



QAD Enterprise Applications

Training Guide

QAD Data Warehouse Designer

70-3123-6.5.6
Data Warehouse Designer Version 6.5.6
March 2012

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 © 2012 by QAD Inc.

DataWarehouseDesigner_TG_v656.pdf/mat

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

Contents

Tutorial 1 - Basic Star Schema Fact Table	5
1.1 Purpose and Roadmap	6
1.2 The First Step	8
1.2.1 Logging In	8
1.2.2 Switching Between Databases	10
1.3 Repository Defaults	10
1.4 Tablespace (FileGroup) Defaults	11
1.5 Table Name Defaults	12
1.6 Creating a Connection	13
1.7 Loading Source Tables	18
1.8 Building Dimensions	22
1.9 Creating Dimension Views	31
1.10 Defining the Staging Table	35
1.11 Including Dimension Links	38
1.12 Creating a Fact Table	48
1.13 Switching to Diagrammatic View	52
1.14 Producing Documentation	55
1.15 Data Store Objects (Optional)	56
Tutorial 2 - Rollup Fact Tables, ASCII File Loads, Aggregates	63
2.1 Purpose and Roadmap	64
2.2 Making a Connection to Windows	66
2.3 Loading Tables from Flat Files	70
2.4 Creating Stage Tables	75
2.5 Creating Fact Tables	77
2.6 Rollup/Combined Fact Table	79
2.7 Aggregate Tables	82
2.8 Creating a Customer Aggregate	84
Tutorial 3 - Scheduling and Dependencies	91
3.1 Purpose and Roadmap	92
3.2 Creating and Scheduling a Job	93
3.3 Adding Tasks	94
3.4 Task Dependencies	97
3.5 Editing a Scheduled Job	98
3.6 Job Results	100
3.7 Diagrammatic View for Jobs	102
Tutorial 4 - Complex Dimensions and Hierarchies	105
4.1 Purpose and Roadmap	106
4.2 Creating a Slowly Changing Dimension	107

4.3 Multiple Source Table Dimension	114
4.4 Creating a Dimension Hierarchy	125

Tutorial 5 - Analysis Services Cubes	129
---	------------

5.1 Purpose and Roadmap	130
5.2 Creating an OLAP Cube Object	131
5.3 Adding a Measure Group	155
5.4 Cube Connections for Other Databases	167

Tutorial 6 - KPI fact table	171
------------------------------------	------------

6.1 Purpose and Roadmap	172
6.2 The KPI Dimension	173
6.3 Defining the KPI Fact Table	174
6.4 Creating the KPI Fact Table and Procedure	181
6.5 The KPI Parameters	184
6.6 Defining a KPI	185
6.7 Defining the Activity KPI	189

Change Summary

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

Date/Version	Description	Reference
March 2012/6.5.6	Numerous updates throughout	
September 2011/QAD BI 3.5	Rebranded for QAD BI 3.5	--

C H A P T E R 1

Tutorial 1 - Basic Star Schema Fact Table

In This Chapter

1.1 Purpose and Roadmap.....	6
1.2 The First Step.....	8
1.3 Repository Defaults	10
1.4 Tablespace (FileGroup) Defaults	11
1.5 Table Name Defaults	12
1.6 Creating a Connection.....	13
1.7 Loading Source Tables.....	18
1.8 Building Dimensions	22
1.9 Creating Dimension Views	31
1.10 Defining the Staging Table	35
1.11 Including Dimension Links.....	38
1.12 Creating a Fact Table	48
1.13 Switching to Diagrammatic View	52
1.14 Producing Documentation.....	55
1.15 Data Store Objects (Optional).....	56


1.1 Purpose and Roadmap

Purpose

This tutorial is designed to introduce you to the basic objects used by QAD Data Warehouse Designer. At the end of the tutorial you will have built a simple dimensional analysis area of a data warehouse.

The tutorial will build the star-schema shown below. This star-schema comprises a central fact table, `fact_sales_detail`, joined to four dimension tables. Data will be loaded from tables in another Sql Server database, Oracle schema or DB2 schema.

In the process of creating this star-schema you will learn to create load, stage, fact and dimension tables. You will also see how data flows from the data source, through the different tables enroute to its fact or dimension table destination.



Tutorial Environment

This tutorial has been completed using Microsoft SQL Server. All of the features illustrated in this tutorial are available in SQL Server, Oracle and DB2 (unless otherwise stated). Any differences in usage of QAD Data Warehouse Designer between these databases are highlighted.

Tutorial Roadmap

This tutorial works through a number of steps. These steps and the relevant section within the chapter are summarized below to assist in guiding you through the tutorial.

Step in Tutorial	Section
Setup the WhereScape tool ▪ Install QAD ▪ Set-up tablespace defaults ▪ Set-up default prefixes for tables ▪ Create connection (to data source)	▪ The First Step ▪ Logging In ▪ Tablespace Defaults ▪ Table Name Defaults ▪ Creating a Connection
Create and load the load tables for ▪ Customer ▪ Product ▪ Order_line ▪ Order_header	Loading Source Tables
Create (and update from load tables) the following dimension tables ▪ Dim_customer ▪ Dim_product	Building Dimensions
Create dimensions for ▪ Dim_order_date ▪ Dim_ship_date These are views on the dim_date table	Creating Dimension Views
Create the stage_sales_detail table ▪ Create stage table using columns from load_order_line and load_order_header ▪ Specify join condition ▪ Include links to the following dimensions (dim_customer, dim_product, dim_order_date, dim_ship_date)	Defining the Staging Table Including Dimension Links
Create the fact_sales_detail table	Creating a Fact Table
View the WhereScape generated documentation	Switching to Diagrammatic View Producing Documentation

1.2 The First Step

The first step

To get started you need to follow the steps in the QAD Setup Administrator to create the required environment. The basic steps in this process are:

Oracle and IBM DB2 data warehouse

- 1 Install the QAD product suite.
- 2 Create a database schema for the QAD metadata repository.
- 3 Install the QAD metadata repository.

SQL Server data warehouse

- 1 Install the QAD product suite onto a PC.
- 2 Use the Quick Start option in the Setup Administrator utility to load the metadata and repository.

Note: See the Installation and Administration Guide for these procedures. For Oracle this tutorial assumes source data resides in the wtutorial schema and that the metadata has been loaded under the dssadm schema which has select access to the tutorial tables. For SQL Server it assumes that the data warehouse is in the WslWarehouse database and that the source data resides in the WslTutorial database. For DB2 it assumes that the data warehouse is in the QAD Data Warehouse Designer schema.

You are now ready to **log on** (see "**1.2.1 Logging In**" on page 8) to the repository you have created.

1.2.1 Logging In

Having completed the first step, and using QAD Data Warehouse Designer, you can now log on to the repository you have created .

To log in:

- 1 Click QAD Data Warehouse Designer from the Start menu. The Access Control screen displays. See sample screen below:

The screenshot shows a 'Repository Login' dialog box for 'QAD BI Data Warehouse Designer'. The title bar includes 'Version 6.5 by QAD Inc', 'Copyright 2012', and 'Licensed to QAD BI DWD Manual and Testing'. The dialog is divided into two sections: 'DATABASE' and 'METADATA REPOSITORY'. Under 'DATABASE', there are fields for 'Data Source' (a dropdown menu showing 'WslWarehouse'), 'Database Login ID' (a text box with 'dbo'), and 'Password' (an empty text box). Under 'METADATA REPOSITORY', there is a 'DWD User Name' field (a text box with 'QAD QuickStart'). At the bottom right, there are three buttons: 'Help', 'Cancel', and 'Connect'.

- 2 For SQL Server, the Data Source, Database Login ID and Password are those required to logon to the database server. If a trusted connection is being used enter dbo as the username and leave the password blank. For Oracle the Database Login ID and Password are those of the user under which the metadata repository has been loaded. See ***Switching Between Databases*** (see "1.2.2 Switching Between Databases" on page 10) for details on logging into IBM DB2.

Note: The metadata repository and data warehouse tables must reside under the same schema for Oracle and DB2. For SQL Server, the metadata repository and data warehouse tables must reside under dbo.

- 3 The RED User Name is the name that will be associated with any procedures, tables, etc, and scheduled jobs that are created from within QAD Data Warehouse Designer. Normally this would be your full name.
- 4 Click Connect. The Builder screen displays.

Note: ODBC is the only supported connection method. This connection must have been established prior to logon. Refer to the Installation and Administrator Guide if no such connection exists.

You are now ready to proceed to the next step where you define the ***Repository Defaults*** (see "1.3 ***Repository Defaults***" on page 10).

1.2.2 Switching Between Databases

The following sample logon screen shows the details entered for IBM DB2 for an operating system authenticated connection:



The screenshot shows the 'Repository Login' dialog box for 'QAD BI Data Warehouse Designer'. The title bar includes 'Version 6.5 by QAD Inc', 'Copyright 2012', and 'Licensed to QAD BI DWD Manual and Testing'. The dialog has a blue gradient background. On the left, the text 'QAD BI Data Warehouse Designer' is displayed. On the right, there are input fields for 'Data Source' (a dropdown menu showing 'DBTWO'), 'Database Login ID', 'Password', 'DWD Schema' (showing 'dssdemo'), and 'DWD User Name' (showing 'QAD QuickStart'). At the bottom right are three buttons: 'Help', 'Cancel', and 'Connect'.

For DB2, the Data Source, Database Login ID and Password as well as the Metadata Schema are those required to logon to the database server


Note: A user name and password will be required if operating system authentication is not being used.

Note: Ignore the Metadata Schema field if connecting to an Oracle or SQL Server repository after successfully connecting to DB2.

1.3 Repository Defaults

Before you begin to create the data warehouse, you can choose the defaults for the repository. You can do this from the **Tools** menu, by either selecting **Options** or **User Preferences**. There is no need to change the defaults for the tutorials.

- 1 From the **Tools** menu, select **User Preferences**.
Select **Show Dimension Views as a separate object group** and set to **True**.




2 Choose OK.

You are now ready to proceed to the next step where you define the **Tablespace (FileGroup) Defaults** (see "1.4 Tablespace (FileGroup) Defaults" on page 11)

1.4 Tablespace (FileGroup) Defaults

Before you begin to create the data warehouse, you can choose the defaults for the tablespaces (filegroups for SQL Server). There is no need to change the defaults for this tutorial.

1 From the Tools menu, select Options and then Storage.



Make the appropriate tablespace/filegroup choice for each option.

Note: The default table space or filegroup for the user will be used if no settings are selected.

You are now ready to proceed to the next step where you define the **Table Name Defaults** (see "**1.5 Table Name Defaults**" on page 12).

1.5 Table Name Defaults

Before you begin to create the data warehouse, you can choose the defaults for the table names. There is no need to change the defaults for this tutorial, and the examples given reflect the default naming convention.

- 1 From the Tools menu, select User Preferences and then Local Naming Conventions. Alter the defaults as required.
- 2 From the Tools menu, select Options and then Global Naming Conventions. Alter the defaults as required.

3 If no changes are made, the default table names will be:

- **load_** load tables with data copied from a source system
- **stage_** tables for manipulating and transforming data prior to publishing
- **dim_** dimension tables
- **fact_** fact tables, detail, rollup and snapshot
- **agg_** aggregate or summary tables built from fact tables
- **olap_** Analysis Services Olap cubes built from stage or fact tables

You are now ready to proceed to the next step ***Creating a Connection*** (see "**1.6 Creating a Connection**" on page 13).

1.6 Creating a Connection

In order to populate the metadata repository, connections need to be made to the source data. There must also be a connection to the data warehouse itself. This section describes how to make two new connections.

Note: The following two connections should have been automatically created. They should however be validated to ensure they are correct for the environment.

The first connection is to the source system. For Oracle this is the user within your Oracle database, for SQL Server the database, that contains the tutorial tables and for DB2 this is another schema within your database.


The second connection will be to the data warehouse tables.



TIP: In order to utilize the drag and drop features there must always be a connection to the data warehouse itself.

How to make a connection

- 1** Click on and highlight the **Connection** object group in the left pane. This selects the object group to be worked on.
- 2** Select **File | New**, or use the right mouse button and choose **New Object**. A dialog box displays with the Object Type defaulted to **Connection**. Name your connection. In this instance type *Tutorial (OLTP)*, and select **Add**. A Properties box will display:



SQL Server:

If running a SQL Server data warehouse then proceed as follows.

In the Properties box, complete the details as below, and then select **Update**:

- The **ODBC source** is the ODBC connection which has been defined to connect to the database. In this case the ODBC connection to the database that holds the tutorial tables.
- The **Provider name** identifies the type of connection that SQL Server will make in the case of a linked server. In this case it is not required as we are using tutorial tables in a SQL Server database on the same server.
- The **Database ID (SID)** is the SQL Server database name of the database being connected to. In this case the SID of the tutorial database.
- The **Database Link ID** is a SQL Server linked server link to connect from the data warehouse database to the source system database.

Note: This link is only required if the source database is on a different server from the data warehouse database. For the purposes of this tutorial, the database link ID is not required as the tutorial data is usually loaded into a database on the same server as the metadata.

- The **Extract User ID** and **Password** are the username and password required to logon to the tutorial database. If a trusted connection is being used then set the Extract User ID to "dbo".
- The **Administrator User ID** and **Password** are the administrator logon to the source location (tutorial). These can be left blank for the tutorial.
- The **Default Load Type** enables you to set the default load type at connection level for ODBC and database connections. Set to Database link load.
- 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. Leave this field blank.
- **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.

Tutorial (OLTP) Definition

Connection Properties

Notes

☐ Data Warehouse

Connection Name: Tutorial (OLTP)

Connection Type: Database

ODBC Source: WsITutorial

Provider Name:

Database ID (SID): WsITutorial

Database Link ID:

Extract User ID: sa Password: xxx

Administrator User ID: Password:

Default Load Type: Database link load

SSIS Connection String: Reset

Data Type Mapping Set: (Default)

OK Cancel Help

Oracle:

If running an Oracle data warehouse then proceed as follows.

In the Properties box, complete the details as below, and then select **Update**:

- The **ODBC source** is the ODBC connection which has been defined to connect to the database. In this case the ODBC connection to the database that holds the tutorial tables.
- The **Database ID (SID)** is the Oracle SID of the database being connected to. In this case the SID of the tutorial database.
- The **Database Link ID** is an Oracle database link to connect from the data warehouse database to the source system database.

Note: This link is only required if the source database is different to the data warehouse database. For the purposes of this tutorial, the database link ID is not required as the tutorial data is usually loaded into the same database as the metadata.

- The **Extract User ID** and **Password** are the username and password for the schema where the source tables reside. For the tutorial this is the user where the tutorial files have been loaded.
- The **Administrator User ID** and **Password** are the administrator logon to the source location (tutorial). These can be left blank for the tutorial.

- The **Default Load Type** enables you to set the default load type at connection level for ODBC and database connections. Set to Database link load.
- **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.

Tutorial (OLTP) Definition

☒ Data Warehouse

Connection Name: Tutorial (OLTP)

Connection Type: Database

ODBC Source: TUTORIAL

Database ID (SID): ORCL

Database Link ID:

Extract User ID: wtutorial Password: xxx

Administrator User ID: Password:

Default Load Type: Database link load

Data Type Mapping Set: (Default)

Notes

OK Cancel Help

IBM DB2:

If running an IBM DB2 data warehouse then proceed as follows.

In the Properties box, complete the details as below, and then select **Update**:

- The **ODBC source** is the ODBC connection which has been defined to connect to the database. In this case the ODBC connection to the database that holds the tutorial tables.
- The **Work directory** is not used.
- The **Database ID (SID)** is not used.
- The **Database Link ID** is not used.
- The **Extract User ID** and password are the username and password required to logon to the tutorial database. If an operating system authenticated connection is being used then leave the Extract User ID and Password blank.

- The **Administrator User ID** and password are the administrator logon to the source location (tutorial). These can be left blank for the tutorial.
- The **Default Load Type** enables you to set the default load type at connection level for ODBC and database connections. Set to Database link load.
- **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.

Tutorial (OLTP) Definition

☐ Data Warehouse

Connection Name: Tutorial (OLTP)

Connection Type: Database

ODBC Source: Wstutorial

Work Directory:

Database ID (SID):

Database Link ID:

Extract User ID: wtutorial Password: ***

Administrator User ID: Password:

Default Load Type: Database link load

Data Type Mapping Set: (Default)

OK Cancel Help

- 3 To confirm that you have connected to the system correctly, select **Source Tables** from the **Browse** menu, or click on one of the browse icons from the main tool bar or right hand pane tool bar.
 - Select the connection you want to view, in this instance **Tutorial (OLTP)**, and click **OK**. For SQL Server the schema must be set to **dbo**. For Oracle the schema should be the tutorial schema.
 - A third pane on the right hand side, displays showing the tables contained under the tutorial source system.
- 4 Repeat steps 1 through 3 to create the connection for the Data Warehouse.
 - The Connection name will be **Data Warehouse**

- Enter an extract user id (we have used **dssadm**) and a password (we have used **WSL**) for the metadata repository. For a SQL Server trusted connection set the extract user id to **dbo**.

