Load Types*" on page 120).

The following load types are **not** enabled by default for newly created Teradata repositories:

- Integration Services (SSIS) load in Database, ODBC and Windows connections.
 - Native ODBC load in an ODBC connection.
 - XML file load in a Windows connection.
-

New Table Default Load Script Template

The default Template used for generating Scripts for new Load Tables created using this connection.

Data Type Mapping Set

Mapping Set to use when converting from a source database data type to a destination database data type. Setting this field to **(Default)** makes RED automatically select the relevant mapping set, otherwise you can choose one of the standard mapping sets from the drop-down list or create a new set.

Once the connection has been set up, you can right-click the connection in the middle pane and view the Properties for that connection.


To test the connection

- Select **Browse > Source Tables** from the top menu.
- In the right pane you should be able to drill down to the area required.

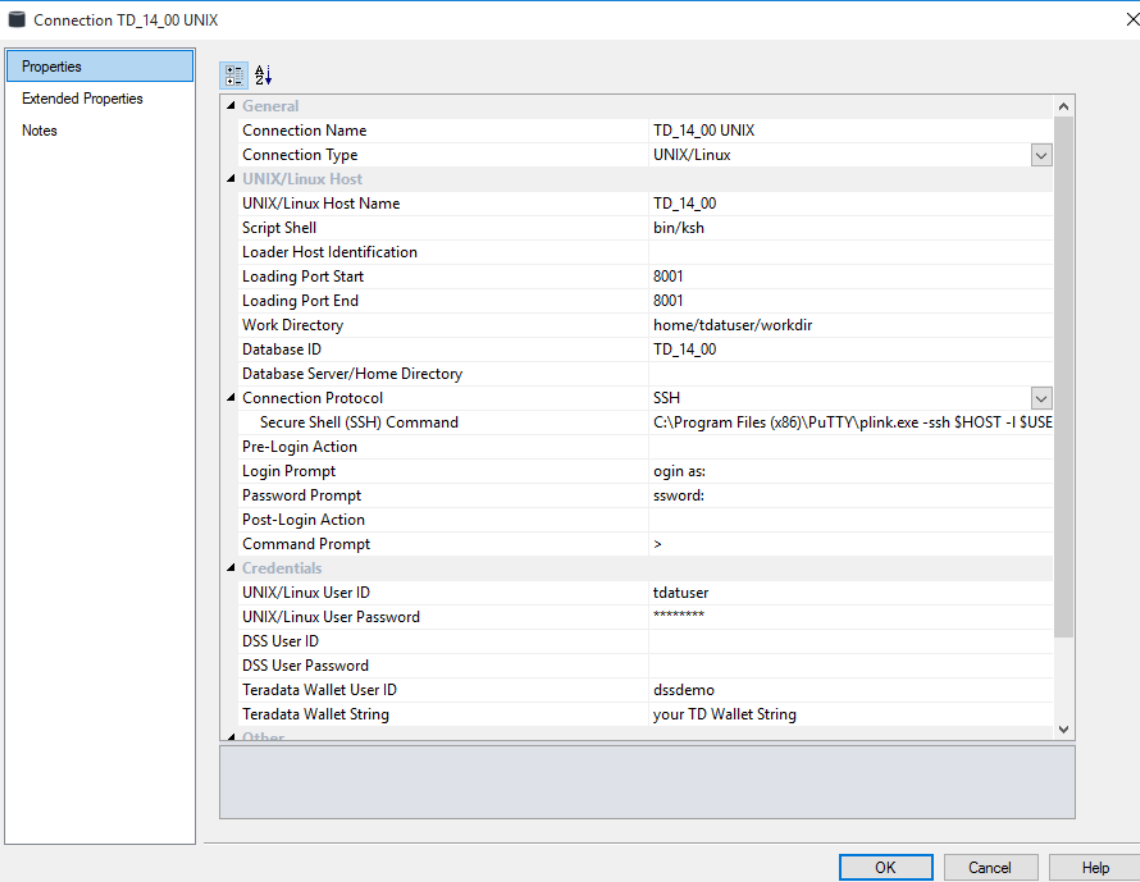
UNIX

This section describes the connection properties as they apply to UNIX connections. From a UNIX connection you can only process flat file loads.

Note: If the **UNIX/Linux** connection returns a blank screen or an error message in the **Results** pane after the connection is browsed, take necessary action through the **Server (SSH)** tab next to the main Builder and Scheduler tabs.

. This tab is displayed after browsing the UNIX connection.

Sample UNIX connection screen:



The screenshot shows a dialog box titled "Connection TD_14_00 UNIX" with a "Properties" tab selected. The dialog is divided into several sections:

- General:**
 - Connection Name: TD_14_00 UNIX
 - Connection Type: UNIX/Linux
- UNIX/Linux Host:**
 - UNIX/Linux Host Name: TD_14_00
 - Script Shell: bin/ksh
 - Loader Host Identification:
 - Loading Port Start: 8001
 - Loading Port End: 8001
 - Work Directory: home/tdatuser/workdir
 - Database ID: TD_14_00
 - Database Server/Home Directory:
- Connection Protocol:**
 - Secure Shell (SSH) Command: C:\Program Files (x86)\PuTTY\plink.exe -ssh \$HOST -l \$USER
 - Pre-Login Action:
 - Login Prompt: ogin as:
 - Password Prompt: ssword:
 - Post-Login Action:
 - Command Prompt: >
- Credentials:**
 - UNIX/Linux User ID: tdatuser
 - UNIX/Linux User Password: *****
 - DSS User ID:
 - DSS User Password:
 - Teradata Wallet User ID: dssdemo
 - Teradata Wallet String: your TD Wallet String
- Other:** (Empty section)

Buttons at the bottom: OK, Cancel, Help.

General

Connection Name

Name used to label the connection within WhereScape RED.

Connection Type

Indicates the connection source type or the connection method such as Database, ODBC, Windows, Unix. Set to **UNIX**.

UNIX/Linux Host

UNIX/Linux Host Name

IP address or host name that identifies the UNIX machine.

Script Shell

Path to the POSIX-compliant UNIX/Linux shell to use for generated scripts. For **UNIX** hosts, set to **/bin/ksh**. For **Linux** hosts set to **/bin/sh**.

If this field is left blank, a default will be chosen based on the name of the connection and the type of database used for the WhereScape RED metadata repository.

Loader Host Identification

IP Address or host name(s) that identifies the Loader/ Multiple hosts can be entered with using a comma (,) to delimit.

Work Directory

Windows directory used by WhereScape RED to create temporary files for minimal logged extracts. The directory must exist and allow write access. There must be a different work directory for each WhereScape RED Scheduler running on the same machine to avoid file conflicts. Typically C:\Temp or a sub-directory of C:\Temp is used.

Database ID

Source Database Identifier (Teradata TDPID).

Database Server/Home Directory

Optional to specify the Database Home Directory if it is different from the standard home directory.

Connection Protocol

Telnet or Secure Shell (SSH) protocol to use to connect to the UNIX/Linux machine. For SSH, the 'Secure Shell (SSH) Command' property is enabled to specify how to connect.

Secure Shell (SSH) Command

Command to execute to connect to a UNIX/Linux machine using the Secure Shell (SSH) protocol such as **C:\Program Files(x86)\PuTTY\plink.exe -ssh \$HOST\$ -l \$USER\$ -pw \$PASSWORD\$**

Pre-Login Action, Login Prompt, Password Prompt, Post-Login Action, and Command Prompt.

These fields are only used to create a Telnet connection to the host machine. WhereScape RED uses the Telnet connection in the drag and drop functionality. **It is not used in the actual production running of the Data Warehouse**, and is only necessary if you wish to use the drag and drop functionality.

Pre-Login Action

Response or command to send BEFORE logging in to the UNIX/Linux machine. Typically this is NOT necessary but it can be used to indicate that the UNIX/Linux Login Prompt is preceded by a line-feed (\n). However it is preferable that the UNIX/Linux login displays the Login Prompt without anything preceding it. [Optional]

Login Prompt

The UNIX login prompt, or the tail end of the login prompt, e.g, **ogin as:**

Password Prompt

The UNIX password prompt, or the tail end of the password prompt, e.g, **ssword:**

Post-Login Action

Not often used but may be necessary to respond to a login question. It is preferable that the UNIX login goes straight to the command prompt.

Command Prompt

Enter the UNIX/Linux command prompt, or the tail end of that prompt, typically **\$**

Note: In order to ascertain some of the above fields you have to log in to the UNIX system.

Credentials**UNIX/Linux User ID**

User Account to login to the UNIX/Linux Host.

UNIX/Linux User Password

Password to login to the UNIX/Linux Host.

DSS User ID

Database User to connect to the WhereScape RED metadata repository.

DSS User Password

Database Password to connect to the WhereScape RED metadata repository.

Teradata Wallet User ID and Teradata Wallet String

Enter the relevant Teradata Wallet credentials instead of the DSS User and Password for the Unix connection if using the Teradata log on method.

Other

Default Path for Browsing

Optional default Path for browser pane filter. When a path has been selected in this field, it becomes the initial point for browsing and it is also expanded on open in the right hand browser pane.

New Table Default Load Type

The default Load type for new tables created using this connection. Select from the **File Load**, **Script based load** or **Externally loaded** options.

Note: The available options in this drop-down list is configured from Tools>Options>**Available Load Types** (see "**Settings - Available Load Types**" on page 120).

New Table Default Load Script Template

The default Template used for generating Scripts for new Load Tables created using this connection.

Data Type Mapping Set

XML files have been created to store mappings from one set of data types to another. Setting this field to **(Default)** makes RED automatically select the relevant mapping set, otherwise you can choose one of the standard mapping sets from the drop-down list or create a new one.

To validate the fields

- Right-click on the connection name
- Select **Telnet window**

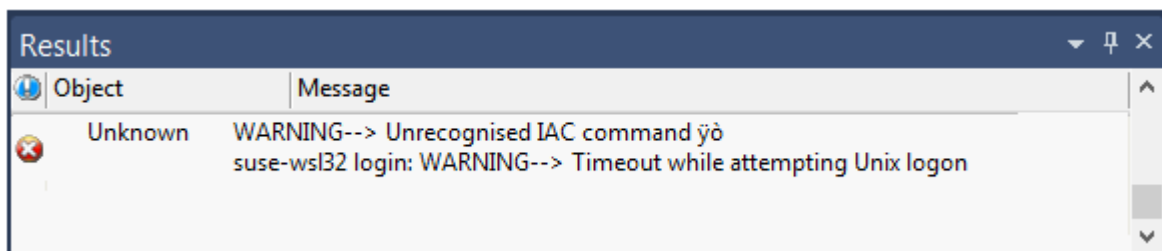
This will provide a telnet window that can be used to log on to the UNIX server.

To test the drag and drop functionality

- From the top menu select **Browse>Source Tables**
- Drill down to the area required
- Select an object from left pane and drag it to the middle pane

Connection Failures

In the event that a telnet connection cannot be established to the UNIX host the following result box will normally appear after approximately 30 seconds.

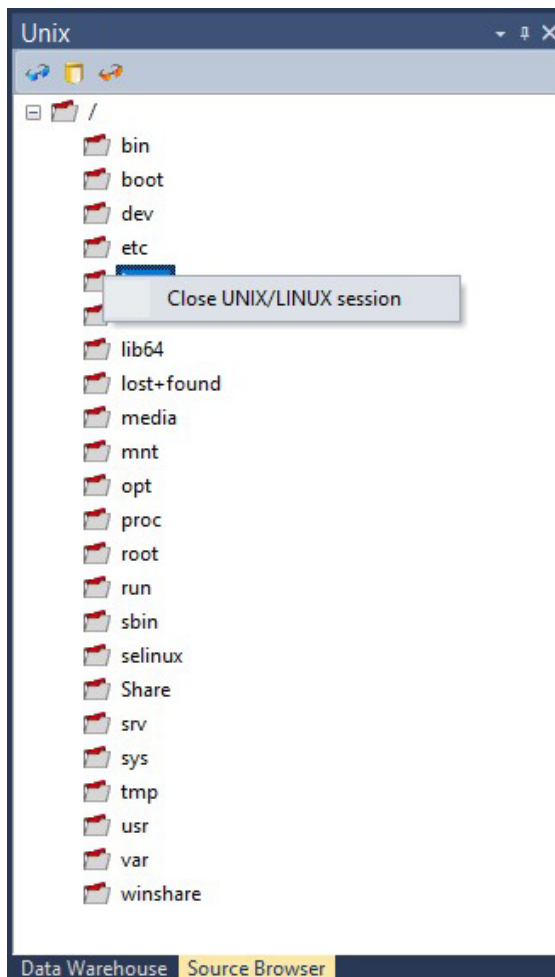


Attempt the connection again, and using the **Window** menu option select the **Telnet** window. This displays the login session, and provides an insight as to why the connection is not being completed.

If the situation cannot be resolved a telnet trace can be acquired. Select the **Tools>Options** menu and click the **Trace all Unix sessions** check box. Try to do the connection or browse again. A log file called WslMedTelnet.log will be created in the WhereScape program directory. Edit the log file and ensure there are no passwords visible and then contact WhereScape support, using the WhereScape forum at <http://www.wherescape.com>.

Closing the Connection

To close the collection, right-click in the browser pane and select **Close UNIX/LINUX session**:



HADOOP

This section describes in greater detail the Connection Properties as they apply to **Hadoop** connections.

Hadoop as a source enables RED users to connect to the Hadoop System and process two load types from a Hadoop source into a Teradata repository. These connections must always be set via a **Secure Shell (SSH)** protocol.

Please note that WhereScape RED only fully supports HDFS as the underlying file system.

The two load types that can be processed from RED are:

- **Native SSH Load** - connections to Hadoop on UNIX/Linux from which users do flat file loads. To process a Native SSH load, select **Hadoop** as the connection type set the remaining connection properties.
- **TPT Load** - connections to Hadoop from which users can do TPT script-based loads. To process Hadoop TPT loads, select **Hadoop** as the connection type and set the remaining connection properties.

For **Hadoop TPT Loads**, users will need to have the following system prerequisites before setting up a connection within RED:

- Install **Hadoop**
- Include Hadoop Client Jar files in Classpath

Example for an Apache Hadoop environment setup

```
#Hadoop
export PATH=$PATH:/opt/hadoop-2.4.1/bin:/opt/hadoop-2.4.1/sbin


#For Teradata TPT Load (Hadoop)
export CLASSPATH=$(find /opt/hadoop-2.4.1/share/hadoop/hdfs -name *.jar -printf '%p:' | sed 's/:$/')
export CLASSPATH=$CLASSPATH:$(find /opt/hadoop-2.4.1/share/hadoop/common -name *.jar -printf '%p:' | sed 's/:$/')
```



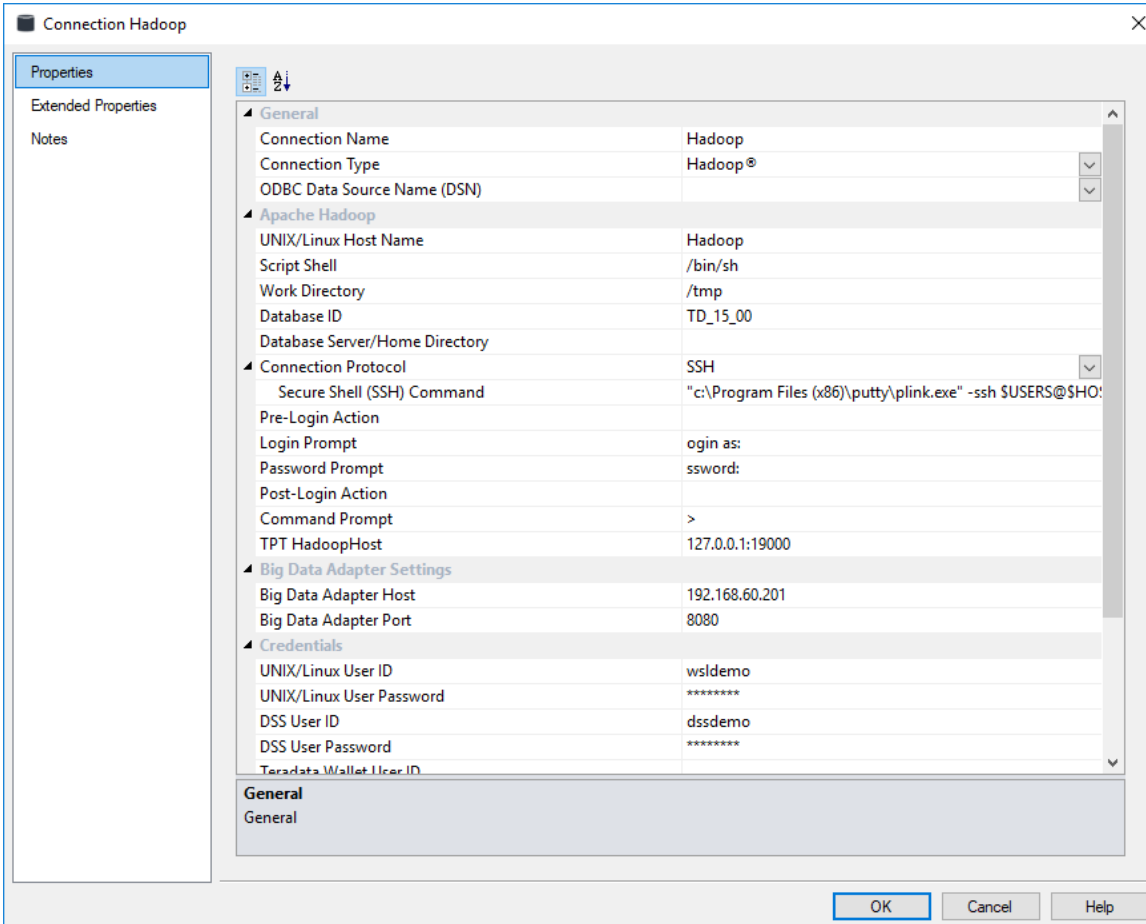
Tip: When the **Big Data Adapter** settings are populated in Hadoop connections, RED can load data from Hadoop into Hive and/or Datawarehouse tables and also perform loads from Hadoop directly into the Datawarehouse using Sqoop through WhereScape RED's Big Data Adapter (BDA).

For more information about these settings, see the **Big Data Adapter Settings** fields description below and see also *Connection to the Data Warehouse* (see "*Database - Data Warehouse/Metadata Repository*" on page 147), *Configuring your database for use by BDA* and *Apache Sqoop Load*.

Note: If the **Hadoop** connection returns a blank screen or an error message in the **Results** pane after the connection is browsed, take necessary action through the **Server (SSH)** tab next to the main Builder and Scheduler tabs.

 This tab is displayed after browsing the Hadoop connection.

Sample Hadoop TPT connection screen:



Connection Hadoop	
Properties	
Extended Properties	
Notes	
General	
Connection Name	Hadoop
Connection Type	Hadoop®
ODBC Data Source Name (DSN)	
Apache Hadoop	
UNIX/Linux Host Name	Hadoop
Script Shell	/bin/sh
Work Directory	/tmp
Database ID	TD_15_00
Database Server/Home Directory	
Connection Protocol	
Connection Protocol	SSH
Secure Shell (SSH) Command	"c:\Program Files (x86)\putty\plink.exe" -ssh SUSERS@SHO:
Pre-Login Action	
Login Prompt	ogin as:
Password Prompt	ssword:
Post-Login Action	
Command Prompt	>
TPT HadoopHost	127.0.0.1:19000
Big Data Adapter Settings	
Big Data Adapter Host	192.168.60.201
Big Data Adapter Port	8080
Credentials	
UNIX/Linux User ID	wsldemo
UNIX/Linux User Password	*****
DSS User ID	dssdemo
DSS User Password	*****
Teradata Wallet User ID	
General	
General	

General

Connection Name

Name used to label the connection within WhereScape RED.

Connection Type

Indicates the connection source type or the connection method such as Database, ODBC, Windows, Unix. Set to **Hadoop**.

Apache Hadoop

UNIX/Linux Host Name

Enter the IP address or host name that identifies the Hadoop server.

Script Shell

Path to the POSIX-compliant UNIX/Linux shell to use for generated scripts. For **UNIX** hosts, set to **/bin/ksh**. For **Linux** hosts set to **/bin/sh**.

If this field is left blank, a default will be chosen based on the name of the connection and the type of database used for the WhereScape RED metadata repository.

Work Directory

Windows directory used by WhereScape RED to create temporary files for minimal logged extracts. The directory must exist and allow write access. There must be a different work directory for each WhereScape RED Scheduler running on the same machine to avoid file conflicts. Typically **C:\Temp** or a sub-directory of **C:\Temp** is used.

Database ID

Enter the Source Database Identifier (Teradata TDPID).

Database Server/Home Directory

Optional to specify the Database Home Directory if it is different from the standard home directory.

Connection Protocol

Telnet or Secure Shell (SSH) protocol to use to connect to the Hadoop machine. For **SSH**, the 'Secure Shell (SSH) Command' property is enabled to specify how to connect.

Secure Shell (SSH) Command

Command to execute to connect to a Hadoop machine using the Secure Shell (SSH) protocol, such as **C:\Program Files\PuTTY\plink.exe -ssh \$HOST\$ -l \$USER\$ -pw \$PASSWORD\$**.

In-built variables **\$HOST\$**, **\$USER\$** and **\$PASSWORD\$** can be used here for the required host, user and password fields.

Pre-Login Action, Login Prompt, Password Prompt, Post-Login Action, and Command Prompt.

These fields are only used to create a Telnet connection to the host machine. WhereScape RED uses the Telnet connection in the drag and drop functionality. **It is not used in the actual production running of the Data Warehouse**, and is only necessary if you wish to use the drag and drop functionality.

Pre-Login Action

Response or command to send BEFORE logging in to the host machine. Typically this is NOT necessary but it can be used to indicate that the host Login Prompt is preceded by a line-feed (\n). However it is preferable that the host login displays the Login Prompt without anything preceding it. [Optional]

Login Prompt

The host login prompt, or the tail end of the login prompt, e.g, **ogin as:**

Password Prompt

The host password prompt, or the tail end of the password prompt, e.g, **ssword:**

Post-Login Action

Not often used but may be necessary to respond to a login question. It is preferable that the UNIX login goes straight to the command prompt.

Command Prompt

Enter the host command prompt, or the tail end of that prompt, typically >.

Note: In order to ascertain some of the above fields you will have to log in to the host system.

TPT HadoopHost

IP address or host name (and optional port number) that identifies the Hadoop Host to a TPT load routine, in order to connect to the Hadoop file system from the machine you run TPT. e.g,

HadoopHost:9000 or **127.0.0.1:9000**.

If this is not specified, then the UNIX/Linux Host Name will be used as the Hadoop Host to the TPT load routine.

Big Data Adapter Settings

Set the two fields below to enable RED to communicate with BDA and enable loading data from Hadoop into Hive and/or into data warehouse tables using Sqoop.

For further information about setting these fields, see *Connection to the Data Warehouse* (see "*Database - Data Warehouse/Metadata Repository*" on page 147) and *Configuring the BDA Server/Configuring your database for use by BDA*.

Big Data Adapter Host

Host machine on which the Big Data Adapter is running its web-server.

Big Data Adapter Port

Port that Tomcat is running. Default is 8080.

Credentials

UNIX/Linux User ID

User Account to login to the UNIX/Linux Host.

UNIX/Linux User Password

Password to login to the UNIX/Linux Host.

DSS User ID

Database User to connect to the WhereScape RED metadata repository.

DSS User Password

Database Password to connect to the WhereScape RED metadata repository.

Teradata Wallet User ID and Teradata Wallet String

Enter the relevant Teradata Wallet credentials instead of the DSS User and Password for the Unix connection if using the Teradata log on method.

Other

Default Path for Browsing

Optional default Path for browser pane filter. When a path has been selected in this field, it becomes the initial point for browsing and it is also expanded on open in the right hand browser pane.

New Table Default Load Type

The default Load type for new tables created using this connection. Select from the **File Load**, **Native SSH** or **Externally loaded** options.

Note: The available options in this drop-down list is configured from Tools>Options>**Available Load Types** (see "**Settings - Available Load Types**" on page 120).

New Table Default Load Script Template

The default Template used for generating Scripts for new Load Tables created using this connection.

Data Type Mapping Set

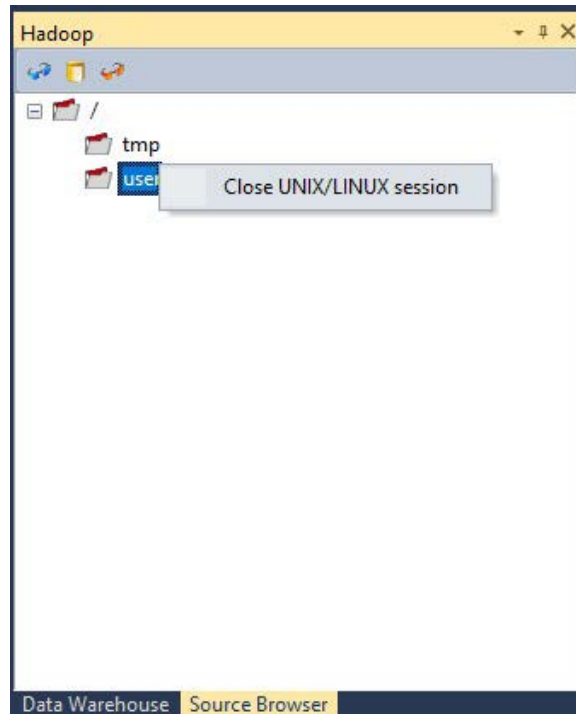
XML files have been created to store mappings from one set of data types to another. Setting this field to **(Default)** makes RED to automatically select the relevant mapping set, otherwise you can choose one of the standard mapping sets from the drop-down list or create a new one.

To test the drag and drop functionality

- From the menu strip select **Browse>Source Tables**
- Drill down to the area required
- Select the object from the left pane and then drag it to the middle pane

Closing the Connection

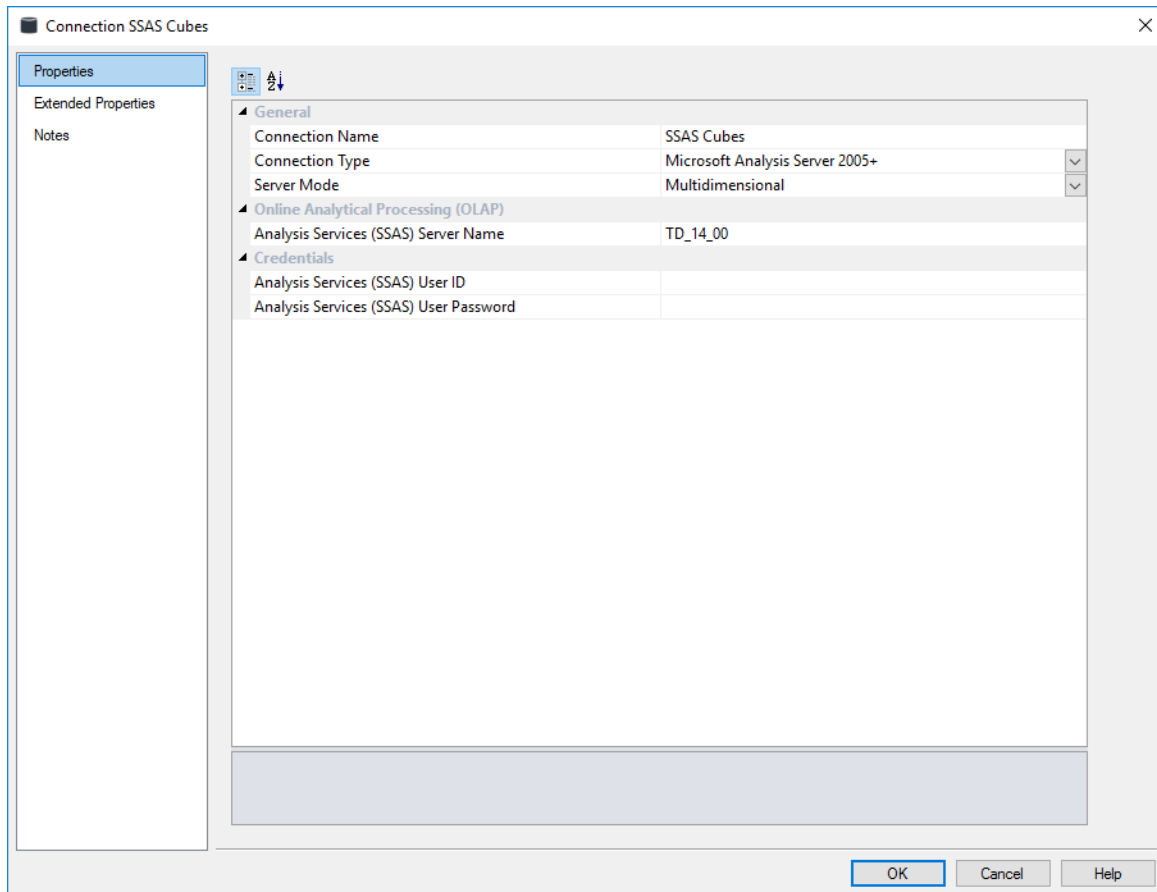
To close the collection, right-click in the browser pane and select **Close UNIX/LINUX session**:



MICROSOFT ANALYSIS SERVER 2005+

A Connection to an **Analysis Services Server** provides the location for cubes defined in the metadata repository.

This connection is used in the creation and processing of cubes. An example screen follows:



Connection Name

Enter a name to identify the connection to the Analysis Services server.

Connection Type

The connection type is chosen from the drop-down list. Select "Microsoft Analysis Server 2005".

Analysis Server

Enter the name of the Analysis Services server you wish to connect to. This must be a valid server name. Contact your system administrator if you do not have a valid server name.

User ID and Password

Not used, leave blank.

MICROSOFT ANALYSIS SERVER 2005+ - OLAP CUBES

A Connection to an **Analysis Services server** provides the location for cubes defined in the metadata repository. This connection is used in the creation and processing of cubes.

A sample screen follows:

The screenshot shows the 'Connection Cubes' dialog box with the following configuration:

Property	Value
Connection Name	Cubes
Connection Type	Microsoft Analysis Server 2005+
Server Mode	Multidimensional
SSAS Client Version	2012
Analysis Services (SSAS) Server Name	WSL-DOC
Analysis Services (SSAS) User ID	
Analysis Services (SSAS) User Password	

Connection Name
Name used to label the connection within WhereScape RED.

General

Connection Name

Enter a name to identify the connection to the Analysis Services server

Connection Type

Indicates the connection source type or method. Select Microsoft Analysis Server 2005+.

Server Mode

The operational mode that Microsoft Analysis Services will use. Select **Multidimensional** from the drop-down list.

SSAS Client Version

Microsoft Analysis Services Client version available for connecting to the SSAS database. It is recommended that the client version matches your database version.

If a 'Fail - missing AMO data provider' message is displayed in the Results pane when attempting to execute the OLAP action, check that the correct SSAS client version is specified and that the respective version of the data provider is installed on the client workstation.

Note: If you have SSAS client version 2008 installed on your computer, WhereScape recommends selecting **2012** for the SSAS client version.

If the required SQL Server Analysis Management Objects (AMO) are missing, see the following article for more information: <https://msdn.microsoft.com/en-us/library/dn141152.aspx>.

Online Analytical Processing (OLAP)

Analysis Services Server Name

Enter the name of the Analysis Services server you wish to connect to.

Credentials

Analysis Services (SSAS) User ID and Password

User Name used to connect to Analysis Services when using SQL Server Authentication. Can be left blank for a trusted connection using Windows Authentication.

Analysis Services (SSAS) User ID and Password

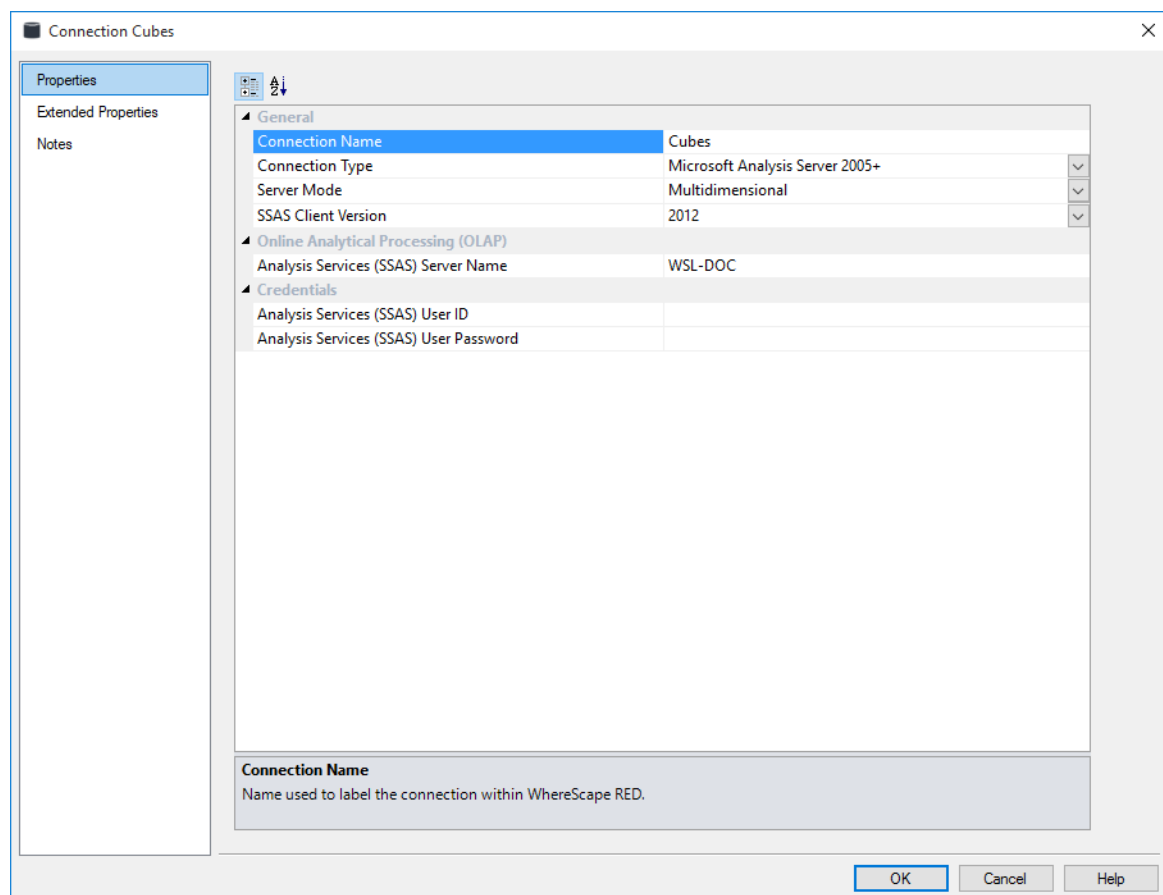
Password used to connect to Analysis Services when using SQL Server Authentication. Can be left blank for a trusted connection using Windows Authentication.

MICROSOFT ANALYSIS SERVER 2005+ - TABULAR MODE

A Connection to an **Analysis Services server** in **Tabular Mode** provides the location for Tabular cubes defined in the metadata repository. This connection is used in the creation and processing of Tabular cubes.

NOTE: Relationships must be created manually for tables stored on a Tabular target, for more information see *Relationship Maintenance* (on page 983).

A sample screen follows:



General

Connection Name

Enter a name to identify the connection to the Analysis Services server

Connection Type

Indicates the connection source type or method. Select Microsoft Analysis Server 2005+.

Server Mode

The operational mode that Microsoft Analysis Services will use. Select **Tabular** from the drop-down list.

SSAS Client Version

Microsoft Analysis Services Client version available for connecting to the SSAS database. It is recommended that the client version matches your database version. If the required SQL Server Analysis Management Objects (AMO) are missing, see the following article for more information: <https://msdn.microsoft.com/en-us/library/dn141152.aspx>.

Online Analytical Processing (OLAP)

Analysis Services Server Name

Enter the name of the Analysis Services server you wish to connect to. You may need to specify the port number of the Analysis Services instance. To find your port number, follow the procedure documented in this Microsoft article: <https://support.microsoft.com/en-us/kb/2466860>.

An example of your Analysis Server (SSAS) Server Name using the port number in RED would be:
VH1D-REDSQL:49449\TABULAR

Credentials

Analysis Services (SSAS) User ID and Password

User Name to connect to Analysis Services with when using SQL Server Authentication. Can be left blank for a trusted connection using Windows Authentication.

Analysis Services (SSAS) User ID and Password

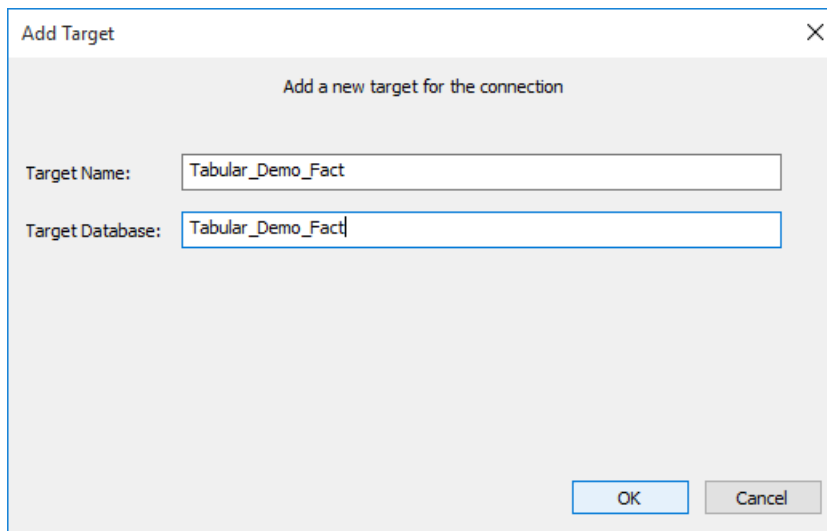
Password to connect to Analysis Services with when using SQL Server Authentication. Can be left blank for a trusted connection using Windows Authentication.

Target Table Location

Add new Target Location

Click the Add new Target Location button to specify the name of the database associated with the targets to be used with this connection. To use the Tabular Mode functionality, it is required to create targets in the Tabular connection.

- The Target Name will be the relevant Tabular Database's name displayed in RED.
- The Target Database will be the relevant Tabular Database's name displayed in Analysis Services.




The screenshot shows a dialog box titled "Add Target" with a close button (X) in the top right corner. The main text inside the dialog reads "Add a new target for the connection". Below this text are two input fields. The first is labeled "Target Name:" and contains the text "Tabular_Demo_Fact". The second is labeled "Target Database:" and also contains the text "Tabular_Demo_Fact". At the bottom right of the dialog are two buttons: "OK" and "Cancel".

BROWSING A CONNECTION

The tables or files associated with a connection can be displayed in the Browser Pane by:


- 1 selecting the **Browse>Data Warehouse** menu option to browse for the data warehouse connection.
- 2 selecting the **Browse>Source Tables** menu option to browse a source system connection.
- 3 right-clicking on a Connection in the Object Pane and selecting **Browse Connection**, or
- 4 clicking on one of the two browser buttons on the toolbar:



The orange button is used to browse the data warehouse connection and the blue button is used to browse a source system connection.

Each button remembers the last connection it browsed, so in this way one button can be used for the Data Warehouse and one for a source system.

Clicking on one of the buttons will display the source tables without first displaying the source browser dialog box. To change the connection, click the small black down arrow next to one of the browser buttons on the toolbar and select Change Connection.

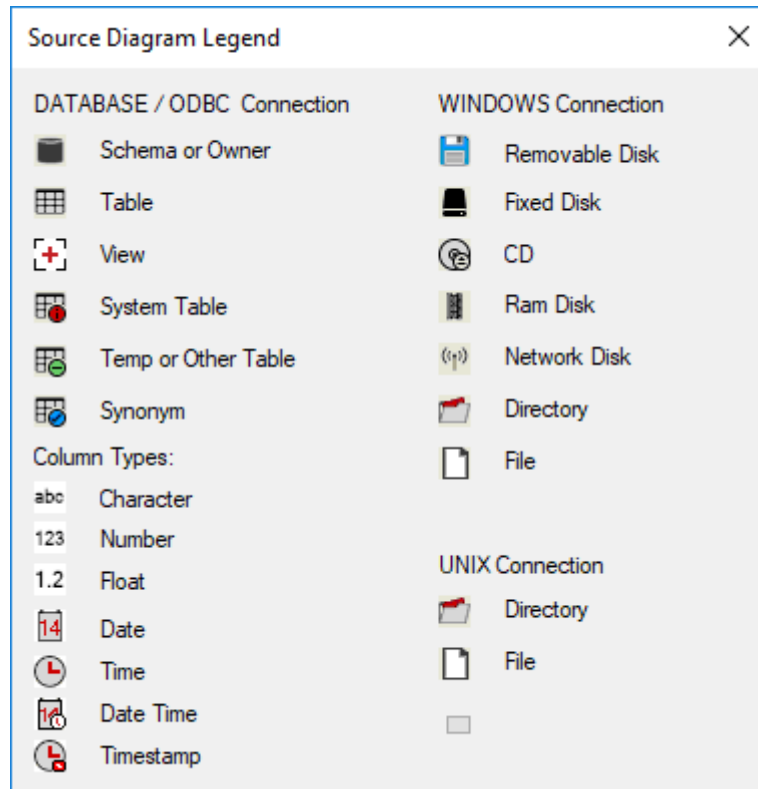


The current connection being browsed is shown in the status bar at the bottom right of the screen.

Browser Icons

When browsing a connection the following legend applies for the source tables and objects.

This legend is displayed via the **Help>Source Legend** menu.



CONNECTION BROWSE PROPERTIES

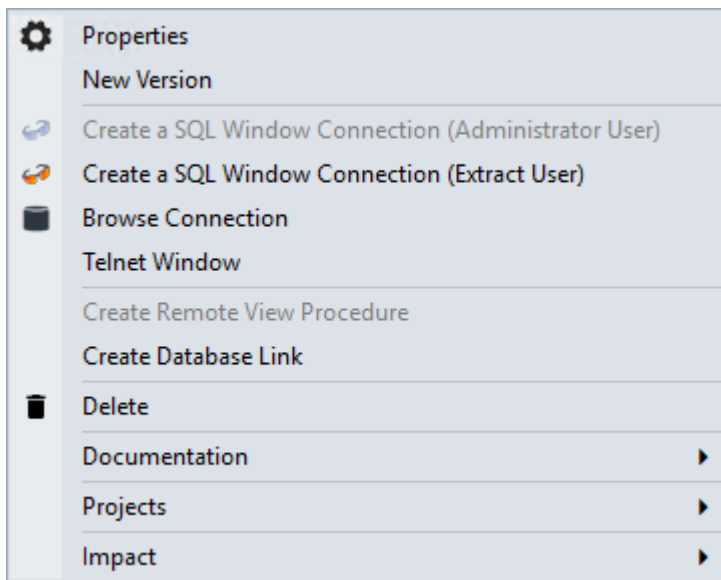


TIP: When browsing to a connection leave the schema field blank in order to see all schemas.


To have RED browsing a specific schema or schemas by default, go to a Connection's Properties screen and enter the schema(s) to browse on the **Default Schema for Browsing** field.

To change the properties of the connection in the Browser Pane:

- Right-click on a connection in the Object Pane and select **Browse Connection**, or



- Click the down arrow next to one of the Browser buttons on the toolbar and select **Change Connection**.



The **List Source Tables Connection** is displayed:

The dialog enables you to change the properties of the connection you are browsing.

The **User ID** and **Password** fields can be changed in order to browse the connection as a different user. A TDWallet string can be supplied as the password by using the **Password is TDWallet string** check-box.

Selecting the **Include Rowcount** check-box displays a row count in brackets next to each source table in the Browser Pane. This is only available for databases which update table statistics.

A **filter** can also be applied when browsing a connection. Filters can be applied to any combination of:

- One or more Schema names (separated by commas),
- a standard SQL table name,
- specific Object Types (Tables, Views or System Objects),
- a Group, or
- a Project

The **Data Type Mapping Set** drop-down can be used to change the data type conversion used during drag and drop operations. If set to (Default) the Data Type Mapping to use for each drag and drop operation is set based on the source and the Target Location selected, but can be changed in the Add New Metadata Object dialog if needed.

CHANGING A CONNECTION'S PROPERTIES

Whenever a Connection's properties are changed, the impact on the objects that use that connection must be considered. Load tables have information from the connection stored within their properties.

This information is stored in the load objects to minimize the complexity of the scheduled tasks. The database link and database name are stored locally in each load table.

When either the database link or database name are changed on a connection, WhereScape RED displays the **Update Associated Load Tables** dialog box.

Click **Yes** to automatically update the database link and/or database name on all associated load tables.

This can also be done manually:

- 1 Double-click on the **Load Tables object group** in the left pane. This displays all Load tables in the middle pane.
- 2 Select those load tables that use the connection that has changed. Use standard Windows selection.
- 3 Right-click to bring up a menu and select **Change Connect/Schema**.
- 4 Select a different connection (e.g. Data warehouse) to change all the selected load tables.
- 5 Repeat step (3) and now change the tables back to the altered connection. This updates all the Load tables with the new connection information.



Note: Whenever a load connection's properties are changed. All load tables that use the connection must be changed. See above. You will be asked if you wish to perform this action when changing the connection.

RESET META DATABASE CONNECTIONS

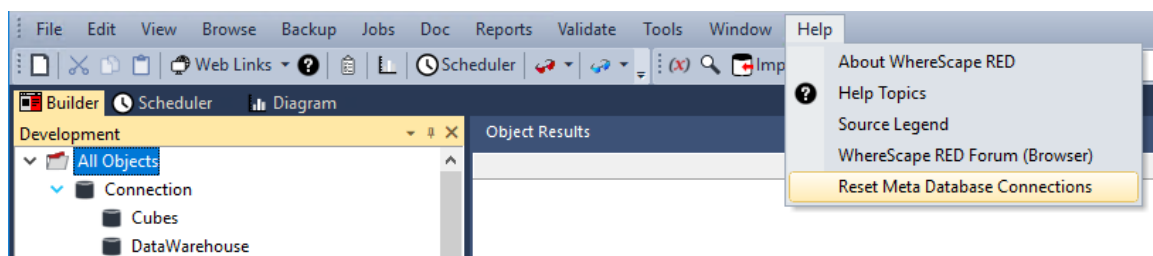
From the **Help** menu, on RED's main top bar, users can select the **Reset Meta Database Connections** option.

This option disconnects and frees most connections that RED has to any existing ODBC connections. This option can be useful for users that are already connected to existing ODBC sources but want to alter the credentials used.


When this option is used, RED attempts to release most existing ODBC connections so the next time an ODBC Connection is used, RED will attempt to reestablish a link.

NOTE: At this stage, not all RED connections are handled via this mechanism and therefore not all connections will be reset when this option is used.

- 1 To reset meta database connections, click the **Help** menu in the main top bar and then click **Reset Meta Database Connections**.



- 2 Click **OK** on the following reset connection dialogs.



CONNECTION SETTINGS FOR BDA

This topic describes the required settings for connections using the WhereScape Big Data Adapter (BDA).

The WhereScape Big Data Adapter (BDA) is designed as an adapter to RED, focused on executing ELT related processing within the Hadoop/Hive eco-system.

For more information about the initial BDA setup, overview of BDA, the prerequisites and step-by-step instructions to set up BDA, please refer to section **18. BDA** of the RED Setup Administrator Guide.

The BDA connection settings are always visible on both Hive and Hadoop connection types, which include the BDA server settings.

BDA settings are also displayed for any other database connection types if there is a Hive or Hadoop connection in the Datawarehouse, but these will only have the JDBC settings displayed.

Important: A Hive connection must exist before BDA settings (**Connection>Properties>Big Data Adapter Settings**) appear in other database connections.

BDA enables RED to use Sqoop as a load method to load data from Hive and HDFS into the Datawarehouse and also to load data into Hive as a target.

Only one BDA server connection is supported per Metadata repository.

CONFIGURING THE BDA SERVER

This topic explores the configuration of the BDA server in Hadoop connections to enable loading data from Hadoop into Hive and/or Datawarehouse tables as a source database into RED Datawarehouse tables using Sqoop.

For RED to be able to load data from Hadoop into Hive and/or Datawarehouse tables and also perform loads from Hadoop directly into the Datawarehouse using Sqoop, the **Big Data Adapter** settings need to be populated in Hadoop and Hive connections.

For more information about what is required for loading data into RED using Sqoop, see the **Big Data Adapter Settings** fields description below and see also *Connection to the Data Warehouse* (see "*Database - Data Warehouse/Metadata Repository*" on page 147), *Connections to Hadoop* and *Apache Sqoop Load*.

Hadoop connection example

The screenshot shows the 'Connection Hadoop' dialog box with the following settings:

Property	Value
Connection Name	Hadoop
Connection Type	Hadoop®
ODBC Data Source Name (DSN)	
Apache Hadoop	
UNIX/Linux Host Name	Hadoop
Script Shell	/bin/sh
Work Directory	/home/wsl_user/temp
Database ID	TD_14_00
Database Server/Home Directory	
Connection Protocol	
Secure Shell (SSH) Command	"C:\Program Files (x86)\putty\plink.exe" -ssh SHOSTS -ISUS
Pre-Login Action	
Login Prompt	ogin as:
Password Prompt	ssword:
Post-Login Action	
Command Prompt	>
TPT HadoopHost	127.0.0.1:19000
Big Data Adapter Settings	
Big Data Adapter Host	192.168.60.24
Big Data Adapter Port	8080
Hadoop Connectors	
Credentials	
UNIX/Linux User ID	wsldemo
UNIX/Linux User Password	*****
DSS User ID	dssdemo
DSS User Password	*****

Buttons: OK, Cancel, Help

BIG DATA ADAPTER SETTINGS

Big Data Adapter Host

Host machine on which the Big Data Adapter is running its web-server.

Big Data Adapter Port

Port that Tomcat is running. Default is 8080.

CONFIGURING YOUR DATABASE FOR USE BY BDA

Connections using BDA enable loading data into Hive as a target database and they also enable loading data from Hive into Datawarehouse tables as a source, using the Apache Sqoop load method.

This connection type needs to include the JDBC connection string (JDBC URL) and related attributes (username, password) for the Hive database. The JDBC User and Password is usually the same as the Extract User ID but users can specify different credentials if necessary.

The Big Data Adapter settings will also need to be populated in the Datawarehouse connection. For more details, see **Connection to the Data Warehouse** (see "**Database - Data Warehouse/Metadata Repository**" on page 147).

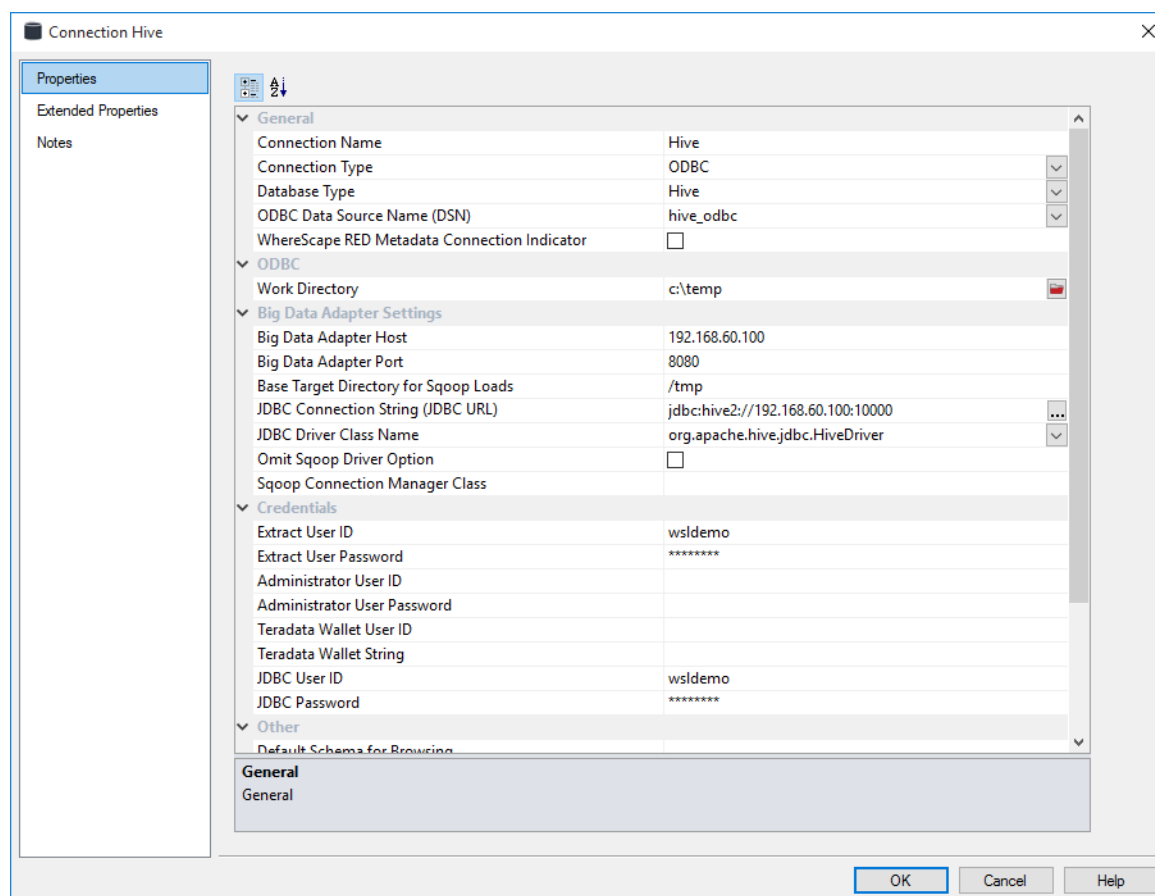
RED can also load data directly into Hive from any database source. This load can also be processed via an Apache Sqoop load and the JDBC settings on the Hive connection will need to be populated. Please see the connection example and field description below for more details about this.

When loading data into Hive as a target, users can also add specific target locations in their Hive ODBC connections, if they have a Hive target license enabled.

Hive connection properties will be the same for any database sources.

Note: When creating objects in a Hive target database from a Teradata repository, artificial keys must be of type 'int', not 'integer'. To correct this, go to **Tools -> Options -> Global Naming Conventions -> Global Naming of Key Columns** and change the relevant Data Types to 'int'.

Example of a Hive ODBC connection



BIG DATA ADAPTER SETTINGS

Big Data Adapter Host

Host machine on which the Big Data Adapter is running its web-server.

Big Data Adapter Port

Port that Tomcat is running. Default is 8080.

Base Target Directory for Sqoop Loads

HDFS directory in which to create target directories for Sqoop loads using the Big Data Adapter.

JDBC Connection String (JDBC URL)

Connection string used by the WhereScape Big Data Adapter to access this database.

JDBC Driver Class Name

JDBC driver class to be used by the WhereScape Big Data Adapter. This field must be set if the JDBC URL is set.

Select the appropriate JDBC Driver class name from the drop-down list. If this is left empty this will not be specified in generated commands.

Omit Sqoop Driver Option

If set, the --driver option to Sqoop will be omitted. This is required for certain connection types such as Oracle connections.

If you select the **Omit Sqoop Driver Option** check-box, the driver parameter will not be used in sqoop command line. This is a requirement for Oracle at the moment, as suggested by Sqoop documentation for 1.4.5.

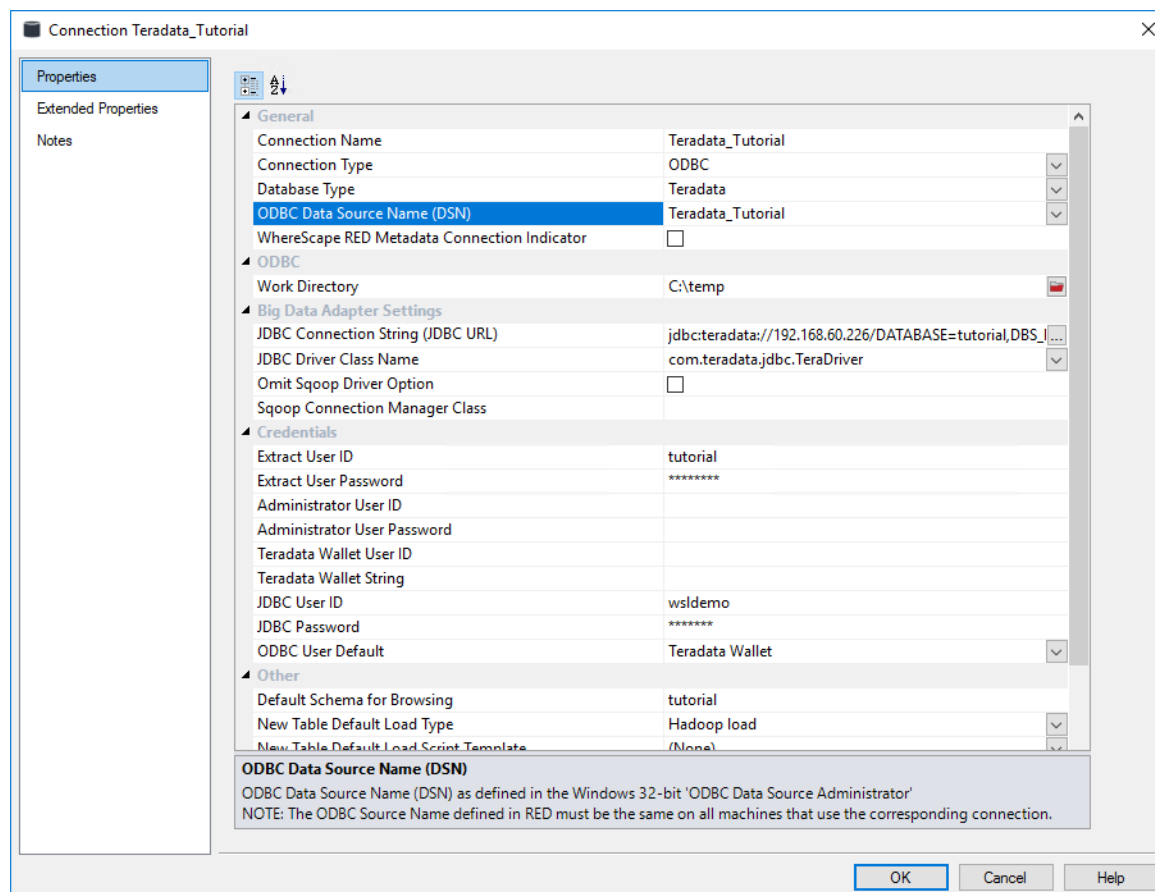
Sqoop Connection Manager Class

Custom Sqoop connection manager class. Corresponds to the --connection-manager command line argument. Leave blank if this is not required.

Example connection from a source database to Hive

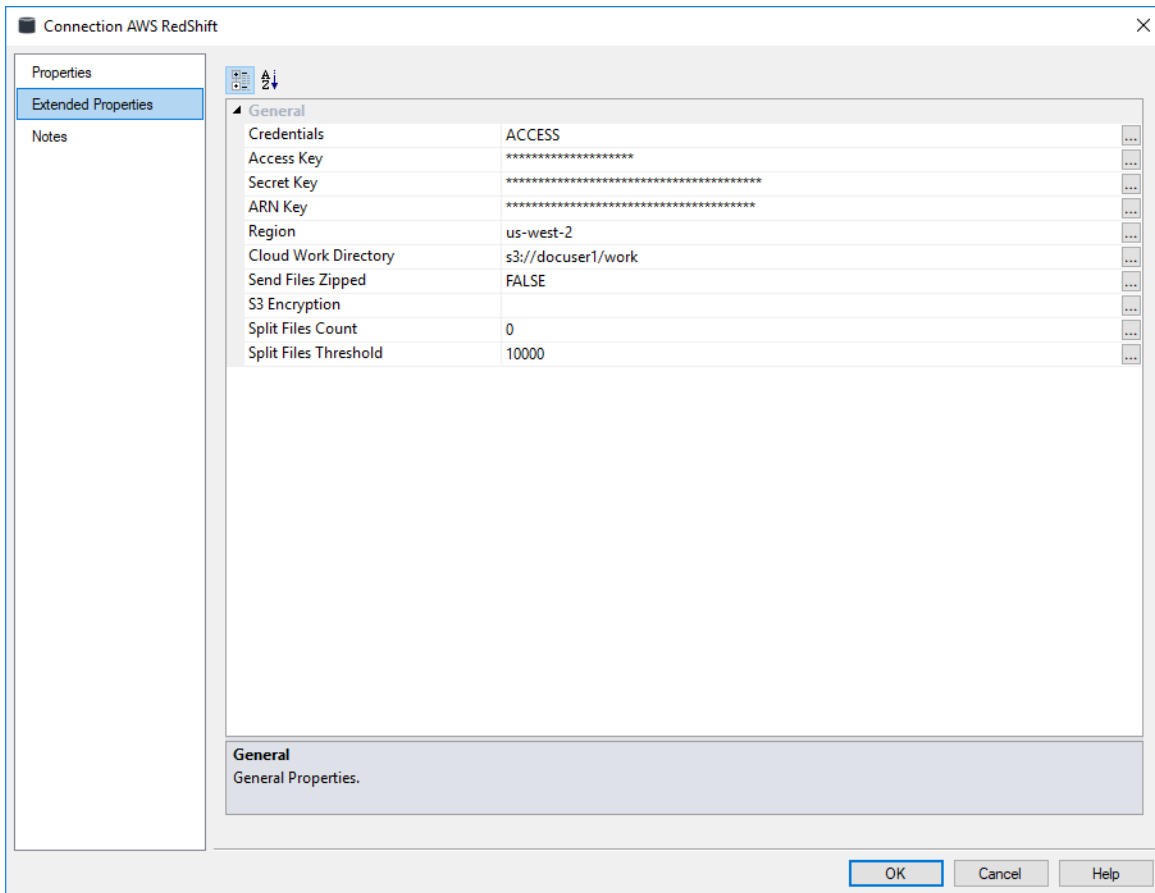
The full JDBC connection string is:

```
jdbc:sqlserver://192.168.60.100/DATABASE=SQL_Hive,DBS_PORT=1025
```



CONNECTION EXTENDED PROPERTIES

Note: These settings enable you to assign extended property values for a connection. Refer to the *Extended Properties* (on page 1065) chapter for more information.



CHAPTER 8

TABLE PROPERTIES

Various properties can be set on all table objects in WhereScape RED. The screens available in the Table Properties Dialog depend on the object type selected and can be a subset of:

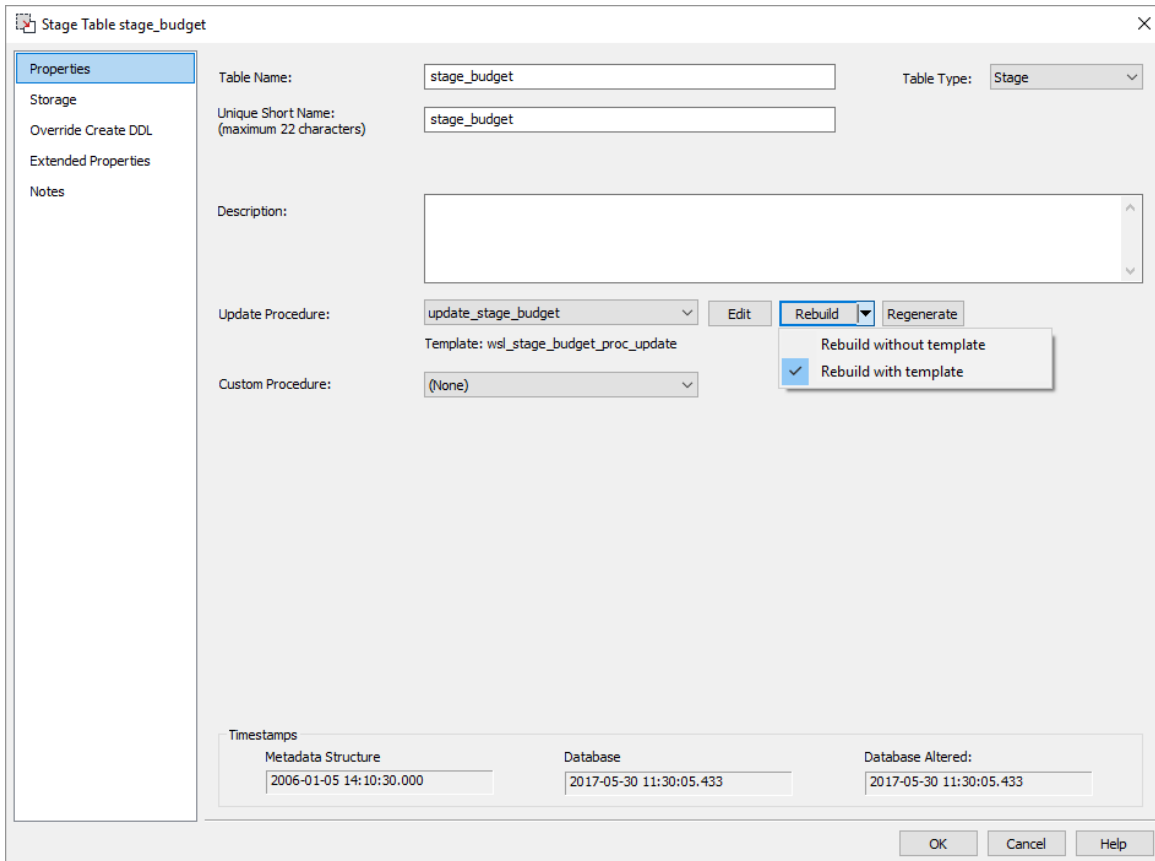
- **Properties** (on page 197)
- **Storage** (on page 202)
- **Override Create DDL** (on page 209)
- **Source** (on page 210)
- **Documentation Fields** (on page 210)
- **Notes** (on page 212)

IN THIS CHAPTER

Properties	197
Storage	202
Override Create DDL	209
Source	210
Documentation Fields	210
Table Extended Properties	211
Notes	212

PROPERTIES

The fields available on the Properties screen depend on the Object Type selected. More specific information is available in the respective chapters describing each Object Type. Rebuilding Tables using Update Procedures and templates is applicable to the objects listed and described in the table below.



Fields	Description
Table Name	The user-name of the selected table.
Unique Short Name	The short name is derived from the Table Name and is used internally by RED.
Table Type	The drop down list provides a list of the available sub types for the selected table type.
Description	Optional, free text.
Update Procedure	The name of the procedure to be used when updating the table.
Custom Procedure	The name of the procedure to be used for custom updates of the table.
Edit	Edit the content of the displayed update procedure.

Fields	Description
Rebuild	Rebuild the Update Procedure. See <i>Rebuilding Update Procedures</i> (on page 199) for more details.
Regenerate	Regenerate the selected Update Procedure using the responses to previously provided information.

REBUILDING UPDATE PROCEDURES

The update procedure for a data warehouse object can be generated by RED using parameters provided by the user, or it can be generated using a prepared template. Templates appropriate to the table type must first be created before they can be selected and used to generate update procedures. The option to select a template is only available if RED detects the presence of an appropriate template.

Note: Update procedures for **Data Vault** and **Source Mapping** objects can only be generated using a template. Refer to *Generating Update Procedures for the Data Vault Stage Table* (on page 436) and *Generating Update Procedures for Source Mapping Objects* (on page 1115) respectively for details.

Generating an Update Procedure without a template

If RED does not detect a template, the drop-down option is not displayed. RED manually regenerates the update procedure when **Rebuild** is clicked, by prompting the user to respond to a series of input requests. Refer to the **Generating Update Procedure** topic in each of the data warehouse object chapters for further information and the steps to manually generate update procedures.

Generating an Update Procedure via a Template

If RED detects the presence of a template, the **Rebuild** button provides additional options, in a drop down list, for rebuilding the selected update procedure. Two possible rebuild options are provided:

- **Rebuild** - RED rebuilds the procedure using the last selected option as a default when the Rebuild button is clicked. If a template has not been previously used then RED prompts you for inputs as required for the rebuild.
- **Rebuild without template** - RED rebuilds the update procedure, but ignores any previously used template. RED prompts you for inputs as required for the rebuild.
- **Rebuild with template** - Prompts you to select from a list, a template that is appropriate to the current object type. Depending on the selected template, the user is prompted to provide responses that will be used by the parameters within the template.

Notes:

1. If a template has been previously used, RED uses this template by default when **Rebuild** is clicked. The name of this template is displayed below the **Update Procedure** field of the table Properties screen.
 2. During procedure generation, a message prompt is displayed if any source tables are missing in the Source Join property of the **Source** tab on the **Update Build Options** screen. Select from the following options:
 - **Yes** to return to the **Source Join** screen and edit the **Source Join** property.
 - **No** to continue with the procedure generation.
 - **Cancel** to cancel the procedure generation.
-

The following table summarizes the conditions that will display the drop-down options in the **Rebuild** button or that require a template. All other conditions will require a manual rebuild.

Table Object	SQL Server	Oracle	Teradata	DB2	Greenplum	Netezza	Hive/PDW/Custom
Stage	Both	Both	Both	Both	Both	Both	Template Only
Data Vault Stage	Template Only	Template Only	Template Only	Template Only	Template Only	Template Only	Template Only
Fact	Both	Both	Both	Both	Both	Both	Template Only
Fact Rollup	Both	Both	Red Automation Only	Both	Red Automation Only	Red Automation Only	Template Only
Fact Kp1	Both	Both	Red Automation Only	Both	Red Automation Only	Red Automation Only	Template Only
Aggregate	Both	Both	Both	Both	Red Automation Only	Red Automation Only	Template Only
Dimension	Both	Both	Both	Both	Both	Both	Template Only
Data Store	Both	Both	Both	Both	Both	Both	Template Only
Normal, Hub, Link, Satellite, Custom1 or 2	Both	Both	Both	Both	Both	Both	Template Only

Both	This object type / database combination supports both RED automation for code generation and template based code generation.
RED Automation only	This object type / database combination uses RED automation for code generation and template based code generation is not currently available.
Template only	This object type / database combination only supports template based code generation.

STORAGE

The Storage tab of the table Properties window displays the options applicable for storing data in the associated RDBMS.

The fields available on the Storage tab depend on the RDBMS on which you are storing the data:

For a Teradata example, see *Table Storage Screen - Teradata*.

For a Tabular example, see *Table Storage Screen - Tabular* (on page 205).

For information on changing storage locations for multiple tables at once, see *Bulk Table Storage Change* (on page 206).

TABLE STORAGE SCREEN - TERADATA

Typical Storage screen for a **Teradata** Table:

Location	
Target Location	DataWarehouse:doc6871_LOAD
Database Type	Teradata
Database	doc6871
Temp Database	
Create DDL Template	None

Storage	
Primary Index	Non-Unique Primary Index (NUPI)
Primary Index Name	load_product_idx_PR
Primary Index Columns	product_code
MultiSet	<input checked="" type="checkbox"/>
Fallback	<input type="checkbox"/>
Data Block Size	
Enable Free Space	<input type="checkbox"/>

Other	
Optional CREATE Clause	

Location
The location of the table

LOCATION

Target Location

The target location that defines the path to the location for the table. Select (local) for a local table (see note below) or select the target schema, if you are placing tables in different schemas.

Note: The **Force Target Usage** setting in the **Tools>Options>Storage>Target Usage** removes the local target option from the drop-down list. See *Target Usage* (on page 112) for details.

To add another database/schema to the list see more details on *Connections to the Data Warehouse/Metadata Repository* (see "Database - Data Warehouse/Metadata Repository" on page 147). To set default target locations for tables see *Settings - Storage: Target Location* (see "Target Location" on page 113).

Database Type

The database type for a connection that is used for target Data Warehouse tables.

Database

The database where the table is located. Leave blank to use the default for the connection or local environment.

Temp Database

The database where temporarily created tables are located. Leave blank to use the default for the connection or local environment. This field is only used in Load and Export processes.

Create DDL Template

Optional. Specify the template to use when creating a new DDL procedure script. This option is only visible if a DDL template is available for this database type. Default value is *None*.

NOTE: Since Teradata does not support the moving of tables, all affected tables will also need to be manually recreated after any storage changes.

WARNING: Please note that changing the Storage for Dimension and Fact tables will need to be handled very carefully as artificial key relationships between Dimension and Fact could become out of sync.

Recreating Fact Tables and large Dimension tables might take a considerable amount of time.

STORAGE

Primary Index

Select the Primary Index Type.

Primary Index Name

Name of the Primary Index.

Primary Index Columns

Columns of the Primary Index.

MultiSet

This options makes this table multiset. A Multi Set table allows for duplicate rows.

Fallback

This option enables Fallback. A fallback table is a duplicate copy of a primary table. Each fallback row in a fallback table is stored on a different AMP to the one used for the corresponding row in the primary table. The default is no fallback.

Data Block Size

The block size for the table. Default is blank.

Enable Free Space

Use Database default Free Space setting. Disabling this will enable setting of the Free Space. Default is blank.

Free Space

The amount of free space to leave in the table for future update. Enter a value between 0 and 75 for a percentage or to use the Database default set the above option.

OTHER

Optional CREATE Clause

Database-specific-and-compliant DDL to append to the generated CREATE TABLE statement.

TABLE STORAGE SCREEN - TABULAR

Typical Storage screen for a **Tabular** Table:

The screenshot shows a configuration window titled "Dimension tab_dim_customer_v12". On the left is a sidebar with a tree view containing: Properties, Storage (selected), Extended Properties, Purpose, Concept, Grain, Examples, Usage, and Notes. The main area displays the following settings:

Location	
Target Location	Tabular:Tabular_Demo_Fact
Database Type	Tabular
Database	Tabular_Demo_Fact
Create DDL Template	None
Other	
Hidden	<input type="checkbox"/>
Processing	
Processing Type	Regular

At the bottom of the main area, there is a section labeled "Location" with the text "The location of the table". At the bottom right of the window are three buttons: "OK", "Cancel", and "Help".

LOCATION

Target Location

The path for the MSAS target. To add another MSAS target see *Microsoft Analysis Server 2005+ - Tabular Mode* (see "*Microsoft Analysis Server 2005+ - OLAP Cubes*" on page 179) . To set default target locations for tables see *Settings - Storage: Target Location* (see "*Target Location*" on page 113).

Note: The **Force Target Usage** setting in the **Tools>Options>Storage>Target Usage** removes the local target option from the drop-down list. See *Target Usage* (on page 112) for details.

Database Type

For information purposes only, this displays Tabular, for MSAS Tabular targets.

Database

For information purposes only, this displays the database name.

OTHER

Hidden

Specifies whether the table is hidden from reporting client field lists.

PROCESSING

Processing Type

Set the database reporting tools processing to Regular or LazyAggregations in the target Analysis Services database.

BULK TABLE STORAGE CHANGE

Table Storage locations can be changed through the **Storage** tab on a table's Properties screen but they can also be **changed in bulk** by using the following process:

- 1 Double-click on the desired **object group** in the left pane. This displays all the tables in that group in the middle pane.
- 2 Select the tables that you wish to change the storage for using standard Windows selection.
- 3 Right-click to bring up a menu and select **Storage**.

Stage Table			
Stage Table Name	Short Name	Stage Table Type	Filegroup
stage_budget	stage_budget	Stage	
stage_customer	stage_customer	Stage	
stage			
stage			
stage			

Properties
Storage

- 4 On the **Target Location Selection** screen, select the new target location from **Target Location** drop-down list to change the location of all the selected tables in bulk.

Target Location Selection

Storage

Current Locations

Target Location	Multiple Target Locations
Database Type	Teradata
Database	WSL_REPOSITORY

Location

Target Location	(Local)
-----------------	---------

Current Storage

(No Change)

Database

(Local)

Temp Database

DataWarehouse:Load

Storage

Hive:Hive

Database

(No Change)

Target Location
The Target Location that defines the path to the location for the table.
Dialog Opening Value: (No Change)

OK Cancel Help

Note: The **Force Target Usage** setting in the **Tools>Options>Storage>Target Usage** removes the local target option from the drop-down list. See **Target Usage** (on page 112) for details.

- 5 Follow the next dialogs to complete the bulk storage change.
Please note that **all procedures** from the affected tables need to be **manually changed or regenerated** after a bulk storage change.
- 6 Since Teradata does not support the moving of tables, all affected tables also need to be manually recreated after any storage changes.

WARNING: Please note that changing the Storage for Dimension and Fact tables need to be handled very carefully as artificial key relationships between Dimension and Fact could become out of sync.

Recreating Fact Tables and large Dimension tables might take a considerable amount of time.

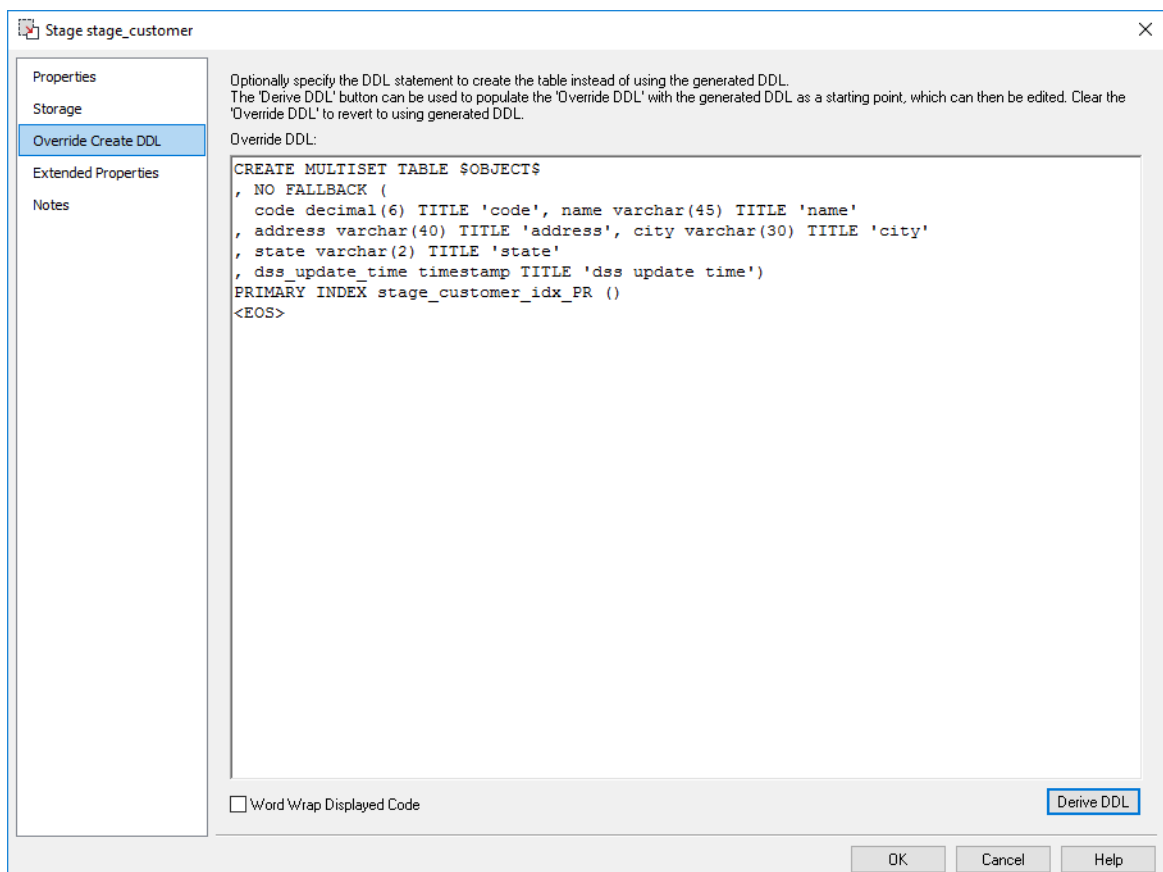
OVERRIDE CREATE DDL

WhereScape RED allows to create tables using a specific DDL statement instead of the RED generated DDL.

You can override and edit the created DDL on load, stage, data store, EDW 3NF, dimension, fact, aggregate, user defined view and retro table's **Override Create DDL** tab.

Creating a table using a specific DDL statement:

- 1 If you are creating a **new** table, drag and drop the table from the right pane to the middle pane.
 - On the Properties dialog, click on the **Override Create DDL tab**.
 - Click on the **Derive DDL** button at the bottom of the dialog to get the generated DDL as your starting point.
 - Edit or clear the generated DDL and enter the desired DDL..
- 2 If you are **editing** an existing table, double click on the table object on the left pane to open the Properties dialog.
 - Click on the **Override Create DDL tab**.
 - Click on the **Derive DDL** button at the bottom of the dialog to get the generated DDL as your starting point.
 - Edit or clear the generated DDL and enter the desired DDL.



- 3 The "end of statement" indicator is <EOS> by default but can be configured in **Tools>Options>Code Generation>General**.



TIP: To revert to RED deriving the original generated DDL at runtime, leave the Override DDL text box blank or clear out the contents. Clicking the **Derive DDL** button at the bottom of the screen pops-up a warning message, asking if you want to replace the current DDL definition with new DDL.

SOURCE

The Source tab is only available for Load Tables. More information is available in the **Loading Data** (on page 215) chapter, or more specifically:

- **Database Link Load - Source Screen** (on page 223)
- **Native ODBC Based Load - Source Screen** (on page 226)
- **TPT Load - Source Screen** (on page 235)
- **Loading Data into RED Load Tables using SSIS** (on page 245)
- **Flat File Load - Source Screen**

DOCUMENTATION FIELDS

Each table has a set of associated documentation fields. These are used to store descriptive metadata about the tables and how they are used. The contents of each field appears in the generated documentation. Their order and placement can be customized in the generated documentation using the **Tools>Options** menu.

The following fields are available to store additional metadata information that will appear in the generated WhereScape RED documentation:

- Purpose
- Concept
- Grain
- Examples
- Usage

NOTE: The names of these fields are completely arbitrary and can be changed in the generated documentation in the **Tools>Options>Documentation** menu.

By default, these fields are not enabled for Load and Stage tables, but can be enabled using the **Tools>Options>Documentation** menu.

Up to 4000 characters of information can be stored in each field. Some or all of these fields can be removed from the documentation via the **Tools>Options>Documentation** menu.

DOCUMENTATION FIELDS SCREEN


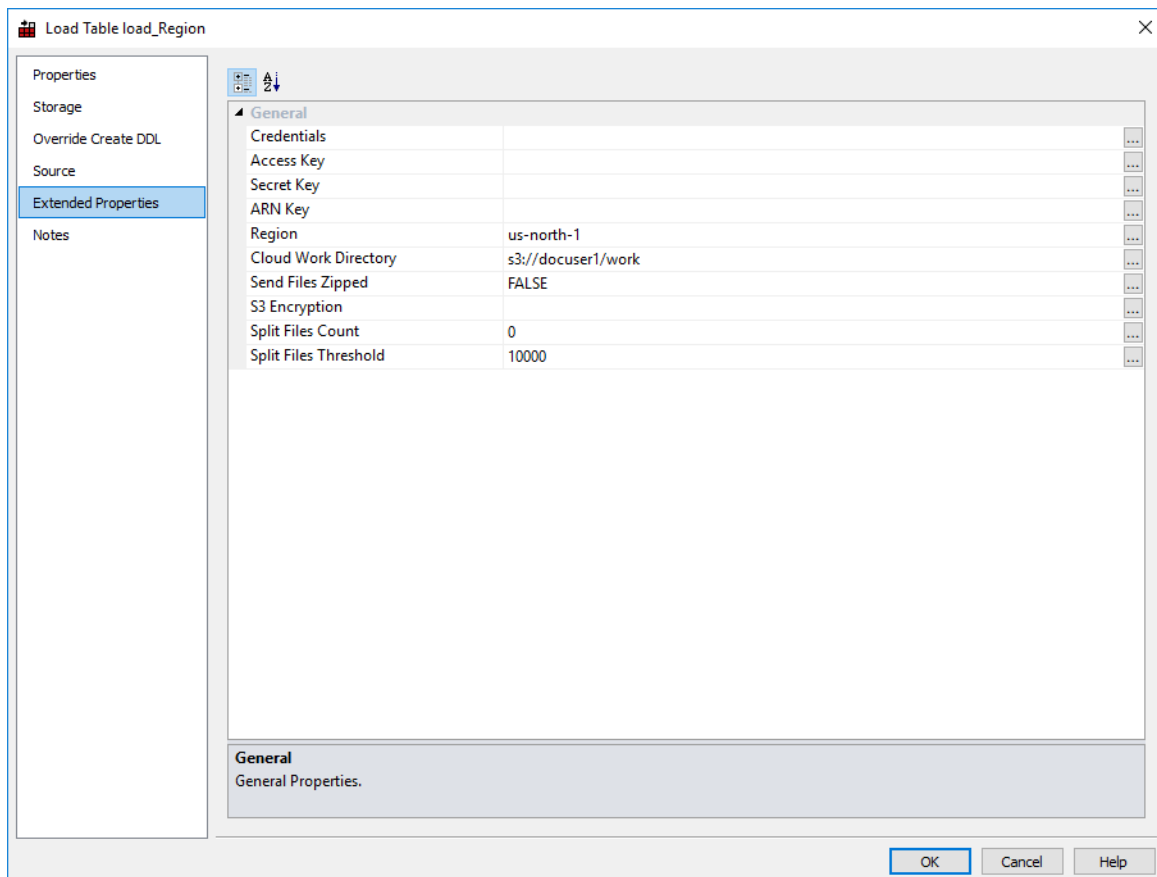


TABLE EXTENDED PROPERTIES

Note: These settings enable you to assign extended property values for a table object. Refer to the *Extended Properties* (on page 1065) chapter for more information.




NOTES

The Notes screen is used to enter notes against WhereScape RED objects.

The notes on an object are included in the generated documentation. Up to 4000 characters of information can be stored in the Notes field.

By default, the Notes field is enabled for all object types.

The Notes field can be removed from the documentation via the **Tools>Options>Documentation** menu.



CHAPTER 9

LOADING DATA

Load tables are the first entry point for all information coming into the data warehouse. There are multiple ways of getting data into load tables.

- Database link load - loaded from another database.
- Externally loaded - the load table is populated by some external process, e.g. an ETL (Extract, Transform and Load) tool or EAI (Enterprise Application Integration) tool, putting the data directly into the load tables.
- ODBC based load - the data is loaded via an ODBC connection, either directly by reading and inserting each row (a Load Type of **ODBC load**), or via a file by reading from the source system using ODBC and writing to a file and then loading the file (a Load Type of **Native ODBC**). ODBC connections require a Windows scheduler.
- File load - a flat file load where most of the processing is managed and controlled by the scheduler.
- Script-based load - a flat file load where a host system, e.g. UNIX or Windows script file is executed to perform the load. Script-based loads on Windows supports both DOS Batch and *PowerShell scripts* (see "*Windows PowerShell Scripts*" on page 653).
- XML file load - loading an XML file from a Windows directory.

Notes:

The following load types are **not** enabled by default for newly created Teradata repositories:

- Integration Services (SSIS) load in Database, ODBC and Windows connections.
- Native ODBC load in an ODBC connection.
- XML file load in a Windows connection.

Refer to Tools>Options>*Available Load Types* (see "*Settings - Available Load Types*" on page 120) for details.

IN THIS CHAPTER

Choosing the Best Load Method	217
Load Drag and Drop	218
Database Link Load	221
ODBC Based Load	225
Native ODBC Based Load	225
TPT Load	233
TPT UNIX/Linux Script Load.....	240
SSIS Loader	243
Flat File Loads	253
XML File Load	277
External Load	281
Apache Sqoop Load	281
Handling Missing Source Columns	288
Load Table Transformations	290
Changing Load Connection and Schema	291

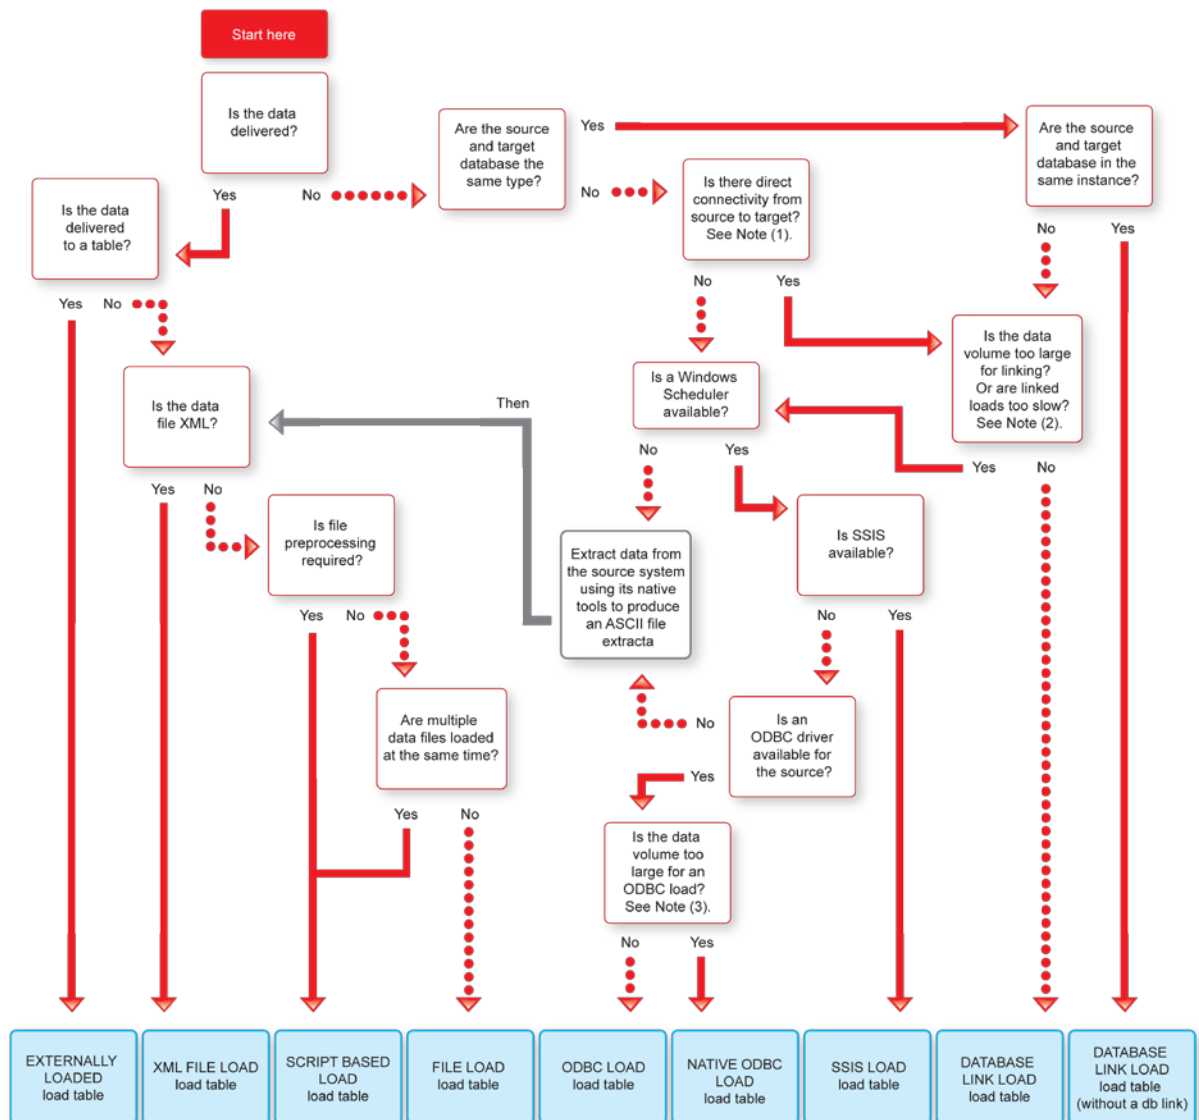
CHOOSING THE BEST LOAD METHOD

Several different factors need to be considered when choosing the best type of load table to use:

- Source and target database types
- Locations of source and target databases
- Available connectivity options and performance
- Amount of data being loaded
- Is the data delivered or fetched?
- For delivered data, what format is it in and does it require processing to make it loadable?

Load Table Decision Tree

The following decision tree can be used when selecting the best type of load table:



LOAD DRAG AND DROP

The simplest way to create a load table is to use the drag and drop functionality of WhereScape RED.

Drag and drop can be used for all connection types and the process is the same in all cases.

- 1 Browse to the source system connection (**Browse>Source Tables**).
- 2 Create a drop target by double-clicking on the Load Table object group in the left pane. The middle pane should have a column heading of **Load Table Name** for the leftmost column.
- 3 Select a table or file in the right pane and drag it into the middle pane. Drop the table or file anywhere in the middle pane.
- 4 Answer the resulting prompts to create the load table. See the tutorials for examples on how to create load tables using drag and drop.
- 5 When using targets, you can also change the predefined **Target Location** settings in **Tools>Options>Storage** (see "**Target Location**" on page 113). At the time of the table's drag and drop, the **Add a New Metadata Object** dialog enables editing the **Target Location** and **Data Type Mapping**, so the table's location can be changed on a table by table basis. The **Data Type Mapping** field is automatically set, based on the source and the **Target Location** selected, but can be changed if needed.

NOTE: The option **Add meta data columns to table** is used for creating load tables that are used in creating Data Vault objects. If this option is selected, two DSS columns (**dss_record_source** and **dss_load_date**) are included in the meta data for the table and are populated by transformations. These two DSS columns could equally be applied to other load tables not used in a Data Vault system but are particularly important to comply with the Data Vault standards. Please refer to the *Data Vaults chapter* (see "**Data Vaults**" on page 413) for more details.


WhereScape RED supports loading tables of up to 2048 columns, however this maximum number of column loading restriction can in fact be lower than the target database or tools permit. The target database will provide users the appropriate warning if the maximum number of columns loaded is breached at runtime.

Notes: When creating a load table in WhereScape RED by dragging over a source table, RED will read the structure of the table on the source system and attempt to construct an equivalent load table in the data warehouse. There are occasions when the load table creation will fail due to incompatible data types in the target data warehouse. The remedy is to change the data types of the particular attributes which are causing the load failure. Once corrected the load table should create. It is important that the table load is tested to ensure that data can be INSERTED into the load table from the source table. If the load fails then the data may need to be explicitly converted to the destination data type using a column transformation that is executed during the load (see *Load Table Column Transformations* (on page 597)). Incompatible data types that cause load table creation errors are typically caused by:

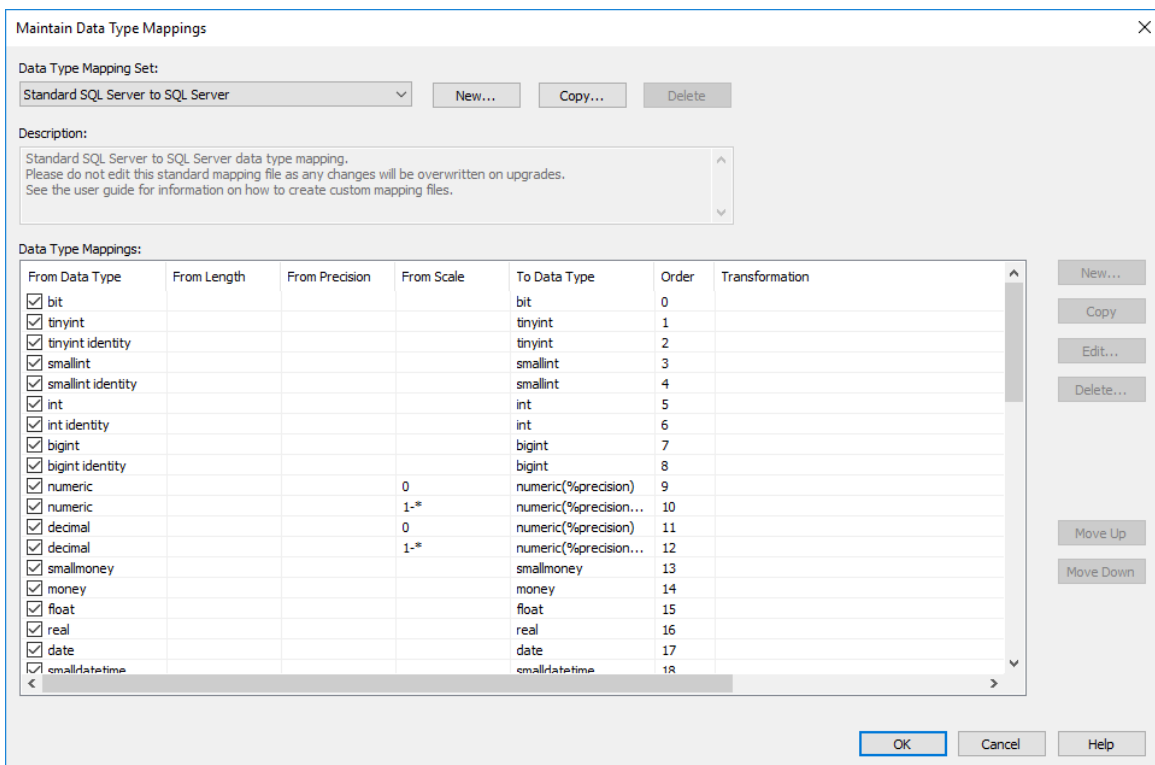
1. User defined data types in the source database.
 2. Incorrect data type mapping during load table definition in WhereScape RED.
 3. Data types that cannot be inserted into (e.g. identity columns on SQL Server)
-

DATA TYPE MAPPINGS

In **Tools>Data Type Mappings**, the first menu option enables you to maintain the data type mappings.



The user interface for maintaining data type mappings is as follows:



See **Data Type Mappings** (on page 991) for a detailed explanation of Data Type Mappings.

DATABASE LINK LOAD

Database Link Load would have the data loaded from another database located on the same Teradata server into the data warehouse. This load type is used infrequently on Teradata.

DATABASE LINK LOAD - PROPERTIES

The fields of the **Load Table Properties** screen for database link loads are described below:

Load Table Name

The table name is limited by Teradata to a maximum of 30 characters and must be unique. Table name defaults can be set in **Tools > Options** to define a prefix or a post fix that can be added in order to identify clearly that this is a load table. Example: load_customer.

By default, RED uses the prefix **load_** for load tables.

Unique Short Name

The table short name is limited in size to 22 characters and must be unique. The short name is used in the naming of indexes, keys and procedures.

Description

Enter here a description of the table. This description appears in the documentation that can be generated once the data warehouse is built.

Connection

Enter the connection being used to get the data. The connections for load tables can be changed on bulk see *Changing load Connection and Schema* (on page 291).

Load Type

The load type is typically defined by the connection, and should not normally be changed. This drop-down shows all valid load types for the connection.

Database Link

This field is not active in WhereScape RED for Teradata.

Script Template

The script template used for a script based load.

Script Name

This field is only active for script based loads.

Pre-load Action

Select an action to be performed on the table before the load occurs. Options are:

- Truncate
- Execute pre-load Sql
- Both Truncate and Execute pre-load Sql
- No action

Pre-load Sql

If a Pre-load Action of **Execute pre-load Sql** was selected, then the Sql statement to execute before the load takes place should be entered in this box.

The contents of pre-load sql can be an SQL statement or a procedural block. If using a procedural block, then the final semi-colon is required. The following examples illustrate the possible values in this field. Note the trailing semi-colon on the procedural block example.

Example of pre-load statements:

```
delete from load_customer where code < 23
```

```
delete table load_customer all
```

Table Properties clauses (e.g. Partition)

These are clauses that are added to the end of the table create statement. Typically used for putting partition information on the table in Teradata.

Post Load Procedure

A procedure that is executed immediately following the load. If you execute an externally loaded table, no load occurs, but a post load procedure can still be executed.

Post load procedures can either be manually generated from a RED provided procedure outline or generated leveraging a RED template—refer to *Rebuilding Update Procedures* (on page 199) for details.

The three fields at the bottom of the Load Table Properties screen display date information:

1. Date table structure last updated
2. Date created in database
3. Date last updated in database

DATABASE LINK LOAD - SOURCE SCREEN

The fields on the **Load Table Source** screen for database link loads are described below:

The screenshot shows the 'Load load_customer' dialog box with the 'Source' tab selected. The configuration is as follows:

Load Type	Database link load
Source Connection	Tutorial (OLTP)
General	
Select Distinct Values	<input type="checkbox"/>
Source Schema	WslTutorial
Derive Source Table(s) and Source Columns	<input type="button" value="Derive"/>
Source Table(s)	customer
Override Source Column/Transformations	<input type="checkbox"/>
Where and Group By Clauses	...
Override Load SQL	...

Select Distinct Values
Include the DISTINCT keyword in the SQL SELECT statement.

Buttons: OK, Cancel, Help

Load Type

Method of loading data into the table. The available options are dependent on the **Source Connection**. Defaults to the 'Default Load Type' of the **Source Connection**. Can be specified via the Properties screen.

Source Connection

The connection that identifies the source database. Can be specified via the Properties page.

General**Select Distinct Values**

Include the DISTINCT keyword in the SQL SELECT statement.

Source Schema

Schema within the source database where the source table resides.

Derive Source Tables(s) and Source Columns

Derive the **Source Table(s)** and **Source Column(s)** properties (of this dialog) from the source details of this table's columns.

Note: The existing property values will be overwritten.

Source Table(s)

Name of the table or tables that the data is sourced from.

Override Source Column/Transformations

Ignore the source and transformation details of this table's columns and instead use the override details specified below. See the section on **Load Table Transformations** (on page 290) below for a fuller explanation.

Where and Group By Clauses

Optional SQL SELECT WHERE-clause and/or GROUP BY-clause. Parameter names can be specified using \$Pparameter_name\$ (with leading \$P and trailing \$), which are replaced at run-time by the parameter's value.



TIP: This is where you can build a statement to handle change data.

Parameter values can be in-line replaced and included in the 'Where' clause. Prefix the parameter name with a '\$P' and add a trailing '\$'. For example, if we have a parameter called SALES_LOCATION_CODE we could construct a statement, such as WHERE location_code = '\$PSALES_LOCATION_CODE\$' AND region_code = 'NY'. When this statement is executed, the value of the parameter will replace the parameter name. For example if the parameter was set to 'New York' then the statement would execute as: WHERE location_code = 'New York' AND region_code = 'NY'.

Override Load SQL

Optional SQL statement to load data into the table, which overrides all other properties. The specified SQL will be executed instead of generated SQL. For a linked database specify a complete INSERT statement. For an ODBC source specify only the SELECT statement.

ODBC BASED LOAD

An ODBC based load provides an extensive option for acquiring data from many sources. It is slower than most other forms of load, but may still be a viable option, depending on the volume of data being loaded.

An ODBC based 'interactive load' when using the WhereScape RED tool will use the established ODBC connection to pull the data back to the local PC, and then push the data to the data warehouse database via a sql ODBC statement.

A scheduler load will perform in the same way, however the data is loaded into the server that is running the scheduler and then pushed to the data warehouse database.

The obvious disadvantage in an ODBC based load is the two network transactions required, and the overhead placed on the Scheduler Server.

The properties screens for an ODBC based load are the same as those of a database link load. Refer to *Database Link Load - Properties* and *Database Link Load - Source Screen* (on page 223) for more details.

NATIVE ODBC BASED LOAD

A Native ODBC based load is similar to an ODBC load, except it provides a faster method to move data from another database into Teradata.

A Native ODBC based 'interactive load' when using the WhereScape RED tool will use the established ODBC connection to pull the data back to a delimited file on the local PC, and then push the data to the data warehouse Teradata database via a Fastload session.

A scheduler load will perform in the same way, however the data is loaded into the server that is running the scheduler and then pushed to the data warehouse database.

For fastload loading to work all dates and times must be a character string of the form 'YYYY-MM-DD HH24:MI:SS'. This is normally achieved via a 'During' load transformation, using the correct casting function for the source database.

The Properties and Storage screens for a Native ODBC based load are the same as those of a database link load. Refer to *Database Link Load - Properties* for more details. Details of the Native ODBC Source screen follow.

Notes:

- The Native ODBC load in an ODBC connection is **not** enabled by default for newly created Teradata repositories. Refer to Tools>Options>*Available Load Types* (see "*Settings - Available Load Types*" on page 120) for details.
 - If you are doing **Native Loads using UNIX and LINUX**, see section *Native Loads using UNIX and LINUX* for more details about this type of load.
-

NATIVE ODBC BASED LOAD - SOURCE SCREEN

The fields on the **Source** tab of the properties screen for Native ODBC loads are described below:

The screenshot shows the 'Load load_customer' dialog box with the 'Source' tab selected. The dialog is divided into several sections:

- Properties:** A sidebar on the left with options: Properties, Storage, Override Create DDL, Source (selected), File Actions, Extended Properties, and Notes.
- General:**
 - Load Type: Native ODBC
 - Source Connection: Tutorial (OLTP)
 - Select Distinct Values:
 - Source Schema: WsITutorial
 - Derive Source Table(s) and Source Columns:
 - Source Table(s): customer
 - Override Source Column/Transformations:
 - Where and Group By Clauses:
 - Override Load SQL:
- Native ODBC Load:**
 - Native ODBC Load Routine: FastLoad
 - Field Delimiter: |
 - UTF-8 Extract File:
 - Load Routine Minimum Sessions: 1
 - Set Load Routine Maximum Sessions:
 - Populate Load Parameters:
 - Insert into dss_load_file Table:

At the bottom, there is a section for 'Select Distinct Values' with the instruction: 'Include the DISTINCT keyword in the SQL SELECT statement.' The dialog has 'OK', 'Cancel', and 'Help' buttons at the bottom right.

Load Type

Method of loading data into the table. The available options are dependent on the **Source Connection**. Defaults to the 'Default Load Type' of the **Source Connection**. Can be specified via the Properties screen.

Source Connection

The connection that identifies the source database. Can be specified via the Properties screen.

Select Distinct Values

Include the DISTINCT keyword in the SQL SELECT statement.

Source Schema

Schema within the source database where the source table resides.

Derive Source Tables(s) and Source Columns

Derive the **Source Table(s)** and **Source Column(s)** properties (of this dialog) from the source details of this table's columns.

Note: The existing property values will be overwritten.

Source Table(s)

Name of the table or tables that the data is sourced from.

Override Source Column/Transformations

Ignore the source and transformation details of this table's columns and instead use the override details specified below.

Where and Group By Clauses

Optional SQL SELECT WHERE-clause and/or GROUP BY-clause. Parameter names can be specified using \$Pparameter_name\$ (with leading \$P and trailing \$), which are replaced at run-time by the parameter's value.



TIP: This is where you can build a statement to handle change data.

Override Load SQL

Optional SQL statement to load data into the table, which overrides all other properties. The specified SQL will be executed instead of generated SQL. For a linked database, specify a complete INSERT statement. For an ODBC source, specify only the SELECT statement.

Native ODBC Load Routine

File Loader utility/mechanism to use to load the generated extract file.

Field Delimiter

Character that separates the fields within each record of the generated extract file. The default value is a | character (pipe). This should be changed if pipes are possible in the source data.

UTF-8 Extract File

Data will be loaded via a UTF-8 format file. The default is unselected.

Populate Load Parameters

Populate any load-related WhereScape RED parameters, which may be used for validation purposes.

Insert into dss_load_file Table

Populate the metadata table dss_load_file in any generated post_load procedure.

FILE ACTIONS

A file action defines a ftp or copy step that happens to files before or after they are loaded.

Creating a File Action

Before creating a ftp file action, ensure a connection exists to the remote server where the files will be transferred.

To create a file action:

- 1 Click on the **File Actions** tab of the load table properties dialog.
- 2 Click the **Add New Action** button.
- 3 Choose the **Action Type** from the drop-down list.
- 4 Enter the **Action Description**.
- 5 Choose the **Action Program** from the drop-down list.
- 6 For a ftp action:
 - Choose the **Action Connection** for the remote server where the file will be transferred.
 - Enter the ftp commands in **ftp command 1** thru **ftp command 9**.
- 7 For a copy action:
 - To move the data file, enter the copy to location into **File Directory**. If the file is to be renamed at the same time, enter a new name into **File Name**, otherwise enter %FILE_NAME%.
 - To move the trigger file, enter the copy to location into **Trigger Directory**. If the trigger file is to be renamed at the same time, enter a new name into **Trigger Name**, otherwise enter %TRIG_NAME%.
- 8 Click on **Save (Update) Action**.
- 9 Repeat if necessary for additional file actions.
- 10 Click **OK**.

Sample ftp file action:

Load load_currency
✕

- Properties
- Storage
- Override Create DDL
- Source
- File Actions
- Extended Properties
- Notes

A number of actions can be defined to take place either before or after the load has occurred. Select an action type to activate the applicable fields.

Action Number: Add New Action Remove Action Save (Update) Action

Action Type:

Action Description:

Action Program:

Action Connection:

ftp command 1:

ftp command 2:

ftp command 3:

ftp command 4:

ftp command 5:

ftp command 6:

ftp command 7:

ftp command 8:

ftp command 9:

Actions

01 FTP Files to down stream system

OK Cancel Help

Sample **copy** file action:

The screenshot shows a dialog box titled "Load load_currency" with a close button (X) in the top right corner. On the left is a vertical navigation pane with the following items: Properties, Storage, Override Create DDL, Source, File Actions (highlighted in blue), Extended Properties, and Notes. The main area contains a text box at the top: "A number of actions can be defined to take place either before or after the load has occurred. Select an action type to activate the applicable fields." Below this are several fields: "Action Number:" with a dropdown menu showing "02" and buttons "Add New Action", "Remove Action", and "Save (Update) Action"; "Action Type:" with a dropdown menu showing "After load"; "Action Description:" with a text box containing "Copy files to another directory after loading"; "Action Program:" with a dropdown menu showing "copy"; "Action Connection:" with a dropdown menu; "File Directory:" with a text box containing "f:\loaded\prod"; "File Name:" with a text box containing "%FILE_NAME%"; "Trigger Directory:" with an empty text box; and "Trigger Name:" with an empty text box. Below these fields is an "Actions" section with a list box containing two entries: "01 FTP Files to down stream system" and "02 Copy files to another directory after loading". At the bottom right of the dialog are three buttons: "OK", "Cancel", and "Help".

NATIVE LOADS USING UNIX AND LINUX

Sometimes the Teradata database server has a LINUX non-tpa node or a MP-RAS UNIX non-tpa node used for loading data from files. This can provide significantly better platform for loading files into Teradata than a network connected Windows Scheduler server.

Native ODBC loads need to extract data from the source system using ODBC on a Windows machine. An additional feature in WhereScape RED can be enabled to transfer the extract file to LINUX or MP-RAS and then load the file on the remote non-tpa node.

The following steps enables this functionality when Native ODBC loads are scheduled:

- 1 Create a unix connection for the non-tpa node (see **Connections** (on page 145) for more information on creating connections).

The "Unix user id" and "Unix password" are used to sign in to the remote machine specified by "Host name".

The work directory is the default location extract files will be transferred to.

The following shows the minimum fields that need to be completed:

The screenshot shows the 'Connection TD_14_00 UNIX' dialog box with the following fields and values:

Field	Value
Connection Name	TD_14_00 UNIX
Connection Type	UNIX/Linux
UNIX/Linux Host Name	TD_14_00
Script Shell	bin/ksh
Loader Host Identification	
Loading Port Start	8001
Loading Port End	8001
Work Directory	home/tdatuser/workdir
Database ID	TD_14_00
Database Server/Home Directory	
Connection Protocol	SSH
Secure Shell (SSH) Command	C:\Program Files (x86)\PuTTY\plink.exe -ssh \$HOST -l \$USER
Pre-Login Action	
Login Prompt	ogin as:
Password Prompt	ssword:
Post-Login Action	
Command Prompt	>
UNIX/Linux User ID	tdatuser
UNIX/Linux User Password	*****
DSS User ID	
DSS User Password	
Teradata Wallet User ID	dssdemo
Teradata Wallet String	your TD Wallet String

- 2 Create a new file action on each load table (see **File Actions** (on page 228) for more information on file actions).

The file action should be look something like this:

The screenshot shows a dialog box titled "Load load_customer" with a close button (X) in the top right corner. On the left is a sidebar with a tree view containing: Properties, Storage, Override Create DDL, Source, File Actions (highlighted), Extended Properties, and Notes. The main area contains the following fields and controls:

- A note: "A number of actions can be defined to take place either before or after the load has occurred. Select an action type to activate the applicable fields."
- Action Number: 02 (dropdown), with buttons "Add New Action", "Remove Action", and "Save (Update) Action".
- Action Type: After load (dropdown).
- Action Description: Move file from MPRAS to load and initiate load (text box).
- Action Program: unixload (dropdown).
- Action Connection: (dropdown).
- ftp command 1: lcd \$WORKDIR\$
- ftp command 2: cd \$UNIXDIR\$
- ftp command 3: as
- ftp command 4: prompt
- ftp command 5: mput \$FILE\$
- ftp command 6: (empty)
- ftp command 7: (empty)
- ftp command 8: (empty)
- ftp command 9: (empty)
- Actions list:
 - 01 Delete all temporary files
 - 02 Move file from MPRAS to load and initiate load

At the bottom right are buttons for "OK", "Cancel", and "Help".

Note: The variables \$WORKDIR\$ and \$UNIXDIR\$ refer to the work directories of the Native ODBC connection and Unix connection respectively. The variable \$FILE\$ represents the extract file being transferred to LINUX or MP-RAS.

- 3 Ensure ws_sched_tera_550.sh script has been set up on the MP-RAS or LINUX non-tpa node and is scheduled to run on the machine every 5 or 10 minutes using cron or any other scheduling software on the non-tpa node (see the appropriate section in the *WhereScape RED Installation and Administration Guide* for more information on ws_sched_tera_550.sh).

TPT LOAD

A TPT ODBC based load is similar to an ODBC load, except it provides a faster method to move data from another physical database server into your Teradata database. This includes moving data from one Teradata server to another and moving data from any other relational database into Teradata.

NOTE: The **TPT Load** type can only be used with a **Windows** scheduler. If using a **UNIX/Linux** scheduler, refer to section 9.7 which also includes details for using the **TPT Script Load** type for UNIX/Linux.

WhereScape RED runs a TPT ODBC load, by generating a TPT script based on your metadata. The TPT script includes the following steps:

- 1 Extract from the source database, using the TPT ODBC operator.
- 2 Import into the Teradata database, using either the TPT LOAD or TPT UPDATE operator.

Cross database platform TPT ODBC loads present some challenges due to incorrect assumptions made by TPT from time to time. You may need to add 'During' load transformations to your Load tables and/or modify numeric data types from their apparent equivalent data type for a load to complete successfully. This is because TPT assumes a data type with the same name in another database is the same size in Teradata; of course this is not always the case.

The Properties and Storage screens for a TPT ODBC based load are the same as those of a **Database Link Load** (on page 221).

NOTE: TPT Loads from ODBC connections in a Teradata repository must have ODBC connections set up, using the appropriate DataDirect ODBC driver.

TPT Loads from ODBC Connections will not work if a non-DataDirect driver is used.

Refer to the previous section for details. Details of the Source mapping screen is provided in the next section.

Load load_customer

Properties

Storage

Override Create DDL

Source

Extended Properties

Notes

Load Table Name: load_customer

Unique Short Name: (maximum 22 characters) load_customer

Description: Customer load table

Connection: Tutorial (OLTP)

Load Type: TPT load

Database Link:

Script Template: (None)

Script Name: (None)

Pre-Load Action: Truncate

Pre-Load SQL:

Post Load Procedure: (None) Rebuild Regenerate

Timestamps

Metadata Structure	Database	Database
2007-11-09 22:08:29.000000	2017-10-19 02:09:06.870000	2017-10-19 02:09:06.870000

OK Cancel Help

TPT LOAD - SOURCE SCREEN

The fields on the **Source** tab of the Properties screen for TPT loads are described below:

The screenshot shows the 'Load load_customer' Properties dialog box with the 'Source' tab selected. The dialog is divided into several sections:

- Properties:** A sidebar on the left with options: Properties, Storage, Override Create DDL, **Source** (selected), Extended Properties, and Notes.
- General:**
 - Load Type: TPT load
 - Source Connection: Tutorial (OLTP)
 - Select Distinct Values:
 - Source Schema: WslTutorial
 - Derive Source Table(s) and Source Columns: Derive
 - Source Table(s): customer
 - Override Source Column/Transformations:
 - Allow Missing Source Columns:
 - Fail when Incomplete Load:
 - Where and Group By Clauses: ...
 - Override Load SQL: ...
- Teradata Parallel Transporter (TPT):**
 - TPT Load Type: Load TPT
 - TPT Character Set: ASCII
 - TPT ODBC Operator Attributes: ...
 - TPT Load/Update Operator Attributes: ...
 - TPT Job Name: ...
 - TPT Build Command Options: ...
 - TPT Shared Memory Size: (Default)
 - TPT Minimum Sessions: 1
 - Set TPT Maximum Sessions:
 - Load Read Instances: 1
 - Load Write Instances: 1
- Select Distinct Values:**

Include the DISTINCT keyword in the SQL SELECT statement.

At the bottom right, there are buttons for OK, Cancel, and Help.

Load Type

Method of loading data into the table. The available options are dependent on the **Source Connection**. Defaults to the 'Default Load Type' of the **Source Connection**. Can be specified via the Properties screen.

Source Connection

The connection that identifies the source database. This can be specified via the Properties screen.

Select Distinct Values

Include the DISTINCT keyword in the SQL SELECT statement.

Source Schema

Schema within the source database where the source table resides.

Derive Source Tables(s) and Source Columns

Derive the **Source Table(s)** and **Source Column(s)** properties (of this dialog) from the source details of this table's columns.

Note: The existing property values will be overwritten.

Source Table(s)

Name of the table or tables that the data is sourced from.

Override Source Column/Transformations

This ignores the source and transformation details of this table's columns and uses the override details specified below instead.

Allow Missing Source Columns

Allow the load to occur when one or more of the source columns do not exist. (see the section on *Handling missing source columns* (on page 288)).

Fail when incomplete Load

Controls whether the load reports failure when all the rows extracted are not loaded. The specified exit status impacts any remaining tasks in the currently running job. When fail is specified, the WhereScape RED Scheduler will stop/fail the job and hold any remaining tasks when the load fails. In contrast, when fail is disabled the scheduler will continue to run any dependent tasks in the job that is running.

Where and Group By Clauses

Optional SQL SELECT WHERE-clause and/or GROUP BY-clause. Parameter names can be specified using \$Pparameter_name\$ (with leading \$P and trailing \$), which are replaced at run-time by the parameter's value.



WhereScape RED TIP: This is where you can build a statement to handle change data.

Override Load SQL

Optional SQL statement to load data into the table, which overrides all other properties. The specified SQL will be executed instead of generated SQL. For a linked database, specify a complete INSERT statement. For an ODBC source, specify only the SELECT statement.

TPT Load Type

Load TPT, Update TPT or Stream TPT.

NOTE: When importing a Model from 3D to RED, select **Load TPT** instead of Fastload as the Default File Loader method. FastLoad is not a valid option for loading Linux files to Teradata. To change the Default File Loader, go to **Tools->Options->Code Generation->General**.

TPT Character Set

Teradata-compliant Character Set Name to use when loading, such as ASCII, UTF8, UTF16.

TPT ODBC Operator Attributes

Optional comma-delimited list of TPT ODBC operator Attributes, e.g. INTEGER DataBlockSize = 2048.

TPT Load/Update Operator Attributes

Optional comma-delimited list of TPT Load Operator or TPT Update Operator Attributes.

TPT Job Name

Job Name for TPT. If not set, will default to Load Name.

TPT Build Command Options

Additional options included as part of the TBuild call.

TPT Shared Memory Size

Shared memory size can be specified in bytes, kilobytes or megabytes. Examples include 2091752,2048K,2M.

TPT Minimum Sessions

Optional minimum number of sessions[1-99]. The default is one session.

Set TPT Maximum Sessions

Enables the 'TPT Maximum Sessions' property. Optional maximum number of sessions[1-99]. The default is one session per available Access Module Processor (AMP). The maximum value cannot exceed the number of available AMP's.

Load Read Instances

Optional number of TPT read instances[1-99]. The default is one instance.

Load Write Instances

Optional number of TPT write instances[1-99]. The default is one instance.

TPT LogView Command Options

Additional options included as part of the TLOGVIEW call. -f** can be added to this field to get enhanced logging to diagnose issues in the event of a failure.

CLEANUP AFTER TPT LOAD FAILURE

Note: If the cause of the failed load persists (for example, an outage of the source database) then the **Reset Failed TPT Load** function from the Load table's context menu can be used. This option uses TPT functionality to clear the checkpoint file and release the database lock on the Load table. See *Working with Objects > Load Tables* (see "*Working with Objects*" on page 15) for details.

TPT runs all jobs in checkpoint mode by default; and if any one of these jobs fail, it can restart based on the last checkpoint taken for the job.

Automatic Restart

An automatic restart means that a job can restart on its own, without you manually having to resubmit the job. With the Start-of-Data and End-of-Data checkpoints, a job can automatically restart itself when a "retry-able" error occurs (such as a database restart or deadlock) before, during, or after the loading of data. You need to consider the following when dealing with automatic restarts:

Jobs can automatically restart as many times as specified by the value of the RETRY option of the TPT job-launching command. By default, a job can restart up to five times.

If no checkpoint interval is specified for a job, and the job fails during processing, the job restarts either at the Start-of-Data checkpoint or at the End-of-Data checkpoint, depending on which one is the last recorded checkpoint in the checkpoint file.

To avoid reloading data from the beginning, you should specify a checkpoint interval when launching a job so that the restart can be done based on the most recent checkpoint taken.

Manual Restart

If a job fails and terminates, a manual restart can be accomplished by resubmitting the same job with the original job-launching command. By default, all TPT jobs are checkpoint restartable, using one of the two checkpoints at Start-of-Data and End-of-Data.

TPT also provides recovery across job steps within a job, thus if a job has multiple steps, a checkpoint will be created for each successful step. This will allow a job to restart from the failed step by skipping the successful steps. If, for example, you have a step to create or drop tables before a data loading step begins, and the job fails in the data loading step; a restart of the job would resume from the data loading step without recreating or dropping the tables. This can be contrasted with some of the utilities, such as Fastload, where the script might contain the statements DROP TABLE and CREATE TABLE. Such a script could not be used across restarts because those DDL statements would be re-issued.

Removing Checkpoint Files

Job checkpoint files are automatically created by TPT and they are deleted if the job completes without an error. You will however need to remove these checkpoint files before they are automatically deleted if you wish to:

- Rerun an interrupted job from the beginning, rather than restarting it from the last checkpoint taken before the interruption occurred.
- Abandon an interrupted job and run another job, but the new checkpoint files will have the same names as the existing checkpoint files, due to the use of the same job name (or the default checkpoint files created based on the logon user ID).

TPT provides a special command for users to remove checkpoint files, based on either the user ID or the job name.

If the "tbuild" command specifies a job name, the "twbrmcp <job name>" command can be used. If the "tbuild" command does not specify a job name, the "twbrmcp <user ID>" can be used. For z/OS, the deletion of checkpoint files can be done through the MVS/ISPF facility.

If you want to delete checkpoint files manually, you can use one of the following commands:

- del %TWB_ROOT%\checkpoint\<<job-name>.*
- del %TWB_ROOT%\checkpoint\<<user-id>.*

If you want to delete checkpoint files from a user-defined directory, you can use one of the following commands:

- del <user-defined directory>\<job-name>.*
- del <user-defined directory>\<user-id>.*

Note: If a manual restart is necessary and the checkpoint files have been removed, the simplest method to clear the load lock on a table is to recreate the table.

TPT UNIX/LINUX SCRIPT LOAD

Note: If a TPT script load is attempted where RED detects that the previous load failed, then the pre-load truncate/delete action is not performed and a TPT restart action is performed.

Create a **Unix** connection as in the example below:

General	
Connection Name	UNIX
Connection Type	UNIX/Linux
UNIX/Linux Host	
UNIX/Linux Host Name	TD15_doc6871
Script Shell	/bin/ksh
Loader Host Identification	
Loading Port Start	8001
Loading Port End	8001
Work Directory	c:\temp
Database ID	TD15_doc6871
Database Server/Home Directory	
Connection Protocol	
Secure Shell (SSH) Command	/home/red/work/dir
Pre-Login Action	
Login Prompt	ogin as:
Password Prompt	ssword:
Post-Login Action	
Command Prompt	>
Credentials	
UNIX/Linux User ID	tdatuser
UNIX/Linux User Password	*****
DSS User ID	doc6871
DSS User Password	
Teradata Wallet User ID	
Teradata Wallet String	
Other	
Command Prompt	UNIX/Linux command prompt or the tail-end of that prompt. [Typically \$].

TPT UNIX SCRIPT LOAD - PROPERTIES

The fields on the TPT UNIX Script Load Properties screen are the following:



TIP: When doing **TPT Script based loads**, it is easy to use the **Regenerate** button to the right of the Script Name field to regenerate the scripts.

Load Table Name

The table name is limited by Teradata to a maximum of 30 characters and must be unique. Table name defaults can be set up (**Tools>Options** then Naming local or global) to define a prefix or a post fix that can be added in order to identify clearly that this is a load table. Example: load_customer. By default RED uses the prefix load_ for load tables.

Unique Short Name

The table short name is limited in size to 22 characters and must be unique. The short name is used in the naming of indexes, keys and procedures.

Description

Enter here a description of the table. This description appears in the documentation that can be generated once the data warehouse is built.

Connection

Enter the connection being used to get the data. The connections for load tables can be changed in bulk. See section on *Changing load Connection and Schema* (on page 291).

Load Type

The load type is typically defined by the connection, and should not normally be changed. This drop-down lists all valid load types for the connection.

ODBC Script Connection

Choose the **Unix** Connection.

Script Template

The script template used for a script based load.

Script Name

This field is only active if this is a script based load.



TIP: Use the **Rebuild** button after selecting the relevant script to be rebuilt on the the Script Name drop-down menu.

Pre-load Action

Select an action to be performed on the table before the load occurs. Options are:

- Truncate
- Execute pre-load Sql
- Both Truncate and Execute pre-load Sql
- No action

Pre-load Sql

If a Pre-load Action of "**Execute pre-load Sql**" was selected, then the Sql statement to execute before the load takes place should be entered in this box.

The contents of pre-load sql can be a sql statement or a procedural block. If using a procedural block, then the final semi-colon is required. The following examples illustrate the possible values in this field. Note the trailing semi-colon on the procedural block example.

Example of pre-load statements:

```
delete from load_customer where code < 23
```

```
delete table load_customer all
```

Post Load Procedure

A procedure that is executed immediately following the load. If you execute an externally loaded table, no load occurs, but a post load procedure can still be executed.

Post load procedures can either be manually generated from a RED provided procedure outline or generated leveraging a RED template—refer to *Rebuilding Update Procedures* (on page 199) for details.

The three fields at the bottom of the Load Table Properties screen display date information:

1. Date table structure was last updated
2. Date created in database
3. Date last updated in database

SSIS LOADER

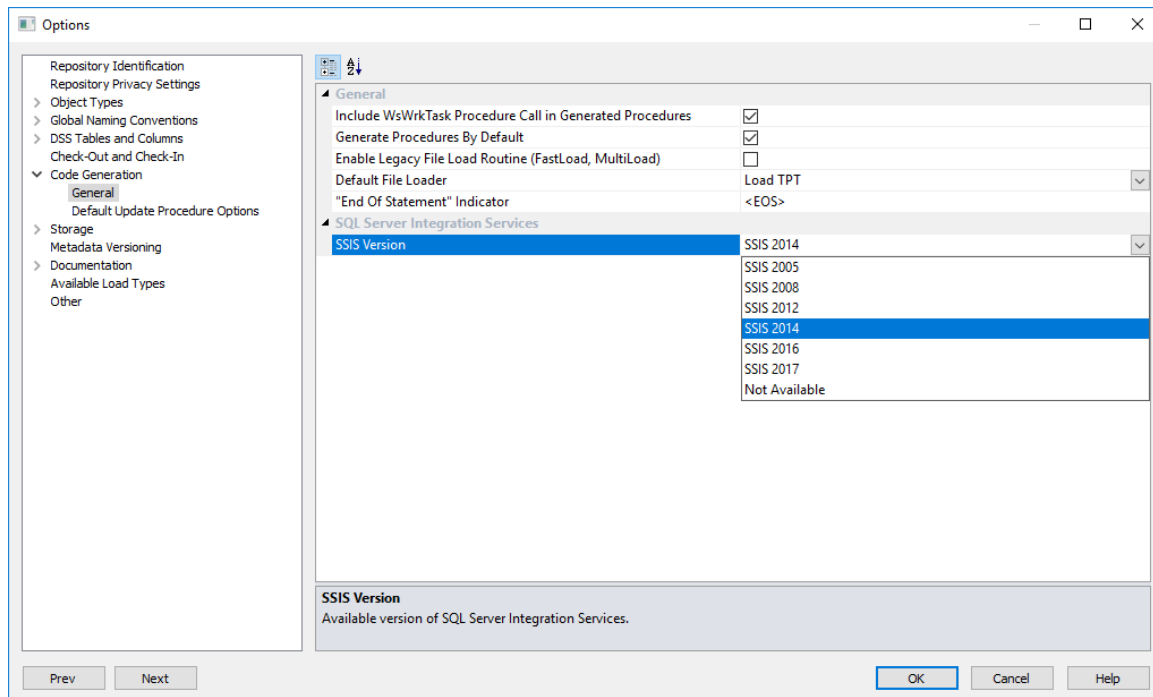
Microsoft SQL Server Integration Services (SSIS) is an Extract Transform and Load (ETL) utility that is supplied and licensed as part of Microsoft SQL Server 2005+. SSIS is built to extract data from data sources, using extraction technologies such as OLEDB, transform the data using its own .NET based language and then load it into data destination. SSIS is primarily used for loading data from various sources into SQL Server, but it can be used to extract from most databases and load into any other database.

Note: The SSIS file load in Database, ODBC and Windows connections is **not** enabled by default for newly created Teradata repositories.

Refer to Tools>Options>**Available Load Types** (see "**Settings - Available Load Types**" on page 120) for details.

WhereScape RED manages the extraction and load of data into the data warehouse as part of data warehouse processing. WhereScape RED attempts to utilize the optimum loading technology for each supported database platform. Optimal loading performance from a database source into Teradata can be achieved, using SSIS. WhereScape RED provides the ability to leverage the extract and load performance available from SSIS as part of the WhereScape RED processing. WhereScape RED does this by generating and executing an SSIS package dynamically at run time. Any errors or messages resulting from execution are passed back into the WhereScape RED workflow metadata to drive subsequent workflow actions. In the event of an error during execution of the SSIS package, the package will be persisted to disk to assist the developer with problem resolution.

To use SSIS to load data, select the relevant version of SSIS in **Tools>Options>Code Generation>General**.



LOADING DATA INTO RED LOAD TABLES USING SSIS

To use SSIS loading, ensure that SSIS loads are enabled and that the **SSIS version** is set correctly in **Tools>Options>Code Generation>General**.

If you are loading from a Flat File source, see *Loading Data from Flat Files using SSIS*.

RED can extract and load data using SSIS from database tables or flat files from a Windows connection. As with any load into RED a connection to the source data needs to be created to provide extraction details.

The **SSIS Connection String** is a valid SSIS connection string that can be used to connect to the data source or destination.

Loading Data via SSIS from a database

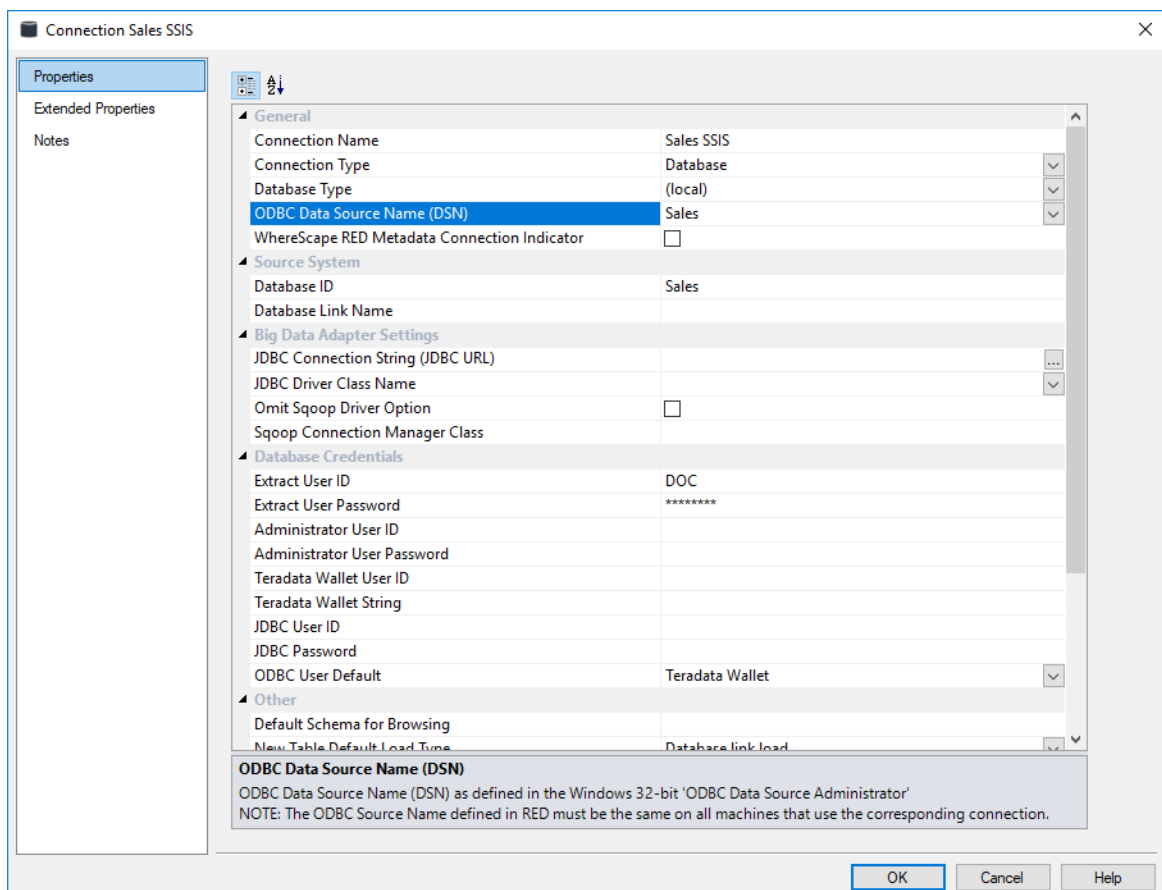
If the connection is a database load then there is additional connection information that should be supplied to use SSIS as a loading option.

This additional information needs to be supplied on both the source connection and the data warehouse connection.

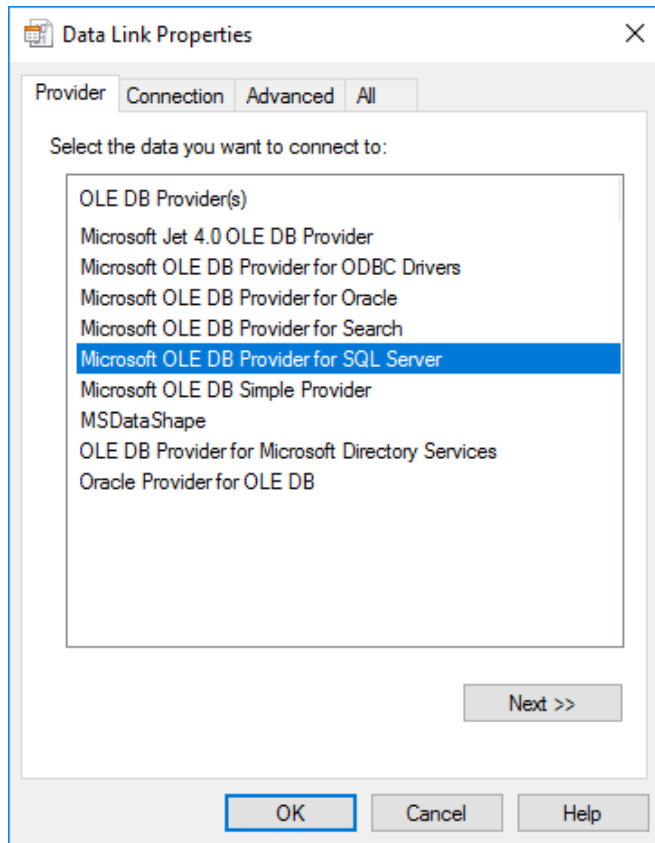
NOTE: SSIS Loads in Teradata can only be processed with a Windows Scheduler.

Creating the SSIS Connection String

- 1 Click on the **ellipsis** button to the right of the SSIS connection string field of the relevant connection:

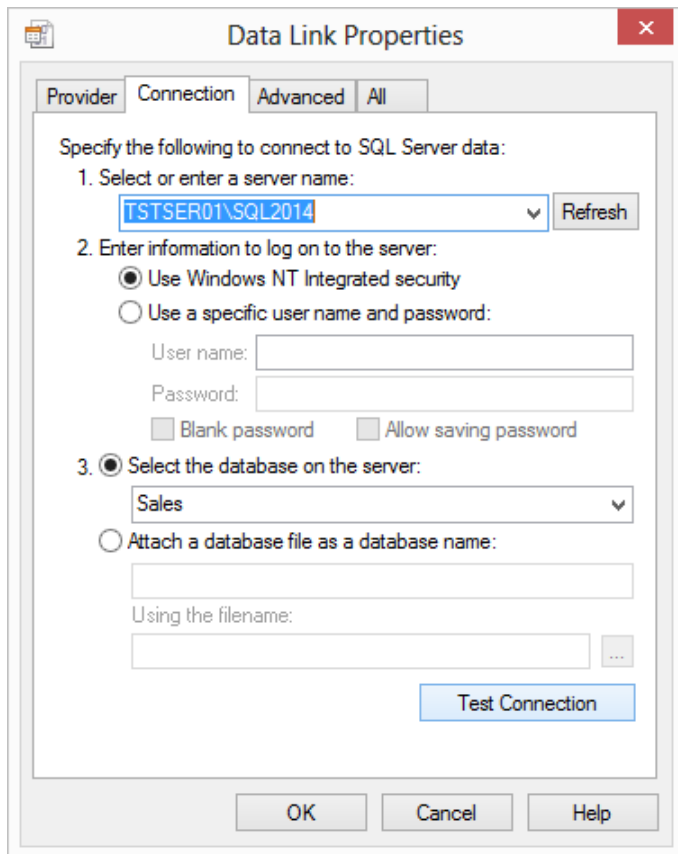


- 2 On the Provider tab, select the relevant **OLE DB Provider** and click **Next**.




- 3 On the **Connection** tab, select the **server name**, enter the information to log on to the server and select the **database** on the server. Click **Test Connection**.

Notes: When using a specific **user name** and **password** to connect to the server instead of using Windows integrated security, the **Allow saving password** check-box must be selected. It is also recommended that the password on the SSIS connection string field that is displayed in the connection properties is replaced with the **\$PASSWORD\$** token that is substituted at runtime.



- 4 Click **OK**.



- 5 Click **OK** on the Data Link Properties dialog to save the connection string settings.

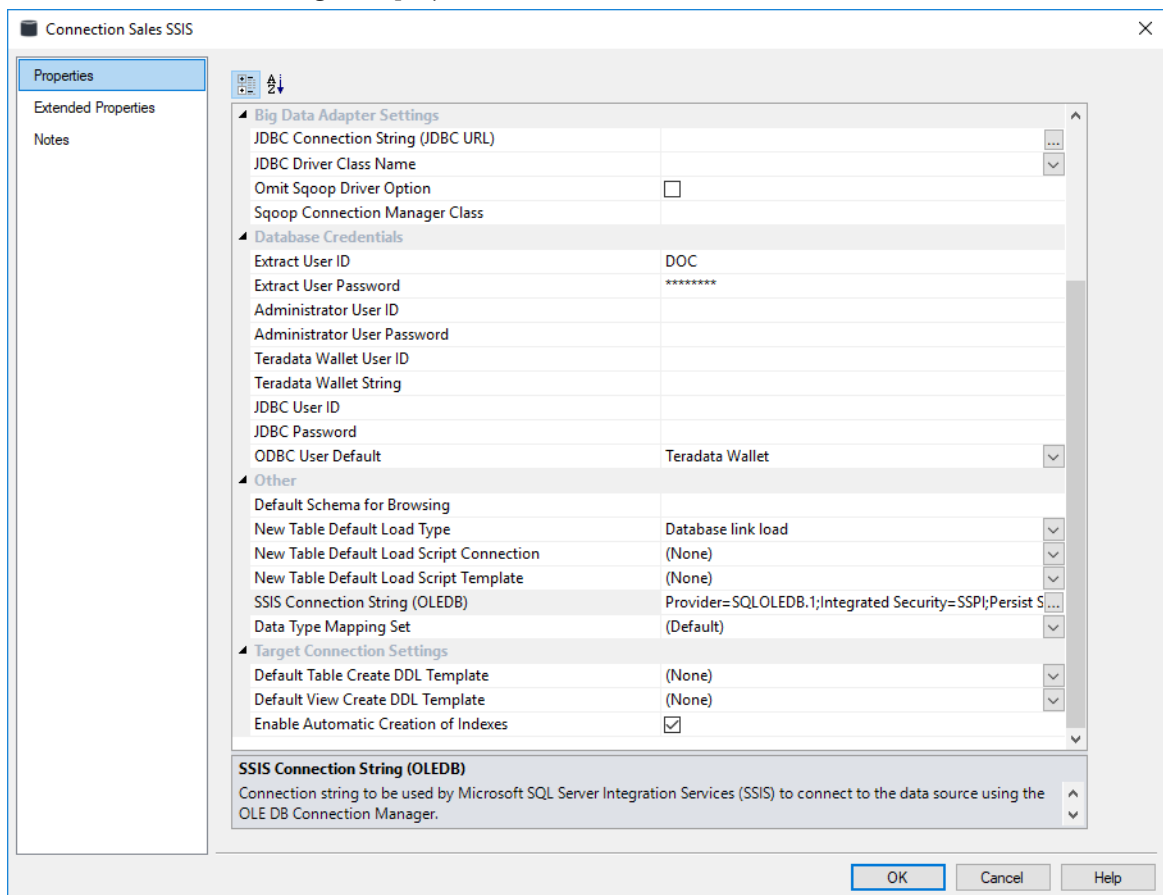
The screenshot shows the 'Data Link Properties' dialog box with the 'Connection' tab selected. The dialog is titled 'Data Link Properties' and has a close button (X) in the top right corner. Below the title bar are four tabs: 'Provider', 'Connection', 'Advanced', and 'All'. The 'Connection' tab is active, displaying the following configuration options:

Specify the following to connect to SQL Server data:

1. Select or enter a server name:
A dropdown menu contains 'TSTSER01\SQLE2014' and a 'Refresh' button is to its right.
2. Enter information to log on to the server:
 Use Windows NT Integrated security
 Use a specific user name and password:
User name: [text box]
Password: [text box]
 Blank password Allow saving password
3. Select the database on the server:
A dropdown menu contains 'Sales'.
 Attach a database file as a database name:
Sales [text box]
Using the filename: [text box] ...

At the bottom of the dialog are three buttons: 'OK', 'Cancel', and 'Help'. A 'Test Connection' button is located in the lower right area of the main configuration area.

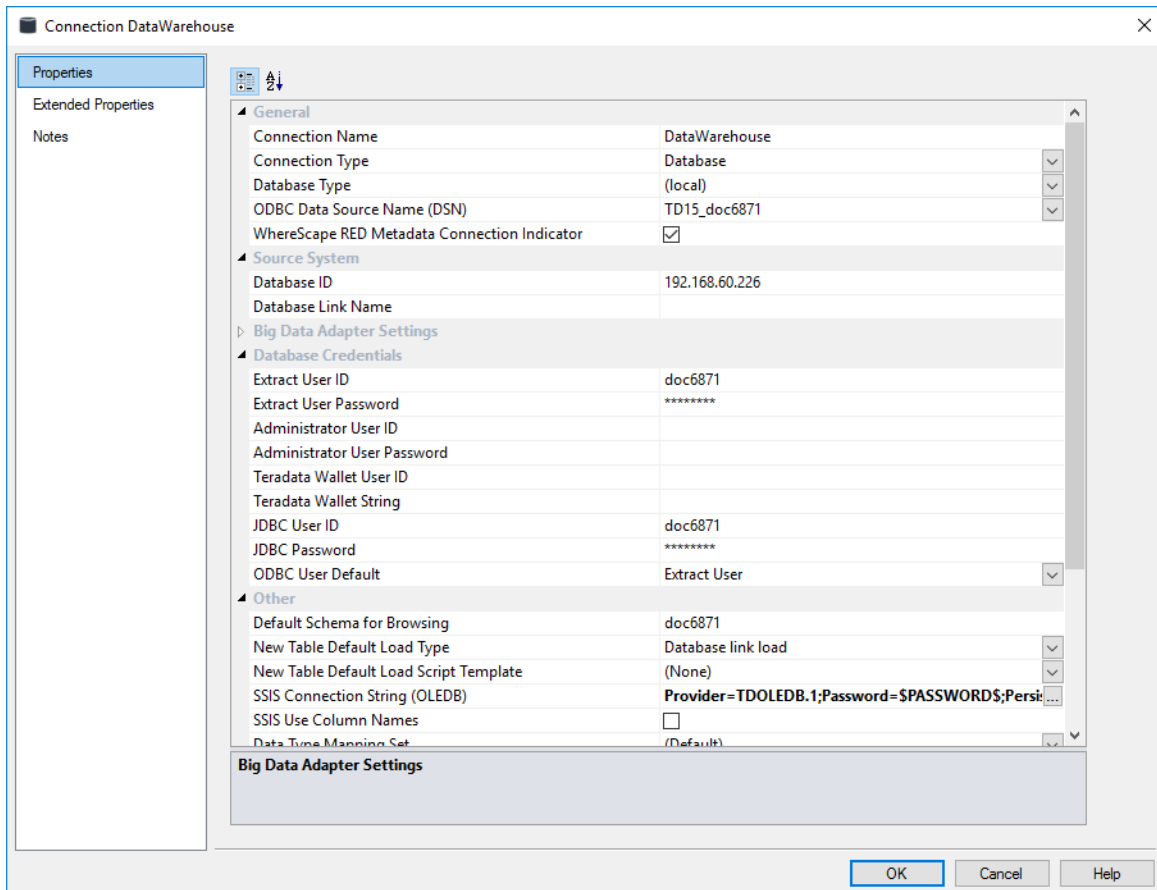
6 The SSIS connection string is displayed.



7 Click OK to save the connection.

- Right-click on **Sales SSIS** and select **Browse Connection**.
- Accept the defaults and click **OK**.

- 8 In SSIS terms, you have now defined your **Source** in **SSIS Connection Manager**. Using the same process, you need to add the SSIS Connection String to the data warehouse connection so you can specify your **Destination** connection:
- Double-click on the **DataWarehouse** connection in the object explorer to open up the Properties dialog.
 - Follow the process above to create the SSIS Connection String, this time selecting the **OLE DB Provider for Teradata**.
 - Click **OK** to save your connection.




Note1: If the connection string is already set, then the ellipsis button will open up an editor dialog.

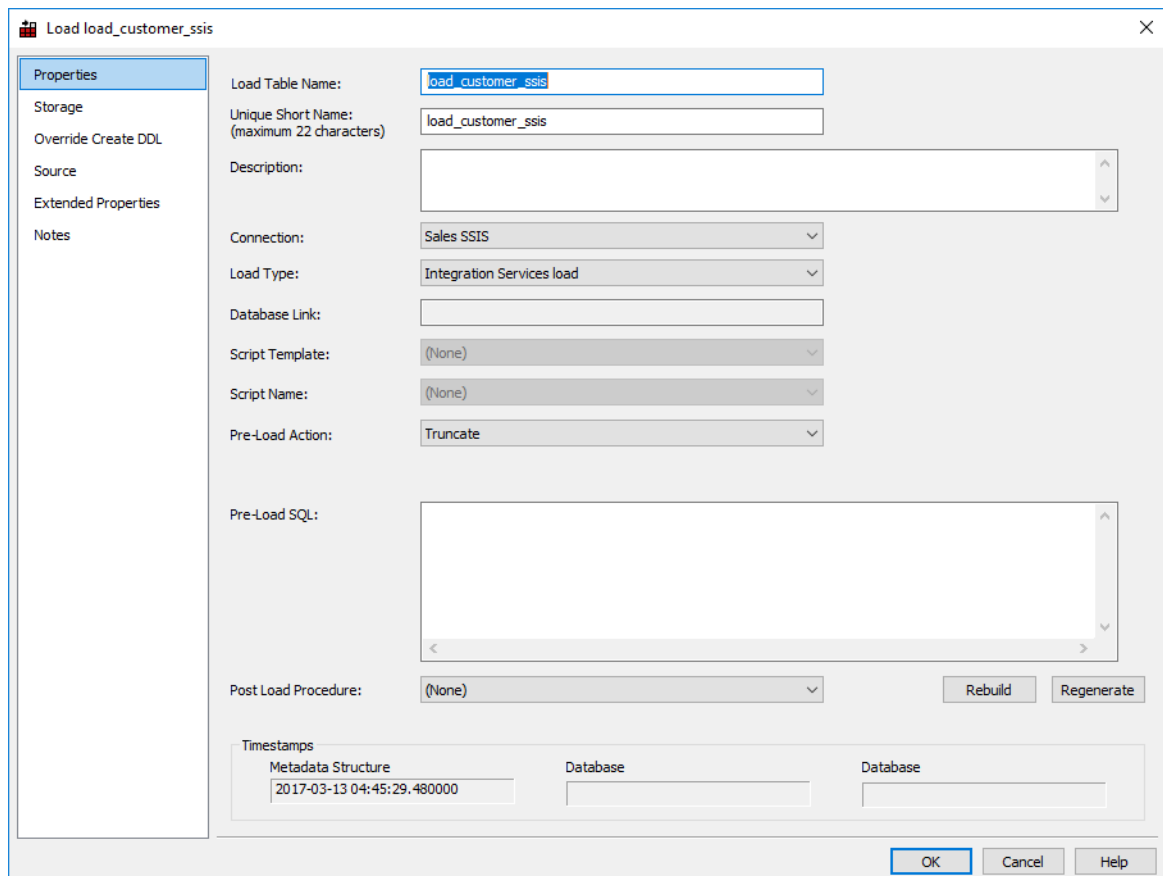
Edit the connection string and click **OK**.

Note2: For connections that require a username and password, the connection string can also be edited to replace the password with the **\$PASSWORD\$** token that is substituted at runtime. If the **\$PASSWORD\$** token is used, RED uses the contents of the masked "Extract User Password" property when making the connection.

E.g. "Provider=TDOLDB.1;Password=**\$PASSWORD\$**;"

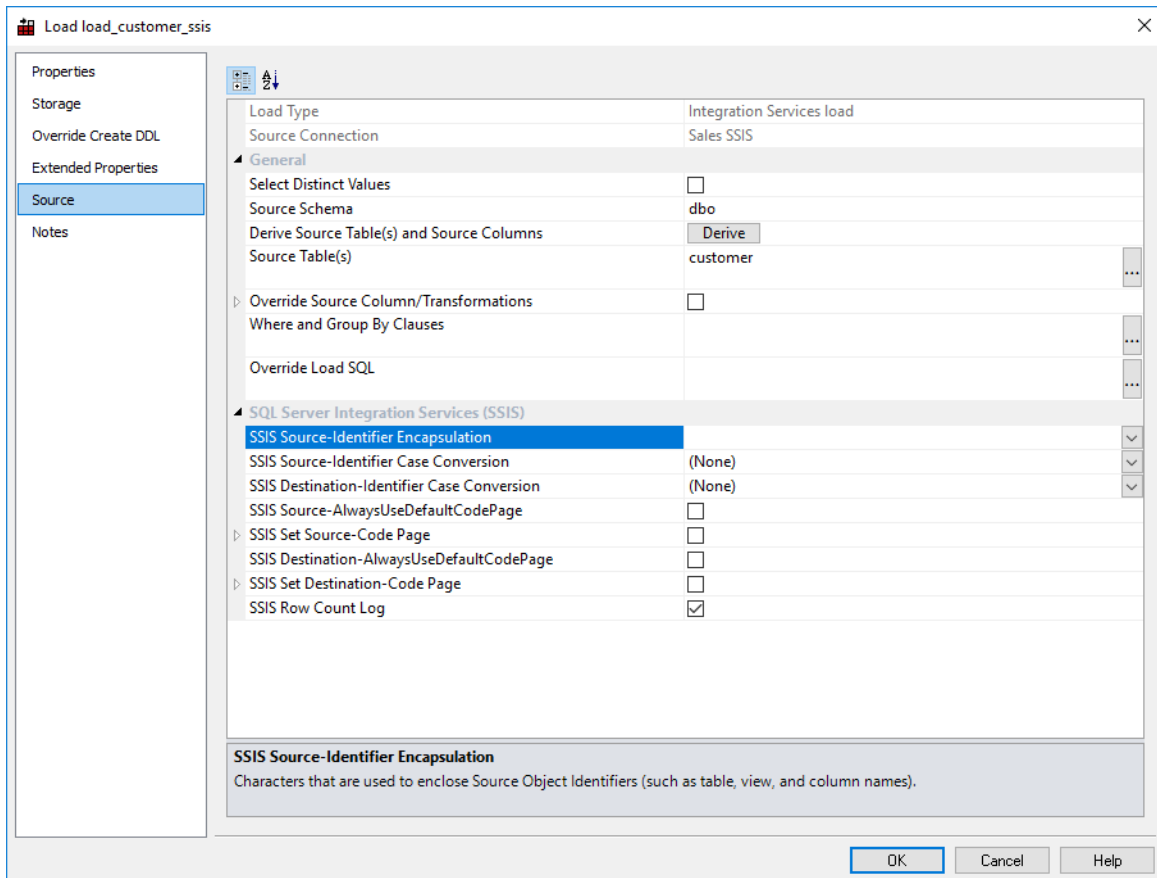


- Once the connection is defined then a load table needs to be created to hold data loaded into the data warehouse by dragging a source table or a flat file to create a load table – (see **Loading Data** (on page 215) or **Loading Data from Flat Files using SSIS**). On the load table properties the Load type can be set to **Integration Services load**. This will create and execute a SSIS package at run time to load data into the data warehouse load table.



10 The configuration options available on an SSIS load are available on the **Source** tab of the load table's Properties. These options are:

- **SSIS Source-Identifier Encapsulation** - Characters that are used to enclose source column names. Options are (None), "", [], ", ``
- **SSIS Source-Identifier Case Conversion** - Case-sensitivity conversion applied to Source Object Identifiers (such as table, view, and column names) in RED-generated SSIS packages. If no conversion is applied then the exact case of the identifier defined in the RED metadata is used in SSIS.
- **SSIS Destination-Identifier Case Conversion** - Case-sensitivity conversion applied to Destination Object Identifiers (such as table, view, and column names) in RED-generated SSIS packages. If no conversion is applied then the exact case of the identifier defined in the RED metadata is used in SSIS.
- **SSIS Destination-AlwaysUseDefaultCodePage** - Forces the use of the DefaultCodePage property value when describing character data.
- **SSIS Set Destination-Code Page** - Enables the SSIS destination code page property.
- **SSIS Row Count Log** - During an SSIS Load include Row Count logging.



FLAT FILE LOADS

Flat file loads can be processed from either a **UNIX/Linux, Windows** or **Hadoop** connection. As with all other load types, it is easier to use the **drag and drop** functionality to create load tables. Flat files can also be loaded using SQL Server Integration Services (SSIS). For Flat File load instructions using SSIS, see the next section—*Loading Data from Flat Files using SSIS*.

The **drag and drop** process for flat files is as follows:

- 1 Browse to the directory and file via the appropriate **UNIX/Linux Connection, Windows Connection** (see "*Windows*" on page 163) or **Hadoop Connection**.
- 2 Double click the **Load Table** object in the left pane to create a drop target.
- 3 **Drag** the file from the right pane and **drop** it into the middle pane.
- 4 The following dialog appears. Rename the file if necessary and click the **ADD** button.

NOTE: The option **Add meta data columns to table** is used for creating load tables that are used in creating Data Vault objects. If this option is selected, two DSS columns (**dss_record_source** and **dss_load_date**) are included in the meta data for the table and are populated by transformations. These two DSS columns could equally be applied to other load tables not used in a Data Vault system but are particularly important to comply with the Data Vault standards. Please refer to the *Data Vaults chapter* (see "*Data Vaults*" on page 413) for more details.

- 5 The following screen displays for file loads from **Windows, UNIX/Linux** and **Hadoop** connections.

Data load Wizard

Load Type:

File parsing:

File Parsing

First Rows from the File

```
product_code,customer_code,budget_quantity,budget_sales_value,budget_date
1002,228,185,409.92,02-JUN-2010
1008,228,80,978.58,02-JUN-2010
1003,227,62,572.42,30-APR-2011
1007,227,98,766.17,30-APR-2011
1004,226,40,218.00,05-NOV-2011
1006,226,40,618.00,05-NOV-2011
1009,225,74,940.24,04-APR-2012
1002,225,74,163.97,04-APR-2012
1006,225,40,618.00,04-APR-2012
1007,225,98,766.17,04-APR-2012
1004,225,74,402.54,04-APR-2012
1003,224,15,134.85,15-NOV-2011
1008,224,15,177.34,15-NOV-2011
1001,224,15,159.50,15-NOV-2011
1001,223,74,812.46,13-AUG-2010
1007,223,80,633.53,13-AUG-2010
```

Column Delimiter: No Column delimiter will initiate width based parsing
CHAR(nn) inserts an ASCII character (e.g. CHAR(9) = tab)

First Row is a Header: Use escaping characters \0 \\ \t \r \n \r\n to identify special characters

Record Delimiter: If no record delimiter is specified a newline or carriage return, newline is assumed.
For a fixed width record enter FIX nnn where nnn is the record width
Use escaping characters \0 \\ \t \r \n \r\n to identify special characters

6

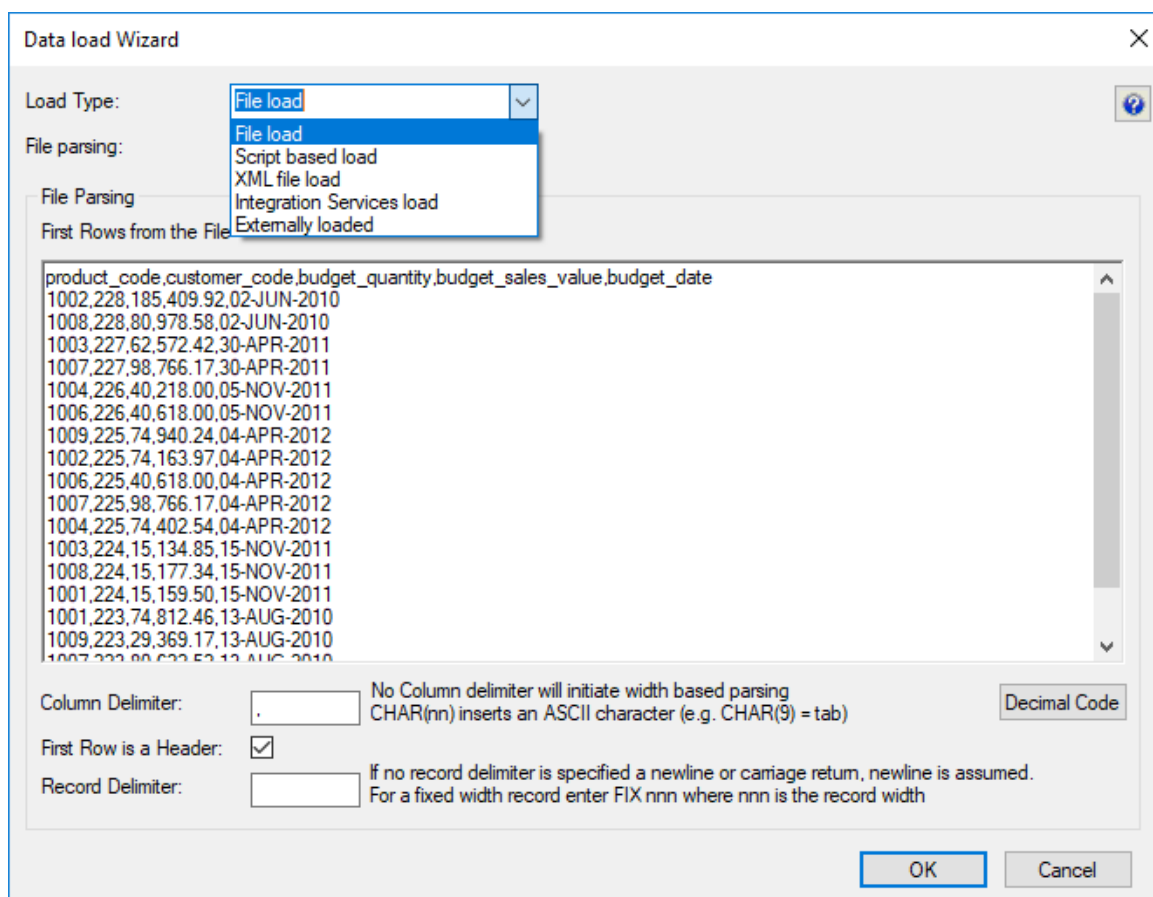
NOTES: First Row is a Header - To make changes in database tables after a table has been defined, users can edit the **First Row is a Header** option in the **Source** tab of the relevant table.

Hive: For File loads into Hive tables, the **First Row is a Header** option is a table option. Please ensure this option is checked in the file load wizard if you want to load files with header rows.

To make any changes after Hive tables have been defined, the **First Row is a Header** option can be found in the **Storage** tab of the relevant Hive table.

- 7 The load type set in the **New Table Default Load Type** field of the Connection Properties screen, is the pre-selected option in the **Load type** drop-down list of the wizard.

To change the desired load type and file parsing, use the **Load type** and **File parsing** drop-down list options.



Load type options

- The **File based load** option results in a load where the bulk of the load management is handled by the scheduler.
- The **Script based load** option will make WhereScape RED generate a host script and the load table will be a **Script based load**. This host script is executed by the scheduler to perform the load. For further details about Script based loads see section **Script based loads**.
- The **XML file load** option will only be an available load option from a Windows connection. To see more details about specific XML loads, see section **XML File load**.
- The **Integration Services load** option will load the file via an Integration Services Package that is generated and executed dynamically at run time. This option is only available from a Windows connection. For specific details about this load option, see section **Loading Data from Flat Files using SSIS**.
- The **Externally loaded** option will not execute an actual load into the table but will process the actions specified in the Post Load procedure property. Any **After** transformations recorded against any of the columns in an Externally loaded table will also be processed.

Post load procedures can either be manually generated from a RED provided procedure outline or generated leveraging a RED template—refer to **Rebuilding Update Procedures** (on page 199) for details.

File parsing options

- **Single data column** - with this option, the majority of the work in terms of parsing the file must occur in a subsequent procedure within the data warehouse. The data is dumped into a single column. The task of coding a procedure to parse the data must then be undertaken. Three columns are created under Oracle. These include the data column, a sequence column (row_sequence) and the file name column (row_file_name). The file name and sequence columns can be deleted if they are not required for a File based load.
- **Columns parsed** - with this option, RED will attempt to parse the columns. You will be asked for details and for the column delimiter. You then step through the columns providing names and data types. RED attempts to guess the data type, but it needs to be checked and the field length will probably need to be adjusted.

The following screen shows the initial file parsed screen.

Data load Wizard

Load Type: File load

File parsing: Columns Parsed

File Parsing

First Rows from the File

```
product_code,customer_code,budget_quantity,budget_sales_value,budget_date
1002,228,185,409.92,02-JUN-2010
1008,228,80,978.58,02-JUN-2010
1003,227,62,572.42,30-APR-2011
1007,227,98,766.17,30-APR-2011
1004,226,40,218.00,05-NOV-2011
1006,226,40,618.00,05-NOV-2011
1009,225,74,940.24,04-APR-2012
1002,225,74,163.97,04-APR-2012
1006,225,40,618.00,04-APR-2012
1007,225,98,766.17,04-APR-2012
1004,225,74,402.54,04-APR-2012
1003,224,15,134.85,15-NOV-2011
1008,224,15,177.34,15-NOV-2011
1001,224,15,159.50,15-NOV-2011
1001,223,74,812.46,13-AUG-2010
1009,223,29,369.17,13-AUG-2010
1007,223,80,633.53,13-AUG-2010
```

Column Delimiter: | No Column delimiter will initiate width based parsing
CHAR(nn) inserts an ASCII character (e.g. CHAR(9) = tab) Use escaping characters \0 \\ \t \r \n \r\n to identify special characters

First Row is a Header:

Record Delimiter: \r\n If no record delimiter is specified a newline or carriage return, newline is assumed.
For a fixed width record enter FIX nnn where nnn is the record width
Use escaping characters \0 \\ \t \r \n \r\n to identify special characters

Decimal Code

OK Cancel

NOTE: The **Decimal Code** button shows the decimal value of each character in the lines retrieved from the source file. These decimal codes is shown below each line and are green.

- 8 Once the screen above is completed, another screen displays to enable the breakdown of the source data into columns.
If **no delimiter is entered** then **width based parsing is assumed** and an addition width size field is prompted for.

If this is a **Fixed Width file** loaded via TPT, the source file format can be specified later in the **File Load Source Screen** of the Load table's Properties dialog.

This following screen is an example of the file parsing technique.

The screenshot shows the 'Data load Wizard - Column Definition' dialog box. It has two main panes: 'Column Data' and 'File'. The 'Column Data' pane lists columns: product_code, 1002, 1008, 1003, 1007, 1004, 1006, 1009, 1002, 1006, 1007, 1004, 1003, 1008. The 'File' pane shows a preview of the source file with columns: product_code, customer_code, budget_quantity, budget_sales_value, budget_date. Below the panes is a 'Display decimal character values' button. The bottom section contains fields for 'Column Name' (product_code), 'Business Display Name' (product code), 'Data Type' (int), a checked 'Nulls' checkbox, 'Conversion', and 'Business Definition' fields. At the bottom right are 'Back', 'Add', and 'Cancel' buttons.

Use the **Back** button to revert to the previous column if an incorrect width or delimiter is entered.

Conversion

During the parsing of the columns a date format conversion string can be used for **date data types only**.

Any other Teradata function can be used with Multi Load or TPT Update for **After Load Transformations only**. Refer to the Multi load and TPT manuals for syntax.

For example, if we are loading a column called 'product_name' we could use the following syntax to bring over only the first 30 characters.

```
substr(product_name, 1, 30)
```

A special variable **%FILE_NAME%** can be used in a File based script load. This will be substituted with the actual source file name.

For example, a transformation such as: '%FILE_NAME%' can be used to store the full file name and path in a database column.

For **After Load** transformations, the conversion string can also be entered in the relevant column conversion field during the parsing of the columns.

However, to process the column conversion, users will need to do the following after the table is created and loaded:

1. Go into the **Properties** of the loaded table's relevant column(s) by double-clicking on the column(s) in the middle pane.
2. Click the **Transformation** tab.
3. Select the **After load** option in the **Transformation Stage** drop-down list.
4. **Recreate** the load table.

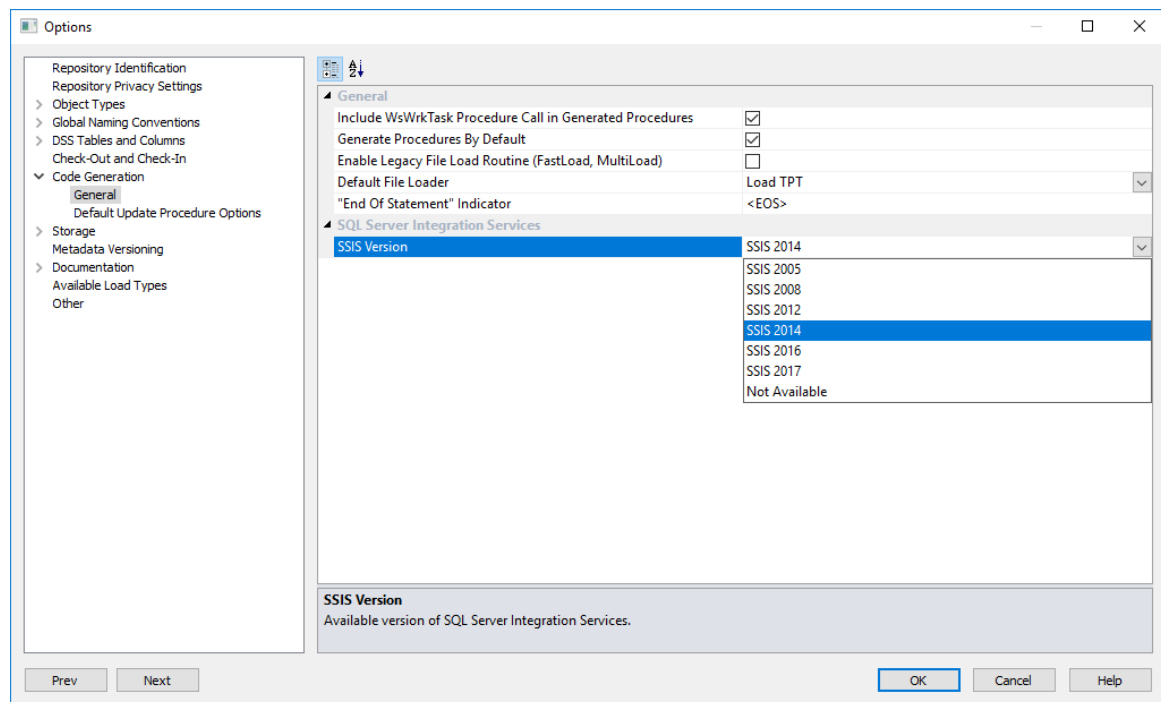
LOADING DATA FROM FLAT FILES USING SSIS

Flat files can be loaded into RED from a **Windows** connection using **SQL Server Integration Services (SSIS)**.


The instructions below detail how to add the SSIS connection string to the data warehouse connection and load flat files, using the drag and drop functionality to create Load tables. To load files via SSIS, the SSIS connection string must be defined in the DataWarehouse connection.

NOTE: SSIS Loads for Teradata only work with a Windows Scheduler.

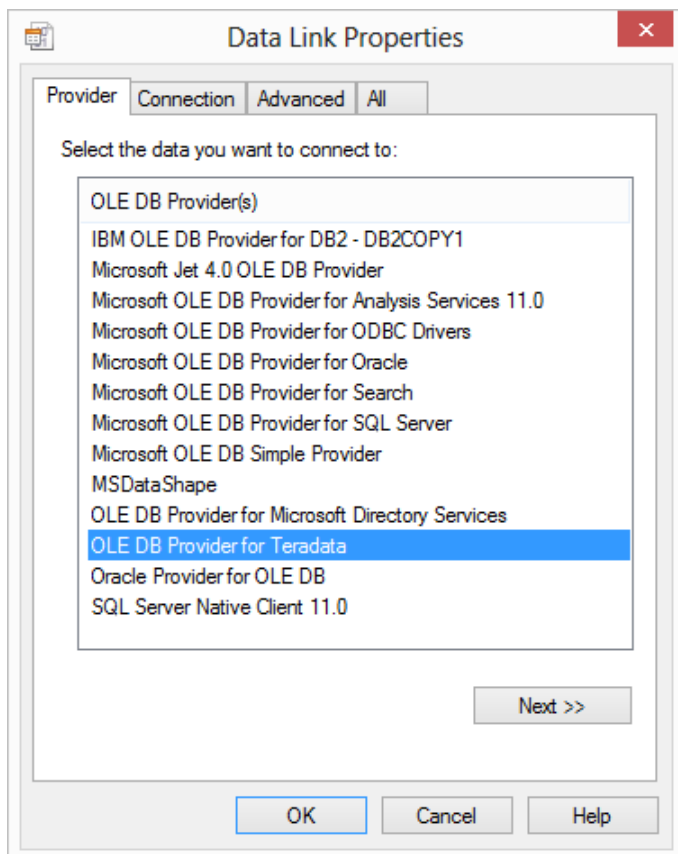
To use SSIS loading, ensure that SSIS loads are enabled and that the SSIS version is set correctly in **Tools>Options>Code Generation>General**.



- 1 To load files via SSIS, the SSIS connection string must be defined in the DataWarehouse connection for the **Destination** connection to be specified:
 - Double-click the **DataWarehouse** connection in the object explorer to launch the Properties screen.
 - Click the ellipsis button to open the wizard and define the SSIS connection string.

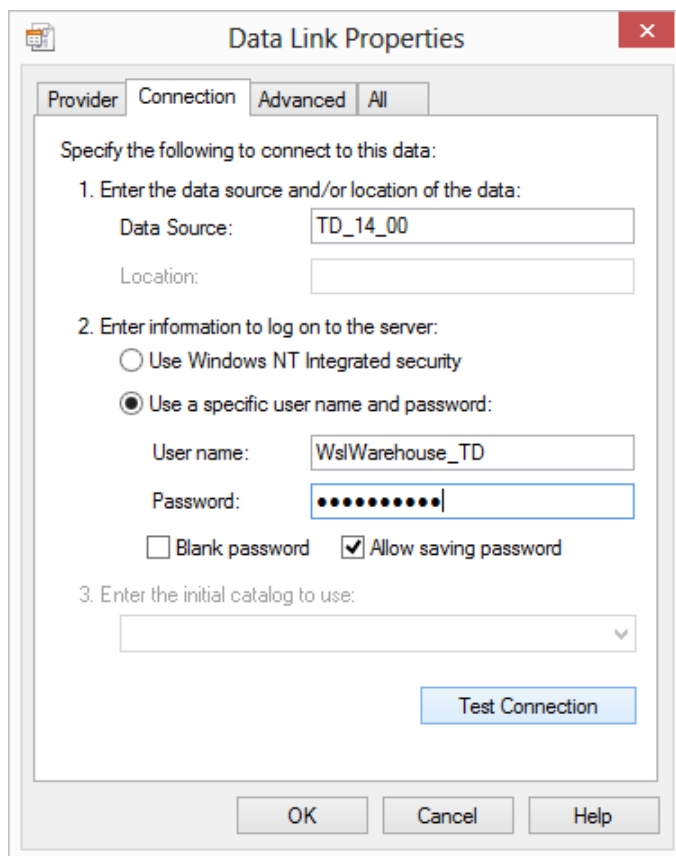


- 2 On the Provider tab, select the **OLE DB Provider for Teradata** and click **Next**.




- 3 On the **Connection** tab, enter the **destination data source name**, enter the **information to log on to the server** and select the **Allow saving password** option. Click **Test Connection**.

NOTE: When using a specific **user name** and **password** to connect to the server, instead of using Windows integrated security, the **Allow saving password** check box must be selected. It is also recommended that the password on the SSIS connection string field that is displayed in the connection properties is replaced with the **\$PASSWORD\$** token that is substituted at runtime.



4 Click **OK**.



- 5 Click **OK** on the Data Link Properties dialog to save the settings.

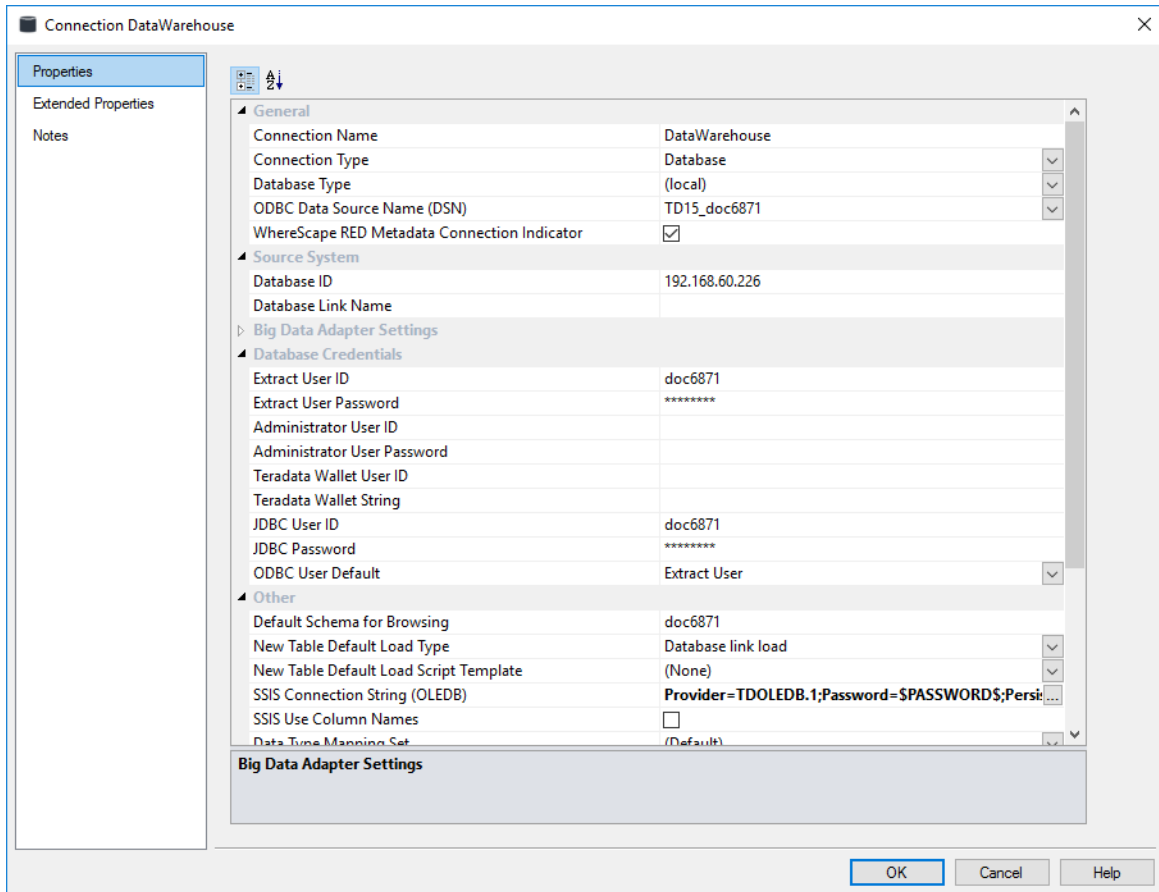
The screenshot shows the 'Data Link Properties' dialog box with the 'Connection' tab selected. The dialog is titled 'Data Link Properties' and has a close button (X) in the top right corner. Below the title bar are four tabs: 'Provider', 'Connection', 'Advanced', and 'All'. The 'Connection' tab is active. The main area contains the following instructions and fields:

Specify the following to connect to this data:

1. Enter the data source and/or location of the data:
Data Source:
Location:
2. Enter information to log on to the server:
 Use Windows NT Integrated security
 Use a specific user name and password:
User name:
Password:
 Blank password Allow saving password
3. Enter the initial catalog to use:

At the bottom of the dialog, there are three buttons: 'OK', 'Cancel', and 'Help'. A 'Test Connection' button is also present, located below the 'Enter the initial catalog to use:' field.

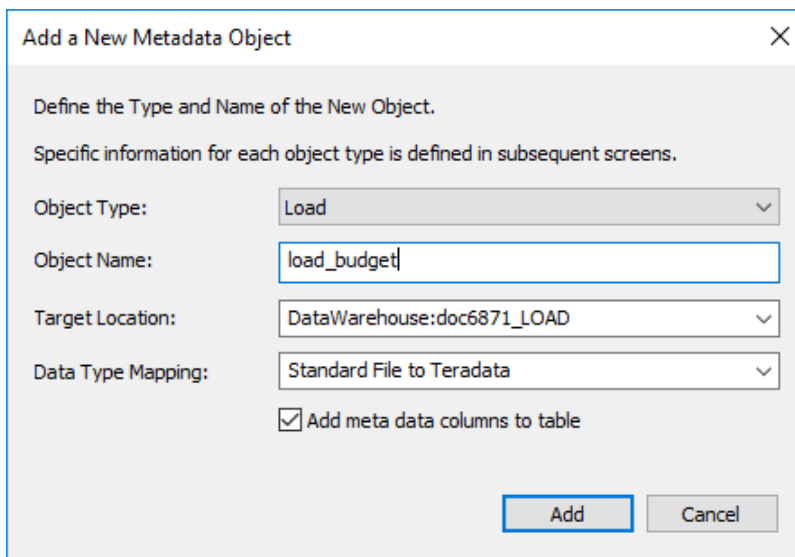
6 Click **OK** to save your connection.



NOTE: It is recommended that the password on the SSIS connection string field that is displayed in the connection properties is replaced with the **\$PASSWORDS** token that is substituted at runtime.

7 Browse to the directory and file from the **Windows** connection.

8 Double-click on the **Load Table** object in the left pane to create a drop target.



- 9 Click the **ADD** button, the following screen appears:

Data load Wizard

Load Type: File load

File parsing: Columns Parsed

File Parsing

First Rows from the File

```
product_code, customer_code, budget_quantity, budget_sales_value, budget_date
1002,228,185,409.92,02-JUN-2010
1008,228,80,978.58,02-JUN-2010
1003,227,62,572.42,30-APR-2011
1007,227,98,766.17,30-APR-2011
1004,226,40,218.00,05-NOV-2011
1006,226,40,618.00,05-NOV-2011
1009,225,74,940.24,04-APR-2012
1002,225,74,163.97,04-APR-2012
1006,225,40,618.00,04-APR-2012
1007,225,98,766.17,04-APR-2012
1004,225,74,402.54,04-APR-2012
1003,224,15,134.85,15-NOV-2011
1008,224,15,177.34,15-NOV-2011
1001,224,15,159.50,15-NOV-2011
1001,223,74,812.46,13-AUG-2010
1009,223,29,369.17,13-AUG-2010
```

Column Delimiter: No Column delimiter will initiate width based parsing
 CHAR(nn) inserts an ASCII character (e.g. CHAR(9) = tab) Decimal Code
 Use escaping characters \0 \\ \t \r \n \r\n to identify special characters

First Row is a Header:

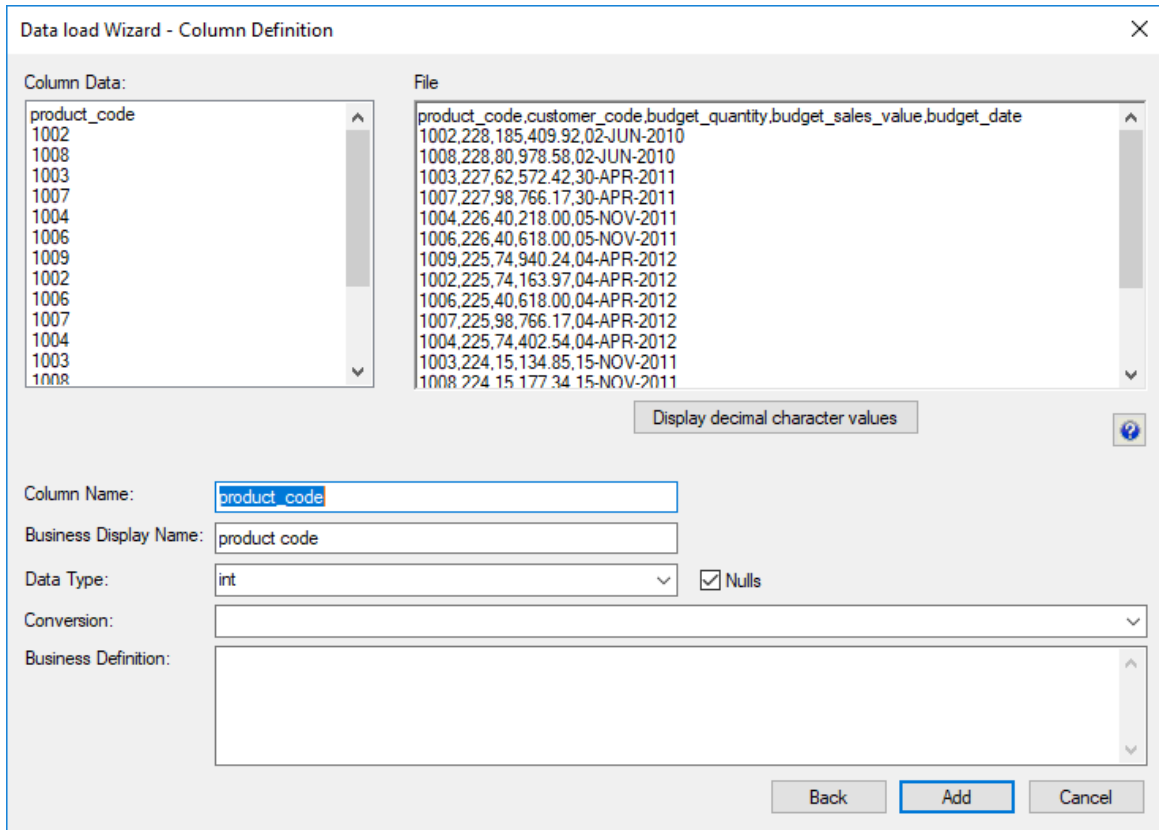
Record Delimiter: If no record delimiter is specified a newline or carriage return, newline is assumed.
 For a fixed width record enter FIX nnn where nnn is the record width
 Use escaping characters \0 \\ \t \r \n \r\n to identify special characters

OK Cancel

- 10 Configure the settings and click **OK** to continue. See the previous section for details about these settings.

Note: The **Decimal Code** button shows the decimal value of each character in the lines retrieved from the source file. These decimal codes are shown in green below each line.

- The next screen enables you to breakdown the source data into columns. If no delimiter is entered then width based parsing is assumed and an addition width size is prompted. The following screen is an example of the file parsing technique.



Use the **Back** button to revert to the previous column if an incorrect width or delimiter is entered.

- On the **Properties** screen for the new load table, select **Integration Services Load** as the Load Type. Click **OK**.

This creates and executes a SSIS package at run time to load data into the data warehouse load table.

NOTES: If the table is changed to an Integration Services load and has been set up using the wizard for the "File load (columns parsed)" flow, some columns might have transformations set up that will not work.

In RED 6.8.4.0 date/time fields have transformations that are invalid for SSIS and will make the load fail.

Since SSIS does not provide any configuration for the parsing of date/time fields, if users have any date/time field special requirements, file or script-based loads provide a better load option instead.

13 Click **Yes** to Create and Load the table.

FLAT FILE LOAD - SOURCE SCREEN

The fields for the **Flat File Source Screen** are described below:



TIP: If the file has been *dragged and dropped* (see "*Flat File Loads*" on page 253) into the load table (middle pane) then some of the fields on this tab are automatically populated.

Property	Value
Load Type	Script based load
Source Connection	Windows
Source File Details	
Source Directory	C:\Program Files (x86)\WhereScape\Tutorial\
Source File Name	budget.txt
Source File Field Delimiter	,
Source File Record Terminator	
Source File has Field Headings/Labels	<input checked="" type="checkbox"/>
Trigger File Details	
Load Configuration	
Check existence of Source File	<input checked="" type="checkbox"/>
Wait for Source File or Trigger File	<input type="checkbox"/>
File Load Routine	FastLoad
MultiLoad/FastLoad Options	...
Character Set	ASCII
Script Load supports File Name Wildcards	<input type="checkbox"/>
Only Load Latest File	<input type="checkbox"/>
Validate using Header/Trailer Data	<input type="checkbox"/>
Populate Audit Log with File Actions	<input checked="" type="checkbox"/>
Populate Detail Log when Successful	<input checked="" type="checkbox"/>
Insert into dss_load_file Table	<input type="checkbox"/>
Archived File Details	

Load Type
Method of loading data into the table. The available options are dependent on the "Source Connection". Defaults to the "Default Load Type" of the "Source Connection". Can be specified via the "Properties" page.

OK Cancel Help

Load Type

Method of loading data into the table. The available options are dependent on the **Source Connection**. Defaults to the 'Default Load Type' of the **Source Connection**. Can be specified via the Properties page.

Source Connection

Connection to the data source (database or file system). Can be specified via the **Properties** page.

Load Script Template

Available for script loads only and only if there is a valid template available. Select the template to use when generating a load script, or select (None) to use RED's built-in load script generator. Only templates with the correct Type and Target DB will appear in this drop-down list. For more information, see **Templates** (on page 673).

Source File Details

Source File identification and definition information.

Source Directory

The full path (absolute path) of the folder/directory containing the Source File on the Windows or UNIX/Linux system.

Source File Name

The name of the source file containing the data to be loaded.

Source File Field Delimiter

Optional character that separates the fields within each record of the Source File. The delimiter identifies the end of each field. Common field delimiters are tab, comma, colon, semi-colon, pipe, tilde.

If no field delimiter is specified the record is regarded as fixed-width.

Note: If an ASCII character value is used this field may show as an unprintable character. To enter a special character enter the uppercase string CHAR with the ASCII value in brackets (e.g. CHAR(9)).

Source Fixed Width File Format

Source File format to be set for TPT loads of Fixed width files. This option is only available for TPT Flat File or Script based loads from Windows and UNIX/Linux connections.

The default option for all Fixed width files is **Text**. To change the file format, select from **Text**, **Binary**, **Formatted** or **Unformatted**.

Source File Record Terminator

Optional string to identify how each line/record in the Source File is ended/terminated/delineated. The system default is used when not specified. On UNIX/Linux systems, end-of-line is typically line-feed (ASCII 10). On Windows systems, end-of-line is typically carriage-return (ASCII 13) and line-feed (ASCII 10).

Source File has Field Headings/Labels

Indicates whether the first line of the Source File contains a heading/label for each field, which is not regarded as data so it should not be loaded.

Trigger File Details

Optional Trigger File identification and definition information. If a Trigger File is specified the Source File will not be loaded until the Trigger File is available.

Trigger File Path

The purpose of the trigger file is to indicate that the copying/loading of the main file has completed and that it is now safe to load the file. Secondly the trigger file may contain control sums to validate the contents of the main load file. This field should contain the full path name to the directory in which a trigger file is located on the Windows or UNIX systems. If this field and/or the **Trigger Name** field is populated then the scheduler will look for this file rather than the actual load file.

Trigger File Name

Refers to the name of the file that is used as a trigger to indicate that the main load file is available. A trigger file typically contains check sums. If the trigger file exists then this is the file that is used in the Wait Seconds, and Action on wait expire processing. (see notes under Trigger Path above).

Trigger File Field Delimiter

If the trigger file provides control information then the delimiter identifies the field separation, e.g, or \n for a return. The data found will be loaded into parameter values whose names will be prefixed by the prefix specified and numbered 0 to n.

Trigger Parameter Name Prefix

If a trigger file and delimiter have been specified then the contents of the trigger file are loaded into parameters. The parameters will be prefixed by the contents of this field and suffixed by an underscore and the parameter number. These parameters can be viewed under **Tools>Parameters** from the WhereScape RED menu bar. The checking of these parameters can be achieved by generating a **Post Load procedure**. An example set of parameters may be budget_0, budget_1 and budget_2 where there are 3 values in the trigger file and the prefix is set to 'budget'.

Load Configuration

Configuration details to control the load processing.

Check existence of Source File

This field controls whether the load process checks if the file exists before performing the load. This check can only be disabled when doing script-based TPT Loads from Windows and UNIX/Linux connections, which can improve the use of built-in TPT functionality to wait for the arrival of the file.

Wait for Source File or Trigger File

Controls whether the load process waits for the file to arrive when it is NOT available to load. This is disabled by default but enabling this allows the wait-related properties to be specified. When a Trigger File is specified the load waits for it rather than the Source File. Expand to set the **Wait Limit** and the **Action when Wait Limit Reached**.

Wait Limit (in seconds)

Maximum duration to wait when no file is available, which is specified in seconds e.g. 1800 seconds to wait up to 30 minutes. A value of 0 equates to no wait. If the wait time expires and the specified file cannot be found, then the load will exit with the status defined in Action, e.g.

Default Action = Error

Action when Wait Limit Reached

Action to take when the Wait Limit has been reached and no file is available. The specified action impacts any remaining tasks in the currently running job. When 'Success' or 'Warning' is specified, the WhereScape RED Scheduler will continue to run any dependent tasks in the running job. In contrast, specifying the 'Error' or 'Fatal Error' actions will cause the scheduler to stop/fail the job and hold any remaining tasks when the "Wait Limit" is reached.

File Load Routine

File Loader utility/mechanism used to load the Source File. Select **MultiLoad**, **FastLoad**, **Load TPT**, **Update TPT** or **Stream TPT**. If the 'No Load' option is specified the file does not load. This can be useful for a Script Load that has "File Actions".

Note: The MultiLoad and FastLoad options are only available, if the **Enable Legacy File Load Routine** option is set in the Tools>Options>Code Generation>**General** (on page 108).

MultiLoad/FastLoad Options

If **MultiLoad** or **FastLoad** is selected, optional comma-delimited list of TPT Operator Attributes. e.g. INTEGER DataBlockSize =2048.

Character Set

Teradata-compliant Character Set Name to use when loading such as ASCII, UTF8, UTF16.

TPT HadoopHost

Field that identifies the Hadoop host to the TPT load routine. This is set to the **TPT HadoopHost** property of the Hadoop connection if specified there. If this not specified in the connection, the field will display the UNIX/Linux Host name property of the Hadoop Connection.

TPT Job Name

Job name for TPT. If this is not set it will default to Load Name.

TPT Build Command Options

Additional options included as part of the TBuild call.

TPT LogView Command Options

Additional options included as part of the TLogView call. For example, adding "-f '*'" can be added to prove greater diagnostics.

TPT Shared Memory Size

Shared memory size can be specified in bytes, kilobytes or megabytes. Examples include 2091752, 2048K, 2M.

TPT Load Routine Minimum Sessions

Optional minimum number of TPT sessions [1-99]. The default is one instance.

Set Load Routine Maximum Sessions

Enables the "Load Routine Maximum sessions" property.

Load Read Instances

Optional number of TPT read instances [1-99]. The default is one instance.

Load Write Instances

Optional number of TPT write instances [1-99]. The default is one instance.

Only Load Latest File

Controls whether only the latest file (or all matching files) is loaded when the "Source File Name" includes a wildcard and multiple matching files are available.

Script Load supports File Name Wildcards

This option is only available for **Update TPT** or **Stream TPT** load routines.

It controls whether the RED generated script supports loading multiple files based on a file name wildcard. When enabled and the "Source File name" contains a wildcard, the RED generated script will loop to load each matching file while preserving the contents of the load table as each file is loaded. In addition the "Archived Source Path" and/or "Archived Source File Name" properties must be specified to allow each successfully loaded file to be archived before loading a subsequent file.

When doing **Hadoop Native TPT loads** where the file load routine is **Update TPT**, users can load multiple files from a Hadoop connection by adding * to the **Source File Name**. e. g.

hadoop_customer.csv*

Fail when incomplete load

Controls whether the load reports failure when ALL the rows in the file are not loaded. The specified exit status impacts any remaining tasks in the currently running job. When fail is specified, The WhereScape RED Scheduler will stop/fail the job and hold any remaining tasks when the load fails. In contrast when fail is disabled, the scheduler will continue to run any dependent tasks in the running job.

Validate using Header/Trailer Data

Controls whether a RED-generated Post-Load procedure validates the integrity of the Source File and the load processing using control totals &/or counts ("check sums") from the FIRST and/of LAST line(s) of the Source File. When enabled, any subsequently RED-generated Post-Load procedure for this table will use the HEADER and/or TRAILER rows as a source of validation information. Alternatively, A Trigger File can be used for the same validation purpose.

Populate Audit Log with File Actions

Controls whether details of any "File Actions" processing are recorded in the WhereScape RED Audit Log.

Populate Detail Log when Successful

Controls whether additional details from the Teradata "File Load Routine" are recorded in the WhereScape RED Error/Detail Log when the load processing is successful.

Insert into dss_load_file Table

Populate the metadata table dss_load_file.

Archived File Details

Optional Archived file identification and definition information. Once a file has been successfully loaded or processed it can be optionally archived by moving it and/or renaming it.

Compress Source File when Archive

Optionally compresses the successfully loaded Source File if it is archived.

Archived Source File Path

Optional full path (absolute path) of the folder/directory to MOVE the successfully loaded Source File on the Windows or UNIX/Linux system.

Archived Source File Name

Optional new name to RENAME the successfully loaded Source File to. By default, the original file name is used it is optional to rename it. However, the in-built variable \$SEQUENCE\$ can be used to include a unique sequence number in the new name. Likewise, the in-built variables \$YYYY\$, \$MM\$, \$DD\$, \$HH\$, \$MI\$ and/or \$SS\$ can be used to include the number of the current year, month, day, hour, minute and/or second respectively. The date/time components can be used separately or can be combined together using one set of enclosing \$ such as \$YYYYMMDD\$.

Archived Trigger File Path

Optional full path (absolute path) of the folder/directory to MOVE the successfully processed Trigger File on the Windows or UNIX/Linux system.

Archived Trigger File Name

Optional new name to RENAME the successfully processed Trigger File to. By default, the original file name is used it is optional to rename it. However, the in-built variable \$SEQUENCE\$ can be used to include a unique sequence number in the new name. Likewise, the in-built variables \$YYYY\$, \$MM\$, \$DD\$, \$HH\$, \$MI\$ and/or \$SS\$ can be used to include the number of the current year, month, day, hour, minute and/or second respectively. The date/time components can be used separately or can be combined together using one set of enclosing \$ such as \$YYYYMMDD\$.

SQL Server Integration Services (SSIS)

SQL Server Integration Services (SSIS) Attributes.

SSIS Destination-Identifier Case Conversion

Case-sensitivity conversion applied to Destination Object Identifiers (such as table, view, and column names) in RED-generated SSIS packages. If no conversion is applied then the exact case of the identifier defined in the RED metadata is used in SSIS.

SSIS Set Source-Code Page

Enables the SSIS source code page property.

SSIS Destination-AlwaysUseDefaultCodePage

Forces the use of the Default CodePage property value when describing character data.

SSIS Set Destination-Code Page

Enables the SSIS destination code page property.

SSIS Row Count Log

During an SSIS Load include Row Count logging.

SCRIPT BASED LOADS

A script based load table must have a Host Script defined. During the load process this host script is executed and the results returned.

During the 'drag and drop' creation of a load table from a Windows file, a script can be generated by selecting one of the 'Script based' load options.

This script can then be edited to fully meet any requirements.

There are a number of conventions that must be followed if these host scripts are to be used by the WhereScape scheduler.

- 1** The first line of data in 'standard out' must contain the resultant status of the script. Valid values are '1' to indicate success, '-1' to indicate a Warning condition occurred but the result is considered a success, '-2' to indicate a handled Error occurred and subsequent dependent tasks should be held, -3 to indicate an unhandled Failure and that subsequent dependent tasks should be held.
- 2** The second line of data in 'standard out' must contain a resultant message of no more than 256 characters.
- 3** Any subsequent lines in 'standard out' are considered informational and are recorded in the audit trail. The normal practice is to place a minimum of information in the audit trail. All bulk information should be output to 'standard error'
- 4** Any data output to 'standard error' will be written to the error/detail log. Both the audit log and detail log can be viewed from the RED tool under the scheduler window.
- 5** When doing **Script based loads**, it is easy to use the **Regenerate** button on the **Script Name** field to regenerate the scripts.

The screenshot shows a dialog box titled "Load load_budget" with a close button (X) in the top right corner. On the left is a sidebar with a tree view containing the following items: Properties (selected), Storage, Override Create DDL, Source, Extended Properties, and Notes. The main area contains the following fields and controls:

- Load Table Name:
- Unique Short Name: (maximum 22 characters)
- Description:
- Connection:
- Load Type:
- Database Link:
- Script Template:
- Script Name:
- Pre-Load Action:
- Pre-Load SQL:
- Post Load Procedure:
- Timestamps section with three sub-sections:
 - Metadata Structure:
 - Database:
 - Database:

At the bottom right are three buttons:

Note: Script-based loads on Windows supports both DOS Batch and *PowerShell scripts* (see "*Windows PowerShell Scripts*" on page 653).

XML FILE LOAD

XML file loads are only supported from a Windows connection. There are multiple formats for data exchange when using XML. See the following section for details on how an XML file is handled.

For more details about XML File load Properties and Source screen fields, please see ***Flat File Loads*** and ***File Load - Source Screen***.

Note: The XML file load in a Windows connection is **not** enabled by default for newly created Teradata repositories.

Refer to Tools>Options>**Available Load Types** (see "**Settings - Available Load Types**" on page 120) for details.

To load an XML file located in a Windows directory proceed as follows:

- 1 Create a connection to the Windows system.
- 2 Browse to the connection and locate the XML file.
- 3 Make Load tables the middle pane drop target by double clicking on the Load Table object group in the left pane.
- 4 Drag the XML file into the middle pane.

The only rules concerning the xml file are that the data element tags are the column names and each row of data is a child of the root element.

For example:

```
<row>
  <dim_customer_key>7</dim_customer_key>
  <code>228</code>
  <name>JOHN AND JOES TOYS</name>
  <address>3700 PARNELL RISE</address>
  <city>BEAVERTON</city>
  <state>OR</state>
  <dss_source_system_key>1</dss_source_system_key>
  <dss_update_time>2003-10-03T10:02:15.310</dss_update_time>
</row>
```

Supported XML Formats

WhereScape RED supports two types of xml file construct. The normal xml standards have the data in the xml file and the table definitions in a separate xsd (xml schema definition) file which is only required when the table is being created or when the xml file is being validated for form. An alternate standard is used by Microsoft. This second standard is an in line definition which produces one file which contains a Schema element in the data stream where the column names and their approximate data types are defined.

Separate XML and XSD files

The normal XML standards have the data in the xml file and the table definitions in a separate xsd (xml schema definition) file which is only required when the table is being created or when the xml file is being validated for form. The xsd file name is found within the xml file in an xsi (xml schema instance) statement which can include a namespace definition; e.g.:

```
<root xmlns="http://www.wherescape.com/wsl-schema"
      xmlns:xsi="http://www.w3.org/2000/10/XMLSchema-instance"
      xsi:schemaLocation="http://www.wherescape.com/load_table.xsd">
```

or no namespace;e.g.

```
<root xmlns="http://www.wherescape.com/wsl-schema"
      xmlns:xsi="http://www.w3.org/2000/10/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="load_table.xsd">
```

The xsd file is an xml file and should be found in the same directory as the xml file that calls it. This xsd file will contain the column definitions for the load table which will be defined during the drag and drop.

The column definitions within the xsd file must be detailed enough to define a load table that the xml file can be loaded into.

The data type mapping between the xsd file and the database have been implemented as below:

XSD	Teradata
string with length	char()
string with maxlength	varchar()
integer	integer
decimal with precision and scale	numeric(x,y)
dateTime (ISO8601)	timestamp
i2	integer
i4	integer
r4	varchar(40)
r8	varchar(40)
float	varchar(40)

These are the ISO-ANSI SQL/XML standards and in the case of integers, dateTime and floats the column can be defined with one line; e.g.:

```
<xsd:element name="Policy_ID" type="xsd:integer"/>
<xsd:element name="Quote_Date" type="xsd:dateTime"/>
<xsd:element name="Quote_Price" type="xsd:r4"/>
```

In the case of strings and decimals the column requires a bit more detail to produce the correct data type. Strings can be fixed length with padded data by using the length attribute. The following will produce a char(1) column called Excess_Waiver:

```
<xsd:element name="Excess_Waiver">
  <xsd:restriction base="xsd:string">
    <xsd:length value="1"/>
  </xsd:restriction>
</xsd:element>
```

Strings can be of variable length by using the maxLength attribute. The following produces a column of varchar(8) called Password:

```
<xsd:element name="Password">
  <xsd:restriction base="xsd:string">
    <xsd:maxLength value="8"/>
  </xsd:restriction>
</xsd:element>
```

Decimal numbers are defined with the precision and scale attributes. If the scale attribute is zero or missing then the column will be a whole number of size precision. The following produces a column of numeric(6):

```
<xsd:element name="code" >
  <xsd:restriction base="xsd:decimal">
    <xsd:precision value="6"/>
    <xsd:scale value="0"/>
  </xsd:restriction>
</xsd:element>
```

The following produces a column of numeric(8,2):

```
<xsd:element name="code" >
  <xsd:restriction base="xsd:decimal">
    <xsd:precision value="8"/>
    <xsd:scale value="2"/>
  </xsd:restriction>
</xsd:element>
```


An example file with most data types would be as follows:

```
<xsd:schema xmlns="http://www.wherescape.com/wsl-schema"
            xmlns:xsd="http://www.wherescape.com/XMLSchema">
<xsd:element name="Col_name1" type="xsd:integer"/>
<xsd:element name="Col_name4" type="xsd:dateTime"/>
<xsd:element name="Col_name5">
  <xsd:restriction base="xsd:string">
    <xsd:maxLength value="100"/>
  </xsd:restriction>
</xsd:element>
<xsd:element name="Col_name6">
  <xsd:restriction base="xsd:string">
    <xsd:length value="100"/>
  </xsd:restriction>
</xsd:element>
<xsd:element name="Col_name7" type="xsd:float"/>
<xsd:element name="Col_name8" >
  <xsd:restriction base="xsd:decimal">
    <xsd:precision value="6"/>
    <xsd:scale value="2"/>
  </xsd:restriction>
</xsd:element>
</xsd:schema>
```

The column order will be the same as the xsd file.

Any columns which are missing from the row will be NULL in the loaded row.

The dateTime format in the xml file is defined as ISO8601 which looks like this:

```
2003-10-03T10:02:15.310
```

WhereScape RED will load this string into Teradata as:

```
CAST('20031003100215' AS TIMESTAMP FORMAT 'YYYYMMDDHHMISS')
```

- The xsd file is only required to create the load table, if the load table is only being loaded then this file is ignored.
- To check that the xml and xsd files are well formed you can open them with any web browser. If the files display with no errors then they are valid xml files.

In line schema definition

The other supported xml construct allows the use of in line schema definitions as produced by the Microsoft FOR XML AUTO, ELEMENTS, XMLDATA query. This will produce one file which contains a Schema element in which the column names and their approximate data types are defined. Because the supplied data types are not concise enough to define the table columns correctly, this method will produce load tables of data type varchar(4000). The column names are taken from the <element type="col_name"/> strings within the Schema element. The data elements will be the same as above with the column names making up the start and end tags and the rows being the children of the root element. The file that is produced by the FOR XML query above needs to be changed slightly to comply with the xml standard. Remove everything before the Schema element and then give the file a starting root element and a closing root element.eg <root> and </root>

The xml files can optionally start with an xml pre process statement.eg

```
<?xml version="1.0"?>
```

They may also contain xml comments.eg

```
<!-- comments -->
```

EXTERNAL LOAD

For an externally loaded table the only property that is executed is the Post Load procedure.

Any **After** transformations recorded against any of the columns in an Externally loaded table will also be processed.

Post load procedures can either be manually generated from a RED provided procedure outline or generated leveraging a RED template—refer to *Rebuilding Update Procedures* (on page 199) for details.

APACHE SQOOP LOAD

The **Apache Sqoop load** type enables loading data directly from Hive/HDFS to any (non-Hive) targets, however, for loading data directly from Hive into Teradata, it is recommended to use the TPT functionality instead.

When processing Teradata loads from Hive/HDFS using Apache Sqoop, WhereScape recommends that you use vendor supplied drivers, such as ‘Hortonworks Connector for Teradata’.

Scheduler loads and Apache Sqoop loads from a Hive connection to a Hive target are not supported.

The following known issues exist when using Sqoop loads with the generic JDBC driver on Teradata:

- 1 Sqoop loads from Hive to Teradata fail if column names and titles are different.
- 2 For Sqoop loads from Hive to Teradata where column names and titles are the same, complete the following to enable the load to work:

In the Load table properties, select the Source tab and on the Generic Hadoop Arguments field enter the following command: `-Dsqoop.export.records.per.statement=1`

To load tables directly from HDFS/Hive into Teradata target databases using the Apache Sqoop load:

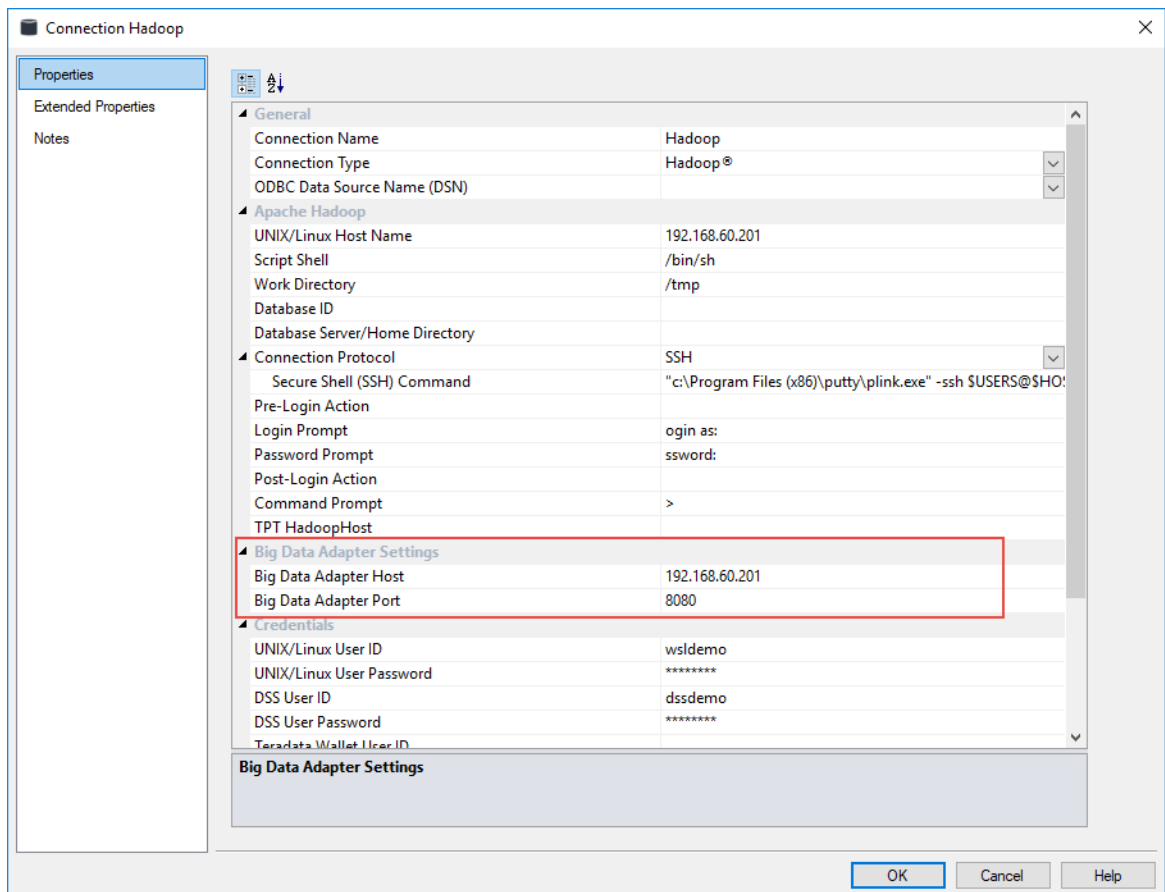
- 1 Ensure the relevant **Data Warehouse connection** has the following fields set:
 - JDBC Connection string (JDBC URL)
 - JDBC Driver Class Name
 - Omit Sqoop Driver Option - tick this check-box for loads into Oracle target databases
 - JDBC User ID
 - JDBC Password

The screenshot shows the 'Connection DataWarehouse' dialog box with the following configuration:

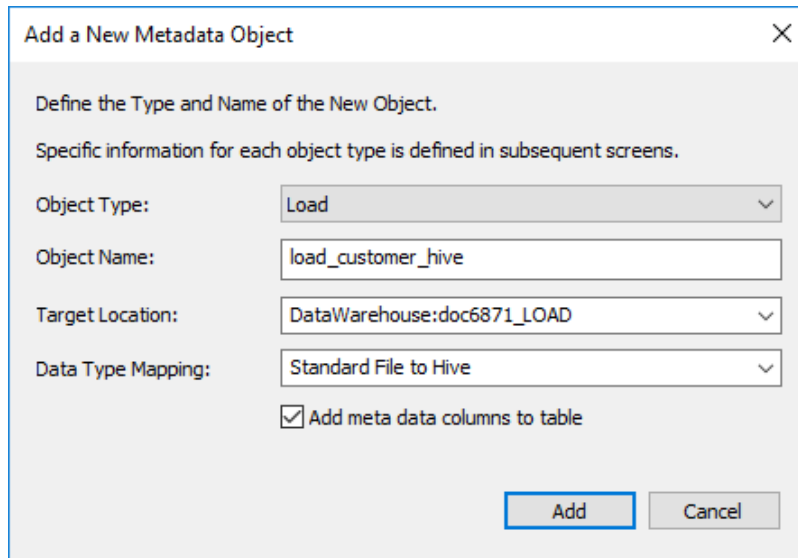
Property	Value
Connection Name	DataWarehouse
Connection Type	Database
Database Type	Teradata
ODBC Data Source Name (DSN)	TD15_doc6871
WhereScape RED Metadata Connection Indicator	<input checked="" type="checkbox"/>
Source System	
Database ID	192.168.60.226
Database Link Name	
Big Data Adapter Settings	
JDBC Connection String (JDBC URL)	jdbc:teradata://192.168.60.226/DATABASE=WSL_REPOS...
JDBC Driver Class Name	com.teradata.jdbc.TeraDriver
Omit Sqoop Driver Option	<input type="checkbox"/>
Sqoop Connection Manager Class	
Include Database/Schema Name in Sqoop Table Option	<input checked="" type="checkbox"/>
Include Sqoop Columns Option	<input checked="" type="checkbox"/>
Database Credentials	
Extract User ID	doc6871
Extract User Password	*****
Administrator User ID	
Administrator User Password	
Teradata Wallet User ID	
Teradata Wallet String	
JDBC User ID	doc6871
JDBC Password	*****
ODBC User Default	Extract User

NOTE: Users loading into Teradata from a Hive or Hadoop connection using the Teradata connection manager for Sqoop who need to load into more than one database will need to add `DATABASE=$OBJECT_DATABASE$` into their JDBC Connection String (JDBC URL) field on their Teradata DataWarehouse connection, e.g.
`jdbc:teradata://192.168.60.226/DATABASE=$OBJECT_DATABASE$`.
 BDA will replace `$OBJECT_DATABASE$` with the database containing their load table when doing an Apache Sqoop load.

- 2 When doing loads from **Hadoop connections**, please ensure the Hadoop connection has its **BDA server host** and **port** fields set in addition to Hive connections.



- 3 Browse the desired **Hadoop/Hive connection**.
- 4 **Drag and Drop** the table from the Hadoop/Hive connection on the right hand-side into the middle pane.
 - Change the table name if necessary and select the relevant target location to place the table from the drop-down list.



Add a New Metadata Object

Define the Type and Name of the New Object.
Specific information for each object type is defined in subsequent screens.

Object Type: Load

Object Name: load_customer_hive

Target Location: DataWarehouse:doc6871_LOAD

Data Type Mapping: Standard File to Hive

Add meta data columns to table

Add Cancel

NOTE: The option **Add meta data columns to table** is used for creating load tables that are used in creating Data Vault objects. If this option is selected, two DSS columns (**dss_record_source** and **dss_load_date**) are included in the meta data for the table and are populated by transformations. These two DSS columns could equally be applied to other load tables not used in a Data Vault system but are particularly important to comply with the Data Vault standards. Please refer to the *Data Vaults chapter* (see "*Data Vaults*" on page 413) for more details.

5 Select Apache Hadoop Load from the Load Type drop-down list.

Load Table load_customer_hive

Properties

Storage

Override Create DDL

Source

Extended Properties

Notes

Load Table Name: load_customer_hive

Unique Short Name: (maximum 22 characters) load_customer_hive

Description:

Connection: Hive

Load Type: Native ODBC

Database Link:

Script Template:

Script Name: (None)

Pre-Load Action: Truncate

Pre-Load SQL:

Post Load Procedure: (None) [Rebuild] [Regenerate]

Timestamps

Metadata Structure	Database	Database
2017-03-13 04:45:29.480000		

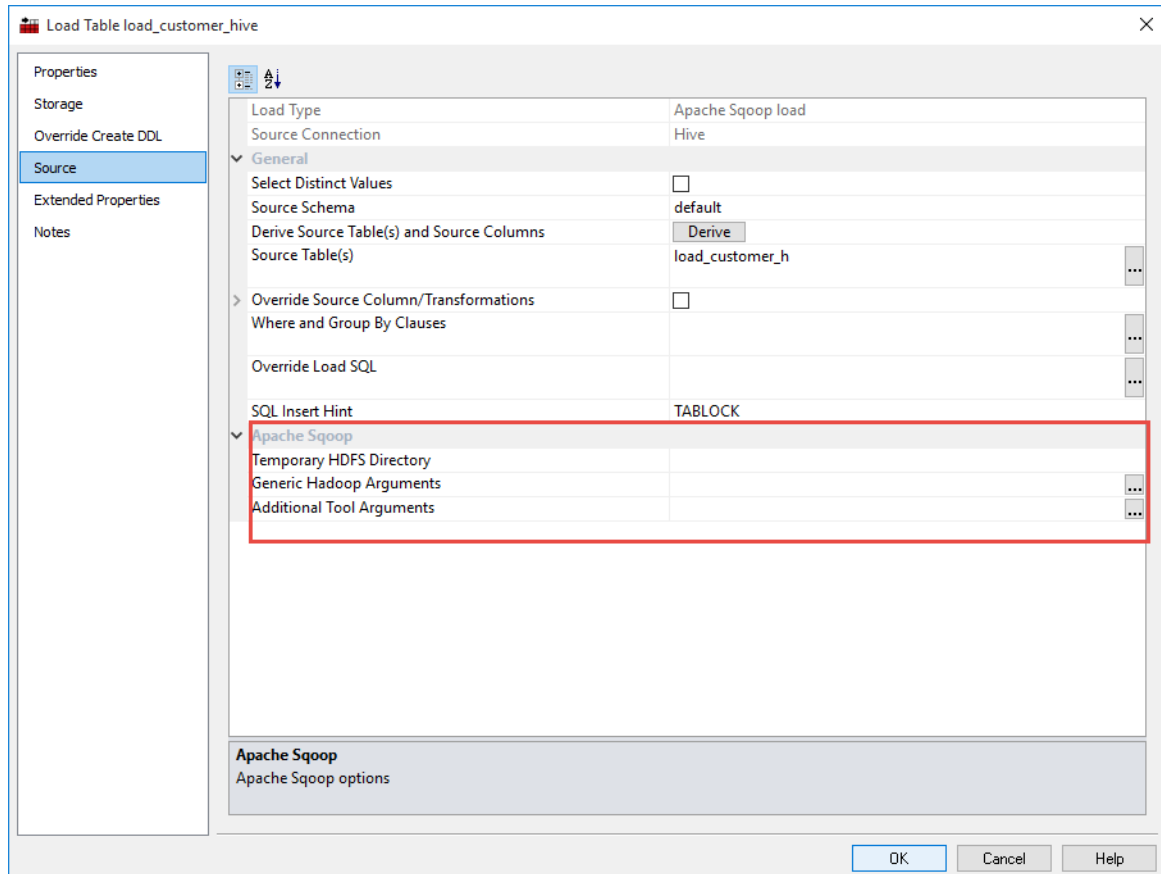
[OK] [Cancel] [Help]

6 Click the Source tab to add any Apache Sqoop specific options:

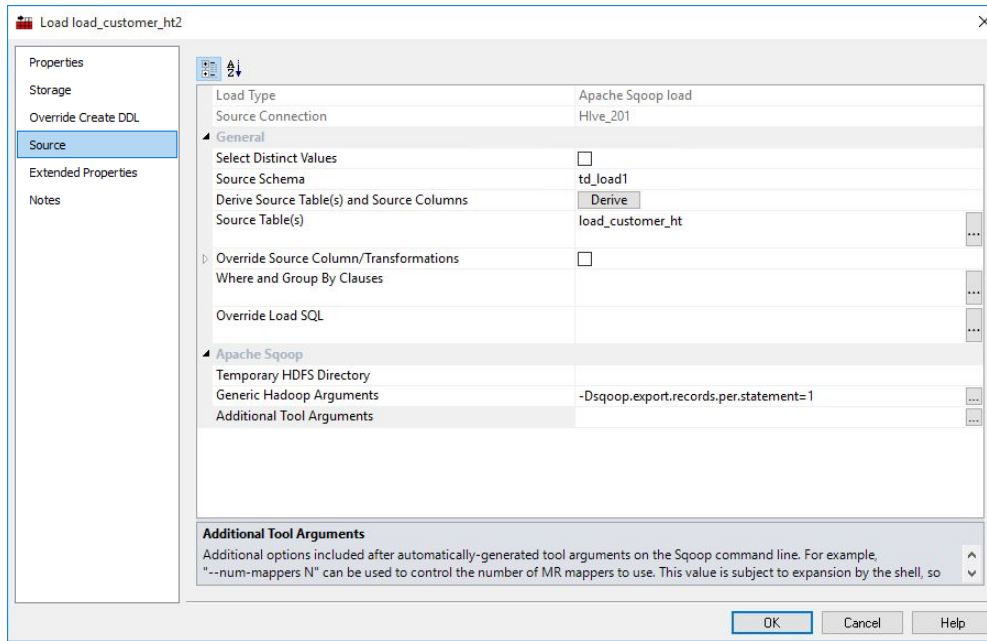
- Temporary HDFS Directory - Loading from a Hive source to the data warehouse is implemented in two steps. First, the data is extracted from the Hive table into a temporary HDFS directory. Then, this temporary directory is loaded using "sqoop export".

The location of the temporary directory can be configured in RED on the source tab of the load table. When this field is left blank, the default is "/tmp".

- Generic Hadoop Arguments - This field allows adding additional arguments just after the Sqoop command keyword, in this case it is import, in the Sqoop command line.
- Additional Tool Arguments - This field allows adding additional arguments after the generated Sqoop command line.



NOTE: Sqoop loads from Hive to Teradata fails if column names and titles are different. Please ensure column names and titles are the same and then, on the Load table properties, go to the Source tab and on the Generic Hadoop Arguments field specify the following command=: -Dsqoop.export.records.per.statement=1.



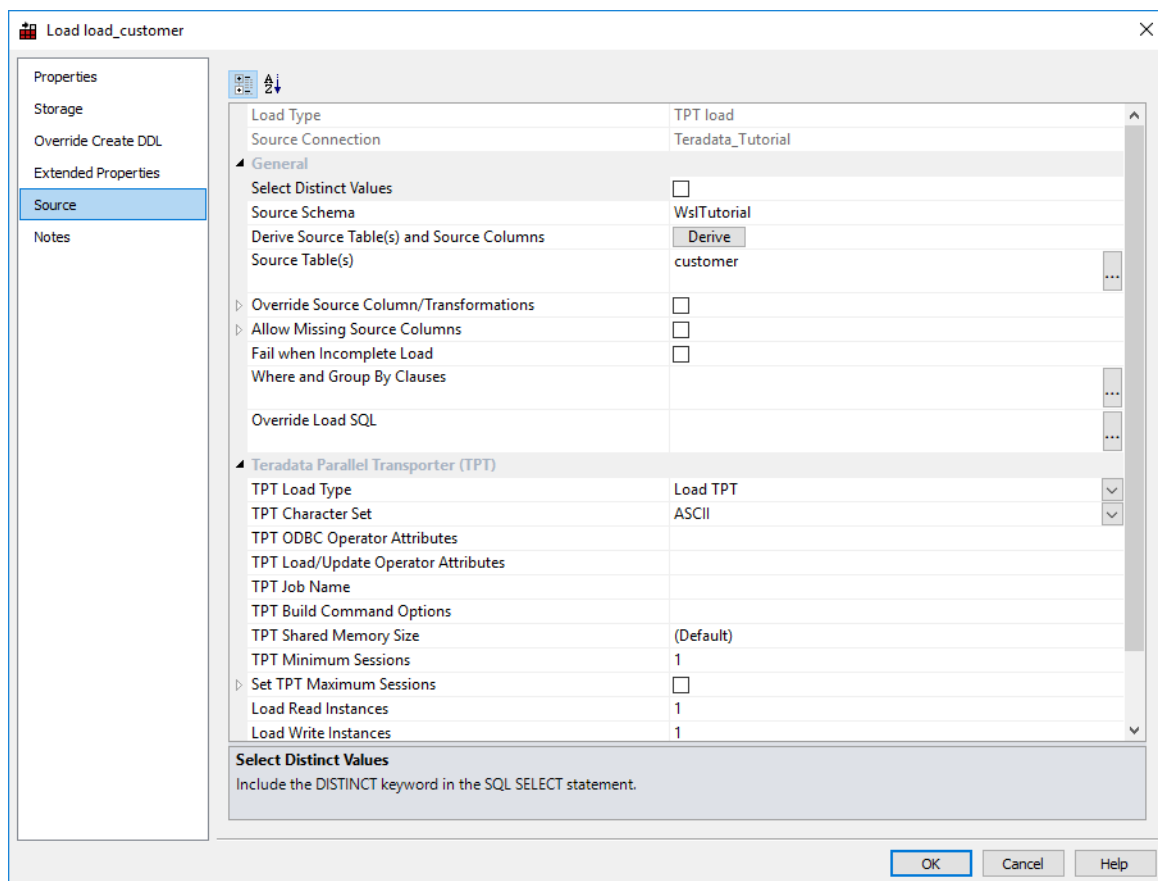
7 Click **Create and Load** to create and load the table.

HANDLING MISSING SOURCE COLUMNS

By default, a load fails if a source column that is to be loaded does not exist. This default action can be modified by using the 'Non mandatory source column' feature of a load table. When this feature is enabled, the load process checks all source columns, and if any are found to be missing—the missing columns are replaced with a Null value (which can be changed, see below).

On the **Source** tab of a load tables properties there are two check boxes and a drop down list that is covered in this section.

See the following example:



Override Source Column/Transformation

When this is enabled the **Source Columns** edit window is enabled. With this enabled, a load table uses the contents of the **Source Columns** to specify which columns are being extracted during a table load.

When this is unselected the load process builds up the statement to select from the source table(s) by looking at the two fields source table and source column and any transformations associated with each column in the load table. If a transformation is present then the transformation overrides the contents of source table/source column. See the following section for more information on transformations.

Note: This check box must be unselected to enable **Allow Missing Source Columns** support.

Allow Missing Source Columns

When this check box is selected the load process will examine every column in the load table. Based on the source table/source column fields associated with each column it will check to see if the column exists in the source database. If the column exists normal processing will occur. If the column does not exist then a Null will be substituted for the column, and the load will proceed.

If one or more columns are found to be missing the load process reports this situation. The status level of this reporting can be set via the **Exit Status** drop-down. See the following topic. In all cases the load will be deemed to have been successful, if no other errors occur.

Often Null values are not desirable in a data warehouse. This Null value can be replaced by some other value by means of a **During** or **After** transformation. For example, a **During** transformation, as shown below, set on a missing column called 'State' will replace the Null with the value 'N/A'.

```
NVL(state,'N/A')
```

Exit Status When Missing Columns

If columns are found to be missing as a result of a 'Non Mandatory' check then a message is recorded in the Audit trail and against the task if running in the Scheduler. The drop-down list of the same name allows the choice of action in the event of a missing column or columns. The choices are:

Choice	Impact
Success	The message informing that columns are missing uses an information status.
Warning	The message issued is a warning. The table load will be identified as successful but with warnings.
Error	The message issued is an error. The table load will still complete successfully albeit with an error message.

This drop-down list is only available if the **Non Mandatory Source Columns** option is set.

Limitations of Missing Column Check

The check for missing columns only examines and check those columns that are recorded in the source table and source column fields of a load table's column properties. Therefore, if a source column is used in a transformation or join, but is not recorded in the two fields mentioned above, it is not checked.

If a column is used in a transformation, it should have the same usage (name and case) as is found in the source column and source table fields.

This check has no upper limit on the number of missing columns. All columns can be missing and the load still succeeds.

LOAD TABLE TRANSFORMATIONS

Each load table column can have a transformation associated with it.

See *Transformations* (on page 593) and *Load Table Column Transformations* (on page 597) for more details.

POST-LOAD PROCEDURES

If a procedure is specified in the **Post Load Procedure** drop-down field of a Load table's Properties, then this procedure will be executed after the load has completed and after any subsequent transformations have occurred.

Click the **Rebuild** button to start the process of generating a new procedure. The **Regenerate** button is used to regenerate the Load table's existing post load procedure.

See *Load Table Column Transformations* (on page 597) for more details.

Post load procedures can either be manually generated from a RED provided procedure outline or generated leveraging a RED template—refer to *Rebuilding Update Procedures* (on page 199) for details.

