Jet Data Manager 2016

User Guide

Jet Enterprise

Welcome

This documentation provides descriptions of the concepts and features of the Jet Data Manager and how to use them.

With the Jet Data Manager you can transform and validate data from different data sources, consolidate them in a data warehouse, and build Online Analytical Processing (OLAP) cubes. When you execute or deploy a project, the Jet Data Manager automatically generates the underlying SQL code of your solution.

With the Jet Data Manager you can access data from a variety of sources including ERP systems, CRM applications, SQL databases, spreadsheets, and plain text files. The data is stored in a data warehouse in Microsoft SQL Server 2005, 2008 or 2012 and can be viewed in your preferred front-end application.

Prerequisites

This document outlines the features and functionality of the Jet Data Manager and how to use this tool to construct or modify your business intelligence project. This document does not cover detailed concepts of data modeling or the tables and fields of any particular relational database. Users of the Jet Data Manager are expected to already be familiar with their source database as well as how they would like to see their data modeled.

Trademarks

Microsoft®, Windows® and other names of Microsoft products are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

All other product names mentioned in this documentation may be trademarks or registered trademarks of their respective companies.

TABLE OF CONTENTS

GETTING STARTED WITH JET DATA MANAGER

The purpose of Jet Data Manager is to enable you to create and maintain a complete data warehouse solution with as little effort as possible.To achieve this, Jet Data Manager employs Data Warehouse Automation techniques to automate the tedious parts of the work.

SQL code for the ETL (Extract, transform, load) process, MDX code for OLAP cubes, indexes etc. are generated automatically. Most tasks can be accomplished using drag-and-drop in the graphical user interface, and the amount of code you need to write is minimized. However, if and when you need to customize the code, you can do it in your favorite development environment.

This user guide outlines the features and functionality of Jet Data Manager. Please note that this document does not cover detailed concepts of data modeling or the tables and fields of any particular relational database. If you need help with this, please explore our training options.

This section provides an overview of the workflow, the components and the user interface you will be working with when you build a data warehouse in Jet Data Manager.

JET DATA MANAGER INSTALLATION

Installing Jet Data Manager is a three-step process. First you need to make sure that the necessary prerequisites are met, second you download and install the software and third, you need to set up the scheduler service if you plan to use it.

PREREQUISITES FOR INSTALLATION

On deploying machine:

- .NET Frame Work 4.0 Extended
- Core XML Services (MSXML) 6.0
- Microsoft SQL Server Native Client
- Microsoft SQL Server System CLR Types
- Microsoft SQL Server Management Objects
- Microsoft SQL Server Integration Services (SSIS)

Components on destination server:

- SQL Server 2008 or later
- Edition should be Standard or above with the following components installed:
    - DB Engine
    - Integration Services (SSIS)
    - Analysis Services (SSAS) (when using cubes)

Often, the deploying machine and destination server will be one and the same.

DOWNLOADING AND INSTALLING JET DATA MANAGER

When you have made sure that your machine has the required prerequisites, you can move on to downloading and installing Jet Data Manager.

1. Download the latest version of Jet Data Manager from http://www.jetreports.com/download/. The software comes in 32 and 64 bit flavors - make sure to download the bit-version that matches your setup. You can choose between a client and a server edition. The client edition comes without the scheduler and server services (used for i.a. multiple environments) and is most suited for use in development. If in doubt, download the server edition.
2. When the appropriate file has been downloaded, unzip the installation package and run the **Setup.exe** file. Select **Typical** from the list of installation types when prompted.
3. Once the installation has completed there will be an icon on the desktop for the new version of Jet Data Manager. Double click this icon to launch the new version. The first time the new version is opened it will prompt you for an Activation Code. Select

Internet – I have a license code.
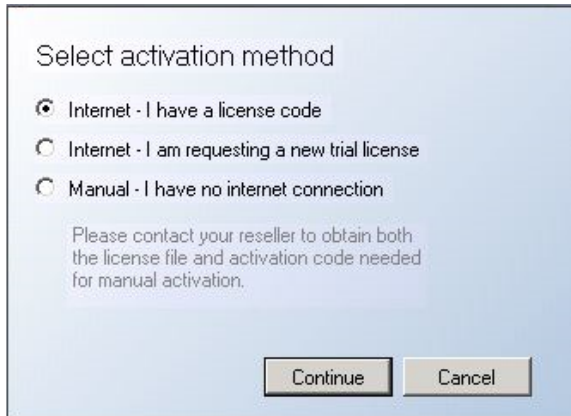
4. On the following screen paste or type in your organization's Activation Code and click **Send**. This will synchronize with our licensing server and enable all your licensed features of Jet Data Manager.

5. After activating the software, you will receive a warning that the project repository settings are invalid.
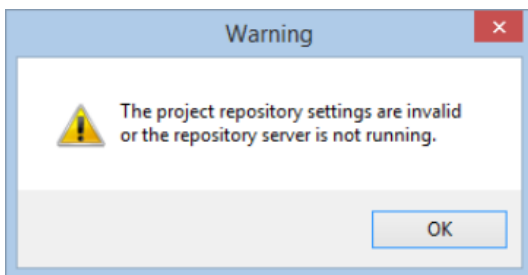
Click **OK** to dismiss the warning. The **General Settings** window opens on the **Project Repository** tab.
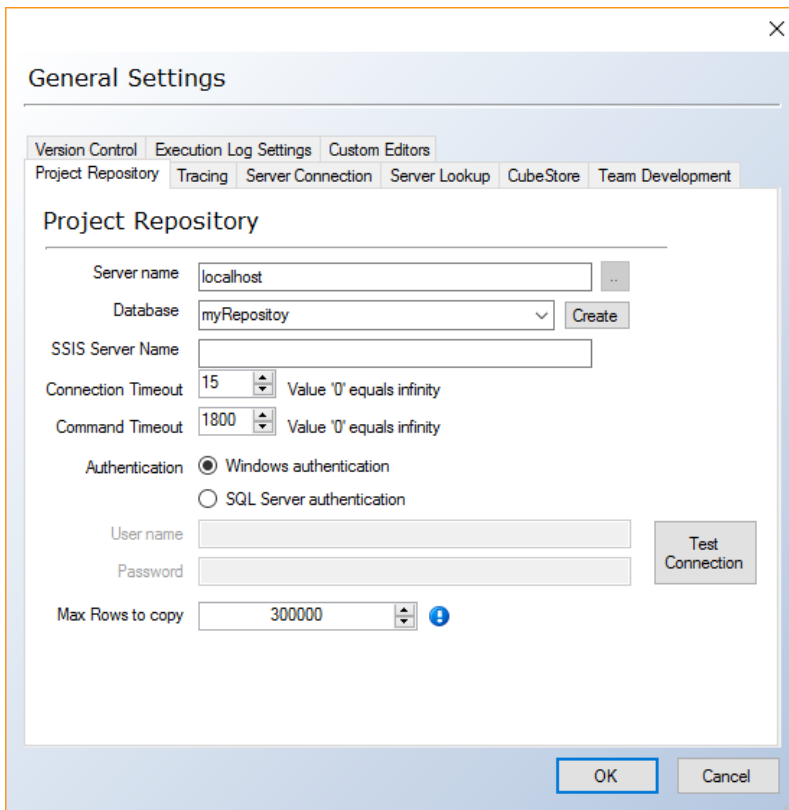
Enter the server name and database you wish to use for your repository. If you are upgrading from an earlier version, enter the same repository settings you used for that earlier version.

6. If there already is a repository in the database you entered, for instance if you are upgrading, Jet Data Manager will recognize the repository and prompt that it needs to be upgraded to be compatible with the new version. Click **Upgrade**.
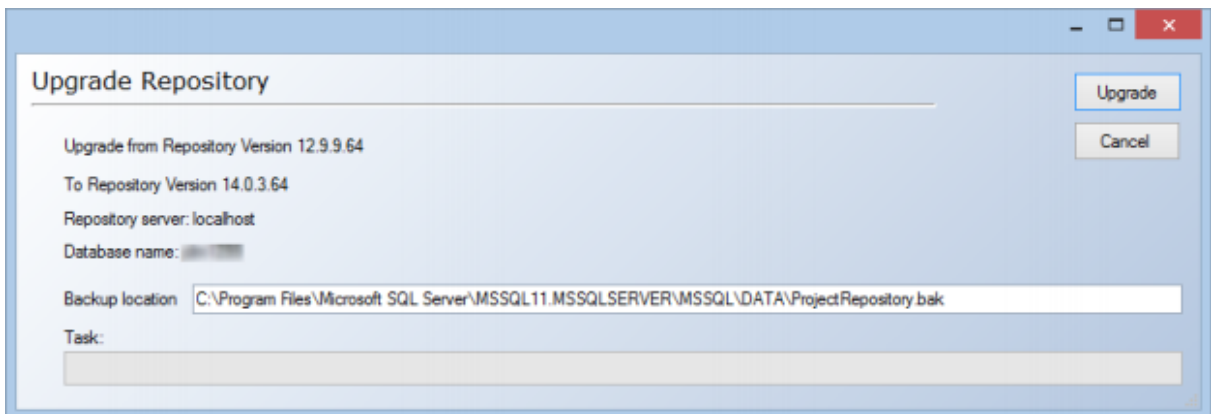
Jet Data Manager will then ask you what upgrade method you wish to use. Select **Web service**.

Select upgrade method

Local files   Web service

When using Local files you might
want to check for script changes.
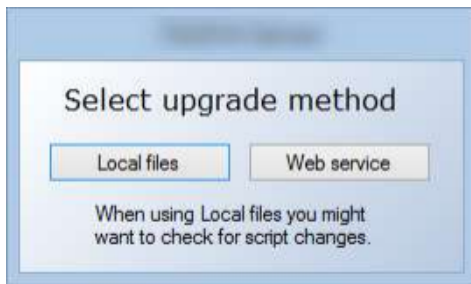
The upgrade scripts will be downloaded and processed.

7. If you have downloaded the server edition of the software, set up the scheduler service by following the instructions below. If you are upgrading from an earlier version, you can uninstall the old version from Programs and Features in the Windows control panel when you have set up the scheduler.

SETTING UP THE SCHEDULER

The execution of scheduled packages is handled by a service on the machine named "Jet Data Manager Scheduler". The service is automatically installed along with the server edition of Jet Data Manager and checks every two minutes to see if any packages are scheduled to be run.

CONFIGURING A USER ACCOUNT FOR USE WITH THE SCHEDULER SERVICE

The service needs to be run by a user with sufficient user rights and configured to use the correct repository database. This can either be the same account that you use to develop your projects, or a different service account. You can skip the steps below if you are using the account you use for development or if you are upgrading Jet Data Manager and already have a user account set up with the correct repository settings.

1. Make sure that the account has sufficient permissions.
   - On Analysis Services, **Administrator** rights (if you use OLAP Cubes).
   - On the **msdb** database, membership of the **db_ssisadmin** role.
   - On repository, staging and data warehouse databases, membership of the **db_owner** role.
2. Log in with the user account, open Jet Data Manager and set up the repository database that contains the project you wish to schedule. This creates the necessary configuration files for the scheduler to run correctly.

CONFIGURING THE SCHEDULER SERVICE

When you have the user account ready, it is time to configure the scheduler service to use the user account. Jet Data Manager includes a tool, Windows Services Management, to view, add and delete the Windows services used by Jet Data Manager.

1. In the ribbon, click the **Tools** tab and in the **Environments** group click **Windows Services Setup**. This will show you the services of the versions of Jet Data Manager you currently have installed.
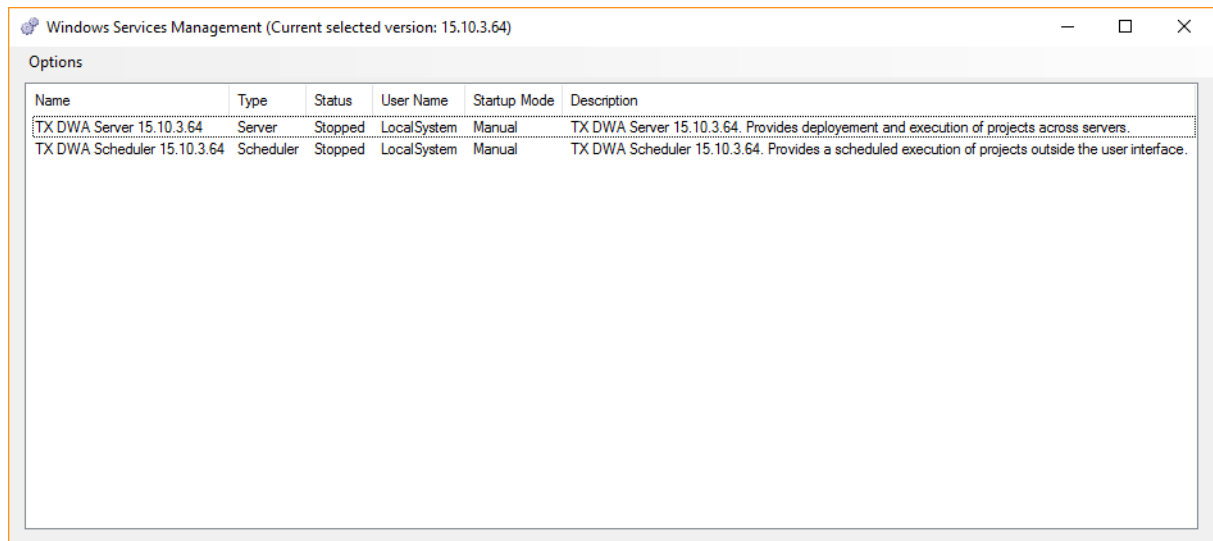
2. Locate the Jet Data Manager **Scheduler <version>** item for the version you have just installed, right-click it, click **Start Mode** and click **Automatic (delayed)**. The scheduler service will now start up together with Windows.
3. Right click the service in the list again and click **Change Username and Password**. In the window that appear, enter the username and password for the user account you configured earlier and click **Update**.
4. Right click the service againt and click **Start Service**.

INSTALLING THE REMOTE EXECUTION OF SSIS PACKAGES FEATURE

The Remote Execution of SSIS Packages feature improves performance when Jet Data Manager is installed seperately from the SQL Server that houses the data warehouse.

In that setup, all traffic from source systems to the staging databases and data warehouses has to pass throught the Jet Data Manager machine where Jet Data Manager opens and executes the SSIS packages it has generated. Even on a high-speed LAN, this is considerably slower than having Jet Data Manager and the data warehouse on the same server.

As the name suggests Remote Execution of SSIS Packages lets you execute SSIS packages on a remote server, typically the server that houses the actual data warehouse, saving the data an extra trip thorught the network cables.

To use the feature, you need to install a service on the remote server and configure Jet Data Manager to use the remote server for executing SSIS packages.

INSTALLING THE REMOTE SSIS EXECUTION SERVICE

To install the Remote SSIS Execution Service on the remote server, follow the steps below.

1. Download the installation package for the service that matches the version of Microsoft SQL Server (2008-2014) on the remote machine from our website: https://support.timextender.com/hc/en-us/articles/210439563
2. Unzip the content of the installation package to a temporary folder and double-click Setup.exe. Follow the installation instructions to install the service.
3. Next, you need to set up the service. Click/Right-click **Start**, click **Run**, type **Services.msc** and click **OK**. Locate the **Remote SSIS Execution** service in the list. Right click the service and click **Properties**.
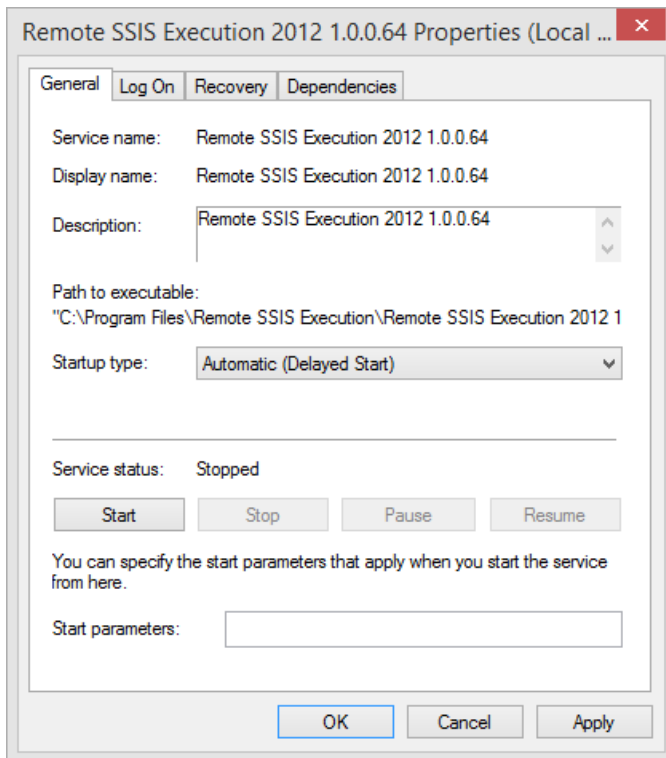
4. On the **General** tab, in the **Startup type** list, click **Automatic (Delayed Start)** to ensure that the service is started after SQL Server.
5. On the **Log On** tab, click **This Account** and enter the name and password for the user account you wish to use. The service should run with a user account that has permission to execute SSIS packages and, if using Integrated Security, read permissions on all data sources and read/write per-missions on the staging database(s) and data warehouse(s). By default, the service requires the user calling the service to be member of a specific Active Directory Group. You can change this - see change advanced settings below.
6. Click **OK**.
7. (Optional) Change advanced settings, such as the port the service listens to, by editing the config file. Start the service once to have the file created. Then open Explorer and navigate to the following folder:
%APPDATA%\Roaming\SSISWindowsService\Remote SSIS Execution Service\<version>
Open **SSISServiceConfig.xml** in Notepad or another editor.

```
SSISServiceConfig.xml - Notepad                                    _  □  ×
File  Edit  Format  View  Help
<?xml version="1.0" encoding="utf-16"?>
<ServiceSettings xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <ServerPort>16500</ServerPort>
  <CheckUserIsInGroup>true</CheckUserIsInGroup>
  <ADGroup>RemoteSSISOperators</ADGroup>
</ServiceSettings>
```

You have the following options:

Option	Description	Default setting
ServerPort	Specify the TCP port number that the service should listen on.	16500
CheckUserIsInGroup	If set to True, the service will only allow members of the AD Security Group specified in <ADGroup> to use the service. If set to False, all users can use the service.	True
ADGroup	Specify which Active Directory Security Group users have to be members of in order to use the service.	RemoteSSISOperators

Save the file and restart the service to apply your changes.

CONFIGURING YOUR PROJECT TO USE REMOTE EXECUTION OF SSIS PACKAGES

When the service is installed and configured, remote SSIS execution can be enabled on the individual staging database or data warehouse in Jet Data Manager. To enable Remote Executing of SSIS Packages:

1. Right click a data warehouse you wish to enable the feature on and click **Edit Data Warehouse**.
   - OR -
   Right click a staging database you wish to enable the feature on and click **Edit Staging Database**.

2. Click **Advanced...**.
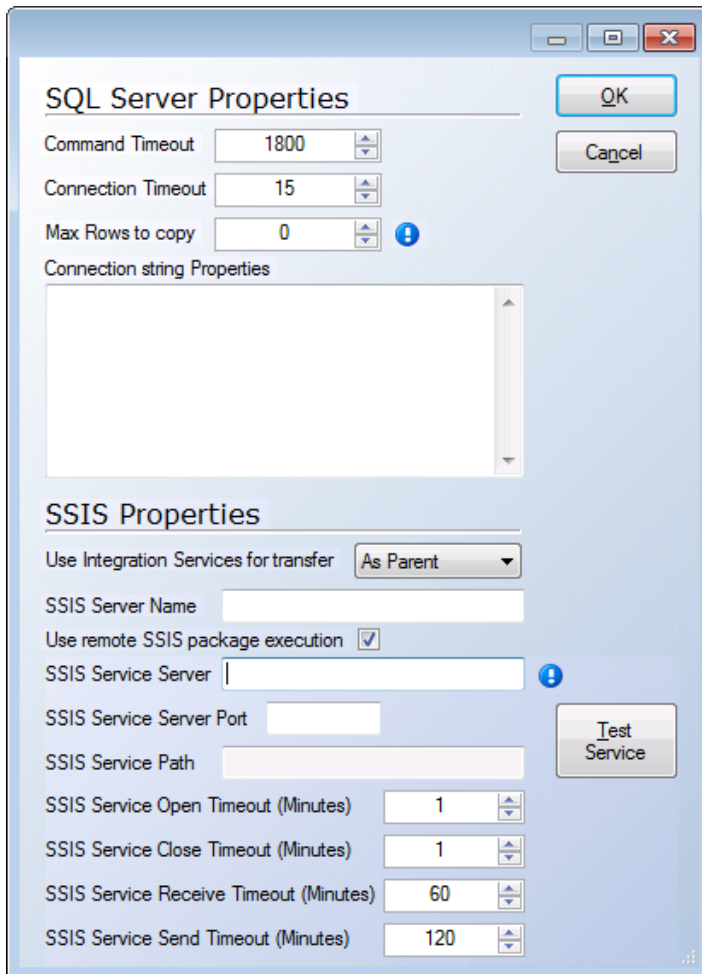
3. Select **Use remote SSIS package execution** to enable the use of the feature.
4. In **SSIS Service Server**, enter the name of the server where the service is running. If you use HTTPS, please write the complete URL including https://.
5. In **SSIS Service Server Port**, enter the port number the service is listening to.
6. **SSIS Service Path** shows the complete URL Jet Data Manager will use to connect to the remote server.

7. (Optional) In **SSIS Service Open Timeout (Minutes)**, enter the timeout for opening connections when no explicit timeout is otherwise specified.

8. (Optional) In **SSIS Service Close Timeout (Minutes)**, enter the timeout for closing channels when no explicit timeout is otherwise specified.

9. **SSIS Service Receive Timeout (Minutes)** is not used in the current implementation, but would control how long Jet Data Manager will wait for an answer from the remote server.

10. (Optional) In **SSIS Service Send Timeout (Minutes)**, enter the value used to initialize the OperationTimeout, which governs the whole process of sending a

message, including receiving a reply message for a request/reply service operation. This timeout also applies when sending reply messages from a callback contract method.

11. Click **Test Service** to test for any connectivity and/or permission issues and click **OK** to save your changes and click **OK** in the data warehouse or staging database settings window.

THE USER INTERFACE

This section provides an introduction to the Jet Data Manager user interface, which is based on a ribbon and a number of tabs that contain the different objects you will be working with. As the number of objects can become quite large, Jet Data Manager includes a feature, Project Perspectives, that lets you hide the objects you are not currently working on.

THE RIBBON

The ribbon presents the tools and actions that are relevant to the current selection in the project tree.
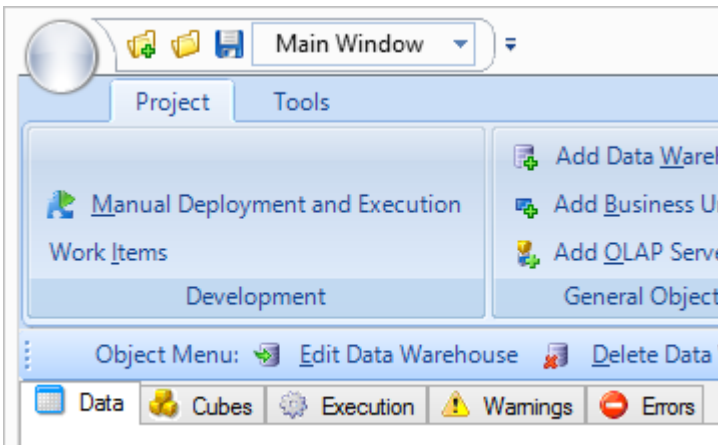
THE TABS

The project is organized in a number of tabs, where each tab contain the tasks, objects and information that relates to a specific part of the project.

THE DATA TAB

The Data tab contains the project tree and all related elements. It is used for specifying business units, data warehouses, and for extracting, transforming, and loading data. When you edit the elements in the project tree, a properties window is displayed to the right when applicable.

THE CUBES TAB

The Cubes tab is used for defining multi-dimensional cubes. The cubes tree contains all the elements you use to define cubes, such as dimensions and measures. Measures and dimensions are listed below the cube they are associated with. A master list of dimensions is listed separately because they can be used in more than one cube.

THE EXECUTION TAB

The Execution tab is used for scheduling the automatic execution, and for keeping track of the execution process. In the Execution Package tree, you can add elements, such as custom actions, checkpoints, notifications, and schedules. The tree lists the element in the order in which they are processed. You can move custom actions up or down in the tree depending on when you want the actions to be executed.

THE WARNINGS TAB

The Warnings tab is used to view warnings resulting from violations of validation rules. You can see which row is affected and which rule has been violated for each warning . Records with warnings will still be promoted to the data warehouse or staging database.

THE ERRORS TAB

The Errors tab is used to view errors resulting from violations of validation rules. For each error, you can see which row is affected and which rule has been violated. Records with errors will not be promoted to the data warehouse or staging database and are excluded from the final data set.

PROJECT PERSPECTIVES

The purpose of Project Perspectives is to make it easier to work with large projects. Working in a big project can make it hard to maintain a good overview and find an individual object quickly.

The idea is that you can create different perspectives on a project. A perspective is a subset of the project objects that relates to a specific area or task. For example, you could create a "finance" perspective that contains all the tables, dimensions and cubes that are related to finance. When this perspective is active, anything else will be hidden in the project tree.

An object can be in any number of perspectives. You can also chose to make a project perspective dynamic. Any object that depend on an object already in the perspective, will then automatically be included in the perspective.

ADDING A PROJECT PERSPECTIVE

To add a project perspective, follow the steps below.

1. On the **Data** tab, right click **Project Perspectives** and click **Add/Edit Project Perspectives**. Note, that if you are adding your first perspective to the project, you need to right click on the project, click **Advanced** and click **Add/Edit Project Perspectives**. The **Project Perspectives** window open.
2. Click **Add Perspective**. A new column is added to the grid with the name "Perspective" in the **Perspective Name** row. Double-click the name to edit it and click outside the field when you have finished typing. Select **Dynamic Perspective** if you

want the new project perspective to be dynamic. You can change this setting at any time.

3. Each OLAP server, data warehouse and business unit is listed in the first column of the grid. Click the + besides the OLAP server, data warehouse or business unit to show the objects child elements. Select the cubes, dimensions and tables you wish to include in your project perspective. If you add an object to a dynamic perspective, dependent objects will be added automatically. You can recognize dynamically added objects by the checkbox with an indeterminate state.

4. Click **OK**. The new project perspective now appears under Project Perspectives in the project tree.

ADDING OBJECTS TO AND REMOVING OBJECTS FROM A PERSPECTIVE

You can add most objects - tables, fields, dimensions, cubes - to a perspective. You can add the same object to as many perspectives as you need to. To add an object to a perspective or remove an object from a perspective, follow the steps below:

1. Right click the object you wish to add and navigate to **Project Perspectives**.
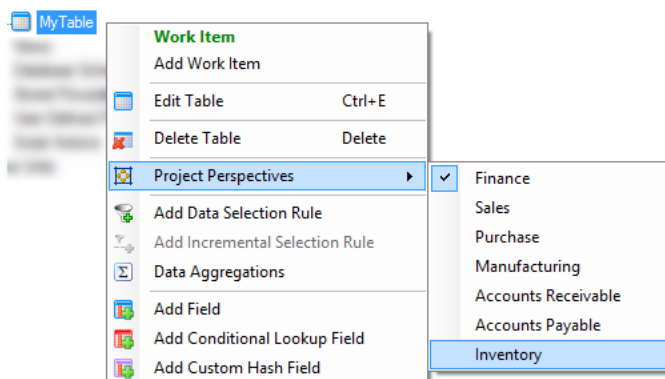
Here, the perspectives that the object is currently a part of have a checkmark next to them.

2. Click the name of an unchecked perspective to add the object to this perspective
- OR -
Click the name of a checked perspective to remove the object from this perspective.

ACTIVATING A PERSPECTIVE

There are three ways of activating a perspective:

- In the project tree, navigate to **Project Perspectives**, right click the name of the perspective you wish to activate and click **Use Project Perspective**.
- In the ribbon, navigate to **Tools**. In the **Project Perspectives** group, click the **Project Perspectives** list and then click the perspective you wish to activate.
- In the Quick Access Toolbar, click the **Project Perspectives** list and then click the perspective you wish to activate.

> **Note:** The Project Perspectives list is not shown in the Quick Access Toolbar when there are no perspectives in the project.

DEACTIVATING ALL PERSPECTIVES

To disable all perspectives and see all objects, you have two options.

- Click the **None** perspective in the **Project Perspectives** list in the ribbon or the Quick Access Toolbar as described in "Activation a perspective" above.
- In the project tree, navigate to **Project Perspectives**, right click the name of the currently active perspective and click **Use Project Perspective**. This will remove the Checkmark from **Use Project Perspective** and change the current perspective to "None".

SORTING OBJECTS IN A PERSPECTIVE

Objects within a perspective can be sorted by the execution order or alphabetically. To change the sort order of the active perspective, follow the steps below:

1. Click **Project Perspectives** in the project tree.
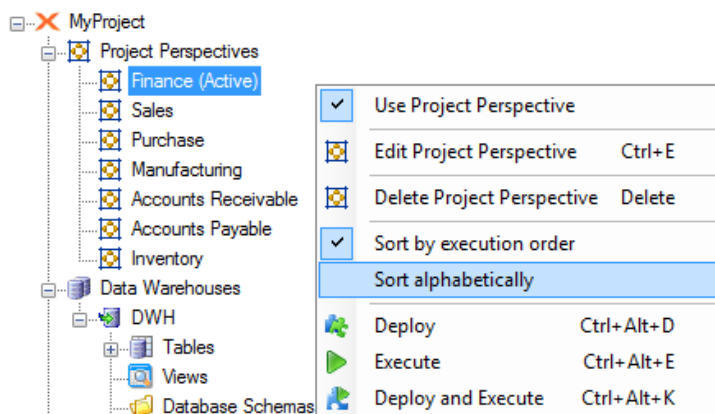
2. Right click the name of the currently active perspective and select either **Sort by execution** order or **Sort alphabetically**. The chosen sort order will be saved for each perspective.

> **Note:** It is not possible change the order of objects manually while a perspective is active.

DEPLOYING AND EXECUTING A PERSPECTIVE

Perspectives can be deployed and executed just as other objects. This enables you to easily work with a subset of you project from source to execution. You can deploy and/or execute a perspective in three ways:

- Right click your project in the project tree, navigate to **Deploy** and click **Deploy Current Perspective**

  - OR -

  Navigate to **Execute** and click **Execute Current Perspective**

  - OR -

  navigate to **Deploy and Execute** and click **Deploy and Execute Current Perspective**
- In the project tree, navigate to **Project Perspectives**, right click the perspective you wish to deploy and/or execute and click **Deploy**, **Execute** or **Deploy and Execute**.
- You can also add a perspective to a Execution Package, for example if you wish to execute the perspective on a schedule.

THE WORKFLOW

In the following, the steps you typically need to go through when you build a data warehouse solution with Jet Data Manager are briefly described. However, you are in complete control. You can base your solution on a prebuilt project from the CubeStore, import parts of a legacy data warehouse and built a cube structure on top of that or start from scratch.

CREATING A JET DATA MANAGER PROJECT

The project is a container for all of the elements of your solution. Thus, creating or importing a project is always the first step when you build a new solution. The key elements of a project are the data warehouse and the business unit.

The data warehouse is a Microsoft SQL Server database. The data warehouse stores all of the extracted and cleansed data that you need for query and analysis.

Business units represent separate units within your organization. For example, if you have a global organization, you could create one business unit for the world headquarters, and separate business units for each subsidiary. A business unit contains a staging database and one or more data sources.

For more information, see [Projects](#).

DESIGNING THE DATA WAREHOUSE

Jet Data Manager is built to utilize the star and snowflake schemas, the two standard schemas for designing a data warehouse. This means that you'll often find yourself building fact and dimension tables. The fact tables contain the transactional data you wish to analyze, while the dimension tables add context to these data.

When you design the data warehouse in the Jet Data Manager, you create relationships between the tables in your data warehouse, assign primary keys, create views and so on.

For more information, see [Designing the Data Warehouse](#).

MODELING CUBES

Some front-end presentation tools connect directly to your data warehouse. Often, however, Online Analytical Processing (OLAP) cubes will serve as the basis for your analysis and reporting solution.

To create cubes, you use dimensions to structure how you want to analyze your data and measures to specify which numerical values you want to analyze. Cubes are stored in an OLAP database, and you can use your preferred front-end application to drill-down or roll-up through the data.

You can also reverse engineer an existing OLAP database and then use Jet Data Manager to maintain and change existing cubes.

For more information, see [Online Analytical Processing (OLAP)](#).

CONNECTING TO DATA SOURCES

The Jet Data Manager supports a wide variety of data sources, such as Microsoft SQL Server databases, Oracle databases, Microsoft Excel files and text files. Application adapters simplify access to all major ERP and CRM systems. Furthermore, access to generic or legacy databases is possible through generic ODBC.

For more information, see [Data Sources](#).

SELECTING AND CLEANSING DATA

Selecting data is the process where you identify which data you need to extract from the data source.

After selecting tables and fields, you further limit your selection by applying Data Selection Rules.

The selected data is moved to a staging database where the cleansing process takes place according to the transformation and validation rules you have specified. The transformation and validation rules ensure consistent data, uniform formatting of data, and removal of duplicate data.

The staging database stores the cleansed data temporarily until it has been cleansed and stored in the data warehouse. Extracting the selected data to a staging database means that the cleansing process has very little effect on the transaction database, and thus on the daily business operations.

The staging database consists of different tables that store the extracted data before and after the cleansing process. To handle the data transformations, a number of views are created. The SQL code that is generated during the cleansing process is stored in the staging database.

For more information, see [Selecting, Validating and Transforming Data](#).

DEPLOYING AND EXECUTING PROJECTS

During deployment, the structure of your data warehouse and of your cubes is created.

When you execute a project, data is loaded into the data warehouse and the cubes are processed.

During deployment, the underlying SQL code is automatically generated and stored in the data warehouse.

For more information, see [Deploying and Executing Projects](#).

SETTING UP YOUR PROJECT

In Jet Data Manager, your work is organized in projects stored in a repository on a SQL Server. Each project is a collection of different objects: data warehouses, tables, fields, business units, execution packages, OLAP servers and so on.

The objects are organized into different tabs in the user interface. In this chapter, we will cover the project itself, data warehouses, business units and staging databases, which you can all find on the Data tab.
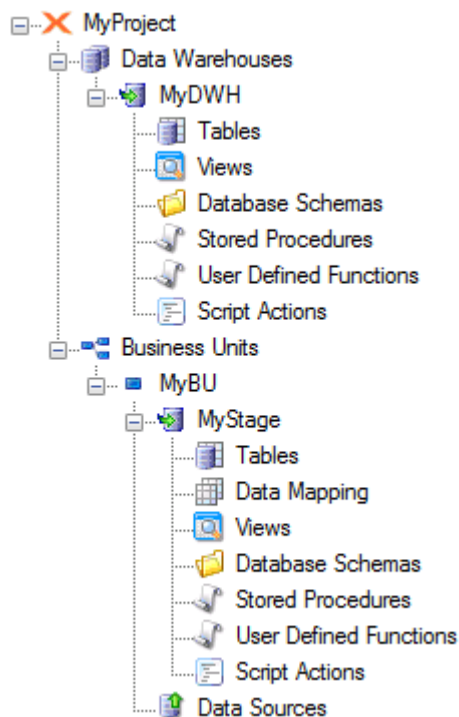
Directly below the project in the tree are Data Warehouses and Business Units. Each can contain a number of data warehouse databases or business units. In turn, each business unit contains a staging database and a number of data sources.

You can learn more about data sources for your project in the chapters about Data Sources and Application Adapters. The content of the Cubes tab is covered in the Online Analytical Processing (OLAP) chapter, while the Execution tab is covered in the Deploying and Executing Projects chapter.
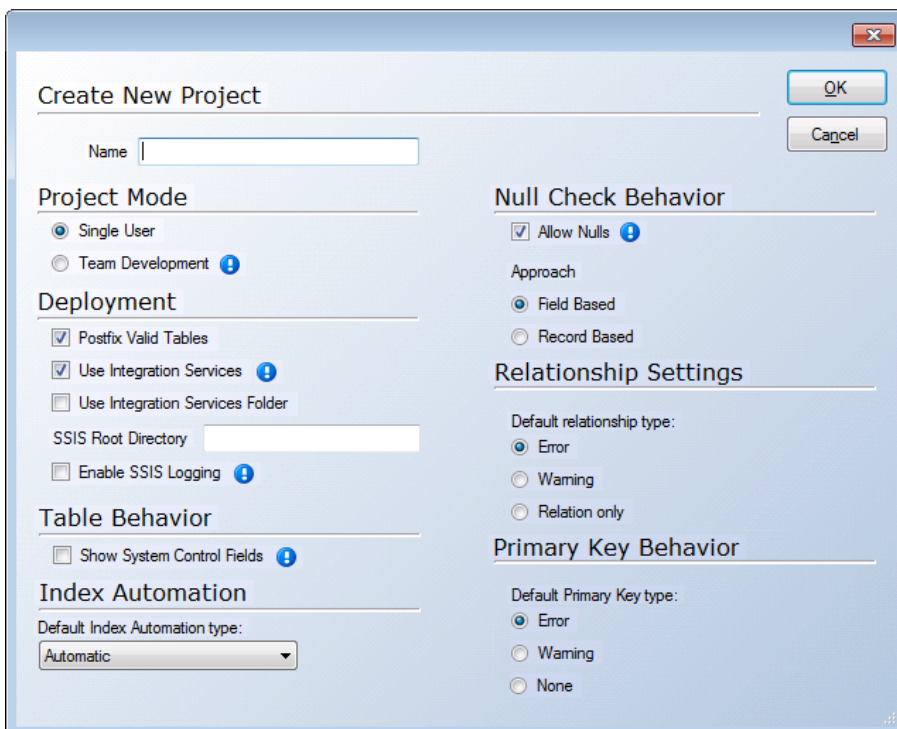
PROJECTS

A Jet Data Manager project contains all other elements in the data warehouse.

You can only have one project open at a time. However, if you want to compare different versions of a project, you can open another instance of Jet Data Manager, and load another version of the project to view side-by-side.

CREATING A PROJECT

To create a projects follow the steps below:

1. Click the application icon and click **New Project...**.

2. In the **Create New Project** window, type a name for the new project.
3. Configure the settings to fit your needs. Depending on what features you have licensed, some settings might not be available.
   - Under **Project Mode**, click **Team Development** if you wish to use the multiple developers feature.
   - Under **Deployment**:
     - Select the **Postfix Valid Tables** checkbox to postfix valid tables with [...].
     - Select **Use Integration Services** to use SQL Server Integration Services (SSIS) for data transfer. If cleared, Jet Data Manager will use ADO.net instead.
     - Select **Use Integration Services Folder** to create a folder for the SSIS packages created. This prevents accidental overwriting of the packages by other projects that contain tables with the same name.
     - Select **Enable SSIS Logging** to...

- Under **Table Behavior**, select **Show System Control Fields** to show system control fields such as DW_ID, DW_Batch, DW_SourceCode and DW_ Timestamp in the GUI.
- Under **Index Automation**, choose the default setting for automatic index generation. Select **Automatic** to automatically create indexes as needed, **Manual** to enable automatic index generation on an a per-table basis or **Disabled** to create indexes as needed during data cleansing as in earlier versions of Jet Data Manager. This legacy behavior comes with a performance penalty, since indexes are always recreated even when an existing index could be used.
- Under **Null Check Bahavior**, select **Allow Nulls** to allow null field values instead of moving the row to the error table when encountering null values. Click **Field Based** to use a field based check or click **Row Based** to use a row based check. A field based check will tell you exactly where the null value is, while a row based check will only tell you that the record has a a null value.
- Under **Relationship Settings**, choose how Jet Data Manager should treat foreign key violations. Click **Error** to move the offending row to the error table, click **Warning** to move the offending row to the warning table or click **Relation only** to ignore the violation.
- Under **Primary Key Behavior**, choose how Jet Data Manager should treat primary key violations. Click **Error** to move the offending row to the error table, click **Warning** to move the offending row to the warning table or click **None** to ignore the violation.

You can only have one project open at a time in Jet Data Manager. If you try to create a new project or load an existing project when you already have one project open, you will be asked to save the current project.

Note: The first time you start working with Jet Data Manager, you must specify a project repository in which all projects are stored. For more information, see Setting Up a Project Repository.

SAVING AND OPENING PROJECTS

Saving a project in the repository and opening a project from the repository works much the same as in other Windows programs.

SAVING A PROJECT

Jet Data Manager includes version control. On every save, a new version of the project is saved, meaning that you can always go back to an earlier version of your project.

To save the project, follow the steps below.

1. From the **File** menu, click **Save** or **Save As**.
2. If you clicked on **Save As**, or you are saving the project for the first time, type a name for the project and click **OK**.

Note: The project is automatically saved after a succesful deployment.

OPENING A PROJECT

1. From the **File** menu, click **Open**.
2. (Optional) Click **Change Version**, click the version you wish to open and click **OK** if you wish to open a version of the project other than the latest version.
3. In the **Project** list, select the project that you want to open, and then click **OK**.

EXPORTING AND IMPORTING PROJECTS

Being able to export a project to an XML document is useful when you want to save a copy of a running project for future reference, or if you want to reuse parts of a project in another project. You can export a project to an XML document, and you can import a project from an XML document.

IMPORTING PROJECTS FROM XML DOCUMENTS

1. From the **File** menu, choose Import/Export, and then click Import Project.
2. In the **Import** File field, click the ellipsis **(…)**, and then navigate to and select the file you want to import.
3. Click **Open**, and then click **OK**.

EXPORTING PROJECTS TO XML DOCUMENTS

1. From the **File** menu, choose **Import/Export**, and then click **Export Project**
2. In the **Export File** field, click the ellipsis (**…**), and then navigate to and select the file you want to export to.
3. Click **Open**, and then click **OK**.

PROJECT REPOSITORIES

Your project is stored in a project repository on a SQL Server. You can have as many or as few projects in a repository as you wish. When you open Jet Data Manager for the first time, you'll be prompted to set up your repository. All projects you create in the future will be saved in the specified repository, unless you change the repository.

If the following message is displayed "Insufficient rights - must be a member ofJet Data Manager administrator group", please contact your system administrator for more information.

Note: When you install a new version of Jet Data Manager, you will be prompted to run an upgrade script that automatically updates the repository to ensure compatibility with the new software version.

SETTING UP A PROJECT REPOSITORY

To change the repository settings, follow the steps below:

1. In the ribbon, on the **Tools** tab, click **General Settings** in the **Administration** group.
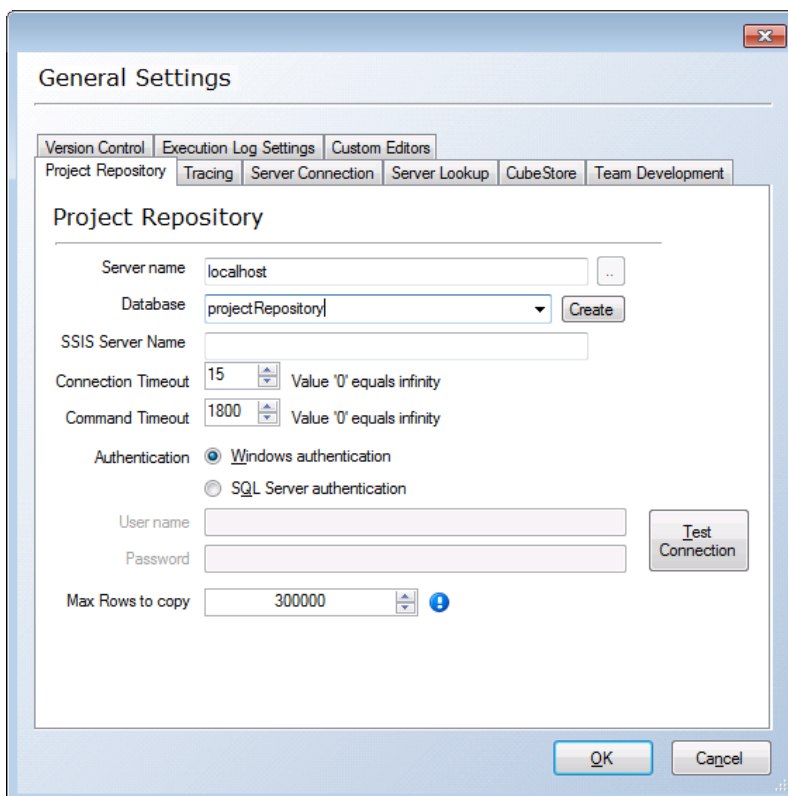
2. In the **Server Name** field, enter the name of the database server on which you want to store the project.
3. In the **Database** list, enter a name for the database, and then click **Create**. Alternatively, you can select an existing database from the list.

4. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server. The default is 15 seconds. A value of zero will disable the timeout.
5. In the **Command Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the database. The default is 1800 seconds. A value of zero will disable the timeout.
6. Specify the authentication mode. The default setting is Windows authentication. If you choose SQL Server authentication, you are prompted for a user name and a password.
7. Click **Test Connection** to verify that the connection is working.
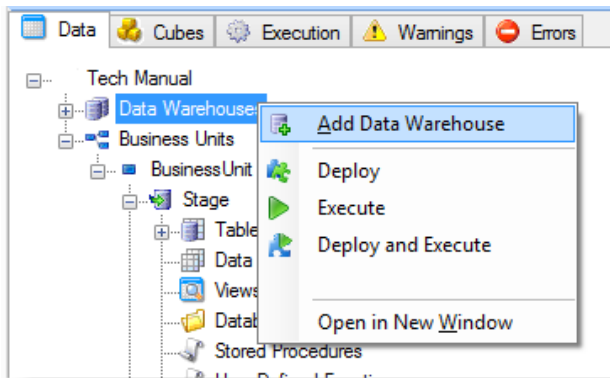
DATA WAREHOUSES

A data warehouse in Jet Data Manager is a SQL Server database where your data is stored for queries and analysis. Most often, a Jet Data Manager project consists of one data warehouse where you consolidate data from one of more staging databases and a number of data sources.

During execution of a project, Jet Data Manager extracts data from the staging database and transfers it to the data warehouse. Initially, the data resides in what is known as raw table instances in the data warehouse. Jet Data Manager applies data transformations and data cleansing rules and saves the resulting data in valid instances of the tables, ready for queries and analysis.

ADDING A DATA WAREHOUSE

1. On the **Data** tab, right-click **Data Warehouse**, and then select **Add Data Warehouse**.
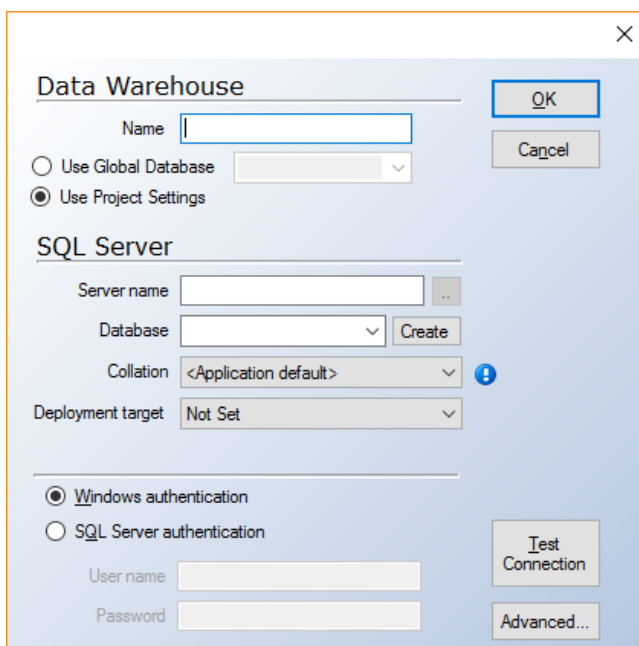
2. In the **Name** field, type a name for the data warehouse. The name cannot exceed 15 characters in length.
3. In the **Server Name** field, type the name of the server that you want to store the database on. If it is a named instance, type the server name and the instance name.
4. In the **Database** field, select an existing database, or type the name of a new database, and then click **Create**.
5. Specify the authentication mode. The default setting is **Windows Authentication**. If you choose **SQL Server Authentication**, you enter the user name and password.
6. Click **Test Connection** to verify that the connection is working.
7. Click **Advanced...** to access the advanced settings for the data warehouse. These are all optional.

8. In the **Command Timeout** field, enter the number of seconds to wait before terminating a command.
9. In the **Connection Timeout** field, enter the number of seconds to wait before terminating the attempt to connect to the server. Set it to 0 to wait indefinitely.
10. If you want to add additional connection strings, enter them in the **Connection String Properties** box.
11. See Adding a Database Role for an explanation of the **Drop role options** setting.
12. In the **Use Integration Services for transfer list** click on **Yes** to enable SSIS data transfer, **No** to disable it or leave it at **As Parent** to respect the project setting.

13. If your SSIS Server is installed under a different name than the database, enter the name in the **SSIS Server Name** box.
14. Select **Use Remote SSIS package execution** to enable the execution of SSIS packages on a remote server and enter the details for the remote server below. Note that you will need to install the Remote SSIS Execution service on the remote server - see [Installing the Remote SSIS Execution Feature](#)

ADDING TABLES TO A DATA WAREHOUSE

1. Expand **Data Warehouses**, and then expand the preferred data warehouse.
2. Select **Tables**, right-click, and then select **Add table**.
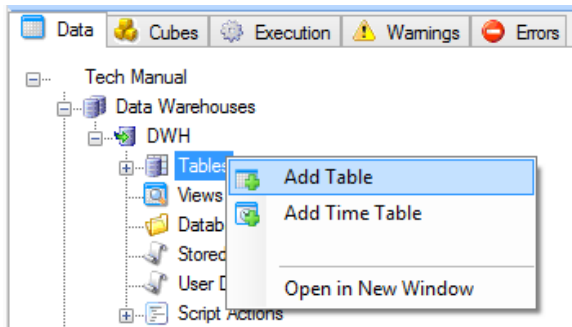
3. In the **Name** field, type a name for the table, and then click **OK**.

> **Note:** You can also move tables from a staging database to the data warehouse by dragging and dropping the tables. See [Moving Data from the Staging Database](#).