You have now created two database connections, one to the source system (**Tutorial**), and one to the **Data Warehouse**.

You are now ready to proceed to the next step - **Loading Source Tables** (see "**1.7 Loading Source Tables**" on page 18).


1.7 Loading Source Tables

In this step you will load data from the tutorial source system into load tables in the data warehouse. Dragging and dropping from the source system (using the previously defined connection) will create the metadata. You will then be prompted to create and load the tables which will create the physical tables in the data warehouse, and then load the data.




TIP: Ensure that your source system is displayed in the right hand pane of the screen, by selecting **Source Tables** from the **Browse** menu, then **Tutorial (OLTP)** from the **Connection List**. For **SQL Server** the schema must be **dbo**. For **Oracle** the schema should be the tutorial schema. Click **OK**.


- 1 Double click on the **Load Table object group** on the **Object Tree** in the left pane. The first column heading in the middle pane should read *Load Table Name*.
- 2 Expand the source table **Object Tree** in the right pane.



- 3 Click on **customer** and drag this table into the middle pane - placing it anywhere in the pane. A dialog box displays with the name of the object defaulted to **load_customer**.



- 4 Click **Add** and the following table definition will display:



The image shows a dialog box titled "load_customer Table Definition". On the left is a sidebar with tabs: "Properties" (selected), "Source Mapping", "Storage", and "Notes". The main area contains the following fields:


- 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 section:
 - Metadata Structure Changed:
 - Database Created:
 - Database Altered:

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

Note: For the purposes of this tutorial, all the necessary details have been automatically created, so select OK. See the Loading Data chapter for explanations of the load parameters.

Note: In IBM DB2, short names are limited to 12 characters.

- 5 A dialog box displays showing that the load table `load_customer` has been defined and asks if you want to create and load the table:



The image shows a dialog box titled "Create Database Table". It contains a message: "Load Table load_customer has been defined". At the bottom are three buttons: "Create", "Create and Load" (with a dropdown arrow), and "Close".

- 6 Select **Create and Load**. This will create the physical tables in the data warehouse and load the data.

- 7 Results will be posted in the results pane. Note that the Load Table object group in the left pane now has a dependent/child.



TIP: Remember to double click on the left pane Load Table object group between loading each of the source tables to ensure that you are reassigning the target, rather than adding to the columns in the middle pane.

- 8 Repeat this process (steps 2 - 7) for the source tables product, order_header, and order_line.
- 9 Double click the Load Table group in the left hand pane. Your QAD Data Warehouse Designer screen should look like this when you have finished:

The screenshot shows the QAD Data Warehouse Designer interface. The main window is titled "Builder Scheduler" and contains several panes:

- Development** (Left Pane): A tree view showing the project structure. The "Load Table" group is expanded, showing sub-items: load_customer, load_order_header, load_order_line, and load_product. The "Load Table" group is selected.
- Load Table Columns** (Middle Pane): A table with two columns: "Column Name" and "Display Name". The table contains the following data:

Column Name	Display Name
code	code
description	description
prod_line	prod line
prod_group	prod group
subgroup	subgroup
- Results** (Bottom Pane): A table with two columns: "Obj..." and "Message". The table contains the following data:

Obj...	Message
	product.code,product.description,pro
	From [WslTutorial].dbo.product
	product
	;
	9 rows loaded
- Tutorial (OLTP)** (Right Pane): A tree view showing the database structure. The "dbo" group is expanded, showing sub-items: budget, code_table, customer, forecast, kpi, order_header, order_line, prod_group, prod_line, prod_subgroup, product, and states. The "product" table is selected.

The status bar at the bottom indicates: "Read Middle Pane: Load Table Columns | Development (WslWarehouse) | UserId: dbo | Browse: Tutorial (OLTP) | CAP | NUM | SCRL | INS ...".


You are now ready to proceed to the next step - **Building Dimensions** (see "1.8 Building Dimensions" on page 22).

1.8 Building Dimensions

The necessary source tables have been loaded into the data warehouse. Now the dimensions of the data warehouse can be built. When building dimensions you will be prompted for how you would like the dimension managed. QAD Data Warehouse Designer generates code for normal, slowly changing, previous value and date ranged dimensions. You will also be prompted for the business (or natural) key of the dimension. This is needed so QAD Data Warehouse Designer knows when to add new dimensional records.

- 1 Change the right pane view to show the Data Warehouse tables by selecting DataWarehouse from the Browse menu OR click the tab along the bottom of the source window.

Note: For SQL Server the Data Warehouse schema must be **dbo**.




Note: From this point on, all work is performed within the data warehouse.

- 2 Double click on the **Dimension** object group in the object tree in the left pane. The first column of the middle pane now reads *Dimension Name*.

Note: You will see that some dimensions have already been created for you.

- 3 Click and drag the **load_customer** table from the data warehouse schema in the right pane into the middle pane. A dialog box displays defaulting the name of the object to **dim_customer**.




Add a New Metadata Object

Define the Type and Name of the New Object.
Specific information for each object type is defined in subsequent screens.

Object Type:

Object Name:

- 4 Click Add. A dialog box displays asking how you want the dimension managed. Select Normal.



Dimension Type

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

1. Normal. The dimension is updated based on a business key, with new records being added if required. All columns except the business key can change.

2. Slowly changing. Changes in the values of selected columns result in new dimensional records being created. In all other respects the same as Type 1.

3. Previous data retained. The previous values of selected columns are stored in additional columns. In all other respects the same as Type 1.

4. Date Ranged. The source system provides a date ranged business key. Similar to Type 2 except that we deal with the record as a whole and the dates are provided.

- 5 A table definition displays with all the necessary defaults completed.
- Make one change - Select **(Build Procedure...)** from the Update Procedure drop-down list box. This will generate procedures to get surrogate (artificial) keys based on the business key and to update the dimension.

dim_customer Table Definition

Properties

Storage

Language Mapping

Purpose

Concept

Grain

Examples

Usage

Notes

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)

Get Key Function: Mnemonic (EUL):

Timestamps


Metadata Structure Changed: Database Created: Database Altered:

- 6 Select OK. A dialog box displays confirming the dimension table **dim_customer** has been defined and asking if you want to create and load the table.


Create Database Table

Dimension dim_customer has been defined.

- 7 Select **Create and Load**.
- 8 A Procedure Build Type dialog will appear. Select **Cursor** and then click OK.



- 9 Define the Business Key (see below). The business (natural) key is the unique identifier for the dimensional record. Select **Code** and **>** and click OK.



- 10 The procedure results display and can be reviewed.

Note: The Dimension Table object group in the left pane now has added `dim_customer` as a dependent/child.

- 11 Repeat this same process (steps 2 through 9) for the load table `load_product`. The business key will be 'code'.



TIP: Remember to double click on the left pane Dimension Table object group between loading each of the above dimension tables.

Double click Dimension in the left pane. The QAD Data Warehouse Designer screen should look like this:

The screenshot shows the QAD Data Warehouse Designer interface. The left pane (Development) shows the 'Dimension' object group selected. The central pane (Dimension) displays a table with the following data:

Dimension Name	Short Name
dim_customer	dim_customer
dim_date	dim_date
dim_product	dim_product
dss_fact_table	dss_fact_table
dss_source_system	dss_source_syste


The right pane (Data Warehouse) shows the 'dbo' schema with the following tables:

- dim_date
- dss_current_day_details
- dss_fact_table
- dss_parameter
- dss_source_system
- load_customer
- load_order_header
- load_order_line
- load_product

The bottom status bar shows: Ready | Middle Pane: Dimension | Development (WslWarehouse) | UserId: dbo | Browse: DataWarehouse | CAP | NUM | SCRL | INS

Note: Analysis Services does not like 'name' as a column name. For dim_customer it will therefore be necessary to change the column name from 'name' to 'cname'.

- 1 Click on the dim_customer object in the left hand pane to display the dim_customer columns in the middle pane.
- 2 When positioned on the column 'name' in the middle pane, right click and select Properties from the drop down menu.



3 Change the column name from 'name' to 'cname' as shown below. Click OK.

dim_customer.name Column Definition

Column Properties
Transformation
Language Mapping

Table Name: dim_customer

Column Name: cname

Business Display Name: cname

Data Type: varchar(45)

Format:

Column Order: 30

Default Value:

Nulls: ☒

End User Layer Display: ☒

Numeric: ☐

Key Type (0,A,B,C,...):

Additive: ☐

Artificial Key: ☐

Attribute: ☒

Primary Business Key: ☐

Source Table: load_customer

Source Column: name


Business Definition:
Source strategy:
The full name of the customer. Forms a hierarchy with city and state.

The value shown below is loaded in this column for the 'Invalid Join' record. When empty NULL is used.

Zero Key Value: Unknown

OK Cancel Help

- 4 Right click on the dim_customer object in the left hand pane and choose Validate against Database.



- 5 The results will show that the metadata has been changed to 'cname' while the column name in the database is still 'name'.


Table validation list	
Table Name	Results
dim_customer	dss-> name - varchar(45), meta-> cname - varchar(45),

- 6 Right click on dim_customer in the bottom pane and select Alter table from the drop down list.

Table validation list	
Table Name	Results
dim_customer	dss-> name - varchar(45), meta-> cname - varchar(45),

Alter Table
Export to CSV File

- 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. Select **OK**.
- 9 Right click on the dim_customer object in the left hand pane and select **Properties** from the drop down menu. Choose the **Rebuild** button.
- 10 A Procedure Build Type dialog will appear. Select **Cursor** and then **OK**.
- 11 Leave the Business key as **Code** and click **OK**.
- 12 Right click on the dim_customer object in the left hand pane and select **Execute Update Procedure**.
- 13 Click in the right hand pane and press F5 to refresh the Data Warehouse table view in the right hand pane. The QAD Data Warehouse Designer screen should look like this:




You are now ready to proceed to the next step - **Creating Dimension Views** (see "1.9 Creating Dimension Views" on page 31)

1.9 Creating Dimension Views

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 this step you will create dimension views from an existing dimension. In many cases 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. This process is essentially the same as creating a dimension, but you are creating a view of an existing dimension, in this instance, `dim_date`.

- 1 After double clicking on **Dimension View** in the left pane, click and drag `dim_date` from the right pane into the center pane.
 - The dialog box that displays defaults the object type to a dimension view, and names the dimension view `dim_date_view`.
 - Because we want to create two dimension views from the same source, `dim_date`, we need to change this dimension view name to one that is more meaningful, specifically `dim_order_date`.



- Make this change, and click ADD.
- 2 The `dim_order_date` view property defaults have all been completed as necessary so click OK.

dim_order_date Table Definition

Properties

Storage

Language Mapping

Purpose

Concept

Grain

Examples

Usage

Notes

View Name: View Type:

Unique Short Name: (maximum 22 characters)

Business Display Name (EUL):

Description:

View Properties Clauses:

Where Clause:


Get Key Function: Mnemonic (EUL):

Timestamps

Metadata Structure Changed:	Database Created:	Database Altered:
2011-07-28 16:10:01.563		

OK Cancel Help

- 3 A dialog box displays to provide a means of re-mapping some of the column names in the view if required. Ensure calendar_date is being renamed to order_date. Click the OK button.



The dialog box is titled "View Column Definition". It contains a text box with the instruction: "The column names for the view being created can be modified by filling in the following form. If the Default button is pressed only the artificial (surrogate) key will be changed." Below this are four input fields with arrows pointing from left to right:


- Remove Column Prefix: [] → Add Column Prefix: order_
- Remove Business Display Prefix: [] → Add Business Display Prefix: order

Below these is a section titled "Change Column Names for Specific Columns" with a table of input fields:

Old Column Name:	New Column Name:
calendar_date	order_date
[]	[]
[]	[]
[]	[]
[]	[]

At the bottom right are two buttons: "OK" and "Default".


- 4 A dialog box displays indicating that the dimension view dim_order_date has been defined and asks if you want to create the view now:




The dialog box is titled "Create Database Table". It contains a text box with the message: "Dimension View dim_order_date has been defined". Below the text box are three buttons: "Create View + Function", "Create View", and "Finish".

Select Create view + Function. A dialog box will confirm that view dim_order_date has been created.

- 5 Select a business (natural) key for the dimension view `dim_order_date`. Select `order_date` and > and click OK.



- 6 Repeat steps (1) to (5) to create the dimension view `dim_ship_date`.
- 7 Click in the right hand pane and press F5 to refresh the Data Warehouse table view in the right hand pane. The QAD Data Warehouse Designer screen should look like this:



You are now ready to proceed to the next step - **Defining the Staging Table** (see "1.10 Defining the Staging Table" on page 35).


1.10 Defining the Staging Table

In this step you will create a stage table from two load tables. A stage table is used to build the format of the fact table, and generally contains changed or new data that will be added to the fact table. As stage tables contain dimensional keys, they should be defined after the dimensions.

Note: The source of data for the stage table will be the load tables `load_order_line` and `load_order_header`.

- 1 Double click on the **Stage Table** object group in the object tree in the left pane to create a stage table target. The first column heading in the middle pane reads *Stage Table Name*.

- Click and drag the `load_order_line` table from the right hand pane data warehouse schema. Drop it in the middle pane. A dialog box displays defaulting the name of the object to `stage_order_line`. To make it a more meaningful name, change the name of the object to `stage_sales_detail` 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:

- A table definition displays with all the necessary defaults completed.



stage_sales_detail Table Definition

Properties

Storage

Notes

Table Name: Table Type:

Unique Short Name: (maximum 22 characters)

Description:

Update Procedure:

Custom Procedure:

Table Properties Clauses: (e.g. Partition)

Timestamps


Metadata Structure Changed: Database Created: Database Altered:

- 4 Select OK.

Note: The Stage Table object group in the left pane now has a dependent/child.

- 5 To add the remaining information from the second load table, click on stage_sales_detail in the left pane. Next drop load_order_header from the right pane into the middle pane. A message is displayed with options to create a "New table" or to "Add columns". Choose **Add columns**.

QAD Data Warehouse Designer detects duplicated columns. As both load_order_header and load_order_line have the *order number* field, the following is displayed:




Click **Skip** to exclude the second instance of order_number.

Note: If the second instance of order_number is required, then click OK.

This combines data from two load tables (load_order_header and load_order_line) into one stage table.

The QAD Data Warehouse Designer screen should now look like the sample below. In the middle pane under *Src Table* you can see the source of each of the columns.




You are now ready to proceed to the next step - **Including Dimension Links** (see "1.11 Including Dimension Links" on page 38).

1.11 Including Dimension Links


The dimension links that allow us to create the fact-like star schema now need to be included:

- 1 In the left pane, click on the **stage_sales_detail** table in the Stage Table object group. The middle pane should display the contents of this stage table.
- 2 Drag each of the following dimensions from the right pane into the stage table in the middle pane:
 - dim_customer
 - dim_order_date
 - dim_product
 - dim_ship_date

This adds the dimension keys from each dimension to the stage table. The QAD Data Warehouse Designer screen should now look like this:



- 3 The stage table metadata has been defined, but the stage table has not been created. To create the stage table in the data warehouse, click on `stage_sales_detail` in the left hand pane, right click and choose Create (ReCreate).



Note: The table must exist in the data warehouse before we can proceed to the next step. If the table has not been physically created then the procedure in step 5 will fail to compile.

- 4 Double click on the stage table to select Properties.
- 5 Under Update Procedure, choose (Build Procedure...)

stage_sales_detail Table Definition

Properties
Storage
Notes

Table Name: stage_sales_detail Table Type: Stage

Unique Short Name: (maximum 22 characters) stage_sales_detail

Description:

Update Procedure: (Build Procedure...) Rebuild Regenerate

Custom Procedure: (None)


Table Properties Clauses: (e.g. Partition)

Timestamps

Metadata Structure Changed: 2011-07-29 11:16:41.230 Database Created: Database Altered:


OK Cancel Help

- 6 This will create an update stage procedure. Click OK.
- 7 Select the **Cursor** based procedure generation from the stage procedure type dialog box:




Note: When building an Oracle data warehouse, this dialog has an additional option for bulk bind procedures. See Staging generating the update procedure for more information.


- 8 Click OK on the parameters screen.
- 9 A dialog box will display indicating that multiple source tables have been detected.



- 10 Select OK and then highlight the source tables `load_order_line` and `load_order_header` which are to be joined by `order_number`.
 - With the two tables highlighted click **Outer Join**. See the chapter on Staging data for an explanation of the join types and options.
 - Select `order number` from the `load_order_line` empty drop-down list box at the bottom of the screen. Then select `order number` from the `load_order_header` drop-down list box.



