



QAD Enterprise Applications

User Guide

QAD Data Warehouse Designer

This document contains proprietary information that is protected by copyright and other intellectual property laws. No part of this document may be reproduced, translated, or modified without the prior written consent of QAD Inc. The information contained in this document is subject to change without notice.

QAD Inc. provides this material as is and makes no warranty of any kind, expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. QAD Inc. shall not be liable for errors contained herein or for incidental or consequential damages (including lost profits) in connection with the furnishing, performance, or use of this material whether based on warranty, contract, or other legal theory.

QAD and MFG/PRO are registered trademarks of QAD Inc. The QAD logo is a trademark of QAD Inc.

Designations used by other companies to distinguish their products are often claimed as trademarks. In this document, the product names appear in initial capital or all capital letters. Contact the appropriate companies for more information regarding trademarks and registration.

Copyright © 2015 by QAD Inc.

DataWarehouseDesigner_UG_v0674.pdf/b3s

QAD Inc.

100 Innovation Place
Santa Barbara, California 93108
Phone (805) 566-6000
<http://www.qad.com>

Contents

Overview	1
Overview of QAD Data Warehouse Designer	2
Documentation Roadmap	9
Design Introduction	11
Objects and Windows	13
Object Types	14
Organizing Objects	17
Adding Objects to Projects	21
Removing Objects from Projects	23
Using Project Object Maintenance	24
Adding Projects to Groups	25
Removing Projects from Groups	26
Moving Projects within Groups	27
Windows and Panes	28
Main Builder Window	29
Scheduler window	32
Diagram window	34
Procedure Editor window	35
Database Error Messages	37
Export Middle Pane Output	38
Tutorials	41
Default Settings	43
Settings - Options	44
Settings - Repository Identification	45
Settings - Object Types	46
Object Type Availability	47
Object Type Names	48
Object Type Ordering	49
Object End User Setting	50
Object Color	51
Settings - Global Naming Conventions	52
Case Conversion	53
Global Naming of Tables	54
Global Naming of Indexes	55
Global Naming of Key Columns	57
Global Naming of Procedures	58
Settings - DSS Tables and Columns	59
DSS Tables	62
DSS Columns	63
Settings - Check-Out and Check-In	64
Settings - Code Generation	65
General	66
Default Update Procedure Options	67
Generated Code Statistics	69

Settings - Storage	71
Table Storage	72
Index Type	73
Settings - Versioning	74
Settings - Scheduler	75
Scheduler General	76
Scheduler Statistics	77
Settings - Documentation	78
Settings - Other	79
Settings - User Preferences	80
Settings - Common	81
Look and Feel	82
Local Naming Conventions	90
Local Paths	95
Outputs	97
Other	99
Settings - Current Repository	100
Look and Feel	101
Settings - Language Options	102

Parameters	105

Connections	107

Browsing a Connection	109
Changing a Connection's Properties	112
Creating a Database Link	113
SQL Server Connections	114
Database	115
Data Warehouse	116
Database Link Load - Current	119
Database Link Load - Other DB	122
Database Link Load - Linked Server	125
Integration Services Load	136
Windows	142
ODBC	144
Native ODBC Load	145
ODBC Load	149
Integration Services Load	153
Analysis Server 2000	157
Microsoft Analysis Server 2005+	158
Oracle Connections	159
Database	160
Data Warehouse	161
Database Link Load - Current	164
Database Link Load - Linked Server	167
UNIX	170
Windows	174
ODBC	176
Native ODBC Load	177
ODBC Load - Single Row Inserts	181
ODBC Load - Direct Path OCI	185
Analysis Server 2000	189
Microsoft Analysis Server 2005+	190
DB2 Connections	191
Database	192
Data Warehouse	193
UNIX	196

Windows	200
ODBC	203
Native ODBC Load	204
ODBC Load - Single Row Inserts	208
ODBC Load - Direct Path OCI	212
Analysis Server 2000	216
Microsoft Analysis Server 2005+	217

Table Properties **219**

Table Storage Properties	220
Table Storage Screen - SQL Server	222
Table Storage Screen - Oracle	223
Table Storage Screen - DB2	225
Table Notes	226
Table Notes Screen	227
Table Documentation Fields	228
Table Documentation Fields Screen	229
Table Statistics in DB2	230
Table Statistics Screen	231

Loading Data **233**

Choosing the Best Load Method	234
Load Drag and Drop	235
Data Type Mappings	236
Database Link Load	237
Database Link Load - Properties	238
Load Table Properties Screen	240
Database Link Load - Source Mapping	241
Load Table Source Mapping Screen	243
Remote View Extract	246
ODBC Based Load	248
SSIS Loader	249
Loading data into Data Warehouse Designer load tables using SSIS	250
Troubleshooting	262
Native ODBC Based Load	264
Native ODBC Based Load - Source Mapping	265
Load Table Source Mapping Screen (Native)	267
File Actions	270
File Action Example	271
Flat File Load	273
File Load - Source Mapping	277
Load Table Source Mapping Screen	280
Host Script Load	283
File Wild Card Load	284
XML File load	285
External Load	289
Handling Missing Source Columns	290
Load Table Transformations	293
Post-Load Procedures	294
Changing Load Connection and Schema	295

Dimensions **297**

Dimensions Overview	298
Building a Dimension	300
Generating the Dimension Update Procedure	308

Dimension Artificial Keys	323
Dimension Get Key Function	324
Dimension Initial Build	326
Dimension Column Properties	327
Dimension Column Transformations	335
Dimension Hierarchies	336
Adding a Dimension Hierarchy	337
Using a Maintained Hierarchy	339
Snowflake	340
Creating a Snowflake	341
Dimension Language Mapping	343

Staging 345

Building the Stage Table	346
Generating the Staging Update Procedure	349
Tuning the Staging Update Process	362
Stage Table Column Properties	363
Stage Table Column Transformations	367
Set Merge Procedure	368

Data Store Objects 377

Data Store Object Overview	378
Building a Data Store Object	380
Generating the Data Store Update Procedure	384
Generating a Procedure	384
Procedure Type and Options	384
Data Store Artificial Keys	397
Data Store Column Properties	399
Data Store Object Column Transformations	403

Normalized Tables 405

Normalized Tables Overview	406
Building a Normalized Table	408
Generating the Normalized Update Procedure	412
Generating a Procedure	412
Procedure Type and Options	412
Normalized Table Artificial Keys	426
Normalized Table Column Properties	428
Normalized Table Column Transformations	432

Fact Tables 433

Detail Fact Tables	434
Creating Detail Fact Tables	435
Generating the Detail Fact Update Procedure	436
Generating a Procedure	436
Define Fact Business Key Columns	436
Provide Cursor Mapping	437
Rollup or Combined Fact Tables	439
Creating Rollup Fact Tables	440
Generating the Rollup Fact Update Procedure	441
Generating a Procedure	441
Define Rollup date dimension and column	441
KPI Fact Tables	443
Creating the KPI Dimension	444

Creating KPI Fact Tables	445
Setup of KPI Fact Tables	446
KPI Setup	447
Generating the KPI Fact Update Procedure	453
Generating a Procedure	453
KPI parameter definition	453
Define Rollup date dimension and column	454
Select the KPI dimension and identifier	455
Snapshot Fact Tables	457
Creating Snapshot Fact Tables	458
Generating the Snapshot Fact Update Procedure	459
Generating a Procedure	459
Slowly Changing Fact Tables	460
Creating Slowly Changing Fact Tables	461
Partitioned Fact Tables	462
Creating a Partitioned Fact Table in SQL Server	463
Creating a Partitioned Fact Table in Oracle	468
Creating a Partitioned Fact Table in DB2	472
Generating the Partitioned Fact Update Procedure	476
Generating a Procedure	476
Define Fact Business Key Columns	476
Provide Cursor Mapping	477
Fact Table Column Properties	479
Fact Table Column Transformations	484
Fact Table Language Mapping	485
Aggregation	487
Creating an Aggregate Table	488
Change an Aggregates Default Look Back Days	491
Creating an Aggregate Summary Table	492
Aggregate Table Column Properties	494
Aggregate Table Column Transformation	499
Views	501
Creating a Dimension View	502
View Column Re-mapping	508
Non-Dimension Object Views	510
Creating a Custom View	514
Dimension View Language Mapping	516
View Aliases	517
Analysis Services OLAP Cubes	519
Overview	520
Defining the Data Source for the OLAP Cube	521
Defining an OLAP Cube	526
Inspecting and Modifying Advanced Cube Properties	530
Creating an OLAP Cube on the Analysis Services Server	531
OLAP Cube Objects	532
OLAP Cube Properties	533
OLAP Cube Measure Groups	538
OLAP Cube Measure Group Processing/Partitions	540
OLAP Cube Measure Group Partitions	544
OLAP Cube Measures	547
OLAP Cube Calculations	550
OLAP Cube Key Performance Indicators	553

OLAP Cube Actions	557
OLAP Cube Dimensions	561
OLAP Cube Measure Group Dimensions	564
OLAP Dimension Objects	567
OLAP Dimension Overview	568
OLAP Dimension Attributes	572
OLAP Dimension Attribute Relationships	576
OLAP Dimension Hierarchies	578
OLAP Dimension User Defined Hierarchy Levels	581
Changing OLAP Cubes	583
Retrofitting an OLAP Object	586

Exporting Data **591**

Building an Export Object	592
File Attributes	595
Export Column Properties	598
Host Script Exports	601

Transformations **603**

Column Transformations	604
Column Transformation Properties	605
Load Table Column Transformations	608
Database Link During Load Transformations	609
File During Load Transformations	612
After Load Transformations	613
Database Functions	614
Re-usable Transformations	615
Creating a New Re-usable Transformation	616
Specify the Name of the Transformation	617
Enter Re-usable Transformation Metadata	619
Define the Transformation Model	620
Completed Re-usable Transformation	622
Changing a Re-usable Transformation	623
Applying Changes to Dependant Transformations	624
Using Re-usable Transformations	626

Procedures and Scripts **627**

Procedure Generation	628
Procedure Editing	631
Procedure Loading and Saving	634
Procedure Comparisons	636
Procedure Compilation	637
Procedure Running	638
Procedure Syntax	639
Procedure Properties	640
Script Generation	642
Script Generation (Windows/SQL Server)	643
Script Generation (UNIX/Oracle)	647
Script Generation (Windows/Oracle)	653
Script Generation (Windows/DB2)	658
Script Editing	662
Script Testing	663
Script Syntax	664
Calling a Batch File from a Script	666
Scheduling Scripts	669

Templates	671
Scheduler	673
Scheduler Options	674
Scheduler States	677
Scheduling a Job	679
Working with Jobs	687
Creating a Job	689
Editing a Job	699
Editing Tasks in a Job	703
Editing Dependencies	711
Inserting a Copy of a Job	716
Deleting a Job	718
Deleting Job Logs	719
Starting a Job	721
Halting a Job	722
Aborting a Job	723
Restarting a Job	724
Creating an Application from a Job	729
Monitoring the Database and Jobs	732
Database Monitoring	733
Job Monitoring	735
Stand Alone Scheduler Maintenance	740
SQL to return Scheduler Status	743
Reset Columns in Job and Task View	744
Scheduling a Data Warehouse Designer Host Script from 3D	745
Indexes	753
Index Definition	754
Documentation and Diagrams	761
Creating Documentation	762
Batch Documentation Creation	765
Reading the Documentation	766
Diagrams	767
Types of Diagrams	768
Schema Diagram	770
Source Diagram	772
Joins Diagram	776
Links Diagram	777
Impact Diagram	778
Dependency Diagram	779
Working with Diagrams	782
Creating a Job from a Diagram	785
Creating an Application from a Diagram	787
Creating a Project from a Diagram	789
Query Tool Access	791
Reports	793
Dimension-Fact Matrix	794
OLAP Dimension-Cube Matrix	795
Dimension Views of a Specified Dimension	796
Column Reports	797
Columns without Comments	798

All Column Transformations	799
Re-Usable Column Transformations	800
Column Track-Back	801
Column Track-Forward	804
Table Reports	807
Tables without Comments	808
Load Tables by Connection	809
Export Objects by Connection	810
Records that failed a Dimension Join	811
External Source Table/files	813
Procedure Reports	814
Modified Procedures	815
Custom Procedures	816
Index Reports	817
Modified Indexes	818
Object Reports	819
Objects-Project Matrix	820
Modified Objects (excluding indexes)	821
Objects Checked-out	823
Loaded or Imported Objects	824
Job Reports	826
Object-Job Matrix	827
Jobs with an Object	828
Tasks of a Job	829
Operational Reports	831
Object Performance History	832
Job Performance History	834
Task Performance History	835
Fragmentation	837

Validate	839

Validate Meta-data	840
Validate Workflow Data	841
Validate Table Create Status	842
Validate Load Table Status	843
Validate Index Status	844
Validate Procedure Status	845
List Meta-data Tables not in the Database	846
List Database Tables not in the Meta-data	847
List Tables with no related Procedures or Scripts	849
List Procedures not related to a Table	851
Compare Meta-data Repository to another	852
Compare Meta-data Indexes to Database	855
List Duplicate Business Key Columns	856
Query Data Warehouse Objects	857

Promoting Between Environments	859

Applications	860
Application Creation	861
Application Loading	866
Importing Object Metadata	867
Importing Language Files	869
Data Warehouse Testing	870
Creating and Loading Applications from the Command Line	874

Backing Up and Restoring Metadata **877**

Backup using DB Routines	878
Restoring DB Backups	881
Unloading Metadata	885
Loading an Unload	888

Altering Metadata **893**

Validating Tables	894
Validating Source (Load) Tables	896
Validating Procedures	897
Altering Tables	898
Validating Indexes	900
Recompiling Procedures	901

Callable Routines **903**

Introduction to Callable Routines	904
Callable Routines API	904
Callable Routines per RDBMS	905
Callable Routines Names Qualifier	907
Callable Routines Common Input	907
Callable Routines Invocation	908
Alternative Invocation Methods	909
Ws_Api_Glossary	912
SQL Server	912
Teradata	913
Oracle	914
DB2	914
Ws_Connect_Replace	916
SQL Server	917
Teradata	918
Oracle	919
DB2	920
Ws_Job_Abort	921
SQL Server	921
Teradata	922
Oracle	922
DB2	923
Ws_Job_Clear_Archive	924
SQL Server	925
Teradata	926
Oracle	927
DB2	928
Ws_Job_Clear_Logs	930
SQL Server	931
Teradata	932
Oracle	933
DB2	934
Ws_Job_Clear_Logs_By_Date	935
SQL Server	936
Teradata	937
Oracle	938
DB2	939
Ws_Job_Create	940
SQL Server	941

Teradata	943
Oracle	945
DB2	946
Ws_Job_CreateWait	948
SQL Server	949
Teradata	951
Oracle	953
DB2	954
Ws_Job_Dependency	956
SQL Server	957
Teradata	958
Oracle	959
DB2	960
Ws_Job_Release	962
SQL Server	963
Teradata	964
Oracle	965
DB2	966
Ws_Job_Restart	967
SQL Server	968
Teradata	969
Oracle	970
DB2	971
Ws_Job_Schedule	972
SQL Server	973
Teradata	974
Oracle	975
DB2	976
Ws_Job_Status	977
SQL Server	979
Teradata	980
Oracle	982
DB2	983
Ws_Load_Change	985
SQL Server	986
Teradata	987
Oracle	988
DB2	989
Ws_Maintain_Indexes	990
SQL Server	991
Teradata	992
Oracle	993
DB2	994
Ws_Version_Clear	996
SQL Server	997
Teradata	998
Oracle	999
DB2	1000
WsParameterRead	1001
SQL Server	1001
Teradata	1002
Oracle	1003
DB2	1003
WsParameterReadF	1004
SQL Server	1004
WsParameterReadG	1006
SQL Server	1007
Teradata	1007
Oracle	1008

DB2	1008
WsParameterReadJ	1010
SQL Server	1010
Teradata	1011
Oracle	1012
DB2	1012
WsParameterWrite	1013
SQL Server	1014
Teradata	1014
Oracle	1015
DB2	1015
WsWrkAudit	1016
SQL Server	1017
Teradata	1018
Oracle	1019
DB2	1020
WsWrkAuditBulk	1022
SQL Server	1023
Teradata	1024
Oracle	1025
DB2	1026
WsWrkError	1028
SQL Server	1029
Teradata	1030
Oracle	1031
DB2	1032
WsWrkErrorBulk	1034
SQL Server	1035
Teradata	1036
Oracle	1037
DB2	1038
WsWrkTask	1040
SQL Server	1041
Teradata	1042
Oracle	1043
DB2	1044

Ws_admin_v Views	1045

Ws_admin_v_audit	1046
Ws_admin_v_dim_col	1047
Ws_admin_v_error	1049
Ws_admin_v_fact_col	1050
Ws_admin_v_fact_join	1053
Ws_admin_v_sched	1055
Ws_admin_v_task	1059

Retrofitting	1063

Migrating the Data Warehouse Database Platform	1064
Importing a Data Model	1072
Re-Targeting Source Tables	1077
Retro Copy Column Properties	1079
Retro Copy Column Properties Screen	1083
Retro Copy Column Transformations	1084

Integrating QAD Data Warehouse Designer into an Existing Warehouse **1085**

Rebuilding	1086
Integrating	1087
Integrating, Host Scripts	1088
Integrating, Selecting a Table Type	1089
Integrating, Questions	1090
Integrating, Procedures	1094
Integrating, Views	1095
Integrating, QAD Tables	1096

Join Information Maintenance **1097**

Upgrading Data Warehouse Designer **1099**

Login Checks **1101**

Data Type Mappings **1105**

Using Data Type Mapping Sets	1106
Maintaining Data Type Mapping Sets	1108
Creating a New Data Type Mapping Set	1110
Copying a Data Type Mapping Set	1114
Editing a Data Type Mapping Set	1118
Deleting a Data Type Mapping Set	1128
Loading Data Type Mapping Sets	1130
Exporting Data Type Mapping Sets	1132
Data Type Mapping Examples	1134

Column Context Menu **1141**

Properties	1142
Change Column(s)	1144
Add Column	1147
Duplicate Column	1149
Delete Column	1151
Re-space Order Number	1152
Impact	1154
Send Columns to Another Object	1156

Database Function Sets	1159
Using Database Function Sets	1160
Maintaining Database Function Sets	1163
Creating a New Database Function Set	1166
Copying a Database Function Set	1169
Editing a Database Function Set	1172
Deleting a Database Function Set	1184
Loading Database Functions	1186
Exporting Database Function Sets	1189
Index	i

Change Summary

The following table summarizes significant differences between this document and the previous version.

Date/Version	Description	Reference
March 2015 / 6.7.4	New version of DWD; changes throughout.	--
September 2012 / 6.6.2	New version of DWD; changes throughout.	--
March 2012 / 6.5.6	Numerous updates throughout	--
September 2011 / QAD BI 3.5	Rebranded for QAD BI 3.5	--

Chapter 1

Overview

In This Chapter

Overview of QAD Data Warehouse Designer 2

Documentation Roadmap 9

Overview of QAD Data Warehouse Designer

Traditionally data warehouses take too long to build and are too hard to change. QAD Data Warehouse Designer is an Integrated Development Environment to support the building and managing of data warehouses. It has the flexibility to enable you to build a variety of architectures including:

- ☐ enterprise data warehouses
- ☐ dimensional data warehouses
- ☐ data marts
- ☐ user facing views, aggregates and summaries

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

- ☐ As a data warehouse specific tool, QAD Data Warehouse Designer embodies a simple, pragmatic approach to building data warehouses. With QAD Data Warehouse Designer you specify what you want to achieve by dragging and dropping objects to create a meta view, and then let QAD Data Warehouse Designer do the heavy lifting of creating the necessary tables and 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 QAD Data Warehouse Designer. 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, QAD's Pragmatic Data Warehousing Methodology can be used.

The basic concepts behind QAD'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 star schemas, 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

QAD Data Warehouse Designer supports these concepts to facilitate very rapid delivery of data warehouses.

QAD Data Warehouse Designer controls the flow of data from the source systems through transforming and modeling layers to analysis areas. Different styles of data warehousing (normalized, dimensional etc) are supported and utilize different objects, but all follow the same basic flow. While some of these stages can be skipped, the expected flow for the common approaches are:


Data Flow - Dimensional models

- 1 source (OLTP) system
- 2 load tables
- 3 stage tables
- 4 detailed (transactional) fact tables and dimensional tables
- 5 roll up fact table(s)
- 6 aggregate and/or KPI fact table(s)
- 7 views

- 8 Microsoft Analysis Services cubes
- 9 export objects

The diagram below shows the objects and the information flow:

Overview: Dimensional Data Warehouse




Data Flow - Normalized models

- 1 source (OLTP) system
- 2 load tables
- 3 stage tables
- 4 normalized tables
- 5 history tables
- 6 views
- 7 Microsoft Analysis Services cubes
- 8 export objects

The diagram below shows the objects and the information flow:

Overview: Normalized Data Warehouse




Data Flow - Data Store models

- 1 source (OLTP) system
- 2 load tables
- 3 stage tables
- 4 data store tables
- 5 history tables
- 6 views
- 7 Microsoft Analysis Services cubes
- 8 export objects

The diagram below shows the objects and the information flow:

Overview: Data Store




Data Flow - Enterprise models

- 1 source (OLTP) system
- 2 load tables
- 3 stage tables
- 4 data store tables
- 5 model tables, dimension tables, or detailed (transactional) fact tables
- 6 roll up fact table(s)
- 7 aggregate and/or KPI fact table(s)
- 8 views
- 9 export objects
- 10 Microsoft Analysis Services cubes

The diagram below shows the objects and the information flow:

Overview: **Enterprise Data Warehouse**



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, model or dimension tables. Data from multiple load tables can be combined at this level.

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

Cubes can be created from transactional tables or views.

Procedural code

QAD Data Warehouse Designer generates procedural code in the target database's native language (e.g. PL/SQL for Oracle) 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 QAD Data Warehouse Designer, 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 QAD Data Warehouse Designer 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 should a problem occur 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. QAD Data Warehouse Designer 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, QAD Data Warehouse Designer 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.

QAD Data Warehouse Designer and Traditional ETL Tools

QAD Data Warehouse Designer's core strength is in the rapid building of data warehouse structures. Organizations that have already purchased traditional ETL tools can use QAD Data Warehouse Designer as a pureplay data warehouse toolset. QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer is not large. QAD Data Warehouse Designer is tightly integrated into the data warehouse database and has an embedded data warehouse building approach. For QAD 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 QAD Data Warehouse Designer's Load, Stage, Dimension, Fact or Model table from which QAD Data Warehouse Designer builds more advanced data warehouse structures.

Documentation Roadmap

The QAD Data Warehouse Designer **Installation and Administration Guide**, the **online help**, and the QAD Data Warehouse Designer **User Guide** assume that you are proficient in the use of the Windows operating system.

QAD Data Warehouse Designer 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 QAD Data Warehouse Designer:

QAD Data Warehouse Designer 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 QAD Data Warehouse Designer software
- ☐ Validate the various software components required by QAD Data Warehouse Designer
- ☐ Install the QAD Data Warehouse Designer metadata
- ☐ Install a scheduler
- ☐ Optionally install third party data warehouse applications
- ☐ Upgrade the QAD Data Warehouse Designer software
- ☐ Create and load language files.

QAD Data Warehouse Designer User Guide

The QAD Data Warehouse Designer **User Guide** is available either as online help, as a PDF, or in a printed manual. The User Guide provides information on how to use QAD Data Warehouse Designer to build a data warehouse.

QAD Data Warehouse Designer Tutorials

The QAD Data Warehouse Designer **Tutorials** are available either as online help or as a PDF. The Tutorials introduce you to the terms and methodologies embodied in QAD Data Warehouse Designer by guiding you through:

- ☐ Building a simple dimensional data warehouse
- ☐ Creating roll-up fact tables and aggregates
- ☐ Loading an ascii file into a new load table
- ☐ Scheduling the data warehouse objects
- ☐ Fine tuning the data warehouse by creating complex dimensional hierarchies
- ☐ Fine tuning the data warehouse by adding Analysis Services cubes
- ☐ Fine tuning the data warehouse by creating KPI fact tables

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 QAD Data Warehouse Designer product.

QAD Forum

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

Chapter 2

Design Introduction

Design Introduction

QAD Data Warehouse Designer can be used to build data warehouses based on any number of design philosophies from normalized 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 QAD 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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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, normalized, star and snowflake schemas, dimensions, fact tables, etc. Refer to the QAD web site for a basic overview of data warehouse design if required.

Chapter 3

Objects and Windows

QAD Data Warehouse Designer 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 QAD Data Warehouse Designer tool are also explained.

In This Chapter

Object Types.....	14
Organizing Objects.....	17
Windows and Panes.....	28
Database Error Messages.....	37
Export Middle Pane Output.....	38

Object Types

QAD Data Warehouse Designer has a concept of objects which are combined to create real world data warehouses and data marts, fast. Each QAD Data Warehouse Designer object has properties that allow the data warehouse developer to change how the object is used.

Note: Some object types may not be available for certain types of QAD Data Warehouse Designer licenses.

QAD Data Warehouse Designer objects include:

Object Type	Purpose
Connection 	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.
Load Table 	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 database link, odbc, 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.
Dimension 	Dimension tables are the constraining elements in the star schema design, and are defined by this object type. QAD Data Warehouse Designer 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). QAD Data Warehouse Designer 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.
Dimension View 	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 QAD Data Warehouse Designer include fact views, other table views, work views and user defined views.
Stage Table 	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, QAD Data Warehouse Designer can generate different types of procedural code based on the complexity and size of the dataset and performance requirements. Examples of procedural types that can be generated are cursor, sorted cursor, set, set + cursor or set merge procedural code. 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.

Object Type	Purpose
Normalized Table 	A normalized table is a data warehouse object used to build third normal form enterprise data warehouses. In QAD Data Warehouse Designer, normalized 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.
Data Store Table 	A Data Store Table is a data warehouse object used to store any type of data for later processing. In QAD Data Warehouse Designer, 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.
Fact Table 	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.
Kpi Fact Table 	This object type supports a special type of fact table. A mandatory KPI (Key Performance Indicator) dimension provides a set of KPIs which are stored and maintained by this object type.
Aggregate 	The aggregate object type provides a means to speed up access by summarizing data to a higher grain. For dimensional models a rollup of the fact data will allow removal of dimensions that are no longer valid.
OLAP Cube 	The OLAP Cube object type uses Analysis Services cubes to deliver OLAP functionality in QAD Data Warehouse Designer. A cube is a set of related measures and dimensions that is used to analyze data from a variety of different front end tools. OLAP Cubes are built from fact objects and aggregate objects in QAD Data Warehouse Designer.
OLAP Dimension 	An OLAP Dimension is built by QAD Data Warehouse Designer 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.
Procedure 	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.

Object Type	Purpose
Host script 	Host script objects are either Windows or UNIX scripts. These scripts are maintained within the QAD Data Warehouse Designer environment and can be scheduled to run in their host environments.
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, Fact, Kpi Fact and Aggregate)
Export 	Exports are used to manage exports from the data repository. In essence, exports are the reverse of load tables, taking data from a table to a flat file.
Retro 	Retros are used to load predefined data models from modeling tools and to retrofit existing tables into the QAD Data Warehouse Designer metadata.
Retro Copy 	Retro Copies are used to retrofit existing data warehouses into the QAD Data Warehouse Designer metadata and to transfer data from the existing data warehouse to the new data warehouse.
Template 	Template objects are used in Version 4 of QAD Data Warehouse Designer to provide example code.

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 through the use of 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.

Organizing Objects

As mentioned in the previous section, there are many object types in QAD Data Warehouse Designer.

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 **F5** or the **Ctrl/R** key.

Object Groups

The objects created in QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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 Detail, Demand Planning and Inventory. In this way projects allow us to restrict the amount of information (objects) we need to deal with to just those relevant to the area being worked on.

Projects can also be used to group objects for an upcoming code promote.

The example below shows a meta repository using two project groups (Sales and Stock Control) and five additional projects (Budgets, Promotions, Sales Detail, 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 on the project name and select **Project Properties**.

New Object
New Object (from Version)
Expand
New Group
New Project
Project Properties
Rename Project
Remove Project from Group
Delete Project
Copy Objects to Another Project
Paste Objects from Another Project
Project Object Maintenance

The following screen shows three checkboxes.

The screenshot shows the 'Project Properties' dialog box. On the left, a tree view shows 'Project' selected. The main area contains the following fields and controls:

- Name:** Order Prep
- Description:** (Empty text box)
- ☒ Include in User Documentation
- ☒ Include in Tech Documentation
- ☐ Local Project
- Date Created in Database:** 2012-02-14 23:37:57.000000
- By:** WhereScape QuickStart
- Date Project Last Modified:** 2012-11-27 00:52:54.660000
- By:** x

At the bottom right are buttons for OK, Cancel, and Help.

Include in User Documentation - When checked, the project will be included in the User Documentation. The default state is checked.

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

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

Project Groups


Projects can in turn be grouped together into 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, type **Ctrl/C**, click the target project (either the project folder or any object group or object within it) and type **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**
- ☐ **Using the Project/Object Maintenance Facility** (see "*Using Project Object Maintenance*" on page 24)


Both options above that use the right-click **Projects/Add to Project** menu result in the following dialog box being displayed:



The screenshot shows a dialog box titled "Select the Project". It has a standard Windows-style title bar with a close button (X) in the top right corner. The dialog contains two dropdown menus. The first is labeled "Group:" and is currently empty. The second is labeled "Project:" and is also empty. Below these dropdowns is a checkbox labeled "Include Associated Objects" which is checked. At the bottom of the dialog are two buttons: "OK" and "Cancel".


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

Choose the required project from the project drop-down list and click **OK**.




This screenshot shows the same "Select the Project" dialog box as before, but with the "Project:" dropdown menu now containing the text "Sales". The "Group:" dropdown remains empty, and the "Include Associated Objects" checkbox is still checked. The "OK" and "Cancel" buttons are at the bottom.

If Groups have been created, both drop-down lists will be visible. If Groups have not been created, only the project drop-down will be visible like this:



To add an object to a dependent project

- 1 Choose the required **Group** from the group drop-down list.
- 2 Choose the required **Project** from the project drop-down list.
- 3 Click **OK**.



Note: The **Include Associated Objects** checkbox on the above dialog boxes, will also add any indexes, procedures and scripts to the selected project.

Removing Objects from Projects

To remove objects from a project


There are several 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**
- ☐ **Using the Project/Object Maintenance Facility** (see "*Using Project Object Maintenance*" on page 24)

Using Project Object Maintenance

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

The following dialog 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 right. 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** will update object/project relationships without having to exit the Project Object Maintenance Facility.

Clicking **Refresh** will refresh the object tree in the left pane.

Adding Projects to Groups

The best way to move a project into a group is to create the project in the group in the first place. 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 27)

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**.

Windows and Panes


QAD Data Warehouse Designer 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. These are the **Builder** window, the **Scheduler** window, the **Diagram** window and the **Procedure Editor** window.

Main Builder Window

The following graphic is a snapshot of the main standard **Builder Window**. This window typically has three active 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. Dimension). 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 38).

The **right pane** is the browser pane and it shows source systems. There are two browser panes available at any one time. This source may be the data warehouse itself. Typically this pane is used as the source of information in the drag and drop operations. The status line at the bottom of the screen displays the connection used in this right pane.


The **bottom pane** is the results pane. It shows the results of any command executed in the tool. The messages are related to an object, multiple messages can show as a result of a command executed for an object. 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 available in all four panes.

F5 or Ctrl/R 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 **diagrammatic view** button (diagram view) switches to the diagrammatic window. This window allows the representation and printing of star schema and track back diagrams.

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

The **two source browse** buttons (one orange and one blue) allow a quick method of invoking the source browser which populates the 'Browser pane'. Each of the two browse buttons, when chosen, 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 on line QAD forum in a new tab. To select or enter other web links, click the down arrow beside the **Web Links** button.

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 three panes. The toolbar provides a quick way to display various job selections in the top pane. Other jobs can be displayed through the Jobs menu. Double-click on the job in the middle pane to see the tasks of the selected job. Double-click on the task to see the audit trail in the bottom pane.

The following is an example of the **Scheduler** window:

The screenshot shows the Scheduler window with the following data:

Job	Status	Seq	Start	Finish	Elapsed	OK	Info	Detail	Warn	Error	Who
Daily Update	On Hold		... 2013-08-16 08:46:21								JS
Process load_budget 6073	Failed		... 2013-08-21 15:24:52	2013-08-21 15:24:54	00:00		2			2	JS
Daily Update	Completed		... 2013-08-16 08:46:21	2013-08-16 08:46:24	00:00	11	26				JS


Task	Action	Status	Seq	Start	Finish	Elap...	Info	De...	Wa...	Result
load_product	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	4			9 rows loaded into load_product
load_customer	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	4			6 rows loaded into load_customer
load_order_line	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	4			21 rows loaded into load_order_line
load_order_header	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	4			9 rows loaded into load_order_header
dim_customer	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	2			dim_customer updated. 0 new records. 6 records updated.
dim_date	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	2			Date Dimension updated from Dec 8 1999 12:00AM to Apr 25 2027 ...
dim_product	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	2			dim_product updated. 0 new records. 9 records updated.
stage_sales_detail	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	2			stage_sales_detail updated. 21 new records. 0 records updated.
fact_sales_detail	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	2			fact_sales_detail updated. 0 new records. 21 records updated.
fact_sales_detail	Analyze	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00				Table fact_sales_detail analyzed

Database time at last action was 17:01:42 | Development | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

The various options and displays are covered in the **Scheduler** (on page 673) chapter.

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 Diagrammatic view window is used to display the tables of the data warehouse in diagrammatic form. For fact and aggregate tables a star schema diagram can be produced. For all table types a track back or track forward diagram can be displayed showing the various sources or targets of information that went into the makeup of the chosen table.

The following is an example of a Source Diagram in the **Diagram** window:



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 QAD 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. Multiple such windows can be open at any one time, each processing a different procedure. An example of the **Procedure Editor** window is as follows:




Comments (identified by a leading double dash --) are displayed in green in this window and the procedural code in black. The font by default is a fixed pitch font 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 627) chapter for more information on the procedure editor.

Toolbar

The procedure editor toolbar is shown below:



The **Save** button will write the procedure to the QAD 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. For Oracle data warehouses, if a procedure fails to compile it is flagged as invalid within the database.

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


Database Error Messages

From time to time, a SQL error may be raised by QAD Data Warehouse Designer.

These are usually caused by the data warehouse database returning an error.

There are many causes of database error messages being returned to Data Warehouse Designer. The most common causes include permission issues and invalid SQL statements being run.

To find the exact SQL statement run by QAD Data Warehouse Designer, go to the **View** menu and select **Last SQL**. The following dialog box will be displayed:




The four fields shown are:

- ☐ The last SQL statement run
- ☐ The second to last SQL statement run
- ☐ The third to last SQL statement run
- ☐ The error message (if any) from the last SQL statement

In the example above, we have tried to drop a procedure that does not exist. IBM DB2 returned an error telling us Data Warehouse Designer.UPDATE_DIM_KPI does not exist in the database.

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 will be 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 will be created.

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



To view the other settings for **Middle Pane File Output**, see **Export Middle Pane Output Settings** (see "**Outputs**" on page 97).

- ☐ 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 97).

Chapter 4

Tutorials

The QAD Data Warehouse Designer tutorials are located in the QAD Data Warehouse Designer **Tutorial** Help.

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 either **Options** or **User Preferences**.

In This Chapter

Settings - Options	44
Settings - User Preferences	80
Settings - Language Options.....	102


Settings - Options

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



Settings - Repository Identification

This option allows you to set the **Repository Identification** settings.



Repository Name

Set the **Name** for the repository. This name appears in the title bar in QAD Data Warehouse Designer. Restart QAD Data Warehouse Designer for repository name changes to take affect.

Repository Type

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


Data Warehouse Schema

Set the **Schema** for the repository. The schema of the data warehouse is used by the QAD Data Warehouse Designer scheduler. You should not normally need to change this value.

Settings - Object Types

Object Type Availability

This option allows you to **enable** or **disable** the various object types within the repository.




Object Type Availability

Enable object types in the data warehouse by setting to **True**.

Object Type Names

This option allows you to set the **names** for the various object types.




Object Type Name

Set the **Name** for each object type.

Object Type Ordering

This option allows you to set the **ordering** in which the object types appear in the object tree.




Object Type Order

Set the **Ordering** of the object types as displayed in the object tree.

Object End User Setting

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




Object Type End User Visible

Set to **True** to make object type an end user object, else False.

Object Color

This option allows you to set the **diagram colors** for each object type.




Object Type Diagram Color

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

Settings - Global Naming Conventions

Case Conversion

This option allows you to set the **case conversion** methods for Tables and Columns.




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 allows 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 (QAD Data Warehouse Designer 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. So, 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 allows you to set the **Global Naming of Indexes** options.



Whenever a new procedure is defined, QAD Data Warehouse Designer 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 a 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 specifications above QAD Data Warehouse Designer 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 "_SC". When a new index is manually added it will have the additional value of "_x" by default. This should be changed. The QAD Data Warehouse Designer naming standard for indexes is described below, but any valid name may be used.

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

Ultimate suffix	meaning
_0	artificial key
_1 thru _99	bitmap key index on dimensional join
_A	primary business key
_B thru _Z	secondary business keys
_SC	key to support slowly changing dimensions

Global Naming of Key Columns

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



During the drag and drop generation of new dimension and fact tables, QAD Data Warehouse Designer will build an artificial (surrogate) key for the table. The naming convention for that key can be set through the same menu option as above. As well as potential pre-fix and suffix values, we also need to choose between the inclusion of the full table name, or the short 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". So our load_customer table example would generate a key called dim_customer_key if it were dragged into a dimension drop target

Global Naming of Procedures

The default naming conventions for generated procedures can be set through the menu option **Tools/Options**. To generate a procedure select the **(Build Procedure...)** option from the **Update Procedure** drop list found in the table Properties screen.

Procedure name defaults

The dialog shown in the screen shot below will appear in response to the procedure name defaults menu option. It provides a means of setting the naming defaults for all types of generated procedure. The values as shown below are the default settings, but may be changed to meet the site requirements. The only restriction is on the size of the resultant name, which is database dependent.

The screenshot shows the 'Options' dialog box with the 'Global Naming of Procedures/Scripts' section selected. The dialog is divided into three main sections: Load Table, Stage Table, and Dimension. Each section contains a table of settings for procedure naming.

Section	Setting	Value	Type
Load Table	Load Table Post Procedure Prefix	post_	
	Load Table Post Procedure Name Type	Short name	▼
	Load Table Post Procedure Suffix		
	Load Table Script Prefix	script_	
	Load Table Script Name Type	Short name	▼
	Load Table Script Suffix		
Stage Table	Stage Table Update Procedure Prefix	update_	
	Stage Table Update Procedure Name Type	Short name	▼
	Stage Table Update Procedure Suffix		
	Stage Table Custom Procedure Prefix	custom_	
	Stage Table Custom Procedure Name Type	Short name	▼
	Stage Table Custom Procedure Suffix		
Dimension	Dimension Update Procedure Prefix	update_	
	Dimension Update Procedure Name Type	Short name	▼
	Dimension Update Procedure Suffix		
	Dimension Custom Procedure Prefix	custom_	
	Dimension Custom Procedure Name Type	Short name	▼
	Dimension Custom Procedure Suffix		
	Dimension GetKey Procedure Prefix	get_	▼

At the bottom of the dialog are buttons for 'Prev', 'Next', 'OK', 'Cancel', and 'Help'.

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 dimension called `dim_product`, then from the example screen above the three possible generated procedures would be called `update_dim_product`, `custom_dim_product` and `get_dim_product_key`.

Settings - DSS Tables and Columns

When building the data warehouse QAD Data Warehouse Designer makes use of a number of special tables and columns. Two tables are used. These are called by default `dss_source_system` and `dss_fact_table` and are discussed in detail in the sections below.

The special columns used are defined in the table below.

Column name	Description
<code>dss_source_system_key</code>	Added to support dimensions that cannot be fully conformed, and the inclusion of subsequent source systems. See the section below for more details.
<code>dss_fact_table_key</code>	Used in composite rollup fact tables to identify the source fact table that contributed the particular row.
<code>dss_create_time</code>	Indicates when a dimension record was created.
<code>dss_update_time</code>	Indicates when the record was last updated in the data warehouse. Used in the updating of rollup fact tables and aggregate tables.
<code>dss_count</code>	Applied to fact tables. Provides a simple row count variable that can be used by end user tools.
<code>dss_current_flag</code>	Used for slowly changing dimensions. This flag identifies the current record where multiple versions exist.
<code>dss_version</code>	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.
<code>dss_start_date</code>	Used for slowly changing dimensions. This column provides a date time stamp when the dimension record began life. If null then it was the first record. It is used to ascertain which dimension record should be used when multiple are available.
<code>dss_end_date</code>	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.

All of these special columns and tables may be renamed through the **Tools/Options/DSS Tables and Columns** menu option. The columns can simply be renamed, but the 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` and `dss_fact_table` 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 is designed to identify a data source for a dimension row. Its purpose is to handle non conformed dimensions, or 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 to be non conformed. We do however, insist on a standard factory name 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 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. So, manual code changes will be required to make use of the functionality that this table offers.

When dimensions are built the `dss_source_system_key` is added by default. It can be deleted if the need for such a feature is not seen.

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 Change the object type from dimension view to dimension and specify the new table name. (see note above on the use of "dss_").
- 3 Rename the `dss_source_system_name` column to match the new table name.
- 4 Delete the last two columns.
- 5 Under the tables Properties, change the table type to **Mapping table**. This will prevent the table from being seen as a dimension in the documentation.
- 6 Change the `dss_source_system` table name in the screen above, via the **Tools/Options/DSS Tables and Columns** defaults menu option.

Dss_fact_table

This pseudo dimension is used internally by the QAD Data Warehouse Designer generated procedures to assist in the updating of rollup fact tables. It provides a means of identifying which fact table contributed a particular row to a composite rollup fact table. See the section on rollup fact tables for an explanation of how the generated procedures operate.

This table can be renamed, or a new table created. The table identified as the 'dss_fact_table' must, however, conform to a number of standards if the generated code is to compile and work correctly.


The following columns must exist, where *table_name* is the new name of the table:

Column name	purpose
<i>table_name_key</i>	The artificial key for the table. Must be a numeric data type.
<i>table_name_name</i>	The name of the fact table. Must be at least a varchar2(64) data type.
<i>table_name_type</i>	The type of fact table. Must be at least a varchar2(24).

The generated procedures will automatically add rows to this table as required. As with the dss_source_system table the table type for this table should be set to **Mapping table** to prevent it from being seen as a dimension in the documentation.

DSS Tables

This option allows you to set the **DSS Tables**.




Tables

Set the **DSS Tables**.

DSS Columns

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



Columns

Set the **DSS Columns**.

Settings - Check-Out and Check-In

This option allows you to set up for the **Checking-out** or **Checking-In** of Procedures.

The screenshot shows a software window titled 'Options'. On the left is a tree view with the following items: Repository Identification, Object Types, Global Naming Conventions, DSS Tables and Columns (expanded), DSS Tables to Include and Naming, DSS Columns to Include and Naming, **Check-Out and Check-In** (selected), Code Generation, Storage, Metadata Versioning, Documentation, and Other. The main area on the right contains settings for 'Check out' and 'Retention Period'. The 'Check out' section has 'Enabled' set to 'False' and 'Mandatory Reason' set to 'True'. The 'Retention Period' section has 'Years' set to '1' and 'Months' set to '0'. At the bottom are buttons for 'Prev', 'Next', 'OK', 'Cancel', and 'Help'.

Check out	
Enabled	False
Mandatory Reason	True

Retention Period	
Years	1
Months	0

Check out

Enabled: Set to **True** to enable procedures to be checked-in or checked-out, else False.

Mandatory Reason: Set to **True** if a reason is mandatory, else False.


Retention Period

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

Settings - Code Generation

General

This option allows you to set some general **Code Generation** settings.



General

Include WsWrkTask Procedure: When set to **True**, this will result 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**.

SQL Server Integration Services

Sets the **version** of SSIS being used, else set the **Not Available** option.

Default Update Procedure Options

This option allows you to set some **default update procedure** settings.

The screenshot shows a software window titled "Options" with a standard Windows-style title bar (minimize, maximize, close buttons). On the left is a tree view with the following structure:

- Repository Identification
 - Object Types
 - Global Naming Conventions
 - DSS Tables and Columns
 - Check-Out and Check-In
 - Code Generation
 - General
 - Default Update Procedure Options** (highlighted)
 - Storage
 - Metadata Versioning
 - Documentation
 - Other

The main area of the dialog is a table of settings. At the top of this area are two icons: a list icon and a refresh icon. The table is organized into sections: "Dimension" and "Data Store". Each section contains several rows of settings, each with a name, a value, and a dropdown arrow.

Section	Setting Name	Value
Dimension	Process Dimension in Batch	False
	Include Initial Load Insert for Dimension	False
	Insert Zero Key Record for Dimension	True
	Enable Parallel DML for Dimension	False
	Distinct Data Select for Dimension	False
	Delete before Insert for Dimension	False
	Include Update Statement for Dimension	True
	Update Changed Rows Only	True
	Update Hint	
	Include Insert Statement for Dimension	True
Data Store	Process Data Store in Batch	False
	Include Initial Load Insert for Data Store	False
	Insert Zero Key Record for Data Store	True
	Enable Parallel DML for Data Store	False
	Distinct Data Select for Data Store	False
	Delete before Insert for Data Store	False
	Include Update Statement for Data Store	True
	Insert New Rows Only	True
	Insert Hint	TABLOCK
	Include Merge Statement for Dimension	False

At the bottom of the dialog are four buttons: "Prev", "Next", "OK", and "Help".

The **Process in Batch** option, when selected, allows users to select a column to drive data processing in a loop based on the distinct ordered values of the selected column.

The **Include Initial Load Insert** option, when selected, adds 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.

The **Insert Zero Key Record** option, when selected, adds an insert statement for an unknown record with an artificial key of zero. Only applicable to tables with an artificial key.

The **Parallel DML** option, when selected, adds all code required to the update procedure for enabling Oracle parallel inserts. **Note:** Oracle only.

The **Distinct Data Select** option, when selected, ensures duplicate rows are not added to the table.

The **Delete before Insert** option, when selected, enables a delete statement to be added to the update procedure before any update or insert statement.

The **Include Update Statement** option, when selected, includes an update statement in the procedure to update changing rows in the table.

The **Update Changed Rows Only** option, when selected, uses change detection to work out what rows require updating.

The **Update Hint** option, when selected, enters a database hint to be used in the UPDATE statement.

The **Include Insert Statement** option, when selected, includes an insert statement in the procedure to insert new rows in the table.


The **Insert New Rows Only** option, when selected, uses change detection to work out what rows require inserting.

The **Insert Hint** option, when selected, enters a database hint to be used in the INSERT statement.

The **Include Merge Statement** option, when selected, includes a merge statement in the procedure to merge new/changed rows in the table.

Generated Code Statistics

This option is only available for **Oracle** databases.



Generated Code Statistics

The **Include Statistics in Generated Stage Update Procedure** option sets the statistics command to be included in the automatically generated procedure for Stage tables. Parameter \$TABLE\$ will be automatically replaced by the table name. The key word \$UNUSED\$ will remove the statistics portion from the generated code in the procedure.

The **Stage Table Inserted Rows Threshold for Statistics** option sets the threshold parameter to be used by the four statistic commands that are inserted into the automatically generated procedures. It indicates how many rows are to be inserted before the statistic commands are executed.

The **Include Statistics in Generated Dimension Update Procedure** option sets the statistics command to be included in the automatically generated procedure for Dimension tables. Parameter \$TABLE\$ will be automatically replaced by the table name. The key word \$UNUSED\$ will remove the statistics portion from the generated code in the procedure.

The **Dimension Table Inserted Rows Threshold for Statistics** option sets the threshold parameter used by the four statistic commands that are inserted into the automatically generated procedures. It indicates how many rows are to be inserted before the statistic commands are executed.

The **Include Statistics in Generated Fact Update Procedure** option sets the statistics command to be included in the automatically generated procedure for Fact tables. Parameter \$TABLE\$ will be automatically replaced by the table name. The key word \$UNUSED\$ will remove the statistics portion from the generated code in the procedure.


The **Fact Table Inserted Rows Threshold for Statistics** option sets the threshold parameter to be used by the four statistic commands that are inserted into the automatically generated procedures. It indicates how many rows are to be inserted before the statistic commands are executed.

The **Include Statistics during Partition Exchange** option sets the statistics command to be included in the automatically generated procedure during partition exchanges. Parameter \$TABLE\$ will be automatically replaced by the table name. The key word \$UNUSED\$ will remove the statistics portion from the generated code in the procedure.

Settings - Storage

Table Storage

This option allows you to set the **Storage** locations and to set the type of indexes built for foreign key columns.




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.

Index Type

This option allows you to set the default type of **Foreign Key Index Type** for each table type.



Foreign Key Index Type

Set the default index types built for foreign key columns of Dimension tables, Kpi Fact tables, Fact tables and Aggregate tables.

The options for SQL Server and DB2 are:

- ☐ (None) - we will not define
- ☐ BTREE - this is a standard index per key column
- ☐ COVERING - this is a single index with all the key columns in it


The options for Oracle are:

- ☐ (None) - will not be defined
- ☐ BITMAP - This is the standard oracle bitmap index
- ☐ BTREE - this is a standard index per key column

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 - Versioning

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




Metadata Versioning

Select **True** or **False** to set when to auto-version the metadata.

Settings - Scheduler

Scheduler General


This option is only available for **Oracle** databases.



On Oracle, this enables you to set whether or not the scheduler logon uses the Authenticated User.

Scheduler Statistics

This option is only available for **Oracle** databases.



Statistics command entered here will be executed when using the scheduler with Quick Stats Action to gather statistics. Parameter \$TABLE\$ will be automatically replaced by the table name, parameter \$INDEX\$ will be replaced by the index name; and parameter \$SCHEMA\$ will be replaced by the owner name.

Below are some examples:

Scheduler Statistics	
Statistics Command for Table Analyze Full	<code>analyze table \$TABLE\$ estimate statistics sample 35 percent</code>
Statistics Command for Table Analyze Quick	<code>analyze table \$TABLE\$ estimate statistics sample 1000 rows</code>
Statistics Command for Table Statistics Full	<code>BEGIN dbms_stats.gather_table_stats(ownname=>'\$SCHEMA',tabname=>'\$TABLE\$',estimate_percent=>'50',degree=>16); END;</code>
Statistics Command for Table Statistics Quick	<code>BEGIN dbms_stats.gather_table_stats(ownname=>'\$SCHEMA',tabname=>'\$TABLE\$',estimate_percent=>'10',degree=>16); END;</code>
Statistics Command for Index Analyze Full	<code>analyze index \$INDEX\$ estimate statistics sample 10 percent</code>
Statistics Command for Index Analyze Quick	<code>analyze index \$INDEX\$ estimate statistics sample 1 percent</code>
Statistics Command for Index Statistics Full	<code>BEGIN dbms_stats.gather_index_stats(ownname=>'\$SCHEMA',indname=>'\$INDEX\$',estimate_percent=>'20',degree=>16); END;</code>
Statistics Command for Index Statistics Quick	<code>BEGIN dbms_stats.gather_index_stats(ownname=>'\$SCHEMA',indname=>'\$INDEX\$',estimate_percent=>'5',degree=>16); END;</code>

Settings - Documentation

This options allows you to alter the **Documentation** settings.

The screenshot shows the 'Options' dialog box with the 'Documentation' section selected. The left pane lists various metadata settings, with 'Project' highlighted under the 'Documentation' category. The right pane displays configuration options for four project documentation fields.

Project Documentation Field 1	
Project Documentation 1 Name	
Project Documentation 1 Label	
Project Documentation 1 User	False
Project Documentation 1 Before Columns	False
Project Documentation 1 Order	1

Project Documentation Field 2	
Project Documentation 2 Name	
Project Documentation 2 Label	
Project Documentation 2 User	False
Project Documentation 2 Before Columns	False
Project Documentation 2 Order	2

Project Documentation Field 3	
Project Documentation 3 Name	
Project Documentation 3 Label	
Project Documentation 3 User	False
Project Documentation 3 Before Columns	False
Project Documentation 3 Order	3

Project Documentation Field 4	
Project Documentation 4 Name	
Project Documentation 4 Label	
Project Documentation 4 User	False

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 dialog tabs.

Settings - Other

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




Repository Identification

Set to **True** to prevent a shadow appearing on all printable diagrams produced in the diagrammatic window; else False.

Settings - User Preferences

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




Settings - Common

Look and Feel

General

This option allows you to set the **Look and Feel** in **General**.



General

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

The **Maximize QAD Data Warehouse Designer on Startup** option, when selected, starts QAD Data Warehouse Designer in full screen mode.

The **Show Dimension Views as a Separate Object Type** option, when selected, displays a new object group called Dimension Views where all view objects are stored.

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

The **Scheduler Results in Color**, when selected, turns on job status color coding in the scheduler.

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

The **Update Column then Previous/Next Wraps to End/Start** option, when selected, 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.

Code Editor

This option allows you to set the **Look and Feel** in **Code Editor**.



Code Editor

The **Show Code as Word Wrapped** option, when selected, defaults to have word wrapping applied to code.

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

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

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


The **Procedure Indent Size** option allows 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 specifies 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 specifies 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 allows you to set the **Look and Feel** in **Confirmation Prompts**.




Confirmation Prompts

Set to true if you want to **Prompt to Regenerate Indexes when Rebuild Procedures**, else False.

Set to true if you want to **Prompt when Truncate Table via Context Menu**, else False.

Diagrams

This option allows you to set the **Look and Feel** in the **Diagrams**.




Diagram

The **Draw Diagrams using the Latest Style** option will use the new diagramming tool.

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

Property Grids

This option allows you to set the **Look and Feel** in **Property Grids**.



Property Grid

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

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

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

The **Default Property Grid Sort Order** option, allows 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, allows you to select how boolean items are to be displayed. The options are Text and Checkbox and the default is Text.

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

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

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


The **Minimum displayed lines for MultiLine items** option, gives the minimum display lines for multiline inputs.

The **Maximum displayed lines for MultiLine items** option, gives the maximum display lines for multiline inputs.

Local Naming Conventions

General

This option allows you to set the **Local Naming Conventions**.




Local Naming General

Set to True if you want to **Use Local Naming Conventions**, else False.

Local Naming of Tables


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



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

Local Naming of Key Columns

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




The **Key Prefix** option sets the prefix that will be 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 that will be used in the default naming convention.

Local Naming of Indexes

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




The **Index Prefix** option allows you to set the prefix that will be used in the default index naming convention.

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

The **Index Suffix** allows you to set the suffix that will be used in the default index naming convention.

Local Paths

This option allows you to set the **Local Paths** for documentation, backup and restore and for versioning to disk.



Documentation Path

Set the **Local Documentation Directory**.

Backup And Restore

The **Backup Executable** option sets the override for backup executable. By default QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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, QAD Data Warehouse Designer will automatically create ascii files containing the applicable ddl or code each time an automated version occurs in the entered directory.

Outputs

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



Middle Pane File Output

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

The **Output File Extension** option allows 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, when set to True, results in files created from the middle pane being opened automatically.

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

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

The **Output File End Of Line** option allows 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 allows 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 allows 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 allows you to set the characters that will replace the delimiter character if it occurs inside a field.

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

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

Other

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




Other

The **Trace Unix Sessions** option, When set to True, this will trace all Unix activity undertaken by the QAD Data Warehouse Designer program until terminated. The file WslMedTelnet.txt will be created in the program directory for QAD Data Warehouse Designer. This option is intended for debugging of specific Unix problems and the setting of this switch would normally be done at the request of QAD when attempting to solve a Telnet issue. This setting is only relevant for the PC on which the setting is made. (i.e. it is not a global setting for the repository).

Settings - Current Repository

Look and Feel

This allows 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.



The screenshot shows the 'Language Options' dialog box. It has a title bar with a close button (X). Inside, there are three main sections: 'Language:' with a dropdown menu showing 'French', 'Language Description:' with a large text area, and 'Analysis Services Language:' with a dropdown menu showing 'French (Belgium)'. To the right of the 'Language:' dropdown are two buttons: 'Add Language' and 'Delete Language'. At the bottom right are 'OK' and 'Cancel' buttons.

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 QAD Data Warehouse Designer environment and procedures. They can be edited within the QAD Data Warehouse Designer environment by selecting the **Tools/Parameters** menu option. A list of parameters is displayed as per the example below:

Parameter listing		
Parameter	Value	Comments
_0	23,3983000	load_sales_quota loaded at 10:09:08, Monday, 12 November, 2012
FACT_SALES_KPI_END	201312	
FACT_SALES_KPI_START	200701	
load_address_0	11000	Rowcount from a native load
load_address_s2_0	11000	Rowcount from native load at 16/11/12 08:26:49
load_address_sales2_0	11000	Rowcount from a native load
load_customer_sales2_0	5000	Rowcount from native load at 20/11/12 11:51:04
load_order_header_0	12500	Rowcount from native load at 16/11/12 08:26:51
load_order_header_s2_0	12500	Rowcount from native load at 16/11/12 08:26:53
load_order_line_0	100000	Rowcount from native load at 16/11/12 08:26:58
load_order_line_s2_0	100000	Rowcount from native load at 16/11/12 08:27:04
load_order_line_sales2_0	100000	Rowcount from a native load
load_prod_sub_grp_s2_0	300	Rowcount from native load at 20/11/12 11:51:14
load_product_0	5000	Rowcount from a native load
load_product_group_0	50	Rowcount from native load at 19/11/12 07:59:33
load_product_group_s2_0	50	Rowcount from native load at 20/11/12 11:51:10
load_product_line_0	725	Rowcount from native load at 19/11/12 07:59:34
load_product_line_s2_0	725	Rowcount from native load at 20/11/12 11:51:13

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

Add Parameter...
Copy Parameter...
Modify Parameter...
Used by ...
Delete Parameter...

Typical parameter usage may be the global definition of how many days should be looked back for changed 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** (on page 241) for more information.

They are also used by **stage table procedures** as variables. See **Generating the Staging Update Procedure** (on page 349) for more information.

IBM DB2: Two procedures are provided to allow procedures to read and write parameters. These procedures are **WsParameterRead** (see "**WsParameterReadF**" on page 1004) and **WsParameterWrite** (on page 1013). Using these procedures a procedure can load and use the contents of a parameter, or modify an existing parameter, or add a new parameter.

Oracle: Two functions are provided to allow procedures to read and write parameters. These functions are **WsParameterRead** (see "**WsParameterReadF**" on page 1004) and **WsParameterWrite** (on page 1013). Using these functions a procedure can load and use the contents of a parameter, or modify an existing parameter, or add a new parameter.

SQL Server: Two procedures are provided to allow procedures to read and write parameters. These procedures are **WsParameterRead** (see "**WsParameterReadF**" on page 1004) and **WsParameterWrite** (on page 1013). Using these procedures a procedure can load and use the contents of a parameter, or modify an existing parameter, or add a new parameter. A function **WsParameterReadF** (on page 1004) is also provided. This function will return the value of a parameter.

Global Parameters

A number of global parameters are provided when loading tables either via QAD Data Warehouse Designer or via the scheduler. These global parameters are accessed by calling the function **WsParameterReadG** (on page 1006). The load table name and the source table name can be acquired by using this function.

See **Callable Routines** (on page 903) for more information on **WsParameterRead** (see "**WsParameterReadF**" on page 1004), **WsParameterWrite** (on page 1013), **WsParameterReadF** (on page 1004) and **WsParameterReadG** (on page 1006).

Chapter 7

Connections

Connection objects serve several purposes in QAD Data Warehouse Designer:

- 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, QAD Data Warehouse Designer:

- 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 flat files, QAD Data Warehouse Designer:

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

Note: *Windows* and *UNIX* connections do not have an **ODBC Source** property. *UNIX* connections are used 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.
- 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. QAD Data Warehouse Designer:
 - 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** of *Analysis Server 2000* if SQL Server 2000 OLAP is being used. The **Connection type** is *Microsoft Analysis Server 2005* for later versions of SQL Server OLAP.
- 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*.

Connection types

Connections can be set up via the following methods:

- ☐ Connection to the Data Warehouse
- ☐ Connections to Another Database
- ☐ ODBC Based Connections
- ☐ Connections to UNIX
- ☐ Connections to Windows
- ☐ Connections to Analysis Services

In This Chapter

Browsing a Connection.....	109
Changing a Connection's Properties	112
Creating a Database Link.....	113
SQL Server Connections	114
Oracle Connections.....	159
DB2 Connections	191

Browsing a Connection

The tables or files associated with a connection are displayed one of the following ways:

- ☐ Selecting the **Browse/Data warehouse** menu option nominally for the data warehouse connection.
- ☐ Selecting the **Browse/Source Tables** menu option nominally for a source system connection.
- ☐ Clicking on one of the two browser buttons on the toolbar:



The orange button is nominally for the data warehouse connection and the blue button is nominally for a source system connection.

- ☐ Clicking on one of the browse source system buttons in the right pane:




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 the glasses to access the source browser dialog box. The current connection being browsed is shown in the status bar at the bottom right of the screen.



TIP: When browsing to a connection, by default QAD Data Warehouse Designer will show a schema name. If you want to see all schemas, blank out the default schema name. You can also enter multiple schemas, separated by commas.

Browser Icons

When browsing a connection the following legend applies for the source tables and objects:



The legend is displayed via the **Help/Source Legend** menu.

Browsing a connection



TIP: When browsing to a connection leave the schema field blank in order to see all schemas.

A **filter** can also be applied when browsing a connection. This filter restricts what is returned and can be:

- ☐ a standard SQL table name
- ☐ a project
- ☐ a group
- ☐ show/hide tables only
- ☐ show/hide views only
- ☐ show/hide system objects
- ☐ any combination of the above options

In addition to the filter, the dialog for browsing a connection allows for the input of a **User ID** and **Password**.

The screenshot shows a dialog box titled "List Source Tables Connection" with a close button (X) in the top right corner. The dialog contains the following fields and controls:

- Connection:** A dropdown menu with "DataWarehouse" selected.
- User ID:** A text input field.
- Password:** A text input field.
- Filter:** A section containing:
 - Schema:** A text input field with "dbo" entered.
 - Name:** A dropdown menu with "(None)" selected.
 - Object Types:** A group box containing three checkboxes: "Tables" (checked), "Views" (checked), and "System Objects" (unchecked).
 - Group:** A dropdown menu with "(All)" selected.
 - Project:** A dropdown menu with "(All)" selected.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

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, QAD Data Warehouse Designer 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 will display 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 will update all the load tables with the new connection information.

Creating a Database Link

Once a connection has been set up, you can right-click on the connection name to see the following menu:



If the connection is a database link connection then the link should be created.

Once selected, the **Create Database Link** option will attempt to create a user database link to the source database. If the link already exists, a prompt will appear asking if it is ok to overwrite the existing link. A number of problems can occur during this action, and your database administrator should be consulted.

For SQL Server a trusted link will be created if the extract user in the connection is set to **dbo**. In all other cases the extract user/password will be used as the logon for the remote server.

For Oracle data warehouses the more common problems are:

- 1 Insufficient privileges to create the link. You need the **Create Database Link** privilege in order to create a new link.
- 2 TNS could not resolve the connection to the remote database. The database link uses the value entered in the **SID** field of the connection to identify the remote database. This value must be a valid entry in the Tnsnames table for the source database.

SQL Server Connections

Database

Database connection types have multiple load table types and locations that influence behavior:

- ☐ Database link load - from a table in the current SQL Server database
- ☐ Database link load - from a table in another database on the same SQL Server instance
- ☐ Database link load - through a linked server
- ☐ Integration Services load

load_product Table Definition

Properties

Load Table Name:

Unique Short Name: (maximum 22 characters)

Description:

Connection:

Load Type:

Database Link:

Script Name:

Pre-Load Action:

Pre-Load SQL:

Table Properties Clauses: (e.g. Partition)

Post Load Procedure:

Timestamps

Metadata Structure Changed:	Database Created:	Database Altered:
2012-11-12 15:55:18.747	2012-11-12 15:55:22.673	2012-11-12 15:55:22.673

OK Cancel Help

Data Warehouse

This section describes the connection to the Data Warehouse. Tutorial 1 gives basic instructions for creating a connection. This topic describes in greater detail the connection Properties as they apply to the Data Warehouse connection. This connection is used in the drag and drop functionality to create the dimension, stage, fact and aggregate tables. It is also used to create cubes.



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.

Connection DataWarehouse

Properties

Notes

General

Connection Name	DataWarehouse
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	R_1
Data Warehouse Connection Indicator	☒

Source System

Database ID	R1
Database Link Name	
Provider Name	SQLOLEDB

Database Credentials

Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
Data Type Mapping Set	(Default)
Default Transform Function Set	(Default)

When Connection is an OLAP Data Source

OLAP Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
Connection Provider/Driver	SQLOLEDB
Data Warehouse Server	DOC
Data Warehouse Database ID	R1

Connection Name

Name used to label the connection

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer. Typically this is **DataWarehouse**

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Distinguishes the special connection that identifies the QAD Data Warehouse Designer metadata repository. Set to **True**.

Note: There should only be one metadata connection in a QAD Data Warehouse Designer repository.

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.

Provider Name

Name of the connection provider/driver used to connect to the database.

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 blank if a trusted login, or the server login password.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Left blank.

Administrator User Password

Left blank.

New Table Default Load Type

The default **Load Type** 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.

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.

Default Transform Function Set

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

Cube 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.

OLAP 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.

For details on how to use the wizard to build the OLAP connection string, see **Define Data Source for OLAP Cube** (see "**Defining the Data Source for the OLAP Cube**" on page 521)

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 **SQLOLEDB**.

Data Warehouse Server

Data Warehouse Server Name, which is used when the data warehouse is used as the data source for OLAP cubes.

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.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Database Link Load - Current

Database link load - from a table in the current SQL Server database:

Connection Sales Link

Properties

Notes

General

Connection Name	Sales Link
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	SQLOLEDB
Database ID	Sales
Database Link Name	

Database Credentials

Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Provider Name

Name of the connection provider/driver used to connect to the database.

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. Since the database is on the same server as the data warehouse, no link needs to be defined. This field can be left blank.

Extract User ID

Database User that has access to SELECT from the source system tables to extract data. For more information see **remote view extract** (on page 246).

Extract User Password

The password of the data warehouse user. For SQL Server blank if a trusted login, or the server login password.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** 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.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Database Link Load - Other DB

When running a SQL Server data warehouse it is possible to create database links (linked server) to other SQL Server instances and most of the other major databases such as Oracle, DB2 etc.

When coming to the Data Warehouse from an OLTP source system, the connection defines the path to get to those source tables. A connection object must be defined for each source system.

Database link load - from a table in another database on the same SQL Server instance:

Connection Sales Link

Properties

Notes

General

Connection Name	Sales Link
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Sales2
Data Warehouse Connection Indicator	False

Source System

Provider Name	SQLOLEDB
Database ID	Sales2
Database Link Name	

Database Credentials

Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Provider Name

Name of the connection provider/driver used to connect to the database.

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. If SQL Server and the database are on the same server as the data warehouse then no link needs to be defined, so this field can be left blank.

Extract User ID

Database User that has access to SELECT from the source system tables to extract data. For more information see **remote view extract** (on page 246).

Extract User Password

The password of the data warehouse user. For SQL Server blank if a trusted login, or the server login password.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Left blank.

Administrator User Password

Left blank.

New Table Default Load Type

The default **Load Type** 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.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Database Link Load - Linked Server

A Database Link Load through a linked server can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel
- ☐ Netezza
- ☐ Greenplum
- ☐ Hive

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts. Typically C:\Temp or a sub-directory of C:\Temp is used.

Provider Name

Name of the connection provider/driver used to connect to the database.

Full Database Path Name

This field is used when you enter **Microsoft.Jet.OLEDB.4.0** as the Provider Name, ie when the source database is from MS Access or Excel. Enter the full database path name.

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

If a database link already exists to this source database then enter that link name in this field. If no link exists then enter a link name. Under some circumstances Oracle insists that this link name is the same as the tnsnames entry for the database, so it may be good practice to use the source SID as the database link name. If this is a new link then see the notes later in this section on **database link creation** (see "**Creating a Database Link**" on page 113).

Enter the defined SQL LINK name that points to the correct server, the link name does not need to be the same name as the server name.

Extract User ID

Database User that has access to SELECT from the source system tables to extract data. For more information see **remote view extract** (on page 246).

Extract User Password

The password of the data warehouse user. For SQL Server blank if a trusted login, or the server login password.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** 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.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Below you will see examples of a Database Link Load through a **linked server** with various source databases:

SQL Server

Connection Sales Link

Properties

Notes

General

Connection Name	Sales Link
Connection Type	Database
Database Type	SQL Server
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	SQLOLEDB
Database ID	Sales
Database Link Name	AppSales

Database Credentials

Extract User ID	dbo
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	...
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK Cancel Help

Oracle

Properties

Notes

Connection Sales Link

General

Connection Name	Sales Link
Connection Type	Database
Database Type	Oracle
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	MSDAORA
Database ID	Sales
Database Link Name	AppSales

Database Credentials

Extract User ID	dbo
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	...
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK

Cancel

Help

DB2 for Linux, UNIX and Windows (the DB2 in which Data Warehouse Designer can build a warehouse)

Connection Sales Link

Properties

Notes

General

Connection Name	Sales Link
Connection Type	Database
Database Type	DB2
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Work Directory	C:\temp
Provider Name	DB2OLEDB
Database ID	Sales
Database Link Name	AppSales

Database Credentials

Extract User ID	dbo
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK Cancel Help

Teradata

Connection Sales Link

Properties

Notes

General

Connection Name	Sales Link
Connection Type	Database
Database Type	Teradata
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	TDOLEDB
Database ID	Sales
Database Link Name	AppSales

Database Credentials

Extract User ID	dbo
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	...
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

Sybase

Connection Sales Link

Properties
Notes

General

Connection Name	Sales Link
Connection Type	Database
Database Type	Other
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	Sybase ASE OLE DB Provider
Database ID	Sales
Database Link Name	AppSales

Database Credentials

Extract User ID	dbo
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	...
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

MySQL

Connection Sales Link

Properties

Notes

General

Connection Name	Sales Link
Connection Type	Database
Database Type	Other
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	MySQLProv
Database ID	Sales
Database Link Name	AppSales

Database Credentials

Extract User ID	dbo
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	...
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)

Connection Sales Link

Properties

Notes

General

Connection Name	Sales Link
Connection Type	Database
Database Type	Other
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	SNAOLEDB
Database ID	Sales
Database Link Name	AppSales

Database Credentials

Extract User ID	dbo
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	...
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

MS Access

Connection Sales Link

Properties

Notes

General

Connection NameSales Link

Connection TypeDatabase

Database TypeOther

ODBC Data Source Name (DSN)Sales

Data Warehouse Connection IndicatorFalse

Source System

Provider NameMicrosoftJet.OLEDB.4.0

Database IDSales

Database Link NameAppSales

Database Credentials

Extract User IDdbo

Extract User Password***

Administrator User ID

Administrator User Password

Other

New Table Default Load TypeDatabase link load

SSIS Connection String...

Data Type Mapping Set(Default)

Connection Name

Name used to label the connection

OK

Cancel

Help

Excel

Connection Sales Link

Properties

Notes

General

Connection Name	Sales Link
Connection Type	Database
Database Type	Other
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	Microsoft.Jet.OLEDB.4.0
Database ID	Sales
Database Link Name	AppSales

Database Credentials

Extract User ID	dbo
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK Cancel Help

Integration Services Load

An integration services load can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Provider Name

Name of the connection provider/driver used to connect to the database.

Full Database Path Name

This field is used when you enter **Microsoft.Jet.OLEDB.4.0** as the Provider Name, ie when the source database is from MS Access or Excel. Enter the full database path name.

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

Leave blank.

Extract User ID

Leave blank.

Extract User Password

Leave blank.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** 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.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Below you will see examples of an **Integration Services load** with the various source databases:

SQL Server

Connection Sales SSIS

Properties

Notes

General

Connection Name	Sales SSIS
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	SQLOLEDB
Database ID	Sales
Database Link Name	

Database Credentials

Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Integration Services load
SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Security Info=False;Initial Catalog=Sales
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

Oracle

Connection Sales SSIS

Properties
Notes

General

Connection Name	Sales SSIS
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	MSDAORA
Database ID	Sales
Database Link Name	

Database Credentials

Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Integration Services load
SSIS Connection String	Provider=MSDAORA.1;Integrated Security=SSPI;Persist Secur...
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

DB2 for Linux, UNIX and Windows (the DB2 in which Data Warehouse Designer can build a warehouse)

Connection Sales SSIS

Properties

Notes

General

Connection Name	Sales SSIS
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	DB2OLEDB
Database ID	Sales
Database Link Name	

Database Credentials

Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Integration Services load
SSIS Connection String	Provider=DB2OLEDB.1;Integrated Security=SSPI;Persist Secur...
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK Cancel Help

Teradata

Connection Sales SSIS

Properties

Notes

General

Connection Name	Sales SSIS
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	TDOLEDB
Database ID	Sales
Database Link Name	

Database Credentials

Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Integration Services load
SSIS Connection String	Provider=TDOLEDB.1;Integrated Security=SSPI;Persist Securit...
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK Cancel Help

Note: For more information on how to create the SSIS Connection String, or how to use SSIS loading in Data Warehouse Designer, see **How to use SSIS loading** (see "**Loading data into Data Warehouse Designer load tables using SSIS**" on page 250)

Windows

This connection is back to the PC that you are working on, or to a host Windows PC. From Windows you can only do flat file loads.

Sample **SQL Server Windows Connection** screen:

The screenshot shows the 'Connection Windows' dialog box. On the left is a sidebar with 'Properties' and 'Notes' tabs. The main area is divided into sections: 'General', 'Windows Host', 'Credentials', and 'Other'. Each section contains a list of properties and their values.

Section	Property	Value
General	Connection Name	Windows
	Connection Type	Windows
Windows Host	Windows Host Name	
	Work Directory	c:\temp
	Database ID	Red1
	Database Server/Home Directory	
Credentials	Windows User ID	
	Windows User Password	
	DSS User ID	dbo
	DSS User Password	***
Other	Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

Connection Type

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

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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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 (e.g. Oracle SID or TNS Name, Teradata TDPID) or Database Name (e.g. as in DB2 or SQL Server).

Database Server/Home Directory

Optional to specify the Database Home Directory if it is different from the standard home directory.

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.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Dss User ID and Password

Enter the relevant details for connecting to the **Data Warehouse**.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

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 Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

To test the connection

- ☐ Select **Browse | Source Tables** from the menu strip
- ☐ In the right pane you should be able to drill down to the area required.

ODBC

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

ODBC connection types have multiple load table types and locations that influence behavior:

- ☐ Externally loaded
- ☐ Integration Services load
- ☐ Native ODBC load
- ☐ ODBC load

Load Table load_product

Properties

Storage

Source Mapping

File Actions

Notes

Load Table Name: load_product

Unique Short Name: (maximum 22 characters) load_product

Description:

Connection: Tutorial (OLTP)

Load Type: Native ODBC (selected), Externally loaded, Integration Services load, ODBC load

Database Link:

Script Name:

Pre-Load Action: Truncate

Pre-Load SQL:

Post Load Procedure: (None)

Timestamps

Metadata Structure Changed: 2013-03-01 14:37:37.900

Database Created: 2013-03-01 14:37:40.677

Database Altered: 2013-03-01 14:37:40.677

OK Cancel Help

Native ODBC Load

A Native ODBC load can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts. Typically C:\Temp or a sub-directory of C:\Temp is used.

Extract User ID

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

Extract User Password

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

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection. Set to **Native ODBC**.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Here are two examples of a Native ODBC Load. If the source database supports windows authentication, then it is not necessary to specify a username and password; as below:

Connection Sales

Properties

Notes

2

General

Connection Name

Sales

Connection Type

ODBC

Database Type

(local)

ODBC Data Source Name (DSN)

Sales

Data Warehouse Connection Indicator

False

ODBC

Work Directory

C:\temp

Credentials

Extract User ID

Extract User Password

Administrator User ID

Administrator User Password

Other

New Table Default Load Type

Native ODBC

SSIS Connection String

Data Type Mapping Set

(Default)

Connection Name

Name used to label the connection

OK

Cancel

Help

If however, the source database does not support windows authentication, then a username and password must be specified; as below:

Connection Sales

Properties

Notes

General

Connection Name	Sales
Connection Type	ODBC
Database Type	((local))
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

ODBC

Work Directory	C:\temp

Credentials

Extract User ID	tutorial
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Native ODBC
SSIS Connection String	
Data Type Mapping Set	((Default))

Connection Name

Name used to label the connection

OK Cancel Help

ODBC Load

An ODBC load can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts. Typically C:\Temp or a sub-directory of C:\Temp is used.

Extract User ID

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

Extract User Password

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

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection. Set to **ODBC load**.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Here are two examples of an ODBC Load. If the source database supports windows authentication, then it is not necessary to specify a username and password; as below:

Connection Sales

Properties

Notes

General

Connection Name	Sales
Connection Type	ODBC
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

ODBC

Work Directory	C:\temp

Credentials

Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	ODBC load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK Cancel Help

If however, the source database does not support windows authentication, then a username and password must be specified; as below:

Connection Sales

Properties

Notes

General

Connection Name	Sales
Connection Type	ODBC
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

ODBC

Work Directory	C:\temp

Credentials

Extract User ID	tutorial
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	ODBC load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

Integration Services Load

An Integration Services Load can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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. 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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts. Typically C:\Temp or a sub-directory of C:\Temp is used.

Extract User ID

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

Extract User Password

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

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection. Set to **Integration Services load**.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Here are two examples of an Integration Services Load. If the source database supports windows authentication, then it is not necessary to specify a username and password; as below:

Connection Sales SSIS

Properties

Notes

General
 Connection Name Sales SSIS
 Connection Type ODBC
 Database Type (local)
 ODBC Data Source Name (DSN) Sales
 Data Warehouse Connection Indicator False

ODBC
 Work Directory C:\temp

Credentials
 Extract User ID
 Extract User Password
 Administrator User ID
 Administrator User Password

Other
 New Table Default Load Type Integration Services load
 SSIS Connection String Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
 Data Type Mapping Set (Default)

Connection Name
 Name used to label the connection

OK Cancel Help

If however, the source database does not support windows authentication, then a username and password must be specified; as below:

Connection Sales SSIS

Properties

Notes

General

Connection Name	Sales SSIS
Connection Type	ODBC
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

ODBC

Work Directory	C:\temp

Credentials

Extract User ID	tutorial
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Integration Services load
SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK

Cancel

Help

Analysis Server 2000

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 shot follows:

Connection Cubes

Properties

Notes

General

Connection Name	Cubes
Connection Type	Analysis Server 2000

Online Analytical Processing (OLAP)

Analysis Services (SSAS) Server Name	DOC

Credentials

Analysis Services (SSAS) User ID	
Analysis Services (SSAS) User Password	

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

Connection Type

Indicates the connection source type or the connection method. Set to **Analysis Server 2000**.

Analysis Services (SSAS) Server Name

Name of the Analysis Services server instance to connect to. This must be a valid server name. Contact your system administrator if you do not have a valid server name.

Analysis Services (SSAS) User ID and Password

Not used, leave blank.

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. A sample screen shot follows:

The screenshot shows the 'Connection Cubes' dialog box. It has a 'Properties' tab and a 'Notes' section. The main area is divided into sections: 'General' with 'Connection Name' (Cubes) and 'Connection Type' (Microsoft Analysis Server 2005+); 'Online Analytical Processing (OLAP)' with 'Analysis Services (SSAS) Server Name' (DOC); and 'Credentials' with 'Analysis Services (SSAS) User ID' and 'Analysis Services (SSAS) User Password' fields. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

Connection Name

Enter a name to identify the connection to the Analysis Services server

Connection Type

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.

Oracle Connections

Database

Database connection types have multiple load table types and locations that influence behavior:

- ☐ Database link load - from a table in the current Oracle database server (includes Oracle gateways and Oracle heterogeneous services)
- ☐ Database link load - through a database link to another Oracle server

Load Table load_customer

Properties

Storage

Source Mapping

Notes

Load Table Name: load_customer

Unique Short Name: (maximum 22 characters) load_customer

Description:

Connection: Tutorial (OLTP)

Load Type: Database link load

Database Link:

Script Name: (None)

Pre-Load Action: Truncate

Pre-Load SQL:

Post Load Procedure: (None)

Timestamps

Metadata Structure Changed: 2013-02-28 16:22:24

Database Created: 2013-02-28 16:22:31

Database Altered: 2013-02-28 16:22:31

OK Cancel Help

Data Warehouse

This section describes the connection to the Data Warehouse. Tutorial 1 gives basic instructions for creating a connection. This topic describes in greater detail the connection Properties as they apply to the Data Warehouse connection. This connection is used in the drag and drop functionality to create the dimension, stage, fact and aggregate tables. It is also used to create cubes.



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.

The screenshot shows the 'Connection DataWarehouse' dialog box. It has a 'Properties' tab on the left and a 'Notes' section below it. The main area contains a table of properties grouped into sections: General, Source System, Database Credentials, Other, and When Connection is an OLAP Data Source. The 'ODBC Data Source Name (DSN)' property is highlighted in blue. At the bottom, there is a section for 'ODBC Data Source Name (DSN)' with a note about its definition in the Windows 32-bit 'ODBC Data Source Administrator'.

Connection DataWarehouse	
General	
Connection Name	DataWarehouse
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	ORCL
Data Warehouse Connection Indicator	☒
Source System	
Database ID	dssdemo
Database Link Name	
Database Credentials	
Extract User ID	dssdemo
Extract User Password	***
Administrator User ID	
Administrator User Password	
Other	
New Table Default Load Type	Database link load
SSIS Connection String	Provider=MSDAORA.1;User ID=dssdemo;Data Source=ORCL
Data Type Mapping Set	(Default)
Default Transform Function Set	(Default)
When Connection is an OLAP Data Source	
OLAP Connection String	Provider=MSDAORA.1;User ID=dssdemo;Data Source=ORCL
Connection Provider/Driver	MSDAORA
Data Warehouse Server	ORCL
Data Warehouse Database ID	dssdemo

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 must be the same on all machines that use the corresponding connection.

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer. Typically this is **DataWarehouse**

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Distinguishes the special connection that identifies the QAD Data Warehouse Designer metadata repository. Set to **True**.

Note: There should only be one metadata connection in a QAD Data Warehouse Designer repository.

Provider Name

Name of the connection provider/driver used to connect to the database.

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.

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.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Left blank.

Administrator User Password

Left blank.

New Table Default Load Type

The default **Load Type** 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.

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.

Default Transform Function Set

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

Cube 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.

OLAP 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.

For details on how to use the wizard to build the OLAP connection string, see **Define Data Source for OLAP Cube** (see "**Defining the Data Source for the OLAP Cube**" on page 521)

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 **MSDAORA**.

Data Warehouse Server

Data Warehouse Server Name, which is used when the data warehouse is used as the data source for OLAP cubes.

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.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Database Link Load - Current

Database link load - from a table in the current ORCL Server database:

Connection Tutorial (OLTP)

Properties

Notes

General

Connection Name	Tutorial (OLTP)
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	ORCL_TUT
Data Warehouse Connection Indicator	False

Source System

Provider Name	
Database ID	ORCL
Database Link Name	

Database Credentials

Extract User ID	Tutorial
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name

Name used to label the connection

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Provider Name

Name of the connection provider/driver used to connect to the database.

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. Since the database is on the same server as the data warehouse, no link needs to be defined. This field can be left blank.

Extract User ID

Database User that has access to SELECT from the source system tables to extract data. For more information see **remote view extract** (on page 246).

Extract User Password

The password of the data warehouse user.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** 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.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Database Link Load - Linked Server

Using the Oracle database it is possible to create a database link to another Oracle database as standard, or to any other database by using Oracles Heterogeneous Services. Tutorial 1 gives basic instructions for creating a connection. This topic describes in greater detail the connection Properties. From this type of connection you can do database link loads.

When coming to the Data Warehouse from an OLTP source system, the connection defines the path to get to those source tables. A connection object must be defined for each source system.

Database link load - through a database link to another Oracle server:

The screenshot shows a Windows-style dialog box titled "Connection Sales Link". On the left is a sidebar with "Properties" and "Notes" tabs. The main area is divided into sections: "General", "Source System", "Database Credentials", and "Other". Each section contains a list of properties and their values. At the bottom, there are "OK", "Cancel", and "Help" buttons.