ASSIGNING PRIMARY KEYS

All tables in your project can have a primary key that uniquely identifies every row in the table. If you are consolidating data from different business units, you must assign primary keys to avoid duplicate values in your dimensions.

1. Expand **Business Units**, expand the preferred business unit, and then expand the staging database.
   or
   Expand **Data Warehouses**, expand the preferred data warehouse, and then select **Tables**.
2. Expand the table that you want to modify.

3. Right-click the field you wish to modify, and then select **Include in Primary Key**. The primary key can be a composite key, i.e. based on more than one field.
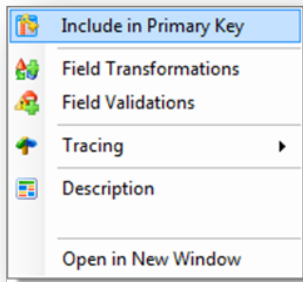
ADDING CUSTOM FIELDS

You can add custom fields to both the staging database and the data warehouse.

1. Expand the preferred business unit, expand the staging database, and then expand Tables.
   or
   Expand the preferred data warehouse, and then expand Tables.
2. Select the table to which you want to add a custom field, right-click, and then select **Add Field**.
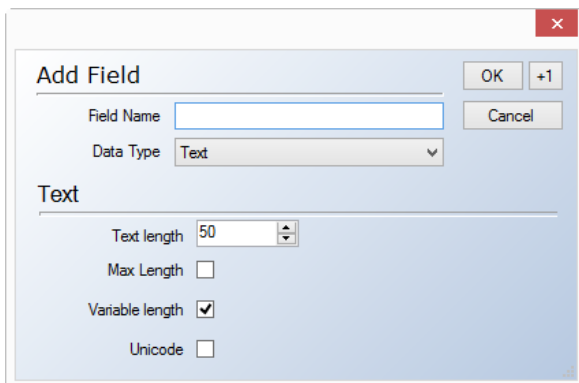
3. In the **Field name** field, type a name for the field.
4. In the **Data type** list, select the data type.
5. Define the attributes of the selected data type. You have the following options:

Data Type	Attributes	Description
Text	Text length	Specifies the maximum number of characters the field can contain
	Variable length	Specifies that the field can be of variable length
	Unicode	Specifies whether to use Unicode character-encoding

Integer	Type	The size of the integer: bigint, int, smallint, tinyint
Numeric	Is float	Specifies if the number is a float
	Number of decimals	Specifies the number of decimals in the field
Binary	Length	The length of the binary field
	Max length	

SPECIFYING THE EXECUTION ORDER OF TABLES

When you create a standard execution package, tables are executed in the order in which they appear in the tree on the **Data** tab. To avoid errors when executing, you must therefore ensure that tables are executed in logical order. For example, an Order table must be executed before the related Order Detail table. You may, therefore, have to move tables up and down in the data warehouse tree.

Note: For advanced execution packages, the order will be determined in the **Execution Setup** window itself. For more information, see Adding Execution Packages.

TO MOVE TABLES UP AND DOWN IN THE TREE

1. On the Data tab, expand data warehouse or staging database, and select the table that you want to move.
2. Left-click and drag and drop the table to the preferred location. Alternatively, use ALT+UP ARROW or ALT + DOWN ARROW to change the order in which the table will appear in the execution order.
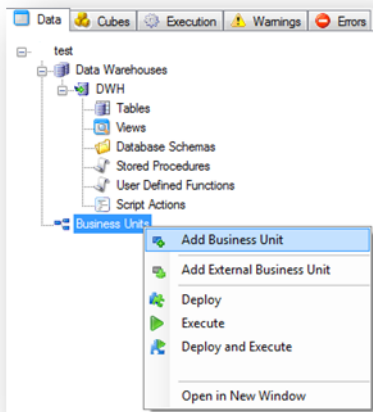
BUSINESS UNITS

In Jet Data Manager, a business unit is any part of your organization that you want to treat as a separate entity in your project. For example, you may want to treat a company headquarters and each of its subsidiaries as separate business units.

Each business unit in your project has its own staging databases and its own data sources.

ADDING A BUSINESS UNIT

To add a new business unit, follow the steps below:

1. On the **Data** tab, right-click **Business Units**, and choose **Add Business Unit**.

2. In the **Add Business Unit** dialog, type a name for the business unit, and then click **OK**.


3. When you create a business unit, you always have to specify a staging database. See Setting Up a Staging Database.

ADDING AN EXTERNAL BUSINESS UNIT

You can add business units from other projects to your project. This enables reuse of business units across different projects. To add an external business unit, follow the steps below:

1. On the **Data** tab, right-click **Business Units**, and then choose **Add External Business Unit**.
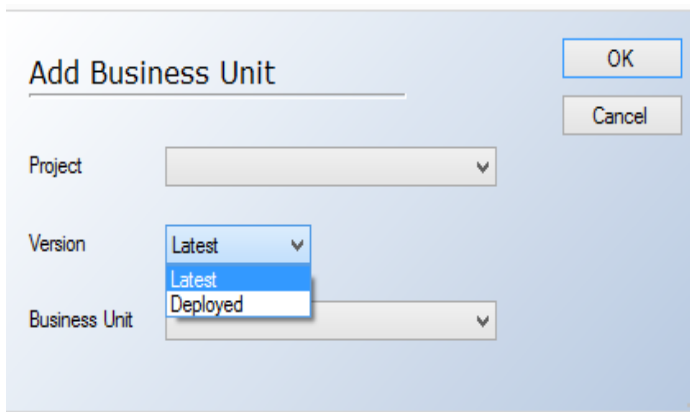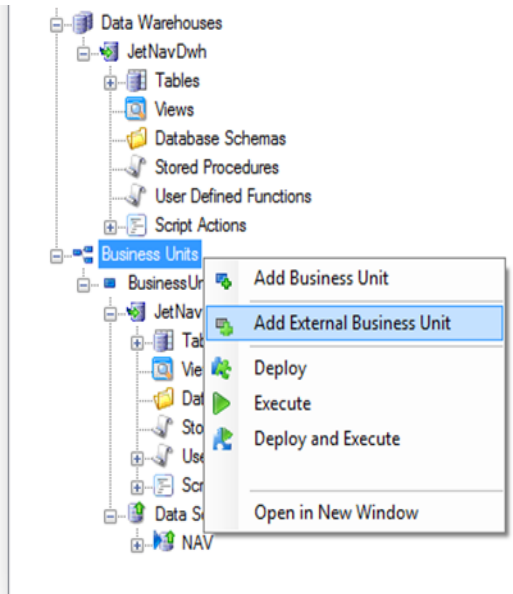
2. In the **Project** list, select the project to amend.
3. Select which version of the business unit you want to import. You have the following options:

Version	Description
Latest	Adds the last saved version of the project, which does not always correspond to the last deployed version
Deployed	Adds the last deployed version of the project

4. In the Business Unit list, select the business unit you want to add, and then click **OK**. The business unit is displayed as a separate entity in the project tree.

If you make changes in the business unit, you might want to synchronize the changes to your project. To synchronize an external business unit, follow the steps below:

1. On the **Data** tab, expand **Business Units**, and then right-click the external business unit you want to synchronize.
2. Select **By ID** to synchronize fields by ID, or select **By Name** to synchronize fields by name.

STAGING DATABASES

A business unit always contains a staging database. The staging database stores the selected data from the data sources. Additionally, many of the validation and transformation processes take place in the staging database. This ensures that the cleansing process has limited impact on the transaction database.

You create table relationships and mapping tables on the staging database. It is also possible to add views, stored procedures, and user defined functions to the staging database. When you execute a project, all invalid rows are stored in either the Warnings table or the Errors table.

SETTING UP A STAGING DATABASE

When you add a Business Unit, the **Add Staging Database** dialog is displayed to let specify a database.

1. In the **Name** field, type a name for the staging database. The name cannot exceed 15 characters in length.
2. In the **Server Name** field, type the name of the database server.

3. In the Database list, select the preferred database from the drop-down list. If you wish to create a new database, then type a name for the new database, and then click **Create**.
4. Specify the authentication mode. The default setting is **Windows authentication**. If you choose**SQL Server authentication**, you are prompted for a user name and a password.
5. Click **Test Connection** to verify that the connection is working.
6. In the **Command Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the database. The recommended value for this is 0 to disable the command timeout.
7. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server.
8. If you want to add additional connection strings, click the **Additional Connection Properties** button. In the **Connection String Properties** window, type the preferred connection strings, and then click **OK.**

TEAM DEVELOPMENT

Jet Data Manager supports multiple developers working on the same project at the same time. Version notes enable developers to share details about their changes to the project and work items allows developers to see which objects other team members are working on to prevent them from modifying the same objects simultaneously.

When collaborating on a project, developers should avoid working on the same object on the same time. If this happens, the outcome depends on the type of change made:

- If multiple team members add or modify different sub-objects (such as fields) within the same object (such as a table), then Jet Data Manager will save these changes and display them the next time that Jet Data Manager is opened. For example, if one team member adds two fields to a table, and another team member adds another two fields to the same table, all four fields will be visible to all team members the next time they close and re-open the project.
- If multiple team members modify the same sub-object at the same time (for instance renaming a field), the changes deployed and executed last will take precedence, and the project will reflect those changes.

The team development feature of Jet Data Manager depends on the developers to take care when working together as the feature itself does not prevent developers from making conflicting changes.

Follow the rules below to ensure that the project stays consistent:

1. Never work on the same object - table, dimension or cube - as someone else.
    - Always create work items before making any changes to an object.
    - If anyone else has the item marked for their use, wait for them to release the work item.
    - Once you are done working on an object, save the project and release the work item.
    - Keep the work items window open when working.
2. Do not rename tables, cubes or dimensions. The only exceptions are if you just created the table and have not saved yet or if nobody else has the project open. You can see who has a project open in the work items dialog.
3. Always save (CTRL + S) and reload (CTRL + F5) before you start a new development task. This ensures that you have the latest version of all objects before you start taking over somebody else's work.
4. Immediately before creating a new table, always save and reload.

ENABLING THE TEAM DEVELOPMENT ENVIRONMENT

You need to enable Team Development for any project you wish to use the feature in.

1. In the ribbon, click the **Tools** tab and then click **Repository Administration**.
2. On the **Projects** tab, right-click the relevant project and click **Enable Team Development.**
3. Close the window. Multiple users can now access the project concurrently.

WORK ITEMS

Work items allow the team to know what its users are working on. This will give a visual indication as to what areas of the project are currently under development. Work items are meant to be created immediately prior to starting work on a set of objects and can be either manually deleted or removed during deployment of the object. Work items can be added to the following objects:

- Project
- Data Warehouse
- Data Warehouse table
- Staging Database
- Staging Database table
- Data Source Table
- OLAP Database
- Cube

ADDING WORK ITEMS

To add a work item, follow the steps below.

1. Right-click the relevant object and click **Add Work Item**.

If the object that the user is adding the work item to already has a work item created by another user, it will display **Add Work Item** in red and identify the other user below. This allows users to know if other users are modifying the same object in order to facilitate collaboration.


Otherwise, the Create Work Item window appears.

2. In **Description**type a short description of the work item, e.g. what task the work item is a part of, and click **Create Work Item**.
3. Click **Create Work Item.**

VIEWING EXISTING WORK ITEMS

To view existing work items, follow the steps below.

1. In the ribbon on the **Project** tab in the **Development** group, click **Work Items**.
2. The **Work Item** window opens. It contains a list of your own work items, **My work items**, and a list of **Other users work items**.
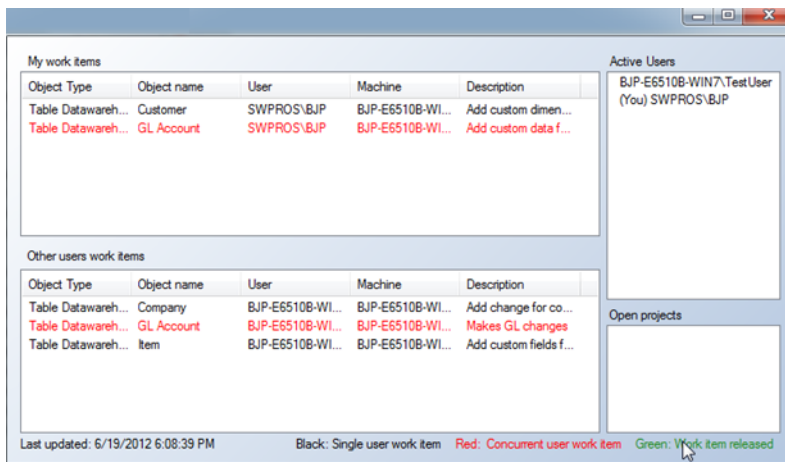
Each list displays the following information:

Column	Description
Object Type	The type of object the work item is associated with (project, database, table, or cube)
Object Name	The name of the project, database, table, or cube associated with the work item
User	The user name of the person who created the work item
Machine	The machine that the work item was created on
Description	Descriptive text defined by the user for the work item

Work items can have three colors depending on their status:
- Black: A work item that is only flagged by a single user.
- Red: Work items that share the same Object Name for different users. This tells you that some collaboration will be needed regarding which user is accessing the data. While Team Development allows multiple users to modify the project concurrently, these users should not modify the same object at the same time.
- Green: Work items that have been completed by another user who has saved or deployed and marked the work item as completed.

In the **Active Users** pane, all users currently accessing the project are shown.

EDITING A WORK ITEM

Work items can only be edited by the user that created the work item. To edit a work item, follow the steps below:

1. In the ribbon on the **Project** tab in the **Development** group, click **Work Items**.
2. In **My Work Items**, right-click the work item to edit and select **Edit Work Item**.
3. Update the work item descriptiop and click**Update Work Item** to save the changes.
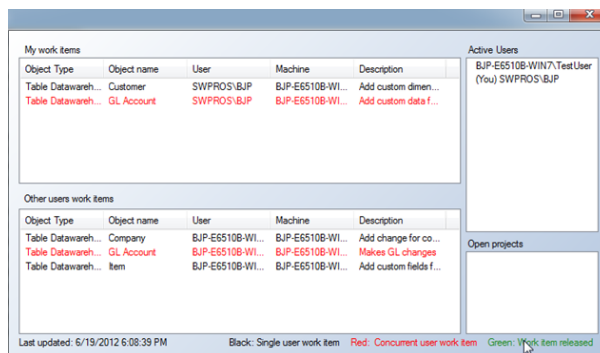
DELETING WORK ITEMS

Work items can only be edited by the user that created the Work item. You can delete a work item in three different ways:

- Right click the relevant object and click **Delete Work Item**

- In the ribbon, in the **Administration** group, click **Repository Administration**. On the **Work items** tab, click a work item an click **Delete**.
- In the ribbon, on the **Project** tab in the **Development** group, click **Work Items**.

In **My work items**, right-click the relevant work item to edit and click **Delete Work Item**. Jet Data Manager will ask you to confirm the deletion. Click **Yes**.

Work items can also be removed when adding version notes, which is discussed below.

VERSION NOTES

Version notes enables you to add comments about changes that you have made to a project. These comments can then be viewed in the future to reference changes that were made.

ADDING VERSION NOTES

1. To ensure version notes are enabled, go to the **Tools** tab in the ribbon and click **General Settings**. The General Settings window appears.
2. Click the **Version Control** tab and click one of the following options in the list:
   - **Prompt on Every Project Save**: You will be prompted to enter a Version Note when you deploy an object or when you click the **Save** button.
   - **Prompt on Project Deployment Save**: You will only be prompted to enter a version note when you deploy an object.
3. When you deploy an object in the project, the **Version Details** window will appear just before the deployment begins.
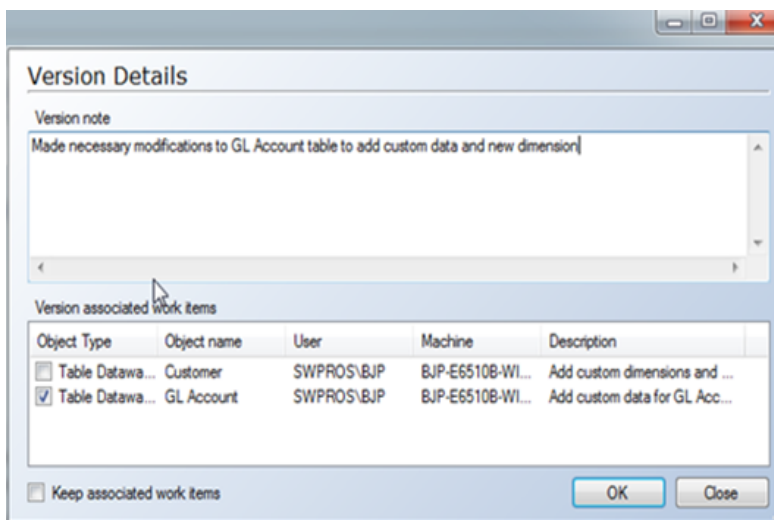

You can type in any details you wish to include in the version note. If relevant, you can also select one of your existing work items to associate with the version note. Selecting a work item will also remove the work item from the list of existing work items. If you wish to associate a work item with the version note, but not remove the work item, select **Keep Associated Work Items**. If you do not wish to save a version note for this deployment, you can click **Close**.This will close the window without saving a version note.

VIEWING VERSION NOTES

Viewing version notes can be useful when you would like to look back and see when particular changes were made. Follow the steps below to see the list of project versions and associated version notes.

1. In the ribbon, on the **Tools** tab, click **Repository Administration**. The Repository Adminstration window appears.
2. Right click the relevant project and click **Show Project Versions**. The Project Versions window appears.
3. If a notepad icon is displayed in the **Note** column, a version note is available for the version. Hover over the notepad icon to see a quick description of the note.
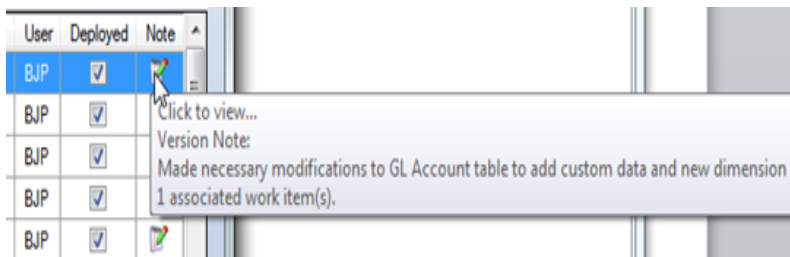
4. Click the notepad icon to see the full version note as well as any work items associated with the version note.
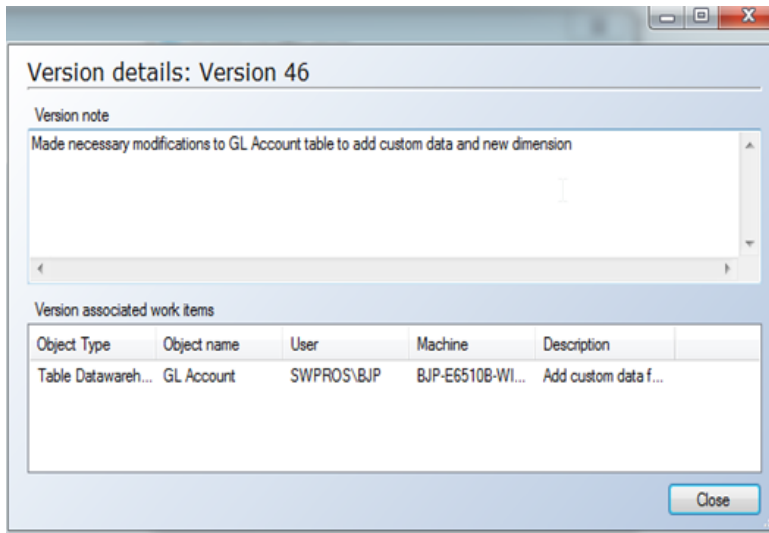
CONNECTING TO DATA SOURCES

Data sources contain the data that you want to load into your data warehouse and use for analysis. The current version of Jet Data Manager connects to the following data sources:

- Microsoft SQL Server
- IBM DB2
- IBM Informix
- Oracle
- Oracle MySQL
- ODBC compliant data sources
- Microsoft Excel files
- Plain text files
- AnySource OLE DB/ADO.NET
- Custom Data Sources

In this chapter, you will learn how to connect to the different data sources.

DATA EXTRACTION

SYNCHRONIZING DATA FROM DATA SOURCES

To see the tables available in a data source, you can synchronize the data in your data source with the data in your project.

1. On the **Data** tab, expand **Business Units**, and then expand the relevant business unit.
2. Expand **Data Sources**, and right-click the relevant data source, and select **Synchronize Data Source**.

All tables and fields that have been removed from, or added to, the data source are listed in a separate window.

SPECIFYING THE DATABASE COLLATION

Jet Data Manager has the ability to change the collation of the staging database, data warehouse, and OLAP databases from within Jet Data Manager itself. To specify the database collation, follow the steps below:

1. Right-click the database you would like to specify the collation for, and select **Edit**.
2. From the **Collation** drop-down, choose the desired collation.
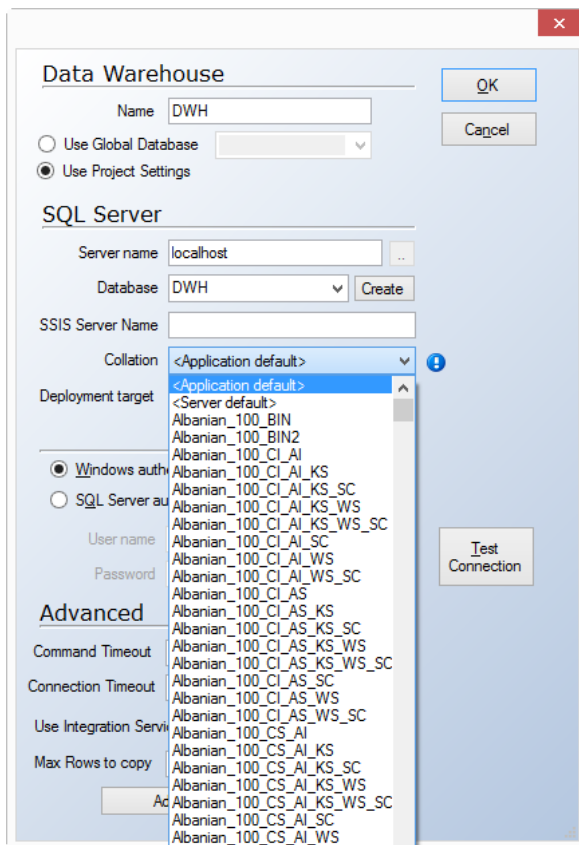
<Application Default> will use the original **Latin1_General_CI_AS** collation.
<Server Default> will inherit the collation from the specified server.

3. Click **OK**.

Note: The best practice is to select the same collation setting for the staging database, data warehouse database and OLAP database.

LIMITING CONCURRENT TRANSFERS ON DATA SOURCE

You can put a limit on the number of concurrent transfers from a specific data source. Some data sources can only handle a certain number of transfers before adding more transfers will actually slow down the overall performance of the transfer rather than speed it up.

To limit the amount of concurrent transfers from a data source, follow the steps below.

1. On the **Data** tab, expand **Business Units**, the relevant business unit and **Data Sources**. Right click the data source in question, click **Advanced** and **Advanced Settings…** The **Advanced Settings** window opens.
2. In **Max Concurrent Transfers**, enter the maximum number of concurrent trans-fers, and click **OK**.

ALLOWING A DATA SOURCE TO FAIL

If you have some source systems in your solution that are less than critical for your report-ing, you can configure your solution so that the entire execution does not stop just because Jet Data Manager cannot reach these noncritical systems. Instead, you can choose to keep the newest data from the system in question until fresh data can be fetched.

Follow the steps below to allow a data source to fail.

1. On the **Data** tab, expand **Business Units**, the relevant business unit and **Data Sources**. Right click the data source in question, click **Advanced** and **Ad-vanced Settings…** The **Advanced Settings** window opens.
2. In the **Transfer Failure Option** list, click the option you wish to use if the transfer should fail. You have the following options:
   - **Fail stop execution**: Stops execution and reports the execution as failed (default setting).
   - **Continue without data**: Continues the execution, pretending that the source contains no data.
   - **Continue with existing data**: Continues execution, retaining the existing data from the source.

Note: If you use additional connectors, you need to configure this setting on the addi-tional data source as well as the template data source.

If a data source, that you have allowed to fail, fails during execution, the following execu-tion message will be recorded: "Execution was successful, but one or more data sources failed".

INTEGRATION SERVICES

On most data source types, the setting, **Use Integration Services For Transfer** can be found. This determines the approach that is used to transfer data between a source and a destination table. There are three settings for this:

- **As Parent** - The setting will be taken from the project setting 'Use Integration Ser-vices'.
- **Yes** - A SQL Server Integration Services (SSIS) package is used. This requires that the SQL Server component Integration Services is installed on the machine that deploys and executes the tables.
- **No** - ADO .NET is used. This does not require any SQL components except SQL Server Management Objects.

Using SSIS packages for transferring data is generally considered to be faster for trans-ferring large amounts of data since SSIS packages are optimized for this.

When transferring data between tables with fewer records, using ADO .NET can sometimes be faster, as it takes time to load the SSIS packages from the SQL Server where they are stored before data transfer can begin.

It takes significantly longer time to deploy SSIS packages than it does to deploy ADO .NET transfer. This setting can have a great impact on the overall deployment time of a project.

SSIS and ADO .NET use different technologies to transfer data. This means that if you get erroneous data through using SSIS, you can sometimes get correct data using ADO .NET.

ADO .NET can be used on trial installations, proof of concepts, or other installations where SSIS is not available.

GUARDING A DATA SOURCE

You can guard a data source in Jet Data Manager, which prevents tables that get their data from the data source from being deployed, executed or both. In general, the guard feature is useful for tables that contain data that never changes, e.g. from a legacy system. You can also guard a single data source - see Guarding a Table

To guard a data source, follow the steps below.

1. On the **Data** tab, right click the data source, click **Advanced** and click **Guard**. The Guard window opens.
2. Select **Guard on deployment** to prevent Jet Data Manager from deploying the table and/or **Guard on execution** to prevent Jet Data Manager from executing the table.
3. Click **OK**.

CHANGING DATA SOURCE PROVIDERS

If you have moved your data sources to a new database, you have to change the provider specified for the data source. You can change provider from Microsoft SQL Server to Oracle, and vice versa.

1. On the **Data** tab, expand **Business Units**, then expand the preferred business unit.
2. Expand **Data Sources**, then select the data source that you want to change the provider for.
3. Right-click the data source, and select **Change Provider**, then select the preferred database.

TO CHANGE THE PROVIDER TO MICROSOFT SQL SERVER

1. In the **Name** field, type a name for the data source. The name cannot exceed 15 characters in length.
2. In the **Server name**field, enter the location of the server.
3. In the **Database** field, enter the name of the database.

4. Specify the authentication mode. The default setting is **Windows authentication**. If you choose **SQL Server authentication**, you are prompted for a user name and a password.
5. Click **Test Connection** to verify that the connection is working, then click **OK**. The data source is added to the Data Sources folder in the project tree.
6. In the **Command Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the database.
7. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server.
8. If you want to add additional connection strings, click the **Additional Connection Properties** button. In the **Connection String Properties** window, type the preferred connection strings, and click **OK**.

TO CHANGE THE PROVIDER TO ORACLE

1. In the **TNS alias** field, type the alias that identifies the database.
2. In the **Owner** list, select the owner of the database.
3. Specify the authentication mode. When you select **Oracle authentication**, you are prompted for a user name and a password.
4. Select **Convert out of range dates to MS SQL min/max** if you want to convert all dates older than January 01, 1753 to 01-01-1753.
5. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server.
6. In the **Command Timeout field**, specify the number of seconds to wait before terminating the attempt to connect to the database.
7. If you want to set the character encoding to either Unicode or Non-Unicode, select **Force Character Setting**, and then select the preferred character encoding in the list.

   Note: Forcing character encoding may affect performance.

8. If you want to add additional connection strings, click the **Additional Connection Properties** button. In the **Connection String Properties** window, type the preferred connection strings, and then click **OK**.

TEMPLATE DATA SOURCES

Template Data Sources allows you to easily connect multiple data sources that have an identical data structure. After the data sources have been connected, the user can add fields or tables. These changes will be incorporated across all of the data sources that have been enabled as a template data source.

## ENABLING TEMPLATE DATA SOURCES

Follow the steps below to enable Template Data Sources.

1. On the **Data** tab, expand **Business Units**, and then expand the preferred business unit. Right-click **Data Sources** and add the data source to be used as a template. For this example, a SQL data source will be added.


Note: Since all of the data sources will have a similar data structure, it does not matter which one is added first.

2. Configure the data source as needed.
3. Right-click the data source and click **Add ... Data Source**, where "..." is the type of data source you added in step 1.
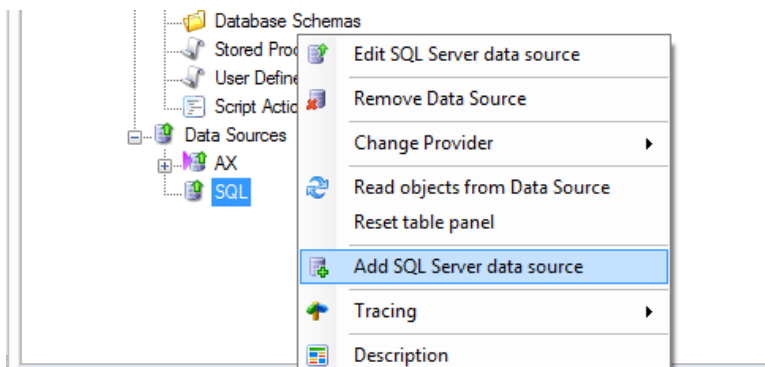
4. Configure this second data source with the needed parameters. This additional data source will then appear under the original data source in an **Additional Connections** folder.
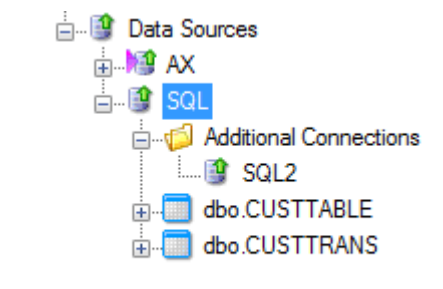
When tables and fields are added to the original data source, these changes will automatically be propagated to all of the data sources that are configured under the **Additional Connections** folder.

Below is the result of adding the table and fields shown above to the SQL data source with the records from the SQL2 data source being brought in automatically.

Table: SQL_dbo_CUSTTABLE

ACCOUNTNUM	NAME	CITY	COUNTY	COUNTRYREGION	DW_Id	DW_Batch	DW_SourceCode
902301	Birch Company	Detroit	LAMAR	US	66	1	SQL
902302	Dolphin Wholesal...	Birnamwood	SHAWANO	US	67	1	SQL
9100	Contoso Europe	Berlin		DE	68	1	SQL
Contoso	Contoso Standar...				69	1	SQL
1101	Forest Wholesales	Bothell	SNOHOMISH	US	70	1	SQL2
1102	Sunset Wholesales	Artesia Wells	LA SALLE	US	71	1	SQL2
1103	Cave Wholesales	Abbeville	WILCOX	US	72	1	SQL2
1104	Desert Wholesales	Washington	DISTRICT O	US	73	1	SQL2
1201	Sparrow Wholes...	Arvada	JEFFERSON	US	74	1	SQL2
1202	Owl Wholesales	Los Angeles	LOS ANGELE	US	75	1	SQL2
1203	Pelican Wholesal...	Burlington	SKAGIT	US	76	1	SQL2
1204	Kingbird Wholesa...	Guangzhou		CN	77	1	SQL2
1301	Whale Wholesales	Springfield	HAMPDEN	US	78	1	SQL2
1302	Turtle Wholesales	Aberdeen	HARFORD	US	79	1	SQL2

QUERY TABLES

When you connect to a data source in Jet Data Manager, you can simply use the read objects feature to list the content of the source and pick the tables and fields you wish to use in your solution. However, some AnySource providers allows you to connect to a data source, but cannot list the contents of it. In other cases, it is simply useful to be able to create a table, that does not already exist on a data source, from a query. To enable you to use these data sources in Jet Data Manager, we have included the Query Tables feature. While the SQL behind ordinary tables is created by Jet Data Manager, you write the query that brings query tables to life.

The following data sources and adapters support the query tables feature:

- AnySource adapter
- Dynamics AX adapter
- Dynamics NAV adapter
- Oracle data source
- SQL data source

ADDING QUERY TABLES

To add a query table, follow the steps below.

1. Connect to a data source using one of the supported data sources or adapters..

2. Right click the data source, click **Advanced** and click **Query Table Setup**. The **Query table Setup** window opens.
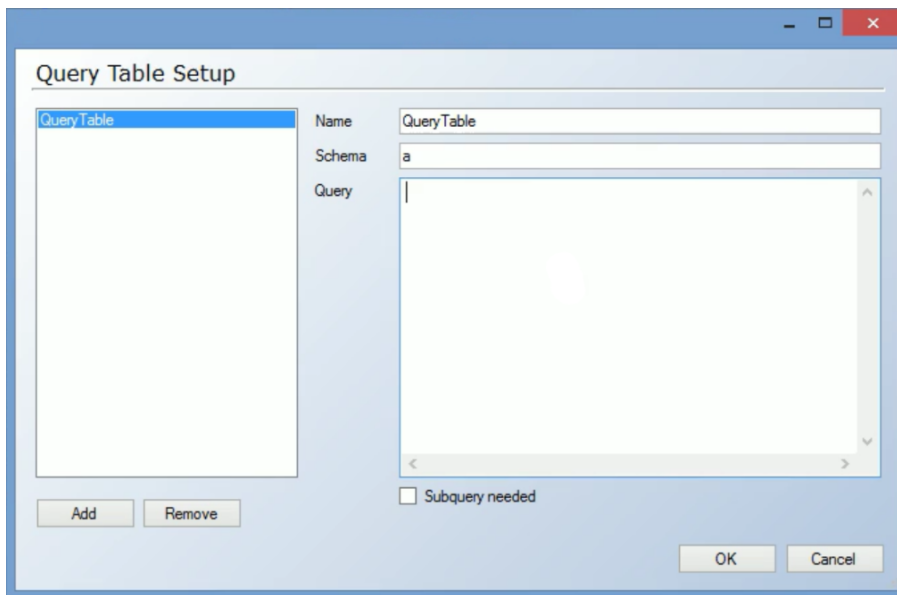
3. Click **Add**. A new table is added to the list.
4. In the **Name** box, type a name for the table.
5. (Optional) In the **Schema** box, type a schema to use.
6. In **Query**, enter the query you wish to use for creating the table. The query should contain a SELECT statement and follow the syntax required by the source.
7. Select **Subquery needed** if you are using an alias in your query. Otherwise, selection rules will fail.

8. Repeat step 3-6 to add the tables you need and click **OK**.
9. Right click the data source and click **Read Objects from Data Source**. The tables are listed in the panel in the right-hand side of Jet Data Manager and can be included in the project like any other table.

HANDLING ACCOUNTS IN DYNAMICS NAV

When you create query tables for Dynamics NAV, you will have to consider how you handle accounts.

To get data from one account, remember the account in the FROM part of your state-ment:

```
SELECT * FROM [dbo].[MyCompany$MyTable]
```

To get data from multiple accounts, in the same way the Dynamics NAV Adapter does it, you can use placeholders:

```
SELECT * {0} FROM [dbo].[{1}$MyTable]
```

Jet Data Manager will replace the digits in curly brackets during execution to create the following statement for each account:

```
SELECT
* ,CAST('MyCompany' AS nvarchar(30)) AS (DW_Account)
FROM
[dbo].[MyCompany$MyTable]
```

SQL SERVER DATA SOURCE

Jet Data Manager supports all versions of Microsoft SQL Server as a source.

ADDING A SQL SERVER DATA SOURCE

To add a new SQL Server data source, follow the steps below:

1. On the **Data** tab, expand the preferred business unit, then right-click Data Sources.
2. Select **Data Sources**, then select **Add SQL Server Data Source**.

3. In the **Name** field, type a name for the data source. The name cannot exceed 15 characters in length.
4. In the **Server Name** field, enter the location of the database server.
5. In the **Database** field, enter the name of the database, or select it from the drop-down list.
6. Specify the authentication mode. The default setting is **Windows authentication**. If you choose **SQLServer authentication**, you are prompted for a user name and a password.
7. Click **Test Connection** to verify that the connection is working, and then click OK. The data source is added to the Data Sources folder in the project tree.

8. In the **Command Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the database. **Note: The recommended setting for this is 0 seconds to disable the timeout.**
9. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server.
10. If you want to add additional connection strings, click the **Additional Connection Properties** button. In the **Connection String Properties** window, type the preferred connection strings, and then click **OK**.

DB2 DATA SOURCE

Jet Data Manager can extract data from IBM DB2 databases.

ADDING A DB2 DATA SOURCE

1. In the project tree, expand **Business Units**, expand the relevant business unit, right click **Data Sources**, click **Data Sources** and click **Add DB2 Data Source**. The

Add DB2 Data Source window opens.
2. Enter the connection information and click **OK**.

INFORMIX DATA SOURCE

Jet Data Manager can extract data from IBM Informix databases.

ADDING A INFORMIX DATA SOURCE

1. In the project tree, expand **Business Units**, expand the relevant business unit, right click **Data Sources**, click **Data Sources** and click **Add Informix Data Source**. The **Add Informix Data Source** window opens.
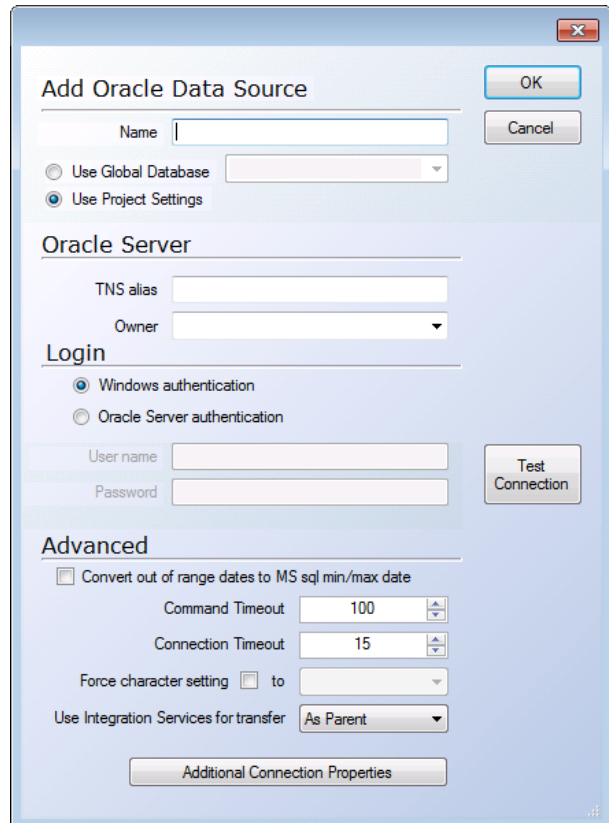

2. Enter the connection information and click **OK**.

ORACLE DATA SOURCE

Jet Data Manager can extract data from Oracle databases.

ADDING A ORACLE DATA SOURCE

1. In the project tree, expand **Business Units**, expand the relevant business unit, right click **Data Sources**, click **Data Sources** and click **Add Oracle Data Source**. The

   Add Oracle Data Source window opens.
2. Type a **Name** used to identify the data source in Jet Data Manager.
3. Type **TNS alias**, type the alias that identifies the database.
4. In the **Owner** list, click the owner of the database.
5. Under **Login**, click **Oracle Server authentication** if you wish to use this login method and enter **User name** and **password**.
6. Under **Advanced**, select **Convert out of range dates to MS SQL min/max** if you want to convert all dates older than January 01, 1753 to 01-01-1753.
7. (Optional) Enter the number of seconds to wait before terminating the attempt to connect to the database in **Command Timeout**.
8. (Optional) Enter the number of seconds to wait before terminating the attempt to connect to the server in **Connection Timeout**.
9. (Optional) Select **Force Character Setting** if you want to set the character encoding to either Unicode or Non-Unicode and click the preferred character encoding in the list.

   Note: Forcing character encoding may affect performance.

10. (Optional) In the **Use Integration Services for transfer** list, you can change the default, **As Parent**, by clicking either **Yes** or **No**.

11. (Optional) Add additional connection strings, click **Additional Connection Properties**. In the **Connection String Properties** window, type the connection strings and click **OK**.

MYSQL DATA SOURCE

Jet Data Manager supports MySQL data sources through ODBC.

ADDING A MYSQL DATA SOURCE

To add a MySQL data source, follow the steps below:

1. On the **Data** tab, expand the preferred business unit, and then right-click **Data Sources**.
2. Point to **Data Source**, select **Application specific ODBC**, and then select the preferred MySQL native database.

3. In the **Name** field, type the name of the data source.
4. In the **System DSN** list, select the Data Source Name.
5. In the **Escape Character** list, select the escape character specific to your ODBC driver.
6. The **Text Type Behavior** fields are used to control how the ODBC driver handles text. These fields are optional. You have the following options:

Option	Definition

Set Length	Specifies an exact text string length
Set Variable Length	True, if you want a variable text string length
Set Unicode	True, if you want to use Unicode

7. In **Set Number of Decimals**, specify a fixed number of decimals. This field is optional.
8. Select **Convert out of range dates to MS SQL min/max** if you want to convert all dates older than January 01, 1753 to 01-01-1753.
9. Select **Use Low Compatibility Mode** if you have trouble retrieving data from the database.
10. In the **Command Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the database.
11. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server.
12. If you want to add additional connection strings, click the **Additional Connection Properties** button. In the **Connection String Properties** window, type the preferred connection strings, and then click **OK**.
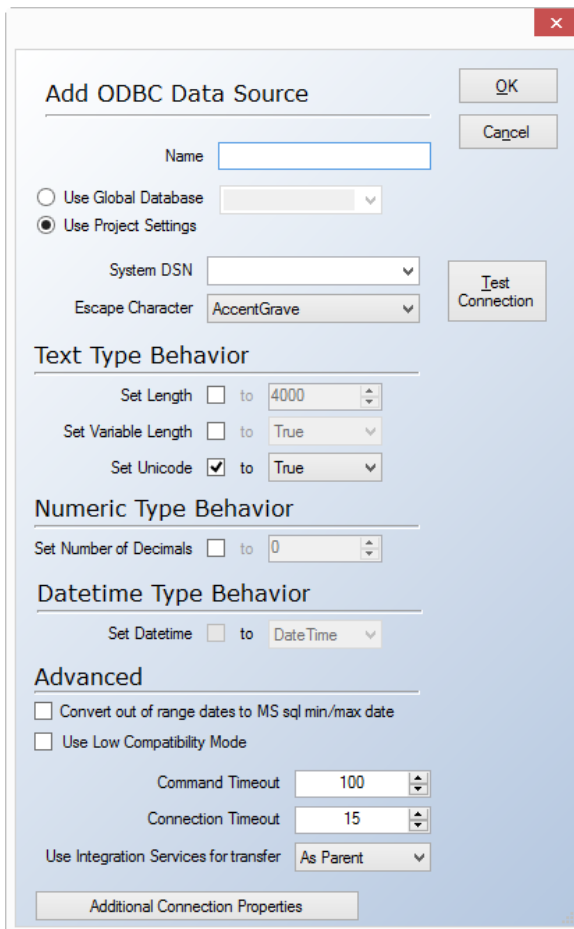
ODBC DATA SOURCES

Jet Data Manager supports ODBC for retrieving data from a wide range of databases.

ADDING AN ODBC DATA SOURCE

To add a new ODBC data source, follow the steps below:

1. On the **Data** tab, expand the preferred business unit, and then right-click **Data Sources**.
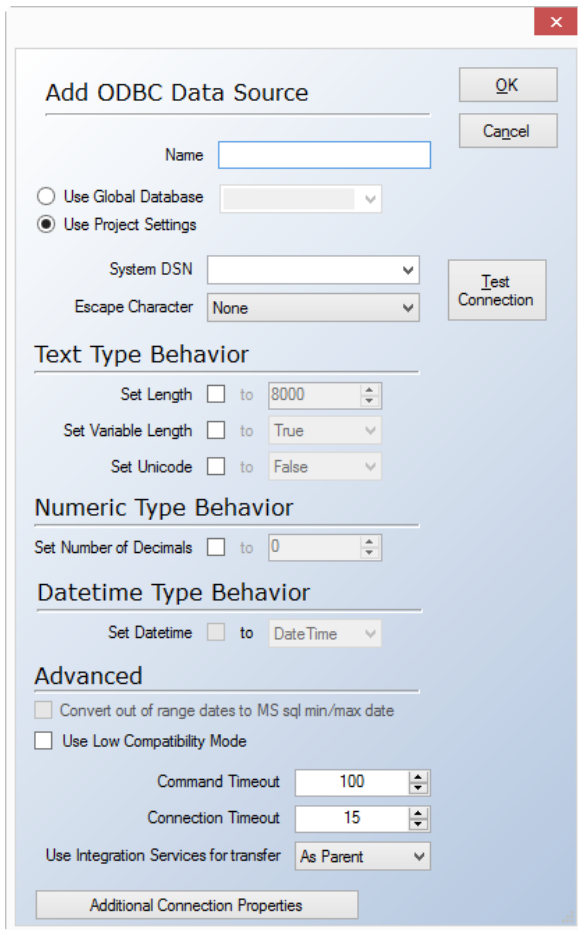
2. Point to **Data Source**, select **Application specific ODBC**, and then select the preferred database.
3. In the **Name** field, type the name of the data source.
4. In the **System DSN** list, select the Data Source Name.
5. In the **Escape Character** list, select the escape character specific to your ODBC driver.
6. The **Text Type Behavior** fields are used to control how the ODBC driver handles text. These fields are optional. You have the following options:

Option	Definition

Set Length	Specifies an exact text string length
Set Variable Length	True, if you want a variable text string length
Set Unicode	True, if you want to use Unicode

7. In **Set Number of Decimals**, specify a fixed number of decimals. This field is optional.
   Note: The **Convert out of range dates to MS SQL/min max** is not available for Navision native databases.

8. Select **Use low compatibility mode** if you have trouble retrieving data from the database.

9. In the **Command Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the database.

10. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server.

11. If you want to add additional connection strings, click the **Additional Connection Properties** button. In the **Connection String Properties** window, type the preferred connection strings, and then click **OK**.
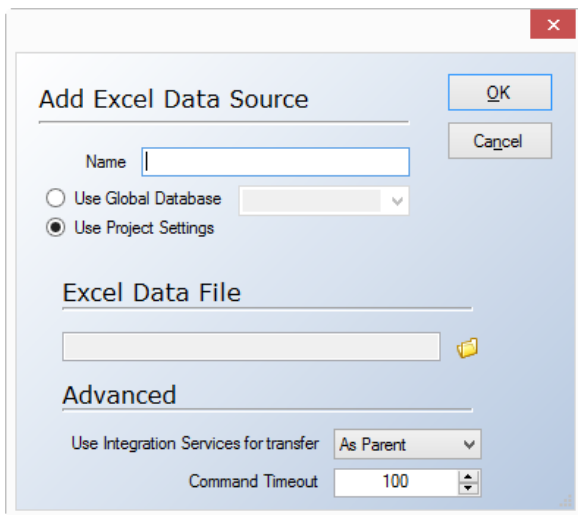
EXCEL DATA SOURCE

TX supports Microsoft Excel spreadsheets as a data source, both the older .xls format and the newer .xlsx format from Excel 2007 and beyond.

To use Excel files as a data source, you must ensure that the worksheet data is in list format. This means that the data must be set up in a database format consisting of one or more named columns. The first row in each column must have a label, and there can be no blank columns or rows within the list.

ADDING AN EXCEL DATA SOURCE

To add a new Excel data source, follow the steps below:

1. On the **Data** tab, expand the preferred business unit, then right-click **Data Sources**.
2. Select **Data Sources**, and then select **Add Excel Data Source**. The **Add Excel Data Source** window appears.

3. Click the folder icon under **Excel Data File** to to display the **Find Files or Folders** window. Navigate to and select the Excel file you wish to use as a data source, and then click **OK**.

TEXT FILE DATA SOURCE

One of the simplest ways to transfer data between systems is through the use of text files. For instance, some industrial equipment produce log files in text format that can be used as a data source in Jet Data Manager to include productivity data in the data warehouse.

One text file data source corresponds to one table in Jet Data Manager . You can load a single text file or multiple identically formatted text files.

ADDING A TEXT FILE DATA SOURCE

While virtually identical, you have to choose between the **Single Text File** and **Multiple Text File** data sources when you set up a text file data source. The **Multiple Text File** data source is usually preferable even if you only have one text file since it enables you to add more text files as sources should you need it sometime in the future. To add a text file data source, follow the steps below:

1. Expand **Business Units**, and then expand the preferred business unit.
2. Right-click **Data Sources**, select **Data Source**, and then select **Add Multiple Text File data source**
   -OR-
   Right-click **Data Sources**, select **Data Source**, and then select **Add Single Text File data source**

The **Add Multiple Text File** or **Add Single Text File** window appears.

3. In the **Name** box, type a name for the data source.
4. In the **Table Name** box, type the name of the table that is created in the staging data-base. The table is prefixed with the data source name. If you wish to set your own pre-fix, clear **Auto Prefix Tables** and type the prefix you wish to use in **Manual Table Prefix** box.
5. (Optional) Click you prefered option in the **Transfer Failure Option** list to change the **Allow Failing Data Source** setting.
6. In the **Format** list, click the format of the text file, e.g. how Jet Data Manager should make sense of the content of the file.
   - Select **Delimited** if rows and fields are seperated by a character and click the relevant characters in the **Header Row delimiter**, **Row delimiter** and **Field delimiter** lists.

- Select **FixedWidth** if the fields have a fixed length and type the lengths in **Field lenghts** in a semicolon-seperated format, e.g. "2;4;8;3".
- Select **RaggedRight** if the last fields is delimited by a character, while the previous fields are fixed width. Click the relevant characters in the **Header Row delimiter**, **Row delimiter** lists and type the lengths in **Field lenghts** in a semicolon-seperated format, e.g. "2;4;8;3".

7. Select **Field names in first data row** if the first row of data contains field name, i.e. not data.

8. Type a **Text Qualifier**, often a quatation mark, if you would like Jet Data Manager to strip from the fields before loading data into the staging data base.