- This will create a join statement in the right hand window:



The dialog box is titled "Provide Cursor Mapping". It has a close button (X) in the top right corner. Below the title bar, there is a text field for "Cursor Name (if applicable):" with the value "Load". To the right of this field is a help icon. Below this, a text label reads "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 main area is divided into two panes. The left pane, titled "Source Tables:", contains a list box with two items: "load_order_line" and "load_order_header", with "load_order_line" selected. The right pane, titled "From and Where Clause:", contains a text area with the following SQL code:

```
FROM load_order_line
LEFT OUTER JOIN load_order_header
ON load_order_line.order_number = load_order_header.order_number
```

 Below the panes are three buttons: "Outer Join", "Simple Join", and a checked checkbox labeled "ANSI join code generated". Below these buttons is a text label: "Select the columns that join the two tables. Select the column from the Master Table first." There are two dropdown menus. The first dropdown is labeled "load_order_line" and the second is labeled "load_order_header". At the bottom left is a checkbox labeled "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_line
- load_order_header

From and Where Clause:

```
FROM load_order_line
LEFT OUTER JOIN load_order_header
ON load_order_line.order_number = load_order_header.order_number
```

Outer Join Simple Join ☒ ANSI join code generated

Select the columns that join the two tables. Select the column from the Master Table first.

load_order_line load_order_header

☐ Word Wrap Displayed Code

OK Cancel

- Click OK.

11 You need to match the dimension business keys with the business keys in the stage table. This associates the correct dimensional record to each stage table record. A dialog box displays for each dimensional join:

- For the `dim_customer`, select `customer_code`. Click > and OK.

Dimension Business Key Definition

Stage Table Column List:

- customer_code
- order_date
- order_line_no
- order_number
- product_code
- quantity
- sales_value
- ship_date
- tax
- unit_sale_price

Business Keys for dim_customer

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:

- customer_code

Dimension Business Keys:

- code

dim_customer

☐ 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)

Add Text Enter a text string and press Add Text to add a static Business Key Value


OK Cancel

- For **dim_product**, select **product_code**. Click > and OK.
 - The business key for **dim_ship_date** has the same column name in the stage table and the dimension view, allowing QAD Data Warehouse Designer to automatically move **ship_date** to the left hand side. Click OK to progress to **dim_order_date**, where **order_date** has also been automatically chosen. Click OK again.
- 12** You must now select the business keys to uniquely identify each record in the staging table itself. This essentially defines the business key we will be using in the fact table, and as such defines the grain of the fact table. For this example the grain is order line. Select **order_number** and **order_line_no**. Click > and OK.




- 13** QAD Data Warehouse Designer now builds and compiles the update procedure. The results pane shows any indexes that were created.

- 14 The final step is the population of the stage table. Click on `stage_sales_detail` in the left hand pane, right click and choose **Execute Update Procedure**.



- 15 Output from the stage table has being updated can now be seen in the **Results** window:



You are now ready to proceed to the next step - **Creating a Fact Table** (see "1.12 Creating a Fact Table" on page 48).


1.12 Creating a Fact Table

In this step you will create a fact table.


Create a drop target by double clicking on the **Fact Table** object group in the left pane.

To create the fact table:

- 1 Browse the data warehouse connection again (or refresh the data warehouse connection):



- 2 Drag the stage table **stage_sales_detail** over from the right pane into the middle pane. The following dialog is displayed:




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:

- 3 Click Add. The fact_sales_detail table properties screen will appear. Select (Build Procedure...) in the update procedure drop down and click OK:



fact_sales_detail Table Definition

Properties

Storage

Language Mapping

Purpose

Concept

Grain

Examples

Usage

Notes

Table Name: Table Type:

Unique Short Name: (maximum 22 characters)

Business Display Name (EUL):

Description:

Update Procedure: Cursor Based Update

Custom Procedure:


Table Properties Clauses: (e.g. Partition)

Get Key Function: Mnemonic (EUL):


Timestamps

Metadata Structure Changed: Database Created: Database Altered:


- 4 Select Create and Load to create and load the table now:



- 5 Select the Business Key for the fact table. Choose `order_number` and `order_line_number`. Click > and then OK.




- 6 Output from the fact table being created and updated can now be seen in the **Results** window. Your QAD Data Warehouse Designer screen should look something like this:



You are now ready to proceed to the next step - **Switching to Diagrammatic View** (see "1.13 Switching to Diagrammatic View" on page 52)


1.13 Switching to Diagrammatic View

QAD Data Warehouse Designer provides the ability to diagrammatically view the data warehouse you have created. To switch to the diagrammatic view, click on the  button. Click the button again to bring up the diagram selector box. Select a fact table to be drawn as a star schema from the drop-down list box – in this instance **fact_sales_detail**:




Click OK.

The diagram looks like this:



TIP: To view the source tracking of the **fact_sales_detail** table, click once more on the  button, and choose the **fact_sales_detail** table from the second drop-down list box.

The diagram looks like this:



To close the diagrammatic view, click on the X at the right hand end of the toolbar.

You are now ready to proceed to the next step - ***Producing Documentation*** (see "***1.14 Producing Documentation***" on page 55).

1.14 Producing Documentation

QAD Data Warehouse Designer also provides the ability to produce user and technical documentation. This is obviously of more value if the descriptive data has been entered against the columns and tables in the data warehouse, which we have not done during this tutorial.

To view the documentation for the components of the data warehouse, select Doc from the menu, then Create Documentation. Select a file path (directory) under which to save the HTML files that will be produced. The next screen allows for the inclusion of a banner and user defined links:

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.

Leave these options unchecked and click OK to proceed.

The documentation runs:



TIP: To view the documentation select Doc, Display Documentation:




You are now ready to proceed to the next step - **Data Store Objects** (see "**1.15 Data Store Objects (Optional)**" on page 56)

1.15 Data Store Objects (Optional)

Data Store objects are used to provide a persistent storage of load tables. These objects are not licensed for every installation and hence this section is optional.

- 1 Browse the Data Warehouse in the right hand pane.
- 2 Double click the data store object in the left hand pane.
- 3 Drag `load_product` from the right hand pane into the middle pane.
- 4 Accept the default name of `ds_product` 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 Select (Build Procedure...) from the Update Procedure drop-down list. Click OK.



ds_product Table Definition

Properties
Storage
Notes

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)


Get Key Function: Mnemonic (EUL):

Timestamps

Metadata Structure Changed:	Database Created:	Database Altered:
2011-07-29 16:14:08.010		

6 Click Create and Load.**7** Select the options as shown below and click OK.

- 8 Select **code** as the business key, select > and click OK.



The dialog box is titled "Define Data Store Business Key Columns". It contains two main sections: "Column List" on the left and "Business Key List" on the right. The "Column List" contains the following items: description, prod_group, prod_line, and subgroup. The "Business Key List" contains the item: code. Between the two lists are two buttons: ">" and "<". Above the "Business Key List" is the instruction: "Select the business keys that uniquely identify each record in the Data Store table. Move them over to the business key list". At the bottom right are "OK" and "Cancel" buttons.

Define Data Store Business Key Columns

Column List:

- description
- prod_group
- prod_line
- subgroup

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:

- code

>

<


OK Cancel

9 Click OK for Source Table mapping.

The dialog box is titled "Source Table Mapping". It contains the following elements:

- Instructions:** "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 list box containing "[load_product]".
- Join Clause:** A text area containing "FROM [load_product] load_product".
- Where Clause:** An empty text area.
- Join Type:** Two buttons, "Outer Join" and "Simple Join".
- Column Selection:** Four dropdown menus labeled "Table A", "Column A", "Table B", and "Column B".
- Group By Clause:** An empty text area.
- Word Wrap:** A checkbox labeled "Word Wrap Displayed Code".
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

10 Data should now be loaded into the ds_product table.



Repeat the exercise for order_line, order_header and customer.

CHAPTER 2

Tutorial 2 - Rollup Fact Tables, ASCII File Loads, Aggregates

Before you start on this chapter you should have:

- Completed *Tutorial 1 - Basic Star Schema Fact Table* (on page 5)
- Successfully completed *Creating a Fact Table* (see "*1.12 Creating a Fact Table*" on page 48)

This chapter deals with fine tuning the data warehouse by creating roll-up fact tables and aggregates. It also includes loading an ascii file into a new load table.

In This Chapter

2.1 Purpose and Roadmap.....	64
2.2 Making a Connection to Windows.....	66
2.3 Loading Tables from Flat Files	70
2.4 Creating Stage Tables	75
2.5 Creating Fact Tables	77
2.6 Rollup/Combined Fact Table	79
2.7 Aggregate Tables	82
2.8 Creating a Customer Aggregate	84

2.1 Purpose and Roadmap

Purpose


This tutorial will walk you through the process to:

- Load data from flat files (in Tutorial 1 source data was obtained via database links)
- Create a rollup fact table that allows users to see budgeted, forecast, and actual sales amounts and quantities broken down by customer, product and month
- Create separate aggregate tables that summarize data in the rollup table by (i) product and (ii) customer


In short, this tutorial loads budget and forecast data from flat files into their own load, stage and fact tables. This data is then combined with data from the fact_sales_detail table (created in Tutorial 1) and summarized to create a new rollup fact table, fact_sales_analysis. Further summarization is done on the rollup table to create two separate aggregate tables.

The following are diagrams showing (i) the rollup table, fact_sales_analysis and (ii) the customer aggregate table, agg_sa_customer that will be created as part of this tutorial.

Rollup:



Aggregate:



Tutorial Environment

This tutorial has been completed using IBM DB2. All of the features illustrated in this tutorial are available in SQL Server, Oracle and DB2 (unless otherwise stated). Any differences in usage of QAD Data Warehouse Designer between these databases are highlighted.

Tutorial Roadmap

This tutorial works through a number of steps. These steps and the relevant section within the chapter are summarized below to assist in guiding you through the tutorial.

Step in Tutorial	Section
Create a new connection to allow data to be loaded in from flat files on C: drive.	Making a Connection to Windows

Step in Tutorial	Section
Create (and load) the data for ▪ load_budget ▪ load_forecast Note: Data is loaded from flat files on C: drive	Loading Tables from Flat Files
Create the following stage tables ▪ stage_budget ▪ stage_forecast Stage table creation entails using corresponding load tables and including links to the following dimensions: (dim_customer, dim_product, dim_date)	Creating Stage Tables
Create the following fact tables ▪ fact_budget ▪ fact_forecast	Creating Fact Tables
Create the rollup fact table, fact_sales_analysis This rollup combines forecast, budget and sales data from fact_budget, fact_forecast and fact_sales_detail tables respectively. The data is rolled-up (grouped) by product, customer and month. Note that dimension keys are used for the rollup.	Rollup/Snapshot Fact Tables
Create product aggregate (agg_sa_product) table. This aggregate summarizes fact_sales_analysis data by product	Aggregate Tables
Create customer aggregate (agg_sa_customer) table. This aggregate summarizes fact_sales_analysis data by customer.	Creating a Customer Aggregate

This tutorial starts with the section ***Making a Connection to Windows*** (see "2.2 ***Making a Connection to Windows***" on page 66).

2.2 Making a Connection to Windows

This follows a similar process to the earlier ***connections*** (see "1.6 ***Creating a Connection***" on page 13) made, but differs in that the connection is within the PC.

Note: The following connection should have been automatically created. It should however be validated to ensure that it is correct for the environment.

- 1 Log on to QAD Data Warehouse Designer.
- 2 In the left pane, double click on the Connection object group.

- 3 Select **File | New**, or right click the **Connection** object and select **New Object**.
- 4 Enter **Windows** in the 'Name of Object' text box and select **Add**.
- 5 Enter the Connection properties as below:

Connection properties

Field	Description
Connection Name	Windows
Connection Type	Windows
Host Name	Not required
Work Directory	Required. Must be an existing valid directory on the PC, e.g, c:\temp
Database id (SID)	For Oracle, the appropriate SID for your metadata installation, e.g. ORCL. For SQL Server and DB2 leave this field blank.
Database home directory	This is only required for Oracle and if the Database SID is in a non-default directory
Windows User ID	Leave blank
Windows Password	Leave blank
Dss User ID	For Oracle, this is dssadm (the data warehouse username). This is the database logon for SQL Server and DB2. It should be dbo if a trusted connection is being used for SQL Server. If an OS authenticated user is being used for DB2 this should be left blank.
Dss password	For Oracle, this is wsl (the data warehouse password). This is the database password for SQL Server and DB2. It should be left blank if a trusted connection is being used for SQL Server or if an OS authenticated user is being used for DB2.

Sample SQL Server properties:

The screenshot shows a 'Windows Definition' dialog box with a 'Connection Properties' tab. The 'Data Warehouse' checkbox is unchecked. The 'Connection Name' is 'Windows' and the 'Connection Type' is 'Windows'. The 'Host Name' is empty. The 'Work Directory' is 'c:\temp'. The 'Database ID (SID)' and 'Database Server' are empty. The 'Windows User ID' and 'Password' are empty. The 'DSS User ID' is 'dbo' and the 'Password' is empty. The 'Data Type Mapping Set' is '(Default)'. The dialog has 'OK', 'Cancel', and 'Help' buttons at the bottom right.

Property	Value
Connection Name	Windows
Connection Type	Windows
Host Name	
Work Directory	c:\temp
Database ID (SID)	
Database Server	
Windows User ID	
Windows Password	
DSS User ID	dbo
DSS Password	
Data Type Mapping Set	(Default)

Sample Oracle properties:

The screenshot shows a 'Windows Definition' dialog box with a 'Connection Properties' tab. The 'Data Warehouse' checkbox is unchecked. The 'Connection Name' is 'Windows', and the 'Connection Type' is 'Windows'. The 'Host Name' is empty. The 'Work Directory' is 'c:\temp', the 'Database ID (SID)' is 'ORCL', and the 'Database Server' is empty. The 'Windows User ID' is empty, and the 'Password' is empty. The 'DSS User ID' is 'dssdemo', and the 'Password' is 'xxxx'. The 'Data Type Mapping Set' is '(Default)'. The dialog has 'OK', 'Cancel', and 'Help' buttons at the bottom right.

Property	Value
Connection Name	Windows
Connection Type	Windows
Host Name	
Work Directory	c:\temp
Database ID (SID)	ORCL
Database Server	
Windows User ID	
Password	
DSS User ID	dssdemo
Password	xxxx
Data Type Mapping Set	(Default)

Sample DB2 properties:


The screenshot shows the 'Windows Definition' dialog box with the 'Connection Properties' tab selected. The 'Data Warehouse' checkbox is checked. The 'Connection Name' is 'Windows' and the 'Connection Type' is 'Windows'. The 'Host Name' is empty. The 'Work Directory' is 'c:\temp'. The 'Database ID (SID)' and 'Database Server' are empty. The 'Windows User ID' and 'DSS User ID' are empty, each with a corresponding 'Password' field. The 'Data Type Mapping Set' is '(Default)'. The 'OK', 'Cancel', and 'Help' buttons are at the bottom right.

You are now ready to proceed to the next step - **Loading Tables from Flat files** (see "2.3 Loading Tables from Flat Files" on page 70)


2.3 Loading Tables from Flat Files

In this step you will parse and load a file from Windows into a load table in the data warehouse.

- 1 Double click on the **Load Table** object group in the left pane. This will list all load tables in the middle pane and make the middle pane a drop target for new load tables.
- 2 Browse to the Windows connection in the right pane by selecting **Browse | Source Tables** from the menu strip at the top of the screen. Select **Windows** as the Connection. Leave the Schema field blank. Click OK.



- 3 Navigate to `c:\Program Files\QAD\Tutorial` folder, click on `budget.txt` and drag it into the middle pane. Rename the object `load_budget` (i.e. remove the `_txt` suffix):



- 4 Click ADD.


Specifying the load type

You must now specify the type of load you require from the four options given:

- File load (single data column)
- File load (columns parsed)
- Script load (single data column)
- Script load (columns parsed)

For more information on load types see the section on flat file loads in the loading data chapter.

- 1 For this tutorial select **Script Load Columns Parsed**:



Note: For DB2, file loads are not available in this dialog.

- 2 From the definition dialog that displays select OK.
- 3 From the data load wizard enter a comma (,) in the Column Delimiter field

- 4 As the first row is a header, place a check in the box and click OK:

The image shows a 'Data load Wizard' dialog box. The title bar says 'Data load Wizard' with a close button (X). The main area is titled 'First Rows from the File' and contains a list of data rows. The first row is a header: 'product_code,customer_code,budget_quantity,budget_sales_value,budget_date'. The subsequent rows are data records with numerical values and dates. Below the list, there are three settings: 'Column Delimiter' set to a comma, 'First Row is a Header' checked, and 'Record Delimiter' set to a newline. To the right of these settings are explanatory text blocks and a 'Decimal Code' button. At the bottom right are 'OK' and 'Cancel' buttons.

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
1007,223,80,622.53,13-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

- 5 QAD Data Warehouse Designer uses the header row as suggested column names. For each column confirm the name and data type:

Data load Wizard - Column Definition

Column Data:

product_code
1002
1008
1003
1007
1004
1006
1009
1002
1006
1007
1004
1003
1008

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		

Display decimal character values

Column Name:

Business Display Name:

Data Type: ☒ Nulls

Conversion:


Business Definition:

Back Add Cancel

- 6 You may have to change it to a more appropriate value. Click **Add**.
- 7 Click **OK** for the DB2 Load cannot skip header rows dialog.
- 8 Click **OK** for the New Script Create dialog.
- 9 Select **Yes** on the prompt to create and load the table now.

Note: Loading files with a header row into DB2 will result in an error message.

- 10 Repeat steps (1) to (9) for the `forecast.txt` file to create and load the table `load_forecast`.
- 11 The results display:



You are now ready to proceed to the next step - **Creating Stage Tables** (see "2.4 Creating Stage Tables" on page 75).