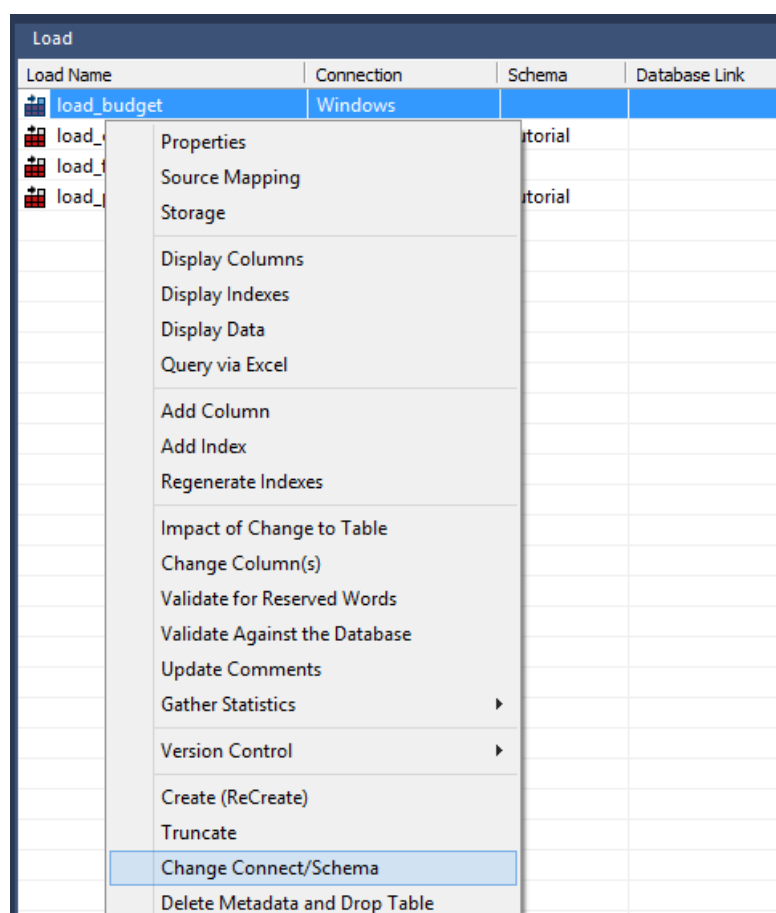
CHANGING LOAD CONNECTION AND SCHEMA

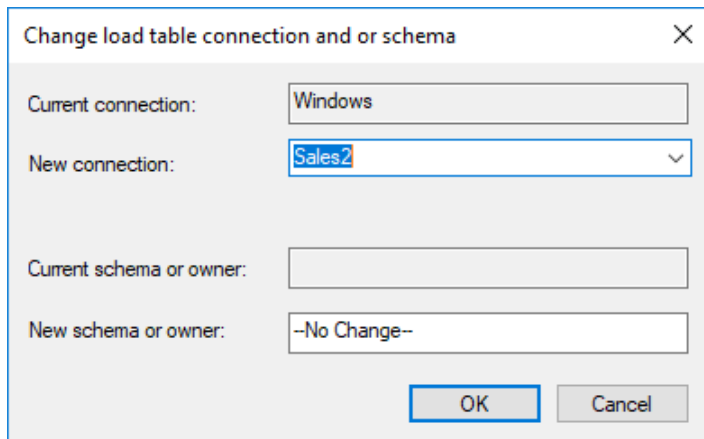
The connection associated with a load table can be changed through the properties of that table.

Connections can also be changed in bulk by using the following process:

- 1 Double click on the **Load Table** object group in the left pane. This displays all load tables in the middle pane.
- 2 Select those load tables that you wish to change using standard Windows selection.
- 3 Right-click to bring up a menu and select **Change Connect/Schema**.
- 4 Select the new connection to change all the selected load tables.

Note: You cannot change the connection **type** but it is possible to change from Database to ODBC connections when the following considerations are taken into account.





Change load table connection and or schema

Current connection: Windows

New connection: Sales2

Current schema or owner:

New schema or owner: -No Change--

OK Cancel

Switching Connection from ODBC to Database and vice versa

This switch should be successful in most cases, although it may not provide the most performant load possible.

By default, ODBC connections use the source table/column transformation loading method as dates and potentially other data types need to be converted.

When switching to a database link load, any transformations still occurs although they may no longer be necessary.

CHAPTER 10

DIMENSIONS

IN THIS CHAPTER

Dimensions Overview	294
Building a Dimension	295
Generating the Dimension Update Procedure	303
Dimension Artificial Keys	314
Dimension Column Properties	316
Dimension Column Transformations	325
Dimension Hierarchies	326
Snowflake.....	330
Dimension Language Mapping	332

DIMENSIONS OVERVIEW

A dimension table is normally defined, for our purposes, as a table that enables us to constrain queries on the fact table.

A dimension is built from the Data Warehouse connection. Unless you are doing a retro-fit of an existing system, dimensions are typically built from one or more load tables.

The normal steps for creating a dimension are defined below and are covered in this chapter. The steps are:

- Identify the source transactional data that will constitute the dimension. If the data is sourced from multiple tables, ascertain if a join between the source tables is possible, or if a series of lookups would be a better option.
- Using the 'drag and drop' functionality drag the load table that is the primary source of information for the dimension into a dimension target. See *Building a Dimension*.
- If only one table is being sourced and most of the columns are to be used (or if prototyping) you can select the auto create option to build and load the dimension and skip the next 4 steps. See *Building a Dimension*.
- Add columns from other load tables if required. See *Building a Dimension*.
- Create the dimension table in the database. See *Building a Dimension*.
- Build the update procedure. See *Generating the Dimension Update Procedure* (on page 303).
- Run the update procedure and analyze the results. See *Dimension Initial Build*.

Modify the update procedure as required. See *Dimension Initial Build*.

Dimension Keys

Dimensions have two types of keys that we will refer to frequently. These are the **Business Key** and the **Artificial Key**. A definition of these two key types follows:

Business Key

The business key is the column or columns that uniquely identify a record within the dimension. Where the dimension maps back to a single or a main table in the source system, it is usually possible to ascertain the business key by looking at the unique keys for that source table. Some people refer to the business key as the 'natural' key. Examples of business keys are:

- The product SKU in a product dimension
- The customer code in a customer dimension
- The calendar date in a date dimension
- The 24 hour time in a time dimension (e.g. HHMM) (e.g. 1710)
- The airport short code in an airport dimension.

It is assumed that business keys will never be NULL. If a null value is possible in a business key then the generated code will need to be modified to handle the null value by assigning some default value. For example, the 'Where' clause in a dimension update may become:

```
Where coalesce(business_key,'N/A') = coalesce(v_LoadRec.business_key,'N/A')
```

Note: Business keys are assumed to never be Null. If they could be null it is best to transform them to some value prior to dimension or stage table update. If this is not done an unmodified update will probably fail with a duplicate key error on the business key index.

Artificial Key

The artificial key is the unique identifier that is used to join a dimension record to a fact table. When joining dimensions to fact tables, it would be possible to perform the join using the business key. For fact tables with a large number of records, this however would result in slow query times and very large indexes. As query time is one of our key drivers in data warehouse implementations, the best answer is always to use some form of artificial key. A price is paid in the additional processing required to build the fact table rows, but this is offset by the reduced query times and index sizes. We can also make use of database specific features, such as bitmap indexes in Oracle.

The artificial key is an integer and is built sequentially from 1 upwards. See the section on artificial keys for a more detailed explanation. An artificial key is sometimes referred to as a "surrogate" key.

BUILDING A DIMENSION

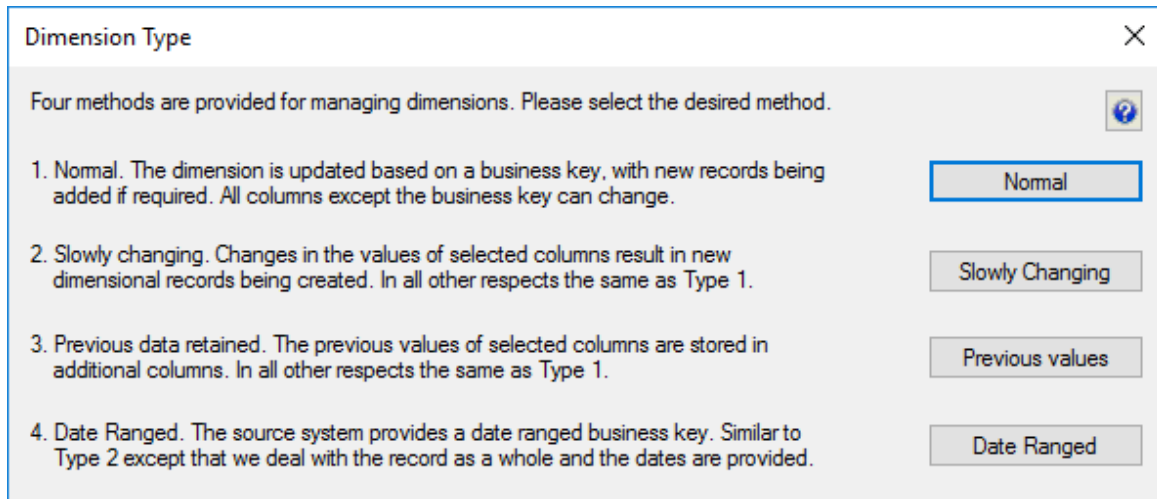
Dimensions are often sourced from one table in the base application. In many cases, there are also codes that require description lookups to complete the de-normalization of the dimensional data. The process for building a dimension is the same for most other tables and begins with the drag and drop of the load table that contains the bulk of the dimensional information.

Drag and Drop

- 1 Create a dimension target by double-clicking on the **Dimension group** in the left pane.
- 2 The middle pane displays a list of all existing dimensions, when this list is displayed in the middle pane, the pane is identified as a target for new dimension tables.
- 3 Browse to the Data Warehouse via the Browse>Source Tables menu option.
- 4 Drag the load table, that contains the bulk of the dimensional columns, into the middle pane.
- 5 Drop the table anywhere in the pane.
- 6 The new object dialog box appears and identifies the new object as a Dimension and provides a default name based on the load table name—either accept this name or enter the name of the dimension.
- 7 Click **OK** to proceed.

Dimension Type

A dialog appears as shown below. There are four choices for the default generation of the dimension table and its update procedure.



- The first choice being a **normal** dimension where a dimensional record is updated and changed whenever any of the non business key information changes (see more details below).
- The second choice is a **slowly changing** dimension where new dimension records are created when certain identified columns in the dimension change (see more details below).

Note: For custom database targets, you can set the initial default values for the DSS columns that are used in the procedure generation for slowly changing dimensions. Please refer to the section *DSS Columns for Custom Targets* (on page 104) for details.

- The third choice is a **Previous values** dimension, which allows the storing of the last values of selected fields in secondary columns.
- The fourth choice is a **Date Ranged** dimension, which supports source systems that provide start and end dates.

With any dimension, we identify a **business key** that uniquely identifies the dimension records.

For example in the case of the product dimension, the product **code** is deemed to be the business key. The code uniquely identifies each product within the dimension. The product may also have a name or description and various other attributes that distinguish it. (e.g. Size, shape, color, etc.).

A common question when handling dimensions is what to do when the name or description changes:

- Do we want to track our fact table records based only on the product code? or
- Do we also want to track records based on different descriptions?

An example :

code	description	product_group	sub_group
1235	15oz can of brussel sprouts	canned goods	sprouts

This product has been sold for many years and we consequently have a very good history of sales and the performance of the product in the market. The company does a '20% extra for free' promotion for 3 months during which time it increases the size of the can to 18oz. The description is also changed to be '15 + 3oz can of brussel sprouts'. At the end of the promotion the product is reverted to its original size and the description changed back to its original name.

The question is do we want to track the sales of the product when it had a different description (slowly changing) , or should the description of the product simply change to reflect its current name (normal). For this scenario a previous value dimension would not provide much advantage, so it is not discussed.

The decision is not a simple one and the advantages and disadvantages of each of the two choices is discussed below.

Slowly Changing

- Allows the most comprehensive analysis capabilities when just using the product dimension.
- Complicates the analysis. Does not allow a continuous analysis of the product called '15oz can of brussel sprouts' when the description is used. This analysis is however still available through the code which has not changed.
- Adds considerable additional processing requirements to the building of the fact tables that utilize this dimension.
- May track data quality improvements rather than real business change.

Normal

- Does not allow specific analysis of the product during its size change. Note, however that this analysis will probably be available through the combination of a 'promotion' dimension.
- Provides a continuous analysis history for the product called '15oz can of brussel sprouts'. An analysis via description and code will produce the same results.
- Simplifies analysis from an end user's perspective.

As mentioned above the choice is never a simple one. Even among experienced data warehouse practitioners there will be a variety of opinions. The decision must be based on the business requirements. In many cases keeping the analysis simple is the best choice, at least in the early stages of a data warehouse development. Slowly changing dimensions do have a place, but there is nearly always an alternate method that provides equal or better results. In the example above a promotion dimension coupled with the product dimension could provide the same analysis results whilst still keeping product only analysis simple and easy to understand.



TIP: Do not over complicate the design of an analysis area. Keep it simple and avoid the unnecessary use of slowly changing dimensions.

Dimension Properties

- Once the dimension type is chosen the object Properties screen appears.
- Change the storage options if required.
- If prototyping, and the dimension is simple (i.e. one source table) then it is possible to create, load and update the dimension in a couple of steps. If you wish to do this, select the **(Build Procedure...)** option from the **Update Procedure** drop-down and answer **Create and Load** to the next question.

The screenshot shows the 'Dimension dim_product' dialog box with the 'Properties' tab selected. The dialog contains the following fields and controls:

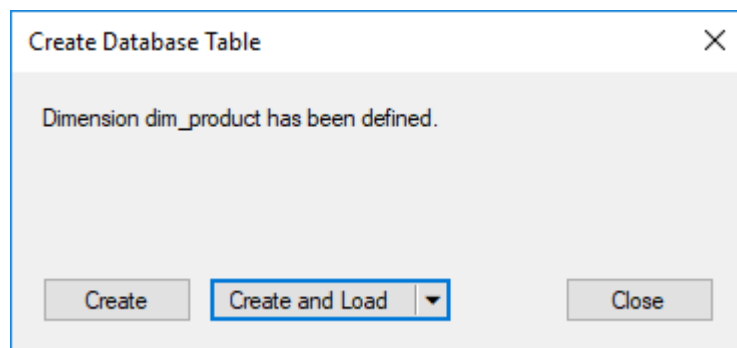
- Table Name:** dim_product
- Table Type:** Dimension
- Unique Short Name:** dim_product (maximum 22 characters)
- Business Display Name (EUL):** dim_product
- Description:** (Empty text area)
- Update Procedure:** (Build Procedure...) with a **Rebuild** button
- Custom Procedure:** (None) with **Rebuild** and **Regenerate** buttons
- Mnemonic (EUL):** (Empty text field)
- Timestamps:**
 - Metadata Structure:** 2017-03-27 02:50:04.640000
 - Database:** 2017-03-27 02:51:59.630000
 - Database Altered:** 2017-03-27 02:51:59.630000

Buttons at the bottom: **OK**, **Cancel**, **Help**.

Note: You can also generate an update procedure via a template, refer to *Rebuilding Update Procedures* (on page 199) for details.

Create and Load

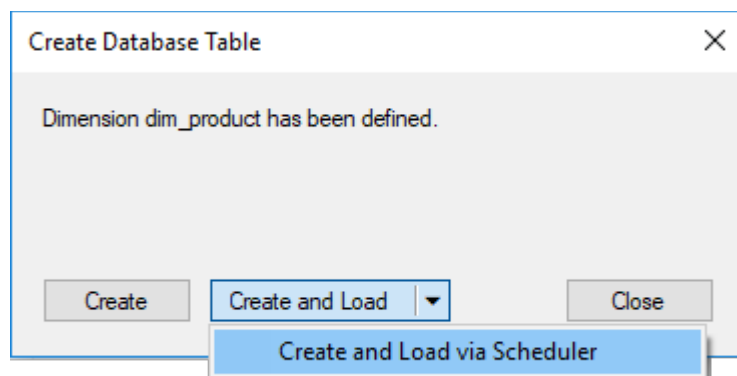
If you chose to build the update procedure, the following dialog appears after clicking **OK** on the Properties screen. This dialog asks if you want to create the Dimension table in the database and execute the update procedure.



If you are satisfied with the columns that will be used and do not wish to add any additional columns, you can select the **Create and Load** button. Alternatively, the **Create** button creates the table in the repository but does not execute an update, allowing you to change columns before loading data into the table.

If **Create** or **Create and Load** is selected and a new procedure creation was chosen on the Properties dialog you can proceed directly to the *Generating the Dimension Update Procedure* (on page 303) section.

Note: It is possible to create and load the table via the Scheduler; by selecting this option from the drop-down list on the **Create and Load** button:



If you have additional columns to add or columns to delete then select **Close** and proceed as follows.

Deleting and Changing columns

The columns defined for the dimension is displayed in the middle pane.

- It is possible to delete any unwanted columns by highlighting a column name or a group of names and choosing the **Delete** key.
- You can also change the name of a column by selecting the column and using the right-click menu to edit its properties. Any new name must conform to the database naming standards.
- Good practice is to use alphanumeric characters and the underscore character. See the section on column properties for a fuller description on what the various fields mean.



TIP: When prototyping and in the initial stages of an analysis area build, it is best not to remove columns nor to change their names to any great extent. This type of activity is best left until after end users have used the data and provided feedback.

Adding additional columns

- With the columns of the dimension table displayed in the middle pane, this pane is considered a drop target for additional columns.
- It is simple to select columns from other load tables and to drag these columns into the middle pane.

The following column list is from the product table as supplied in the tutorial data set.

Column Name	Display Name	Data Type	Source Table
dim_product_key	dim product key	integer	
code	code	number(6)	load_product
description	description	varchar2(64)	load_product
prod_line	prod line	varchar2(24)	load_product
line_description	line description	varchar2(64)	load_prod_line
prod_group	prod group	varchar2(24)	load_product
group_description	group descript...	varchar2(64)	load_prod_group
subgroup	subgroup	varchar2(24)	load_product
subgroup_description	subgroup desc...	varchar2(64)	load_prod_subgro...
dss_start_date	dss start date	date	
dss_end_date	dss end date	date	
dss_current_flag	dss current flag	varchar2(1)	
dss_version	dss version	integer	
dss_update_time	dss update time	date	

The source table shows where each column was dragged from. Although not the case in the tutorial, it is often common to have columns of the same name coming from different tables. In the example above the description column is acquired from the load_product, load_prod_group and load_prod_subgroup tables. In order that the dimension table be created we need to assign these columns unique names, so for this example the last two columns in question have been renamed to group_description and subgroup_description.

There are a number of columns that do not have a source table. These columns have been added by WhereScape RED, and are added depending on earlier choices.

A description of these columns follows.

Column name	description
dim_product_key	The unique identifier (artificial key) for the dimension. This key is used in the joins to the fact table. It is generated via a sequence associated with the table, except for the date dimension where it has the form YYYYMMDD
dss_start_date	Used for slowly changing dimensions. This column provides a date time stamp when the dimension record came into existence. It is used to ascertain which dimension record should be used when multiple are available.
dss_end_date	Used for slowly changing dimensions. This column provides a date time stamp when the dimension record ceased to be the current record. It is used to ascertain which dimension record should be used when multiple are available.
dss_current_flag	Used for slowly changing dimensions. This flag identifies the current record where multiple versions exist.
dss_source_system_key	Added to support dimensions that cannot be fully conformed, and the inclusion of subsequent source systems. See the ancillary settings section for more details.
dss_version	Used for slowly changing dimensions. This column contains the version number of a dimension record. Numbered from 1 upwards with the highest number being the latest or current version. It forms part of the unique constraint for the business key of a slowly changing dimension.
dss_update_time	Indicates when the record was last updated in the data warehouse.
dss_create_time	Indicates when the record was first created in the data warehouse

Manually adding previous value columns

If a **Previous value** type of dimension is chosen, or in fact if the dimension is converted to this type, it is possible to manually add any required columns that were not defined as part of the create. The steps are:

- 1 Add a new column by dragging in the column that is to have a previous value stored.
- 2 Change the name to a unique name. Typically by adding the prefix 'prev_' to the column name.
- 3 Change the source table, to be that of the dimension we are building.
- 4 Set the Key Type to 4.
- 5 Having performed these actions WhereScape RED will detect the column and build the appropriate code during the procedure generation phase.

Create the table

Once the dimension has been defined in the metadata you need to physically create the table in the database.

- This is done by right-clicking on the dimension name and selecting **Create (ReCreate)** from the pop-up menu.
- The results dialog box display the results of the creation.
- The contents of this dialog are a message to the effect that the dimension table was created. A copy of the actual database create statement, and if defined the results of any index create statements is listed. For the initial create, no indexes are defined.
- If the table was not created then ascertain and fix the problem. A common problem is a 'Duplicate column' where a column has the same name in two of the source tables.
- The best way of finding this a column is to double click on the list heading 'Col name'. This sorts the column names into alphabetic order. Another double click on the heading will sort the columns back into their create order.

The next section covers the *Generating the Dimension Update Procedure* (on page 303).

GENERATING THE DIMENSION UPDATE PROCEDURE

Once a dimension has been defined in the meta data and created in the database, an update procedure can be generated to handle the joining of any tables and the update of the dimension records.

Notes: You can also generate an update procedure via a template, refer to *Rebuilding Update Procedures* (on page 199) for details.

The zero key row

WhereScape RED always inserts a record into the dimension with an artificial key value of zero by default. This record is used to link any fact records that do not have valid dimension joins. The values of the various columns in this record are acquired from the contents of the field **Zero Key Value** which is set in the properties screen of each dimension column.


Generating a Procedure

- 1 Double click on the dimension object to edit the properties for the dimension.
- 2 From the **Update Procedure** drop-down list, select (**Build Procedure...**).
- 3 Click **OK** to update the properties and start the process of generating the new procedure.
- 4 The **Update Build Options** screen displays, requesting you to select the relevant build options.
- 5 There are other steps to complete during the procedure generation based on the type of load information. These steps are described below.

Processing tab

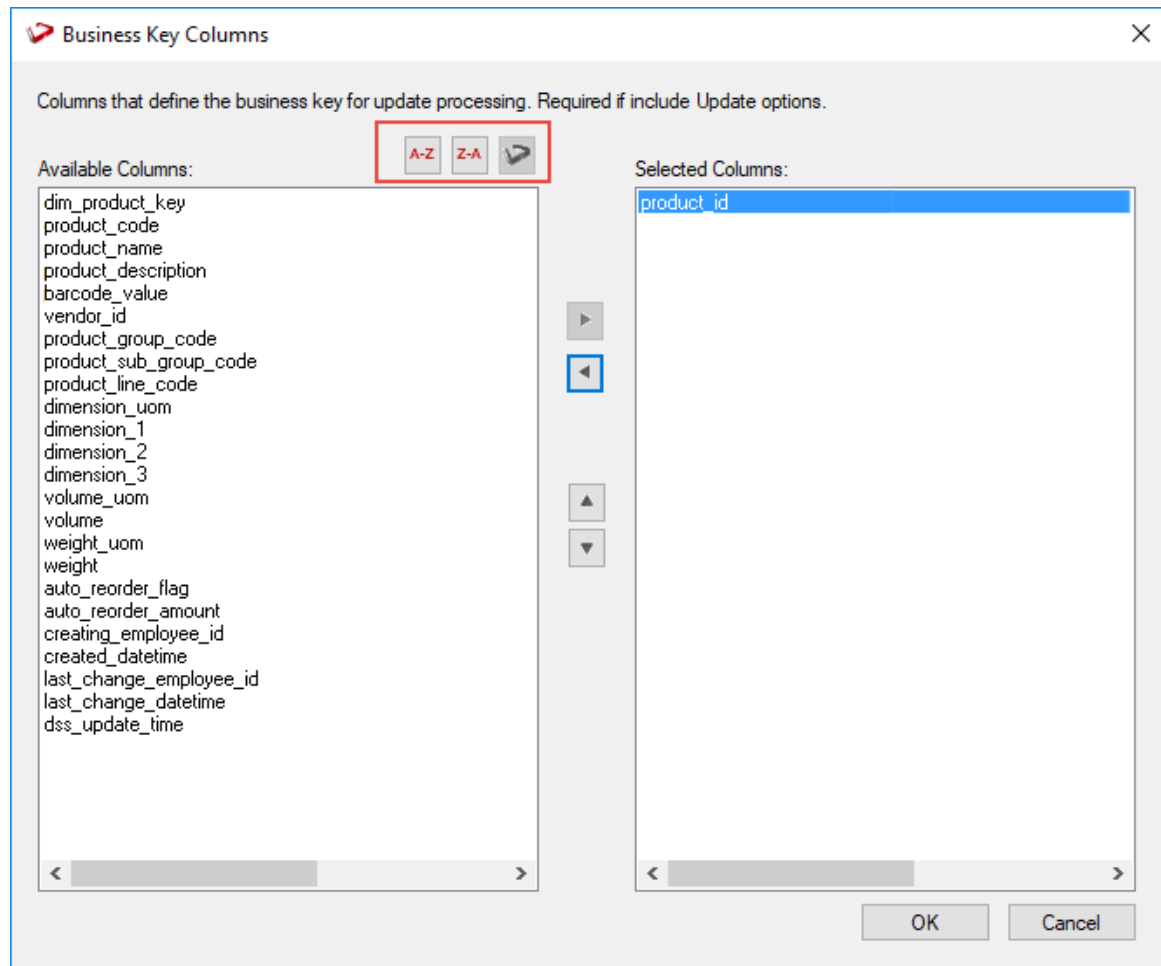
Option	Value
Business Key Columns	dim_product_key
Parameters	
Insert Zero Key Record	<input checked="" type="checkbox"/>
Include Initial Load Insert	<input type="checkbox"/>
Create Storage Index After CTAS	<input type="checkbox"/>
Batch Processing	
Process by Batch	<input type="checkbox"/>
Explicit Locking	
Include Explicit Lock	<input checked="" type="checkbox"/>
Delete Processing	
Delete Before Insert	No
Update Processing	
Process Method	Insert/Update
Source Table Locking	
dim_product	ACCESS
Insert Method	
Include Insert Statement	<input checked="" type="checkbox"/>
Insert Hint	
Insert New Rows Only	<input checked="" type="checkbox"/>
New Row Identification Method	Join
Update Method	
Include Update Statement	<input checked="" type="checkbox"/>
Update Hint	
Update Changed Rows Only	<input checked="" type="checkbox"/>
Change Row Identification Method	Join

Business Key Columns: Columns that define the business key for update processing. This is required for include Update options. The source table from which the dimension is derived would normally have some form of unique constraint applied. In most cases, this is the business key. In the example below **product_id** is selected as the business key.

- Clicking on the ellipsis button  on the rightmost of the **Business Key Columns** field will bring up the Business Key selection screen.
- Select the relevant **Business Key**.
- Click **OK** on the Update Build Options dialog.



TIP: Use the column name **ascending/descending** buttons to sort column names. To revert to the meta column order, click on the **meta column order** button.

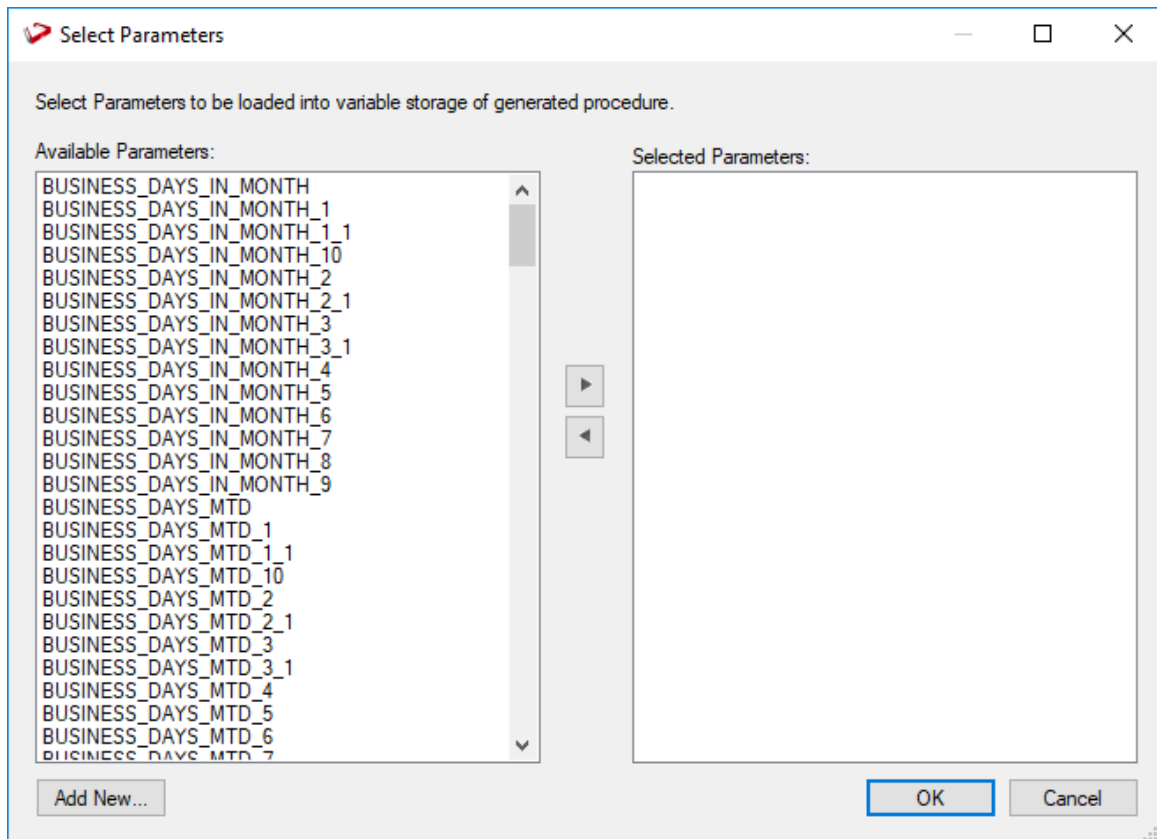


A business key uniquely identifies each dimension record and it can be made up of multiple columns, but it must provide a unique identifier. Where multiple columns uniquely and separately identify the dimension, choose one to act as the primary business key. For example, a source table may have a unique constraint on both a product code and a product description. Therefore the description as well as the code must be unique. It is of course possible to combine the two columns, but the normal practice would be to choose the code as the business key.

NULL Values: None of the columns chosen as the business key should ever contain a NULL value. See the note at the start of this chapter.

Parameters: Any parameters selected are included in the generated update procedure as variables. The procedure will include code to retrieve the value of the parameter at run time and store it in the declared variable.

- Clicking the ellipsis button brings up the Parameters selection screen.



The variables can also be used in column transformations and in the from/where clause for the update procedure. Some databases have a 30 character limit for variable names. WhereScape RED ensures the variables added for any parameters are less than 30 characters long by creating variable names in the form v_ followed by the first 28 characters of the parameter name.

For example, a parameter called MINIMUM_ORDER_NUMBER_SINCE_LAST_SOURCE_LOAD will be available as the variable v_MINIMUM_ORDER_NUMBER_SINCE_L.



TIP1: WhereScape RED parameters should be unique within the first 28 characters to avoid conflicting variables names.



TIP2: If the desired parameter doesn't exist in the metadata yet, a new parameter can be added by clicking the **Add New** button on the bottom leftmost corner of the Select Parameters dialog.

See *Parameters* (on page 143) for more information on WhereScape RED Parameters.

Include Initial Load Insert: adds an additional insert statement to the update procedure that runs if the target Dimension is empty. The benefit of this is improved performance inserting into an empty table without performing any checks to see if rows already exist. The default for this field is not set (i.e. an initial insert statement is not added to the procedure).

Process by Batch: enables users to select a column to drive data processing in a loop based on the distinct ordered values of the selected Business Key columns. The update procedure loops on this column and performs the delete, update and/or insert for each value. If the column chosen is a date datatype (date, datetime or timestamp), then the user is able to specify yearly, monthly, daily or column level looping. The default for this field is not set (do not do batch processing).

Delete Before Insert: enables selecting how to process deletes. It enables a delete statement to be added to the update procedure before any update or insert statement. This is a particularly useful option for purging old data and for updates based on a source system batch number. When this option is selected, it enables the **Issue Warning if a Delete occurs** and the **Delete Where Clause** Fields.

Issue Warning if a Delete occurs: this option sets the procedure to a warning state if any deletes occur.

Delete Where Clause: the delete where clause is appended to the generated delete statement to constrain the rows deleted.

Process Method: enables updating the Dimension with either an **Insert/Update** or a **Merge** statement. **Merge** allows you to use one Merge statement instead of two separate Insert and update statements.

Source Table Locking: enables a locking request modifier to be specified for each source table. The specified locking request modifier is applied to each source table during generated update procedures. By default this is set to 'ACCESS' which locks each row being accessed, a blank entry will result in no locking clause in the generated procedure. This option may also be presented in a separate dialog.

Insert Method

Include Insert Statement: set this field to include the insert statement in the procedure. This allows inserting new rows in the Dimension.

Insert New Rows Only: uses change detection to work out which rows will require inserting.

New Row Identification Method: method used to identify that records in source are not currently recorded in the target table. Select **Join** or **Minus**.

Include Update Statement: set this field to include an update statement in the procedure. This allows updating the changing rows in the Dimension. If this is set, the **Update Changed Rows Only** option is available.

Update Changed Rows Only: uses change detection to work out what rows require updating. When set, this option enables the **Change Row Identification Method**.

Change Row Identification Method: method used to identify that records in source have changed from what is currently recorded in the target table. Select **Join** or **Minus**.

Merge Method

Merge Changed Rows only: uses change detection to work out what rows require merging. When the option is set, it enables the **New Row Identification Method**.

New Row Identification Method: method used to identify which records in the source are not recorded or are recorded differently in the target table. Select between **Join** and **Minus**.

If **non identity columns** are used as artificial keys the only new row identification method is **Join**.



Dimension Update procedures usually perform faster when you use the **Join** method for new row identification.

SOURCE TAB

Field	Value
Distinct Data Select	<input type="checkbox"/>
Source Join	FROM [load_product] load_product JOIN [load_product_line] load_product_line ON load_product.product_line_code = load_product_line.product_line_code JOIN [load_product_group] load_product_group ON load_product.product_group_code = load_product_group.product_group_code LEFT OUTER JOIN [load_product_sub_group] load_product_sub_group ON load_product.product_group_code = load_product_sub_group.product_group_code AND load_product.product_sub_group_code = load_product_sub_group.product_sub_group_code
Where Clause	
Group By	

Distinct Data Select: ensures duplicate rows are not added to the Dimension. This is achieved by adding the word DISTINCT to the source select in the update procedure. The default for this field is not set.

Source Join: The From clause, including Source Join information. See example below for Joining multiple source tables.

Where Clause: The Where clause. Use as a filter to extract only the necessary records that fulfill a specified criteria.

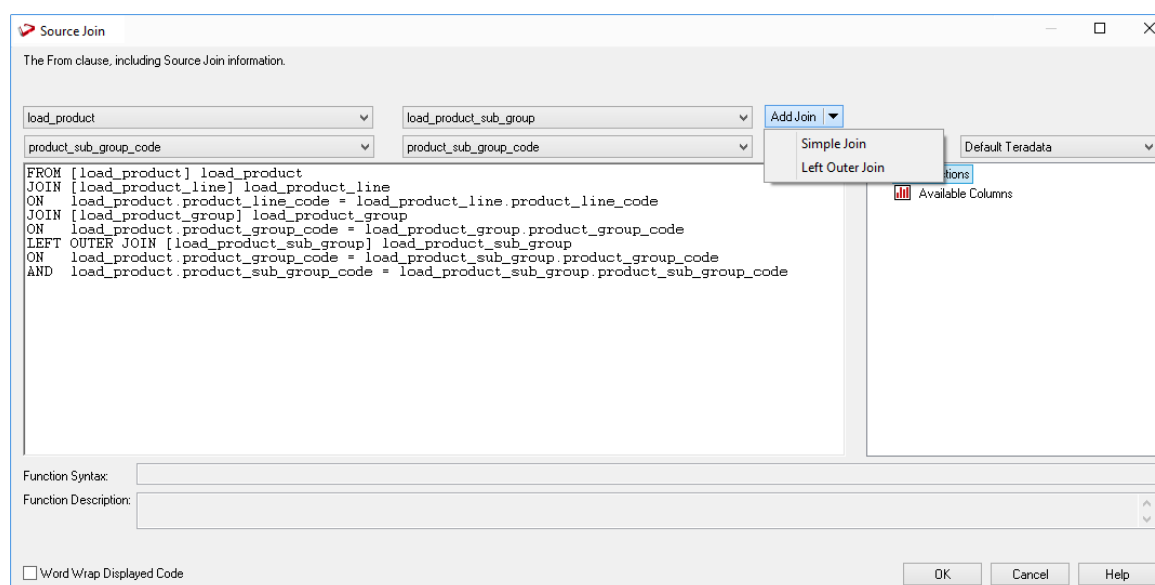
Group By: The Group By clause. Use in collaboration with the SELECT statement to arrange identical data into groups.

Joining multiple source tables

If multiple source tables were used to build the dimension, the tables will need to be joined on the **Source** tab.

- Select the first two tables from the top drop-down list.
- Select one column for each of the two tables from the bottom drop-down lists to effect the join between the selected tables.

In the example below, the `load_product` and `load_prod_subgroup` tables are joined by two columns, `prod_group` and `subgroup`. In this case, two joins are actioned for these two tables so both columns can be selected.



Simple Join

Either a 'Where' or from clause join can be generated. A simple join only returns rows where data is matched in both tables. So for example if table A has 100 rows and table B has a subset of 24 rows. If all the rows in table B can be joined to table A then 24 rows will be returned. The other 76 rows from table A will not be returned.

Outer Join

Either a 'Where' or from clause join can be generated. The ANSI standard 'from clause' join is recommended. The outer join returns all rows in the master table, regardless of whether or not they are found in the second table. Therefore, if the example above was executed with table A as the master table, then 100 rows would be returned. 76 of those rows would have null values for the table B columns. When WhereScape RED builds up a 'Where' clause join, it must place the outer join indicator next to the appropriate column. As this indicator goes on the master, WhereScape RED needs to know which table is master and which subordinate. Select the join column from the master table first.

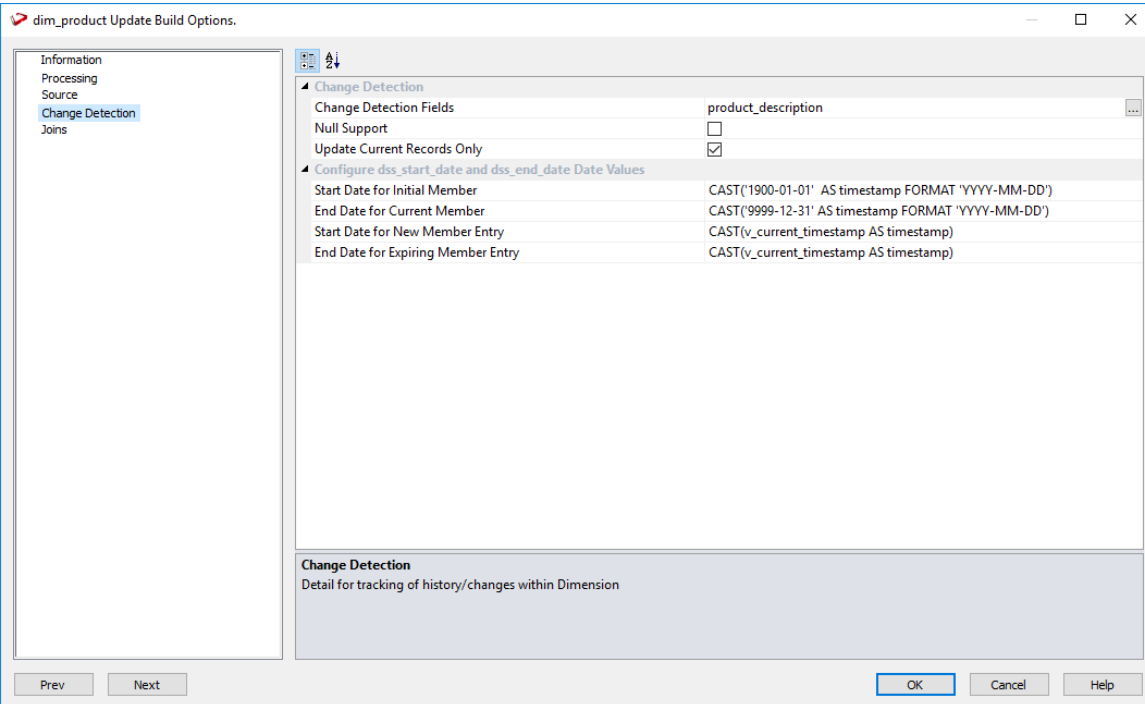
In the example screen above, the table 'load_product' has had its column chosen and the column for the table 'load_prod_subgroup' is currently being chosen. This will result in the 'load_product' table being defined as the master, as per the example statement as shown in the 'Where' clause

edit window above. The results of this example select are that a row will be added containing product information, regardless of whether or not a corresponding prod_subgroup entry exists.

As the join columns are selected, the join statement is built up in the large edit window above. Once all joins have been made, the contents of this window can be changed if the join statement is not correct.

- When you are happy with the join clause click the **OK** button to proceed to the next step. This clause will be either the 'Where' clause or a combined from and 'Where' clause depending on the option chosen.
- This clause can be edited in the procedure that is generated if not correct.
- For Teradata, you have the choice between 'Where' statement joins and ANSI standard joins.

Change Detection Tab



dim_product Update Build Options.

Information
Processing
Source
Change Detection
Joins

Change Detection

Change Detection Fields	product_description
Null Support	<input type="checkbox"/>
Update Current Records Only	<input checked="" type="checkbox"/>

Configure dss_start_date and dss_end_date Date Values

Start Date for Initial Member	CAST('1900-01-01' AS timestamp FORMAT 'YYYY-MM-DD')
End Date for Current Member	CAST('9999-12-31' AS timestamp FORMAT 'YYYY-MM-DD')
Start Date for New Member Entry	CAST(v_current_timestamp AS timestamp)
End Date for Expiring Member Entry	CAST(v_current_timestamp AS timestamp)


Change Detection
Detail for tracking of history/changes within Dimension

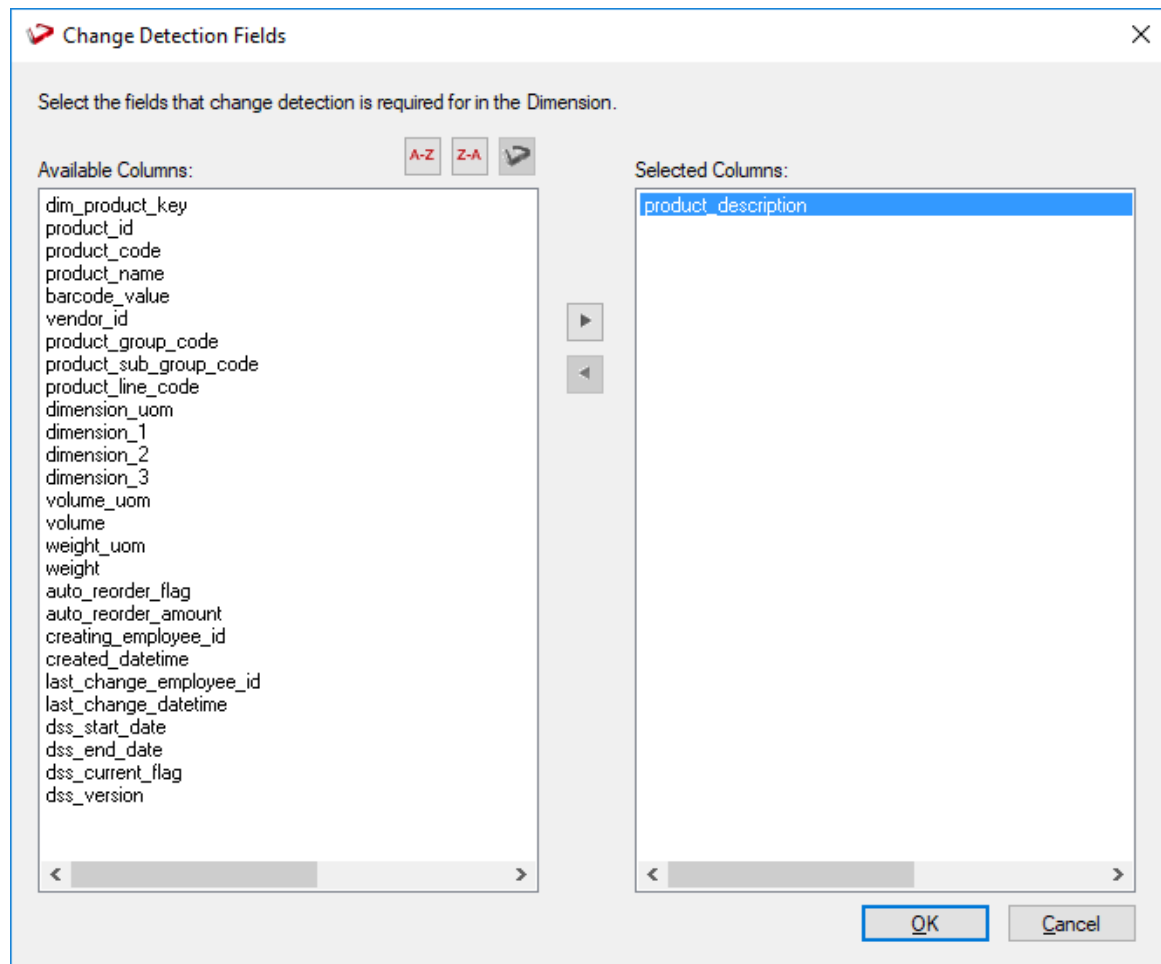
Prev Next OK Cancel Help

Change Detection Fields: if the dimension was defined as a **Changing Dimension** you have to select the change detection fields required for the Dimension on the **Change Detection Tab**. This enables you to select the columns to be managed as Slowly Changing dimension columns.

Note: For custom database targets, you can set the initial default values for the DSS columns that are used in the procedure generation for slowly changing dimensions. Please refer to the section **DSS Columns for Custom Targets** (on page 104) for details. The default values set are populated in the corresponding fields of the **Change Detection** tab.

The advantages and disadvantages of **changing** dimensions are discussed earlier in this chapter, but as a general rule try to minimize their usage. They invariably complicate the processing and end user queries.

- Click the **Change Detection** tab in the Update Build Options dialog.
- Click the ellipsis button  at the rightmost corner of the **Change Detection** field.
- Select the required columns to be managed as slowly changing dimension columns and click **OK** to continue.
- In the example below, the `product_description` is to be managed as a slowly changing column.



Null Support: if this option is set, the change detect column management will cater for Null values in any changing columns. If this is not set and there are Null values in the changing columns there may be errors while running the update procedure. The default for this option is not set (Nulls are not catered for).

Null values are the enemy of a successful data warehouse. They result in unreliable query results and can often lead to a lack of confidence in the data. If a column is considered important enough to be managed as a slowly changing column then it should not normally contain null values. It is often best to ensure that a Null cannot occur by using a `Coalesce()` transformation when loading the column.

Configuring `dss_start_date` and `dss_end_date` for Date Detection Values:

Reset Dates to Initial Values: resets `dss_start_date` and `ds_end_date` Values to original values.

Start Date for Initial Member: the start date for initial member field contains the start date for the first version of a particular business key. The value should be specified in an appropriate form, taking into account the default date format in the databases. The date may need to be explicitly cast to the current data type. The default value provided will usually be cast to the correct database and can be treated as a template. The default for this field is 1 January 1900.

End Date for Current Member: the end date for current member field contains the start date for the current version (the row with a current flag of Y and the maximum version number) of a particular business key. The value should be specified in an appropriate form, taking into account the default date format in the databases. The date may need to be explicitly cast to the current data type. The default value provided will usually be cast to the correct database and can be treated as a template. The default for this field is 31 December 2999.

Start Date for New Member Entry: the start date for new member entry field contains the start date for any subsequent rows added to the history table (not the first row for a particular business key i.e. not version 1). The value should be specified in an appropriate form, taking into account the default date format in the databases. The date may need to be explicitly cast to the current data type. The default value provided will usually be cast to the correct database and can be treated as a template. The default for this field is the current date and time.

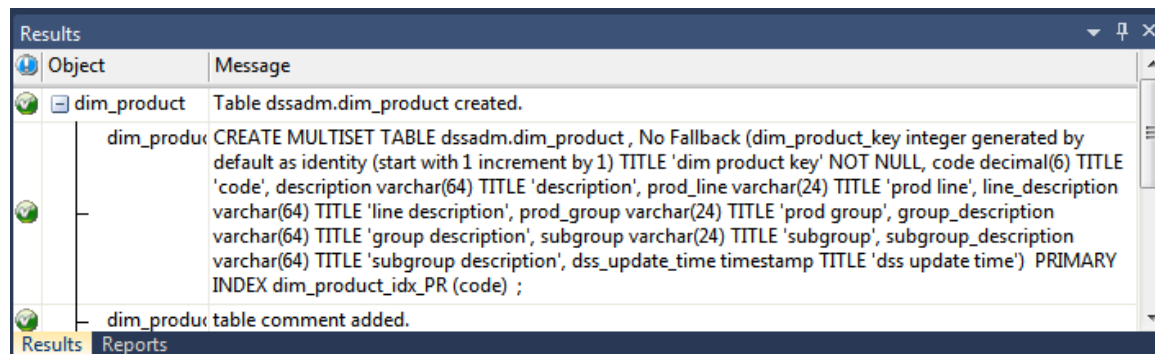
End Date for Expiring Member Entry: the end date for the expiring member entry field contains the end date for any rows updated no longer to no longer be the current row in the history table (i.e. rows that are replaced by a new current row). The value should be specified in an appropriate form, taking into account the default date format in the databases. The date may need to be explicitly cast to the current data type. The default value provided will usually be cast to the correct database and can be treated as a template. The default for this field is the current date and time less an arbitrary small amount (for SQL Server this is 0.00000005 of a day, or about 4 thousandth of a second).

Building and Compiling the Procedure

- When the steps above are completed the procedure is built and compiled automatically. This will display in the **Results** pane.
- If the compile **fails**, an error will be displayed along with the first few lines of error messages. Compile fails typically occur when the physical creation of the table was not done.
- If the compile fails for some other reason the best approach is to use the procedure editor to edit and compile the procedure. The procedure editor will highlight all the errors within the context of the procedure.
- Once the procedure has been successfully compiled it can either be executed interactively or passed to the scheduler.

Indexes

By default a number of indexes will be created to support the dimension. These indexes will be added once the procedure has been built. An example of the type of indexes created is as follows:



This example shows three indexes being created. They are:

- 1 A primary key constraint placed on the artificial key for the dimension.
- 2 A unique index placed on the business key for the dimension.
- 3 A unique index placed on the business key and a slowly changing column from the dimension.

This third index is only created when a **Slowly Changing** dimension is chosen.

Additional indexes can be added, or these indexes changed. See the chapter on indexes for further details.

DIMENSION ARTIFICIAL KEYS

The artificial (surrogate) key for a dimension is set via an identity column. This artificial key normally, and by default, starts at one and progresses as far as is required.

A WhereScape standard for the creation of special rows in the dimension is as follows:

Key value	Usage
1 upwards	The normal dimension artificial keys are numbered from 1 upwards, with a new number assigned for each distinct dimension record.
0	Used as a join to the dimension when no valid join existed. It is the normal convention in the WhereScape generated code that any dimension business key that either does not exist or does not match is assigned to key 0.
-1 through -9	Used for special cases. The most common being where a dimension is not appropriate for the record. For example we may have a sales system that has a promotion dimension. Not all sales have promotions. In this situation it is best to create a specific record in the dimension that indicates that a fact table record does not have a promotion. The stage table procedure would be modified to assign such records to this specific key. A new key is used rather than 0 as we want to distinguish between records that are invalid and not appropriate.
-10 backward	Pseudo records. In many cases we have to deal with different granularities in our fact data. For example, we may have a fact table that contains actual sales at a product SKU level and budget information at a product group level. The product dimension only contains SKU based information. To be able to map the budget records to the dimension, we need to create these pseudo keys that relate to product groups. The values -10 and backwards are normally used for such keys. A template called 'Pseudo' is shipped with WhereScape RED to illustrate the generation of these pseudo records in the dimension table.

Surrogate keys for a Dimension set via a non identity column:

Normal, Slowly Changing and **Date Ranged Dimension** Tables can have non identity columns as surrogate keys.

The generation of the update procedure automatically adds logic to the code which will associate a sequential number to the artificial key of the dimension when a new row is inserted into the Dimension table.

The order of these sequential numbers is determined by the business key of the source table. The value of the first newly inserted artificial key will be the value of the highest artificial key in the dimension table plus 1.

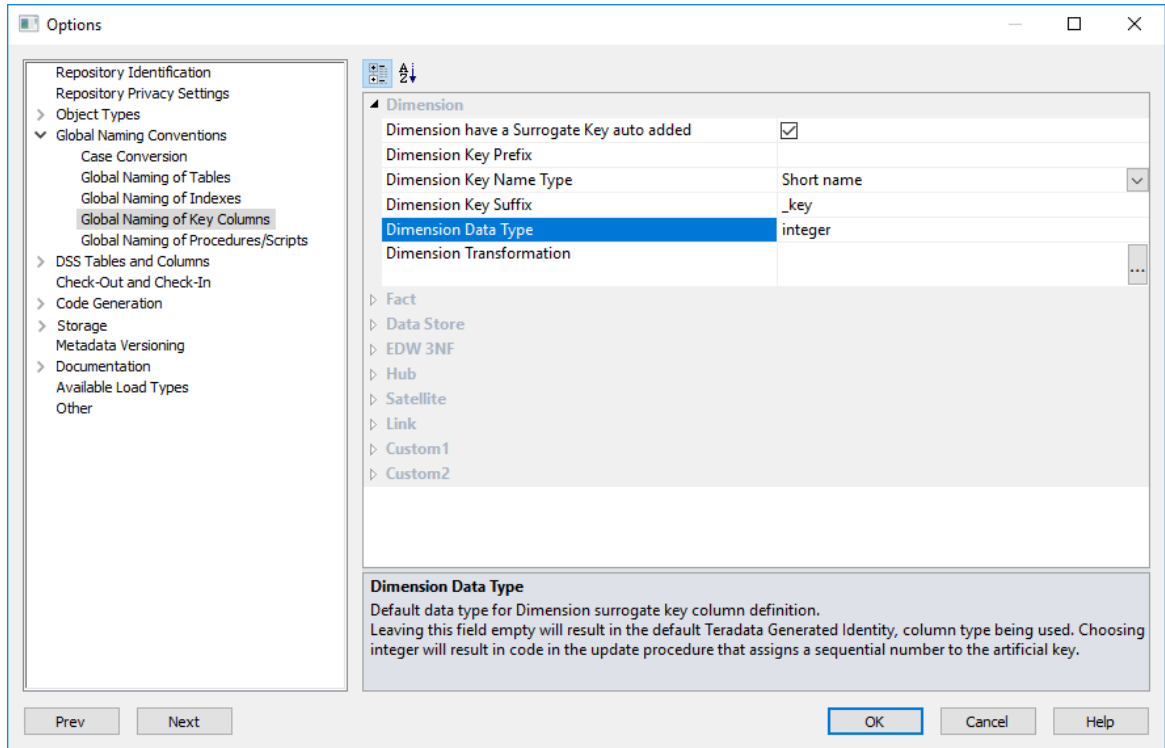
This automatically generated logic can be overwritten by defining a user specific logic on the **Dimension Transformation** field on the **Tools>Options** menu or in the transformation column of the artificial key.

To have a Dimension with a non identity column as a surrogate key, you can **set the Dimension Data Type to integer** in the **Tools>Options** menu.

The old logic for dimensions can be retained if an identity column is chosen as surrogate key.

To allow for non identity surrogate keys on Dimensions:

- 1 Go to **Tools>Options>Global Naming Conventions>Global Name of Key Columns**.
- 2 Set the **Dimension Data Type** to be **integer** and click **OK**.
- 3 If your tables have been created previously, you have to **Recreate** the tables after you set this option in the **Tools** menu.



DIMENSION COLUMN PROPERTIES

Each dimension column has a set of associated properties. The definition of each property is described below:



TIP: If a database table's definition is changed in the metadata then the table will need to be altered in the database. Use the **Validate>Validate Table Create Status** to compare metadata definitions to physical database tables. The option also provides the ability to alter the database table, through a pop-up menu option from the validated table name. See the example below.

A sample **Properties** screen is as follows:

The two special update keys allow you to update the column and step either forward or backward to the next column's properties. **ALT-Left Arrow** and **ALT-Right Arrow** can also be used instead of the two special update keys.

Table Name

Database-compliant name of the table that contains the column. [Read-only].

Column Name

Database-compliant name of the column. A good practice is to only use alphanumeric characters, and the underscore character. Changing this field alters the table's definition.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Title

Name that the business uses to refer to the column. It does not affect the physical table definition, but rather provides input to the documentation and to the view **ws_admin_v_dim_col** which can be used to assist in the population of a end user tool's end user layer. As such it is a free form entry and any characters are valid.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Description

This field contains the description for the column. It might contain information on where and how the column was acquired. For example, if the column is sourced from multiple tables or is a composite or derived column then this definition would normally describe the process used to populate the column. This field is used in the documentation and is available via the view **ws_admin_v_dim_col**. This field is also stored as a comment against the column in the database.

Column Order

Numeric value that controls the relative order of columns in the database create statement. The lowest numbered column will appear first in the table. Although this affects the physical table definition no action will be taken unless the table is re-created in the database. The columns can be re-numbered based on the existing order by choosing the **Respace order number** pop-up menu option when positioned over any column in the table. This action will number the columns in increments of 10 starting at 10. In addition to a simple change of the order field, the column order can be changed by first displaying the columns in the middle pane and then using drag and drop to move the columns around. This drag and drop process will automatically renumber the columns as required.

Data Type

Database-compliant data type that must be valid for the target database. Typical Teradata databases often have integer, numeric(), varchar(), char(), date and timestamp data types. See the database documentation for a description of the data types available. Changing this field alters the table's definition.

Null Values Allowed

Determines whether the table column can hold NULL values or whether a value is always mandatory.

Default Value

Initial value that is assigned to the column when a row is inserted into the table but no value is specified for the column.

Character Set

Database-compliant table column character-set used for storage. Select Latin or Unicode.

Format

Database-compliant table column format. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. As such it is a free form entry and any characters are valid. Typically format masks are only used on numeric fields. Example: #,###0.00. It is not worth the effort of populating this field unless it can be utilized by the end user tools in use.

Character Comparison/Sorting

Determines how the column character values are treated for comparison and sorting operations. Choose from: case specific, not case specific, uppercase case specific or uppercase not case specific.

Compress

Indicates whether the table column values are compressed when stored.

Compress/Compress Value

Optional list of values to be compressed. By default, only NULL is compressed if no list of values is specified.

Numeric

Indicates whether the table column holds values that are numeric. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Additive

Indicates whether the table column holds values that are additive. This implies that the column can be summed when performing data grouping in a query. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Attribute

Indicates whether the table column holds values that are descriptive, and/or are used for grouping/summing. An attribute is defined as a column that is non factual. For example we may have an order number, or a invoice number stored in the fact table. Such columns are considered attributes, rather than facts. This checkbox is therefore normally only relevant for fact tables. This checkbox does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tools end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

End User Layer display

Indicates whether the table column is available/visible to end users. If set the documentation will include the column in the glossary and in the user documentation. It is also used to decide what columns appear in the view **ws_admin_v_dim_col**. Typically columns such as the artificial key would not be enabled for end user display.

Business Key

Indicates whether the column is part of the BUSINESS key, which is defined during the update procedure build. [Normally maintained automatically]. Multiple columns can form the primary business key. This indicator is set and cleared by WhereScape RED during the dimension update procedure generation process. This checkbox should not normally be altered.

Key Type

Key type that is assigned and used when generating the table's update procedure and indexes. [Normally maintained automatically]. It can be altered here, but this should only be done if the consequences are fully understood and tested. The supported values are:

Key type	Meaning
0	The artificial key. Set when the key is added during drag and drop table generation.
1	Component of all business keys. Indicates that this column is used as part of any business key. For example: By default the dss_source_system_key is added to every dimension table. It is considered part of any lookup on that table and has the key type set to 1. Set when the column is added during drag and drop table generation.
2	Indicates that this column is a dimensional join. Used on fact tables to indicate the dimension keys. Results in an index being built for the column (Bitmap in Oracle). Set during the update procedure generation for a fact table, based on information from the staging table.
3	Slowly changing column indicator. Used on dimension tables to indicate that the column is being managed as a slowly changing column within the context of a slowly changing dimension. Set when a column is identified during the dimension update procedure generation.
4	Previous value column indicator. Used on dimension tables to indicate that the column is being managed as a previous value column. The source column identifies the parent column. Set during the dimension creation.
5	Start date of a date ranged dimension. Used on dimension tables to indicate that the column is defined as the starting date for a source system date ranged dimension. Forms part of the business key. Set during the dimension creation.
6	End date of a date ranged dimension. Used on dimension tables to indicate that the column is defined as the ending date for a source system date ranged dimension. Forms part of the business key. Set during the dimension creation.
A	Indicates that the column is part of the primary business key. Set whenever a business key is defined as part of an update procedure generation.
B-Z	Indicates that the column is part of a secondary business key. Only used during index generation and not normally set.

Zero Key Value

Determines the value populated for the column in the "Invalid Join" or "Unknown" record. By default, NULL is used when a value is not specified.

Source Table

Identifies the source table where the column's data comes from. This source table is normally a load table within the data warehouse. If the column was sourced from multiple tables, then the normal practice is to record one of the tables in this field and a comment listing all of the other tables in the **Source Strategy** field. This field is used when generating a procedure to update the dimension. It is also used in the track back diagrams and in the documentation.

Source Column

Identifies the source column where the column's data comes from. Such a column is normally a load table column, which in turn may have been a transformation or the combination of multiple columns. For previous value managed columns the source column is the column in the table whose previous value is being recorded.

Transformation

See *Dimension Column Transformations* (on page 325). [Read-only].

Join

Indicates whether the table column is used in a table join. Normally this is maintained automatically but can be optionally changed to override the default logic used in the generated update procedure. The default for this option is not set.

Changing a Column Name

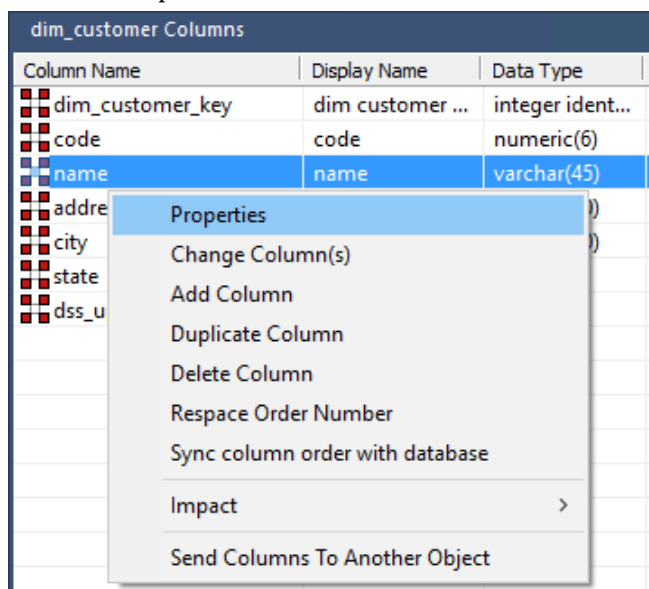
If the **Column name** or **Data type** is changed for a column then the metadata will differ from the table as recorded in the database.

- Use the **Validate/Validate Table Create Status** menu option or the right-click menu to compare the metadata to the table in the database.
- A right-click menu option of **Alter table** is available when positioned on the table name after the validate has completed. This option will alter the database table to match the metadata definition.

For example: Analysis Services does not like **name** as a column name.

For dim_customer it will therefore be necessary to change the column name from **name** to **cname**.

- 1 Click on the **dim_customer** object in the left pane to display the dim_customer columns in the middle pane.
- 2 When positioned on the column **name** in the middle pane, right-click and select **Properties** from the drop-down menu.



- 3 Change the column name from **name** to **cname** as shown below. Click **OK** to leave the Properties page.

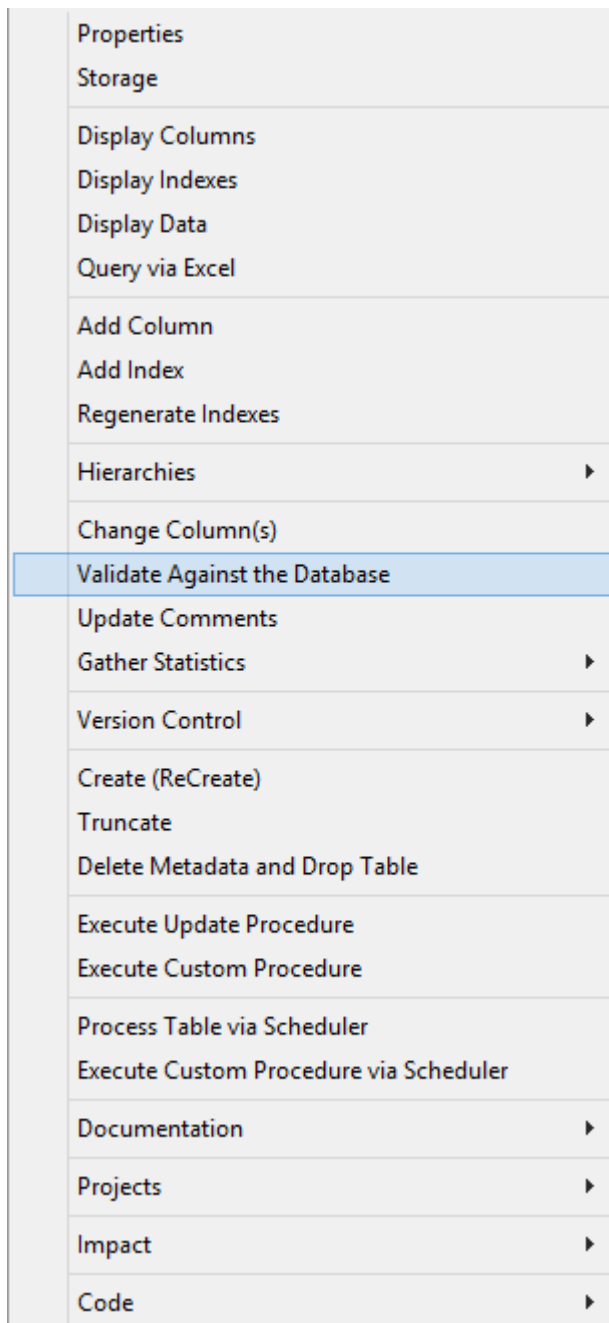
The screenshot shows the 'Dimension Column dim_customer.name' Properties dialog box. The 'Column Name' field is highlighted in blue and contains the text 'cname'. The dialog is organized into several sections: General, Physical Definition, Meta Definition, Code Generation, and Source Details. The 'Column Name' field is located in the General section. Below the main table, there is a 'Column Name' section with a description: 'Database-compliant name of the column. Dialog Opening Value: name'. At the bottom of the dialog, there are five buttons: '< Update', 'Update >', 'OK', 'Cancel', and 'Help'. The 'OK' button is highlighted with a blue border.

Section	Property	Value
General	Table Name	dim_customer
	Column Name	cname
	Column Title	cname
	Column Description	The full name of the customer. Forms a hierarchy with city and state.
Physical Definition	Column Order	30
	Data Type	varchar(45)
	Null Values Allowed	<input checked="" type="checkbox"/>
	Default Value	
	Character Set	
	Format	
	Character Comparison/Sorting	
Meta Definition	Compress	<input type="checkbox"/>
	Numeric	<input type="checkbox"/>
	Additive	<input type="checkbox"/>
	Attribute	<input checked="" type="checkbox"/>
	End User Layer Display	<input checked="" type="checkbox"/>
	Business Key	<input type="checkbox"/>
Code Generation	Artificial Key	<input type="checkbox"/>
	Key Type (0,A,B,C,...)	
Source Details	Zero Key Value	Unknown

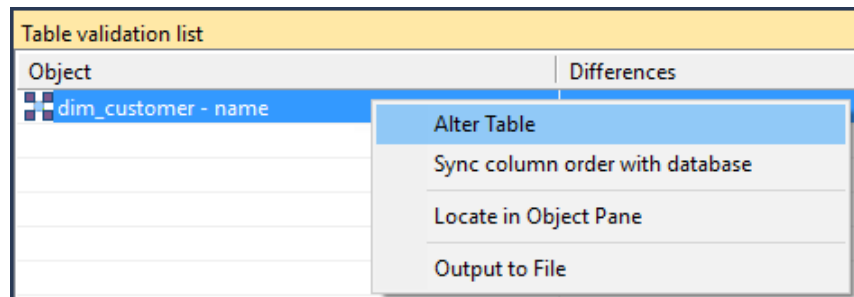
Column Name
Database-compliant name of the column.
Dialog Opening Value: name

< Update Update > **OK** Cancel Help

- 4 Right-click on the **dim_customer** object in the left pane and select **Validate against Database**.



- 5 The results in the middle pane will show that the metadata has changed to **cname** while the column name in the database is still **name**.
- 6 Right-click on `dim_customer` in the middle pane and select **Alter table** from the drop-down list.



- 7 A warning appears, displaying the table and column name to be altered. Select **Alter Table**.
- 8 A dialog appears confirming that `dim_customer` has been altered. Click **OK**.

DIMENSION COLUMN TRANSFORMATIONS

Each dimension table column can have a transformation associated with it. The transformation will be included in the generated procedure and will be executed as part of the procedure update. The transformation must therefore be a valid SQL construct that can be included in a **Select** statement. For example, we could have a transformation of 'load_order_line.qty * 0.125' to calculate a tax column of 12.5%.

Click the **Transformation** tab to enter a transformation.

The transformation screen is as follows:

Note: Transformations are only put into effect when the procedure is re-generated.

Microsoft Analysis Services 2005+ Tabular Mode Tables: For Tabular Mode table column transformations, **Default DAX** is the only applicable Function Set for **after load** transformations.

See *Transformations* (on page 593) for more details.

DIMENSION HIERARCHIES

The various hierarchies associated with a dimension can be recorded in the WhereScape RED metadata. These hierarchies are often not used in any form, except to provide documentary completeness. They can however be used in conjunction with the WhereScape hierarchy maintenance utility to allow the maintenance of hierarchy elements externally to a production source system.

When used in conjunction with the hierarchy maintenance utility, these dimension hierarchies add a powerful method of enriching the analysis capabilities of the data warehouse. For example, we may have a source system that has a dimension called sales_person. This dimension has no information apart from employee_code, name and title. We could add additional columns of sales_manager, territory, state and region to this dimension. A hierarchy could then be formed from the salesperson name, sales_manager, territory, state and region. The hierarchy maintenance utility allows the maintenance of this hierarchy externally to the data warehouse. This external hierarchy can then become a source system to enrich the data in the warehouse.

Two areas will be covered. Firstly the creation of hierarchies using WhereScape RED, and secondly the process required to setup and use externally maintained hierarchies as source systems to the data warehouse.

ADDING A DIMENSION HIERARCHY

Any number of hierarchies can be created against a dimension. There is no restriction on the form of the hierarchy.

- 1 To add a new hierarchy, position on the dimension table in the left hand pane and using the right-click menu select **Hierarchies>Add Hierarchy**.

The following dialog appears:

Move columns from the column list into the hierarchy. The hierarchy is a top down list. For example a date hierarchy may be year, month, day. Year will be the first column shown.

Column Name
dim_customer_key
code
name
address
city
state
dss_update_time

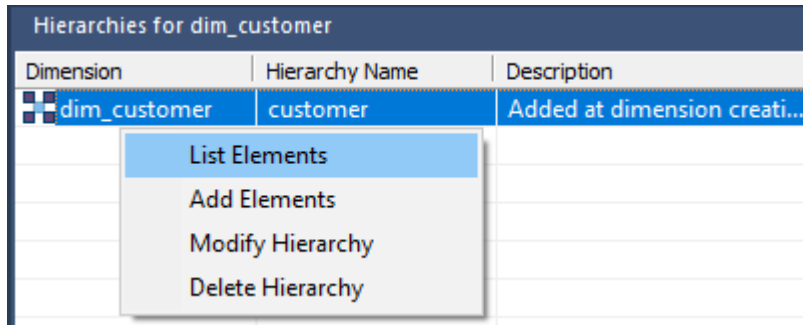
Hierarchy: Levels

- 2 Enter a meaningful name for the hierarchy.
- 3 Enter a meaningful description for the hierarchy. This description is carried through into the Hierarchy Description field of any OLAP Dimensions that are built from the original Dimension object.

Note: The description text is automatically set to "Added at dimension creation for cube support" but this can be edited to match the user's intended description.

The hierarchy is built with the highest level at the top; for example a customer dimension may have **state** at the highest level, then **city**, then **address** and finally **code** at the lowest level.

- To enter the hierarchy elements, select them in the required order, from the left pane and click the right arrow (>) to add them to the right column.
- Once all the hierarchy elements have been added, click **OK**.
- A hierarchy and its elements can be edited by listing the hierarchies associated with a dimension and using the right-click menu options available in the middle pane.



The screenshot shows a table titled "Hierarchies for dim_customer" with three columns: "Dimension", "Hierarchy Name", and "Description". The first row is highlighted in blue and contains the values "dim_customer", "customer", and "Added at dimension creati...". A context menu is open over the "customer" cell, listing four options: "List Elements", "Add Elements", "Modify Hierarchy", and "Delete Hierarchy".

Dimension	Hierarchy Name	Description
dim_customer	customer	Added at dimension creati...

Copying Dimension Hierarchies from Source

Hierarchies are automatically copied from a source table when the source table is dragged into the middle pane to create a new Dimension.

To copy the hierarchies from the source objects manually, right-click on a dimension in the Object Pane and select **Hierarchies>Copy Hierarchies from Source**. This feature is useful when the source for a Dimension has been updated to contain new hierarchies.

USING A MAINTAINED HIERARCHY

Once a hierarchy has been added to a dimension, it can be maintained externally by the hierarchy maintenance utility if the **Maintained** attribute is set.

The process for maintaining this hierarchy externally and using it as a source system to the data warehouse is as follows.

- 1 Grant the user that is to undertake the maintenance select access to the following tables: ws_dim_tab, ws_dim_col, ws_hie_header, ws_hie_link, ws_meta_tables.
- 2 Grant the user that is to undertake the maintenance select access to the dimension that has the hierarchy.
- 3 Using the hierarchy maintenance utility (see online documentation) log on to the maintenance area and create and populate/edit the maintenance table.
- 4 In WhereScape RED, create a connection to the hierarchy maintenance schema/database.
- 5 Browse the hierarchy connection.
- 6 Using drag and drop create a load table from the hierarchy maintenance table.
- 7 Edit the columns in the dimension that are to be sourced from the maintenance table and change their source table and column to that of the load table and columns created in step 6.
- 8 Generate a new update procedure for the Dimension and either do a lookup of the maintenance table or a join based on the business key.
- 9 Run the update procedure.

SNOWFLAKE

Snowflake schemas normalize dimensions to eliminate redundancy. That is, the dimension data has been grouped into multiple tables instead of one large table. For example, a product dimension table in a star schema might be normalized into a products table, a product_category table, and a product_manufacturer table in a snowflake schema. While this saves space, it increases the number of dimension tables and requires more foreign key joins. The result is more complex queries and reduced query performance.

CREATING A SNOWFLAKE

A snowflake dimensional structure is supported by WhereScape RED. A snowflake can be created for EDW 3NF or partially EDW 3NF dimension tables. It is created by including the surrogate key of the parent dimension in the child dimension. In the example below, the dim_state table represents the parent dimension. The column dim_state_key is added to the child dimension dim_customer. Any fact tables that include the dim_customer dimension will inherently have a link to the dim_state dimension.

The process for creating a snowflake is as follows:

- 1 Build both dimensions (see previous sections).
- 2 Expand **dimensions** in the left pane.
- 3 Click on the child dimension table in the left pane to display its columns in the middle pane.
- 4 Browse the data warehouse connection in the right pane.
- 5 Expand the parent dimension table in the right pane.
- 6 Drag the surrogate key of the parent dimension table from the right pane to the child dimension's column list in the middle pane.
- 7 Create/Recreate the child dimension.
- 8 Rebuild the child dimension's update procedure.
- 9 A dialog will now appear asking for the business key column(s) in the child dimension that matches the business key for the parent dimension:

- 10 Add the business key column(s) and click **OK**.

DIMENSION LANGUAGE MAPPING

The Dimension Properties screen has a tab called **Language Mapping**.

- 1 Select the language from the drop-down list and then enter the translations for the **Business Display Name** and the **Description** in the chosen language.
- 2 The translations for these fields can then be pushed through into OLAP cubes.

The screenshot shows a dialog box titled "Dimension dim_customer" with a close button (X) in the top right corner. On the left is a vertical menu with the following items: Properties, Storage, Override Create DDL, Language Mapping (highlighted in blue), Extended Properties, Purpose, Concept, Grain, Examples, Usage, and Notes. The main area of the dialog is divided into sections:

- Language :** A dropdown menu showing "French".
- Language Settings:** A section containing:
 - Business Display Name:** A text box containing "dim_customer".
 - A text box containing "Translated Value of Business Display Name".
- Description:** A text area containing "This is the master Customer List for the company." with up and down arrow icons on the right side.
- Below the description, another text box containing "Translated Value of Business Display Name" with up and down arrow icons on the right side.

At the bottom right of the dialog are three buttons: "OK", "Cancel", and "Help".

CHAPTER 11

STAGING

Stage tables are used to transform the data to a star schema or third normal form model. A stage table can be a fact or EDW 3NF table that only contains change data or a work table. In star schema data warehouses, the stage table brings all the dimensional joins together in preparation for publishing into the fact table.

A stage table is built from the Data Warehouse connection. Unless you are retrofitting an existing system, stage tables are typically built from one or more load tables. They can utilize the surrogate keys from a number of dimension tables.

The normal steps for creating a stage table are defined below and are covered in this chapter. As the stage table is essentially a subset of the fact table, the design and creation of the stage table is essentially the design and creation of the model table. The steps are:

- 1 Identify the source transactional data that will ultimately constitute the model table. If the data is sourced from multiple tables ascertain if a join between the source tables is possible, or if a series of passes will be required to populate the stage table. If the latter option is chosen then bespoke code is needed.
- 2 Using the 'drag and drop' functionality drag the table with the lowest granular data into a stage target. See *Building the Stage Table* (on page 334).
- 3 Add columns from other source tables. See *Building the Stage Table* (on page 334).
- 4 Create the stage table in the database. See *Building the Stage Table* (on page 334).
- 5 Build the update procedure. See *Generating the Staging Update Procedure* (on page 337).

NOTE: If you are building a Data Vault system, a **Stage** table with sub type of **Data Vault Stage** can be created to generate hash keys that are used in building Data Vault objects (Hub, Link or Satellite tables). Refer to *Data Vaults* (on page 413) for details.

IN THIS CHAPTER

Building the Stage Table	334
Generating the Staging Update Procedure	337
Stage Table Custom Procedure	345
Stage Table Column Properties	345
Stage Table Column Transformations	351
Permanent Stage Tables	352
Generating the Permanent Staging Update Procedure.....	352
Set Merge Procedure.....	358

BUILDING THE STAGE TABLE

Building the stage table is potentially the most challenging part of the overall task of building a data warehouse analysis area. Most of the effort required is in the design phase, in terms of knowing what data needs to come into the model table that will be ultimately built. This section assumes that the decision as to what to include has been made.

Multiple data sources

A stage table typically contains the change data for a detail fact table. As such it normally maps to a specific function within the business and in many cases relates back to one main OLTP table. In many cases however it may be necessary to combine information from a number of tables. One of the decisions required is whether or not it is practical or even possible to join the data from the different source tables.

We may have to build a model table containing data from invoice headers, invoice lines, order headers and order lines source tables. There are three basic options:

- 1 Join all four tables in one large join in our staging table.
- 2 Update the staging table in two passes. One pass updating the order information and one pass updating the invoice information.
- 3 Generate two stage tables, one for order and one for invoice. Use these two staging tables to update the one sales_detail model table.

Although all three options are viable and a normal situation in the WhereScape RED environment, options (2) and (3) will require specific coding and modifications to the generated procedures from the outset (or the use of a custom procedure). Given the example provided option (2) would be the normal approach, although in some cases option (3) would be valid.

Drag and Drop

The best approach in creating a stage table is to choose the source table that contains the most fields that we will be using and drag this table into the stage target. Then drag specific columns from the other source tables until we have all the source data that is required.

The process for defining the metadata is as follows:

- 1 Double click on the **Stage Table** object group in the left pane. This will result in all existing stage tables being displayed in the middle pane. This also sets the middle pane as a **stage drop target**.
- 2 Browse the DataWarehouse connection to display our load tables in the right pane. This is achieved via the **Browse>Source Tables** menu option and then choosing the **DataWarehouse connection**.
- 3 Drag the **primary load table** (e.g. the one with the most columns, or the lowest data granularity) from the right pane and drop it in the middle pane. A dialog appears to create the new staging object. Leave the object type as **Stage Table** and change the name to reflect what is being done. For example in the tutorial the load_order_line table is dropped and a stage table called stage_sales_detail defined.
- 4 Once a valid name is entered the properties for the new stage table are displayed. Normally these would be left unchanged except perhaps for storage settings.
- 5 Once the Properties dialog is closed the columns for the new stage table are displayed in the middle pane. This middle pane is now considered a *drop target for this specific stage table*.

Any additional columns or tables dropped into the middle pane are considered additions to this stage table definition. Any columns that are not required can be deleted at this stage.

- 6 Drag and drop **additional columns** from other source tables if appropriate. In the tutorial we would now drag the `customer_code`, `order_date` and `ship_date` from the `load_order_header` table.
- 7 If using surrogate keys, then drag in the **model artificial key** from each model table that is to be joined to the stage table. We can only join a model if a business key exists amongst the stage table columns or it is possible to derive that business key in some way from the columns or other model tables.

Note: If a column is being used to join information from two or more source tables, that column must only appear once in the stage table. It is irrelevant which table is used to create the column in the new stage table.

Once completed our list of columns for the stage table should look something like the list below. Note the source table for each column.

Column Name	Display Name	Data Type	Source Table	Source Column
order_number	order number	decimal(12)	load_order_header	order_number
order_date	order date	timestamp	load_order_header	order_date
customer_code	customer code	integer	load_order_header	customer_code
ship_date	ship date	date	load_order_header	ship_date
model_customer_key	model custom...	integer	model_customer	model_customer_...
model_ship_date_key	model ship dat...	integer	model_ship_date	model_ship_date_...
dss_update_time	dss update time	timestamp		dss_update_time

The source table (src table) reflects where each column was dragged from. In the example above, the bulk of the columns came from the `load_order_header` table. Each model artificial key was dragged from its appropriate table. The final column `'dss_update_time'` was generated by RED and has no source.

Create the table

Once the stage table has been defined in the metadata we need to physically create the table in the database.

- 1 Right-click on the stage table name and selecting **Create (ReCreate)** from the pop up menu.
- 2 A results box will appear to show the results of the creation. The following example shows a successful creation.

Results	
Object	Message
stage_customer	Create of Stage dssdemo.stage_customer completed successfully.
stage_customer	CREATE MULTISET TABLE dssdemo.stage_customer , NO FALLBACK (code decimal(6) TITLE 'code', name varchar(45) TITLE 'name' , address varchar(40) TITLE 'address', city varchar(30) TITLE 'city' , state varchar(2) TITLE 'state' , dss_update_time timestamp TITLE 'dss update time') PRIMARY INDEX stage_customer_idx_PR (code) ;
stage_customer	COMMENT ON COLUMN dssdemo.stage_customer.dss_update_time IS 'Date and time the row was updated in the data warehouse.';

The contents of this dialog are a message to the effect that the table was created followed by a copy of the actual database create statement, and if defined the results of any index creates. For the initial create no indexes will be defined.

If the table was not created then ascertain and fix the problem. A common problem is a 'Duplicate column' where a column has been accidentally added twice.

The best way of finding such a column is to double click on the list heading 'Col name'. This will sort the column names into alphabetic order.

Another double click on the heading will sort the columns back into their create order. Column ordering can be changed by altering the column order value against a column's properties.



TIP: Double clicking on the heading of a column in a list sorts the list into alphabetical order based on the column chosen.

The next section covers *Generating the Staging Update Procedure* (on page 337).

GENERATING THE STAGING UPDATE PROCEDURE

Once a stage table has been defined in the metadata and created in the database, an update procedure can be generated to handle the joining of any tables and the lookup of the model table artificial keys.

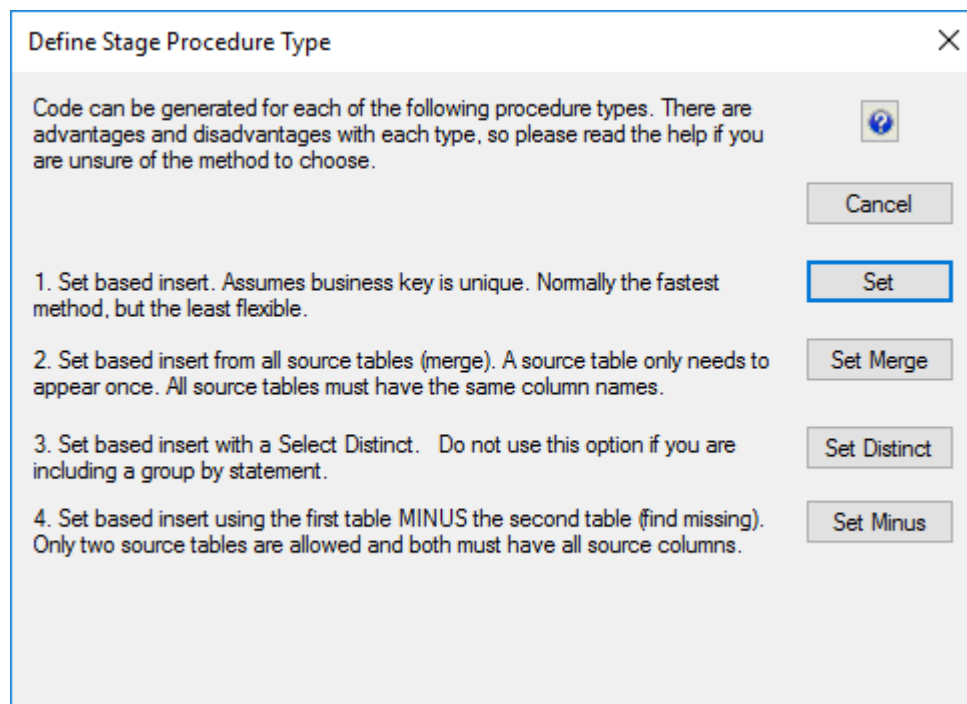
Note: You can also generate an update procedure via a template, refer to *Rebuilding Update Procedures* (on page 199) for details.

GENERATING A PROCEDURE

- 1 To generate a procedure, right-click on the stage table in the left pane and select **Properties**.
- 2 From the **Update Procedure** drop-down list, select (**Build Procedure...**).
- 3 Click **OK** to update the properties and start the process of generating the new procedure.
- 4 A series of prompts are displayed during the procedure generation to join the tables and link in the model tables.

PROCEDURE TYPE

The first dialog box asks for the type of procedure that is to be generated:



A number of different types of procedure generation are available. Each type is discussed below.

The check box at the bottom of the dialog box is only visible if advanced procedure building features are enabled. The check box enables editing of the 'Where' clause when no table joining is being performed, and hence the 'Where' clause would not be exposed.

Set based procedure

A set based procedure performs one SQL statement to join all the model and source tables together and then insert this data into the stage table. This is normally the fastest method of building the stage table. Caution must be taken in regards to NULL values in the business keys that are used to make the table joins. The generated code deliberately does not handle such null values.

Set Merge Procedure

This option is to allow the merging of two or more identical tables. The tables to be merged must have exactly the same number of columns and column names. If necessary, additional blank columns could be added to one or other of the tables to make them identical. To use this procedure, you must simply have the tables to be merged mentioned at least once in the **Source Table** field of a columns properties. For more details, see *Set Merge Procedure* (on page 358).

Set Distinct

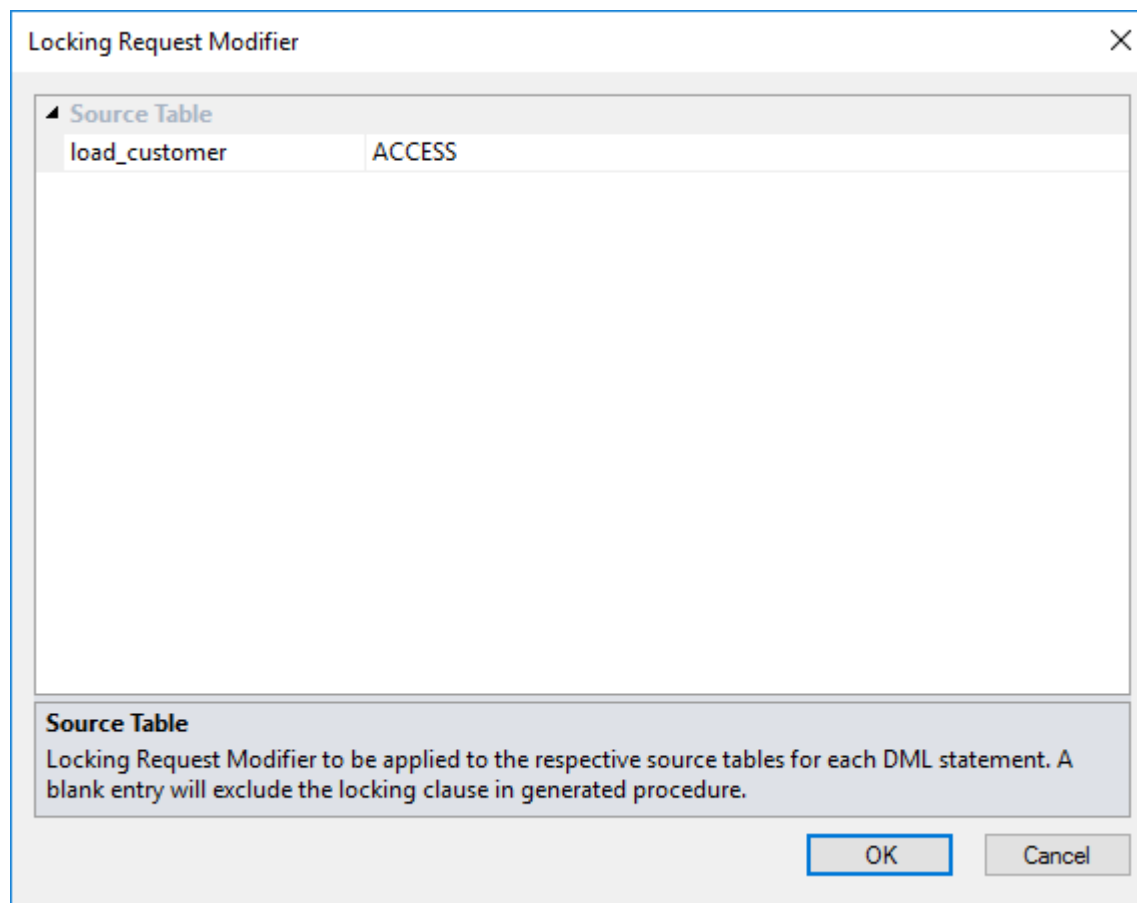
Essentially the same as Set, except for the DISTINCT key word being added to the SELECT statement. This option therefore removes duplicate rows.

Set Minus

The **Set Minus** option can be used to determine change data or for programmatic referential integrity checking. This option works in a similar way to Set Merge. It generates SQL code in this form: `SELECT ... FROM source_table1 {where} MINUS SELECT ... FROM source_table2 {where}`. It requires exactly two source tables to be specified. All source columns must exist in both source tables.

LOCKING REQUEST MODIFIER

Source Table: Specify a locking request modifier to be applied to each source table during generated update procedures. By default, this is set to 'ACCESS' which locks each row being accessed, a blank entry will result in no locking clause in the generated procedure.



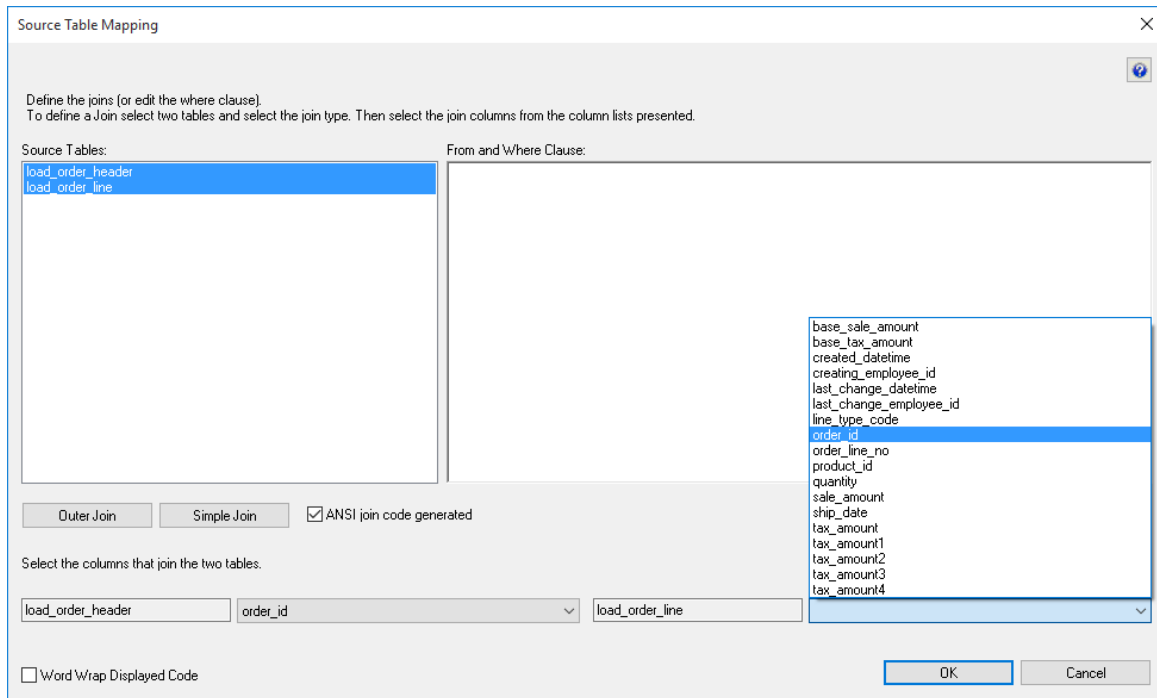
SOURCE TABLE MAPPING

If multiple source tables are present, then the definition of the joins between the various tables is required. Note that at this stage we are not joining the model tables. These joins only relate to the source tables.

The joining of the tables will provide part of the construct of the set based update in the update procedure.

Only two tables may be joined at a time. To join two tables select the two tables in the left box and click either the outer join or simple join button. Column lists for both tables will appear at the bottom of the dialog box. Select the column (or one of the columns) that allows the two tables to be joined. If an outer join is being used, the column for the master table must be chosen first. If there are multiple columns joining two tables then this action must be repeated for each column. Continue to perform all joins between all tables. The example below only has two tables with one join column so is a relatively simple case. An additional option is available to allow either an

ANSI standard join or a 'Where' clause based join. The ANSI standard join should be chosen in most situations. See the example screen in the following section.



Simple Join

A simple join joins the two tables, and only returns rows where data is matched in both tables. So for example if table A has 100 rows and table B has a subset of 24 rows. If all the rows in table B can be joined to table A then 24 rows will be returned. The other 76 rows from table A will not be returned.

Outer Join

An outer join joins the two tables, and returns all rows in the master table regardless of whether or not they are found in the second table. Therefore, if the example above was executed with table A as the master table, then 100 rows would be returned. 76 of those rows would have null values for the table B columns. In the example screen above, the table 'load_order_line' has had its column chosen and the column for the table 'load_order_header' is currently being chosen. This will result in the statement as shown in the 'Where' clause edit window. The results of this select are that a row will be added containing order_line information regardless of whether or not an order_header exists.

As the join columns are selected, the 'Where' statement is built up in the large edit window on the right side. Once all joins have been made, the contents of this window can be changed if the join statement is not correct.

Once satisfied with the 'Where' statement, click the **OK** button to proceed to the next step. As indicated in its description, this statement is the 'Where' clause that will be applied to the select statement of the cursor to allow the joining of the various source tables. It can of course be edited in the procedure that is generated if not correct.

You have the choice between 'Where' statement joins and ANSI standard joins.

Note: 'Where' joins are not available if using outer joins in Teradata.

The example below shows the result of an ANSI standard join which takes place in the 'From' statement.

Source Table Mapping

Define the joins (or edit the where clause).
To define a Join select two tables and select the join type. Then select the join columns from the column lists presented.

Source Tables:

- load_order_header
- load_address
- load_order_line

From and Where Clause:

```
FROM [TABLEOWNER].[load_order_header] load_order_header  
JOIN [TABLEOWNER].[load_order_line] load_order_line  
ON load_order_header.order_id = load_order_line.order_id  
LEFT OUTER JOIN [TABLEOWNER].[load_address] load_address  
ON load_order_header.sold_to_address_id = load_address.address_id
```

ANSI join code generated

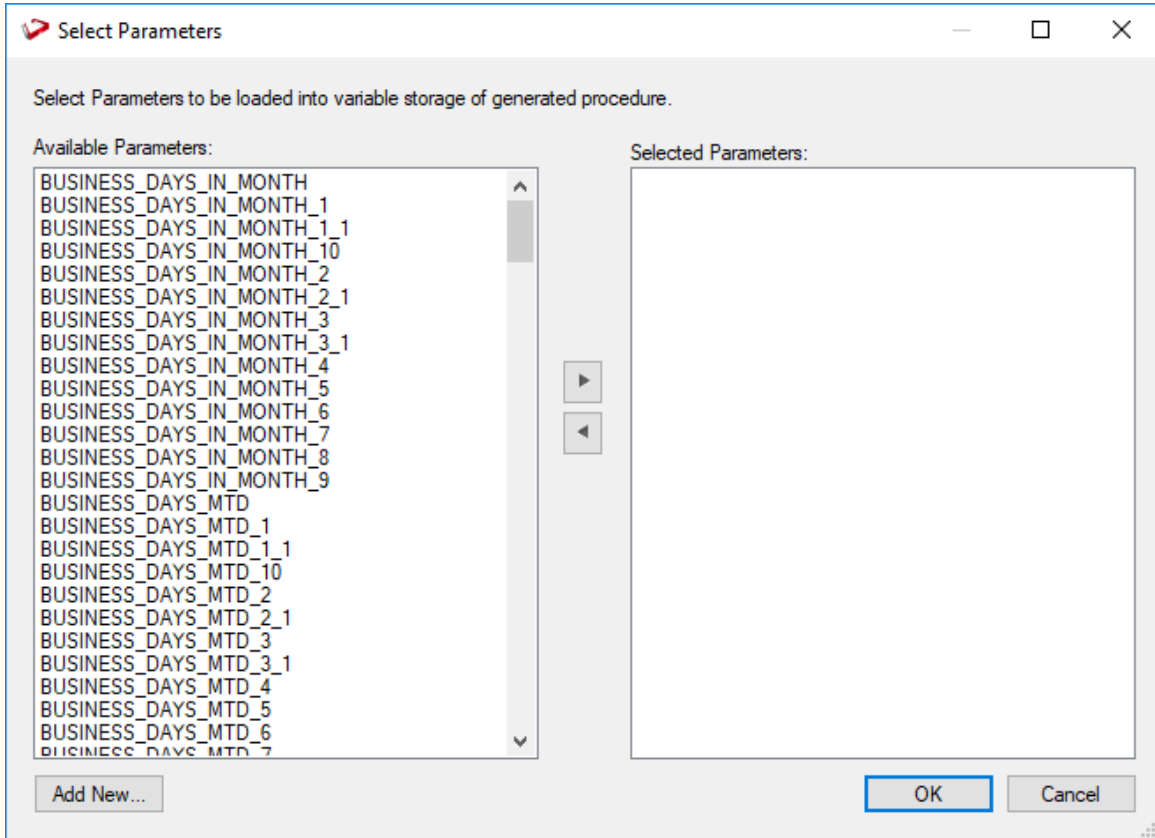
Select the columns that join the two tables. Select the column from the Master Table first.

load_order_header load_address

Word Wrap Displayed Code

PARAMETER SELECTION

The parameter selection dialog is used for choosing WhereScape RED **parameters** to be included as variables in the stage table procedure. Each parameter chosen is included in the procedure as `v_parameter_name`, limited to the first 30 characters. For example, parameter **THE_DAY_OF_THE_WEEK_FOR_LOADING_MONTH_END_DATA** will be available in the procedure as `v_THE_DAY_OF_THE_WEEK_FOR_LOAD`. RED automatically declares this variable and assigns it the current value of the parameter every time the procedure is run. Parameter variables can be used in column transformations or 'Where' clauses. A sample dialog box follows:



Note: If the desired parameter doesn't exist in the metadata yet, a new parameter can be added by clicking on the **Add New** button on the bottom leftmost corner of the Select Parameters dialog.

DIMENSION JOINS

For each dimension key, a dialog appears asking for the business key from the stage table that matches the business key for the dimension.

In the example below, we are asked for the stage table business key for the customer dimension table. The dimension name is shown both on the first prompt line and at the lower left side of the dialog box.

The customer dimension has a unique business key named **customer_code**.

We must provide the corresponding business key from the staging table. In the case of our example this is the **customer_code** column.

Note: The **Add Text** button and the associated message and edit box are only shown if the user possesses a full license, thus enabling the advanced procedure build options. When the **Add Text** button is clicked any data in the edit box to the right of the button is placed in the stage table column list. In this way a number or string can be assigned as part or all of a dimension join.

- 1 Click the **OK** button after the correct business key has been entered.
- 2 If the business key does not exist and will be derived from another dimension or from some form of lookup then enter any column and edit the procedure once produced.

DIMENSION HISTORY INFORMATION

If the dimension table being joined was defined as a dimension history table, then an additional dialog will appear asking for a date in the dimension table that allows the coordination of the dimension record changes. This dialog asks for a date field in the dimension table that enables RED to determine which version of the tracked column (the **customer_name** field, below) to use based on the specified date field.

Select the appropriate date field for your business needs and click **OK**. If you wish to take the last (or current) version for the tracked column, select **No Date**.

Tracking a History Table [X]

This model table has been identified as a history model table.
The columns listed below are tracked for changes.

If there is a date in the staging table that correlates to when these changes occur select the date and enter OK.

Model columns tracked for changes

customer_name

Staging Table Dates:

order_date

OK No Date

For Example:

As shown in the screen above, we have defined the **customer_name** as an column that we expect to have versions for over time i.e. each time the data warehouse processing sees a new **customer_name** value, the dimension table will record the date range for that version's validity—even though the business key (customer_code in this example) remains the same. This implies we want to create a new record in the dimension table, whenever a customer name is changed even though the customer_code remains the same.

Let's say a customer changes their name on the 5th of the month. If the **Staging Table Dates** field is set to **order_date**, any order received before the 5th of the month is identified under the old customer name and any order received on or after the 5th has the new customer name.

Alternatively, by setting the **Staging Table Dates** to **ship_date**, we can specify that any order *shipped* on or after the 5th of the month is shipped with the new name.

BUILDING AND COMPILING THE PROCEDURE

- Once the above questions are completed the procedure is built and compiled automatically.
- If the compile fails, an error is displayed along with the first few lines of error messages.
- Compilation failures typically occur when the physical creation of the table was not done.
- If the compile fails for some other reason, the best approach is to use the procedure editor to edit and compile the procedure.
- The procedure editor will highlight all the errors within the context of the procedure.
- Once the procedure has been successfully compiled it can either be executed interactively or passed to the scheduler.

STAGE TABLE CUSTOM PROCEDURE

A second procedure can be created on every stage table. This is called the custom procedure. Rather than modifying the generated procedure, it is often more practical to make additions to the generated code in a separate procedure. This allows for regeneration of the staging table's update procedure without losing changes (and having to reapply them).

The generated procedure for a custom procedure is template code. That is, a procedure that declares and initializes variables, does nothing and returns the correct return code and message for the WhereScape RED scheduler.

STAGE TABLE COLUMN PROPERTIES

Each stage table column has a set of associated properties. The definition of each property is defined below:

If the **Column name** or **Data type** is changed for a column then the metadata will differ from the table as recorded in the database. Use the **Validate>Validate Table Create Status** menu option to compare the metadata to the table in the database. A right-click menu option of **Alter Table** is available when positioned on the table name after the validate has completed. This option will alter the database table to match the metadata definition.



TIP: If a database table's definition is changed in the metadata then the table need to be altered in the database. Use the **Validate>Validate Table Create Status** to compare metadata definitions to physical database tables. The option also provides the ability to alter the database table, through a pop-up menu option from the validated table name.

A sample **Properties** screen is as follows:

General	
Table Name	stage_sales_detail
Column Name	customer_code
Column Title	customer code
Column Description	A unique code which identifies the customer. Forms a hierarchy with city and state.
Physical Definition	
Column Order	30
Data Type	integer
Null Values Allowed	<input checked="" type="checkbox"/>
Default Value	
Character Set	
Format	
Character Comparison/Sorting	
Compress	<input type="checkbox"/>
Meta Definition	
Numeric	<input checked="" type="checkbox"/>
Additive	<input checked="" type="checkbox"/>
Attribute	<input type="checkbox"/>
Business Key	<input type="checkbox"/>
Key Type (0,A,B,C,...)	
Source Details	
Source Table	load_order_header
Source Column	customer_code
Transformation	CAST(load_order_header.customer_code AS INTEGER)
Join	False

Column Name
Database-compliant name of the column.
Dialog Opening Value: customer_code

The two special update keys allow you to update the column and step either forward or backward to the next column's properties. **ALT-Left Arrow** and **ALT-Right Arrow** can also be used instead of the two special update keys.

Table Name

Database-compliant name of the table that contains the column. [Read-only].

Column Name

Database-compliant name of the column. A good practice is to only use alphanumeric characters, and the underscore character. Changing this field alters the table's definition.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Title

Name that the business uses to refer to the column. It does not affect the physical table definition, but rather provides input to the documentation and to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. As such it is a free form entry and any characters are valid.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Description

This field contains the description for the column. It might contain information on where and how the column was acquired. For example if the column is sourced from multiple tables or is a composite or derived column then this definition would normally describe the process used to populate the column. This field is used in the documentation and is available via the view **ws_admin_v_dim_col** . This field is also stored as a comment against the column in the database.

Column Order

Numeric value that controls the relative order of columns in the database create statement. The lowest numbered column will appear first in the table. Although this affects the physical table definition no action will be taken unless the table is re-created in the database. The columns can be re-numbered based on the existing order by choosing the **Respace order number** pop-up menu option when positioned over any column in the table. This action will number the columns in increments of 10 starting at 10. In addition to a simple change of the order field, the column order can be changed by first displaying the columns in the middle pane and then using drag and drop to move the columns around. This drag and drop process will automatically renumber the columns as required.

Data Type

Database-compliant data type that must be valid for the target database. Typical Teradata databases often have integer, numeric(), varchar(), char(), date and timestamp data types. See the database documentation for a description of the data types available. Changing this field alters the table's definition.

Null Values Allowed

Determines whether the table column can hold NULL values or whether a value is always mandatory.

Default Value

Initial value that is assigned to the column when a row is inserted into the table but no value is specified for the column.

Character Set

Database-compliant table column character-set used for storage. Select Latin or Unicode.

Format

Database-compliant table column format. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. As such it is a free form entry and any characters are valid. Typically format masks are only used on numeric fields. Example: #,###0.00. It is not worth the effort of populating this field unless it can be utilized by the end user tools in use.

Character Comparison/Sorting

Determines how the column character values are treated for comparison and sorting operations. Choose from: case specific, not case specific, uppercase case specific or uppercase not case specific.

Compress

Indicates whether the table column values are compressed when stored.

Compress/Compress Value

Optional list of values to be compressed. By default, only NULL is compressed if no list of values is specified.

Numeric

Indicates whether the table column holds values that are numeric. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Additive

Indicates whether the table column holds values that are additive. This implies that the column can be summed when performing data grouping in a query. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Attribute

Indicates whether the table column holds values that are descriptive, and/or are used for grouping/summing. An attribute is defined as a column that is non factual. For example, we may have an order number, or an invoice number stored in the fact table. Such columns are considered attributes, rather than facts. This checkbox is therefore normally only relevant for fact tables. This checkbox does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tools end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Business Key

Indicates whether the column is part of the BUSINESS key, which is defined during the update procedure build. [Normally maintained automatically]. Multiple columns can form the primary business key. This indicator is set and cleared by WhereScape RED during the dimension update procedure generation process. This check box should not normally be altered.

Key Type

Key type that is assigned and used when generating the table's update procedure and indexes. [Normally maintained automatically]. It can be altered here, but this should only be done if the consequences are fully understood and tested.

The supported values are:

Key type	Meaning
0	The artificial key. Set when the key is added during drag and drop table generation.
1	Component of all business keys. Indicates that this column is used as part of any business key. It is considered part of any lookup on that table and has the key type set to 1. Set when the column is added during drag and drop table generation.

Key type	Meaning
2	Indicates that this column is a dimension table join. Used on dimension tables to indicate the dimension keys to other dimension tables. Results in indexes being defined for the columns. Set during the update procedure generation for a dimension table, based on information from the staging table.
3	Not used in WhereScape RED for Teradata.
4	Not used in WhereScape RED for Teradata.
5	Indicates a column is a start date column.
6	Indicates a column is an end date column.
7	History column indicator. Used on dimension history tables to indicate that the column is being managed as a history column within the context of a dimension history table. Set when a column is identified during the dimension history update procedure generation.
c	Change Hash Key column indicator. Used in Data Vault tables to indicate the differences in the descriptive columns of a Satellite table. Refer to Data Vaults (on page 413) for details.
h	Hub Hash Key column indicator. Used in Data Vault tables to indicate the hash key column of a Hub Table. Refer to Data Vaults (on page 413) for details.
l	Link Hash Key column indicator. Used in Data Vault tables to indicate the hash key column of a Link Table. Refer to Data Vaults (on page 413) for details.
A	Indicates that the column is part of the primary business key. Set whenever a business key is defined as part of an update procedure generation.
B-Z	Indicates that the column is part of a secondary business key. Only used during index generation and not normally set.

Hash Key Sources

This field is only displayed for Hash key types. Displays the hash source columns that are used to generate the selected Hub, Link or Change hash key.

Hash Key Source For

This field is only displayed for Hash key types. Displays the hash keys columns that use the displayed hash key sources.

Source Table

Identifies the source table where the column's data comes from. This source table is normally a load table or a dimension table within the data warehouse. If the column was sourced from multiple tables, then the normal practice is to record one of the tables in this field and a comment listing all of the other tables in the **Source strategy** field. This field is used when generating a procedure to update the dimension. It is also used in the track back diagrams and in the documentation.

Source Column

Identifies the source column where the column's data comes from. Such a column is normally a load table column, which in turn may have been a transformation or the combination of multiple columns. This may also be a dimension table key where a dimension is being joined.

Transformation

See *Stage Table Column Transformations* (on page 351). [Read-only].

Join

Indicates whether the table column is used in a table join. [Normally maintained automatically but can be optionally changed to override the default join logic used in the generated update procedure]. The **Source table** and **Source column** fields will provide the dimension table's side of the join. The options for this field are: False, True, Manual and Pre Join.

Setting this field to Manual changes the way the dimension table is looked up during the staging table update procedure build. It allows you to join the dimension manually in the Cursor mapping dialog (where the 'Where' clause is built). The usual dialog for matching the dimension business key to a column or columns in the staging table is not displayed if this option is enabled.

Setting this field to Pre Join activates the **Pre Join Source Table** field and allows you to select a table from the drop-down list.

Pre Join Source Table

Indicates the table from which the pre joined column was sourced. When the Join option is set to False, this field becomes inactive. When the Join option is set to True or Manual, this field is set to the current table name. When the Join option is set to Pre Join, then you can select the required table from the drop-down list.

STAGE TABLE COLUMN TRANSFORMATIONS

Each stage table column can have a transformation associated with it. The transformation will be included in the generated procedure and will be executed as part of the procedure update.

The transformation must therefore be a valid SQL construct that can be included in a **Select** statement.

For example we could have a transformation of 'load_order_line.qty * 0.125' to calculate a tax column of 12.5%. Click the **Transformation** tab to enter a transformation.

The transformation screen is as follows:

Stage Column stage_sales_detail.customer_code

Properties

Transformation

Target: customer_code Paste

Source: load_order_header.customer_code Paste

Column Transformation Code (must execute within a SQL SELECT statement):

```
CAST(load_order_header.customer_code AS INTEGER)
```

Function Set: Default Teradata

- Functions
 - Number
 - String
 - Date
 - Conversion
 - Case Statement (CASE)
 - Cast as Type (CAST)
 - If Null Value (COALESCE)
 - Aggregate
 - Operator
 - User
 - Available Columns

Word Wrap Displayed Code

Function Syntax: CAST(column AS data_type)

Function Desc.: Returns the column converted to the specified data type and size.
Example: CAST(load_inventory.size AS VARCHAR(20))

<- Update Update -> OK Cancel Help

Note: Transformations are only put into effect when the procedure is re-generated.

See *Transformations* (on page 593) for more details.

PERMANENT STAGE TABLES

Note: WhereScape RED version 8.1.1.0 and above no longer supports the Permanent Stage table object sub type. However, users that have existing Permanent Stage table objects in their data warehouse can retain and continue to use them.

Normally Stage tables have all data removed from them at the start of their update procedure.

Sometimes it's necessary to hold more than a single set of data in a stage table. Permanent Stage tables allow this.

By default, a Permanent Stage table update does not start with all data being removed. However, options are available to selectively remove data.

For example, a Permanent Stage table may be used to hold the last three months source data provide functionality to remove all data for a day if new data for that day arrives.

GENERATING THE PERMANENT STAGING UPDATE PROCEDURE

Once a permanent stage table has been defined in the metadata and created in the data base an update procedure can be generated to handle the joining of any tables.


Note: You can also generate an update procedure via a template, refer to *Rebuilding Update Procedures* (on page 199) for details.

Generating a Procedure

- 1 To generate a procedure, right-click on the permanent stage table name in the left pane and select **Properties**.
- 2 From the **Update Procedure** drop-down list select **(Build Procedure...)**.
- 3 Click **OK** to update the properties and start the process of generating the new procedure.
- 4 A series of prompts are displayed during the procedure generation to join the tables and link in the model tables.

Procedure type

The first dialog box asks for the type of procedure that is to be generated:



A number of different types of procedure generation are available. Each type is discussed below. A check box appears at the bottom of the dialog if advanced procedure building features are enabled in the **Tools>Options** screen. This check box enables the editing of the 'Where' clause when no table joining is being performed, and hence the 'Where' clause would not be exposed.

Set based procedure

A set based procedure performs one SQL statement to join all the source tables together and then insert this data into the stage table. This is normally the fastest method of building the stage table. Caution must be taken in regards to NULL values in the business keys that are used to make the table joins. The generated code deliberately does not handle such null values.

Set Merge Procedure

This option is to allow the merging of two or more identical tables. The tables to be merged must have exactly the same number of columns and column names. If necessary additional blank columns could be added to one or other of the tables to make them identical. To use this procedure you must simply have the tables to be merged mentioned at least once in the **Source Table** field of a columns properties.

Set Distinct

Essentially the same as Set, except for the DISTINCT key word being added to the SELECT statement. This options therefore removes duplicate rows.

Set Minus

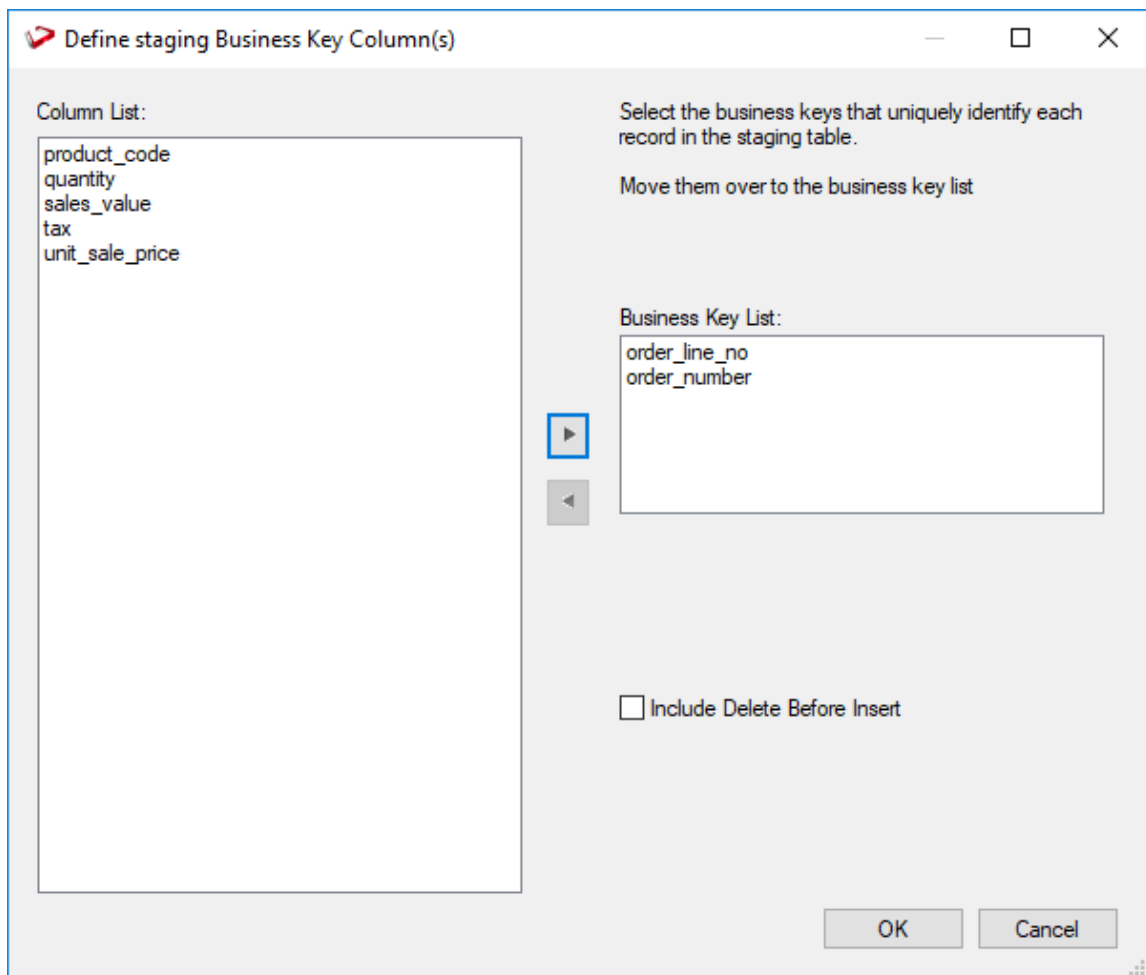
The **Set Minus** option can be used to determine change data or for programmatic referential integrity checking. This option works in a similar way to Set Merge. It generates SQL code in this form: `SELECT... FROM source_table1 {where} MINUS SELECT... FROM source_table2 {where}`. It requires exactly two source tables to be specified. All source columns must exist in both source tables.

Business Key definition

A dialog will appear asking for the business key that will uniquely identify each permanent stage record.

The source table from which the permanent stage is derived would normally have some form of unique constraint applied. In most cases this will be the business key.

In the example below **order_number** and **order_line_no** are selected as the business key.



A business key can be made up of multiple columns, but it must provide a unique identifier. Where multiple columns uniquely and separately identify the permanent stage, choose one to act as the primary business key. For example a source table may have a unique constraint on both a product code and a product description. Therefore the description as well as the code must be unique. It is of course possible to combine the two columns, but the normal practice would be to choose the code as the business key.

NULL Values: none of the columns chosen as the business key should ever contain a NULL value. See the note at the start of this chapter.

The **Include delete before insert** check will display an additional dialog wizard for building the 'Where' clause of the delete statement; at the start of the update procedure.

Definition for a Delete Statement Before the Stage Table Insert

Select the permanent stage table column on which the delete will be based. All values in the source (load or stage) table for the selected column will be used to delete the corresponding records in the permanent stage table.

Permanent Stage Table Column: Source (load or stage) Table: Source (load or stage) Table Column:

Issue Warning if a Delete Occurs

Truncate Table Before Insert

Delete Where Clause: The Where Clause can be edited, or directly entered.

```
WHERE stage_order_line_perm.order_number IN ( SELECT load_order_line.order_number FROM [TABLEOWNER].
```

Word Wrap Displayed Code

OK Cancel

Note: If a business key is specified, the generated procedure includes both the update and insert statement; otherwise only an insert statement is supplied.

Source table joins

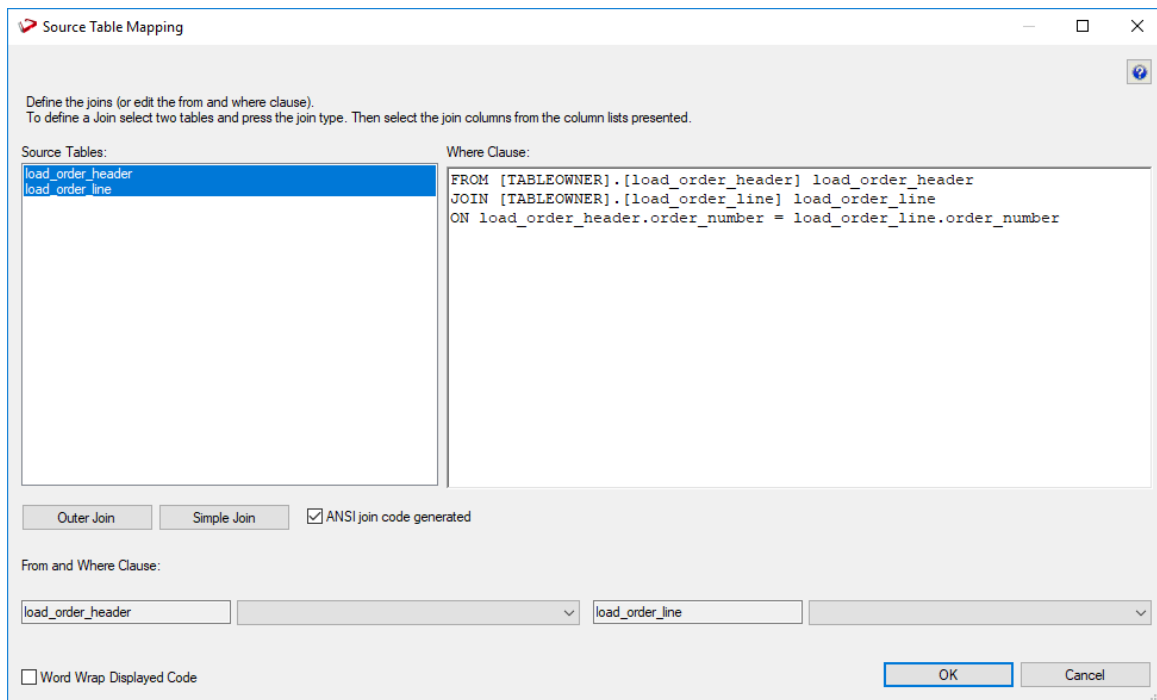
If multiple source tables are present then the definition of the joins between the various tables is required.

The joining of the tables will provide part of the construct of the set based update in the update procedure.

Only two tables may be joined at a time. To join two tables select the two tables in the left box and click either the outer join or simple join button. Column lists for both tables will appear at the bottom of the dialog box.

Select the column (or one of the columns) that allows the two tables to be joined. If an outer join is being used, the column for the master table must be chosen first. If there are multiple columns joining two tables then this action must be repeated for each column. Continue to perform all joins between all tables.

The example below only has two tables with one join column so is a relatively simple case. An additional option is available to allow either an ANSI standard join or a 'Where' clause based join. The ANSI standard join should be chosen in most situations. See the example screen in the following section.



Simple Join

A simple join joins the two tables, and only returns rows where data is matched in both tables. So for example if table A has 100 rows and table B has a subset of 24 rows. If all the rows in table B can be joined to table A then 24 rows will be returned. The other 76 rows from table A will not be returned.

Outer Join

An outer join joins the two tables, and returns all rows in the master table regardless of whether or not they are found in the second table. Therefore, if the example above was executed with table A as the master table, then 100 rows would be returned. 76 of those rows would have null values for the table B columns. In the example screen above the table 'load_order_line' has had its column chosen and the column for the table 'load_order_header' is currently being chosen. This will result in the statement as shown in the 'Where' clause edit window. The results of this select are that a row will be added containing order_line information, regardless of whether or not an order_header exists.

As the join columns are selected, the 'Where' statement is built up in the large edit window on the right. Once all joins have been made, the contents of this window can be changed if the join statement is not correct.

Once satisfied with the 'Where' statement click **OK** to proceed to the next step. As indicated in its description, this statement is the 'Where' clause that will be applied to the select statement of the cursor to allow the joining of the various source tables. It can of course be edited in the procedure that is generated if not correct.

You have the choice between 'Where' statement joins and ANSI standard joins.

Note: 'Where' joins are not available if using outer joins in Teradata.

The example below shows the result of an ANSI standard join which takes place in the 'From' statement.

Source Table Mapping

Define the joins (or edit the where clause).
To define a Join select two tables and select the join type. Then select the join columns from the column lists presented.

Source Tables:

- load_order_header
- load_order_line

From and Where Clause:

```
FROM [TABLEOWNER].[load_order_header] load_order_header
LEFT OUTER JOIN [TABLEOWNER].[load_order_line] load_order_line
ON load_order_header.order_id = load_order_line.order_id
```

Outer Join Simple Join ANSI join code generated

Select the columns that join the two tables. Select the column from the Master Table first.

load_order_header load_order_line

Word Wrap Displayed Code

OK Cancel

Building and Compiling the Procedure

- Once the above questions are completed the procedure is built and compiled automatically.
- If the compile fails an error will be displayed along with the first few lines of error messages.
- Compilation failures typically occur when the physical creation of the table was not done.
- If the compile fails for some other reason the best approach is to use the procedure editor to edit and compile the procedure.
- The procedure editor will highlight all the errors within the context of the procedure.
- Once the procedure has been successfully compiled it can either be executed interactively or passed to the scheduler.

SET MERGE PROCEDURE


The set merge procedure type allows the merging of two or more identical tables. The tables to be merged must have exactly the same number of columns and column names.

If necessary additional blank columns could be added to one or other of the tables to make them identical.

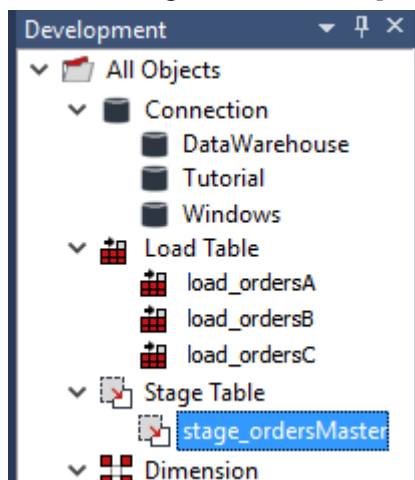
To use this procedure you must simply have the tables to be merged mentioned at least once in the **Source Table** field of a columns Properties.

In this example, we will merge three load tables into a single stage table.

- 1 The browser pane shows the three load tables:








- 1 Double click on the stage table object group and then drag one of these load tables from the source pane, into the Stage Object work area.
- 2 Name the stage table, for example **stage_ordersMaster**.

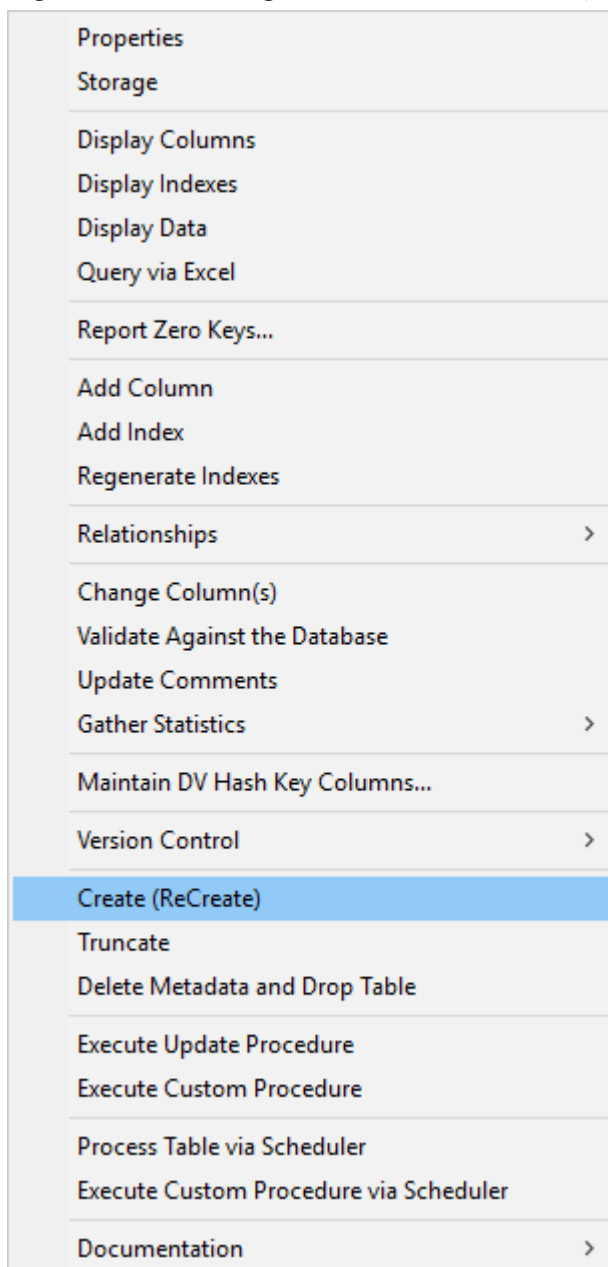


- 3 Next, modify the source table column to include one instance of each of the three load tables; the order does not matter.

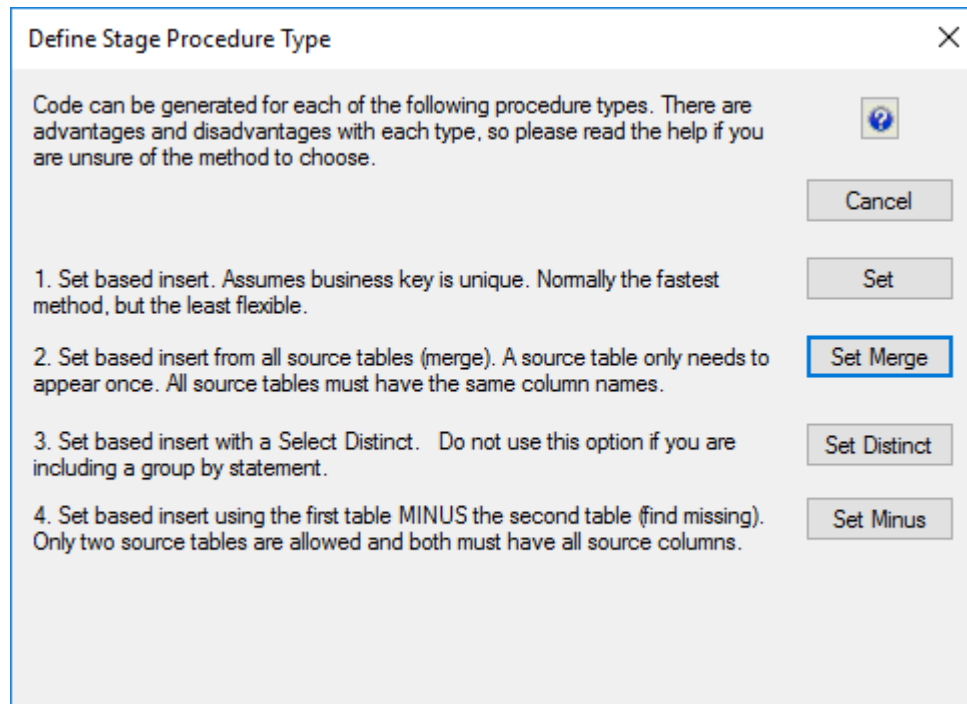
You can do this either by typing in the table directly, or by going to each of the column's properties screen.

stage_ordersMaster Columns			
Column Name	Display Name	Data Type	Source Table
 order_number	order number	numeric(12)	load_ordersA
 order_date	order date	datetime	load_ordersB
 customer_code	customer code	numeric(6)	load_ordersC
 ship_date	ship date	datetime	load_ordersA
 dss_update_time	dss update time	datetime	

- 4 Right-click on the stage table and select **Create (ReCreate)**.



- 5 Double click on the stage table in the left pane to bring up the **Properties** dialog.
 - Click the **Rebuild** button to rebuild the stored procedure.
 - Select **Set Merge** as the procedure type.



6 The stored procedure is rebuilt, as follows:

```

CREATE PROCEDURE [METABASE].update_stage_ordersMaster
(
  IN p_sequence      integer,
  IN p_job_name      varchar(256),
  IN p_task_name     varchar(256),
  IN p_job_id        integer,
  IN p_task_id       integer,
  OUT p_return_msg   varchar(256),
  OUT p_status       integer
)
BEGIN
  -- =====
  -- DBMS Name      : Teradata
  -- Script Name    : update_stage_ordersMaster
  -- Description    : Build the table stage_ordersMaster
  -- Generated by   : Version 6.5.5.2 (build 111221)
  -- Generated for  : Testing, Tutorials and Documentation
  -- Generated on   : Wednesday, February 08, 2012 at 13:52:09
  -- Author        : QuickStart
  -- =====
  -- Notes / History
  --
  -- =====
  -- Control variables used in most programs
  -- =====
  DECLARE v_msgtext      varchar(255); -- Text for audit_trail
  DECLARE v_sql          varchar(255); -- Text for SQL statements
  DECLARE v_set          integer;      -- commit set
  DECLARE v_analyze_flag integer;      -- analyze flag
  DECLARE v_step         integer;      -- return code
  DECLARE v_update_count integer;      -- no of records updated
  DECLARE v_insert_count integer;      -- no of records inserted
  DECLARE v_count        integer;      -- General counter
  DECLARE v_sql_code     integer;      -- SQL Error Code for Audit Trail
  DECLARE v_sql_error    varchar(255); -- SQL Error Code for Audit Trail as varchar

  -- =====
  -- Variables
  -- =====
  DECLARE v_dss_update_time timestamp; -- Used for date insert

  -- =====
  -- Exceptions
  -- =====
  DECLARE EXIT HANDLER
  FOR SQLEXCEPTION
  BEGIN
    SET v_sql_code = SQLCODE;
    LOCK ROW FOR ACCESS
    SELECT SUBSTR(ErrorText,1,255)
    INTO   :v_sql_error
    FROM   dbc.ErrorMsgs
    WHERE  ErrorCode = :v_sql_code;
    SET v_msgtext = SUBSTR('Unhandled Exception in stage_ordersMaster. ' ||
      'Step ' || CAST(v_step AS VARCHAR(64)) ||
      ' SQL Error Code: ' || CAST(v_sql_code AS VARCHAR(10)) || ' - ' || v_sql_error,1,255);
    SET p_return_msg = v_msgtext;
    CALL [METABASE].WsWrkAudit('F', :p_job_name, :p_task_name, :p_sequence
      , :v_msgtext, :v_sql_code, :v_sql_error, :p_task_id, :p_job_id);
    SET p_status = -3;
  END;

```

```
-----  
-- Main  
-----  
SET v_step = 100;  
SET v_dss_update_time = CURRENT_TIMESTAMP;  
SET v_update_count = 0;  
SET v_insert_count = 0;  
  
-----  
-- Delete all existing data from the table  
-----  
DELETE FROM [stage_ordersMaster] ALL;  
  
SET v_step = 200;  
  
-----  
-- Insert input records into stage_ordersMaster  
-- from load_ordersA  
-----  
INSERT INTO [stage_ordersMaster]  
( order_number  
, order_date  
, customer_code  
, ship_date  
, dss_update_time  
)  
SELECT load_ordersA.order_number  
, load_ordersA.order_date  
, load_ordersA.customer_code  
, load_ordersA.ship_date  
, :v_dss_update_time  
FROM [load_ordersA] load_ordersA  
;  
  
SET v_insert_count = v_insert_count + ACTIVITY_COUNT;  
  
SET v_step = 300;  
  
-----  
-- Insert input records into stage_ordersMaster  
-- from load_ordersB  
-----  
INSERT INTO [stage_ordersMaster]  
( order_number  
, order_date  
, customer_code  
, ship_date  
, dss_update_time  
)  
SELECT load_ordersB.order_number  
, load_ordersB.order_date  
, load_ordersB.customer_code  
, load_ordersB.ship_date  
, :v_dss_update_time  
FROM [load_ordersB] load_ordersB  
;  
  
SET v_insert_count = v_insert_count + ACTIVITY_COUNT;  
  
SET v_step = 400;
```

```
-----  
-- Insert input records into stage_ordersMaster  
-- from load_ordersC  
-----  
INSERT INTO [stage_ordersMaster]  
( order_number  
, order_date  
, customer_code  
, ship_date  
, dss_update_time  
)  
SELECT load_ordersC.order_number  
, load_ordersC.order_date  
, load_ordersC.customer_code  
, load_ordersC.ship_date  
, :v_dss_update_time  
FROM [load_ordersC] load_ordersC  
;  
  
SET v_insert_count = v_insert_count + ACTIVITY_COUNT;  
  
SET v_step = 500;  
  
-----  
-- All Done report the results  
-----  
-- Work out the return message  
SET p_status = 1;  
SET v_msgtext = 'stage_ordersMaster updated. '  
|| CAST(v_insert_count AS VARCHAR(64)) || ' new records. '  
|| CAST(v_update_count AS VARCHAR(64)) || ' records updated.';  
SET p_return_msg = v_msgtext;  
  
END;
```

CHAPTER 12

DATA STORE OBJECTS

IN THIS CHAPTER

Data Store Objects Overview	365
Building a Data Store Object.....	367
Generating the Data Store Update Procedure.....	371
Data Store Artificial Keys.....	378
Data Store Column Properties	382
Data Store Column Transformations	387

DATA STORE OBJECTS OVERVIEW

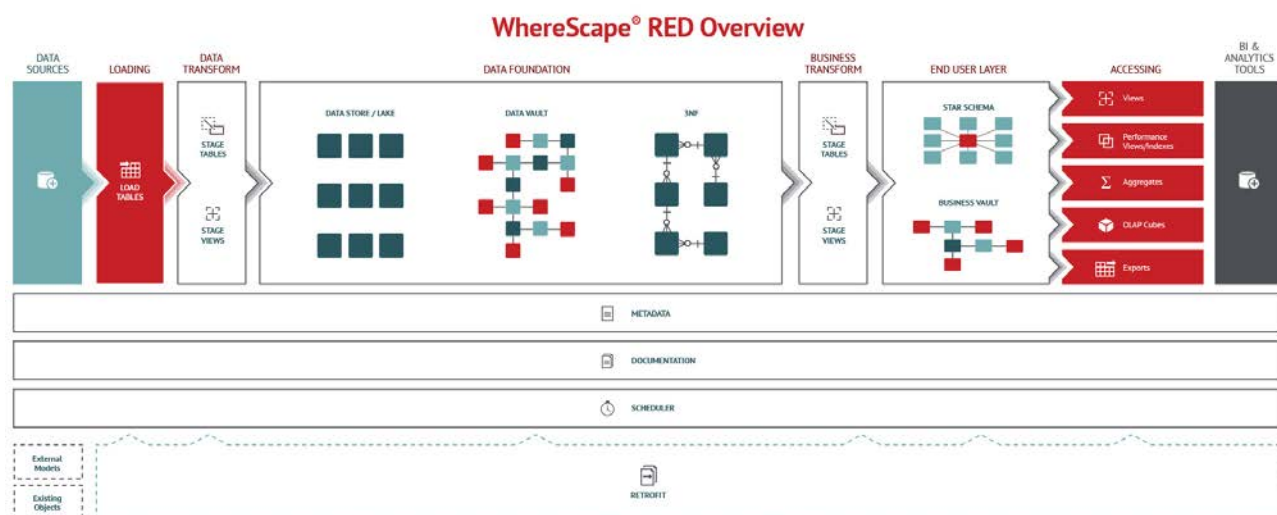
A Data Store Object is a data warehouse object used to store any type of data for later processing. In WhereScape RED, Data Store objects have many of the code generating attributes of stage, dimension and fact tables. Data Store objects can be thought of as a source system for the data warehouse. Alternatively, they may be reported off directly by users and reporting tools. Data Store Objects can be considered either reference or transactional in nature.

A Data Store Object is built from the Data Warehouse connection. Unless you are retrofitting an existing system, Data Store Objects are typically built from one or more load or stage tables. The Data Store model may be retrofitted from an enterprise modeling tool. See *Importing a Data Model* (on page 957) for more details.

The usual steps for creating a Data Store model are defined below and are covered in this chapter. The steps are:

- 1 Identify the source reference or transactional data that will constitute the Data Store Object. If the data is sourced from multiple tables ascertain if a join between the source tables is possible, or if one or more intermediate stage (work) tables would be a better option.
- 2 Using the 'drag and drop' functionality drag the load or stage table that is the primary source of information for the Data Store Object into a Data Store target. See *Building a Data Store Object* (on page 367).
- 3 If there's only one source table and all of the columns from it are being used, you can select the auto create option to build and load the table. This automatically completes the next four steps. See .
- 4 Add columns from other load and/or stage tables if required. See *Building a Data Store Object* (on page 367).
- 5 Create the Data Store Object in the database. See *Building a Data Store Object* (on page 367).
- 6 Build the update procedure. See *Generating the Data Store Update Procedure* (on page 371).
- 7 Run the update procedure and analyze the results.

If necessary, modify the update procedure or create a custom procedure.



Data Store Object Keys

Data Store Objects have Business Keys, they do not usually have Artificial Keys.

Business Key

The business key is the column or columns that uniquely identify a record within a Data Store Object. Where the Data Store Object maps back to a single or a main table in the source system, it is usually possible to ascertain the business key by looking at the unique keys for that source table. The business key is sometimes referred to as the 'natural' key. Examples of business keys are:

- The product SKU in a product table
- The customer code in a customer table
- The IATA airport code in an airport table.

It is assumed that business keys will never be NULL. If a null value is possible in a business key then the generated code will need to be modified to handle the null value by assigning some default value. In the following examples, the business key column is modified by using a database function and default value:

- `COALESCE(business_key,'N/A')`

Note: Business keys are assumed to never be Null. If they can be null it is best to transform them to some value prior to the Data Store or stage table update. If this is not done, an un-modified update procedure will probably fail with a duplicate key error on the business key index.

BUILDING A DATA STORE OBJECT

Data Store Objects are often sourced from one table in the base application. The process for building a Data Store Object begins with the drag and drop of the load or stage table that contains the bulk of the Data Store Object's information.

Drag and Drop

- 1 Create a Data Store Object target by double clicking on the **Data Store** group in the left pane. The middle pane will display a list of all existing Data Store Objects in the current project. When this list is displayed in the middle pane, the pane is identified as a target for new Data Store Objects.
- 2 Browse to the Data Warehouse via the **Browse Connection** menu option.
- 3 Drag the load or stage table that contains the bulk of the Data Store Object columns into the middle pane.
- 4 Drop the table anywhere in the pane. The new object dialog box will appear identifying the new object as a Data Store Object and providing a default name based on the load or stage table name.
- 5 Either accept this name or enter a name for Data Store Object and click **ADD** to proceed:

The screenshot shows a dialog box titled "Add a New Metadata Object" with a close button (X) in the top right corner. The dialog contains the following text and fields:

- Define the Type and Name of the New Object.
- Specific information for each object type is defined in subsequent screens.
- Object Type: Data Store (dropdown menu)
- Object Name: ds_customer (text input field)
- Target Location: DataWarehouse:doc6871_DATA_STORE (dropdown menu)
- Buttons: Add and Cancel

Data Store Object Properties

The table properties dialog for the new table displays.

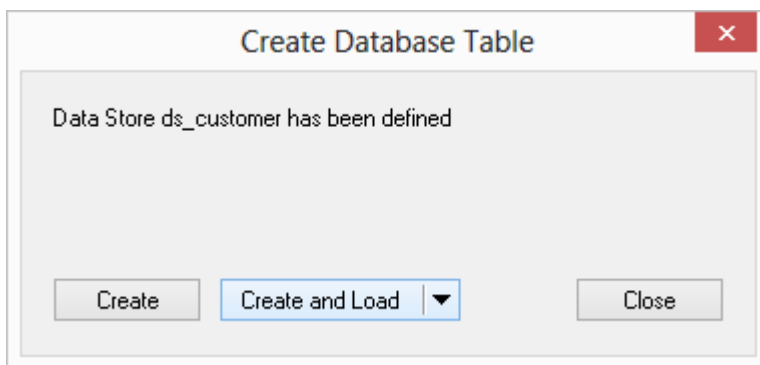
- If required, the Data Store Object can be changed to be a history table by choosing History from the table type drop-down list on the right side of the dialog. History tables are like

slowly changing dimensions in dimensional data warehouses. See **Building a Dimension** for more information. Change the storage options if desired.

- If prototyping, and the Data Store Object is simple (i.e. one source table) then it is possible to create, load and update the Data Store Object in just a couple of steps.
- If you want to do this, select the **(Build Procedure...)** option from the **Update Procedure** drop-down, and click **Create and Load** to the next screen.

Create and Load

If you chose to build the update procedure the following dialog appears after clicking OK on the Properties page. This dialog asks if you want to create the Data Store table in the database and execute the update procedure.

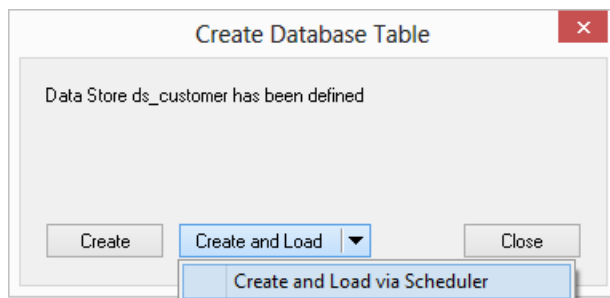


If you are satisfied with the columns that will be used and do not wish to add any columns you can select the **Create and Load** button. Alternatively, the **Create** button creates the table in the repository but does not execute an update, allowing you to change columns before loading data into the table.

If **Create** or **Create and Load** is selected and a new procedure creation was chosen proceed directly to the **Generating the Data Store Update Procedure** (on page 371) section.

If you have additional columns to add or columns to delete, click **Close** and proceed as below (Deleting and Changing Columns).

Note: It is possible to create and load the table via the Scheduler; by selecting this option from the drop-down list on the **Create and Load** button:



Deleting and Changing columns

The columns defined for the Data Store Object will be displayed in the middle pane.

- It is possible to delete any unwanted columns by highlighting a column name or a group of names and clicking the **Delete** key.
- You can change a column name by selecting the column and using the right-click context menu to edit its properties.
- Any new name must conform to the database naming standards. Good practice is to use alphanumeric characters and the underscore character.
- See the section *Data Store Column Properties* (on page 382) for a full description of fields.



TIP: When prototyping, and in the initial stages of an analysis area build, it is best not to remove columns, nor to change their names to any great extent. This type of activity is best left until after end users have used the data and provided feedback.

Adding additional columns

With the Data Store Object columns displayed in the middle pane, this pane is considered a drop target for additional columns.

- It is simple to select columns from other load and/or stage tables and drag these columns into the middle pane.
- The **source table column** in the middle pane shows where each column was dragged from.
- The column **description** could be acquired from three different tables.
- **Best practice** is to rename at least two of the columns, perhaps also adding context to the column name. For example, description could become group_description, and so forth.
- There are a number of WhereScape RED ancillary columns that do not have a source table. These columns have been added by WhereScape RED, and are added depending on earlier choices.

A description of these columns follows.

Column name	Description
dss_start_date	Used for history tables. This column provides a date time stamp when the Data Store Object record came into existence. It is used to ascertain which Data Store Object record should be used when multiple are available.
dss_end_date	Used for history tables. This column provides a date time stamp when the Data Store Object record ceased to be the current record. It is used to ascertain which Data Store Object record should be used when multiple are available.
dss_current_flag	Used for Data Store history tables. This flag identifies the current record where multiple versions exist.

Column name	Description
dss_source_system_key	Added to support history tables that cannot be fully conformed, and the inclusion of subsequent source systems. See the ancillary settings section for more details.
dss_version	Used for Data Store history tables. This column contains the version number of a Data Store history tables record. Numbered from 1 upwards with the highest number being the latest or current version. It forms part of the unique constraint for the business key of a Data Store history tables.
dss_update_time	Indicates when the record was last updated in the data warehouse.
dss_create_time	Indicates when the record was first created in the data warehouse

Create the table

Once the Data Store Object has been defined in the metadata you need to physically create the table in the database.

- 1 To do this, right-click on the Data Store Object name and select **Create (ReCreate)** from the pop-up menu.
- 2 The Results pane will display the results of the creation: a copy of the actual database create statement and if defined the results of any index create statements will be listed. For the initial create no indexes will be defined.
- 3 If the table was not created then ascertain and fix the problem. A common problem is a 'Duplicate column' where a column has the same name in two of the source tables. The best way of finding such a column is to double click on the list heading **Col name**. This will sort the column names into alphabetic order.
- 4 Another double click on the heading will sort the columns back into their create order.

The next section covers *Generating the Data Store Update Procedure* (on page 371).

GENERATING THE DATA STORE UPDATE PROCEDURE

Once a Data Store Object has been defined in the metadata and created in the database, an update procedure can be generated to handle the joining of any tables and the update of the Data Store Object.

Note: You can also generate an update procedure via a template, refer to *Rebuilding Update Procedures* (on page 199) for details.

Generating a Procedure

- 1 To generate a procedure, right-click on the Data Store Object in the left pane and select **Properties**.
- 2 Click on the **Rebuild** button to start the process of generating the new procedure.
- 3 A series of options are available.

PROCESSING TAB

The screenshot shows the 'ds_customer Update Build Options' dialog box. The 'Processing' tab is selected in the left pane. The main area contains the following options:

Template	None
Business Key Columns	code
Parameters	
Include Initial Load Insert	<input type="checkbox"/>
Create Storage Index After CTAS	<input type="checkbox"/>
Batch Processing	
Process by Batch	<input type="checkbox"/>
Explicit Locking	
Include Explicit Lock	<input checked="" type="checkbox"/>
Delete Processing	
Delete Before Insert	No
Update Processing	
Process Method	Insert/Update
Source Table Locking	
load_customer	ACCESS
Insert Method	
Include Insert Statement	<input checked="" type="checkbox"/>
Insert Hint	
Insert New Rows Only	<input checked="" type="checkbox"/>
New Row Identification Method	Join
Update Method	
Include Update Statement	<input checked="" type="checkbox"/>
Update Hint	
Update Changed Rows Only	<input checked="" type="checkbox"/>
Change Row Identification Method	Join

Business Key Columns
Columns that define the business key for update processing. Required if include Update options.

Buttons: Prev, Next, OK, Cancel, Help

Template: Enables you to generate update procedures via a *template* (see "*Rebuilding Update Procedures*" on page 199).

Business Key Columns: Columns that define the business key for update processing. Required for include Update options.

Clicking on the ellipsis button will bring up the Business Key selection screen.



TIP: Use the column name **ascending/descending** buttons to sort column names. To revert to the metadata column order, click on the **meta column order** button.

Business Key Columns

Columns that define the business key for update processing. Required if include Update options.

Available Columns: A-Z Z-A ⋮

name
address
city
state

Selected Columns:

code

OK Cancel

A business key can be made up of multiple columns, but it must provide a unique identifier. Where multiple columns separately uniquely identify rows in the Data Store Object, choose one to act as the primary business key.

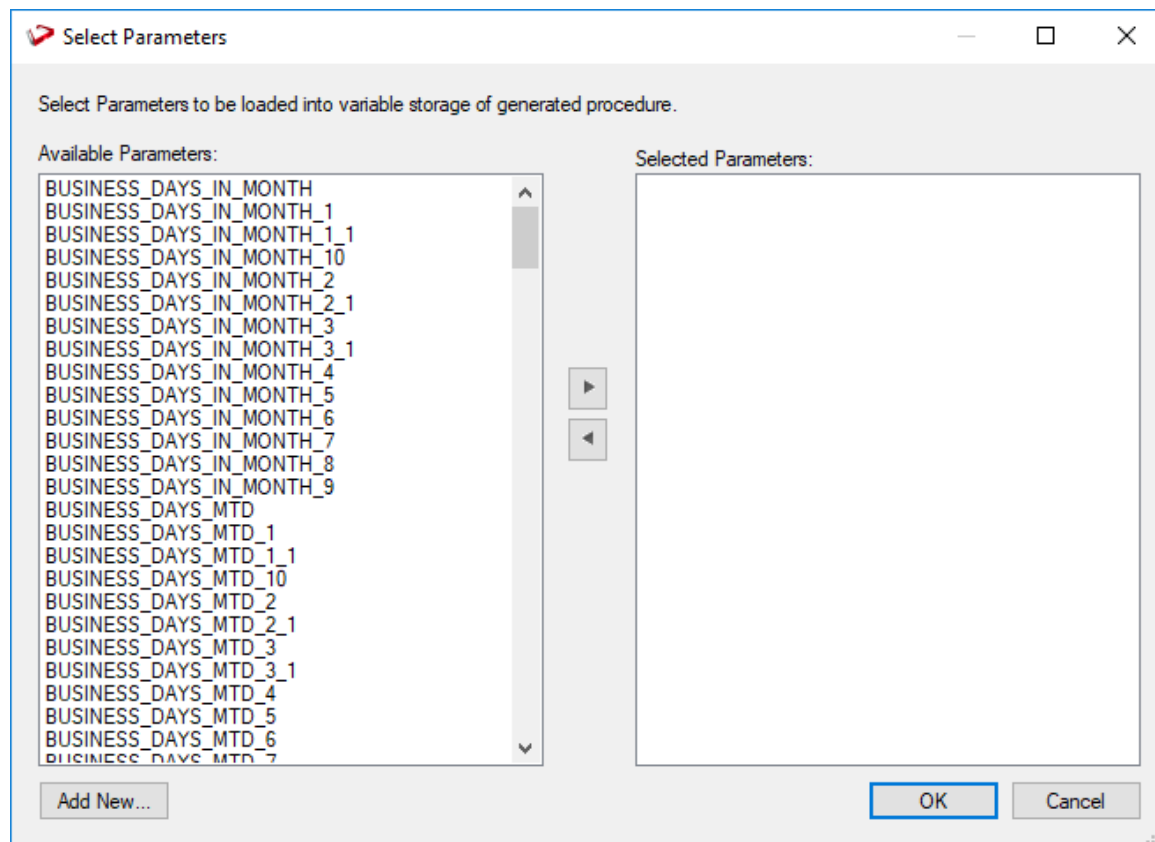
For example, a source table may have a unique constraint on both a product code and a product description. Therefore, the description as well as the code must be unique. It is of course possible to combine the two columns, but the normal practice would be to choose the code as the business key.

NULL Values: none of the columns chosen as the business key should ever contain a NULL value.

Parameters: Any parameters selected are included in the generated update procedure as variables.

The procedure will include code to retrieve the value of the parameter at run time and store it in the declared variable.

Clicking on the ellipsis button opens the Parameters selection screen.



The variables can also be used in column transformations and in the from/where clause for the update procedure. Some databases have a 30 character limit for variable names. WhereScape RED ensures the variables added for any parameters are less than 30 characters long by creating variable names in the form v_ followed by the first 28 characters of the parameter name.

For example, a parameter called `MINIMUM_ORDER_NUMBER_SINCE_LAST_SOURCE_LOAD` will be available as the variable `v_MINIMUM_ORDER_NUMBER_SINCE_L`.



TIP1: WhereScape RED parameters should be unique within the first 28 characters to avoid conflicting variable names.



TIP2: If the desired parameter doesn't exist in the metadata yet, a new parameter can be added by clicking on the **Add New** button on the bottom leftmost corner of the Select Parameters dialog.

See *Parameters* (on page 143) for more information on WhereScape RED Parameters

Include Initial Load Insert: adds an additional insert statement to the update procedure that runs if the target Data Store is empty. The benefit of this is improved performance inserting into an empty table without performing any checks to see if rows already exist. The default for this field is not set (i.e. an initial insert statement is not added to the procedure).

Process by Batch: allows users to select a column to drive data processing in a loop based on the distinct ordered values of the selected Business Key columns. The update procedure loops on this column and performs the delete, update and/or insert for each value. If the column chosen is a date datatype (date, datetime or timestamp), then the user is able to specify yearly, monthly, daily or column level looping. The default for this field is not set (do not do batch processing).

Delete Before Insert: allows selecting how to process deletes. It enables a delete statement to be added to the update procedure before any update or insert statement. This is a particularly useful option for purging old data and for updates based on a source system batch number. When this option is selected, it enables the **Issue Warning if a Delete occurs** and the **Delete Where Clause** Fields.

Issue Warning if a Delete occurs: this option sets the procedure to a warning state if any deletes occur.

Delete Where Clause: the delete where clause is appended to the generated delete statement to constrain the rows deleted.

Process Method: allows updating the Dimension with either an **Insert/Update** or a **Merge** statement. **Merge** allows you to use one Merge statement instead of two separate Insert and update statements.

Source Table Locking: allows a locking request modifier to be specified for each source table. The specified locking request modifier is applied to each source table during generated update procedures. By default this is set to 'ACCESS' which locks each row being accessed, a blank entry will result in no locking clause in the generated procedure.

Insert Method

Include Insert Statement: set this field to include the insert statement in the procedure. This allows inserting new rows in the Data Store.

Insert New Rows Only: uses change detection to work out which rows will require inserting.

New Row Identification Method: method used to identify that records in source are not currently recorded in the target table. Select **Join** or **Minus**.

Include Update Statement: set this field to include an update statement in the procedure. This allows updating the changing rows in the Data Store. If this is set, the **Update Changed Rows Only** option is available.

Update Changed Rows Only: uses change detection to work out what rows require updating. When set, this option enables the **Change Row Identification Method**.

Change Row Identification Method: method used to identify that records in source have changed from what is currently recorded in the target table. Select **Join** or **Minus**.

Merge Method

Merge Changed Rows only: uses change detection to work out what rows require merging. When the option is set, it enables the **New Row Identification Method**.

New Row Identification Method: method used to identify which records in the source are not recorded or are recorded differently in the target table. Select between **Join** and **Minus**.

If **non identity columns** are used as artificial keys the only new row identification method is **Join**.



Data Store Update procedures usually perform faster when you use the **Join** method for new row identification.

SOURCE TAB

ds_customer	
Distinct Data Select	<input type="checkbox"/>
Select Hint	
Source Join	FROM [TABLEOWNER].[load_customer] load_customer
Where Clause	
Group By	

ds_customer

Distinct Data Select: ensures duplicate rows are not added to the Data Store. This is achieved by adding the word DISTINCT to the source select in the update procedure. The default for this field is not set.

Source Join: The From clause, including Source Join information.

Simple Join

A simple join only returns rows where data is matched in both tables. So for example if table A has 100 rows and table B has a subset of 24 rows. If all the rows in table B can be joined to table A then 24 rows will be returned. The other 76 rows from table A will not be returned.

Outer Join

The outer join returns all rows in the master table regardless of whether or not they are found in the second table. Therefore, if the example above was executed with table A as the master table, then 100 rows would be returned. 76 of those rows would have null values for the table B columns.

Note: When WhereScape RED builds up an outer join, it needs to know which table is the master table and which is subordinate. Select the join column from the master table first. In the example screen above the table 'load_order_header' has had its column chosen and the column for the table 'load_order_line' is currently being chosen. This will result in the 'load_order_header' table being defined as the master, as per the example statement above. The results of this example select are that a row will be added containing order information regardless of whether or not a corresponding load_order_line entry exists.

Where Clause: The Where clause. Use as a filter to extract only the necessary records that fulfill a specified criteria.

Group By: The Group By clause. Use in collaboration with the SELECT statement to arrange identical data into groups.

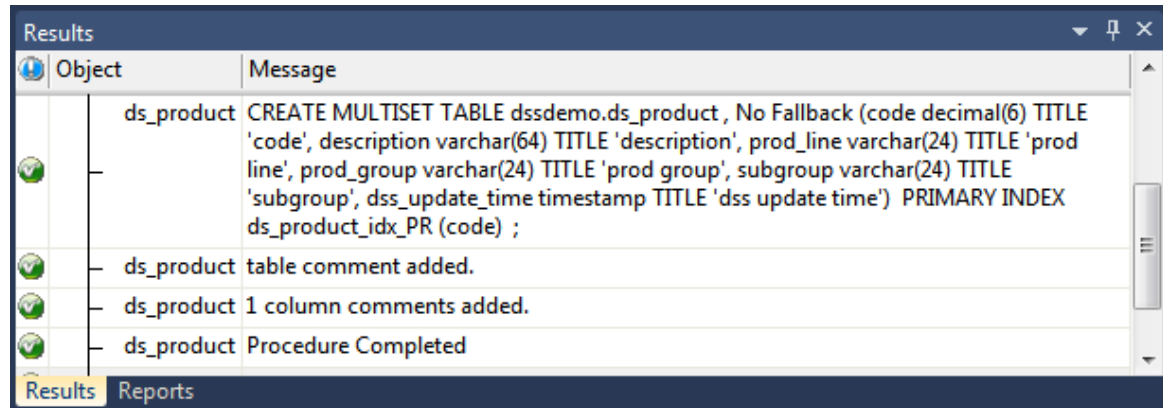
Building and Compiling the Procedure

- Once the relevant options are completed, click **OK**. The procedure will be built and compiled.
- If the compile fails an error will be displayed along with the first few lines of error messages. Compile fails typically occur when the physical creation of the table was not done.
- If the compile fails for some other reason the best approach is to use the procedure editor to edit and compile the procedure. The procedure editor will highlight all the errors within the context of the procedure.
- Once the procedure has been successfully compiled it can either be executed interactively or passed to the scheduler.

Indexes

By default, a number of indexes will be created to support each Data Store Object. These indexes will be added once the procedure has been built.

An example of the type of indexes created is as follows:



The screenshot shows a 'Results' window with a table containing four rows of messages. The first row shows the creation of a table with a primary index. The subsequent rows show the addition of comments to the table and columns, and finally, the completion of the procedure.

Object	Message
ds_product	CREATE MULTISET TABLE dssdemo.ds_product, No Fallback (code decimal(6) TITLE 'code', description varchar(64) TITLE 'description', prod_line varchar(24) TITLE 'prod line', prod_group varchar(24) TITLE 'prod group', subgroup varchar(24) TITLE 'subgroup', dss_update_time timestamp TITLE 'dss update time') PRIMARY INDEX ds_product_idx_PR (code) ;
ds_product	table comment added.
ds_product	1 column comments added.
ds_product	Procedure Completed

Additional indexes can be added, or these indexes changed. See the chapter on indexes for further details.

DATA STORE ARTIFICIAL KEYS

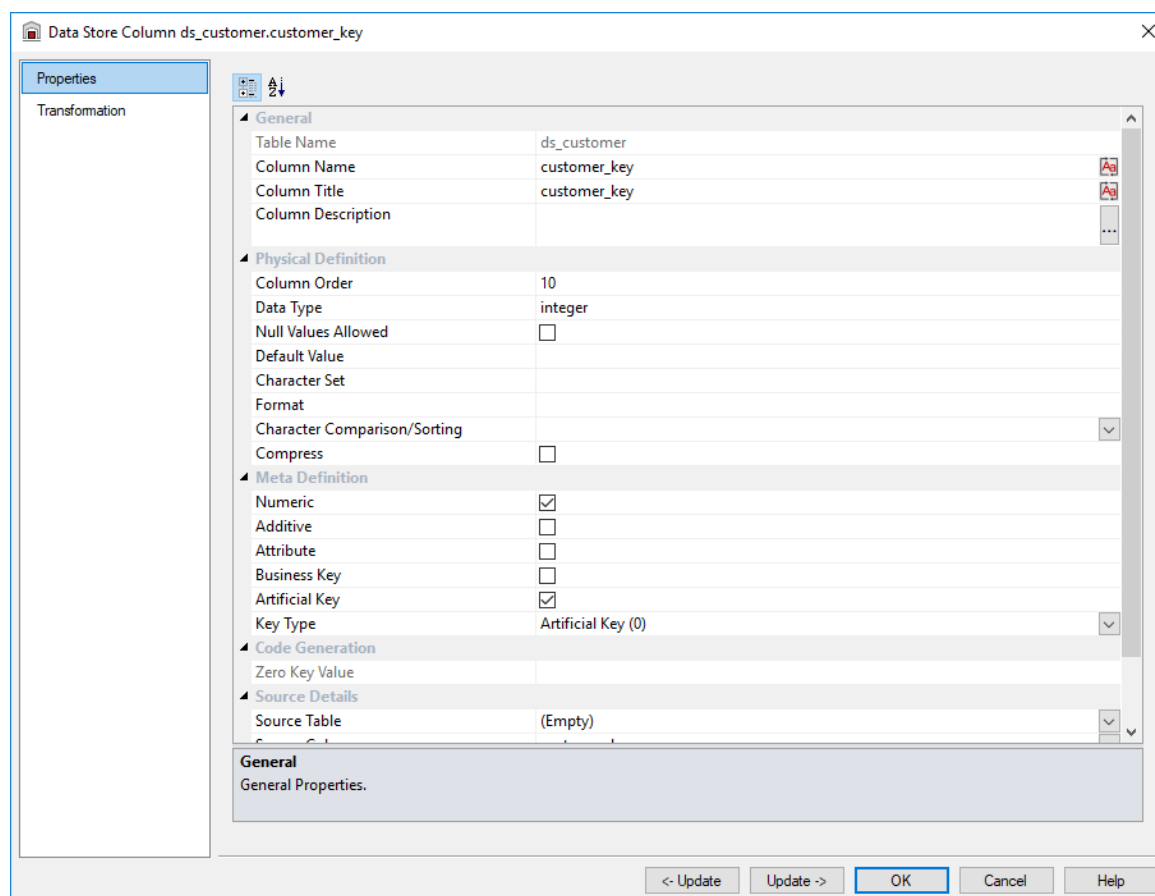
By default, Data Store Objects in WhereScape RED do not have an artificial (surrogate) key. Artificial keys can be added manually, but needing to do so could indicate Data Store Objects are not the correct WhereScape RED object for this table (perhaps an EDW 3NF Table would be more appropriate).

There is also an option for artificial keys to be automatically added to Data Store tables through an option in the **Tools** menu (see below - *Allowing for non identity surrogate keys on Data Store tables*).

To manually add an extra artificial key column to a Data Store table:

- 1 Right click in the middle pane and click either **Add Column** or **Duplicate Column**.
- 2 Edit the properties of the new column to have the correct name and order, source table and column, datatype, key type and flags as below.
- 3 The **Column Name** and **Source Column** should be the same.
- 4 The **Source Table** should be empty.
- 5 The **Data Type** should be integer.
- 6 The **Key Type** should be 0.
- 7 Only the **Numeric** and **Artificial Key** flags should be set.

The following example shows a manually added artificial key column:



The artificial key for a Data Store Object is set via an identity column. This artificial key normally, and by default, starts at one and progresses as far as is required.

A WhereScape standard for the creation of special rows in the EDW 3NF tables is as follows:

Key value	Usage
1 upwards	The standard artificial keys are numbered from 1 upwards, with a new number assigned for each distinct Data Store Object record.
0	Used as a join to the Data Store Object when no valid join existed. It is the convention in the WhereScape generated code that any EDW 3NF table business key that either does not exist or does not match is assigned to key 0.
-1 through -9	Used for special cases. The most common being where an EDW 3NF table is not appropriate for the record. A new key is used rather than 0 as we want to distinguish between records that are invalid and not appropriate.

Key value	Usage
-10 backward	Pseudo records. In many cases we have to deal with different granularities in our data. For example we may have a table that contains actual sales at a product SKU level and budget information at a product group level. The product table only contains SKU based information. To be able to map the budget records to the same table, we need to create these pseudo keys that relate to product groups. The values -10 and backwards are normally used for such keys.

Artificial keys set via a non identity column:

Data Store Tables can have non identity columns as surrogate keys.

The generation of the update procedure will automatically add logic to the code which will associate a sequential number to the artificial key when a new row is inserted into the EDW 3NF table.

The order of these sequential numbers is determined by the business key of the source table. The value of the first newly inserted artificial key will be the value of the highest artificial key in the dimension table plus 1.

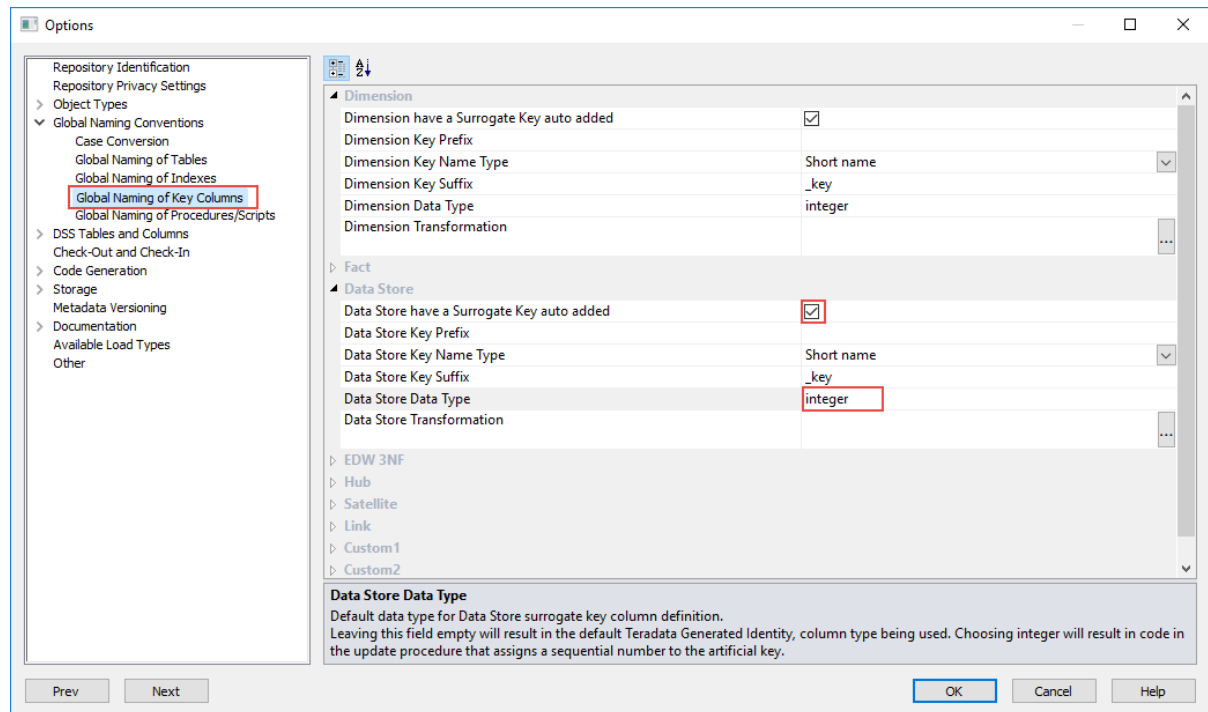
This automatically generated logic can be overwritten by defining a user specific logic on the **Dimension Transformation** field on the **Tools>Options** menu or in the transformation column of the artificial key.

To have an EDW 3NF table with a non identity column as a surrogate key, you can set the table **Data Type** to **integer** in the **Tools>Options** menu.

The old logic for dimensions can be retained if an identity column is chosen as surrogate key.

Allowing for non identity surrogate keys on Data Store tables:

- 1 Go to **Tools -> Options -> Global Naming Conventions -> Global Name of Key Columns**.
- 2 Expand the **Data Store** section.
- 3 Set the **Data Store have a Surrogate Key auto added** field if you want a surrogate key added by default to all Data Store tables.
- 4 Set the **Data Store Data Type** to be **integer** and click **OK**.
- 5 If your tables have been created previously, you will have to **Recreate** the tables after you set this option in the **Tools** menu.



DATA STORE COLUMN PROPERTIES

Each Data Store Object column has a set of associated properties. The definition of each property is described below:

If the **Column name** or **Data type** is changed for a column then the metadata will differ from the table as recorded in the database.

Use the **Validate>Validate Table Create Status** menu option or the right-click context menu to compare the metadata to the table in the database.

The **Alter table** option is available from the context menu when you right-click the table name, after the validate has completed. This option alters the database table to match the metadata definition.



TIP: If a database table's definition is changed in the metadata then the table will need to be altered in the database. Use the **Validate>Validate Table Create Status** to compare metadata definitions to physical database tables. The option also provides the ability to alter the database table, through a pop-up menu option from the validated table name.

A sample **Properties** screen is as follows:

Data Store Column ds_customer.code

Properties

Transformation

General

Table Name	ds_customer
Column Name	code
Column Title	code
Column Description	

Physical Definition

Column Order	10
Data Type	decimal(6)
Null Values Allowed	<input checked="" type="checkbox"/>
Default Value	
Character Set	
Format	
Character Comparison/Sorting	
Compress	<input type="checkbox"/>

Meta Definition

Numeric	<input checked="" type="checkbox"/>
Additive	<input checked="" type="checkbox"/>
Attribute	<input type="checkbox"/>
Business Key	<input checked="" type="checkbox"/>
Artificial Key	<input type="checkbox"/>
Key Type	Primary Business Key (A)

Code Generation

Zero Key Value	
----------------	--

Source Details

Source Table	load_customer
--------------	---------------

General

General Properties.

<- Update Update -> OK Cancel Help

The two special update keys allow you to update the column and step either forward or backward to the next column's properties.

ALT-Left Arrow and **ALT-Right Arrow** can also be used instead of the two special update keys.

Table Name

Database-compliant name of the table that contains the column. [Read-only].

Column Name

Database-compliant name of the column. Typically column-naming standards exclude the use of spaces etc. A good practice is to only use alphanumerics, and the underscore character. Changing this field alters the table's definition.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Business Display Name

Name that the business uses to refer to the column, which is included in the RED-generated documentation and can be used in the end user layer of other tools. [Does NOT affect the physical database table]. As such it is a free form entry and any characters are valid.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Description

This field contains the description for the column. It may be a description from a business user's point of view. This field might additionally contain information on where and how the column was acquired. For example if the column is sourced from multiple tables or is a composite or derived column then this definition would normally describe the process used to populate the column. This field is used in the documentation and is available via the view **ws_admin_v_dim_col** . This field is also stored as a comment against the column in the database.

Column Order

Numeric value that controls the relative order of columns in the database create statement. The lowest numbered column will appear first in the table. Although this affects the physical table definition no action will be taken unless the table is re-created in the database. The columns can be re-numbered based on the existing order by choosing the **Respace Order Number** pop-up menu option when positioned over any column in the table. This action will number the columns in increments of 10 starting at 10. In addition to a simple change of the order field, the column order can be changed by first displaying the columns in the middle pane and then using drag and drop to move the columns around. This drag and drop process will automatically renumber the columns as required.

Data Type

Database-compliant data type that must be valid for the target database. Typical Teradata databases often have integer, numeric(), varchar(), char(), date and timestamp data types. See the database documentation for a description of the data types available. Changing this field alters the table's definition.

Null Values Allowed

Determines whether the table column can hold NULL values or whether a value is always mandatory.

Default Value

Initial value that is assigned to the column when a row is inserted into the table but no value is specified for the column.

Character Set

Database-compliant table column character-set used for storage. Select Latin or Unicode.

Format

Optional format mask that can be used in end user tools. [Does NOT affect the physical database table]. As such it is a free form entry and any characters are valid. Typically format masks are only used on numeric fields. Example: #,###0.00. It is not worth the effort of populating this field unless it can be utilized by the end user tools in use.

Numeric

Indicates whether the table column holds values that are numeric. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Additive

Indicates whether the table column holds values that are additive. This implies that the column can be summed when performing data grouping in a query. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Attribute

Indicates whether the table column holds values that are descriptive, and/or are used for grouping/summing. An attribute is defined as a column that is non factual. For example we may have an order number, or a invoice number stored in the fact table. Such columns are considered attributes, rather than facts. This checkbox is therefore normally only relevant for fact tables. This checkbox does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tools end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Business Key

Indicates whether the column is part of the BUSINESS key, which is defined during the update procedure build. [Normally maintained automatically]. Multiple columns can form the primary business key.

Artificial Key

Indicates whether the column is the artificial key. Only one artificial key is supported. This indicator is set by WhereScape RED during the initial drag and drop creation of a table, and should not normally be altered.

Key Type

Key type that is assigned and used when generating the table's update procedure and indexes. [Normally maintained automatically]. It can be altered here, but this should only be done if the consequences are fully understood and tested. The supported values are:

Key type	Meaning
0	The artificial key. Set when the key is added during drag and drop table generation.
1	Component of all business keys. Indicates that this column is used as part of any business key.
A	Indicates that the column is part of the primary business key. Set whenever a business key is defined as part of an update procedure generation.
B-Z	Indicates that the column is part of a secondary business key. Only used during index generation and not normally set.

Source Table

Identifies the source table where the column's data comes from. This source table is normally a load table within the data warehouse. If the column was sourced from multiple tables, then the normal practice is to record one of the tables in this field and a comment listing all of the other tables in the **Source strategy** field. This field is used when generating a procedure to update the Data Store object. It is also used in the track back diagrams and in the documentation.

Source Column

Identifies the source column where the column's data comes from. Such a column is normally a load table column, which in turn may have been a transformation or the combination of multiple columns. This may also be a dimensional key where a dimension is being joined.

Transformation

See *Data Store Column Transformation* (see "*Data Store Column Transformations*" on page 387). [Read-only].

Join

Indicates whether the table column is used in a table join. [Normally maintained automatically but can be optionally changed to override the default join logic used in the generated update procedure]. The **Source Table** and **Source Column** fields will provide the dimension table's side of the join. The options for this field are: False, True, Manual and Pre Join.

Setting this field to Manual changes the way the dimension table is looked up during the update procedure build. It allows you to join the dimension manually in the Cursor mapping dialog (where the 'Where' clause is built).

Setting this field to Pre Join activates the **Pre Join Source Table** field and allows you to select a table from the drop-down list.

Pre Join Source Table

Indicates the table from which the pre joined column was sourced. When the Join option is set to False, this field becomes inactive. When the Join option is set to True or Manual, this field is set to the current table name. When the Join option is set to Pre Join, then you can select the required table from the drop-down list.

DATA STORE COLUMN TRANSFORMATIONS

Each Data Store Object column can have a transformation associated with it. The transformation will be included in the generated procedure and will be executed as part of the procedure update. The transformation must therefore be a valid SQL construct that can be included in a **Select** statement.

For example, we could have a transformation of 'load_order_line.qty * 0.125' to calculate a tax column of 12.5%.

Click the **Transformation** tab to enter a transformation.

The transformation screen is as follows:

Note: Transformations only take effect when the procedure is re-generated.

See *Transformations* (on page 593) for more details.

CHAPTER 13

EDW 3NF TABLES

NOTE: EDW 3NF/Normalized Table rename

Former WhereScape RED Normalized Tables have been renamed to EDW 3NF from RED version 6.8.4.3.

However, this change applies only for new metadata repositories, existing metadata repositories will not be affected and will not have its table's naming modified.

All references to Normalized tables have been updated in the RED documentation from version 6.8.4.3 onwards, however, some screenshots of the RED left pane browser might still show instances of the Normalized object type instead of the new EDW 3NF type.

To modify table naming from Normalized to EDW 3NF in existing repositories see **Object Type Names** and **Global Naming Conventions**.

Please note that short name and table prefixes can be overwritten by the **Local Naming conventions** setting in the **User Preferences**. If this is the case, you can disable this option here: **Local Naming conventions**.

IN THIS CHAPTER

EDW 3NF Tables Overview	389
Building EDW 3NF Table	391
Generating the EDW 3NF Update Procedure	395
EDW 3NF Table Artificial Keys	403
EDW 3NF Table Column Properties	406
EDW 3NF Table Column Transformations.....	411

EDW 3NF TABLES OVERVIEW

An EDW 3NF Table is a data warehouse object used to build third normal form enterprise data warehouses. In WhereScape RED, EDW 3NF objects have many of the code generating attributes of stage, dimension and fact tables. Third normal form enterprise data warehouses can be thought of as a source system for star schema data marts. Alternatively, they may be reported off directly by users and reporting tools. EDW 3NF Tables can be considered either reference or transactional in nature.

An EDW 3NF Table is built from the Data Warehouse connection. Unless you are retrofitting an existing system, EDW 3NF Tables are typically built from one or more load or stage tables. The EDW 3NF model may be retrofitted from an enterprise modeling tool. See *Importing a Data Model* (on page 957) for more details.

The usual steps for creating an EDW 3NF model are defined below and are covered in this chapter. The steps are:

- 1 Identify the source reference or transactional data that will constitute the EDW 3NF Table. If the data is sourced from multiple tables ascertain if a join between the source tables is possible, or if one or more intermediate stage (work) tables would be a better option.
- 2 Using the 'drag and drop' functionality drag the load or stage table that is the primary source of information for the EDW 3NF Table into an EDW 3NF target. See *Building an EDW 3NF Table* (see "*Building EDW 3NF Table*" on page 391)
- 3 If there's only one source table and all of the columns from it are being used, you can select the auto create option to build and load the table. This automatically completes the next four steps. See *Building an EDW 3NF Table* (see "*Building EDW 3NF Table*" on page 391)
- 4 Add columns from other load and/or stage tables if required. See *Building a EDW 3NF Table* (see "*Building EDW 3NF Table*" on page 391)
- 5 Create the EDW 3NF Table in the database. See *Building an EDW 3NF Table* (see "*Building EDW 3NF Table*" on page 391)
- 6 Build the update procedure. See *Generating the EDW 3NF Update Procedure* (on page 395)
- 7 Run the update procedure and analyze the results.

If necessary, modify the update procedure or create a custom procedure.

EDW 3NF Table Keys

EDW 3NF Tables have two types of keys that we will refer to frequently. These are the Business Key and the Artificial Key. A definition of these two key types follows:

Business Key

The business key is the column or columns that uniquely identify a record within an EDW 3NF Table. Where the EDW 3NF Table maps back to a single or a main table in the source system, it is usually possible to ascertain the business key by looking at the unique keys for that source table. The business key is sometimes referred to as the 'natural' key. Examples of business keys are:

- The product SKU in a product table
- The customer code in a customer table
- The IATA airport code in an airport table.

It is assumed that business keys will never be NULL. If a null value is possible in a business key then the generated code will need to be modified to handle the null value by assigning some default value. In the following examples, the business key column is modified by using a database function and default value:

- COALESCE(business_key,'N/A')

Note: Business keys are assumed to never be Null. If they can be null it is best to transform them to some value prior to the EDW 3NF or stage table update. If this is not done, an un-modified update procedure will probably fail with a duplicate key error on the business key index.

Artificial Key

By default, EDW 3NF Tables in WhereScape RED do not have an artificial key (artificial keys can be added manually or set to be added by default through the Tools menu. See **EDW 3NF Table Artificial Keys** for more details.

An artificial key is the unique identifier that can be used to join an EDW 3NF Table record to other EDW 3NF Tables. When joining EDW 3NF Tables it would be possible to perform the join using the business key. For EDW 3NF Tables that satisfy one or more of the following conditions, joining with business keys could result in slow query times and excessive use of database storage:

- Multiple column business keys (excessive storage and multiple column joins)
- One or more large character business key columns (excessive storage)
- Very large tables (excessive storage - integer artificial keys often use less space than one small character field)
- History EDW 3NF Tables (complex joins involving a between dates construct)

As query time is one of our key drivers in data warehouse implementations the best answer is often to use some form of artificial key. A price is paid in the additional processing required doing key lookups, but this is offset by the reduced query times and reduced complexity.

The artificial key is an integer and is built sequentially from 1 upwards. See the section on artificial keys for a more detailed explanation. An artificial key is sometimes referred to as a "surrogate" key.

BUILDING EDW 3NF TABLE

EDW 3NF tables are often sourced from one table in the base application. The process for building an EDW 3NF table begins with the drag and drop of the load or stage table that contains the bulk of the EDW 3NF table's information.

Drag and Drop

- 1 Create an EDW 3NF table target by double clicking on the **EDW 3NF** group in the left pane. The middle pane will display a list of all existing EDW 3NF tables in the current project. When such a list is displayed in the middle pane, the pane is identified as a target for new EDW 3NF tables.
- 2 Browse to the Data Warehouse via the Browse>Source Tables menu option.
- 3 Drag the load or stage table, that contains the bulk of the EDW 3NF table columns, into the middle pane. Drop the table anywhere in the pane.
- 4 The new object dialog box appears and identifies the new object as an EDW 3NF table and will provide a default name based on the load or stage table name.
- 5 Either accept this name or enter the name of the EDW 3NF table and click ADD to proceed.

EDW 3NF Table Properties

The table properties dialog for the new table is displayed.

If required, the EDW 3NF table can be changed to be a history table by choosing History from the table type drop-down list on the right side of the dialog.

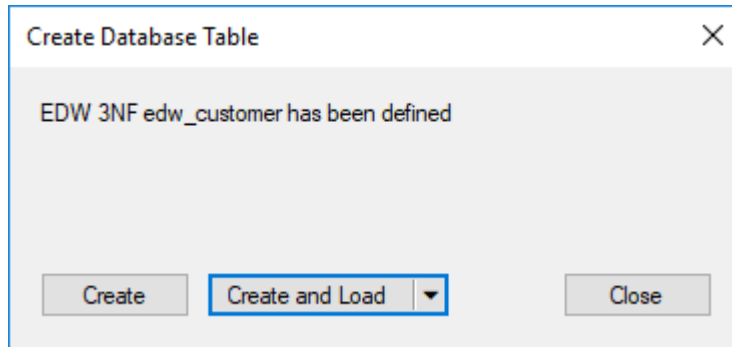
History tables are like slowly changing dimensions in dimensional data warehouses. See **EDW 3NF Table Column Properties** (on page 406) for more information. Change the storage options if desired.

If prototyping, and the EDW 3NF table is simple (i.e. one source table) then it is possible to create, load and update the EDW 3NF table in a couple of steps.

If you wish to do this select the **(Build Procedure...)** option from the **Update Procedure** drop-down, and answer **Create and Load** to the next question.

Create and Load

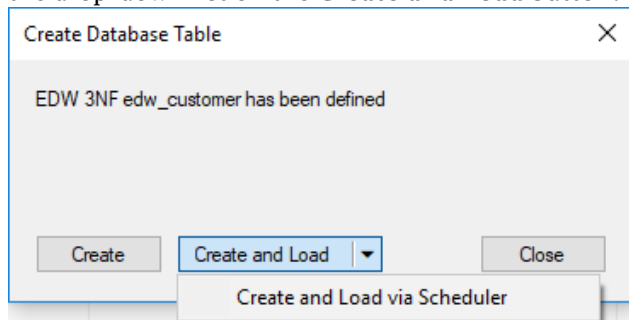
If you chose to build the update procedure the following dialog appears after clicking OK on the Properties page. This dialog asks if you want to create the EDW 3NF table in the database and execute the update procedure.



If you are satisfied with the columns that will be used and do not wish to add any columns you can select the **Create and Load** button. Alternatively, the **Create** button creates the table in the repository but does not execute an update, allowing you to change columns before loading data into the table.

If **Create** or **Create and Load** is selected and a new procedure creation was chosen proceed directly to the *Generating the EDW 3NF Update Procedure* (on page 395) section.

Note: It is possible to create and load the table via the Scheduler; by selecting this option from the drop-down list on the **Create and Load** button:



If you have additional columns to add or columns to delete, then select **Finish** and proceed as follows.

Deleting and Changing columns

The columns defined for the EDW 3NF table will be displayed in the middle pane. It is possible to delete any unwanted columns by highlighting a column name or a group of names and clicking the **Delete** key.

The name of a column can also be changed by selecting the column and using the right-click menu to edit its properties. Any new name must conform to the database naming standards. Good practice is to use alphanumerics and the underscore character. See the section on column properties for a fuller description on what the various fields mean.



TIP: When prototyping, and in the initial stages of an analysis area build it is best not to remove columns, nor to change their names to any great extent. This type of activity is best left until after end users have used the data and provided feedback.

Adding additional columns

With the columns of the EDW 3NF table displayed in the middle pane, this pane is considered a drop target for additional columns.

It is a simple matter to select columns from other load and/or stage tables and drag these columns into the middle pane. The source table column in the middle pane shows where each column was dragged from. The column **description** could be acquired from three different tables. The best practice is to rename at least two of the columns, perhaps also adding context to the column name. For example, description could become group_description, and so forth. There are a number of WhereScape RED ancillary columns that do not have a source table. These columns have been added by WhereScape RED, and are added depending on earlier choices.

A description of these columns follows.

Column name	Description
dss_start_date	Used for history tables. This column provides a date time stamp when the EDW 3NF table record came into existence. It is used to ascertain which EDW 3NF table record should be used when multiple are available.
dss_end_date	Used for history tables. This column provides a date time stamp when the EDW 3NF table record ceased to be the current record. It is used to ascertain which EDW 3NF table record should be used when multiple are available.
dss_current_flag	Used for EDW 3NF history tables. This flag identifies the current record where multiple versions exist.
dss_source_system_key	Added to support history tables that cannot be fully conformed, and the inclusion of subsequent source systems. See the ancillary settings section for more details.
dss_version	Used for EDW 3NF history tables. This column contains the version number of an EDW 3NF history tables record. Numbered from 1 upwards with the highest number being the latest or current version. It forms part of the unique constraint for the business key of an EDW 3NF history tables.
dss_update_time	Indicates when the record was last updated in the data warehouse.
dss_create_time	Indicates when the record was first created in the data warehouse

Create the table

Once the EDW 3NF table has been defined in the metadata we need to physically create the table in the database.

This is achieved by right-clicking on the EDW 3NF table name and selecting **Create (ReCreate)** from the pop-up menu.

A results dialog box will appear to show the results of the creation. The contents of this dialog are a message to the effect that the EDW 3NF table was created.

A copy of the actual database create statement and if defined the results of any index create statements will be listed. For the initial create no indexes will be defined.

If the table was not created then ascertain and fix the problem. A common problem is a 'Duplicate column' where a column has the same name in two of the source tables. The best way of finding such a column is to double click on the list heading 'Col name'. This will sort the column names into alphabetic order.

Another double click on the heading will sort the columns back into their create order.

The next section covers *Generating the EDW 3NF Update Procedure* (on page 395).

GENERATING THE EDW 3NF UPDATE PROCEDURE

Once an EDW 3NF Object has been defined in the metadata and created in the database, an update procedure can be generated to handle the joining of any tables and the update of the EDW 3NF Object.

Note: You can also generate an update procedure via a template, refer to *Rebuilding Update Procedures* (on page 199) for details.

Generating a Procedure

- 1 To generate a procedure, right-click on the EDW 3NF Object in the left pane and select **Properties**.
- 2 Click the **Rebuild** button to start the process of generating the new procedure.
- 3 A series of options are available.

PROCESSING TAB

Option	Value
Template	None
Business Key Columns	
Parameters	
Include Initial Load Insert	<input type="checkbox"/>
Create Storage Index After CTAS	<input type="checkbox"/>
Batch Processing	
Process by Batch	<input type="checkbox"/>
Explicit Locking	
Include Explicit Lock	<input checked="" type="checkbox"/>
Delete Processing	
Delete Before Insert	No
Update Processing	
Process Method	Insert/Update
Source Table Locking	
stage_customer	ACCESS
Insert Method	
Include Insert Statement	<input checked="" type="checkbox"/>
Insert Hint	
Insert New Rows Only	<input checked="" type="checkbox"/>
New Row Identification Method	Join
Update Method	
Include Update Statement	<input checked="" type="checkbox"/>
Update Hint	
Update Changed Rows Only	<input checked="" type="checkbox"/>
Change Row Identification Method	Join

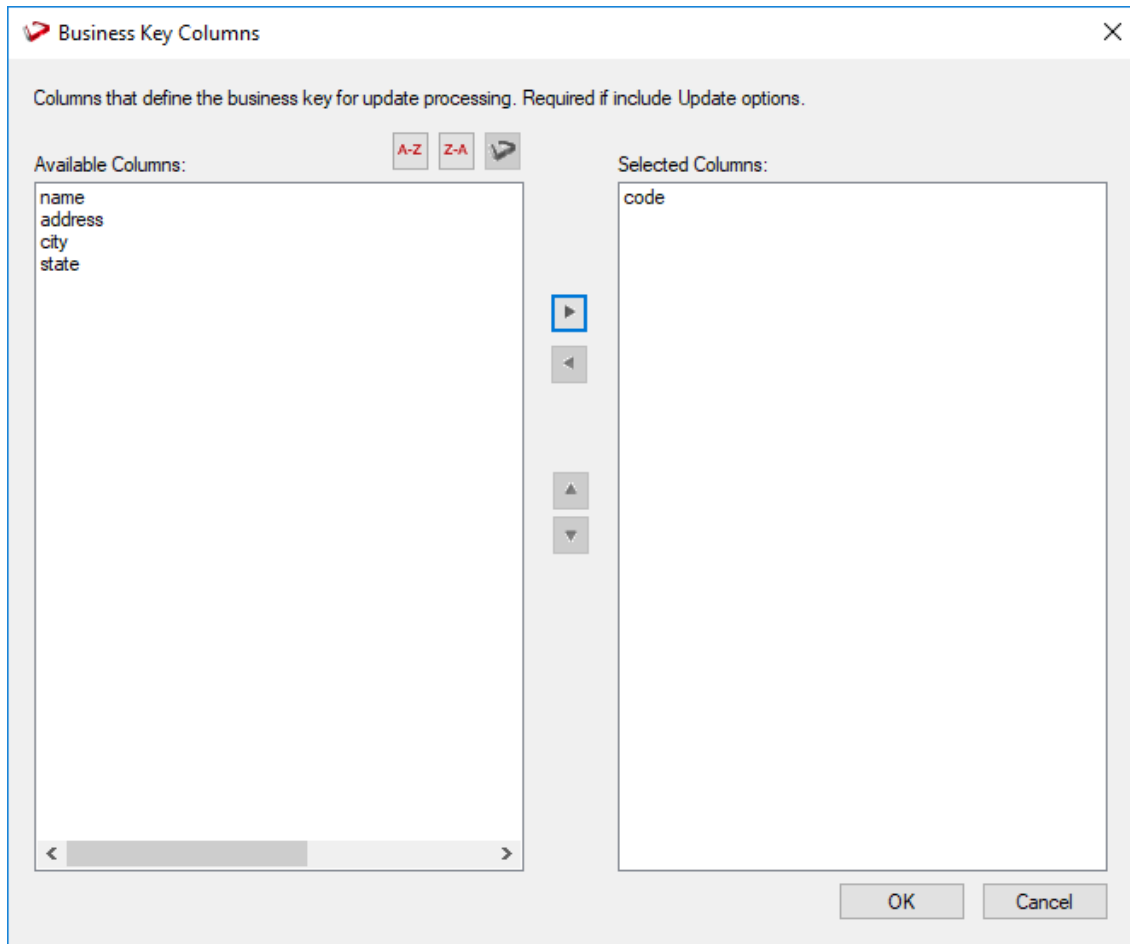
Template: Enables you to generate update procedures via a *template* (see "*Rebuilding Update Procedures*" on page 199).

Business Key Columns: Columns that define the business key for update processing. Required for include Update options.

Clicking the ellipsis button brings up the Business Key selection screen.



TIP: Use the column name **ascending/descending** buttons to sort column names. To revert to the metadata column order, click on the **meta column order** button.



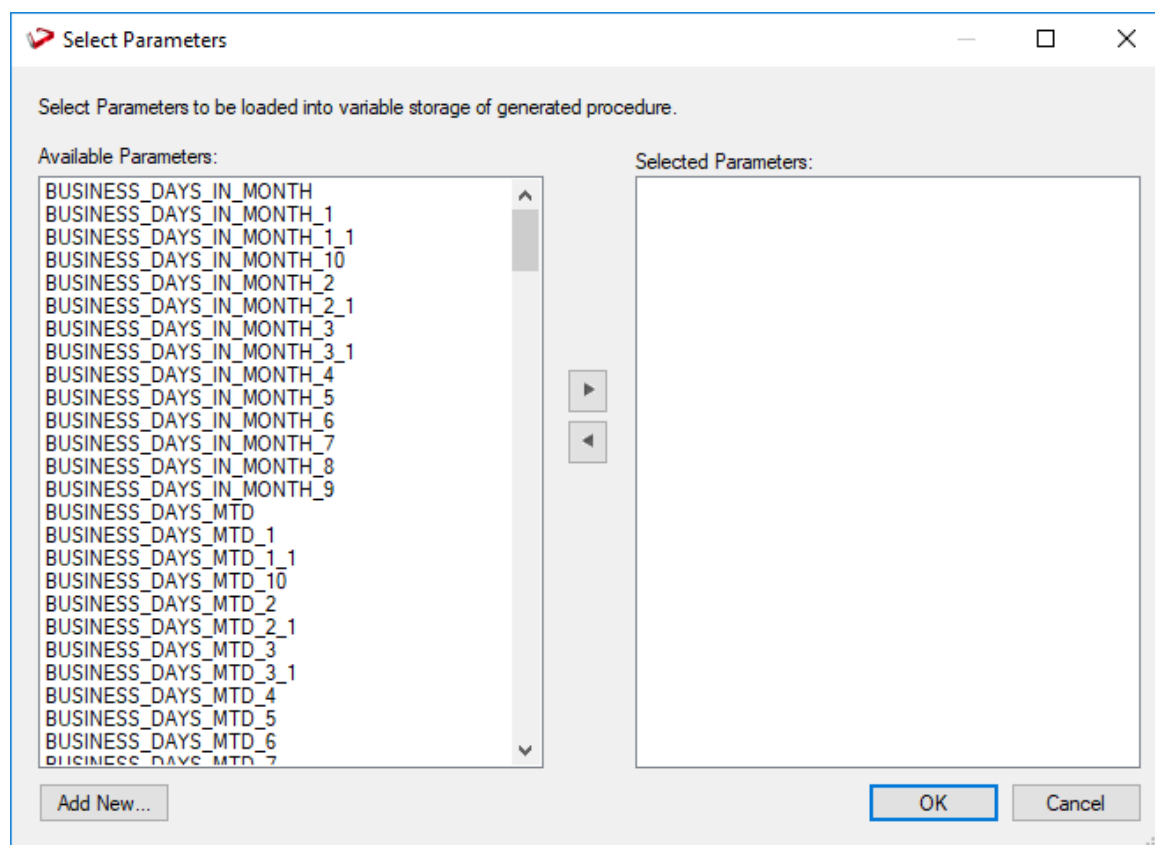
A business key can be made up of multiple columns, but it must provide a unique identifier. Where multiple columns separately uniquely identify rows in the EDW 3NF Object, choose one to act as the primary business key. For example a source table may have a unique constraint on both a product code and a product description. Therefore the description as well as the code must be unique. It is of course possible to combine the two columns, but the normal practice would be to choose the code as the business key.

NULL Values: none of the columns chosen as the business key should ever contain a NULL value.

Parameters: Any parameters selected are included in the generated update procedure as variables. The procedure will include code to retrieve the value of the parameter at run time and store it in the declared variable.

Clicking the ellipsis button will bring up the **Select Parameters** screen.

If the desired parameter doesn't exist in the metadata yet, a new parameter can be added by clicking on the **Add New** button on the bottom leftmost corner of the **Select Parameters** screen.



The variables can also be used in column transformations and in the from/where clause for the update procedure. Some databases have a 30 character limit for variable names. WhereScape RED ensures the variables added for any parameters are less than 30 characters long by creating variable names in the form v_ followed by the first 28 characters of the parameter name.

For example, a parameter called `MINIMUM_ORDER_NUMBER_SINCE_LAST_SOURCE_LOAD` will be available as the variable `v_MINIMUM_ORDER_NUMBER_SINCE_L`.



TIP: WhereScape RED parameters should be unique within the first 28 characters to avoid conflicting variable names.

See **Parameters** (on page 143) for more information on WhereScape RED Parameters

Include Initial Load Insert: adds an additional insert statement to the update procedure that runs if the target EDW 3NF table is empty. The benefit of this is improved performance inserting into an empty table without performing any checks to see if rows already exist. The default for this field is not set (i.e. an initial insert statement is not added to the procedure).

Process by Batch: allows users to select a column to drive data processing in a loop based on the distinct ordered values of the selected Business Key columns. The update procedure loops on this column and performs the delete, update and/or insert for each value. If the column chosen is a date datatype (date, datetime or timestamp), then the user is able to specify yearly, monthly, daily or column level looping. The default for this field is not set (do not do batch processing).

Delete Before Insert: allows selecting how to process deletes. It enables a delete statement to be added to the update procedure before any update or insert statement. This is a particularly useful option for purging old data and for updates based on a source system batch number. When this option is selected, it enables the **Issue Warning if a Delete occurs** and the **Delete Where Clause** Fields.

Issue Warning if a Delete occurs: this option sets the procedure to a warning state if any deletes occur.

Delete Where Clause: the delete where clause is appended to the generated delete statement to constrain the rows deleted.

Process Method: allows updating the EDW 3NF table with either an **Insert/Update** or a **Merge** statement. **Merge** allows you to use one Merge statement instead of two separate Insert and update statements.

Source Table Locking: allows a locking request modifier to be specified for each source table. The specified locking request modifier is applied to each source table during generated update procedures. By default, this is set to 'ACCESS' which locks each row being accessed, a blank entry will result in no locking clause in the generated procedure.

Insert Method

Include Insert Statement: set this field to include the insert statement in the procedure. This allows inserting new rows in the EDW 3NF table.

Insert New Rows Only: uses change detection to work out which rows will require inserting.

New Row Identification Method: method used to identify that records in source are not currently recorded in the target table. Select **Join** or **Minus**.

Include Update Statement: set this field to include an update statement in the procedure. This allows updating the changing rows in the EDW 3NF table. If this is set, the **Update Changed Rows Only** option is available.

Update Changed Rows Only: uses change detection to work out which rows require updating. When set, this option enables the **Change Row Identification Method**.

Change Row Identification Method: method used to identify that records in source have changed from what is currently recorded in the target table. Select **Join** or **Minus**.

Merge Method

Merge Changed Rows only: uses change detection to work out what rows require merging. When the option is set, it enables the **New Row Identification Method**.

New Row Identification Method: method used to identify which records in the source are not recorded or are recorded differently in the target table. Select between **Join** and **Minus**.

If **non identity columns** are used as artificial keys the only new row identification method is **Join**.



EDW 3NF Update procedures usually perform faster when you use the **Join** method for new row identification.

SOURCE TAB

Property	Value
Distinct Data Select	<input type="checkbox"/>
Select Hint	
Source Join	FROM [TABLEOWNER].[stage_customer] stage_customer
Where Clause	
Group By	

Distinct Data Select: ensures duplicate rows are not added to the EDW 3NF Object. This is achieved by the word DISTINCT being added to the source select in the update procedure. The default for this field is off.

Source Join: The From clause, including Source Join information.

Simple Join

A simple join only returns rows where data is matched in both tables. So for example if table A has 100 rows and table B has a subset of 24 rows. If all the rows in table B can be joined to table A, then 24 rows will be returned. The other 76 rows from table A will not be returned.

Outer Join

The outer join returns all rows in the master table, regardless of whether or not they are found in the second table. Therefore, if the example above was executed with table A as the master table, then 100 rows would be returned. 76 of those rows would have null values for the table B columns.

Note: When WhereScape RED builds up an outer join, it needs to know which table is the master table and which is subordinate. Select the join column from the master table first. In the example screen above, the table 'load_order_header' has had its column chosen and the column for the table 'load_order_line' is currently being chosen. This will result in the 'load_order_header' table being defined as the master, as per the example statement above. The results of this example select are that a row will be added containing order information regardless of whether or not a corresponding load_order_line entry exists.

Where Clause: The Where clause.

Group By: The Group By clause.

Building and Compiling the Procedure

- Once the relevant options are completed, click **OK**. The procedure will be built and compiled.
- If the compile fails, an error is displayed along with the first few lines of error messages. Compile fails typically occur when the physical creation of the table was not done.
- If the compile fails for some other reason the best approach is to use the procedure editor to edit and compile the procedure. The procedure editor will highlight all the errors within the context of the procedure.
- Once the procedure has been successfully compiled it can either be executed interactively or passed to the scheduler.

INDEXES

By default a number of indexes will be created to support each EDW 3NF table. An example of the type of indexes created is as follows:

Results	
Object	Message
	PRIMARY INDEX edw_customer_idx_PR (customer_code) ;
edw_customer	COMMENT ON COLUMN WSL_REPOSITORY.edw_customer.dss_update_time IS 'Date and time the row was updated in the data warehouse.';
edw_customer	Procedure Completed
edw_customer	1 edw_customer updated. 5000 records added. 0 records updated.

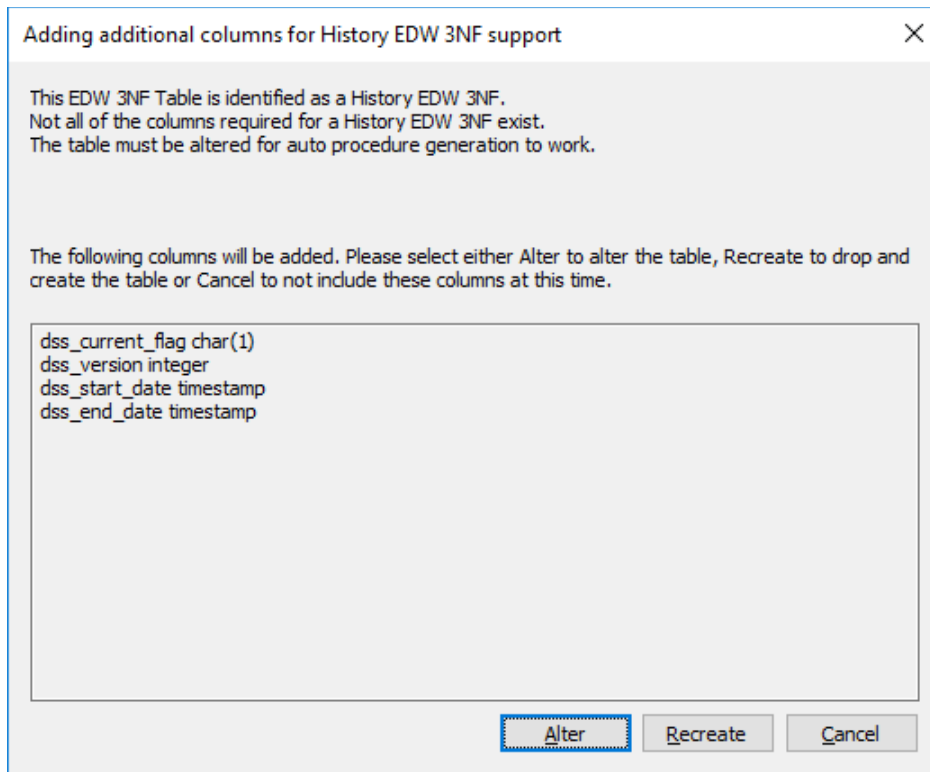
Additional indexes can be added, or these indexes changed. See the chapter on indexes for further details.

Converting an existing EDW 3NF Table to a EDW 3NF History Table


To convert a EDW 3NF table to a EDW 3NF history table, change the table type to **History**, select **(Build Procedure...)** from the **Update Procedure** drop-down list and click **OK**.

If the existing EDW 3NF table is **NOT** to be dropped and recreated, then the following process is followed:

- 1 Click **Alter** on the Adding Additional Columns dialog.



- 2 Click **Alter Table** on the Alter Table Commands dialog:



Note: The SQL commands in this dialog can be edited if required.

- 3 Click **OK** on the succeeding dialogs.
- 4 Define required fields on the succeeding dialogs and then click **OK** to finish.

EDW 3NF TABLE ARTIFICIAL KEYS

By default, EDW 3NF tables in WhereScape RED do not have an artificial (surrogate) key. Artificial keys can be added manually but you can also set an option in the **Tools** menu to have them automatically added to new tables as well (see below - *Allowing for non identity surrogate keys on EDW 3NF tables*).

To manually add an extra artificial key column to an EDW 3NF table:

- 1 Right click in the middle pane and click either **Add Column** or **Duplicate Column**.
- 2 Edit the properties of the new column to have the correct name and order, source table and column, data type, key type and flags as below.
- 3 The **Column Name** and **Source Column** must be the same.
- 4 The **Source Table** must be empty.
- 5 The **Data Type** must be integer.
- 6 The **Key Type** must be 0.
- 7 Only the **Numeric** and **Artificial Key** flags must be set on.

The following example shows a manually added artificial key column:

The screenshot shows the 'EDW 3NF Column' dialog box with the following configuration:

Section	Property	Value
General	Table Name	edw_customer
	Column Name	customer_key
	Column Title	customer_key
	Column Description	Manually added artificial key
Physical Definition	Column Order	10
	Data Type	integer
	Null Values Allowed	<input type="checkbox"/>
	Default Value	
	Character Set	
	Format	
	Character Comparison/Sorting	
	Compress	<input type="checkbox"/>
Meta Definition	Numeric	<input checked="" type="checkbox"/>
	Additive	<input type="checkbox"/>
	Attribute	<input type="checkbox"/>
	End User Layer Display	<input checked="" type="checkbox"/>
	Business Key	<input type="checkbox"/>
	Artificial Key	<input checked="" type="checkbox"/>
Key Type	Artificial Key (0)	
Source Details	Source Table	(Empty)
	Source Column	customer_key
	Transformation	

At the bottom of the dialog, there are buttons for 'OK', 'Cancel', and 'Help'.

Artificial Keys set via identity columns:

The artificial key for an EDW 3NF table is set via an identity column. This artificial key normally, and by default, starts at one and progresses as far as is required.

A WhereScape standard for the creation of special rows in the EDW 3NF tables is as follows:

Key value	Usage
1 upwards	The standard artificial keys are numbered from 1 upwards, with a new number assigned for each distinct EDW 3NF table record.
0	Used as a join to the EDW 3NF table when no valid join existed. It is the convention in the WhereScape generated code that any EDW 3NF table business key that either does not exist or does not match is assigned to key 0.
-1 through -9	Used for special cases. The most common being where an EDW 3NF table is not appropriate for the record. A new key is used rather than 0 as we want to distinguish between records that are invalid and not appropriate.
-10 backward	Pseudo records. In many cases we have to deal with different granularities in our data. For example, we may have a table that contains actual sales at a product SKU level and budget information at a product group level. The product table only contains SKU based information. To be able to map the budget records to the same table, we need to create these pseudo keys that relate to product groups. The values -10 and backwards are normally used for such keys.

Artificial keys set via a non identity column:

EDW 3NF Tables can have non identity columns as surrogate keys.

The generation of the update procedure will automatically add logic to the code which will associate a sequential number to the artificial key when a new row is inserted into the EDW 3NF table.

The order of these sequential numbers is determined by the business key of the source table. The value of the first newly inserted artificial key will be the value of the highest artificial key in the dimension table plus 1.

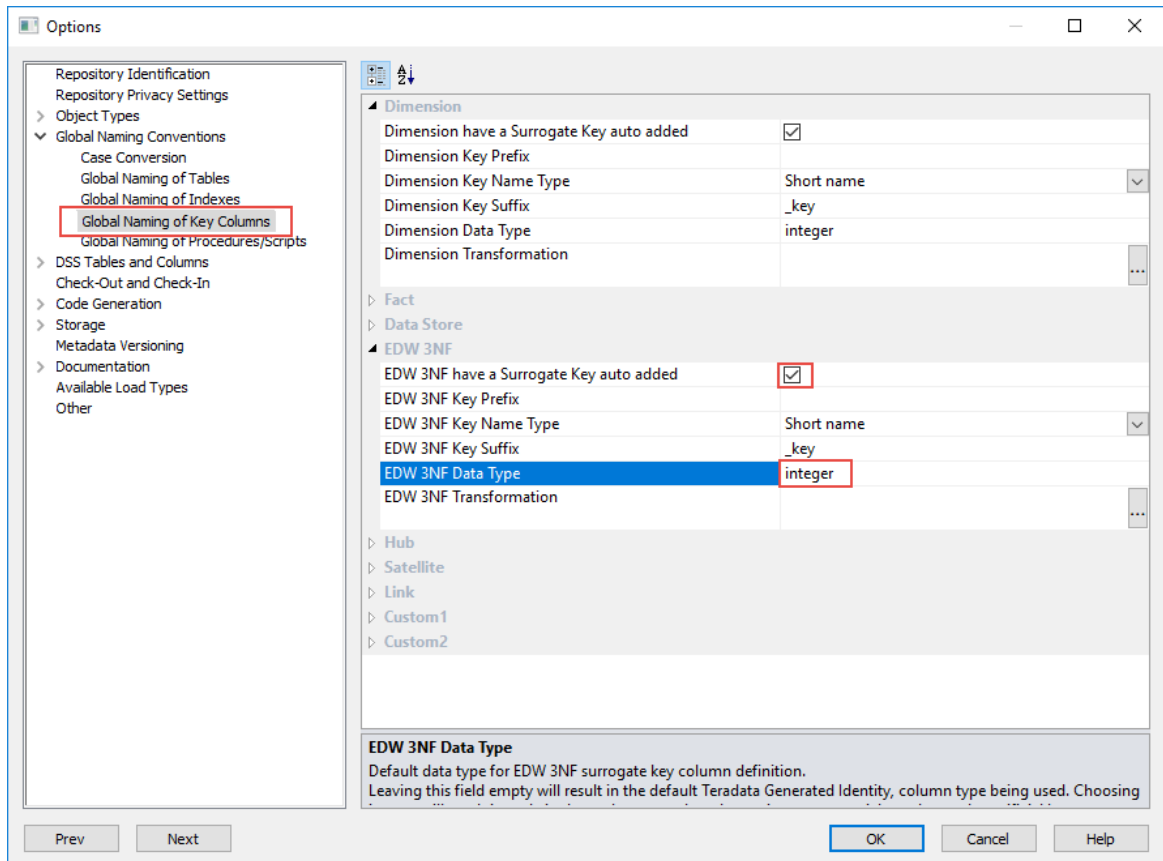
This automatically generated logic can be overwritten by defining a user specific logic on the **Dimension Transformation** field on the **Tools>Options** menu or in the transformation column of the artificial key.

To have an EDW 3NF table with a non identity column as a surrogate key, you can set the table **Data Type** to **integer** in the **Tools>Options** menu.

The old logic for dimensions can be retained, if an identity column is chosen as surrogate key.

Allowing for non identity surrogate keys on EDW 3NF tables:

- Go to **Tools>Options>Global Naming Conventions>Global Name of Key Columns**.
- Expand the **EDW 3NF** section.
- Set the **EDW 3NF table to have a Surrogate Key auto added** if you want a surrogate key added by default to all EDW 3NF tables.
- Set the **EDW 3NF Data Type** to be **integer** and click **OK**.
- If your tables have been created previously, you will have to **Recreate** the tables after you set this option in the **Tools** menu.



EDW 3NF TABLE COLUMN PROPERTIES

Each EDW 3NF table column has a set of associated properties. The definition of each property is described below:

If the **Column name** or **Data type** is changed for a column then the metadata will differ from the table as recorded in the database. Use the **Validate>Validate Table Create Status** menu option or the right-click menu to compare the metadata to the table in the database. A right-click menu option of **Alter table** is available when positioned on the table name after the validate has completed. This option will alter the database table to match the metadata definition.



TIP: If a database table's definition is changed in the metadata then the table will need to be altered in the database. Use the **Validate>Validate Table Create Status** to compare metadata definitions to physical database tables. The option also provides the ability to alter the database table, through a pop-up menu option from the validated table name.

A sample **Properties** screen is as follows:

The screenshot shows the 'Properties' dialog for the column 'code' in the table 'edw_customer'. The dialog is organized into several sections:

- General:** Table Name: edw_customer; Column Name: code; Column Title: code; Column Description: The unique code which identifies the customer.
- Physical Definition:** Column Order: 10; Data Type: decimal(6); Null Values Allowed: ; Default Value: ; Character Set: ; Format: ; Character Comparison/Sorting: ; Compress: .
- Meta Definition:** Numeric: ; Additive: ; Attribute: ; End User Layer Display: ; Business Key: ; Artificial Key: ; Key Type: Primary Business Key (A).
- Code Generation:** Zero Key Value: ;
- Source Details:** ;

At the bottom of the dialog, there are buttons for '<- Update', 'Update ->', 'OK', 'Cancel', and 'Help'. The 'OK' button is highlighted.

The two special update keys enable you to update the column and step either forward or backward to the next column's properties.

ALT-Left Arrow and **ALT-Right Arrow** can also be used instead of the two special update keys.

Table Name

Database-compliant name of the table that contains the column. [Read-only].

Column Name

Database-compliant name of the column. A good practice is to only use alphanumeric, and the underscore character. Changing this field alters the table's definition.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Title

Name that the business uses to refer to the column. It does not affect the physical table definition, but rather provides input to the documentation and to the view **ws_admin_v_dim_col** which can be used to assist in the population of a end user tool's end user layer. As such it is a free form entry and any characters are valid.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Description

This field contains the description for the column. It might contain information on where and how the column was acquired. For example if the column is sourced from multiple tables or is a composite or derived column then this definition would normally describe the process used to populate the column. This field is used in the documentation and is available via the view **ws_admin_v_dim_col** . This field is also stored as a comment against the column in the database.

Column Order

Numeric value that controls the relative order of columns in the database create statement. The lowest numbered column will appear first in the table. Although this affects the physical table definition no action will be taken unless the table is re-created in the database. The columns can be re-numbered based on the existing order by choosing the **Respace order number** pop-up menu option when positioned over any column in the table. This action will number the columns in increments of 10 starting at 10. In addition to a simple change of the order field, the column order can be changed by first displaying the columns in the middle pane and then using drag and drop to move the columns around. This drag and drop process will automatically renumber the columns as required.

Data Type

Database-compliant data type that must be valid for the target database. Typical Teradata databases often have integer, numeric(), varchar(), char(), date and timestamp data types. See the database documentation for a description of the data types available. Changing this field alters the table's definition.

Null Values Allowed

Determines whether the table column can hold NULL values or whether a value is always mandatory.

Default Value

Initial value that is assigned to the column when a row is inserted into the table but no value is specified for the column.

Character Set

Database-compliant table column character-set used for storage. Select Latin or Unicode.

Format

Database-compliant table column format. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. As such it is a free form entry and any characters are valid. Typically format masks are only used on numeric fields. Example: #,###0.00. It is not worth the effort of populating this field unless it can be utilized by the end user tools in use.

Character Comparison/Sorting

Determines how the column character values are treated for comparison and sorting operations. Choose from: case specific, not case specific, uppercase case specific or uppercase not case specific.

Compress

Indicates whether the table column values are compressed when stored.

Compress/Compress Value

Optional list of values to be compressed. By default, only NULL is compressed if no list of values is specified.

Numeric

Indicates whether the table column holds values that are numeric. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Additive

Indicates whether the table column holds values that are additive. This implies that the column can be summed when performing data grouping in a query. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Attribute

Indicates whether the table column holds values that are descriptive, and/or are used for grouping/summing. An attribute is defined as a column that is non factual. For example we may have an order number, or a invoice number stored in the fact table. Such columns are considered attributes, rather than facts. This checkbox is therefore normally only relevant for fact tables. This checkbox does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tools end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

End User Layer display

Indicates whether the table column is available/visible to end users. If set the documentation will include the column in the glossary and in the user documentation. It is also used to decide what columns appear in the view `ws_admin_v_dim_col`. Typically columns such as the artificial key would not be enabled for end user display.

Business Key

Indicates whether the column is part of the BUSINESS key, which is defined during the update procedure build. [Normally maintained automatically]. Multiple columns can form the primary business key. This indicator is set and cleared by WhereScape RED during the dimension update procedure generation process. This checkbox should not normally be altered.

Artificial Key

Indicates whether the column is the artificial key. Only one artificial key is supported. This indicator is set by WhereScape RED during the initial drag and drop creation of a table, and should not normally be altered.

Key Type

Key type that is assigned and used when generating the table's update procedure and indexes. [Normally maintained automatically]. It can be altered here, but this should only be done if the consequences are fully understood and tested. The supported values are:

Key type	Meaning
0	The artificial key. Set when the key is added during drag and drop table generation.
1	Component of all business keys. Indicates that this column is used as part of any business key.
A	Indicates that the column is part of the primary business key. Set whenever a business key is defined as part of an update procedure generation.
B-Z	Indicates that the column is part of a secondary business key. Only used during index generation and not normally set.

Source Table

Identifies the source table where the column's data comes from. This source table is normally a load table within the data warehouse. If the column was sourced from multiple tables, then the normal practice is to record one of the tables in this field and a comment listing all of the other tables in the **Source strategy** field. This field is used when generating a procedure to update the data store object. It is also used in the track back diagrams and in the documentation.

Source Column

Identifies the source column where the column's data comes from. Such a column is normally a load table column, which in turn may have been a transformation or the combination of multiple columns. This may also be a model table key where a model is being joined.

Transformation

See *EDW 3NF Table Column Transformations* (on page 411). [Read-only].

Join

Indicates whether the table column is used in a table join. [Normally maintained automatically but can be optionally changed to override the default join logic used in the generated update procedure]. The **Source table** and **Source column** fields will provide the other EDW 3NF table's side of the join. The options for this field are: False, True, Manual and Pre Join.

Setting this field to Manual changes the way the other EDW 3NF table is looked up during the update procedure build. It allows you to join the other EDW 3NF table manually in the Cursor mapping dialog (where the 'Where' clause is built). The usual dialog for matching the other EDW 3NF table's business key to a column or columns in the table is not displayed if this option is enabled.

Setting this field to Pre Join activates the **Join Source** field and allows you to select a table from the drop-down list.

Pre Join Source Table

Indicates the table from which the pre joined column was sourced. When the Join option is set to False, this field becomes inactive. When the Join option is set to True or Manual, this field is set to the current table name. When the Join option is set to Pre Join, then you can select the required table from the drop-down list.

EDW 3NF TABLE COLUMN TRANSFORMATIONS

Each EDW 3NF table column can have a transformation associated with it. The transformation will be included in the generated procedure and will be executed as part of the procedure update. The transformation must therefore be a valid SQL construct that can be included in a **Select** statement. For example we could have a transformation of 'load_order_line.qty * 0.125' to calculate a tax column of 12.5%.

Click the **Transformation** tab to enter a transformation.
The transformation screen is displayed as follows:

Note: Transformations are only put into effect when the procedure is re-generated.

Microsoft Analysis Services 2005+ Tabular Mode Tables: For Tabular Mode table column transformations, **Default DAX** is the only applicable Function Set for **after load** transformations.

See *Transformations* (on page 593) for more details.

CHAPTER 14

DATA VAULTS

The Data Vault system is an alternative approach to modelling an enterprise data warehouse that has been gaining popularity among organizations.

The Data Vault data warehouse architecture was invented by Dan Linstedt to provide an alternative to the traditional data warehouse modelling approach that includes developing 3rd Normal Form (3NF) type models or dimensional star schema models. The data vault methodology seeks to improve the efficiency of data ingestion and the flexibility of structure changes. For more information about Data Vaults, please refer to Dan Linstedt's website (<http://danlinstedt.com/solutions-2/data-vault-basics/>).

WhereScape RED has been enhanced to expand its current Data Vault functionality and provide improved automation for creating and managing Data Vault objects in WhereScape RED managed Data Warehouses. The enhancement includes the following:

- New DSS columns for Load tables
- New Wizard for Hash key generation
- New Wizard for building Hub, Link and Satellite tables
- New Templates for Procedure generation

All these enhancements are designed to be compliant with Data Vault 2.0 standards and are described in the succeeding sections below.

Note that previous releases of WhereScape RED used a workflow for Data Vault objects that is similar to the workflow for creating Data Store Objects.

The Hub, Link and Satellite tables are based on standard Load or Stage tables (that do not include the hash key column type flags) then WhereScape RED reverts to this behaviour and the resulting procedures are generated by internal WhereScape RED automation and not via templates.

If the above legacy method of creating Data Vault Objects is required, please refer to the Data Store Objects chapter of this User Guide. You can also refer to the previous version of this User Guide, for details on the previous process used to import Data Vault objects into WhereScape RED.

For more information on Data Vault design, refer to *Building a Scalable Data Warehouse With Data Vault 2.0* by Daniel Linstedt and Michael Olschimke.

IN THIS CHAPTER

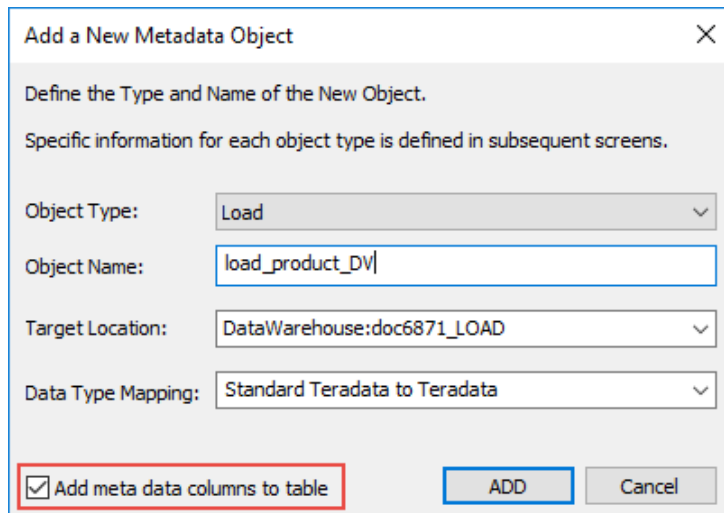
Data Vault Functions and Features	414
Building Data Vault Objects	425

DATA VAULT FUNCTIONS AND FEATURES

The following describe the WhereScape RED functions and features that are used for building Data Vault objects (Hub, Link and Satellite) to create a Data Vault model.

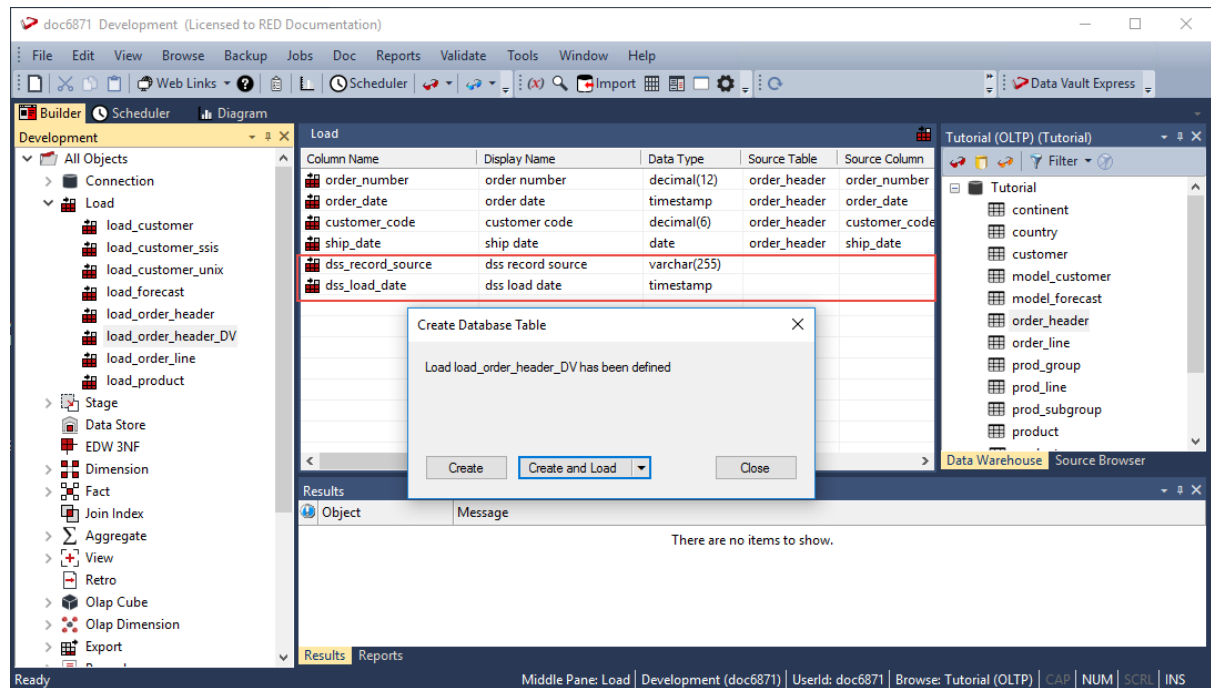
LOAD TABLE META DATA COLUMNS

The option to add default meta data columns to a new Load table object in WhereScape RED.



The screenshot shows a dialog box titled "Add a New Metadata Object" with a close button (X) in the top right corner. Below the title, there is a subtitle "Define the Type and Name of the New Object." and a note: "Specific information for each object type is defined in subsequent screens." The dialog contains four dropdown menus: "Object Type" set to "Load", "Object Name" containing "load_product_DV", "Target Location" set to "DataWarehouse:doc6871_LOAD", and "Data Type Mapping" set to "Standard Teradata to Teradata". At the bottom, there is a checkbox labeled "Add meta data columns to table" which is checked and highlighted with a red rectangular box. To the right of the checkbox are two buttons: "ADD" (highlighted with a blue rectangular box) and "Cancel".

If the option **Add meta data columns to table** is selected, two DSS columns (**dss_record_source** and **dss_load_date**) are included in the meta data for the table and are populated by transformations. Note that these two DSS columns could equally be applied to other Load tables not used in a Data Vault model but are particularly important to comply with the Data Vault standards.



Note: Custom aliases for the `dss_load_date` and `dss_record_source` columns can be defined in the **Tools>Options>DSS Tables and Columns>DSS Columns to Include and Naming** window. If custom DSS aliases are defined then the aliases are used for the column names that are added to the new load tables. Refer to **DSS Columns** (on page 103) for more details.

DATA VAULT STAGE TABLE

WhereScape RED provides the option to create **Data Vault Stage** objects that are used to define the source columns for the Hub, Link and Change hash key columns.

The screenshot shows the 'Stage stage_product_DV' dialog box. The 'Table Type' dropdown menu is open, displaying the following options: Data Vault Stage, Stage, Data Vault Stage, Permanent Stage, and Work Table. The 'Data Vault Stage' option is highlighted. The dialog also includes fields for Table Name, Unique Short Name, Description, Update Procedure, Custom Procedure, and Timestamps.

Metadata Structure	Database	Database Altered:
2017-05-28 23:09:50.730000	2017-05-28 23:32:45.030000	2017-05-28 23:32:45.030000

The **Data Vault Stage** option can be selected from **Table Type** drop-down of the **Stage Table Properties** screen. This object type is created via a Wizard which is described in the section **Creating Data Vault Stage Tables** (on page 427) below.