9. If you are adding a multiple text files data source, enter the path to the files you wish to process separated by semicolon (;) in the **File** box. You can also use wildcards. Use "*" for any number of characters and "?" for a single character. You can also click the folder icon next to the **File** box to choose the file to process.
-OR-
If you are adding a single text file data source, click the folder icon next to the **File** box to choose the file to process.

10. In the **Culture** list, click the language of the text file.

11. Select **Unicode** if Jet Data Manager should treat your file as Unicode.

12. In the **Post processing** list, click the action you wish Jet Data Manager to perform when the file has been processes.
- Select **Backup** to move the file to a backup folder and click the folder icon next to the **Backup folder** field to select the folder.
- Select **Delete** to delete the file.
- Select **None** to leave the file as it is.

13. In the **Use Integration Services for transfer** list, you can click Yes or No to change the setting from the default As Parent.

14. Click the **Columns** tab and click **Get Fields** to load the fields, which will then be displayed in a list in the left-hand side of the window. You can select one or more fields in the list and adjust different settings for them:
- **Column name**
- **Data type**
- **Text length**: Enter the maximum number of characters in the field.
- **Variable length**: Select if you do not want the field to have a fixed length.
- **Unicode**: Select to convert data to Unicode
- **Number of decimals**: Enter the maximum number of decimals allowed in the field.

15. Click **Update** to show a preview of the data as Jet Data Manager understands it with the settings you have chosen. You have the option of adjusting the **Number of rows** to see more or less rows.

16. Click **OK** to add the data source.

ANYSOURCE DATA SOURCE

With the AnySource Adapter, you can connect to any data source that you have an OLE DB or ADO provider installed for on your server. Instead of waiting on us to create an adapter that enables you to connect to a specific system, you can now acquire a provider from the system developer or a third party vendor.

ADDING A ANYSOURCE OLE DB OR ADO DATA SOURCE

The setup for the AnySource adapter is similar for the OLE DB and ADO based adapter. To add an AnySource OLE DB or ADO data source, follow the steps below.

1. On the **Data** tab, right click **Data Source**, click **Data Sources** and click **Add AnySource OLE DB** -OR- **Add AnySource ADO**. The **AnySource OLE DB**-OR-**AnySource ADO** window opens.
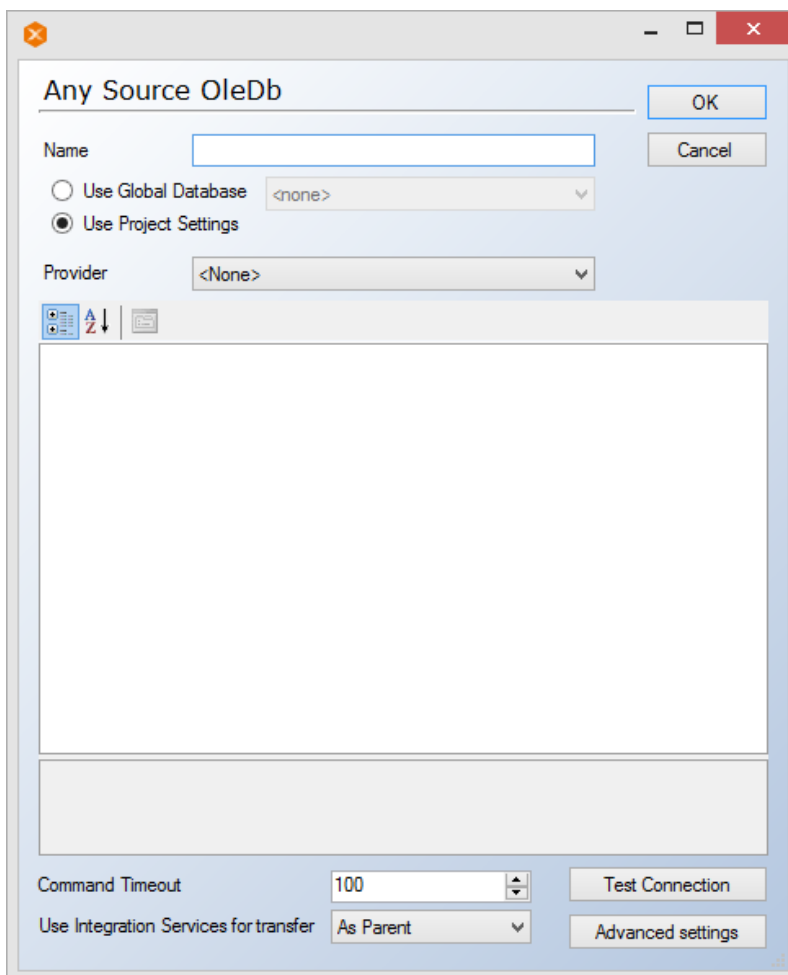
2. In the **Name** box, type a name for the data source.
3. In the **Provider** list, click the provider you wish to use.
4. In the property sheet, edit the settings for the connection. See the documentation for the provider you are using for more information.

5. (Optional) In **Command Timeout**, enter the number of seconds after which a command should time out.
6. (Optional) In the **Use Integration Services for Transfer** list, click **As Parent** to use the same setting as the parent business unit or **Yes** or **No** to use or not use Integration Services for transfer, respectively.
7. Click **Test Connection** to verify that the connection is working.
8. Click **Advanced settings** to access additional settings. The **Advanced Data Source Properties** window opens.
9. In the **Query Formatting** list, type the **Prefix**, **Suffix** and **Separator** used in the source. Click **Read Value** to fill in the values automatically if possible.
10. In the **Character Replacements** list, you can type a **Replace Character** to re-place with the **Replace Value** in data from the data source.
11. In the **Schema Properties** list, type the **Schema Name**, **Table Name**, **Column Name** etc.
12. In the **Object Filtering** list, you can choose to filter the tables you receive from the provider using regular expressions or different string comparisons. Click **Table** or **Schema** in the **Object Type** list, click a filter type in the **Filter Type** list and type a value in the **Filter Value** List.
13. Click **OK** to close the **Advanced Data Source Properties** window.
14. Click **OK** to close the **AnySource OLE DB** -OR- **AnySource ADO** window and add the data source.

Note: On other data sources, you can use the preview feature in Jet Data Manager to view the content of a table on the source. This might not work on all AnySource data sources since the Jet Data Manager does not know exactly what the source is and what syntax to use. However, you can use the query tool to explore the content of a table. See Query Tool for more information.

CUSTOM DATA SOURCE

The Custom Data Source works in conjunction with a separate provider - or driver - to enable access to data sources that are not supported by the core Jet Data Manager product.

ADDING A CUSTOM DATA SOURCE

The setup for a Custom Data Source depends on the provider you are using. While the general steps to add a Custom Data Source is explained below, you should consult the documentation for the provider to learn more about the specific settings.

1. On the **Data** tab, right click **Data Source**, click **Data Sources** and click **Add Custom Data Source** The **Add Custom Data Source** window appears.
2. In the **Name** box, type a name for the data source.
3. In the **Provider** list, click the provider you wish to use.
4. In the **Setup Property** list, click the property you wish to edit. Go through all properties in the list and edit the settings for the connection in the property list below.
5. Click **Test Connection** to verify that the connection is working.
6. Click **OK** to close the window and add the data source.

CONNECTING WITH APPLICATION ADAPTERS

Jet Data Manager supports two different approaches for connecting to source systems: simple data sources and intelligent application adapters.

A data source connection simply connects to the source and enables you to browse the content of the source.

An adapter is a component that enables you to easily extract and synchronize data from different source systems. The adapter knows how a given system organizes and stores data, which enables the adapter to simplify the table structure you see in Jet Data Manager. For instance, data for each company in a Dynamics NAV system is stored in a separate set of tables. The Dynamics NAV adapter merges these tables together and lets you select companies on a global level.

Jet Data Manager includes Application Adapters for the following systems:

- [Microsoft Dynamics NAV](#)
- [Microsoft Dynamics AX](#)
- [Microsoft Dynamics GP](#)
- [Microsoft Dynamics CRM](#)
- [SAP](#)
- [Sun Systems](#)
- [Salesforce](#)
- [Infor Movex/M3](#)
- [UNIT4 Agresso](#)

MICROSOFT DYNAMICS AX ADAPTER

This adapter simplifies the extraction of data from Microsoft Dynamics AX.

If you connect to a Dynamics AX database as a regular data source, you will have to apply and maintain selection rules on all tables. With Jet Data ManagerDynamics AX adapter, you can select company accounts at a global level. You can, however, override this behavior on a table by table basis.

The adapter also extracts any virtual company accounts, including, table collections, and tables that are set up in the source database. The information can then be used in dimensions and cubes.

Furthermore, the adapter extracts all Base Enumerations and their associated labels and supports synchronization with the back-end application.

IMPORTING XPO FILES INTO DYNAMICS AX

The Dynamics AX adapter is only available if the .xpo file has been imported into Dynamics AX.

1. Import the .xpo file into Dynamics AX.
2. Compile the imported project within Dynamics AX.
3. Run all four classes in Dynamics AX to populate the tables.
4. Add a Dynamics AX adapter to your Jet Data Manager project. For more information, see Add Dynamics AX Adapters.

For detailed instructions on how to import files, compile projects, and run classes in Dynamics AX, see the Installation documentation on the JetReports website.

ADDING A DYNAMICS AX ADAPTER

Use the Dynamics AX Adapter to load data from separate Dynamics AX company accounts tables in a single table.

1. On the **Data** tab, expand the preferred business unit, and then right-click **Data Sources**.
2. Select **Add Adapter Data Sources**, and then select **Add Dynamics AX Adapter**.
3. Enter a name for the adapter, and then click **OK**.

You can now choose the provider which contains the data source you want to connect to.

ADDING A MICROSOFT SQL SERVER PROVIDER

1. Right-click the adapter, and select **Source Providers**. Then select **Add MS SQL Provider**.
2. In the **Server Name** field, enter the location of the server.
3. In the **Database** field, enter the name of the database.

4. Specify the authentication mode. The default setting is **Windows authentication**. If you choose **SQL Server authentication**, you are prompted for a user name and a password.
5. In the **Command Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the database.
6. In the Connection Timeout field, specify the number of seconds to wait before terminating the attempt to connect to the server, and then click OK.

ADDING AN ORACLE PROVIDER

1. Right-click the adapter, and select **Source Providers**. Then select **Oracle Provider**.
2. In the **TNS alias** field, type the alias that identifies the database.
3. In the Owner list, select the owner of the database.
4. Specify the authentication mode. When you select **Oracle authentication,** you are prompted for a user name and a password.
5. Select **Convert Out of Range Dates to MS SQL min/max** if you want to convert all dates older than January 01, 1753 to 01-01-1753.
6. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server.
7. In the **Command Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the database.
8. If you want to set the character encoding to either Unicode or Non-Unicode, select **Force Character Setting**. Then select the preferred character encoding in the list.

> **Note:** Forcing character encoding may affect performance.

If you want to add additional connection strings, click the **Additional Connection Properties** button. In the Connection String Properties window, type the preferred connection strings, and then click **OK**.

SETTING THE ACCOUNT TABLE

Before you can continue with setting up accounts, you have to verify that the company account table is correct.

1. Right-click the adapter, and then choose **Set Account Table**. The table DATAAREA and the field ID are selected by default.
2. Click **OK**.

SETTING UP DYNAMICS AX COMPANIES

When you have added a Dynamics AX adapter and specified a provider, you need to set up the accounts. To set up accounts:

1. Right-click the adapter, and select **Set Up Accounts**. An information message is displayed, which lists the accounts that have been added. Click **OK**. The **Accounts**

window appears.
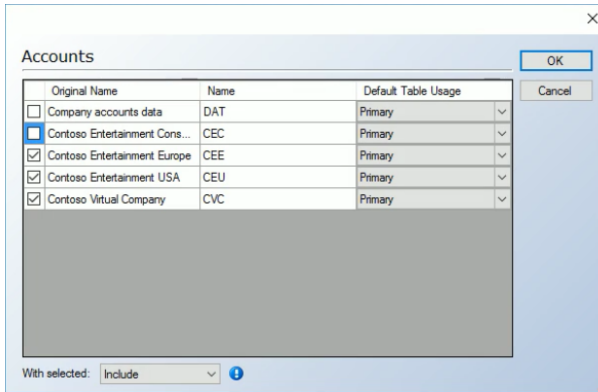
2. Select the accounts from which you want to retrieve data. In the **With selected**list, click **Include** to use the selected accounts in the data warehouse or click **Exclude** to exclude the selected accounts and include all other accounts. The last option is useful if you often add new accounts in Dynamics AX and want to make sure that all accounts are included in the data warehouse as soon as they exist in the ERP system.

3. In the **Default Table Usage** list, specify the order in which data from the tables is retrieved and read. You have the following options:

Option	Definition
Primary	Data from this company account is read and retrieved first
Secondary	Data from this company account is read and retrieved after the primary account if they have not already been retrieved from the primary account
None	Tables from this company are not retrieved unless you specify at the table level that you want to retrieve data from a specific table. For more information, see To modify table usage in Dynamics AX tables.

4. Click **OK**.

LOADING AND SELECTING DATA FROM DYNAMICS AX DATA SOURCES

1. On the **Data** tab, expand **Business Units**, expand the preferred business units, and then expand **Data Sources**.
2. Right-click the Dynamics AX adapter you want to select data from, and then select **Read objects**. The Tables pane displays all tables and fields.
3. In the **Tables**, select the tables and fields you want to extract to your staging database.
4. There are two ways of viewing the data: Alphabetical view, which displays all tables alphabetically, and Group view, where you specify how many tables each group must contain. To view data in groups, enter the number of tables in each group in the Group view field, and then click **Group view**. You can then group the tables alpha-

betically or by specifying the number of tables you want in each group. To view data alphabetically, click **Alphabetical view**.

The tables and fields are displayed in the data source tree and in the staging database tree.

ADDING DYNAMICS AX VIRTUAL TABLE REFERENCES

1. On the **Data** tab, expand **Business Units**, expand the preferred business units, and then expand **Data Sources.**
2. Expand the AX adapter that contains the table to which you want to add a virtual table reference, right-click the table, and then choose **Add Virtual Table Reference**.
3. In the Add Virtual Table Reference window, select the preferred virtual tables, and then click **OK**.

VIEWING DYNAMICS AX TABLE INFORMATION

Jet Data Manager can retrieve table information directly from your Dynamics AX database.

1. On the **Data** tab, expand **Business Units**, expand the preferred business units, and then expand **Data Sources**.
2. Expand the Dynamics AX adapter that contains the table you want to view inform-ation about, right- click the table, and then select **View Table Information**.
   The three tabs in the View Table Information dialog contain the following information:

Fields Tab	Description
Name	Specifies the name of the field as it appears in the database
Label	Specifies the name of the field as it appears in the user interface
Help Text	Contains the help text for the field
EDT Name	Specifies the name of the extended Data Type if applicable
Enum Name	Specifies the name of the enumeration if applicable
System	Specifies whether the table is a system table or visible in the user inter-face

Relations Tab	Description
External Table	Specifies the name of the table the selected table is related to
Directions	Specifies whether the selected table is the child or the parent in the relation
Field	Specifies which field in the selected table that relates to a field in the related table

External Field	Specifies the field on the related table
Relation Type	Specifies the type of relation. **Field** specifies relation fields without conditions. **ThisFixed** specifies relation fields to restrict the records in the primary table. **ExternFixed** specifies relation fields that restrict the records in the related table

Virtual Company References	Description
Company	The name of the company account
Virtual Company	The name of the Virtual Company that contains tables shared by several company accounts

VIEWING DYNAMICS AX ENUM TABLE INFORMATION

All Enum values in Dynamics AX are represented as integers in the tables. However, you can see the corresponding literal values by viewing the enumeration table information.

1. On the Data tab, expand Business Units, expand the preferred business units, and then expand Data Sources.
2. Expand the Dynamics AX adapter that contains the table you want to view information about, right- click the table, and then select Preview Enum Table.

CHANGING DYNAMICS AX SCHEMAS

1. On the **Data** tab, expand **Business Units**, and then expand the preferred business unit.
2. Expand **Data Sources**, right-click the AX adapter that contains the table whose priority you want to change, and then select **Change Schema**.
3. In the **Select Schema To Change list**, select the schema you want to change.
4. In the **New Schema Name** field, enter a name for the schema.

MODIFYING TABLE USAGE ON DYNAMICS AX TABLES

When you set up accounts, you specify the default order in which data is retrieved from the individual accounts. However, it is possible to specify a different order of priority for individual tables.

1. On the **Data** tab, expand **Business Units**, and then expand the preferred business unit.
2. Expand **Data Sources**, right-click the AX adapter that contains the table whose priority you want to change, and then select **Modify Table Usage**. The company accounts and the usage of all tables will be displayed.
3. Right-click the table and account field with the setting you wish to change the order of

priority on for data retrieval. You have the following options:

Priority	Definition
Default	Data from this table is read and retrieved first
Primary	Data from this table is read and retrieved first
Secondary	Data from this table are read and retrieved after the primary table if they have not already been retrieved from the primary table
None	Data from this table is not retrieved
1-9	Specify the order priority in the range from 1-9
Enter priority	If the order of priority exceeds the numbers 1-9, you can specify additional numbers

4. Click **OK**.

MODIFYING THE USAGE OF A SINGLE DYNAMICS AX TABLE

If you want to change the order of priority in which data is retrieved on a single table, you can do so from the individual table.

1. On the **Data** tab, expand **Business Units**, and then expand the preferred business unit.
2. Expand **Data Sources**, right-click the AX adapter that contains the table whose priority you want to change, and then select the preferred table.
3. Right-click the table, and select **Modify Single Table Usage**. The company accounts and the usage specified in Setup Company Accounts will be displayed.
4. Right-click the field that contains the setting for the table, and then specify the table usage. You have the following options:

Priority	Definition
Default	Data is retrieved based on the settings specified when you set up the company accounts
Primary	Data from this table is read and retrieved first
Secondary	Data from this table is read and retrieved after the primary table if they have not already been retrieved from the primary table
None	Data from this table is not retrieved
1-9	Specify the order of priority in the range form 1-9
Enter priority	If the order of priority exceeds the numbers 1-9, you can specify additional numbers here.

MICROSOFT DYNAMICS NAV ADAPTER

This adapter simplifies the extraction of data from Microsoft Dynamics NAV.

 If you connect to a Dynamics NAV database as a regular data source, you will have to apply and maintain selection rules on all tables because different companies are stored in separate tables. With Dynamics NAV adapter, you can select company accounts at a global level and apply only one set of selection rules. It is, however, also possible to over-rule this behavior on a table by table basis.

TO ADD DYNAMICS NAV ADAPTERS

Use the Dynamics NAV Adapter to load data from separate Dynamics company account tables in a single table.

1. On the **Data** tab, expand the preferred business unit, and then right-click **Data Sources**.
2. Point to **Add Adapter Data Sources**, and then select **Add Dynamics NAV Adapter**.
3. Enter a name for the adapter. Optionally, select **Read Aggregation Tables - SIFT** if you need to include Sum Index Flow Technology (SIFT) tables, and then click **OK**.

You can now choose the provider which contains the data sources you want to connect to.

TO ADD AN MS SQL PROVIDER

1. Right-click the adapter and select **Source Providers**. Then select **Microsoft SQL Provider**.
2. In the **Server Name** field, enter the location of the server.
3. In the **Database** field, enter the name of the database.
4. Specify the authentication mode. The default setting is **Windows authentication**. If you choose **SQL Server authentication**, you are prompted for a user name and a password.
5. In the **Command Timeout** field, specify the number of seconds to wait before ter-minating the attempt to connect to the database. The recommended value is 0 to dis-able the timeout. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server.
6. If you want to add additional connection strings, click the **Additional Connection Properties** button. In the **Connection String Properties** window, type the preferred connection strings, and then click **OK**.

TO ADD NAVISION NATIVE DATABASE SOURCES

When you want to retrieve data from Navision databases hosted in a Native Navision server environment, you will have to use ODBC. The Navision ODBC driver must be installed and configured prior to adding the Native NAV data source.

Note: If you are connecting to NAV through an ODBC connection, you must be using the 32-bit version of **Jet Data Manager** as the NAV ODBC driver only supports 32-bit connections.

1. On the **Data** tab, expand the preferred business unit, and then right-click **Data Sources**.
2. Point to **Adapter Data Sources**, and select Add **Dynamics NAV Adapter**.
3. Select Wizard Setup.
4. Select Navision Native.
5. In the **Name** field, type the name of the data source.
6. In the **DSN Name**, select the ODBC connection that you have configured for the data source. In the Escape Character list, select the escape character specific to your ODBC driver. The Text Type Behavior fields are used to control how the ODBC driver handles text. These fields are optional. You have the following options:

Option	Definition
Set Length	Specifies an exact text string length
Set Variable Length	True, if you want a variable text string length
Set Unicode	True, if you want to use Unicode

7. In **Set Number of Decimals**, specify a fixed number of decimals. This field is optional.
   Note: The Convert Out of Range Dates to MS SQL/min max is not available for Navision native databases.
8. Select **Use low compatibility mode** if you have trouble retrieving data from the database.
9. In the **Command Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the database.
10. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server.
11. If you want to add additional connection strings, click the Additional Connection Properties button. In the Connection String Properties window, type the preferred connection strings, and then click **OK**.

TO CHANGE THE DYNAMICS NAV COMPANY TABLE

By default, when you add a Dynamics NAV adapter, the company account table is set to dbo.Company.

However, it is possible to change the account table. This is generally not recommended.

1. On the Data tab, expand **Business Units**, expand the preferred business unit, and then expand **Data Sources**.

2. Right-click the preferred Dynamics NAV adapter, and then select **Edit Account Table**. A message is displayed saying ,"Retrieving database structure".
3. In the **Table** list, select the account to table that you want to use.
4. In the **Name Field** list, select the field that contains the account name, and then click **OK**.

TO LOAD AND SELECT DATA FROM DYNAMICS NAV DATA SOURCES

1. On the **Data** tab, expand **Business Units**, expand the preferred business units, and then expand **Data Sources.**
2. Right-click the Dynamics NAV adapter containing the desired data, and then select **Read objects**. The Tables pane displays all tables, fields, and views.
3. In the **Tables** pane, select the tables, fields, and views you want to extract to your staging database.
4. There are two ways of viewing the data: Alphabetical view, which displays all tables alphabetically, and Group view, where you specify how many tables each group must contain.

To view data in groups, enter the number of tables in each group in the **Group View** field, and then click **Group View**. You can then group the tables alphabetically or by specifying the number of tables you want in each group. To view data alphabetically, click **Alphabetical View**.

The tables and fields are displayed in the data source tree and in the staging database tree.

TO SET UP DYNAMICS NAV COMPANIES

When you have added a Dynamics NAV adapter and specified a provider, you need to set up accounts representing the companies requiring the extracted data.

To set up accounts:

1. Right-click the adapter, and then choose **SetUp Accounts**. A dialog is displayed that shows all companies in the database.
2. In the **Template** list, select the company account you want to use as template for the table and column structure. If you are only selecting one company, then the template company must match the company that is selected.
3. Select **Use** to specify whether to retrieve data from the company.
4. In the **Default Table Usage** list, specify the order in which tables are retrieved and read. You have the following options:

Option	Definition
Primary	Data from this company account is read and retrieved first
Secondary	Data from this company account is read and retrieved after the primary

account if they have not already been retrieved from the primary account

None	Tables from this company are not retrieved, unless you specify at table level that you want to retrieve data from a specific table

TO CHANGE DYNAMICS NAV SCHEMAS

You can change the schema for the entire Dynamics NAV adapter or for individual tables that belong to the adapter.

1. On the **Data** tab, expand **Business Units**, expand preferred business unit, and then expand Data Sources.
2. Right-click the NAV adapter whose schema you want to change, and then select **Change Schema**.
   - OR -
   Expand the NAV adapter, and then select the table whose schema you want to change.
3. In the **Select Schema** to change list, select the schema you want to change, and then select **Change Schema**.
4. In the **New Schema Name** field, enter a name for the schema, and then click **OK.**

TO MODIFY TABLE USAGE ON DYNAMICS NAV TABLES

When you set up accounts, you specify the default order in which data is retrieved from the individual accounts. However, it is possible to specify a different order of priority for individual tables.

1. On the **Data** tab, expand **Business Units**, and then expand the preferred business unit.
2. Expand **Data Sources**, right-click the NAV adapter that contains the tables whose priority you want to change, and then select **Modify Table Usage**. The company accounts and the usage of all tables will be displayed.
3. Right-click the field that contains the setting for the table and the account for which you want to change priority of data retrieval. You have the following options:

Priority	Definition
Default	Data from this table is read and retrieved first
Primary	Data from this table is read and retrieved first
Secondary	Data from this table is read and retrieved after the primary table if they have not already been retrieved from the primary table
None	Data from this table is not retrieved
1-9	Specify the order of priority in the range from 1-9.
Enter priority	If the order of priority exceeds the numbers 1-9, you can specify additional numbers.

4. Click **OK**.

TO MODIFY THE USAGE OF A SINGLE DYNAMICS NAV TABLE

You can change the order in which data is retrieved from individual tables.

1. On the **Data** tab, expand **Business Units**, and then expand the preferred business unit.
2. Expand **Data Sources**, right-click the NAV adapter that contains the table whose priority you want to change, and then select the preferred table.
3. Right-click the table and select **Modify Single Table Usage**. The company accounts and the usage specified in Setup Company Accounts will be displayed.
4. Right-click the field containing the settings for the table, and then specify the table usage. You have the following options:

Priority	Definition
Default	Data is retrieved based on the settings specified when you set up the company accounts.
Primary	Data from this table is read and retrieved first
Secondary	Data from this table is read and retrieved after the primary table if they have not already been retrieved from the primary table
None	Specify the order of priority in the range from 1-9
Enter priority	If the order of priority exceeds the numbers 1-9 you can specify additional numbers here

5. Click **OK**.

MICROSOFT DYNAMICS GP ADAPTER

This adapter simplifies the extraction of data from Microsoft Dynamics GP.

If you connect to a Dynamics GP database as a regular data source, you will have to apply and maintain selection rules on all tables because different companies are stored in separate databases. WithJet Data Manager Dynamics GP Adapter, you can select company accounts at a global level and apply only one set of selection rules. It is, however, also possible to overrule this behavior on a table by table basis.

TO ADD A DYNAMICS GP ADAPTER

Use the Dynamics GP Adapter to load data from separate Dynamics company account databases in a single table.

1. On the **Data** tab, expand the preferred business unit, and then right-click **Data Sources**.
2. Point to **Add Adapter Data Sources**, and then select **Add Dynamics GPAdapter**.
3. Enter a name for the adapter.
4. In the **Server name** field, enter the location of the server where Dynamics GP resides.
5. In the **Database** field, enter the name of the database (this should be the DYNAMICS database).
6. Specify the authentication mode. The default setting is **Windows authentication**. If you choose **SQL Server authentication**, you are prompted for a user name and a password.
7. In the **Command Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the database. The recommended value is 0 to disable the timeout. In the **Connection Timeout** field, specify the number of seconds to wait before terminating the attempt to connect to the server.
8. If you want to add additional connection strings, click the **Additional Connection Properties** button. In the **Connection String Properties** window, type the preferred connection strings, and then click **OK**.
9. The next window will be the GP Company Table Setup. All settings should be left as default and click **OK.**

TO SET UP DYNAMICS GP COMPANIES

After you have added a Dynamics GP adapter and specified a provider, you will set up the accounts which represent the companies for which data will be extracted from the data source.

To set up accounts:

1. Right-click the adapter, and then choose **Read Dynamics GP Companies**. A dialog is displayed that shows all companies in the database.

2. In the **Template** list, select the company account you want to use as template for the table and column structure. If you are only selecting one company, then the template company must match the company that is selected.
3. Select **Use** to specify whether to retrieve data from the company.

TO LOAD AND SELECT DATA FROM DYNAMICS GP DATA SOURCES

1. On the **Data** tab, expand **Business Units**, expand the preferred business units, and then expand **Data Sources.**
2. Right-click the Dynamics GP Adapter from which you want to select data, and then select **ReadObjects from Data Source**. The Tables pane on the right displays all tables, fields, and views.
3. In the **Tables** pane, select the tables, fields, and views you want to extract to your staging database.
4. There are two ways of viewing the data: Alphabetical view, which displays all tables alphabetically, and Group view, where you specify how many tables each group must contain.
5. To view data in groups, enter the number of tables in each group in the **Group view** field, and then click **Group view**. You can then group the tables alphabetically or by specifying the number of tables you want in each group. To view data alphabetically, click **Alphabetical view**.
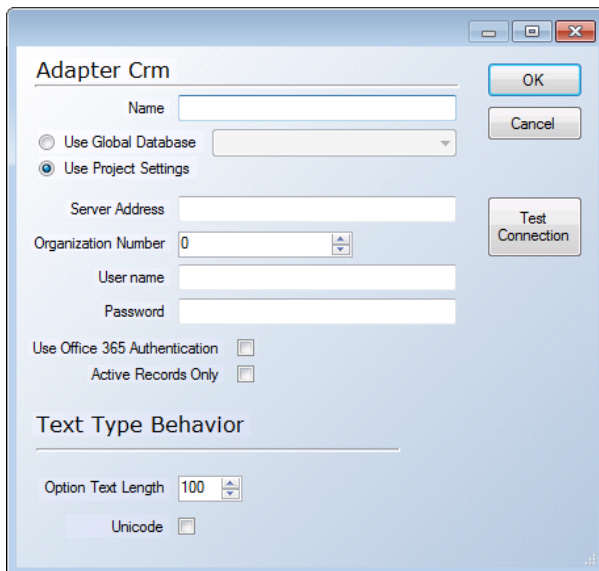
The tables and fields are displayed in the data source tree and in the staging database tree.

MICROSOFT DYNAMICS CRM ONLINE ADAPTER

This adapter simplifies extraction of data from Microsoft Dynamics CRM Online.

ADDING A DYNAMICS CRM APPLICATION ADAPTER

1. In the project tree, expand **Business Units**, expand the relevant business unit, right click **Data Sources**, click **Adapter Data Sources** and click **Add Dynamics CRM Adapter**. The **Add CRM Adapter** window opens.

2. Fill in the settings that matches your setup and click **OK**.

SALESFORCE ADAPTER

The Salesforce application adapter enables you to extract data stored in your Salesforce Sales Cloud CRM. You will need an edition of Salesforce Sales Cloud that enables the use of the Salesforce AP. At the time of writing, this means at least the Enterprise edition.

The adapter uses the Salesforce REST API to extract data. The data types of the extracted data is converted from Salesforce data types to their SQL Server equivalents.

Note: The Salesforce adapter requires .NET Framework 4.5 or higher. This is necessary because the Salesforce REST API uses the TLS 1.1 protocol, which is not supported by earlier versions of the .NET Framework. Jet Data Manager only requires .NET Framework 4.0, but if you do not have the .NET Framework 4.5 installed, Jet Data Manager will use TLS 1.0, which will result in connection errors.

ADDING A SALESFORCE ADAPTER

To connect to a Salesforce data source using the application adapter, follow the steps below.

1. On the **Data** tab, in the project tree, expand **Business Units**, expand the relevant **Business Unit**, right click **Data Source**, click **Adapter Data Sources** and click **Add Salesforce Adapter**.

   The **Add Salesforce Adapter** window opens.

2. Type the **Name** you wish to use for the adapter.
3. Under login, enter your **Username**, **Password** and **Token**. The token is provided by Salesforce and will change if the password is changed. Salesforce can be configured not to use security tokens. In that case, simply leave **Token** empty.
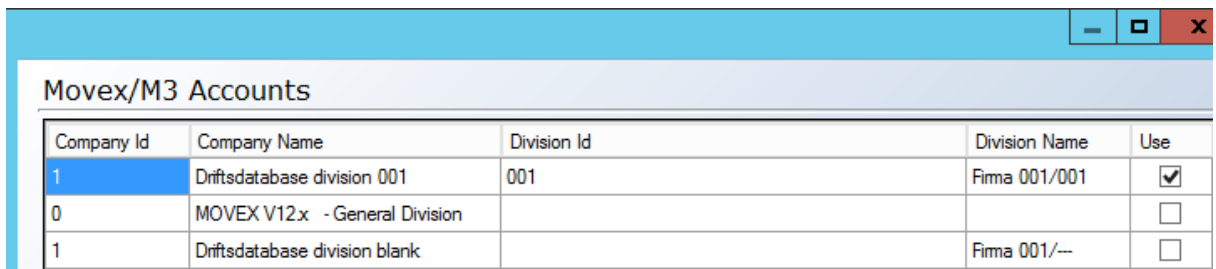4. In the **API Version** list, click the Salesforce API version you wish to use. Usually, you should use the newest version. However, if Salesforce makes changes in the API that breaks the adapter, you have the option of using an earlier version.
5. Select **Use Label as Name** to use labels as names instead of the physical system names.
6. Select **Unicode** to make string based fields Unicode ready.
7. Enter a **Custom URL** if you wish to connect to your Salesforce sandbox instead of the default URL. The **Resolved Path** box shows you the entire URL the adapter will use to connect.
8. In **Max. Concurrent HTTP Request Per Table**, type or select the maximum number of HTTP requests you wish to allow per table. For each request, the adapter fetches 1000 rows.

SAP APPLICATION ADAPTER

The SAP Application Adapter allows you to connect to SAP systems. The Adapter utilizes the **XTract IS** component from Theobald Software. Make sure that the latest version of the component is downloaded and installed on the server where Jet Data Manager is installed. The component can be downloaded from Theobald Software: [http://theobald-soft-ware.com/en/product-downloads.html](http://theobald-software.com/en/product-downloads.html)

Due to some issues in the standard SAP meta-data layer, you will need to install a function module in SAP. For more information, see [https://my.theobald-soft-ware.com/index.php?/Knowledgebase/Article/View/56/3/installing-z_xtract_is_table](https://my.theobald-software.com/index.php?/Knowledgebase/Article/View/56/3/installing-z_xtract_is_table)

The function module will make it possible to:

- Extract tables/table columns with an overall width greater than 512.
- Extract tables that contain at least one column of type F (floating point).
- Extract table TCURR which has some meta data problems in the Data Dictionary.

ADDING A SAP APPLICATION ADAPTER

1. In the project tree, expand **Business Units**, expand the relevant business unit, right click **Data Sources**, click **Adapter Data Sources** and click **Add SAP Table Adapter**. The **Add SAP Table Adapter** window opens.

2. Fill in the settings that matches your setup and click **OK**.

SUN SYSTEM ADAPTER

The Sun System adapter simplifies the extraction of data from a Sun System using a Microsoft SQL Server.

ADDING A SUN SYSTEM ADAPTER

1. In the project tree, expand **Business Units**, expand the relevant business unit, right click **Data Sources**, click **Adapter Data Sources** and click **Add Sun System Adapter**. The **Add Sun System Adapter** window opens.
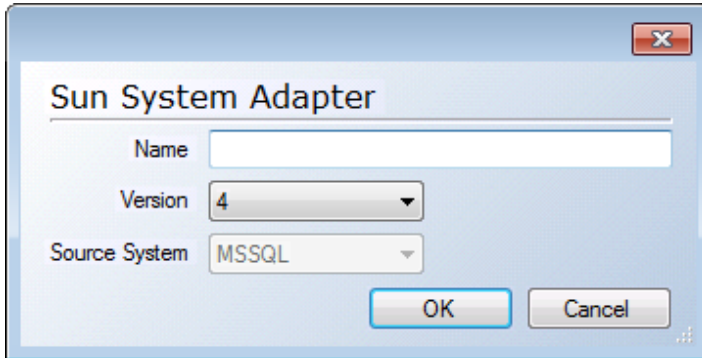
2. Type a **Name** for the adapter to easily identify it in the project tree.
3. In the **Version** list, click the version of Sun System you are connecting to.
4. Click **OK**.
5. The standard window for adding a Microsoft SQL Server data source opens. Enter the connection details for your data source - see Adding a SQL Server Data Source for more information.

MOVEX/M3 ADAPTER

The Movex/M3 adapter simplies extraction of data from Infor M3, an ERP system also known by its previous name Movex.

ADDING A MOVEX/M3 ADAPTER

1. In the project tree, expand **Business Units**, expand the relevant business unit, right click **Data Sources**, click **Adapter Data Sources** and click **Add Movex/M3 Adapter**. The **Add Movex/M3 Adapter** window opens.

2. Click the database type of your M3 installation. You have the following options:
    - DB2
    - Oracle
    - Microsoft SQL Server
3. Click **OK**.
4. The standard window for adding a data source of the type you chose opens. Enter the connection details. See Adding a DB2 Data Source, Adding a Oracle Data Source and Adding a SQL Server Data Source for more information.
5. In the project tree, right click the M3 data source you just added and click **Change System Schema**. The **Setup System Schema** window opens.
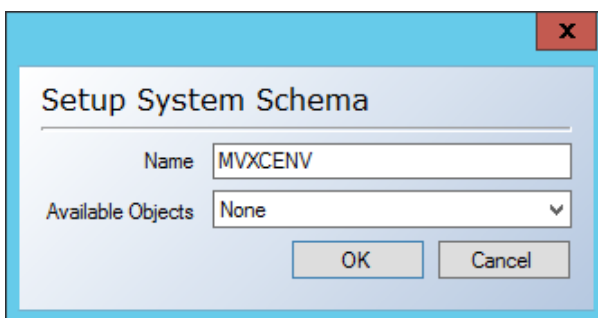
6. Type the environment schema name in the **Name** box.
7. In the **Available Objects** list, click **None** to make no tables available for use in the project, **Company Tables Only** to make only tables from the chosen company available or **All Tables and Views** to make every table and view in the M3 database available. Click**OK**.

8. In the project tree, right click **Setup Accounts**. The **M3/Movex Accounts** windows opens.

Company Id	Company Name	Division Id	Division Name	Use
1	Driftsdatabase division 001	001	Firma 001/001	☑
0	MOVEX V12.x - General Division			☐
1	Driftsdatabase division blank		Firma 001/---	☐

9. Select the companies and divisions to include by checking the box in the **Use** column. Click **OK**.

AGRESSO ADAPTER

The Agresso Adapter enables the use of the UNIT4 Business World ERP system in Jet Data Manager. Agresso is the former and more well known name of the system and as such, it is used in Jet Data Manager.

ADDING AN AGRESSO ADAPTER DATA SOURCE

Adding an Agresso Adapter consists of two major parts. Unlike other application adapters, the Agresso Adapter comes with a Jet Data Manager project based on Agressos standard dimensional setup. Downloading and configuring this project is the first part, while the second part is configuring Agresso-specific settings for client, dimensions, translations and flexi-tables.

To download and configure the Agresso project, follow the steps below.

1. Click the application button in the top left corner of Jet Data Manager and click **CubeStore**. The Cubestore window opens.
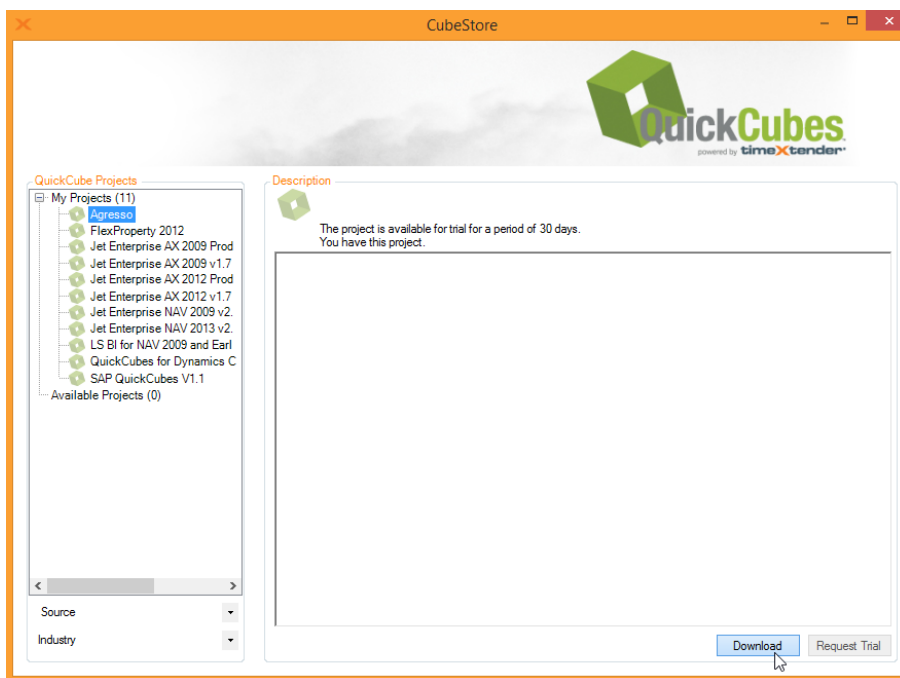
2. In the **QuickCube Projects** list, click **Agresso** and then click **Download**.

   Note: The project is only available if you have a license for Jet Data Manager that includes the Agresso Adapter.

3. One the project has been downloaded, Jet Data Manager will ask you if you wish to run the Connection Manager. Click **Yes**.

4. The Connection Manager window opens.
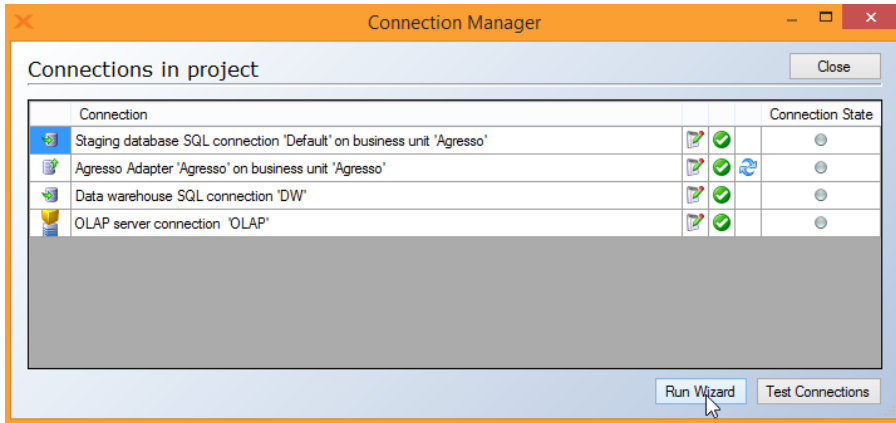
Click **Run Wizard**. Four windows will now open in tour, allowing you to quickly configure the different connections.

5. Create a staging database for use with the project and click **OK**. For more information, see Setting Up a Staging Database.
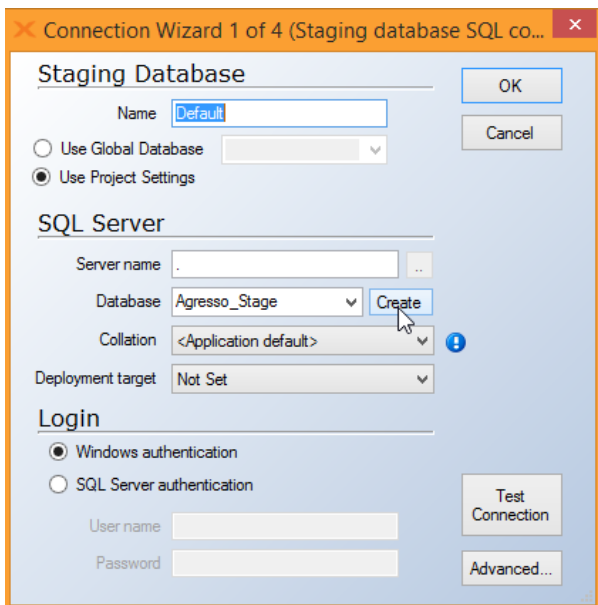
6. Enter the connection information for your Agresso database and click **OK**. For more information, see Adding a SQL Server Data Source.

7. Create a data warehouse database and click **OK**. For more information, see Adding a Data Warehouse.
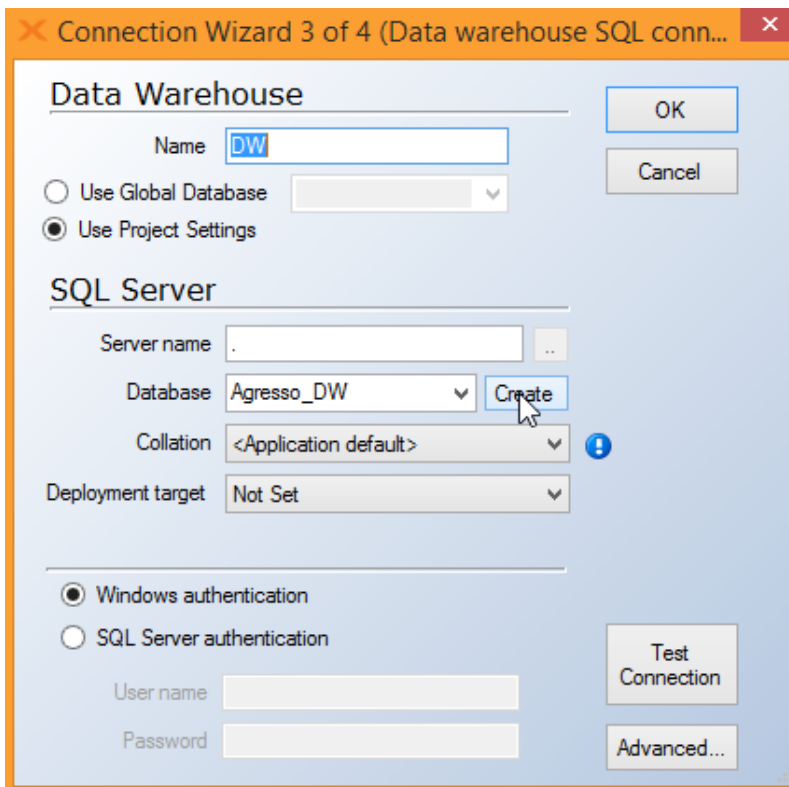
8. Set up an OLAP server for the project. For more information, see Adding an OLAP Server.
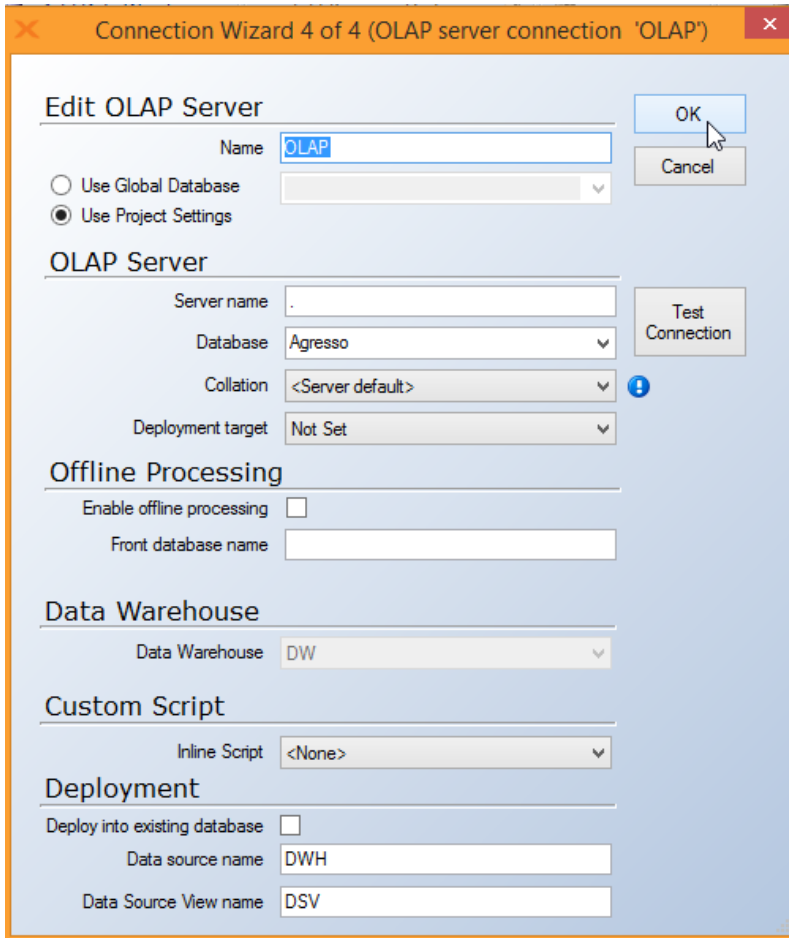
9. Back in the **Connection Manager**, click **Test Connections** to ensure that all connections work and click **Close**.
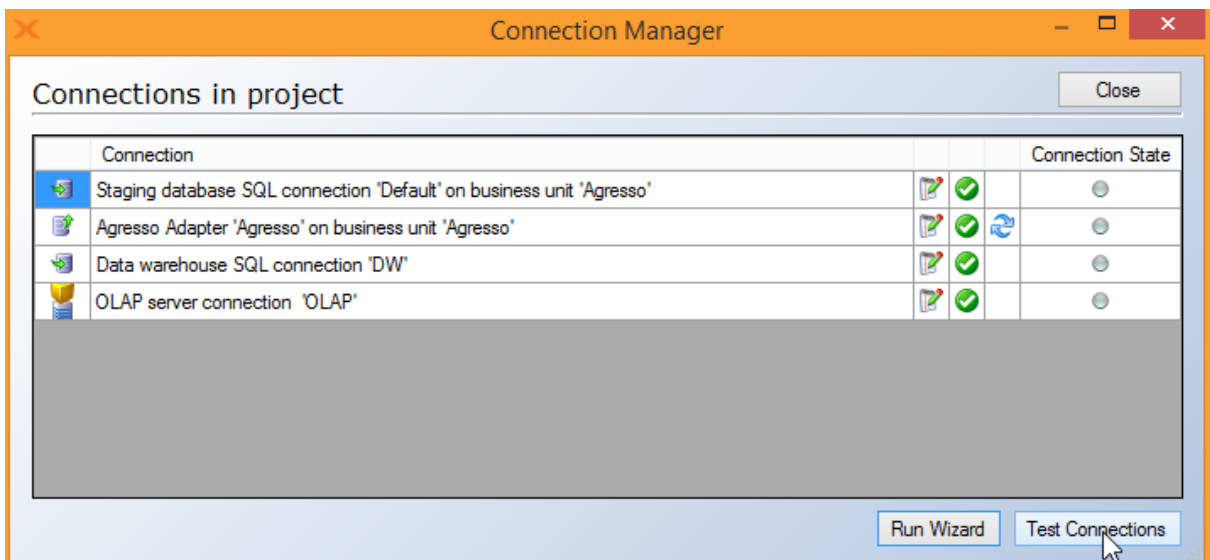
As mentioned earlier, the second part of adding an Agresso Adapter involves setting up the clients, dimensions, flexi-tables and translations that are part of the Agresso system.