2.4 Creating Stage Tables

Two separate stage tables need to be created for `load_budget` and `load_forecast`. This is the same as the procedure from the first tutorial for **Defining the Staging Table** (see "1.10 Defining the Staging Table" on page 35).


- 1 Double click the **stage table** object group in the left hand pane. This will list the existing stage table in the middle pane.
- 2 Browse to the data warehouse **Browse/Source Tables** or click on the orange glasses in the toolbar.
- 3 Drag the table `load_budget` from the right hand pane to the middle pane and drop.
- 4 Click **Add** to add the stage table called `stage_budget`. Click **OK** on the properties screen.
- 5 Repeat steps (1) through (4) for the table `load_forecast` to create `stage_forecast`.

Once the stage tables have been defined, bring the following keys from the right pane into both new tables. Click the stage table name in the left hand pane to list the stage table columns in the middle pane. This also makes the middle pane a drop target for new columns:

Note: Repeat the following procedure for `stage_budget` and `stage_forecast`.

- `dim_customer_key`
 - `dim_product_key`
 - `dim_date_key`
- 1** Now in the left pane, right click on `stage_budget`, then select **Create (ReCreate)**.
 - 2** Again, in the left pane, right click on `stage_budget`, then select **Properties**.
 - 3** In the Update Procedure field select **(Build Procedure...)**. Click OK.
 - 4** Select **Cursor** as the update procedure type.
 - 5** SQL Server data warehouse users will now see an additional join screen. This screen is presented even though no joins are required. This screen allows the selection of either a 'Where' based join or an ANSI standard join. The default will be ANSI standard join. Click OK to proceed.
 - 6** Select the dimension keys:
 - `dim_customer` - customer code
 - `dim_product` - product code
 - `dim_date` - budget/forecast date (depending on which stage table you are working on).Click <--Add, then OK for each one.
 - 7** Now define the business keys. Add `customer_code`, `product_code`, and `budget/forecast_date` to the business key list, and click OK.
 - 8** For each stage table, right click in the left pane, and select **Execute Update Procedure**.

Your QAD Data Warehouse Designer screen should look like this when you have finished:



You are now ready to proceed to the next step - **Creating Fact Tables** (see "2.5 Creating Fact Tables" on page 77) for these two new stage tables.

2.5 Creating Fact Tables

- 1 Double click on the **Fact Table** object group in the left pane.
- 2 Refresh the Data Warehouse source tables in the right pane (F5).
- 3 Click and drag **stage_budget** into the middle pane. Accept the name **fact_budget**. Click **Add**.
- 4 Choose **(Build Procedure...)** from the update procedure drop down list and click **OK**.
- 5 Click **Create and Load** when asked if you wish to create and load the table now.
- 6 Select the Business Key definitions. Add **customer_code**, **product_code**, and **budget_date**.

Define Fact Business Key Columns

Column List:

- budget_quantity
- budget_sales_value
- dim_customer_key
- dim_date_key
- dim_product_key

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:

- customer_code
- product_code
- budget_date

☐ 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

- 7 Click OK.
- 8 Repeat steps (1) to (7) `stage_forecast` to create `fact_forecast` (with a business key of `customer_code`, `product_code`, and `forecast_date`).

Your QAD Data Warehouse Designer screen should look like this when you have finished:



You are now ready to proceed to the next step - **Rollup/Snapshot Fact Table** (see "2.6 Rollup/Combined Fact Table" on page 79)

2.6 Rollup/Combined Fact Table

A rollup table enables the viewing and combining of different levels of granularity in the data, such as sales, budget and forecast detail. The result is that the end user can compare, for example, sales against budget against forecast on a monthly basis.

- 1 Double click on the Fact Table object group in the left pane and refresh the Data Warehouse connection in the right pane.
- 2 Click and drag `fact_sales_detail` to the middle pane, and change the new object name to `fact_sales_analysis`.


Note: Do not make any changes to the table definition and click **Finish** when asked if you want to create and load the table now.

- 3 Because this level of granularity is no longer required, *delete* the following columns:
 - `customer_code`
 - `product_code`


- order_date
- ship_date
- dim_order_date_key
- unit_sale_price
- order_number
- order_line_no

Note: A new column has appeared - dss_fact_table_key. This is used to identify which fact table has populated a row in the rollup fact table and should not be removed. The dss_update_time field must also be present to record the time that the record was updated in the data warehouse


- 4 In the left pane click the **fact_sales_analysis** table. In the right pane open **fact_budget** and drag **budget_quantity** and **budget_sales_value** into the middle pane (within **fact_sales_analysis**).
- 5 In the right pane open **fact_forecast** and drag **forecast_quantity** and **forecast_sales_value** into the middle pane (within **fact_sales_analysis**).
- 6 In the left pane, right click on **fact_sales_analysis** and select **Create (ReCreate)**.
- 7 In the left pane, again right click on **fact_sales_analysis** and select **Properties**. In the Update Procedure field select **(Build Procedure...)** and then click **OK**.
- 8 Rollup tables are rolled up via a dimensional hierarchy. You will be given the opportunity to specify what to roll up on. From the dialog "Define rollup date dimension and column" select the following then click **OK**:
 - Date dimension - **dim_ship_date**
 - Rollup column - **ship_cal_month**:



- 9 Now select the date dimension for each of the detail tables. For **fact_sales_detail**, choose **dim_ship_date_key** and click **OK**:




For fact_budget choose dim_date_key and click OK:



Finally for fact_forecast again choose **dim_date_key** and click OK. This is required so QAD Data Warehouse Designer knows which dimension to use to rollup each of these detail fact tables.

- 10 Populate the fact rollup table by clicking the right mouse button on **fact_sales_analysis** and choosing **Execute Update Procedure**.

Your QAD Data Warehouse Designer screen should look like this when you have finished:



You are now ready to proceed to the next step - **Aggregate Tables** (see "2.7 Aggregate Tables" on page 82)


2.7 Aggregate Tables

Aggregate tables are used to improve performance. They provide a subset of the main fact table which the end user tools can navigate for a faster query time. An aggregate is typically created by the deletion of items that don't make sense when summarized and by deleting one or more of the dimension keys.




TIP: It is common practice to create two or more aggregate tables for large fact tables.

- 1 Double Click on the **Aggregate** object group in the left pane. Refresh the Data Warehouse source table in the right pane (F5).
- 2 From the right pane drag **fact_sales_analysis** into the middle pane, changing the name to **agg_sa_product**.



- 3 Click Add, then click OK from the Properties dialog.
- 4 Delete `dss_fact_table_key` so that data can be summarized from various source fact tables. Also delete `dim_customer_key` and `dss_update_time`.
- 5 Create the aggregate table.
- 6 In the left pane right click `agg_sa_product` and select Properties. Select (Build Procedure...) in the Update Procedure field and click OK on the Properties screen.
- 7 Select `dim_ship_date_key` as the date dimension key and click OK.



- 8 Update the table.




TIP: For Oracle data warehouses. If you receive an "insufficient privileges" notification in the Procedure Results, you need to grant the following privileges to the data warehouse user:

- * Create materialized view
- * Query rewrite

If you are unable to do this for any reason, contact your database administrator.

Your QAD Data Warehouse Designer screen should look like this when you have finished:




You are now ready to proceed to the next step - **Creating a Customer Aggregate** (see "2.8 Creating a Customer Aggregate" on page 84)

2.8 Creating a Customer Aggregate

This aggregate uses an alternative process to that described in **Aggregate Tables**. For this process we will create a version of the product aggregate table's metadata and create a new aggregate from this version.

- 1 With the right mouse button, click on **agg_sa_product** and select Version Control/New Version.
- 2 The following screen displays.



The image shows a 'Version Definition' dialog box. At the top, it says 'Create version for Aggregate agg_sa_product'. Below this is a checkbox labeled 'Include Associated Procedures/Scripts' which is currently unchecked. There are two text input fields: 'Version Name or Short Description:' containing the text 'copy of agg_sa_product for tutorial 2', and 'Detailed Description:' which is empty. At the bottom left, there is a 'Retain Until:' label followed by a date picker showing 'Tuesday , August 03, 2021'. At the bottom right are 'OK' and 'Cancel' buttons.

Version Definition

Create version for Aggregate agg_sa_product

☐ Include Associated Procedures/Scripts


Version Name or Short Description: copy of agg_sa_product for tutorial 2

Detailed Description:

Retain Until: Tuesday , August 03, 2021

OK Cancel

- 3 Enter a name for the new version and click OK.
- 4 In the left pane, right click on **Aggregate** and select **New Object (from Version)**:



- 5 Double click on the copy of `agg_sa_product` to select it.

Create an Object From a Version

Object Type: Aggregate

Original Object Name	Description of Version	Versioned Date
 agg_sa_product	copy of agg_sa_product for tutorial 2	2011-08-03 21:55:28.270
 agg_sa_product	Auto version on create by WhereScape Quickstart	2011-08-03 21:42:46.347


Click on the original object name of the version to be used as the basis of the create.
Then change the object name if the object still exists, and Create

New Object Name:

Unique Short Name (maximum 22 characters)

Create Cancel


- 6 Change the name and short name to `agg_sa_customer`. Click Create.



Note: Short names are used by QAD Data Warehouse Designer to derive names for associated objects (such as index, procedures, cursor, etc). The table short name is limited in size to 22 characters in Oracle and SQL Server and to twelve characters in DB2. It must be unique.

- 7 Select the new **agg_sa_customer** table in the left pane.
- 8 Delete the **dim_product_key** column, as this will be a customer and not a product based aggregate.
- 9 Browse to the Data Warehouse and from the **fact_sales_analysis** table, drag **dim_customer_key** into the middle pane.
- 10 In the left pane right click **agg_sa_customer** and select **Create (ReCreate)**.
- 11 In the left pane right click **agg_sa_customer** and select **Properties**. For the Update Procedure field select **(Build Procedure...)** and click OK.
- 12 Select **dim_ship_date_key** as the date dimension and click OK.
- 13 Right click on the table in the left hand pane and choose **Execute Update Procedure**.

Your QAD Data Warehouse Designer screen should look like this when you have finished:



CHAPTER 3

Tutorial 3 - Scheduling and Dependencies

Before you start on this chapter you should have:

- Completed ***Tutorial 1 - Basic Star Schema Fact Table*** (on page 5)
- Successfully completed ***Creating a Fact Table*** (see "**1.12 Creating a Fact Table**" on page 48)

This chapter deals with the scheduling of the data warehouse objects created in the first tutorial.

We will cover the scheduling of a job and the editing of both the dependencies and the job.

In This Chapter

3.1 Purpose and Roadmap.....	92
3.2 Creating and Scheduling a Job.....	93
3.3 Adding Tasks	94
3.4 Task Dependencies	97
3.5 Editing a Scheduled Job.....	98
3.6 Job Results	100
3.7 Diagrammatic View for Jobs	102

3.1 Purpose and Roadmap

Purpose

The scheduler allows jobs (e.g. data loads and updates) to be run in background mode and/or at a pre-determined time.

In this tutorial you will learn (i) how to set-up jobs and their associated job tasks (ii) create task dependencies, and (iii) view job results.

This tutorial focuses on creating a job to update the fact_sales_detail star-schema created in Tutorial 1.

Tutorial Environment


This tutorial has been completed using Oracle. All of the features illustrated in this tutorial are available in SQL Server, Oracle and DB2 (unless otherwise stated). Any differences in usage of QAD Data Warehouse Designer between these databases are highlighted.

Tutorial Roadmap

Step in Tutorial	Section
Create a new job for 'Daily Update'.	Creating and Scheduling a Job
Add tasks to ▪ Load load_customer ▪ Load load_order_header ▪ Load load_order_line ▪ Load load_product ▪ Update dim_customer ▪ Update dim_date ▪ Update dim_product ▪ Update stage_sales_detail ▪ Update fact_sales_detail ▪ Analyze fact_sales_detail.	Creating and Scheduling Tasks
Setup task dependencies so that an analyze of fact_sales_detail occurs after the table has been updated.	Task Dependencies
Modify scheduling and runtime options (that is, edit the job properties).	Editing a Scheduled Job
Check job results.	Job Results


The tutorial starts with the **Creating and Scheduling a Job** (see "3.2 Creating and Scheduling a Job" on page 93) section.

3.2 Creating and Scheduling a Job

To schedule a job click on the scheduler window button . This will open the scheduler window. A new job can be initiated by selecting the File/New Job menu option.

The new job dialog will appear.

- 1 Change the job name to 'Daily Update' and enter in the description.



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:

- 2 Click OK.

You are now ready to proceed to the next step - **Adding Tasks** (see "3.3 Adding Tasks" on page 94)

3.3 Adding Tasks

The task selection window contains an object tree in the right hand pane. Objects are selected from this tree and added to the scheduled list of tasks in the left hand pane. Perform the following actions to schedule an update of our fact table and dimensions.


- 1** Open the object tree by double clicking on the **All Objects** project in the left hand pane.
- 2** Double click on the **Load Table** object group.
- 3** Double click on **load_product**, **load_customer**, **load_order_line** and **load_order_header**. Note that as each object is double clicked it is added to the right hand pane.
- 4** Double click on the **Dimension** object group.
- 5** Double click on **dim_customer**, **dim_date** and **dim_product**. As each object is double clicked it is added to the right hand pane. We do not add the date views since they do not alter, only the underlying date dimension does.
- 6** Double click on the **Stage Table** object group to expand it and then double click on **stage_sales_detail** to add this object to the right hand pane.
- 7** Double click on the **Fact Table** object group to expand it and then double click on **fact_sales_detail** to add this object to the right hand pane.
- 8** Double click again on **fact_sales_detail** to add a second copy of this object to the right hand pane.
- 9** Position on the second **fact_sales_detail** object in the right hand pane and use the right mouse to bring up the pop-up menu. Select the **Analyze** option from this menu:

 load_customer	Process	11.10
 load_order	Drop	12.10
 load_order	Create	13.10
 dim_customer	Truncate	14.10
 dim_date	Initial Build	15.10
 dim_product		16.10
 stage_sales	Drop All Indexes	17.10
 fact_sales	Pre Drop Indexes	18.10
 fact_sales	Load	19.10
	Custom	
	Update	
	Execute	
	Process	
	Build Indexes	
	Build All Indexes	
	Stats	
	Quick Stats	
	Analyze	
	Quick Analyze	

The 'Order' column defines the basic dependencies of the tasks. If the two numbers are the same, then the tasks can run at the same time. In this example no tasks will run at the same time. The job will process the tasks sequentially.

Click the 'Group object types' button. You will note that the order number for tasks of the same type now have the same two numbers. This will allow objects of the same type to run concurrently. (i.e. all the load tables can be processed at the same time if there are sufficient processing threads).

Your task selection window should now look like the following:



Notice that the tasks all have an action of Process with the exception of the last task which is set to 'Analyze'.

The fact table `fact_sales_detail` has two actions. The first will process and update the table, the second will analyze the table. At present these two actions can run at the same time. They should however be sequential. We could alter the order of the second task by using the right mouse menu option 'Increase the order'. This would be the normal method, but we will leave these two tasks with the same order and address the sequence of events in the next section.

Choose OK to close.

You are now ready to proceed to the next step - **Task Dependencies** (see "**3.4 Task Dependencies**" on page 97).

3.4 Task Dependencies

A scheduled job that is in a 'Hold' or 'Waiting' state can have its task dependencies altered. To alter the dependencies for our newly defined job proceed as follows:

- 1 Click on the 'All Jobs' button in the toolbar. Our daily update job will be displayed in the top pane.
- 2 Position over the job name 'Daily Update' and using the right mouse pop-up menu select **Edit Dependencies**. A list of the current task dependencies will be displayed. You will see that the final two dependencies are from stage_sales_detail to each of the fact table tasks:

[illegible]

- 3** Position on the 'Parent task' for the last dependency and using the pop-up right mouse menu select **Modify Dependency**:

load_order_header	Process	-->	dim_date	Process
load_order_header	Process	-->	dim_product	Process
dim_customer	Process	-->	stage_sales_detail	Process
dim_date	Process	-->	stage_sales_detail	Process
dim_product	Process	-->	stage_sales_detail	Process
stage_sales_detail	Process	-->	fact_sales_detail	Process
stage_sales_detail	Process	-->	fact_sales_detail	Analyze

Add Dependency
Modify Dependency
Delete Dependency

- Change the Parent task from `stage_sales_detail (Process)` to `fact_sales_detail (Process)`. Click OK to record the change. An example of this change is shown in the screen shot below:

Task Dependency

Enter the Parent and Child Dependency. The child task will not be actioned until the parent has successfully completed.

Parent: --> Child:

OK Cancel

- Examine the new dependency list and see that the fact processing will now occur after the stage processing and the fact analyze will occur after the fact processing.

We are now ready to release the job, which is done in the next section - **Editing a Scheduled Job** (see "3.5 Editing a Scheduled Job" on page 98).