HASH KEY GENERATION WIZARD

This Wizard is launched when building a **Stage** table with **Table Type** of **Data Vault Stage**. It enables you to specify the source columns to be used in defining the **Hub**, **Link** and **Change** hash key columns.

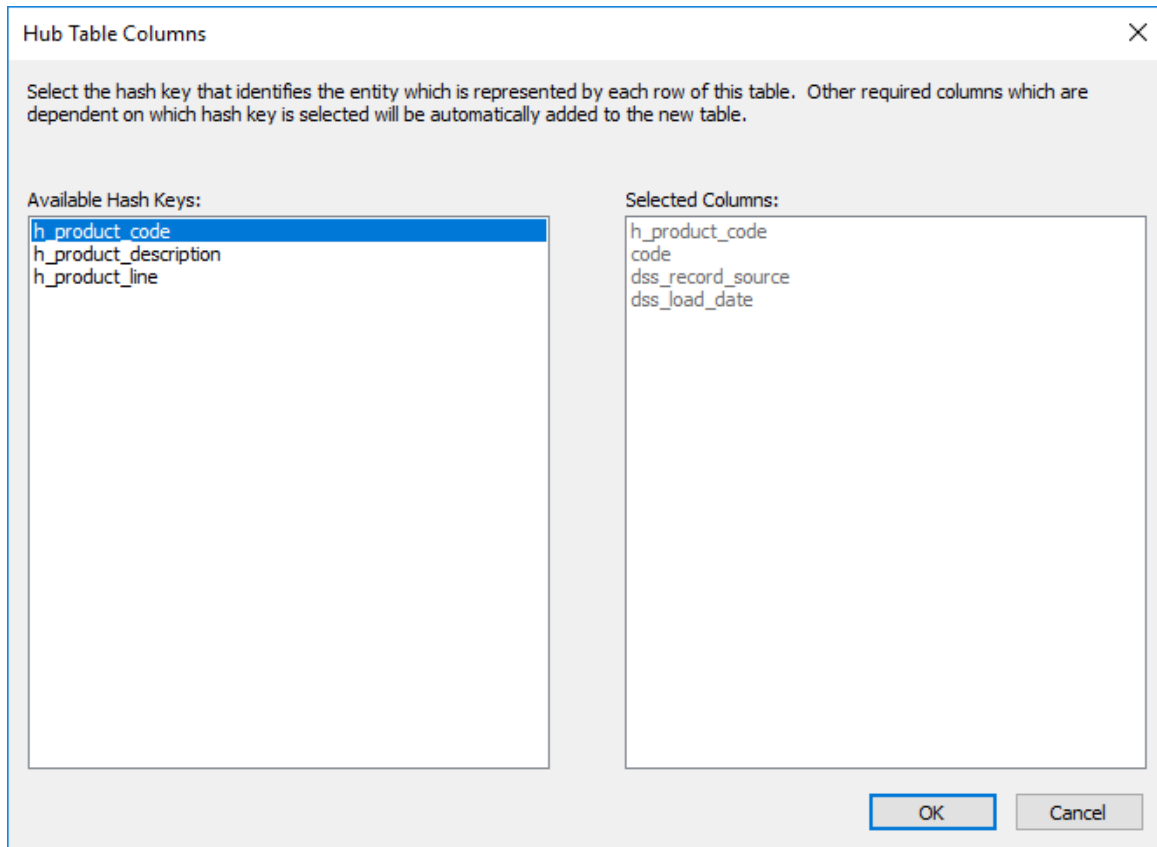
The generated Hash Keys are used to build the **Hub**, **Link** and **Satellite** objects in WhereScapeRED. The detailed steps for using this Wizard is described in the succeeding section, **Creating Data Vault Stage Tables** (on page 427).

Notes:

1. Creating a **Change Hash Key** via the Wizard in WhereScape RED is limited to one key per **Data Vault Stage** table. This limitation does not apply for Data Vault models imported from WhereScape 3D.
 2. Custom aliases for the **dss_change_hash** column can be defined in the **Tools>Options>DSS Tables and Columns>DSS Columns to Include and Naming** window. If custom DSS aliases are defined then the aliases are used for the column names that are added to the new Load tables. Refer to **DSS Columns** (on page 103) for more details.
 3. The data type defined for the **dss_change_hash** column in **Tools>Options>DSS Tables and Columns>DSS Columns for Custom Targets** window is applied when creating a hub hash key or a link hash key for tables in custom database targets. Refer to **DSS Columns for Custom Targets** (on page 104) for details.
 4. Refer to **Changing the Data Vault Hash Key Function in WhereScape RED 6.9.1.0 and above** (on page 455), if you need to amend the Hash key function on the shipped WhereScape RED **Data Vault Templates** (on page 418).
-

HUB, LINK AND SATELLITE CREATION WIZARD

This Wizard is launched when building a Data Vault object (**Hub**, **Link** or **Satellite** table) by dragging and dropping the source **Data Vault Stage** table from the right pane to the middle pane.



The screenshot shows a dialog box titled "Hub Table Columns" with a close button (X) in the top right corner. The dialog contains the following text:

Select the hash key that identifies the entity which is represented by each row of this table. Other required columns which are dependent on which hash key is selected will be automatically added to the new table.

Available Hash Keys:

- h_product_code
- h_product_description
- h_product_line

Selected Columns:

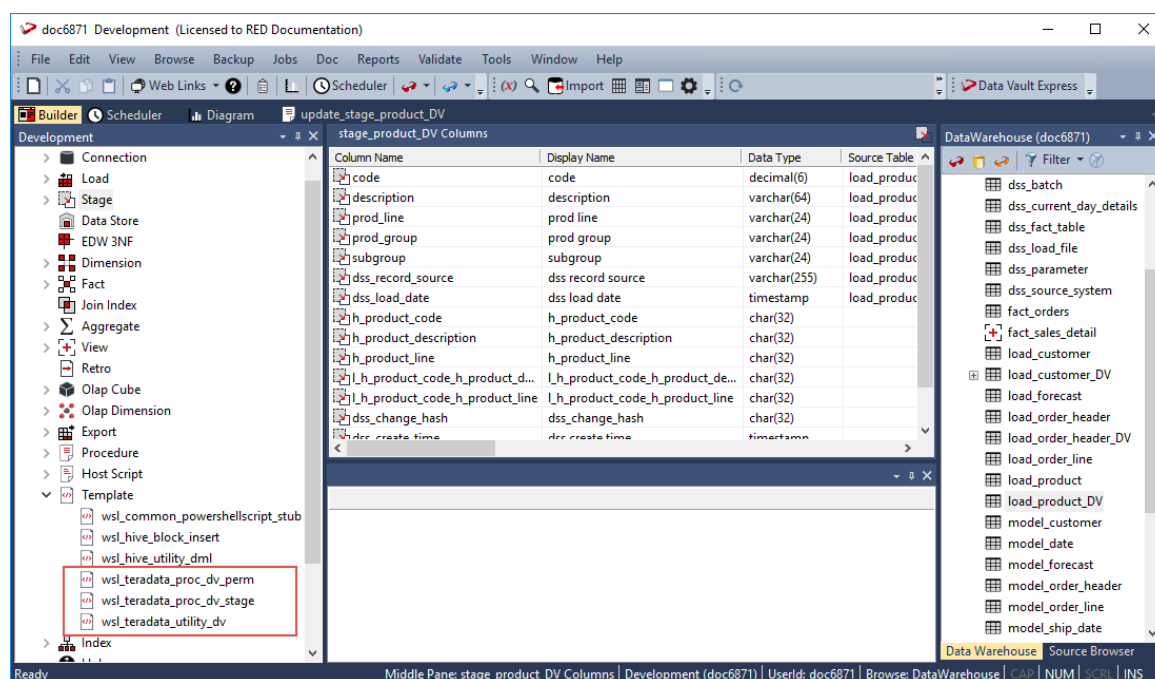
- h_product_code
- code
- dss_record_source
- dss_load_date

At the bottom right of the dialog are two buttons: "OK" and "Cancel".

The detailed steps for using this Wizard is described in the succeeding section, *Creating the Hub, Link and Satellite Tables* (on page 440).

DATA VAULT TEMPLATES

Templates are used to generate update procedures for **Data Vault** objects. Users must select a template to use when generating the update procedure for **Data Vault** objects created in WhereScape RED.



The **Teradata Data Vault** templates are described below:

- **wsl_teradata_proc_dv_stage** – this template creates a Teradata procedure for updating WhereScapeRED **Data Vault Stage** tables.
- **wsl_teradata_proc_dv_perm** – this template creates a Teradata procedure for updating WhereScapeRED **Data Vault** objects (**Hub, Link** and **Satellite** tables).
- **wsl_teradata_utility_dv** – this utility template contains generic Teradata macros that are used by the other two templates above.

The above templates are available in WhereScape RED version 6.9.1.0 and above. If the templates are not visible in **Template** objects list after installing/upgrading WhereScape RED, use the WhereScape RED Setup Administrator to Validate the MetaData Repository.

Note: Templates for SQL Server and Oracle are also supplied by WhereScape. For other database types, users need to create/provide their own templates.

Important: The above templates have changed in WhereScape RED version 8.0.1.0. The Data Vault hash key functions for Teradata and Oracle repositories have been changed in the templates (`wsl_teradata_utility_dv` and `wsl_oracle_utility_dv`). Data Vault customers running on Teradata or Oracle who were using these templates from the RED 6.9.1.0 release are advised to retain that version of the template, unless they migrate any key generated using the previous hash key function.

Refer to *Changing the Data Vault Hash Key Function in WhereScape RED 6.9.1.0 and above* (on page 455) if you need to amend the Hash key function on the shipped WhereScape RED Data Vault Templates.

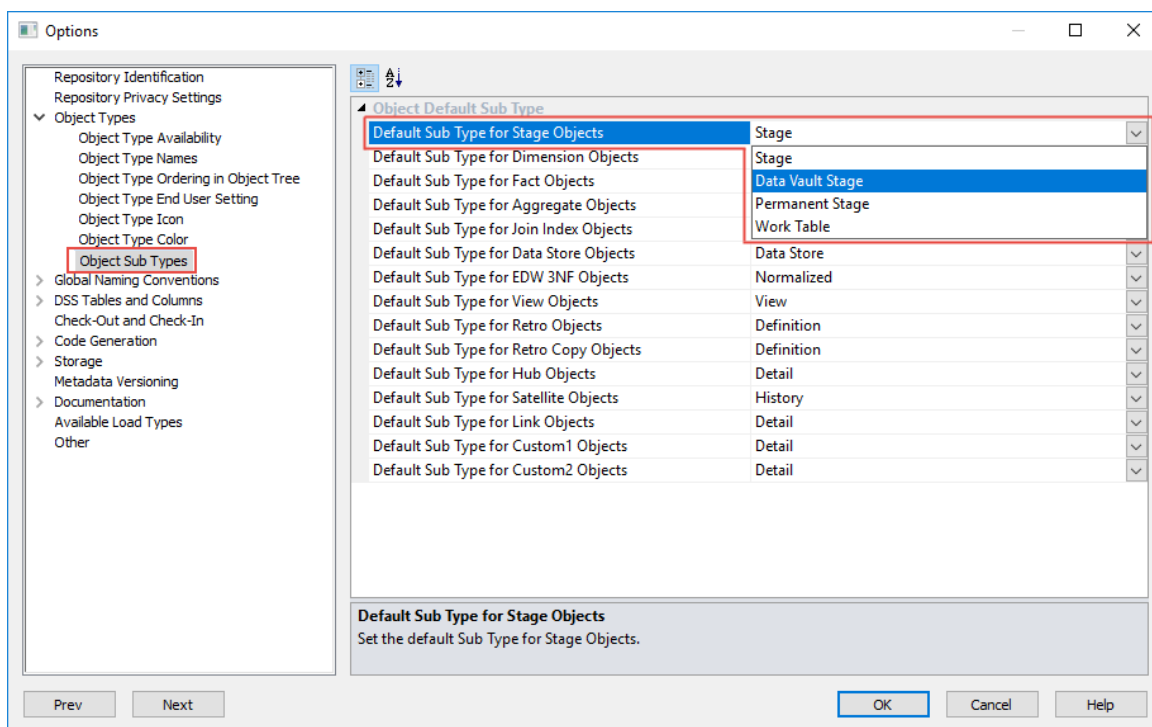
A Wizard to generate update procedures via templates is used to populate **Data Vault** tables. The detailed steps for using this Wizard is described in the succeeding section, *Building Data Vault Objects* (on page 425).

DATA VAULT SETTINGS

Settings for Data Vault objects can be configured from the **Tools > Options** screen.

OBJECT TYPES SETTINGS:

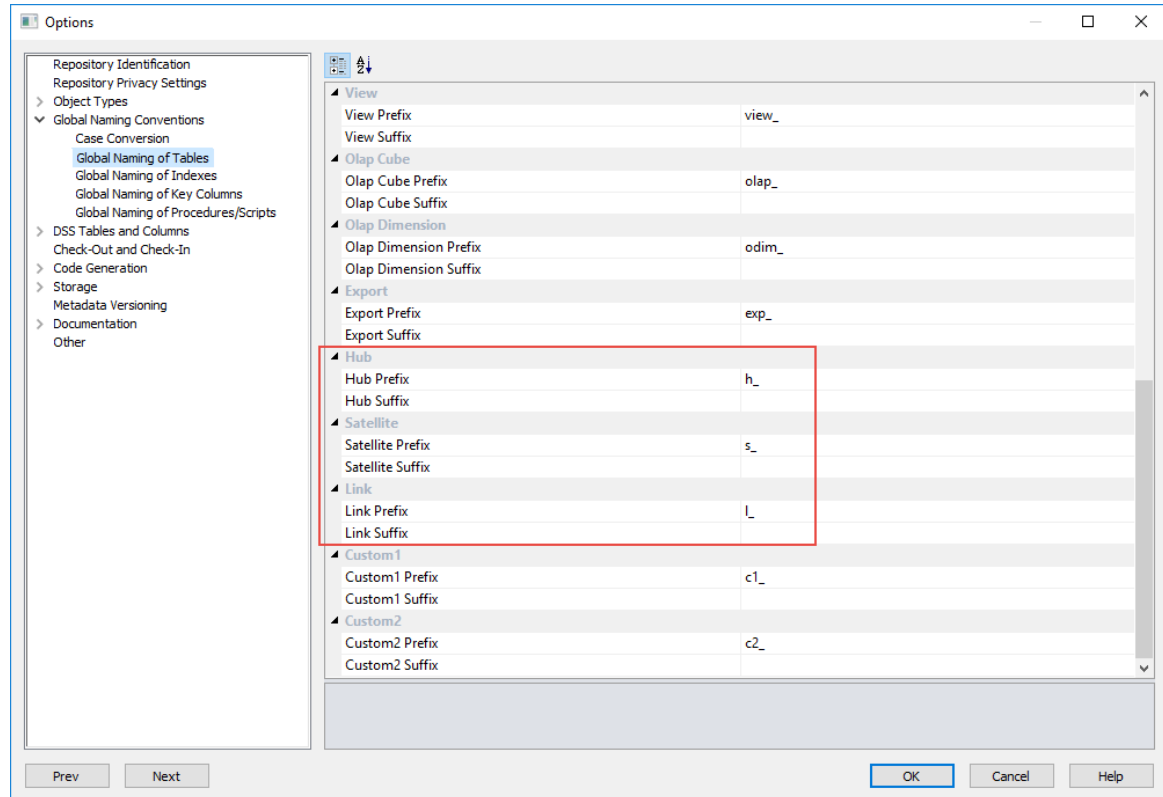
The **Default Sub Type for Stage Table Objects** drop-down includes the **Data Vault Stage** option.



Configure this setting, if you want to set **Data Vault Stage** to be the default **Table Type** in the **Stage Table Properties** screen.

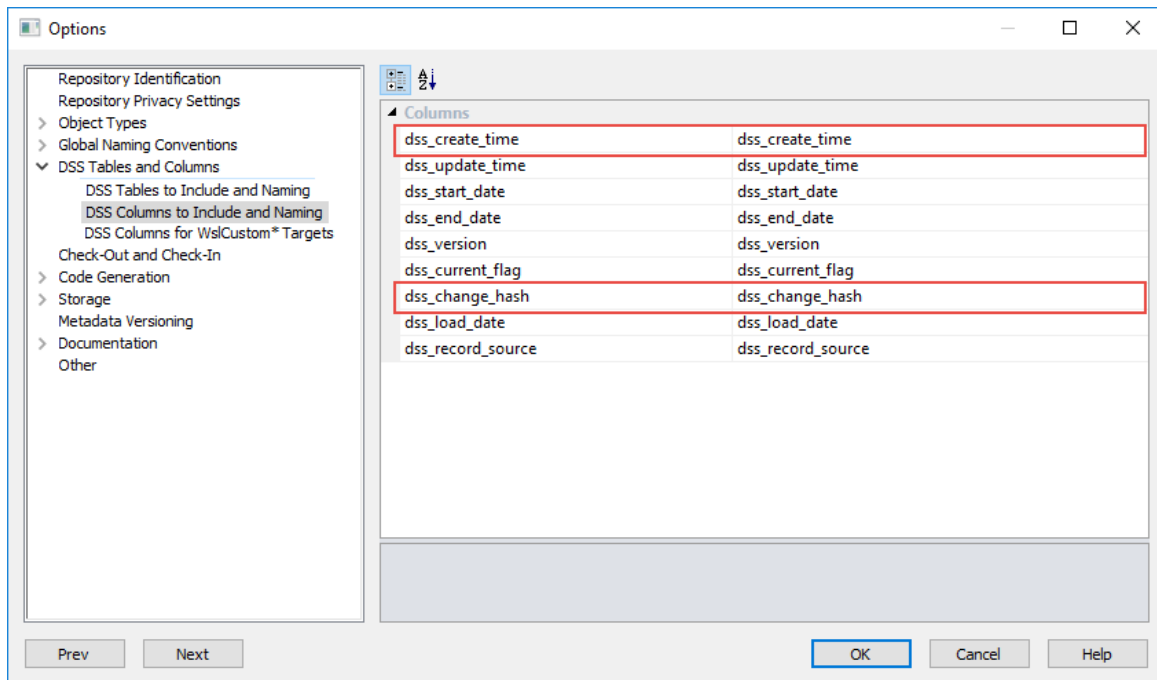
GLOBAL NAMING CONVENTIONS SETTINGS:

The **Global Naming of Tables** setting for **Hub**, **Satellite** and **Link** tables have been set to comply with the recommended standard naming convention for **Data Vault** tables. You can edit this setting to suit your requirements.



DSS TABLES AND COLUMNS SETTINGS:

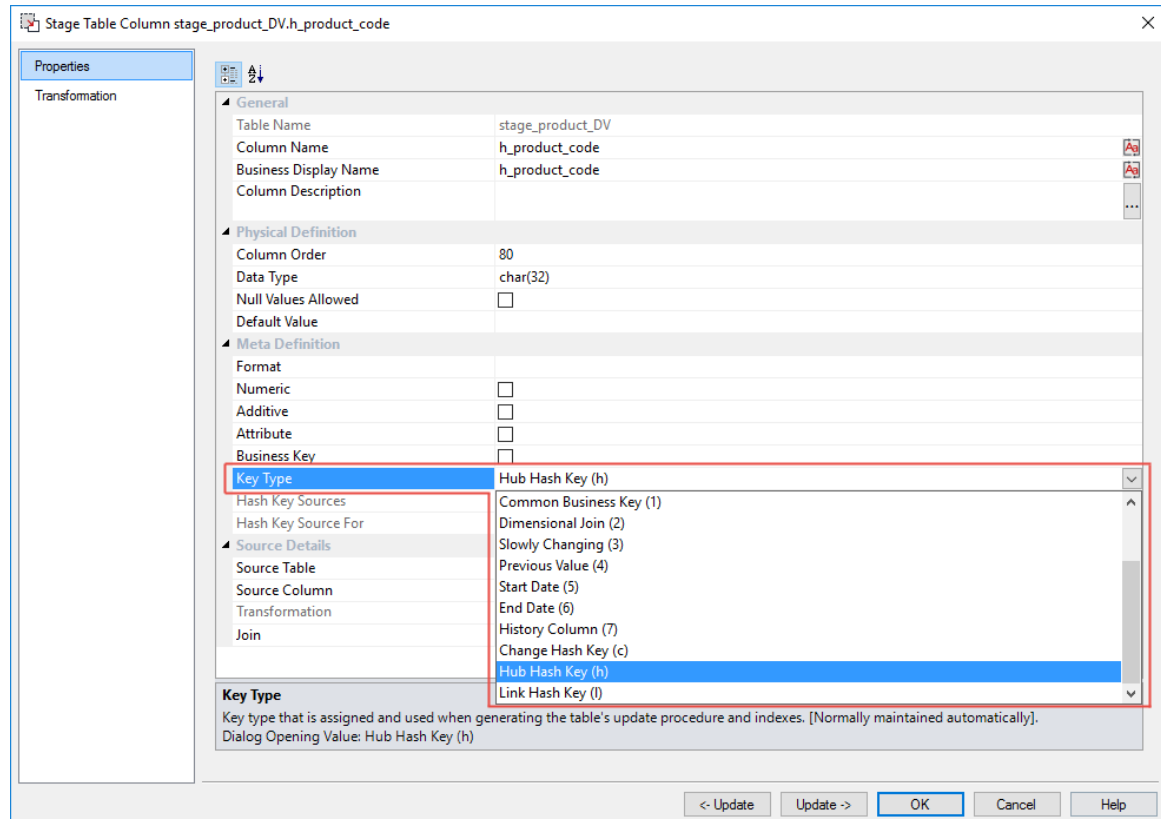
The **DSS Columns to Include and Naming** setting includes two additional columns which is described below:



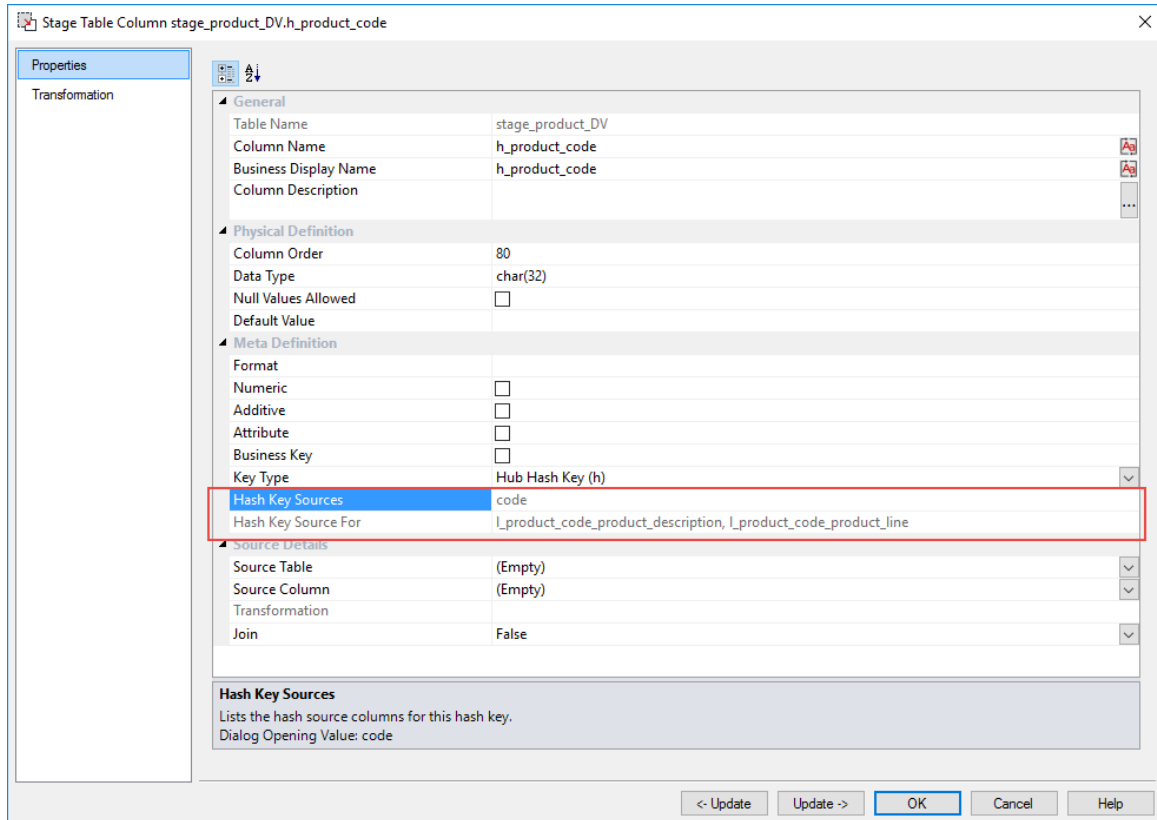
- **dss_create_time** – this column is added to all Stage, ODS, Normalized, Dimension, Fact and Aggregate tables for information purposes only. Leave blank to deactivate.
- **dss_change_hash** – this column is used to identify the differences in the descriptive columns of a Satellite table which is required for generating the change hash key for a Satellite table.

TABLE COLUMN PROPERTIES

The **Key Type** field drop-down include options for **Data Vault** hash keys, e.g. **Change Hash Key (c)**, **Hub Hash Key (h)** and **Link Hash Key (l)**.



Hash Key source information are also displayed for these key types.



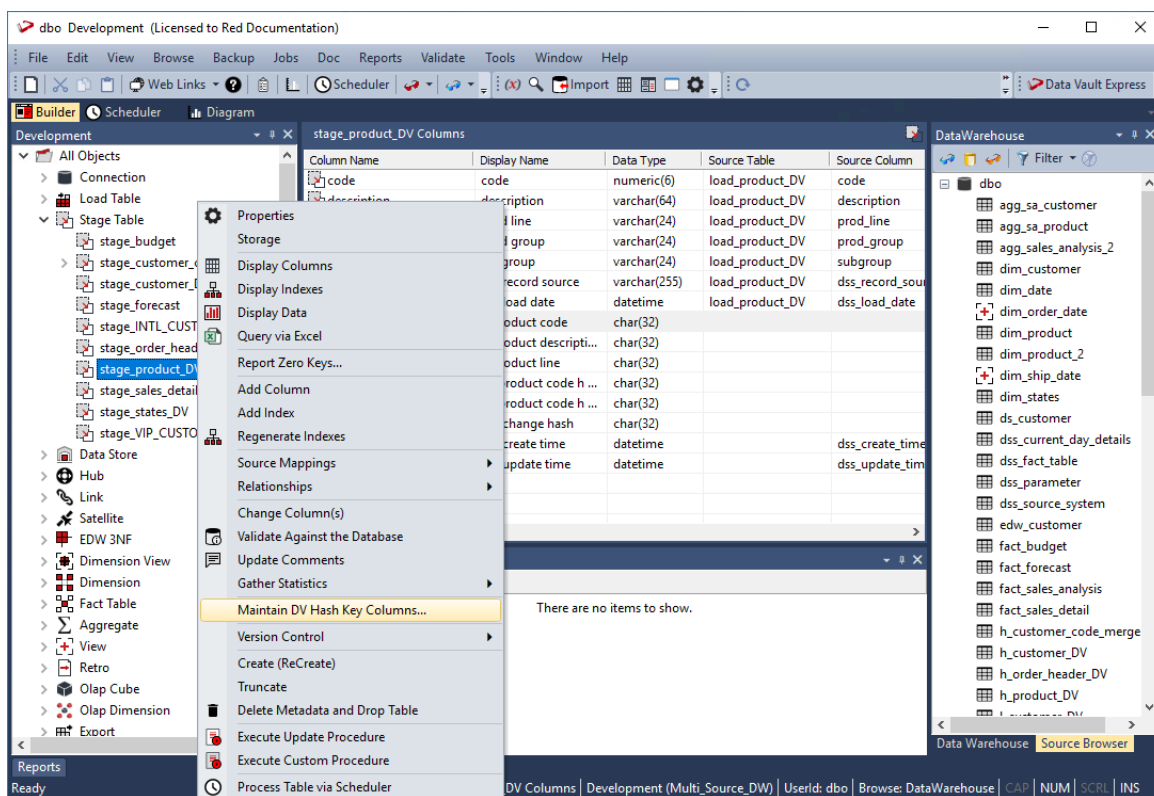
- **Hash Key Sources** – displays the source columns that are used to generate the selected **Hub**, **Link** or **Change** hash key.
- **Hash Key Source For** – displays the hash keys columns that use the displayed hash key sources.

The hash key generation Wizard enables you to specify the source columns to be used in defining the **Hub**, **Link** and **Change** hash key columns. The detailed steps for using this Wizard is described in the succeeding section, *Creating Data Vault Stage Tables* (on page 427).

Once the hash key columns have been defined, another Wizard is used to generate the procedure to populate the columns. The detailed steps for using this Wizard is described in the section, *Generating Update Procedures for the Data Vault Stage Table* (on page 436).

MAINTAIN HASH KEY COLUMNS

The context menu for **Stage Table** objects, listed in the left pane provides an option for maintaining **Data Vault** hash key columns.



You can review the composition of existing hash keys for a **Data Vault Stage** table (Hub, Link and Satellite) and create additional hash keys by selecting the **Maintain DV Hash Key Columns** option from the selected **Stage Table**'s context menu. This launches the hash key generation Wizard which enables you to maintain the source columns defined for the hash keys.

Note: To remove or change a hash key column, you need to delete it first, e.g. right click the column listed in the middle pane and then select **Delete Column** from the context menu.

BUILDING DATA VAULT OBJECTS

To build Data Vault objects (Hub, Link and Satellite) in WhereScape RED, involves the following procedures.

- 1 Creating Load Tables with the required DSS columns.
- 2 Creating Data Vault Stage tables.
- 3 Generating update procedures for the Stage table via templates.
- 4 Creating the Hub, Link and Satellite tables.
- 5 Generating update procedures for the Hub, Link and Satellite tables via templates.

The detailed steps for each procedure are outlined in the following sections.

CREATING LOAD TABLES

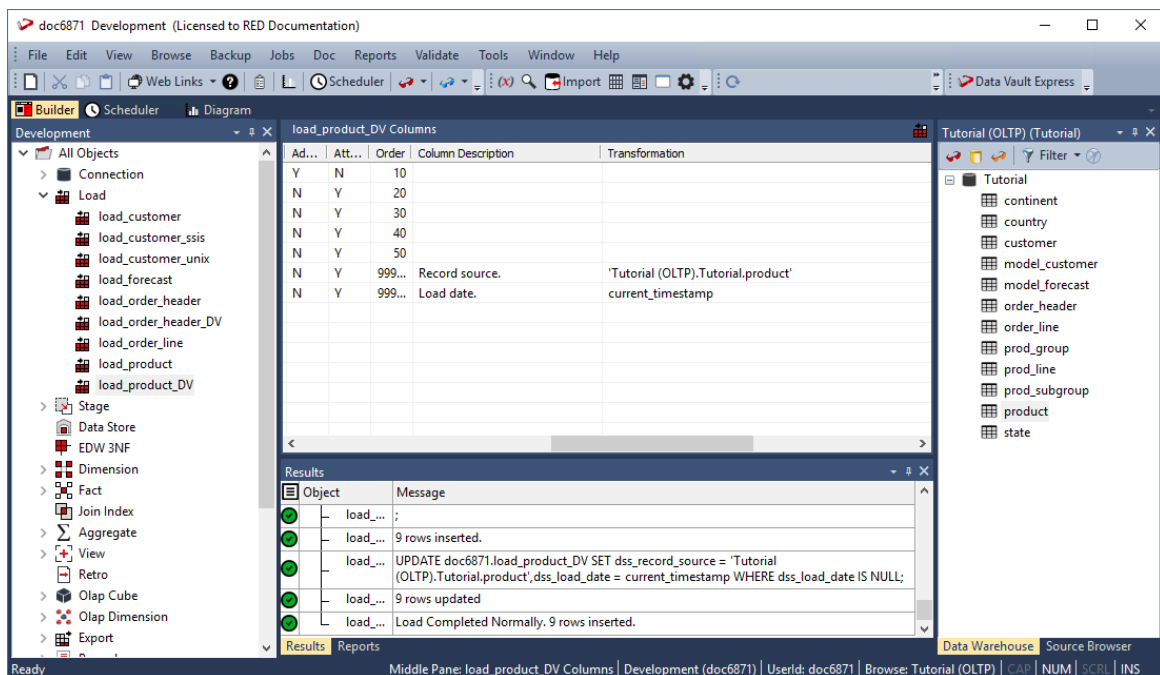
The following describe the steps for creating a **Load** table:

- 1 Browse to the source system connection required (Browse>Source Tables).
- 2 Double-click the **Load Table** object group in the left pane, the middle pane displays a list of existing Load tables.
- 3 Click the source table from the right pane and drag it to the middle pane. You need to create the **Load** table with the required DSS columns—the option to add default meta data columns to the load table must be selected:

When the build table is performed, the load table created has the two additional columns, which are populated by transformations:

- **dss_record_source** – the connection or source for the load table.
- **dss_load_date** – the date when the data was loaded to the table. This is updated every time a load operation is performed.

These DSS columns added, include column description and transformation information.



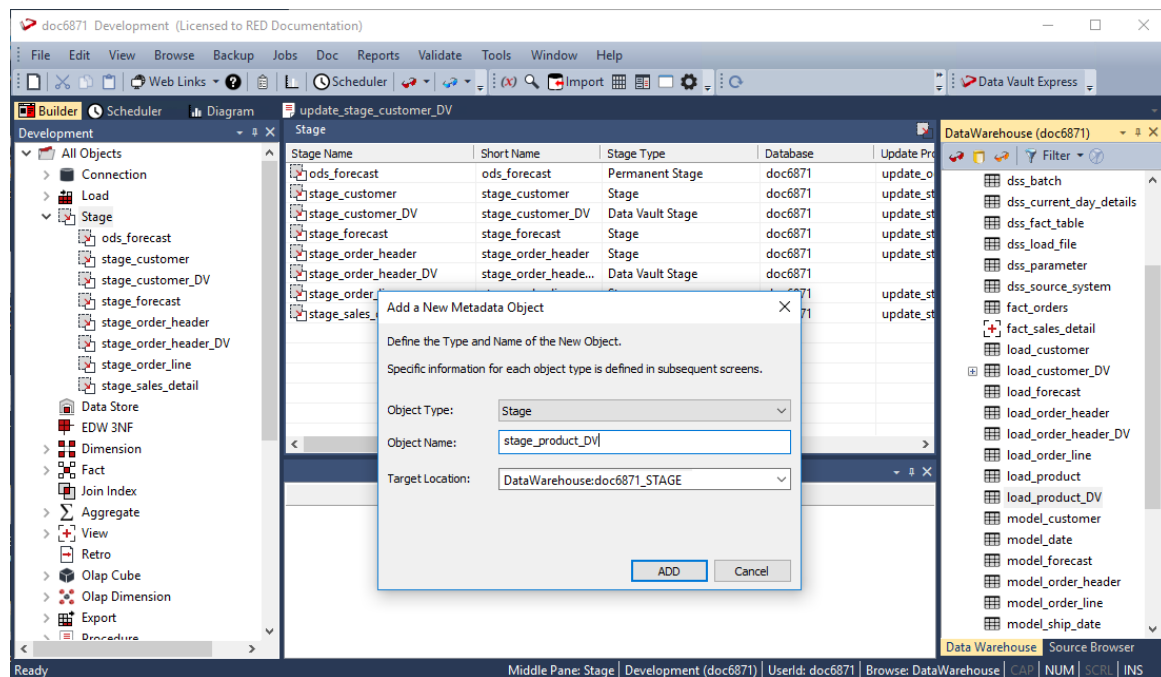
Notes:

- Custom aliases for the **dss_load_date** and **dss_record_source** columns can be defined in the **Tools>Options>DSS Tables and Columns>DSS Columns to Include and Naming** window. If custom DSS aliases are defined then the aliases are used for the column names that are added to the new load tables. Refer to **DSS Columns** (on page 103) for more details.
 - The **dss_record_source** column transformation is set to the name of the source connection when the table is first created. If the source connection changes, the transformation value does not update automatically. You can manually update the transformation—refer to **Transformations** (on page 593) for more details.
-

CREATING DATA VAULT STAGE TABLES

The following describe the steps for creating a **Data Vault Stage** table:

- 1 Browse to the Data Warehouse connection to create the **Stage** table.
- 2 Double-click the **Stage Table** object group in the left pane, the middle pane displays a list of existing **Stage** tables.
- 3 Click the source **Load** table from the right pane and drag it to the middle pane. The selected **Load** table must have the required DSS columns (**dss_record_source** and **dss_load_date**).



- 4 The **Add a New Metadata Object** screen appears and classifies the new object as a **Stage** table. It provides a default name based on the **Load** table name. Accept the name or enter a new name for the **Stage** table and click **ADD** to continue.
- 5 On the **Table Properties** screen, select the **Data Vault Stage** option from the **Table Type** drop-down.

Stage stage_product_DV

Properties

Storage

Override Create DDL

Extended Properties

Notes

Table Name: stage_product_DV

Unique Short Name: (maximum 22 characters) stage_product_DV

Description:

Update Procedure: update_stage_product_DV Edit Rebuild Regenerate

Template: wsl_teradata_proc_dv_stage

Custom Procedure: (None) Rebuild Regenerate

Table Type: Data Vault Stage

Stage

Data Vault Stage

Permanent Stage

Work Table

Timestamps

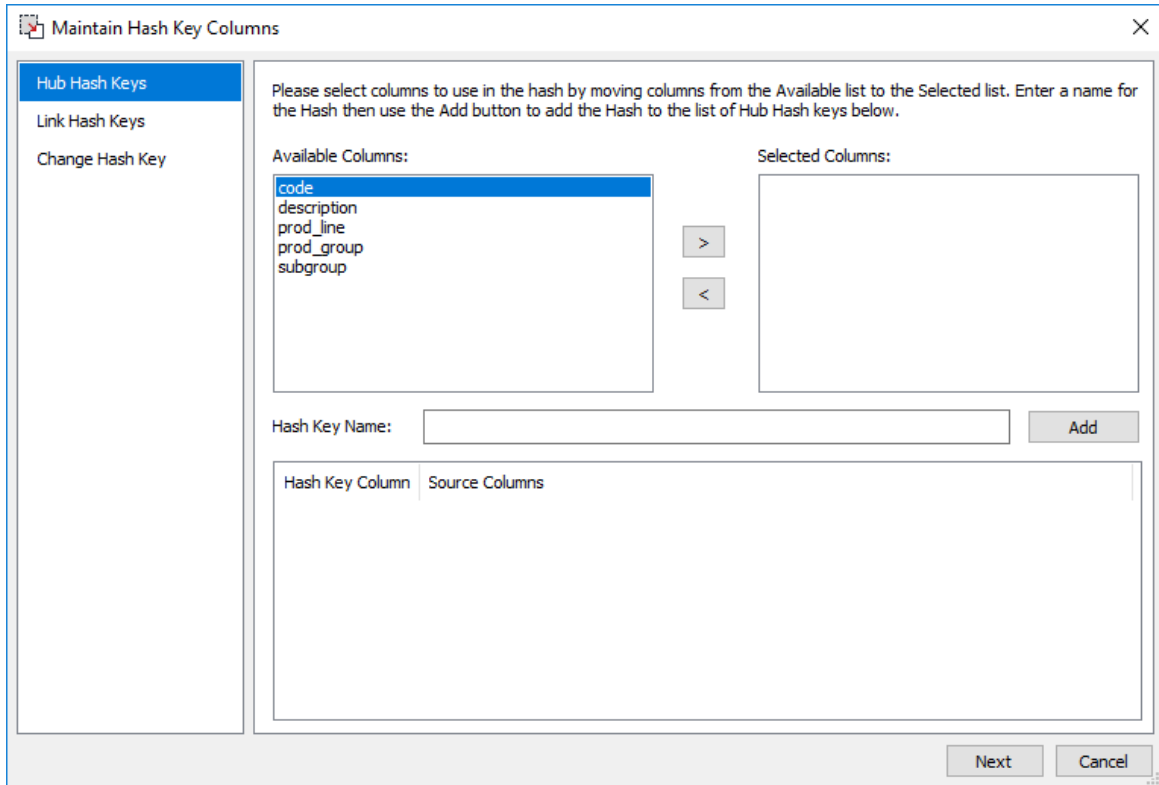
Metadata Structure	Database	Database Altered:
2017-05-28 23:09:50.730000	2017-05-28 23:32:45.030000	2017-05-28 23:32:45.030000

OK Cancel Help

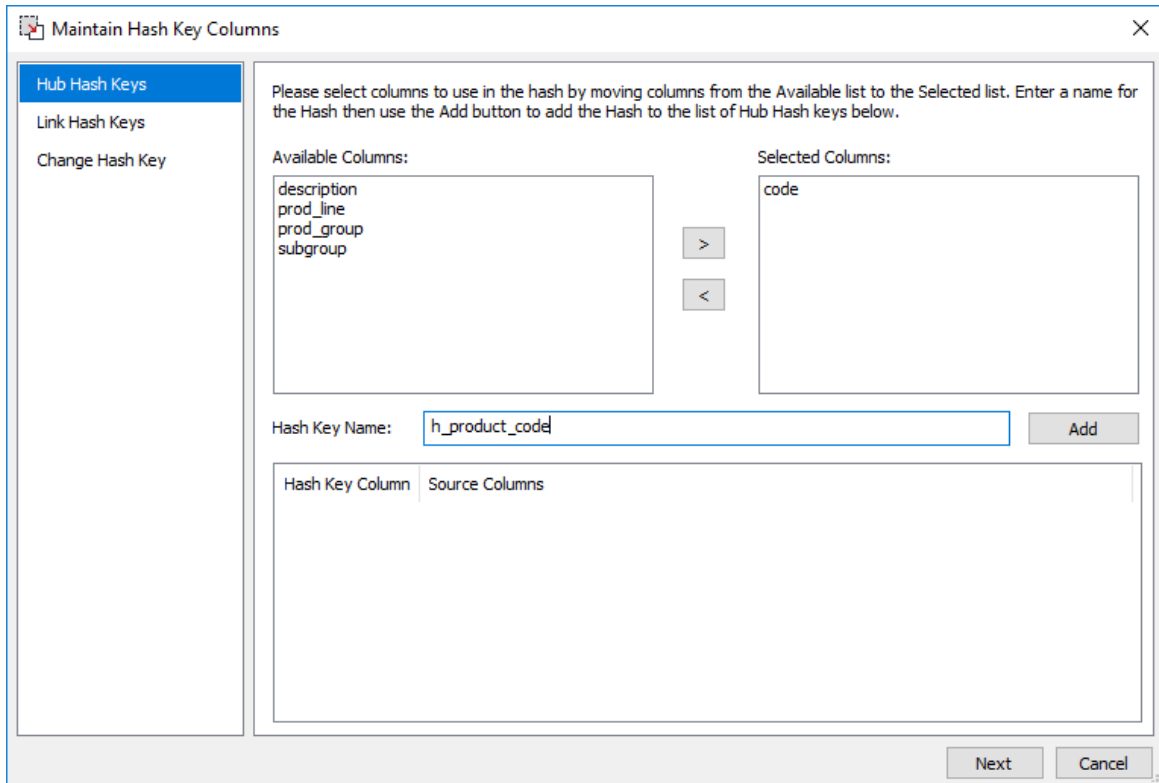
Notes:

You can set this table type to be the default by configuring the **Default Sub Type for Stage Table Objects** setting in the **Tools > Options > Object Types > Object Sub Types** screen. RED displays the name of the previously used update procedure template by default, below the **Update Procedure** drop-down.

- Click **OK** on the **Table Properties** screen to launch the Wizard that enables you to define the source columns for the **Hub**, **Link** and **Change** hash key columns.



- Select the source column(s) to use in defining the first Hub hash key. The **Hash Key Name** is formed based on the prefix (defined in the **Tools>Options>Global Naming Conventions** settings) and the source column(s) name. You can manually amend the name if required.



- 8 Click **Add** to create the first Hub hash key. Repeat the same steps as required. Once all the required **Hub Hash Keys** are created, click **Next** to progress to the **Link** hash key generation screen.

Maintain Hash Key Columns

Hub Hash Keys
Link Hash Keys
Change Hash Key

Please select columns to use in the hash by moving columns from the Available list to the Selected list. Enter a name for the Hash then use the Add button to add the Hash to the list of Hub Hash keys below.

Available Columns: code, description, prod_line, prod_group, subgroup

Selected Columns:

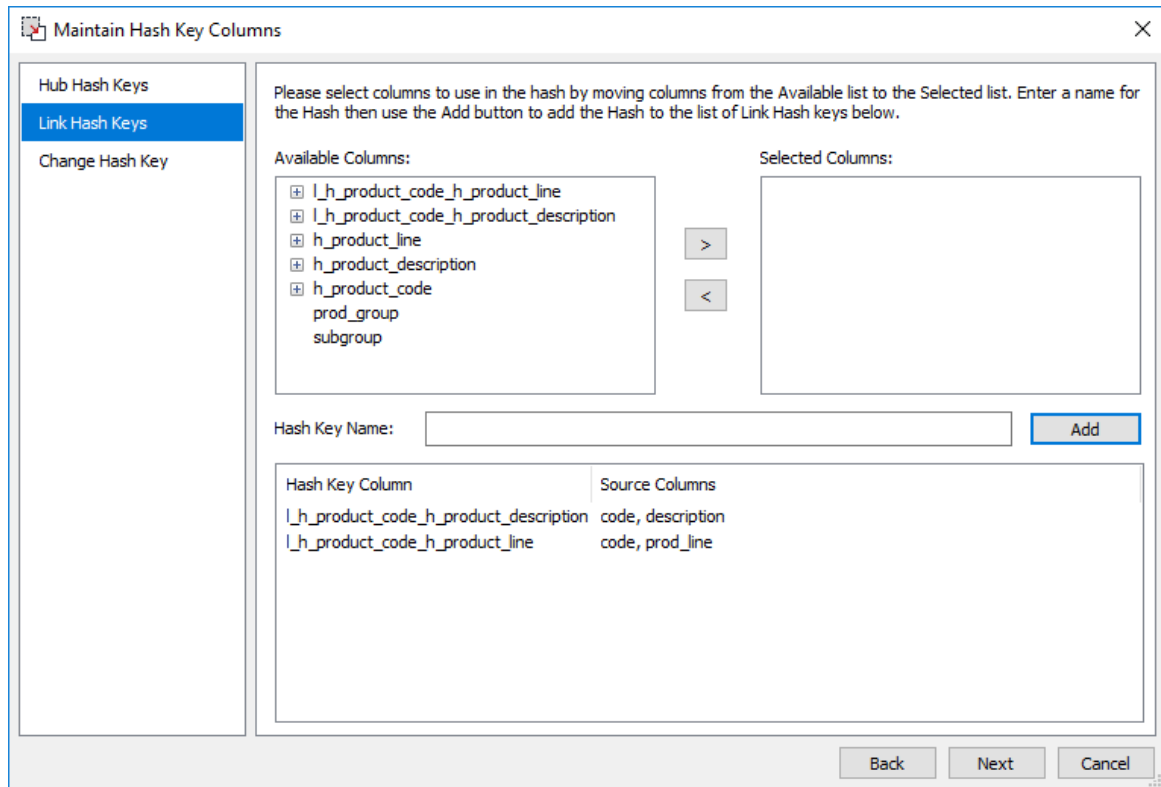
Hash Key Name: **Add**

Hash Key Column	Source Columns
h_product_code	code
h_product_description	description
h_product_line	prod_line

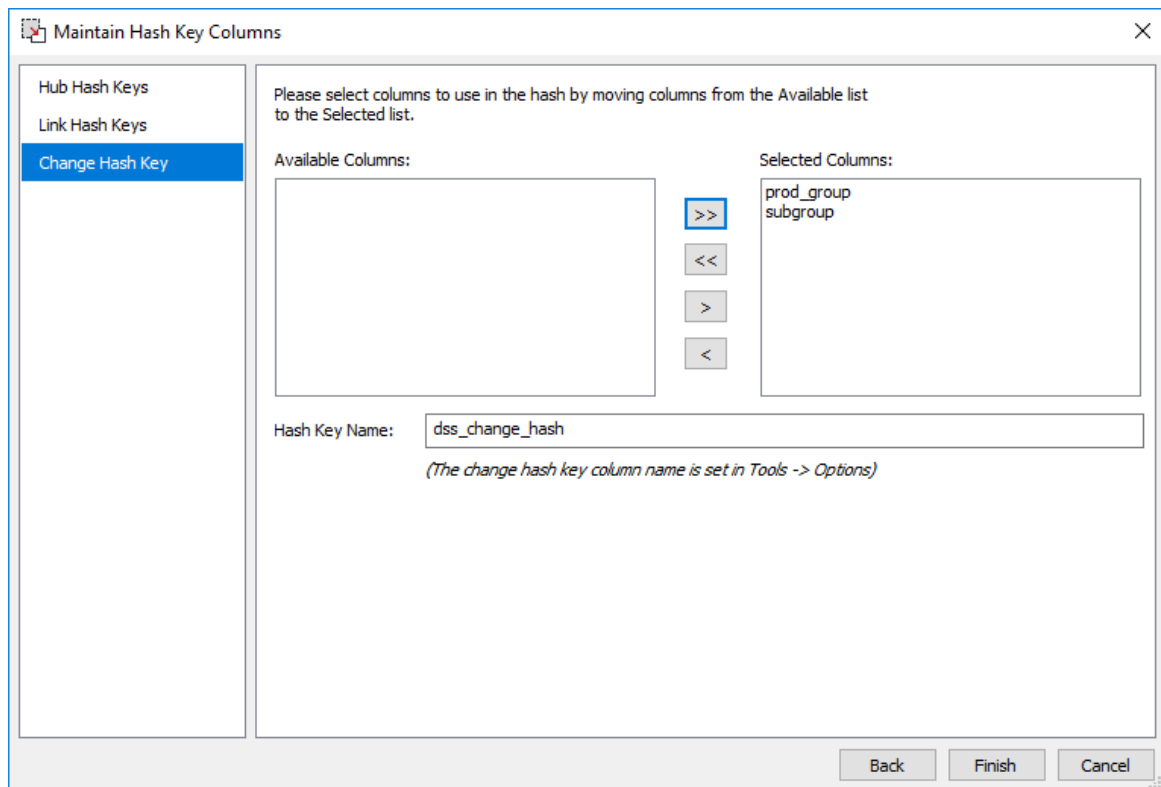
Next **Cancel**

- 9 Creating the **Link Hash Keys** involves the same process, follow the previous steps (6 and 7) to select *multiple* source columns to combine to create the **Link** hash key. Click **Add** to create the first Link hash key. Repeat the same steps for any subsequent keys. Once all the required

Link Hash Keys are created, click **Next** to progress to the **Change** hash key generation screen.



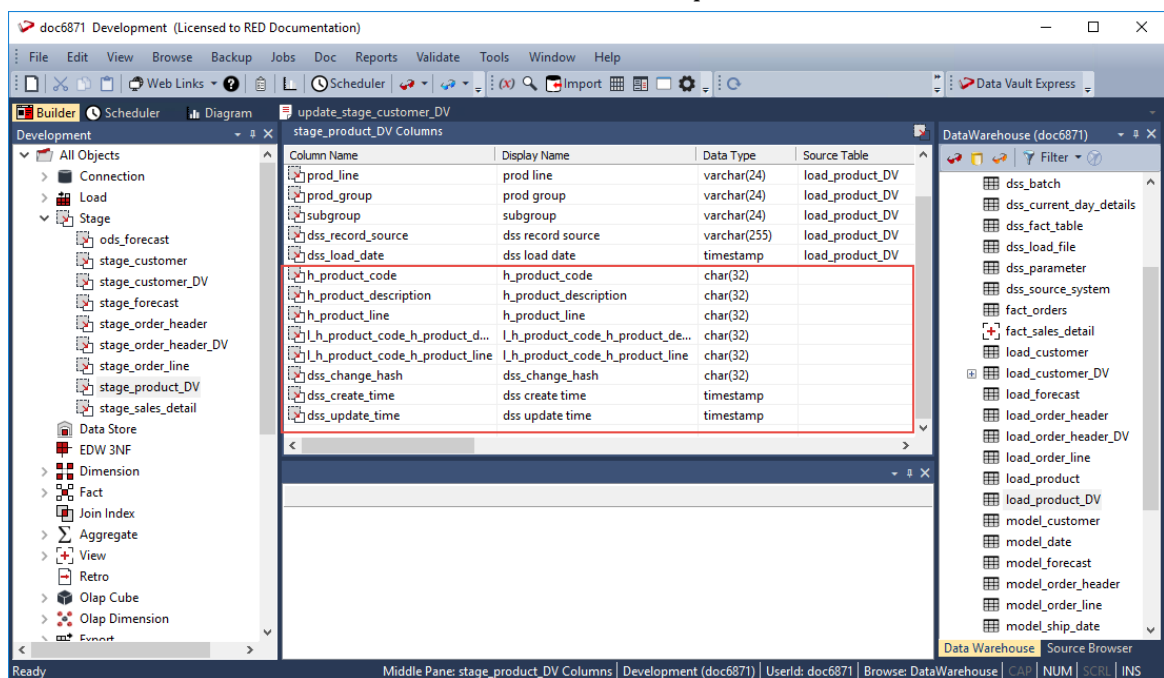
10 Follow the same steps to select the columns to use for the **Change Hash Key**.



Notes:

1. Creating a **Change Hash Key** via the new Wizard in WhereScape RED is limited to one key per **Data Vault Stage** table. This limitation does not apply for Data Vault models imported from WhereScape 3D.
2. Custom aliases for the **dss_load_date** and **dss_record_source** columns can be defined in the **Tools>Options>DSS Tables and Columns>DSS Columns to Include and Naming** window. If custom DSS aliases are defined then the aliases are used for the column names that are added to the new Load tables. Refer to **DSS Columns** (on page 103) for more details.
3. The data type defined for the **dss_change_hash column** in **Tools>Options>DSS Tables and Columns>DSS Columns for Custom Targets** window is applied when creating a hub hash key or a link hash key for tables in custom database targets. Refer to **DSS Columns for Custom Targets** (on page 104) for details.

- 11 Click **Finish**, once you have defined the required list of descriptive columns for the Change hash key. The new **Stage** table is added to the **Stage Table** objects list in the left pane and the columns included in the table are listed in the middle pane.



In addition to the columns defined from the Load table, the following columns and their metadata have been added to the Data Vault Stage table:

- The Hub Hash Keys
- The Link Hash Keys
- The Change Hash Key
- The DSS_CREATE_TIME column
- The DSS_UPDATE_TIME column

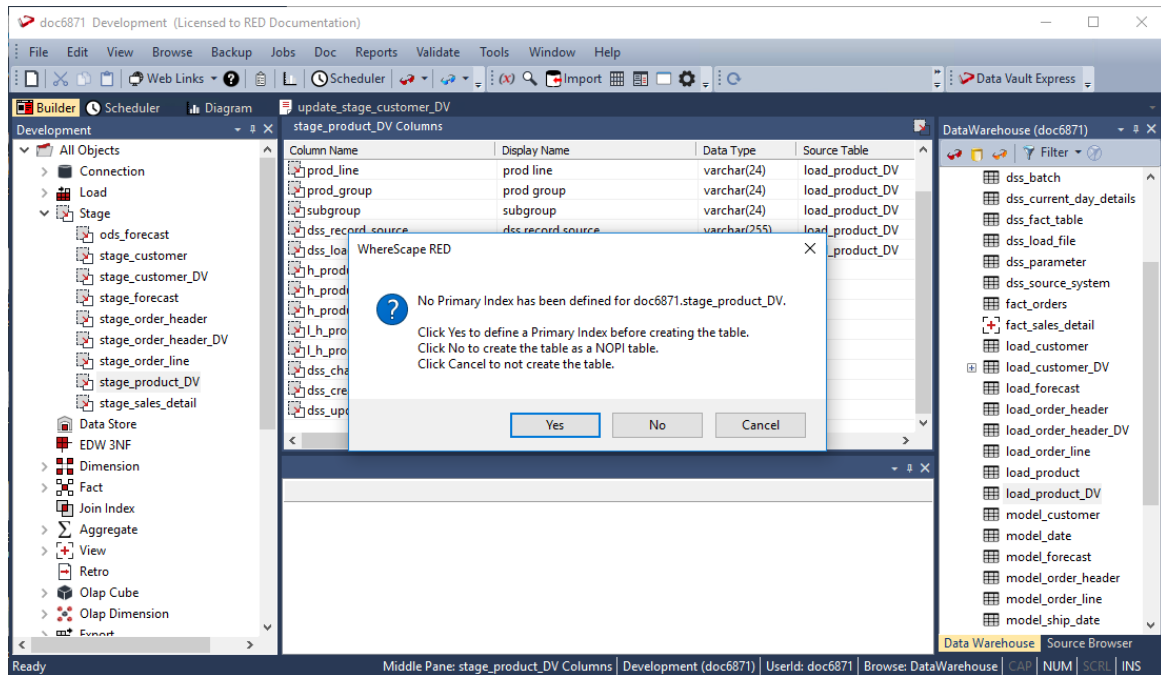
The metadata for the hash columns include the source columns that were used to create them (used to generate the hash keys).

The hash keys created are used in the subsequent Data Vault object (Hub, Links and Satellites) creation Wizards.

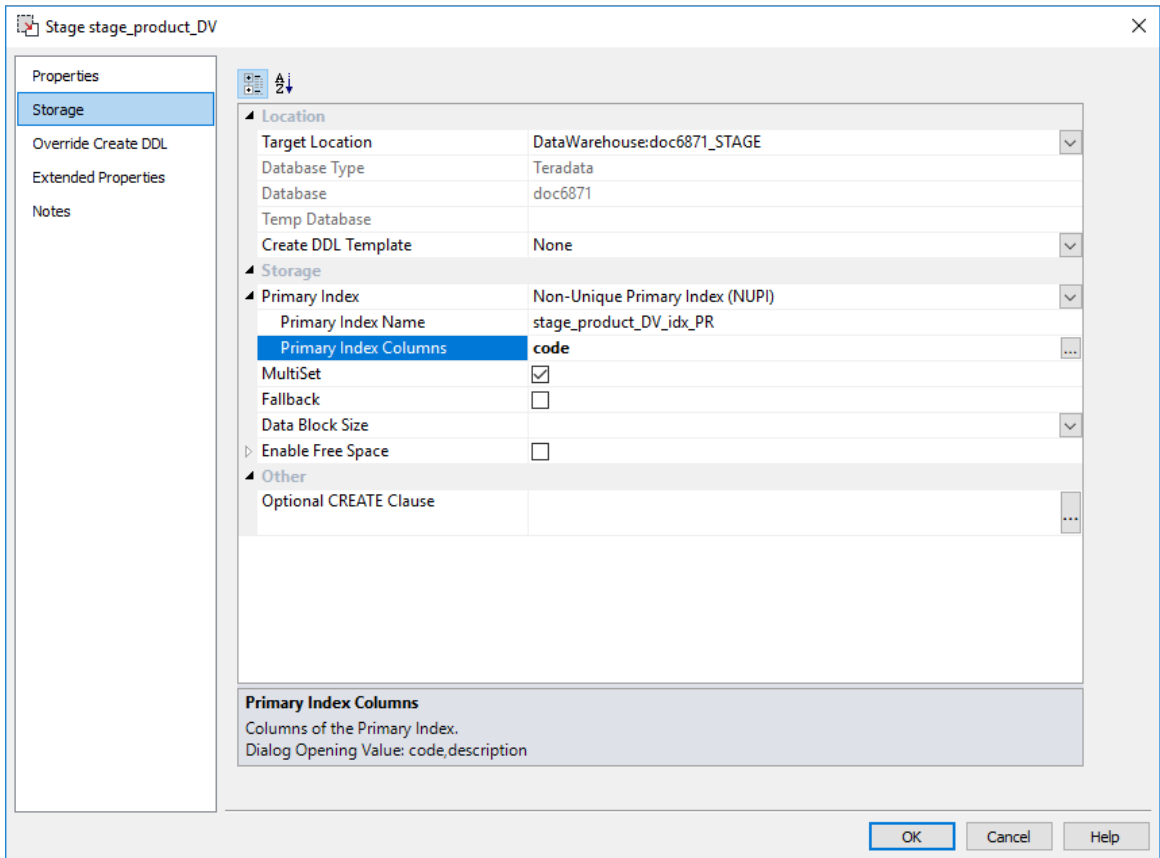
Tip: You can review the composition of existing hash keys for a Data Vault Stage table (Hub, Link and Satellite) and create additional hash keys by selecting the **Maintain DV Hash Key Columns** option from the selected **Data Vault Stage** table's context menu. This launches the hash key generation Wizard which enables you to maintain the source columns defined for the hash keys.

Note: To remove or change a hash key column, you need to delete it first, e.g. right click the column listed in the middle pane and then select **Delete Column** from the context menu.

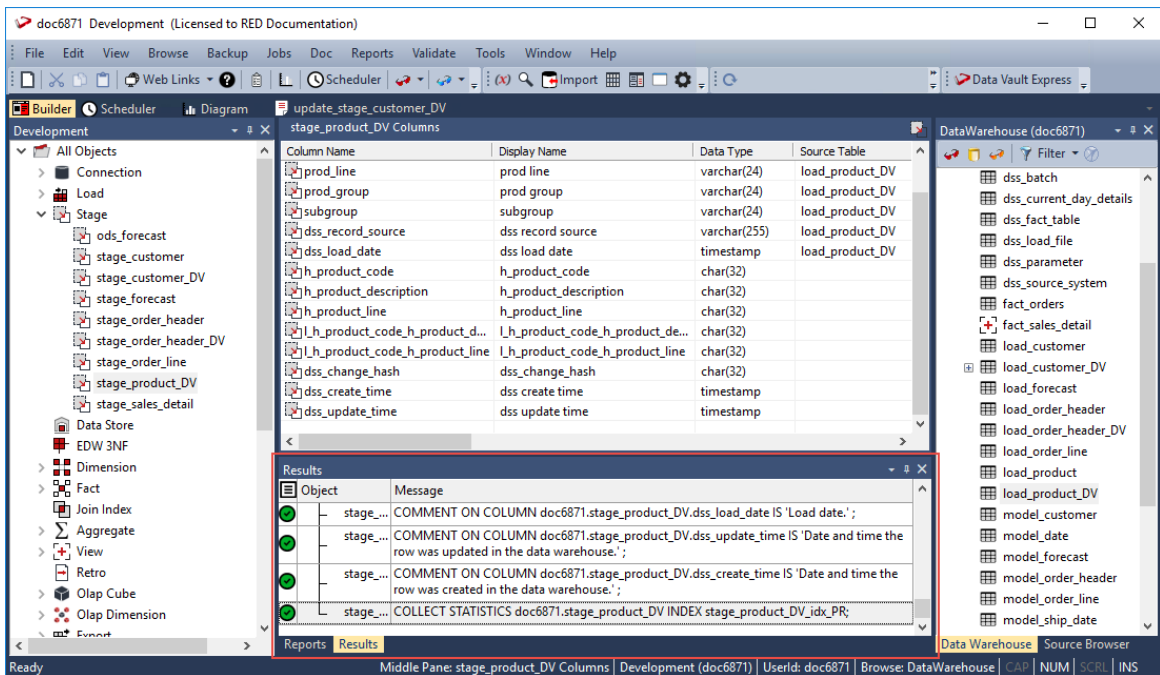
- 12 Right-click the new **Data Vault Stage** table you defined from the left pane, under the **Stage Table** objects list and select **Create (ReCreate)** from the context menu to create the table.
- 13 Click Yes on the Primary Index prompt.



14 Specify **code** on the **Primary Index Columns** field and then click **OK**.



15 The **Results** pane displays confirmation that the **Data Vault Stage** table was successfully created.



Once the new Data Vault Stage table is defined and created, clicking the **Rebuild** button on the **Table Properties** screen launches the Wizard to generate the procedure to populate the table. This Wizard utilizes a template to create the procedure.

The detailed steps for using this Wizard is described in the next section, Generating Update Procedures for the Data Vault Stage Table.

GENERATING UPDATE PROCEDURES FOR THE DATA VAULT STAGE TABLE

After successfully defining and creating the Stage table, you can generate the update procedure via a template to populate the table.

Note: Please ensure that you have installed the WhereScape supplied *templates* (see "*Data Vault Templates*" on page 418) or created your own Data Vault templates, before performing the steps below.

Important: The Data Vault hash key functions for Teradata and Oracle repositories have been changed in the templates (wsl_teradata_utility_dv and wsl_oracle_utility_dv). Data Vault customers running on Teradata or Oracle who were using these templates from the RED 6.9.1.0 release are advised to retain that version of the template, unless they migrate any key generated using the previous hash key function.


- 1 Click the **Rebuild** button on the **Table Properties** screen to launch the procedure generation Wizard to populate the table.

The screenshot shows the 'Table Properties' dialog for a 'Data Vault Stage' table named 'stage_product_DV'. The 'Update Procedure' is set to 'update_stage_product_DV' and the 'Custom Procedure' is set to '(None)'. The 'Rebuild' button is highlighted in blue. The 'Timestamps' section at the bottom shows the following values:

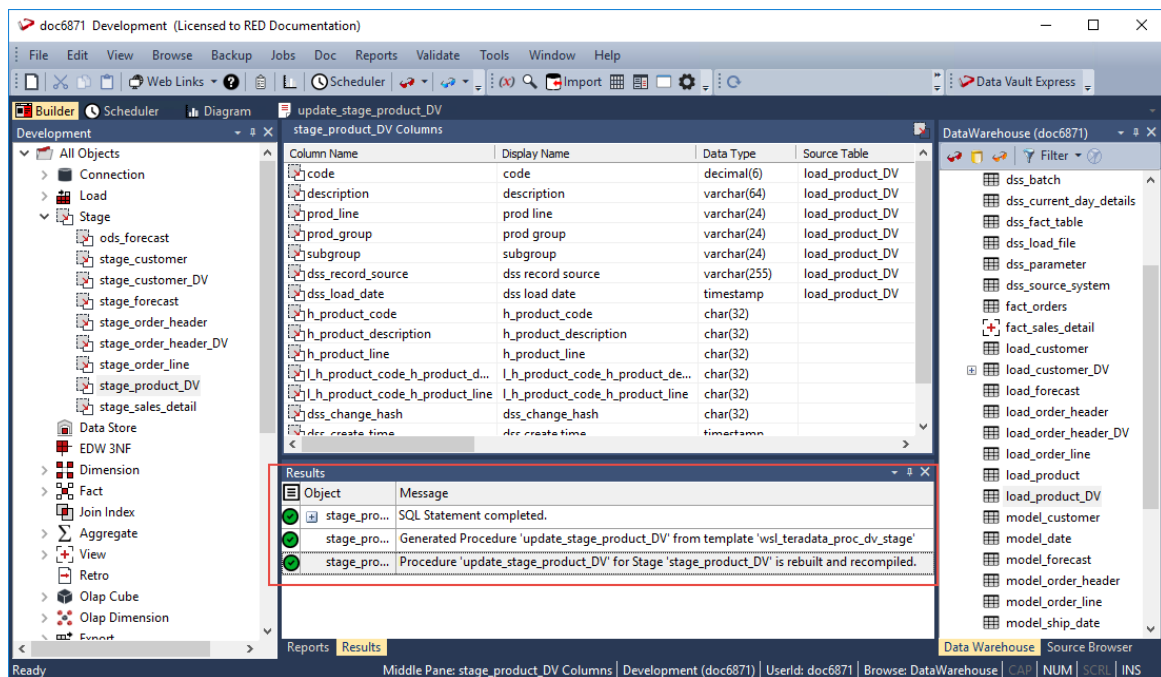
Timestamps	Metadata Structure	Database	Database Altered:
	2017-05-28 23:09:50.730000	2017-05-28 23:32:45.030000	2017-05-28 23:32:45.030000

Note: RED displays the name of the previously used update procedure template below the **Update Procedure** drop-down field by default.

- 2 On the **Processing** tab of **Table Update Build Options** screen, select the template to use from the **Template** drop-down or use the previous update procedure template.



- 3 Click **OK** to proceed with the procedure generation. The **Results** pane displays confirmation that the procedure was generated.



- Right-click the **Data Vault Stage** table in the left pane, under the **Stage Table** objects list and select **Code>View update** from the context menu to view the contents of the update procedure generated.

```

doc6871 Development (Licensed to RED Documentation)
File Edit View Format Compile Execute Tools Window Help
Builder Scheduler Diagram update_stage_product_DV x
-----
-- DBMS Name      : Teradata
-- Procedure Name : update_stage_product_DV
-- Template       : wsl_teradata_proc_dv_stage
-- Template Version : 8.0.2.0
-- Description    : Update the Stage table stage_product_DV
-- Generated by   : WhereScape RED Version 8.1.1.0 (build 180318-214019 pre-release)
-- Generated for  : RED Documentation
-- Generated on   : Monday, March 19, 2018 at 11:21:56
-- Author        : WhereScape Documentation
-----
-- Notes / History
--
CREATE PROCEDURE [METABASE].update_stage_product_DV
(
  IN p_sequence      INTEGER
, IN p_job_name     VARCHAR(256)
, IN p_task_name    VARCHAR(256)
, IN p_job_id       INTEGER
, IN p_task_id      INTEGER
, OUT p_return_msg  VARCHAR(256)
, OUT p_status      INTEGER
)
BEGIN
-----
-- Control variables used in most programs
-----
DECLARE v_msgtext    VARCHAR(255); -- Text for audit trail
DECLARE v_sql       VARCHAR(255); -- Text for SQL statements
DECLARE v_step      INTEGER;      -- return code
DECLARE v_update_count INTEGER;   -- no of records updated
DECLARE v_insert_count INTEGER;   -- no of records inserted
  
```

- Right-click the **Data Vault Stage** table in the left pane and select **Execute Update Procedure** from the context menu to run the procedure. The **Results** pane displays the number of records created.

The screenshot shows the 'Results' pane with the following messages:

Object	Message
stage_pro...	SQL Statement completed.
stage_pro...	Generated Procedure 'update_stage_product_DV' from template 'wsl_teradata_proc_dv_stage'
stage_pro...	Procedure 'update_stage_product_DV' for Stage 'stage_product_DV' is rebuilt and recompiled.
stage_pro...	Procedure Completed
stage_...	1 stage_product_DV updated. 9 new records.0 records updated.

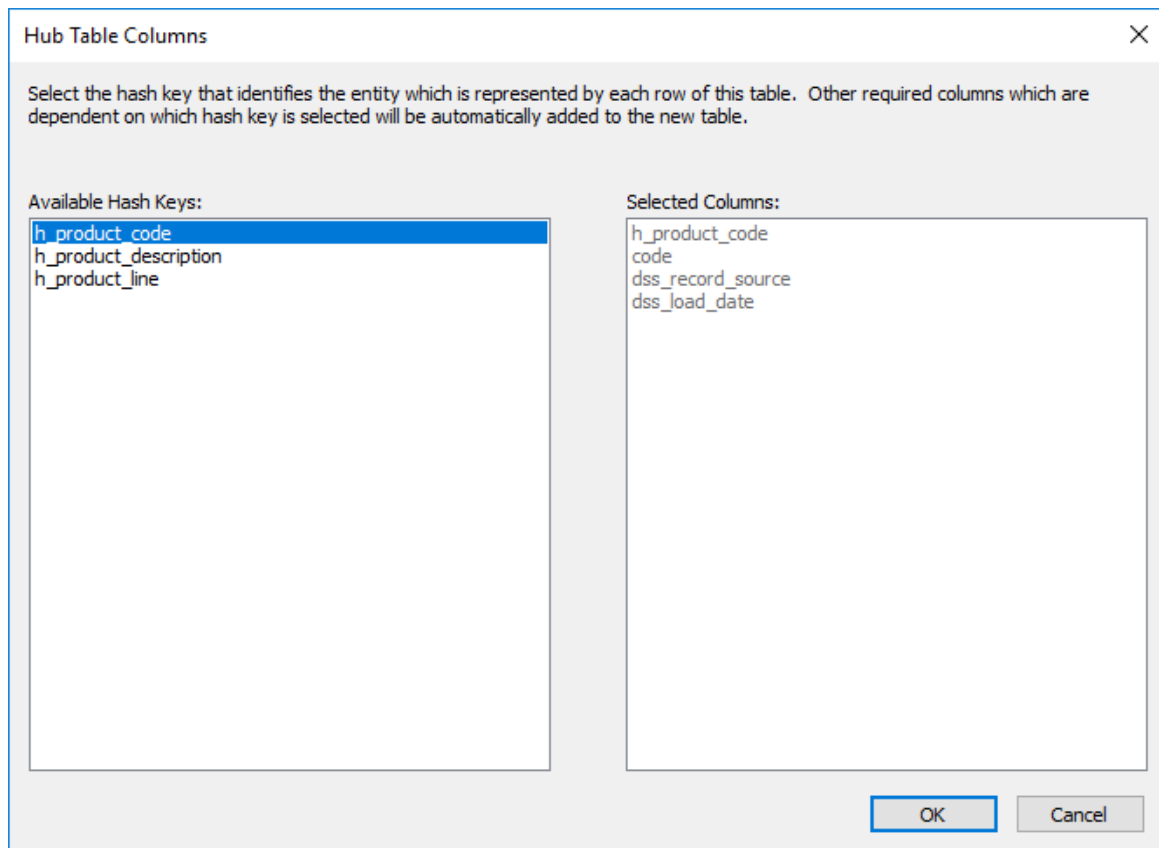
CREATING THE HUB, LINK AND SATELLITE TABLES

After successfully creating and populating the **Data Vault Stage** table you can now create the **Hub**, **Link** and **Satellite** tables. The **Hub**, **Link** and **Change** hash keys information stored in the **Data Vault Stage** table is used by the Wizard for building these **Data Vault** objects.

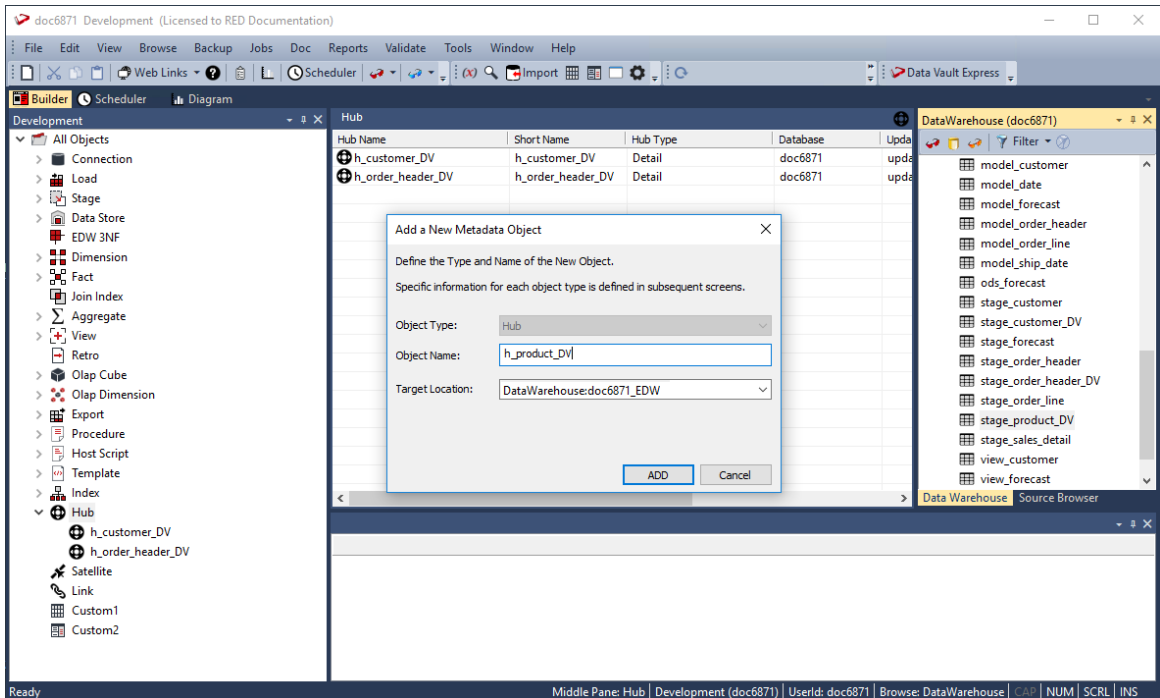
CREATING THE HUB TABLE

The following describe the steps for creating a Hub table:

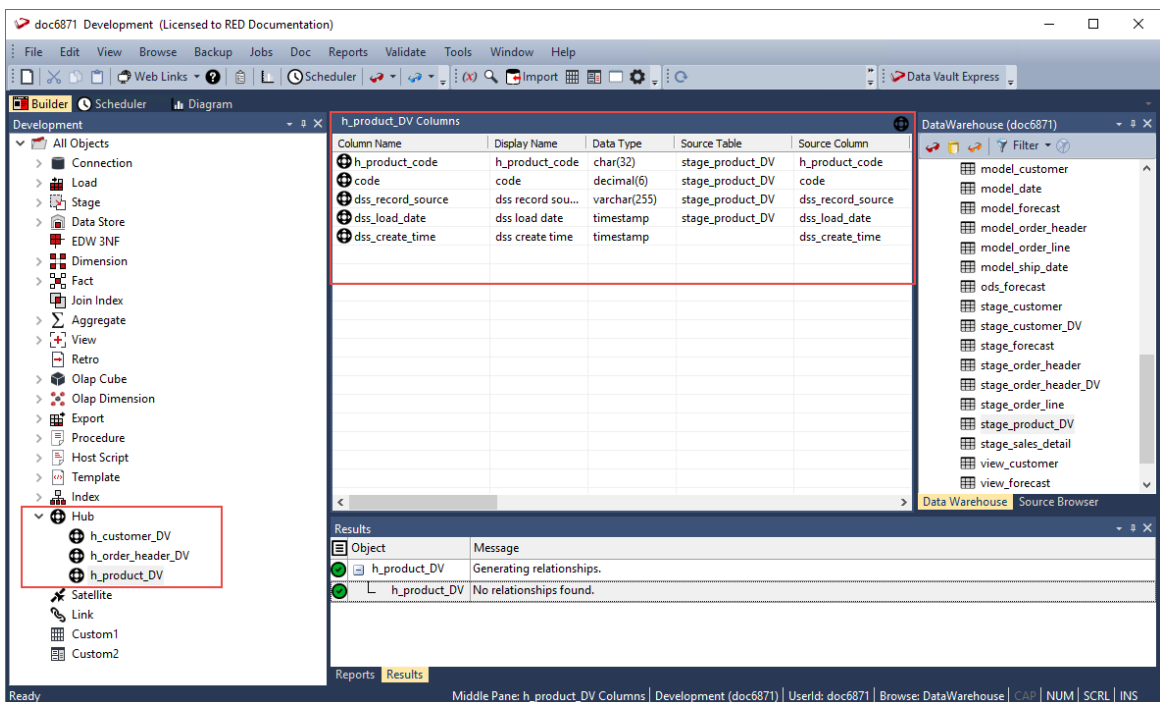
- 1 Browse to the Data Warehouse connection to create the Hub table.
- 2 Double-click the **Hub** object group in the left pane, the middle pane displays a list of existing Hub tables.
- 3 Click the source **Data Vault Stage** table from the right pane and drag it to the middle pane.
- 4 The Hub table creation Wizard appears and prompts you to select the Hash Key to use from the **Available Hash Keys** pane. The columns that comprise the selected Hash Key are displayed under the **Selected Columns** pane—these are the columns that will be populated by the Wizard on the new **Hub** table.



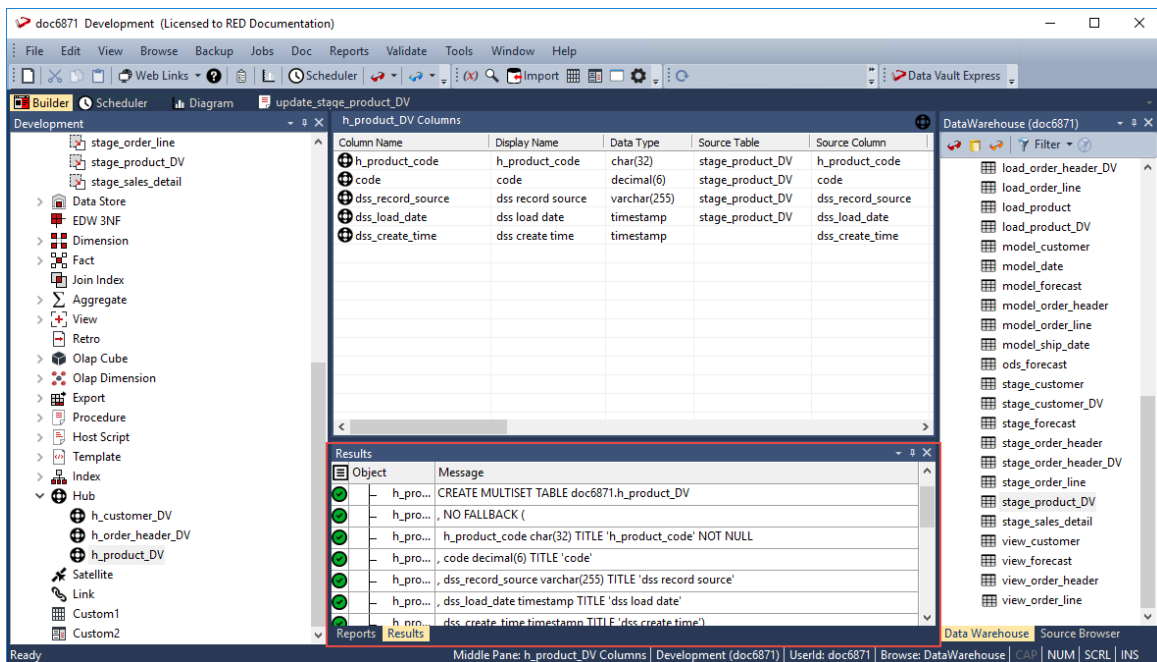
- 5 Select the Hash Key you want to use from the **Available Hash Keys** pane to see the columns that will be included in your Hub table under the **Selected Columns** pane. Click **OK** to continue.
- 6 The **Add a New Metadata Object** screen appears and classifies the new object as a Hub table. It provides a default name based on the source **Data Vault Stage** table name. Accept the name or enter a new name for the Hub table and click **ADD** to continue.



- 7 Click the **OK** button on the **Table Properties** screen to finish defining the meta data for the Hub table. The new Hub table is added to the **Hub** objects list in the left pane and the columns included in the table are listed in the middle pane.



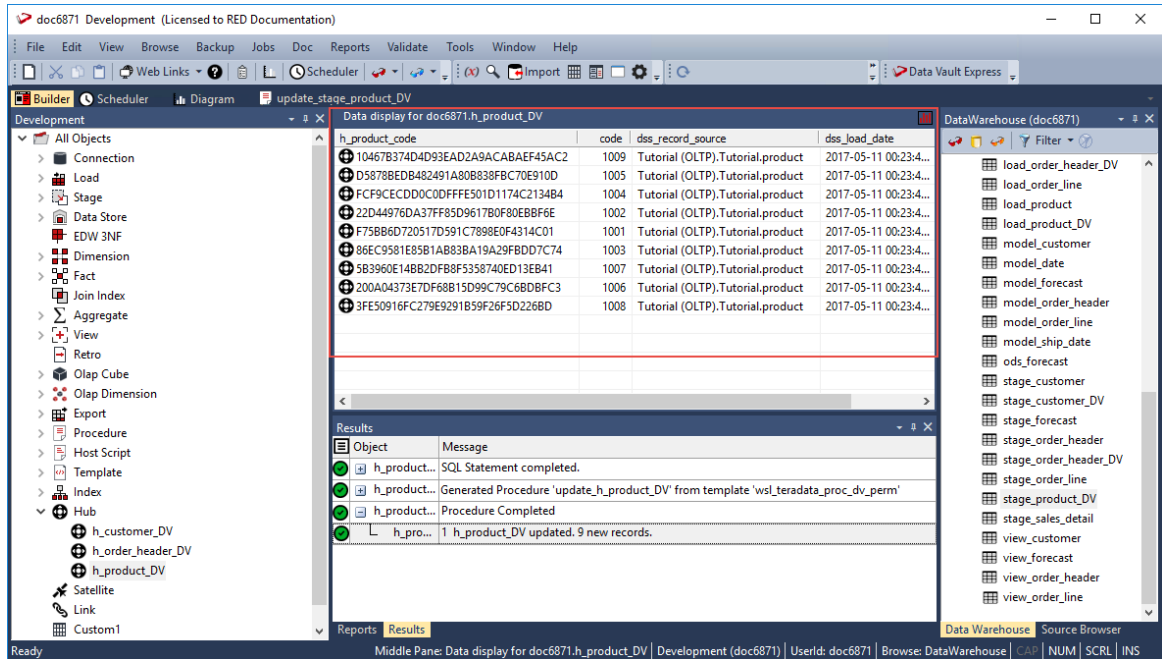
- 8 Right-click the new Hub table you defined in the left pane and select **Create (ReCreate)** from the context menu to create the table.
- 9 Click Yes on the Primary Index prompt and specify **code** on the **Primary Index Columns** field and then click **OK**. The **Results** pane displays confirmation that the Hub table was successfully created.



After the new Hub table is defined and created, clicking the **Rebuild** button on the **Table Properties** screen launches the Wizard to generate the update procedure to populate the table. This Wizard utilizes a template to create the procedure.

The detailed steps for using this Wizard is described in the section, *Generating Update Procedures for the Hub, Link and Satellite Tables* (see "Generating Update Procedures for Hub, Link and Satellite Tables" on page 451).

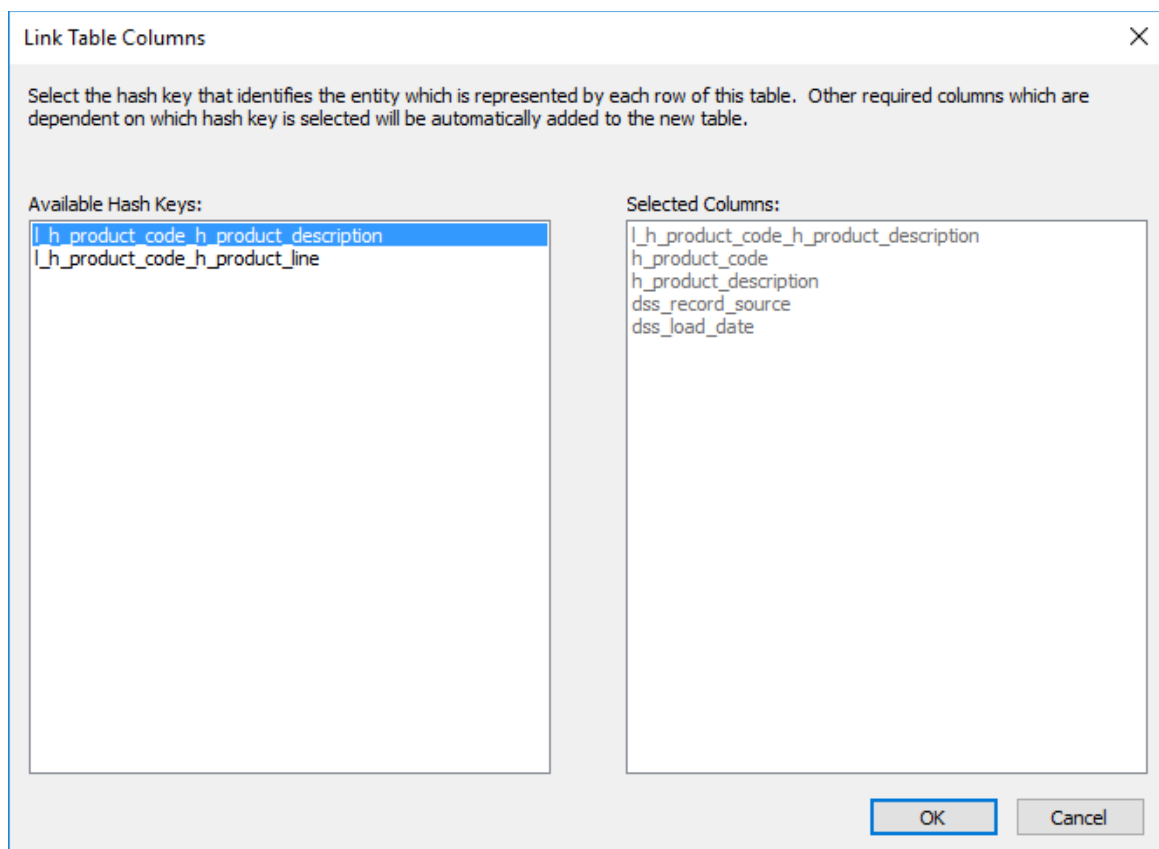
Once you have run the generated update procedures, you can view the generated Hub Hash keys, by right-clicking the new **Hub** table you created in the left pane and then selecting **Display Data** from the context menu:



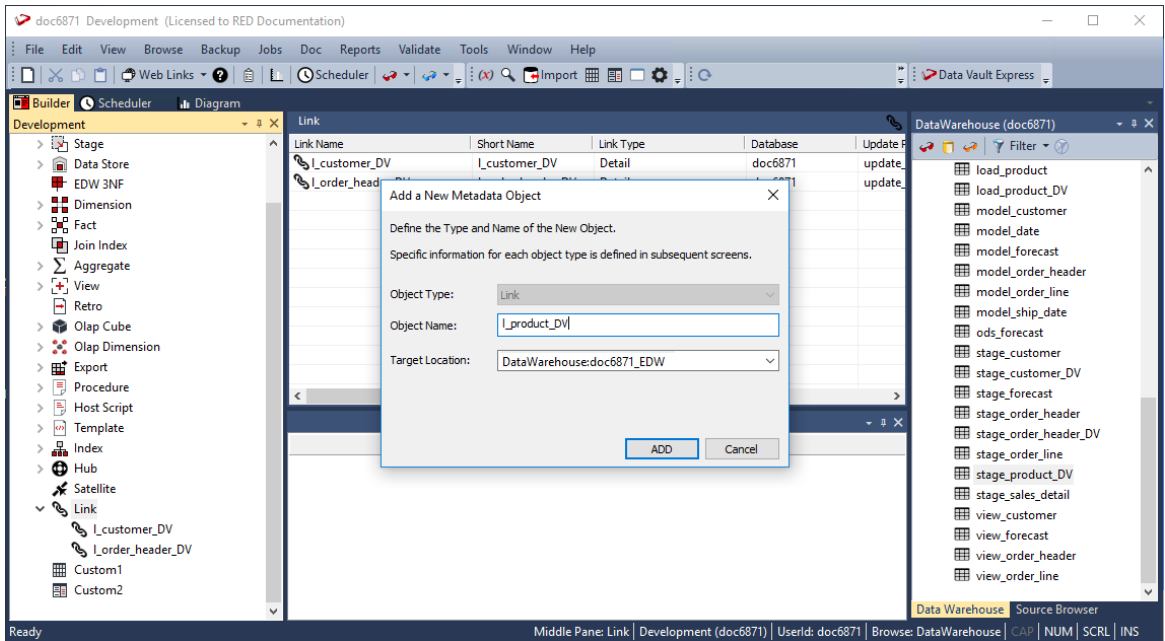
CREATING THE LINK TABLE

The steps for creating the Link table is similar to the steps used in creating the Hub table:

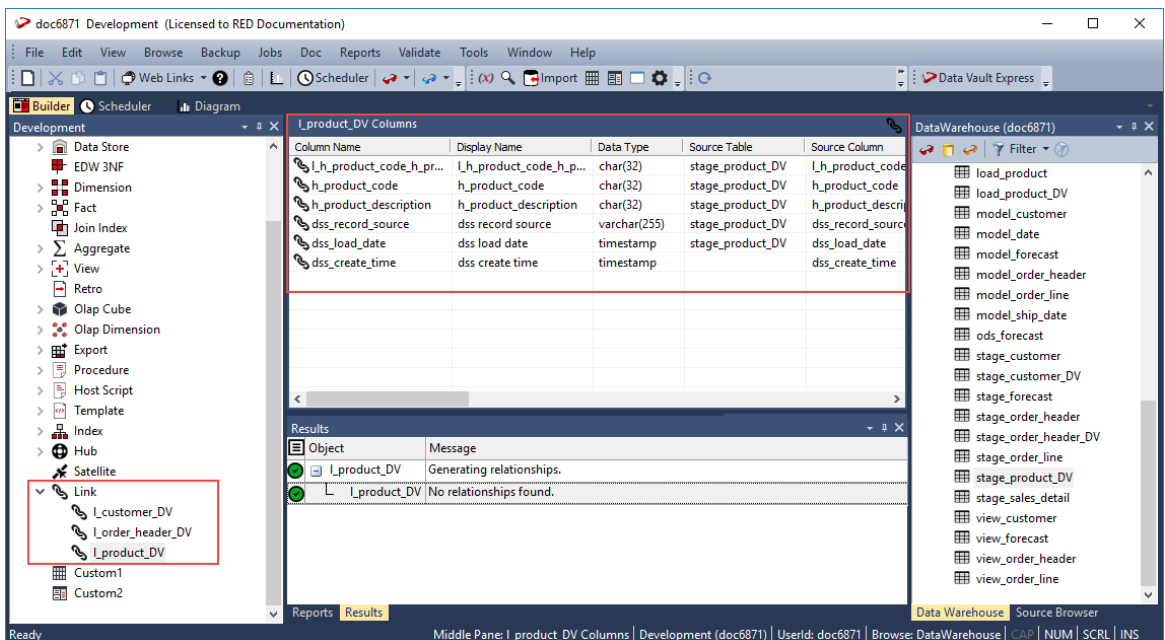
- 1 Browse to the Data Warehouse connection to create the Link table.
- 2 Double-click the **Link** object group in the left pane, the middle pane displays a list of existing Link tables.
- 3 Click the source **Data Vault Stage** table from the right pane and drag it to the middle pane.
- 4 The Link table creation Wizard appears and prompts you to select the Hash Key to use from the **Available Hash Keys** pane. The columns that comprise the selected Hash Key are displayed under the **Selected Columns** pane—these are the columns that will be populated by the Wizard on the new **Link** table.



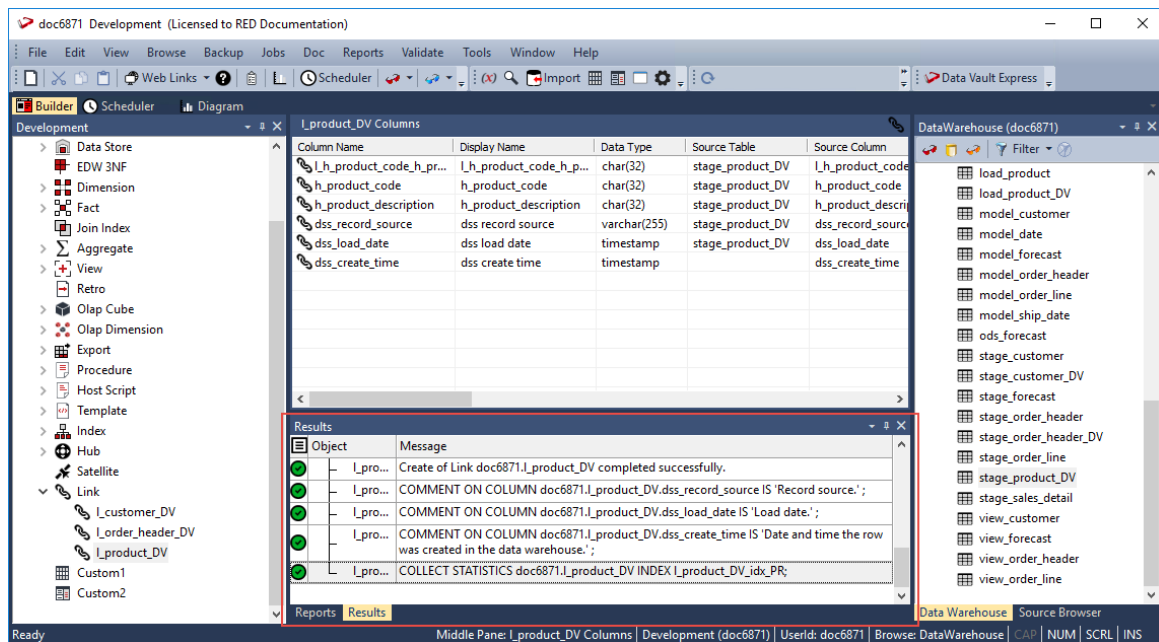
- 5 Select the Hash Key you want to use from the **Available Hash Keys** pane to see the columns that will be included in your Link table under the **Selected Columns** pane. Click **OK** to continue.
- 6 The **Add a New Metadata Object** screen appears and classifies the new object as a Link table. It provides a default name based on the source **Data Vault Stage** table name. Accept the name or enter a new name for the Link table and click **ADD** to continue.



- Click the **OK** button on the **Table Properties** screen to finish defining the meta data for the Link table. The new Link table is added to the **Link** objects list in the left pane and the columns included in the table are listed in the middle pane.



- 8 Right-click the new Link table you defined in the left pane and select **Create (ReCreate)** from the context menu to create the table.
- 9 Click Yes on the Primary Index prompt and specify **code** on the **Primary Index Columns** field and then click **OK**. The **Results** pane displays confirmation that the Link table was successfully created.



After the new Link table is defined and created, clicking the **Rebuild** button on the **Table Properties** screen launches the Wizard to generate the update procedure to populate the table. This Wizard utilizes a template to create the procedure.

The detailed steps for using this Wizard is described in the succeeding section, *Generating Update Procedures for the Hub, Link and Satellite Tables* (see "Generating Update Procedures for Hub, Link and Satellite Tables" on page 451).

Once you have run the generated update procedures, you can view the generated Link Hash keys, by right-clicking the new Link table you created in the left pane and then selecting **Display Data** from the context menu:

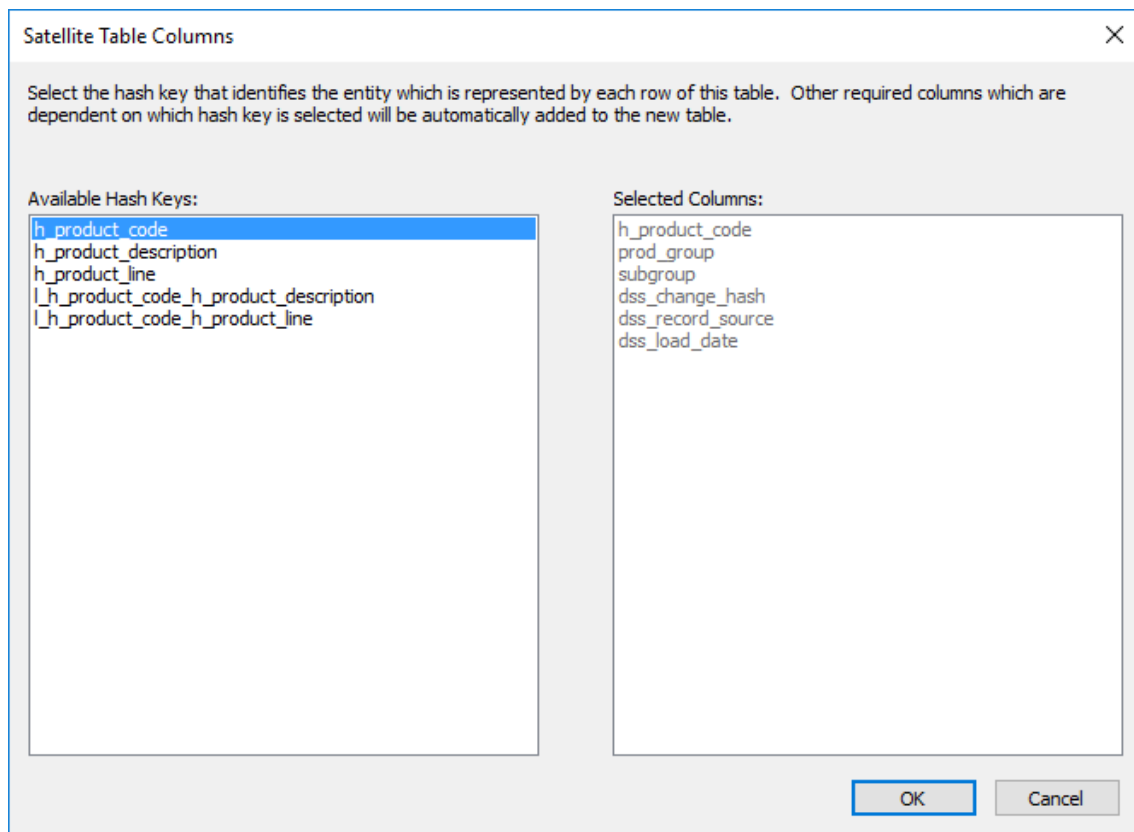
The screenshot shows the WhereScape Development tool interface. The main window displays a data display for the Link table 'L_product_DV'. The data is organized into columns: 'l_h_product_code', 'h_product_code', 'l_h_product_description', 'h_product_description', and 'dss_record'. The 'dss_record' column contains 'Tutorial (C)' for all rows. Below the data table, a 'Results' pane shows a list of messages, including 'Generating relationships.', 'SQL Statement completed.', 'Generated Procedure 'update_L_product_DV' from template 'wsl_teradata_proc_dv_perm'', 'Procedure 'update_L_product_DV' for Link 'L_product_DV' is rebuilt and recompiled.', 'Procedure Completed', and '1 L_product_DV updated. 9 new records.'.

l_h_product_code	h_product_code	l_h_product_description	h_product_description	dss_record
AC1D029E9329BA0B136D21C901D...	FCF9CECDD0C0FFFE501D...		941CE791B85E03FBEF82A2B...	Tutorial (C)
6A37CF3C938AA7C3EB46499D718...	10467B374D4D93EAD2A9A...		7BB6907A56374AB1FDD68C...	Tutorial (C)
EA2DA748CAD54555FF47A2073939...	3FE50916FC279E9291859F2...		252A2BC5364C81F2C8E1B9...	Tutorial (C)
E647B20A475C9960328D87CEE960F...	22D44976DA37FF85D9617B...		8F5DC24D18B4EF757A58D0...	Tutorial (C)
82D7A8EDF464040C485F1A8949D0...	5B3960E148B2DFB8F535874...		54C9BF562345C001C469E0...	Tutorial (C)
2032F22480A841F8DD10935CB159E...	86C9581E85B1A883BA19A...		BCDC0CC25837EC0B4F83C...	Tutorial (C)
7F0D271A155A4EA32B92AFA6A4C...	200A04373E7DF68B15D99C...		D2B10115699F38ABAB38E2...	Tutorial (C)
BDCCD8CEB4C459200FAAFEF0C6...	D58788EDB482491A80B838...		C8EB4492A2D8DA8647E31...	Tutorial (C)
5CF725956C3D540289D025EC606F...	F75BB6D720517D591C7898...		6C46B634A4D9930058E34F...	Tutorial (C)

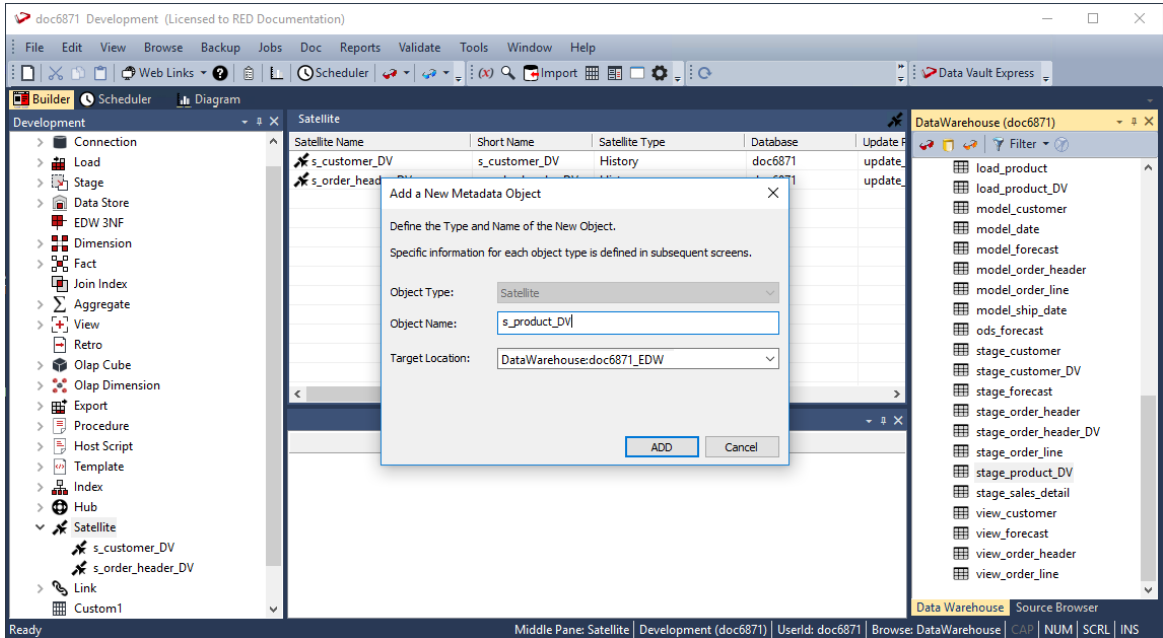
CREATING THE SATELLITE TABLE

The following describe the steps for creating a Satellite table:

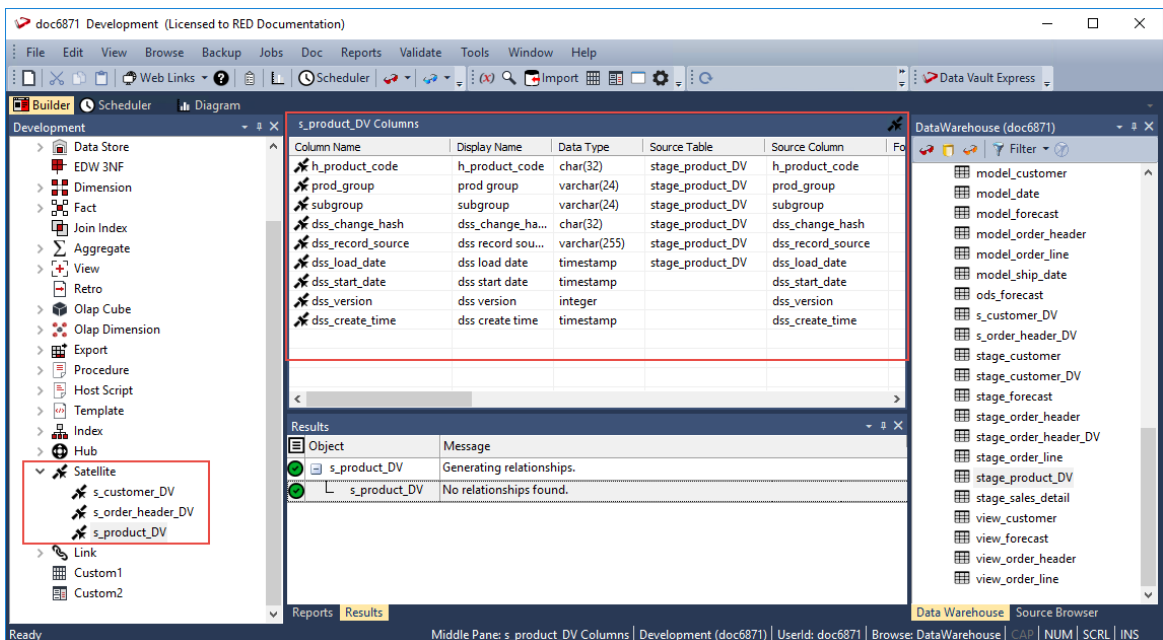
- 1 Browse to the Data Warehouse connection to create the Satellite tables.
- 2 Double-click the **Satellite** object group in the left pane, the middle pane displays a list of existing Satellite tables.
- 3 Click the source **Data Vault Stage** table from the right pane and drag it to the middle pane.
- 4 The Satellite table creation Wizard appears and prompts you to select the Hash Key to use from the **Available Hash Keys** pane. The columns that comprise the selected Hash Key are displayed under the **Selected Columns** pane—these are the columns that will be populated by the Wizard on the new Satellite table.



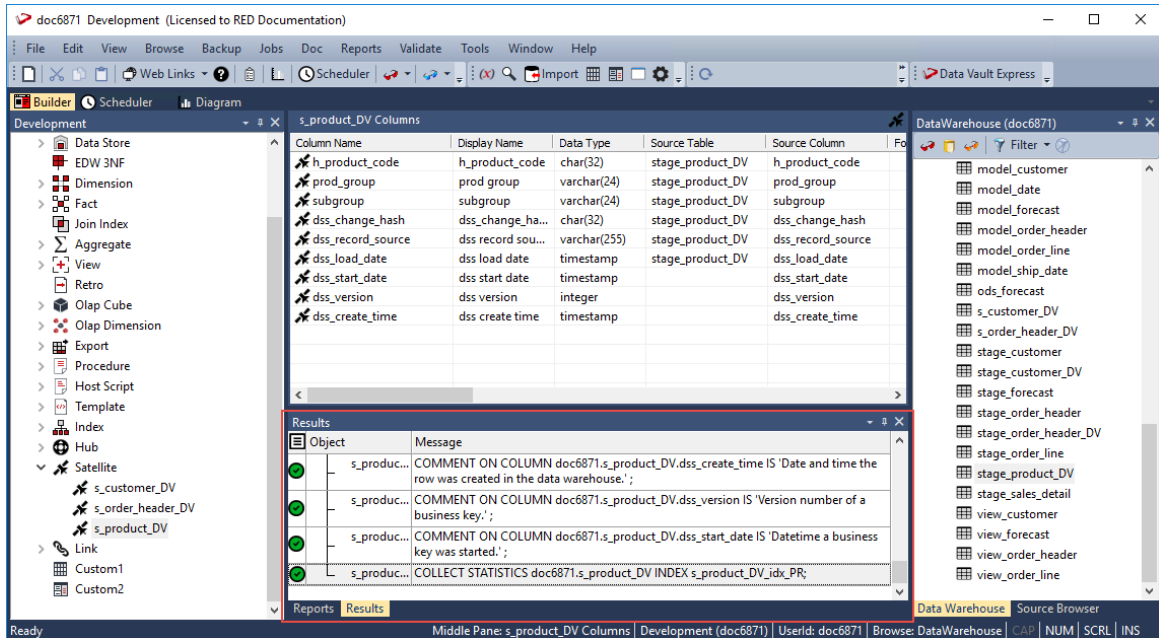
- 5 Select the Hash Key you want to use from the **Available Hash Keys** pane to see the columns that will be included in your Satellite table under the **Selected Columns** pane. Click **OK** to continue.
- 6 The **Add a New Metadata Object** screen appears and classifies the new object as a Satellite table. It provides a default name based on the source **Data Vault Stage** table name. Accept the name or enter a new name for the Satellite table and click **ADD** to continue.



- 7 Click the **OK** button on the **Table Properties** screen to finish defining the meta data for the Satellite table. The new Satellite table is added to the **Satellite** objects list in the left pane and the columns included in the table are listed in the middle pane.



- 8 Right-click the new Satellite table you defined in the left pane and select **Create (ReCreate)** from the context menu to create the table.
- 9 Click Yes on the Primary Index prompt and specify **code** on the **Primary Index Columns** field and then click **OK**. The **Results** pane displays confirmation that the Satellite table was successfully created.



After the new Satellite table is defined and created, clicking the **Rebuild** button on the **Table Properties** screen launches the Wizard to generate the update procedure to populate the table. This Wizard utilizes a template to create the procedure.

The detailed steps for using this Wizard is described in the succeeding section, *Generating Update Procedures for the Hub, Link and Satellite Tables* (see "Generating Update Procedures for Hub, Link and Satellite Tables" on page 451).

Once you have run the generated update procedures, you can view the generated Satellite Hash keys, by right-clicking the new Satellite table you created in the left pane and then selecting **Display Data** from the context menu:

The screenshot shows the WhereScape Data Warehouse interface. The left pane displays a tree view of the development environment, including a 'Satellite' folder with a newly created 's_product_DV' table. The main pane shows a 'Data display for doc6871.s_product_DV' table with the following data:

h_product_code	prod_group	subgroup	dss_change_hash	dss_record
FCF9CECCDD0C0DFFFE501D1174C2134B4	Talking	Jar	D779157DA6DCD3A44D960...	Tutorial (
10467B374D493EAD2A9ACABAEF45AC2	Equipment	Tools	9CC56C6C1C8A3540677AF8...	Tutorial (
3FE50916FC279E9291B59F26F5D226BD	Blocks	Plastic	267ABA4851962DDA164362...	Tutorial (
22D44976DA37FF85D961780F80EBBF6E	Equipment	Tractor	DEBDFDE7079F8BC310705C...	Tutorial (
5B3960E14BB2DFB8F5358740ED13EB41	Clubs	Golf	C052CAA992083D13F994D0...	Tutorial (
86EC9581E85B1A883BA19A29FBDD7C74	Blocks	Foam	8E122A4ACB1F21D5426F504...	Tutorial (
200A04373E7DF68B15D99C79C6BDBFC3	Balls	Golf	925E40DD792E307857C893E...	Tutorial (
D5878BEDB482491A80B838FBC70E910D	Balls	Golf	925E40DD792E307857C893E...	Tutorial (
F75B86D720517D591C7898E0F4314C01	Balls	Large	D18DA6354C35274324BEDD...	Tutorial (

The 'Results' pane at the bottom shows the following messages:

Object	Message
s_product_DV	Generating relationships.
s_product_DV	SQL Statement completed.
s_product_DV	Generated Procedure 'update_s_product_DV' from template 'wsl_teradata_proc_dv_perm'
s_product_DV	Procedure 'update_s_product_DV' for Satellite 's_product_DV' is rebuilt and recompiled.
s_product_DV	Procedure Completed
s_product...	1 s_product_DV updated, 9 new records.

The right pane shows a list of objects in the Data Warehouse, including 'model_customer', 'model_date', 'model_forecast', 'model_order_header', 'model_order_line', 'model_ship_date', 'ods_forecast', 's_customer_DV', 's_order_header_DV', 'stage_customer', 'stage_customer_DV', 'stage_forecast', 'stage_order_header', 'stage_order_header_DV', 'stage_order_line', 'stage_product_DV', 'stage_sales_detail', 'view_customer', 'view_forecast', 'view_order_header', and 'view_order_line'.

GENERATING UPDATE PROCEDURES FOR HUB, LINK AND SATELLITE TABLES

The following describe the steps for generating update procedures via a template.

HUB TABLE

After successfully defining and creating the Hub table, you can generate the update procedure via a template to populate the table.

- 1 Click the **Rebuild** button beside the **Update Procedure** drop-down in the **Table Properties** screen, to launch the procedure generation Wizard to populate the table.

Hub h_product_DV

Properties

Storage

Override Create DDL

Extended Properties

Purpose

Concept

Grain

Examples

Usage

Notes

Table Name: h_product_DV Table Type: Detail

Unique Short Name: (maximum 22 characters) h_product_DV

Business Display Name (EUL): h_product_DV

Description:

Update Procedure: update_h_product_DV Edit Rebuild Regenerate

Template: wsl_teradata_proc_dv_perm

Custom Procedure: (None) Rebuild Regenerate

Get Key Function: (None) Mnemonic (EUL):

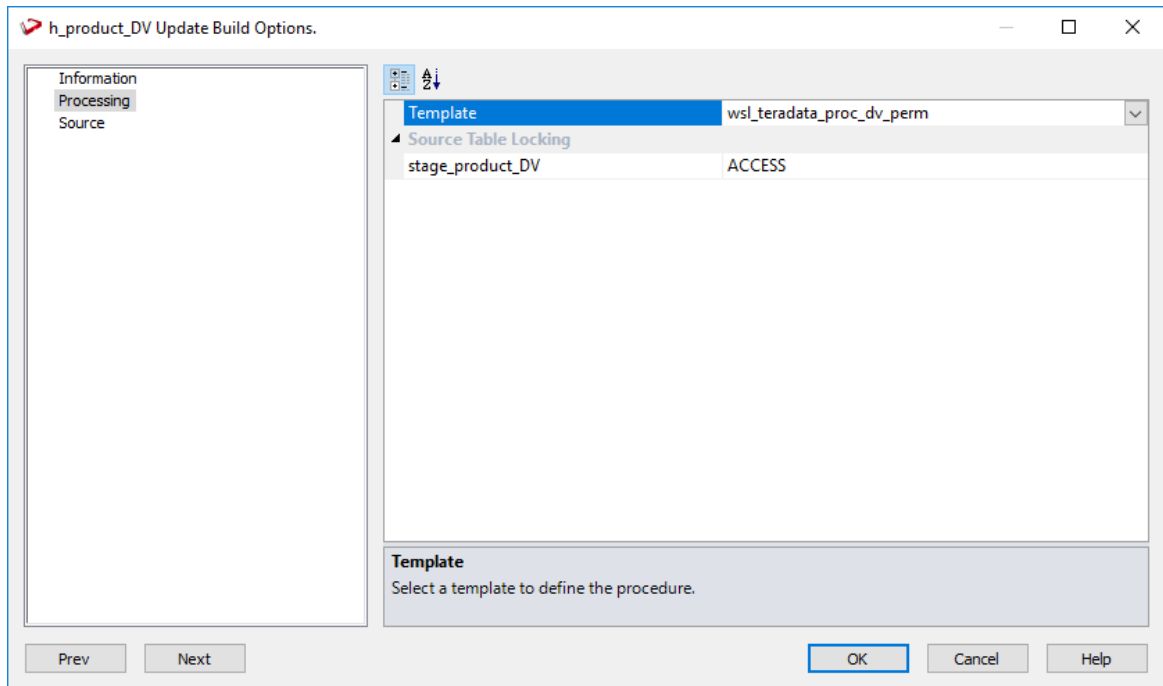
Timestamps

Metadata Structure	Database	Database Altered:
2017-05-29 02:25:33.080000	2017-05-29 02:26:07.280000	2017-05-29 02:26:07.280000

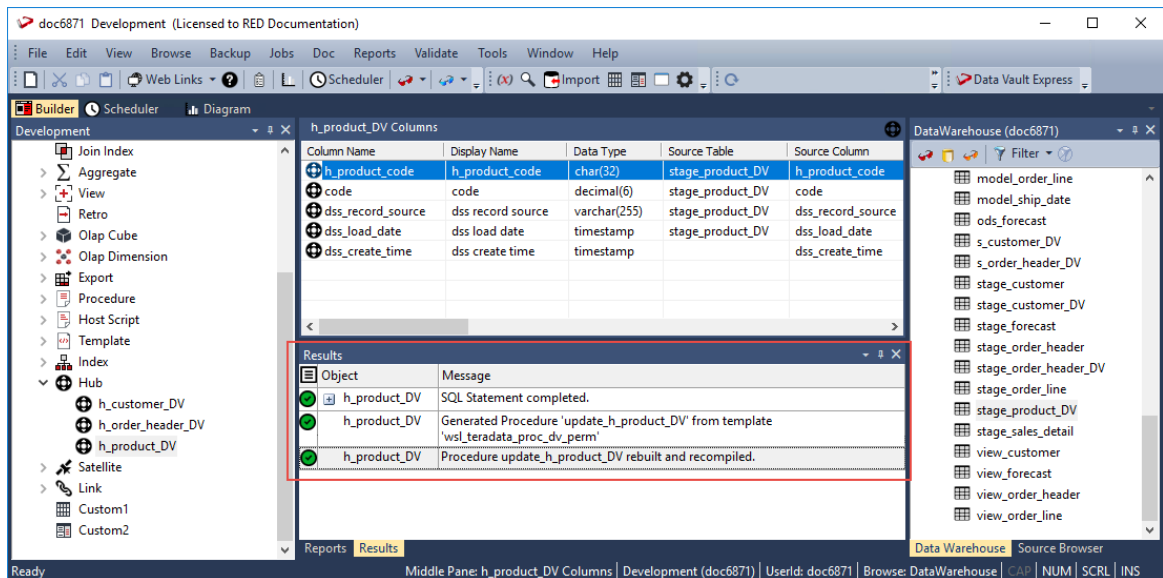
OK Cancel Help

Note: RED displays the name of the previously used update procedure template below the **Update Procedure** drop-down by default.

- 2 On the **Processing** tab of **Table Update Build Options** screen, select the template to use from the **Template** drop-down or use the previous update procedure template.



- 3 Click **OK** to proceed with the procedure generation. The **Results** pane displays confirmation that the procedure was generated.



- Right-click the Hub table in the left pane, under the **Hub** objects list and select **Code>View update** from the context menu to view the contents of the update procedure generated.

```

doc6871 Development (Licensed to RED Documentation) -
File Edit View Format Compile Execute Tools Window Help
Builder Scheduler Diagram update_h_product_DV x
-----
-- DBMS Name      : Teradata
-- Procedure Name  : update_h_product_DV
-- Template       : wsl_teradata_proc_dv_perm
-- Template Version : 6.9.1.0
-- Description    : Update the Hub table h_product_DV
-- Generated by   : WhereScape RED Version 8.0.2.0 (build 171212-222554 pre-release)
-- Generated for  : RED Documentation
-- Generated on   : Monday, December 18, 2017 at 12:02:45
-- Author        : WhereScape Documentation
-----
-- Notes / History
--
CREATE PROCEDURE [METABASE].update_h_product_DV
(
  IN p_sequence      INTEGER
, IN p_job_name      VARCHAR(256)
, IN p_task_name     VARCHAR(256)
, IN p_job_id       INTEGER
, IN p_task_id      INTEGER
, OUT p_return_msg  VARCHAR(256)
, OUT p_status      INTEGER
)
BEGIN
-----
-- Control variables used in most programs
-----
DECLARE v_msgtext    VARCHAR(255); -- Text for audit trail
DECLARE v_sql       VARCHAR(255); -- Text for SQL statements
DECLARE v_step      INTEGER;      -- return code
DECLARE v_update_count INTEGER;   -- no of records updated
DECLARE v_insert_count INTEGER;   -- no of records inserted
-----
  
```

- Right-click the Hub table in the left pane and select **Execute Update Procedure** from the context menu to run the procedure. The **Results** pane displays the number of records created.

Column Name	Display Name	Data Type	Source Table	Source Column
h_product_code	h_product_code	char(32)	stage_product_DV	h_product_code
code	code	decimal(6)	stage_product_DV	code
dss_record_source	dss record source	varchar(255)	stage_product_DV	dss_record_source
dss_load_date	dss load date	timestamp	stage_product_DV	dss_load_date
dss_create_time	dss create time	timestamp		dss_create_time

Object	Message
h_product_DV	SQL Statement completed.
h_product_DV	Generated Procedure 'update_h_product_DV' from template 'wsl_teradata_proc_dv_perm'
h_product_DV	Procedure 'update_h_product_DV' for Hub 'h_product_DV' is rebuilt and recompiled.
h_product_DV	Procedure Completed
h_product_...	1 h_product_DV updated. 9 new records.

LINK AND SATELLITE TABLES

Follow the same steps described above to create and execute the update procedures for the Link and Satellite tables, via a template to populate the tables.

CHANGING THE DATA VAULT HASH KEY FUNCTION IN WHERESCAPE RED 6.9.1.0 AND ABOVE

WhereScape RED ships **templates** (see "**Data Vault Templates**" on page 418) to generate the procedures for **Data Vault** objects if the customer's license includes **Data Vault** object support and for database platforms SQL Server, Oracle or Teradata.

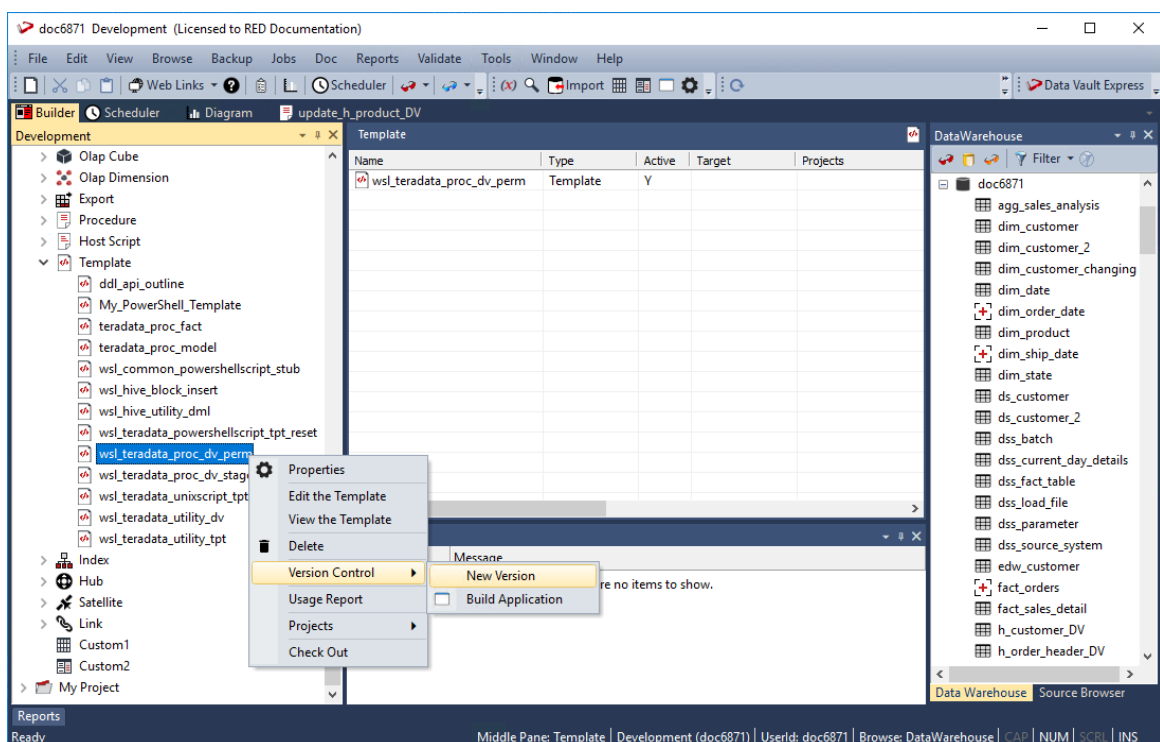
Included in the **Data Vault** utility (wsl_<database>_utility_dv) template is an **initial Hash key function**. This Hash key function must be evaluated before deployment, and if an alternate function is preferable for the customer's operating environment, then this function must be amended before Hash keys are created in a production environment.

The following outline the suggested steps to amend the shipped Hash key function, if required.

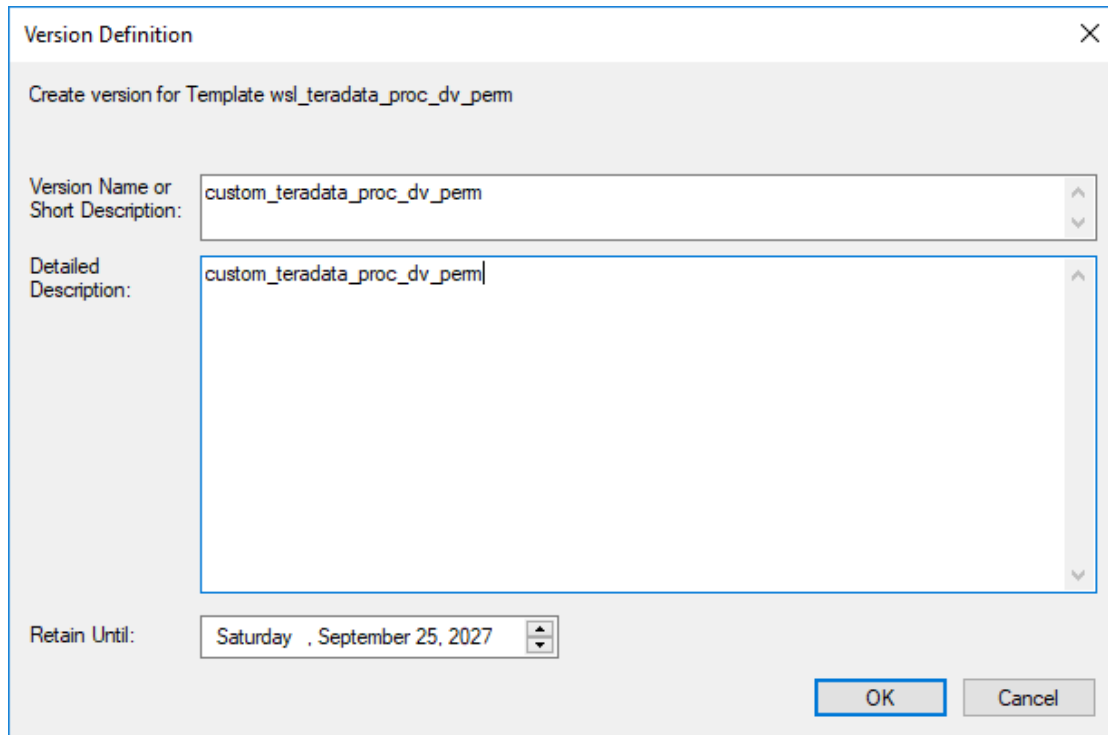
WhereScape **Data Vault** templates are Read-Only (to differentiate from custom templates), the first step is to make a copy of each of the three templates:

- wsl_<database>_proc_dv_perm
- wsl_<database>_proc_dv_stage
- wsl_<database>_utility_dv

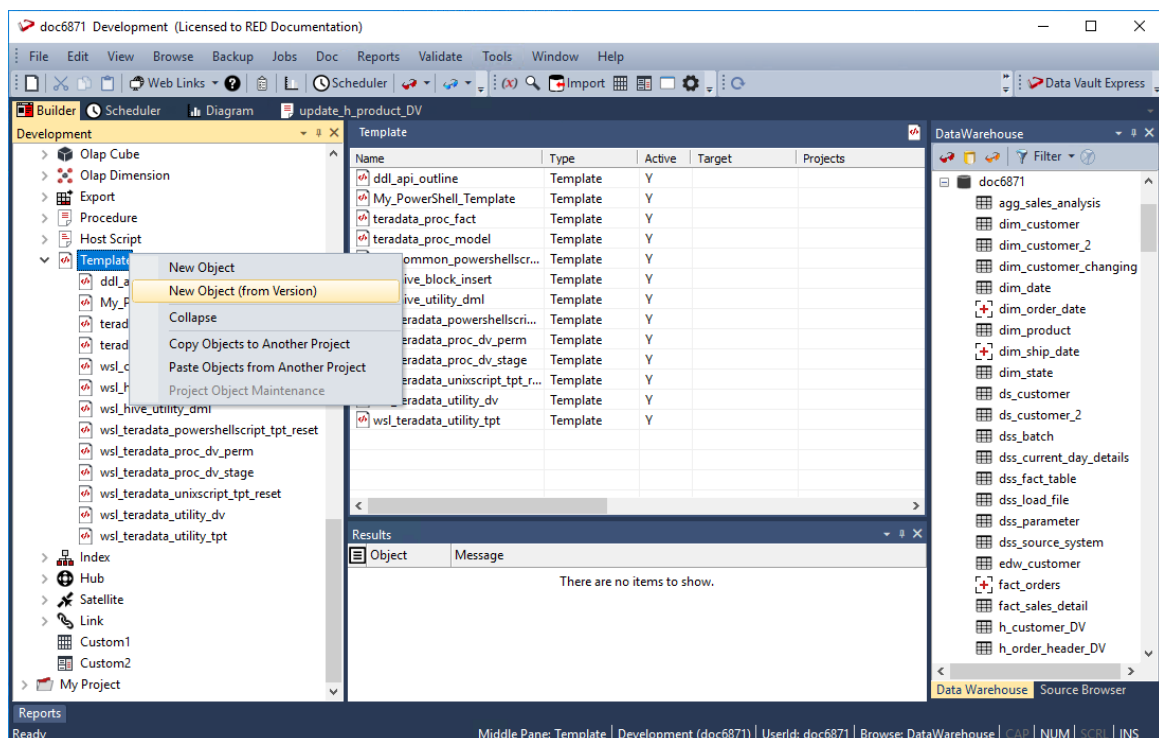
1 Right-click the template from the **Template** objects list and select **Version Control**>**New Version** from the context menu.



- 2 Enter the version name and other details on the **Version Definition** window and then click **OK** to save the new version of the template.



- 3 Right-click the **Template** object type and select **New Object (from version)** from the context menu.



- Select the new version of the template you created from the list in the **Create an Object From a Version** window and then enter the new object name.

Object Type:

Original Object Name	Description of Version	Versioned Date
wsl_teradata_proc_dv_perm	custom_teradata_proc_dv_perm	2017-09-25 04:41:13.83...
wsl_teradata_utility_dv	Auto Version on Application Load by Teresa	2017-09-18 03:55:47.82...
wsl_teradata_proc_dv_sta...	Auto Version on Application Load by Teresa	2017-09-18 03:55:47.30...
wsl_teradata_proc_dv_perm	Auto Version on Application Load by Teresa	2017-09-18 03:55:46.57...
wsl_common_powershells...	Auto Version on Application Load by Teresa	2017-09-18 03:55:39.17...
wsl_hive_utility_dml	Auto Version on Application Load by Teresa	2017-09-18 03:55:38.67...
wsl_hive_block_insert	Auto Version on Application Load by Teresa	2017-09-18 03:55:38.05...
wsl_teradata_utility_dv	Auto Version on Application Load by Teresa	2017-09-13 03:40:27.47...
wsl_teradata_proc_dv_sta...	Auto Version on Application Load by Teresa	2017-09-13 03:40:26.91...
wsl_teradata_proc_dv_perm	Auto Version on Application Load by Teresa	2017-09-13 03:40:26.34...
wsl_common_powershells...	Auto Version on Application Load by Teresa	2017-09-13 03:40:18.77...
wsl_hive_utility_dml	Auto Version on Application Load by Teresa	2017-09-13 03:40:18.36...
wsl_hive_block_insert	Auto Version on Application Load by Teresa	2017-09-13 03:40:17.41...
wsl_teradata_utility_dv	Auto Version on Application Load by Teresa	2017-08-29 06:18:47.05...
wsl_teradata_proc_dv_sta...	Auto Version on Application Load by Teresa	2017-08-29 06:18:46.62...
wsl_teradata_proc_dv_perm	Auto Version on Application Load by Teresa	2017-08-29 06:18:46.03...
wsl_common_powershells...	Auto Version on Application Load by Teresa	2017-08-29 06:18:39.27...
wsl_hive_utility_dml	Auto Version on Application Load by Teresa	2017-08-29 06:18:38.49...
wsl_hive_block_insert	Auto Version on Application Load by Teresa	2017-08-29 06:18:37.63...
wsl_teradata_utility_dv	Auto Version on Application Load by Teresa	2017-07-28 01:41:42.19...
wsl_teradata_proc_dv_sta...	Auto Version on Application Load by Teresa	2017-07-28 01:41:41.64...
wsl_teradata_proc_dv_perm	Auto Version on Application Load by Teresa	2017-07-28 01:41:41.10...

Click on the original object name of the version to be used as the basis of the create. Then change the object name if the object still exists, and Create

New Object Name:

Unique Short Name (maximum 22 characters)

- Click **Create** to create the new Template object.
- Follow the same steps to create the other new templates with unique names.
- Modify the replacement utility template with new function in this section:

```
{%- set hashFunctionPatBeg = "CAST(HASHBYTES('md5', "
-%}
{%- set hashFunctionPatEnd = ") AS CHAR(32))"
-%}
{%- set hashColTransPattern = "COALESCE(CAST([SRCCOL] AS
VARCHAR(MAX)), 'null')" -%}
{%- set hashConcatPattern = " + ' || '+'
-%}
```

In each of the addHubHashKey, addLinkHashKey and addChangeHashKey macros.

- In the two replacement proc_dv templates, change the import command to reference the new utility template:

```
{% import "custom_sqlserver_utility_dv" %}
```

- 9 Update existing **Data Vault Stage** table objects to reference the replacement `proc_dv_stage` template (the **Usage Report** on the context menu of the old template objects can help locate these).
- 10 After successfully testing the new templates, it is recommended to delete the original **WSL Data Vault** templates to avoid confusion.

Note: The above workflow does not provide allowance for the migration of any existing **Data Vault** records, created with a previous **Data Vault** function. This workflow can be used for new **Data Vault** implementations. In an existing **Data Vault** system, changing the Hash key function requires migration of existing Hash keys that is not provided in this workflow.

CHAPTER 15

CUSTOM OBJECTS

Custom1 and **Custom2** objects are user defined objects. These Object Types can be renamed in the **Tools>Options>Object Types>Object Names** menu.

A Custom object license is required for these object types.

Custom objects have the same options and properties as EDW 3NF tables, for more information see ***EDW 3NF Tables*** (on page 388).

CHAPTER 16

MODEL TABLES

IN THIS CHAPTER

Model Table Overview	462
Building a Model Table	464
Generating the Model Table Update Procedure.....	469
Model Table Artificial Keys.....	475
Model Table Custom Procedure.....	475
Model History Tables	476
Generating History Table Update Procedures	478
Model Table Column Properties	481
Model Table Column Transformations.....	486

MODEL TABLE OVERVIEW

Note: WhereScape RED version 8.1.1.0 and above no longer supports the Model Table object type. However, users that have existing Model table objects in their data warehouse can retain and continue to use them.

Model objects are used to create EDW 3NF models in an enterprise data warehouse. They can contain surrogate keys to other model tables.

A model is built from the Data Warehouse connection. Unless you are doing a retro-fit of an existing system, model tables are typically built from one or more stage tables.

The normal steps for creating a model table are defined below and are covered in this chapter. The steps are:

- Identify the source transactional or reference data that will constitute the model table. If the data is sourced from multiple tables ascertain if a join between the source tables is possible.
- Using the 'drag and drop' functionality drag the load table that is the primary source of information for the model table into a model target. See *Building a Model Table* (on page 464).
- If only one table is being sourced and most of the columns are to be used (or if prototyping) you can select the auto create option to build and load the model table and skip the next 4 steps. See *Building a Model Table* (on page 464).
- Add columns from other load tables if required. See *Building a Model Table* (on page 464).
- Create the model table in the database. See *Building a Model Table* (on page 464).
- Build the update procedure. See *Generating the Model Table Update Procedure* (on page 469).
- Run the update procedure and analyze the results. See *Generating the Model Table Update Procedure* (on page 469).

Modify the update procedure as required.

Model Keys

Model tables have up to two types of keys that we will refer to frequently.

These are the Business Key and the Artificial Key. A Definition of these two key types follows:

Business Key

The business key is the column or columns that uniquely identify a record within the model table. Where the model maps back to a single or a main table in the source system, it is usually possible to ascertain the business key by looking at the unique keys for that source table. Some people refer to the business key as the 'natural' key. Examples of business keys are:

- The product SKU in a product model table
- The customer code in a customer model table
- The calendar date in a date model table
- The 24 hour time in a time model table (i.e.HHMM) (e.g.1710)
- The airport short code in an airport model table.

It is assumed that business keys will never be NULL. If a null value is possible in a business key then the generated code will need to be modified to handle the null value by assigning some default value. For example the 'Where' clause in a model update may become:

```
Where coalesce(business_key,'N/A') = coalesce(v_LoadRec.business_key,'N/A')
```

Note: Business keys are assumed to never be Null. If they could be null it is best to transform them to some value prior to model or stage table update. If this is not done an unmodified update will probably fail with a duplicate key error on the business key index.

Artificial Key

The artificial key is the unique identifier that is used to join a model table record to another model table. When joining model tables to other model tables it would be possible to perform the join using the business key. For model tables with a large number of records this however would result in slow query times and very large indexes. As query time is one of our key drivers in data warehouse implementations the best answer is often to use some form of artificial key. A price is paid in the additional processing required to build the model table rows (particularly high volume transaction model tables rows), but this is offset by the reduced query times, storage and index sizes.

The artificial key is an integer and is built sequentially from 1 upwards. See the section on artificial keys for a more detailed explanation. An artificial key is sometimes referred to as a "surrogate" key.

Note: The default behavior of RED can be changed to not automatically add surrogate keys. See *Settings - Repository Attributes and Options* (see "*Settings - Repository Identification*" on page 79).

BUILDING A MODEL TABLE

Model Tables are often sourced from one table in the base application. The process for building a model table is the same for most other tables and begins with the drag and drop of the stage table that contains the bulk of the model information.

DRAG AND DROP

Create a model target by double clicking on the **Dimension object group** in the left pane. The middle pane displays a list of all existing Dimension tables.

Browse to the Data Warehouse via the **Browse>Source Tables** menu option.

Drag the table that contains the bulk of the model table columns, into the middle pane. Drop the table anywhere in the pane.

The new object dialog box will appear and will identify the new object as a Dimension table and will provide a default name based on the load table name.

Either accept this name or **change the name to reflect the new model table** and click OK to proceed.

MODEL TABLE PROPERTIES

At this stage change the Table type to **Model Table** and change any other storage options if desired.

If prototyping and the model table is simple (e.g. one source table), then it is possible to create, load and update the model table in a couple of steps.

If you want to do this, select the **(Build Procedure...)** option from the 'Update Procedure' drop-down and then click **Create and Load** on the next dialog.

Dimension model_customer

Properties

Storage

Override Create DDL

Language Mapping

Extended Properties

Purpose

Concept

Grain

Examples

Usage

Notes

Table Name: model_customer

Table Type: Model Table

Unique Short Name: (maximum 22 characters) model_customer

Business Display Name (EUL): model_customer

Description:

Update Procedure: (None) Rebuild

Custom Procedure: (None) Rebuild Regenerate

Mnemonic (EUL):

Timestamps

Metadata Structure: 2017-12-18 01:05:28.800000

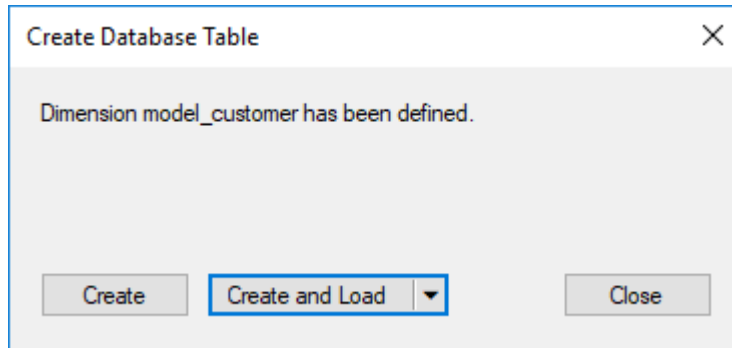
Database:

Database Altered:

OK Cancel Help

CREATE AND LOAD

If you chose to build the update procedure the following dialog appears after clicking OK on the Properties page. This dialog asks if you want to create the Model table in the database and execute the update procedure.

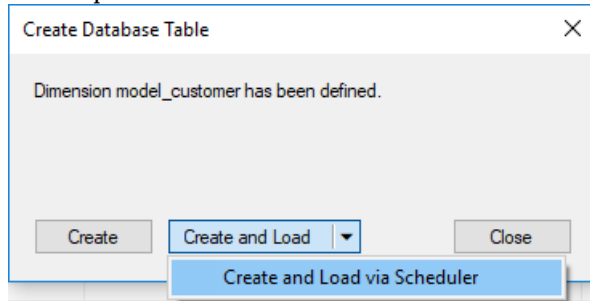


If you are satisfied with the columns that will be used and do not wish to add any columns you can select the **Create and Load** button. Alternatively, the **Create** button creates the table in the repository but does not execute an update, allowing you to change columns before loading data into the table.

If **Create** or **Create and Load** is selected and a new procedure creation was chosen proceed directly to the *Generating the Model Table Update Procedure* (on page 469) section.

If you have additional columns to add or columns to delete then select **Close** and proceed as follows.

Note: It is possible to create and load the table via the Scheduler; by selecting this option from the drop-down list on the **Create and Load** button:



DELETING AND CHANGING COLUMNS

The columns defined for the model table will be displayed in the middle pane. At this stage it is possible to delete any unwanted columns by highlighting a column name or a group of names and clicking the **Delete** key.

The name of a column can also be changed at this stage by selecting the column and using the right-click menu to edit its properties. Any new name must conform to the database naming standards. Good practice is to use alphanumeric characters and the underscore character. See the section on column properties for a fuller description on what the various fields mean.

Note: When prototyping, and in the initial stages of an analysis area build it is best not to remove columns, nor to change their names to any great extent. This type of activity is best left until after end users have used the data and provided feedback.

ADDING ADDITIONAL COLUMNS

With the columns of the model table displayed in the middle pane, this pane is considered a drop target for additional columns.

It is a simple matter therefore to select columns from other load tables and drag these columns into the middle pane.

The source table shows where each column was dragged from. Although not the case in the tutorial, it is often common to have columns of the same name coming from different tables. In the example above the description column is acquired from the load_product, load_prod_group and load_prod_subgroup tables. In order that the model table can be created we need to assign these columns unique names, so for this example the last two columns in question have been renamed to group_description and subgroup_description.

There are a number of columns that do not have a source table. These columns have been added by WhereScape RED, and are added depending on earlier choices. A description of these columns follows.

Column name	Description
model_customer_key	The unique identifier (artificial key) for the model table. This key is used in the joins to the fact table. It is generated via an identity associated with the table, except for the date model table where it has the form YYYYMMDD. If you have changed the default behavior of RED not automatically add surrogate keys, this column will not have been added. See Settings - Repository Attributes and Options (see " Settings - Repository Identification " on page 79).
dss_start_date	Used for model history tables. This column provides a date time stamp when the model table record came into existence. It is used to ascertain which model table record should be used when multiple are available.

Column name	Description
dss_end_date	Used for model history tables. This column provides a date time stamp when the model table record ceased to be the current record. It is used to ascertain which model table record should be used when multiple are available.
dss_current_flag	Used for model history tables. This flag identifies the current record where multiple versions exist.
dss_version	Used for model history tables. This column contains the version number of a model table record. Numbered from 1 upwards with the highest number being the latest or current version. It forms part of the unique constraint for the business key of a model history table.
dss_update_time	Indicates when the record was last updated in the data warehouse.
dss_create_time	Indicates when the record was first created in the data warehouse

CREATE THE TABLE

Once the model table has been defined in the metadata we need to physically create the table in the database. This is achieved by right-clicking on the model name and selecting **Create (ReCreate)** from the pop up menu.

A results dialog box will appear to show the results of the creation.

The contents of this dialog are a message to the effect that the model table was created. A copy of the actual database create statement and if defined the results of any index create statements will be listed. For the initial create no indexes will be defined.

If the table was not created then ascertain and fix the problem. A common problem is a 'Duplicate column' where a column has the same name in two of the source tables. The best way of finding such a column is to double click on the list heading 'Col name'. This will sort the column names into alphabetic order. Another double click on the heading will sort the columns back into their create order.

The next section covers the *Generating the Model Table Update Procedure* (on page 469).

GENERATING THE MODEL TABLE UPDATE PROCEDURE

Once a model table has been defined in the metadata and created in the data base an update procedure can be generated to handle the joining of any tables and the update of the model table records.

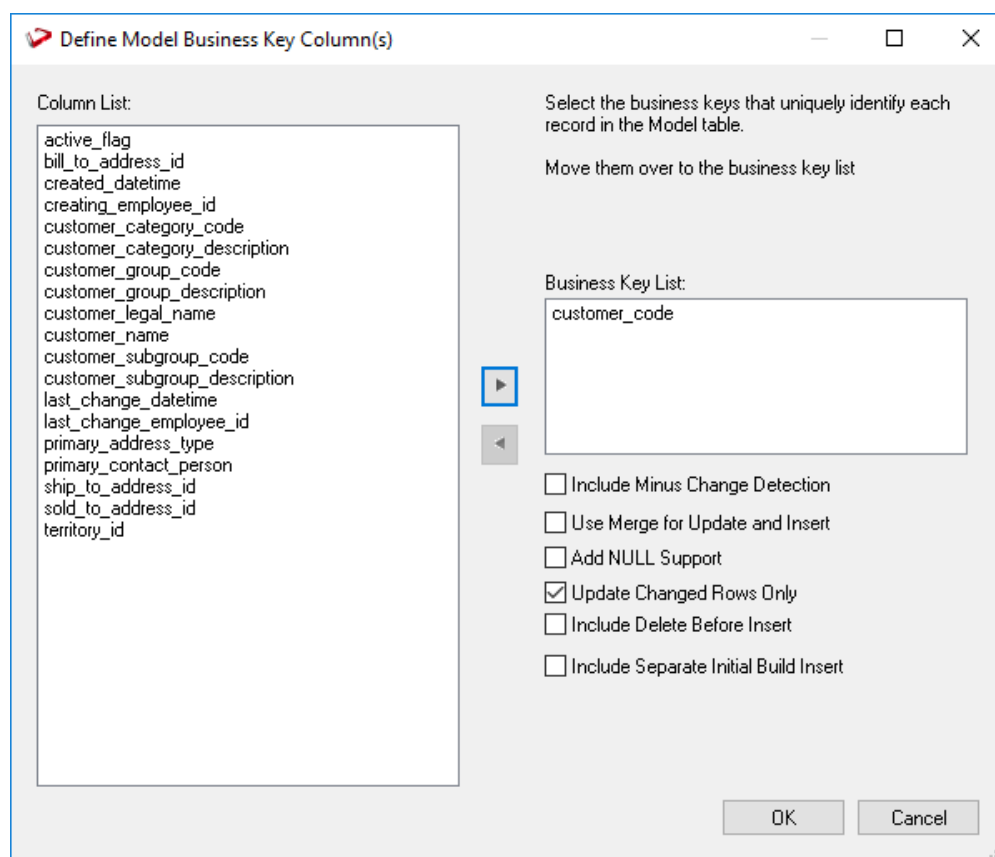
Note: You can also generate an update procedure via a template, refer to *Rebuilding Update Procedures* (on page 199) for details.

GENERATING A PROCEDURE

To generate a procedure, right-click on the model table in the left pane and select **Properties**. Click on the **Rebuild** button to start the process of generating the new procedure. A series of questions will be asked during the procedure generation based on the type of load information.

BUSINESS KEY DEFINITION

A dialog will appear asking for the business key that will uniquely identify each model table record. The source table from which the model table is derived would normally have some form of unique constraint applied. In most cases, this will be the business key. In the example below, the customer code is selected as the business key.



A business key can be made up of multiple columns, but it must provide a unique identifier. Where multiple columns uniquely and separately identify the model table, choose one to act as the primary business key. For example, a source table may have a unique constraint on both a product code and a product description. Therefore, the description as well as the code must be unique. It is of course possible to combine the two columns, but the normal practice would be to choose the code as the business key.

None of the columns chosen as the business key should ever contain a NULL value. See the note at the start of this chapter.

The **Include Minus Change Detection** checkbox will detect new rows using a minus sub-query rather than the default where not exists query. Enabling this option can significantly improve performance.

The **Use Merge for Update and Insert** checkbox will generate merge syntax. This option is only available for non-history model tables.

The **Include Separate Initial Build Insert** adds a second insert to the procedure to separately insert all data if the target model table is empty. This significantly improves performance with a large model table being loaded the first time.

LOCKING REQUEST MODIFIER

Source Table: Specify a locking request modifier to be applied to each source table during generated update procedures. By default, this is set to 'ACCESS' which locks each row being accessed, a blank entry will result in no locking clause in the generated procedure.

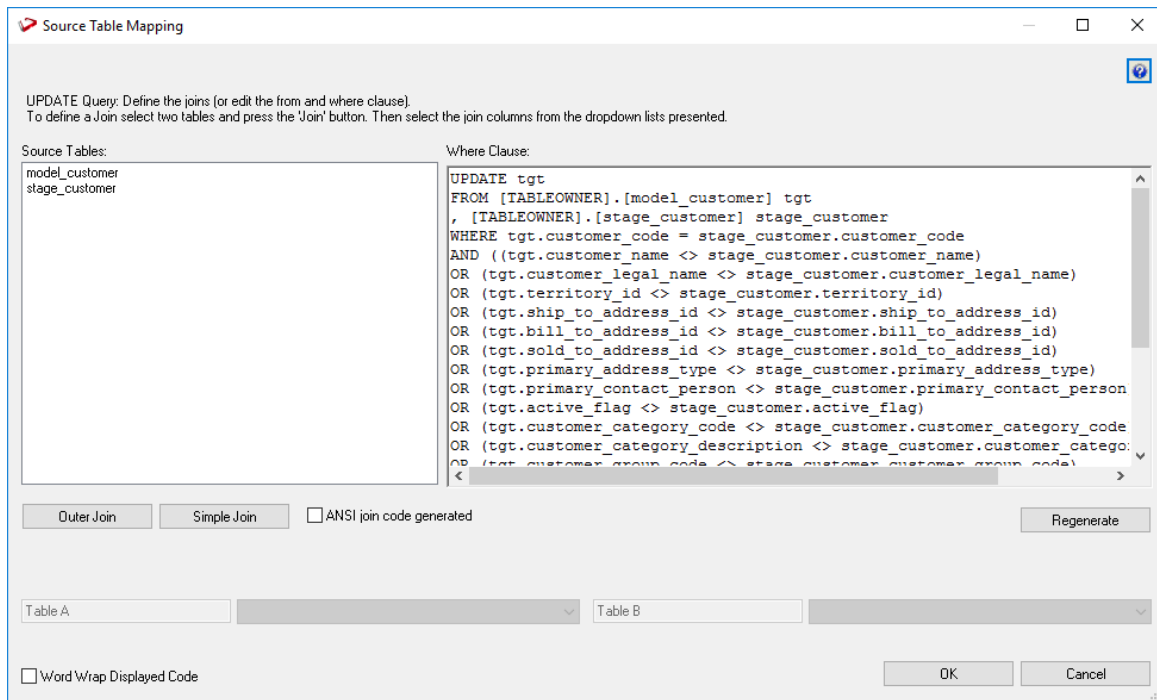
Source Table	
stage_customer	ACCESS

Source Table
Locking Request Modifier to be applied to the respective source tables for each DML statement. A blank entry will exclude the locking clause in generated procedure.

OK Cancel

SOURCE TABLE MAPPING

WhereScape RED generates a default update statement with a 'Where' clause to join source and target tables together. The default update statement can be edited via the following dialog.



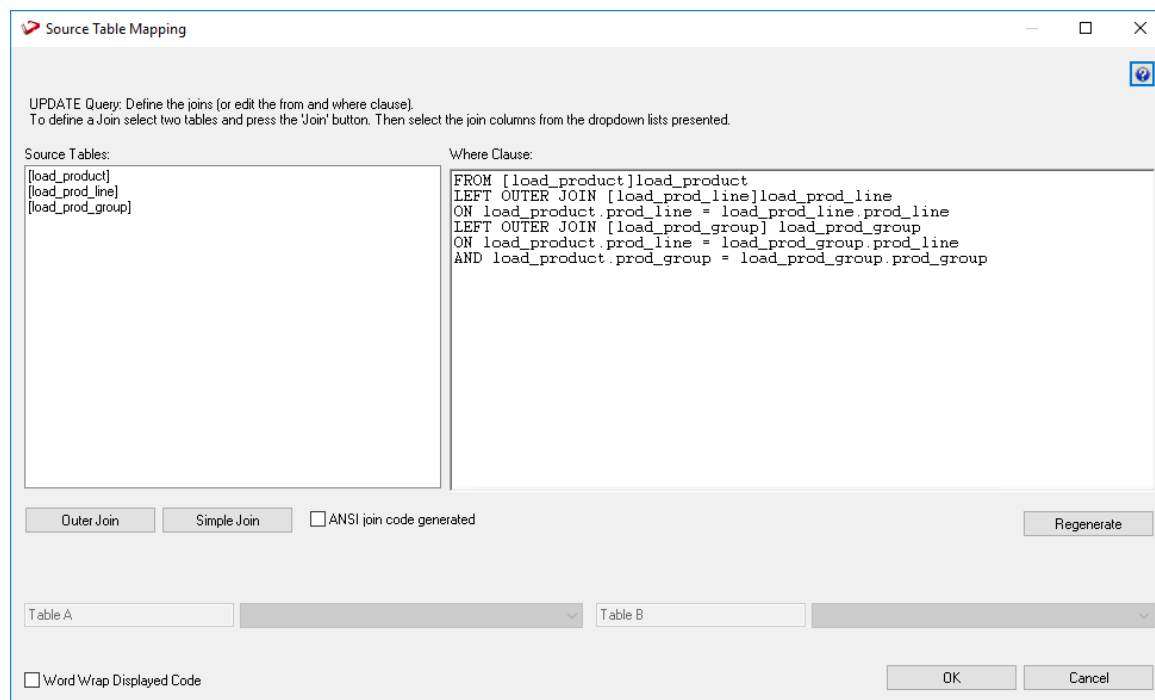
If there is more than one source table, additional joins will also have to be created. See Joining multiple source tables below.

Insert query where clause

If multiple source tables were used to build the model table, then a dialog box will appear prompting for the joins. This applies to the insert statement only in the generated procedure. See Joining multiple source tables for more information.

Joining multiple source tables

The example below shows the joining of the product, prod_line and prod_group tables as supplied with the tutorial data set.



Select two tables in the left box and then click one of the join buttons. The columns for the two tables then appear at the bottom of the dialog and one column is selected from each drop-down list to effect the join between the selected tables. In the example above, the load_product and load_prod_group tables are joined by two columns namely prod_line and group. In such a case, two joins are actioned for these two tables, so that both columns can be selected.

Simple Join

A simple join joins the two tables via either a 'Where' clause or from clause join (ANSI). A simple join only returns rows where data is matched in both tables. So for example, if table A has 100 rows and table B has a subset of 24 rows. If all the rows in table B can be joined to table A, then 24 rows will be returned. The other 76 rows from table A will not be returned.

Outer Join

An outer join joins the two tables, and returns all rows in the master table, regardless of whether or not they are found in the second table. Therefore, if the example above was executed with table A as the master table then 100 rows would be returned. 76 of those rows would have null values for the table B columns. When RED builds up a 'Where' clause join, it must place the outer join indicator next to the appropriate column. RED needs to know which table is master and which subordinate.

Select the join column from the master table first. In the example screen above, the table 'load_product' has had its column chosen and the column for the table 'load_prod_subgroup' is currently being chosen. This will result in the 'load_product' table being defined as the master, as per the example statement as shown in the 'Where' clause edit window above.

The results of this example select are that a row will be added containing product information, regardless of whether or not a corresponding prod_subgroup entry exists.

As the join columns are selected, the join statement is built up in the large edit window on the right. Once all joins have been made, the contents of this window can be changed if the join statement is not correct.

Once satisfied with the join clause click the **OK** button to proceed to the next step. This clause will be a combined from and 'Where' clause. This clause can of course be edited in the procedure that is generated if not correct.

Only ANSI Outer Joins are available in Teradata.

BUILDING AND COMPILING THE PROCEDURE

Once the above questions are completed, the procedure is built and compiled automatically. If the compile fails an error will be displayed along with the first few lines of error messages. Compile fails typically occur when the physical creation of the table was not done. If the compile fails for some other reason, the best approach is to use the procedure editor to edit and compile the procedure. The procedure editor will highlight all the errors within the context of the procedure.

Once the procedure has been successfully compiled it can either be executed interactively or passed to the scheduler.

Indexes

By default, a number of indexes will be created in the RED meta repository to support the model table. The primary index is the only Active index. Secondary indexes can add significant performance cost during updates in Teradata, so these are defined in the RED meta repository but are not active (so are not created on the table). An example of the type of indexes defined is as follows:

Index Name	Table	Columns	Type	Key Type	Rebuild Frequ...	Unique	Foreign	Pre Drop	Active	Partition	Parallel	Columns
model_product_idx_0	model_product	1	Index	Artificial	Never	Y	N	N	N		0	model_product_key
model_product_idx_A	model_product	1	Index	Business	Never	Y	N	N	N		0	code
model_product_idx_PR	model_product	1	Primary	Primary	Never	N	N	N	Y		0	code

This example shows three indexes being created. They are:

- 1 A primary key constraint placed on the artificial key for the model table.
- 2 A unique index placed on the business key for the model table.
- 3 The primary index of the model table.

Only the third kind of index is active. To activate one of the other indexes as a secondary index, click the active check box in the index properties dialog.

Additional indexes can be added, or these indexes changed. See the chapter on indexes for further details.

MODEL TABLE ARTIFICIAL KEYS

The artificial (surrogate) key for a model table is set via an identity column. This artificial key normally, and by default, starts at one and progresses as far as is required.

A WhereScape standard for the creation of special rows in the model table is as follows:

Key value	Usage
1 upwards	The normal model table artificial keys are numbered from 1 upwards, with a new number assigned for each distinct model table record.
0	Used as a join to the model table when no valid join existed. It is the normal convention in the WhereScape generated code that any model table business key that either does not exist or does not match is assigned to key 0.
-1 through -9	Used for special cases. The most common being where a model table is not appropriate for the record. For example, we may have a sales system that has a promotion model table. Not all sales have promotions. In this situation it is best to create a specific record in the model table that indicates that a fact table record does not have a promotion. The stage table procedure would be modified to assign such records to this specific key. A new key is used rather than 0 as we want to distinguish between records that are invalid and not appropriate.
-10 backward	Pseudo records. In many cases, we have to deal with different granularities in our fact data. For example, we may have a fact table that contains actual sales at a product SKU level and budget information at a product group level. The product model table only contains SKU based information. To be able to map the budget records to the model table we need to create these pseudo keys that relate to product groups. The values -10 and backwards are normally used for such keys. A template called 'Pseudo' is shipped with WhereScape RED to illustrate the generation of these pseudo records in the model table.

MODEL TABLE CUSTOM PROCEDURE

A second procedure can be created on every model table. This is called the custom procedure. Rather than modifying the generated procedure, it is often more practical to make additions to the generated code in a separate procedure. This allows for regeneration of the model table's update procedure without losing changes (and having to reapply them).

The generated procedure for a custom procedure is template code. That is, a procedure that declares and initializes a variable, does nothing and returns the correct return code and message for the WhereScape RED scheduler.

MODEL HISTORY TABLES

Model history tables are a special type of model table where new records are created when certain identified columns in the model table change.

With any model table we identify a business key that uniquely identifies the model table records. For example in the case of the product model table from the tutorial the product code is deemed to be the business key. The code uniquely identifies each product within the model table. The product may also have a name or description and various other attributes that distinguish it. (e.g. size, shape, color etc.). A common question when handling model tables is what do we do when the name or description changes. Do we want to track transactional records in other model table based only on the product code or do we also want to track records based on different descriptions.

An example:

code	description	product_group	sub_group
1235	15oz can of brussel sprouts	canned goods	sprouts

This product is sold for many years and we consequently have a very good history of sales and the performance of the product in the market. The company does a '20% extra for free' promotion for 3 months during which time it increases the size of the can to 18oz. The description is also changed to be '15 + 3oz can of brussel sprouts'. At the end of the promotion the product is reverted to its original size and the description changed back to its original name.

The question is do we want to track the sales of the product when it had a different description (using a model history table), or should the description of the product simply change to reflect its current name (a standard model table).

The decision is not a simple one and the advantages and disadvantages of each of the two choices is discussed below.

Model History Table

- Allows the most comprehensive analysis capabilities when just using the product model table.
- Complicates the analysis. Does not allow a continuous analysis of the product called '15oz can of brussel sprouts' when the description is used. This analysis is however still available through the code which has not changed.
- Adds considerable additional processing requirements to the building of any other model tables joined to this model table.

Model Table

- Does not allow specific analysis of the product during its size change. Note, however that this analysis will probably be available through the combination of a 'promotion' model table.
- Provides a continuous analysis history for the product called '15oz can of brussel sprouts'. An analysis via description and code will produce the same results.
- Simplifies analysis from an end user's perspective.

As mentioned above the choice is never a simple one. Even amongst experienced data warehouse practitioners there will be a variety of opinions. The decision must be based on the business requirements. In many cases keeping the analysis simple is the best choice, at least in the early stages of a data warehouse development. Model history tables do have a place, but there is nearly always an alternate method that provides equal or better results. In the example above a promotion model history table coupled with the product model table could provide the same analysis results whilst still keeping product only analysis simple and easy to understand.

GENERATING HISTORY TABLE UPDATE PROCEDURES

Business Key definition

This is almost identical to standard model tables.

Define Model Business Key Column(s)

Column List:

- customer_address
- customer_city
- customer_name
- customer_state

Select the business keys that uniquely identify each record in the Model table.

Move them over to the business key list

Business Key List:

- customer_code

Source System Supplied Start and End Dates

Include Delete Before Insert

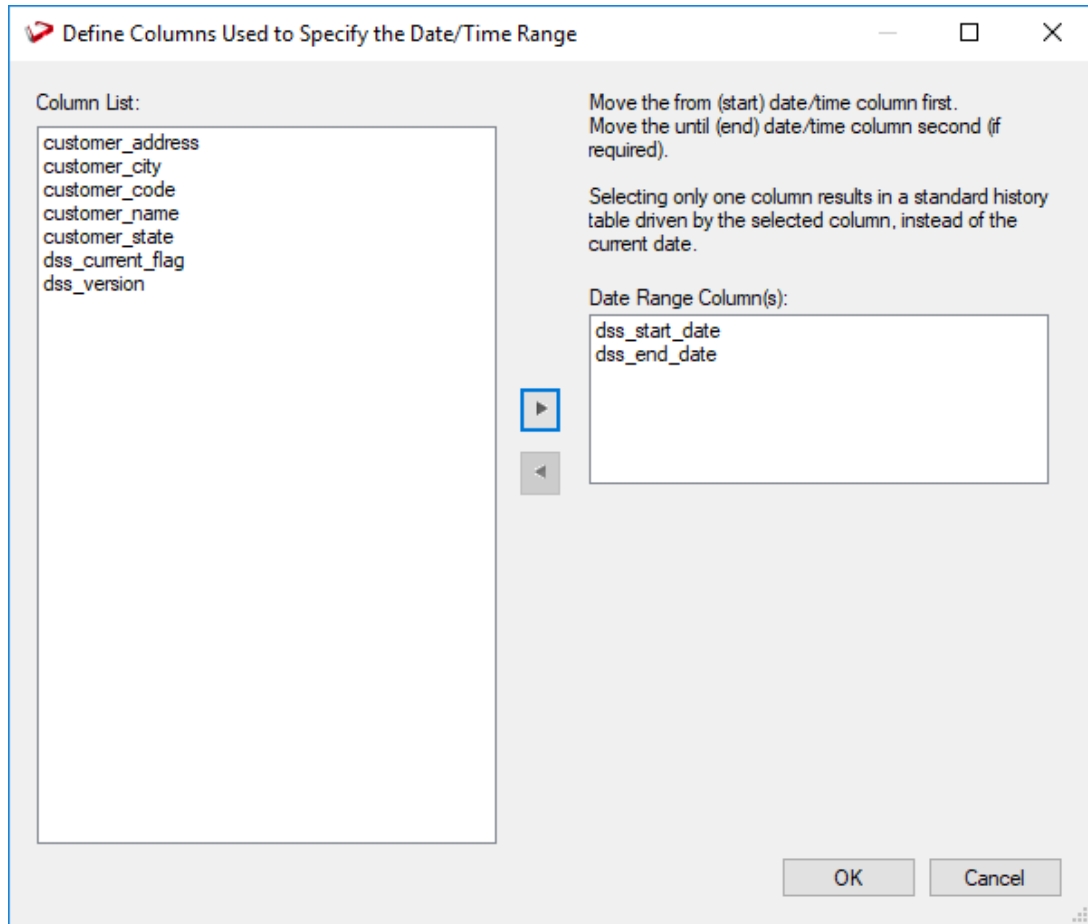
Include Separate Initial Build Insert

OK Cancel

An additional check-box, **Source System supplied Start and End dates** is supplied. This enables you to specify a start and end dates in the source data.

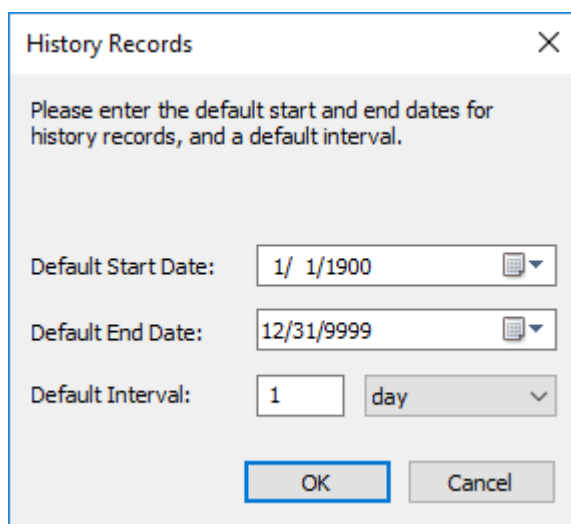
Source System supplied Start and End dates

If Source System supplied Start and End dates are enabled, the define start and end date dialog is displayed. Select the **start date** first, then the **end date** and click **OK**.



Default Start and End Dates

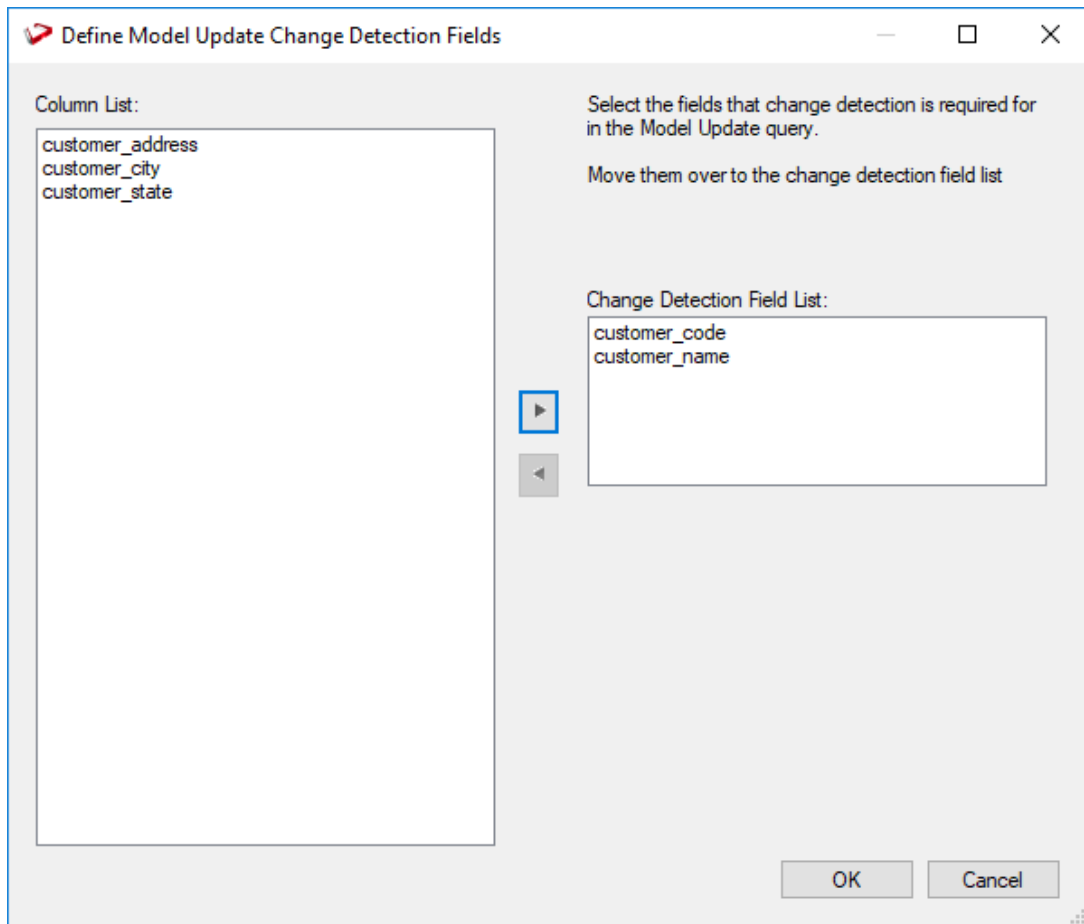
The `dss_start_date` and `dss_end_date` columns are used to track the start and end dates of history records. Defaults must be entered to ensure the null values are not placed in these columns. Use the history records dialog to enter these.



The default interval and interval type should also be chosen. These provide the difference between the end of the previous record and the start of the new record.

History columns

The change detection fields dialog requests the selection of the columns to be managed as history columns. Select the required columns and click **OK** to proceed. In the example below; columns `customer_code` and `customer_name` are to be managed as history columns.



Update query statement

This is the same as standard model tables.

Insert query where clause

This is the same as standard model tables.

Joining multiple source tables

This is the same as standard model tables.

MODEL TABLE COLUMN PROPERTIES

Each model table column has a set of associated properties. The definition of each property is described below:

If the Column name or Data type is changed for a column then the metadata will differ from the table as recorded in the database. Use the **Validate>Validate Table Create Status** menu option or the right-click menu to compare the metadata to the table in the database. A right-click menu option of **Alter table** is available when positioned on the table name after the validate has completed. This option will alter the database table to match the metadata definition.



TIP: If a database table's definition is changed in the metadata then the table will need to be altered in the database. Use the **Validate>Validate Table Create Status** to compare metadata definitions to physical database tables. The option also provides the ability to alter the database table, through a pop-up menu option from the validated table name.

A sample **Properties** screen is as follows:

Dimension Column model_date.calendar_date

Properties

- Transformation
- Language Mapping

General

Table Name	model_date
Column Name	calendar_date
Column Title	calendar date
Column Description	The calendar date for this row.

Physical Definition

Column Order	1949
Data Type	date
Null Values Allowed	<input checked="" type="checkbox"/>
Default Value	
Character Set	
Format	
Character Comparison/Sorting	
Compress	<input type="checkbox"/>

Meta Definition

Numeric	<input type="checkbox"/>
Additive	<input type="checkbox"/>
Attribute	<input checked="" type="checkbox"/>
End User Layer Display	<input checked="" type="checkbox"/>
Business Key	<input checked="" type="checkbox"/>
Artificial Key	<input type="checkbox"/>
Key Type (0,A,B,C...)	A

Source Details

Source Table	model_date
Source Column	calendar_date

Column Name

Database-compliant name of the column.
Dialog Opening Value: calendar_date

<- Update Update -> OK Cancel Help

The two special update keys enable you to update the column and step either forward or backward to the next columns properties.

Table Name

Database-compliant name of the table that contains the column. [Read-only].

Column Name

Database-compliant name of the column. A good practice is to only use alphanumerics, and the underscore character. Changing this field alters the table's definition.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Title

Name that the business uses to refer to the column. It does not affect the physical table definition, but rather provides input to the documentation and to the view **ws_admin_v_dim_col** which can be used to assist in the population of a end user tool's end user layer. As such it is a free form entry and any characters are valid.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Description

This field contains the description for the column. It might contain information on where and how the column was acquired. For example if the column is sourced from multiple tables or is a composite or derived column then this definition would normally describe the process used to populate the column. This field is used in the documentation and is available via the view **ws_admin_v_dim_col** . This field is also stored as a comment against the column in the database.

Column Order

Numeric value that controls the relative order of columns in the database create statement. The lowest numbered column will appear first in the table. Although this affects the physical table definition no action will be taken unless the table is re-created in the database. The columns can be re-numbered based on the existing order by choosing the **Respace order number** pop-up menu option when positioned over any column in the table. This action will number the columns in increments of 10 starting at 10. In addition to a simple change of the order field, the column order can be changed by first displaying the columns in the middle pane and then using drag and drop to move the columns around. This drag and drop process will automatically renumber the columns as required.

Data Type

Database-compliant data type that must be valid for the target database. Typical Teradata databases often have integer, numeric(), varchar(), char(), date and timestamp data types. See the database documentation for a description of the data types available. Changing this field alters the table's definition.

Null Values Allowed

Determines whether the table column can hold NULL values or whether a value is always mandatory.

Default Value

Initial value that is assigned to the column when a row is inserted into the table but no value is specified for the column.

Character Set

Database-compliant table column character-set used for storage. Select Latin or Unicode.

Format

Database-compliant table column format. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. As such it is a free form entry and any characters are valid. Typically format masks are only used on numeric fields. Example: #,###0.00. It is not worth the effort of populating this field unless it can be utilized by the end user tools in use.

Character Comparison/Sorting

Determines how the column character values are treated for comparison and sorting operations. Choose from: case specific, not case specific, uppercase case specific or uppercase not case specific.

Compress

Indicates whether the table column values are compressed when stored.

Compress/Compress Value

Optional list of values to be compressed. By default, only NULL is compressed if no list of values is specified.

Numeric

Indicates whether the table column holds values that are numeric. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Additive

Indicates whether the table column holds values that are additive. This implies that the column can be summed when performing data grouping in a query. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Attribute

Indicates whether the table column holds values that are descriptive, and/or are used for grouping/summing. An attribute is defined as a column that is non factual. For example we may have an order number, or a invoice number stored in the fact table. Such columns are considered attributes, rather than facts. This checkbox is therefore normally only relevant for fact tables. This checkbox does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tools end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

End User Layer display

Indicates whether the table column is available/visible to end users. If set the documentation will include the column in the glossary and in the user documentation. It is also used to decide what columns appear in the view **ws_admin_v_dim_col**. Typically columns such as the artificial key would not be enabled for end user display.

Business Key

Indicates whether the column is part of the BUSINESS key, which is defined during the update procedure build. [Normally maintained automatically]. Multiple columns can form the primary business key. This indicator is set and cleared by WhereScape RED during the dimension update procedure generation process. This checkbox should not normally be altered.

Key Type

Key type that is assigned and used when generating the table's update procedure and indexes. [Normally maintained automatically]. It can be altered here, but this should only be done if the consequences are fully understood and tested. The supported values are:

Key type	Meaning
0	The artificial key. Set when the key is added during drag and drop table generation.
1	Component of all business keys. Indicates that this column is used as part of any business key. It is considered part of any lookup on that table and has the key type set to 1. Set when the column is added during drag and drop table generation.
2	Indicates that this column is a model table join. Used on model tables to indicate the model keys to other model tables. Results in indexes being defined for the columns. Set during the update procedure generation for a model table, based on information from the staging table.
3	Not used in WhereScape RED for Teradata.
4	Not used in WhereScape RED for Teradata.
5	Indicates a column is a start date column.
6	Indicates a column is a end date column.
7	History column indicator. Used on dimension history tables to indicate that the column is being managed as a history column within the context of a dimension history table. Set when a column is identified during the dimension history update procedure generation.
A	Indicates that the column is part of the primary business key. Set whenever a business key is defined as part of an update procedure generation.
B-Z	Indicates that the column is part of a secondary business key. Only used during index generation and not normally set.

Source Table

Identifies the source table where the column's data comes from. This source table is normally a load table, stage table or another model table within the data warehouse. If the column was sourced from multiple tables, then the normal practice is to record one of the tables in this field and a comment listing all of the other tables in the **Source strategy** field. This field is used when

generating a procedure to update the model table. It is also used in the track back diagrams and in the documentation.

Source Column

Identifies the source column where the column's data comes from. Such a column is normally a load table column, which in turn may have been a transformation or the combination of multiple columns. This may also be a model table key where a model table is being joined.

Transformation

See *Model Table Column Transformations* (on page 486). [Read-only].

MODEL TABLE COLUMN TRANSFORMATIONS

Each model table column can have a transformation associated with it. The transformation will be included in the generated procedure and will be executed as part of the procedure update. The transformation must therefore be a valid SQL construct that can be included in a **Select** statement. For example we could have a transformation of 'load_order_line.qty * 0.125' to calculate a tax column of 12.5%.

Click the **Transformation** tab to enter a transformation.

It is possible to do transformations on model table columns. It is recommended that transformations are not performed on columns that are model keys or the business keys for the table.

The transformation screen is as follows:

The screenshot shows a configuration window for a transformation. The title bar reads "Dimension Column model_date.calendar_date". On the left, a sidebar contains "Properties", "Transformation" (selected), and "Language Mapping". The main area has a "Target:" field with "calendar_date" and a "Paste" button. Below it, a "Source:" field has "model_date.calendar_date" and another "Paste" button. A large text area is labeled "Column Transformation Code (must execute within a SQL SELECT statement)". To the right of this area is a "Function Set:" dropdown menu currently set to "Default Teradata". Below the dropdown is a tree view showing "Functions" and "Available Columns". At the bottom left, there is a checkbox for "Word Wrap Displayed Code". Below that are two input fields: "Function Syntax:" and "Function Desc:". At the bottom right, there are five buttons: "<- Update", "Update ->", "OK", "Cancel", and "Help".

Note: Transformations are only put into effect when the procedure is re-generated.

See *Transformations* (on page 593) for more details.

CHAPTER 17

FACT TABLES

A Fact Table is normally defined, for our purposes, as a table with facts (measures) and dimensional keys that allow the linking of multiple dimensions. It is normally illustrated in the form of a Star Schema with the central fact table and the outlying dimensions.

The ultimate goal of the Fact Table is to provide business information to the end user community. In many cases, different types of fact tables are required to address different end user requirements.

IN THIS CHAPTER

Detail Fact Tables	488
Fact Table Column Properties	495
Fact Table Column Transformations	500
Fact Table Language Mapping	501

DETAIL FACT TABLES

A detail fact table is normally a transactional table that represents the business data at its lowest level of granularity. In many ways these tables reflect the business processes within the organization. Such fact tables are usually large and are focused on a specified analysis area.

There may be quite a large number of detail fact tables in a data warehouse implementation, of which only a few are used on any regular basis by the end user community. The disadvantage of such fact tables is that they provide isolated pools of information. Although joined by conformed dimensions, it is still often difficult to answer queries across the various analysis areas. They do however provide the ultimate drill down for all information and also the platform on which to build higher level and more complex fact tables. In terms of the time dimension detail fact tables are typically at a daily or even hourly granular level.

An example of a detail fact table may be the sales, or orders fact tables that have a daily granularity, and show all sales by product, by customer etc. on a given day.

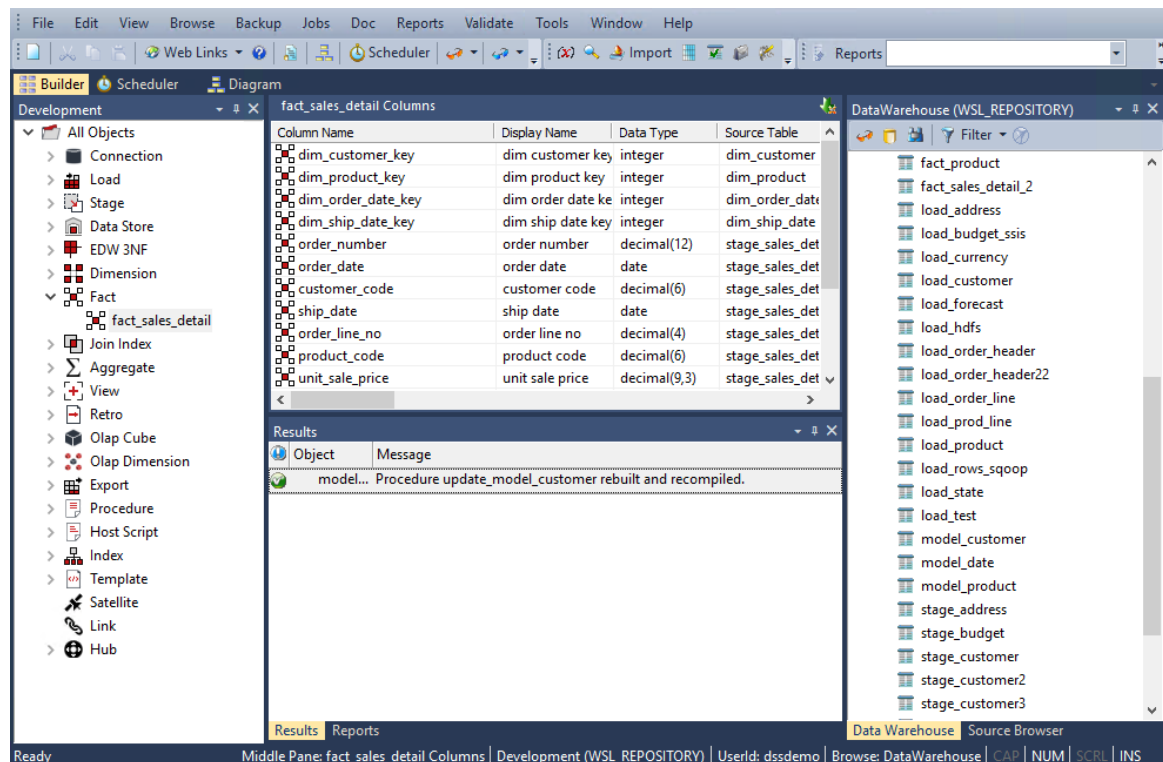
CREATING DETAIL FACT TABLES

A detail fact table is typically created by dragging a staging table onto a fact table list.

In the following example screen the fact table list has been produced by double clicking on the Fact Table object group under the Sales project. A list is produced showing the existing fact tables.

A **fact_sales_detail** detail fact table can be created by selecting the stage_sales_detail name in the right pane of the builder window, holding down the left mouse button and dragging the table into the middle (fact table list) pane.

Once released, the dialog to create the fact table will start.



When a staging table is dragged into the list window (middle pane) all fact tables are by default *detail fact tables*. If manually creating a table, then the table type can be selected under the Properties of the fact table.

Detail Fact Columns

The columns for a detail fact table are typically those of the staging table. Such fact tables typically contain a wide range of measures and attributes as well as the business keys used to look up the artificial dimension keys. These business keys should be included whenever possible as they provide a means of rebuilding the dimensional link. If size prohibits their inclusion it will probably be necessary to backup or archive all source data to ensure that the fact table can be rebuilt.

These fact tables normally contain a large number of attributes such as dates, which have not been converted to dimensions. Also contained would be information such as order numbers, invoice numbers etc.

See the first tutorial for an example of a fact table creation.

GENERATING THE DETAIL FACT UPDATE PROCEDURE

Once a detail fact table has been defined in the metadata and created in the database, an update procedure can be generated to handle the update of the fact table.

Note: You can also generate an update procedure via a template, refer to *Rebuilding Update Procedures* (on page 199) for details.

Generating a Procedure

- To generate a procedure, right-click on the fact table in the left pane and select **Properties**.
- Click on the **Rebuild** button to start the process of generating the new procedure.

Define Fact Procedure Type and Options

- The first dialog displayed when generating a detail fact table update procedure is the define Fact Procedure Type and Options dialog.
- Several other fields need to be set or adjusted on this dialog to ensure the required type of update procedure is generated.
- Once the required options have been selected, click **OK** to proceed to the next dialog.

Option	Value
Template	None
Business Key Columns	order_number, order_line_no
Parameters	
Include Initial Load Insert	<input type="checkbox"/>
Create Storage Index After CTAS	<input type="checkbox"/>
Batch Processing	
Process by Batch	<input type="checkbox"/>
Explicit Locking	
Include Explicit Lock	<input checked="" type="checkbox"/>
Delete Processing	
Delete Before Insert	No
Update Processing	
Process Method	Insert/Update
Source Table Locking	
stage_sales_detail	ACCESS
Insert Method	
Include Insert Statement	<input type="checkbox"/>
Update Method	
Include Update Statement	<input type="checkbox"/>

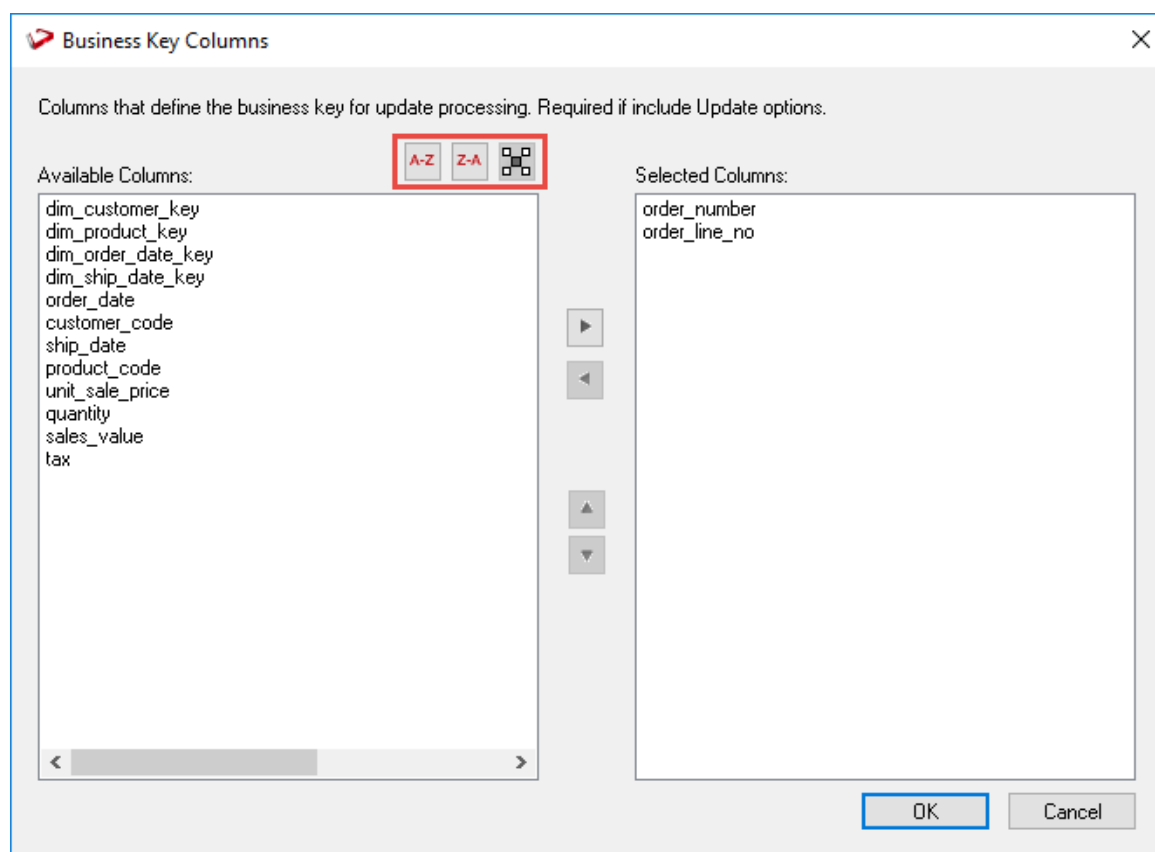
Buttons: Prev, Next, OK, Cancel, Help

Template

Enables you to generate update procedures via a *template* (see "*Rebuilding Update Procedures*" on page 199).

Define Fact Business Key Columns

The next dialog displayed is the define fact business key columns dialog, asking for the business key that will uniquely identify each fact table record. The source table from which the fact table is derived would normally have some form of unique constraint applied. In most cases this will be the business key. In the example below the `order_number` and `order_line_no` are selected for the business key list.



TIP: Use the column name **ascending/descending** buttons to sort column names. To revert to the meta column order, click on the **meta column order** button.

A business key can be made up of multiple columns, but it must provide a unique identifier. Where multiple columns separately uniquely identify rows in the fact table, choose one to act as the primary business key. For example a source table may have a unique constraint on both a product code and a product description. Therefore the description as well as the code must be unique.

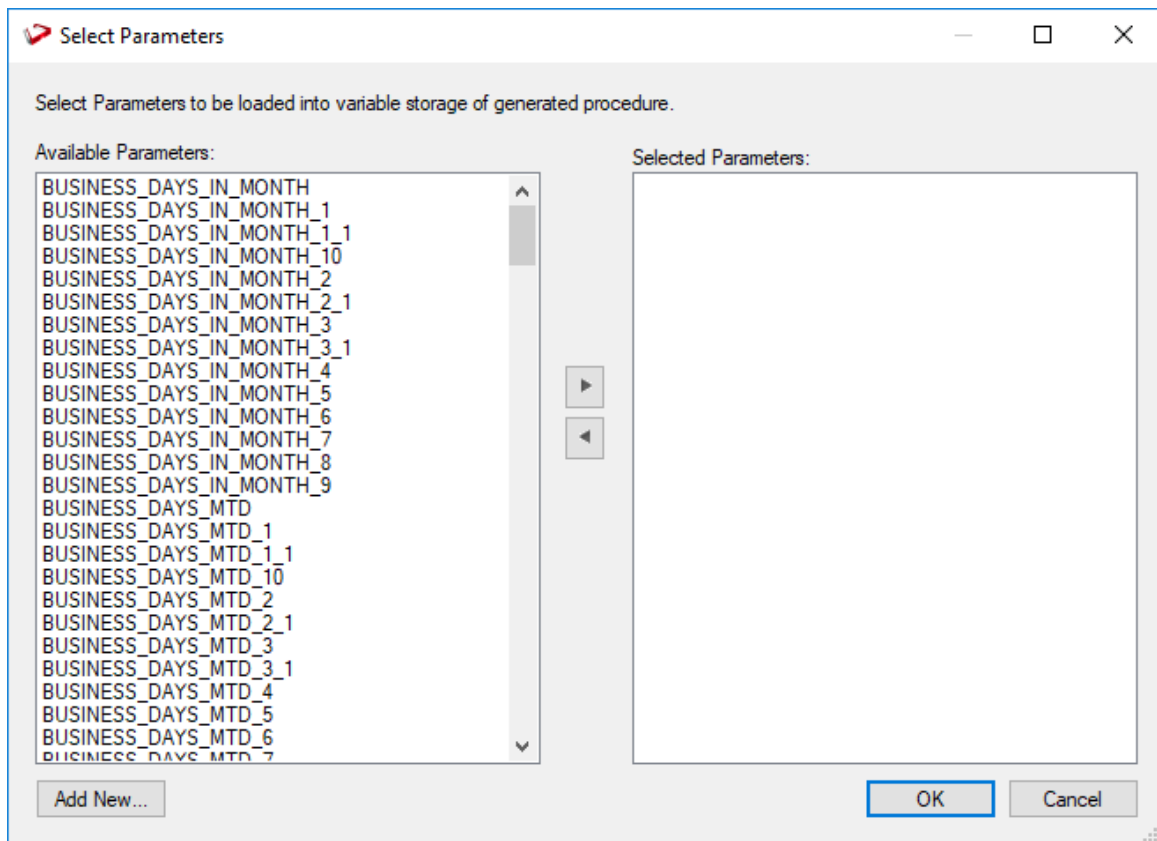
NULL Values: none of the columns chosen as the business key should ever contain a NULL value.

Select Parameters

The next dialog displayed is the Select Parameters dialog. If WhereScape RED parameters exist in the metadata, the following dialog is displayed. Any parameters selected in this dialog (by moving them to right side), are included in the generated update procedure as variables. The procedure will include code to retrieve the value of the parameter at run time and store it in the declared variable.

Select a parameter by clicking on the parameter and then the > arrow to move it to the right column; then click **OK**.

If the desired parameter doesn't exist in the metadata yet, a new parameter can be added by clicking on the **Add New** button on the bottom leftmost corner of the Select Parameters dialog.



The variables can also be used in column transformations and in the from/where clause for the update procedure.

Some databases have a 30 character limit for variable names. WhereScape RED ensures the variables added for any parameters are less than 30 characters long by creating variable names in the form v_ followed by the first 28 characters of the parameter name.

For example, a parameter called MINIMUM_ORDER_NUMBER_SINCE_LAST_SOURCE_LOAD will be available as the variable v_MINIMUM_ORDER_NUMBER_SINCE_L.



TIP: WhereScape RED parameters should be unique within the first 28 characters to avoid conflicting variables names.

See *Parameters* (on page 143) for more information on WhereScape RED Parameters.

Include initial load insert

The **include initial load insert** option adds an additional insert statement to the update procedure that runs if the target Data Store Object is empty. The benefit of this is improved performance inserting into an empty table without performing any checks to see if rows already exist. The default for this field is off (i.e. an initial insert statement is not added to the procedure).

Process by batch

The **process by batch** field allows the user to select a column to use to break up the data being processed in a loop based on the distinct values in the column. The update procedure loops on this column and performs the delete, update and/or insert for each value. If the column chosen is a date datatype (date, datetime or timestamp), then the user is able to specify yearly, monthly, daily or column level looping. The default for this field is off (do not do batch processing).

Batch Processing Field - allows selecting a field to batch process on. If you select a date field you will have the ability to process by date part. If you select a join field to process by you can choose an attribute of that related table to group by.

Delete before insert

The **delete before insert** option enables a delete statement to be added to the update procedure before any update or insert statement. This is a particularly useful option for purging old data and for updates based on a source system batch number. If this option is set Issue a Warning if a Delete Occurs and Delete Where Clause fields are enabled. The default for this field is off (i.e. a delete statement is not added to the procedure).

Truncate - if this option is chosen, the delete is ignored and the Fact Object is truncated.

Issue warning if a delete occurs: sets the procedure to a warning state if any deletes occur.

Delete 'Where' clause: the delete where clause that is appended to the generated delete statement to constrain the rows deleted.

Update processing

The update **process method** option allows you to use the Merge statement instead of two separate Insert and Update statements. The default value for this option is Insert/Update.

Source table locking

This section allows a locking request modifier to be specified for each source table. The specified locking request modifier is applied to each source table during generated update procedures. By default this is set to 'ACCESS' which locks each row being accessed, a blank entry will result in no locking clause in the generated procedure.

Insert method

The **include insert statement** option includes an insert statement in the procedure to insert new rows in the Fact Object. If this option is chosen, then the **New rows only** option is available. Choosing this option, displays the **code type** drop-down, enabling the generated update statement to use either a sub-select with *except* change detection or *checksum* change detection to work out what rows in the source table(s) are new. If this option is turned off, the update procedure will not contain an insert statement. The default for this field is on (i.e. an insert statement is included).

Update method

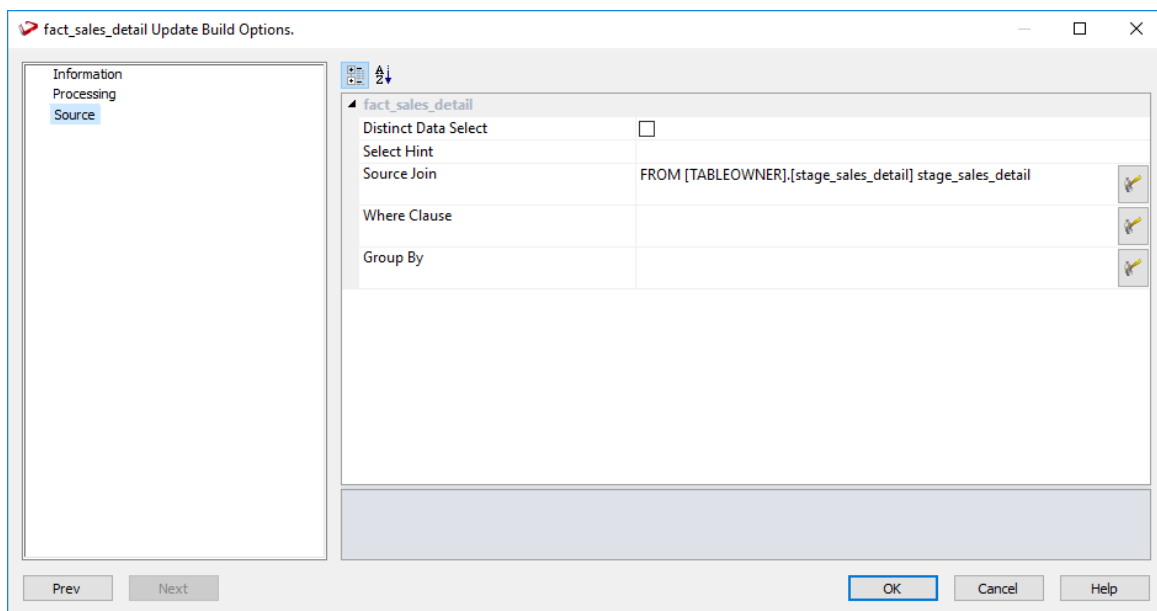
The **include update statement** option includes an update statement in the procedure to update changing rows in the Fact Object. If this option is chosen, then the **Changed rows only** option is available. Choosing this option, displays the **code type** drop-down, enabling the generated update statement to use either a sub-select with *except* change detection or *checksum* change detection to work out what rows in the Fact Object require updating. If this option is turned off, the update procedure will not contain an update statement. The default for this field is on (i.e. an update statement is included).

Changed and New Rows Only

Choosing this option, displays the **code type** drop-down, enabling the generated merge statement to use either a sub-select with **Join** change detection or **Minus** change detection to work out what rows in the Fact Object are new or which require updating.

SOURCE TAB

This dialog is used to join source tables, add 'Where' clauses and specify group by clauses.



Distinct data select

The **distinct data select** option ensures duplicate rows are not added to the Data Store Object. This is achieved by the word **DISTINCT** being added to the source select in the update procedure.

The default for this field is not set (i.e. duplicates are not removed).

As source table joins should have been performed in the stage table, see *Generating the Staging Update Procedure* (on page 337) for more details.

FACT TABLE COLUMN PROPERTIES

Each fact table column has a set of associated properties. The definition of each property is defined below.

If the **Column name** or **Data type** is changed for a column then the metadata will differ from the table as recorded in the database.

Use the **Validate>Validate Table Create Status** menu option to compare the metadata to the table in the database.

A right-click menu option of **Alter table** is available when positioned on the table name after the validate has completed. This option will alter the database table to match the metadata definition.



TIP: If a database table's definition is changed in the metadata then the table will need to be altered in the database.

Use the **Validate>Validate Table Create Status** to compare metadata definitions to physical database tables. The option also provides the ability to alter the database table, through a pop-up menu option from the validated table name.

A sample **Properties** screen is as follows:

The two special update keys allow you to update the column and step either forward or backward to the next column's properties.

ALT-Left Arrow and **ALT-Right Arrow** can also be used instead of the two special update keys.

Table Name

Database-compliant name of the table that contains the column. [Read-only].

Column Name

Database-compliant name of the column. Typically column-naming standards exclude the use of spaces etc. A good practice is to only use alphanumeric, and the underscore character. Changing this field alters the table's definition.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Title

Name that the business uses to refer to the column, which is included in the RED-generated documentation and can be used in the end user layer of other tools. [Does NOT affect the physical database table]. As such it is a free form entry and any characters are valid.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Description

This field contains the description for the column. It may be a description from a business user's point of view. This field might additionally contain information on where and how the column was acquired. For example if the column is sourced from multiple tables or is a composite or derived column then this definition would normally describe the process used to populate the column. This field is used in the documentation and is available via the view **ws_admin_v_dim_col**. This field is also stored as a comment against the column in the database.

Column Order

Numeric value that controls the relative order of columns in the database create statement. The lowest numbered column will appear first in the table. Although this affects the physical table definition no action will be taken unless the table is re-created in the database. The columns can be re-numbered based on the existing order by choosing the **Respace Order Number** pop-up menu option when positioned over any column in the table. This action will number the columns in increments of 10 starting at 10. In addition to a simple change of the order field, the column order can be changed by first displaying the columns in the middle pane and then using drag and drop to move the columns around. This drag and drop process will automatically renumber the columns as required.

Data Type

Database-compliant data type that must be valid for the target database. Typical Teradata databases often have integer, numeric(), varchar(), char(), date and timestamp data types. See the database documentation for a description of the data types available. Changing this field alters the table's definition.

Null Values Allowed

Determines whether the table column can hold NULL values or whether a value is always mandatory.

Default Value

Initial value that is assigned to the column when a row is inserted into the table but no value is specified for the column.

Format

Optional format mask that can be used in end user tools. [Does NOT affect the physical database table]. As such it is a free form entry and any characters are valid. Typically format masks are only used on numeric fields. Example: #,###0.00. It is not worth the effort of populating this field unless it can be utilized by the end user tools in use.

Numeric

Indicates whether the table column holds values that are numeric. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Additive

Indicates whether the table column holds values that are additive. This implies that the column can be summed when performing data grouping in a query. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Attribute

Indicates whether the table column holds values that are descriptive, and/or are used for grouping/summing. An attribute is defined as a column that is non factual. For example we may have an order number, or a invoice number stored in the fact table. Such columns are considered attributes, rather than facts. This checkbox is therefore normally only relevant for fact tables. This checkbox does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tools end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

End User Layer Display

Indicates whether the table column is available/visible to end users. If set the documentation will include the column in the glossary and in the user documentation. It is also used to decide what columns appear in the view **ws_admin_v_dim_col**. Typically columns such as the artificial key would not be enabled for end user display.

Business Key

Indicates whether the column is part of the BUSINESS key, which is defined during the update procedure build. [Normally maintained automatically]. Multiple columns can form the primary business key.

Key Type

Key type that is assigned and used when generating the table's update procedure and indexes. [Normally maintained automatically]. It can be altered here, but this should only be done if the consequences are fully understood and tested. The supported values are:

Key type	Meaning
0	The artificial key. Set when the key is added during drag and drop table generation.
1	Component of all business keys. Indicates that this column is used as part of any business key. For example: By default the dss_source_system_key is added to every dimension table. It is considered part of any lookup on that table and has the key type set to 1. Set when the column is added during drag and drop table generation.
2	Indicates that this column is a dimensional join. Used on fact tables to indicate the dimension keys. Results in bitmap indexes being built for the columns. Set during the update procedure generation for a fact table, based on information from the staging table.

Key type	Meaning
3	Slowly changing column indicator. Used on dimension tables to indicate that the column is being managed as a slowly changing column within the context of a slowly changing dimension. Set when a column is identified during the dimension update procedure generation.
4	Previous value column indicator. Used on dimension tables to indicate that the column is being managed as a previous value column. The source column identifies the parent column. Set during the dimension creation.
A	Indicates that the column is part of the primary business key. Set whenever a business key is defined as part of an update procedure generation.
B-Z	Indicates that the column is part of a secondary business key. Only used during index generation and not normally set.

KPI Column Type

Only used by **KPI** fact tables. This field defines the column type for the KPI Fact Table. Refer to the KPI table creation section for more details on this field.

Source Table

Identifies the source table where the column's data comes from. This source table is normally a stage table within the data warehouse. If the column was sourced from multiple tables, then the normal practice is to record one of the tables in this field and a comment listing all of the other tables in the **Source strategy** field. This field is used when generating a procedure to update the fact table. It is also used in the track back diagrams and in the documentation.

Source Column

Identifies the source column where the column's data comes from. Such a column is normally a stage table column, which in turn may have been a transformation or the combination of multiple columns. This may also be a dimensional key where a dimension is being joined.

Transformation

See **Fact Table Column Transformations** (on page 500). [Read-only].

Join

Indicates whether the table column is used in a table join. [Normally maintained automatically but can be optionally changed to override the default join logic used in the generated update procedure]. The **Source Table** and **Source Column** fields will provide the dimension table's side of the join. The options for this field are: False, True, Manual and Pre Join.

Setting this field to Manual changes the way the dimension table is looked up during the update procedure build. It allows you to join the dimension manually in the Cursor mapping dialog (where the 'Where' clause is built).

Setting this field to Pre Join activates the **Pre Join Source Table** field and allows you to select a table from the drop-down list.

Pre Join Source Table

Indicates the table from which the pre joined column was sourced. When the Join option is set to False, this field becomes inactive. When the Join option is set to True or Manual, this field is set to the current table name. When the Join option is set to Pre Join, then you can select the required table from the drop-down list.

FACT TABLE COLUMN TRANSFORMATIONS

Each fact table column can have a transformation associated with it. The transformation will be included in the generated procedure and will be executed as part of the procedure update.

The transformation must therefore be a valid SQL construct that can be included in a **Select** statement.

For example we could have a transformation of 'load_order_line.qty * 0.125' to calculate a tax column of 12.5%. Click the **Transformation** tab to enter a transformation.

The screenshot shows a configuration window for a fact column transformation. The window title is "Fact Column fact_customer.customer_code". On the left, there is a sidebar with tabs: "Properties", "Transformation" (selected), and "Language Mapping". The main area contains the following fields and controls:

- Target:** A text box containing "customer_code" with a "Paste" button to its right.
- Source:** A text box containing "load_customer.customer_code" with a "Paste" button to its right.
- Column Transformation Code (must execute within a SQL SELECT statement):** A large empty text area.
- Function Set:** A dropdown menu currently set to "Default Teradata".
- Functions Panel:** A panel on the right showing a tree view with "Functions" and "Available Columns".
- Word Wrap Displayed Code:** A checkbox that is currently unchecked.
- Function Syntax:** A text box.
- Function Desc.:** A text box.
- Buttons:** At the bottom, there are five buttons: "<- Update", "Update ->", "OK" (highlighted in blue), "Cancel", and "Help".

Note: Transformations are only put into effect when the procedure is re-generated.

Microsoft Analysis Services 2005+ Tabular Mode Tables: For Tabular Mode table column transformations, **Default DAX** is the only applicable Function Set for **after load** transformations.

See *Transformations* (on page 593) for more details.

FACT TABLE LANGUAGE MAPPING

The Fact Properties screen has a tab called **Language Mapping**.

Select the language from the drop-down list and then enter the translations for the **Business Display Name** and the **Description** in the chosen language.

The translations for these fields can then be pushed through into OLAP cubes.

The screenshot shows a dialog box titled "Fact Table Column fact_sales_detail.quantity" with a close button (X) in the top right corner. On the left, there is a vertical navigation pane with three tabs: "Properties", "Transformation", and "Language Mapping", with "Language Mapping" selected and highlighted in blue. The main area of the dialog is divided into three sections:

- Language :** A dropdown menu currently set to "French".
- Language Settings:** A section containing two text input fields. The first field is labeled "Business Display Name:" and contains the text "quantity". The second field is unlabeled and also contains "quantity".
- Business Definition:** A section containing two text input fields. Both fields contain the text "Quantity of product sold (i.e. number of product units)." and have small up/down arrow icons on their right sides.

At the bottom of the dialog, there are five buttons: "<- Update", "Update ->", "OK", "Cancel", and "Help". The "OK" button is highlighted with a blue border.

CHAPTER 18

AGGREGATION

Two types of aggregate tables are discussed.

The first is where all non-additive facts and one or more dimensions are removed from a fact table. Typically this results in a smaller table that can answer a subset of the queries that could be posed against the fact table. This aggregate table still maintains full integrity to the remaining dimensions, and consequently reflects all changes to those dimensions.

The second type, we will call an aggregate summary, or summary table. This table includes additive measures and in some cases hierarchical elements of one or more of the dimensions providing a rolled-up summary of the fact table data. For example we may choose to deal at product group level rather than product SKU which is the granularity of the dimension.

IN THIS CHAPTER

Creating an Aggregate Table	504
Creating an Aggregate Summary Table	504
Aggregate Table Column Properties.....	504
Aggregate Table Column Transformations	509

CREATING AN AGGREGATE TABLE

- 1 In the left pane double click on the aggregate group to list the aggregates in the middle pane and set aggregates as the drop target.
- 2 From the Data Warehouse browse (right) pane drag a fact table into the middle pane. Remove any columns that will not make sense at an aggregated level. For example, `dss_fact_table_key`, any business keys, any non-additive facts, any measures that relate to detail (e.g. unit price).
- 3 Create the aggregate table in the database by right-clicking on the aggregate and selecting **Create(ReCreate)**.
- 4 Create a procedure to update the aggregate by right-clicking on the aggregate, selecting **Properties** and selecting (**Build Procedure...**) in the Update Procedure field. You will be asked for the date in the fact table that is to be used as the basis for rebuilding changes in the fact table. The aggregate update process looks at any records that have been updated in the fact table in the last 7 days (by default). It then rebuilds all the information for the dates that have been altered.
You will also be asked to specify a locking request modifier to be applied to each source table during generated update procedures. By default this is set to 'ACCESS' which locks each row being accessed, a blank entry will result in no locking clause in the generated procedure.

CREATING AN AGGREGATE SUMMARY TABLE

The creation of a summary table proceeds initially in the same way as an aggregate table.

- 1 In the left pane double click on the aggregate group to list the aggregates in the middle pane and set aggregates as the drop target.
- 2 From the Data Warehouse browse (right) pane drag a fact table into the middle pane. Remove any columns that will not make sense at an aggregated level. For example, `dss_fact_table_key`, any business keys, any non-additive facts, any measures that relate to detail (e.g. unit price).
- 3 Drag over columns from dimensions that are linked to the fact table. Delete the dimension keys to allow a rollup to the level of the dimension elements.
- 4 In the properties of the aggregate table change the **Table Type** to **Summary**.
- 5 Create the aggregate summary table in the database by right-clicking on the aggregate and selecting **Create(ReCreate)**.
- 6 Create a procedure to update the aggregate summary by right-clicking on the aggregate, selecting **Properties** and selecting (**Build Procedure...**) in the Update Procedure field. The aggregate summary table is totally rebuilt each time the procedure is executed.

AGGREGATE TABLE COLUMN PROPERTIES

Each aggregate table column has a set of associated properties. The definition of each property is defined below:

If the **Column name** or **Data type** is changed for a column then the metadata will differ from the table as recorded in the database. Use the **Validate>Validate Table Create Status** menu option to compare the metadata to the table in the database. A right-click menu option of **Alter Table** is available when positioned on the table name after the validate has completed. This option will alter the database table to match the metadata definition.



TIP: If a database table's definition is changed in the metadata, then the table need to be altered in the database. Use the **Validate>Validate Table Create Status** to compare metadata definitions to physical database tables. The option also provides the ability to alter the database table, through a pop-up menu option from the validated table name.

A sample **Properties** screen is as follows:

The screenshot shows the 'Properties' dialog box for an aggregate column named 'agg_sales_analysis.quantity'. The dialog is organized into several sections:

- General:** Table Name (agg_sales_analysis), Column Name (quantity), Column Title (quantity), Column Description (The quantity of the product ordered.).
- Physical Definition:** Column Order (90), Data Type (decimal(8)), Null Values Allowed (checked), Default Value, Character Set, Format, Character Comparison/Sorting, Compress (unchecked).
- Meta Definition:** Numeric (checked), Additive (checked), Attribute (unchecked), End User Layer Display (checked), Key Type (None).
- Code Generation:** Aggregate (Sum).
- Source Details:** Source Table (model_order_line), Source Column (quantity).

At the bottom, there are buttons for '< Update', 'Update >', 'OK', 'Cancel', and 'Help'. A 'Column Name' section at the bottom provides a description: 'Database-compliant name of the column. Dialog Opening Value: quantity'.

The two special update keys enable you to update the column and step either forward or backward to the next column's properties. **ALT-Left Arrow** and **ALT-Right Arrow** can also be used instead of the two special update keys.

Table Name

Database-compliant name of the table that contains the column. [Read-only].

Column Name

Database-compliant name of the column. Typically, column-naming standards exclude the use of spaces etc. A good practice is to only use alphanumerics, and the underscore character. Changing this field alters the table's definition.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Title

Name that the business uses to refer to the column, which is included in the RED-generated documentation and can be used in the end user layer of other tools. [Does NOT affect the physical database table]. As such it is a free form entry and any characters are valid.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Description

This field contains the description for the column. It may be a description from a business user's point of view. This field might additionally contain information on where and how the column was acquired. For example, if the column is sourced from multiple tables or is a composite or derived column then this definition would normally describe the process used to populate the column. This field is used in the documentation and is available via the view `ws_admin_v_dim_col`. This field is also stored as a comment against the column in the database.

Column Order

Numeric value that controls the relative order of columns in the database create statement. The lowest numbered column appears first in the table. Although this affects the physical table definition no action is taken unless the table is re-created in the database. The columns can be re-numbered based on the existing order by choosing the **Respace Order Number** pop-up menu option when positioned over any column in the table. This action numbers the columns in increments of 10 starting at 10. In addition to a simple change of the order field, the column order can be changed by first displaying the columns in the middle pane and then using drag and drop to move the columns around. This drag and drop process automatically renumbers the columns as required.

Data Type

Database-compliant data type that must be valid for the target database. Typical Teradata databases often have integer, numeric(), varchar(), char(), date and timestamp data types. See the database documentation for a description of the data types available. Changing this field alters the table's definition.

Null Values Allowed

Determines whether the table column can hold NULL values or whether a value is always mandatory.

Default Value

Initial value that is assigned to the column when a row is inserted into the table but no value is specified for the column.

Character Set

Database-compliant table column character-set used for storage, Latin or Unicode.

Format

Optional format mask that can be used in end user tools. [Does NOT affect the physical database table]. As such it is a free form entry and any characters are valid. Typically, format masks are only used on numeric fields. Example: #,###0.00. It is not worth the effort of populating this field unless it can be utilized by the end user tools in use.

Character Comparison/Sorting

Determines how the column character values are treated for comparison and sorting operations. Choose from: case specific, not case specific, uppercase case specific or uppercase not case specific.

Compress

Indicates whether the table column values are compressed when stored.

Compress/Compress Value

Optional list of values to be compressed. By default, only NULL is compressed if no list of values is specified.

Numeric

Indicates whether the table column holds values that are numeric. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Additive

Indicates whether the table column holds values that are additive. This implies that the column can be summed when performing data grouping in a query. This is normally only relevant for fact tables. It does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tool's end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

Attribute

Indicates whether the table column holds values that are descriptive, and/or are used for grouping/summing. An attribute is defined as a column that is non factual. For example we may have an order number, or a invoice number stored in the fact table. Such columns are considered attributes, rather than facts. This check box is therefore normally only relevant for fact tables. This check box does not affect the physical table definition, but rather provides input to the view **ws_admin_v_dim_col** which can be used to assist in the population of an end user tools end user layer. The use of this field is not relevant unless it can be utilized by the end user tools.

End User Layer Display

Indicates whether the table column is available/visible to end users. If set, the documentation will include the column in the glossary and in the user documentation. It is also used to decide what columns appear in the view **ws_admin_v_dim_col**. Typically, columns such as the artificial key would not be enabled for end user display.

Key Type

Key type that is assigned and used when generating the table's update procedure and indexes. [Normally maintained automatically]. It can be altered here, but this should only be done if the consequences are fully understood and tested. The supported values are:

Key type	Meaning
0	The artificial key. Set when the key is added during drag and drop table generation.
1	Component of all business keys. Indicates that this column is used as part of any business key. For example: By default the <code>dss_source_system_key</code> is added to every dimension table. It is considered part of any lookup on that table and has the key type set to 1. Set when the column is added during drag and drop table generation.
2	Indicates that this column is a dimensional join. Used on fact tables to indicate the dimension keys. Results in bitmap indexes being built for the columns. Set during the update procedure generation for a fact table, based on information from the staging table.
3	Slowly changing column indicator. Used on dimension tables to indicate that the column is being managed as a slowly changing column within the context of a slowly changing dimension. Set when a column is identified during the dimension update procedure generation.
4	Previous value column indicator. Used on dimension tables to indicate that the column is being managed as a previous value column. The source column identifies the parent column. Set during the dimension creation.
A	Indicates that the column is part of the primary business key. Set whenever a business key is defined as part of an update procedure generation.
B-Z	Indicates that the column is part of a secondary business key. Only used during index generation and not normally set.

Aggregate

Indicates whether aggregation is required and the aggregation method used, e.g. Sum, Count or Group by.

Source Table

Identifies the source table where the column's data comes from. This source table is normally a fact table or a dimension table within the data warehouse. If the column was sourced from multiple tables, then the normal practice is to record one of the tables in this field and a comment listing all of the other tables in the **Source Strategy** field. This field is used when

generating a procedure to update the aggregate table. It is also used in the track back diagrams and in the documentation.

Source Column

Identifies the source column where the column's data comes from. Such a column is normally a fact table column or a dimension table column, which in turn may have been a transformation or the combination of multiple columns.

Transformation

See *Aggregate Table Column Transformations* (on page 509). [Read-only].

Join

Indicates whether the table column is used in a table join. [Normally maintained automatically but can be optionally changed to override the default join logic used in the generated update procedure]. The **Source Table** and **Source Column** fields will provide the dimension table's side of the join. The options for this field are: False, True, Manual and Pre Join.

Setting this field to Manual changes the way the dimension table is looked up during the update procedure build. It allows you to join the dimension manually in the Cursor mapping dialog (where the 'Where' clause is built).

Setting this field to Pre Join activates the **Pre Join Source Table** field and allows you to select a table from the drop-down list.

Pre Join Source Table

Indicates the table from which the pre joined column was sourced. When the Join option is set to False, this field becomes inactive. When the Join option is set to True or Manual, this field is set to the current table name. When the Join option is set to Pre Join, then you can select the required table from the drop-down list.

AGGREGATE TABLE COLUMN TRANSFORMATIONS

Each aggregate table column can have a transformation associated with it. The transformation will be included in the generated procedure and will be executed as part of the procedure update. The transformation must therefore be a valid SQL construct that can be included in a **Select** statement. For example we could have a transformation of 'load_order_line.qty * 0.125' to calculate a tax column of 12.5%. Click the **Transformation** tab to enter a transformation.

The screenshot shows a dialog box titled "Aggregate Column agg_product.budget_quantity". On the left, there is a sidebar with "Properties" and "Transformation" tabs, with "Transformation" selected. The main area contains the following fields and controls:

- Target:** budget_quantity (with a Paste button)
- Source:** fact_sales_analysis.budget_quantity (with a Paste button)
- Column Transformation Code (must execute within a SQL SELECT statement):** A large empty text area.
- Function Set:** Default SQL Server (dropdown menu)
- Functions:** A tree view showing "Functions" and "Available Columns".
- Word Wrap Displayed Code
- Function Syntax:** An empty text field.
- Function Desc.:** An empty text field with a scroll bar.

At the bottom of the dialog, there are buttons: "< Update", "Update >", "OK", "Cancel", and "Help".

Note: Transformations are only put into effect when the procedure is re-generated.

See **Transformations** (on page 593) for more details.

CHAPTER 19

JOIN INDEXES

Join indexes are used to perform one or more of the following tasks in Teradata:

- 1 Replicate all or part of a single table using a new primary index
- 2 Join multiple tables in a pre-join table
- 3 Aggregates one or more columns of one or more tables

WhereScape RED supports all of these uses.

IN THIS CHAPTER

Creating a Join Index	512
-----------------------------	-----

CREATING A JOIN INDEX

Drag and Drop

- 1 In the left pane double click on the **join index** group to list the join indexes in the middle pane and set join indexes as the drop target.
- 2 From the Data Warehouse browse (right) pane drag a table into the middle pane.
- 3 The new object dialog box will appear and will identify the new object as a Join index and will provide a default name based on the join index name.
- 4 Either accept this name or enter the name of the join index and click **OK** to proceed.

Join Index Properties

At this stage change the storage options if desired and click on **OK**.

If prototyping, and the join index is simple (i.e. one source table) then you can create the join index automatically by answering **Create and Load** to the next question and specifying the primary index. Otherwise proceed to the next section.

Pre-joined Join Indexes

If joining multiple tables in a pre-join table, add the columns from the other tables.

Define the join between the source tables using the 'Where' clause builder (right-click on the join index and select **Build From/Where clause**).

Define source table joins for Join Index

Define the Joins (or edit the where clause).

load_order_header | load_order_line | Add Join

order_number | order_number | Function Set: Default Teradata

```
FROM [TABLEOWNER].[load_order_header] load_order_header
JOIN [TABLEOWNER].[load_order_line] load_order_line
ON load_order_header.order_number = load_order_line.order_number
```

> Functions
Available Columns

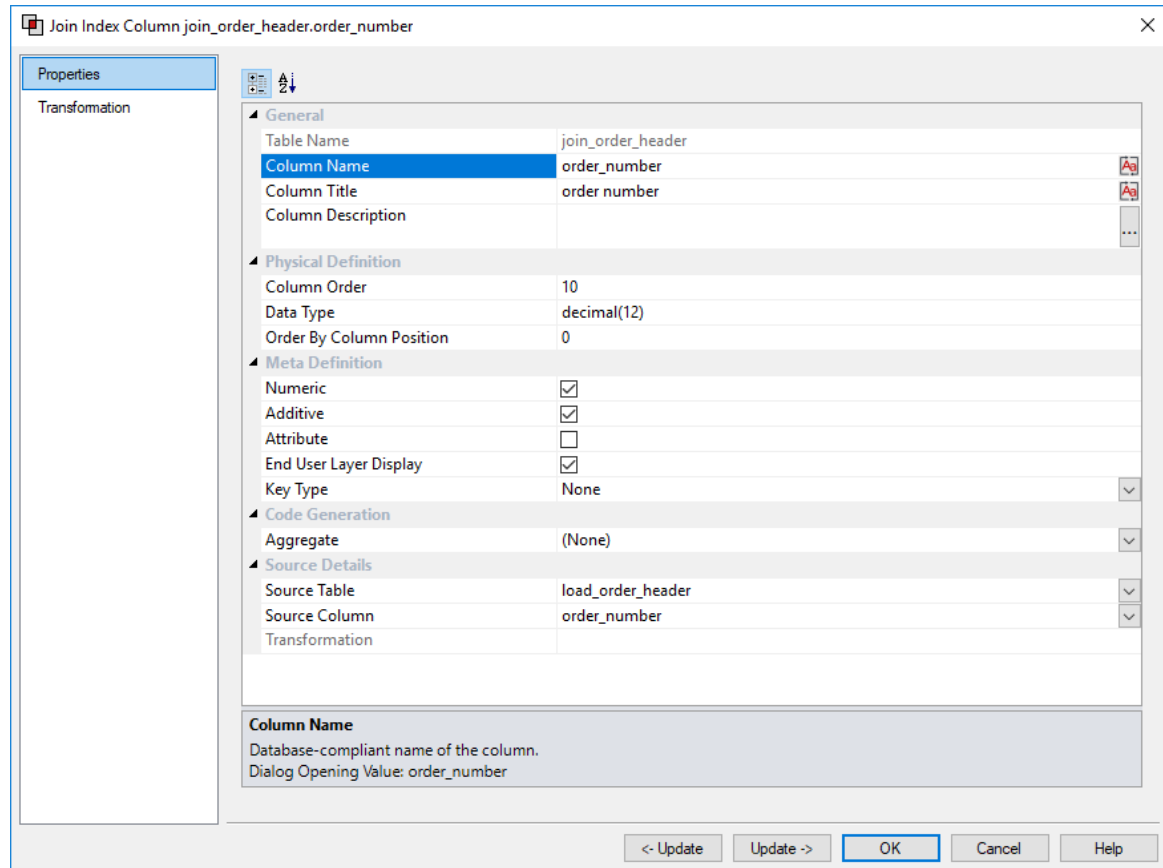
Function Syntax:

Function Description:

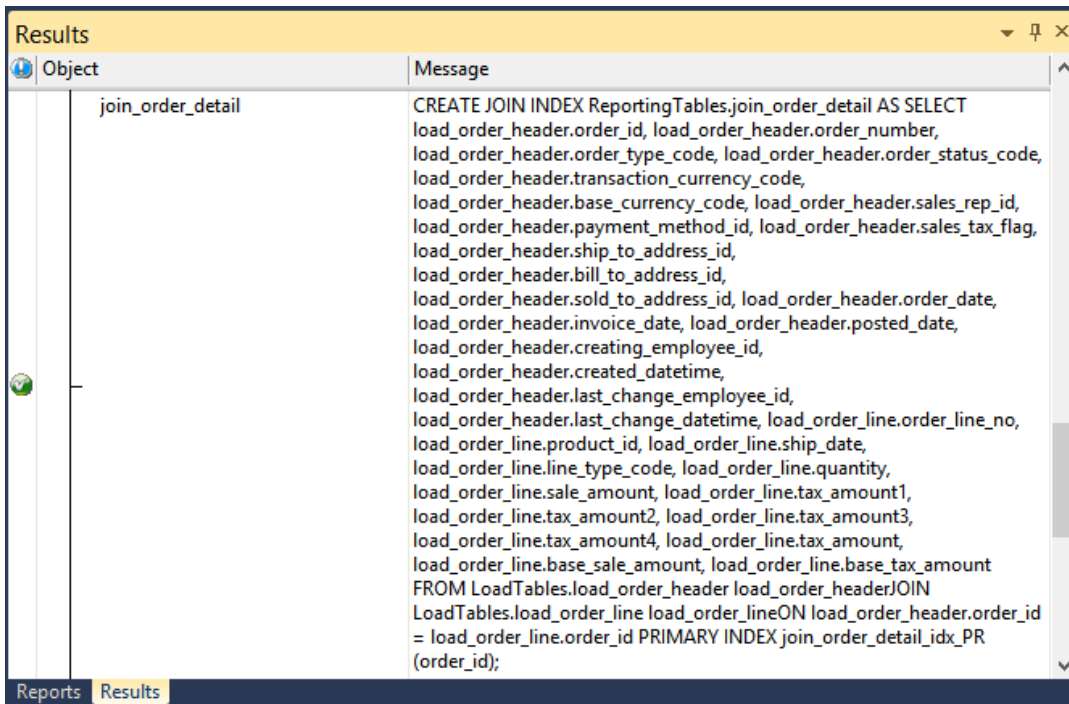
Word wrap displayed code

OK Cancel Help

If aggregation is required, proceed to the next section, otherwise remove any checks from the **Sum**, **Count** and **Group by** check boxes on each column of the join index.