To configure the Agresso-specific settings, follow the steps below.

1. In the project tree, right-click the Agresso Adapter you just added and click **Setup Clients**.
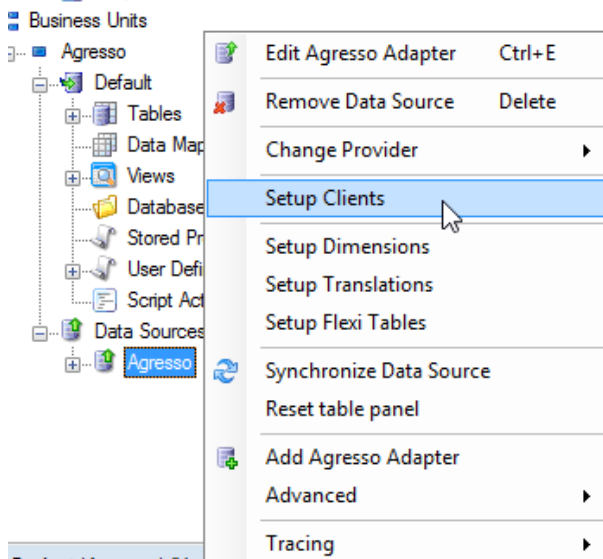
A window containing a list of clients opens.
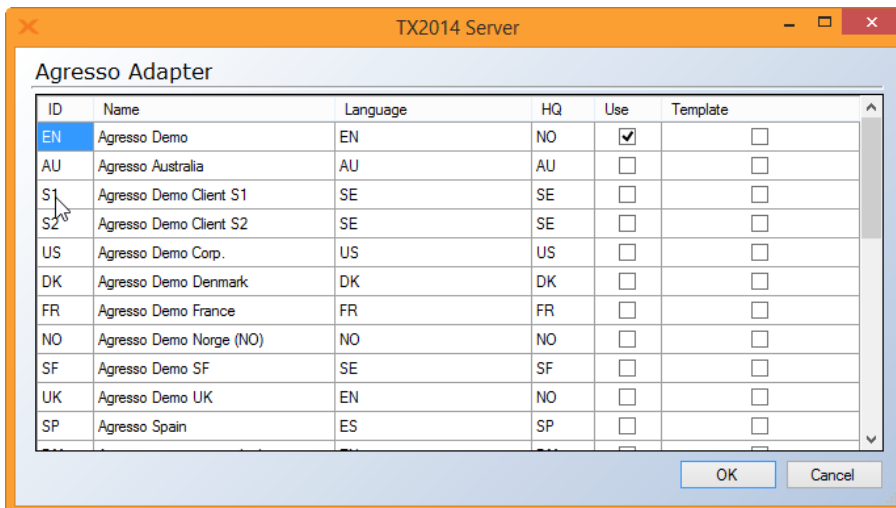
2. Choose the clients you wish to extract data from by selecting the check boxes in the **Use** column. If you want a client to serve as a template for the dimensional setup (see below), select the check box in the **Template** column. Click **OK**.
3. In the project tree, right-click the Agresso Adapter and click **Synchronize Data Source**.
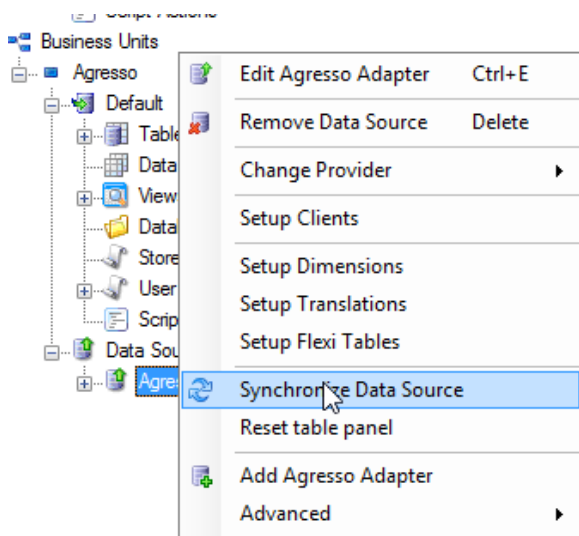
e synchronization window opens. Wait for the process to finish and click **Close**.

4. Before the next step, the project needs to be deployed and executed. In the ribbon, on the Project tab in the Development group, click **Manual Deployment and Execution**. The **Deploy and Execute** window opens. Click **Start**, wait for the process to finish, and click **Close**.

5. In the project tree, right-click the Agresso Adapter and click **Setup Dimensions**. The **Agresso Dimension Setup** window opens.
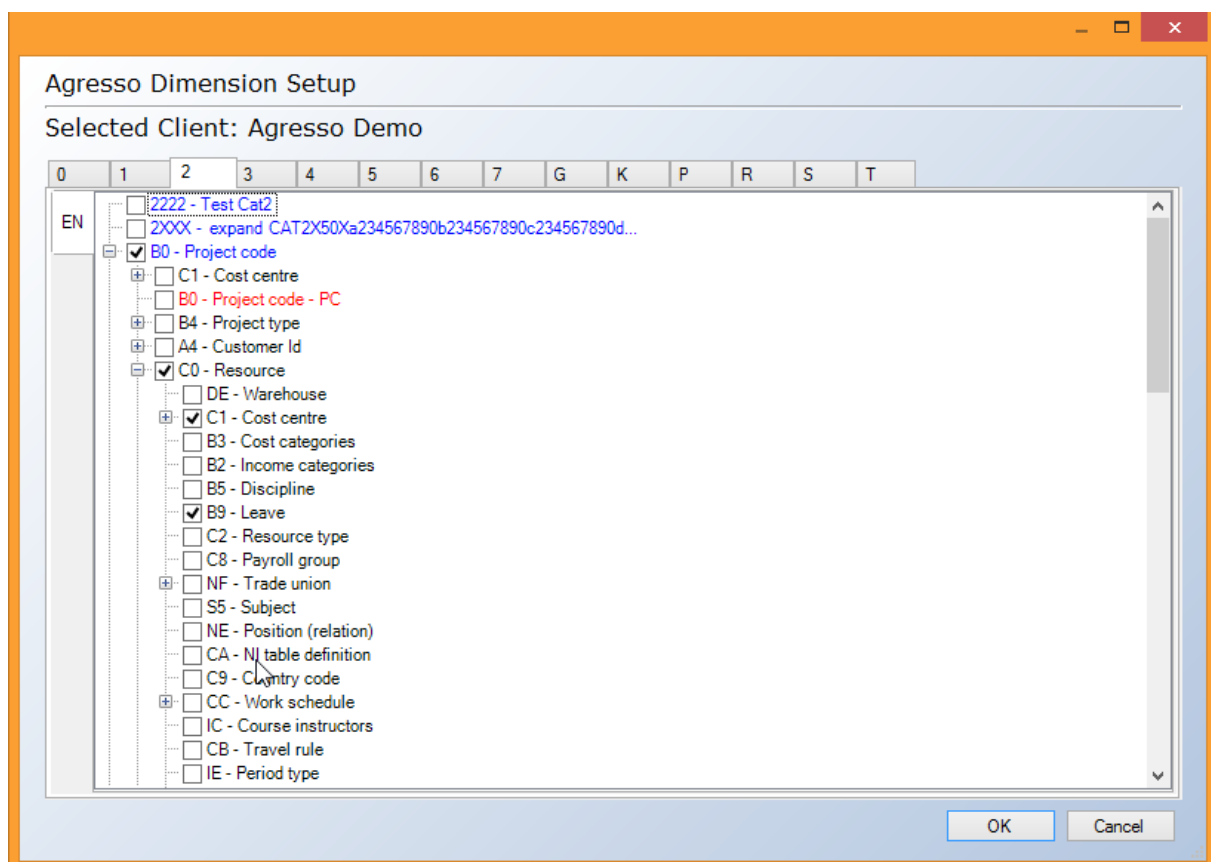
6. Click **Load Dimension Values** (Default setting only shows previously selected dimension attributes)

7. Select the dimension attributes you wish to use in the project. If you did not select a template in the client setup, you need to do this for each client.
8. Click **OK**. The selected attributes are automatically added to the dimension table and can then be used in the project.

9. In the project tree, right-click the Agresso Adapter and click **Setup Flexi Tables**. The **Agresso Flexi Table Setup** window opens. The content of the **Flexi Tables** and **Clients** list depends on the clients and dimensions you have chosen previously.
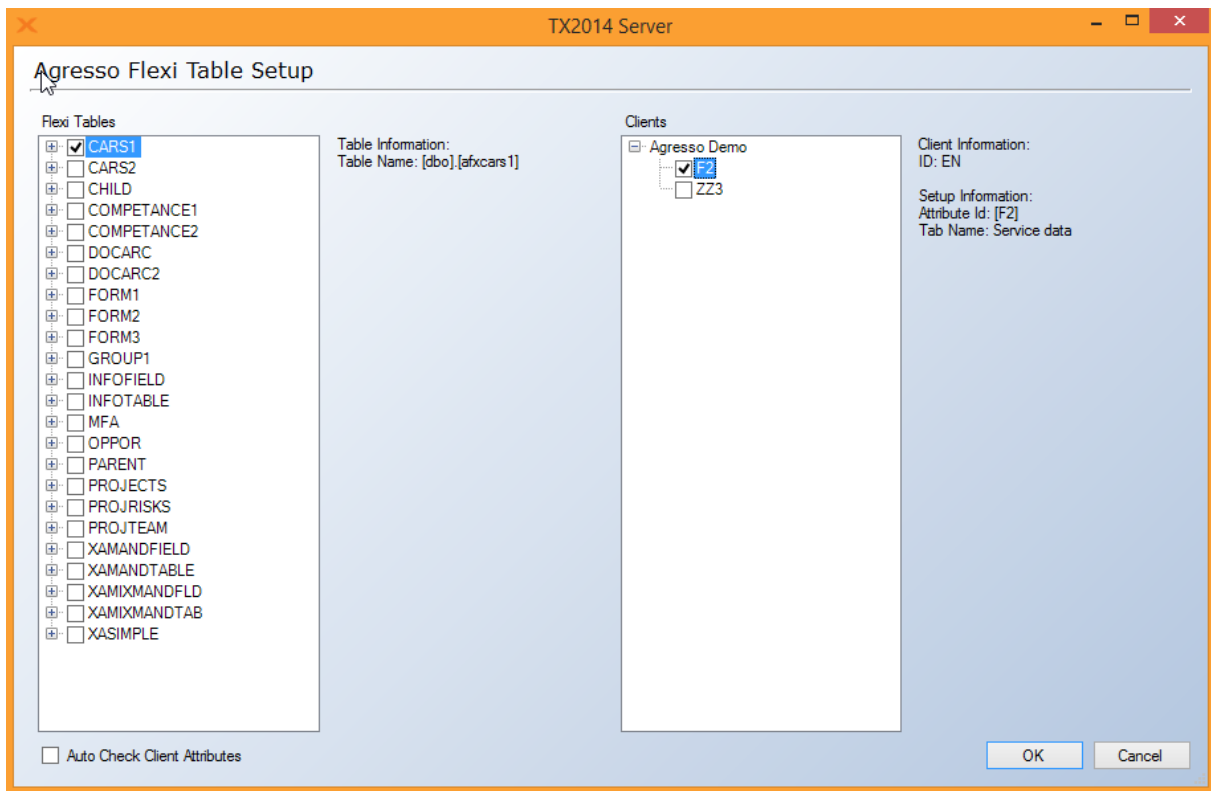
10. Select the combinations of flexi tables and client attributes that you wish to use in your project.

11. Click **OK**. The selected Flexi Tables are automatically added to the source tables and must be integrated into the solution.

12. In the project tree, right-click the Agresso Adapter and click **Setup Translations**. The **Agresso Translation Setup** window opens.
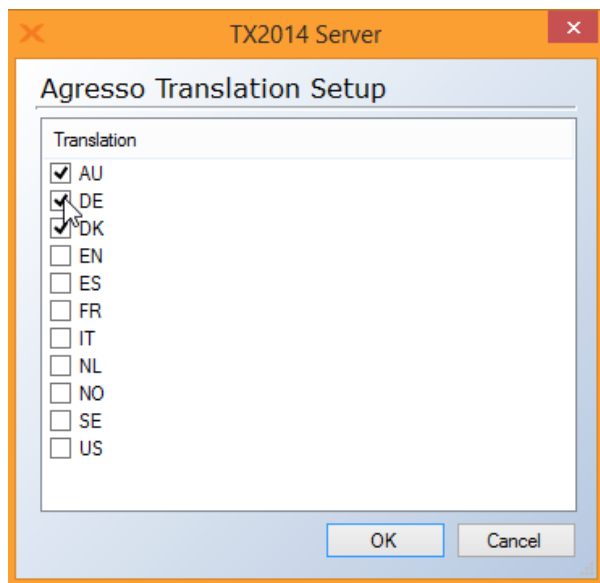
13. Select the translations you wish to use in your project from the **Translation** list.
14. Click **OK**. The selected translations are automatically added to the source tables for use in the project.

DESIGNING THE DATA WAREHOUSE

Once you have set up your data warehouses and business units and established connections to the sources you wish to use, it is time to start designing your data warehouse. This involves transferring data from the source systems to the data warehouse(s) via a staging database and applying transformations and data cleansing rules along the way.

DIMENSIONAL MODELING

Dimensional modeling is the technique and methodology used in data warehouse design. A dimensional model typically consists of a fact table and a number of dimension tables taking the form of a star schema or a snowflake schema.

The fact table defines what you are going to analyze and contains numerical values. Fact tables are more commonly known as transaction tables and generally may contain a large number of records.

Dimension tables define how to analyze the information in the fact table. These tables are much smaller as they do not contain transactional information. Instead, these tables contain summary or attribute information about things such as: customers, items, general ledger accounts, etc.

STAR SCHEMA

A star schema represents a fact table that is directly joined to all dimension tables. The schema resembles a star with the fact table at the center and the dimension tables as the points of the star. In a star schema, all of the descriptive information for a particular dimension is contained in a single table. This is the methodology used in most of the standard projects. Below is an example of a standard star schema.
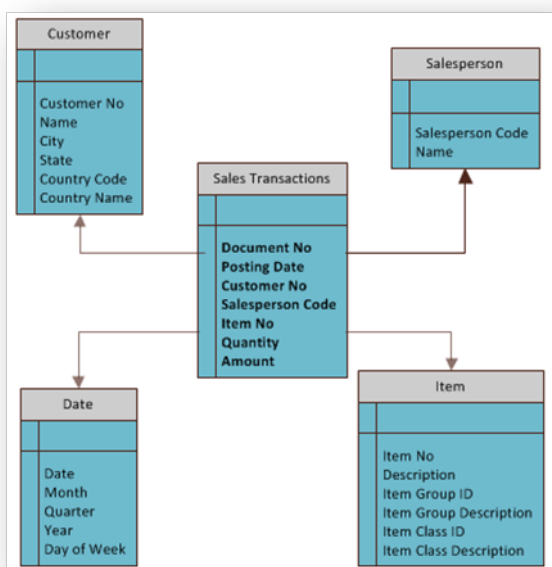
SNOWFLAKE SCHEMA

A snowflake schema is based on a variation of the star schema. The snowflake schema is more normalized with a fact table at the center and dimension tables surrounding this fact table. One of the main differences between the snowflake schema and star schema is that

all the descriptive information in a snowflake schema can be stored in multiple tables, whereas with a star schema, all of the descriptive information is in a single table. Notice in the diagram below, the Address, Item Class, and Item Group tables, whereas in the star schema all of these details were stored in the Customer and Item tables.
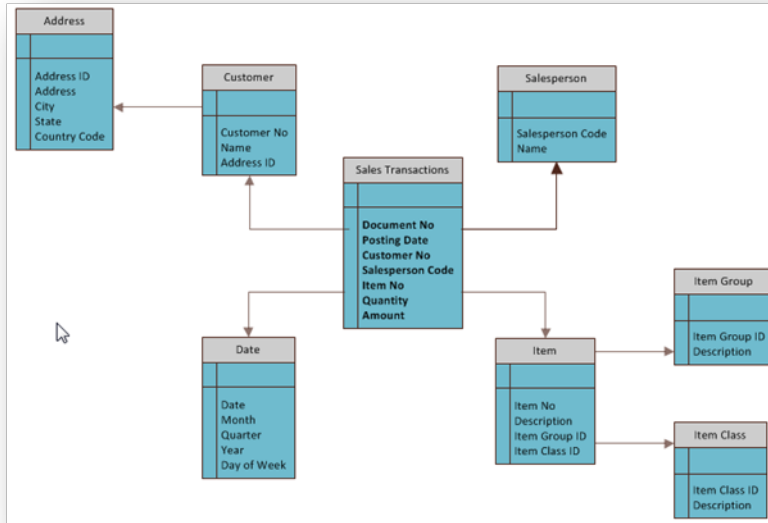
DATA LINEAGE AND IMPACT ANALYSIS

As your project grows, the increasing complexity can make it hard to keep track of all objects and their dependencies. To help you with this, Jet Data Manager contains two tracing features, data lineage and impact analysis, that you can use on most objects in Jet Data Manager:

- Data warehouses
- Business units
- Staging databases
- Data sources
- Tables
- Fields
- OLAP servers
- Cubes
- Dimensions

The purpose of data lineage is to show you where the object in question gets its data, while the impact analysis feature shows you where the data is used.

TRACING AN OBJECT WITH DATA LINEAGE OR IMPACT ANALYSIS

To use the tracing features on an object, follow the steps below.

1. Right-click the object, click **Tracing** and click **Data lineage** or **Impact analysis**, depending on what you wish to see. The **Tracing** window opens with a tracing diagram:
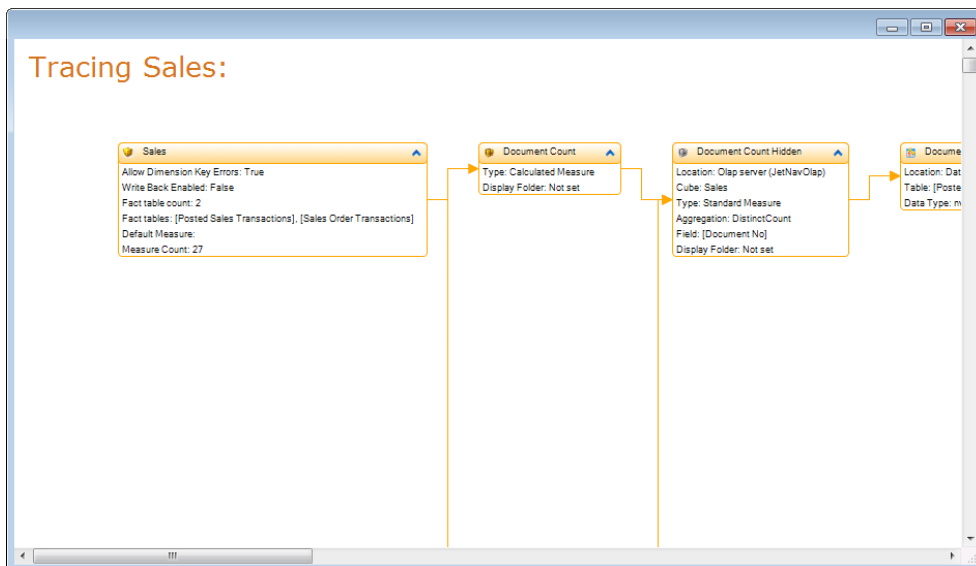
2. Click an object to see what objects deliver data to the object (marked in green - and what objects it delivers data to (marked in red).
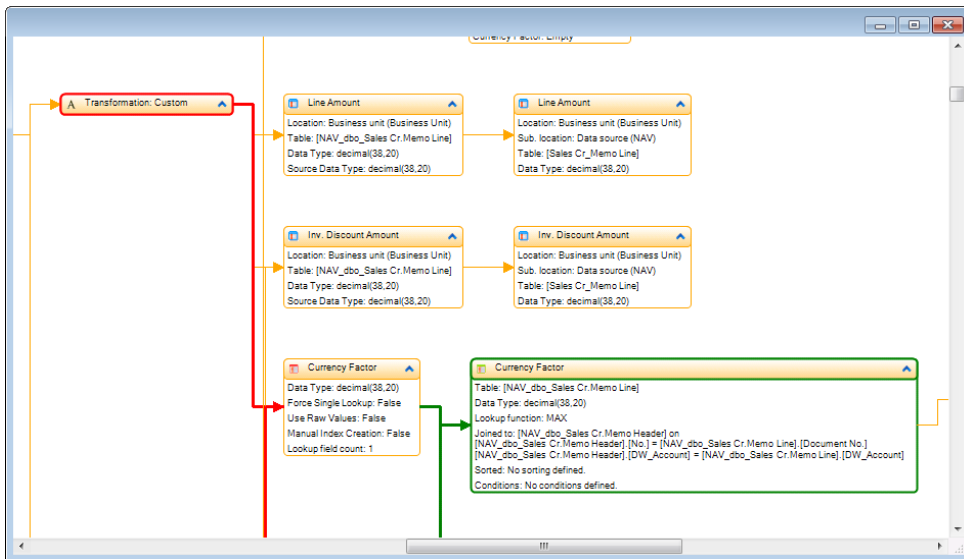
3. Click and hold on a object to move it around in the diagram and release the mouse button to place it.
4. Right-click the background in the **Tracing** window to bring up a menu.

5. Click **Auto Layout Vertical** or **Auto Layout Horizontal** to make Jet Data Manager reorder the objects.
6. Click **Collapse All** or **Expand All** to either collapse the objects to create a better over-view, or expand the objects to see more details.
7. Click **Print** to bring up a **Print** window, where you can select print options and print the diagram.

DOCUMENTATION

For audit and other purposes, you might need a printable documentation of your project. Jet Data Manager lets you create full documentation of a data warehouse, a business unit or an entire project with a few clicks. The documentation contains names, settings and descriptions for every object as well as the code where applicable.

GENERATING DOCUMENTATION

To generate documentation of a part or your entire project, follow the steps below.

1. If you wish to generate documentation for the entire project, on the **Data** tab, right-click the project node and click **Documentation**.
   -OR-
   If you wish to generate documentation for a data warehouse, on the **Data** tab, expand your project, expand **Data Warehouses**, right-click the relevant data warehouse and click **Documentation**.
   -OR-
   If you wish to generate documentation for a business unit, on the **Data** tab, expand your project, expand **Business Units**, right-click the relevant business unit and click **Documentation**.

   The **Documentation** window opens.
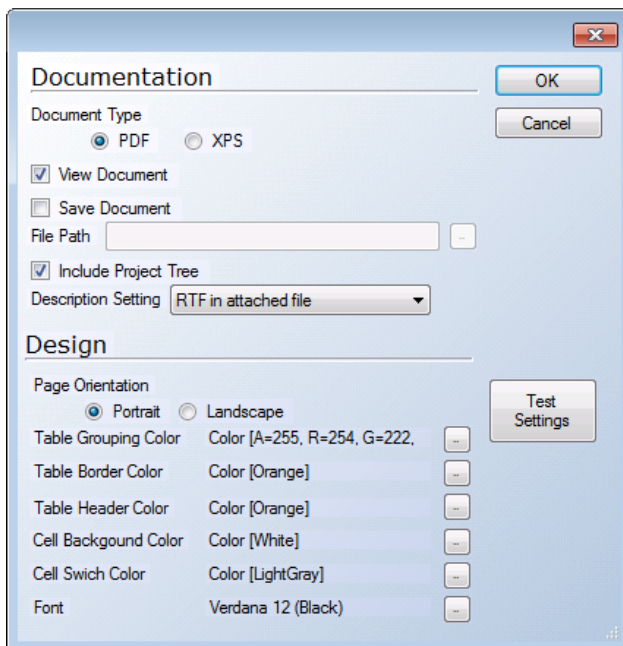
2. Under **Document type**, click the file format you wish to use.
3. Select **View document** to view the document when Jet Data Manager has generated it. Jet Data Manager shows the document using external viewers that needs to be present on the machine.

4. Select **Save document** and click the ellipsis next to the **File path** box to input a file path and name. The document will be saved here when Jet Data Manager has generated it.

5. Select **Include project tree** to include a rendering of the project tree in the document.

6. In the **Description setting** list, click **RTF in Attached File** to have descriptions of objects in your project attached to the document as RTF files, **Only Text Placed in Document** to have the descriptions included in the document as plain text and **No description** to not include the descriptions in the generated documentation.

7. Under **Design** you can adjust the layout and design of the documentation. Under **Page Orientation**, click the page orientation you wish the documentation to use.

8. Click the ellipsis next to the different colors to choose a color for the color in question

9. Click the ellipsis next to **Font** to choose a font for the documentation.

10. Click **Test settings** to generate and open a sample file with the colors you have chosen.

11. Click **OK** to generate the documentation.

TABLES

In addition to the standard tables, Jet Data Manager has three special table types: Date tables, hierarchy tables and aggregate tables, a version of an ordinary table with aggregated data.

CHANGING SETTINGS FOR A TABLE

Most settings for tables are consolidated in the Table Settings window.

- To open the table settings window, right click a table and click **Table settings**.
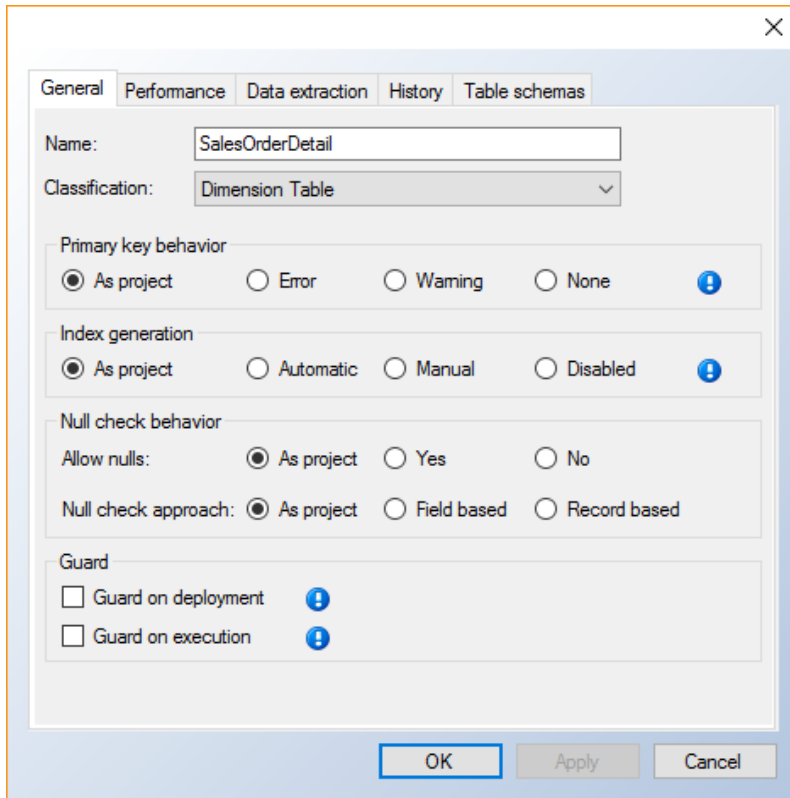
Depending on the table type, not all settings are available. For instance, incremental load does not make sense for date and hierarchy tables, so incremental load settings are disabled for these table types.

GUARDING A TABLE

Guarding a table tells Jet Data Manager to skip the table on execution or deployment. This is useful if, for instance, the table contains old data from a legacy system that is no longer running.

- To guard a table, right click the table, click **Table settings** and then select **Guard on deployment** and/or **Guard on execution**.

ENABLING BATCH DATA CLEANSING TO IMPROVE DATA CLEANSING PERFORMANCE

You can choose to split the INSERT statement up in batches during data cleansing, i.e. when copying data from the transformation view for table to the valid table. This saves log space on the SQL Server which gives you better performance on large tables with 100,000s or millions of rows.

To enable batch data cleansing, follow the steps below.

1. On the Data tab, in a data warehouse, right click the table you wish to use batch data cleansing on and click **Table settings**.
2. Click the **Performance** tab and select **Enable batch data cleansing**.
3. (Optional) Enter the number of records you would like each batch to contain in **Batch size**. The default is 100,000.
4. Click **OK**.

DATE TABLES

You will typically use date tables when you build OLAP cubes on top of the data warehouse. Most often, the cubes you create will contain a date dimension to make it possible to analyze data over time. For example, you may report data on a daily, weekly, or monthly basis. In Jet Data Manager you use date tables, stored in the data warehouse, as the basis for your time dimensions.

In addition to day of month, day of quarter, week of year and other ordinary information about each date, date tables also contain indexes. An index is a column in the table that tells you something about the date's relation to the current date. Data tables contain year, quarter, month and week indexes. All index values for today's date are 0, while for instance, any day last year would have a year index of -1. This makes it trivial to compare e.g. the same month, day or quarter across years.

Data tables can also contain custom periods, special periods of time, for instance holidays or yearly sales campaigns, that enables you to easily track data across these reoccurring time periods.

ADDING A DATE TABLE

To add a date table, follow the steps below.

1. On the **Data** tab, on a data warehouse, right click **Tables** and click **Add Date Table**. The **Add Date Table** window opens.

2. In the **Name** box, type a name for the table.
3. Select a **Date range** by entering a **Start date** and an **End date**. Instead of entering an end date, you can enter a number of days to add to the current date in **Days ahead**. This way, your date table will effectively never end.
4. Under **Date display**, select the **Format** you wish to use for dates. You have the fol-lowing options:
   - YYYY-MM-DD
   - DD-MM-YYYY
   - MM-DD-YYYY
5. Select which **Separator** to use in the format you chose. You have the following options:
   - - (dash)
   - / (slash)
   - . (dot)
6. Under **Week numbering**, click the **First day of the week**. You have the fol-lowing options:

- Sunday
- Monday

7. Select the how to define the First week of the year. You have the following op-tions:
    - First 4-day week (following the ISO 8601 standard, common in Europe)
    - Starts on Jan 1 (common in North American)
    - First full week

8. Under Fiscal year, click **Staggered** to use a staggered fiscal year and click the first month of the staggered fiscal year in the **First month** list.

9. (Optional) Click **Add** under **Custom periods** if you wish to add a custom period. In the custom period window that opens, you can type a **Name** for the custom period and **Name**, **Start date** and **End date** for the included periods. You can also import and export custom periods by clicking **Import** and **Export** respectively. Click **OK** when you are done.

10. (Optional) Click **Custom names** if you wish to change the names used for days, quarters and months. In the **Date Table Custom Names** window that opens, you can type the names you wish to use. The default is derived from the regional settings on the deploying machine. Click **OK** when you are done.

11. Click **OK** to add the date table.

HIERARCHY TABLES

A hierarchy table is used to select data from a table and create a new reporting structure which is different from the structure in the data source. You will typically use a hierarchy table for financial reporting where you want to consolidate data from a number of different accounts, such as ledger accounts.

A hierarchy table is used in conjuction with a parent-child dimension. First, you create the hierarchy table and specify the contents of the table.Then you create a parent-child dimension and add it to a cube. When you build the structure, be sure to choose names that are meaningful to the end-user.
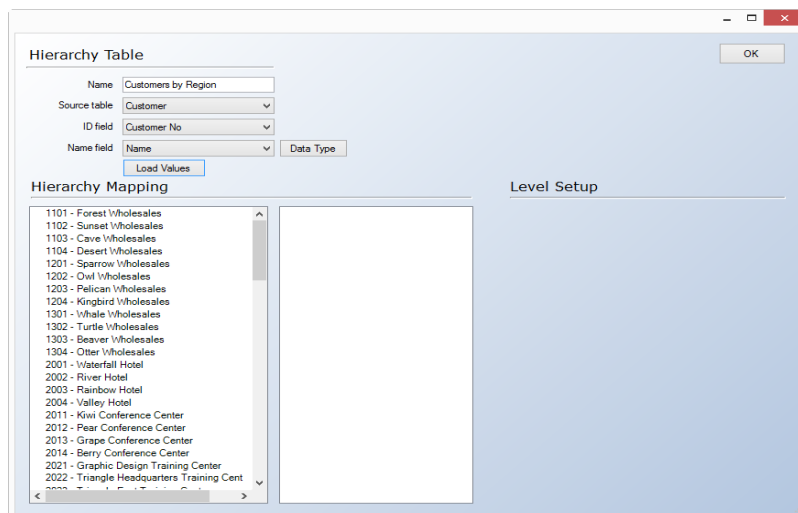
ADDING A HIERARCHY TABLE

1. On the **Data** tab, expand **Data Warehouses**, and expand the relevant data warehouse.

2. Right-click **Tables** and then click **Add Hierarchy Table**.

The **Hierachy Table** window opens:

3. In the **Name** box, type a name for the table.
4. In the **Source table** list, click the table containing the desired data.
5. In the **ID field**, click the field that identifies the individual entries in the table; for example, the customer number. If you need more than one field to identify the entries, you have to create a concatenated field before you create the hierarchy table.
6. In the **Name field**, enter the name that identifies the individual entries; for example, account name, and then click **Load**. The Hierarchy Mapping pane is now populated with the entries of the source table.
7. You can now create the report structure. The structure you create corresponds to the structure of the report that is displayed to the end-user.
8. Right-click in the **Blank** pane, and then select **Add Root Heading**. The root headings become root nodes in the final report.
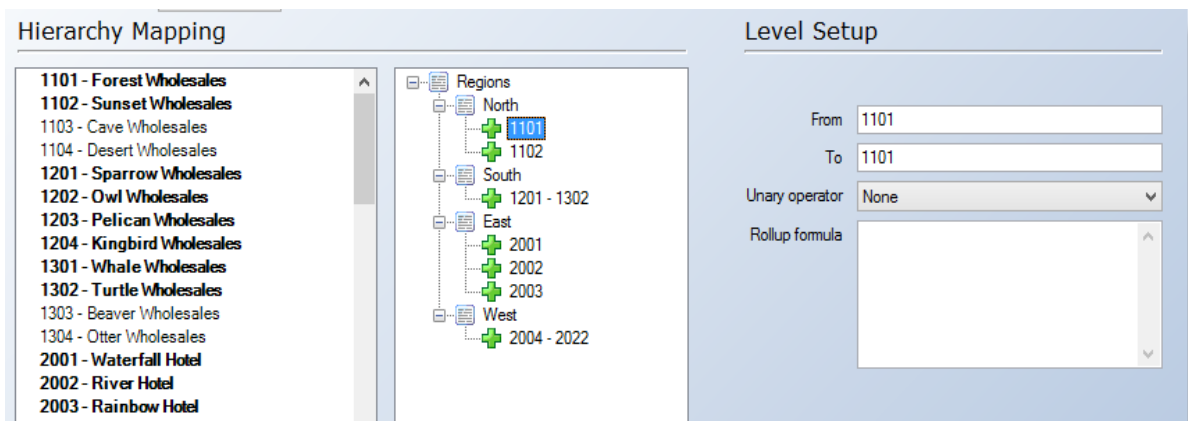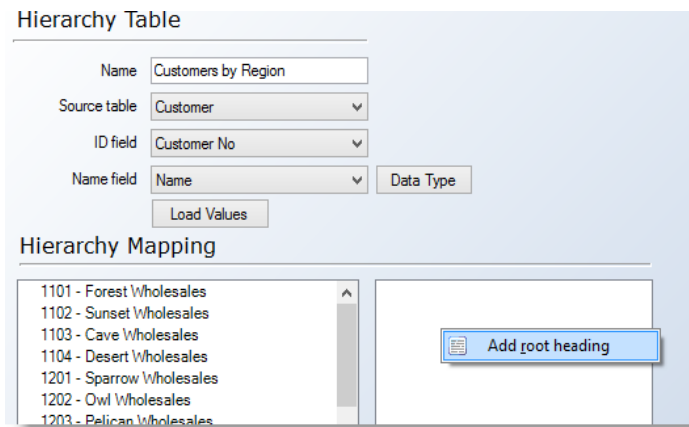
9. In **Level Setup**, type a name for the heading in the **Name** field.
10. Right-click the root heading, and select **Add Sub Heading** to add a child node to the structure.
11. In the **Name** box, type a name for the child node.
12. In the **Unary Operator** list, you specify how you want the value of the child node to be aggregated to the sum of all the values in the subheading. The unary operator ensures that the values are aggregated properly in the final report. You have the following options:

Operator	Definition
None	The value is ignored
Add	The value is added to the sum of the values
Subtract	The value is subtracted from the sum of the values
Multiply	The value is multiplied by the sum of values
Divide	The value is divided by the sum of the values

Repeat steps 8-12 for all root headings and subheadings you want to add.

13. Click and hold an entry in **Hierarchy Mapping** , and drag it to the preferred sub-heading in the **Blank** pane. Alternatively, you can specify a range of entries by typing the relevant numbers in **From** and **To**.
14. To exclude an entry from a given range, right-click the relevant subheading, click **Add Exclude**, and then specify a range by typing the relevant numbers in **From** and **To**. Alternatively, right-click the specific entry and select **Change to Exclude**.
15. If a root heading or subheading represents the sum of other subheadings, such as Contribution Margin, you can use a formula to determine the content of the heading. Type a formula in the **Roll-up formula**. Formulas are written in MDX.
16. Click **OK**when you have completed the structure. You can now create the parent-child dimension where the consolidation table will be used.

Note: When you create the parent-child dimension you will typically use Sort By Attribute. You therefore need to create a Sort order dimension level where the key column is Sort order. It is also necessary to enable Unary column and Roll-up column on the dimension. You can then set the parent-child dimension to Sort By Attribute.

AGGREGATE TABLES

An aggregated table is an aggregated version of another table in you project. Often, you will not need the transactional level in financial or sales reports, but only data grouped by business unit or sales team. This makes the aggregated tables feature very useful if you are doing reporting directly from you data warehouse as opposed to using, for instance, OLAP cubes.

ADDING AN AGGREGATED TABLE

To add an aggregated table, follow the steps below.

1. On the **Data** tab, under **Tables** in a data warehouse, right click the table, you wish to add an aggregated version of, click **Advanced** and click **Add Aggregate Table**. The **Add Aggregate Data Table** window opens.

2. Under **GroupBy**, you can choose what columns on the table the aggregated table should use for grouping the aggregated data. Click the column you wish to use in the empty list under **Table: [table name]**. Type a name for the field in **Field** name. In the column you have chosen contain date values, click the list under **GroupBy Type** to adjust the granularity of the grouping. You can choose second, minute, hour and all the way up to year. Note that you can use the same date column multiple times with different **GroupBy** types. For other data types, the **GroupBy** type will always be **Value**.

3. Under **Aggregate**, you can choose what columns from the table you wish to have aggregated. Click the column you wish to use in the empty list under **Table: [table name]**. Type a name for the field in **Field** name. Click the list in the **Aggregation Type** column and click the method you wish to use for calculating the aggregation. You have the following options:

   - **Min**: The lowest value of the field in question.
   - **Max**: The highest value of the field in question.
   - **Count**: The number of rows.
   - **Count_Big**: Same as count, but is able to count higher than 2^31, be-cause it uses the bigint data type instead of the int data type.
   - **DistinctCount**: The number of unique values in the field.
   - **Sum**: The sum of all row values.
   - **Average**: The average of all row values.

4. Click **OK**. The aggregated table is added under **Tables** and can be recognized by its yellow icon.

PREVIEWING DATA

During development, it is often useful to be able to see what data is present in different tables. For instance, you might want to check that a transformation works as intended. Jet Data Manager provides two different ways of viewing the content of a table.

PREVIEW TABLE

The Preview Table feature gives you a basic overview of the content of a table. To preview the content of a table, follow the step below.

1. On the **Data** tab, expand the relevant business unit, expand **Data Sources**, expand the data source and right-click the table.
   - OR -
   On the **Data** tab, expand **Data Warehouses**, expand the relevant data warehouse, expand **Tables** and right-click the table.

   The Preview Table window opens.

2. You have a number of options for previewing the data:
   - In **Select Top**, type the number of rows you wish to fetch and display. Click **Apply Top** to apply the setting. Please notice that the select top is applied before any sorting of data.
   - In the **Date format** list, click the date format you wish to use for dates in the data. Click **Apply Format** to apply the settings.
   - In the **Instance** list, click the instance of the table you wish to preview.
   - Click **Count Rows** to
3. When you are done, click **Close** to close the window.

QUERY TOOL

The Query Tool is a powerful supplement to the Preview Table feature that gives you more flexibility in exploring the content of a table. You can execute any SQL query to see the data you wish to see the way you wish to see it.

OPENING THE QUERY TOOL

You can open the Query Tool in three different ways.

- Right click a table, click **Preview Table** and click **Query Tool** in the **Preview Table** window.
- Right click a table, click **Advanced** and click **Query Tool**.
- Click table and press **F8** on your keyboard.

The Query Tool opens with a query that selects the content of the currently selected table, similar to the query that is executed to get the content for the **Preview Table** window.

EXECUTING QUERIES

To execute a query, follow the steps below.

1. Open the **Query Tool** using one of the options described above.
2. If available, choose the **Source** and **Account** you wish to query. **Account** is only displayed when using an adapter with multiple possible accounts.
3. Enter your query in the top text box of the **Query Tool** window. You can enter multiple queries that will be executed in sequence by Jet Data Manager.
   Adjust **Max no. of rows** to the maximum number of rows you wish to have returned.

4. Click **Execute**

- OR -

Press **F5** on your keyboard.
5. If you wish to stop the query before it completes, click **Stop**.

When they query is complete, you can see the result in the **Result** tab. If you have entered multiple queries, you can select the query result you wish to see in **Result set**. If your query resulted in a message, for example because of a syntax error, the Message tab will display this message.

DRAG-AND-DROP AND THE QUERY TOOL WINDOW

The Query Tool supports drag and drop of tables and fields.

- You can drag a table or a field from the project tree into the query. This places the table name in the query.
- If you drag a table to an empty query, the default query is generated. The default query fetches everything in the table.
- If you drag a table from another source into the window, you will be asked if you want to change connection and generate a default query. If you answer **No**, the name is simply added to the query.

SORTING AND FILTERING DATA

The Query Tool enables you to sort and filter the results.

Note: Only the rows returned by the query are available for sorting and filtering in the **Results** tab. If you wish to sort or filter all the rows in a table, the most efficient way is to include the conditions in the query, e.g. by using "order by" or "where" clauses. Fetching thousands of rows and sorting them using the tools provided in the **Results** tab can be very slow.

To sort the data, follow the steps below.

1. Open the **Query Tool** and execute a query as described above.
2. In the **Result** tab, click a column heading to sort the rows on the value in that column. Click again to switch between ascending or descending order.

To apply a filter, follow the steps below.

1. Open the **Query Tool** and execute a query as described above.
2. In the Result tab, click the filter icon besides the name of the column you wish to filter on. You have five filtering options:
   - **(All)** is equal to no filtering.
   - **(Custom)** opens the **Custom Filter** window, where you can add conditions for filtering.

Each condition evaluates the value of the row field compared to the possible values in the column. The comparison can be made on **Equals**, **Does not equal**, **Less than**, **Less than or equals to**, **Greater than** and **Greater than or equal to**. Click **Add** to add an additional filter and click **Delete** the currently selected condition. You can choose to filter on **Any** or **All** conditions, i.e. stringing the conditions together with "or" or "and". Click **OK** to activate the filter.

- **(Blanks)** shows rows where the column in question is blank, i.e. empty.
- **(NonBlanks)** shows rows where the column in question is not blank.
- A specific value. All unique values in the column is listed and can be chosen as a filter.

MOVING AND RELATING DATA

You can move tables, fields, and views from:

- A data source to a staging database
- A staging database to a data warehouse

This way tables and fields from different business units can be consolidated in one data warehouse. You do not have to move an entire table from the staging database. Rather, you can simply move individual fields and add them to the desired table in the data warehouse.

When you move fields, you specify whether to copy the fields, split the fields, or concatenate the fields.

Concatenate Field: Two or more fields are combined into one field.

Copy Field: A copy of the field without any changes.

Split Field: One field must be divided into two or more fields.

MOVING DATA INTO THE STAGING DATABASE

In order to add data to the staging database, you will select the tables and fields that you would like to extract from the desired data source. To move data from the data source to the staging database, follow the steps below.

1. Expand the **Data Sources** node, and click the preferred data source.
2. On the right-hand side, you will see a list of all of the tables and fields in the selected data source.
3. Expand the desired table by double-clicking the table name or by clicking on the table and pressing the right-arrow on the keyboard.
4. Click the check boxes for all of the fields that you wish to move from the data sources into the staging database.
5. Deploy and execute the tables to create the structure in the staging database, and move the data over.

MOVING DATA FROM THE STAGING DATABASE

You can move tables, fields, and views from a staging database to a data warehouse. This way, tables and fields from different business units can be consolidated into a single data warehouse. To move data from the staging database to the data warehouse, follow the steps below.

1. Expand the business unit containing the desired table to be copied, and then expand the staging database.
2. Expand the destination data warehouse where you will send the table.

3. Expand Tables, left-click the preferred table, and then drag and drop it to **Tables** in the data warehouse.

The table and all fields belonging to the table are added. If you expand a field, you can see which field has been copied.

MOVING AND CONCATENATING A FIELD FROM A TABLE ON THE STAGING DATABASE

To move and concatenate a field from a table on a staging database to a field on a data warehouse, follow the steps below.
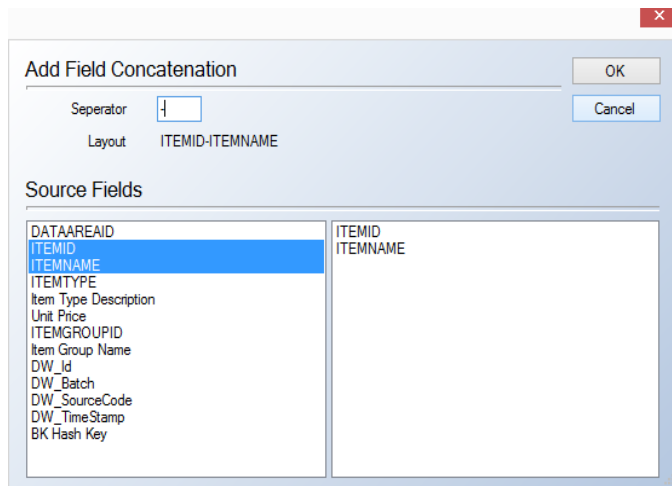
1. Expand the preferred business unit, expand the staging database, and then expand **Tables**.
2. Expand the preferred data warehouse, expand **Tables**, and then expand the table to be modified.

   Note: The destination table must already contain a custom field, or the table must contain a field that has been copied from the staging database.

3. In the staging database, expand the table containing the field you wish to send.
4. Left-click the field, and then drag and drop it to the preferred field in the data warehouse to display the **Data Movement** window.
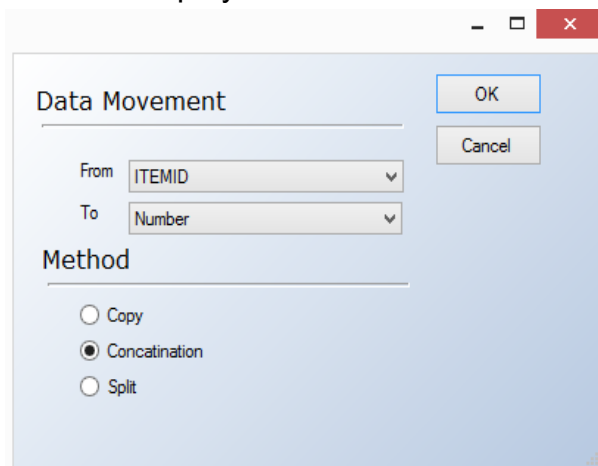
5. In the **From** list, the field you want to move is selected by default.
6. In the **To** list, the field to which you have moved the field is selected by default.
7. Check **Concatenate** to specify that you want to concatenate the selected field with one or more fields, and then click **OK**. The **Add Field Concatenation** window opens.

8. In the **Source Fields** list, select the fields you want to concatenate.
9. In the **Separator** field, specify how the data must be separated in the field, for example, using commas, periods, etc.
10. Click **OK**. Below the field, you can now see where the field originated.

You can also move and concatenate fields from a view to a view or a table in the data warehouse.

MOVING AND CONCATENATING A FIELD FROM A VIEW ON THE STAGING DATABASE

To move and concatenate a field from a view on a staging database to a field on a data warehouse, follow the steps below.

1. Expand the preferred business unit, expand the staging database, and then expand **Views**.
2. Expand the preferred data warehouse, expand **Views**, and then expand the desired view.
   - OR -
   Expand **Tables**, and then expand the preferred table.
3. Left-click the field in the staging database, and then drag and drop it to the preferred field in the data warehouse to display the Data Movement Dialog.
4. In the **From** list, the field you want to move is selected by default.
5. In the **To** list, the field to which you have moved the field is selected by default.
6. Check **Concatenate** to specify that you want to concatenate the selected field with one or more fields, and then click **OK**.
7. In the **Source Fields** list, select the fields you want to concatenate.
8. In the **Separator** field, specify how the data must be separated in the field, for example, using commas, periods, etc.
9. Click **OK**. You can now see where the field originated below the field.

MOVING AND SPLITTING A FIELD FROM A TABLE ON THE STAGING DATABASE

To move and split a field from a table on a staging database to a field on a data warehouse, follow the steps below.

1. Expand the preferred business unit, expand the staging database, and then expand **Tables**.
2. Expand the preferred data warehouse, expand **Tables**, and then expand the desired table.

   Note: The destination table must already contain a custom field, or the table must contain a field that has been copied from the staging database.

3. In the staging database, expand the table from which you want to move a field.
4. Left-click the field, and then drag and drop it to the preferred field in the data warehouse. The Data Movement dialog is displayed.
5. In the **From** list, the field you want to move is selected by default.
6. In the **To** list, the field to which you have moved the field is selected by default.
7. Check **Split** to specify that you want to split the selected field into two or more fields, and then click **OK**.
8. In the **Separator** field, specify the separator that is used to separate the data in the field.
9. In the **Source Field** list, the field you want to split is selected by default.
10. In the **Destination Fields** list, select two or more fields that you want to split the field into.
11. Click **OK**. Below the field, you can now see the originating field in the staging database that the field has been moved and split from.

MOVING AND SPLITTING A FIELD FROM A VIEW ON THE STAGING DATABASE

To move and concatenate a field from a view on a staging database to a field on a data warehouse, follow the steps below.