General	
Connection Name	Sales Link
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System	
Provider Name	MSDAORA
Database ID	Sales
Database Link Name	AppSales

Database Credentials	
Extract User ID	dssdemo
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other	
New Table Default Load Type	Database link load
SSIS Connection String	...
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts. Typically **C:\Temp** or a **sub-directory of C:\Temp** is used.

Provider Name

Name of the connection provider/driver used to connect to the database.

Full Database Path Name

This field is used when you enter **Microsoft.Jet.OLEDB.4.0** as the Provider Name, ie when the source database is from MS Access or Excel. Enter the full database path name.

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

If a database link already exists to this source database then enter that link name in this field. If no link exists then enter a link name. Under some circumstances Oracle insists that this link name is the same as the tnsnames entry for the database, so it may be good practice to use the source SID as the database link name. If this is a new link then see the notes later in this section on **database link creation** (see "**Creating a Database Link**" on page 113).

Enter the defined SQL LINK name that points to the correct server, the link name does not need to be the same name as the server name.

Extract User ID

Database User that has access to SELECT from the source system tables to extract data. For more information see **remote view extract** (on page 246).

Extract User Password

The password of the data warehouse user.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** 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.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

UNIX

Note: UNIX connections are only supported from Oracle and DB2 data warehouses.

This section describes connections to UNIX. Tutorial 1 gives basic instructions for creating a connection. This topic describes in greater detail the connection Properties as they apply to UNIX connections.

From a UNIX connection you can only do flat file loads.



TIP: If the connection type is Linux then the connection name must contain the word **Linux**. This forces QAD Data Warehouse Designer to use /bin/sh shell instead of /bin/ksh shell.

Sample UNIX connection screen:

The screenshot shows the 'Connection UNIX' dialog box with the following properties:

Section	Property	Value
General	Connection Name	UNIX
	Connection Type	Unix
UNIX/Linux Host	UNIX/Linux Host Name	wsl-linux
	Work Directory	/tmp
	Database ID	ORCL
	Database Server/Home Directory	
Connection Protocol	Connection Protocol	SSH
	Secure Shell (SSH) Command	
	Pre-Login Action	
	Login Prompt	ogin as:
	Password Prompt	ssword:
	Post-Login Action	
Credentials	Command Prompt	\$
	UNIX/Linux User ID	dssdemo
	UNIX/Linux User Password	***
	DSS User ID	dssdemo
Other	DSS User Password	***
	Data Type Mapping Set	(Default)

At the bottom of the dialog, there is a section for 'Connection Name' with the description 'Name used to label the connection'.

Buttons: OK, Cancel, Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

Connection Type

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

UNIX/Linux Host Name

IP address or host name that identifies the UNIX machine.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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 (e.g. Oracle SID or TNS Name, Teradata TDPID) or Database Name (e.g. as in DB2 or SQL Server).

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\PuTTY\plink.exe -ssh some_host_name**

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. QAD Data Warehouse Designer 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 will have to log in to the UNIX system.

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.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

DSS User ID

Database User to connect to the QAD Data Warehouse Designer metadata repository.

DSS User Password

Database Password to connect to the QAD Data Warehouse Designer metadata repository.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

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)** will cause Data Warehouse Designer 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 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 menu strip select **Browse | Source Tables**
- ☐ Drill down to the area required
- ☐ Drag an item to the middle pane, (having first selected the object in the left 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 will display the login session, and should provide 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 menu option **Tools/Options** and click on the **Trace all Unix sessions** checkbox. Then try to do the connection or browse again. A log file called WslMedTelnet.log will be created in the QAD program directory. Edit the log file and ensure there are no passwords visible and then contact QAD support using the QAD forum at <http://www.qad.com>.

Closing the Connection

To close the collection, right-click in the browser pane and select **Close UNIX/LINUX session**:



Windows

This connection is back to the PC that you are working on, or to a host Windows PC. From Windows you can only do flat file loads.

Connection Windows

Properties

Notes

General

Connection Name	Windows
Connection Type	Windows

Windows Host

Windows Host Name	
Work Directory	c:\temp
Database ID	ORCL
Database Server/Home Directory	

Credentials

Windows User ID	
Windows User Password	
DSS User ID	dssdemo
DSS User Password	***

Other

Data Type Mapping Set	(Default)

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

Connection Type

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

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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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 (e.g. Oracle SID or TNS Name, Teradata TDPID) or Database Name (e.g. as in DB2 or SQL Server).

Database Server/Home Directory

Optional to specify the Database Home Directory if it is different from the standard home directory.

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.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Dss User ID and Password

Enter the relevant details for connecting to the **Data Warehouse**.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

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 Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

To test the connection

- ☐ Select **Browse | Source Tables** from the menu strip
- ☐ In the right pane you should be able to drill down to the area required.

ODBC

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

ODBC connection types have multiple load table types and locations that influence behavior:

- ☐ Externally loaded
- ☐ Integration Services load
- ☐ Native ODBC load
- ☐ ODBC load

Load Table load_customer

Properties
Storage
Source Mapping
Notes

Load Table Name: load_customer

Unique Short Name: (maximum 22 characters) load_customer

Description:

Connection: Tutorial (OLTP)

Load Type: ODBC load
Externally loaded
Integration Services load
Native ODBC
ODBC load

Database Link:

Script Name:

Pre-Load Action: Truncate

Pre-Load SQL:

Post Load Procedure: (None)

Timestamps

Metadata Structure Changed:	Database Created:	Database Altered:
2013-05-17 13:31:31	2013-05-17 13:31:42	2013-05-17 13:31:42

OK Cancel Help

Native ODBC Load

A Native ODBC load can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts. Typically **C:\Temp** or a **sub-directory of C:\Temp** is used.

Extract User ID

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

Extract User Password

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

Note: Special chars such as `<>!%$|&` should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection. Set to **Native ODBC**.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Here are two examples of a **Native ODBC Load**. If the source database supports windows authentication, then it is not necessary to specify a username and password; as below:

The screenshot shows the 'Connection Sales' dialog box. On the left, there is a sidebar with 'Properties' and 'Notes' tabs. The 'Properties' tab is active, displaying a tree view of connection properties. The tree has four main sections: 'General', 'ODBC', 'Credentials', and 'Other'. The 'General' section is expanded, showing a list of properties: 'Connection Name' (Sales), 'Connection Type' (ODBC), 'Database Type' ((local)), 'ODBC Data Source Name (DSN)' (Sales), and 'Data Warehouse Connection Indicator' (False). The 'ODBC' section is also expanded, showing 'Work Directory' (C:\temp). The 'Credentials' and 'Other' sections are collapsed. At the bottom of the dialog, there is a 'Connection Name' label and a description: 'Name used to label the connection'. The 'OK', 'Cancel', and 'Help' buttons are located at the bottom right of the dialog.

General	
Connection Name	Sales
Connection Type	ODBC
Database Type	((local))
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

ODBC	
Work Directory	C:\temp

Credentials	
Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other	
New Table Default Load Type	Native ODBC
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

If however, the source database does not support windows authentication, then a username and password must be specified; as below:

The screenshot shows the 'Connection Sales' dialog box. On the left is a 'Properties' tab with a 'Notes' section. The main area contains a tree view with the following sections:

- General**
 - Connection Name: Sales
 - Connection Type: ODBC
 - Database Type: (local)
 - ODBC Data Source Name (DSN): Sales
 - Data Warehouse Connection Indicator: False
- ODBC**
 - Work Directory: C:\temp
- Credentials**
 - Extract User ID: tutorial
 - Extract User Password: ***
 - Administrator User ID:
 - Administrator User Password:
- Other**
 - New Table Default Load Type: Native ODBC
 - SSIS Connection String:
 - Data Type Mapping Set: (Default)

At the bottom, the 'Connection Name' field is highlighted with the description: 'Name used to label the connection'. The 'OK', 'Cancel', and 'Help' buttons are located at the bottom right.

ODBC Load - Single Row Inserts

ODBC load - standard (single row inserts) in the scheduler

An ODBC load can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts. Typically **C:\Temp** or a **sub-directory of C:\Temp** is used.

Extract User ID

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

Extract User Password

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

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection. Set to **ODBC load**.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Here are two examples of an **ODBC Load**. If the source database supports windows authentication, then it is not necessary to specify a username and password; as below:

Connection Sales

Properties

Notes

2

↓

General

Connection Name

Sales

Connection Type

ODBC

▼

Database Type

(local)

▼

ODBC Data Source Name (DSN)

Sales

▼

Data Warehouse Connection Indicator

False

▼

ODBC

Work Directory

C:\temp

...

Credentials

Extract User ID

Extract User Password

Administrator User ID

Administrator User Password

Other

New Table Default Load Type

ODBC load

▼

SSIS Connection String

...

Data Type Mapping Set

(Default)

▼

Connection Name

Name used to label the connection

OK

Cancel

Help

If however, the source database does not support windows authentication, then a username and password must be specified; as below:

Connection Sales

Properties

Notes

General

Connection Name	Sales
Connection Type	ODBC
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

ODBC

Work Directory	C:\temp

Credentials

Extract User ID	tutorial
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	ODBC load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

ODBC Load - Direct Path OCI

ODBC load - using direct path OCI in the scheduler

An ODBC load can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts. Typically **C:\Temp** or a **sub-directory of C:\Temp** is used.

Extract User ID

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

Extract User Password

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

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection. Set to **ODBC load**.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.


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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Here are two examples of an ODBC Load. If the source database supports windows authentication, then it is not necessary to specify a username and password; as below:



If however, the source database does not support windows authentication, then a username and password must be specified; as below:

Connection Sales

Properties

Notes

General

Connection Name	Sales
Connection Type	ODBC
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

ODBC

Work Directory	C:\temp

Credentials

Extract User ID	tutorial
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	ODBC load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

Analysis Server 2000

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 shot follows:

Connection Cubes

Properties

Notes

General

Connection Name	Cubes
Connection Type	Analysis Server 2000

Online Analytical Processing (OLAP)

Analysis Services (SSAS) Server Name	DOC

Credentials

Analysis Services (SSAS) User ID	
Analysis Services (SSAS) User Password	

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

Connection Type

Indicates the connection source type or the connection method. Set to **Analysis Server 2000**.

Analysis Services (SSAS) Server Name

Name of the Analysis Services server instance to connect to. This must be a valid server name. Contact your system administrator if you do not have a valid server name.

Analysis Services (SSAS) User ID and Password

Not used, leave blank.

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. A sample screen shot follows:

The screenshot shows the 'Connection Cubes' dialog box. It has a 'Properties' tab and a 'Notes' section. The main area is divided into sections: 'General' with 'Connection Name' (Cubes) and 'Connection Type' (Microsoft Analysis Server 2005+); 'Online Analytical Processing (OLAP)' with 'Analysis Services (SSAS) Server Name' (DOC); and 'Credentials' with 'Analysis Services (SSAS) User ID' and 'Analysis Services (SSAS) User Password' fields. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

Connection Name

Enter a name to identify the connection to the Analysis Services server

Connection Type

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.

DB2 Connections

Database

Database connection types have multiple load table types and locations that influence behavior:

- ☐ Database link load - from another table in the same DB2 system (includes federation)
- ☐ Database link load - via CLI from a DB2 table in another system

Load Table load_customer

Properties

Load Table Name: load_customer

Unique Short Name: (maximum 12 characters) load_custome

Description:

Connection: Tutorial (OLTP)

Load Type: Database link load

Database Link:

Script Name: (None)

Pre-Load Action: Truncate

Pre-Load SQL:

Post Load Procedure: (None)

Timestamps

Metadata Structure Changed:	Database Created:	Database Altered:
2013-04-08 15:42:14.209000	2013-04-08 15:42:19.553000	2013-04-08 15:42:19.553000

OK Cancel Help

Data Warehouse

This section describes the connection to the Data Warehouse. Tutorial 1 gives basic instructions for creating a connection. This topic describes in greater detail the connection Properties as they apply to the Data Warehouse connection. This connection is used in the drag and drop functionality to create the dimension, stage, fact and aggregate tables. It is also used to create cubes.



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.

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

General	
Connection Name	DataWarehouse
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	WS
Data Warehouse Connection Indicator	☒

Source System	
Database ID	DSSDEMO
Database Link Name	
Work Directory	...

Database Credentials	
Extract User ID	wsl
Extract User Password	*****
Administrator User ID	
Administrator User Password	

Other	
New Table Default Load Type	Database link load
SSIS Connection String	...
Data Type Mapping Set	(Default)
Default Transform Function Set	(Default)

When Connection is an OLAP Data Source	
OLAP Connection String	...
Connection Provider/Driver	IBMDADB2 (OLE DB for IBM DB2)
Data Warehouse Server	
Data Warehouse Database ID	WS

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 must be the same on all machines that use the corresponding connection.

Buttons: OK, Cancel, Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer. Typically this is **DataWarehouse**

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Distinguishes the special connection that identifies the QAD Data Warehouse Designer metadata repository. Set to **True**.

Note: There should only be one metadata connection in a QAD Data Warehouse Designer repository.

Provider Name

Name of the connection provider/driver used to connect to the database.

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.

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.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Left blank.

Administrator User Password

Left blank.

New Table Default Load Type

The default **Load Type** 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.

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.

Default Transform Function Set

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

Cube 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.

OLAP 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.

For details on how to use the wizard to build the OLAP connection string, see **Define Data Source for OLAP Cube** (see "**Defining the Data Source for the OLAP Cube**" on page 521)

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 **IBMDADB2**.

Data Warehouse Server

Leave **Server** field empty for DB2

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.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

UNIX

Note: UNIX connections are only supported from Oracle and DB2 data warehouses.

This section describes connections to UNIX. Tutorial 1 gives basic instructions for creating a connection. This topic describes in greater detail the connection Properties as they apply to UNIX connections.

From a UNIX connection you can only do flat file loads.



TIP: If the connection type is Linux then the connection name must contain the word **Linux**. This forces QAD Data Warehouse Designer to use /bin/sh shell instead of /bin/ksh shell.

Sample UNIX connection screen:

Connection UNIX

Properties

Notes

General

Connection Name	UNIX
Connection Type	UNIX/Linux

UNIX/Linux Host

UNIX/Linux Host Name	LinuxHost
Work Directory	/tmp
Database ID	DB2
Database Server/Home Directory	

Connection Protocol

Connection Protocol	SSH
Secure Shell (SSH) Command	
Pre-Login Action	
Login Prompt	ogin as:
Password Prompt	ssword:
Post-Login Action	
Command Prompt	\$

Credentials

UNIX/Linux User ID	dssdemo
UNIX/Linux User Password	***
DSS User ID	dssdemo
DSS User Password	***

Other

Data Type Mapping Set	(Default)

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

Connection Type

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

UNIX/Linux Host Name

IP address or host name that identifies the UNIX machine.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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 (e.g. Oracle SID or TNS Name, Teradata TDPID) or Database Name (e.g. as in DB2 or SQL Server).

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\PuTTY\plink.exe -ssh some_host_name**

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. QAD Data Warehouse Designer 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 will have to log in to the UNIX system.

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.

Note: Special chars such as `<>!%$|&` should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

DSS User ID

Database User to connect to the QAD Data Warehouse Designer metadata repository.

DSS User Password

Database Password to connect to the QAD Data Warehouse Designer metadata repository.

Note: Special chars such as `<>!%$|&` should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

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)** will cause Data Warehouse Designer 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 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 menu strip select **Browse | Source Tables**
- ☐ Drill down to the area required
- ☐ Drag an item to the middle pane, (having first selected the object in the left 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 will display the login session, and should provide 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 menu option **Tools/Options** and click on the checkbox **Trace all Unix sessions**. Then try to do the connection or browse again. A log file called WslMedTelnet.log will be created in the QAD program directory. Edit the log file and ensure there are no passwords visible and then contact QAD support using the QAD forum at <http://www.qad.com>.

Closing the Connection

To close the collection, right-click in the browser pane and select **Close UNIX/LINUX session**:



Windows

This connection is back to the PC that you are working on, or to a host Windows PC. From Windows you can only do flat file loads.

Connection Windows

Properties

Notes

General

Connection Name	Windows
Connection Type	Windows

Windows Host

Windows Host Name	
Work Directory	c:\temp
Database ID	...
Database Server/Home Directory	...

Credentials

Windows User ID	
Windows User Password	
DSS User ID	wsl
DSS User Password	*****

Other

Data Type Mapping Set	(Default)

OK Cancel Help

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

Connection Type

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

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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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 (e.g. Oracle SID or TNS Name, Teradata TDPID) or Database Name (e.g. as in DB2 or SQL Server).

Database Server/Home Directory

Optional to specify the Database Home Directory if it is different from the standard home directory.

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.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Dss User ID and Password

Enter the relevant details for connecting to the **Data Warehouse**.

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

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 Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

To test the connection

- ☐ Select **Browse | Source Tables** from the menu strip
- ☐ In the right pane you should be able to drill down to the area required.

ODBC

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

ODBC connection types have multiple load table types and locations that influence behavior:

- ☐ Native ODBC load
- ☐ ODBC load - standard (single row inserts) in the scheduler
- ☐ ODBC load - using direct path OCI in the scheduler

Load Table load_customer

Properties

Load Table Name: load_customer

Unique Short Name: load_custome (maximum 12 characters)

Description:

Connection: Tutorial (OLTP)

Load Type: ODBC load

Database Link: Native ODBC

Script Name: [None]

Pre-Load Action: Truncate

Pre-Load SQL:

Post Load Procedure: [None]

Timestamps

Metadata Structure Changed:	Database Created:	Database Altered:
2013-06-12 21:12:23.761000	2013-06-12 21:12:28.620000	2013-06-12 21:12:28.620000

OK Cancel Help

Native ODBC Load

A Native ODBC load can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts. Typically **C:\Temp** or a **sub-directory of C:\Temp** is used.

Extract User ID

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

Extract User Password

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

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection. Set to **Native ODBC**.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Here are two examples of a Native ODBC Load. If the source database supports windows authentication, then it is not necessary to specify a username and password; as below:

Connection Sales

Properties

Notes

General

Connection NameSales

Connection TypeODBC

Database Type(local)

ODBC Data Source Name (DSN)Sales

Data Warehouse Connection IndicatorFalse

ODBC

Work DirectoryC:\temp

Credentials

Extract User ID

Extract User Password

Administrator User ID

Administrator User Password

Other

New Table Default Load TypeNative ODBC

SSIS Connection String

Data Type Mapping Set(Default)

Connection Name

Name used to label the connection

OK

Cancel

Help

If however, the source database does not support windows authentication, then a username and password must be specified; as below:

Connection Sales

Properties

Notes

General

Connection Name	Sales
Connection Type	ODBC
Database Type	((local))
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

ODBC

Work Directory	C:\temp

Credentials

Extract User ID	tutorial
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Native ODBC
SSIS Connection String	
Data Type Mapping Set	((Default))

Connection Name

Name used to label the connection

OK Cancel Help

ODBC Load - Single Row Inserts

ODBC load - standard (single row inserts) in the scheduler

An ODBC load can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts.

Typically **C:\Temp** or a **sub-directory of C:\Temp** is used.

Extract User ID

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

Extract User Password

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

Note: Special chars such as <>!%\$|& should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection. Set to **ODBC load**.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Here are two examples of an ODBC Load. If the source database supports windows authentication, then it is not necessary to specify a username and password; as below:

Connection Sales

Properties

Notes

General

Connection NameSales

Connection TypeODBC

Database Type(local)

ODBC Data Source Name (DSN)Sales

Data Warehouse Connection IndicatorFalse

ODBC

Work DirectoryC:\temp

Credentials

Extract User ID

Extract User Password

Administrator User ID

Administrator User Password

Other

New Table Default Load TypeODBC load

SSIS Connection String

Data Type Mapping Set(Default)

Connection Name

Name used to label the connection

OK

Cancel

Help

If however, the source database does not support windows authentication, then a username and password must be specified; as below:

Connection Sales

Properties

Notes

General

Connection Name	Sales
Connection Type	ODBC
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

ODBC

Work Directory	C:\temp

Credentials

Extract User ID	tutorial
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	ODBC load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

ODBC Load - Direct Path OCI

ODBC load - using direct path OCI in the scheduler

An ODBC load can have a number of source databases:

- ☐ SQL Server
- ☐ Oracle
- ☐ DB2 for Linux, UNIX and Windows (the DB2 Data Warehouse Designer can build a warehouse in)
- ☐ Teradata
- ☐ Sybase
- ☐ MySQL
- ☐ DB2/400 (the DB2 that runs on AS/400 machines - quite different to above)
- ☐ MS Access
- ☐ Excel

The connection object Properties window has the following fields:

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

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.

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 Data Warehouse Designer must be the same on all machines that use the corresponding connection.

Data Warehouse Connection Indicator

Set to **False**.

Work Directory

Windows directory used by QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler running on the same machine to avoid file conflicts.

Typically **C:\Temp** or a **sub-directory of C:\Temp** is used.

Extract User ID

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

Extract User Password

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

Note: Special chars such as **<>!%\$|&** should not be used in passwords on connections as they break certain command line calls for some load tables and export objects.

Administrator User ID

Leave blank.

Administrator User Password

Leave blank.

New Table Default Load Type

The default **Load Type** for new Load tables created using this connection. Set to **ODBC load**.

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.

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)** will cause Data Warehouse Designer to automatically select the relevant mapping set.

Once the connection has been set up, you can right-click on the connection in the middle pane and view the Properties for that connection.

Here are two examples of an ODBC Load. If the source database supports windows authentication, then it is not necessary to specify a username and password; as below:

Connection Sales

Properties

Notes

General

Connection NameSales

Connection TypeODBC

Database Type(local)

ODBC Data Source Name (DSN)Sales

Data Warehouse Connection IndicatorFalse

ODBC

Work DirectoryC:\temp

Credentials

Extract User ID

Extract User Password

Administrator User ID

Administrator User Password

Other

New Table Default Load TypeODBC load

SSIS Connection String

Data Type Mapping Set(Default)

Connection Name

Name used to label the connection

OK

Cancel

Help

If however, the source database does not support windows authentication, then a username and password must be specified; as below:

The screenshot shows the 'Connection Sales' dialog box. On the left is a 'Notes' pane. The main area is divided into sections: 'General', 'ODBC', 'Credentials', and 'Other'. The 'General' section contains 'Connection Name' (Sales), 'Connection Type' (ODBC), 'Database Type' ((local)), 'ODBC Data Source Name (DSN)' (Sales), and 'Data Warehouse Connection Indicator' (False). The 'ODBC' section contains 'Work Directory' (C:\temp). The 'Credentials' section contains 'Extract User ID' (tutorial), 'Extract User Password' (***), 'Administrator User ID' (empty), and 'Administrator User Password' (empty). The 'Other' section contains 'New Table Default Load Type' (ODBC load), 'SSIS Connection String' (empty), and 'Data Type Mapping Set' ((Default)). At the bottom, there is a 'Connection Name' label and a description 'Name used to label the connection'. The 'OK', 'Cancel', and 'Help' buttons are at the bottom right.

General	
Connection Name	Sales
Connection Type	ODBC
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

ODBC	
Work Directory	C:\temp

Credentials	
Extract User ID	tutorial
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other	
New Table Default Load Type	ODBC load
SSIS Connection String	
Data Type Mapping Set	(Default)

Connection Name
Name used to label the connection

OK Cancel Help

Analysis Server 2000

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 shot follows:

The screenshot shows the 'Connection Cubes' dialog box. It has a 'Properties' tab and a 'Notes' section. The main area is divided into three sections: 'General', 'Online Analytical Processing (OLAP)', and 'Credentials'. The 'General' section contains 'Connection Name' (Cubes) and 'Connection Type' (Analysis Server 2000). The 'Online Analytical Processing (OLAP)' section contains 'Analysis Services (SSAS) Server Name' (DOC). The 'Credentials' section contains 'Analysis Services (SSAS) User ID' and 'Analysis Services (SSAS) User Password'. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

General	
Connection Name	Cubes
Connection Type	Analysis Server 2000

Online Analytical Processing (OLAP)	
Analysis Services (SSAS) Server Name	DOC

Credentials	
Analysis Services (SSAS) User ID	
Analysis Services (SSAS) User Password	

Connection Name

Name used to label the connection within QAD Data Warehouse Designer.

Connection Type

Indicates the connection source type or the connection method. Set to **Analysis Server 2000**.

Analysis Services (SSAS) Server Name

Name of the Analysis Services server instance to connect to. This must be a valid server name. Contact your system administrator if you do not have a valid server name.

Analysis Services (SSAS) User ID and Password

Not used, leave blank.

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 shot follows:

The screenshot shows the 'Connection Cubes' dialog box. It has a title bar with a close button. On the left is a 'Properties' tab and a 'Notes' section. The main area contains the following fields:

General	
Connection Name	Cubes
Connection Type	Microsoft Analysis Server 2005+
Online Analytical Processing (OLAP)	
Analysis Services (SSAS) Server Name	DOC
Credentials	
Analysis Services (SSAS) User ID	
Analysis Services (SSAS) User Password	

At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

Connection Name

Enter a name to identify the connection to the Analysis Services server

Connection Type

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.

Chapter 8

Table Properties

Several generic properties can be set on all table objects in QAD Data Warehouse Designer. These are:

- ☐ Storage Properties
- ☐ Notes
- ☐ Documentation Fields

This Chapter discusses these properties.

In This Chapter

Table Storage Properties	220
Table Notes	226
Table Documentation Fields	228
Table Statistics in DB2	230

Table Storage Properties

The fields on the **Storage** tab of the Properties screen for tables are described below. The properties available depend on the data warehouse database environment.

- ☐ For a SQL Server example, see **Table Storage Screen - SQL Server** (on page 222).
- ☐ For an Oracle example, see **Table Storage Screen - Oracle** (on page 223).
- ☐ For a DB2 example, see **Table Storage Screen - DB2** (on page 225).

Tablespace/Filegroup

Select a name from the drop-down list box. If no tablespace is selected in Oracle or DB2, then the default tablespace for the schema is used. If no filegroup name is selected for SQL Server, the default for the database is used.

Compress

This field enables you to compress a particular table/index.

Compression options are:

- ☐ NONE
- ☐ PAGE
- ☐ ROW

VarDecimal Storage Format

Enables variable numeric storage in SQL Server for the table. The default is off - no variable numeric storage. To enable variable number storage, the following command must also be run for the SQL Server database: `sp_db_vardecimal_storage_format 'DatabaseName', 'ON'`

Optional CREATE Clause

Database-specific-and-compliant DDL to append to the generated CREATE TABLE statement.

Initial extent (KB)

Oracle only. The size of the first extent of a created database table. The default is **0** kilobytes - use the tablespace's default value. For locally managed tablespaces, this field is ignored.

Next extent (KB)

Oracle only. The size of the next (second) extent of a created database table, unless percent increase is specified. The default is **0** kilobytes - use the tablespace's default value. For locally managed tablespaces, this field is ignored.

Minimum extents

Oracle only. The number of extents the table is created with. The default is **0** - use the tablespace's default value. For locally managed tablespaces, this field is ignored.

Maximum extents

Oracle only. The maximum number of extents the table can have. The default is **0** - use the tablespace's default value. For locally managed tablespaces, this field is ignored.

Percent increase

Oracle only. The size the next extent is as a percent of the previous extent added. For example, a value of 5 will create the next extent 5% bigger than the previous extent added. The default is **0** - use the tablespace's default value. For locally managed tablespaces, this field is ignored.

Percent free

Oracle and DB2 only. The amount of free space to leave in each extent in Oracle and each page in DB2. The default is **0** - use the database default.

Buffer Pool attribute

Oracle Dimension tables only. Mark the table to be retained in a special buffer pool. Possible values are **Default**, **Keep**, **Recycle** and not set (blank). The default is not set.

Compress

DB2 only. Enable table compression. Default is off - no compression.

Volatile

DB2 only. Enable this option to alter the table after it is created to indicate size volatility to the DB2 optimizer.

Append Only

DB2 only. Enable this option to alter the table after it is created to append new data to the end of the table, rather than using available free space in the table.

Index build logging

DB2 only. Specifies the degree of logging performed by DB2 when creating, recreating and re-organizing indexes. Value is one of not selected, **NULL**, **OFF** or **ON**. Not selected and NULL are the same. The table is created without this option and then altered to add this property. The default is not set (no alter takes place). The following describes the values:

- ☐ **NULL** - Database parameter **logindexbuild** determines whether logging is minimal or full.
- ☐ **OFF** - only minimal logging occurs.
- ☐ **ON** - all operations are logged.

Table Storage Screen - SQL Server

Properties

Storage

Source Mapping

Notes

Load Table load_product

Storage

Filegroup	LoadTables	▼
Compress	(Not Defined)	▼
VarDecimal Storage Format	False	▼

Other

Optional CREATE Clause		...

OK

Cancel

Help

Table Storage Screen - Oracle

Typical Storage Tab for an **Oracle** Table:

The screenshot shows a dialog box titled "load_customer_xml Table Definition". On the left is a sidebar with four tabs: "Properties", "File Attributes", "Storage" (which is selected and highlighted in blue), and "Notes". The main area of the dialog contains several configuration fields for table storage. The "Tablespace Name" is set to "LOADTABLESDATA" in a dropdown menu. Below it are six text input fields, each with a label and a value of "0": "Initial Extent (KB)", "Next Extent (KB)", "Minimum Extents", "Maximum Extents", "Percent Increase", and "Percent Free". At the bottom right of the dialog are three buttons: "OK", "Cancel", and "Help".

Property	Value
Tablespace Name	LOADTABLESDATA
Initial Extent (KB)	0
Next Extent (KB)	0
Minimum Extents	0
Maximum Extents	0
Percent Increase	0
Percent Free	0

Storage Tab for an **Oracle Dimension Table**:

The screenshot shows a dialog box titled "dim_customer Table Definition" with a close button (X) in the top right corner. On the left is a vertical sidebar with a list of tabs: Properties, Storage (highlighted in blue), Language Mapping, Purpose, Concept, Grain, Examples, Usage, and Notes. The main area of the dialog is for the "Storage" tab and contains the following fields:

Tablespace Name:	REPORTINGTABLESDATA
Initial Extent (KB):	0
Next Extent (KB):	0
Minimum Extents:	0
Maximum Extents:	0
Percent Increase:	0
Percent Free:	0
Buffer Pool Attribute:	

At the bottom right of the dialog are three buttons: OK, Cancel, and Help.

Table Storage Screen - DB2

The screenshot shows the 'Load Table load_customer' dialog box. On the left is a sidebar with 'Properties', 'Storage' (selected), 'Source Mapping', 'Statistics', and 'Notes'. The main area is divided into 'Storage' and 'Other' sections. The 'Storage' section contains a table of properties with dropdown menus for 'Table space' and 'Index Table space', and checkboxes for 'Enable Percent Free', 'Compress', 'Volatile', 'Append New Data', and 'Index Build Logging'. The 'Other' section has an 'Optional CREATE Clause' field with a dropdown arrow. At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

Storage	
Table space	(Default) ▼
Index Table space	(Default) ▼
Enable Percent Free	False ▼
Compress	False ▼
Volatile	False ▼
Append New Data	False ▼
Index Build Logging	▼

Other	
Optional CREATE Clause	...

Table Notes

The notes screen is used to enter notes against QAD Data Warehouse Designer objects. The notes on an object are included in the QAD Data Warehouse Designer 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.

Notes can be removed from the documentation via the **Tools/Options** menu.

Table Notes Screen




Table Documentation Fields

Each table has a set of associated documentation tabs. These are used to store descriptive metadata about the fact table and how it is used. The contents of each tab appears in the generated documentation. Their order and placement can be customized in the generated documentation using the **Tools/Options** menu.

The following tabs are available to store additional informational metadata that will appear in the generated QAD Data Warehouse Designer documentation:

- ☐ Purpose
- ☐ Concept
- ☐ Grain
- ☐ Examples
- ☐ Usage

Note: These names of the tabs are completely arbitrary and can be changed in the generated documentation using the **Tools/Options** menu, Documentation tab. By default, these fields are not enabled for load and stage tables, but can be enabled using the **Tools/Options** menu, **Documentation** tab.

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** menu.

Table Documentation Fields Screen

The screenshot shows a software window titled "fact_order_detail Table Definition". On the left is a vertical sidebar with a list of tabs: Properties, Storage, Language Mapping, Purpose (which is highlighted with a blue background), Concept, Grain, Examples, Usage, and Notes. The main area of the window is titled "Documentation: Purpose" and contains a large, empty rectangular text field with a vertical scrollbar on the right side. At the bottom right of the window, there are three buttons: "OK", "Cancel", and "Help".

Table Statistics in DB2

QAD Data Warehouse Designer for DB2 allows for table level statistics definition. Four different statistics commands may be defined for a table:

- ☐ Stats Full
- ☐ Stats Quick
- ☐ Analyze Full
- ☐ Analyze Quick

These are arbitrary names, used as **actions** in the scheduler.

To set the command that is run for one of the actions above, enter the command to be run using the \$SCHEMA\$ and \$TABLE\$ variables.

Note: The variable \$SCHEMA\$ allows for metadata portability between environments. Variable \$TABLE\$ avoids any requirement to change the command if the table is renamed.

For example: enter *runstats on table \$SCHEMA\$. \$TABLE\$* to do basic statistics for the table.

If a particular statistics field is left blank, then the QAD Data Warehouse Designer default statistics options are used.

The default statistics method is the statistics method used in the scheduler if the chosen action is **Process with default stats**.

Table Statistics Screen

Dimension dim_customer

Select the statistics method to use with the Scheduler 'Process and Statistics' task action.

Process and Statistics Method:

The following commands can be used to optionally override the defaults for gathering table statistics via the Scheduler.
If a command is left blank then the corresponding Repository command or WhereScape RED default will be used for that method.
Use \$SCHEMA\$ to denote the table owner and \$TABLE\$ to denote the table name.

Stats:

Quick Stats:

Analyze:

Quick Analyze:

OK Cancel Help

This screen shot shows the default values; but if these fields are left blank, then the default values will be used.

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, i.e. UNIX or Windows script file is executed to perform the load.
- ☐ XML file load - loading an XML file from a Windows directory.

In This Chapter

Choosing the Best Load Method	234
Load Drag and Drop.....	235
Database Link Load	237
ODBC Based Load	248
SSIS Loader	249
Native ODBC Based Load.....	264
Flat File Load	273
XML File load.....	285
External Load	289
Handling Missing Source Columns	290
Load Table Transformations	293
Changing Load Connection and Schema.....	295


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 QAD Data Warehouse Designer. 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.

Note: When creating a load table in QAD Data Warehouse Designer by dragging over a source table, Data Warehouse Designer 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 Transformations** (on page 293)). 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 QAD Data Warehouse Designer.
 3. Data types that cannot be inserted into (eg identity columns on SQL Server)
-

Data Type Mappings

In **Tools/Data Type Mappings**, the first menu option will allow you to maintain the data type mappings.



The user interface for maintaining data type mappings is as follows:



See **Data Type Mappings** (on page 1105) for a detailed explanation of Data Type Mappings.

Database Link Load

Database Link Load - Properties

The fields on the **Database Link Load Properties Screen** (see "**Load Table Properties Screen**" on page 240) are described below:

Load Table Name

The table name is limited by most relational databases to a maximum of 30 characters and must be unique. Table name defaults can be set up 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 QAD Data Warehouse Designer uses the prefix **load_** for load tables.

Unique Short Name

The table short name is limited in size to 22 characters in Oracle and SQL Server and 12 characters in DB2. It must be unique. The short name is used in the naming of indexes, keys, procedures and scripts.

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 295).

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

The database link will display in this field.

Script Name

This field is only active if this is a script based load.

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 single statement then the trailing semi-colon is not required for Oracle. 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
truncate table load_customer
DECLARE v_status number; v_result varchar2(256); BEGIN update_dim_customer(1,'a','b',0,0,v_result, v_status); END;

Note: Not all pre load SQL commands will work in a SQL Server database. These commands are executed as a batch. For example: create trigger as the second command will fail as it is not supported in batch mode unless it is the first command.

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 Oracle, this can also be used to add a parallel clause or compression clause to the table.

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.

Note: At the bottom of the Load Table Properties screen are three fields that display information on that load table:

1. Date table structure last updated
 2. Date created in database
 3. Date last updated in database
-

Load Table Properties Screen

Load load_customer

Properties

Storage

Source

Notes

Load Table Name:

load_customer

Unique Short Name:
(maximum 22 characters)

load_customer

Description:

Connection:

Tutorial(OLTP)

Load Type:

Database link load

Database Link:

Script Name:

(None)

Pre-Load Action:

Truncate

Pre-Load SQL:

Post Load Procedure:

(None)

Timestamps

Metadata Structure Changed:

2013-08-26 00:01:02.800000

Database Created:

2013-08-26 00:13:20.030000

Database Altered:

2013-08-26 00:13:20.030000

OK

Cancel

Help

Database Link Load - Source Mapping

The fields on the **Database Link Load Source Mapping Screen** (see "**Load Table Source Mapping Screen**" on page 243) are described below:

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

The connection that identifies the source database. Can be specified via the Properties page.

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 293) below for a fuller explanation.

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 290)).

Note: This option is only available in Oracle.

Exit Status when Missing Source Columns

Load exit status when source columns are missing. Only enabled when **Allow Missing Source Columns** is set to True.

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.

SQL Insert Hint

Optional hint to include in the SQL insert statement.

DB2 Load Logging

Controls the level of logging during the load (Minimal Logging or Recoverable Load). The default is to reduce logging, which does NOT provide for full recovery.

DB2 Load Modifier

Optional Modifiers for the Load utility.

Source View Name

Name of the database view to be created by QAD Data Warehouse Designer (if necessary) in the source system that will be used to select the source data. See **Remote View Extract** (on page 246) topic.

Note: This option is only available in Oracle.

Load Table Source Mapping Screen

DB2 Source Screen:

Load Table load_customer

Properties
Storage
Source
Statistics
Notes

Load Type: Database link load
Source Connection: Tutorial (OLTP)

General

Select Distinct Values: False
Source Schema: tutorial
Derive Source Table(s) and Source Columns: Derive
Source Table(s): customer
Override Source Column/Transformations: False
Where and Group By Clauses:
Override Load SQL:
DB2 Load Logging: Minimal Logging
DB2 Load Modifier:

Select Distinct Values
Include the DISTINCT keyword in the SQL SELECT statement.

OK Cancel Help

Oracle **Source Screen:**

Load Table load_customer

Properties

Storage

Source

Notes

Load Type

Source Connection

General

Select Distinct Values

Source Schema

Derive Source Table(s) and Source Columns

Source Table(s)

Override Source Column/Transformations

Allow Missing Source Columns

Exit Status when Missing Source Columns

Where and Group By Clauses

Override Load SQL

Source View Name

Database link load

Tutorial (OLTP)

False

tutorial

Derive

customer

False

True

Warning

...

...

...

...

...

...

...

Select Distinct Values

Include the DISTINCT keyword in the SQL SELECT statement.

OK

Cancel

Help

SQL Server **Source Screen:**

The screenshot shows the 'Load load_customer' dialog box. On the left is a sidebar with 'Properties', 'Storage', 'Source' (selected), and 'Notes'. The main area contains a table of properties and a 'Select Distinct Values' section at the bottom.

Load load_customer	
Load Type	Database link load
Source Connection	Tutorial(OLTP)
General	
Select Distinct Values	False
Source Schema	tutorial
Derive Source Table(s) and Source Columns	Derive
Source Table(s)	customer
Override Source Column/Transformations	False
Where and Group By Clauses	
Override Load SQL	

Select Distinct Values
Include the DISTINCT keyword in the SQL SELECT statement.

OK Cancel Help

Remote View Extract

Note: To use Remote View Extracts, both the source database and target data warehouse database must be **Oracle**.

Remote view extracts are used to force a remote Oracle source system database to perform complex extraction logic in the remote source system database and return the results to the data warehouse.

When specifying a join in an extract (load) over a database link the Oracle cost based optimizer decides whether it is better to join the tables on the remote machine, or if the tables should be copied onto the local machine and the join performed locally. If the decision to copy is made for large tables this can result in huge extract times. The simplest method of forcing the join to occur on the remote machine is to create a view on the remote machine from which the extract is done.

The view is re-created every time the load table is run. Then the load table selects all columns from the remote view and inserts the results into the load table over the defined database link.

The steps for creating a remote view load table are:

- 1 Create a connection object for the remote system, populating the following fields:
 - Connection Name
 - Connection Type => **Database**
 - ODBC Source
 - Database id (SID)
 - Database link id
 - Extract user name
 - Extract password
- 2 Ensure the user name and password being used to connect to the remote system have the following privileges:
 - create view (granted as a specific privilege and not part of a role)
 - compile procedure (granted as a specific privilege and not part of a role)
 - execute on sys.dbms_sql
- 3 Right-click on the connection to the remote system and select **Create Database Link**. When the dialog with the database link has been created message is displayed, click **OK**.
- 4 Right-click on the connection to the remote system and select **Create Remote View Procedure**. When the dialog with the Procedure dss_view_create compiled is displayed, click **OK**.
- 5 Browse the connection in the right pane.
- 6 Drag and drop a source table or columns of the source tables to create the required load table.
- 7 Create the load table in the data warehouse database by selecting **Create (Recreate)** from the right-click menu.
- 8 Right-click on the load table and select **Source Mapping**.
- 9 Enter a name for the remote view to be created into the **Remote View Name** box. For example, use the name of the load table.
- 10 Cut or copy any existing 'Where' clause from the **Where, Group etc. clauses** box and paste it into the **Where clause (for Remote view)** box. Alternatively, enter the 'Where' clause (including any table joins, if multiple source tables are being extracted from) into the **Where clause (for Remote view)** box.

Note: Once the load table has a remote view name, the 'Where' clause in the **Where, Group etc. clauses** box is ignored and the 'Where' clause in the **Where clause (for Remote view)** box is used instead.

- 11 Click **OK**.
- 12 Load or schedule the load table.

Example remote view load table source mapping properties:

Source View Name	load_customer
Source View WHERE-clause	WHERE code <= 999

Note: During Load column transformations are ignored with Remote View Extract load tables. After Load column transformations can be used. To create a derived column on a Remote View Extract load table, add the column leaving its source table blank and setting its source column to NULL. Then create an After Load transformation to populate the derived column.

ODBC Based Load

An ODBC based load provides an extensive option for acquiring data from many sources. It will be slower than most other forms of loads, but may still be a viable option depending on the volume of data being loaded.

An ODBC based 'interactive load' when using the QAD Data Warehouse Designer 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 except that the data is loaded into the server that is running the scheduler and then pushed to the data warehouse database. For Oracle, this push to the data warehouse can be performed via Oracle Direct Path loading, which can be considerably faster than the ODBC sql method.

For Oracle Direct Path 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 such as `TO_CHAR(source_date_column, 'YYYY-MM-DD HH24:MI:SS')`

Note: Dates must be in the correct format for Oracle Direct Path loading to occur.

The obvious disadvantage in an ODBC based load is the two network transactions required and the overhead placed on the Scheduler Server. The UNIX scheduler does not support ODBC based loads, therefore a Windows scheduler must be available to handle ODBC based loads.

Note: A windows scheduler must be available to handle ODBC based loads.


The Properties screens for an ODBC based load are the same as those of a database link load. Refer to the previous section for details.

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.


QAD Data Warehouse Designer manages the extraction and load of data into the data warehouse as part of data warehouse processing. QAD Data Warehouse Designer attempts to utilize the optimum loading technology for each supported database platform. Optimal loading performance from a database source into SQL Server is achieved using SSIS. QAD Data Warehouse Designer provides the ability to leverage the extract and load performance available from SSIS as part of the QAD Data Warehouse Designer processing. QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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.

Select the relevant version of SSIS in **Tools/Options/Code Generation/General**.



Loading data into Data Warehouse Designer load tables using SSIS

Note: To use SSIS loading, ensure that SSIS loads are enabled and that the SSIS version is set correctly in **Tools/Options/Code Generation /General**.



As with any load into Data Warehouse Designer a connection to the source data needs to be created to provide extraction details. If the connection is a database load then there is additional connection information that should be supplied to utilize SSIS as a loading option. This additional information needs to be supplied on both the source connection and the data warehouse connection:

Connection Sales SSIS

Properties

Notes

General

Connection Name	Sales SSIS
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	SQLOLEDB
Database ID	
Database Link Name	

Database Credentials

Extract User ID	Sales
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Integration Services load
SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Security Info=False;Initi...
Data Type Mapping Set	(Default)

SSIS Connection String


Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.
NOTE: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited.

OK Cancel Help

The **SSIS Connection String** is a valid SSIS connection string that can be used to connect to the data source or destination. The **Reset** button will attempt to construct a valid connection string from the connection information supplied in the connection details consisting of the Database ID, Database Link ID (Instance name), Provider Name, Extract User details.

Steps to create the SSIS Connection String

- ☐ Click on the **expander** button.



Connection Sales SSIS

Properties


Notes

General	
Connection Name	Sales SSIS
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False
Source System	
Provider Name	SQLOLEDB
Database ID	
Database Link Name	
Database Credentials	
Extract User ID	Sales
Extract User Password	***
Administrator User ID	
Administrator User Password	
Other	
New Table Default Load Type	Integration Services load
SSIS Connection String	
Data Type Mapping Set	(Default)


SSIS Connection String
Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.
NOTE: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited.

OK Cancel Help

- On the Provider tab, select the **OLE DB Provider** and click **Next**.




- On the **Connection** tab, select the **server name**, enter information to log on to the server; and select the **database** on the server. Click **Test Connection**.



The image shows the 'Data Link Properties' dialog box with the 'Connection' tab selected. The dialog has four tabs: 'Provider', 'Connection', 'Advanced', and 'All'. The 'Connection' tab contains the following fields and options:

- Specify the following to connect to SQL Server data:
- 1. Select or enter a server name: A dropdown menu showing 'DOC\SQL2012' and a 'Refresh' button.
- 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 showing 'Sales'.
- ☐ Attach a database file as a database name:
 - [text box]
 - Using the filename: [text box] [button ...]
- [button Test Connection]
- [button OK] [button Cancel] [button Help]

- Click **OK**.



The image shows the 'Microsoft Data Link' dialog box. It has a title bar with a close button. The main area contains an information icon (i) and the text 'Test connection succeeded.' At the bottom right, there is an 'OK' button.

- Click **OK** on the Data Link Properties dialog.

The screenshot shows the 'Data Link Properties' dialog box with the 'Connection' tab selected. The dialog has a title bar with a close button (X) and a tab bar with 'Provider', 'Connection', 'Advanced', and 'All'. The 'Connection' tab contains the following elements:

- Text: 'Specify the following to connect to SQL Server data:'
- Step 1: '1. Select or enter a server name:' followed by a dropdown menu showing 'DOC\SQL2012' and a 'Refresh' button.
- Step 2: '2. Enter information to log on to the server:' followed by two radio buttons: 'Use Windows NT Integrated security' (selected) and 'Use a specific user name and password:'. Below the second radio button are 'User name:' and 'Password:' text boxes, and two checkboxes: 'Blank password' and 'Allow saving password'.
- Step 3: '3. Select the database on the server:' followed by a radio button (selected) and a dropdown menu showing 'Sales'. Below this is another radio button: 'Attach a database file as a database name:' followed by a text box containing 'Sales' and a label 'Using the filename:' with an empty text box and a browse button (...).
- A 'Test Connection' button is located below the database selection options.
- At the bottom of the dialog are 'OK', 'Cancel', and 'Help' buttons.

- The SSIS connection string will now be displayed.

Connection Sales SSIS

Properties

Notes

General

Connection Name	Sales SSIS
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	False

Source System

Provider Name	SQLOLEDB
Database ID	
Database Link Name	

Database Credentials

Extract User ID	Sales
Extract User Password	***
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Integration Services load
SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Security Info=False;Initi...
Data Type Mapping Set	(Default)

SSIS Connection String
Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.
NOTE: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited.

OK Cancel Help

- ☐ Click **OK** to save the connection.
- ☐ Right-click on **Sales SSIS** and select **Browse Connection**.
- ☐ Accept the defaults and click **OK**.

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 that 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.

Connection DataWarehouse

Properties

Notes

General

Connection Name	DataWarehouse
Connection Type	Database
Database Type	((local))
ODBC Data Source Name (DSN)	WslWarehouse
Data Warehouse Connection Indicator	True

Source System

Provider Name	SQLOLEDB
Database ID	WslWarehouse
Database Link Name	

Database Credentials

Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
Data Type Mapping Set	(Default)
Default Transform Function Set	(Default)

When Connection is an OLAP Data Source

Connection Provider/Driver	SQLOLEDB
Data Warehouse Server	DOC\SQL2012
Data Warehouse Database ID	WslWarehouse


SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.
NOTE: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited.

OK Cancel Help

☐ Click **OK** to save your connection.

Note: If the connection string is already set, then the expander will open up an editor dialog.



The image shows a Windows-style dialog box titled "SSIS Connection String". It has a standard title bar with a close button (X) in the top right corner. The main area of the dialog contains a text box with the following text: "Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination. NOTE: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited." Below this text is a large text area where a connection string is entered: "Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Security Info=False;Initial Catalog=WsWarehouse;Data Source=DOC\SQL2012". At the bottom of the dialog, there is a checkbox labeled "Word Wrap" which is currently unchecked. To the right of the checkbox are three buttons: "OK", "Cancel", and "Help".

SSIS Connection String

Connection string to be used by Microsoft SQL Server Integration Services (SSIS) to connect to the data source or destination.
NOTE: A connection string is typically composed of multiple property name/value pairs that are semi-colon delimited.

Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Security Info=False;Initial Catalog=WsWarehouse;Data Source=DOC\SQL2012

☐ Word Wrap

OK Cancel Help

Edit the connection string and click **OK**.

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 to create a load table. (see **Loading data** (on page 233)) 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.

Load Table load_customer_ssis

Properties

Storage

Source

Notes

Load Table Name: load_customer_ssis

Unique Short Name: (maximum 22 characters) load_customer_ssis

Description:

Connection: Sales SSIS

Load Type: Integration Services load

Database Link:

Script Name: (None)

Pre-Load Action: Truncate

Pre-Load SQL:

Post Load Procedure: (None)

Timestamps

Metadata Structure Changed:	Database Created:	Database Altered:
2013-08-22 14:05:56.443	2013-08-22 14:05:59.230	2013-08-22 14:05:59.230

OK Cancel Help

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 Data Warehouse Designer-generated SSIS packages. If no conversion is applied then the exact case of the identifier defined in the Data Warehouse Designer 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 Data Warehouse Designer-generated SSIS packages. If no conversion is applied then the exact case of the identifier defined in the Data Warehouse Designer metadata is used in SSIS.
- ☐ **SSIS Row Count Log** - During an SSIS Load include Row Count logging.
- ☐ **SSIS Acquire Table Lock** - Specify whether the destination table is locked during the load. This allows SQL Server to reduce lock escalation overheads during loading and will promote minimal logging in SQL 2008+ when a Bulk Logged recovery model is used for the database. The default value of this property is true.
- ☐ **SSIS Commit Interval** - Specify the batch size that the OLE DB destination tries to commit during fast load operations. The default value of 2147483647 indicates that all data is committed in a single batch after all rows have been processed. If you provide a value for this property, the destination commits rows in batches that are the smaller of (a) the Maximum insert commit size, or (b) the remaining rows in the buffer that is currently being processed.
- ☐ **Set SSIS Batch Size** - Specify the number of rows in a batch. The default value of this property is empty, which indicates that no value has been assigned.

Load Table load_customer_ssis

Properties
Storage
Source
Notes

Load Type: Integration Services load
Source Connection: Sales SSIS

General

Select Distinct Values: False
Source Schema: dbo
Derive Source Table(s) and Source Columns: Derive
Source Table(s): customer

Override Source Column/Transformations: False
Where and Group By Clauses:
Override Load SQL:

SQL Server Integration Services (SSIS)

SSIS Source-Identifier Encapsulation:
SSIS Source-Identifier Case Conversion: (None)
SSIS Destination-Identifier Case Conversion: (None)
SSIS Row Count Log: True
SSIS Acquire Table Lock: True
SSIS Commit Interval: 2147483647
Set SSIS Batch Size: False


Select Distinct Values
Include the DISTINCT keyword in the SQL SELECT statement.

OK Cancel Help

Troubleshooting


If an SSIS load fails due to incorrect connection strings then the SSIS package will fail to create at run time. The connection string information in both the source connection and the data warehouse connection (the destination) should be checked to ensure that they are correct.

If the SSIS load fails due to incorrect extraction SQL, due to mismatched destination load table structure or due to constraint violation then the SSIS package will be created in the file system for the developer to check. The SSIS package file will be created in the QAD Data Warehouse Designer Work Directory; shown in the dialog opened from **Help/About QAD Data Warehouse Designer** (refer to dialog below). This SSIS package will have a DTSX file extension and can be opened using the appropriate version of Microsoft's BI Development Studio add-in to Visual Studio.



The **Integration Services load** in Data Warehouse Designer does not currently automatically handle spaces in source field names. You need to put [] around the source column on the column Properties of the load table column that contains spaces in Excel. Failure to do so results in SSIS (runtime) crashing and returning an error message. The SSIS package file will be created in the QAD Data Warehouse Designer Work Directory; shown in the dialog opened from **Help/About** QAD Data Warehouse Designer (refer to dialog above).

On the **Source** tab for a load_table, the following drop-down list gives the options for **SSIS Source-Identifier Encapsulation**. We use encapsulation to ensure that column aliases are not changed during an SSIS load:



Load Table load_customer_ssis	
Load Type	Integration Services load
Source Connection	Sales SSIS
General	
Select Distinct Values	False
Source Schema	dbo
Derive Source Table(s) and Source Columns	Derive
Source Table(s)	customer
Override Source Column/Transformations	False
Where and Group By Clauses	
Override Load SQL	
SQL Server Integration Services (SSIS)	
SSIS Source-Identifier Encapsulation	[]
SSIS Source-Identifier Case Conversion	(None)
SSIS Destination-Identifier Case Conversion	""
SSIS Row Count Log	""
SSIS Acquire Table Lock	[]
SSIS Commit Interval	""
Set SSIS Batch Size	""
SSIS Source-Identifier Encapsulation Characters that are used to enclose Source Object Identifiers (such as table, view, and column names).	

Native ODBC Based Load

A Native ODBC based load is similar to an ODBC load, except it provides a faster method to load data from non-linked databases.

A Native ODBC based **interactive load** when using the QAD Data Warehouse Designer 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 database via a Bulk Insert in SQL Server, SQL*LOADER in Oracle or LOAD statement in DB2.

A scheduler load will perform in the same way except that the data is loaded into the server that is running the scheduler and then pushed to the data warehouse database.

All loaders require dates in the default format of the target data warehouse database. 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 the previous section for details. Details of the source mapping screen follow.

Native Loads in Oracle

Some versions of Oracle do not return the ORACLE_SID correctly when an ODBC call requests it. To ensure that you don't have any issues because of this, please ensure that the following steps are followed:

If you are doing interactive Native Loads in Data Warehouse Designer into Oracle:

- 1 Ensure that the **Database id (SID)** field on the Data Warehouse connection Properties screen is set to the actual ORACLE_SID for your database.

This value is used in the sqlldr connection during the native load IF the ORACLE_SID is not returned correctly by the ODBC driver.

If you are using the windows scheduler for Native Loads into Oracle:

- 1 Change the scheduler to be a direct path load (via maintain scheduler in the QAD Administrator).
- 2 Ensure that the TNS Service is set to the actual ORACLE_SID for your database.
- 3 Restart the scheduler.

This value is used in the sqlldr connection IF the SID is not returned by the ODBC driver and only IF the scheduler is set to do direct path loads.

Native ODBC Based Load - Source Mapping

The fields on the **Native ODBC Load Source Mapping Screen** (see "**Load Table Source Mapping Screen (Native)**" on page 267) are described below:

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

The connection that identifies the source database. Can be specified via the Properties page.

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.

Source View Name

Name of the database view to be created by QAD Data Warehouse Designer (if necessary) in the source system that will be used to select the source data. See **Remote View Extract** (on page 246) topic.

Note: This option is only available in Oracle.

SQL Insert Hint

Optional hint to include in the SQL insert 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.

Unicode Extract File

Data will be loaded via a Unicode format file. The default is unchecked and this field is not available on DB2.

Populate Load Parameters

Populate any load-related QAD Data Warehouse Designer parameters, which may be used for validation purposes.

Load Table Source Mapping Screen (Native)

DB2 Source Screen:

Properties

Storage

Source

File Actions

Statistics

Notes

Load Table load_customer

Load Type

Native ODBC

Source Connection

Tutorial (OLTP)

General

Select Distinct Values

False

Source Schema

tutorial

Derive Source Table(s) and Source Columns

Derive

Source Table(s)

customer

Override Source Column/Transformations

False

Where and Group By Clauses

Override Load SQL

Native ODBC Load

Native ODBC Load Routine

Load

Field Delimiter

|

Populate Load Parameters

True

Select Distinct Values

Include the DISTINCT keyword in the SQL SELECT statement.

OK

Cancel

Help

Oracle **Source Screen:**

Properties

Storage

Source

File Actions

Notes

Load Table load_customer

Load Type

Source Connection

General

Select Distinct Values

Source Schema

Derive Source Table(s) and Source Columns

Source Table(s)

Override Source Column/Transformations

Where and Group By Clauses

Override Load SQL

Source View Name

Native ODBC Load

Native ODBC Load Routine

Field Delimiter

Unicode Extract File

Populate Load Parameters

Native ODBC

Tutorial (OLTP)

False

tutorial

Derive

customer

False

SQL*Loader

|

False

True

Select Distinct Values

Include the DISTINCT keyword in the SQL SELECT statement.

OK

Cancel

Help

SQL Server **Source Screen:**

The screenshot shows the 'Load Table load_customer' dialog box. On the left is a sidebar with 'Source' selected. The main area contains a table of settings. The 'General' section is expanded, showing options like 'Select Distinct Values' (False), 'Source Schema' (dbo), and 'Source Table(s)' (customer). The 'Native ODBC Load' section is also expanded, showing 'Native ODBC Load Routine' (BULK INSERT) and 'Populate Load Parameters' (True). At the bottom, there is a 'Select Distinct Values' section with a description: 'Include the DISTINCT keyword in the SQL SELECT statement.' The dialog has 'OK', 'Cancel', and 'Help' buttons at the bottom right.

Property	Value
Load Type	Native ODBC
Source Connection	Tutorial (OLTP)
General	
Select Distinct Values	False
Source Schema	dbo
Derive Source Table(s) and Source Columns	Derive
Source Table(s)	customer
Override Source Column/Transformations	False
Where and Group By Clauses	
Override Load SQL	
SQL Insert Hint	TABLOCK
Native ODBC Load	
Native ODBC Load Routine	BULK INSERT
Field Delimiter	
Unicode Extract File	False
Populate Load Parameters	True

Select Distinct Values
Include the DISTINCT keyword in the SQL SELECT statement.

OK Cancel Help

File Actions

A file action defines a copy, move or delete step that happens to **Native ODBC** temporary files after a load is run. There are six different file action programs:

File Action	Description
copy all files	Copies all temporary files for the load table (including the temporary data file) to another specified directory.
copy data files	Copies the Native ODBC load table's temporary data file to another specified directory. Also allows a new file name to be specified.
delete all files	Deletes all temporary files for the load table (including the temporary data file).
delete data files	Deletes the Native ODBC load table's temporary data file.
move all files	Moves all temporary files for the load table (including the temporary data file) to another specified directory.
move data files	Moves the Native ODBC load table's temporary data file to another specified directory. Also allows a new file name to be specified.

Note: A load table may have more than one file action. File actions are processed in order, based on the action number.

Creating a File Action

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 *After load* from the **Action Type** drop-down list.
- 4 Enter the **Action Description**.
- 5 Choose the **Action Program** from the drop-down list (see above for the options).
- 6 For action programs of copy... and move..., enter the target directory name.
- 7 For copy data files and move data files, optionally enter the file's new name.
- 8 Click on **Save/update action**.
- 9 Repeat if necessary for additional file actions.
- 10 Click **OK**.

File Action Example

The following example moves the Native ODBC Load data file to another directory then deletes all temporary files.

Note: The data file is not deleted as it has been moved prior to the delete action.

Sample **move data file** file action:

Load Table load_customer

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 Add New Action Remove Action Save (Update) Action

Action Type: After load

Action Description: Move the data file to the data file backup directory

Action Program: move data file

Action Connection: Tutorial (OLTP)

Target directory: e:\backup\load_customer

Target file name (optional):

Actions

01 Delete all temporary files

OK Cancel Help

Sample **delete all files** file action:

Properties

Storage

Source

File Actions

Notes

Load Table load_customer

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:

01

Add New Action

Remove Action

Save (Update) Action

Action Type:

After load

Action Description:

Delete all temporary files

Action Program:

delete all files

Action Connection:

Tutorial (OLTP)

Actions

01 Delete all temporary files

OK


Cancel

Help

Flat File Load

Flat file loads can be performed from either a Windows or UNIX connection. As with all other load types it is easier to use the drag and drop functionality to create the load table. The drag and drop process for flat files is as follows:

- 1 Browse to the directory and file via the appropriate **Connections** (on page 107).
- 2 Double-click on 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. The following dialog appears:



The dialog box is titled "Add a New Metadata Object" with a red close button in the top right corner. It contains the following text and controls:

Define the Type and Name of the New Object.

Specific information for each object type is defined in subsequent screens.


Object Type:

Object Name:

At the bottom right, there are two buttons: "ADD" and "Cancel".

Click the **ADD** button.

- 4 The following dialog appears.



The dialog box is titled "Specify the type of Load required" with a red close button in the top right corner. It contains the following text and controls:

A File based load places control and management of the extract process in the Scheduler. It is the best method for most simple data loads.

A single data column file based load results in the majority of the work occurring in the Data Warehouse. The column parsing can be performed in the post load procedure

A column parsed file based load results in the column parsing being performed on the host system as part of the data load process.

File load (single data column)

File load (columns parsed)

A Script based load places control and management of the extract process in the script which can be tailored. It is the best method for more complex data loads.

A single data column script based load results in the majority of the work occurring in the Data Warehouse. The column parsing can be performed in the post load procedure.

A column parsed script based load results in the column parsing being performed on the host system in the defined script as part of the load process.

Script load (single data column)

Script load (columns parsed)

Cancel

There are four options on this screen (buttons at right).

- The first two options result in a **File based load** where the bulk of the load management is handled by the scheduler.
- If you select either of the last two options, then QAD Data Warehouse Designer will generate a host script and the load table will be a **Script based load**. This host script is executed by the scheduler to effect the load.

Single data column

As stated beside the button, 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

QAD Data Warehouse Designer attempts 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. QAD Data Warehouse Designer 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 shot shows the initial file parsed screen.

Data load Wizard

First Rows from the File

product_code,customer_code,budget_quantity,budget_sales_value,budget_date
1002,228,185,409.92,02-JUN-2000
1008,228,80,978.58,02-JUN-2000
1003,227,62,572.42,30-APR-2001
1007,227,98,766.17,30-APR-2001
1004,226,40,218.00,05-NOV-2001
1006,226,40,618.00,05-NOV-2001
1009,225,74,940.24,04-APR-2002
1002,225,74,163.97,04-APR-2002
1006,225,40,618.00,04-APR-2002
1007,225,98,766.17,04-APR-2002
1004,225,74,402.54,04-APR-2002
1003,224,15,134.85,15-NOV-2001
1008,224,15,177.34,15-NOV-2001
1001,224,15,159.50,15-NOV-2001
1001,223,74,812.46,13-AUG-2000
1009,223,29,369.17,13-AUG-2000
1007,223,98,633.53,13-AUG-2000

Column Delimiter: No Column delimiter will initiate width based parsing
CHAR(nn) inserts an ASCII character (e.g. CHAR(9) = tab) Decimal Code

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

OK Cancel

The **Decimal Code** button will show the decimal value of each character in the lines retrieved from the source file. These decimal codes will be shown below each line and are green. Once the screen above is completed a screen will appear to allow the breakdown of the source data into columns. If no delimiter is entered then width based parsing is assumed and an addition width size is prompted for. Use the **Back** button to revert to the previous column if an incorrect width or delimiter is entered. The following screen is an example of the file parsing technique.

Data load Wizard - Column Definition

Column Data:	File
product_code	product_code,customer_code,budget_quantity,budget_sales_value,budget_date
1002	1002,228,185,409.92,02-JUN-2000
1008	1008,228,80,978.58,02-JUN-2000
1003	1003,227,62,572.42,30-APR-2001
1007	1007,227,98,766.17,30-APR-2001
1004	1004,226,40,218.00,05-NOV-2001
1006	1006,226,40,618.00,05-NOV-2001
1009	1009,225,74,940.24,04-APR-2002
1002	1002,225,74,163.97,04-APR-2002
1006	1006,225,40,618.00,04-APR-2002
1007	1007,225,98,766.17,04-APR-2002
1004	1004,225,74,402.54,04-APR-2002
1003	1003,224,15,134.85,15-NOV-2001
1008	1008,224,15,177.34,15-NOV-2001

Display decimal character values

Column Name:

Business Display Name:

Data Type: ☒ Nulls

Conversion:

Business Definition:

Back Add Cancel

Conversion

During the parsing of the columns an Oracle SQL*Loader or SQL Server BULK Insert conversion string can be used and is required in the case of dates for Oracle. Refer to the Oracle Sql loader manual for syntax. Any Oracle function can be used if enclosed in quotes and with the column name prefixed by a colon. For example: when loading a column called 'product_name', use the following syntax to bring over only the first 30 characters:

```
product_name "substr(:product_name,1,30)"
```

A special variable %FILE_NAME% can be used in an Oracle File based load. This will be substituted with the actual source file name.

To load the full file name (including the path) into a field, enter this transformation:

```
"RTRIM( '%FILE_NAME%' )"
```

To remove the path from the full file name on a UNIX or Linux system, enter this transformation:

```
"SUBSTR( '%FILE_NAME%', INSTR( '%FILE_NAME%', '/', -1, 1)+1) "
```

Note: If running from the windows scheduler then a wildcard must be used in the file name if you wish to use the %FILE_NAME% variable. If a wild card is not used then the literal file name appears enclosed in double quotes. When a wild card is used Windows looks up the file and returns the actual file name with no enclosing double quotes.

Understanding SQL*Loader

When using either type of loading method SQL*Loader is called to perform the actual load. There is therefore a requirement that the executable called 'sqlldr' be in the path.

Under UNIX this means that the UNIX user must be able to execute 'sqlldr' from the command prompt under a normal log on.

Under Windows 'sqlldr' must execute when the command is issued from a DOS prompt.

In both the Windows and UNIX environments sqlldr returns a result code indicating if the load worked fully, partially or failed. QAD Data Warehouse Designer makes use of this returned code to report on the results of the load.

File Load - Source Mapping

The fields for the **Flat File Load Source Mapping Screen** (see "**Load Table Source Mapping Screen**" on page 280) are described below:



TIP: If the file has been **dragged and dropped** (see "**Flat File Load**" on page 273) into the load table (middle pane) then some of the fields on this tab are automatically populated.

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.

Source File Path

The full path (absolute path) of the folder/directory containing the source File is located 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 File Field Enclosure Delimiter

Oracle only. Optional character that delimits BOTH the start & end of a field value i.e. encapsulates a value. A double-quote(") is a common enclosure delimiter. QAD Data Warehouse Designer allows the specification of one enclosure delimiter that is used for both the start & end delimiters. Everything between the start & end enclosure delimiters is considered part of the field value. Enclosing a field's value is OPTIONAL but it must either be enclosed by BOTH a start and an end enclosure delimiter or BOTH of them must be absent. A field value can contain the enclosure delimiter character as embedded data if each embedded instance is "escaped" by the same character.

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 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 QAD Data Warehouse Designer 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'.

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. 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.

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 QAD Data Warehouse Designer 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.

BULK INSERT Options

SQL Server only. Optional comma-delimited list of options to control the behavior of the SQL Server BULK INSERT command. For example the following options will respectively lock the table; set the batch size to 100,000 rows; and fail the load when the first error occurs: TABLOCK, BATCHSIZE= 100000, MAXERRORS=0.

SQL*Loader Options

Oracle only. Optional Oracle SQL*Loader OPTIONS clause and/or WHEN clause. Multiple options can be specified via the OPTIONS clause by separating each option by a comma. The primary intent of this property is to allow the specification of the direct path load method and/or a WHEN-clause to selectively load rows. Example: OPTIONS(DIRECT=TRUE) specifies a direct-path load via an OPTIONS clause.

Note: There are some limitations on the **When statement**. At the time of writing the not symbol (!) will not work for a Windows based load but will work under UNIX.

DB2 Load Command Parameters

DB2 only. Optional Command Parameter Names/Keywords to control the behavior of the DB2 LOAD command/utility. For example, NONRECOVERABLE specifies that the load transaction is to be marked as un-recoverable and it will not be possible to recover it by a subsequent roll-forward action.

DB2 Load Command File Type Modifiers

DB2 only. Optional File Type Modifiers to control the behavior of the DB2 LOAD command/utility. For example, `TIMESTAMPFORMAT="YYYY-MM-DD HH:MM:SS.UUUUUU"`

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 `$YYYYMMDD$`.

Load Table Source Mapping Screen

SQL Server **Source Screen:**

Load Table load_budget

Properties

Storage

Source

Notes

Load Type

File load

Source Connection

Windows

Source File Details

Source File Path

C:\Program Files (x86)\WhereScape_67101\Tutorial\

Source File Name

budget.txt

Source File Field Delimiter

,

Source File Record Terminator

Source File has Field Headings/Labels

True

Trigger File Details

Trigger File Path

Trigger File Name

Trigger File Field Delimiter

Trigger Parameter Name Prefix

Load Configuration

Wait for Source File or Trigger File

False

Wait Limit (in seconds)

1

Action when Wait Limit Reached

Warning

BULK INSERT Options

TABLOCK

Archived File Details

Compress Source File when Archive

False

Archived Source File Path

Archived Source File Name

Archived Trigger File Path

Archived Trigger File Name

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

Oracle **Source Screen:**

Properties

Storage

Source

Notes

Load Table load_budget

Load Type

File load

Source Connection

Windows

Source File Details

Source File Path

Source File Name

Source File Field Delimiter

Source File Field Enclosure Delimiter

Source File Record Terminator

Source File has Field Headings/Labels

Trigger File Details

Trigger File Path

Trigger File Name

Trigger File Field Delimiter

Trigger Parameter Name Prefix

Load Configuration

Wait for Source File or Trigger File

Wait Limit (in seconds)

Action when Wait Limit Reached

SQL*Loader Options

Archived File Details

Compress Source File when Archive

Archived Source File Path

Archived Source File Name

Archived Trigger File Path

Archived Trigger File Name

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

DB2 Source Screen:

Load Table load_budget

Properties
Storage
Source
Statistics
Notes

Load Type	Script based load
Source Connection	Windows
Source File Details	
Source File Path	C:\Program Files (x86)\WhereScape_67102\Tutorial\
Source File Name	budget.txt
Source File Field Delimiter	,
Source File Record Terminator	
Source File has Field Headings/Labels	True
Trigger File Details	
Trigger File Path	
Trigger File Name	
Trigger File Field Delimiter	
Trigger Parameter Name Prefix	
Load Configuration	
Wait for Source File or Trigger File	False
Wait Limit (in seconds)	1
Action when Wait Limit Reached	Warning
DB2 Load Command Parameters	NONRECOVERABLE
DB2 Load Command File Type Modifiers	timestampformat="DD-MMM-YYYY"
Archived File Details	
Compress Source File when Archive	False
Archived Source File Path	
Archived Source File Name	
Archived Trigger File Path	
Archived Trigger File Name	

Trigger File Details
Optional Trigger File identification and definition information. If a Trigger File is specified the Sourcm_pstLoad->bWaite File will not be loaded until the Trigger File is available.

OK Cancel Help

Host Script Load

A script based load table will 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 UNIX or Windows file a script can be generated by selecting one of the 'Script based' load options. This script can then be edited to more fully meet any requirements.

Note: If the UNIX Scheduler is not being used and a UNIX script is generated and executed interactively the following two variables will be setup as part of the environment.

DSS_USER=username

DSS_PWD=password

Where 'username' and 'password' are the entries for the data warehouse database.

There are a number of conventions that must be followed if these host scripts are to be used by the QAD 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 QAD Data Warehouse Designer tool under the scheduler window.

File Wild Card Load

The loading of multiple files in one pass is best undertaken as a host script.

A generated host script contains the code to handle wild card loads and the subsequent loading and renaming of multiple files. The process for setting up wild card loads is as follows:

- 1 Drag one of the files into a **Load Table** target and select one of the **Script** based options.
- 2 If necessary edit the generated script and replace the load file name with the required wild card version. (For example test_result_234.txt renamed to test_result_*.txt. This load file name is defined as a variable at the start of the script.
- 3 If renaming or moving to another directory is required, then you may need to uncomment the rename phase and enter the desired directory or rename syntax.
- 4 If only one file at a time is to be processed, but still based on a wild card file name, then you may need to uncomment the break statement at the bottom of the script. This will result in the loading of only the first file encountered that matches the wild card criteria.

In every event check the generated script as it will contain instructions on how to handle wild cards for the appropriate script type (Windows or UNIX).

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.

To load an XML file located in a Windows directory proceed as follows:

- ☐ Create a connection to the Windows system.
- ☐ Browse to the connection and locate the XML file.
- ☐ Make Load tables the middle pane drop target by double-clicking on the **Load Table** object group in the left pane.
- ☐ Drag the XML file into the middle pane.

Supported XML Formats

QAD Data Warehouse Designer 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.qad.com/wsl-schema"
      xmlns:xsi="http://www.w3.org/2000/10/XMLSchema-instance"
      xsi:schemaLocation="http://www.qad.com/load_table.xsd">
```

or no namespace;e.g.