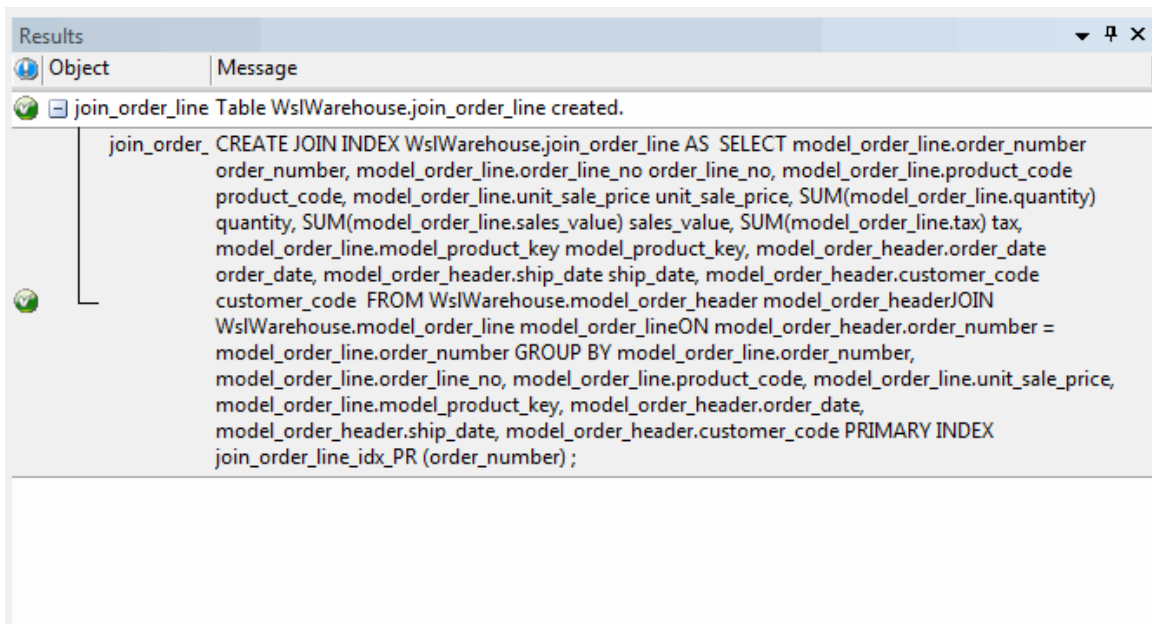
Create the join index in the database by right-clicking on the join index and selecting **Create(ReCreate)**.



Aggregated Join Indexes

To build an aggregated join index, edit each column and set the Sum, Count and Group by check-boxes as appropriate.

Create the aggregate table in the database by right-clicking on the aggregate and selecting **Create(ReCreate)**.



CHAPTER 20

VIEWS

Views are normally created to manage locking in Teradata, to join tables together for presentation to users or to provide additional methods for updating underlying tables. In our tutorials we create a view on each dimension table to provide an access path to the dimension tables with a built in locking clause.

Views can be created from any table type.

Note: WhereScape RED version 8.1.1.0 and above no longer supports the Dimension View object type when creating new objects. However, users that have existing Dimension View objects in their data warehouse can retain and continue to use them.

IN THIS CHAPTER

One to One Views	516
Dimension Views for Aliasing.....	518
Compound Views, Facts and Dimensions	519
Creating a Custom View	526
View Aliases	527

ONE TO ONE VIEWS

A one to one view is a database view of a dimension table. It may be a full or partial view. It is typically used to provide an access path to the dimension tables with a built in locking clause.

In many data warehouses views are built as part of the end user layer, but creating them in the data warehouse means they are available regardless of the end user tools used.

The process for creating a view is as follows:

- 1 Double click on **View** in the left pane.
- 2 Browse to the data warehouse in the right pane.
- 3 Drag a table from the right pane into the center pane.
 - The dialog box that displays defaults the object type to a view.
 - Change the view name as required, and click **ADD**.
- 4 The View properties dialog displays:

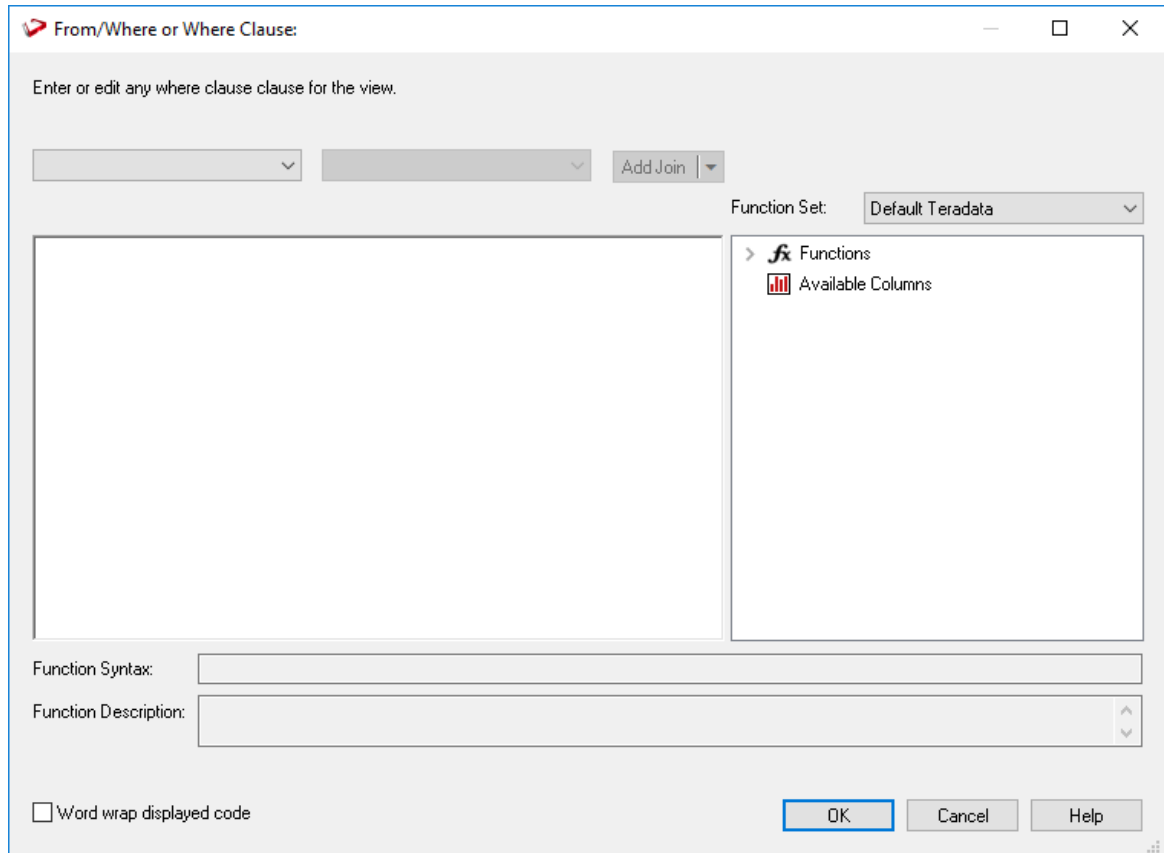
The screenshot shows the 'View view_customer' dialog box with the following fields and controls:

- View Name:** view_customer
- View Type:** View
- Unique Short Name:** view_customer (maximum 22 characters)
- Business Display Name (EUL):** view_customer
- Description:** View of customer table
- Update Procedure:** update_view_customer (with Edit and Regenerate buttons)
- Custom Procedure:** (None) (with Rebuild and Regenerate buttons)
- Distinct Data Select:** (highlighted with a red box)
- From/Where or Where Clause:** (empty text area, highlighted with a red box)
- Table Locking Mode:** LOCK ROW FOR ACCESS (dropdown menu, highlighted with a red box)
- Mnemonic (EUL):** (empty text field)
- Timestamps:**
 - Metadata Structure: 2011-07-04 11:18:44.000000
 - Database: 2017-03-14 21:48:48.250000
 - Database Altered: 2017-03-14 21:48:48.250000

Buttons at the bottom: OK, Cancel, Help.

Change the following properties, if desired:

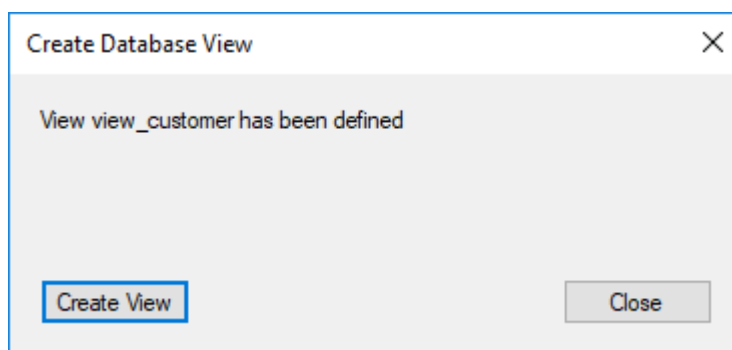
- Tick the **Distinct Data Select** check box if you want the view to return only distinct values.
- You can enter a **From/Where** clause, but this can be done later. The ellipses button can be used to open the From/Where Clause editor dialog:



- The **Table Locking Mode** can be changed at this point. The default is LOCK ROW FOR ACCESS, all other Teradata locking modes are available for use.

5 Click **OK**.

6 A dialog box displays indicating that the view has been defined. Click **Create view** to create the view in Teradata.



DIMENSION VIEWS FOR ALIASING

A view of a Dimension table used to create an alternative logical object for a Dimension table. It may have some or all columns renamed to match the aliased name of the view. Dimension Views may be used to look up surrogate keys when building staging tables.

The process for creating a dimension view is as follows:

- 1 Double click **View** in the left pane.
- 2 Browse to the data warehouse in the right pane.
- 3 Drag a dimension table from the right pane into the middle pane.
 - The dialog box defaults the object type to **View**.
 - Change the view name as required, and click **Add**.
- 4 The **View Column Definition** dialog appears to provide a means of re-mapping some of the column names in the View if required. A description of the fields in this dialog are provided below.
- 5 The View Properties window appears—change the Table Type to **Dimension View**.
- 6 Select the **Distinct Data Select** check box if you want the Dimension View to return only different values.
- 7 Click **OK**. A dialog box confirms that the View has been defined—click **Create View** to create the view in Teradata.

View Column Re-mapping

The View Column Definition dialog is used for the automated re-mapping of certain column names. It provides an easy method for changing the column names for a large number of columns when creating a view. The various actions undertaken as a result of entries in this dialog can all be done or reversed manually by changing the individual column properties. The various fields are described below:

View Column Definition

The column names for the view being created can be modified by filling in the following form. If the Default button is pressed nothing will be changed.

Remove Column Prefix: Add Column Prefix:
 --> order_

Remove Business Display Prefix: Add Business Display Prefix:
 --> order

Change Column Names for Specific Columns

Old Column Name:	New Column Name:
dim_date_key	dim_order_date_key
calendar_date	order_date
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>

OK Default

Remove Column Prefix

If the columns in the source table are prefixed then that prefix can be removed by entering it in this field. An example may be a date dimension which has each column prefixed with date_ (e.g. date_month, date_year, etc.). If this field is left blank then no removal action is taken.

Add Column Prefix

If required a prefix can be added to each column name. This is particularly useful when defining a date dimension view where you would like each column to be prefixed by the date type.

Remove Business Display Prefix

As per the column names it may be required to remove a prefix from the business display fields. If so enter the prefix to remove in this column.

Add Business Display Prefix

The business display fields are used in the creation of the glossary. It is therefore quite useful to prefix these display fields with a quire identifier for the view being created. It is assumed that these business display names will be carried forward to the end user layer. Enter a value in this field to prefix the business display name fields for each column. It is normal to include a space at the end of this field.

Old Column Name

Up to five individual column names can be re mapped. Enter the column name as it appears in the source table in one of the 'old column name' fields in order to re map that column name. The business display name is also changed to match.

New Column Name

Place a new column name alongside any existing column name you wish to re map. in the example dialog above a column named 'calendar_date' is being renamed to 'order_date' in the view.

COMPOUND VIEWS, FACTS AND DIMENSIONS

A compound view is used to join tables and views together for presentation to users. They simplify user access, particularly if multiple dimension history tables need to be accessed in a single query.

The process for creating a compound view is as follows:

- 1 Double click on **View** in the left pane.
- 2 Browse to the data warehouse in the right pane.
- 3 Drag an existing one-to-one view from the right pane into the center pane. Change the view name as required, click **ADD**.
- 4 The view property defaults dialog will appear. A 'Where' clause could be entered, but this can be done later using the 'Where' clause builder. The table locking mode can also be changed at this point. The default is lock for access, all other Teradata locking modes are available for use. Change the table type to be a **dimension view** or **fact view** (as appropriate) if a star schema presentation layer is being built.

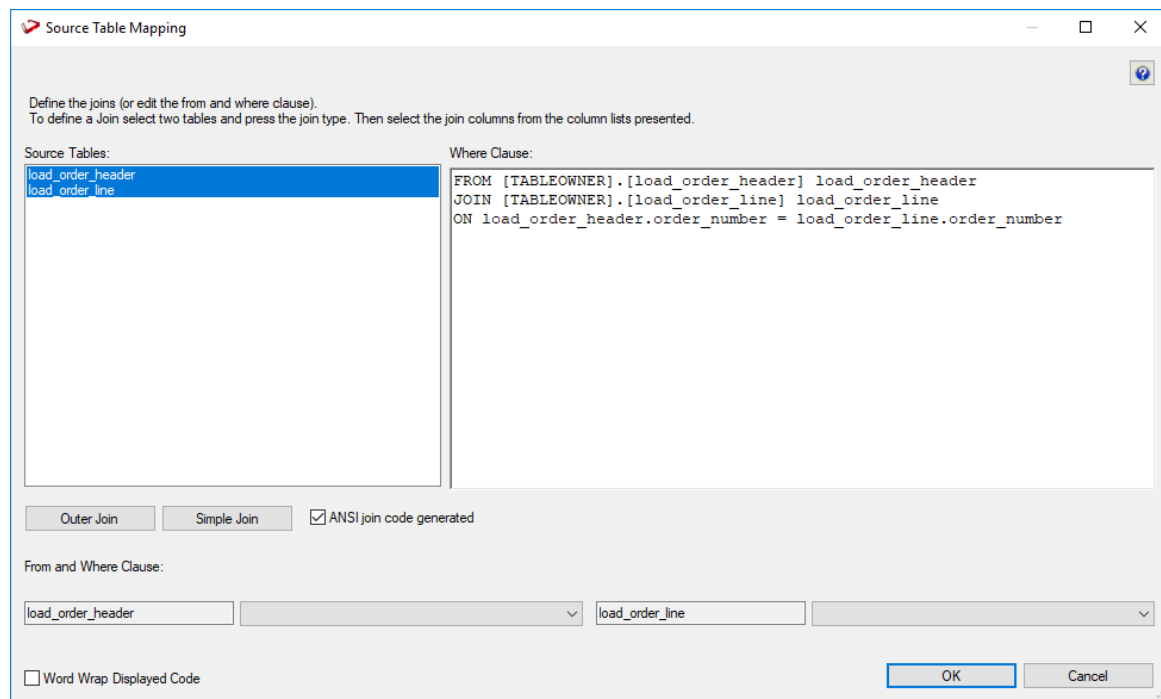
Note: You will need to create **fact** and **dimension** views to use **Analysis Services Cubes**.

- 5 Tick the **Distinct Data Select** check-box if you want the view to return only distinct values.
- 6 Click **OK** and then **Finish**, not creating the view for now.
- 7 Add additional columns from other existing one-to-one views from the right pane into the center pane.

Join the source objects

If columns from more than one table or view have been added, then we have to define the joins between the source objects. This is achieved by right-clicking on the view name and selecting **Build From/Where** clause from the pop up menu.

To join tables, select the tables in the left box and click either the **Outer Join** or **Simple Join** button. Column lists for both tables will appear at the bottom of the dialog box. Select the column (or one of the columns) that allows the two tables to be joined. If an outer join is being used, the column for the master table must be chosen first. If there are multiple columns joining two tables then this action must be repeated for each column. Continue to perform all joins between all tables. The example below only has two tables with one join column so is a relatively simple case. An additional option is available to allow either an ANSI standard join or a 'Where clause' based join. The ANSI standard join should be chosen in most situations. See the example screen in the following section.



Simple Join

A simple join joins the two tables, and only returns rows where data is matched in both tables. So for example if table A has 100 rows and table B has a subset of 24 rows. If all the rows in table B can be joined to table A then 24 rows will be returned. The other 76 rows from table A will not be returned.

Outer Join

An outer join joins the two tables, and returns all rows in the master table regardless of whether or not they are found in the second table. So if the example above was executed with table A as the master table then 100 rows would be returned. 76 of those rows would have null values for the table B columns. In the example screen above the table 'load_order_line' has had its column chosen and the column for the table 'load_order_header' is currently being chosen. This will result in the statement as shown in the 'Where' clause edit window. The results of this select are that a row will be added containing order_line information regardless of whether or not an order_header exists.

As the join columns are selected the 'Where' statement is built up in the large edit window on the right. Once all joins have been made the contents of this window can be changed if the join statement is not correct.

Once satisfied with the 'Where' statement click the **OK** button to proceed to the next step. As indicated in its description this statement is the 'Where' clause that will be applied to the select statement of the cursor to allow the joining of the various source tables. It can of course be edited in the procedure that is generated if not correct.

You have the choice between 'Where' statement joins and ANSI standard joins.

Note: 'Where' joins are not available if using outer joins in Teradata.

The example below shows the result of an ANSI standard join which takes place in the 'From' statement.

Define source table joins for View

Define the Joins (or edit the where clause).
To define a Join select two tables and select the join type. Then select the join columns from the column lists presented.

Source Tables:

- load_order_header
- load_order_line

From and Where Clause:

```
FROM [TABLEOWNER].[view_order_header] view_order_header  
LEFT OUTER JOIN [TABLEOWNER].[view_order_line] view_order_line  
ON view_order_header.order_number = view_order_line.order_number
```

Outer Join Simple Join ANSI join code generated

Select the columns that join the two tables. Select the column from the Master Table first.

[load_order_header] [load_order_line]

Word Wrap Displayed Code

OK Cancel

Create the view

Once the view has been defined in the metadata we need to physically create the view in the database. This is achieved by using the right-clicking on the view name and selecting **Create (ReCreate)** from the pop up menu.

The output from the creation are visible in the output windows. The following example shows a successful creation.

Object	Message
view_sales_analysis	, tax
view_sales_analysis	, dss_update_time)
view_sales_analysis	As Select
view_sales_analysis	view_order_header.model_order_header_key
view_sales_analysis	, view_order_line.model_order_line_key
view_sales_analysis	, view_order_header.model_customer_key
view_sales_analysis	, view_order_line.model_product_key
view_sales_analysis	, view_order_header.order_number
view_sales_analysis	, view_order_header.order_date
view_sales_analysis	, view_order_header.customer_code
view_sales_analysis	, view_order_header.ship_date
view_sales_analysis	, view_order_line.order_line_no
view_sales_analysis	, view_order_line.product_code
view_sales_analysis	, view_order_line.unit_sale_price
view_sales_analysis	, view_order_line.quantity
view_sales_analysis	, view_order_line.sales_value
view_sales_analysis	, view_order_line.tax
view_sales_analysis	, view_order_header.dss_update_time
view_sales_analysis	FROM WslWarehouse.view_order_header view_order_header
view_sales_analysis	LEFT OUTER JOIN WslWarehouse.view_order_line view_order_line
view_sales_analysis	ON view_order_header.order_number = view_order_line.order_number
view_sales_analysis	Table comments added.
view_sales_analysis	5 column comments added.

The contents of this window are a message to the effect that the view was created followed by a copy of the actual database create statement.

If the view was not created then ascertain and fix the problem. A common problem is a 'Duplicate column' where a column has been accidentally added twice. The best way of finding such a column is to double click on the list heading **Col name**. This will sort the column names into alphabetic order. Another double click on the heading will sort the columns back into their create order. Column ordering can be changed by altering the column order value against a column's properties.



TIP: Double clicking on the heading of a column in a list sorts the list into alphabetical order based on the column chosen.

DIMENSION VIEW HIERARCHIES

The various hierarchies associated with a dimension view can be recorded in the WhereScape RED metadata. These hierarchies are often not used in any form, except to provide documentary completeness and for creating **Analysis Services OLAP Cubes**.

ADDING A DIMENSION VIEW HIERARCHY

Any number of hierarchies can be created against a dimension view. There is no restriction on the form of the hierarchy. To add a new hierarchy, position on the dimension view in the left pane and using the right-click menu, select **Hierarchies>Add Hierarchy**. The following dialog appears.

The screenshot shows the 'Add Hierarchy' dialog box. The 'Hierarchy Name' field contains the text 'customer'. The 'Description' field is empty. Below the description, there is a note: 'Move columns from the column list into the hierarchy. The hierarchy is a top down list. For example a date hierarchy may be year, month, day. Year will be the first column shown.' The 'Available Columns' list contains the following items: 'dim_customer_key', 'code', 'name', 'address', 'city', 'state', and 'dss_update_time'. The 'Hierarchy: Levels' list is empty. The dialog has 'OK' and 'Cancel' buttons at the bottom right.

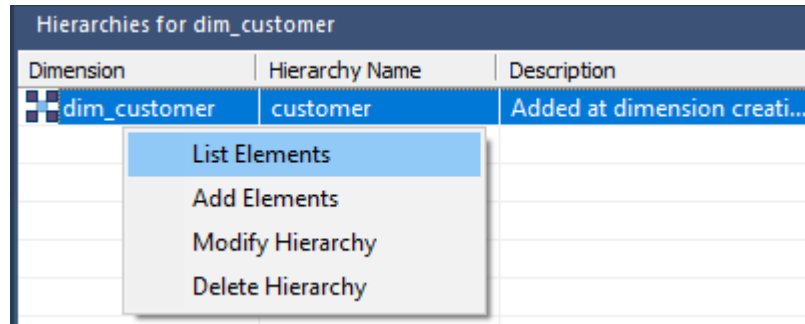
- 1 Enter a meaningful name for the hierarchy.
- 2 Enter a meaningful description for the hierarchy. This description is carried through into the Hierarchy Description field of any OLAP Dimensions that are built from the original Dimension object.

Note: The description text is automatically set to "Added at dimension creation for cube support" but this can be edited to match the user's intended description.

The hierarchy is built with the highest level at the top; for example a customer dimension view may have **state** at the highest level, then **city**, then **address** and finally **code** at the lowest level.

To enter the hierarchy elements, select them in the required order, from the left pane and click > to add them to the right column. Once all the hierarchy elements have been added, click **OK**.

A hierarchy and its elements can be edited by listing the hierarchies associated with a dimension and using the right-click menu options available in the middle pane.



The screenshot shows a table with the following data:

Dimension	Hierarchy Name	Description
dim_customer	customer	Added at dimension creati...

A right-click context menu is open over the first row, with the following options:

- List Elements
- Add Elements
- Modify Hierarchy
- Delete Hierarchy

CREATING A CUSTOM VIEW

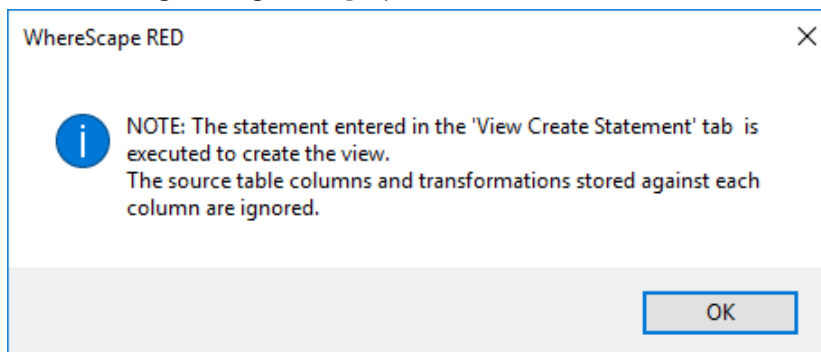
A custom view can be created within RED to handle views that are not strictly one to one, such as where multiple tables are joined or where a complex condition is placed on the view. There are two options for custom views, the first where the columns are defined in RED and the 'Select' component of the view is customized. The second option is where the view is totally custom and no columns need to be defined in RED, although it is good practice to still define the columns for documentation purposes.

To create a Custom or 'User Defined' view proceed as follows:

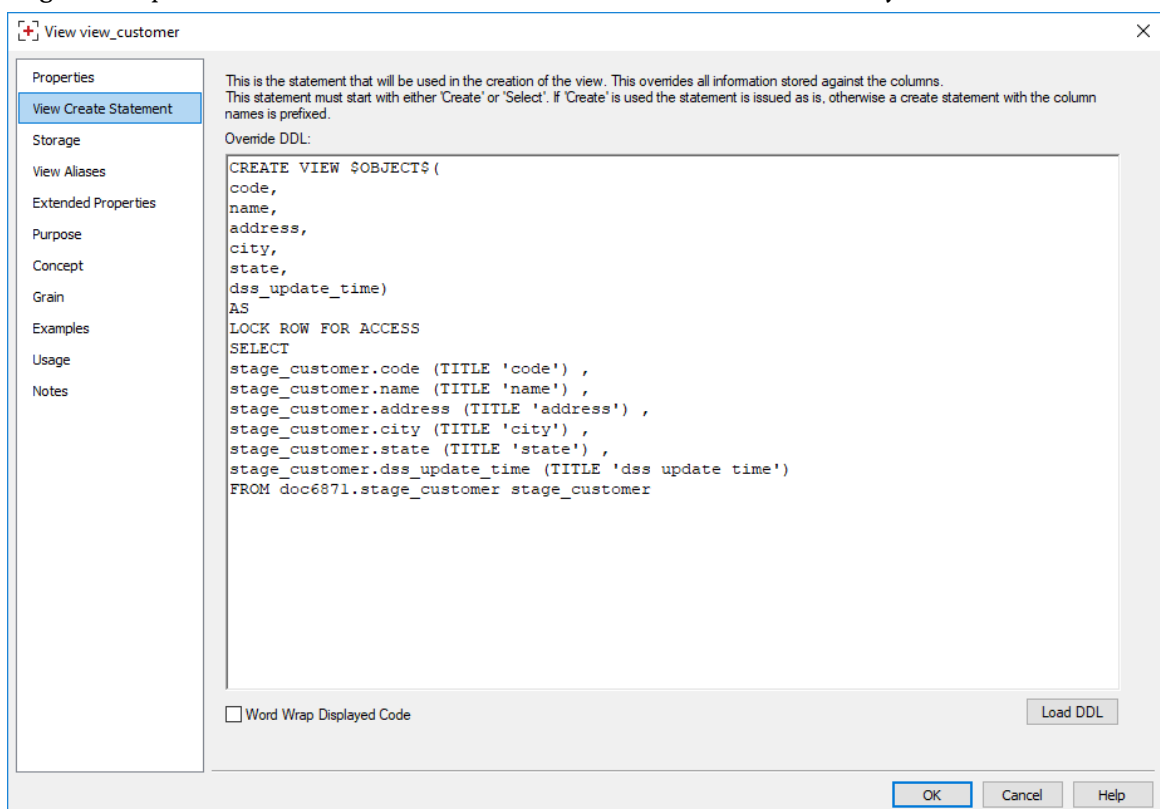
- 1 Create a view in the normal manner either by dragging a table in or adding a new object.
- 2 Change the **View Type** to **User Defined View** in the properties of the view.

The screenshot shows the 'View view_customer' properties dialog. The 'View Type' dropdown is set to 'User Defined View'. The 'View Name' is 'view_customer', 'Unique Short Name' is 'view_customer', and 'Business Display Name (EUL)' is 'view_customer'. The 'Description' field is empty. The 'Update Procedure' and 'Custom Procedure' are both set to '(None)'. The 'Distinct Data Select' checkbox is unchecked. The 'From/Where or Where Clause' field is empty. The 'Table Locking Mode' is set to 'LOCK ROW FOR ACCESS'. The 'Mnemonic (EUL)' field is empty. The 'Timestamps' section shows 'Metadata Structure' as 2017-12-14 23:53:11.020000, 'Database' as 2017-12-14 23:59:08.560000, and 'Database Altered' as 2017-12-14 23:59:08.560000. The 'OK', 'Cancel', and 'Help' buttons are at the bottom right.

- 3 The following message is displayed. Click **OK**.



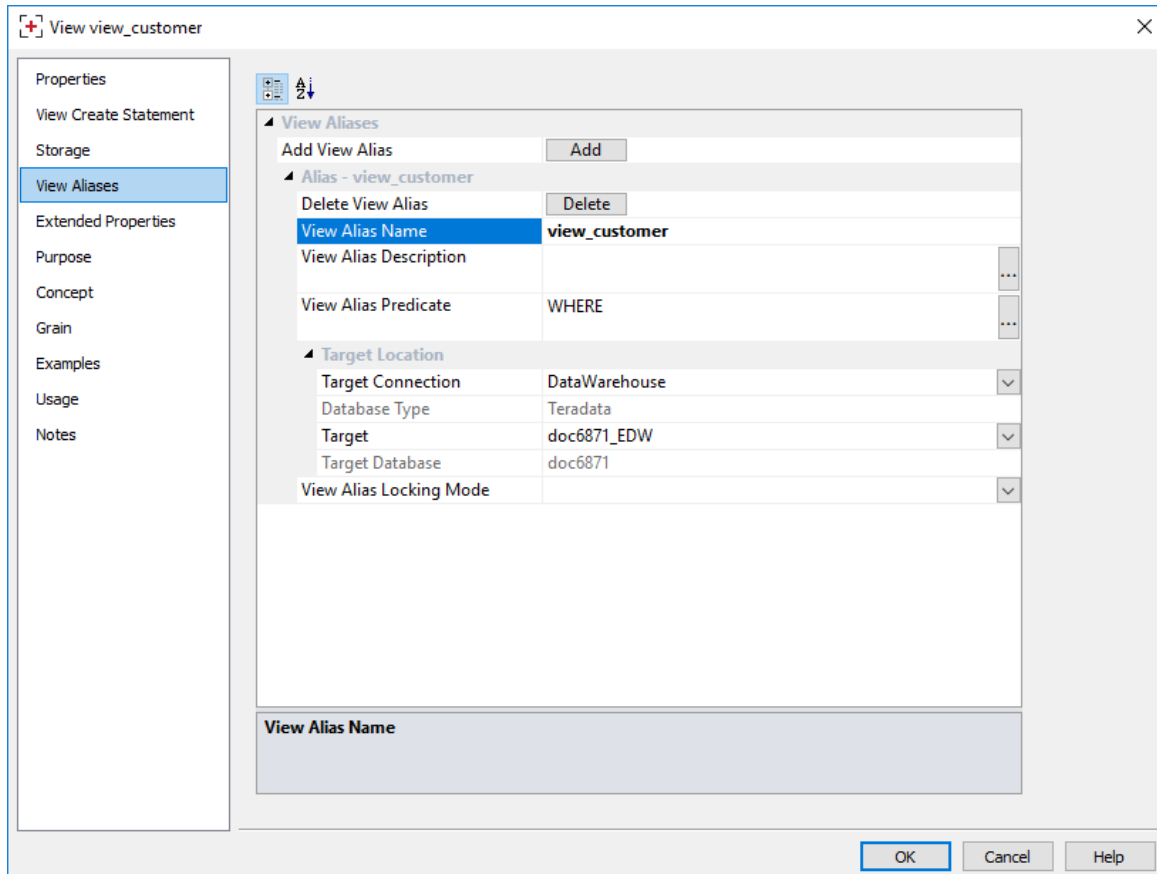
- 4 Edit the new tab **View Create Statement** and insert the SQL Statement that will be used to create the view. This SQL Statement must start with either 'Create' or 'Select'. If 'Create' is used then the columns in the view are ignored and the statement will be issued to create the view. If the statement starts with 'Select' then RED will build up a view create statement from the column names and the supplied **Select** clause. A button 'Load from Columns' is available to get a sample **Select** statement based on the columns in the view and any transformations.



VIEW ALIASES

View Aliases provide multiple deployments of the same view into different Teradata presentation layers.

The **View Aliases** tab enables you to define additional/replica views.



Add button

Enables you to add a View Alias which adds a new copy of the view. This enables you to define new View Alias which is an alternate definition that will ultimately exist as another database view with a different name, predicate and/or target location.

Delete button

Enables you to delete a View Alias and remove its metadata details after closing the dialog.

View Alias Name

The view alias name that is used as the alternate database view name.

View Alias Description

Description of the view alias.

View Alias Predicate

Optional 'Where' clause to include in the alternate view definition.

Target Connection

The name of the WhereScape RED connection that identifies the target location to create the alternate view.

Database Type [Read Only]

The database type of the selected target connection.

Target

The Target that defines the database and schema for the table.

Target Database [Read Only]

Target database identifier (e.g. Oracle SID or TNS Name, Teradata TDPID) or Database Name (e.g. as in DB2 or SQL Server) to reference the alternate view.

View Alias Locking Mode

The locking mode applied to the alternate view.

CHAPTER 21

ANALYSIS SERVICES OLAP CUBES

IN THIS CHAPTER

OLAP Overview	532
OLAP Defining the Data Source for the OLAP Cube.....	532
OLAP Defining an OLAP Cube.....	535
OLAP Inspecting and Modifying Advanced Cube Properties	539
OLAP Creating an OLAP Cube on the Analysis Services Server	540
OLAP Cube Objects.....	541
OLAP Dimension Objects	573
OLAP Changing OLAP Cubes.....	586
OLAP Retrofitting an OLAP Object	588

OLAP OVERVIEW

A **cube** is a set of related measures and dimensions that is used to analyze data.

- A **measure** is a transactional value or measurement that a user may want to aggregate. The source of measures are usually columns in one or more source tables. Measures are grouped into **measure groups**.
- A **dimension** is a group of attributes that represent an area of interest related to the measures in the cube and which are used to analyze the measures in the cube. For example, a customer dimension might include the attributes:
 - Customer Name
 - Customer Gender
 - Customer City

These would enable measures in the cube to be analyzed by Customer Name, Customer Gender, and Customer City. The source of attributes are usually columns in one or more source tables. The attributes within each dimension can be organized into hierarchies to provide paths for analysis.

A cube is then augmented with **calculations**, key performance indicators (generally known as **KPIs**), **actions**, **partitions**, **perspectives**, and **translations**.

The information required to build and support an Analysis Services cube and its surrounding structure is reasonably complex and diverse. In attempting to automate the building of Analysis Services cubes WhereScape RED has simplified and restricted many of the functions available to the cube designer. WhereScape RED includes most of the commonly used capabilities and the components that logically fit into the methodology incorporated within WhereScape RED.

WhereScape RED broadly provides functionality to manage all of the above, except for perspectives and translations. These can be created outside of WhereScape RED, scripted in xmla and executed from within WhereScape RED. Features of cubes that are not supported in WhereScape RED can be added to the cube via the Microsoft tools. These altered cubes can still be processed through the WhereScape RED scheduler, and the cube should be documented within WhereScape RED to explain the post creation phases required.

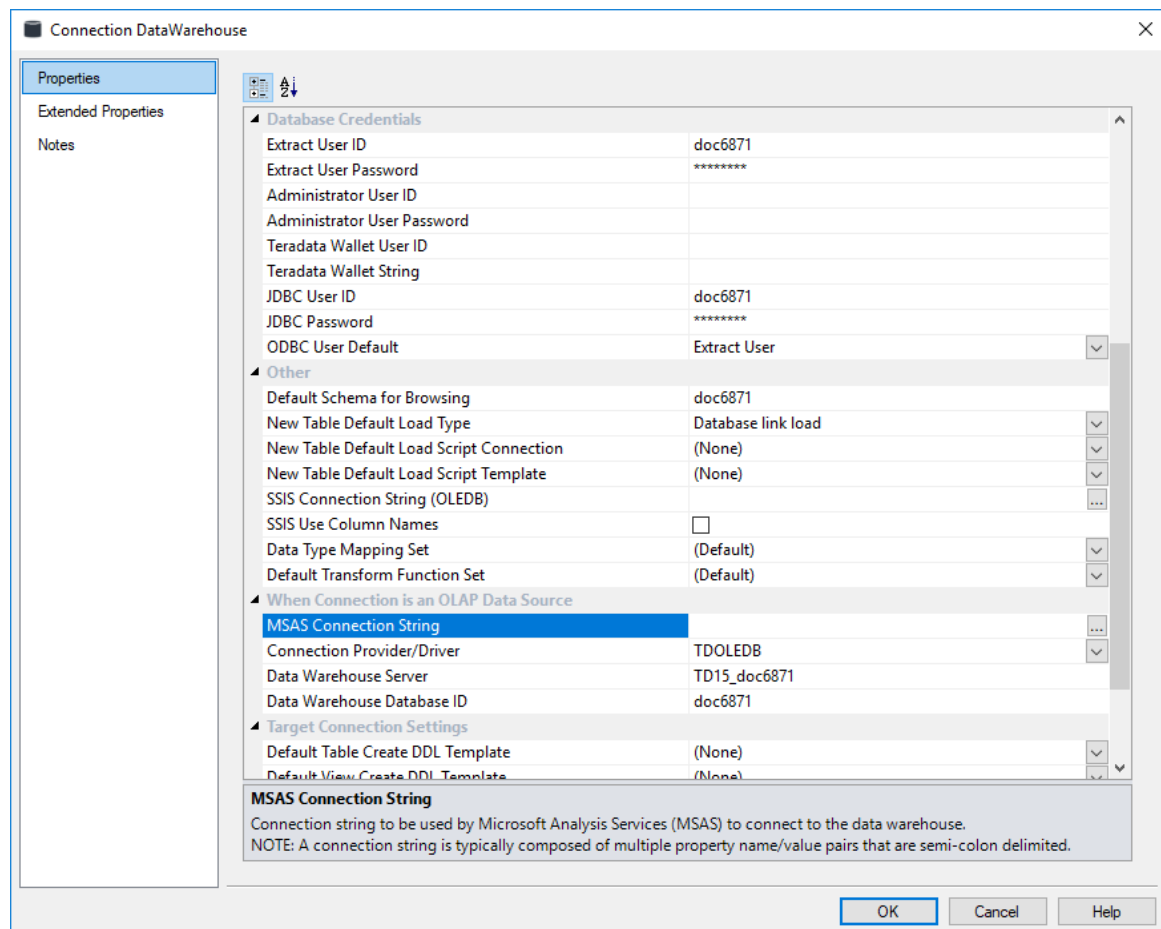
As a general rule, once a cube or a component of a cube is created on the Analysis Services server it cannot be altered through WhereScape RED. The OLAP object can be dropped and recreated easily using RED. New OLAP objects defined in RED (e.g. additional calculations or measures) can be added by recreating the cube.

WhereScape RED supports cubes in Microsoft SQL Server Analysis Services versions 2005 and 2008.

OLAP DEFINING THE DATA SOURCE FOR THE OLAP CUBE

Before we can create an OLAP cube, we first need to set up the data warehouse to be used as a source for Analysis Services cubes.

On the **Datawarehouse** Properties screen, the fields in the section **When Connection is an OLAP Data Source** are required.



MSAS Connection String

Connection string to be used by Microsoft Analysis Services (MSAS) to connect to the data warehouse.

Note: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited.

Connection Provider/Driver

Name of the Connection Provider/Driver to use to connect to the data warehouse database when it is used as the data source for OLAP cubes. Set to **TDOLEDB**.

Data Warehouse Server

Data Warehouse Server Name, which is used when the data warehouse is used as the data source for OLAP cubes. Set this to the Teradata TDPID.

Data Warehouse Database ID

Data Warehouse Database Identifier (e.g. Oracle SID or TNS Name, Teradata TDPID) or Database Name (e.g. as in DB2 or SQL Server), which is used when the data warehouse is used as the data source for OLAP cubes.

OLAP DEFINING AN OLAP CUBE

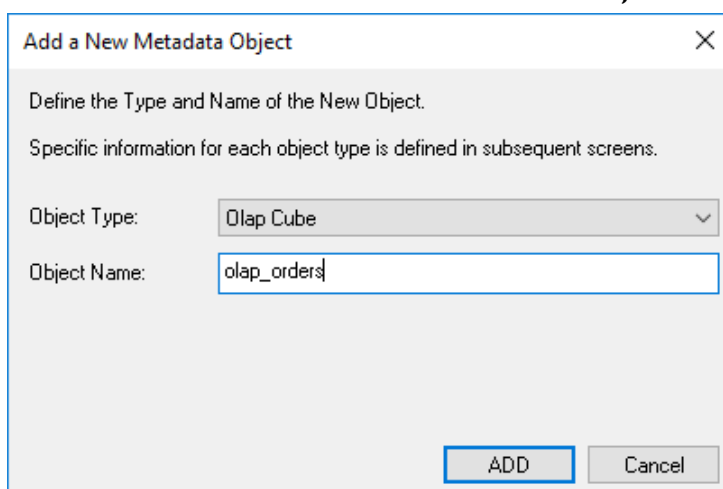
OLAP Cubes can be created from fact views. A single cube can contain data from multiple source star schemas, each defined with a measure group. An OLAP Cube consists of many parts namely, measure groups, measures, calculations, actions, dimensions, dimension hierarchies, dimension attributes and dimension attribute relationships. It is strongly recommended that drag and drop is used to create an OLAP Cube in order that all the components are set up correctly. OLAP Cubes can utilize a hierarchical structure in the dimensions to facilitate user queries. Therefore, each dimension present in an OLAP Cube should have either a hierarchy of levels or attributes and relationships. The hierarchies are defined against the underlying dimensional attributes which can be inherited from the source dimension metadata. Individual attributes can be added to the dimension after the OLAP Cube or OLAP Dimension metadata has been created.

Note: Analysis Services does not like **name** as a column name. For dim_customer it will therefore be necessary to change the column name from **name** to **cname** for example.

Building a New OLAP Cube

To create an OLAP Cube proceed as follows:

- 1 Double click on the **OLAP Cube object group** to make the middle pane a cube drop target.
- 2 Select the data warehouse connection to browse in the source pane. The connection can be selected by right-clicking the Data Warehouse connection in the Object pane and choosing Browse Source System.
- 3 Drag a fact view from the source pane into the target pane.
- 4 Set the cube name in the **Create new metadata Object** dialog box and click **ADD**.



Add a New Metadata Object

Define the Type and Name of the New Object.

Specific information for each object type is defined in subsequent screens.

Object Type: Olap Cube

Object Name: olap_orders

ADD Cancel

- 5 A dialog box will prompt for any OLAP Dimensions that do not already exist that are required for this cube (based on the fact view metadata). Set the dimension name in the **Add a new metadata Object** dialog box and click **ADD**. Repeat this for each dimension as required.

Note: If you wish to include Attribute Relationships in Analysis Services for this dimension, click on the **Include Attribute Relationships** checkbox.

- 6 A dialog appears, prompting you to select the attributes to be included in the *Customer* OLAP dimension.

The attributes available for selection are in the left column. To select an attribute, click on the attribute in the left column and click **>**. This will move the attribute to the right column. To de-select an attribute, click on the attribute in the right column and click **<**. This will move the attribute to the left column.

Repeat Step 5 for each dimension as required.

- 7 A dialog will appear with a list of the fields that are to be added as measures. Remove any columns that are not measures. A measure is a column that uses the sum, count, min or max of the column. Calculations can be chosen if required at this point. A date dimension must be present along with a hierarchy to allow the definition of these calculated members.

Manage Cube Measures

The following measures have been added to the cube. Select any measures that are not appropriate and click the Remove button if required.

Pre-built calculated members will be generated for each measure if the check boxes to the right are set. These calculated members can be deleted later if not required.

Measure	Calculated Members:
<input type="checkbox"/> quantity	<input checked="" type="checkbox"/> Month To Date
<input type="checkbox"/> gross sale amount	<input checked="" type="checkbox"/> Year To Date
<input type="checkbox"/> net sale amount	<input type="checkbox"/> Moving Quarter
<input type="checkbox"/> tax amount	<input checked="" type="checkbox"/> Moving Year
	<input checked="" type="checkbox"/> Same Month Previous Year
	<input checked="" type="checkbox"/> Previous Year To date

Use OLAP Functions

Dimension and level information for calculated members:

Date Dimension:

Date Hierarchy:

Month Level:

Year Level:

- 8 During cube creation the **Adding Cube Measures** dialog is shown. In this dialog the following options are provided:
 - Measure, provides a list of measures that can be aggregated (e.g. using Sum, Count, Min, Max or Distinct Count). By default WhereScape RED will show all attributes in the fact view that are defined as numeric and additive. Those attributes that should not be considered measures can be removed using the **Remove** button.
 - Calculated Member options will allow the user to add some predefined date based calculated member definitions to be built against the cube. The standard calculations:
 - Month to date
 - Year to date
 - Moving Quarter
 - Moving Year

Same Month Previous Year

Previous Year to date

These will define a calculated measures based on the associated drop-down boxes. There are 2 different ways that WhereScape RED will implement these calculations which is dictated by the Use OLAP functions checkbox:

Using OLAP Functions - will implement the calculations using MDX Expressions within the cube using date based MDX functions. These calculations are efficiently executed by Analysis Services.

Without using OLAP functions - will implement the calculations using an MDX Filter function built over date dimension attributes. This option leverages the flags from the relational date dimension and ensures that a query using the calculations in the Cube will match an equivalent query against the star schema and is particularly useful if non-standard date periods are used.

- 9 The cube and dimensions will be created in WhereScape RED metadata and the cube measures will be displayed.

Setting Cube Properties

The properties of the cube must be completed before we can create the cube in the Analysis services database. Most of the elements in the properties screen will be defaulted, but each of the following columns will probably need to be completed.

- 1 The **Connection** to the Analysis services server must be defined within the cube properties. This connection is a connection object. If no such connection exists then a new connection of type must be created and configured. SQL Server 2005 or 2008 Analysis Services use a connection type of "Analysis Server 2005+". This connection name must then be chosen in the cubes properties.
- 2 A **Cube Database Name** must be selected. A new database name can be created by selecting **(Define New Cube Database...)** from the drop-down list. This database name is the database that the cubes will reside in on the Analysis services server.
- 3 The **Data Source Connection** must be defined and the three derived values shown under this connection must be present. If there is nothing in the three fields below the data source connection then the connection object will need to be modified.

OLAP INSPECTING AND MODIFYING ADVANCED CUBE PROPERTIES

Now that the basic OLAP Cube has been defined, various properties of the OLAP Cube can be inspected or modified:

Measure Groups

- 1 Display the Measure Groups by right-clicking on the cube name and selecting **Display Measure Groups**.
- 2 Change the Measure Group properties by right-clicking on the measure group and selecting **Properties**.

Measures

- 1 Display all of the Measures associated with a cube by right-clicking on the cube name and selecting **Display Measures**.
- 2 Change the measure properties by right-clicking on the measure name and selecting **Properties**.

Calculations

- 1 Display all of the Calculated members defined on the cube by right-clicking on the cube name and selecting **Display Calculations**.
- 2 Change the Calculated members by right-clicking on a calculation and selecting **Properties**.

KPIs

- 1 Display all of the KPIs defined on the cube by right-clicking on the cube name and selecting **Display KPIs**.
- 2 Change the KPIs by right-clicking on the KPI name and selecting **Properties**.

Actions

- 1 Display all of the Actions defined on the cube by right-clicking on the cube name and selecting **Display Actions**.
- 2 Actions can be changed by right-clicking on the Action name and selecting **Properties**.

Partitions

- 1 Display all of the Partitions defined on the Measure Groups that are associated with the cube by right-clicking on the cube name and selecting **Display Partitions**.
- 2 Change Partitions by right-clicking on the Partition name and selecting **Properties**.

Dimensions

- 1 Display all of the OLAP Dimensions associated with the cube by right-clicking on the cube name and selecting **Display Dimensions**.
- 2 Change the customizable OLAP Dimension properties by right-clicking on the OLAP Dimension name and selecting **Properties**.

Measure Group Dimensions

- 1 Display the relationship of OLAP Dimensions to Measure Groups defined against the cube by right-clicking on the cube name and selecting **Display Measure Group Dimensions**.
- 2 Change the customizable properties of the relationship of the OLAP Dimension to Measure Group by right-clicking on the OLAP Dimension name and selecting **Properties**.

OLAP CREATING AN OLAP CUBE ON THE ANALYSIS SERVICES SERVER

If all the tasks above are completed, then it should be possible to create the cube on the Analysis Services server. When positioned on the **OLAP Cube** name, right-click and select **Create (Alter) Cube**. WhereScape RED checks that key components of the cube are correct before it proceeds to issue the create command.

The create cube menu option performs the following tasks on the Analysis Services server:

- Create an Analysis Services database if the name specified is not already present.
- Create a Data Source with connection information back to the data warehouse based on the cube source information in the Data Warehouse connection.
- Create a Data Source View to support the cube objects defined
- Create the dimensions used by the cube as database shared dimensions if they do not already exist.
- Create the cube if it does not exist
- Create a partition for the cube.

OLAP CUBE OBJECTS

OLAP CUBE PROPERTIES

The properties associated with a cube are described below. These properties relate both to the cube environment and the cube itself.

There are seven tabs in the cube properties screen.

The first is the main properties, the second the processing and partitioning options and the rest are for documentation stored in the WhereScape RED metadata and displayed in the generated WhereScape RED documentation. In order to see the cube properties, right-click on the cube and select **Properties**.

Internal Cube Name

This is the name by which the cube is known within WhereScape RED. This name is limited to 64 characters in length and may contain only alphanumeric characters and underscores.

Cube Publish Name

This is the name that the cube will have in the Analysis Services server. It is not constrained in its content except by the limitations imposed by Analysis Services.

Cube Description

A description of the cube. This is used both in the WhereScape RED documentation and is also stored against the cube in Analysis Services.

Cube Database Connection

This field allows the selection of one of the existing connections defined within WhereScape RED. The connection must be of type 'Microsoft Analysis Server 2005+'. If no such connection exists, then a new connection object must be created. This connection object is used to point to the Analysis Services server.

Cube Database Name

An Analysis Services server must have databases associated with it. Each database can contain one or more cubes, dimensions, data sources etc. Select the name of an existing database on the server from the drop-down list. To create a new Database name, select '(Define New Cube Database)' from the drop-down list and the dialog that follows will allow you to register the name within the WhereScape RED metadata. Once registered, the name can then be selected as the database.

Data Source Connection

In Analysis Services the data source is the location of the data to be used to build the cube. It also defines the path for any drill through operations. This field provides a drop-down of the existing connections. The Data Warehouse connection must be chosen.

Data Source Provider Type

This field essentially defines what type of database the Data Warehouse is. This field is a read only field in the properties screen. Its value is set in the properties of the data source connection.

Data Source Server

The data source server is also a read only field being sourced from the properties of the data source connection. For SQL Server, it defines the server on which the data warehouse database runs.

Data Source Database

The data source database is also a read only field being sourced from the properties of the data source connection. For SQL Server, it defines the database in which the data warehouse runs.

Post Create XML/A Script

This is an XML/A script that is run on the cube database when the cube is created. This script allows Analysis Services features to be added to the cube or cube database that have been built outside of WhereScape RED—for example security roles that has been defined for the cube can be recreated from the script when the cube is created (or recreated).

Post Update XML/A Script

This is an XML/A script that is run on the cube database when the cube is updated or processed via the scheduler. This script allows Analysis Services features to be added to the cube or cube database that have been built outside of WhereScape RED—for example security roles that has been defined for the cube can be recreated from the script when the cube is updated or processed.

Processing Mode

Gets or sets the index and aggregation settings for cube processing. The value indicates whether processed data is available after all required data aggregation has been completed (Regular) or immediately after the data has been loaded (Lazy Aggregations). This setting will be used as the default for new measure groups and partitions created for the cube.

Processing Priority

Gets or sets the processing priority for the cube.


Partition Processing Mode

This option determines how partitions are updated when a cube is updated.

All Partitions - Sequential will update each cube partition sequentially.

All Partitions - Parallel will have all the cube partitions updated in parallel.


Measure Group Partitions - Parallel will process each measure group sequentially with all partitions of that same measure group being updated in parallel.



Process Cube Dimensions

Will determine whether to process the OLAP Dimensions related to the cube as part of the processing of the cube. The options are to process the **enabled dimensions** only, to process **no dimensions**, or to process **all the dimensions**.

Processing of specific Dimensions with the Cube can be enabled or disabled on the **Process Cube Dimension With Cube** checkbox of each Dimension's Properties screen. See **OLAP Cube Dimensions**.



Process Selected Cube Dimensions in Parallel

Selecting this check-box will allow for all dimensions within the cube to be updated in parallel instead of being updated sequentially.

Storage Mode

This field allows two options; MOLAP - Multidimensional OLAP or ROLAP.- Relational OLAP. At the cube properties level, setting this field will determine the defaults for the Storage Mode field on its related Measure Groups and partitions.

Default Measure

Specifies the measure used to resolve MDX expressions if a measure is not explicitly referenced in the expression. If no default measure is specified an arbitrary measure is used as the default measure.

Estimated Rows

Specifies the estimated number of rows in the fact views. Enter the size of the fact view if known, otherwise leave as zero.

Visible

Indicates whether the cube is visible to client applications.

ROLAP stands for Relational Online Analytical Processing.

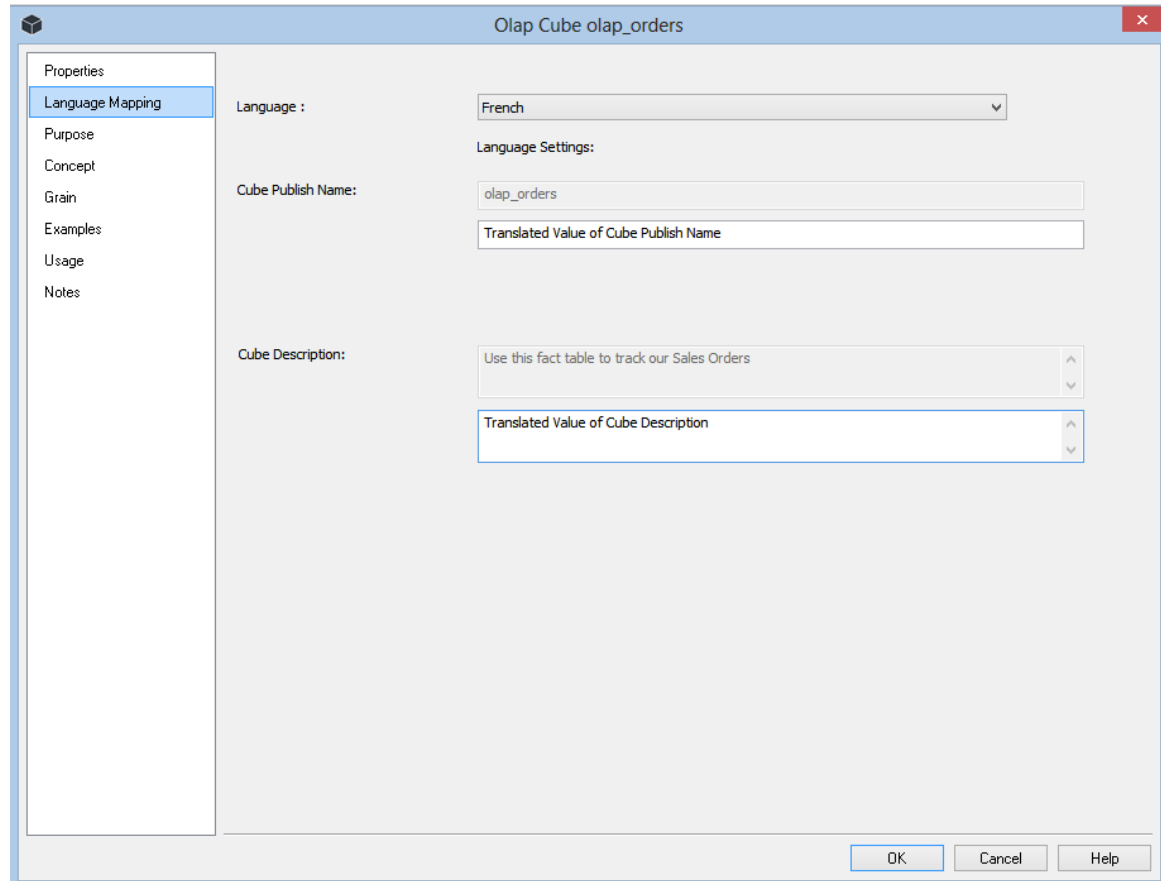
ROLAP is an alternative to the MOLAP (Multidimensional OLAP) technology. While both ROLAP and MOLAP analytic tools are designed to allow analysis of data through the use of a multidimensional data model, ROLAP differs significantly in that it does not require the pre-computation and storage of information. Instead, ROLAP tools access the data in a relational database and generate SQL queries to calculate information at the appropriate level when an end user requests it. With ROLAP, it is possible to create additional database tables (summary tables or aggregations) which summarize the data at any desired combination of dimensions.

While ROLAP uses a relational database source, generally the database must be carefully designed for ROLAP use. A database which was designed for OLTP will not function well as a ROLAP database. Therefore, ROLAP still involves creating an additional copy of the data. However, since it is a database, a variety of technologies can be used to populate the database.

Language Mapping

The OLAP Cube Properties screen has a tab called **Language Mapping**.

Select the language from the drop-down list and then enter the translations for the **Cube Publish Name** and the **Cube Description** in the chosen language.



The screenshot shows a dialog box titled "Olap Cube olap_orders" with a close button (X) in the top right corner. On the left, there is a vertical list of tabs: Properties, Language Mapping (selected), Purpose, Concept, Grain, Examples, Usage, and Notes. The main area contains the following fields:

- Language :** A dropdown menu with "French" selected.
- Language Settings:** A section header.
- Cube Publish Name:** A text box containing "olap_orders".
- Translated Value of Cube Publish Name:** An empty text box.
- Cube Description:** A text box containing "Use this fact table to track our Sales Orders".
- Translated Value of Cube Description:** An empty text box.

At the bottom right, there are three buttons: OK, Cancel, and Help.

See **Language Settings** (see "**Settings - Language Options**" on page 140) for more details on how to add languages for translation.

OLAP CUBE MEASURE GROUPS

A cube can contain multiple measure groups. In WhereScape RED each measure group can belong to a single cube, and each measure group relates to a single star schema. The Measure Groups Cube processing is defined on the Processing a tab of the **Measure Group Properties**.

The Measure Group's Properties are shown by right-clicking on the Measure Group and choosing **Properties**. The Measure Group Properties associated with cubes are described below.

The screenshot shows the 'Measure Group order_detail Definition' dialog box. On the left is a sidebar with tabs: 'Measure Group Properties' (selected), 'Processing / Partitions', 'Language Mapping', 'Purpose', 'Concept', 'Grain', 'Examples', 'Usage', and 'Notes'. The main area contains the following fields:

- Measure Group Name:** Text box containing 'order_detail'. Buttons: '<- Update' and 'Update ->'.
- Measure Group Description:** Text area containing 'Use this fact table to track our Sales Orders'.
- Source Table Type:** Dropdown menu set to 'Fact Table'.
- Source Table:** Dropdown menu set to 'fact_order_detail'.
- Estimated Rows:** Text box containing '0'.
- Storage Mode:** Dropdown menu set to 'MOLAP'.
- Ignore Unrelated Dimension:** Dropdown menu set to 'True'.
- Measure Group Type:** Dropdown menu set to 'Regular'.
- Processing Mode:** Dropdown menu set to 'Regular'.
- Processing Priority:** Text box containing '0'.

At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

Measure Group Name

Specifies the name of the Measure Group in Analysis Services

Measure Group Description

Specifies the Metadata description of the Measure Group. This description is stored in Analysis Services and is also used in the WhereScape RED auto-generated documentation.

Update Buttons

The Update Buttons: **Update** <- and **Update** -> are used to move from the current Measure Group to the previous and next Measure Group respectively for the current OLAP cube. The alternative is to exit the Measure Group and choose the next Measure Group Properties screen.

Source Table Type

Specifies the type of table from which the Measure Group has been built in WhereScape RED.

Source Table

Specifies the table from which the Measure Group has been built and populated in WhereScape RED.

Estimated Rows

Specifies the estimated number of rows in the source table. If it is unknown then leave this value as 0.

Storage Mode

This field allows two options; MOLAP - Multidimensional OLAP or ROLAP.- Relational OLAP. The default value is inherited from the value set at the Cube level. Again, setting this field at the Measure Group level will determine the default for the Storage Mode field on the related Partitions.

Ignore Unrelated Dimensions

Indicates whether dimensions that are unrelated to the Measure Group are forced to their top level when their members are included in a query.

Measure Group Type

Specifies the type of information the Measure Group contains. In some cases this enables specific treatment by the server and client applications.

Processing Mode

Indicates whether processed data is available after all the aggregations have been computed (Regular) or immediately after the data has been loaded (Lazy Aggregations).

Processing Priority

Specifies the priority for processing the measure group.

OLAP CUBE MEASURE GROUP PROCESSING/PARTITIONS

Partitions define separately manageable data slices of the measure group data. Partitions can be created, processed and deleted independently within a Measure Group. Each Measure Group needs at least one partition to be defined to allow the Measure Group to be processed.

Partitions are managed through the **Processing/Partitions tab** of the Measure Group's Properties within WhereScape RED.

The screenshot shows the 'Olap Cube Measure Group olap_orders.sales_detail' dialog box. The 'Processing / Partitions' tab is selected in the left-hand menu. The main area contains the following settings:

- Process Method:** Full process
- Increment Filter:** (without the WHERE keyword)
- Partitioning Method:** Auto create partitions as required (one numeric level)
- Partition By Dimension:** dim_order_date
- Partition By Attribute:** order cal month
- Partition Value Type:** YYYYMM
- Fact Partition Lookup Clause:** (without the WHERE keyword) dss_update_date < GETDATE()-14
- Max # Auto Create Partitions:** 12 (0 = unlimited)
- # Historic Partitions Updated:** 3

Buttons for 'OK', 'Cancel', and 'Help' are located at the bottom right of the dialog.

Process Method

The processing or updating method of a Measure Group is an area that requires careful consideration. The default option is 'Full process' which will result in the Measure Group being rebuilt. This is in many ways the safest option, but processing time may mean other options must be chosen. The valid options are:

The following describes the processing methods that are available in Analysis Services for Measure Groups:

- Default Process**
 Detects the process state of the measure group, and performs processing necessary to deliver a fully processed state.
- Full Process**
 Processes an Analysis Services Measure Group regardless of state. When Process Full is executed against a Measure Group that has already been processed, Analysis Services drops all data, and then processes it.

- **Incremental Process**

Adds newly available fact data and process only to the relevant partitions. In order to use this option you must set the 'incremental filter' field with a statement that will result in only new data being selected. Failure to do so will result in duplicated data in the cube. In many data warehouses transactional data undergoes changes as well as the addition of new data, and so the incremental update option is not possible. A validation regime should be put in place to compare cube data to the transactional data if incremental is used. This validation regime should be used to notify the administrator in the event that duplicate data is inserted into the cube.

- **Update Data**

Processes data only without building aggregations or indexes. If there is data in the partitions, it will be dropped before re-populating the partition with source data. This processing option is supported for dimensions, cubes, measure groups, and partitions.

- **Build Structure**

If the cube is unprocessed, Analysis Services will process, if it is necessary, all the cube's dimensions. After that, Analysis Services will create only cube definitions. The build structure option just builds the cube structure without populating it. This can be useful if you have a very large cube and want to validate the design.

Increment Filter

If an incremental processing option is chosen then a filter statement must be selected to only return those rows that are to be added to the Measure Group. As mentioned above, care must be taken when using incremental updates. For example, if the Measure Group is accidentally processed twice and the filter is based on date, then duplicate data will be inserted into the Measure Group without any warning.

Partitioning

Partitioning is useful for handling large datasets where it is impractical to reprocess the entire Measure Group. In such a case the full process option would probably be chosen but only selected partitions would be processed. See the section on partitioning for more information. The default process will perform a full process on the first pass followed by incremental updates on subsequent processing runs. Care should be taken when choosing default for the cube.

Partitioning Method

Three options are provided for handling Measure Group partitions. They are:

- 1 One partition only**

When this option is selected the partition information for the Measure Group is ignored and one partition is created and processed for the Measure Group. This would be the normal situation unless performance issues require an alternate strategy.

- 2 Manually managed multiple partitions**

With this option the partition information stored for the Measure Group is used in the creation and processing of the Measure Group.

- 3 Automatic partition handling**

This option is available if the Measure Group is to be partitioned by one numeric value. The partitioning should preferably be on something like day, month or year. (i.e. YYYY, YYYYMM or YYYYDDD). If this option is chosen together with one of the date formats described above,

then WhereScape RED will automatically create partitions as required and process only those partitions that are marked for processing.

Partition by Dimension

This field is only available if automatic partition handling is chosen. Select the dimension in the Measure Group that we will partition by. This would normally be a date dimension.

Partition by Attribute

This field is only available if automatic partition handling is chosen. Select the attribute that we are to partition by. This would normally be a year or maybe a month level. (e.g. cal_year, fin_year from the WhereScape date dimension).

Partition by Value Type

This field is only available if automatic partition handling is chosen. Select the type of level we are dealing with. Choose YYYY for a year partition and YYYYMM for a month partition. This format must correspond with the column in the date dimension. WhereScape RED only supports partitioning by Year, Quarter, month or day.

Fact Partition Lookup Clause

This field is only available if automatic partition handling is chosen. In order to know when to create a new partition WhereScape RED executes a query against the fact view and the date dimension to acquire each unique period. When dealing with a large fact view, such a query may take a long time to complete. This field can be used to include the components of the 'Where' clause to restrict the amount of data examined. For example we may enter 'dss_update_time < GETDATE()-14' to only look at fact view records that have been inserted or updated in the last 14 days. This should still allow us to catch any new partitions and add them. The first time a cube is converted to auto partitioning handling, a full pass of the fact view should occur to allow inclusion of every partition. This field should therefore only be populated once the cube has been initially built with all partitions intact.

Max Number of Auto Created Partitions

This field is only available if automatic partition handling is chosen. You can specify an upper limit for automatically created partitions. The default is zero, or no limit. This limit may be useful if your source system can get erroneous data. If set, then the processing of the Measure Group will fail if a new partition will exceed the counter.

Number of Historic Partitions Updated

This field is only available if automatic partition handling is chosen. This field allows you to restrict the partition updating to the latest nnn partitions. If for example, we were partitioning by year and we set this value to 2, we would process the current and previous years only.

WhereScape RED turns off partition processing after it does a partition update, so the first pass will still update all partitions.

To Display Measure Groups

To display a list of measure groups defined against a cube, right-click on a cube and select **Display Measure Groups**.

To Add a Measure Group

To add a measure group, display the measure groups in the middle pane and either:

- Drag over a new fact view into the target pane - this will automatically create a new measure group in the cube. Any additional dimensions required to support analysis of the Measure Group will be added to the cube.
- Right-click on the cube in the object pane and select Add Measure Group and fill in the Measure Group properties.

To Delete a Measure Group

To delete a measure group, display the measure groups in the middle pane and right-click, select delete measure group.

Displaying Measures

Measures can be displayed or added while viewing measure groups in the middle pane. Right-click on a measure group and select the appropriate option.

Displaying Partitions

Partitions can be displayed or added while viewing measure groups in the middle pane. Right-click on a measure group and select the appropriate option.

OLAP CUBE MEASURE GROUP PARTITIONS

The Measure Group Partition's properties are shown by right-clicking on the Measure Group Partition and selecting **Properties**. The partition's properties associated with a measure group are described below.

Measure Group Partition Definition

Partition Properties

Cube Name:

Data Source:

Fact Table:

Partition Name:

Partition Description:

Data Slice Formular:
[cube dimension].[level].[slice]
 Example: [ship_date].[cal year].[2003]

Aggregation Prefix:

Filter Statement:

Storage Mode:

Partition Type:

Remote Server:

Processing Method:

Cube Name

The name of the cube that the partition belongs to.

Data Source

The data source for the partitions. This will be inherited from the cube and cannot be changed. You cannot have partitions with different data sources or different fact tables in WhereScape RED. If you need to support either scenario then the partition must be created directly within Analysis Services. In such a case it can still be managed in terms of processing through WhereScape RED.

Fact Table

The fact table from which the data is derived. This is inherited from the cube and cannot be changed. See the notes above under data source.

Update Buttons

The Update Buttons: **Update** <- and **Update** -> are used to move from the current Measure Group Partition to the previous and next Measure Group Partition respectively for the current Measure Group. The alternative is to exit the Measure Group Partition and choose the next Measure Group Partition Properties screen.

Partition Name

Where only one partition exists it is normally given the same name as the cube. If manually creating then a unique name must be assigned for each partition. If auto partitioning is chosen then WhereScape RED will use the cube name plus the level value to make the partition name.

Partition Description

A description of the partition for documentation purposes.

Data Slice Formula

This field defines the range of data stored in the partition. It is a very simplified version of what can be done in Analysis Services. If a more complex partitioning algorithm is required then the partition will need to be created in Analysis Services. The format for the formula is as follows:

The brackets must surround each name and a full stop must separate the three parts of the formula. For example a cube that is partitioned by year on its date dimension would have the following formula for the 2003 year. [dim_date].[cal_year].[2003]

Aggregation Prefix

By default any cube aggregation will be prefixed with the partition name. An alternate name can be entered here. See Analysis Services for more details.

Filter Statement

Not implemented.

Storage Mode

This field allows two options; MOLAP - Multidimensional OLAP or ROLAP.- Relational OLAP. This determines how the OLAP cube is processed. The default value is inherited from the value set at the Measure Group level.

Partition Type

The partition can be either Local or Remote. Local means that the partition resides on the same Analysis Services server as the cube. If Remote is chosen, then a server must be specified where the partition will be located.

Remote Server

If a Remote partition is chosen, then the name of the remote Analysis Services server must be entered here.

Processing Method

A partition is either enabled for processing or disabled. This field can be set to Always process or Never process. If left as default and in automatic mode then WhereScape RED will disable the processing once the partition has been aged out.

OLAP CUBE MEASURES

Measures represent the numeric attributes of a fact table that are aggregated using an OLAP aggregate function defined against each Measure. Each Measure is defined against a Measure Group, which is defined against a cube. The properties of a measure are shown by right-clicking on a Measure and choosing **Properties**. In more detail:

The screenshot shows the 'Cube Measure Definition' dialog box. On the left, there is a sidebar with 'Measure Properties' selected and 'Language Mapping' below it. The main area contains the following fields:

- Cube Name: olap_orders
- Measure Name: quantity
- Measure Group: order_detail
- Source Table: fact_order_detail
- Source Column: quantity
- Data Type: int
- Aggregation Method: Sum
- Display Format: #, #
- Null Processing: Automatic
- Order Number: 160
- Visible: True
- Display Folder: (empty)
- Description: Quantity of product sold (i.e. number of product units)

Buttons for '<- Update', 'Update ->', 'OK', 'Cancel', and 'Help' are visible.

Cube Name

A Read Only field that indicates against which OLAP Cube the measure is defined.

Measure Name

Specifies the Analysis Services name defined for the measure.

Measure Group

Specifies against which Measure Group the Measure is defined. This is related to the fact view of which the Measure is an attribute.

Update Buttons

The Update Buttons: **Update** <- and **Update** -> are used to move from the current Measure to the previous and next Measure respectively for the current OLAP Cube Measure Group. The alternative is to exit the Measure and choose the next Measure Properties screen.

Fact Table

A read only field that indicates the fact view that is related to the Measure Group.

Source Column

Specifies from which numeric attribute of the underlying fact view the Measure is built. This is a drop-down list populated from REDs metadata definition of the fact view associated with the Measure Group above.

Data Type

Specifies the data type used by Analysis Services. The data type specified for this property is inherited from the fact view attribute defined in WhereScape RED metadata but can be different (typically a larger data type is used in the cube to cope with the larger numbers generated by aggregating the source data).

Aggregation Method

Specifies the OLAP function used to aggregate measure values. The default options are:

- Sum
- Count
- Distinct Count - only one distinct count is allowed per Measure Group and can have query performance implications.
- Min - Minimum
- Max - Maximum

Display Format

Specifies the format used by clients when displaying the measure value.

Null Processing

Specifies the processing null values. Setting this property to Automatic means that Analysis Services uses default behavior.

Order Number

The order in which the measures appear in the cube is dictated by their order number.

Visible

Measures are normally visible, but some measures used in calculations may be hidden. In such a case clear this checkbox.

Display Folder

Cube Measures can be organized into user-defined folders to view and manage these attributes within the Analysis Services user interface more easily. Enter the display folder name.

Note: One object can be in multiple display folders, for example:

Display Folder: ▼

Description

A description of the measure which is stored in the cube. This description will by default be acquired from the source column.

To View measures

The measures can be viewed by clicking on an OLAP Cube in the left pane which will display the measures in the middle pane.

To Add a New Measure

To add a new measure, view the measures in the middle pane, right-click in the middle pane and select add measure. This can also be done when viewing measure groups in the middle pane. Alternatively to create measure which is very similar to an existing measure, view the measures in the middle pane and right-click, select **Copy Measure**. The same dialog box appears as for **Add Measure** with most of the fields filled in. Notice that the measure name has the suffix "- Copy". Change the name in the Measure Name field and make any other alterations and click **OK**.

To Delete a measure

Display the measures in the middle pane, select the measure to delete, right-click and select **Delete**.

OLAP CUBE CALCULATIONS

Calculations provide the ability to define the derivation of a value at query time within a cube. The calculation is most typically a numeric derivation based on measures, but can be defined against any dimension. The calculation is defined in MDX (Multi-Dimensional eXpressions). The definition of a Calculation is shown by right-clicking a **Calculation** and choosing **Properties**. The following Cube Calculated Member Definition dialog is shown:

Cube Name

A Read Only field that indicates against which OLAP Cube the Calculation is defined.

Calculated Member Name

Specifies the Analysis Services name defined for the Calculation.

Update Buttons

The Update Buttons: **Update <-** and **Update ->** are used to move from the current Calculation to the previous and next Calculation respectively for the current OLAP Cube. The alternative is to exit the Calculation and choose the next Calculation Properties screen.

Description

A business description of the Calculation that is used to populate WhereScape RED documentation.

Expression

Specifies the MDX expression that defines the calculation.

Parent Hierarchy

Specifies where the calculation is displayed for use. By default and in most cases, this will be 'Measures'. This means that the calculated member will be displayed to the end user of the cube as a measure, otherwise known as a calculated measure. Alternatively, you can include the calculated member in a dimension instead of in the measures. Hierarchies are descriptive categories of a dimension by which the measures in a cube can be separated for analysis. A calculated member provides a new label in the parent dimension you select.

Parent Member

This is not available if you select **Measures** as your parent hierarchy, or if you select a one-level hierarchy. Hierarchies are divided into levels that contain members. Each member produces a heading in the cube. While browsing a cube, users can drill down to subordinate headings. The heading for the calculated member will be added directly below the selected Parent Member.

Associated Measure Group

Specifies against which Measure Group the Measure is defined. This is related to the fact view of which the Measure is an attribute.

Note: SSAS 2008+ uses the property **Associated Measure Group** but previous versions of SSAS do not and can result in errors when creating the cube. It is possible to set this attribute in RED to **(Undefined)** for previous versions of RED, but this is not necessary for the current version of RED as this attribute will only be used when appropriate.

Display Folder

Cube Calculations can be organized into user-defined folders to view and manage these attributes within the Analysis Services user interface more easily. Enter the display folder name.

Note: One object can be in multiple display folders, for example:



Display Folder:

Display Format

Specifies the format used by clients when displaying the measure value.

Visible

Specifies whether the calculation is visible to client tools.

Non Empty Behavior

Determines the non-empty behavior associated with the calculation

Order Number


The create order of the member in the dimension hierarchy.

To View Calculations

To view the list of calculations (sometimes called a calculated measure), right-click on an OLAP Cube in the left pane and select **Display Calculations**.

To Add a Calculation

- 1 To add a calculation, right-click on an OLAP Cube in the left pane, select **Add Calculation** and then choose between:
 - Add Calculated Member
 - Add Named Set
 - Add Script



This can also be achieved by displaying calculations in the middle pane, right-clicking and selecting **Add Calculated Member**, **Add Named Set** or **Add Script**. Fill out the dialog box with the relevant details.

- 2 To create a calculation that is similar to an existing calculation, display the calculations in the middle pane and select **Copy Calculation**. The same dialog box appears as for **Add Calculation** with most of the fields filled in. Notice that the calculated member name has the suffix "- Copy". Change the name in the Measure Name field and make any other alterations and click **OK**.

To Delete a calculation

Display the calculations in the middle pane, select the calculation to delete, right-click and select **Delete**.

OLAP CUBE KEY PERFORMANCE INDICATORS

In Analysis Services, a KPI is a collection of calculations that are associated with a measure group in a cube that are used to evaluate business success. Typically, these calculations are a combination of Multidimensional Expressions (MDX) expressions or calculated members. KPIs also have additional metadata that provides information about how client applications should display the results of the KPI's calculations. The definition of a Calculation is shown by right-clicking a **Calculation** and choosing **Properties**. The following Cube KPI Definition dialog is shown below:

Cube Name

A Read Only field that indicates against which OLAP Cube the KPI is defined.

KPI Name

Specifies the name of the KPI defined in Analysis Services.

Description

A business description of the KPI that is used to populate WhereScape RED documentation.

Update Buttons

The Update Buttons: **Update** <- and **Update** -> are used to move from the current KPI to the previous and next KPI respectively for the current OLAP Cube. The alternative is to exit the KPI and choose the next KPI Properties screen.

Display Folder

KPIs can be organized into user-defined folders to view and manage these attributes within the Analysis Services user interface more easily. Enter the display folder name.

Note: One object can be in multiple display folders, for example:

Display Folder:	<input type="text" value="financials;sales"/>
-----------------	---

Associated Measure Group

Specifies the Measure Group against which the KPI is defined.

Value Expression

An MDX numeric expression that returns the actual value of the KPI.

Goal Expression

An MDX numeric expression or a calculation that returns the target value of the KPI.

Status Expression

An MDX expression that represents the state of the KPI at a specified point in time.

The status MDX expression should return a normalized value between -1 and 1. Values equal to or less than -1 will be interpreted as "bad" or "low." A value of zero (0) is interpreted as "acceptable" or "medium." Values equal to or greater than 1 will be interpreted as "good" or "high."

An unlimited number of intermediate values can optionally be returned and can be used to display any number of additional states, if supported by the client application.

Status Indicator

A visual element that provides a quick indication of the status for a KPI. The display of the element is determined by the value of the MDX expression that evaluates status.

Trend Indicator

A visual element that provides a quick indication of the trend for a KPI. The display of the element is determined by the value of the MDX expression that evaluates trend.

Trend Expression

An MDX expression that evaluates the value of the KPI over time. The trend can be any time-based criterion that is useful in a specific business context.

The trend MDX expression enables a business user to determine whether the KPI is improving over time or degrading over time.

Display Folder

The folder in which the KPI will appear when a user is browsing the cube.

Parent KPI

A reference to an existing KPI that uses the value of the child KPI as part of computation of the parent KPI. Sometimes, a single KPI will be a computation that consists of the values for other KPIs. This property facilitates the correct display of the child KPIs underneath the parent KPI in client applications.

Current Time Member

An MDX expression that returns the member that identifies the temporal context of the KPI.

Weight

An MDX numeric expression that assigns a relative importance to a KPI. If the KPI is assigned to a parent KPI, the weight is used to proportionally adjust the results of the child KPI value when calculating the value of the parent KPI.

To View KPIs

To view the list of KPIs, right-click on an OLAP Cube in the left pane and select **Display KPIs**.

To Add a KPI

- 1 To add a calculation, right-click on an OLAP Cube in the left pane and select **Add KPI**. This can also be achieved by displaying KPIs in the middle pane, right-clicking and selecting **Add KPI**. Fill out the dialog box with the relevant details.
- 2 To create a KPI that is similar to an existing KPI, display the KPIs in the middle pane and select **Copy KPI**. The same dialog box appears as for **Add KPI** with most of the fields filled in. Notice that the KPI name has the suffix "- Copy". Change the name in the KPI Name field and make any other alterations and click OK.

To Delete a KPI

Display the KPIs in the middle pane, select the KPI to delete, right-click and select **Delete**.

OLAP CUBE ACTIONS

An action provides information to a client application to allow an action to occur based on the property of a clicked dimensional member. Actions can be of different types and they have to be created accordingly. To view the definition of an Action, right-click on the **Action** and select **Properties**. The following Cube Action Definition will be shown:

Cube Name

A Read Only field that indicates against which OLAP Cube the Calculation is defined.

Action Type

Actions can be of the following types:

- **Drill through actions** which return the set of rows that represents the underlying data of the selected cells of the cube where the action occurs. When this option is chosen an additional tab is enabled that allows drill through columns to be chosen.
- **Reporting actions** which return a report from Reporting Services that is associated with the selected section of the cube where the action occurs.
- **Standard actions** (Action), which return the action element (URL, HTML, DataSet, RowSet, and other elements) that is associated with the selected section of the cube where the action occurs.

The action type chosen determines which action specific fields are enabled within the dialog.

Action Name

Defines the name of the Action in the Cube.

Update Buttons

The Update Buttons: **Update** <- and **Update** -> are used to move from the current Action to the previous and next Action respectively for the current OLAP Cube. The alternative is to exit the Action and choose the next Action Properties screen.

Target Type

Select the object to which the action is attached. Generally, in client applications, the action is displayed when end users select the target object; however, the client application determines which end-user operation displays actions. For Target type, select from the following objects:

- Attribute members
- Cells
- Cube
- Dimension members
- Hierarchy
- Hierarchy members
- Level
- Level members

Target Object:

The cube object of the designated target type against which the action is defined.

Condition

Specify an optional Multidimensional Expressions (MDX) expression that resolves to a Boolean value. If the value is True, the action is performed on the specified target. If the value is False, the action is not performed.

Report Action: Server

The name of the computer running report server.

Report Action: Report Path

The path exposed by report server.

Format/Content Type

Select the type of action. The following table summarizes the available types.

- Data Set - Retrieves a dataset
- Proprietary - Performs an operation by using an interface other than those listed in this table.
- Row Set - Retrieves a rowset.
- Statement - Runs an OLE DB command.
- URL - Displays a page in an Internet Browser.

Expression

Specifies the parameters that are passed when the action is run. The syntax must evaluate to a string, and you must include an expression written in MDX. MDX expressions are evaluated before the parameters are passed.

Default

An additional true/false drop list is enabled for Drill through actions. This gets or sets the current DrillThroughAction as the default action when multiple drill through actions are defined. This is important for users of Excell 2007 to browse the Olap Cube because Excell will only invoke the Drill through Action marked as the default.

Invocation

Specifies how the action is run. Interactive, the default, specifies that the action is run when a user accesses an object. The possible settings are:

- Batch
- Interactive
- On Open

Application

Describes the application of the action.

Description

Describes the action.

Caption

Provides a caption that is displayed for the action.

Caption is MDX

If the caption is MDX, specify True, if not specify False.

To View Actions

To view the list of Actions, right-click on an OLAP Cube in the left pane and select **Display Actions**.

To Add an Action

- 1 To add an Action, right-click on an OLAP Cube in the left pane and select **Add Action**. This can also be achieved by displaying Actions in the middle pane, right-clicking and selecting **Add Action**. Fill out the dialog box with the relevant details.
- 2 To create an Action that is similar to an existing Action, display the Actions in the middle pane and select **Copy Action**. The same dialog box appears as for **Add Action** with most of the fields filled in. Notice that the Action name has the suffix "- Copy". Change the name in the Action Name field and make any other alterations and click **OK**.

To Delete an Action

Display the Actions in the middle pane, select the Action to delete, right-click and select **Delete**.

OLAP CUBE DIMENSIONS

OLAP Dimensions are associated automatically with a cube when a cube is created in WhereScape RED based on the underlying star schema. OLAP Dimensions that are associated with a cube can be displayed, or additional OLAP Dimensions can be manually added from the list of OLAP Dimensions defined in WhereScape RED.

Once an OLAP Dimension is associated with a cube a relationship is created with the relevant Measure Groups within the cube - these relationships are defined automatically with WhereScape RED, and they can also be added. The properties of an OLAP Dimension associated with a cube are shown by right-clicking the cube Dimensions listed in the middle pane and selecting **Properties** from the right-click menu. The Properties are shown below:

The screenshot shows the 'Cube Dimension dim_customer Definition' dialog box. It has a title bar with a cube icon and a close button. The main area is divided into two panes: 'Cube Dimension Properties' (selected) and 'Language Mapping'. The 'Cube Dimension Properties' pane contains the following fields:

- Internal Dimension Name: odim_customer (dropdown)
- OLAP Dimension Name: dim_customer (text)
- Cube Dimension Name: dim_customer (text)
- Dimension Description: This is the master Customer List for the Company (text area)
- Order Number: 10 (text)
- Visible: True (dropdown)
- All Member Aggregation Usage: Default (dropdown)
- Hierarchy Unique Name Style: IncludeDimensionName (dropdown)
- Member Unique Name Style: Native (dropdown)
- Source Table Type: Dimension (dropdown)
- Source Table: dim_customer (dropdown)
- Source Table Key: dim_customer_key (dropdown)
- Processing Mode: Regular (dropdown)
- All Caption: All customer (text)
- OLAP Dimension Type: Regular (dropdown)
- Unknown Member Action: None (dropdown)
- Unknown Member Name: Unknown (text)

Buttons include '<- Update', 'Update ->', 'OK', 'Cancel', and 'Help'. A checkbox 'Process Cube Dimension With Cube' is checked.

Internal Dimension Name

A read only field displaying the name of the OLAP Dimension in WhereScape RED.

OLAP Dimension Name

Specifies the name of the OLAP Dimension as a Dimension in Analysis Services.

Cube Dimension Name

Specifies the exposed name of the Dimension when associated with a cube (this can be different from the OLAP Dimension Name).

Update Buttons

The Update Buttons: **Update** <- and **Update** -> are used to move from the current OLAP Dimension to the previous and next OLAP Dimension respectively for the current OLAP Cube. The alternative is to exit the OLAP Dimension and choose the next OLAP Dimension Properties screen.

Dimension Description

A description of the dimension when associated with the cube.

Order Number

The order number.

Process Cube Dimension With Cube

Using this checkbox you can enable the dimension to be processed with the OLAP Cube.

Visible

Determines whether or not the dimension is visible to client applications.

All Member Aggregation Usage

Specifies how aggregations will be designed by the BIDS Storage Design Wizard if it is used to design cube aggregations.

Hierarchy Unique Name Style

Indicates whether the dimension name will be included in the name of the hierarchies. If set to Default then the system will apply a default behavior and will include the dimension name in the case where there is more than one usage of the same dimension.

Member Unique Name Style

Indicates how member unique names will be formed.

Other Read only fields that are displayed in this dialog are configurable against the OLAP Dimensions' properties and cannot be changed in this dialog, including:

- Source Table Type
- Source Table
- Source Table Key
- Processing Mode
- All caption
- OLAP Dimension Type
- Unknown Member Action
- Unknown Member Name

To View Cube Dimensions

To view the list of Dimensions associated with a cube, right-click on an OLAP Cube in the left pane and select **Display Dimensions**.

To Add a Cube Dimension

To add an existing OLAP Dimension, right-click on an OLAP Cube in the left pane and select **Add Dimension**. This can also be achieved by displaying Dimensions in the middle pane, right-clicking and selecting **Add Dimension**. Fill out the dialog box with the relevant details.

To Remove a Cube Dimension

Display the Cube Dimensions in the middle pane, select the Dimension to remove, right-click and select **Remove Dimension from Cube**. This action removes the association of the OLAP Dimension from the OLAP Cube.

OLAP CUBE MEASURE GROUP DIMENSIONS

Measure group dimensions are the relationships between cube Measure Groups and OLAP Dimensions. In WhereScape RED this equates to the relationships between fact views and dimensions in the underlying star schema.

The **Properties** are shown below:

The screenshot shows the 'Measure Group Dimension dim_customer Definition' dialog box. The 'Internal Dimension Name' is 'odim_customer', 'OLAP Dimension Name' is 'dim_customer', and 'Cube Dimension Name' is 'dim_customer'. The 'Dimension Description' is 'This is the master Customer List for the Company'. The 'Order Number' is '10' and 'Process Cube Dimension With Cube' is checked. Other settings include 'Visible: True', 'All Member Aggregation Usage: Default', 'Hierarchy Unique Name Style: IncludeDimensionName', 'Member Unique Name Style: Native', 'Measure Group: order_detail', 'Measure Group Column: dim_customer_key', 'Relationship Type: Regular', 'Cardinality: Many', 'Source Table Type: Dimension', 'Source Table: dim_customer', 'Source Table Key: dim_customer_key', 'Processing Mode: Regular', 'All Caption: All customer', 'OLAP Dimension Type: Regular', 'Unknown Member Action: None', and 'Unknown Member Name: Unknown'. Buttons for '<- Update', 'Update ->', 'OK', 'Cancel', and 'Help' are visible.

Internal Dimension Name

A read only field displaying the name of the OLAP Dimension in WhereScape RED.

OLAP Dimension Name

A ready only field displaying the name of the OLAP Dimension as a Dimension in Analysis Services.

Cube Dimension Name

A read only field displaying the name of the Dimension when associated with a cube (this can be different from the OLAP Dimension Name).

Update Buttons

The Update Buttons: **Update** <- and **Update** -> are used to move from the current OLAP Measure Group Dimension to the previous and next OLAP Measure Group Dimension respectively for the current OLAP Cube. The alternative is to exit the OLAP Measure Group Dimension and choose the next OLAP Measure Group Dimension Properties screen.

Dimension Description

A description of the dimension when it is associated with the cube.

Order Number

The order number.

Update Dimension with Cube

A read only checkbox showing whether the dimension is processed when the cube is processed.

Visible

A read only field displaying the visibility of the dimension on the cube.

All Member Aggregation Usage

A read only field displaying how aggregations will be designed by the BIDS Storage Design Wizard if it is used to design cube aggregations.

Hierarchy Unique Name Style

A read only field which indicates whether the dimension name will be included in the name of the hierarchies.

Member Unique Name Style

A read only field which indicates how member unique names will be formed. A read only field.

Measure Group

A read only field displaying the Measure Group which is being referenced by the Measure Group Dimension relationship.

Measure Group Column

Specifies which fact view key joins to the dimension key.

Relationship Type

Defines the relationship type for the relationship between the Dimension and Measure Group. In WhereScape RED this option can be Regular, which means that the relationship is based on a dimension key join, or No Relationship between the Measure Group and Dimension.

Cardinality

Indicates whether the measure group has a many to one or one to one relationship with the dimension.

Source Table Type

A read only field that displays the type of table from which the OLAP Dimension was created.

Source Table

A read only field that displays the name of the table from which the OLAP Dimension was created.

Source Table Key

A read only field that displays the key (typically the primary key) that relates the dimension to the fact in the underlying star schema.

Processing Mode

A read only field that indicates whether processed data is available after all aggregations have been computed or immediately after the data has been loaded.

All Caption

A read only field that displays the name of the (All) member. This applies to all hierarchies in the dimension that have an (All) member.

OLAP Dimension Type

A read only field that displays the type of information contained by the dimension.

Unknown Member Action

A read only field that displays the existence of an Unknown member and whether that member is visible or hidden.

Unknown Member Name

A read only field that displays the caption for the unknown member.

OLAP DIMENSION OBJECTS

OLAP DIMENSION OVERVIEW

OLAP Dimensions are dimensions that get created in an Analysis Services database.

An OLAP Dimension is a collection of related attributes which can be used to provide information about fact data in one or more cubes. By default attributes are visible as attribute hierarchies and can be used to understand the fact data in a cube. Attributes can be organized into user-defined hierarchies that provide navigational paths to assist users when browsing the data in a cube.

They are typically created and populated from a relational dimension.

One or more OLAP Dimensions are defined automatically by WhereScape RED when a fact view is dragged over to create a cube or measure group. WhereScape RED will take the relational dimension tables and related metadata (including hierarchies) defined in the star schemas and create OLAP Dimensions automatically. They can also be defined manually in WhereScape RED.

The properties of an OLAP Cube dimension are shown by right-clicking on the **OLAP Dimension** and choosing **Properties**. The following dialog is shown:

The screenshot shows the 'Olap Dimension odim_customer' dialog box with the following properties:

- Internal Dimension Name:** odim_customer
- Dimension Publish Name:** dim_customer
- Dimension Description:** (Empty text area)
- Default Database Connection: (Analysis Services):** SSAS Cubes
- OLAP Database Name:** OrderAnalysis
- Data Source Connection: (Data Warehouse):** DataWarehouse
- Data Source Provider Type:** TDOLADB
- Data Source Server:** TD_14_00
- Data Source Database:** dssdemo
- Post Create XML/A Script:** (None) [Edit]
- Source Table Type:** Dimension
- Source Table:** dim_customer
- Source Table Key:** dim_customer_key
- Processing Group:** ByAttribute
- Processing Mode:** Regular
- Processing Method:** Default process
- Storage Mode:** MOLAP
- All Caption:** All customer
- OLAP Dimension Type:** Regular
- Unknown Member Action:** None
- Unknown Member Name:** Unknown

Buttons at the bottom: OK, Cancel, Help.

Internal Dimension Name

Specifies the name of the dimension in WhereScape RED.

Dimension Publish Name

Specifies the name of the dimension as created in Analysis Services.

Dimension Description

A business description of the OLAP Dimension for use in documentation - this description also gets created in the analysis services metadata.

Default Database Connection and OLAP Database Name

The WhereScape RED connection that is an OLAP connection to an analysis services server. These fields only need to be populated when the OLAP Dimension needs to be created in Analysis Services separately from a cube. If these fields are blank this dimension can only be created in the same Analysis Services server and database as the related cubes when the cubes get created.

Data Source Connection

Defines the WhereScape RED connection that points to the relational dimensional table(s) used to populate the OLAP Dimension - typically the Data Warehouse connection. When the connection is defined the following read only fields are populated with the connection information:

- Data Source Provider Type
- Data Source Server
- Data Source Database

Source Table Type

Specifies the type of table from which the OLAP Dimension was created.

Source Table

Specifies the name of the table from which the OLAP Dimension was created.

Source Table Key

Specifies the key (typically the primary key) that relates the dimension to the fact in the underlying star schema.

Processing Group

Specifies the processing group for processing the dimension. This determines how much data is read into memory during dimension processing at any one time.

Processing Mode

Indicates whether processed data is available after all aggregations have been computed (Regular) or immediately after the data has been loaded (Lazy aggregations).

Processing Method

Indicates which processing method should be used for populating the dimension:

- **Process Default** - Detects the process state of an object, and performs processing necessary to deliver unprocessed or partially processed objects to a fully processed state.
- **Process Full** - Processes an Analysis Services object and all the objects that it contains. When Process Full is executed against an object that has already been processed, Analysis Services drops all data in the object, and then processes the object. This kind of processing is required when a structural change has been made.
- **Rebuild Data** - Processes data only without building aggregations or indexes.

Storage Mode

This field allows two options; MOLAP - Multidimensional OLAP or ROLAP.- Relational OLAP. This determines how the OLAP dimension is processed.

All Caption

Specifies the name of the (All) member. This applies to all hierarchies in the dimension that have an (All) member.

OLAP Dimension Type

Specifies the type of information contained by the dimension. Some client tools can treat the dimension differently based on this information.

Unknown member Action

Specifies the existence of an Unknown member and whether that member is visible or hidden. Fact data not associated with a member can be associated with the unknown member.

Unknown Member Name

Specifies the caption for the unknown member.

Language Mapping

The OLAP Dimension Properties screen has a tab called **Language Mapping**.

Select the language from the drop-down list and then enter the translations for the **Dimension Publish Name**, **All Caption**, **Dimension Description** and the **Unknown Member Name**.

Dimension odim_customer Definition

Dimension Properties

- Language Mapping
- Purpose
- Concept
- Grain
- Examples
- Usage
- Notes

Language : French

Language Settings:

Dimension Publish Name: dim_customer
Translated Value of Dimension Publish Name

All Caption: All customer
Translated Value of All Caption

Dimension Description: This is the master Customer List for the Company
Translated Value of Customer Dimension

Unknown Member Name: Unknown
Translated Value of Unknown Member Name

OK Cancel Help

See **Language Settings** (see "**Settings - Language Options**" on page 140) for more details on how to add languages for translation.

OLAP DIMENSION ATTRIBUTES

Dimensional attributes contain information about the Dimension object. Attributes are exposed in the cube to provide the ability to navigate and aggregate the data in the cube.

User defined hierarchies can be built over attributes to provide drill paths through the data and to aid aggregation.

The properties of an attribute can be displayed by right-clicking an attribute in the middle pane and choosing **Properties**. The following dialog is displayed:

The screenshot shows a dialog box titled "Olap Dimension Attribute odim_order_date.dim_order_date_key". The dialog is divided into a left sidebar and a main configuration area. The sidebar has "Properties" selected, with "Language Mapping" below it. The main area contains the following fields and controls:

- Dimension Name:** odim_order_date
- Internal Attribute Name:** dim_order_date_key
- Published Name:** order_date
- Description:** Key for dim_order_date
- Estimated Count:** 1
- Member Names Unique:** False
- Hierarchy Visible:** True
- Hierarchy Enabled:** True
- Hierarchy Optimized State:** FullyOptimized
- Hierarchy Display Folder:** (empty)
- Order By:** Key
- Order By Attribute:** (empty)
- Type:** Regular
- Usage:** Key
- Key Column:** dim_order_date (left), dim_order_date_key (right)
- Name Column:** dim_order_date (left), order_date (right)
- Value Column:** dim_order_date (left), order_date (right)

Buttons at the top right include "<- Update" and "Update ->". Buttons at the bottom right include "OK", "Cancel", and "Help".

Dimension Name

A read only field to display the dimension which the attribute is related to.

Internal Attribute Name

The name of the attribute in WhereScape RED.

Published Name

The name of the attribute created in Analysis Services.

Description

A business name that is stored in WhereScape RED for documentation and stored in the Analysis Services metadata.

Estimated Count

Specifies the number of members in the attribute. This number is either the amount last counted by Analysis Services or a user provided estimate of the member count.

Member Names Unique

Indicates whether member names are unique for this attribute.

Hierarchy Visible

Indicates whether the attribute hierarchy is visible to client applications. Even if the attribute hierarchy is not visible it can still be used in a user defined hierarchy.

Hierarchy Enabled

Indicates whether an attribute hierarchy is enabled for this attribute. If the attribute hierarchy is not enabled, then the attribute cannot be used in a user defined hierarchy.

Hierarchy Optimized state

Specifies the level of optimization applied to the attribute hierarchy.

Order by

Specifies the method used to order the members of the attribute.

Order by attribute

Specifies the attribute used to order the members of the attribute hierarchy

If the 'Order by' property is set to 'AttributeKey' or 'AttributeName' then 'Order by attribute' cannot be empty. It must be populated with values from attribute relationships.

Order By:	Key
Order By Attribute:	
Type :	order_current_cal_year order_current_cal_ytd order_moving_cal_year Regular
Usage:	
Key Column:	dim_order_date
Name Column:	dim_order_date
Value Column:	dim_order_date

Type

Specifies the type of information contained by the attribute.

Usage

Specifies the usage of the attribute.

Key Column

Specifies the details of the binding to the column containing the member key.

Name Column

Specifies the details of the binding to the column containing the member name.

Value Column

Specifies the details of the binding to the column containing the member value.

Using the Value Column OLAP cube attribute setting for Excel date filtering

In the relevant OLAP Date dimension ensure the OLAP Dimension Type property is set to "Time", then for the Key Attribute of the OLAP Date Dimension (e.g. dim_date_key) set the Value Column property to a date data type column (e.g. calendar_date). Usually it will be useful to set the Name Value property for the Key Attribute to a column containing a textual date format (e.g. dates presented in dd/mm/yy format). After publishing and processing the OLAP cube use Microsoft Office Excel PivotTables to expose date-specific filtering options for this dimension's hierarchies instead of label filtering options.

To View Attributes

To view the list of Attributes, right-click on an OLAP Dimension in the left pane and select **Display Attributes..**

To Add an Attribute

- 1 Display the attributes of an OLAP Dimension in the middle pane. Display the columns of the dimension in the right pane then drag over a column from the underlying relational dimension into the middle pane.
- 2 To add an Attribute, right-click on an OLAP Dimension in the left pane and select **Add Attribute**. This can also be achieved by displaying Attributes in the middle pane, right-clicking and selecting **Add Attribute**. Fill out the dialog box with the relevant details.
- 3 To create an Attribute that is similar to an existing Attribute, display the Attributes in the middle pane and select **Copy Attribute**. The same dialog box appears as for **Add Attribute** with most of the fields filled in. Notice that the Attribute name has the suffix "- Copy". Change the name in the Attribute Name field and make any other alterations and click **OK**.

To Delete an Attribute

Display the Attributes in the middle pane, select the Attribute to delete, right-click and select **Delete**.

OLAP DIMENSION ATTRIBUTE RELATIONSHIPS

Attribute relationships define functional dependencies between attributes. If attribute A has a related attribute B in a hierarchical relationship, then an attribute relationship can be defined that assists Analysis Services to aggregate data.

Note: Attribute relationships require unique members in order to function correctly—if there are duplicate values participating in an attribute relationship then cube data can be incorrectly aggregated and displayed.

To view the **Properties** of an attribute relationship right-click on an attribute relationship and select **Properties** to show the dialog below:

The screenshot shows a dialog box titled "Attribute Relationship Properties". It contains the following fields and values:

- Dimension Name: odim_customer
- Attribute: customer_name
- Relationship Name: Customer Name
- Related Attribute: customer_name
- Relationship Type: Rigid
- Cardinality: One
- Visible: True

Buttons for "OK" and "Cancel" are located at the bottom right of the dialog.

Dimension Name

A read only field indicating the dimension associated with the attribute relationship.

Attribute

Specifies the attribute on which the attribute relationship is based.

Relationship Name

Specifies the name of the attribute relationship.

Related Attribute

Specifies the name of the related attribute.

Relationship Type

Indicates whether the relationship between attributes can change over time.

Cardinality

Indicates if the related attribute has a many to one or a one to one relationship with this attribute.

Visible

Indicates whether the attribute relationship is visible to the client applications.

To View Attribute Relationships

To view the list of Attribute Relationships, right-click on an OLAP Dimension in the left pane and select **Display Attribute Relationships**.

To Add an Attribute Relationship

- 1 Attribute relationships are defined automatically when a User Defined hierarchy is inherited from an underlying relational dimension.
- 2 To add an Attribute Relationship, right-click on an OLAP Dimension in the left pane and select **Add Attribute Relationship**. Fill out the dialog box with the relevant details.

To Delete an Attribute Relationship

Display the Attribute Relationships in the middle pane, select the Attribute Relationship to delete, right-click and select **Delete**.

OLAP DIMENSION HIERARCHIES

User defined hierarchies define hierarchical relationships between related dimensional attributes (e.g. Geographical or time based attributes). These related attributes are defined as levels within the hierarchy.

To view the properties of a user defined Hierarchy right-click on the **Hierarchy** and select **Properties**.

The screenshot shows a dialog box titled "Dimension odim_customer Hierarchy Definition". On the left, there is a sidebar with two options: "Hierarchy Properties" (which is selected) and "Language Mapping". The main area of the dialog contains the following fields and controls:

- Dimension Name:** A text box containing "odim_customer".
- Internal Hierarchy Name:** A text box containing "customer".
- Hierarchy Publish Name:** A text box containing "customer".
- Hierarchy Description:** A text area containing "Added at dimension creation for cube support".
- All Member Name:** An empty text box.
- Allow Duplicate Names :** A dropdown menu set to "True".
- Member Keys Unique:** A dropdown menu set to "NotUnique".
- Member Names Unique:** A dropdown menu set to "False".
- Display Folder:** A dropdown menu.

At the bottom right of the dialog, there are three buttons: "OK", "Cancel", and "Help".

Dimension Name

A read only field indicating the dimension associated with the attribute relationship.

Internal Hierarchy Name

The name of the hierarchy within WhereScape RED.

Hierarchy Publish Name

The name of the hierarchy as created within Analysis Services.

Hierarchy Description

A business name that is stored in WhereScape RED for documentation and stored in the Analysis Services metadata.

All Member Name

Specifies the name of the member in the All level

Allow Duplicate Names

Indicates whether the members under a common parent can have the same name.

Member Keys Unique

Indicates whether member keys are unique for this hierarchy.

Note: If you are using a version of Analysis Services earlier than service pack 2, the only value allowed for **Member Keys Unique** is 'NotUnique'. If the value 'Unique' is used, Analysis Services will return an error and the cube will not be created.

Member Names Unique

Indicates whether member names are unique for this hierarchy.

To view User Defined Hierarchies

To view the list of User Defined Hierarchies, right-click on an OLAP Dimension in the left pane and select **Display Hierarchies**.

To Add a User Defined Hierarchy

To add an Attribute Relationship, right-click on an OLAP Dimension in the left pane and select **Add Hierarchy**. Fill out the dialog box with the relevant details.

To Delete a User Defined Hierarchy

Display the User Defined Hierarchies in the middle pane, select the User Defined Hierarchy to delete, right-click and select **Delete**.

OLAP DIMENSION USER DEFINED HIERARCHY LEVELS

The levels specify the drill path over a set of related attributes. The classic hierarchy levels are Year, Month, Date in a Calendar based hierarchy in the date dimension.

To view the Properties of a user defined Hierarchy Level right-click on a user defined Hierarchy Level and select **Properties**.

Dimension Name

A read only field indicating the dimension associated with the attribute relationship.

Hierarchy

The User Defined Hierarchy that contains the level.

Level Number

The number of levels from the top most hierarchy level. These level numbers must start at 1 for the top level and provide continuous numbering to the bottom level.

Internal Level Name

The name of the Level within WhereScape RED.

Level Publish Name

The name of the Level as created within Analysis Services.

Level Description

A business name that is stored in WhereScape RED for documentation and stored in the Analysis Services metadata.

Source Attribute

Specifies the source attribute on which the level is based.

Hide If

Specifies which members are hidden. This property supports ragged hierarchies contain logical gaps between members.

To View User Defined Hierarchy Levels

To view the list of User Defined Hierarchy Levels, right-click on an OLAP Dimension in the left pane and select **Display Hierarchy Levels**.

To Add a User Defined Hierarchy

- 1** To add a User Defined Hierarchy level, right-click on an OLAP Dimension in the left pane and select **Add Hierarchy Level**. Fill out the dialog box with the relevant details.
- 2** Alternatively right-click on a User Defined Hierarchy in the middle pane and select **Add Hierarchy Level**. Fill out the dialog box with the relevant details.

To Delete a User Defined Hierarchy Level

Display the User Defined Hierarchy Levels in the middle pane, select the User Defined Hierarchy Level to delete, right-click and select **Delete**.

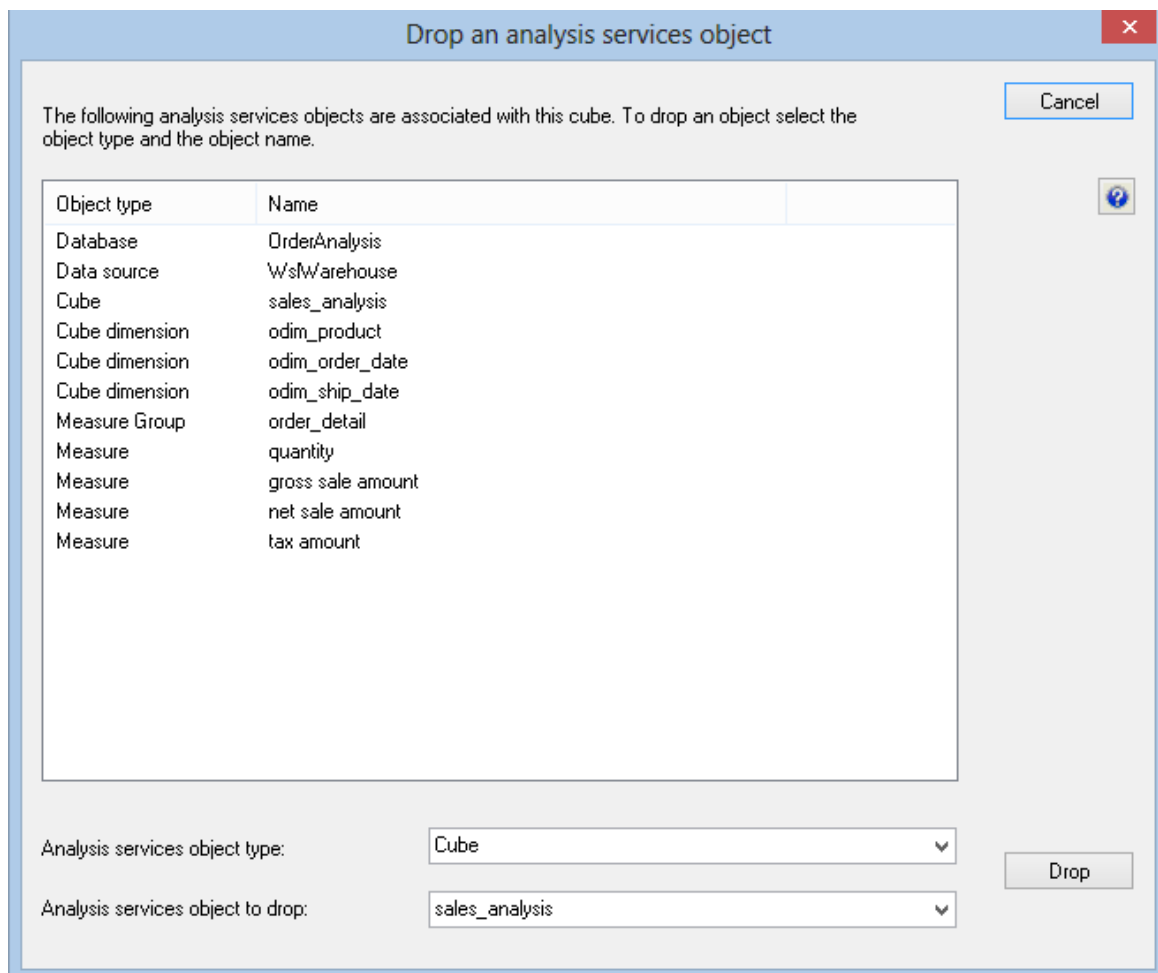
OLAP CHANGING OLAP CUBES

An understanding of the dependency of objects within Analysis Services is the key to figuring out what needs to be dropped or recreated in a cube database using WhereScape RED.

Changes to the underlying relational star schema can cause cube processing to fail as the star schema is frozen in the Data Source View (DSV) of the cube database. Minor changes such as the addition of table columns, or altered data types (e.g. changing a char to varchar) will not break the cube, but renaming a source table or column that is used as a source to the cube will invalidate the DSV and cause processing to fail.

The solution to this issue is to drop and recreate the Cube database from RED to recreate the DSV or manually update the DSV using Microsoft BIDS.

If an object needs to be dropped and recreated in RED then this is two separate actions. For example to drop the OLAP database, right-click an OLAP cube within that database in RED and select **Drop Analysis Services Object**, then using the drop-down boxes in the Drop Analysis Services Object dialog choose the object to drop, and click **Drop**. This will drop the object from Analysis Services.



A Create action on an Analysis Services object in RED will be different depending on whether or not the object already exists in Analysis Services:

- If the object does not already exist in Analysis Services then RED will create the object (and any related objects e.g. OLAP database and DSV).
- If the object does already exist in Analysis Services then RED will try to detect any changes or additional features that need to be added to the object and add or alter the existing Analysis Services object.

Some objects need to be dropped and recreated in order to be changed (eg dimension structures), and some only need to be recreated (eg calculations).

Changes to cube runtime objects do not require the cube database to be dropped. For example a new or changed definition of a calculation or KPI will not require the cube to be dropped and recreated (so data is retained). By Recreating the cube the definition of these runtime objects will be updated and available immediately to cube users.

A brief summary of the hierarchy of objects and the remedial action is shown below:

Cube Object	Change	Action
Data Source	This changes the source database connection. It is defined in the Data Warehouse connection in RED.	OLAP database needs to be dropped and recreated.
Data Source View (underlying relational star)	The DSV reflects the design of the relational star. Therefore, the DSV would need to be updated if any changes are made to tables or views that are used to build OLAP objects.	Changes to the underlying relational star that affect an existing OLAP Object requires that the OLAP Database is dropped and recreated to regenerate the DSV.
OLAP Dimension	The addition or deletion of attributes or hierarchies to an existing OLAP dimension.	The OLAP dimension plus any OLAP cubes associated with the dimension need to be dropped and recreated.
OLAP Cube Measure Group	Delete or Add a Measure Group based on a fact that already exists in the DSV.	Recreate the cube in RED and reprocess.
OLAP Cube Measure Group	Add a Measure Group based on a fact that does not exist in the DSV.	Recreate the OLAP cube database and reprocess.
OLAP Cube Measures	Delete or Add measures based on columns that already exist in the DSV.	Recreate the cube in RED and reprocess.
OLAP Cube Measures	Add measures that are based on new columns that do not exist in the DSV.	Recreate the OLAP cube database and reprocess.
OLAP Cube Calculations, KPIs, Actions	Add, change or delete definition on the cube.	Recreate the cube in RED (a reprocess is not necessary because just the calculation definition is stored in the cube - the result is calculated at query time).

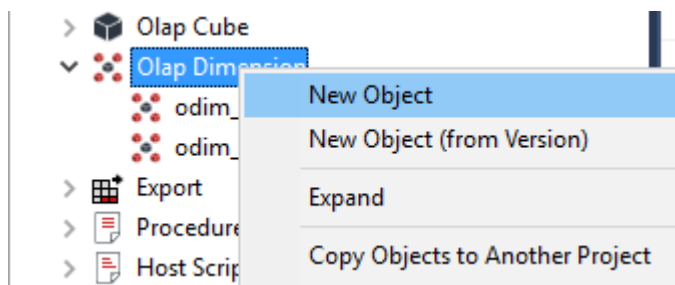
OLAP RETROFITTING AN OLAP OBJECT

WhereScape RED provides the functionality for retrofitting OLAP cubes from Analysis Services.

Note: Before you can retrofit an OLAP cube, you must first retrofit any and all of the OLAP dimensions used by the OLAP cube.

The process to retrofit an OLAP dimension is as follows:

- 1 Right-click on the **OLAP Dimension** object heading in the left pane and select **New Object**.



- 2 Enter any name for the object name and click **ADD**.

 A screenshot of the 'Add a New Metadata Object' dialog box. The dialog has a title bar with a close button. Below the title bar, there is a text area with the instruction: 'Define the Type and Name of the New Object. Specific information for each object type is defined in subsequent screens.' There are two input fields: 'Object Type:' with a dropdown menu showing 'Olap Dimension', and 'Object Name:' with a text box containing 'odim_customer'. At the bottom left, there is a checkbox labeled 'Include Attribute Relationships' which is unchecked. At the bottom right, there are two buttons: 'ADD' (highlighted in blue) and 'Cancel'.

- 3 On the Properties dialog:
 - For the **Internal dimension name** enter the name that matches exactly the name of the OLAP Dimension in Analysis Services that you want to retrofit.
 - For the **Dimension publish name** enter a name for the dimension.

For the **Dimension description** enter a description for the dimension.

For the **Default database connection** select the required Analysis Services connection.

For the **OLAP database name** select the database name in Analysis Services.


For the **Data source connection** select the relevant data source connection.

Olap Dimension odim_customer

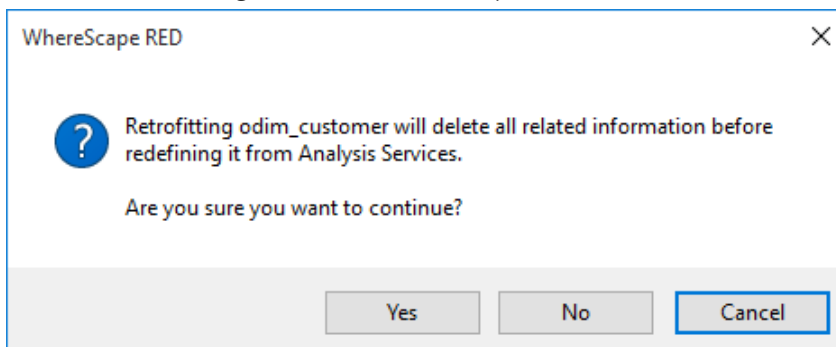
Internal Dimension Name:
 Dimension Publish Name:
 Dimension Description:

Default Database Connection: (Analysis Services)
 OLAP Database Name:
 Data Source Connection: (Data Warehouse)
 Data Source Provider Type:
 Data Source Server:
 Data Source Database:
 Post Create XML/A Script:
 Source Table Type:
 Source Table: Source Table Key:
 Processing Group:
 Processing Mode:
 Processing Method:
 Storage Mode:
 All Caption:
 OLAP Dimension Type:
 Unknown Member Action: Unknown Member Name:

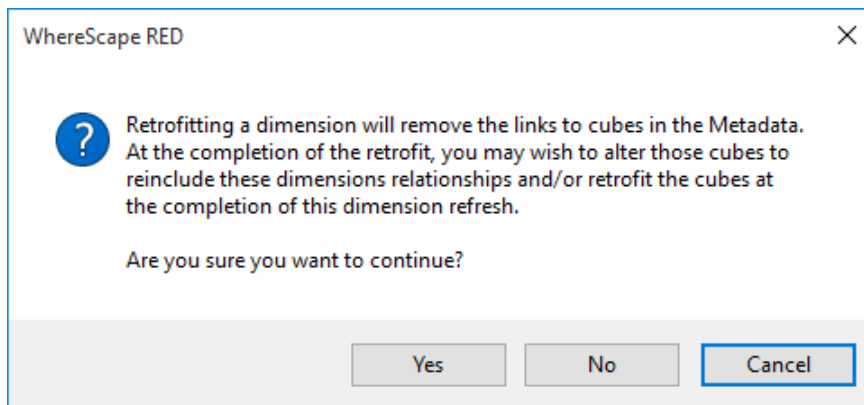
- The OLAP Dimension can now be retrofitted by right-clicking the object in the left pane and choosing **Retrofit Dimension**.



- Two warning dialogs now appear. The first dialog warns that the existing information will be deleted before being redefined from Analysis Services. Select **Yes**.



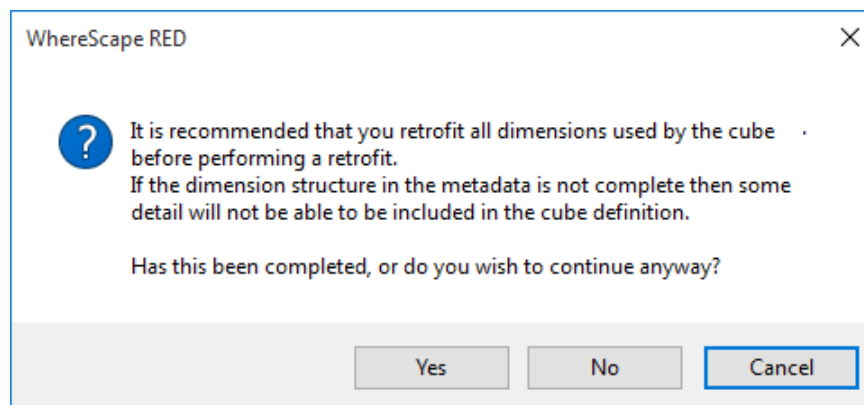
- 6 The second dialog warns that retrofitting a dimension will remove the links to cubes in the Metadata. Select **Yes**.



- 7 The results of the retrofit are displayed in the results panel at the bottom of the screen. If it was successful the new object should now have all the attributes, hierarchies and hierarchy levels (as well as all of their properties) as set in Analysis Services.

Once all of the dimensions have been retrofitted you can retrofit the OLAP Cube. Follow steps 1-6 above to retrofit the OLAP Cube as for the OLAP dimensions.

The final dialog is a reminder to only retrofit an OLAP cube once the dimensions used by the cube have been retrofitted. If all relevant dimensions have been retrofitted, select **Yes**.



Once again, the results of the retrofit are displayed in the results panel at the bottom of the screen.

CHAPTER 22

TRANSFORMATIONS

Standard **column transformations** (on page 594) can be used in WhereScape RED to perform calculations, change data types or format data.

Re-using complex transformations can save a significant amount of time. These can be achieved two ways in WhereScape RED:

- Teradata User Defined Functions (UDFs)
- WhereScape RED User Defined Transformations

IN THIS CHAPTER

Column Transformations.....	594
Teradata User Defined Functions	601
Re-usable Transformations	605

COLUMN TRANSFORMATIONS

Each table, view, join index or export object column can have a transformation associated with it. For all table types, except for load tables, views and join indexes, the transformation will be included in the generated procedure for the table. These are executed as part of the procedure update. The transformation must therefore be a valid SQL construct that can be included in a **Select** statement. For example, we could have a transformation of 'load_order_line.qty * 0.125' to calculate a tax column of 12.5%. Click the **Transformation** tab on the column properties to enter a transformation.

Note: Transformations added to an existing table that have an update procedure are only put into effect when the procedure is re-generated and re-compiled.

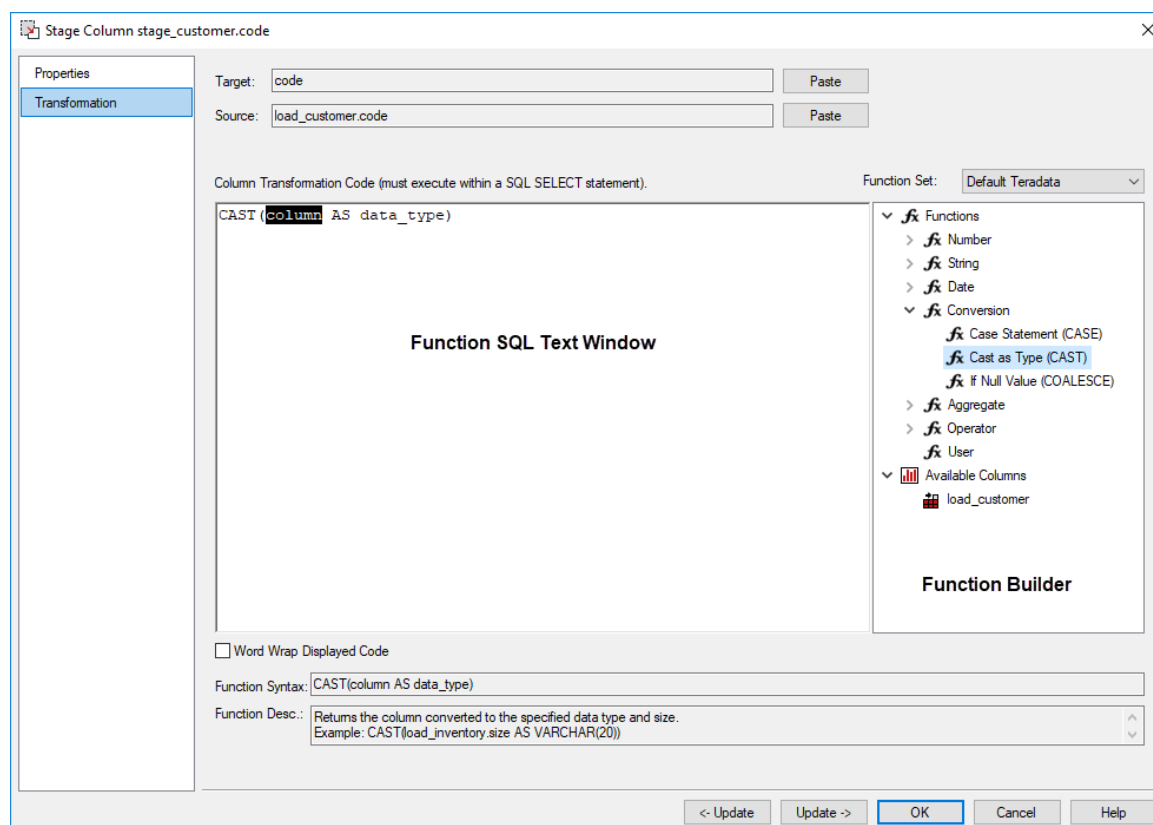
Column transformations on load tables are more complex, due to the unique nature of load tables. See *Load Table Column Transformations* (on page 597) for more details.

View and Join Index transformations are included in the Database Definition Language (DDL) that creates the object in the Teradata database. Any changes to transformations require these object types to be dropped and recreated.

Export object column transformations are dynamically applied for file loads. If the export object is executed via a host script, then the script needs to be regenerated for changes to transformations to take effect.

COLUMN TRANSFORMATION PROPERTIES

An example below shows the transformation property screen with a simple transformation:



The two special update keys allow you to update the column and step either forward or backward to the next column's properties. **ALT-Left Arrow** and **ALT-Right Arrow** can also be used instead of the two special update keys

Function SQL Text Window

The **Function SQL Text Window** contains the SQL used in the transformation. It can be directly entered, built up using the **Function Builder** and **Add** buttons or a combination of both.

Function Builder

The **Function Builder** contains a list of standard database functions, operators, user defined functions and all columns belonging to all source tables.

Expanding the **Function Heading** displays the **Function Groups** (Number, String, Data, Conversion, etc) and the **User Defined Function Heading**. Similarly, expanding the **Data Heading** displays **Source Tables**.

Each function group or source table can in turn be expanded to show individual **Functions** and **Source Columns**.

Double clicking on a function adds the **Function Model** to the Function SQL Text Window. The first variable (almost always the source column) is left highlighted in the Function SQL Text Window. This allows additional Functions or Source Columns to be added to the correct place in

the Function SQL Text Window by double clicking on the required Function or Source Column in the Function Builder.

Target Paste Button

The **Target Paste** button adds the current column in the form ColumnName to the Function SQL Text Window at the location of the cursor.

Source Paste Button

The **Source Paste** button adds the source table and column in the form TableName.ColumnName to the Function SQL Text Window at the location of the cursor.

Transform Stage

Only visible on load table column transformations. See *Load Table Column Transformations* (on page 597) for more information.

Function Set

This drop-down list enables the user to select which set of functions are to be displayed in the tree view when creating a transformation on a column of a table.

Update Buttons

The Update Buttons: **Update <-** and **Update ->** are used to move from the current column to previous and next columns respectively in the current table. The alternative is to exit the Column Transformation Properties, choose the next column, re-enter the Column Properties and choose the **Transformation** tab.

Function Syntax

The syntax guide for the **Function** visible when the function is clicked. Essentially the same as the function model loaded into the **Function SQL Text Window** when the function is double-clicked. Read only.

Function Desc

The description of the **Function** visible when the function is clicked. Read only.

Function Model

The model (template SQL code) for a **User Defined Transformations**. This is only visible for User Defined Transformations. Read only.

Localize Transformation

The **Localize transformation** button breaks the link between a **column transformation** and the **User Defined Transformation** it's based on. If this button is clicked, changes to the underlying user defined transformation cannot be automatically propagated to the column transformation. This is only visible for User Defined Transformations.

LOAD TABLE COLUMN TRANSFORMATIONS


Overview

Data entering the data warehouse can be manipulated if required. This manipulation can occur at any stage, but is supported via a number of methods during the **Load** stage. Load tables provide options to transform data. If multiple pass transformations are required then a load table can be created from another load table, i.e. multiple load tables can be supported in the data flow path.

The options available differ depending on the type of load but in most cases the **after** transformation and post load procedure can be utilized. Specifically:

Database Link Load	<ul style="list-style-type: none"> • During Load transformations • After Load transformations • Post Load procedure
ODBC Based Load	<ul style="list-style-type: none"> • During Load transformations • After Load transformations • Post load procedure
File Based Load	<ul style="list-style-type: none"> • During Load transformations • After Load transformations • Post load procedure
Integration Services Load	<ul style="list-style-type: none"> • During Load transformations • After Load transformations
Script Based Load	<ul style="list-style-type: none"> • During Load transformations • After Load transformations • Post load procedure
Externally Loaded	<ul style="list-style-type: none"> • After Load transformations • Post Load procedure

The **Transformation** tab of a column's properties is used to define **during** and **after** load transformations. It can only be one or the other for a specific column. One column can be used to build another, so an **after** can be based on the results of a **during**, if different columns are used.



Note: The **During** transformations use **Source Table** columns. The **After** transformations use the **Load Table** columns.

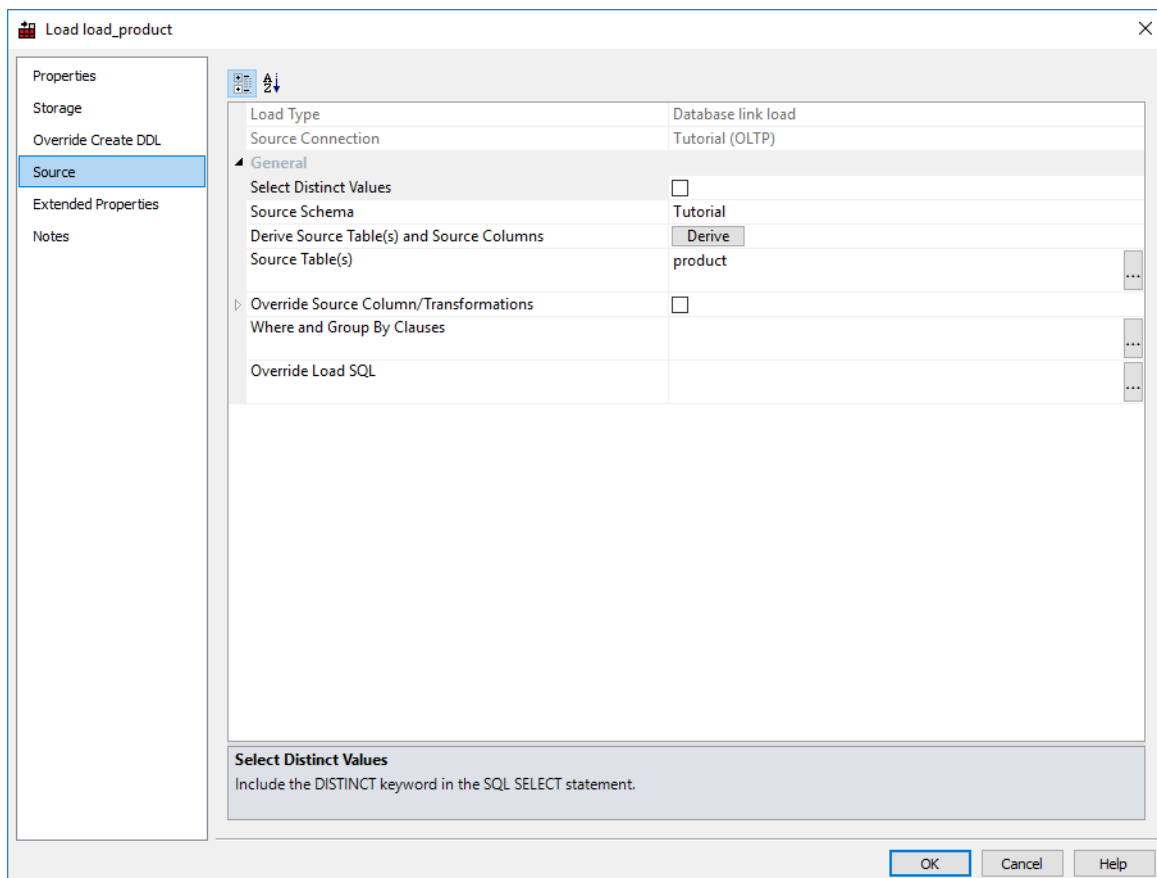
All **After** transformations take place in the native Teradata SQL language. The **During** transformations differ in terms of which language is used. This is particularly true for file based loads. Normally the **During** transformation will occur in the native SQL language of the source database.

For **Flat file loads** using SSIS, **After Load transformations** use the SQL syntax of the target database but **During Load transformations** use SSIS expression syntax that can be referred to on the Microsoft Developer Network website : [https://msdn.microsoft.com/en-us/library/ms137547\(v=sql.110\).aspx](https://msdn.microsoft.com/en-us/library/ms137547(v=sql.110).aspx).

DATABASE LINK DURING LOAD TRANSFORMATIONS

The **during** load transformation allows the manipulation of column data as it enters the data warehouse.

By default, **Database link loads** and **ODBC based loads** have **During** and **After** transformations enabled.



When **transformations** are enabled, the contents of a source table/source column for each column are used as the basis of the loading statement.

- If source table and source column are null, then a null is used.
- If data exists in the **Transformation** tab of a column's properties, then this transformation data is used instead of the source table/source column combination.

Example

The following load table columns will generate the load sql statement if no transformation data is present against these columns.

load_product Columns				
Column Name	Display Name	Data Type	Source Table	Source Column
code	code	numeric(6)	product	code
description	description	varchar(64)	product	description
prod_line	prod line	varchar(24)	product	prod_line
prod_group	prod group	varchar(24)	product	prod_group
subgroup	subgroup	varchar(24)	product	subgroup

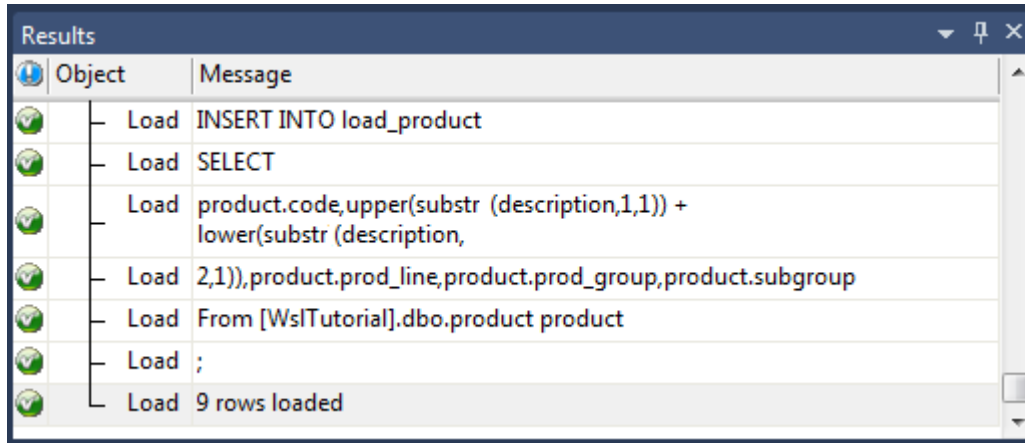
The SQL code from the load results is:

Results	
Object	Message
load_proc	Table load_product truncated.
Load	INSERT INTO load_product
Load	SELECT
Load	product.code,product.description,product.prod_line,product.prod_group,pr
Load	From [WsITutorial].dbo.product product
Load	;
Load	9 rows loaded

If the column 'description' has a transformation defined as follows:

upper(substr(description,1,1)||lower(substr(description,2,1))

then the following SQL statement will be executed.



FILE DURING LOAD TRANSFORMATIONS

The loading of flat files is performed using Fastload, Multi-load or TPT. The contents of the **Transformation** tab in a column's Properties are the functions and conversions supported by the database loader.

Example

Multiload performs transformations such as:

- FORMAT 'DD-MMM-YYYY' converts from a value such as 23-Mar-1999 to a Teradata date.
- CAST(:COL AS NUMERIC(18,4)) converts the data to a numeric(18,4)
- COALESCE(LTRIM(RTRIM(:COL)),0) trims leading and trailing white characters and inserts zero if null.

AFTER LOAD TRANSFORMATIONS

After transformations will initiate a pass of the load table after the load has occurred. They allow manipulation of the load table columns using the database and SQL functions.

Example

The following **after** transformation set for the column `code` in the table `load_product`

`substr(code,1,5)`

would result in the following SQL statement being executed after the load:

`update load_product set code = substr(code,1,5);`

TERADATA USER DEFINED FUNCTIONS

Teradata User Defined Functions (UDFs) can be built using the C programming language and compile directly onto the Teradata server and registered via a SQL command.

This can be daunting for many data warehouse developers, who may not know C well, and may not have access to compile UDFs on the Teradata server.

TERADATA UDF EXAMPLE

A standard example given for Teradata UDFs is building a REPLACE function similar to the Oracle database's REPLACE function.

Here is a simplified example:

Step 1: Write c code for UDF as replace.c:

```
void replace ( VARCHAR_LATIN *inputStr,
              VARCHAR_LATIN *inputFrom,
              VARCHAR_LATIN *inputTo,
              VARCHAR_LATIN *result,
              int *inputStrIsNull,
              int *inputFromIsNull,
              int *inputToIsNull,
              int *resultIsNull,
              char sqlstate[6],
              SQL_TEXT extname[129],
              SQL_TEXT specific_name[129],
              SQL_TEXT error_message[257])
{
  int linputFrom;
  int linputTo;
  char *pinputStr;
  strcpy(sqlstate, NoSqlError);
  strcpy((char *) error_message, " ");
  *resultIsNull = IsNotNull;
  *result = '\0';
  if (*inputStrIsNull == IsNull)
  {
    strcpy(sqlstate, "22004");
    strcpy((char *) error_message, "First argument cannot be null.");
    *resultIsNull = IsNull;
    return;
  }
  if (*inputFromIsNull == IsNull || (linputFrom=strlen((const char *)inputFrom)) == 0)
  {
    strcpy((char *)result, (const char *)inputStr);
    return;
  }
  linputTo = (*inputToIsNull == IsNull ? 0 : strlen((const char *)inputTo));
  for (pinputStr = (char *)inputStr; *pinputStr != '\0'; )
  {
    if (strncmp((const char *)pinputStr, (const char *)inputFrom, linputFrom) == 0)
    {
      strncat((char *)result, (const char *)inputStr, (int)(pinputStr-(char *)inputStr));
      if (linputTo != 0) strcat((char *)result, (const char *)inputTo);
      pinputStr = (char *)inputStr = pinputStr + linputFrom;
    }
    else
      ++pinputStr;
  }
  strcat((char *)result, (const char *)inputStr);
}
```

Step 2: Put c file onto the Server

Copy the c program file onto the data base server into the following location:

C:\program files\ncr\tdat\tdconfig\tdbs_udf\replace.c

Step 3: Register and Compile UDF via SQL:

```
CREATE FUNCTION replace
( Str VARCHAR(512)
, aFrom VARCHAR(512)
, aTo VARCHAR(512)
)
RETURNS VARCHAR(512)
LANGUAGE C
NO SQL
SPECIFIC replace2
EXTERNAL NAME 'SS!replace!C:\program files\ncr\tdat\tdconfig\tdbs_udf\replace.c'
PARAMETER STYLE SQL;
```

Step 4: Use the Function

Simply use the UDF in a RED column transformation or 'Where' clause like any other SQL function.

RE-USABLE TRANSFORMATIONS

WhereScape RED Re-usable Transformations allow a complex SQL transformation using standard Teradata functions (including UDFs) to be built once and reused in multiple column transformations. Functionality is included to manage and propagate changes to user defined transformations.

CREATING A NEW RE-USABLE TRANSFORMATION

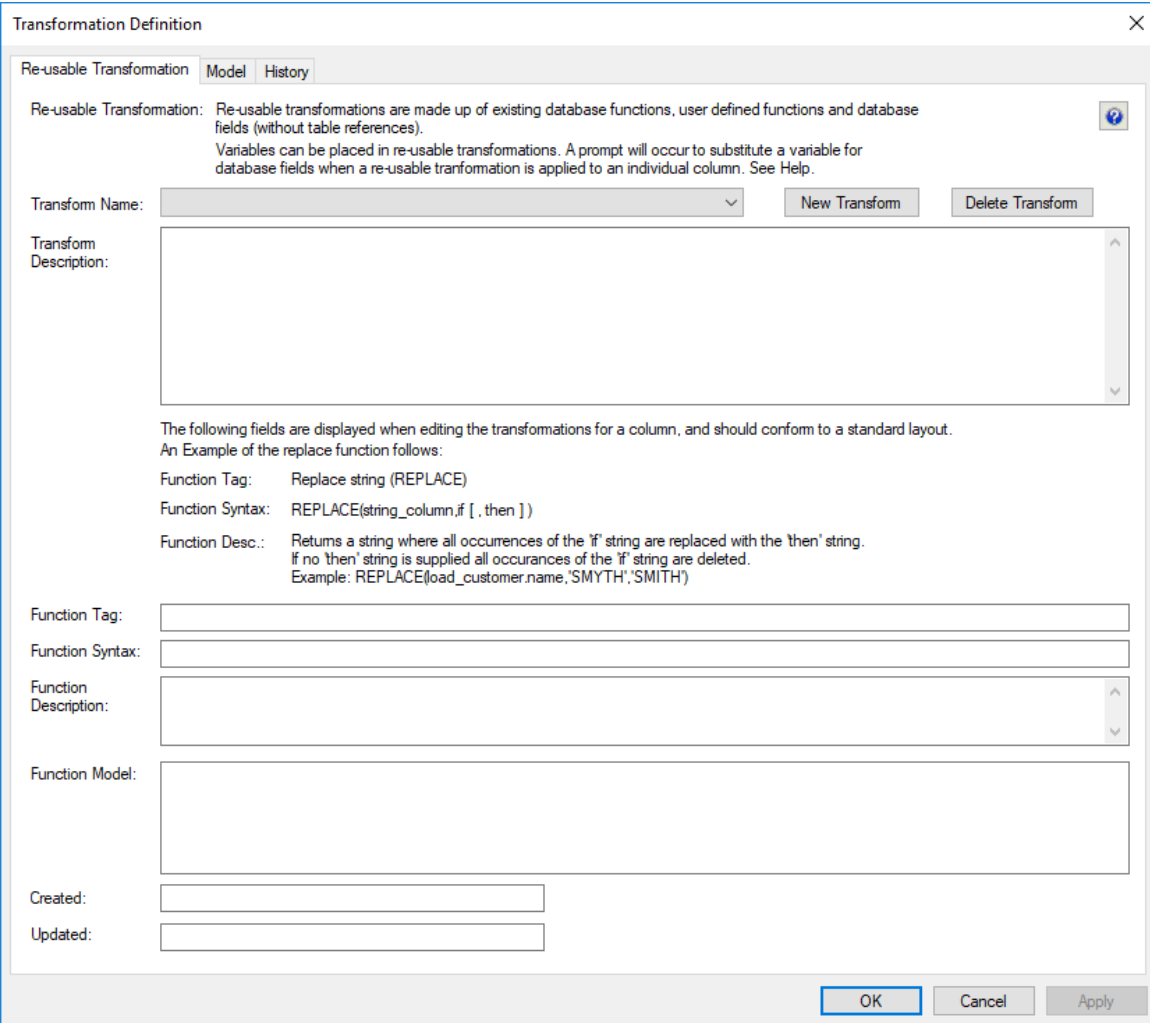
New **re-usable transformations** are created from the **Tools>Define Re-Usable Transformations** menu.

Creating a new re-usable transformation is a three step process:

- Specify the name of the transformation
- Enter metadata for the transformation
- Define the transformation

SPECIFY THE NAME OF THE TRANSFORMATION

Selecting **Define Re-Usable Transformations** from the **Tools** menu displays the following screen:

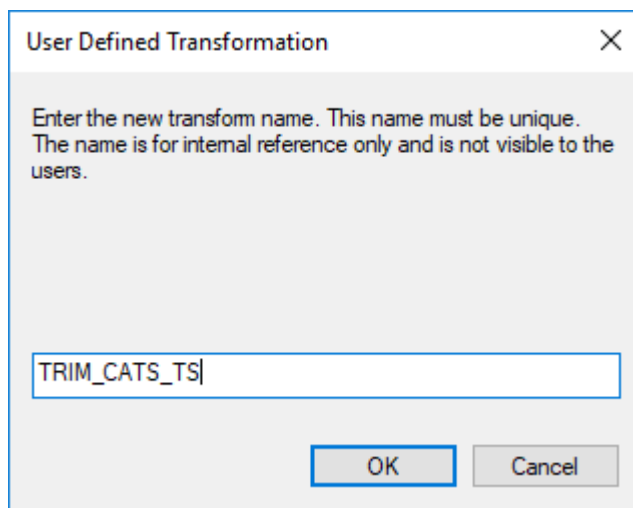


The image shows a 'Transformation Definition' dialog box with the following fields and controls:

- Re-usable Transformation:** A text area containing the text: "Re-usable transformations are made up of existing database functions, user defined functions and database fields (without table references). Variables can be placed in re-usable transformations. A prompt will occur to substitute a variable for database fields when a re-usable transformation is applied to an individual column. See Help." A help icon is visible to the right.
- Transform Name:** A dropdown menu with a 'New Transform' button to its right and a 'Delete Transform' button to its left.
- Transform Description:** A large text area for entering the description of the transformation.
- Example text:** "The following fields are displayed when editing the transformations for a column, and should conform to a standard layout. An Example of the replace function follows:"
- Function Tag:** A text field with the value "Replace string (REPLACE)".
- Function Syntax:** A text field with the value "REPLACE(string_column,if [, then])".
- Function Desc.:** A text field with the value "Returns a string where all occurrences of the 'if' string are replaced with the 'then' string. If no 'then' string is supplied all occurrences of the 'if' string are deleted. Example: REPLACE(load_customer.name,'SMYTH','SMITH)".
- Function Tag:** An empty text field.
- Function Syntax:** An empty text field.
- Function Description:** An empty text area.
- Function Model:** An empty text area.
- Created:** An empty text field.
- Updated:** An empty text field.

At the bottom right, there are three buttons: **OK**, **Cancel**, and **Apply**.

Click **New Transform** and enter a name for the User defined transformation:



A dialog box titled "User Defined Transformation" with a close button (X) in the top right corner. The dialog contains the following text: "Enter the new transform name. This name must be unique. The name is for internal reference only and is not visible to the users." Below the text is a text input field containing the text "TRIM_CATS_TS". At the bottom of the dialog are two buttons: "OK" and "Cancel".

Note: This is the internal WhereScape RED name for the transformation, not the name developers reference when utilizing the transformation on column transformations.

Click **OK**.

ENTER RE-USABLE TRANSFORMATION METADATA

Enter the following metadata for the transformation to describe the transformation for developers.

Transform Description

A general description of the transformation.

Function Tag

This is the name the function will appear as for users to select from the function builder when building column transformations.

Function Syntax

The syntax guide for the function. This is visible in the function builder when clicking on the User defined function.

Function Description

The description of the function, visible in the function builder when clicking on the User defined function.

Function Model

The model (template SQL code) for a User Defined Transformations. This is only visible for User Defined Transformations. [Read only].

DEFINE THE TRANSFORMATION MODEL

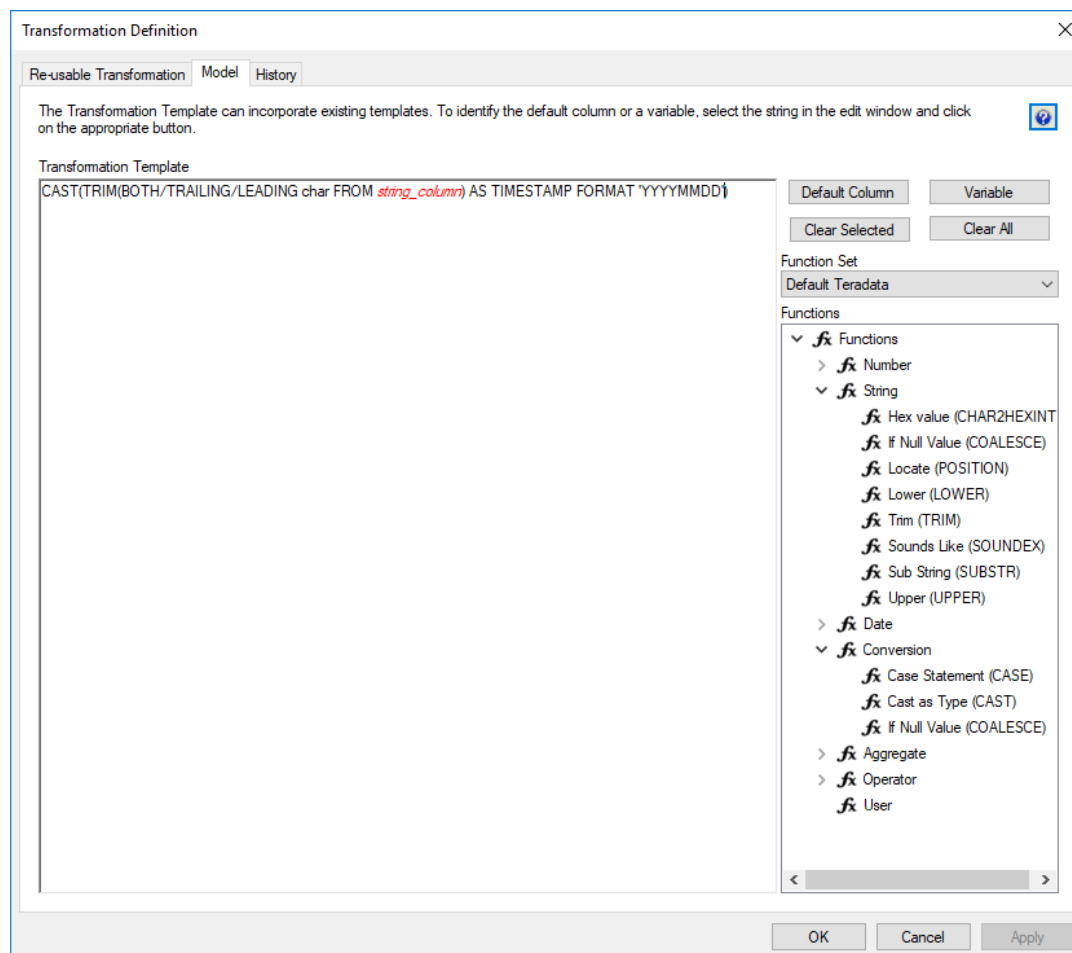
Once the transformation has been created and the metadata entered, the actual SQL code used by the transformation can be defined. The SQL code can be entered directly or via the **Function Builder** on the right side. To use the function builder, expand the tree to find the required function.

Example of Building a Transformation Model

To build a model to CAST a trimmed string in YYYYMMDD format to a timestamp, do the following:

- Click on the **Model** tab
- Expand **Functions** heading
- Expand the **Conversion** heading
- Double click on **Cast as Type (CAST)**
- Expand the **String** heading
- Double click on **Trim (TRIM)**
- Highlight **data_type** and type **TIMESTAMP FORMAT 'YYYYMMDD'**

You should see the following:



Click **OK**.

COMPLETED RE-USABLE TRANSFORMATION

Transformation Definition ✕

Re-usable Transformation Model History

Re-usable Transformation: Re-usable transformations are made up of existing database functions, user defined functions and database fields (without table references).
 Variables can be placed in re-usable transformations. A prompt will occur to substitute a variable for database fields when a re-usable transformation is applied to an individual column. See Help.

Transform Name: New Transform Delete Transform

Transform Description:

The following fields are displayed when editing the transformations for a column, and should conform to a standard layout.
 An Example of the replace function follows:

Function Tag: Replace string (REPLACE)
 Function Syntax: REPLACE(string_column,if [, then])
 Function Desc.: Returns a string where all occurrences of the 'if' string are replaced with the 'then' string. If no 'then' string is supplied all occurrences of the 'if' string are deleted.
 Example: REPLACE(load_customer.name,'SMYTH','SMITH')

Function Tag:

Function Syntax:

Function Description:

Function Model:

Created:

Updated:

OK Cancel Apply

CHANGING A RE-USABLE TRANSFORMATION

To change a re-usable transformation:

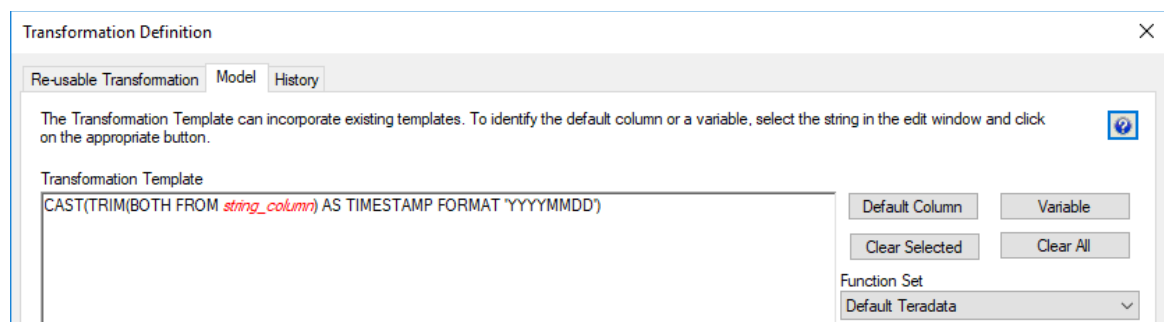
- Select **Re-usable transformations** from the Tools menu.
- Choose the transformation from the **Transform Name** drop-down list.
- Click on the **Model** Tab.
- Change the SQL as required.
- Click **OK**.

Example of a change to the Model SQL

In the example used in *Creating a New User Defined Transformation* (see "*Creating a New Re-usable Transformation*" on page 605), the SQL was:

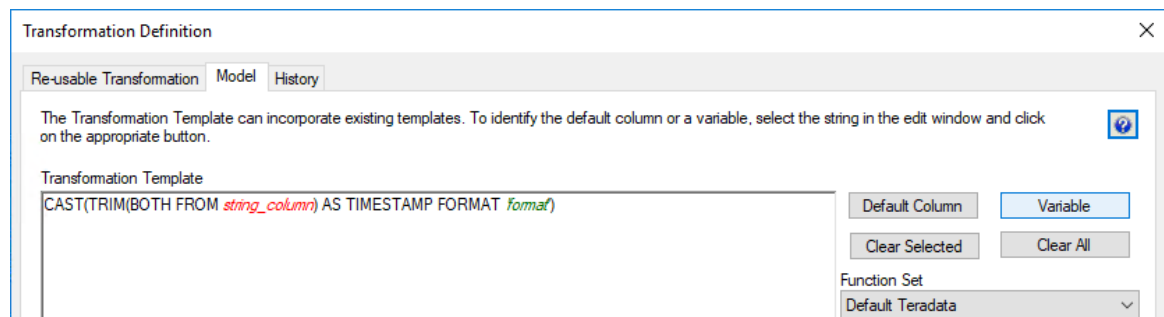
CAST(TRIM(BOTH FROM string_column) AS TIMESTAMP FORMAT 'YYYYMMDD')

Change the SQL to allow the format to be specified when the transformation is used by changing **YYYYMMDD** to **format**.



Then highlight the word **format** and click the **Variable** button. This makes the word **format** a variable than can be substituted when the User Defined Transformation is used.

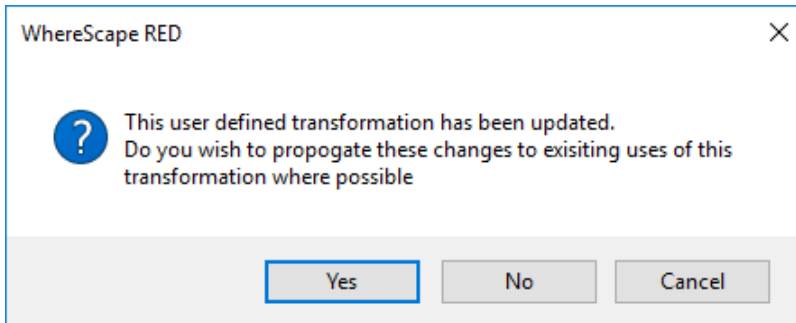
Now **format** is green and in italics:



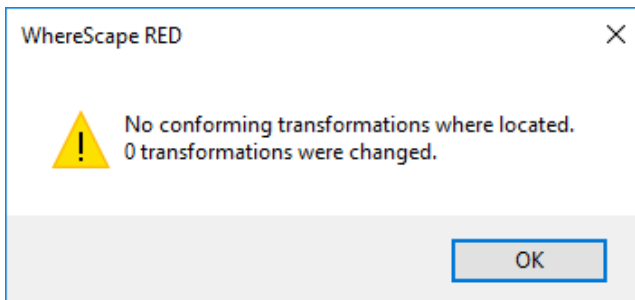
Click **OK**.

APPLYING CHANGES TO DEPENDANT TRANSFORMATIONS

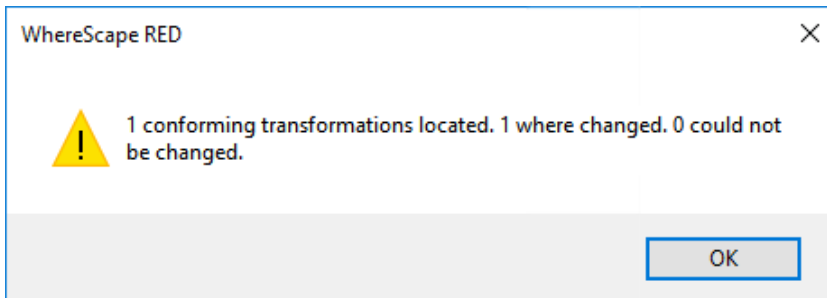
After changing a **Re-usable Transformation**, a dialog appears asking to confirm that changes should be applied to individual columns using the transformation, where possible:



If the **Re-usable Transformation** doesn't have any dependant columns, then the following message is displayed:

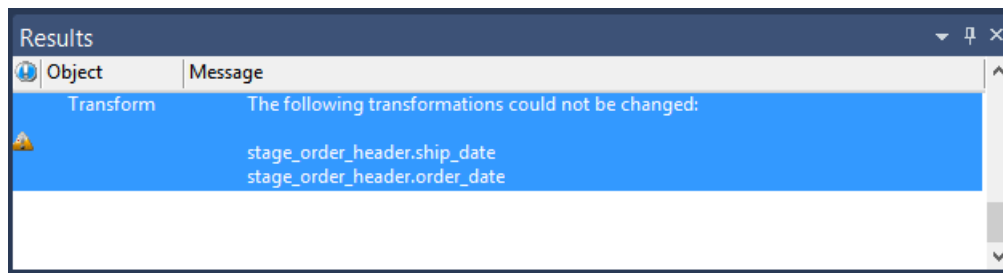


If the **Re-usable Transformation** has been used for dependant columns then this message is displayed:



When an attempt is made to update a dependant Transformation, and the transformation has been modified in such a way as to make it impossible for the changes to the User Defined Transformation to be applied, the error message above will include a count of the failures.

The results pane detail the columns (and tables) where the update failures have occurred:



USING RE-USABLE TRANSFORMATIONS

User defined transformations are used exactly the same way as any standard database **Function**. They can be used on any object type. See **Column Transformation Properties** (on page 595)

CHAPTER 23

EXPORTING DATA

Export objects are used in RED to produce ascii files from a single database table or view for a downstream feed. Some or all of the columns in a table or view can be exported.

There are three ways of performing exports.

- **File export** - an export where most of the processing is managed and controlled by the scheduler.
- **Script-based export** - an export where a Windows or a Unix/Linux script file is executed to perform the export. Script-based exports on Windows supports both DOS Batch and *PowerShell scripts* (see "*Windows PowerShell Scripts*" on page 653).
- **Integration Services export** - an export processed using a Windows connection where the processing is handled via an Integration Services Package that is generated and executed dynamically at run time.
SSIS exports to UNIX/Linux connections and processed via the UNIX/Linux scheduler are currently not supported.

IN THIS CHAPTER

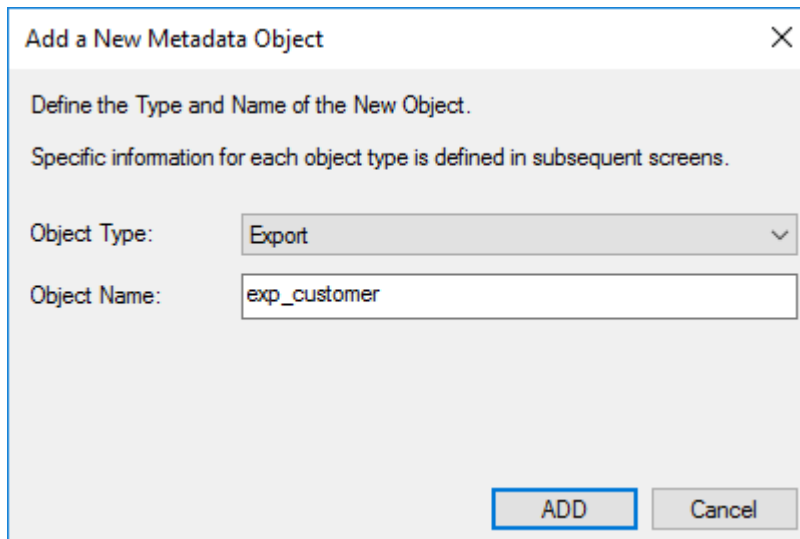
Building an Export Object.....	615
File Attributes	619
Export Column Properties	624
Script based Exports.....	626

BUILDING AN EXPORT OBJECT

Creating an Export:

- 1 Browse the Data Warehouse connection.
- 2 Create a drop target by double clicking on the **Export object group** in the left pane. The middle pane should have a column heading of **Export Objects** for the leftmost column.
- 3 Select a table or view in the right pane and drag it into the middle pane. Drop the table or view anywhere in the middle pane.

The following dialog appears:



Add a New Metadata Object

Define the Type and Name of the New Object.

Specific information for each object type is defined in subsequent screens.

Object Type:

Object Name:

- 4 Rename the export object if it needs to be renamed and click **ADD**.

The Properties dialog displays:

Export exp_customer

Properties

File Attributes

Storage

Extended Properties

Purpose

Concept

Grain

Examples

Usage

Notes

Export Object Name: exp_customer

Unique Short Name: (maximum 22 characters) exp_customer

Description:

Connection:

Export Type: File export

Database Link:

Script Template: (None)

Script Name: (None)

Pre-Export Action: No action

Pre-Export Sql:

Where Clause: Allows filtering of the export data.

Post Export Procedure: (None)

Timestamps

Metadata Structure Changed: 2017-03-15 00:30:05.100000

Last Exported:

OK Cancel Help

5 If exporting the object via **Windows**, change the Connection to **Windows**.

From a Windows connection, **File**, **Script based** and **Integration Services** export methods are supported.

Export exp_customer

Properties

File Attributes

Storage

Extended Properties

Purpose

Concept

Grain

Examples

Usage

Notes

Export Object Name: exp_customer

Unique Short Name: (maximum 22 characters) exp_customer

Description:

Connection: Windows

Export Type: File export

Database Link:

Script Template: (None)

Script Name: (None)

Pre-Export Action: No action

Pre-Export Sql:

Where Clause: Allows filtering of the export data.

Post Export Procedure: (None)

Timestamps

Metadata Structure Changed: 2017-03-15 00:30:05.100000

Last Exported:

OK Cancel Help

6 If exporting the object via **Unix/Linux**, change the connection to **Unix**.

From a **Unix/Linux** connection **File export** and **Script based** exports methods are supported.



TIP: When doing a **Script based export** on Windows or Unix/Linux, use the **Rebuild** button after selecting the relevant script to be rebuilt on the the Script Name drop-down menu.

- Click the **File Attributes** tab to fill in the fields described in the next section to set the relevant location, name and contents of the exported data file.
- When the File Attributes fields are filled in, go back to the **Properties** tab, select (**Build Script**) from the **Script Name** drop-down menu and click **OK**.
- Export the object after filling in the Files Attributes fields by right-clicking on the export object on the left pane and selecting **Export**.

FILE ATTRIBUTES

The following fields are available to define the location, name and contents of the exported data file:

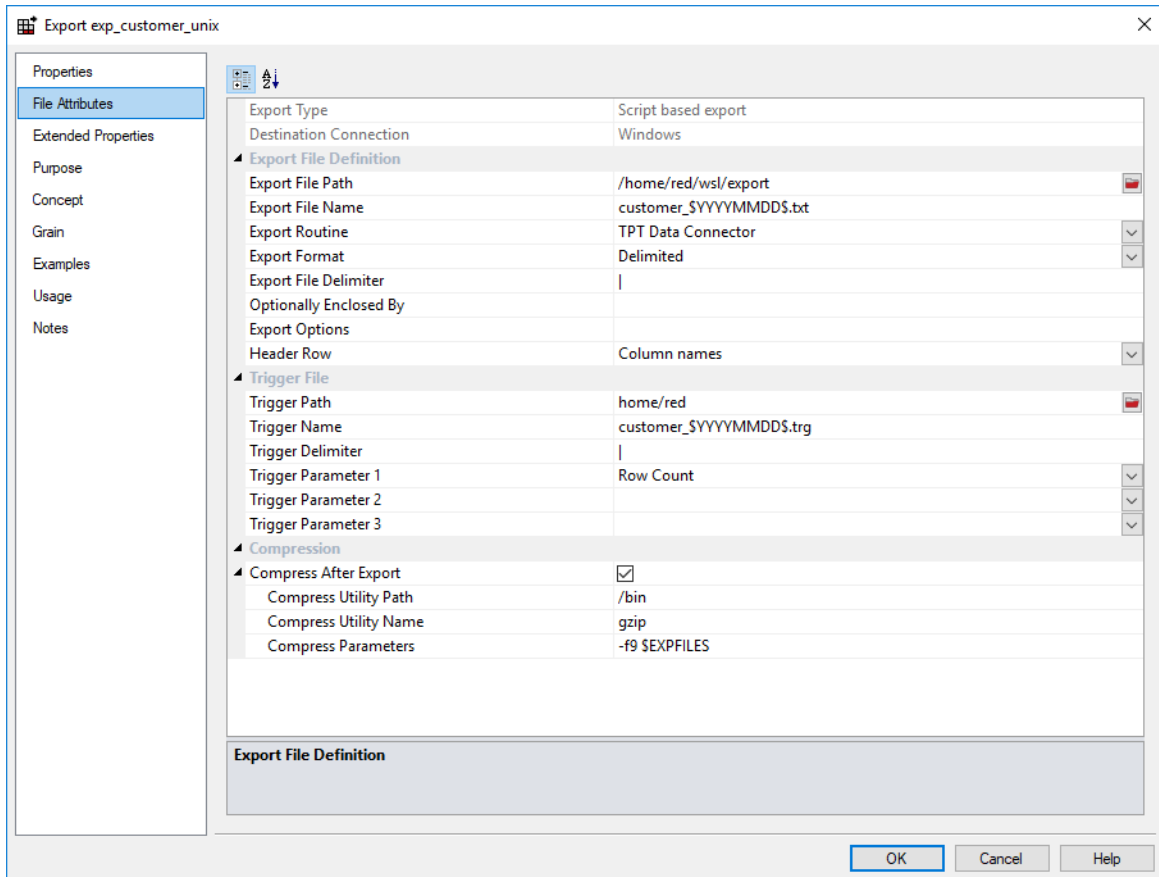
- **Export Type** - method of exporting data from the table. The available options are dependent on the Destination connection that can be specified via the Properties page.
- **Destination Connection** - destination to the file system to which data will be exported. The destination connection can be specified on the Properties screen.
- **Export File Path** - the full path (absolute path) of the folder/directory where the File is to be created on the Windows or UNIX/Linux system.
- **Export File Name** - name of the Export File to which the data will be exported. The variable \$SEQUENCE\$ can be used to provide a unique sequence number for the export file. Also, the data/file components YYYY, MM, HH, MI, SS can be used when enclosed with the \$ character. For example, an export file name might be customer_YYYYMMDD\$.txt which would result in a file name like customer_20150520.txt.
- **Export Routine** - database specific routine to use to export the data. FastExport or TPT Data Connector for Windows exports, or TPT Data Connector for Unix/Linux exports.
- **Export Format** - routine-specific format to use to export the data. Select one of the default Teradata formats: Text, Unformatted, Delimited, Binary, Formatted, FastLoad or a special WhereScape RED derived format: Delimited or Width Fixed Text.
- **Export File Delimiter** - character that separates the fields within each record of the Export File for Delimited formats. The delimiter identifies the end of which field. Common field delimiters are tab, comma, colon, semi-colon, pipe. To enter a special character enter the uppercase string CHAR with the ASCII value in brackets (e.g. CHAR(9)). This is only available if the Export Format is Delimited text.
- **Optionally Enclosed by** - character that brackets text fields within each record of the Export File for Delimited formats. A common example is ". This is only available if the Export Format is Delimited Text.
- **Export Options** - allows the entry of export utility options. If more than one option is required a semi-colon should be used between options.
- **Header Row in Export** - if a header line is required, choose business names or column names from this drop-down list.
- **Trigger Path** - the trigger file indicates that the export to the main file has completed and it is now safe to load the file. Secondly, the trigger file may contain control sums to validate the contents of the main load file. This field should contain the full path name to the directory in which a trigger file is to be generated on the destination system.
- **Trigger Name** - this refers to the name of the file that is to be created as a trigger file. A trigger file typically contains check sums (row count or the sum of a numeric column). The variable \$SEQUENCE\$ can be used to provide a unique sequence number for the trigger file. Also, the data/file components YYYY, MM, HH, MI, SS can be used when enclosed with the \$

character. For example, a trigger file name might be customer_YYYYMMDD\$.txt which would result in a file name like customer_20150520.txt.

- **Trigger Delimiter** - multiple fields in the trigger file need to be separated by the trigger delimiter.
- **Trigger Parameter 1,2,3** - checksums to be put in the trigger file. One of the row count and the sum of any numeric fields in the source data.
- **Compress After Export** - tick this check-box if you want to compress the export file after it has been created.
- **Compress Utility Path** - directory in which the compress utility exists.
- **Compress Utility Name** - name of the compression utility executable.
- **Compress Parameters** - name of the file to be compressed (using the RED variable \$EXPFIL\$) and any commands or switches required to make the compression utility work. These parameters will depend on the compression utility used.

For Windows **File** or **Script based** Exports, you can select the between the **FastExport** or **TPT Data Connector** export routines.

1 Example screen for a Windows Script based Export using the TPT Data Connector.



2 Example screen for a **Unix/Linux** File/Script based Export using the **TPT Data Connector**.

Export Type	File export
Destination Connection	UNIX
▲ Export File Definition	
Export File Path	c:\data\
Export File Name	customer_SYYYYMMDD\$.txt
Export Routine	TPT Data Connector
Export Format	Delimited
Export File Delimiter	
Optionally Enclosed By	
Export Operator Attr.	
Header Row	Column names
▲ Trigger File	
Trigger Path	c:\data
Trigger Name	customer_SYYYYMMDD\$.txt
Trigger Delimiter	
Trigger Parameter 1	Row Count
Trigger Parameter 2	
Trigger Parameter 3	
▲ Compression	
Compress After Export	<input checked="" type="checkbox"/>
Compress Utility Path	c:\Program Files\7-Zip\
Compress Utility Name	7z.exe
Compress Parameters	a -y\$EXPFILES.zip \$EXPFILES

Export File Path
Full path (absolute path) of the folder/directory containing the Export File on the Windows or UNIX/Linux system.



TIP: With all the relevant fields filled in, remember to go back the Properties tab and select **(Build Script)** from the Script Name drop-down menu.

Right-click on the export object on the left pane and select **Export**.

FILE ATTRIBUTES - SSIS EXPORTS

The following fields below are available to define the location, name and definitions of the exported data file:

Integration Services File Attributes screen:

The screenshot shows the 'Export exp_customer' dialog box with the 'File Attributes' tab selected. The properties are as follows:

Export Type	Integration Services export
Destination Connection	Windows
Export File Definition	
Export File Path	c:\data\
Export File Name	customer_SYYYYMMDDS.txt
Export Format	Delimited
Export File Delimiter	
Optionally Enclosed By	
Header Row	Column names
SQL Server Integration Services (SSIS)	
SSIS Row Count Log	<input checked="" type="checkbox"/>

Export File Path
Full path (absolute path) of the folder/directory containing the Export File on the Windows or UNIX/Linux system.

Buttons: OK, Cancel, Help

Export Type

Method of exporting data from the table. The available options are dependent on the Destination connection that can be specified via the Properties page.

Destination Connection

Destination to the file system to which data will be exported. The destination connection can be specified on the Properties screen.

Export File Definition

Export File Path

The full path (absolute path) of the folder/directory/ where the File is to be created on the Windows or UNIX/Linux system.

Export File Name

Name of the Export File to which the data will be exported. The variable \$SEQUENCE\$ can be used to provide a unique sequence number for the export file. Also the data/file components YYYY, MM, HH, MI, SS can be used when enclosed with the \$ character. For example an export file name might be customer_YYYYMMDD\$.txt which would result in a file name like customer_20150520.txt.

Export Format

Routine-specific format to use to export the data.

Export File Delimiter

Character that separates the fields within each record of the Export File for Delimited formats. The delimiter identifies the end of which field. Common field delimiters are tab, comma, colon, semi-colon, pipe. To enter a special character enter the uppercase string CHAR with the ASCII value in brackets (e.g. CHAR(9)). This is only available if the Export Format is Delimited Text.

Optionally Enclosed by

Character that brackets text fields within each record of the Export File for Delimited formats. A common example is ". This is only available if the Export Format is Delimited Text.

Header Row

If a header line is required, choose business names or column names from this drop-down list. This option is not available in DB2.

SQL Server Integration Services (SSIS)

SSIS Row Count Log

When enabled, this option includes Row Count logging on SSIS exports.

EXPORT COLUMN PROPERTIES

Export Column exp_product.prod_group

Properties

Transformation

General

Table Name	exp_product
Column Name	prod_group
Column Title	prod group
Column Description	

Physical Definition

Column Order	40
Data Type	varchar(24)
Format	

Meta Definition

Numeric	<input type="checkbox"/>
Additive	<input type="checkbox"/>
Attribute	<input checked="" type="checkbox"/>

Source Details

Source Table	load_product
Source Column	prod_group
Transformation	

Column Name

Database-compliant name of the column.
Dialog Opening Value: prod_group

<- Update Update -> OK Cancel Help

Table Name

Database-compliant name of the table that contains the column. [Read-only field].

Column Name

Database-compliant name of the column. Typically, column-naming standards exclude the use of spaces etc. A good practice is to only use alphanumeric, and the underscore character. Changing this field alters the table's definition.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Title

Name that the business uses to refer to the column, which is included in the RED-generated documentation and can be used in the end user layer of other tools. [Does NOT affect the physical database table]. As such, it is a free form entry and any characters are valid.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Description

This field contains the description for the column. It may be a description from a business user's point of view. This field might additionally contain information on where and how the column was acquired. For example, if the column is sourced from multiple tables or is a composite or derived column, then this definition would normally describe the process used to populate the column. This field is used in the documentation and is available via the view **ws_admin_v_dim_col**. It is also stored as a comment against the column in the database.

Column Order

Numeric value that controls the relative order of columns in the database create statement. The lowest numbered column will appear first in the table. Although this affects the physical table definition, no action is taken unless the table is re-created in the database. The columns can be re-numbered based on the existing order by choosing the **Respace Order Number** pop-up menu option when positioned over any column in the table. This action numbers the columns in increments of 10, starting at 10. In addition to a simple change of the order field, the column order can be changed by first displaying the columns in the middle pane and then using drag and drop to move the columns around. This drag and drop process automatically rennumbers the columns as required.

Data Type

Database-compliant data type that must be valid for the target database. Typical Teradata databases often have integer, numeric(), varchar(), char(), date and timestamp data types. See the database documentation for a description of the data types available. Changing this field alters the table's definition.

Format

Not relevant for Export Objects

Numeric

Not relevant for Export Objects

Additive

Not relevant for Export Objects

Attribute

Not relevant for Export Objects

Source Table

Identifies the source table where the column's data comes from. This source table is normally a load table within the data warehouse. If the column was sourced from multiple tables, then the normal practice is to record one of the tables in this field and a comment listing all of the other tables in the **Source Strategy** field. This field is used when generating a procedure to update the dimension. It is also used in the track back diagrams and in the documentation.

Source Column

Identifies the source column where the column's data comes from. Such a column is normally a load table column, which in turn may have been a transformation or the combination of multiple columns. For previous value managed columns the source column is the column in the table whose previous value is being recorded.

Transformation

See *Column Transformations* (on page 594). [Read-only].

SCRIPT BASED EXPORTS

A script based export object will have a Host Script defined. During the export process, this host script is executed and the results returned.

During the **drag and drop** creation of an export object from a single table or view, a script can be generated by selecting one of the 'Script based' export options. This script can then be edited to more fully meet any requirements.

There are a number of conventions that must be followed if these host scripts are to be used by the WhereScape scheduler.

- 1 The first line of data in **standard out** must contain the resultant status of the script. Valid values are '1' to indicate success, '-1' to indicate a Warning condition occurred but the result is considered a success, '-2' to indicate a handled Error occurred and subsequent dependent tasks should be held, -3 to indicate an unhandled Failure and that subsequent dependent tasks should be held.
- 2 The second line of data in **standard out** must contain a resultant message of no more than 256 characters.
- 3 Any subsequent lines in **standard out** are considered informational and are recorded in the audit trail. The normal practice is to place a minimum of information in the audit trail. All bulk information should be output to **standard error**.
- 4 Any data output to **standard error** will be written to the error/detail log. Both the audit log and detail log can be viewed from the RED tool under the scheduler window.
- 5 When doing **Script based exports**, it is easy to use the **Rebuild** button to the right of the Script-name field to rebuild the scripts.

The screenshot shows the 'Export exp_customer' dialog box with the following fields and values:

- Export Object Name: exp_customer
- Unique Short Name (maximum 22 characters): exp_customer
- Description: (empty)
- Connection: Windows
- Export Type: Script based export
- Database Link: (empty)
- Script Template: (None)
- Script Name: exp_exp_customer (with Edit and Rebuild buttons)
- Pre-Export Action: No action
- Pre-Export Sql: (empty)
- Where Clause: Allows filtering of the export data. (empty)
- Post Export Procedure: (None)
- Timestamps:
 - Metadata Structure Changed: 2017-03-15 00:30:05.100000
 - Last Exported: (empty)

Note: Script-based exports on Windows supports both DOS Batch and **PowerShell scripts** (see "**Windows PowerShell Scripts**" on page 653).

CHAPTER 24

PROCEDURES AND SCRIPTS

WhereScape RED has a **Procedure object group** for database stored procedures and a **Script** object group for host system scripts, such as Windows batch files.

Procedures

RED generates the bulk of the procedures during a prototype build, but these procedures can be customized. In fact it would be normal practice once the prototype phase is completed to modify these procedures to meet specific requirements. The procedure object group refers to the concept of database stored procedures. Specific objects may in fact be functions, procedures or packages. In this chapter the generation, editing and compilation of procedures is covered.

Scripts

Host scripts are generated when a script based file load is chosen during a file drag and drop from a Windows connection. Scripts can also be created manually provided the rules for their inclusion into the scheduling process are followed. This chapter covers the generation, editing and testing of host scripts as well as explaining the components required to allow them to work in the scheduler environment.

Notes:

- Windows script supports both DOS Batch and **PowerShell scripts** (see "**Windows PowerShell Scripts**" on page 653).
 - While Windows scripts are available on Load table objects, currently only Windows Scripts are available on other object types that are created in a Custom Database Connection.
-

IN THIS CHAPTER

Procedure Generation	631
Procedure Editing	637
Procedure Loading and Saving.....	640
Procedure Comparisons	642
Procedure Compilation	643
Procedure Running	643
Procedure Syntax	644
Procedure Properties	645
Macros.....	647
BTEQ Scripts	647
Script Generation.....	647
Script Editing	658
Script Testing.....	659
Script Syntax.....	659
Script Environment Variables.....	661
Calling a Batch File from a Script	667
Scheduling Scripts	669
Manually created scripts.....	671

PROCEDURE GENERATION

WhereScape RED generates template procedures to assist in the various phases of the data warehouse build process. A procedure is generated by selecting the **(Build Procedure...)** option from the drop-down list field in a table's Properties window to configure the update, initial build, or post load procedures.

For **Load** tables a post load procedure can be generated by selecting the option above. This post load procedure is designed to assist in the management of file loads where a trigger file has been used.

Post load procedures can either be manually generated from a RED provided procedure outline or generated leveraging a RED template—refer to *Rebuilding Update Procedures* (on page 199) for details.

For **Stage, Dimension, Fact** and **Aggregate** tables, an Update or Initial Build procedure can be generated by selecting from the corresponding procedure drop-down list field in the table's Properties.

If a new procedure is created from scratch (i.e. not auto generated), then an outline of the syntax required by the WhereScape scheduler can be generated by selecting the **Tools>Create Procedure Outline** menu option in the procedure editor.

Wrapper procedures

In some cases, multiple procedures is required to update a table. In such cases, it is best to create a top level procedure that is seen by the scheduler as the 'Update' procedure. This procedure can in turn call other procedures.

Example

We may have a model table that is updated from multiple stage tables. This wrapper procedure calls two child procedures, one for each stage table that is to update the model table. A status is reported back to the audit trail for each stage and an overall status ascertained for the model table update.

The wrapper procedure may look as follows:

```
CREATE PROCEDURE [METABASE].update_model_order_header
(
  IN   p_sequence      integer,
  IN   p_job_name      varchar(256),
  IN   p_task_name     varchar(256),
  IN   p_job_id        integer,
  IN   p_task_id       integer,
  OUT  p_return_msg    varchar(256),
  OUT  p_status        integer
)
BEGIN
  -----
  -- DBMS Name       : Teradata
  -- Script Name     : update_model_order_header
  -- Description     : Build the table model_order_header
  -- Generated by    : WhereScape RED
  -- Generated for   : WhereScape Limited
  -- Author         : WhereScape Developer
  -----
  -- Notes / History
  -----
  -----
  -- Control variables used in most programs
  -----
  DECLARE v_msgtext    varchar(255); -- Text for audit_trail
  DECLARE v_sql       varchar(255); -- Text for sql statements
  DECLARE v_set       integer;      -- commit set
  DECLARE v_analyze_flag integer;   -- analyze flag
  DECLARE v_step      integer;      -- return code
  DECLARE v_update_count integer;   -- no of records updated
  DECLARE v_insert_count integer;   -- no of records inserted
  DECLARE v_count     integer;      -- General counter
  DECLARE v_sql_code  integer;      -- SQL Error Code for Audit Trail
  DECLARE v_sql_error varchar(255); -- SQL Error Code for Audit Trail as varchar
  -----
  -- Variables
  -----
  DECLARE v_dss_update_time timestamp; -- Used for date insert
  DECLARE v_order_msg      varchar(256);
  DECLARE v_order_status  integer;
  DECLARE v_retcode       varchar(256);
  DECLARE v_invoice_msg   varchar(256);
  DECLARE v_invoice_status integer;
  DECLARE v_status        integer;
  -----
  -- Exceptions
  -----
  DECLARE EXIT HANDLER
  FOR SQLEXCEPTION
  BEGIN
    SET v_sql_code = SQLCODE;
    SELECT ErrorText
    INTO   :v_sql_error
    FROM   dbc.ErrorMessage
    WHERE  ErrorCode = :v_sql_code;
    SET v_msgtext = 'Unhandled Exception in update_model_order_header. ' ||
      ' Step ' || CAST(v_step AS VARCHAR(64)) ||
      ' SQL Error Code: ' || CAST(v_sql_code AS VARCHAR(10)) || ' - ' || v_sql_error;
    SET p_return_msg = v_msgtext;
    CALL [METABASE].WsWrkAudit('F', :p_job_name, :p_task_name, :p_sequence
      :v_msgtext, :v_sql_code, :v_sql_error, :p_task_id, :p_job_id);
    SET p_status = -3;
  END;
```

```

-----
-- Main
-----
SET v_step = 100;
SET v_dss_update_time = CURRENT_TIMESTAMP;
SET v_update_count = 0;
SET v_insert_count = 0;

-----
-- Update model order header table with info from the orders
-- stage table. Report the results to the audit trail.
-----
CALL [METABASE].update_model_order_header_o(:p_sequence, :p_job_name, :p_task_name
, :p_job_id, :p_task_id, :v_order_msg, :v_order_status);

SET v_step = 200;

IF ( v_order_status = 1 ) THEN
  SET v_retcode = 'S';
ELSEIF ( v_order_status = -1 ) THEN
  SET v_retcode = 'W';
ELSEIF ( v_order_status = -2 ) THEN
  SET v_retcode = 'E';
ELSEIF ( v_order_status = -3 ) THEN
  SET v_retcode = 'F';
END IF;

SET v_step = 300;

CALL [METABASE].WsWrkAudit(:v_retcode, :p_job_name, :p_task_name, :p_sequence
, :v_order_msg, NULL, NULL, :p_task_id, :p_job_id);

SET v_step = 400;

-----
-- If the previous update (orders) was successful, then
-- update model order header table with info from the invoices
-- stage table. Report the results to the audit trail.
-----
IF ( v_order_status > -2 ) THEN

  CALL [METABASE].update_model_order_header_o(:p_sequence, :p_job_name, :p_task_name
, :p_job_id, :p_task_id, :v_invoice_msg, :v_invoice_status);

  SET v_step = 500;

  IF ( v_invoice_status = 1 ) THEN
    SET v_retcode = 'S';
  ELSEIF ( v_invoice_status = -1 ) THEN
    SET v_retcode = 'W';
  ELSEIF ( v_invoice_status = -2 ) THEN
    SET v_retcode = 'E';
  ELSEIF ( v_invoice_status = -3 ) THEN
    SET v_retcode = 'F';
  END IF;

  SET v_step = 600;

  CALL [METABASE].WsWrkAudit(:v_retcode, :p_job_name, :p_task_name, :p_sequence
, :v_invoice_msg, NULL, NULL, :p_task_id, :p_job_id);

  SET v_step = 700;
ELSE
  SET v_invoice_status = 0;

```

```

END IF;

-----
-- All Done report the results
-----
SET v_step = 800;

IF ( v_invoice_status < v_order_status ) THEN
  SET v_status = v_invoice_status;
ELSE
  SET v_status = v_order_status;
END IF;

SET v_msgtext = 'model_order_header update ';

IF ( v_status = 1 ) THEN
  SET v_msgtext = v_msgtext || 'completed.';
ELSEIF ( v_status = -1 ) THEN
  SET v_msgtext = v_msgtext || 'completed with warnings.';
ELSE
  SET v_msgtext = v_msgtext || 'failed.';
END IF;

SET p_status = v_status;
SET p_return_msg = v_msgtext;

END;

```

Procedure Placeholders

Procedure placeholders can help in moving procedures between environments without the necessity of regenerating those same procedures. In <PRODUCT, the purpose of these placeholders is to automatically substitute the corresponding strings, which is needed for a specific environment.

The following procedure placeholders described below can be found in the update_xxxx_xxxx procedure.

[TABLEOWNER] is used as a placeholder to replace the schema name defined in the connection or target.

For targets, the **[TABLEOWNER]** placeholder is derived from the Target Location Database/Schema in the connection. The target can be changed in the table's Storage screen, on the Target drop-down list. For more information about Target Location Database/Schema in connections and table's storage screens, see *Connection to the Data Warehouse* (see "*Database - Data Warehouse/Metadata Repository*" on page 147) and *Storage* (on page 202).

When moving tables between environments, the **[TABLEOWNER]** placeholder is determined by the individual connection of the target environment.

Example

During the compilation process of the procedure, [TABLEOWNER.tablename] is replaced with PRODUSER.tablename, if the table owner is PRODUSER in the destination environment.

```

-----
-- Main
-----
SET v_step = 100;
SET v_dss_update_time = CURRENT_TIMESTAMP;
SET v_update_count = 0;
SET v_insert_count = 0;

-----
-- Delete all existing data from the table
-----
DELETE FROM [TABLEOWNER].[stage_customer] ALL;

SET v_step = 200;

-----
-- Insert input records into stage_customer
-----
INSERT INTO [TABLEOWNER].[stage_customer]
(
    customer_code,
    customer_name,
    customer_legal_name,
    territory_id,

```

**WhereScape RED Tip: dim_date**

The TABLEOWNER placeholder is especially useful in update procedures when the related table is moved to a different schema or environment. For example, when moving dim_date to other schemas, [TABLEOWNER] will be replaced with the schema of the table when the procedure is compiled.

To do this, prefix the table name in the procedure to [TABLEOWNER].table_name (e.g. [TABLEOWNER].dim_date). Then it is only necessary to do a recompile instead of rebuilding or regenerating the procedure.

[METABASE] is used as a placeholder for the Teradata database metadata repository to enable the deployment between environments without regenerating the procedures.

```
CREATE PROCEDURE [METABASE].update_edw_customer
(
  IN  p_sequence      integer,
  IN  p_job_name      varchar(256),
  IN  p_task_name     varchar(256),
  IN  p_job_id        integer,
  IN  p_task_id       integer,
  OUT p_return_msg    varchar(256),
  OUT p_status        integer
)
BEGIN

-----
-- DBMS Name       : Teradata
-- Script Name     : update_edw_customer
-- Description     : Update the EDW 3NF table edw_customer
-- Generated by    : WhereScape RED Version 6.8.4.3 (build 150908-203700 pre-release)
-- Generated for   : RED Documentation
-- Generated on    : Thursday, September 17, 2015 at 12:37:07
-- Author         : WhereScape Documentation
-----
-- Notes / History
--
-----
-- Control variables used in most programs
-----
```

PROCEDURE EDITING

WhereScape RED includes a procedure editor which allows the maintenance of the various procedures, functions and packages within the data warehouse. The editor is invoked by double-clicking on a procedure name in the left pane or by right-clicking on the procedure name and selecting **Edit the Procedure**.

A procedure can be compiled by selecting the **Compile>Compile** menu option. See the section on compiling procedures for more information.

This section will discuss some of the features of the procedure editor.

In the following sections, reference is made to a selected block of text. A selected block of text is a normal Windows selection where the text in question is highlighted. Normally achieved by holding down the left mouse button and moving the cursor.

Indenting code

The **tab** character inserts two spaces into the text. A **shift/tab** removes two spaces.

Cut, Copy, Paste and Delete

The normal Windows cut, copy, paste and delete functions are available either through the toolbar or via the right-click context menu.

Indenting a block of text

A selected block of text can be indented by four spaces by depressing the **tab** character. Each tab indents by a further two spaces. A **shift/tab** will remove two spaces from the front of each line in the selected block.

Commenting out a block of text

A selected block of text can be turned into a comment by using the **Comment** option in the right-click pop-up menu.

The editor will place two dashes '--' at the front of each line in the selected block.

In the same way a block of text can be un commented by choosing the **Uncomment** option.

Note: Only lines that start with two dashes in the left most column can be uncommented.

Inserting Steps

The right-click pop-up menu has an option to **insert step**. This option inserts a code line of the format 'SET v_step = 1000;'. Each subsequent insert adds 100 to the step value. The Alt/Z key can also be used to insert a step line.

The v_step variable is used in the event of an unhandled exception. In such a case, the current v_step value is reported, and it may be possible to ascertain the code that failed based on the value of this v_step variable.

Note: If a step is inserted via this method then the step numbering is automatically reset for all steps numbering from 100 in increments of 100 when the procedure is compiled or saved.

Inserting Audit and Detail Messages

The Alt/A key can be used to insert an extra audit message to be written to the audit log while the procedure is running. The default message can be changed as appropriate. The inserted code is:

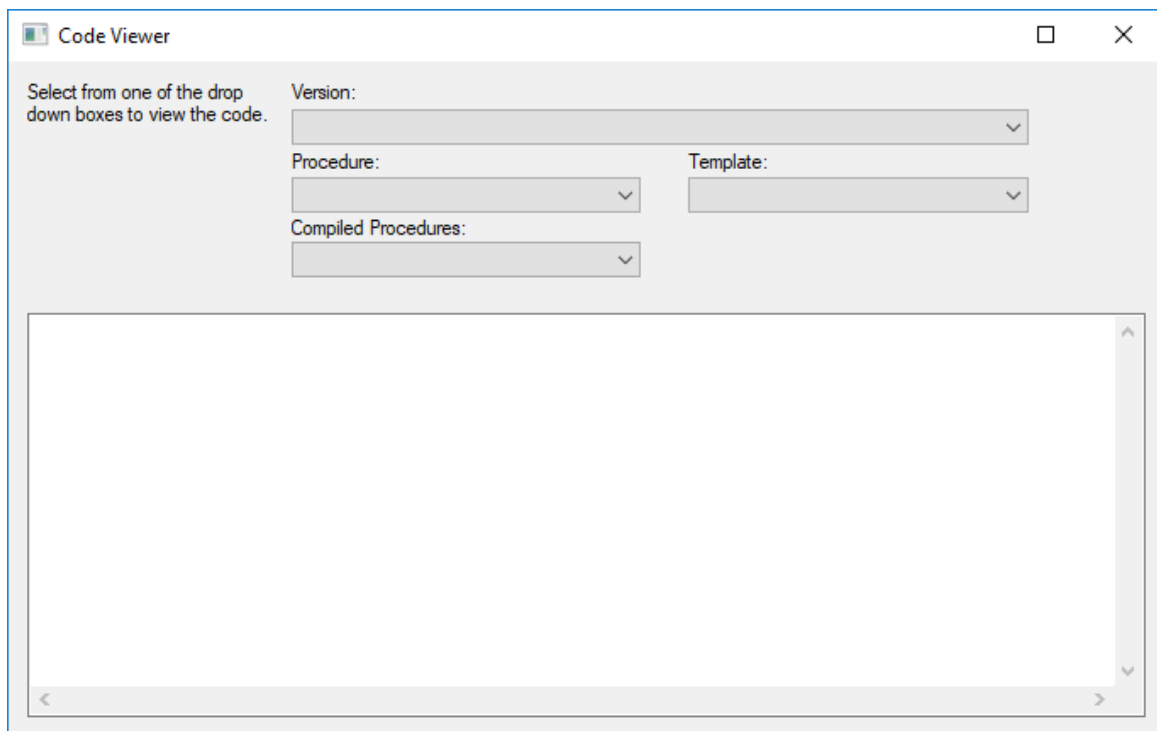
```
SET v_msgtext = ' Step ' || CAST(v_step AS VARCHAR(64))
CALL [METABASE].WsWrkAudit('I', :p_job_name, :p_task_name, :p_sequence
, :v_msgtext, '', '', :p_task_id, :p_job_id);
```

The Alt/D key can be used to insert an extra detail (or error) message to be written to the detail log while the procedure is running. The default message can be changed as appropriate. The inserted code is:

```
SET v_msgtext = ' Step ' || CAST(v_step AS VARCHAR(64))
CALL [METABASE].WsWrkError('I', :p_job_name, :p_task_name, :p_sequence
, :v_msgtext, '', '', :p_task_id, :p_job_id, NULL);
```

Viewing other procedural code

During the editing process it is possible to pop up a window containing other procedural code. This window enables cut and paste operations. In such a way, it can be used as a work area or as a source of code. Select the **Tools>View Procedure or Template** menu option to bring the viewer window up. A dialog appears:



A number of drop-down lists can be chosen from. Once an item is selected, the viewer loads the code and moves to the right side of the edit window. The various options are:

- Version: A previously saved version of this procedure. The description and time are shown.
- Procedure: Some other procedure stored within the WhereScape metadata.
- Template: A template procedure as defined in the WhereScape metadata.
- Compiled procedure: One of the currently compiled procedures existing in the database.

Once an item is chosen, the viewer appears on the right side of the edit window. The viewer or the main edit window can receive the focus by clicking the mouse within the appropriate window. Code can be cut from one window and pasted into the other. Changes made in the viewer window cannot be saved. An example of an active view window is as follows:

The screenshot shows a development environment window titled 'doc6871 Development (Licensed to RED Documentation)'. It contains two panes. The left pane shows the source code for a procedure named 'update_dim_customer'. The right pane, titled 'Code Viewer', shows the same code but with a metadata section below it.

```

CREATE PROCEDURE [METABASE].update_dim_customer
(
  IN p_sequence      integer,
  IN p_job_name      varchar(256),
  IN p_task_name     varchar(256),
  IN p_job_id        integer,
  IN p_task_id       integer,
  OUT p_return_msg   varchar(256),
  OUT p_status       integer
)
BEGIN

  -----
  -- DBMS Name       : Teradata
  -- Script Name     : update_dim_customer
  -- Description     : Update the Dimension tabl
  -- Generated by    : WhereScape RED Version 8.
  -- Generated for   : RED Documentation
  -- Generated on    : Thursday, November 23, 20
  -- Author         : WhereScape Documentation
  -----

  -- Notes / History
  --

  -----
  -- Control variables used in most programs
  -----

  DECLARE v_msgtext  varchar(255); -- Te
  DECLARE v_sql      varchar(255); -- Te
  DECLARE v_set      integer;      -- co
  DECLARE v_analyze_flag integer;  -- an
  DECLARE v_step     integer;      -- re
  DECLARE v_insert_count integer;  -- no
  DECLARE v_update_count integer;  -- no
  DECLARE v_count    integer;      -- Ge

```

The 'Code Viewer' window shows the same code as the left pane, followed by a metadata section:

```

-----
-- DBMS Name       : Teradata
-- Script Name     : update_dim_customer
-- Description     : Update the Dimension table dim_customer
-- Generated by    : WhereScape RED Version 8.0.2.0 (build 171
-- Generated for   : RED Documentation
-- Generated on    : Thursday, November 23, 2017 at 12:21:57
-- Author         : WhereScape Documentation
-----

-- Notes / History
--

-----
-- Control variables used in most programs
-----

DECLARE v_msgtext  varchar(255); -- Text for audit_tra
DECLARE v_sql      varchar(255); -- Text for SQL state
DECLARE v_set      integer;      -- commit set
DECLARE v_analyze_flag integer;  -- analyze flag
DECLARE v_step     integer;      -- return code
DECLARE v_insert_count integer;  -- no of records inse
DECLARE v_update_count integer;  -- no of records upda
DECLARE v_count    integer;      -- General counter

```

NOTE: Editing, deleting or compiling Locked for Edit or opened procedures/scripts

Procedures or scripts cannot be deleted if they are Locked for Edit by any user, checked out by another user or if there is another object that has the same associated procedures or scripts. Saving or Compiling in the procedure or script edit window cannot be performed if the procedures or scripts become Locked for Edit by other users after the edit window was opened. Procedures or scripts cannot be deleted or modified after the edit window has been opened, unless the Edit Lock has been released. Edit Locks can be released by any user in the Script or Procedure Properties screen.

To prevent updates, deletes and modifications to certain procedures or scripts, it is best to use the Check Out functionality instead. For more information about this functionality, please see section *Check Outs and Check Ins*.

PROCEDURE LOADING AND SAVING

Procedures are normally stored in the WhereScape RED metadata tables. When a procedure is opened within RED then the data is retrieved from the meta tables. Likewise when the procedure is saved it overwrites the existing data in the meta tables.

When a procedure is compiled it is also written to the database system tables. In Teradata, a compiled procedure can be viewed via the Teradata Administrator by doing a Show Definition of the procedure located in the meta repository database.

Loading data

As mentioned above, when a procedure is opened for editing the information stored in the metadata is loaded. Additional text can be loaded into the procedure edit window by selecting the **File>Insert from File** menu option which allows the browsing and inserting of a PC based file. Also, paste buffer data can be inserted in the normal manner. In the previous 'Editing' section, the viewer window was discussed. This window can also be a source of data via cut and paste.

Saving the Procedure

The default **File>Save Procedure** menu option overwrites the existing procedure in the metadata. In addition a procedure can be saved to a Windows disk by selecting the **File>Write procedure to disk** menu option. All procedures can be written to disk from the main Builder menu option **Backup>Save Procedures to Disk**. This option allows the selection of a directory. The procedures are then written individually into that directory.

Versions

Multiple versions of a procedure can be stored. Once a version is created that version may be read but may not be updated. Only the current procedure can be edited. There are a number of ways to create a version of a procedure. These are:

- 1 By setting **Auto-Version before Procedure Compile** under **Tools>Options>Metadata Versioning**. If set, a new version of a procedure is created whenever the procedure is compiled.
- 2 The Procedure Editor menu option **File>Save and Version** saves a procedure and creates a version at the same time.
- 3 Selecting the **Version Control>New Version** menu option from the context menu when right clicking a procedure from right pane in the main Builder window.
- 4 By selecting the **Tools>Version Objects>Version All** menu option.

When a version is created via method (2) or (3) above, the following screen appears to enable the definition of the version. If an auto version is created, then the person creating the version is recorded along with the reason for the version. (e.g. Version on compile, Version on procedure delete).

Version Definition

Create version for Procedure update_ds_customer

Version Name or Short Description:

Detailed Description:

Retain Until: Tuesday , December 7, 2027

OK Cancel

- The version name/description appears, when the versions are subsequently browsed.
- The Retain until date is set ten years in the future by default.
- The automated deletion of versions is not supported at this stage.

PROCEDURE COMPARISONS

A procedure can be compared to either an earlier version, or to the currently running code as compiled/stored in the database. The menu option **Tools>Compare to Compiled Source** enables the comparison of the procedure being edited with the code currently compiled and running in the database. If a viewer window is open (see the procedure editing section), then the **Tools>Compare to Viewer** menu option compares the contents of the viewer window with the current code. Therefore, to compare against an older version, we first load the viewer window with the older version and perform a **Compare to Viewer**.

The comparison will highlight the differences, as shown in the example below:

```
-->>-- set v_step = 200;
| SELECT COUNT(*) --<<--
  INTO   :v_count --<<--
  FROM   [model_customer] model_customer --<<--
  WHERE  dss_update_time >= CURRENT_TIMESTAMP - INTERVAL '7' DAY; --<<--
```

In this example the line `SET v_step = 100;` has been removed from the current code in the edit window and the remaining three lines have been inserted.

Once the comparison has been completed you can either remove the compare comments or accept the compare changes. The menu option **Tools>Remove Compare Comments** removes the added blue comments and code. The menu option **Tools>Accept Compare Changes** implements the changes highlighted. For the above example the line `'SET v_step = 100;'` would be added and the following three lines deleted.

PROCEDURE COMPILATION

From within the procedure editor a procedure can be compiled by selecting the menu option **Compile>Compile** or by clicking the **Compile** icon. If the procedure compiles successfully a dialog box appears notifying of a successful compile. If the compile fails, then error message comments is inserted into the procedure code. In the following example, the error messages are in red and begin with --E--.

```

-----
-- Exceptions
-----
DECLARE EXIT HANDLER
FOR SQLEXCEPTION
BEGIN
    SET v_sql_code = SQLCODE;
    SELECT SUBSTR(ErrorText,1,255)
    INTO   :v_sql_error
    FROM   dbc.ErrorMsgs
    WHERE  ErrorCode = :v_sql_code;
    SET v_msgtext = SUBSTR('Unhandled Exception in model_customer: ' ||
    ' Step ' || CAST(v_step AS VARCHAR(64)) ||
    ' SQL Error Code: ' || CAST(v_sql_code AS VARCHAR(10)) || ' - ' || v_sql_error.1,255);
    SET p_return_msg = v_msgtext;
    CALL [METABASE].WsWrkAudit('F', :p_job_name, :p_task_name, :p_sequence
    , :v_msgtext, :v_sql_code, :v_sql_error, :p_task_id, :p_job_id);
    SET p_status = -3;
END;

-----
-- Main
-----
SET v_step = 100;
SET v_update_count = 0;
SET v_insert_count = 0;

SELECT COUNT(*)
INTO   :v_count
FROM   [model_customer] model_customer
|--line 72-- SPL5000:W(L72), E(5628):Column dss_update_date not found in WslWarehouse.model_customer.
WHERE  dss_update_date >= CURRENT_TIMESTAMP - INTERVAL '7' DAY;

```

Error comments is inserted at each error point. *A compile deletes any previous error comments.* Error comments can also be removed through the menu option **Compile>Delete Error messages**.

Note: In some instances, the error comments may not be positioned on the correct line. This can occur as the result of one or more procedure lines being wrapped. Therefore, ensure the procedure editor window is maximized when dealing with compile errors.

PROCEDURE RUNNING

Only procedures that conform to the WhereScape scheduler syntax can be executed from within the procedure editor. Select the **Execute>Execute** menu option or click the **Execute the procedure** icon to run the procedure. A procedure must have been compiled in order to run.

The results of the procedure is displayed in a dialog box. The result code and result messages is displayed, as well as any additional messages.

PROCEDURE SYNTAX

The procedures managed by the WhereScape scheduler require the following standards. If a function or procedure is being developed that is not called directly by the scheduler then it does not need to conform with this standard. If however such a procedure or function wants to log messages to the audit or error logs then it will need the input parameters included in its parameter list.

Parameters

The procedure must have the following parameters in the following order:

Parameter name	Input or Output	Data Type
p_sequence	Input	Integer
p_job_name	Input	Varchar(256)
p_task_name	Input	Varchar(256)
p_job_id	Input	Integer
p_task_id	Input	Integer
p_return_msg	Output	Varchar(256)
p_status	Output	Integer

The input parameters are passed to the procedure by the scheduler. If the procedure is called outside the scheduler, then the normal practice is to pass zero (0) in the sequence, job_id and task_id. A description of the run can be passed in the job name and the task name is typically the name of the procedure.

The output parameters must be populated by the procedure on completion. The return_msg can be any string up to 256 characters long that describes the result of the procedure's execution. The status must be one of the following values:

Status	Meaning	Description
1	Success	Normal completion
-1	Warning	Completion with warnings
-2	Error	Hold subsequent tasks dependent on this task
-3	Fatal Error	Hold all subsequent tasks

Note: Multiple SQL statements can be separated using the "end of statement" indicator. This is <EOS> by default but can be configured in **Tools>Options>Code Generation>General**.

PROCEDURE PROPERTIES

The properties screen for procedures and scripts is the same. A procedure can be renamed by changing the name field.

If a procedure is renamed, then it is also necessary to change the procedure name within the actual code. The purpose and owner fields are purely informational.

The screenshot shows the 'Procedure update_dim_product' dialog box with the following details:

- Name:** update_dim_product
- Type:** Procedure
- Purpose:** update_dim_product
- Owner:** auto created
- Delete Lock:**
- Last Update By:** (empty)
- Default Connect:** DataWarehouse
- Edit Lock:**
 - Locked For Edit By:** WhereScape Documentation
 - Edit Lock Reason or Last Update:** new procedure
 - Release Edit Lock** button
- Timestamps:**
 - Created:** 2017-03-27 02:53:52
 - Last Update:** (empty)
 - Compiled:** 2017-03-27 02:53:57

In the example above, the **Delete Lock** check box is not selected. Selecting this check box prevents the procedure from being deleted through the **Delete** menu option. It also prevents the procedure from being overwritten, if a new procedure generation is requested.


Also in the example above, the procedure is currently being edited and is shown as being Locked for Edit by "WhereScape Documentation". If procedures or scripts have already been opened for editing, they can only subsequently be opened for viewing.

Edit locks and delete operations for procedures and scripts, as well as the regeneration and drop of procedures are not permitted if the object is currently Locked For Edit by another user.

The **Release Edit Lock** button to the right of the edit lock message, clears the edit lock. If WhereScape RED, the database or the PC crashes when a procedure is open, then the check out needs to be cleared through this screen.

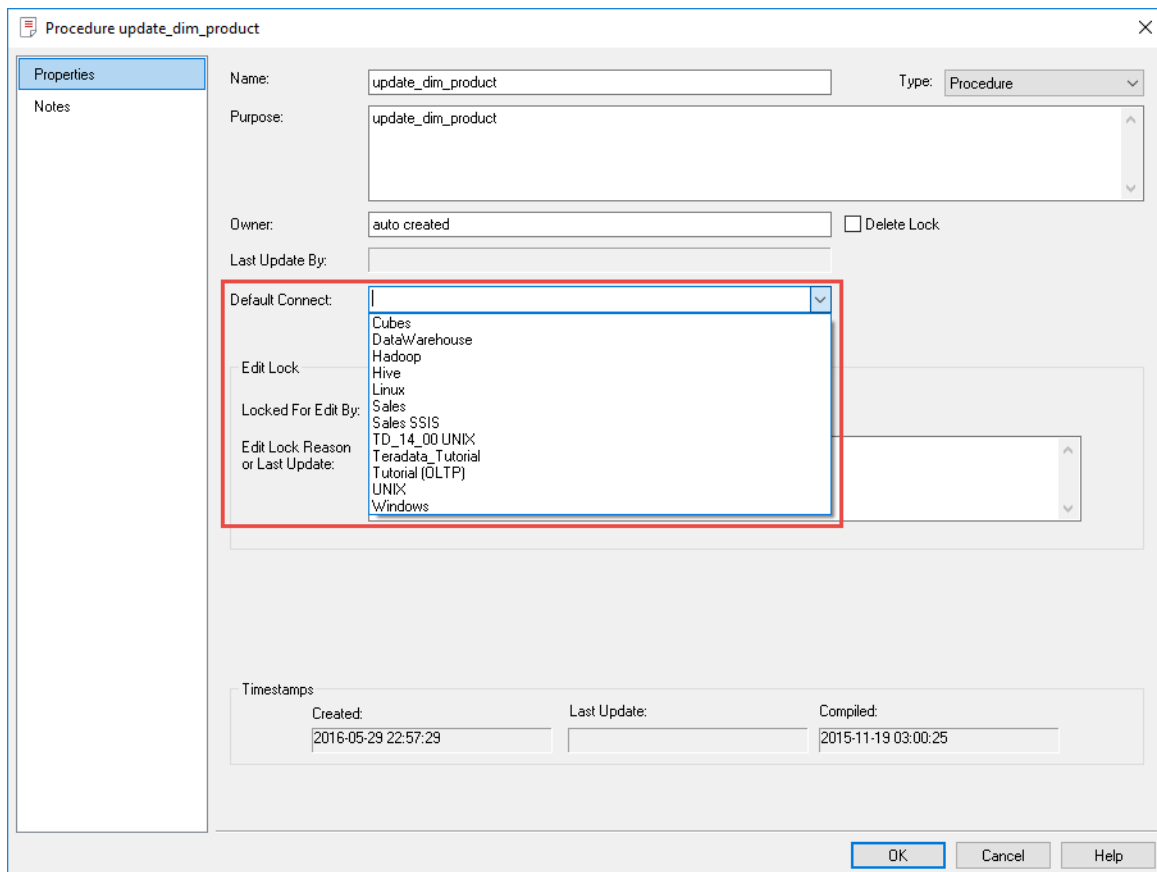
The **Edit Lock Reason** is for information only, and can be used as another comment field if desired.

The **Type** drop list enables the selection of **Block**, **Function** or **Procedure**:



Selecting a type of **Block** enables you to execute a SQL block against another connection.

An additional field on the Properties screen—**Default Connect** then enables you to select the connection against which the SQL block will be executed.



MACROS

WhereScape RED can also retrofit, run, schedule and generate Teradata macros.

Template **macros** can be generated for moving data into simple staging tables without surrogate keys. A macro is generated by selecting the ***** Create New Macro ***** option from a drop-down box that is used to display the update procedure in a staging table's properties.

A utility is available from WhereScape to mass retrofit large numbers of existing Teradata macros into the WhereScape RED metadata.

BTEQ SCRIPTS

WhereScape RED can also retrofit, run, schedule and generate Teradata BTEQ scripts.

Template **bteq scripts** can be generated for moving data into simple staging tables without surrogate keys. A BTEQ script is generated by selecting the ***** Create New Bteq ***** option from a drop-down box that is used to display the update procedure in a staging table's properties.

SCRIPT GENERATION

WhereScape RED generates template scripts to assist in the loading of textual files from Windows. These scripts are generated when a Windows file is dragged into a Load table target and one of the two 'Script based' file load options is chosen. Typically, script loads are used when some complexity prevents the use of the standard file based load. In such a case, the script needs modification.

Note: Windows script supports both DOS Batch and *PowerShell scripts* (see "*Windows PowerShell Scripts*" on page 653).

SCRIPT GENERATION (WINDOWS/TERADATA)

A sample Windows script for Teradata is as follows. The key components of the script are described below:

```

@echo off
setlocal enabledelayedexpansion
setlocal enableextensions
REM *****
REM ***** DBMS Name      : Teradata
REM ***** Description   : File Load/Movement for C:\temp\forecast.txt into load_forecast
REM ***** Generated by  : WhereScape RED
REM ***** Generated for  : WhereScape Limited
REM ***** Author        : WhereScape Developer
REM *****
REM ***** Notes / History
REM *****
REM *****
REM ***** NOTE: The following environment variables will be set
REM ***** WSL_LOAD_TABLE = the load table name
REM ***** WSL_LOAD_DB    = the load table names database
REM ***** WSL_TEMP_DB    = the load table names database for temporary tables
REM ***** WSL_SEQUENCE  = a unique sequence number for the scheduler
REM ***** WSL_WORKDIR    = the work directory defined in the connection
REM ***** WSL_SCHEDLOG   = the directory and file name of the scheduler
REM ***** log file which is used to defined the parameter batch file
REM ***** when this is scheduled
REM ***** WSL_SERVER     = the server defined in the connection
REM ***** WSL_DATABASE   = the database defined in the connection
REM ***** WSL_METABASE   = the database storing the meta data
REM ***** WSL_USER       = the dss user defined in the connection
REM ***** WSL_PWD        = the dss password defined in the connection
REM ***** WSL_REMOTEHOSTn = the nth remote server for sending files to
REM ***** WSL_REMOTEUSERn = the username for the nth remote server
REM ***** WSL_REMOTEPWdn = the password for the nth remote server
REM ***** WSL_PARAMnnn  = Any parameters that start with the load
REM ***** table name. For example: A table called load_abc has a
REM ***** parameter called load_abc_server defined. In this case
REM ***** a variable called WSL_PARAM_SERVER will be created
REM *****
REM *****
if defined WSL_LOAD_TABLE (SET LOAD_TABLE=%WSL_LOAD_TABLE%) else SET LOAD_TABLE=load_forecast
if defined WSL_LOAD_DB (SET LOAD_DB=%WSL_LOAD_DB%) else SET LOAD_DB=WslWarehouse
if defined WSL_TEMP_DB (SET TEMP_DB=%WSL_TEMP_DB%) else SET TEMP_DB=WslWarehouse
if defined WSL_SERVER (SET SERVER=%WSL_SERVER%) else SET SERVER=
if defined WSL_DATABASE (SET DATABASE=%WSL_DATABASE%) else SET DATABASE=
if defined WSL_METABASE (SET METABASE=%WSL_METABASE%) else SET METABASE=WslWarehouse.
if defined WSL_USER (SET USER=%WSL_USER%) else SET USER=WslWarehouse
if defined WSL_PWD (SET PWD=%WSL_PWD%) else SET PWD=wsl
if defined WSL_WORKDIR (SET WORKDIR=%WSL_WORKDIR%) else goto LABEL_FILE_DEFAULT
if not defined WSL_SCHEDLOG (SET PARAMBAT="%WORKDIR%\parameters.bat") else (SET PARAMBAT="%WSL_SCHEDLOG:log=bat%")
if defined WSL_SEQUENCE (SET SEQUENCE=%WSL_SEQUENCE%) else goto LABEL_FILE_DEFAULT
SET FILECTL=%WORKDIR%\wsl_%LOAD_TABLE%%SEQUENCE%.ctl
SET FILEAUD=%WORKDIR%\wsl_%LOAD_TABLE%%SEQUENCE%.aud
goto LABEL_WAIT
:LABEL_FILE_DEFAULT
SET FILECTL=wsl_load_forecast100002.ct1
SET FILEAUD=wsl_load_forecast100002.aud
SET SEQUENCE=100002
SET PARAMBAT=parameters.bat
:LABEL_WAIT
ECHO Starting > %FILEAUD%
IF EXIST %PARAMBAT% (CALL %PARAMBAT%) ELSE (ECHO Parameters batch file does not exist >> %FILEAUD%)
REM
REM If a wild card is used in the file name then %FILE_NAME%
REM in the control file must be enclosed in single quotes,
REM otherwise it must not be enclosed in quotes.
REM
SET LOAD_FILE=C:\temp\forecast.txt
REM
REM Remove Existing Parameters
echo DELETE FROM %METABASE%dss_parameter WHERE dss_parameter_name like '____' | bteq .Logon %DATABASE%/%USER%,%PWD% >> %FILEAUD%

```

```

REM
REM ***** W A I T       T I M E R *****
REM *****
REM If a wait time specified then loop looking for the file
REM or trigger file to arrive until the wait time expires.
SET /A WAITSECS=0
:LABEL_WAITLOOP
if %WAITSECS% LEQ 0 goto LABEL_FILECHECK
REM see if the file exists and if so skip the wait check
if exist "%LOAD_FILE%" goto LABEL_FILECHECK
REM Call WslSleep to wait for 30 seconds. It will return remaining seconds
"C:\Program Files\WhereScape\WslSleep" %WAITSECS% 30
SET /A WAITSECS=%errorlevel%
goto LABEL_WAITLOOP
REM
REM Finished our wait loop. See if we have the file
REM ***** F I L E   C H E C K *****
REM *****
:LABEL_FILECHECK
if exist "%LOAD_FILE%" goto LABEL_FOUND
echo -1
echo File "%LOAD_FILE%" was not found.
exit
REM
REM ***** L O A D       T H E   D A T A *****
REM *****
:LABEL_FOUND
REM
echo Load of "%LOAD_FILE%" >> %FILEAUD%
SET /A RESULT_CODE=1
SET RESULT_MSG=Load Completed Normally.
:LABEL_LOAD
REM
REM ***** L O A D I N G *****
REM *****
SET FILE_NAME=NO MORE FILES
for %%A in ("%LOAD_FILE%") do if "!FILE_NAME!"=="NO_MORE_FILES" SET FILE_NAME=%%A
if "!FILE_NAME!"=="NO_MORE_FILES" goto LABEL_EXIT
SET FIRST_CHAR=%FILE_NAME:~0,1%
if "%FIRST_CHAR%=" SET FILE_NAME=%FILE_NAME:~1,-1%
FOR /F "usebackq delims=" %%i IN ("%FILE_NAME%") do SET FILE_SHORT="%%~nx"
echo Loading %FILE_NAME% >> %FILEAUD%
echo LOGON %DATABASE%\%USER%\%PWD% >> %FILECTL%
echo DROP TABLE %TEMP_DB%.load_forecast_e1; >> %FILECTL%
echo DROP TABLE %TEMP_DB%.load_forecast_e2; >> %FILECTL%
echo SET RECORD VARTXT " " >> %FILECTL%
echo RECORD 2; >> %FILECTL%
echo DEFINE >> %FILECTL%
echo product_code (varchar(50)) >> %FILECTL%
echo , customer_code (varchar(50)) >> %FILECTL%
echo , forecast_quantity (varchar(50)) >> %FILECTL%
echo , forecast_sales_value (varchar(50)) >> %FILECTL%
echo , forecast_date (varchar(50)) >> %FILECTL%
echo file=%FILE_NAME%; >> %FILECTL%
echo SHOW; >> %FILECTL%
echo BEGIN LOADING %LOAD_DB%.load_forecast ERRORFILES %TEMP_DB%.load_forecast_e1, %TEMP_DB%.load_forecast_e2; >> %FILECTL%
echo INSERT INTO %LOAD_DB%.load_forecast >> %FILECTL%
echo ( product_code >> %FILECTL%
echo , customer_code >> %FILECTL%
echo , forecast_quantity >> %FILECTL%
echo , forecast_sales_value >> %FILECTL%
echo , forecast_date >> %FILECTL%
echo ) >> %FILECTL%

echo END LOADING; >> %FILECTL%
echo LOGOFF; >> %FILECTL%
fastload < %FILECTL% >> %FILEAUD%
SET ERRLEV=%errorlevel%
IF %ERRLEV% EQU 0 GOTO LABEL_OKAY
IF %ERRLEV% LEQ 4 GOTO LABEL_WARNING
:LABEL_FAIL
echo -3
IF %ERRLEV% EQU 8 echo Load Failed. A user error occurred in the loader.
IF %ERRLEV% EQU 12 echo Load Failed. A fatal error occurred in the loader.
IF %ERRLEV% EQU 16 echo Load Failed. No message destination available from the loader.
type %FILEAUD% >&2
:ERR_EXIT
exit
:LABEL_WARNING
echo -1
echo Load Completed Normally, with warnings from the loader.
:LABEL_OKAY
REM
REM ***** R E N A M E *****
REM *****
:LABEL_RENAME
REM
REM ***** N E X T   F I L E *****
REM *****
REM If multiple files are being processed via a wildcard
REM uncomment the goto LABEL_LOAD statement
REM WARNING: Do not loop back unless the file name contains a wildcard
REM WARNING: You must rename the file that has been loaded if
REM looping back otherwise the script will loop for
REM ever loading the same file.
:LABEL_NEXTFILE
REM goto LABEL_LOAD
REM
REM ***** E X I T *****
REM *****
:LABEL_EXIT
echo %RESULT_CODE%
echo %RESULT_MSG%
exit

```

**WhereScape RED TIP: Parameters**

Parameters can also be added to scripts to facilitate deployment processes or environment changes without the need to regenerate scripts. They can be added to scripts of Load and Export tables.

For example: add `$P<ParameterName>$` to the script where `P` is the parameter indicator as shown below.

Before adding a Parameter to the script, create the desired parameter(s) in **Tools>Parameters>Add Parameter**.

```
SET LOAD_FILE="$PSourcePath$\forecast.txt"
SET TRIG_FILE="$PSourcePath$\forecast.trg"
```

The script makes use of a number of environmental variables. These variables are acquired from both the table and connection properties. These variables are established in the environment by either RED or the scheduler. If the script is to be executed outside of RED or scheduler control then these variables will need to be assigned.

The first section of the script defines the variables. The second section provides a timed wait for the load file to arrive. By default, the `WAITSECS` variable is set to zero, so that no wait occurs. This can be set to a number of seconds that the script is to wait for the file to arrive.

```
REM
REM *****
REM ***** W A I T   T I M E R *****
REM *****
REM If a wait time specified then loop looking for the file
REM or trigger file to arrive until the wait time expires.
SET /A WAITSECS=0
:LABEL_WAITLOOP
if %WAITSECS% LEQ 0 goto LABEL_FILECHECK
REM see if the file exists and if so skip the wait check
if exist "%TRIG_FILE%" goto LABEL_FILECHECK
REM Call WslSleep to wait for 30 seconds. It will return remaining seconds
"C:\Program Files\WhereScape\WslSleep" %WAITSECS% 30
SET /A WAITSECS=%errorlevel%
goto LABEL_WAITLOOP
REM
REM Finished our wait loop. See if we have the file
REM *****
REM ***** F I L E   C H E C K *****
REM *****
:LABEL_FILECHECK
if exist "%TRIG_FILE%" goto LABEL_FOUND
echo -1
echo File %TRIG_FILE% was not found.
exit
REM
```

Once the wait has completed, either through a time expiry or through the location of the file, we check that the file is present, and if not found report back a warning. This warning can be changed to an error by changing the first echo statement to "-2". See the syntax section for more information.

When a trigger file is specified the script looks for a trigger file, and will exit with the specified status if the file is not found. The following code is included if a trigger file is present.

```

REM
REM *****
REM ***** T R I G G E R *****
REM *****
REM
REM First clear any existing parameters
REM
.LABEL_TRIGGER
SET TRIG_NAME=NO_MORE_FILES
for %%A in ("%TRIG_FILE%") do if "!TRIG_NAME!"=="NO_MORE_FILES" SET TRIG_NAME=%%A
if "!TRIG_NAME!"=="NO_MORE_FILES" goto LABEL_TRIG_EXIT
FOR /F "usebackq delims==" %%I IN ("%TRIG_NAME%") do SET TRIG_SHORT=%%I
REM
REM Get the first four and all remaining parameters
REM
REM /F "tokens=1,2,3,4,5,6,7,8,9* usebackq delims=|" %%a in (%TRIG_NAME%) do SET A=%%a&SET B=%%b&SET C=%%c&SET D=%%d&SET E=%%e&SET F=%%f&SET G=%%g&SET H=
REM
REM For each parameter write it to the parameter table
REM
if defined A echo call %METABASE%ParameterWrite( FORECAST_0 '%%A' 'script load') | bteq .Logon %DATABASE%\%USER%\%PWD% >> %FILEAUD%
if defined B echo call %METABASE%ParameterWrite( FORECAST_1 '%%B' 'script load') | bteq .Logon %DATABASE%\%USER%\%PWD% >> %FILEAUD%
if defined C echo call %METABASE%ParameterWrite( FORECAST_2 '%%C' 'script load') | bteq .Logon %DATABASE%\%USER%\%PWD% >> %FILEAUD%
if defined D echo call %METABASE%ParameterWrite( FORECAST_3 '%%D' 'script load') | bteq .Logon %DATABASE%\%USER%\%PWD% >> %FILEAUD%
if defined E echo call %METABASE%ParameterWrite( FORECAST_4 '%%E' 'script load') | bteq .Logon %DATABASE%\%USER%\%PWD% >> %FILEAUD%
if defined F echo call %METABASE%ParameterWrite( FORECAST_5 '%%F' 'script load') | bteq .Logon %DATABASE%\%USER%\%PWD% >> %FILEAUD%
if defined G echo call %METABASE%ParameterWrite( FORECAST_6 '%%G' 'script load') | bteq .Logon %DATABASE%\%USER%\%PWD% >> %FILEAUD%
if defined H echo call %METABASE%ParameterWrite( FORECAST_7 '%%H' 'script load') | bteq .Logon %DATABASE%\%USER%\%PWD% >> %FILEAUD%
if defined I echo call %METABASE%ParameterWrite( FORECAST_8 '%%I' 'script load') | bteq .Logon %DATABASE%\%USER%\%PWD% >> %FILEAUD%
if defined J echo call %METABASE%ParameterWrite( FORECAST_9 '%%J' 'script load') | bteq .Logon %DATABASE%\%USER%\%PWD% >> %FILEAUD%
.LABEL_TRIG_EXIT
REM

```

Such a file (trigger) contains control information about the main file to be loaded and arrives after the main file to indicate that the main file transfer has completed and that it is okay to load. This section loads the contents of the trigger file into the Parameters table, so that the table can be validated. See the section on Post Load procedures for an explanation on how trigger files are used to validate a load file.

```

REM
REM *****
REM ***** L O A D I N G *****
REM *****
REM
SET FILE_NAME=NO_MORE_FILES
for %%A in ("%LOAD_FILE%") do if "!FILE_NAME!"=="NO_MORE_FILES" SET FILE_NAME=%%A
if "!FILE_NAME!"=="NO_MORE_FILES" goto LABEL_EXIT
SET FIRST_CHAR=%FILE_NAME:~0,1%
if %FIRST_CHAR%==" " SET FILE_NAME=%FILE_NAME:~1,-1%
FOR /F "usebackq delims==" %%I IN ("%FILE_NAME%") do SET FILE_SHORT=%%I
echo Loading %FILE_NAME% >> %FILEAUD%
echo LOGON %DATABASE%\%USER%\%PWD% >> %FILECTL%
echo DROP TABLE %TEMP_DB%.load_forecast_e1; >> %FILECTL%
echo DROP TABLE %TEMP_DB%.load_forecast_e2; >> %FILECTL%
echo SET RECORD VARTEXT " "; >> %FILECTL%
echo RECORD 2; >> %FILECTL%
echo DEFINE >> %FILECTL%
echo product_code (varchar(50)) >> %FILECTL%
echo customer_code (varchar(50)) >> %FILECTL%
echo forecast_quantity (varchar(50)) >> %FILECTL%
echo forecast_sales_value (varchar(50)) >> %FILECTL%
echo forecast_date (varchar(50)) >> %FILECTL%
echo file=%FILE_NAME%; >> %FILECTL%
echo SHOW; >> %FILECTL%
echo BEGIN LOADING %LOAD_DB%.load_forecast ERRORFILES %TEMP_DB%.load_forecast_e1, %TEMP_DB%.load_forecast_e2; >> %FILECTL%
echo INSERT INTO %LOAD_DB%.load_forecast >> %FILECTL%
echo ( product_code >> %FILECTL%
echo customer_code >> %FILECTL%
echo forecast_quantity >> %FILECTL%
echo forecast_sales_value >> %FILECTL%
echo forecast_date >> %FILECTL%
echo ) >> %FILECTL%
echo VALUES >> %FILECTL%
echo ( :product_code >> %FILECTL%
echo :customer_code >> %FILECTL%
echo :forecast_quantity >> %FILECTL%
echo :forecast_sales_value >> %FILECTL%
echo :forecast_date ( FORMAT 'DD-MMM-YYYY' ) >> %FILECTL%
echo ); >> %FILECTL%
echo END LOADING; >> %FILECTL%
echo LOGOFF; >> %FILECTL%
fastload < %FILECTL% >> %FILEAUD%
SET ERRLEV=%errorlevel%

```

This section calls Fastload, Multi-Load or TPT to load the file. It makes use of a temporary file to build as a control file and then calls Fastload, Multi-Load or TPT to load the data. Note that the load is actually in a for loop. Wild card file names can be used to load multiple files. Each file to be loaded must have the same format.

Note that the data being loaded is appended to the database table. As part of the scheduler run the load table is truncated if the property for truncation is set. In this way multiple files can be loaded into the database table.



If this script is to be executed outside the control of the WhereScape RED scheduler then a truncate statement may need to be performed on the database load table. This would normally be placed before the 'for loop' and would look something like the following:

```
echo DELETE FROM %LOAD_DB%.%LOAD_TABLE% ALL; | bteq .Logon
%DATABASE%/ %USER%, %PWD% >> %FILEAUD%
```

This next section handles the rename and potential looping. The first block of code renames the file and also the trigger file if appropriate. This code is only generated if the rename fields in the file attributes are populated.