3.5 Editing a Scheduled Job

Our job is now all set-up and ready to be released. We need to edit the job and change it from a held job to one that the scheduler can action. Proceed as follows:

- Click on the 'All Jobs' button in the toolbar. Our daily update job will be displayed in the top pane. Note that it is in an 'On Hold' state.
- Position over the job name 'Daily Update' and using the right mouse pop-up menu select **Edit Job**. The job definition screen will re-appear.
- Change the **Frequency** to 'Once and Hold'. This will result in the job being run and then a copy of the job being placed back in an 'On Hold' state so that it may be rescheduled for some future processing. Note that other options exist under Frequency including 'Daily', 'Custom' etc.
- Change the **Start Time** to be 2 minutes from now.
- Change the **Max Threads** counter to 2. This will allow two tasks to run concurrently. This may not be a big help here, as the run should be very quick.

Job Definition

Job Name:

Description:

Frequency:

Start Date:

Start Time:

Maximum Threads: Inactive Wait Interval (seconds):

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:

- 6 Click OK to save the changes.
- 7 Click on the 'All Jobs' button in the toolbar. Our daily update job will be displayed in the top pane. Note that its state should now be 'Waiting' or maybe 'Running'. If the job is in the 'Running' state we can double click on the Job name to see the state of the individual tasks.



TIP: If you don't need to change a job and wish to run it immediately, choose **Start the Job** from the job's popup menu.

- 8 If the job does not go into a 'Running' state after 30 seconds, check that a scheduler is running by clicking on the 'Scheduler Status' in the scheduler menu:

Audit and Scheduler log									
Type	Name	Host	Status	Started	Last Status	Stopped	Sampl...	Version	Message
Wind...	WIN0001	SSERVER	Running	2011-08-04 08:...	2011-08-04 09:...		30	6001000	

9 If no schedulers are running, refer to the Setup and Administration Guide on how to start a scheduler.

We are now ready to proceed to the next section - **Job Results** (see "3.6 Job Results" on page 100)

3.6 Job Results

Once a job has completed, or in fact while it is running, we can check on the results of each of the tasks by proceeding as follows:

- 1 Click on the 'All Jobs' button in the toolbar. Our daily update job will be displayed in the top pane. Note that if the job has started or is completed there will be two entries. One is in an 'On Hold' state and one is in a 'Completed', 'Running' or 'Failed' state.
- 2 Position over the job with name 'Daily Update' and in a 'Completed', 'Running' or 'Failed' state and double click the left mouse button. The individual tasks within the job will be displayed.
- 3 Position on the `fact_sales_detail` task with action `Process` and double click the left mouse button. The messages returned from this task will now be displayed. These messages should include information on any indexes dropped and created.

Your QAD Data Warehouse Designer screen should look like this when you have finished:

The screenshot shows the QAD Data Warehouse Designer Scheduler window. The interface includes a menu bar (File, View, Auto, Monitor, Logs, Jobs, Scheduler, Tools, Window, Help) and a toolbar with icons for various actions. The 'Jobs' tab is active, displaying a table of jobs. Below it, the 'Tasks' tab is active, displaying a table of tasks. The status bar at the bottom shows the database time, development environment, user ID, and various system parameters.

Jobs Table:

Job	Status	Seq	Start	Finish	Elap...	OK	Info	Det...	Warn	Error	Who
Daily Update	On Hold	80	2011-08-04 09:5...								WQ
Daily Update	Comple...	81	2011-08-04 09:5...	2011-08-04 09:5...	00:00	11	26				WQ

Tasks Table:


Task	Acti...	Status	Seq	Start	Finish	El...	I...	D...	W...	Result
load_product	Proc...	Com...	81	2011-08-0...	2011-08-0...	00:...	4			9 rows loaded into load_product
load_customer	Proc...	Com...	81	2011-08-0...	2011-08-0...	00:...	4			6 rows loaded into load_customer
load_order_line	Proc...	Com...	81	2011-08-0...	2011-08-0...	00:...	4			21 rows loaded into load_order_line
load_order_h...	Proc...	Com...	81	2011-08-0...	2011-08-0...	00:...	4			9 rows loaded into load_order_header
dim_customer	Proc...	Com...	81	2011-08-0...	2011-08-0...	00:...	2			dim_customer updated. 0 new records. 6 records ...
dim_date	Proc...	Com...	81	2011-08-0...	2011-08-0...	00:...	2			Date Dimension updated from Nov 25 1997 12:00A...
dim_product	Proc...	Com...	81	2011-08-0...	2011-08-0...	00:...	2			dim_product updated. 0 new records. 9 records up...
stage_sales_d...	Proc...	Com...	81	2011-08-0...	2011-08-0...	00:...	2			stage_sales_detail updated. 21 new records. 0 reco...
fact_sales_det...	Proc...	Com...	81	2011-08-0...	2011-08-0...	00:...	2			fact_sales_detail updated. 0 new records. 21 record...
fact_sales_det...	Anal...	Com...	81	2011-08-0...	2011-08-0...	00:...				Table fact_sales_detail analyzed

Status Bar: Database time at last action was 09:58:28 | Development (WslWarehouse) | UserId: dbo | Browse: No source system | CAP | NUM | SCRL

We are now ready to proceed to the next section - **Diagrammatic View for Jobs** (see "3.7 Diagrammatic View for Jobs" on page 102)

3.7 Diagrammatic View for Jobs

QAD Data Warehouse Designer provides the ability to diagrammatically view the job dependencies for the job you have created. To switch to the diagrammatic view, click on the  button. Click the button again to bring up the diagram selector box. Select the job name from the **Select a Job to show dependencies** drop-down list box - in this instance **Daily Update**:



The **Diagram Selection** dialog box is shown. It contains several options for selecting a diagram to view:

- Select a Fact Table to be Drawn as a Star Schema:** (Empty dropdown)
- Select a Table to Draw Source Tracking:** (Empty dropdown)
- Select a Table to Draw Data Join Track Back:** (Empty dropdown)
- Select a Table to Draw a Track Forward Diagram:** (Empty dropdown)
- Select a Table to Show Linked Tables:** (Empty dropdown)
- Select a Job to Show Dependencies:** (Dropdown with **Daily Update** selected)


On the right side, there are additional options:

- Group:** (Empty dropdown)
- Project:** (Empty dropdown)
- Hops:** (Empty dropdown)

Buttons for **OK** and **Cancel** are located in the top right corner.

Click OK.

The diagram looks like this:



CHAPTER 4

Tutorial 4 - Complex Dimensions and Hierarchies

Before you start on this chapter you should have:

- Completed ***Tutorial 1 - Basic Star Schema Fact Table*** (on page 5)
- Successfully completed ***Creating a Fact Table*** (see "**1.12 Creating a Fact Table**" on page 48)

This chapter deals with fine tuning the data warehouse by creating complex dimensions and hierarchies.

In This Chapter

4.1 Purpose and Roadmap.....	106
4.2 Creating a Slowly Changing Dimension.....	107
4.3 Multiple Source Table Dimension	114
4.4 Creating a Dimension Hierarchy.....	125

4.1 Purpose and Roadmap

Purpose

This tutorial will walk you through the process to:

- Create a slowly changing dimension
- Creating a complex dimension with multiple table sources
- Adding hierarchies to a dimension for external maintenance and for use in Analysis Services cubes.

In short, this tutorial alters the existing customer and product dimensions. The customer dimension is converted to a slowly changing dimension and the product dimension has its content enriched from additional data sources. Hierarchies are built on all dimensions that will be used in the next tutorial.

Tutorial Environment

This tutorial has been completed using Oracle. All of the features illustrated in this tutorial are available in SQL Server, Oracle and DB2 (unless otherwise stated). Any differences in usage of QAD Data Warehouse Designer between these databases are highlighted.

Tutorial Roadmap

This tutorial works through a number of steps. These steps and the relevant section within the chapter are summarized below to assist in guiding you through the tutorial.

Step in Tutorial	Section
Convert the customer dimension to a slowly changing dimension.	Creating a slowly changing dimension
Add additional data sources to the product dimension	Multiple source table dimension
Create hierarchies for the following tables: ▪ dim_date ▪ dim_customer ▪ dim_product	Creating a dimension hierarchy

This tutorial starts with the section ***Creating a Slowly Changing Dimension*** (see "***4.2 Creating a Slowly Changing Dimension***" on page 107)

4.2 Creating a Slowly Changing Dimension

The process of creating a slowly changing dimension is largely the same as creating a normal dimension. Two additional questions are asked during the dimension creation process when the 'Slowly changing dimension' button is chosen during the dimension create. In this section we will cover the more common scenario of changing an existing normal dimension to a slowly changing dimension.

The dimension `dim_customer` created in tutorial one will be changed to a slowly changing dimension.

- 1 Using the right mouse menu when positioned on the object `dim_customer` select **Properties**.
- 2 On the dimension properties change the Update procedure drop down to select **(Build Procedure...)**.
- 3 Use the **Table Type** drop down to select **Changing Dimension**:

dim_customer Table Definition

Properties

Storage

Language Mapping

Purpose

Concept

Grain

Examples

Usage

Notes

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)

Get Key Function:


Mnemonic (EUL):

Timestamps


Metadata Structure Changed: Database Created: Database Altered:

- 4 Click OK to leave the properties page.


- 5 You will be asked if you wish to version and replace the update and get key procedures. Answer **Yes** to both prompts.



- 6 A Procedure Build Type dialog will appear. Select **Cursor**.




- 7 The 'Define dimension business key(s)' dialog will be presented. The existing business key 'code' should already be the default value so click **OK** to proceed to the next screen.
- 8 The 'Define slowly changing dimension columns' dialog will now be presented. Multiple columns can be selected to be handled as slowly changing. Select 'cname' and add it to the list of slowly changing columns. Double click on the column 'cname':




- 9 Click OK. Refer to the Dimensions chapter for an explanation of slowly changing dimensions.


- 10** A dialog box will appear indicating that a number of additional columns will need to be added to the dimension table in order to support a slowly changing model. The table can be Altered (i.e. the columns added to the table) or re-created. As we have fact tables that use this dimension we cannot re-create the table. To do so would make all of the joins in the fact tables to this dimension invalid. Therefore we will alter the dimension. Click the 'Alter' button.



- 11 A dialog will now be presented with the SQL commands that will be executed to add the new columns and set default values for the `dss_version` and `dss_current_flag` columns. Click the 'Alter Table' button to alter the table in the database. It is worth noting that whenever a database table is altered from within QAD Data Warehouse Designer the SQL commands should be reviewed. These commands can be changed if a different result is required.



- 12 A message box will appear informing you that the table was altered in the database. Click OK.
- 13 A dialog box will appear asking if you want to allow NULL support for the slowly changing columns. The normal response would be No. Refer to the Dimensions chapter for a detailed explanation. Click the 'No' button.




- 14** The results dialog will show the indexes that were created/re-created for the dimension.

Note: there is a new index with a suffix of _SC to support the slowly changing dimension.

The dimension has now been converted to a slowly changing dimension. If a customer now has a name change a new version of the customer record will be created to allow the tracking of the customer by both the old and new name.

Your QAD Data Warehouse Designer screen should look like this when you have finished:




Note: This conversion to a slowly changing dimension has changed the get key function for the dimension. This function is called from the stage tables that use this dimension. We must now rebuild the update procedures for the stage tables that use this dimension.

Rebuild Stage Update Procedures

As the customer dimension is used in the stage tables stage_sales_detail, stage_budget and stage_forecast we will need to rebuild these procedures. Proceed as follows:

- 1 Select the **stage_sales_detail** table and using the right mouse menu choose **Properties**.
- 2 Use the Update procedure drop down to select (**Build Procedure...**).
- 3 Click OK to leave the properties page.
- 4 Answer YES to the procedure versioning question.
- 5 Select **Cursor** for the procedure type.
- 6 Click OK on the screen informing of multiple source tables.
- 7 Accept the previous entries by clicking OK on the cursor mapping screen.

- 8 Click OK on the business key join for each of the dimensions. The previously chosen business key should be the default value provided.
- 9 An additional dialog will appear after the customer dimension join. Select **ship_date** from the date list and click OK.



- 10 Click OK on any other dimension joins and the business key from the stage table.
- 11 The update procedure will now be rebuilt to handle the fact that the customer dimension is now slowly changing. Repeat steps (1) through (10) for the other stage tables **stage_budget** and **stage_forecast**. Use the dates **budget_date** and **forecast_date** to track the dimension changes.

You have successfully converted the dimension **dim_customer** to a slowly changing dimension, and made all dependent changes.

You are now ready to proceed to the next section - **Multiple Source Table Dimension** (see "4.3 Multiple Source Table Dimension" on page 114)


4.3 Multiple Source Table Dimension

Dimensions typically get their information from multiple sources. A common scenario is to have a series of codes that relate to the dimension. The descriptions of these codes are often stored in a code lookup table. The following example will clarify the practice of producing a dimension from multiple source tables.

The dimension **dim_product** created in tutorial one will be enhanced to provide additional descriptions for the code values already present.


- 1 In the right hand pane, browse to the tutorial connection that was created in tutorial one. Click on the glasses or select the Browse/Source tables menu option. The tutorial tables should now be shown in the right hand pane.

- 2 We need to acquire a number of additional source tables from the tutorial database. Double click on the **Load Table object group** in the left hand pane. This will list all the load tables in the middle pane and make the middle pane a drop target for additional load tables.
- 3 Select the **prod_group** table from the right hand pane and holding the left mouse down drag to the middle pane. Click **Add** to add the new object and then **OK** on the properties screen. Select the 'Create and Load' button.
- 4 Repeat steps (2) and (3) above to bring in and load from the **prod_subgroup** and **prod_line** tables.
- 5 In the left hand pane click on **dim_product**. This will display all of the dimension columns in the middle pane and make the middle pane a drop target for additional dimension columns. We will be adding descriptions to it.
- 6 Browse to the Data Warehouse connection. The right hand pane should now show the data warehouse tables. You can position the mouse in this right hand pane and select the right menu option 'Filter - Load' to restrict the display to just load tables.
- 7 Expand **load_prod_group**, **load_prod_subgroup** and **load_prod_line** by double clicking on each table name.



- 8 Drag **group_description** from **load_prod_group** into the middle pane. This will add it to the product dimension columns.
- 9 Drag **subgroup_description** from **load_prod_subgroup** into the middle pane. This will add it to the product dimension columns.

- 10 Drag `line_description` from `load_prod_line` into the middle pane. This will add it to the product dimension columns.
- 11 Review the product columns displayed in the middle pane. Expand the 'source table' column to see that we now have four different tables that contribute to the dimension:



Column Name	Display Name	Data Type	Source Table
dim_product_key	dim product key	integer ident...	
code	code	numeric(6)	load_product
description	description	varchar(64)	load_product
prod_line	prod line	varchar(24)	load_product
line_description	line description	varchar(64)	load_prod_line
prod_group	prod group	varchar(24)	load_product
group_description	group descript...	varchar(64)	load_prod_group
subgroup	subgroup	varchar(24)	load_product
subgroup_description	subgroup desc...	varchar(64)	load_prod_subgro...
dss_update_time	dss update time	datetime	

- 12 Using the right mouse menu when positioned on the object `dim_product` in the left hand pane select **Validate against database**. This option will compare the metadata as displayed within QAD Data Warehouse Designer with the physical table `dim_product` as it exists in the database. We have added new columns to the metadata in the steps above, but the physical table has not been changed. A message will appear in the middle pane as follows:






Table Name	Results
dim_product	meta->line_description - varchar(64), meta->group_description - varchar(64), meta->subgroup_description -


- 13 We now need to alter the physical table in the database. The message in the middle pane shows that the metadata has additional columns not present in the dss table. We alter the table by positioning on the `dim_product` name in the middle pane and using the right mouse 'Alter Table' menu option:




- 14 An alter table commands dialog will appear with the SQL commands that will be used to alter the database table. Click the 'Alter Table' button.
- 15 A message will display advising that the table was altered. Click OK to clear the message.
- 16 Using the right mouse menu when positioned on the object `dim_product` in the left hand pane select Properties.
- 17 Click Rebuild.
- 18 A Procedure Build Type dialog will appear. Select Cursor.



- 19 A dialog will now appear asking you to define the primary source table and to choose between a join of the source tables or a series of lookups. Refer to the Dimensions chapter for an explanation on these two choices. The `load_product` should have been selected as the primary source table. If it has not been selected then select it. Click the **Lookup** button.



- 20 The 'Define dimension business key(s)' dialog will be presented. The existing business key 'code' should already be the default value so click **OK** to proceed to the next screen.
- 21 A message is displayed, asking you to ensure that any join syntax be removed from the statement in the dialog to follow. Click **OK**.



22 Click OK on the cursor mapping screen.

Provide Cursor Mapping

Cursor Name (if applicable):

Edit the where clause if required.

Source Tables:

- load_product
- load_prod_line
- load_prod_group
- load_prod_subgroup

From and Where Clause:


```
FROM
  load_product
```

☒ ANSI join code generated

Table A: Table B:

☐ Word Wrap Displayed Code

- 23** The Dimension lookup definition screen should now appear. Two drop down boxes are available. The left hand drop down shows the columns for the table chosen as the primary source table. The right hand drop down shows the columns for the source table we wish to lookup. There is a reminder above the left hand drop down informing you which table is being looked up.