```
<root xmlns="http://www.qad.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	SQL Server	Oracle	IBM DB2
string with length	char()	char()	char()
string with maxlength	varchar()	varchar2()	varchar()
integer	integer	integer	integer
decimal with precision and scale	numeric(x,y)	number(x,y)	decimal(x,y)
dateTime (ISO8601)	datetime	date	timestamp
i2	integer	integer	integer

XSD	SQL Server	Oracle	IBM DB2
i4	integer	integer	integer
r4	float	number	float
r8	float	number	float
float	float	number	float

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. i.e.

```
<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 SQL Server varchar(8) or Oracle varchar2(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 SQL Server column of numeric(6) or an Oracle column of number(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 SQL Server column of numeric(8,2) or an Oracle column of number(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.qad.com/wsl-schema"
  xmlns:xsd="http://www.qad.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.

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. eg

```

<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>

```

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

QAD Data Warehouse Designer will load this string into Oracle as:

TO_DATE('20031003100215', 'YYYYMMDDHH24MI')

or for SQL Server:

CONVERT(DATETIME, '2003-10-03T10:02:15.310', 126)

or for DB2:

TIMESTAMP_FORMAT('20031003100215', 'YYYYMMDDHH24MISS')

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 either data type text for SQL Server or varchar2(4000) for Oracle. 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.

Handling Missing Source Columns

Note: This option is only available when running an **Oracle** data warehouse.

By default a load will fail if a source column that is to be loaded does not exist. This default action can be modified by using the **Allow Missing Source Columns** feature of a load table. When this feature is enabled the load process will check all source columns, and if any are found to be missing will replace those missing columns with a Null value (which can be changed, see below).

On the **Source Mapping** tab of a load tables Properties there are two checkboxes and a drop-down list that we will cover 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 box must be unselected to enable **Allow Missing Source Columns** support.

Allow Missing Source Columns

When this checkbox is checked 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 provided that 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** checkbox is checked.

Limitations of Missing Column check

The check for missing columns will only examine 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 will not be 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 will still succeed.

Load Table Transformations

Each load table column can have a transformation associated with it. See ***Transformations*** (on page 603) and ***Load Table Column Transformations*** (on page 608) for more details.

Post-Load Procedures

If a procedure name is entered in the post-load procedure field of a load table's Properties, then this procedure will be executed after the load has completed and after any **after transformations** have occurred.

See *Load Table Column Transformations* (on page 608) for more 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 will display 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.

Switching Connection from ODBC to Database

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 will still occur although they may no longer be necessary.

Switching Connection from Database to ODBC

Although it is possible to switch from a Database connection to an ODBC connection, the resultant extract may not function correctly. The main issue is around the conversion of date fields under Oracle. SQL Server tends to handle the date conversion better in most cases. When an extract occurs over an ODBC connection the date is converted in the extract from the source system into a standard ASCII format. This ASCII formatted date is then converted back when it is loaded into the load table. To resolve this problem place a transformation on any date fields. Examples of typical Oracle transformations are:

```
to_date("'" || to_char(ORDER_DATE,'YYYYMMDD') || "'", 'YYYYMMDD')
```

```
to_char(ORDER_DATE,'YYYY-MM-DD HH24:MI:SS')
```

where ORDER_DATE is the name of the column being extracted. In the first example the conversion string is enclosed within quote marks (') so that it is passed as a literal to the select statement that will be built. Quote marks within the conversion string are protected from premature evaluation by the inclusion of a second quote mark alongside the original.

Note: If a load table is created via drag and drop from an ODBC based connection QAD Data Warehouse Designer will build all the required date transformations.

There are several supplied APIs for changing schema and connections programmatically. See **Callable Routines** (on page 903) for more details.

Chapter 10

Dimensions

In This Chapter

Dimensions Overview	298
Building a Dimension	300
Generating the Dimension Update Procedure	308
Dimension Artificial Keys	323
Dimension Get Key Function	324
Dimension Initial Build	326
Dimension Column Properties.....	327
Dimension Column Transformations	335
Dimension Hierarchies	336
Snowflake	340
Dimension Language Mapping.....	343

Dimensions Overview

A dimension table is normally defined, for our purposes, as a table that allows 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*** (on page 300).
- ☐ 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*** (on page 300).
- ☐ Add columns from other load tables if required. See ***Building a Dimension*** (on page 300).
- ☐ Create the dimension table in the database. See ***Building a Dimension*** (on page 300).
- ☐ Build the update procedure. See ***Generating the Dimension Update Procedure*** (on page 308).
- ☐ Run the update procedure and analyze the results. See ***Dimension Initial Build*** (on page 326).

Modify the update procedure as required. See ***Dimension Initial Build*** (on page 326).

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 (i.e. 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:

SQL Server: Where isnull(business_key,'N/A') = isnull(v_LoadRec.business_key,'N/A')

Oracle: Where nvl(business_key,'N/A') = nvl(v_LoadRec.business_key,'N/A')

DB2: 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


Create a dimension target by double-clicking on the **Dimension group** in the left pane. The middle pane will display a list of all existing dimensions. When such a list is displayed in the middle pane, the pane is identified as a target for new dimension tables.

Browse to the Data Warehouse via the Browse/Source Data menu option. Drag the load table, that contains the bulk of the dimensional 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 and will provide a default name based on the load table name. Either accept this name or enter the name of the dimension and click **OK** to proceed.

Dimension Type

A dialog will appear 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.
- ☐ The second choice is a **slowly changing** dimension where new dimension records are created when certain identified columns in the dimension change.
- ☐ The third choice is a **Previous values** dimension, which allows the storing of the last value 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 from the tutorial 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 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. 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 Properties screen will appear. Change the storage options if desired.

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.

Previous Values Dimension Type

If you selected a 'Previous values' dimension type then the following questions will be asked. If one of the other two types were chosen then proceed directly to **Create and Load**.

The following dialog will appear:

Define Previous Value Retained Dimension Columns

Column List:

- customer_category_code
- customer_code
- customer_group_code
- customer_subgroup_code
- sales_source

Select any columns to be managed by retaining a previous value in a new column.

Move any columns over to the previous value retained list, or leave the list empty

Previous Value Retained Column List:

- customer_name

>

<

OK Cancel

It requests the definition of each column that will be managed by storing the previous value in a separate additional column. Note that the business key cannot be used in this manner, as a change to the business key will result in a new record.

Select the columns to be managed and click **OK**. The following dialog will then appear for each column that has been selected. Either accept the name of the additional column or enter a new name.

Enter previous value column name

customer_name


The following column name will be used as the previous value for the column 'customer_name'. Change the name if required.

prev_customer_name

OK Cancel

Create and Load


The following dialog appears after the Properties screen and 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 additional columns you can select the **Create and Load** button. In such an event it would be normal practice to select the **(Build Procedure...)** option from the 'Update Procedure' drop-down box in the previous Properties page. If no procedure has been selected then select the **Create** button as no update will occur in any event.

If **Create** or **Create and Load** is selected and a new procedure creation was chosen proceed directly to the **Generating the Dimension Update Procedure** (on page 308) 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 dimension 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 choosing 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 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 a simple matter therefore to select columns from other 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.

dim_product Columns			
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 QAD Data Warehouse Designer, 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 (Oracle) or identity (SQL Server) 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.

Having performed these actions QAD Data Warehouse Designer 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 we need to physically create the table in the database. This is achieved by right-clicking on the dimension name and choosing **Create (ReCreate)** from the popup 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 dimension table was created. A copy of the actual database create statement, the sequence create statement (Oracle only), 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 to find such a column is to double-click on the list heading **Col name**; which will sort the column names into alphabetical order. Another double-click on the heading will sort the columns back into their create order.

For Oracle data warehouses the sequence created as part of the table create is normally called *table_name_seq* unless the default has been overridden. If this sequence already exists a warning is issued. This sequence is used in the assigning of a unique artificial key to the dimension. SQL Server makes use of an identity attribute on the artificial key column.

The next section covers the **Generating the Dimension Update Procedure** (on page 308).

Generating the Dimension Update Procedure

Once a dimension 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 dimension records.


The zero key row

QAD Data Warehouse Designer will by default always insert a record into the dimension with an artificial key value of zero. 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

To generate a procedure, right-click on the dimension name to edit the properties for the dimension. From the **Update Procedure** drop-down list select **(Build Procedure...)**. Click **OK** to update the properties and start the process of generating the new procedure.

A **Procedure Build Type** dialog will appear, requesting you to choose between Cursor and Set.



If **Set** is chosen then another dialog will appear, allowing you to select the relevant build options.



Define Dimension Procedure Type and Options

Procedure Types

Procedure Type: **Set** Previous build was : Set

Build Options

☐ Process in Batch

☐ Distinct Data Select

☐ Include Delete Before Insert

☐ Include Initial Load Insert

☒ Include Update Statement

☒ Include Insert Statement

☒ Use Merge Statement for Update and Insert

☒ Insert Zero Key Record

☐ Changed Rows Only

Update Hint: (e.g. TABLOCK)

☒ New Rows Only **Except**

Insert Hint: (e.g. TABLOCK)

☐ Changed and New Rows Only

☐ Delete Rows Not in Source

Merge Hint: (e.g. TABLOCK)

OK Cancel

Several different questions need to be answered 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.

Procedure Types

A number of different types of procedure generation are available. The procedure type is chosen from the procedure type drop-down list. A **set** based procedure performs one SQL statement to join all the source tables together and then insert this data into the Dimension Object. This is normally the fastest method of populating a table.


Build Options

The following build options are available:

Process in Batch

The **Process in Batch** checkbox 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).

The following dialog is the default for a non-date column:




The image shows a Windows-style dialog box titled "Batch Processing". It has a light blue title bar with a red close button in the top right corner. The main area is light gray and contains the following text:

Select 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 then choose an attribute of that related table to group by.

Below the text is a white text box with a blue border and a small downward arrow on the right. The text "customer_code" is entered in the box.

At the bottom right of the dialog are two buttons: "OK" and "Cancel".

The following dialog is displayed for date fields:



The image shows a 'Batch Processing' dialog box with a light blue border and a red close button in the top right corner. The dialog contains the following elements:

- Title Bar:** 'Batch Processing' with a red 'X' icon.
- Instructions:**

Select 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 then choose an attribute of that related table to group by.
- Field Selection:** A dropdown menu showing 'created_datetime'.
- Date Format Selection:** A dropdown menu with the following options:
 - created_datetime
 - Day
 - Month (highlighted in blue)
 - Year
- Buttons:** 'OK' and 'Cancel' buttons at the bottom right.

Distinct Data Select

The **Distinct Data Select** option ensures duplicate rows are not added to the Dimension 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 (i.e. duplicates are not removed).

Include Delete Before Insert

The **Include 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 selected, the following dialog is displayed:

Include Initial Load Insert

The **Include Initial Load Insert** option adds an additional insert statement to the update procedure that runs if the target Dimension 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).

Insert Zero Key Record

This option applies only to datastore and normalized tables; and set based dimensions. It allows you to include the zerokey (unknown record) for tables with an artificial key.

Include Update Statement with/without Minus Select using except/checksum

The **Include Update Statement** option includes an update statement in the procedure to update changing rows in the Dimension 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 Dimension 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).

Include Insert Statement with/without Minus Select using except/checksum

The **Include Insert Statement** option includes an insert statement in the procedure to insert new rows in the Dimension 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 Hint.

Enter a database hint to be used in the UPDATE statement.

Note: SQL Server only.

Insert Hint.

Enter a database hint to be used in the INSERT statement.

Note: SQL Server only.

Use Merge Statement for Update and Insert

The **Use Merge Statement for Update and Insert** option allows you to use the Merge statement instead of two separate Update and Insert statements.

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 *except* change detection or *checksum* change detection to work out what rows in the Dimension Object are new or require updating

Delete Rows Not in Source

The **Delete Rows Not in Source** option allows you to delete all the rows which are not in the source but in the target table as part of the Merge statement. To avoid any accidental deletions, this option is de-activated if **Process in Batch** or **Changed and New Rows only** are selected.

Merge Hint


Enter a database hint to be used in the MERGE statement.

Note: SQL Server only.

A series of questions will be asked during the procedure generation based on the type of load information. If only one source table was used then proceed directly to the **Business Key** section and skip the subsequent sections on Joining and Lookups.

Join Or Lookup

If columns from multiple source tables were used to create the dimension then the following dialog will appear at the start of the procedure creation process.



The dialog box is titled "Define Primary Dimension Source Table" with a red close button in the top right corner. The main text area contains the following information:

- Multiple source tables have been detected.
- For dimensions either a multi table join procedure can be created or a procedure with a series of lookups on the secondary tables can be created.
- If a lookup model is chosen then the primary load table must be selected below.

Below the text, there is a label "Primary" followed by a dropdown menu. The dropdown menu currently displays "load_product2". At the bottom of the dialog, there are three buttons: "Lookup", "Join", and "Cancel".

The choice provided is either to create a join that will combine all the information from the various tables into one large select statement, or to do a series of lookups.


The normal process would be to join the tables to create one large select statement. There are however situations where this is not desirable or possible. When a large number of tables are involved then a join will be slower than a series of lookups. In some situations we may not have all the information to do an initial join. A series of lookups may be required to build up the information.

For complex situations it will be necessary to edit and enhance the generated procedure code.

If Lookup is chosen then the primary table for the dimension must be chosen. The philosophy here is that a series of lookups will be conducted against the secondary tables using information from this primary table.

Business Key definition

A dialog will appear asking for the business key that will uniquely identify each dimension record. The source table from which the dimension 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 product code is selected as the business key.



The dialog box is titled "Define Dimension Business Key Columns". It features a "Column List" on the left containing the following items: description, group_description, line_description, prod_group, prod_line, subgroup, and subgroup_description. To the right of this list are two buttons: ">" and "<". Further right is a "Business Key List" containing the item "code". Above this list is the instruction: "Select the business keys that uniquely identify each record in the Dimension table. Move them over to the business key list". Below the "Business Key List" are three checkboxes: "Allow Where Clause Editing", "Not Exists Select", and "Add a Hierarchy for Cubes". At the bottom right are "OK" and "Cancel" buttons.

Column List	Business Key List
description	code
group_description	
line_description	
prod_group	
prod_line	
subgroup	
subgroup_description	

☐ Allow Where Clause Editing
☐ Not Exists Select
☐ Add a Hierarchy for Cubes

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 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.

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 Dimension Select** checkbox (available in SQL Server as **Not Exists Select**) will add additional code to the main select statement. This code will exclude every row in the source table that has not changed in the dimension table and can result in faster execution of the update procedure for large dimensions.

The **Add a hierarchy for Cubes** checkbox will produce another dialog that allows the definition of a simple hierarchy structure. This is useful if the hierarchy of the dimension is known and the dimension will be used in the creation of Analysis services cubes. Hierarchies may be defined now or later when required. They may also be defined or modified after the dimension has been created.

Slowly Changing columns

If the dimension was defined as a slowly changing dimension then the following additional dialogs will appear. The first dialog requests the selection of the columns to be managed as slowly changing dimension columns. Select the required columns and click **OK** to proceed. In the example below the customer name is to be managed as a slowly changing column in a customer dimension.

Define Dimension Change Detection Fields

Column List:

- customer_category_code
- customer_code
- customer_group_code
- customer_subgroup_code
- dim_customer_key
- dss_current_flag
- dss_end_date
- dss_start_date
- dss_version
- sales_source

Select the fields that change detection is required for in the Dimension Update query.
Move them over to the change detection field list

Change Detection Field List:


- customer_name

> <

OK Cancel

The advantages and disadvantages of slowly changing dimensions are discussed earlier in this chapter, but as a general rule try to minimize their usage, as they invariably complicate the processing and end user queries.


The following dialog will appear if one or more columns are chosen for management.



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 IsNull() (SQL Server), Nvl() (Oracle) or Coalesce() (DB2) transformation when loading the column. If, however, this situation is unavoidable then answering **Yes** to this question will result in additional code to test for Nulls during the update of the dimension records. If **No** is answered and Nulls are encountered then a unique constraint violation will occur during the dimension update.

Joining multiple source tables

If multiple source tables were used to build the dimension then an option would have appeared in an earlier dialog to choose between a join or a lookup. If **Join** was chosen then a dialog box will appear, prompting for the joins.



The dialog box is titled "Provide Cursor Mapping" and has a close button (X) in the top right corner. It contains the following elements:

- Cursor Name (if applicable):** A text box containing the word "Load".
- Define the Joins (or edit the where clause):** A section with a help icon (i) and a description: "To define a Join select two tables and select the join type. Then select the join columns from the column lists presented."
- Source Tables:** A list box containing four items: "load_product2", "load_prod_line2", "load_prod_group2", and "load_prod_subgroup2". The first three items are highlighted in blue.
- From and Where Clause:** A text area containing the following SQL code:

```
FROM load_product2
JOIN load_prod_line2
ON load_product2.prod_line = load_prod_line2.prod_line
JOIN load_prod_group2
ON load_product2.prod_group = load_prod_group2.prod_group
LEFT OUTER JOIN load_prod_subgroup2
ON load_product2.prod_group = load_prod_subgroup2.prod_group
AND load_product2.subgroup = load_prod_subgroup2.subgroup
```
- Join Options:** Three buttons: "Outer Join", "Simple Join", and a checked checkbox labeled "ANSI join code generated".
- Select the columns that join the two tables. Select the column from the Master Table first.** A section with two dropdown menus. The first dropdown menu is set to "load_product2" and the second is set to "load_prod_subgroup2".
- Word Wrap:** A checkbox labeled "Word Wrap Displayed Code" which is currently unchecked.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

Two tables are selected in the left box and one of the join buttons clicked. 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_subgroup tables are joined by two columns namely prod_group and subgroup. 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 a 'Where' clause in Oracle. In SQL Server either a where or from clause join can be generated. In DB2, only a 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

An outer join joins the two tables via a 'Where' clause in Oracle. In SQL Server either a 'Where' or from clause join can be generated. In DB2, only a from clause join can be generated. The ANSI standard **From clause** join is recommended as multiple 'Where' clause outer joins will return incorrect results under SQL Server. The outer join 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. When QAD Data Warehouse Designer builds up a 'Where' clause join it must place the outer join indicator next to the appropriate column. As this indicator goes on the subordinate table in Oracle, and the master in SQL Server, QAD Data Warehouse Designer 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. For example the Oracle outer join indicator (+) could be removed if required.

Once satisfied with the join clause click **OK** to proceed to the next step. For Oracle this clause is the 'Where' clause that will be applied to the select statement of the cursor to allow the joining of the various tables. For SQL Server this clause will be either the 'Where' clause or a combined from and 'Where' clause depending on the option chosen. This clause can of course be edited in the procedure that is generated if not correct.

For SQL Server data warehouses you have the choice between 'Where' statement joins and ANSI standard joins. For DB2 data warehouses, only ANSI standard joins are available.

Note: 'Where' joins will not work correctly in SQL Server if more than two tables are used in an outer join.

The example below shows the result of a SQL Server ANSI standard join which takes place in the 'From' statement.

Provide Cursor Mapping

✕

Cursor Name (if applicable):

?

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_product2
load_prod_line2
load_prod_group2
load_prod_subgroup2

From and Where Clause:

```

FROM load_product2
LEFT OUTER JOIN load_prod_line2
ON load_product2.prod_line = load_prod_line2.prod_line
LEFT OUTER JOIN load_prod_group2
ON load_product2.prod_group = load_prod_group2.prod_group
LEFT OUTER JOIN load_prod_subgroup2
ON load_product2.prod_group = load_prod_subgroup2.prod_group
AND load_product2.subgroup = load_prod_subgroup2.subgroup

```

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_product2

▼

load_prod_subgroup2

▼

☐ Word Wrap Displayed Code


OK

Cancel

Lookup

If multiple source tables were used to build the dimension then an option would have appeared in an earlier dialog to choose between a join or a lookup. If **Lookup** was chosen, then a dialog box will appear prompting for the joins between the primary table and each secondary (lookup) table.

The example below shows the lookup dialog for two tables from the tutorial data set.




The dialog box is titled "Dimension Lookup Table Definition" and contains the following elements:

- Instructions: "Create the condition that defines the lookup from our primary table to the lookup table. Select columns from each table that are used in the join and add them to the where clause."
- Text: "Lookup for group_description"
- Checkboxes: ☒ "Multiple Lookups Required on This Table" and ☐ "Post Update"
- Dropdowns: "load_product" with "product_group_code" selected, and "load_prod_group2" with "prod_group" selected.
- Buttons: "Add Join" and "OK" (highlighted with a blue border).
- Text Area: "Where Clause for the Lookup:" with the content "WHERE prod_group = @v_load_product_group_code".
- Buttons: "Cancel" and "OK" (highlighted with a blue border).

The load_product table is the main table and a lookup is being performed on the load_prod_line table. Columns are selected from the drop-down column list for each table and then the **Add Join** button clicked to add these two columns to the 'Where' clause as shown in the bottom box. This 'Where' clause may be edited and goes to form a **lookup** select statement within the main cursor loop of the update procedure. In this example the 'Where' clause has had the second line manually added to identify the type of lookup in the code table.

Multiple lookups can occur on the same table. Where we have a generic code table, for example, we may do multiple lookups to get descriptions. The checkbox **Multiple Lookups Required on This Table** will result in repeated lookups against the second table until all columns sourced from this table have been utilized in the lookup statements. If the **Multiple Lookup** checkbox is ticked then the following dialog will appear to allow selection of the columns that are to be retrieved by the lookup.




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 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:



Results		
	Object	Message
	dim_product	Create indexes on Dimension dbo.dim_product completed successfully.
	dim_product	ALTER TABLE dbo.dim_product ADD CONSTRAINT dim_product_idx_0 PRIMARY KEY CLUSTERED (dim_product_key) WITH (SORT_IN_TEMPDB = OFF);
	dim_product	CREATE UNIQUE NONCLUSTERED INDEX dim_product_idx_A ON dbo.dim_product (code,dss_current_flag,dss_version) WITH (SORT_IN_TEMPDB = OFF);
	dim_product	CREATE UNIQUE NONCLUSTERED INDEX dim_product_idx_SC ON dbo.dim_product (code,description,dss_current_flag,dss_version) WITH (SORT_IN_TEMPDB = OFF);
	dim_product	Procedure Completed
	dim_product	1 dim_product updated. 9 records added. 0 records updated.

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 a sequence in Oracle and an identity column in SQL Server and DB2. This artificial key normally, and by default, starts at one and progresses as far as is required.

A QAD 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 QAD 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 QAD Data Warehouse Designer to illustrate the generation of these pseudo records in the dimension table.

Dimension Get Key Function

When QAD Data Warehouse Designer is asked to generate a new procedure to support the updating of a dimension it also generates a **Get Key** function (procedure in SQL Server and DB2). Such a function is also generated when dimension views are created, and can be generated for retro-fitted dimensions.

This **Get Key** function is used to return an artificial key when supplied a business key. The normal syntax is to call the function passing the business key and be returned a status indicating the result of the lookup. On a successful lookup the artificial key for the dimension record is also returned. If the lookup fails because the business key is not found in the dimension then an appropriate status is returned and the artificial key is set to 0.

Note: Two parameters exist on the generated **Get Key** function that allow manipulation of the process. It is possible to automatically add a new dimension record when a match is not made. It is also possible to record the lookup failure in the detail/error log. By default both these options are disabled.

Modifying the 'Get Key' function

There may be a situation where a dimension record can have multiple business keys.

For example: You may be able to look up a 'factory' based on the factory code, the name of the factory or its address. The generated **Get Key** function will only support the primary business key which in this case may be the factory code. However, we may have a fact record that only contains the factory name. In such a case we need to modify the **Get Key** function or add a second function. The advantages of adding a second function are that any code generated in the future that uses this dimension will work without modification. The disadvantages are the multiple points of maintenance required when changing the dimension. Both options are discussed below:

Adding a second 'Get Key' function for a dimension

Choose a name similar to the existing function. If say we have a 'get_dim_factory_key' function then we may choose a name like 'get_factory_name_key'. Proceed to create and populate the function as follows:

- 1 Right-click on the **Procedure** group in the left pane and select **New Object** to create a new procedure. Enter the name and click **ADD**.
- 2 A Properties dialog will appear. Change the type to **Function** for Oracle, leave as a **Procedure** for SQL Server. Enter a description under the **Purpose** column and click the **UPDATE** button.
- 3 Expand the **Procedure** group in the left pane and double-click on the new object just created. This will start the editor.
- 4 Select the **Tools/View Procedure or Template** menu option. This will start a dialog box.
- 5 Select the existing **Get Key** function from the procedure drop-down. A new window will appear on the right showing the existing **Get Key** function.
- 6 Click the mouse in the existing function window. Right-click and choose **Select All** and then right-click again and select **Copy**.
- 7 Click the mouse in the main editor window and paste the existing function into the window. The existing function window can now be closed if desired.
- 8 Using the menu option **Edit/Replace**, change all occurrences of the old function name to the new name.
- 9 Change the various occurrences of the business key name and variable to match the new one and test it all.

Once a new **Get Key** function has been produced it will obviously be necessary to modify the stage update procedure to utilize the new function.



TIP: An alternative approach is to create a dimension view specifically to generate a different kind of get key function. Then use this view's surrogate key in the stage table instead of the dimension table's surrogate key. This approach avoids modifying the staging table procedure.

Extending the 'Get Key' function

If you want to keep all the information for the lookup of dimension keys in one location then extend the functionality of the existing **Get Key** function. The normal practice in this case is to include all potential business keys as parameters. Multiple select statements can be employed to select the keys based on the appropriate business key. To choose the correct statement either use Nulls in the unused business keys or include an additional parameter to indicate which business key is being passed.

If this method is used then any stage update procedures will need to be modified to handle the new calling syntax.

Dimension Initial Build

The initial population of a dimension with data can be achieved by generating a custom procedure and then; using the right-click on the dimension, choosing **Execute Custom Procedure via Scheduler**. The dimension should be analyzed once the custom procedure is completed so that the database query optimizer can make use of the indexes.

For smaller dimensions (i.e. less than 500,000 rows) run the normal **Update** procedure against an empty dimension table. There is however a danger in this action in that the query optimizer will not have any information on the table and hence will do a full table pass when adding each new row. For a very small dimension this will not be an issue, but it will rapidly become a major problem as the row count increases. The problem with the initial update is that the database does not know to use the index that has been created for the business key, and hence does a full table pass when attempting to update/insert a record.

To prevent this problem the generated code will issue an analyze statement after 1000 rows have been added to the table. The statement used is as follows:

SQL Server: update statistics *dim_table* with sample 5 percent

Oracle: analyze table *dim_table* estimate statistics sample 5 percent for all indexed columns

DB2: runstats on table *dim_table*

where *dim_table* is the name of the dimension.

This command will be issued whenever 1000 or more new rows are added to a dimension. If this is undesirable then the code should be removed from the update procedure. It may be undesirable if a plan has been locked in for the table.

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 screenshot shows the 'Dimension Column dim_customer.code' dialog box. It has a left sidebar with 'Properties' selected, and a main area with several sections: General, Physical Definition, Meta Definition, Code Generation, and Source Details. The 'Column Name' field is highlighted in blue. At the bottom, there are buttons for '< Update', 'Update ->', 'OK', 'Cancel', and 'Help'.

Dimension Column dim_customer.code	
General	
Table Name	dim_customer
Column Name	code
Business Display Name	code
Column Description	Unique code to identify the customer. Forms a hierarchy with city and state.
Physical Definition	
Column Order	20
Data Type	numeric(5)
Null Values Allowed	True
Default Value	
Meta Definition	
Format	#,##0
Numeric	True
Additive	True
Attribute	False
End User Layer Display	True
Business Key	True
Artificial Key	False
Key Type (0,A,B,C,...)	A
Code Generation	
Zero Key Value	
Source Details	
Source Table	load_customer
Source Column	code
Transformation	
Column Name Database-compliant name of the column. Dialog Opening Value: 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. 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 Data Warehouse Designer-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 data types for Oracle are integer, number, char(), varchar2() and date. For SQL Server common types are integer, numeric, varchar() and datetime. 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.

Artificial Key

Indicates whether the column is the generated artificial/surrogate key (unique identifier) for the table. Only one artificial key per table is supported. [Normally maintained automatically].

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

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. All dimensions that use standard QAD Data Warehouse Designer generated procedures have a row with an artificial key of zero. This row is used to link to the fact records when no match on the dimension can be found. For example we may have some sales records that come through without a product defined. QAD Data Warehouse Designer will by default associate these fact records with the zero key product dimension entry. So, we might set the zero key value to 'Unknown product' for the name column in the product dimension.

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

Transformation. [Read-only].


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.

- When positioned on the column **name** in the middle pane, right-click and select **Properties** from the drop-down menu.




- Change the column name from **name** to **cname** as shown below. Click **OK** to leave the Properties page.



- 4 Right-click on the **dim_customer** object in the left pane and select **Validate against Database**.

Properties
Storage
Display Columns
Display Indexes
Display Data
Query via Excel
Add Column
Add Index
Regenerate Indexes
Hierarchies ▶
Change Column(s)
Validate Against the Database
Update Comments
Gather Statistics ▶
Version Control ▶
Create (ReCreate)
Truncate
Delete Metadata and Drop Table
Execute Update Procedure
Execute Custom Procedure
Process Table via Scheduler
Execute Custom Procedure via Scheduler
Documentation ▶
Projects ▶
Impact ▶
Code ▶

- 5 The results in the middle pane will show that the metadata has been 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 dialog will appear, displaying the table and column name to be altered. Select **Alter Table**.
- 8 A dialog will appear 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:

The screenshot shows a software window titled "Dimension Column dim_customer.code". On the left is a sidebar with three tabs: "Properties", "Transformation" (which is selected and highlighted in blue), and "Language Mapping". The main area of the window contains the following elements:

- Target:** A text box containing "code" and a "Paste" button to its right.
- Source:** A text box containing "load_customer.code" and a "Paste" button to its right.
- Column Transformation Code (must execute within a SQL SELECT statement):** A large, empty text area for entering the transformation code.
- Function Set:** A dropdown menu currently set to "SQL Server Standard Funct".
- Functions and Available Columns:** A panel on the right side of the code area, containing a tree view with "Functions" and "Available Columns" (indicated by a small grid icon).
- Word Wrap:** A checkbox labeled "Word Wrap Displayed Code" which is currently unchecked.
- Function Syntax:** A text box for defining the function's syntax.
- Function Desc.:** A text box for defining the function's description.
- Buttons:** At the bottom right, there are five buttons: "<- Update", "Update ->", "OK" (highlighted with a blue border), "Cancel", and "Help".

Note: Transformations are only put into effect when the procedure is re-generated.

See **Transformations** (on page 603) for more details.

Dimension Hierarchies


The various hierarchies associated with a dimension can be recorded in the QAD Data Warehouse Designer metadata. These hierarchies are often not used in any form, except to provide documentary completeness. They can however be used in conjunction with the QAD 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 QAD Data Warehouse Designer, 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. To add a new hierarchy, right-click on the dimension table in the left pane and select **Hierarchies/Add Hierarchy**. The following dialog will appear.



The dialog box is titled "Add Hierarchy" and has a close button (X) in the top right corner. It contains a text field for "Hierarchy Name:" with the value "customer". Below this is a descriptive text: "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 dialog is divided into two main sections: "Dimension Columns:" on the left and "Hierarchy:" on the right. The "Dimension Columns:" section contains a list of columns with a scroll bar at the bottom: dim_customer_key, customer_code, customer_name, customer_category_code, customer_group_code, customer_subgroup_code, sales_source, dss_start_date, dss_end_date, dss_current_flag, dss_version, and dss_update_time. The "Hierarchy:" section contains a list of levels with a scroll bar at the bottom: customer_category_code, customer_group_code, customer_subgroup_code, and customer_name. Between these two lists are four buttons: a right arrow (>), a left arrow (<), an up arrow (↑), and a down arrow (↓). At the bottom right of the dialog are "OK" and "Cancel" buttons.

Add Hierarchy

Hierarchy Name:

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.

Dimension Columns:

- dim_customer_key
- customer_code
- customer_name
- customer_category_code
- customer_group_code
- customer_subgroup_code
- sales_source
- dss_start_date
- dss_end_date
- dss_current_flag
- dss_version
- dss_update_time

Hierarchy:

- customer_category_code
- customer_group_code
- customer_subgroup_code
- customer_name

> < ↑ ↓


OK Cancel

Enter a meaningful name for the hierarchy.

The hierarchy is built with the highest level at the top; for example a customer dimension may have **category_code** at the highest level, then **group_code**, then **subgroup_code** and finally **name** 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 pane. 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 window titled "Hierarchies for dim_customer". Inside is a table with four columns: "Dimension Name", "Hierarchy Name", "Description", and "Maintained". The first row is selected and has a context menu open over it. The context menu contains four options: "List Elements", "Add Elements", "Modify Hierarchy", and "Delete Hierarchy".

Dimension Name	Hierarchy Name	Description	Maintained
dim_customer	customer	Added at dimension creati...	N

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 QAD Data Warehouse Designer, 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 QAD Data Warehouse Designer. A snowflake can be created for normalized or partially normalized dimension tables, 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:

Dimension Business Key Definition

Dimension Table Column List:

- address
- city
- code
- dim_customer_key
- name
- state

Business Keys for snowflake dimension dim_state

Select the business key from the dimension table that matches each business key for this snowflake dimension.

Move them over to the business key list. They must be in the correct order to match the snowflake dimension business keys.

Dimension Business Key List:

- code

Snowflake Business Keys:

- state_code

dim_state

☐ No Warning (if dimension joins do not succeed)

☐ Automatically Add a New Entry (if no valid dimension record)

☐ Write Detail Log Message (when no dimension record match)

Source Table Column List:


Add Text

Enter a text string and press Add Text to add a static Business Key Value

OK Cancel

- 10 Add the business key column(s) and click **OK**.

The QAD Data Warehouse Designer screen should look like this:



Dimension Language Mapping

The Dimension 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 "Dimension dim_customer" with a close button (X) in the top right corner. On the left is a vertical tab menu with the following items: Properties, Storage, Language Mapping (which is highlighted with a blue background), Purpose, Concept, Grain, Examples, Usage, and Notes. The main area of the dialog is divided into two sections. The top section is labeled "Language :" and contains a dropdown menu currently set to "French". Below this is a section labeled "Language Settings:" which contains two text input fields. The first field is labeled "Business Display Name:" and contains the text "dim_customer". The second field is labeled "Translated Value of Business Display Name" and is currently empty. The bottom section is labeled "Description:" and contains two text input fields. The first field contains the text "This is the master Customer List for the Company" and has up and down arrow icons on its right side. The second field is labeled "Translated Value of Dimension Description" and is currently empty, also with up and down arrow icons on its 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 normalized 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 or stage 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. The steps are:

- ☐ Identify the source transactional data that will ultimately constitute the fact or normalized 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.
- ☐ 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 346).
- ☐ Add columns from other source tables. See ***Building the Stage Table*** (on page 346).
- ☐ Add any relevant dimension table or normalized table keys. See ***Building the Stage Table*** (on page 346).
- ☐ Create the stage table in the database. See ***Building the Stage Table*** (on page 346).
- ☐ Build the update procedure. See ***Generating the Staging Update Procedure*** (on page 349).
- ☐ Test the update procedure and analyze the results. See ***Tuning the Staging Update Process*** (on page 362).
- ☐ Modify the update procedure as required. See ***Tuning the Staging Update Process*** (on page 362).

In This Chapter

Building the Stage Table.....	346
Generating the Staging Update Procedure	349
Tuning the Staging Update Process	362
Stage Table Column Properties.....	363
Stage Table Column Transformations	367
Set Merge Procedure	368

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 fact table that will ultimately be 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.

In the tutorial the `stage_sales_detail` stage table is built by combining information from the `order_header` table and the `order_line` table. These two tables are connected via the `order_number` field which is common to both tables. The tutorial shows the building of a stage table and the joining of these two tables.

We could however also include two additional source tables being `invoice_header` and `invoice_line` which contain specific information relating to what was on the invoice. We may want our fact table to contain information from these tables as well. Although these two tables may contain the `order_number` and hence potentially allow a join with the order tables we may choose not to perform such a join for performance reasons. In such a situation we have three obvious choices in terms of how we ultimately update our fact table. These are:

- 1 Join all four tables using one large join in our staging table.
- 2 Update the staging table in two passes; one pass updating the order information, the other 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` fact table.

Although all three options are viable and a normal situation in the QAD Data Warehouse Designer environment, options (2) and (3) will require specific coding and modifications to the generated procedures from the outset. 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. We can 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 Data warehouse connection to display our load tables in the right pane. This is achieved via the **Browse/Source Tables** menu option, choosing the **Data warehouse connection**.
- 3 Drag the **primary load table** (i.e. the one with the most columns, or the lowest data granularity) from the right pane and drop it into the middle pane. A dialog will appear 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` is 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.
- 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 Drag in the **dimension artificial key** from each dimension that is to be joined to the stage/fact table. We can only join a dimension if a business key exists amongst the stage table columns or if it is possible to derive that business key in some way from the columns or other dimensions.

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	numeric(12)	load_order_line	order_number
🔗 order_line_no	order line no	numeric(4)	load_order_line	order_line_no
🔗 product_code	product code	numeric(6)	load_order_line	product_code
🔗 unit_sale_price	unit sale price	numeric(9,3)	load_order_line	unit_sale_price
🔗 quantity	quantity	numeric(8)	load_order_line	quantity
🔗 sales_value	sales value	numeric(13,2)	load_order_line	sales_value
🔗 tax	tax	numeric(9,2)	load_order_line	tax
🔗 order_date	order date	datetime	load_order_header	order_date
🔗 customer_code	customer code	numeric(6)	load_order_header	customer_code
🔗 ship_date	ship date	datetime	load_order_header	ship_date
🔗 dim_customer_key	dim customer ...	integer	dim_customer	dim_customer_key
🔗 dim_order_date_key	dim time key	integer	dim_order_date	dim_order_date_key
🔗 dim_product_key	dim product key	integer	dim_product	dim_product_key
🔗 dim_ship_date_key	dim time key	integer	dim_ship_date	dim_ship_date_key
🔗 dss_update_time	dss update time	datetime		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_line table, and the customer_code, order_date and ship_date came from the load_order_header columns. These two load tables will be joined via the order_number column. This order_number column appears in both load tables but is sourced, in this example, from the load_order_line table. Each dimension artificial key was dragged from its appropriate table. The final column 'dss_update_time' was generated by QAD Data Warehouse Designer 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. This is achieved by right-clicking on the stage table in the left pane and selecting **Create (ReCreate)** from the popup menu.

A results box will appear to show the results of the creation. The following example shows a successful creation.



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 to find such a column is to double-click on the list heading **Column Name**, which will sort the column names into alphabetical 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 349).

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 dimension table artificial keys.

Generating a Procedure


To generate a procedure, right-click on the stage table in the left pane and select **Properties**. From the **Update Procedure** drop-down list, select **(Build Procedure...)**. Click **OK** to update the properties and start the process of generating the new procedure.

A series of questions will be asked during the procedure generation to join the tables and link the dimensions.

Procedure type

The first dialog box asks for the type of procedure that is to be generated. An example of this follows:

DB2:



Oracle:

Define Stage Procedure Type

Code can be generated for each of the following procedure types. There are advantages and disadvantages with each type, so please read the help if you are unsure of the method to choose.

1. Standard cursor based update and insert. Allows detection of non unique business key.

2. Cursor based sorted by the dimension business keys with the lowest cardinality. Otherwise as per option 1.

3. Set based insert. Assumes business key is unique. Normally the fastest method, but the least flexible.

4. Set based insert followed by cursor on missing dimension keys. Assumes business key is unique.

5. Set based insert from all source tables (merge). A source table only needs to appear once. All source tables must have the same column names.

6. Bulk bind insert. Often the fastest method of resolving multiple dimension joins. However, uses a lot of memory. Requires Oracle 9.2 or greater.

Cancel

Cursor

Sorted Cursor

Set

Set + Cursor

Set Merge

Bulk Bind

☐ Distinct Data Select

☐ Allow Modification of the 'Where Clause

350

SQL Server:

Define Stage Procedure Type

Code can be generated for each of the following procedure types. There are advantages and disadvantages with each type, so please read the help if you are unsure of the method to choose.

1. Standard cursor based update and insert. Allows detection of non unique business key.

2. Cursor based sorted by the dimension business keys with the lowest cardinality. Otherwise as per option 1.

3. Set based insert. Assumes business key is unique. Normally the fastest method, but the least flexible.

4. Set based insert followed by cursor on missing dimension keys. Assumes business key is unique.

5. Set based insert from all source tables (merge). A source table only needs to appear once. All source tables must have the same column names.

Cancel

Cursor

Sorted Cursor

Set

Set + Cursor

Set Merge

Insert Hint: TABLOCK

Update Hint:

(e.g. TABLOCK)

(e.g. TABLOCK)

☐ Distinct Data Select

☐ Allow Modification of the Where Clause

A number of different types of procedure generation are available. Each type is discussed below. There is a checkbox at the bottom of the dialog box. This checkbox will only be present if advanced procedure building features are enabled. This checkbox allows the editing of the 'Where' clause when no table joining is being performed, and hence the 'Where' clause would not be exposed.

Cursor based procedure

This option creates a procedure that loops through each source transaction performing specific tasks on each transaction before writing the transaction to the stage table. The tasks performed include the lookup of each dimension key, and any bespoke value assignments or adjustments.

Sorted cursor procedure

This is a variant on the regular cursor procedure. In this case dimension key lookups only occur when there is a change in the business key for a given dimension. This can improve the performance of the procedure particularly when one or more of the dimensions has a low cardinality. Where slowly changing dimensions are used this advantage is lessened as the relative date of the dimension records also comes into play.

Set based procedure

A set based procedure performs one SQL statement to join all the dimensions and source tables together and then insert this data into the stage table. This is normally the fastest method of building the stage table. It is also the least flexible method and is easily prone to incorrect syntax. Care and experience must be used with this option to ensure that all source transactions are processed and that they are processed correctly. This option is only recommended for the experienced practitioner where performance is the overriding requirement. Particular care must be taken with 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 + Cursor Procedure

This option is normally used where a number of dimension keys will fail to join and we wish to automatically add new dimension records. The first phase of the generated code will perform a set based join and insert of the data. The dimensions that failed to join will be assigned the zero key value. The second phase of the procedure will use a cursor to cycle through all the stage records that have a zero key dimension join. By setting the appropriate flags in the dimension joins when building this procedure we can request that any missing dimension records be automatically added.

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 368)

Bulk Bind Procedure

The bulk bind procedure is only available for Oracle. It uses memory arrays to join all the information before writing it to the table. It can be the fastest method of performing a stage table update when a lot of dimension joins are being used.

Note: The bulk bind procedure is only available for Oracle.


There are also some options available in this dialog:

- ☐ **Allow modification of the Where clause.** The 'Where' clause dialog is displayed automatically if the table has more than one source table, but not for a single source table stage table unless this option is chosen.

- ☐ **Distinct data select.** Choosing this option adds the DISTINCT key word to the SELECT statement in the staging procedure.
- ☐ Use Minimal Logging (for set insert). **DB2** only. The staging procedure uses the **LOAD FROM CURSOR** option in DB2 to minimize database logging.
- ☐ Insert hint. **SQL Server** only. Enter a database hint to be used in the INSERT statement.
- ☐ Update hint. **SQL Server** only. Enter a database hint to be used in the UPDATE statement.

Parameters variables

If QAD Data Warehouse Designer parameters exist in the metadata, the following dialog is displayed. Any parameters selected in this dialog (by moving them to the right), 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.



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. QAD Data Warehouse Designer 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: QAD Data Warehouse Designer parameters should be unique within the first 28 characters to avoid conflicting variables names.

See **Parameters** for more information on QAD Data Warehouse Designer Parameters

Multiple source tables

If multiple tables are being joined and we did not choose a **Set** based procedure then an initial dialog will display (as follows) advising that multiple source tables have been detected and that QAD Data Warehouse Designer will only generate one cursor (one join) to handle these tables. If multiple cursors or passes of the data are required, as in scenario (2) in the previous section then it is best to proceed with one set of joins and manually do the second set after the generation of the procedure has completed.




Source table joins

If multiple source tables are detected then a dialog will appear to allow the definition of the joins between the various tables.

Note: Joins to dimension tables are not defined in this step. These joins only relate to the source tables.

The joining of the tables will provide either the data for a cursor loop, or part of the construct of the set based update in the update procedure. If a cursor based procedure is being generated then the default name for the cursor will be 'Load' unless altered. If a set based procedure then the cursor name is ignored.

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. For SQL Server an additional option is available to allow either an ANSI standard join or a 'Where' clause based join. In DB2, only a from clause join can be generated. The ANSI standard join should be chosen in most situations. See the example screen in the following section.



The dialog box is titled "Provide Cursor Mapping". It has a "Cursor Name (if applicable):" field with the value "Load". Below this is a text area with instructions: "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." The "Source Tables:" list on the left contains "load_order_header" and "load_order_line". The "From and Where Clause:" area is empty. Below the source tables are three buttons: "Outer Join", "Simple Join", and a checked checkbox for "ANSI join code generated". Below these is the instruction "Select the columns that join the two tables." followed by two dropdown menus. The first dropdown is set to "load_order_header" and the second to "order_id". To the right of these is a large list of column names, with "order_id" highlighted. At the bottom left is a checkbox for "Word Wrap Displayed Code". At the bottom right are "OK" and "Cancel" buttons.

Cursor Name (if applicable): Load

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:

Outer Join Simple Join ☒ ANSI join code generated

Select the columns that join the two tables.

load_order_header order_id load_order_line

base_sale_amount
base_tax_amount
created_datetime
creating_employee_id
last_change_datetime
last_change_employee_id
line_type_code
order_id
order_line_no
product_id
quantity
sale_amount
ship_date
tax_amount
tax_amount1
tax_amount2
tax_amount3
tax_amount4

☐ Word Wrap Displayed Code

OK Cancel

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. When QAD Data Warehouse Designer builds up the 'Where' clause for Oracle it places the outer join indicator (+) at the end of each line. As this indicator goes on the subordinate table, it must be that table whose column is selected last. 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. For example the Oracle outer join indicator (+) could be removed if required.

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. If it is not correct, it can of course be edited in the generated procedure.

For SQL Server data warehouses you have the choice between 'Where' statement joins and ANSI standard joins. For DB2 data warehouses, only ANSI standard joins are available.

Note: 'Where' joins will not work correctly in SQL Server if more than two tables are used in an outer join.

The example below shows the result of an ANSI standard join which takes place in the 'From' statement.

Provide Cursor Mapping

Cursor Name (if applicable): Load

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 load_order_header
JOIN 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.

load_order_header

load_order_line

☐ Word Wrap Displayed Code

OK

Cancel

Chapter 11 Staging - 357

Dimension Joins

For each dimension key a dialog will now appear 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. 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 **code**. We must provide the corresponding business key from the staging table. In the case of our example, this is the `customer_code` column.

The screenshot shows a dialog box titled "Dimension Business Key Definition" with a red close button in the top right corner. The dialog is divided into several sections:


- Stage Table Column List:** A list box containing 30 columns from a staging table, including `base_currency_code`, `base_sale_amount`, `base_tax_amount`, `bill_to_address_id`, `created_datetime`, `creating_employee_id`, `customer_code`, `gross_sale_amount`, `invoice_date`, `last_change_datetime`, `last_change_employee_id`, `line_type_code`, `net_sale_amount`, `order_date`, `order_id`, `order_line_no`, `order_number`, `order_status_code`, `order_type_code`, `payment_method_id`, `posted_date`, `product_id`, `quantity`, `sales_rep_id`, `sales_source`, `sales_tax_flag`, `ship_date`, `ship_to_address_id`, and `sold_to_address_id`.
- Business Keys for dim_customer:** A text area with instructions: "Select the business key from the staging table that matches each business key for this dimension. Move them over to the business key list. They must be in the correct order to match the dimension business keys."
- Stage Business Key List:** A list box containing the column `customer_code`.
- Dimension Business Keys:** A list box containing the column `customer_code`.
- dim_customer:** The dimension name is displayed at the bottom left.
- No Warning:** A checkbox labeled "No Warning (if dimension joins do not succeed)".
- Add Text:** A button next to a text input field labeled "Source Table Column List".
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

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.

Click **OK** after the correct business key has been entered. 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.

Slowly Changing Dimension information

If the dimension being joined was defined as a slowly changing dimension, then an additional dialog will appear asking for a date in the stage table that allows the coordination of the dimension record changes. Select the date and click **OK** or click **No Date** if a date does not apply.



The dialog box is titled "Tracking a Slowly Changing Dimension" and has a red close button in the top right corner. The main text area contains the following instructions: "This dimension has been identified as a slowly changing dimension. The columns listed below are tracked for changes." and "If there is a date in the staging table that correlates to when these changes occur select the date and enter OK." Below this, it says "dim_customer columns tracked for changes" and shows a text box containing "customer_name". Under the heading "Staging Table Dates:", there is a dropdown menu currently showing "order_date". At the bottom, there are two buttons: "OK" and "No Date".

For Example:

Say we have defined the `product_name` as a slowly changing column. This implies we want to create a new dimensional record whenever a product name is changed even though the SKU for the product remains the same.

Customer A orders 100 units of product SKU 123 on the 1st of the month and is shipped the order 3 days later on the 4th of the month.

The product has a name change on the 5th of the month. We may therefore decide that any of the product shipped before the 5th of the month were shipped under the old name and any product shipped on or after the 5th has the new product name.

If the example order shown above is not processed in the data warehouse until the 7th of the month we need to ensure it acquires the correct version of its product dimension record. We would therefore assign the 'Ship date' as the date that corresponds to the correct dimension record.

The product dimension record would therefore be assigned as the record in force on the 4th day of the month or more specifically the old product name.

Staging Business Key

Once all the dimensional joins have been defined we will be asked to define the unique business key for the staging table. This is the column or columns that allow us to uniquely identify each record in the staging table. In the example below and in our tutorial the unique business key is a combination of the `order_number` and the `order_line_no`.

Define Staging Business Key Columns

Column List:

- customer_code
- order_date
- product_code
- quantity
- sales_value
- ship_date
- tax
- unit_sale_price

Select the business keys that uniquely identify each record in the Staging table.

Move them over to the business key list

Business Key List:

- order_number
- order_line_no

☒ Include Update Statement

OK Cancel


None of the columns chosen as the business key should ever contain a NULL value. See the note at the start of the Dimension chapter.

Note: The order of the columns in the business key list is set to the order that the columns appear in the table. This is done to optimize performance for any indexes.

The **Include update statement** checkbox provides two options in the generated procedure. If checked then an update/insert combination will be included. If clear then only an Insert statement will be included in the procedure. If cleared you must be sure that the business key for the table is unique, otherwise either a duplicate key error will occur or records will be lost depending on the type of procedure.

Index re-creation

Finally, we're asked if we want to drop and re-create any indexes defined on the table:



Click **Yes** to drop and re-create, **No** to leave existing indexes in place.

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.

Tuning the Staging Update Process

This section is for **Cursor** based procedures only. It is not relevant for Set based or other types.

When generating the cursor based procedures to update the staging table QAD Data Warehouse Designer includes Update/Insert logic. This logic first attempts to update an existing staging record and then inserts a new record if the record does not exist. This build option should be altered once an understanding of the data you are dealing with is acquired. For the initial prototype and testing phases of a data warehouse analysis area we need to find out if our business keys for the fact table are genuinely unique. The update/insert logic in the staging procedure will only perform an 'Update' if our business key is not unique. In such a situation we need to decide if this is a failure in the business key, or some form of additive record situation we need to handle and code accordingly.

Removing the Update Component

Once we have resolved any data issues we would normally remove the update component to speed up the processing of the stage data. There are two ways of doing this. Firstly we can clear the **Include Update Statement** checkbox when defining the table business key. This will result in a procedure with only an insert statement. The other option is to do it manually as follows:

There are normally three steps to removing the update component. These being:

- 1 The update makes use of an index that can also be deleted in most situations.
- 2 The actual **Update** statement can be removed from the code such that the insert is always executed.
- 3 The procedure also includes an analyze of the table after 1000 rows have been inserted. This code is present to ensure that the index is used in the update. We can now remove this code. It would normally look like the following (Oracle version):

```
--=====
-- If 1000 or more rows inserted then perform an analyze
--=====
IF v_insert_count >= 1000 And v_analyze_flag = 0 THEN
    v_analyze_flag := 1;
    v_sql := 'analyze table stage estimate statistics sample 5 percent for all indexed columns';
    execute immediate v_sql;
END IF;
```

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 mouse 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 'Stage Table Column stage_forecast.product_code' dialog box. The 'Properties' tab is selected, showing a tree view with 'General', 'Physical Definition', 'Meta Definition', and 'Source Details'. The 'Column Name' property is highlighted in the 'General' section. The 'Column Name' is 'product_code', 'Business Display Name' is 'product code', and 'Column Description' is empty. The 'Physical Definition' section shows 'Column Order' as 40, 'Data Type' as integer, 'Null Values Allowed' as True, and 'Default Value' as empty. The 'Meta Definition' section shows 'Format' as '#,##0', 'Numeric' as True, 'Additive' as True, 'Attribute' as False, 'Business Key' as True, and 'Key Type (0,A,B,C,...)' as A. The 'Source Details' section shows 'Source Table' as load_forecast, 'Source Column' as product_code, 'Transformation' as empty, and 'Join' as False. At the bottom, there is a 'Column Name' section with the text 'Database-compliant name of the column. Dialog Opening Value: product_code'. The dialog has buttons for '< Update', 'Update >', 'OK', 'Cancel', and 'Help'.

Stage Table Column stage_forecast.product_code	
Properties	
Transformation	
General	
Table Name	stage_forecast
Column Name	product_code
Business Display Name	product code
Column Description	
Physical Definition	
Column Order	40
Data Type	integer
Null Values Allowed	True
Default Value	
Meta Definition	
Format	#,##0
Numeric	True
Additive	True
Attribute	False
Business Key	True
Key Type (0,A,B,C,...)	A
Source Details	
Source Table	load_forecast
Source Column	product_code
Transformation	
Join	False
Column Name	
Database-compliant name of the column. Dialog Opening Value: product_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. 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.

Business Display Name

Name that the business uses to refer to the column, which is included in the Data Warehouse Designer-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 data types for Oracle are integer, number, char(), varchar2() and date. For SQL Server common types are integer, numeric, varchar() and datetime. 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.

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.

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 stage 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 dimensional key where a dimension is being joined.

Transformation

Transformation. [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.

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:

The screenshot shows a software window titled "Stage Table Column stage_customer.customer_name". On the left is a "Properties" sidebar with a "Transformation" tab selected. The main area contains fields for "Target:" (customer_name) and "Source:" (load_customer.customer_name), each with a "Paste" button. Below these is a large text area for "Column Transformation Code (must execute within a SQL SELECT statement)". To the right of the code area is a "Function Set" dropdown menu set to "SQL Server Standard Funct" and a "Functions" panel with an "Available Columns" icon. At the bottom left of the main area are checkboxes for "Word Wrap Displayed Code", "Function Syntax:", and "Function Desc:". The bottom of the window features a row of buttons: "< Update", "Update >", "OK", "Cancel", and "Help".

Note: Transformations are only put into effect when the procedure is re-generated.


See **Transformations** (on page 603) for more details.

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.

The browser pane shows the three load tables:



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. Name the stage table, for example **stage_orders_Master**.




Next, modify the source table column to include one instance of each of the three load tables; the order does not matter. You can achieve this either by typing in the table directly, or by visiting each of the column Properties.

Stage Table Columns			
Column Name	Display Name	Data Type	Source Table
order_number	order number	numeric(12)	load_orders_A
order_date	order date	datetime	load_orders_B
customer_code	customer code	numeric(6)	load_orders_C
ship_date	ship date	datetime	load_orders_A
dss_update_time	dss update time	datetime	

Right-click on the stage table and select **Create (ReCreate)**.

Properties
Storage
Display Columns
Display Indexes
Display Data
Query via Excel
Report Zero Keys...
Add Column
Add Index
Regenerate Indexes
Change Column(s)
Validate Against the Database
Update Comments
Gather Statistics ▶
Version Control ▶
Create (ReCreate)
Truncate
Delete Metadata and Drop Table
Execute Update Procedure
Execute Custom Procedure
Process Table via Scheduler
Execute Custom Procedure via Scheduler
Documentation ▶
Projects ▶
Impact ▶
Code ▶

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.



The dialog box is titled "Define Stage Procedure Type" and has a red close button in the top right corner. It contains a help icon in the top right. The main text area explains that code can be generated for five procedure types and advises reading the help if unsure. A note states "NOTE: Last build was a type 5. SET MERGE". The five options are listed on the left, each with a corresponding button on the right. The "Set Merge" button is highlighted with a blue border. Below the list are two hint fields: "Insert Hint:" with the value "TABLOCK" and "Update Hint:" which is empty. At the bottom are two checkboxes: "Distinct Data Select" and "Allow Modification of the Where Clause", both of which are unchecked.

Code can be generated for each of the following procedure types. There are advantages and disadvantages with each type, so please read the help if you are unsure of the method to choose.

NOTE: Last build was a type 5. SET MERGE

1. Standard cursor based update and insert. Allows detection of non unique business key.
2. Cursor based sorted by the dimension business keys with the lowest cardinality. Otherwise as per option 1.
3. Set based insert. Assumes business key is unique. Normally the fastest method, but the least flexible.
4. Set based insert followed by cursor on missing dimension keys. Assumes business key is unique.
5. Set based insert from all source tables (merge). A source table only needs to appear once. All source tables must have the same column names.

Buttons: Cancel, Cursor, Sorted Cursor, Set, Set + Cursor, Set Merge

Insert Hint: (e.g. TABLOCK)

Update Hint: (e.g. TABLOCK)

☐ Distinct Data Select

☐ Allow Modification of the Where Clause

The stored procedure is rebuilt, as follows:


```

|-----
-- DBMS Name      :      SQL Server
-- Script Name    :      update_stage_orders_Master
-- Description    :      Build the staging table stage_orders_Master
-- Generated by   :      Version 6.5.5.2 (build 111221)
-- Generated for  :      Testing, Tutorials and Documentation
-- Generated on   :      Tuesday, February 07, 2012 at 20:43:10
-- Author        :      QuickStart
|-----
-- Notes / History
--
CREATE PROCEDURE update_stage_orders_Master
    @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
    , @v_insert_count integer      -- no of records inserted
    , @v_count       integer       -- General counter
    , @v_db_code     varchar(10)   -- Database error code
    , @v_db_msg      varchar(100)  -- Database error message

|-----
-- General Variables
|-----
DECLARE
    @v_return_status integer      -- Update result status
    , @v_getkey_status integer    -- GetKey procedure status
    , @v_row_count    integer     -- General row count
    , @v_truncate_date datetime   -- Used to set date to midnight
    , @v_dss_source_system_key integer -- Used for dimension joins
    , @v_dss_current_flag char(1) -- Used for dimension joins
    , @v_dss_update_time datetime -- Used for date insert

|-----
-- MAIN
|-----
SELECT @v_step = 0
BEGIN TRY

-- Set initial variable values
SELECT
    @v_set          = 0
    , @v_analyze_flag = 0
    , @v_step        = 1
    , @v_update_count = 0
    , @v_insert_count = 0
    , @v_dss_source_system_key = 1
    , @v_dss_current_flag = 'Y'
    , @v_dss_update_time = GETDATE()

```

```

-----
-- Delete all existing data from the staging table
-----
SET @v_sql = N'TRUNCATE TABLE stage_orders_Master'
EXEC @v_return_status = sp_executesql @v_sql
IF @v_return_status <> 0
BEGIN
    SET @v_db_code = CONVERT(varchar, @v_return_status)
    SELECT @v_db_msg = description FROM master.dbo.sysmessages
    WHERE error = @v_return_status
    EXEC WsWrkAudit
        @p_status_code = 'W'
        , @p_job_name   = @p_job_name
        , @p_task_name  = @p_task_name
        , @p_sequence   = @p_sequence
        , @p_message    = 'Failure when truncating table'
        , @p_db_code    = @v_db_code
        , @p_db_msg     = @v_db_msg
        , @p_task_key   = @p_task_id
        , @p_job_key    = @p_job_id

    SET @p_return_msg = ' Failure to truncate stage table. '
    + @v_db_msg
    SET @p_status = -2
    RETURN 0
END

-----
-- Insert input records into stage_orders_Master
-----
SET @v_step = 100

BEGIN TRANSACTION

INSERT INTO stage_orders_Master WITH ( TABLOCK )
(
    order_number,
    order_date,
    customer_code,
    ship_date,
    dss_update_time
)
SELECT
    load_orders_A.order_number,
    load_orders_A.order_date,
    load_orders_A.customer_code,
    load_orders_A.ship_date,
    @v_dss_update_time
FROM
    load_orders_A

SELECT
    @v_return_status = @@ERROR
    , @v_insert_count = @v_insert_count + @@ROWCOUNT

IF @v_return_status <> 0
BEGIN
    ROLLBACK
    SET @p_return_msg = 'Failed Insert into table stage_orders_Master'
    SET @v_db_code = CONVERT(varchar, @v_return_status)
    SELECT @v_db_msg = description FROM master.dbo.sysmessages
    WHERE error = @v_return_status
    EXEC WsWrkAudit 'E', @p_job_name, @p_task_name, @p_sequence,
        @p_return_msg, @v_db_code, @v_db_msg, @p_task_id, @p_job_id
    SET @p_status = -2
    RETURN 0
END

```

```

COMMIT

-----
-- Insert input records into stage_orders_Master
-----
SET @v_step = 100

BEGIN TRANSACTION

INSERT INTO stage_orders_Master WITH ( TABLOCK )
(
    order_number,
    order_date,
    customer_code,
    ship_date,
    dss_update_time
)
SELECT
    load_orders_B.order_number,
    load_orders_B.order_date,
    load_orders_B.customer_code,
    load_orders_B.ship_date,
    @v_dss_update_time
FROM
    load_orders_B

SELECT
    @v_return_status = @@ERROR
, @v_insert_count = @v_insert_count + @@ROWCOUNT

IF @v_return_status <> 0
BEGIN
    ROLLBACK
    SET @p_return_msg = 'Failed Insert into table stage_orders_Master'
    SET @v_db_code = CONVERT(varchar, @v_return_status)
    SELECT @v_db_msg = description FROM master.dbo.sysmessages
    WHERE error = @v_return_status
    EXEC WsWrkAudit 'E', @p_job_name, @p_task_name, @p_sequence,
        @p_return_msg, @v_db_code, @v_db_msg, @p_task_id, @p_job_id
    SET @p_status = -2
    RETURN 0
END

COMMIT

```

```

-----
-- Insert input records into stage_orders_Master
-----
SET @v_step = 100

BEGIN TRANSACTION

INSERT INTO stage_orders_Master WITH ( TABLOCK )
(
    order_number,
    order_date,
    customer_code,
    ship_date,
    dss_update_time
)
SELECT
    load_orders_C.order_number,
    load_orders_C.order_date,
    load_orders_C.customer_code,
    load_orders_C.ship_date,
    @v_dss_update_time
FROM
    load_orders_C

SELECT
    @v_return_status = @@ERROR
, @v_insert_count = @v_insert_count + @@ROWCOUNT

IF @v_return_status <> 0
BEGIN
    ROLLBACK
    SET @p_return_msg = 'Failed Insert into table stage_orders_Master'
    SET @v_db_code = CONVERT(varchar, @v_return_status)
    SELECT @v_db_msg = description FROM master.dbo.sysmessages
    WHERE error = @v_return_status
    EXEC WsWrkAudit 'E', @p_job_name, @p_task_name, @p_sequence,
        @p_return_msg, @v_db_code, @v_db_msg, @p_task_id, @p_job_id
    SET @p_status = -2
    RETURN 0
END

COMMIT

-----
-- All Done report the results
-----

```

Chapter 12

Data Store Objects

In This Chapter

Data Store Object Overview	378
Building a Data Store Object.....	380
Generating the Data Store Update Procedure	384
Data Store Artificial Keys.....	397
Data Store Column Properties	399
Data Store Object Column Transformations.....	403

Data Store Object Overview

A Data Store Object is a data warehouse object used to store any type of data for later processing. In QAD Data Warehouse Designer, 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 1072) for more details.

The usual steps for creating a Data Store model are defined below and are covered in this chapter. The steps are:

- ☐ 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.
- ☐ 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 380).
- ☐ 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 a Data Store Object** (on page 380).
- ☐ Add columns from other load and/or stage tables if required. See **Building a Data Store Object** (on page 380).
- ☐ Create the Data Store Object in the database. See **Building a Data Store Object** (on page 380).
- ☐ Build the update procedure. See **Generating the Data Store Update Procedure** (on page 384).
- ☐ Run the update procedure and analyze the results.

If necessary, modify the update procedure or create a custom procedure.

Overview: Data Store



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:

- ☐ **DB2:** COALESCE(business_key,'N/A')
- ☐ **SQL Server:** ISNULL(business_key,'N/A')
- ☐ **Oracle:** NVL(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

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 such a list is displayed in the middle pane, the pane is identified as a target for new Data Store Objects.

Browse to the Data Warehouse via the **Browse/Source Data** menu option. Drag the load or stage table, that contains the bulk of the Data Store Object 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 Data Store Object and will provide a default name based on the load or stage table name. Either accept this name or enter the name of the Data Store Object and click **ADD** to proceed:



Add a New Metadata Object [X]

Define the Type and Name of the New Object.

Specific information for each object type is defined in subsequent screens.

Object Type: Data Store ▼

Object Name: ds_customer

[ADD] [Cancel]


Data Store Object Properties

The table Properties dialog for the new table is now displayed. 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** (on page 300) 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 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


The following dialog appears after the Properties page and asks if you want to create the Data Store Object 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 additional columns you can select the **Create and Load** button. In such an event it would be normal practice to select the **(Build Procedure...)** option from the **Update Procedure** drop-down box in the previous Properties page. If no procedure has been selected then select the **Create** button as no update will occur in any event.

If **Create** or **Create and Load** is selected and a new procedure creation was chosen proceed directly to the **Generating the Update Procedure** (see "**Generating the Normalized Update Procedure**" on page 412) 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 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 **Delete**.

The name of a column can also be changed by selecting the column and using the right mouse 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 Data Store Object 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 QAD Data Warehouse Designer ancillary columns that do not have a source table. These columns have been added by QAD Data Warehouse Designer, 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.
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 we need to physically create the table in the database. This is achieved by right-clicking on the Data Store Object in the left pane and selecting **Create (ReCreate)** from the popup 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 Data Store Object 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 to find such a column is to double-click on the list heading **Col name**, which will sort the column names into alphabetical order. 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 384).

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.


Generating a Procedure

To generate a procedure, right-click on the Data Store Object 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.

Procedure Type and Options

The first dialog displayed when generating a Data Store Object update procedure is the procedure types and options dialog:



The dialog box is titled "Define Data Store Procedure Type and Options". It contains the following sections and controls:

- Procedure Types**: A section with a "Procedure Type:" dropdown menu set to "Set" and a "Previous build was : Set" label.
- Build Options**: A large section containing several checkboxes and input fields:
 - ☐ Process in Batch
 - ☐ Distinct Data Select
 - ☐ Include Delete Before Insert
 - ☐ Include Initial Load Insert
 - ☒ Include Update Statement
 - ☒ Changed Rows Only (with a dropdown menu set to "Except")
 - Update Hint: (e.g. TABLOCK)
 - ☒ Include Insert Statement
 - ☒ New Rows Only (with a dropdown menu set to "Except")
 - Insert Hint: (e.g. TABLOCK)
 - ☒ Use Merge Statement for Update and Insert
 - ☐ Changed and New Rows Only
 - ☐ Delete Rows Not in Source
 - Merge Hint: (e.g. TABLOCK)

At the bottom right, there are "OK" and "Cancel" buttons.

Several different questions need to be answered 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.

Procedure Types

A number of different types of procedure generation are available. The procedure type is chosen from the procedure type drop-down list. A **set** based procedure performs one SQL statement to join all the source tables together and then insert this data into the Data Store Object. This is normally the fastest method of populating a table.


Build Options

The following build options are available:

Process in Batch

The **Process in Batch** checkbox 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).

The following dialog is the default for a non-date column:




The image shows a Windows-style dialog box titled "Batch Processing". It has a light blue header bar with a red close button in the top right corner. The main area is light gray and contains the following text:

Select 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 then choose an attribute of
that related table to group by.

Below the text is a dropdown menu with "customer_code" selected and a small downward arrow on the right. At the bottom right of the dialog are two buttons: "OK" and "Cancel".

The following dialog is displayed for date fields:



The image shows a 'Batch Processing' dialog box with a light blue title bar and a red close button. The main area is light gray. It contains two instructions: 'Select a field to batch process on.' and '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 then choose an attribute of that related table to group by.' Below these is a dropdown menu showing 'created_datetime'. Another instruction, 'Select the date format that will be used to group by for the processing', is followed by a second dropdown menu. This menu is open, showing options: 'created_datetime', 'Day', 'Month' (which is highlighted in blue), and 'Year'. At the bottom right are 'OK' and 'Cancel' buttons.

Batch Processing

Select 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 then choose an attribute of that related table to group by.

created_datetime

Select the date format that will be used to group by for the processing

created_datetime
Day
Month
Year


OK Cancel

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 off (i.e. duplicates are not removed).

Include Delete Before Insert

The **include 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 selected, the following dialog is displayed:



The dialog box is titled "Definition for a Delete Statement Before the Data Store Table Insert". It contains the following elements:

- Instructions:** "Select the Data Store Table column on which the delete will be based. All values in the source table for the selected column will be used to delete the corresponding records in the Data Store Table."
- Data Store Table Column:** A dropdown menu with "customer_name" selected.
- Source Table:** A dropdown menu with "load_customer_sales2" selected.
- Source Table Column:** A dropdown menu with "customer_name" selected.
- Options:** Three checkboxes are present:
 - ☐ Issue Warning if a Delete Occurs
 - ☐ Multiple Pass Delete (smaller transactions)
 - ☐ Truncate Table Before Insert
- Delete Where Clause:** A text area containing the SQL code:

```
WHERE ds_customer.customer_name IN ( SELECT DISTINCT load_customer_sales2.customer_name FROM load_cv
```
- Help Text:** "The Where Clause can be edited, or directly entered."
- Word Wrap:** A checkbox labeled "Word Wrap Displayed Code" which is currently unchecked.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

Use the wizard to build a 'Where' clause for the delete statement. Additional options included in this dialog are:

- ☐ Issue warning if a delete occurs
- ☐ Multiple pass delete - if this option is chosen, the delete is performed in a loop, 50000 rows at a time, until the required rows have been deleted.
- ☐ Truncate table before insert - if this option is chosen, the delete is ignore and the Data Store Object is truncated.

The default for this field is off (i.e. a delete statement is not added to the procedure).

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).

Insert Zero Key Record

This option applies only to datastore, normalized tables and set based dimensions. It allows you to include the zerokey (unknown record) for tables with an artificial key.

Include Update Statement with/without Minus Select using except/checksum

The **include update statement** option includes an update statement in the procedure to update changing rows in the Data Store 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 Data Store 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).

Include Insert Statement with/without Minus Select using except/checksum

The **include insert statement** option includes an insert statement in the procedure to insert new rows in the Data Store 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 Hint.

Enter a database hint to be used in the UPDATE statement.

Note: SQL Server only.

Insert Hint.

Enter a database hint to be used in the INSERT statement.

Note: SQL Server only.

Use Merge Statement for Update and Insert

The **Use Merge Statement for Update and Insert** option allows you to use the Merge statement instead of two separate Update and Insert statements.

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 *except* change detection or *checksum* change detection to work out what rows in the Data Store Object are new or require updating

Delete Rows Not in Source

The **Delete Rows not in Source** option allows you to delete all the rows which are not in the source but in the target table as part of the Merge statement. To avoid any accidental deletions, this option is de-activated if **Process in Batch** or **Changed and New Rows only** are selected.


Merge Hint

Enter a database hint to be used in the MERGE statement.

Note: SQL Server only.

Parameters variables

If QAD Data Warehouse Designer parameters exist in the metadata, the following dialog is displayed. Any parameters selected in this dialog (by moving them to the right), 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.



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. QAD Data Warehouse Designer 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: QAD Data Warehouse Designer parameters should be unique within the first 28 characters to avoid conflicting variables names.

See **Parameters** for more information on QAD Data Warehouse Designer Parameters

Business Key definition

The define business key dialog will appear asking for the business key that will uniquely identify each Data Store Object record. The source table from which the Data Store Object 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.

Define Data Store Business Key Columns

Column List:

- active_flag
- bill_to_address_id
- created_datetime
- creating_employee_id
- customer_category_code
- customer_category_description
- customer_group_code
- customer_group_description
- customer_legal_name
- customer_name
- customer_subgroup_code
- customer_subgroup_description
- last_change_datetime
- last_change_employee_id
- primary_address_type
- primary_contact_person
- ship_to_address_id
- sold_to_address_id
- territory_id

Select the business keys that uniquely identify each record in the Data Store table.

Move them over to the business key list

Business Key List:

- customer_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.

None of the columns chosen as the business key should ever contain a NULL value. See the note at the start of this chapter.

Change Detection and History Tables

If the Data Store Object was defined as a history table then the following additional dialogs will appear. Data Store history tables are like slowly changing dimensions. The first dialog requests the selection of the columns to be managed as change detection columns. Select the required columns and click **OK** to proceed. In the example below the customer name is to be managed as a change detection column.

Define Data Store Change Detection Fields

Column List:

- active_flag
- bill_to_address_id
- created_datetime
- creating_employee_id
- customer_category_code
- customer_category_description
- customer_code
- customer_group_code
- customer_group_description
- customer_legal_name
- customer_subgroup_code
- customer_subgroup_description
- dss_current_flag
- dss_end_date
- dss_start_date
- dss_version
- last_change_datetime
- last_change_employee_id
- primary_address_type
- primary_contact_person
- ship_to_address_id
- sold_to_address_id
- territory_id

Select the fields that change detection is required for in the Data Store Update query.

Move them over to the change detection field list

Change Detection Field List:

- customer_name

>

<

OK Cancel

The following dialog will appear if one or more columns are chosen as change detection columns. When completed, click **OK** to continue.

History record settings

☐ Add NULL Support

☒ Update Current Records Only

Please enter the default start and end dates to be used.
Correctly fomatted dates or variables may be used.
Pressing Reset will reset the dates to the default settings.

Reset

Start Date for Initial Member (A):

CAST('01-JAN-1900' AS datetime)

End Date for Current Member (B):

CAST('31-DEC-2999' AS datetime)

Start Date for New Member Entry (C):

@v_current_date

End Date for Expiring Member Entry (D):

@v_current_date -0.00000005

Business ID	History Field	Start Date	End Date	Version	Current Flag
Project RED	Started	(A)	(D)	1	N
Project RED	Complete	(C)	(B)	2	Y

OK

Cancel

Add NULL Support

If the **Add null support** option is selected, then change detect column management will cater for null values in any change detect columns. If this option is not selected, and null values are present, errors may occur running the update procedure. The default for this value is off (i.e. nulls are not catered for).

Update Current Records Only

The **Update Current record only** option will only apply changes to non-change detect columns on the **current** record. If this option is not selected, all past (non-current) rows will also be updated to reflect changes to non-change detect columns. The default for this value is on (i.e. only the current record is updated).

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 correct data type. The default value provided will usually be cast to the correct datatype and can be treated as a template. The default value for this field is 1 January 1900.

End Date for Current Member

The **End date for current member** field contains the end 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 correct data type. The default value provided will usually be cast to the correct datatype and can be treated as a template. The default value 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 correct data type. The default value provided will usually be cast to the correct datatype and can be treated as a template. The default value for this field is the current date and time.

End Date for Expiring Member Entry

The **End date for expiring member entry** field contains the end date for any rows updated 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 correct data type. The default value provided will usually be cast to the correct datatype and can be treated as a template. The default value 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).

Reset

Click on the **Reset** button to return all fields in this dialog to their default values.

Joining multiple source tables, where and group by clauses

Next the **Source table mapping** dialog appears. This dialog is used to join source tables, add 'Where' clauses and specify group by clauses.

Source Table Mapping

Define the Joins (or edit the join 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_customer_sales2]

Join Clause:

FROM [load_customer_sales2] load_customer_sales2

Where Clause:

Outer Join

Simple Join

Table A

Column A

Table B

Column B

Group By Clause:

☐ Word Wrap Displayed Code

OK

Cancel

To join two tables, start by selecting the two tables in the left box and then clicking 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_order_header` and `load_order_line` tables are joined by the column `order_no`. For multiple column joins, repeat the join process.

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. 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.

Note: When QAD Data Warehouse Designer 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.

As the join columns are selected the join 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. For example the Oracle outer join indicator (+) could be removed if required.

Where Clause

A 'Where' clause can be manually added to the **Where Clause** field in this dialog.

Group By Clause

A group by clause (and any other clauses that come after the 'Where' clause) can be added to the group by clause in this dialog.

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 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:

Results		
Object	Message	
ds_customer	EXECUTE sp_addextendedproperty N'Comment', N'Date and time the row was updated in the data warehouse.', N'user', N'dbo', N'table', 'ds_customer', N'column', N'dss_update_time';	
ds_customer	Create indexes on Data Store dbo.ds_customer completed successfully.	
ds_customer	CREATE UNIQUE NONCLUSTERED INDEX ds_customer_idx_A ON dbo.ds_customer (code) WITH (SORT_IN_TEMPDB = OFF);	
ds_customer	Procedure Completed	
ds_customer	1 ds_customer updated. 6 records added. 0 records updated.	

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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer object for this table (perhaps a Normalized Table would be more appropriate). The quickest way to do this is to add an extra column to the Data Store Object by using either **Add Column** or **Copy Column**. Edit the properties of the new column to have the correct name and order, source table and column, datatype, key type and flags. Specifically:

- ☐ The **Column Name** and **Source Column** name should be the same.
- ☐ The **Source Table** should be empty.
- ☐ The **Data Type** should be:
 - DB2: integer generated by default as identity (start with 1, increment by 1)
 - Oracle: integer
 - SQL Server: integer identity(0,1)
- ☐ The **Key Type** should be 0.
- ☐ Only the **Numeric** and **Artificial Key** flags should be set on.

The following example shows a manually added artificial key column:

The screenshot shows the 'Data Store Column ds_customer.customer_key' dialog box. The 'General' tab is selected, displaying the following properties:

Property	Value
Table Name	ds_customer
Column Name	customer_key
Business Display Name	customer_key
Column Description	Manually added artificial key
Physical Definition	
Column Order	10
Data Type	integer
Null Values Allowed	False
Default Value	
Meta Definition	
Format	
Numeric	True
Additive	False
Attribute	False
Business Key	False
Artificial Key	True
Key Type (0,A,B,C,...)	0
Code Generation	
Zero Key Value	0
Source Details	
Source Table	(Empty)
Source Column	customer_key
Transformation	
Join	False

Buttons at the bottom: OK, Cancel, Help. Top right buttons: < Update, Update >.

The artificial key for a Data Store Object is set via a sequence in Oracle and an identity column in SQL Server and DB2. This artificial key normally, and by default, starts at one and progresses as far as is required.

A QAD standard for the creation of special rows in the data store 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 QAD generated code that any normalized 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 normalized 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.

Note: To have a surrogate key auto added for Data Store tables, see ***Global Naming of Key Columns*** (on page 57).

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 mouse menu to compare the metadata to the table in the database. A right mouse 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 'Data Store Column ds_customer.code' dialog box. It has a left sidebar with 'Properties' and 'Transformation' tabs. The main area is divided into sections: General, Physical Definition, Meta Definition, and Source Details. The 'Column Name' field is highlighted in blue. At the bottom, there are buttons for '< Update', 'Update >', 'OK', 'Cancel', and 'Help'.

Data Store Column ds_customer.code	
General	
Table Name	ds_customer
Column Name	code
Business Display Name	code
Column Description	Unique code to identify the customer. Forms a hierarchy with city and state.
Physical Definition	
Column Order	10
Data Type	numeric(6)
Null Values Allowed	True
Default Value	
Meta Definition	
Format	#,##0
Numeric	True
Additive	True
Attribute	False
Business Key	True
Artificial Key	False
Key Type (0,A,B,C,...)	A
Source Details	
Source Table	load_customer
Source Column	code
Transformation	
Join	False

Column Name
Database-compliant name of the column.
Dialog Opening Value: code

< 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 Data Warehouse Designer-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 data types for Oracle are integer, number, char(), varchar2() and date. For SQL Server common types are integer, numeric, varchar() and datetime. 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.

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 QAD Data Warehouse Designer 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.

Key type	Meaning
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

Transformation. [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 Object 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:

The screenshot shows a dialog box titled "Data Store Column ds_customer.code". On the left is a "Properties" sidebar with a "Transformation" tab selected. The main area contains fields for "Target:" (code) and "Source:" (load_customer.code), each with a "Paste" button. Below these is a large text area for "Column Transformation Code (must execute within a SQL SELECT statement)". To the right of this area is a "Function Set" dropdown menu set to "SQL Server Standard Funct" and a list of "Available Columns". At the bottom left, there is a checkbox for "Word Wrap Displayed Code", and fields for "Function Syntax:" and "Function Desc:". At the bottom right are buttons for "<- Update", "Update ->", "OK", "Cancel", and "Help".

Note: Transformations are only put into effect when the procedure is re-generated.

See **Transformations** (on page 603) for more details.

Chapter 13

Normalized Tables

In This Chapter

Normalized Tables Overview	406
Building a Normalized Table	408
Generating the Normalized Update Procedure	412
Normalized Table Artificial Keys	426
Normalized Table Column Properties	428
Normalized Table Column Transformations	432

Normalized Tables Overview

A Normalized Table is a data warehouse object used to build third normal form enterprise data warehouses. In QAD Data Warehouse Designer, normalized 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. Normalized Tables can be considered either reference or transactional in nature.

A Normalized Table is built from the Data Warehouse connection. Unless you are retrofitting an existing system, Normalized Tables are typically built from one or more load or stage tables. The normalized model may be retrofitted from an enterprise modeling tool. See **Importing a Data Model** (on page 1072) for more details.

The usual steps for creating a normalized model are defined below and are covered in this chapter. The steps are:

- ☐ Identify the source reference or transactional data that will constitute the Normalized 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.
- ☐ Using the 'drag and drop' functionality drag the load or stage table that is the primary source of information for the Normalized Table into a normalized target. See **Building a Normalized Table** (on page 408).
- ☐ 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 a Normalized Table** (on page 408).
- ☐ Add columns from other load and/or stage tables if required. See **Building a Normalized Table** (on page 408).
- ☐ Create the Normalized Table in the database. See **Building a Normalized Table** (on page 408).
- ☐ Build the update procedure. See **Generating the Update Procedure** (see "**Generating the Normalized Update Procedure**" on page 412).
- ☐ Run the update procedure and analyze the results.

If necessary, modify the update procedure or create a custom procedure.

Normalized Table Keys

Normalized 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 a Normalized Table. Where the Normalized 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:

- ☐ **DB2:** COALESCE(business_key,'N/A')
- ☐ **SQL Server:** ISNULL(business_key,'N/A')
- ☐ **Oracle:** NVL(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 normalized 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, Normalized Tables in QAD Data Warehouse Designer do not have an artificial key (artificial keys can be added manually, see **Normalized Table Artificial Keys** (on page 426) for more details).

An artificial key is the unique identifier that can be used to join a Normalized Table record to other Normalized Tables. When joining Normalized Tables it would be possible to perform the join using the business key. For Normalized Tables that satisfy one of 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 Normalized 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 a Normalized Table


Normalized tables are often sourced from one table in the base application. The process for building a normalized table begins with the drag and drop of the load or stage table that contains the bulk of the normalized table's information.

Drag and Drop

Create a normalized table target by double-clicking on the **Normalized** group in the left pane. The middle pane will display a list of all existing normalized tables in the current project. When such a list is displayed in the middle pane, the pane is identified as a target for new normalized tables.

Browse to the Data Warehouse via the Browse/Source Data menu option. Drag the load or stage table, that contains the bulk of the normalized 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 Normalized table and will provide a default name based on the load or stage table name. Either accept this name or enter the name of the normalized table and click ADD to proceed:



Add a New Metadata Object [X]

Define the Type and Name of the New Object.

Specific information for each object type is defined in subsequent screens.

Object Type:

Object Name:


Normalized Table Properties

The table Properties dialog for the new table is now displayed. If required, the normalized 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 **Building a Dimension** (on page 300) for more information. Change the storage options if desired.

If prototyping, and the normalized table is simple (i.e. one source table) then it is possible to create, load and update the normalized 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


The following dialog appears after the Properties page and asks if you want to create the normalized 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 additional columns you can select the **Create and Load** button. In such an event it would be normal practice to select the **(Build Procedure...)** option from the **Update Procedure** drop-down box in the previous Properties page. If no procedure has been selected then select the **Create** button as no update will occur in any event.

If **Create** or **Create and Load** is selected and a new procedure creation was chosen proceed directly to the **Generating the Update Procedure** (see "**Generating the Normalized Update Procedure**" on page 412) 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 normalized 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 **Delete**.

The name of a column can also be changed by right-clicking on the column and choosing **Properties** 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 normalized 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 QAD Data Warehouse Designer ancillary columns that do not have a source table. These columns have been added by QAD Data Warehouse Designer, 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 normalized table record came into existence. It is used to ascertain which normalized 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 normalized table record ceased to be the current record. It is used to ascertain which normalized table record should be used when multiple are available.
dss_current_flag	Used for normalized 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 normalized history tables. This column contains the version number of a normalized 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 normalized 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 normalized 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 normalized table in the left pane and selecting **Create (ReCreate)**.

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 normalized 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 to find such a column is to double-click on the list heading **Col name**; which will sort the column names into alphabetical order. Another double-click on the heading will sort the columns back into their create order.

The next section covers **Generating the Update Procedure** (see "**Generating the Normalized Update Procedure**" on page 412).

Generating the Normalized Update Procedure

Once a normalized 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 update of the normalized table.


Generating a Procedure

To generate a procedure, right-click on the normalized 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.

Procedure Type and Options

The first dialog displayed when generating a normalized table update procedure is the procedure types and options dialog:



The dialog box is titled "Define Normalized Procedure Type and Options". It contains the following sections and controls:

- Procedure Types**: A section with a "Procedure Type:" dropdown menu set to "Set" and a label "Previous build was : Set".
- Build Options**: A section containing several checkboxes and input fields:
 - ☐ Process in Batch
 - ☐ Distinct Data Select
 - ☐ Include Delete Before Insert
 - ☐ Include Initial Load Insert
 - ☒ Include Update Statement
 - ☒ Changed Rows Only (with a dropdown menu set to "Except")
 - Update Hint: (e.g. TABLOCK)
 - ☒ Include Insert Statement
 - ☒ New Rows Only (with a dropdown menu set to "Except")
 - Insert Hint: (e.g. TABLOCK)
 - ☒ Use Merge Statement for Update and Insert
 - ☐ Changed and New Rows Only
 - ☐ Delete Rows Not in Source
 - Merge Hint: (e.g. TABLOCK)

At the bottom right, there are "OK" and "Cancel" buttons.

Several different questions need to be answered 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.

Procedure Types

A number of different types of procedure generation are available. The procedure type is chosen from the procedure type drop-down list. A set based procedure performs one SQL statement to join all the source tables together and then insert this data into the normalized table. This is normally the fastest method of populating a table.


Build Options

The following build options are available:

Process in Batch


The **Process in Batch** checkbox 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).

The following dialog is the default for a non-date column:



The image shows a Windows-style dialog box titled "Batch Processing" with a red close button in the top right corner. The dialog has a light gray background. Inside, there is instructional text: "Select 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 then choose an attribute of that related table to group by." Below this text is a dropdown menu with "customer_code" selected and a downward arrow on the right. At the bottom right of the dialog are two buttons: "OK" and "Cancel".

The following dialog is displayed for date fields:



The image shows a 'Batch Processing' dialog box with a light blue title bar and a red close button. The main area is light gray. It contains two instructions: 'Select a field to batch process on.' and '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 then choose an attribute of that related table to group by.' Below these is a dropdown menu showing 'created_datetime'. Another instruction, 'Select the date format that will be used to group by for the processing', is followed by a second dropdown menu. This menu is open, showing options: 'created_datetime', 'Day', 'Month' (which is highlighted in blue), and 'Year'. At the bottom right are 'OK' and 'Cancel' buttons.

Batch Processing

Select 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 then choose an attribute of that related table to group by.

created_datetime

Select the date format that will be used to group by for the processing

created_datetime
Day
Month
Year

OK Cancel

Distinct Data Select

The **distinct data select** option ensures duplicate rows are not added to the normalized table. This is achieved by the word DISTINCT being added to the source select in the update procedure. The default for this field is off (i.e. duplicates are not removed).

Include Delete Before Insert

The **include 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 selected, the following dialog is displayed:

The dialog box is titled "Definition for a Delete Statement Before the Normalized Table Insert". It contains the following elements:

- Instructions:** "Select the Normalized Table column on which the delete will be based. All values in the source table for the selected column will be used to delete the corresponding records in the Normalized Table."
- Normalized Table Column:** A dropdown menu with "customer_code" selected.
- Source Table:** A dropdown menu with "stage_customer" selected.
- Source Table Column:** A dropdown menu with "customer_code" selected.
- Options:** Three checkboxes are present, all of which are unchecked:
 - ☐ Issue Warning if a Delete Occurs
 - ☐ Multiple Pass Delete (smaller transactions)
 - ☐ Truncate Table Before Insert
- Delete Where Clause:** A text area containing the SQL code: `WHERE normal_customer.customer_code IN ( SELECT DISTINCT stage_customer.customer_code FROM stage_cus`. To the right of the text area is the text: "The Where Clause can be edited, or directly entered."
- Word Wrap:** A checkbox labeled "Word Wrap Displayed Code" is unchecked.
- Buttons:** "OK" and "Cancel" buttons are located at the bottom right.

Use the wizard to build a 'Where' clause for the delete statement. Additional options included in this dialog are:

- ☐ Issue warning if a delete occurs
- ☐ Multiple pass delete - if this option is chosen, the delete is performed in a loop, 50000 rows at a time, until the required rows have been deleted.
- ☐ Truncate table before insert - if this option is chosen, the delete is ignore and the normalized table is truncated.

The default for this field is off (i.e. a delete statement is not added to the procedure).

Include Initial Load Insert

The **include initial load insert** option adds an additional insert statement to the update procedure that runs if the target normalized 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 off (i.e. an initial insert statement is not added to the procedure).

Insert Zero Key Record

This option applies only to datastore, normalized tables and set based dimensions. It allows you to include the zerokey (unknown record) for tables with an artificial key.

Include Update Statement with/without Minus Select using except/checksum

The **include update statement** option includes an update statement in the procedure to update changing rows in the normalized table. 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 normalized table 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).

Include Insert Statement with/without Minus Select using except/checksum

The **include insert statement** option includes an insert statement in the procedure to insert new rows in the normalized table. 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 Hint.

Enter a database hint to be used in the UPDATE statement.

Note: SQL Server only.

Insert Hint.

Enter a database hint to be used in the INSERT statement.

Note: SQL Server only.

Use Merge Statement for Update and Insert

The **Use Merge Statement for Update and Insert** option allows you to use the Merge statement instead of two separate Update and Insert statements.

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 *except* change detection or *checksum* change detection to work out what rows in the Normalized Object are new or require updating

Delete Rows Not in Source

The **Delete Rows not in Source** option allows you to delete all the rows which are not in the source but in the target table as part of the Merge statement. To avoid any accidental deletions, this option is de-activated if **Process in Batch** or **Changed and New Rows only** are selected.


Merge Hint

Enter a database hint to be used in the MERGE statement.

Note: SQL Server only.

Parameters variables

If QAD Data Warehouse Designer parameters exist in the metadata, the following dialog is displayed. Any parameters selected in this dialog (by moving them to the right), 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.



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. QAD Data Warehouse Designer 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: QAD Data Warehouse Designer parameters should be unique within the first 28 characters to avoid conflicting variables names.

See **Parameters** for more information on QAD Data Warehouse Designer Parameters

Business Key definition

The define business key dialog will appear asking for the business key that will uniquely identify each normalized table record. The source table from which the normalized 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.

Define Normalized Business Key Columns

Column List:

- customer_category_code
- customer_group_code
- customer_name
- customer_subgroup_code
- sales_source

Select the business keys that uniquely identify each record in the Normalized table.

Move them over to the business key list

Business Key List:

- customer_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 normalized 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.

Change Detection and History Tables

If the normalized table was defined as a history table then the following additional dialogs will appear. Normalized **history** tables are like slowly changing dimensions. The first dialog requests the selection of the columns to be managed as change detection columns. Select the required columns and click **OK** to proceed. In the example below the customer name is to be managed as a change detection column.

Define Normalized Change Detection Fields

Column List:

- customer_category_code
- customer_code
- customer_group_code
- customer_subgroup_code
- dss_current_flag
- dss_end_date
- dss_start_date
- dss_version
- sales_source

Select the fields that change detection is required for in the Normalized Update query.
Move them over to the change detection field list

Change Detection Field List:

- customer_name

>

<

OK Cancel

The following dialog will appear if one or more columns are chosen as change detection columns and if the user possesses a full license, enabling the additional procedure build options.

If one or more columns are chosen as change detection columns but the user only possesses a trial license, thus disabling the additional procedure build options, then the default values shown below will be used and the dialog will not be displayed.

When completed, click **OK** to continue.

History record settings

☐ Add NULL Support

☒ Update Current Records Only

Please enter the default start and end dates to be used.
Correctly fomatted dates or variables may be used.
Pressing Reset will reset the dates to the default settings.

Reset

Start Date for Initial Member (A):

CAST('01-JAN-1900' AS datetime)

End Date for Current Member (B):

CAST('31-DEC-2999' AS datetime)

Start Date for New Member Entry (C):

@v_current_date

End Date for Expiring Member Entry (D):

@v_current_date -0.00000005

Business ID	History Field	Start Date	End Date	Version	Current Flag
Project RED	Started	(A)	(D)	1	N
Project RED	Complete	(C)	(B)	2	Y

OK

Cancel

Add NULL Support

If the **Add null support** option is selected, then change detect column management will cater for null values in any change detect columns. If this option is not selected, and null values are present, errors may occur running the update procedure. The default for this value is off (i.e. nulls are not catered for).

Update Current Records Only

The **Update Current record only** option will only apply changes to non-change detect columns on the **current** record. If this option is not selected, all past (non-current) rows will also be updated to reflect changes to non-change detect columns. The default for this value is on (i.e. only the current record is updated).

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 correct data type. The default value provided will usually be cast to the correct datatype and can be treated as a template. The default value for this field is 1 January 1900.

End Date for Current Member

The **End date for current member** field contains the end 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 correct data type. The default value provided will usually be cast to the correct datatype and can be treated as a template. The default value 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 correct data type. The default value provided will usually be cast to the correct datatype and can be treated as a template. The default value for this field is the current date and time.

End Date for Expiring Member Entry

The **End date for expiring member entry** field contains the end date for any rows updated 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 correct data type. The default value provided will usually be cast to the correct datatype and can be treated as a template. The default value 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).

Reset


Click on the **Reset** button to return all fields in this dialog to their default values.

Converting an existing Normalized Table to a Normalized History Table

To convert a normalized table to a normalized history table, change the table type to **History**, select **(Build Procedure...)** from the **Update Procedure** drop-down list and click **OK**.

If the existing Normalized table is **NOT** to be dropped and recreated, then the following process should be followed:

- 1 Click **OK** on the Procedure Type dialog.
- 2 Click **OK** on the Business Key dialog.
- 3 Click **Alter** on the Adding Additional Columns dialog.
- 4 Click **Alter Table** on the Alter Table Commands dialog:




Note: The SQL in this dialog can be edited if required.

- 5 Click **OK** on all remaining dialogs.

Joining multiple source tables, where and group by clauses

Next the **Source table mapping** dialog appears. This dialog is used to join source tables, add 'Where' clauses and specify group by clauses.



The dialog box is titled "Source Table Mapping" and contains the following elements:

- Define the Joins (or edit the join 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:** A text area containing "[stage_customer]".
- Join Clause:** A text area containing "FROM [stage_customer] stage_customer".
- Where Clause:** An empty text area.
- Outer Join** and **Simple Join** buttons.
- Table A** and **Table B** dropdown menus, both set to "Table A".
- Column A** and **Column B** dropdown menus, both set to "Column A".
- Group By Clause:** An empty text area.
- ☐ Word Wrap Displayed Code checkbox.
- OK** and **Cancel** buttons.

To join two tables, start by selecting the two tables in the left box and then clicking 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_prod_group and load_prod_subgroup tables are joined by the column prod_group. For multiple column joins, repeat the join process.

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. 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.

Note: When QAD Data Warehouse Designer 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_prod_group' has had its column chosen and the column for the table 'load_prod_subgroup' is currently being chosen. This will result in the 'load_prod_group' 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 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. For example the Oracle outer join indicator (+) could be removed if required.

Where Clause

A 'Where' clause can be manually added to the **Where Clause** field in this dialog.

Group By Clause

A group by clause (and any other clauses that come after the 'Where' clause) can be added to the group by clause in this dialog.

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 to support each normalized table. These indexes will be added once the procedure has been built. An example of the type of indexes created is as follows:

Results			▼	🔍	×
📁	Object	Message	^		
📁	normal_customer	EXECUTE sp_addextendedproperty N'Comment', N'Date and time the row was updated in the data warehouse.', N'user', N'dbo', N'table', 'normal_customer', N'column', N'dss_update_time';			
📁	normal_customer	Create indexes on Normalized dbo.normal_customer completed successfully.			
📁	normal_customer	CREATE UNIQUE NONCLUSTERED INDEX normal_customer_idx_A ON dbo.normal_customer (code) WITH (SORT_IN_TEMPDB = OFF);			
📁	normal_customer	Procedure Completed			
📁	normal_customer	1 normal_customer updated. 6 records added. 0 records updated.			
			▼		
Results			Reports		

Additional indexes can be added, or these indexes changed. See the chapter on indexes for further details.

Normalized Table Artificial Keys

By default, normalized tables in QAD Data Warehouse Designer do not have an artificial (surrogate) key. Artificial keys can be added manually. The quickest way to do this is to add an extra column to the normalized table by using either **Add Column** or **Copy Column**. Edit the properties of the new column to have the correct name and order, source table and column, datatype, key type and flags. Specifically:

- ☐ The **Column Name** and **Source Column** should be the same.
- ☐ The **Source Table** should be empty.
- ☐ The **Data Type** should be:
 - DB2: integer generated by default as identity (start with 1, increment by 1)
 - Oracle: integer
 - SQL Server: integer identity(0,1)
- ☐ The **Key Type** should be 0.
- ☐ Only the **Numeric** and **Artificial Key** flags should be set on.

The following example shows a manually added artificial key column:

Normalized Column normal_customer.customer_key

Properties
Transformation

General

Table Name	normal_customer
Column Name	customer_key
Business Display Name	customer_key
Column Description	Manually added artificial key

Physical Definition

Column Order	10
Data Type	integer
Null Values Allowed	False
Default Value	

Meta Definition

Format	
Numeric	True
Additive	False
Attribute	False
End User Layer Display	False
Business Key	False
Artificial Key	True
Key Type (0,A,B,C,...)	0

Code Generation

Zero Key Value	0

Source Details

Source Table	(Empty)
Source Column	customer_key
Transformation	
Join	False

OK Cancel Help

The artificial key for a normalized table is set via a sequence in Oracle and an identity column in SQL Server and DB2. This artificial key normally, and by default, starts at one and progresses as far as is required.

A QAD standard for the creation of special rows in the normalized 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 normalized table record.
0	Used as a join to the normalized table when no valid join existed. It is the convention in the QAD generated code that any normalized 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 normalized 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.

Note: To have a surrogate key auto added for Normalized tables, see ***Global Naming of Key Columns*** (on page 57).

Normalized Table Column Properties

Each normalized 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 a dialog box titled "Normalized Column normal_customer.code". On the left is a "Properties" sidebar with a "Transformation" section. The main area displays a list of properties grouped into sections:

- General**
 - Table Name: normal_customer
 - Column Name**: code (highlighted)
 - Business Display Name: code
 - Column Description: Unique code to identify the customer. Forms a hierarchy with city and state.
- Physical Definition**
 - Column Order: 10
 - Data Type: numeric(6)
 - Null Values Allowed: True
 - Default Value:
- Meta Definition**
 - Format: #,##0
 - Numeric: True
 - Additive: True
 - Attribute: False
 - End User Layer Display: True
 - Business Key: True
 - Artificial Key: False
 - Key Type (0,A,B,C,...): A
- Source Details**
 - Source Table: load_customer
 - Source Column: code
 - Transformation:
 - Join: False

At the bottom, there is a "Column Name" section with the text: "Database-compliant name of the column. Dialog Opening Value: code".

At the bottom right are buttons: "< Update", "Update >", "OK", "Cancel", and "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 Data Warehouse Designer-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 data types for Oracle are integer, number, char(), varchar2() and date. For SQL Server common types are integer, numeric, varchar() and datetime. 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.

Artificial Key

Indicates whether the column is the artificial key. Only one artificial key is supported. This indicator is set by QAD Data Warehouse Designer 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

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 normalized 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. For previous value managed columns the source column is the column in the table whose previous value is being recorded.

Transformation

Transformation. [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 normalized 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.

Normalized Table Column Transformations

Each normalized 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:

Normalized Column normal_customer.code

Properties
Transformation

Target: code Paste

Source: load_customer.code Paste

Column Transformation Code (must execute within a SQL SELECT statement)

Function Set: SQL Server Standard Funct

Functions
Available Columns

☐ Word Wrap Displayed Code

Function Syntax:

Function Desc.:

< Update Update > OK Cancel Help

Note: Transformations are only put into effect when the procedure is re-generated.

See **Transformations** (on page 603) for more details.

Chapter 14

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. For simplicity the different types of fact table will be grouped together under the following headings:

- ☐ **Detail** (see "**Detail Fact Tables**" on page 434)
- ☐ **Rollup** (see "**Rollup or Combined Fact Tables**" on page 439)
- ☐ **KPI** (see "**KPI Fact Tables**" on page 443)
- ☐ **Snapshot** (see "**Snapshot Fact Tables**" on page 457) (essentially a type of Rollup fact table)
- ☐ **Slowly Changing** (see "**Slowly Changing Fact Tables**" on page 460)
- ☐ **Partitioned** (see "**Partitioned Fact Tables**" on page 462)

In This Chapter

Detail Fact Tables.....	434
Rollup or Combined Fact Tables	439
KPI Fact Tables	443
Snapshot Fact Tables.....	457
Slowly Changing Fact Tables.....	460
Partitioned Fact Tables	462
Fact Table Column Properties	479
Fact Table Column Transformations.....	484
Fact Table Language Mapping	485

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_forecast** detail fact table can be created by selecting the stage_forecast 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 commence.



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.

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.

A series of questions will be asked during the procedure generation based on the type of load information.

Define Fact Business Key Columns

The first dialog displayed when generating a detail fact table update procedure 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_id and order_line_no are selected for the business key list.

Define Fact Business Key Columns

Column List:

- base_currency_code
- base_sale_amount
- base_tax_amount
- bill_to_address_id
- created_datetime
- creating_employee_id
- customer_code
- dim_customer_key
- dim_order_date_key
- dim_product_key
- dim_ship_date_key
- gross_sale_amount
- invoice_date
- last_change_datetime
- last_change_employee_id
- line_type_code
- net_sale_amount
- order_date
- order_number
- order_status_code
- order_type_code
- payment_method_id
- posted_date
- product_id
- quantity
- sales_rep_id
- sales_source
- sales_tax_flag
- ship_date
- ship_to_address_id
- sold_to_address_id
- tax_amount
- tax_amount1
- tax_amount2
- tax_amount3
- tax_amount4
- transaction_employee_id

Select the business keys that uniquely identify each record in the Fact table. Move them over to the Business key list.

NOTE: Set based updates can only be selected if no business keys are defined.

Business Key List:

- order_id
- order_line_no

☐ Set Based Insert

☐ Allow Where Clause Editing

☐ Group By Dimension Keys

☐ Include Delete Before Insert

Insert Hint: (e.g. TABLOCK)

Update Hint: (e.g. TABLOCK)

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 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.

None of the columns chosen as the business key should ever contain a NULL value.

Allow Where Clause Editing

If the **Allow Where Clause Editing** option is selected, then the next dialog to be displayed is the **Provide Cursor Mapping** dialog, else this dialog is skipped.

Insert Hint

Enter a database hint to be used in the INSERT statement.

Note: SQL Server only.


Update Hint

Enter a database hint to be used in the UPDATE statement.

Note: SQL Server only.

Provide Cursor Mapping

This dialog is used to join source tables, add 'Where' clauses and specify group by clauses.



The 'Provide Cursor Mapping' dialog box is shown. It has a title bar with a close button. Inside, there is a text field for 'Cursor Name (if applicable):' with the value 'Load'. Below it is a label 'Edit the where clause if required.' followed by a large empty text area. To the left of this area is a list box labeled 'Source Tables:' containing the text 'stage_order_detail'. To the right is a large empty text area labeled 'From and Where Clause:'. Below these areas are three buttons: 'Outer Join', 'Simple Join', and a checked checkbox labeled 'ANSI join code generated'. At the bottom, there are two sets of dropdown menus: 'Table A' and 'Column A', and 'Table B' and 'Column B'. At the very bottom, there is a checkbox labeled 'Word Wrap Displayed Code' and two buttons: 'OK' and 'Cancel'.

As source table joins should have been performed in the stage table, see ***Generating the Staging Update Procedure*** (on page 349) for more details.

Since most of the real work is done in the stage table, this is all that is needed to build the Detail Fact Update Procedure.

Rollup or Combined Fact Tables


Rollup fact tables are typically the most heavily used fact tables in a data warehouse implementation. They provide a smaller set of information at a higher level of granularity. These fact tables typically have a monthly granularity in the time dimension and reflect figures for the month or month-to-date. They often combine multiple analysis areas (detail fact tables) into one super set of information. This allows simpler comparisons or joining of otherwise isolated pools of information.

An example may be the combination of daily sales, weekly demand forecasts, and monthly budgets, into one rollup fact table. Such a fact table provides the ability to compare monthly or month-to-date sales against budget and forecast.

These rollup/combined fact tables may and often do contain different levels of granularity, as well as different sets of dimensions. For example it may be possible to look at sales and forecasts by promotion, but the budget information does not have a populated promotion dimension, so we will not get a budgeted value for the promotion.

Creating Rollup Fact Tables

A rollup or combined fact table is typically created by dragging a fact table onto a fact table list. In the example screenshot below a double-click on the Fact Table object group under the Sales project produced a list in the middle pane showing the existing fact tables. A **fact_sales_analysis** rollup fact table is being created after having selected the **fact_sales_detail** table in the right panel, holding down the left mouse button and dragging the table into the middle pane. The initial object create dialog will appear. Note that the name implies that this is a rollup fact table. The name will be changed to the one chosen above. A list of columns is displayed and those not required are deleted (everything except additive measures and dimensions).



The screenshot shows a dialog box titled "Add a New Metadata Object". It contains instructions to define the type and name of the new object. The "Object Type" is set to "Fact Table" and the "Object Name" is "fact_sales_analysis". There are "ADD" and "Cancel" buttons at the bottom.

Drag and drop additional columns from other fact tables into the middle pane if required. Where columns have the same name these will need to be changed. For example if we acquired a quantity column from **fact_sales_detail** and a column of the same name from **fact_forecast** we may choose to change the names to **actual_quantity** and **forecast_quantity**.

Once all columns have been defined:

- 1 Create the table.
- 2 Edit the properties of the table and request that a new procedure be created.

Click **OK** and the dialog to create a template procedure will commence. Included in this process will be the definition of the date dimension granularity. Monthly rollup is the norm. To just combine fact tables with no rollup select the date dimension key as directed in the dialog. Other forms of dimension rollup will require some alteration to the produced template procedure.

Rollup Fact Columns

The columns for a rollup fact table are typically the dimensions, and only those key measures that are added. For example in the sales world, measures may simply be quantity, kilograms, \$value, forecast_quantity, forecast_\$value, budget_quantity, budget_\$value.

See the second tutorial for an example of rollup fact table creation.

Generating the Rollup Fact Update Procedure

Once a rollup 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.


Generating a Procedure

To generate a procedure, right-click on the fact table in the left pane and select **Properties** to edit the properties for the fact table. 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.

Define Rollup date dimension and column


The first dialog displayed when generating a detail fact table update procedure is the define rollup date dimension and column dialog:



The dialog box is titled "Define Rollup date dimension and column" with a red close button in the top right corner. It contains the following text: "Identify the Date dimension that will be used as the basis for the rollup. Then select the column that this fact table rolls up to (e.g. Cal_Month). If no date rollup is required. Select the date dimension key as the rollup column." Below this text are two dropdown menus. The first is labeled "Date Dimension" and has "dim_order_date" selected. The second is labeled "Rollup Column" and has "order_cal_month" selected. At the bottom are "OK" and "Cancel" buttons.


Data Warehouse Designer needs to know which date is driving the rollup and to which level of detail we wish to set as our base. In this example we have chosen **dim_order_date** as our date dimension and **order_cal_month** as our rollup column to set the level to summarize all transactions to. Click **OK**.

Data Warehouse Designer also needs to understand which date field in the actuals should be used to base the rollup; we have chosen **dim_order_date_key** and click **OK**.



The dialog box is titled "Define rollup date dimension and column" with a red close button in the top right corner. It contains the following text: "Fact: fact_order_detail. Select the date dimension key to be used for the rollup of this fact table." Below this text is a dropdown menu labeled "Date dimension Key" with "dim_order_date_key" selected. At the bottom are "OK" and "Cancel" buttons.

Once again, select the date dimension key and click **OK**.



The image shows a Windows-style dialog box titled "Define rollup date dimension and column" with a red close button in the top right corner. The dialog has a light gray background. Inside, the text "Fact: fact_forecast" is displayed. Below it, the instruction "Select the date dimension key to be used for the rollup of this fact table" is shown. A label "Date dimension Key" is positioned above a dropdown menu. The dropdown menu is open, showing the selected item "dim_forecast_date_key" with a small downward arrow on the right. At the bottom right of the dialog, there are two buttons: "OK" and "Cancel".

KPI Fact Tables

Key Performance Indicator (KPI) fact tables share a lot of similarities with the rollup fact tables in that they are typically at a monthly level of granularity and they combine information from multiple detail fact tables. They differ in that they provide specific sets of pre-defined information (KPIs). Such KPIs are often difficult to answer and require the implementation of some form of business rule. Also they often provide pre-calculated *year-to-date*, *same month last year*, and *previous year-to-date* values for easier comparisons. They will have a KPI dimension which provides a list of all the specific statistics or performance indicators we are tracking.

An example from our previous sales examples may be a set of KPIs dealing with customer acquisition and loss. We may have a specific KPI that records the number of new customers during the period and the \$value of their business during that period. Another KPI may record the percentage growth of all customers over the last six months. Another may show the number of customers who closed their accounts in the period, and the value of their business over the preceding 12 months.

KPI fact tables require the specific definition of each KPI.

It is strongly recommended that the KPI tutorial be undertaken. This tutorial will give an understanding of what can be achieved via a KPI fact table.

Creating the KPI Dimension

The KPI dimension essentially provides a record of each KPI or statistic we are tracking. There should be one row in this table for each such element. The KPI fact table requires that this dimension contain at the least an artificial dimension key and a business key. We will refer to the business key as the KPI Identifier.

The following example covers the manual creation of a basic KPI dimension:

- 1 Right-click on the **Dimension** object group in the left pane and select **New Object**.
- 2 Give the new dimension a name for example **dim_kpi**.
- 3 Select a **Normal** dimension and click **OK** on the Properties screen to define the dimension table.
- 4 Right-click on the dimension name in the left pane and select **Add Column**, add the three columns defined in the following table.

Field	Column 1	Column 2	Column 3
Column Name	dim_kpi_key	kpi_identifier	kpi_name
Business display	kpi artificial key	kpi unique id	kpi name
Datatype - Oracle - SQL Server/DB2	integer integer	varchar2(12) varchar(12)	varchar2(256) varchar(256)
Nulls	Clear	Clear	Checked
Key Type	0	A	Clear
Artificial Key	Checked	Clear	Clear
Primary Business	Clear	Checked	Clear


The first two columns are required, and should be set up as shown above, although a different data type can be used for the second column, and the names can be chosen to suit the requirements.

Creating KPI Fact Tables

A KPI fact table is normally created manually.

- 1 Right-click on the Fact Table object group in the left pane and select **New Object**.
- 2 A dialog displays as shown below.

Note: The default object type is **Fact Table**. This must be changed to **KPI Fact Table** and a table name entered.



Adding columns

Once created the KPI fact table will need to have columns added. New columns can be added by dragging columns from the browser window, or by manually adding from the right mouse menu. Dimension keys should normally be dragged to ensure that all the properties are correctly set.

The columns for a KPI fact table are typically the dimensions and generic measures. For example valid measures could be **kpi_count**, **kpi_quantity** and **kpi_value**. The interpretation of these measures will depend on the KPI or group of KPIs being queried. As with the rollup fact table it is normal and desirable to keep the number of measures as small as possible. Attributes should not be included in these fact tables. These measures must allow Null values for the generated update procedure to function correctly.

Notes:

1. The kpi measure column order is important. For example if we have mtd count, quantity and value in that order then the prev month, prev ytd and ytd figures also need count first followed by quantity and finally value. If this order is not adhered to then the generated procedure will load the wrong values into the wrong columns.
 2. A kpi fact table should also have the column dss_update_time added. This column is required if the kpi table is to be used to populate higher level fact tables, aggregates or cubes. This column should be populated with the current date time whenever the particular row is updated.
-

Setup of KPI Fact Tables

The set-up of a KPI fact table is somewhat more complex than the other types of fact table. Once all columns have been defined the following steps need to be taken:

- 1 Create and populate a KPI dimension that provides at least an id (unique identifier) for each KPI and a description. The key from this dimension must be one of the columns in the fact table.
- 2 Physically create the KPI fact table in the database once all columns including the KPI dimension key have been added.
- 3 Define the column type and default value for ALL columns in the table. There must be at least one date dimension key and one KPI dimension key. These attributes are set by selecting a column and modifying its properties.
- 4 Create the update procedure for the fact table.
- 5 Set-up each individual KPI.

KPI Setup

KPI Dimension

The KPI dimension provides a means of selecting a specific KPI. This dimension needs to be built in advance of the KPI fact table. It does not need to be populated, as the KPI fact table will add entries to this table as required. A key and some form of unique identifier for the KPI is required, e.g., they may be numbered or given short names.

Create the KPI Fact Table

Right-click on the KPI table and select **Create/Recreate** to physically create the table in the database.

Defining column types and default values

- 1 List all the columns for the KPI fact table.
- 2 Amend the properties for each column. The column types can be one of the following:

Type	Description
Dimension Key	Any dimension that is not the time or KPI dimension
Time Dimension Key	Time dimension key
KPI Dimension Key	KPI dimension key
Month-to-date Measure	Used to identify the normal KPI measures. These are calculated to populate the KPI during each run.
Year-to-date Measure	Used to identify year-to-date measures for the KPI. Optionally calculated or derived from existing rows and measures in the KPI fact table.
Last year same month Measure	Used to identify measures for same month last year comparisons. Derived from existing rows and measures in the KPI fact table.
Previous year-to-date Measure	Used to identify previous year-to-date measures. Derived from existing rows and measures in the KPI fact table.
Fact Table Key	Not normally used
Date	Not normally used
Attribute	Not normally used
Measure	Not normally used




TIP: If an **Attribute** or **Measure** column is specified then this column will just be set to its default value. In this way an update_time column can be set by making its default value **sysdate**.

- 3 For each column set a default value.
 - For measures and dimension keys this would normally be zero (0). Nulls can be represented as two single quotation marks, i.e. ''
 - **sysdate** can be used for a default date value
 - These default values populate the columns where and when no other information is available. For example if a KPI only makes use of 4 of the 6 dimensions available then the remaining 2 dimensions will be set to their default values.

Create the 'Update' procedure


Right-click on the KPI fact table name and select **Properties**. In the Properties dialog box, set the update procedure list to the **(Build Procedure...)** option. Click **OK** to initiate a series of dialogs to create a procedure used to update the kpi fact table.

This first two dialogs ask for parameter names to define the start and end periods for which the KPIs will be updated. The default naming convention is fact_table_name_START and fact_table_name_END. Accept these default names if possible. These two parameters will need to be created, and will need to be updated once the KPIs have been defined and tested. Under normal operation these parameters should probably be set daily from the daily_time_roll procedure.



A dialog box titled "fact_sales_kpi KPI parameter definition" with a red close button in the top right corner. The main text area contains the instruction: "Enter the parameter that will identify the update start period. The daily_time_roll procedure will need to include this parameter." Below the text is a text input field containing the text "FACT_SALES_KPI_START". At the bottom of the dialog are two buttons: "OK" and "Cancel".

The date dimension and column for rollup is required. Normally this would be cal_month or fin_month. This is based on the assumption that most KPIs are monthly or annual. Subsequently a rollup year and date will be required from the same date dimension.



A dialog box titled "Define Rollup date dimension and column" with a red close button in the top right corner. The main text area contains the instruction: "Identify the date dimension that will be used as the basis for the rollup. Then select the column that this fact table rolls up to (e.g. Cal_Month)". Below the text are two dropdown menus. The first dropdown is labeled "Date Dimension" and has "dim_order_date" selected. The second dropdown is labeled "Rollup Column" and has "order_cal_month" selected. At the bottom of the dialog are two buttons: "OK" and "Cancel".

The KPI dimension is required. Also the identifier column (not the key) that uniquely identifies the KPI. This would normally be a number or short name (max 64 chars).

Select the KPI dimension and identifier

Select the KPI dimension, and the column that provides the identifier for the KPI

KPI Dimension: dim_sales_kpi

KPI Identifier: KPI_ID

OK Cancel

The procedure is then created and compiled.

Define each KPI

Once the procedure has successfully compiled we are ready to start setting up each of the KPIs. Right-click on the KPI table name to get the popup menu. This menu is similar to the other fact table menus except for the additional options to display and add KPIs. A display of the KPIs will list any KPIs that are currently defined. The result set is as appears below:

KPI	Name	Active	Method	Description
Activity	Customer activit...	Enabled	Statement	Provides a count of active customer in a particular month, t
Sales	Customer sales (...)	Enabled	Statement	Provides a count and dollar value of sales in a particular mo

The 'Id' is the same unique identifier that will be found in the kpi dimension. The Name is a name entered in the properties of the KPI and does not relate to the dimension in any way. The Active flag indicates if the KPI will be processed when the **Update** or **Process** menu options are selected. Each KPI can be defined as a SQL Statement, a Procedural Block or a Procedure. The current method used to update the KPI is one of these three and is displayed under Method. The description is a comment or description field entered in the properties of the KPI and in no way relates in the KPI dimension.

A right-click anywhere in the List Window will allow a KPI add. If a right mouse button is depressed while positioned over the 'Id' then a menu popup will provide options to copy, modify, delete, add and view the KPI statement.

When the add KPI option is selected a Properties screen will appear, as shown below. The id and name fields are required. If the **Active** checkbox is checked the procedure will be enabled for update.

The **Select Columns** button allows the selection of the columns that this KPI will update during the month to date update pass. The date and KPI dimensions will be updated for you so are not provided as choices. The only choices provided are the 'Month to date measures' and the other 'Dimensions'. Select those measures and dimensions that will be updated. All non selected measures and dimensions will be set to their default values. The resultant choice is shown in the Update Columns window. This window should only be updated via the **Select Columns** button.

KPI Definition

Properties

MTD Statement

YTD Statement

Procedural Block

KPI Identifier: SaleVol

KPI Name: Sales Volumes

Active: ☒

Description: These are the sales volumes / quantities that are ordered for the products

Month to Date Update Columns: dim_product_key,sales_quantity_mtd

Select Columns

Year to Date Update Columns: dim_product_key,sales_quantity_ytd

Select YTD Cols

OK Cancel Help

Once these initial values are defined the **Statement** tabs can be selected to define the update SQL statement. If a procedural block or procedure is required for the update (i.e. too complex for a single SQL statement) then select the Procedural Block tab.

The KPI Statement

The **Statement** tabs present edit windows that allow the input of a select statement that will return the values for the columns being updated. They must return the values for the columns selected in the Properties tab in the order shown in the **Update Columns** window. An example of a completed KPI statement follows:

The screenshot shows a window titled "KPI Definition" with a sidebar on the left containing four tabs: "Properties", "MTD Statement", "YTD Statement", and "Procedural Block". The "MTD Statement" tab is selected. The main area of the window contains a text box with the following SQL statement:

```
-- Insert into fact_sales_kpi(  
-- dim_product_key,sales_quantity_mtd)  
Select dim_product_key,  
       sum(quantity)  
From fact_order_detail  
Where dim_order_date_key >= $KPI_START_YYYYMMDD$  
And dim_order_date_key <= $KPI_END_YYYYMMDD$  
group by dim_product_key
```

Below the text box are three buttons: "OK", "Cancel", and "Help".

The first two lines are comments inserted (when modifying) to provide a reminder of the columns and order required by the select statement. Anything prior to the initial select statement is ignored.

The statement must make use of at least one of the variables provided to select the appropriate period. This variable will be replaced with the values for the period being updated at run time. These values are derived from the start and end parameters defined for the KPI fact table.

In order to group by the non selected columns and any attributes etc.QAD Data Warehouse Designer must be able to break this statement into its component Select, From, Where, Group By, and Having components (where present). If a complex statement is used that has more than one of these key words then the primary one must be in UPPER case. For example if a sub query is being used then the main SELECT and FROM key words must be in upper case and must be the only occurrence of these key words in uppercase. QAD Data Warehouse Designer will check and issue a warning if it is unable to resolve the main key words.

To test a statement, right-click on the KPI Id and select **View Statement**. This will provide a window with the full statement with the default columns inserted and the parameters converted. This statement can be cut and run in Sql Admin (shipped with QAD Data Warehouse Designer) or in a database specific utility for testing.

Procedural Block

The **Procedural Block** tab presents an edit window that allows the input of an Oracle anonymous block. If there is anything in this window it will be actioned instead of the statement. A button in the lower right allows a rudimentary template block to be built. If used this block must do the actual insert statement rather than just the select. It must also handle ALL dimension, month to date measure and non standard columns. In addition, it will need to put default values in columns not actively being updated. The KPI dimension and date dimension keys must also be updated. The output from a 'View Statement' option for another KPI will give an indication of what the block must do. The same month last year, year to date and previous year to date measures will still be handled automatically. The same \$KPI_ variables as shown above the statement edit window can and must be used.

Procedure

If for some reason it is not possible to implement the KPI in either the statement or procedural block then a procedure may be created. The name of such a procedure can be entered in the window at the bottom of the Procedural Block tab. If any value is in this window then the procedure will take precedence over both the statement and block.

This procedure must adhere to the standard set of parameters as with all other procedures to be executed by the scheduler. It must undertake the same tasks as defined above in the Procedural Block section, and as with those tasks it does not need to set the year to date, last year same month and previous year to date values.

Copying KPIs

Right-click on the KPI Id and select **Insert Copy** to insert an exact duplicate of the KPI. This copy can then have its Id and Name changed along with its statement to create a new KPI. This is particularly useful when only small differences exist between a number of KPIs.

Testing/Development.

During testing it would be normal to DISABLE all KPIs except the one currently being developed. In this way the update or process options will only utilize that KPI. As mentioned above the 'View Statement' option of the KPI popup menu displays, and allows the cutting of, the statement that will be executed to acquire the KPI information. This statement can be tested in sql*plus or some other environment if problems are encountered.

Generating the KPI 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.


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.

A series of questions will be asked during the procedure generation based on the type of load information.

KPI parameter definition

The first dialog displayed when generating a detail fact table update procedure is the KPI parameter definition dialog:



fact_sales_kpi KPI parameter definition

Enter the parameter that will identify the update start period
The daily_time_roll procedure will need to include this parameter

FACT_SALES_KPI_START


OK Cancel

The **KPI parameter definition** dialog box will prompt for a parameter to hold the start period that we are interested in. This period sets the first period that will be processed when the KPI fact table is updated. Once entered, click **OK**.

NOTE: If this parameter is not set, the default current period will be set from dim_date. The format is YYYYMM.

TIP: The **dim_date** custom update procedure **daily_date_roll** can be updated to set these parameters automatically if you do not wish to control them manually. This is good practice


Similarly, a **KPI parameter definition** dialog will prompt for a parameter to hold the end period. Once entered, click **OK**.



The dialog box is titled "fact_sales_kpi KPI parameter definition" with a red close button. It contains the text: "Enter the parameter that will identify the update end period. The daily_time_roll procedure will need to include this parameter." Below the text is a text input field containing "FACT_SALES_KPI_END". At the bottom are "OK" and "Cancel" buttons.


Define Rollup date dimension and column

As with Rollup tables, we now need to specify which date in the transactional data we will use to control the rollups and to which level – in this example, for the Monthly level rollup, we choose **dim_order_date** as our **Date Dimension** and **order_cal_month** as our **Rollup Column** and click **OK**.




The dialog box is titled "Define Rollup date dimension and column" with a red close button. It contains the text: "Identify the date dimension that will be used as the basis for the rollup. Then select the column that this fact table rolls up to (e.g. Cal_Month)". Below the text are two dropdown menus. The "Date Dimension" dropdown is set to "dim_order_date" and the "Rollup Column" dropdown is set to "order_cal_month". At the bottom are "OK" and "Cancel" buttons.

Once again for our Yearly rollup, in this example we choose **dim_order_date** as our **Date Dimension** and **order_cal_year** as our **Rollup Year** and click **OK**.



The dialog box is titled "Define Rollup date dimension and column" with a red close button in the top right corner. It contains a text label "Select the column used to roll up the year (e.g. Cal_Year)". Below this, there are two dropdown menus. The first is labeled "Date Dimension" and has "dim_order_date" selected. The second is labeled "Rollup year" and has "order_cal_year" selected. At the bottom right, there are two buttons: "OK" and "Cancel".


The next screen sets the Daily level for the update procedure to understand the lower level detail, we choose **dim_order_date** as our **Date Dimension** and **order_date** as our **Date Column** and click **OK**.



The dialog box is titled "Define Rollup date dimension and column" with a red close button in the top right corner. It contains a text label "Select the column that provides a day date (e.g. Calendar_date)". Below this, there are two dropdown menus. The first is labeled "Date Dimension" and has "dim_order_date" selected. The second is labeled "Date column" and has "order_date" selected. At the bottom right, there are two buttons: "OK" and "Cancel".

Select the KPI dimension and identifier

Finally we need to specify where the KPIs are defined and how they are defined, so in this example we choose **dim_sales_kpi** as the **KPI Dimension** and **KPI_ID** as the **KPI Identifier**. Click **OK**.



The dialog box is titled "Select the KPI dimension and identifier" with a red close button in the top right corner. It contains a text label "Select the KPI dimension, and the column that provides the identifier for the KPI". Below this, there are two dropdown menus. The first is labeled "KPI Dimension" and has "dim_sales_kpi" selected. The second is labeled "KPI Identifier" and has "KPI_ID" selected. At the bottom right, there are two buttons: "OK" and "Cancel".

Snapshot Fact Tables

Snapshot fact tables provide an 'as at' view of some part of the business. They always have a time dimension that represents the 'as at' date. For example, we may have a weekly or monthly snapshot of our inventory information. Such a snapshot would allow us to see what inventory we had at a particular point in time, and allow us to compare our holdings over time. Snapshots are often used when dealing with people profiles or client/customer profiles to enable us to view the state of our relationship with an entity at a given point in time e.g. to see a monthly trend of a client's outstanding balance.

Creating Snapshot Fact Tables

The creation of snapshot fact tables is very dependent on the method used to create the snapshot. In some cases they may be created in exactly the same way as detail fact tables. Such cases are typically where the OLTP application provides the 'as at' state. For example an inventory system may provide a weekly stock take table which is the basis for our snapshot fact table. In such cases the table is created as per a detail fact table, but with additional tailoring around the time dimension. In other situations the snapshot is built from an existing transactional fact table, where all information up to a specific date is summarized or frozen to create the snapshot. In these instances the snapshot can be created as a rollup fact table. Regardless of whether a detail or rollup method is used for the creation of the snapshot table, a significant amount of work will be required to tailor the update procedure to reflect the particular requirements of the snapshot table.

Snapshot Fact Columns

The columns for a snapshot fact table vary to a large degree. A date dimension is always present to reflect the 'as at' date. Some snapshots may have a large number of measures reflecting \$values quantities etc, whereas others may be rich in attribute (non additive) facts.

Generating the Snapshot Fact Update Procedure

Once a snapshot 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.

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.

A series of questions will be asked during the procedure generation based on the type of load information.

If a **detail** method was used for the creation of the snapshot table, refer to ***Generating the Detail Fact Update Procedure*** (on page 436)

However if a **rollup** method was used for the creation of the snapshot table, refer to ***Generating the Rollup Fact Update Procedure*** (on page 441)

Slowly Changing Fact Tables

Slowly changing fact tables provide a complete history of changes for some part of the business. Typically such fact tables are used for situations where a relatively small number of changes occur over time, but we need to track all of these changes. A good example is an inventory stock holding that does not change very often. In such a case a slowly changing fact table can record both the current inventory levels as well as providing a means of seeing the state of inventory holdings at any point in time.

Creating Slowly Changing Fact Tables

QAD Data Warehouse Designer does not provide any automated way for creating a slowly changing fact table. The best method is to proceed as follows.

- 1 Drag and drop the stage table into a dimension target and create a dummy slowly changing dimension with the name we will be using for our fact table. (i.e. fact_...).
- 2 Answer the questions to the pop-up dialog boxes. Select a join if asked, but do not bother to join any of the tables. Select the columns to be managed as slowly changing. Normally these will be the measures such as quantity and all of the dimension keys. The generated procedure will fail to compile.
- 3 Inspect the additional columns added for the dimension and make a note of them. Specifically dss_end_date, dss_start_date, dss_current_flag and dss_version.
- 4 Inspect the indexes created to help manage the slowly changing dimension and make a note of their columns.
- 5 Delete the dummy dimension you have created.
- 6 Delete the get_..._key procedure created for the dimension. Do not delete the update procedure.
- 7 Create a detail fact table in the normal manner.
- 8 Add in the special columns as created for the dimension. Namely dss_end_date, dss_start_date, dss_current_flag and dss_version.
- 9 Recreate the table
- 10 Create indexes as per the dimension table to help maintain the slowly changing fact table.
- 11 Modify the update procedure for the fact table. Bring up the old version of the procedure that was created for the dimension of the same name. This will have been automatically versioned out when the latest procedure was created and can be seen through the procedure viewer.
- 12 This new procedure will need to merge the code from the old procedure, which will provide a guide in how to build the code for a slowly changing fact table.

Partitioned Fact Tables

Partitioned fact tables are normally used for performance reasons. They allow us to break a large fact table into a number of smaller tables (partitions).

Note: Partitioned fact tables are not directly supported in SQL Server 2000. A view must be created over a number of tables to create the concept of a logical partitioned fact table.

QAD Data Warehouse Designer supports partitioned fact tables. The following section explains their creation and support.

Creating a Partitioned Fact Table in SQL Server

In QAD Data Warehouse Designer a partitioned fact table can only be created from an existing fact table. The process therefore is to create a fact table normally, including the update procedure and indexes. The table type is then changed to a partitioned table type and QAD Data Warehouse Designer will assist in creating the exchange table, modifying the indexes and building the procedure to handle the update of the partitions.

Note: Partitioned fact tables are not directly supported in SQL Server 2000. A view must be created over a number of tables to create the concept of a logical partitioned fact table.

Create a detail or rollup/combined fact table

Create the fact table that you wish to partition in the normal way. You must create an update procedure but do not need to load any data. The process of creating the update procedure will also build any indexes for the fact table. We will be using a partition exchange table to move data to and from the partitioned fact table. This exchange table must have the same columns, column order and indexes as the fact table. It is therefore well worth getting the fact table into its final form before partitioning it. Add any indexes to avoid extra work later.

Convert to a partitioned fact table.

Edit the properties of the fact table to be converted and change the table type to either **Detail partitioned** or **Rollup partitioned**. See example below. These two options reflect two different methodologies in the updating of a partitioned fact table. The first (Detail) method moves a partition from the fact table into the exchange table. The data in the exchange table is then updated with any changes or additions. When the update process has completed the exchange table data is moved back into position in the fact table. During the update process the partition in the fact table will be empty. The second (Rollup) method builds all information for a period in the exchange table. The exchange table is then used to replace the relevant period in the fact table. The fact table remains in a queryable state throughout the process.

Fact Table fact_order_detail

Properties

Storage

Statistics

Language Mapping

Purpose

Concept

Grain

Examples

Usage

Notes

Table Name: fact_order_detail

Table Type: Detail

Unique Short Name: (maximum 12 characters) fact_order_d

Business Display Name (EUL): fact_order_detail

Description:

Update Procedure: update_fact_order_detail Edit Rebuild Cursor Based Update

Custom Procedure: custom_fact_order_detail Edit

Get Key Function: [None] Edit

Mnemonic (EUL):

Timestamps

Metadata Structure Changed:	Database Created:	Database Altered:
2013-04-11 15:31:53.381000	2013-04-11 15:31:57.975000	2013-04-11 15:31:57.975000

OK Cancel Help


A dialog will pop-up as follows, asking for confirmation that the table is to be converted to a partitioned fact table. If **Yes** is chosen then the existing fact table will be renamed to become the partition exchange table. Any current data will remain in this table.

A change to a partitioned Fact Table has been detected
WARNING: The table must be re-created to change to a partitioned table.

If assistance is chosen the existing table will be renamed to allow creation of the new Fact Table.
 Would you like assistance in converting this table to a partitioned table and creating a partition exchange table (containing the existing data)?

Yes No Cancel

The parameters for the partitioned fact table are then prompted for:



The existing fact table will be renamed to act as an exchange partition.
If no exchange partition is required then this table can be deleted.

Note: The existing fact table data will be available under this new name.
The update procedure may however clear this table.
Enter the name for the exchange partition table:

part_order_detail

Automatic code generation is only supported on date dimension keys.
Select a date dimension key to base the partitioning on:

dim_order_date_key

Select either day, month or year as the partition granularity:

Month

Select the corresponding date dimension column.
Choose the date for day grain. (E.g. cal_year, cal_month, or calendar_date):

order_cal_month

The first partition will cater for all periods prior to and including the specified period.
New partitions will be created dynamically as required, so the final partition can be the same as the initial partition or perhaps the current period.

First partition (YYYYMM): 200101

Final partition (YYYYMM): 200201

OK Cancel

The partition exchange table name is prompted for. This table is identical to the fact table and is used to swap partitions to and from the fact table. Select a name that indicates that the table is not a fact table as such.

The fact table will be partitioned by the date dimension key. You need to specify the date dimension key that will be used as the basis for the partitioning of the fact table. It is assumed that this key has the standard QAD format of YYYYMMDD.

A partition granularity of day, month or year must then be selected along with the corresponding column in the date dimension. For a daily partition a column with a date data type must be chosen. The date dimension key will actually be used for the partitioning but the date column is required. For a monthly partition a column with the format YYYYMM must be chosen. In the standard QAD date dimension this is the column cal_month. For a yearly partition a column with the format YYYY must be chosen. In the standard QAD date dimension this is the column cal_year.

The first and last partition should then be selected. The last partition is not that important as additional partitions will be created as required. Normally just select a partition a few on from the first. The first partition is however important as any data that is dated prior to the first partition will be loaded into this first partition. The partition must be entered in the format shown. For example monthly partitions require the format YYYYMM.

Once the **OK** button is clicked the conversion process will populate the **Table Properties Clauses** field of the tables properties. Note that this field is limited in size, so if too many partitions are chosen only the first few will be added to this field. The other partitions will be created dynamically by the update procedure.

fact_order_detail Table Definition

Properties | Storage | Language Mapping | Purpose | Concept | Grain | Examples | Usage | Notes

Table Name: fact_order_detail Table Type: Detail partitioned

Unique Short Name: (maximum 22 characters) fact_order_detail

Business Display Name (EUL): fact_order_detail

Description: Use this fact table to track our Sales Orders

Update Procedure: update_fact_order_detail Edit Rebuild Cursor Based Update

Custom Procedure: custom_fact_order_detail Edit

Table Properties Clauses: (e.g. Partition)

Partition Function Command: CREATE PARTITION FUNCTION function_part_order_detail(int) AS RANGE LEFT FOR VALUES (20010131,20010231,20010331,20010431,20010531,20010631,20010731,20010831,20010931,20011031,20011131,20011231)

Partition Scheme Command: CREATE PARTITION SCHEME scheme_part_order_detail AS PARTITION function_part_order_detail ALL TO (ReportingTables)

Get Key Function: (None) Mnemonic (EUL):

Timestamps

Metadata Structure Changed:	Database Created:	Database Altered:
2013-01-18 15:54:12.187	2013-01-18 15:54:18.850	2013-01-18 15:54:18.850

OK Cancel Help

Click the **OK** button on the table **Properties** dialog to begin the process of converting the fact table, creating the new fact table and building the new update procedure. Examine the results screen to see what has been done.

The partition table will now have any current data, to populate the fact table, copy the data from partition table to the fact table, drop and recreate fact index. The date dimension key used to base the partitioning on is not allowed to have nulls. QAD Data Warehouse Designer has modified the table property, however the partition table needs to be validated and recreated to be valid. Now the fact table can be processed.

The partitioned fact table update procedure

The generated procedure will handle the creation of new partitions as required. Any data for periods prior to the first partition is placed in the first partition. If data is received for a period after the last partition then new partitions are created. Empty partitions are created if necessary to skip periods where no data is present. For example, if we have a table partitioned by month and the latest partition is 200204 (april/2002). If data is then received for say september 2002 the update procedure will build partitions for may, june, july, august and september. There is a variable at the start of the update procedure called `v_max_skip_periods`. This variable is by default set to 12. This defines the maximum number of continuous partitions we will leave empty. From our previous example if our latest partition was april 2002 and we received data for july 2003 with no interim data then the update procedure will fail as it would have to skip 14 partitions. This check is present to prevent erroneous data from creating hundreds of partitions into the future.

Creating a Partitioned Fact Table in Oracle

In QAD Data Warehouse Designer a partitioned fact table can only be created from an existing fact table. The process therefore is to create a fact table normally, including the update procedure and indexes. The table type is then changed to a partitioned table type and QAD Data Warehouse Designer will assist in creating the exchange table, modifying the indexes and building the procedure to handle the update of the partitions.

Create a detail or rollup/combined fact table

Create the fact table that you wish to partition in the normal way. You must create an update procedure but do not need to load any data. The process of creating the update procedure will also build any indexes for the fact table. We will be using a partition exchange table to move data to and from the partitioned fact table. This exchange table must have the same columns, column order and indexes as the fact table. It is therefore well worth getting the fact table into its final form before partitioning it. Add any indexes to avoid extra work later.

Convert to a partitioned fact table.


Edit the properties of the fact table to be converted and change the table type to either **Detail partitioned** or **Rollup partitioned**. See example below. These two options reflect two different methodologies in the updating of a partitioned fact table. The first (Detail) method moves a partition from the fact table into the exchange table. The data in the exchange table is then updated with any changes or additions. When the update process has completed the exchange table data is moved back into position in the fact table. During the update process the partition in the fact table will be empty. The second (Rollup) method builds all information for a period in the exchange table. The exchange table is then used to replace the relevant period in the fact table. The fact table remains in a queryable state throughout the process.

The screenshot shows the 'Fact Table fact_order_detail' properties dialog. The 'Table Type' dropdown is open, showing the following options: Detail, Rollup, Work table, Detail partitioned (selected), and Rollup partitioned. The dialog includes the following fields and controls:

- Table Name:** fact_order_detail
- Unique Short Name:** fact_order_d (maximum 12 characters)
- Business Display Name (EUL):** fact_order_detail
- Description:** (empty text area)
- Update Procedure:** update_fact_order_detail (dropdown) with **Edit** and **Rebuild** buttons.
- Custom Procedure:** custom_fact_order_detail (dropdown) with **Edit** button.
- Cursor Based Update:** (checkbox, currently unchecked)
- Get Key Function:** (None) (dropdown) with **Edit** button.
- Mnemonic (EUL):** (empty text field)
- Timestamps:**
 - Metadata Structure Changed:** 2013-04-11 15:31:53.381000
 - Database Created:** 2013-04-11 15:31:57.975000
 - Database Altered:** 2013-04-11 15:31:57.975000

At the bottom of the dialog are **OK**, **Cancel**, and **Help** buttons.

A dialog will pop-up as follows, asking for confirmation that the table is to be converted to a partitioned fact table. If **Yes** is chosen then the existing fact table will be renamed to become the partition exchange table. Any current data will remain in this table.



The parameters for the partitioned fact table are then prompted for:

A dialog box titled "Partitioning attributes" with a blue border and a red close button in the top right corner. It contains the following text: "The existing fact table will be renamed to act as an exchange partition. If no exchange partition is required then this table can be deleted.", "Note: The existing fact table data will be available under this new name. The update procedure may however clear this table.", and "Enter the name for the exchange partition table:". Below this is a text input field containing "part_order_detail". Then it says: "Automatic code generation is only supported on date dimension keys. Select a date dimension key to base the partitioning on:". Below this is a dropdown menu showing "dim_order_date_key". Then it says: "Select either day, month or year as the partition granularity:" and "Select the corresponding date dimension column. Choose the date for day grain. (E.g. cal_year, cal_month, or calendar_date):". Below these are two dropdown menus: the first shows "Month" and the second shows "order_cal_month". Then it says: "The first partition will cater for all periods prior to and including the specified period. New partitions will be created dynamically as required, so the final partition can be the same as the initial partition or perhaps the current period." Below this are two text input fields: "First partition (YYYYMM):" with "200101" and "Final partition (YYYYMM):" with "200201". At the bottom are two buttons: "OK" and "Cancel".

The partition exchange table name is prompted for. This table is identical to the fact table and is used to swap partitions to and from the fact table. Select a name that indicates that the table is not a fact table as such.

The fact table will be partitioned by the date dimension key. You need to specify the date dimension key that will be used as the basis for the partitioning of the fact table. It is assumed that this key has the standard QAD format of YYYYMMDD.

A partition granularity of day, month or year must then be selected along with the corresponding column in the date dimension. For a daily partition a column with a date data type must be chosen. The date dimension key will actually be used for the partitioning but the date column is required. For a monthly partition a column with the format YYYYMM must be chosen. In the standard QAD date dimension this is the column cal_month. For a yearly partition a column with the format YYYY must be chosen. In the standard QAD date dimension this is the column cal_year.

The first and last partition should then be selected. The last partition is not that important as additional partitions will be created as required. Normally just select a partition a few on from the first. The first partition is however important as any data that is dated prior to the first partition will be loaded into this first partition. The partition must be entered in the format shown. For example monthly partitions require the format YYYYMM.

Once the **OK** button is clicked the conversion process will populate the **Table Properties Clauses** field of the table's properties. Note that this field is limited in size, so if too many partitions are chosen only the first few will be added to this field. The other partitions will be created dynamically by the update procedure.

Fact Table fact_order_detail

Table Name: fact_order_detail Table Type: Detail partitioned

Unique Short Name: (maximum 22 characters) fact_order_detail

Business Display Name (EUL): fact_order_detail

Description:

Update Procedure: update_fact_order_detail Edit Rebuild Cursor Based Update

Custom Procedure: custom_fact_order_detail Edit

Table Properties Clauses: (e.g. Partition)

PARTITION BY RANGE (dim_order_date_key)
[partition p_200101 values less than (20010201)
, partition p_200102 values less than (20010301)
, partition p_200103 values less than (20010401)
, partition p_200104 values less than (20010501)
, partition p_200105 values less than (20010601)

Partition Function Command:

Partition Scheme Command:

Get Key Function: (None) Edit Mnemonic (EUL):

Timestamps

Metadata Structure Changed: 2013-04-11 14:46:27 Database Created: 2013-04-11 14:46:38 Database Altered: 2013-04-11 14:46:38

OK Cancel Help

Click the **OK** button on the table **Properties** dialog to begin the process of converting the fact table, creating the new fact table and building the new update procedure. Examine the results screen to see what has been done.

The partitioned fact table update procedure

The generated procedure will handle the creation of new partitions as required. Any data for periods prior to the first partition is placed in the first partition. If data is received for a period after the last partition then new partitions are created. Empty partitions are created if necessary to skip periods where no data is present. For example, if we have a table partitioned by month and the latest partition is 200204 (april/2002). If data is then received for say september 2002 the update procedure will build partitions for may, june, july, august and september. There is a variable at the start of the update procedure called `v_max_skip_periods`. This variable is by default set to 12. This defines the maximum number of continuous partitions we will leave empty. From our previous example if our latest partition was april 2002 and we received data for july 2003 with no interim data then the update procedure will fail as it would have to skip 14 partitions. This check is present to prevent erroneous data from creating hundreds of partitions into the future.

Creating a Partitioned Fact Table in DB2

In QAD Data Warehouse Designer a partitioned fact table can only be created from an existing fact table. The process therefore is to create a fact table normally, including the update procedure and indexes. The table type is then changed to a partitioned table type and QAD Data Warehouse Designer will assist in creating the exchange table, modifying the indexes and building the procedure to handle the update of the partitions.

Create a detail or rollup/combined fact table


Create the fact table that you wish to partition in the normal way. You must create an update procedure but do not need to load any data. The process of creating the update procedure will also build any indexes for the fact table.

Convert to a partitioned fact table.

Edit the properties of the fact table to be converted and change the table type to either **Detail partitioned** or **Rollup partitioned**. See example below.

The screenshot shows the 'Fact Table fact_order_detail' dialog box. The 'Table Type' dropdown is open, showing the following options: Detail, Rollup, Work table, Detail partitioned (selected), and Rollup partitioned. The 'Table Name' is 'fact_order_detail', 'Unique Short Name' is 'fact_order_d', and 'Business Display Name (EUL)' is 'fact_order_detail'. The 'Update Procedure' is 'update_fact_order_detail' and the 'Custom Procedure' is 'custom_fact_order_detail'. The 'Get Key Function' is '(None)'. The 'Mnemonic (EUL)' is empty. The 'Timestamps' section shows 'Metadata Structure Changed' as '2013-04-11 15:31:53.381000', 'Database Created' as '2013-04-11 15:31:57.975000', and 'Database Altered' as '2013-04-11 15:31:57.975000'. The 'Cursor Based Update' checkbox is checked.

A dialog will pop-up as follows, asking for confirmation that the table is to be converted to a partitioned fact table. If **Yes** is chosen, then the existing fact table will be dropped and any current data will be lost.



The parameters for the partitioned fact table are then prompted for:

A dialog box titled "Partitioning attributes" with a blue border and a red close button in the top right corner. It contains the following text: "Automatic code generation is only supported on date dimension keys. Select a date dimension key to base the partitioning on:". Below this is a dropdown menu with "dim_order_date_key" selected. Then, there are two columns of instructions. The left column says "Select either day, month or year as the partition granularity:" and has a dropdown menu with "Month" selected. The right column says "Select the corresponding date dimension column. Choose the date for day grain. (E.g. cal_year, cal_month, or calendar_date):" and has a dropdown menu with "order_cal_month" selected. Below these is a paragraph: "The first partition will cater for all periods prior to and including the specified period. New partitions will be created dynamically as required, so the final partition can be the same as the initial partition or perhaps the current period." At the bottom, there are two text input fields: "First partition (YYYYMM) :" with the value "200101" and "Final partition (YYYYMM) :" with the value "200201". At the very bottom, there are two buttons: "OK" and "Cancel".

The fact table will be partitioned by the date dimension key. You need to specify the date dimension key that will be used as the basis for the partitioning of the fact table. It is assumed that this key has the standard QAD format of YYYYMMDD.

A partition granularity of day, month or year must then be selected along with the corresponding column in the date dimension. For a daily partition a column with a date data type must be chosen. The date dimension key will actually be used for the partitioning but the date column is required. For a monthly partition a column with the format YYYYMM must be chosen. In the standard QAD date dimension this is the column cal_month. For a yearly partition a column with the format YYYY must be chosen. In the standard QAD date dimension this is the column cal_year.

The first and last partition should then be selected. The last partition is not that important as additional partitions will be created as required. Normally just select a partition a few on from the first. The first partition is however important as any data that is dated prior to the first partition will be loaded into this first partition. The partition must be entered in the format shown. For example monthly partitions require the format YYYYMM.

Once the **OK** button is clicked the conversion process will populate the **Table Properties Clauses** field of the tables properties. Note that this field is limited in size, so if too many partitions are chosen only the first few will be added to this field. The other partitions will be created dynamically by the update procedure.

Fact Table fact_order_detail

Properties

Storage

Statistics

Language Mapping

Purpose

Concept

Grain

Examples

Usage

Notes

Table Name: fact_order_detail

Table Type: Detail partitioned

Unique Short Name: fact_order_detail

Business Display Name (EUL): fact_order_detail

Description:

Update Procedure: update_fact_order_detail

Custom Procedure: custom_fact_order_detail

Table Properties Clauses: PARTITION BY RANGE (dim_order_date_key) (PARTITION P_200101 STARTING FROM (20010101) INCLUSIVE ENDING AT (20010201) EXCLUSIVE , PARTITION P_200102 STARTING FROM (20010201) INCLUSIVE ENDING AT (20010301) EXCLUSIVE , PARTITION P_200103 STARTING FROM (20010301) INCLUSIVE ENDING AT (20010401) EXCLUSIVE , PARTITION P_200104 STARTING FROM (20010401) INCLUSIVE ENDING AT (20010501) EXCLUSIVE , PARTITION P_200105 STARTING FROM (20010501) INCLUSIVE ENDING AT (20010601) EXCLUSIVE

Get Key Function: [None]

Mnemonic (EUL):

Timestamps

Metadata Structure Changed: 2013-04-11 15:31:53.381000

Database Created: 2013-04-11 15:31:57.975000

Database Altered: 2013-04-11 15:31:57.975000

OK Cancel Help

Click the **OK** button on the table **Properties** dialog to begin the process of converting the fact table to be partitioned. This drops and re-creates the fact table in DB2. The new update procedure is also generated and compiled. Examine the results screen to see what has been done.

The partitioned fact table update procedure

The generated procedure will handle the creation of new partitions as required. Any data for periods prior to the first partition is placed in the first partition. If data is received for a period after the last partition then new partitions are created. Empty partitions are created if necessary to skip periods where no data is present. For example, if we have a table partitioned by month and the latest partition is 200204 (april/2002). If data is then received for say september 2002 the update procedure will build partitions for may, june, july, august and september. There is a variable at the start of the update procedure called `v_max_skip_periods`. This variable is by default set to 12. This defines the maximum number of continuous partitions we will leave empty. From our previous example if our latest partition was april 2002 and we received data for july 2003 with no interim data then the update procedure will fail as it would have to skip 14 partitions. This check is present to prevent erroneous data from creating hundreds of partitions into the future.

Generating the Partitioned Fact Update Procedure

Once a partitioned 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.

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.

A series of questions will be asked during the procedure generation based on the type of load information.

Define Fact Business Key Columns

The first dialog displayed when generating a partitioned fact table update procedure 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_id and order_line_no are selected for the business key list.

Define Fact Business Key Columns

Column List:

- base_currency_code
- base_sale_amount
- base_tax_amount
- bill_to_address_id
- created_datetime
- creating_employee_id
- customer_code
- dim_customer_key
- dim_order_date_key
- dim_product_key
- dim_ship_date_key
- gross_sale_amount
- invoice_date
- last_change_datetime
- last_change_employee_id
- line_type_code
- net_sale_amount
- order_date
- order_number
- order_status_code
- order_type_code
- payment_method_id
- posted_date
- product_id
- quantity
- sales_rep_id
- sales_source
- sales_tax_flag
- ship_date
- ship_to_address_id
- sold_to_address_id
- tax_amount
- tax_amount1
- tax_amount2
- tax_amount3
- tax_amount4
- transaction_employee_code

Select the business keys that uniquely identify each record in the Fact table. Move them over to the Business key list.

NOTE: Set based updates can only be selected if no business keys are defined.

Business Key List:

- order_id
- order_line_no

☐ Set Based Insert

☐ Allow Where Clause Editing

☐ Group By Dimension Keys

☐ Include Delete Before Insert

Insert Hint: (e.g. TABLOCK)

Update Hint: (e.g. TABLOCK)

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 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.

None of the columns chosen as the business key should ever contain a NULL value.

Allow Where Clause Editing

If the **Allow Where Clause Editing** option is selected, then the next dialog to be displayed is the **Provide Cursor Mapping** dialog, else this dialog is skipped.

Insert Hint

Enter a database hint to be used in the INSERT statement.

Note: SQL Server only.


Update Hint

Enter a database hint to be used in the UPDATE statement.

Note: SQL Server only.

Provide Cursor Mapping

This dialog is used to join source tables, add 'Where' clauses and specify group by clauses.



The 'Provide Cursor Mapping' dialog box is shown. It has a title bar with a close button. Inside, there is a text field for 'Cursor Name (if applicable):' with the value 'Load'. Below it is a label 'Edit the where clause if required.' followed by a large empty text area. To the left of this area is a list box labeled 'Source Tables:' containing the text 'stage_order_detail'. To the right of the list box is a label 'From and Where Clause:' followed by another large empty text area. Below these areas are three buttons: 'Outer Join', 'Simple Join', and a checked checkbox labeled 'ANSI join code generated'. At the bottom, there are two sets of dropdown menus: 'Table A' and 'Column A', and 'Table B' and 'Column B'. At the very bottom, there is a checkbox labeled 'Word Wrap Displayed Code' and two buttons: 'OK' and 'Cancel'.

As source table joins should have been performed in the stage table, see ***Generating the Staging Update Procedure*** (on page 349) for more details.

Since most of the real work is done in the stage table, this is all that is needed to build the Partitioned Fact Update Procedure.

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 mouse 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 a dialog box titled "Fact Table Column fact_budget.product_code". On the left is a sidebar with "Properties" selected, and sub-items "Transformation" and "Language Mapping". The main area is divided into sections: "General", "Physical Definition", "Meta Definition", and "Source Details".

General	
Table Name	fact_budget
Column Name	product_code
Business Display Name	product code
Column Description	

Physical Definition	
Column Order	40
Data Type	integer
Null Values Allowed	True
Default Value	

Meta Definition	
Format	#,##0
Numeric	True
Additive	True
Attribute	False
End User Layer Display	True
Business Key	True
Artificial Key	False
Key Type (0,A,B,C,...)	A

Source Details	
Source Table	stage_budget
Source Column	product_code
Transformation	
Join	False

Column Name
Database-compliant name of the column.
Dialog Opening Value: product_code

Buttons at the bottom: <- 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 Data Warehouse Designer-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 data types for Oracle are integer, number, char(), varchar2() and date. For SQL Server common types are integer, numeric, varchar() and datetime. 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.

Key type	Meaning
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.
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

Transformation. [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.

Fact Table Column fact_sales_detail.quantity

Properties
Transformation
Language Mapping

Target: quantity Paste

Source: stage_sales_detail.quantity Paste

Column Transformation Code (must execute within a SQL SELECT statement):

Function Set: SQL Server Standard Funct

ABS(stage_sales_detail.order_number)

dim_product
dim_ship_date
stage_sales_detail
customer_code (numeric)
dim_customer_key (integ
dim_order_date_key (inte
dim_product_key (integer
dim_ship_date_key (integ
dss_update_time (datetim
order_date (datetime)
order_line_no (numeric(4
order_number (numeric(1
product_code (numeric(6
quantity (numeric(8))
sales_source (integer)
sales_value (numeric(13,
ship_date (datetime)
tax (numeric(9,2))
unit_sale_price (numeric

☐ Word Wrap Displayed Code

Function Syntax:

Function Desc.:

<- Update Update -> OK Cancel Help

Note: Transformations are only put into effect when the procedure is re-generated.

See **Transformations** (on page 603) 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". On the left is a sidebar with three tabs: "Properties", "Transformation", and "Language Mapping", with the last one selected. The main area is divided into sections. The "Language :" section has a dropdown menu set to "French". Below it, the "Language Settings:" section contains two text input fields, both containing the word "quantity". The "Business Display Name:" section is empty. The "Business Definition:" section contains two text input fields, both containing the text "Quantity of product sold (i.e. number of product units)". At the bottom right are five buttons: "<- Update", "Update ->", "OK", "Cancel", and "Help".

Chapter 15

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.


When using the Oracle database the aggregate tables are seen as pre-built materialized views with reduced precision and query rewrite enabled.

In This Chapter

Creating an Aggregate Table.....	488
Creating an Aggregate Summary Table	492
Aggregate Table Column Properties	494
Aggregate Table Column Transformation.....	499

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 and enter the aggregate name. Click **ADD**.



The dialog box is titled "Add a New Metadata Object" with a close button (X) in the top right corner. It 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:

Object Name:

At the bottom right, there are two buttons: "ADD" and "Cancel".

- 3 The aggregate properties are displayed. Click **OK**.




The dialog box is titled "agg_sa_product Table Definition" with a close button (X) in the top right corner. It has a left sidebar with a tree view containing the following items: Properties (selected), Storage, Purpose, Concept, Grain, Examples, Usage, and Notes.

The main area contains the following fields and controls:

- Table Name:
- Table Type:
- Unique Short Name: (maximum 22 characters)
- Business Display Name (EUL):
- Description:
- Update Procedure:
- Custom Procedure:
- Table Properties Clauses: (e.g. Partition)
- Mnemonic (EUL):
- Timestamps section with three sub-headers:
 - Metadata Structure Changed:
 - Database Created:
 - Database Altered:


At the bottom right, there are three buttons: "OK", "Cancel", and "Help".

- 4 The list of columns in the Aggregate is displayed in the middle pane:



Column Name	Display Name	Data Type	Source Table
order_number	order number	numeric(12)	fact_sales_detail
order_line_no	order line no	numeric(4)	fact_sales_detail
product_code	product code	numeric(6)	fact_sales_detail
unit_sale_price	unit sale price	numeric(9,3)	fact_sales_detail
quantity	quantity	numeric(8)	fact_sales_detail
sales_value	sales value	numeric(13,2)	fact_sales_detail
tax	tax	numeric(9,2)	fact_sales_detail
order_date	order date	datetime	fact_sales_detail
customer_code	customer code	numeric(6)	fact_sales_detail
ship_date	ship date	datetime	fact_sales_detail
dim_customer_key	dim customer ...	integer	dim_customer
dim_order_date_key	dim time key	integer	dim_order_date
dim_product_key	dim product key	integer	dim_product
dim_ship_date_key	dim time key	integer	dim_ship_date
dss_update_time	dss update time	datetime	

- 5 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) and any dimension keys not required:




Column Name	Display Name	Data Type	Source Table
quantity	quantity	numeric(8)	fact_sales_detail
sales_value	sales value	numeric(13,2)	fact_sales_detail
tax	tax	numeric(9,2)	fact_sales_detail
dim_order_date_key	dim time key	integer	dim_order_date
dim_product_key	dim product key	integer	dim_product
dim_ship_date_key	dim time key	integer	dim_ship_date

Note: All aggregate table columns (other than dss columns) should always have a source table specified, even if a column transformation is in use.

- 6 Create the aggregate table in the database by right-clicking on the aggregate and selecting **Create(ReCreate)**.
- 7 Create a procedure to update the aggregate by right-clicking on the aggregate, selecting **Properties** and selecting (**Build Procedure...**) in the Update Procedure field.

Note: QAD Data Warehouse Designer generated update procedures for aggregates are incremental. Incremental updates are based on a date dimension key and number of look back days. The aggregate update process looks at any records that have been updated in the fact table in the last 7 days (by default).

To support this, a date dimension key must be selected. The columns chosen must be in both the source fact table and the aggregate. Select this column and click **OK**.



8 Update the table by right-clicking and choosing **Execute Update Procedure**.

Note: Any column transformations added to a measure column of an aggregate table must always include an aggregate function, usually **SUM**. For example, an ISNULL added to forecast_quantity should be entered as: SUM(ISNULL(forecast_quantity,0)).

Results	
Object	Message
agg_sa_f	Procedure update_agg_sa_product rebuilt and recompiled.
agg_sa_f	Procedure Completed
agg_1	agg_sa_product updated. 21 new records. 0 records deleted.
agg_I	Date: 20000602 Deleted: 0 Inserted: 2.
agg_I	Date: 20010220 Deleted: 0 Inserted: 4.
agg_I	Date: 20010223 Deleted: 0 Inserted: 1.
agg_I	Date: 20010430 Deleted: 0 Inserted: 2.
agg_I	Date: 20011105 Deleted: 0 Inserted: 2.
agg_I	Date: 20011115 Deleted: 0 Inserted: 3.
agg_I	Date: 20011120 Deleted: 0 Inserted: 1.
agg_I	Date: 20020130 Deleted: 0 Inserted: 1.
agg_I	Date: 20020404 Deleted: 0 Inserted: 5.

Oracle: When the procedure runs it first removes the materialized view, then updates any changes from the fact table into the aggregate. It then re-establishes the materialized view, enabling query rewrite.

DB2: When the procedure runs it first removes the materialized query table, then updates any changes from the fact table into the aggregate. It then re-establishes the materialized query table, enabling query rewrite.


Change an Aggregates Default Look Back Days

QAD Data Warehouse Designer generate update procedures for aggregates that are incremental. Incremental updates are based on a date dimension key and number of look back days.

The aggregate update process looks at any records that have been updated in the fact table in the last number of look back days.

The default number of look back days is 7. The default is set in each update procedure. To change the number of look back days for an aggregate table, create a QAD Data Warehouse Designer parameter called AggregateName_lookback and enter the required number of look back days as the parameter value.

For Example, add the following parameter to change the look back days to 5 for agg_sa_product:


A screenshot of a 'Parameter Maintenance' dialog box. The dialog has a title bar with the text 'Parameter Maintenance' and a red close button. Inside the dialog, there are three input fields: 'Parameter:' with the value 'agg_sa_product_lookback', 'Value:' with the value '5', and 'Comments:' with the text 'The Look Back days for updating agg_sa_product'. At the bottom right, there are two buttons: 'OK' and 'Cancel'.

See **Parameters** for more information on QAD Data Warehouse Designer parameters.

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.



Column Name	Display Name	Data Type	Source Table
quantity	quantity	numeric(8)	fact_sales_detail
sales_value	sales value	numeric(13,2)	fact_sales_detail
tax	tax	numeric(9,2)	fact_sales_detail
dim_order_date_key	dim time key	integer	dim_order_date
dim_product_key	dim product key	integer	dim_product
dim_ship_date_key	dim time key	integer	dim_ship_date

Note: All aggregate summary table columns (other than dss columns) should always have a source table specified, even if a column transformation is in use.

- 4 In the Properties of the aggregate table change the **Table Type** to **Summary**:

agg_sa_product Table Definition

Properties | Storage | Purpose | Concept | Grain | Examples | Usage | Notes

Table Name: Table Type: **Aggregate** (dropdown menu with options: Aggregate, Summary, Work Table)

Unique Short Name: (maximum 22 characters)

Business Display Name (EUL):

Description:

Update Procedure:

Custom Procedure:

Table Properties Clauses: (e.g. Partition)

Mnemonic (EUL):

Timestamps

Metadata Structure Changed: Database Created: Database Altered:

- 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.

Note: Any column transformations added to a measure column of an aggregate summary table must always include an aggregate function, usually **SUM**. For example, an ISNULL added to forecast_quantity should be entered as: SUM(ISNULL(forecast_quantity,0)).

Results		
Object	Message	
agg_sa_product	EXECUTE sp_addextendedproperty N'Comment', N'The tax component of the sales value.', N'user', N'dbo', N'table', 'agg_sa_product', N'column', N'tax';	
agg_sa_product	CREATE NONCLUSTERED INDEX agg_sa_product_idx_1 ON dbo.agg_sa_product (dim_product_key) WITH (SORT_IN_TEMPDB = OFF);	
agg_sa_product	CREATE NONCLUSTERED INDEX agg_sa_product_idx_2 ON dbo.agg_sa_product (dim_ship_date_key) WITH (SORT_IN_TEMPDB = OFF);	
agg_sa_product	update_agg_sa_product regenerated.	
agg_sa_product	update_agg_sa_product compiled	

Results Reports

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 mouse 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 a dialog box titled "Aggregate Column agg_sa_customer.quantity". On the left is a sidebar with "Properties" and "Transformation" tabs. The main area is divided into sections: General, Physical Definition, Meta Definition, and Source Details. The "Column Name" field is highlighted in blue.

General	
Table Name	agg_sa_customer
Column Name	quantity
Business Display Name	quantity
Column Description	Quantity of product sold (i.e. number of product units).

Physical Definition	
Column Order	50
Data Type	numeric(8)
Null Values Allowed	True
Default Value	

Meta Definition	
Format	#,##0
Numeric	True
Additive	True
Attribute	False
End User Layer Display	True
Key Type (0,A,B,C,...)	

Source Details	
Source Table	fact_sales_analysis
Source Column	quantity
Transformation	
Join	False

Column Name
Database-compliant name of the column.
Dialog Opening Value: quantity

Buttons at the bottom: <- 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 Data Warehouse Designer-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 data types for Oracle are integer, number, char(), varchar2() and date. For SQL Server common types are integer, numeric, varchar() and datetime. 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.

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.

Key type	Meaning
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

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

Transformation. [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 Transformation

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 software window titled "Aggregate Column agg_sa_customer.budget_quantity". On the left is a "Properties" sidebar with a "Transformation" tab selected. The main area contains the following fields and controls:

- Target:** A text box containing "budget_quantity" with a "Paste" button to its right.
- Source:** A text box containing "fact_sales_analysis.budget_quantity" with a "Paste" button to its right.
- Column Transformation Code (must execute within a SQL SELECT statement):** A large empty text area for entering the transformation logic.
- Function Set:** A dropdown menu currently set to "SQL Server Standard Funct".
- Functions and Available Columns:** A panel on the right with expandable sections for "Functions" and "Available Columns".
- Word Wrap Displayed Code:** An unchecked checkbox.
- Function Syntax:** An empty text box.
- Function Desc.:** An empty text box.
- Buttons:** At the bottom right are buttons for "<- Update", "Update ->", "OK", "Cancel", and "Help".

Note: Transformations are only put into effect when the procedure is re-generated.

See **Transformations** (on page 603) for more details.

Chapter 16

Views

Views are normally created from dimensions. These views are then used to define dimensions in a fact table where the dimension appears multiple times. A classic example is that of the date dimension. In our tutorials we create a fact table that has two date dimensions. An order date and a ship date dimension. To facilitate the creation of the fact table we create two dimension views to allow the date dimension to be joined to the fact table twice.

Views can also be created from any other table type. By default views are visualized as part of the dimension object group. If fact table views or other views are being used then views can be visualized as a separate object type by checking the appropriate box under the **Tools/Options** menu item.

In This Chapter


Creating a Dimension View	502
Non-Dimension Object Views	510
Creating a Custom View	514
Dimension View Language Mapping	516
View Aliases.....	517

Creating a Dimension View


A dimension view is a database view of a dimension table. It may be a full or partial view. It is typically used in such cases as date dimensions where multiple date dimensions exist for one fact table.

In many data warehouses dimension 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.


Views can be visualized in QAD Data Warehouse Designer either as part of the Dimension object group or in a separate object group called View. The default is to view dimensions and dimension views together:



This visualization option is set from the **Tools/User Preferences** menu item:




This example shows **Dimensions** and **Dimension Views** separately:



The process for creating a view is as follows:

- 1 Double-click on Dimension (or 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 middle pane.
 - The dialog box defaults the object type to a dimension view
 - Change the view name as required, 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:

Object Name:

- The view property defaults dialog will appear. A 'Where' clause could be entered, but this can be done later. Normally we would just click **OK**.



Dimension View dim_ship_date

Properties

- Language Mapping
- Purpose
- Concept
- Grain
- Examples
- Usage
- Notes

View Name: View Type:

Unique Short Name: (maximum 22 characters)

Business Display Name (EUL):

Description:


Where Clause:

Get Key Function: Mnemonic (EUL):

Timestamps


Metadata Structure Changed:	Database Created:	Database Altered:
2006-01-05 14:02:57.000	2013-08-16 08:29:00.833	2013-08-16 08:29:00.833

- 5 A dialog box displays 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 in the section that follows.



The dialog box is titled "View Column Definition" and contains a text area with the following text: "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." Below the text area are four input fields with labels: "Remove Column Prefix:", "Add Column Prefix:", "Remove Business Display Prefix:", and "Add Business Display Prefix:". The "Add Column Prefix" field contains the text "ship_". Below these are five pairs of input fields labeled "Old Column Name:" and "New Column Name:". The first pair has "dim_date_key" in the old field and "dim_ship_date_key" in the new field. The second pair has "calendar_date" in the old field and "ship_date" in the new field. The other three pairs have empty fields. At the bottom right are "OK" and "Default" buttons.

- 6 A dialog box indicates that the view has been defined. If this is a dimension view then you will be given an option to create both the view and an associated Get Key function. For any other view type, a Get Key function is not created.



The dialog box is titled "Create Database View" and contains a text area with the following text: "Dimension View dim_ship_date has been defined". Below the text area are three buttons: "Create View", "Create View + Function", and "Close". The "Create View + Function" button is highlighted with a blue border.

- 7 If the **Create View + Function** button is clicked, the following dialog will display to allow the selection of a business (natural) key for the dimension view.

Define Dimension Business Key Columns

Column List:

ship_cal_day_in_month

ship_cal_day_in_week

ship_cal_day_in_week_no

ship_cal_day_in_year

ship_cal_month

ship_cal_month_name

ship_cal_month_no

ship_cal_quarter

ship_cal_quarter_no

ship_cal_week_in_year

ship_cal_year

ship_current_cal_day

ship_current_cal_month

ship_current_cal_mtd

ship_current_cal_week

ship_current_cal_year

ship_current_cal_ytd

ship_current_fin_day

ship_current_fin_month

ship_current_fin_mtd

ship_current_fin_quarter

ship_current_fin_week

ship_current_fin_year

ship_current_fin_ytd

ship_fin_day_in_month

ship_fin_day_in_week

ship_fin_day_in_week_no

ship_fin_day_in_year

ship_fin_month

ship_fin_month_name

ship_fin_month_no

ship_fin_period

ship_fin_quarter

ship_fin_quarter_no

ship_fin_week_in_month

ship_fin_week_in_year

ship_fin_year

Select the business keys that uniquely identify each record in the Dimension table.

Move them over to the business key list

Business Key List:

ship_date

>

<

OK

Cancel

View Column Re-mapping

The view column definition dialog allows 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:

Remove Business Display Prefix: ---> Add Business Display Prefix:

Change Column Names for Specific Columns

Old Column Name:		New Column Name:
dim_date_key	--->	dim_ship_date_key
calendar_date	--->	ship_date
	--->	
	--->	
	--->	

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 an 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.

Non-Dimension Object Views


QAD Data Warehouse Designer also supports creating views of objects other than dimensions (load tables, facts, stage tables, aggregates, ...).

A fact view is a database view of a fact table. It may be a full, partial view, or a view incorporating both the fact data and some dimensional attributes. It is typically used in cases where a subset of information is to be provided for specific end user communities. It is common that the 'Where' clause be populated in fact views, either to join the dimension tables or restrict the data that the view represents.

Similarly, views may be created over staging tables or any other **table** object.

The creation of non-dimension table views is the same process as that for dimension views:

- 1 Drag a table from the right pane into the middle pane.
- 2 The dialog box defaults the object type to a view.
- 3 Change the view name as required, and click **Add**.



The screenshot shows a dialog box titled "Add a New Metadata Object". The dialog contains the following elements:

- Instructions: "Define the Type and Name of the New Object." and "Specific information for each object type is defined in subsequent screens."
- Object Type: A dropdown menu with "View" selected.
- Object Name: A text input field containing "view_budget".
- Buttons: "ADD" and "Cancel" at the bottom right.

- 4 The view property defaults dialog will appear. A 'Where' clause could be entered, or done later. Click **OK**.

View view_budget

Properties

View Aliases

Purpose

Concept

Grain

Examples

Usage

Notes

View Name: View Type:

Unique Short Name: (maximum 22 characters)

Business Display Name (EUL):

Description:

Custom Procedure:

From/Where or Where Clause:

Mnemonic (EUL):

Timestamps

Metadata Structure Changed:

Database Created:

Database Altered:

- 5 The view column definition dialog allows 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.

View Column Definition [X]

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:	bgt_
Remove Business Display Prefix:		→	Add Business Display Prefix:	bgt_

Change Column Names for Specific Columns

Old Column Name:		→	New Column Name:	bgt_product_code
product_code		→	bgt_customer_code	
customer_code		→	bgt_budget_date	
budget_date		→		
		→		
		→		

OK Default


6 A dialog box indicates that the view has been defined. Create the view.

Create Database View [X]

View view_budget has been defined

Create View Create View + Function Close

7 The columns of the view table will be displayed in the middle pane.



Note: If dimension attributes are being added (eg: to a fact view) then drag these attributes into the views column list, create the view and edit the 'Where' clause for the view.

Creating a Custom View

A custom view can be created within QAD Data Warehouse Designer 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 QAD Data Warehouse Designer and the 'Select' component of the view is customized. The second option is where the view is totally customized and no columns need to be defined in QAD Data Warehouse Designer, 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 **Table Type** to **User Defined View** in the Properties of the view.

The screenshot shows the 'View view_budget' dialog box. On the left is a 'Properties' sidebar with options: Properties, View Aliases, Purpose, Concept, Grain, Examples, Usage, and Notes. The main area contains the following fields:


- View Name:** view_budget
- Unique Short Name: (maximum 22 characters):** view_budget
- Business Display Name (EUL):** view_budget
- Description:** (empty text area)
- Custom Procedure:** (None)
- From/Where or Where Clause:** (empty text area)
- Mnemonic (EUL):** (empty text field)
- Timestamps:**
 - Metadata Structure Changed: 2013-08-19 15:01:49.630
 - Database Created: (empty text field)
 - Database Altered: (empty text field)

At the bottom right are buttons for OK, Cancel, and Help. The 'View Type' dropdown menu is open, showing 'View', 'Work View', and 'User Defined View'.

- 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 QAD Data Warehouse Designer will build up a view create statement from the column names and the supplied **Select** clause. A **Load From Columns** button is available to get a sample **Select** statement based on the columns in the view and any transformations.



Dimension View Language Mapping

The Dimension View 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 "Dimension View dim_ship_date". On the left is a vertical sidebar with a list of tabs: Properties, Language Mapping (which is selected and highlighted in blue), Purpose, Concept, Grain, Examples, Usage, and Notes. The main area of the dialog is divided into sections. The "Language :" section has a dropdown menu currently set to "French". Below this is the "Language Settings:" section. The "Business Display Name:" section contains a text field with the value "dim_ship_date" and an empty text input field below it. The "Description:" section contains two empty text input fields, each with up and down arrow buttons on its right side. At the bottom right of the dialog are three buttons: "OK", "Cancel", and "Help".

Business Display Name

Translated Value of Business Display Name

Description

Translated Value of Description

View Aliases

The **View Aliases** tab enables you to define additional/replica views.

The screenshot shows a software window titled "View view_customer" with a red close button in the top right corner. On the left is a sidebar with a list of tabs: Properties, View Aliases (selected), Purpose, Concept, Grain, Examples, Usage, and Notes. The main area displays the "View Aliases" configuration for the "view_customer" view. It includes a tree view on the left with "View Aliases" expanded, showing "Alias - view_customer" selected. To the right of the tree are buttons for "Add View Alias" and "Delete View Alias". Below the tree is a table with fields for "View Alias Name" (containing "view_customer"), "View Alias Description", and "View Alias Predicate", each with a corresponding "..." button for editing. At the bottom of the table is a section for "Target Location" with a "Target Schema" field. At the very bottom of the window are "OK", "Cancel", and "Help" buttons.

View Aliases	
Add View Alias Add	
Alias - view_customer	
Delete View Alias	Delete
View Alias Name	view_customer
View Alias Description	...
View Alias Predicate	...
Target Location	
Target Schema	

Add button

Enables you to add a View Alias.

Delete button

Enable you to delete a View Alias.

View Alias Name

The view alias name.

View Alias Description

Description of the view alias.

View Alias Predicate

Optional 'Where' clause to include in the alternate view definition.

Target Schema

The target schema

Chapter 17

Analysis Services OLAP Cubes

In This Chapter

- Overview 520
- Defining the Data Source for the OLAP Cube 521
- Defining an OLAP Cube..... 526
- Inspecting and Modifying Advanced Cube Properties 530
- Creating an OLAP Cube on the Analysis Services Server .. 531
- OLAP Cube Objects 532
- OLAP Dimension Objects..... 567
- Changing OLAP Cubes 583
- Retrofitting an OLAP Object 586

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 QAD Data Warehouse Designer has simplified and restricted many of the functions available to the cube designer. QAD Data Warehouse Designer includes most of the commonly used capabilities and the components that logically fit into the methodology incorporated within QAD Data Warehouse Designer.

QAD Data Warehouse Designer broadly provides functionality to manage all of the above, except for perspectives and translations. These can be created outside of QAD Data Warehouse Designer, scripted in xmla and executed from within QAD Data Warehouse Designer. Features of cubes that are not supported in QAD Data Warehouse Designer can be added to the cube via the Microsoft tools. These altered cubes can still be processed through the QAD Data Warehouse Designer scheduler, and the cube should be documented within QAD Data Warehouse Designer 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 QAD Data Warehouse Designer. The OLAP object can be dropped and recreated easily using Data Warehouse Designer. New OLAP objects defined in Data Warehouse Designer (e.g. additional calculations or measures) can be added by recreating the cube.

QAD Data Warehouse Designer supports cubes in Microsoft Sql Server Analysis Services versions 2005 and 2008.

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.

The screenshot shows the 'Connection DataWarehouse' dialog box with the 'When Connection is an OLAP Data Source' section expanded. The 'OLAP Connection String' field is highlighted in blue. The 'Connection Provider/Driver' is set to 'SQLOLEDB', the 'Data Warehouse Server' is 'DOC', and the 'Data Warehouse Database ID' is 'R1'.

Connection DataWarehouse	
General	
Connection Name	DataWarehouse
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	R_1
Data Warehouse Connection Indicator	True
Source System	
Database ID	R1
Database Link Name	
Provider Name	SQLOLEDB
Database Credentials	
Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	
Other	
New Table Default Load Type	Database link load
SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
Data Type Mapping Set	(Default)
Default Transform Function Set	(Default)
When Connection is an OLAP Data Source	
OLAP Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
Connection Provider/Driver	SQLOLEDB
Data Warehouse Server	DOC
Data Warehouse Database ID	R1

OLAP 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.

OK Cancel Help

OLAP 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 **SQLOLEDB**.

Data Warehouse Server

Data Warehouse Server Name, which is used when the data warehouse is used as the data source for OLAP cubes.


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.


The **OLAP Connection String** is built using a wizard. To activate the wizard, click on the expander button.

When Connection is an OLAP Data Source	
OLAP Connection String	
Connection Provider/Driver	SQLLEDB 
Data Warehouse Server	DOC
Data Warehouse Database ID	WslWarehouse

On the Provider tab, select the **OLE DB Provider** and click **Next**.




Enter the connection details and click Test Connection.



The image shows a 'Data Link Properties' dialog box with a close button (X) in the top right corner. It has four tabs: 'Provider', 'Connection', 'Advanced', and 'All'. The 'Connection' tab is selected. The dialog contains the following elements:

- Text: 'Specify the following to connect to SQL Server data:'
- Section 1: '1. Select or enter a server name:'
 - Dropdown menu: '(local)' with a downward arrow.
 - Button: 'Refresh'.
- Section 2: '2. Enter information to log on to the server:'
 - Radio button (selected): 'Use Windows NT Integrated security'.
 - Radio button: 'Use a specific user name and password:'
 - Text: 'User name:' followed by a text input field.
 - Text: 'Password:' followed by a text input field.
 - Checkboxes: 'Blank password' and 'Allow saving password'.
- Section 3: '3. Select the database on the server:'
 - Radio button (selected): 'Select the database on the server:' followed by a dropdown menu showing 'WslWarehouse' with a downward arrow.
 - Radio button: 'Attach a database file as a database name:'
 - Text input field.
 - Text: 'Using the filename:' followed by a text input field and a browse button ('...').
- Button: 'Test Connection'.
- Bottom buttons: 'OK', 'Cancel', and 'Help'.

Click **OK** on the success message and then **OK** again on the Data Link Properties screen.



The OLAP connection string will be displayed.


Defining an OLAP Cube

OLAP Cubes can be created from fact, fact view or aggregate objects. A single cube can contain dates 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.

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 table** 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:

Object Name:

- 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 table metadata). Set the dimension name in the **Add a new metadata Object** dialog box and click **ADD**.




The dialog box is titled "Add a New Metadata Object" with a red close button in the top right corner. It contains the following elements:

- Text: "Define the Type and Name of the New Object."
- Text: "Specific information for each object type is defined in subsequent screens."
- Object Type: A dropdown menu showing "Olap Dimension".
- Object Name: A text box containing "odim_customer".
- Include Attribute Relationships: An unchecked checkbox.
- Buttons: "ADD" and "Cancel".

Note: If you wish to include Attribute Relationships in Analysis Services for this dimension, click on the **Include Attribute relationships** checkbox.

The following dialog appears, prompting you to select the attributes to be included in the *Customer* OLAP dimension.



The dialog box is titled "Select OLAP Dimension Attributes" with a red close button in the top right corner. It contains the following elements:

- Text: "Select end user layer attributes to be included as OLAP dimension attributes"
- Available Attributes: A list box containing "dss_current_flag", "dss_end_date", "dss_start_date", "dss_update_time", and "dss_version".
- Selected Attributes: A list box containing "customer_category_code", "customer_code", "customer_group_code", "customer_name", "customer_subgroup_code", and "sales_source".
- Navigation: ">" and "<" buttons between the two list boxes.
- Buttons: "OK", "Cancel", and "Help" at the bottom right.

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.

- 6 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

- quantity
- gross sale amount
- net sale amount
- tax amount

Remove -->

Calculated Members:

- ☒ Month To Date
- ☒ Year To Date
- ☐ Moving Quarter
- ☒ Moving Year
- ☒ Same Month Previous Year
- ☒ Previous Year To date

☒ Use OLAP Functions

Dimension and level information for calculated members:

Date Dimension: dim_order_date

Date Hierarchy: calendar

Month Level: order cal month

Year Level: order cal year

OK Cancel

- 7 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 QAD Data Warehouse Designer will show all attributes in the fact table 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 measure based on the associated drop-down boxes. There are 2 different ways that QAD Data Warehouse Designer 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.

- 8 The cube and dimensions will be created in QAD Data Warehouse Designer 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. The provider type will normally be MSDAORA if the data warehouse is an Oracle database or SQLOLEDB if the data warehouse is SQL Server. The Server is either the SQL Server database server, or the Oracle TNS name. The database is the data warehouse database name.

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**.

Creating an OLAP Cube on the Analysis Services Server

If all of the tasks above are completed then it should be possible to now create the cube on the Analysis Services server. When positioned on the **OLAP Cube** name, right-click and select **Create (Alter) Cube**. QAD Data Warehouse Designer will check that key components of the cube are correct before it proceeds to issue the create command.

The create cube menu option will perform 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 to 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 QAD Data Warehouse Designer metadata and displayed in the generated QAD Data Warehouse Designer documentation. In order to see the cube properties, right-click on the cube name and select **Properties**.

The screenshot shows the 'Cube olap_orders Definition' dialog box. On the left is a tabbed interface with the following tabs: 'Cube Properties' (selected), 'Language Mapping', 'Purpose', 'Concept', 'Grain', 'Examples', 'Usage', and 'Notes'. The main area contains the following fields and controls:

- Internal Cube Name:
- Cube Publish Name:
- Cube Description:
- Cube Database Connection: (Analysis Services)
- Cube Database Name:
- Data Source Connection: (Data Warehouse)
- Data Source Provider Type:
- Data Source Server:
- Data Source Database:
- Post Create XML/A Script:
- Post Update XML/A Script:
- Processing Mode:
- Processing Priority:
- Process Cube Dimensions:
- Storage Mode:
- Default Measure:
- Estimated Rows:
- Visible:

At the bottom right are three buttons: 'OK', 'Cancel', and 'Help'.

Internal Cube Name

This is the name by which the cube is known within QAD Data Warehouse Designer. 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 for the limitations imposed on it by Analysis Services.

Cube Description

A description of the cube. This is used both in the QAD Data Warehouse Designer 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 QAD Data Warehouse Designer. 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 QAD Data Warehouse Designer 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. The three options are MSDAORA where the data source is an Oracle database. OLESQLDB where the data source is a SQL Server database, or MSDASQL to define an ODBC source. The MSDASQL provider type could be used for either Oracle or SQL Server data warehouses, but performance may be compromised. 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. For Oracle it is the TNS names entry for the database.

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. For Oracle it is also the database name.

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 QAD Data Warehouse Designer - 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 QAD Data Warehouse Designer - 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.

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**, to process **no dimensions**, or to process **all the dimensions**.



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 tables. Enter the size of the fact or source table 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.

Cube olap_orders Definition

Cube Properties
Language Mapping
Purpose
Concept
Grain
Examples
Usage
Notes

Language : French

Language Settings:

Cube Publish Name: olap_orders

Translated Value of Cube Publish Name

Cube Description: Use this fact table to track our Sales Orders

Translated Value of Cube Description

OK Cancel Help

See **Language Settings** (see "**Settings - Language Options**" on page 102) for more details on how to add languages for translation.

OLAP Cube Measure Groups

A cube can contain multiple measure groups. In QAD Data Warehouse Designer 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 tab entitled **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 displays the 'Measure Group order_detail Definition' dialog box. On the left is a sidebar with a tree view containing the following items: 'Measure Group Properties' (selected), 'Processing / Partitions', 'Language Mapping', 'Purpose', 'Concept', 'Grain', 'Examples', 'Usage', and 'Notes'. The main area of the dialog is titled 'Measure Group order_detail Definition' and contains the following fields and controls:

- Measure Group Name:** A text box containing 'order_detail'.
- Measure Group Description:** A large text area containing 'Use this fact table to track our Sales Orders'.
- Source Table Type:** A dropdown menu set to 'Fact Table'.
- Source Table:** A dropdown menu set to 'fact_order_detail'.
- Estimated Rows:** A text box containing '0'.
- Storage Mode:** A dropdown menu set to 'MOLAP'.
- Ignore Unrelated Dimension:** A dropdown menu set to 'True'.
- Measure Group Type:** A dropdown menu set to 'Regular'.
- Processing Mode:** A dropdown menu set to 'Regular'.
- Processing Priority:** A text box containing '0'.

At the top right of the main area are two buttons: '< Update' and 'Update >'. At the bottom right are three buttons: 'OK', 'Cancel', and 'Help'.

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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer.

Source Table

Specifies the table from which the Measure Group has been built and populated in QAD Data Warehouse Designer.

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 QAD Data Warehouse Designer.

Measure Group order_detail Definition

Measure Group Properties

Processing / Partitions

Language Mapping

Purpose

Concept

Grain

Examples

Usage

Notes

The measure group processing method is, by default, inherited by all partitions. Partitions can be individually changed to utilize a different method if required.

Process Method: Full process

Increment Filter:
(without the WHERE keyword)

The partitions for the measure group can be viewed. These partitions can be managed manually or automatically by WhereScope RED. Select the desired partition management strategy.

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 < GETDATE0-14

Max # Auto Create Partitions:
(0 = unlimited) 12

Historic Partitions Updated: 3

OK Cancel Help

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 a fact table undergoes changes to the data 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 fact 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 QAD Data Warehouse Designer 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 QAD 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. QAD Data Warehouse Designer 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 QAD Data Warehouse Designer executes a query against the fact table and the date dimension to acquire each unique period. When dealing with a large fact table, 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 table 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 table 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. QAD Data Warehouse Designer 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 table 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 Formulas:

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 QAD Data Warehouse Designer. 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 QAD Data Warehouse Designer.

Fact Table

The fact table that the data is derived from. 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 QAD Data Warehouse Designer 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 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 QAD Data Warehouse Designer 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:

Cube Measure Definition

Measure Properties

Language Mapping

Cube Name:

Measure Name:

Measure Group:

Source Table:

Source Column:

Data Type:

Aggregation Method:

Display Format:

Null Processing:

Order Number:

Visible:

Display Folder:

Description:

< - Update

Update ->

OK Cancel Help

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 table 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.

Source Table

A read only field that indicates the fact table that is related to the Measure Group.

Source Column

Specifies from which numeric attribute of the underlying fact table the Measure is built. This is a drop-down list populated from Data Warehouse Designers metadata definition of the fact table 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 table attribute defined in QAD Data Warehouse Designer 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:

financials;sales

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 Calculated Member Definition

Calculated Member Properties
Language Mapping

Cube Name: <- Update

Calculated Member Name: Update ->

Description:

Expression:

Parent Hierarchy:

Parent Member:

Associated Measure Group:

Display Folder:

Display Format:

Font: ☐ Bold

Visible:

Non Empty Behavior:

Order Number:

OK Cancel Help

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 QAD Data Warehouse Designer 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 table 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 Data Warehouse Designer to **(Undefined)** for previous versions of Data Warehouse Designer, but this is not necessary for the current version of Data Warehouse Designer 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:

financials;sales ▼

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 KPI Definition

KPI Definition

Language Mapping

Cube Name: olap_orders

KPI Name: sales_growth

Description: The KPI shows growth against the same period last year. The goal is to attain 10% growth year on year.

Display Folder:

Associated Measure Group: order_detail

Value Expression: [Measures].[order_count]

Goal Expression: [Measures].[prev year month order count]*1.1

Status Indicator: Smiley Face

Status Expression: KpiValue("Sales Growth")/KpiGoal("Sales Growth") >= 0.85

Trend Indicator: Standard Arrow

Trend Expression: KpiValue("Sales Growth")/KpiGoal("Sales Growth")*1.1

Parent KPI:

Weight:

OK Cancel Help

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 QAD Data Warehouse Designer 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:

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 Indicator

A visual element that provides a quick indication of the status of a KPI. The display of the element is determined by the value of the MDX expression that evaluates status.

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.

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.

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.

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 Action Definition

Action Properties

Drill Columns

Language Mapping

Cube Name:

Action Type:

Action Name:

Target Type:

Measure Group Members:

Condition:

Server:

Report Path:

Format / Content Type:

Expression:

Default:

Invocation:

Application:

Description:

Caption:

Caption is MDX:

<- Update

Update ->

OK

Cancel

Help

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-down list is enabled for Drill through actions. This gets or sets the current DrillThroughAction as the default action when multiple drillthrough actions are defined. This is important for users of Excell 2007 to browse the Olap Cube because Excel 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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer.

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 QAD Data Warehouse Designer, 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**. The properties are shown below:

Cube Dimension dim_customer Definition

Cube Dimension Properties

Language Mapping

Internal Dimension Name:

OLAP Dimension Name:

Cube Dimension Name:

Dimension Description:

Order Number: ☒ Process Cube Dimension With Cube

Visible:

All Member Aggregation Usage:

Hierarchy Unique Name Style:

Member Unique Name Style:

Source Table Type:

Source Table:

Source Table Key:

Processing Mode:

All Caption:

OLAP Dimension Type:

Unknown Member Action:

Unknown Member Name:

Internal Dimension Name

A read only field displaying the name of the OLAP Dimension in QAD Data Warehouse Designer.

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 QAD Data Warehouse Designer this equates to the relationships between fact tables and dimensions in the underlying star schema.

The **Properties** are shown below:

Measure Group Dimension dim_customer Definition

Cube Dimension Properties

Internal Dimension Name:

OLAP Dimension Name:

Cube Dimension Name:

Dimension Description:

Order Number: ☒ Process Cube Dimension With Cube

Visible:

All Member Aggregation Usage:

Hierarchy Unique Name Style:

Member Unique Name Style:

Measure Group:

Measure Group Column:

Relationship Type:

Cardinality:

Source Table Type:

Source Table:

Source Table Key:

Processing Mode:

All Caption:

OLAP Dimension Type:

Unknown Member Action:

Unknown Member Name:

< - Update Update ->

OK Cancel Help

Internal Dimension Name

A read only field displaying the name of the OLAP Dimension in QAD Data Warehouse Designer.

OLAP Dimension Name

A read 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 table key joins to the dimension key.

Relationship Type

Defines the relationship type for the relationship between the Dimension and Measure Group. In QAD Data Warehouse Designer 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 QAD Data Warehouse Designer when a fact table is dragged over to create a cube or measure group. QAD Data Warehouse Designer 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 QAD Data Warehouse Designer.

The properties of an OLAP Cube dimension are shown by right-clicking on the **OLAP Dimension** and choosing **Properties**. The following dialog is shown:

Olap Dimension odim_customer

Properties

Language Mapping

Purpose

Concept

Grain

Examples

Usage

Notes

Internal Dimension Name: odim_customer

Dimension Publish Name: dim_customer

Dimension Description: This is the master Customer List for the Company

Default Database Connection: (Analysis Services)

OLAP Database Name:

Data Source Connection: (Data Warehouse)

Data Source Provider Type: SQLOLEDB

Data Source Server: WSL-SHUTTLE-8\SQL2012

Data Source Database: WslWarehouse

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

OK Cancel Help

Internal Dimension Name

Specifies the name of the dimension in QAD Data Warehouse Designer.

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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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

Post Create XML/A Script

Specifies XML/A script that is executed after the object is created.

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 cube 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

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 102) 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:

Dimension odim_order_date Attribute Definition

Attribute Properties

Language Mapping

Dimension Name: odim_order_date

Internal Attribute Name: order_cal_year

Published Name: order cal year

Description: The calendar year. Format YYYY

Estimated Count: 1

Member Names Unique: False

Hierarchy Visible: True

Hierarchy Enabled: True

Hierarchy Optimized State: FullyOptimized

Hierarchy Display Folder:

Order By: Key

Order By Attribute:

Type : Years

Usage: Regular

Key Column: dim_order_date

Name Column: dim_order_date

order_cal_year

order_cal_year

<- Update

Update ->

OK

Cancel

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 QAD Data Warehouse Designer.

Published Name

The name of the attribute created in Analysis Services.

Description

A business name that is stored in QAD Data Warehouse Designer 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.

Hierarchy Display Folder

Attribute hierarchies 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.

Notes:

1. One object can be in multiple display folders, for example:

Display Folder:

financials;sales




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: AttributeKey

Order By Attribute: order_current_cal_year, order_current_cal_ytd, order_moving_cal_year, Regular

Type : Regular

Usage: Regular

Key Column: dim_order_date

Name 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.

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 image shows a dialog box titled "Attribute Relationship Properties" with a close button (X) in the top right corner. The dialog contains several input fields and dropdown menus for configuring an attribute relationship. The fields are as follows:

Property	Value
Dimension Name:	odim_customer
Attribute:	customer_name
Relationship Name:	Customer Name
Related Attribute :	customer_name
Relationship Type:	Rigid
Cardinality:	One
Visible:	True

At the bottom right of the dialog are two buttons: "OK" and "Cancel".

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**.

Dimension odim_customer Hierarchy Definition

Hierarchy Properties

Language Mapping

Dimension Name: odim_customer

Internal Hierarchy Name: customer

Hierarchy Publish Name: customer

Hierarchy Description: Added at dimension creation for cube support

All Member Name:

Allow Duplicate Names : True

Member Keys Unique: NotUnique

Member Names Unique: False

Display Folder:

OK Cancel Help

Dimension Name

A read only field indicating the dimension associated with the attribute relationship.

Internal Hierarchy Name

The name of the hierarchy within QAD Data Warehouse Designer.

Hierarchy Publish Name

The name of the hierarchy as created within Analysis Services.

Hierarchy Description

A business name that is stored in QAD Data Warehouse Designer 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.

Display Folder

Dimension hierarchies 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.

Notes:

1. One object can be in multiple display folders, for example:

Display Folder:

financials;sales



2. Dimension Attributes and Dimension Hierarchies have the same display folder drop-down list.

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 a User Defined Hierarchy, 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 odim_customer Hierarchy Level Definition

Hierarchy Level Properties

Language Mapping

Dimension Name: odim_customer

Hierarchy: customer

Level Number: 2

Internal Level Name: Group Code

Level Publish Name: Group Code

Level Description: The group code for the customer. Forms a hierarchy with customer_code

Source Attribute : customer_group_code

Hide If: Never

OK Cancel Help

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 QAD Data Warehouse Designer.

Level Publish Name

The name of the Level as created within Analysis Services.

Level Description

A business name that is stored in QAD Data Warehouse Designer 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 containing 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**.


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 QAD Data Warehouse Designer.

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 Data Warehouse Designer to recreate the DSV or manually update the DSV using Microsoft BIDS.

If an object needs to be dropped and recreated in Data Warehouse Designer then these are two separate actions. For example to drop the OLAP database, right-click an OLAP cube within that database in Data Warehouse Designer and select **Drop Analysis Services Object**, then either double-click on an object in the list or use the drop-down boxes in the Drop Analysis Services Object dialog to 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 Data Warehouse Designer 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 Data Warehouse Designer will create the object (and any related objects e.g. OLAP database and DSV).
- ☐ If the object does already exist in Analysis Services then Data Warehouse Designer 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 Data Warehouse Designer.	OLAP database needs to be dropped and recreated.
Data Source View (underlying relational star)	The DSV reflects the design of the relational star. So 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 Data Warehouse Designer 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 Data Warehouse Designer 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 Data Warehouse Designer (a reprocess is not necessary because just the calculation definition is stored in the cube - the result is calculated at query time).


Retrofitting an OLAP Object

QAD Data Warehouse Designer 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**.



- 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_customer2

Properties

Language Mapping

Purpose

Concept

Grain

Examples

Usage

Notes

Internal Dimension Name: odim_customer

Dimension Publish Name: odim_customer

Dimension Description: Customer Dimension

Default Database Connection: (Analysis Services) Cubes

OLAP Database Name: Tutorial5

Data Source Connection: (Data Warehouse) DataWarehouse

Data Source Provider Type:

Data Source Server:

Data Source Database:

Post Create XML/A Script: (None) Edit

Source Table Type:

Source Table:

Source Table Key:

Processing Group: ByAttribute

Processing Mode: Regular

Processing Method: Default process

Storage Mode: MOLAP

All Caption: All odim_customer


OLAP Dimension Type:

Unknown Member Action: None


Unknown Member Name: Unknown

OK Cancel Help


- 4 The OLAP Dimension can now be retrofitted by right-clicking the object in the left pane and choosing **Retrofit Dimension**.



- 5 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 18

Exporting Data

Export objects are used in QAD Data Warehouse Designer 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 two 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 script file is executed to perform the load.


In This Chapter

Building an Export Object	592
File Attributes	595
Export Column Properties	598
Host Script Exports	601

Building an Export Object

The simplest way to create an export object is to use the drag and drop functionality of QAD Data Warehouse Designer.

- 1 Browse to the data warehouse connection (Browse/Source Tables).
- 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 [X]


Define the Type and Name of the New Object.
Specific information for each object type is defined in subsequent screens.

Object Type: Export ▼

Object Name: exp_customer

ADD Cancel

- 4 If the export object needs to be renamed, do so and then click the **ADD** button. The following dialog appears.



The dialog box is titled "exp_customer Table Definition". It features a left-hand sidebar with a list of tabs: Properties, File Attributes, Purpose, Concept, Grain, Examples, Usage, and Notes. The "Properties" tab is currently selected. The main area of the dialog contains several fields and dropdown menus for configuring the export object. The fields are as follows:

Field Label	Value
Export Object Name:	exp_customer
Unique Short Name: (maximum 22 characters)	exp_customer
Description:	
Connection:	
Export Type:	File export
Database Link:	
Script Name:	(None)
Pre-Export Action:	No action
Pre-Export Sql:	
Where Clause: Allows filtering of the export data.	
Post Export Procedure:	(None)

At the bottom of the dialog, there is a "Timestamps" section with two sub-fields:

Timestamps	
Metadata Structure Changed: 2013-01-31 12:59:37.230	Last Exported:

The dialog concludes with three buttons at the bottom right: "OK", "Cancel", and "Help".

5 Change the Connection to **Windows**.

The screenshot shows a dialog box titled "exp_customer Table Definition". On the left is a sidebar with a tree view containing the following items: Properties (selected), File Attributes, Purpose, Concept, Grain, Examples, Usage, and Notes. The main area of the dialog contains the following fields and controls:

- Export Object Name:** exp_customer
- Unique Short Name:** (maximum 22 characters) exp_customer
- Description:** (empty text box)
- Connection:** Windows (dropdown menu)
- Export Type:** File export (dropdown menu)
- Database Link:** (empty text box)
- Script Name:** [None] (dropdown menu)
- Pre-Export Action:** No action (dropdown menu)
- Pre-Export Sql:** (empty text box)
- Where Clause:** Allows filtering of the export data. (empty text box)
- Post Export Procedure:** [None] (dropdown menu)
- Timestamps:**
 - Metadata Structure Changed:** 2013-01-31 12:59:37.230
 - Last Exported:** (empty text box)


At the bottom right of the dialog are three buttons: OK, Cancel, and Help.

6 Click on the **File Attributes** tab and fill in the fields as described in the next section.

7 Now run the export by right-clicking on it and selecting **Export**.

File Attributes

The following fields are available to define the location, name and contents of the exported data file:



The screenshot shows a dialog box titled "Export exp_budget" with a sidebar on the left containing the following menu items: Properties, File Attributes (selected), Purpose, Concept, Grain, Examples, Usage, and Notes. The main area is titled "Export File Definition" and contains the following fields and options:

- Export File Path:** C:\data\
- Export File Name:** budget_YYYYMMDD\$.txt
- Export Routine:** sqlcmd
- Export Format:** Delimited Text
- Export File Delimiter:** | **Optionally Enclosed by (E.g. '');**
- Export Options:**
- Header Row in Export:** Column names

If a trigger file is specified then this file is created after the export file. It will normally contain a row count and check sum.

- Trigger Path:** C:\data\
- Trigger Name:** budget_YYYYMMDD\$.trg
- Trigger Delimiter:** |
- Trigger Parameter 1:** Row Count
- Trigger Parameter 2:**
- Trigger Parameter 3:**

The variable \$EXPFIL\$ can be used in 'Compress Parameters' to substitute in the export file path and name.

- Compress After Export:** ☒
- Compress Utility Path:** C:\Program Files\7-Zip\
- Compress Utility Name:** 7z.exe
- Compress Parameters:** a -y \$EXPFIL\$.zip \$EXPFIL\$

At the bottom right are buttons for OK, Cancel, and Help.

Export File Path

The full path name to the directory in which the file is to be created on the Windows system.

Export File Name

The name of the file in the directory.

Export Routine

For Oracle this is SQL*Plus. In SQL Server 2000 this is BCP or isql. In later versions of SQL Server this is BCP or sqlcmd. For DB2 the supported export routine is EXPORT.

Export Format

This is the format of the data in the export file and is either Delimited Text or Width Fixed Text in Oracle and SQL Server. In DB2, the supported formats are DEL (delimited), WSF (worksheet format) and IXF (integrated exchange format).

Export File Delimiter

Any delimiter in the file that splits out the columns, e.g., a comma (,). 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 to enclose each text field with. This is only available if the Export Format is Delimited Text.

Header Row in Export

If a header line is required, choose business names or column names from this drop-down list. This option is not available in DB2.

Export Options

Allows the entry of export utility options. If more than one option is required, then a semi-colon should be used between options. This option is not available in DB2.

Export Modifier

Allows the entry of EXPORT modifiers in DB2.

Trigger 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 to be generated on the Windows systems.

Trigger Name

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).

Trigger Delimiter

Multiple fields in the trigger file are to be separated by the trigger delimiter.

Trigger Parameter 1,2,3

The 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

Check this box if you want to compress the export file after it has been created.

Compress Utility Path

The directory in which the compress utility exists.

Compress Utility Name

The name of the compression utility executable.

Compress Parameters

The name of the file to be compressed (using the QAD Data Warehouse Designer variable \$EXPFILE\$) and any commands or switches required to make the compression utility work. These parameters will depend on the compression utility used.

Export Column Properties

Export Column exp_budget.product_code

Properties

Transformation

General

Table Nameexp_budget

Column Nameproduct_code

Business Display Nameproduct code

Column Description

Physical Definition

Column Order40

Data Typeinteger

Meta Definition

Format#,##0

NumericTrue

AdditiveTrue

AttributeFalse

Source Details

Source Tablefact_budget

Source Columnproduct_code

Transformation

Column Name

Database-compliant name of the column.

Dialog Opening Value: product_code

<- Update

Update ->

OK

Cancel

Help

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 Data Warehouse Designer-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 data types for Oracle are integer, number, char(), varchar2() and date. For SQL Server common types are integer, numeric, varchar() and datetime. 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 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

Transformation. [Read-only].

Host Script 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 a 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 QAD 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 QAD Data Warehouse Designer tool under the scheduler window.

Chapter 19

Transformations

Standard **column transformations** (on page 604) can be used in QAD Data Warehouse Designer to perform calculations, change data types or format data.

Re-using complex transformations can save a significant amount of time. These can be achieved in two ways with QAD Data Warehouse Designer:

- ☐ Database Functions
- ☐ QAD Data Warehouse Designer Re-usable Transformations

In This Chapter

Column Transformations.....	604
Database Functions	614
Re-usable Transformations	615

Column Transformations

Each table or export object column can have a transformation associated with it. For all table types, except for load tables, 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 608) for more details.

Dimension View transformations are included in the Database Definition Language (DDL) that creates the view in the database. Any changes to Dimension View column transformations require the view to be dropped and recreated.

Export object column transformations are dynamically applied for file exports. 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 **Re-usable 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 608) 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 you click on the function. Essentially the same as the Function model; loaded into the **Function SQL Text Window** when you click on the function. 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 **Re-usable Transformations**. This is only visible for Re-usable Transformations. Read only.

Localize Transformation

The **Localize transformation** button breaks the link between a **column transformation** and the **Re-usable Transformation** it is based on. If this button is clicked, changes to the underlying re-usable transformation cannot be automatically propagated to the column transformation. This is only visible for Re-usable 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	☐ During Load transformations ☐ After Load transformations ☐ Post Load procedure
ODBC Based Load	☐ During Load transformations ☐ After Load transformations ☐ Post load procedure
File Based Load	☐ During Load transformations ☐ After Load transformations ☐ Post load procedure
Script Based Load	☐ During Load transformations ☐ After Load transformations ☐ Post load procedure
Externally Loaded	☐ 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 SQL language of the data warehouse database. 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.

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.

Load Table load_product

Properties
Storage
Source
Notes

Load Type: Database link load
Source Connection: Tutorial (OLTP)

General

Select Distinct Values: False
Source Schema: dbo
Derive Source Table(s) and Source Columns: Derive
Source Table(s): product
Override Source Column/Transformations: False
Where and Group By Clauses: ...
Override Load SQL: ...
SQL Insert Hint: TABLOCK

Select Distinct Values
Include the DISTINCT keyword in the SQL SELECT statement.

OK Cancel Help

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_product	Load Complete
load_product	Table dbo.load_product truncated.
Load	INSERT INTO dbo.load_product WITH (TABLOCK)
Load	SELECT
Load	product.code,product.description,product.prod_line,product.prod_group,product....
Load	From [Tutorial].dbo.product product
Load	;
Load	9 rows loaded

If the Column has a transformation defined as follows:

UPPER(substring(description,1,1)) + LOWER(substring(description,2,1))

then the following SQL statement will be executed.

Results			▼	🔍	×
Object	Message				
load_product	Table dbo.load_product truncated.		✓		
Load	INSERT INTO dbo.load_product WITH (TABLOCK)		✓		
Load	SELECT		✓		
Load	product.code,UPPER(substring(description,1,1)) + LOWER(substring(description,		✓		
Load	2,1)),product.prod_line,product.prod_group,product.subgroup		✓		
Load	From [Tutorial].dbo.product product		✓		
Load	;		✓		
Load	9 rows loaded		✓		
Results			Reports		

File During Load Transformations

The loading of flat files is performed using the database specific loader. For Oracle, this is SQL*LOADER. For SQL Server it is the Bulk Insert Transact SQL statement. The contents of the **Transformation** tab in the column's Properties are the functions and conversions supported by the database loader.

Example

SQL*LOADER performs transformations such as:

- ☐ DATE 'DD-MMM-YYYY' converts from a value such as 23-Mar-1999 to an Oracle date.
- ☐ "TO_NUMBER(:COL)" converts the data to a number
- ☐ "NVL(TRIM(: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

lpad(code,5,'0')

would result in the following SQL statement being executed after the load:

update load_product set code = lpad(code,5,'0');

Database Functions

Database functions can be created in the data warehouse database (as procedure objects in QAD Data Warehouse Designer) and used in column transformation in QAD Data Warehouse Designer. See the Oracle and SQL Server documentation on functions and ***Procedures and Scripts*** (on page 627) for more information on creating functions inside QAD Data Warehouse Designer.

Re-usable Transformations

QAD Data Warehouse Designer Re-usable Transformations allow a complex SQL transformation using standard database functions (including User Defined Functions) to be built once and reused in multiple column transformations. Functionality is included to manage and propagate changes to re-usable 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

After selecting **Define Re-Usable Transformations** from the **Tools** menu the following dialog is displayed:



The image shows a 'Transformation Definition' dialog box with a blue title bar and a close button. It has three tabs: 'Re-usable Transformation' (selected), 'Model', and 'History'. The main area contains instructions on re-usable transformations, a 'Transform Name' dropdown, and 'New Transform' and 'Delete Transform' buttons. Below is a large text area for the 'Transform Description'. Further down, it provides an example of a 'replace' function with its tag, syntax, and description. At the bottom, there are input fields for 'Function Tag', 'Function Syntax', 'Function Description', 'Function Model', 'Created', and 'Updated', along with 'OK', 'Cancel', and 'Apply' buttons.

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

Click **New Transform** and enter a name for the re-usable transformation:

A screenshot of a 'User Defined Transformation' dialog box. The dialog has a blue title bar with the text 'User Defined Transformation' and a red close button. The main area is light gray and contains the 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 this text is a text input field containing the text 'TRIM_TODATE'. At the bottom of the dialog are two buttons: 'OK' and 'Cancel'.

Note: This is the internal QAD Data Warehouse Designer 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 function name visible to the users, which can be selected 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.

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.

Oracle Example of Building a Transformation Model

To build a model TO_DATE, converting a trimmed string in YYYYMMDD format to a date, do the following:

- ☐ Click on the **Model** tab
- ☐ Expand **Functions** heading
- ☐ Expand the **Conversion** heading
- ☐ Double-click on **To Date (TO_DATE)**
- ☐ Highlight **numeric_column**
- ☐ Expand the **String** heading
- ☐ Double-click on **Right Trim (RTRIM)**
- ☐ Highlight **[format]** and type 'YYYYMMDD'

You should see the following:



Now 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: TRIM_TODATE

New Transform

Delete Transform

Transform Description:

Trim and convert to a date

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: TDATEYYYYMMDD

Function Syntax: TO_DATE(RTRIM(string_column),'YYYYMMDD')

Function Description: Trim and then convert to a date from a string in the format YYYY-MM-DD

Function Model: TO_DATE(RTRIM(string_column [,character']), 'YYYYMMDD')

Created: 2013-01-31 14:15:37

Updated: 2013-02-20 16:03:57

OK

Cancel

Apply

622

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 Re-usable Transformation* (on page 616), the SQL was:


TO_DATE(RTRIM(string_column),'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 on the **Variable** button. This makes the word **format** a variable than can be substituted when the Re-usable 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 Re-usable Transformation to be applied, the error message above will include a count of the failures. The results window will detail the columns (and tables) where the update failures have occurred:



Object	Message
Transform	The following transformations could not be changed: stage_order_header.ship_date stage_order_header.order_date

Using Re-usable Transformations

Re-usable 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 605)

Chapter 20

Procedures and Scripts

QAD Data Warehouse Designer has a **Procedure** object group for database stored procedures and a **Script** object group for host system scripts, such as UNIX shell scripts or Windows batch files.

Procedures

QAD Data Warehouse Designer 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 or UNIX 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.

In This Chapter

Procedure Generation.....	628
Procedure Editing	631
Procedure Loading and Saving	634
Procedure Comparisons	636
Procedure Compilation.....	637
Procedure Running.....	638
Procedure Syntax	639
Procedure Properties.....	640
Script Generation	642
Script Editing	662
Script Testing	663
Script Syntax	664
Calling a Batch File from a Script	666
Scheduling Scripts.....	669
Templates	671

Procedure Generation

QAD Data Warehouse Designer 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 a drop-down box that is used to display the update, initial build, or post load procedure in a table's properties.

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.

For **Dimension**, **Stage**, **Fact** and **Aggregate** tables an Update or Initial Build procedure can be generated by selecting the option above from the appropriate drop-down box. When an 'Update' procedure generation is selected for a Dimension a **get_key function** is also automatically created.

If a new procedure is created from scratch (i.e. not auto generated) then an outline of the syntax required by the QAD scheduler can be generated by selecting the **Tools/Create Procedure Outline** menu option when in the procedure editor.

Wrapper procedures

In some cases multiple procedures will be 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.

For example: We may have a fact table that is updated from multiple stage tables. The Oracle wrapper procedure may look as follows:


```

CREATE OR REPLACE procedure update_fact_sales_detail
(
  v_sequence          IN   number,
  v_job_name          IN   varchar2,
  v_task_name         IN   varchar2,
  v_job_id            IN   number,
  v_task_id           IN   number,
  v_return_msg        OUT  varchar2,
  v_status            OUT  number
)
AS
-----
-- Control variables used in most programs
-----
v_MsgText          varchar2(255); -- Text for audit_trail
v_step             integer := 0;  -- return code
v_invoice_msg      varchar2(256); -- from invoice update
v_order_msg        varchar2(256); -- from order update
v_invoice_status    number  := 0;  -- from invoice update
v_order_status     number  := 0;  -- from order update
v_retcode          varchar2(1);

BEGIN

-----
-- Update fact table with info from orders stage table
-- Report the results to the audit trail.
-----
v_step := 10;
update_fact_sales_detail_o(v_sequence,v_job_name,v_task_name,
                          v_job_id, v_task_id,v_order_msg, v_order_status);
If v_order_status = 1 Then
  v_retcode := 'S';
ElsIf v_order_status = -1 Then
  v_retcode := 'W';
ElsIf v_order_status = -2 Then
  v_retcode := 'E';
ElsIf v_order_status = -3 Then
  v_retcode := 'F';
End If;
v_status := WsWrkAudit(v_retcode,v_job_name,v_task_name, v_sequence,
                      v_order_msg,NULL,NULL,v_task_id, v_job_id);

-----
-- If The previous update (orders) was successful
-- Update fact table with info from invoice stage table
-----
If v_order_status > -2 Then

```

```

        v_step := 20;
        update_fact_sales_detail_i(v_sequence,v_job_name,v_task_name,
                                   v_job_id, v_task_id,v_invoice_msg, v_invoice_status);
        If v_invoice_status = 1 Then
            v_retcode := 'S';
        ElseIf v_invoice_status = -1 Then
            v_retcode := 'W';
        ElseIf v_invoice_status = -2 Then
            v_retcode := 'E';
        ElseIf v_invoice_status = -3 Then
            v_retcode := 'F';
        End If;
        v_status := WsWrkAudit(v_retcode,v_job_name,v_task_name, v_sequence,
                               v_invoice_msg,NULL,NULL,v_task_id, v_job_id);
    Else
        v_invoice_status := 0;
    End If;

    -----
    -- All Done report the results
    -- If any of our called procedures failed or gave
    -- warnings then report the final status as such
    -----
    v_step := 90;
    v_status := v_order_status;
    If v_invoice_status < v_order_status Then
        v_status := v_invoice_status;
    Else
        v_status := v_order_status;
    End If;
    v_return_msg := 'fact_sales_detail update ';
    If v_status = 1 Then
        v_return_msg := v_return_msg||'completed.';
    ElseIf v_status = -1 Then
        v_return_msg := v_return_msg||'completed with warnings.';
    Else
        v_return_msg := v_return_msg||'Failed.';
    End If;

    RETURN;

EXCEPTION
    WHEN OTHERS THEN
        v_return_msg := 'Unhandled Exception in '||'update_fact_sales_detail. '||
                        ' Step '||v_step||' '||SQLERRM;
        v_status := WsWrkAudit('F',v_job_name,v_task_name, v_sequence,
                               v_return_msg,SQLCODE,SQLERRM,v_task_id, v_job_id);
        v_status := -3;

        RETURN;
END update_fact_sales_detail;

```

This Oracle wrapper procedure calls two child procedures, one for each stage table that is to update the fact table. A status is reported back to the audit trail for each stage and an overall status ascertained for the fact table update.

Procedure Editing

QAD Data Warehouse Designer 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 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 mouse popup menu.

Indenting a block of text

A selected block of text can be indented by four spaces by depressing the **tab** character. Each tab will indent by a further four spaces. A **shift/tab** will remove four 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 right-clicking and selecting **Comment**. 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 uncommented by choosing the **Uncomment** option. Note that only lines that start with two dashes in the left most column can be uncommented.

Inserting Steps

The right-click menu option **Insert Step** will insert a code line of the format 'v_step := 1000;'. Each subsequent insert will add 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 will be automatically reset for all steps numbering from 100 in increments of 100 when the procedure is compiled or saved.

Inserting Audit Message


The right-click menu option **Insert Audit Log Call** inserts several lines of code into the procedure which calls WsWrkAudit. This API is used to log messages to the audit log in the QAD Data Warehouse Designer scheduler. By default, the message is generated to log the current step number. It can be changed by altering the first added line. The Alt/A key can also be used instead of the right-click menu option.

Inserting Detail/Error Message

The right-click menu option **Insert Detail/Error Log Call** inserts several lines of code into the procedure which calls WsWrkError. This API is used to log messages to the detail/error log in the QAD Data Warehouse Designer scheduler. By default, the message is generated to log the current step number. It can be changed by altering the first added line. The Alt/D key can also be used instead of the right-click menu option.

Viewing other procedural code

During the editing process it is possible to pop up a window containing other procedural code. This window allows 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 will appear as follows:



There are a number of drop-down list boxes to choose 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 QAD metadata.
- ☐ Template: A template procedure as defined in the QAD 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. Any changes made in the viewer window cannot be saved. An example of an active view window is as follows:

```

-- DBMS Name      :      SQL Server
-- Script Name     :      update_dim_customer
-- Description      :      Update the dimension d
-- Generated by    :
-- Generated for    :
-- Generated on     :      Monday, February 11, 2
-- Author          :      John Smith
-- 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
SET NOCOUNT ON   -- Turn off row count m

-- Control variables used in most programs
--
DECLARE
    @v_msgtext      varchar(255) -- Text
    , @v_sql         nvarchar(255) -- Text
    , @v_set         integer -- commi
    , @v_analyze_flag integer -- analy
    , @v_step        integer -- retur
    , @v_update_count integer -- no of
    , @v_insert_count integer -- no of
    , @v_count        integer -- Gener
    , @v_db_code     varchar(10) -- Datab
    , @v_db_msg      varchar(100) -- Datab

-- General Variables
--
DECLARE
    @v_return_status integer -- Updat
    , @v_row_count    integer -- Gener
    , @v_getkey_status integer -- GetKe
    , @v_dss_current_flag char(1) -- Used
    , @v_dss_source_system_key integer -- Used

```

Procedure Loading and Saving

Procedures are normally stored in the QAD Data Warehouse Designer metadata tables. When a procedure is opened within QAD Data Warehouse Designer 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 the case of Oracle a compiled procedure can be viewed via the 'User_source' view (e.g. select text from user_source where name = 'MY_PROCEDURE' order by line;). For SQL Server the command is sp_helptext 'MY_PROCEDURE'.

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. All procedures are also written to one text file as part of the **Backup/Export the metadata** menu option.

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 of creating a version of a procedure. These are:

- 1 By setting **Auto-Version before Procedure Compile** under **Tools/Options/Versioning**. If set to true, a new version of a procedure will be created whenever the procedure is compiled.
- 2 The Procedure Editor menu option **File/Save and Version** will save a procedure and create a version at the same time.
- 3 By selecting the **Version Control/New Version** menu option from the pop-up menu when positioned on a procedure in the main Builder window.
- 4 By selecting the **Tools/Version All** menu option.

When a version is created via method (2) or (3) above the following screen will appear, to allow 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)




The image shows a 'Version Definition' dialog box with a blue title bar and a red close button. The main area is light gray and contains the text 'Create version for Procedure update_dim_date'. Below this, there are two input fields: 'Version Name or Short Description:' with a text box and a vertical scrollbar, and 'Detailed Description:' with a larger text box and a vertical scrollbar. At the bottom left, there is a 'Retain Until:' label followed by a date picker showing 'Tuesday , January 31, 2023'. At the bottom right, there are 'OK' and 'Cancel' buttons.

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.

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** allows 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 will compare 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:



```
--=====
-- MAIN
--=====

-->-- select @v_step = 0
BEGIN TRY

-- Set initial variable values
SELECT
    @v_set          = 0
  , @v_step         = 1
  , @v_analyze_flag = 0
  , @v_update_count = 0
  , @v_insert_count = 0
  , @v_insert_ind   = 0
  , @v_version_ind  = 0
  , @v_dss_version  = 0

-- Set values for unknown dimension record
SELECT
    @v_step          = 1
  , @v_dim_code       = NULL
  , @v_dim_cname      = SUBSTRING('Unknown',1,45)
  , @v_dim_address    = SUBSTRING('Unknown',1,40)
  , @v_dim_city       = SUBSTRING('Unknown',1,30)
  , @v_dim_state      = SUBSTRING('Unknown',1,2)
  , @v_dim_dss_current_flag = 'Y'
  , @v_dim_dss_version = 1
  , @v_dim_dss_start_date = NULL
  , @v_dim_dss_end_date = NULL
  , @v_dim_dss_update_time = GETDATE()

--=====
-- Every dimension table should have a 0 row for
-- when the dimension is null
--=====

-->-- set @v_step = 10
SET @v_count = 0


SELECT @v_count = 1
FROM dim_customer
WHERE dim_customer key = 0
```

In this example the lines **select @v_step = 0** and **set @v_step = 10** have been removed from the current code in the edit window.

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** will remove the added blue comments and code. The menu option **Tools/Accept Compare Changes** will implement the changes highlighted. For the above example the line **select @v_step = 0** and **set @v_step = 10** would be added.

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 will appear notifying of a successful compile. If the compile fails then error message comments will be inserted into the procedure code. In the following example the error messages are in red and begin with **--E--**.



```
-----
-- MAIN
-----
SELECT @v_step = 0
BEGIN TRY

-- Set initial variable values
SELECT
    @v_set           = 0
    , @v_step        = 1
    , @v_analyze_flag = 0
    , @v_update_count = 0
    , @v_insert_count = 0
    , @v_insert_ind   = 0
    , @v_version_ind  = 0
    , @v_dss_version  = 0

-- Set values for unknown dimension record
SELECT
    @v_step           = 1
    , @v_dim_code      = NULL
    , @v_dim_cname     = SUBSTRING('Unknown',1,45)
    , @v_dim_address   = SUBSTRING('Unknown',1,40)
    , @v_dim_city      = SUBSTRING('Unknown',1,30)
    , @v_dim_state     = SUBSTRING('Unknown',1,2)
    , @v_dim_dss_current_flag = 'Y'
    , @v_dim_dss_version = 1
    , @v_dim_dss_start_date = NULL
    , @v_dim_dss_end_date = NULL
    , @v_dim_dss_update_time = GETDATE()

-----
-- Every dimension table should have a 0 row for
-- when the dimension is null
-----
--E--line 125-- Incorrect syntax near ':'.'
SET @v_step := 10
SET @v_count := 0

SELECT @v_count = 1
FROM dim_customer
WHERE dim_customer_key = 0
```

Error comments will be inserted at each error point. *A compile will delete 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.

If a procedure fails to compile then it is invalidated in the database, and will not run until successfully compiled.

Procedure Running

Only procedures that conform to the QAD scheduler syntax can be executed from within the procedure editor. Select the **Execute/Execute** menu option or click the **Execute** icon to run the procedure. A procedure must have been compiled in order to run.

The results of the procedure will be displayed in a dialog box. The result code and result message will be displayed as well as any additional messages.

Procedure Syntax

The procedures managed by the QAD 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/Output	Oracle Type	SQL Server/DB2 Type
p_sequence	Input	Number	Integer
p_job_name	Input	Varchar2	Varchar(256)
p_task_name	Input	Varchar2	Varchar(256)
p_job_id	Input	Number	Integer
p_task_id	Input	Number	Integer
p_return_msg	Output	Varchar2	Varchar(256)
p_status	Output	Number	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

Procedure Properties

The properties screen for procedures and scripts is the same. A procedure can be renamed by changing the name field. If renamed then it will also be necessary to change the procedure name within the actual code. The purpose and owner fields are purely informational.

The screenshot shows a dialog box titled "update_dim_product Definition". On the left is a sidebar with "Procedure Properties" and "Notes". The main area contains the following fields:

- Name:** update_dim_product
- Type:** Procedure (dropdown menu)
- Purpose:** update_dim_product (text area)
- Owner:** auto created
- Last Update By:** (empty text field)
- Delete Lock:** ☐ (checkbox)
- Checkout section:**
 - Checked Out By:** John Smith (text field)
 - Clear Checkout:** (button)
 - Checkout Reason or Last Update:** testing (text area)
- Timestamps section:**
 - Created:** 2013-01-16 17:05:18
 - Last Update:** (empty text field)
 - Compiled:** 2013-01-16 17:05:18

At the bottom are buttons for OK, Cancel, and Help.

In the example above the **Delete Lock** checkbox is checked. This will prevent the procedure being deleted through the **Delete** menu option. It will also prevent the procedure 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 checked out. A **Clear Checkout** button to the right of the check out message allows the check out to be cleared. If QAD Data Warehouse Designer, the database or the PC crashes when a procedure is open, then the check out will need to be cleared through this screen.

The check out reason is for information only, and can be used as another comment field if desired.

The **Type** drop-down list allows the selection of **Block**, **Function** or **Procedure**:

The screenshot shows a dropdown menu for the "Type" field. The menu is open, showing three options: "Block" (selected), "Function", and "Procedure".

Selecting a type of **Block** will allow you to execute a SQL block against another connection. An additional field on the Properties screen; namely **Default Connect** then allows you to select the connection against which the SQL block will be executed.

The screenshot shows the 'update_dim_product Definition' dialog box. The 'Name' field is 'update_dim_product', 'Type' is 'Block', and 'Purpose' is 'update_dim_product'. The 'Default Connect' dropdown menu is open, showing a list of connections: DataWarehouse, Sales, Sales Forecasts, Sales Link, Sales SSIS, Sales2, SSAS Cubes, Tutorial (DLTP), and Windows. The 'Owner' field is 'auto created' and the 'Delete Lock' checkbox is unchecked. The 'Last Update By' field is empty. The 'Timestamps' section shows 'Created: 2013-01-16 17:05:18', 'Last Update:' (empty), and 'Compiled: 2013-01-16 17:05:18'. The 'OK', 'Cancel', and 'Help' buttons are at the bottom right.

Timestamps		
Created:	Last Update:	Compiled:
2013-01-16 17:05:18		2013-01-16 17:05:18

Script Generation

QAD Data Warehouse Designer generates template scripts to assist in the loading of textual files from UNIX or Windows. These scripts are generated when a UNIX or 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 will probably need modification.

Note: UNIX scripts are only supported for an Oracle data warehouse running on a UNIX platform.

The following sample scripts are covered in the remainder of this section:

- 1 Windows script for SQL.
- 2 UNIX script for Oracle.
- 3 Windows script for Oracle.
- 4 Windows script for DB2.

Script Generation (Windows/SQL Server)

A sample Windows script for SQL Server is as follows. The key components of the script are described below:

```
@echo off
setlocal enabledelayedexpansion
setlocal enableextensions
REM *****
REM ***** Load of D:\forecast.txt into load_forecast.txt
REM *****

REM *****
REM ***** NOTE: The following environment variables will be set
REM ***** WSL_LOAD_TABLE = the load table name
REM ***** WSL_SEQUENCE = a unique sequence number for the scheduler
REM ***** WSL_WORKDIR = the work directory defined in the connection
REM ***** WSL_SERVER = the server defined in the connection
REM ***** WSL_DATABASE = the database defined in the connection
REM ***** WSL_USER = the dss user defined in the connection
REM ***** WSL_PWD = the dss password defined in the connection
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 *****

SET LOAD_FILE=D:\forecast.txt
SET TRIG_FILE=D:\forecast.trg

if defined WSL_LOAD_TABLE (SET LOAD_TABLE=%WSL_LOAD_TABLE%) else SET LOAD_TABLE=load_forecast_txt
if defined WSL_SERVER (SET SERVER=%WSL_SERVER%) else SET SERVER=PMSDEV\PMSDEV
if defined WSL_DATABASE (SET DATABASE=%WSL_DATABASE%) else SET DATABASE=DataWarehouse
if defined WSL_USER (SET USER=%WSL_USER%) else SET USER=sa
if defined WSL_PWD (SET PWD=%WSL_PWD%) else SET PWD=9986
if defined WSL_WORKDIR (SET WORKDIR=%WSL_WORKDIR%) else goto LABEL_FILE_DEFAULT
if defined WSL_SEQUENCE (SET SEQUENCE=%WSL_SEQUENCE%) else goto LABEL_FILE_DEFAULT
SET FILECTL=%WORKDIR%\wsl%SEQUENCE%.ctl
SET FILELOG=%WORKDIR%\wsl%SEQUENCE%.log
SET FILEAUD=%WORKDIR%\wsl%SEQUENCE%.aud
goto LABEL_WAIT
:LABEL_FILE_DEFAULT
SET FILECTL=c:\temp\wsl39.ctl
SET FILELOG=c:\temp\wsl39.log
SET FILEAUD=c:\temp\wsl39.aud
:LABEL_WAIT
```

The script makes use of a number of environmental variables. These variable are acquired from both the table and connection properties. These variables are established in the environment by either QAD Data Warehouse Designer or the scheduler. If the script is to be executed outside of QAD Data Warehouse Designer 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 ***** 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
Call WslSleep to wait for 30 seconds. It will return remaining seconds
D:\Pms\Release\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_TRIGGER
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 *****
:LABEL_TRIGGER
REM
REM First clear any existing parameters
REM
isql -S %SERVER% -d %DATABASE% -U%USER% -P%PWD% -n -Q 'UPDATE dss_parameter SET dss_parameter_value = NULL WHERE dss_parameter_name like 'FORECAST_%' -o %FILELOG%
REM
REM Get the first four and all remaining parameters
REM
for /F "tokens=1,2,3,4* delims=" %a in (%TRIG_FILE%) do SET A=%a&SET B=%b&SET C=%c&SET D=%d&SET E=%e
REM
REM For each parameter write it to the parameter table
REM
if defined A isql -S %SERVER% -d %DATABASE% -U%USER% -P%PWD% -n -Q "EXEC VsParameterWrite 'FORECAST_0', '%A%', 'script load'"
if defined B isql -S %SERVER% -d %DATABASE% -U%USER% -P%PWD% -n -Q "EXEC VsParameterWrite 'FORECAST_1', '%B%', 'script load'"
if defined C isql -S %SERVER% -d %DATABASE% -U%USER% -P%PWD% -n -Q "EXEC VsParameterWrite 'FORECAST_2', '%C%', 'script load'"
if defined D isql -S %SERVER% -d %DATABASE% -U%USER% -P%PWD% -n -Q "EXEC VsParameterWrite 'FORECAST_3', '%D%', 'script load'"
if defined E isql -S %SERVER% -d %DATABASE% -U%USER% -P%PWD% -n -Q "EXEC VsParameterWrite 'FORECAST_4', '%E%', 'script load'"
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 OK 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 ***** L O A D   T H E   D A T A *****
REM *****
echo load of %LOAD_TABLE% > %FILEAUD%
SET /A RESULT_CODE=1
SET RESULT_MSG=Load Completed
:LABEL_LOAD
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
echo loading %FILE_NAME% >> %FILEAUD%
echo exit( >> %FILECTL%
echo BULK INSERT >> %FILECTL%
echo %LOAD_TABLE% >> %FILECTL%
echo FROM >> %FILECTL%
echo '%FILE_NAME%' >> %FILECTL%
echo WITH >> %FILECTL%
echo ( >> %FILECTL%
echo DATAFILETYPE='char' >> %FILECTL%
echo , CODEPAGE='raw' >> %FILECTL%
echo , FIELDTERMINATOR = ',' >> %FILECTL%
echo , FIRSTROW = 2 >> %FILECTL%
echo ) >> %FILECTL%
echo ) >> %FILECTL%
rem
rem In the following isql statement use either -E for a trusted connection
rem or -Uxxx -Pyyy where xxx is the username and ppp the password
rem
isql -S %SERVER% -d %DATABASE% -U%USER% -P%PWD% -n -i %FILECTL% -o %FILELOG%
SET ERRLEV=%errorlevel%
IF %errorlevel% EQU 0 GOTO LABEL_OKAY
:LABEL_FAIL
echo -2
echo Load failed with return status %ERRLEV% See error log for details
type %FILEAUD% >&2
type %FILELOG% >&2
exit
:LABEL_OKAY
type %FILELOG% >> %FILEAUD%
REM

```

This section calls isql to invoke Bulk Insert to load the file. It makes use of a temporary file to build as a control file and then calls isql 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.



TIP: If this script is to be executed outside the control of the QAD Data Warehouse Designer 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:

```
isql -S %SERVER% -d %DATABASE% -U%USER% -P%PWD% -n -Q "truncate table load_forecast" -o %FILELOG%
```

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 ***** R E N A M E *****
REM *****
:LABEL_RENAME
REM
REM rename the file
move %FILE_NAME% D:\loaded\
echo %FILE_NAME% moved to D:\loaded\ >> %FILEAUD%
REM rename the trigger file
move %TRIG_FILE% D:\loaded\
echo %TRIG_FILE% moved to D:\loaded\ >> %FILEAUD%
REM
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%
type %FILEAUD%
exit
```

Script Generation (UNIX/Oracle)

A sample UNIX script for Oracle is as follows. The key components of the script are described below:

```
#!/bin/sh
#####
##### Load of /home/dssadm/budget.txt into table load_budget_txt
#####
#
LOAD_FILE=/home/dssadm/budget.txt
#
#####
##### W A I T      T I M E R #####
#####
# If a wait time specified then loop looking for the file or trigger
# to arrive until the wait period expires.
WAITSECS=0
if [ "$WAITSECS" -gt "0" ]
then
    # Get the current time and work out the end of the wait
    #
    NOW_HH=`date +%H`
    NOW_MM=`date +%M`
    NOW_SS=`date +%S`
    let NOW_HH="$NOW_HH * 3600"
    let NOW_MM="$NOW_MM * 60"
    let NOW_TIME="$NOW_HH + $NOW_MM + $NOW_SS"
    let TILL_TIME="$NOW_TIME + $WAITSECS"
    let START_TIME="$NOW_TIME"
    #
    # Loop checking for the file until the wait is complete
    #
    while [ "$NOW_TIME" -lt "$TILL_TIME" ]
    do
        FCOUNT=`find $LOAD_FILE -print 2>/dev/null | wc -l`
        if [ "$FCOUNT" -gt "0" ]
        then
            let NOW_TIME="$TILL_TIME"
        else
            #
            # if the file is not present sleep 30 seconds
            # and recalculate the time. If less than the start
            # then we have gone past midnight so add 86400
            sleep 30
        fi
    done
fi
```

```

NOW_HH=`date +%H`
NOW_MM=`date +%M`
NOW_SS=`date +%S`
let NOW_HH="$NOW_HH * 3600"
let NOW_MM="$NOW_MM * 60"
let NOW_TIME="$NOW_HH + $NOW_MM + $NOW_SS"
if [ "$NOW_TIME" -lt "$START_TIME" ]
then
    let NOW_TIME="$NOW_TIME + 86400"
fi
done
fi
#
# Finished our wait loop. See if we have the file
#
#*****
#***** F I L E   C H E C K   *****
#*****
FCOUNT=`find $LOAD_FILE -print 2>/dev/null | wc -l`
if [ "$FCOUNT" -eq "0" ]
then
    #
    # file not present so exit with the requested return code
    #
    echo "-1"
    echo "File /home/dssadm/budget.txt was not found"
    exit
fi

```

The script makes use of two environmental variables called DSS_USER and DSS_PWD. These variables are established in the environment by the scheduler. If the script is to be executed outside of scheduler control then these two variables will need to be assigned after the LOAD_FILE variable.

The first section of the script defines the load_file variable. 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.

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 the file based load is used instead of the script based load, the scheduler looks for a trigger file if defined, and will exit with the specified status if the file is not found.

```
#####
##### T R I G G E R #####
#####
#
# Clear any existing Trigger file parameters
#
### sqlplus -s <<EOF >/dev/null
### $DSS_USER/$DSS_PWD
### set sqlprompt "";
### set heading off;
### set pagesize 0
### set linesize 256
### set trimspool on
### set echo off;
### update dss_parameter set dss_parameter_value = NULL
### where dss_parameter_name like 'budget.txt_%';
### exit;
### EOF
### if [ "$?" -ne "0" ]
### then
###     echo "-1"
###     echo "Parameter clear Sqlplus returned a non standard return code of $?"
### fi
#
# Load the contents of the trigger file into parameters
#
### cat /home/dssadm/budget.trg | tr "\012" "\012" > /tmp/wsl2.trg
### TRIGNO=`cat /tmp/wsl2.trg | wc -l | tr -d " "`
### if [ "$TRIGNO" != "0" ]
### then
###     ROWNUM=0
###     while [ "$ROWNUM" -lt "$TRIGNO" ]
###     do
###         let ROWNUM="$ROWNUM+1"
###         TRIG_NO="budget.txt_$ROWNUM"
###         TRIG_MSG=`cat /tmp/wsl2.trg | head -$ROWNUM | tail -1 | sed "s/'/'/g"`
###         sqlplus -s <<EOF1 >/dev/null
###         $DSS_USER/$DSS_PWD
###         variable x number;
###         exec :x := WsParameterWrite('$TRIG_NO','$TRIG_MSG','load of /home/dssadm/budget.txt');
###         exit;
###     EOF1
###     done
### fi
```

This section of the script is commented out by default. It is used when a trigger file is present. Such a file 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 OK 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.

```
#####
##### L O A D   T H E   D A T A #####
#####
#
# Loop through all files found loading them into the table
# Will need to use append in the sql load in case there are multiple files
#
FLIST=`find $LOAD_FILE -print`
for ACTFILE in $FLIST
do
    # Build the sql loader control script
    echo "load data" >>/tmp/wsl2.ct1
    echo "infile '$ACTFILE'" >>/tmp/wsl2.ct1
    echo "badfile '/tmp/wsl2.bad'" >>/tmp/wsl2.ct1
    echo "append" >>/tmp/wsl2.ct1
    echo "into table load_budget_txt" >>/tmp/wsl2.ct1
    echo "fields terminated by \",\"" >>/tmp/wsl2.ct1
    echo "optionally enclosed by '\"'" >>/tmp/wsl2.ct1
    echo "trailing nullcols" >>/tmp/wsl2.ct1
    echo "(" >>/tmp/wsl2.ct1
    echo "order_number," >>/tmp/wsl2.ct1
    echo "order_line_number," >>/tmp/wsl2.ct1
    echo "product_code," >>/tmp/wsl2.ct1
    echo "customer_code," >>/tmp/wsl2.ct1
    echo "budget_quantity," >>/tmp/wsl2.ct1
    echo "budget_sales," >>/tmp/wsl2.ct1
    echo "budget_date date 'yyyy-mm-dd'" >>/tmp/wsl2.ct1
    echo ")" >>/tmp/wsl2.ct1
    ORACLE_SID=0817
    export ORACLE_SID
    sqlldr userid=$DSS_USER/$DSS_PWD control=/tmp/wsl2.ct1 skip=1 silent=header,feedback log=/tmp/wsl2.log
    retcode=`echo $?`
    rows=`cat /tmp/wsl2.log| grep Rows | grep success`
done
```

This section calls sql*loader to load the file. It makes use of a temporary file to build as a control file and then calls sqlldr 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.



TIP: If this script is to be executed outside the control of the QAD Data Warehouse Designer 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 "truncate table load_budget_txt;" >/tmp/wsl2.sql
echo "commit;" >>/tmp/wsl2.sql
echo "exit" >>/tmp/wsl2.sql
sqlplus $DSS_USER/$DSS_PWD @/tmp/wsl2.sql >/dev/null
```

Note that the first echo has a single output '>' whereas the subsequent lines have two '>>'. The first output command creates or re-creates the file and the subsequent lines append to it. We put the output of the command to the null device as we do not want data in the output stream that does not match our syntax.

```
case "$retcode" in
0) echo "1";
   echo "load_budget_txt completed. $rows from $ACTFILE"
   ;;
2) echo "-1";
   echo "Not all rows loaded. $rows. see error trail" ;
   cat /tmp/wsl2.log| grep Rows
   cat /tmp/wsl2.log>&2
   ;;
1) echo "-2";
   echo "SQL*Loader execution exited with EX_FAIL, see error trail";
   cat /tmp/wsl2.log>&2
   ;;
3) echo "-3";
   echo "SQL*Loader execution encountered a fatal error"
   cat /tmp/wsl2.log>&2
   ;;
*) echo "-3";
   echo "SQL*Loader unknown return code"
   cat /tmp/wsl2.log>&2
   ;;
esac
```

This section of the script checks the return code from the sql*loader command. Depending on the code it either identifies success, a warning or an error. In any non success situation it puts the log file out to the error stream '>&2'. This will result in the information ending up in the Detail/Error log when this script is executed through the scheduler.

```
#
#####
##### R E N A M E #####
#####
#
# If we had a success or a warning then do a rename
### if [ "$retcode" -eq "0" -o "$retcode" -eq "2" ]
### then
#
# Rename the file
#
### mv $ACTFILE /home/dssadm/budget.txt
### echo "$ACTFILE renamed to /home/dssadm/budget.txt"
#
# Rename the trigger
#
### mv /home/dssadm/budget.trg /home/dssadm/budget.trg
### echo "/home/dssadm/budget.trg renamed to /home/dssadm/budget.trg"
### fi
#####
##### E N D   F I L E   L O A D   L O O P #####
#####
#
# Uncomment the break statement if you only wish to process one file at
# a time in a wildcard file load. The break will case the script to finish
# after the first file that matches the wildcard has been loaded
#
# break
done
exit
```

This final section is largely commented out. It ends with the closing of the for loop that processes each file in a wild card file load.

The first block of commented out code renames the file and also the trigger file if appropriate. If this action is required enter the rename path etc. and un-comment this code.

The break statement 3 lines from the end can be used if just the first file in a wild card file load is required. Simply uncomment this break statement and the script will end after the first file has been loaded.

Script Generation (Windows/Oracle)

A sample Windows script for Oracle is as follows. The key components of the script are described below:

```
@echo off
setlocal enabledelayedexpansion
setlocal enableextensions
REM *****
REM ***** Load of D:\forecast.txt into load_forecast_txt
REM *****

REM *****
REM ***** NOTE: The following environment variables will be set
REM ***** WSL_LOAD_TABLE = the load table name
REM ***** WSL_SEQUENCE = a unique sequence number for the scheduler
REM ***** WSL_WORKDIR = the work directory defined in the connection
REM ***** WSL_DBHOME = the database home defined in the connection
REM ***** WSL_DATABASE = the database defined in the connection
REM ***** WSL_USER = the dss user defined in the connection
REM ***** WSL_PWD = the dss password defined in the connection
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 *****

SET LOAD_FILE=D:\forecast.txt
SET TRIG_FILE=D:\forecast.trg

if defined WSL_LOAD_TABLE (SET LOAD_TABLE=%WSL_LOAD_TABLE%) else SET LOAD_TABLE=load_forecast_txt
if defined WSL_DBHOME (SET DBHOME=%WSL_DBHOME%) else SET DBHOME=
if defined WSL_DATABASE (SET DATABASE=%WSL_DATABASE%) else SET DATABASE=
if defined WSL_USER (SET USER=%WSL_USER%) else SET USER=dssadm
if defined WSL_PWD (SET PWD=%WSL_PWD%) else SET PWD=ws1
if defined WSL_WORKDIR (SET WORKDIR=%WSL_WORKDIR%) else goto LABEL_FILE_DEFAULT
if defined WSL_SEQUENCE (SET SEQUENCE=%WSL_SEQUENCE%) else goto LABEL_FILE_DEFAULT
SET FILECMD=%WORKDIR%\ws1%SEQUENCE%.cmd
SET FILECTL=%WORKDIR%\ws1%SEQUENCE%.ctl
SET FILELOG=%WORKDIR%\ws1%SEQUENCE%.log
SET FILEAUD=%WORKDIR%\ws1%SEQUENCE%.aud
SET FILEBAD=%WORKDIR%\ws1%SEQUENCE%.bad
SET FILELO2=%WORKDIR%\ws1%SEQUENCE%.lo2
goto LABEL_WAIT
:LABEL_FILE_DEFAULT
SET FILECMD=c:\temp\ws122.cmd
SET FILECTL=c:\temp\ws122.ctl
SET FILELOG=c:\temp\ws122.log
SET FILEAUD=c:\temp\ws122.aud
SET FILEBAD=c:\temp\ws122.bad
SET FILELO2=c:\temp\ws122.lo2
:LABEL_WAIT
```

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 QAD Data Warehouse Designer or the scheduler. If the script is to be executed outside of QAD Data Warehouse Designer 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 ***** 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
Call WslSleep to wait for 30 seconds. It will return remaining seconds
D:\Pms\Release\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_TRIGGER
echo -1
echo File %TRIG_FILE% was not found.
exit

```

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 *****
:LABEL_TRIGGER
REM
REM First clear any existing parameters
REM
echo UPDATE dss_parameter SET dss_parameter_value = NULL > %FILECMD%
echo WHERE dss_parameter_name like 'FORECAST_%'; >> %FILECMD%
echo variable X varchar2 >> %FILECMD%
REM
REM Get the first four and all remaining parameters
REM
for /F "tokens=1,2,3,4* delims=," %%a in (%TRIG_FILE%) do SET A=%%a&SET B=%%b&SET C=%%c&SET D=%%d&SET E=%%e
REM
REM For each parameter write it to the parameter table
REM
if defined A echo EXEC :X := WsParameterWrite('FORECAST_0', '%A%', 'script load'); >> %FILECMD%
if defined B echo EXEC :X := WsParameterWrite('FORECAST_1', '%B%', 'script load'); >> %FILECMD%
if defined C echo EXEC :X := WsParameterWrite('FORECAST_2', '%C%', 'script load'); >> %FILECMD%
if defined D echo EXEC :X := WsParameterWrite('FORECAST_3', '%D%', 'script load'); >> %FILECMD%
if defined E echo EXEC :X := WsParameterWrite('FORECAST_4', '%E%', 'script load'); >> %FILECMD%
REM
REM Execute the sql command
REM
echo exit; >> %FILECMD%
sqlplus -s %USER%/%PWD%@%DATABASE% @%FILECMD% > %FILELO2% 2>&1
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 OK 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 ***** L O A D   T H E   D A T A *****
REM *****
echo load of %LOAD_TABLE% > %FILEAUD%
SET /A RESULT_CODE=1
SET RESULT_MSG=Load Completed
:LABEL_LOAD
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
echo loading %FILE_NAME% >> %FILEAUD%
echo load data > %FILECTL%
echo infile '%FILE_NAME%' >> %FILECTL%
echo badfile '%FILEBAD%' >> %FILECTL%
echo into table %LOAD_TABLE% >> %FILECTL%
echo fields terminated by "," >> %FILECTL%
echo optionally enclosed by '"' >> %FILECTL%
echo trailing nullcols >> %FILECTL%
echo ( >> %FILECTL%
echo   product_code,>> %FILECTL%
echo   customer_code,>> %FILECTL%
echo   forecast_quantity,>> %FILECTL%
echo   forecast_sales_value,>> %FILECTL%
echo   forecast_date date 'dd-mon-yyyy'>> %FILECTL%
echo ) >> %FILECTL%
sqlldr userid=%USER%/%PWD%@%DATABASE% control=%FILECTL% skip=1 silent=header,feedback log=%FILELOG%
SET /A ERRLEV=%errorlevel%
```

This section calls sql*loader (bulk insert for SQL Server) to load the file. It makes use of a temporary file to build as a control file and then calls sqlldr 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.



TIP: If this script is to be executed outside the control of the QAD Data Warehouse Designer 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 for Oracle would look something like the following:

```
echo truncate table load_budget_txt; >/tmp/wsl2.sql
echo commit; >>/tmp/wsl2.sql
echo exit >>/tmp/wsl2.sql
sqlplus %USER%/%PWD%@%DATABASE% @/tmp/wsl2.sql > %FILELO2 2>&1
```

Note that the first echo has a single output '>' whereas the subsequent lines have two '>>'. The first output command creates or re-creates the file and the subsequent lines append to it. We put the output of the command to the null device as we do not want data in the output stream that does not match our syntax.

This next section handles the return status from the load. It returns the appropriate status and messages to the scheduler.

```
if %ERRLEV% GTR 3 goto LABEL_BADRESULT
if %ERRLEV% LSS 0 goto LABEL_BADRESULT
if %ERRLEV% EQU 0 goto LABEL_OKAY
if %ERRLEV% EQU 1 goto LABEL_WARNING
if %ERRLEV% EQU 2 goto LABEL_ERROR
if %ERRLEV% EQU 3 goto LABEL_FATAL
:LABEL_BADRESULT
echo -3
echo SQL*Loader execution encountered an unexpected fatal error
type %FILELOG% >&2
exit
:LABEL_ERROR
echo -2
echo SQL*Loader execution exited with EX_FAIL, see error trail
type %FILELOG% >&2
exit
:LABEL_FATAL
echo -3
echo SQL*Loader execution encountered a fatal error
type %FILELOG% >&2
exit
:LABEL_WARNING
SET /A RESULT_CODE=-1
SET RESULT_MSG=WARNING: Not all rows loaded. See error trail
type %FILELOG% >>&2
:LABEL_OKAY
type %FILELOG% >>&2
REM
```

This final section is as follows:

```
REM *****
REM ***** R E N A M E *****
REM *****
:LABEL_RENAME
REM
REM rename the file
move %FILE_NAME% D:\old_data
echo %FILE_NAME% moved to D:\old_data >> %FILEAUD%
REM rename the trigger file
move %TRIG_FILE% D:\old_data
echo %TRIG_FILE% moved to D:\old_data >> %FILEAUD%
REM
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%
type %FILEAUD%
exit
```

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.

Script Generation (Windows/DB2)

A sample Windows script for DB2 is as follows. The key components of the script are described below:

```
@echo off
setlocal enabledelayedexpansion
setlocal enableextensions
REM *****
REM ***** Load of c:\temp\forecast.txt into load_forecast
REM *****

REM *****
REM ***** NOTE: The following environment variables will be set
REM ***** WSL_LOAD_TABLE = the load table name
REM ***** WSL_LOAD_SCHEMA = the load table schema
REM ***** WSL_SEQUENCE = a unique sequence number for the scheduler
REM ***** WSL_WORKDIR = the work directory defined in the connection
REM ***** WSL_SERVER = the server defined in the connection
REM ***** WSL_DATABASE = the database defined in the connection
REM ***** WSL_USER = the dss user defined in the connection
REM ***** WSL_PWD = the dss password defined in the connection
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 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.

SET LOAD_FILE="c:\temp\forecast.txt"
SET TRIG_FILE="c:\temp\forecast.trg"

if defined WSL_LOAD_TABLE (SET LOAD_TABLE=%WSL_LOAD_TABLE%) else SET LOAD_TABLE=load_forecast
if defined WSL_LOAD_SCHEMA (SET LOAD_SCHEMA=%WSL_LOAD_SCHEMA%) else SET LOAD_SCHEMA=RED
if defined WSL_SERVER (SET SERVER=%WSL_SERVER%) else SET SERVER=
if defined WSL_DATABASE (SET DATABASE=%WSL_DATABASE%) else SET DATABASE=RED
if defined WSL_METABASE (SET METABASE=%WSL_METABASE%) else SET METABASE=RED
if defined WSL_USER (SET USER=%WSL_USER%) else SET USER=
if defined WSL_PWD (SET PWD=%WSL_PWD%) else SET PWD=
if defined WSL_WORKDIR (SET WORKDIR=%WSL_WORKDIR%) else goto LABEL_FILE_DEFAULT
if defined WSL_SEQUENCE (SET SEQUENCE=%WSL_SEQUENCE%) else goto LABEL_FILE_DEFAULT
SET FILECTL=%WORKDIR%\wsl%SEQUENCE%.ctl
SET FILELOG=%WORKDIR%\wsl%SEQUENCE%.log
SET FILEAUD=%WORKDIR%\wsl%SEQUENCE%.aud
goto LABEL_WAIT
:LABEL_FILE_DEFAULT
SET FILECTL=c:\temp\wsl74.ctl
SET FILELOG=c:\temp\wsl74.log
SET FILEAUD=c:\temp\wsl74.aud
:LABEL_WAIT
```

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 QAD Data Warehouse Designer or the scheduler. If the script is to be executed outside of QAD Data Warehouse Designer 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 ***** 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_TRIGGER
echo -1
echo File %TRIG_FILE% was not found.
exit
```

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 *****
:LABEL_TRIGGER
REM
REM First clear any existing parameters
REM
SET SQL=UPDATE %METABASE%.dss_parameter SET dss_parameter_value = NULL WHERE dss_parameter_name like 'FORECAST_%';
echo %SQL% > %FILECTL%
REM
REM Get the first four and all remaining parameters
REM
for /F "tokens=1,2,3,4* delims=," %a in ('type %TRIG_FILE%') do SET A=%a&SET B=%b&SET C=%c&SET D=%d&SET E=%e
REM
REM For each parameter write it to the parameter table
REM
if defined A SET SQL=CALL %METABASE%.WsParameterWrite ('FORECAST_0','%A%','script load');
if defined A echo %SQL% >> %FILECTL%
if defined B SET SQL=CALL %METABASE%.WsParameterWrite ('FORECAST_1','%B%','script load');
if defined B echo %SQL% >> %FILECTL%
if defined C SET SQL=CALL %METABASE%.WsParameterWrite ('FORECAST_2','%C%','script load');
if defined C echo %SQL% >> %FILECTL%
if defined D SET SQL=CALL %METABASE%.WsParameterWrite ('FORECAST_3','%D%','script load');
if defined D echo %SQL% >> %FILECTL%
if defined E SET SQL=CALL %METABASE%.WsParameterWrite ('FORECAST_4','%E%','script load');
if defined E echo %SQL% >> %FILECTL%
rem
db2batch -d %DATABASE% -f %FILECTL% -r %FILELOG% -time off -o p 0
SET ERRLEV=%errorlevel%
```

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 OK 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 ***** L O A D   T H E   D A T A *****
REM *****
echo load of %LOAD_TABLE% > %FILEAUD%
SET /A RESULT_CODE=1
SET RESULT_MSG=Load Completed
:LABEL_LOAD
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
echo loading %FILE_NAME% >> %FILEAUD%
echo CALL SYSPROC.ADMIN_CMD('LOAD FROM %FILE_NAME% OF DEL > %FILECTL%
echo MODIFIED BY COLDEL, timestampformat="DD-MMM-YYYY" >> %FILECTL%
echo METHOD P (1, 2, 3, 4, 5) >> %FILECTL%
echo MESSAGES ON SERVER INSERT INTO %LOAD_SCHEMA%.%LOAD_TABLE%( >> %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 NONRECOVERABLE'); >> %FILECTL%
rem
rem In the following sql statement no username/password is required for a trusted connection
rem or -axxx/yyy where xxx is the username and ppp the password. The slash is required
rem
db2batch -d %DATABASE% -f %FILECTL% -r %FILELOG% -time off -o p 0
SET ERRLEV=%errorlevel%

```

This section calls the DB2batch command to load the file. It makes use of a temporary file to build as a control file which includes a call to the LOAD statement. It then runs the temporary files using DB2batch 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.

This next section handles the return status from the load. It returns the appropriate status and messages to the scheduler.

```

SET ERRLEV=%errorlevel%
IF %errorlevel% EQU 0 GOTO LABEL_OKAY
:LABEL_FAIL
echo -2
echo Load failed with return status %ERRLEV% See error log for details
type %FILEAUD% >&2
type %FILELOG% >&2
exit
:LABEL_OKAY
type %FILELOG% >> %FILEAUD%

```


This final section is as follows:

```
REM *****
REM ***** R E N A M E *****
REM *****
:LABEL_RENAME
rem
SET YYYY=%DATE:~-4%
SET MM=%DATE:~-10,2%
SET DD=%DATE:~-7,2%
rem rename the file
move %FILE_NAME% "c:\temp\loadedforecast_%YYYY%MM%DD%.txt" >> %FILEAUD%
rem
SET ERRLEV=%errorlevel%
rem
IF %ERRLEV% EQU 0 GOTO LABEL_MOVE3
echo -3
echo %FILE_NAME% NOT MOVED >> %FILEAUD%
exit
GOTO LABEL_AFTER_MOVE3
:LABEL_MOVE3
rem
echo %FILE_NAME% moved to c:\temp\loadedforecast_%YYYY%MM%DD%.txt >> %FILEAUD%
:LABEL_AFTER_MOVE3
rem
rem rename the trigger file
move %TRIG_FILE% "c:\temp\loadedforecast_%YYYY%MM%DD%.trg" >> %FILEAUD%
rem
SET ERRLEV=%errorlevel%
IF %ERRLEV% EQU 0 GOTO LABEL_MOVE4
echo -3
rem
echo %TRIG_FILE% NOT MOVED >> %FILEAUD%
exit
GOTO LABEL_AFTER_MOVE4
:LABEL_MOVE4
rem
echo %TRIG_FILE% moved to c:\temp\loadedforecast_%YYYY%MM%DD%.trg >> %FILEAUD%
:LABEL_AFTER_MOVE4
REM
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%
type %FILEAUD%
exit
```

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.

Script Editing

QAD Data Warehouse Designer includes a script editor which allows 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 mouse popup menu.

Indenting a block of text

A selected block of text can be indented by four spaces by depressing the tab character. Each tab will indent by a further four spaces. A shift/tab will remove 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 allows 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 will appear.

A number of drop-down list boxes 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 QAD metadata.
- ☐ Template: A template script as defined in the QAD 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.

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. For a Windows script this is achieved by running the script on the current PC. For a UNIX script QAD Data Warehouse Designer will attempt to start a telnet window to the UNIX host (as defined in the default connection for the script) and run the script within this telnet window.

A script is invoked via the **Execute/Execute the Script** menu option. The output from the script is shown in a pop-up dialog box.

UNIX Scripts

UNIX scripts have the additional execution option being the menu option **Execute/ Execute and retain**. This option will create a Telnet window to the UNIX host, execute the script, display the results and leave the Telnet window up. Additional commands or a review of what occurred can then be undertaken directly in the Telnet window. When completed the Telnet window should be closed to return to the script editor.

Note: A connection must be correctly set-up to support UNIX script testing. See the **Installation and Administration Guide** and the section on UNIX connections.

Script Syntax

There are a number of conventions that must be followed if a host script is to be used by the QAD 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 QAD Data Warehouse Designer tool under the scheduler window.

Windows Example for Oracle:

In the following example the first line '@echo off' prevents unwanted information from being reported to standard out. A Sql*loader script file is built up (echo statements). The sqlldr command is then executed with the 'silent-header,feedback' option once again to prevent unwanted output.