The goto label_load statement 9 lines from the end can be used if all the files in a wild card file load are required. Simply uncomment this goto statement and the script will load each file in the wild card.

```
REM
REM *****
REM ***** R E N A M E *****
REM *****
:LABEL_RENAME
SET YYYY=%DATE:~-4%
SET MM=%DATE:~-10,2%
SET DD=%DATE:~-7,2%
SET HOUR1=%TIME:~0,1%K
IF "%HOUR1%" EQU "X X" GOTO LABEL_LEAD_ZERO
SET HH=%TIME:~0,2%
GOTO LABEL_LEAD_END
:LABEL_LEAD_ZERO
SET HH=0%TIME:~1,1%
:LABEL_LEAD_END
SET MI=%TIME:~3,2%
SET SS=%TIME:~6,2%
REM rename the file
move "%FILE_NAME%" "C:\temp\loaded\forecast_%YYYY%%MM%%DD%%HH%%MI%%SS%.txt" >> %FILEAUD%
SET ERRLEV=%errorlevel%
IF %ERRLEV% EQU 0 GOTO LABEL_MOVE3
echo -3
echo %FILE_NAME% NOT MOVED >> %FILEAUD%
SET MSG=The file %FILE_NAME% has NOT been moved to C:\temp\loaded\forecast_%YYYY%%MM%%DD%%HH%%MI%%SS%.txt
echo call %METABASE%\$WrkAuditL('W', '%LOAD_TABLE%', '%SEQUENCE%', '%MSG%', NULL, NULL); | bteq .Logon %DATABASE%/ %USER%, %PWD% >> %FILEAUD%
exit
GOTO LABEL_AFTER_MOVE3
:LABEL_MOVE3
echo %FILE_NAME% moved to C:\temp\loaded\forecast_%YYYY%%MM%%DD%%HH%%MI%%SS%.txt >> %FILEAUD%
SET MSG=The file %FILE_NAME% has been moved to C:\temp\loaded\forecast_%YYYY%%MM%%DD%%HH%%MI%%SS%.txt
echo call %METABASE%\$WrkAuditL('I', '%LOAD_TABLE%', '%SEQUENCE%', '%MSG%', NULL, NULL); | bteq .Logon %DATABASE%/ %USER%, %PWD% >> %FILEAUD%
:LABEL_AFTER_MOVE3
REM rename the trigger file
move %TRIG_NAME% "C:\temp\loaded\forecast_%YYYY%%MM%%DD%%HH%%MI%%SS%.trg" >> %FILEAUD%
SET ERRLEV=%errorlevel%
IF %ERRLEV% EQU 0 GOTO LABEL_MOVE4
echo -3
echo %TRIG_NAME% NOT MOVED >> %FILEAUD%
SET MSG=The file %TRIG_NAME% has NOT been moved to C:\temp\loaded\forecast_%YYYY%%MM%%DD%%HH%%MI%%SS%.trg
echo call %METABASE%\$WrkAuditL('W', '%LOAD_TABLE%', '%SEQUENCE%', '%MSG%', NULL, NULL); | bteq .Logon %DATABASE%/ %USER%, %PWD% >> %FILEAUD%
exit
GOTO LABEL_AFTER_MOVE4
:LABEL_MOVE4
echo %TRIG_NAME% moved to C:\temp\loaded\forecast_%YYYY%%MM%%DD%%HH%%MI%%SS%.trg >> %FILEAUD%
SET MSG=The file %TRIG_NAME% has been moved to C:\temp\loaded\forecast_%YYYY%%MM%%DD%%HH%%MI%%SS%.trg
echo call %METABASE%\$WrkAuditL('I', '%LOAD_TABLE%', '%SEQUENCE%', '%MSG%', NULL, NULL); | bteq .Logon %DATABASE%/ %USER%, %PWD% >> %FILEAUD%
:LABEL_AFTER_MOVE4
REM
REM
```

```
IF %ERRLEV% EQU 0 GOTO LABEL_OKAY
IF %ERRLEV% LEQ 4 GOTO LABEL_WARNING
:LABEL_FAIL
echo -3
IF %ERRLEV% EQU 8 echo Load Failed.  A user error occurred in the loader.
IF %ERRLEV% EQU 12 echo Load Failed.  A fatal error occurred in the loader.
IF %ERRLEV% EQU 16 echo Load Failed.  No message destination available from the loader.
type %FILEAUD% >&2
:ERR_EXIT
exit
:LABEL_WARNING
echo -1
echo Load Completed Normally, with warnings from the loader.
:LABEL_OKAY

```


WINDOWS POWERSHELL SCRIPTS


In addition to the conventional Windows scripting and other tools, WhereScape RED also supports Windows PowerShell scripts for loading data into a WhereScape RED managed Data Warehouse, as well as for exporting data from a WhereScape RED managed Data Warehouse.

The Windows PowerShell command line and scripting environment was introduced by Microsoft in Windows 7. For more information about PowerShell, please refer to the Microsoft TechNet website (<https://technet.microsoft.com/en-us/library/dd742419.aspx>).


Note that WhereScape RED does not automatically create the PowerShell scripts, you must write them either directly via the WhereScape RED script editor or via a template that generates a PowerShell script. Each method is described below:

Via the WhereScape RED Script Editor

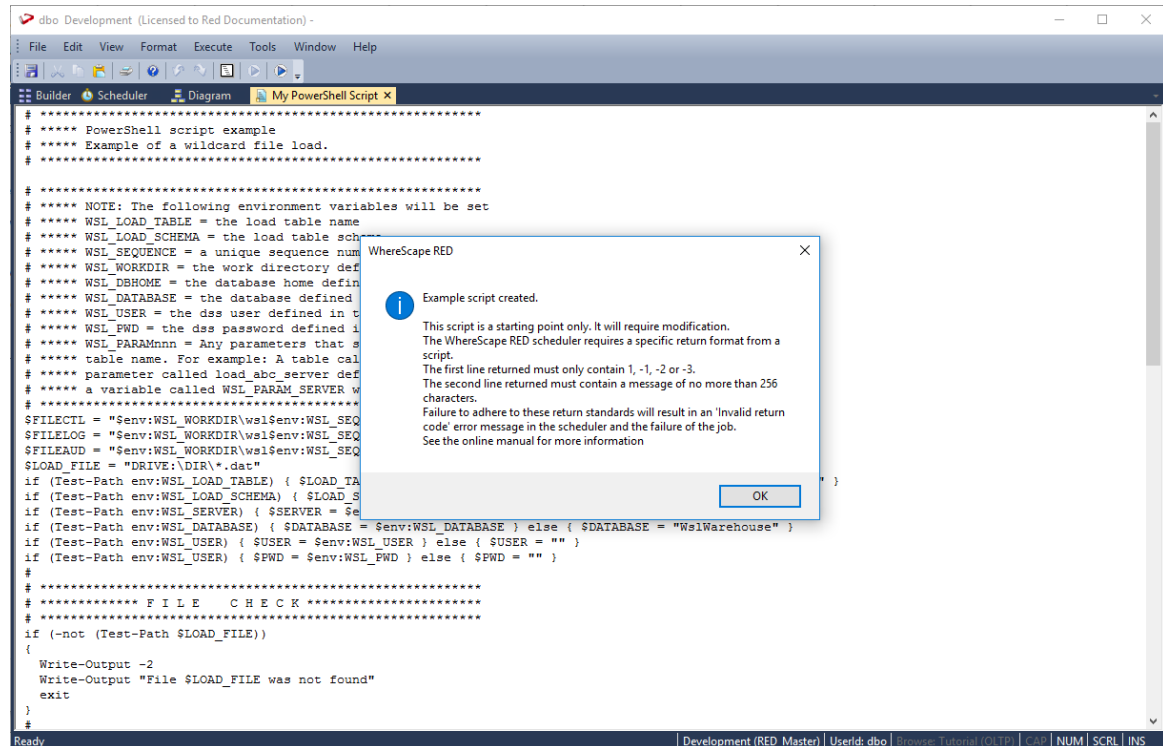
You can manually write the PowerShell script from scratch, using the Script Editor. You need to create a Host Script object and set the script type to PowerShell script from the **Host Script Properties** screen.



You can then open the blank Host Script created of type PowerShell Script to write your own script, or you can use the sample PowerShell scripts in WhereScape RED, which can be generated by selecting **Tools > Create Template Script** from the Script Editor screen.



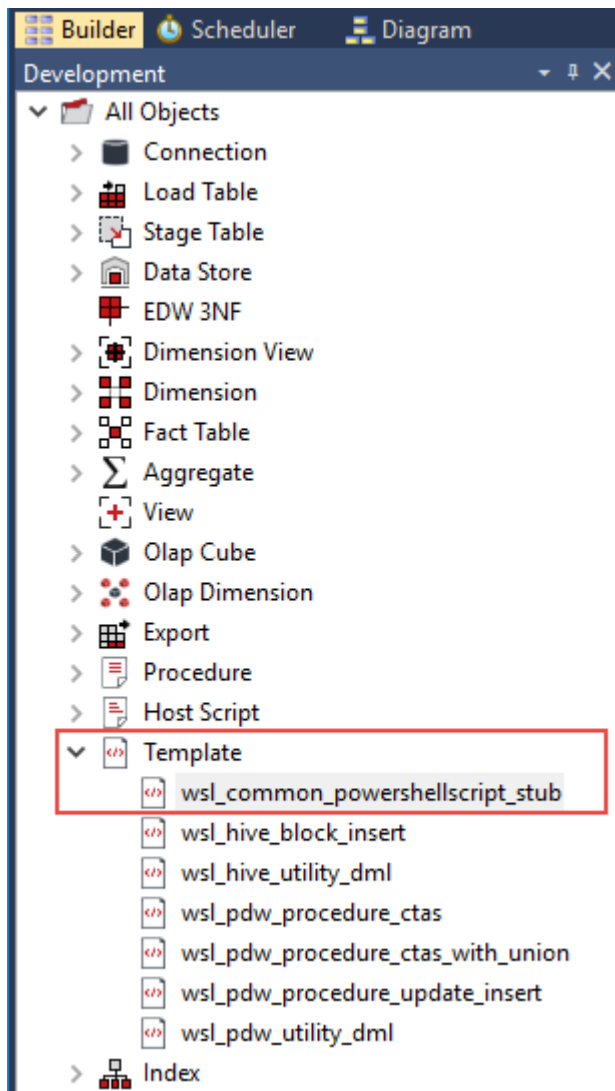
The contents generated is an actual PowerShell script syntax that you can use to load data into a WhereScape RED managed Data Warehouse. You can also modify it and use it as a starting point to write your own PowerShell script for Load or Export objects.



Via the Stub Template

WhereScape RED also provides the Template object **wsl_common_powershellscript_stub** which is a basic PowerShell script type stub template that serves as guide on the use of a template to generate a PowerShell script.

Additional PowerShell Templates can be downloaded from the WhereScape website (<https://www.wherescape.com/support/software-downloads-documentation/templates/>).



You can also create your own Template object and select PowerShell Script under the **Type** drop-down of the **Template Properties** screen.

The screenshot shows a dialog box titled "Template My_PowerShell_Template". On the left, there is a "Properties" tab. The main area contains several fields: "Name" (My_PowerShell_Template), "Purpose" (empty text area), "Author" (WhereScape Documentation), "Target DB" (Common), and "Type" (PowerShell Script). The "Type" dropdown is open, showing a list of options: DDL, DDL (XML/A) Script, PowerShell Script (highlighted), Unix Script, Utility, and Windows Script. Below these fields are "Edit Lock" and "Locked For Edit By:" sections, followed by "Edit Lock Reason or Last Update:" (New Template) and "Timestamps" (Created and Last Update). At the bottom right, there are "OK", "Cancel", and "Help" buttons.

Similar to the WhereScape RED Script Editor method described above, you can edit the template created to write your own script or open an existing PowerShell script to copy its contents to your template.

Note:

WhereScape RED uses the 32-bit version of PowerShell by default. To use the 64-bit version, the following code must be added as the first line of the script. The script is relaunched and will be run in 64-bit version.

```
if ($pshome -like "*syswow64*") {
    & (join-path ($pshome -replace "syswow64", "sysnative") powershell.exe) -
file (join-path $psrootroot $myinvocation.mycommand) @args
    exit
}
```

SCRIPT EDITING

WhereScape RED includes a script editor which enables the maintenance of any host scripts within the data warehouse. The editor is invoked by double-clicking on a script name in the left pane or by right-clicking on the script name and selecting **Edit the Script**.

Indenting code

The tab character inserts four spaces into the text. A shift/tab removes four spaces.

Cut, Copy, Paste and Delete

The normal Windows cut, copy, paste and delete functions are available either through the toolbar or via the right-click pop up menu.

Indenting a block of text

A selected block of text can be indented by four spaces by depressing the tab character. Each tab indents by a further four spaces. A shift/tab removes four spaces from the front of each line in the selected block.

Viewing other scripts

During the editing process it is possible to pop up a window containing other scripts. This window enables cut and paste operations. In such a way, it can be used as a work area or as a source of code. Select the **Tools>View Script or Template** menu option to bring the viewer window up. A dialog appears.

A number of drop-down lists can be chosen from. Once an item is selected the viewer loads the code and moves to the right side of the edit window. The various options are:

- Version: A previously saved version of this script. The description and time are shown.
- Script: Some other script stored within the WhereScape metadata.
- Template: A template script as defined in the WhereScape metadata.

Once an item is chosen the viewer appears on the right side of the edit window. The viewer or the main edit window can receive the focus by clicking the mouse within the appropriate window. Code can be cut from one window and pasted into the other. Any changes made in the viewer window cannot be saved.

NOTE: Editing, deleting or compiling Locked for Edit or opened procedures/scripts

Procedures or scripts cannot be deleted if they are Locked for Edit by any user, checked out by another user or if there is another object that has the same associated procedures or scripts. Saving or Compiling in the procedure or script edit window cannot be performed if the procedures or scripts become Locked for Edit by other users after the edit window was opened. Procedures or scripts cannot be deleted or modified after the edit window has been opened, unless the Edit Lock has been released. Edit Locks can be released by any user in the Script or Procedure Properties screen.

To prevent updates, deletes and modifications to certain procedures or scripts, it is best to use the Check Out functionality instead. For more information about this functionality, please see section **Check Outs and Check Ins**.

SCRIPT TESTING

When a host script is scheduled, it is run in the scheduler environment. Therefore, a UNIX scheduler must be available to run a UNIX script and only a Windows scheduler can run a Windows script.

It is possible to test a script interactively.

A script is invoked via the **Execute>Execute the Script** menu option. You can also right click the script from the **Host Script** objects list on the left pane and select **Execute Script** from the context menu.

The output from the script is shown in a pop-up dialog box.

SCRIPT SYNTAX

There are a number of conventions that must be followed if a host script is to be used by the WhereScape scheduler. These conventions are:

- 1 The first line of data in 'standard out' must contain the resultant status of the script. Valid values are '1' to indicate success, '-1' to indicate a warning condition occurred but the result is considered a success, '-2' to indicate a handled error occurred and subsequent dependent tasks should be held, -3 to indicate an unhandled Failure and that subsequent dependent tasks should be held.
- 2 The second line of data in 'standard out' must contain a resultant message of no more than 256 characters.
- 3 Any subsequent lines in 'standard out' are considered informational and are recorded in the audit trail. The normal practice is to place a minimum of information in the audit trail. All bulk information should be output to 'standard error'.
- 4 Any data output to 'standard error' will be written to the error/detail log. Both the audit log and detail log can be viewed from the WhereScape RED tool under the scheduler window.

Example:

In the following example the first line '@echo off' prevents unwanted information from being reported to standard out. A Multi-Load script file is built up (echo statements). The mload command is then executed to run the load.

```

@echo off
echo Loading c:\temp\budget.txt >> wsl_load_budget100001.aud
echo .Logtable WslWarehouse.load_budget_e1; > wsl_load_budget100001.ct1
echo .Logon DemoTDAT/WslWarehouse,wsl; >> wsl_load_budget100001.ct1
echo DROP TABLE WslWarehouse.ET_load_budget; >>
wsl_load_budget100001.ct1
echo DROP TABLE WslWarehouse.UV_load_budget; >>
wsl_load_budget100001.ct1
echo DROP TABLE WslWarehouse.WT_load_budget; >>
wsl_load_budget100001.ct1
echo .Begin Import Mload tables WslWarehouse.load_budget >>
wsl_load_budget100001.ct1
echo Worktables WslWarehouse.WT_load_budget >>
wsl_load_budget100001.ct1
echo Errortables WslWarehouse.ET_load_budget >>
wsl_load_budget100001.ct1
echo WslWarehouse.UV_load_budget; >>
wsl_load_budget100001.ct1
echo .Layout Transaction; >> wsl_load_budget100001.ct1
echo .Field product_code * VARCHAR(50); >> wsl_load_budget100001.ct1
echo .Field customer_code * VARCHAR(33); >> wsl_load_budget100001.ct1
echo .Field budget_quantity * VARCHAR(50); >>
wsl_load_budget100001.ct1
echo .Field budget_sales_value * VARCHAR(50); >>
wsl_load_budget100001.ct1
echo .Field budget_date * VARCHAR(50); >> wsl_load_budget100001.ct1
echo .DML Label Inserts; >> wsl_load_budget100001.ct1
echo INSERT INTO WslWarehouse.load_budget >> wsl_load_budget100001.ct1
echo ( product_code>> wsl_load_budget100001.ct1
echo , customer_code>> wsl_load_budget100001.ct1
echo , budget_quantity>> wsl_load_budget100001.ct1
echo , budget_sales_value>> wsl_load_budget100001.ct1
echo , budget_date>> wsl_load_budget100001.ct1
echo ) >> wsl_load_budget100001.ct1
echo VALUES >> wsl_load_budget100001.ct1
echo ( :product_code>> wsl_load_budget100001.ct1
echo , :customer_code>> wsl_load_budget100001.ct1
echo , :budget_quantity>> wsl_load_budget100001.ct1
echo , :budget_sales_value>> wsl_load_budget100001.ct1
echo , :budget_date ( FORMAT 'DD-MMM-YYYY') >>
wsl_load_budget100001.ct1
echo ); >> wsl_load_budget100001.ct1
echo .Import Infile 'c:\temp\budget.txt' >> wsl_load_budget100001.ct1
echo From 2 >> wsl_load_budget100001.ct1
echo Format Vartext '^,' >> wsl_load_budget100001.ct1
echo Layout Transaction >> wsl_load_budget100001.ct1
echo Apply Inserts; >> wsl_load_budget100001.ct1
echo .End Mload; >> wsl_load_budget100001.ct1
echo .Logoff; >> wsl_load_budget100001.ct1
mload -b < wsl_load_budget100001.ct1 >> wsl_load_budget100001.aud
SET ERRLEV=%errorlevel%
IF %ERRLEV% EQU 0 GOTO LABEL_OKAY
IF %ERRLEV% LEQ 4 GOTO LABEL_WARNING
:LABEL_FAIL
echo -3
IF %ERRLEV% EQU 8 echo Load Failed. A user error occurred in the
loader.
IF %ERRLEV% EQU 12 echo Load Failed. A fatal error occurred in the
loader.

```

```

IF %ERRLEV% EQU 16 echo Load Failed. No message destination available
from the loader.
type wsl_load_budget100001.aud >&2
:ERR_EXIT
exit
:LABEL_WARNING
echo -1
echo Load Completed Normally, with warnings from the loader.
:LABEL_OKAY
echo 1
echo Load Completed Normally.
type wsl_load_budget100001.aud
exit

```

SCRIPT ENVIRONMENT VARIABLES

The following environment variables are available for all script loads and script exports, both Windows and UNIX/Linux.

All load scripts

The following variables are available in all load scripts:

Windows variable	UNIX/Linux variable	Description									
WSL_LOAD_FULLNAME	LOAD_FULLNAME	The fully-qualified load table name.									
WSL_LOAD_TABLE	LOAD_TABLE	The unqualified load table name.									
WSL_LOAD_SCHEMA	LOAD_SCHEMA	<p>The schema for the load table.</p> <p>Note: A trailing dot is appended for SQL Server or Oracle due to historical usage. A trailing dot is not appended for any other database type due to historical usage.</p> <p>However, it is better not to assume the trailing dot is or isn't appended by using the variable like this, when it is not empty:</p> <table border="1"> <thead> <tr> <th>OS</th> <th>If no trailing dot is wanted</th> <th>If a trailing dot is wanted</th> </tr> </thead> <tbody> <tr> <td>Windows</td> <td>!WSL_LOAD_SCHEMA : . = !</td> <td>!WSL_LOAD_SCHEMA : . = ! .</td> </tr> <tr> <td>UNIX/Linux</td> <td>`\${LOAD_SCHEMA%.}`</td> <td>`\${LOAD_SCHEMA%.}` .</td> </tr> </tbody> </table>	OS	If no trailing dot is wanted	If a trailing dot is wanted	Windows	!WSL_LOAD_SCHEMA : . = !	!WSL_LOAD_SCHEMA : . = ! .	UNIX/Linux	`\${LOAD_SCHEMA%.}`	`\${LOAD_SCHEMA%.}` .
OS	If no trailing dot is wanted	If a trailing dot is wanted									
Windows	!WSL_LOAD_SCHEMA : . = !	!WSL_LOAD_SCHEMA : . = ! .									
UNIX/Linux	`\${LOAD_SCHEMA%.}`	`\${LOAD_SCHEMA%.}` .									

Windows variable	UNIX/Linux variable	Description
WSL_LOAD_DB	LOAD_DB	The name of the database for the load table.
WSL_TEMP_DB	TEMP_DB	Teradata: The name of the database for load temporary tables. PDW: The name of the staging database for the load. Others: Not Used.
WSL_TGT_DSN	TGT_DSN	The ODBC data source name (DSN) for the load table's storage connection.
WSL_TGT_SERVER	TGT_SERVER	The server for the load table's storage connection.
WSL_TGT_DBPORT	TGT_DBPORT	The database port for the load table's storage connection.
WSL_TGT_DBID	TGT_DBID	The <i>Database ID</i> property of the load table's storage connection. For Teradata this is the Teradata Director Program ID (TDPID). For Oracle this is the Oracle SID or TNS Name.
WSL_TGT_USER	TGT_USER	The user id for the load table's storage connection.
WSL_TGT_PWD	TGT_PWD	The password for the load table's storage connection.

All load scripts from Database or ODBC connections

In **addition** to the variables in the previous table, the following variables are available in all load scripts from Database or ODBC connections:

Windows variable	UNIX/Linux variable	Description
WSL_SRC_DSN	SRC_DSN	The ODBC data source name (DSN) for the source connection.
WSL_SRC_SERVER	SRC_SERVER	The server for the source connection.
WSL_SRC_DBPORT	SRC_DBPORT	The database port for the source connection.
WSL_SRC_DBID	SRC_DBID	The <i>Database ID</i> property of the source connection. For Teradata this is the Teradata Director Program ID (TDPID). For Oracle this is the Oracle SID or TNS Name.
WSL_SRC_DB	SRC_DB	The name of the database for the source connection.

Windows variable	UNIX/Linux variable	Description									
WSL_SRC_SCHEMA	SRC_SCHEMA	<p>The Source Schema property of the load.</p> <p>Note: The property is fetched without modification, so there may or may not be a trailing dot depending on how it is configured.</p> <p>However, it is better not to assume the trailing dot is or isn't appended by using the variable like this, when it is not empty:</p> <table border="1"> <thead> <tr> <th>OS</th> <th>If no trailing dot is wanted</th> <th>If a trailing dot is wanted</th> </tr> </thead> <tbody> <tr> <td>Windows</td> <td>!WSL_SRC_SCHEMA:. =!</td> <td>!WSL_SRC_SCHEMA:. =!</td> </tr> <tr> <td>UNIX/Linux</td> <td>\${SRC_SCHEMA%.}</td> <td>\${SRC_SCHEMA%.}</td> </tr> </tbody> </table>	OS	If no trailing dot is wanted	If a trailing dot is wanted	Windows	!WSL_SRC_SCHEMA:. =!	!WSL_SRC_SCHEMA:. =!	UNIX/Linux	\${SRC_SCHEMA%.}	\${SRC_SCHEMA%.}
OS	If no trailing dot is wanted	If a trailing dot is wanted									
Windows	!WSL_SRC_SCHEMA:. =!	!WSL_SRC_SCHEMA:. =!									
UNIX/Linux	\${SRC_SCHEMA%.}	\${SRC_SCHEMA%.}									
WSL_SRC_USER	SRC_USER	The user id for the source connection.									
WSL_SRC_PWD	SRC_PWD	The password for the source connection.									

All export scripts

The following variables are available in all export scripts:

Windows variable	UNIX/Linux variable	Description
WSL_EXP_NAME	EXP_NAME	The export object name.
WSL_EXP_FULLNAME	EXP_FULLNAME	The fully-qualified export table name.
WSL_EXP_TABLE	EXP_TABLE	<p>The unqualified export table name.</p> <p>Note: For Windows script exports from PDW, this variable is initialized with the fully-qualified export table name, due to historical usage.</p> <p>To enable this variable to be uniformly described and used as the unqualified export table name, an additional variable WSL_EXP_SIMPLENAME is created.</p> <p>This allows the following command to be explicitly added to the top of the script by the script author, to be executed before all other processing:</p> <pre>if defined WSL_EXP_SIMPLENAME (SET WSL_EXP_TABLE=!WSL_EXP_SIMPLENAME!) else SET WSL_EXP_TABLE=</pre>

Windows variable	UNIX/Linux variable	Description									
		After such a command is executed, the variable WSL_EXP_TABLE will contain the unqualified export table name.									
WSL_EXP_SCHEMA	EXP_SCHEMA	<p>The schema for the export table.</p> <p>Note: A trailing dot is appended for SQL Server or Oracle due to historical usage. A trailing dot is not appended for any other database type due to historical usage.</p> <p>However, it is better not to assume the trailing dot is or isn't appended by using the variable like this, when it is not empty:</p> <table border="1" data-bbox="683 846 1401 1120"> <thead> <tr> <th>OS</th> <th>If no trailing dot is wanted</th> <th>If a trailing dot is wanted</th> </tr> </thead> <tbody> <tr> <td>Windows</td> <td>!WSL_EXP_SCHEMA : . = !</td> <td>!WSL_EXP_SCHEMA : . = ! .</td> </tr> <tr> <td>UNIX/Linux</td> <td>\${EXP_SCHEMA%.}</td> <td>\${EXP_SCHEMA%.} .</td> </tr> </tbody> </table>	OS	If no trailing dot is wanted	If a trailing dot is wanted	Windows	!WSL_EXP_SCHEMA : . = !	!WSL_EXP_SCHEMA : . = ! .	UNIX/Linux	\${EXP_SCHEMA%.}	\${EXP_SCHEMA%.} .
OS	If no trailing dot is wanted	If a trailing dot is wanted									
Windows	!WSL_EXP_SCHEMA : . = !	!WSL_EXP_SCHEMA : . = ! .									
UNIX/Linux	\${EXP_SCHEMA%.}	\${EXP_SCHEMA%.} .									
WSL_EXP_DB	EXP_DB	The name of the database for the export table.									
WSL_TEMP_DB	TEMP_DB	<p>Teradata: The name of the database for export temporary tables.</p> <p>Others: Not used.</p>									
WSL_SRC_DSN	SRC_DSN	The ODBC data source name (DSN) for the export table's storage connection.									
WSL_SRC_SERVER	SRC_SERVER	The server for the export table's storage connection.									
WSL_SRC_DBPORT	SRC_DBPORT	The database port for the export table's storage connection.									
WSL_SRC_DBID	SRC_DBID	<p>The <i>Database ID</i> property of the export table's storage connection.</p> <p>For Teradata this is the Teradata Director Program ID (TDPID).</p> <p>For Oracle this is the Oracle SID or TNS Name.</p>									
WSL_SRC_USER	SRC_USER	The user id for the export table's storage connection.									
WSL_SRC_PWD	SRC_PWD	The password for the export table's storage connection.									

All scripts

In **addition** to the specific variables in the previous tables, the following variables are available in all scripts:

Windows variable	UNIX/Linux variable	Description									
WSL_META_DSN	META_DSN	The ODBC data source name (DSN) for the meta-repository connection.									
WSL_META_SERVER	META_SERVER	The server for the meta-repository connection.									
WSL_META_DBID	META_DBID	The <i>Database ID</i> property of the meta-repository connection. For Teradata this is the Teradata Director Program ID (TDPID). For Oracle this is the Oracle SID or TNS Name.									
WSL_META_DB	META_DB	The name of the database for the meta-repository connection.									
WSL_META_SCHEMA	META_SCHEMA	The meta-repository table qualifier, with a trailing dot. For SQL Server and DB2 this is the schema for the meta-repository. For Teradata and Oracle this is not actually a schema name. Note: A trailing dot is appended due to historical usage. However, it is better not to assume the trailing dot is or isn't appended by using the variable like this, when it is not empty: <table border="1" data-bbox="730 1375 1447 1648"> <thead> <tr> <th>OS</th> <th>If no trailing dot is wanted</th> <th>If a trailing dot is wanted</th> </tr> </thead> <tbody> <tr> <td>Windows</td> <td>!WSL_META_SCHEMA : . = !</td> <td>!WSL_META_SCHEMA : . = ! .</td> </tr> <tr> <td>UNIX/Linux</td> <td>\$(META_SCHEMA%.)</td> <td>\$(META_SCHEMA%.) .</td> </tr> </tbody> </table>	OS	If no trailing dot is wanted	If a trailing dot is wanted	Windows	!WSL_META_SCHEMA : . = !	!WSL_META_SCHEMA : . = ! .	UNIX/Linux	\$(META_SCHEMA%.)	\$(META_SCHEMA%.) .
OS	If no trailing dot is wanted	If a trailing dot is wanted									
Windows	!WSL_META_SCHEMA : . = !	!WSL_META_SCHEMA : . = ! .									
UNIX/Linux	\$(META_SCHEMA%.)	\$(META_SCHEMA%.) .									
WSL_META_USER	META_USER	The user id for the meta-repository connection.									
WSL_META_PWD	META_PWD	The password for the meta-repository connection.									
WSL_WORKDIR	WORKDIR	Windows: The work directory defined in the Windows connection. UNIX/Linux: The work directory defined in the UNIX/Linux or Hadoop connection.									
WSL_SEQUENCE	SEQ	A unique sequence number for the load or export task.									

Windows variable	UNIX/Linux variable	Description
WSL_PARAMnnn	PARAMnnn	<p>Any parameters that start with the load table or export object name.</p> <p>Example:</p> <p>A table called <i>load_abc</i> has a parameter called <i>load_abc_server</i> defined.</p> <p>In this case, a variable called <i>WSL_PARAM_SERVER</i> (Windows) or <i>PARAM_SERVER</i> (UNIX/Linux) will be created.</p>

CALLING A BATCH FILE FROM A SCRIPT

Below is an example RED host script which calls a batch file:

```
@ECHO OFF
SETLOCAL ENABLEDELAYEDEXPANSION
SETLOCAL ENABLEEXTENSIONS
CALL c:\temp\MyBatchFile.bat > c:\temp\MyBatchFile.log 2>&1
IF %ERRORLEVEL% EQU 0 GOTO LABEL_OKAY
ECHO -2
ECHO Batch file returned an error code of %ERRORLEVEL%
TYPE c:\temp\MyBatchFile.log
EXIT
:LABEL_OKAY
ECHO 1
ECHO Batch file completed successfully
TYPE c:\temp\MyBatchFile.log
```

Where "c:\temp\MyBatchFile.bat" contains this:

```
ECHO Hello
SET ERRORLEVEL=0
```

Create the Host Script in RED:

Host Script Call_Batch_File

Properties

Name: Call_Batch_File Type: Windows script

Purpose: Calls a batch file

Owner: dssdemo Delete Lock

Last Update By:

Default Connect: Windows

Edit Lock

Locked For Edit By:

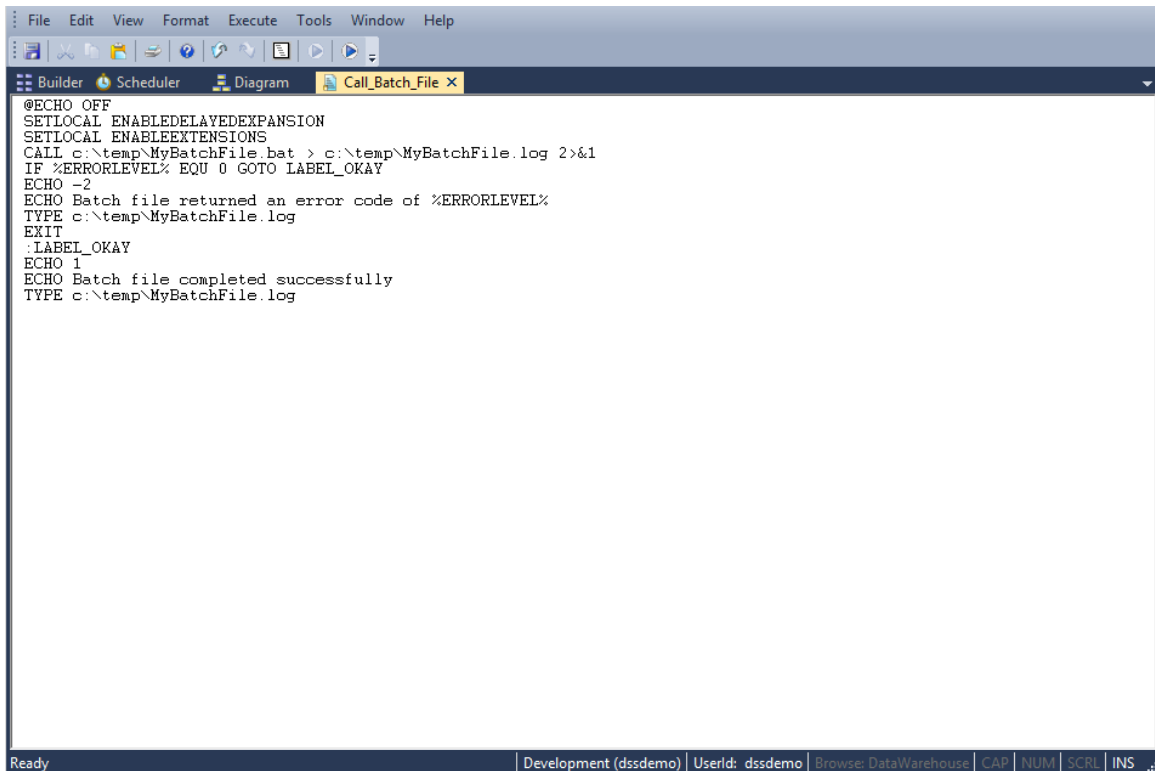
Edit Lock Reason or Last Update: New Script

Timestamps

Created: 2015-11-19 03:12:07 Last Update: Compiled:

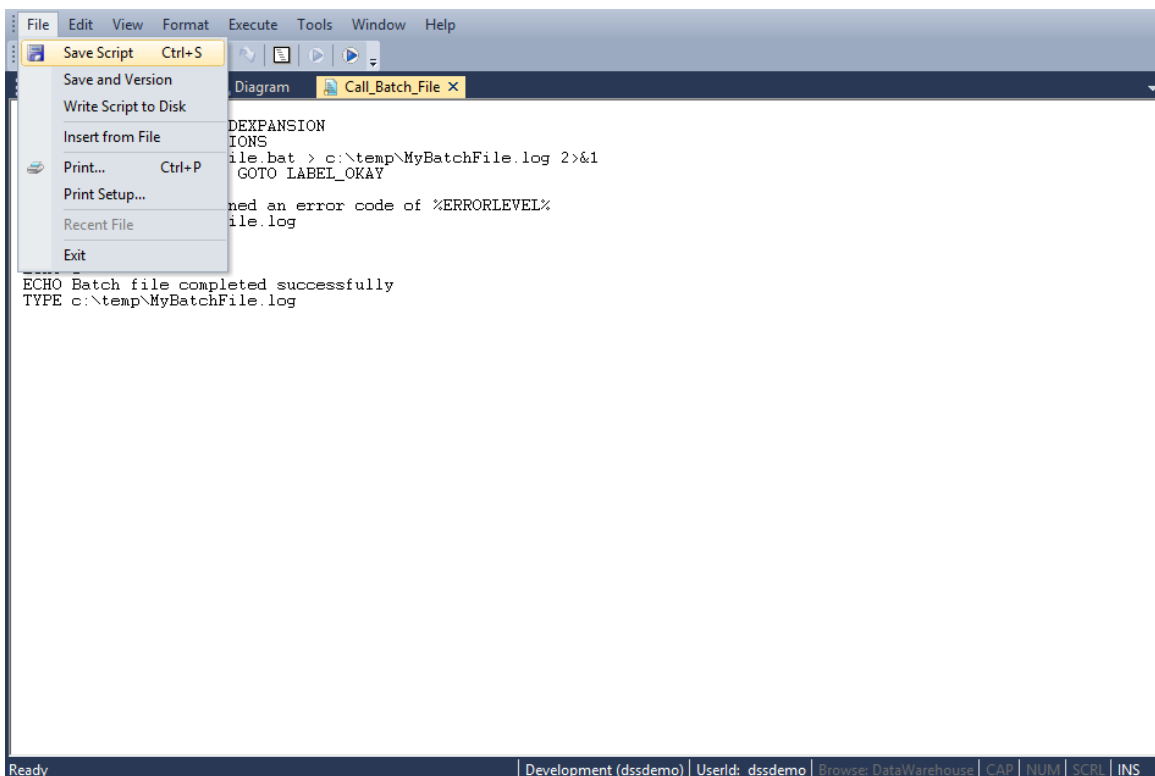
OK Cancel Help

Edit the Script and enter the following:



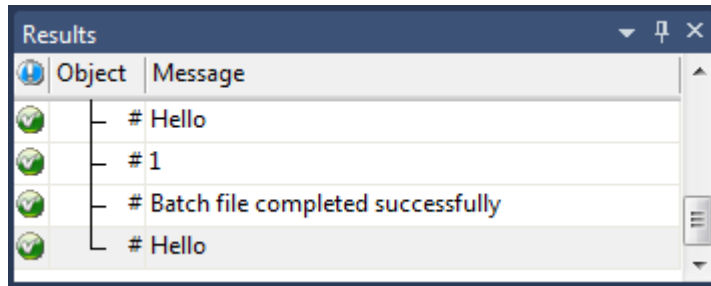
```
@ECHO OFF
SETLOCAL ENABLEDELAYEDEXPANSION
SETLOCAL ENABLEEXTENSIONS
CALL c:\temp\MyBatchFile.bat > c:\temp\MyBatchFile.log 2>&1
IF %ERRORLEVEL% EQU 0 GOTO LABEL_OKAY
ECHO -2
ECHO Batch file returned an error code of %ERRORLEVEL%
TYPE c:\temp\MyBatchFile.log
EXIT
:LABEL_OKAY
ECHO 1
ECHO Batch file completed successfully
TYPE c:\temp\MyBatchFile.log
```

Save the Script:



```
File Edit View Format Execute Tools Window Help
Save Script Ctrl+S
Save and Version
Write Script to Disk
Insert from File
Print... Ctrl+P
Print Setup...
Recent File
Exit
Diagram Call_Batch_File X
EXPANSION
IONS
file.bat > c:\temp\MyBatchFile.log 2>&1
GOTO LABEL_OKAY
ned an error code of %ERRORLEVEL%
file.log
ECHO Batch file completed successfully
TYPE c:\temp\MyBatchFile.log
```

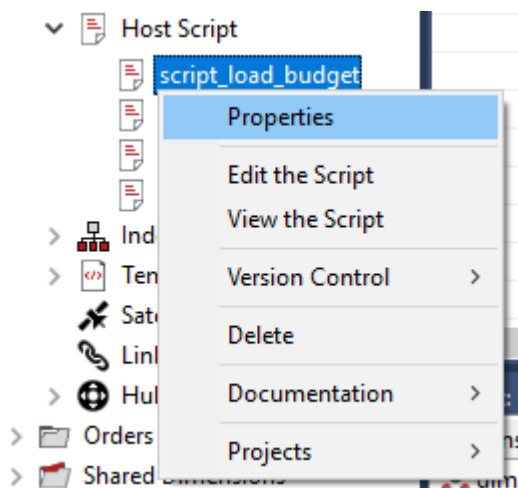
When the script is executed, you will see the following results:



SCHEDULING SCRIPTS

When a host script is scheduled, it is run in the scheduler environment. Therefore a UNIX scheduler must be available to run a UNIX script and only a Windows scheduler can run a Windows script.

It is important to set the **default connection** on the Properties screen for that script. Right-click on the host script in the left pane and select **Properties**.



Set **Default Connect** to either **Windows** or **Unix** and click **OK**.

Host Script script_load_budget

Properties

Name: Type:

Purpose:

Owner: Delete Lock

Last Update By:

Default Connect:

Edit Lock

Locked For Edit By:

Edit Lock Reason or Last Update:

Timestamps

Created: Last Update: Compiled:

OK Cancel Help

Note: If you fail to set the default connection for the host script, you will receive a return message of **Invalid Host Type** when the host script is executed.

There are a number of conventions that must be followed if a host script is to be used by the WhereScape scheduler. These conventions are:

- 1 The first line of data in 'standard out' must contain the resultant status of the script. Valid values are '1' to indicate success, '-1' to indicate a warning condition occurred but the result is considered a success, '-2' to indicate a handled error occurred and subsequent dependent tasks should be held, -3 to indicate an unhandled Failure and that subsequent dependent tasks should be held.
- 2 The second line of data in 'standard out' must contain a resultant message of no more than 256 characters.
- 3 Any subsequent lines in 'standard out' are considered informational and are recorded in the audit trail. The normal practice is to place a minimum of information in the audit trail. All bulk information should be output to 'standard error'.
- 4 Any data output to 'standard error' will be written to the error/detail log. Both the audit log and detail log can be viewed from the WhereScape RED tool under the scheduler window.

MANUALLY CREATED SCRIPTS

Individual scripts can also be manually created in RED to perform and schedule tasks that are not related to load tables.

The example below shows a minimal script that will run successfully.

```
@echo off

REM *****
REM ***** Parameter Example
REM *****

SET /A RESULT_CODE=1
SET RESULT_MESSAGE="Success. "

SET FILEAUD=%WORKDIR%\ws1%SCRIPT%%SEQUENCE%.aud
SET FILE_PROCESS_LOG=c:\temp\process.log

echo $PDS_Customer_Process_Date$ >> %FILE_PROCESS_LOG%

echo %RESULT_CODE%
echo %RESULT_MSG%

exit
```

Please note that you need to use the following codes to determine the script's results meaning. It is important that one of these codes is the first output of the script.

Output	Description
Result Number	Output Result Number:
	1 Success.
	-1 Warning.
	-2 Error.
	-3 Fatal/Unexpected Error.

CHAPTER 25

TEMPLATES

Templates provide the ability to customize automatically generated code within RED. This feature is most suited to users that would like to customize automatically generated code or would like to expand RED to support non-native database platforms.

Creating templates is an advanced function that requires intimate knowledge of RED operations and metadata structure. WhereScape recommends that you contact our consulting team to assist with this feature. However, should you wish to use this feature independently, example templates and up-to-date reference information is available on our website:

<https://www.wherescape.com/support/software-downloads-documentation/wherescape-red/templates/>

Some templates may be included in your RED installation, depending on your license.

Each template is assigned a type and a target database, these properties are used to assist with filtering when associating table operations to templates. RED supports templates for the following operations:

Operation	Database	Template Type
Create DDL	All database types	DDL
Export Script	All database types	Windows Script
		PowerShell Script
		Unix Script
		Olap/XMLA Script
Load Script	All database types	Windows Script
		PowerShell Script
		Unix Script
		Olap/XMLA Script
Update Procedure Custom Procedure Post Load Procedure	Custom	Block
		PowerShell Script
	Hive	Block
	SQL, Teradata, Oracle and PDW	Block
		Procedure

Notes:

- Script-based loads and exports on Windows supports both DOS Batch and **PowerShell scripts** (see "**Windows PowerShell Scripts**" on page 653).
 - Refer to **Rebuilding Update Procedures** (on page 199) for details on generating an update procedure via a template.
-

Utility type templates can contain common code for use by other templates.

Templates are written in the Pebble template language, for more information on Pebble see <http://www.mitchellbosecke.com/pebble/documentation/guide/basic-usage>.



TIP: Detailed logs can be produced during template evaluation by typing *FULLLOG* in the Notes of the relevant connection.

IN THIS CHAPTER

Template Properties	674
Template Editor	676
Template Usage.....	681

TEMPLATE PROPERTIES

The properties screen for a template is shown below.

The screenshot shows a dialog box titled "Template wsl_common_powershellscript_stub- NO UPDATE: This template is READ ONLY." with a close button (X) in the top right corner. The dialog is divided into several sections:

- Properties:** A sidebar on the left with a "Properties" tab selected.
- Name:** A text field containing "wsl_common_powershellscript_stub".
- Purpose:** A text area containing "This template creates a PowerShell script stub".
- Author:** A text field containing "WhereScape Ltd".
- Target DB:** A dropdown menu set to "Common".
- Type:** A dropdown menu set to "PowerShell Script".
- Edit Lock:** A section containing:
 - Locked For Edit By:** A greyed-out text field.
 - Edit Lock Reason or Last Update:** A text area.
- Timestamps:** A section containing:
 - Created:** A text field with the value "2017-12-12 23:10:34.000000".
 - Last Update:** A text field with the value "2017-12-12 23:10:34.000000".

At the bottom right, there are three buttons: "OK", "Cancel", and "Help".

Name, **Purpose** and **Author** fields should be completed to provide background information on the template, these fields are purely informational.

Created and **Last Update** fields provide date information on the template.

The **Target DB** sets the type of database connections for this template. The template will only be selectable for an operation when the **Target DB** field matches that of the object. Target DB is restricted based on your license.

- Common (applies to all databases)
- Custom
- DB2
- Greenplum
- Hive
- Netezza
- Oracle
- PDW
- SQL Server
- Tabular
- Teradata

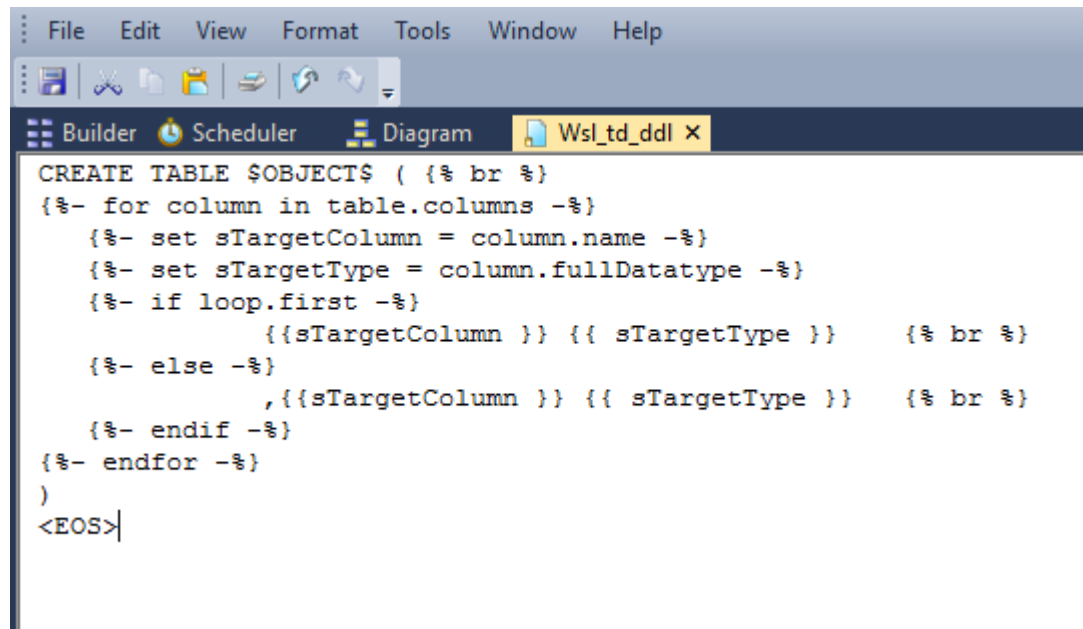
NOTE: Hive and Custom update procedure templates only support Block update procedures so you should create a block template for these.

The **Type** field informs RED what this template can be used for. This can be set to one of the following:

- Block
- DDL
- Function
- OLAP (XML/A) Script
- PowerShell Script
- Procedure
- Unix Script
- Utility
- Windows Script

TEMPLATE EDITOR

Right-click on a template and select **Edit Template** or **View Template** to open the Template Editor.



```
File Edit View Format Tools Window Help
Builder Scheduler Diagram Wsl_td_ddl x
CREATE TABLE $OBJECT$ ( {% br %}
{%- for column in table.columns -%}
  {%- set sTargetColumn = column.name -%}
  {%- set sTargetType = column.fullDatatype -%}
  {%- if loop.first -%}
    {{sTargetColumn }} {{ sTargetType }}   {% br %}
  {%- else -%}
    ,{{sTargetColumn }} {{ sTargetType }}   {% br %}
  {%- endif -%}
{%- endfor -%}
)
<EOS>
```

EVALUATING AN API OUTLINE TEMPLATE

An API Outline Template is available to output all object properties relevant to the current object. Upon evaluation of this template, the status of each property is generated and printed to the script or procedure file.

NOTE: Template evaluation usually generates a script or procedure file, but the API Outline Template generates plain text. The output of this template is intended to be viewed or copied to a text file, it cannot be executed as a script.

To evaluate an API Outline Template:

- 1 Create a **new template**. The template can normally be of any type, but in this example, we will use the DDL template type because viewing the evaluation of DDL templates is simple. Set **Target DB** to your source connection database type. In this example, we will set the Target DB to SQL Server because the load table we are evaluating is stored on a SQL Server connection.

The screenshot shows a dialog box titled "Template ddl_api_outline". On the left, there is a "Properties" tab. The main area contains the following fields:

- Name: ddl_api_outline
- Purpose: API Outline generation as a DDL Template. Dumps metadata information on an object.
- Author: DOCS
- Target DB: SQL Server
- Type: DDL

Below these fields is an "Edit Lock" section with the following fields:

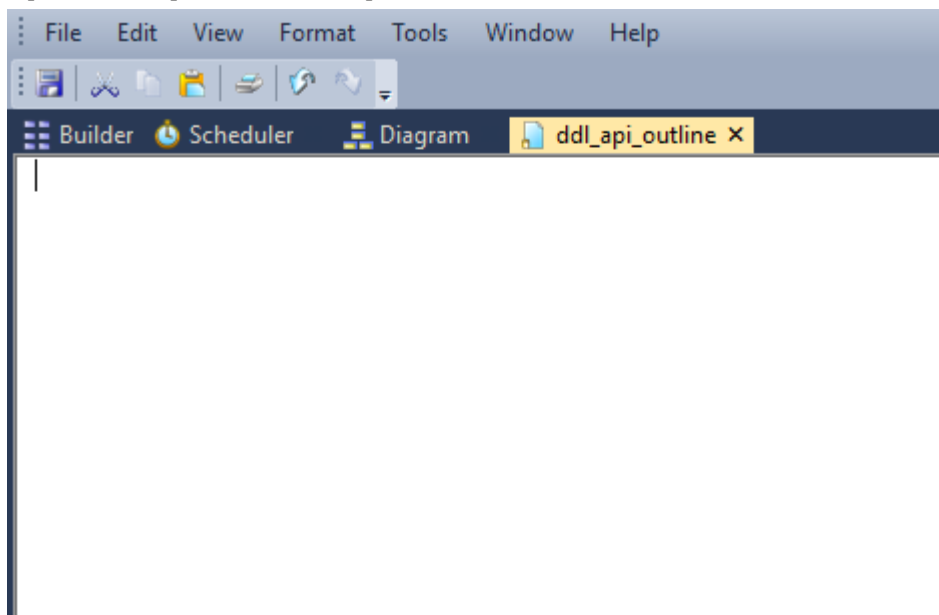
- Locked For Edit By: [Redacted]
- Edit Lock Reason or Last Update: New Template

At the bottom, there is a "Timestamps" section with the following fields:

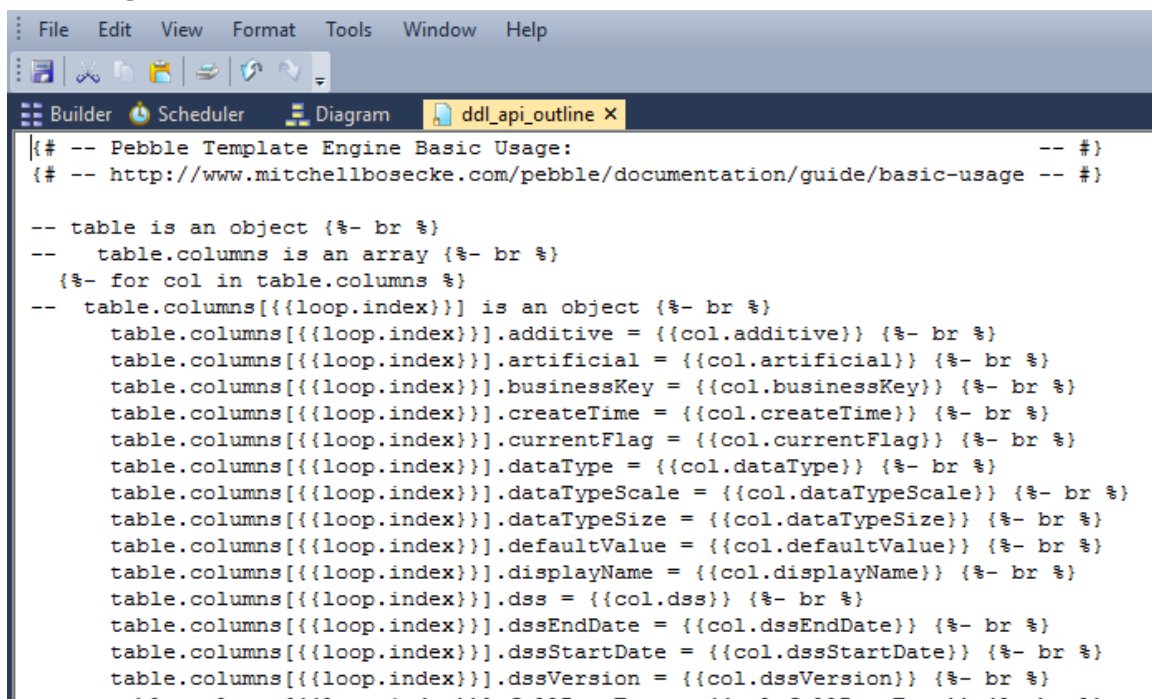
- Created: [Empty]
- Last Update: [Empty]

At the bottom right, there are three buttons: "OK", "Cancel", and "Help".

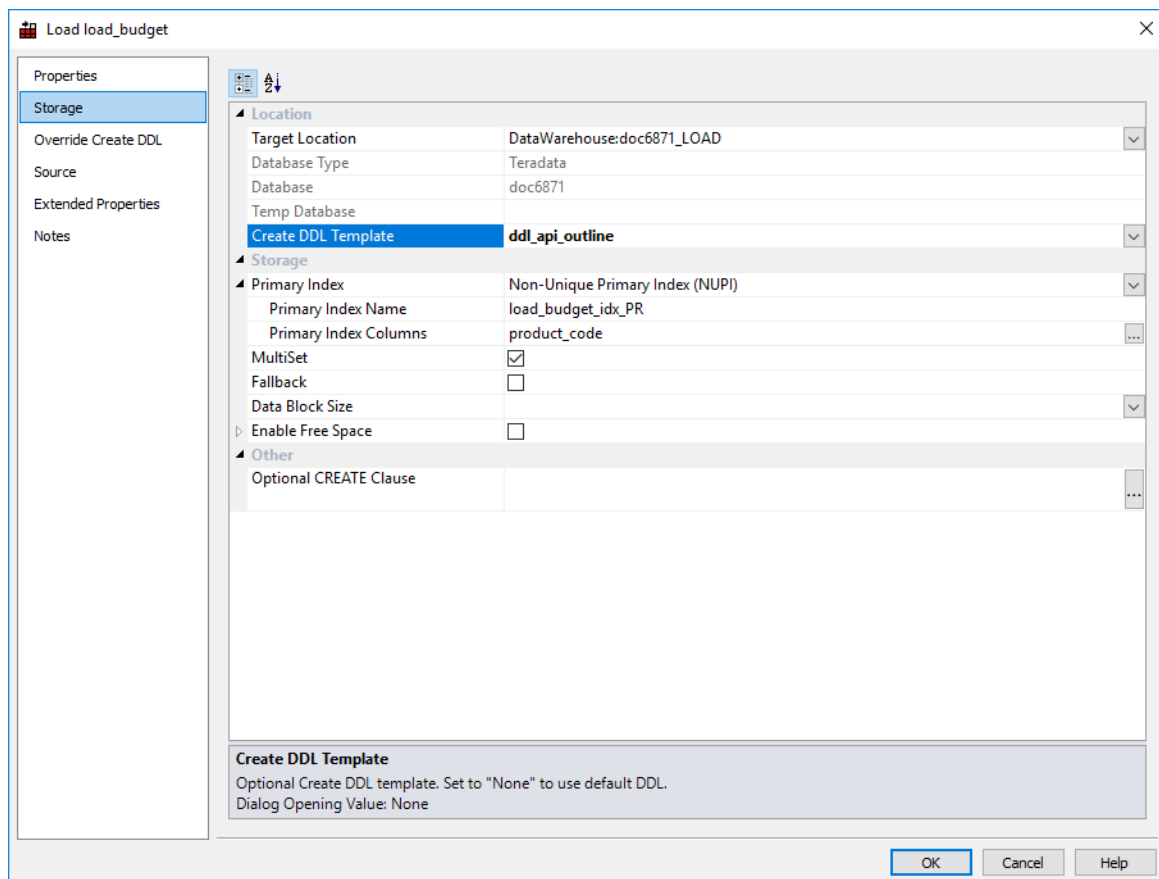
- Open the template in the Template Editor.



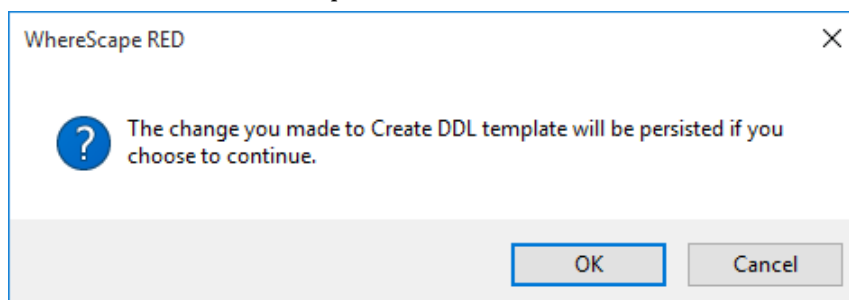
- Click **Tools > Create API Example Outline**. The API Example Outline text is added to the blank template.



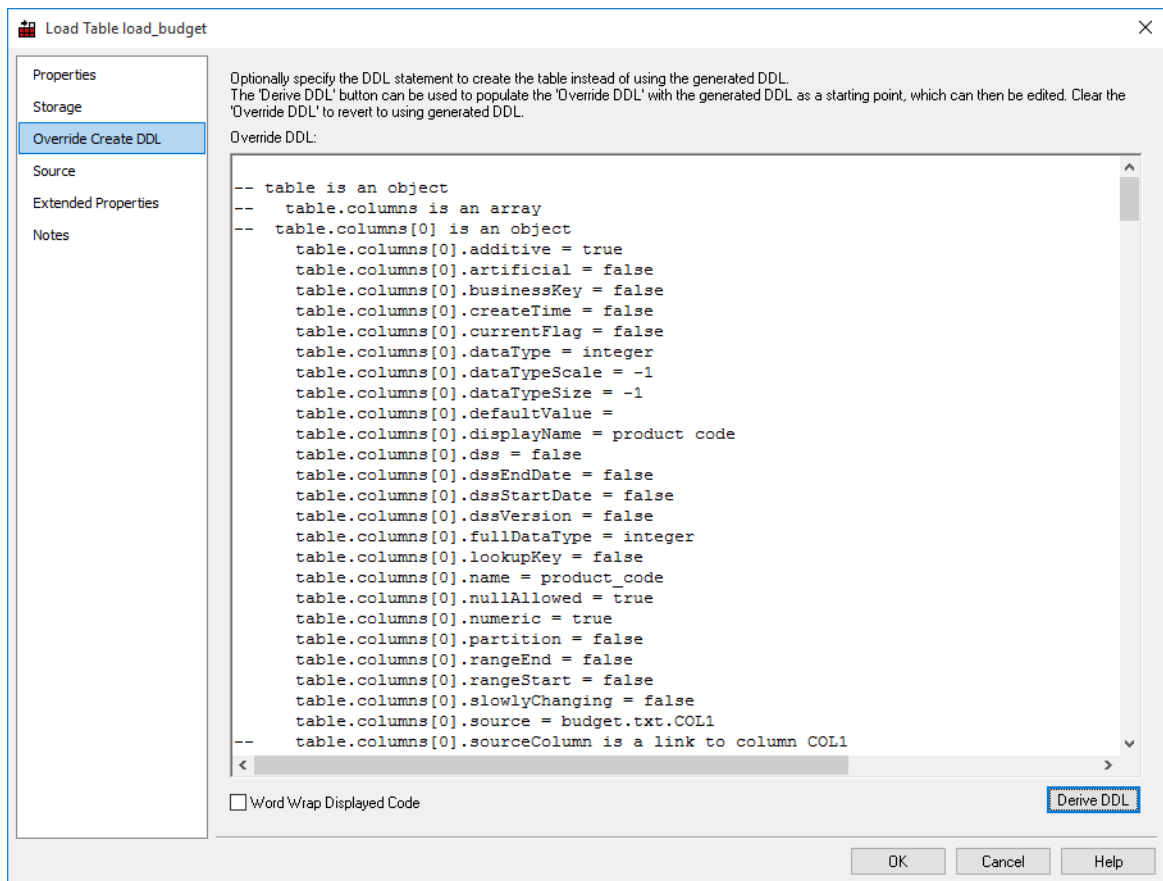
- 4 Save and close the template.
- 5 Open the Properties dialog for the **Load Table** you wish to evaluate.
- 6 In the **Storage** tab, select the template you created in the **Create DDL Template** drop-down box.



- 7 Open the **Override Create DDL** tab. If the **Override DDL** field is populated with a custom DDL statement, copy and paste this statement to a text file for backup purposes.
- 8 Click the **Derive DDL** button. A warning dialog box displays informing you that the association of the DDL template with this table will be saved to the metadata, click **OK**.



- 9 The results of the API Outline Template are printed to the **Override DDL** text box. Cut or copy this text to a text editor and save as a text file for reference purposes.



- 10 Click **Cancel** in the Load Table properties dialog.

NOTE: Ensure the Load Table properties are returned to their previous state. The default value for **Storage>Create DDL Template** is None and the **Override Create DDL>Override DDL** field is left blank to use the automatically generated DDL statement or ensure your custom DDL statement remains.

TEMPLATE USAGE

If a template exists of the correct **Type** and **Target DB**, templates can be specified and evaluated as follows:

Operation	Location to Specify the Template	Notes
Create DDL	Table Properties>Storage>Create DDL Template	DDL will be evaluated for the object at runtime if a template is specified in the Storage tab. Alternatively, clicking the Derive DDL button in the Override Create DDL tab will generate Override DDL based on the specified template. IMPORTANT: If Override DDL is specified, the Override DDL will be used at runtime.
Export Script	Export Object Properties>File Attributes>Script Template	
Load Script	Load Table Properties>Source>Script Template	
Update Procedure	When building the Update Procedure, specify the Template field on the Update Build Options>Processing tab .	If a Block template is specified, a SQL block will be generated. If a Procedure template is specified an update procedure will be generated.

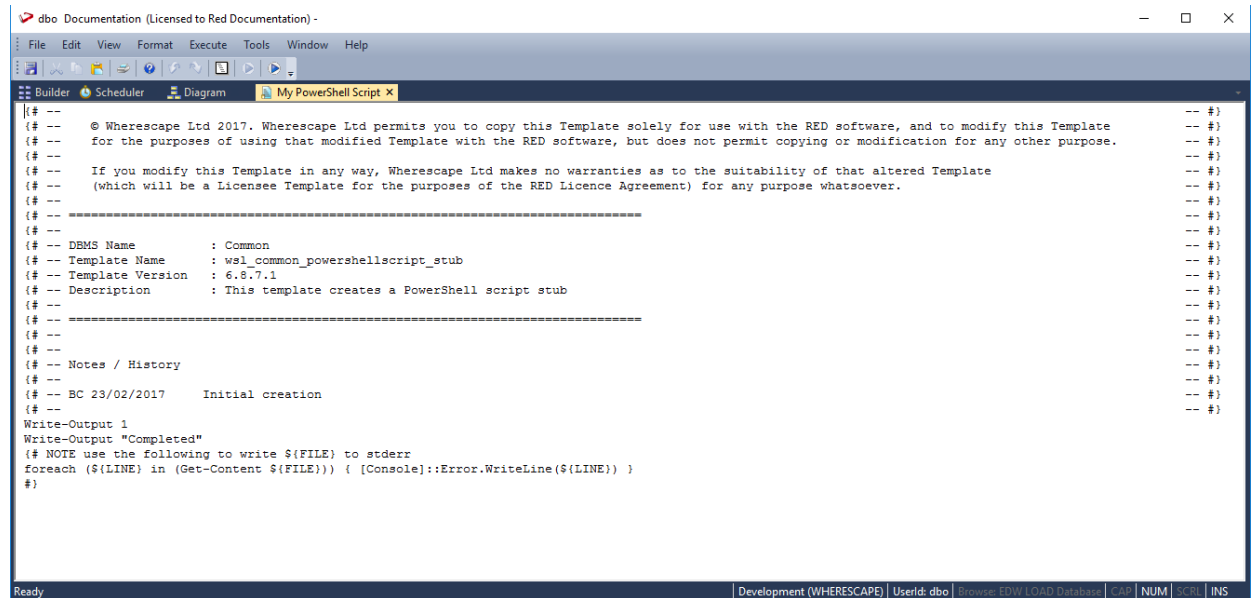
To check which Operations are supported by Templates on your Target DB, see **Templates** (on page 673).

WINDOWS POWERSHELL TEMPLATES

PowerShell Template (wsl_common_powershellscript_stub):

You can use the basic PowerShell stub template available in WhereScape RED that serves as guide on the use of a template to generate a PowerShell script.

Additional PowerShell Templates can be downloaded from the WhereScape website (<https://www.wherescape.com/support/software-downloads-documentation/templates/>).



```

|## --
|## -- © Wherescape Ltd 2017. Wherescape Ltd permits you to copy this Template solely for use with the RED software, and to modify this Template
|## -- for the purposes of using that modified Template with the RED software, but does not permit copying or modification for any other purpose.
|## --
|## -- If you modify this Template in any way, Wherescape Ltd makes no warranties as to the suitability of that altered Template
|## -- (which will be a Licensee Template for the purposes of the RED Licence Agreement) for any purpose whatsoever.
|## --
|## -----
|## -- DBMS Name          : Common
|## -- Template Name     : wsl_common_powershellscript_stub
|## -- Template Version  : 6.8.7.1
|## -- Description      : This template creates a PowerShell script stub
|## --
|## -----
|## --
|## -- Notes / History
|## --
|## -- BC 23/02/2017    Initial creation
|## --
|## --
Write-Output 1
Write-Output "Completed"
|# NOTE use the following to write $(FILE) to stderr
foreach ($LINE in (Get-Content $(FILE))) { [Console]::Error.WriteLine($LINE) }
|#
  
```

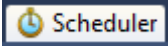
Notes:

The PowerShell stub template (wsl_common_powershellscript_stub) is listed under the Template objects pane, along with the other templates available in WhereScape RED. If this stub template is not visible after installing/upgrading WhereScape RED, use the WhereScape RED Setup Administrator to Validate the Metadata Repository. For more information, please refer to the WhereScape RED Installation Guide.

WhereScape provides a Windows application extension called WslMetadataServiceClient that can be used from Powershell templates to simplify access to RED metadata content. This DDL is accessed from WSL provided templates in custom database enablement packs. Please contact your WhereScape customer representative, if you require details on how to use this DDL.

CHAPTER 26

SCHEDULER

The scheduler is accessible by clicking the  Scheduler button from the toolbar. It is also available as a stand alone utility. In this way operators can be given access to the scheduler without gaining full access to the data warehouse.

The scheduler runs under Windows (as a system service). It processes pre-defined jobs, recording the success, failure or otherwise of the job.

Audit trail

Specific information relating to the tasks in the job are recorded to the audit trail. Generally only summary information is written to this audit trail. The contents of the audit trail are maintained even after a job is deleted.

Error trail

Detail or error information is written to the error trail. The contents of the error trail are deleted when the job is deleted.

Administration Views

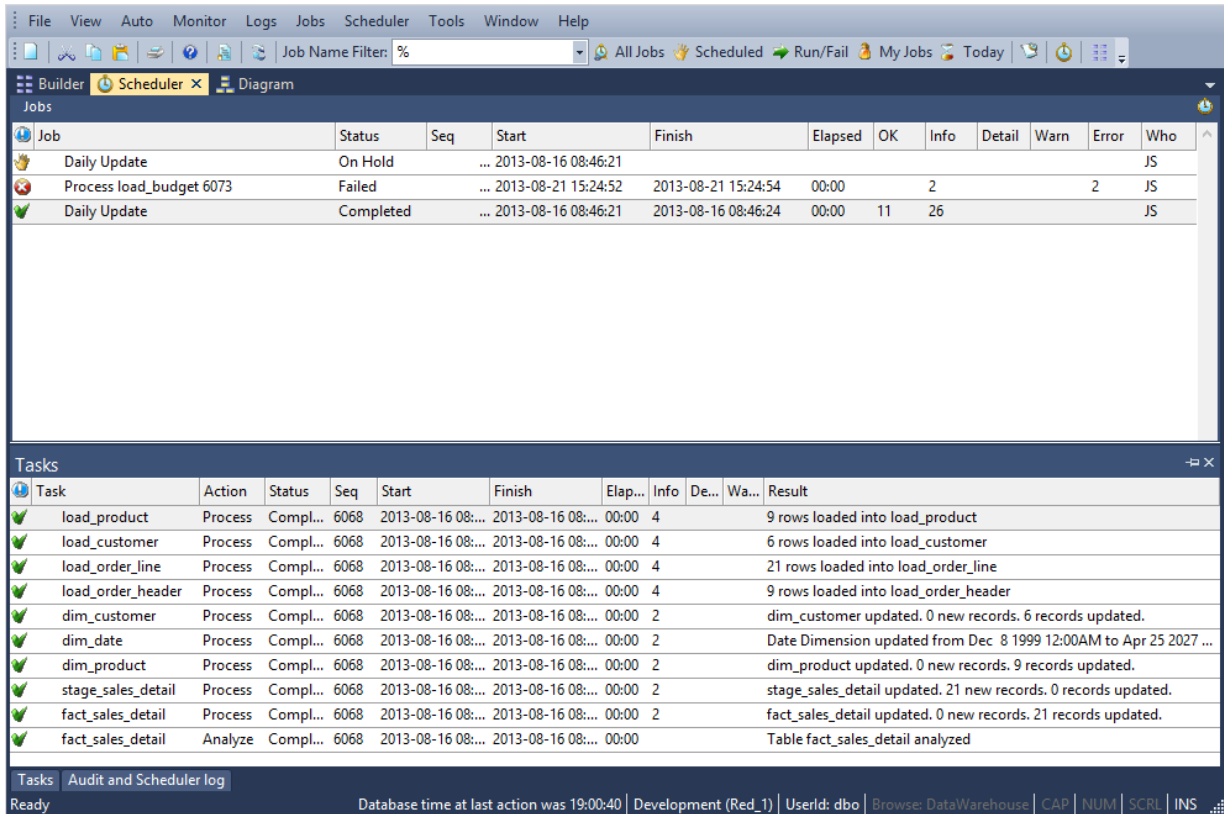
It is possible to view the status of a job without using the WhereScape RED product. Three views are provided to assist in this undertaking. They are `ws_admin_v_audit`, `ws_admin_v_error` and `ws_admin_v_sched`. Queries can be issued using these views to see the results or status of a job.

IN THIS CHAPTER

Scheduler Options	686
Scheduler States	691
Scheduling a Job	693
Working with Jobs.....	699
Stand Alone Scheduler Maintenance.....	745
SQL to return Scheduler Status	748
Reset Columns in Job and Task View.....	749
Stopping a Linux/UNIX Scheduler from within RED	750

SCHEDULER OPTIONS

An example of the **Scheduler** screen is shown below.



Toolbar/Jobs menu

Quick access to some job categories are in the toolbar. The complete options are listed under the Jobs menu and while most are self-explanatory they are described below:

Object	Description
All jobs	All jobs are listed in the middle pane
Scheduled jobs	In the middle pane lists those jobs that are waiting to run or are on hold.
Last 24 hours	Lists the jobs that have run or started to run during the last 24 hours.
Prior 24 hours	Lists the jobs that ran during the previous 24 hours.
This weeks Jobs	Lists the jobs that have run or are scheduled to run during the current week
Last weeks Jobs	Lists the jobs that ran during the last week.
My Jobs	Lists the jobs you have scheduled or have run.

Object	Description
Job Name Filter	Lists the jobs whose names match the filter supplied. This filter only works as an appended filter to the main filter selected under Jobs , e.g first enter a filter for Job Name Filter or select a filter from the drop-down list; and then choose your main filter under Jobs , e.g. All Jobs, Scheduled Jobs etc.
Recent audit trail Today's audit trail	Provides listings from the audit trail. The information is useful when a job fails to start or enters some other unknown state. Generally the audit trail entries for a job can be found by drilling down into the job itself.
Scheduler status	Lists all schedulers and displays their current status. The status is updated every few minutes, and a right menu option allows the polling of a scheduler for status, and the termination of a scheduler.

Top pane

The top pane shows the details of the jobs. Information covers:

Column	Description
Job name	The name given to the job when created.
Status	The status of the job. Refer to the following section for the various status values.
Sequence	This is a unique number assigned to each job iteration and job. If you enter a new job it will acquire a new sequence number. In normal daily processing when no new jobs have been created sequence numbers will be sequential.
Start and Finish times	As the names suggest, the start and finish dates and times for the job.
Elapsed time	The time that elapses from start to finish of a job.
Ok	These are success messages written to the audit trail.
Inf	Informational messages written to the audit trail about the the running of the job.
Det	Lines written to the detail or error logs
Warn	The number of warnings written to the audit trail.
Err	The number of error messages written to the audit trail.
Who	The initials of the person who scheduled the job.

Additional fields can be added via the **Tools>Select Job Report Fields** option

Middle pane

The middle pane shows the tasks related to a selected job. Task information available includes:

Column	Description
Task	The object name
Action	The action to be done to the object
Status	Status of the task
Sequence	This is a unique number assigned to each job iteration and job. If you enter a new job it will acquire a new sequence number. In normal daily processing when no new jobs have been created sequence numbers will be sequential.
Start and Finish times	As the names suggest, the start and finish dates and times for the task
Elapsed time	The time that elapses from start to finish of a task.
Info	Informational messages written to the audit trail about the the running of the job.
Detail	Lines written to the detail or error logs
Warning	The number of warnings written to the audit trail.
Result	Result of the task

Additional fields can be added via the **Tools>Select Task Report Fields** option

Bottom pane


The bottom pane shows the audit trail of a selected task/job. Audit trail information includes:

Column	Description
Task	The object name
Status	The status of the message
Sequence	This is a unique number assigned to each job iteration and job. If you enter a new job it will acquire a new sequence number. In normal daily processing when no new jobs have been created sequence numbers will be sequential.
Timestamp	Time of message output
Message	The message
DB Message	Message from database
Job	Job relating to this message

AUTO


Auto Refresh

Use auto refresh to automatically refresh all jobs.



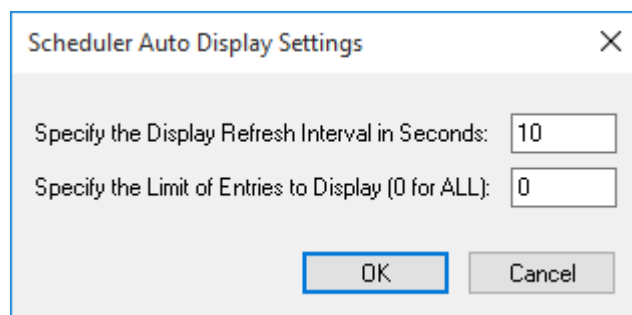
Set Refresh Options

Use this setting to control the maximum number of rows that are displayed via the Auto Refresh option as well as the display refresh interval.



When clicking the Set Refresh Options, the settings dialog allows adding a display limit. The display of jobs when on auto refresh will stop at the selected number of rows, so when set to 100, the refresh will stop after first 100 rows (jobs) are returned.

If 0 is selected in the Specify the Limit of Entries to Display, then all rows (jobs) will be displayed.




TOOLS

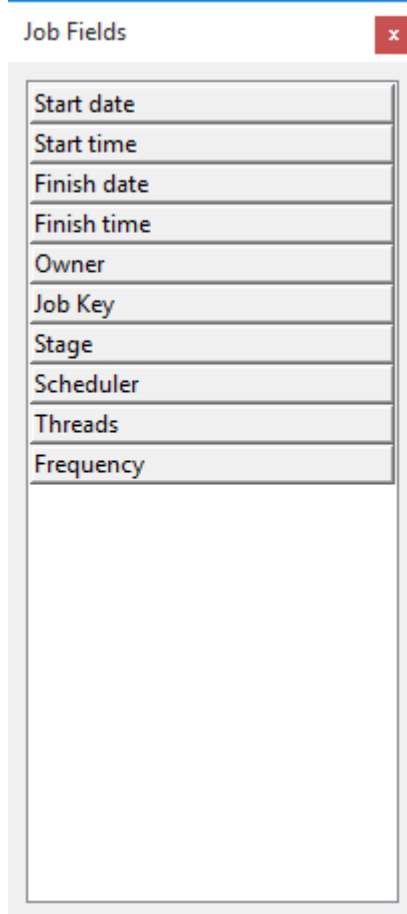
SELECT JOB REPORT FIELDS

The **Select Job Report Fields** menu option in the scheduler pane enables users to select some extra fields such as Scheduler, Threads and Frequency fields into the job report.

- 1 To make these fields available, click **Tools->Select Job Report Fields** from the top pane.



- 2 Select the fields you want to add to your job report from the menu and drag them to where you want to place them on the report.



NOTE: The "Frequency" field is only populated for scheduled jobs. If selected, it will return blank results for running or completed jobs.

SCHEDULER STATES

A scheduled **job** can have the following states:

- Hold
- Waiting
- Blocked
- Pending
- Running
- Failed
- Failed - Aborted
- Completed

State	Description
Hold	The job is on hold. It can be edited and its state changed in order to release the job.
Waiting	The job is waiting to start, or waiting for its scheduled time to arrive, or waiting for a scheduler to become available.
Blocked	The job is blocked as a previous instance of the same job is still running.
Pending	This is the first stage of a running job. The scheduler has identified the job as ready to start and has allocated a thread, or sub task to process the job. A job is in this state until the thread or sub tasks begins processing. If a job stays in this state then the scheduler thread has failed for some reason. The logs can be checked on either the Windows server on which the scheduler is running.
Running	The job is currently running. Double-click on the job name in the right pane to drill down into the specific tasks.
Failed	A failed job is one that had a problem. It can be restarted from the point of failure and is considered to be running unless subsequently aborted.
Failed - Aborted	The job has been aborted after having failed. Once in this state a job cannot be restarted. The job exists then only as a log of what occurred and is no longer regarded as a job.
Completed	The job has successfully completed, possibly with warnings. Once in this state, a job cannot be restarted. The job exists then only as a log of what occurred and is no longer regarded as a job.

Note: When a job fails and drilling down does not show any errors against the tasks, right-click on the job and **View Audit Trail**. The job may have failed because of an error in the JOB level.

A scheduled **task** can have the following states:

- Waiting or Blank
- Held
- Running
- Failed
- Completed
- Error Completion
- Bad Return Status

State	Description
Held	The task has been held due to a prior dependency failure. The problem must be rectified and the job restarted.
Waiting (Blank)	Tasks that are waiting to run either due to a shortage of threads, or prior dependencies normally have a blank status.
Running	The task is currently running.
Failed	The task has had a fatal error. Any dependencies on this task will be held. Double click on the task to see more detail error information or review the audit and error/detail log for the job.
Completed	The task has completed successfully.
Error Completion	The task has completed with a handled Error. Any dependent tasks will be held, and the job must be restarted when the problem is rectified.
Bad Return Status	The task has returned an unknown status. This normally occurs with script files that produce unexpected information. The rule for scripts is that the first line returned must be a status of either 1, -1, -2, or -3. The second line is a message detailing the result. If the first line does not contain one of these four values, then this status will be returned and dependent tasks held. Run the script manually to view the output or check the logs.

SCHEDULING A JOB

To schedule a job

Firstly access the scheduler by clicking the **Scheduler** button  on the toolbar.

Select **File | New Job** from the menu strip at the top of the screen, or click the **New Job** button on the toolbar.



The following **Job Definition** dialog is displayed.

Job Definition
✕

Job Name:

Description:

Frequency:

Start Date:

Start Time:

Maximum Threads:

Scheduler:

Dependent On:

Logs Retained: This field lets you set the number of logs that are retained for this job before an automatic delete and archive occurs. 0 = keep all logs (default action)

The following two fields are optional. They are executed after the job completes and therefore need to reflect the scheduler environment. (i.e. Unix or Windows). The special variables \$JOB_KEY\$, \$JOB_SEQ\$ and \$JOB_NAME\$ can be used to return the associated values.
The Success command will be executed if a successful completion, the failure command will be executed if a job fails to complete:

Success Command:

Failure Command:

Execute Failure Command in event of dependency failure

Custom Settings

Interval Between Jobs: (minutes)

Start At or After HHMM (e.g. 0800):


Do Not Start After HHMM (e.g. 1700):

Active Days:

Mon Tue Wed Thu Fri Sat Sun

See **Creating a Job** (on page 701) for more details on how to create a job.

Once the job has been created, click on the **All Jobs** button on the toolbar. The newly created job will now be displayed in the scheduler window.



To create a job within a job

It is possible to schedule one job from another job. There are however some limitations and rules that must be understood when doing this.

- 1 A job that is called from another job is only ever allocated one thread. All tasks within the called job will therefore run sequentially.
- 2 A job can only have one running iteration. Therefore, a called job will be blocked if that job is already running independently or as part of another job.
- 3 Any job dependencies for the called job are ignored. The parent's job dependencies are the only ones that are used.
- 4 A called job essentially runs as a separate job, so that if it fails both it and the parent job will show in a failed state. Once the problem is fixed the parent job should be restarted which will restart the called job.

To create a job dependency

It is possible to make a job dependent on another job, using the **Dependent On** field in the **Job Definition** dialog.

The screenshot shows the 'Job Definition' dialog box with the following fields and settings:

- Job Name:** Enterprise Reporting Daily Refresh
- Description:** Full Daily Refresh of all Presentation Tables for the Enterprise Reporting Layer
- Frequency:** Custom
- Start Date:** Thursday, March 3, 2016
- Start Time:** 3:00:00 AM
- Maximum Threads:** 4
- Inactive Wait Interval (seconds):** 30
- Scheduler:** Windows Only
- Dependent On:** Parent job (highlighted in red)
- Fail:** Fail
- Look back (minutes):** (empty)
- Maximum wait (minutes):** (empty)
- Buttons:** Add Parent Job, Remove Parent
- Custom Settings:**
 - Interval Between Jobs: (minutes) 1440
 - Start At or After HHMM (e.g. 0800): 0300
 - Do Not Start After HHMM (e.g. 1700): 0500
 - Active Days: Mon Tue Wed Thu Fri Sat Sun
- Logs Retained:** 0
- Success Command:** (empty)
- Failure Command:** (empty)
- Buttons:** OK, Cancel

The following two fields are optional. They are executed after the job completes and therefore need to reflect the scheduler environment. (i.e. Unix or Windows). The special variables \$JOB_KEY\$, \$JOB_SEQ\$ and \$JOB_NAME\$ can be used to return the associated values.
The Success command will be executed if a successful completion, the failure command will be executed if a job fails to complete:

Click on the **Add Parent Job** button.

Job Definition
✕

Job Name:

Description:

Frequency: Custom

Start Date: Thursday, March 3, 2016

Start Time: 3:00:00 AM

Maximum Threads: Inactive Wait Interval (seconds):

Scheduler: Windows Only

Dependent On: Parent job Fail Look back (minutes) Maximum wait (minutes) Add Parent Job Remove Parent

Logs Retained: This field lets you set the number of logs that are retained for this job before an automatic delete and archive occurs. 0 = keep all logs (default action)

The following two fields are optional. They are executed after the job completes and therefore need to reflect the scheduler environment. (i.e. Unix or Windows). The special variables \$JOB_KEY\$, \$JOB_SEQ\$ and \$JOB_NAME\$ can be used to return the associated values.
The Success command will be executed if a successful completion, the failure command will be executed if a job fails to complete:

Success Command:

Failure Command:

OK
Cancel

In the dialog that follows, select the Parent Job from the drop-down list. In our case we will choose the job **Shared Dimensions Daily Refresh**.

Add a Parent Job [X]

Select a Job on which this Job is dependent. Specify the amount of time to look back for the Parent Job completion. Normally this would be in the 8 to 10 hour range to prevent seeing a previous days run.

Next specify a maximum time to wait for this Parent Job to complete. A failure will occur if the Parent Job does not complete in the specified time.

Parent Job:

Shared Dimensions Daily Refresh [v]

Maximum Time to Look Back for the Parent Job Completion (minutes):

60

Maximum Time to Wait for the Parent Job to Complete (minutes):

20

Fail if Parent Job Does Not Complete Successfully in the Required Time:

Add Cancel

The **Maximum Time to Look Back for the Parent Job Completion** field prevents older iterations of the parent job as being identified as a completion. In our example, we are starting both jobs at 3am, so we don't need to look too far back to ensure that the dimension refresh has run. We have therefore set the look back minutes to 60 to allow for any delays in starting this job.

The **Maximum Time to Wait for the Parent Job to Complete** specifies how long to await a successful completion of the parent job. In our example we know that the dimension refresh only takes a few minutes, but we should allow for the occasional slow network or resource drains making the dimension refresh take longer; so we have set the maximum wait to 20 minutes. This means that our job will wait 20 minutes from its own scheduled start time for the parent job to complete.

The checkbox to fail if the parent job does not complete in time will prevent this job from running if the parent job (dimension refresh) does not complete successfully. As we do not wish for the transactional data in our fact deliveries to be flagged with 'Unknown' dimensional item(s), we can leave this checkbox checked to ensure that this job does not run.

Click **Add**.

NOTE: Clearing the checkbox to fail if the parent fails will simply ensure that this job waits for the completion of the dimension refresh and, irrespective of the dimensions refresh's success or failure, starts.

Click **OK** to link this data job to the parent dimensional job. In this way, the job **Enterprise Reporting Daily Refresh** cannot run until the parent job **Shared Dimensions Daily Refresh** has completed successfully; thus the facts will have the latest dimensional keys associated with them.

Job Definition
✕

Job Name:

Description:

Frequency: Custom ▾

Start Date:

Start Time:

Maximum Threads: Inactive Wait Interval (seconds):

Scheduler: Windows Only ▾

Parent job	Fail	Look back (minutes)	Maximum wait (minutes)	
Shared Dimensions Daily Refresh	Y	60	20	<input type="button" value="Add Parent Job"/> <input type="button" value="Remove Parent"/>

Logs Retained: This field lets you set the number of logs that are retained for this job before an automatic delete and archive occurs. 0 = keep all logs (default action)

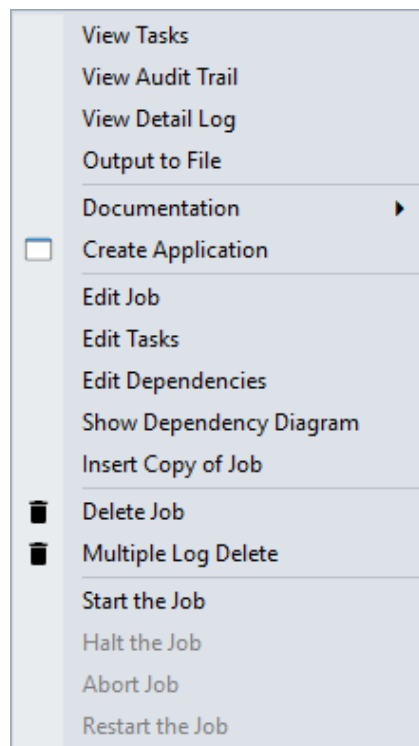
The following two fields are optional. They are executed after the job completes and therefore need to reflect the scheduler environment. (i.e. Unix or Windows). The special variables \$JOB_KEY\$, \$JOB_SEQ\$ and \$JOB_NAME\$ can be used to return the associated values.
The Success command will be executed if a successful completion, the failure command will be executed if a job fails to complete:

Success Command:

Failure Command:

WORKING WITH JOBS

When positioned on a Job in the scheduler window, the right-click pop-up menu provides a number of options for working with the job. Some of the options are discussed in more detail in the chapters that follow, however a brief overview of the menu options follows:



The **View Tasks** menu option enables you to view the tasks of a job.

The **View Audit Trail** option enables you to view the audit trail of a job.

The **View Detail Log** option enables you to view a detail log of a job.

The **Export to CSV File** option enables you to export a job to a CSV file.

The **Documentation** option enables you to create documentation for a job.

The **Edit Job** option enables you to edit a job. See *Editing a Job* (on page 711)

The **Edit Tasks** option enables you to edit the tasks of a job. See *Editing Tasks* (see "*Editing Tasks in a Job*" on page 715)

The **Edit Dependencies** option enables you to edit the task dependencies of a job. See *Editing Dependencies* (see "*Editing Task Dependencies*" on page 722)

The **Insert Copy of Job** option enables you to insert a copy of a job. See *Inserting a Copy of a Job* (on page 729)

The **Delete Job** option enables you to delete a job. See *Deleting a Job* (on page 731)

The **Multiple Log Delete** option enables you to delete multiple logs of a job. See *Deleting Job Logs* (on page 732)

The **Start the Job** option enables you to start a job. See *Starting a Job* (on page 734)

The **Halt the Job** option enables you to halt a job. See *Halting a Job* (on page 735)


The **Abort the Job** option enables you to abort a job. See *Aborting a Job* (on page 736)

The **Restart the Job** option enables you to restart a job. See *Restarting a Job* (on page 737)


CREATING A JOB

To create a job

Click on the **Scheduler** tab to open the scheduler window.



Click on the **New Job** button to create a new job.



A **Job Definition** dialog is displayed.

Job Definition ✕

Job Name:

Description:

Frequency:

Start Date:

Start Time:

Maximum Threads:

Scheduler:

Dependent On:

Logs Retained: This field lets you set the number of logs that are retained for this job before an automatic delete and archive occurs. 0 = keep all logs (default action)

The following two fields are optional. They are executed after the job completes and therefore need to reflect the scheduler environment. (i.e. Unix or Windows). The special variables \$JOB_KEY\$, \$JOB_SEQ\$ and \$JOB_NAME\$ can be used to return the associated values.
 The Success command will be executed if a successful completion, the failure command will be executed if a job fails to complete:

Success Command:

Failure Command:

Execute Failure Command in event of dependency failure

Complete the fields and click **OK**. The main fields are described in the following table:

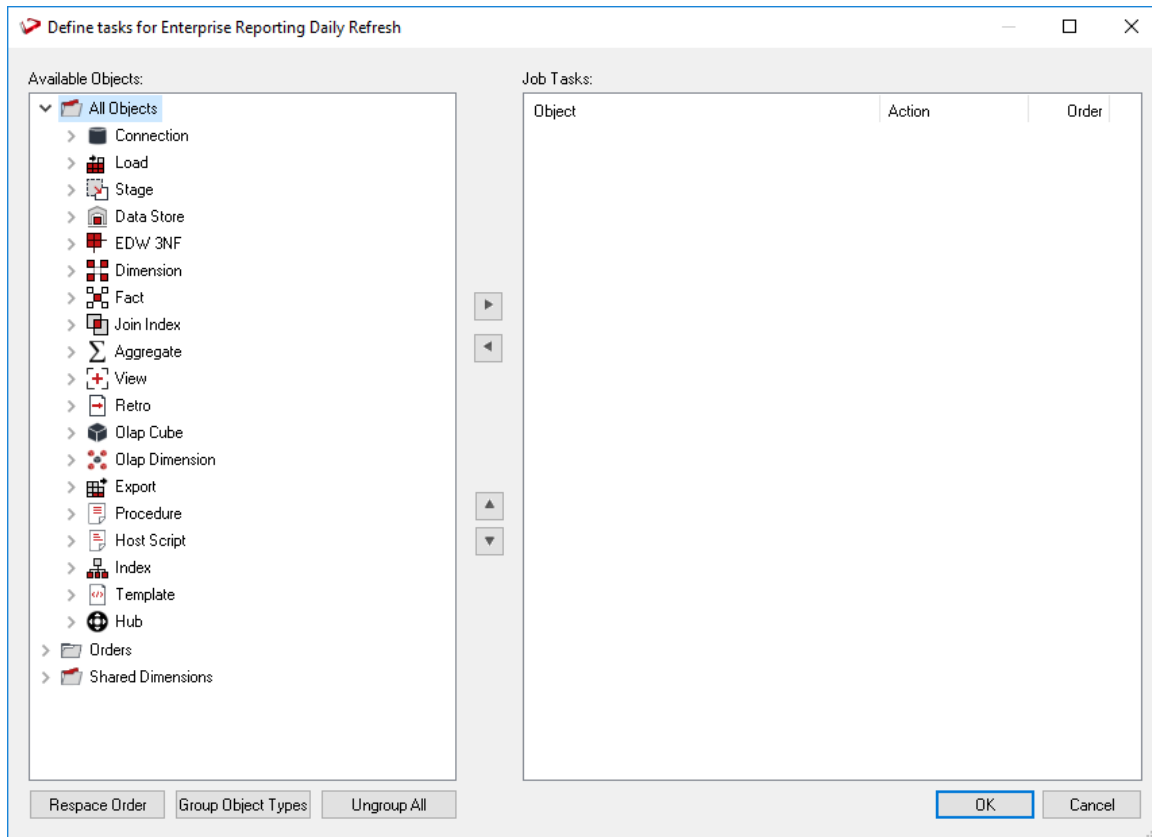
Field	Description
Job Name	The Scheduler defaults to the next job number in the sequence. You can alter this to any alphanumeric. Note: Only alphanumerics, spaces and the underscore are supported in the name.
Description	A description of the job
Frequency	When the job runs. The options available in the drop-down list are: <ul style="list-style-type: none"> • Once Only - job is deleted on completion • Once and Hold - runs and puts another copy of the job on hold • Hold - puts the job on hold for manual release • Daily - runs the job daily • Custom - enables custom definition • Weekly - runs the job weekly • Monthly - runs the job monthly • Annually - runs the job annually
Start Date and Start Time	<ul style="list-style-type: none"> • The date and time for the job to start.
Max Threads	The maximum number of threads allocated to run the job, e.g. if some tasks can run in parallel then if more than one thread is allocated then they will run in parallel.
Scheduler	It is possible to have multiple schedulers running. Select the desired scheduler from this drop-down. The valid options are: Windows Preferred, Windows Only, or the name of a specific scheduler can be entered (e.g. WIN0002)
Dependent On	A job can be dependent on the successful completion of one or more other jobs. Click the Add Parent Job button to select a job that this job will be dependent on. The maximum time to look back for parent job completion field prevents older iterations of the parent job as being identified as a completion. The maximum time to wait specifies how long to await a successful completion of the parent job. The action if that wait expires can also be set. See the Dependency example in <i>Scheduling a Job</i> (on page 693)
Logs Retained	Specify the number of logs to retain for the job. By default all logs are retained. This field can be used to reduce the build up of scheduler logs by specifying a number of logs to retain.

Field	Description
Okay command and Failure command	<p>These are Windows shell commands depending on which scheduler is used. They are executed if the condition is met. Typically, these commands would mail or page on success or failure.</p> <p>Notes:</p> <ol style="list-style-type: none"> 1. The RED scheduler does not check return codes from called commands, scripts and programs. 2. It is recommended that all output from commands, scripts and programs is redirected to a log file. For example, add this to the end of any SUCCESS/FAILURE commands: >> c:\scheduler\success_failure_prod.log 2>&1

The following fields are available if a frequency of **Custom** is chosen:

Field	Description
Interval between jobs (Minutes)	Specify the number of minutes between iterations of the job. For example to run a job every 30 minutes set this value to 30. If a job is to run only once but on selected days set this value to 1440 (daily)
Start at or after HHMM	The time that the job may run from. To run anytime set to 0000.
Do not start after HHMM	If multiple iterations are being done then this is the time after which a new iteration will not be started. For example if a job is running every 10 minutes it will continue until this time is reached. To run till the end of day set to 2400.
Active on the days	Select each day of the week that the custom job is to be active on.

Once the job itself has been defined, tasks then need to be added to the job. The **Define tasks** window is shown below.



The screen has two main areas. The right pane shows the tasks to be run for this job and the left pane lists all the objects.

To add a task

Double click on an object to add it to the left pane. Normally objects such as load or fact tables are scheduled rather than procedures.

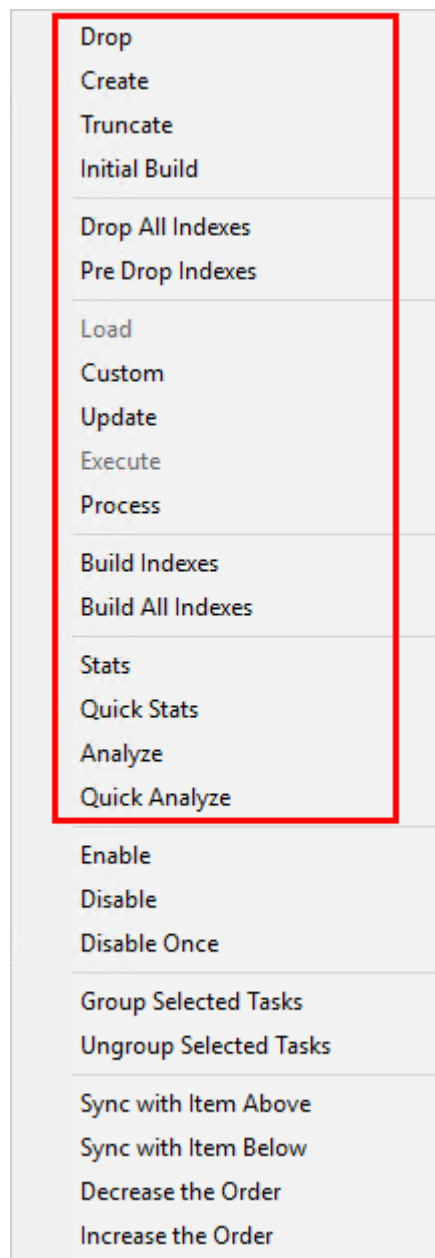
To set the action on a task

Each task can have a specific action that is to be performed on its object.

The default action for **load tables** is **process**. This means that when the task is actioned it will drop any indexes that are due to be dropped, or have **pre-drop** set, then load the table and perform any post-load procedures or transformations and then re-create any dropped indexes.

The default action for all other tables is the same as above, except it will execute the **update** procedure rather than loading the table.

You can change the action on a task by right-clicking on the task in the right pane. The menu options are shown below.



The following task actions are available:

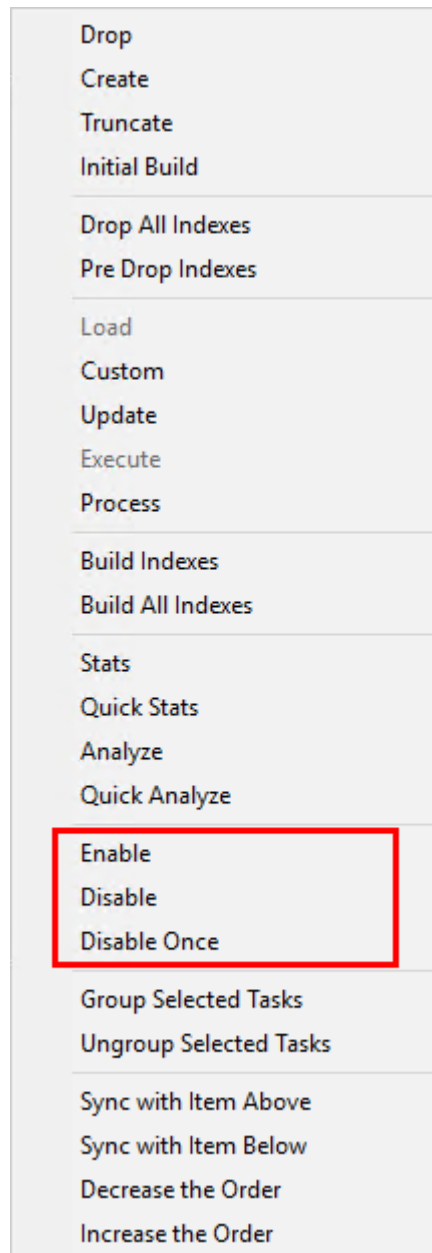
Action	Description
Drop	Drop table, view, join index or index.
Create	Create table, view, join index or index.
Truncate	Delete all rows from the table.
Initial Build	Drop All Indexes then Custom then Build All Indexes .

Action	Description
Drop All Indexes	Drop all indexes on the table.
Pre Drop Indexes	Drop all indexes on the table marked as "Pre Drop".
Load	Load the table (Load tables only).
Custom	Run the custom procedure on the table.
Update	Run the update procedure on the table.
Execute	Execute the procedure or host scripts.
Process	Pre Drop Indexes then Update and then Build Indexes .
Process and Statistics	Process then Default Stats as defined on Table Properties/Statistics/Process and statistics method (DB2 only).
Build Indexes	Build the indexes on the table marked as "Pre Drop".
Build All Indexes	Build all indexes on the table.
Stats	Refreshes predefined statistics on a table or index: COLLECT STATISTICS ON DatabaseName.TableName; COLLECT STATISTICS ON DatabaseName.TableName INDEX IndexName;
Quick Stats	Refreshes predefined statistics on an index using sampling: COLLECT STATISTICS USING SAMPLE ON DatabaseName.TableName INDEX IndexName;
Analyze	Refreshes predefined statistics on a table or index: COLLECT STATISTICS ON DatabaseName.TableName; COLLECT STATISTICS ON DatabaseName.TableName INDEX IndexName;
Quick Analyze	Refreshes predefined statistics on an index using sampling: COLLECT STATISTICS USING SAMPLE ON DatabaseName.TableName INDEX IndexName;

Note: Not all actions are available on all object types.

To set the state of a task

Each task can be set to a state:

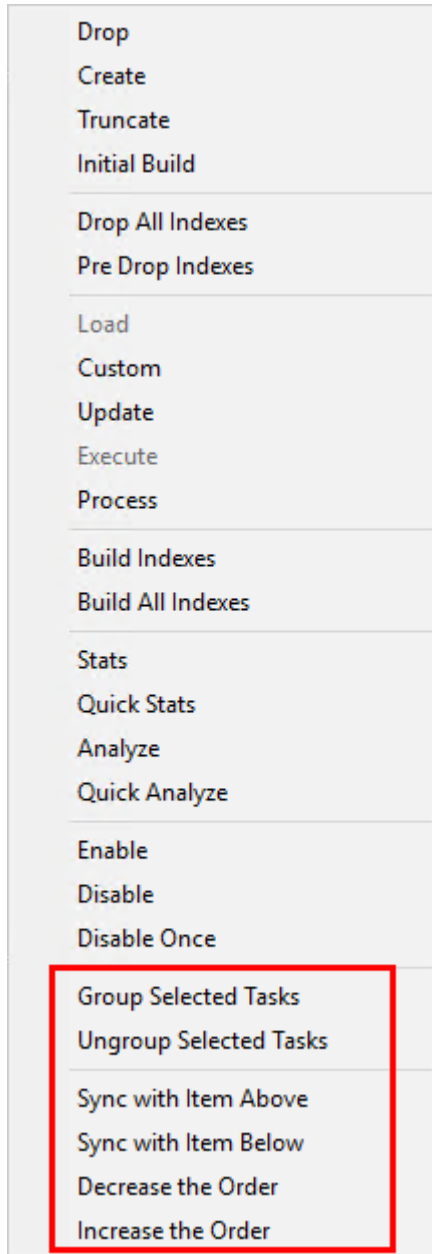


The following states are available:

State	Description
Enable	Job Task is enabled.
Disable	Job Task is disabled.
Disable Once	Job Task is disabled once and reverts to enabled next time the Job is released by the Scheduler.

To create dependencies between tasks

It is possible to create dependencies between tasks in the list by selecting one or more tasks and right-clicking to bring up the dependency options.





The following task dependency options are available from the menu:

Task Option	Description
Group Selected Tasks	Groups two or more selected tasks to have the same order value, allowing them to run in parallel if the maximum threads setting allows.
Ungroup Selected Tasks	Un-group selected tasks.

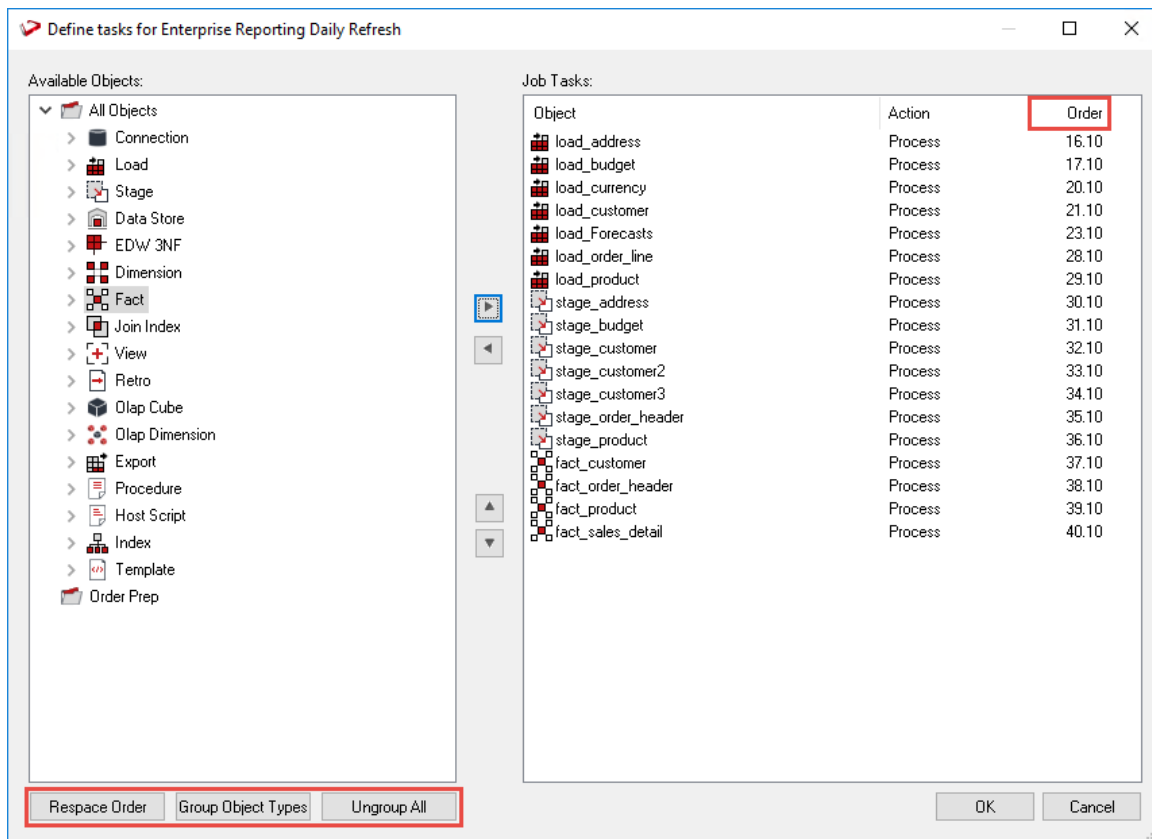
Task Option	Description
Sync with Item Above	Changes a selected task to have the same order value as the task above it, allowing them to run in parallel if the maximum threads setting allows.
Sync with Item Below	Changes a selected task to have the same order value as the task below it, allowing them to run in parallel if the maximum threads setting allows.
Decrease the Order	Changes a selected task to an order number one less than its current value. The task will now run immediately before it would have previously.
Increase the Order	Changes a selected task to an order number one more than its current value. The task will now run immediately after it would have previously.

To order or group the tasks

The **Order** column shows the order in which the tasks are to be run, e.g. 20.20 If the two numbers are the same as another task then those tasks can run in parallel. If the two numbers are different then those tasks run sequentially. This is an initial definition of dependencies. These dependencies can be altered specifically once the job has been created.

Tasks can be moved up or down by selecting the task and clicking the **Move Up**  or **Move Down**  buttons.

To re-space the order of the tasks; to group or un-group object types, use the **buttons** at the bottom of the **Define tasks** dialog.



- **Respace Order**

This button will respace the order numbers. The existing dependency structure and groupings are retained. The purpose of this button is simply to allow room between tasks to fit new tasks. So for example if we have two tasks that have an order of 20.19.5 and 20.20.6 and we want to add a task between these two tasks we can click the **Respace Order** button to open up a gap between the two tasks.

- **Group Object Types**

This option will put all objects of the same type into groups. For example all load tables will be able to run in parallel, all dimensions etc.

- **Ungroup All**

This button will remove all groupings and make all tasks sequential. New groupings can be made by selecting a range of sequentially listed tasks in the left pane and using the right-click menu option **Group Selected Tasks**. Tasks that are grouped have the same first two numbers in the order and can execute at the same time if the job has multiple threads.

Upon completion of adding tasks, click **OK**.

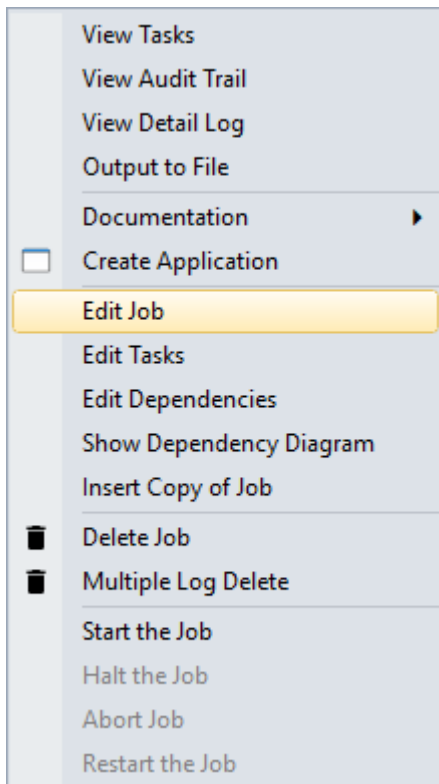
EDITING A JOB

Once jobs have been created they can be edited.

Note: A job can only be edited when it is not in a running state and only if the job is a scheduled job. Completed jobs remain in the list but only logs remain.

To edit a job

Select the job from the scheduler middle pane. Right-click on the job and select **Edit Job** from the drop-down list.



The **Job Definition** is displayed.

Edit the fields as required and click **OK**. The main fields are described in the following table:

Field	Description
Job Name	The Scheduler defaults to the next job number in the sequence. You can alter this to any alphanumeric. Note: Only alphanumeric, spaces and the underscore are supported in the name.
Description	A description of the job

Field	Description
Frequency	<p>When the job runs. The options available in the drop-down list are:</p> <ul style="list-style-type: none"> • Once Only - job is deleted on completion • Once and Hold - runs and puts another copy of the job on hold • Hold - puts the job on hold for manual release • Daily - runs the job daily • Custom - enables custom definition • Weekly - runs the job weekly • Monthly - runs the job monthly • Annually - runs the job annually
Start Date and Start Time	<ul style="list-style-type: none"> • The date and time for the job to start.
Max Threads	<p>The maximum number of threads allocated to run the job, e.g. if some tasks can run in parallel then if more than one thread is allocated then they will run in parallel.</p> <p>Caution: Users must be aware that adding additional threads to a job increases the communication overhead with the DB server containing the Data Warehouse repository. Each job thread holds a dedicated connection to the server for its lifetime.</p>
Scheduler	<p>It is possible to have multiple schedulers running. Select the desired scheduler from this drop-down. The valid options are:</p> <p>Windows Preferred, Windows Only, or the name of a specific scheduler can be entered (e.g. WIN0002)</p>
Dependent On	<p>A job can be dependent on the successful completion of one or more other jobs. Click the Add Parent Job button to select a job that this job will be dependent on. The maximum time to look back for parent job completion field prevents older iterations of the parent job as being identified as a completion. The maximum time to wait specifies how long to await a successful completion of the parent job. The action if that wait expires can also be set.</p> <p>See the Job Dependency example in <i>Scheduling a Job</i> (on page 693)</p>
Logs Retained	<p>Specify the number of logs to retain for the job. By default all logs are retained. This field can be used to reduce the build up of scheduler logs by specifying a number of logs to retain.</p>

Field	Description
Okay command and Failure command	<p>These are Windows shell commands depending on which scheduler is used. They are executed if the condition is met. Typically, these commands would mail or page on success or failure.</p> <p>Notes:</p> <ol style="list-style-type: none"> 1. The RED scheduler does not check return codes from called commands, scripts and programs. 2. It is recommended that all output from commands, scripts and programs is redirected to a log file. For example, add this to the end of any SUCCESS/FAILURE commands: <pre>>> c:\scheduler\success_failure_prod.log 2>&1</pre>

The following fields are available if a frequency of **Custom** is chosen:

Field	Description
Interval between jobs (Minutes)	Specify the number of minutes between iterations of the job. For example to run a job every 30 minutes set this value to 30. If a job is to run only once but on selected days set this value to 1440 (daily)
Start at or after HHMM	The time that the job may run from. To run anytime set to 0000.
Do not start after HHMM	If multiple iterations are being done then this is the time after which a new iteration will not be started. For example if a job is running every 10 minutes it will continue until this time is reached. To run till the end of day set to 2400.
Active on the days	Select each day of the week that the custom job is to be active on.

EDITING TASKS IN A JOB

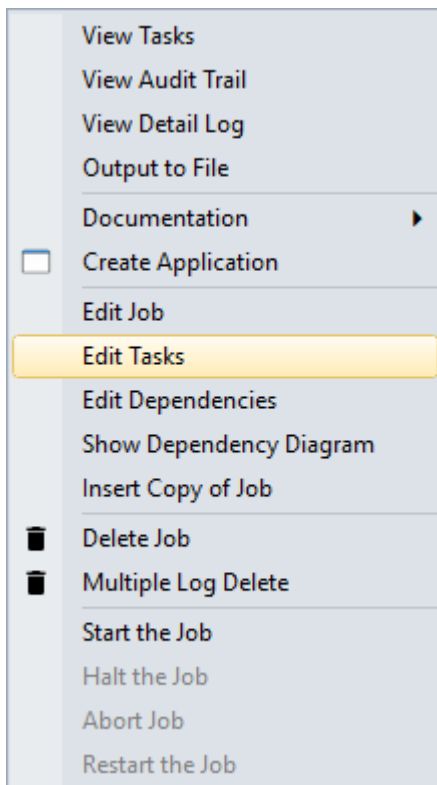
Once jobs have been created, you can edit their tasks.

Note: A job can only be edited when it is not in a running state and only if the job is a scheduled job. Completed jobs remain in the list but only logs remain.

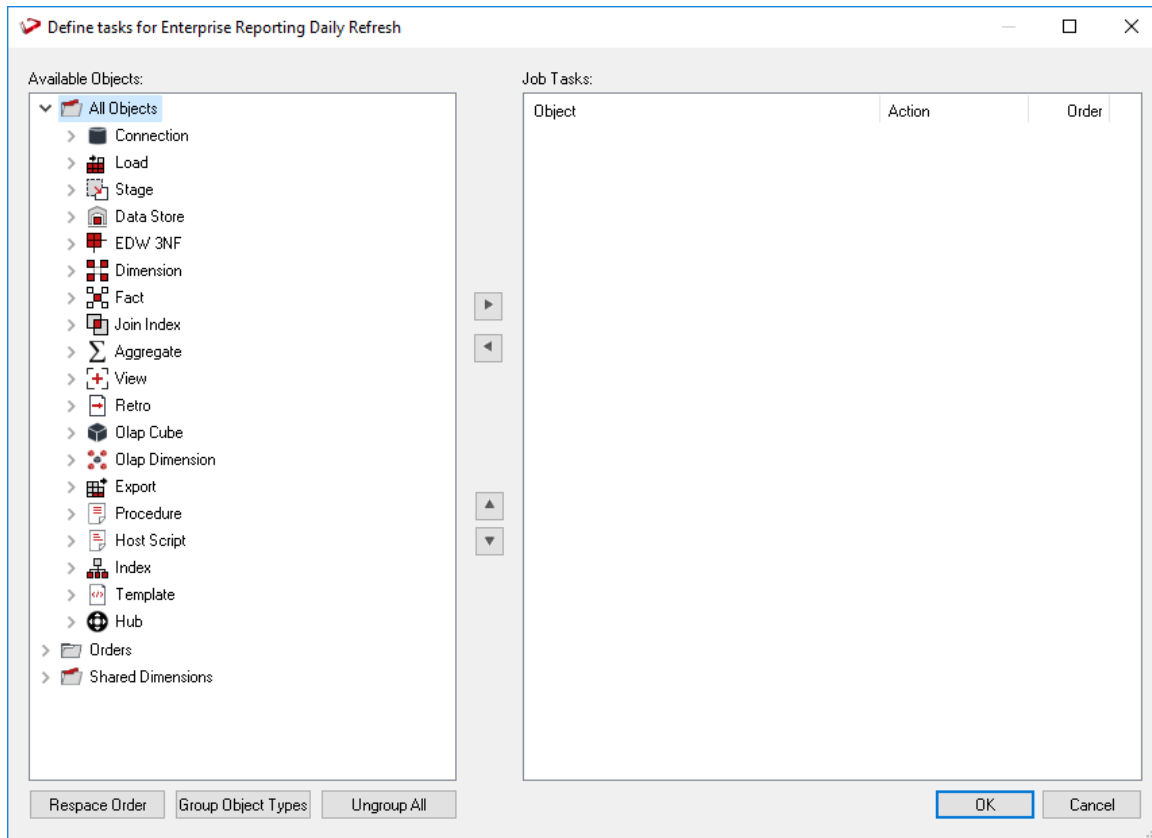
JOB TASK LIMIT - There is maximum number of **999 tasks** that can be added to a job.

To edit the tasks of a job

Select the job from the scheduler middle pane. Right-click on the job and select **Edit Tasks** from the drop-down list.



The **Define tasks** window is shown below.



The screen has two main areas. The right pane shows the tasks to be run for this job and the left pane lists all the objects.

To add a task

Double click on an object to add it to the left pane. Normally objects such as load or fact tables are scheduled rather than procedures.

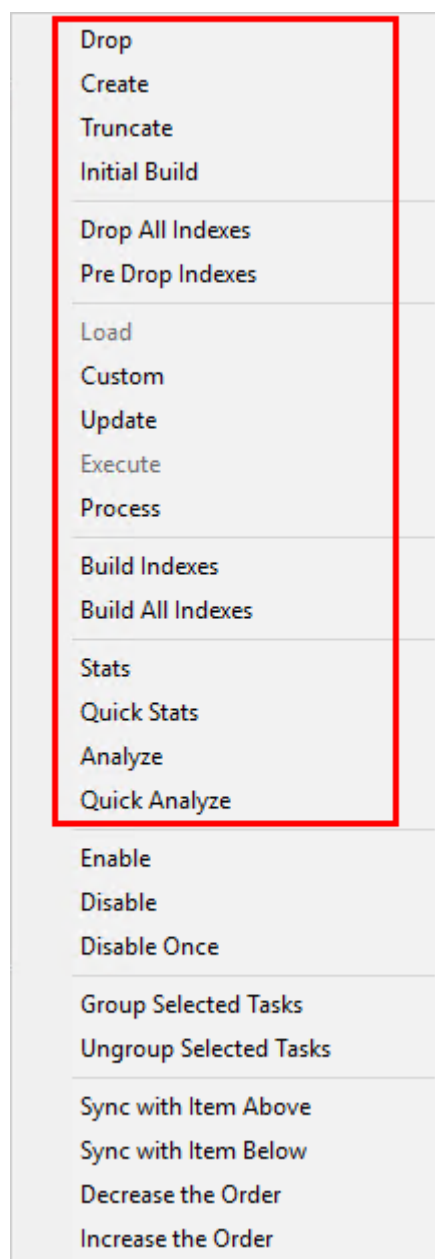
To set the action on a task

Each task can have a specific action that is to be performed on its object.

The default action for **load tables** is **process**. This means that when the task is actioned it will drop any indexes that are due to be dropped, or have **pre-drop** set, then load the table and perform any post-load procedures or transformations and then re-create any dropped indexes.

The default action for all other tables is the same as above, except it will execute the *update* procedure rather than loading the table.

You can change the action on a task by right-clicking on the task in the right pane. The menu options are shown below.



The following task actions are available:

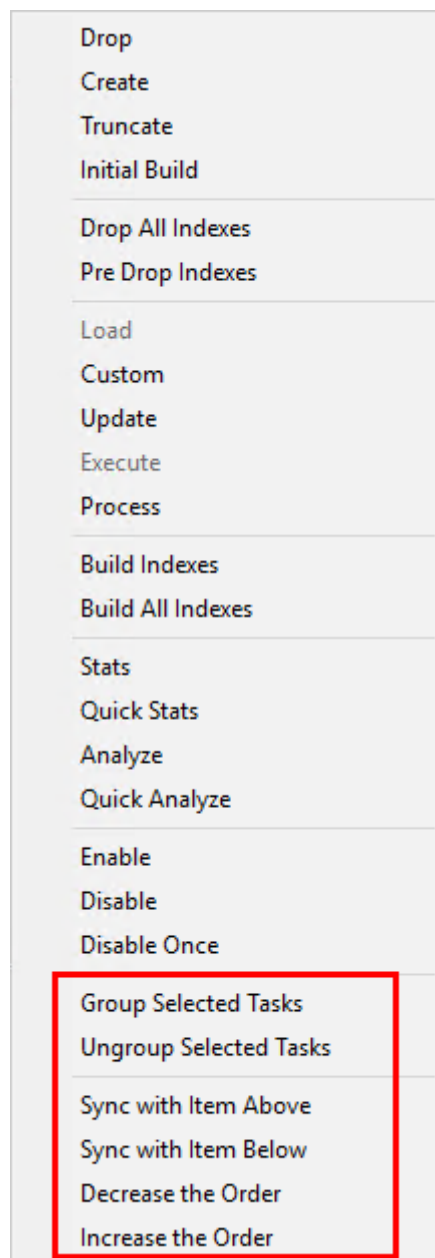
Action	Description
Drop	Drop table, view, join index or index.
Create	Create table, view, join index or index.
Truncate	Delete all rows from the table.
Initial Build	Drop All Indexes then Custom then Build All Indexes .

Action	Description
Drop All Indexes	Drop all indexes on the table.
Pre Drop Indexes	Drop all indexes on the table marked as "Pre Drop".
Load	Load the table (Load tables only).
Custom	Run the custom procedure on the table.
Update	Run the update procedure on the table.
Execute	Execute the procedure or host scripts.
Process	Pre Drop Indexes then Update and then Build Indexes .
Process and Statistics	Process then Default Stats as defined on Table Properties/Statistics/Process and statistics method (DB2 only).
Build Indexes	Build the indexes on the table marked as "Pre Drop".
Build All Indexes	Build all indexes on the table.
Stats	Refreshes predefined statistics on a table or index: COLLECT STATISTICS ON DatabaseName.TableName; COLLECT STATISTICS ON DatabaseName.TableName INDEX IndexName;
Quick Stats	Refreshes predefined statistics on an index using sampling: COLLECT STATISTICS USING SAMPLE ON DatabaseName.TableName INDEX IndexName;
Analyze	Refreshes predefined statistics on a table or index: COLLECT STATISTICS ON DatabaseName.TableName; COLLECT STATISTICS ON DatabaseName.TableName INDEX IndexName;
Quick Analyze	Refreshes predefined statistics on an index using sampling: COLLECT STATISTICS USING SAMPLE ON DatabaseName.TableName INDEX IndexName;

Note: Not all actions are available on all object types.

To create dependencies between tasks

It is possible to create dependencies between tasks in the list by selecting one or more tasks and right-clicking to bring up the dependency options.



The following task dependency options are available from the menu:

Task Option	Description
Group Selected Tasks	Groups two or more selected tasks to have the same order value, allowing them to run in parallel if the maximum threads setting allows.
Ungroup Selected Tasks	Un-group selected tasks.

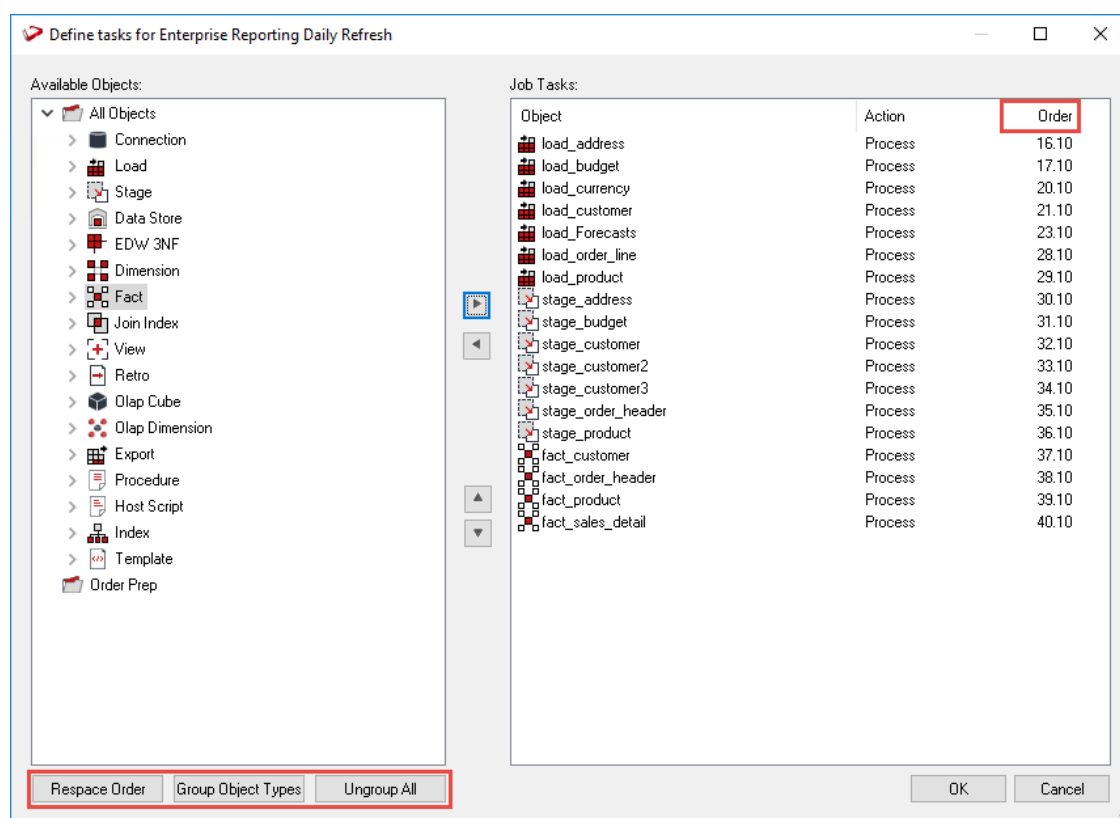
Task Option	Description
Sync with Item Above	Changes a selected task to have the same order value as the task above it, allowing them to run in parallel if the maximum threads setting allows.
Sync with Item Below	Changes a selected task to have the same order value as the task below it, allowing them to run in parallel if the maximum threads setting allows.
Decrease the Order	Changes a selected task to an order number one less than its current value. The task will now run immediately before it would have previously.
Increase the Order	Changes a selected task to an order number one more than its current value. The task will now run immediately after it would have previously.

To order or group the tasks

The **Order** column shows the order in which the tasks are to be run, e.g. 20.20 If the two numbers are the same as another task then those tasks can run in parallel. If the two numbers are different then those tasks run sequentially. This is an initial definition of dependencies. These dependencies can be altered specifically once the job has been created.

Tasks can be moved up or down by selecting the task and clicking the **Move Up** or **Move Down** buttons.

To respace the order of the tasks; to group or ungroup object types, use the **buttons** at the bottom of the **Define tasks** dialog.



- **Respace Order**

This button will respace the order numbers. The existing dependency structure and groupings are retained. The purpose of this button is simply to allow room between tasks to fit new tasks. So for example if we have two tasks that have an order of 20.19.5 and 20.20.6 and we want to add a task between these two tasks we can click the **Respace Order** button to open up a gap between the two tasks.

- **Group Object Types**

This option will put all objects of the same type into groups. For example all load tables will be able to run in parallel, all dimensions etc.

- **Ungroup All**

This button will remove all groupings and make all tasks sequential. New groupings can be made by selecting a range of sequentially listed tasks in the left pane and using the right-click menu option **Group Selected Tasks**. Tasks that are grouped have the same first two numbers in the order and can execute at the same time if the job has multiple threads.

Upon completion of editing tasks, click **OK**.

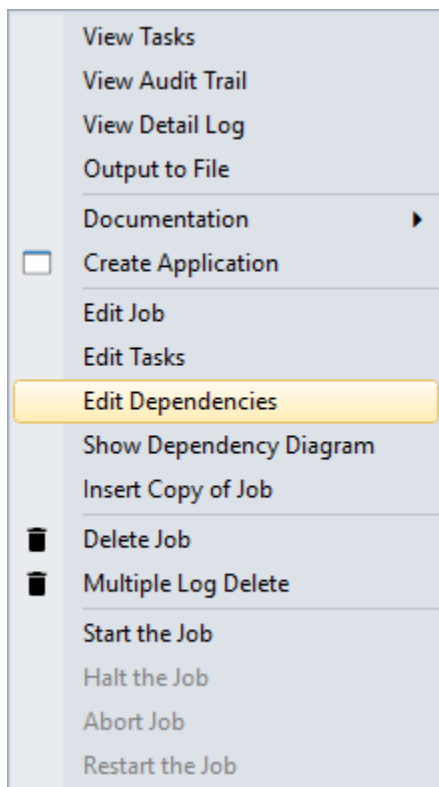
EDITING TASK DEPENDENCIES

Once jobs have been created they can be edited.

Note: A job can only be edited when it is not in a running state and only if the job is a scheduled job. Completed jobs remain in the list but only logs remain.

To edit task dependencies

Select the job from the scheduler middle pane. Right-click on the job and select **Edit Dependencies** from the drop-down list.



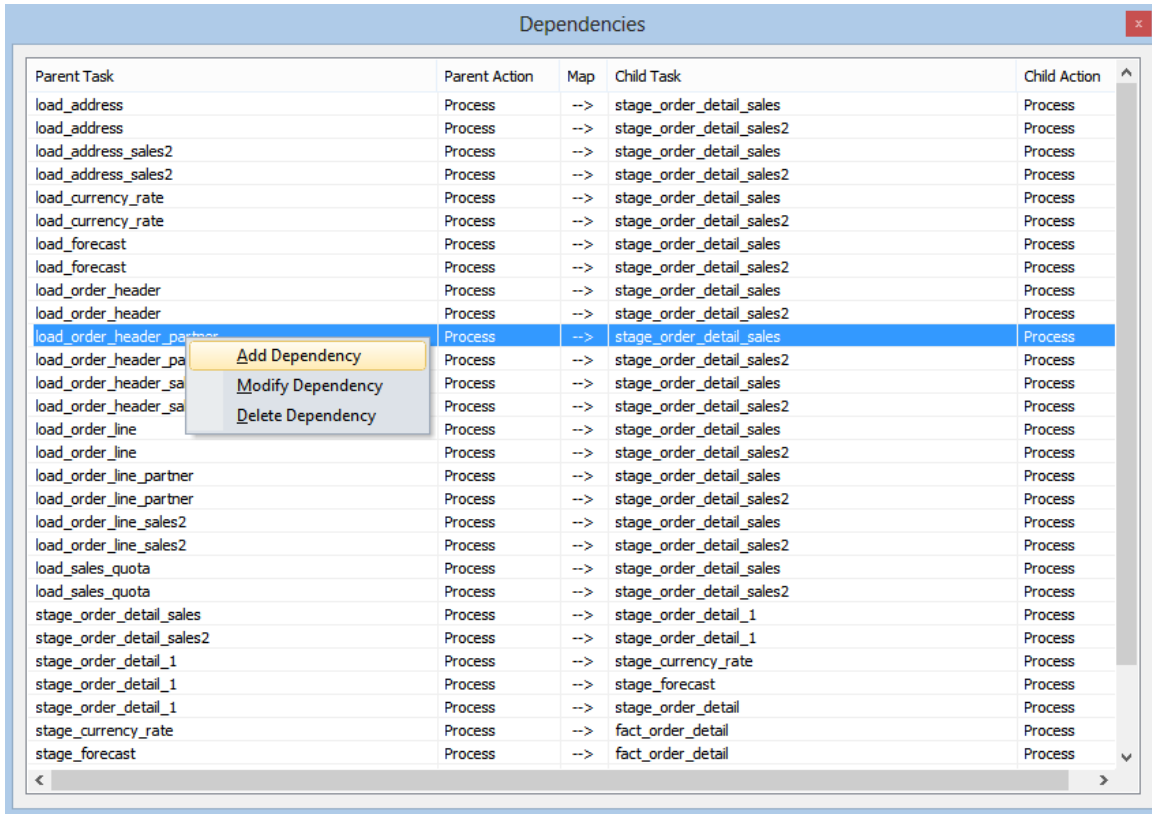
The **Dependencies** dialog will be displayed, showing the dependencies between the tasks of the job. The list consists of **Parent Tasks** on the left and **Child Tasks** on the right. A child task is thus dependent on its parent task in that it cannot run until its parent has run.

Parent Task	Parent Action	Map	Child Task	Child Action
load_address	Process	-->	stage_order_detail_sales	Process
load_address	Process	-->	stage_order_detail_sales2	Process
load_address_sales2	Process	-->	stage_order_detail_sales	Process
load_address_sales2	Process	-->	stage_order_detail_sales2	Process
load_currency_rate	Process	-->	stage_order_detail_sales	Process
load_currency_rate	Process	-->	stage_order_detail_sales2	Process
load_forecast	Process	-->	stage_order_detail_sales	Process
load_order_header	Process	-->	stage_order_detail_sales2	Process
load_order_header	Process	-->	stage_order_detail_sales	Process
load_order_header_partner	Process	-->	stage_order_detail_sales	Process
load_order_header_partner	Process	-->	stage_order_detail_sales2	Process
load_order_header_sales2	Process	-->	stage_order_detail_sales	Process
load_order_header_sales2	Process	-->	stage_order_detail_sales2	Process
load_order_line	Process	-->	stage_order_detail_sales	Process
load_order_line	Process	-->	stage_order_detail_sales2	Process
load_order_line_partner	Process	-->	stage_order_detail_sales	Process
load_order_line_partner	Process	-->	stage_order_detail_sales2	Process
load_order_line_sales2	Process	-->	stage_order_detail_sales	Process
load_order_line_sales2	Process	-->	stage_order_detail_sales2	Process
load_sales_quota	Process	-->	stage_order_detail_sales	Process
load_sales_quota	Process	-->	stage_order_detail_sales2	Process
stage_order_detail_sales	Process	-->	stage_order_detail_1	Process
stage_order_detail_sales2	Process	-->	stage_order_detail_1	Process
stage_order_detail_1	Process	-->	stage_currency_rate	Process
stage_order_detail_1	Process	-->	stage_forecast	Process
stage_order_detail_1	Process	-->	stage_order_detail	Process
stage_currency_rate	Process	-->	fact_order_detail	Process
stage_forecast	Process	-->	fact_order_detail	Process

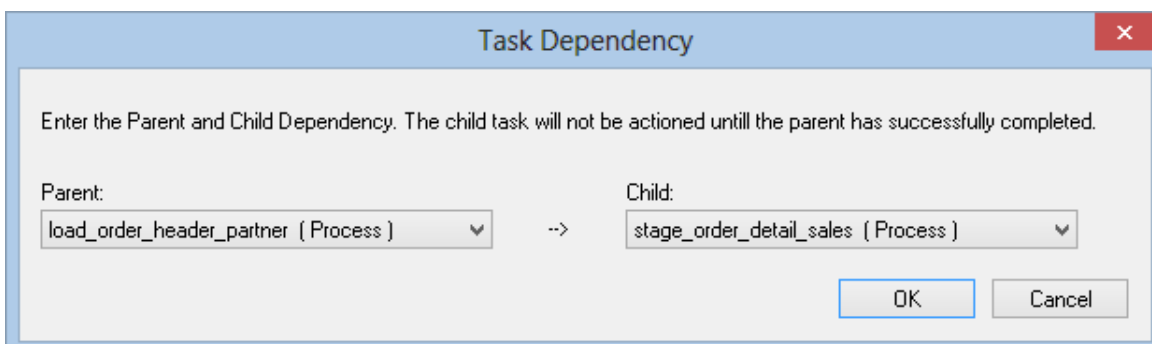
Edit the dependencies and close the dialog.

To add a task dependency

To add a task dependency, right-click anywhere in the Dependencies pane and select **Add Dependency**.

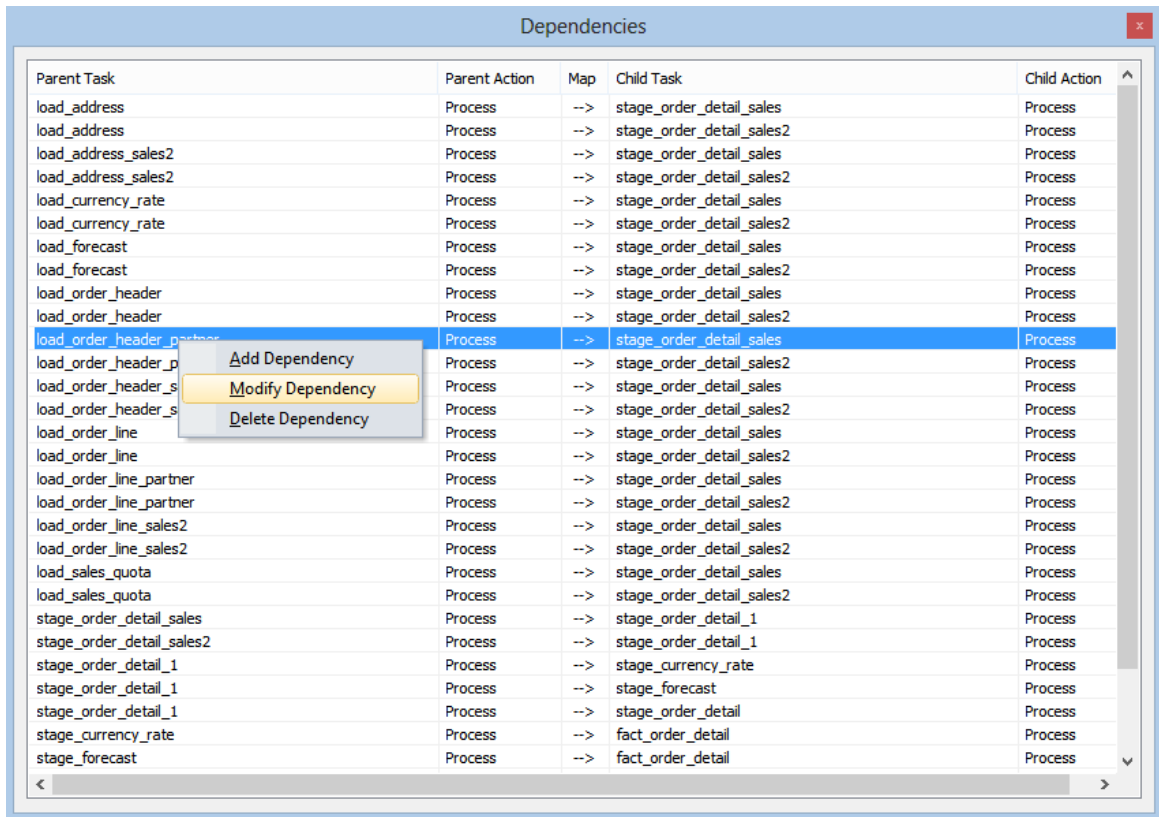


Select the **Parent** and **Child** tasks from the drop-down lists to create the dependency and click **OK**.

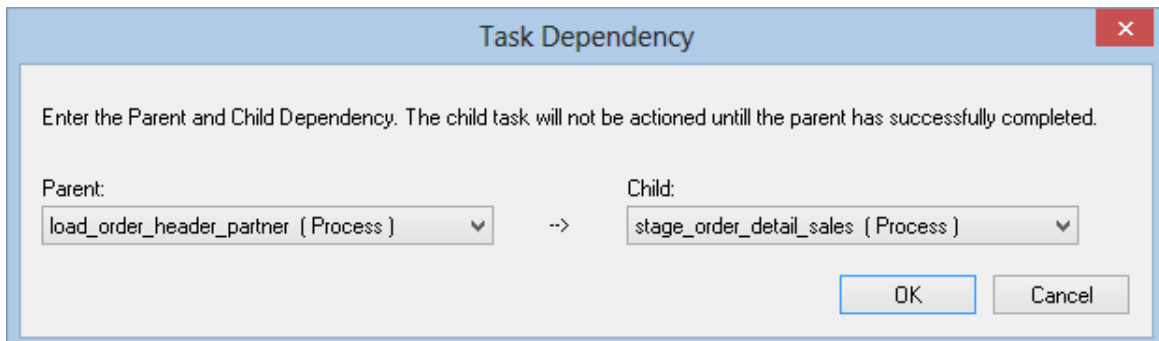


To modify a task dependency

To modify a task dependency, right-click on the dependency in the Dependencies pane and select **Modify Dependency**.

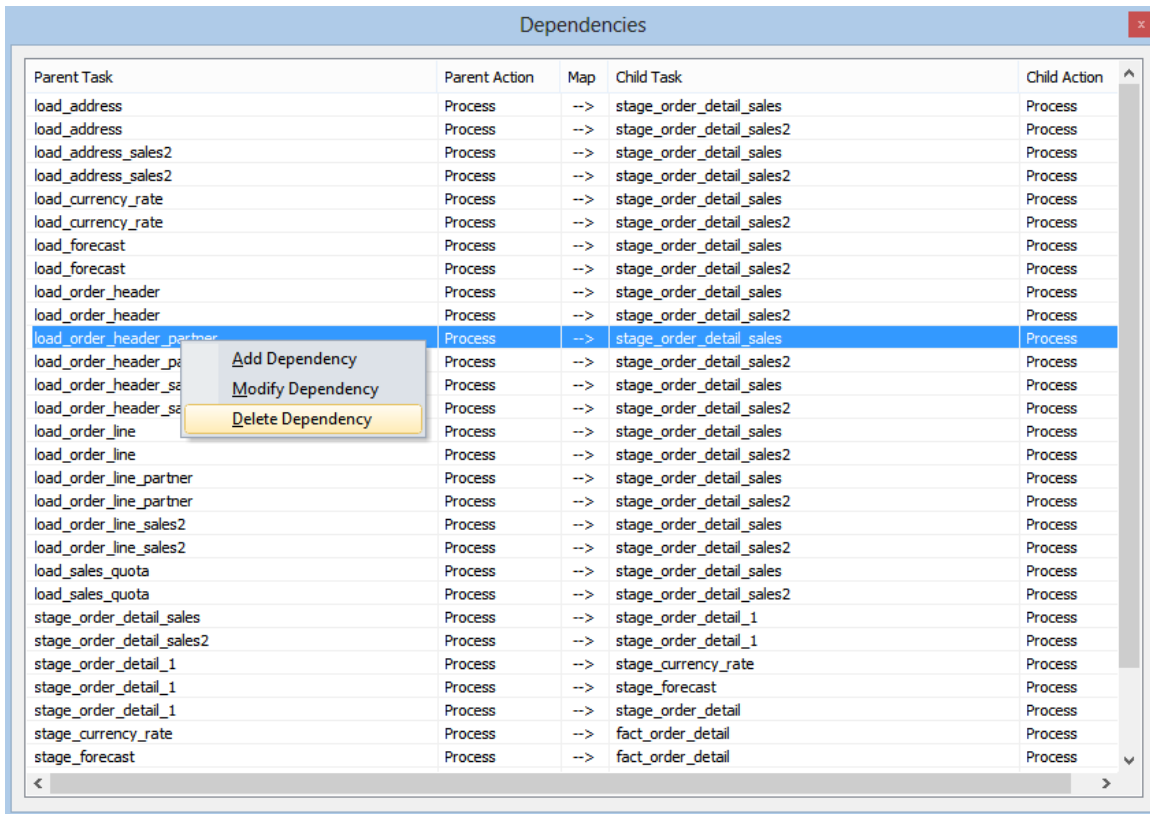


Change the **Parent** and **Child** tasks to modify the dependency and click **OK**.



To delete a task dependency

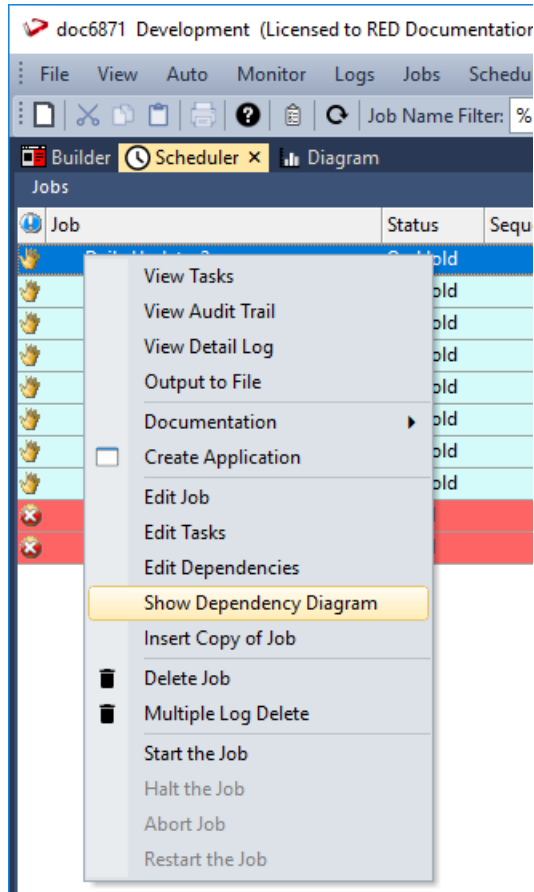
To delete a task dependency, right-click on the dependency in the Dependencies pane and select **Delete Dependency**.



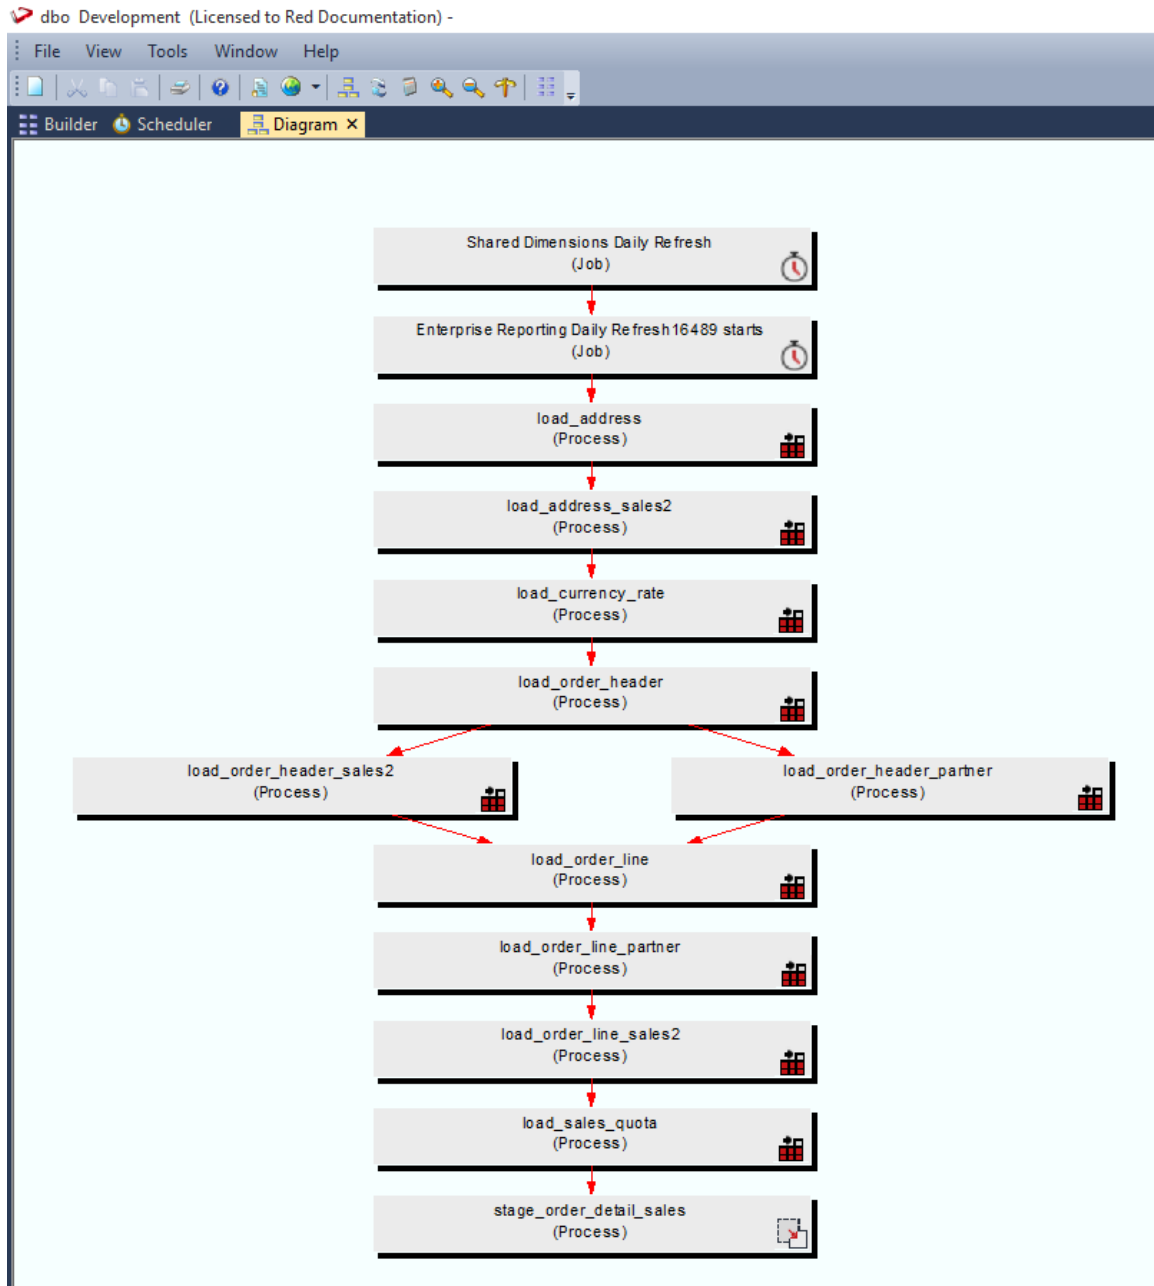
The dependency will be deleted without warning.

SHOW DEPENDENCIES DIAGRAM

Select the **Show Dependency Diagram** option from the right-click menu of any job to see all job dependencies displayed as a Diagram from RED's Diagram view tab.



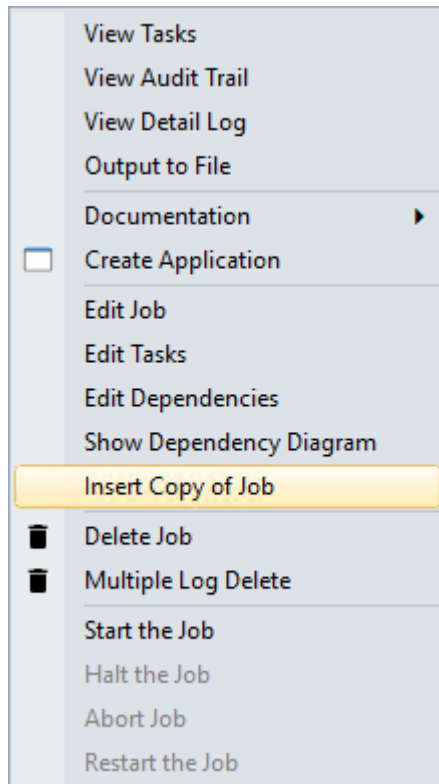
Job Dependency Diagram view



INSERTING A COPY OF A JOB

To insert a copy of a job

A copy of a job can be inserted by right-clicking on the job and choosing **Insert Copy of Job**.



The new job will immediately be visible and the Status will be **On Hold**.

The screenshot shows the 'Scheduler' window in the WhereScape interface. The 'Jobs' tab is active, displaying a table of scheduled jobs. The table has the following columns: Job, Status, Seq, Start, Finish, Elapsed, OK, Info, Detail, Warn, Error, and Who. The 'Daily Update' job is highlighted in blue and has a status of 'On Hold'. The 'Job_10456078' job is also highlighted in blue and has a status of 'On Hold'. The 'Daily Update' job has a status of 'Completed'.

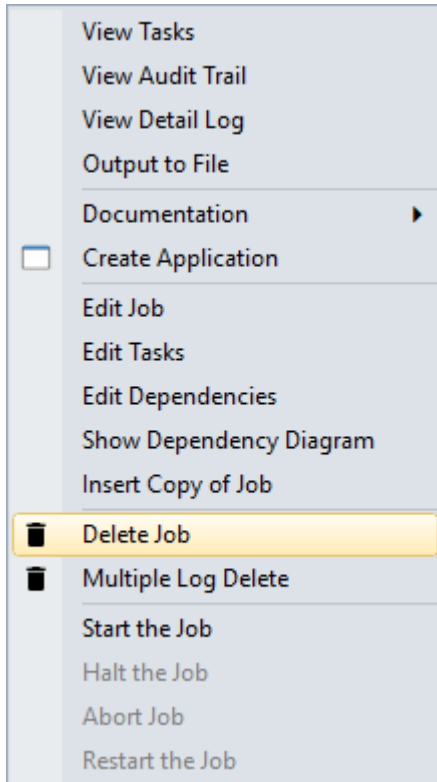
Job	Status	Seq	Start	Finish	Elapsed	OK	Info	Detail	Warn	Error	Who
Job_1045	On Hold	...	2013-08-22 07:47:00								JS
Job_10456078	On Hold	...	2013-08-22 07:47:00								JS
Daily Update	On Hold	...	2013-08-16 08:46:21								JS
Daily Update	Completed	...	2013-08-16 08:46:21	2013-08-16 08:46:24	00:00	11	26				JS

Tasks: Audit and Scheduler log
 Ready Database time at last action was 08:34:36 | Development (Red_1) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS


DELETING A JOB

To delete a job

A job can be deleted by right-clicking on a job in the scheduler window and choosing **Delete Job**.



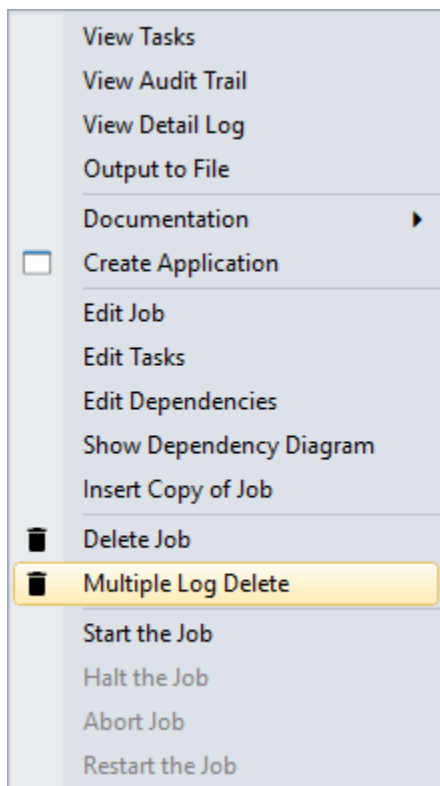
A warning message will be displayed; click **Yes** to delete.



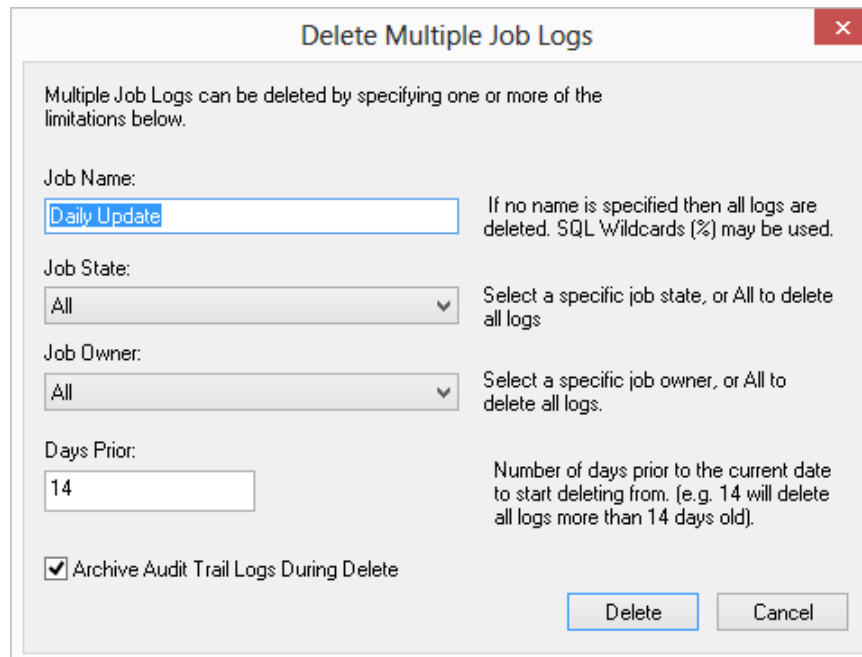
DELETING JOB LOGS

To delete multiple job logs

Multiple job logs can be deleted by right-clicking on a job in the scheduler window and choosing **Multiple Log Delete**.



The **Delete Multiple Job Logs** dialog is displayed. Select or enter the appropriate options to delete the range of job logs required.



Delete Multiple Job Logs

Multiple Job Logs can be deleted by specifying one or more of the limitations below.

Job Name: If no name is specified then all logs are deleted. SQL Wildcards (%) may be used.

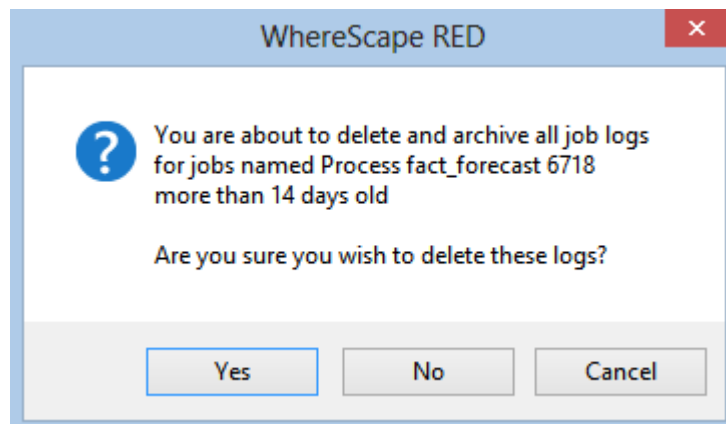
Job State: Select a specific job state, or All to delete all logs

Job Owner: Select a specific job owner, or All to delete all logs.

Days Prior: Number of days prior to the current date to start deleting from. (e.g. 14 will delete all logs more than 14 days old).

Archive Audit Trail Logs During Delete

A warning message is displayed. Click **Yes** to delete.



WhereScape RED

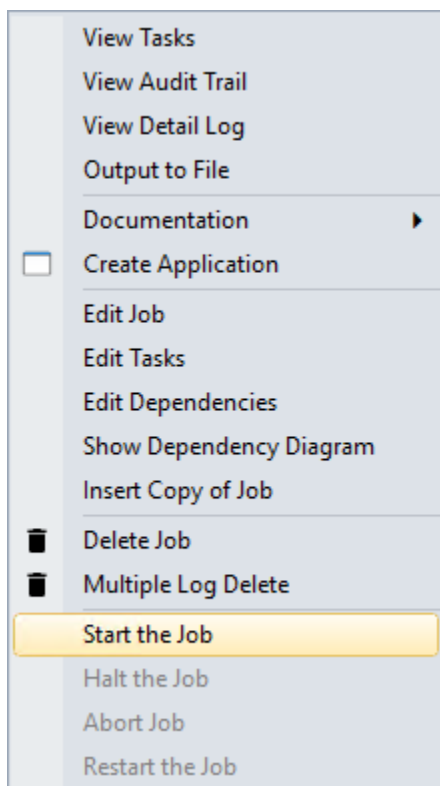
? You are about to delete and archive all job logs for jobs named Process fact_forecast 6718 more than 14 days old

Are you sure you wish to delete these logs?

STARTING A JOB

To start a job

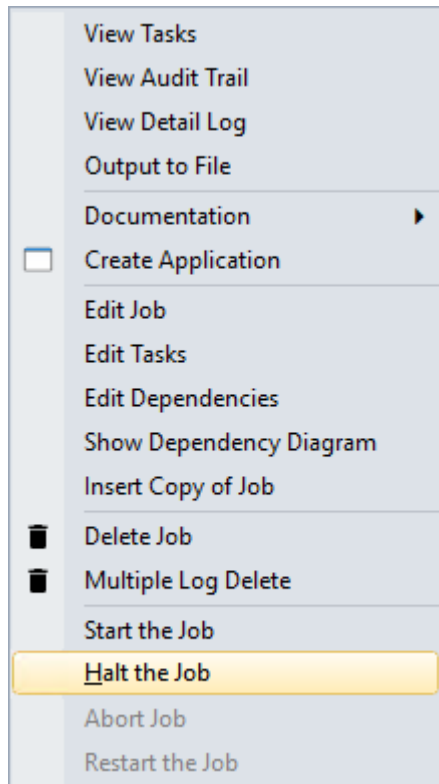
Multiple job logs can be deleted by right-clicking on a job in the scheduler window and choosing **Multiple Log Delete**.



HALTING A JOB

To halt a job

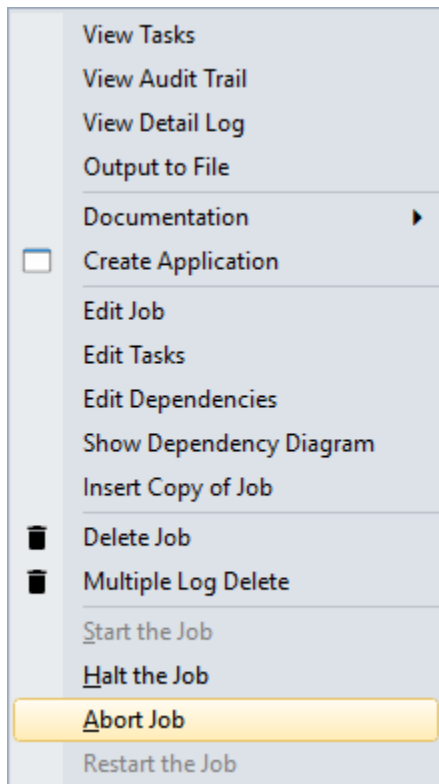
A job can be halted by right-clicking on a job in the scheduler window and choosing **Halt the Job**.



ABORTING A JOB

To abort a job

A job can be aborted by right-clicking on the job in the scheduler window and choosing **Abort Job**.



Once in this state, a job cannot be restarted. The job now exists only as a log of what occurred and is no longer regarded as a job.

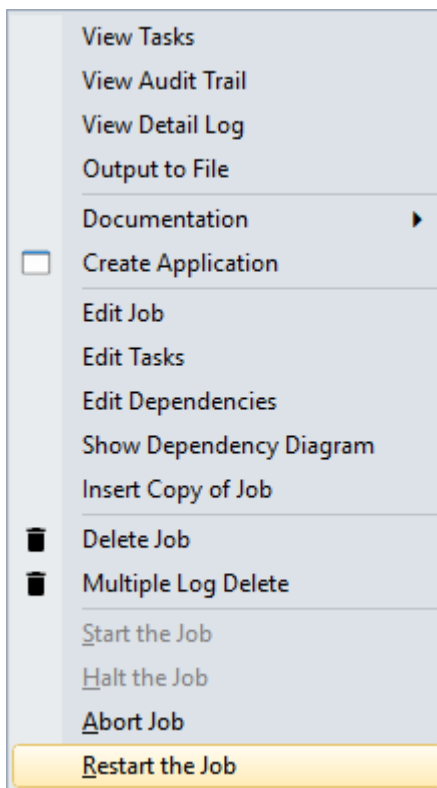
Effects of aborting a job

Load and update processes are not stopped for all objects in Teradata repositories.

RESTARTING A JOB

To restart a job

A job can be restarted by right-clicking on a job in the scheduler window and choosing **Restart the Job**.



Before restarting a job, it is possible to edit the **status** of the job tasks so that only certain tasks will be **run again** or be **skipped over**.

To run a task again

View the job tasks by double-clicking on the failed job. The tasks will be displayed in the bottom pane.

Job	Status	Seq	Start	Fini...	Elapsed	OK	Info	Detail	Warn	Error	Who
Process fact_forecast 6718	On Hold	6718	2013-02-14 15:36:00								JS
Refresh Order Details	On Hold	7248	2013-02-14 12:39:36								JS
Shared Dimensions Daily Refresh	On Hold	1150	2013-02-14 03:00:00								WQ
Enterprise Reporting Daily Refresh	On Hold	6280	2013-01-31 17:18:32								JS
Process_to_fact_order_detail	On Hold	6730	2013-01-31 16:20:00								JS
Refresh Order Details	Failed	7258	2013-02-14 12:39:36	201...	00:00	3	11			3	JS
Refresh Order Details	Completed	7256	2013-02-14 12:29:05	201...	00:00	5	12				JS

Task	Action	Status	Seq	Start	Finish	Ela...	I...	D...	W...	Result
load_order_head...	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	4			9 rows loaded into load_order_header
load_order_line	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	4			21 rows loaded into load_order_line
stage_order_detail	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	2			stage_order_detail updated. 21 new records. 0 records up...
fact_order_detail	Process	Failed	7258	2013-02-14 ...	2013-02-14 ...	00:...	1			Ws_Act_Update(6.5.5.1) step 1600: update_fact_order_det...

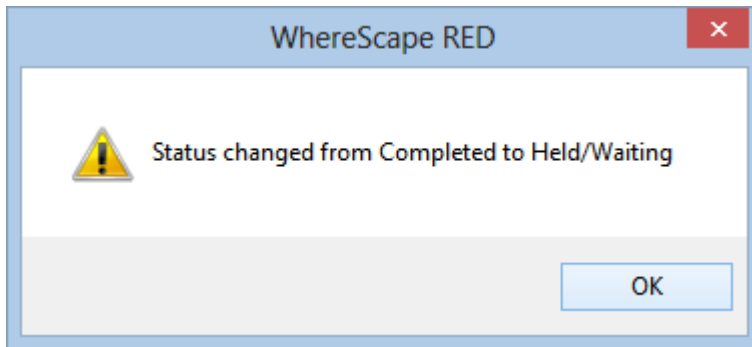
Tasks Audit and Scheduler log
Ready Database time at last action was 13:03:10 | Development (WsWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

To rerun a task, right-click on the completed task and select **Change to On Hold**.

Task	Action	Status	Seq	Start	Finish	Ela...	I...	D...	W...	Result
load_order_head...	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	4			9 rows loaded into load_order_header
load_order_line	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	4			21 rows loaded into load_order_line
stage_order_detail	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	2			stage_order_detail updated. 21 new records. 0 records up...
fact_order_detail	Process	Failed	7258	2013-02-14 ...	2013-02-14 ...	00:...	1			Ws_Act_Update(6.5.5.1) step 1600: update_fact_order_det...

Tasks Audit and Scheduler log
Ready Database time at last action was 13:03:10 | Development (WsWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Click **OK** on the message dialog.



Double-click on the job again to display the tasks. You will see that the selected task now has a status of **Hold** and will thus be rerun when you restart the job.

Task	Action	Status	Seq	Start	Finish	Ela...	I...	D...	W...	Result
load_order_head...	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	4			9 rows loaded into load_order_header
load_order_line	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	4			21 rows loaded into load_order_line
fact_order_detail	Process	Failed	7258	2013-02-14 ...	2013-02-14 ...	00:...	1			Ws_Act_Update(6.5.5.1) step 1600: update_fact_order_det...
stage_order_detail	Process	Held	7258	2013-02-14 ...	2013-02-14 ...	00:...	2	1		stage_order_detail updated. 21 new records. 0 records up...

Tasks Audit and Scheduler log
Ready Database time at last action was 13:15:07 | Development (WslWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

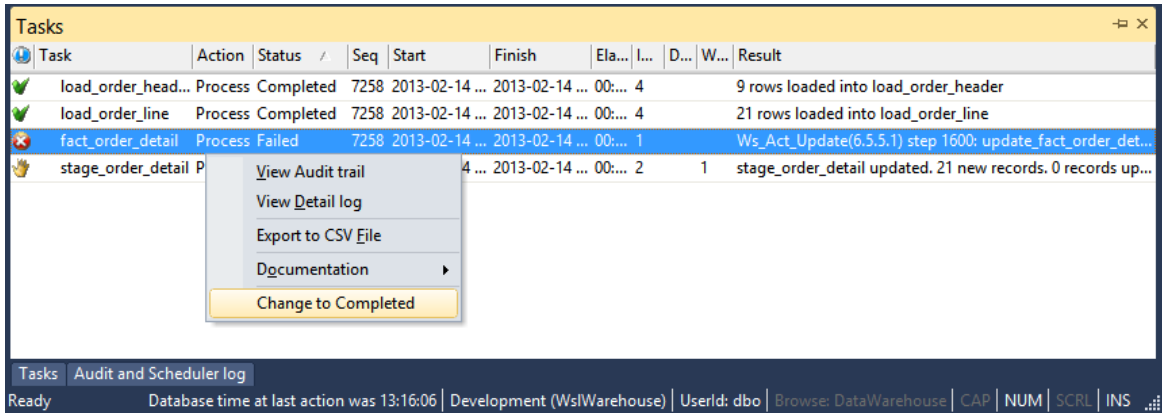
To skip over a task

View the job tasks by double-clicking on the failed job. The tasks will be displayed in the bottom pane.


Task	Action	Status	Seq	Start	Finish	Ela...	I...	D...	W...	Result
load_order_head...	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	4			9 rows loaded into load_order_header
load_order_line	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	4			21 rows loaded into load_order_line
fact_order_detail	Process	Failed	7258	2013-02-14 ...	2013-02-14 ...	00:...	1			Ws_Act_Update(6.5.5.1) step 1600: update_fact_order_det...
stage_order_detail	Process	Held	7258	2013-02-14 ...	2013-02-14 ...	00:...	2	1		stage_order_detail updated. 21 new records. 0 records up...

Tasks Audit and Scheduler log
Ready Database time at last action was 13:15:07 | Development (WslWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

To skip over a task, right-click on the task and select **Change to Completed**.



Click **OK** on the message dialog.



Double-click on the job again to display the tasks. You will see that the selected task now has a status of **Completed** and will thus be skipped when you restart the job.

The screenshot displays the 'Jobs' and 'Tasks' sections of the WhereScape interface. The 'Jobs' table lists several jobs, with the 'Refresh Order Details' job (Seq 7256) highlighted in green, indicating it is 'Completed'. The 'Tasks' table below shows the details of this job, with the 'fact_order_detail' task highlighted in blue and marked as 'Completed'.

Job	Status	Seq	Start	Fin...	Elapsed	OK	Info	Detail	Warn	Error	Who
Process fact_forecast 6718	On Hold	6718	2013-02-14 15:36:00								JS
Refresh Order Details	On Hold	7248	2013-02-14 12:39:36								JS
Shared Dimensions Daily Refresh	On Hold	1150	2013-02-14 03:00:00								WQ
Enterprise Reporting Daily Refresh	On Hold	6280	2013-01-31 17:18:32								JS
Process_to_fact_order_detail	On Hold	6730	2013-01-31 16:20:00								JS
Refresh Order Details	Failed	7258	2013-02-14 12:39:36	201...	00:00	3	11			3	JS
Refresh Order Details	Completed	7256	2013-02-14 12:29:05	201...	00:00	5	12				JS

Task	Action	Status	Seq	Start	Finish	Ela...	I...	D...	W...	Result
load_order_head...	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	4			9 rows loaded into load_order_header
load_order_line	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	4			21 rows loaded into load_order_line
fact_order_detail	Process	Completed	7258	2013-02-14 ...	2013-02-14 ...	00:...	1	1		Ws_Act_Update(6.5.5.1) step 1600: update_fact_order_det...
stage_order_detail	Process	Held	7258	2013-02-14 ...	2013-02-14 ...	00:...	2	1		stage_order_detail updated. 21 new records. 0 records up...

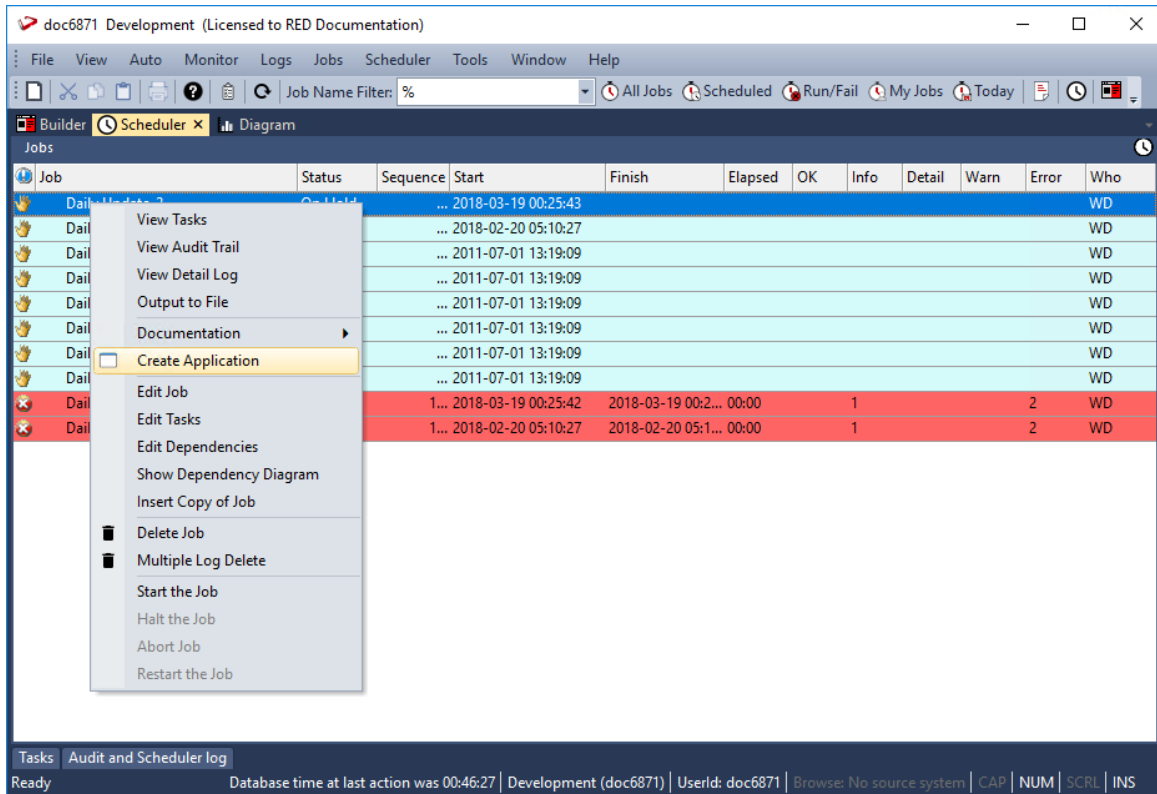
Tasks | Audit and Scheduler log

Ready Database time at last action was 13:18:12 | Development (WsWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

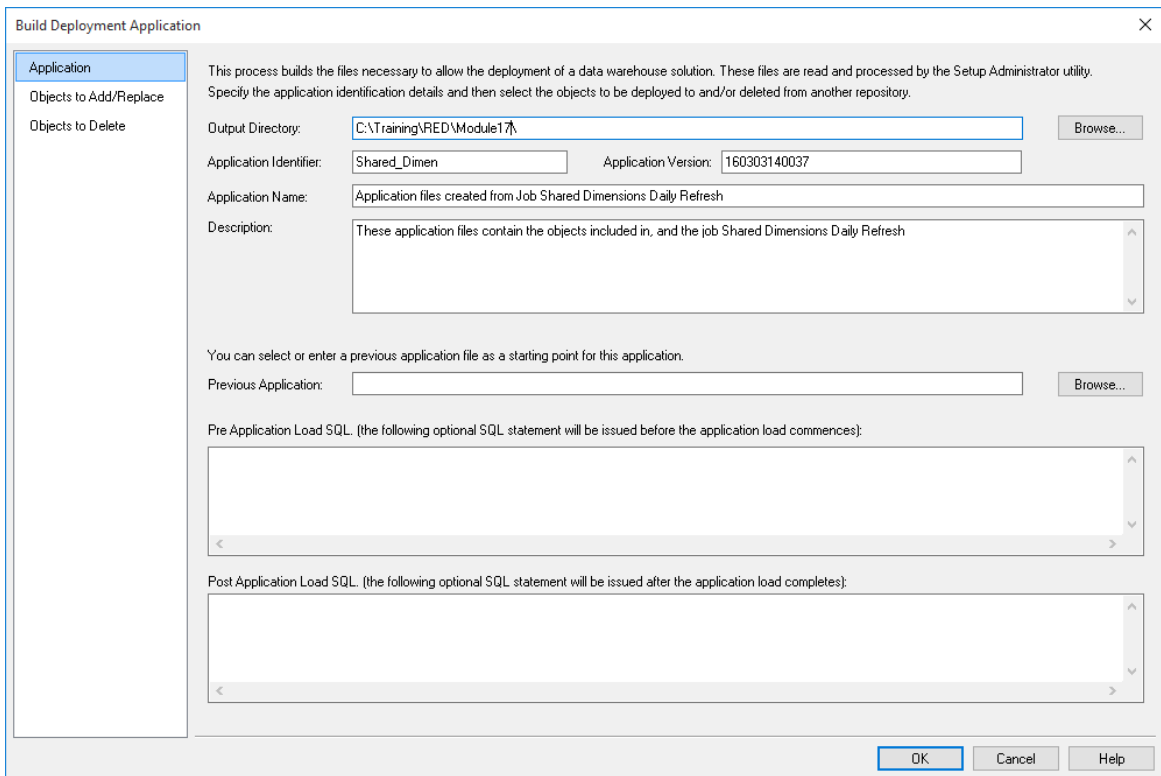
CREATING AN APPLICATION FROM A JOB

To Create an Application from a Job

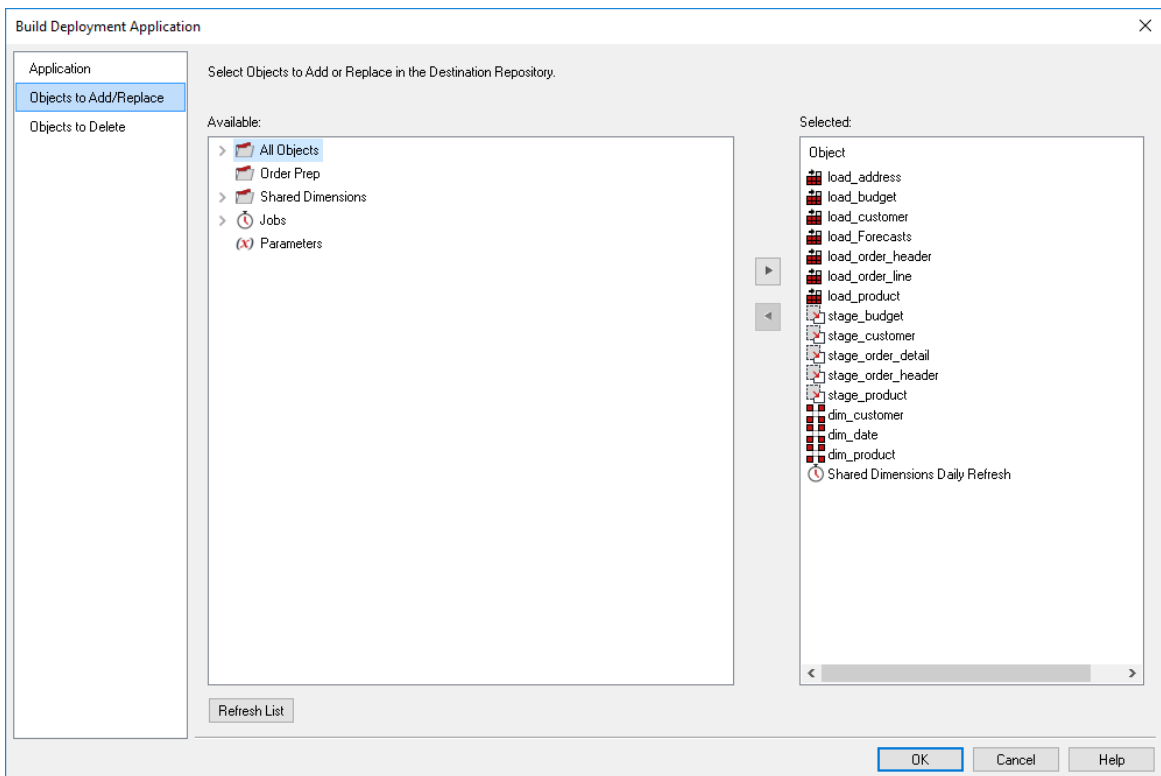
- 1 Right-click on the job in the scheduler window and select **Create Application**.



2 Edit the application as required.



3 Edit the objects to add or replace as required.



Note: Creating an application from a job will save the objects in the job and the job, but not the associated objects.

- 4 Click **OK** when finished.
- 5 A dialog will display, confirming the creation of the application files. Click **OK**.

STAND ALONE SCHEDULER MAINTENANCE

WhereScape RED includes a stand alone scheduler maintenance screen. This screen provides all the scheduler control functionality found in the main RED utility, but with no access to the main metadata repository.

Scheduler maintenance logon

The logon screen differs in that the user name and password do not have to be that of a valid metadata repository. This user name/password combination can be any valid database user.

Scheduler maintenance grants

Statement	Reason
grant select on ws_dbc_connect to dsssched;	Repository access
grant select on ws_meta to dsssched;	Repository access
grant select on ws_meta_tables to dsssched;	Repository access
grant select on ws_meta_names to dsssched;	Repository access
grant select on ws_obj_type to dsssched;	Repository access
grant select on ws_obj_object to dsssched;	Object access (job create)
grant select on ws_obj_pro_map to dsssched;	Object access (job create)
grant select on ws_obj_project to dsssched;	Object access (job create)
grant select on ws_obj_group to dsssched;	Object access (job create)

Statement	Reason
grant select on ws_pro_gro_map to dsssched;	Object access (job create)
grant select on ws_wrk_audit_log to dsssched;	Scheduler status
grant select,insert,update on ws_user_adm to dsssched;	Repository access
grant select,delete on ws_wrk_error_log to dsssched;	Scheduler status, and job deletion
grant select,update on ws_wrk_scheduler to dsssched;	Scheduler status, poll
grant select,insert,update,delete on ws_wrk_dependency to dsssched;	Job creation, maintenance
grant select,insert,update,delete on ws_wrk_job_ctrl to dsssched;	Job creation, maintenance
grant select,insert,delete on ws_wrk_job_log to dsssched;	Job maintenance
grant select,update,delete on ws_wrk_job_run to dsssched;	Job maintenance
grant select,insert,update,delete on ws_wrk_dependency to dsssched;	Job maintenance
grant select,insert,update,delete on ws_wrk_job_dependency to dsssched;	Job maintenance
grant select,delete on ws_wrk_job_thread to dssched;	Job maintenance
grant select,insert on ws_wrk_sequence to dssched;	Job creation
grant select,insert,update,delete on ws_wrk_task_ctrl to dssched;	Task maintenance
grant select,update,delete on ws_wrk_task_run to dssched;	Task maintenance
grant select,insert,delete on ws_wrk_task_log to dssched;	Task maintenance
grant select,insert,update on dss_parameter to dssched;	Task maintenance
grant select on ws_pro_header to dssched;	Right-click used by option in parameters listing

Statement	Reason
grant select on ws_pro_line to dsssched;	Right-click used by option in parameters listing
grant select on ws_scr_header to dsssched;	Right-click used by option in parameters listing
grant select on ws_scr_line to dsssched;	Right-click used by option in parameters listing
grant select on ws_load_tab to dsssched;	Right-click used by option in parameters listing
grant select on ws_load_col to dsssched;	Right-click used by option in parameters listing
grant select on ws_stage_tab to dsssched;	Right-click used by option in parameters listing
grant select on ws_stage_col to dsssched;	Right-click used by option in parameters listing
grant select on ws_dim_tab to dsssched;	Right-click used by option in parameters listing
grant select on ws_dim_col to dsssched;	Right-click used by option in parameters listing
grant select on ws_agg_tab to dsssched;	Right-click used by option in parameters listing
grant select on ws_agg_col to dsssched;	Right-click used by option in parameters listing

A sample script to grant these privileges is shipped with WhereScape RED. This script is called 'grant_sched_access.sql' and can be found in the WhereScape program directory.

The scheduler maintenance utility does not require a WhereScape license key. The WhereScape RED software can be installed onto a PC, and this utility utilized without having to use the WhereScape 'Setup Administrator' utility.

SQL TO RETURN SCHEDULER STATUS


This SQL returns the scheduler status:

```
SELECT CASE
WHEN ws_stop_date IS NOT NULL
THEN 'STOPPED'
WHEN ((DATEDIFF(mi,ws_active_date,GETDATE()) -
CONVERT(INTEGER,DATEDIFF(mi,ws_active_date,GETDATE())/60)*60) > 15 )
OR (CONVERT(INTEGER,DATEDIFF(mi,ws_active_date,GETDATE())/60)>0)
THEN 'NOT ACTIVE'
WHEN (((DATEDIFF(mi,ws_active_date,GETDATE()) -
CONVERT(INTEGER,DATEDIFF(mi,ws_active_date,GETDATE())/60)*60)>((ws_inter
val/60)+10)
OR CONVERT(INTEGER,DATEDIFF(mi,ws_active_date,GETDATE())/60)>0)
AND ws_poll_flag=1)
THEN 'NOT ACTIVE'
ELSE 'Running'
END
FROM dbo.ws_wrk_scheduler
WHERE ws_name = 'YourSchedulerName'
```

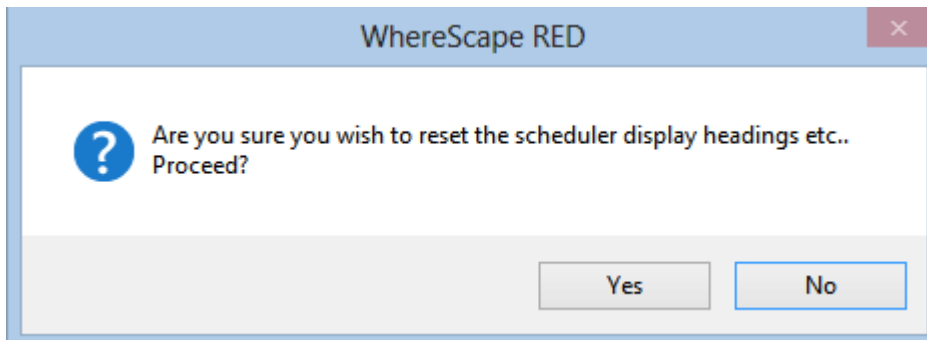
The procedure sets the status in the metadata.

RESET COLUMNS IN JOB AND TASK VIEW

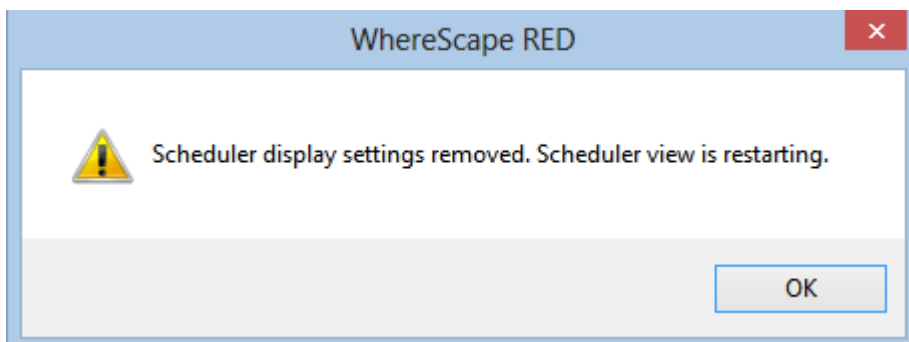
Job and Task Report headings can be reset by selecting the **View/Reset Display Headings** menu option from the scheduler window. The short-cut keys are Alt+V-R.



A dialog will ask you to confirm the request.



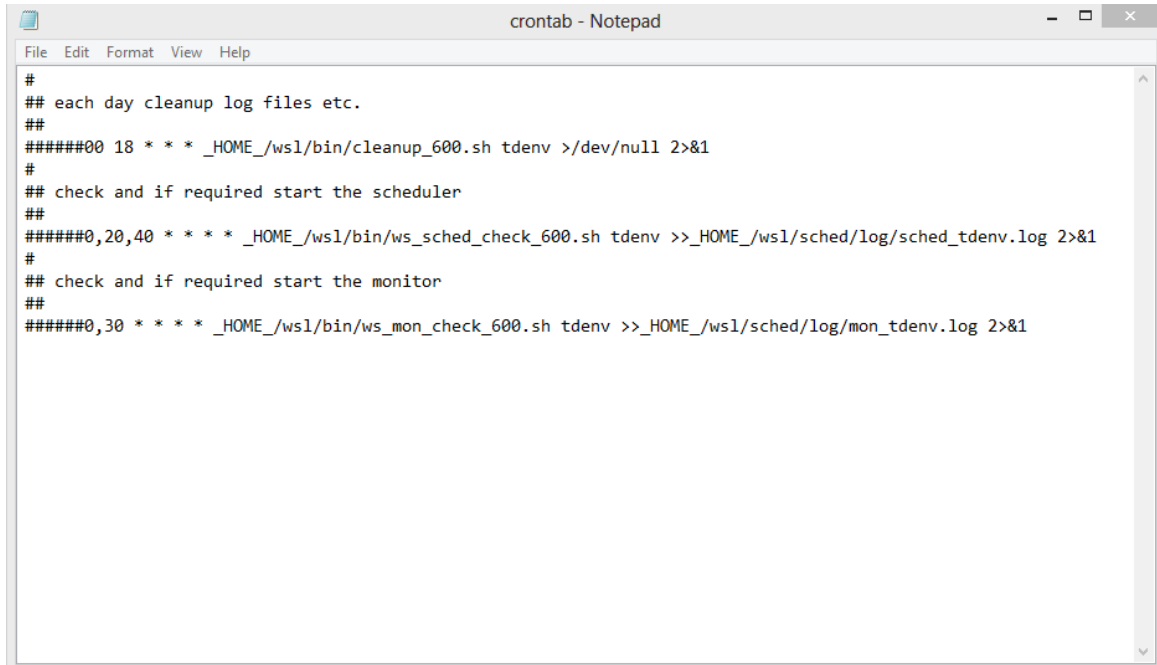
If you selected **Yes** to reset the display settings, then a dialog will confirm once the reset has occurred.



STOPPING A LINUX/UNIX SCHEDULER FROM WITHIN RED

To stop a Linux/UNIX Scheduler from within RED, follow the steps below:

- 1 Edit the **crontab** and comment out the `ws_sched_check_nnn.sh` entry. This will stop the scheduler restarting within the next 20 minutes.




```

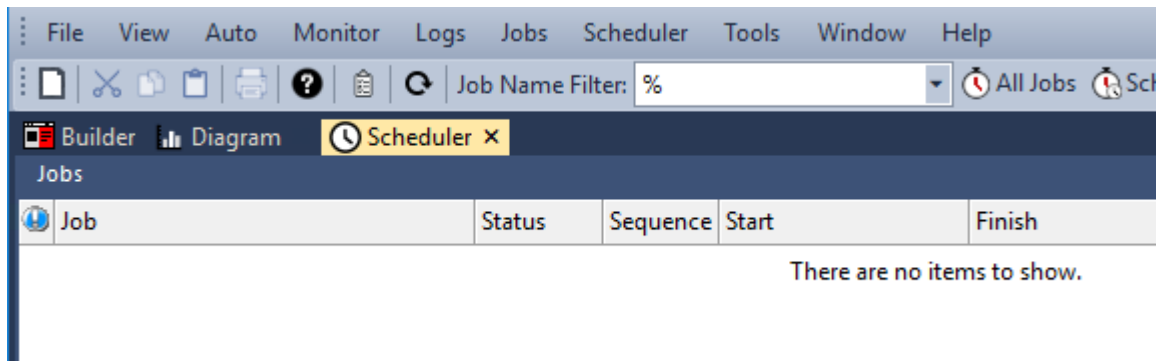
#
## each day cleanup log files etc.
##
#####00 18 * * * _HOME_/wsl/bin/cleanup_600.sh tdenv >/dev/null 2>&1
#
## check and if required start the scheduler
##
#####0,20,40 * * * * _HOME_/wsl/bin/ws_sched_check_600.sh tdenv >>_HOME_/wsl/sched/log/sched_tdenv.log 2>&1
#
## check and if required start the monitor
##
#####0,30 * * * * _HOME_/wsl/bin/ws_mon_check_600.sh tdenv >>_HOME_/wsl/sched/log/mon_tdenv.log 2>&1

```

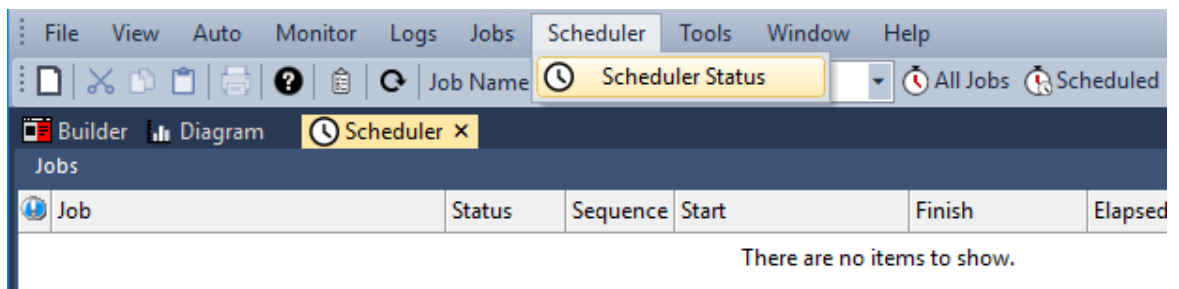
Note: There may be several different versions of the scheduler files for a given database and platform (UNIX or Linux). For example, there may be different folders in ...\\WhereScape\\Teradata\\Linux\\: Version_550 and Version_600. The highest version number script less than or equal to the version of RED in use should always be used.




- 2 Start RED and click on the **Scheduler** tab.



- 3 Click on Scheduler in the toolbar and then select **Scheduler Status**.



- 4 Right-click on the displayed UNIX/Linux scheduler entry and choose **Shutdown Scheduler**.



Sometime within the next poll interval of the scheduler, the scheduler will gracefully stop.

CHAPTER 27

INDEXES

Indexes may exist on any table defined in RED. By default, RED will auto-generate a number of indexes during the drag and drop process and when building procedures.

The **Enable Automatic Creation of Indexes** option in the **Connection Properties** window enables you to turn ON/OFF the automatic creation of metadata for indexes when creating objects. Neither the RED metadata nor the physical index is created, if this setting is not enabled.

These indexes can be altered or deleted. New indexes can be created as desired.

NOTE: The maintenance of the indexes is performed as part of the normal scheduler processing.

In the left pane, right-click on a table to:

- Display indexes
- Add indexes

IN THIS CHAPTER




Index Definition 754

INDEX DEFINITION

By right-clicking on a table in the left pane and selecting **Display Indexes** the middle pane will display the indexes for that table. Alternatively, you can double-click on the Index object type in the left pane to display all indexes in the repository or a specific group or project.

In the middle pane, right-click on an index and the following options are available:

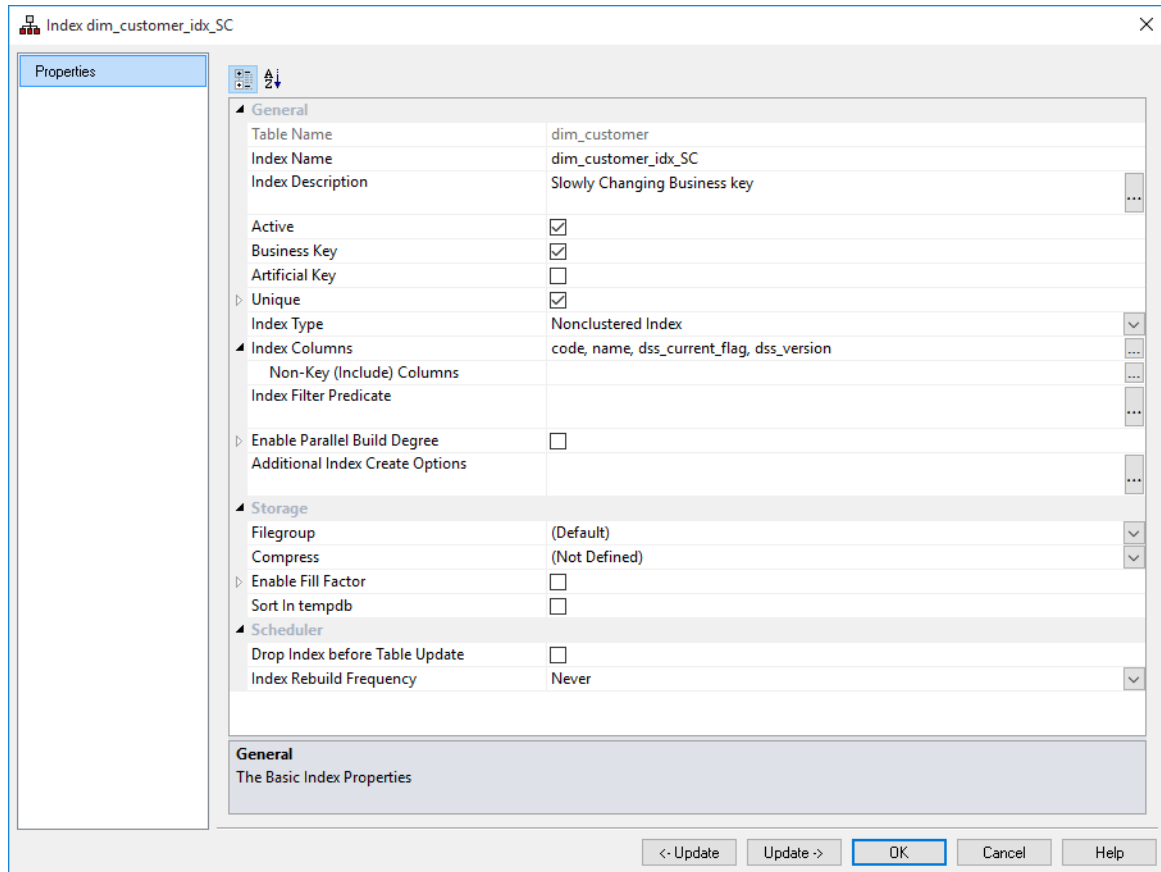
- Properties
- Create Index
- Drop Index
- Delete Metadata and Drop Index
- Create via Scheduler
- Projects

Index list for dim_customer						
Index Name	Table	Columns	Type	Key Type	Rebuild ...	
 dim_customer_idx_0	dim_customer	1	Clustered	Artificial	Never	
 dim_customer_idx_A	dim_customer	3	Composite ...	Business	Never	
 dim_customer_idx_SC	dim_customer	4	Composite ...	Business	Never	

- Properties
- Create Index
- Drop Index
- Delete Metadata and Drop Index
-
- Create via Scheduler
-
- Projects >

Properties

The properties screen (see example below) can be selected via the right-click menu when positioned on an index name in the middle pane. The Update Buttons: **Update <-** and **Update ->** are used to move to the previous and next index respectively. The Update Buttons are not available when browsing all indexes in a group, project or repository.



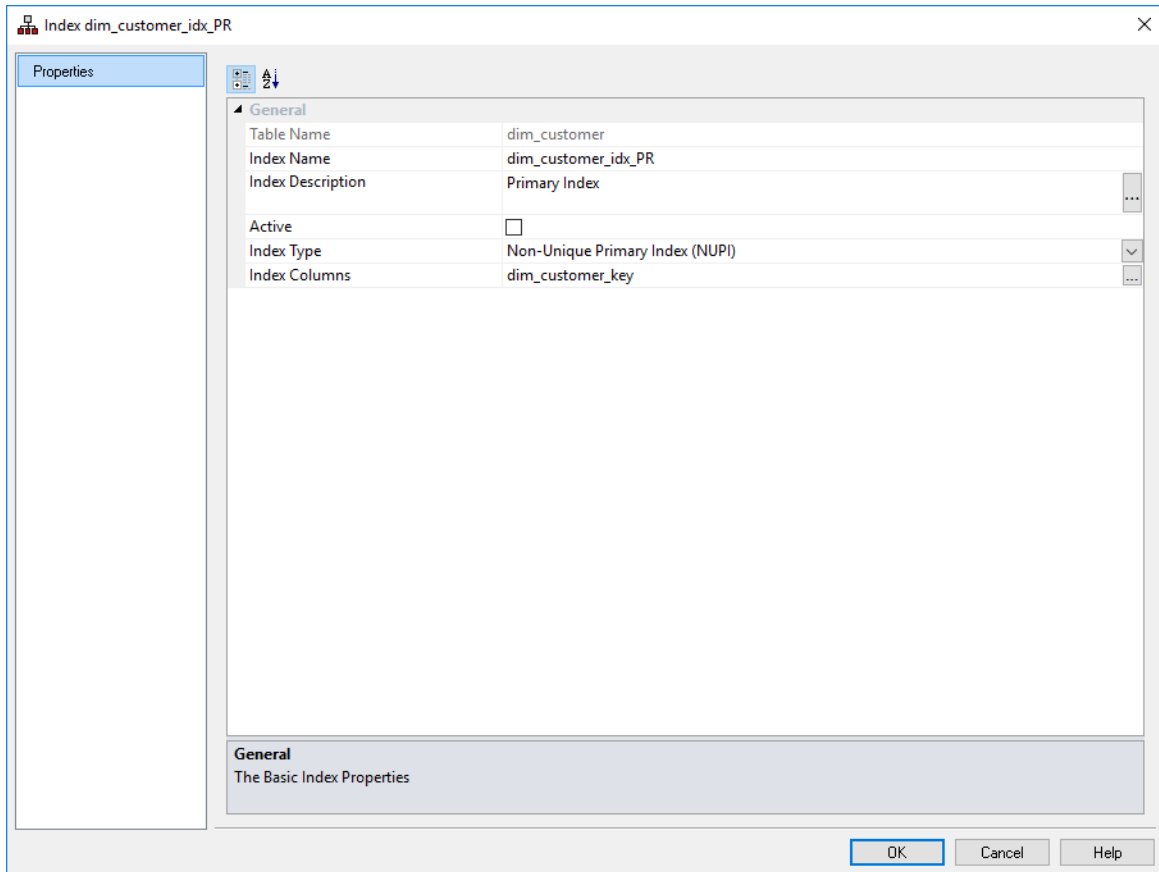
The fields are described below:

Field	Description
Index name	Typically, the table short name followed by: <ul style="list-style-type: none"> • _idx_0 indicating primary key • _idx_x where x = any letter a to z indicating business keys • _idx_PR indicating primary index
Index description	Free flow description of the index

Field	Description
Rebuild frequency	When the index is rebuilt by the scheduler. Select an option from the drop-down list: <ul style="list-style-type: none"> • Never (default) • Daily • End Month • Start Month • Saturday • Sunday • Monday • Tuesday • Wednesday • Thursday • Friday
Active checkbox	<ul style="list-style-type: none"> • When selected means the index is in use. • When not selected means the index is not managed by the scheduler. • The active checkbox on a primary index may not be turned off as all tables in Teradata require a primary index.
Artificial Key	When checked indicates that this is the surrogate (artificial) key generated by the system.
Business Key	Denotes a business key.
Unique	Specifies that the index is a unique index Note: If both unique and artificial are set it is assumed to be a primary key constraint and it is added as such.
Primary Index	Specifies that the index is a Primary Index
Drop before Update	The index is dropped before the update procedure runs on the table and is reinstated after the update is completed. The Drop before update checkbox on a primary index may not be turned on as all tables in Teradata require a primary index.
Hash Index	Defines the index as a Teradata hash index. This limits the hash-ordering to one column, rather than all columns of the index (the default).
Index columns	Shows the columns in the order that will be applied to the index. The order can be changed using the up/down buttons on the left. For a primary index without any indexed columns, the table is created as NOT PRIMARY INDEX table. See example below.
Table columns	Shows all columns in the table that can be indexed. These table columns can be added or removed by highlighting the column and checking the appropriate button.

Indexes are normally managed by the scheduler as part of the normal processing of a table.

Below is an example of a **NO PRIMARY INDEX** index definition:



CHAPTER 28

DOCUMENTATION AND DIAGRAMS

WhereScape RED includes the ability to document the data warehouse, based on the information stored against the metadata for all the tables and columns in the data warehouse.

The documentation will only be meaningful if information is stored in the meta data. The business definition and a description should be stored against all columns that will be visible to end users.

The following sections describe how to **generate** (see "**Creating Documentation**" on page 760) and **read** (see "**Reading the Documentation**" on page 764) the documentation.

IN THIS CHAPTER

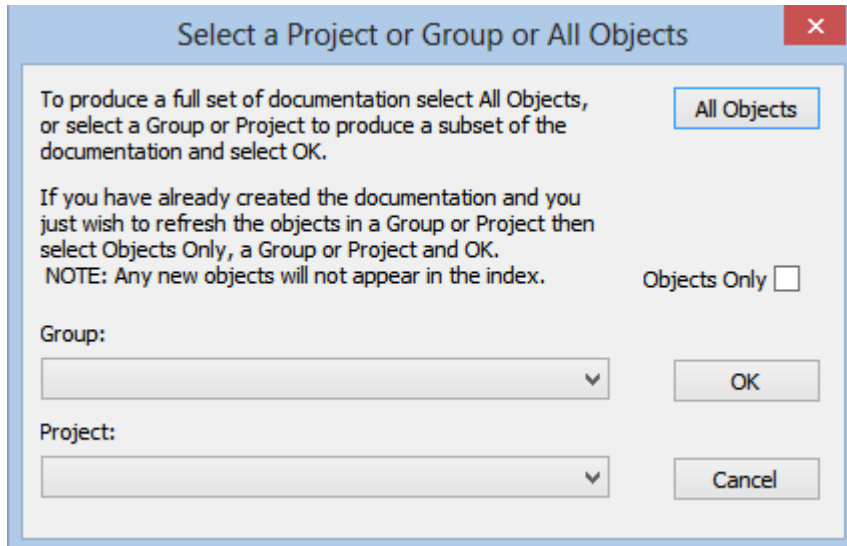
Creating Documentation	760
Batch Documentation Creation	763
Reading the Documentation.....	764
Diagrams	765

CREATING DOCUMENTATION

Create documentation

To create the documentation for the components of the data warehouse, select **Doc** from the builder menu bar, then **Create Documentation**.

If the repository has projects or groups in use, then the following dialog appears to enable the selection of a specific group or project. The default is all objects in the repository.



A file dialog appears next. Select a file path (directory) where the saved the HTML files that will be produced.

A style sheet called **mainstyle.css** is created if it does not exist. If this style sheet exists then it is not overwritten. The style sheet can therefore be modified to match an existing intranet standard.

The next screen allows for the inclusion of a banner and user defined links. It also provides some build options:

The Sizes checkbox instructs the documentation generator to examine the size of all tables and indexes in the database. This process may be slow in some situations, so should normally only be used for final documentation.

The sorted checkbox sorts the columns within the tables into alphabetical order. By default, the columns are in the order that they appear within the tables.

Creating a header

If you check the banner frame option, then a banner (heading) is created at the top of each page in the documentation. You will be prompted for height of the banner frame in pixels (default is 60), an image file (jpg or gif) and any banner text. It is recommended that any image be relatively small (60 pixels high or approximately 1/2 an inch) as it will appear on every page.

Adding Links

Custom information can be linked into the generated documentation.

This means that every time the documentation is regenerated, custom information will be included. In this way the complete documentation for the data warehouse can be viewed in one location.

If you select the add Links option, then you will be prompted to include personalized links from index pages. These links must be to previously created HTML files.

Index pages (linkage points) are available at three points:

- index - initial page
- techindex - technical documentation initial page
- indexuser - user documentation initial page

Document Links

To include any personalized links from the index page complete the details below.

Please enter the text which will form the html link

Please enter the html filename to link to (Either a filename which is relative to the documentation directory or use the browse button for an absolute filename.)

Finished

More

Cancel

Browse

Delete

Multiple links can be added to each index page by using the More option.

Adding glossary elements

As part of the user documentation a glossary is produced defining all the business terms and column names used in the data warehouse. This glossary is primarily based off the columns in the model tables. Additional information can, however, be added via the 'Ws_Api_Glossary' procedure defined in the procedures chapter. This procedure allows the manual inclusion of glossary elements that will be stored in the metadata repository and added to the glossary whenever the documentation is recreated.

BATCH DOCUMENTATION CREATION

WhereScape RED includes the ability to document the data warehouse based on the information stored against the metadata for all the tables and columns in the data warehouse. In a larger environment, it may be a good idea to generate documentation in batch mode.

The following syntax chart illustrates the options available:

```
med.exe /BD { /U UserName { /P Password } } /C OdbcSource /M
MetaDatabase { /N FullName } /D Directory { /G GroupName | ProjectName }
/S NumHops { /I- }
```

Note: {} indicates an optional parameter and | indicates alternative values.

Parameter Descriptions

The following parameters are available:

Parameter	Specify Value?	Mandatory?	Description
BD	No	Yes	Indicates batch documentation mode.
U	Yes	Sometimes *1	Username parameter.
P	Yes	Sometimes *1	Password parameter.
C	Yes	Yes	ODBC data source parameter.
M	Yes	Yes	Metadata database parameter.
N	Yes	No	Full user name parameter, only for logging purposes.
D	Yes	Yes	Directory name where documentation is created.
G	Yes	No	Group or Project name if specified. All Objects if not included.
S	Yes	Yes	Number of processes/hops in the source diagrams.
I-	No	No	Exclude Impact Analysis on load tables (reduces the time required to generate documentation)

Note: User Name and Password are not required when using a trusted connection or operating system authentication.

Example

The following example connects to a Teradata repository using the WslWarehouse ODBC data source, a username of Jack, a password of fly1nG, a metadata database of ProdMeta and generates documentation into the C:\Temp\my_doco directory with 4 hops in diagrams:

```
med.exe /BD /UJack /P fly1nG /C "WslWarehouse" /M "ProdMeta" /D  
"C:\Temp\my_doco" /S "4"
```

READING THE DOCUMENTATION


To read the documentation you have created, select **Doc** from the builder menu bar, then **Read Documentation**. This will launch a browser and display the contents of index.html. Alternatively you can access the HTML pages directly from their saved location.

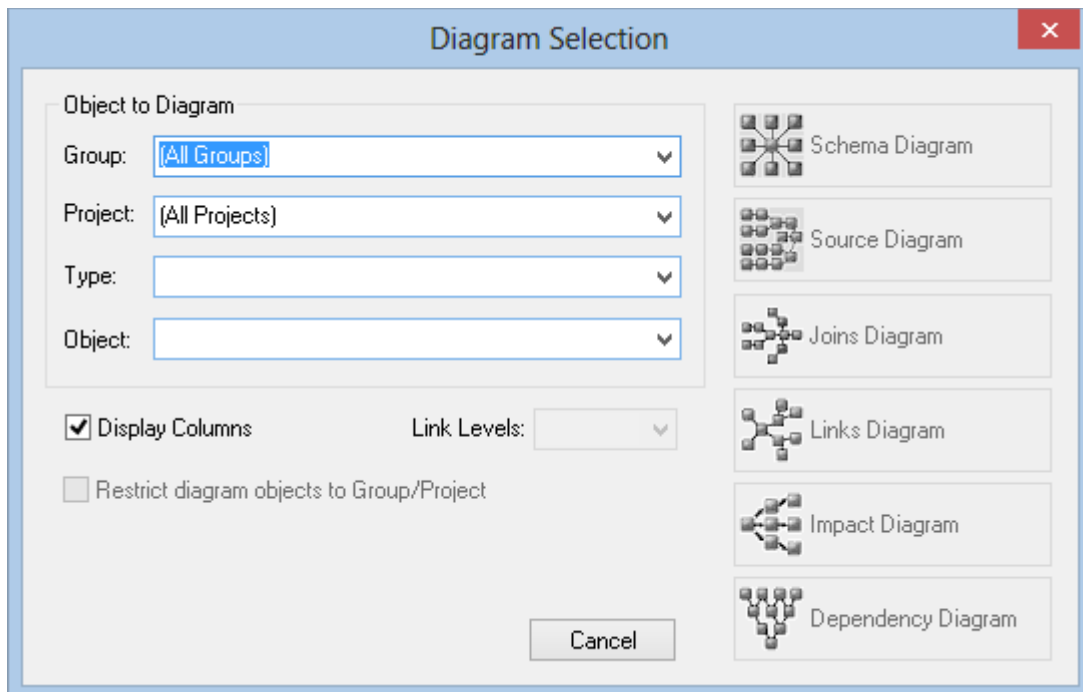
DIAGRAMS

TYPES OF DIAGRAMS


Six types of diagrams are provided to give visual representation of what has been created. These are:

- The **Schema Diagram** (on page 767)
- The **Source Diagram** (on page 769)
- The **Joins Diagram** (on page 773)
- The **Links Diagram** (on page 774)
- The **Impact Diagram** (on page 775)
- The **Dependency Diagram** (on page 776)

- 1 To display the **Diagram Selection** dialog, click on the diagrammatic view button 



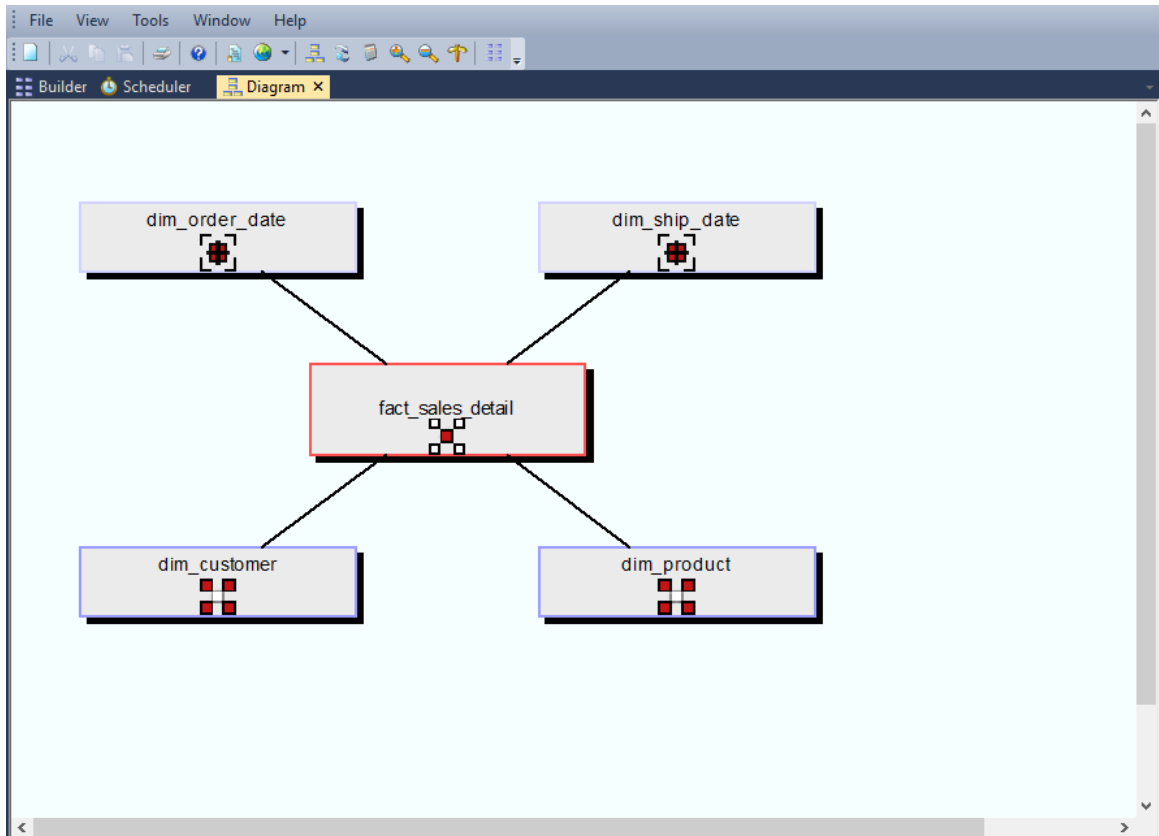
- 2 Choose the object to diagram by optionally choosing the **Type** to limit the selection list; and then selecting the **Object**. The diagram type buttons on the right will then become active and you can choose the type of diagram to display.




SCHEMA DIAGRAM

A star schema diagram can be displayed for a fact_table, aggregate table, fact view or OLAP cube. It shows the central table with the outlying dimensions.

An example of a **Schema** diagram in **Standard Diagram** format is displayed below.



An example of a **Schema diagram in Detail Diagram** format is displayed below.



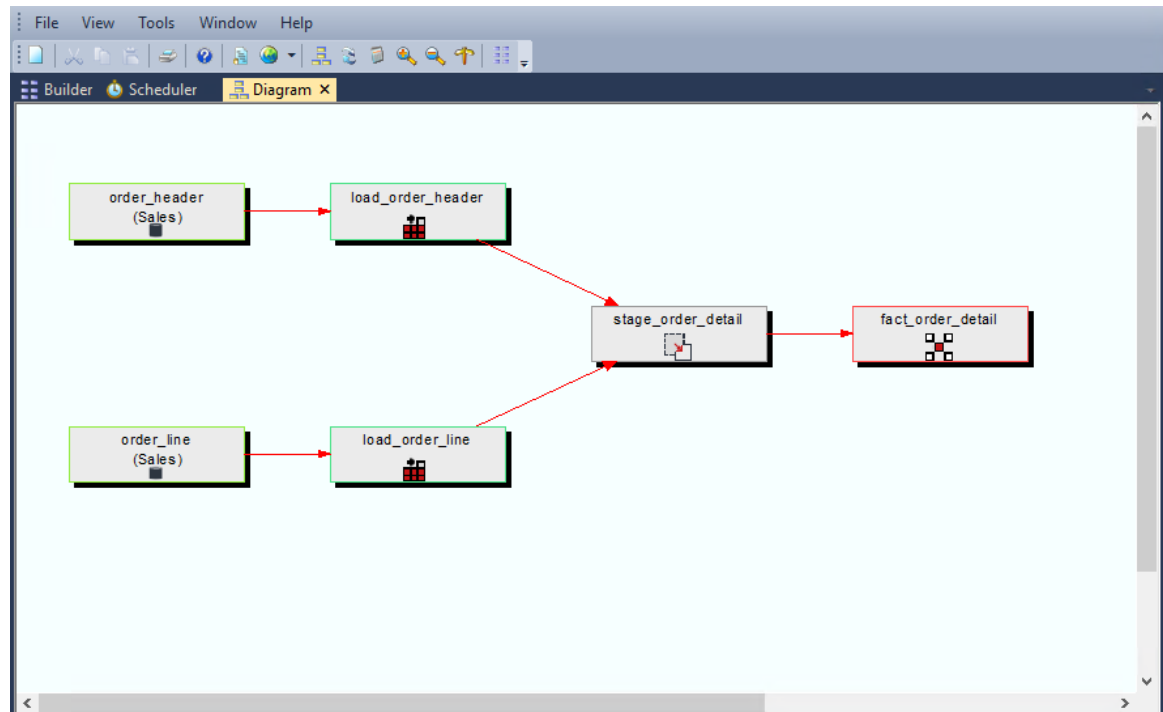
Note: A star schema diagram for a fact view will display only the selected fact view and related dimension views.

The star schema diagrams are produced in **Standard Diagram** format as part of the user and technical documentation when you select **Doc > Create Documentation** from the main builder window.

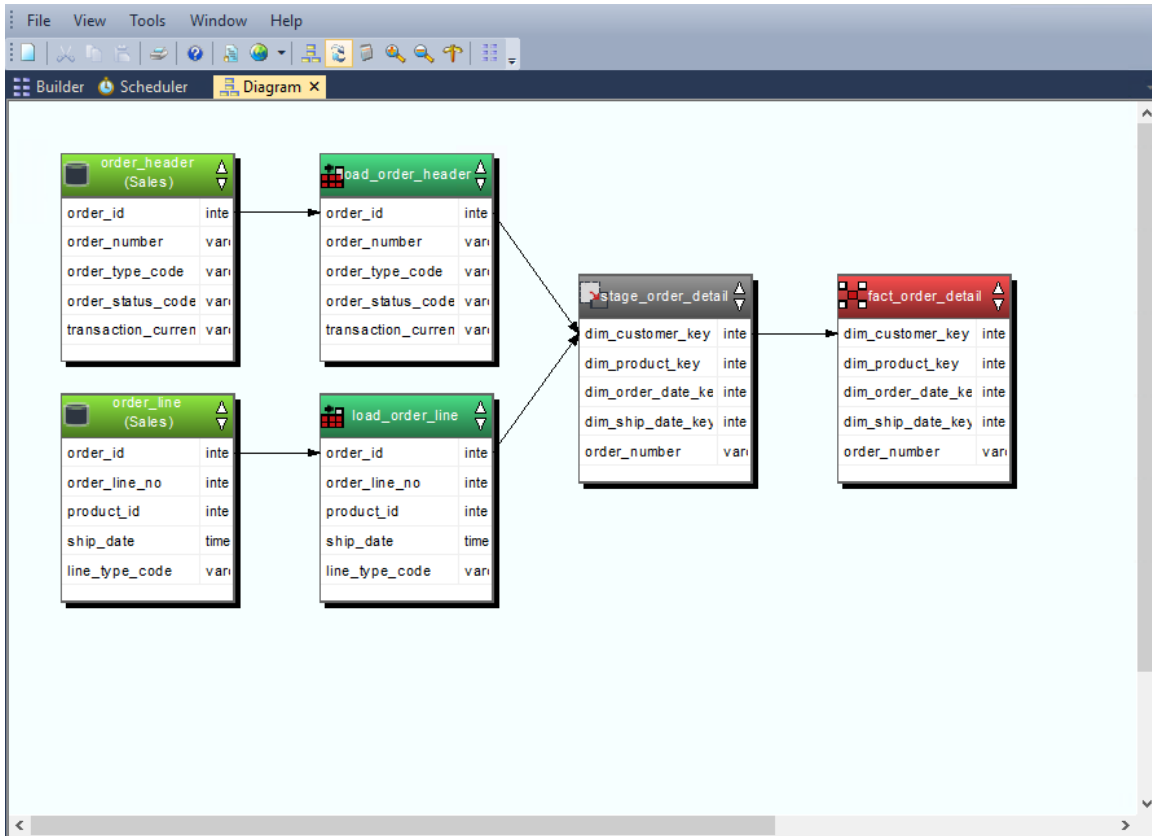
SOURCE DIAGRAM

A source tracking diagram can be displayed for any table. It shows connections back from the chosen table to the source tables from which information was derived. Hovering the cursor over a line shows additional information. For lines going into load tables, the source of the data will be displayed; while for other lines in the diagram, the procedure used to move data between two tables is displayed.

An example of a **Source** diagram in **Standard Diagram** format is displayed below.



An example of a **Source** diagram in **Detail Diagram** format is displayed below.

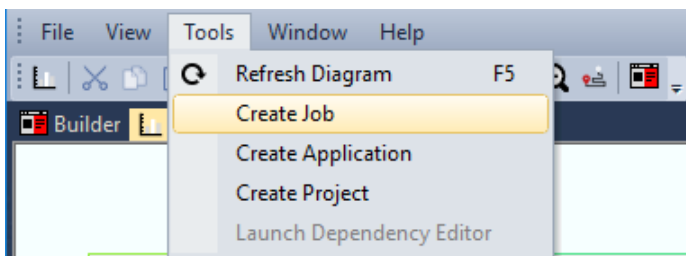


The Source diagrams are produced in **Standard Diagram** format as part of the technical documentation when you select **Doc > Create Documentation** from the main builder window.

Creating a Job from a Source Tracking Diagram

Once a source tracking diagram has been created for a table, a scheduler job can be generated from the diagram. This job will be called **Process_to_table_name**, where `table_name` is the name of the table the track back diagram was run for.

To create a Job, select **Create Job** from the **Tools** menu after the diagram is displayed:



The Job Definition is then displayed:

Job Definition [X]

Job Name:

Description:

Frequency:

Start Date:

Start Time:

Maximum Threads:

Scheduler:

Dependent On:

Logs Retained: This field lets you set the number of logs that are retained for this job before an automatic delete and archive occurs. 0 = keep all logs (default action)

The following two fields are optional. They are executed after the job completes and therefore need to reflect the scheduler environment. (i.e. Unix or Windows). The special variables \$JOB_KEY\$, \$JOB_SEQ\$ and \$JOB_NAME\$ can be used to return the associated values.
The Success command will be executed if a successful completion, the failure command will be executed if a job fails to complete:

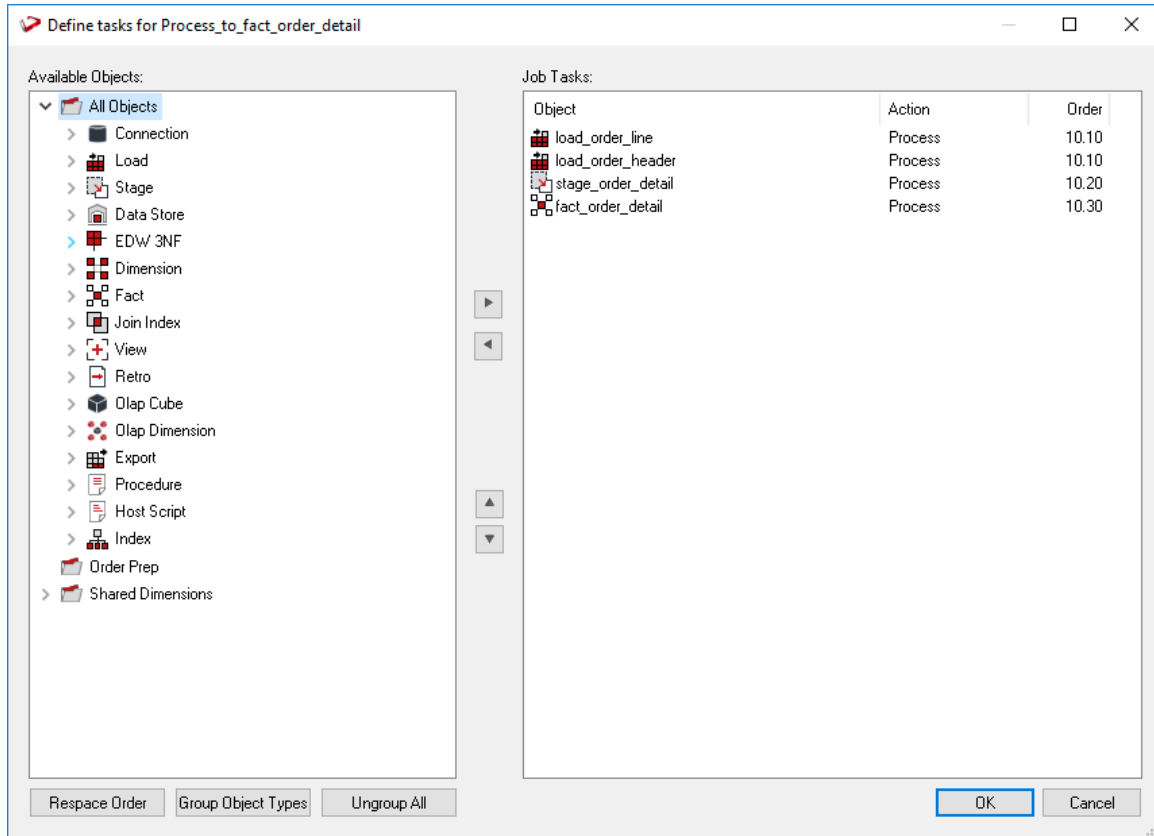
Success Command:

Failure Command:


Execute Failure Command in event of dependency failure

Make any changes here that are required and click **OK**.

For the diagram above, a job is created with the following **tasks**:



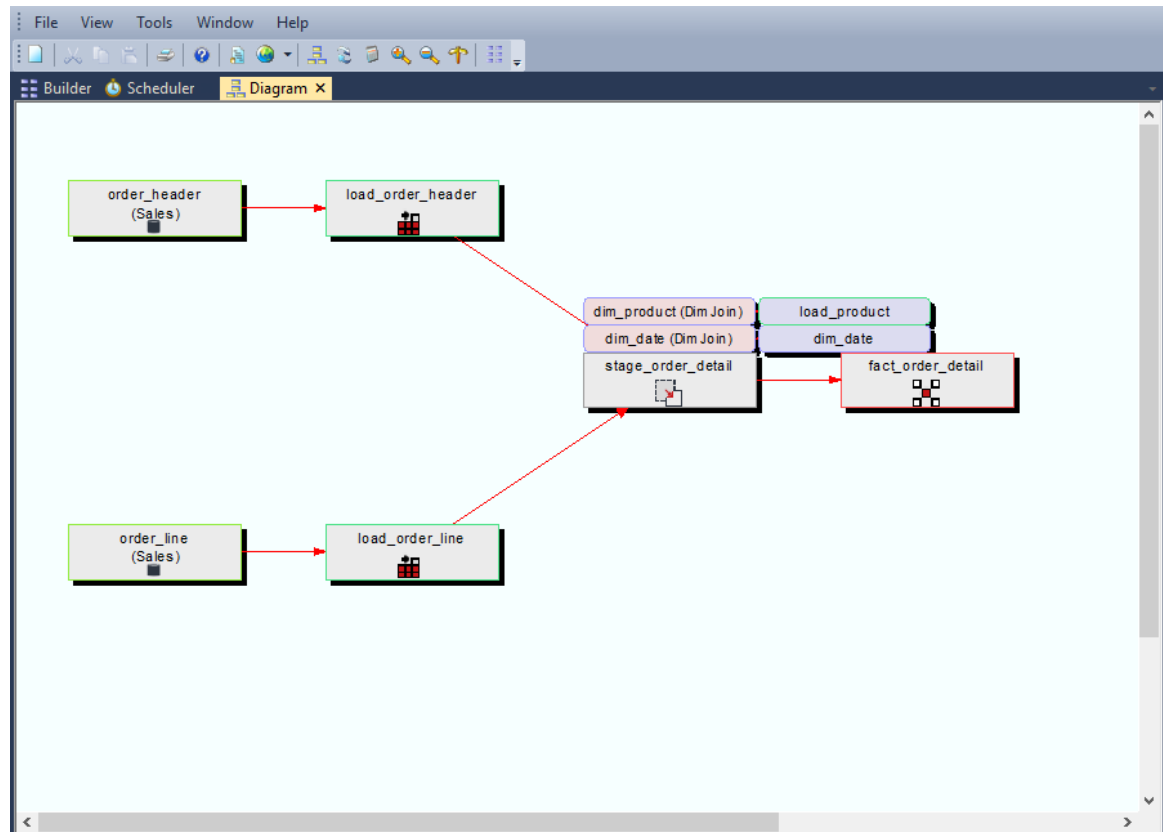
Note: It is also possible to display the source tracking diagram by right-clicking on a table and choosing **Impact/Track Back Diagram**:



JOINS DIAGRAM

A data join track back diagram can be displayed for any table. It shows connections back from the chosen table to the source tables from which information was derived and includes dimension table joins. Hovering the cursor over a line shows additional information. For lines going into load tables, the source of the data will be displayed; while for other lines in the diagram, the procedure used to move data between two tables is displayed.