The image shows a software dialog box titled "Dimension Lookup Table Definition". It contains instructions to create a lookup condition, a label for the lookup ("Lookup for line_description"), two checkboxes for "Multiple Lookups Required on This Table" and "Post Update", and two dropdown menus. The left dropdown is labeled "load_product" and the right is labeled "load_prod_line", both showing "prod_line" as the selected option. Below these is a text area for the "Where Clause for the Lookup" containing the SQL snippet "WHERE prod_line = @v_load_prod_line". An "Add Join" button is positioned above the text area. At the bottom right are "OK" and "Cancel" buttons.

Dimension Lookup Table Definition

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.

Lookup for line_description

☐ Multiple Lookups Required on This Table

☐ Post Update

load_product

load_prod_line

prod_line

prod_line

Where Clause for the Lookup:

WHERE prod_line = @v_load_prod_line

Add Join

OK Cancel

Dimension Lookup Table Definition

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.

Lookup for group_description

☐ Multiple Lookups Required on This Table

☐ Post Update

load_product

prod_group

load_prod_group

prod_group

Where Clause for the Lookup:

WHERE prod_group = @v_load_prod_group

Add Join

OK

Cancel

Dimension Lookup Table Definition

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.

Lookup for subgroup_description

☐ Multiple Lookups Required on This Table

☐ Post Update

load_product: subgroup

load_prod_subgroup: subgroup

Add Join

Where Clause for the Lookup:

WHERE subgroup = @v_load_subgroup

OK Cancel

Although not being used at this point there are a number of features in these dialogs which are worth noting. The first is that there is a check box to allow a table to be the source of multiple lookups. This would be used where a generic code table was used to lookup descriptions and we would need to make multiple lookups on the same table to get different descriptions. The second is that procedure variables are available in the drop down list of the primary load table. Using this feature it is possible to make lookups that are dependent on the results from previous lookups. This is done by selecting the columns that would have been populated by the earlier lookups.

The lookups require that a join be defined for the two tables involved. The relevant joins for our example are as follows. Step through the dialog boxes making these joins. First select the column from each drop down list and then click the 'Add Join' button.

Lookup table	Primary src (load_product)	Lookup column
load_prod_group	prod_group	prod_group
load_prod_subgroup	prod_group prod_subgroup	prod_group prod_subgroup
load_prod_line	prod_line	prod_line

Note: The load_prod_subgroup lookup will require two joins. After each join is made click the OK button to move to the next dialog.


- 24 Once the lookups have been completed the procedure will be generated and the results dialog will show the indexes that were created/re-created.
- 25 Using the right mouse menu when positioned on the object **dim_product** in the left hand pane select **Execute Update Procedure**. The dimension will now be refreshed.
- 26 Using the right mouse menu when positioned on the object **dim_product** in the left hand pane select **Display Data**. Provided the lookups were done correctly, the descriptions should now be populated.

Data display for dim_product					
dim_produc...	code	description	prod_line	prod_group	subgroup
 0		Unknown	Unknown	Unknown	Unknown
 1	1001	Stripped ball	Toy	Balls	Large
 2	1002	Front load...	Toy	Equipment	Tractor
 3	1003	Building bl...	Edu	Blocks	Foam
 4	1004	Talking co...	Toy	Talking	Jar
 5	1005	yellow golf...	Toy	Balls	Golf
 6	1006	red golf ball	Toy	Balls	Golf
 7	1007	blue gold c...	Toy	Clubs	Golf
 8	1008	plastic buil...	Edu	Blocks	Plastic
 9	1009	tool case	Toy	Equipment	Tools

Oracle Procedure Invalidation

If running this as a tutorial in an *Oracle Data Warehouse* then the procedures `update_stage_sales_detail`, `update_stage_budget` and `update_stage_forecast` will have been invalidated. These procedures all use the function `get_dim_product_key`. This function was re-compiled as part of the process above, but has not changed in structure.

Select the menu option Tools/Re-compile Invalid Procedures:



Click **Yes** to proceed. The three procedures will be re-validated and the results shown in the middle pane.


You are now ready to proceed to the next section - ***Creating a Dimension Hierarchy*** (see "***4.4 Creating a Dimension Hierarchy***" on page 125)

4.4 Creating a Dimension Hierarchy

Dimensions often have a number of hierarchical levels. An example is the date dimension which has a natural hierarchy of year, month and day. QAD Data Warehouse Designer allows the definition of a hierarchy against a dimension. A dimension may have multiple hierarchies defined. These dimension hierarchies are used in the generation of Analysis Services cubes. Every dimension in a cube must have a hierarchy with at least one element in it. Hierarchies can also be accessed and the underlying data maintained with the Hierarchy maintenance utility shipped with the enterprise version of QAD Data Warehouse Designer. This utility allows a user to maintain a hierarchy structure in a different database or schema so that the data warehouse can use this structure as input to the dimension. See the hierarchy section in the Dimension chapter for more details.


We will add a hierarchy to the dimensions `dim_product` and `dim_customer` created in the first tutorial. We will also check that a hierarchy exists for the date dimension and if not present, add one.

- 1 Using the right mouse menu when positioned on the object `dim_customer` in the left hand pane and select **Hierarchies/List Hierarchies**.




- 2 A list of any existing hierarchies will be shown in the middle pane. Initially there are no hierarchies defined.
- 3 Using the right mouse menu when positioned on the object `dim_customer` in the left hand pane, select **Hierarchies/Add Hierarchy**.

- 4 A dialog will appear asking for a hierarchy name and the hierarchy elements. Enter **customer** for the name and add **state**, **city**, **address** and **code** as hierarchy elements. Click OK to exit the add hierarchy dialog.



- 5 Repeat steps (1) through (4) for **dim_product** defining a top down hierarchy of **line_description**, then **group_description**, **subgroup_description**, and **code**.
- 6 Check to see if a hierarchy exists for **dim_date**. If no hierarchy exists then add a top down hierarchy of **cal_year**, **cal_month** and **calendar_date**.



The screenshot shows a window titled "Hierarchies for dim_date". Inside is a table with four columns: "Dimension Name", "Hierarchy Name", "Description", and "Maintained". The first row contains the data: "dim_date" (with a cube icon), "calendar", "Default hierarchy used for ...", and "N". The table has several empty rows below it. At the bottom of the window is a scrollbar.

Dimension Name	Hierarchy Name	Description	Maintained
 dim_date	calendar	Default hierarchy used for ...	N

As mentioned above these hierarchies are utilized by other processes within the data warehouse. In this case we will use the hierarchies created in the building of Analysis Services cubes in the next tutorial.

CHAPTER 5

Tutorial 5 - Analysis Services Cubes

Before you start on this chapter you should have:

- Completed ***Tutorial 1 - Basic Star Schema Fact Table*** (on page 5)
- Completed ***Tutorial 2 - Rollup Fact Tables, ASCII File Loads, Aggregates*** (on page 63)
- Completed ***Tutorial 4 - Complex Dimensions and Hierarchies*** (on page 105)

This chapter deals with fine tuning the data warehouse by adding Analysis Services cubes.

In This Chapter

5.1 Purpose and Roadmap.....	130
5.2 Creating an OLAP Cube Object.....	131
5.3 Adding a Measure Group	155
5.4 Cube Connections for Other Databases	167

5.1 Purpose and Roadmap

Purpose

This tutorial will walk you through the process to:

- Create an analysis services cube.

In short, this tutorial uses existing fact tables to generate a "multi-measure group cube" based on fact_sales_detail and fact_budget. It also shows you how to query the cube.

Tutorial Environment

This tutorial has been completed using Microsoft SQL Server. All of the features illustrated in this tutorial are available in SQL Server, Oracle and DB2 (unless otherwise stated) using Analysis Services Cubes. Any differences in usage of QAD Data Warehouse Designer between these databases are highlighted. Refer to ***Cube Connections for Other Databases*** (see "**5.4 Cube Connections for Other Databases**" on page 167) for more information on configuring the DataWarehouse connection for Oracle or DB2 data warehouses.

Tutorial Roadmap

This tutorial works through a number of steps. These steps and the relevant section within the chapter are summarized below to assist in guiding you through the tutorial.

Step in Tutorial	Section
Create an Analysis Services cube from fact_sales_detail	Creating a Cube Object
Add a measure group to the cube based on fact_budget	Adding a Measure Group object

This tutorial starts with the section ***Creating an OLAP Cube Object*** (see "**5.2 Creating an OLAP Cube Object**" on page 131)

5.2 Creating an OLAP Cube Object

The process of creating a cube object is largely the same as creating any other data warehouse object. By dragging a fact or aggregate table into an OLAP Cube drop target an OLAP Cube and OLAP Dimensions are defined.

OLAP Cubes are more complex than some other objects in QAD Data Warehouse Designer, primarily due to the functionality available in Microsoft Analysis Services.

This tutorial covers creating a basic cube. Refer to the chapter on Analysis Services Cubes in the QAD Data Warehouse Designer User Guide for more information.


In this step we will create an OLAP Cube from the *fact_sales_detail* table. Before we start we need to have access to a Microsoft Analysis Services server with appropriate security rights to create an OLAP database.

Analysis Services is supplied with Microsoft SQL Server. You will also need to know the name of this server.


Note: This tutorial is only possible if an Analysis Services server is available

If the Analysis Services server is located on another machine then the Analysis Services client software will need to be loaded onto the computer running QAD Data Warehouse Designer.

- 1 Create a new **connection** for the Analysis Services server. Position on the 'Connection' object group in the left hand pane and using the right mouse menu select **New Object** as follows:



Enter a name of *Cubes* for the connection and click **Add**:



- 2 The connection properties screen will now appear. Select the connection type: **Microsoft Analysis Server 2005+**. Enter the **name** of the server for the Analysis Server. Leave the username and password blank. Enter the server and database details for your Analysis Services server. A sample is shown next:

The screenshot shows the 'Cubes Definition' dialog box with the 'Connection Properties' tab selected. On the left is a 'Notes' text area. The main area contains the following fields:


- ☐ Data Warehouse
- Connection Name: Cubes
- Connection Type: Microsoft Analysis Server 2005+ (dropdown menu)
- Analysis Server 2005+: SERVER\SQLSERVER
- Analysis Server User ID: (empty text box) Password: (empty text box)

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

Then click OK.

Note: Microsoft Analysis Services 2005 and 2008 use "Microsoft Analysis Server 2005+".

- 3 Expand the Connection object group in the left hand pane and double click on the DataWarehouse connection to bring up the Properties dialog.



- 4 The properties dialog will now show some additional columns at the bottom of the screen. These are required so that the data warehouse can be used as a source for the Analysis Services cubes. Some other fields are determined by the type of relation database being used for the data warehouse. These fields are:


- Provider type
- Server
- Database

For a SQL Server data warehouse:

- Set **Provider Type** to *SQLOLEDB*
- Set **Server** to the SQL Server *server name*
- Set **Database** to the SQL Server *database name*

See **Cube Connections for Other Databases** (see "5.4 Cube Connections for Other Databases" on page 167) for more details on Oracle and DB2.

A SQL Server sample is shown next:




The image shows a 'DataWarehouse Definition' dialog box. On the left is a 'Notes' pane. The main area contains the following fields and controls:

- ☒ Data Warehouse
- Connection Name: DataWarehouse
- Connection Type: Database
- ODBC Source: WstWarehouse
- Provider Name:
- Database ID (SID):
- Database Link ID:
- Extract User ID: Password:
- Administrator User ID: Password:
- Default Load Type: Database link load
- SSIS Connection String: Reset
- Data Type Mapping Set: (Default) As a Cube Data Source
- Provider Type: SQLOLEDB (OLE DB for SQL Server)
- Server: SSERVER
- Database: WstWarehouse

Buttons at the bottom: OK, Cancel, Help.

Then click OK.

- 5 Now you need to browse the *DataWarehouse* connection to see available fact tables. From the **Browse** menu, select *DataWarehouse* then **Change Connection**. The following dialog displayed:



The image shows a 'List Source Tables Connection' dialog box. It has a title bar with a close button. Inside, there are fields for 'Connection:' (set to 'DataWarehouse'), 'User ID:', and 'Password:'. Below these is a 'Filter' section containing 'Schema:' (set to 'dbo'), 'Name:' (set to '(None)'), 'Object Types' (with checkboxes for 'Tables' and 'Views' checked, and 'System Objects' unchecked), 'Group:' (set to '(All)'), and 'Project:' (set to '(All)'). At the bottom are buttons for 'Refresh Current', 'OK', and 'Cancel'.

List Source Tables Connection

Connection: DataWarehouse

User ID:

Password:

Filter

Schema: dbo

Name: (None)

Object Types

☒ Tables ☒ Views ☐ System Objects

Group: (All)


Project: (All)

Refresh Current OK Cancel


Choose *DataWarehouse* from the **Connection** drop down and click OK.

Note: For SQL Server the Data Warehouse schema must be *dbo*.

The following browse pane is displayed:




- 6 Double click on the OLAP Cube object group in the left hand pane to list all cubes in the middle pane. This makes the middle pane a drop target for cubes. Your screen should now look like this:




Note: As OLAP Cubes have not been created yet, the middle pane is empty.

- 7 Drag *fact_sales_detail* from the browser (right hand pane) and drop it into the middle to create a new OLAP Cube object. Give the OLAP Cube a name of *sales_analysis* as follows:



Click ADD.

- 8 QAD Data Warehouse Designer will now cycle through each dimension associated with that fact table and will create an OLAP Dimension object for each, displaying the following dialog first:




The screenshot shows a dialog box titled "Add a New Metadata Object" with a close button in the top right corner. Inside the dialog, there is instructional text: "Define the Type and Name of the New Object." and "Specific information for each object type is defined in subsequent screens." Below this, there are two input fields: "Object Type:" with a dropdown menu showing "Olap Dimension", and "Object Name:" with a text box containing "odim_customer". There is also a checkbox labeled "Include Attribute Relationships" which is currently unchecked. At the bottom right, there are two buttons: "ADD" and "Cancel".

Note: If you wish to include Attribute Relationships in Analysis Services for this dimension, click on the Include Attribute relationships check box.

Click ADD for the *Customer* dimension.

The following dialog appears, prompting you to select the attributes to be included in the *Customer* OLAP dimension.



The attributes available for selection are in the left hand column. To select an attribute, click on the attribute in the left hand column and click >. This will move the attribute to the right hand column.

To de-select an attribute, click on the attribute in the right hand column and click <. This will move the attribute to the left hand column.

- 9 Repeat step 8 for the remaining OLAP Dimensions:

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:

☐ Include Attribute Relationships

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:

☐ Include Attribute Relationships

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:


☐ Include Attribute Relationships

- 10** The Manage cube measures dialog is displayed next. Remove all non-measure columns (that is, columns that cannot be aggregated) from the **Measure** list by highlighting them and clicking the **Remove** button.

The columns to remove are:

- Order Number
- Order Line No
- Product Code
- Customer Code


At this stage, your dialog should look like this:



- 11** Choose to create some date based calculated members using OLAP functions by selecting the **Month to date**, **Year to date**, and the **Use OLAP functions** check boxes. The calculated measures drop downs are now displayed. Select the required **date hierarchy** information. Specifically:
- for Date dim, select *dim_order_date*


- for Dim hie, select *calendar*
- for Month level, select *order cal month*
- for Year level, select *order cal year*

After selecting the Month level, you will be asked to confirm changing the attribute type to Months for the the *order cal month* attribute in Analysis Services as follows:



Click Yes.

Similarly, selecting the Year level, you will be asked to confirm changing the attribute type to Years for the the *order cal year* attribute in Analysis Services as follows:



Again, click Yes.

Your Manage cube measures dialog should now look like this:

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
unit sale price
quantity
sales value
tax

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:


Date Hierarchy:

Month Level:


Year Level:

Click OK to complete definition of the cube.

- 12 Right click on the OLAP Cube *sales_analysis* in the left pane and choose Properties as follows:



- 13 Click on the down arrow for the Cube Database name field and choose the option (Define New Cube Database), which will bring up the Cube database dialog box. Enter a new cube Database name of *Tutorial5* as follows:



Now click OK.

Change the Cube publish name to be *sales_analysis* and the Cube database connection to be *Cubes* as follows:

Cube sales_analysis Definition

Internal Cube Name: sales_analysis

Cube Publish Name: sales_analysis

Cube Description:

Cube Database Connection: (Analysis Services) Cubes

Cube Database Name: Tutorial5

Data Source Connection: (Data Warehouse) DataWarehouse

Data Source Provider Type: SQLOLEDB

Data Source Server: SSERVER

Data Source Database: WslWarehouse

Post Create XML/A Script: (None) Edit

Post Update XML/A Script: (None) Edit

Processing Mode: Regular

Processing Priority: 0

☒ Update Dimensions

Storage Mode: MOLAP

Default Measure:


Estimated Rows: 0

Visible: True


OK Cancel Help

Click OK to close the *Cube sales_analysis definition* dialog and save the changes you've made.