```
@echo off
echo load data >c:\temp\wsl1.ctl
echo infile 'C:\Temp\budget.txt' >>c:\temp\wsl1.ctl
echo badfile 'c:\temp\wsl1.bad' >>c:\temp\wsl1.ctl
echo into table load_budget_txt >>c:\temp\wsl1.ctl
echo fields terminated by "," >>c:\temp\wsl1.ctl
echo optionally enclosed by '^"' >>c:\temp\wsl1.ctl
echo trailing nullcols >>c:\temp\wsl1.ctl
echo ( >>c:\temp\wsl1.ctl
echo  product_code,>>c:\temp\wsl1.ctl
echo  customer_code,>>c:\temp\wsl1.ctl
echo  budget_quantity,>>c:\temp\wsl1.ctl
echo  budget_sales_value,>>c:\temp\wsl1.ctl
echo  budget_date date 'dd-mon-yyyy'>>c:\temp\wsl1.ctl
echo ) >>c:\temp\wsl1.ctl
sqlldr userid=dssadm/wsl@ORCL control=c:\temp\wsl1.ctl skip=1
silent=header,feedback log=c:\temp\wsl1.log
goto answer%errorlevel%
:answer0
echo 1
echo Load completed successfully
type c:\temp\wsl1.log >&2
exit
:answer2
echo -1
echo Not all rows loaded.  See error trail
type c:\temp\wsl1.log >&2
exit
:answer1
echo -2
echo SQL*Loader execution exited with EX_FAIL, see error trail
type c:\temp\wsl1.log >&2
exit
:answer3
echo -3
echo SQL*Loader execution encountered a fatal error
type c:\temp\wsl1.log >&2
exit
```

Calling a Batch File from a Script

Below is an example Data Warehouse Designer 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 Data Warehouse Designer.

The screenshot shows the 'Host Script Call_Batch_File' dialog box. On the left is a sidebar with 'Properties' and 'Notes' tabs. The main area contains the following fields:

- Name:** Call_Batch_File
- Type:** Windows script (dropdown)
- Purpose:** Calls a batch file (text area)
- Owner:** dbo
- Last Update By:** (empty text box)
- Default Connect:** Windows (dropdown)
- Checkout:**
 - Checked Out By:** (greyed out text box)
 - Checkout Reason or Last Update:** New Script (text area)
- Timestamps:**
 - Created:** (empty text box)
 - Last Update:** (empty text box)
 - Compiled:** (empty text box)

At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

Edit the Script and enter the following:




The screenshot shows the SQL Enterprise Manager script editor with the following script content:

```
@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
```

The status bar at the bottom indicates: Ready | Development (Red_1) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS


Save the Script.



The screenshot shows the SQL Enterprise Manager script editor with the File menu open. The 'Save Script' option (Ctrl+S) is highlighted. The script content is partially visible in the background:

```
...EXPANSION
...IONS
...ile.bat > c:\temp\MyBatchFile.log 2>&1
...GOTO LABEL_OKAY
...ned an error code of %ERRORLEVEL%
...ile.log
ECHO Batch file completed successfully
TYPE c:\temp\MyBatchFile.log
```

When the script is executed, you will see the following results:




Object	Message
✓	# Hello
✓	# 1
✓	# Batch file completed successfully
✓	# Hello

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**.

A screenshot of the 'Host Script script_load_budget' Properties dialog box. The dialog has a title bar with the text 'Host Script script_load_budget' and a close button. On the left is a 'Properties' tab and a 'Notes' section. The main area contains the following fields:

- Name:** script_load_budget
- Type:** Windows script (dropdown menu)
- Purpose:** Auto generated file load (text area)
- Owner:** (empty text field)
- Last Update By:** (empty text field)
- Default Connect:** Windows (dropdown menu)
- Checkout:** section containing:
 - Checked Out By:** (empty text field)
 - Checkout Reason or Last Update:** (empty text area)
- Timestamps:** section containing:
 - Created:** 2013-08-22 15:22:27
 - Last Update:** (empty text field)
 - Compiled:** (empty text field)

At the bottom right are three buttons: 'OK', 'Cancel', and '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 QAD 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 QAD Data Warehouse Designer tool under the scheduler window.

Templates

Templates serve two main purposes. Firstly they provide a structure for the building of procedures and scripts. Secondly they provide example code to illustrate how to handle a specific application problem.

A simple template editor is provided to allow the maintenance of templates. Templates are normally either procedure or host script examples. The normal practice is to develop the template within either the procedure or script editor and then paste it into the template for storage.

Procedure Structure Templates


These templates are normally only defined at initial setup, and are in fact not visible or accessible in this version of QAD Data Warehouse Designer.

Submitting Example Templates

QAD welcomes the supply of new templates. If you develop custom code that you believe may be of value to other QAD Data Warehouse Designer users, please feel free to email a copy to <http://support.qad.com>. Author identities will be acknowledged.

Chapter 21

Scheduler

The scheduler is accessible by clicking the **Scheduler** button  on 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 on either a UNIX host or under Windows (as a system service). It processes predefined 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 QAD Data Warehouse Designer 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.....	674
Scheduler States	677
Scheduling a Job	679
Working with Jobs	687
Monitoring the Database and Jobs	732
Stand Alone Scheduler Maintenance.....	740
SQL to return Scheduler Status.....	743
Reset Columns in Job and Task View.....	744
Scheduling a Data Warehouse Designer Host Script from 3D	745

Scheduler Options

An example of the **Scheduler** screen is shown below.



The screenshot shows the Scheduler interface with 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 a refresh icon. The main area is divided into two panes. The top pane, titled 'Jobs', contains a table with columns: Job, Status, Seq, Start, Finish, Elapsed, OK, Info, Detail, Warn, Error, and Who. It lists three jobs: 'Daily Update' (On Hold), 'Process load_budget 6073' (Failed), and 'Daily Update' (Completed). The bottom pane, titled 'Tasks', contains a table with columns: Task, Action, Status, Seq, Start, Finish, Elap..., Info, De..., Wa..., and Result. It lists ten tasks, including loading data into tables like 'load_product', 'load_customer', 'load_order_line', and 'load_order_header', and updating dimensions like 'dim_customer' and 'dim_date'. The status bar at the bottom shows 'Ready', 'Database time at last action was 19:00:40 | Development', and 'Userid: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS'.

Job	Status	Seq	Start	Finish	Elapsed	OK	Info	Detail	Warn	Error	Who
Daily Update	On Hold	...	2013-08-16 08:46:21								JS
Process load_budget 6073	Failed	...	2013-08-21 15:24:52	2013-08-21 15:24:54	00:00		2			2	JS
Daily Update	Completed	...	2013-08-16 08:46:21	2013-08-16 08:46:24	00:00	11	26				JS

Task	Action	Status	Seq	Start	Finish	Elap...	Info	De...	Wa...	Result
load_product	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	4			9 rows loaded into load_product
load_customer	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	4			6 rows loaded into load_customer
load_order_line	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	4			21 rows loaded into load_order_line
load_order_header	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	4			9 rows loaded into load_order_header
dim_customer	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	2			dim_customer updated. 0 new records. 6 records updated.
dim_date	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	2			Date Dimension updated from Dec 8 1999 12:00AM to Apr 25 2027 ...
dim_product	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	2			dim_product updated. 0 new records. 9 records updated.
stage_sales_detail	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	2			stage_sales_detail updated. 21 new records. 0 records updated.
fact_sales_detail	Process	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00	2			fact_sales_detail updated. 0 new records. 21 records updated.
fact_sales_detail	Analyze	Compl...	6068	2013-08-16 08:...	2013-08-16 08:...	00:00				Table fact_sales_detail analyzed

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 . ie 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 - eg 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

Column	Description
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

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. When the currently running instance completes this job will start.
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 task 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 UNIX or 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.




The **Job Definition** dialog box is shown with the following fields and options:

- Job Name:** Job_1045
- Description:** Job number 1045
- Frequency:** Hold
- Start Date:** Thursday, August 22, 2013
- Start Time:** 7:47:00 AM
- Maximum Threads:** 1
- Scheduler:** Windows Only
- Dependent On:** Parent job, Fail, Look back (minutes), Maximum wait (minutes). Includes **Add Parent Job** and **Remove Parent** buttons.
- Logs Retained:** 0. Text: "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)".
- Custom Settings:**
 - Interval Between Jobs: (minutes)
 - Start At or After HHMM (e.g. 0800): 0000
 - Do Not Start After HHMM (e.g. 1700): 0000
 - Active Days: ☐ Mon ☐ Tue ☐ Wed ☐ Thu ☐ Fri ☐ Sat ☐ Sun
- Success Command:**
- Failure Command:**
- 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:

See **Creating a Job** (on page 689) for more details on how to create a new 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 the **Dependent On** field in the **Job Definition** dialog.



The image shows a '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:** Friday, February 1, 2013
- Start Time:** 3:00:00 AM
- Maximum Threads:** 4
- Inactive Wait Interval (seconds):** 30
- Scheduler:** Windows Only
- Dependent On:** (This field is highlighted with a red box in the original image. It contains a table with columns: Parent job, Fail, Look back (minutes), Maximum wait (minutes), and buttons for Add Parent Job and 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. 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)
- Success Command:** (Empty text box)
- Failure Command:** (Empty text box)
- Buttons:** OK, Cancel

Click on the **Add Parent Job** button.

Job Definition

Job Name:Enterprise Reporting Daily Refresh

Description:Full Daily Refresh of all Presentation Tables for the Enterprise Reporting Layer

Frequency:Custom

Start Date:Friday, February 1, 2013

Start Time:3:00:00 AM

Maximum Threads:4Inactive Wait Interval (seconds):30

Scheduler:Windows Only

Dependent On:Parent jobFailLook back (minutes)Maximum wait (minutes)

Logs Retained:0This 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:

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

Add Parent JobRemove Parent

OKCancel

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**.



The dialog box is titled "Add a Parent Job" and has a red close button in the top right corner. It contains two paragraphs of instructional text. The first paragraph explains that the user should select a parent job and specify a look-back time. The second paragraph explains that the user should specify a maximum wait time and that a failure will occur if the parent job does not complete in time. Below the text, there is a "Parent Job:" label followed by a drop-down menu showing "Shared Dimensions Daily Refresh". Below that is a label for "Maximum Time to Look Back for the Parent Job Completion (minutes):" followed by a text box containing "60". Then, another label for "Maximum Time to Wait for the Parent Job to Complete (minutes):" followed by a text box containing "20". At the bottom left, there is a checked checkbox with the label "Fail if Parent Job Does Not Complete Successfully in the Required Time:". At the bottom right, there are two buttons: "Add" and "Cancel".

Add a Parent Job

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

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:

Enterprise Reporting Daily Refresh

Description:

Full Daily Refresh of all Presentation Tables for the Enterprise Reporting Layer

Frequency:

Custom

Start Date:

Friday, February 1, 2013

Start Time:

3:00:00 AM

Maximum Threads:

4

Inactive Wait Interval (seconds):

30

Scheduler:

Windows Only

Dependent On:

Parent job	Fail	Look back (minutes)	Maximum wait (minutes)	
Shared Dimensions Daily Refresh	Y	60	20	

Add Parent Job

Remove Parent

Logs Retained:

0

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

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 699)

The **Edit Tasks** option enables you to edit the tasks of a job. See ***Editing Tasks*** (see "***Editing Tasks in a Job***" on page 703)

The **Edit Dependencies** option enables you to edit the task dependencies of a job. See ***Editing Dependencies*** (on page 711)

The **Insert Copy of Job** option enables you to insert a copy of a job. See ***Inserting a Copy of a Job*** (on page 716)

The **Delete Job** option enables you to delete a job. See ***Deleting a Job*** (on page 718)

The **Multiple Log Delete** option enables you to delete multiple logs of a job. See ***Deleting Job Logs*** (on page 719)

The **Start the Job** option enables you to start a job. See ***Starting a Job*** (on page 721)

The **Halt the Job** option enables you to halt a job. See ***Halting a Job*** (on page 722)


The **Abort the Job** option enables you to abort a job. See ***Aborting a Job*** (on page 723)

The **Restart the Job** option enables you to restart a job. See ***Restarting a Job*** (on page 724)


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.

The 'Job Definition' dialog box is shown with the following fields and options:

- Job Name:** Job_6709
- Description:** Job number 6709
- Frequency:** Hold
- Start Date:** Thursday, January 31, 2013
- Start Time:** 2:41:00 PM
- Maximum Threads:** 1
- Scheduler:** Windows Only
- Dependent On:** Parent job, Fail, Look back (minutes), Maximum wait (minutes). Includes 'Add Parent Job' and 'Remove Parent' buttons.
- Custom Settings:**
 - Interval Between Jobs: (minutes)
 - Start At or After HHMM (e.g. 0800): 0000
 - Do Not Start After HHMM (e.g. 1700): 0000
 - Active Days: Mon, Tue, Wed, Thu, Fri, Sat, Sun (all unchecked)
- Logs Retained:** 0. Text: '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)'
- Success Command:** (empty text area)
- Failure Command:** (empty text area)
- Buttons:** OK, Cancel

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. WARNING: on some UNIX systems long job names can cause jobs to be canceled (see Knowledge Base article 67), so where possible keep the name short.
Description	A description of the job
Frequency	When the job runs. The options available in the drop-down list box are: ☐ 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	☐ 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	Certain types of jobs will only run in a specific environment. For example ODBC based loads can only be handled by the Windows' scheduler. It is possible to have multiple schedulers running. Select the desired scheduler from this drop-down. The valid options are: UNIX Preferred, UNIX Only, 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 Job Dependency example in Scheduling a Job (on page 679).
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
OK command and Failure command	These are either UNIX or 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. Notes: 1. The Data Warehouse Designer 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.

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 in the left pane to add it to the task list in the right 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 or index.
Create	Create table, view or index.
Truncate	Delete all rows from the table.
Initial Build	Drop All Indexes then Custom then Build All Indexes .
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 script.
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.
DB2: Analyze	Performs a custom analyze if defined in: ☐ Table Properties/Statistics/Analyze else for tables performs: ☐ RUNSTATS ON TABLE tablename WITH DISTRIBUTION AND DETAILED INDEXES ALL otherwise (for indexes): ☐ RUNSTATS ON TABLE indname FOR INDEXES
DB2: Quick Analyze	Performs a custom analyze if defined in: ☐ Table Properties/Statistics/Quick Analyze else for tables performs: ☐ RUNSTATS ON TABLE tablename otherwise (for indexes): ☐ RUNSTATS ON TABLE indname FOR INDEXES
DB2: Stats	Performs a custom stats if defined in: ☐ Table Properties/Statistics/Stats else for tables performs: ☐ RUNSTATS ON TABLE tablename WITH DISTRIBUTION AND DETAILED INDEXES ALL otherwise (for indexes): ☐ RUNSTATS ON TABLE indname FOR INDEXES

DB2: Quick Stats	Performs a custom analyze if defined in: ☐ Table Properties/Statistics/Quick Stats else for tables performs: ☐ RUNSTATS ON TABLE tablename otherwise (for indexes): ☐ RUNSTATS ON TABLE indname FOR INDEXES
Oracle: Analyze	Performs a custom analyze if defined in: ☐ Tools-Options-Statistics-Table/Index Analyze Full otherwise performs: ☐ ANALYZE TABLE / INDEX name COMPUTE STATISTICS
Oracle: Quick Analyze	Performs a custom analyze if defined in: ☐ Tools-Options-Statistics-Table/Index Analyze Quick otherwise performs: ☐ ANALYZE TABLE / INDEX name ESTIMATE STATISTICS SAMPLE 3 PERCENT
Oracle: Stats	Performs a custom stats if defined in: ☐ Tools-Options-Statistics-Table/Index Stats Full otherwise performs: ☐ DBMS_STATS.GATHER_ TABLE / INDEX _STATS() using the ownname and tablename / indname parameters with cascade.
Oracle: Quick Stats	Performs a custom analyze if defined in: ☐ Tools-Options-Statistics-Table/Index Stats Quick otherwise performs: ☐ DBMS_STATS.GATHER_ TABLE / INDEX _STATS() using the ownname and tablename/indname parameters with cascade and estimate_percent = 3.
SQL Server: Analyze	Performs: ☐ UPDATE STATISTICS name WITH FULLSCAN
SQL Server: Quick Analyze	Performs: ☐ UPDATE STATISTICS name WITH SAMPLE 3 PERCENT
SQL Server: Stats	Same as Analyze.
SQL Server: Quick Stats	Same as Quick Analyze.

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.
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 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** will be 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 alphanumerics, spaces and the underscore are supported in the name. WARNING: on some UNIX systems long job names can cause jobs to be canceled (see Knowledge Base article 67), so where possible keep the name short.
Description	A description of the job

Field	Description
Frequency	When the job runs. The options available in the drop-down list box are: ☐ 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	☐ 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	Certain types of job will only run in a specific environment. For example ODBC based loads can only be handled by the Windows scheduler. It is possible to have multiple schedulers running. Select the desired scheduler from this drop-down. The valid options are: UNIX Preferred, UNIX Only, 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 Job Dependency example in <i>Scheduling a Job</i> (on page 679).
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.
OK command and Failure command	These are either UNIX or 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. Notes: 1. The Data Warehouse Designer 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.

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 in the left pane to add it to the task list in the right 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 or index.
Create	Create table, view or index.
Truncate	Delete all rows from the table.
Initial Build	Drop All Indexes then Custom then Build All Indexes .
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.
DB2: Analyze	Performs a custom analyze if defined in: ☐ Table Properties/Statistics/Analyze else for tables performs: ☐ RUNSTATS ON TABLE tablename WITH DISTRIBUTION AND DETAILED INDEXES ALL otherwise (for indexes): ☐ RUNSTATS ON TABLE indname FOR INDEXES
DB2: Quick Analyze	Performs a custom analyze if defined in: ☐ Table Properties/Statistics/Quick Analyze else for tables performs: ☐ RUNSTATS ON TABLE tablename otherwise (for indexes): ☐ RUNSTATS ON TABLE indname FOR INDEXES
DB2: Stats	Performs a custom stats if defined in: ☐ Table Properties/Statistics/Stats else for tables performs: ☐ RUNSTATS ON TABLE tablename WITH DISTRIBUTION AND DETAILED INDEXES ALL otherwise (for indexes): ☐ RUNSTATS ON TABLE indname FOR INDEXES

DB2: Quick Stats	Performs a custom analyze if defined in: ☐ Table Properties/Statistics/Quick Stats else for tables performs: ☐ RUNSTATS ON TABLE tablename otherwise (for indexes): ☐ RUNSTATS ON TABLE indname FOR INDEXES
Oracle: Analyze	Performs a custom analyze if defined in: ☐ Tools-Options-Statistics-Table/Index Analyze Full otherwise performs: ☐ ANALYZE TABLE / INDEX name COMPUTE STATISTICS
Oracle: Quick Analyze	Performs a custom analyze if defined in: ☐ Tools-Options-Statistics-Table/Index Analyze Quick otherwise performs: ☐ ANALYZE TABLE / INDEX name ESTIMATE STATISTICS SAMPLE 3 PERCENT
Oracle: Stats	Performs a custom stats if defined in: ☐ Tools-Options-Statistics-Table/Index Stats Full otherwise performs: ☐ DBMS_STATS.GATHER_ TABLE / INDEX _STATS() using the ownname and tablename / indname parameters with cascade.
Oracle: Quick Stats	Performs a custom analyze if defined in: ☐ Tools-Options-Statistics-Table/Index Stats Quick otherwise performs: ☐ DBMS_STATS.GATHER_ TABLE / INDEX _STATS() using the ownname and tablename/indname parameters with cascade and estimate_percent = 3.
SQL Server: Analyze	Performs: ☐ UPDATE STATISTICS name WITH FULLSCAN
SQL Server: Quick Analyze	Performs: ☐ UPDATE STATISTICS name WITH SAMPLE 3 PERCENT
SQL Server: Stats	Same as Analyze.
SQL Server: Quick Stats	Same as Quick Analyze.

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.


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 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_customer_sales2	Process	-->	stage_customer_ssis	Process
load_customer_sales2	Process	-->	stage_customer_xml	Process
load_customer_sales2	Process	-->	stage_product_sales	Process
load_customer_sales2	Process	-->	stage_product_sales2	Process
load_customer_ssis	Process	-->	stage_customer_ssis	Process
load_customer_ssis	Process	-->	stage_customer_xml	Process
load_customer_ssis	Process	-->	stage_product_sales	Process
load_customer_ssis	Process	-->	stage_product_sales2	Process
load_customer_xml	Process	-->	stage_customer_ssis	Process
load_customer_xml	Process	-->	stage_customer_xml	Process
load_customer_xml	Process	-->	stage_product_sales	Process
load_customer_xml	Process	-->	stage_product_sales2	Process
load_product2	Process	-->	stage_customer_ssis	Process
load_product2	Process	-->	stage_customer_xml	Process
load_product2	Process	-->	stage_product_sales	Process
load_product2	Process	-->	stage_product_sales2	Process
load_product_group	Process	-->	stage_customer_ssis	Process
load_product_group	Process	-->	stage_customer_xml	Process
load_product_group	Process	-->	stage_product_sales	Process
load_product_group	Process	-->	stage_product_sales2	Process
load_product_group_sales2	Process	-->	stage_customer_ssis	Process
load_product_group_sales2	Process	-->	stage_customer_xml	Process
load_product_group_sales2	Process	-->	stage_product_sales	Process
load_product_group_sales2	Process	-->	stage_product_sales2	Process
load_product_line	Process	-->	stage_customer_ssis	Process
load_product_line	Process	-->	stage_customer_xml	Process
load_product_line	Process	-->	stage_product_sales	Process
load_product_line	Process	-->	stage_product_sales2	Process
load_product_line_sales2	Process	-->	stage_customer_ssis	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.


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 displays the SQL Server Enterprise Manager interface, specifically the Scheduler tab. The 'Jobs' folder is expanded, showing a list of jobs. The jobs are as follows:


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

The interface includes a menu bar (File, View, Auto, Monitor, Logs, Jobs, Scheduler, Tools, Window, Help) and a toolbar with various icons. The bottom status bar shows 'Ready', 'Database time at last action was 08:34:36', 'Development', 'UserId: dbo', 'Browse: DataWarehouse', 'CAP', 'NUM', 'SCRL', and '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 will be displayed. Select or enter the appropriate options to delete the range of job logs required.



The dialog box is titled "Delete Multiple Job Logs" and contains the following fields and options:

- Job Name:** A text box containing "Daily Update". To the right, a note states: "If no name is specified then all logs are deleted. SQL Wildcards (%) may be used."
- Job State:** A dropdown menu set to "All". To the right, a note states: "Select a specific job state, or All to delete all logs"
- Job Owner:** A dropdown menu set to "All". To the right, a note states: "Select a specific job owner, or All to delete all logs."
- Days Prior:** A text box containing "14". To the right, a note states: "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 checked checkbox.
- Buttons:** "Delete" and "Cancel" buttons at the bottom right.

A warning message will be displayed. Click **Yes** to delete.




The warning dialog box contains the following text and buttons:

- Icon:** A blue question mark icon.
- Text:** "You are about to delete and archive all job logs for jobs named Process fact_forecast 6718 more than 14 days old"
- Question:** "Are you sure you wish to delete these logs?"
- Buttons:** "Yes", "No", and "Cancel" buttons at the bottom.

Starting a Job

To start a job


A job can be started by right-clicking on a job in the scheduler window and choosing **Start the Job**.



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 exists then only as a log of what occurred and is no longer regarded as a job.

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.

Jobs												
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	

Tasks												
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 (WslWarehouse) Userid: dbo Browse: DataWarehouse CAP NUM SCRL INS

To rerun a task, right-click on the completed task and select **Change to On Hold**

Tasks												
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							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 (WslWarehouse) 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.

Tasks											
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.

Tasks											
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.

Jobs

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

Tasks

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 (WslWarehouse) | 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.

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.

Application
Objects to Add/Replace
Objects to Delete

Output Directory: C:\Training\QAD\Module18\ Browse...

Application Identifier: Shared_Dimen Application Version: 130228105414

Application Name: Application files created from Job Shared Dimensions Daily Refresh

Description: These application files contain the objects included in, and the job Shared Dimensions Daily Refresh

You can select or enter a previous application file as a starting point for this application.

Previous Application: Browse...

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):

OK Cancel Help

3 Edit the objects to add or replace as required.

Build Deployment Application

Select Objects to Add or Replace in the Destination Repository.

Application
Objects to Add/Replace
Objects to Delete

Available:

- All Objects
- Orders
- Shared Dimensions
- Customers
- Jobs
- Parameters

Refresh List


Selected:

- load_customer_sales2
- load_customer_ssis
- load_customer_xml
- load_product2
- load_product_group
- load_product_group_sales2
- load_product_line
- load_product_line_sales2
- load_product_sales2
- load_product_sub_group
- load_product_sub_group_sales2
- stage_customer_ssis
- stage_customer_xml
- stage_product_sales
- stage_product_sales2
- stage_customer
- stage_product
- dim_currency_rate
- dim_date
- dim_product
- dim_sales_kpi
- Shared Dimensions Daily Refresh

OK Cancel Help

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**.



Monitoring the Database and Jobs

Note: Oracle on UNIX and Linux only.


QAD Data Warehouse Designer provides a methodology for monitoring both the database and selected jobs. In the event that a job fails or the database is not available a notification message can be sent to an administrator. The monitoring process requires the establishment of an external process. This process must run on a UNIX platform. It does not have to run on the data warehouse database platform, and in fact it is often preferable that it runs on a separate system. The monitor requires network access to the database. Refer to the Setup/Administrator guide for information on how to establish the external monitor process. This external monitor process must be running in order to utilize the features discussed in the following sections.

The following sections cover both database monitoring and job monitoring. The normal practice would be to setup database monitoring to cover those times when critical jobs are running. The individual monitoring of critical jobs is then defined.

Note: Although it is possible to setup job monitoring without database monitoring this is not recommended. If database monitoring is not in place and the database fails, then no job monitoring will occur and no notifications regarding the jobs will be issued. Therefore, database monitoring is required to provide a comprehensive job monitoring regime.

Database Monitoring

Database monitoring is defined by selecting the **Monitor/Database monitoring** menu option from the scheduler window, or from the stand-alone scheduler maintenance program. When this menu option is selected a check is made to ensure that the external monitor process is active. If the external process has not reported in the last 24 hours the following dialog will appear:



In such a case the external monitoring process on the UNIX platform needs to be checked or started. Refer to the **Setup/Administrator** guide for instructions on how to maintain the external monitoring process.

The following example dialog will then appear:

A dialog box titled "Set monitoring parameters for Database and Scheduler availability". It contains a table for scheduling parameters and a section for script names and parameters.

	HHMM		Scheduler			HHMM		Scheduler			HHMM		Scheduler		
	From	Until	Database	Windows	Unix	From	Until	Database	Windows	Unix	From	Until	Database	Windows	Unix
Monday:	2100	2400	☒	☐	☐	0	900	☒	☐	☐	0	0	☐	☐	☐
Tuesday:	2100	2400	☒	☐	☐	0	900	☒	☐	☐	0	0	☐	☐	☐
Wednesday:	2100	2400	☒	☐	☐	0	900	☒	☐	☐	0	0	☐	☐	☐
Thursday:	2100	2400	☒	☐	☐	0	900	☒	☐	☐	0	0	☐	☐	☐
Friday:	0	0	☐	☐	☐	0	900	☒	☐	☐	0	0	☐	☐	☐
Saturday:	0	0	☐	☐	☐	0	0	☐	☐	☐	0	0	☐	☐	☐
Sunday:	2100	2400	☒	☐	☐	0	0	☐	☐	☐	0	0	☐	☐	☐

	Script name	Period 1 Parameter	Period 2 Parameter	Period 3 Parameter
Database down:	monitor_db_mail	"database not available" dssadm	"database not available" dssadm	
Windows scheduler down:				
Unix scheduler down:				

Database Monitoring help ?

Save Exit


In this example the database is being monitored between the hours of 9:00pm and 9:00am to cover processing for business days monday through friday. If the database is not available during these times the host script called 'monitor_db_mail' will be called and it will be passed two positional parameters. The first parameter being "database not available" and the second being "dssadm". The parameters required are specific to the script being invoked.

In this example we have defined two monitoring periods. The first is for 9:00pm until midnight and the second is from midnight to 9:00am. As a result there are two sets of parameters defined for a database notification, one for the first period and one for the second period. In this case both are the same.

Job Monitoring

The external monitoring process regularly monitors jobs running in the scheduler and provides notifications when certain conditions occur. By default the interval between the monitoring of jobs is 15 minutes. This interval can be overridden during the setup of the external job monitor. See the **Setup/Administration** guide.

Job monitoring is defined by selecting the **Monitor/Job monitoring** menu option from the scheduler window, or from the stand-alone scheduler maintenance program. When this menu option is selected a check is made to ensure that the external monitor process is active. If the external process has not reported in the last 24 hours the following dialog will appear:



In such a case the external monitoring process on the UNIX platform needs to be checked or started. Refer to the Setup/Administrator guide for instructions on how to maintain the external monitoring process.

The following example dialog will then appear:

Define job monitoring parameters

Job Name: Enterprise Reporting Daily Refresh ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☐ Sat ☐ Sun ☐ 1st Month ☐ Last day of Month

Nominal start time (HHMM): 1515 Earliest possible start before nominal start time. Number of minutes before [e.g. 120 is 2 hours before nominal start time]. (minutes): 0

Periodicity of job (HHMM): 2400

	Action script	Parameter to script (or command)	Active
Notification on failure to start or failure to finish within specified time frame			
Must start within (minutes): 60	monitor_db_mail	"Enterprise Reporting Daily Refresh" "Failed to start on time" dssadm	☒
Must finish in (minutes): 240	monitor_db_mail	"Enterprise Reporting Daily Refresh" "Failed to finish on time" dssadm	☒
Notification on warning or error messages occurring while a job is running			
Notify on RUN warning count > 0			☐
Notify on RUN error count > 0	monitor_db_mail	"Enterprise Reporting Daily Refresh" "Errors during run" dssadm	☒
Notification on warning or error messages at the completion of a job			
Notify on FINISH warning count > 0			☐
Notify on FINISH error count > 0	monitor_db_mail	"Enterprise Reporting Daily Refresh" "Failed with errors" dssadm	☒
Notify on FINISH Success			☐
Notification on failure to reach a specified number of actions by a checkpoint time			
Checkpoint 1 by (minutes): 0			☐
Completed task count >= 0	OK + Information message count >= 0		
Checkpoint 2 by (minutes): 0			☐
Completed task count >= 0	OK + Information message count >= 0		
Checkpoint 3 by (minutes): 0			☐
Completed task count >= 0	OK + Information message count >= 0		

Job monitoring help 

Save Exit

In this example the job called 'Sales Daily' is being monitored. The days Monday through Friday are checked, so monitoring will occur on these days. The nominal start time is 04:00am. Notifications will occur if the following conditions occur.

- 1 The job fails to start within 60 minutes of this nominal start time.
- 2 The job fails to finish within 240 minutes (4 hours) of the nominal start time.
- 3 Errors occur while the job is running
- 4 The job fails with errors.

In all conditions that could cause a notification we will invoke the host script 'monitor_job_mail'. This script requires three parameters. The first parameter is the job name, the second a message, and the third the user or users to mail.

Monitor Days

When defining monitoring for a job one or more days need to be selected. As well as the normal week days there is also the first and last day of the month. Using these two special days it is possible to monitor a job that is critical at the end or start of a month. If Monday through Sunday are checked then it would not be necessary to set the special days. The monitoring day is the day on which the job starts (or could start).

Note: The monitoring day is the day of the absolute minimum start time of a job. For example if we wish to monitor a job that starts at 00:01, and which has an earliest possible start of 5 minutes earlier (i.e. 23:56) then we need to set monitoring for the day of the earliest possible start. Therefore to monitor the Monday morning run we would need to have the Sunday checkbox set.

Nominal Start Time

The nominal start time is the time at which the job will normally start. This field must be populated for all jobs being monitored even if not used. Monitoring is optimized for daily jobs, so this field is normally set to the start time of the job.

Earliest Possible Start

The earliest possible start field allows for the setting of a possible start time before the nominal start time. In most situations this field will be set to 0. When setting to a value other than zero the impact of the potential start time on the days being monitored needs to be considered. See the note under the Monitor Days section above. The nominal start time less the earliest possible start time provides an **Absolute Minimum Start Time** which is used by most of the notification conditions, as well as the monitor days.

Periodicity of Job

The periodicity defines the interval between iterations of the job. For a daily job this would be 24 hours and 0 minutes or 2400. This periodicity is used by the monitoring software to ascertain if a job being examined is the current iteration of the job or a previous iteration.

Notification Order

Notifications are processed in the following order. If a notification is sent for a job then no further notification checks are made for that job.

- 1 Finish Error
- 2 Finish Warning
- 3 Finish Success
- 4 Must Finish In
- 5 Run Error
- 6 Run Warning
- 7 Checkpoint 3
- 8 Checkpoint 2
- 9 Checkpoint 1
- 10 Must Start Within

To fully cover all possible events it is usually necessary to set multiple notifications.

Must Start Within (Notification)

The **Must Start Within** notification allows notification when a job fails to start within a specified number of minutes of the absolute minimum start time. In the example dialog above the job must start within 60 minutes of the start time.

If this criteria is not met then the monitoring software will look at the **Action Script** and **Parameter to Script** fields to ascertain how to process the notification. If an action script has been defined then this script will be executed in the UNIX environment and will be passed the defined parameters. If no script is defined then the **Parameter to Script** field will be executed in the UNIX environment.

An example notification script is included as a template called 'unix_mon_job_mail'.

Must Finish In (Notification)

The **Must Finish In** notification allows notification when a job fails to finish within a specified number of minutes of the absolute minimum start time. In the example dialog above the job must finish within 4 hours of the start time. See the notes under **Must Start Within** for the action taken if the criteria is met.

Run Warning Count (Notification)

The **Run Warning Count** notification allows notification when a job exceeds a specified number of warning messages whilst running. This notification *will only occur when a job is running*. If the job has failed to start, has finished or failed before the external monitor checks the number of warnings this notification will be ignored. See the notes under **Must Start Within** for the action taken if the criteria is met.

Run Error Count (Notification)

The **Run Error Count** notification allows notification when a job exceeds a specified number of error messages whilst running. This notification *will only occur when a job is running*. If the job has failed to start, has finished or failed before the external monitor checks the number of errors this notification will be ignored. See the notes under **Must Start Within** for the action taken if the criteria is met.

Finish Warning Count (Notification)

The **Finish Warning Count** notification allows notification when a job has exceeded a specified number of warning messages and has completed or failed. This notification *will only occur when a job has completed or failed*. If the job has failed to start or is still running when the external monitor checks the number of warnings this notification will be ignored. See the notes under **Must Start Within** for the action taken if the criteria is met.

Finish Error Count (Notification)

The **Finish Error Count** notification allows notification when a job has exceeded a specified number of error messages and has failed. This notification *will only occur when a job has failed*. If the job is still running when the external monitor checks the number of warnings, this notification will be ignored. If the job has been restarted and gone on to complete normally, this notification will be ignored even when errors have occurred. See the notes under **Must Start Within** for the action taken if the criteria is met.

Finish Success (Notification)

The **Finish Success** notification allows notification when a job completes. This notification can be used for example to mail a log of the completed job to a specified user. See the notes under **Must Start Within** for the action taken if the criteria is met.

Checkpoint 1 (Notification)

The **Checkpoint 1** notification allows notification when a job fails to achieve either a specified number of task completions, or a specified number of information/OK messages within a specified elapsed time. The elapsed time is the time from the absolute minimum start time. This notification *will only occur when the job is running*. If the job has failed to start, completed or failed when the external monitor checks the checkpoint, this notification will be ignored. See the notes under **Must Start Within** for the action taken if the criteria is not met.

Checkpoint 2 (Notification)

The **Checkpoint 2** notification allows notification when a job fails to achieve either a specified number of task completions, or a specified number of information/OK messages within a specified elapsed time. The elapsed time is the time from the absolute minimum start time. This notification *will only occur when the job is running*. If the job has failed to start, completed or failed when the external monitor checks the checkpoint, this notification will be ignored. See the notes under **Must Start Within** for the action taken if the criteria is not met.

Checkpoint 3 (Notification)

The **Checkpoint 3** notification allows notification when a job fails to achieve either a specified number of task completions, or a specified number of information/OK messages within a specified elapsed time. The elapsed time is the time from the absolute minimum start time. This notification *will only occur when the job is running*. If the job has failed to start, completed or failed when the external monitor checks the checkpoint, this notification will be ignored. See the notes under **Must Start Within** for the action taken if the criteria is not met.

Stand Alone Scheduler Maintenance

QAD Data Warehouse Designer includes a stand alone scheduler maintenance screen. This screen provides all the scheduler control functionality found in the main QAD Data Warehouse Designer 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. There is an additional entry required being **Meta database**. This is the schema of the metadata repository that you wish to connect to in terms of scheduler maintenance.

DATABASE

Data Source: WslWarehouse

Database Login ID: dbo

Password:

METADATA REPOSITORY

RED User Name: John Smith

Help Cancel Connect

For **SQL Server** the Meta database must be set to **dbo**.

For **Oracle** data warehouses the 'Meta database' owner must have made the following grants to the specified user before the scheduler maintenance utility can be used.

Scheduler maintenance grants

Repository access grants:

- ☐ grant select on ws_meta to dsssched;
- ☐ grant select on ws_meta_tables to dsssched;
- ☐ grant select on ws_obj_type to dsssched;
- ☐ grant select,insert on ws_user_adm to dsssched;
- ☐ grant select,insert,update on ws_user_adm to dsssched;
- ☐ grant select,alter on ws_wrk_job_dependency_seq to dsssched;
- ☐ grant select on ws_table_attributes to dsssched;

Object access (job create) grants:

- ☐ grant select on ws_obj_object to dsssched;
- ☐ grant select on ws_obj_pro_map to dsssched;
- ☐ grant select on ws_obj_project to dsssched;
- ☐ grant select on ws_obj_group to dsssched;
- ☐ grant select on ws_pro_gro_map to dsssched;

Scheduler status grants:

- ☐ grant select on ws_wrk_audit_log to dsssched;

Scheduler status, and job deletion grants:

- ☐ grant select,delete on ws_wrk_error_log to dsssched;

Scheduler status, poll grants:

- ☐ grant select,update on ws_wrk_scheduler to dsssched;

Job creation grants:

- ☐ grant select,insert,update,delete on ws_wrk_dependency to dsssched;
- ☐ grant select,insert,update,delete on ws_wrk_job_ctrl to dsssched;
- ☐ grant select,alter on ws_job_seq to dsssched;
- ☐ grant select,alter on ws_task_seq to dsssched;

Job maintenance grants:

- ☐ grant select,insert,delete on ws_wrk_job_log to dsssched;
- ☐ grant select,update,delete on ws_wrk_job_run to dsssched;
- ☐ grant select,insert,update,delete on ws_wrk_dependency to dsssched;
- ☐ grant select,insert,update,delete on ws_wrk_job_dependency to dsssched;
- ☐ grant select,delete on ws_wrk_job_thread to dsssched;

Task maintenance grants:

- ☐ grant select,insert,update,delete on ws_wrk_task_ctrl to dsssched;
- ☐ grant select,update,delete on ws_wrk_task_run to dsssched;
- ☐ grant select,insert,delete on ws_wrk_task_log to dsssched;
- ☐ grant select,insert,update on dss_parameter to dsssched;

Right-click **Used by** option in parameters listing grants:

- ☐ grant select on ws_pro_header to dsssched;
- ☐ grant select on ws_pro_line to dsssched;

- ☐ grant select on ws_scr_header to dsssched;
- ☐ grant select on ws_scr_line to dsssched;
- ☐ grant select on ws_load_tab to dsssched;
- ☐ grant select on ws_load_col to dsssched;
- ☐ grant select on ws_stage_tab to dsssched;
- ☐ grant select on ws_stage_col to dsssched;
- ☐ grant select on ws_dim_tab to dsssched;
- ☐ grant select on ws_dim_col to dsssched;
- ☐ grant select on ws_fact_tab to dsssched;
- ☐ grant select on ws_fact_col to dsssched;
- ☐ grant select on ws_agg_tab to dsssched;
- ☐ grant select on ws_agg_col to dsssched;

A sample script to grant these privileges is shipped with QAD Data Warehouse Designer. This script is called 'grant_sched_access.sql' and can be found in the QAD program directory.

The scheduler maintenance utility does not require a QAD license key. The QAD Data Warehouse Designer software can be installed onto a PC, and this utility utilized without having to use the QAD **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_interval/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.



Scheduling a Data Warehouse Designer Host Script from 3D

Create the Data Warehouse Designer host script in 3D. Use Cntrl-C to copy the script.

Memory: 512 MB

RED host script: ☒ Windows

Command:


```
@ECHO OFF
SETLOCAL ENABLEDELAYEDEXPANSION
SETLOCAL ENABLEEXTENSIONS
REM
REM =====
REM Purpose:      WhereScape RED Windows Host Script for scheduling
REM              WhereScape 3D processing.
REM Generated by: WhereScape 3D
REM =====
REM
SET LOGFILE=C:\temp\3Ddisc.aud
SET ERRFILE=C:\temp\3Ddisc.err
SET PATH3D=C:\Program Files (x86)\WhereScape\WhereScape 3D\WhereScape-3D-HEAD-bundle.jar
SET REPO=Sales Analysis Repository
```

Cron expression: Calculate


Schtasks expression: Calculate

OK

In Data Warehouse Designer, create a host script by right-clicking on the **Host Script** object group and choosing **New Object**.



Enter a name for the host script 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: Host Script

Object Name: 3d

ADD Cancel

Set the default connection to Windows and click **OK**.



Host Script 3d

Properties Notes

Name: 3d Type: Windows script

Purpose:

Owner: dbo ☐ Delete Lock

Last Update By:

Default Connect: Windows

Checkout

Checked Out By:


Checkout Reason:
or Last Update New Script

Timestamps


Created: Last Update: Compiled:

OK Cancel Help

Double-click on the 3d host script in the left pane.



Now paste the script from 3D and click **Save**; then close the host script window.



The screenshot shows a script editor window with a menu bar (File, Edit, View, Format, Execute, Tools, Window, Help) and a toolbar. The '3d' tab is active. The script content is as follows:

```
@ECHO OFF
SETLOCAL ENABLEDELAYEDEXPANSION
SETLOCAL ENABLEEXTENSIONS
REM
REM =====
REM Purpose:          WhereScape RED Windows Host Script for scheduling
REM                  WhereScape 3D processing.
REM Generated by:     WhereScape 3D
REM =====
REM


SET "LOGFILE=C:\temp\3Ddisc.aud"
SET "ERRFILE=C:\temp\3Ddisc.err"
SET "PATH3D=C:\Program Files (x86)\WhereScape\WhereScape 3D\WhereScape-3D-HEAD-bundle.jar"
SET "REPO=Sales Analysis Repository"
SET "CATEGORY=Source System"
SET "MODEL=Tutorial Source"
SET "VERSION=Version 1"
SET "LOCATION=C:\Users\User\WhereScape\3D\Documentation\001.docx"
SET "TITLE=Source System Design"

java -Xmx512m -XX:MaxPermSize=256m -splash: -jar "%PATH3D%" document "%REPO%" "%CATEGORY%" "%MOD

ECHO 1
ECHO WhereScape 3D process has completed
TYPE "%LOGFILE%"
TYPE "%ERRFILE%" 1>&2
```


The status bar at the bottom shows: Ready | Development (Red_1) | UserId: dbo | Browse: No source system | CAP | NUM | SCRL | INS

Go to **Control Panel > Administrative Tools > Services** and right-click on the relevant service name. Choose **Properties** and then the **Log On** tab. Change the scheduler to run as the current user. Click **Apply** and then **OK**.



Stop and start the scheduler and then close the **Services** window.

Click on **Scheduler Maintenance** in QAD Administrator; right-click on the scheduler and choose **Configure**.



Set the **Service logged on under User** field to the current user. Click **OK** and close QAD Administrator.

Windows Scheduler Definition

OK
Cancel

Unique Identifier (e.g. SID_user):	Red1
Scheduler Name (e.g. \WINnnnn):	\WIN0005
Service Display Name:	WhereScape Red1
Service logged on under User:	User
Work Directory:	C:\ProgramData\WhereScape\Work\
Log File:	C:\ProgramData\WhereScape\Scheduler\WslSched_Red1.log
Poll Interval (default 30) seconds:	30
Log Level (0=none, 9=debug)	2
Direct Path Load or ODBC	OdBC sql load
TNS Service (Oracle):	Red_1
ODBC Connection:	Red_1
Database User id:	dbc
Database Password:	xxx
MetaData Schema (DB2):	
MetaData Database (Teradata):	

In Data Warehouse Designer, click on the **Scheduler** tab and create a new job. Click **OK**.

Job Definition

Job Name:

3d doc generation

Description:

Job number 8129

Frequency:

Once and Hold

Start Date:

Thursday , December 5, 2013

Start Time:

11:32:00 AM

Maximum Threads:

1

Scheduler:

Windows Only

Dependent On:

Parent job

Fail

Look back (minutes)

Maximum wait (minutes)

Add Parent Job

Remove Parent

Logs Retained:

0

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

Select the host script to be run and click **OK**.



Run the job.

Chapter 22

Indexes

Indexes may exist on any table defined in QAD Data Warehouse Designer. By default QAD Data Warehouse Designer will auto-generate a number of indexes during the drag and drop process and when building procedures. 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.

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



The screenshot shows a window titled "Index list for dim_customer". It contains a table with the following data:

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


A context menu is open over the third row (dim_customer_idx_SC), displaying the following options:

- 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.

❑ SQL Server example:



□ Oracle example:


Index dim_date_idx_0

Properties

Index Description	Primary Artificial key	...
Active	True	▼
Business Key	False	▼
Artificial Key	True	▼
Unique	True	▼
Primary Key Constraint	True	
Index Type	B-Tree Index	▼
Index Columns	dim_date_key	...
Parallel	False	▼
Compute Statistics	False	▼
Storage		
Tablespace	DSS_INDEX	▼
Initial Extent (KB)	0	
Next Extent (KB)	0	
Minimum Extents	0	
Maximum Extents	0	
Percent Increase	0	
Percent Free	0	
Buffer Pool	(None)	▼
Compress	False	▼
Scheduler		
Drop Index before Table Update	False	▼
Index Rebuild Frequency	Never	▼

OK Cancel Help

❑ IBM DB2 example:



The fields in the **General** section are described below:

Field	Description
Index Name	Typically the table short name followed by: ☐ _idx_0 indicating primary key ☐ _idx_n where n = any number from 1 indicating dimensional keys ☐ _idx_x where x = any letter a thru z indicating business keys.
Index Description	☐ Free flow description of the index
Active	☐ When enabled means the index is in use. ☐ When disabled means the index is not managed by the scheduler.
Business Key	Denotes a business key.
Artificial Key	When enabled indicates that this is the surrogate (artificial) key generated by the system.
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.

Field	Description
Index Type	Database-specific type of index On SQL Server the options are: ☐ Nonclustered Index ☐ Clustered Index ☐ ColumnStore Index available for SQL2012 On Oracle the options are: ☐ B-Tree Index ☐ Bitmap Index ☐ Function-based Index
Index Columns	Shows the columns in the order they will be applied to the index. The order can be changed using the up/down buttons at the left.
Non-Key (Include) Columns	Extend a Nonclustered Index to include Non-Key columns in addition to the key columns.
Index Filter Predicate	A filtered index is an optimized Nonclustered index, which uses a filter predicate to
Enable Parallel Build Degree	Parallelize the creation /build of the index. For unsupported database versions (Pre SQL2005) leave disabled.
Additional Index Create Options	Database-specific-and-compliant options to include in the generated CREATE INDEX statement.
Compute Statistics	For Oracle. When enabled, Data Warehouse Designer automatically collects statistics during index creation and rebuild. NOTE: Deprecated since 10g (but supported for backwards compatibility) as now statistics are always collected during index creation and rebuild.

The fields in the **Storage** section are described below:

Field	Description
Filegroup/Tablespace	Default Filegroup/Tablespace of the index that determines the storage location it is created in. Select (Default) to use the default.
Initial Extent (KB)	For Oracle. Size of the first extent allocated to the index when it is created. The default value of 0 will instead use the value from the Tablespace in which the index is stored.
Next Extent (KB)	For Oracle. Size of the next extent allocated to the index, unless Percent Increase is specified.
Minimum Extents	For Oracle. Number of extents to allocate when the index is created. The default value of 0 will instead use the value from the Tablespace in which the index is stored.
Maximum Extents	For Oracle. Total number of extents that can be allocated to the index. The default value of 0 will instead use the value from the Tablespace in which the index is stored.
Percent Increase	For Oracle. Percent by which the third and subsequent grow over the preceding extent. The default value of 0 will instead use the value from the Tablespace in which the index is stored.

Field	Description
Percent Free	For Oracle. Percentage [0-99] of space in each data block of the index to reserve for future updates. The default value of 0 will instead use the value from the Tablespace in which the index is stored.
Buffer Pool	For Oracle. Default buffer pool used to cache the data blocks of the index.
Compress	Compress index to reduce disk and memory use. For unsupported database versions (Pre SQL2008) select (Not Defined).
Enable Fill Factor	Setting this will enable setting of the Fill Factor. For unsupported database versions (Pre SQL2005) leave disabled.
Sort In tempdb	Use the tempdb to store the intermediate sort results that are used to build the index, which may reduce the index build time if tempdb is on different disks to the index.
Enable Percent Free	For DB2. Setting this will enable setting of the Percent Free to something other than the Database defaults.
Percent Free	For DB2. The amount of free space to leave in each extent/page of the index.
Enable Level2 Percent Free	For DB2. Setting this will enable setting of the Level2 Percent Free to something other than the Database defaults.
Level2 Percent Free	For DB2. The amount of free space in each index level 2 page.
Enable Minimum Percent Used	For DB2. Setting this will enable setting of the Minimum Percent Used to something other than the Database defaults.
Minimum Percent Used	For DB2. The threshold of the minimum percentage of space used on an index leaf page.

The fields in the **Scheduler** section are described below:

Field	Description
Drop Index before Table Update	Drop the index before running the table update procedure and recreate the index after the update procedure has completed.
Index Rebuild Frequency	Optional frequency that the index is automatically rebuilt by the Scheduler.

Indexes are normally managed by the scheduler as part of the normal processing of a table.

Refer to **Ws_Maintain_Indexes** (on page 990). This function allows the control of index drop and creation from within a procedure. Typically this function is called when using partitioned tables.

Chapter 23

Documentation and Diagrams

QAD Data Warehouse Designer 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 metadata. 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 762) and **read** (see "**Reading the Documentation**" on page 766) the documentation.

Also included is a section on information that is available to assist in the connection of end user **query tools** (see "**Query Tool Access**" on page 791).

In This Chapter


Creating Documentation.....	762
Batch Documentation Creation.....	765
Reading the Documentation	766
Diagrams.....	767
Query Tool Access	791

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 will appear to allow the selection of a specific group or project. The default is all objects in the repository.



The dialog box is titled "Select a Project or Group or All Objects" and has a red close button in the top right corner. It contains the following text and controls:

To produce a full set of documentation select All Objects, or select a Group or Project to produce a subset of the documentation and select OK.

If you have already created the documentation and you just wish to refresh the objects in a Group or Project then select Objects Only, a Group or Project and OK.

NOTE: Any new objects will not appear in the index.

Objects Only ☐

Group:


Project:

Buttons: All Objects, OK, Cancel

A file dialog will appear next. Select a file path (directory) under which to save 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:



Documentation Creation Options

User and technical documentation will be created in HTML format in the destination directory. OK

To use a custom look and feel add your own MainStyle.css file into the destination directory. Cancel

Documentation Title (e.g. Data Warehouse):

Do you want to link in any custom HTML pages? ☐ Links

Do you want to include current file group usage?
It will take longer to create the documentation. ☒ Sizes

How would you like the columns sorted? ☒ Column Order ☐ Column Name ☐ Business Name

Do you want to include shadows on the diagram boxes? ☒ Shadow

Do you want to create impact analysis on load tables?
It will take longer to create the documentation. ☒ Impact

Do you want to replace the existing style sheet?
Do not tick this box if you utilize a custom style sheet. ☐ Replace Style Sheet

Do you wish to limit the complexity of the diagrams?
Select the maximum number of process steps to display in the source diagrams.

The size 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, thus should normally be used only for final documentation.

The sorted checkbox sorts the columns within the tables into alphabetical order. By default the columns are in the order in which they appear within the tables.

Creating a header

If you check the banner frame option then a banner (heading) will be 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 check 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



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 on the columns in the dimension and fact 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

QAD Data Warehouse Designer 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 Schema } { /N FullName } /D Directory { /G GroupName | ProjectName } /S NumHops
```

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	Sometimes *2	Metadata schema 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.

Notes: 1. User Name and Password are not required when using a trusted connection or operating system authentication.

2. Metadata Schema is required for DB2, but is not required for Oracle and SQL Server.

Example

The following example connects to a SQL Server repository called WslWarehouse, using a trusted connection into the C:\Temp\my_doco directory with 4 hops in diagrams:

```
med /BD /C "WslWarehouse" /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 **Display 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 770)
- ☐ The **Source Diagram** (on page 772)
- ☐ The **Joins Diagram** (on page 776)
- ☐ The **Links Diagram** (on page 777)
- ☐ The **Impact Diagram** (on page 778)
- ☐ The **Dependency Diagram** (on page 779)

- 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 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.




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 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 Properties are then displayed:

Job Definition

Job Name:

Process to fact_order_detail

Description:

Job created from source diagram for object fact_order_detail

Frequency:

Hold

Start Date:

Thursday , January 31, 2013

Start Time:

4:20:00 PM

Maximum Threads:

1

Scheduler:

Windows Only

Dependent On:

Parent job

Fail

Look back (minutes)

Maximum wait (minutes)

Add Parent Job

Remove Parent

Logs Retained:

0

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

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 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 shown 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.



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. This diagram relies on the source table and source column fields in each table type. If these fields do not reflect all of the base tables then the track back and track forward diagrams will not return the correct results.

An example of an **Impact** diagram in **Standard Diagram** format is displayed below.




Note: It is also possible to display the track forward diagram by right-clicking on a table and choosing **Impact > Track Forward Diagram**:

Execute Custom Procedure via Scheduler		Track Back Report
Documentation	▶	Track Forward Report
Projects	▶	Track Back Diagram
Impact	▶	Track Forward Diagram
		Dependent Jobs Report

Dependency Diagram


A job dependency diagram can be displayed for any job defined in the QAD Data Warehouse Designer scheduler. It shows the parent and child relationships between tasks within a job.table.

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:

[illegible]

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 QAD Data Warehouse Designer.

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 from QAD Data Warehouse Designer 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:

Process to fact_order_detail

Description:

Job created from source diagram for object fact_order_detail

Frequency:

Hold

Start Date:

Thursday , August 22, 2013

Start Time:

9:40:00 AM

Maximum Threads:

1

Scheduler:

Windows Only

Dependent On:

Parent job

Fail

Look back (minutes)

Maximum wait (minutes)

Add Parent Job

Remove Parent

Logs Retained:

0

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

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 click **OK**.




The 'Build Deployment Application' dialog box is shown. It has a sidebar on the left with 'Application' selected. The main area contains the following fields:

- Output Directory:** C:\ProgramData\QAD\Application\ (with a 'Browse...' button)
- Application Identifier:** fact_order_d
- Application Version:** 130822101354
- Application Name:** Application files created from a Source/Joins Diagram for object fact_order_detail
- Description:** These application files form a process to create and load the object fact_order_detail
- Previous Application:** (empty field with a 'Browse...' button)
- Pre Application Load SQL:** (empty text area)
- Post Application Load SQL:** (empty text area)

At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

3 A dialog will display, confirming the creation of the application files. Click **OK**.



The confirmation dialog box displays the following information:

- Icon:** Information icon (i)
- Text:** The following files have been created in:
C:\ProgramData\QAD\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
- Text:** These files can be moved to another location or system if required.
All files must be moved for the application to load.

An 'OK' button is located at the bottom right.


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 default for this checkbox is on.

Query Tool Access

Where the business definition has been set in the properties of either a fact or dimension column QAD Data Warehouse Designer will attempt to record this information in the database. For the Oracle database this business definition is stored as a column comment. Many end user query tools can make use of these Oracle column comments. They can also be examined via the Oracle database system views `all_col_comments` and `user_col_comments`. For the SQL Server database the business definition is also stored as a column comment via the `sp_addextendedproperty` procedure with a property name of **Comment**.

In addition a number of metadata views exist to make the task of connecting an end user query tool to the data warehouse simpler. Many end user tools can make use of external data to populate their end-user-layer (eul) definitions. Using these views end user query tools can utilize QAD Data Warehouse Designer's business metadata. The views supplied are:

Dimension tables and columns

The view **ws_admin_v_dim_col** provides a listing of all columns in all dimensions that have the 'Eul visible' flag set. Included in the information provided is the business definition and description of the column as defined within QAD Data Warehouse Designer.

Fact tables and columns

The view **ws_admin_v_fact_col** provides a listing of all columns in all fact tables that have the 'Eul visible' flag set. Included in the information provided is the business definition and description of the column as defined within QAD Data Warehouse Designer.

Fact - Dimension Joins

The view **ws_admin_v_fact_join** provides a listing of all fact tables and all the dimensions that each fact table joins to. Both the fact table dimension key column and the dimension key column are included in the view.

Chapter 24

Reports

QAD Data Warehouse Designer 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 Data Warehouse Designer screen.

The following sections describe the purpose, parameters and results for each report.

In This Chapter

Dimension-Fact Matrix	794
OLAP Dimension-Cube Matrix	795
Dimension Views of a Specified Dimension	796
Column Reports	797
Table Reports	807
Procedure Reports	814
Index Reports	817
Object Reports	819
Job Reports	826
Operational Reports	831

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 QAD Data Warehouse Designer object types are included in this report:

- ☐ Dimension Tables
- ☐ Fact Tables
- ☐ Aggregate Tables

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



The screenshot shows a report window titled "Report: Dimension-Fact Matrix". The report displays a crosstab where the rows represent dimensions and the columns represent fact and aggregate tables. The dimensions listed are dim_custo..., dim_date, dim_order..., dim_product, and dim_ship_d... The fact and aggregate tables listed are fact_budget, fact_forecast, fact_sales_anal..., fact_sales_detail, agg_sa_custo..., and agg_sa_product. The cells contain the value 1 if the dimension is used by the fact or aggregate table, and blank otherwise.

Dimensions	fact_budget	fact_forecast	fact_sales_anal...	fact_sales_detail	agg_sa_custo...	agg_sa_product
dim_custo...	1	1	1	1	1	
dim_date	1	1				
dim_order...				1		
dim_product	1	1	1	1		1
dim_ship_d...			1	1	1	1

Below the table, there is a "Reports" section with a status bar showing: "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 **Excel it** 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 QAD Data Warehouse Designer 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	sales_analysis (...)	sales_analysis (...)	
odim_order...	dim_order_date	No Relationship	
odim_produ...	dim_product	dim_product	
dim_ship_d...	dim_ship_date	No Relationship	
odim_budg...	No Relationship	Date dimension	
odim_custo...			

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 **Excel it** to send the results of this report to Microsoft Excel.

Dimension Views of a Specified Dimension

This report shows dimension views built on a specified dimension table.

Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Dimensions (Dimension View table type only)

Parameters

This report has one parameter:

- ☐ Dimension table name

Select a Dimension Table (Dimension Views Report)

Select a Dimension Table to list the views of this table

Dimension Table:

dim_date

OK

Cancel

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

Report: Views of a Specified Dimension

Dimension Name	Dimension View Names
dim_date	dim_order_date
dim_date	dim_ship_date

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 **Excel it** 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 user facing **table** objects columns in the metadata that don't have descriptions.

Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate Tables

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 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 **Excel it** to send the results of this report to Microsoft Excel.

All Column Transformations

This report shows all columns that have a column transformation on them and the details of the transformation.

Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Load Tables
- ☐ Stage Tables
- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate Tables
- ☐ Exports

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_product	description	upper(substring(descripti...	Description of the product. Normally unique.

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 **Excel it** to send the results of this report to Microsoft Excel.

Re-Usable Column Transformations

This report shows all reusable transformations as defined via **Tools/Reusable 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

Report: Re-Usable Column Transformations		
Template Name	Description	Model
TO_DATE	converts a trimmed string in YYYYMMDD format to a date	CASE WHEN column THEN expres
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 **Excel it** to send the results of this report to Microsoft Excel.

Column Track-Back

This report lists the origins of the selected objects.


Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Load Tables
- ☐ Stage Tables
- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate Tables
- ☐ Exports

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:

Report: Column Track Back			
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
fact_sales_analysis	tax	stage_sales_detail	tax
fact_sales_analysis	tax	load_order_line	tax
fact_sales_analysis	tax	order_line	tax

Reports

F Middle Pane: load_budget Columns | Development (WslWarehouse) | UserId: dbo | Browse: No source system | CAP | NUM | SCRL | INS

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:

Report: Column Track Back (Excluding intermediate steps)			
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_date	dim_date_key
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	dss_fact_table_key	dss_fact_table	dss_fact_table_key
fact_sales_analysis	dss_update_time		dss_update_time
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

Reports

F Middle Pane: load_budget Columns | Development (WslWarehouse) | UserId: dbo | Browse: No source system | CAP | NUM | SCRL | INS

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 **Excel it** 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 QAD Data Warehouse Designer object types are included in this report:

- ☐ Load Tables
- ☐ Stage Tables
- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate Tables
- ☐ Exports

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

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:

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:

Chapter 24 **Reports** - 805

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 **Excel it** to send the results of this report to Microsoft Excel.

Table Reports

There are five reports for analyzing Tables:

- ☐ Tables Without Comments
- ☐ Load Tables by Connection
- ☐ Export Objects by Connection
- ☐ Records that Failed a Dimension Join
- ☐ External source Table/Files

Tables without Comments

This report shows user facing **table** objects in the metadata that don't have descriptions.

Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate 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 missing comments with the following columns:

- ☐ Table name (*the name of the table*)
- ☐ Table type

Report Example



The screenshot shows a window titled "Report: Tables Without Comments". It contains a table with two columns: "Table Name" and "Table Type". The table lists several objects: "agg_sa_customer" (Aggregate), "agg_sa_product" (Aggregate), "dim_customer" (Dimension), "dim_order_date" (Dimension View), "dim_product" (Dimension), "dim_ship_date" (Dimension View), and "fact_budget" (Fact Table). Each row has a small icon to the left of the table name. Below the table, there is a "Reports" section with a "Middle Pane: dim_order_date Columns" and a status bar showing "Development (WslWarehouse) | UserId: dbo | Browse: No source system | CAP | NUM | SCRL | INS".

Table Name	Table Type
agg_sa_customer	Aggregate
agg_sa_product	Aggregate
dim_customer	Dimension
dim_order_date	Dimension View
dim_product	Dimension
dim_ship_date	Dimension View
fact_budget	Fact Table

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** 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 QAD Data Warehouse Designer 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*)
- ☐ Connections
- ☐ 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_order_header	Tutorial (OLTP)		dbo
load_order_line	Tutorial (OLTP)		dbo
load_prod_group	Tutorial (OLTP)		dbo
load_prod_line	Tutorial (OLTP)		dbo
load_prod_subgroup	Tutorial (OLTP)		dbo
load_product	Tutorial (OLTP)		dbo
load_customer	Tutorial (OLTP)		dbo

Reports

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 **Excel it** 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 QAD Data Warehouse Designer 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

Report: Export Tables by Connection						
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 **Excel it** 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.

Objects Included


The following QAD Data Warehouse Designer 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 of the dimension key in the staging table selected first*)



Select Stage and Dimension Tables (Zero Keys Report) ✕

Select a Stage Table in the top combo and then select a Dimension table to list zero key joins

Stage Table: ▼

Dimension Table: ▼

Results

The results of this report are displayed as a list of refreshed objects in the metadata repository with the following columns:

- ☐ Business Key Column 1 (from staging table)
- ☐ Business Key Column 2, if applicable (from staging table)
- ☐ ...
- ☐ Business Key Column *n*, if applicable (from staging table)

The result set is sortable by clicking on the appropriate column heading.

Report Example

Report: Zero key joins						
budget_quantity	budget_sales_value	dim_customer_key	dim_date_key	dim_product_key	dss_update_time	customer_key
 185	409.92	0	20000602	2	2007-11-23 00:2...	228
 80	978.58	0	20000602	8	2007-11-23 00:2...	228
 62	572.42	0	20010430	3	2007-11-23 00:2...	227
 98	766.17	0	20010430	7	2007-11-23 00:2...	227
 40	218.00	0	20011105	4	2007-11-23 00:2...	226
 40	618.00	0	20011105	6	2007-11-23 00:2...	226
 74	940.24	0	20020404	9	2007-11-23 00:2...	225
 74	163.97	0	20020404	2	2007-11-23 00:2...	225
 40	618.00	0	20020404	6	2007-11-23 00:2...	225
 98	766.17	0	20020404	7	2007-11-23 00:2...	225
 74	402.54	0	20020404	4	2007-11-23 00:2...	225
 15	134.85	0	20011115	3	2007-11-23 00:2...	224
 15	177.34	0	20011115	8	2007-11-23 00:2...	224
 15	159.50	0	20011115	1	2007-11-23 00:2...	224
 74	812.46	0	20000813	1	2007-11-23 00:2...	223
 29	369.17	0	20000813	9	2007-11-23 00:2...	223
 80	622.52	0	20000813	7	2007-11-23 00:2...	223
 29	269.70	0	20000813	3	2007-11-23 00:2...	223

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** to send the results of this report to Microsoft Excel.

External Source Table/files

This report shows external sources for load tables in the metadata repository.

Objects Included

The following QAD Data Warehouse Designer 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
budget.txt	load_budget	File	Windows	..budget.txt
customer	load_customer	Table	Tutorial (OLTP)	[WslTutorial].d
forecast.txt	load_forecast	File	Windows	..forecast.txt
order_header	load_order_header	Table	Tutorial (OLTP)	[WslTutorial].d
order_line	load_order_line	Table	Tutorial (OLTP)	[WslTutorial].d
prod_group	load_prod_group	Table	Tutorial (OLTP)	[WslTutorial].d

Reports

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 **Excel it** 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 QAD Data Warehouse Designer 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:

- ☐ Procedure 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 **Excel it** 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 QAD Data Warehouse Designer 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:

- ☐ Procedure 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



Report: Custom Procedures			
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

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 **Excel it** to send the results of this report to Microsoft Excel.

Index Reports

There is one report for analyzing Indexes:

- ☐ Modified Indexes

Modified Indexes

This report shows indexes in the metadata repository with their creation and modification dates.

Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Indexes

Parameters

There are no parameters for this report.


Results

The results of this report are displayed as a list of indexes in the metadata repository with the following columns:

- ☐ Index Name (*the name of the index*)
- ☐ Dated Created
- ☐ Date Modified (*last modification date*)

The result set is sortable by clicking on the appropriate column heading.

Report Example



The screenshot shows a software window titled "Report: Modified Indexes". It contains a table with three columns: "Index Name", "Date Created", and "Date Modified". The table lists seven indexes, each preceded by a yellow folder icon. The data is as follows:

Index Name	Date Created	Date Modified
dim_product_idx_0	2006-01-05 14:28:51.000	2006-01-05 14:28:51.000
dim_product_idx_A	2006-01-05 14:28:51.000	2006-01-05 14:28:51.000
stage_forecast_idx_A	2006-01-05 14:25:01.000	2006-01-05 14:25:01.000
stage_budget_idx_A	2006-01-05 14:24:44.000	2006-01-05 14:24:44.000
stage_sales_detail_idx_A	2006-01-05 14:24:25.000	2006-01-05 14:24:25.000
dim_customer_idx_0	2006-01-05 14:23:41.000	2006-01-05 14:23:41.000
dim_customer_idx_A	2006-01-05 14:23:41.000	2006-01-05 14:23:41.000

Below the table, there is a "Reports" tab and a status bar. The status bar shows "Ready" and "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 **Excel it** to send the results of this report to Microsoft Excel.

Object Reports

There are four reports for analyzing Objects:

- ☐ Objects-Project Matrix
- ☐ Modified Objects (excluding indexes)
- ☐ Objects Checked-out
- ☐ Loaded or Imported Objects

Objects-Project Matrix

This report shows the objects that exist in projects other than **All Objects**.

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



The screenshot shows a window titled "Report: Objects-Project Matrix". It contains a table with four columns: "Objects", "Dimensions", "Sales", and an empty column. The table lists various objects and their presence in different projects, indicated by a '1' in the project columns.

Objects	Dimensions	Sales	
dim_custo...	1		
dim_custo...	1		
dim_custo...	1		
dim_custo...	1		
dim_product	1		
dim_produc...	1		
dim_produc...	1		
get_dim_cu...	1		
get_dim_pr...	1		
load_budget	1	1	
load_custo...	1	1	
load_forecast	1	1	
load_order_...	1	1	
load_order_l...	1	1	
load_prod_g...	1	1	
load_prod_li...	1	1	
load_prod_s...	1	1	

At the bottom of the window, there is a status bar with the following text: "Reports | Middle Pane: load_order_line 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 **Excel it** to send the results of this report to Microsoft Excel.

Modified Objects (excluding indexes)

This report shows objects in the metadata repository with their creation and modification dates and indicates if they have been modified.

Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Load Tables
- ☐ Stage Tables
- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate Tables
- ☐ 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*)
- ☐ Modified (*a star for modified objects, blank for objects that have not been modified*)

The result set is sortable by clicking on the appropriate column heading.

Report Example

Report: Modified Objects (excluding Indexes)					
Name	Object Type	Date Created	Date Modified	Modified	
 update_dim_customer	Procedure	2006-01-05 14:23:41.000	2011-08-18 18:34:18.670	*	
 update_stage_budget	Procedure	2006-01-05 14:24:44.000	2011-08-18 18:34:04.173	*	
 exp_customer	Export table	2011-08-17 20:05:40.840	2011-08-17 20:05:40.840		
 get_dim_product_key	Procedure	2006-01-05 14:28:50.000	2006-01-05 14:28:50.000		
 update_dim_product	Procedure	2006-01-05 14:28:50.000	2006-01-05 14:28:50.000		
 dim_product	Dimension t...	2006-01-05 14:02:04.000	2006-01-05 14:27:12.000	*	
 load_prod_subgroup	Load table	2006-01-05 14:25:51.000	2006-01-05 14:25:51.000		
 load_prod_line	Load table	2006-01-05 14:25:45.000	2006-01-05 14:25:45.000		
 load_prod_group	Load table	2006-01-05 14:25:38.000	2006-01-05 14:25:38.000		
 update_stage_forecast	Procedure	2006-01-05 14:25:01.000	2006-01-05 14:25:01.000		
 update_stage_sales_d...	Procedure	2006-01-05 14:24:25.000	2006-01-05 14:24:25.000		
 get_dim_customer_key	Procedure	2006-01-05 14:23:41.000	2006-01-05 14:23:41.000		
 update_agg_sa_custo...	Procedure	2006-01-05 14:17:59.000	2006-01-05 14:17:59.000		
 agg_sa_customer	Aggregate ta...	2006-01-05 14:16:07.000	2006-01-05 14:17:51.000	*	
 update_agg_sa_produ...	Procedure	2006-01-05 14:16:42.000	2006-01-05 14:16:42.000		
 agg_sa_product	Aggregate ta...	2006-01-05 14:16:07.000	2006-01-05 14:16:24.000	*	
 update_fact_sales_an...	Procedure	2006-01-05 14:15:19.000	2006-01-05 14:15:19.000		

Reports
Middle Pane: load_order_line 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 **Excel it** 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 QAD Data Warehouse Designer 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

[illegible]

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** to send the results of this report to Microsoft Excel.

Loaded or Imported Objects

This report shows objects in the metadata repository that have been refreshed or imported from another repository.

Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Load Tables
- ☐ Stage Tables
- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate Tables
- ☐ Join Indexes
- ☐ Views
- ☐ Cubes
- ☐ Exports
- ☐ Procedures
- ☐ Host Scripts

Note: Indexes are not included.

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
agg_sa_customer	2011-07-27 15:13:52.087	Auto version on delete/upgrade by A...	
agg_sa_product	2011-07-27 15:13:52.120	Auto version on delete/upgrade by A...	
daily_date_roll	2011-03-25 10:17:27.793	Auto version on delete/upgrade by A...	
daily_date_roll	2011-07-27 15:13:49.893	Auto version on delete/upgrade by A...	
daily_date_roll	2011-07-27 17:09:44.903	Auto version on delete/upgrade by A...	
daily_date_roll	2011-08-17 16:21:18.980	Auto version on delete/upgrade by A...	
dim_customer	2011-07-27 15:13:50.597	Auto version on delete/upgrade by A...	
dim_date	2011-03-25 10:17:29.257	Auto version on delete/upgrade by A...	
dim_date	2011-07-27 15:13:50.747	Auto version on delete/upgrade by A...	
dim_date	2011-07-27 17:09:45.187	Auto version on delete/upgrade by A...	
dim_date_idx_0	2011-03-25 10:17:32.607	Auto version on delete/upgrade by A...	
dim_date_idx_A	2011-03-25 10:17:32.653	Auto version on delete/upgrade by A...	
dim_order_date	2011-07-27 15:13:50.833	Auto version on delete/upgrade by A...	
dim_product	2011-07-27 15:13:50.863	Auto version on delete/upgrade by A...	
dim_ship_date	2011-07-27 15:13:50.983	Auto version on delete/upgrade by A...	
dss_fact_table	2011-07-27 15:13:51.010	Auto version on delete/upgrade by A...	

Reports

Rea Middle Pane: fact_budget 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 **Excel it** to send the results of this report to Microsoft Excel.

Job Reports

There are three reports for analyzing Jobs:

- ☐ Object-Job Matrix
- ☐ 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

Object Name	Action	Job Name	Job Status	Job Last Run
daily_date_roll		No Job		
get_dim_customer_key		No Job		
get_dim_date_key		No Job		
get_dim_order_date_key		No Job		
get_dim_product_key		No Job		
get_dim_ship_date_key		No Job		
update_agg_sa_customer		No Job		
update_agg_sa_product		No Job		
update_dim_customer		No Job		
update_dim_date		No Job		
update_dim_product		No Job		
update_dss_fact_table		No Job		
update_dss_source_system		No Job		
update_fact_budget		No Job		
update_fact_forecast		No Job		
update_fact_sales_analysis		No Job		

Rea Middle Pane: fact_budget 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 **Excel it** to send the results of this report to Microsoft Excel.

Jobs with an Object

This report shows all jobs containing a selected object and lists the job action.

Objects Included

All object types are included in this report.

Parameters

This report has one parameter:

- ☐ Object Name

Select an Object (Jobs Including Object Report)

Select an Object name to list the jobs which include this object

Object Name:

dim_customer

OK

Cancel

Results

The results of this report are displayed as a list of jobs with the following columns:

- ☐ Jobs including **object_name**
- ☐ Action

Report Example

Report: Jobs With an Object		
Jobs including dim_customer	Action	
Daily Update	Process	
Reports		
Rea Middle Pane: fact_budget 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 **Excel it** 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



Select Job (Job Tasks Report)

Select a Job to display tasks and dependencies

Job:


OK Cancel

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



The screenshot shows a window titled "Report: Tasks of a Job" with a table of task dependencies. The table has four columns: Task Name, Action, Order, and Depends On. The tasks are listed in ascending order of their Order number. The first four tasks (load_customer, load_order_header, load_order_line, load_product) are of type "Process" and have no dependencies. The subsequent tasks (dim_customer, dim_date, dim_product) are of type "Process" and have dependencies on the first four tasks.

Task Name	Action	Order	Depends On
load_customer	Process	10.10.1	
load_order_header	Process	10.10.2	
load_order_line	Process	10.10.3	
load_product	Process	10.10.4	
dim_customer	Process	11.10.5	load_customer (10.10.1)
dim_customer	Process	11.10.5	load_order_header (10.10.2)
dim_customer	Process	11.10.5	load_order_line (10.10.3)
dim_customer	Process	11.10.5	load_product (10.10.4)
dim_date	Process	11.10.6	load_customer (10.10.1)
dim_date	Process	11.10.6	load_order_header (10.10.2)
dim_date	Process	11.10.6	load_order_line (10.10.3)
dim_date	Process	11.10.6	load_product (10.10.4)
dim_product	Process	11.10.7	load_customer (10.10.1)
dim_product	Process	11.10.7	load_order_header (10.10.2)
dim_product	Process	11.10.7	load_order_line (10.10.3)
dim_product	Process	11.10.7	load_product (10.10.4)

At the bottom of the window, there is a "Reports" tab and a status bar showing the current context: "Rea Middle Pane: fact_budget 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 **Excel it** to send the results of this report to Microsoft Excel.

Operational Reports

There are four Operational Reports:

- ☐ Object Performance History
- ☐ Job Performance History
- ☐ Task Performance History
- ☐ Fragmentation

Object Performance History

This report shows the audit trail from the scheduler logs for a selected object.


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



The screenshot shows a window titled "Report: Object Performance History". It contains a table with the following columns: Sta, Time, Seq, Message, Task, Database ..., and Db Code. The table lists 16 audit log entries. The first 10 entries have a Seq of 119 and a Date of 2011-08-17 16:24:0... The next 5 entries have a Seq of 151 and a Date of 2011-09-16 14:52:1... The last entry has a Seq of 166 and a Date of 2011-09-16 16:25:5... The Message column contains details about deleted and inserted rows for various objects. The Task column shows "update_ag...". The Database and Db Code columns are empty.

Sta	Time	Seq	Message	Task	Database ...	Db Code
I	2011-08-17 16:24:0...	119	Date: 20000601 Deleted: 0 Inserted: 1.	update_ag...		
I	2011-08-17 16:24:0...	119	Date: 20000801 Deleted: 0 Inserted: 1.	update_ag...		
I	2011-08-17 16:24:0...	119	Date: 20010401 Deleted: 0 Inserted: 1.	update_ag...		
I	2011-08-17 16:24:0...	119	Date: 20011101 Deleted: 0 Inserted: 2.	update_ag...		
I	2011-08-17 16:24:0...	119	Date: 20020401 Deleted: 0 Inserted: 1.	update_ag...		
I	2011-09-16 14:52:1...	151	Date: 20000601 Deleted: 0 Inserted: 1.	update_ag...		
I	2011-09-16 14:52:1...	151	Date: 20000801 Deleted: 0 Inserted: 1.	update_ag...		
I	2011-09-16 14:52:1...	151	Date: 20010201 Deleted: 0 Inserted: 2.	update_ag...		
I	2011-09-16 14:52:1...	151	Date: 20010401 Deleted: 0 Inserted: 1.	update_ag...		
I	2011-09-16 14:52:1...	151	Date: 20011101 Deleted: 0 Inserted: 2.	update_ag...		
I	2011-09-16 14:52:1...	151	Date: 20020101 Deleted: 0 Inserted: 1.	update_ag...		
I	2011-09-16 14:52:1...	151	Date: 20020401 Deleted: 0 Inserted: 1.	update_ag...		
I	2011-09-16 16:25:5...	166	Date: 20000601 Deleted: 1 Inserted: 1.	update_ag...		
I	2011-09-16 16:25:5...	166	Date: 20000801 Deleted: 1 Inserted: 1.	update_ag...		
I	2011-09-16 16:25:5...	166	Date: 20010201 Deleted: 2 Inserted: 2.	update_ag...		
I	2011-09-16 16:25:5...	166	Date: 20010401 Deleted: 1 Inserted: 1.	update_ag...		

Below the table, there is a "Reports" section with a "Middle Pane: agg_sa_customer Columns | Development (WslWarehouse) | UserId: dbo | Browse: No source system | CAP | NUM | SCRL".

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** to send the results of this report to Microsoft Excel.

Job Performance History

This report shows the performance (duration) over time of a selected job.


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: Daily Update ▼

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

Report: Job Performance History				
Job Name	Start time	Elapsed hh:mi		
Daily Update	2011-08-18 20:14:28.127	00:00		

Reports

Rea Middle Pane: fact_budget 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 **Excel it** to send the results of this report to Microsoft Excel.

Task Performance History

This report shows the performance (duration) over time of a selected task within a selected job.


Objects Included

All object types are included in this report.

Parameters

This report has two parameters:

- ☐ Job Name
- ☐ Task Name (including action)




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
dim_customer	Process	2011-08-18 20:14:31.150	00:00

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** to send the results of this report to Microsoft Excel.

Fragmentation

This report is only available on a SQL Server database.

Objects Included

All object types are included in this report.

Parameters

This report has no parameters.

Results

The results of this report are displayed as a list of tables within the database with the following columns:

- ☐ Database
- ☐ Table Name
- ☐ Index ID
- ☐ Index Name
- ☐ Logical Fragmentation
- ☐ Avg Page Space Used
- ☐ Avg Fragment Size in Pages
- ☐ Fragment Count
- ☐ Page Count
- ☐ Size in GB

Report Example

Database ...	Table Name	Index ID	Index Name	Logical Fra...	Avg Page S...	Avg Fragm...	Fragment ...	Page Count	Si
WslWar...	dim_date	2	dim_date_i...	31	97	2	8	19	0
WslWar...	dim_date	3	dim_date_i...	21	97	3	9	28	0
WslWar...	dim_product	1	dim_produ...	0	14	1	1	1	0
WslWar...	dim_product	2	dim_produ...	0	1	1	1	1	0
WslWar...	dss_fact_ta...	0		0	2	0	1	1	0
WslWar...	dss_source...	0		0	1	0	1	1	0
WslWar...	stage_bud...	0		0	13	0	1	1	0
WslWar...	stage_bud...	2	stage_bud...	0	6	1	1	1	0
WslWar...	stage_fore...	0		0	13	0	1	1	0
WslWar...	stage_fore...	2	stage_fore...	0	6	1	1	1	0
WslWar...	stage_sales...	0		0	25	0	1	1	0
WslWar...	stage_sales...	2	stage_sales...	0	7	1	1	1	0
WslWar...	fact_budget	0		0	13	0	1	1	0
WslWar...	fact_budget	2	fact_budg...	0	3	1	1	1	0
WslWar...	fact_budget	3	fact_budg...	0	3	1	1	1	0

Reports

Rea Middle Pane: fact_budget 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 **Excel it** to send the results of this report to Microsoft Excel.

Chapter 25

Validate

When these Validate processes are run, the results are displayed in the middle pane of the Data Warehouse Designer screen; the results of the other reports are displayed in a separate tab in the bottom pane of the Data Warehouse Designer screen.

In This Chapter

Validate Meta-data	840
Validate Workflow Data.....	841
Validate Table Create Status	842
Validate Load Table Status.....	843
Validate Index Status.....	844
Validate Procedure Status.....	845
List Meta-data Tables not in the Database	846
List Database Tables not in the Meta-data	847
List Tables with no related Procedures or Scripts	849
List Procedures not related to a Table.....	851
Compare Meta-data Repository to another.....	852
Compare Meta-data Indexes to Database.....	855
List Duplicate Business Key Columns	856
Query Data Warehouse Objects	857

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 **Excel it** 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 **Excel it** 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.

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.

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.

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** 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 **Excel it** to send the results of this report to Microsoft Excel.

Validate Index Status

This process validates an index structure in the meta data against the index in the database. Select one or more indexes and click the **Validate Selected** button or click the **Validate All** button to validate all indexes.

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** 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 **Excel it** 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 QAD Data Warehouse Designer object types are included in this report:

- ☐ Load Tables
- ☐ Stage Tables
- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate Tables
- ☐ 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

[illegible]

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** 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 exist in the data warehouse database but not in the metadata.

Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Load Tables
- ☐ Stage Tables
- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate Tables
- ☐ 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 data warehouse database not in the metadata.

Report Example

Report: Tables in Database but not in Metadata

Table Name
load_promotions

Ready Middle Pane: Object Results 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 **Excel it** 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 QAD Data Warehouse Designer object types are included in this report:

- ☐ Load Tables (*script based loads only*)
- ☐ Stage Tables
- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate Tables
- ☐ Data Store Tables
- ☐ Normalized 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
 fact_budget

Read: Middle Pane: Fact Table Columns | Development (WslWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** 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 QAD Data Warehouse Designer 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

[illegible]

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** to send the results of this report to Microsoft Excel.

Compare Meta-data Repository to another

This report compares the current metadata repository to a remote repository.

Objects Included

All object types are included in this report.

Parameters

This report requires connection information for the remote repository to be entered. Specifically:

- ☐ Odbc Connect (*the odbc source for the other repository*)
- ☐ User Name (*user name for the remote repository*)
- ☐ Password (*password for the remote repository*)
- ☐ Meta Repository (*database/user the remote repository metadata is stored in*)

Four additional parameters may be specified:

- ☐ Optional filter on Local Groups
- ☐ Optional filter on Local Projects
- ☐ Do detail report
- ☐ Only validate procedures that have been modified

Compare two metadata repositories

Compare the current metadata repository with the repository selected below. Enter the connection and logon information for the repository to be compared and press OK.

Odbc Connect

Wsl/warehouse

User Name

dbo

Password

Meta Repository

dbo

Local Groups:

Local Projects:

This is usually a two pass process. The first pass will report each object with validation errors. A more detailed report for listed objects can be obtained by selecting the objects and right clicking detail.

You can however go straight to the detail report by selecting the detail option below.

☐ Detail Report Only

☐ Validate Modified Procedures Only

NOTE: No sizing, tablespace, file group, date, description or documentation comparisons are performed.

WARNING: This will take a very long time for large repositories.

OK

Cancel

Results

The results of this report are displayed as a list of differences with the following columns:

- ☐ Object Name
- ☐ Comments (*Summary difference between current and selected repository*)

The result set is sortable by clicking on the appropriate column heading.

Report Example

Object Name	Comments
Cubes	Connection name not found
DataWarehouse	5 validation errors
Tutorial (OLTP)	3 validation errors
Windows	2 validation errors
load_budget	Load table not found in remote repository
load_customer	Load table not found in remote repository
load_forecast	Load table not found in remote repository
load_order_header	Load table not found in remote repository
load_order_line	Load table not found in remote repository
load_prod_group	Load table not found in remote repository
load_prod_line	Load table not found in remote repository
load_prod_subgroup	Load table not found in remote repository
load_product	Load table not found in remote repository
stage_budget	Stage table not found in remote repository
stage_forecast	Stage table not found in remote repository
stage_sales_detail	Stage table not found in remote repository

Reports

Read: Middle Pane: Fact Table Columns | Development (WslWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Note: Right-clicking on an object and choosing **Detail** will rerun the report for just the selected object, returning more detail in the **Comments** column

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** to send the results of this report to Microsoft Excel.

Compare Meta-data Indexes to Database

This report shows database indexes in the data warehouse database that don't exist in the metadata.

Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Indexes

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of indexes in the data warehouse database not in the metadata with the following columns:

- ☐ Index Name (*the name of the index*)
- ☐ Table Name

Report Example

Index Name	Table Name
idx_prods	load_promotions

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** to send the results of this report to Microsoft Excel.

List Duplicate Business Key Columns

This report shows any columns that are the business (natural) key of more than one table.

Objects Included

The following QAD Data Warehouse Designer object types are included in this report:

- ☐ Stage Tables
- ☐ Dimension Tables and Views
- ☐ Fact Tables
- ☐ Aggregate Tables

Parameters

There are no parameters for this report.

Results

The results of this report are displayed as a list of table columns with the following columns:

- ☐ Table name
- ☐ Column name

The result set is sortable by clicking on the appropriate column heading.

Report Example

[illegible]

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** to send the results of this report to Microsoft Excel.

Query Data Warehouse Objects

This report allows 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 COUNT(*)  
FROM fact_sales_detail
```

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 the "Data Warehouse Query" window. The title bar includes standard window controls (minimize, maximize, close). Below the title bar is a light gray header area. The main area is a large white grid with horizontal lines, representing a query editor. On the left side of the grid, there is a small blue icon and the number "21". At the bottom of the window, there is a dark blue status bar containing the following information: "Read Middle Pane: Fact Table Columns | Development (WslWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS".

Sending Results to Microsoft Excel

Right-click on the result set and click **Excel it** to send the results of this report to Microsoft Excel.

Chapter 26

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. QAD Data Warehouse Designer 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 (see **Backup and Restore** (see "**Backing Up and Restoring Metadata**" on page 877)).

In This Chapter

Applications	860
Importing Object Metadata	867
Importing Language Files	869
Data Warehouse Testing	870
Creating and Loading Applications from the Command Line	874

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. An Oracle application cannot be loaded into a SQL Server metadata repository).

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 in which they are listed in the application. If procedure dependencies exist, ensure that 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 with the 'Application' tab selected. The dialog is titled 'Build Deployment Application' and has a close button (X) in the top right corner. On the left, there is a sidebar with three tabs: 'Application' (selected), 'Objects to Add/Replace', and 'Objects to Delete'. The main area contains the following fields and sections:


- Output Directory:** A text box containing 'C:\Training\QAD\Module18\' 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 large 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: 'Pre Application Load SQL. (the following optional SQL statement will be issued before the application load commences):'.
- Post Application Load SQL:** A text area with the instruction: 'Post Application Load SQL. (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 QAD program directory.