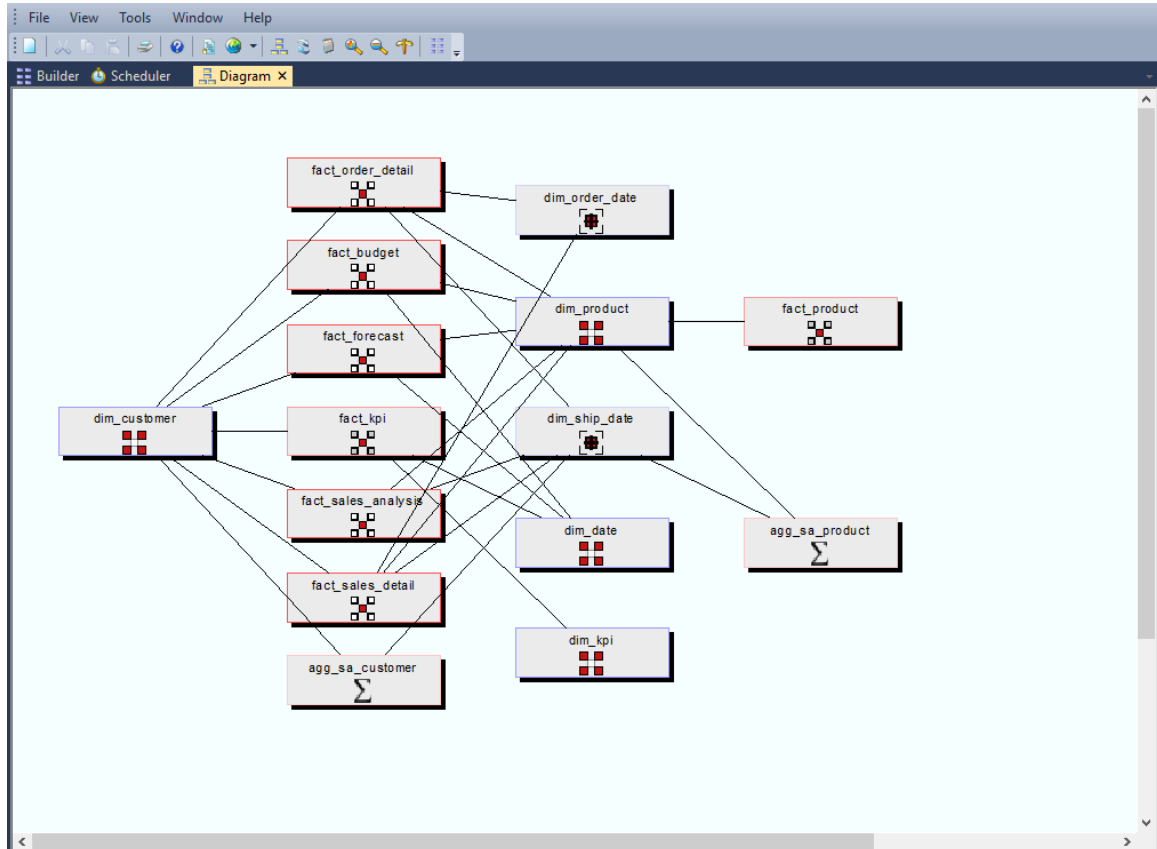
An example of a **Joins** diagram in **Standard Diagram** format is displayed below.



LINKS DIAGRAM

A linked tables diagram can be displayed for any table. It shows relationships between tables, looking out from the chosen table a selected number of hops. The number of hops is determined by table relationships and source and target relationships.

An example of a **linked tables** diagram in **Standard Diagram** format is displayed below.



Notes:

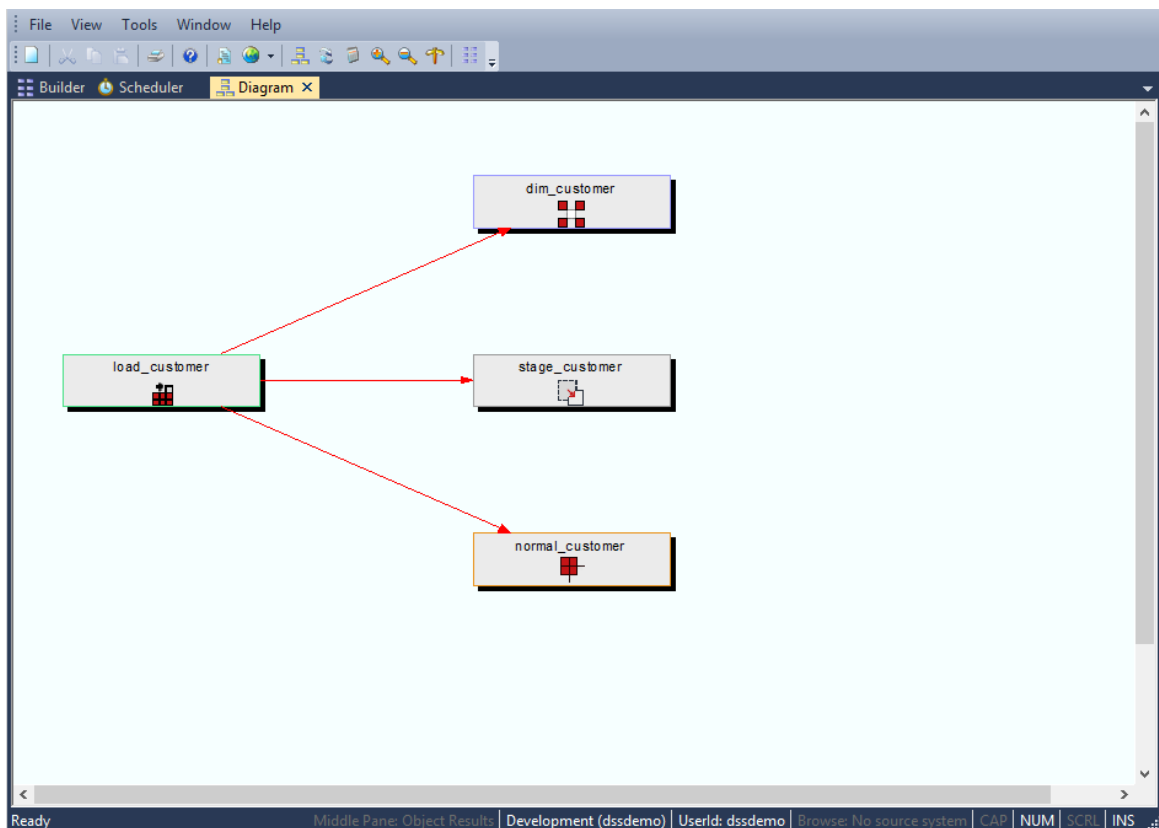
1. A linked table diagram for a model object will display only related model objects, aggregate objects and views of type "Model View".
2. A linked table diagram for an aggregate object will display only related model objects, aggregate objects and views of type "Model View".
3. A linked table diagram for a view of type "Model View" will display only related model objects, aggregate objects and views of type "Model View".
4. A linked table diagram for a view of type "View" will display only related views of type "View".

IMPACT DIAGRAM


A track forward impact diagram can be displayed for any table. It shows connections forward from the chosen table to the subsequent tables built with columns from this table.

A track back impact diagram can be displayed for any table. It shows connections backwards from the chosen table to the previous tables.

An example of an Impact diagram in **Standard Diagram** format is displayed below.



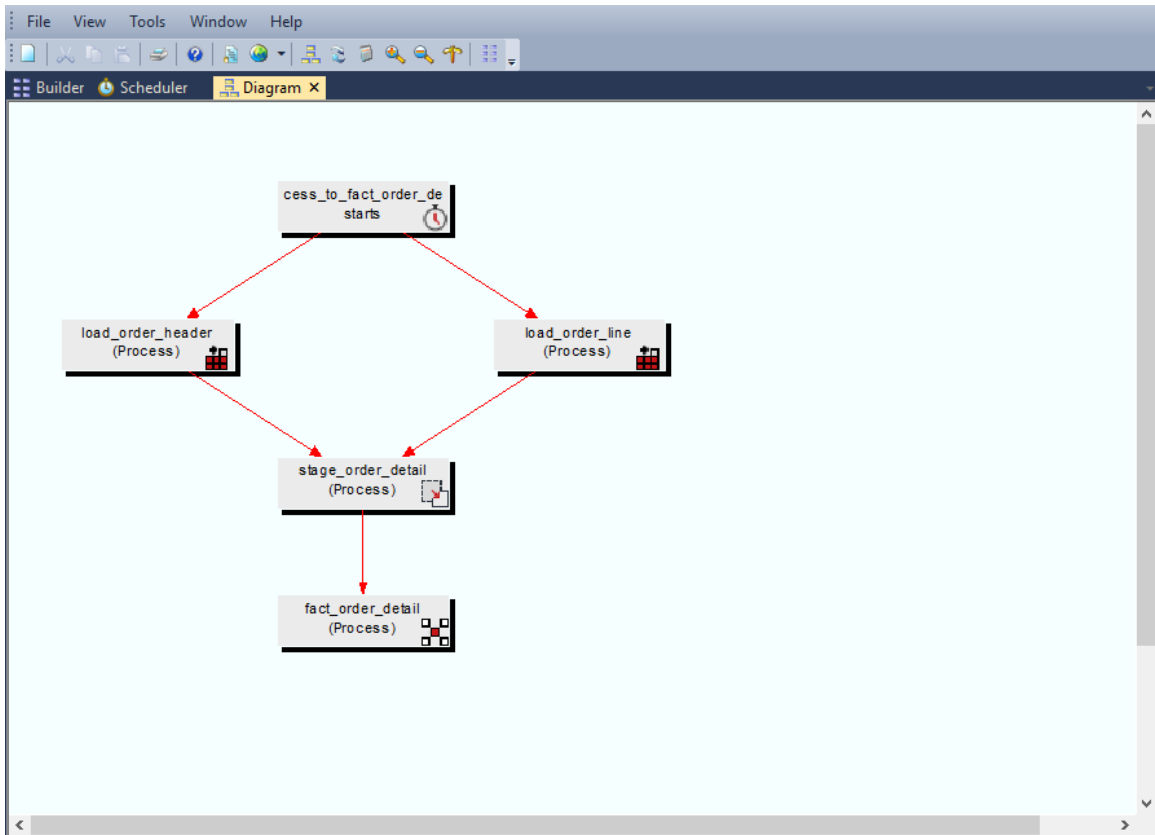
Note: It is also possible to display the track back / forward diagram by right-clicking on a table and choosing **Impact > Track Back Diagram** or **Impact > Track Forward Diagram**:



DEPENDENCY DIAGRAM

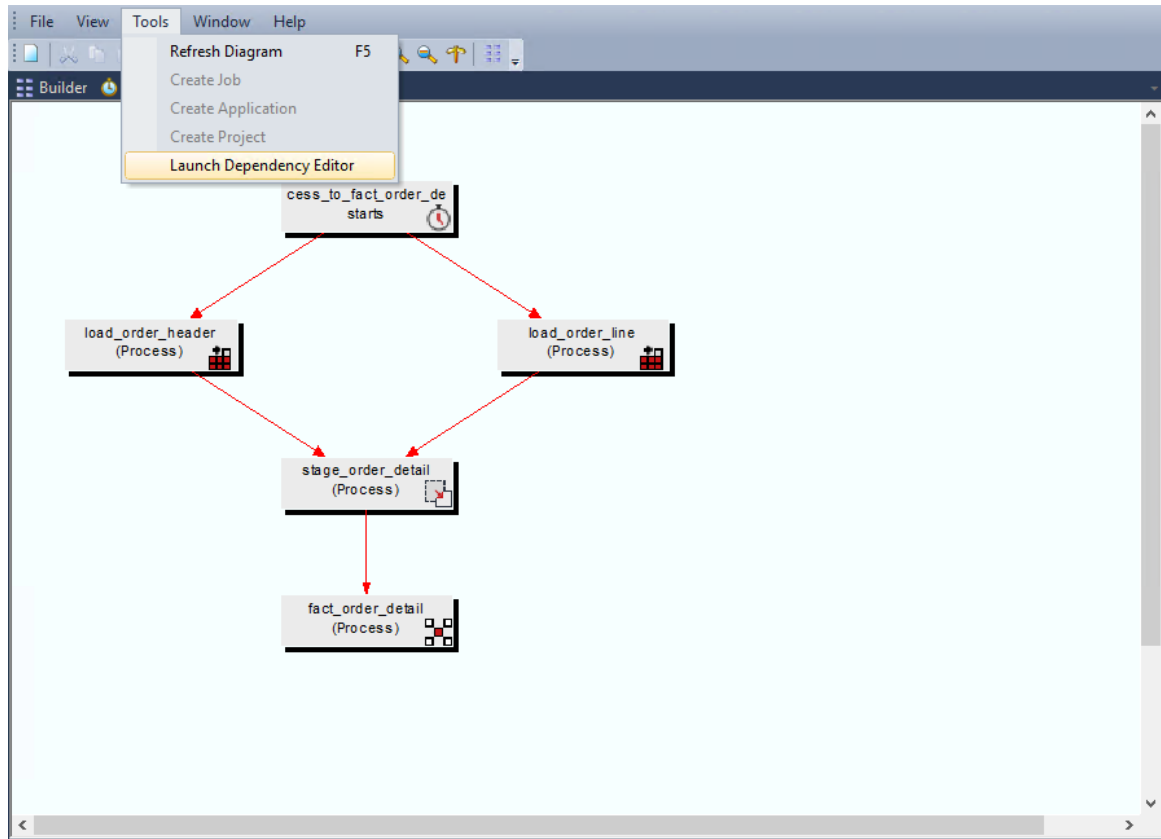
A job dependency diagram can be displayed for any job defined in the WhereScape RED scheduler. It shows the parent and child relationships between tasks within a job.

An example of a **Dependency** diagram in **Standard Diagram** format is displayed below.



Editing a Job's Dependencies from a Job Dependency diagram

Once a Job Dependency diagram has been created for a job, its dependencies can be edited from the diagram. To do this, select **Launch Dependency Editor** from the Tools menu:



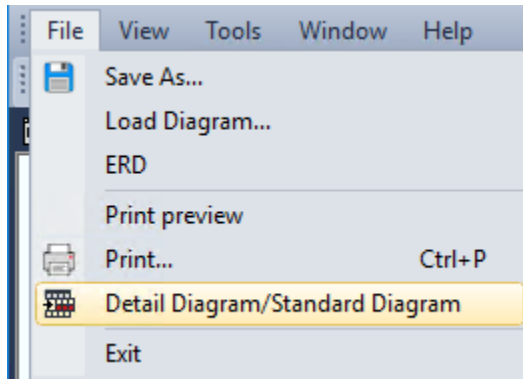
The dialog that follows allows you to edit the dependencies within a job:

Dependencies				
Parent Task	Parent Action	Map	Child Task	Child Action
load_product	Process	-->	dim_customer	Process
load_product	Process	-->	dim_date	Process
load_product	Process	-->	dim_product	Process
load_customer	Process	-->	dim_customer	Process
load_customer	Process	-->	dim_date	Process
load_customer	Process	-->	dim_product	Process
load_order_line	Process	-->	dim_customer	Process
load_order_line	Process	-->	dim_date	Process
load_order_line	Process	-->	dim_product	Process
load_order_header	Process	-->	dim_customer	Process
load_order_header	Process	-->	dim_date	Process
load_order_header	Process	-->	dim_product	Process
dim_customer	Process	-->	stage_sales_detail	Process
dim_date	Process	-->	stage_sales_detail	Process
dim_product	Process	-->	stage_sales_detail	Process
stage_sales_detail	Process	-->	fact_sales_detail	Process
fact_sales_detail	Process	-->	fact_sales_detail	Analyze

WORKING WITH DIAGRAMS

Diagram Display

Once displayed there are two modes for diagram display; standard and detailed. To move between displays, select **File > Detail Diagram / Standard Diagram (Toggle)**



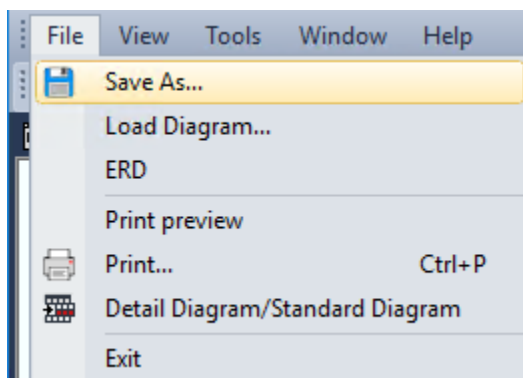
or use the **toggle** button.



Diagram Save

A diagram can be saved either as a meta file or as a jpeg image. If saved as a meta file, it can be subsequently reloaded and re-edited. A jpeg cannot be reloaded into WhereScape RED.

The diagram can be saved by selecting **File > Save As ...**



Note: By default the diagram is saved as a meta file.

Diagram Load

If a diagram has been saved as a meta file, it can be reloaded. To reload a saved meta file, switch to diagrammatic view and select the menu option **File > Load Diagram**. A dialog box will allow you to choose a windows meta file (*.wmf). If the meta file had previously been saved in WhereScape RED, then the diagram will be loaded.

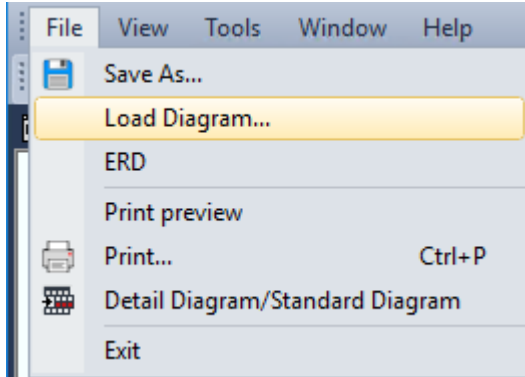


Diagram Print

A diagram can be printed by selecting **File > Print ...**

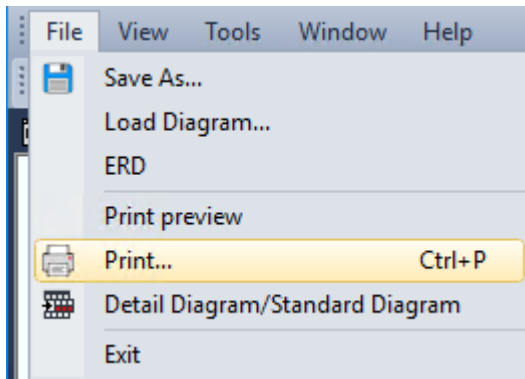
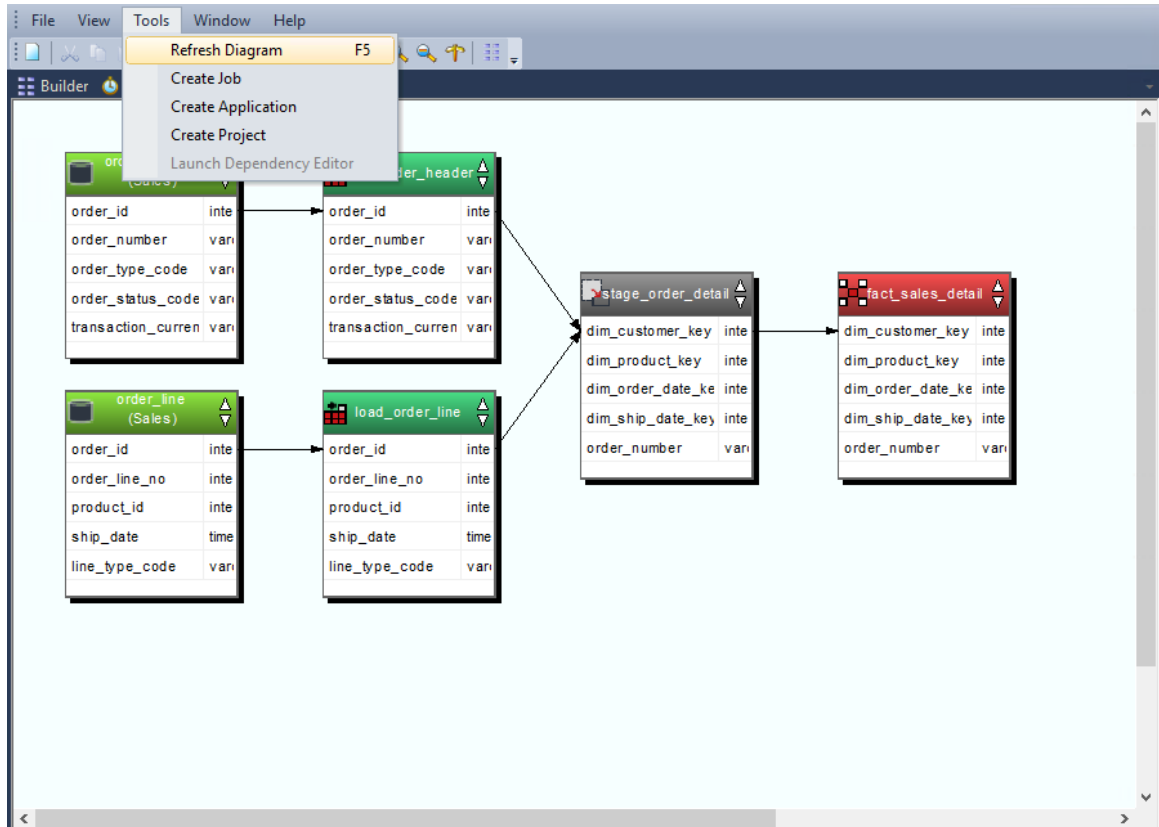


Diagram Refresh

Once a diagram has been displayed, it can be refreshed by choosing **Tools > Refresh Diagram**, or by pressing **F5**.

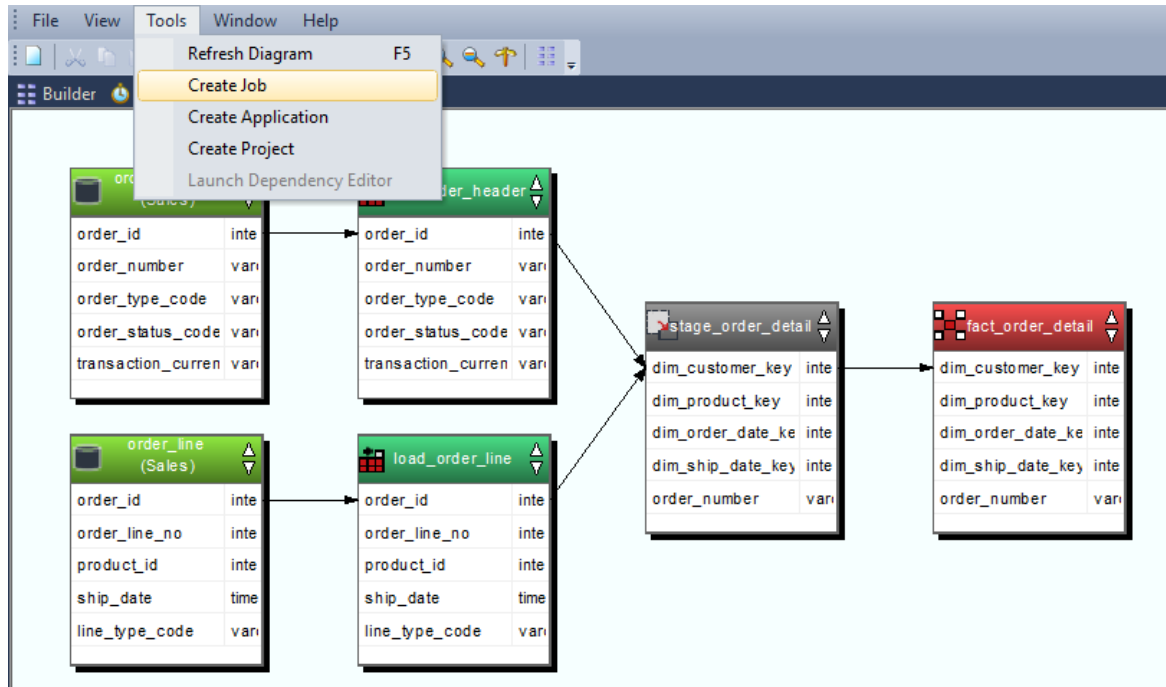


CREATING A JOB FROM A DIAGRAM

A job can be created from a **Source** diagram or a **Joins** diagram.

To Create a Job from a Diagram

- 1 Once the diagram is displayed, select **Tools/Create Job**.



2 Edit the job as required and click OK.

Job Definition ✕

Job Name:

Description:

Frequency:

Start Date:

Start Time:

Maximum Threads:

Scheduler:

Dependent On:

Logs Retained: This field lets you set the number of logs that are retained for this job before an automatic delete and archive occurs. 0 = keep all logs (default action)

The following two fields are optional. They are executed after the job completes and therefore need to reflect the scheduler environment. (i.e. Unix or Windows). The special variables \$JOB_KEY\$, \$JOB_SEQ\$ and \$JOB_NAME\$ can be used to return the associated values.
The Success command will be executed if a successful completion, the failure command will be executed if a job fails to complete:

Success Command:

Failure Command:

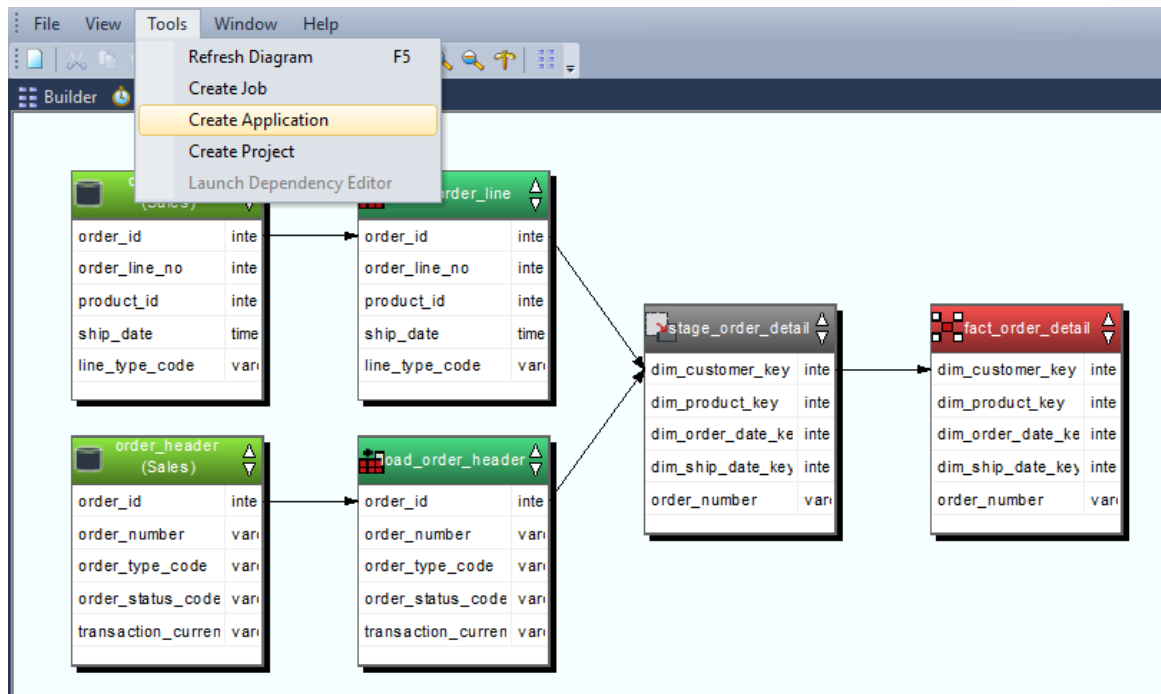
Execute Failure Command in event of dependency failure

CREATING AN APPLICATION FROM A DIAGRAM

An application can be created from a **Source** diagram or a **Joins** diagram.

To Create an Application from a Diagram

- 1 Once the diagram is displayed, select **Tools/Create Application**.



2 Edit the application as required and then click **OK**.

Build Deployment Application

This process builds the files necessary to allow the deployment of a data warehouse solution. These files are read and processed by the Setup Administrator utility. Specify the application identification details and then select the objects to be deployed to and/or deleted from another repository.

Output Directory:

Application Identifier: Application Version:

Application Name:

Description:

You can select or enter a previous application file as a starting point for this application.

Previous Application:

Pre Application Load SQL. (the following optional SQL statement will be issued before the application load commences):

Post Application Load SQL. (the following optional SQL statement will be issued after the application load completes):

3 A dialog will display, confirming the creation of the application files. Click **OK**.

WhereScape RED

i The following files have been created in:
C:\ProgramData\WhereScape\Application\

app_obj_fact_order_d_130822101354.wst
app_data_fact_order_d_130822101354.wst
app_id_fact_order_d_130822101354.wst
app_con_fact_order_d_130822101354.wst
app_map_fact_order_d_130822101354.wst
app_del_fact_order_d_130822101354.wst
app_ext_fact_order_d_130822101354.wst

These files can be moved to another location or system if required.
All files must be moved for the application to load.

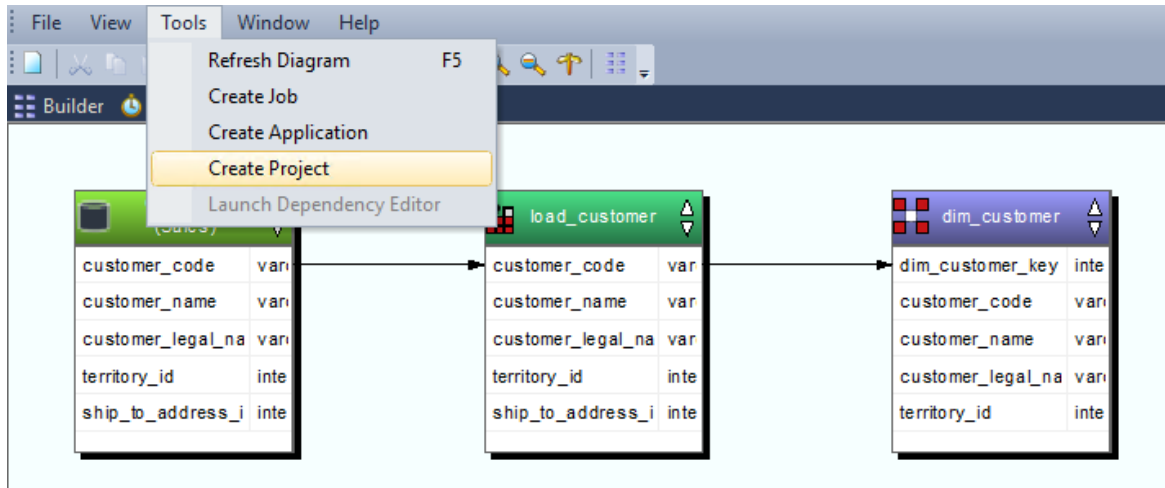
Note: Creating an application from a diagram will save the objects in the diagram and the associated objects, including indexes.

CREATING A PROJECT FROM A DIAGRAM

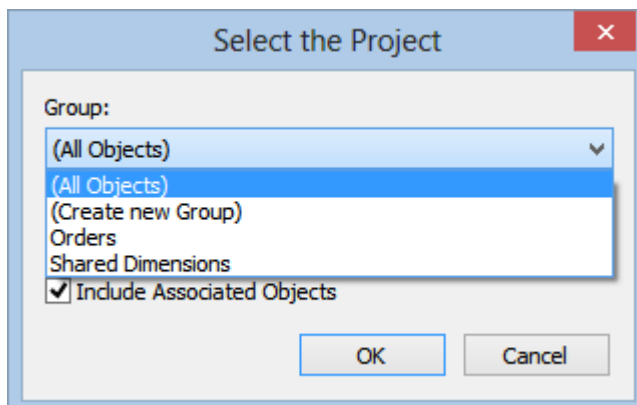
A project can be created from a **Source** diagram or a **Joins** diagram.

To Create a Project from a Diagram


- 1 Once the diagram is displayed, select **Tools/Create Project**.



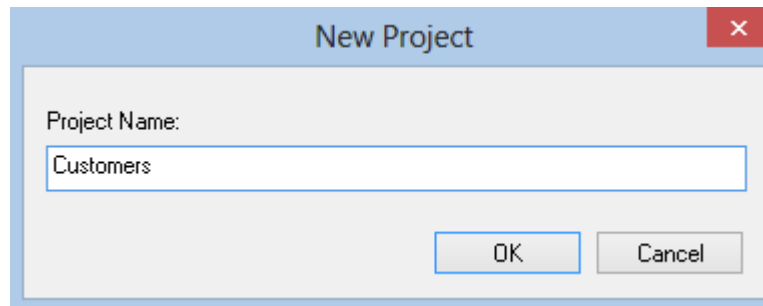
- 2 Select an existing group or create a new group.



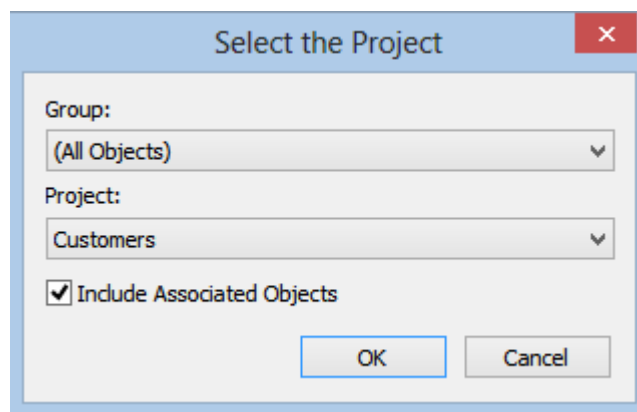
- 3 Select to **Create new Project**.



- 4 Enter the name of the new project and click **OK**.



- 5 Click **OK**.



The objects in the diagram will be moved into the selected project. If the **Include Associated Objects** checkbox is selected, this will include all associated procedures, scripts and indexes. The checkbox is selected by default.

CHAPTER 29

REPORTS

WhereScape RED includes reports for analyzing columns, tables, procedures, indexes, objects and jobs.

When these reports are run, the results are displayed in a separate tab in the bottom pane of the RED screen.

The following sections describe the purpose, parameters and results for each report.

IN THIS CHAPTER

Dimension-Fact Matrix.....	788
OLAP Dimension-Cube Matrix	789
Dimension Views for a Specified Dimension.....	790
Column Reports	791
Table Reports.....	799
Procedure Reports	806
Object Reports	808
Job Reports.....	816
Operational Reports.....	820

DIMENSION-FACT MATRIX

This report shows dimension tables used by each fact and aggregate table in the metadata as a matrix.

Objects Included

The following WhereScape RED object types are included in this report:

- Dimension Tables
- Fact Tables
- Aggregate Tables
- Dimension Views

Note: WhereScape RED version 8.1.1.0 and above no longer supports the Dimension View object type when creating new objects. However, users that have existing Dimension View objects in their data warehouse can retain and continue to use them.

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of refreshed objects in the metadata repository with the following columns:

- Dimensions (*the dimension name*)
- Fact/Aggregate Table 1
- Fact/Aggregate Table 2
- ...
- Fact/Aggregate Table *n*

The cells in the crosstab have a 1 to indicate the dimension is used by the fact or aggregate table and blank otherwise. The result set is not sortable.

Report Example

Dimension	fact_budget	fact_forecast	fact_sales_anal...	fact_sales_detail	agg_sa_custo...	agg_sa_product
dim_customer	1	1	1	1	1	
dim_date	1	1				
dim_order_date				1		
dim_product	1	1	1	1		1
dim_ship_date			1	1	1	1

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

OLAP DIMENSION-CUBE MATRIX

This report shows the relationships between cube measure groups and OLAP dimensions in the metadata as a matrix.

Objects Included

The following WhereScape RED object types are included in this report:

- OLAP Dimensions
- OLAP Measure Groups

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of refreshed objects in the metadata repository with the following columns:

- Dimensions (*the dimension name*)
- Measure Group 1
- Measure Group 2
- ...
- Measure Group *n*

The cells in the crosstab have a value to indicate the relationship else 'No Relationship' if no relationship exists between the Measure Group and the OLAP Dimension. The result set is not sortable.

Report Example

Dimensions	fact_customer ...	fact_sales_deta...	fact_orders (fa...
dim_customer			dim_customer
dim_ship_date			dim_ship_date

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

DIMENSION VIEWS FOR A SPECIFIED DIMENSION

Note: WhereScape RED version 8.1.1.0 and above no longer supports the Dimension View object type when creating new objects. However, users that have existing Dimension View objects in their data warehouse can retain and continue to use them.

This report shows Views built on a specified Dimension table.

Objects Included

The following WhereScape RED object types are included in this report:

- Views (Dimension View table type only)

Parameters

This report has one parameter:

- Dimension table name

Results

The results of this report are displayed as a list of objects in the metadata repository with the following columns:

- Dimension Name (*the name of the Dimension table*)
- Dimension View Names

The result set is sortable by clicking on the appropriate column heading.

Report Example

Dimension Name	Dimension View Names
dim_date	dim_order_date
dim_date	dim_ship_date

Sending Results to Microsoft Excel

Right-click the result set and click **Output to File** to send the results of this report to Microsoft Excel.


COLUMN REPORTS

There are five reports for analyzing Columns:

- Columns Without Comments
- All Column Transformations
- Re-Usable Column Transformations
- Column Track-Back
- Column Track-Forward

COLUMNS WITHOUT COMMENTS

This report shows **table** object columns in the metadata that don't have descriptions. The report prompts the user with the option to only display user facing table object columns or all table object columns.



Objects Included

The following WhereScape RED object types are included in this report:

- Views
- Aggregate Tables
- Join Indexes
- Hubs, Satellites and Links

Note: The EUL object are set in **Tools>Options>Object Types>Object Type End User Setting** (see "**Object Type End User Setting**" on page 86).

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of objects and columns in the metadata that are missing comments with the following columns:

- Table name (*the name of the table*)
- Column Name
- Table type

Report Example

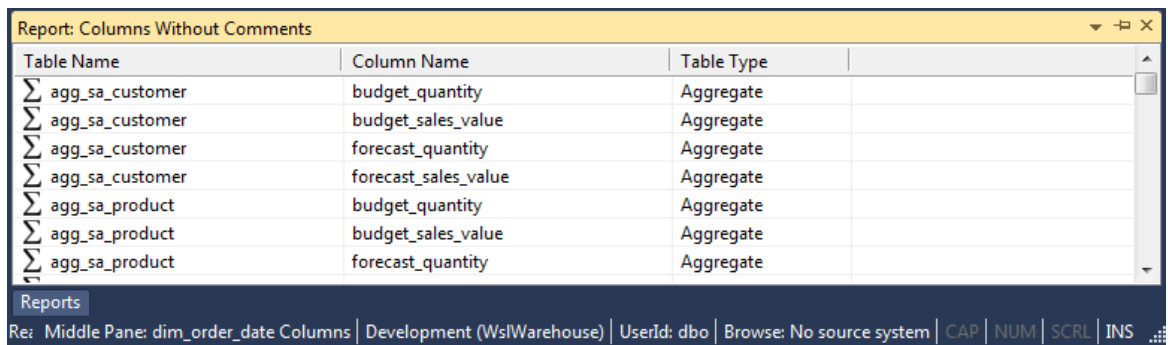


Table Name	Column Name	Table Type
agg_sa_customer	budget_quantity	Aggregate
agg_sa_customer	budget_sales_value	Aggregate
agg_sa_customer	forecast_quantity	Aggregate
agg_sa_customer	forecast_sales_value	Aggregate
agg_sa_product	budget_quantity	Aggregate
agg_sa_product	budget_sales_value	Aggregate
agg_sa_product	forecast_quantity	Aggregate

Reports

Re: Middle Pane: dim_order_date Columns | Development (WslWarehouse) | UserId: dbo | Browse: No source system | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

ALL COLUMN TRANSFORMATIONS

This report shows all columns that have a column transformation and the details of the transformation.

Objects Included

The following WhereScape RED object types are included in this report:

- Load Tables
- Stage Tables
- Aggregate Tables
- Join Indexes
- Views
- Exports
- Hubs, Satellites and Links

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of objects in the metadata repository with the following columns:

- Table name (*the name of the table*)
- Column name
- Transformation

The result set is sortable by clicking on the appropriate column heading.

Report Example

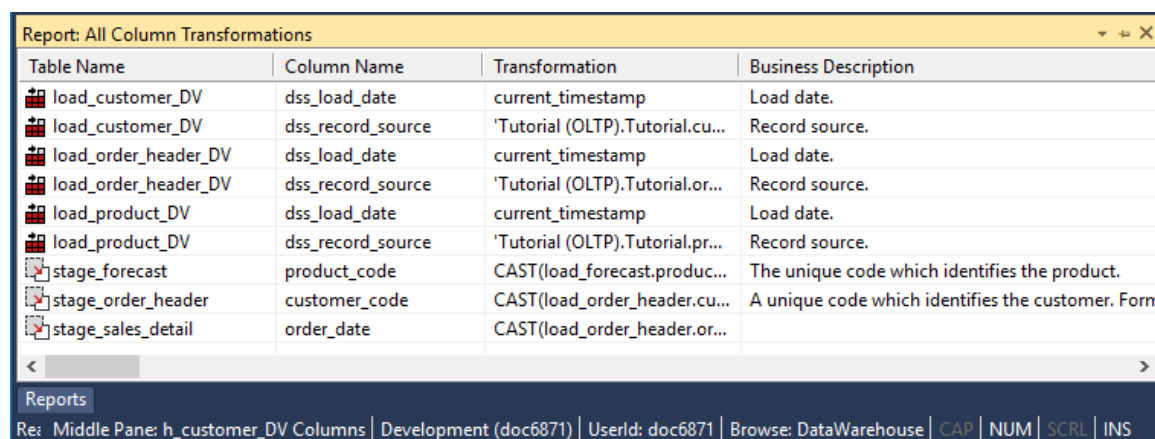


Table Name	Column Name	Transformation	Business Description
load_customer_DV	dss_load_date	current_timestamp	Load date.
load_customer_DV	dss_record_source	'Tutorial (OLTP).Tutorial.cu...	Record source.
load_order_header_DV	dss_load_date	current_timestamp	Load date.
load_order_header_DV	dss_record_source	'Tutorial (OLTP).Tutorial.or...	Record source.
load_product_DV	dss_load_date	current_timestamp	Load date.
load_product_DV	dss_record_source	'Tutorial (OLTP).Tutorial.pr...	Record source.
stage_forecast	product_code	CAST(load_forecast.produc...	The unique code which identifies the product.
stage_order_header	customer_code	CAST(load_order_header.cu...	A unique code which identifies the customer. Form
stage_sales_detail	order_date	CAST(load_order_header.or...	

Reports
Re: Middle Pane: h_customer_DV Columns | Development (doc6871) | UserId: doc6871 | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

RE-USABLE COLUMN TRANSFORMATIONS

This report shows all reusable transformations as defined via **Tools>Define Re-usable Transformations**.

Parameters

There are no parameters for this report.

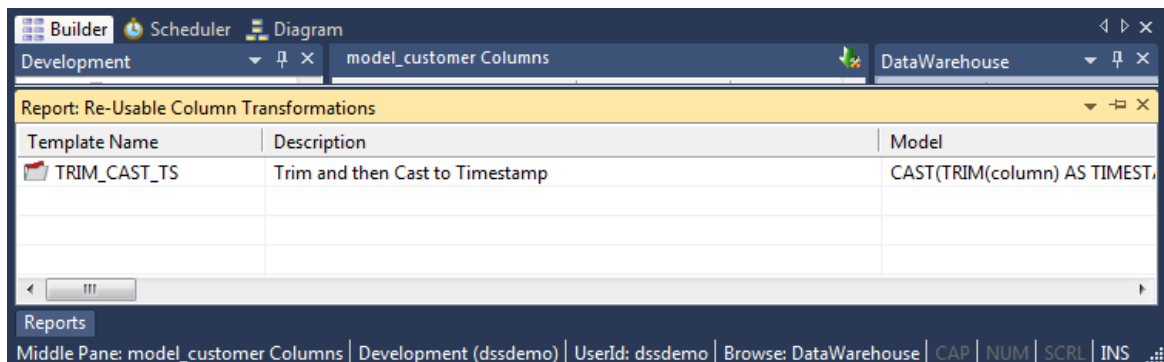
Results

The results of this report are displayed as a list of Re-Usable transformations with the following columns :

- Template Name
- Description

The result set is not sortable.

Report Example



Template Name	Description	Model
TRIM_CAST_TS	Trim and then Cast to Timestamp	CAST(TRIM(column) AS TIMEST...

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

COLUMN TRACK-BACK

This report shows the lineage of a specified column in a specified table, including any transformations.

Objects Included

The following WhereScape RED object types are included in this report:

- Load Tables
- Stage Tables
- Model Tables
- Aggregate Tables
- Join Indexes
- Views
- Exports
- Hubs, Satellites and Links

Parameters

This report has the following parameters:

- Groups
- Projects
- Table
- Column

Column Track Back Report Parameters

Select a Group, Project or Table and Column to list the origins of the objects.

Groups:

Projects:

Table:

Column:

Exclude Intermediate Steps

OK Cancel

Results

If you left the **Exclude Intermediate Steps** checkbox **unchecked**, then the results screen will be as follows, showing the line of origins for the selected tables:

Tables (fact_sales_analysis)	Columns	Source Tables	Source Columns
fact_sales_analysis	dim_customer_key	dim_customer	dim_customer_key
fact_sales_analysis	dim_product_key	dim_product	dim_product_key
fact_sales_analysis	dim_ship_date_key	dim_ship_date	dim_ship_date_key
fact_sales_analysis	dim_ship_date_key	dim_date	dim_date_key
fact_sales_analysis	quantity	fact_sales_detail	quantity
fact_sales_analysis	quantity	stage_sales_detail	quantity
fact_sales_analysis	quantity	load_order_line	quantity
fact_sales_analysis	quantity	order_line	quantity
fact_sales_analysis	sales_value	fact_sales_detail	sales_value
fact_sales_analysis	sales_value	stage_sales_detail	sales_value
fact_sales_analysis	sales_value	load_order_line	sales_value
fact_sales_analysis	sales_value	order_line	sales_value
fact_sales_analysis	tax	fact_sales_detail	tax

If however, you **selected Exclude Intermediate Steps**, then the results screen will be as follows, showing only the original source table for the selected tables:

Tables (fact_sales_analysis)	Columns	Source Tables	Source Columns
fact_sales_analysis	quantity	order_line	quantity
fact_sales_analysis	sales_value	order_line	sales_value
fact_sales_analysis	tax	order_line	tax
fact_sales_analysis	budget_quantity	budget.txt	COL3
fact_sales_analysis	budget_sales_value	budget.txt	COL4
fact_sales_analysis	forecast_quantity	forecast.txt	COL3
fact_sales_analysis	forecast_sales_value	forecast.txt	COL4

The results of this report are displayed as a list of source tables and columns, the order of the result set determining the immediate lineage. The report includes the following columns:

- Tables (*the name of a selected table*)
- Columns (*the name of a selected column*)
- Source Tables
- Source Columns

The result set is **NOT** sortable, as the order of the result set determines the immediate lineage.

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

COLUMN TRACK-FORWARD

This report lists the columns derived from the selected objects.

Objects Included

The following WhereScape RED object types are included in this report:

- Load Tables
- Stage Tables
- Dimension Tables and Views
- Fact Tables
- Aggregate Tables
- Exports
- Hubs, Satellites and Links

Parameters

- Groups
- Projects
- Table
- Column

Column Track Forward Report Parameters

Select a Group, Project or Table and Column to list the columns derived from the objects.

Groups:

Projects:

Table:

Column:

Exclude Intermediate Steps

OK Cancel

Results

If you left the **Exclude Intermediate Steps** checkbox **unchecked**, then the results screen will be as follows, showing the impacted tables for the selected tables:

Report: Column Track Forward

Tables (load_customer)	Columns	Impact Tables	Impact Columns
load_customer	code	dim_customer	code
load_customer	code	stage_customer	customer_code
load_customer	code	customer	customer_code
load_customer	name	dim_customer	name
load_customer	name	stage_customer	customer_name
load_customer	name	customer	name
load_customer	name	odim_customer	name
load_customer	address	dim_customer	address
load_customer	address	stage_customer	customer_address
load_customer	address	customer	address
load_customer	city	dim_customer	city
load_customer	city	stage_customer	customer_city
load_customer	city	customer	city
load_customer	city	odim_customer	city

Reports Results

If however, you **selected Exclude Intermediate Steps**, then the results screen will be as follows, showing only the final impacted table for the selected tables:

Report: Column Track Forward (Excluding intermediate steps)

Tables (load_customer)	Columns	Impact Tables	Impact Columns
load_customer	code	customer	customer_code
load_customer	name	odim_customer	name
load_customer	address	customer	address
load_customer	city	odim_customer	city
load_customer	state	odim_customer	state

Reports Results

The results of this report are displayed as a list of impacted tables and columns, the order of the result set determining the immediate impact. The report includes the following columns:

- Tables (*the names of selected tables*)
- Columns (*the names of selected columns*)
- Impact Tables
- Impact Columns

The result set is **NOT** sortable, as the order of the result set determines the immediate lineage.

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.


TABLE REPORTS

There are four reports for analyzing Tables:

- Tables Without Comments
- Load Tables by Connection
- Export Objects by Connection
- External Source Table/Files

TABLES WITHOUT COMMENTS

This report shows **table** objects in the metadata that don't have descriptions. The report prompts the user with the option to only include end user facing table objects or all table objects.



Objects Included

The following WhereScape RED object types can be included in this report:

- All Tables objects or
- Only End User Layer (EUL) objects

Note: The EUL object are set in *Tools>Options>Object Types>Object Type End User Setting* (see "*Object Type End User Setting*" on page 86).

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of objects in the metadata that are missing comments with the following columns:

- Table name
- Table type

Report Example

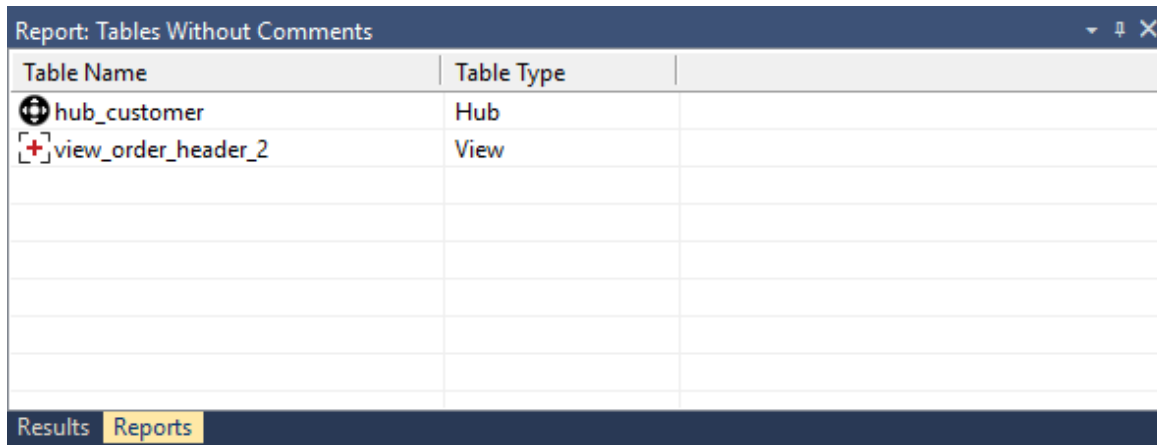


Table Name	Table Type	
hub_customer	Hub	
view_order_header_2	View	

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

LOAD TABLES BY CONNECTION

This report shows load tables in the metadata repository with their Connection and Source schema or database.

Objects Included

The following WhereScape RED object types are included in this report:

- Load Tables

Parameters

There are no parameters for this report.

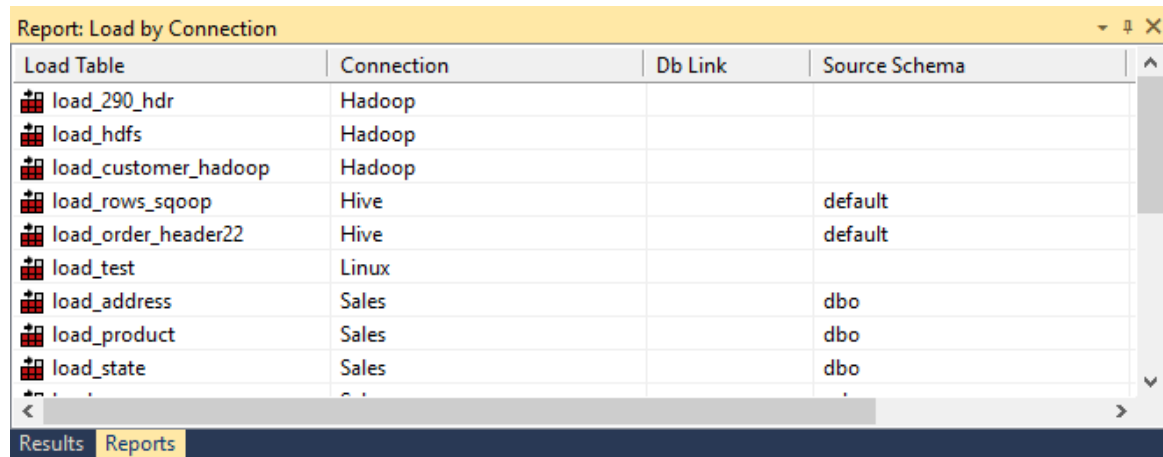
Results

The results of this report are displayed as a list of objects in the metadata repository with the following columns:

- Load table (*the name of the load table*)
- Connection
- Source schema (*the name the database or schema the load table is source from - blank for files*)

The result set is sortable by clicking on the appropriate column heading.

Report Example



Load Table	Connection	Db Link	Source Schema
load_290_hdr	Hadoop		
load_hdfs	Hadoop		
load_customer_hadoop	Hadoop		
load_rows_sqoop	Hive		default
load_order_header22	Hive		default
load_test	Linux		
load_address	Sales		dbo
load_product	Sales		dbo
load_state	Sales		dbo

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

EXPORT OBJECTS BY CONNECTION

This report shows export tables in the metadata repository with their Connection and Source schema or database.

Objects Included

The following WhereScape RED object types are included in this report:

- Export Tables

Parameters

There are no parameters for this report.

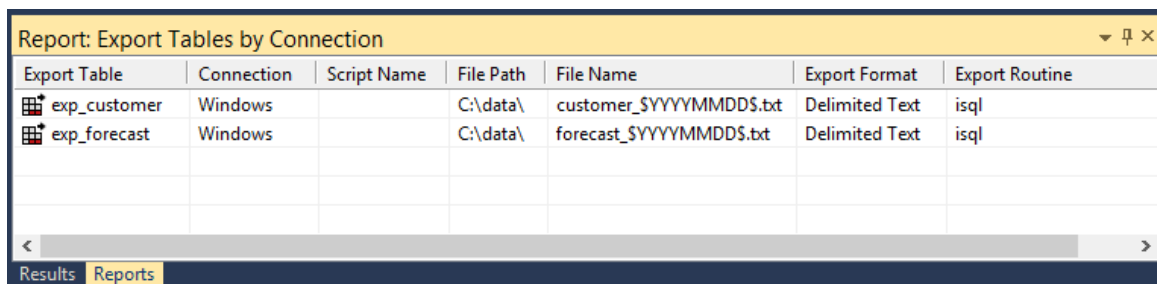
Results

The results of this report are displayed as a list of objects in the metadata repository with the following columns:

- Export table (*the name of the export table*)
- Connection
- Script Name
- File Path
- File Name
- Export Format
- Export Routine

The result set is sortable by clicking on the appropriate column heading.

Report Example



Export Table	Connection	Script Name	File Path	File Name	Export Format	Export Routine
exp_customer	Windows		C:\data\	customer_YYYYMMDD\$.txt	Delimited Text	isql
exp_forecast	Windows		C:\data\	forecast_YYYYMMDD\$.txt	Delimited Text	isql

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

RECORDS THAT FAILED A DIMENSION JOIN

This report shows the dimension business key(s) that could not be found in a specified dimension when a specified staging table was last updated. This report will show null values, blank values and business keys not found in the dimension.

Objects Included

The following WhereScape RED object types are included in this report:

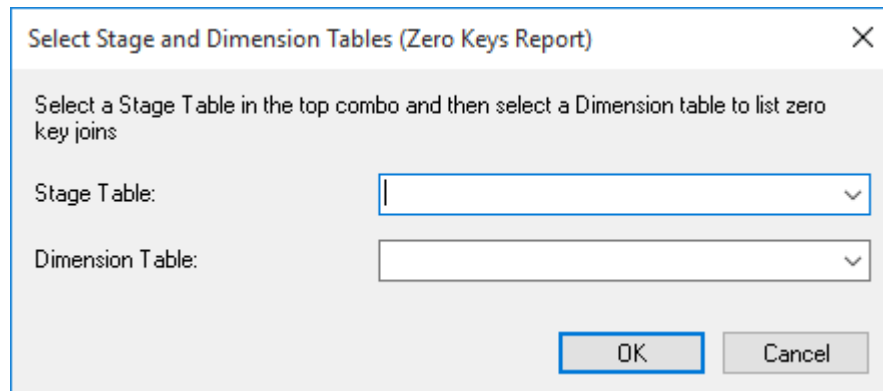
- Stage Tables

Parameters

This report requires two parameters to be specified:

- Stage Table Name (*the staging table to be checked*)
- Dimension Table Name (*the dimension table or dimension view of the dimension key in the selected staging table*)

Note: WhereScape RED version 8.1.1.0 and above no longer supports the Dimension View object type when creating new objects. However, users that have existing Dimension View objects in their data warehouse can retain and continue to use them.



Select Stage and Dimension Tables (Zero Keys Report) X

Select a Stage Table in the top combo and then select a Dimension table to list zero key joins

Stage Table: []

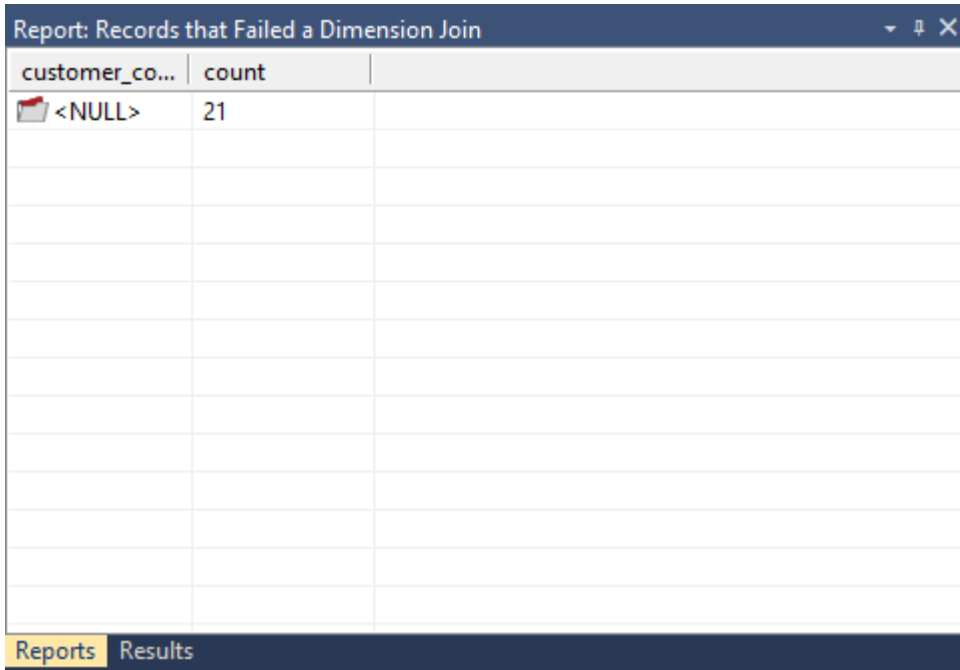
Dimension Table: []

OK Cancel

Results

The report contains a list of values not found in the dimension and a count for each value. Each column is sortable by clicking on the appropriate column heading.

Report Example



Report: Records that Failed a Dimension Join

customer_co...	count	
<NULL>	21	

Reports Results

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

EXTERNAL SOURCE TABLES/FILES

This report shows external sources for load tables in the metadata repository.

Objects Included

The following WhereScape RED object types are included in this report:

- Load Tables

Parameters

There are no parameters for this report.

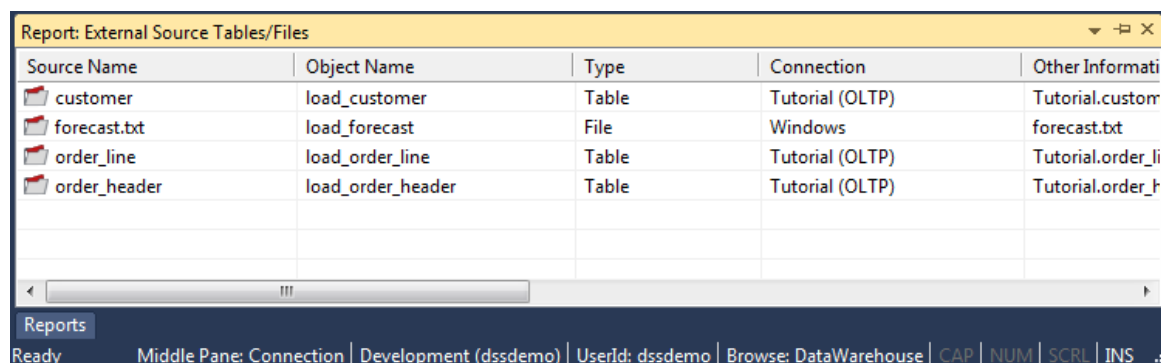
Results

The results of this report are displayed as a list of objects in the metadata repository with the following columns:

- Source name (*the name of the object's source*)
- Object name (*the name of the object*)
- Type (*the type of object the source is: Table or File*)
- Connection
- Other information (*for tables, the source schema/database.source table; for files, the file name*)

The result set is sortable by clicking on the appropriate column heading.

Report Example



Source Name	Object Name	Type	Connection	Other Information
customer	load_customer	Table	Tutorial (OLTP)	Tutorial.custom
forecast.txt	load_forecast	File	Windows	forecast.txt
order_line	load_order_line	Table	Tutorial (OLTP)	Tutorial.order_li
order_header	load_order_header	Table	Tutorial (OLTP)	Tutorial.order_h

Reports
Ready Middle Pane: Connection | Development (dssdemo) | UserId: dssdemo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

PROCEDURE REPORTS

There are two reports for analyzing Procedures:

- Modified Procedures
- Custom Procedures

MODIFIED PROCEDURES

This report shows **modified** procedures in the metadata repository with their creation and modification dates.

Objects Included

The following WhereScape RED object types are included in this report:

- Procedures

Parameters

There are no parameters for this report.

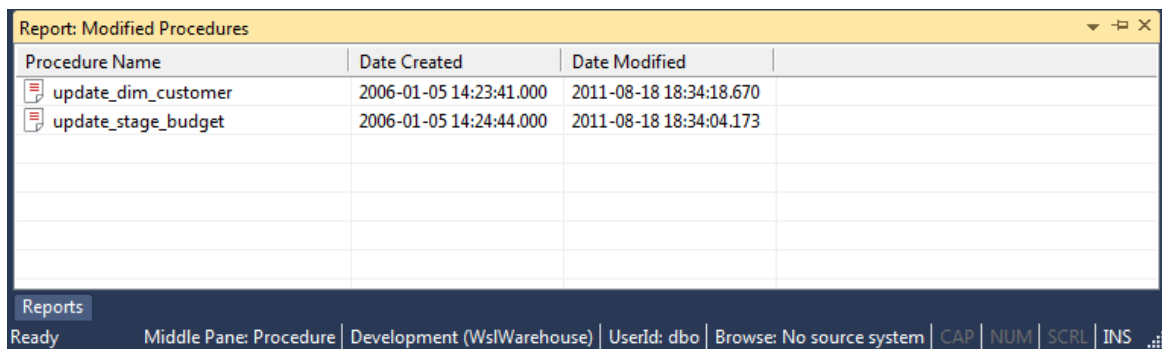
Results

The results of this report are displayed as a list of modified procedures in the metadata repository with the following columns:

- Name (*the name of the object*)
- Dated Created
- Date Modified (*last modification date*)

The result set is sortable by clicking on the appropriate column heading.

Report Example



Procedure Name	Date Created	Date Modified
update_dim_customer	2006-01-05 14:23:41.000	2011-08-18 18:34:18.670
update_stage_budget	2006-01-05 14:24:44.000	2011-08-18 18:34:04.173

Reports
Ready Middle Pane: Procedure | Development (WslWarehouse) | UserId: dbo | Browse: No source system | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

CUSTOM PROCEDURES

This report shows **custom** procedures in the metadata repository with their creation and modification dates.

Note: Custom procedures are procedures attached to any table object as a **Custom Procedure**.

Objects Included

The following WhereScape RED object types are included in this report:

- Procedures

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of custom procedures in the metadata repository with the following columns:

- Name (*the name of the object*)
- Table Name (*the table object the procedure is attached to*)
- Dated Created
- Date Modified (*last modification date*)

The result set is sortable by clicking on the appropriate column heading.

Report Example

Procedure Name	Table Name	Date Created	Date Modified
daily_date_roll	dim_date	2006-01-05 13:55:03.000	2006-01-05 13:55:03.000

Ready Middle Pane: Procedure | Development (WsIWarehouse) | UserId: dbo | Browse: No source system | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

OBJECT REPORTS

There is five reports for analyzing Objects:

- Objects-Project Matrix
- Modified Objects (excluding indexes)
- Objects Checked-out
- Loaded or Imported Objects
- Objects with Extended Properties

OBJECTS-PROJECTS MATRIX

This report lists all objects that exist in one or more projects (other than **All Objects**) and the project(s) they exist in.

Objects Included

All object types are included in this report.

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of objects in the metadata repository that are in one or more projects (other than All Objects) with the following columns as a grid:

- Objects (*the name of the object*)
- Project Name 1 (*heading is the name of the first project, value is a 1 to indicate the object is in this project, blank otherwise*)
- Project Name 2 (*heading is the name of the second project, value is a 1 to indicate the object is in this project, blank otherwise*)
- ...
- Project Name n (*heading is the name of the nth project, value is a 1 to indicate the object is in this project, blank otherwise*)

The result set is sortable by clicking on the appropriate column heading.

Report Example

Objects	Order Presenta...	Shared Dimens...	Included in Pr...	Object Type
Cubes			0	Connection
daily_date_roll		1	1	Procedure
daily_model...			0	Procedure
DataWareho...	1		1	Connection
dim_customer			0	Dimension
dim_custom...			0	Index
dim_custom...			0	Index
dim_custom...			0	Index
dim_date		1	1	Dimension
dim_date_id...		1	1	Index
dim_date_id...		1	1	Index
dim_product		1	1	Dimension

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

MODIFIED OBJECTS (EXCLUDING INDEXES)

This report lists objects in the metadata repository with their creation and modification dates.

Objects Included

The following WhereScape RED object types are included in this report:

- Load Tables
- Stage Tables
- Model Tables
- Aggregate Tables
- Join Indexes
- Views
- Exports
- Procedures
- Host Scripts

Parameters

There are no parameters for this report.

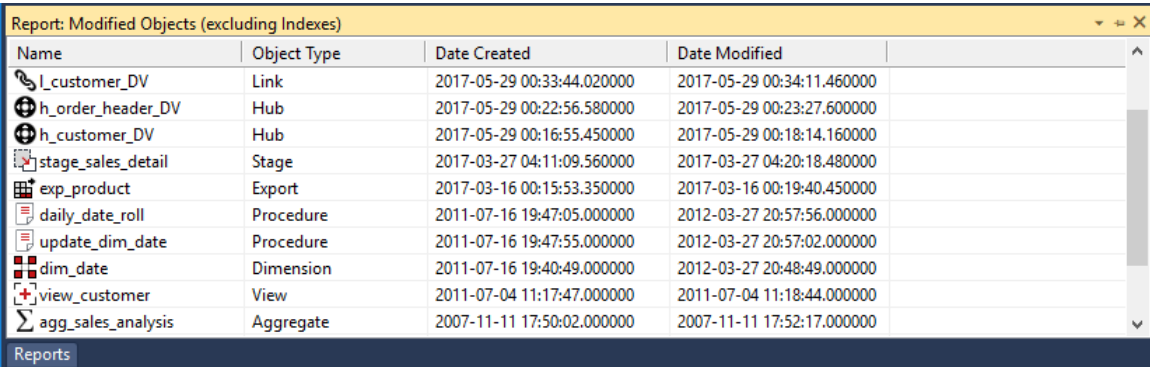
Results

The results of this report are displayed as a list of objects in the metadata repository with the following columns:

- Name (*the name of the object*)
- Object Type
- Dated Created
- Date Modified (*last modification date*)

The result set is sortable by clicking on the appropriate column heading.

Report Example



Name	Object Type	Date Created	Date Modified
l_customer_DV	Link	2017-05-29 00:33:44.020000	2017-05-29 00:34:11.460000
h_order_header_DV	Hub	2017-05-29 00:22:56.580000	2017-05-29 00:23:27.600000
h_customer_DV	Hub	2017-05-29 00:16:55.450000	2017-05-29 00:18:14.160000
stage_sales_detail	Stage	2017-03-27 04:11:09.560000	2017-03-27 04:20:18.480000
exp_product	Export	2017-03-16 00:15:53.350000	2017-03-16 00:19:40.450000
daily_date_roll	Procedure	2011-07-16 19:47:05.000000	2012-03-27 20:57:56.000000
update_dim_date	Procedure	2011-07-16 19:47:55.000000	2012-03-27 20:57:02.000000
dim_date	Dimension	2011-07-16 19:40:49.000000	2012-03-27 20:48:49.000000
view_customer	View	2011-07-04 11:17:47.000000	2011-07-04 11:18:44.000000
agg_sales_analysis	Aggregate	2007-11-11 17:50:02.000000	2007-11-11 17:52:17.000000

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

OBJECTS CHECKED-OUT

This report lists all objects currently checked out.

Objects Included

All object types and data warehouse tables can be included in this report.

Parameters

There are no parameters for this report.

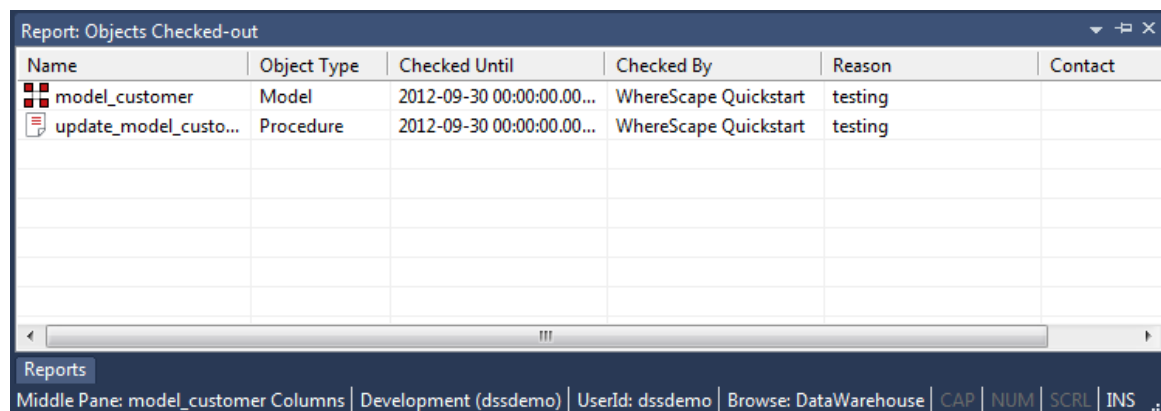
Results

The results of this report are displayed with the following columns:

- Name (*The name of the object*)
- Object Type (*The type of the object, e.g. Fact, Dimension, etc.*)
- Checked Until (*The date the object will be automatically checked back in*)
- Checked By (*The name of the WhereScape RED user who checked out the object*)
- Reason (*The reason provided for checking out the object*)
- Contact (*The contact details provided when the object was checked out*)

The result set is sortable by clicking on the appropriate column heading.

Report Example



Name	Object Type	Checked Until	Checked By	Reason	Contact
model_customer	Model	2012-09-30 00:00:00.00...	WhereScape Quickstart	testing	
update_model_custo...	Procedure	2012-09-30 00:00:00.00...	WhereScape Quickstart	testing	

Reports
Middle Pane: model_customer Columns | Development (dssdemo) | UserId: dssdemo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

LOADED OR IMPORTED OBJECTS

This report lists objects in the metadata repository that have been refreshed or imported from another repository.

Objects Included

The following WhereScape RED object types are included in this report:

- Load Tables
- Stage Tables
- Model Tables
- Aggregate Tables
- Join Indexes
- Views
- Cubes
- Exports
- Procedures
- Host Scripts

Notes:

- Indexes are not included.
 - Objects created via Application Load must have Auto Version enabled to be included in the report.
-

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of refreshed objects in the metadata repository with the following columns:

- Object Name
- Date (*of last refresh or import*)
- Description (*the kind of import or refresh*)
- Detail (*not currently used*)

The result set is sortable by clicking on the appropriate column heading.

Report Example

Object Name	Date	Description	Detail
odim_customer	2017-03-14 21:35:33.900000	Auto Version on Application Load by ...	
odim_ship_date	2017-03-14 21:35:35.480000	Auto Version on Application Load by ...	
ods_forecast	2017-03-14 21:35:09.480000	Auto Version on Application Load by ...	
sales_analysis	2017-03-14 21:35:46.580000	Auto Version on Application Load by ...	
stage_forecast	2017-03-14 21:35:10.400000	Auto Version on Application Load by ...	
stage_forecast	2017-06-09 00:19:42.600000	Auto Version on Application Load by ...	
stage_forecast	2017-09-18 04:39:30.890000	Auto Version on Application Load by ...	
stage_forecast	2017-10-02 00:08:48.250000	Auto Version on Application Load by ...	
stage_forecast	2017-10-02 05:50:30.240000	Auto Version on Application Load by ...	

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

OBJECTS WITH EXTENDED PROPERTIES

This report lists objects in the metadata repository with Extended Properties defined.

Objects Included

The following WhereScape RED object types are included in this report:

- Connections
- Load Tables
- Stage Tables
- Data Store Tables
- Hub Tables
- Link Tables
- Satellite Tables
- EDW 3NF Tables
- Dimension Tables and Views
- Fact Tables
- Aggregate Tables
- Join Indexes
- Views
- Exports

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of refreshed objects in the metadata repository with the following columns:

- Object Name
- Variable Name *(The unique variable name in the format it is accessed from the RED metadata service and included into scripts)*
- Display Name *(The extended property name displayed in the Connection or Table properties)*

The result set is sortable by clicking on the appropriate column heading.

Report Example

Object Name	Variable Name	Display Name
AWS RedShift	S3_ENCRYPTION	S3 Encryption
AWS RedShift	SECRET_KEY	Secret Key
load_budget	SECRET_KEY	Secret Key
AWS RedShift	SEND_FILES_ZIPPED	Send Files Zipped
load_budget	SEND_FILES_ZIPPED	Send Files Zipped
AWS RedShift	SPLIT_FILES_INTO	Split Files Count
load_budget	SPLIT_FILES_INTO	Split Files Count
AWS RedShift	SPLIT_FILES_THRESHOLD	Split Files Threshold
load_budget	SPLIT_FILES_THRESHOLD	Split Files Threshold

Middle Pane: Connection | Development (doc6871) | UserId: doc6871 | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

JOB REPORTS

There are three reports for analyzing Jobs:

- Object by Scheduler Job
- Jobs with an Object
- Tasks of a Job

OBJECT-JOB MATRIX

This report shows all jobs and objects, as well as the object actions. Table and Cube objects not in any jobs are displayed with a job name of **No Job**.

Objects Included

All object types are included in this report.

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of jobs with the following columns:

- Table Name (*the name of the table or cube*)
- Action
- Job Name (*the name of the job*)
- Job Status
- Job Last Run
- Job Next Run

Report Example

Report: Objects by Scheduler Job

Object Name	Action	Job Name	Job Status	Job La:
daily_model_date_roll		No Job		
update_model_customer		No Job		
update_model_date		No Job		
update_model_forecast		No Job		
update_model_order_header		No Job		
update_model_order_line		No Job		
update_ods_forecast		No Job		
update_stage_customer		No Job		
update_stage_forecast		No Job		
update_stage_order_header		No Job		
update_stage_order_line		No Job		
test		No Job		
update_dim_time		No Job		
model_customer		No Job		
model_date		No Job		
model_forecast	Statistics	Daily Update	On Hold	2011-0
model_forecast	Statistics	Daily Update_1	On Hold	2011-0
model_forecast	Process	Daily Update_1	On Hold	2011-0
model_forecast	Process	Daily Update	On Hold	2011-0
model_order_header		No Job		
model_order_line		No Job		
ods_forecast		No Job		

Reports
Middle Pane: model_customer Columns | Development (dssdemo) | UserId: dssdemo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

JOBS WITH AN OBJECT

This report shows all jobs a specified object appears in and its action.

Objects Included

All object types are included in this report.

Parameters

This report has one parameter:

- Object Name

Results

The results of this report are displayed as a list of jobs with the following columns:

- Jobs including **object_name**
- Action

Report Example

Jobs including load_customer	Action
Daily Update	Process
Daily Update_1	Process

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

TASKS OF A JOB

This report shows all tasks for a selected job including dependencies.

Objects Included

All object types are included in this report.

Parameters

This report has one parameter:

- Job Name

Results

The results of this report are displayed as a list of task dependencies with the following columns:

- Task name (*the table, Index, procedure or script name*)
- Action
- Order (*the order number as shown in the edit tasks dialog in the scheduler*)
- Depends On (*the task(s) and order number this task depends on*)

Report Example

Task Name	Action	Order	Depends On
load_customer	Process	10.10.1	
load_forecast	Process	10.10.2	
stage_forecast	Process	12.10.4	load_forecast (10.10.2)
model_forecast	Process	13.10.5	stage_forecast (12.10.4)
model_forecast	Stats	14.10.6	model_forecast (13.10.5)

Report: Tasks of a Job

Middle Pane: model_customer Columns | Development (dssdemo) | UserId: dssdemo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

OPERATIONAL REPORTS

There are three Operational Reports:

- Object Performance History
- Job Performance History
- Task Performance History

OBJECT PERFORMANCE HISTORY

This report shows the audit trail for a selected object from the scheduler logs.

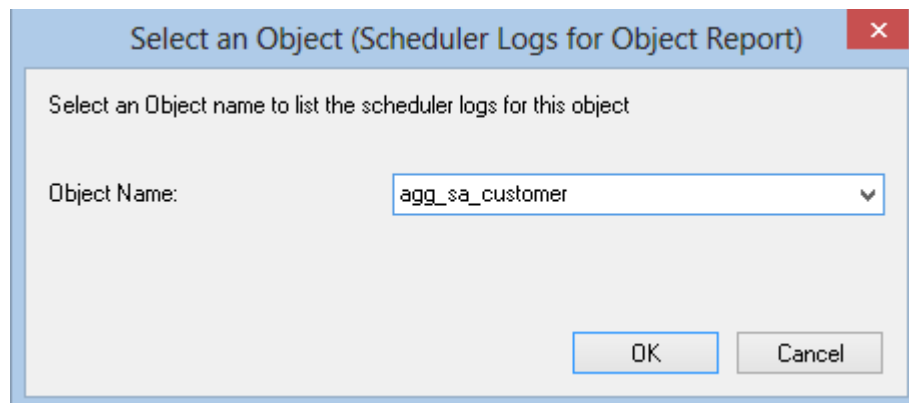
Objects Included

All object types are included in this report.

Parameters

This report has one parameter:

- Object Name



Select an Object (Scheduler Logs for Object Report)

Select an Object name to list the scheduler logs for this object

Object Name:

OK Cancel

Results

The results of this report are displayed as a list of audit log entries with the following columns:

- Sta (*the type of audit log entry*)
- Time (*the date and time the audit log entry was written*)
- Seq (*the job sequence of the job writing the audit log entry*)
- Message (*the message in the Audit log*)
- Task (*the object name*)
- Job (*the name of the job that ran the task*)

Report Example

Report: Object Performance History

Sta	Time	Seq	Message	Task	Database ...	Db Code
F	2011-05-19 10:33:4...	1500...	Unhandled Exception in model_order_header. Step 1...	update_m...	Column/P...	381
F	2011-07-01 13:53:2...	2000...	Unhandled Exception in model_order_header. Step 1...	update_m...	Column/P...	381
F	2011-09-27 15:43:2...	2700...	Unhandled Exception in model_order_header. Step 1...	update_m...	Column/P...	381
F	2011-09-30 13:46:1...	2800...	Unhandled Exception in model_order_header. Step 1...	update_m...	Column/P...	381
F	2011-09-30 13:46:5...	2800...	Unhandled Exception in model_order_header. Step 2...	update_m...	Column/P...	381

Reports

Middle Pane: model_customer Columns | Development (dssdemo) | UserId: dssdemo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

JOB PERFORMANCE HISTORY

This report shows the performance (duration) of a specified job over time.

Objects Included

All object types are included in this report.

Parameters

This report has one parameter:

- Job Name

Select Job (Job Performance Report)

Select a Job to display scheduler performance for a Job.

Job:

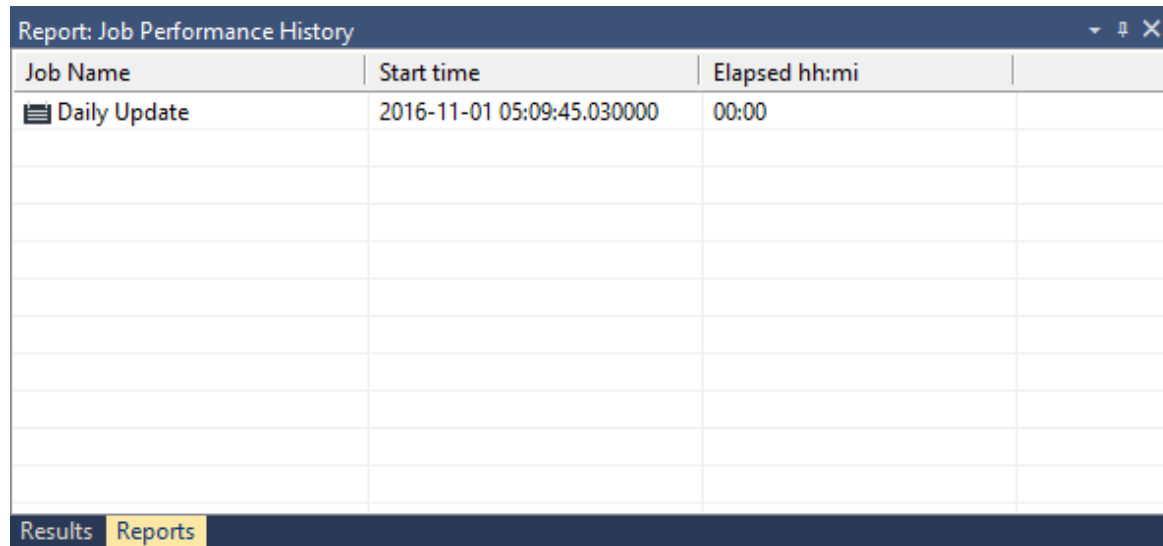
OK Cancel

Results

The results of this report are displayed as a list of job instances with the following columns:

- Job name (*the name of the job*)
- Start time (*the date and time the job started*)
- Elapsed hh:mi (*the duration of the job*)

Report Example



Job Name	Start time	Elapsed hh:mi
☰ Daily Update	2016-11-01 05:09:45.030000	00:00

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

TASK PERFORMANCE HISTORY

This report shows the performance (duration) of a specified task within a specified job over time.

Objects Included

All object types are included in this report.

Parameters

This report has two parameters:

- Job Name
- Task Name (including action)

Select Job and Task (Task Performance Report)

Select a Job and Task to display scheduler performance for a Task.

Job:

Task:

OK Cancel

Results

The results of this report are displayed as a list of task instances for the selected job with the following columns:

- Task name (*the table, Index, procedure or script name*)
- Action
- Start time (*the date and time the task started*)
- Elapsed hh:mi (*the duration of the task*)

Report Example

Task Name	Action	Start time	Elapsed hh:mi
load_customer	Process	2016-11-01 05:09:45.330000	00:00

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

CHAPTER 30

VALIDATE

When these Validate processes are run, the results are displayed in the middle pane of the RED screen; the results of the other reports are displayed in a separate tab in the bottom pane of the RED screen.

IN THIS CHAPTER

Validate Meta-data	826
Validate Workflow Data	826
Validate Table Create Status	826
Validate Load Table Status	827
Validate Procedure Status	827
List Meta-data Tables not in the Database	827
List Database Tables not in the Meta-data	828
List Tables with no related Procedures or Scripts	829
List Procedures not related to a Table	831
Teradata: View of Model Validate	832
Query Data Warehouse Objects	833

VALIDATE META-DATA

This process validates the Meta data. If problems are encountered, the results are displayed in the middle pane.

Use the right-click option against each identified issue to apply a repair.

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

VALIDATE WORKFLOW DATA

This process validates the Workflow data. If problems are encountered, the results are displayed in the middle pane.

Use the right-click option against each identified issue to apply a repair.

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel

VALIDATE TABLE CREATE STATUS

This process validates a table structure in the meta data against the table in the database.

- Select one or more tables and click the **Validate Selected** button or click the **Validate All** button to validate all the tables.
- 1 If a table is found to be different then it can be altered by using the right-click menu option when positioned over the table name.
 - 2 If the update date and the modified in database date imply a change that is not apparent then these dates can be re-synced in the same way.

Sync Column order with database

Right click on the result set and select **Sync Column order with database** to reorder the metadata columns to match the column order in the database table.

Sending Results to Microsoft Excel

Right click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

VALIDATE LOAD TABLE STATUS

This process compares a load table in the meta data with the table in the source system. It compares the columns and column data types.

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

VALIDATE PROCEDURE STATUS

This process compares a procedure in the meta data with the compiled version of the same procedure stored within the database. The subsequent display will report either a match or a difference.

If a procedure is found to differ then you can use the procedure editor to examine the exact differences by selecting the **Tools>Compare to Compiled Source** option.

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

LIST META-DATA TABLES NOT IN THE DATABASE

This report shows database **table** objects in the metadata that don't exist in the data warehouse database.

Objects Included

The following WhereScape RED object types are included in this report:

- Load Tables
- Stage Tables
- Model Tables
- Aggregate Tables
- Join Indexes
- Views
- Retro Copies (but not Retro Definitions)

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of objects in the metadata not in the data warehouse database.

Report Example

Table Name
dbo.stage_customer

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

LIST DATABASE TABLES NOT IN THE META-DATA

This report shows database **table** objects that exists in the Teradata database but not in the metadata.

Objects Included

The following WhereScape RED object types are included in this report:

- Load Tables
- Stage Tables
- Model Tables
- Aggregate Tables
- Join Indexes
- Views
- Retro Copies (but not Retro Definitions)

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of objects in the Teradata database not in the metadata.

Report Example

Report: Tables in Database but not in Metadata	
Table Name	
load_promotions	

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

LIST TABLES WITH NO RELATED PROCEDURES OR SCRIPTS

This report shows all table objects (certain types of objects only - see below) in the metadata repository that don't have an associated update procedure.

Objects Included

The following WhereScape RED object types are included in this report:

- Load Tables (*script based loads only*)
- Stage Tables
- Model Tables
- Aggregate Tables

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of table objects in the metadata repository with the following columns:

- Table name

The result set is sortable by clicking on the appropriate column heading.

Report Example

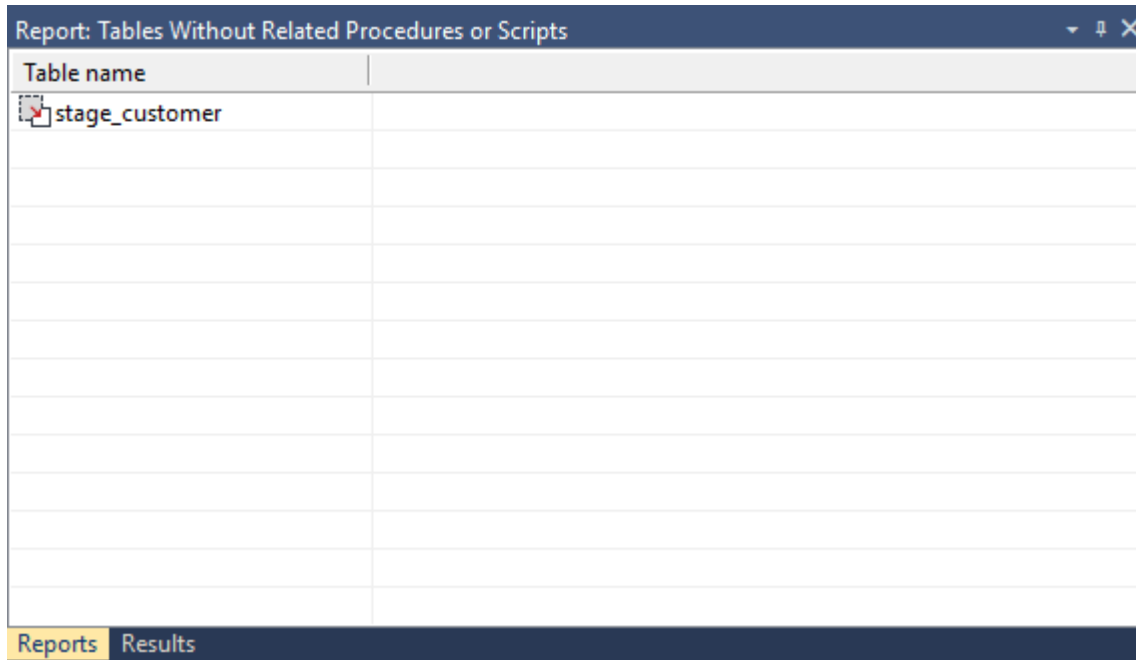


Table name
stage_customer

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

LIST PROCEDURES NOT RELATED TO A TABLE

This report shows all procedures and host scripts in the metadata repository that aren't associated with a table object.

Objects Included

The following WhereScape RED object types are included in this report:

- Procedures
- Host Scripts

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of code objects in the metadata repository with the following columns:

- Procedure/Script name

The result set is sortable by clicking on the appropriate column heading.

Report Example



The screenshot shows a report window with a dark blue header and a white table. The window title is 'Report: Procedures Not Related to a Table'. The table has two columns and two rows of data. The first row has a small icon in the first column and the text 'daily_date_roll'. The second row has a small icon in the first column and the text 'get_dim_date_key'. The table is followed by several empty rows. At the bottom of the window, there is a navigation bar with 'Reports' and 'Results' tabs.

Procedure/Script name	
daily_date_roll	
get_dim_date_key	

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

TERADATA: VIEW OF MODEL VALIDATE

This report shows any views built on model tables that do not have the same columns or column properties.

Objects Included

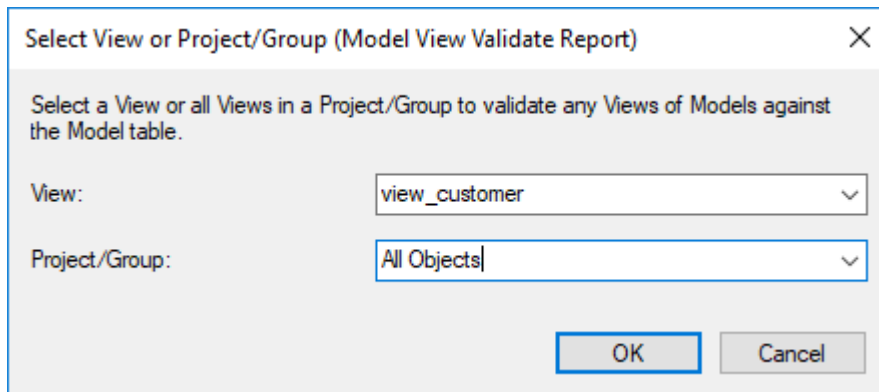
The following WhereScape RED object types are included in this report:

- Views

Parameters

This report has two alternate parameters:

- View Name **OR**
- Project/Group/All Objects



Select View or Project/Group (Model View Validate Report) X

Select a View or all Views in a Project/Group to validate any Views of Models against the Model table.

View: view_customer

Project/Group: All Objects

OK Cancel

Results

The results of this report are displayed as a list of differing view columns with the following columns:

- View Name
- Model Name (*the name of the model table the view is based on*)
- Column Name (*the column where a difference exists*)
- Validates (*further details of the difference*)

The result set is sortable by clicking on the appropriate column heading.

Report Example

View Name	Model Table Name	Column Name	Validates
view_customer	model_customer	model_customer_key	eul, does not validate

Reports
Middle Pane: model_customer2 Columns | Development (dssdemo) | UserId: dssdemo | Browse: DataWarehouse | CAP | NUM | SCRL

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

QUERY DATA WAREHOUSE OBJECTS

This report enables SQL queries to be run as the user signed into the repository.

Objects Included

All object types and data warehouse tables can be included in this report.

Parameters

This report has one parameter:

- The SQL Query

Data Warehouse Object Query

Enter a SQL query (or multiple queries separated by a semi-colon) below. Press the 'Run Query' button to execute. Press the 'Close' button to close this window when all queries have been executed.

```
SELECT DISTINCT customer_code
FROM load_order_header
ORDER BY 1
```

Max Rows to Return: 1000

Run Query Close

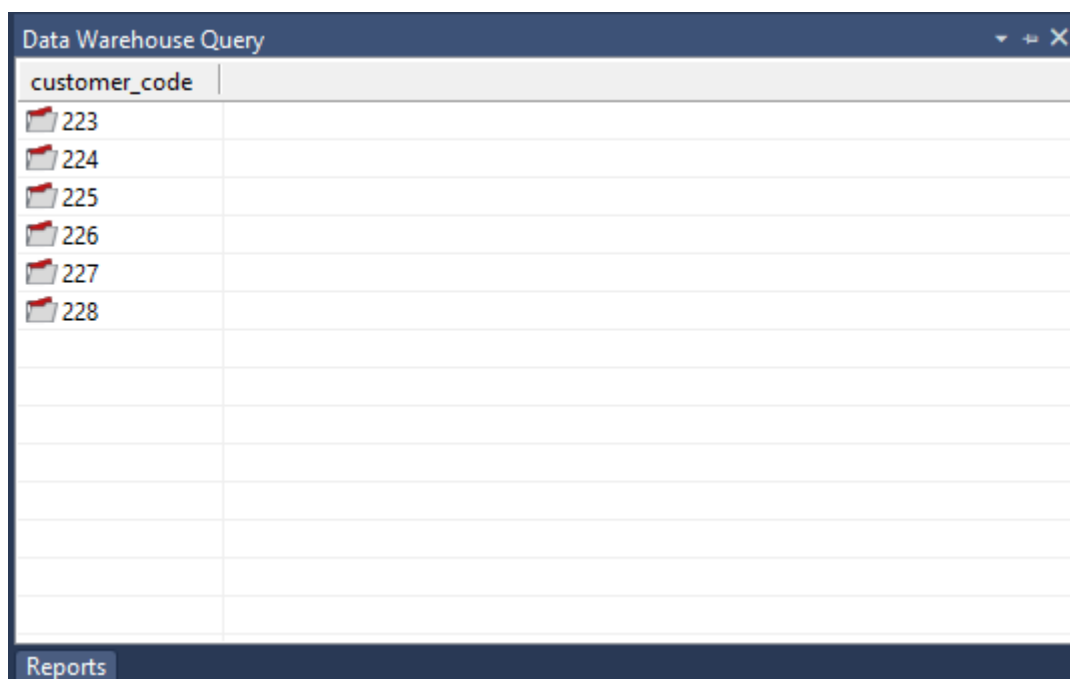
Results

The results of this report are displayed with the following columns:

- First SQL column SELECTed
- Second SQL column SELECTed
- ...
- nth SQL column SELECTed

The result set is sortable by clicking on the appropriate column heading.

Report Example



The screenshot shows a window titled "Data Warehouse Query" with a table. The table has one column labeled "customer_code". The rows contain the values 223, 224, 225, 226, 227, and 228. Each row has a small folder icon to the left of the value. The window has a dark blue header and footer. The footer contains a button labeled "Reports".

customer_code
223
224
225
226
227
228

Sending Results to Microsoft Excel

Right-click on the result set and click **Output to File** to send the results of this report to Microsoft Excel.

CHAPTER 31

PROMOTING BETWEEN ENVIRONMENTS

This chapter covers the promotion of metadata objects between environments. Various methods exist for getting new or changed metadata from one repository into another repository.

It is of course possible and in fact desirable to have multiple metadata repositories. At the very least we would normally have a development and a production environment.

In some situations it may also be desirable to have multiple child development repositories with one master repository where all elements are brought together. WhereScape RED supports this type of structure but does not include source control or co-ordination of the repositories. It is up to the data warehouse manager to manually ensure that the various objects are kept in sync and coordinated.

As with any software system there are issues around how code is moved from a development environment into a testing or production environment.

This promotion of objects can be achieved via a number of different methods. Each is discussed below. In summary they are:

- 1 Updating a repository with an application or application patch.
- 2 Importing objects from another repository.
- 3 Restoring a full metadata set into a repository

IN THIS CHAPTER

Applications	836
Importing Object Metadata	843
Importing Language Files	845
Data Warehouse Testing	846

APPLICATIONS

The definition of an application is discussed in the following section on applications and the loading and updating of applications is discussed at some length in the **Installation and Administration Guide**. Only the concepts of the use of applications will be covered here.

An application is defined for our purposes as a group of objects. An application is a method of loading objects into a metadata repository. It can be used to upgrade or provide patches to an existing metadata repository. As such an application can be used to distribute and remotely maintain a specific data warehousing solution.

An application consists of a series of Windows' files, which can be distributed to remote sites.

A list of the applications that have been applied to the metadata repository can be acquired via the **Tools>List Loaded Deployment Applications** menu option.

An application is created through the **Tools>Build Deployment Application** menu option. This application can then update a metadata repository through the Setup>Administration utility. In this manner the application model can be used to update a metadata repository in an ordered and controlled fashion. Loading an application inserts various objects into the chosen metadata repository. An application is best defined as a set of objects that are shipped to allow inclusion of those objects in a remote repository.

Note: An application can only be loaded into a metadata repository running on the same database type as that of the application creator (e.g. a Teradata application can only be loaded into a Teradata metadata repository, etc.).

APPLICATION CREATION

Creating an Application

An application is created by selecting the **Tools>Build Deployment Application** menu option. The following dialog box is displayed. Once the application is defined and the objects selected, the application files are generated when the **OK** button is clicked.

If procedures are compiled as part of the subsequent application load, the compiles occur in the order they are listed in the application. This way if there are procedure dependencies, ensure their ordering in the application object list is correct.

There are three tabs in the **Build Deployment Application** screen.

The first tab defines the application, the second lists the objects to add or replace in the destination repository and the third tab lists the objects to delete in the destination repository.

Define an Application distribution

The screenshot shows the 'Build Deployment Application' dialog box. It has a title bar with a close button (X) and a left-hand navigation pane with three tabs: 'Application' (selected), 'Objects to Add/Replace', and 'Objects to Delete'. The main area contains the following fields and controls:

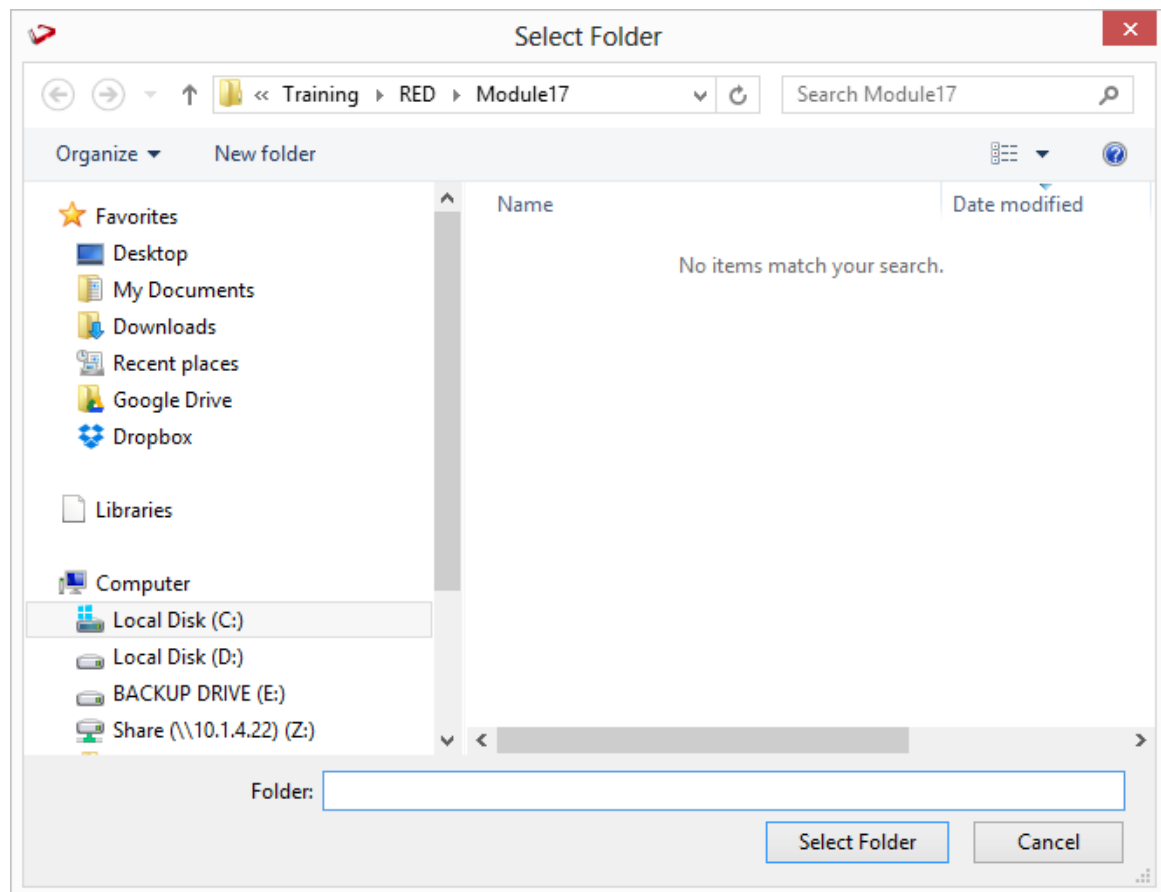
- Output Directory:** A text box containing 'C:\Training\RED\Module17' and a 'Browse...' button.
- Application Identifier:** A text box containing 'Orders'.
- Application Version:** A text box containing '1'.
- Application Name:** A text box containing 'Sales Orders Project Release'.
- Description:** A text area containing 'First Release of Orders into Production'.
- Previous Application:** A text box and a 'Browse...' button, with the instruction: 'You can select or enter a previous application file as a starting point for this application.'
- Pre-Application Load SQL:** A text area with the instruction: '(the following optional SQL statement will be issued before the application load commences):'
- Post-Application Load SQL:** A text area with the instruction: '(the following optional SQL statement will be issued after the application load completes):'

At the bottom right, there are three buttons: 'OK', 'Cancel', and 'Help'.

Output Directory

The directory to which the application files will be written. By default this will be the WhereScape program directory.

- To browse for the required folder, click on the **Browse...** button.
- The **Make New Folder** button allows you to create a new folder in the currently selected directory.



Application Identifier

The application identifier is a four character code used to uniquely identify the application. This identifier is used in the naming of the files that are created to contain the application data.

Application Version

The version is a character string that provides a version number for reference purposes. This version number is displayed when applications are being loaded, and is used in the naming of the files that are created to contain the application data. As such it must contain characters that are valid in a Windows file name.

Application Name

The name by which the application is known. This name is displayed during the choosing of an application to load and is recorded in the metadata of the repository into which an application is loaded. It is not used apart from documentation purposes.

Description

This description is displayed during the choosing of an application to load. It is not used at any other point apart from documentation purposes.

Application Files

When an application is created the following files are built, where XXXX is the application identifier and NNNNN is the application version.

File	Purpose
App_data_XXXX_NNNNN.wst	This file contains the scripts and data required to rebuild the objects in the new metadata repository.
App_id_XXXX_NNNNN.wst	This control file identifies the application and its version.
App_obj_XXXX_NNNNN.wst	This file contains control information and each object in the application.
App_con_XXXX_NNNNN.wst	A list of all the connections either in the application or used by objects in the application.
App_map_XXXX_NNNNN.wst	A list of all the project and group mappings for the objects in the application.

Previous Application

Click on the **Browse** button next to **Previous application** to choose a previously built application to use as a list of objects to include in the new application. After using a previous application as a starting point for this application, additional objects can be added or removed from the application.

Pre Application Load SQL

This box allows the entry of a SQL Statement that will be executed before the application is loaded. For example, we may wish to drop the date dimension before loading the application because we have changed the primary key constraint. In such a case, we would enter 'drop table dim_date' in this field to have the table dropped before the application is loaded.

Post Application Load SQL

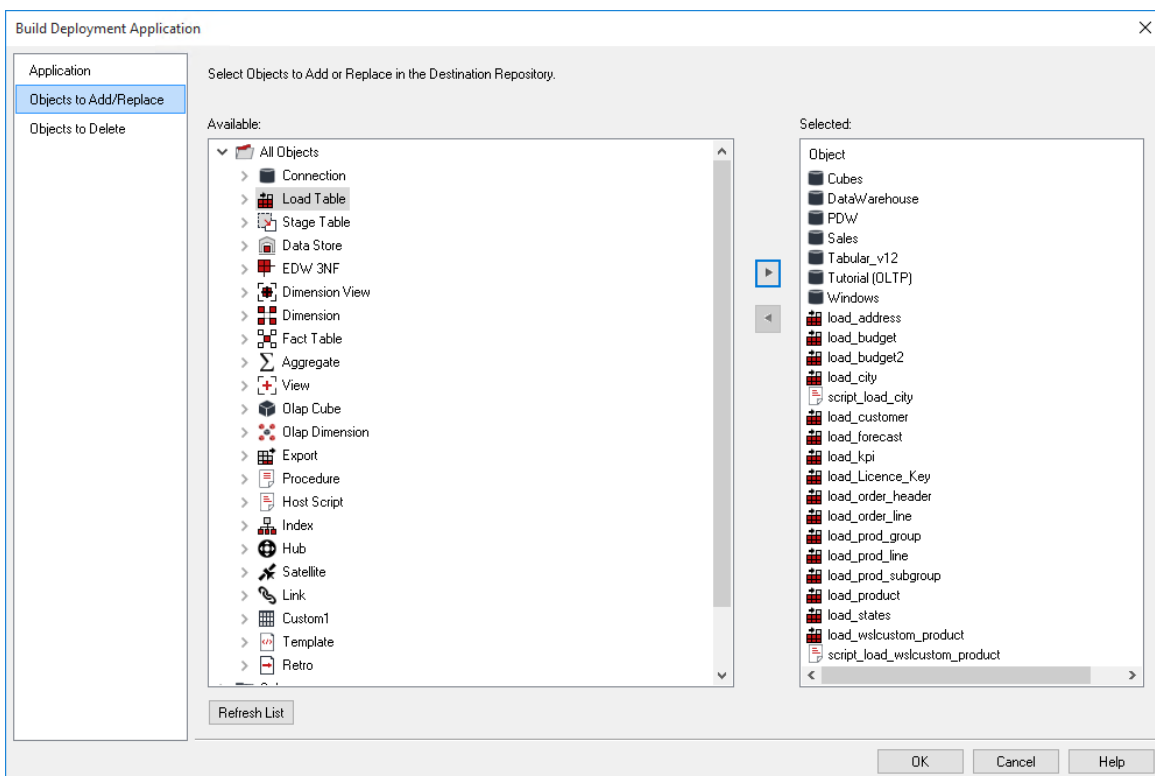
This box allows the entry of a SQL Statement that will be executed after the application is loaded. For example, we could execute a function to populate a table.

Objects to Add/Replace

Objects can be moved from the left object tree by double clicking on an object name or by using the > button. This tab allows you to select the objects to add or replace in the destination repository.

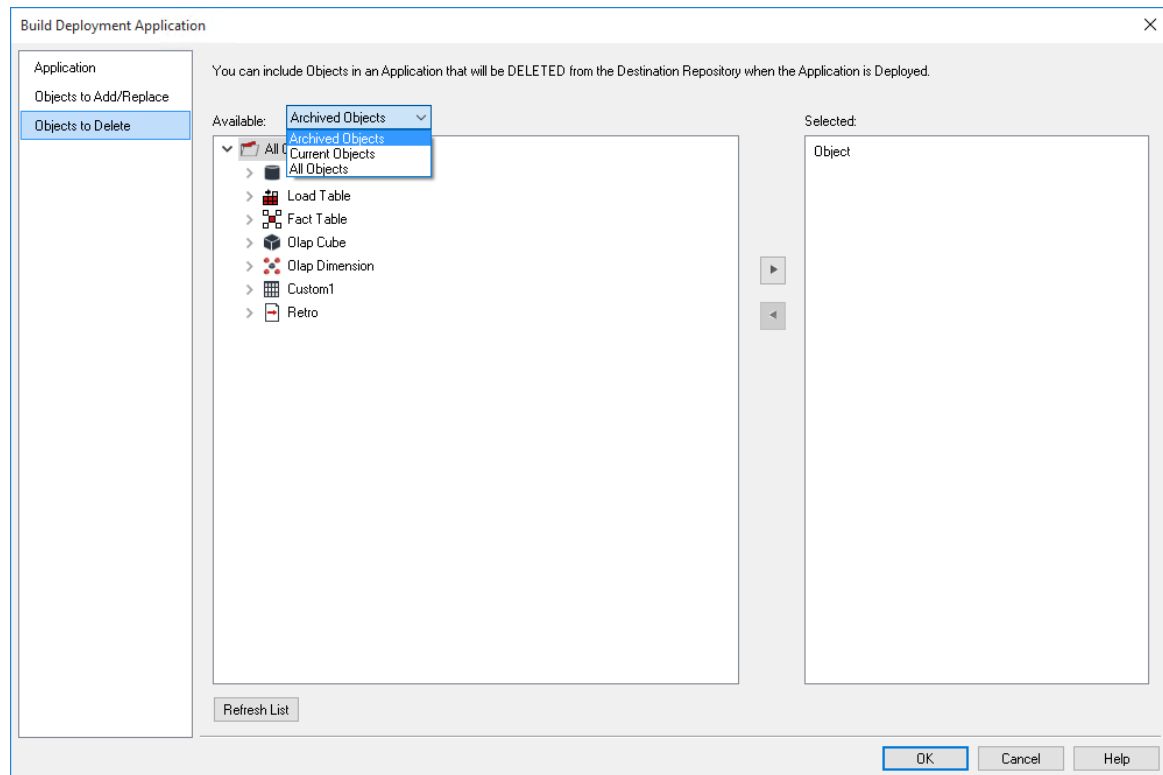
NOTE: Maximum number of objects in an application

- 5000 objects (including jobs)
 - 2000 source views of views
 - 1000 jobs
-



Objects to Delete

Objects can be moved from the left object tree by double clicking on an object name or by using the > button. This tab allows you to select the objects to delete in the destination repository.



NOTE: To set the objects available for selection, choose from the **Available** drop-down list. The options are Archived Objects, Current Objects and All Objects. The default is Archived Objects.

APPLICATION LOADING

Note: Applications can only be loaded into the same relational database type from which they were created. (e.g. a Teradata application can only be loaded into a Teradata database).

Applications are loaded via the Setup Administrator utility. The normal process for implementing an application would be as follows:

- 1 Run the **Setup Administrator** utility.
- 2 Change the application directory to the application's location.
- 3 Turn on logging in the Setup Administrator utility using Tools/Start logging.
- 4 Load the application via the Setup Administrator utility.
- 5 Choose the level of metadata application. There are several levels, from load metadata only through to load metadata and apply changes to all tables.
- 6 Resolve any connections and tablespaces to those present in the target environment.
- 7 Create/Re-create/Alter database tables, if selected in (5).
- 8 Compile database procedures, if selected in (5).
- 9 Turn off logging.
- 10 Review the output in the Setup Administrator utility.
- 11 Review the log file.

Note: Some database operations, such as converting an existing non-partitioned table to a partitioned table, cannot be done using a deployment application. In these cases, some manual intervention may be required to update the target databases to match the new metadata.



Refer to the Setup Administrator manual for more information about loading an application.

CREATING AND LOADING APPLICATIONS FROM THE COMMAND LINE

It is possible to create and load applications from the command line by running a bat file. For more detailed instructions, please see section **12.2 Creating and Loading Applications from the Command Line** in the **RED Installation Guide**.

IMPORTING OBJECT METADATA

Any group of objects can be imported into the current metadata repository from another repository. If an object already exists in the target repository then it is either skipped or replaced depending on the type of import undertaken. If an object is to be replaced as part of an import, a version of the object is created prior to its replacement.

To import an object or group of objects select the **Tools>Import Metadata Objects** menu option. A dialog as below will appear. The two options are IMPORT or REFRESH. An import will not replace an existing object of the same name. A refresh will version and replace any existing object of the same name.

Import Objects from Another Meta Repository [X]

You have chosen to import objects from another MetaRepository.
Two methods are supported, and described below:

IMPORT, imports objects without overwriting any existing objects.
You are given the option of renaming the new objects if in conflict.

REFRESH, imports objects replacing any existing object of the same name.
A version of the replaced object is created.

Enter the uppercase word **IMPORT** or **REFRESH** in the first window,
a username and password with access rights to the source repository
and the source MetaRepository to import from.

Odbc Connect

User Name

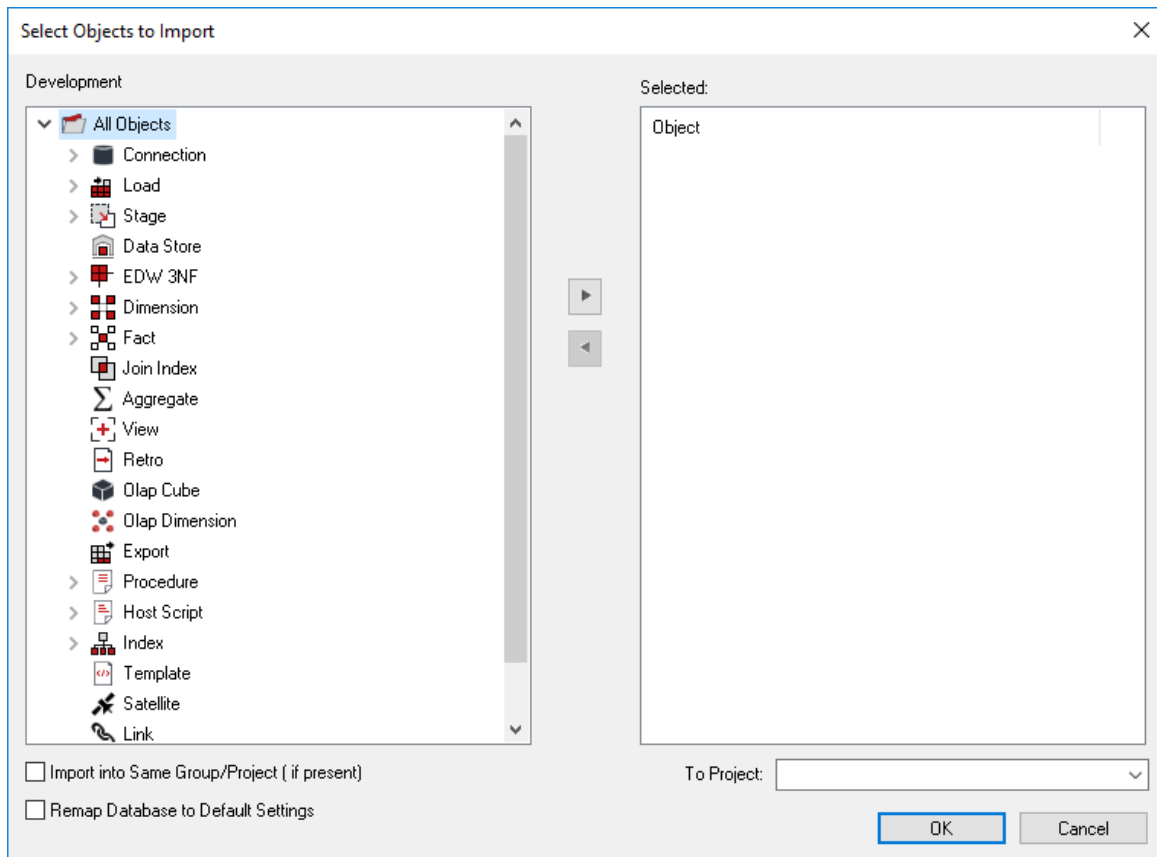
Password

Meta Repository

Enter the connection, and a database user name and password that has access to the source metadata repository. Finally, enter the user name of the metadata repository you want to import from. In most situations the 'user name' and 'meta repository' would be the same. However, if you only have read access to a meta repository then it may be necessary to login to the database under a different user name from that of the repository you are trying to import from.

You are not permitted to select the current meta repository in the **Meta Repository** field, but are permitted to login using the existing repository username. Once a successful logon is completed the contents of the source repository are loaded and the following dialog appears.

Select an object by double-clicking on it or by selecting and using the > button. If an object such as a table is chosen then any related scripts, procedures and indexes are also selected. They can be removed if required. A target project can be selected.



Once all required objects are selected the import will commence when the **OK** button is clicked.

On completion, a dialog box appears, notifying the number of each type of object imported, and skipped.

Note: The repository from which you are importing should be the same metadata version as the target repository.

IMPORTING LANGUAGE FILES

Note: Applications can only be loaded into the same relational database type from which they were created. (For example, a Teradata application can only be loaded into a Teradata database).

Language Files are loaded via the Setup Administrator utility. The normal process for implementing a Language file would be as follows:

- 1 Run the **Setup Administrator** utility.
- 2 Go to the Languages menu item in the top command bar and select **Load Languages**.
- 3 Right-click on the Language file to be loaded and select **Install Language**.
- 4 Select the ODBC data source and Log on to the target meta repository.
- 5 Select the language to be updated.
- 6 Review the output in the Setup Administrator utility.

Refer to the **Setup Administrator** manual for more information about loading a Language File.

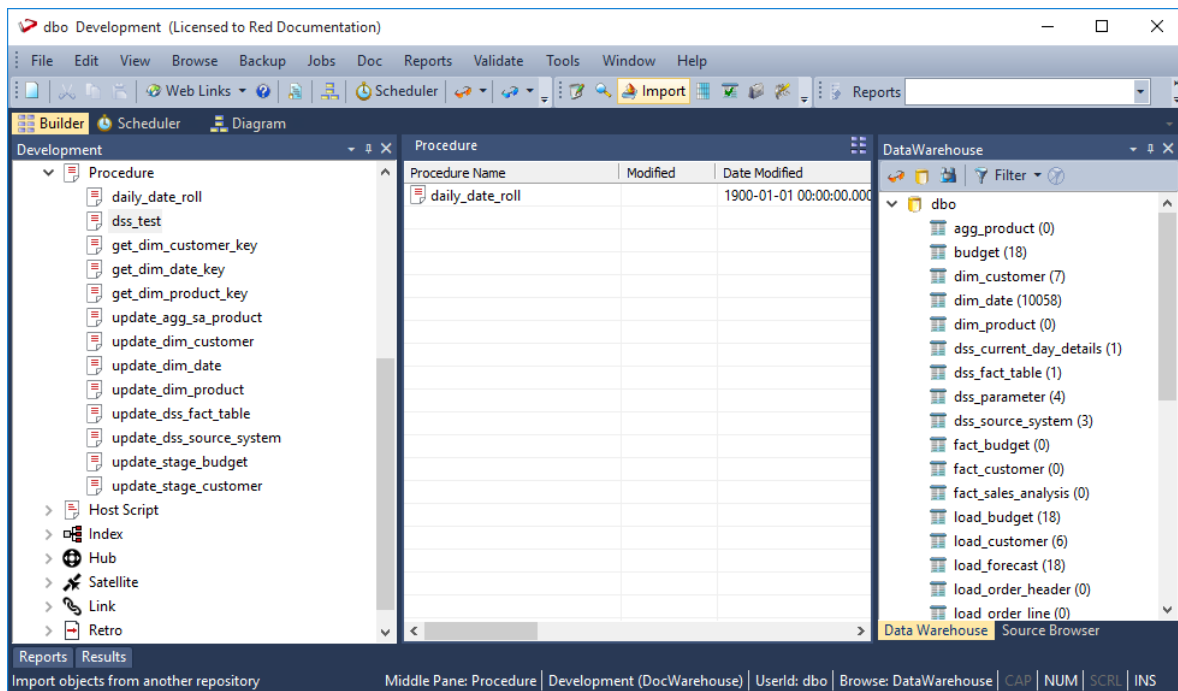
DATA WAREHOUSE TESTING

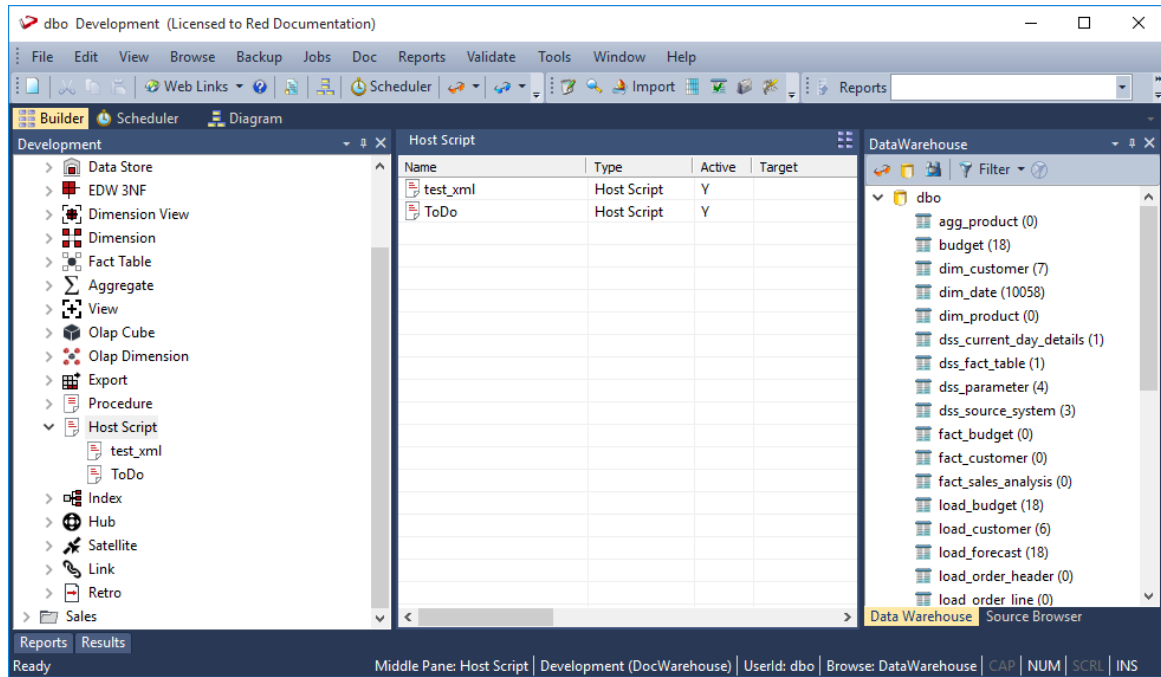
Note: Applications can only be loaded into the same relational database type from which they were created. (For example, a Teradata application can only be loaded into a Teradata database).

Testing applications are loaded via the Setup Administrator utility. Refer to the Setup Administrator manual for more information on how to load an application.

A testing application set consists of a Procedure and an XML script and provides the ability to define a series of tests against data warehouse objects; either comparing them to an expected value or to the results of a query.

Once the application set has been loaded, the Procedure and the XML script will be visible in the left pane.





The XML script contains the test definitions. Each test is a new XML node in the comparison query. The procedure simply runs the test and determines whether the tests are passed or not. This is most likely to be run as a scheduled job within WhereScape RED.

To create a job

- 1 Click the **Scheduler** Button.
- 2 Choose **File** and then **New Job**.
- 3 Enter the definition of the job.

Job Definition
✕

Job Name:

Description:

This job runs a procedure to perform a series of tests on the Data Warehouse.

Frequency:

Start Date:

Start Time:

Maximum Threads:

Scheduler:

Dependent On:

Parent job	Fail	Look back (minutes)	Maximum wait (minutes)	
------------	------	---------------------	------------------------	--

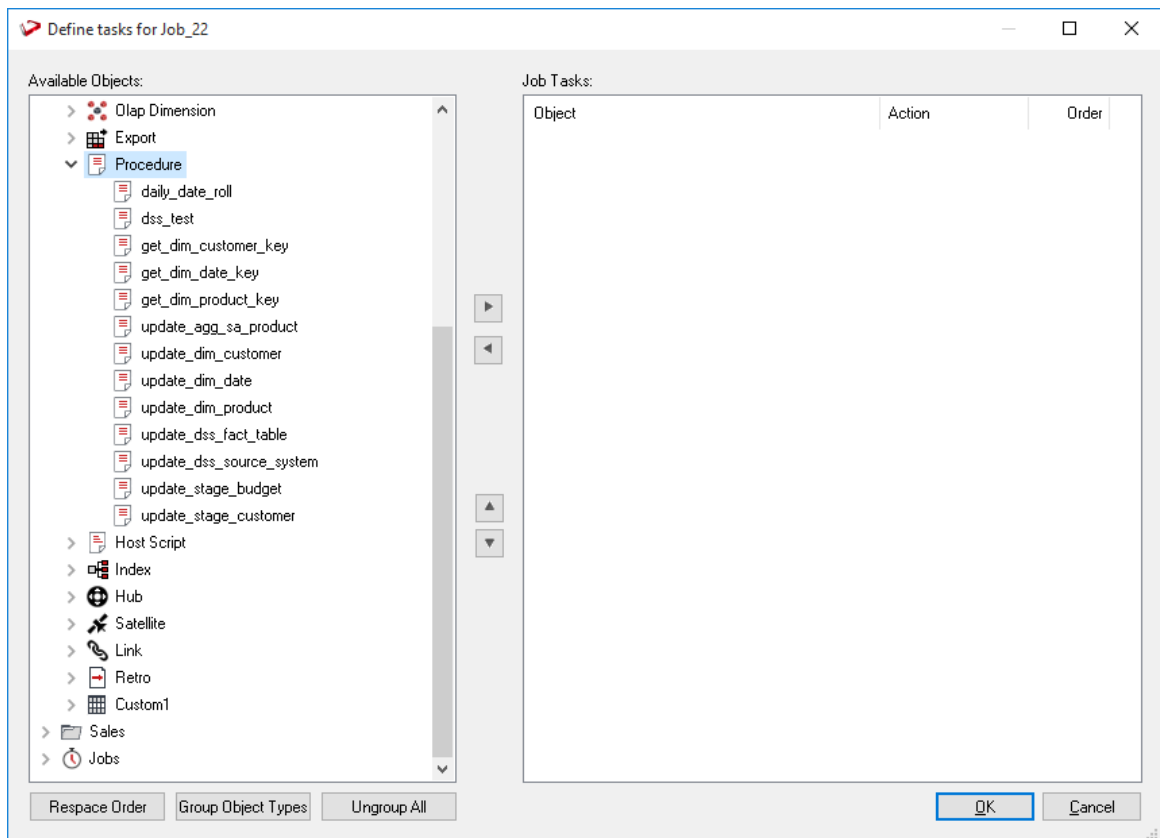
Logs Retained: This field lets you set the number of logs that are retained for this job before an automatic delete and archive occurs. 0 = keep all logs (default action)

The following two fields are optional. They are executed after the job completes and therefore need to reflect the scheduler environment. (i.e. Unix or Windows). The special variables \$JOB_KEY\$, \$JOB_SEQ\$ and \$JOB_NAME\$ can be used to return the associated values.
The Success command will be executed if a successful completion, the failure command will be executed if a job fails to complete:

Success Command:

Failure Command:

- 4 To select the test procedure as a task, open the Procedure object heading in the left pane. Choose **dss_test** and the > button. Click **OK**.



- 5 To run the job, click on the **All Jobs** button and then right-click on the job and select **Start the Job**.

CHAPTER 32

BACKING UP AND RESTORING METADATA

This chapter covers the moving, saving and reloading of metadata repository objects. The backup section describes the methods for backing up the metadata repository. It can also be backed up via normal database backup procedures. The load/restore section covers the metadata restoration functions available.

Various methods exist for getting new or changed metadata from one repository into another repository. These methods are covered in the applications, and managing multiple repositories sections.

IN THIS CHAPTER

Backup using DB Routines	852
Unloading Metadata	853
Loading an Unload	856

BACKUP USING DB ROUTINES

The backup of the metadata repository can be undertaken as a separate exercise from the general backup of the data warehouse. The metadata backup does not backup any of the actual data warehouse tables. It only saves the table definitions, index definitions, procedures etc., so is not normally large.

The backup includes any tables that begin with "dss_". In this way, key metadata tables such as dss_parameter, dss_source_system and dss_fact_table are included in the backup. These tables are required if the restored metadata is to function correctly.

It is recommended that the metadata back up is used to augment existing database back up strategies that are implemented independently from RED.

Windows Backup - WhereScape Unload and Load

WhereScape RED provides a generic unload and load utility for backing up and restoring the metadata. This is a Windows backup option that exists within WhereScape RED, which is a database independent backup that is designed for the restoration/migration of RED metadata but does not offer protection/security of data warehouse objects.

The unload and load utility can also be used to share metadata with WhereScape support, while troubleshooting issues. These options have the advantage of being database, and database version independent, so can be used to backup the metadata regardless of the version of the database client running on the PC.

It is possible to Load the metadata from a different database. For example, the metadata from a Teradata unload can be loaded into an Oracle or SQL Server database. There is, however, additional work in order to successfully move the metadata in this fashion. If such a move is required, please contact WhereScape support.

The unload and load process is detailed in the succeeding sections.

UNLOADING METADATA

WhereScape RED provides a generic unload utility for backing up the metadata. The advantage of this backup is that it is database, and database version independent. It can be used to backup the metadata regardless of the version of the database client running on the PC. It is possible to transport the metadata from one database platform to another using unload and load. For example, the metadata from a SQL Server unload can be loaded into an Oracle database.

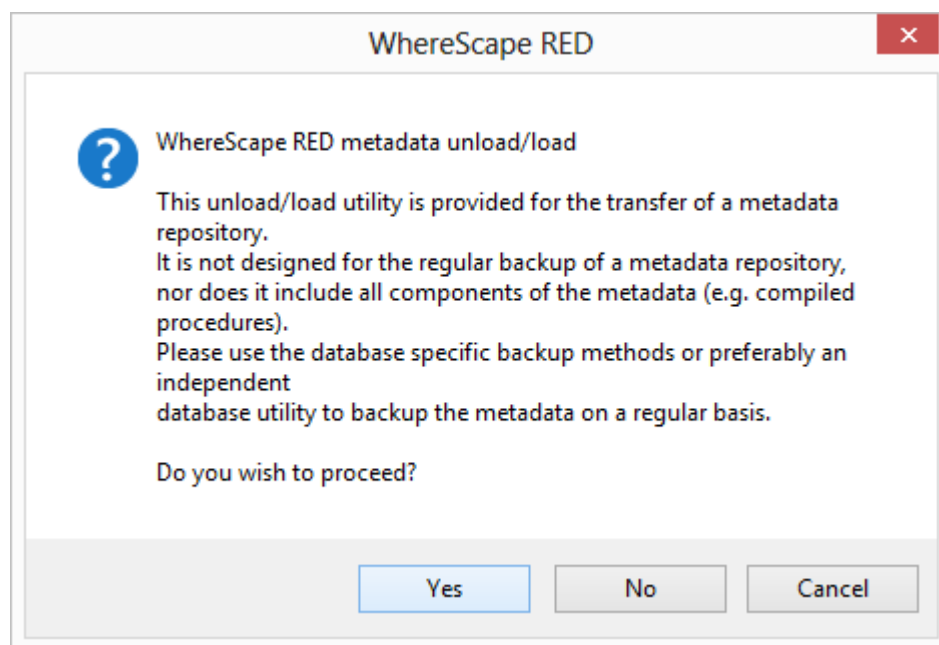
After performing a cross platform unload and load:

- Transformations must be altered manually to use the correct syntax for the new database.
- Generated procedures regenerated.
- Modified and custom procedures must be changed to use the procedure language of the new database platform.

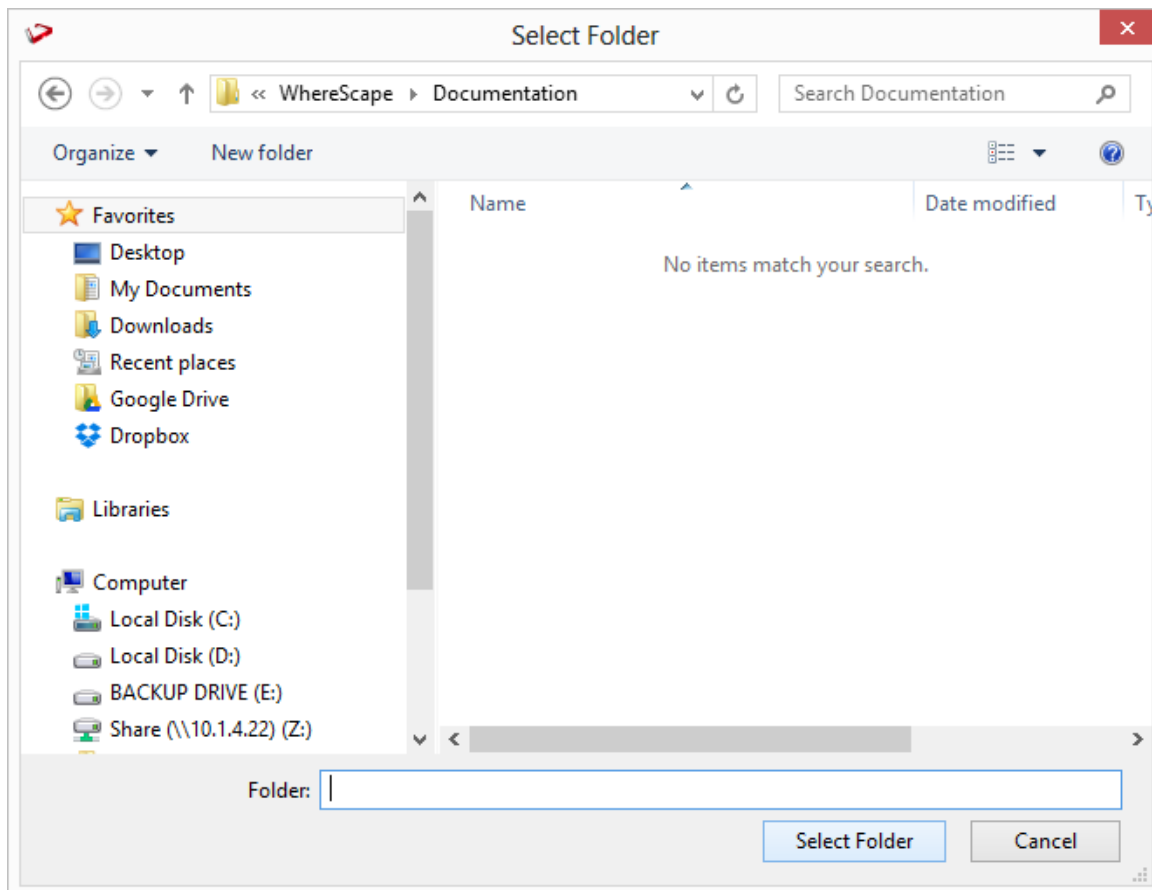
Performing an Unload

An unload can be performed within the WhereScape RED tool by selecting the **Backup>Unload the metadata to disk** menu option.


A popup window asks for confirmation to proceed with the unload. Click **Yes**.



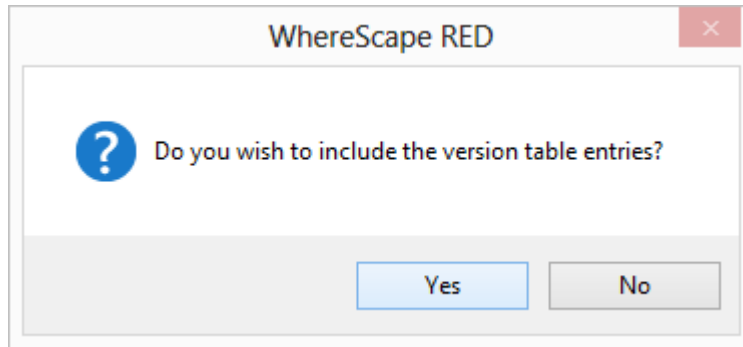
An additional pop-up window asks for a file name for the export. A directory should be chosen and a name entered. Click **Save**.



Enter a comment for the unload and click **OK**.



Finally, click **Yes** or **No** on the include version history dialog:



This either includes or excludes version history metadata in the unload.
The unload starts, and indicates progress with a progress bar.

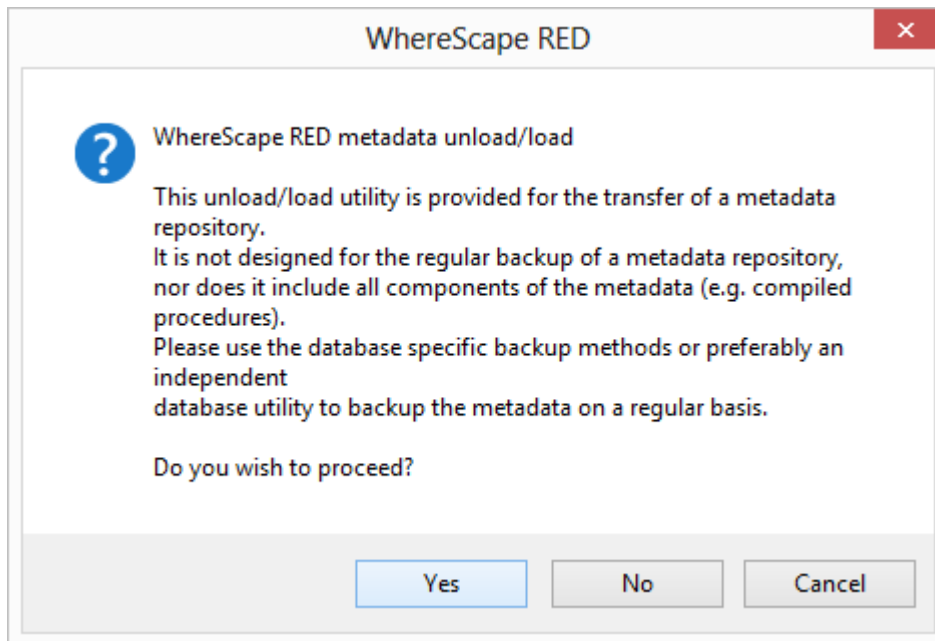
LOADING AN UNLOAD

WhereScape RED metadata can be restored from a prior unload.

Performing a Load

An unload of a metadata repository can be loaded over the top of an existing repository. This action replaces the existing repository in its entirety. To load an unload, select the menu option **Backup>Load the Metadata from disk** to begin the load process.


A popup window asked for confirmation a load is intended. Click **Yes**.





A dialog box appears. The word **RESTORE** needs to be entered. The ODBC connection needs to be chosen, along with the username and password where the metadata is to be restored to. The username does not have to be that of the current metadata repository, but it must be a valid repository. Click **OK**.


Restore Metadata from Backup ✕

WARNING: You have chosen to restore the Metadata
This will remove all existing Metadata and replace it with the
selected backup. Enter RESTORE in the window below and
then the username and password that you wish to restore into.
Click the ok button to proceed.

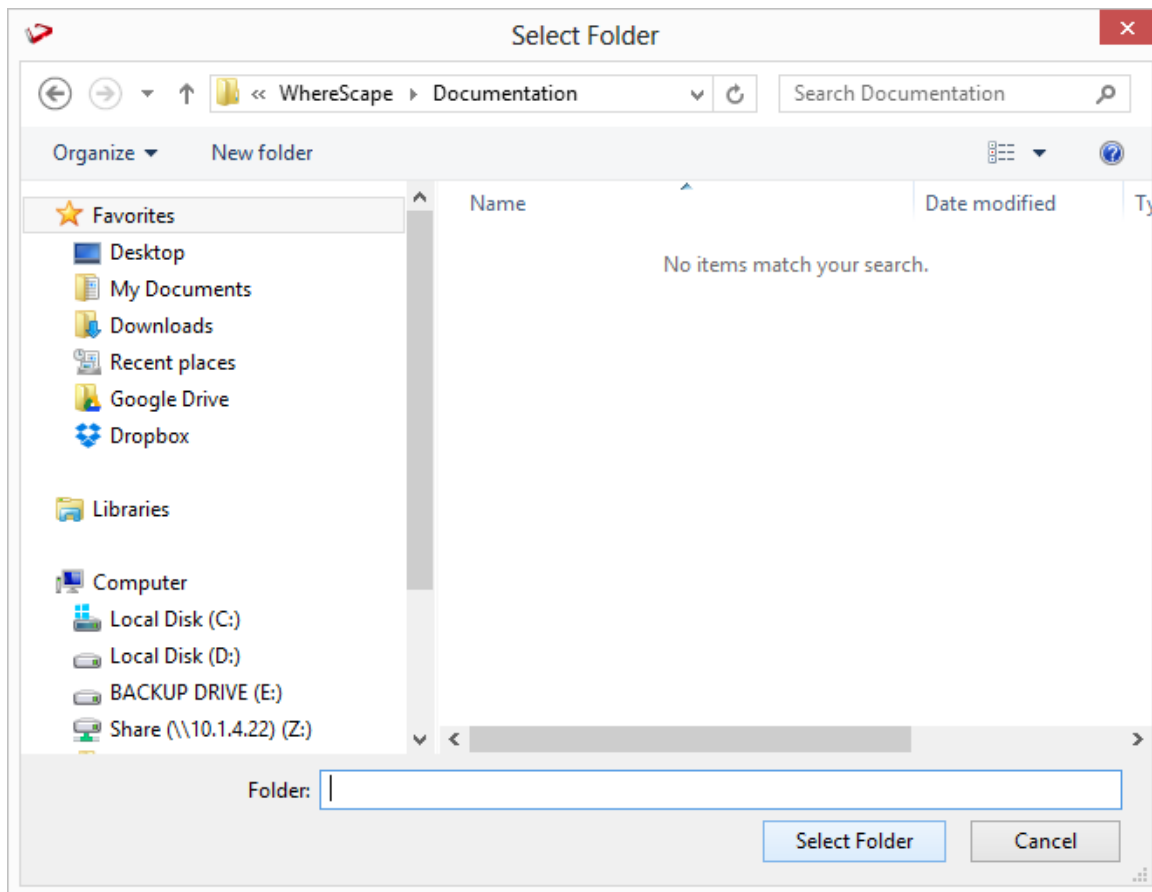
 **Odbc Connect**

 **User Name**

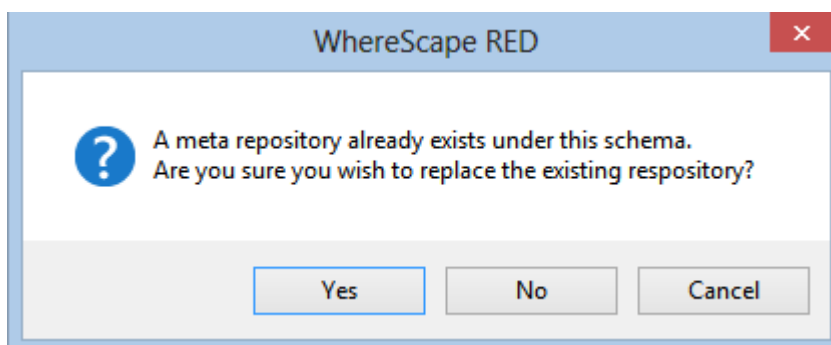
 **Password**

 **Meta Database**

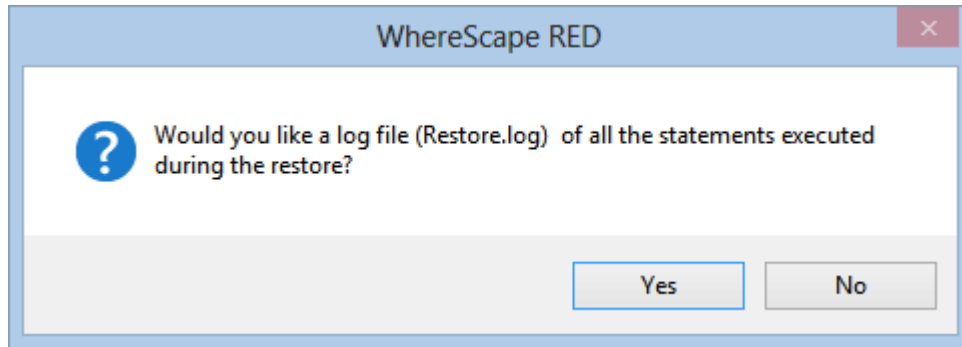
The next dialog box asks for the folder for the metadata to be loaded from. Browse to the contents of the directory where the unload is located and click **Select Folder**.




Confirm the load will overwrite the existing repository by clicking **Yes**:



A pop-up window asks if a restore log is required. Click **Yes** for a log, or No otherwise.



If a cross platform load is being performed, the following dialog is displayed. Click **Yes**.



The metadata load now runs. Once the load has completed, start WhereScape Administrator and validate the metadata repository that has just been loaded.

CHAPTER 33

ALTERING METADATA

This chapter provides information on how to change and manipulate the data warehouse once it has been established.

New source columns or changes to the source systems from which the data warehouse is built, will require modifications to both the metadata and the data warehouse tables and procedures.

IN THIS CHAPTER

Validating Tables	862
Validating Source (Load) Tables	864
Validating Procedures.....	865
Altering Tables	866
Validating Indexes	867
Recompiling Procedures	867

VALIDATING TABLES

The metadata as stored and maintained by RED does not necessarily reflect the actual tables and procedures in use in the data warehouse. For example, if a new column is added to the metadata for a table then that change is not automatically made in the actual physical table residing in the data warehouse. Likewise if a column is deleted from the metadata then that column may still exist in the physical database table.

This situation may be particularly apparent after an application patch or upgrade. The menu option **Validate>Validate Table Create Status**, and the right-click menu options in either the left or middle panes all provide a means of comparing the metadata to the physical tables in the database. A table, range of tables or all tables can be chosen. Each chosen table is a table in the metadata and it is compared against the physical database table if it exists.

The following example is the output from a validation.

Table Name	Results
load_customer	Validates OK
load_forecast	Validates OK
load_order_header	Validates OK
load_order_line	Validates OK
load_state	Table:meta->load_state column data not found,
model_customer	dss-> address - varchar(40),
model_date	Validates OK
model_forecast	meta->product_line - char(20),
model_order_header	Validates OK
model_order_line	Validates OK
ods_forecast	Validates OK
stage_customer	Validates OK - but different column order
stage_forecast	Validates OK
stage_order_header	Validates OK
stage_order_line	Validates OK

Reports
Middle Pane: model_customer Columns | Development (dssdemo) | UserId: dssdemo | Browse: No source system | CAP | NUM | SCRL

In this example we see five different scenarios.

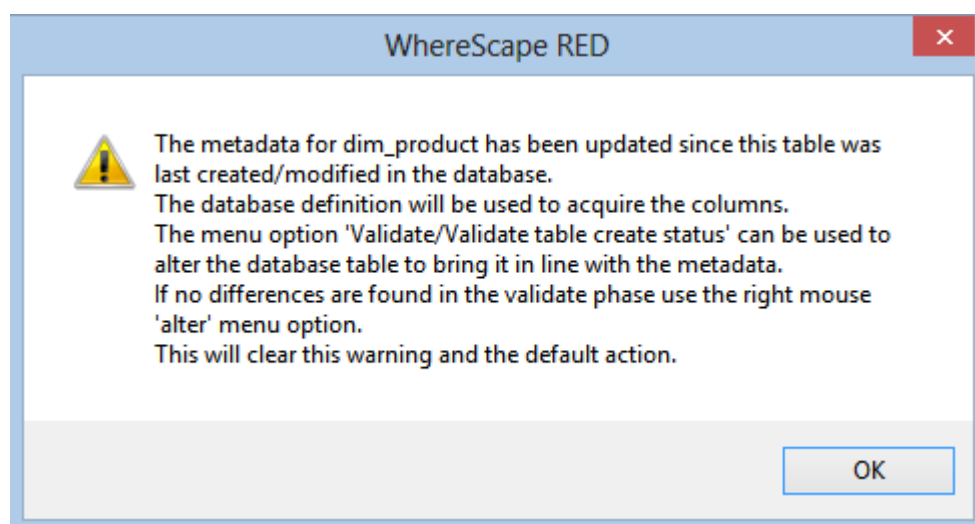
- 1 The metadata for table **load_customer** matches the physical table in the database.
- 2 The metadata for the table **model_forecast** does not match the physical table. The metadata has an additional column called 'product_line'. This column was not found in the physical table. The table can be altered if desired. See the next section on Altering Tables.
- 3 The physical database table **model_customer** has an additional column not found in the metadata. The column is 'address'. The table can be altered if desired. See the next section on Altering Tables.
- 4 The table **stage_customer** has the same columns in both the metadata and the physical table, but the column order is different. This is probably not an issue for most tables, but may be a

problem for some type of load tables, where the column order is important. This could be the result of a previous alter of the table. The table must be re-created if the order is important.

- 5 The table **load_state** is defined in the metadata but has not been physically created in the database. The table can be created in the normal manner.

Using outdated metadata in drag and drop

When dragging from a data warehouse table to create another data warehouse table (e.g. model_table to create view_table) a check is made to ensure that the metadata matches the database table. If the two are found to be out of sync the following message will appear:



If a subsequent validate of the table in question shows that it validates, this message will mean that the dates are somehow out of sync. This can occur for example after an import where the metadata has been replaced, but the underlying table still matches the metadata. Another common occurrence is where a new column is added and then deleted. To prevent the message from re-occurring in such a situation, proceed as follows.

Use the right-click menu to select **Alter Table** when positioned on the table name in the validate results screen (event though the table validates OK). The metadata update time will be set back to that of the last database table create.

VALIDATING SOURCE (LOAD) TABLES

Changes to the source systems from which the data warehouse is built can be detected to a limited degree. The menu option **Validate>Validate Load Table Status** allows a comparison between load tables and the source tables from which they were built. This comparison is not available for flat file or script based loads. A load table or group of load tables are selected and the results are displayed in the middle pane. An example screen from a load table validate is as follows:

Load Table	Results
load_customer	Source-> Additional source columns;
load_forecast	Source is non database, unable to validate;
load_order_header	Validates OK
load_order_line	Validates OK
load_state	MetaData-> column data not found;

Reports
 Middle Pane: load_order_line Columns | Development (dssdemo) | UserId: dssdemo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

The table **load_forecast** is a Windows file load and as such cannot be validated.

The table **load_customer** shows additional columns in the source table. Such a scenario will not cause problems for the continued operation of the data warehouse. It simply means that more columns are present in the source table than have been loaded into the data warehouse. This may have been the result of an initial partial selection or as a result of new columns. Further investigation of the source table would be required to ascertain if there was new information available.

The table **load_state** reflects a problem for the continued operation of the data warehouse. The source table does not have a column that was previously identified as having come from that table. This will probably cause the load of that table to fail. This scenario would also require an investigation into the source table. The resolution may be to delete the column. The potential impact on later tables (stage and model) and procedures in the data warehouse can be ascertained by using the right-click menu when positioned over a load table name.

VALIDATING PROCEDURES

The menu option **Validate>Validate Procedure Status** compares procedures as stored in the metadata with those compiled and running in the data warehouse. This option provides a listing in the middle pane of each selected procedure and its status. The status will be **Validates OK**, **Not compiled**, or **Compare failed**.

Where a procedure is marked as **Not compiled** this means that the procedure exists in the metadata but has not been compiled into the database.

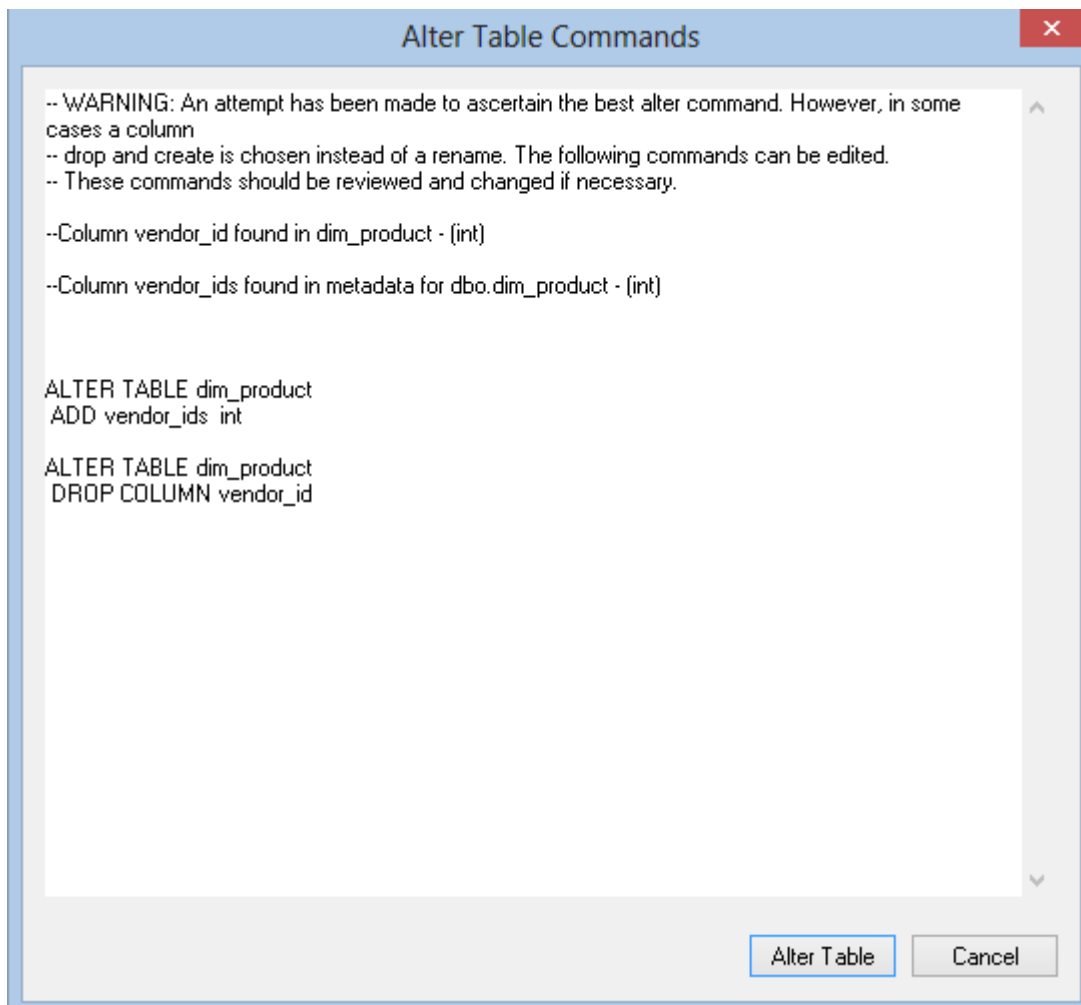
Where a procedure fails to compare the Procedure, the Editor must be used to find the actual differences. Start the editor by double clicking on the procedure name in the left pane. Use the **Tools>Compare to Compiled Source** menu option to display the differences between the procedure in the metadata and that compiled in the database.

ALTERING TABLES

The previous section covered the process of validating a table as defined in the metadata with the physical table as defined in the database. If a table is found to be different it is possible to alter the table.

Note: Care should be taken when altering large data warehouse tables. A number of factors such as the time required to perform the alter, access to the table and the optimum storage of the table come into play.

To alter a table first validate the table through the **Validate>Validate Table Create Status** menu option, or the right-click menu option from the object name. Then in the middle pane (the validation listing) select the table that has not validated. Position on the table name and using the right mouse, select the **Alter Table** pop-up menu option. A screen similar to the one below will appear advising of the planned changes.



In this example, the **dim_product** table is to be altered. The new column 'State' will be added to the table. Comments at the top of the screen, show the reason(s) for the alteration and the actual alter command(s) follow.

The alter table window is an edit window. The command to be executed can be changed or additional commands entered. The command may also be cut to be executed in some other environment or at some later stage.

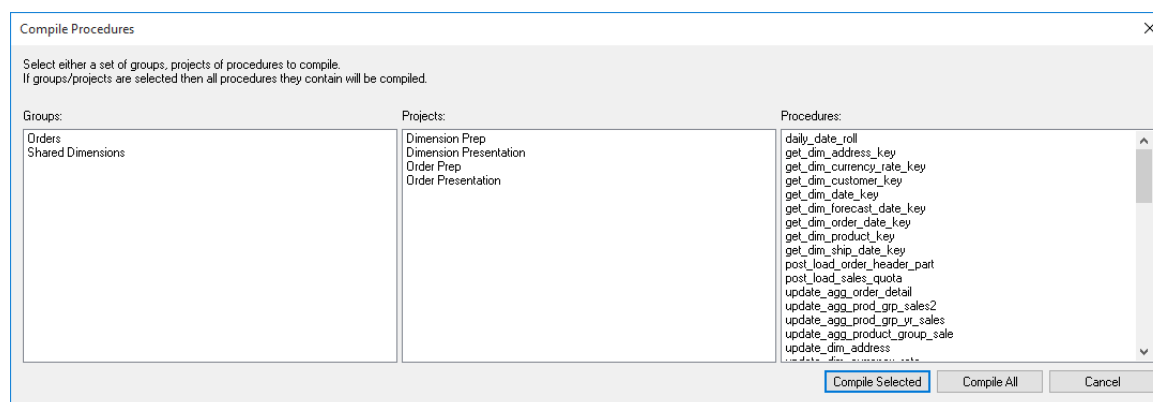
Clicking the **Alter Table** button will proceed to alter the table. In effect, it will execute any command in the window.

VALIDATING INDEXES

Index validate is not available on Teradata.

RECOMPILING PROCEDURES

Procedures can be invalidated as a result of changes to underlying tables, or child procedures. A procedure can be recompiled through the procedure editor or via the menu option **Tools>Compile Procedures**. This menu option displays the following dialog box:



As each procedure is compiled, it is displayed in the middle pane of the builder window. If a procedure fails to compile, a failure message is displayed along side the procedure name. Procedures that fail to compile will need to be investigated through the procedure editor, as no specific error information is provided by this bulk compiler.

Compile Selected

A selected group of procedures may be compiled. Selection is done via standard Windows multi selection using the shift and CTRL keys. By selecting projects or groups, it is possible to compile all procedures associated with the projects or groups. Click on the **Compile Selected** button to compile/re-compile the selected procedures.

Compile All

Click on the **Compile All** button to compile/re-compile all procedures in the metadata repository.

Compile Missing

The **Compile Missing** button compiles all procedures in the metadata repository that do not exist in the metadata database in Teradata.

CHAPTER 34

CALLABLE ROUTINES

IN THIS CHAPTER

Introduction to Callable Routines	870
Ws_Api_Glossary	877
Ws_Connect_Replace	879
Ws_Job_Abort	880
Ws_Job_Clear_Archive	882
Ws_Job_Clear_Logs	885
Ws_Job_Clear_Logs_By_Date	887
Ws_Job_Create	889
Ws_Job_CreateWait	893
Ws_Job_Dependency	897
Ws_Job_Release	900
Ws_Job_Restart	902
Ws_Job_Schedule	905
Ws_Job_Status	907
Ws_Load_Change	913
Ws_Maintain_Indexes	915
Ws_Version_Clear	918
WsParameterRead	921
WsParameterReadF	922
WsParameterReadG	923
WsParameterWrite	925
WsWrkAudit	926
WsWrkAuditBulk	929
WsWrkError	932
WsWrkErrorBulk	935
WsWrkTask	938

INTRODUCTION TO CALLABLE ROUTINES

CALLABLE ROUTINES API

WhereScape RED callable routines provide an Application Program Interface (API) to the WhereScape RED metadata using the following SQL-invoked routines:

Routine Name	Description
<i>Ws_Api_Glossary</i> (on page 877)	Adds an entry to the documentation glossary.
<i>Ws_Connect_Replace</i> (on page 879)	Replaces the contents of a connection with details from another connection.
<i>Ws_Job_Abort</i> (on page 880)	Aborts a job if it is in a running state.
<i>Ws_Job_Clear_Archive</i> (on page 882)	Purges archived job logs that are older than the specified age in days.
<i>Ws_Job_Clear_Logs</i> (on page 885)	Archives job logs when the maximum number of current logs to retain is exceeded.
<i>Ws_Job_Clear_Logs_By_Date</i> (on page 887)	Archives job logs that are older than the specified age in days.
<i>Ws_Job_Create</i> (on page 889)	Creates a job based on an existing job and optionally starts it immediately.
<i>Ws_Job_CreateWait</i> (on page 893)	Creates a job based on an existing job and schedules it to start later.
<i>Ws_Job_Dependency</i> (on page 897)	Adds or removes a child-to-parent dependency between two jobs to control the child job.
<i>Ws_Job_Release</i> (on page 900)	Starts a job if it is in a holding or waiting state.
<i>Ws_Job_Restart</i> (on page 902)	Starts a job if it is in a failed state.
<i>Ws_Job_Schedule</i> (on page 905)	Schedules a job if it is in a holding or waiting state.
<i>Ws_Job_Status</i> (on page 907)	Returns the current status of a job.
<i>Ws_Load_Change</i> (on page 913)	Changes the Connection or Schema of a load table.
<i>Ws_Maintain_Indexes</i> (on page 915)	Drops and/or builds database indexes that are defined in the WhereScape RED metadata.
<i>Ws_Version_Clear</i> (on page 918)	Purges metadata versions for all objects that do not meet the specified retention criteria.

Routine Name	Description
<i>WsParameterRead</i> (on page 921)	Returns the value and comment (for most RDBMS) of a WhereScape RED metadata Parameter.
<i>WsParameterReadF</i> (on page 922)	Returns the value of a WhereScape RED metadata Parameter. [SQL Server only]
<i>WsParameterReadG</i> (on page 923)	Returns the value of a "global" WhereScape RED metadata Parameter that relates to a load table.
<i>WsParameterWrite</i> (on page 925)	Updates the value and comment of a WhereScape RED metadata Parameter or creates it.
<i>WsWrkAudit</i> (on page 926)	Records a message in the Audit Log.
<i>WsWrkAuditBulk</i> (on page 929)	Records multiple messages in the Audit Log.
<i>WsWrkError</i> (on page 932)	Records a message in the Error/Detail Log.
<i>WsWrkErrorBulk</i> (on page 935)	Records multiple messages in the Error/Detail Log.
<i>WsWrkTask</i> (on page 938)	Updates row counts for a task in the Task Log.

CALLABLE ROUTINES PER RDBMS

Each WhereScape RED Callable Routine exists in the database as either a Stored Procedure or a User-defined Function that can be invoked from SQL. A Callable Procedure is invoked by a SQL call/execute statement and a Callable Function is invoked in a SQL SELECT statement or a value expression. The Callable Routines are typically implemented as **procedures in most database systems** due to RDBMS limitations of user-defined functions when the WhereScape RED API was originally developed.

Routine Name	Teradata
<i>Ws_Api_Glossary</i> (on page 877)	Procedure
<i>Ws_Connect_Replace</i> (on page 879)	Procedure
<i>Ws_Job_Abort</i> (on page 880)	Procedure
<i>Ws_Job_Clear_Archive</i> (on page 882)	Procedure
<i>Ws_Job_Clear_Logs</i> (on page 885)	Procedure
<i>Ws_Job_Clear_Logs_By_Date</i> (on page 887)	Procedure
<i>Ws_Job_Create</i> (on page 889)	Procedure
<i>Ws_Job_CreateWait</i> (on page 893)	Procedure
<i>Ws_Job_Dependency</i> (on page 897)	Procedure
<i>Ws_Job_Release</i> (on page 900)	Procedure
<i>Ws_Job_Restart</i> (on page 902)	Procedure
<i>Ws_Job_Schedule</i> (on page 905)	Procedure
<i>Ws_Job_Status</i> (on page 907)	Procedure
<i>Ws_Load_Change</i> (on page 913)	Procedure
<i>Ws_Maintain_Indexes</i> (on page 915)	Procedure
<i>Ws_Version_Clear</i> (on page 918)	Procedure
<i>WsParameterRead</i> (on page 921)	Procedure
<i>WsParameterReadF</i> (on page 922)	Procedure
<i>WsParameterReadG</i> (on page 923)	Procedure
<i>WsParameterWrite</i> (on page 925)	Procedure
<i>WsWrkAudit</i> (on page 926)	Procedure
<i>WsWrkAuditBulk</i> (on page 929)	Procedure
<i>WsWrkError</i> (on page 932)	Procedure
<i>WsWrkErrorBulk</i> (on page 935)	Procedure
<i>WsWrkTask</i> (on page 938)	Procedure

CALLABLE ROUTINES NAMES QUALIFIER

The WhereScape RED Callable Routines can be invoked using the unqualified routine name for SQL Server and Oracle. However, for Teradata and DB2 it is necessary to qualify the routine name with the owner/schema of the WhereScape RED metadata repository. All RED-generated procedures in a Teradata or DB2 repository that invoke a WhereScape RED Callable Routine do so by qualifying the routine name with **[METABASE]** e.g. **[METABASE].routine_name**. When RED creates/compiles a procedure in a Teradata or DB2 database it automatically replaces the **[METABASE]** "token" with the repository owner/schema.

When you edit a RED-generated procedure or create your own custom procedure in Teradata or DB2 you can qualify each callable routine name by either hard-coding the owner/schema or by using the **[METABASE]** "token" that RED will replace when the procedure is created/compiled in the database. The Teradata and DB2 examples in this chapter use the **[METABASE]** "token" but if you invoke a WhereScape RED Callable Routine interactively or from outside a RED-compiled procedure/function then the actual owner/schema must be specified because RED won't have the chance to replace the "token".

CALLABLE ROUTINES COMMON INPUT

The following input parameters are common to most of the WhereScape RED Callable Routines, which are primarily used for integration with the WhereScape RED Scheduler.

Input	Parameter Name	Description
Job Instance Identifier	p_sequence	<p>Unique identifier of the running job (i.e. the running instance of a held or scheduled job) that executed the routine.</p> <p>When invoked from a WhereScape RED Scheduler the routine will be passed the parameter argument.</p> <p>When invoked manually or externally to the WhereScape RED Scheduler, then any integer value can be used.</p>
Job Name	p_job_name	<p>Name of the running job that executed the routine.</p> <p>When invoked from a WhereScape RED Scheduler the routine will be passed the parameter argument.</p> <p>When invoked manually or externally to the WhereScape RED Scheduler, then any name can be used.</p>

Input	Parameter Name	Description
Task Name	p_task_name	Name of the running task (of a running job) that executed the routine. When invoked from a WhereScape RED Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the WhereScape RED Scheduler, then any name can be used.
Job Identifier	p_job_id	Unique identifier of the held or scheduled job that the running job is a specific instance of. When invoked from a WhereScape RED Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the WhereScape RED Scheduler, it is recommended to use 0 (zero).
Task Identifier	p_task_id	Unique identifier of the running task (of a running job) that executed the routine. When invoked from a WhereScape RED Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the WhereScape RED Scheduler, it is recommended to use 0 (zero).

Note: Typically, the **parameter names** of the WhereScape RED Callable Routines use a p_ prefix as indicated but in some routines a v_ prefix is used instead. In addition, for SQL Server all parameter names are also prefixed by @. The RDBMS-specific parameter names are included in the subsequent details of each Callable Routine.

CALLABLE ROUTINES INVOCATION

This chapter includes examples of each WhereScape RED Callable Routine to illustrate how to invoke the routine for each RDBMS, along with the necessary variable declarations. Typically, the routines are executed from a Stored Procedure so the routine-specific examples illustrate that scenario and use parameters/arguments that are passed by position (as in a RED-generated procedure). However, it can be useful to execute a Callable Routine using other methods such as the following.

ALTERNATIVE INVOCATION METHODS

Invocation via ODBC

The following examples illustrate how to invoke a WhereScape RED Callable Routine via an ODBC connection (using a tool such as WhereScape SQL Admin), which uses both input and output parameters. The examples work for a SQL Server RED repository but for another RDBMS, the routine name may need to be qualified with the appropriate owner/schema. Refer to the detailed description of the `Ws_Connect_Replace` routine for an explanation of the parameters/arguments.

The following invocations via ODBC are equivalent and the only difference is the formatting as the first command uses a single line while the second command is formatted across multiple lines:

-- ODBC Example 1 (single line).

```
{ CALL Ws_Connect_Replace( 0, 'Test Job Name', 'Test Task Name', 0, 0,
'REPLACE', 'Connection1', 'Connection2', ?, ?, ?) };
```

-- ODBC Example 2 (same command formatted using multiple lines).

```
{ CALL Ws_Connect_Replace
( 0
, 'Test Job Name'
, 'Test Task Name'
, 0
, 0
, 'REPLACE'
, 'Connection1'
, 'Connection2'
, ?
, ?
, ?
)
};
```

The result of the `Ws_Connect_Replace` invocation can be confirmed by checking the target connection. In addition, a log entry is created using the specified job and task names that can be viewed via the Logs – Recent Audit Trail Logs menu item of the WhereScape RED Scheduler.

Invocation via the Command-Line

Each WhereScape RED Callable Routine can also be invoked from the command-line, using an RDBMS-specific tool such as:

RDBMS	SQL Command-Line-Interface (CLI) Tool
Teradata	Basic Teradata Query (BTEQ) i.e. bteq.

These tools can be used to invoke a WhereScape RED Callable Routine by connecting to the database, executing SQL statements, and disconnecting from the database. Refer to the RDBMS-specific documentation for details of how to connect and execute SQL via the relevant CLI tool as well as details of tool's return codes.

Typically, the CLI command will include options to specify the database, connection credentials, and the SQL to execute. Multiple SQL statements can typically be executed by terminating/delimiting each statement using a semi-colon (;). In most cases, the SQL needs to be quoted (typically double-quoted) when specified on the command-line and if the SQL statements include embedded quotes (such as single-quotes around literals) then they may need to be "escaped" depending on the CLI tool and the platform. Most of the tools also allow the SQL commands to be read from a file instead of from standard input.

WS_API_GLOSSARY

Synopsis

Adds an entry to the documentation glossary.

Description

Adds the specified entry to the documentation glossary, which is included in documentation that is subsequently generated by WhereScape RED.

Input

Input	Description
Object Name	Item that appears in the left column of the glossary, which normally represents the business name of a column in a dimension or fact table (although it can be used for other purposes).
Glossary Term	Item that appears under the analysis area heading of the glossary, which normally represents a dimension or fact table name (although it can be used for other purposes).
Glossary Description	Description of the term being defined.
Action	Either ADD or DELETE the specified glossary entry.

Output

Output	Description
Result Text	A message to indicate whether or not the glossary term was successfully added/deleted.

Teradata Parameters: Ws_Api_Glossary

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_object_name	VARCHAR(64)	IN
p_term	VARCHAR(256)	IN
p_comments	VARCHAR(4000)	IN

Parameter Name	Datatype	Mode
p_option	VARCHAR(64)	IN
p_result	VARCHAR(256)	OUT

Teradata Examples: Ws_Api_Glossary

```
DECLARE v_result_txt          varchar(256);
CALL [METABASE].Ws_Api_Glossary
( 'Data Warehouse'
, 'Overview'
, 'A repository of business information'
, 'ADD'
, v_result_txt
);
```

WS_CONNECT_REPLACE

Synopsis

Replaces the contents of a connection with the details from another connection.

Description

Copies the details of the specified Source connection to the specified Target connection.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Action	REPLACE the details of the specified Target connection.
Source Connection	The name of the Source connection whose details will be copied.
Target Connection	The name of the Target connection whose details will be changed.

Output

Output	Description
Return Code	Output Return Code: S Success. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.
Result Number	Output Result Number: 1 Success. -2 Error. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Connect_Replace

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_action	VARCHAR(64)	IN
p_source	VARCHAR(64)	IN
p_target	VARCHAR(64)	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Connect_Replace

```
-- The p_ VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence          integer;
DECLARE p_job_name         varchar(256);
DECLARE p_task_name       varchar(256);
DECLARE p_job_id          integer;
DECLARE p_task_id         integer;
DECLARE p_return_msg      varchar(256);
DECLARE p_status          integer;
DECLARE v_result_num      integer;
DECLARE v_return_code     varchar(1);
DECLARE v_return_msg      varchar(256);
CALL [METABASE].Ws_Connect_Replace
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'REPLACE', 'Connection1', 'Connection2'
, v_return_code
, v_return_msg
, v_result_num
);
```

WS_JOB_ABORT

Synopsis

Aborts a job if it is in a running state.

Description

Aborts the specified job if it is in a running state, which changes it to a failed state, fails all running tasks, and holds all waiting tasks.

Input

Input	Description
Abort Job Name	The name of the job to be aborted. The specified name must exactly match the job name as displayed by the WhereScape RED Scheduler. The specified job must be in a running or failed state in order to be aborted.
Job Instance Identifier	Unique identifier of the running job (that may be in a failed state).
Abort Job Message Text	Custom message text to be recorded in the WhereScape RED Audit Log for the aborted job and the WhereScape RED Task Log for each aborted task.

Teradata Parameters: Ws_Job_Abort

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_job_name	VARCHAR(64)	IN
p_job_sequence	INTEGER	IN
p_job_msg	VARCHAR(256)	IN

Teradata Examples: Ws_Job_Abort

```
CALL [METABASE].Ws_Job_Abort
( 'Daily Run'
, 1234
, 'Job aborted via manual execution of Ws_Job_Abort.'
);
```

WS_JOB_CLEAR_ARCHIVE

Synopsis

Purges archived job logs that are older than the specified age in days.

Description

Deletes job-related logs that were previously archived (into the WX_WRK_AUDIT_ARCHIVE and WX_WRK_ERROR_ARCHIVE tables via a RED Scheduler and/or RED callable routines such as Ws_Job_Clear_Logs and Ws_Job_Clear_Logs_By_Date) depending on their age in days.

When the maximum age of the archived logs to retain is exceeded all the older logs are deleted. For example, if 90 days are retained then all the archived logs that are older than 90 days are deleted. If a maximum age of 0 days is specified then all the archived logs are deleted.

Alternatively, the **TRUNCATE** option can be used to remove all the archived logs, which overrides all other criteria.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Maximum Days to Retain	The maximum age (in days) of the archived logs to retain. If 90 days are retained then all the archived logs that are older than 90 days are purged/deleted . If 0 days are retained then all the archived logs are purged/deleted .
Job Name to Purge	The name of the job whose archived logs are to be purged . Wild cards are supported. Specifying % will match ALL jobs.
Options	The TRUNCATE option can be used to remove ALL the archived logs, which overrides all other criteria i.e. irrespective of the days to retain or job name.

Output

Output	Description
Return Code	Output Return Code: S Success. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.

Output	Description
Result Number	Output Result Number: 1 Success. -2 Error. e.g. Due to invalid job name or job not running. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Job_Clear_Archive

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_day_count	INTEGER	IN
p_job	VARCHAR(64)	IN
p_options	VARCHAR(256)	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Job_Clear_Archive

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
DECLARE v_return_code        varchar(1);
DECLARE v_return_msg        varchar(256);
CALL [METABASE].Ws_Job_Clear_Archive
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 90, 'Daily Run', ''
, v_return_code
, v_return_msg
, v_result_num
);
```


WS_JOB_CLEAR_LOGS

Synopsis

Archives job logs when the maximum number of current logs to retain is exceeded.

Description

Moves job-related logs from the current log tables (such as WS_WRK_AUDIT_LOG and WS_WRK_ERROR_LOG) to the corresponding archive log tables (such as WX_WRK_AUDIT_ARCHIVE and WX_WRK_ERROR_ARCHIVE) depending on the number of logs to retain.

When the maximum number of current logs to retain is exceeded the oldest logs are archived for the specified job(s) to reduce the number of current logs to the specified retention limit. For example, if 10 is specified then only the latest 10 logs are retained. If a retained count of 0 is specified then all the current logs are archived for the specified job(s).

Note: Equivalent functionality is available via a WhereScape RED Scheduler and the "Logs Retained" property of a job.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Job Name(s) to Archive	The name of the job(s) whose current logs are to be archived . Wild cards are supported. Specifying % will match ALL jobs.
Maximum Logs to Retain	The maximum number of logs to retain. When the maximum is exceeded the oldest logs are archived to reduce the number of current logs to the specified retention limit.

Output

Output	Description
Return Code	Output Return Code: S Success. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.
Result Number	Output Result Number: 1 Success. -2 Error. e.g. Due to invalid job name or job not running. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Job_Clear_Logs

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_job_to_clean	VARCHAR(64)	IN
p_log_keep	INTEGER	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Job_Clear_Logs

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
DECLARE v_return_code        varchar(1);
DECLARE v_return_msg        varchar(256);
CALL [METABASE].Ws_Job_Clear_Logs
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'Daily Run', 10
, v_return_code
, v_return_msg
, v_result_num
);
```

WS_JOB_CLEAR_LOGS_BY_DATE

Synopsis

Archives job logs that are older than the specified age in days.

Description

Moves job-related logs from the current log tables (such as WS_WRK_AUDIT_LOG and WS_WRK_ERROR_LOG) to the corresponding archive log tables (such as WX_WRK_AUDIT_ARCHIVE and WX_WRK_ERROR_ARCHIVE) depending on their age in days.

When the maximum age of the current logs to retain is exceeded all the older logs are archived for the specified job(s). For example, if 90 days are retained then all the current logs that are older than 90 days are archived.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Job Name(s) to Archive	The name of the job(s) whose current logs are to be archived . Wild cards are supported. Specifying % will match ALL jobs.
Maximum Days to Retain	The maximum age (in days) of the current logs to retain. If 90 days are retained, then all the current logs that are older than 90 days are archived.

Output

Output	Description
Return Code	Output Return Code: S Success. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.
Result Number	Output Result Number: 1 Success. -2 Error. e.g. Due to invalid job name or job not running. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Job_Clear_Logs_By_Date

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_job_to_clean	VARCHAR(64)	IN
p_day_count	INTEGER	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Job_Clear_Logs_By_Date

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
DECLARE v_return_code       varchar(1);
DECLARE v_return_msg        varchar(256);
CALL [METABASE].Ws_Job_Clear_Logs_By_Date
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'Daily Run', 90
, v_return_code
, v_return_msg
, v_result_num
);
```

WS_JOB_CREATE

Synopsis

Creates a job based on an existing job and optionally starts it immediately.

Description

Creates a job from the specified existing job, if it is in either a holding or waiting state. The new job can be started immediately. Typically, this routine is used to create & start a job from within another job. Only jobs that are in a holding or waiting state can be used as a template for the new job.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Template Job Name	The name of the job to be used as a template for the new job. The specified name must exactly match the job name as displayed by the WhereScape RED Scheduler. The specified job must be in a holding or waiting state in order to be used as a template.
New Job Name	The name of the job to be created. The new job name cannot already exist.
Job Description	A description of the new job. When not specified the setting of the Template job is copied.
New Job Status	The initial status/frequency of the new job: HOLD - The status of the new job will show as 'On Hold'. The job will not run until it is subsequently released. ONCE - The new job will start immediately and upon successful completion the new job will be deleted. ONCE+HOLD - The new job will start immediately and upon successful completion the status will show as 'On Hold'.
Thread Count	The number of threads for the new job. When not specified the setting of the Template job is copied.
Scheduler Preference	A scheduler type or a specific scheduler name that is allowed to run the job. When not specified the setting of the Template job is copied. Note: Some jobs/tasks can only run in a specific environment such as Windows or UNIX/Linux
Maximum Logs to Retain	The maximum number of logs to retain. When not specified the setting of the Template job is copied.

Input	Description
Success Command	A command-line action to execute upon successful completion of the new job. When not specified the setting of the Template job is copied. The command must be executable within the context of the scheduler that runs the job so it must be a valid Windows/UNIX/Linux command that is appropriate to the scheduler environment.
Failure Command	A command-line action to execute upon failure of the new job. When not specified the setting of the Template job is copied. The command must be executable within the context of the scheduler that runs the job so it must be a valid Windows/UNIX/Linux command that is appropriate to the scheduler environment.

Output

Output	Description
Return Code	Output Return Code: S Success. N No action because the Template Job is not in a holding or waiting state. P No action because the New Job name already exists. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.
Result Number	Output Result Number: 1 Success. -1 Template Job is not in a holding/waiting state or the New Job name already exists. -2 Error. e.g. Due to invalid job name or job not running. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Job_Create

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN

Parameter Name	Datatype	Mode
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_template_job	VARCHAR(64)	IN
p_new_job	VARCHAR(64)	IN
p_description	VARCHAR(256)	IN
p_state	VARCHAR(10)	IN
p_threads	INTEGER	IN
p_scheduler	VARCHAR(8)	IN
p_logs	INTEGER	IN
p_okay	VARCHAR(256)	IN
p_fail	VARCHAR(256)	IN
p_att1	VARCHAR(4000)	IN
p_att2	VARCHAR(4000)	IN
p_att3	VARCHAR(4000)	IN
p_att4	VARCHAR(4000)	IN
p_att5	VARCHAR(4000)	IN
p_att6	VARCHAR(4000)	IN
p_att7	VARCHAR(4000)	IN
p_att8	VARCHAR(4000)	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Job_Create

```
-- The p_ VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence          integer;
DECLARE p_job_name         varchar(256);
DECLARE p_task_name        varchar(256);
DECLARE p_job_id           integer;
DECLARE p_task_id          integer;
DECLARE p_return_msg       varchar(256);
DECLARE p_status           integer;
DECLARE v_result_num       integer;
DECLARE v_return_code      varchar(1);
DECLARE v_return_msg       varchar(256);
CALL [METABASE].Ws_Job_Create
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'Daily Run', 'New Daily Run', 'This is the New Daily Run job.', 'ONCE'
, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL,
NULL, NULL
, v_return_code
, v_return_msg
, v_result_num
);
```


WS_JOB_CREATEWAIT

Synopsis

Creates a job based on an existing job and schedules it to start later.

Description

Creates a job from the specified existing job, if it is in either a holding or waiting state. The new job is scheduled to start later at the specified release time. Typically, this routine is used to create & schedule a job from within another job. Only jobs that are in a holding or waiting state can be used as a template for the new job.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Template Job Name	The name of the job to be used as a template for the new job. The specified name must exactly match the job name as displayed by the WhereScape RED Scheduler. The specified job must be in a holding or waiting state in order to be used as a template.
New Job Name	The name of the job to be created. The new job name cannot already exist.
Job Description	A description of the new job. When not specified the setting of the Template job is copied.
New Job Status	The initial status/frequency of the new job: HOLD - The status of the new job will show as 'On Hold'. The job will not run until it is subsequently released. ONCE - The new job will start immediately and upon successful completion the new job will be deleted. ONCE+HOLD - The new job will start immediately and upon successful completion the status will show as 'On Hold'.
Scheduled Release Date/Time	The date/time when the new job is scheduled to be run.
Thread Count	The number of threads for the new job. When not specified the setting of the Template job is copied.
Scheduler Preference	A scheduler type or a specific scheduler name that is allowed to run the job. When not specified the setting of the Template job is copied. Note: Some jobs/tasks can only run in a specific environment such as Windows or UNIX/Linux

Input	Description
Maximum Logs to Retain	The maximum number of logs to retain. When not specified the setting of the Template job is copied.
Success Command	A command-line action to execute upon successful completion of the new job. When not specified the setting of the Template job is copied. The command must be executable within the context of the scheduler that runs the job so it must be a valid Windows/UNIX/Linux command that is appropriate to the scheduler environment.
Failure Command	A command-line action to execute upon failure of the new job. When not specified the setting of the Template job is copied. The command must be executable within the context of the scheduler that runs the job so it must be a valid Windows/UNIX/Linux command that is appropriate to the scheduler environment.

Output

Output	Description
Return Code	Output Return Code: S Success. N No action because the Template Job is not in a holding or waiting state. P No action because the New Job name already exists. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.
Result Number	Output Result Number: 1 Success. -1 Template Job is not in a holding/waiting state or the New Job name already exists. -2 Error. e.g. Due to invalid job name or job not running. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Job_CreateWait

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_template_job	VARCHAR(64)	IN
p_new_job	VARCHAR(64)	IN
p_description	VARCHAR(256)	IN
p_state	VARCHAR(10)	IN
p_release_time	TIMESTAMP	IN
p_threads	INTEGER	IN
p_scheduler	VARCHAR(8)	IN
p_logs	INTEGER	IN
p_okay	VARCHAR(256)	IN
p_fail	VARCHAR(256)	IN
p_att1	VARCHAR(4000)	IN
p_att2	VARCHAR(4000)	IN
p_att3	VARCHAR(4000)	IN
p_att4	VARCHAR(4000)	IN
p_att5	VARCHAR(4000)	IN
p_att6	VARCHAR(4000)	IN
p_att7	VARCHAR(4000)	IN
p_att8	VARCHAR(4000)	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Job_CreateWait

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence          integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name        varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id          integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
DECLARE v_return_code        varchar(1);
DECLARE v_return_msg        varchar(256);
CALL [METABASE].Ws_Job_CreateWait
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'Daily Run', 'New Daily Run', 'This is the New Daily Run job.', 'ONCE'
, (CURRENT_TIMESTAMP + INTERVAL '1' MONTH)
, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL,
NULL, NULL
, v_return_code
, v_return_msg
, v_result_num
);
```

WS_JOB_DEPENDENCY

Synopsis

Adds or removes a child-to-parent dependency between two jobs to control the child job.

Description

Adds or removes a child-to-parent dependency between two jobs to control the child job. The dependent child job can be defined to fail (if necessary) when the parent job does not complete successfully in the required timeframe. The acceptable timeframe can be defined in terms of the maximum minutes in the past to look back and the maximum minutes in the future to wait for successful completion of the parent job.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Action	Either ADD or DELETE the job dependency.
Parent Job Name	The name of Parent Job that the Child Job will depend on.
Child Job Name	The name of the Child Job that will be dependent on the Parent Job.
Parent Success Required	Indicates whether or not the Child Job will fail when the Parent Job does not complete successfully in the required time frame.
Maximum Look Back Minutes	The Maximum minutes in the past to look back for successful completion of the Parent Job.
Maximum Wait Minutes	The Maximum minutes in the future to wait for successful completion of the Parent Job.

Output

Output	Description
Return Code	Output Return Code: S Success. W Warning. Dependency already exists (ADD action) or does not exist (DELETE action). E Error. F Fatal/Unexpected Error.

Output	Description
Return Message	Output message indicating the action applied or the reason for no action.
Result Number	Output Result Number: 1 Success. -1 Warning. Dependency already exists (ADD action) or does not exist (DELETE action). -2 Error. e.g. Due to invalid job name or job not running. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Job_Dependency

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_action	VARCHAR(10)	IN
p_parent	VARCHAR(64)	IN
p_child	VARCHAR(64)	IN
p_required	VARCHAR(1)	IN
p_look_back	INTEGER	IN
p_max_wait	INTEGER	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Job_Dependency

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
DECLARE v_return_code        varchar(1);
DECLARE v_return_msg        varchar(256);
CALL [METABASE].Ws_Job_Dependency
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'ADD', 'Daily Run', 'Daily Run Part2', 'Y', 60, 60
, v_return_code
, v_return_msg
, v_result_num
);
```

WS_JOB_RELEASE

Synopsis

Starts a job if it is in a holding or waiting state.

Description

Releases the specified job if it is in a holding or waiting state, which sets the start time to the current time so that it starts immediately. Typically, this routine is used to start a job from within another job or via a third-party scheduler (rather than a WhereScape RED Scheduler).

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Release Job Name	The name of the job to be started/released. The specified name must exactly match the job name as displayed by the WhereScape RED Scheduler. The specified job must be in a holding or waiting state in order to be released.

Output

Output	Description
Return Code	Output Return Code: S Success. N No action because the Template Job is not in a holding or waiting state. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.
Result Number	Output Result Number: 1 Success. -1 No action because the job is not in a holding or waiting state. -2 Error. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Job_Release

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_release_job	VARCHAR(64)	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Job_Release

```
-- The p_ VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
DECLARE v_return_code       varchar(1);
DECLARE v_return_msg        varchar(256);
CALL [METABASE].Ws_Job_Release
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'Daily Run'
, v_return_code
, v_return_msg
, v_result_num
);
```

WS_JOB_RESTART

Synopsis

Starts a job if it is in a failed state.

Description

Releases the specified job if it is in a failed state, which sets the start time to the current time so that it starts immediately. This routine can be executed as part of a database start-up sequence to restart each failed job that may have stopped due to an earlier database shutdown.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Restart Job Name	The name of the job to be restarted/released. The specified name must exactly match the job name, as displayed by the WhereScape RED Scheduler. The specified job must be in a failed state in order to be restarted.

Output

Output	Description
Return Code	Output Return Code: S Success. N No action because the job is not in a failed state. R No action because the job is currently running. U No action because the job is in an unusual state due to an error (result number -2). The job is classified as running but it is NOT actually running or failed so it cannot be restarted. This may occur if the scheduler has failed and the job is in a pending state. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.

Output	Description
Result Number	Output Result Number: 1 Success. -1 No action because the job is not in a failed state or it is currently running. -2 Error. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Job_Restart

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_restart_job	VARCHAR(64)	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Job_Restart

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
DECLARE v_return_code        varchar(1);
DECLARE v_return_msg        varchar(256);
CALL [METABASE].Ws_Job_Restart
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'Daily Run'
, v_return_code
, v_return_msg
, v_result_num
);
```

WS_JOB_SCHEDULE

Synopsis

Schedules a job if it is in a holding or waiting state.

Description

Schedules the specified job if it is in a holding or waiting state, which will start at the specified time. Typically, this routine is used to schedule a job from within another job.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Schedule Job Name	The name of the job to be scheduled. The specified name must exactly match the job name as displayed by the WhereScape RED Scheduler. The specified job must be in a holding or waiting state in order to be scheduled.
Scheduled Release Time	The date/time that the job is to be scheduled to be released/started.

Output

Output	Description
Return Code	Output Return Code: S Success. N No action because the job is not in a holding or waiting state. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.
Result Number	Output Result Number: 1 Success. -1 No action because the job is not in a holding or waiting state. -2 Error. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Job_Schedule

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_release_job	VARCHAR(64)	IN
p_release_time	TIMESTAMP	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Job_Schedule

```
-- The p_ VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
DECLARE v_return_code        varchar(1);
DECLARE v_return_msg        varchar(256);
CALL [METABASE].Ws_Job_Schedule
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'Daily Run', CURRENT_TIMESTAMP + INTERVAL '1' DAY
, v_return_code
, v_return_msg
, v_result_num
);
```

WS_JOB_STATUS

Synopsis

Returns the current status of a job.

Description

Returns the current status of the specified job as recorded by a WhereScape RED Scheduler. Typically, this routine is used by a third-party scheduler or a user-defined procedure/script to check on a job.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Job Sequence	The unique integer identifier of the job to return the status of. This input is optional but when it is specified, the started within and started after inputs should not be specified.
Job Name	The name of the job to return the status of. The specified name must exactly match the job name as displayed by the WhereScape RED Scheduler.
Started Within Last Minutes	The maximum minutes [0-148599] (up to ~103.1 days) in the past to look back for the job to have started. This input is optional but when it is specified, the job name must be specified and the job sequence and started after inputs should not be specified. Note: If multiple instances of the job have started in the specified time frame then the last job to start is returned (i.e. the job with the highest sequence number).
Started After Time	The date/time after which to look for the job to have started. This input is optional but when it is specified, the job name must be specified and the job sequence and started within inputs should not be specified. Note: If multiple instances of the job have started in the specified time frame then the last job to start is returned (i.e. the job with the highest sequence number).

Output

Output	Description
Return Code	Output Return Code: S - Success. N - The job exists but it was NOT started within the specified time frame. E - Error. F - Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.
Result Number	Output Result Number: 1 - Success. -1 - The job exists but it was NOT started within the specified time frame. -2 - Error. -3 - Fatal/Unexpected Error. 0 - see note below.
Simplified Job Status Code	Simplified Job Status Code: N - Not Running. R - Running. F - Failed. C - Completed. 0 - see note below.

Output	Description
Standard Job Status Code	<p>Standard Job Status Code:</p> <p>H - On Hold. The job is on hold. A held job can be edited and/or started.</p> <p>W - Waiting. The job is waiting to start (it is either waiting for the scheduled time to arrive or is waiting for an available scheduler).</p> <p>B - Blocked. The job is blocked because a previous instance of the same job is still running.</p> <p>P - Pending. This is the initial interim status of an "about to start running" job. The scheduler has identified that the job is ready to start and is preparing to run it. A job should only be pending for a brief period so if it remains pending for a prolonged period then an unexpected error has occurred.</p> <p>R - Running. The job is currently running.</p> <p>F - Failed. The job failed due to an error.</p> <p>C - Completed. The job completed successfully (but it may have warnings). A completed job cannot be restarted.</p> <p>G - Failed - Aborted. The job failed and was subsequently aborted. An aborted job cannot be restarted.</p> <p>E - Error Completion.</p> <p>0 - see note below.</p>

Output	Description
Enhanced Job Status Number	<p>Enhanced Job Status Number that returns an integer rather than the standard alphabetic code. The running and completed statuses are enhanced to distinguish errors or warnings.</p> <p>1 - On Hold. The job is on hold. A held job can be edited and/or started.</p> <p>2 - Waiting. The job is waiting to start (it is either waiting for the scheduled time to arrive or is waiting for an available scheduler).</p> <p>3 - Blocked. The job is blocked because a previous instance of the same job is still running.</p> <p>4 - Pending. This is the initial interim status of an "about to start running" job. The scheduler has identified that the job is ready to start and is preparing to run it. A job should only be pending for a brief period so if it remains pending for a prolonged period then an unexpected error has occurred.</p> <p>5 - Running. The job is currently running and no tasks have failed or produced warnings.</p> <p>6 - Running with Errors. The job is currently running but some tasks have failed. The job will ultimately fail when all the tasks that are NOT dependent on the failed tasks have finished.</p> <p>7 - Running with Warnings. The job is currently running and some tasks have produced warnings.</p> <p>8 - Failed. The job failed due to an error.</p> <p>9 - Completed. The job completed without warnings. A completed job cannot be restarted.</p> <p>10 - Completed with Warnings. The job completed with warnings. A completed job cannot be restarted.</p> <p>11 - Failed - Aborted. The job failed and it was subsequently aborted. An aborted job cannot be restarted.</p> <p>12 - Error Completion.</p>

*** Note:***

All three returned status values can also return '0' in any of the following situations:

- Illegal combination of parameters specified.
 - Unable to locate specified job sequence.
 - Unable to locate specified job name.
 - Job Not Found having started in the last *SpecifiedMinutes* minutes.
 - Job Not Found having started after *SpecifiedDateTime*.
-

Teradata Parameters: Ws_Job_Status

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_check_sequence	INTEGER	IN
p_check_job	VARCHAR(64)	IN
p_started_in_last_mi	INTEGER	IN
p_started_after_dt	TIMESTAMP	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT
p_job_status_simple	VARCHAR(1)	OUT
p_job_status_standard	VARCHAR(1)	OUT
p_job_status_enhanced	VARCHAR(2)	OUT

Teradata Examples: Ws_Job_Status

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence          integer;
DECLARE p_job_name         varchar(256);
DECLARE p_task_name        varchar(256);
DECLARE p_job_id           integer;
DECLARE p_task_id          integer;
DECLARE p_return_msg       varchar(256);
DECLARE p_status           integer;
DECLARE v_result_num       integer;
DECLARE v_return_code      varchar(1);
DECLARE v_return_msg       varchar(256);
DECLARE v_job_status_simple varchar(1);
DECLARE v_job_status_standard varchar(1);
DECLARE v_job_status_enhanced varchar(2);
CALL [METABASE].Ws_Job_Status
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, NULL, 'Daily Run', 10, NULL
, v_return_code
, v_return_msg
, v_result_num
, v_job_status_simple, v_job_status_standard, v_job_status_enhanced
);
```

WS_LOAD_CHANGE

Synopsis

Changes the Connection or Schema of a load table.

Description

Changes either the Connection or the Schema of the specified load table. Only the Connection or Schema can be changed so two calls are required to change both.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Change Property	Change either the SCHEMA or the CONNECTION of the specified load table. Separate calls must be made if both the schema and connection need to be changed.
Load Table Name	The name of the load table to be changed.
New Property Value	Either the new schema name or the new connection name.

Output

Output	Description
Return Code	Output Return Code: S Success. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.
Result Number	Output Result Number: 1 Success. -2 Error. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Load_Change

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_action	VARCHAR(64)	IN
p_table	VARCHAR(64)	IN
p_new_value	VARCHAR(255)	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Load_Change

```
-- The p_ VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence          integer;
DECLARE p_job_name         varchar(256);
DECLARE p_task_name        varchar(256);
DECLARE p_job_id           integer;
DECLARE p_task_id          integer;
DECLARE p_return_msg       varchar(256);
DECLARE p_status           integer;
DECLARE v_result_num       integer;
DECLARE v_return_code      varchar(1);
DECLARE v_return_msg       varchar(256);
CALL [METABASE].Ws_Load_Change
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'CONNECTION', 'load_customer', 'Connection2'
, v_return_code
, v_return_msg
, v_result_num
);
```

WS_MAINTAIN_INDEXES

Synopsis

Drops and/or builds database indexes that are defined in the WhereScape RED metadata.

Description

Drops and/or builds indexes for a specified table or a specified index. Only indexes that are defined in the WhereScape RED metadata are supported. Typically, this routine is used by a WhereScape RED Scheduler and RED-generated procedures to automatically maintain indexes. However, it is also valid for user-defined custom procedures/scripts to execute this routine to control when indexes are dropped and/or created.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Table Name	Table Name to process the relevant indexes of. Note: The Table Name is ignored when the optional Index Name is specified.
Index Name	Optional Index Name to only process the specified index. When NOT specified all the relevant indexes of the table are processed. Note: The Table Name is ignored when the Index Name is specified. Note: Must be specified to use the DROP or BUILD index actions.
Index Action	Action that specifies whether indexes are dropped or built and what types of indexes are applicable: DROP Drop the specified index (Index Name must be specified). DROP ALL Drop ALL the indexes of the table. PRE DROP Drop the indexes of the table that are defined as pre-drop. BUILD Build the specified index (Index Name must be specified). Otherwise, build all the indexes of the table that were pre-dropped. BUILD ALL Build ALL the indexes of the table.

Output

Output	Description
Result Number	Output Result Number: 1 Success. -1 Warning. -2 Error. -3 Fatal/Unexpected Error.

Note: Ws_Maintain_Indexes does NOT include a Return Code or Return Message like most of the WhereScape RED Callable routines but it does output a Result Number.

Teradata Parameters: Ws_Maintain_Indexes

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_table_name	VARCHAR(64)	IN
p_parameter	VARCHAR(4000)	IN
p_index_name	VARCHAR(64)	IN
p_option	VARCHAR(20)	IN
p_result	INTEGER	OUT

Teradata Examples: Ws_Maintain_Indexes

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
CALL [METABASE].Ws_Maintain_Indexes
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'load_customer', NULL, NULL, 'DROP ALL'
, v_result_num
);
```

WS_VERSION_CLEAR

Synopsis

Purges metadata versions for all objects that do not meet the specified retention criteria.

Description

Deletes metadata versions for all objects that do not meet the specified retention criteria, which can be specified as the maximum number of versions to retain per object and/or the maximum age (in days) of versions to retain. For example, it is possible to specify that a maximum of 5 versions are retained for each object and/or that versions are retained for a maximum of 90 days.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Maximum Days to Retain	The maximum age (in days) of the versions to retain. If 90 days are retained, then all the versions that are older than 90 days are purged/deleted to reduce the number of versions to the specified maximum number of versions per object. If not specified, then the retention date of each version determines whether it is deleted.
Maximum Versions per Object to Retain	The maximum number of versions to retain for each object. If 5 is specified, then the last 5 versions are retained per object regardless of the specified maximum age to retain.
Options	Currently NOT used.

Output

Output	Description
Return Code	Output Return Code: S Success. E Error. F Fatal/Unexpected Error.
Return Message	Output message indicating the action applied or the reason for no action.

Output	Description
Result Number	Output Result Number: 1 Success. -2 Error. -3 Fatal/Unexpected Error.

Teradata Parameters: Ws_Version_Clear

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_sequence	INTEGER	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_day_count	INTEGER	IN
p_keep_count	INTEGER	IN
p_options	VARCHAR(256)	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

Teradata Examples: Ws_Version_Clear

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
DECLARE v_return_code        varchar(1);
DECLARE v_return_msg        varchar(256);
CALL [METABASE].Ws_Version_Clear
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 90, 5, NULL
, v_return_code
, v_return_msg
, v_result_num
);
```

WSPARAMETERREAD

Synopsis

Returns the value and comment (for most RDBMS) of a WhereScape RED metadata Parameter.

Description

Returns the value and comment (for most RDBMS) of the specified parameter from the DSS_PARAMETER metadata table. For SQL Server, Teradata, and DB2 this routine is a PROCEDURE that returns both the parameter value and comment. However, for Oracle this routine is a FUNCTION that only returns the parameter value. For SQL Server, there is also a WsParameterReadF FUNCTION.

Typically, this routine is used by procedures to read information that is written by another process (automatically or manually via the RED Tools>Parameters menu item), which is external to the procedure.

Input

Input	Description
Parameter Name	The case-sensitive name of the WhereScape RED metadata parameter to be retrieved. The name must exactly match an existing parameter, otherwise a NULL value is returned.

Output

Output	Description
Parameter Value	The retrieved value of the parameter. Corresponds to the "Value" property that is visible and maintainable via Tools >Parameters.
Parameter Comments	The retrieved comments of the parameter. Corresponds to the "Comments" property that is visible and maintainable via Tools>Parameters.

Teradata Parameters: WsParameterRead

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_parameter	VARCHAR(64)	IN
p_value	VARCHAR(2000)	OUT
p_comment	VARCHAR(256)	OUT

Teradata Examples: WsParameterRead

```

DECLARE v_current_date          varchar(4000); -- Same length as
DSS_PARAMETER.dss_parameter_value.
DECLARE v_comment              varchar(256);
CALL
[METABASE].WsParameterRead( 'CURRENT_DATE',v_current_date,v_comment );

```

WSPARAMETERREADF

Synopsis

Returns the value of a WhereScape RED metadata Parameter **[SQL Server only]**.

Description

Returns the value of the specified parameter from the DSS_PARAMETER metadata table. This routine is a FUNCTION that is only available for SQL Server. For SQL Server, Teradata, and DB2 there is a WsParameterRead PROCEDURE that returns both the parameter value and comment. For Oracle, the WsParameterRead FUNCTION is equivalent to the SQL Server WsParameterReadF FUNCTION.

Typically, this routine is used by procedures to read information that is written by another process (automatically or manually via the RED Tools - Parameters menu item), which is external to the procedure.

Input

Input	Description
Parameter Name	The name of the WhereScape RED metadata parameter to be retrieved. The case-sensitive name must exactly match an existing parameter, otherwise a NULL value is returned.

Output

Output	Description
Parameter Value	The value of the parameter. Corresponds to the "Value" property that is visible and maintainable via Tools - Parameters.

WSPARAMETERREADG

Synopsis

Returns the value of a "global" WhereScape RED metadata Parameter that relates to a load table.

Description

Returns the value of an internal parameter that is defined and populated by WhereScape RED, which is available to a procedure that is currently processing a load table.

The supported parameters are \$\$TABLE_NAME and \$\$SOURCE_TABLE.

Input

Input	Description
Global Parameter Name	The supported global parameters are: \$\$TABLE_NAME returns the Load Table Name that the procedure is executing against. Only available for load tables. \$\$SOURCE_TABLE returns the maximum value of the Source Table property from the columns of the Load Table that the procedure is executing against. Only available for load tables. Note: Typically, a Load Table has a single Source Table but if it has multiple sources then the maximum (alphabetically) Source Table Name will be returned.
Job Identifier	Unique identifier of the held or scheduled job that the running job is a specific instance of. When invoked from a WhereScape RED Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the WhereScape RED Scheduler, it is recommended to use 0 (zero).
Task or Object Identifier	Unique identifier of the running task (of a running job) that executed the routine. When invoked from a WhereScape RED Scheduler, the routine will be passed the parameter argument. When invoked manually or externally to the WhereScape RED Scheduler, it should be the object key .

Output

Output	Description
Result Table Name	The requested Load Table Name or Source Table Name.

Teradata Parameters: WsParameterReadG

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_parameter	VARCHAR(64)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_value	VARCHAR(2000)	OUT

Teradata Examples: WsParameterReadG

```
-- The p_ VARIABLES are normally PARAMETERS in a RED-generated  
Procedure.  
DECLARE p_job_id          integer;  
DECLARE p_task_id        integer;  
DECLARE v_source_name    varchar(256);  
CALL [METABASE].WsParameterReadG('$$SOURCE_TABLE', p_job_id, p_task_id,  
v_source_name);
```


WSPARAMETERWRITE

Synopsis

Updates the value and comment of a WhereScape RED metadata Parameter or creates it.

Description

Updates the value and comment of the specified parameter in the DSS_PARAMETER metadata table. If the specified parameter is not found then it is added.

Typically, this routine is used by procedures to write information that is read by another process (automatically or manually via the RED Tools>Parameters menu item), which is external to the procedure.

Input

Input	Description
Parameter Name	The case-sensitive name of the WhereScape RED metadata parameter to be updated or added.
Parameter Value	The new value of the parameter to be assigned. Corresponds to the "Value" property that is visible and maintainable via Tools>Parameters.
Parameter Comments	The new comments of the parameter to be assigned. Corresponds to the "Comments" property that is visible and maintainable via Tools>Parameters. The parameter comments will not be modified if a NULL value is specified.

Teradata Parameters: WsParameterWrite

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_parameter	VARCHAR(64)	IN
p_value	VARCHAR(2000)	IN
p_comment	VARCHAR(256)	IN

Teradata Examples: WsParameterWrite

```
CALL [METABASE].WsParameterWrite('LAST_INVOICE_ID', '123456', 'The last invoice ID loaded');
```

WSWRKAUDIT

Synopsis

Records a message in the Audit Log.

Description

Adds the specified message to the WS_WRK_AUDIT_LOG workflow metadata table, which is referred to as the Audit Log or Audit Trail. A variety of message types are supported such as Information, Warning, and Error that are included in the corresponding message type counts for the task and job. Audit Log messages are accessible via the "Scheduler" tab/window and/or the WS_ADMIN_V_AUDIT view of the WS_WRK_AUDIT_LOG table.

NOTE: Both the Audit Log and Error/Detail Log support similar information and in user-defined custom procedures either or both logs can be used. However, in RED-generated procedures/scripts the Audit Log is used for higher-level or summary messages while the Error/Detail Log is used for more detailed supporting information.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873) Note: Refer to the RDBMS-specific parameters for the relative positions (they are NOT declared as the first parameters)
Audit Message Type Code	Audit Message Type Code: B Beginning of a Job or Task. I Information. S Success. W Warning. E Error. F Fatal Error.
Audit Message Text	Custom message text to be recorded in the WhereScape RED Audit Log.
RDBMS Code	RDBMS-specific message code. e.g. The Oracle special variable SQLCODE. It is optional but recommended to populate this when an error occurs.
RDBMS Message	RDBMS-specific message. e.g. The Oracle special variable SQLERRM. It is optional but recommended to populate this when an error occurs.

Output

Output	Description
Result Number	Output Result Number: 1 Success. -3 Error.

Teradata Parameters: WsWrkAudit

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_status_code	VARCHAR(1)	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_sequence	INTEGER	IN
p_message	VARCHAR(255)	IN
p_db_code	VARCHAR(10)	IN
p_db_msg	VARCHAR(255)	IN
p_task_key	INTEGER	IN
p_job_key	INTEGER	IN

Teradata Examples: WsWrkAudit

```
-- The p_ VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence          integer;
DECLARE p_job_name         varchar(256);
DECLARE p_task_name        varchar(256);
DECLARE p_job_id           integer;
DECLARE p_task_id          integer;
DECLARE p_return_msg       varchar(256);
DECLARE p_status           integer;
CALL [METABASE].WsWrkAudit
( 'I', p_job_name, p_task_name, p_sequence
, 'The task has started.'
, NULL
, NULL
, p_task_id
, p_job_id
);
```

WSWRKAUDITBULK

Synopsis

Records multiple messages in the Audit Log.

Description

Adds the specified multiple messages to the WS_WRK_AUDIT_LOG workflow metadata table, which is referred to as the Audit Log or Audit Trail. A variety of message types are supported, such as Information, Warning, and Error that are included in the corresponding message type counts for the task and job. Audit Log messages are accessible via the "Scheduler" tab/window and/or the WS_ADMIN_V_AUDIT view of the WS_WRK_AUDIT_LOG table.

NOTE: Both the Audit Log and Error/Detail Log support similar information and in user-defined custom procedures, either or both logs can be used. However, in RED-generated procedures/scripts the Audit Log is used for higher-level or summary messages while the Error/Detail Log is used for more detailed supporting information.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873) Note: Refer to the RDBMS-specific parameters for the relative positions (they are NOT declared as the first parameters)
Audit Message Type Code	Audit Message Type Code: B Beginning of a Job or Task. I Information. S Success. W Warning. E Error. F Fatal Error.
Audit Message(s) Text	Custom message(s) text to be recorded in the WhereScape RED Audit Log. Multiple messages can be specified but each is limited to 256 characters. Each message must be separated by either a new-line (ASCII 10) or tilde (~) character. e.g. Message1~Message2~Message3 will create 3 messages.
RDBMS Code	RDBMS-specific message code. e.g. The Oracle special variable SQLCODE. It is optional but recommended to populate this when an error occurs.
RDBMS Message	RDBMS-specific message. e.g. The Oracle special variable SQLERRM. It is optional but recommended to populate this when an error occurs.

Output

Output	Description
Result Number	<p>Note: Not provided for all RDBMS.</p> <p>Output Result Number:</p> <p>1 Success.</p> <p>-2 Error</p> <p>-3 Fatal/Unexpected Error.</p>

Teradata Parameters: WsWrkAuditBulk

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_status_code	VARCHAR(64)	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_sequence	INTEGER	IN
p_message	VARCHAR(61440)	IN
p_db_code	VARCHAR(10)	IN
p_db_msg	VARCHAR(256)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_result	INTEGER	OUT

Teradata Examples: WsWrkAuditBulk

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_result_num        integer;
CALL [METABASE].WsWrkAuditBulk
( 'I', p_job_name, p_task_name, p_sequence
, 'Message1~Message2~Message3'
, NULL
, NULL
, p_job_id  --### NOTE order.
, p_task_id --### NOTE order.
, v_result_num
);
```

WSWRKERROR

Synopsis

Records a message in the Error/Detail Log.

Description

Adds the specified message to the WS_WRK_ERROR_LOG workflow metadata table, which is referred to as the Error Log or Detail Log. A variety of message types are supported such as Information, Warning, and Error that are included in the "detail" message counts for the task and job (viewable via the "Scheduler" tab/window). Error/Detail Log messages are accessible via the "Scheduler" tab/window and/or the WS_ADMIN_V_ERROR view of the WS_WRK_ERROR_LOG table.

NOTE: Both the Audit Log and Error/Detail Log support similar information and in user-defined custom procedures either or both logs can be used. However, in RED-generated procedures/scripts the Audit Log is used for higher-level or summary messages while the Error/Detail Log is used for more detailed supporting information.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873) Note: Refer to the RDBMS-specific parameters for the relative positions (they are NOT declared as the first parameters)
Error/Detail Message Type Code	Error/Detail Message Type Code: E Error. I Information. W Warning.
Error/Detail Message Text	Custom message text to be recorded in the WhereScape RED Error/Detail Log.
RDBMS Code	RDBMS-specific message code. e.g. The Oracle special variable SQLCODE. It is optional but recommended to populate this when an error occurs.
RDBMS Message	RDBMS-specific message. e.g. The Oracle special variable SQLERRM. It is optional but recommended to populate this when an error occurs.
Custom Message Type Code	Custom Message Type Code. For custom usage and has no meaning within the WhereScape RED metadata.

Output

Output	Description
Result Number	Output Result Number: 1 Success. -3 Error.

Teradata Parameters: WsWrkError

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_status_code	VARCHAR(1)	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_sequence	INTEGER	IN
p_message	VARCHAR(255)	IN
p_db_code	VARCHAR(10)	IN
p_db_msg	VARCHAR(255)	IN
p_task_key	INTEGER	IN
p_job_key	INTEGER	IN
p_msg_type	VARCHAR(10)	IN

Teradata Examples: WsWrkError

```
-- The p_ VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence          integer;
DECLARE p_job_name         varchar(256);
DECLARE p_task_name       varchar(256);
DECLARE p_job_id          integer;
DECLARE p_task_id         integer;
DECLARE p_return_msg      varchar(256);
DECLARE p_status          integer;
CALL [METABASE].WsWrkError
( 'I', p_job_name, p_task_name, p_sequence
, 'This is an INFO message in the Error/Detail Log.'
, NULL
, NULL
, p_task_id
, p_job_id
, NULL
);
```

WSWRKERRORBULK

Synopsis

Records multiple messages in the Error/Detail Log.

Description

Adds the specified multiple messages to the WS_WRK_ERROR_LOG workflow metadata table, which is referred to as the Error Log or Detail Log. A variety of message types are supported such as Information, Warning, and Error that are included in the "detail" message counts for the task and job (viewable via the "Scheduler" tab/window). Error/Detail Log messages are accessible via the "Scheduler" tab/window and/or the WS_ADMIN_V_ERROR view of the WS_WRK_ERROR_LOG table.

NOTE: Both the Audit Log and Error/Detail Log support similar information and in user-defined custom procedures, either or both logs can be used. However, in RED-generated procedures/scripts the Audit Log is used for higher-level or summary messages while the Error/Detail Log is used for more detailed supporting information.

Input

Input	Description
Common Input	Includes all 5 inputs of the <i>Callable Routines Common Input</i> (on page 873) Note: Refer to the RDBMS-specific parameters for the relative positions (they are NOT declared as the first parameters)
Error/Detail Message Type Code	Error/Detail Message Type Code: E Error. I Information. W Warning.
Error/Detail Message(s) Text	Custom message(s) text to be recorded in the WhereScape RED Error/Detail Log. Multiple messages can be specified but each is limited to 256 characters. Each message must be separated by either a new-line (ASCII 10) or tilde (~) character. e.g. Message1~Message2~Message3 will create 3 messages.
RDBMS Code	RDBMS-specific message code. e.g. The Oracle special variable SQLCODE. It is optional but recommended to populate this when an error occurs.
RDBMS Message	RDBMS-specific message. e.g. The Oracle special variable SQLERRM. It is optional but recommended to populate this when an error occurs.
Custom Message Type Code	Custom Message Type Code. For custom usage and has no meaning within the WhereScape RED metadata.

Output

Output	Description
Result Number	<p>Note: NOT provided for all RDBMS.</p> <p>Output Result Number:</p> <p>1 Success.</p> <p>-2 Error.</p> <p>-3 Fatal/Unexpected Error.</p>

Teradata Parameters: WsWrkErrorBulk

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_status_code	VARCHAR(64)	IN
p_job_name	VARCHAR(64)	IN
p_task_name	VARCHAR(64)	IN
p_sequence	INTEGER	IN
p_message	VARCHAR(61440)	IN
p_db_code	VARCHAR(10)	IN
p_db_msg	VARCHAR(256)	IN
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_msg_type	VARCHAR(10)	IN
p_result	INTEGER	OUT

Teradata Examples: WsWrkErrorBulk

```
-- The p_VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence                integer;
DECLARE p_job_name                varchar(256);
DECLARE p_task_name              varchar(256);
DECLARE p_job_id                 integer;
DECLARE p_task_id                integer;
DECLARE p_return_msg             varchar(256);
DECLARE p_status                 integer;
DECLARE v_result_num             integer;
CALL [METABASE].WsWrkErrorBulk
( 'I', p_job_name, p_task_name, p_sequence
, 'Message1~Message2~Message3'
, NULL
, NULL
, p_job_id  --### NOTE order.
, p_task_id --### NOTE order.
, NULL
, v_result_num
);
```

WSWRKTASK

Synopsis

Updates row counts for a task in the Task Log.

Description

Updates row counts for the specified task in the Task Log. Task Log messages (and row counts) are accessible via the "Scheduler" tab/window and/or the WS_ADMIN_V_TASK view of the WS_WRK_TASK_RUN and WS_WRK_TASK_LOG tables.

This routine is intended to be executed by a task of a job since it requires a valid job, task, and job sequence number that are provided by a WhereScape RED Scheduler.

Input

Input	Description
Common Input	Includes 3 inputs of the <i>Callable Routines Common Input</i> (on page 873)
Inserted Row Count	The number of rows inserted by the task.
Updated Row Count	The number of rows updated by the task.
Replaced Row Count	The number of rows replaced by the task.
Deleted Row Count	The number of rows deleted by the task.
Discarded Row Count	The number of rows discarded by the task.
Rejected Row Count	The number of rows rejected by the task.
Error Row Count	The number of rows with an error that were failed by the task.

Output

Output	Description
Result Number	Output Result Number: 0 Success. -1 Warning. -3 Fatal/Unexpected Error.

Teradata Parameters: WsWrkTask

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_job_key	INTEGER	IN
p_task_key	INTEGER	IN
p_sequence	INTEGER	IN
p_inserted	INTEGER	IN
p_updated	INTEGER	IN
p_replaced	INTEGER	IN
p_deleted	INTEGER	IN
p_discarded	INTEGER	IN
p_rejected	INTEGER	IN
p_errored	INTEGER	IN

Teradata Examples: WsWrkTask

```
-- The p_ VARIABLES are normally PARAMETERS in a RED-generated
Procedure.
DECLARE p_sequence           integer;
DECLARE p_job_name          varchar(256);
DECLARE p_task_name         varchar(256);
DECLARE p_job_id            integer;
DECLARE p_task_id           integer;
DECLARE p_return_msg        varchar(256);
DECLARE p_status            integer;
DECLARE v_insert_count      integer;
DECLARE v_update_count      integer;
CALL [METABASE].WsWrkTask
( p_job_id, p_task_id, p_sequence
, v_insert_count, v_update_count, 0, 0, 0, 0, 0
);
```

CHAPTER 35

WS_ADMIN_V VIEWS

Admin views provide a means of interacting with the WhereScape RED metadata from within your chosen reporting tools.

The following admin views are available:

Name	Description
ws_admin_v_audit	Created from the ws_wrk_audit_log table
ws_admin_v_error	Created from the ws_wrk_error_log table
ws_admin_v_sched	Created from the ws_wrk_job_log table
ws_admin_v_task	Created from the ws_wrk_task_run and ws_wrk_task_log tables

IN THIS CHAPTER

Ws_admin_v_audit	942
Ws_admin_v_error.....	943
Ws_admin_v_sched	944
Ws_admin_v_task	947

WS_ADMIN_V_AUDIT

This Audit view is created using columns from the ws_wrk_audit_log table.

Columns

The following columns are created:

Column	Description
wa_time_stamp	the date or time at which this view was created
wa_sequence	See <i>Callable Routines Common Input</i> (on page 873)
wa_job	See <i>Callable Routines Common Input</i> (on page 873)
wa_task	See <i>Callable Routines Common Input</i> (on page 873)
wa_status	See <i>Callable Routines Common Input</i> (on page 873)
wa_message	the message associated with this audit log
wa_db_msg_desc	the database message associated with this audit log

WS_ADMIN_V_ERROR

This Error view is created using columns from the ws_wrk_error_log table.

Columns

The following columns are created:

Column	Description
wd_time_stamp	the date or time at which this view was created
wd_sequence	See <i>Callable Routines Common Input</i> (on page 873)
wd_job	See <i>Callable Routines Common Input</i> (on page 873)
wd_task	See <i>Callable Routines Common Input</i> (on page 873)
wd_status	See <i>Callable Routines Common Input</i> (on page 873)
wd_message	the message associated with this audit log
wd_db_msg_desc	the database message associated with this audit log

WS_ADMIN_V_SCHED

This Scheduled Job view is created from the ws_wrk_job_log table.

Columns

The following columns are created:

Column	Description
type	"Waiting"
job_name	the name of the job
status	a character value indicating the job status H - on hold R - running P - pending W - waiting C - completed B - blocked F - failed G - failed - aborted E - error completion else unknown
sequence	sequence number of the job
started_scheduled	
completed	date completed
hours_elapsed	hours elapsed since job started
minutes_elapsed	minutes elapsed since job started
okay	okay count
info	info count
warn	warning count
detail	detail count
error	error count

WS_ADMIN_V_TASK

This Task view is created from the ws_wrk_task_run and ws_wrk_task_log tables.

Columns

The following columns are created:

Column	Description
result	
task_name	the name of the task
status	a character value indicating the task status H - on hold R - running P - pending W - waiting C - completed B - blocked F - failed G - failed - aborted E - error completion else unknown
sequence	sequence number of the task
started	
completed	
hours_elapsed	hours elapsed since the task started
minutes_elapsed	minutes elapsed since the task started
info	info count
warn	warning count
detail	detail count
inserted	record inserted
updated	record updated
replaced	record replaced
deleted	record deleted
discarded	record discarded
rejected	record rejected
errored	record errored

CHAPTER 36

RETROFITTING

WhereScape RED includes an advanced retrofit capability that can be used to:

- 1 Migrate an existing data warehouse from one relational database to another (known as fork-lifting).
- 2 Load a data model from a modeling tool.

Retrofitting is achieved using the **Retro** object type in WhereScape RED and the **Retrofit tables wizard**.

For information on migrating an existing data warehouse, see *Migrating the Data Warehouse Database Platform* (see "*OLAP Retrofitting an OLAP Object*" on page 588, on page 948).

For information on importing a data model, see *Importing a Data Model* (on page 957).

IN THIS CHAPTER

Migrating the Data Warehouse Database Platform	948
Importing a Data Model.....	957
Re-Targeting Source Tables.....	964
Retro Column Properties	966

MIGRATING THE DATA WAREHOUSE DATABASE PLATFORM

WhereScape RED has an advanced retrofitting wizard for migrating an existing data warehouse from one relational database to another.

The process to migrate an existing data warehouse is:

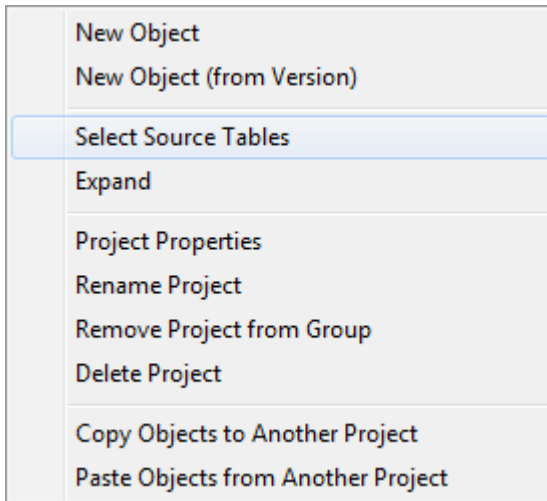
- 1 Create a connection object to the existing warehouse database.
- 2 Create Retro objects based on the source tables in the existing warehouse database.
- 3 Set the Retro objects as Retro Copy type objects.
- 4 Run a Scheduler task to build the Retro Copy objects from the source tables.
- 5 Set the Retro objects back to Retro (Retro Definition) type objects.
- 6 Convert the Retro objects to the Target Object types.

The steps to use this wizard are:

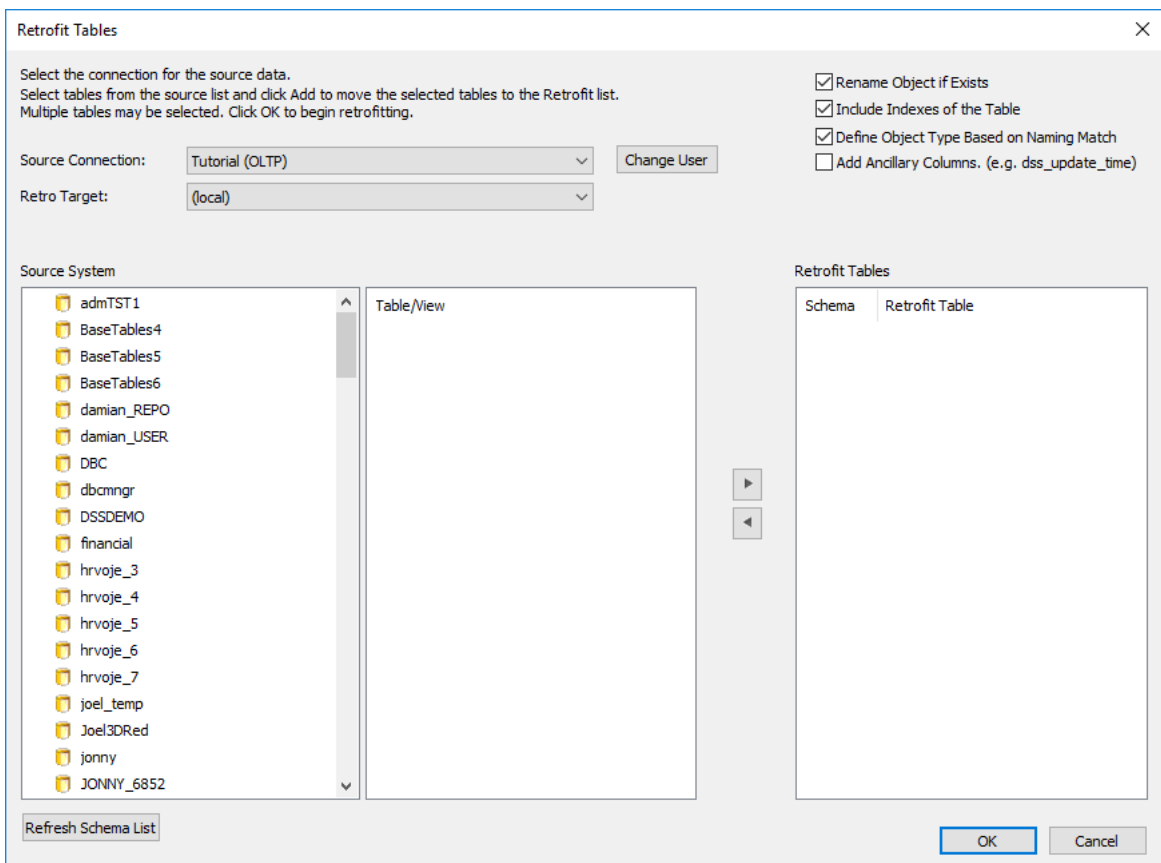
- 1 Create a connection object for the old data warehouse database, populating the following fields:
 - Connection Name
 - Connection Type => **ODBC**
 - ODBC Source
 - Work Directory
 - Extract user name
 - Extract passwordOR
 - Teradata Wallet User ID / Teradata Wallet String

Note: The extract user/ Teradata Wallet user must be able to select from all tables to be migrated.

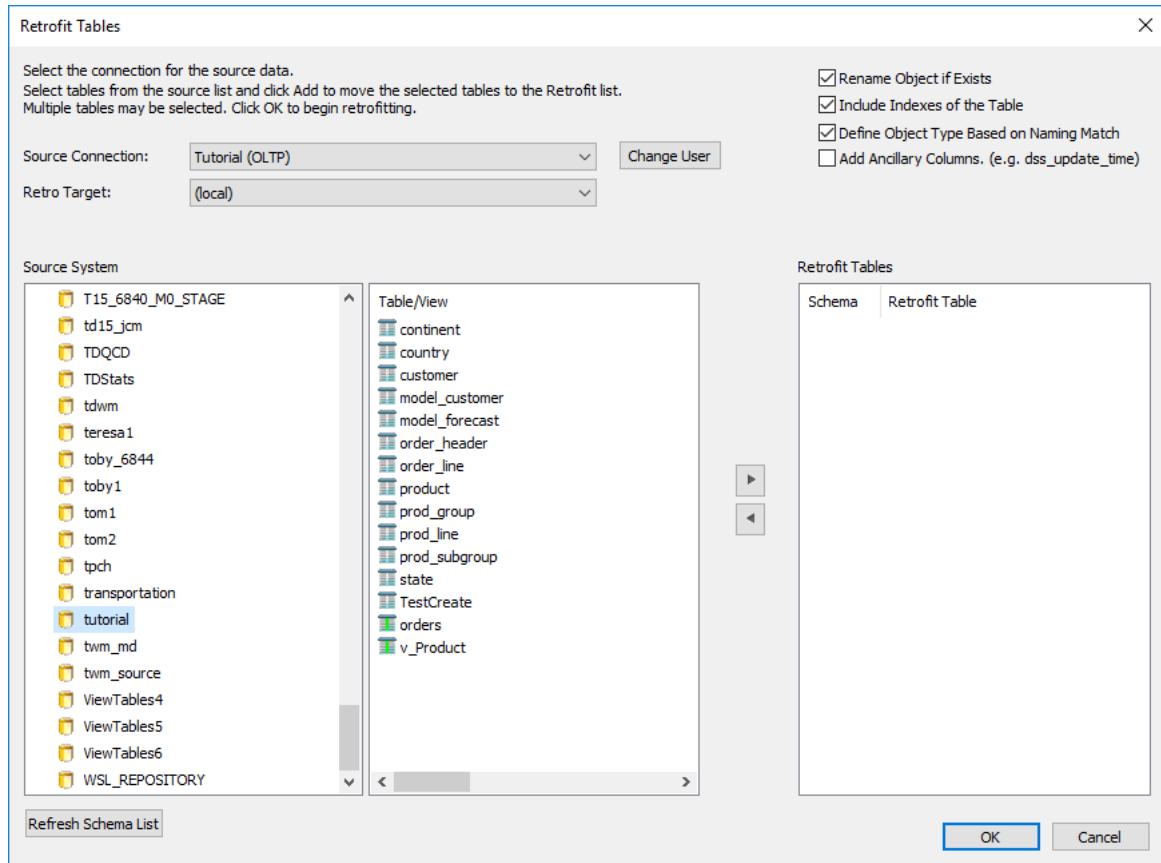
- 2 Ensure all naming standards in **Tools>Options** are set to match the objects being retrofitted. This saves work later.
- 3 Ensure **Enable Retro** is selected in the **Tools>Options>Object Types** menu.
- 4 Right-click on the **Retro object group** in the object tree in the left pane and then click **Select Source Tables**.