1. Expand the preferred business unit, expand the staging database, and then expand **Views**.
2. Expand the preferred data warehouse, expand **Views**, and then expand the desired view.
   - OR –
   Expand Tables, and then expand the preferred table.
3. In the staging database, expand the desired view.
4. Left-click the field, and then drag and drop it to the preferred field in the data warehouse to display the Data Movement dialog.
5. In the **From** list, the field you want to move is selected by default.
6. In the **To** list, the field to which you have moved the field is selected by default.
7. Check **Split** to specify that you want to split the selected field into two or more fields, and then click **OK**.

8. In the **Separator** field, specify the separator that is used to separate the data in the field.
9. In the **Source Field** list, the field you want to split is selected by default.
10. In the **Destination Fields** list, select two or more fields that you want to split the field into.
11. Click **OK**. Below the field, you can now see the originating field in the staging database.

COPYING FIELDS FROM THE STAGING DATABASE TO THE DATA WAREHOUSE

To copy a field from a table on a staging database to a field on a data warehouse, follow the steps below.

1. Expand the preferred business unit, expand the staging database, and then expand **Tables**.
2. Expand the preferred data warehouse, expand tables, and then expand the desired table.
3. In the staging database, expand the table from which you want to move a field.
4. Left-click the field, and then drag and drop it to the preferred table field in the data warehouse.

When you copy a field, one data type is sometimes converted automatically to another data type, if required.

TABLE RELATIONS

For Jet Data Manager to know how tables are related, you have to specify relations. Among other things, the relations between tables you have defined are used for a referential integrity check on execution, for the default join when you create conditional lookup fields and for relating dimensions in an OLAP cube.

Under **Relations** under each table in the project tree, you can see the relations that this particular table has to other tables.

Relations are grouped by the foreign table and the relation name defaults to "[foreign table name]_[foreign field name]".

Each relation contains one or more field relations. On each field relation, the part on the left side of the equals sign is a field on the foreign table, while the part on the right side is a field on the table that has the relation.

ADDING A NEW RELATION

To add a new relation, follow the steps below.

1. On the **Data** tab, navigate to the table you wish to relate to another table.

2. Click on the field you wish to base the relation on and then drag and drop the field on a field on another table. Note that the fields must be of the same data type to create a relation.
3. Jet Data Manager will ask you if you want to create a relation. Click on **Yes**. The relation is created on the table on which you drop the field.

RENAMING A RELATION

The relation name defaults to "[foreign table name]_[foreign field name]", but you can rename the relation if you wish. To rename a relation, follow the steps below.

1. Locate the relation in the project tree and click on it.
2. Press **F2** to make the relation name editable.
3. Type the new name for the relation and press **Enter**.

SETTING A DEFAULT RELATION

You can set one of the relations between two tables to be the default relation. A default relation is useful if you have more than one relation between two tables, for instance as join for lookup fields that do not have a specific join set and for auto-relation when you add a dimension to an OLAP cube.

- To set a relation as the default relation, right click the relation and click **Set as Default Relation**.

SETTING THE SEVERITY OF VIOLATING THE REFERENTIAL INTEGRITY CHECK

Jet Data Manager users the relations you have defined to perform a referential integrity check, or foreign key constraint check, when the table is executed. Jet Data Manager checks the value of any field that are part of a relation to see if the value exists in the related field in the related table. If not, Jet Data Manager considers the record invalid.

For instance, a Sales Order table might contain a Customer ID field that is related to a Customer ID field in a Customer table. If a specific sales order record contains a Customer ID that is not in the Customer table, Jet Data Manager considers that record to be invalid.

You can define the severity of the violation on each relation. To set the severity of a violations on a particular relation, follow the steps below.

1. Locate the relation in the project tree.
2. Right click the relation, click **Relation Type** and click your preferred type. You have the following options:
    - **Error**: Jet Data Manager moves the invalid record to the error table. This means data will be missing from the valid instance of the table.
    - **Error with physical relation**: The relation is stored in the database for other database tools to see. The behavior is otherwise the same as error. Note that

the table needs to have a primary key and a unique index set. If Index Automation is disabled on the table, you will have to create the index yourself.

- **Warning**: Jet Data Manager copies the invalid record to the warnings table and the valid instance of the table. You will not be missing data from the valid table. However, you might need to handle the violated rule in some way.
- **Relation only**: Jet Data Manager ignores any violations of the check.

UNUSED FIELDS

Jet Data Managerhas the ability to display all unused fields in the project. This feature is useful for removing unnecessary objects from the project which decreases clutter and improves performance.

DISPLAYING UNUSED FIELDS

From the data tab, right-click either the Business Unit or the Data Warehouse, and select Find Unused Fields.
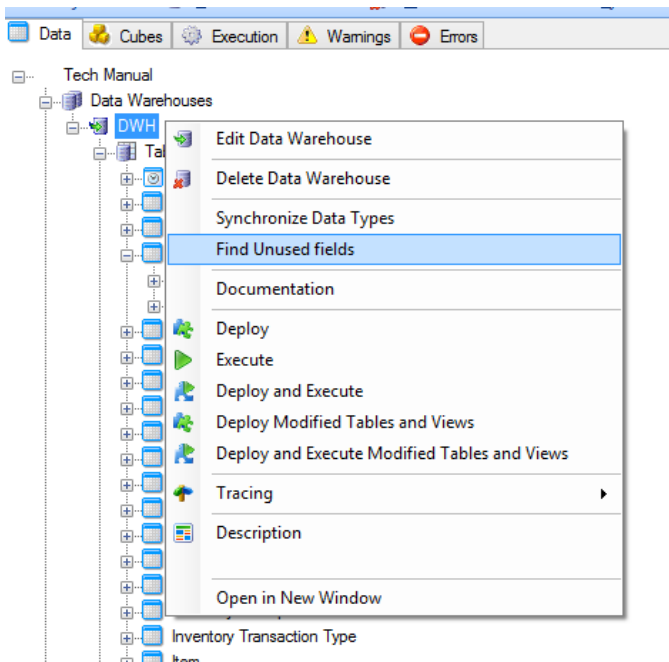
A dialog appears that shows unused fields:
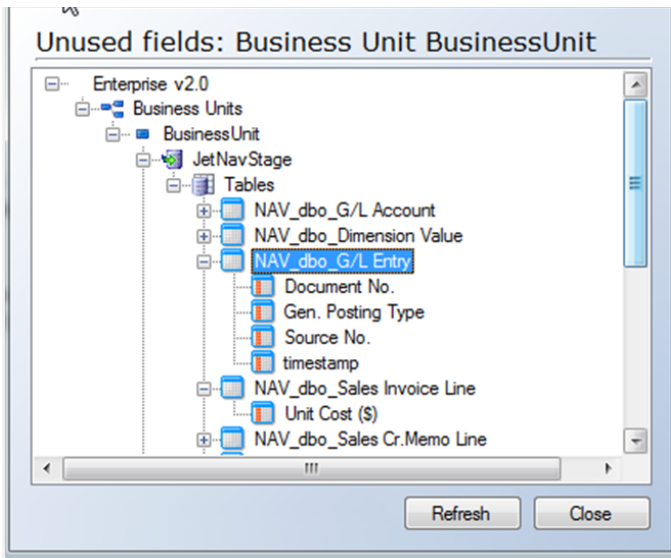
In the staging database, this would show fields that exist but are not:

- Promoted to the data warehouse
- Used as a data selection rule
- Used as an incremental selection rule
- Used in a strongly typed custom table
- Used in a SQL snippet
- Used as a conditional lookup in another table

In the data warehouse database, this would show fields that exist but are not:

- Promoted to the OLAP Cubes as measures or dimensions
- Used as a data selection rule
- Used as an incremental selection rule
- Used in a strongly typed custom table
- Used in a SQL snippet
- Used as a conditional lookup in another table

EXTERNAL TABLES

External tables is a way to incorporate tables from an existing data warehouse into a Jet Data Manager project. This is useful if you, for instance, have a legacy data warehouse humming along that you would like to use data from without remodeling it in Jet Data Manager.

An external table will initially not be deployed or executed, but will be available for data movement to data warehouses and data marts, can be used in views and scripts and for cubes and dimensions. By default Jet Data Manager will create views and read data from

the external connection, but you can also chose to move the data into you Jet Data Manager data warehouse. You can also add a custom SSIS package to the table, which can then be executed.
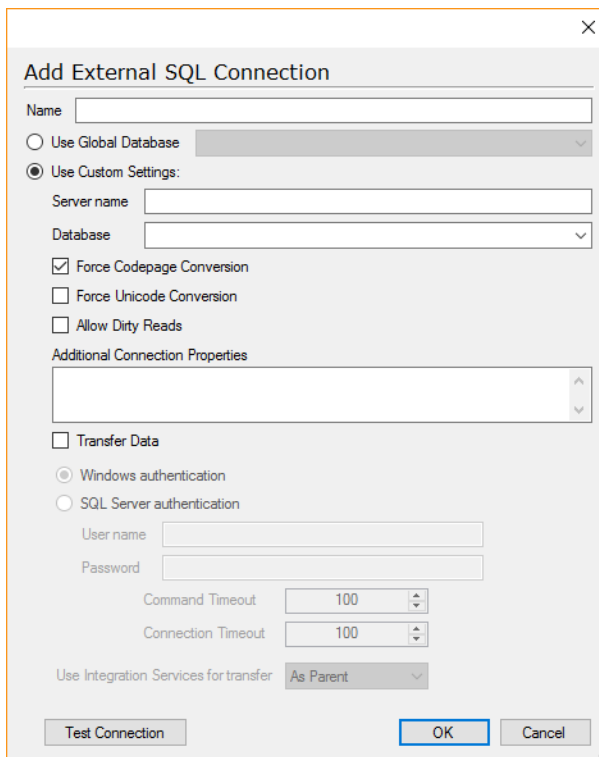
ADDING AN EXTERNAL SQL CONNECTION

To add an external table, you first need to add an external SQL connection.

> **Note:** The SQL Server needs to be on the same physical SQL Server instance as your data warehouse or on a linked SQL Server, unless you enable the Transfer Data option.

To add an external SQL connection, follow the steps below:

1. Right click your data warehouse or a business unit, navigate to **Advanced** and click **Add External SQL Connection**. The **Add External SQL Connection** window appears.

2. Type a **Name** for the connection.
3. Click **Use Global Database** and choose a global database in the list
   - OR -
   Click **Use Custom Settings** and customize the settings:
   - Type the name of the server in the **Server Name** box.
   - Type the name of the database you want to use in the **Database** box.
   - Select **Force Codepage Conversion** to convert all fields to the collation of the data warehouse.
   - Select **Force Unicode Conversion** to declare all alphanumeric fields as **nvarchar**.

- Select **Allow Dirty Reads** to allow reading from the source without locking the table.
- (Optional) Enter additional connection properties in the **Additional Connection Properties** box.
- Select **Transfer Data** to move the data from the external SQL Server to your local data warehouse, much like a regular SQL data source.

4. Click **OK** to add the source and close the window.

ADDING AN EXTERNAL TABLE

To add an external table, follow the steps below.

1. Navigate to **External SQL Connections** under your data warehouse or business unit in the project tree, right click the connection you just created and click **Read Objects from Data Source**.
2. When Jet Data Manager has finished reading objects from the data source, the source explorer pane in the right hand side of the window is populated with the objects from the source. Select the tables, views and fields you wish to use in you data warehouse.

WORKING WITH EXTERNAL TABLES

The external tables in your project are shown in the project tree alongside the standard tables and you can use them in the same way. External tables can be used in dimensions and cubes, for reporting, in Qlikview models etc. You can recognize an external table in the project tree on the black table icon.

Some of the transformations and data cleansing you can do with standard tables, can be done with external tables as well. You can add custom fields, but not lookup fields. For instance, you can add a custom field to the external table and apply a transformation to the field to concatenate two other fields on the table.

You can also add custom data to an external table.

DEPLOYING AN EXTERNAL TABLE

To deploy an external table, right click the table and click **Deploy**. A View will be created that selects from the external table.

EXECUTING AN EXTERNAL TABLE WITH AN SSIS PACKAGE

Since an external table is set up outside Jet Data Manager, Jet Data Manager expects it to be executed separately from your project. This means that you initially will not find any execute command on an external table. However, if you have a SSIS Package that is used to populate the table, you can add this package to the table and get the ability to execute the table.

1. Right click an external table, navigate to advanced and click **Customize code**. The **Customize Code** window appears.
2. Click the **Add** button to the right of **SSIS Package**. The **Custom Editor** window appears.
3. In the **Editor Name** list, click you editor of choice and click **OK**. The **Custom SSIS** window appears.
4. Make sure **Existing Package** is selected and click **OK**. The **Pick SSIS Package** window appears.
5. Type the server name in the **Server** box. Optionally, you can select **Use SQL Server Authentication** and type your credentials in the **User Name** and **Password** boxes as appropriate. In the **Location** list, click **File system** or **SQL Server** and then click **…** next to the **Package Name** box to browse for the SSIS package. When you have found the package and clicked **Open** in the **Open** window, click **OK** and the editor of your choice opens.
6. Make any changes you wish to make in the editor, save the package and close the editor.
7. While you edit the SSIS package, Jet Data Managerdisplays the **Custom Code Editor** dialog. When you return to Jet Data Manager, click **Import** to import the changes you made to the SSIS package.
8. In the **Customize Code** window you'll notice that the Add command next to **SSIS Package** has changed to **Edit** and that you can now click **Parameters** and **Delete** as well. Click **Close**.
9. Right click the table and choose **Execute** to run the SSIS package. You can also execute the table by including it in an execution package, executing the entire project etc.

CLEANING UP THE DATABASE

To prevent accidental data loss, deleting a table in Jet Data Manager does not delete the physical table on the SQL Server. The downside is that tables deleted in Jet Data Manager still takes up space in the database. The SQL database cleanup feature enables you to identify tables left behind by Jet Data Manager and delete - drop - them to free up space. Note that database schemas are not deleted from the database. You will need to drop those manually in SQL Server after deleting them in Jet Data Manager.

Warning: Tables deleted with the SQL Database Cleanup feature cannot be recovered. Take care when you use the feature.

IDENTIFYING AND DELETING TABLES WITH SQL DATABASE CLEANUP

To use the database cleanup feature, follow the steps below.

1. In the project tree, right click a data warehouse, click **Advanced** and click **SQL Database Cleanup Wizard**. Jet Data Manager will read the objects from the database and open the **SQL Database Cleanup** window.
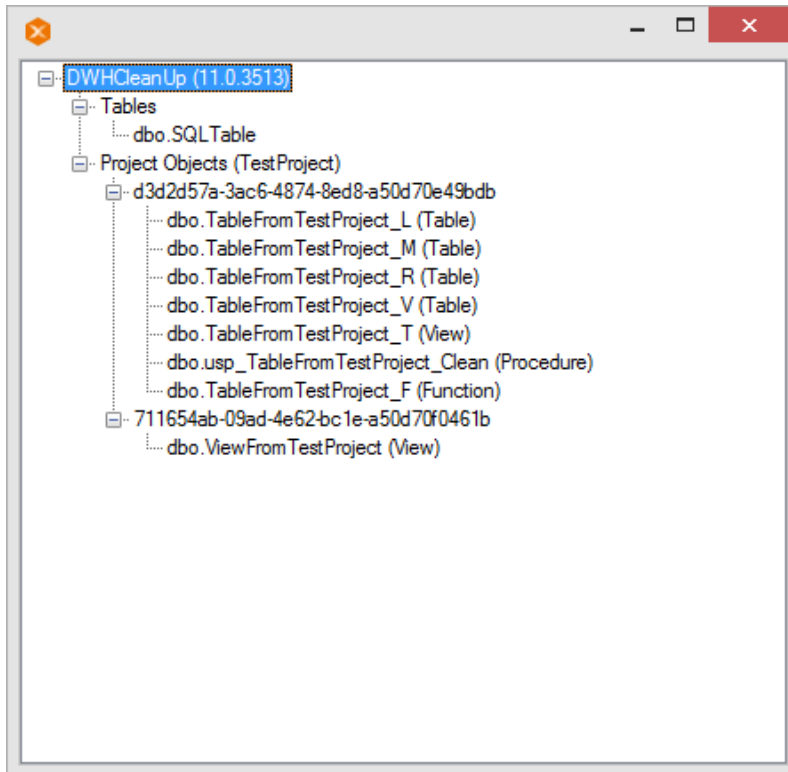
2. In the **SQL Database Cleanup** window, the content of the database is listed. Expand the root node, **[database name] ([SQL Server Version])**, to display the categories of content found on the second level. Depending on the content, the following categories are displayed:

   - **Tables**: Tables in the database unrelated to any Jet Data Manager project.
   - **Views**: Views in the database unrelated to any project.
   - **Procedures**: Procedures in the database unrelated to any project.
   - **Functions**: Functions in the database unrelated to any project.
   - **Deleted Project Object**: Objects deleted from the project, but not from the database. These should be safe to drop.
   - **Project Objects (<Unresolved Project: [project id]>)**: Objects in the database related to an unknown project, i.e. a project not in the current repository.

- **Project Objects ([name of project])**: Objects in the database related to a project in the current repository.

3. Expand Project Objects to display a list of object ids in the project. If you expand an object id, the tables, views etc. related to the object are listed. For example, if the object is a table, the valid, raw and other instances of the table are listed.

4. (Optional) Right click a table, view, procedure or function and click **Script** to display the SQL script behind the object.

5. Right click a table, view, procedure or table and click **Drop** to drop the object from the database. Click Yes, when Jet Data Manager asks you to confirm the drop. A message will tell you if the action succeeded or failed.

6. Right click an object id or a Project Objects item and click **Drop** to drop the object and all objects on the levels below. A window will open with a list of the objects that will be dropped. Clear the selection for any tables you wish to keep and click **Drop**. Note: Jet Data Manager will automatically clear the selection for any incrementally loaded tables to prevent accidental data loss. Jet Data Manager will ask you to confirm if you wish to drop an incrementally loaded table.

7. When you have dropped the all the objects you wish to delete from the database, close the window.

SELECTING, VALIDATING AND TRANSFORMING DATA

Selecting the right data from the source, validating it and transforming the data if needed are central parts of the data warehouse process.

In Jet Data Manager you specify data selection rules to ensure that only the data needed for your analysis is extracted from the data source to the staging database.

On the staging database, you perform data cleansing by applying validation and transformation rules to the data. This ensures that only valid data is loaded into the data warehouse.

However, you can also apply selection, validation and transformation rules on a data warehouse. This is useful when you have moved data from different business units into the data warehouse and want to ensure the validity of the consolidated data.

OPERATORS FOR SELECTING AND VALIDATING DATA

When defining a data selection or validation rule, you can use the operators listed below.

Values must be either integers or letters. You can also specify a list of values by entering comma-separated values.

Operator	Definition
Not Empty	Selects records where the value of a field is not empty or NULL
Equal	Selects records where the value of a field is equal to the specified value
Greater Than	Selects records where the value of a field is greater than the specified value
Less Than	Selects records where the value of a field is less than the specified value
Not Equal	Selects records where the value of a field is not equal to the specified value
Greater or Equal	Selects records where the value of a field is greater than or equal to the specified value
Less or Equal	Selects records where the value of a field is less than or equal to the specified value
Min. Length	Selects records that contain at least the specified number of characters
Max. Length	Selects records that contain no more than the specified number of characters
List	Selects records where the value of a field is equal to one of the specified comma separated values
Empty	Selects records where the value of a field is empty or NULL

Not in List	Selects records where the value of a field is not equal to one of the specified comma separated values
Like	Selects records where the value of a field is similar to the specified value. A percent sign (%) can be used as a wildcard. ABC% will return all records where the value in the specified field starts with ABC.
Not Like	Selects records where the value of a field is not similar to the specified value. A percent sign (%) can be used as a wildcard. ABC% will return all records where the value in the specified field does not start with ABC.

DATA SELECTION RULES

Data selection rules are used to specify a set of conditions that data extracted from the data source must satisfy. By applying selection rules, only the subset of data that you actually need is loaded into the staging database. You can also define data selection rules on data warehouse tables to further control the data loaded into the data warehouse.

You can add usage conditions to selection rules based on project parameters. This enables you to e.g. load less data in a development environment than in the production enviroment.

ADDING A DATA SELECTION RULE

1. In the project tree, expand **Business Units**, expand the business unit that contains the data source you wish to apply the selection rule to, expand **Data Sources** and then expand the relevant data source.
   - OR -
   In the project tree, expand **Data Warehouse**, expand the data warehouse that contains the table you wish to apply the selection rule to, and expand **Tables**.
2. Right-click the table you wish to add the selection rule to and click **Add Data Selection Rule**.
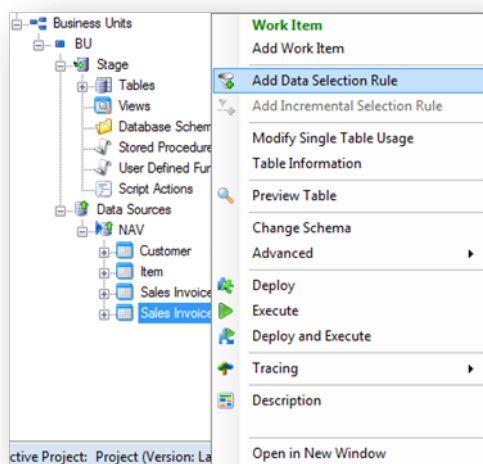
The **Data Selection** pane appears in the right hand side of the window.
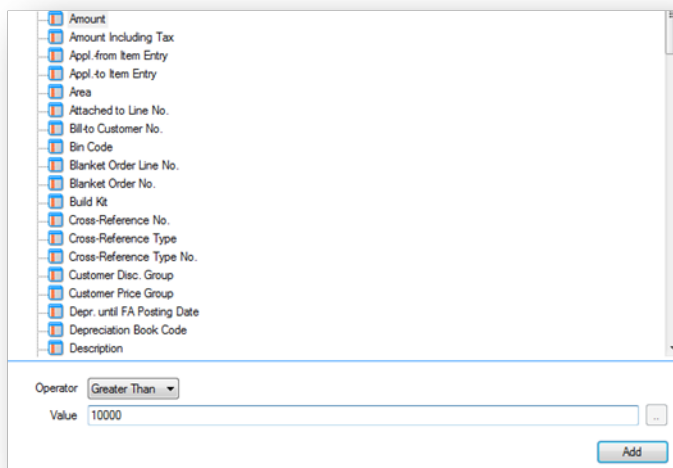
- Click the field you wish to use in the selection rule.
- In the **Operator** list, click the operator you wish to use. See Operators for Selecting and Validating Data
- If applicable, type a value for the operator in the **Value**box.

All selection rules that you have applied to a table are displayed in the project tree below the relevant table.

ADDING A USAGE CONDITION TO A SELECTION RULE

To add a usage condition to a selection rule based on a project variable, follow the steps below.

1. Right click a selection rule in the project tree and click **Add Usage Condition**. The **Usage Condition** panel is displayed in the right hand side of the appli-cation window.
2. In the **Usage Condition** panel, click the variable you wish to use.
3. In the **Operator** list, click the operator you wish to use. You have the following options:
    - Equal
    - NotEqual
    - GreaterThan
    - LessThan
    - GreaterEqual (Greater than or Equal to)
    - LessEqual (Less than or Equal to)
4. In the **Comparer** list, click the general data type of the variable, which Jet Data Man-ager will use in the comparison. You have the following options:
    - String

- Date
- Numeric

5. Type the value you wish to compare the parameter with in the **Value** box.
6. Click **Add** to add the usage condition.

For more information about project parameters, see [Project Variables](#).

DATA VALIDATION RULES

Validation rules ensure a high level of accuracy and reliability of the data in the data warehouse and are used to discover invalid data. You can apply validation rules at the field level in the staging database or at field level in the data warehouse.

While data is cleansed on the staging database, it often has to be cleansed again if you have consolidated data from different business units in the data warehouse.

You can make a validation rule conditional if you want the rule to apply in specific situations only.

For each validation rule you apply to a field, you must also classify the severity of a violation. The following classifications are available:

Severity	Definition
Warning	The violation is not critical to the data quality and does not require immediate attention. The data is considered valid and will still be made available to the end users.
Error	The violation is critical to the data quality and requires immediate attention. The data is considered invalid and will not be made available to the end users.

ADDING DATA VALIDATION RULES

You can add any number of validation rules to a field.

1. In the project tree, expand **Business Units**, expand the business unit that contains the data source you wish to apply the validation rule to, expand **Data Sources**, expand the relevant data source and expand the relevant table.
   - OR -
   In the project tree, expand **Data Warehouse**, expand the data warehouse that contains the table you wish to apply the selection rule to, expand **Tables** and expand the relevant table.
2. Right-click the field, you wish to apply the validation rule to, and click **Field Validations**.
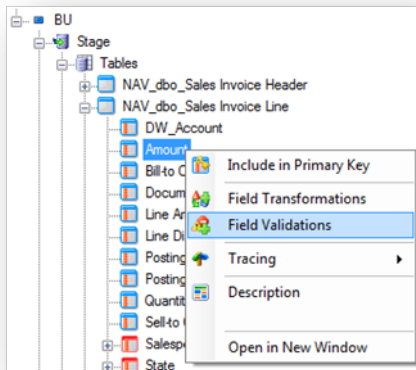
The **Field Validations** pane appears in the right-hand side of the window.

- Click the field you wish to use in the validation rule.
- In the **Operator** list, click the operator you wish to use. See Operators for Selecting and Validating Data
- If appliate, type a value for the operator in the **Value** box.
- Click **Error** to specify that as the severity level or leave it at **Warning**.

3. Click **Add** to add the rule.

ADDING CONDITIONS

You can add any number of conditions to your validation rules. Follow the steps below to add a validation rule.

1. Locate the rule you wish to modify.
2. Right-click the rule and then click **Add Condition**.
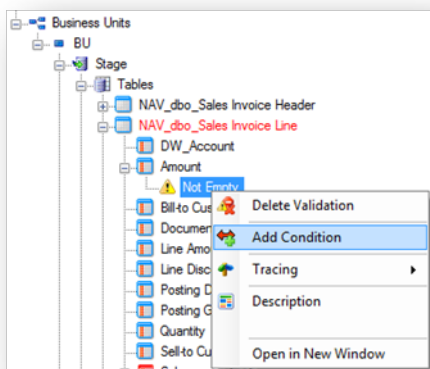
The **Condition** pane is displayed.
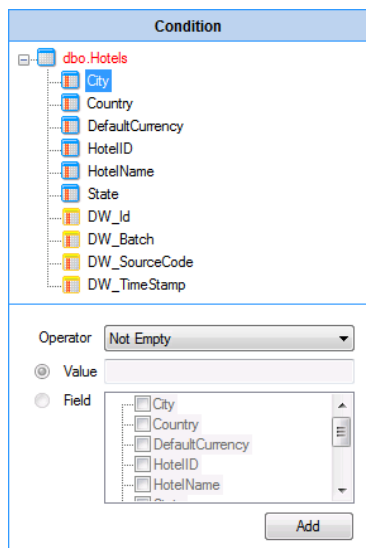
- In the **Operator** list, click the operator you wish to use. See Operators for Selecting and Validating Data
- In the **Value** field, type the value you wish to use in the comparison.
  - OR -

  Click **Field** and click the field, you wish to use in the comparison.
3. Click **Add** to add the condition to the rule.

The condition is displayed in the project tree below the validation rule or transformation rule it belongs to.

TO VIEW VALIDATION ERRORS OR WARNINGS

1. Click the **Errors** or the **Warnings** tab.
2. In the **Database** list, click the database that contains the table you wish to view errors or warnings for.
3. In the **Table** list, click the relevant table. The **No. of rows** box displays the number of errors or warnings in the table and the rows that violate the rules are displayed in the pane below.
4. Click any row to display the error or warning message in the **Error Message** or **Warning Message** box.

DATA TRANSFORMATION

Fields transformations lets you modify existing data in a number of ways. You can, for example, easily reverse the sign of numeric values, trim fields or return a specified number of characters from the original field value.

ADDING FIELD TRANSFORMATION RULES

1. In the project tree, expand **Business Units**, expand the business unit that contains the data source you wish to apply the validation rule to, expand **Data Sources**, expand the relevant data source and expand the relevant table.
   - OR -

In the project tree, expand **Data Warehouse**, expand the data warehouse that contains the table you wish to apply the selection rule to, expand **Tables** and expand the relevant table.

2. Right-click the field, you wish to add a transformation rule to, and then click **Field Transformations**.
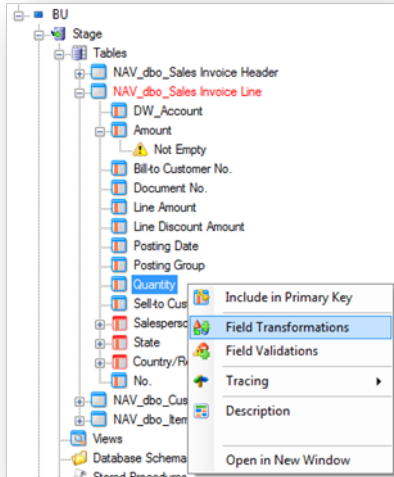
3. In the **Field Transformation** pane, click the field you wish to add a transformation to.
4. In the **Operator** list, click the operator, you wish to use, and then click **Add**.

Operator	Description
Upper	Converts all text values to upper-case
Lower	Converts all text values to lower-case
First	Returns the number of beginning characters specified by the user
Last	Returns the number of ending characters specified by the user
TrimLeft	Trims padded spaces from the left of the data
TrimRight	Trims padded spaces from the right of the data
Trim	Trims padded spaces from the left and right of the data
Fixed	Inserts a fixed value that is specified by the user
Custom	Allows for custom SQL code to be executed
ReverseSign	Reverses the sign for numeric values
TimeOnly	Returns only the time portion of a datetime field
DateOnly	Returns only the date portion of a datetime field
Replace	Replaces one set of characters with another

5. If you have selected **First** or **Last** as the operator, enter how many characters you wish to include in the **Length** field.

ADDING CONDITIONS

You can add conditions to transformation rules in the same way that you add conditions to validation rules. " Adding Conditions" on page 137

STRONGLY TYPED VALUES IN CUSTOM FIELDS

Strongly Typed Values are available for use in Custom Fields within a project. Strongly typed values closely coincide with other objects within the project. For example, if there is a field named **Sales Amount** that is used in a custom field that is not strongly typed, the custom field will cause an error if the **Sales Amount** field is renamed to something else. The reason for this is that the values are simply stored as text. Strongly typed values are inherently linked to the object it refers to, so if the object is renamed, it will be dynamically changed in the custom field as well.

USING A STRONGLY TYPED VALUE IN A CUSTOM FIELD

To use a strongly typed value in a custom field, follow the steps below.

1. Insert a custom field in a table. Right-click a table and click **Add Custom Field** (in a staging database) or **Add Field** (in a data warehouse).

    The **Add Field** window opens.
2. Type a Name for the field.
3. In the **Data Type** list, click the data type you wish the field to have and the adjust any settings for the data type you have selected.
4. Click**OK**.
5. Right-click the custom field you just added and click **Field Transformations**.

6. In the **Field Transformation** pan, click **Custom** in the **Operator** list and click**Add**.
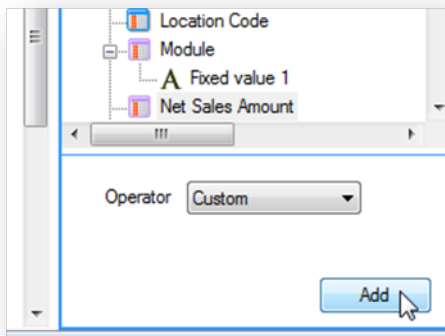
The **Transformation Custom SQL** window opens.
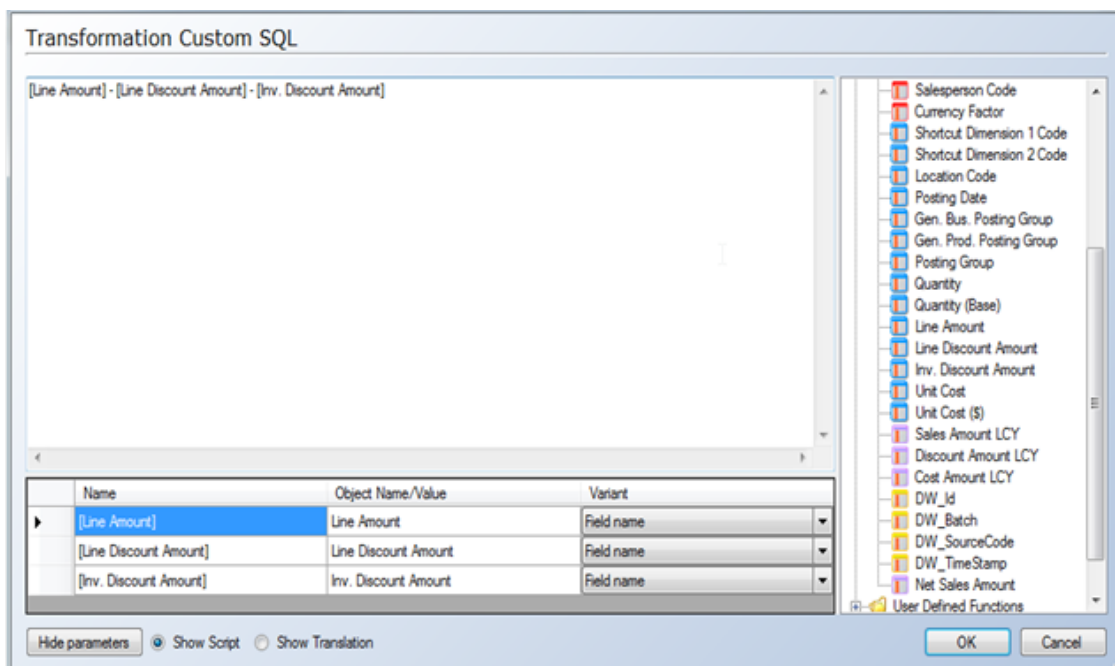
7. Drag fields from the pane in the right hand side of the window and drop them in the box in the middle of the window to add them to the script that calculates the field's value. As the fields are dropped, the parameters are automatically created at the bottom. Click **OK.**

If the underlying objects are renamed, they will dynamically be updated in the custom field without any user interaction required.

CONDITIONAL LOOKUP AND CUSTOM HASH FIELDS

In addition to the standard fields, Jet Data Manager has two special field types: conditional lookup fields and custom hash fields.

CONDITIONAL LOOKUP FIELDS

Lookup fields are used to add a field to a table in order to retrieve the value of the field in another table. The process of adding a conditional lookup field consists of a number of steps described below.

ADDING A CONDITIONAL LOOKUP FIELD

1. Expand **Business Units**, expand the preferred business units, and then expand the staging database.
2. Expand **Tables**, and then select the table you wish to modify.
3. Right-click the table, and then select **Add Conditional Lookup Field.**
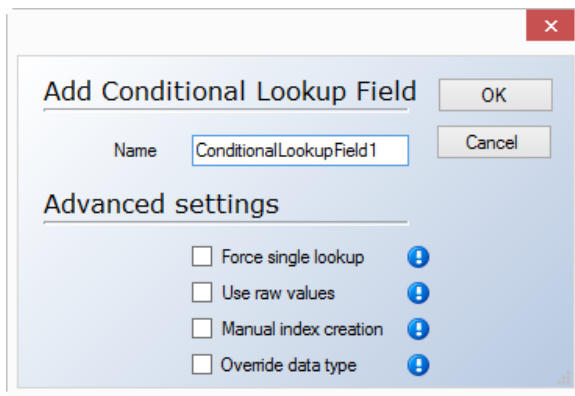
4. In the **Name** field, type a name for the lookup field.
5. Select **Force single lookup** to overwrite the default behavior where multiple different conditional lookups from the same table with the same join and aggregation type is grouped together and retrieved in one SQL statement. When this option is selected, lookups will always be performed in separate update statements. This carriers a performance penalty and as such, you should only enable this option for troubleshooting purposes.
6. Select **Use raw values** to perform the lookup on the raw values of the source table instead of the valid values, i.e. before any transformations or other cleansing tasks are performed. Lookups are always inserted into the raw destination table, and this setting does not affect that.
7. Select **Manual index creation** to prevent Jet Data Manager from automatically creating indexes on the source table on the fields specified in the join, with the lookup field(s) in the include clause. Enable this option if you wish create indexes manually on the source table.

8. Select **Override data type** to set the data type of the lookup field manually. As a default, Jet Data Manager will set the data type of the conditional lookup field to the same data type as the source field of first lookup field. If you enable this option, you can right-click the conditional lookup field, when it has been added, and click **Edit Data Type**. This option is useful if, for example, you are doing one conditional lookup on multiple fields with different data types that must have a target data type that is different than the data type of the first lookup field.
9. Click **OK**. The field is added to the project tree under the table.

SPECIFYING THE LOOKUP FIELD

The next step is to specify the lookup field that contains the values to be used in the field you just created. You can add multiple lookup fields to one conditional lookup field.

1. Expand the field, and then right-click **Lookup Field**.
2. Select **Add Lookup Field**.
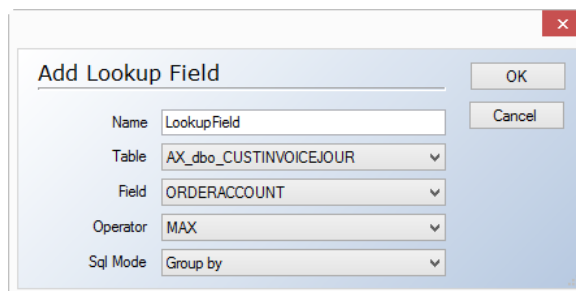
3. In the **Name** field, type a name for the field.
4. In the **Table** list, select the table containing the field you wish to use.
5. In the **Field** list, select the field you wish to use.
6. In the **Operator** list, specify how to return the values. You have the following options:

Option	Description
Top	Returns the value from the first record that matches the join criteria. When you select this operator, a **Sorting** node will be added to the project tree under the lookup field. Right click this and click **Add Sorting** to define how the matching values are sorted before they are retrieved from the source table.
Sum	Returns a sum of all the values that match the join criteria. This will only work on numeric values. Null values are ignored.

Count	Returns a count of all the values that match the join criteria. Null values are ignored.
Maximum	Returns the highest value of the values that match the join criteria. For strings, it will find the highest value in the collating sequence. Null values are ignored.
Minimum	Returns the lowest value of the values that match the join criteria. For strings, it will find the lowest value in the collating sequence. Null values are ignored.
Average	Returns the average value of the values that matches the join criteria. This will only work on numeric values. Null values are ignored.

7. (Optional) In the **SQL Mode** list, click the the mode you wish the generated SQL script to use. Usually, the default value will give you the best performance.
8. Click **OK**.

> **Note:** You can also drag a field from one table and drop it on the name of another table. This will automatically create the Conditional Lookup field with the exception of the joins, which are covered below.

ADDING JOINS

Next you have to add a join that specifies which join criteria must be met in the source table. Less complex joins will make the lookup perform faster. Complexity is a combination of the number of fields in the join and the data type. To get the best performance, use one single numeric field for the join.

1. Expand the lookup field, right-click **Joins**, and then select **Add Join**.
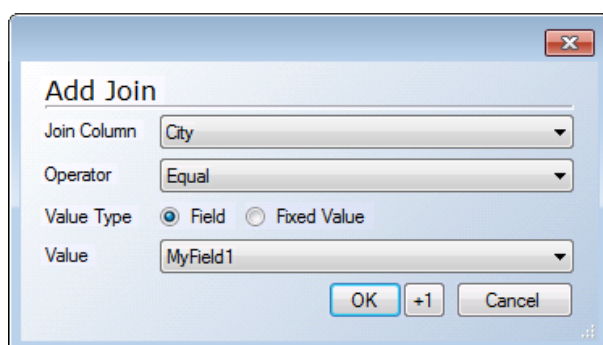
2. In the **Join Column** list, select the field that uses the lookup
3. In the **Operator** field, specify when to look up a value.
4. Click **Field** or **Fixed Value** to specify if you wish to compare the field selected in the join colum list to a field on the destination table or a fixed value. The **Value** box changes to fit your choice.
5. Depending on your chosen value type, click the relevant field in the **Value** list or enter

a value in the **Value** box.

6. Click **OK**.

SPECIFYING CONDITIONS

You can now specify conditions for when to lookup. The lookup will only be performed when the condition evaluates to true. For example, if you can determine that the lookup will only find related values when a certain field in the destination table has a certain value, apply a condition to avoid the lookup being performed on many records without finding a matching record. Conditions must also be used when having multiple lookup fields within one conditional lookup field to determine which lookup field to use. The first lookup field where the condition evaluates to true will be used, even if it returns a NULL value or finds no matching records. If no conditions are specified, the first lookup field will always be used and any subsequent lookup fields will be ignored.

1. Expand the lookup field, right-click**Conditions**, and click **Add Condition**.
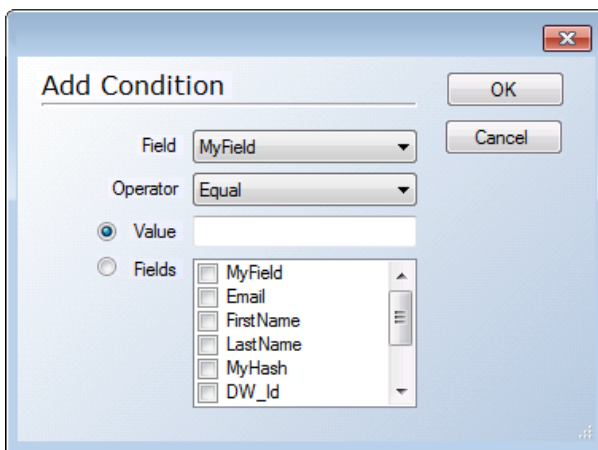
2. In the **Field** list, click the field that you wish to use in the comparison.
3. In the **Operator** list, click the operator you wish to use.
4. Click **Value** and enter a value to use in the comparison in the box
   - OR -
   Click **Fields** and select a field to use for the comparison in the list.
5. Click **OK** to add the condition.

CUSTOM HASH FIELDS

In scenarios with multiple fields making up the primary key, hashing the values of those fields into a single field can improve lookup performance. You can also use the field to easily investigate whether changes has been made to a record or not. In Jet Data Manager, such a field is called a custom hash field.

ADDING A CUSTOM HASH FIELD

1. Right-click a table and click **Add Custom Hash Field**. A custom hash field is added to the table and selected in the project tree. The Custom Hash Field is displayed in

the right hand side of the window.

2. In the **Custom Hash Key** pane, select the fields you wish to include in the custom hash field.
3. (Optional) Under **Sequence**, you can reorder the fields using drag-and-drop or by selecting a field and pressing **ALT** + **Up** or **Down**. The order of the fields is important since the port

VIEWS

A view is a virtual table in your data warehouse or in your staging database where you can group together information from two or more tables in your data source. Views can, for example, be used to provide a user with a simplified view of a table and to limit access to sensitive data. Creating views in Jet Data Managerfollows the same methodology as creating standard SQL views.

In Jet Data Manager you can create two types of views: views that consist of a subset of columns or rows in one or more tables and views that are joins of one or more tables.

CREATING VIEWS BASED ON LOOKUP FIELDS

Creating a view based on a lookup field consists of the following steps. The steps are all carried out in the View dialog.

CREATING A LOOKUP FIELD

1. On the **Data** tab, expand the chosen data warehouse or staging database. Right-click**Views** and then click **Add View**.

2. The Add View window appears.

3. In the **Name** field, type a name for the view. You can also click **Schema** to select a schema and **Outer Join** to use outer join in the view.
4. In the **Field Type** list, select **Lookup Field**. The dialog changes so that you can create and specify the properties of the lookup field:
   - Click the **Table** list and select the table that holds the lookup field.
   - Click the **Field/Function** list, select the field or function you wish to use and then specify which values to return. You have the following options:
     - **TOP** returns the value of the first record in the column.
     - **SUM** returns the sum of all field values in the column.
     - **COUNT** returns the number of records.
     - **MAX** returns the maximum value of the records in the column.
     - **MIN** returns the minimum value of the records in the column.
   - In the **Alias** field, type a name for the lookup field if you want the name to be different from the source field name.
5. Click **OK**. The field is displayed in the **View** pane of the window.
6. Click **New** if you want to create another field.
7. You have to specify a join between the view table and the lookup tables.

8. Click the **Lookup field** list and select the field to look up.
9. In the Operator field, select the operator that determines how you want the columns compared.

Operator	Description
Equal	Returns values that are equal
Greater	Returns values that are greater than the value of the lookup field or the specified fixed value
Greater or Equal	Returns values that are greater than or equal to the value of the lookup field or the specified fixed value
Less or Equal	Returns values that are less than or equal to the value of the lookup field or the specified fixed value
Less Than	Returns values that are less than the value of the lookup field or the specified fixed value
Not Equal	Returns values that are different from that of the lookup field or the specified fixed value

Note: A default inner join is created which only returns results from the rows common to the two joined tables. For the complete set of records from the joined tables, check the Outer Join box at the top of the dialog.

10. Click **OK**, and then click **New** if you want to specify a new join.
11. You also have to define the sort order.
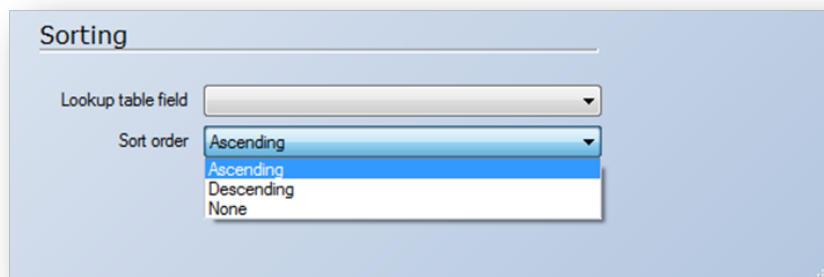
12. In **Lookup Table Field** list, click the field you prefer.
13. In the **Sort Order** list, select how you want the results sorted. Results can be sorted in either ascending or descending order.
14. Click **OK**, and then click **New** if you want to specify a sort order for another field in the view.
15. Once you have completed all steps, and created all the joins you need, click **Ok** in the upper right corner of the dialog to create the view.

TO CREATE VIEWS BASED ON STANDARD FIELDS

1. 1. On the **Data** tab, expand the chosen data warehouse or staging database , right-click Views, and then select **Add View**.
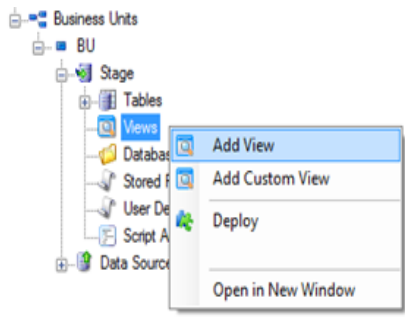
   The **New View** windows opens.
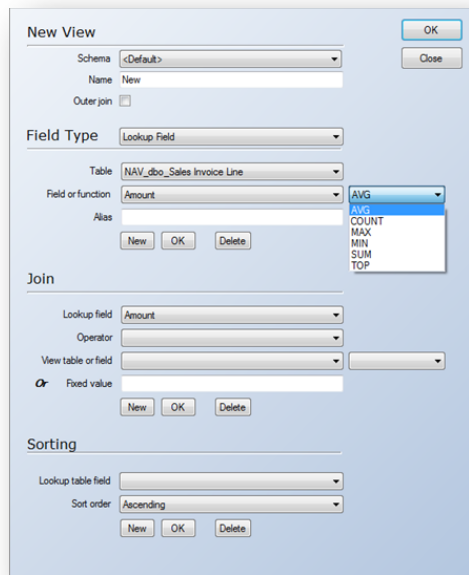
2. In the **Name**field, type a name for the view.
3. In the **Field** type field, select **Standard Table Field**.
4. In the **Table** list, select the table you want to retrieve data from.
5. In the **Field** list, select the field you want to use in the view.
6. In the **Alias** field, type a name for the view, and then click **OK**. The selected field is displayed in the **View** pane.
7. To add more fields from the same table or a field from another table, click **New**, and then repeat steps 3-6. Do this for all the tables you want in the view.

Note: If you want to add fields from more than one table, the tables must be related.

INDEXES

To achieve the optimal performance on your data warehouse, it is important to have the right indexes on your tables. Jet Data Manager can generate the necessary indexes automatically or assist you in creating indexes manually. You can also choose to use a legacy approach to indexes instead.

With the Index Automation feature, you can let Jet Data Manager handle all index creation and maintenance. Index Automation considers the following when designing indexes for the project:

- Relations between tables with relationship type set to Error or Warning
- Joins on conditional lookup fields
- Primary Key fields (On Raw Table)
- Selection Rules on the Data Warehouse
- Incremental Selection Rule on the Data Warehouse
- Partitioning fields (DW_Partitionkey, DW_TimeStamp)

Index Automation will try to minimize the number of indexes. If two lookups can use the same index, Jet Data Manager will take advantage of that. In addition to that, Jet Data Manager takes any manually created indexes into consideration. It will not change your manually created indexes, but it will use them instead of creating similar indexes. The indexes created by Index Automation will be named AutoIndex and postfixed with a number for uniqueness within each table.

SETTING UP INDEX AUTOMATION

Index Automation is configured on the project level, but can be overwritten on the individual table. The following options are available:

- **Automatic** (default): Index automation updates the indexes whenever the user changes the project in a way that could trigger a new or altered index.
- **Manual**: The user can have Jet Data Manager create indexes on selected tables. However, these indexes are not managed by Jet Data Manager. Nothing happens automatically if the table is changed in a way that impacts the indexes.
- **Disabled**: Jet Data Manager will use the legacy index generation behavior. Indexes will be generated during execution when needed by a data cleansing procedure. However, the same index might be created multiple times, since the index generation behavior is not tuned for performance. In addition to that, these auto-generated indexes are not visible for the end user in the project tree.

CONFIGURING INDEX AUTOMATION FOR THE PROJECT

To configure the Index Automation setting on the entire project, follow the steps below.

1. On the **Data** tab, right click the project node, and click **Edit Project**. The **Edit Project** window appears.
2. In the **Default Index Automation type** list, click the option you wish to use.
3. Click **OK**.

CONFIGURING INDEX AUTOMATION FOR A TABLE

To configure the Index Automation setting on a specific table, follow the steps below.