Note: 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.

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. (For Example an Oracle application can only be loaded into an Oracle 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/filegroups 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.

Refer to the Setup Administrator manual for more information about loading an application.

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.



The dialog box is titled "Import Objects from Another Meta Repository" and contains the following text and fields:

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.

IMPORT

 **Odbc Connect** ORCL

 **User Name** \sqlwarehouse

 **Password** xxx


 **Meta Repository** \sqlwarehouse

OK Cancel

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 will appear notifying of 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 an Oracle application can only be loaded into an Oracle 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 an Oracle application can only be loaded into an Oracle 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 QAD Data Warehouse Designer. 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:testing

Description:This job runs a procedure to perform a series of tests on the Data Warehouse

Frequency:Once and Hold

Start Date:Thursday , February 14, 2013

Start Time:4:37:00 PM

Maximum Threads:2

Scheduler:Windows Only

Dependent On:Parent jobFailLook back (minutes)Maximum wait (minutes)

Logs Retained:0

Success Command:

Failure Command:

Custom Settings

Interval Between Jobs: (minutes)

Start At or After HHMM (e.g. 0800):0000

Do Not Start After HHMM (e.g. 1700):0000

Active Days:

☐ Mon☐ Tue☐ Wed☐ Thu☐ Fri☐ Sat☐ Sun

Add Parent Job

Remove Parent

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:

OK

Cancel

- 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**.

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. The QAD Data Warehouse Designer Application directory contains an example bat file WSL_Application_Create_Restore_Point_and_Load.bat for creating and loading.

If you right-click on this file and select **Edit** you will see the steps outlined as well as the details on the options available.

The first step **creates** a restore point application (R) based on objects about to be loaded. This process calls the command line functionality of Data Warehouse Designer and creates the application file.

The second step **loads** an application (A) into a test QAD Data Warehouse Designer repository. It uses an xml file to specify various options, calls the command line automation functionality of Setup Administrator and loads the application (A). The QAD Data Warehouse Designer Application directory contains example xml files:

- ☐ WSL_Application_Load_SQL.xml for SQL Server
- ☐ WSL_Application_Load_DB2.xml for DB2
- ☐ WSL_Application_Load_ORA.xml for Oracle

The tags in the xml file need to be edited as the login/connection details etc need setting.

The Batch Application Create options and the Batch Application Load options are listed at the end of the bat file. The values for these variables will also need to be customized before running this file.

The process typically involves the following steps:

- 1 Create an application (A) in Data Warehouse Designer containing your data warehouse changes.
- 2 RUN WSL_Application_Create_Restore_Point_and_Load.bat from the command line; which creates a restore point application (R) and applies application (A) to a test QAD Data Warehouse Designer repository.
- 3 If the changes are incorrect, they can be undone by loading the restore point application (R).

Note: You can only create a RESTORE application through the command line; deployment applications are still built through Data Warehouse Designer.

Batch Application CREATE options

/BA selects batch application create.
/U username.
/P password.
/C DSN name.
/M metaschema/database for DB2/Teradata logon.
/D directory to save application files to.
/I new application files identifier.
/V new application files version.
/AP Project name - all objects in a Project; and associated jobs.
/AG Group name - all objects in a Group; and associated jobs.
/ALL All objects - all objects, jobs and parameters.
/RC Remove connections - drop all connections from the application.
/RJ Remove jobs - drop all jobs from the application.
/RP Remove parameters - drop all parameters from the application.
/AF absolute application id file name which restore point is being created for.
If a trusted connection, the user and password are not necessary.

Batch Application LOAD options

/AL select application load
/AF absolute application id file name
/LF absolute log file name
/PF absolute xml parameter file name

Chapter 27

Backing Up and Restoring Metadata

This chapter covers the 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 restore section covers the metadata restoration functions available.

QAD Data Warehouse Designer provides two different methods for backing up and restoring metadata:

- 1 Backup and Restore using the native backup tools available in the data warehouse database.
- 2 QAD Data Warehouse Designer unload and load.

In This Chapter

Backup using DB Routines.....	878
Restoring DB Backups.....	881
Unloading Metadata.....	885
Loading an Unload	888

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 is backed up at least daily when possible. There are two methods for backing up the meta repository. These are through a scheduled UNIX task, and through the main QAD Data Warehouse Designer tool.

UNIX Backups

Where the data warehouse database is running on a UNIX server, it is possible to schedule a regular backup of the metadata. The UNIX scripts shipped with QAD Data Warehouse Designer include a script called from cron to daily backup the metadata. These backups are in turn compressed and saved for 60 days by default. Refer to the **Installation and Administration Guide** for details on these scripts and their setup.

Windows Backups

Two main methods of Windows backup exist within QAD Data Warehouse Designer. The first is a database independent backup which is designed purely for moving the meta repository, and should not be used for regular backups. The other method is the database specific backup. See the following sections for the appropriate database.

Oracle Windows Backups

An Oracle export based backup can be taken from the QAD Data Warehouse Designer tool by selecting the **Backup/Export the metadata (Oracle Export)** menu option. This option assumes that the Oracle Client on the PC is the same version as that of the Oracle database where the QAD Data Warehouse Designer meta repository is stored. The backup may not work if the versions differ.

When executed this menu option attempts to locate the Oracle export utility. This utility is normally called 'exp' and resides in the Oracle bin directory. If QAD Data Warehouse Designer cannot locate this utility or it has not been loaded onto the PC then the export does not proceed and an error message will be displayed.

A pop-up window asks for a file name for the export. A directory should be chosen and a name entered. When the **Save** button is clicked the export will start. The following files are created (where 'file' is the name of the chosen file). A dialog box appears to show the results of the backup.

File name	Purpose
file.bat	Windows command file containing the export command.
file.dbl	Command file for the re-creation of any database links.
file.drp	Command file for dropping all metadata tables. (Used in the restore prior to a reload).
file.exp	Oracle export file.
file.log	Log file of the export.
file.par	Parameter file for the export. Contains an entry for each meta table to be exported.

File name	Purpose
file.pro	A file containing each procedure stored in the metadata.
file.seq	Command file to re-create the sequences used by both the metadata and the data warehouse tables.

If problems are encountered with the backup it may be possible to manually run the generated .bat script file to ascertain what has gone wrong.

SQL Server Windows Backups

A SQL Server bcp based backup can be taken from the QAD Data Warehouse Designer tool by selecting the **Backup/Export the metadata (Sql server bcp)** menu option. This option assumes that the SQL Server bcp utility is available on the PC.

When executed this menu option attempts to locate the SQL Server bcp utility. This utility is normally called 'bcp' and resides in the SQL Server bin directory. If QAD Data Warehouse Designer cannot locate this utility or it has not been loaded onto the PC then the export does not proceed and an error message will be displayed.

A pop-up window asks for a directory name for the export. A directory should be chosen or created. When the **Save** button is clicked the export will start. The following files are created. A dialog box appears to show the results of the backup.

File name	Purpose
wsl_bcp_out.bat	Windows batch file used to run all the bcp table exports
wsl_bcp_out.err	Errors from the bcp export of each table
wsl_bcp_out.log	Log file of the bcp table exports
*.bcp	Each metadata table is stored with a .bcp extension. This is the export of each table.
wsl_bcp.par	Parameter file for the export. Contains an entry for each meta table to be exported. Unused by SQL Server.
wsl_bcp.drp	Command file for dropping all metadata tables. (Used in the restore prior to a reload).
wsl_bcp.pro	A file containing each procedure stored in the metadata.
wsl_bcp.seq	Unused by SQL Server.
wsl_bcp.dbl	Command file for the rebuild of all the linked servers. Used during the restore.
wsl_bcp.bld	Basis of the command file that will be used to restore the metadata. This file is modified to include the target directory prior to restore.

If problems are encountered with the backup it may be possible to manually run the generated .bat script file to ascertain what has gone wrong.

DB2 Windows Backups

A DB2 IMPORT based backup can be taken from the QAD Data Warehouse Designer tool by selecting the **Backup/Export the metadata (DB2 export)** menu option. This option assumes that the IBM DB2 EXPORT command is available on the PC.

When executed this menu option attempts to locate the SQL Server bcp utility. This utility is normally called **bcp** and resides in the SQL Server bin directory. If QAD Data Warehouse Designer cannot locate this utility or it has not been loaded onto the PC then the export does not proceed and an error message will be displayed.

A pop-up window asks for a directory name for the export. A directory should be chosen or created. When the **Save** button is clicked the export will start. The following files are created. A dialog box appears to show the results of the backup.

File name	Purpose
Data Warehouse Designer.bat	Windows batch file used to run the DB2 EXPORT for each metadata table.
Data Warehouse Designer.cmd	Control file for the export. Contains an entry for each meta table to be exported.
Data Warehouse Designer.log	Log file of the EXPORT.
Data Warehouse Designer.pro	A file containing each procedure stored in the metadata.
*.ixf	Each metadata table is stored in a file with a ixf extension. This is the export of each table.
restore.drp	Command file for dropping all metadata tables. (Used in the restore prior to a reload).
restore.lst	Basis of the command file that will be used to restore the metadata. This file is modified to include the target directory prior to restore.

If problems are encountered with the backup it may be possible to manually run the generated Data Warehouse Designer.bat script file to ascertain what has gone wrong.

Restoring DB Backups

QAD Data Warehouse Designer metadata can be restored from a prior backup. For Oracle data warehouses the restore can take place either from UNIX or through the QAD Data Warehouse Designer tool. The same set of files as defined in the backup section are used for both restores and can be moved between the Windows and UNIX environments. For SQL Server data warehouses the restore must take place in a Windows environment either through the QAD Data Warehouse Designer tool or manually.

Note: If transferring the export files via ftp, ensure that the file.exp and .bcp files are transferred in binary mode and that the other files are transferred in ASCII mode.

Restoring metadata


A backup of a metadata repository can be restored over the top of an existing repository. This action replaces the existing repository in its entirety. The restore of repositories is covered in an earlier section, but its uses in the promoting of code are discussed here. This restore process does not affect any existing database tables or any compiled procedures, so it can be used as a means of updating a metadata repository, albeit by replacing that repository. It is often a good method to choose when first establishing a new repository. For example if we have a development repository and want to create a production environment the steps may be:

- 1 Backup the development repository, using either the UNIX or Windows repository backup. See the previous sections.
- 2 Create a new repository using the **Setup/Administrator** utility.
- 3 Log on to the newly created repository.
- 4 Restore the backup of the development repository into the new repository. For Oracle this requires the database privilege 'imp_full_database', or the dba role.
- 5 Modify any connections, set any default values, alter any table sizes or extents.
- 6 Create the tables.
- 7 Compile all procedures.

Oracle Windows Restore

Select the menu option **Backup/Restore Metadata (from Oracle Export)** to begin the restore process.

A dialog box will appear. The word **RESTORE** needs to be entered along with the **User Name** and **Password** where the metadata is to be restored to. The **User Name** does not have to be that of the current metadata repository, but it must be a valid repository. If a restore is being performed to a repository that differs from that of the backup then the 'imp_full_database' database privilege may be required.



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.

RESTORE

 **Odbc Connect**

 **User Name**

 **Password**

OK Cancel

Once the **OK** button is clicked a new dialog box will appear asking for a selection of the export file. Browse to the directory where the export is located and select the export file. Once the export file is selected the import will begin. A dialog box will appear to show the results of the import.

Oracle UNIX restore


The UNIX restore is actioned via the meta_restore_010.sh shell script. This script expects the export file to be located in the wsl/export directory, so it may need to be un compressed and moved to this directory if not already present.

Execute the meta_restore_010.sh shell script. You will be prompted to enter the backup file name. Once successfully entered a username and password to be restored into are requested and the restore commences.

SQL Server Windows Restore

Select the menu option **Backup/Restore the metadata (from Sql server bcp)** to begin the restore process.

A dialog box will appear. The word **RESTORE** needs to be entered along with the connection, **User Name** and **Password** where the metadata is to be restored to. The metadata will always be restored to the dbo schema. The connection does not have to be that of the current metadata repository, but it must be a valid repository.



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.

RESTORE

Odbc Connect wsl/warehouse

User Name dbo

Password


OK Cancel

Once the **OK** button is clicked a new dialog box will appear asking for a selection of the export directory. Browse to the contents of the directory where the export is located. Once the export directory is selected the import will begin. A dialog box will appear to show the results of the import.

DB2 Windows Restore

Select the menu option **Backup/Restore metadata (from DB2 export)** to begin the restore process.

A dialog box will appear. The word **RESTORE** needs to be entered with the **User Name** and **Password** for DB2. If using a trusted connection to DB2, the username and password should be left blank. The metadata schema should also be entered. The schema does not have to be that of the current metadata repository, but it must contain a valid repository.



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.

RESTORE

Odbc Connect DBTWO

User Name

Password

Meta Schema dssdemo

OK Cancel

Once the **OK** button is clicked a new dialog box will appear asking for a selection of the export file. Browse to the directory where the export is located and select the export file. Once the export file is selected the import will begin. A dialog box will appear to show the results of the import.

Unloading Metadata

QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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

QAD Data Warehouse Designer 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 **User Name** 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**.



The dialog box is titled "Restore Metadata from Backup" and features a red close button in the top right corner. The main content area contains a warning message: "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." Below the warning is a text input field containing the word "RESTORE". Further down, there are three labeled input fields: "Odbc Connect" with a dropdown menu showing "wslwarehouse", "User Name" with a text field containing "dbo", and "Password" with an empty text field. At the bottom right, there are "OK" and "Cancel" buttons.

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.

RESTORE

Odbc Connect wslwarehouse


User Name dbo

Password


OK Cancel

Note: The dialog box above will include an additional field for *Schema* if loading into DB2.


The next dialog box asks for a selection of the unload directory. Browse to the contents of the directory where the unload is located and click **Open**.



Confirm the load will overwrite the existing repository by clicking **Yes**:




Confirm the metadata and dimension tablespaces/filegroups for the metadata load and click **OK**:




Note: The dialog above was displayed when loading into SQL Server. It may differ slightly if loading into Oracle or DB2.

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 QAD Administrator and validate the metadata repository that has just been loaded.

Chapter 28

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.....	894
Validating Source (Load) Tables.....	896
Validating Procedures	897
Altering Tables	898
Validating Indexes	900
Recompiling Procedures	901

Validating Tables

The metadata as stored and maintained by QAD Data Warehouse Designer 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 validation list	
Table Name	Results
 agg_sa_customer	Validates OK
 agg_sa_product	Validates OK
 dim_date	Validates OK - but different column order
 dim_order_date	Validates OK
 dim_product	dss->description - varchar(64), dss->subgroup_description - varchar(64), dss->line_description - varchar(6
 dim_ship_date	Validates OK
 dss_fact_table	Validates OK
 dss_source_system	Validates OK
 fact_budget	Validates OK
 fact_sales_analysis	Validates OK
 fact_sales_detail	Validates OK
 load_budget	dss->prod_code - int, meta->product_code - integer,
 load_customer	Validates OK
 load_forecast	Validates OK
 load_order_header	Validates OK
 load_order_line	Validates OK
 load_prod_group	Validates OK
 load_prod_line	Validates OK
 load_prod_subgroup	Validates OK
 load_product	meta->descr2 - varchar(24),
 load_state	Table:meta->load_state column data not found,
 stage_budget	Validates OK

Reports

Middle Pane: dim_customer Columns | Development (WslWarehouse) | UserId: dbo | Browse: No source system | CAP | NUM | SCRL

In this example we see five different scenarios.

The metadata for table **agg_sa_customer** matches the physical table in the database.

The table **dim_date** 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 altering of the table. The table must be re-created if the order is important.


The physical database table **dim_product** has an additional columns not found in the metadata. The columns are 'subgroup_description' and 'line_description'. The table can be altered if desired. See the next section on Altering Tables.

The metadata for the table **load_product** does not match the physical table. The metadata has an additional column called 'state'. This column was not found in the physical table. The table can be altered if desired. See the next section on Altering Tables.

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. load table to create dim 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 reoccurring in such a situation proceed as follows. Use the right-click menu and 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 validation list	
Load Table	Results
load_budget	Source is non database, unable to validate;
load_customer	Validates OK
load_forecast	Source is non database, unable to validate;
load_order_header	Validates OK
load_order_line	Validates OK
load_prod_group	Validates OK
load_prod_line	Validates OK
load_prod_subgroup	Validates OK
load_product	Validates OK
load_state	MetaData-> column data not found;

The tables **load_budget** and **load_forecast** are Windows file loads and as such cannot be validated.

If a table 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 (dimension, stage and fact) 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 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.



The screenshot shows a window titled "Procedure Validation List" with a table containing two columns: "Procedure Name" and "Results". The table lists 19 procedures. Most have a status of "Validates OK", but "update_stage_budget" has a status of "Compare failed". The window includes a search bar at the bottom left and a status bar at the bottom with various system information.

Procedure Name	Results
daily_date_roll	Validates OK
dss_test	Validates OK
get_dim_date_key	Validates OK
get_dim_order_date_k...	Validates OK
get_dim_product_key	Validates OK
get_dim_ship_date_key	Validates OK
update_agg_sa_custo...	Validates OK
update_agg_sa_produ...	Validates OK
update_dim_date	Validates OK when spaces compressed. Edit procedure and load viewer with compile code and compare to
update_dim_product	Validates OK
update_dss_fact_table	Validates OK
update_dss_source_sy...	Validates OK
update_fact_budget	Validates OK
update_fact_sales_an...	Validates OK
update_fact_sales_det...	Validates OK
update_stage_budget	Compare failed
update_stage_forecast	Validates OK


Read: Middle Pane: Fact Table Columns | Development (WslWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

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 use the right-click menu option from the object name. Then in the middle pane (the validation listing) right-click on the table that has not validated and select **Alter table**. 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. Comments at the top of the screen shown 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.


Pressing the **Alter Table** button will proceed to alter the table. In effect it will execute any command in the window.

Validating Indexes

The menu option **Validate/Validate Index Status** checks the metadata definition of an index against that in use in the database. It checks to ensure that the index exists and that the index type and columns are the same.

The results are listed in the middle pane with the status of each selected index shown.

If an index differs, then the only option is to drop and create the index either via the scheduler or through the right-click menu options for the index.




Index Name	Results
agg_sa_customer_idx_1	Validates OK
agg_sa_customer_idx_2	Validates OK
agg_sa_product_idx_1	Validates OK
agg_sa_product_idx_2	Validates OK
dim_date_idx_0	Validates OK
dim_date_idx_A	Validates OK
dim_product_idx_0	Validates OK
dim_product_idx_A	Validates OK
fact_product_idx_1	Validates OK
fact_sales_analysis_idx_1	Validates OK
fact_sales_analysis_idx_2	Validates OK
fact_sales_analysis_idx_3	Validates OK
fact_sales_detail_idx_1	Validates OK
fact_sales_detail_idx_2	Validates OK
fact_sales_detail_idx_3	Validates OK
fact_sales_detail_idx_4	Validates OK
fact_sales_detail_idx_A	Validates OK

Reports

Re Middle Pane: dim_product Columns | Development (WslWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

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 will load the following dialog box:



All procedures in the metadata can be compiled, or a selected group of procedures may be compiled. Selection is done via standard Windows multi selection using the shift and CTRL keys. By selecting a project or group it is possible to compile all procedures associated with the project or group.

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.

Re-validate Procedures (Oracle only)

The menu option **Tools/Re-compile Invalid Procedures** will attempt to validate the procedure by issuing an 'alter procedure xxx compile' command. This command will be executed for any procedure or function that is in an invalid state in the schema. Any procedures or functions that are still invalid after the commands have been completed are shown in the results dialog. QAD Data Warehouse Designer does not drop procedures from the database even when their metadata is deleted. Therefore, there may be old procedures present that are in an invalid state. These procedures may be redundant.

Chapter 29

Callable Routines

In This Chapter

Introduction to Callable Routines	904
Ws_Api_Glossary	912
Ws_Connect_Replace	916
Ws_Job_Abort	921
Ws_Job_Clear_Archive	924
Ws_Job_Clear_Logs	930
Ws_Job_Clear_Logs_By_Date	935
Ws_Job_Create	940
Ws_Job_CreateWait	948
Ws_Job_Dependency	956
Ws_Job_Release	962
Ws_Job_Restart	967
Ws_Job_Schedule	972
Ws_Job_Status	977
Ws_Load_Change	985
Ws_Maintain_Indexes	990
Ws_Version_Clear	996
WsParameterRead	1001
WsParameterReadF	1004
WsParameterReadG	1006
WsParameterReadJ	1010
WsParameterWrite	1013
WsWrkAudit	1016
WsWrkAuditBulk	1022
WsWrkError	1028
WsWrkErrorBulk	1034
WsWrkTask	1040

Introduction to Callable Routines

Callable Routines API

QAD Data Warehouse Designer callable routines provide an Application Program Interface (API) to the QAD Data Warehouse Designer metadata using the following SQL-invoked routines:

Routine Name	Description
Ws_Api_Glossary (on page 912)	Adds an entry to the documentation glossary.
Ws_Connect_Replace (on page 916)	Replaces the contents of a connection with details from another connection.
Ws_Job_Abort (on page 921)	Aborts a job if it is in a running state.
Ws_Job_Clear_Archive (on page 924)	Purges archived job logs that are older than the specified age in days.
Ws_Job_Clear_Logs (on page 930)	Archives job logs when the maximum number of current logs to retain is exceeded.
Ws_Job_Clear_Logs_By_Date (on page 935)	Archives job logs that are older than the specified age in days.
Ws_Job_Create (on page 940)	Creates a job based on an existing job and optionally starts it immediately.
Ws_Job_CreateWait (on page 948)	Creates a job based on an existing job and schedules it to start later.
Ws_Job_Dependency (on page 956)	Adds or removes a child-to-parent dependency between two jobs to control the child job.
Ws_Job_Release (on page 962)	Starts a job if it is in a holding or waiting state.
Ws_Job_Restart (on page 967)	Starts a job if it is in a failed state.
Ws_Job_Schedule (on page 972)	Schedules a job if it is in a holding or waiting state.
Ws_Job_Status (on page 977)	Returns the current status of a job.
Ws_Load_Change (on page 985)	Changes the Connection or Schema of a load table.
Ws_Maintain_Indexes (on page 990)	Drops and/or builds database indexes that are defined in the QAD Data Warehouse Designer metadata.
Ws_Version_Clear (on page 996)	Purges metadata versions for all objects that do not meet the specified retention criteria.
WsParameterRead (see "WsParameterReadF" on page 1004)	Returns the value and comment (for most RDBMS) of a QAD Data Warehouse Designer metadata Parameter.
WsParameterReadF (on page 1004)	Returns the value of a QAD Data Warehouse Designer metadata Parameter. [SQL Server only]

Routine Name	Description
WsParameterReadG (on page 1006)	Returns the value of a "global" QAD Data Warehouse Designer metadata Parameter that relates to a load table.
WsParameterReadJ (on page 1010)	Returns information about the job that initiated a running procedure.
WsParameterWrite (on page 1013)	Updates the value and comment of a QAD Data Warehouse Designer metadata Parameter or creates it.
WsWrkAudit (on page 1016)	Records a message in the Audit Log.
WsWrkAuditBulk (on page 1022)	Records multiple messages in the Audit Log.
WsWrkError (on page 1028)	Records a message in the Error/Detail Log.
WsWrkErrorBulk (on page 1034)	Records multiple messages in the Error/Detail Log.
WsWrkTask (on page 1040)	Updates row counts for a task in the Task Log.

Callable Routines per RDBMS

Each QAD Data Warehouse Designer 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 QAD Data Warehouse Designer API was originally developed.

Routine Name	SQLServer	Teradata	Oracle	DB2
Ws_Api_Glossary (on page 912)	Procedure	Procedure	Function	Procedure
Ws_Connect_Replace (on page 916)	Procedure	Procedure	Function	Procedure
Ws_Job_Abort (on page 921)	Procedure	Procedure	Function	Procedure
Ws_Job_Clear_Archive (on page 924)	Procedure	Procedure	Function	Procedure
Ws_Job_Clear_Logs (on page 930)	Procedure	Procedure	Function	Procedure
Ws_Job_Clear_Logs_By_Date (on page 935)	Procedure	Procedure	Function	Procedure
Ws_Job_Create (on page 940)	Procedure	Procedure	Function	Procedure
Ws_Job_CreateWait (on page 948)	Procedure	Procedure	Function	Procedure
Ws_Job_Dependency (on page 956)	Procedure	Procedure	Function	Procedure

Routine Name	SQLServer	Teradata	Oracle	DB2
Ws_Job_Release (on page 962)	Procedure	Procedure	Function	Procedure
Ws_Job_Restart (on page 967)	Procedure	Procedure	Function	Procedure
Ws_Job_Schedule (on page 972)	Procedure	Procedure	Function	Procedure
Ws_Job_Status (on page 977)	Procedure	Procedure	Function	Procedure
Ws_Load_Change (on page 985)	Procedure	Procedure	Function	Procedure
Ws_Maintain_Indexes (on page 990)	Procedure	Procedure	Function	Procedure
Ws_Version_Clear (on page 996)	Procedure	Procedure	Function	Procedure
WsParameterRead (see "WsParameterReadF" on page 1004)	Procedure	Procedure	Function	Procedure
WsParameterReadF (on page 1004)	Procedure	Procedure	Function	Procedure
WsParameterReadG (on page 1006)	Procedure	Procedure	Function	Procedure
WsParameterReadJ (on page 1010)	Procedure	Procedure	Function	Procedure
WsParameterWrite (on page 1013)	Procedure	Procedure	Function	Procedure
WsWrkAudit (on page 1016)	Procedure	Procedure	Function	Procedure
WsWrkAuditBulk (on page 1022)	Procedure	Procedure	Function	Procedure
WsWrkError (on page 1028)	Procedure	Procedure	Function	Procedure
WsWrkErrorBulk (on page 1034)	Procedure	Procedure	Function	Procedure
WsWrkTask (on page 1040)	Procedure	Procedure	Function	Procedure

Callable Routines Names Qualifier

The QAD Data Warehouse Designer 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 QAD Data Warehouse Designer metadata repository. All Data Warehouse Designer-generated procedures in a Teradata or DB2 repository that invoke a QAD Data Warehouse Designer Callable Routine do so by qualifying the routine name with **[METABASE]** e.g.

[METABASE].routine_name. When Data Warehouse Designer 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 Data Warehouse Designer-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 Data Warehouse Designer 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 QAD Data Warehouse Designer Callable Routine interactively or from outside a Data Warehouse Designer-compiled procedure/function then the actual owner/schema must be specified because Data Warehouse Designer won't have the chance to replace the "token".

Callable Routines Common Input

The following input parameters are common to most of the QAD Data Warehouse Designer Callable Routines, which are primarily used for integration with the QAD Data Warehouse Designer Scheduler.

Input	Parameter Name	Description
Job Instance Identifier	p_sequence	Unique identifier of the running job (i.e. the running instance of a held or scheduled job) that executed the routine. When invoked from a QAD Data Warehouse Designer Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the QAD Data Warehouse Designer Scheduler, then any integer value can be used.
Job Name	p_job_name	Name of the running job that executed the routine. When invoked from a QAD Data Warehouse Designer Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the QAD Data Warehouse Designer Scheduler, then any name can be used.
Task Name	p_task_name	Name of the running task (of a running job) that executed the routine. When invoked from a QAD Data Warehouse Designer Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the QAD Data Warehouse Designer Scheduler, then any name can be used.

Input	Parameter Name	Description
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 QAD Data Warehouse Designer Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the QAD Data Warehouse Designer Scheduler, it is recommended to use 0 (zero).

Note: Typically, the **parameter names** of the QAD Data Warehouse Designer 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 QAD Data Warehouse Designer 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 Data Warehouse Designer-generated procedure). However, it can be useful to execute a Callable Routine using other methods such as the following.

Alternative Invocation Methods

Invocation using Named Parameters

SQL Server, Oracle, and DB2 support the passing of parameter arguments by **name** as an alternative to passing them by position. Typically, for most RDBMS when one parameter argument is supplied by name then all subsequent parameter arguments must be supplied in the same manner. The order of the parameter names doesn't matter but for clarity the following alternative examples pass the parameter names in the same order that they are declared in the routine. When parameter arguments are not passed by name the values must be supplied in the identical order (left to right) as the parameters are declared in the routine.

SQL Server example of invoking the `Ws_Connect_Replace` PROCEDURE using named parameters:

```
DECLARE @v_return_code      varchar(1)
DECLARE @v_return_msg       varchar(256)
DECLARE @v_result_num       integer
EXEC Ws_Connect_Replace
    @p_sequence      = 0
  , @p_job_name      = 'Test Job Name'
  , @p_task_name     = 'Test Task Name'
  , @p_job_id        = 0
  , @p_task_id       = 0
  , @p_action        = 'REPLACE'
  , @p_source        = 'Connection1'
  , @p_target        = 'Connection2'
  , @p_return_code   = @v_return_code OUTPUT
  , @p_return_msg    = @v_return_msg  OUTPUT
  , @p_result        = @v_result_num  OUTPUT
```

Oracle example of invoking the `Ws_Connect_Replace` FUNCTION using named parameters via an anonymous PL/SQL block:

```
DECLARE
    v_return_code      varchar2(1);
    v_return_msg       varchar2(256);
    v_result_num       integer;
BEGIN
    v_result_num := Ws_Connect_Replace
( p_sequence      => 0
  , p_job_name     => 'Test Job Name'
  , p_task_name    => 'Test Task Name'
  , p_job_id       => 0
  , p_task_id      => 0
  , p_action       => 'REPLACE'
  , p_source       => 'Connection1'
  , p_target       => 'Connection2'
  , p_return_code  => v_return_code
  , p_return_msg   => v_return_msg
  );
END;
/
```

DB2 example of invoking the `Ws_Connect_Replace` PROCEDURE using named parameters via another procedure:

```
-- The owner/schema is DSSDEMO.
DROP PROCEDURE DSSDEMO.invoke_DB2_Ws_Connect_Replace;
CREATE PROCEDURE DSSDEMO.invoke_DB2_Ws_Connect_Replace
    LANGUAGE SQL
```

```

BEGIN
  DECLARE v_return_code      varchar(1);
  DECLARE v_return_msg      varchar(256);
  DECLARE v_result_num      integer;
  CALL DSSDEMO.Ws_Connect_Replace
  ( p_sequence      => 0
  , p_job_name      => 'Test Job Name'
  , p_task_name     => 'Test Task Name'
  , p_job_id        => 0
  , p_task_id       => 0
  , p_action        => 'REPLACE'
  , p_source        => 'Connection1'
  , p_target        => 'Connection2'
  , p_return_code   => v_return_code
  , p_return_msg    => v_return_msg
  , p_result        => v_result_num
  );
END
/
CALL DSSDEMO.invoke_DB2_Ws_Connect_Replace();

```

Invocation via ODBC

The following examples illustrate how to invoke a QAD Data Warehouse Designer Callable Routine via an ODBC connection (using a tool such as QAD SQL Admin), which uses both input and output parameters. The examples work for a SQL Server Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler.

Invocation via the Command-Line

Each QAD Data Warehouse Designer Callable Routine can also be invoked from the command-line using an RDBMS-specific tool such as:

RDBMS	SQL Command-Line-Interface (CLI) Tool
SQL Server	SQL Server sqlcmd Utility i.e. sqlcmd.exe.
Teradata	Basic Teradata Query (BTEQ) i.e. bteq.
Oracle	Oracle SQL*Plus i.e. sqlplus.
DB2	DB2 Command Line Processor (CLP) i.e. db2 or db2cmd.

These tools can be used to invoke a QAD Data Warehouse Designer 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 QAD Data Warehouse Designer.

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.

SQL Server

SQL Server 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
@p_option	VARCHAR(12)	IN

Parameter Name	Datatype	Mode
@p_result	VARCHAR(256)	OUT

SQL Server Examples: Ws_Api_Glossary

```

DECLARE @v_result_txt          varchar(256)
EXEC Ws_Api_Glossary
    'Data Warehouse'
, 'Overview'
, 'A repository of business information'
, 'ADD'
, @v_result_txt OUTPUT

```

Teradata

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
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
);

```

Oracle

Oracle Parameters: Ws_Api_Glossary

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
v_object_name	VARCHAR2	IN
v_term	VARCHAR2	IN
v_comments	VARCHAR2	IN
v_option	VARCHAR2	IN
FUNCTION Return Value	VARCHAR2	OUT-Function

Oracle Examples: Ws_Api_Glossary

```
v_result_txt          varchar2(256);  
v_result_txt := Ws_Api_Glossary  
( 'Data Warehouse'  
  , 'Overview'  
  , 'A repository of business information'  
  , 'ADD'  
);
```

DB2

DB2 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
p_option	VARCHAR(64)	IN
p_result	VARCHAR(256)	OUT

DB2 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 Callable Routines Common Input (on page 907).
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.

SQL Server

SQL Server 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

SQL Server Examples: Ws_Connect_Replace

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC Ws_Connect_Replace
    @p_sequence, @p_job_name, @p_task_name, @p_job_id, @p_task_id
, 'REPLACE', 'Connection1', 'Connection2'
, @v_return_code OUTPUT
, @v_return_msg OUTPUT
, @v_result_num OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: Ws_Connect_Replace

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_action	VARCHAR2	IN
p_source	VARCHAR2	IN
p_target	VARCHAR2	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Connect_Replace

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated
Procedure.
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_return_code        varchar2(1);
v_return_msg        varchar2(256);
v_result_num := 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
);
```

DB2

DB2 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

DB2 Examples: Ws_Connect_Replace

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Audit Log for the aborted job and the QAD Data Warehouse Designer Task Log for each aborted task.

Output

Output	Description
Result Number [ORACLE only]	Output Result Number [ORACLE only]: 1 Success. -2 Error. e.g. Due to invalid job name or job not running. -3 Fatal/Unexpected Error.

SQL Server

SQL Server Parameters: Ws_Job_Abort

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
@p_job_name	VARCHAR(64)	IN

Parameter Name	Datatype	Mode
@p_job_sequence	INTEGER	IN
@p_job_msg	VARCHAR(256)	IN

SQL Server Examples: Ws_Job_Abort

```
EXEC Ws_Job_Abort
    'Daily Run'
    , 1234
    , 'Job aborted via manual execution of Ws_Job_Abort.'
```

Teradata

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.'
);
```

Oracle

Oracle Parameters: Ws_Job_Abort

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
v_job_name	VARCHAR2	IN

Parameter Name	Datatype	Mode
v_job_sequence	NUMBER	IN
v_job_msg	VARCHAR2	IN
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Job_Abort

```

v_result_num          number;
v_result_num := Ws_Job_Abort
( 'Daily Run'
, 1234
, 'Job aborted via manual execution of Ws_Job_Abort.'
);

```

DB2

DB2 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

DB2 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 Data Warehouse Designer Scheduler and/or Data Warehouse Designer 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 Callable Routines Common Input (on page 907).
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.

SQL Server

SQL Server 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

SQL Server Examples: Ws_Job_Clear_Archive

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC Ws_Job_Clear_Archive
    @p_sequence, @p_job_name, @p_task_name, @p_job_id, @p_task_id
, 90, 'Daily Run', ''
, @v_return_code OUTPUT
, @v_return_msg OUTPUT
, @v_result_num OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: Ws_Job_Clear_Archive

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_day_count	NUMBER	IN
p_job	VARCHAR2	IN
p_options	VARCHAR2	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Job_Clear_Archive

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_return_code       varchar2(1);
v_return_msg        varchar2(256);
v_result_num := 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
);
```

DB2

DB2 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

DB2 Examples: Ws_Job_Clear_Archive

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 QAD Data Warehouse Designer Scheduler and the "Logs Retained" property of a job.

Input

Input	Description
Common Input	Includes all 5 inputs of the Callable Routines Common Input (on page 907).
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.

SQL Server

SQL Server 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_keep_count	INTEGER	IN
@p_return_code	VARCHAR(1)	OUT
@p_return_msg	VARCHAR(256)	OUT
@p_result	INTEGER	OUT

SQL Server Examples: Ws_Job_Clear_Logs

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC Ws_Job_Clear_Logs
    @p_sequence, @p_job_name, @p_task_name, @p_job_id, @p_task_id
, 'Daily Run', 10
, @v_return_code OUTPUT
, @v_return_msg  OUTPUT
, @v_result_num  OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```


Oracle

Oracle Parameters: Ws_Job_Clear_Logs

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_job_to_clean	VARCHAR2	IN
p_log_keep	NUMBER	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Job_Clear_Logs

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_return_code       varchar2(1);
v_return_msg        varchar2(256);
v_result_num := 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
);
```

DB2

DB2 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_keep_count	INTEGER	IN
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

DB2 Examples: Ws_Job_Clear_Logs

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 Callable Routines Common Input (on page 907).
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.

SQL Server

SQL Server 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

SQL Server Examples: Ws_Job_Clear_Logs_By_Date

-- The p_VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC 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 OUTPUT
, @v_return_msg  OUTPUT
, @v_result_num  OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: Ws_Job_Clear_Logs_By_Date

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_job_to_clean	VARCHAR2	IN
p_day_count	NUMBER	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
p_result	NUMBER	OUT-Function

Oracle Examples: Ws_Job_Clear_Logs_By_Date

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_return_code       varchar2(1);
v_return_msg        varchar2(256);
v_result_num := 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
);
```

DB2

DB2 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

DB2 Examples: Ws_Job_Clear_Logs_By_Date

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 Callable Routines Common Input (on page 907).
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 QAD Data Warehouse Designer 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.
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.

Input	Description
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.
Attributes 1-8	Optional attributes to associate with the new job. Such attributes can be retrieved by tasks while they are running within the job via the WsParameterReadJ callable routine.

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.

SQL Server

SQL Server Parameters: Ws_Job_Create

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

Parameter Name	Datatype	Mode
@p_template_job	VARCHAR(64)	IN
@p_new_job	VARCHAR(64)	IN
@p_description	VARCHAR(256)	IN
@p_state	VARCHAR(64)	IN
@p_threads	INTEGER	IN
@p_scheduler	VARCHAR(64)	IN
@p_logs	INTEGER	IN
@p_okay	VARCHAR(256)	IN
@p_fail	VARCHAR(256)	IN
@p_att1	VARCHAR(64)	IN
@p_att2	VARCHAR(64)	IN
@p_att3	VARCHAR(64)	IN
@p_att4	VARCHAR(64)	IN
@p_att5	VARCHAR(64)	IN
@p_att6	VARCHAR(64)	IN
@p_att7	VARCHAR(64)	IN
@p_att8	VARCHAR(64)	IN
@p_return_code	VARCHAR(1)	OUT
@p_return_msg	VARCHAR(256)	OUT
@p_result	INTEGER	OUT

SQL Server Examples: Ws_Job_Create

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC 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
, @v_return_code OUTPUT
, @v_return_msg OUTPUT
, @v_result_num OUTPUT
```

Teradata

Teradata Parameters: Ws_Job_Create

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_threads	INTEGER	IN
p_scheduler	VARCHAR(8)	IN
p_logs	INTEGER	IN
p_okay	VARCHAR(256)	IN

Parameter Name	Datatype	Mode
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 Data Warehouse Designer-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
, v_return_code
, v_return_msg
, v_result_num
);

```

Oracle

Oracle Parameters: Ws_Job_Create

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_template_job	VARCHAR2	IN
p_new_job	VARCHAR2	IN
p_description	VARCHAR2	IN
p_state	VARCHAR2	IN
p_threads	NUMBER	IN
p_scheduler	VARCHAR2	IN
p_logs	NUMBER	IN
p_okay	VARCHAR2	IN
p_fail	VARCHAR2	IN
p_att1	VARCHAR2	IN
p_att2	VARCHAR2	IN
p_att3	VARCHAR2	IN
p_att4	VARCHAR2	IN
p_att5	VARCHAR2	IN
p_att6	VARCHAR2	IN
p_att7	VARCHAR2	IN
p_att8	VARCHAR2	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Job_Create

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated
Procedure.
p_sequence                number;
p_job_name                varchar2(256);
p_task_name               varchar2(256);
p_job_id                  number;
p_task_id                 number;
p_return_msg              varchar2(256);
p_status                  number;
v_result_num              number;
v_return_code              varchar2(1);
v_return_msg              varchar2(256);
v_result_num := 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
, v_return_code
, v_return_msg
);
```

DB2

DB2 Parameters: Ws_Job_Create

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_threads	INTEGER	IN
p_scheduler	VARCHAR(8)	IN
p_logs	INTEGER	IN

Parameter Name	Datatype	Mode
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

DB2 Examples: Ws_Job_Create

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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
, 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 Callable Routines Common Input (on page 907).
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 QAD Data Warehouse Designer 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
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.
Attributes 1-8	Optional attributes to associate with the new job. Such attributes can be retrieved by tasks while they are running within the job via the WsParameterReadJ callable routine.

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.

SQL Server

SQL Server Parameters: Ws_Job_CreateWait

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
@p_sequence	INTEGER	IN

Parameter Name	Datatype	Mode
@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(64)	IN
@p_release_time	DATETIME	IN
@p_threads	INTEGER	IN
@p_scheduler	VARCHAR(64)	IN
@p_logs	INTEGER	IN
@p_okay	VARCHAR(256)	IN
@p_fail	VARCHAR(256)	IN
@p_att1	VARCHAR(64)	IN
@p_att2	VARCHAR(64)	IN
@p_att3	VARCHAR(64)	IN
@p_att4	VARCHAR(64)	IN
@p_att5	VARCHAR(64)	IN
@p_att6	VARCHAR(64)	IN
@p_att7	VARCHAR(64)	IN
@p_att8	VARCHAR(64)	IN
@p_return_code	VARCHAR(1)	OUT
@p_return_msg	VARCHAR(256)	OUT
@p_result	INTEGER	OUT

SQL Server Examples: Ws_Job_CreateWait

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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_run_date            datetime
SET @v_run_date = DATEADD(MONTH, 1, GETDATE())
EXEC Ws_Job_CreateWait
    @p_sequence, @p_job_name, @p_task_name, @p_job_id, @p_task_id
, 'Daily Run', 'New Daily Run2', 'This is the New Daily Run job.', 'ONCE'
, @v_run_date
, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL
, @v_return_code OUTPUT
, @v_return_msg OUTPUT
, @v_result_num OUTPUT
```

Teradata

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

Parameter Name	Datatype	Mode
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 Data Warehouse Designer-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
, v_return_code
, v_return_msg
, v_result_num
);

```

Oracle

Oracle Parameters: Ws_Job_CreateWait

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_template_job	VARCHAR2	IN
p_new_job	VARCHAR2	IN
p_description	VARCHAR2	IN
p_state	VARCHAR2	IN
p_release_time	DATE	IN
p_threads	NUMBER	IN
p_scheduler	VARCHAR2	IN
p_logs	NUMBER	IN
p_okay	VARCHAR2	IN
p_fail	VARCHAR2	IN
p_att1	VARCHAR2	IN
p_att2	VARCHAR2	IN
p_att3	VARCHAR2	IN
p_att4	VARCHAR2	IN
p_att5	VARCHAR2	IN
p_att6	VARCHAR2	IN
p_att7	VARCHAR2	IN
p_att8	VARCHAR2	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Job_CreateWait

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_return_code       varchar2(1);
v_return_msg        varchar2(256);
v_result_num := 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'
, ADD_MONTHS(SYSDATE, 1)
, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL
, v_return_code
, v_return_msg
);
```

DB2

DB2 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

Parameter Name	Datatype	Mode
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

DB2 Examples: Ws_Job_CreateWait

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 + 1 MONTH)
, 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 Callable Routines Common Input (on page 907).
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.
Return Message	Output message indicating the action applied or the reason for no action.

Output	Description
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.

SQL Server

SQL Server 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(64)	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

SQL Server Examples: Ws_Job_Dependency

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC 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 OUTPUT
, @v_return_msg  OUTPUT
, @v_result_num  OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: Ws_Job_Dependency

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_action	VARCHAR2	IN
p_parent	VARCHAR2	IN
p_child	VARCHAR2	IN
p_required	VARCHAR2	IN
p_look_back	NUMBER	IN
p_max_wait	NUMBER	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT

Parameter Name	Datatype	Mode
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Job_Dependency

-- The p_VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```

p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_return_code        varchar2(1);
v_return_msg        varchar2(256);
v_result_num := 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
);

```

DB2

DB2 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

Parameter Name	Datatype	Mode
p_return_code	VARCHAR(1)	OUT
p_return_msg	VARCHAR(256)	OUT
p_result	INTEGER	OUT

DB2 Examples: Ws_Job_Dependency

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 QAD Data Warehouse Designer Scheduler).

Input

Input	Description
Common Input	Includes all 5 inputs of the Callable Routines Common Input (on page 907).
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 QAD Data Warehouse Designer 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.

SQL Server

SQL Server 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

SQL Server Examples: Ws_Job_Release

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC Ws_Job_Release
    @p_sequence, @p_job_name, @p_task_name, @p_job_id, @p_task_id
, 'Daily Run'
, @v_return_code OUTPUT
, @v_return_msg OUTPUT
, @v_result_num OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```


Oracle

Oracle Parameters: Ws_Job_Release

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_release_job	VARCHAR2	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Job_Release

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name           varchar2(256);
p_task_name          varchar2(256);
p_job_id             number;
p_task_id            number;
p_return_msg         varchar2(256);
p_status             number;
v_result_num         number;
v_return_code        varchar2(1);
v_return_msg         varchar2(256);
v_result_num := 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
);
```

DB2

DB2 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

DB2 Examples: Ws_Job_Release

-- The p_VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 Callable Routines Common Input (on page 907).
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 QAD Data Warehouse Designer 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.
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.

SQL Server

SQL Server 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

SQL Server Examples: Ws_Job_Restart

-- The p_VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC Ws_Job_Restart
    @p_sequence, @p_job_name, @p_task_name, @p_job_id, @p_task_id
, 'Daily Run'
, @v_return_code OUTPUT
, @v_return_msg  OUTPUT
, @v_result_num  OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: Ws_Job_Restart

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_restart_job	VARCHAR2	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Job_Restart

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name           varchar2(256);
p_task_name          varchar2(256);
p_job_id             number;
p_task_id            number;
p_return_msg         varchar2(256);
p_status             number;
v_result_num         number;
v_return_code        varchar2(1);
v_return_msg         varchar2(256);
v_result_num := 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
);
```

DB2

DB2 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

DB2 Examples: Ws_Job_Restart

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 Callable Routines Common Input (on page 907).
Schedule Job Name	The name of the job to be scheduled. The specified name must exactly match the job name as displayed by the QAD Data Warehouse Designer 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.

SQL Server

SQL Server 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	DATETIME	IN
@p_return_code	VARCHAR(1)	OUT
@p_return_msg	VARCHAR(256)	OUT
@p_result	INTEGER	OUT

SQL Server Examples: Ws_Job_Schedule

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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_run_date           datetime
SET @v_run_date = GETDATE() + 1
EXEC Ws_Job_Schedule
    @p_sequence, @p_job_name, @p_task_name, @p_job_id, @p_task_id
, 'Daily Run', @v_run_date
, @v_return_code OUTPUT
, @v_return_msg OUTPUT
, @v_result_num OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: Ws_Job_Schedule

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_release_job	VARCHAR2	IN
p_release_time	DATE	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Job_Schedule

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_return_code       varchar2(1);
v_return_msg        varchar2(256);
v_result_num := Ws_Job_Schedule
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'Daily Run', SYSDATE+1
, v_return_code
, v_return_msg
);
```

DB2

DB2 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

DB2 Examples: Ws_Job_Schedule

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 + 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 QAD Data Warehouse Designer 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 907).
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 QAD Data Warehouse Designer 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.

Output	Description
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.
Simplified Job Status Code	Simplified Job Status Code: N Not Running. R Running. F Failed. C Completed. O Job did not start in the specified time.
Standard Job Status Code	Standard Job Status Code: H On Hold. The job is on hold. A held job can be edited and/or started. 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). B Blocked. The job is blocked because a previous instance of the same job is still running. 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. R Running. The job is currently running. F Failed. The job failed due to an error. C Completed. The job completed successfully (but it may have warnings). A completed job cannot be restarted. G Failed & Aborted. The job failed and was subsequently aborted. An aborted job cannot be restarted. E Error.

Output	Description
Enhanced Job Status Number	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. 1 On Hold. The job is on hold. A held job can be edited and/or started. 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). 3 Blocked. The job is blocked because a previous instance of the same job is still running. 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. 5 Running. The job is currently running and no tasks have failed or produced warnings. 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. 7 Running with Warnings. The job is currently running and some tasks have produced warnings. 8 Failed. The job failed due to an error. 9 Completed. The job completed without warnings. A completed job cannot be restarted. 10 Completed with Warnings. The job completed with warnings. A completed job cannot be restarted. 11 Failed & Aborted. The job failed and it was subsequently aborted. An aborted job cannot be restarted. 12 Error.

SQL Server

SQL Server 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	DATETIME	IN

Parameter Name	Datatype	Mode
@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

SQL Server Examples: Ws_Job_Status

-- The p_VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC 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 OUTPUT
, @v_return_msg  OUTPUT
, @v_result_num  OUTPUT
, @v_job_status_simple OUTPUT
, @v_job_status_standard OUTPUT
, @v_job_status_enhanced OUTPUT

```

Teradata

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

Parameter Name	Datatype	Mode
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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: Ws_Job_Status

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_check_sequence	NUMBER	IN
p_check_job	VARCHAR2	IN
p_started_in_last_mi	NUMBER	IN
p_started_after_dt	DATE	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT

Parameter Name	Datatype	Mode
p_job_status_simple	VARCHAR2	OUT
p_job_status_standard	VARCHAR2	OUT
p_job_status_enhanced	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Job_Status

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```

p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_return_code       varchar2(1);
v_return_msg        varchar2(256);
v_job_status_simple varchar2(1);
v_job_status_standard varchar2(1);
v_job_status_enhanced varchar2(2);
v_result_num := 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_job_status_simple, v_job_status_standard, v_job_status_enhanced
);

```

DB2

DB2 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

Parameter Name	Datatype	Mode
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

DB2 Examples: Ws_Job_Status

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 Callable Routines Common Input (on page 907).
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.

SQL Server

SQL Server 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(64)	IN
@p_return_code	VARCHAR(1)	OUT
@p_return_msg	VARCHAR(256)	OUT
@p_result	INTEGER	OUT

SQL Server Examples: Ws_Load_Change

-- The p_VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC Ws_Load_Change
    @p_sequence, @p_job_name, @p_task_name, @p_job_id, @p_task_id
, 'CONNECTION', 'load_customer', 'Connection2'
, @v_return_code OUTPUT
, @v_return_msg  OUTPUT
, @v_result_num  OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: Ws_Load_Change

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_action	VARCHAR2	IN
p_table	VARCHAR2	IN
p_new_value	VARCHAR2	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Load_Change

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated
Procedure.
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_return_code        varchar2(1);
v_return_msg        varchar2(256);
v_result_num := 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
);
```


DB2

DB2 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

DB2 Examples: Ws_Load_Change

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 QAD Data Warehouse Designer metadata.

Description

Drops and/or builds indexes for a specified table or a specified index. Only indexes that are defined in the QAD Data Warehouse Designer metadata are supported. Typically, this routine is used by a QAD Data Warehouse Designer Scheduler and Data Warehouse Designer-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 Callable Routines Common Input (on page 907).
Table Name	Table Name to process the relevant indexes of. Note: The Table Name is ignored when the optional Index Name is specified.
Optional Build Arguments [ORACLE only]	Optional arguments related to building indexes [ORACLE only] : TABLESPACE=XYZ will override the default tablespace name and use the specified tablespace name. e.g. 'TABLESPACE=non_default_tablespace'. Typically the TABLESPACE argument is used to build indexes in a partition-specific tablespace.
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 QAD Data Warehouse Designer Callable routines but it does output a Result Number.

SQL Server

SQL Server 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(64)	IN
@p_index_name	VARCHAR(64)	IN
@p_option	VARCHAR(64)	IN
@p_result	INTEGER	OUT

SQL Server Examples: Ws_Maintain_Indexes

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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
EXEC 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 OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: Ws_Maintain_Indexes

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_table_name	VARCHAR2	IN
p_parameter	VARCHAR2	IN
p_index_name	VARCHAR2	IN
p_option	VARCHAR2	IN
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Maintain_Indexes

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_result_num := Ws_Maintain_Indexes
( p_sequence, p_job_name, p_task_name, p_job_id, p_task_id
, 'load_customer', NULL, NULL, 'DROP ALL'
);
```

DB2

DB2 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

DB2 Examples: Ws_Maintain_Indexes

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 Callable Routines Common Input (on page 907).
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.
Result Number	Output Result Number: 1 Success. -2 Error. -3 Fatal/Unexpected Error.

SQL Server

SQL Server 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

SQL Server Examples: Ws_Version_Clear

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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)
EXEC Ws_Version_Clear
    @p_sequence, @p_job_name, @p_task_name, @p_job_id, @p_task_id
, 90, 5, NULL
, @v_return_code OUTPUT
, @v_return_msg  OUTPUT
, @v_result_num  OUTPUT
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: Ws_Version_Clear

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_sequence	NUMBER	IN
p_job_name	VARCHAR2	IN
p_task_name	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
p_day_count	NUMBER	IN
p_keep_count	NUMBER	IN
p_options	VARCHAR2	IN
p_return_code	VARCHAR2	OUT
p_return_msg	VARCHAR2	OUT
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: Ws_Version_Clear

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated
Procedure.
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_return_code        varchar2(1);
v_return_msg        varchar2(256);
v_result_num := 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
);
```

DB2

DB2 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

DB2 Examples: Ws_Version_Clear

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 QAD Data Warehouse Designer 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 Data Warehouse Designer **Tools/Parameters** menu item), which is external to the procedure.

Input

Input	Description
Parameter Name	The case-sensitive name of the QAD Data Warehouse Designer 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 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.

SQL Server

SQL Server Parameters: WsParameterRead

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
@p_parameter	VARCHAR(64)	IN

Parameter Name	Datatype	Mode
@p_value	VARCHAR(2000)	OUT
@p_comment	VARCHAR(256)	OUT

Note: There is also a corresponding WsParameterReadF FUNCTION available for SQL Server.

SQL Server Examples: WsParameterRead

```

DECLARE @v_current_date          varchar(4000) -- Same length as
DSS_PARAMETER.dss_parameter_value.
DECLARE @v_comment               varchar(256)
EXEC WsParameterRead 'CURRENT_DATE', @v_current_date OUTPUT, @v_comment OUTPUT;

```

Teradata

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);

```

Oracle

Oracle Parameters: WsParameterRead

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
v_parameter	VARCHAR2	IN
FUNCTION Return Value	VARCHAR2	OUT-Function

Oracle Examples: WsParameterRead

```
v_current_date          varchar2(4000);    -- Same length as
DSS_PARAMETER.dss_parameter_value.
v_current_date := WsParameterRead('CURRENT_DATE'); -- NOTE: Does NOT support comment
output.
```

DB2

DB2 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

DB2 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 QAD Data Warehouse Designer 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 Data Warehouse Designer **Tools/Parameters** menu item), which is external to the procedure.

Input

Input	Description
Parameter Name	The name of the QAD Data Warehouse Designer 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 .

SQL Server

SQL Server Parameters: WsParameterReadF

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
@p_parameter	VARCHAR(64)	IN
FUNCTION Return Value	VARCHAR(4000)	OUT-Function

Note: There is also a corresponding WsParameterRead PROCEDURE available for SQL Server.

SQL Server Examples: WsParameterReadF

```
DECLARE @v_current_date          varchar(4000)
SELECT @v_current_date = dbo.WsParameterReadF('CURRENT_DATE')
```

WsParameterReadG

Synopsis

Returns the value of a "global" QAD Data Warehouse Designer metadata Parameter that relates to a load table.

Description

Returns the value of an internal parameter that is defined and populated by QAD Data Warehouse Designer, 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 QAD Data Warehouse Designer Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the QAD Data Warehouse Designer 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 QAD Data Warehouse Designer Scheduler the routine will be passed the parameter argument. When invoked manually or externally to the QAD Data Warehouse Designer Scheduler, it should be the object key .

Output

Output	Description
Result Table Name	The requested Load Table Name or Source Table Name.

SQL Server

SQL Server Parameters: WsParameterReadG

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
@p_parameter	VARCHAR(64)	IN
@p_job_id	INTEGER	IN
@p_task_id	INTEGER	IN
FUNCTION Return Value	VARCHAR(4000)	OUT-Function

SQL Server Examples: WsParameterReadG

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.  
DECLARE @p_job_id          integer  
DECLARE @p_task_id        integer  
DECLARE @v_source_name    varchar(64)  
SELECT @v_source_name = dbo.WsParameterReadG('$$SOURCE_TABLE', @p_job_id,  
@p_task_id) -- Needs dbo.
```

Teradata

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 Data Warehouse Designer-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);
```

Oracle

Oracle Parameters: WsParameterReadG

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_parameter	VARCHAR2	IN
p_job_id	NUMBER	IN
p_task_id	NUMBER	IN
FUNCTION Return Value	VARCHAR2	OUT-Function

Oracle Examples: WsParameterReadG

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.  
p_job_id            number;  
p_task_id           number;  
v_source_name       varchar2(64);  
v_source_name := WsParameterReadG('$$SOURCE_TABLE', p_job_id, p_task_id);
```

DB2

DB2 Parameters: WsParameterReadG

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_parameter	VARCHAR(64)	IN

Parameter Name	Datatype	Mode
p_job_id	INTEGER	IN
p_task_id	INTEGER	IN
p_value	VARCHAR(2000)	OUT

DB2 Examples: WsParameterReadG

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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);

```

WsParameterReadJ

Synopsis

Returns information about the job that initiated a running procedure.

Description

Returns the name, description, or an attribute of the job that executed the current procedure i.e. the job of the running procedure that called WsParameterReadJ.

Input

Input	Description
Job Property	The name of the job property to be retrieved: NAME returns the Job Name. DESCRIPTION returns the Job Description. ATTRIBUTE1...ATTRIBUTE8 returns Job Attribute 1, 2, 3, 4, 5, 6, 7, or 8. Note: Job Attributes are not available for all RDBMS.
Job Identifier	Unique identifier of the held or scheduled job that the running job is a specific instance of. When invoked from a QAD Data Warehouse Designer Scheduler the routine will be passed the parameter argument. Without a valid job identifier nothing is returned.

Output

Output	Description
Job Property Value	The value of the requested Job Property.

SQL Server

SQL Server Parameters: WsParameterReadJ

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
@p_parameter	VARCHAR(64)	IN
@p_job_id	INTEGER	IN

Parameter Name	Datatype	Mode
FUNCTION Return Value	VARCHAR(4000)	OUT-Function

SQL Server Examples: WsParameterReadJ

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.
DECLARE @p_job_id                integer
DECLARE @v_job_name              varchar(256)
SET @p_job_id = 3
SELECT @v_job_name = dbo.WsParameterReadJ('NAME', @p_job_id) -- Needs dbo.
```

Teradata

Teradata Parameters: WsParameterReadJ

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_parameter	VARCHAR(256)	IN
p_job_id	INTEGER	IN
p_result	VARCHAR(2000)	OUT

Teradata Examples: WsParameterReadJ

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.
DECLARE p_job_id                integer;
DECLARE v_job_name              varchar(256);
CALL [METABASE].WsParameterReadJ('NAME', p_job_id, v_job_name);
```

Oracle

Oracle Parameters: WsParameterReadJ

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
p_parameter	VARCHAR2	IN
p_job_id	NUMBER	IN
FUNCTION Return Value	VARCHAR2	OUT-Function

Oracle Examples: WsParameterReadJ

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_job_id          number;
v_job_name        varchar2(256);
v_job_name := WsParameterReadJ('NAME', p_job_id);
```

DB2

DB2 Parameters: WsParameterReadJ

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
p_parameter	VARCHAR(256)	IN
p_job_id	INTEGER	IN
p_result	VARCHAR(2000)	OUT

DB2 Examples: WsParameterReadJ

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
DECLARE p_job_id          integer;
DECLARE v_job_name        varchar(256);
CALL [METABASE].WsParameterReadJ('NAME', p_job_id, v_job_name);
```


WsParameterWrite

Synopsis

Updates the value and comment of a QAD Data Warehouse Designer 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 Data Warehouse Designer **Tools/Parameters** menu item), which is external to the procedure.

Input

Input	Description
Parameter Name	The case-sensitive name of the QAD Data Warehouse Designer 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.

Output

Output	Description
Result Number [SQL SERVER & ORACLE only]	Output Result Number [SQL SERVER & ORACLE only]: 1 Metadata Parameter Updated. 2 Metadata Parameter Inserted. -3 Fatal/Unexpected Error.

SQL Server

SQL Server 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

SQL Server Examples: WsParameterWrite

```
-- Example 1: Result Number (procedure return code) NOT captured.
EXEC WsParameterWrite 'LAST_INVOICE_ID', '123456', 'The last invoice id loaded'

-- Example 2: Result Number (procedure return code) captured.
DECLARE @v_result_num          integer
EXEC @v_result_num = WsParameterWrite 'LAST_INVOICE_ID', '123456', 'The last invoice
id loaded'
```

Teradata

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');
```

Oracle

Oracle Parameters: WsParameterWrite

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
v_parameter	VARCHAR2	IN
v_value	VARCHAR2	IN
v_comment	VARCHAR2	IN
FUNCTION Return Value	VARCHAR2	OUT-Function

Oracle Examples: WsParameterWrite

```
v_result_num          number;  
v_result_num := WsParameterWrite('LAST_INVOICE_ID', '123456', 'The last invoice id  
loaded');
```

DB2

DB2 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

DB2 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 Data Warehouse Designer-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 Callable Routines Common Input (on page 907). 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 QAD Data Warehouse Designer 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.

SQL Server

SQL Server 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(256)	IN
@p_db_code	VARCHAR(10)	IN
@p_db_msg	VARCHAR(256)	IN
@p_task_key	INTEGER	IN
@p_job_key	INTEGER	IN

SQL Server Examples: WsWrkAudit

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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
EXEC @v_result_num = WsWrkAudit
    'I', @p_job_name, @p_task_name, @p_sequence
, 'The task has started.'
, NULL
, NULL
, @p_task_id
, @p_job_id
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: WsWrkAudit

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
v_status_code	VARCHAR2	IN
v_job_name	VARCHAR2	IN
v_task_name	VARCHAR2	IN
v_sequence	NUMBER	IN
v_message	VARCHAR2	IN
v_db_code	VARCHAR2	IN
v_db_msg	VARCHAR2	IN
v_task_key	NUMBER	IN
v_job_key	NUMBER	IN
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: WsWrkAudit

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_result_num := WsWrkAudit
( 'I', p_job_name, p_task_name, p_sequence
, 'The task has started.'
, NULL
, NULL
, p_task_id
, p_job_id
);
```

DB2

DB2 Parameters: WsWrkAudit

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
v_status_code	CHARACTER	IN
v_job_name	VARCHAR(64)	IN
v_task_name	VARCHAR(64)	IN
v_sequence	INTEGER	IN
v_message	VARCHAR(256)	IN
v_db_code	VARCHAR(10)	IN
v_db_msg	VARCHAR(256)	IN
v_task_key	INTEGER	IN
v_job_key	INTEGER	IN

DB2 Examples: WsWrkAudit

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 Data Warehouse Designer-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 Callable Routines Common Input (on page 907). 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 QAD Data Warehouse Designer 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	Note: Not provided for all RDBMS. Output Result Number: 1 Success. -2 Error -3 Fatal/Unexpected Error.

SQL Server

SQL Server Parameters: WsWrkAuditBulk

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(4000)	IN
@p_db_code	VARCHAR(10)	IN
@p_db_msg	VARCHAR(256)	IN
@p_task_key	INTEGER	IN
@p_job_key	INTEGER	IN

SQL Server Examples: WsWrkAuditBulk

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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
EXEC @v_result_num = WsWrkAuditBulk
    'I', @p_job_name, @p_task_name, @p_sequence
, 'Message1~Message2~Message3'
, NULL
, NULL
, @p_job_id
, @p_task_id
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: WsWrkAuditBulk

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
v_status_code	VARCHAR2	IN
v_job_name	VARCHAR2	IN
v_task_name	VARCHAR2	IN
v_sequence	NUMBER	IN
v_message	VARCHAR2	IN
v_db_code	VARCHAR2	IN
v_db_msg	VARCHAR2	IN
v_task_key	NUMBER	IN
v_job_key	NUMBER	IN
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: WsWrkAuditBulk

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_result_num := WsWrkAuditBulk
( 'I', p_job_name, p_task_name, p_sequence
, 'Message1~Message2~Message3'
, NULL
, NULL
, p_task_id
, p_job_id
);
```

DB2

DB2 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(32672)	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

DB2 Examples: WsWrkAuditBulk

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 Data Warehouse Designer-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 Callable Routines Common Input (on page 907). 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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer metadata.

Output

Output	Description
Result Number	Output Result Number: 1 Success. -3 Error.

SQL Server

SQL Server 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(256)	IN
@p_db_code	VARCHAR(10)	IN
@p_db_msg	VARCHAR(256)	IN
@p_task_key	INTEGER	IN
@p_job_key	INTEGER	IN
@p_msg_type	VARCHAR(10)	IN

SQL Server Examples: WsWrkError

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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
EXEC @v_result_num = 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
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: WsWrkError

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
v_status_code	VARCHAR2	IN
v_job_name	VARCHAR2	IN
v_task_name	VARCHAR2	IN
v_sequence	NUMBER	IN
v_message	VARCHAR2	IN
v_db_code	VARCHAR2	IN
v_db_msg	VARCHAR2	IN
v_task_key	NUMBER	IN
v_job_key	NUMBER	IN
v_msg_type	VARCHAR2	IN
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: WsWrkError

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_result_num := 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
);
```

DB2

DB2 Parameters: WsWrkError

Callable Routine Type: PROCEDURE.

Parameter Name	Datatype	Mode
v_status_code	CHARACTER	IN
v_job_name	VARCHAR(64)	IN
v_task_name	VARCHAR(64)	IN
v_sequence	INTEGER	IN
v_message	VARCHAR(256)	IN
v_db_code	VARCHAR(10)	IN
v_db_msg	VARCHAR(256)	IN
v_task_key	INTEGER	IN
v_job_key	INTEGER	IN
v_msg_type	VARCHAR(10)	IN

DB2 Examples: WsWrkError

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 Data Warehouse Designer-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 Callable Routines Common Input (on page 907). 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 QAD Data Warehouse Designer 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 QAD Data Warehouse Designer metadata.

Output

Output	Description
Result Number	Note: NOT provided for all RDBMS. Output Result Number: 1 Success. -2 Error. -3 Fatal/Unexpected Error.

SQL Server

SQL Server Parameters: WsWrkErrorBulk

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(4000)	IN
@p_db_code	VARCHAR(10)	IN
@p_db_msg	VARCHAR(256)	IN
@p_task_key	INTEGER	IN
@p_job_key	INTEGER	IN
@p_msg_type	VARCHAR(10)	IN

SQL Server Examples: WsWrkErrorBulk

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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
EXEC @v_result_num = WsWrkErrorBulk
    'I', @p_job_name, @p_task_name, @p_sequence
, 'Message1~Message2~Message3'
, NULL
, NULL
, @p_job_id
, @p_task_id
, NULL
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: WsWrkErrorBulk

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
v_status_code	VARCHAR2	IN
v_job_name	VARCHAR2	IN
v_task_name	VARCHAR2	IN
v_sequence	NUMBER	IN
v_message	VARCHAR2	IN
v_db_code	VARCHAR2	IN
v_db_msg	VARCHAR2	IN
v_task_key	NUMBER	IN
v_job_key	NUMBER	IN
v_msg_type	VARCHAR2	IN
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: WsWrkErrorBulk

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated Procedure.

```
p_sequence          number;
p_job_name          varchar2(256);
p_task_name         varchar2(256);
p_job_id            number;
p_task_id           number;
p_return_msg        varchar2(256);
p_status            number;
v_result_num        number;
v_result_num := WsWrkErrorBulk
( 'I', p_job_name, p_task_name, p_sequence
, 'Message1-Message2-Message3'
, NULL
, NULL
, p_task_id
, p_job_id
, NULL
);
```

DB2

DB2 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(32672)	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

DB2 Examples: WsWrkErrorBulk

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 QAD Data Warehouse Designer Scheduler.

Input

Input	Description
Common Input	Includes 3 inputs of the <i>Callable Routines Common Input</i> (on page 907).
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.

SQL Server

SQL Server 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

SQL Server Examples: WsWrkTask

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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
EXEC WsWrkTask
    @p_job_id, @p_task_id, @p_sequence
, 27, 30, 0, 0, 0, 0, 0
```

Teradata

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 Data Warehouse Designer-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
);
```

Oracle

Oracle Parameters: WsWrkTask

Callable Routine Type: FUNCTION.

Parameter Name	Datatype	Mode
v_job_key	NUMBER	IN
v_task_key	NUMBER	IN
v_sequence	NUMBER	IN
v_inserted	NUMBER	IN
v_updated	NUMBER	IN
v_replaced	NUMBER	IN
v_deleted	NUMBER	IN
v_discarded	NUMBER	IN
v_rejected	NUMBER	IN
v_errored	NUMBER	IN
FUNCTION Return Value	NUMBER	OUT-Function

Oracle Examples: WsWrkTask

```
-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-generated
Procedure.
p_sequence                number;
p_job_name                varchar2(256);
p_task_name               varchar2(256);
p_job_id                 number;
p_task_id                number;
p_return_msg              varchar2(256);
p_status                  number;
v_result_num              number;
v_insert_count            number;
v_update_count            number;
v_result_num := WsWrkTask
( p_job_id, p_task_id, p_sequence
, v_insert_count, v_update_count, 0, 0, 0, 0, 0
);
```

DB2

DB2 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

DB2 Examples: WsWrkTask

-- The p_ VARIABLES are normally PARAMETERS in a Data Warehouse Designer-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 30

Ws_admin_v Views

Admin views provide a means of interacting with the QAD Data Warehouse Designer 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_dim_col	Created from the ws_dim_tab and ws_dim_col tables
ws_admin_v_error	Created from the ws_wrk_error_log table
ws_admin_v_fact_col	Created from the ws_fact_tab and ws_fact_col tables
ws_admin_v_fact_join	Created from the ws_fact_tab and ws_fact_col tables
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	1046
Ws_admin_v_dim_col	1047
Ws_admin_v_error	1049
Ws_admin_v_fact_col	1050
Ws_admin_v_fact_join	1053
Ws_admin_v_sched	1055
Ws_admin_v_task	1059

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 Callable Routines Common Input (on page 907)
wa_job	See Callable Routines Common Input (on page 907)
wa_task	See Callable Routines Common Input (on page 907)
wa_status	See Callable Routines Common Input (on page 907)
wa_message	the message associated with this audit log
wa_db_msg_desc	the database message associated with this audit log

Screen shot of Sample Data returned

SQL Script

```
SQLQuery3.sql - L:\TITLE-B\User (53))
USE [Warehouse]
GO

/***** Object: View [dbo].[ws_admin_v_audit]    Script Date: 03/23/2
SET ANSI_NULLS ON
GO

SET QUOTED_IDENTIFIER ON
GO

create view [dbo].[ws_admin_v_audit] (
    wa_time_stamp, wa_sequence, wa_job, wa_task, wa_status,
    wa_message, wa_db_msg_desc )
as
select wa_time_stamp, wa_sequence, wa_job, wa_task, wa_status,
    wa_message, wa_db_msg_desc
from ws_wrk_audit_log
--order by isnull(wa_row_number,0), wa_time_stamp, wa_job, wa_task
GO
```

Ws_admin_v_dim_col

This Dimension Column view is created from the ws_dim_tab and ws_dim_col tables.

Columns

The following columns are created:

Column	Description
dimension_table	name of the dimension table
dim_display_value	display value of the dimension
dim_description	description of the dimension
dim_type	a character value indicating the dimension type T - time dimension V - dimension view D - dimension C - slowly changing dimension P - previous value dimension else - unknown
column_name	column name
business_name	business name chosen for the column
description	description of the dimension column
data_type	data type of the dimension column
source_table	source table of the dimension table
source_column	source column of the dimension column
nulls_allowed	flag indicating whether or not nulls are allowed
numeric	flag indicating whether or not the column is numeric
additive	flag indicating whether or not the column is additive
attribute	flag indicating whether or not the column is an attribute
format	the dimension column format
column_order	the dimension column order

Screen shot of Sample Data returned

SQL Script

```
SQLQuery4.sql - L...TTLE-B\User (53))
create view [dbo].[ws_admin_v_dim_col] (
    dimension_table,
    dim_display_value,
    dim_description,
    dim_type,
    column_name,
    business_name,
    description,
    data_Type,
    source_table,
    source_column,
    nulls_allowed,
    numeric,
    additive,
    attribute,Format,column_order)
AS
select
    dt_table_name,
    dt_display_name,
    dt_description,
    CASE dt_type_ind
        WHEN 'T' THEN 'Time dimension'
        WHEN 'V' THEN 'Dimension View'
        WHEN 'D' THEN 'Dimension'
        WHEN 'C' THEN 'Slowly changing dimension'
        WHEN 'P' THEN 'Previous value dimension'
        ELSE 'Unknown'
    END,
    dc_col_name ,
    dc_display_name,
    dc_src_strategy,
    dc_data_type,
    dc_src_table,
    dc_src_column,
    dc_nulls_flag,
    dc_numeric_flag,
    dc_additive_flag,
    dc_attribute_flag,
    dc_format,
    dc_order
from ws_dim_tab LEFT OUTER JOIN ws_dim_col
on dt_obj_key = dc_obj_key
where dt_type_ind in ('D','C','T','V','P')
and isnull(dc_eul_flag,'N') = 'Y'
GO
```

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 907)
wd_job	See <i>Callable Routines Common Input</i> (on page 907)
wd_task	See <i>Callable Routines Common Input</i> (on page 907)
wd_status	See <i>Callable Routines Common Input</i> (on page 907)
wd_message	the message associated with this audit log
wd_db_msg_desc	the database message associated with this audit log

Screen shot of Sample Data returned

SQL Script

```
SQLQuery5.sql - L:\TITLE-B\User (53))
USE [Warehouse]
GO

/***** Object:  View [dbo].[ws_admin_v_error]      Script Date: 03/23/2012 11:56:46 *****/
SET ANSI_NULLS ON
GO

SET QUOTED_IDENTIFIER ON
GO

create view [dbo].[ws_admin_v_error] ( wd_time_stamp, wd_sequence,wd_job, wd_task, wd_status,
wd_message,wd_db_msg_desc )
as
select  wd_time_stamp, wd_sequence,wd_job, wd_task, wd_status,
wd_message,wd_db_msg_desc
from ws_wrk_error_log
--order by isnull(wd_row_number,0),wd_time_stamp, wd_job, wd_task
GO
```

Ws_admin_v_fact_col

This Fact Column view is created from the ws_fact_tab and ws_fact_col tables.

Columns

The following columns are created:

Column	Description
fact_table	name of the fact table
fact_display_value	display value of the fact table
fact_description	description of the fact table
fact_type	a character value indicating the fact type D - detail R - rollup / combined S - kpi P - detail partitioned Q - rollup partitioned else unknown
column_name	column name
business_name	business name chosen for the column
description	description of the fact column
data_type	data type of the fact column
source_table	source table of the fact table
source_column	source column of the dimension column
nulls_allowed	flag indicating whether or not nulls are allowed
numeric	flag indicating whether or not the column is numeric
additive	flag indicating whether or not the column is additive
attribute	flag indicating whether or not the column is an attribute
format	the fact column format
column_order	the fact column order

Screen shot of Sample Data returned

SQL Script

```
SQLQuery6.sql - L...TTLE-B\User (53))
create view [dbo].[ws_admin_v_fact_col] (
    "fact_table",
    "fact_display_value",
    "fact_description",
    "fact_type",
    "column_name",
    "business_name",
    "description",
    "data_Type",
    "source_table",
    "source_column",
    "nulls_allowed",
    "numeric",
    "additive",
    "attribute",
    "format",
    "column_order")
as
select
    ft_table_name,
    ft_display_name,
    ft_description,
    CASE ft_type_ind
        WHEN 'D' THEN 'Detail'
        WHEN 'R' THEN 'Rollup/Combined'
        WHEN 'S' THEN 'KPI'
        WHEN 'P' THEN 'Detail partitioned'
        WHEN 'Q' THEN 'Rollup partitioned'
        ELSE 'Unknown'
    END,
    fc_col_name,
    fc_display_name,
    fc_src_strategy,
    fc_data_type,
    fc_src_table,
    fc_src_column,
    fc_nulls_flag,
    fc_numeric_flag,
    fc_additive_flag,
    fc_attribute_flag,
    fc_format,
    fc_order
from ws_fact_tab LEFT OUTER JOIN ws_fact_col
on ft_obj_key = fc_obj_key
where ft_type_ind in ('D','S','R','P','Q')
and isnull(fc_eul_flag,'N') = 'Y'
GO
```


Ws_admin_v_fact_join

This Fact Column view is created from the ws_fact_tab and ws_fact_col tables.

Columns

The following columns are created:

Column	Description
fact_table	fact table name
fact_column	fact table column name
dimension_table	dimension table name
dimension_column	dimension table column name

Screen shot of Sample Data returned

SQL Script

```
SQLQuery7.sql - L:\TTLE-B\User (53))
USE [Warehouse]
GO

/***** Object:  View [dbo].[ws_admin_v_fact_join]
SET ANSI_NULLS ON
GO

SET QUOTED_IDENTIFIER ON
GO

create view [dbo].[ws_admin_v_fact_join] (
    fact_table,
    fact_column,
    dimension_table,
    dimension_column)
as
select
    ft_table_name ,
    fc_col_name,
    fc_src_table,
    fc_src_column
from ws_fact_tab LEFT OUTER JOIN ws_fact_col
on ft_obj_key = fc_obj_key
where ft_type_ind in ('D','S','R','Q','P')
and isnull(fc_join_flag,'N') = 'Y'
GO
```


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

Screen shot of Sample Data returned

SQL Script

SQLQuery8.sql - L...TTLE-B\User (53))

```
USE [Warehouse]
GO

/***** Object: View [dbo].[ws_admin_v_sched]      Script
SET ANSI_NULLS ON
GO

SET QUOTED_IDENTIFIER ON
GO

create view [dbo].[ws_admin_v_sched] (
    "type",
    "job_name",
    "status",
    "sequence",
    "started_scheduled",
    "completed",
    "hours_elapsed" ,
    "minutes_elapsed",
    "okay",
    "info",
    "warn",
    "detail",
    "error")
as
select
    'Waiting',
    wjc_name,
    CASE wjc_status
        WHEN 'H' THEN 'On Hold'
        WHEN 'R' THEN 'Running'
        WHEN 'P' THEN 'Pending'
        WHEN 'W' THEN 'Waiting'
        WHEN 'C' THEN 'Completed'
        WHEN 'B' THEN 'Blocked'
        WHEN 'F' THEN 'Failed'
        WHEN 'G' THEN 'Failed - Aborted'
        WHEN 'E' THEN 'Error Completion'
        ELSE 'Unknown'
    END,
```

```

wjc_sequence,
wjc_start_after,
GETDATE(),
0,
0,
0,
0,
0,
0,
0
    from ws_wrk_job_ctrl
union
    select 'Runnnning',wjr_name,
        CASE wjr_status
            WHEN 'H' THEN 'On Hold'
            WHEN 'R' THEN 'Running'
            WHEN 'P' THEN 'Pending'
            WHEN 'W' THEN 'Waiting'
            WHEN 'C' THEN 'Completed'
            WHEN 'B' THEN 'Blocked'
            WHEN 'F' THEN 'Failed'
            WHEN 'G' THEN 'Failed - Aborted'
            WHEN 'E' THEN 'Error Completion'
            ELSE 'Unknown'
        END,
        wjr_sequence, wjr_started, wjr_completed,
        CONVERT(integer,DATEDIFF(mi,wjr_started,GETDATE())/60),
        DATEDIFF(mi,wjr_started,GETDATE()) -
        CONVERT(integer,DATEDIFF(mi,wjr_started,GETDATE())/60)*60,
        wjr_okay_count,
        wjr_info_count, wjr_warning_count, wjr_detail_count,
        wjr_error_count
    from ws_wrk_job_run
union
    select 'Finished',wjl_name,
        CASE wjl_status
            WHEN 'H' THEN 'On Hold'
            WHEN 'R' THEN 'Running'
            WHEN 'P' THEN 'Pending'
            WHEN 'W' THEN 'Waiting'
            WHEN 'C' THEN 'Completed'
            WHEN 'B' THEN 'Blocked'
            WHEN 'F' THEN 'Failed'
            WHEN 'G' THEN 'Failed - Aborted'
            WHEN 'E' THEN 'Error Completion'
            ELSE 'Unknown'
        END,
        wjl_sequence, wjl_started, wjl_completed,
        wjl_elapsed_hh, wjl_elapsed_mi, wjl_okay_count,
        wjl_info_count, wjl_warning_count, wjl_detail_count,
        wjl_error_count
    from ws_wrk_job_log
GO

```


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

Screen shot of Sample Data returned

SQL Script

SQLQuery10.sql -...TTLE-B\User (53))

```
USE [Warehouse]
GO

/***** Object:  View [dbo].[ws_admin_v_task]      Script Date: 11/11/2015 1:58:58 PM *****/
SET ANSI_NULLS ON
GO

SET QUOTED_IDENTIFIER ON
GO

create view [dbo].[ws_admin_v_task] (
    "Result",
    "task_name",
    "status",
    "sequence",
    "started",
    "completed",
    "hours_elapsed" ,
    "minutes_elapsed",
    "info",
    "warn",
    "detail",
    "inserted",
    "updated",
    "replaced",
    "deleted",
    "discarded",
    "rejected",
    "errored")
as
Select
    'Runnnning',wtr_name,
        CASE wtr_run_status
            WHEN 'H' THEN 'On Hold'
            WHEN 'R' THEN 'Running'
            WHEN 'P' THEN 'Pending'
            WHEN 'W' THEN 'Waiting'
            WHEN 'C' THEN 'Completed'
            WHEN 'B' THEN 'Blocked'
            WHEN 'F' THEN 'Failed'
            WHEN 'G' THEN 'Failed - Aborted'
            WHEN 'E' THEN 'Error Completion'
            ELSE 'Unknown'
        END,
```



```

wtr_sequence, wtr_started, wtr_completed,
CONVERT(integer,DATEDIFF(mi,wtr_started,GETDATE())/60),
DATEDIFF(mi,wtr_started,GETDATE()) -
CONVERT(integer,DATEDIFF(mi,wtr_started,GETDATE())/60)*60,
wtr_info_count, wtr_warning_count, wtr_detail_count,
wtr_rec_inserted, wtr_rec_updated,wtr_rec_replaced,wtr_rec_deleted,wtr_rec_discarded,
wtr_rec_rejected,wtr_rec_errored
from ws_wrk_task_run
union
select wtl_return_msg,wtl_name,
CASE wtl_run_status
WHEN 'H' THEN 'On Hold'
WHEN 'R' THEN 'Running'
WHEN 'P' THEN 'Pending'
WHEN 'W' THEN 'Waiting'
WHEN 'C' THEN 'Completed'
WHEN 'B' THEN 'Blocked'
WHEN 'F' THEN 'Failed'
WHEN 'G' THEN 'Failed - Aborted'
WHEN 'E' THEN 'Error Completion'
ELSE 'Unknown'
END,
wtl_sequence, wtl_started, wtl_completed,
wtl_elapsed_hh, wtl_elapsed_mi,
wtl_info_count, wtl_warning_count, wtl_detail_count,
wtl_rec_inserted, wtl_rec_updated,wtl_rec_replaced,wtl_rec_deleted,wtl_rec_discarded,
wtl_rec_rejected,wtl_rec_errored
from ws_wrk_task_log
GO

```


Chapter 31

Retrofitting

QAD Data Warehouse Designer 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 Copy** object type in QAD Data Warehouse Designer and the **Retrofit tables wizard**.

See **Retrofitting** (see "**Migrating the Data Warehouse Database Platform**" on page 1064, "**Retrofitting an OLAP Object**" on page 586)

In This Chapter

Migrating the Data Warehouse Database Platform	1064
Importing a Data Model	1072
Re-Targeting Source Tables	1077
Retro Copy Column Properties	1079

Migrating the Data Warehouse Database Platform


QAD Data Warehouse Designer has an advanced retrofitting wizard for migrating an existing data warehouse from one relational database to another.

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 password

Note: The extract 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 Right-click on the **Retro Copy object group** in the object tree in the left pane and select **Select Source Tables**.