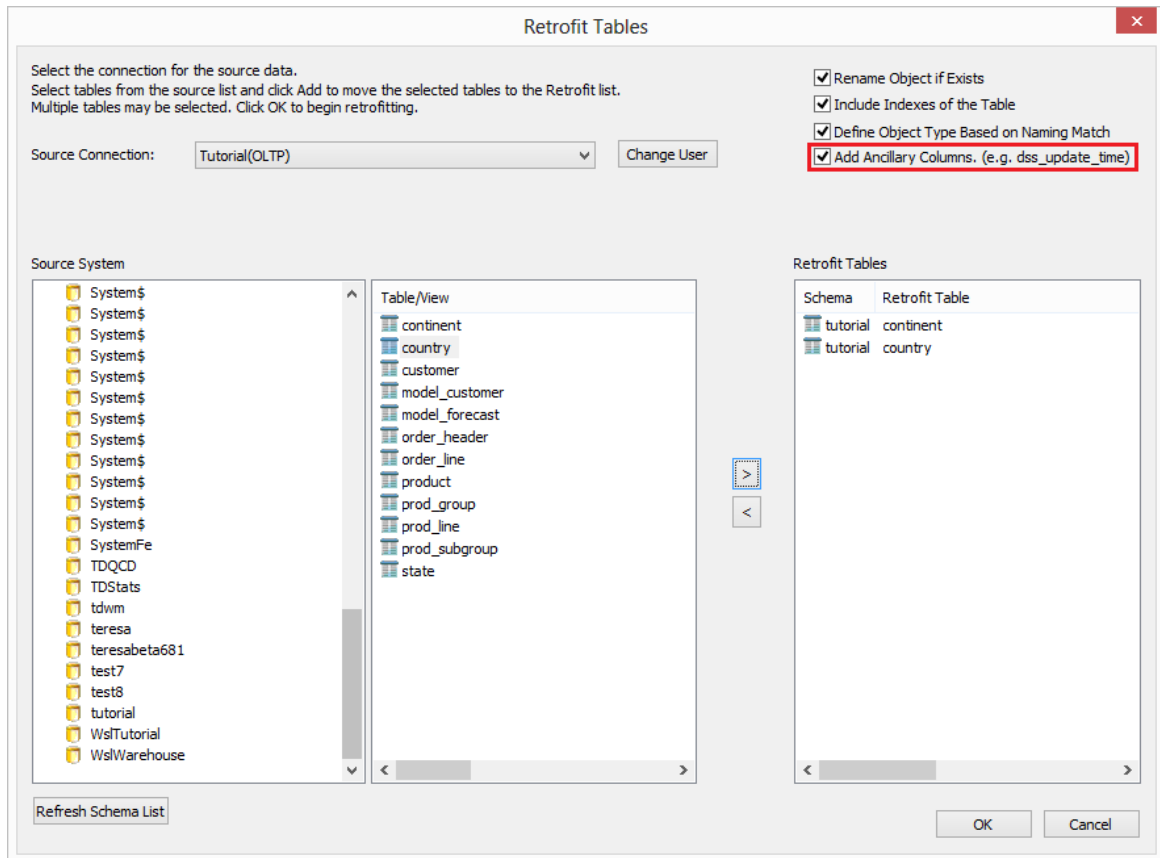
- 5 The Retrofit Tables dialog appears. In the **Source Connection** drop-down list choose the connection set up in step 1. A list of **databases** appears in the left pane.



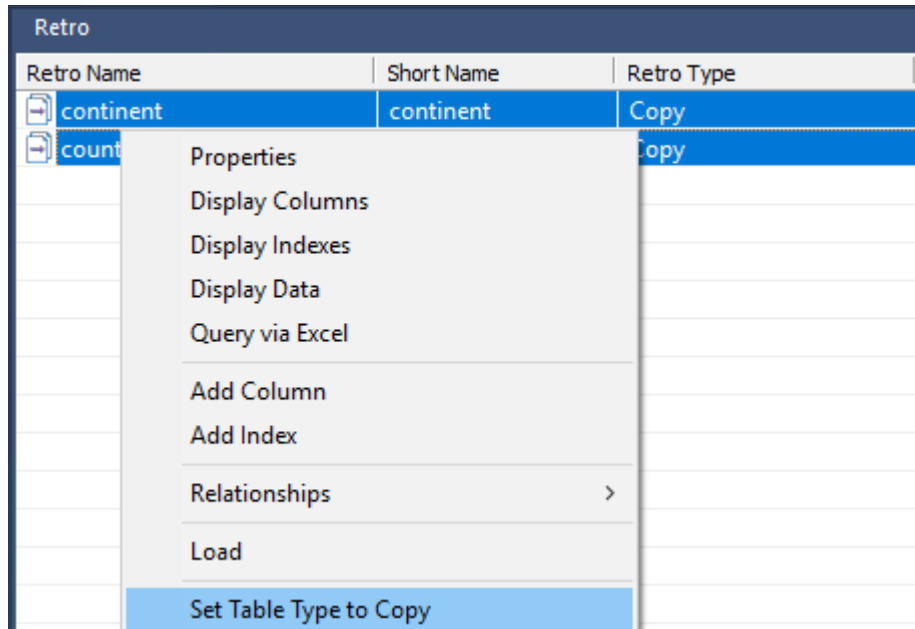
- Double-click on the database/user/schema in the left pane. A list of **tables** in the database is displayed in the middle pane.




- 7 Select all the required tables from the middle pane list and click > to move them to the right pane. Then click the **Add Ancillary Columns (e.g. dss_update_time)** check box and then click **OK**.



- 8 WhereScape RED acquires the metadata for the tables being migrated and creates a new WhereScape RED Retro object for each.
- 9 Double-click on the **Retro object group** in the left pane. Select all Retros in the middle pane. Right-click and select **Set Table Type to Copy**. This allows the data in the legacy data warehouse to be copied across to the new data warehouse.



- 10 Click on the **Scheduler** button on the toolbar.



- 11 Create a new job to run straight away.

Note: A scheduler must be running on the data warehouse connection for this job to complete, please refer to section 17. **Scheduler Installation and Configuration** of the RED Setup Administrator Guide.

Job Definition
✕

Job Name:

Description:

Frequency: Once and Hold ▾

Start Date: Wednesday, March 2, 2016 ▾

Start Time: 7:13:00 PM ▾

Maximum Threads:

Scheduler: Windows Only ▾

Dependent On: Parent job Fail Look back (minutes) Maximum wait (minutes) Add Parent Job Remove Parent

Logs Retained: This field lets you set the number of logs that are retained for this job before an automatic delete and archive occurs. 0 = keep all logs (default action)

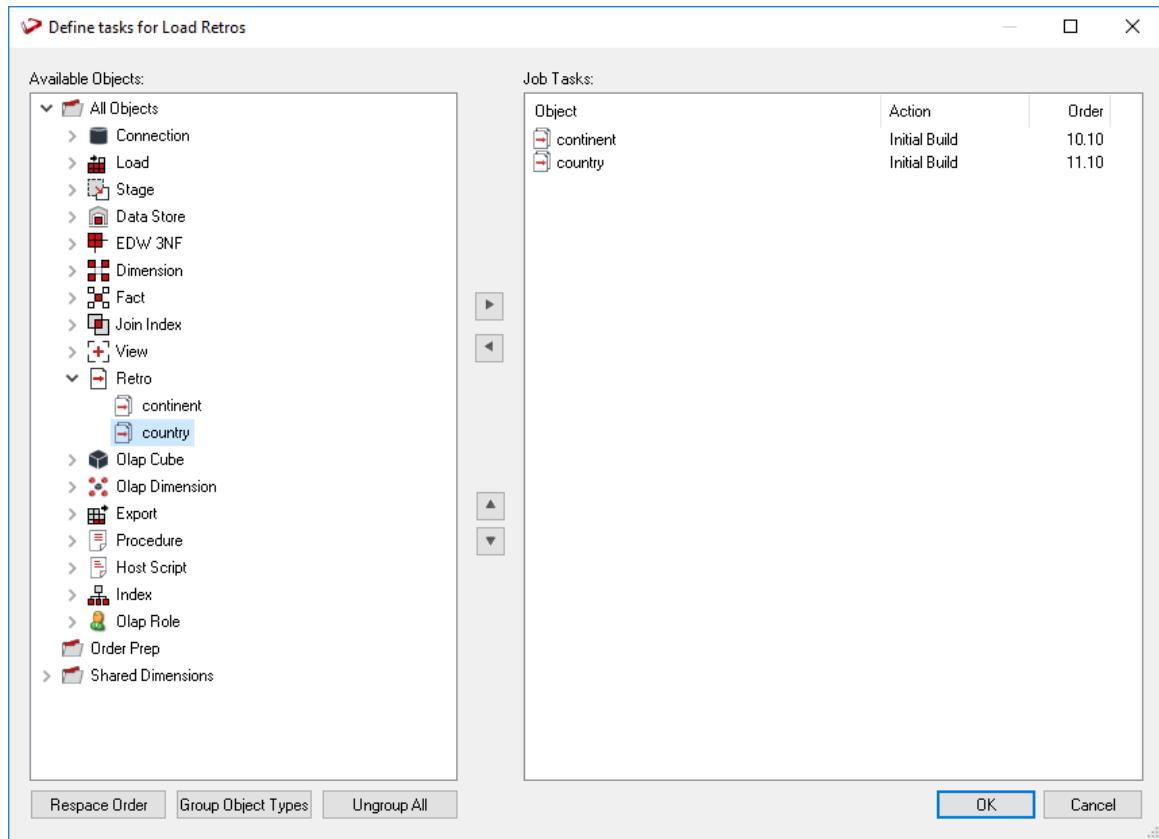
The following two fields are optional. They are executed after the job completes and therefore need to reflect the scheduler environment. (i.e. Unix or Windows). The special variables \$JOB_KEY\$, \$JOB_SEQ\$ and \$JOB_NAME\$ can be used to return the associated values. The Success command will be executed if a successful completion, the failure command will be executed if a job fails to complete:

Success Command:

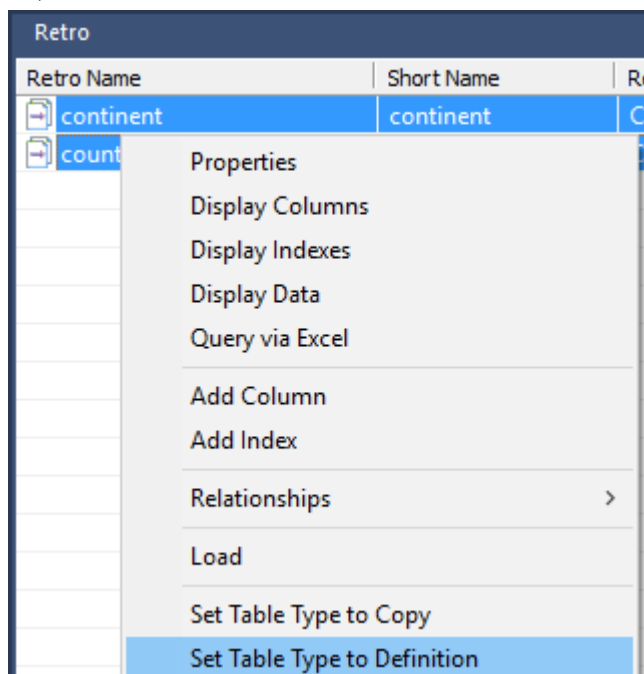
Failure Command:

OK
Cancel

12 Add all Retro objects created in steps (3) to (9) and click on **Group Object Types**.

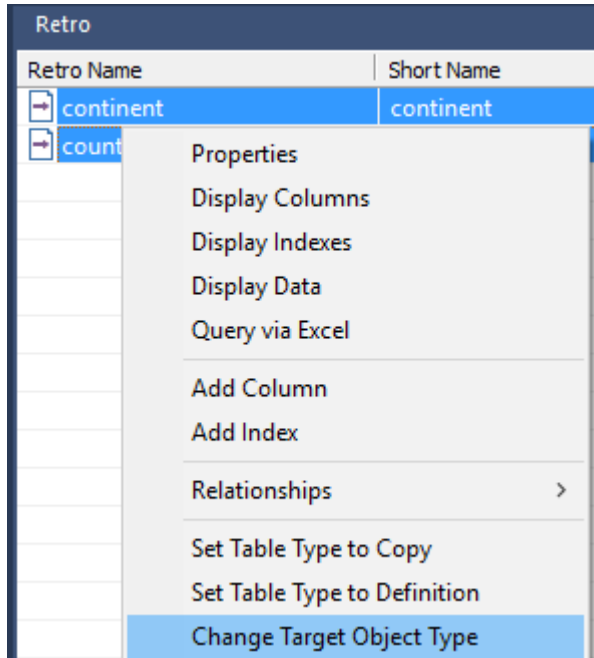


13 Once the job has completed, return to the WhereScape RED builder. Double-click on the **Retro Object group**. Select all objects in the middle pane and from the right-click menu select **Set Table Type to Definition**. This indicates the data has been copied into the Retro objects and the Retros can now be converted to the target objects.

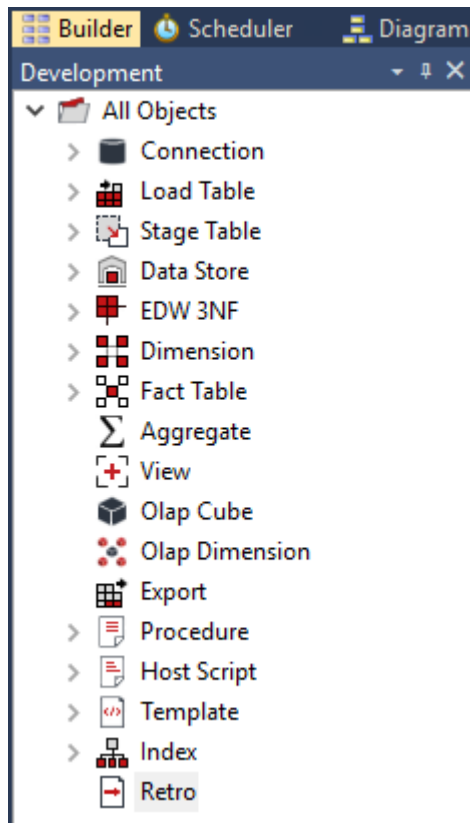


- 14 In the middle pane, select all objects. Right-click and select **Convert to Target Object**. WhereScape RED now converts the **Retro** objects to the appropriate object types.

Note: If the appropriate Target Object Type has not been set for one or more Retro objects; in the right click menu select **Change Target Object Type** and select the correct Object Type.



- 15 There are no longer any Retro objects. They have been converted to Load, Stage, Dimension or Fact objects.



- 16 Change the source table and source column values on all of the retrofitted objects using either the Re-target source table dialog, or by editing column properties. See **Re-target source tables** (see "**Re-Targeting Source Tables**" on page 964) for more information.
- 17 Convert the old data warehouses code to WhereScape RED procedure in the new data warehouse database. See **Integrate Procedures** (see "**Integrating, Procedures**" on page 981) for more information.
- 18 If necessary, create new connections to be used with any migrated load tables. Attach a connection to each load table. See **Loading Data** (on page 215) for more information.

IMPORTING A DATA MODEL

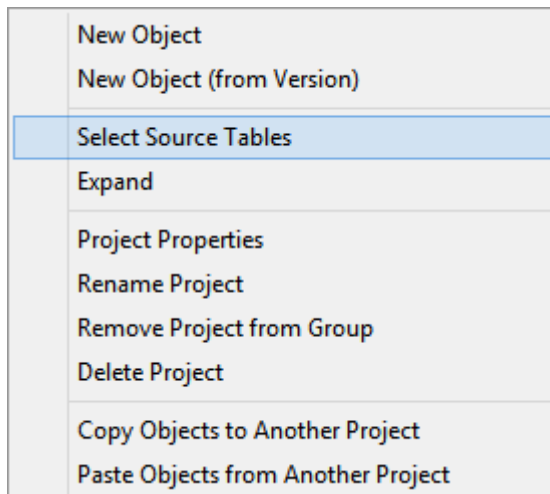
WhereScape RED provides functionality for importing data models from modeling tools.

The process to import a model is:

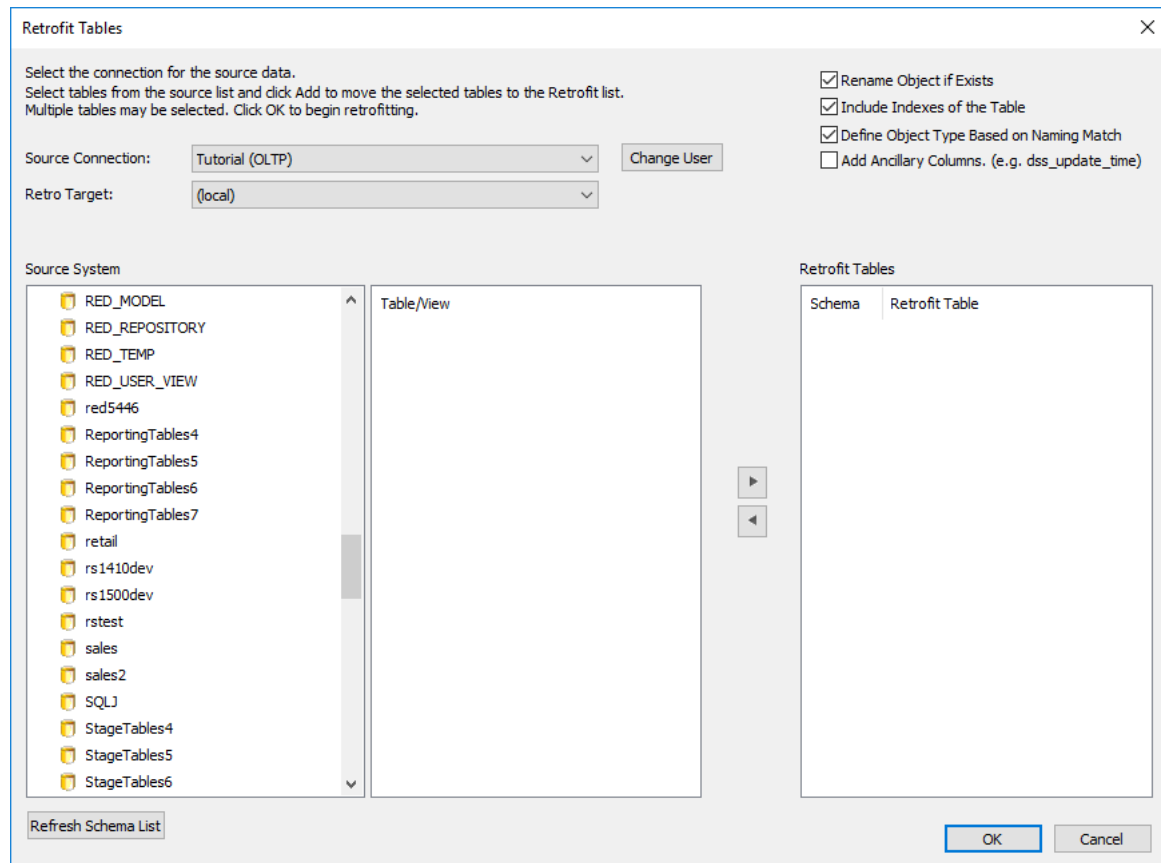
- 1 Create the physical data model in the modeling tool.
- 2 Generate DDL for the physical model in the modeling tool.
- 3 Run the DDL in the data warehouse database to create empty versions of the model tables.
- 4 Retrofit the tables in the dummy database into the WhereScape RED metadata as Retro objects.
- 5 Convert the Retro objects to Dimensions and Facts.

The following instructions outline steps 4 and 5 above:

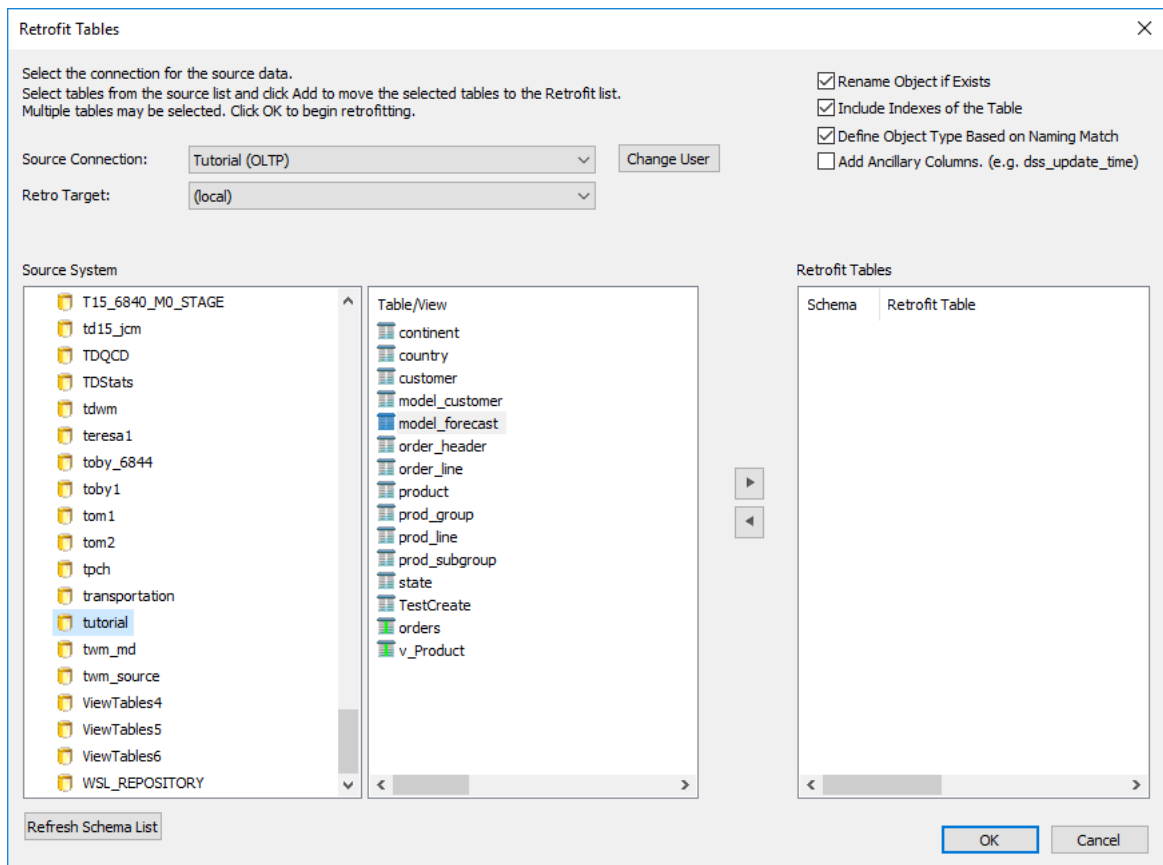
- 1 Right-click on the **Retro object group** in the object tree in the left pane and then click **Select Source Tables**.



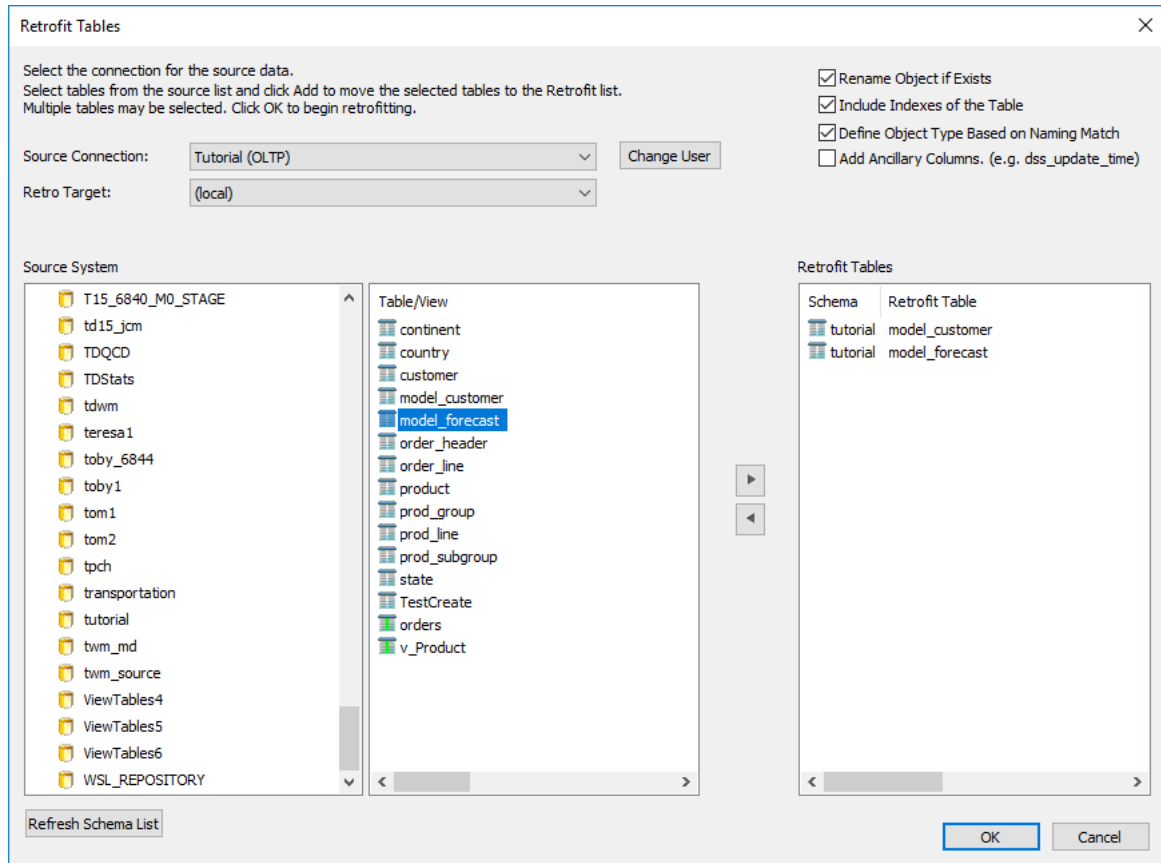
- The Retrofit tables dialog is displayed. In the **Source Connection** drop-down list choose the Data Warehouse connection. A list of **databases** appears in the left pane (your list will be different).



- 3 Double click on the required database in the left pane list. A list of **tables** in the database is displayed in the middle pane.



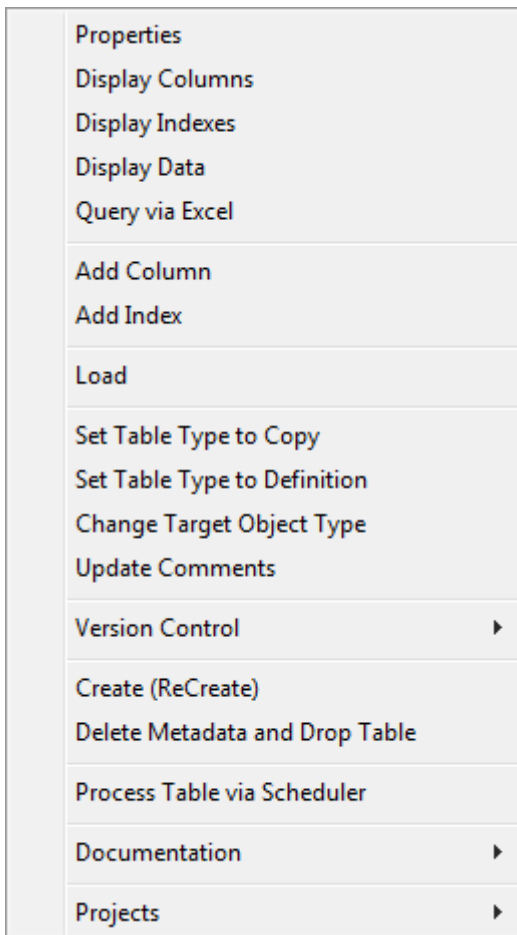
- Click on the required tables in the middle pane list and click > to move them to the right pane. Then click the **Add Ancillary columns (e.g. dss_update_time)** check box. Click **OK**.



- Double click on the **Retro object group** in the object tree in the left pane. You should see something like this.

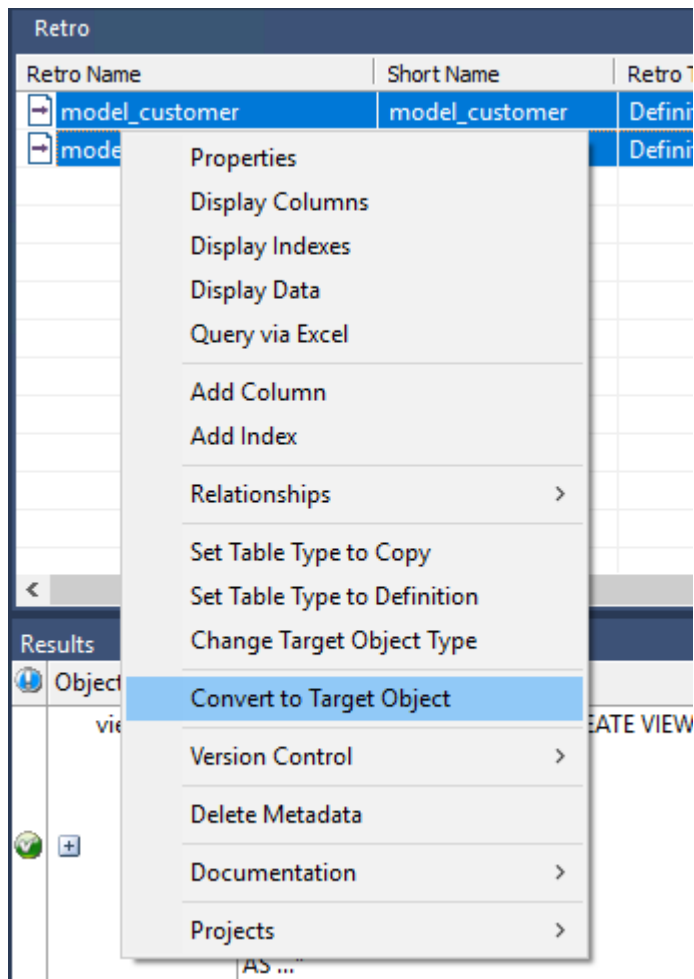
Retro	
Retro Name	Short Name
model_customer	model_customer
model_forecast	model_forecast

- 6 In the middle pane, select the tables. Right-click and select **Set table type to Definition**.

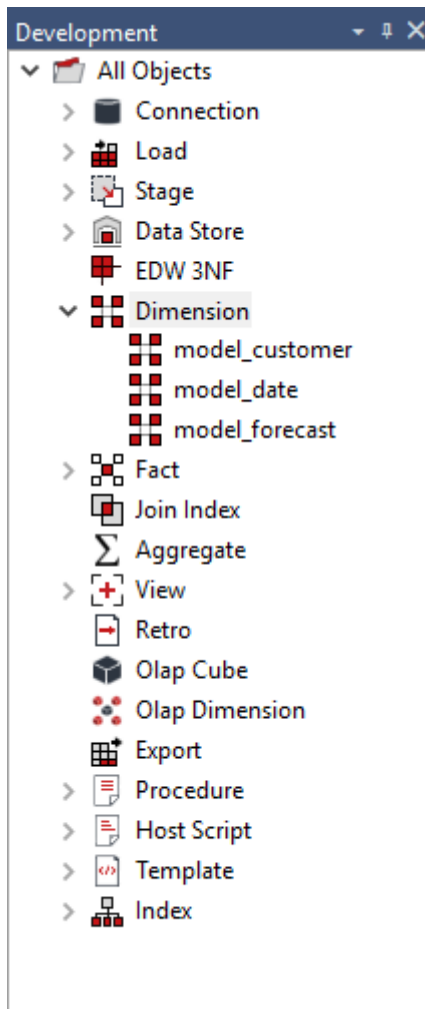


- 7 In the middle pane, select the tables. Right-click and select **Convert to Target Object**.

Note: If the appropriate Target Object Type has not been set for one or more Retro objects; in the right click menu select **Change Target Object Type** and select the correct Object Type.



- 8 The new tables have been imported.



- 9 At this stage you have created the table metadata only. To create the tables in the data warehouse, double click on the object group in the object tree in the left pane. In the middle pane highlight the tables, then right-click and select **Create (ReCreate)**.

RE-TARGETING SOURCE TABLES

Objects that have been retrofitted into the WhereScape RED metadata have themselves as their source table:

Column Name	Display Name	Data Type	Source Table	Source Column
dim_product_key	dim product key	integer ident...	dim_product	dim_product_key
description	description	varchar(24)	dim_product	description
code	code	numeric(6)	dim_product	code
prod_line	prod line	varchar(24)	dim_product	prod_line
prod_group	prod group	varchar(24)	dim_product	prod_group
subgroup	subgroup	varchar(24)	dim_product	subgroup
dss_update_time	dss update time	datetime		

They can be re-targeted to the correct source table(s), using the WhereScape RED re-target wizard as follows.


- 1 Right-click on a table object in the left pane and select **Change Column(s)**.
- 2 Select the **Column Source Table** checkbox. Select **dim_product** from the Original Value drop-down list. Select **load_product** from the New Value drop-down list. Click **Apply**.

Change Columns ✕

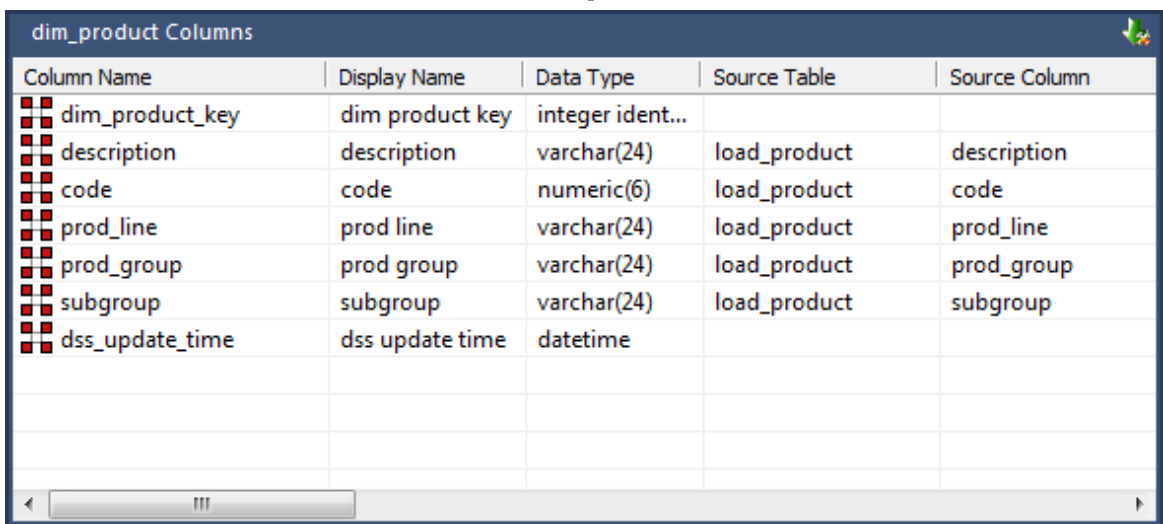
Select the changes to be applied. Note that only the selected columns will have the selected changes applied.

<p>Apply Changes to:</p> <p><input checked="" type="checkbox"/> Column Source Table</p> <p><input type="checkbox"/> Data Type</p> <p><input type="checkbox"/> Nulls</p> <p><input type="checkbox"/> Character Set</p> <p><input type="checkbox"/> Format String</p> <p><input type="checkbox"/> Numeric</p> <p><input type="checkbox"/> Additive</p> <p><input type="checkbox"/> Attribute</p> <p><input type="checkbox"/> EUL</p>	<p>Original Value:</p> <p>dim_product</p> <p> </p> <p> </p> <p> </p> <p> </p> <p> </p> <p> </p> <p> </p>	<p>New Value:</p> <p>load_product</p> <p> </p> <p> </p> <p> </p> <p> </p> <p> </p> <p> </p> <p> </p>	<p><input type="checkbox"/> Update matched columns only</p>
--	--	--	---

- 3 A message will be displayed to show that the source columns have been changed. Click **OK**.



- 4 Click **Close**.
- 5 Confirm the **Source Table** column in the middle pane.



Column Name	Display Name	Data Type	Source Table	Source Column
dim_product_key	dim product key	integer ident...		
description	description	varchar(24)	load_product	description
code	code	numeric(6)	load_product	code
prod_line	prod line	varchar(24)	load_product	prod_line
prod_group	prod group	varchar(24)	load_product	prod_group
subgroup	subgroup	varchar(24)	load_product	subgroup
dss_update_time	dss update time	datetime		

RETRO COLUMN PROPERTIES

Each Retro column has a set of associated properties. The definition of each property is defined below.

If the **Column name** or **Data type** is changed for a column, then the metadata will differ from the table as recorded in the database. Use the **Validate>Validate Table Create Status** menu option to compare the metadata to the table in the database. When positioned on the table name after the validate has completed, a right-click menu option **Alter Table** will alter the database table to match the metadata definition.



TIP: If a database table's definition is changed in the metadata then the table will need to be altered in the database. Use the **Validate>Validate Table Create Status** to compare metadata definitions to physical database tables. The option also provides the ability to alter the database table, through a pop-up menu option from the validated table name.

Table Name

The table to which the column belongs. Read only.

Column Name

The database name of the column. This name must conform to database column naming standards. Typically, such standards exclude the use of spaces etc. A good practice is to only use alphanumeric characters and the underscore character. Changing this field alters the table's definition.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Title

This defines how the business refers to this column. Although it would not be normal practice to provide end user access to Retro tables, the contents of this field are passed through to the fact table during the drag and drop operation, so this field can be commented for later use in the fact table if desired. This field does not affect the physical table definition.

Note: A case conversion button on the right converts the text between different cases: UPPER CASE, Capitalized Case and lower case. The mode cycles to the next case in the sequence each conversion.

Column Description

Normally this field will contain the strategy for acquiring or populating a column. It is a useful place to record specific problems with the source data.

In the case of dimension keys, this field is used to show the join between the Retro table and the dimension, once it has been defined as part of the Retro table update procedure generation.

Column Order

This numeric field provides the ordering of the columns in the database table create statement. The lowest numbered column will appear first in the table. Although this affects the physical

table definition no action will be taken unless the table is re-created in the database. The columns can be re-numbered based on the existing order by choosing the **Respace Order Number** pop-up menu option when positioned over any column in the table. This action will number the columns in increments of 10 starting at 10. In addition to a simple change of the order field, the column order can be changed by first displaying the columns in the middle pane and then using drag and drop to move the columns around. This drag and drop process will automatically renumber the columns as required.

Data Type

This is the database specific data type. It must be a valid data type for the data warehouse database platform. Typical Teradata databases often have integer, numeric(), varchar(), char(), date and timestamp data types. See the database documentation for a description of the data types available. Changing this field alters the table's definition.

Null Values Allowed

This check box when set allows nulls to exist in this column. If cleared then nulls are not permitted in the column. Although this affects the physical table definition no action or comparison is made on this field. If you wish to change this attribute for a column in an existing table, then the only supported way is to re-create the table, which may not be desirable.

Default Value

The default value Teradata puts in the column, if the column is not included in the insert statement populating the table.

Character Set

Select Latin or Unicode.

Format

This field is to allow the definition of a format mask for end user tools. It does not affect the physical table definition. This field would not normally be populated.

Character Comparison/Sorting

Indicates the Teradata case of the column; one of: case specific, not case specific, uppercase case specific or uppercase not case specific.

Compress/Compress Value

Indicates the column is compressed and enables the compress value text box. In the compress value text box, you can enter the list of values to use when compressing the column.

Numeric

Indicates that the column in question is numeric when set. This is normally only relevant for fact tables and would not normally be used for Retro tables.

Additive

Indicates that the column in question is additive when set. This is normally only relevant for fact tables and would not normally be used for Retro tables.

Attribute

Indicates that the column in question is an attribute when set. This is normally only relevant for fact tables and would not normally be used for Retro tables.

Business Key

Indicates that the column is part of the primary business key when set. Multiple columns can form the primary business key. This indicator is set and cleared by WhereScape RED during the Retro update procedure generation process. This should not normally be altered.

Key Type

The key type is used internally by WhereScape RED in the building of the update procedure and the generation of indexes. It can be altered here, but this should only be done if the consequences are fully understood and tested. The supported values are:

Key type	Meaning
0	The artificial key. Set when the key is added during drag and drop table generation.
1	Component of all business keys. Indicates that this column is used as part of any business key. For example: By default the <code>dss_source_system_key</code> is added to every dimension table. It is considered part of any lookup on that table and has the key type set to 1. Set when the column is added during drag and drop table generation.
2	Indicates that this column is a dimensional join. Used on fact tables to indicate the dimension keys. Results in bitmap indexes being built for the columns. Set during the update procedure generation for a fact table, based on information from the staging table.
3	Slowly changing column indicator. Used on dimension tables to indicate that the column is being managed as a slowly changing column within the context of a slowly changing dimension. Set when a column is identified during the dimension update procedure generation.
4	Previous value column indicator. Used on dimension tables to indicate that the column is being managed as a previous value column. The source column identifies the parent column. Set during the dimension creation.
A	Indicates that the column is part of the primary business key. Set whenever a business key is defined as part of an update procedure generation.
B-Z	Indicates that the column is part of a secondary business key. Only used during index generation and not normally set.

Source Table

Indicates the table from which the column was sourced. This source table is normally a load table, or a dimension table within the data warehouse. If the column was sourced from multiple tables, then the normal practice is to record one of the tables in this field and a comment listing all of the other tables in the **Source Strategy** field. This field is used when generating a procedure to update the Retro table. It is also used in the track back diagrams and in the documentation.

Source Column

Indicates the column in the **Source table** from which the column was sourced. Such a column is normally a load table column, which in turn may have been a transformation or the combination of multiple columns. This may also be a dimensional key where a dimension is being joined.

Join

This field is set by WhereScape RED during the building of the update procedure and should not normally be altered. It indicates that the column in question provides a join to a dimension table. The **Source Table** and **Source Column** fields will provide the dimension's side of the join. The options for this field are: False, True, Manual and Pre Join.

Setting this field to Manual changes the way the dimension table is looked up, during the update procedure build. It allows you to join the dimension manually in the Cursor mapping dialog (where the 'Where' clause is built). The usual dialog for matching the dimension business key to a column or columns in the retro table is not displayed if this option is enabled.

Setting this field to Pre Join activates the **Join Source** field and allows you to select a table from the drop-down list.

Join Source

When the Join option is set to False, this field becomes inactive. When the Join option is set to True or Manual, this field is set to the current table name. When the Join option is set to Pre Join, then you can select the required table from the drop-down list.

RETRO COLUMN PROPERTIES SCREEN

Retro Column continent.continent

Properties

Transformation

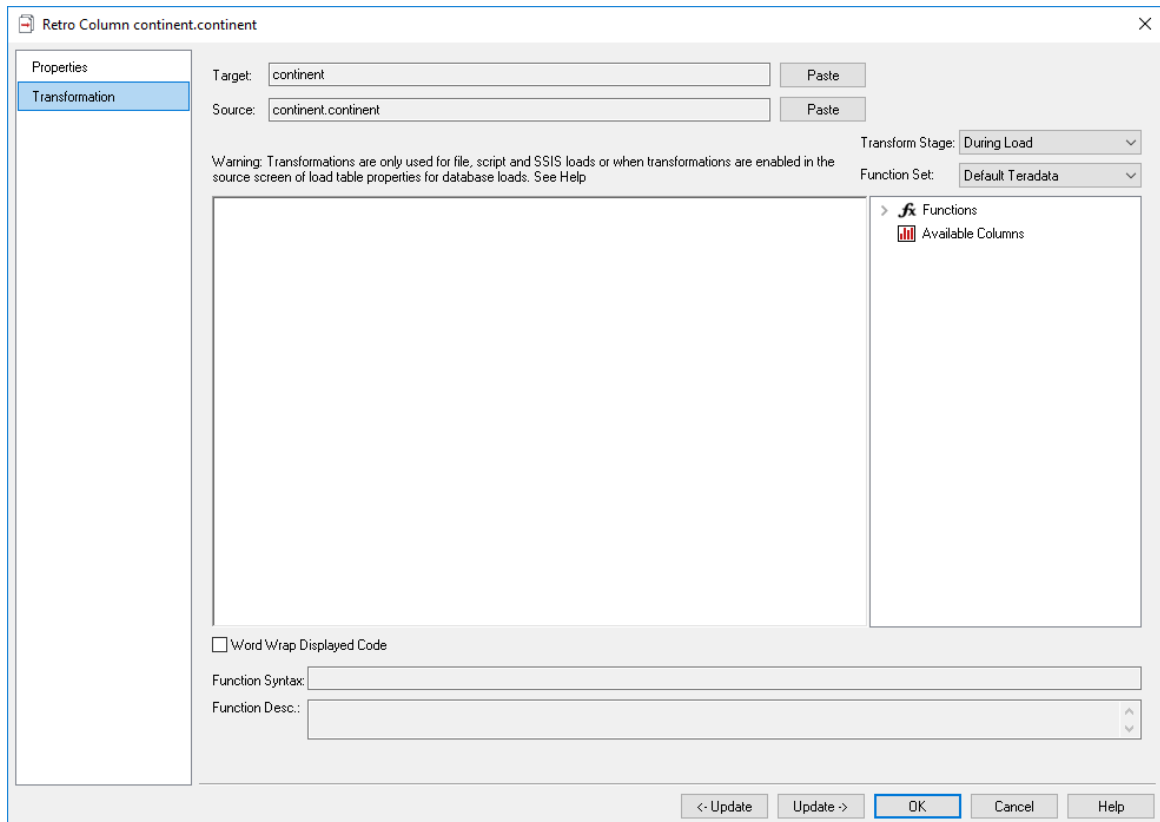
General	
Table Name	continent
Column Name	continent
Column Title	continent
Column Description	
Physical Definition	
Column Order	20
Data Type	varchar(256)
Null Values Allowed	<input checked="" type="checkbox"/>
Default Value	
Character Set	LATIN
Format	X(256)
Character Comparison/Sorting	
Compress	<input type="checkbox"/>
Meta Definition	
Numeric	<input type="checkbox"/>
Additive	<input type="checkbox"/>
Attribute	<input checked="" type="checkbox"/>
Business Key	<input type="checkbox"/>
Key Type (0,A,B,C...)	
Source Details	
Source Table	continent
Source Column	continent
Transformation	
Join	False

Column Name
Database-compliant name of the column.
Dialog Opening Value: continent

<- Update Update -> **OK** Cancel Help

RETRO COLUMN TRANSFORMATIONS

It is possible to do transformations on Retro table columns. It is recommended that transformations are not performed on columns that are dimension keys or the business keys for the table. The **Transformation** screen is as follows:



Note: Transformations are only put into effect when the procedure is re-generated.

See *Transformations* (on page 593) for more details.

CHAPTER 37

INTEGRATING WHERESCAPE RED INTO AN EXISTING WAREHOUSE

Two main options exist in terms of bringing WhereScape RED into an existing data warehouse environment:

- 1 Rebuild tables and procedures with WhereScape RED.
- 2 Integrating existing tables and procedures into WhereScape RED.

Both options require manual coding changes to stored procedures. The main advantages and disadvantages of these two options are discussed below, and in detail in the following sections.

Rebuilding

The rebuild option is essentially a redevelopment of the existing data warehouse utilizing the knowledge acquired in the initial development and the rapid development capabilities of WhereScape RED. A rebuild will take more time and effort than just integrating existing tables and procedures but will provide a better platform on which to extend the data warehouse. See the *Rebuild* (see "*Rebuilding*" on page 974) section for the approach to achieve this option.

Integrate

Existing data warehouse tables can be identified to WhereScape RED. The tables are seen and can be managed to a degree. The main disadvantage is the increased difficulty in utilizing these tables when trying to extend the data warehouse. This option is however significantly quicker and easier than a rebuild. It is discussed in detail in the *Integrate* (see "*Integrating*" on page 975) section of this chapter.

The decision as to which option to choose will depend on the size and complexity of the existing data warehouse. Another important factor is the degree to which the existing data warehouse is to be extended. If future enhancements revolve around new analysis areas that have little overlap with the existing environment then an integrate may be the best answer. If the data warehouse is small and relatively simple than a rebuild may be worth considering. In any event the best plan may be to do a test integrate and then re-evaluate the situation.

IN THIS CHAPTER

Rebuilding.....	974
Integrating.....	975

REBUILDING

The rebuild process essentially is a total re-creation of the data warehouse. One of the major impacts of such an approach is the 'end user layer', or rather the effect on the end user tools and saved queries and reports that are currently in use. The redesign or redeployment of this interface to the end users of the data warehouse may be too large a task to undertake. The problem can be circumvented to some degree though the use of views to make the new data warehouse environment look the same as the previous. But it is this impact and the subsequent re-testing process that must be considered when deciding to undertake a rebuild.

The advantages of a rebuild is the seamless integration of future analysis areas into the data warehouse and the single point of management that is provided. The major steps in the rebuild process will depend very much on the environment being replaced. As a guideline the following steps may be worth considering.

The rebuild process

- 1 Load a copy of the WhereScape metadata repository into an otherwise empty test environment. See the **Installation and Administration Guide** for instructions on how to load a metadata repository.
- 2 Ensure there are no public synonyms that point to existing table names, if the rebuild process is to use the same names as some or all of the existing tables.
Working within the WhereScape RED tool proceed to:
- 3 Create connections to the new test data warehouse and to all the source systems.
- 4 Using the source system knowledge from the existing data warehouse, create the appropriate load tables in the data warehouse based on the existing extract or load methodology.
- 5 Build up the stage and model tables using the same column and table names where possible.
- 6 Examine the existing procedures or update methodology and include this into the generated stored procedures.
- 7 Test the new environment.
- 8 Work out a plan to convert the existing data into the new data warehouse. Where possible it is best to keep existing key values and re-assign sequences to match these existing key values where appropriate.
- 9 Convert and test the old data warehouse data in the new environment.
- 10 Redeploy the end user tool access.

INTEGRATING

The integrate process

The steps in the integrate process are:

- 1 Create a test environment (database user) with the existing data warehouse tables loaded.
- 2 Load a copy of the WhereScape metadata repository into this test environment. See the **Installation and Administration Guide** for instructions on how to load a metadata repository.
Working within the WhereScape RED tool proceed to:
- 3 Create any connections to Windows servers where host scripts are currently executed. See **Windows connection** (see "**Windows**" on page 163).
- 4 Create a Data Warehouse connection mapping back to the test environment. See **connection to the data warehouse** (see "**Database - Data Warehouse/Metadata Repository**" on page 147).
- 5 Incorporate any Host system scripts currently used. See **incorporating host scripts** (see "**Integrating, Host Scripts**" on page 976).
- 6 Browse the Data Warehouse connection (Browse>Source Tables).
- 7 Drag and drop each existing data warehouse table into an appropriate object type. See **Selecting an appropriate table type** (see "**Integrating, Selecting a Table Type**" on page 977)
- 8 Answer the retrofit questions, and build any required procedures. See **integrate questions** (see "**Integrating, Questions**" on page 978).
- 9 Edit and amend all generated procedures, or create new procedures to handle the existing processing methodology. See **procedure changes** (see "**Integrating, Procedures**" on page 981).
- 10 Test the new environment.

Removing the metadata for a table

It is possible to delete the metadata for a table without deleting the table itself. For example, if the integrate process is incorrectly undertaken, the metadata for the specific table can be removed.

To delete the metadata only: First, select the table and using the right-click menu select **Delete**. A dialog box asks if you wish to delete the object and drop the table. Answer **No**. A second dialog box appears asking if you just want to delete the metadata. Answer **Yes** to this question and only the metadata is removed.

INTEGRATING, HOST SCRIPTS

Existing windows host scripts can be brought into the WhereScape RED meta data. To incorporate an existing script the process is as follows:

- 1 Create a **Host Script** object using RED. In the left pane click on a project or the **All Objects** project and using the right-click menu select **New Object**. The new object dialog box will appear. From the Object Type drop-down select **Host Script** and enter a name for the new script.
- 2 The following properties dialog will appear. Select the script type. Either UNIX or Windows script. Select the appropriate connection from the **Default Connect** drop-down. Fill in the Purpose field to define the role of the script and then click Update to store the changes.

The screenshot shows a dialog box titled "Host Script monitor_db_mail". It has a "Properties" tab selected on the left. The main area contains the following fields:

- Name:** monitor_db_mail
- Type:** Windows script (dropdown)
- Purpose:** (empty text area)
- Owner:** dssdemc
- Delete Lock:**
- Last Update By:** (empty text field)
- Default Connect:** (empty dropdown)
- Edit Lock:**
 - Locked For Edit By:** (empty text field)
 - Edit Lock Reason or Last Update:** New Script (dropdown)
- Timestamps:**
 - Created:** 2015-11-19 03:12:07
 - Last Update:** (empty text field)
 - Compiled:** (empty text field)

Buttons for "OK", "Cancel", and "Help" are located at the bottom right of the dialog.

- 3 Double-click on the new script or right-click the new script and select **Edit the Script**.
- 4 Within the script editor, either paste the script or if it is available on the PC, select the **File/Insert from file** option and load the file.
- 5 The script will need to be modified to handle the standards required by the scheduler. See **Loading via a host script** (see "**Script based loads**" on page 275) for more details.

INTEGRATING, SELECTING A TABLE TYPE

When integrating existing tables there may not be a clear decision as to which table type to use. As a guideline, the following groupings can be considered.

Temporary tables:

- Load tables
- Stage tables

Permanent tables:

- Model tables

Although these table groups have very distinct names in terms of data warehousing, they do not impose any restrictions on the types of tables they contain. The table groupings are most relevant in the automatic generation of procedures, and in the sequencing for the scheduler.

Typically, a mapping table may be stored in the Staging table group.

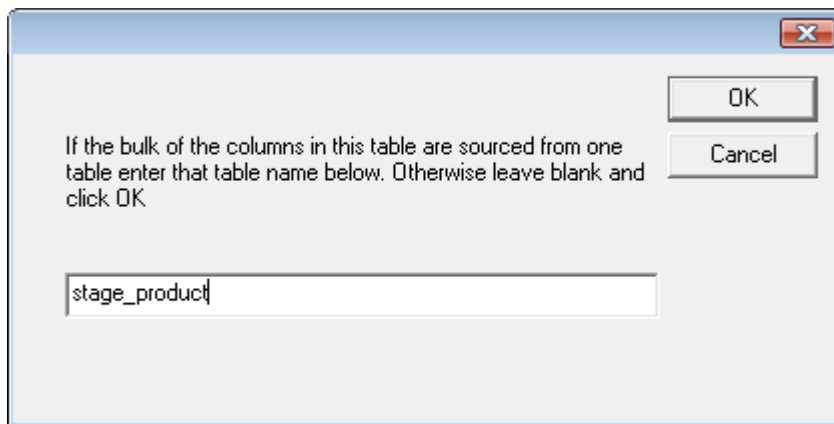
INTEGRATING, QUESTIONS

When a table within the data warehouse schema that is unknown to RED is dropped onto a table target, the following dialog appears.



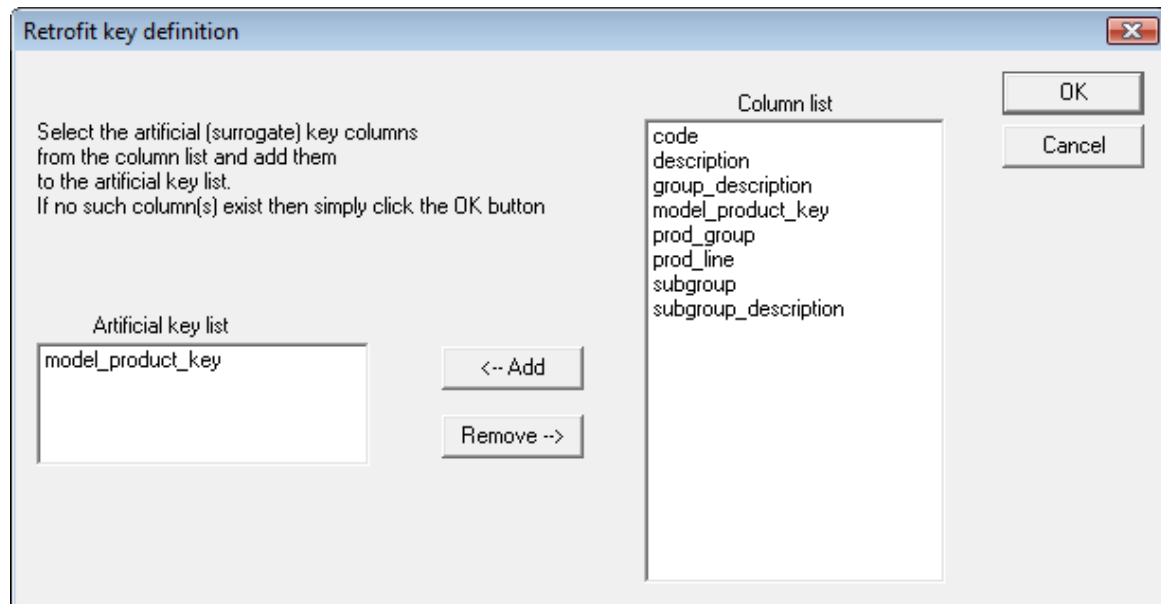
If this is a retro-fit then click Yes to proceed with the retro-fit process. The standard **New Object** dialog will appear and it would be advisable to leave the name of the object unchanged so that it matches the existing table.

A dialog asks if the bulk of the columns in this table are derived from another table. If they are enter the table from which these columns derive at this stage. The purpose of this dialog is simply to set the 'source table' field against each of the columns for the table.



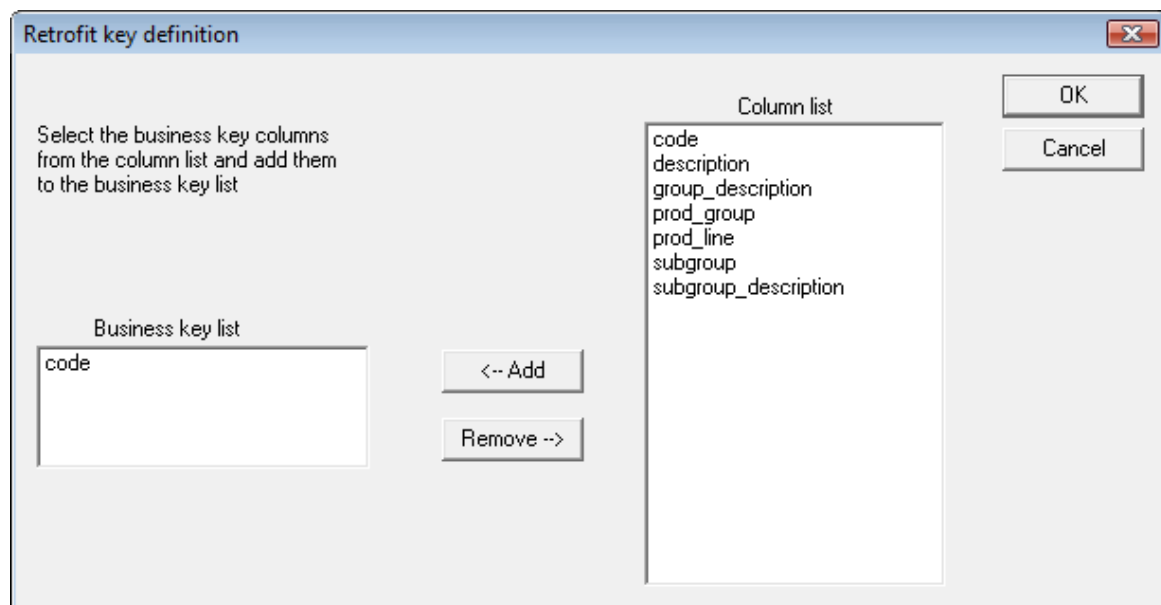
Artificial Key definition

If the table target is a model table, then the following dialog appears to enable the definition of any artificial or surrogate key. If no such key exists, then simply proceed to the next question.



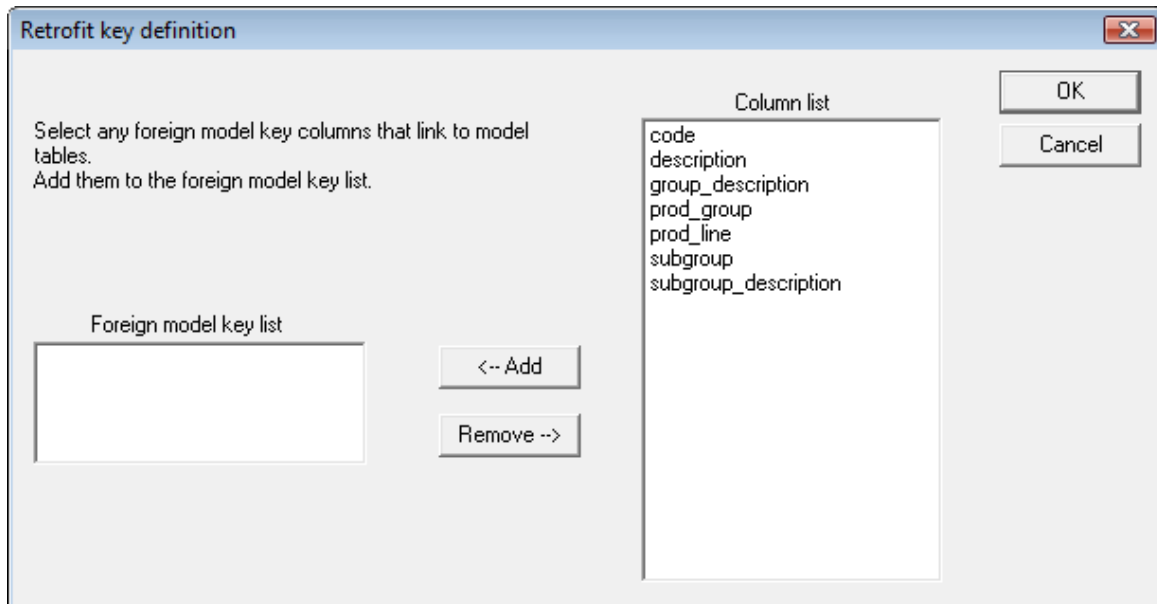
Business Key definition

A dialog prompts for the selection of the business key from the table. Multiple columns may constitute the business key, but they must uniquely identify each record in the table.



Foreign Key definition

A dialog will now prompt for the selection of the foreign keys from the table.



Update Procedure

If the target table is a stage, model or aggregate a dialog box will ask if an update procedure is required. If selected a subsequent dialog will define the structure and content of this procedure. If an existing procedure is to be modified to include the scheduler imposed procedure standards it is best to select this option and use the existing procedure name to cut down on the amount of work required. The existing procedure can then be loaded into the generated procedure using the procedure editor.

Index definition

Any indexes associated with the table is automatically defined and loaded into the metadata. Changes may need to be made in terms of when the indexes are dropped and to set the **Drop Before Update** check box, if appropriate and the scheduler is to be used to manage these indexes.

Procedure creation

If defined, the get key and update procedures is generated. They need to be manually edited and compiled. See the section on *retro-fitting procedures* (see "*Integrating, Procedures*" on page 981).

INTEGRATING PROCEDURES

The procedures managed by the WhereScape scheduler require the following standards.

Parameters

The procedure must have the following parameters in the following order:

Parameter name	Input or Output	Data Type
p_sequence	Input	Integer
p_job_name	Input	Varchar(256)
p_task_name	Input	Varchar(256)
p_job_id	Input	Integer
p_task_id	Input	Integer
p_return_msg	Output	Varchar(256)
p_status	Output	Integer

The input parameters are passed to the procedure by the scheduler. If the procedure is called outside the scheduler then the normal practice is to pass zero (0) in the sequence, job_id and task_id. A description of the run can be passed in the job name and the task name is typically the name of the procedure.

The output parameters must be populated by the procedure on completion. The return_msg can be any string up to 256 characters long that describes the result of the procedures execution. The status must be one of the following values:

Status	Meaning	Description
1	Success	Normal completion
-1	Warning	Completion with warnings
-2	Error	Hold subsequent tasks dependent on this task
-3	Fatal Error	Hold all subsequent tasks

The major task in retro-fitting a procedure is in adapting it to the WhereScape scheduler standards and work flow.

INTEGRATING, VIEWS

When integrating views an additional step is required if you want WhereScape RED to be able to recreate the view.

The view will be mapped correctly and the **Get Key** function can still be built. This step is only required if the view is to be re-created.

Note: The **Get Key** function is not available for dimension objects that are stored in custom database targets.

Change the source column on the artificial key to match the artificial key in the table from which the view was created.

INTEGRATING, WHERESCAPE TABLES

When retro-fitting WhereScape generated tables and views a number of additional considerations need to be taken.

Change the properties of all such columns. The key type should be set to 1, and the primary business key checkbox should be set.


CHAPTER 38

RELATIONSHIP MAINTENANCE

Relationship Maintenance is available for the maintenance of joins between tables; providing a way to record joins between tables when surrogate keys are not being used. This functionality then enables the generation of **Links Diagrams** for these tables.

NOTE: It is necessary to explicitly specify relationships for tables on **Tabular** target databases for the relationships to be created in the Tabular database.

Relationship Maintenance options are available in the Relationships sub-menu when right-clicking on an object in the **Object Pane**.



ADD RELATIONSHIP

To add a relationship, right-click on the object in the **Object Pane** and select **Relationships->Add Relationships**. The following dialog appears.

For each object in the relationship, enter in the following details:

Join Type

For the **Join Type**, choose between the following:

- Undefined
- Many to One
- One to One
- One to Many
- Many to Many

Primary Join

If this is a primary join, select the **Primary Join** check-box. For MSAS Tabular tables this defines an active relationship. At most, one relationship between two specific tables can have this enabled.

Object Type

Enter the **Object Type** from the drop-down box (Data Store, Load table, Stage table, etc.).

Table/View

Enter the name of the object in the relationship.

Column

Enter the **Column** to join in each object.

Once you have entered the details for the join, the joined columns are displayed in the list of **Joins** at the bottom of the dialog box. Erroneous joins can be removed by right-clicking on the join and selecting **Remove Join**. All joins can be removed by clicking the **Reset** button.

To add all the relationships shown in the list of **Joins**, click the **Add** button.


LIST RELATIONSHIPS

To view relationship for an object, right-click on the object in the **Object Pane** and select **Relationships->List Relationships**. The relationships for the selected object are displayed in the **Drop Target Pane** (middle pane).

Relationships for ds_customer							
Object	Column	Join Type	Object	Column	Primary Join	Part	
ds_customer	code	One to One	dim_customer	code	Y	1 of 1	

Multi-column joins are shaded when one join is selected.

Right-clicking on a column or join displays the following menu:

**Modify Relationship**

The Modify Relationship option shows the following dialog, allowing the editing of joins (including multi-column joins) between the two objects in the selected relationship.

Modify Relationship ✕

Specify the Relationship by selecting a table and column from both the left hand side and the right hand side, and then clicking the Modify button.
For a multiple column Relationship, select each pair of columns before clicking the Modify button.

Join Type: One to One Primary Join

Object Type: Data Store Load

Table/View: ds_test load_customer

Column:

Joins:

code	code
cname	name

Relationships are edited in this dialog in the same way as the **Add Relationship** dialog above. **Object Types** and **Table names** cannot be modified.

Delete Relationship

Deletes the selected relationship.

GENERATE RELATIONSHIPS

To generate relationships in metadata for an object, right-click on the object in the **Object Pane** and select **Relationships->Generate Relationships**. Results are shown in the **Results Pane**.

Results	
Object	Message
fact_customer	Generating relationships.
fact_customer	Relationship FK fact_customer.dim_customer_key to PK dim_customer.dim_customer_key has been created as a primary join relationship.

CHAPTER 39

UPGRADING RED

The upgrading of RED is discussed at some length in the **Installation and Administration Guide**.

CHAPTER 40

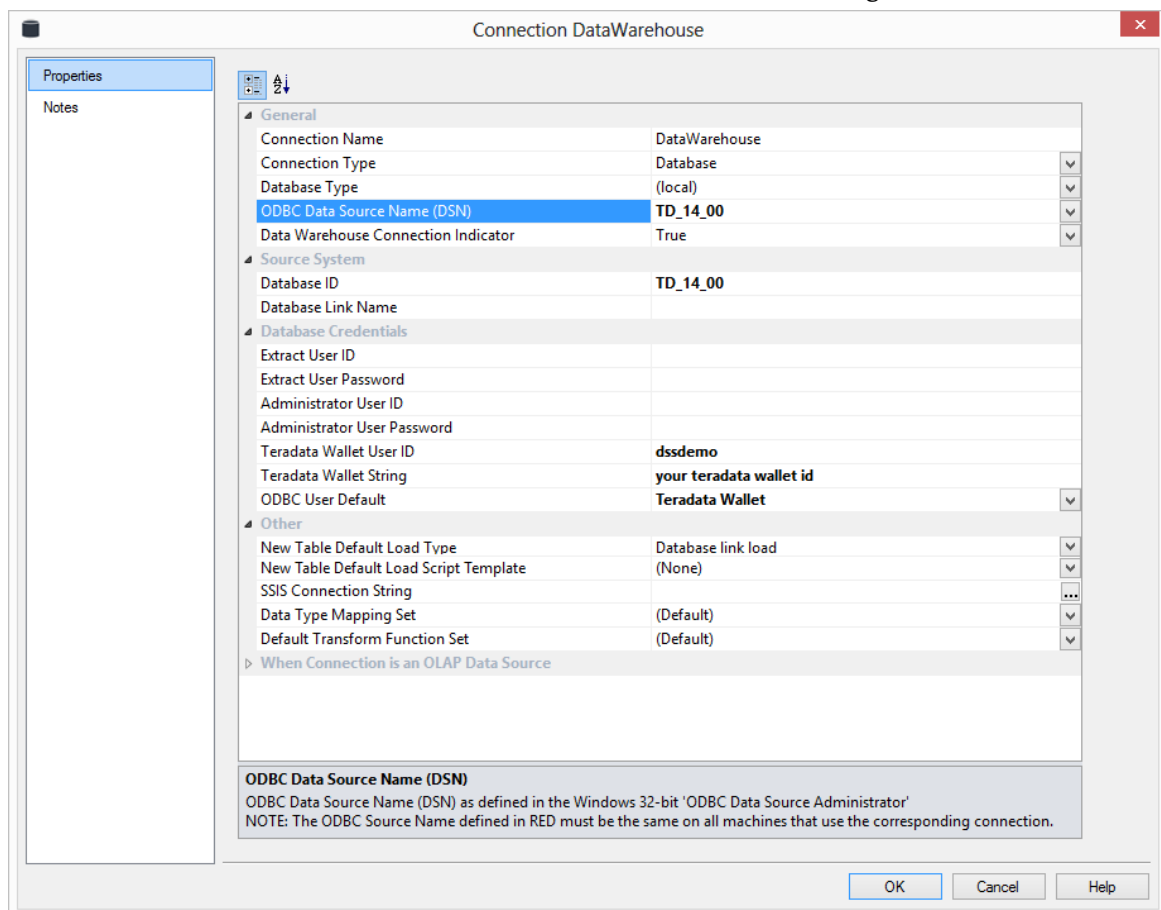
LOGIN CHECKS

The following checks are performed during login; and if necessary, warning messages are displayed:

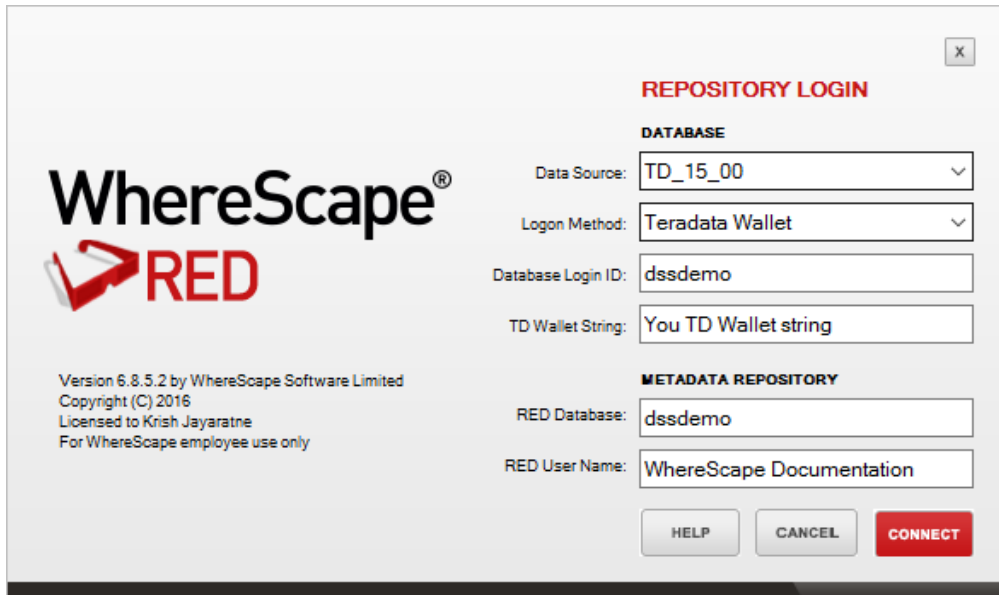
- 1 Warning if the login Data Source does not match the data warehouse connection's ODBC DSN.

You can correct this issue by performing one of the following actions:

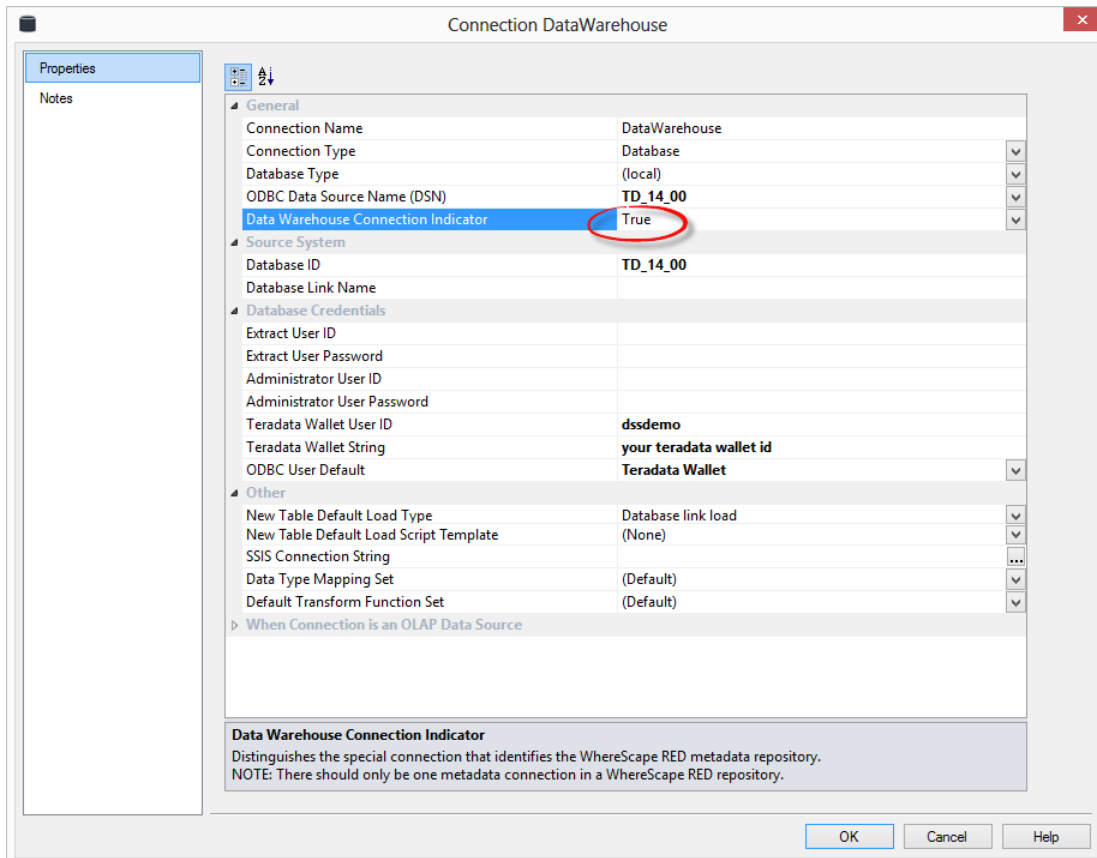
- Alter the Data Warehouse Connection **ODBC Source** to match the login **Data Source**.



- Log off and log back in using the **Data Source** with the same name as the Data Warehouse Connection **ODBC Source**.



- 2 Warning if more than one connection has the data warehouse check-box on:
You can correct this issue by editing all the connections and making sure that only one is set to Data Warehouse:



CHAPTER 41

DATA TYPE MAPPINGS

IN THIS CHAPTER

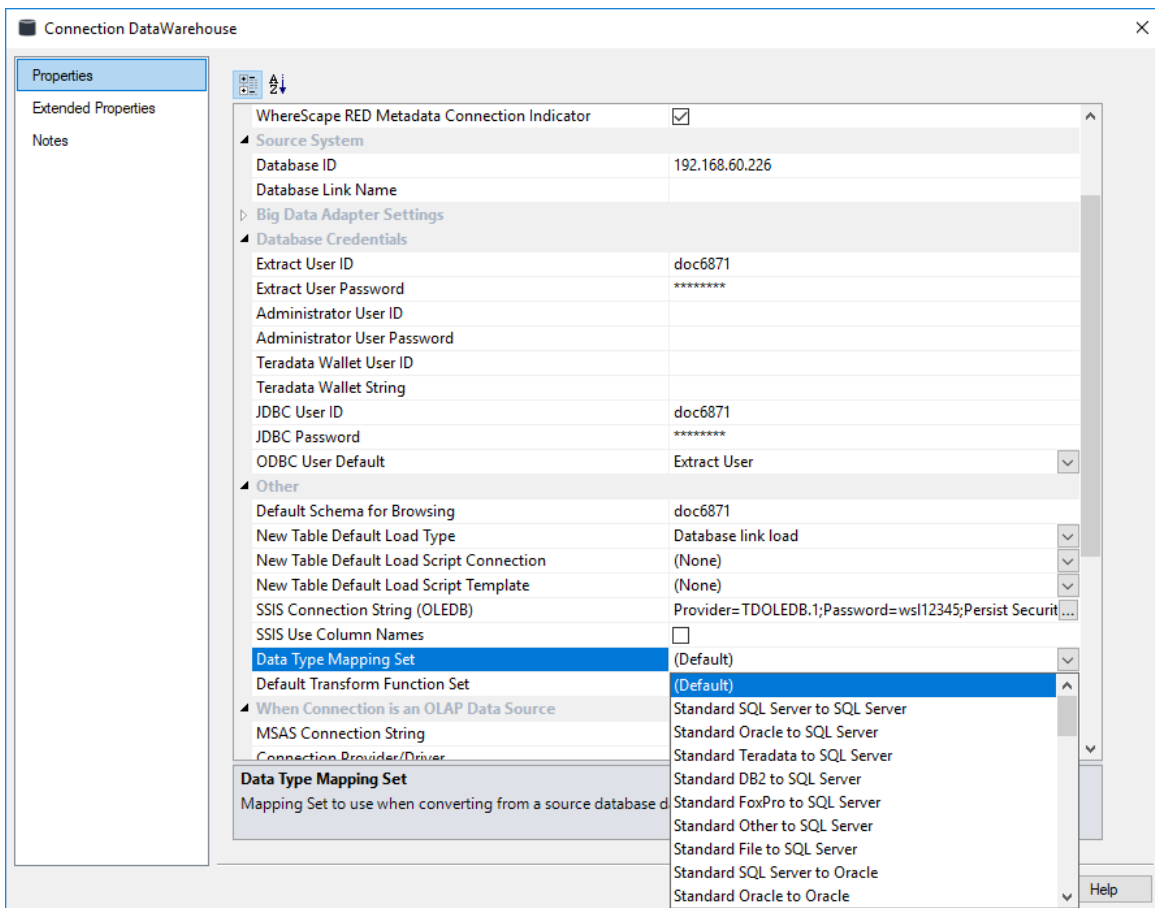
Using Data Type Mapping Sets	992
Maintaining Data Type Mapping Sets	994
Loading Custom Data Type Mapping Sets	1008
Exporting Custom Data Type Mapping Sets	1010
Custom Data Type Mapping Examples	1012

USING DATA TYPE MAPPING SETS

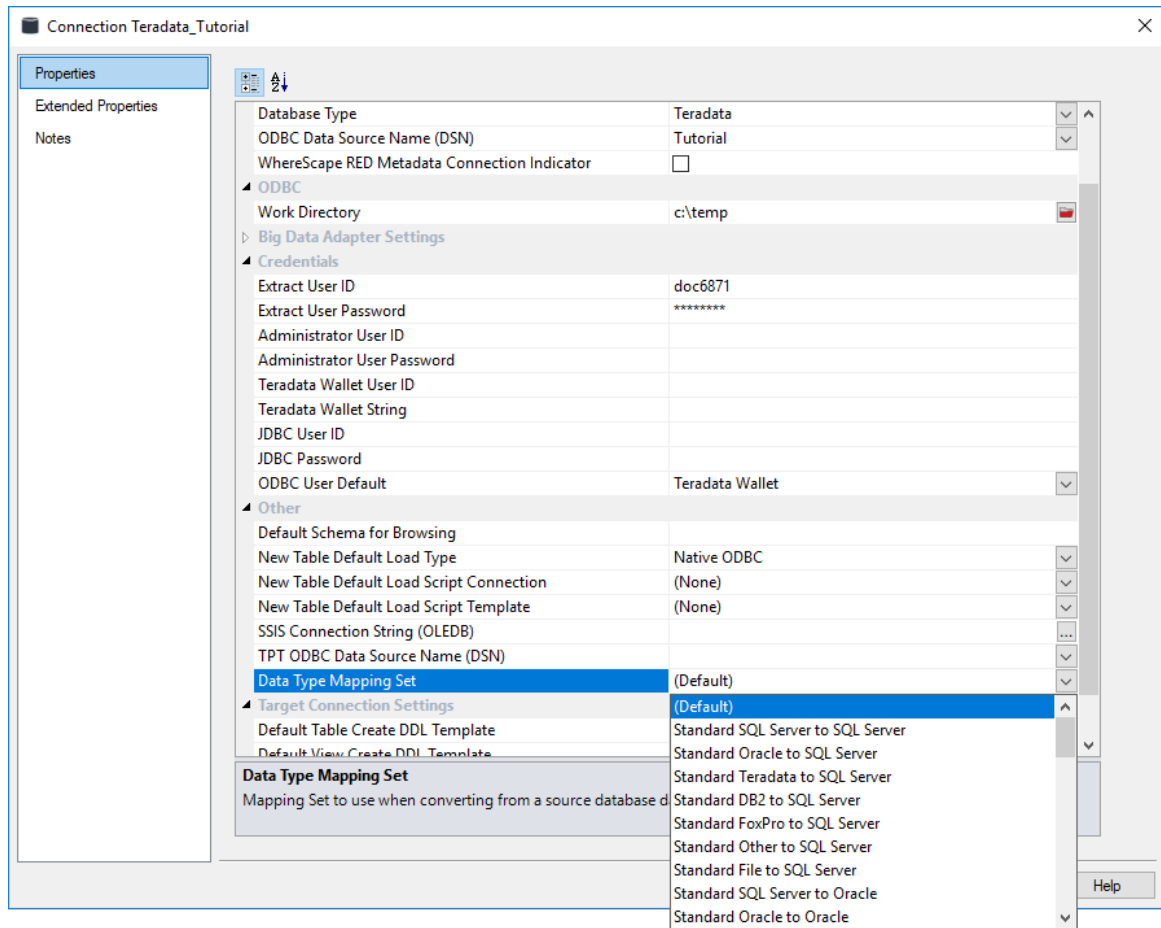
Data type mapping sets contain a list of mappings that are used when loading tables into the data warehouse.

Custom data type mapping sets give you the ability to automatically change the data type of any column or to add column transformations when dragging and dropping new load tables. These mapping sets may be created, edited, deleted, imported and exported using the **Data Type Mappings** options on the **Tools** menu.

Data Warehouse Connection Properties Dialog



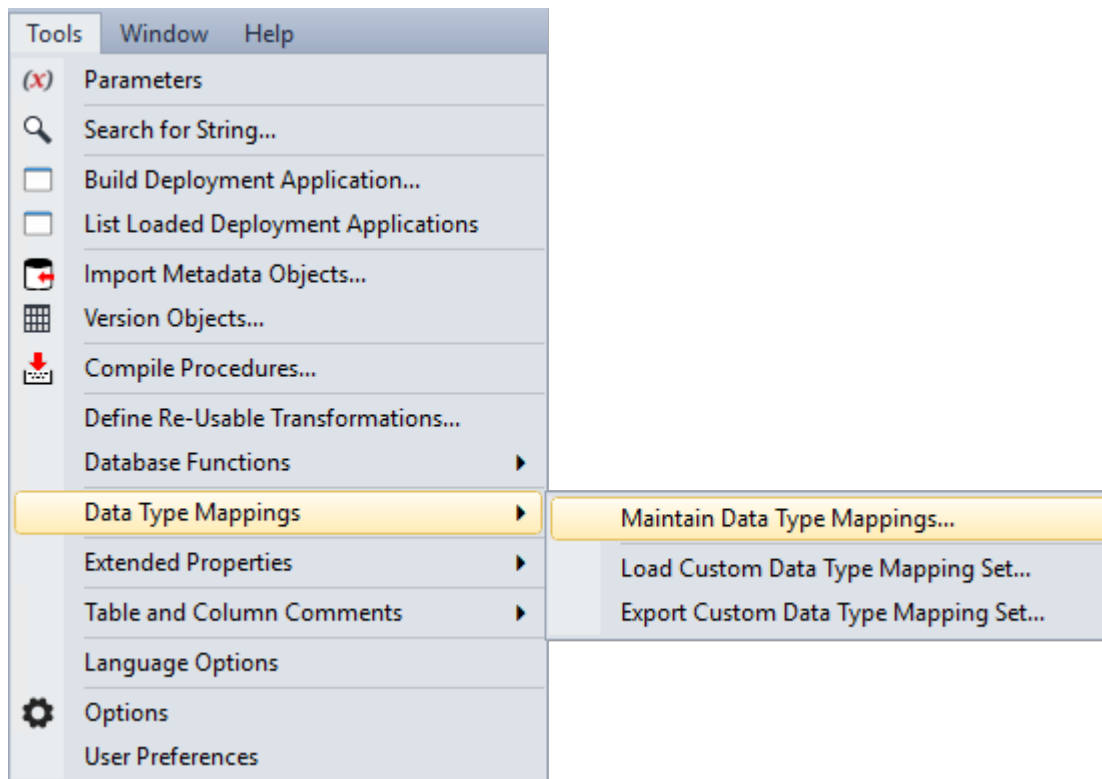
Non-Data Warehouse Connection Properties Dialog



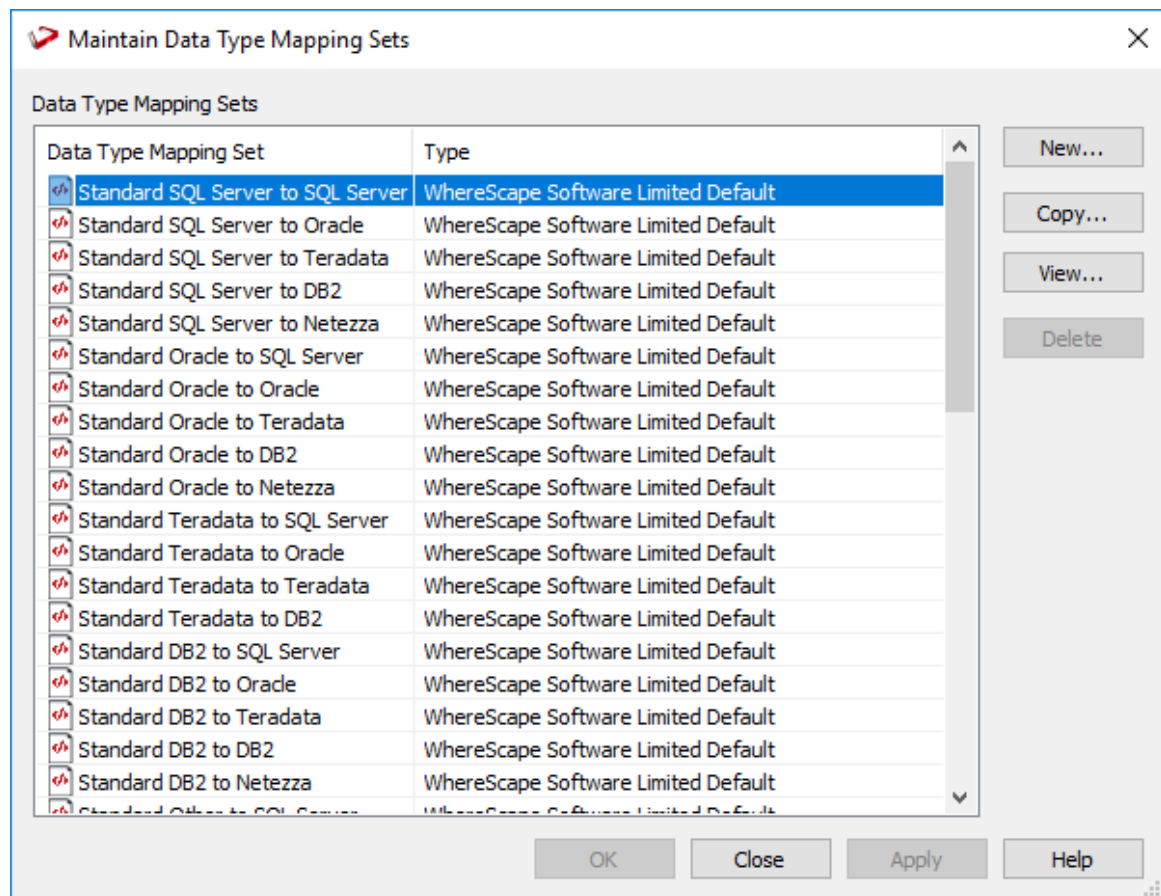
In the Connection Properties dialogs, a drop-down list **Data Type Mapping Set** is displayed. This is the default set that will be used when loading tables into the data warehouse.

MAINTAINING DATA TYPE MAPPING SETS

To maintain the data type mapping sets, select **Tools>Data Type Mappings>Maintain Data Type Mappings**.



The **Maintain Data Type Mappings** window is displayed.



Select a data type mapping set from the **Data Type Mapping Set** drop-down list.

Of the standard files, only the files relevant to the database that you are on will be displayed in the list.

To create a data type mapping set, see *Creating a New Data Type Mapping Set* (on page 996)

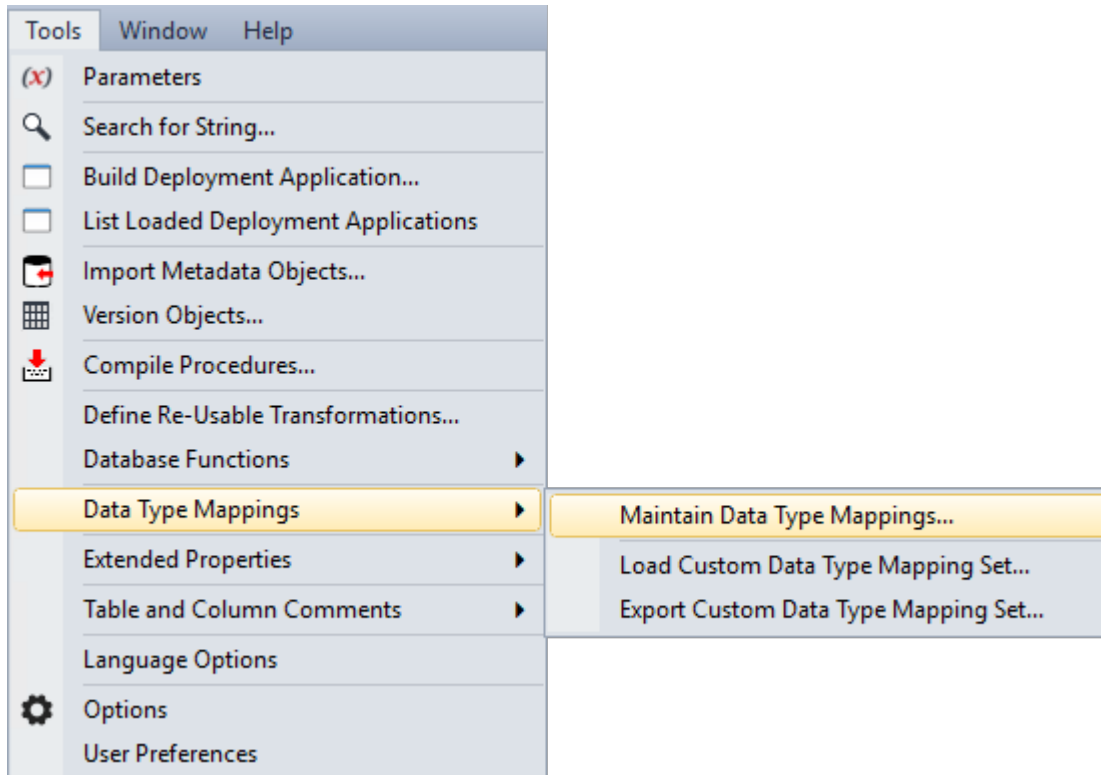
To copy a data type mapping set, see *Copying a Data Type Mapping Set* (on page 1000)

To edit a data type mapping set, see *Editing a Data Type Mapping Set* (on page 1003)

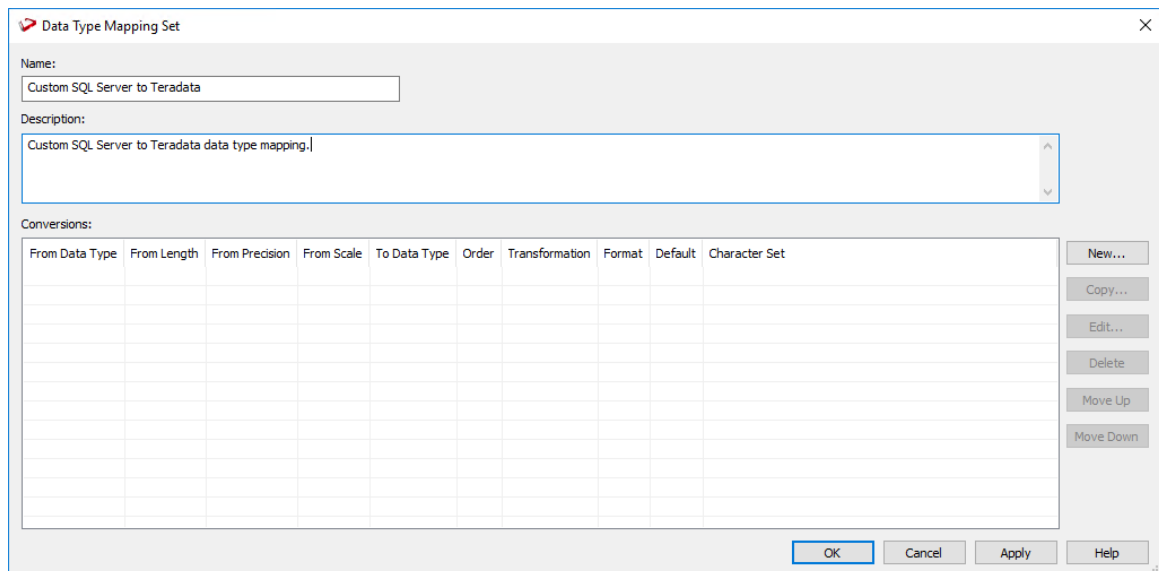
To delete a data type mapping set, see *Deleting a Data Type Mapping Set* (on page 1006)

CREATING A NEW DATA TYPE MAPPING SET

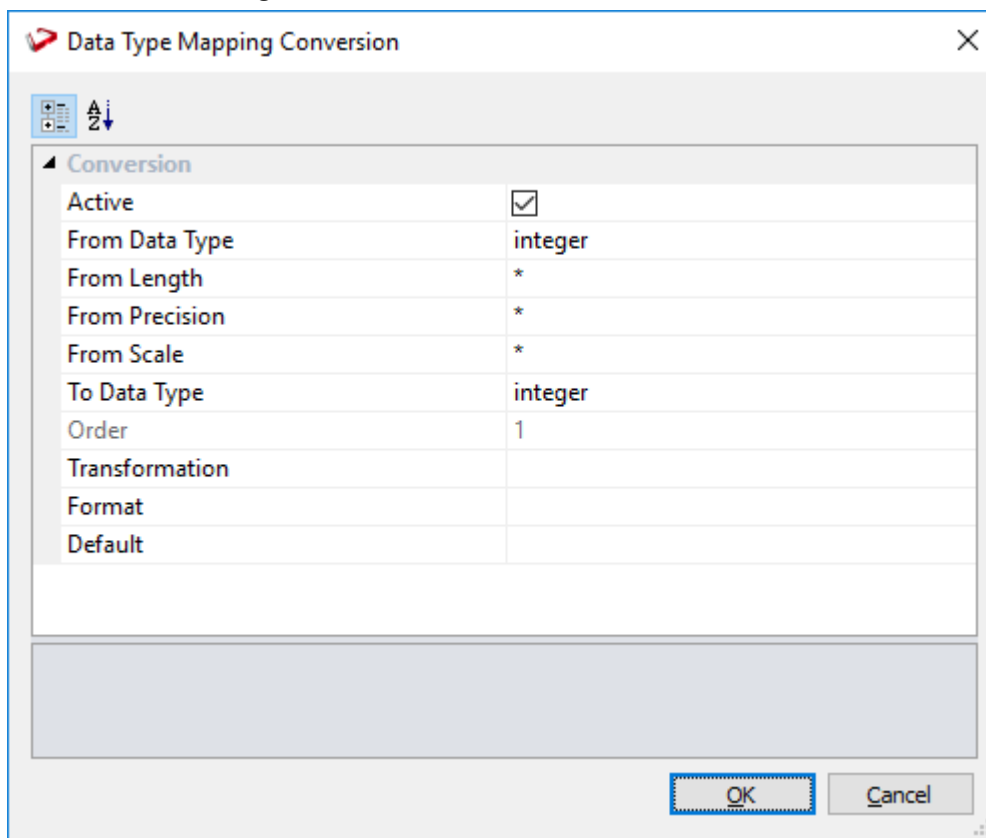
To create a new data type mapping set, select **Tools>Data Type Mappings>Maintain Data Type Mappings...**



- 1 Click the **New** button from the **Maintain Data Type Mapping Sets** window to enter the name and description of the new Mapping Set.



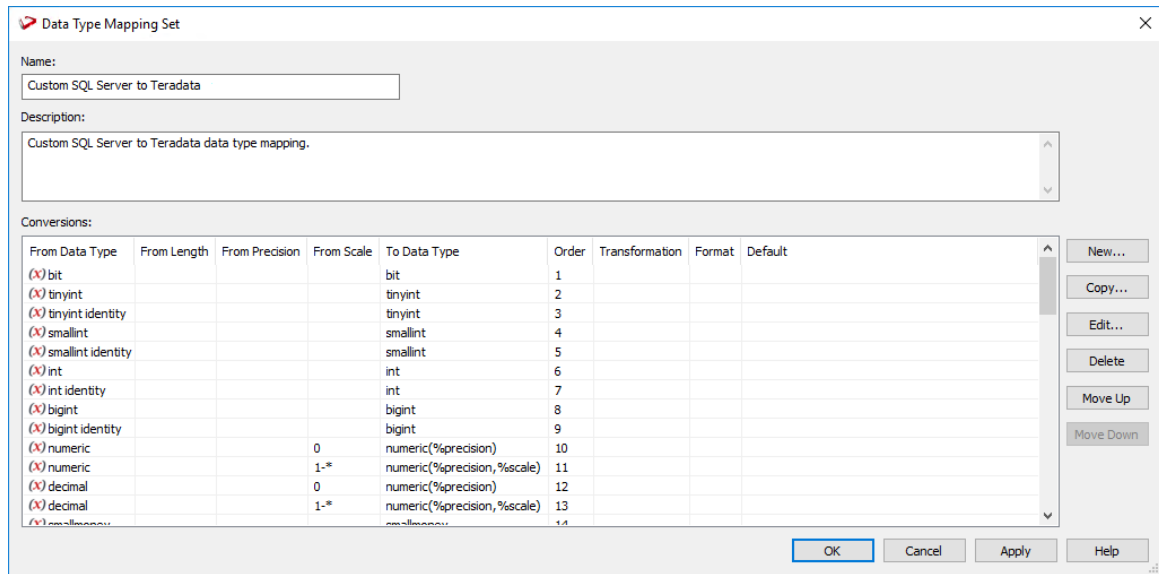
- 2 You can then enter the individual data type mapping conversion parameters by clicking the **New** button on the right side of the window.



The following describe the conversion parameters:

- Active - Indicates if the conversion parameter can be used when searching for a match.
- From Data Type - The source Data Type to that is used for the conversion to the Target Data Type. It can be any type that source RDBMS uses.
- From Length - The length parameter to match. (the number of characters or the number of bytes used to store the number).
- From Precision - The precision parameter to match. (the number of digits in a number).
- From Scale - The scale parameter to match. (the number of digits to the right of the decimal point in a number).
- To Data Type - The data type, including parameters, which is the output of this mapping. Available tokens are %data_type, %length, %precision and %scale. See **Custom Data Type Mapping Examples** (on page 1012).
- Order - The order of the parameter within the enclosing set of the conversion.
- Transformation - The transformation to apply when the conversion is applied. See **Custom Data Type Mapping Examples** (see "**Custom Data Type Mapping Examples**" on page 1012).
- Format - The format that values using the data type will use. Typically used for numeric values.
- Default - The default value to use for values that use the data type.

3 Click **OK** to save the parameter. Repeat to define more conversion parameters.

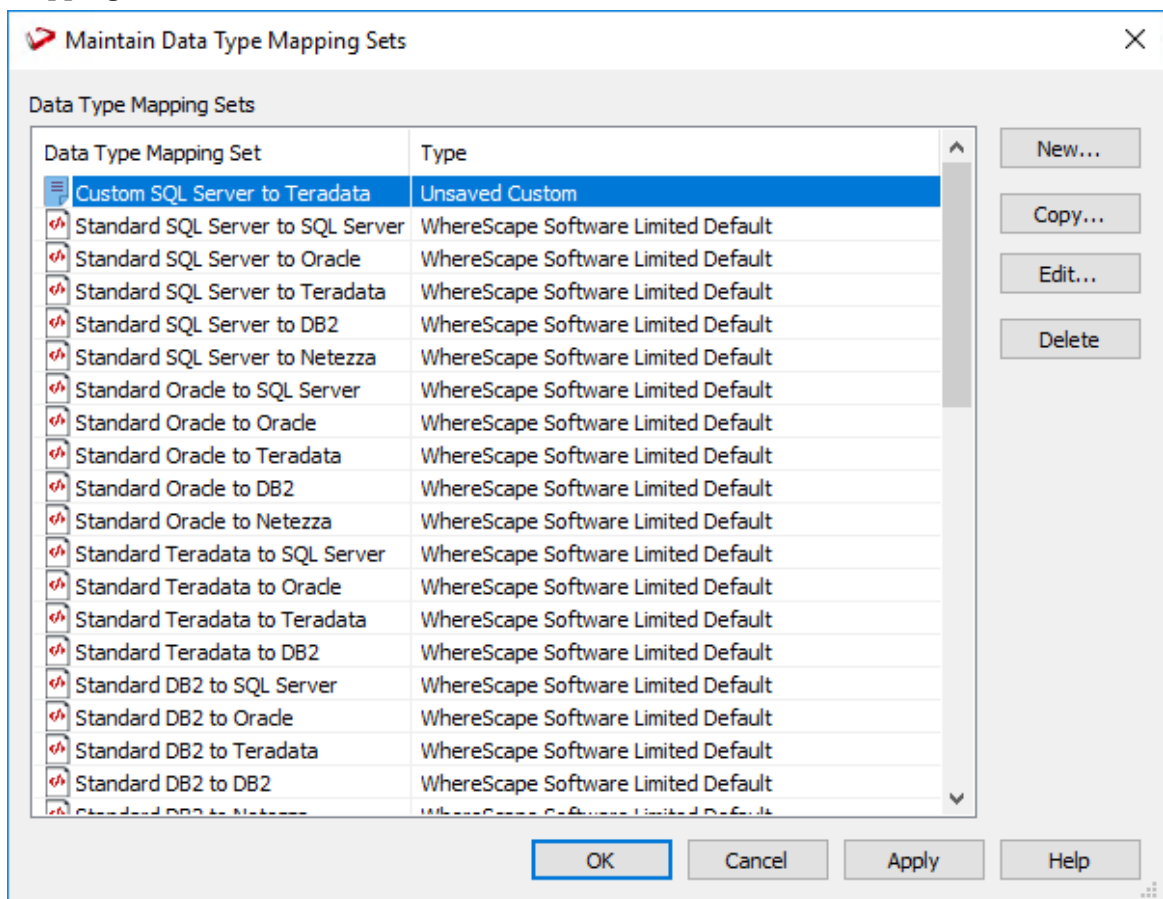


Notes: Use the buttons on the right side of the **Data Type Mapping Set** window to edit, copy or delete a selected parameter.

To move a parameter up or down in the list, select it and then click the **Move Up** or **Move Down** button.

The order of the data type mapping conversion parameters is significant because when loading a table, the procedure checks the data type mappings from top to bottom and stops when a data type and its parameters are correctly matched. A blank parameter means that it will match to anything.

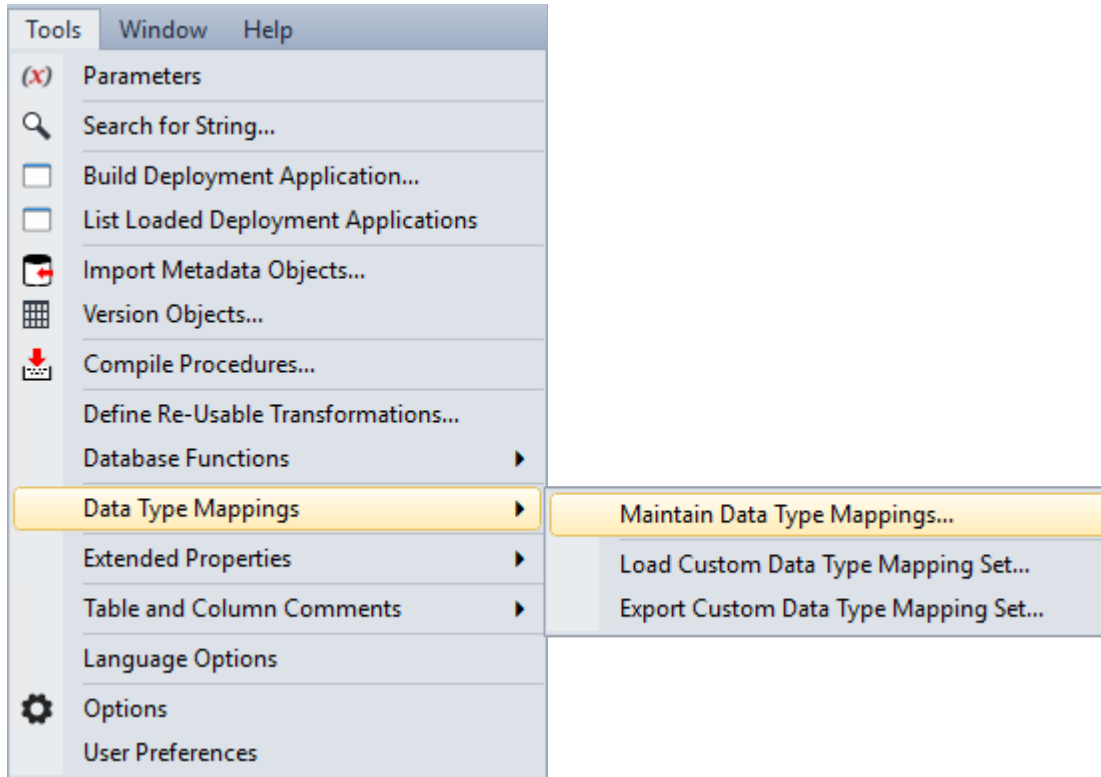
- Click **OK** when you're done. The new data type mapping set is listed in the **Maintain Data Mapping Sets** window.



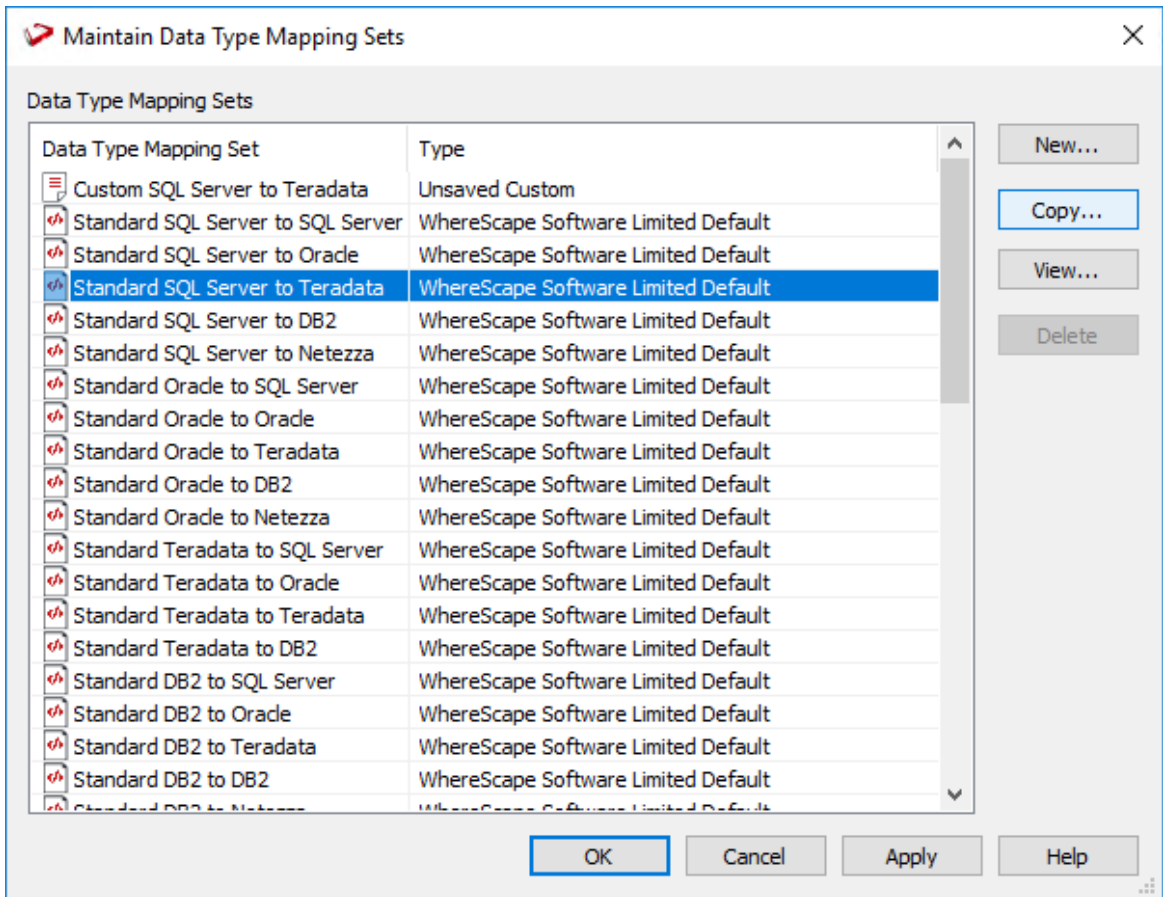
- Click **OK** or **Apply** to save the new data type mapping set.

COPYING A DATA TYPE MAPPING SET

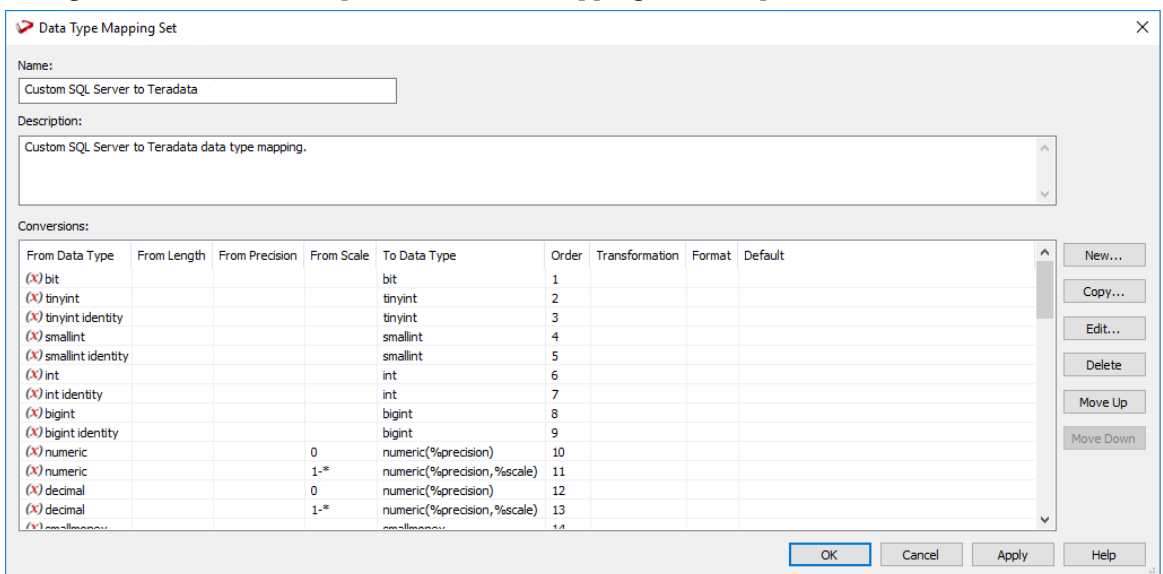
To copy an existing data type mapping set, select **Tools>Data Type Mappings>Maintain Data Type Mappings...**



- 1 Select the Mapping Set to be copied from the list in the **Maintain Data Type Mapping Sets** window and then click the **Copy** button.



- 2 Change the name and description of the new Mapping Set as required.

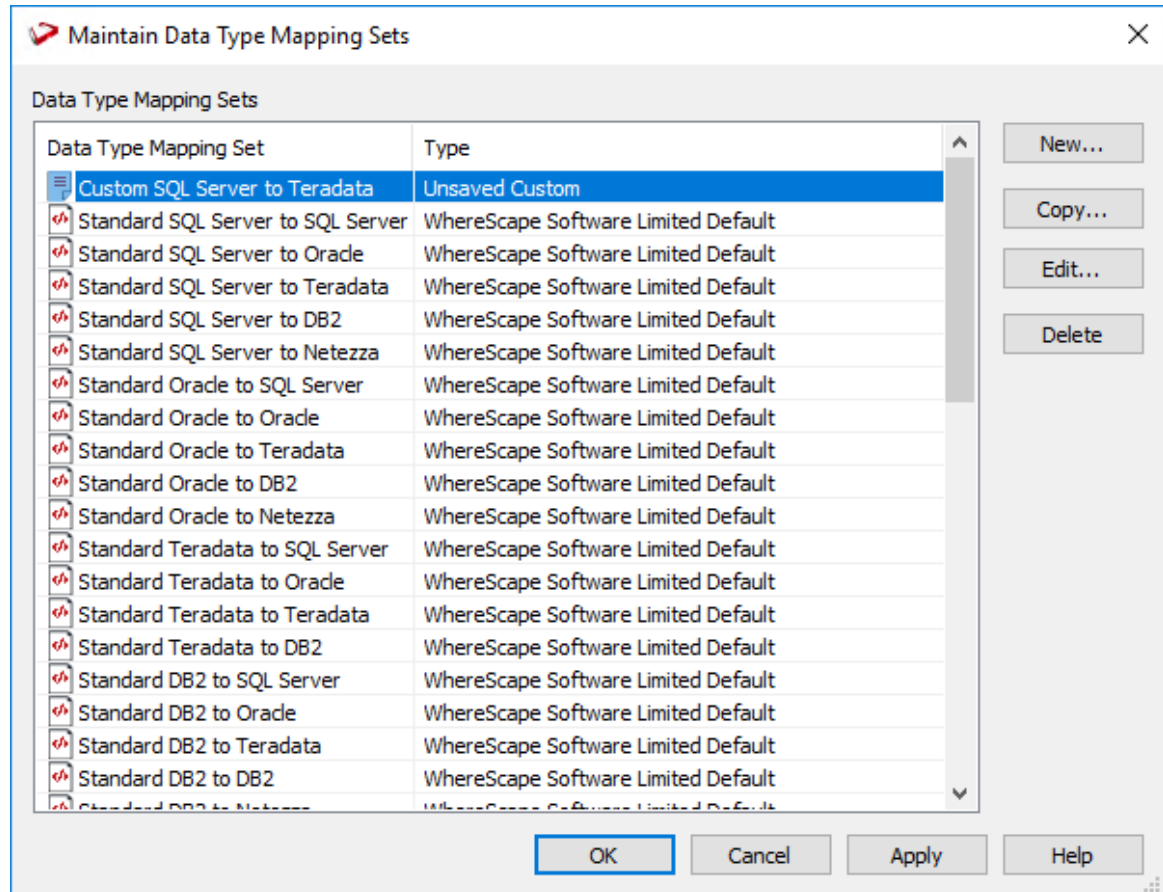


- 3 Use the buttons on the right side of the **Data Type Mapping Set** window to add new conversion parameters, edit, copy or delete a selected parameter.

Notes: To move a parameter up or down in the list, select it and then click the **Move Up** or **Move Down** button.

The order of the data type mapping conversion parameters is significant because when loading a table, the procedure checks the data type mappings from top to bottom and stops when a data type and its parameters are correctly matched. A blank parameter means that it will match to anything.

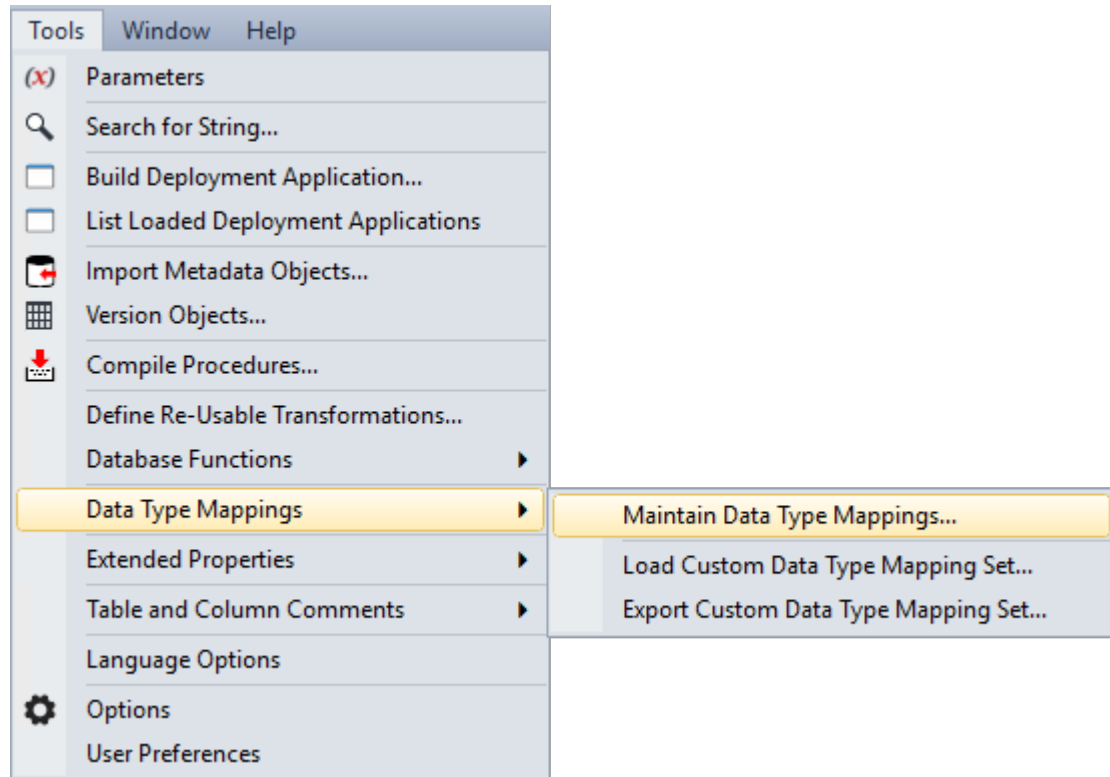
- 4 Click **OK** when you're done. The new data type mapping set is listed in the **Maintain Data Mapping Sets** window.



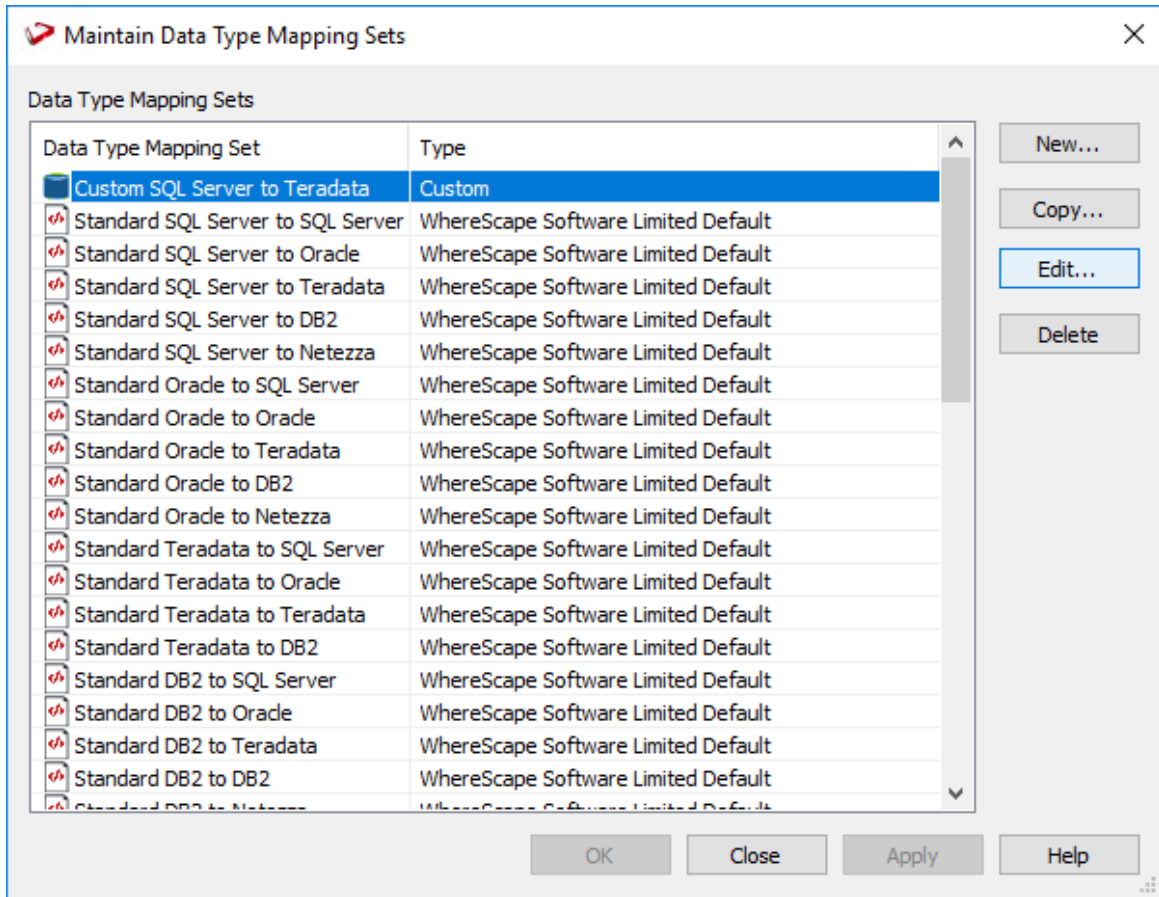
- 5 Click **OK** or **Apply** to save the new data type mapping set.

EDITING A DATA TYPE MAPPING SET

To edit a data type mapping set, select **Tools>Data Type Mappings>Maintain Data Type Mappings**.

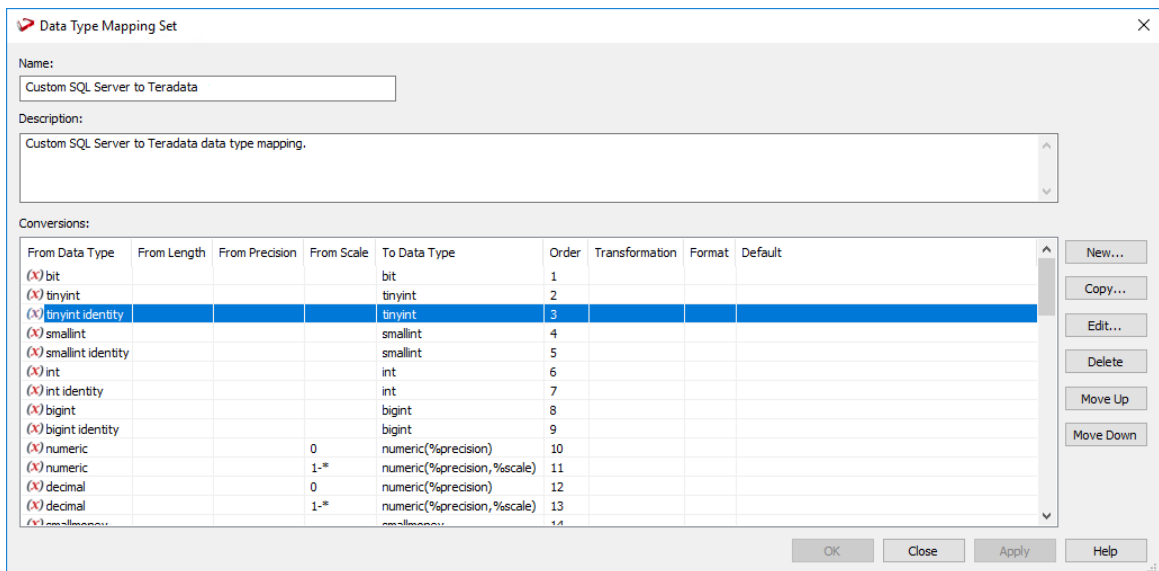


- 1 Select the Mapping Set to be modified from the list in the **Maintain Data Type Mapping Sets** window and then click the **Edit** button.



Note: Standard Data Type Mapping sets are not editable, they can only be viewed.

- 2 In the **Data Type Mapping Set** window, select the parameter you want to change from the **Conversions** pane. Use the scroll bar to see all the defined parameters.



- 3 Use the buttons on the right side of the **Data Type Mapping Set** window to add new conversion parameters, edit, copy or delete a selected parameter. These buttons are not available for standard mapping sets. Only **user defined** mapping sets are editable.

Note: To move a parameter up or down in the list, select it and then click the **Move Up** or **Move Down** button.

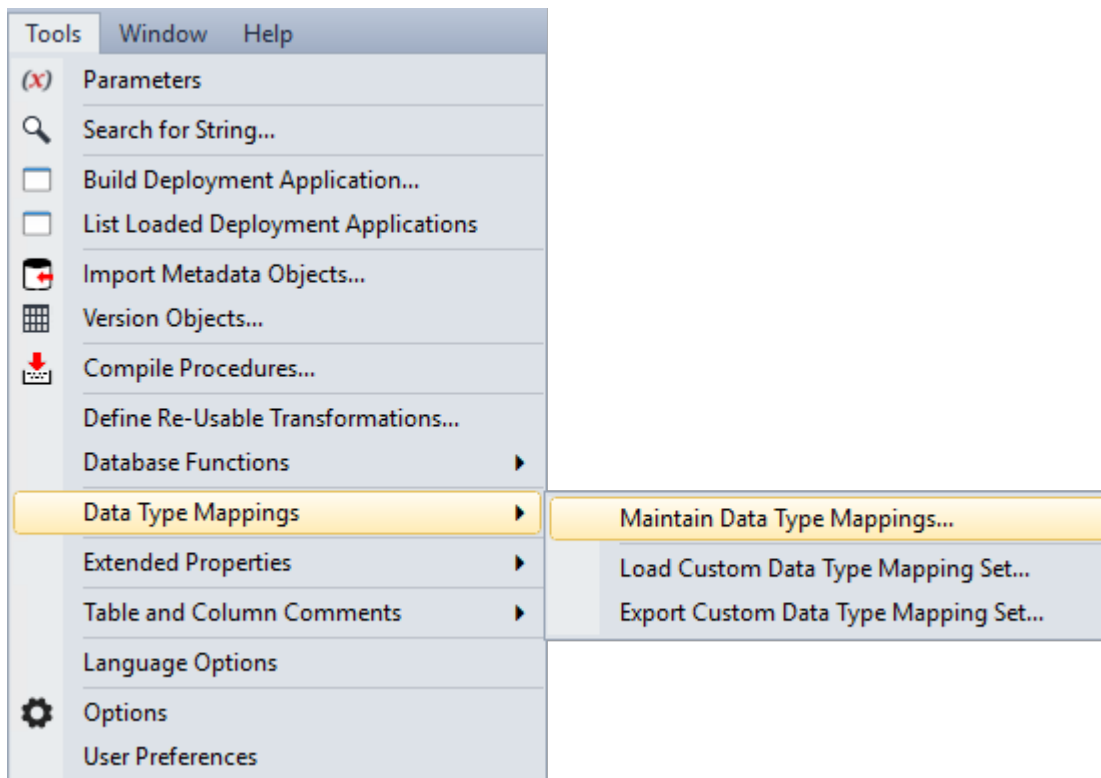
The order of the data type mapping conversion parameters is significant because when loading a table, the procedure checks the data type mappings from top to bottom and stops when a data type and its parameters are correctly matched. A blank parameter means that it will match to anything.

- 4 Click **OK** when you're done.

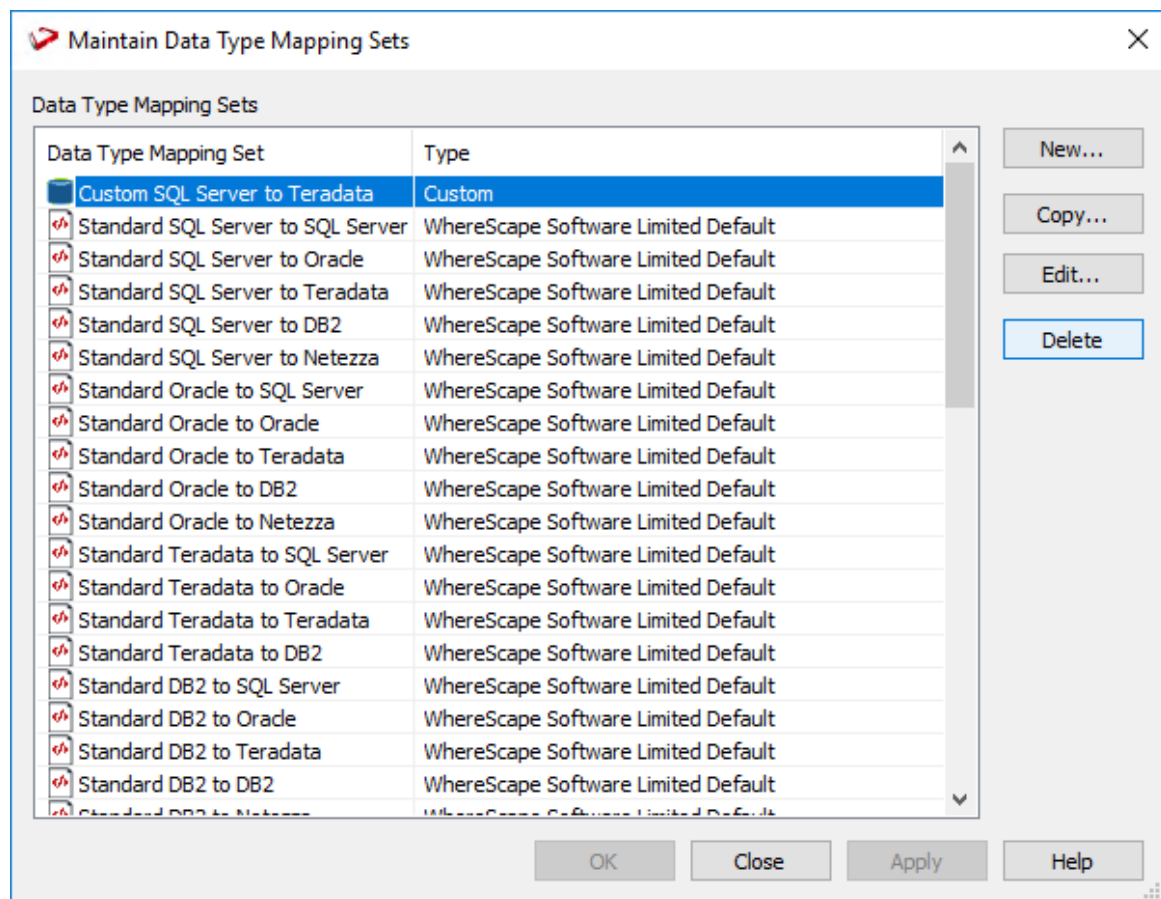
See previous sections for details on adding or copying a new data type mapping set.

DELETING A DATA TYPE MAPPING SET

To delete a data type mapping set, select **Tools>Data Type Mappings>Maintain Data Type Mappings...**



Select the Mapping Set to be deleted and click the **Delete** button.

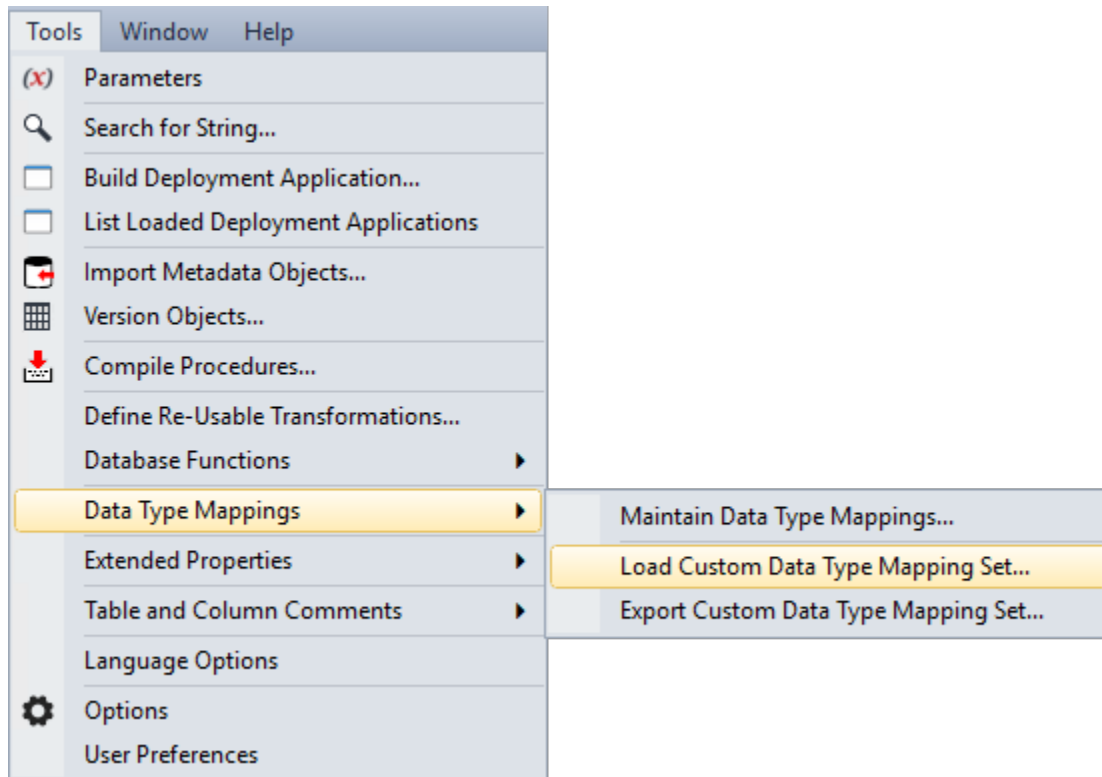


Note: The **Delete** button is disabled for standard data type mapping sets.

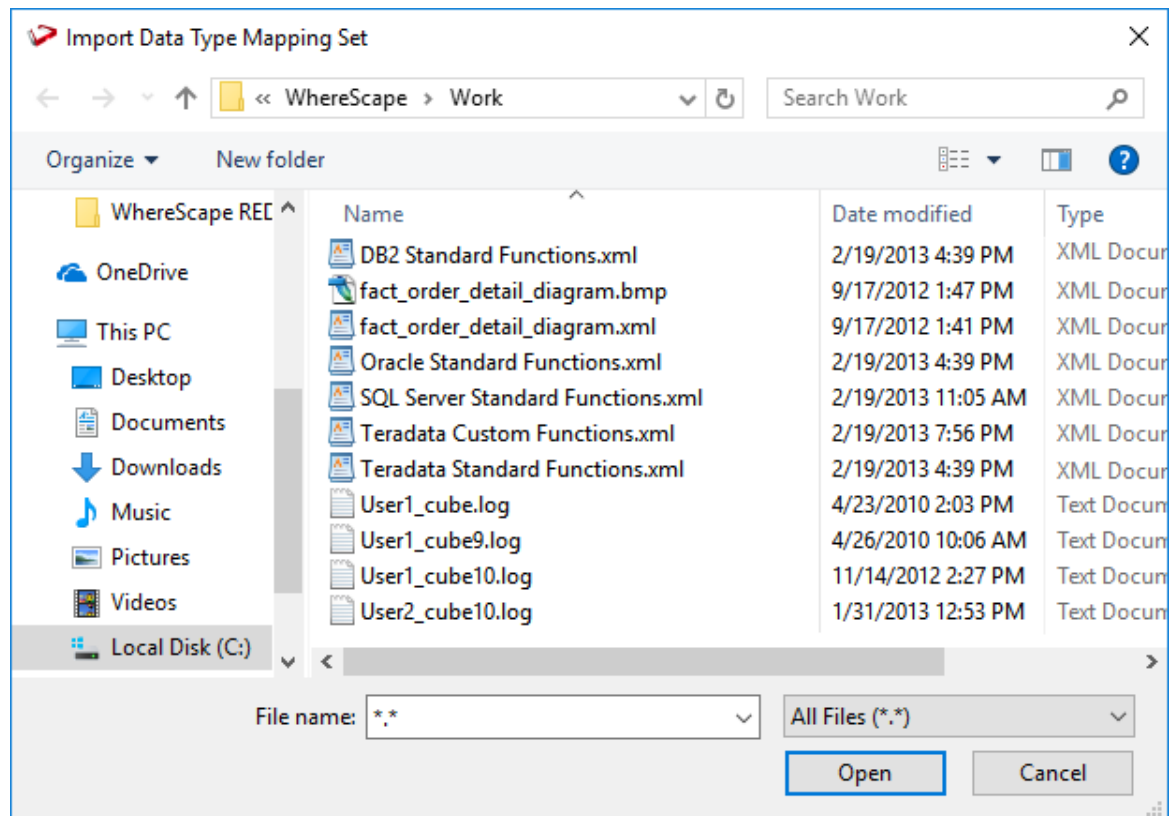
LOADING CUSTOM DATA TYPE MAPPING SETS

The **Load Custom Data Type Mapping Set** menu option allows you to load a custom data type mapping set from an XML file into the metadata repository.

To load a data type mapping set, select **Tools>Data Type Mappings>Load Custom Data Type Mapping Set**.



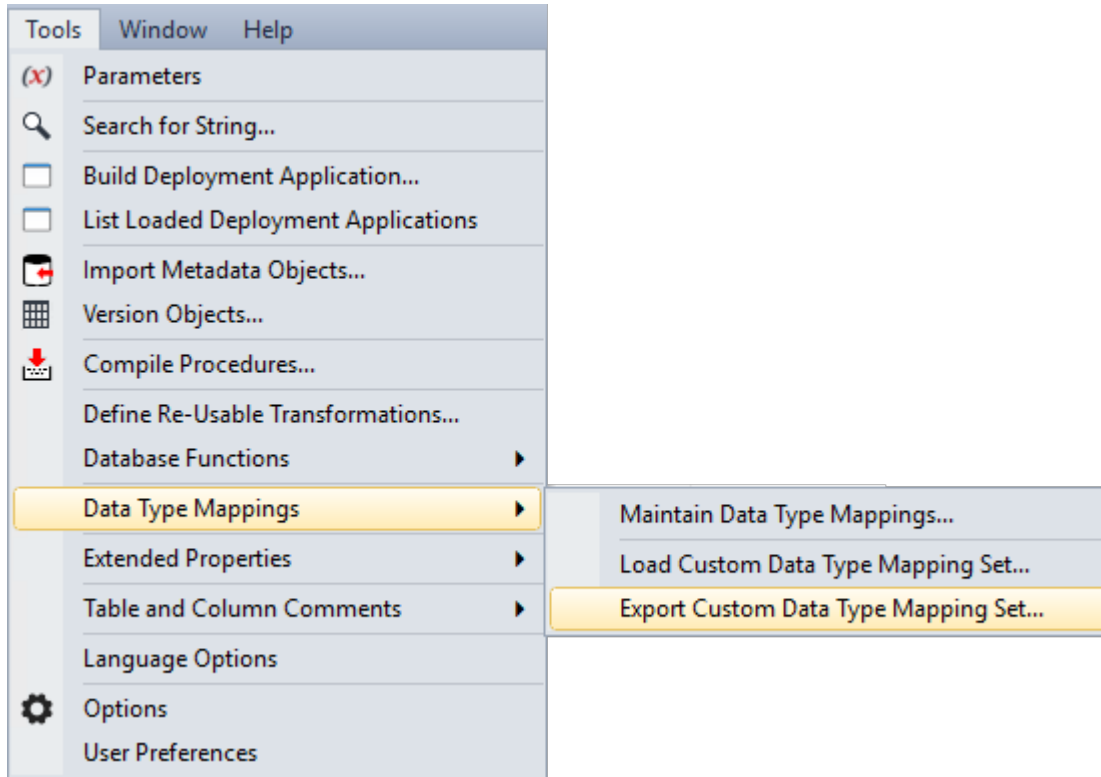
The following window is displayed. Select the xml file to load the data type mappings. By default RED expects the xml files to be in **ProgramData\WhereScape\Work**



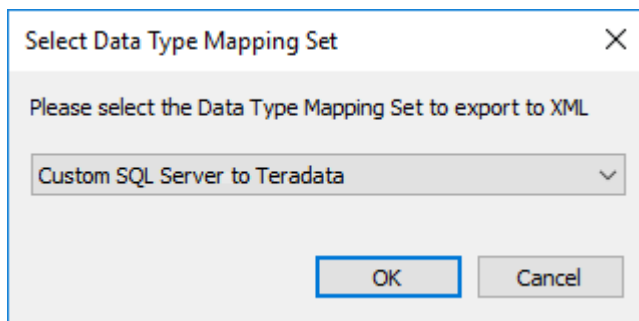
EXPORTING CUSTOM DATA TYPE MAPPING SETS

The **Export Custom Data Type Mapping Set** menu option allows you to export a custom data type mapping set from the metadata repository to an XML file.

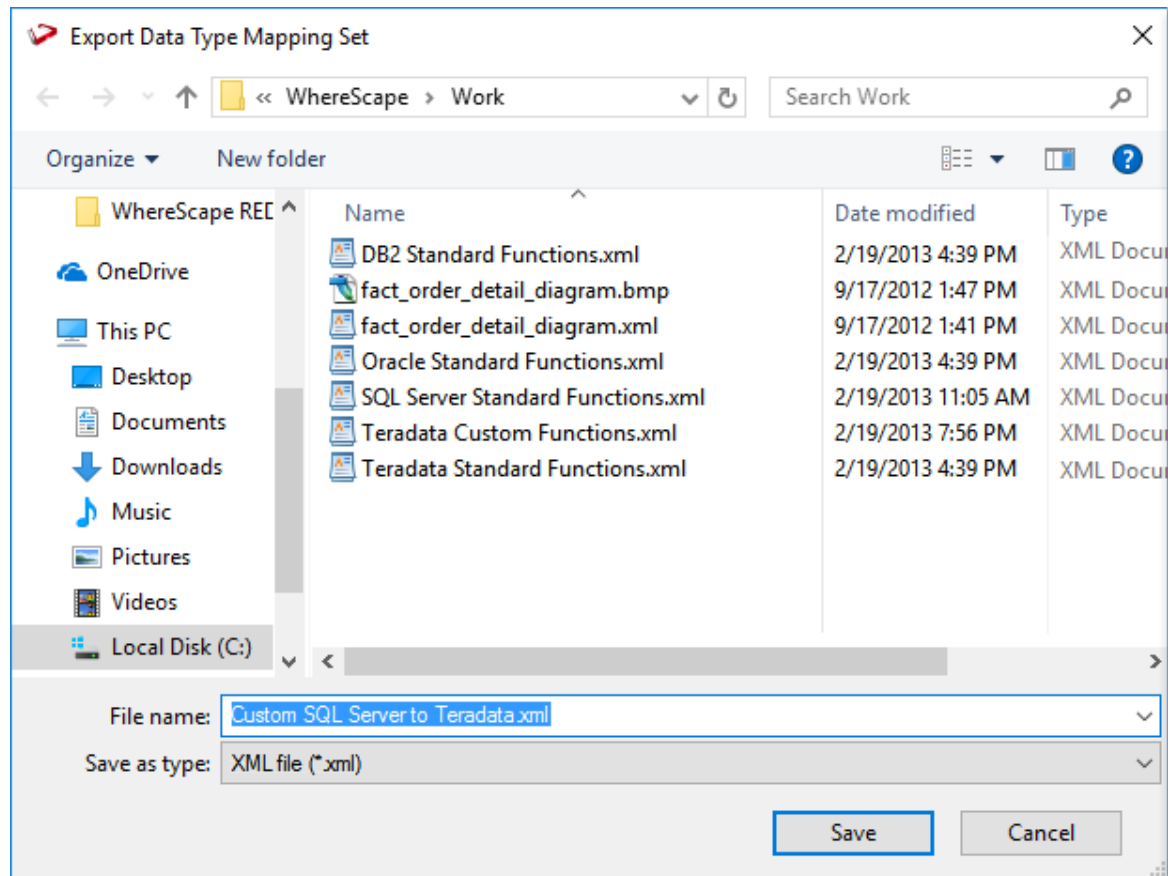
To export a data type mapping set, select **Tools>Data Type Mappings>Export Custom Data Type Mapping Set**.



Select the data type mapping set to export from the drop-down list. Click **OK**.



By default, RED exports the xml file to **ProgramData\WhereScape\Work**, but this can be changed. Change the File name if necessary and click **Save**.



CUSTOM DATA TYPE MAPPING EXAMPLES

WhereScape RED allows you to create **Custom** Data Type Mapping Sets. These gives you the ability to automatically change the data type of any column or to add column transformations, when dragging and dropping new load tables.

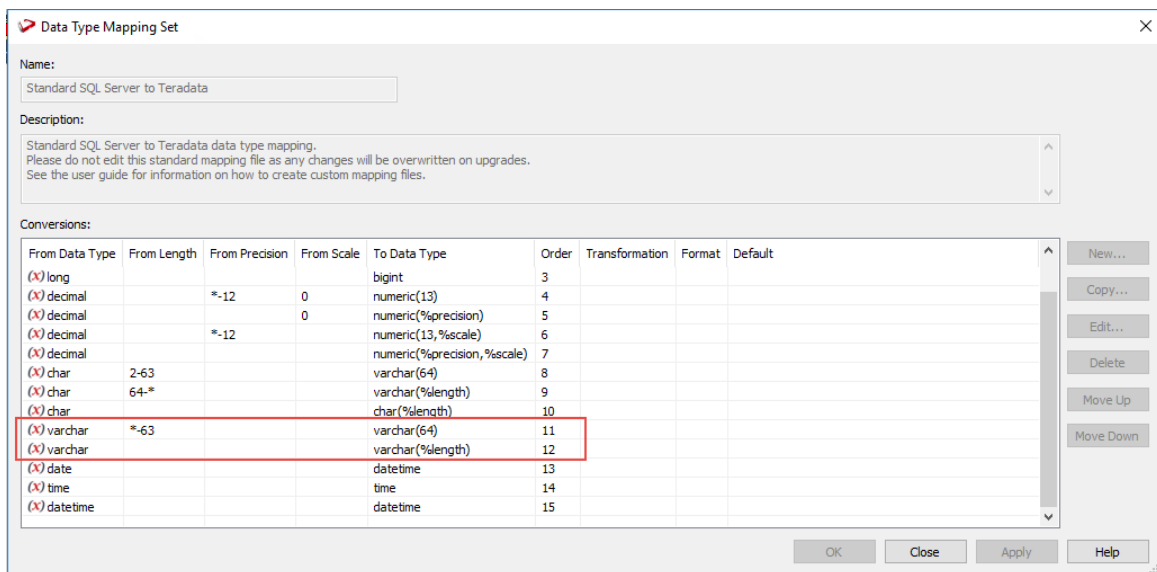
The examples in this topic demonstrate how **Custom** Data Type Mapping Sets can be configured using the following variables:

- %length
- %scale
- %precision
- %table_name
- %column_name
- %format
- %prompt

%length

In the example below, when converting a varchar in a file to Teradata format, we follow the following steps in the given order:

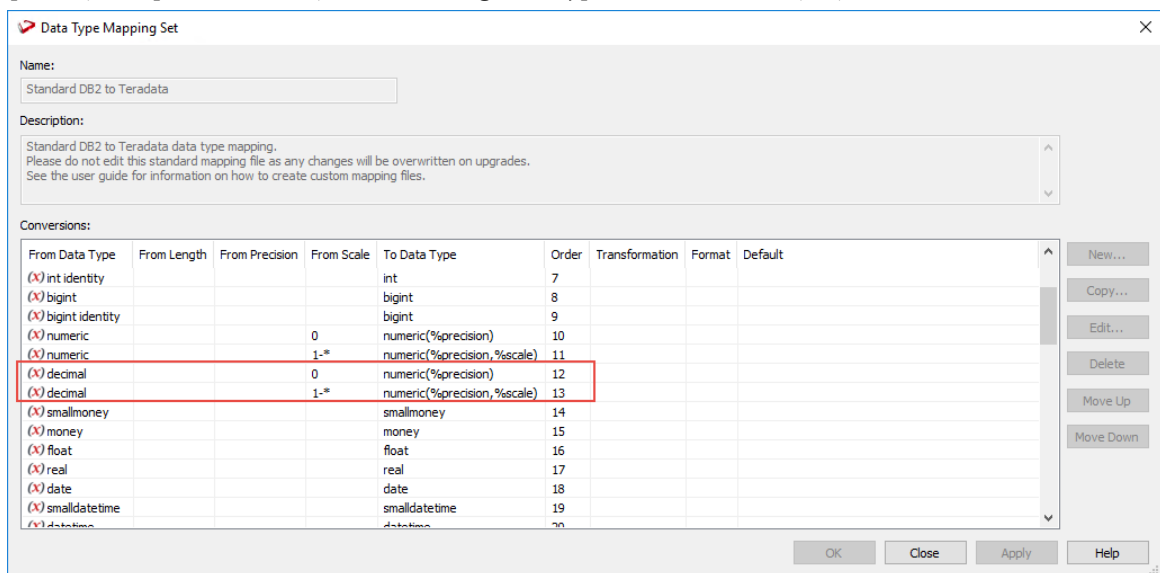
- If the varchar is of a length less than or equal to 63, the **data type** will become varchar(64).
- If the first step was NOT applied, i.e. the varchar is of a length greater than 63, then the **data type** will become varchar(%length); where we substitute the length for the variable '%length'. Thus if the varchar is of length 64 then the resulting data type will be varchar(64), but if the varchar is of length 123 then the resulting data type will be varchar(123).



%scale

In the example below, when converting a decimal in DB2 to Teradata format, we follow the following steps in the given order:

- If the decimal has a scale of zero, the **data type** will become decimal(%precision); where we substitute the number of digits in the number for the variable '%precision'. Thus if the decimal has 8 digits then the resulting data type will be decimal(8).
- If the first step was NOT applied, i.e. the decimal has a scale of 1 or greater, then the **data type** will become decimal(%precision,%scale); where we not only substitute the number of digits in the number for the variable '%precision', but we also substitute the scale for the variable '%scale'. Thus if the decimal is made up of 8 digits and has 3 digits after the decimal point (example 12345,678), the resulting data type will be decimal(8,3).

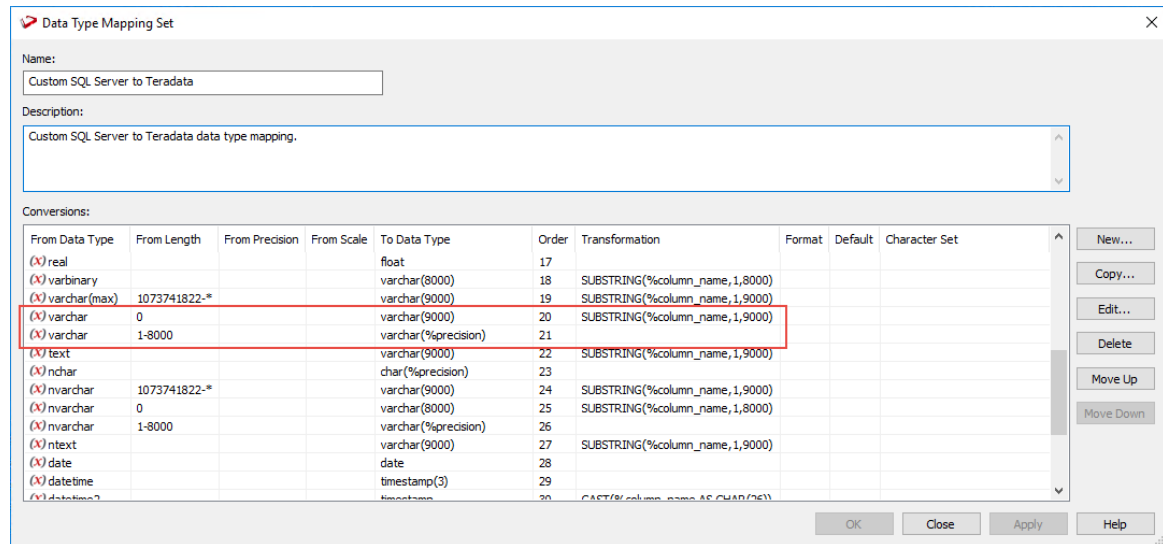


NOTE: The **Scale** is the number of digits to the right of the decimal point in a number.

%precision

In the example below, when converting a varchar in SQL Server to Teradata format, we follow the following steps in the given order:

- If the varchar has a length of 1 or greater, the **data type** will become varchar(%precision); where we substitute the number of digits in the varchar for the variable '%precision'. Thus if the varchar has a length of 14, the resulting data type will be varchar(14).
- If the varchar has a length of 0, the data type will be become varchar(9000) where the transformation will ensure that the correct length of the source string.

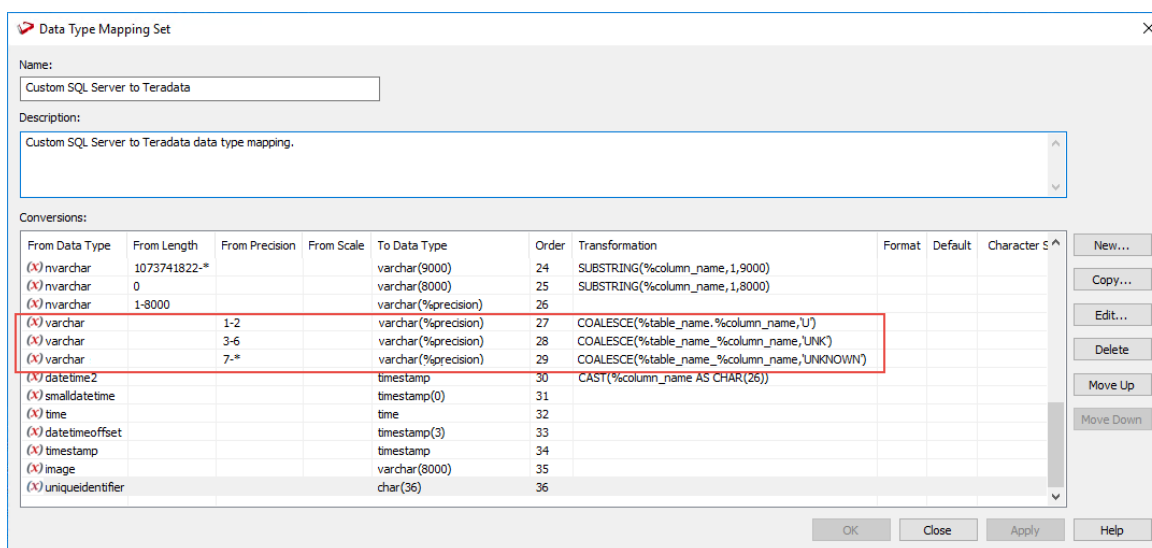


NOTE: The **Precision** is the total number of digits in a number.

%table_name and %column_name

In the example below, we use the following transformations to handle NULL for different lengths of varchar:

- If the varchar is 1 or 2 digits/chars long, the **data type** will become varchar(%precision); where we substitute the number of digits/chars in the varchar for the variable '%precision'. Secondly, the **value** of the column will become the column value (if it is not null), else it will become 'U'.
- If the varchar is 3-6 digits/chars long, the **data type** will become varchar(%precision); where we substitute the number of digits/chars in the varchar for the variable '%precision'. Secondly, the **value** of the column will become the column value (if it is not null), else it will become 'UNK'.
- If the varchar is 7 or more digits/chars long, the **data type** will become varchar(%precision); where we substitute the number of digits/chars in the varchar for the variable '%precision'. Secondly, the **value** of the column will become the column value (if it is not null), else it will become 'UNKNOWN'.



Data Type Mapping Set

Name: Custom SQL Server to Teradata

Description: Custom SQL Server to Teradata data type mapping.

Conversions:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation	Format	Default	Character S
(X) nvarchar	1073741822-*			varchar(9000)	24	SUBSTRING(%column_name, 1,9000)			
(X) nvarchar	0			varchar(8000)	25	SUBSTRING(%column_name, 1,8000)			
(X) nvarchar	1-8000			varchar(%precision)	26				
(X) varchar		1-2		varchar(%precision)	27	COALESCE(%table_name, %column_name, 'U')			
(X) varchar		3-6		varchar(%precision)	28	COALESCE(%table_name, %column_name, 'UNK')			
(X) varchar		7-*		varchar(%precision)	29	COALESCE(%table_name, %column_name, 'UNKNOWN')			
(X) datetime2				timestamp	30	CAST(%column_name AS CHAR(26))			
(X) smalldatetime				timestamp(0)	31				
(X) time				time	32				
(X) datetimeoffset				timestamp(3)	33				
(X) timestamp				timestamp	34				
(X) image				varchar(8000)	35				
(X) uniqueidentifier				char(36)	36				

Buttons: New..., Copy..., Edit..., Delete, Move Up, Move Down, OK, Close, Apply, Help

%format

In the example below, we use the following transformations to convert a certain character field to a date:

- If the varchar has a length of 1-10, the **data type** will become date and the **value** of the column will become the date 20131212 (a chosen date in the future).
- If the varchar has a length of 11, the **data type** will become date and the **value** of the column will use the transformation `TO_DATE(%table_name.%column_name,%format)`; where we substitute 'YYYYMMDD' for the variable '%format'. Thus the value of the column will be converted to a date of format 'YYYYMMDD'.
- If the varchar has a length of 12 or greater, the **data type** will become date and the **value** of the column will become the date 20131212 (a chosen date in the future).

Data Type Mapping Set

Name: Standard File to Teradata

Description: Standard File to Teradata data type mapping. Please do not edit this standard mapping file as any changes will be overwritten on upgrades. See the user guide for information on how to create custom mapping files.

Conversions:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation	Format	Default	Character S
(X) long				bigint	3				
(X) decimal		*-12	0	numeric(13)	4				
(X) decimal			0	numeric(%precision)	5				
(X) decimal		*-12		numeric(13,%scale)	6				
(X) decimal				numeric(%precision,%scale)	7				
(X) char	2-63			varchar(64)	8				
(X) char	64-*			varchar(%length)	9				
(X) char				char(%length)	10				
(X) varchar	*-63			varchar(64)	11				
(X) varchar				varchar(%length)	12				
(X) date				date	13	%format			
(X) time				time	14				
(X) datetime				timestamp	15	%format			

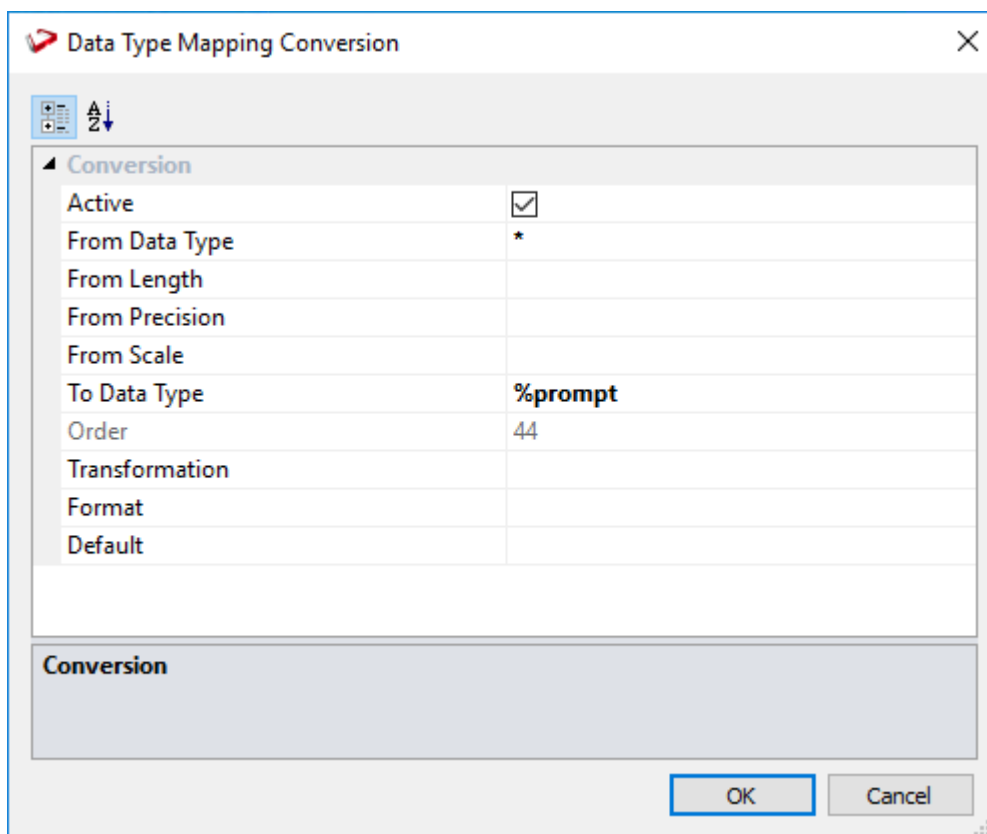
Buttons: New..., Copy..., Edit..., Delete, Move Up, Move Down, OK, Close, Apply, Help

%prompt

In the example below we use %prompt to help the user to define a mapping for an unknown datatype that is not already mapped in the previous mapping rules.

This placeholder must be used with a **custom** Data Type mapping set, as described in the following steps:

- Create a new custom set or copy from an existing set.
- Create a new Data Type mapping with a **From Data Type** of star (*) and a **To Data Type** of **%prompt**. Click OK to save the New Data Type Mapping to the Custom set.

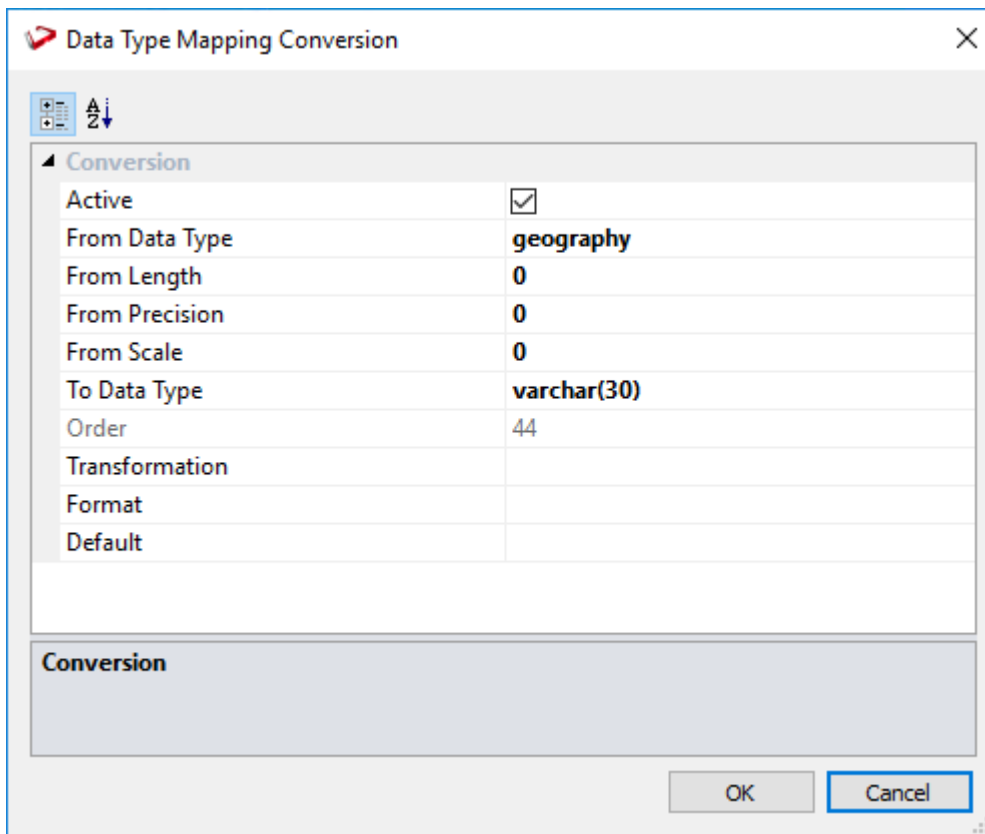


Conversion	
Active	<input checked="" type="checkbox"/>
From Data Type	*
From Length	
From Precision	
From Scale	
To Data Type	%prompt
Order	44
Transformation	
Format	
Default	

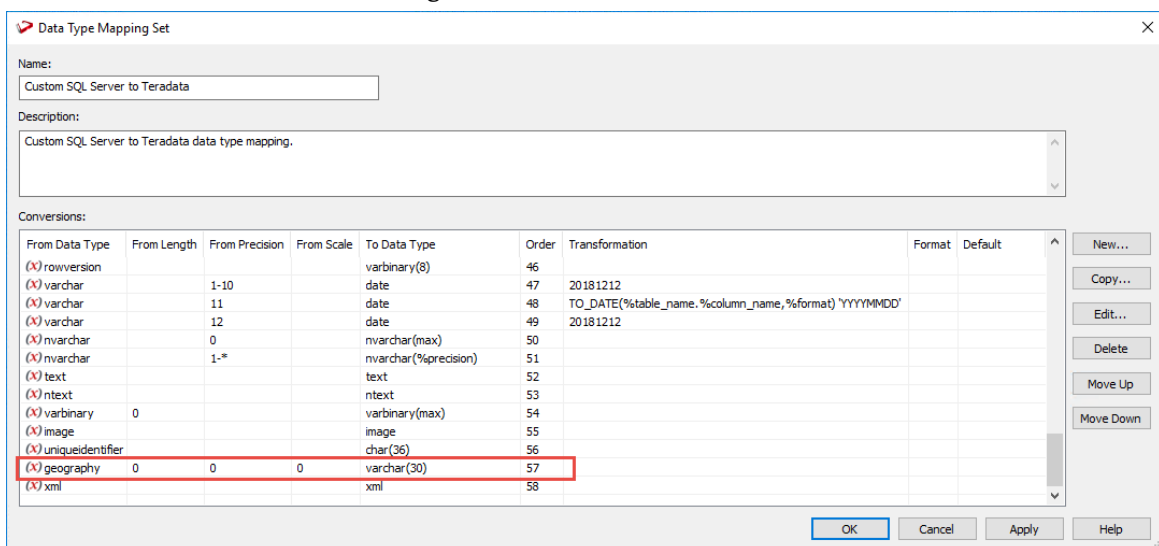
Conversion

OK Cancel

- When browsing a connection to Load a table, set the **Data Type Mapping Set** to the new **Custom SQL Server to Teradata** mapping set. This can be set on the List Sources Tables Connection window or on the connection Properties window.
- As the table is dragged and dropped to the middle pane, RED will prompt to have the new data type mapping defined.
- In the example below, just before loading the table, users can map the unknown **geography** SQL Server data type mapping to a **varchar(30)** in Teradata.



- Clicking **OK** or **Apply** in the **Data Type Mapping Set** window, saves the parameter to the custom set that was used for loading the table.



CHAPTER 42

COLUMN CONTEXT MENU

To view the column context menu, click on an object in the left pane to display the columns in the middle pane. When positioned on a column in the middle pane, right-click on the column to bring up the menu.

The screenshot shows a window titled "stage_customer_code_merge Columns" with a table of columns. A context menu is open over the "state" column. The table has the following data:

Column Name	Display Name	Data Type	Format	Nulls	Nu...	Ad...	Att...	Ke...
code	code	numeric(6)	#,##0	Y	Y	Y	N	
name	name	varchar(45)		Y	N	N	Y	
cat_code	cat_code	varchar(20)		Y	N	N	Y	
grp_code	grp_code	varchar(20)		Y	N	N	Y	
subgrp	subgrp	char(20)		Y	N	N	Y	
dss_record_source	dss record source	varchar(255)		Y	N	N	Y	
dss_load_date	dss load date	datetime		Y	N	N	Y	
address	address	varchar(40)						
city	city	varchar(30)						
state	state	varchar(2)						
hk_h		char(32)		N	N	N	N	h
dss_cl		char(32)		N	N	N	N	c
dss_co		datetime		Y	N	N	Y	
dss_u		datetime		Y	N	N	Y	

The context menu is open over the "state" column and contains the following items:

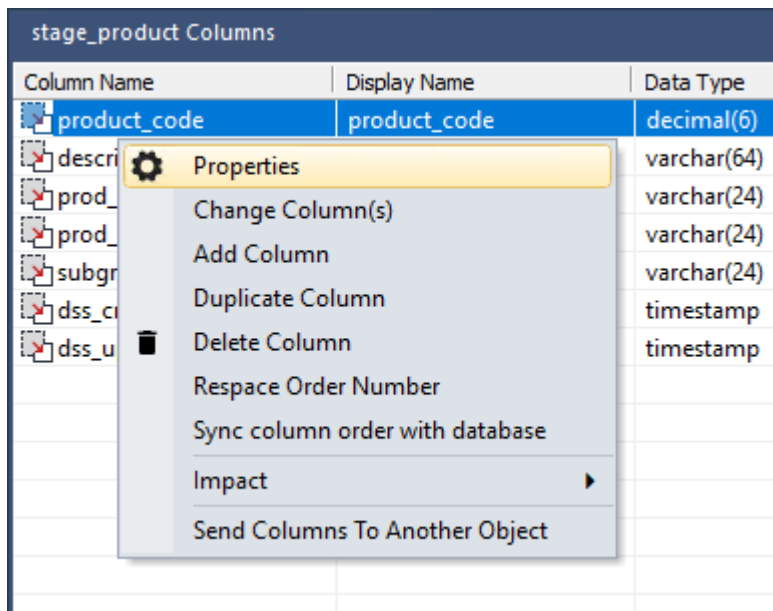
- Properties
- Change Column(s)
- Add Column
- Duplicate Column
- Delete Column
- Respace Order Number
- Sync column order with database
- Impact
- Send Columns To Another Object

IN THIS CHAPTER

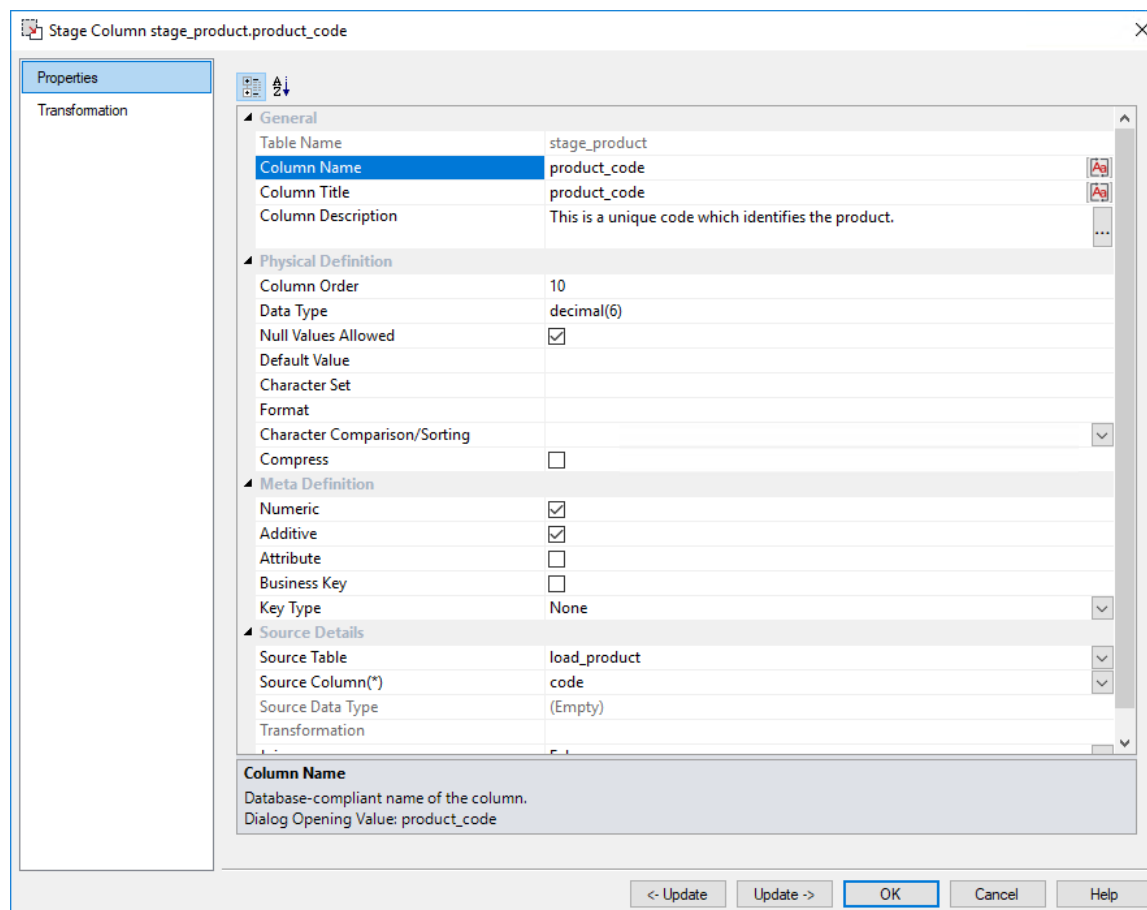
Properties.....	1020
Change Column(s)	1023
Add Column	1025
Duplicate Column	1026
Delete Column	1027
Re-space Order Number.....	1028
Impact	1029
Sync Column order with database	1030
Send Columns to Another Object	1031

PROPERTIES

To display the column Properties, right-click on a column in the middle pane and select **Properties**.




Edit any field as required and then click **OK** to close.



Warning: WhereScape RED does not support the following characters in **Column Names**:

- leading and trailing white spaces
- internal white spaces
- symbols other than #, \$ and _

If users attempt to enter any of the above characters in Column Names, the following dialog will be displayed, advising users to review changes made by RED to correct any unsupported column name characters:



Note: There is a variation in column Properties, depending on the object type.

For Dimension tables, see *Dimension Column Properties* (on page 316).

For Stage tables, see *Stage Table Column Properties* (on page 345).

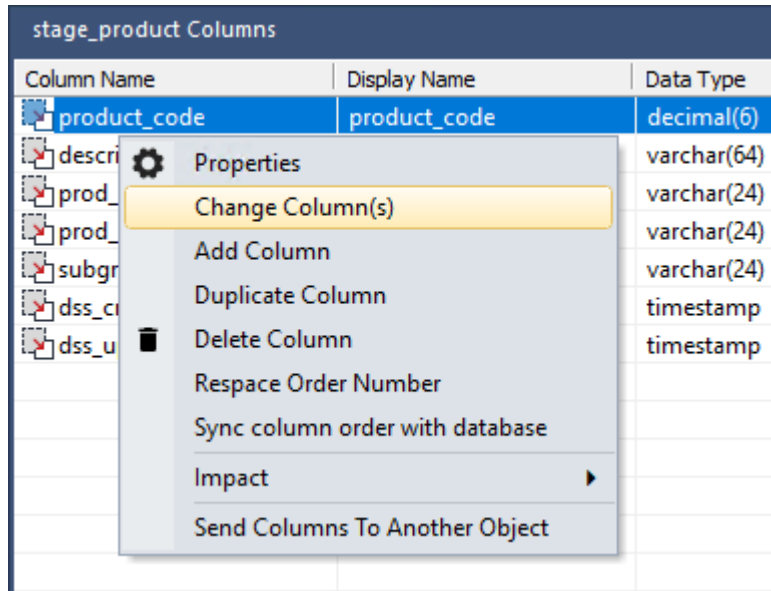
For Model tables, see *Model Column Properties* (see "*Model Table Column Properties*" on page 481).

For Data Store tables, see *Data Store Column Properties* (on page 382).

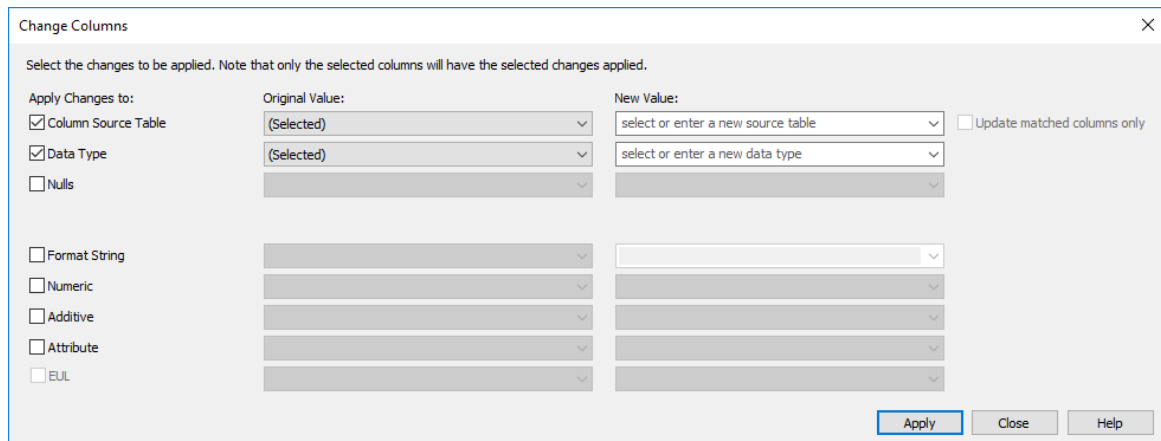
For EDW 3NF tables, see *EDW 3NF Column Properties* (see "*EDW 3NF Table Column Properties*" on page 406).

CHANGE COLUMN(S)

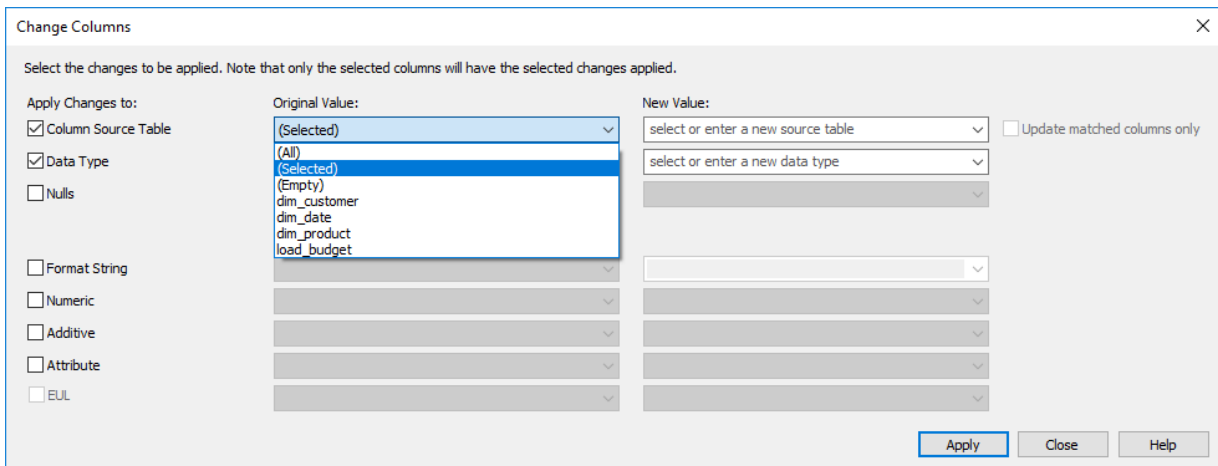
To change the properties for multiple columns, right-click on a column in the middle pane and select **Change Column(s)**.



To change a column property, you first need to select the relevant **check boxes** on the left. Each check box, when selected, allows you to change the value for that field in the column properties.



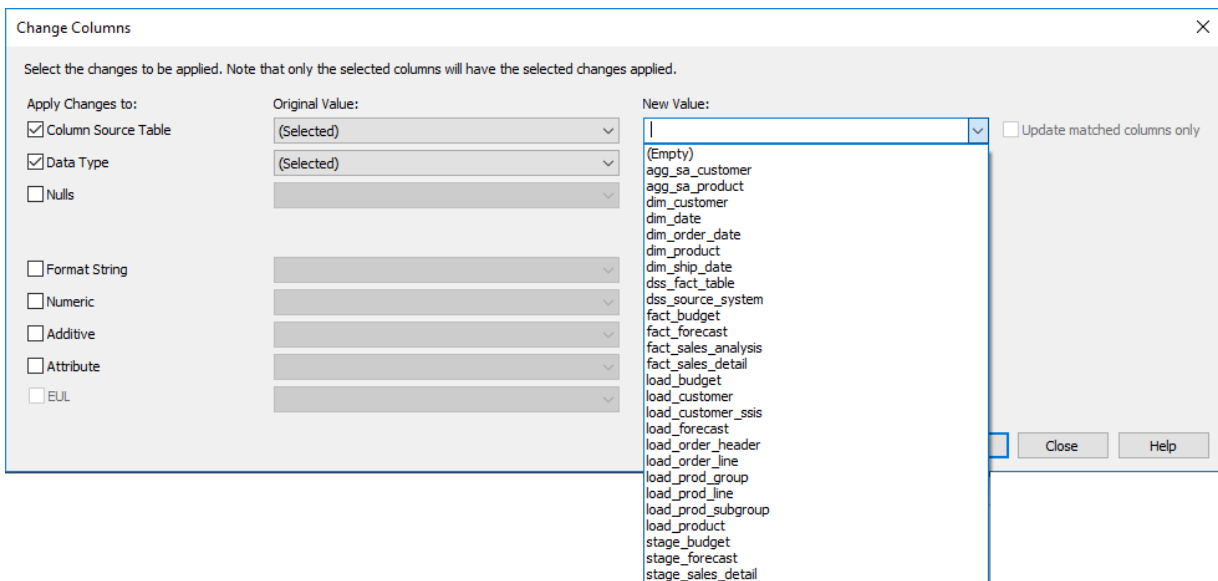
In the **Original Value** column, select the value/s to be changed.



- Choosing **(All)** will change the selected property for all of the columns in the table.
- Choosing **(Selected)** will change the selected property for the selected column in the table.
- Choosing **(Empty)** will change the selected property for all of the columns where that property field is empty. This option is only available if there is a column where this property is empty.
- Choosing one of the **other** options will change the selected property for all the columns in the table having that value.

Note: **(Selected)** is the default for the **Original Value** column.

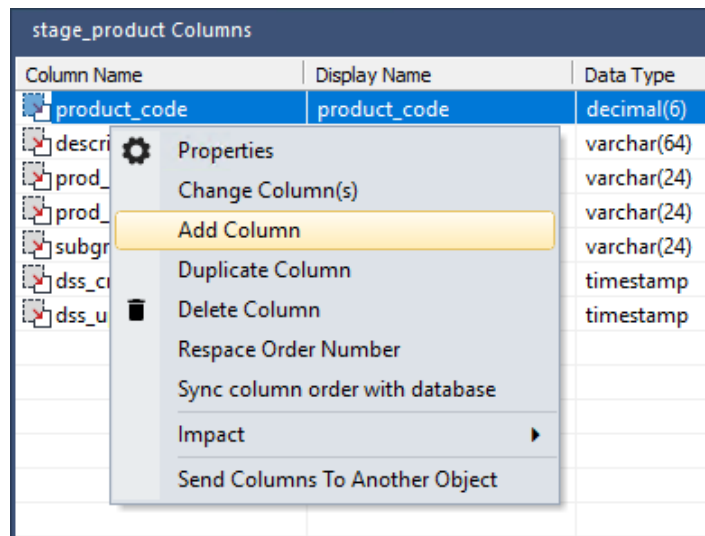
In the **New Value** column, select the new value to be assigned; or key in the new value.



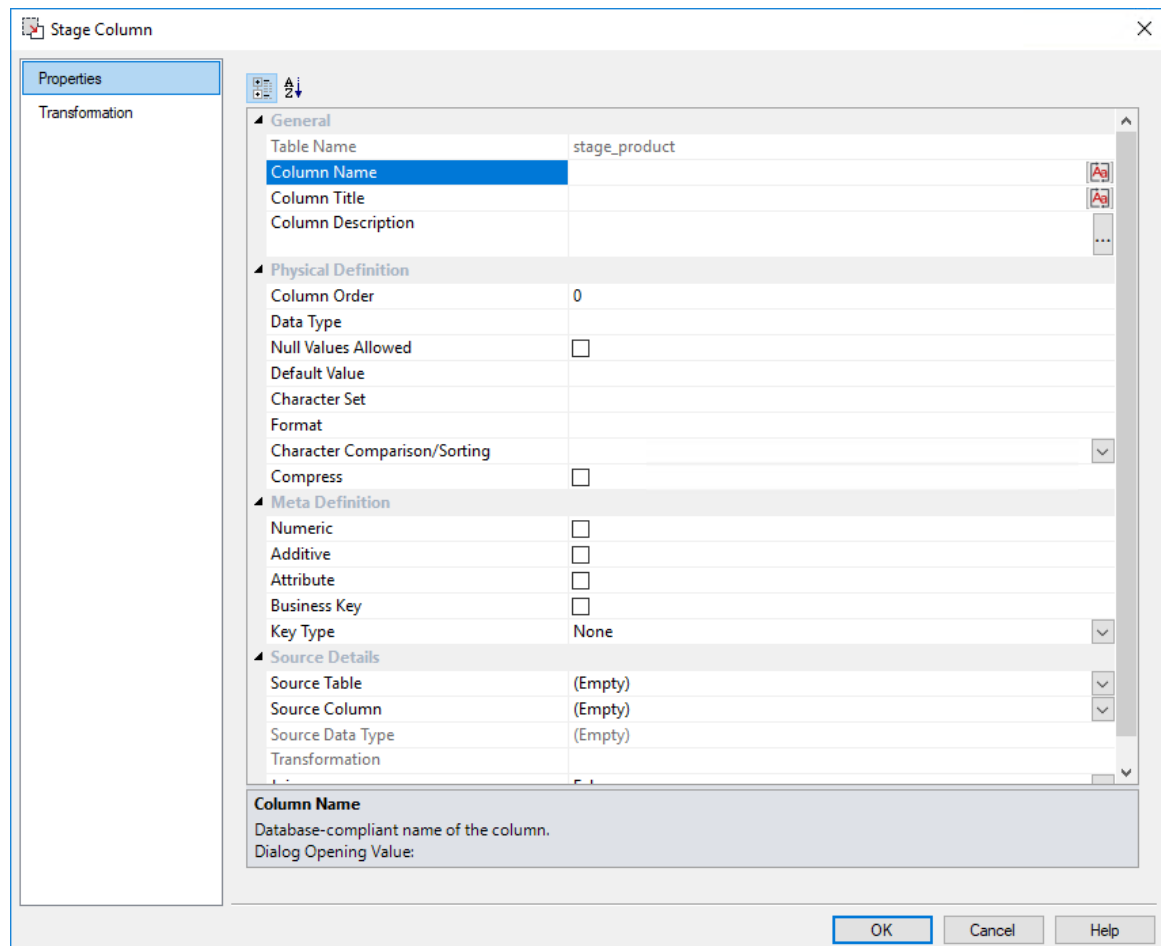
Note: When editing the property **Column Source Table**, selecting the **Update matched columns only** check-box will validate that the selected column name exists in the new source table. If it does not exist, then the update will not take place.

ADD COLUMN

To add a column, right-click on one of the columns in the middle pane and select **Add Column**.

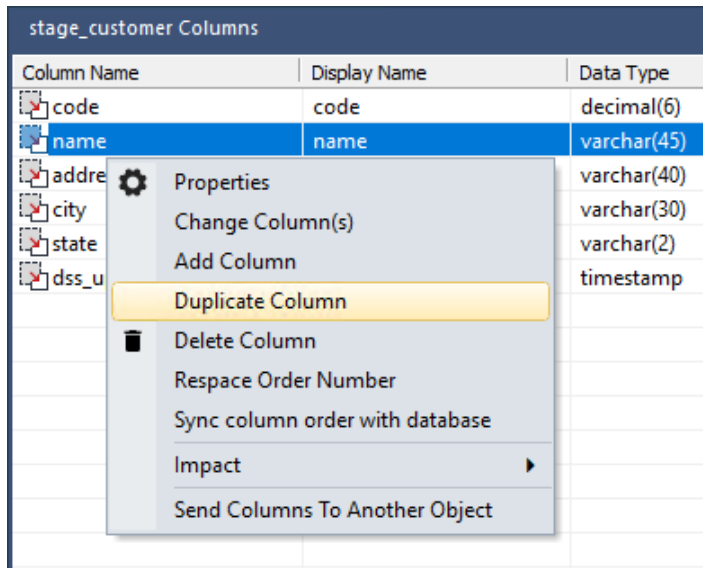


Enter the details to define a new column and click **OK**.

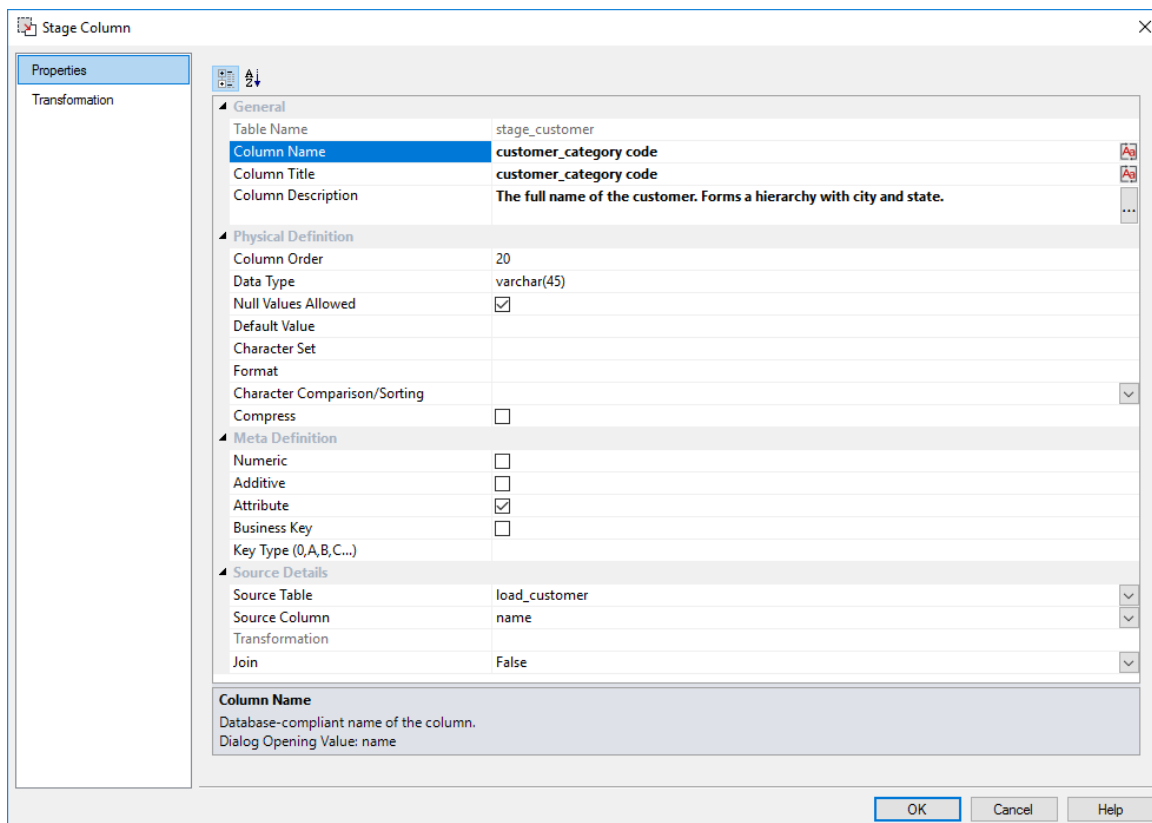


DUPLICATE COLUMN

To duplicate/copy a column, right-click on one of the columns in the middle pane and select **Duplicate Column**.

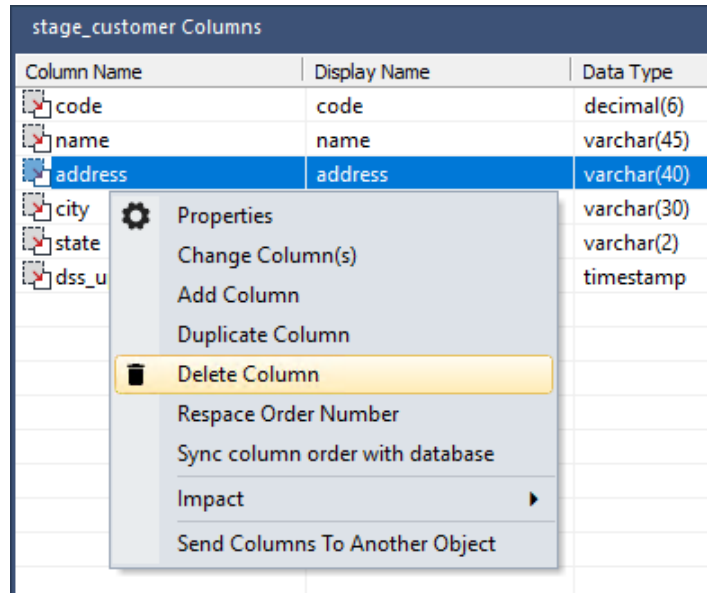


Change the **Column Name** and the **Column Title** and any other properties to define a new column and click **OK**.




DELETE COLUMN

To delete a column, right-click on one of the columns in the middle pane and select **Delete Column**.

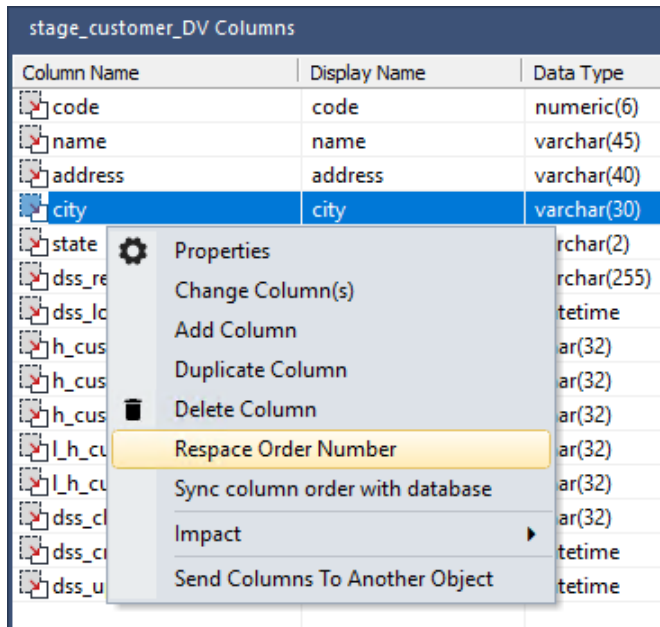


Click **Yes** to continue with the delete.

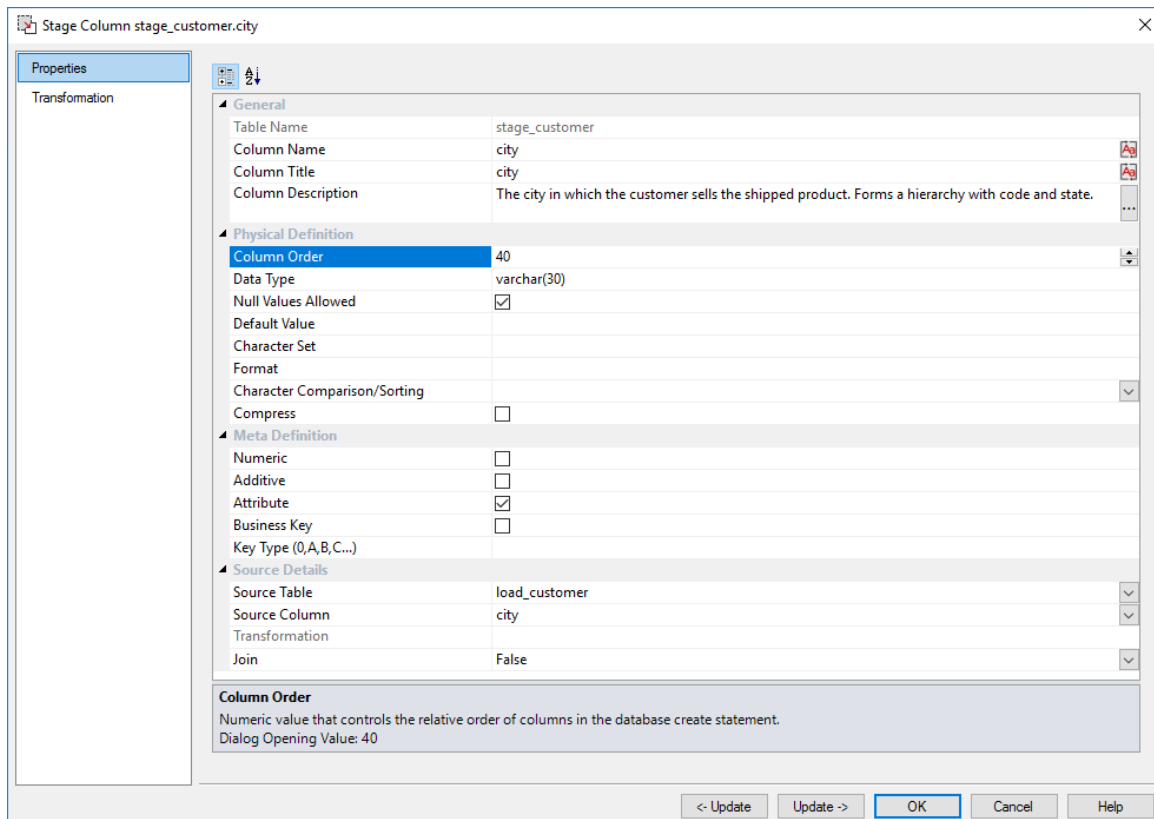


RE-SPACE ORDER NUMBER

To re-space the column order, right-click on any column in the middle pane and select **Respace Order Number**.

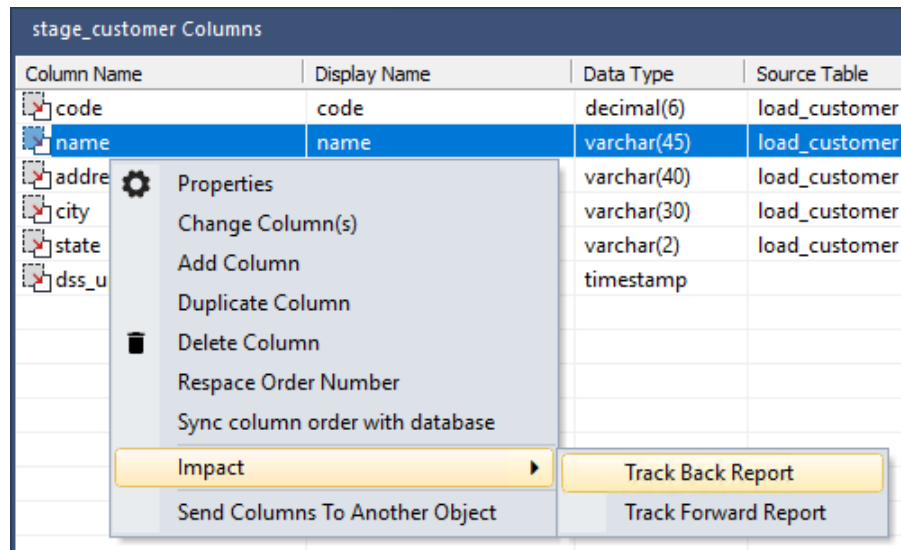


The **Column Order** number for each column will be adjusted so that the column order numbers are evenly spaced.

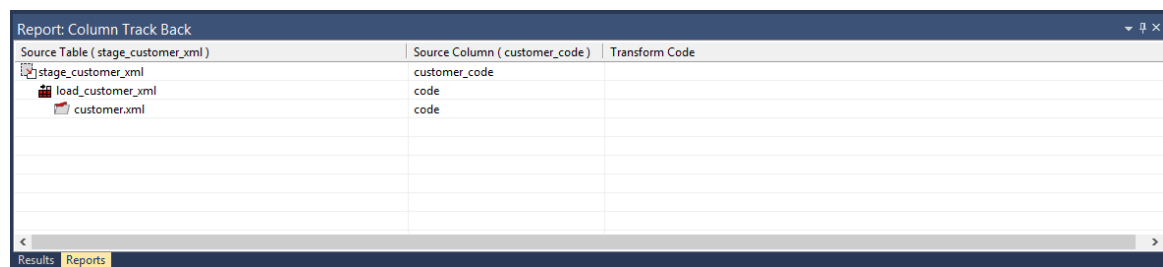


IMPACT

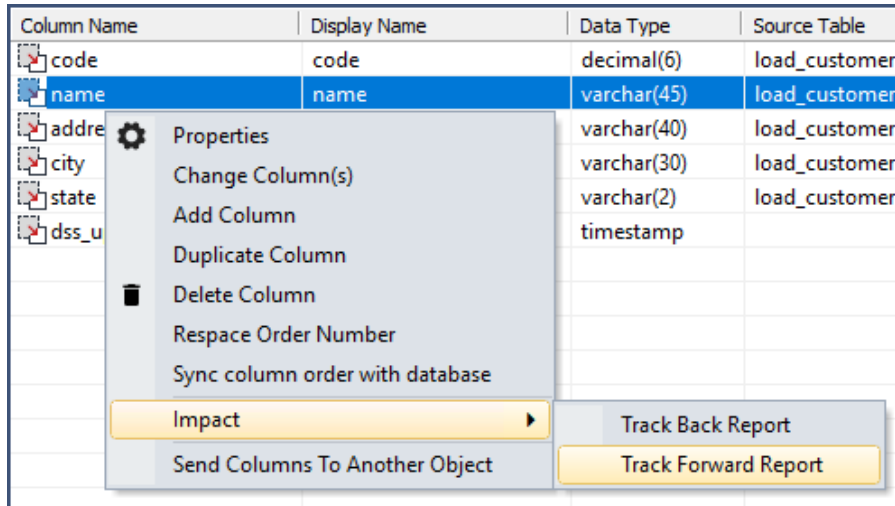
To display a Track Back Report, right-click on a column in the middle pane and select **Impact > Track Back Report**.



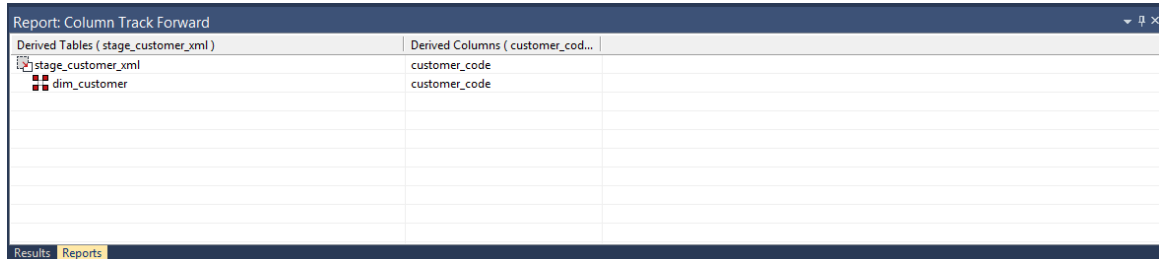
The report will be displayed in the bottom middle pane. This report lists the origins of the selected column. See **Track Back Report** (see "Column Track-Back" on page 795).



To display a Track Forward Report, right-click on a column in the middle pane and select **Impact** > **Track Forward Report**.



The report will be displayed in the bottom middle pane. This report lists the columns derived from the selected column. See **Track Forward Report** (see "**Column Track-Forward**" on page 797).

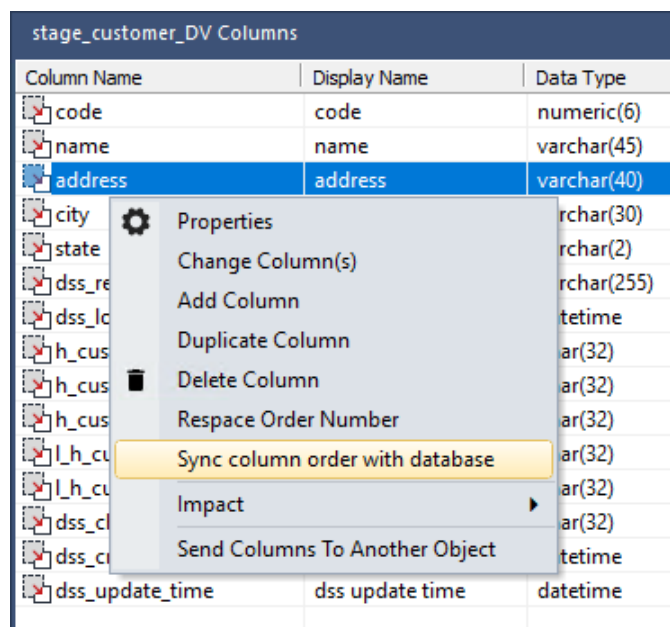


SYNC COLUMN ORDER WITH DATABASE

To synchronize the metadata's column order to match the same order in the physical table in the database:

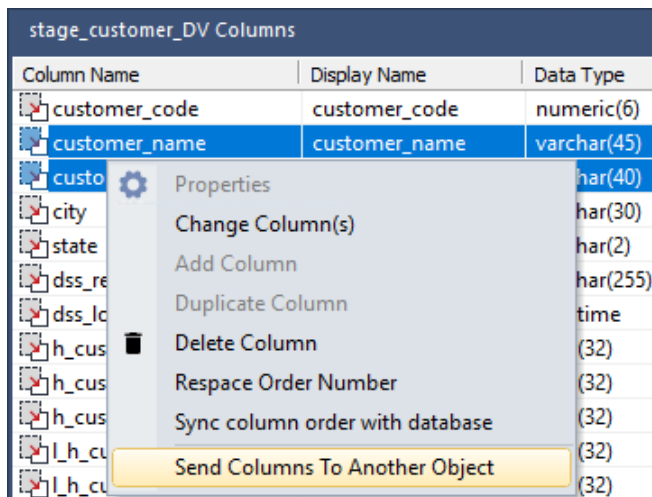
Right-click on one of the columns in the middle pane and select **Sync the column order with the database**.

This will reorder the metadata columns to match the column order in the database table.

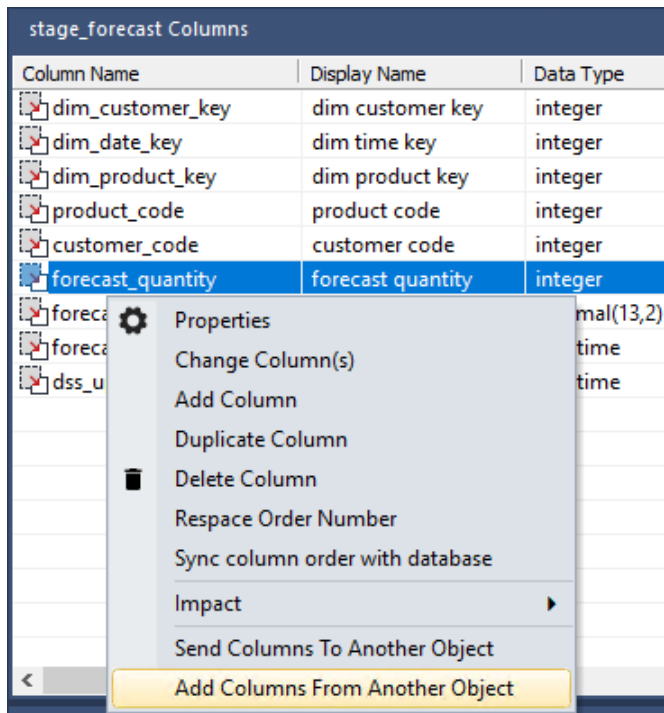


SEND COLUMNS TO ANOTHER OBJECT

To send/copy columns to another object, right-click on a column in the middle pane and select **Send Columns To Another Object**.



Click on the destination table in the left pane, then right-click in the middle pane and select **Add Columns From Another Object**.



The columns will be added to the destination table using the same functionality and settings as drag and drop.

Column Name	Display Name	Data Type	Source Table
dim_customer_key	dim customer key	integer	dim_customer
dim_date_key	dim time key	integer	dim_date
dim_product_key	dim product key	integer	dim_product
product_code	product code	integer	load_forecast
customer_name	customer_name	varchar(45)	stage_customer_DV
customer_address	customer_address	varchar(40)	stage_customer_DV
customer_code	customer code	integer	load_forecast
forecast_quantity	forecast quantity	integer	load_forecast
forecast_sales_value	forecast sales value	decimal(13,2)	load_forecast
forecast_date	forecast date	datetime	load_forecast
dss_update_time	dss update time	datetime	

CHAPTER 43

DATABASE FUNCTIONS

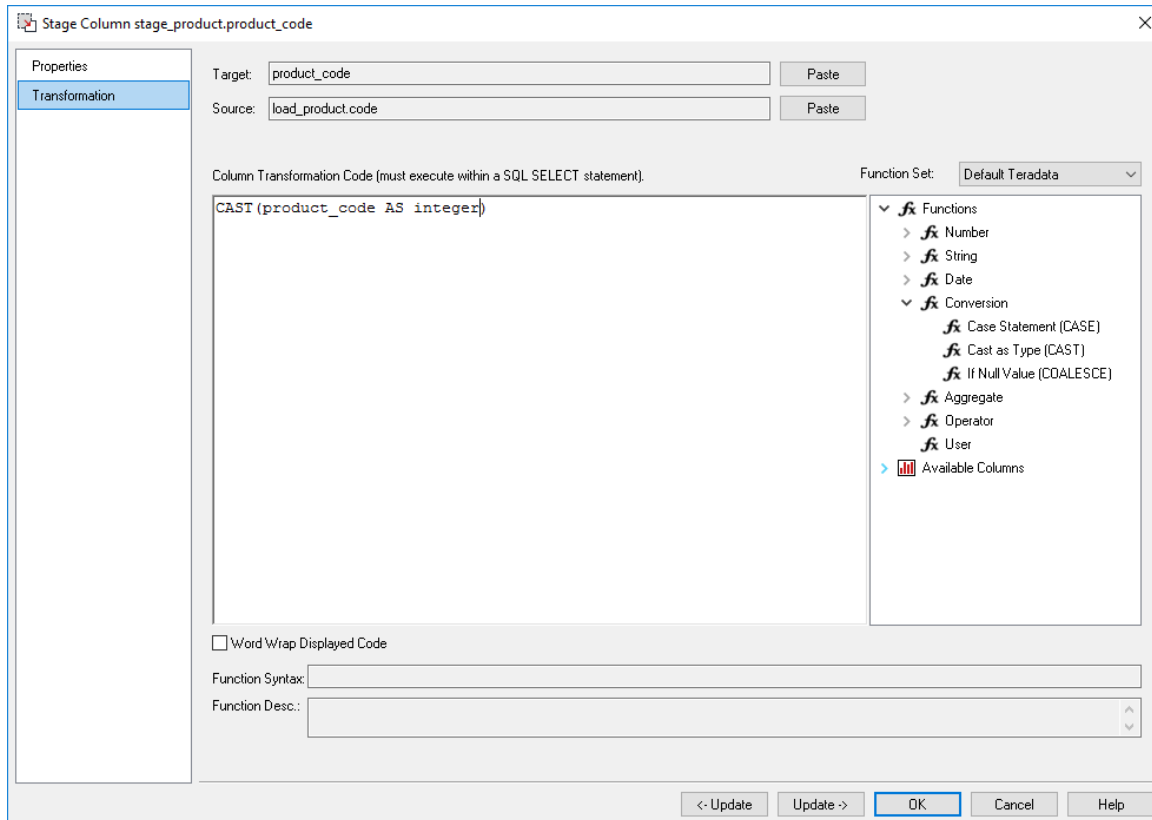
IN THIS CHAPTER

Using Database Function Sets	1036
Maintaining Database Function Sets.....	1039
Loading Database Function Sets.....	1058
Exporting Database Function Sets.....	1061

USING DATABASE FUNCTION SETS

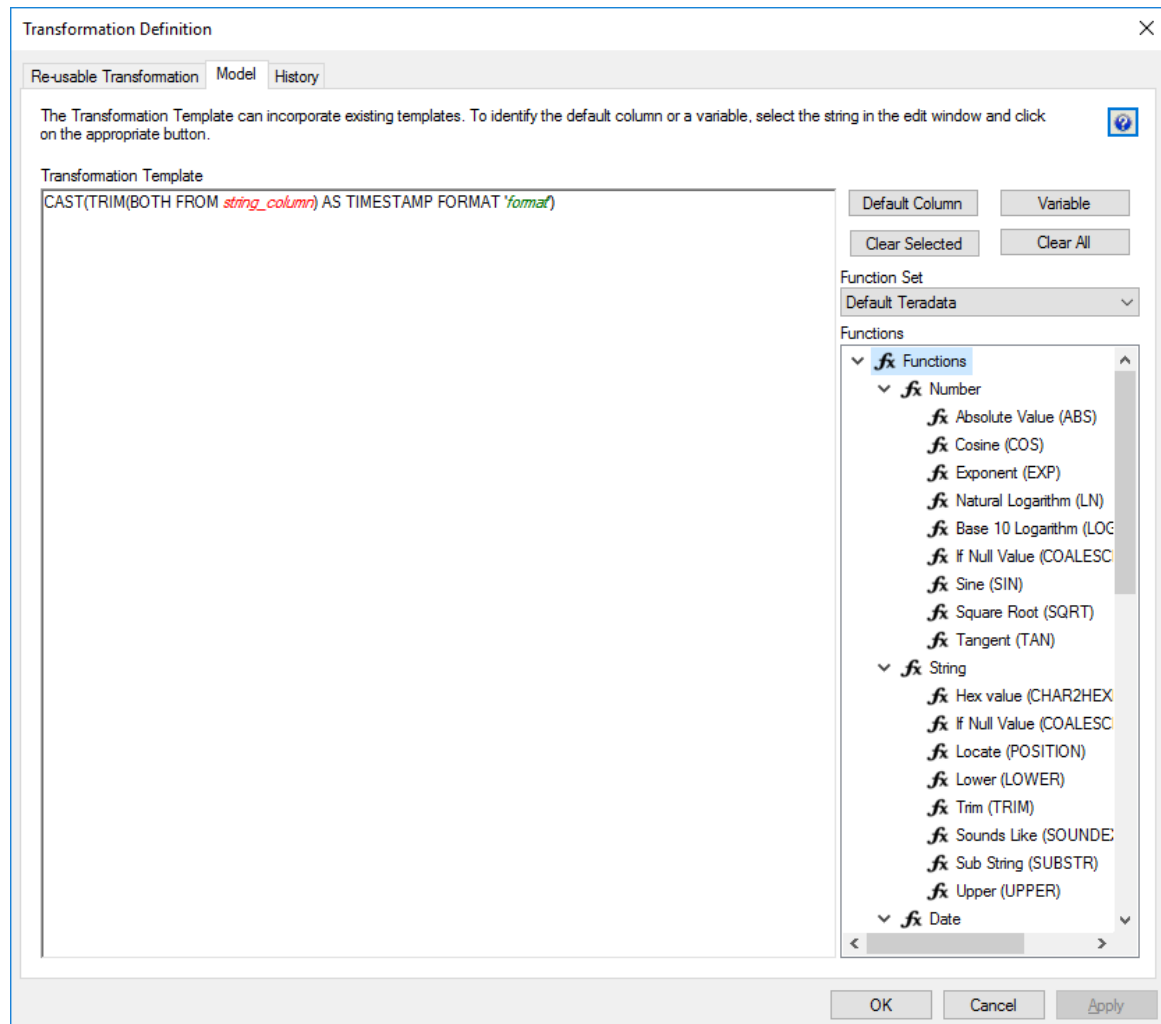
Database function sets contain a list of functions and operators that can be used for building transformations. These function sets may be created, edited, deleted, imported and exported using the **Database Functions** options from the **Tools** menu.

COLUMN TRANSFORMATION PROPERTIES DIALOG



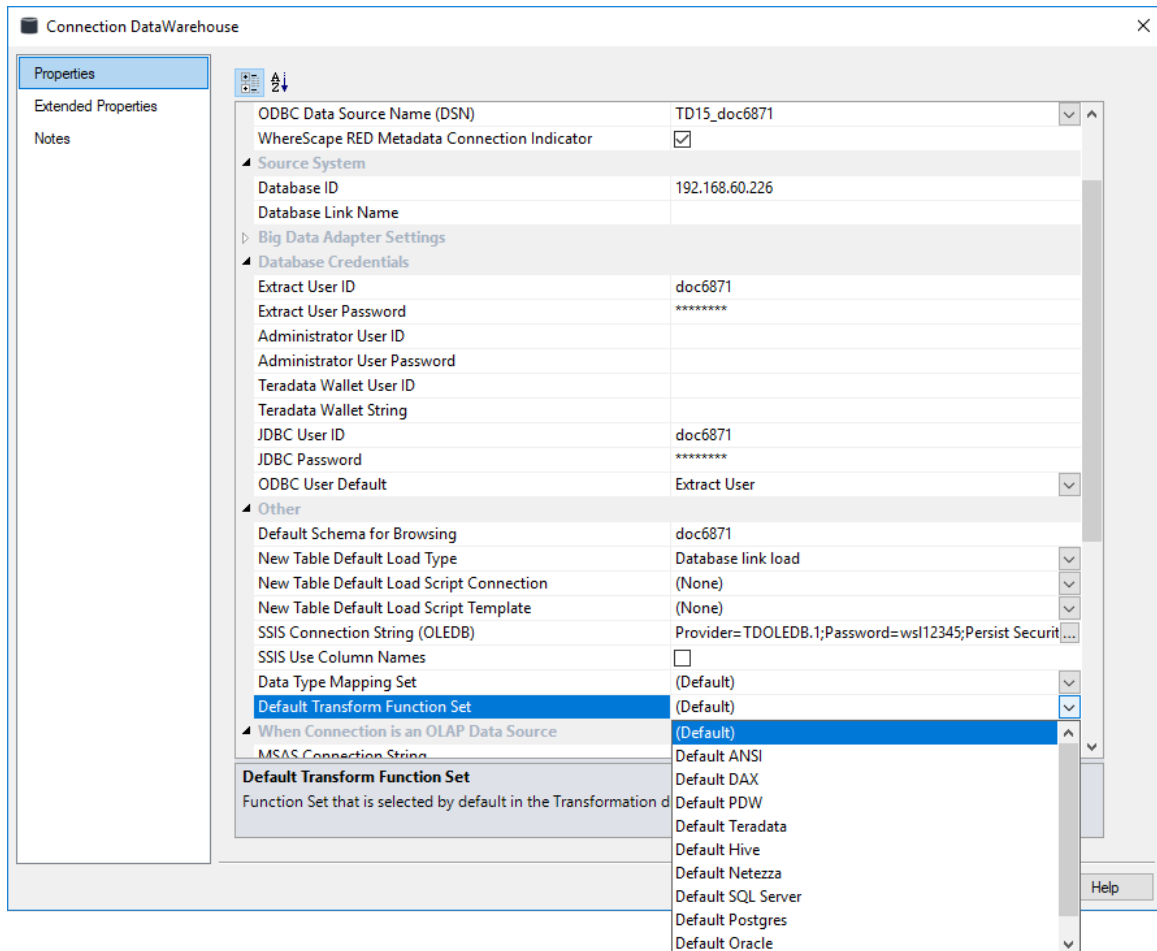
A drop-down list enables you to select which set of functions are to be displayed in the tree view when creating a transformation on a column of a table.

TRANSFORMATION DEFINITION SCREEN



A drop-down list allows the user to select which set of functions are to be displayed in the tree view when creating a re-usable transformation in **Tools>Define Re-Usable Transformations**.

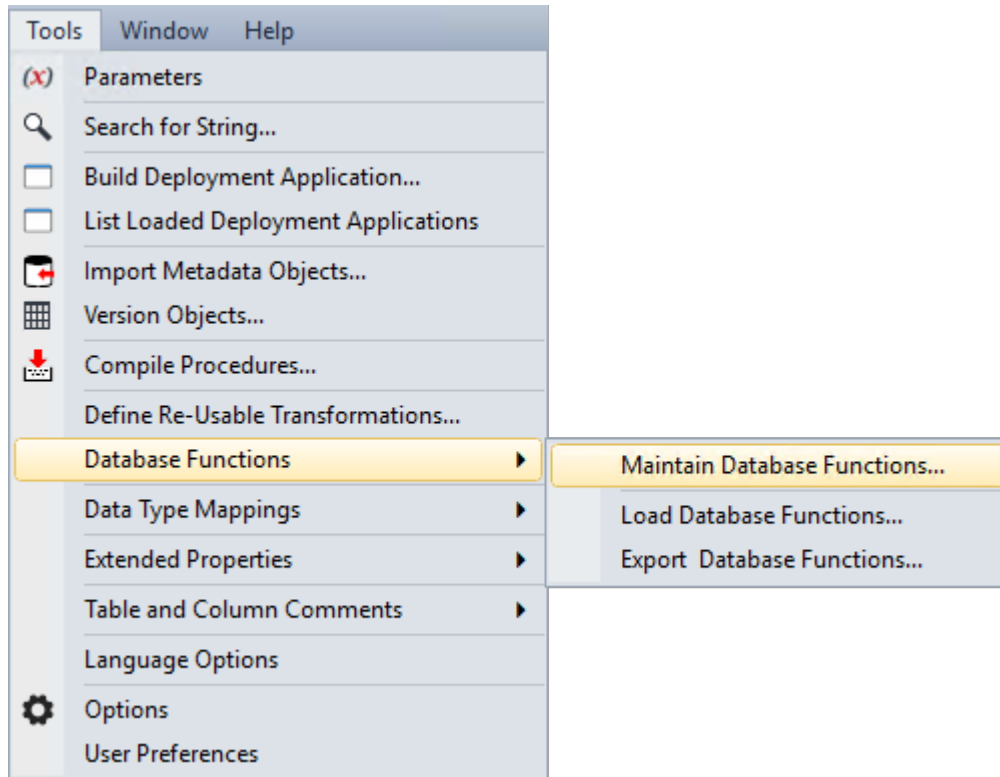
CONNECTION PROPERTIES SCREEN



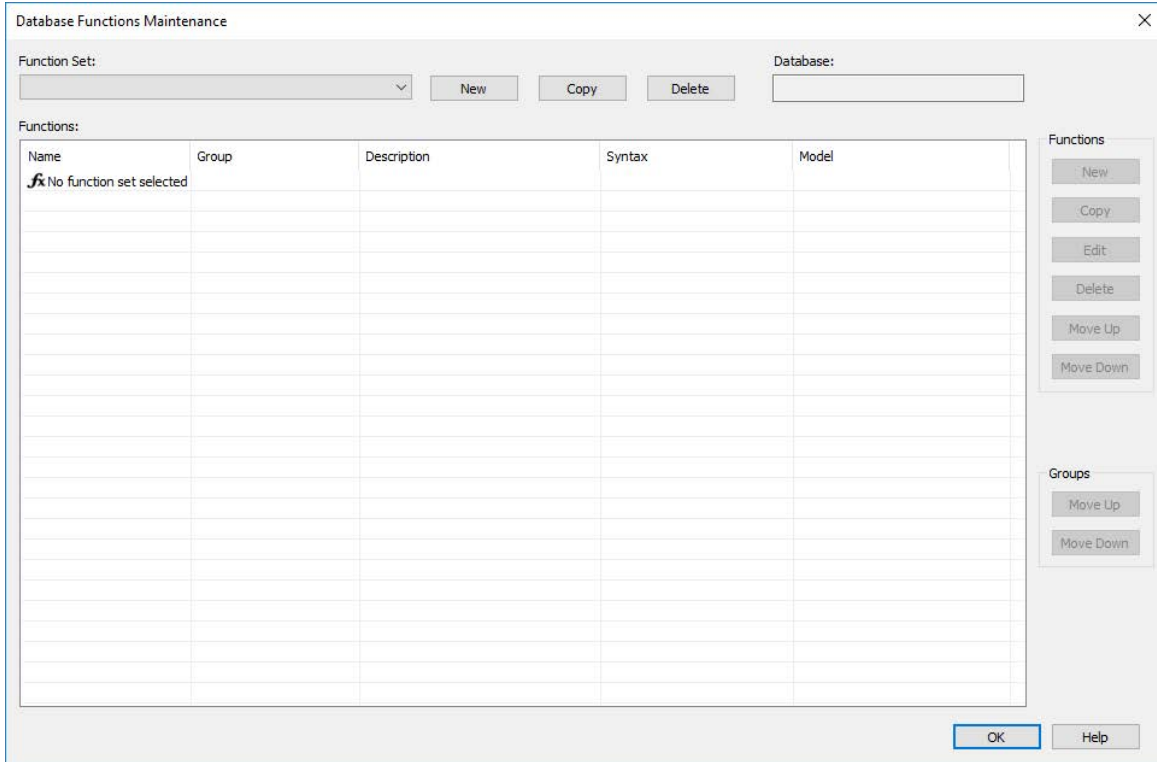
In the connection Properties dialog, when the **Data Warehouse** check box is selected, a drop-down list **Default Transform Function Set** is displayed. This is the default set that is selected in the transformation dialogs above.