1. On the **Data** tab, right click the table and click **Table Settings**. The table settings window appears.
2. On the **General**tab, in the **Index Automation** group, click the option you wish to use.
3. Click **OK**.

MANUAL INDEX GENERATION

Setting the index automation setting to manual, makes it possible for you to use the index generation features of Jet Data Manager while maintaining complete control over the indexes in your project. When you run manual index generation on a table, data warehouse, staging database or the project, Jet Data Manager creates any indexes Index Automation finds necessary. However, you can delete and edit indexes as you see fit. Jet Data Manager will not create new indexes on the tables unless you run manual index generation again.

GENERATE INDEXES MANUALLY ON THE PROJECT, A DATA WAREHOUSE OR A STAGING DATABASE

To generate indexes on a data warehouse, staging database or the entire project, follow the steps below.

1. Set or make sure Index Automation is set to manual. See Configuring Index Automation for the Project.
2. Right-click the project, data warehouse or staging database you wish to use manual index generation on, click **Advanced** and click **Index Automation(manual)**.
3. A message will appear to tell you how many tables Jet Data Manager checked. Click **OK**.

GENERATE INDEXES MANUALLY ON A TABLE

To generate indexes on a specific table, follow the steps below.

1. Set or make sure Index Automation is set to manual. See Configuring Index Automation for the Project.
2. Right-click the table you wish to use manual index generation on, click **Advanced** and click **Index Automation(manual)**.

Note: If the table you wish to add an automatic index to already has one or more indexes, the **Index Automation (manual)** option is not available in the **Advanced** menu. Instead, expand the table, right click **Indexes** and click **Index Automation (manual)**.

SCRIPTING

Jet Data Manager generates most of the code you need, but you can extend the functionality of Jet Data Manager by writing your own scripts. When you need to include custom SQL code in your project, you have different options depending on what you would like to achieve.

- **User Defined Functions** and **Stored Procedures** are used to create reuseable code on SQL Server. Jet Data Manager uses them when it generates the code for executing your project. You can create you own User Defined Functions and Stored Procedures and call them from [Execution Packages](#) or Script Actions.
- **Script Actions** enables you to add snippets of SQL code to be run before or after each step in the deployment or execution of a table.
- **SQL Snippets** are small snippets of parameterized SQL code you can use on the field level.
- **Custom Code** lets you replace the code generated by Jet Data Manager for deployment and execution with your own code written in your favorite IDE.

You can also create global **Project Variables** for use in your scripts.

USER DEFINED FUNCTIONS AND STORED PROCEDURES

User Defined Functions allow you to define your own Transact SQL functions. A user defined function returns a table or a single data value, also known as a scalar value. You can, for example, create a function that can be used to perform complex calculations.

Stored Procedures allow you to define you own Transact SQL stored procedures. You can, for example, create a stored procedure that can be called from the execution package.

ADDING A USER DEFINED FUNCTION

You can add user defined functions to both the staging database and to the data warehouse.

1. Under a data warehouse or a staging database in the project tree, right-click**User Defined Functions** and click **Add User Defined Function**.

2. The **Add User Defined Function** window will open:
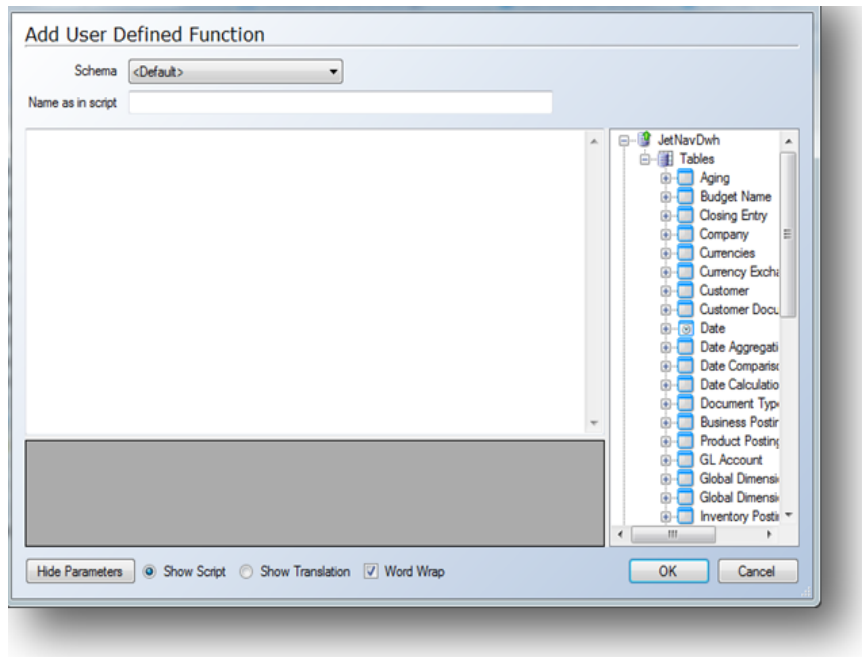
- In the **Name as in script** box, type a name for the function.
- In the text box, enter the SQL code for the user defined function. You can include tables from the data warehouse or staging database in question by dragging them in from the pane to the right.

3. A **Script Action** can then be created, if necessary, to call the User Defined Function.

ADDING A STORED PROCEDURE

You can add stored procedures to both the staging database and to the data warehouse.

1. Under a data warehouse or a staging database in the project tree, right-click **Stored Procedures** and then click **Add Stored Procedure**.

2. A window with the following options will appear:
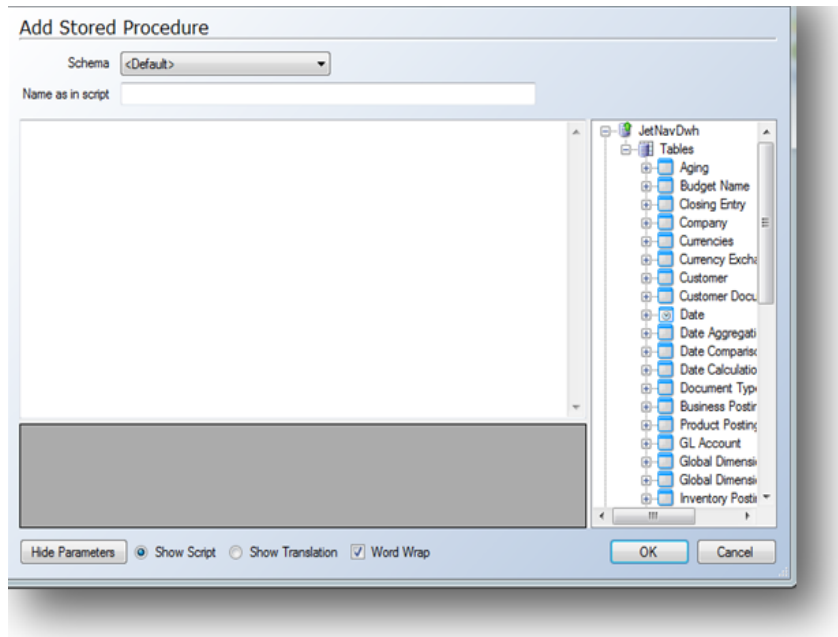
- In the **Name as in script** box, type a name for the function.
- In the text box, enter the SQL code for the stored procedure. You can include tables from the data warehouse or staging database in question by dragging them in from the pane to the right.

3. A **Script Action** can then be created, if necessary, to call the stored procedure.

SCRIPT ACTIONS

Script Actions are SQL scripts that can be executed along with deployment or execution of a table to complete a number of different tasks. A Script Action can utilize the User Defined Functions and Stored Procedures that you have already created.

ADDING A SCRIPT ACTION

1. Under a **Data Warehouse** or a **Staging Database** in the project tree, right-click **Script Actions** and click **Add Custom Step.**

2. The **Edit Custom SQL Script** window will open. Type a **Name** and type or paste your SQL script into the window and click **OK.**
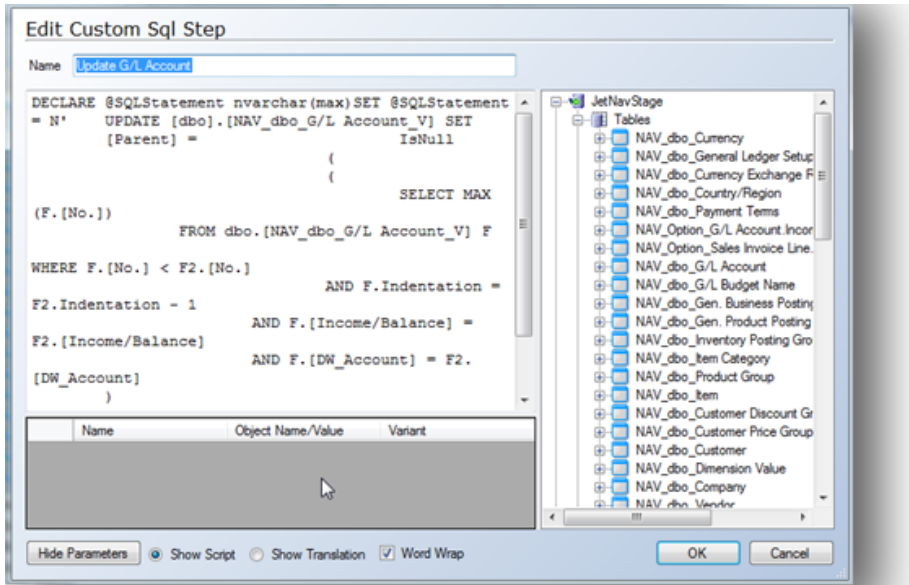
ASSOCIATING A SCRIPT ACTION WITH A TABLE

1. Right-click the table you wish to associate with the **Script Action**, click **Advanced** and click **Set Pre- and Post Scripts**.
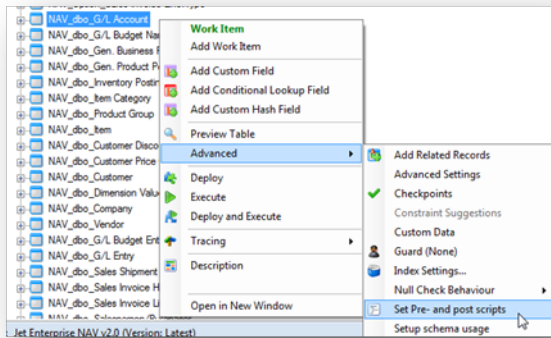
2. The **Set Pre- and Post Scripts** window will open. The deployment and execution steps are listed. Click **< None >** in the **Pre Step** or **Post Step** column of the step you wish to add you script to and click the name of the script. Click **OK.**

The possible execution steps at which to call the script are:
- **Transfer Pre Step:** This will cause the script to be called prior to the beginning of the data transfer process which will move data into the table.
- **Transfer Post Step:** This will cause the script to be called after the transfer of data into the table has been complete but prior to the beginning of the data cleansing process
- **Data Cleansing Pre Step:** This will cause the script to be called after the transfer of data into the table has been complete but prior to the beginning of the data cleansing process
- **Data Cleansing Post Step:** This will cause the script to be called after the data cleansing process has completed.

SQL SNIPPETS
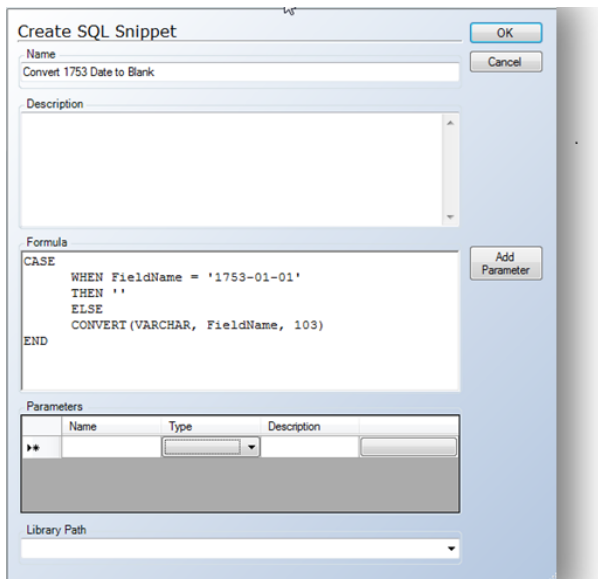
With SQL Snippets, you can create reusable, parameterized pieces of SQL code for use in in field transformations. This enables you to write the code once and use it on multiple fields, saving you the trouble of maintaining the same functionality on a many fields.

HOW TO CREATE A SQL SNIPPET

1. In the ribbon, click the **Tools** tab, and click **Create SQL Snippet** in the **Business Functions and SQL Snippets** group.


2. A window with the following options will appear:

- Type a **Name** and a **Description** for the SQL Snippet. The **Description** is optional but lets other users know what the Snippet does.
- Type the SQL command in the **Formula** box.
- For any variables (in this example, **FieldName**), highlight the variable, and click **Add Parameter**. This will add the highlighted text as a parameter name under

Parameters.
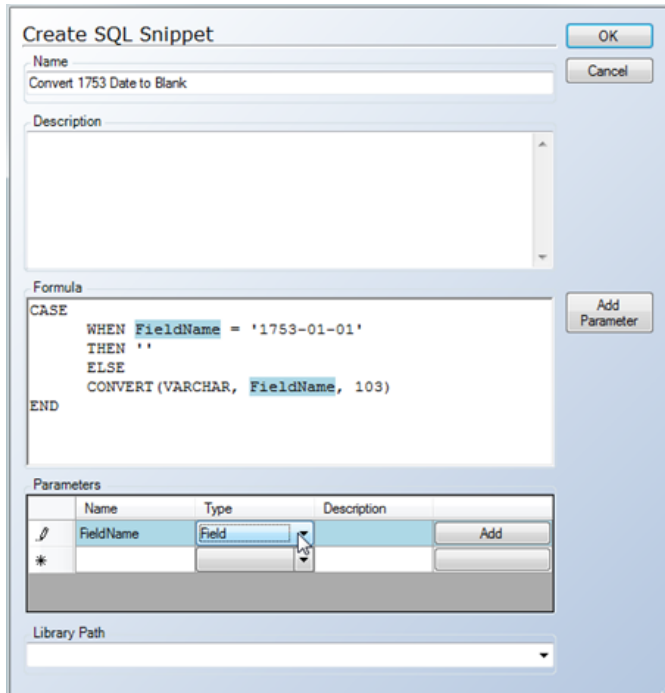
- Under Parameters, change the **Type** to match what the variable represents. The available options are:
  - Table
  - Field
  - Database
  - User Defined Function
  - Stored Procedure
  - Value

3. Click **OK** to save the SQL Snippet.

HOW TO EDIT A SQL SNIPPET

1. In the ribbon, click the **Tools** tab, and click **Edit SQL Snippet** in the **Business Functions and SQL Snippets** group.

2. Select the SQL Snippet to be edited from the list, and click **OK**.
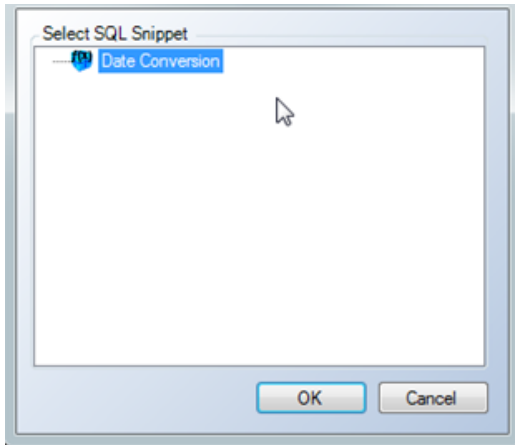
3. Adjust the SQL Snippet as needed.
4. Click **OK** to save the SQL Snippet.

HOW TO DELETE A SQL SNIPPET

1. In the ribbon, click the **Tools** tab, and click **Delete SQL Snippet** in the **Business Functions and SQL Snippets** group.

2. Select the SQL Snippet(s) to be deleted. Alternatively, you can select **Select All** to select all SQL Snippets.
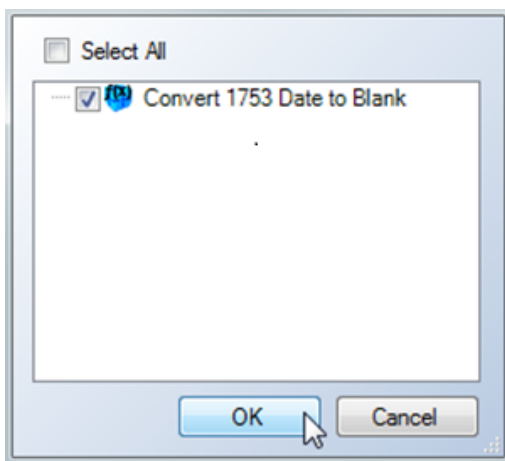
3. Click **OK.**
4. Click **Yes** to permanently delete the SQL Snippet(s).

HOW TO USE A SQL SNIPPET

1. Right-click the field to add the SQL Snippet to. Go to **Add SQL Snippet Transformation**, and select the desired SQL Snippet from the available list.
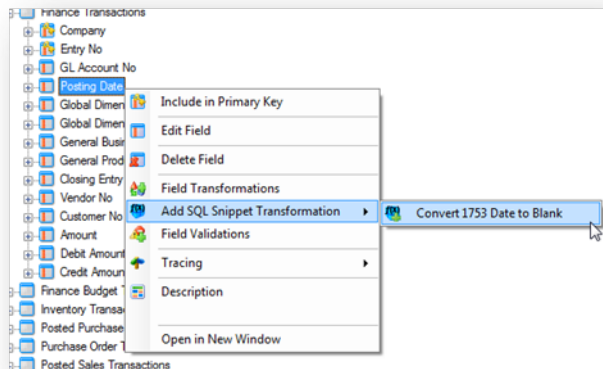
2. Drag the desired field(s) from the **Data Fields** pane on the right, and drop the field on the **Object Name/Value** column for the desired variable. The **Object Name/Value** column and **Variant** column will populate automatically.
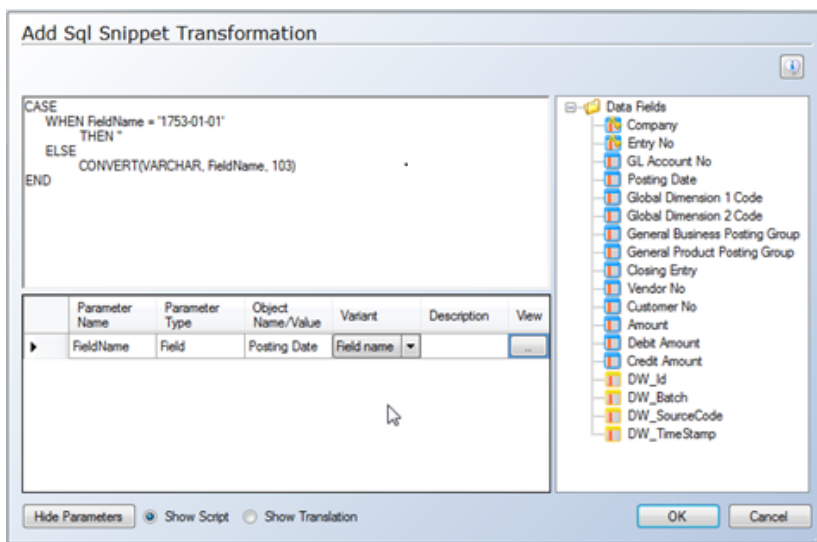
3. Click **OK**.

CUSTOMIZED CODE

Jet Data Managerenables you to integrate "hand-written" code into a project by customizing the data cleansing procedure, transformation view and SSIS package on a given table.

ADDING CUSTOMIZED CODE TO A TABLE

To customize the code on a given table, follow the steps below:

1. Right click the table in question, navigate to **Advanced** and click **Customize code**. The **Customize Code** window opens.
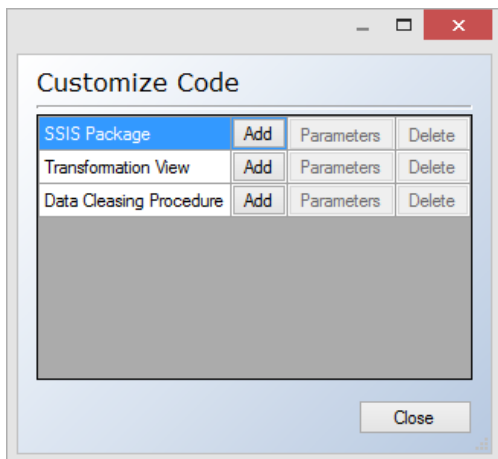
2. Click the **Add** button to the right of the step you wish to customize. The **Choose Editor** window opens.

3. In the **Editor Name** list, you have the following options:
   - **Standard** is the basic built-in editor in TX.
   - **Default File Program** is the program that is set to open files of the type in question. For the data cleansing procedure and the transformation view, the filename extension is .sql. For SSIS packages, the filename extension is .dtsx.
   - Any custom editors you have added (see [Managing Custom Editors](#)).

   If you are adding a SSIS package, the **Custom SSIS** window appears. Chose **Create Default Package** to edit the standard package, **Create Destination Only** to create a package that only contains the destination and **Existing Package** to import an existing package from the file system or an SQL Server.

   > **Note:** Some tables use multiple SSIS packages. When creating the Default package, Jet Data Manager will create the first SSIS package only. Examples of tables that will have multiple SSIS packages as default: Data warehouse tables that receive data from multiple staging tables, data source tables from NAV adapters with multiple companies, any data source table when template data sources are used.

4. If you chose the **Standard** editor, the **Edit** window opens. When you have finished editing the code, click **OK** to confirm you edits.
   If you chose a custom editor, Jet Data Manager will open the code in editor you

chose. When you have finished editing the code, save your changes and close the editor. Back in Jet Data Manager, the **Custom Code Editor** window is open.
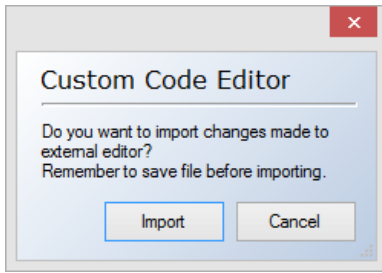

Click **Import** to import the changes you have made in the custom editor into your project.

5. When you return to the **Customize Code** window you will notice that you can now click **Parameters** (if applicate) and **Delete**. Click **Delete** to remove the customization and return to having Jet Data Manager generate the code. Click **Parameters** to decide which parameters are sent to the code on execution.

6. Click Close to close the window.

Note: When editing the data cleansing procedure or the transformation view, make sure to have a "create procedure" or "create view" declaration in the code with the exact same name as Jet Data Manager would have sued. This is what is called during execution. To be sure, simply keep the first line of the code generated by Jet Data Manager.

MANAGING CUSTOM EDITORS

To add, edit or delete a custom editor, click the **Tools** tab in the ribbon and then click **General Settings** in the **Administration** group. The **General Settings** window appears. Click **Custom Editors**.
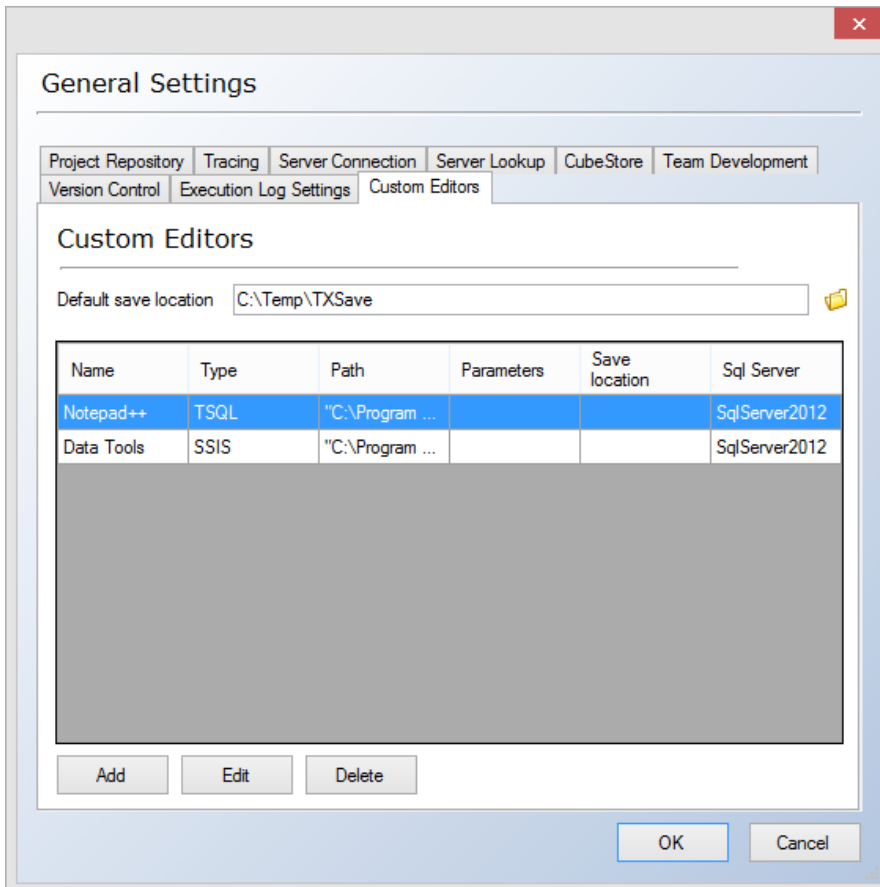
The list of custom editors is displayed. In the **Default save location** box, you can type the path to the folder where the custom code files are temporary stored (or click the folder icon to open a browse dialog).

To edit the settings for a custom editor, select the editor in the list and click **Edit**.

To remove a custom editor from the list, select the editor and click **Delete**.

To add a custom editor, follow the steps below:

1. Click **Add**. The **Add custom editor** window appears.
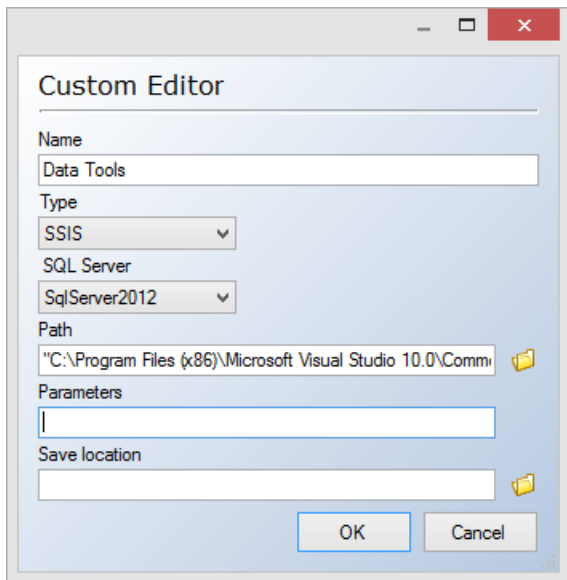
2. In the **Name** box, type a name for the editor.
3. In the **Type** list, click the type of editor you wish to add. Choose TSQL if you wish to use the editor with data cleansing procedures and transformation views and SSIS if you wish to use it with SSIS packages.
4. In the **SQL Server** list, select the SQL Server version that you are using. Currently, this setting is only used for custom editors for SSIS packages. When you want to customize the code for a SSIS package, Jet Data Managerchecks what version of SQL Server the table is stored on. You will only be able to select editors that are marked compatible with that version of SQL Server.
5. In the **Path** box, type the path that Jet Data Manager should call to start the program (or click the folder icon to open a browse dialog)
6. In the **Parameters** box, type any additional parameters for the program.
7. Optionally, in the **Save Location** box, type as save location for the editor (or click the folder icon to open a browse dialog).
8. Click **OK** to add the custom editor.

PROJECT VARIABLES

Project variables allows you to save information in project-wide variables. This is useful when you need to distinguish different environments in a script or

The value of a given variable is determined when you deploy the object that uses the variable. As such, when you have changed a variable it is important to deploy the objects that uses this variable. The exception is when you use project variables with customized code. Here, the value of the variables is determined on execution.

The variable does not have a specific data type. If, for instance, you wish to use the variable as a string, make sure to enclose the variable in quotes in the script.

ADDING A PROJECT VARIABLE

To add a new project variable, follow the steps below.

1. In the project tree, right-click your project and click **Project Variables**. The **Project Variables** window opens.
2. Click **Add**. The **Add New Variable** window appear.
3. In the **Name** column, type a name for the variable.
4. In the **Type** list, click the variable type you wish to use. You have the following options:
    - **Fixed:** A fixed string.
    - **System:** One of the following system properties:
        - MachineName
        - EnvironmentName
        - UserName
        - UserDomainName
    - **Source Scope:** A property of the source of the current object. For instance, if you use a source scope variable in a custom transformation rule on a table in the data warehouse, the variable will have the value of property on the relevant staging database. Since different possible sources have different properties, the variable might not always have a value. Examples of properties include Database Name, API Version, Host, File Name.
    - **Destination Scope:** A property of the destination of the current object, similar to source scope..
    - **Contextual Scope:** A property of one specific element in the project, such as database name on a particular staging database.
    - **Dynamic:** The value of the variable is generated by a custom script you written.
5. If you are adding a dynamic variable, in the **Resolve Type** list, select when you want to resolve the value of the variable. You have the following options:
    - **Every Time:** Resolve the value every time the value is used.
    - **One Time:** Resolve the value when the variable is used for the first time and reuse the resolved value for the following uses until the project is closed.
    - **Each Batch:** Resolve the value once for each batch, e.g. an execution.
6. If you are creating a Contextual Scope variable, click the object you wish to use in the **Context** list.
7. If you are adding a Source or Destination Scope variable, click the Value Filter list and click the type of object you wish to see available properties for in the **Value** field.
8. If you creating a fixed variable, enter the value of the variable in the **Value** field.
   - OR -
   If you are creating a dynamic variable, click **Script Editor** to open the standard script editor in Jet Data Manager and write the script that generates the value.
   - OR -
   If you are creating a variable of a type other than fixed or dynamic, in the **Value** list, click the property you wish to use as a value for the variable.
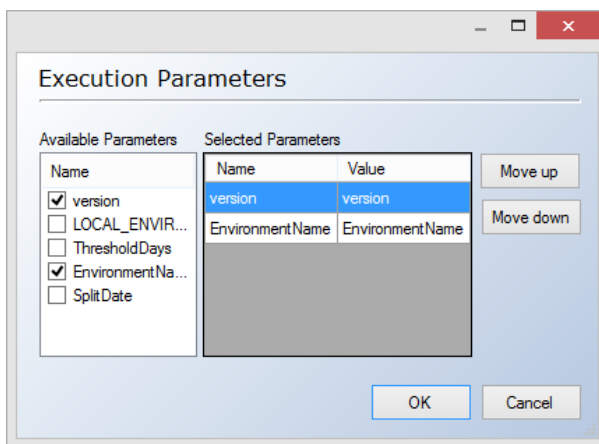9. Click **OK**.

USING A PROJECT VARIABLE IN A SCRIPT ACTION, CUSTOM TRANSFORMATIONS AND CUSTOM VIEWS

Project variables are available for use when writing script actions as well as custom transformation rules and custom views. The available variables are listed in the tree view in the right hand side of the editor window. Simply drag the variable in from the tree to use it.

USING A PROJECT VARIABLE WITH CUSTOMIZED CODE

Project variables are available as parameters when using the customized code feature. To make a project parameter available in a customized step, follow the steps below.

1. On the **Data** tab, right click a table with customized code. Click **Parameters** next to the step in which you would like to use a project parameter. The **Execution Parameters** window opens.

2. In the **Available Parameters** list, select the variables you wish to have available.
3. Click **OK**.

DATABASE SCHEMAS

Database Schemas allow you to apply a certain schema to a table or a group of tables. The main purpose of this is to restrict access to tables that report designers do not need, thereby making reporting off of the data warehouse easier.

CREATING A DATABASE SCHEMA

1. On a data warehouse or business unit, right-click **Database Schemas** and click **Add Database Schema**.
2. In the **Name** box, enter a name for the new schema. In the **Owner** box, you can enter the owning role for the schema. The default is dbo. Click **OK** to create the schema.
3. Assign a Schema Behavior by right-clicking on the newly created schema. You have the following options:
   - **None**: The schema will be applied to the tables you manually assign it to.
   - **Main default schema**: The schema will be applied to all tables and views in the region (data warehouse or staging).
   - **Main Raw default schema**: The schema will be applied to all Raw (_R postfix) tables in the region (data warehouse or staging).
   - **Main Transfer default schema**: The schema will be applied to all Transfer (_T postfix) views in the region (data warehouse or staging).
   - **Main Valid default schema**: The schema will be applied to all Valid (_V postfix) tables and views in the region (data warehouse or staging).
   - **Main Error/Warning default schema**: The schema will be applied to all Link and Message (_L and _M postfix) tables in the region (data warehouse or staging).
4. If you have selected **None** as the **Schema Behavior**, you need to assign the schema manually. Right-click the table, click **Table Settings** and click the **Table Schemas** tab. Here, you can then select a schema as **Default** (all instances of this table), **Raw**, **Transformation**, **Valid** or **Error/Warning**.
5. Assign user rights to the schema. This can be done through SQL Server Management Studio or T-SQL. See this article on the Microsoft website for details on how to grant user rights using T-SQL: http://msdn.microsoft.com/en-us/library/ms187940.aspx

DATA SECURITY

In Jet Data Manager, you can control access directly on the data warehouse. You can restrict access to specific views, schemas, tables and columns on tables - object level permissions - or specific data in a table - data level permissions.

The access control features can be found under the **Security** heading in the project tree on the **Data** tab.

ADDING A DATABASE ROLE

Object level security is based on SQL Server database roles. A user has access to an object if he is a member of a database role that has access to that object. To add a database role, follow the steps below.

1. On the **Data** tab, in the project tree, under Data Warehouses, the relevant data warehouse and **Security**, right click **Database Roles** and click **Add Database Role**. The **Database Role Setup** window opens.
2. In **Name**, type a name for the role.
3. If you are using the Multiple Environments feature: In the **Member** setup list, click **Environment Specific Role Members** if you wish to have a different setup for different environments. The different environments will then each have a tab in the list below.
4. Next, you should add users to the role. Click **Add login** to role to add Active Directory or SQL Server logins that are known to the SQL Server that the data warehouse resides on. The **Select Login(s)** window opens. Select the logins you wish to add in the list and click **OK** to add the user(s).
5. Click **Add manually** to add Active Directory or SQL Server users or groups that the SQL Server does not know, e.g. users on a production server. The **Enter User or Group ID** window opens. In **ID**, type the user or group id. Under **Type**, click **AD user-/group** or **SQL user/group** depending on the type of ID you entered. Click **OK** to add the user.
6. Click **OK** to close the window and add the database role, which is listed under **Database Roles** in the project tree.

If you need to add a new login on the SQL Server, you can right click **Security** in the project tree and click **SQL Server Logins**. Here, you can add logins if you have the necessary permissions on the SQL Server. However, for safety reasons, you cannot delete users here.

On each deployment, Jet Data Manager drops existing roles on the database before recreating them. By default, Jet Data Manager only drops database roles related to the deploying data warehouse. However, you can also set Jet Data Manager to drop more database roles with a setting on the data warehouse. To access the setting, right click the data warehouse, click **Edit Data Warehouse** and click **Advanced…** In the **Drop Role Option** list,

click **Roles Created by Application** to drop al roles created by Jet Data Manager or **All Roles on Database** to drop all roles altogether.

ASSIGNING OBJECT LEVEL PERMISSIONS TO DATABASE ROLES

You can assign permissions to database roles on the object level. Jet Data Manager uses the same allow/deny concept as SQL Server with three possible states:

- **Not set** (grey dot): The database role is not allowed to access the object, but are not explicitly denied access.
- **Grant** (green with white checkmark): The database role is granted access to the object. However, if a user is a member of another database role that is denied access, he will not be able to access the object.
- **Deny** (red with white bar): The database role is denied access to the object. Even if a user is a member of another database role that is allowed access, he will still be denied access.

In addition to the three states described above, a table can have different mixed states depending on the column level permissions set on the table. The mixed states are:

- **Partially Granted** (green and grey icon). The database role is granted access to some columns on the table. Note that you will also see this icon if the database role is granted access to all columns on a table since this will not automatically set Allow on the table level.
- **Partially Denied** (red and grey icon): The database role is denied access to some columns on the table. Note that you will also see this icon if the database role is denied access to all columns on a table since this will not automatically set Deny on the table level.
- **Mixed Grant/Deny** (red and green icon): The database role is granted access to some columns and denied access to other columns on the table.

To assign object level permissions, or column level permissions on tables, to database roles, follow the steps below.

1. On the **Data** tab, in the project tree, under **Data Warehouses** and the relevant data warehouse, right click **Security** and click **Object Security Setup**. The **Object Security Setup** window opens.
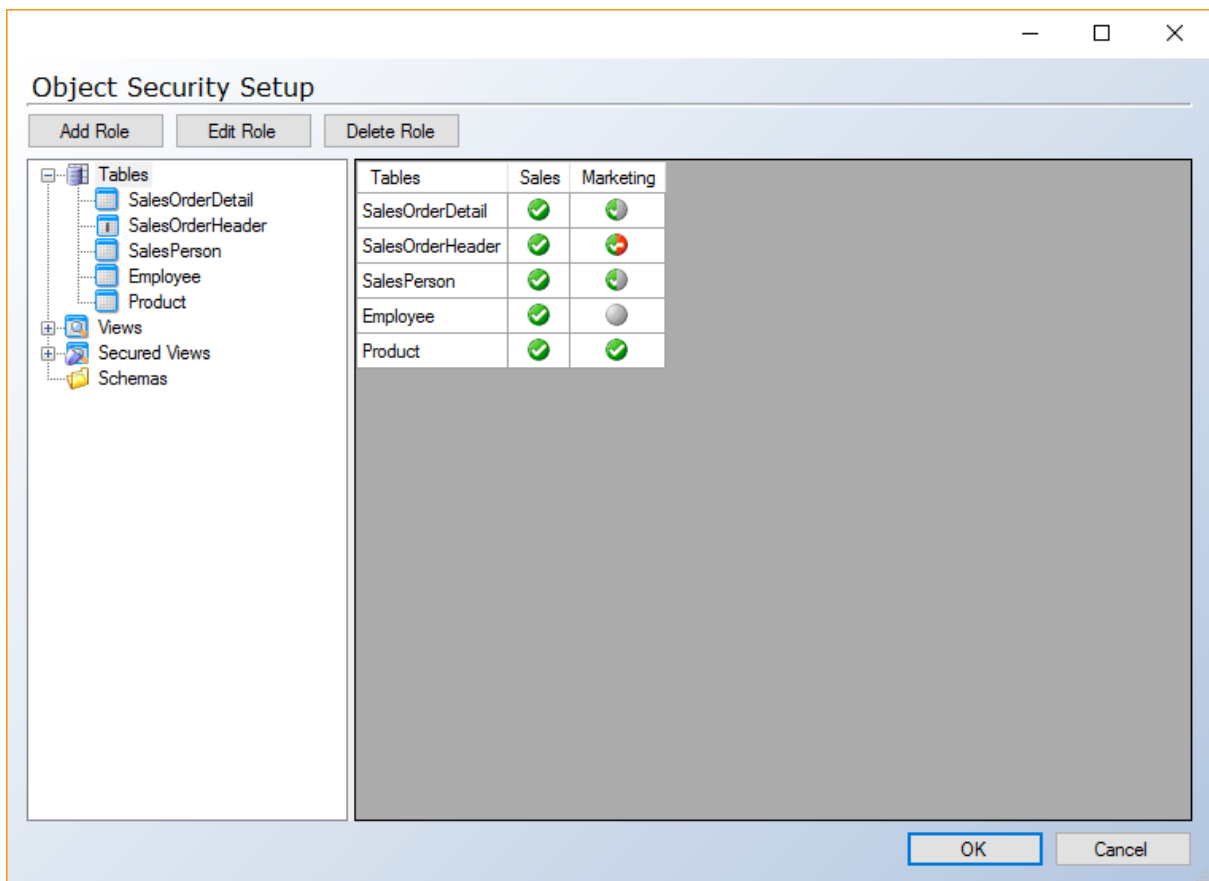
2. Click **Tables**, **Views** or **Schemas** in the left-hand column to choose the type of object you wish to set up access for. Expand **Tables** and click an individual table to assign column level permissions for that table.

3. In the right-hand column, the table shows object names in the left-most column and database roles in the following columns. Click icon in the intersection between the object name and the database role to change the permission for the database role on that object. If you set column level permissions on a table, this will overwrite any current object level permissions set and vice versa.

4. (Optional) Click **Add Role**, **Edit Role** or **Delete Role** to add, edit or delete database roles as needed.

5. Click **OK** to save changes and close the window.

ASSIGNING DATA LEVEL PERMISSIONS

In addition to configuring access on the object level, you can filter the data available to individual Active Directory users or SQL Server database roles. You might, for instance, want a sales person to be able to see all sales data, but only in his or her own region.

Data level security in Jet Data Manager is based on the concepts of securable columns, securable column setups, secured columns and secured views. This design allows you to create one security model and reuse it on any number of tables.

- A securable column contains the values that we wish to use in a filter. Continuing the example above, it could be "sales region id" in a "sales regions" table.
- A securable column setup is a mapping between securable column values and users or database roles, e.g. what "sales region id" does the sales person have access to. Each securable column can have multiple securable column setups.
- A secured column is a column on the table containing the data we wish to filter. This could be a "sales region id" column on a "sales transactions" table.
- A secured view is a view where all the data, that the user does not have access to, is filtered out. For instance, all the "sales transactions" rows where the "sales region id" does not match the "sales region id" the sales person has access to. When using data level security, the secured view should be used for reporting instead of the table it secures.

To assign data level security to a table, follow the steps below.

1. On the **Data** tab, navigate to and expand the table that contains the column you wish base the permissions on, right click the field and click **Add Securable Column**. The **Add Securable Column** window opens.
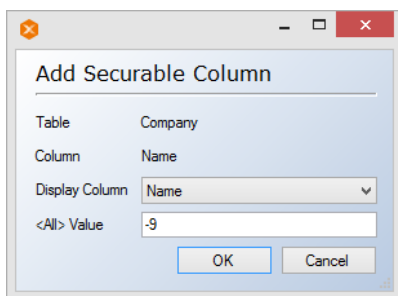
2. (Optional) In the **Display Column** list, click the column value you wish to dis-play instead of the column you are adding as a securable column. If the securable column contains e.g. an ID, it might be helpful to choose something that is easier to under-stand, e.g. a name, as the display column.
3. (Optional) In the **<All> Value** box, type a value that will be used to indicate all values. This value should be a value that is guaranteed not to be among the values in the securable column.

4. Click **OK**. The **Add Securable Column Setup** window opens to let you add your first **Securable Column Setup**.
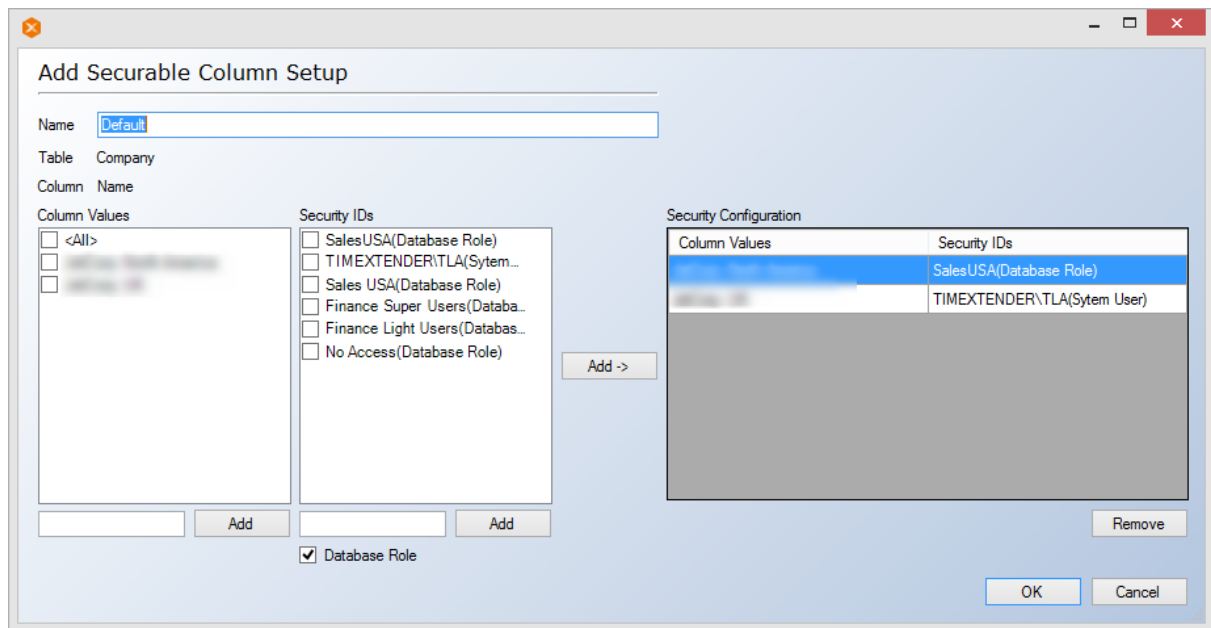
5. In **Name**, type a name for the securable column setup.
6. In the left-hand side of the window, you combine the values in the securable column with the users or database roles that should have permission to access the data. Select a number of values in the **Column Values** list and one or more users or database roles in the **Security ID** list and click **Add->**. The resulting pairs are displayed in the **Security Configuration** list.
7. (Optional) If you need to assign permissions to a value or a security ID that is not in either list, type the value or name in the box under the appropriate list and click **Add**. Select **Database Role** if you wish to add a database role as opposed to an Active Directory user. For more information on database roles, see Adding a Database Role.
8. Click **OK** when you have finished configuring the securable column setup. The securable column setup can be found in the project tree under **Security**, **Securable Columns**, **[table name]**, **[securable column name]**.

APPLYING DATA LEVEL PERMISSIONS

When you have created a securable column setup, you are ready to use it to apply data level permissions to a table. To do so, follow the steps below.

1. Drag and drop a securable column setup on a field in the table in the data warehouse that you wish to secure. A secured view is created and can be found under **Security**, **Secured Views**, **[table name]_SV**.
2. If you wish to add further permissions to the view, you can drag and drop a securable column setup on the view. The **Add Field** window opens.
3. In the **Field Name** list, select the field that contains the values you wish to use in the filter with the securable column setup.

4. Click **OK**. The field is added to the secured view.
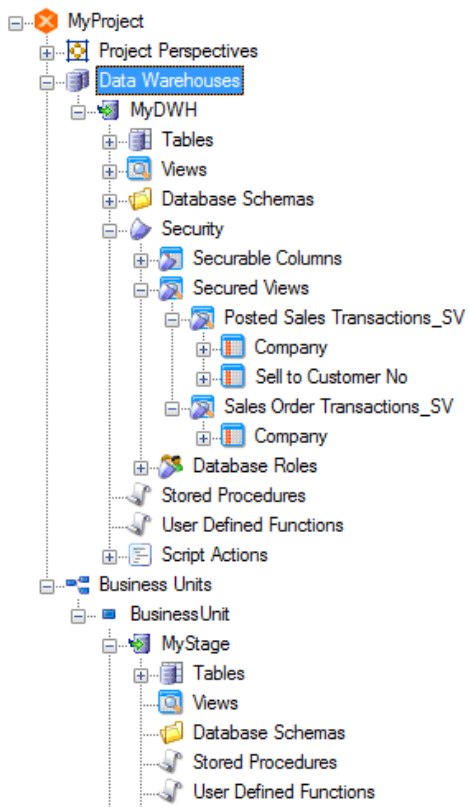
EXPORTING DATA

The nature of a data warehouse is to be a means to an end, for instance getting up-to-date sales numbers every morning or the data needed for financial reporting.

Jet Data Manager supports these scenarios through OLAP cubes or data visualization tools such as QlikView or Qlik Sense. Through the Data Export feature, however, you also have the ability to push parts of - or the entire data warehouse - to another destination, such as an Oracle database or text files.

The feature uses the same external provider concept as the custom data source, which means that we have a framework that makes it possible to add new providers without releasing a new version of Jet Data Manager.

DATA EXPORT

Data export allows you to push all or some of the content of a data warehouse to another destination, for instance another database.

For each type of destination, you will need a data export provider that can be downloaded and installed through the Custom Component Setup application. Please see the support site for more information: https://support.timextender.com/hc/en-us/articles/209604866

At the time of writing, the following data export destinations are supported:

- Oracle
- SQL Server
- Text files

A data export, or data export destination, and the tables and fields it contains, supports a subset of the features found on regular data warehouses.

Table features:

- Selection rules
- Guard table
- Preview table and the query tool
- Tracing
- Pre- and postscripts
- Description

Field features:

- Include in primary key
- Edit name and data type
- Tracing
- Description

ADDING A NEW DATA EXPORT

To add a new data export, follow the steps below.

1. Click the **Data Export** tab, right click **Data Exports** and click **Add Data Export**.
2. In **Name**, type a name for the data export destination.
3. In the **Provider** list, click the provider you wish to use.
4. In the **Setup Property** list, click **Setup Properties**. In the grid below, enter the setup properties required by the selected data export destination.
5. (Optional) Click **Test connection** to test that the connection properties you entered allow you to connect.

ADDING A TABLE TO A DATA EXPORT

To add a table to a data export, follow the steps below.

1. On the **Data Export** tab, expand **Data Warehouses** until you have located the table you want to add. Drag the table from **Tables** under **Data Warehouses** and drop it on **Tables** under the relevant data export under **Data Exports**.
   The **Add Data Export Table** window opens.
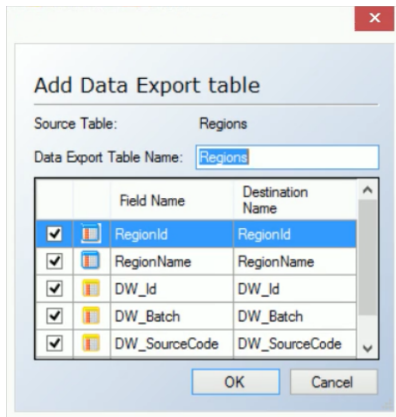
2. (Optional) In **Data Export Table Name** type a name for the table when used in the data export.
3. Clear the selection for fields you do not want to include in the table on the data export. Click a name in the **Destination Name** to edit the name as it appears on the **Data Export**.
4. Click **OK**.