- 14** You now need to check that the date dimension supports using OLAP functions. Expand the OLAP Dimension object type in the left hand pane to display the four OLAP Dimensions. Your object tree should contain the following OLAP Dimensions:



- 15** Click on the *odim_order_date* OLAP Dimension in the left pane to show the dimension attributes in the middle pane as below:



- 16 Right click on the *order_cal_month* attribute choose **Properties**. This displays the **Attribute Properties** dialog. Check the **Type** is set to *Months*. If it is not set to *Months*, change it. It needs to look like this:

Dimension odim_order_date Attribute Definition

Attribute Properties

Language Mapping

Dimension Name: odim_order_date

Internal Attribute Name: order_cal_month

Published Name: order cal month

Description: The calendar month representation. Format YYYYMM. Example: 200206.

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: Months

Usage: Regular

Key Column: dim_order_date

Name Column: dim_order_date

order_cal_month

order_cal_month

OK Cancel Help

Click OK.

Repeat the process for the *order_cal_year* attribute, checking the *Type* is set to the *Years* and changing it if it is not. It needs to look like this:

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

OK Cancel Help


- 17** To create the cube in Analysis Services, right Click the *sales_analysis* cube in the left hand pane and choose Create (Alter) Cube as follows:




This will open the WSL Cube executable that connects to Analysis Services to create the *sales_analysis* cube structure. You will see the successful completion in the results pane:




- 18** To process the cube *sales_analysis* in Analysis Services right click the *sales_analysis* cube in the left hand pane and choose **Update** as follows:



This will open the WSL Cube executable to process data into the cube structure. Once the cube is processed it can be viewed. Again, you will see the successful completion in the results pane:




- 19 To view the cube in Excel right click the *sales_analysis* cube in the left hand pane and choose Query cube Via Excel as follows:




This will open Excel if it is installed.

Note: If Excel displays the following dialog box, click **Enable**:




This message may not be displayed.

Excel then opens a connection to the cube for querying in a pivot table:




Note: In order for Excel to open the cube the OQY file extension in Windows explorer needs to be associated with Excel.

- 20** In addition, the cube structure can be opened in the Microsoft Business Intelligence Development Studio (BIDS). Open BIDS from the start menu in Windows. Then choose Open, Analysis Services Database as follows:




- 21** Choose the server and Tutorial5 cube database. A sample is shown:



Click OK.

This now displays the cube database definition from the Analysis Services server:



Close BIDS when finished.


Note for Oracle: In the latest version of Data Warehouse Designer, the default 0 key date for all databases has been set to 1st Jan 1753. Prior to 6.5.4.2 however, procedure update_dim_date uses sysdate-700000 which gives a date in year 0095. SSAS does not support dates before year 0100 using the oracle OLEDB driver. If using a previous version of Data Warehouse Designer, or upgrading from a previous version, it will therefore be necessary for you to replace 'sysdate - 700000' with 'TO_DATE('17530101','YYYYMMDD')' in procedure update_dim_date and to recreate the dim_date table.

You are now ready to proceed to the next section - **Adding a Measure Group** (see "5.3 Adding a Measure Group" on page 155)


5.3 Adding a Measure Group

A cube can contain multiple **Measure Groups**. In QAD Data Warehouse Designer, a Measure Group relates to a relational star schema. The cube ties Measure Groups together through **shared dimensions** that are conformed in the relational data warehouse. Measure Groups can be added to an existing cube by dragging additional fact tables into the cube.

- 1 Click on the OLAP Cube *sales_analysis* in the left hand pane to show the Measures associated with the cube. Your screen should look like this:



- Now you need to browse the *DataWarehouse* connection to see available fact tables. From the **Browse** menu, select **DataWarehouse** then **Change Connection**. The following dialog displayed:




The dialog box is titled "List Source Tables Connection" and features a close button in the top right corner. It contains the following fields and controls:

- Connection:** A dropdown menu with "DataWarehouse" selected.
- User ID:** An empty text input field.
- Password:** An empty 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:** "Refresh Current", "OK", and "Cancel" are located at the bottom of the dialog.


Choose *DataWarehouse* from the **Connection** drop down and click OK.

Note: For SQL Server the Data Warehouse schema must be **dbo**.

The following browse pane is displayed:




- 3 Drag the **fact_budget** table into the middle pane. The following dialog will appear:



Select Cancel.

- 4 QAD Data Warehouse Designer will confirm that it is OK to add a new OLAP Dimension for the date dimension. Change the Name of object to *odim_budget_date*.




Then click ADD.

Note: If you wish to include Attribute Relationships in Analysis Services for this dimension, click on the Include Attribute relationships check box.

- 5 The Adding cube measures dialog is displayed next. Remove all non-measure columns (that is, columns that cannot be aggregated) from the Measure list by highlighting them and clicking the Remove button. The columns to remove are:


- Product Code
- Customer Code

At this stage, your dialog should look like this:




- 6 Choose to create some date based calculated members using OLAP functions by selecting the **Month to date**, **Year to date**, and the **Use OLAP functions** check boxes. The calculated measures drop downs are now displayed. Select the required **date hierarchy** information. Specifically:
- for **Date dim**, select *Date Dimension*
 - for **Dim hie**, select *calendar*
 - for **Month level**, select *cal month*
 - for **Year level**, select *cal year*

After selecting the **Month level**, you will be asked to confirm changing the attribute type to Months for the the *cal month* attribute in Analysis Services as follows:




Click Yes.

Similarly, selecting the Year level, you will be asked to confirm changing the attribute type to *Years* for the the *cal year* attribute in Analysis Services as below:




Again, click Yes.

Your Manage cube measures dialog should now look like this:




Click OK to complete definition of the cube.


- Click on the OLAP Cube *sales_analysis* in the left hand pane to show the updated Measures associated with the cube. You should now see measures from both fact tables in the middle pane:




- 8 To apply the changes made in the QAD Data Warehouse Designer metadata to Analysis Services, right click the **sales_analysis** OLAP Cube in the left hand pane and choose **Create (Alter) Cube** as below:



This will open the WSL Cube executable that connects to Analysis Services to create the **sales_analysis** cube structure. You will see the successful completion in the results pane:




- 9 Next, reload the OLAP Cube in Analysis Services. Do this by right clicking on the *sales_analysis* cube in the left hand pane and choosing **Update** as follows:



This will open the WSL Cube executable to process data into the cube structure. Once the cube is processed it can be viewed. Again, you will see the successful completion in the results pane:


Results		
Object	Message	
sales	Processing Cube sales_analysis, database Tutorial5, server SSERVER	
sales	Processing partitions (One partition only/Full process) for cube measuregroup sales_detail in database Tutorial5 on server SSERVER	
sales	Processing (One partition only/Full process) partition fact_sales_detail of cube sales_analysis.	
sales	Processing partitions (One partition only/Full process) for cube measuregroup budget in database Tutorial5 on server SSERVER	
sales	Processing (One partition only/Full process) partition fact_budget of cube sales_analysis.	

- 10 To view the cube in Excel, right click the *sales_analysis* cube in the left hand pane and choose Query cube Via Excel like this:




This will open Excel if it is installed.

Note: If Excel displays the following dialog box, click **Enable**:



This message may not be displayed.

Excel then opens a connection to the cube for querying in a pivot table:



Note: the two measure groups are now displayed in the Field list.

You are now ready to proceed to the next section - **Cube Connections for Other Databases** (see "5.4 Cube Connections for Other Databases" on page 167)


5.4 Cube Connections for Other Databases

This tutorial has been completed using Microsoft SQL Server as the data warehouse database. If you wish to use Oracle or DB2, the data warehouse connections need to be set differently as follows.

This section shows how to configure the DataWarehouse connection for either an Oracle or a DB2 data warehouse database.

Oracle Configuration

- Select *MSDAORA* for the Provider Type
- Enter the oracle "TNS Server Name" for the **Server**, specifically the **Server** field in the ODBC set up screen shown:



The screenshot shows the "Microsoft ODBC for Oracle Setup" dialog box. It has a title bar with a close button. Inside, there are four text input fields and four buttons. The "Data Source Name" field contains "ORCL". The "Description" field is empty. The "User Name" field is empty. The "Server" field contains "ORCL". The buttons are "OK", "Cancel", "Help", and "Options >>".

Field	Value
Data Source Name:	ORCL
Description:	
User Name:	
Server:	ORCL

- Enter the data warehouse "schema" for the Database field


Here is an Oracle example:

The image shows a "DataWarehouse Definition" dialog box with a "Connection Properties" tab selected. The dialog is used to configure a data warehouse connection. The "Data Warehouse" checkbox is checked. The "Connection Name" is "DataWarehouse". The "Connection Type" is "Database". The "ODBC Source" is "ORCL". The "Provider Name" is empty. The "Database ID (SID)" is "ORCL". The "Database Link ID" is empty. The "Extract User ID" is "dssdemo" and the "Password" is "xxxx". The "Administrator User ID" and "Password" are empty. The "Default Load Type" is "Database link load". The "SSIS Connection String" is empty. The "Data Type Mapping Set" is "(Default)". Below this is a section labeled "As a Cube Data Source" with a line for configuration. The "Provider Type" is "MSDAORA (OLE DB for Oracle)". The "Server" is "ORCL" and the "Database" is "dssdemo". At the bottom are "OK", "Cancel", and "Help" buttons. A "Reset" button is next to the "SSIS Connection String" field.

☒ Data Warehouse	
Connection Name:	DataWarehouse
Connection Type:	Database
ODBC Source:	ORCL
Provider Name:	
Database ID (SID):	ORCL
Database Link ID:	
Extract User ID:	dssdemo
Password:	xxxx
Administrator User ID:	
Password:	
Default Load Type:	Database link load
SSIS Connection String:	
Data Type Mapping Set:	(Default)
As a Cube Data Source _____	
Provider Type:	MSDAORA (OLE DB for Oracle)
Server:	ORCL
Database:	dssdemo

DB2 Configuration

- Select *IBMDADB2* for the Provider Type
- The **Server** field is left empty for DB2
- Enter the database server alias for the **Database** field, specifically the alias field for the IBM DB2 connection in the following DB2 Configuration Assistant screen:



Here is an IBM DB2 example:

The image shows a 'DataWarehouse Definition' dialog box with a 'Connection Properties' tab. The 'Data Warehouse' checkbox is checked. The 'Connection Name' is 'DataWarehouse', 'Connection Type' is 'Database', and 'ODBC Source' is 'DBTW0'. The 'Work Directory', 'Database ID (SID)', and 'Database Link ID' fields are empty. The 'Extract User ID' and 'Administrator User ID' fields are empty, each with a corresponding 'Password' field. The 'Default Load Type' is 'Database link load'. The 'Data Type Mapping Set' is '(Default)' with a dropdown arrow. Below it is a line 'As a Cube Data Source'. The 'Provider Type' is 'IBMDADB2 (OLE DB for IBM DB2)', 'Server' is empty, and 'Database' is 'QAD'. At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

☒ Data Warehouse	
Connection Name:	DataWarehouse
Connection Type:	Database
ODBC Source:	DBTW0
Work Directory:	
Database ID (SID):	
Database Link ID:	
Extract User ID:	
Administrator User ID:	
Default Load Type:	Database link load
Data Type Mapping Set:	(Default)
As a Cube Data Source	
Provider Type:	IBMDADB2 (OLE DB for IBM DB2)
Server:	
Database:	QAD

CHAPTER 6

Tutorial 6 - KPI fact table

Before you start on this chapter you should have:

- Completed ***Tutorial 2 - Rollup fact tables, ASCII file loads, aggregates*** (on page 63)

This chapter will cover the definition of a Key Performance Indicator (KPI) fact table. The fact tables built as part of the first two tutorials are used in this process.

In This Chapter

6.1 Purpose and Roadmap.....	172
6.2 The KPI Dimension	173
6.3 Defining the KPI Fact Table	174
6.4 Creating the KPI Fact Table and Procedure.....	181
6.5 The KPI Parameters	184
6.6 Defining a KPI	185
6.7 Defining the Activity KPI.....	189

6.1 Purpose and Roadmap

Purpose

KPI fact tables are designed to assist in the answering of difficult business questions. The star schema design does not support an easy means of answering questions that wish to compare current month values with the same period last year, or current month with year to date values. In many cases it is possible to build such measures into monthly rollup fact tables, but in some instances even that design does not provide an adequate solution. KPI fact tables are a very specific type of fact table and are narrowly focused to handle these complex comparative queries.

In this tutorial you will learn how to:

- (i) set-up a KPI dimension
- (ii) create a KPI fact table
- (iii) define individual KPIs

This tutorial relies upon the tables created in the previous tutorials.

Tutorial Environment

This tutorial has been completed using IBM DB2. All of the features illustrated in this tutorial are available in SQL Server, Oracle and DB2 (unless otherwise stated). Any differences in usage of QAD Data Warehouse Designer between these databases are highlighted.

Tutorial Roadmap

Step in Tutorial	Section
Create a dimension to support the KPI fact table	The KPI dimension
Define the columns of the KPI fact table including the measures and the specific dimensions required.	Defining the KPI fact table
Create the KPI fact table and build a procedure used to update the fact table.	Creating the KPI fact table
Create and initialize the parameters used to establish the update window for the fact table.	The KPI parameters
Define a simple KPI	Defining a KPI
Define a KPI that does year to date calculations for each month.	Defining the Activity KPI

The tutorial starts with '*The KPI Dimension*' (see "**6.2 The KPI Dimension**" on page 173).


6.2 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 KPI dimension can be created from the KPI table distributed with the other tutorial tables:

- 1 Double click on the **Load Table** object group in the left hand pane to create a drop target in the middle pane.
- 2 Browse to the tutorial tables and drag the table 'KPI' into the middle pane. Accept the default name of load_kpi. Click OK on the Properties screen. **Create and load the table.**
- 3 Double click on the **Dimension** object group in the left hand pane to create a dimension drop target in the middle pane.
- 4 Browse to the data warehouse and drag the newly created load_kpi table into the middle pane. Accept the default name of dim_kpi and proceed to create and populate a **Normal** dimension as per the examples in the first tutorial. The business key is kpi_identifier.
- 5 Right click on dim_kpi and select **Execute Update Procedure.**

Your QAD Data Warehouse Designer screen should look like this when you have finished:




You are now ready to proceed to the next section - *Defining the KPI Fact Table* (see "6.3 Defining the KPI Fact Table" on page 174).

6.3 Defining the KPI Fact Table

A KPI fact table is normally created manually. We will, however, use the drag and drop functionality to acquire the dimension key settings. Proceed as follows:

- 1 Double click on the **Fact Table object group** in the left hand pane. This will display a list of fact tables in the middle pane and create a fact table drop target.
- 2 Browse to the data warehouse and double click on **dim_date** to expand the tree and display the columns.
- 3 Drag the column **dim_date_key** from the right hand (browse) pane into the middle pane. The "Add a new object" dialog will appear defaulting to a fact table called fact_date. Change the object type to Kpi fact table and the table name to fact_kpi.



- 4 Click ADD to add the new object and OK on the properties screen.
- 5 Use the right mouse menu on the column **dim_date_key** in the middle pane and select the **Properties** of this column. Under the column type select **date dimension key** and enter 0 for the default value. Click OK. This identifies this column as the date dimension column to be used during the definition of the individual KPIs.

Fact Table.dim_date_key Column Definition

Column Properties

Table Name: Fact Table

Column Name: dim_date_key

Business Display Name: dim time key

Data Type: integer

Format: ###0

Column Order: 10

Nulls: ☐

Numeric: ☒

Additive: ☒

Attribute: ☐

Transformation

Source Table: dim_date

Source Column: dim_date_key

Business Definition: Generated key for the date dimension. Nominal format is YYYYMMDD, but DD equal to 0 and greater than 31 permitted for special date handling.

Column Type (KPIs only): Date dimension key

Default Value (KPIs only): 0

End User Layer Display: ☐

Dimension Join: ☒ Manual Dimension Join: ☐

Key Type (0,A,B,C,.): 2

Artificial Key: ☐

Primary Business Key: ☐

Note: When the DimensionJoin flag is set then the contents of Source column denotes the foreign key. Business definition denotes the join.

OK Cancel Help

- 6 Drag the column `dim_kpi_key` from the `dim_kpi` table in the right hand pane into the column list in the middle pane.
- 7 Use the right mouse menu on the column `dim_kpi_key` in the middle pane and select the Properties of this column. Under the column type select Kpi dimension key and enter 0 for the default value.

Fact Table.dim_kpi_key Column Definition

Column Properties

Transformation

Table Name: Fact Table

Column Name: dim_kpi_key

Business Display Name: dim kpi key

Data Type: integer

Format: ###0

Column Order: 20

Nulls: ☒

Numeric: ☒

Additive: ☐

Attribute: ☐

Column Type (KPIs only): Kpi dimension key

Default Value (KPIs only): 0

End User Layer Display: ☐

Dimension Join: ☒ Manual Dimension Join: ☐

Key Type (O,A,B,C,...): 2

Artificial Key: ☐

Primary Business Key: ☐

Note: When the Dimension Join flag is set then the contents of Source column denotes the foreign key. Business definition denotes the join.

Source Table: dim_kpi

Source Column: dim_kpi_key

Business Definition: Generated artificial key

Source strategy:

OK Cancel Help

- 8 Drag the column **dim_customer_key** from the **dim_customer** table in the right hand pane into the column list in the middle pane.
- 9 Use the right mouse menu on the column **dim_customer_key** in the middle pane and select the **Properties** of this column. Under the column type select **Dimension key** and enter 0 for the default value.