MAINTAINING DATABASE FUNCTION SETS

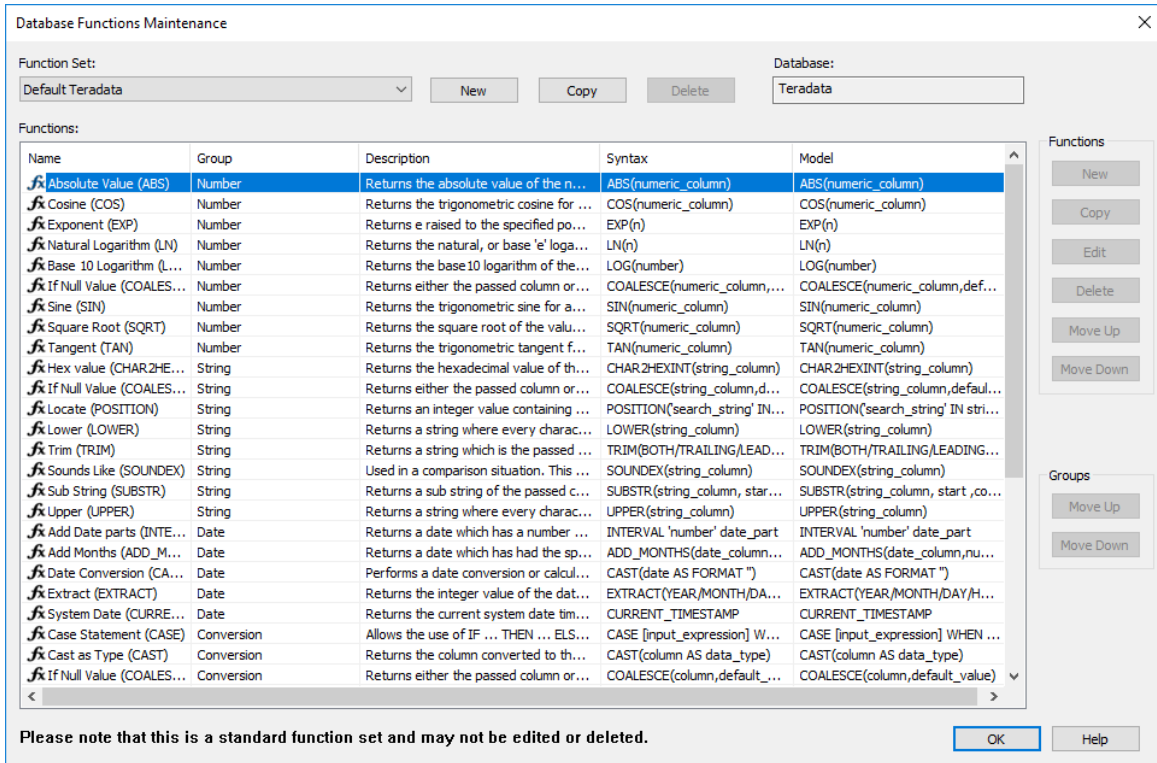
To maintain the database function sets, select **Tools>Database Functions>Maintain Database Functions...**



The following screen is displayed.



Select a database function set from the **Function Set** drop-down list.



To create a database function set, see *Creating a New Database Function Set* (on page 1041)

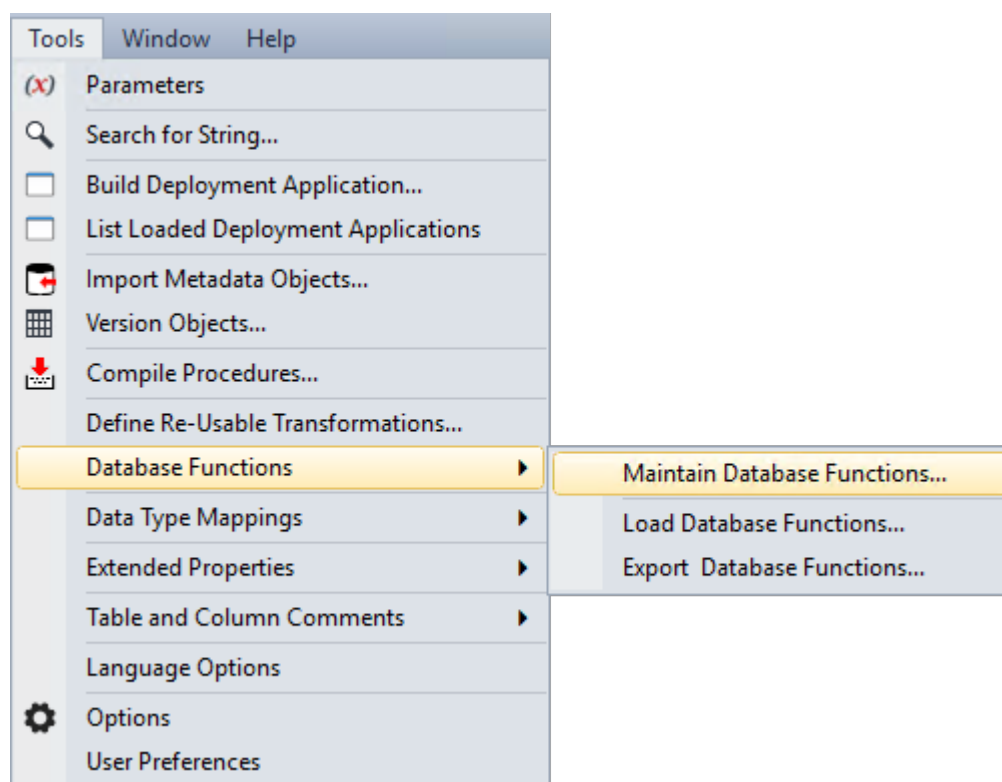
To copy a database function set, see *Copying a Database Function Set* (on page 1043)

To edit a database function set, see *Editing a Database Function Set* (on page 1045)

To delete a database function set, see *Deleting a Database Function Set* (on page 1056)

CREATING A NEW DATABASE FUNCTION SET

To create a new database function set, select **Tools>Database Functions>Maintain Database Functions...**



Click on the **New** button.

Database Functions Maintenance

Function Set: Default Teradata **New** Copy Delete Database: Teradata

Name	Group	Description	Syntax	Model
ABS	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
COS	Number	Returns the trigonometric cosine for ...	COS(numeric_column)	COS(numeric_column)
EXP	Number	Returns e raised to the specified po...	EXP(n)	EXP(n)
LN	Number	Returns the natural, or base 'e' loga...	LN(n)	LN(n)
LOG	Number	Returns the base10 logarithm of the...	LOG(number)	LOG(number)
COALESCE	Number	Returns either the passed column or...	COALESCE(numeric_column,...	COALESCE(numeric_column,def...
SIN	Number	Returns the trigonometric sine for a...	SIN(numeric_column)	SIN(numeric_column)
SQRT	Number	Returns the square root of the valu...	SQRT(numeric_column)	SQRT(numeric_column)
TAN	Number	Returns the trigonometric tangent f...	TAN(numeric_column)	TAN(numeric_column)
CHAR2HEXINT	String	Returns the hexadecimal value of th...	CHAR2HEXINT(string_column)	CHAR2HEXINT(string_column)
COALESCE	String	Returns either the passed column or...	COALESCE(string_column,d...	COALESCE(string_column,default...
POSITION	String	Returns an integer value containing ...	POSITION(search_string' IN...	POSITION(search_string' IN stri...
LOWER	String	Returns a string where every charac...	LOWER(string_column)	LOWER(string_column)
TRIM	String	Returns a string which is the passed ...	TRIM(BOTH/TRAILING/LEAD...	TRIM(BOTH/TRAILING/LEADING...
SOUNDEX	String	Used in a comparison situation. This ...	SOUNDEX(string_column)	SOUNDEX(string_column)
SUBSTR	String	Returns a sub string of the passed c...	SUBSTR(string_column, star...	SUBSTR(string_column, start_co...
UPPER	String	Returns a string where every charac...	UPPER(string_column)	UPPER(string_column)
INTERVAL	Date	Returns a date which has a number ...	INTERVAL 'number' date_part	INTERVAL 'number' date_part
ADD_MONTHS	Date	Returns a date which has had the sp...	ADD_MONTHS(date_column,nu...	ADD_MONTHS(date_column,nu...
CAST	Date	Performs a date conversion or calcul...	CAST(date AS FORMAT '')	CAST(date AS FORMAT '')
EXTRACT	Date	Returns the integer value of the dat...	EXTRACT(YEAR/MONTH/DA...	EXTRACT(YEAR/MONTH/DAY/H...
CURRENT_TIMESTAMP	Date	Returns the current system date tim...	CURRENT_TIMESTAMP	CURRENT_TIMESTAMP
CASE	Conversion	Allows the use of IF ... THEN ... ELS...	CASE [input_expression] W...	CASE [input_expression] WHEN ...
CAST	Conversion	Returns the column converted to th...	CAST(column AS data_type)	CAST(column AS data_type)
COALESCE	Conversion	Returns either the passed column or...	COALESCE(column,default_...	COALESCE(column,default_value)

Please note that this is a standard function set and may not be edited or deleted.

OK Help

Enter a **FUNCTION SET NAME** and select a **Database** from the drop-down list. Click **OK**.

New Database Function Set

Function Set Name:

Database:

Please note that the new function set will only be saved once a function has been added to the set.

OK Cancel

Function Set Name

Enter a unique function set name.

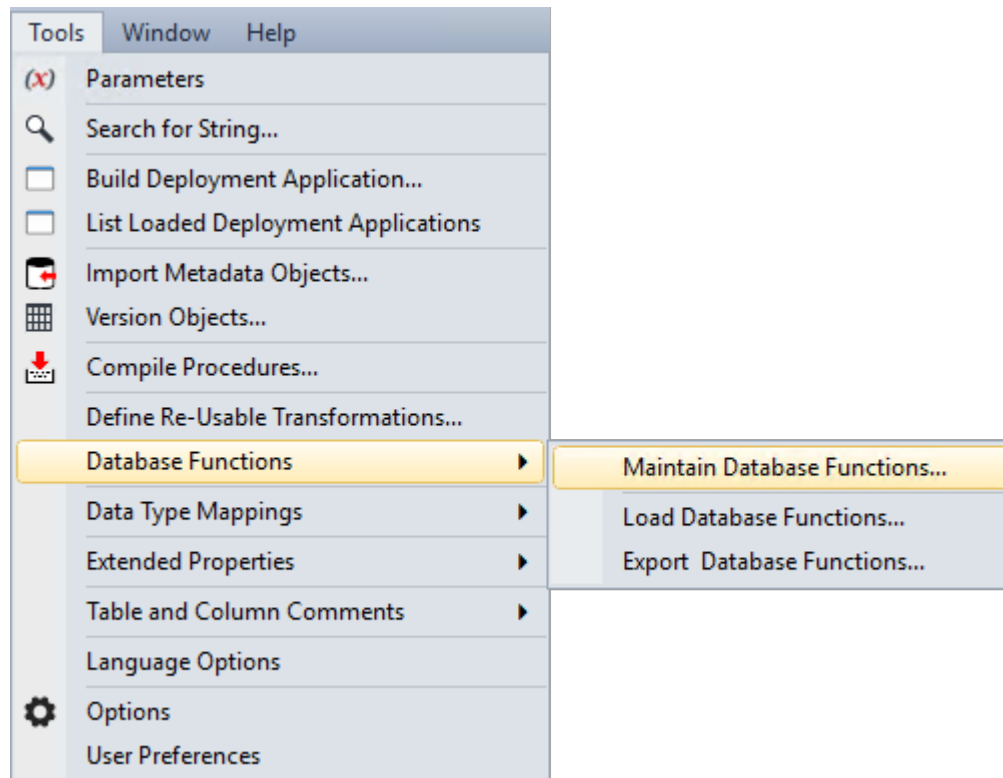
Note: The function set name may not contain the phrase "Standard Functions" as this is reserved for WhereScape supplied function sets.

Database

Select the database from the drop-down list or type in the name of the database if not already in the list. This field is mandatory.

COPYING A DATABASE FUNCTION SET

To copy an existing database function set, select **Tools>Database Functions>Maintain Database Functions...**



Select a **Function Set** from the drop-down list and click on the **Copy** button.

Database Functions Maintenance

Function Set: Default Teradata [New] **Copy** [Delete] Database: Teradata

Name	Group	Description	Syntax	Model
ABS	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
COS	Number	Returns the trigonometric cosine for ...	COS(numeric_column)	COS(numeric_column)
EXP	Number	Returns e raised to the specified po...	EXP(n)	EXP(n)
LN	Number	Returns the natural, or base 'e' loga...	LN(n)	LN(n)
LOG	Number	Returns the base10 logarithm of the...	LOG(number)	LOG(number)
COALESCE	Number	Returns either the passed column or...	COALESCE(numeric_column,...	COALESCE(numeric_column,defau...
SIN	Number	Returns the trigonometric sine for a...	SIN(numeric_column)	SIN(numeric_column)
SQRT	Number	Returns the square root of the valu...	SQRT(numeric_column)	SQRT(numeric_column)
TAN	Number	Returns the trigonometric tangent f...	TAN(numeric_column)	TAN(numeric_column)
CHAR2HEXINT	String	Returns the hexadecimal value of th...	CHAR2HEXINT(string_column)	CHAR2HEXINT(string_column)
COALESCE	String	Returns either the passed column or...	COALESCE(string_column,d...	COALESCE(string_column,default...
POSITION	String	Returns an integer value containing ...	POSITION(search_string' IN...	POSITION(search_string' IN stri...
LOWER	String	Returns a string where every charac...	LOWER(string_column)	LOWER(string_column)
TRIM	String	Returns a string which is the passed ...	TRIM(BOTH/TRAILING/LEAD...	TRIM(BOTH/TRAILING/LEADING...
SOUNDEX	String	Used in a comparison situation. This ...	SOUNDEX(string_column)	SOUNDEX(string_column)
SUBSTR	String	Returns a sub string of the passed c...	SUBSTR(string_column, star...	SUBSTR(string_column, start ,co...
UPPER	String	Returns a string where every charac...	UPPER(string_column)	UPPER(string_column)
INTERVAL	Date	Returns a date which has a number ...	INTERVAL 'number' date_part	INTERVAL 'number' date_part
ADD_MONTHS	Date	Returns a date which has had the sp...	ADD_MONTHS(date_column,nu...	ADD_MONTHS(date_column,nu...
CAST	Date	Performs a date conversion or calcul...	CAST(date AS FORMAT '')	CAST(date AS FORMAT '')
EXTRACT	Date	Returns the integer value of the dat...	EXTRACT(YEAR/MONTH/DA...	EXTRACT(YEAR/MONTH/DAY/H...
CURRENT_TIMESTAMP	Date	Returns the current system date tim...	CURRENT_TIMESTAMP	CURRENT_TIMESTAMP
CASE	Conversion	Allows the use of IF ... THEN ... ELS...	CASE [input_expression] W...	CASE [input_expression] WHEN ...
CAST	Conversion	Returns the column converted to th...	CAST(column AS data_type)	CAST(column AS data_type)
COALESCE	Conversion	Returns either the passed column or...	COALESCE(column,default_...	COALESCE(column,default_value)

Please note that this is a standard function set and may not be edited or deleted.

OK Help

Enter the new **Function Set Name** and select a **Database** from the drop-down list. Click **OK**.

New Database Function Set

Function Set Name:
Teradata Custom Functions

Database:
Teradata

Please note that the new function set will only be saved once a function has been added to the set.

OK Cancel

Function Set Name

Enter a unique function set name.

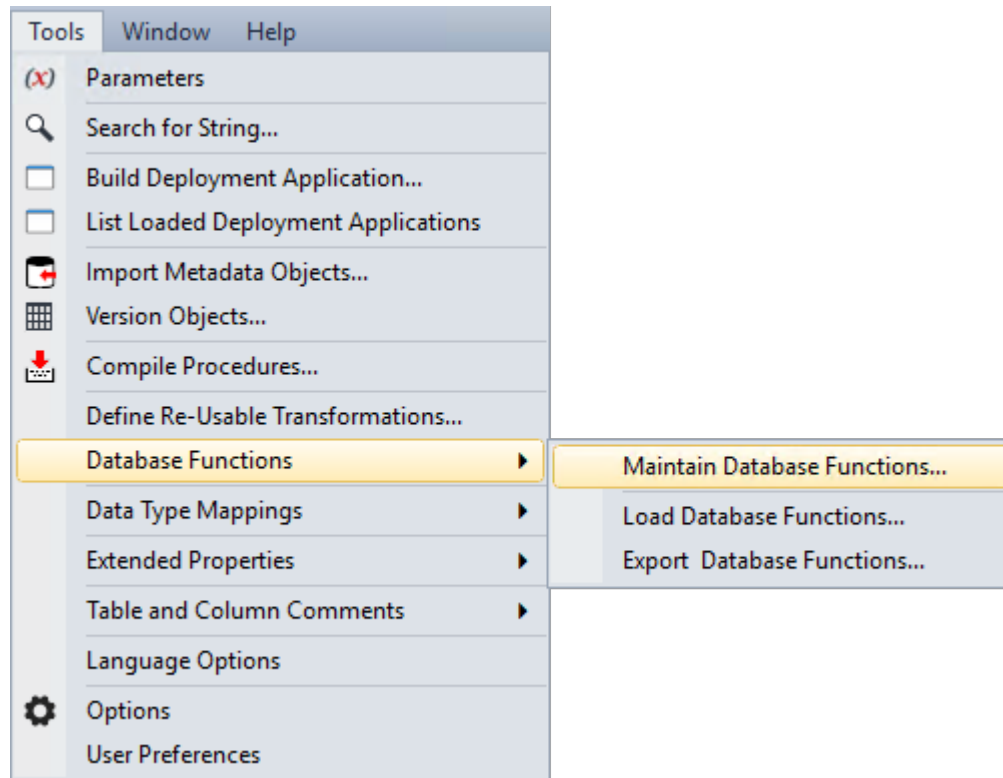
Note: The function set name may not contain the phrase "Standard Functions" as this is reserved for WhereScape supplied function sets.

Database

Select the database from the drop-down list or type in the name of the database if not already in the list. This field is mandatory.

EDITING A DATABASE FUNCTION SET

To edit a database function set, select **Tools>Database Functions>Maintain Database Functions...**



Select a database function set from the **Function Set** drop-down list.

Database Functions Maintenance

Function Set: Teradata Custom Functions New Copy Delete Database: Teradata

Functions:

Name	Group	Description	Syntax	Model
fx Absolute Value (ABS)	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
fx Cosine (COS)	Number	Returns the trigonometric cosine for ...	COS(numeric_column)	COS(numeric_column)
fx Exponent (EXP)	Number	Returns e raised to the specified po...	EXP(n)	EXP(n)
fx Natural Logarithm (LN)	Number	Returns the natural, or base 'e' loga...	LN(n)	LN(n)
fx Base 10 Logarithm (L...	Number	Returns the base10 logarithm of the...	LOG(number)	LOG(number)
fx If Null Value (COALES...	Number	Returns either the passed column or ...	COALESCE(numeric_column,...	COALESCE(numeric_column,def...
fx Sine (SIN)	Number	Returns the trigonometric sine for a...	SIN(numeric_column)	SIN(numeric_column)
fx Square Root (SQRT)	Number	Returns the square root of the valu...	SQRT(numeric_column)	SQRT(numeric_column)
fx Tangent (TAN)	Number	Returns the trigonometric tangent f...	TAN(numeric_column)	TAN(numeric_column)
fx Hex value (CHAR2HE...	String	Returns the hexadecimal value of th...	CHAR2HEXINT(string_column)	CHAR2HEXINT(string_column)
fx If Null Value (COALES...	String	Returns either the passed column or ...	COALESCE(string_column,d...	COALESCE(string_column,default...
fx Locate (POSITION)	String	Returns an integer value containing ...	POSITION('search_string' IN...	POSITION('search_string' IN stri...
fx Lower (LOWER)	String	Returns a string where every charac...	LOWER(string_column)	LOWER(string_column)
fx Trim (TRIM)	String	Returns a string which is the passed ...	TRIM(BOTH/TRAILING/LEAD...	TRIM(BOTH/TRAILING/LEADING...
fx Sounds Like (SOUNDEX)	String	Used in a comparison situation. This ...	SOUNDEX(string_column)	SOUNDEX(string_column)
fx Sub String (SUBSTR)	String	Returns a sub string of the passed c...	SUBSTR(string_column, star...	SUBSTR(string_column, start ,co...
fx Upper (UPPER)	String	Returns a string where every charac...	UPPER(string_column)	UPPER(string_column)
fx Add Date parts (INTE...	Date	Returns a date which has a number ...	INTERVAL 'number' date_part	INTERVAL 'number' date_part
fx Add Months (ADD_M...	Date	Returns a date which has had the sp...	ADD_MONTHS(date_column...	ADD_MONTHS(date_column,nu...
fx Date Conversion (CA...	Date	Performs a date conversion or calcul...	CAST(date AS FORMAT '')	CAST(date AS FORMAT '')
fx Extract (EXTRACT)	Date	Returns the integer value of the dat...	EXTRACT(YEAR/MONTH/DA...	EXTRACT(YEAR/MONTH/DAY/H...
fx System Date (CURRE...	Date	Returns the current system date tim...	CURRENT_TIMESTAMP	CURRENT_TIMESTAMP
fx Case Statement (CASE)	Conversion	Allows the use of IF ... THEN ... ELS...	CASE [input_expression] W...	CASE [input_expression] WHEN ...
fx Cast as Type (CAST)	Conversion	Returns the column converted to th...	CAST(column AS data_type)	CAST(column AS data_type)
fx If Null Value (COALES...	Conversion	Returns either the passed column or ...	COALESCE(column,default_...	COALESCE(column,default_value)

Functions: New Copy Edit Delete Move Up Move Down

Groups: Move Up Move Down

OK Help

On the right is a group of buttons used to maintain the list of functions in a function set. These buttons are not available for standard function sets.

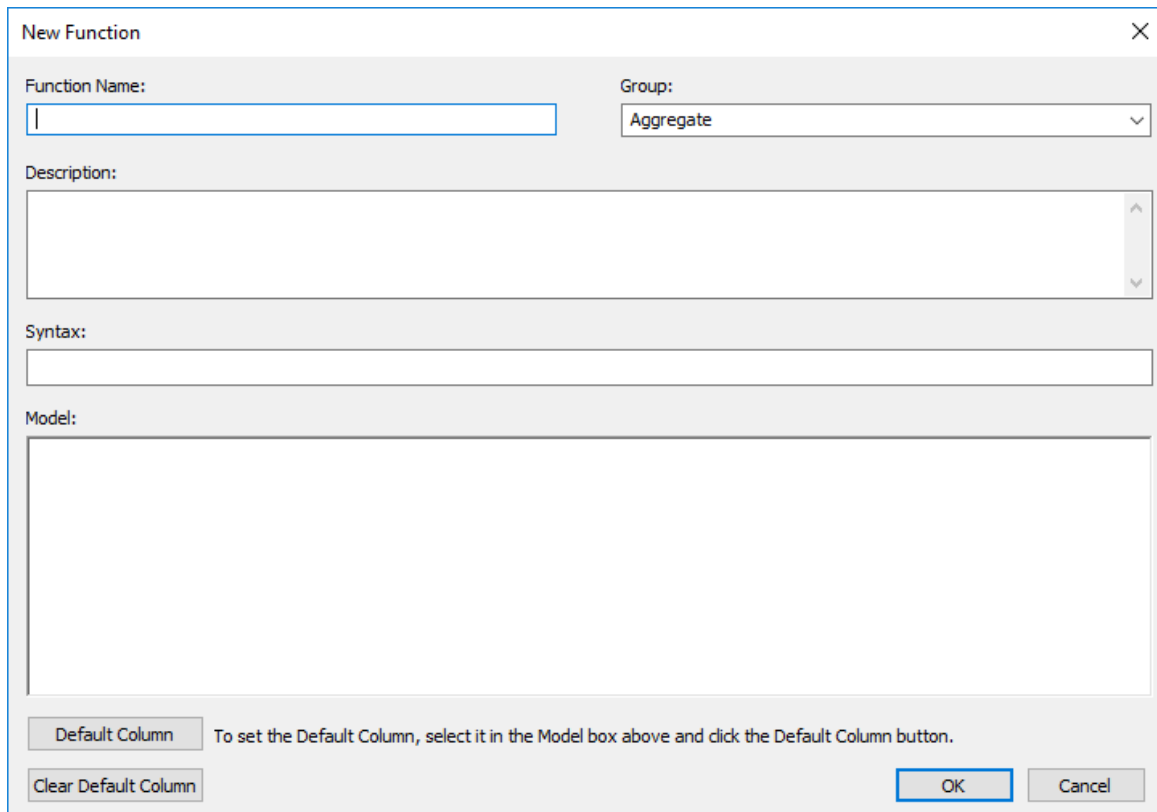
To add a new function to the database function set:

Click on the **New** button on the right.

The screenshot shows the 'Database Functions Maintenance' window. At the top, there is a 'Function Set:' dropdown menu set to 'Teradata Custom Functions' and a 'Database:' dropdown menu set to 'Teradata'. Between these menus are three buttons: 'New', 'Copy', and 'Delete'. Below the menus is a table of functions with columns for Name, Group, Description, Syntax, and Model. The 'New' button is highlighted in the 'Functions' panel on the right side of the window.

Name	Group	Description	Syntax	Model
Absolute Value (ABS)	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
Cosine (COS)	Number	Returns the trigonometric cosine for ...	COS(numeric_column)	COS(numeric_column)
Exponent (EXP)	Number	Returns e raised to the specified po...	EXP(n)	EXP(n)
Natural Logarithm (LN)	Number	Returns the natural, or base 'e' loga...	LN(n)	LN(n)
Base 10 Logarithm (L...	Number	Returns the base10 logarithm of the...	LOG(number)	LOG(number)
If Null Value (COALES...	Number	Returns either the passed column or...	COALESCE(numeric_column,...	COALESCE(numeric_column,defau...
Sine (SIN)	Number	Returns the trigonometric sine for a...	SIN(numeric_column)	SIN(numeric_column)
Square Root (SQRT)	Number	Returns the square root of the valu...	SQRT(numeric_column)	SQRT(numeric_column)
Tangent (TAN)	Number	Returns the trigonometric tangent f...	TAN(numeric_column)	TAN(numeric_column)
Hex value (CHAR2HE...	String	Returns the hexadecimal value of th...	CHAR2HEXINT(string_column)	CHAR2HEXINT(string_column)
If Null Value (COALES...	String	Returns either the passed column or...	COALESCE(string_column,d...	COALESCE(string_column,defaul...
Locate (POSITION)	String	Returns an integer value containing ...	POSITION('search_string' IN...	POSITION('search_string' IN stri...
Lower (LOWER)	String	Returns a string where every charac...	LOWER(string_column)	LOWER(string_column)
Trim (TRIM)	String	Returns a string which is the passed ...	TRIM(BOTH/TRAILING/LEAD...	TRIM(BOTH/TRAILING/LEADING...
Sounds Like (SOUNDEX)	String	Used in a comparison situation. This ...	SOUNDEX(string_column)	SOUNDEX(string_column)
Sub String (SUBSTR)	String	Returns a sub string of the passed c...	SUBSTR(string_column, star...	SUBSTR(string_column, start ,co...
Upper (UPPER)	String	Returns a string where every charac...	UPPER(string_column)	UPPER(string_column)
Add Date parts (INTE...	Date	Returns a date which has a number ...	INTERVAL 'number' date_part	INTERVAL 'number' date_part
Add Months (ADD_M...	Date	Returns a date which has had the sp...	ADD_MONTHS(date_column...	ADD_MONTHS(date_column,nu...
Date Conversion (CA...	Date	Performs a date conversion or calcul...	CAST(date AS FORMAT ")	CAST(date AS FORMAT ")
Extract (EXTRACT)	Date	Returns the integer value of the dat...	EXTRACT(YEAR/MONTH/DA...	EXTRACT(YEAR/MONTH/DAY/H...
System Date (CURRE...	Date	Returns the current system date tim...	CURRENT_TIMESTAMP	CURRENT_TIMESTAMP
Case Statement (CASE)	Conversion	Allows the use of IF ... THEN ... ELS...	CASE [input_expression] W...	CASE [input_expression] WHEN ...
Cast as Type (CAST)	Conversion	Returns the column converted to th...	CAST(column AS data_type)	CAST(column AS data_type)
If Null Value (COALES...	Conversion	Returns either the passed column or...	COALESCE(column,default_...	COALESCE(column,default_value)

Enter the details for the new function and click **OK**.



The screenshot shows a 'New Function' dialog box with the following fields and controls:

- Function Name:** A text input field.
- Group:** A dropdown menu with 'Aggregate' selected.
- Description:** A text area.
- Syntax:** A text input field.
- Model:** A text area.
- Default Column:** A button with a tooltip: 'To set the Default Column, select it in the Model box above and click the Default Column button.'
- Clear Default Column:** A button.
- OK:** A button.
- Cancel:** A button.

Function Name

This field is mandatory and must be unique within the group.

Group

This field is mandatory. Select a group from the drop-down list or add a new group name.

Description

Enter a description for the function.

Syntax


Enter the syntax for the function.

Model

This field is mandatory. This is the text that will be pasted into the transformation/ model fields when the function is selected.

Default Column

This is the text that will automatically be highlighted when the function is used in the transformation/model dialogs. To set the default column, highlight it in the model field and click the **Default Column** button. The default column will now show in red in the Model field.



Clear Default Column

Click this button to clear the default column.

To copy an existing function in the database function set:

Select the function and then click on the **Copy** button on the right.

Database Functions Maintenance ✕

Function Set: Teradata Custom Functions New Copy Delete Database: Teradata

Functions:

Name	Group	Description	Syntax	Model
✖ Absolute Value (ABS)	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
✖ Cosine (COS)	Number	Returns the trigonometric cosine for ...	COS(numeric_column)	COS(numeric_column)
✖ Exponent (EXP)	Number	Returns e raised to the specified po...	EXP(n)	EXP(n)
✖ Natural Logarithm (LN)	Number	Returns the natural, or base 'e' loga...	LN(n)	LN(n)
✖ Base 10 Logarithm (L...	Number	Returns the base10 logarithm of the...	LOG(number)	LOG(number)
✖ If Null Value (COALES...	Number	Returns either the passed column or ...	COALESCE(numeric_column,...	COALESCE(numeric_column,def...
✖ Sine (SIN)	Number	Returns the trigonometric sine for a...	SIN(numeric_column)	SIN(numeric_column)
✖ Square Root (SQRT)	Number	Returns the square root of the valu...	SQRT(numeric_column)	SQRT(numeric_column)
✖ Tangent (TAN)	Number	Returns the trigonometric tangent f...	TAN(numeric_column)	TAN(numeric_column)
✖ Hex value (CHAR2HE...	String	Returns the hexadecimal value of th...	CHAR2HEXINT(string_column)	CHAR2HEXINT(string_column)
✖ If Null Value (COALES...	String	Returns either the passed column or ...	COALESCE(string_column,d...	COALESCE(string_column,default...
✖ Locate (POSITION)	String	Returns an integer value containing ...	POSITION(search_string' IN...	POSITION(search_string' IN stri...
✖ Lower (LOWER)	String	Returns a string where every charac...	LOWER(string_column)	LOWER(string_column)
✖ Trim (TRIM)	String	Returns a string which is the passed ...	TRIM(BOTH/TRAILING/LEAD...	TRIM(BOTH/TRAILING/LEADING...
✖ Sounds Like (SOUNDEX)	String	Used in a comparison situation. This ...	SOUNDEX(string_column)	SOUNDEX(string_column)
✖ Sub String (SUBSTR)	String	Returns a sub string of the passed c...	SUBSTR(string_column, star...	SUBSTR(string_column, start ,co...
✖ Upper (UPPER)	String	Returns a string where every charac...	UPPER(string_column)	UPPER(string_column)
✖ Add Date parts (INTE...	Date	Returns a date which has a number ...	INTERVAL 'number' date_part	INTERVAL 'number' date_part
✖ Add Months (ADD_M...	Date	Returns a date which has had the sp...	ADD_MONTHS(date_column...	ADD_MONTHS(date_column,nu...
✖ Date Conversion (CA...	Date	Performs a date conversion or calcul...	CAST(date AS FORMAT ")	CAST(date AS FORMAT ")
✖ Extract (EXTRACT)	Date	Returns the integer value of the dat...	EXTRACT(YEAR/MONTH/DA...	EXTRACT(YEAR/MONTH/DAY/H...
✖ System Date (CURRE...	Date	Returns the current system date tim...	CURRENT_TIMESTAMP	CURRENT_TIMESTAMP
✖ Case Statement (CASE)	Conversion	Allows the use of IF ... THEN ... ELS...	CASE [input_expression] W...	CASE [input_expression] WHEN ...
✖ Cast as Type (CAST)	Conversion	Returns the column converted to th...	CAST(column AS data_type)	CAST(column AS data_type)
✖ If Null Value (COALES...	Conversion	Returns either the passed column or ...	COALESCE(column,default_...	COALESCE(column,default_value)

Functions

New

Copy

Edit

Delete

Move Up

Move Down

Groups

Move Up

Move Down

OK Help

Enter a new **Function Name** and change any other details; then click **OK**.

New Function ✕

Function Name: Group:

Description:
Returns the absolute value of the number. Essentially converts all values to positive
Example: ABS(load_sales.tax)

Syntax:

Model:

To set the Default Column, select it in the Model box above and click the Default Column button.

To edit an existing function in the database function set:

Select the function and then click on the **Edit** button on the right.

Database Functions Maintenance

Function Set: Teradata Custom Functions New Copy Delete Database: Teradata

Name	Group	Description	Syntax	Model
Absolute Value (ABS)	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
Cosine (COS)	Number	Returns the trigonometric cosine for ...	COS(numeric_column)	COS(numeric_column)
Exponent (EXP)	Number	Returns e raised to the specified po...	EXP(n)	EXP(n)
Natural Logarithm (LN)	Number	Returns the natural, or base 'e' loga...	LN(n)	LN(n)
Base 10 Logarithm (L...	Number	Returns the base10 logarithm of the...	LOG(number)	LOG(number)
If Null Value (COALES...	Number	Returns either the passed column or...	COALESCE(numeric_column,...	COALESCE(numeric_column,defau...
Sine (SIN)	Number	Returns the trigonometric sine for a...	SIN(numeric_column)	SIN(numeric_column)
Square Root (SQRT)	Number	Returns the square root of the valu...	SQRT(numeric_column)	SQRT(numeric_column)
Tangent (TAN)	Number	Returns the trigonometric tangent f...	TAN(numeric_column)	TAN(numeric_column)
Hex value (CHAR2HE...	String	Returns the hexadecimal value of th...	CHAR2HEXINT(string_column)	CHAR2HEXINT(string_column)
If Null Value (COALES...	String	Returns either the passed column or...	COALESCE(string_column,d...	COALESCE(string_column,defaul...
Locate (POSITION)	String	Returns an integer value containing ...	POSITION('search_string' IN...	POSITION('search_string' IN stri...
Lower (LOWER)	String	Returns a string where every charac...	LOWER(string_column)	LOWER(string_column)
Trim (TRIM)	String	Returns a string which is the passed ...	TRIM(BOTH/TRAILING/LEAD...	TRIM(BOTH/TRAILING/LEADING...
Sounds Like (SOUNDEX)	String	Used in a comparison situation. This ...	SOUNDEX(string_column)	SOUNDEX(string_column)
Sub String (SUBSTR)	String	Returns a sub string of the passed c...	SUBSTR(string_column, star...	SUBSTR(string_column, start ,co...
Upper (UPPER)	String	Returns a string where every charac...	UPPER(string_column)	UPPER(string_column)
Add Date parts (INTE...	Date	Returns a date which has a number ...	INTERVAL 'number' date_part	INTERVAL 'number' date_part
Add Months (ADD_M...	Date	Returns a date which has had the sp...	ADD_MONTHS(date_column...	ADD_MONTHS(date_column,nu...
Date Conversion (CA...	Date	Performs a date conversion or calcul...	CAST(date AS FORMAT ")	CAST(date AS FORMAT ")
Extract (EXTRACT)	Date	Returns the integer value of the dat...	EXTRACT(YEAR/MONTH/DA...	EXTRACT(YEAR/MONTH/DAY/H...
System Date (CURRE...	Date	Returns the current system date tim...	CURRENT_TIMESTAMP	CURRENT_TIMESTAMP
Case Statement (CASE)	Conversion	Allows the use of IF ... THEN ... ELS...	CASE [input_expression] W...	CASE [input_expression] WHEN ...
Cast as Type (CAST)	Conversion	Returns the column converted to th...	CAST(column AS data_type)	CAST(column AS data_type)
If Null Value (COALES...	Conversion	Returns either the passed column or...	COALESCE(column,default_...	COALESCE(column,default_value)

Functions: New Copy Edit Delete Move Up Move Down

Groups: Move Up Move Down

OK Help

Change any fields as required and then click **OK**.

✕

Function Name: Group:

Description:
Returns an integer value containing the starting position of the search string, or zero if the search string is not found.
Example: POSITION('MITH' IN load_customer.name)

Syntax:

Model:

To set the Default Column, select it in the Model box above and click the Default Column button.

Note: A function can also be edited by double-clicking on the function.

To delete an existing function in the database function set:

Select the function and then click on the **Delete** button.

Database Functions Maintenance

Function Set: Teradata Custom Functions [New] [Copy] [Delete] Database: Teradata

Name	Group	Description	Syntax	Model
✖ Absolute Value (ABS)	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
✖ Cosine (COS)	Number	Returns the trigonometric cosine for ...	COS(numeric_column)	COS(numeric_column)
✖ Exponent (EXP)	Number	Returns e raised to the specified po...	EXP(n)	EXP(n)
✖ Natural Logarithm (LN)	Number	Returns the natural, or base 'e' loga...	LN(n)	LN(n)
✖ Base 10 Logarithm (L...	Number	Returns the base10 logarithm of the...	LOG(number)	LOG(number)
✖ If Null Value (COALES...	Number	Returns either the passed column or...	COALESCE(numeric_column,...	COALESCE(numeric_column,def...
✖ Sine (SIN)	Number	Returns the trigonometric sine for a...	SIN(numeric_column)	SIN(numeric_column)
✖ Square Root (SQRT)	Number	Returns the square root of the valu...	SQRT(numeric_column)	SQRT(numeric_column)
✖ Tangent (TAN)	Number	Returns the trigonometric tangent f...	TAN(numeric_column)	TAN(numeric_column)
✖ Absolute Value (ABS)2	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
✖ Hex value (CHARZHE...	String	Returns the hexadecimal value of th...	CHAR2HEXINT(string_column)	CHAR2HEXINT(string_column)
✖ If Null Value (COALES...	String	Returns either the passed column or...	COALESCE(string_column,d...	COALESCE(string_column,default...
✖ Locate (POSITION)	String	Returns an integer value containing ...	POSITION(search_string' IN...	POSITION(search_string' IN stri...
✖ Lower (LOWER)	String	Returns a string where every charac...	LOWER(string_column)	LOWER(string_column)
✖ Trim (TRIM)	String	Returns a string which is the passed ...	TRIM(BOTH/TRAILING/LEAD...	TRIM(BOTH/TRAILING/LEADING...
✖ Sounds Like (SOUNDEX)	String	Used in a comparison situation. This ...	SOUNDEX(string_column)	SOUNDEX(string_column)
✖ Sub String (SUBSTR)	String	Returns a sub string of the passed c...	SUBSTR(string_column, star...	SUBSTR(string_column, start ,co...
✖ Upper (UPPER)	String	Returns a string where every charac...	UPPER(string_column)	UPPER(string_column)
✖ Add Date parts (INTE...	Date	Returns a date which has a number ...	INTERVAL 'number' date_part	INTERVAL 'number' date_part
✖ Add Months (ADD_M...	Date	Returns a date which has had the sp...	ADD_MONTHS(date_column...	ADD_MONTHS(date_column,nu...
✖ Date Conversion (CA...	Date	Performs a date conversion or calcul...	CAST(date AS FORMAT ")	CAST(date AS FORMAT ")
✖ Extract (EXTRACT)	Date	Returns the integer value of the dat...	EXTRACT(YEAR/MONTH/DA...	EXTRACT(YEAR/MONTH/DAY/H...
✖ System Date (CURRE...	Date	Returns the current system date tim...	CURRENT_TIMESTAMP	CURRENT_TIMESTAMP
✖ Case Statement (CASE)	Conversion	Allows the use of IF ... THEN ... ELS...	CASE [input_expression] W...	CASE [input_expression] WHEN ...
✖ Cast as Type (CAST)	Conversion	Returns the column converted to th...	CAST(column AS data_type)	CAST(column AS data_type)

Functions: [New] [Copy] [Edit] [Delete] [Move Up] [Move Down]

Groups: [Move Up] [Move Down]

[OK] [Help]

Click **Yes** to confirm the delete.

WhereScape RED

Are you sure you want to delete the function Absolute Value (ABS)2

[Yes] [No]

To move a function in the database function set up or down in the list:

Select a function and then click on the **Move Up** button on the right to move the function up in the function list, within its group; else click on the **Move Down** button on the right to move the function down in the function list, within its group.

Database Functions Maintenance ✕

Function Set: Teradata Custom Functions New Copy Delete Database: Teradata

Name	Group	Description	Syntax	Model
Absolute Value (ABS)	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
Cosine (COS)	Number	Returns the trigonometric cosine for ...	COS(numeric_column)	COS(numeric_column)
Exponent (EXP)	Number	Returns e raised to the specified po...	EXP(n)	EXP(n)
Natural Logarithm (LN)	Number	Returns the natural, or base 'e' loga...	LN(n)	LN(n)
Base 10 Logarithm (L...	Number	Returns the base10 logarithm of the...	LOG(number)	LOG(number)
If Null Value (COALES...	Number	Returns either the passed column or...	COALESCE(numeric_column,...	COALESCE(numeric_column,def...
Sine (SIN)	Number	Returns the trigonometric sine for a...	SIN(numeric_column)	SIN(numeric_column)
Square Root (SQRT)	Number	Returns the square root of the valu...	SQRT(numeric_column)	SQRT(numeric_column)
Tangent (TAN)	Number	Returns the trigonometric tangent f...	TAN(numeric_column)	TAN(numeric_column)
Hex value (CHARZHE...	String	Returns the hexadecimal value of th...	CHAR2HEXINT(string_column)	CHAR2HEXINT(string_column)
If Null Value (COALES...	String	Returns either the passed column or...	COALESCE(string_column,d...	COALESCE(string_column,default...
Locate (POSITION)	String	Returns an integer value containing ...	POSITION('search_string' IN...	POSITION('search_string' IN stri...
Lower (LOWER)	String	Returns a string where every charac...	LOWER(string_column)	LOWER(string_column)
Trim (TRIM)	String	Returns a string which is the passed ...	TRIM(BOTH/TRAILING/LEAD...	TRIM(BOTH/TRAILING/LEADING...
Sounds Like (SOUNDEX)	String	Used in a comparison situation. This ...	SOUNDEX(string_column)	SOUNDEX(string_column)
Sub String (SUBSTR)	String	Returns a sub string of the passed c...	SUBSTR(string_column, star...	SUBSTR(string_column, start ,co...
Upper (UPPER)	String	Returns a string where every charac...	UPPER(string_column)	UPPER(string_column)
Add Date parts (INTE...	Date	Returns a date which has a number ...	INTERVAL 'number' date_part	INTERVAL 'number' date_part
Add Months (ADD_M...	Date	Returns a date which has had the sp...	ADD_MONTHS(date_column...	ADD_MONTHS(date_column,nu...
Date Conversion (CA...	Date	Performs a date conversion or calcul...	CAST(date AS FORMAT ")	CAST(date AS FORMAT ")
Extract (EXTRACT)	Date	Returns the integer value of the dat...	EXTRACT(YEAR/MONTH/DA...	EXTRACT(YEAR/MONTH/DAY/H...
System Date (CURRE...	Date	Returns the current system date tim...	CURRENT_TIMESTAMP	CURRENT_TIMESTAMP
Case Statement (CASE)	Conversion	Allows the use of IF ... THEN ... ELS...	CASE [input_expression] W...	CASE [input_expression] WHEN ...
Cast as Type (CAST)	Conversion	Returns the column converted to th...	CAST(column AS data_type)	CAST(column AS data_type)
If Null Type (COALES...	Conversion	Returns either the passed column or...	COALESCE(column,default_...	COALESCE(column,default_value)

Functions

New

Copy

Edit

Delete

Move Up

Move Down

Groups

Move Up

Move Down

OK Help

To move a group of functions in the database function set up or down in the list:

Using the Group column, select any function in a particular group and then click on the **Move Up** button under the **Groups** heading on the right to move the function group up in the function list. Similarly, use the **Move Down** button under the Groups heading to move a function group down in the function list.

Database Functions Maintenance

Function Set: Teradata Custom Functions New Copy Delete Database: Teradata

Name	Group	Description	Syntax	Model
Absolute Value (ABS)	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
Cosine (COS)	Number	Returns the trigonometric cosine for ...	COS(numeric_column)	COS(numeric_column)
Exponent (EXP)	Number	Returns e raised to the specified po...	EXP(n)	EXP(n)
Natural Logarithm (LN)	Number	Returns the natural, or base 'e' loga...	LN(n)	LN(n)
Base 10 Logarithm (L...	Number	Returns the base10 logarithm of the...	LOG(number)	LOG(number)
If Null Value (COALES...	Number	Returns either the passed column or...	COALESCE(numeric_column,...	COALESCE(numeric_column,def...
Sine (SIN)	Number	Returns the trigonometric sine for a...	SIN(numeric_column)	SIN(numeric_column)
Square Root (SQRT)	Number	Returns the square root of the valu...	SQRT(numeric_column)	SQRT(numeric_column)
Tangent (TAN)	Number	Returns the trigonometric tangent f...	TAN(numeric_column)	TAN(numeric_column)
Hex value (CHARZHE...	String	Returns the hexadecimal value of th...	CHARZHEXINT(string_column)	CHARZHEXINT(string_column)
If Null Value (COALES...	String	Returns either the passed column or...	COALESCE(string_column,d...	COALESCE(string_column,default...
Locate (POSITION)	String	Returns an integer value containing ...	POSITION(search_string' IN...	POSITION(search_string' IN stri...
Lower (LOWER)	String	Returns a string where every charac...	LOWER(string_column)	LOWER(string_column)
Trim (TRIM)	String	Returns a string which is the passed ...	TRIM(BOTH/TRAILING/LEAD...	TRIM(BOTH/TRAILING/LEADING...
Sounds Like (SOUNDEX)	String	Used in a comparison situation. This ...	SOUNDEX(string_column)	SOUNDEX(string_column)
Sub String (SUBSTR)	String	Returns a sub string of the passed c...	SUBSTR(string_column, star...	SUBSTR(string_column, start,co...
Upper (UPPER)	String	Returns a string where every charac...	UPPER(string_column)	UPPER(string_column)
Add Date parts (INTE...	Date	Returns a date which has a number ...	INTERVAL 'number' date_part	INTERVAL 'number' date_part
Add Months (ADD_M...	Date	Returns a date which has had the sp...	ADD_MONTHS(date_column...	ADD_MONTHS(date_column,nu...
Date Conversion (CA...	Date	Performs a date conversion or calcul...	CAST(date AS FORMAT ")	CAST(date AS FORMAT ")
Extract (EXTRACT)	Date	Returns the integer value of the dat...	EXTRACT(YEAR/MONTH/DA...	EXTRACT(YEAR/MONTH/DAY/H...
System Date (CURRE...	Date	Returns the current system date tim...	CURRENT_TIMESTAMP	CURRENT_TIMESTAMP
Case Statement (CASE)	Conversion	Allows the use of IF ... THEN ... ELS...	CASE [input_expression] W...	CASE [input_expression] WHEN ...
Cast as Type (CAST)	Conversion	Returns the column converted to th...	CAST(column AS data_type)	CAST(column AS data_type)
If Null Value (COALES...	Conversion	Returns either the passed column or...	COALESCE(column,default_...	COALESCE(column,default_value)

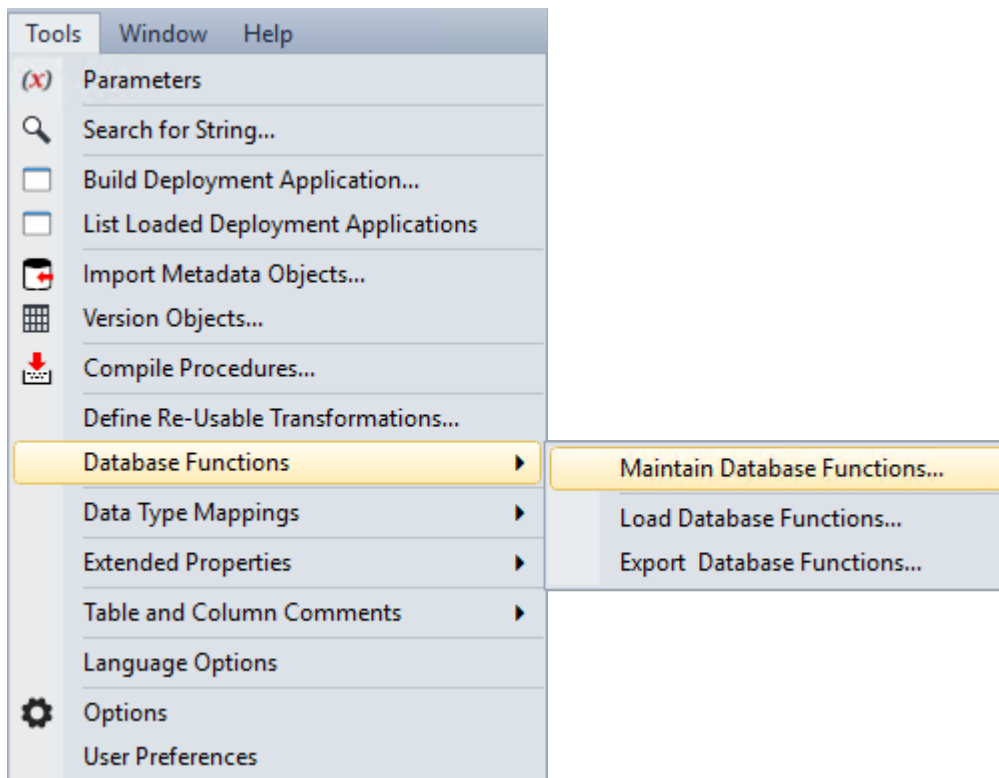
Functions: New Copy Edit Delete Move Up Move Down

Groups: Move Up Move Down

OK Help

DELETING A DATABASE FUNCTION SET

To delete a database function set, select **Tools>Database Functions>Maintain Database Functions...**



Select a **Function Set** from the drop-down list and click on the **Delete** button.

Database Functions Maintenance

Function Set: Teradata Custom Functions [v] New Copy **Delete** Database: Teradata

Name	Group	Description	Syntax	Model
f Absolute Value (ABS)	Number	Returns the absolute value of the n...	ABS(numeric_column)	ABS(numeric_column)
f Cosine (COS)	Number	Returns the trigonometric cosine for ...	COS(numeric_column)	COS(numeric_column)
f Exponent (EXP)	Number	Returns e raised to the specified po...	EXP(n)	EXP(n)
f Natural Logarithm (LN)	Number	Returns the natural, or base 'e' loga...	LN(n)	LN(n)
f Base 10 Logarithm (L...	Number	Returns the base10 logarithm of the...	LOG(number)	LOG(number)
f If Null Value (COALES...	Number	Returns either the passed column or ...	COALESCE(numeric_column,...	COALESCE(numeric_column,def...
f Sine (SIN)	Number	Returns the trigonometric sine for a...	SIN(numeric_column)	SIN(numeric_column)
f Square Root (SQRT)	Number	Returns the square root of the valu...	SQRT(numeric_column)	SQRT(numeric_column)
f Tangent (TAN)	Number	Returns the trigonometric tangent f...	TAN(numeric_column)	TAN(numeric_column)
f Hex value (CHAR2HE...	String	Returns the hexadecimal value of th...	CHAR2HEXINT(string_column)	CHAR2HEXINT(string_column)
f If Null Value (COALES...	String	Returns either the passed column or ...	COALESCE(string_column,d...	COALESCE(string_column,default...
f Locate (POSITION)	String	Returns an integer value containing ...	POSITION('search_string' IN...	POSITION('search_string' IN stri...
f Lower (LOWER)	String	Returns a string where every charac...	LOWER(string_column)	LOWER(string_column)
f Trim (TRIM)	String	Returns a string which is the passed ...	TRIM(BOTH/TRAILING/LEAD...	TRIM(BOTH/TRAILING/LEADING...
f Sounds Like (SOUNDEX)	String	Used in a comparison situation. This ...	SOUNDEX(string_column)	SOUNDEX(string_column)
f Sub String (SUBSTR)	String	Returns a sub string of the passed c...	SUBSTR(string_column, star...	SUBSTR(string_column, start ,co...
f Upper (UPPER)	String	Returns a string where every charac...	UPPER(string_column)	UPPER(string_column)
f Add Date parts (INTE...	Date	Returns a date which has a number ...	INTERVAL 'number' date_part	INTERVAL 'number' date_part
f Add Months (ADD_M...	Date	Returns a date which has had the sp...	ADD_MONTHS(date_column,nu...	ADD_MONTHS(date_column,nu...
f Date Conversion (CA...	Date	Performs a date conversion or calcul...	CAST(date AS FORMAT ")	CAST(date AS FORMAT ")
f Extract (EXTRACT)	Date	Returns the integer value of the dat...	EXTRACT(YEAR/MONTH/DA...	EXTRACT(YEAR/MONTH/DAY/H...
f System Date (CURRE...	Date	Returns the current system date tim...	CURRENT_TIMESTAMP	CURRENT_TIMESTAMP
f Case Statement (CASE)	Conversion	Allows the use of IF ... THEN ... ELS...	CASE [input_expression] W...	CASE [input_expression] WHEN ...
f Cast as Type (CAST)	Conversion	Returns the column converted to th...	CAST(column AS data_type)	CAST(column AS data_type)
f If Null Value (COALES...	Conversion	Returns either the passed column or ...	COALESCE(column,default_...	COALESCE(column,default_value)

Functions: New, Copy, Edit, Delete, Move Up, Move Down

Groups: Move Up, Move Down

OK Help

Note: The **Delete** button is disabled for standard function sets.

Click **Yes** to confirm the delete.

WhereScape RED

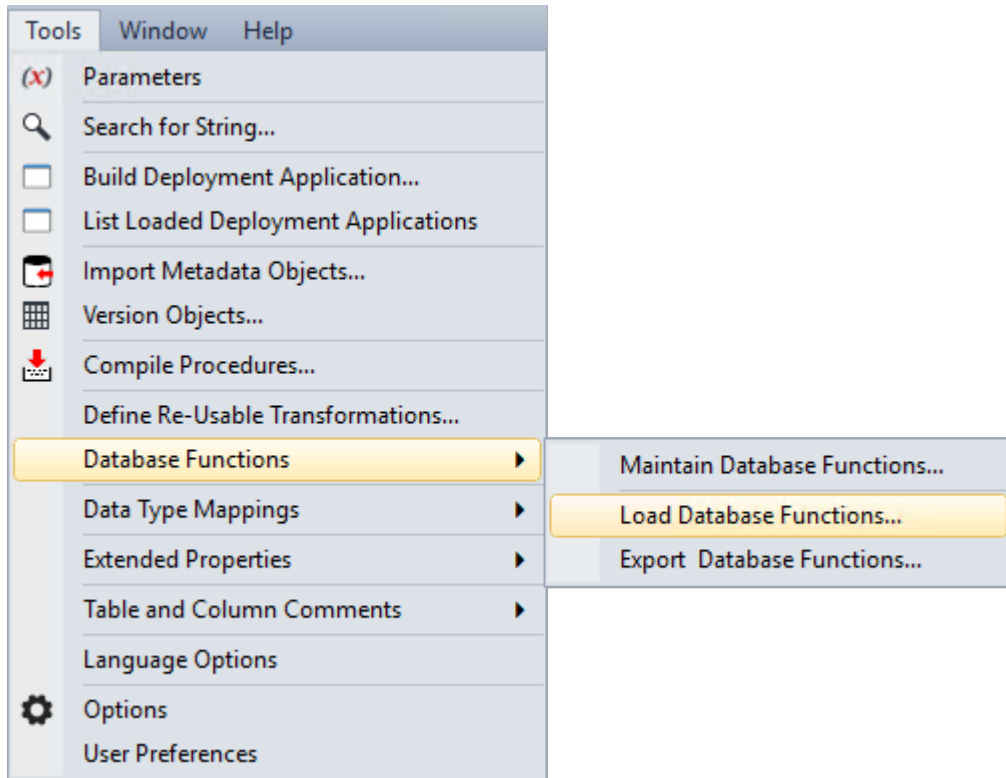
Are you sure you want to delete the function set Teradata Custom Functions

Yes No

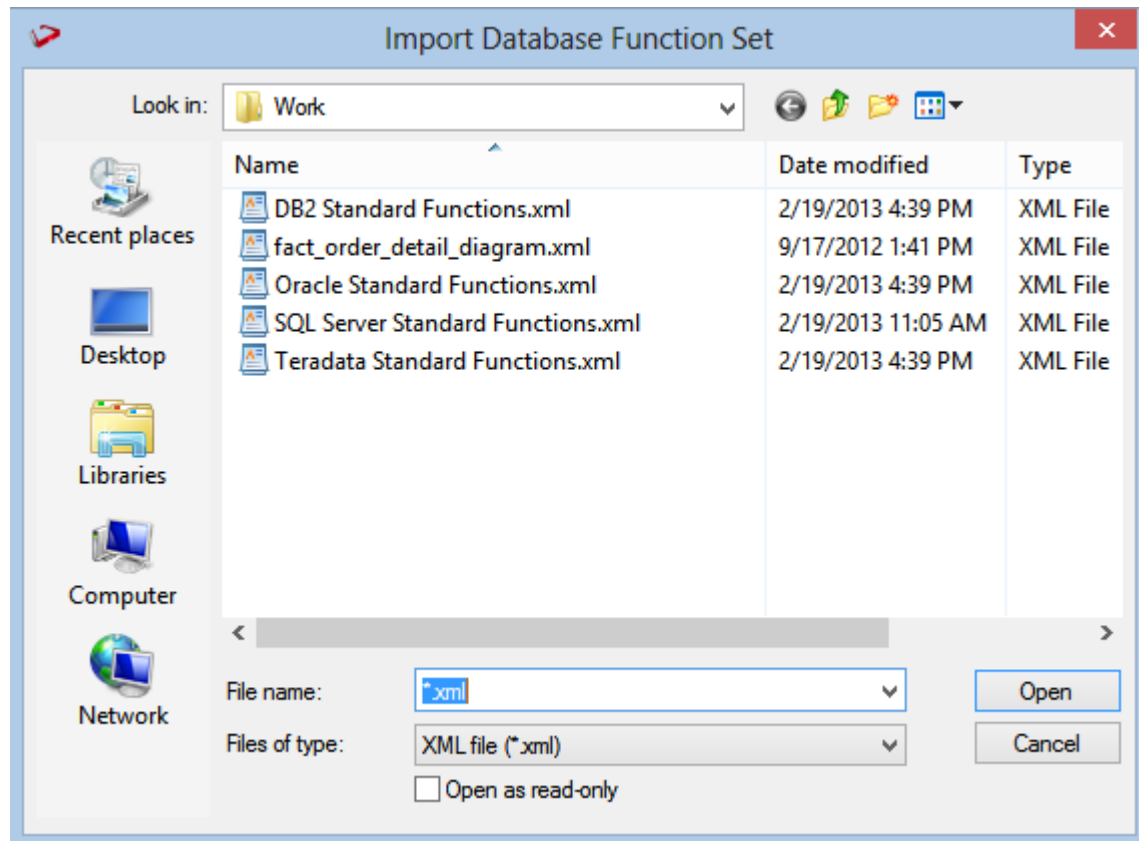
When all functions are deleted, the function set ceases to exist.

LOADING DATABASE FUNCTION SETS

To load a database function set, select **Tools>Database Functions>Load Database Functions...**

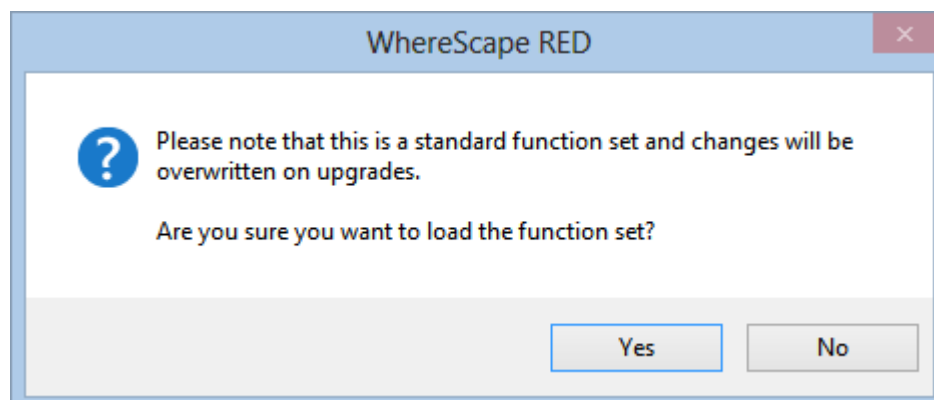


The following dialog is displayed. Select the xml file to load the database functions. By default RED expects the xml files to be in **ProgramData\WhereScape\Work** directory.

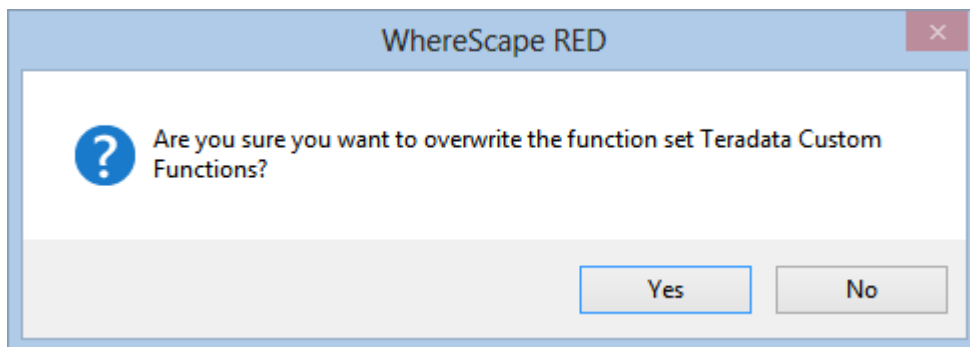


The xml file is validated using the schema definition file at <install directory>\Administrator\Function Sets\Database Function Set.xsd


If a function set containing the phrase **Standard Function** is loaded, a warning is displayed:



If an existing function set (not a standard set) is about to be overwritten, a warning is displayed:

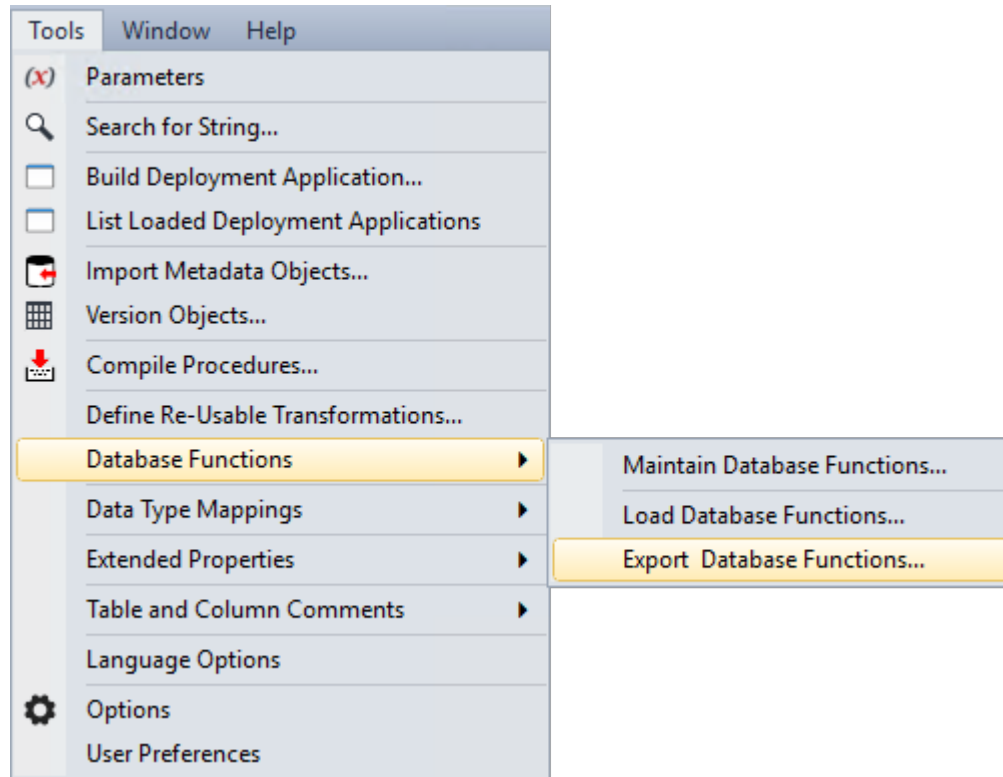


A message is displayed in the results pane.

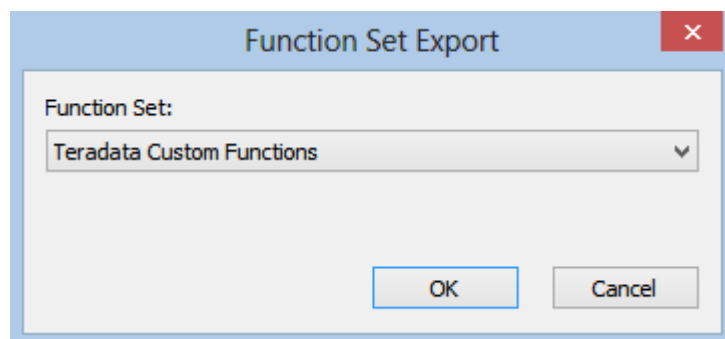


EXPORTING DATABASE FUNCTION SETS

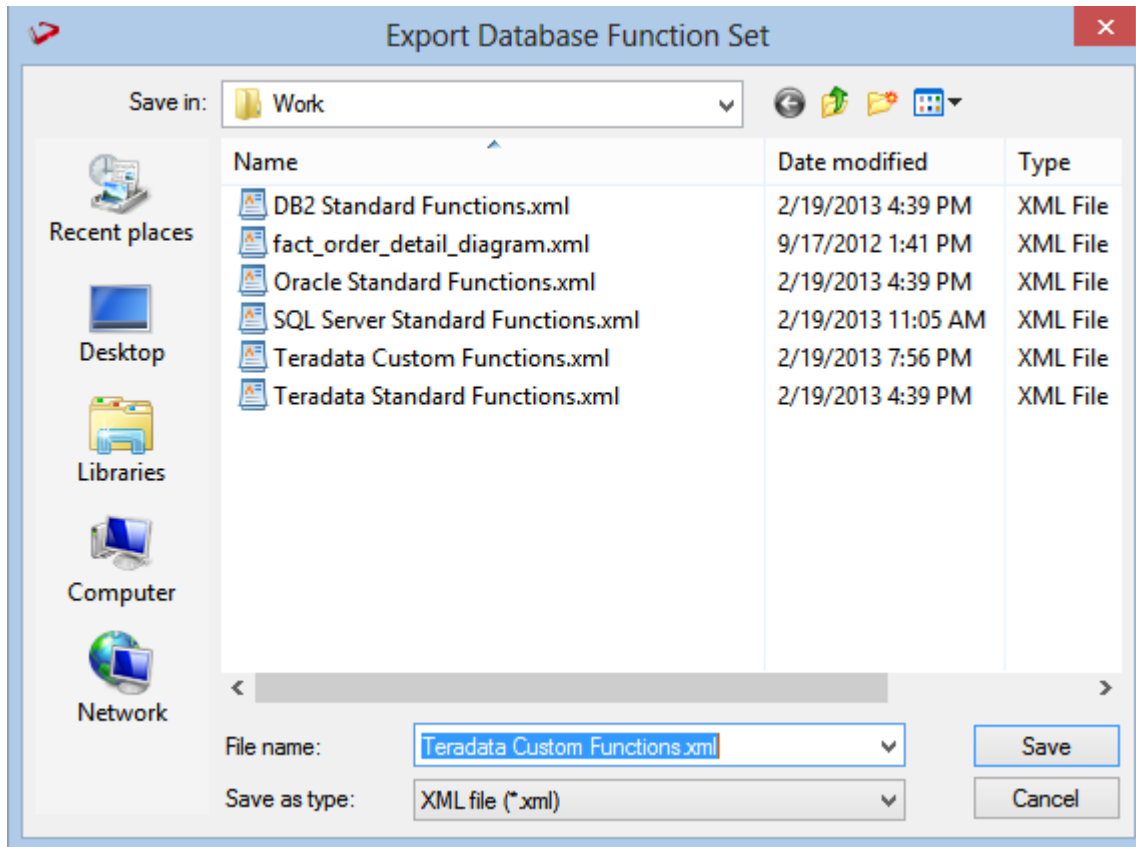
To export Database Function Sets, select **Tools>Database Functions>Export Database Functions...**



Select the **Function Set** to export from the drop-down list. Click **OK**.



By default, RED exports the xml file to **ProgramData\WhereScape\Work**, but this can be changed. Change the File name if necessary and click **Save**.



CHAPTER 44

GATHER STATISTICS

Gathering Statistics on a table will enable the underlying database to optimize each query based on the statistics collected about the data that is being accessed.

Users can either chose to **Define Statistics** or to **Refresh Full Statistics**.

Gathering statistics can be performed on any table by selecting this option from a table's right click menu, or to automate this process, by adding a statistics task to a job being processed by the scheduler (Stats, Quick Stats, Analyze or Quick Analyze). For more information about adding statistics tasks to jobs see *Editing Tasks* (see "*Editing Tasks in a Job*" on page 715).

IN THIS CHAPTER

Define Statistics 1063


DEFINE STATISTICS

In a multiple database environment it is required not to have the database name hard-coded into the COLLECT STATS command.

When tables are deployed to a new environment using RED, these hard coded database names are not translated to the database names in the target environment in the COLLECT STATS commands.

RED can automatically add the Database Name (storage location) prefix to the table_name in the COLLECT STATS DDL.

- 1 Right click on any table object on the left pane to **Gather Statistics**.
- 2 Find **Gather Statics** on the drop-down menu.
- 3 Options allow choosing either to **Define Statistics** or to **Refresh Full Statistics**.
- 4 Select **Gather Statistics -> Define Statistics**.
- 5 On the Define Statistics Collection dialog, enter the "COLLECT STATS dbname.tablename COLUMN(column_name,);" command by using \$OBJECT\$ COLUMN (table_name).



- 6 Click **Execute** if adding the DDL statement for the first time.

When the DDL is executed, RED substitutes the Database name from the Storage Properties of the object. This improves the portability of objects between environments.

- 7 Click **OK** on the **Define Statistics Collection** dialog to close it.
- 8 Right-click on the table and select **Refresh Full Statistics** to manually refresh the statistics on that table.
- 9 To automate gathering statistics on a table, users can add a statistics task to a job to be processed by the scheduler. For more information about adding statistics tasks to jobs see *Editing Tasks* (see "*Editing Tasks in a Job*" on page 715).

CHAPTER 45

EXTENDED PROPERTIES

WhereScape RED provides a facility for defining and managing extended database properties which can be used in a range of scenarios, one of the motives for their introduction was to aid interaction with the variety of new database technologies that RED customers want to exploit for their Business Intelligence infrastructure.

To support these emerging database technologies, WhereScape provides enablement packs for a variety of target “custom” database types, which includes Snowflake, Amazon Redshift and Microsoft Azure SQL data warehouse.

One component of these packs are a set of extended properties to facilitate connecting to a specific database technology.

This feature also supports importing a set of extended property definitions into RED from an external source, and exporting extended properties defined in RED to a file. This enables you to migrate and reuse extended property definitions between different repositories.

Using the extended properties feature involves defining the extended properties to specify the variables and their scope; and then assigning values to the defined extended properties.

Extended Property Definitions

Extended property definitions primarily consist of a unique name and a scope (database type and object type), which specifies where the extended property will be available.

The extended property definitions are created and maintained in the **Extended Property Definition screen** (see "*Defining Extended Properties*" on page 1066).

Extended Property Value Assignment

Once an extended property has been defined for a connection or table object, users can then assign values to each of the variable attributes available in the extended property.

The extended property values are specified and maintained in the **Extended Properties settings** (see "*Extended Properties Value Assignment*" on page 1072) in the Connection or Table object's Properties screen.

Extended Property Lookup

The extended property variables and their assigned values are referenced in PowerShell scripts. PowerShell script based processing is available for all objects in a custom database target. However, in a non-custom database target; PowerShell scripts are only available to process load and export objects.

WhereScape provides an application that enables access to extended property definitions and values. For more information, please contact your WhereScape representative.

IN THIS CHAPTER

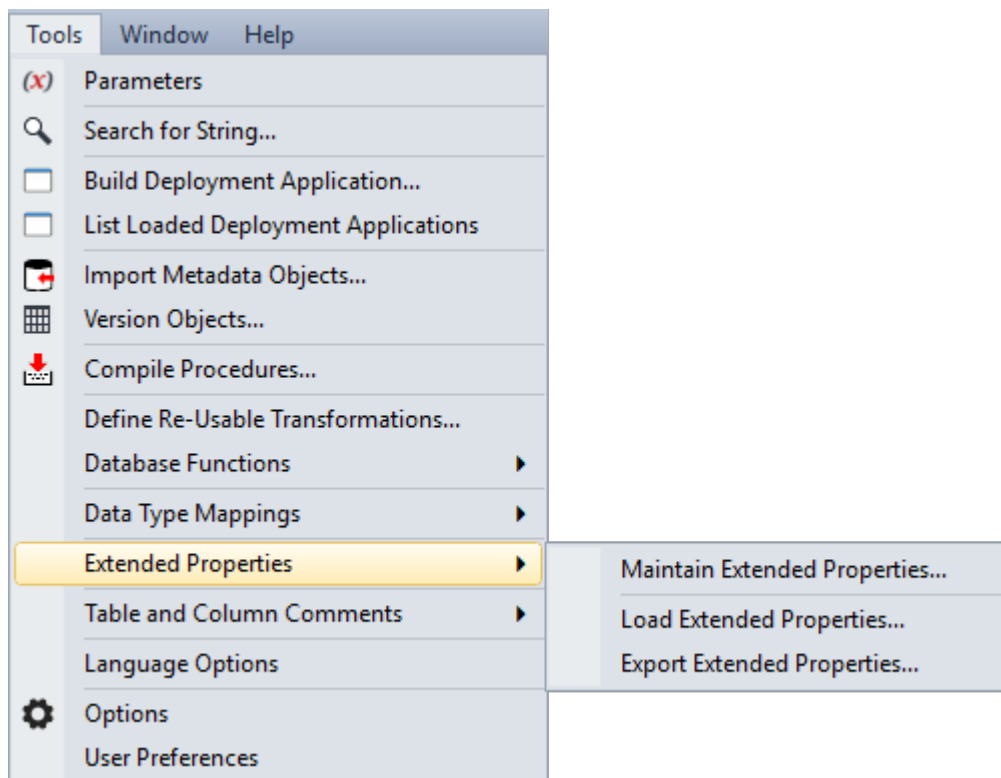
Defining Extended Properties.....	1066
Extended Properties Value Assignment	1072
Extended Properties Lookup.....	1073
Extended Properties Data Migration Between Repositories.....	1074

DEFINING EXTENDED PROPERTIES

An extended property definition mainly consists of the following attributes:

- **Variable** - the name of the extended property in the format it is accessed from the RED and included into scripts.
- **Scope** - the scope is a sub-set definition of an extended property variable which specifies the subset of database types and objects that will have access to the variable.

These user defined extended properties can be created and managed, using the **Extended Properties** facility on the **Tools** menu.



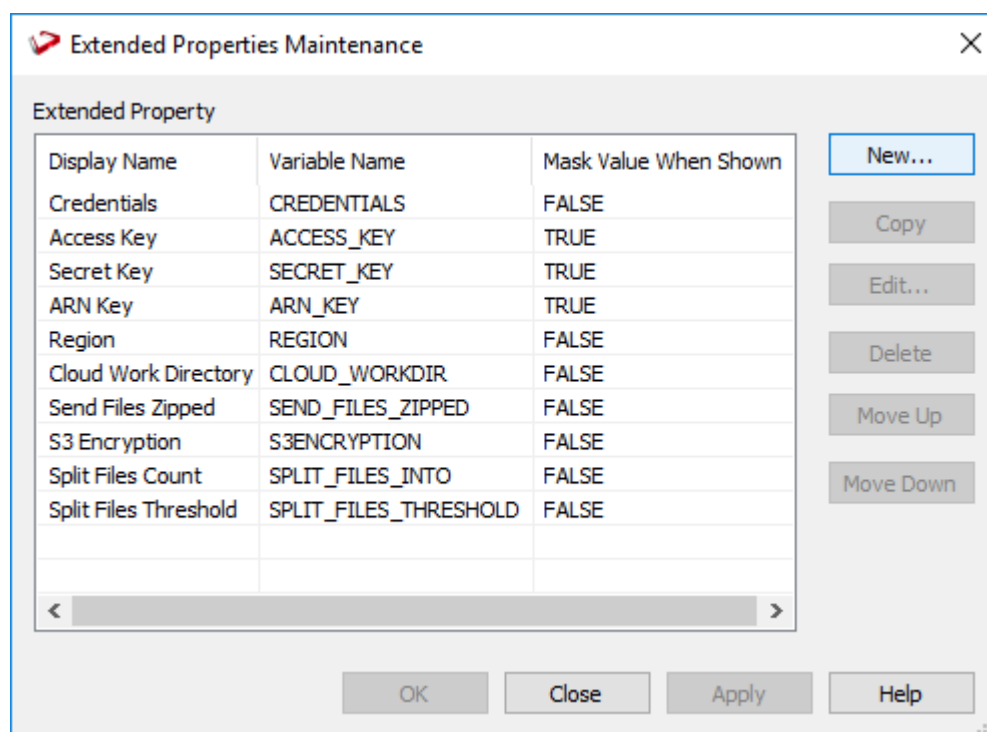
The following options are available:

- **Maintain Extended Properties** is used to add new or edit existing extended property definition settings.
- **Load Extended Properties** is used to load a set of extended property definitions into RED from an external source. This option is typically used with a WhereScape provided enablement pack.
- **Export Extended Properties** is used to export the available extended properties to an external properties file (.extprop).

CREATING AN EXTENDED PROPERTY DEFINITION

To create a new extended property definition, select **Tools>Extended Properties>Maintain Extended Properties**.

On the **Extended Properties Maintenance** screen, click **New**.



On the **Extended Property Definition** screen, enter the display and variable names and configure the settings as described below.

VARIABLE SETTINGS

Display Name

Enter a display name for the variable. This is the extended property name displayed in the Connection or Table properties.

Variable Name

Enter a unique variable name in the format it is accessed from the RED metadata service and included into scripts.

Mask Value when shown

Select this check box if you want to mask the value of the extended property, when displayed in the Connection or Table properties. The extended property value is also masked in the RED generated technical documentation.

Note: During an application creation process (via the **Tools>Build Deployment Application** menu), any extended property variables/assignments applied to a connection or table object are included in the generated application files. However, a connection or table object extended property with a masked value will not be included in the generated application files.

Description

Enter a description for the variable and its purpose/use.

SCOPE SETTINGS

Connection Visibility

Enables you to control the scope of the extended property definition, by selecting a **Connection File System** and/or a **Database** type that can access the extended property variable and their assigned value.

For example, if only the SQL Server check box is selected under the **Databases** option, then the variable will only appear in the Properties screen of connections with an SQL Server database type and not in the Properties screen of any other connections that is not an SQL Server database.

Object Visibility

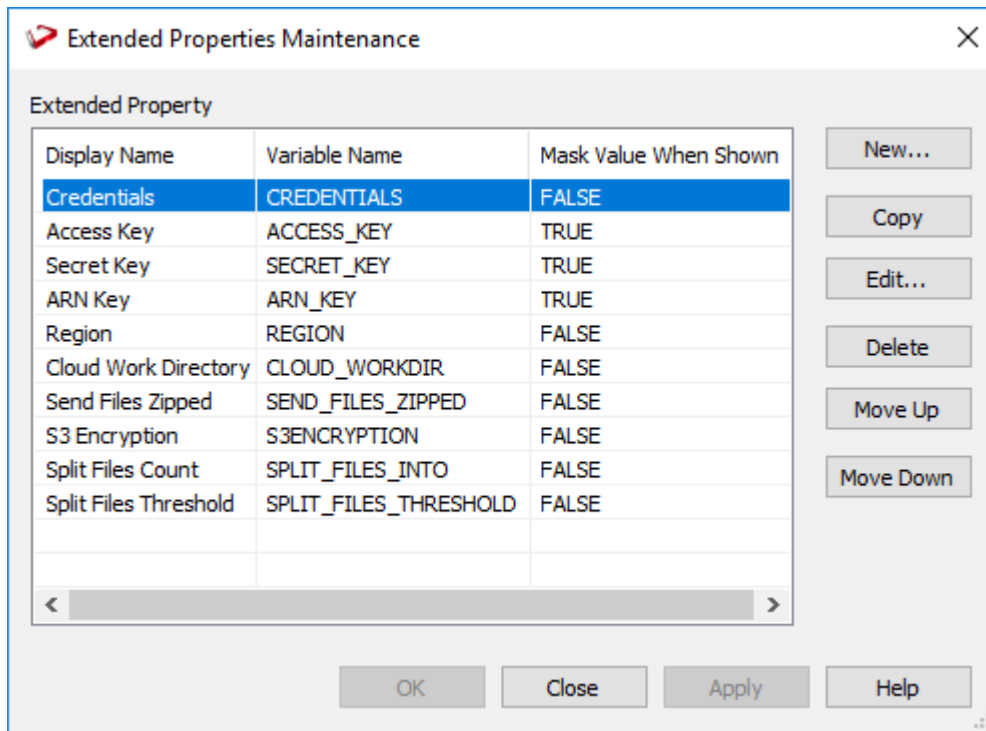
Enables you to control the scope of the extended property definition, by selecting an **Object** type that can access the extended property variable and their assigned value.

For example, if only the Load Table check box is selected under the **Objects** option, then the variable will only appear in the Properties screen of Load tables and not in the Properties screen of any other table objects.

MAINTAINING EXTENDED PROPERTY DEFINITIONS

The **Extended Properties Maintenance** screen displays a list of existing extended property definitions and enables you to add a **New** extended property definition. Clicking an extended property definition in the list enables you to **Copy**, **Edit** or **Delete** the selected item.

In addition, you can change the position of a selected extended property definition by moving it **Up** or **Down** within the list to specify a preferred order or grouping arrangement, when the extended property definitions are displayed in the Connection or Table properties screens. The display order has no impact in the operation.



The following describe the elements of the **Extended Properties Maintenance** screen.

Extended Property

This pane lists the extended property definitions that have been setup in RED. It details the display and variable name and indicates if the value of the extended property is masked, when displayed in the Connection or Table properties.

Function Buttons

These buttons enable you to create and manage the extended property definitions.

Button	Description
New	Launches the Extended Property Definition screen which enables you to create a new variable and specify the subset of database types and objects that will have access to the variable. Refer to <i>Creating an Extended Property Definition</i> (on page 1067) for details.
Copy	Duplicates the selected extended property and adds a _Copy1 suffix to the Display Name and Property Name . All fields are populated with the values of the original extended property. These can be edited as required.
Edit	Launches the Extended Property Definition screen which enables you to edit the selected extended property.
Delete	Deletes the selected extended property, if it is not currently in use. The user is warned if it is in use and asked to confirm or cancel the delete action.
Move Up/Move Down	Moves the selected extended property Up or Down the list. The display order of the extended property has no impact on its use.

Notes:

1. If an extended property is deleted, both the definition and all corresponding value assignments are deleted from the meta data repository. It is removed from the **Extended Properties** settings of the Connection or Table object Properties screen.
 2. If an extended property definition variable name is changed, the change is applied in the **Extended Properties** settings in the Connection or Table object Properties screen and its corresponding value assignment is retained.
 3. If an extended property definition scope is changed, its variable and corresponding value assignments are not deleted from the meta data repository. It is hidden from the **Extended Properties** settings in the Connection or Table object Properties screen. For example, an extended property variable has the Load object type set as in-scope and a corresponding value assignment has also been set in the Extended Properties setting in the Table Properties screen. If the extended property definition is edited to remove the Load object type from its scope, then the value assignment is not removed but becomes inactive and RED displays a warning that the scope change impacts a value assignment.
-

EXTENDED PROPERTIES VALUE ASSIGNMENT

After defining the extended property, you can assign values to each of the variable available in the Extended Properties tab of in-scope objects.

One extended property can have several connections and objects in its scope and therefore can have several values and assignments—one per connection and object.

An extended property assignment establishes the relationship between a variable and its value, which is specific for a particular connection or object.

SETTING UP EXTENDED PROPERTY VALUES FOR A CONNECTION

The extended property values for a connection are assigned and maintained in the **Extended Properties** tab of the **Connection Properties** screen.

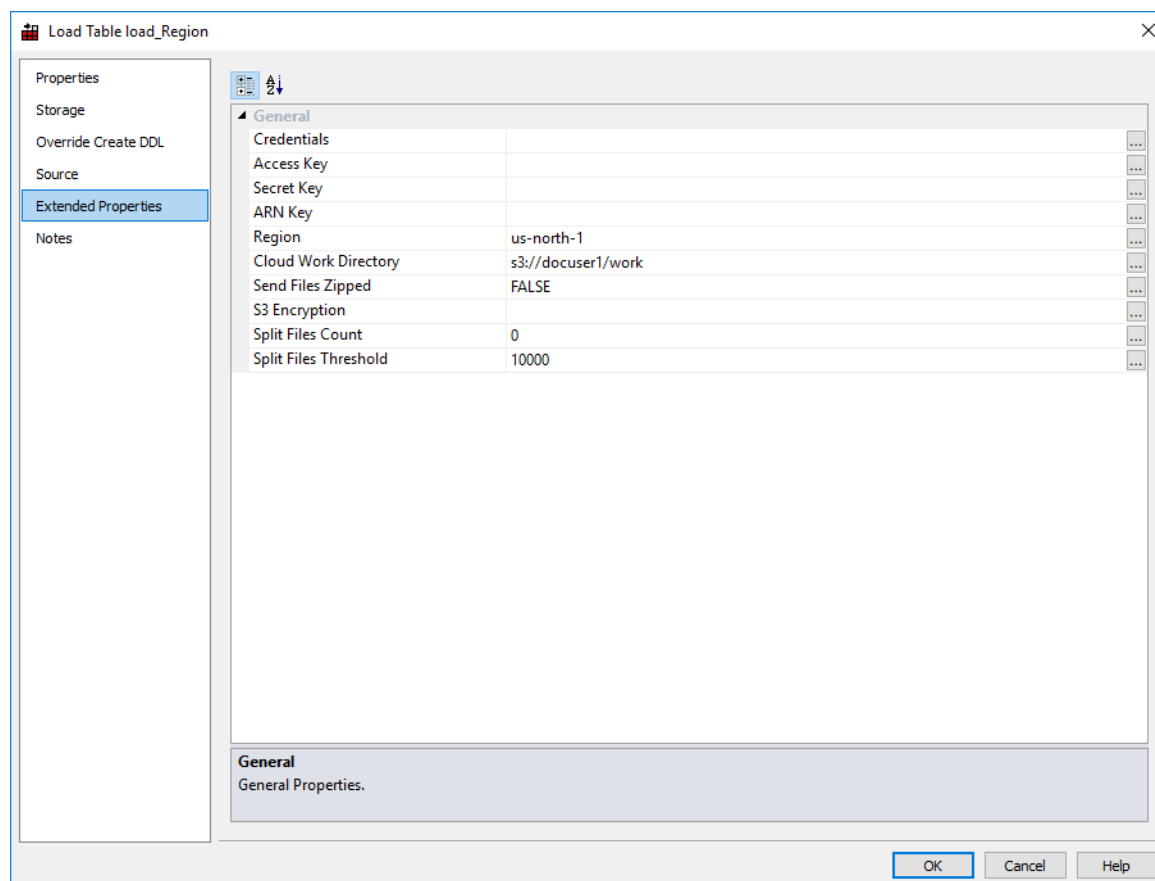
Connection Extended Properties

General	
Credentials	ACCESS
Access Key	*****
Secret Key	*****
ARN Key	*****
Region	us-west-2
Cloud Work Directory	s3://docuser1/work
Send Files Zipped	FALSE
S3 Encryption	
Split Files Count	0
Split Files Threshold	10000

SETTING UP EXTENDED PROPERTY VALUES FOR AN OBJECT

The extended property values for an object are assigned and maintained in the **Extended Properties** tab of the **Table Properties** screen.

WhereScape RED Object Extended Properties



The screenshot shows a dialog box titled "Load Table load_Region" with a sidebar on the left containing tabs: Properties, Storage, Override Create DDL, Source, Extended Properties (selected), and Notes. The main area displays the "General" section of the extended properties, which is expanded to show a table of properties and their values. Each row has a three-dot menu icon on the right.

Credentials		...
Access Key		...
Secret Key		...
ARN Key		...
Region	us-north-1	...
Cloud Work Directory	s3://docuser1/work	...
Send Files Zipped	FALSE	...
S3 Encryption		...
Split Files Count	0	...
Split Files Threshold	10000	...

At the bottom of the dialog, there are three buttons: OK, Cancel, and Help.

EXTENDED PROPERTIES LOOKUP

The extended property variables and their assigned values are referenced in PowerShell scripts. PowerShell script based processing is available for all objects in a custom database target. However, in a non-custom database target; PowerShell scripts are only available to process load and export objects.

WhereScape provides an application that enables access to extended property definitions and values. For more information, please contact your WhereScape representative.

EXTENDED PROPERTIES DATA MIGRATION BETWEEN REPOSITORIES

Once extended properties are defined and set, they can be propagated to the other RED repositories as follows:

The definitions are exported from the source repository and imported to the target repository via the **Tools>Extended Properties** menu.

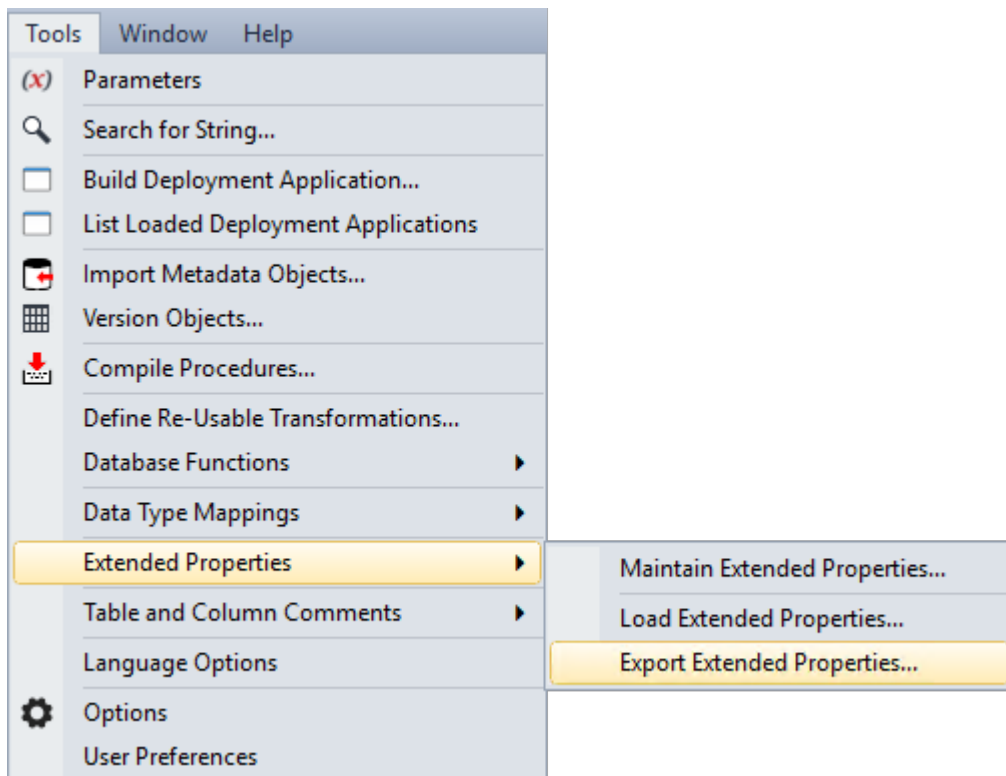
The values assigned to extended properties are included in a RED application that includes the parent Connection or Table object and are imported to the target repository, during the application load process.

Notes:

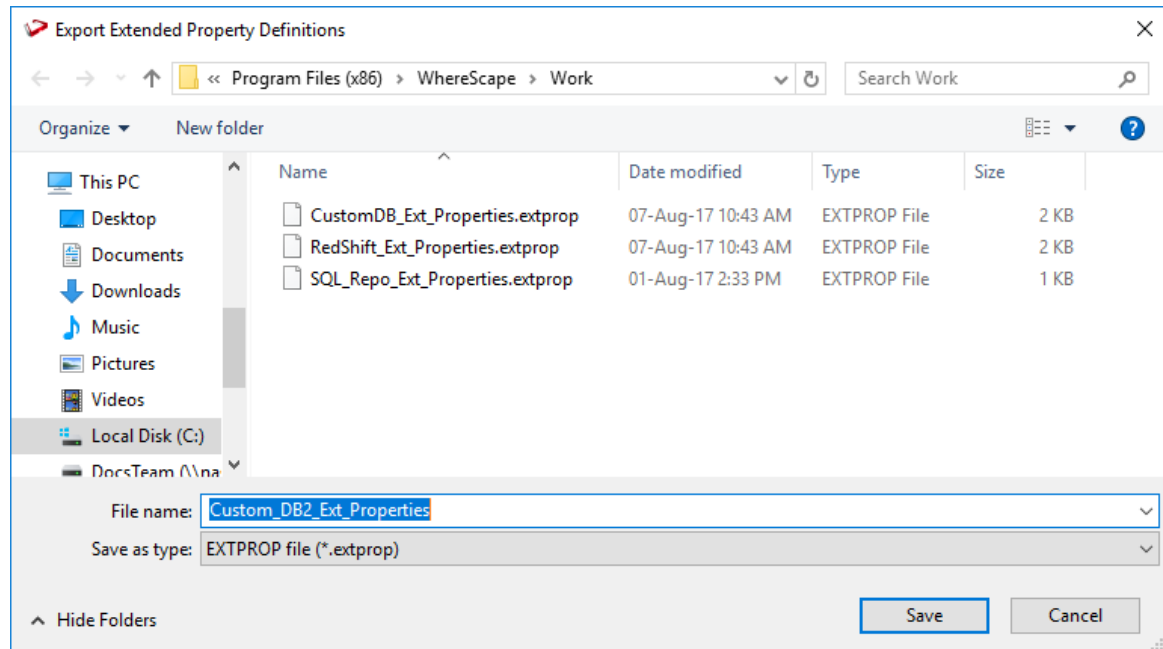
1. Any new or changed extended property definition must first be exported and then imported into the downstream repository, prior to loading an application that references them. Failing to load the definition first, will result in the value assignment not being loaded for the object.
 2. Extended properties are designed to be accessed from Windows PowerShell scripts, via the WhereScape RED metadata service. To aid this access, WhereScape provides a metadata service client library that customers are suggested to use to access the metadata service.
-

EXPORTING EXTENDED PROPERTIES

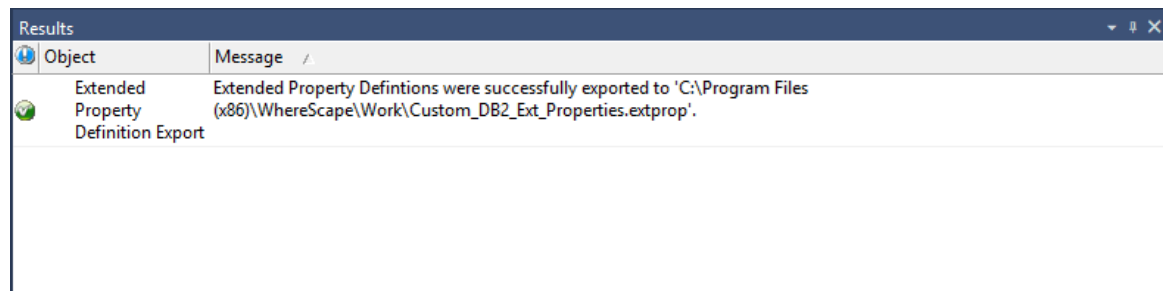
To export extended properties from RED, select **Tools>Extended Properties>Export Extended Properties**.



The following dialog is displayed. By default, RED exports the extended property (EXTPROP) file to **Program Files\WhereScape\Work**, but this can be changed. Enter the File name and click **Save**.

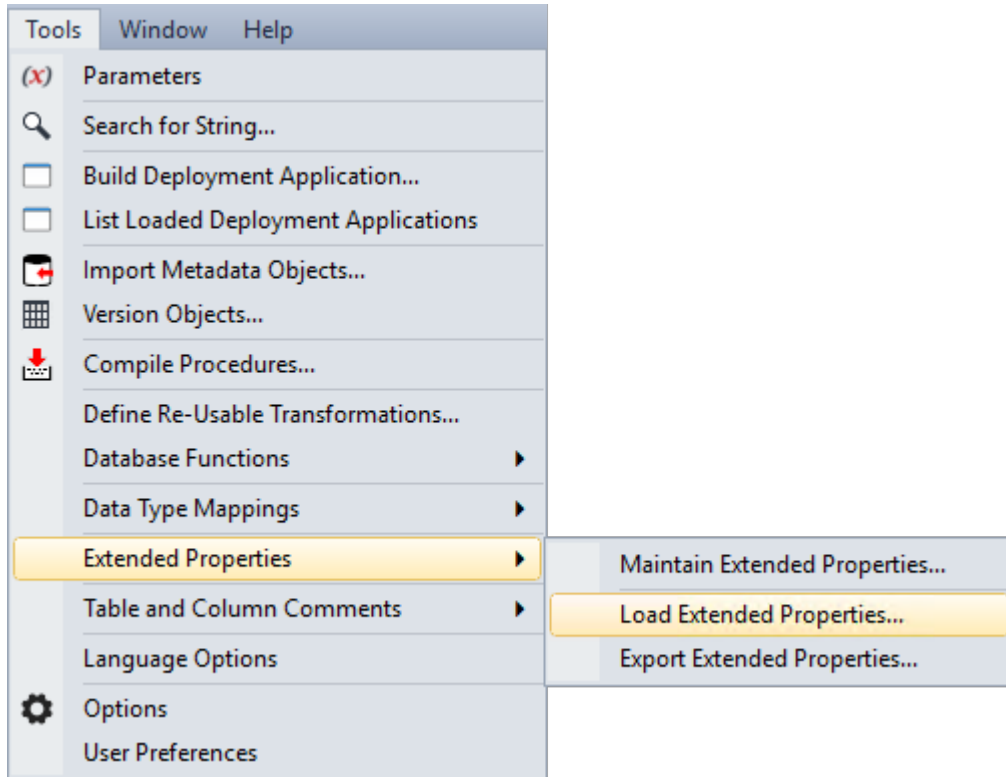


A confirmation message is displayed in the results pane.



LOADING EXTENDED PROPERTIES

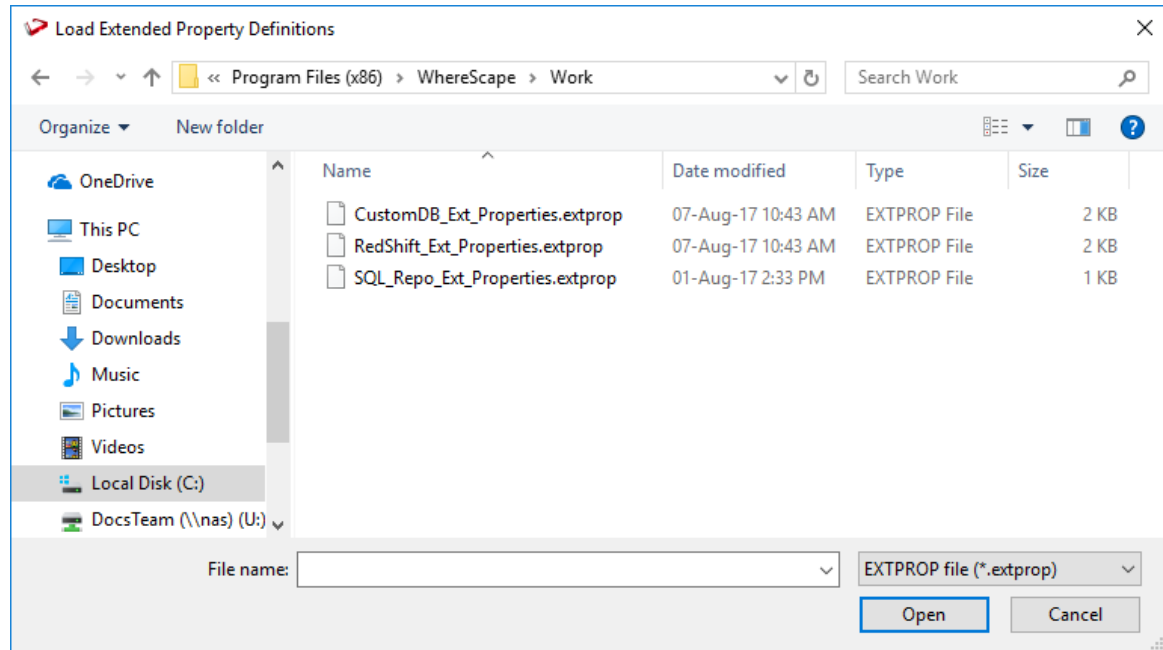
To load an extended property file, select **Tools>Extended Properties>Load Extended Properties**.



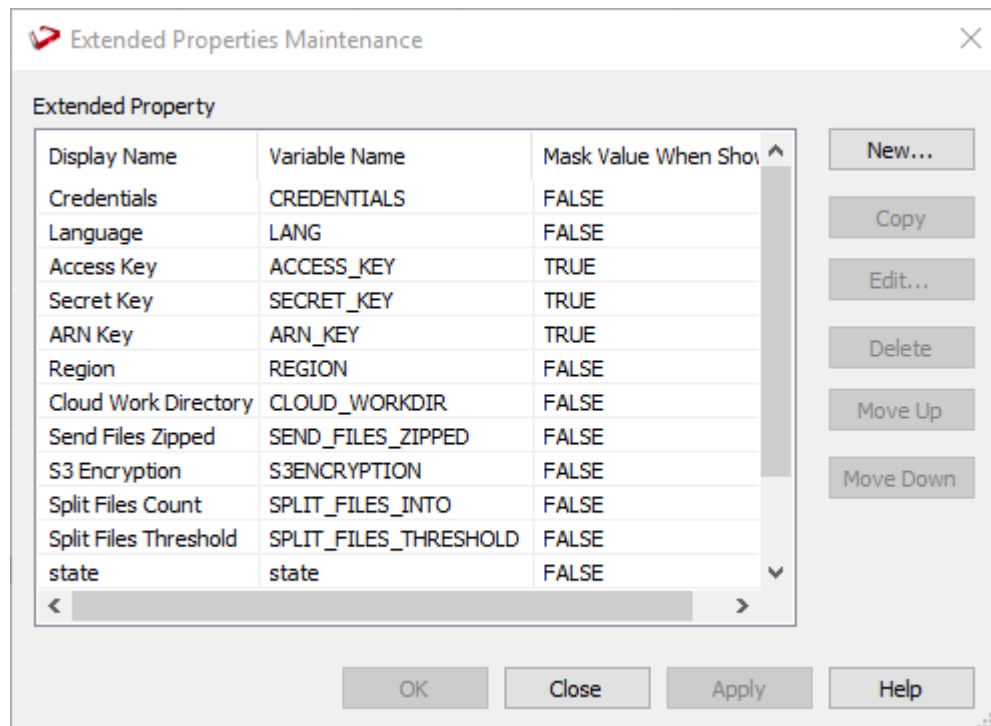
Note: The values assigned to extended properties are included in a RED application that includes the parent connection or table object and are imported to the target repository during the application load process.

An application including an object with an extended property value set, is dependent on the extended property definition being present in the target repository.

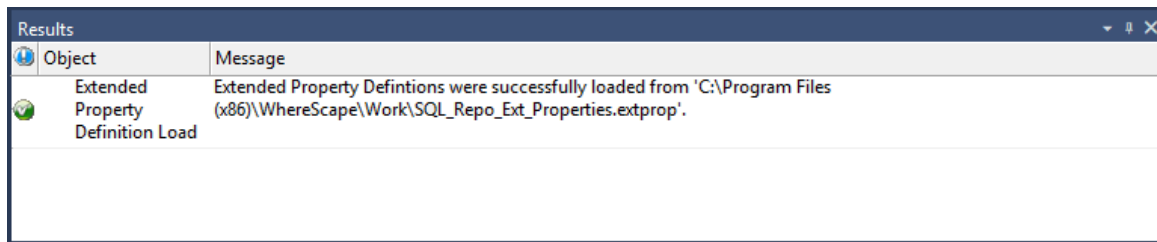
The following dialog is displayed. Select the extended property (EXTPROP) file to load the extended properties. By default RED expects the extprop files to be in **Program Files\WhereScape\Work** directory.



The selected extprop file is loaded and the attributes are listed in the **Extended Properties Maintenance** window.



A confirmation message is displayed in the results pane.



The screenshot shows a window titled 'Results' with a table containing one row of data. The table has two columns: 'Object' and 'Message'. The 'Object' column contains 'Extended Property Definition Load' and the 'Message' column contains 'Extended Property Defintions were successfully loaded from 'C:\Program Files (x86)\WhereScape\Work\SQL_Repo_Ext_Properties.extprop''. There is a green checkmark icon to the left of the object name.

Object	Message
Extended Property Definition Load	Extended Property Defintions were successfully loaded from 'C:\Program Files (x86)\WhereScape\Work\SQL_Repo_Ext_Properties.extprop'.

CHAPTER 46

TABLE AND COLUMN COMMENTS

WhereScape RED enables you to manage table and column comments outside your data warehouse environment and then load the updated comments back into the metadata repository; and subsequently copy them to the target database.

The comments can be exported from the metadata repository and imported back, via the **Tools>Table and Column Comments** menu option.

The export function enables you to output existing table or column comments into a Microsoft Excel format (.csv) file. You can open the file to view, modify or delete existing comments, as well as create new comments.

The import function enables you to load back the updated file into RED and perform a bulk update of the associated metadata.

The **Update Comments** context menu command enables you to copy comments from the metadata repository to the target database, e.g. updates the description and business display name fields (EUL table objects only) in the table and column Properties screen.

The following Object types are supported:

- Normal
- Hub
- Satellite
- Link
- Custom1
- Custom2
- Datastore
- Load
- Dimension
- Dimension View
- Stage
- Fact
- Fact KPI
- Aggregate
- Export
- Retro
- Retro Copy
- View
- Join

Notes:

- The table and column comments export/import feature only supports UTF-8 character encoded files.

- The exported data includes the Business Display Name (EUL table objects only) which can also be updated if required.

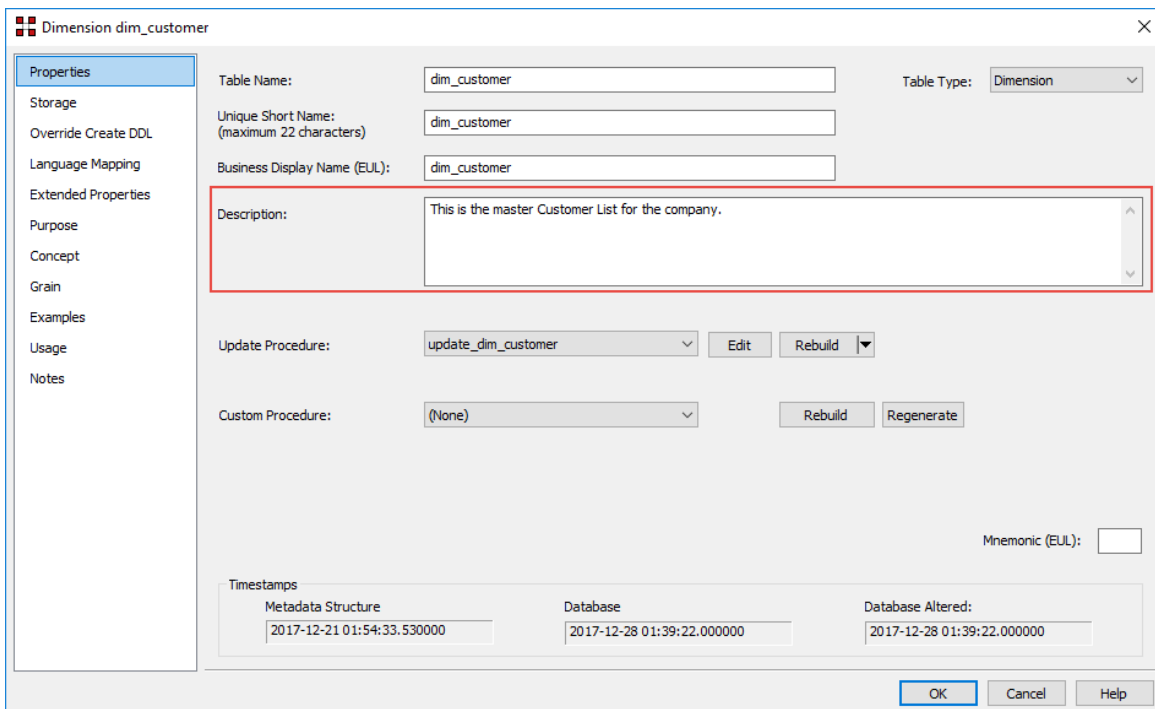
IN THIS CHAPTER

- Defining Table Comments1080
- Exporting Table Comments1081
- Loading Table Comments1084
- Defining Column Comments1087
- Exporting Column Comments1088
- Loading Column Comments1092
- Importing Comments from an External Source1096
- Viewing the Import/Export Logs.....1098

DEFINING TABLE COMMENTS

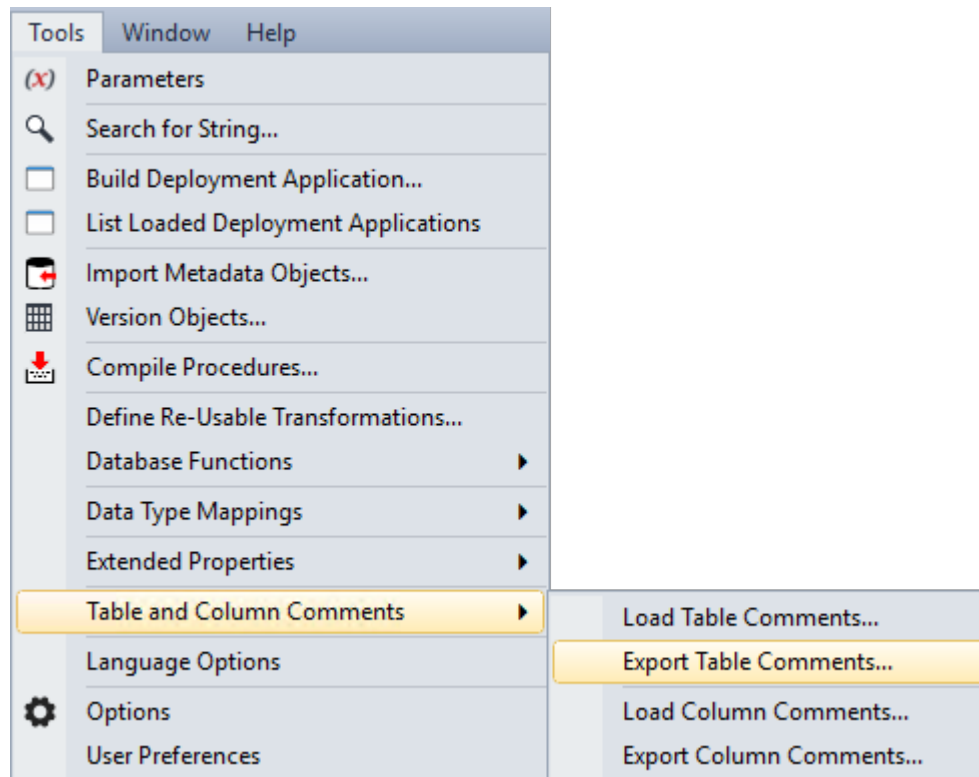
Table Properties

The table comments are defined in the **Description** field of the table's Properties. The comments in this field, along with the **Table Name** and **Business Display Name (EUL)** fields are used when exporting/importing table comments.




EXPORTING TABLE COMMENTS

To export table comments from RED, select **Tools>Table and Column Comments>Export Table Comments**.



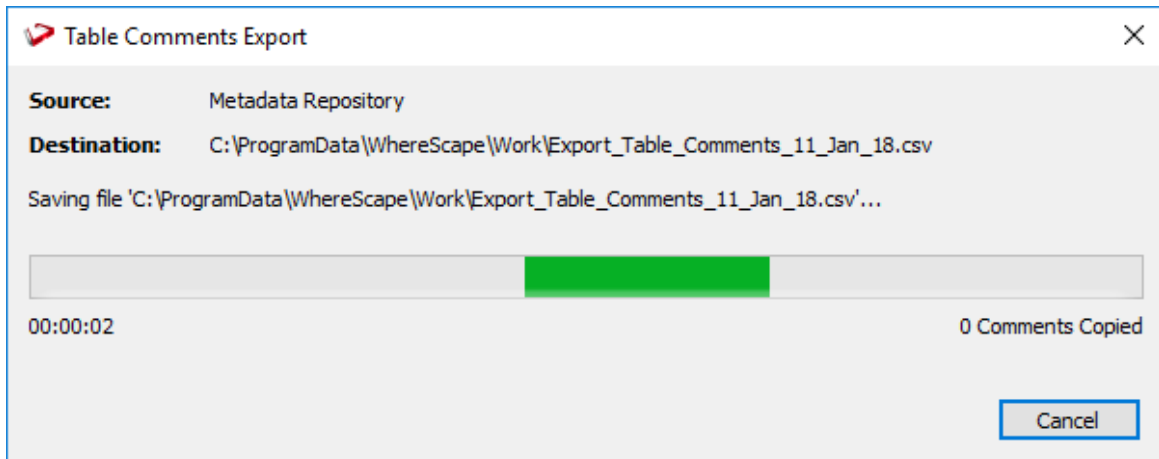
The **Table Comments Export** pop-up window is displayed:



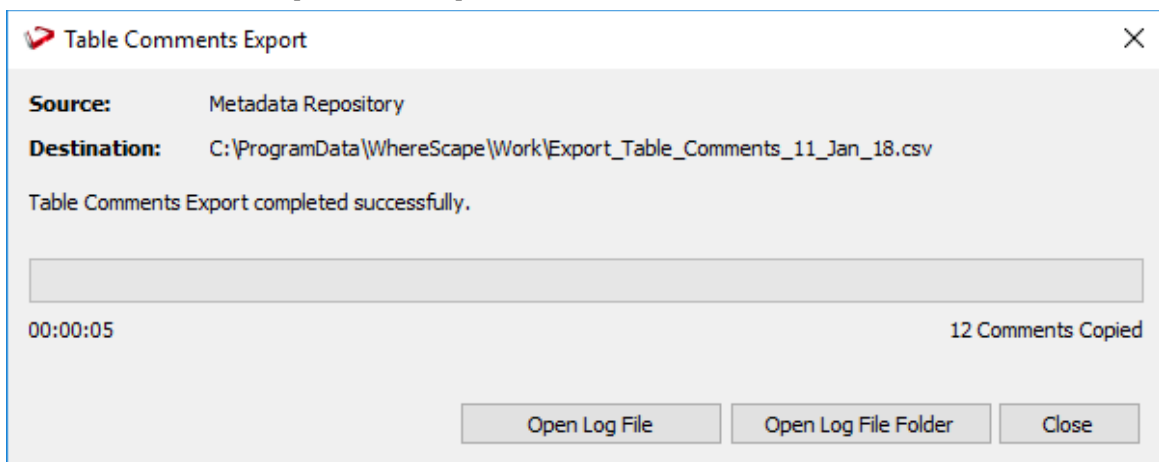
- 1 Type in the directory and file name of the export file or click the folder icon to navigate to the required directory.
- 2 Select or de-select the required export options.
 - **Only export comments from tables with a Business Display Name** - This option enables you to only process tables that have a **Business Display Name (EUL)** defined in the table's Properties screen. Clear the check box if you want to process all tables in your data warehouse.
 - **Only export comments that are not empty** - This option enables you to only process tables that have comments defined in **Description** field of the table's Properties screen. Clear the check box if you want to process all tables in your data warehouse.
 - **Only export comments from End User Layer (EUL) tables** - This option enables you to only process tables that are visible to the end user. Clear the check box if you want to process all tables in your data warehouse.
The visibility of table objects to end users in RED are set in *Tools>Options>Object Types>Object Type End User Setting* (see "*Object Type End User Setting*" on page 86).

All the above options are selected by default.

- 3 Click **OK** to proceed with the export. The **Table Comments Export** pop-up window displays the progress of your export. You can click **Cancel** to abort the process.

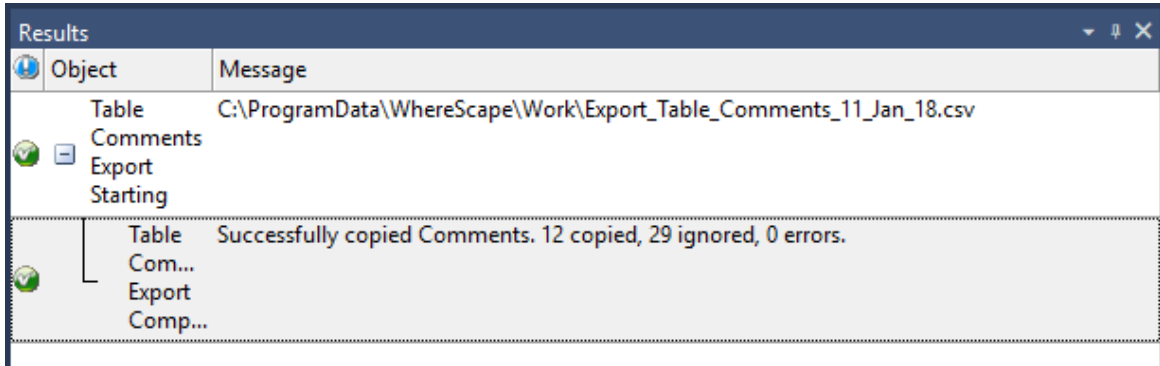


- 4 Once the export process is completed, the pop-window displays the processing time and the number of comments copied to the export file.

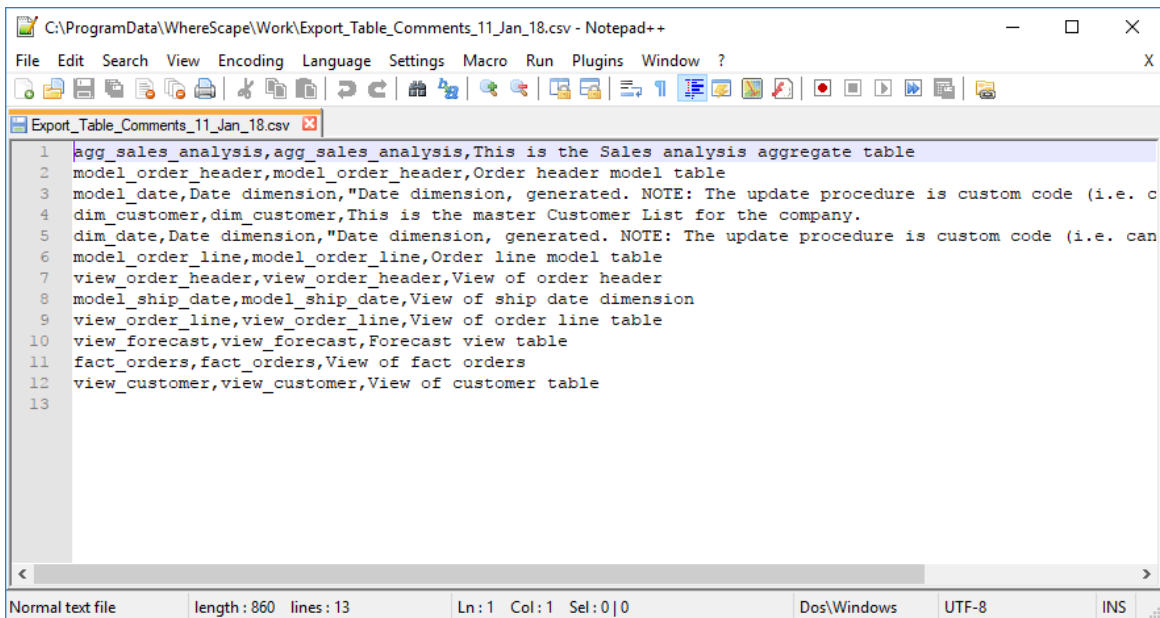


Note: After completing the export process, you can click **Open Log File** or **Open Log File Folder** to view the log file generated. Refer to *Viewing the Import/Export Logs* (on page 1098) for more details.

- 5 Click **Close** to exit the pop-up window. The **Results** pane displays a summary about the comments exported and confirms successful completion of the export.



- 6 Use Microsoft Excel 2013 or Notepad to view/edit the exported comments.



Each row of data includes the Table Name, Business Display Name (EUL table objects only) and Table Comments (Description).

Notes:

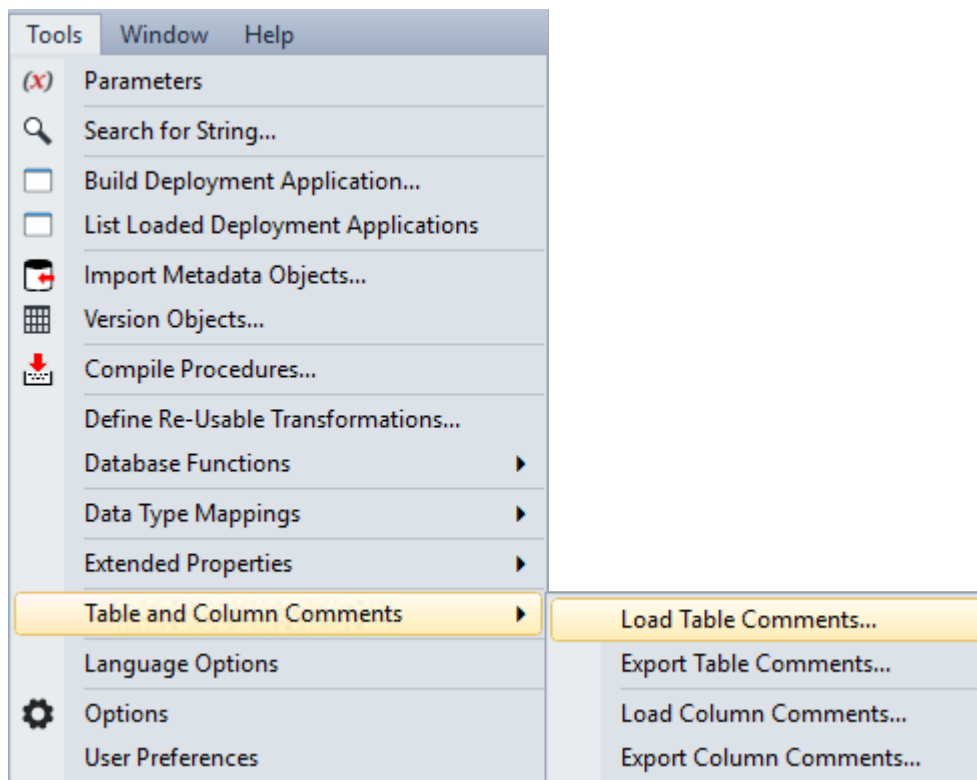
The table and column comments import/export only supports UTF-8 character encoded files. You also need to be aware of the following points, if you are going to use another tool other than Microsoft Excel 2013 to edit the comments:

- The comment portion can be surrounded in double quotes.
- Double quotes can be escaped within the comment by prefixing the quote with another quote.
- Comma characters can be used within the comment, but only if the comment is inside double quotes.
- Comments can span multiple lines, but only if the comment is inside double quotes.

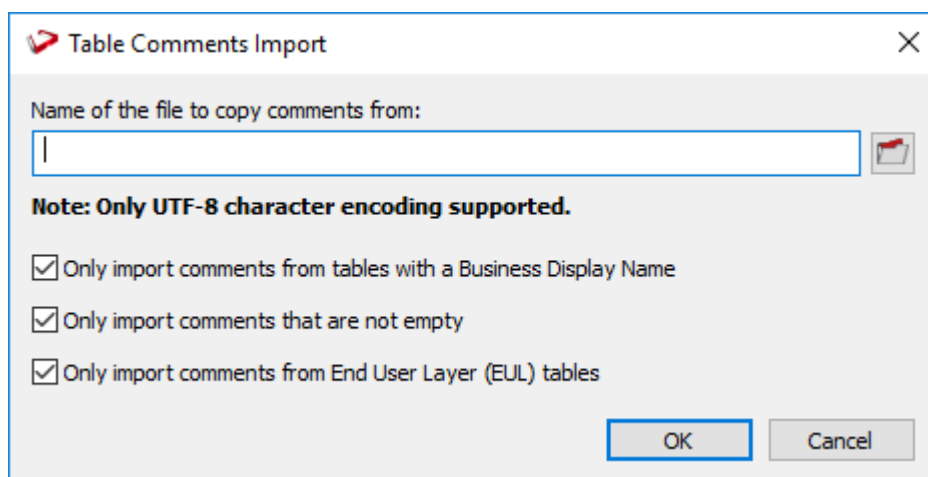
See *Importing Comments from an External Source* (on page 1096) for more details.

LOADING TABLE COMMENTS

To load/import table comments to RED from an external (.csv) file, select **Tools>Table and Column Comments>Load Table Comments**.



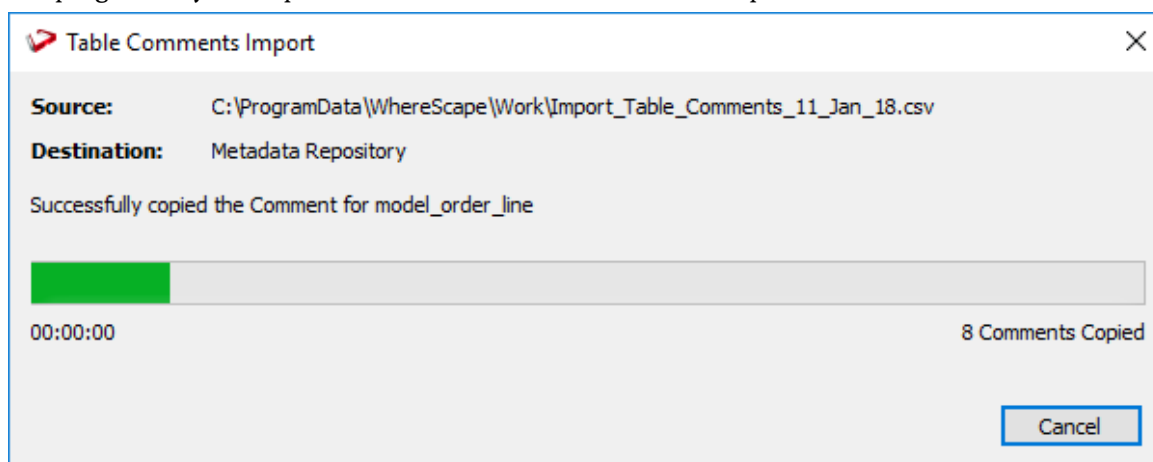
The **Table Comments Import** pop-up window is displayed:



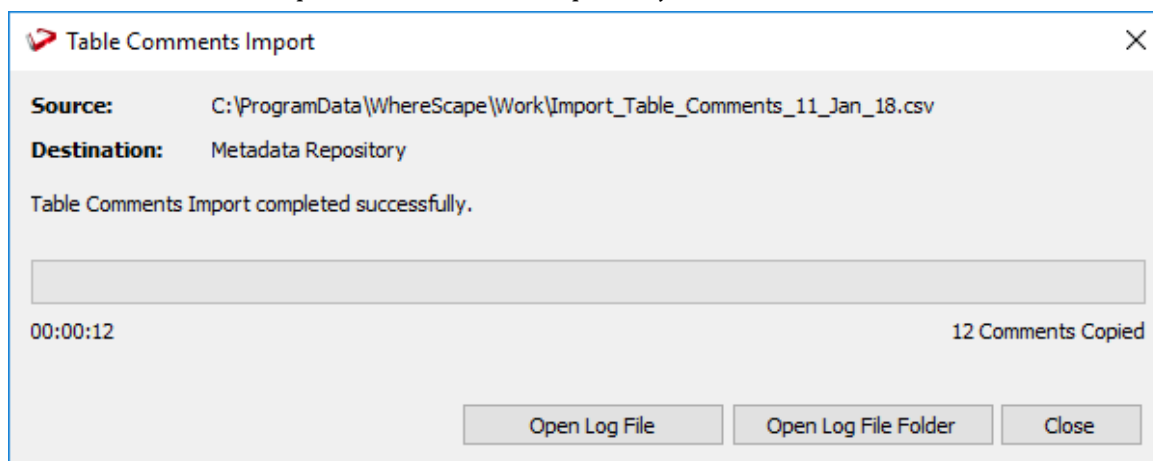
- 1 Type in the directory and file name of the source file or click the folder icon to navigate to the required directory.
- 2 Select or de-select the check box of the required import options.
 - **Only import comments from tables with a Business Display Name** - This option enables you to only process tables in the import file that have a Business Display Name defined. Clear the check box if you want to process all tables from your import file.
 - **Only import comments that are not empty** - This option enables you to only process tables that have comments defined in the import file. Clear the check box if you want to process all tables from your import file.
 - **Only import comments from End User Layer (EUL) tables** - This option enables you to only process tables in the import file that are visible to the end user. Clear the check box if you want to process all tables from your import file.
The visibility of table objects to end users in RED are set in *Tools>Options>Object Types>Object Type End User Setting* (see "*Object Type End User Setting*" on page 86).

All the above options are selected by default.

- 3 Click **OK** to proceed with the import. The **Table Comments Import** pop-up window displays the progress of your import. You can click **Cancel** to abort the process.



- 4 Once the import process is completed, the pop-window displays the processing time and the number of comments copied to the metadata repository.



Note: After completing the import process, you can click **Open Log File** or **Open Log File Folder** to view the log file generated. Refer to *Viewing the Import/Export Logs* (on page 1098) for more details.

- 5 Click **Close** to exit the pop-up window. The **Results** pane displays a summary about the comments imported and confirms successful completion of the import.

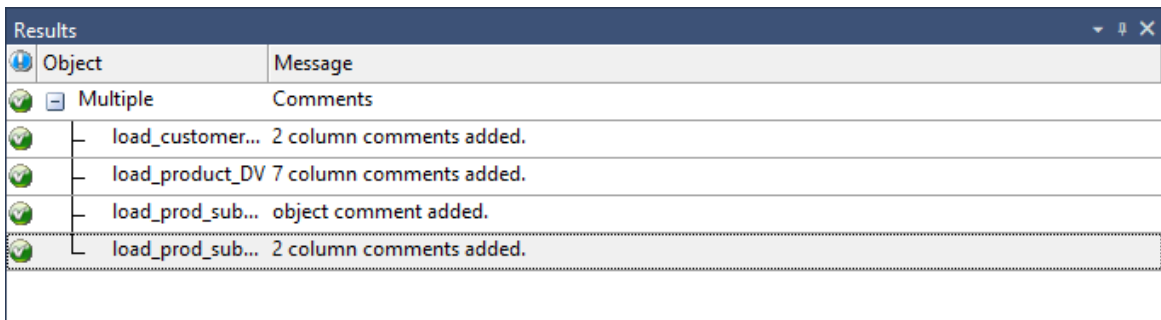
The screenshot shows a 'Results' window with a table containing two rows of information. The first row indicates the start of the import process for a table named 'Comments'. The second row shows the successful completion of the import, with 12 comments copied and no errors.

Object	Message
Table Comments Import Starting	C:\ProgramData\WhereScape\Work\Import_Table_Comments_11_Jan_18.csv
Table Com... Import Comp...	Successfully copied Comments. 12 copied, 0 ignored, 0 errors.

- 6 Use the **Update Comments** command from the context menu to copy the comments from the metadata repository to the corresponding tables in the target database.

The screenshot displays the WhereScape Data Warehouse interface. A context menu is open over a table in the 'DataWarehouse' pane, with the 'Update Comments' option highlighted. The interface includes a 'Development' pane on the left showing a tree view of objects, a central 'Load' pane, and a 'DataWarehouse' pane on the right showing a list of tables. The status bar at the bottom indicates the user is logged in as 'doc6871' and is browsing the 'DataWarehouse' source.

The **Results** pane displays the table(s) and column(s) that have added comments.



Object	Message
Multiple	Comments
load_customer...	2 column comments added.
load_product_DV	7 column comments added.
load_prod_sub...	object comment added.
load_prod_sub...	2 column comments added.

Notes:

The table and column comments import/export only supports UTF-8 character encoded files. You also need to be aware of the following points, if you are going to use another tool other than Microsoft Excel 2013 to edit the comments:

- The comment portion can be surrounded in double quotes.
- Double quotes can be escaped within the comment by prefixing the quote with another quote.
- Comma characters can be used within the comment, but only if the comment is inside double quotes.
- Comments can span multiple lines, but only if the comment is inside double quotes.

See *Importing Comments from an External Source* (on page 1096) for more details.

DEFINING COLUMN COMMENTS

Column Properties

The column comments are defined in the **Column Description** field of the column's Properties. The comments in this field, along with the **Table Name**, **Column Name**, and **Column Title/Business Display Name** fields are used when exporting/importing column comments.

The screenshot shows the 'Dimension Column dim_customer.dss_update_time' dialog box. The 'Column Description' field is highlighted with a red box and contains the text 'Date and time the row was updated in the data warehouse.' The dialog box is divided into several sections: General, Physical Definition, Meta Definition, Code Generation, and Source Details. The 'Column Description' field is located in the General section. Below the dialog box, there is a 'Column Description' section with a text area containing the same text as the dialog box.

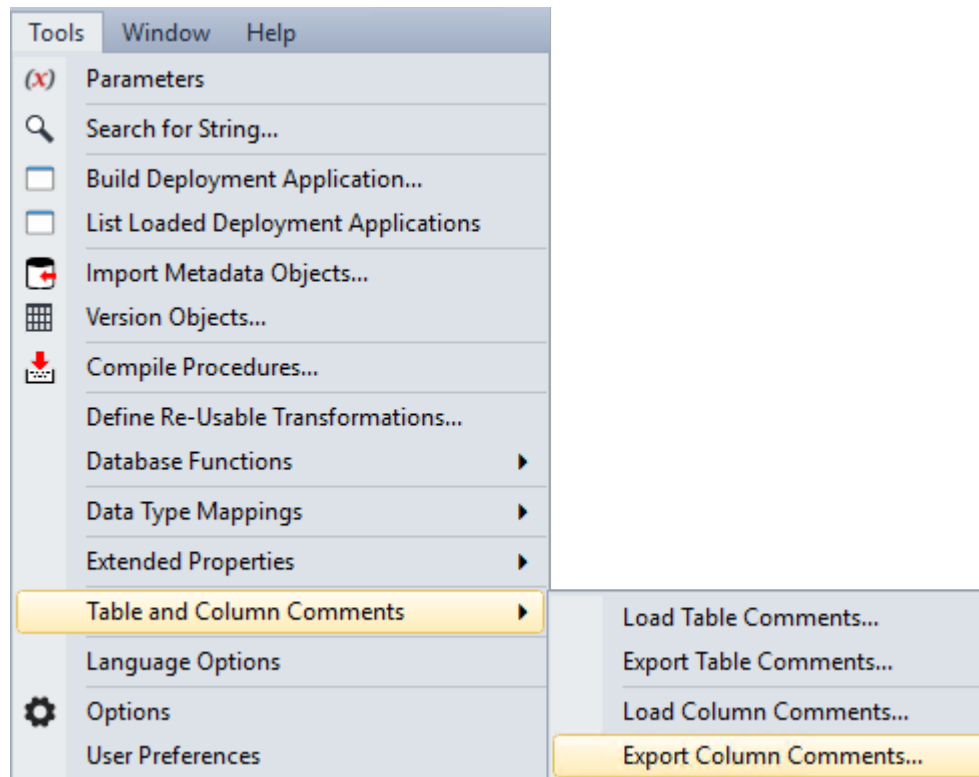
Section	Field	Value
General	Table Name	dim_customer
	Column Name	dss_update_time
	Column Title	dss update time
	Column Description	Date and time the row was updated in the data warehouse.
Physical Definition	Column Order	999999999
	Data Type	timestamp
	Null Values Allowed	<input checked="" type="checkbox"/>
	Default Value	
	Character Set	
	Format	
	Character Comparison/Sorting	
Meta Definition	Compress	<input type="checkbox"/>
	Numeric	<input type="checkbox"/>
	Additive	<input type="checkbox"/>
	Attribute	<input checked="" type="checkbox"/>
	End User Layer Display	<input type="checkbox"/>
	Business Key	<input type="checkbox"/>
Code Generation	Artificial Key	<input type="checkbox"/>
	Key Type	None
Source Details	Zero Key Value	

Column Description
This field contains the description for the column.
Dialog Opening Value: Date and time the row was updated in the data warehouse.

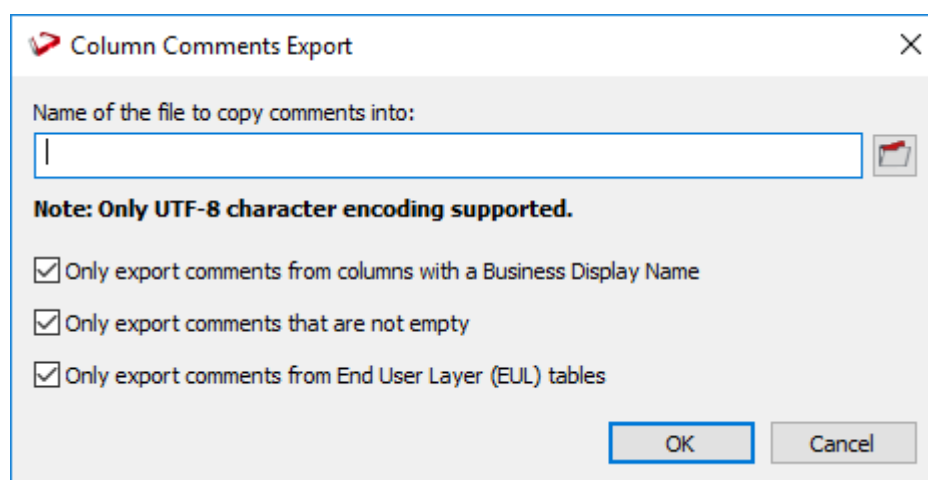
<- Update Update -> **OK** Cancel Help

EXPORTING COLUMN COMMENTS

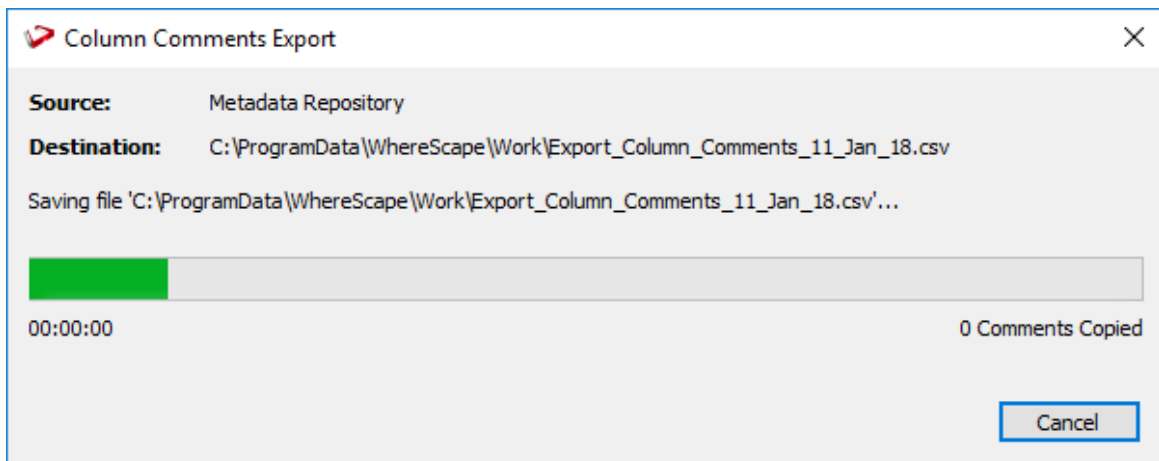
To export table comments from RED, select **Tools>Table and Column Comments>Export Column Comments**.



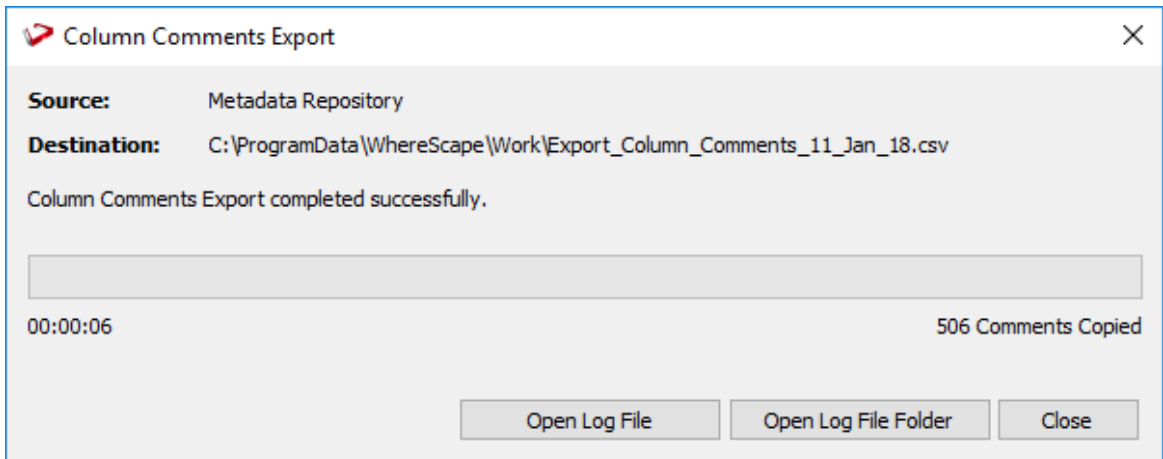
The **Column Comments Export** pop-up window is displayed:



- 1 Type in the directory and file name of the export file or click the folder icon to navigate to the required directory.
- 2 Select or de-select the required export options.
 - **Only export comments from columns with a Business Display Name** - This option enables you to only process columns that have a **Business Display Name (Column Title)** defined in the column's Properties screen. Clear the check box if you want to process all columns in your data warehouse.
 - **Only export comments that are not empty** - This option enables you to only process columns that have comments defined in **Column Description** field of the column's Properties screen. Clear the check box if you want to process all columns in your data warehouse.
 - **Only export comments from End User Layer (EUL) tables** - This option enables you to only process columns from tables that are visible to the end user. Clear the check box if you want to process all columns in your data warehouse.
The visibility of table objects to end users in RED are set in **Tools>Options>Object Types>Object Type End User Setting** (see "**Object Type End User Setting**" on page 86).All the above options are selected by default.
- 3 Click **OK** to proceed with the export. The **Column Comments Export** pop-up window displays the progress of your export. You can click **Cancel** to abort the process.

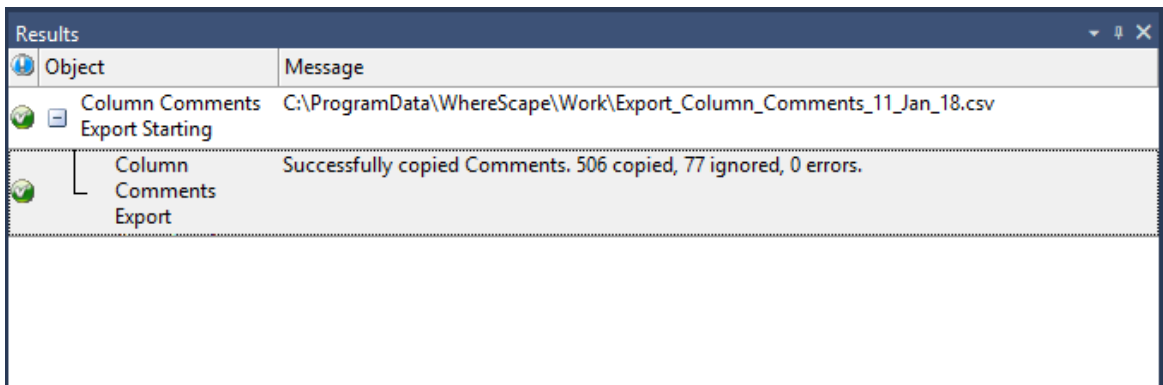


- Once the export process is completed, the pop-window displays the processing time and the number of comments copied to the export file.



Note: After completing the export process, you can click **Open Log File** or **Open Log File Folder** to view the log file generated. Refer to *Viewing the Import/Export Logs* (on page 1098) for more details.

- Click **Close** to exit the pop-up window. The **Results** pane displays a summary about the comments exported and confirms successful completion of the export.



6 Use Microsoft Excel 2013 or Notepad to view/edit the exported comments.

```

1 h_product_DV,code,code,Product code (SKU) that uniquely identifies the product.
2 ds_customer_2,code,code,The unique code which identifies the customer.
3 ds_customer_2,name,name,The full name of the customer.
4 ds_customer_2,address,address,The physical shipping address of the customer.
5 ds_customer_2,city,city,The city in which the customer sells the shipped product. Forms a hierarchy w
6 ds_customer_2,state,state,The state in which the customer sells the shipped product. Forms a hierarch
7 ds_customer_2,dss_update_time,dss update time,Date and time the row was updated in the data warehouse
8 ds_customer_2,dss_create_time,dss create time,Date and time the row was created in the data warehouse
9 dim_customer,code,code,This is the customer code that is used to identify each customer.
10 dim_ship_date,ship_date,ship date,The calendar date for this row.
11 dim_ship_date,s_cal_day_in_week,scal day in week,"The day in the week. Format DDD. Example: mon, tue,
12 dim_ship_date,s_cal_day_in_week_no,scal day in week no,The day number in the week. 1-7 where Sunday i
13 dim_ship_date,s_cal_day_in_month,scal day in month,The day in the month 1-31.
14 dim_ship_date,s_cal_day_in_year,scal day in year,The day in the year 1-366.
15 dim_ship_date,s_cal_week_in_year,scal week in year,The week in the year 0-53.
16 dim_ship_date,s_cal_month_no,scal month no,The calendar month number 1-12.
17 dim_ship_date,s_cal_month,scal month,The calendar month representation. Format YYYYMM. Example: 20020
18 dim_ship_date,s_cal_month_name,scal month name,"The calendar month name. Format MON. Examples: jan, f
19 dim_ship_date,s_cal_quarter_no,scal quarter no,The calendar quarter number 1-4.
20 dim_ship_date,s_cal_quarter,scal quarter,The calendar quarter representation. Format YYYYQQ. Example

```

Each row of data includes the Table Name, Column Name, Column Title/Business Display Name and Column Comments (Description).

Notes:

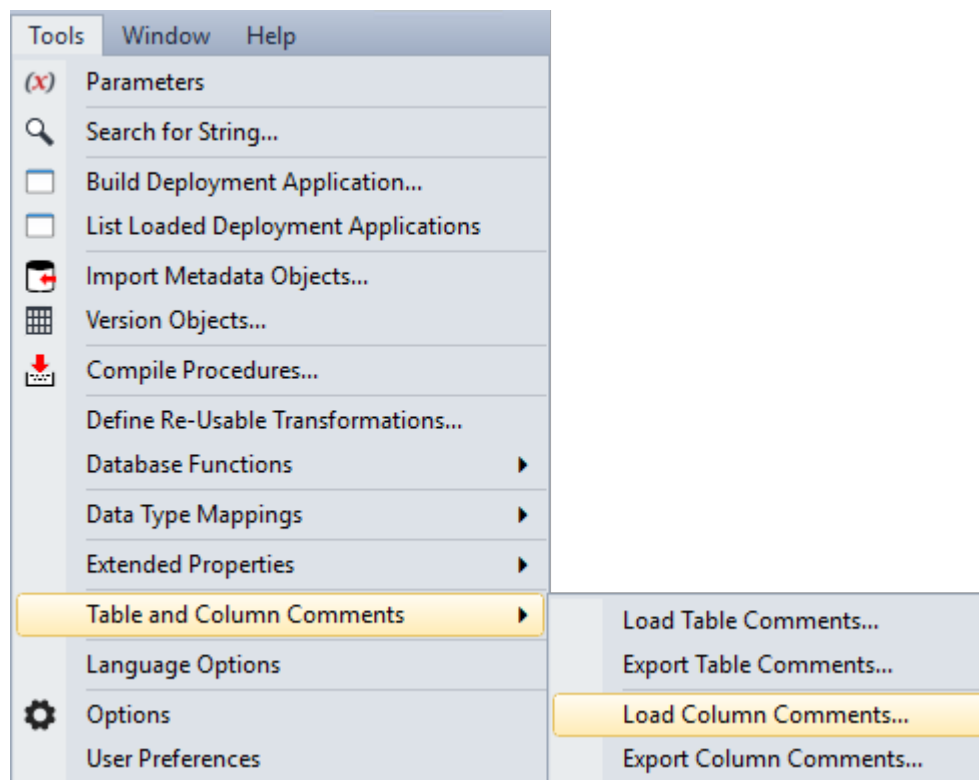
The table and column comments import/export only supports UTF-8 character encoded files. You also need to be aware of the following points, if you are going to use another tool other than Microsoft Excel 2013 to edit the comments:

- The comment portion can be surrounded in double quotes.
- Double quotes can be escaped within the comment by prefixing the quote with another quote.
- Comma characters can be used within the comment, but only if the comment is inside double quotes.
- Comments can span multiple lines, but only if the comment is inside double quotes.

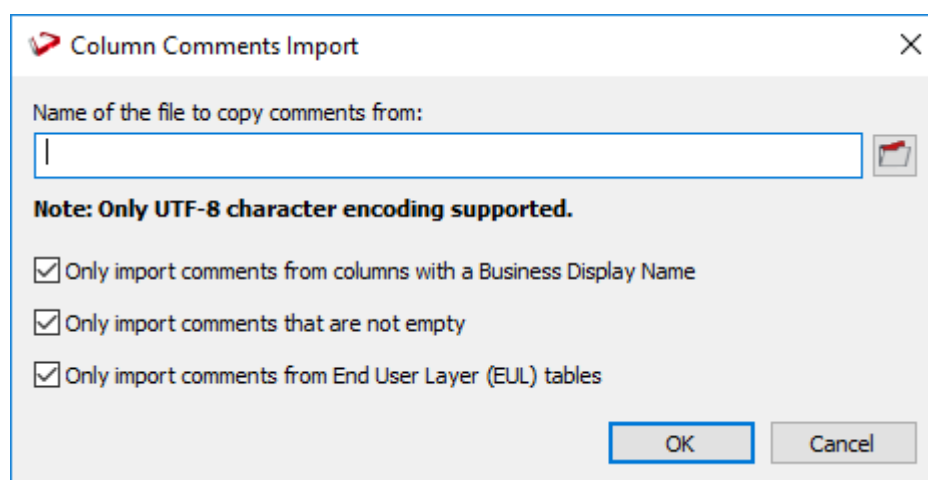
See *Importing Comments from an External Source* (on page 1096) for more details.

LOADING COLUMN COMMENTS

To load/import column comments to RED from an external (.csv) file, select **Tools>Table and Column Comments>Load Column Comments**.



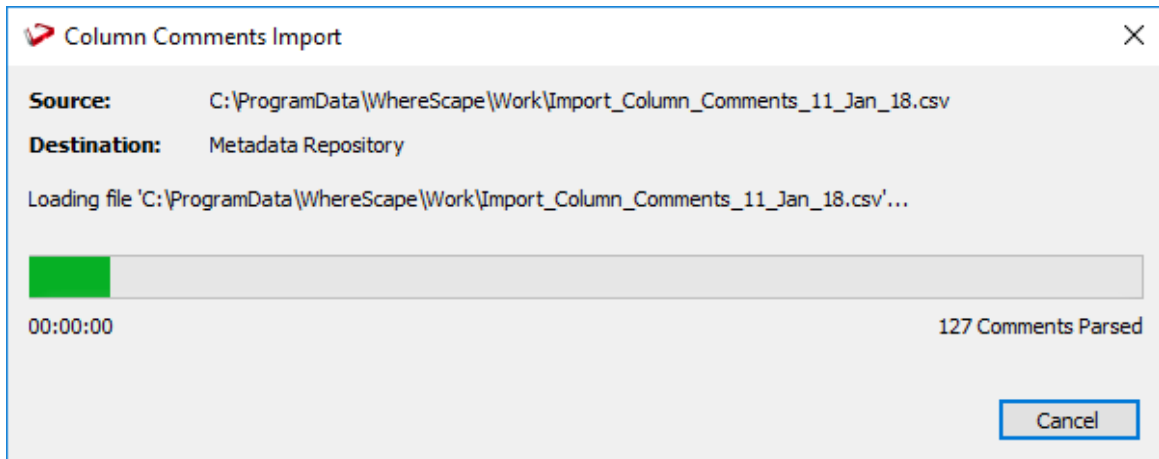
The **Column Comments Import** pop-up window is displayed:




- 1 Type in the directory and file name of the source file or click the folder icon to navigate to the required directory.
- 2 Select or de-select the check box of the required import options.
 - **Only import comments from tables with a Business Display Name** - This option enables you to only process columns in the import file that have a Business Display Name/Column Title defined. Clear the check box if you want to process all tables from your import file.
 - **Only import comments that are not empty** - This option enables you to only process columns that have comments defined in the import file. Clear the check box if you want to process all tables from your import file.
 - **Only import comments from End User Layer (EUL) tables** - This option enables you to only process columns from tables in the import file that are visible to the end user. Clear the check box if you want to process all tables from your import file.
The visibility of table objects to end users in RED are set in **Tools>Options>Object Types>Object Type End User Setting** (see "**Object Type End User Setting**" on page 86).

All the above options are selected by default.

- 3 Click **OK** to proceed with the import. The **Column Comments Import** pop-up window displays the progress of your import. You can click **Cancel** to abort the process.

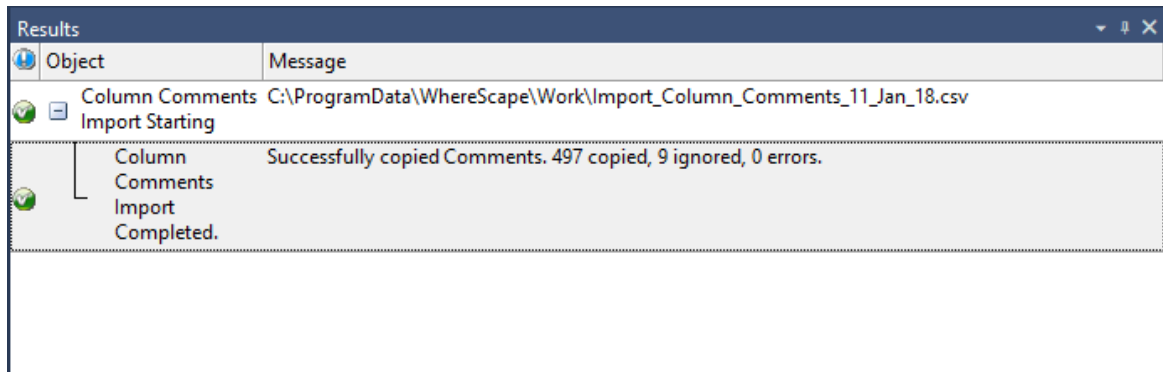


- 4 Once the import process is completed, the pop-window displays the processing time and the number of comments copied to the metadata repository.

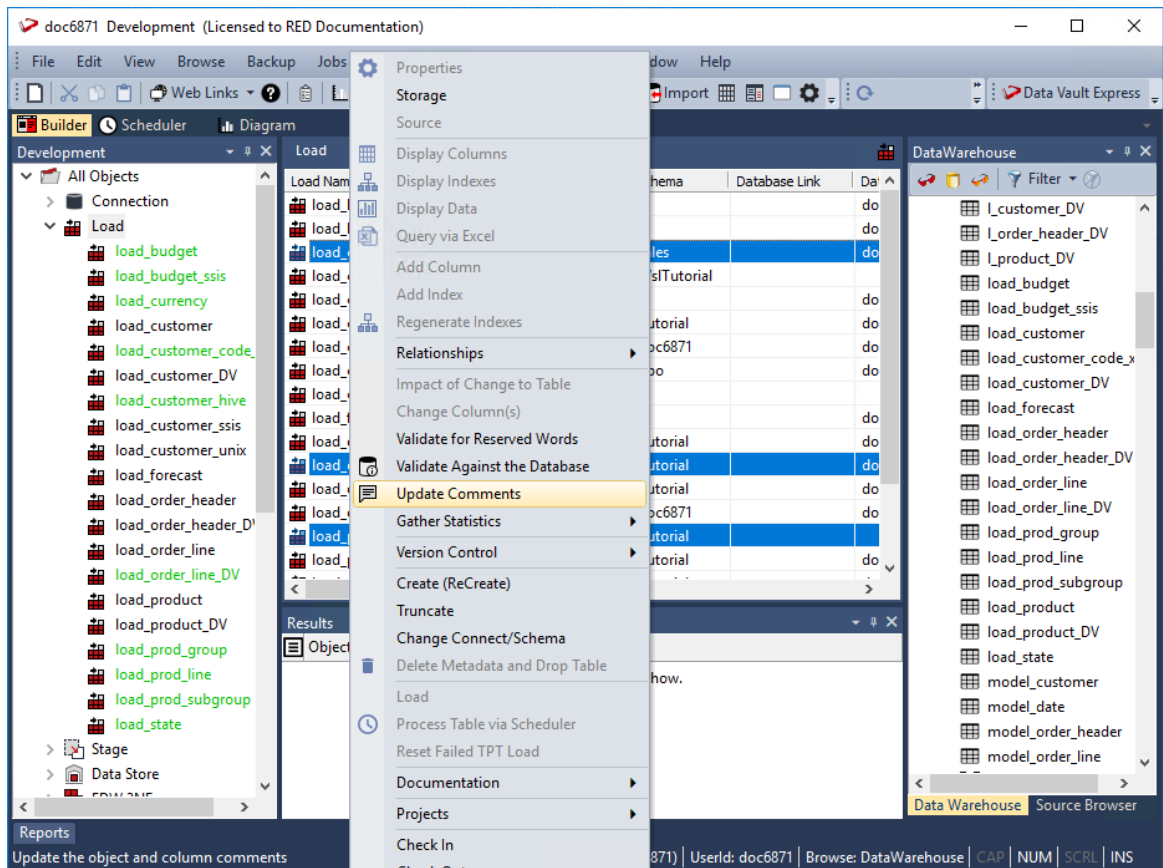


Note: After completing the import process, you can click **Open Log File** or **Open Log File Folder** to view the log file generated. Refer to *Viewing the Import/Export Logs* (on page 1098) for more details.

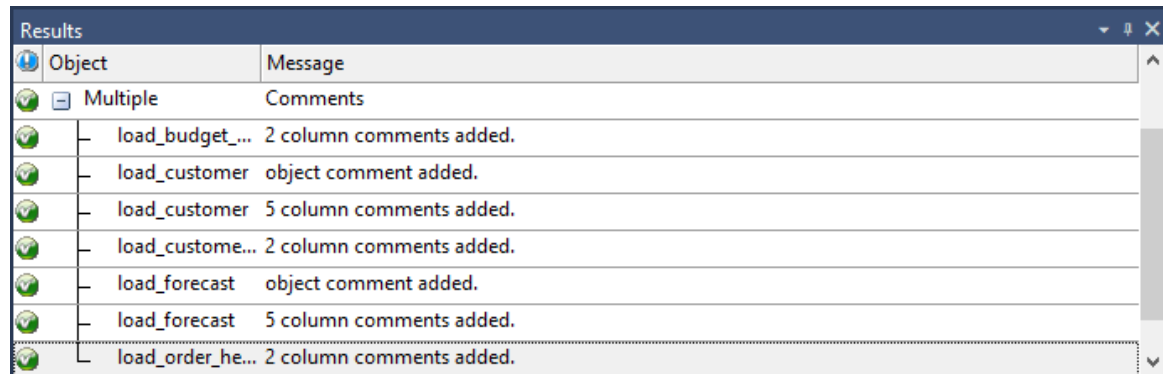
- 5 Click **Close** to exit the pop-up window. The **Results** pane displays a summary about the comments imported and confirms successful completion of the import.



- 6 Use the **Update Comments** command from the context menu, to copy the comments from the metadata repository to the corresponding columns in the target database.



The **Results** pane displays the table(s) and column(s) that have added comments.



Object	Message
Multiple	Comments
load_budget_...	2 column comments added.
load_customer	object comment added.
load_customer	5 column comments added.
load_custome...	2 column comments added.
load_forecast	object comment added.
load_forecast	5 column comments added.
load_order_he...	2 column comments added.

Notes:

The table and column comments import/export only supports UTF-8 character encoded files. You also need to be aware of the following points, if you are going to use another tool other than Microsoft Excel 2013 to edit the comments:

- The comment portion can be surrounded in double quotes.
- Double quotes can be escaped within the comment by prefixing the quote with another quote.
- Comma characters can be used within the comment, but only if the comment is inside double quotes.
- Comments can span multiple lines, but only if the comment is inside double quotes.

See *Importing Comments from an External Source* (on page 1096) for more details.

IMPORTING COMMENTS FROM AN EXTERNAL SOURCE

The import function supports Microsoft Excel 2013 comma separated value (.csv) file format for importing comments to RED.

Table Comments

Each row of data includes the Table Name, Business Display Name (EUL table objects only) and Table Comments (Description).

```
view_order_header,view_order_header,View of order header.
view_order_line,view_order_line,View of order line table.
view_forecast,view_forecast,View of Forecast table.
```

Column Comments

Each row of data includes the Table Name, Column Name, Column Title/Business Display Name and Column Comments (Description).

```
view_order_header,order_number,order_number,A code which identifies the order.
view_order_line,order_line_no,order_line_no,A number which identifies a line on
the order.
view_forecast,forecast_quantity,forecast_quantity,Quantity of product
forecasted (i.e. number of product units forecasted to be sold).
```

Notes:

- The comment portion can be surrounded in double quotes. For example:

```
view_order_header,view_order_header,"View of order header."
view_order_line,view_order_line,"View of order line table."
view_forecast,view_forecast,"View of Forecast table."
```

- Double quotes can be escaped within the comment by prefixing the quote with another quote. For example:

```
view_order_header,view_order_header,"View ""of"" order header."
view_order_line,view_order_line,View of order line table.
view_forecast,view_forecast,View of Forecast table.
```


- Comma characters can be used within the comment, but only if the comment is inside double quotes. For example:

```
view_order_header,view_order_header,"View of order header,comment."  
view_order_line,view_order_line,"View of order line table,comment."  
view_forecast,view_forecast,"View of Forecast table,comment."
```

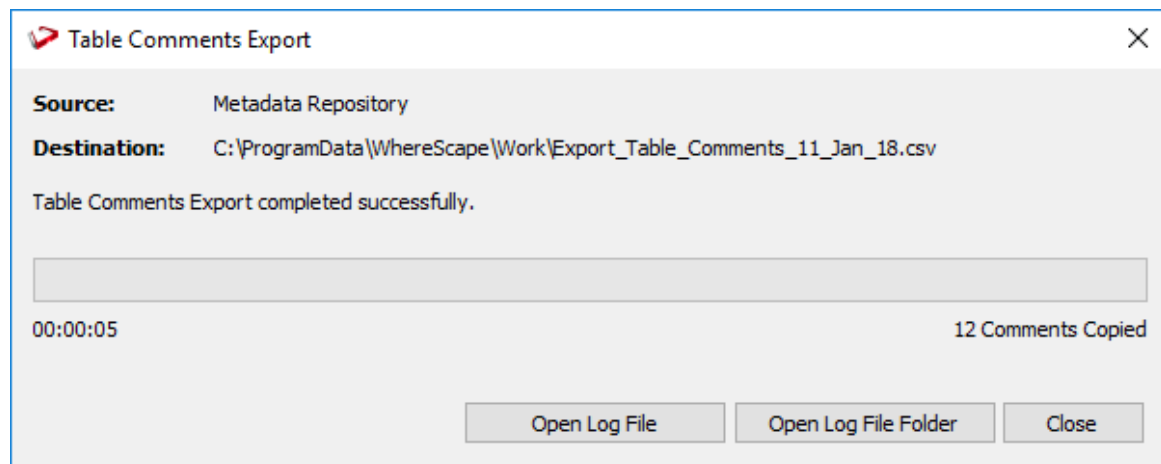
- Comments can span multiple lines, but only if the comment is inside double quotes. For example:

```
view_order_header,view_order_header,"View of order header comment."  
view_order_line,view_order_line,"View of order line  
table comment."  
view_forecast,view_forecast,"View of Forecast  
table comment."
```

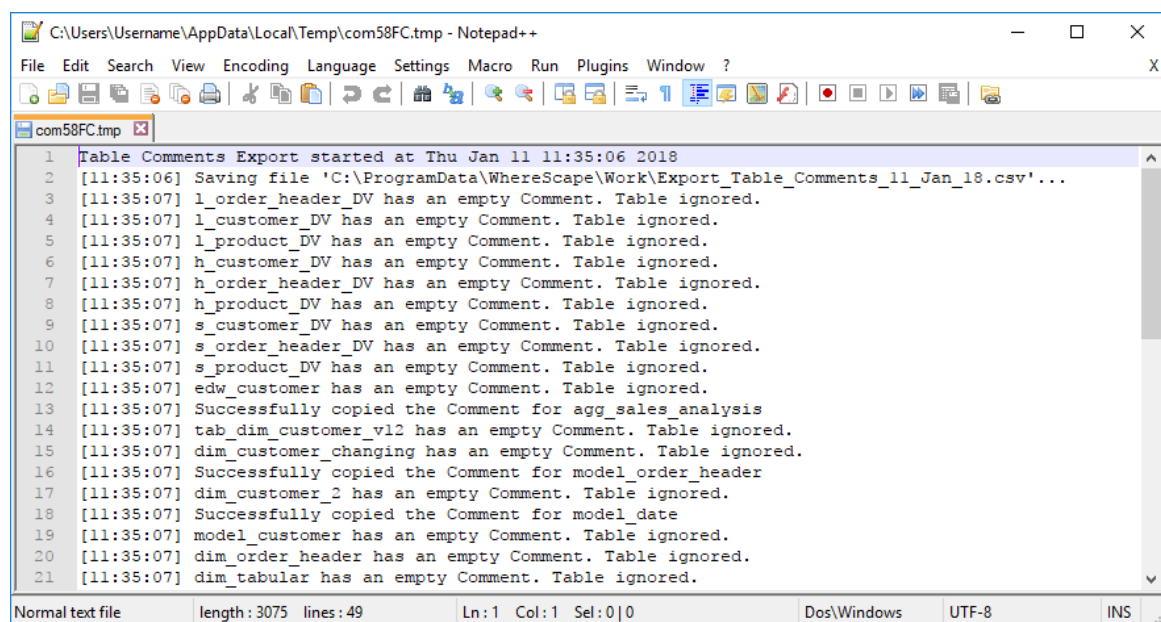
VIEWING THE IMPORT/EXPORT LOGS

RED creates a log to aid in troubleshooting, in case issues are encountered during the import/export process.

After completing the import/export process, you can click **Open Log File** or **Open Log File Folder** from the **Table Comments Import/Export** pop-up window to view the log file generated.



The import/export logs are stored in your Windows Temp directory (c:\Users\Username\AppData\Local\Temp).



CHAPTER 47

MULTI SOURCE PROCESSING

WhereScape RED enables you to process data from more than one source query into a single consolidated table. The multi-source processing capability is designed to provide users with flexibility in the definition of component source statements and the ability to independently process multi source tables that are associated with the target table.

IN THIS CHAPTER

Multi Source Functions and Features	1101
Adding Source Mapping to Objects	1106
Maintaining Source Column Mappings	1112
Generating Update Procedures for Source Mapping Objects.....	1115
Executing Update Procedures via Scheduler	1119
Reverting to Non Source Mapping Object	1121

MULTI SOURCE FUNCTIONS AND FEATURES

The following describes the WhereScape RED functions and features that are used to support multi source processing of data warehouse objects created in RED.

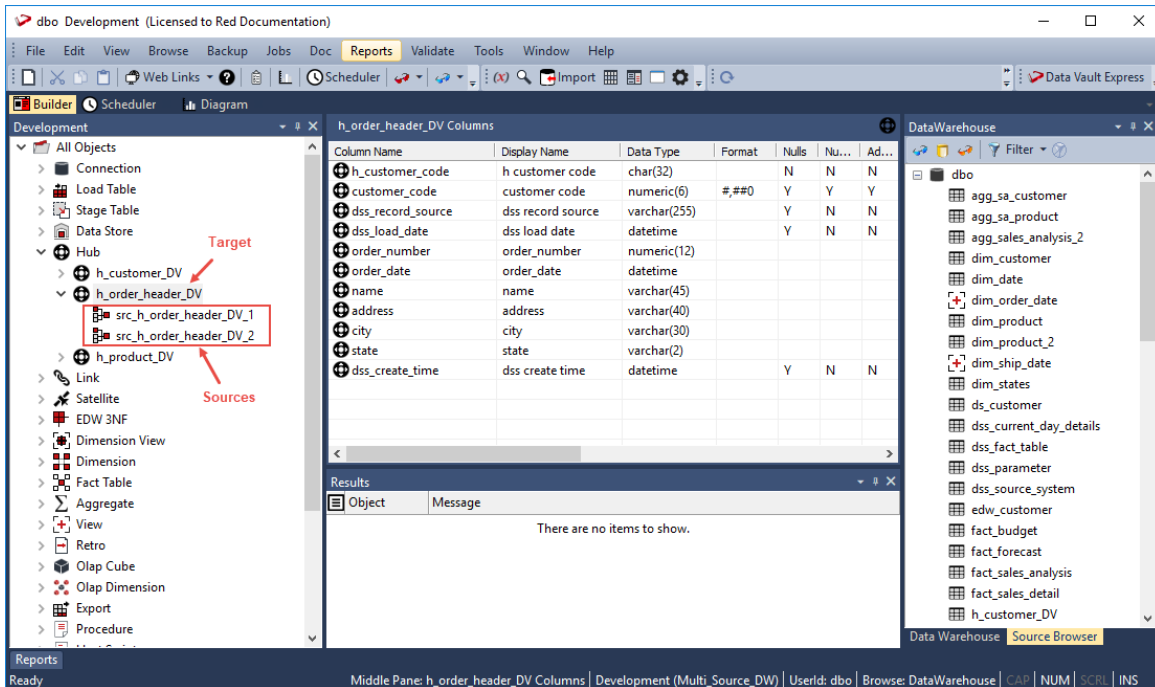
Note: The source mapping feature can be used with all RED table objects, except for Tabular objects, partitioned tables and Load tables.

SOURCE MAPPING OBJECT

Source Mapping objects are child objects that are used in WhereScape RED to map columns from one or more source tables to an existing target table in RED.

A Source Mapping object is built from the Data Warehouse connection and can be created from one or more source tables.

Refer to the section *Adding Source Mapping to Objects* (see "47.2 Adding Source Mapping to Objects" on page 1106) for details.

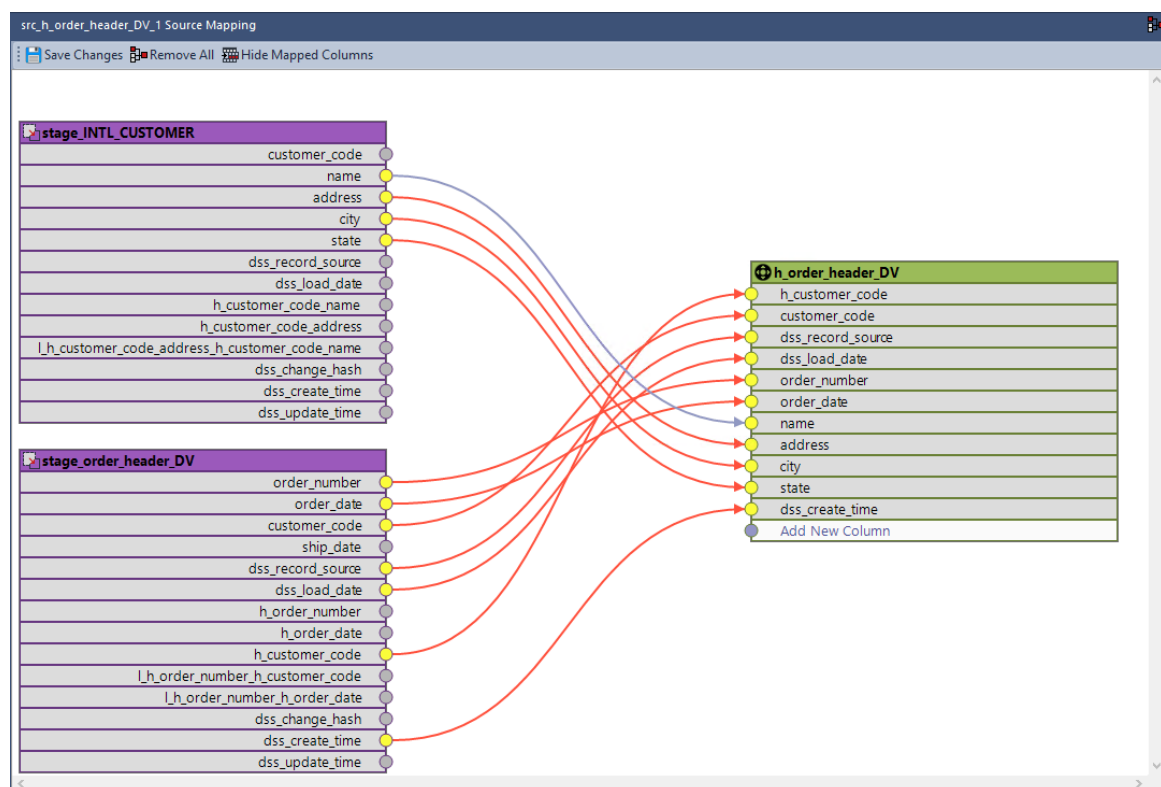


SOURCE MAPPING TOOL

A Source Mapping tool is available in WhereScape RED which enables you to graphically map columns from one or more source tables to an existing target table in RED.

This tool provides a graphical representation of the mapped columns between the source tables and the target table. The columns from the source table that exists in the target table are automatically mapped. You can edit the mappings by clicking a source column connection point and dragging the line that appears to the target column's connection point.

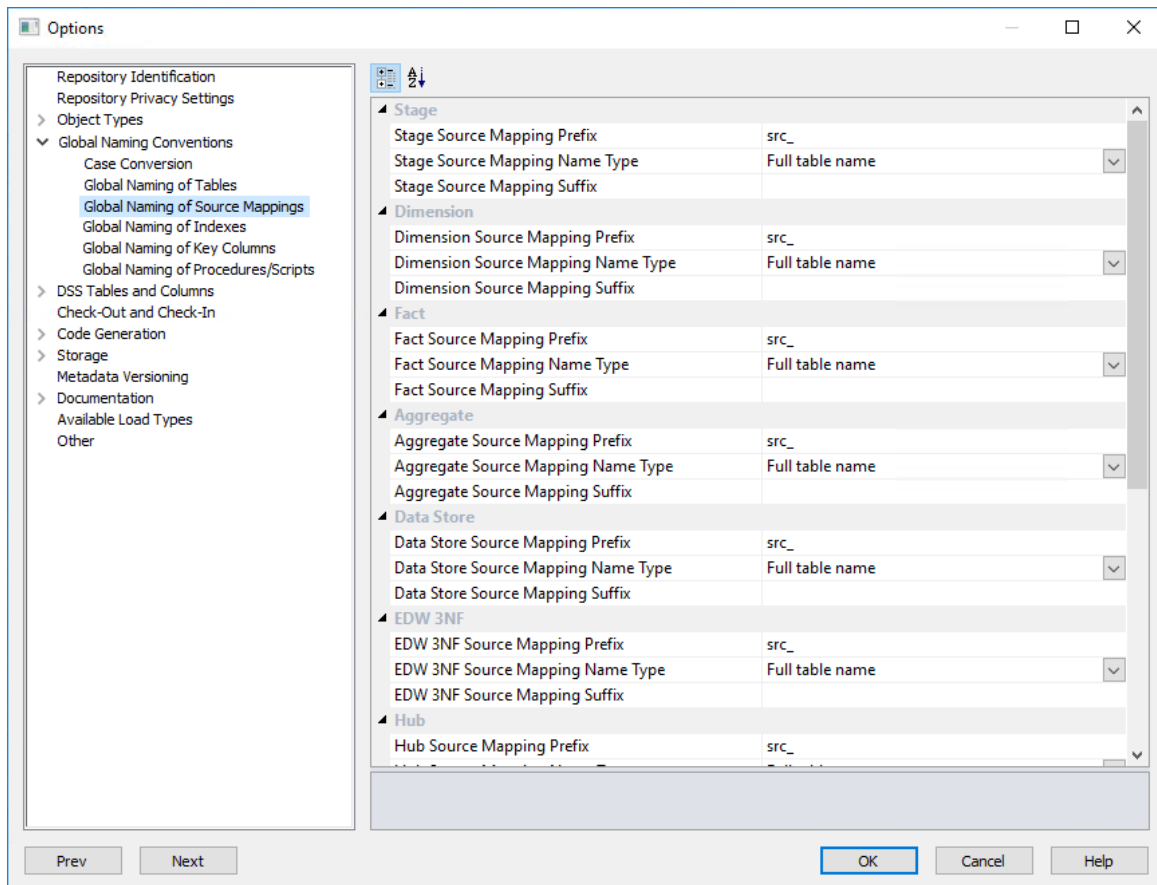
Refer to the section *Maintaining Source Column Mappings* (on page 1112) for details.



GLOBAL NAMING OF SOURCE MAPPINGS

The **Global Naming Conventions** option includes the **Global Naming of Source Mappings** setting that enables you to define naming standards for the source mapping objects in RED.

Refer to the section *Global Naming Conventions > Global Naming of Source Mappings* (see "5.1.4.3 Global Naming of Source Mappings" on page 94) for details.



INDEPENDENT EXECUTION OF UPDATE PROCEDURES

The update procedures associated with the target table and the source mapping objects can be executed individually, separate from each other. Similar to other table objects in RED, the update procedures can be executed on demand or via the **Scheduler** (on page 685). This provides the flexibility of scheduling the update procedures for each individual source mapping object as required, e.g. based on the update frequency of the tables where the source columns are obtained.

The update procedures for source mapping objects can only be generated using a template.

Refer to the section **Generating Update Procedures for Source Mapping Objects** (on page 1115) for details.

The screenshot displays the 'Scheduler' window in the Red Data Warehouse environment. It is divided into two main sections: 'Jobs' and 'Tasks'.

Jobs Section:

Job	Status	Sequen...	Start	Finish	Elapsed	OK	Info	Detail	Warn	Error	Who
Daily Update	On Hold	...	2018-02-15 11:05:21								D
Daily Update	Completed	...	2018-02-15 11:05:21	2018-02-1...	00:00	11	26				D
Job_Customer_Code_Merge	Waiting	...	2018-02-15 12:09:29								D
Job_Customer_Code_Merge	Completed	...	2018-02-15 12:09:27	2018-02-1...	00:00	4					D

Tasks Section:

Task	Action	Status	Seq	Start	Finish	Ela...	Info	Detail	W...	Result
stage_customer_code_merge	Truncate	Completed	22	2018-02-15 12:09:27	2018-02...	00:00				Table truncated stage_customer_code_merge
src_stage_customer_code_merge_1	Update	Completed	22	2018-02-15 12:09:27	2018-02...	00:00				stage_customer_code_merge updated. 6 new re...
src_stage_customer_code_merge_2	Update	Completed	22	2018-02-15 12:09:27	2018-02...	00:00				stage_customer_code_merge updated. 6 new re...

The interface also shows a status bar at the bottom with the text: 'Database time at last action was 12:10:56 | Development (Multi_Source) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS'.

TABLE COLUMN PROPERTIES

The **Source Details** table column properties is used to specify where to obtain the source data and provides options for processing the source data (transformation and join settings). You can also define and manage column source mappings from this screen.

Source Mapping Column src_h_customer_ms_smt_1.city

Table Name	src_h_customer_ms_smt_1
Column Name	city
Column Title	city
Column Description	
Physical Definition	
Column Order	40
Data Type	varchar(30)
Null Values Allowed	<input checked="" type="checkbox"/>
Default Value	
Character Set	
Format	
Character Comparison/Sorting	
Compress	<input type="checkbox"/>
Meta Definition	
Numeric	<input type="checkbox"/>
Additive	<input type="checkbox"/>
Attribute	<input checked="" type="checkbox"/>
Business Key	<input type="checkbox"/>
Key Type	None
Source Details	
Source Table	load_customer_td
Source Column	city
Source Data Type	varchar(30)
Transformation	
Join	False

Source Details
Properties that specify where and how to obtain source data.

<- Update Update -> **OK** Cancel Help

Refer to the relevant Table Column Properties sections in this user guide for details.

ADDING SOURCE MAPPING TO OBJECTS

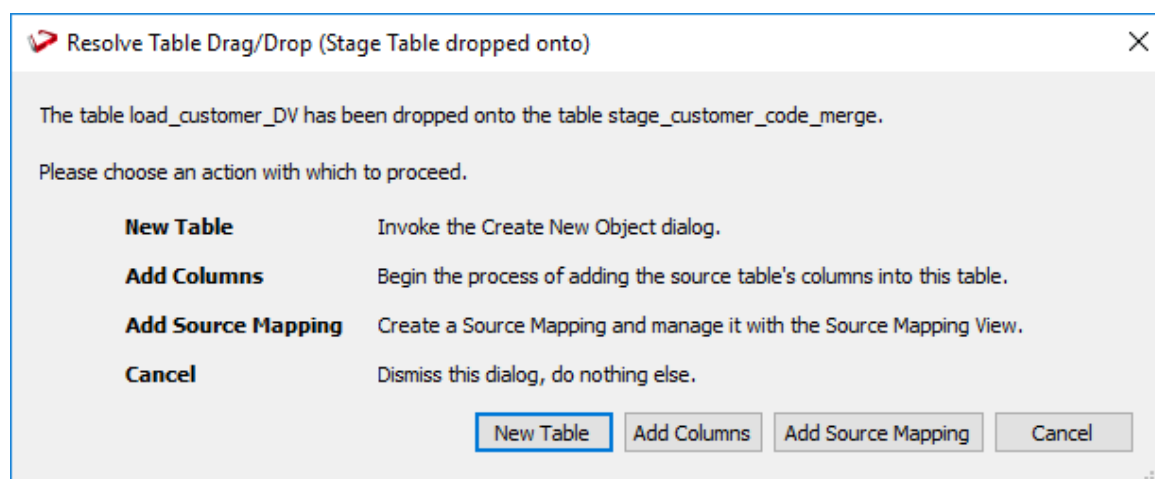
When adding source columns to a table object in RED, the target table's current source columns are used to create the first source mapping object and then proceeds to create the second source mapping object which contains the additional source columns that need to be mapped to the target table. The columns from the source table that exists in the target table are automatically mapped.

The process of adding Source Mappings is initiated when a source table is dragged and dropped to an existing table object in RED.

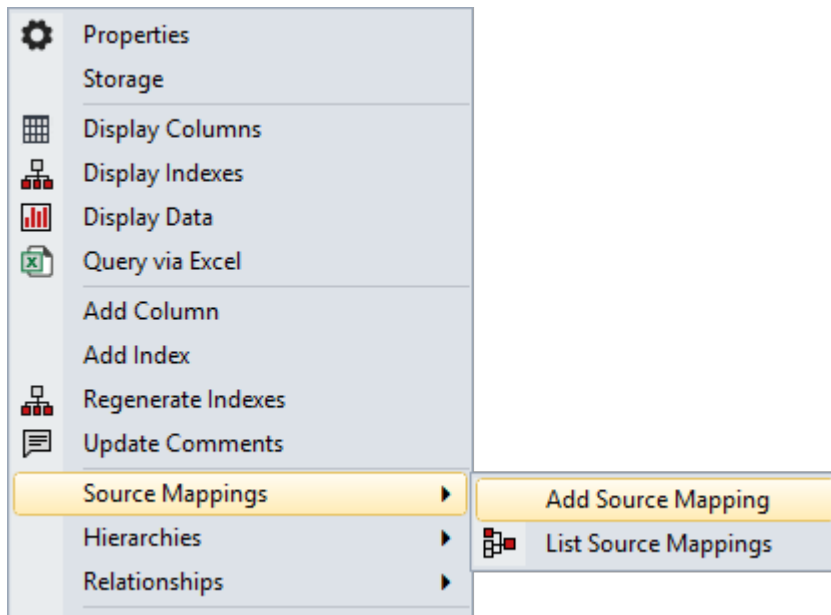
The user is prompted with the followings options:

- **New Table** - Opens a new dialog that enables you to create the dropped table as a new object in the metadata repository.
- **Add Columns** - Inserts the columns of the dropped table into the target table.
- **Add Source Mapping** - Inserts a new source mapping into the target table.

Note: If you are creating the first source mapping, the target table is first converted to a source mappable object, and then a new source mapping is inserted. If it is not the first source mapping, a new source mapping is just inserted into the target table.



The alternative method is to right click the target table (Parent) in the **Objects** list pane and then select the **Add Source Mapping** option from the context menu:



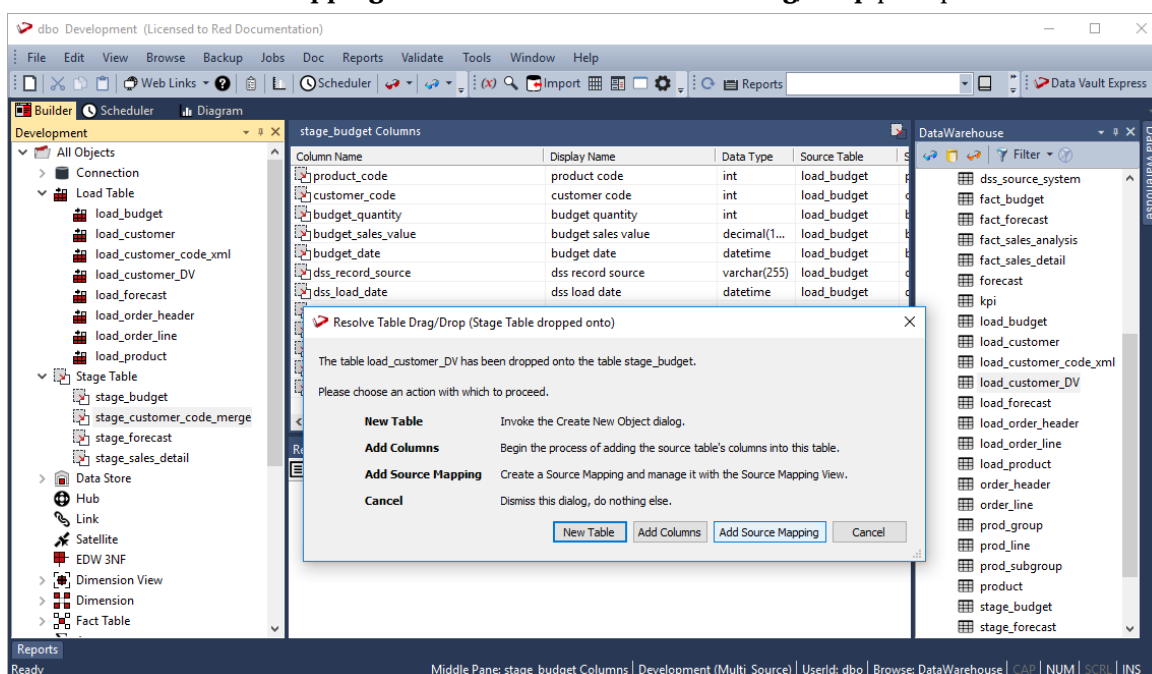
DRAG AND DROP

The common approach to create source mappings is to select the source table that contains the columns that you want to add and then drag this table into the target table.

This process creates two source mapping objects, the first one contains the original source mapping of the target table and second one is for the additional columns from the source table selected in the **Browser** pane.

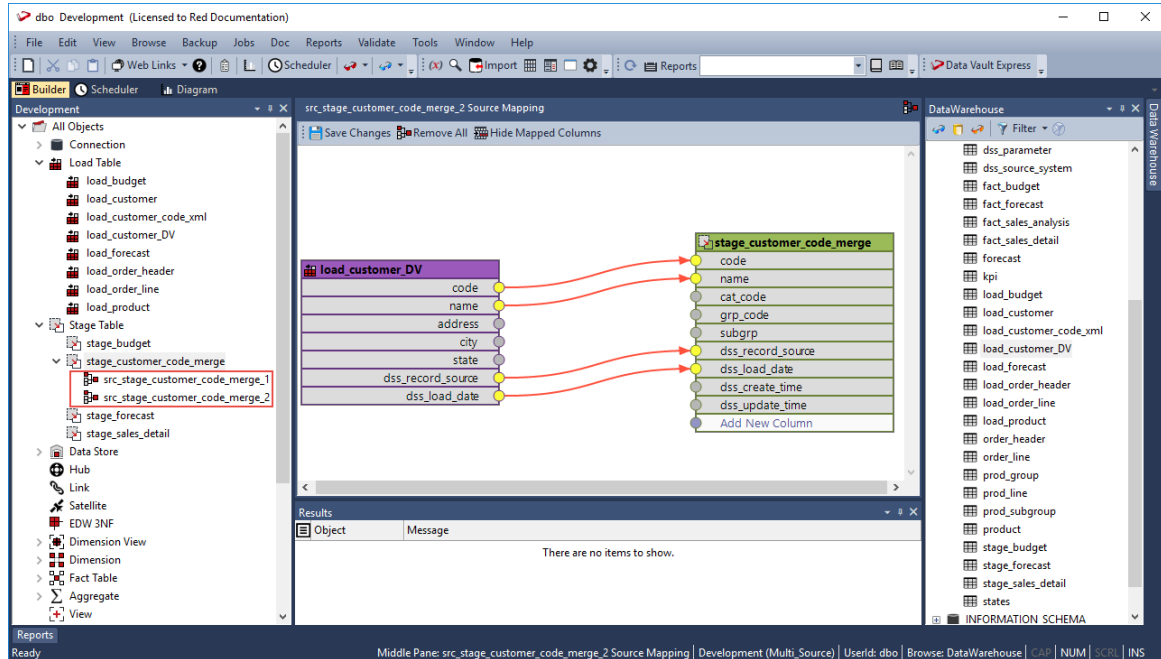
You can drag additional source tables to create additional source mappings objects until you have all the source data required for your target table.

- 1 Click the target table from the **Objects** list in the left pane.
The middle pane displays the properties of the selected target table, the pane is identified as a target for new Source Mapping objects.
- 2 Browse to the Data Warehouse via the **Browse > Source Table** menu option.
- 3 Click the source table that contains the columns you want to include as sources from the **Browser** pane, and drag it to the middle pane.
- 4 Click the **Add Source Mapping** button in the **Resolve Table Drag/Drop** prompt.

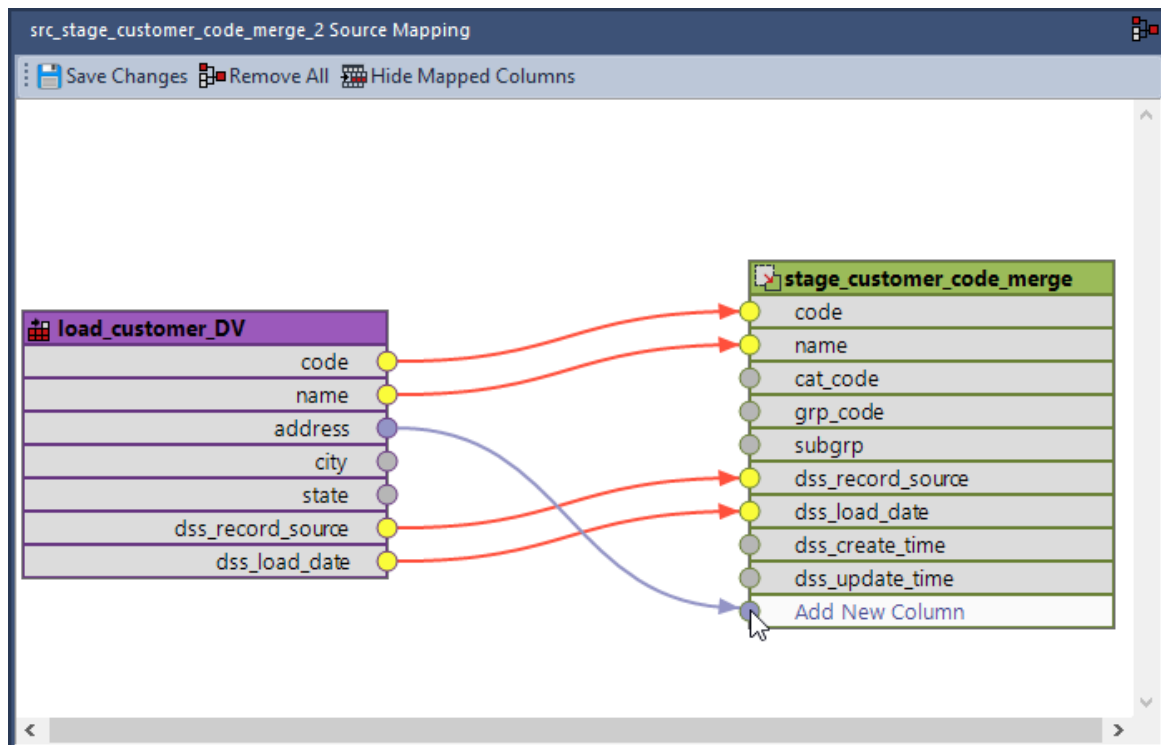


- 5 The new Source Mapping objects are created and displayed under the target table in the **Objects** list pane.

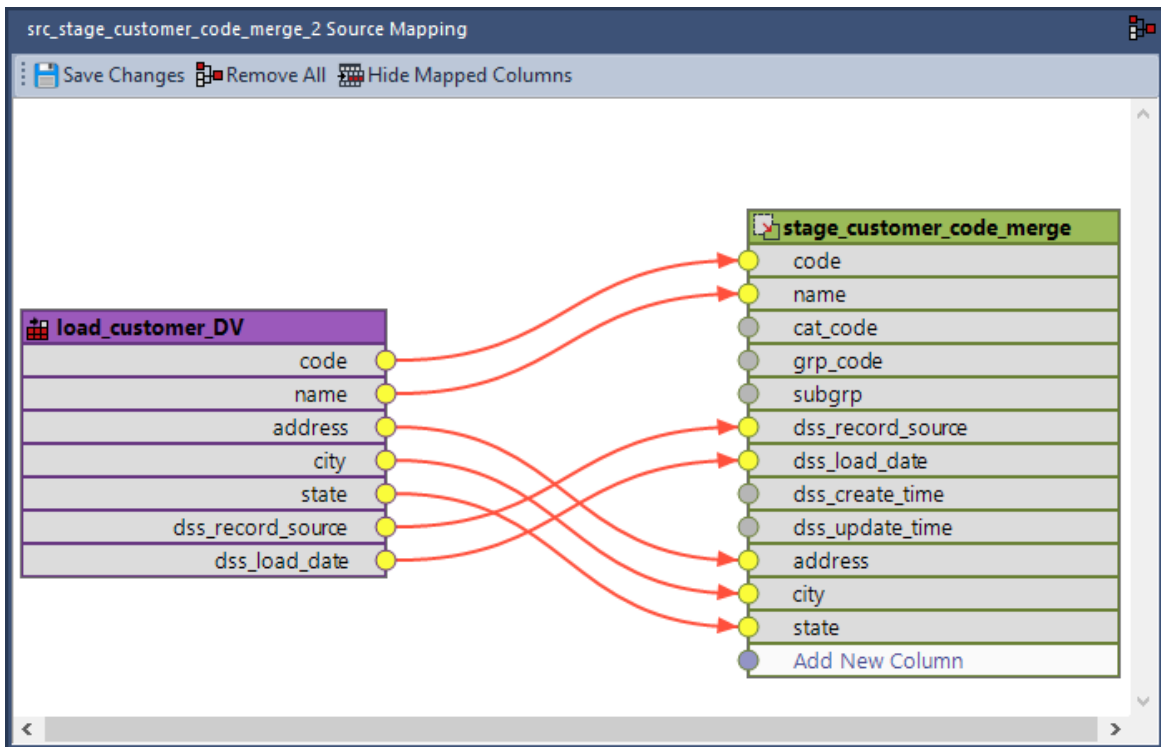
The middle pane displays the Source Mapping tool which provides a graphical representation of the mapped columns between the second source table and the target table. The columns from the source table that exists in the target table are automatically mapped.



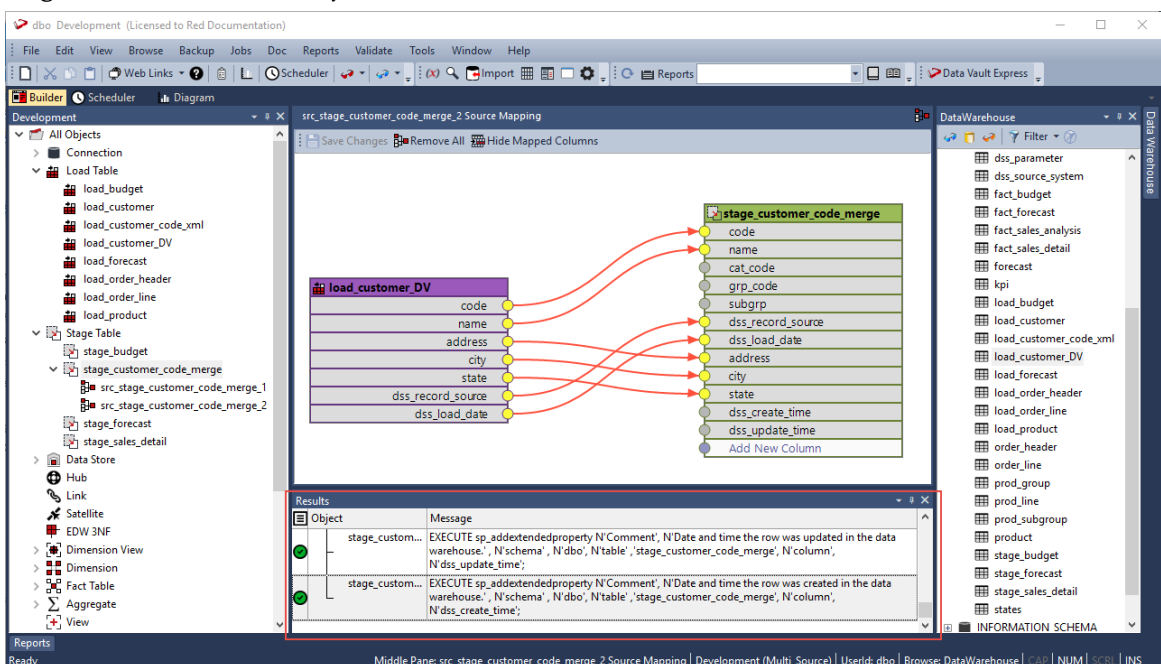
- To map the additional source columns that does not exist in the target table, click the connection point of the source column you want to map and drag the line that appears to the connection point of the **Add New Column** row in the target table.



- Repeat the same steps to map all the other source columns to the target table.



- Click **Save Changes** to save the defined source mappings.
- Right-click the target table from the **Objects** list in the left pane and select **Create(ReCreate)** from the context menu to recreate the table. The **Results** pane displays confirmation that the target table was successfully recreated.

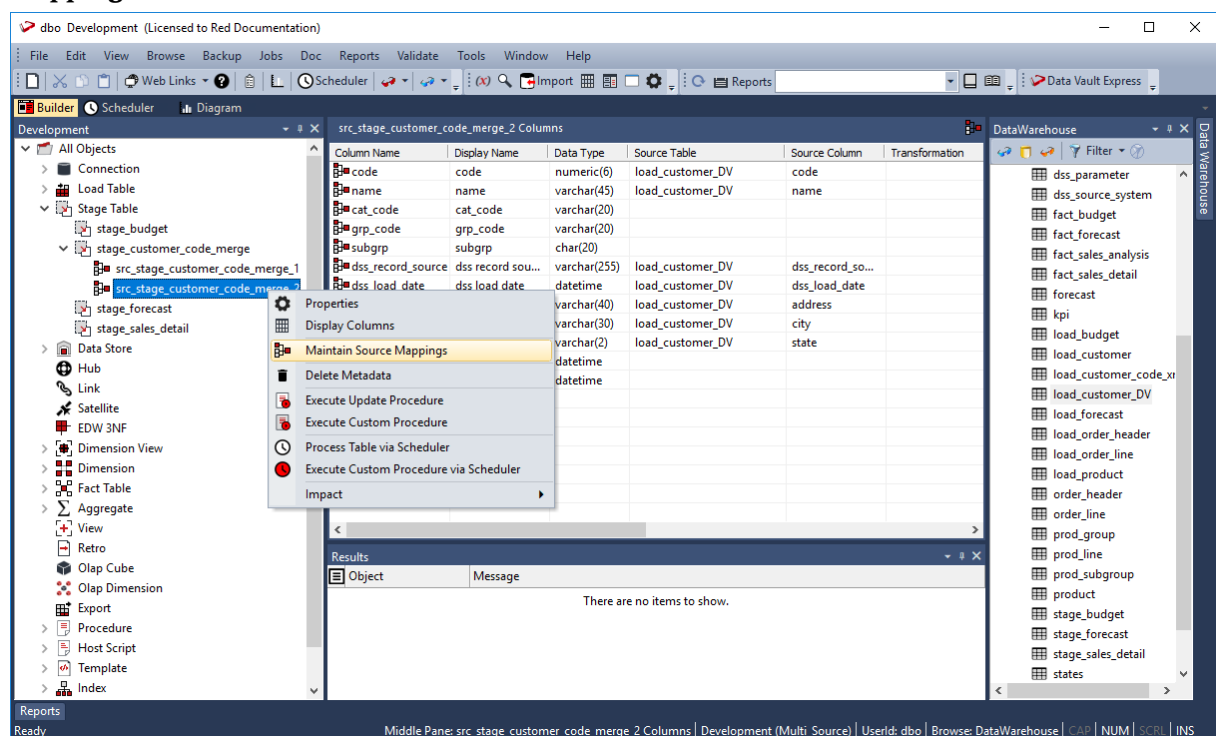


Once the target table is defined and created, you need to generate the update procedure for each source mapping object to populate the target table. The update procedure is created using a template—refer to *Generating Update Procedures for Source Mapping Objects* (on page 1115) for details.

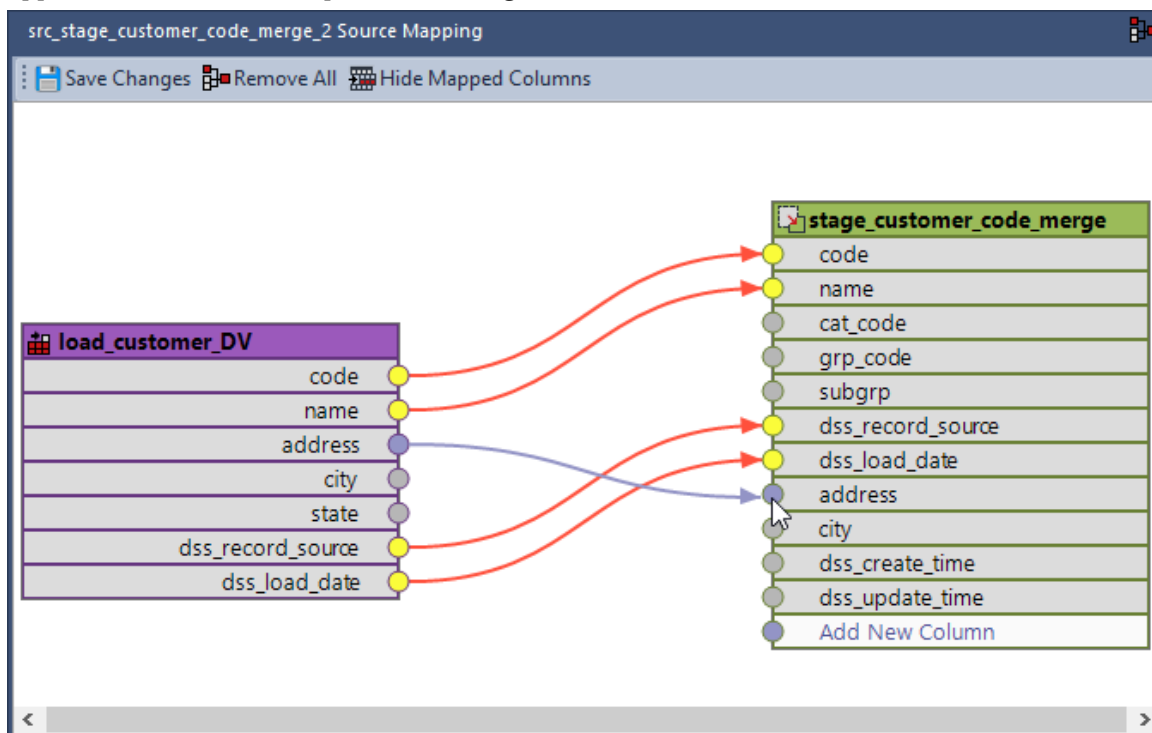
MAINTAINING SOURCE COLUMN MAPPINGS

The following describe the process for managing source mappings between the source table(s) and the target table.

- 1 Right click the Source Mapping object from the right-pane and select **Maintain Source Mappings** from the context menu.



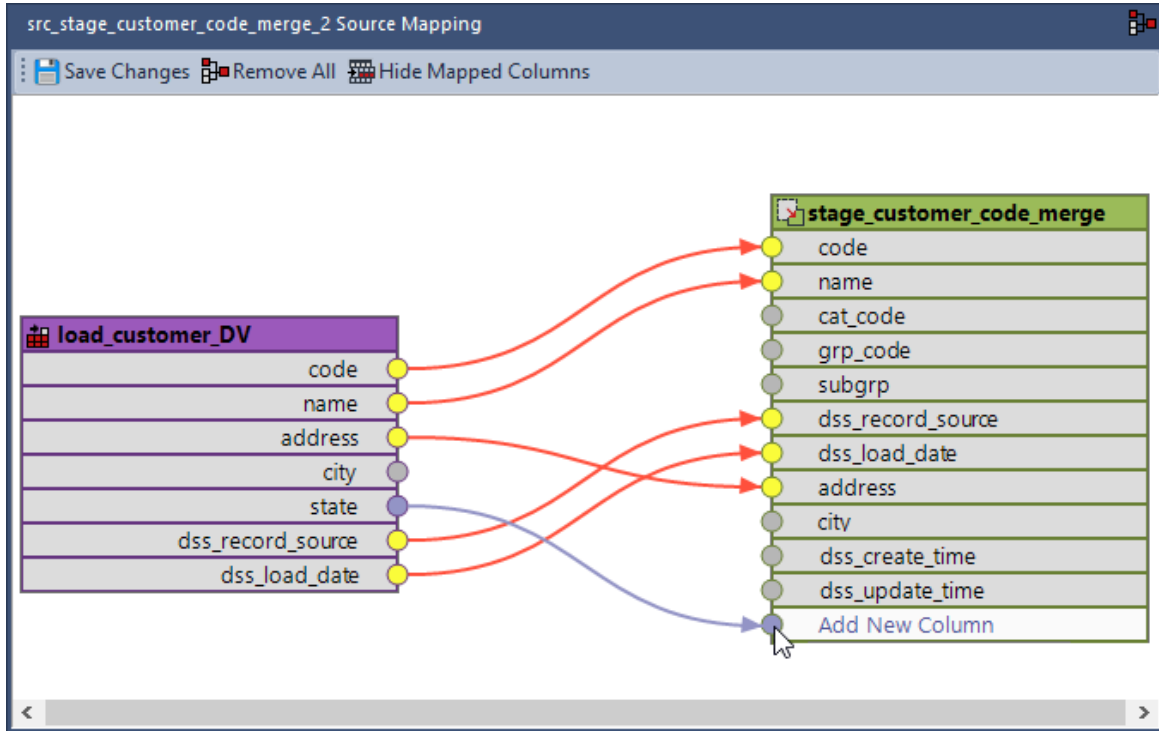
- Click the connection point of the source column you want to map and drag the line that appears to the connection point of the target column.



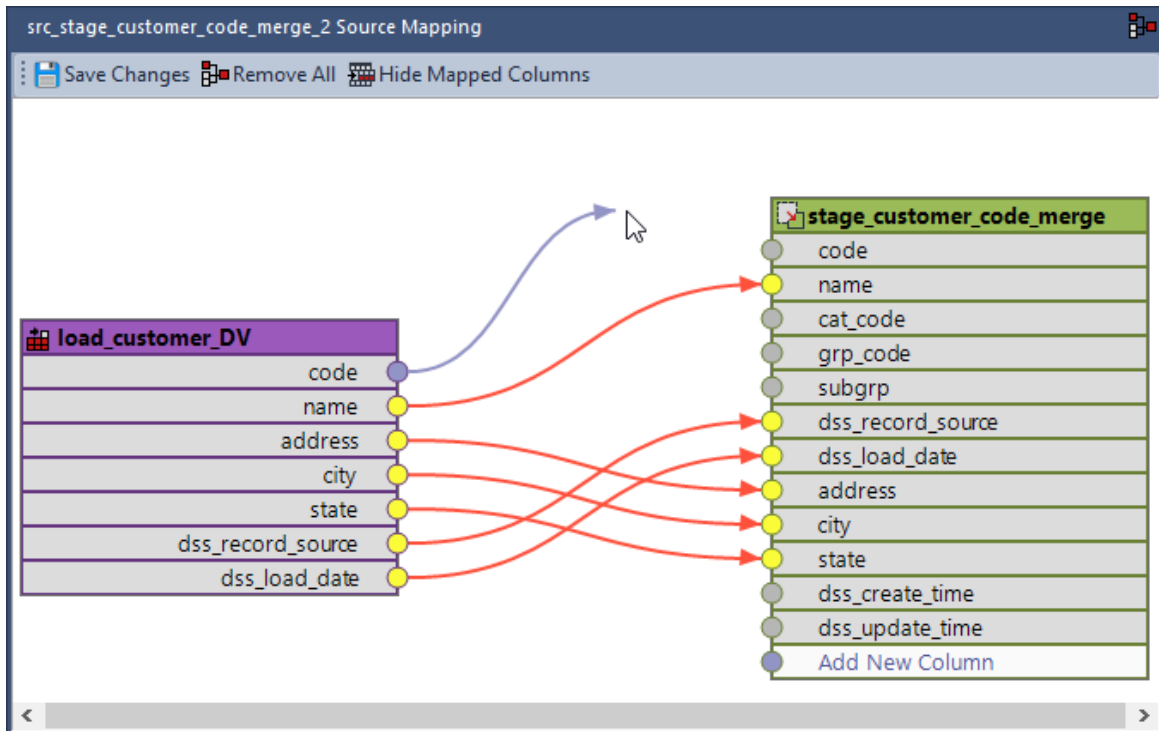
TIPS:

- The target table can be clicked and dragged vertically (up/down) within the diagram view. This is useful when there are more columns in the target table that can fit in the visible area of the diagram view. You can also use the mouse wheel for scrolling up/down the diagram view.
 - You can undock any of the panes (**Object, Browser** and **Results** pane) in the **Builder window** (on page 61), to gain a bigger work area for the Source Mapping tool.
-

- To map a source column that does not exist in target table, click the connection point of the source column you want to map and drag the line that appears to the connection point of the **Add New Column** row in the target table.



- To remove an existing source mapping, click the connection point on the target table and drag away the line that appears to disconnect.



Notes:

- Clicking **Remove All** removes all the current mappings between the source table and the target table.
 - Clicking **Hide Mapped Columns** hides all the current mappings and only the unmapped columns are displayed. This is useful when there are more columns in the target table that can fit in the visible area of the diagram view and you only want to work on the columns that are not mapped.
-

GENERATING UPDATE PROCEDURES FOR SOURCE MAPPING OBJECTS

After successfully defining and creating the Source Mapping objects, you can generate the update procedure for each object via a template to populate the target table.

Generating update and custom procedures for Source Mapping objects is completed using the same workflow as all objects in RED that support template based procedure/script generation.

Executing a procedure associated with the Source Mapping object can be performed both interactively from the object's right-click context menu or via the Scheduler.

Notes:

- Please ensure that you have created your own Source Mapping object templates, before performing the steps below.
- You can also have an update or custom procedure defined on the target table, if a single consolidated procedure to process the table is required.

- 1 Click the **Rebuild** button beside the **Update Procedure** drop-down in the **Properties** screen of the Source Mapping object, to launch the procedure generation Wizard.

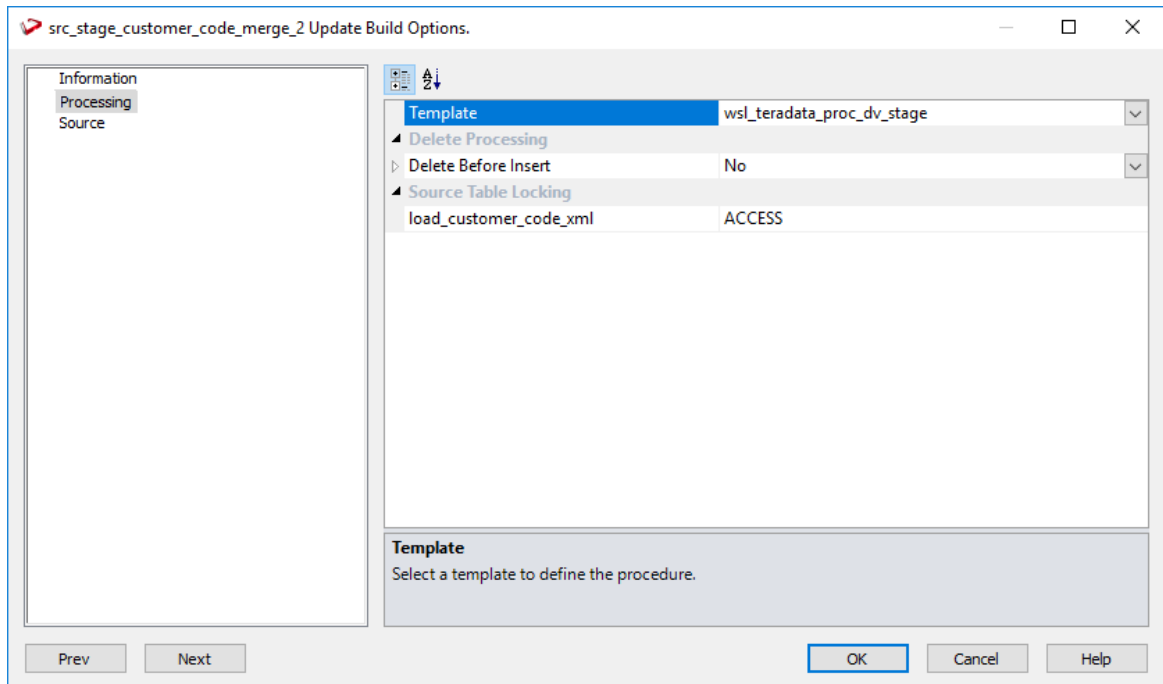
The screenshot shows the 'Source Mapping src_stage_customer_code_merge_2' dialog box. On the left is a sidebar with tabs: Properties (selected), Purpose, Concept, Grain, Examples, Usage, and Notes. The main area contains the following fields and controls:

- Table Name:** src_stage_customer_code_merge_2
- Unique Short Name:** (maximum 22 characters) src_stage_customer_c_2
- Description:** A large empty text area.
- Update Procedure:** (None) with a dropdown arrow. To its right are 'Rebuild' and 'Regenerate' buttons.
- Template:** wsl_sqlserver_proc_dv_stage
- Custom Procedure:** (None) with a dropdown arrow.
- Timestamps:** A section with three sub-fields: 'Metadata Structure', 'Database', and 'Database Altered', each with an empty text box.

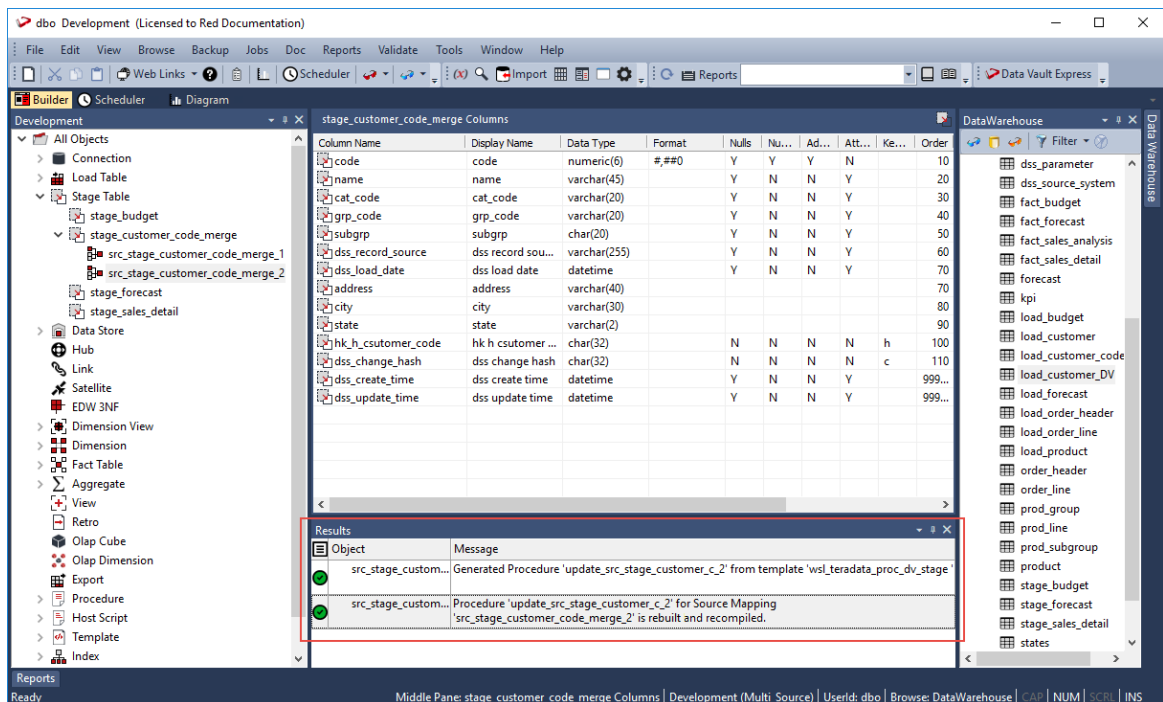
At the bottom right of the dialog are 'OK', 'Cancel', and 'Help' buttons.

Note: RED displays the name of the previously used update procedure template below the **Update Procedure** drop-down field by default.

- On the **Processing** tab of the **Update Build Options** screen, select the template to use from the **Template** drop-down or use the previous update procedure template. Configure the other options available as required.



- Click **OK** to proceed with the procedure generation. The **Results** pane displays confirmation that the procedure was generated.



- 4 Right-click the source table in the left pane, and select **Code>View update** from the context menu to view the contents of the update procedure generated.

```

-- DBMS Name      : Teradata
-- Procedure Name  : update_src_stage_customer_c_2
-- Template       : wsl_teradata_proc_dv_stage
-- Template Version: 8.0.2.0
-- Description    : Update the Stage table stage_customer_DV
-- Generated by   : WhereScape RED Version 8.1.1.0 (build 180219-214048 pre-release)
-- Generated for  : RED Documentation
-- Generated on   : Wednesday, February 21, 2018 at 17:26:02
-- Author        : WhereScape Documentation

-- Notes / History
--
CREATE PROCEDURE [METABASE].update_src_stage_customer_c_2
(
  IN p_sequence      INTEGER
  , IN p_job_name    VARCHAR(256)
  , IN p_task_name   VARCHAR(256)
  , IN p_job_id      INTEGER
  , IN p_task_id     INTEGER
  , OUT p_return_msg VARCHAR(256)
  , OUT p_status     INTEGER
)
BEGIN
  -----
  -- Control variables used in most programs
  -----
  DECLARE v_msgtext      VARCHAR(255); -- Text for audit_trail
  DECLARE v_sql         VARCHAR(255); -- Text for SQL statements
  DECLARE v_step        INTEGER;      -- return code
  DECLARE v_update_count INTEGER;     -- no of records updated
  DECLARE v_insert_count INTEGER;     -- no of records inserted
  DECLARE v_count       INTEGER;     -- General Counter
  DECLARE v_sql_code    INTEGER;     -- SQL Error Code for Audit Trail
  DECLARE v_sql_error   VARCHAR(255); -- SQL Error Code for Audit Trail as varchar
  
```

- 5 Right-click the table in the left pane and select **Execute Update Procedure** from the context menu to run the procedure. The **Results** pane displays the number of records created.

Object	Message
src_stage_custom...	Generated Procedure 'update_src_stage_customer_c_2' from template 'wsl_teradata_proc_dv_stage'
src_stage_custom...	Procedure 'update_src_stage_customer_c_2' for Source Mapping 'src_stage_customer_code_merge_2' is rebuilt and recompiled.
src_stage_custom...	Procedure Completed
src_stage_cus...	1 stage_customer_code_merge updated: 6 new records.

Repeat the same steps described above to create update procedures for the other source mapping objects, defined for the target table.

EXECUTING UPDATE PROCEDURES VIA SCHEDULER

The update procedures associated with the target table and the source mapping objects can be executed individually, separate from each other. Similar to other table objects in RED, the update procedures can be executed via the Scheduler.

- 1 Click the **Scheduler** tab to open the **Scheduler** window.
- 2 Click the **New Job** button from the **Scheduler** window toolbar to create a new job
- 3 Complete the fields on the **Job Definition** window. Refer to **Creating a Job** (on page 701) for the descriptions of the fields.

Job Definition
✕

Job Name:

Description:

Frequency:

Start Date:

Start Time:

Maximum Threads:

Scheduler:

Dependent On:

Logs Retained: This field lets you set the number of logs that are retained for this job before an automatic delete and archive occurs. 0 = keep all logs (default action)

The following two fields are optional. They are executed after the job completes and therefore need to reflect the scheduler environment. (i.e. Unix or Windows). The special variables \$JOB_KEY\$, \$JOB_SEQ\$ and \$JOB_NAME\$ can be used to return the associated values.
The Success command will be executed if a successful completion, the failure command will be executed if a job fails to complete:

Success Command:

Failure Command:

Execute Failure Command in event of dependency failure

Custom Settings

Interval Between Jobs: (minutes)

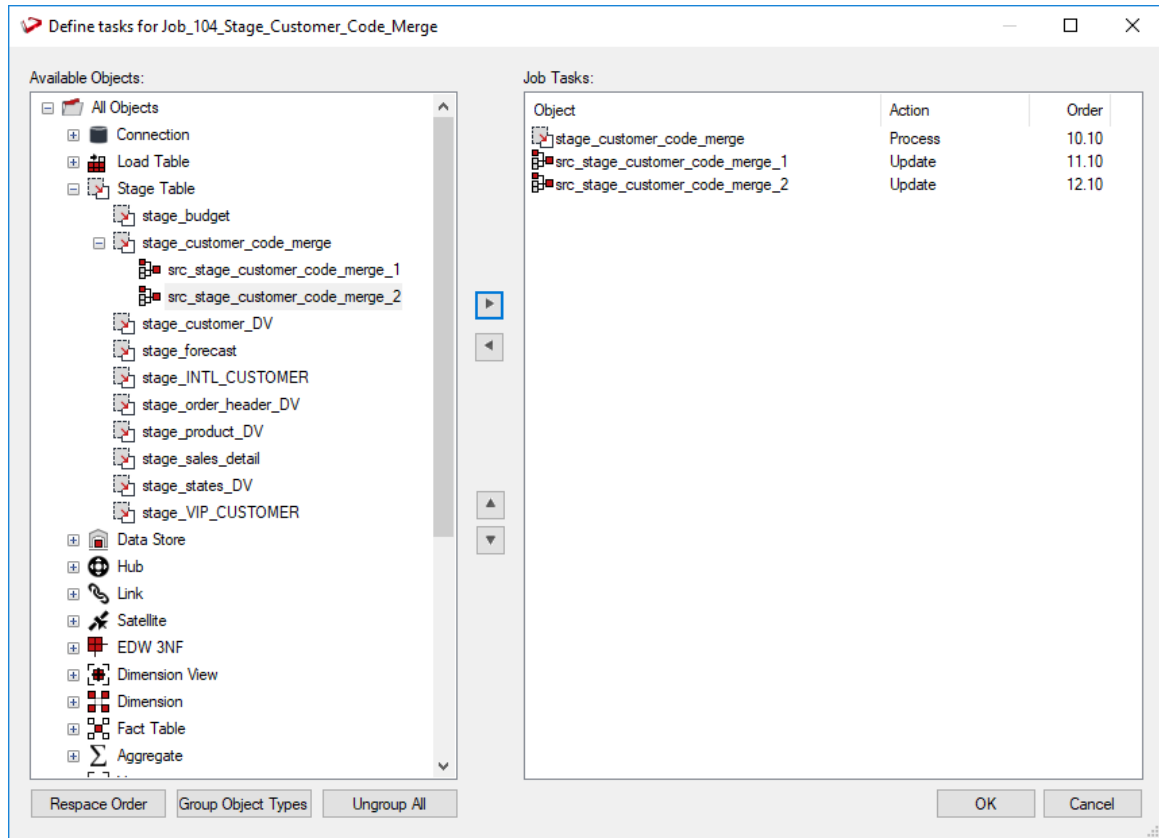
Start At or After HHMM (e.g. 0800):

Do Not Start After HHMM (e.g. 1700):

Active Days:

Mon Tue Wed Thu Fri Sat Sun

- 4 On the **Define tasks** window, select the target object and the its associated source mapping objects from the **Available Objects** pane and add them to **Job Tasks** pane to define the tasks to be run for the job.



Note: The update procedures associated with the target table and the source mapping objects can be executed individually, separate from each other.

5 Click **OK** to save and exit to the **Scheduler** window.

The screenshot displays the SQL Server Enterprise Manager Scheduler window. The window title is "dbo Development (Licensed to Red Documentation)". The interface includes a menu bar (File, View, Auto, Monitor, Logs, Jobs, Scheduler, Tools, Window, Help) and a toolbar with icons for Job Name Filter, All Jobs, Scheduled, Run/Fail, My Jobs, Today, and other functions. The main area is divided into two panes: "Jobs" and "Tasks".

Jobs Table:

Job	Status	Sequence	Start	Finish	Elapsed	OK	Info	Detail	Warn	Error	Who
Job_104_Stage_Customer_Code_Merge	On Hold	1...	2018-02-15 16:00:00								RD
Daily Update_1	On Hold	4	2018-02-15 11:42:44								RD
Daily Update_2	On Hold	7...	2018-02-15 11:42:43								RD
Job_75	On Hold	7...	2018-02-15 11:42:42								RD
Daily Update	On Hold	8...	2006-01-05 14:21:38								RD

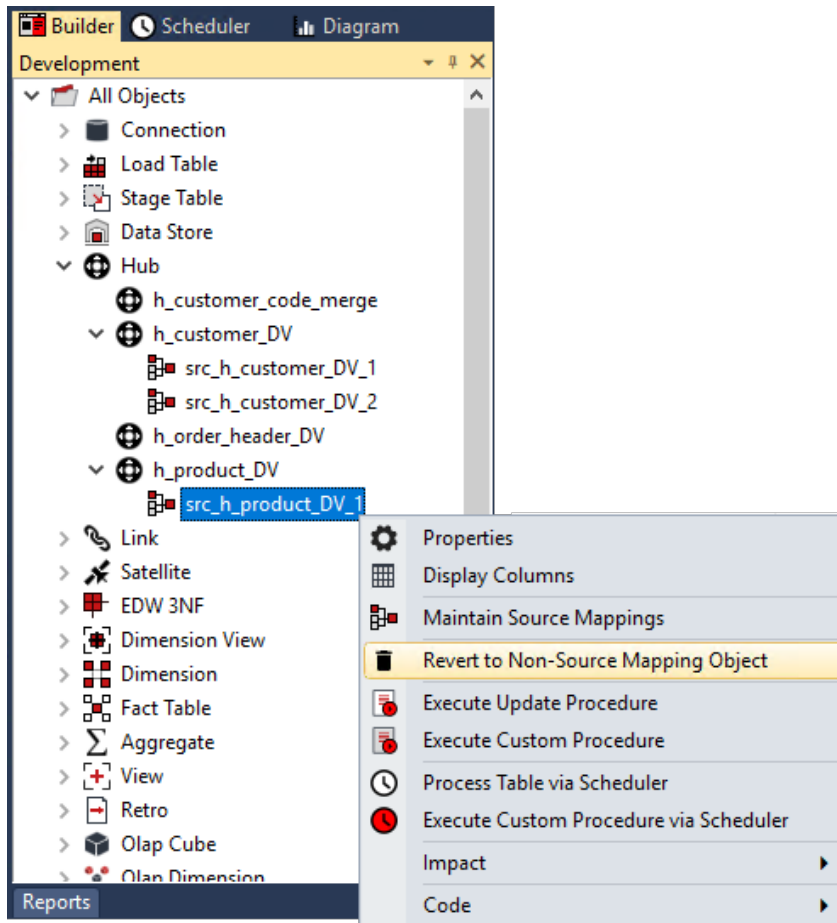
Tasks Table:

Task	Action	Status	Seq	Start	Finish	Elapsed	Info	De...	Warn	Result
stage_customer_code_merge	Truncate		104							
src_stage_customer_code_merge_1	Update		104							
src_stage_customer_code_merge_2	Update		104							

The bottom status bar shows "Ready", "Database time at last action was 15:35:29", "Development (Multi_Source_DW)", "Userid: dbo", "Browse: DataWarehouse", "CAP", "NUM", "SCRL", "INS".

REVERTING TO NON SOURCE MAPPING OBJECT

You can revert the target table to a non-source mapping object if you delete its associated source mapping objects. If only a single source mapping object is left, right clicking the remaining source mapping object provides the option **Revert to Non-Source Mapping Object**.



Note: Deleting the target table table, deletes all its source mapping objects, including its associated meta data.