ADDING ALL TABLES AND/OR VIEWS FROM A DATA WAREHOUSE TO A DATA EXPORT

To add all tables and/or views from a data warehouse to a data export, do as follows. On the **Data Export** tab, expand **Data Warehouses** until you have located the data warehouse that contains the tables or views you want to add. Drag **Tables** or **Views** under the data warehouse – or the entire data warehouse – and drop it on the relevant data export under **Data Exports**.

Jet Data Manager will add all views and all tables with all fields to the data export with the following exceptions:

- Tables or views that have already been added to the data export will be ignored.
- Tables or views that are not visible on the data warehouse in the currently selected project perspective will not be added.
- Tables or views that are not visible on the data export in the currently selected perspective, will be ignored if you add them again.

PREVIEWING A TABLE

Once the data export has been executed, the regular Preview Table command can be used to view the content of the table.

Click on **Destination** in the **Instance** list to see the data stored in the data export and **Source** to view the data in the table in the data warehouse.

ONLINE ANALYTICAL PROCESSING (OLAP)

Jet Data Manager includes support for designing Online Analytical Processing (OLAP) cubes that end users can browse through Excel or specialized business intelligence front-ends. OLAP utilizes the fact- and dimension tables you have created in your data warehouse. See [Dimensional Modeling](#) to learn more about how you structure a data warehouse.

Cubes allow you to present data in a multidimensional model. You can break down data in your data warehouse into smaller units, enabling you to drill-down, or roll-up through data, depending on the level of detail you want to view. You can, for example, create a sales cube, a production cube, a finance cube, and so on.

A cube consists of a number of dimensions and measures. The dimensions determine the structure of the cube, and the measures represent the numerical values. You can, furthermore, define hierarchies within a dimension by using dimension levels.

Dimensions define how a user looks at data in a cube. Dimensions are created independently of a particular cube and can be used within several cubes at the same time. The same dimension can also be utilized several times within the same cube, which is referred to as a role-playing dimension. A common example of this would be the Date dimension, which can represent both the Document Date and Posting Date in a cube, thus having a single dimension play two roles.

ADDING AN OLAP SERVER

To use OLAP, you will first need to add an OLAP Server.

1. On the **Cubes** tab, right-click **OLAP Servers**, and then select **Add OLAP Server**.

A window with the following selections will appear:

2. In the **Name** field, type a name for the OLAP server. The name cannot exceed 15 characters in length.
3. In the **Server Name** box, type the name of the OLAP database server.
4. In the **Database** box, type a name for the database.
5. In the **Collation** list, click the database collation to use for the OLAP database. **<Server Default>** will inherit the default collation currently set in Analysis Services. **<Application Default>** will use Latin1_General_CI_AS. This collation should correspond with the collation that is set for SQL Analysis Services.
6. Select a data warehouse from the **Data Warehouse** list, and then click **OK**. Each OLAP database can pull from a single data warehouse database.

Note: If you choose to delete the OLAP database later, it will remove the database from the project, but it will not delete the physical database on the OLAP Server itself. This must be done manually through SQL Management Studio.

CUBES

Cubes are the cornerstone when presenting data through OLAP. Cubes are typically built around central functions in the company, such as sales, finance, inventory etc.

ADDING OLAP CUBES

To add a cube, follow the steps below.

1. On the **Cubes** tab, expand the OLAP server, right-click **Cubes** and click **Add Cube**.

   The **Add Cube** window opens.
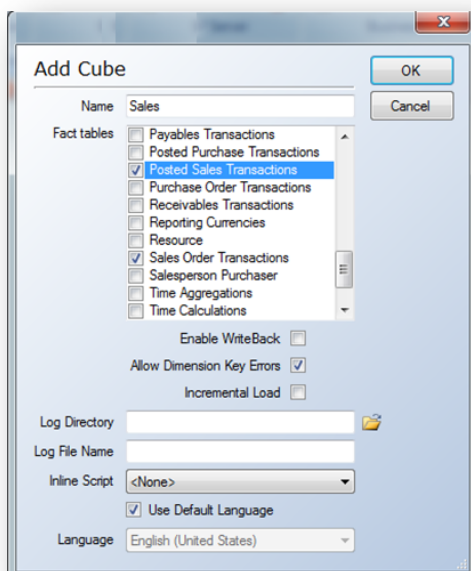
2. In the **Name** field, type a name for the cube.
3. In the **Fact Table** list, select on the table(s) from the data warehouse you want to use for the cube.
4. If you want end users to be able to change cube data while they browse it, select **Enable Writeback**. Any changes the end users make are saved in the write-back table.

   Note: You can only utilize write-back if the front-end application supports it.

5. If you want to continue processing the cube even if dimension key errors occur, select **Allow Dimension Key Errors**. When you allow dimension key errors, all errors are reported to a log and you need to specify where you want to store the log. Click the folder icon besides **Log Directory**. This opens a new window. Navigate to the folder you wish to store the log in and click **OK**. In the **Log File Name** box, type a name for the log file.
6. Click**OK** to add the cube.

ADDING A SINGLE DIMENSION TO A CUBE

1. Expand **Dimensions** and drag-and-drop the dimension to the cube.
2. Set the relationship to the fact table in the cube. See Adding Dimension Relationships.

ADDING MULTIPLE DIMENSIONS TO A CUBE

You can also add multiple dimensions to a cube in one operation.

1. On the **Cubes** tab, expand the relevant OLAP server, expand **Cubes**, and then expand the cube you wish to add dimensions to.
2. Right-click**Dimensions**, and then select **Add Dimension to Cube**.
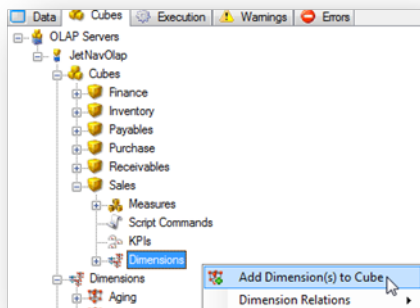
The **Add Dimensions to Cube**window opens with all the dimensions in your project listed.
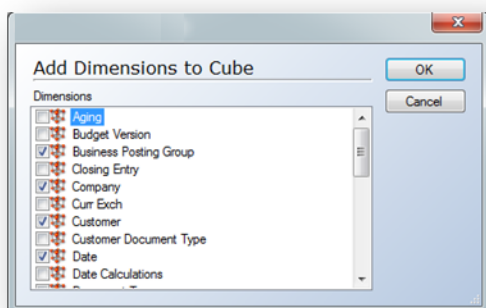
3. Select the dimension or dimensions you wish to add, and click **OK**.

You must then set the relationship to the fact table in the cube. See [Adding Dimension Relationships](#).

ADDING ROLE-PLAYING DIMENSIONS

Role-playing dimensions are dimensions that are used more than once in the same cube. For example, you can use a Customer dimension more than once in the same cube.

When you have added a role-playing dimension to a cube, you should follow the steps described above to add the dimension and then specify a new name for the dimension to distinguish it from the other dimensions of the same type. Then define the relationship to the proper field in the fact table.

For example, the Customer dimension may be used as the Bill-to Customer dimension, which relates to the Bill-to Customer No. field in the fact table, while the Sell-to Customer dimension may relate to the Sell-to Customer No. field in the same fact table.

ADDING DIMENSION RELATIONSHIPS

Dimension relationships specify how a dimension is related to a fact table. You must define how each level in a dimension is related to a fact table.

1. On the **Cubes** tab, expand **OLAP Servers**, and then expand the OLAP server that contains the relevant cube. Expand the cube, right-click **Dimensions**, click **Dimension Relations** and click **All Fact Tables**.
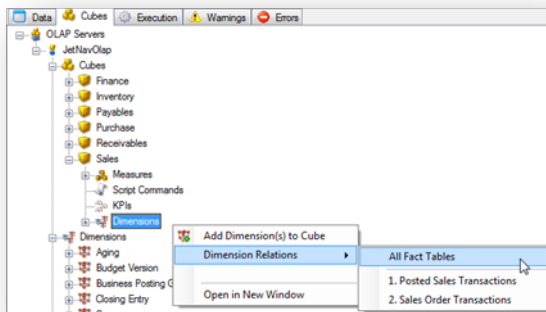
2. The **Dimension Relations** window opens.
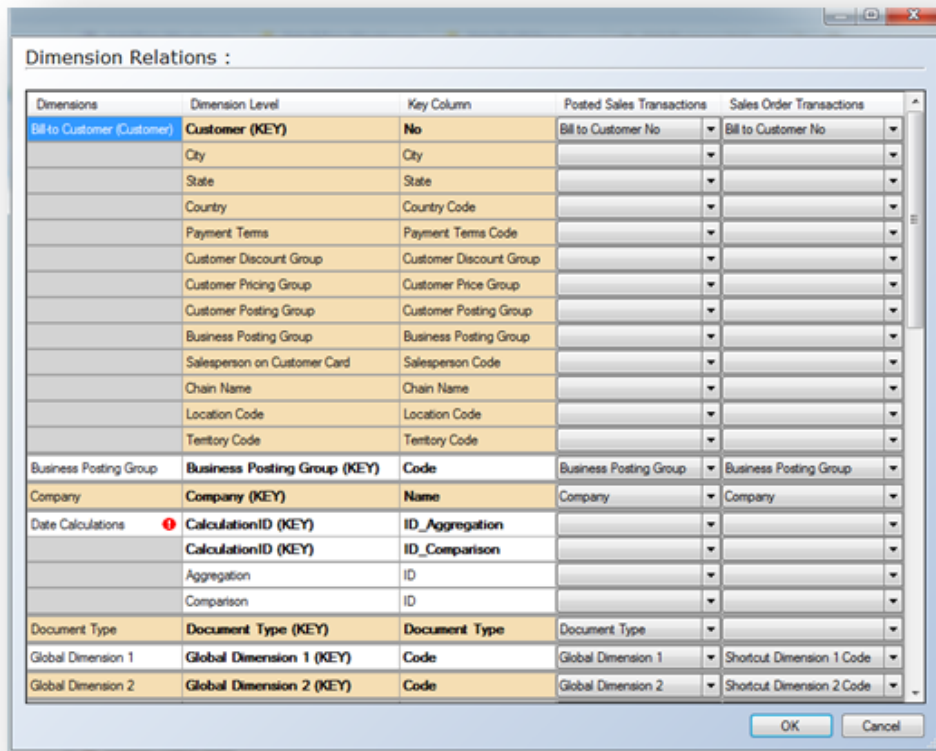
The table contains the following columns:

Column	Description
Dimension	Displays all dimensions in the cube
Dimension Level	Displays all dimension levels associated with each dimension
Key Column	Displays the key column for each dimension level
<Fact Table Name>	Displays the name of the fact table

3. In the Fact Table column(s), select the field where you will join the dimension key with the dimension. The dimension levels in **Bold** are the required values to be set for each dimension.
4. Click **OK** when you have created all of the necessary relationships.

VALIDATING CUBE AND DIMENSIONS

To validate a cube or a dimension, follow the steps below.

1. On the **Cubes** tab, expand the OLAP server that contains the cubes you want to validate, and then expand **Cubes**.
- OR -
On the **Cubes** tab, expand the OLAP server that contains the dimensions you want to validate and then expand **Dimensions**

2. Right-click the cube or dimensions you wishto validate, and then select **Validate Cube/Dimension**.
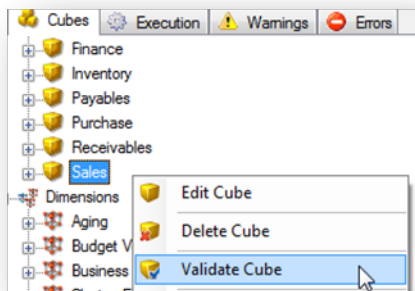
If the cube or dimension is valid, an OK message is displayed. Click **OK** to close the message dialog. If the cube or dimension is invalid, a message is displayed outlining the changes that you need to make.

THE CUBE BROWSER

The Cube Browser allows aJet Data Manager user to browse a cube from within Jet Data Manager without first leaving to go into another application such as Excel. The Cube Browser is not meant to replace a proper front-end tool, such as Excel, but is an easy way to browse the cube structure without having to navigate away from Jet Data Manager.

LAUNCHING THE CUBE BROWSER

Locate the cube that will be browsed on the **Cube** tab, right-click the cube, and select **Browse Cube**.
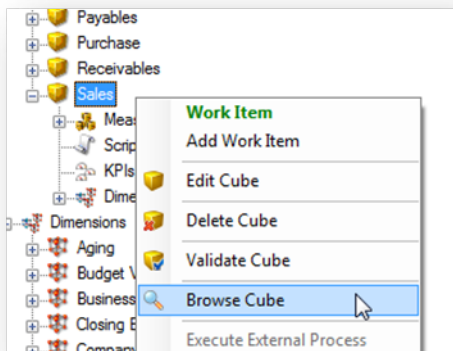
USING THE CUBE BROWSER

The user interface is similar to pivot tables in Excel. Measures and dimensions are dragged from the list on the right and dropped in either the boxes in the lower right-hand corner or directly onto the workspace pane on the left.
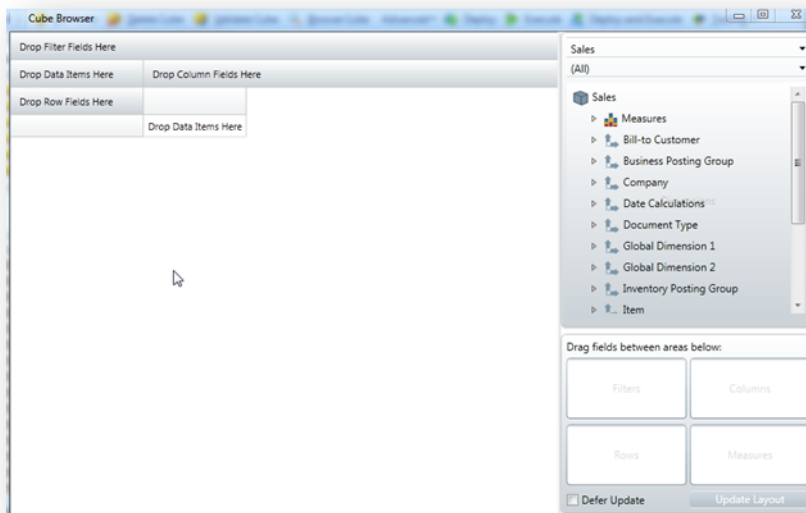
The following report was created from the Sales cube for NAV by dragging the **Date** dimension into the **Rows** box, the**Salesperson on Document** dimension into the **Columns** box, and the **Sales Amount** measure into the **Measures** box.
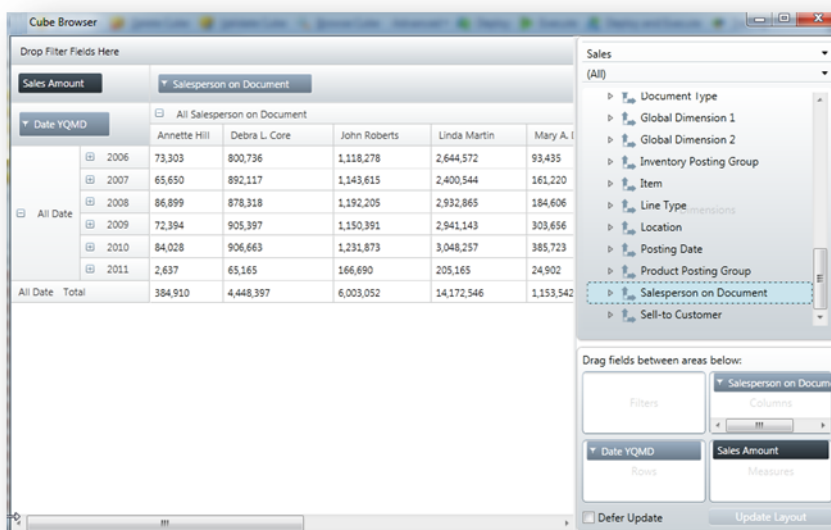
Filters for the rows and columns can be added by clicking the **Filter** icon to the left of the row and the column labels in the workspace pane on the left.
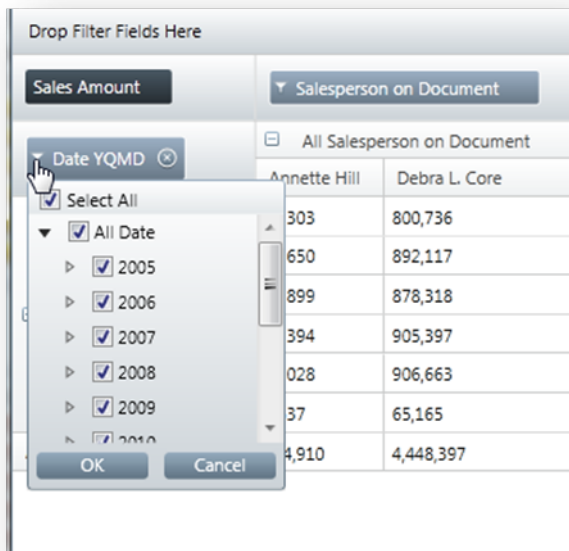
Dimensions can also be expanded and collapsed by clicking the plus and minus signs respectively.

CUBE WRITEBACK

Cube Writeback is a feature that allows users to update or add data to the cubes through the front-end.

When Writeback is enabled for a cube, it will be enabled for all measure groups in that cube.

The structure of the fact tables will be maintained and updated when the structure of the cube changes.

It is important to understand how writeback works before implementing it. While this document does not target its full scope, it warrants a thorough understanding. Otherwise, users will find themselves performing actions that are incomplete.

To enable this feature, right-click the cube you wish to utilize writeback on, and click **Edit Cube**.

In the dialog that opens, check **Enable WriteBack**, and click **OK** to close the dialog.
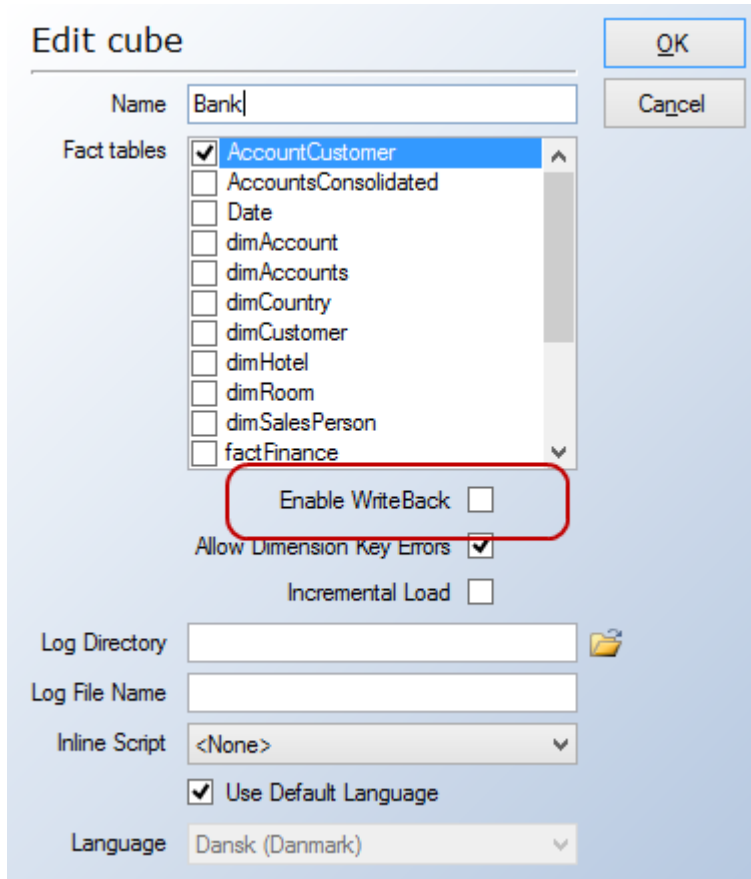
When writeback has been enabled, a writeback table will be created in the data warehouse on the next deployment. Whenever anything is written to a measure, it will be added to the writeback table. Data will never be written directly to the original fact table.

For the best results, design a separate cube specifically targeted for cube writeback.

OFFLINE CUBE PROCESSING

Jet Data Manager supports functionality that allows cubes to be processed without being taken offline. Normally, when a cube is processed and rebuilt in SQL Server Analysis Services, it is taken offline during the duration of the processing and is made unavailable to end-users.

SQL Server Enterprise Edition allows the cube to be kept online, but if you do not run this edition of SQL Server, the Offline Cube Processing feature of Jet Data Manager allows you achieve the same result.

Enabling Offline Cube Processing means that the cubes can be updated throughout the business day without disturbing end-users. While the cube is being processed, the users will have access to the original version of the cube, which is replaced with the new version of the cube once processing has been completed.

ENABLING HIGH AVAILABILITY CUBE PROCESSING

1. On the **Cubes** tab, right-click the name of the OLAP database, and select **Edit OLAP Server**.

The Edit OLAP Server window opens.
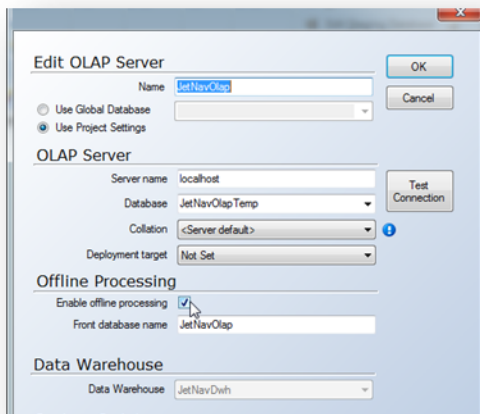
3. Rename the **Database** to represent a temporary OLAP database that will be used during processing. This is *not* the database that will be used by end-users.
4. Select **Enable Offline Processing**.
5. Type in the name of the database that will be used by end-users in **Front Database Name**.
6. Click **OK.**
7. Right-click the OLAP database and click **Deploy and Execute**.

7. You will notice two new steps under **Executed OLAP Server** that handles Offline Cube Processing. These are **Initialize Offline Processing** and **Finalize Offline Processing.**
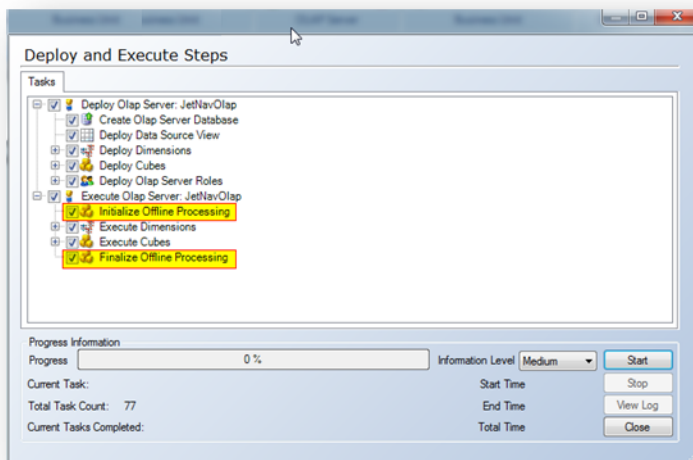
8. Click the **Start** button to begin processing the cubes. Users will now be able to access the cubes as they are being processed.

DIMENSIONS

Dimensions define how a user looks at data in a cube. Dimensions are created independently of a particular cube and can be used within several cubes at the same time. The same dimension can also be utilized several times within the same cube, which is referred to as a role-playing dimension. A common example of this would be the Date dimension, which can represent both the Document Date and Posting Date in a cube, thus having a single dimension play two roles.

REGULAR DIMENSIONS

Regular dimensions are based on a snowflake or a star schema, and are used to create balanced or ragged hierarchies.

CREATING REGULAR DIMENSIONS

1. On the Cubes tab, expand the relevant OLAP server, right-click**Dimensions**, and then select **Add Dimension**.

2. In the **Name** field, type a name for the dimension.
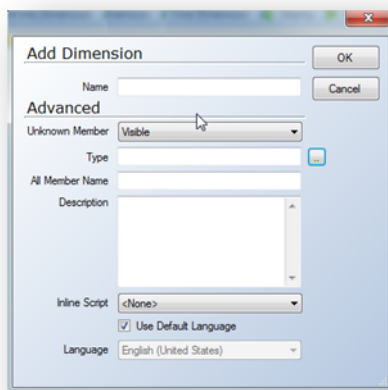
3. In the **Unknown Member** list, select **Visible** to apply an Unknown Member to dimension keys in the fact table with no matching dimension members. This will allow dimension values that exist in the fact table, but not the dimension table to be combined together and displayed to the user as an Unknown value. This is the default and recommended setting in Jet Data Manager. An example could be a Salesperson

Code that exists in the sales transactions fact table but does not exist in the Salesperson dimension table.

4. Next to the **Type** box, click the ellipsis **(...)** and select the type of dimension you wish to create. You can leave the **Type** blank for regular dimension types, which are the most common type.

Type	Description
Regular	Default for dimensions that are not set to a specified type
Time	Used for dimensions whose attributes represent time periods
Geography	Used for dimensions whose attributes represent geographical information
Organization	Used for dimensions whose attributes represent organizational information
BillOfMaterials	Used for dimensions whose attributes represent inventory and manufacturing information
Accounts	Used for dimensions whose attributes represent information used for financial reporting
Customers	Used for dimensions whose attributes represent information about customers
Products	Used for dimensions whose attributes represent information about products
Scenario	Used for dimensions whose attributes represent information about plans and strategies
Quantitative	Used for dimensions whose attributes represent quantitative information
Utility	Used for dimensions whose attributes represent utility information
Currency	Used for dimensions whose attributes represent currency information
Rates	Used for dimensions whose attributes represent currency rate information
Channel	Used for dimensions whose attributes represent channel information
Promotion	Used for dimensions whose attributes represent marketing promotion information

5. In the **All Member Name** box, type a name for the **All** Member. This is left blank by default, which means that Analysis Services creates the **All Member Name** automatically. The **All Member** is the dimension value which represents all members of

the dimension. An example would be a dimension value of "All Customers", which would represent every customer in the Customer dimension.

6. Optional: In the **Description** box, type a describtion for the dimension.
7. Click **OK** to add the dimension.

When you have added a dimension, you also have to add at least one dimension level. The **Add Dimension Level** dialog is displayed when you click **OK** to add a dimension. For more information, see [Adding Dimension Levels](#).

Once you have created a dimension, you can use the dimension in several cubes at the same time. Dimensions that are used in more than one cubes at a time are known as role playing dimensions.

PARENT-CHILD DIMENSIONS

Parent-child dimensions are used to create unbalanced hierarchies where the branches descend to different levels, and where the parent and the child exist in the same table. Typically, parent-child hierarchies are used for creating organizational hierarchies.

CREATING PARENT-CHILD DIMENSIONS

A parent-child dimension is a hierarchy that is defined by a parent column and a child column in the same table. A member of the hierarchy can appear more than once in the hierarchy.

To create a parent-child dimension, follow the steps below.

1. Expand **OLAP Servers**, expand the relevant OLAP server, right-click **Dimensions** and click **Add Parent-Child Dimension**.

The **Add Parent-Child Dimension** window opens.
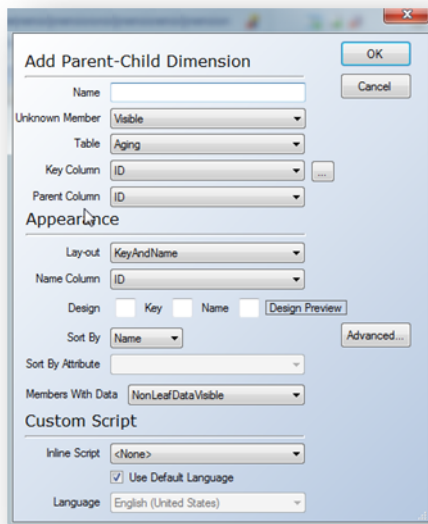
2. In the **Name** box, type a name for the dimension.

3. In the **Unknown Member** list, select **Visible** to apply an Unknown Member to dimension keys in the fact table with no matching dimension members.

4. In the **Table** list, select the main table of the dimension.

5. In the **Key Column** list, select the key column of the child table. This column identifies each member of the dimension.

6. In the **Parent Column** list, select the key column of the parent field. This column identifies the parent of each member.

7. In the **Lay-out** list, select how you want the dimension level displayed to the end user. The following options are available:

Setting	Description
Key	Displays only key column values
Name	Displays only name column values
KeyandName	Displays the key column first and then name column values
NameandKey	Displays the name column values first and then the key column values

8. In the **Name column** list, select the column that provides a meaningful value to the user. This field is only available if you have selected the **Name**,**KeyAndName**,or**NameAndKey** layout.

9. In the **Design** fields, specify which separator to use in the front-end application to separate Key and Name. This field is only enabled if you have selected KeyAndName or NameAndKey. The order in which the Key and Name text fields appear depends on your selection in the Layout list. To preview the design of the layout, move the pointer over **Design Preview**.

10. In the **Sort By**list, select whether you want the values sorted by Key or Name.

11. In the **Sort by Attribute** list, select the specific attribute key or name that you want the values sorted by. This list is only available when you are working with key levels, and Sort By is set to **AttributeKey** or **AttributeName**.

> **Note:** When you create the parent-child dimension based on a consolidation table, you will typically use a **Sort By Attribute**. You therefore need to create a Sort order dimension level where the key column is **Sort Order**. Then, you must enable Unary column and Roll up column on the dimension. You can then set the parent-child dimension to **Sort By Attribute**.

DEFINING ADVANCED PARENT-CHILD DIMENSION SETTINGS

1. To access advanced settings for Parent-Child dimensions, click **Advanced...** in the **Add/Edit Parent-Child Dimension** window.
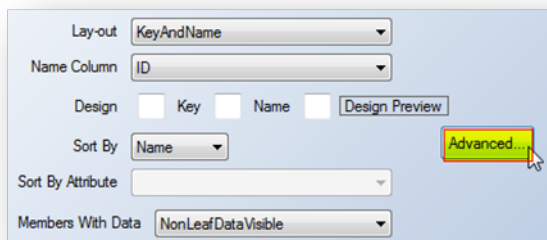
The **Advanced** window opens.
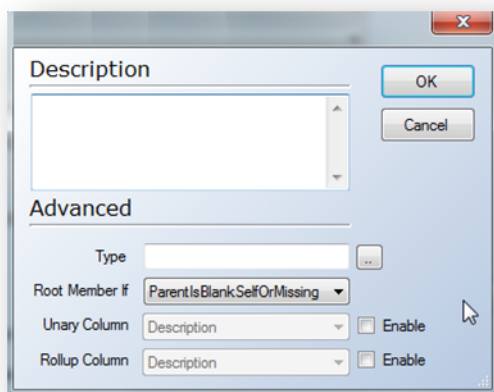
2. Next to the **Type** box, click the ellipsis **(...)** and select the type of dimension you wish to create. You can leave the **Type** blank for regular dimension types, which are the most common type. For a list of possible dimension types, see Creating Regular Dimensions.

3. In the **Root Member If** list, select one of the following options that controls when the dimension is the root member:

Option	Description
ParentIsBlank	Hides the root if the member is a null, a zero or an empty string
ParentIsBlankSelfOrMissing	Hides the root if the member is a null, a zero, an empty string, if the parent is missing, and if the member itself is a parent
ParentIsMissing	Hides the root if the parent is missing
ParentIsSelf	Hides the root if the member itself is a parent

4. In the **Unary Column** list, select the column that contains the unary operators that are used in this dimension level. If you You have to select **Enable** to select from this list.
5. In the **Roll-up Column**list, select the column that contains the roll-up values used in this dimension level. You have to select **Enable** to select from this list.

ADDING DIMENSION LEVELS

Dimension levels are used to create a dimension attribute within a cube, which enables a user to drill down or roll-up through data. A dimension must contain at least one dimension level.

1. On the **Cubes** tab, expand the OLAP server that contains the dimension you wish to add a level to. Then expand **Dimensions**, right-click the dimension and click **Add Dimension Level**.
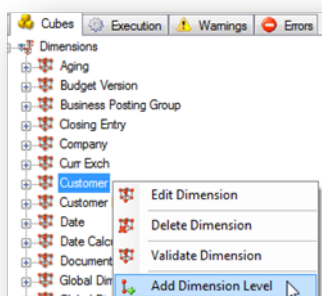
The **Dimension Level** window opens.
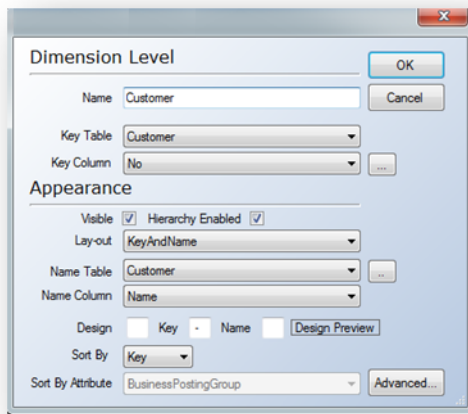
2. In the **Name** box, type a name for the dimension level.

3. In the **Key Table** list, select the table in the data warehouse to add the dimension from.

4. In the **Key Column** list, select the column which uniquely identifies the records in the table. If the key is a composite key, click the ellipsis button (…), to select the attributes on which the Key column is based.

5. Select **Visible** if you want the level to be displayed in the front-end application.

6. In the **Layout** drop-down, select the way that the dimension level should be displayed to users. The options are:

Setting	Description
Key	Displays only key column values
Name	Displays only name column values
KeyandName	Displays the key column values first and then name column values
NameandKey	Displays the name column values first and then the key column values

7. In the **Name Table** list, select the table that provides a meaningful name to the user. To set the Name Table value to the same value as the Key Table value, click the ellipsis button.

8. In the **Name Column** list, select the column that provides a meaningful value to the user.

9. In the **Design** fields, type the separators to use in the front-end application to separate key and name. This field is only enabled if you have selected KeyAndName or NameAndKey. The order in which the Key and Name text fields appear depends on your selection in the Layout list. To review the design of the layout, move the pointer over **Design Preview**.

10. In the **Sort By** list, select if you want the values sorted by **Key** or **Name**. If you are adding a key level, you can also sort by **AttributeKey** and **AttributeName**.

11. In the **Sort By Attribute** list, select the specific attribute key or name that you want the value sorted by. This list is only available when you are working with key levels, and **Sort By** is set to AttributeKey or AttributeName.

12. Click **OK**. The dimension level is added to the **Dimensions** tree below the dimension it belongs to. The OLAP database must be deployed and executed before this change takes effect in the front-end.

ADDING QUICK LEVELS

Quick Levels provide an easy way to add new dimension levels. It will automatically set defaults which can later be changed by editing the dimension level.

1. On the **Cubes** tab, expand the relevant OLAP server, expand **Dimensions**, right-click the parent-child dimension or key level on the dimension to which you wish to add a level and click **Add Quick Levels**.
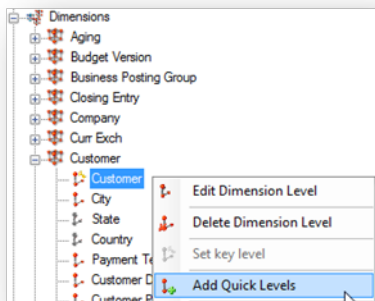
2. Select the columns you wish to use as levels from the source table in the data warehouse, and then click **OK**.
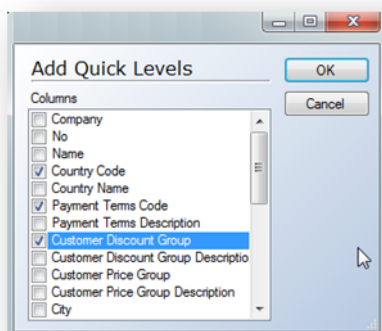
The levels you have added are now available when you create a hierarchy. The OLAP database must be deployed and executed before this change takes effect in the front-end.

ADDING DIMENSION JOINS

Dimension joins are joins between two tables that are not directly related to the dimension's fact table. You only use dimension joins in snowflake schemas where you want more than one table in a dimension. The join is a one-to-many join. The use of dimension joins is not common in the standard Jet Data Manager projects as they utilize a star schema which makes the dimension joins unnecessary.

1. On the **Cubes** tab, expand the OLAP server that contains the dimension you wish to modify, expand **Dimensions**, expand the dimension to which you want to add a join, right-click**Dimension Joins**, and then select **Add Olap join**.
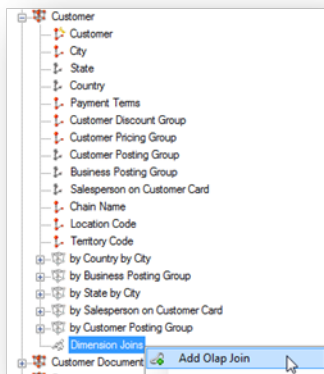
The **Add Join** window opens.

2. In the **From Table**list, select the table *from* which you want to create a join.
3. In the **To Table** list, select the table *to* which you want to create a join.
4. In the **From Column** list, select the column from which you want to create a join. The column's data type is displayed next to the list. You can only make joins between fields with compatible data types.
5. In the **To Column** list, select the column to which you want to create a join. If the column's data type is not compatible with the data type of the **From Column**, the data type is displayed in red.
6. If you want to reverse the direction of the join, click the **Reverse Join** button.
7. Click **OK**. The dimension join is displayed in the **Dimension Joins** folder in the **Dimensions** tree.

ADDING DIMENSION HIERARCHIES

Once you have added dimension levels to a dimension, you can create a dimension hierarchy. Dimension hierarchies make it easier for users to look at commonly used dimension groupings by only having to drag one icon into a report. An example of this could be Customers by Country or Items by Item Category.

1. Expand the relevant OLAP server, expand **Dimensions**, right-click the dimension to which you wish to add a hierarchy and then select **Add Hierarchy**.
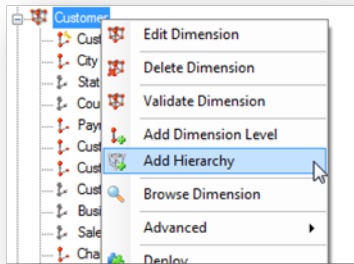
The **Hierarchy** window opens.
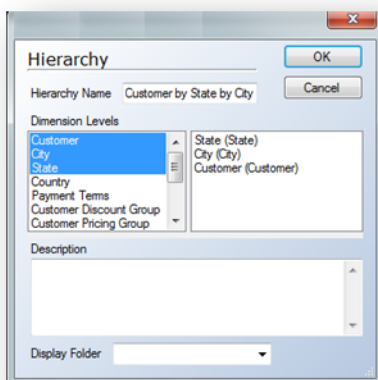
2. In the **Hierarchy Name** field, enter a name for the hierarchy. The name cannot be the same as the name of a dimension level.
3. In the **Dimension Levels** pane, click the levels you want to be part of the hierarchy. The hierarchy elements are then listed in the right pane. You can drag the dimension levels in the right pane up and down to specify the order they should exist in from top to bottom
4. In the **Description** field, type a description of the hierarchy. This field is optional.
5. In the **Display Folder** list, select the folder where the hierarchy is displayed by the front-end application. This is optional.
6. Click **Ok.** The OLAP database must be deployed and executed for this to be finalized for the end users.

Note: Since the hierarchy is associated with the dimension itself, once the dimensions and cubes are deployed and executed, the hierarchy will automatically show up in all cubes in which the dimension exists.

ADDING A TIME DIMENSION

Date dimensions are based on time or date tables, so you have to create a time or date table in the data warehouse before you can create a time dimension. For more information, see Adding Date Tables.

To add a date dimension on your OLAP server based on a date table, follow the steps below:

1. On the cubes tab, on you OLAP server, right click **Dimensions** and click **Add Time Dimension**. The **Add Date Dimension** window opens.
2. Type a **Name** for your Date Dimension, click the date table you wish to use as a basis for the date dimension in the Table list and click **OK**.
3. The date dimension will then be created based on the data table and appear in the OLAP tree under **Dimensions**. It includes date, week, month, quarter, half year and year dimension levels, in both fiscal year and non-fiscal year variations, as well as any custom periods and **Calendar** and **Fiscal Calendar** hierarchies.
4. (Optional) It might be useful to configure any custom period dimension levels to use the name of the custom period as a key. If, for instance, you have defined a number of national holidays across many years, these will then be grouped as op-posed to a having "unique" yearly holidays for each year. To do so, right click any custom period level and click **Edit Dimension Level** and click **[custom period name]Name** in the **Key Column list**.

When you expand the date dimension, you can see that the levels which correspond to fields on the date table in the data warehouse have already been added. However, you can add more levels simply by adding quick levels or by adding regular levels. You can add hierarchies as well.

MEASURES

Measures determine the numerical values of a cube and a cube must contain at least one measure. You can define the following three types of measures:

- **Standard measures** obtain their values directly from a column in a source fact table.
- **Derived measures** are derived before aggregation or summing of columns. This means that they are calculated when the cube is processed and are stored in the fact table. You can use standard arithmetic operators and MDX statements to create derived measures.
- **Calculated measures** are calculated after aggregation and summing. They are calculated at query time and are never stored. You can create calculated measures using standard arithmetic operators and MDX statements and can also combine them with other measures.

ADDING STANDARD MEASURES

1. On the **Cubes** tab, expand **OLAP Servers**, expand the OLAP server that contains the cube to which you wish to add a measure, expand the cube, right-click **Measures** and click **Add Standard Measure**.

The Add Measure window opens.
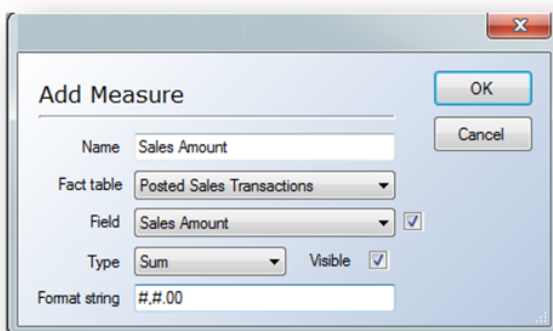
2. In the **Name** box, type a name for the measure.
3. In the **Fact Table**list, select the fact table that you want to use for the measure.
4. In the **Field** list, select the field that you want as the measure. Disable this field by clicking the box next to it.
5. In the **Type** list, select the preferred aggregation method. You have the following options:

Aggregation Method	Description
SUM	Returns the sum of all values
COUNT	Counts all rows and returns the total number of rows
MIN	Returns the lowest value
MAX	Returns the highest value
DistinctCount	Returns the number of unique values

6. Select **Visible** if you want the measure to be displayed in the front-end application.
7. In the Format string field, specify how you want the numeric results displayed. You have the following options:

Format String	Description
None	Applies no formatting
0	Displays a digit if the value has a digit where the zero (0) appears in the string, otherwise a zero is displayed
#	Displays a digit if the value has a digit where the number sign (#) appears in the string, otherwise nothing is displayed
.	Determines the number of digits displayed to the left and right of the decimal separator.
%	Is a percentage placeholder
,	Separates thousands from hundreds
Percent	Typing **Percent** will default the measure to showing as a percentage with two decimal places

Below are examples of what the output will look like for various combinations of the format strings:

Format String	Output
None	1234567.89
#,#	1,234,567

#,#.00	1,234,567.89
#,#%	1%
#,#.0%	1.2%
Percent	1.23%

ADDING DERIVED MEASURES

1. On the **Cubes** tab, expand **OLAP Servers**, expand the OLAP server that contains the cube to which you wish to add a measure, expand the cube, right-click **Measures** and

click **Add Derived Measure**.
The **Add Derived Measure** window will appear.

2. In the **Fact Table** list, select the fact table that you want to use for the measure.
3. In the **Name** box, type a name for the derived measure.
4. Select **Visible** if you want the measure to be displayed in the front-end application.
5. In the **Format String** box, type how you want the numeric results displayed. The format is the same as for standard measures. See Adding Standard Measures.
6. 7. In the **Expression** box, enter an MDX statement
   or
   In the **Measures** list, select the measures from the fact table that you want to use for

the derived measure and click**Add**. Your selections will be added to the expression, where they can be combined with mathematical operators to achieve the outcome you desire.

7. Click **OK** to add the derived measure.

ADDING CALCULATED MEASURES

1. On the **Cubes** tab, expand **OLAP Servers**, expand the OLAP server that contains the cube to which you wish to add a measure, expand the cube, right-click **Measures** and click **Add Calculated Measure**.

The Add Calculated Measure will open.
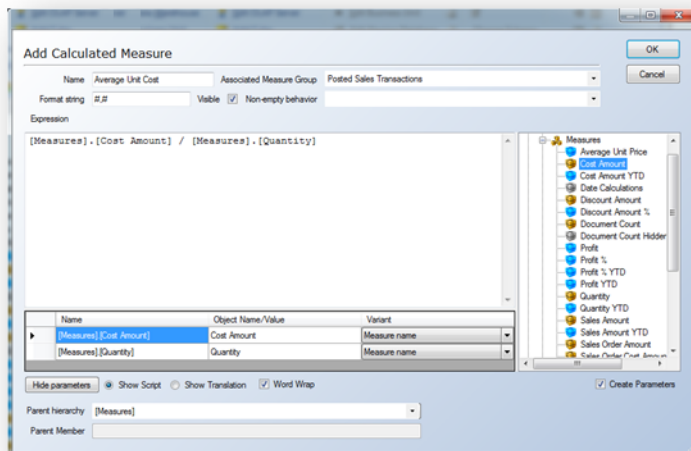
2. In the **Name** box, type a name for the calculated measure.
3. Select **Visible** if you want the value to be displayed in the front-end application.
4. In the **Format string** field, specify how you want the numeric results displayed. The format is the same as for standard measures. See Adding Standard Measures.
5. Optional: In the **Non-empty Behavior** list, select the measure or measures used to resolve NON EMPTY queries in MDX.
6. In the **Expression** field, write an MDX statement or, in the **Measures** list, drag the measures to be used for the calculated measure into the workspace in the middle.

HANDLING EARLY ARRIVING FACTS

In a live working environment, it is possible that transactional data may contain values that have not yet been added to the source database in the corresponding dimension table. An example of this could be a Sales Invoice that has a Salesperson Code where the Salesperson Code does not yet exist in the Salesperson table. When the data warehouse is updated and the cubes are processed, the values for this salesperson will fall under the "Unknown" member for the Salesperson dimension. This happens because the cube does not see the Salesperson Code on the transaction as being a known value when compared to the list of salespeople in the Salesperson dimension.

In Jet Data Manager,it is possible to handle these "early arriving facts" in such a manner that they will show at least partial information until the data source is properly updated with all of the normal dimension information. This prevents information from being placed into the "Unknown" member when the data is consumed by end-users. Once the dimension value is properly added to the ERP system or data source by a user, all fields for the previously missing record will then be populated according to the values in the data source.

ENABLING EARLY ARRIVING FACTS

1. On the **Data** tab, identify the dimension table to which relevant values from the transaction table should be added, right-click the table name, and go to **Advanced -> Add Related Records.**

   The **Add Related Records** window opens.

2. In **Name**, type a descriptive name ofr the **Add Related Records** rule that is currently being created.

3. In the **Create Records from Table**list, select the transaction table that will identify the table from which to bring in potential new values. A window may appear stating that all mappings and conditions will be cleared. Click **Yes**.

4. In the **Record Condition** list, select the option to determine when data will be inserted into the dimension table if new values are found in the transaction table. The most common option is **Not Exist**, which will add in values that do not currently exist in the dimension table.

5. Select the **Data Destination Table** to insert the values into. The default option is the Raw table.

6. In the **Field Mapping** table, specify the fields to be mapped from the transaction table and inserted into the dimension table. In the example below, the DW_Account field (Company) and Salesperson Code fields will be extracted from the transaction table and inserted into the dimension table.

7. It is possible to add in fixed values for fields in the dimension table that the transaction may not have data for. In the example below, the fixed value "Missing Salesperson" will be added in the **Name** field for all Salesperson Codes added from the transaction table. This is achieved by selecting the **Fixed Value** option in the **Mapping** column for the **Name** field and typing the desired fixed value in the **Fixed Value** column.

Field Mapping

	Mapping		Fixed Value
DATAAREAID	DATAAREAID	⌄	
EMPLID	SALESADMINISTRA...	⌄	
Name	Fixed Value	⌄	Missing Salesperson

8. If desired, a default value can be inserted instead of bringing in the values that exist in the transaction table. This could be used to assign fixed values to all data brought in for early arriving facts. This is achieved by clicking the checkbox in the **Allow Default Value** column and typing the corresponding fixed value in the **Default Value** column. This is not common.

Field Mapping

	Mapping		Fixed Value	Allow Default Value	Default Value
DATAAREAID	DATAAREAID	⌄		☐	
▶ EMPLID	SALESADMINISTRA...	⌄		☑	Missing
Name	Fixed Value	⌄	Missing Salesperson	☐	

9. The last step is to define the relationship between the two tables. Click the **Add** button in the **Conditions** section.
10. Select the first field to join in the dimension table (**Code**), and click **OK.**

Conditions

1. [⌄] Emp
  - AX_dbo_EMPLTABl
    - DATAAREAID
    - EMPLID
    - Name
    - DW_Id
    - DW_Batch
    - DW_SourceCod
    - DW_TimeStamp

  OK   Cancel
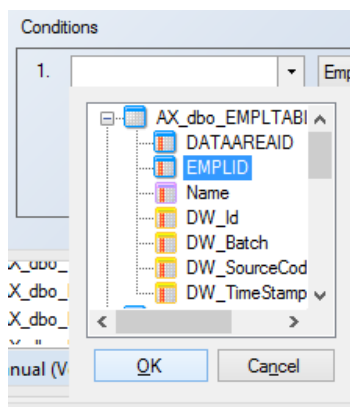
11. Select the operator to be used for the join. The most common operator is **Equal**.
12. Select the matching field in the transaction table (**Salesperson Code**), and click **OK.**

Conditions   Add

| 1. | EMPLID | ▾ | Equal | ⌄ | Fixed Value ☐ | SALESADMINISTRATOR ▾ | Remove |

13. Repeat steps 9 through 12 for any additional joins that need to be made (such as Company). The final result will look similar to the screenshot below.

Click **OK** when finished to save the settings, and close the **Add Related Records** window.

A folder for **Table Transformations** will be added to the bottom of the dimension table. The selection criteria that were previously set can be edited by right-clicking the transformation and selecting **Edit Related Record**.

14. Deploy and execute the dimension table. Any records that exist in the transaction table, but not in the dimension table, will be added during the data cleansing process. A screen-shot of the result based on the example in this document is shown below.