Fact Table.dim_customer_key Column Definition

Column Properties
Transformation

Table Name: Fact Table

Column Name: dim_customer_key

Business Display Name: dim customer key

Data Type: integer

Format: ###0

Column Order: 30

Nulls: ☒

Numeric: ☒

Additive: ☐

Attribute: ☐

Column Type (KPIs only): Dimension key

Default Value (KPIs only): 0

End User Layer Display: ☐

Dimension Join: ☒ Manual Dimension Join: ☐

Key Type (O,A,B,C.): 2

Artificial Key: ☐

Primary Business Key: ☐

Note: When the DimensionJoin flag is set then the contents of Source column denotes the foreign key. Business definition denotes the join.

Source Table: dim_customer

Source Column: dim_customer_key

Business Definition:
Source strategy:
Generated artificial key

<- Update
Update ->

OK Cancel Help


At this point we have defined the new KPI fact table and have added three columns to it, namely `dim_date_key`, `dim_kpi_key` and `dim_customer_key`. For the purposes of this tutorial these will be the only dimensions we will utilize.

We will now define the measures of the KPI fact table. KPI fact tables contain four types of measure. These measure types are:

- month to date
- year to date
- same month last year
- previous year to date.


These measures will have different meanings depending on the KPI being evaluated. They are by necessity, therefore very generic. Typically just a count and a dollar value. Note, that all measures must allow nulls. Although the resultant record will not contain nulls, the KPI update procedure makes multiple passes in order to update the fact table and needs the ability to insert nulls during the process. Proceed as follows to create quantity and sales measures for each of the types defined above.

- 1 Click on the `fact_kpi` table name to refresh the column display in the middle pane. These columns are listed:



Column Name	Display Name	Data Type	Source Table
 dim_date_key	dim time key	integer	dim_date
 dim_kpi_key	dim kpi key	integer	dim_kpi
 dim_customer_key	dim customer ...	integer	dim_customer

- 2 Use the right mouse menu on `fact_kpi` and select Add Column.



- 3 Enter `mtd_quantity` for the column and business names. Enter `integer` for the data type. Select `Month to date measure` for the column type and set a default value of 0. The Nulls check box must be set. Update the column properties:

Column Definition

Column Properties

Transformation

Table Name:

Column Name:

Business Display Name:

Data Type:

Format:

Column Order:

Nulls: ☒

Numeric: ☐

Additive: ☐

Attribute: ☐

Column Type (KPIs only):

Default Value (KPIs only):

End User Layer Display: ☒

Dimension Join: ☐ Manual Dimension Join: ☐

Key Type (0,A,B,C.):

Artificial Key: ☐

Primary Business Key: ☐

Note: When the Dimension Join flag is set then the contents of Source column denotes the foreign key. Business definition denotes the join.

Source Table:

Source Column:


Business Definition:
Source strategy:

OK Cancel Help

- 4 Repeat step 2 and 3 using the column name `mtd_sales` with a data type of `number(13,2)` for Oracle, or `numeric(13,2)` for SQL Server and DB2.

The screenshot shows the 'Column Definition' dialog box for the table 'fact_kpi'. The 'Column Properties' tab is active. The 'Transformation' section is empty. The 'Column Name' is 'mtd_sales' and the 'Business Display Name' is also 'mtd_sales'. The 'Data Type' is 'numeric(13,2)'. The 'Format' is empty. The 'Column Order' is '40'. The 'Nulls' checkbox is checked. The 'Numeric', 'Additive', and 'Attribute' checkboxes are unchecked. The 'Column Type (KPIs only)' is 'Month to date measure'. The 'Default Value (KPIs only)' is '0'. The 'End User Layer Display' checkbox is checked. The 'Dimension Join' checkbox is unchecked, and the 'Manual Dimension Join' checkbox is also unchecked. The 'Key Type (0,A,B,C,...)' is empty. The 'Artificial Key' checkbox is unchecked. The 'Primary Business Key' checkbox is unchecked. A note states: 'Note: When the Dimension Join flag is set then the contents of Source column denotes the foreign key. Business definition denotes the join.' The 'Source Table' and 'Source Column' fields are empty. The 'Business Definition' and 'Source strategy' fields are empty. The 'OK', 'Cancel', and 'Help' buttons are at the bottom right.

- 5 Position on the `mtd_quantity` column and use the right mouse menu option **Copy Column** to copy the column.



Change the properties of the copied column to be **ytd_quantity** and the column type to Year to date measure.

- 6 Repeat step 5 for mtd_sales to create **ytd_sales**.
- 7 Repeat steps 5 and 6 to create **prev_month_quantity** and **prev_month_sales** using a column type of Last year same month measure
- 8 Repeat steps 5 and 6 to create **prev_ytd_quantity** and **prev_ytd_sales** using a column type of Prev year to date measure.

The resultant set of columns should be named as follows with the column type as shown (SQL Server).

Note: Oracle should have number(13,2) instead of numeric(13,2).

Column Name	Display Name	Data Type	Source Table	Source Column	Nulls	Column Type
dim_kpi_key	dim kpi key	integer	dim_kpi	dim_kpi_key	Y	Kpi dimension key
dim_customer_key	dim custom...	integer	dim_custo...	dim_custom...	Y	Dimension key
dim_date_key	dim time key	integer	dim_date	dim_date_key	N	Date dimension key
mtd_quantity	mtd_quantity	integer			Y	Month to date me...
mtd_sales	mtd_sales	numeric(13,2)			Y	Month to date me...
ytd_quantity	ytd_quantity	integer			Y	Year to date meas...
ytd_sales	ytd_sales	numeric(13,2)			Y	Year to date meas...
prev_month_quantity	prev_month...	integer			Y	Last year same mo...
prev_month_sales	prev_month...	numeric(13,2)			Y	Last year same mo...
prev_ytd_quantity	prev_ytd_qu...	integer			Y	Previous year to d...
prev_ytd_sales	prev_ytd_sales	numeric(13,2)			Y	Previous year to d...

Note: The column order is important. For example if we have mtd quantity and sales in that order then the prev month, prev ytd and ytd figures also need quantity first followed by sales. If this order is not adhered to then the generated procedure will load the wrong values into the wrong columns.

The KPI fact table has now been defined and can be created.


You are now ready to proceed to the next section - '**Creating the KPI Fact Table and Procedure**' (see "6.4 Creating the KPI Fact Table and Procedure" on page 181)'.

6.4 Creating the KPI Fact Table and Procedure

Once defined the KPI fact table can be created and the update procedure built.

- 1 Create the KPI fact table by selecting the **Create/ReCreate** menu option when positioned over the fact_kpi table name.
- 2 Alter the properties of the table and select **(Build Procedure...)** from the Update procedure drop list. Enter OK to update the properties.

- 3 A dialog box will prompt for a period start parameter. This parameter along with the subsequent period end parameter define what period(s) will be processed when the KPI fact table is updated. These parameters must be defined, so accept the default FACT_KPI_START and FACT_KPI_END parameter names. See **The KPI Parameters** (see "6.5 The KPI Parameters" on page 184) for a description of these parameters.
- 4 A date dimension and month rollup column are prompted for. Select `dim_date` and `cal_month`:



Define Rollup date dimension and column


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)

Date Dimension: `dim_date`

Rollup Column: `cal_month`

OK Cancel

- 5 Select `cal_year` for the rollup year column:



Define Rollup date dimension and column


Select the column used to roll up the year (e.g. Cal_Year)

Date Dimension: `dim_date`

Rollup year: `cal_year`


OK Cancel

and then `calendar_date` for a date column:



The dialog box is titled "Define Rollup date dimension and column". 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_date" selected. The second is labeled "Date column" and has "calendar_date" selected. At the bottom right, there are "OK" and "Cancel" buttons.

- 6 Enter the KPI dimension and identifier as `dim_kpi` and `kpi_identifier`:



The dialog box is titled "Select the KPI dimension and identifier". 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_kpi" selected. The second is labeled "KPI Identifier" and has "kpi_identifier" selected. At the bottom right, there are "OK" and "Cancel" buttons.

- 7 The procedure will now be built and compiled.

You are now ready to proceed to the next section - *The KPI Parameters* (see "6.5 The KPI Parameters" on page 184).


6.5 The KPI Parameters

During the procedure creation phase of the KPI fact table we are prompted for parameters used to define the period start and period end. These parameters are read by the update procedure whenever a KPI is updated, and are used to ascertain what time periods are being processed.

If the parameters do not exist then the default of the current and previous months is used. It would however be normal practice to create and maintain these parameters. A procedure called 'daily_date_roll' is shipped with the metadata. This procedure updates a number of parameters to reflect the current date. It is designed to be run every night. Normally we would alter this procedure and include the setting of our KPI period parameters to perhaps the current and previous months.

For this tutorial we will create these parameters and set dates to cover the full range of tutorial data.


- 1 Select the **Tools/Parameters** menu option. This displays the current parameters in the middle pane.
- 2 Right mouse in the middle pane and choose the **'Add parameter'** menu option:



- 3 Enter the values from the following table for the period start parameter and then repeat step 2 and enter the values for the period end parameter.

Parameter	Value	Comments
FACT_KPI_START	200006	Period start value. We have chosen cal_month as the rollup column. This column has a format of YYYYMM so we enter 200006 for June/2000
FACT_KPI_END	200204	Period end value. April/2002 for tutorial

Sample parameter:



The screenshot shows a 'Parameter Maintenance' dialog box with the following fields:

- Parameter:** FACT_KPI_START
- Value:** 200006
- Comments:** Period start value. We have chosen cal_month as the rollup column. This column has a format of YYYYMM so we en

Buttons: OK, Cancel

Once the parameters have been created we can define the KPI.

You are now ready to proceed to the next section - **Defining a KPI** (see "6.6 Defining a KPI" on page 185).


6.6 Defining a KPI

We will now define two KPIs. The dimension we have created dim_kpi has two KPIs. These are firstly an Activity KPI, which shows the number of active customers and the value of sales during a month. Secondly a Sales KPI, which shows the number of sales and the value of the sales by customer by month. The process for creating the two KPIs is similar, but there is a difference in complexity. In the Sales KPI we can simply add up all the monthly sales counts and values to achieve year to date values. However, we cannot add the number of active customers in each month to acquire the year to date activity count, as a customer may have been active in more than one month, but we only want to recognize that customer once.

We will do the simpler Sales KPI first followed by the Activity KPI.

Sales KPI

- 1 Use the right mouse menu when positioned over the fact_kpi table and select Add Kpi:



- 2 Populate the dialog box as per the following example. The Kpi Identifier must match the value in the KPI dimension. In this case it must be 'Sales'. The Kpi Name and Description are not matched to the dimension, and can be used here for internal definition of the KPI. The **active** check box should be set, so that this particular KPI will be processed when we update the KPI fact table. For the 'Month to date update columns' we select the `mtd_quantity`, `mtd_sales`, and `dim_customer_key`. The 'Year to date update columns' are left blank to allow the update procedure to automatically handle these values.

KPI Definition

Properties

MTD Statement

YTD Statement

Procedural Block

KPI Identifier: Sales

KPI Name: Customer sales (quantity and value)

Active: ☒

Description: Provides a count and dollar value of sales in a particular month. The customer dimension is used to allow analysis by customer if required.

Month to Date Update Columns: mtd_quantity,mtd_sales,dim_customer_key

Select Columns

Year to Date Update Columns:

Select YTD Cols

OK Cancel Help

- 3 Once the dialog above is completed, click OK to add the KPI. Selecting the Display Kpi's right mouse menu option when positioned on the fact_kpi table will now display our newly defined KPI:


KPI	Name	Active	Method	Description
Sales	Customer sales (...	Enabled	Statement	Provides a count an

There are three additional tabs on the dialog box as shown in the screen shot in point 2. These tabs allow the definition of a **month to date statement**, a **year to date statement** and a **procedural block**. In this example we will be defining a month to date statement. A procedural block can be used when the KPI is too complex to be performed as a single SQL statement and procedural language is required.

We will now proceed to create the SQL statement required to perform the month to date update.


- 1 Use the right mouse menu when positioned on the **Sales** KPI and select **Modify Kpi**. Then choose the **Mtd Statement** tab. You will note that there are two comment lines at the top of the edit window. These comments show the order in which information is to be added, and hence the order in which the Mtd select statement must return its information.
- 2 Enter the statement as shown below. Note, that no trailing semi-colon is required. When done click OK to complete the definition of the KPI.

```
-- Insert into fact_kpi(
-- mtd_quantity,mtd_sales,dim_customer_key)
Select sum(quantity),sum(sales_value),dim_customer_key
From fact_sales_analysis
Where dim_ship_date_key >= $KPI_START_YYYYMMDD$
And dim_ship_date_key <= $KPI_END_YYYYMMDD$
group by dim_customer_key
```




For DB2, enter the statement as shown below.


```
-- Insert into fact_kpi(
-- mtd_quantity,mtd_sales,dim_customer_key)
Select sum(quantity),sum(sales_value),dim_customer_key
From [TABLEOWNER].[fact_sales_analysis]
Where dim_ship_date_key >= $KPI_START_YYYYMMDD$
And dim_ship_date_key <= $KPI_END_YYYYMMDD$
group by dim_customer_key
```



- The KPI has now been defined and we can update the KPI fact table in the normal manner, or by selecting **Update (run)** when positioned over the KPI in the middle pane. This will execute the KPI table update procedure in foreground. If all goes well the procedure will report the processing of 23 months, and will add a number of records.



- 4 We would normally examine the fact table through an end user tool, but the following SQL statement can be used to see the results. The screen shot below shows the normal output from this select statement.



You are now ready to proceed to the next section - '**Defining the Activity KPI**' (see "**6.7 Defining the Activity KPI**" on page 189)

6.7 Defining the Activity KPI

The Activity KPI is similar to the Sales KPI, except that it will require the definition of a year to date statement. We will create the KPI by copying the first KPI to show the use of this copy capability.

Activity KPI

- 1 Use the right mouse menu when positioned over the Sales KPI and select **Insert copy**.
- 2 Position over the second (copy of) Sales KPI and select **Properties** from the right mouse menu.
- 3 Change the Kpi Identifier to **Activity**, and the name and definition as per the example screen below.
- 4 Click the **Select Columns** button and remove the **dim_customer_key** entry.
- 5 Click the **Select YTD Cols** button and add **ytd_quantity** and **ytd_sales**.

The completed properties screen should look as follows:

KPI Definition

Properties

MTD Statement

YTD Statement

Procedural Block

KPI Identifier: Activity

KPI Name: Customer activity (counts and value)

Active: ☒

Description: Provides a count of active customer in a particular month, together with total sales.

Month to Date Update Columns: mtd_quantity, mtd_sales

Select Columns

Year to Date Update Columns: ytd_quantity, ytd_sales

Select YTD Cols

OK Cancel Help


Once the dialog above is completed click OK to modify the KPI.

We will now proceed to modify the SQL statement required to perform the month to date update.

- 1 Use the right mouse menu when positioned on the **Activity** KPI and select **Modify Kpi**. Then choose the **Mtd Statement** tab.
- 2 Replace the statement with the statement as shown below. Note, that no trailing semi-colon is required.


For **SQL Server** enter the statement as:

```
-- Insert into fact_kpi(
-- mtd_quantity, mtd_sales)
Select count (distinct dim_customer_key), coalesce(sum(sales_value),0)
From fact_sales_analysis
Where dim_ship_date_key >= $KPI_START_YYYYMMDD$
And dim_ship_date_key <= $KPI_END_YYYYMMDD$
```



For DB2 enter the statement as:

```
-- Insert into fact_kpi(
-- mtd_quantity,mtd_sales)
Select count (distinct dim_customer_key), coalesce(sum(sales_value),0)
From [TABLEOWNER].[fact_sales_analysis]
Where dim_ship_date_key >= $KPI_START_YYYYMMDD$
And dim_ship_date_key <= $KPI_END_YYYYMMDD$
```




We now need to define the year to date statement, to ensure we do not count the same customer twice. Proceed as follows:

- 1 Choose the Ytd Statement tab.
- 2 Insert the statement as shown below. Note, that no trailing semi-colon is required. When done click OK to complete the definition of the KPI.


For SQL Server enter the statement as:

```
-- Insert into fact_kpi(
-- ytd_quantity,ytd_sales)
Select count (distinct dim_customer_key), coalesce(sum(sales_value),0)
From fact_sales_analysis
Where dim_ship_date_key >= $KPI_START_YYYYMMDD$
And dim_ship_date_key <= $KPI_END_YYYYMMDD$
```



For DB2 enter the statement as:

```
-- Insert into fact_kpi(
-- ytd_quantity,ytd_sales)
Select count (distinct dim_customer_key), coalesce(sum(sales_value),0)
From [TABLEOWNER].[fact_sales_analysis]
Where dim_ship_date_key >= $KPI_START_YYYYMMDD$
And dim_ship_date_key <= $KPI_END_YYYYMMDD$
```



The KPI has now been defined and can be updated as per the Sales KPI. An end user tool is required to adequately examine the results.