- 4 The Retrofit tables dialog appears. In the **Source Connection** drop-down list choose the connection set up in step 1. Enter a **User ID** and **Password**. A list of **databases** appears in the left pane.

Select the connection for the source data.
Select tables from the source list and click Add to move the selected tables to the Retrofit list.
Multiple tables may be selected. Click OK to begin retrofitting.

Source Connection: Tutorial (OLTP) ▼

User ID: tutorial Password: ●●●

☒ Rename Object if Exists
☒ Include Indexes of the Table
☒ Define Object Type Based on Naming Match
☐ Add Ancillary Columns, (e.g. dss_update_time)

Source System

- dbo
- INFORMATION_SCHEMA
- sys
- System\$

Table/View


Retrofit Tables

Schema	Retrofit Table


Refresh Schema List

OK Cancel


- 5 Double-click on the database/user/schema in the left pane list. A list of **tables** in the database is displayed in the middle pane.




- 6 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)** checkbox and click **OK**.



- 7 QAD Data Warehouse Designer now acquires the metadata for the tables being migrated and creates a new QAD Data Warehouse Designer Retro Copy object for each.
- 8 Double-click on the **Retro Copy object group** in the left pane. Select all retro copies 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.



9 Click on the **Scheduler** button on the toolbar.



10 Create a new job to run straight away.

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:

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


- 11 Add all Retro Copy objects created in steps (3) to (8) and click on **Group Object Types**.




- 12 Once the job has completed, return to the QAD Data Warehouse Designer builder. Double-click on the **Retro Copy object group**. Select all objects in the middle pane and from the right-click menu select **Set Table Type to Definition**.



- 13 In the middle pane, select all objects. Right-click and select **Convert to Target Object**. QAD Data Warehouse Designer now converts the **Retro Copy** objects to the appropriate object types.



- 14 There are no longer any Retro copy objects. They have been converted to Load, Stage, Dimension or Fact objects.



- 15 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 1077) for more information.
- 16 Convert the old data warehouses code to QAD Data Warehouse Designer procedure in the new data warehouse database. See **Integrate Procedures** (see "**Integrating, Procedures**" on page 1094) for more information.
- 17 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 233) for more information.

Importing a Data Model


QAD Data Warehouse Designer 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 QAD Data Warehouse Designer metadata as Retro Copy objects.
- 5 Convert the Retro Copy 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 select **Select Source Tables**.



- 2 The **Retrofit tables** dialog is displayed. In the **Source Connection** drop-down list choose the **DataWarehouse** connection. Enter the **User ID** and **Password**. A list of **databases** appears in the left pane (yours list will be different).

Select the connection for the source data.
Select tables from the source list and click Add to move the selected tables to the Retrofit list.
Multiple tables may be selected. Click OK to begin retrofitting.

Source Connection: DataWarehouse ▼

User ID: dbo Password: ●●●

☒ Rename Object if Exists
☒ Include Indexes of the Table
☒ Define Object Type Based on Naming Match
☐ Add Ancillary Columns. (e.g. dss_update_time)

Source System

- dbo
- INFORMATION_SCHEMA
- sys
- System\$
- System\$

Table/View

Retrofit Tables

Schema	Retrofit Table

Refresh Schema List

OK Cancel

- 3 Double-click on the data warehouse database/schema in the left pane list. A list of **tables** in the database is displayed in the middle pane.

Retrofit Tables

Select the connection for the source data.
Select tables from the source list and click Add to move the selected tables to the Retrofit list.
Multiple tables may be selected. Click OK to begin retrofitting.

Source Connection:

User ID: Password:

☒ Rename Object if Exists
☒ Include Indexes of the Table
☒ Define Object Type Based on Naming Match
☐ Add Ancillary Columns. (e.g. dss_update_time)

Source System

- dbo - INFORMATION_SCHEMA - sys - System\$ - System\$	<table border="1"><thead><tr><th>Table/View</th></tr></thead><tbody><tr><td>agg_sa_customer</td></tr><tr><td>agg_sa_product</td></tr><tr><td>dim_customer</td></tr><tr><td>dim_date</td></tr><tr><td>dim_kpi</td></tr><tr><td>dim_product</td></tr><tr><td>dim_state</td></tr><tr><td>ds_customer</td></tr><tr><td>ds_order_header</td></tr><tr><td>ds_order_line</td></tr><tr><td>ds_product</td></tr><tr><td>dss_current_day_details</td></tr><tr><td>dss_fact_table</td></tr><tr><td>dss_parameter</td></tr><tr><td>dss_source_system</td></tr><tr><td>fact_budget</td></tr><tr><td>fact_forecast</td></tr><tr><td>fact_kpi</td></tr><tr><td>fact_sales_analysis</td></tr></tbody></table>	Table/View	agg_sa_customer	agg_sa_product	dim_customer	dim_date	dim_kpi	dim_product	dim_state	ds_customer	ds_order_header	ds_order_line	ds_product	dss_current_day_details	dss_fact_table	dss_parameter	dss_source_system	fact_budget	fact_forecast	fact_kpi	fact_sales_analysis
Table/View																					
agg_sa_customer																					
agg_sa_product																					
dim_customer																					
dim_date																					
dim_kpi																					
dim_product																					
dim_state																					
ds_customer																					
ds_order_header																					
ds_order_line																					
ds_product																					
dss_current_day_details																					
dss_fact_table																					
dss_parameter																					
dss_source_system																					
fact_budget																					
fact_forecast																					
fact_kpi																					
fact_sales_analysis																					

Refresh Schema List

Retrofit Tables

Schema	Retrofit Table

OK Cancel

- 4 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)** checkbox. Click **OK**.

Retrofit Tables

Select the connection for the source data.
Select tables from the source list and click Add to move the selected tables to the Retrofit list.
Multiple tables may be selected. Click OK to begin retrofitting.

Source Connection: DataWarehouse
User ID: dbo Password: ...

☒ Rename Object if Exists
☒ Include Indexes of the Table
☒ Define Object Type Based on Naming Match
☒ Add Ancillary Columns. (e.g. dss_update_time)

Source System

Table/View


Retrofit Tables

Refresh Schema List


OK Cancel

Schema	Retrofit Table
dbo	agg_sa_customer
dbo	agg_sa_product

- 5 Double-click on the **Retro Copy object group** in the object tree in the left pane. Select all objects in the middle pane and right-click select **Convert to Target Object**.



The new aggregate tables have been imported.



Re-Targeting Source Tables

Objects that have been retrofitted into the QAD Data Warehouse Designer metadata have themselves as their source table:

dim_product Columns				
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 QAD Data Warehouse Designer 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.

Apply Changes to:

☒ Column Source Table

☐ Data Type

☐ Nulls

☐ Format String

☐ Numeric

☐ Additive

☐ Attribute

☐ EUL

Original Value:

dim_product

New Value:

load_product


☐ Update matched columns only

Apply

Close

Help

- 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.

dim_product Columns				
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 Copy Column Properties

Each Retro Copy 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

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 Data Warehouse Designer-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. In the case of dimension keys this field is used to show the join between the Retro Copy table and the dimension, once it has been defined as part of the Retro Copy table update procedure generation.

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 data types for Oracle are integer, number, char(), varchar2() and date. For SQL Server common types are integer, numeric, varchar() and datetime. 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.

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.

Key type	Meaning
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.
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

Identifies the source table where the column's data comes from.

Source Column

Identifies the source column where the column's data comes from.

Transformation

Transformation. [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.

Retro Copy Column Properties Screen

A sample **Properties** screen is as follows:

The screenshot shows a software window titled "Retro Column customer.code". On the left is a sidebar with "Properties" and "Transformation" tabs, with "Properties" selected. The main area displays a form with several sections: "General", "Physical Definition", "Meta Definition", and "Source Details".

General	
Table Name	customer
Column Name	code
Business Display Name	code
Column Description	Unique code to identify the customer. Forms a hierarchy with city and state.

Physical Definition	
Column Order	20
Data Type	numeric(6)
Null Values Allowed	True
Default Value	

Meta Definition	
Format	#,##0
Numeric	True
Additive	True
Attribute	False
Business Key	False
Key Type (0,A,B,C,...)	

Source Details	
Source Table	dim_customer
Source Column	code
Transformation	
Join	False

Column Name
Database-compliant name of the column.
Dialog Opening Value: code

At the bottom right are five buttons: "<- Update", "Update ->", "OK", "Cancel", and "Help".

Retro Copy Column Transformations

It is possible to do transformations on Retro Copy 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:

The screenshot shows a software window titled "Retro Column customer.code". On the left is a sidebar with "Properties" and "Transformation" tabs. The main area contains a "Target:" field with "code" and a "Source:" field with "ds_customer.code", each with a "Paste" button. Below these is a warning message: "Warning: Transformations are only used for file and script loads or when transformations are enabled in the source screen of load table properties for database loads. See Help". To the right of the warning are two dropdown menus: "Transform Stage:" set to "During Load" and "Function Set:" set to "SQL Server Standard Funct". Below the warning is a large empty text area for the transformation code. To the right of this area is a "Functions" panel with a tree view showing "Available Columns". At the bottom left of the main area are checkboxes for "Word Wrap Displayed Code", "Function Syntax:", and "Function Desc:". The bottom of the window has a row of buttons: "< Update", "Update ->", "OK", "Cancel", and "Help".

Note: Transformations are only put into effect when the procedure is re-generated.

See **Transformations** (on page 603) for more details.

Chapter 32

Integrating QAD Data Warehouse Designer into an Existing Warehouse

Two main options exist in terms of bringing QAD Data Warehouse Designer into an existing data warehouse environment:

- 1 Rebuild tables and procedures with QAD Data Warehouse Designer.
- 2 Integrating existing tables and procedures into QAD Data Warehouse Designer.

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 QAD Data Warehouse Designer. 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 1086) section for the approach to achieve this option.

Integrate

Existing data warehouse tables can be identified to QAD Data Warehouse Designer. 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 1087) 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 a 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	1086
Integrating	1087

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 QAD 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 QAD Data Warehouse Designer 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 dimension, stage and fact 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 QAD metadata repository into this test environment. See the **Installation and Administration Guide** for instructions on how to load a metadata repository.
Working within the QAD Data Warehouse Designer tool proceed to:
- 3 Create any connections to UNIX or Windows servers where host scripts are currently executed. See either **creating a Windows connection** (see "**Windows**" on page 142) or **creating a Unix connection** (see "**UNIX**" on page 170).
- 4 Create a Data Warehouse connection mapping back to the test environment. See **creating a connection to the data warehouse** (see "**Data Warehouse**" on page 116).
- 5 Incorporate any Host system scripts currently used. See **incorporating host scripts** (see "**Integrating, Host Scripts**" on page 1088).
- 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 1089).
- 8 Answer the retrofit questions, and build any required procedures. See **integrate questions** (see "**Integrating, Questions**" on page 1090).
- 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 1094).
- 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 will ask if you wish to delete the object and drop the table. Answer **No**. A second dialog box will now appear asking if just the metadata is to be deleted. Answer **Yes** to this question and only the metadata will be removed.

Integrating, Host Scripts

Existing host scripts (either UNIX or Windows) can be brought into the QAD Data Warehouse Designer metadata. To incorporate an existing script the process is as follows:

- 1 Create a **Host Script** object using QAD Data Warehouse Designer. 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 the 'monitor_db_mail Definition' dialog box. The 'Script Properties' tab is selected on the left. The main area contains the following fields and controls:

- Name:** monitor_db_mail
- Type:** Windows script (dropdown menu)
- Purpose:** (empty text area)
- Owner:** dbo
- Last Update By:** (empty text field)
- Default Connect:** (empty dropdown menu)
- Checkout:**
 - Checked Out By:** (empty text field)
 - Checkout Reason or Last Update:** New Script
- Timestamps:**
 - Created:** (empty text field)
 - Last Update:** (empty text field)
 - Compiled:** (empty text field)

Buttons at the bottom: OK, Cancel, Help.

- 3 Double-click on the new script or right-click on the 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 "**Host Script Load**" on page 283) 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:


- ☐ Dimension tables
- ☐ Fact tables
- ☐ Aggregate 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 Dimension table group.


Integrating, Questions

When a table within the data warehouse schema, that is unknown to QAD Data Warehouse Designer, is dropped onto a table target the following dialog appears.




If this is an integrate then click **Yes** to proceed with the integrate 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 will ask 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 dimension or fact table then the following dialog will appear to enable the definition of any artificial or surrogate key. If no such key exists then simply proceed to the next question.



The dialog box is titled "Retrofit key definition" and has a close button (X) in the top right corner. It contains the following elements:

- Instructions:** "Select the artificial (surrogate) key columns from the column list and add them to the artificial key list. If no such column(s) exist then simply click the OK button."
- Artificial key list:** A text box on the left containing the text "dim_customer_key".
- Column list:** A list box on the right containing the following items: address, city, code, dim_customer_key, dss_current_flag, dss_end_date, dss_start_date, dss_version, and name.
- Buttons:**
 - "<-- Add" button between the Artificial key list and the Column list.
 - "Remove -->" button between the Artificial key list and the Column list.
 - "OK" button in the top right corner.
 - "Cancel" button below the "OK" button.

Get Key Function


If the target table is a dimension you will be asked if you want to generate a get key function. The get key function is used by QAD Data Warehouse Designer for rapid prototyping. It returns the artificial key when passed the business key. Although this function is possibly not necessary in the existing environment it is useful if the dimension is to be used in any expansion of the data warehouse. If an existing get key type function exists, then this can be used, although changes will probably need to be made to the generated procedures to utilize this existing get key function.

Update Procedure

If the target table is a dimension, stage, fact or aggregate a dialog box will ask if an update procedure or update template 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.

Artificial Key Sequence

If a get key function or update procedure is to be built and an artificial key is in use for the table then a dialog will prompt for the artificial key sequence. A drop-down list provides all the sequences defined for this database user. This sequence is an Oracle sequence that is used in the automatic generation of the artificial key. If no such sequence is used then select the *** no sequence *** option.



Select the Artificial key sequence

Select the sequence used to allocate the next artificial key number. Select *** No Sequence *** if a sequence is not used.

Artificial key sequence


DIM_CUSTOMER_SEQ

OK

Cancel

Update Procedure Definition

If an update procedure was selected then the following dialog will display to help define how the update procedure is generated.



The dialog box titled "Define update procedure" contains the following text and controls:

The scheduler needs specific parameters in order to control the running of the tasks. To incorporate existing PL/SQL code used to update this table two scenarios are offered. These are:

- (1) Generate an update procedure and then manually cut and paste any specific code from the current update procedure.
- (2) Select the existing procedure and generate a wrapper procedure that calls the existing procedure.

For both cases the generated procedure will not be compiled, and will need to be manually checked and compiled.

New update procedure name:

To adopt scenario (2) enter the existing procedure name, or press the 'Lookup' button and select from the resulting list.

Existing procedure name:

Buttons: Build, Cancel

As defined in this dialog box the scheduler needs specific parameters in order to control and run a procedure. An existing procedure will not have these parameters. There are basically two forms of generated procedure. The first option is to generate a procedure that looks much like a standard QAD Data Warehouse Designer procedure and then manually insert the changes from the existing procedure. The second option is to generate a wrapper procedure that calls the existing procedure. This second option will still require manual changes to the procedure to return a valid return code and message.

Business Key definition

If a get key or update procedure has been selected a dialog will now prompt 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.

Index definition

Any indexes associated with the table will now be 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** checkbox if appropriate and the scheduler is to be used to manage these indexes.

Procedure creation

If defined the get key and update procedures will now be generated. They will need to be manually edited and compiled. See the section on **integrating procedures** (see "**Integrating, Procedures**" on page 1094).

Integrating, Procedures

The procedures managed by the QAD scheduler require the following standards.

Parameters

The procedure must have the following parameters in the following order:

Parameter name	Input/Output	Oracle Type	SQL Server/DB2 Type
p_sequence	Input	Number	Integer
p_job_name	Input	Varchar2	Varchar(256)
p_task_name	Input	Varchar2	Varchar(256)
p_job_id	Input	Number	Integer
p_task_id	Input	Number	Integer
p_return_msg	Output	Varchar2	Varchar(256)
p_status	Output	Number	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 integrating a procedure will be in adapting it to the QAD scheduler standards and work flow.

Integrating, Views

When integrating views an additional step is required if you want QAD Data Warehouse Designer 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.

Change the source column on the artificial key to match the artificial key in the table from which the view was created.

Integrating, QAD Tables

When integrating QAD generated tables and views a number of additional considerations need to be taken.


The Dimension tables and views make use of secondary business key columns such as dss_source_system_key, dss_current_flag, dss_end_date, dss_version to assist in handling various forms of slowly changing dimensions and non conformance. These secondary business keys need to be flagged as such against the dimension.

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 33

Join Information Maintenance


Join Information Maintenance is available for the maintenance of joins between normalized 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. To access this tool, use the **Tools/Join Information Maintenance** menu option.



Filter: All Objects					
Object	Column	Join Type	Object	Column	Part
normal_customer	customer_code	n - 1	normal_forecast	customer_code	1
normal_product	product_code	n - 1	normal_forecast	product_code	1

Filter


Allows you to filter on the tables in the data warehouse. You can choose **All Objects** or **Normalized**.



Filter: All Objects					
Object	Column	Join Type	Object	Column	Part
normal_customer	customer_code	n - 1	normal_forecast	customer_code	1
normal_product	product_code	n - 1	normal_forecast	product_code	1

The table joins are displayed; showing the tables, columns and join types.

To add a join, click on the **Add Join** button.



Specify the join by selecting a table and column from the left hand side and then the joined table details on the right hand side. For multiple column joins check the 'multi column join' box.

Object Type: Normalized	☐ Multiple Column Join	
Table/View: normal_product	☒ Primary Join	Table/View: normal_forecast
Column: product_code	Join Type: Many to One	Column: product_code

Finish Cancel Add Join

For each table, enter in the following details:

Object Type

Enter the table type.

Table/View

Enter the name of the table.

Column

Enter the column for the join.

For multiple column joins, select the **Multiple Column Join** checkbox.

If this is a primary join, select the **Primary Join** checkbox.

For the **Join Type**, choose between the following:

- ☐ Undefined
- ☐ Many to One
- ☐ One to One
- ☐ One to Many
- ☐ Many to Many

Once you have entered the details for the join, click the **Add Join** button to add the join. Click **Finish** to close the Add Join dialog.

Chapter 34

Upgrading Data Warehouse Designer

The upgrading of Data Warehouse Designer is discussed at some length in the **Installation and Administration Guide**.

Chapter 35

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 Data Source Name (DSN)** to match the login **Data Source**.

The screenshot shows the 'Connection DataWarehouse' dialog box with the 'General' tab selected. The 'ODBC Data Source Name (DSN)' is set to 'WslWarehouse'. The 'Data Warehouse Connection Indicator' is checked. The 'Source System' section shows 'Database ID' as 'WslWarehouse' and 'Provider Name' as 'SQLOLEDB'. The 'Database Credentials' section is empty. The 'Other' section shows 'New Table Default Load Type' as 'Database link load', 'SSIS Connection String' as 'Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...', 'Data Type Mapping Set' as '(Default)', and 'Default Transform Function Set' as '(Default)'. The 'When Connection is an OLAP Data Source' section shows 'OLAP Connection String' as 'Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...', 'Connection Provider/Driver' as 'SQLOLEDB', 'Data Warehouse Server' as 'DOC', and 'Data Warehouse Database ID' as 'WslWarehouse'.

General	
Connection Name	DataWarehouse
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	WslWarehouse
Data Warehouse Connection Indicator	☒

Source System	
Database ID	WslWarehouse
Database Link Name	
Provider Name	SQLOLEDB

Database Credentials	
Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other	
New Table Default Load Type	Database link load
SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
Data Type Mapping Set	(Default)
Default Transform Function Set	(Default)

When Connection is an OLAP Data Source	
OLAP Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
Connection Provider/Driver	SQLOLEDB
Data Warehouse Server	DOC
Data Warehouse Database ID	WslWarehouse

- Logoff and log back in using the **Data Source** with the same name as the Data Warehouse Connection **ODBC Source**.



The image shows a 'Repository Login' dialog box for 'QAD BI Data Warehouse Designer'. The dialog has a blue header with the title 'Repository Login' and a close button. Below the title, it says 'Version 6.5 by QAD Inc', 'Copyright: 2012', and 'Licensed to QAD BI DWD Manual and Testing'. The main area is divided into two sections. The first section, labeled 'DATABASE' in red, contains three fields: 'Data Source:' with a dropdown menu showing 'WslWarehouse', 'Database Login ID:' with a text box containing 'dbo', and 'Password:' with an empty text box. The second section, labeled 'METADATA REPOSITORY' in red, contains one field: 'DWD User Name:' with a text box containing 'QAD QuickStart'. At the bottom right, there are three buttons: 'Help' (disabled), 'Cancel' (disabled), and 'Connect' (active, highlighted in red).

QAD BI Data Warehouse Designer

Repository Login
Version 6.5 by QAD Inc
Copyright: 2012
Licensed to QAD BI DWD Manual and Testing

DATABASE

Data Source: WslWarehouse

Database Login ID: dbo

Password:

METADATA REPOSITORY

DWD User Name: QAD QuickStart

Help Cancel Connect

- 2 Warning if more than one connection has the data warehouse checkbox on:

You can correct this issue by editing all the connections and making sure that only one is set to Data Warehouse:

The screenshot shows the 'Connection DataWarehouse' dialog box. The 'Data Warehouse Connection Indicator' checkbox is checked and highlighted with a red arrow. The 'Source System' section shows 'Database ID' as 'WslWarehouse' and 'Provider Name' as 'SQLOLEDB'.

General	
Connection Name	DataWarehouse
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	WslWarehouse
Data Warehouse Connection Indicator	☒
Source System	
Database ID	WslWarehouse
Database Link Name	
Provider Name	SQLOLEDB
Database Credentials	
Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	
Other	
New Table Default Load Type	Database link load
SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
Data Type Mapping Set	(Default)
Default Transform Function Set	(Default)
When Connection is an OLAP Data Source	
OLAP Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
Connection Provider/Driver	SQLOLEDB
Data Warehouse Server	DOC
Data Warehouse Database ID	WslWarehouse

OK Cancel Help

Chapter 36

Data Type Mappings

In This Chapter

- Using Data Type Mapping Sets 1106
- Maintaining Data Type Mapping Sets 1108
- Loading Data Type Mapping Sets 1130
- Exporting Data Type Mapping Sets 1132
- Data Type Mapping Examples 1134

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

Connection DataWarehouse

Properties

Notes

General

Connection Name	DataWarehouse
Connection Type	Database
Database Type	(local)
ODBC Data Source Name (DSN)	WslWarehouse
Data Warehouse Connection Indicator	☒

Source System

Database ID	WslWarehouse
Database Link Name	
Provider Name	SQLOLEDB

Database Credentials

Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other

New Table Default Load Type	Database link load
SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
Data Type Mapping Set	(Default)
Default Transform Function Set	(Default)

When Connection is an OLAP Data Source

OLAP Connection String	Standard SQL Server to SQL Server
Connection Provider/Driver	Standard Oracle to SQL Server
Data Warehouse Server	Standard Teradata to SQL Server
Data Warehouse Database ID	Standard DB2 to SQL Server
	Standard FoxPro to SQL Server
	Standard Other to SQL Server
	Standard File to SQL Server

Data Type Mapping Set

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

OK Cancel Help

Non-Data Warehouse Connection Properties Dialog

The screenshot shows the 'Connection Sales' dialog box. On the left is a sidebar with 'Properties' and 'Notes' tabs. The main area is divided into sections: 'General', 'ODBC', 'Credentials', and 'Other'. The 'Data Type Mapping Set' dropdown in the 'Other' section is open, showing a list of mapping sets. Below the list, a description for the 'Data Type Mapping Set' is provided. At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

General	
Connection Name	Sales
Connection Type	ODBC
Database Type	(local)
ODBC Data Source Name (DSN)	Sales
Data Warehouse Connection Indicator	☐

ODBC	
Work Directory	C:\temp

Credentials	
Extract User ID	
Extract User Password	
Administrator User ID	
Administrator User Password	

Other	
New Table Default Load Type	Native ODBC
SSIS Connection String	...
Data Type Mapping Set	(Default)

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


- (Default)
- Standard SQL Server to SQL Server
- Standard Oracle to SQL Server
- Standard Teradata to SQL Server
- Standard DB2 to SQL Server
- Standard FoxPro to SQL Server
- Standard Other to SQL Server
- Standard File to SQL Server

OK Cancel Help


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** dialog 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 1110)


To copy a data type mapping set, see **Copying a Data Type Mapping Set** (on page 1114)

To edit a data type mapping set, see **Editing a Data Type Mapping Set** (on page 1118)


To delete a data type mapping set, see **Deleting a Data Type Mapping Set** (on page 1128)

Creating a New Data Type Mapping Set

To create a new data type mapping set, select **Tools/Data Type Mappings/Maintain Data Type Mappings...**



Click on the **New** button, enter the name of the new Mapping Set and click **OK**.



Description

Enter the description for the data type mapping set.

You can then enter the individual mappings by clicking the **New** button on the right side of the dialog.

[illegible]

Similarly, you can delete an individual mapping, using the **Delete** button on the right side of the dialog; or you can edit a mapping, using the **Edit** button on the right side of the dialog.

To move a mapping up or down in the list, select a mapping and then click on the **Move Up** or **Move Down** button on the right side of the dialog.

Maintain Data Type Mappings

Data Type Mapping Set: Custom SQL Server to SQL Server

Description:


From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ bit				bit	0	
☒ tinyint				tinyint	1	
☒ tinyint identity				tinyint	2	
☒ smallint				smallint	3	
☒ smallint identity				smallint	4	
☒ int				int	5	
☒ int identity				int	6	
☒ bigint				bigint	7	
☒ bigint identity				bigint	8	
☒ numeric			0	numeric(%precision)	9	
☒ numeric			1-*	numeric(%precision...)	10	
☒ decimal			0	numeric(%precision)	11	
☒ decimal			1-*	numeric(%precision...)	12	
☒ smallmoney				smallmoney	13	
☒ money				money	14	
☒ float				float	15	
☒ real				real	16	
☒ date				date	17	
☒ smalldatetime				smalldatetime	18	

Buttons: New..., Copy, Edit..., Delete..., Move Up, Move Down, OK, Cancel


Note: The order of the data type mappings is significant as 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.

Copying a Data Type Mapping Set

To copy an existing data type mapping set, select **Tools/Data Type Mappings/Maintain Data Type Mappings...**



Select the Mapping Set to be copied from the drop-down list and then click on the **Copy** button. Enter the name of the new Mapping Set and click **OK**.



Description

Enter the description for the data type mapping set.

To edit a data type mapping, select the mapping and then click on the **Edit** button on the right side of the dialog.

Maintain Data Type Mappings

Data Type Mapping Set:
Custom SQL Server to SQL Server

New...Copy...Delete

Description:

Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ bit				bit	0	
☒ tinyint				tinyint	1	
☒ tinyint identity				tinyint	2	
☒ smallint				smallint	3	
☒ smallint identity				smallint	4	
☒ int				int	5	
☒ int identity				int	6	
☒ bigint				bigint	7	
☒ bigint identity				bigint	8	
☒ numeric			0	numeric(%precision)	9	
☒ numeric			1-*	numeric(%precision...	10	
☒ decimal			0	numeric(%precision)	11	
☒ decimal			1-*	numeric(%precision...	12	
☒ smallmoney				smallmoney	13	
☒ money				money	14	
☒ float				float	15	
☒ real				real	16	
☒ date				date	17	
☒ smalldatetime				smalldatetime	18	

New...
Copy
Edit...
Delete...

Move Up
Move Down

OK

Cancel

Similarly, you can delete a mapping, using the **Delete** button on the right side of the dialog; or you can add a new mapping, using the **New** button on the right side of the dialog.

To move a mapping up or down in the list, select the mapping and then use the **Move Up** or **Move Down** button on the right side of the dialog.

Maintain Data Type Mappings

Data Type Mapping Set: Custom SQL Server to SQL Server

Description:


From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ bit				bit	0	
☒ tinyint				tinyint	1	
☒ tinyint identity				tinyint	2	
☒ smallint				smallint	3	
☒ smallint identity				smallint	4	
☒ int				int	5	
☒ int identity				int	6	
☒ bigint				bigint	7	
☒ bigint identity				bigint	8	
☒ numeric			0	numeric(%precision)	9	
☒ numeric			1-*	numeric(%precision...)	10	
☒ decimal			0	numeric(%precision)	11	
☒ decimal			1-*	numeric(%precision...)	12	
☒ smallmoney				smallmoney	13	
☒ money				money	14	
☒ float				float	15	
☒ real				real	16	
☒ date				date	17	
☒ smalldatetime				smalldatetime	18	

Buttons: New..., Copy, Edit..., Delete..., Move Up, Move Down, OK, Cancel

Note: The order of the data type mappings is significant as 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.

Editing a Data Type Mapping Set

To edit a data type mapping set, select **Tools/Data Type Mappings/Maintain Data Type Mappings...**



Select a data type mapping set from the **Data Type Mapping Set** drop-down list.

Maintain Data Type Mappings

Data Type Mapping Set:

Custom SQL Server to SQL Server

New...

Copy...

Delete

Description:

Custom SQL Server to SQL Server 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.

Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ bit				bit	0	
☒ tinyint				tinyint	1	
☒ tinyint identity				tinyint	2	
☒ smallint				smallint	3	
☒ smallint identity				smallint	4	
☒ int				int	5	
☒ int identity				int	6	
☒ bigint				bigint	7	
☒ bigint identity				bigint	8	
☒ numeric			0	numeric(%precision)	9	
☒ numeric			1-*	numeric(%precision...)	10	
☒ decimal			0	numeric(%precision...)	11	
☒ smallmoney				smallmoney	12	
☒ money				money	13	
☒ float				float	14	
☒ real				real	15	
☒ date				date	16	
☒ smalldatetime				smalldatetime	17	
☒ datetime				datetime	18	

New...

Copy

Edit...

Delete...

Move Up

Move Down

OK

Cancel

Chapter 36 **Data Type** Mappings - 1119

On the right is a group of buttons used to maintain the list of mappings in a data type mapping set. These buttons are not available for standard mapping sets. Only **user defined** mapping sets are editable.

Note: The order of the mappings within a set is significant as 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.

To add a new mapping to the data type mapping set:

Click on the **New** button on the right.

Maintain Data Type Mappings

Data Type Mapping Set: Custom SQL Server to SQL Server [v] [New...] [Copy...] [Delete]


Description:
Custom SQL Server to SQL Server 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.

Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ bit				bit	0	
☒ tinyint				tinyint	1	
☒ tinyint identity				tinyint	2	
☒ smallint				smallint	3	
☒ smallint identity				smallint	4	
☒ int				int	5	
☒ int identity				int	6	
☒ bigint				bigint	7	
☒ bigint identity				bigint	8	
☒ numeric			0	numeric(%precision)	9	
☒ numeric			1-*	numeric(%precision...	10	
☒ decimal			0	numeric(%precision...	11	
☒ smallmoney				smallmoney	12	
☒ money				money	13	
☒ float				float	14	
☒ real				real	15	
☒ date				date	16	
☒ smalldatetime				smalldatetime	17	
☒ datetime				datetime	18	

[New...] [Copy] [Edit...] [Delete...] [Move Up] [Move Down] [OK] [Cancel]

Enter the details and click **OK**.



The image shows a 'Data Type Mapping' dialog box with a blue title bar and a red close button. It contains eight text input fields for mapping data types, a checkbox for 'Save to Set', and 'OK' and 'Cancel' buttons.

Field	Value
From Data Type:	
From Length:	
From Precision:	
From Scale:	
To Data Type:	
Transform Code:	
Format Value:	
Default Value:	

☒ Save to Set

OK Cancel

From Data Type

The source Data Type to map to the Target Data Type.

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 **Data Type Mapping Examples** (on page 1134)

Transform Code

The code to transform values into the mapped Data Type. Available tokens are %table_name, %column_name, %format, %data_type, %length, %precision and %scale. See **Data Type Mapping Examples** (on page 1134)

Format Value

The format of this Conversion's Data Type output.

Default Value

The default value of this Conversion's Data Type output.

To copy an existing mapping in the data type mapping set:

Select the mapping and then click on the **Copy** button on the right.

Maintain Data Type Mappings

Data Type Mapping Set:

Custom SQL Server to SQL Server

New...

Copy...

Delete

Description:

Custom SQL Server to SQL Server 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.

Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ bit				bit	0	
☒ tinyint				tinyint	1	
☒ tinyint identity				tinyint	2	
☒ smallint				smallint	3	
☒ smallint identity				smallint	4	
☒ int				int	5	
☒ int identity				int	6	
☒ bigint				bigint	7	
☒ bigint identity				bigint	8	
☒ numeric			0	numeric(%precision)	9	
☒ numeric			1-*	numeric(%precision...)	10	
☒ decimal			0	numeric(%precision...)	11	
☒ smallmoney				smallmoney	12	
☒ money				money	13	
☒ float				float	14	
☒ real				real	15	
☒ date				date	16	
☒ smalldatetime				smalldatetime	17	
☒ datetime				datetime	18	

New...

Copy

Edit...

Delete...

Move Up

Move Down

OK

Cancel

Chapter 36 **Data Type** Mappings - 1123

A copy of the mapping will be added to the bottom of the list. Select the mapping and click on the **Edit** button to change the details.

To edit an existing mapping in the data type mapping set:

Select the mapping you wish to edit and click the **Edit** button.

Maintain Data Type Mappings

Data Type Mapping Set:
Custom SQL Server to SQL Server

New...Copy...Delete

Description:


Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ bit				bit	0	
☒ tinyint				tinyint	1	
☒ tinyint identity				tinyint	2	
☒ smallint				smallint	3	
☒ smallint identity				smallint	4	
☒ int				int	5	
☒ int identity				int	6	
☒ bigint				bigint	7	
☒ bigint identity				bigint	8	
☒ numeric			0	numeric(%precision)	9	
☒ numeric			1-*	numeric(%precision)	10	
☒ decimal			0	numeric(%precision)	11	
☒ decimal			1-*	numeric(%precision)	12	
☒ smallmoney				smallmoney	13	
☒ money				money	14	
☒ float				float	15	
☒ real				real	16	
☒ date				date	17	
☒ smalldatetime				smalldatetime	18	

New...CopyEdit...Delete...Move UpMove Down

OKCancel

Change any fields as required and click **OK**.



The image shows a 'Data Type Mapping' dialog box with a blue title bar and a red close button. It contains several input fields for mapping data types. The 'From Data Type' field is set to 'numeric'. The 'From Length' and 'From Precision' fields are empty. The 'From Scale' field is set to '1-*'. The 'To Data Type' field is set to 'numeric(%precision,%scale)'. The 'Transform Code', 'Format Value', and 'Default Value' fields are empty. There is a 'Save to Set' checkbox which is checked. At the bottom right are 'OK' and 'Cancel' buttons.

From Data Type:	numeric
From Length:	
From Precision:	
From Scale:	1-*
To Data Type:	numeric(%precision,%scale)
Transform Code:	
Format Value:	
Default Value:	

☒ Save to Set

OK Cancel

To delete an existing mapping in the data type mapping set:

Select the mapping and then click on the **Delete** button.

Maintain Data Type Mappings

Data Type Mapping Set:
Custom SQL Server to SQL Server

Description:
Custom SQL Server to SQL Server 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.

Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ datetimeoffset		31		datetimeoffset(4)	39	
☒ datetimeoffset		32		datetimeoffset(5)	40	
☒ datetimeoffset		33		datetimeoffset(6)	41	
☒ datetimeoffset		34		datetimeoffset	42	
☒ timestamp				timestamp	43	
☒ varchar		1-2		varchar(%precision)	44	COALESCE(%table_name.%column_name,'U')
☒ varchar		3-6		varchar(%precision)	45	COALESCE(%table_name.%column_name,'UNK')
☒ varchar		7-*		varchar(%precision)	46	COALESCE(%table_name.%column_name,'UN...')
☒ nvarchar		0		nvarchar(max)	47	
☒ nvarchar		1-*		nvarchar(%precision)	48	
☒ text				text	49	
☒ ntext				ntext	50	
☒ varbinary	0			varbinary(max)	51	
☒ image				image	52	
☒ uniqueidentifier				char(36)	53	
☒ xml				xml	54	
☒ text				text	55	
☒ (New)					56	

OK Cancel

Click **OK** to delete.

Are you sure you want to delete this conversion?

OK Cancel

To move a mapping in the data type mapping set up or down in the list:

Select a mapping and then click on the **Move Up** button on the right to move the mapping up in the list; else click on the **Move Down** button on the right to move the mapping down in the list.

Maintain Data Type Mappings

Data Type Mapping Set:
Custom SQL Server to SQL Server

Description:
Custom SQL Server to SQL Server 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.


Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ float				float	14	
☒ real				real	15	
☒ date				date	16	
☒ smalldatetime				smalldatetime	17	
☒ datetime				datetime	18	
☒ time		8		time(0)	19	
☒ time		10		time(1)	20	
☒ time		11		time(2)	21	
☒ time		12		time(3)	22	
☒ time		13		time(4)	23	
☒ time		14		time(5)	24	
☒ time		15		time(6)	25	
☒ time		16		time	26	
☒ datetime2		19		datetime2(0)	27	
☒ datetime2		21		datetime2(1)	28	
☒ datetime2		22		datetime2(2)	29	
☒ datetime2		23		datetime2(3)	30	
☒ datetime2		24		datetime2(4)	31	
☒ datetime2		25		datetime2(5)	32	

Buttons: New..., Copy, Edit..., Delete..., Move Up, Move Down, OK, Cancel

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.

Maintain Data Type Mappings

Data Type Mapping Set:

Custom SQL Server to SQL Server

New...

Copy...

Delete

Description:

Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ bit				bit	0	
☒ tinyint				tinyint	1	
☒ tinyint identity				tinyint	2	
☒ smallint				smallint	3	
☒ smallint identity				smallint	4	
☒ int				int	5	
☒ int identity				int	6	
☒ bigint				bigint	7	
☒ bigint identity				bigint	8	
☒ numeric			0	numeric(%precision)	9	
☒ numeric			1-*	numeric(%precision...	10	
☒ decimal			0	numeric(%precision)	11	
☒ decimal			1-*	numeric(%precision...	12	
☒ smallmoney				smallmoney	13	
☒ money				money	14	
☒ float				float	15	
☒ real				real	16	
☒ date				date	17	
☒ smalldatetime				smalldatetime	18	

New...

Copy

Edit...

Delete...

Move Up

Move Down

OK


Cancel

Note: The **Delete** button is disabled for standard data type mapping sets.


Loading 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 dialog is displayed. Select the xml file to load the data type mappings. By default Data Warehouse Designer expects the xml files to be in **ProgramData\QAD\Work**.




Exporting 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, Data Warehouse Designer exports the xml file to **ProgramData\QAD\Work**, but this can be changed. Change the File name if necessary and click **Save**.



Data Type Mapping Examples

QAD Data Warehouse Designer allows you to create **Custom** Data Type Mapping Sets. These 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.

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 standard file to SQL Server 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).

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ boolean				bit	0	
☒ int				int	1	
☒ long				bigint	2	
☒ decimal		*-12	0	numeric(13)	3	
☒ decimal			0	numeric(%precision)	4	
☒ decimal		*-12		numeric(13,%scale)	5	
☒ decimal				numeric(%precision...	6	
☒ char	2-63			varchar(64)	7	
☒ char	64-*			varchar(%length)	8	
☒ char				char(%length)	9	
☒ varchar	*-63			varchar(64)	10	
☒ varchar				varchar(%length)	11	
☒ date				date	12	
☒ time				time	13	
☒ datetime				datetime	14	

%scale

In the example below, when converting a decimal in SQL Server to SQL Server format, we follow the following steps in the given order:

- ❑ If the decimal has a scale of zero, the **data type** will become `numeric(%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 `numeric(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 `numeric(%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 `numeric(8,3)`.

Maintain Data Type Mappings

Data Type Mapping Set:
Standard SQL Server to SQL Server

Description:
Standard SQL Server to SQL Server 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.

Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ bit				bit	0	
☒ tinyint				tinyint	1	
☒ tinyint identity				tinyint	2	
☒ smallint				smallint	3	
☒ smallint identity				smallint	4	
☒ int				int	5	
☒ int identity				int	6	
☒ bigint				bigint	7	
☒ bigint identity				bigint	8	
☒ numeric			0	numeric(%precision)	9	
☒ numeric			1-*	numeric(%precision,%scale)	10	
☒ decimal			0	numeric(%precision)	11	
☒ decimal			1-*	numeric(%precision,%scale)	12	
☒ smallmoney				smallmoney	13	
☒ money				money	14	
☒ float				float	15	
☒ real				real	16	
☒ date				date	17	
☒ smalldatetime				smalldatetime	18	


OK Cancel

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 number in Oracle to SQL Server format, we follow the following steps in the given order:

- ☐ If the number is 1 or 2 digits long with a scale of zero, the **data type** will become tinyint.
- ☐ If the number is 3 or 4 digits long with a scale of zero, the **data type** will become smallint.
- ☐ If the number is 5-9 digits long with a scale of zero, the **data type** will become int.
- ☐ If the number is 10-18 digits long with a scale of zero, the **data type** will become bigint.
- ☐ If the number is 19 or greater digits long with a scale of zero, the **data type** will become numeric(%precision); where we substitute the number of digits in the number for the variable '%precision'. Thus if the number is made up of 22 digits, the resulting data type will be numeric(22).
- ☐ If however, the scale is not zero and there are digits to the right of the decimal point in the number, then the **data type** will become numeric(%precision,%scale); where we substitute both the number of digits in the number for the variable '%precision' and the number of digits to the right of the decimal point for the variable '%scale'. Thus if the number is made up of 8 digits and has 3 digits after the decimal point (example 12345,678), the resulting data type will be numeric(8,3).



Maintain Data Type Mappings

Data Type Mapping Set:
Standard Oracle to SQL Server

Description:
Standard Oracle to SQL Server 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.

Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation
☒ number		1-2	0	tinyint	0	
☒ number		3-4	0	smallint	1	
☒ number		5-9	0	int	2	
☒ number		10-18	0	bigint	3	
☒ number		19-*	0	numeric(%precision)	4	
☒ number			1-*	numeric(%precision,%scale)	5	
☒ binary_float				float	6	
☒ binary_double				float	7	
☒ float				float	8	
☒ date				datetime	9	
☒ timestamp				datetime	10	
☒ timestamp with ti...				datetime	11	
☒ timestamp with lo...				datetime	12	
☒ interval year to ...				varchar(20)	13	
☒ interval day to se...				varchar(20)	14	
☒ varchar2				varchar(%precision)	15	
☒ nvarchar2				nvarchar(%precision)	16	
☒ long				text	17	
☒ clob				varchar(max)	18	

OK Cancel

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'.

Maintain Data Type Mappings

Data Type Mapping Set: Custom SQL Server to SQL Server [New... Copy... Delete]

Description: Custom SQL Server to SQL Server data type mapping.

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation	Format
☒ datetimeoffset		28		datetimeoffset(1)	36		
☒ datetimeoffset		29		datetimeoffset(2)	37		
☒ datetimeoffset		30		datetimeoffset(3)	38		
☒ datetimeoffset		31		datetimeoffset(4)	39		
☒ datetimeoffset		32		datetimeoffset(5)	40		
☒ datetimeoffset		33		datetimeoffset(6)	41		
☒ datetimeoffset		34		datetimeoffset	42		
☒ timestamp				timestamp	43		
☒ varchar		1-2		varchar(%precision)	44	COALESCE('%table_name.%column_name','U')	
☒ varchar		3-6		varchar(%precision)	45	COALESCE('%table_name.%column_name','UNK')	
☒ varchar		7-*		varchar(%precision)	46	COALESCE('%table_name.%column_name','UNKNOWN')	
☒ nvarchar		0		nvarchar(max)	47		
☒ nvarchar		1-*		nvarchar(%precision)	48		
☒ text				text	49		
☒ ntext				ntext	50		
☒ varbinary	0			varbinary(max)	51		
☒ image				image	52		
☒ uniqueidentifier				char(36)	53		
☒ xml				xml	54		

[New... Copy Edit... Delete... Move Up Move Down]

[OK Cancel]

%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).

Maintain Data Type Mappings

Data Type Mapping Set:
Custom SQL Server to SQL Server

Description:
Custom SQL Server to SQL Server data type mapping.

Data Type Mappings:

From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation	Format
☒ datetimeoffset		28		datetimeoffset(1)	36		
☒ datetimeoffset		29		datetimeoffset(2)	37		
☒ datetimeoffset		30		datetimeoffset(3)	38		
☒ datetimeoffset		31		datetimeoffset(4)	39		
☒ datetimeoffset		32		datetimeoffset(5)	40		
☒ datetimeoffset		33		datetimeoffset(6)	41		
☒ datetimeoffset		34		datetimeoffset	42		
☒ timestamp				timestamp	43		
☒ varchar	1-10			date	44	20131212	
☒ varchar	11			date	45	TO_DATE(%table_name.%column_name,%format)	'YYYYMMDD'
☒ varchar	12-*			date	46	20131212	
☒ nvarchar		0		nvarchar(max)	47		
☒ nvarchar		1-*		nvarchar(%precision)	48		
☒ text				text	49		
☒ ntext				ntext	50		
☒ varbinary	0			varbinary(max)	51		
☒ image				image	52		
☒ uniqueidentifier				char(36)	53		
☒ xml				xml	54		

OK Cancel

%prompt

In the example below, we use the following transformation on columns of type varchar, to replace any null values with a value entered by the user.

- 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 the value entered by the user when prompted.

Maintain Data Type Mappings

Data Type Mapping Set: Custom SQL Server to SQL Server

Description: Custom SQL Server to SQL Server data type mapping.

Data Type Mappings:


From Data Type	From Length	From Precision	From Scale	To Data Type	Order	Transformation	Format
☒ datetimeoffset		28		datetimeoffset(1)	36		
☒ datetimeoffset		29		datetimeoffset(2)	37		
☒ datetimeoffset		30		datetimeoffset(3)	38		
☒ datetimeoffset		31		datetimeoffset(4)	39		
☒ datetimeoffset		32		datetimeoffset(5)	40		
☒ datetimeoffset		33		datetimeoffset(6)	41		
☒ datetimeoffset		34		datetimeoffset	42		
☒ timestamp				timestamp	43		
☒ varchar				varchar(%precision)	44	COALESCE(%table_name.%column_name,%prompt)	
☒ nvarchar		0		nvarchar(max)	45		
☒ nvarchar		1-*		nvarchar(%precision)	46		
☒ text				text	47		
☒ ntext				ntext	48		
☒ varbinary	0			varbinary(max)	49		
☒ image				image	50		
☒ uniqueidentifier				char(36)	51		
☒ xml				xml	52		
☒ (New)					53		

OK Cancel

Chapter 37

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.




In This Chapter

Properties	1142
Change Column(s)	1144
Add Column	1147
Duplicate Column	1149
Delete Column	1151
Re-space Order Number	1152
Impact	1154
Send Columns to Another Object	1156

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.

Stage Table Column stage_forecast.product_code

Properties	
Transformation	

A _i	
General	
Table Name	stage_forecast
Column Name	product_code
Business Display Name	product code
Column Description	
Physical Definition	
Column Order	40
Data Type	integer
Null Values Allowed	True
Default Value	
Meta Definition	
Format	#,##0
Numeric	True
Additive	True
Attribute	False
Business Key	True
Key Type (0,A,B,C,...)	A
Source Details	
Source Table	load_forecast
Source Column	product_code
Transformation	
Join	False

Column Name
Database-compliant name of the column.
Dialog Opening Value: product_code

Note: There is a variation in column Properties, depending on the object type.

For Dimension tables, see **Dimension Column Properties** (on page 327).

For Stage tables, see **Stage Table Column Properties** (on page 363).


For Data Store tables, see **Data Store Column Properties** (on page 399).

For Normalized tables, see **Normalized Table Column Properties** (on page 428).


For Fact tables, see **Fact Table Column Properties** (on page 479).

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 **checkboxes** on the left. Each checkbox, 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.

Change Columns

Select the changes to be applied. Note that only the selected columns will have the selected changes applied.

Apply Changes to:

- ☒ Column Source Table
- ☒ Data Type
- ☐ Nulls
- ☐ Format String
- ☐ Numeric
- ☐ Additive
- ☐ Attribute
- ☐ EUL

Original Value:

- (Selected)
- (All)
- (Selected)
- (Empty)
- stage_customer_sales2
- stage_customer_ssis
- stage_customer_xml

New Value:

- select or enter a new source table
- select or enter a new data type

☐ Update matched columns only

Apply Close Help

- ☐ 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.

Change Columns

Select the changes to be applied. Note that only the selected columns will have the selected changes applied.

Apply Changes to:

- ☒ Column Source Table
- ☒ Data Type
- ☐ Nulls
- ☐ Format String
- ☐ Numeric
- ☐ Additive
- ☐ Attribute
- ☐ EUL

Original Value:

- (Selected)

New Value:

- (Empty)
- agg_product_group_sales
- agg_product_sales
- agg_sa_product
- dim_currency_rate
- dim_customer
- dim_date
- dim_forecast_date
- dim_order_date
- dim_product
- dim_sales_kpi
- dim_ship_date
- dim_state
- ds_customer
- ds_customer_sales2
- ds_customer_ssis
- ds_customer_xml
- dss_fact_table
- dss_source_system
- fact_forecast
- fact_order_detail
- fact_order_vs_forecast
- fact_sales_kpi
- load_address
- load_address_sales2
- load_currency_rate
- load_customer
- load_customer_sales2
- load_customer_ssis
- load_customer_xml


☐ Update matched columns only

Close Help

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**.

Properties

Transformation

Stage Table Column

2

General

Table Name

stage_forecast

Column Name

Business Display Name

Column Description

Physical Definition

Column Order

0

Data Type

Null Values Allowed

False

Default Value

Meta Definition

Format

Numeric

False

Additive

False

Attribute

False

Business Key

False

Key Type (0,A,B,C,...)

Source Details

Source Table

(Empty)

Source Column

(Empty)

Transformation

Join

False

Column Name

Database-compliant name of the column.
Dialog Opening Value:


OK

Cancel

Help

Duplicate Column

To copy a column, right-click on one of the columns in the middle pane and select **Duplicate Column**.



Change the **Column Name** and the **Business Display Name** and any other properties to define a new column and click **OK**.

Properties

Transformation

Stage Table Column

2

General

Table Name

stage_customer

Column Name

customer_category_code

Business Display Name

customer_category_code

Column Description

The full name of the customer. Forms a hierarchy with city and state.

Physical Definition

Column Order

20

Data Type

varchar(45)

Null Values Allowed

True

Default Value

Meta Definition

Format

Numeric

False

Additive

False

Attribute

True

Business Key

False

Key Type (0,A,B,C,...)

Source Details

Source Table

load_customer

Source Column

name

Transformation

Join

False

Column Name

Database-compliant name of the column.
Dialog Opening Value: name


OK

Cancel


Help

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.

Stage Table Column stage_customer.city

Properties
Transformation

General

Table Name	stage_customer
Column Name	city
Business Display Name	city
Column Description	The city in which the customer sells shipped product.Forms a hierarchy with code and state.

Physical Definition

Column Order	40
Data Type	varchar(30)
Null Values Allowed	True
Default Value	

Meta Definition

Format	
Numeric	False
Additive	False
Attribute	True
Business Key	False
Key Type (0,A,B,C,...)	

Source Details


Source Table	load_customer
Source Column	city
Transformation	
Join	False

Column Order
Numeric value that controls the relative order of columns in the database create statement.
Dialog Opening Value: 40


<- Update Update -> OK Cancel Help

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 801).



To display a Track Forward Report, right-click on a column in the middle pane and select **Impact > Track Forward Report**.

Column Name	Display Name	Data Type	Source Table	Source Column
customer_code	customer_code	numeric(6)	load_customer	code
customer_name	customer_name	varchar(45)	load_customer	name
customer_address	customer_address	varchar(45)	load_customer	address
customer_city	customer_city	varchar(45)	load_customer	city
customer_state	customer_state	varchar(45)	load_customer	state
dss_update_time	dss_update_time	timestamp		dss_update_time


- Properties
- Change Column(s)
- Add Column
- Duplicate Column
- Delete Column
- Respace Order Number
- Impact
 - Track Back Report
 - Track Forward Report
- Send Columns To Another Object

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 804)

[illegible]

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**.

stage_forecast Columns				
Column Name	Display Name	Data Type	Source Table	Source Column
dim_customer_key	dim customer ...	integer	dim_customer	dim_customer_key
dim_product_key	dim product key	integer	dim_product	dim_product_key
dim_date_key	dim time key	integer	dim_date	dim_date_key
customer_code	customer code	integer	load_forecast	customer_code
forecast_quantity	forecast quanti...	integer	load_forecast	forecast_quantity
product_code	product code	integer	load_forecast	product_code
forecast_sales_value	forecast sales v...	numeric(13,2)	load_forecast	forecast_sales_value
forecast_date	forecast date	datetime	load_forecast	forecast_date
dss_update_time	dss update time	datetime		dss_update_time

The columns will be copied to the destination table.

stage_forecast Columns				
Column Name	Display Name	Data Type	Source Table	Source Column
dim_customer_key	dim customer ...	integer	dim_customer	dim_customer_key
dim_product_key	dim product key	integer	dim_product	dim_product_key
dim_date_key	dim time key	integer	dim_date	dim_date_key
customer_code	customer code	integer	load_forecast	customer_code
forecast_quantity	forecast quanti...	integer	load_forecast	forecast_quantity
product_code	product code	integer	load_forecast	product_code
forecast_sales_value	forecast sales v...	numeric(13,2)	load_forecast	forecast_sales_value
customer_name	customer_name	varchar(45)	stage_customer	customer_name
customer_address	customer_addr...	varchar(40)	stage_customer	customer_address
forecast_date	forecast date	datetime	load_forecast	forecast_date
dss_update_time	dss update time	datetime		dss_update_time

Chapter 38

Database Function Sets

In This Chapter

- Using Database Function Sets..... 1160
- Maintaining Database Function Sets 1163
- Loading Database Functions 1186
- Exporting Database Function Sets..... 1189

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 on the **Tools** menu.

Column Transformation Properties Dialog

Stage Table Column stage_customer.code

Properties
Transformation

Target: code Paste

Source: load_customer.code Paste

Column Transformation Code (must execute within a SQL SELECT statement).

Function Set: SQL Server Standard Functions

- Function
 - Number
 - String
 - Date
 - Conversion
 - Aggregate
 - Operator
 - User
- Available Columns

☐ Word Wrap Displayed Code


Function Syntax:

Function Desc.:

<- Update Update -> OK Cancel Help

A drop-down list allows 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.

Transformation Definition Dialog



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 Dialog

The screenshot shows the 'Connection DataWarehouse' dialog box. On the left is a 'Properties' sidebar with a 'Notes' section. The main area is divided into several sections: 'General', 'Source System', 'Database Credentials', 'Other', and 'When Connection is an OLAP Data Source'. The 'Default Transform Function Set' is highlighted in the 'Other' section. Below this, a list of function sets is shown: Default DB2, Default Netezza, Default Oracle, Default SQL Server, Default Teradata, Oracle Standard Functions, and SQL Server Standard Functions. The 'Default Transform Function Set' is selected.

Section	Property	Value
General	Connection Name	DataWarehouse
	Connection Type	Database
	Database Type	(local)
	ODBC Data Source Name (DSN)	WslWarehouse
	Data Warehouse Connection Indicator	☒
Source System	Database ID	WslWarehouse
	Database Link Name	
	Provider Name	SQLOLEDB
Database Credentials	Extract User ID	
	Extract User Password	
	Administrator User ID	
	Administrator User Password	
Other	New Table Default Load Type	Database link load
	SSIS Connection String	Provider=SQLOLEDB.1;Integrated Security=SSPI;Persist Secur...
	Data Type Mapping Set	(Default)
	Default Transform Function Set	(Default)
	When Connection is an OLAP Data Source	(Default)
When Connection is an OLAP Data Source	OLAP Connection String	Default DB2
	Connection Provider/Driver	Default Netezza
	Data Warehouse Server	Default Oracle
	Data Warehouse Database ID	Default SQL Server
		Default Teradata
		Oracle Standard Functions
		SQL Server Standard Functions


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

OK Cancel Help

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 will be selected in the transformation dialogs above.

Maintaining Database Function Sets


To maintain the database function sets, select **Tools/Database Functions/Maintain Database Functions...**



The **Maintain Database Functions** dialog is displayed.

[illegible]

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 1166)


To copy a database function set, see **Copying a Database Function Set** (on page 1169)

To edit a database function set, see **Editing a Database Function Set** (on page 1172)


To delete a database function set, see **Deleting a Database Function Set** (on page 1184)

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.



Enter a **Function Set Name** and select a **Database** from the drop-down list. Click **OK**.



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 QAD 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.



Enter the new **Function Set Name** and select a **Database** from the drop-down list. Click **OK**.



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 QAD 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.




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.



Enter the details for the new function and click **OK**.



The image shows a 'New Function' dialog box with a light blue title bar and a red close button. The dialog is divided into several sections: 'Function Name' with a text input field, 'Group' with a dropdown menu showing 'Aggregate', 'Description' with a large text area, 'Syntax' with a text input field, and 'Model' with a large text area. At the bottom, there are three buttons: 'Default Column', 'Clear Default Column', and 'OK', followed by a 'Cancel' button. A help text line reads: 'To set the Default Column, select it in the Model box above and click the Default Column button.'

New Function

Function Name:

Group:

Description:

Syntax:

Model:

To set the Default Column, select it in the Model box above and click the Default Column 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.



The screenshot shows a software interface for defining a function model. It consists of a text input area labeled 'Model:' and a button labeled 'Default Column'.

The 'Model:' text area contains the text `CHAR(integer)`. The word `integer` is highlighted in red, indicating it is the current default column.


The 'Default Column' button is located at the bottom of the 'Model:' text area.

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.



Enter a new **Function Name** and change any other details; then click **OK**.

New Function

Function Name:
Absolute Value (ABS)2

Group:
Number

Description:
Returns the absolute value of the number. Essentially converts all values to positive.
Example: ABS(load_sales.tax)

Syntax:
ABS(numeric_column)

Model:
ABS(*numeric_column*)

Default Column

To set the Default Column, select it in the Model box above and click the Default Column button.


Clear Default Column

OK

Cancel

To edit an existing function in the database function set:

Select the function and then click on the **Edit** button on the right.



Change any fields as required and then click **OK**.

Edit Function

Function Name:
Round (ROUND)

Group:
Number

Description:
Returns the value rounded to a given number of digits of precision.
Example: ROUND(12.399,1) = 12.4

Syntax:
ROUND(numeric_column,precision)

Model:
ROUND(numeric_column,precision)

Default Column

To set the Default Column, select it in the Model box above and click the Default Column button.

Clear Default Column

<- Update

Update ->


OK

Cancel


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.




Click **Yes** to confirm the delete.




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.




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.




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.



Note: The **Delete** button is disabled for standard function sets.


Click **Yes** to confirm the delete.




When all functions are deleted, the function set ceases to exist.

Loading Database Functions

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 Data Warehouse Designer expects the xml files to be in **ProgramData\QAD\Work**.




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, Data Warehouse Designer exports the xml file to **ProgramData\QAD\Work**, but this can be changed. Change the File name if necessary and click **Save**.



Index

A

Aborting a Job 688, 723
 Add Column 298, 345, 444, 508, 1147
 Add Index 753
 Add Table 605
 Adding a Dimension Hierarchy 337
 Adding Objects to Projects 21
 Adding Projects to Groups 25
 After Load Transformations 613
 Aggregate Table Column Properties 494
 Aggregate Table Column Transformation 499
 Aggregation 487
 All Column Transformations 799
 Altering Metadata 893
 Altering Tables 894, 898
 Alternative Invocation Methods 909
 Analysis Server 2000 157, 189, 216
 Analysis Services 2, 14, 107, 157, 158, 308, 520, 526, 530, 531, 532, 533, 538, 540, 547, 550, 553, 557, 561, 564, 567, 568, 572, 576, 578, 581, 583, 586
 Analysis Services OLAP Cubes 519
 Ancillary Columns 1064, 1072
 API 2, 295, 326, 762, 904, 912, 1085, 1090
 Application Creation 861
 Application Loading 866
 Applications 860, 861, 866, 870, 874
 Applying Changes 624
 Applying Changes to Dependant Transformations 624
 Artificial Keys 298, 323, 349
 ASCII 2, 277, 295, 591, 595, 881, 1022, 1034

B

Backing Up and Restoring Metadata 859, 877
 Backup using DB Routines 878
 Batch Documentation Creation 765
 Browse Source System 107
 Browsing a Connection 109
 Building a Data Store Object 378, 380
 Building a Dimension 298, 300, 381, 409
 Building a Normalized Table 406, 408
 Building an Export Object 592
 Building the Stage Table 345, 346
 Business Keys 298, 324, 327, 341, 349, 362, 363, 435, 479, 488, 492, 754, 1079, 1096

C

Calculated Members 526, 532, 547, 550
 Callable Routines 106, 295, 903, 904

Callable Routines API 904
 Callable Routines Common Input 907, 916, 924, 930, 935, 940, 948, 956, 962, 967, 972, 977, 985, 990, 996, 1016, 1022, 1028, 1034, 1040, 1046, 1049
 Callable Routines Invocation 908
 Callable Routines Names Qualifier 907
 Callable Routines per RDBMS 905
 Calling a Batch File from a Script 666
 Case Conversion 53
 Change an Aggregates Default Look Back Days 491
 Change Column(s) 1144
 Changing a Connection's Properties 112
 Changing a Re-usable Transformation 623
 Changing Load Connection and Schema 238, 295
 Changing OLAP Cubes 583
 Choosing the Best Load Method 234
 Code Editor 85
 Column Attributes 598
 Column Comments 791
 Column Context Menu 1141
 Column Properties 290, 300, 327, 363, 479, 508, 604, 605, 1006, 1064, 1079, 1083
 Column Reports 797
 Column Track-Back 801, 1154
 Column Track-Forward 804, 1155
 Column Transformation Properties 605, 626
 Column Transformations 603, 604
 Columns 2, 14, 241, 246, 265, 273, 277, 285, 289, 290, 298, 300, 308, 326, 327, 336, 337, 339, 345, 346, 349, 362, 363, 435, 440, 444, 445, 446, 447, 458, 461, 463, 468, 479, 488, 492, 508, 514, 526, 532, 533, 595, 598, 605, 608, 609, 613, 624, 754, 761, 762, 767
 Columns without Comments 798
 Combined Fact Table 439, 440, 463, 468, 472
 Comments 35, 285, 447, 636, 637, 791, 798, 808, 898, 912, 1013
 Common Parameters 916, 924, 930, 935, 940, 948, 956, 962, 967, 972, 977, 985, 990, 996, 1016, 1022, 1028, 1034, 1040
 Compare Meta-data Indexes to Database 855
 Compare Meta-data Repository to another 852
 Compile Procedure 246
 Completed Re-usable Transformation 622
 Confirmation Prompts 86
 Connections 107, 273
 Copy Column 1079, 1083
 Copy Object 1063, 1064, 1072
 Copying a Data Type Mapping Set 1109, 1114

Copying a Database Function Set	1165, 1169
Creating a Custom View	514
Creating a Database Link	113, 126, 168
Creating a Dimension View	502
Creating a Job	680, 689
Creating a Job from a Diagram	785
Creating a New Data Type Mapping Set	1109, 1110
Creating a New Database Function Set	1165, 1166
Creating a New Re-usable Transformation	616, 623
Creating a Partitioned Fact Table in DB2	472
Creating a Partitioned Fact Table in Oracle	468
Creating a Partitioned Fact Table in SQL Server	463
Creating a Project from a Diagram	789
Creating a Snowflake	341
Creating an Aggregate Summary Table	492
Creating an Aggregate Table	488
Creating an Application from a Diagram	787
Creating an Application from a Job	729
Creating an OLAP Cube on the Analysis Services Server	531
Creating and Loading Applications from the Command Line	874
Creating Detail Fact Tables	435
Creating Documentation	761, 762
Creating KPI Fact Tables	445
Creating Rollup Fact Tables	440
Creating Slowly Changing Fact Tables	461
Creating Snapshot Fact Tables	458
Creating the KPI Dimension	444
Creating Views	510
Custom Procedures	14, 814, 816, 885

D

Data Store Artificial Keys	397
Data Store Column Properties	399, 1143
Data Store Object Column Transformations	403
Data Store Object Overview	378
Data Store Objects	377
Data Type Mapping Examples	1122, 1134
Data Type Mappings	236, 1105
Data Warehouse	116, 161, 193, 1087
Data Warehouse Testing	870
Database	115, 160, 192
Database Error Messages	37
Database Function Sets	1159
Database Functions	614
Database Link During Load Transformations	609
Database Link Load	237
Database Link Load - Current	119, 164
Database Link Load - Linked Server	125, 167
Database Link Load - Other DB	122

Database Link Load - Properties	238
Database Link Load - Source Mapping	106, 241
Database Monitoring	733
DB2914, 920, 923, 928, 934, 939, 946, 954, 960, 966, 971, 976, 983, 989, 994, 1000, 1003, 1008, 1012, 1015, 1020, 1026, 1032, 1038, 1044	
DB2 Connections	191
Default Settings	43
Default Update Procedure Options	67
Define the Transformation Model	620
Defining an OLAP Cube	526
Defining the Data Source for the OLAP Cube	118, 163, 195, 521
Delete Column	1151
Delete Metadata	754
Deleting a Data Type Mapping Set	1109, 1128
Deleting a Database Function Set	1165, 1184
Deleting a Job	688, 718
Deleting Job Logs	688, 719
Dependencies	671, 677, 679, 699, 767, 861, 904, 956
Dependency Diagram	768, 779
Design 2, 11, 14, 300, 345, 346, 520, 533, 628, 878, 940, 948, 1086	
Design Introduction	11
Detail Fact Tables	433, 434, 435, 439, 443, 458
Diagram window	34
Diagrams	2, 34, 87, 327, 363, 479, 767
Dimension Artificial Keys	323
Dimension Column Properties	327, 1143
Dimension Column Transformations	335
Dimension Get Key Function	324
Dimension Hierarchies	336
Dimension Initial Build	298, 326
Dimension Language Mapping	343
Dimension View Language Mapping	516
Dimension Views of a Specified Dimension	796
Dimension-Fact Matrix	794
Dimensions	11, 14, 116, 297, 298, 300, 308, 324, 326, 327, 340, 341, 345, 346, 349, 433, 434, 435, 439, 440, 445, 447, 487, 492, 501, 502, 510, 526, 561, 568, 572, 576, 578, 581, 679, 767, 791, 794, 1072, 1096
Dimensions Overview	298
Documentation 2, 14, 219, 226, 228, 229, 238, 327, 339, 363, 479, 514, 598, 614, 761, 762, 766, 767, 861, 904, 912, 1079	
Documentation and Diagrams	761
Documentation Roadmap	9
Drill Through Columns	526, 533, 557
Drop Table	327, 363, 479, 861, 1079
DSS Columns	63
DSS Tables	62
Duplicate Column	1149
During Load Transform	608, 609, 612

E

Editing a Data Type Mapping Set	1109, 1118
Editing a Database Function Set	1165, 1172
Editing a Job	688, 699
Editing Dependencies	688, 711
Editing Jobs	699
Editing Tasks in a Job	688, 703
Enter Re-usable Transformation Metadata	619
Existing Warehouse	1085
Export Column Properties	598
Export Middle Pane Output	31, 38
Export Objects	107, 591, 592, 598
Export Objects by Connection	810
Exporting Data	591
Exporting Data Type Mapping Sets	1132
Exporting Database Function Sets	1189
External Load	289
External Source Table/files	813

F

Fact Table Column Properties	479, 1143
Fact Table Column Transformations	484
Fact Table Language Mapping	485
Fact Tables	433
File Action Example	271
File Actions	270
File Attributes	595
File During Load Transformations	612
File Load - Source Mapping	277
File Wild Card Load	284
Filegroup	220, 866, 888
Flat File Load	273, 277
Fragmentation	837
FTP	881
Functions	14, 520, 603, 605, 612, 613, 614, 615, 620, 627, 631, 662, 877, 901, 904

G

General	66, 83, 91
Generated Code Statistics	69
Generating the Data Store Update Procedure	378, 383, 384
Generating the Detail Fact Update Procedure	436, 459
Generating the Dimension Update Procedure	298, 304, 307, 308
Generating the KPI Fact Update Procedure	453
Generating the Normalized Update Procedure	381, 406, 409, 411, 412
Generating the Partitioned Fact Update Procedure	476
Generating the Rollup Fact Update Procedure	441, 459

Generating the Snapshot Fact Update Procedure	459

Generating the Staging Update Procedure	106, 345, 348, 349, 438, 478
Get Key	324, 502, 510, 1090, 1095
Global Naming of Indexes	55
Global Naming of Key Columns	57, 398, 427
Global Naming of Procedures	58
Global Naming of Tables	54

H

Halting a Job	688, 722
Handling Missing Source Columns	241, 290
Hierarchy	308, 336, 337, 339, 533, 578, 581
Host Script Exports	601
Host Script Load	283, 1088

I

Impact	1154
Impact Diagram	768, 778
Importing	859, 867, 1072
Importing a Data Model	378, 406, 1072
Importing Language Files	869
Importing Object Metadata	867
Index Definition	754
Index Reports	817
Index Type	73
Indexes	753
Inserting a Copy of a Job	688, 716
Inspecting and Modifying Advanced Cube Properties	530
Integrating	1085, 1087
Integrating QAD Data Warehouse Designer into an Existing Warehouse	1085
Integrating, Host Scripts	1087, 1088
Integrating, Procedures	1071, 1087, 1093, 1094
Integrating, QAD Tables	1096
Integrating, Questions	1087, 1090
Integrating, Selecting a Table Type	1087, 1089
Integrating, Views	1095
Integration Services Load	136, 153
Introduction to Callable Routines	904
Invocation using Named Parameters	909
Invocation via ODBC	910
Invocation via the Command-Line	911

J

Job Dependencies	671, 679, 767, 904, 956
Job Monitoring	735
Job Performance History	834
Job Reports	826
Jobs	2, 32, 673, 674, 679, 699, 732, 735, 826, 827, 828, 940, 948, 962, 972
Jobs with an Object	828

Join Information Maintenance	1097
Joins Diagram	768, 776

K

KPI Dimensions	447
KPI Fact Tables	433, 443
KPI Setup	447

L

Language Mapping	343, 516, 533, 568, 869
Links Diagram	768, 777
List Database Tables not in the Meta-data	847
List Duplicate Business Key Columns	856
List Meta-data Tables not in the Database	846
List Procedures not related to a Table	851
List Tables with no related Procedures or Scripts	849
Load Drag and Drop	235
Load Table Column Transformations	293, 294, 604, 606, 608
Load Table Properties Screen	238, 240
Load Table Source Mapping Screen	241, 243, 277, 280
Load Table Source Mapping Screen (Native)	265, 267
Load Table Transformations	235, 241, 293
Load Tables by Connection	809
Loaded or Imported Objects	824
Loading an Unload	888
Loading Data	233, 259, 1071
Loading data into Data Warehouse Designer load tables using SSIS	141, 250
Loading Data Type Mapping Sets	1130
Loading Database Functions	1186
Local Naming Conventions	90
Local Naming of Indexes	94
Local Naming of Key Columns	93
Local Naming of Tables	92
Local Paths	95
Login Checks	1101
Look and Feel	82, 101

M

Main Builder Window	29
Maintaining Data Type Mapping Sets	1108
Maintaining Database Function Sets	1163
Mapping	14, 241, 243, 246, 264, 265, 267, 285, 290, 343, 363, 502, 508, 516, 533, 861, 1079, 1087, 1089
Measures	11, 433, 435, 440, 445, 447, 458, 461, 487, 488, 492, 520, 532, 533, 538, 540, 547
Members	107, 526, 538, 540, 547
Microsoft Analysis Server 2005+	158, 190, 217
Migrating	1064

Migrating the Data Warehouse Database Platform	1063, 1064

Modified	14, 290, 298, 308, 323, 324, 463, 520, 624, 699, 762, 767, 815, 818, 821, 852, 878, 885, 1013, 1088, 1090

Modified Indexes	818
Modified Objects (excluding indexes)	821
Modified Procedures	815
Monitoring	732, 733, 735
Monitoring the Database and Jobs	732
Moving Projects within Groups	25, 27

N

Native	2, 17, 233, 246, 264, 265, 267, 324, 605, 608, 766, 877
Native ODBC Based Load	264
Native ODBC Based Load - Source Mapping	265
Native ODBC Load	145, 177, 204
Non-Dimension Object Views	510
Normalized Table Artificial Keys	407, 426
Normalized Table Column Properties	428, 1143
Normalized Table Column Transformations	432
Normalized Tables	405
Normalized Tables Overview	406

O

Object Color	51
Object End User Setting	50
Object Performance History	832
Object Reports	819
Object Type Availability	47
Object Type Names	48
Object Type Ordering	49
Object Types	14
Object-Job Matrix	827
Objects and Windows	13
Objects Checked-out	823
Objects-Project Matrix	820
ODBC	2, 14, 107, 119, 122, 125, 144, 149, 176, 203, 233, 246, 248, 264, 265, 295, 533, 608, 609, 679, 888, 1064
ODBC Based Load	248
ODBC Load	149
ODBC Load - Direct Path OCI	185, 212
ODBC Load - Single Row Inserts	181, 208
OLAP Cube Actions	557
OLAP Cube Calculations	550
OLAP Cube Dimensions	561
OLAP Cube Key Performance Indicators	553
OLAP Cube Measure Group Dimensions	564
OLAP Cube Measure Group Partitions	544
OLAP Cube Measure Group Processing/Partitions	540
OLAP Cube Measure Groups	538

OLAP Cube Measures	547
OLAP Cube Objects	532
OLAP Cube Properties	533
OLAP Dimension Attribute Relationships	576
OLAP Dimension Attributes	572
OLAP Dimension Hierarchies	578
OLAP Dimension Objects	567
OLAP Dimension Overview	568
OLAP Dimension User Defined Hierarchy Levels	581
OLAP Dimension-Cube Matrix	795
OLAP Retrofitting	586
Operational Reports	831
Oracle2, 35, 159, 160, 161, 164, 167, 170, 174, 176, 177, 181, 185, 189, 190, 220, 223, 238, 241, 246, 248, 264, 273, 277, 285, 290, 295, 298, 300, 308, 323, 324, 326, 327, 346, 349, 362, 363, 444, 447, 462, 468, 479, 487, 488, 526, 532, 533, 578, 595, 612, 614, 732, 740, 791, 860, 866, 878, 881, 888, 901, 914, 919, 922, 927, 933, 938, 945, 953, 959, 965, 970, 975, 982, 988, 993, 999, 1003, 1008, 1012, 1015, 1019, 1025, 1031, 1037, 1043, 1079, 1090	
Oracle Connections	159
Organizing Objects	17
Other	99
Outputs	39, 97
Overview	1, 520
Overview of QAD Data Warehouse Designer 2	

P

Parameters105, 277, 324, 349, 354, 390, 418, 447, 463, 468, 472, 491, 595, 639, 643, 647, 653, 658, 733, 735, 740, 912, 916, 924, 930, 935, 940, 948, 956, 962, 967, 972, 977, 985, 990, 996, 1001, 1006, 1010, 1013, 1016, 1022, 1028, 1034, 1040, 1090, 1094	
Partition 14, 238, 433, 462, 463, 468, 472, 531, 533, 538, 544, 754, 990	
Partitioned Fact Tables	433, 462
Post-Load	294, 679
Post-Load Procedures	294
Procedure Comparison	636
Procedure Comparisons	636
Procedure Compilation	637
Procedure Editing	631
Procedure Editor window	35
Procedure Generation	628
Procedure Loading and Saving	634
Procedure Properties	640
Procedure Reports	814
Procedure Running	638
Procedure Syntax	639

Procedures and Scripts	35, 614, 627
Promoting	859, 881
Promoting Between Environments	859
Properties	1142
Property Grids	88

Q

Query Data Warehouse Objects	857
Query Tool	761, 791
Query Tool Access	761, 791

R

Reading the Documentation	761, 766
Rebuilding	1085, 1086
Recompiling Procedures	901
Records that failed a Dimension Join	811
Refreshed	17
Re-Mapping	502, 508
Remote View Extract 120, 123, 126, 165, 168, 242, 246, 265	
Removing Objects from Projects	23
Removing Projects from Groups	26
Reports	793
Reset Columns in Job and Task View	744
Re-space Order Number	1152
Restarting a Job	688, 724
Restoring DB Backups	881
Re-Target	1064, 1077
Re-Targeting Source Tables	1071, 1077
Retro Copy Column Properties	1079
Retro Copy Column Properties Screen	1083
Retro Copy Column Transformations	1084
Retrofitting	586, 1063, 1064
Retrofitting an OLAP Object	586, 1063
Re-Usable Column Transformations	800
Re-usable Transformations	615
Rollup or Combined Fact Tables	433, 439

S

Scheduler2, 28, 32, 142, 144, 145, 149, 170, 230, 233, 248, 264, 273, 277, 283, 290, 308, 326, 349, 447, 520, 533, 591, 601, 627, 628, 638, 639, 643, 647, 653, 658, 663, 664, 673, 674, 677, 679, 699, 733, 735, 740, 753, 754, 767, 826, 827, 828, 829, 900, 930, 940, 948, 962, 967, 972, 977, 990, 1006, 1040, 1064, 1088, 1089, 1090, 1094	
Scheduler General	76
Scheduler Options	674
Scheduler States	677
Scheduler Statistics	77
Scheduler window	32
Scheduling a Data Warehouse Designer Host Script from 3D	745

Scheduling a Job 679, 690, 701
 Scheduling Scripts 669
 Schema 2, 11, 14, 34, 107, 122, 220, 230, 238,
 241, 265, 285, 295, 339, 340, 345, 433, 532,
 578, 740, 767, 881, 888, 901, 904, 912, 916,
 924, 930, 935, 940, 948, 956, 962, 967, 972,
 977, 985, 990, 996, 1001, 1006, 1010, 1013,
 1016, 1022, 1028, 1034, 1040, 1064, 1090
 Schema Diagram 768, 770
 Script Editing 662
 Script Generation 642
 Script Generation (UNIX/Oracle) 647
 Script Generation (Windows/DB2) 658
 Script Generation (Windows/Oracle) 653
 Script Generation (Windows/SQL Server) 643
 Script Syntax 664
 Script Testing 663
 Send Columns to Another Object 1156
 Set Merge Procedure 352, 368
 Settings - Check-Out and Check-In 64
 Settings - Code Generation 65
 Settings - Common 81
 Settings - Current Repository 100
 Settings - Documentation 78
 Settings - DSS Tables and Columns 59
 Settings - Global Naming Conventions 52
 Settings - Language Options 102, 537, 571
 Settings - Object Types 46
 Settings - Options 44
 Settings - Other 79
 Settings - Repository Identification 45
 Settings - Scheduler 75
 Settings - Storage 71
 Settings - User Preferences 80
 Settings - Versioning 74
 Setup of KPI Fact Tables 446
 Slowly Changing Dimension 14, 300, 308, 327,
 349, 363, 461, 479, 1079, 1096
 Slowly Changing Fact Tables 433, 460
 Snapshot Fact Tables 433, 457
 Snowflake 11, 340, 341
 Source Diagram 768, 772
 Specify the Name of the Transformation 617
 SQL 2, 107, 113, 116, 119, 122, 142, 157, 158,
 220, 222, 238, 241, 246, 248, 264, 265, 273,
 277, 285, 295, 298, 300, 308, 323, 324, 326,
 327, 335, 349, 362, 363, 367, 444, 447, 462,
 463, 479, 514, 526, 532, 533, 547, 578, 595,
 604, 605, 608, 609, 612, 613, 614, 615, 620,
 623, 634, 639, 642, 643, 647, 653, 664, 740,
 754, 791, 801, 857, 860, 861, 878, 881, 885,
 888, 912, 916, 924, 930, 935, 940, 948, 956,
 962, 967, 972, 977, 985, 990, 996, 1001, 1006,
 1010, 1013, 1016, 1022, 1028, 1034, 1040,
 1079

SQL Server 912, 917, 921, 925, 931, 936, 941,
 949, 957, 963, 968, 973, 979, 986, 991, 997,
 1001, 1004, 1007, 1010, 1014, 1017, 1023,
 1029, 1035, 1041
 SQL Server Connections 114
 SQL to return Scheduler Status 743
 SSIS Loader 249
 Stage Table Column Properties 363, 1143
 Stage Table Column Transformations 367
 Staging 345
 Stand Alone Scheduler Maintenance 740
 Starting a Job 688, 721
 Statistics 230, 231, 326, 362, 443
 Storage 219, 220, 222, 223, 225, 264, 300, 346,
 671, 754, 898

T

Table Documentation Fields 228
 Table Documentation Fields Screen 229
 Table Notes 226
 Table Notes Screen 227
 Table Properties 219
 Table Reports 807
 Table Statistics in DB2 230
 Table Statistics Screen 231
 Table Storage 72
 Table Storage Properties 220
 Table Storage Screen - DB2 220, 225
 Table Storage Screen - Oracle 220, 223
 Table Storage Screen - SQL Server 220, 222
 Tables without Comments 808
 Tablespace 220, 866, 888, 990
 Task Performance History 835
 Tasks of a Job 829
 Templates 671
 Teradata 913, 918, 922, 926, 932, 937, 943, 951,
 958, 964, 969, 974, 980, 987, 992, 998, 1002,
 1007, 1011, 1014, 1018, 1024, 1030, 1036,
 1042
 Transformations 2, 14, 246, 289, 290, 293, 294,
 295, 335, 349, 363, 367, 403, 432, 484, 499,
 514, 603, 604, 605, 608, 609, 612, 613, 615,
 616, 617, 619, 623, 624, 626, 679, 797, 801,
 885, 1079, 1084
 Transforming 345
 Troubleshooting 262
 Tuning the Staging Update Process 345, 362
 Tutorials 41
 Types of Diagrams 768

U

Unix 14, 107, 170, 233, 248, 273, 277, 283, 284,
 627, 640, 642, 647, 663, 673, 677, 679, 732,
 733, 735, 878, 881, 1087, 1088

UNIX	170, 196, 1087
Unloading Metadata	885
Upgrading Data Warehouse Designer	1099
Using a Maintained Hierarchy	339
Using Data Type Mapping Sets	1106
Using Database Function Sets	1160
Using Project Object Maintenance	21, 23, 24
Using Re-usable Transformations	626

V

Validate	839
Validate Index Status	844
Validate Load Table Status	843
Validate Meta-data	840
Validate Procedure Status	845
Validate Table Create Status	842
Validate Workflow Data	841
Validating Indexes	900
Validating Procedures	897
Validating Source (Load) Tables	896
Validating Tables	894
View Aliases	517
View Column Re-mapping	508
Views	501

W

Wild Cards	284, 924, 935
Windows	142, 174, 200, 1087
Windows and Panes	28
Working with Diagrams	782
Working with Jobs	687
Ws_admin_v Views	1045
Ws_admin_v_audit	1046
Ws_admin_v_dim_col	1047
Ws_admin_v_error	1049
Ws_admin_v_fact_col	1050
Ws_admin_v_fact_join	1053
Ws_admin_v_sched	1055
Ws_admin_v_task	1059
Ws_Api_Glossary	904, 905, 912
Ws_Connect_Replace	904, 905, 916
Ws_Job_Abort	904, 905, 921
Ws_Job_Clear_Archive	904, 905, 924
Ws_Job_Clear_Logs	904, 905, 930
Ws_Job_Clear_Logs_By_Date	904, 905, 935
Ws_Job_Create	904, 905, 940
Ws_Job_CreateWait	904, 905, 948
Ws_Job_Dependency	904, 905, 956
Ws_Job_Release	904, 906, 962
Ws_Job_Restart	904, 906, 967
Ws_Job_Schedule	904, 906, 972
Ws_Job_Status	904, 906, 977
Ws_Load_Change	904, 906, 985
Ws_Maintain_Indexes	759, 904, 906, 990

Ws_Version_Clear	904, 906, 996
WsParameterRead	1001
WsParameterReadF	106, 904, 906, 1004
WsParameterReadG	106, 905, 906, 1006
WsParameterReadJ	905, 906, 1010
WsParameterWrite	106, 905, 906, 1013
WsWrkAudit	905, 906, 1016
WsWrkAuditBulk	905, 906, 1022
WsWrkError	905, 906, 1028
WsWrkErrorBulk	905, 906, 1034
WsWrkTask	905, 906, 1040

X

XML	14, 233, 285
XML File load	285