Table: AX_dbo_EMPLTABLE

DATAAREAID	EMPLID	Name	D
ceu	7225	Null	83
ceu	7230	Null	84
ceu	7231	Null	85
ceu	7232	Null	86
ceu	7240	Null	87
ceu	9001	Null	88
ceu	9002	Null	89
ceu		Missing Salesperson	90

The salesperson code "BP" existed on a sales document, but no corresponding Salesperson Code existed in the Salesperson table. Once the salesperson is properly added to the ERP system and the table is refreshed, all proper information will be pulled in from the ERP system, and the name will no longer say "Missing Salesperson."

SLOWLY CHANGING DIMENSIONS

Slowly Changing Dimensions (SCD) enable an organization to track how dimension attributes change over time. For example, it is possible that an item may be associated with a particular product group code but that it is later reclassified into a different product group. The organization wants to be able to analyze the historical sales data that occurred when the item was assigned to the original product group as well as more recent sales data that has occurred after the item was reclassified to the new product group.

DIFFERENT TYPES OF SLOWLY CHANGING DIMENSIONS

TYPE I

Type I dimensions will automatically overwrite old data with updated data from the data source. An example of this would be a change in a customer name. In the data source, the name for a particular customer is changed from ABC Consulting to Acme Consulting. The next time that the data warehouse is updated the customer name will be changed from ABC Consulting to Acme Consulting and no historical record of the change is kept. All historical, current, and future transactions will be displayed under the new customer name of Acme Consulting. This is the default methodology of updating data in the data warehouse and no setup is required.

TYPE II

Type II dimensions will enable the tracking of dimension attributes historically by inserting additional records into the table as the values in specified fields are changed. Jet Data Manager will administer the tracking of the dimension values as well as the updating of the table. Each record for a particular value, such as an item number, can be viewed as a different version of this item. The transaction table can then be linked to this table to display which version of the item was associated with the transaction based on the transaction date.

The follow example illustrates what the Customer dimension table would resemble if the example above was tracked using Type II functionality.

| Customer No	Name	City	State	Version	From Date	To Date
123	ABC Consulting	Portland	OR	1	1/1/1900	9/18/2012
123	Acme Consulting	Portland	OR	2	9/18/2012	12/31/9999

IMPLEMENTING TYPE II SLOWLY CHANGING DIMENSIONS

The steps below will explain how to utilize slowly changing dimensions - also known as History in Jet Data Manager - in a project. In the example, there will be an item named "Bicycle" that has historically had an Inventory Posting Group of "Finished". Recently, however, this item has been reclassified and is now associated with the Inventory Posting

Group "Resale". The organization wishes to track sales for this item under both the historical Inventory Posting Group as well as the new one.

ENABLING SLOWLY CHANGING DIMENSIONS ON THE DIMENSION TABLE

1. Identify the table where historical changes need to be tracked. Right-click the table, click **Table Settings**.
2. Click the **History** tab and select **Enable history**. If an error icon appears next to the setting, it means that another setting needs to be changed to enable source based incremental loading. Move you mouse over the error icon to see the error message.
3. Click **OK**.
4. The table icon will now be overlaid with an "H" to make it easy for you to identify it as a table with history/slowly changing dimensions enabled. Expanding the table will show a node named **History Settings**.
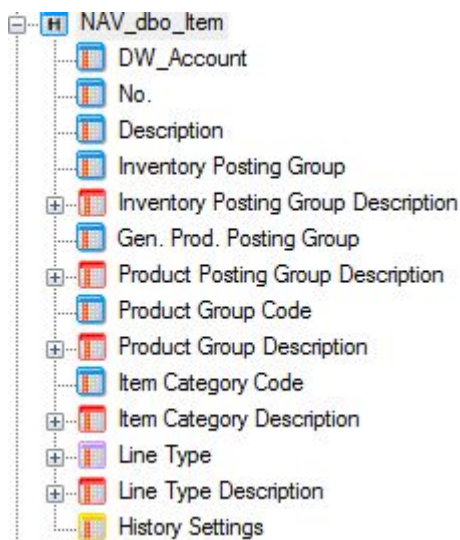
5. Click **History Settings**. The Slowly Changing Dimensions configuration options appear in the right-hand side of the screen.
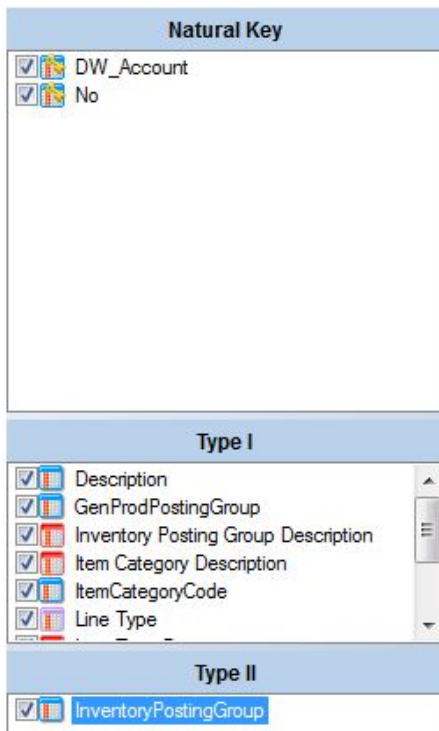
6. This screen is broken down into three sections:
   - **Natural Key:** Selected all fields that represent the natural key of the table. Any primary keys defined on the table is automatically selected. In the example this is the "DW_Account" and "No." fields
   - **Type I:** Here, all fields that should not have history tracking enabled are listed. If values in these fields are changed in the data source, the old values will be overwritten by the new values when the table is next executed.
   - **Type II:** Select all fields that should have history tracking enabled.
7. Deploy and execute the table to have the initial history data stored.

Note: A primary key comprising at least one field must be set on the table in order for the table to deploy and execute properly. The primary key fields for a table can be set by right-clicking the desired fields within the table and selecting **Include in Primary Key**. This can be done for multiple fields in the same table.

EXAMPLE

The screen-shot below illustrates what the Item table currently looks like for the item that is being used in this example. There is one record for the item and currently the Inventory Posting Group is set to "Finished".

DW_Account	No.	Description	Inventory Posting	Inventory Posting
CRONUS EXT U...	1000	Bicycle	FINISHED	FINISHED

The item is now reclassified into the Inventory Posting Group Code of "Resale.". An invoice is then posted that reflect a sale of 100 of the bicycles with the new Inventory Posting Group. To illustrate how the Item table now looks in the staging database, the table is executed to reflect the changes and the results are displayed below:

DW_Account	No.	Description	Inventory Posting	Inventory Posting	Gen. Prod.	Product Posting	Product Group	Product Group	Item Category	Item Category	Line Type	Line Type	DW_Id
CRONUS EXT U...	1000	Bicycle	FINISHED	FINISHED	RETAIL	Retail		Null		Null	2	Item	1
CRONUS EXT U...	1000	Bicycle	RESALE	RESALE	RETAIL	Retail		Null		Null	2	Item	145

The DW_Account and No. fields are the same, but the Inventory Posting Groups now reflect the new value. The DW_ID, which represents a unique record number in the table, is also different as illustrated on the right in the screenshot above.

There are a few more fields that pertain to the historical values that are useful as well. When the table is executed and notices a change in one of the Type II fields, it will automatically add in the dates for which the old value ended and the new value begins. These are the "SCD From Datetime", "SCD To Datetime", and "SCD "IsCurrent" fields.

SCD From DateTime	SCD To DateTime	SCD Is Current
01-Jan-00	25-Sep-12	0
25-Sep-12	31-Dec-99	1
01-Jan-00	31-Dec-99	1

The To and From field represent the date ranges that this version of the dimension was used in and which record is the current record.

BRINGING THE SURROGATE KEY TO THE TRANSACTION TABLE

In a standard transaction table, the transaction itself will only be linked to the Item No. This is problematic, as the item number alone does not identify which version of the item record the transaction applies to. In order to see this detail, the surrogate key from the Item table will be brought into the transaction table. This surrogate key is based on the To and From dates in the Item table, as compared with the Posting Date of the transaction. In the screenshot above, all transactions between January 1, 1900 and September 25, 2012 will be associated with the first version of the Bicycle item where the Inventory Posting Group is "Finished". Any transaction after September 25, 2012 will be associated with the latest version of the Bicycle item where the Inventory Posting Group is "Resale".

A surrogate key is a substitution for the primary key in a table. The surrogate key most often represents the unique row number in the table. It can be used in one table to refer back to a specific record in another table without having to utilize the natural primary key. In Jet Data Manager, all tables in the staging database and data warehouse have a called named "DW_ID" which represents the surrogate key in each respective table.

In order to see the DW_ID field in Jet Data Manager, follow the steps below.

1. Right-click the table, click **Advanced** and click **Show System Control Fields**.
2. Move the **DW_ID** field from the Item table, and add it to the transaction table.
3. Rename the field to "Item Surrogate Key" to make it easily understandable to other users.
4. Add Standard joins between the two tables for DW_Account and the Item No.
5. Add Additional joins for "SCD From Date Time" Less Than or Equal to "Posting Date" and "SCD To Date Time" Greater Than or Equal to "Posting Date". This will capture the correction version of the item based on the Posting Date of the transaction.
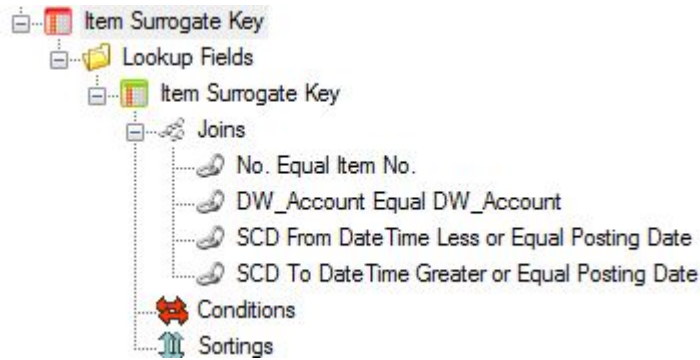
6. Deploy and execute the transaction table to have the new field added and populated.

The process above should be repeated for any additional transaction tables where history needs to be tracked.

MOVING THE SURROGATE KEY FROM THE TRANSACTION TABLE IN THE STAGING DATABASE TO THE DATA WAREHOUSE

Now that the surrogate key has been added to the transaction table(s), this field needs to be added to the relevant transaction tables in the data warehouse.

To accomplish this, follow the steps below.

1. Drag the surrogate key field (in this case "Item Surrogate Key") from the table(s) in the staging database and drop them onto the relevant tables on the data warehouse.
2. Deploy and execute the transaction table in the data warehouse for the changes to take effect.

```
└─⊟ Posted Sales Transactions
   ├─ 🔻 Module Equal 1
   ├─⊞ 🔑 Company
   ├─⊞ 🔑 Value Entry No
   ├─⊞ 🔑 Document Type
   ├─⊞ 🔑 Document No
   ├─⊞ 🔑 Document Line No
   ├─⊞ ▥ Module
   ├─⊞ ▥ Posting Date
   ├─⊞ ▥ Type
   ├─⊞ ▥ No
   ├─⊞ ▥ Sell to Customer No
   ├─⊞ ▥ Bill to Customer No
   ├─⊞ ▥ Salesperson Code
   ├─⊞ ▥ Global Dimension 1 Code
   ├─⊞ ▥ Global Dimension 2 Code
   ├─⊞ ▥ Business Posting Group
   ├─⊞ ▥ Product Posting Group
   ├─⊞ ▥ Inventory Posting Group
   ├─⊞ ▥ Location Code
   ├─⊞ ▥ Sales Quantity
   ├─⊞ ▥ Sales Amount
   ├─⊞ ▥ Discount Amount
   ├─⊞ ▥ Cost Amount
   └─⊟ ▥ Item Surrogate Key
      ├─ 📋 Copy From [Business Unit]: [NAV_dbo_Value Entry].[Item Surrogate Key]
      ├─ 📋 Copy From [Business Unit]: [NAV_dbo_Sales Cr.Memo Line].[Item Surrogate Key]
      └─ 📋 Copy From [Business Unit]: [NAV_dbo_Sales Invoice Line].[Item Surrogate Key]
```

ADDING THE SURROGATE KEY TO THE DIMENSION TABLE IN THE DATA WAREHOUSE

The **DW_ID** field now needs to be added to the related dimension table in the data warehouse. This ensures that the proper mapping will be made between the dimension table and the transaction table in the cubes.

To accomplish this, follow the steps below.

1. Drag the DW_ID field from the dimension table in the staging data (in this example, the Item table) to the related dimension table in the data warehouse (in this example, the Item table).
2. Deploy and execute the dimension table for the changes to take effect.
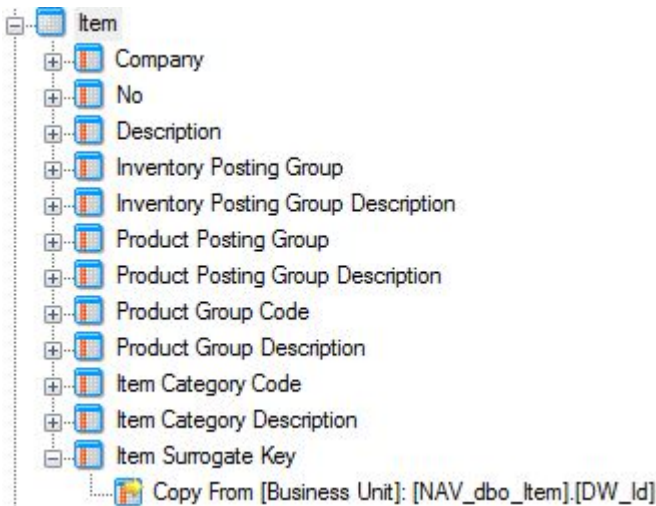
UPDATING THE DIMENSION KEY

The dimension key should now be updated to include this surrogate key. This will ensure that Analysis Services sees the uniqueness of the dimension, not as the natural key (in this example Item No.), but as the combination of Item No. and the surrogate key.

To accomplish this, follow the steps below.

1. On the Cubes tab, expand **Dimensions**, expand the dimension, and edit the key level (in this example "Item").

2. To the right of the Key Column click the ellipsis **(...)**, and add the surrogate key to the dimension key (in this example, it is the "Item")

UPDATING THE DIMENSION RELATIONSHIPS IN THE CUBE

The relationships between the dimension and the transaction table should be updated in the cube to reflect the change made to the dimension key in the previous step.

1. Right-click the relevant dimension in the cube(s), click **Dimension Relations** and click **All Fact Tables**.
2. Set the dimension relationship to use the surrogate key that was added to the transaction table in a previous step.

3. Deploy and execute the OLAP database for the changes to take effect.

The final result is that users can see data based on the historical attributes that may no longer exist in their ERP system because the information has been overwritten. In the screen-shot below, the "Bicycle" item shows up twice with the sales amounts associated with the various Inventory Posting Groups that have been used for the item over time.

Row Labels	Inventory Posting Group	Sales Amount
1000 - Bicycle	FINISHED	150,000
1000 - Bicycle	RESALE	300,000
1896-S - ATHENS Desk	RESALE	1,327,390
1900-S - PARIS Guest Chair, black	RESALE	1,780,730
1906-S - ATHENS Mobile Pedestal	RESALE	136,797

Row Labels	Inventory Posting Group	Sales Amount
1000 - Bicycle	FINISHED	150,000
1000 - Bicycle		300,000
1896-S - ATHENS Desk		
1900-S - PARIS Guest Chair, black		
1906-S - ATHENS Mobile Pedestal		

DEPLOYING AND EXECUTING

Before a Jet Data Manager project is deployed and executed, it is simply a meta data model of your data warehouse. Deployment and execution generates and runs the code for extracting, transferring and loading your data as well as creating any OLAP cubes in the project. During development it can also be a good idea to deploy and execute the project to see if everything works as expected.

DEPLOYING A PROJECT

Deploying a project, or a part of a project, is the process of generating the structure of the staging database and the data warehouse, processing cubes and generating SQL code.

No data is loaded into the staging database or the data warehouse, and no cubes are processed at this time. When you successfully deploy a project, the project is automatically saved in the project repository.

Deployment in Jet Data Manager is optimized in two ways: It is managed, i.e. objects are deployed after any objects they depend on, and differential, i.e. only the steps that have changed since the last deployment are deployed again.

EXECUTING A PROJECT

Executing a project is the process of loading data into the staging database, the data warehouse, and then processing any OLAP cubes.

Executing a project involves the following steps:

1. **Transferring data**: The process of transferring data from the data source to the raw table of the staging database.
2. **Processing data**: The process of cleansing data; that is, validating the data against the business rules, and moving the validated data to the valid table. Status information is also generated at this point.
3. **Verifying data against checkpoints:** The process of checking the data that is being processed against the checkpoints you have specified. You can specify rules that will end the execution process if not met. This way, you avoid overwriting the data in your data warehouse with non-valuable data.
4. **Moving data:** The process of moving data from a business unit to a data warehouse, or from a data warehouse to a cube.
5. **Processing cubes:** The process of creating dimension hierarchies and retrieving values from the fact tables to populate the cubes with measures, including derived and calculated measures.

Jet Data Manager supports managed and threaded execution. This means that Jet Data Manager can execute a project in multiple threads while managing dependencies between objects and optimizing the execution to take the shortest amount of time.

MANUAL DEPLOYMENT AND EXECUTION

While you are developing and maintaining your project, you may want to deploy or execute the entire project or individual objects to confirm that everything works as expected.

Manual execution is configured in the default execution package, that can be configured just like any other execution package. See [Execution Packages](#) for more information.

DEPLOYING AND/OR EXECUTING INDIVIDUAL OBJECTS

Whether you want to deploy, execute or deploy and execute an object, the steps are similar.

1. Right-click the project or the project element you want to deploy and/or execute, and click **Deploy**, **Execute** or **Deploy and Execute**
- OR -
Click the object at press **CTRL + ALT+ D** to deploy, **CTRL + ALT + E** to execute and **CTRL + ALT + K** to deploy and execute. The **Deploy and/or Execute** window appears.
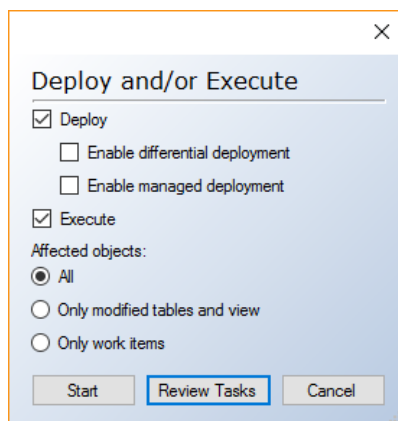
The settings in the window depends on how you initialized the window – choosing deploy, execute or deploy and execute – and your default settings for deployment.
2. Select **Deploy** to deploy the selected objects. You have two deployment options. Their initial setting is based on your project settings.
   1. Select **Enable differential deployment** to take advantage of Jet Data Manager's differential deployment feature that calculates what steps have changed and need to be deployed and selects only those steps for deployment. When differential deployment is disabled, all steps are deployed as in previous versions ofJet Data Manager.
   2. Select **Enable managed deployment** to have Jet Data Manager calculate dependencies and deploy the objects in the optimal order. There is no difference in performance or otherwise if you are only deploying one object. When managed deployment is disabled, you will have to make sure that objects are deployed in the correct order, just like in previous versions of Jet Data Manager.

3. Select **Execute** to execute the selected objects. The execution settings are governed by the default execution package.

4. Under **Affected objects**, select which of the objects you selected for deployment and/or execution you want to deploy and/or execute. Your options are **All**, **Only modified tables and views** and **Only work items**.

3. Click **Start** to begin the deployment and/or execution process as soon as Jet Data Manager is ready or **Preview Tasks** to review the tasks and settings before you start the process.

4. If you use the differential deployment method, it will take Jet Data Manager a few moments to calculate what steps need to be deployed. If Jet Data Manager does not find any steps that needs to be deployed, you will be asked if you wish to save the project as the deployed version. This has to do with the way the scheduler works. It will execute the last deployed version of a project, i.e. if you wish the current version to be the one that is executed by the scheduler, it needs to be marked at such.

5. The **Deploy**, **Execute** or **Deploy and Execute** window opens. If you clicked on **Start** earlier, Jet Data Manager will begin the deployment and/or execution process immediately. Otherwise, review the tasks below b
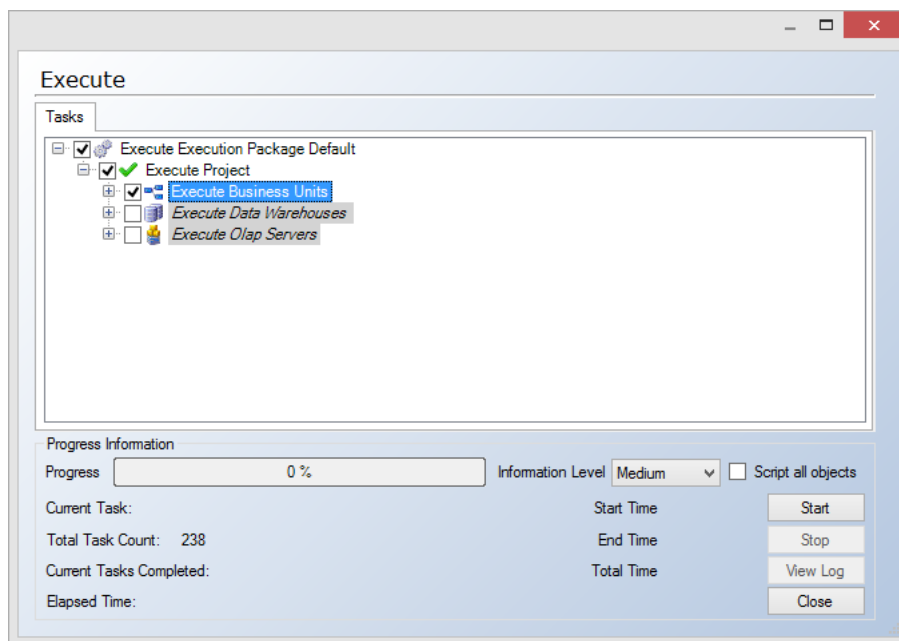
6. Clear the selection for any objects you do not want to deploy, execute or deploy and execute.

7. (Optional) In the **Information level** list, click your preferred level of information during the completion of the tasks. The following options are available:

   1. **None:** Displays no progress information.
   2. **Low:** Displays current task, the total task count, start time, end time, and total time.
   3. **Medium (default):** Displays progress information, name of the current task that is being deployed, number of completed tasks, the total number of tasks that have to be completed, start time, end time, and total time.

4. **High:** Displays all deployment steps in the task window, progress information, name of the current task that is being deployed, number of completed tasks, the total number of tasks that have to be completed, start time, end time, and total time.

8. (Optional) Select **Script all objects** to make all parts of the SQL script available in the log for debug purposes. This option disables the use of Shared Management Objects to create tables instead of executing "CREATE TABLE" statements against SQL Server since statements executed through Shared Management Objects are not in the log.

9. Click **Start** to deploy, execute or deploy and execute the objects you have chosen. Click **Stop**if you wish to halt the processing prematurely.

If there are any errors during deployment or execution, the **Error** window is displayed. Click **Yes** to view the log. The object on which the deployment or execution failed, has an **Error Information** node. Double-click **Error Information** to view an error description.

You will have to wait for Jet Data Manager to finish the tasks before you can continue work on the projects. However, with the Execution Queue, you can execute objects in the background which enables you to continue working on your project while objects are executed. See [Executing Objects with the Execution Queue](#) below.

Customized Code Warnings

Jet Data Manager will display a warning message if you try to deploy a table with customized code where the table has changed in Jet Data Manager since you last customized the code. Often, you need to update the customized code when you have made other changes to the table and the purpose of this feature is to help you catch some errors earlier.

Jet Data Manager displays the warning message when the deployment window opens. Click **OK** to close the message. On the **Customize Code Warning** tab, you can see what tables might have outdated customized code.

If you click **Start** after closing the message, Jet Data Manager will assume that you have things under control and not display the message again for the same issue.

EXPORTING DEPLOYMENT STATUS

You can export a deployment status report that contain a list of what needs to be deployed. You can generate the report for projects, data warehouses, business units and OLAP servers. You can also generate the report for a remote environment. The report is formatted as a CSV file.

To generate a deployment status report for a project, data warehouse, business unit or OLAP server, right click the object, click **Advanced** and click **Export Deployment Status**.

To generate a deployment status report for a remote environment, open **Multiple Envir-onment Transfer**, right click the remote environment and click **Export Deployment Status**.

Jet Data Managerwill ask you to choose a location for the file. When the report has been success-fully generated, Jet Data Manager will ask you if you want to open it.

REVIEWING THE EXECUTION LOG

On each execution, the execution diagram, message and setup is saved to the execution log.

- To view the execution log for an execution package, click the **Execution** tab, right click the execution package and then click **View Execution History Log**.

RESUMING A FAILED EXECUTION

Executions that fail are a fact of life, but restarting an execution that has failed can be very time consuming if the error occured two hours into the execution. Therefore, Jet Data Man-ager enables you to resume a failed execution from the point of failure.

This is useful since, with managed execution, you cannot always determine in what order Jet Data Manager will execute tables. This means that if an execution fails, a complete restart of the execution is usually the only way to ensure that everything is executed cor-rectly.

You can also configure Jet Data Manager to allow some non-essential data sources to fail without failing the entire execution at the same time. See [Allowing a Data Source to Fail](#)

When an execution fails, resume the execution by following the steps below.

1. First, you need to identify the error. On the **Execution** tab, right click the failed exe-cution package, and click **View Execution History Log**. The **Execution Log** window opens.

Start	End	Total Time	Succeede	Gantt Chart	Execution Message	Execution Package
2016-07-06 14:49:29	2016-07-06 14:49:33	00:00:04	False	View	View	View
2016-07-06 14:48:59	2016-07-06 14:49:03	00:00:04	True	View	View	View
2016-07-06 14:48:36	2016-07-06 14:48:40	00:00:04	True	View	View	View
2016-07-06 14:46:34	2016-07-06 14:46:39	00:00:04	True	View	View	View
2016-06-21 10:46:06	2016-06-21 10:46:10	00:00:04	True	View	View	View
2016-04-07 12:39:59	2016-04-07 12:40:10	00:00:10	True	View	View	View
2016-03-01 12:28:13	2016-03-01 12:28:19	00:00:06	True	View	View	View
2016-02-22 11:34:03	2016-02-22 11:34:09	00:00:06	True	View	View	View
2016-02-22 11:28:45	2016-02-22 11:28:51	00:00:06	True	View	View	View

2. Click **View** in the **Execution Message** column next to the failed execution to display the error message. Close the message and the Execution Log window.
3. Second, you need to solve the error you identified in the execution message. Remember to deploy any changed objects as necessary.
4. Now you are ready to resume the execution. Open the **Execution Log** again as described in step 1.
5. Right click the execution you wish to resume and click **Resume Execution**. The **Execute** window opens.
6. Click **Start** to begin the execution.
7. Repeat steps 1-6 if the execution package fails again.

EXECUTING OBJECTS WITH THE EXECUTION QUEUE

The Execution Queue enables you to continue working while tables or you entire project is executed in the background.

OPENING THE EXECUTION QUEUE WINDOW

Navigate to the **Project** ribbon tab, locate the **Development** group and click **Execution Queue** to open the Execution Queue window.
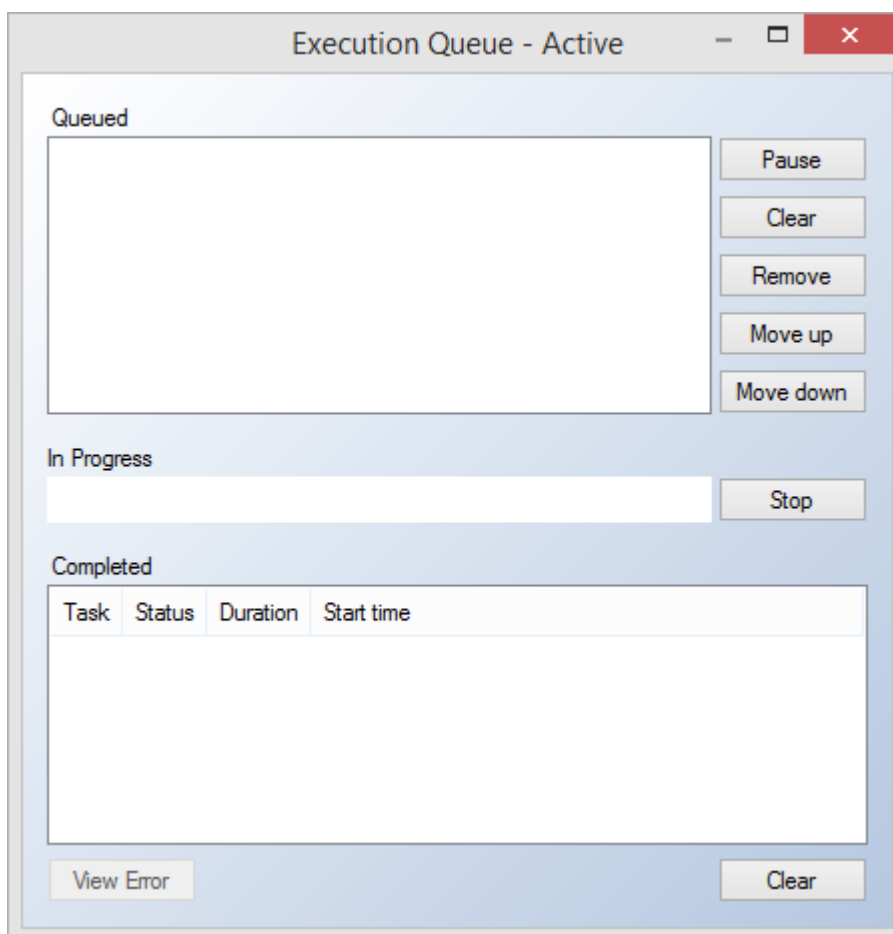
ADDING AN OBJECT TO THE EXECUTION QUEUE

Adding an object to the Execution Queue is a simple drag-and-drop operation.

1. Drag-and-drop a table, a perspective, a data warehouse, a business unit, an execution package or another executable object to the Execution Queue window. A window appears to let you choose which execution steps from the object to add to the queue.
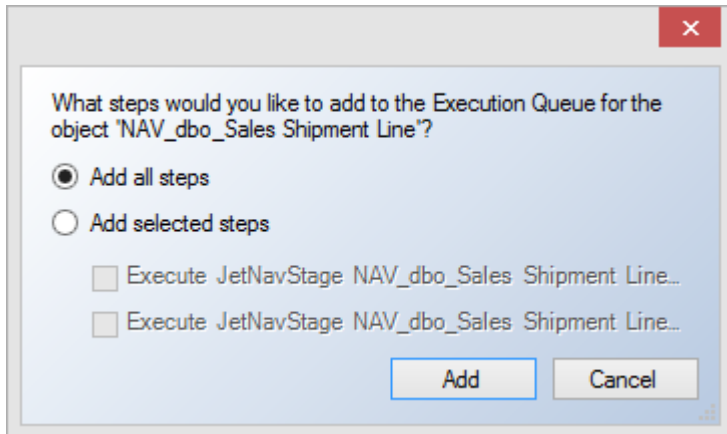
2. Select **Add all steps** or **Add selected steps** and choose which steps you would like to add to the queue. Click **Add** to add the object to the queue.
3. The object is now queued in the Execution Queue. If there is no other items in the queue, the object will be moved to **In Progress** and begin executing immediately.

PAUSING AND RUNNING THE QUEUE

The Execution Queued mode can de toggled between running and pause using the button in the top right corner of the window. When the queue is running, the button is called **Pause**. Clicking the button prevents further objects from being executed and changes the button text to **Resume**. Pressing **Pause** does not stop an object that is currently in progress. Pressing the **Resume** button resumes executing of the queue.

MOVING AND REMOVING QUEUED ITEMS

The **Queued** list shows the items waiting to be executed.

The queued objects can be moved up and down in the list by selecting the item and using the **Move up** and **Move down** buttons. The top item in the list is the next to be execut-ed.

An object can be removed from list by selecting it and clicking **Remove**. Clicking **Clear** removes all items from the list.

STOPPING CURRENT EXECUTION

In Progress shows the object currently being executed. Pressing **Stop** halts the execution of the object and pauses execution of the queue.

REMOVING EXECUTED ITEMS AND VIEWING ERRORS

The **Completed** list shows the objects that have been executed. It lists the **Status** of the individual items, the **Duration** and the **Start Time**. The items can have one of four statuses:

- **Success**: The object was executed without errors.
- **Failed**: The execution was ended prematurely by an error.
- **Stopped**: The execution was stopped by the user before it completed.

You can view error messages for failed objects by selecting the object in the list and clicking **View Error**. This brings up a message box displaying the error message.

CLOSING THE EXECUTION QUEUE WINDOW

You can close the Execution Queue window by clicking the X in the top right corner.

Closing the Execution Queue window or closing the entire project does not stop or pause the execution of the queued objects. It only hides the window, while the Execution Queue will continue working in the background. You can open the Execution Queue window again to review the status of the objects in the queue or to add more objects to the queue.

When you close Jet Data Manager, the Execution Queue will be stopped along with the rest of the application.

SCHEDULED EXECUTION

When your project is ready, you would typically like the entire project, or part of it, to be executed according to a schedule. One common scenario is to execute the project every night to ensure that the business users have updated numbers in the morning.

No matter what you needs are, the answer in Jet Data Manager is the execution package. An execution package contains tables, cubes and other objects that will be executed when the package is run. You can schedule the package to be run at specific times and set up notifications to e.g. alert you if a scheduled execution fails. You can also set up conditions that need to be fulfilled for a execution package to run.

> **Note:** An execution package will only execute the objects in the package, not deploy them. If changes have been made to the project, but have not been successfully deployed, the scheduled execution package will most likely fail.

ADDING AN EXECUTION PACKAGE

An execution package determines which objects in a project will be executed and how. To create an execution package, follow the steps below.

1. On the **Execution** tab, right-click **Execution Packages** and then click **Add Execution Package**.

   The **Add Execution Package** window opens:
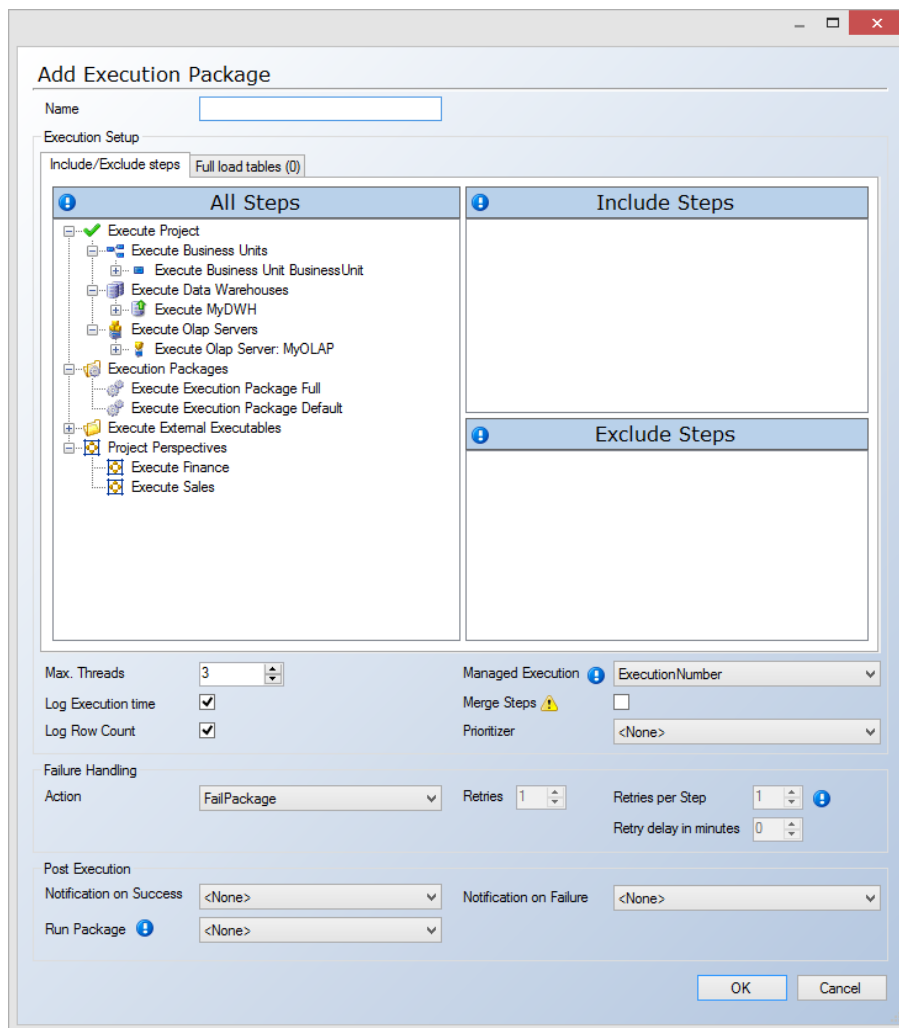
2. Type a **Name** for the execution package.
3. In the **Include/Exclude Steps** tab, choose the steps you wish to include in or exclude from the execution package by dragging objects from **All Steps** to **Include Steps** or **Exclude Steps**, respectively. For instance, simply drag the **Execute Project** step to **Included Steps** if you wish to execute the entire project when the package is executed. If you wish to exclude the entire project except one or more steps, simply drag those steps to **Exclude Steps**. Right click an object and click **Remove Step** to remove a step from the **Include Steps** or **Exclude Steps**.
4. (Optional) In the **Full load tables** tab, drag any incrementally loaded tables, you wish to have full loaded in this execution package, from **All Tables** to **Full Load Tables**. You can also drag business units, data warehouses, OLAP servers or the entire project.
5. Enter the maximal number of threads you wish to utilize during execution in **Max. Threads**.
6. Clear the **Log Execution time** and/or the **Log Row Count**check boxes if you do not wish to log this information.

7. Select a setting for **Managed Execution**. You have the following options:
   - **Disabled**: Managed execution is disabled. Objects will be executed in the order specified in the project tree.
   - **Execution Number**: When more than one object is ready to executed, Jet Data Manager prioritizes the objects based on their position in the project tree from top to bottom.
   - **Classification**: When more than one object is ready to executed, Jet Data Manager prioritizes the objects based on their table classification. The order will be "Fact Table – Large", "Dimension Table – Large", "Fact Table", "Dimension Table". If two tables have the same classification Jet Data Manager will use the execution number as the secondary criteria.
   - **Execution Time**: When more than one object is ready to executed, Jet Data Manager prioritizes the objects based on their average execution time so that the object with the longest execution time is executed first. If two tables have the same execution time (e.g. in case of new objects), Jet Data Manager will use the execution number as the secondary criteria. When execution time of the objects in the project are known, this option will result in the shortest execution time in most cases.
8. Check **Merge Steps** if you wish to treat all individual sub-steps of the chosen steps as one big collection. This can speed up execution.
9. (Optional) In the **Prioritizer** list, click the prioritization you wish the execution package to use. For more information on prioritization, see Adding a Prioritization.
10. Under **Failure Handling**, select what **Action**Jet Data Manager should perform if the execution fails. You have the following options:
    - **Fail Package**: When a step fails, the execution is stopped and the package is declared failed.
    - **Retry Step**: When a step fails, the step will be retried until the maximum number of retries for the entire package or the individual step is reached. Enter the maximum number of retries allowed for the package in the **Retries** box and the maximum number of retries allowed for an individual package in **Retries per Step**. Enter the amount of time to wait between retries in **Retry delay in minutes**.
11. Under **Post Execution**, select a Notification on Success and a Notification on Failure. You have to create a notification before it is available from the list. See Adding Notifications below.
12. If you wish to run a package after the execution, select the package in **Run Package**.
13. Click **OK** to add the execution package.

ADDING AN EXECUTION SCHEDULE

When you specify an execution schedule, the execution process is started automatically at the specified time.

Note: Because of the way the scheduler service works, you should make sure to schedule packages at least two minutes apart. If you schedule two packages to be executed at the same time, only one of the packages will be executed.

To add an execution schedule, follow the steps below.

1. On the **Execution** tab, expand **Execution Packages**, right-click the relevant execution package, and click **Add Schedule**.
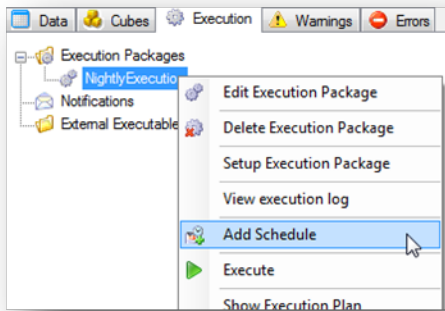
2. You have three different options for specifying the schedule:
   A. To set up multiple daily executions in a specific time frame, click **Daily** and then click **Frequency**. In the **Start time** field, enter the start time of the time interval. In the **End time** field, enter the end time of the interval. In the **Run every** field, specify the number of hours and minutes between each project execution.
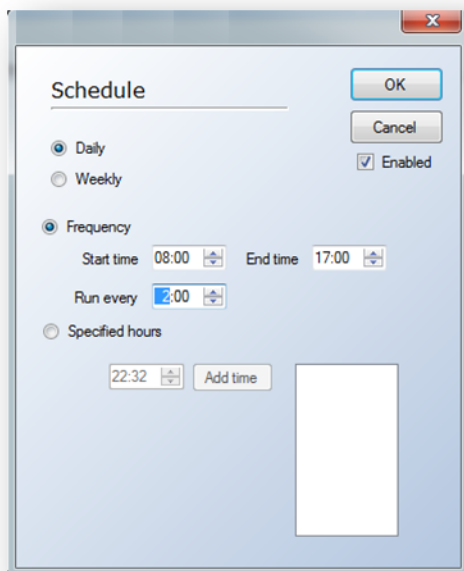
   B. To set up one or more daily executions on specified hours, click **Daily** and then click **Specified Hours**. Enter the exact time for the execution to run and then

click **Add Time**.
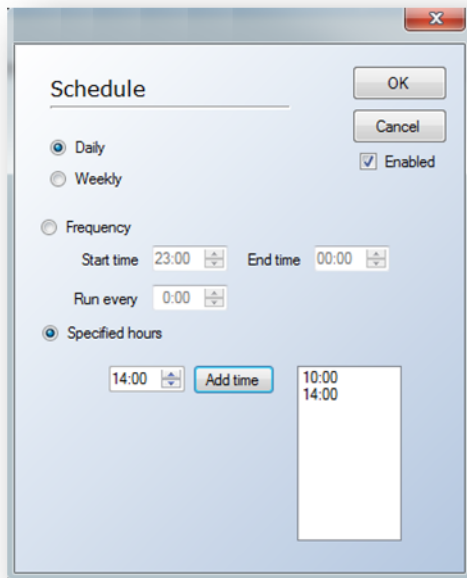
C. To set up weekly executions, click **Weekly**. Select the day(s) when the project
    should execute, enter the **Start time** for the execution, and then click **Add Time**
    to add the time to the schedule. Repeat this step for each day and time that you
    want to execute the project. A project can be executed several times a day and
    several times during the week. To view the entire weekly schedule, click **Show
    All**.
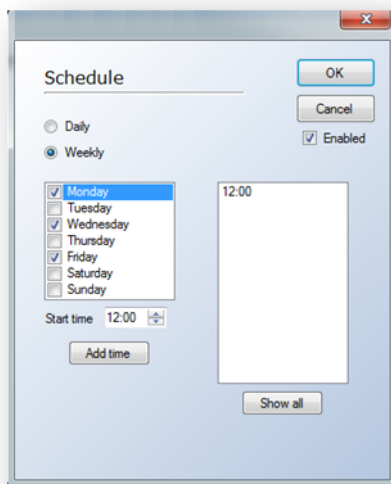
3. Select **Enabled**to activate the schedule, and then click **OK**.

Note: You can get an overview of the previous executions of an object by right-clicking on it an clicking on **View Execution overview Log**. Additionally, all events that happen during execution are registered in the Windows Event Log.
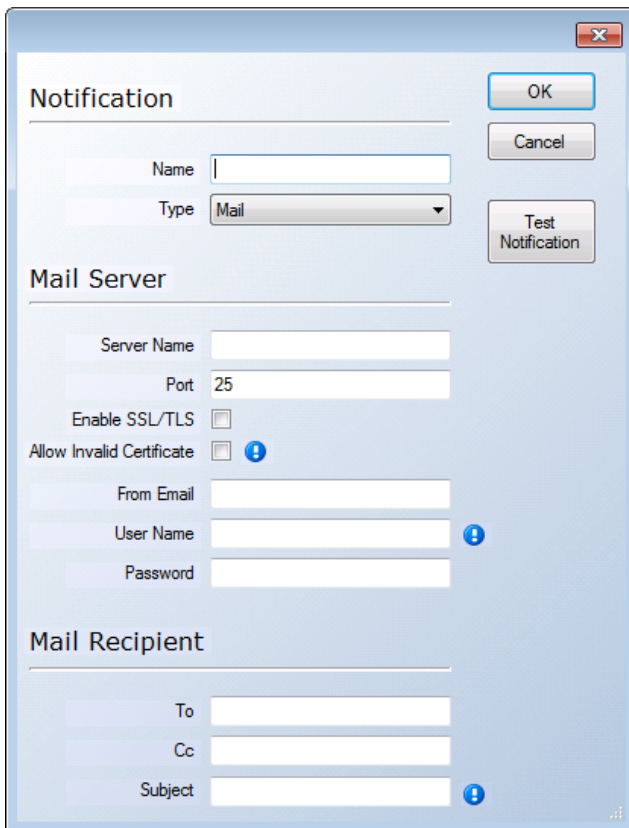
ADDING A NOTIFICATION

Notifications can be used to alert specified individuals when the execution package was successfully run or in case something caused it to fail. Notifications are most commonly set up as email alerts, but can be saved to the Event Log as well.

1. On the **Execution** tab, right-click **Notifications**, and click **Add Notification**.

   A window opens to let you add a notification.

2. Enter a **Name** for the notification and select the **Type** the type of notification you want to create. You have the following options:

Option	Description
Mail	Creates an email notification
EventLog	Writes a notification to the event log
Both	Creates both an email notification and writes to the event log

3. Enter the information for the SMTP server you wish to use under **Mail Server**.
4. Enter the e-mail addresses of the recipients and a **Subject** under **Mail Recipient**. In the subject, you can use the following variables:
    - **%Project%**: The name of the project .
    - **%Status%**: The status of the execution (Success / Fail).
    - **%ExecutionPackage%**: The name of the execution package.
5. Click **OK**.

The notification can now be selected when you create an execution package.

ADDING A PRIORITIZATION

You can choose what objects to prioritize during a managed execution. This is useful if you, for instance, only have a small timeslot for extracting data from a source, or I you would like to have a certain cube ready for the users as early as possible during execution.

To set up a new prioritization for use in an execution package, follow the steps below.

1. On the **Execution** tab, right click **Prioritizations** and click **Add Prioritization**. The **Add Prioritization** window opens.
2. In **Name**, type a name for the prioritization.
3. Drag and drop an object from the **Available Objects** tree to **Selected Objects** to add this object to prioritized objects.
4. Click an object in the **Selected Objects** list to have the execution steps for that object displayed under **Object Settings**. Select or clear the individual steps under **Selected steps** to configure what steps will be prioritized in the execution. For instance, you might only want to prioritize the transfer of data from a specific table because your priority is to have the transfer completed as soon as possible.
5. Select **Blocking** if you wish for the execution to halt all other execution tasks until the selected steps for the selected object has been completed.
6. Click **OK** to close the window and create the new prioritization, which will be added to the tree.

ADDING A USAGE CONDITION

If you wish to execute an execution package only under certain conditions, you can add a usage condition to the execution package. For instance, if you use multiple environments, you could have an execution package execute only in your production environment. To add a usage condition to an execution package, follow the steps below.

1. Add the project variable you wish to use in your usage condition. See [Adding a Project Variable](#).
2. On the **Execution** tab, expand **Execution Packages**, right-click the relevant execution package, and click **Add Usage Condition**. The **Usage Condition** window appears:
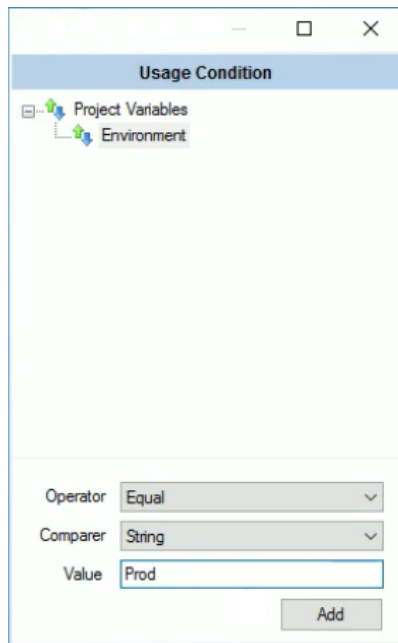
3. The available project variables are listed in the window. Click the variable you wish to use. In the example above, the "environment" variable has been selected.
4. In the **Operator** list, click the operator you wish to use.
5. In the **Comparer** list, click the data type you wish to use when comparing the variable to the value.
6. In the **Value** box, type the value you wishg to compare the variable against.
7. Click **Add** to add the usage condition.

If try to manually execute a package with a usage condition that evaluates to "true", a warning message will pop up. The same message will be written to the log if you try to execute the package in a scheduled execution.

INCREMENTAL LOADING

Incremental loading facilitates faster load times by enabling you to load only new transactional data into the data warehouse and staging databases. This is useful when the volume of transactional data in the data sources causes scheduled execution to take too long.

You should consider enabling incremental loading if your scheduled executions run longer than your execution strategy allows. For instance, if you load data during the night when business is closed to have fresh data in the morning, the load obviously needs to be completed before the users begin their day. If you wish to transfer data from sources that are in use, you might also like to use incremental loading to minimize the impact on performance data transfer can have on users.

The default load plan in Jet Data Manager is a full load. During a full load, the existing tables in the staging database and data warehouse are truncated, which removes all of the existing data, and new data is subsequently loaded from the data sources.

While it can be slow, the advantage of full loading is that you can be sure that any changes made in the source systems are carried over to the data warehouse. Therefore, it is common strategy to schedule periodical full loads while otherwise utilizing incremental loading. For instance, you could set up a full load during the weekend to make absolutely sure that your data warehouse is up to date, and schedule incremental loads during the workweek.

Incremental loading can be either source-based or target-based. Source-based incremental loading is the most efficient form, but it requires data that contains fields that can be used to identify which records are new and should be transferred. If the data do not contain such fields, target-based incremental loading offers a less powerful alternative.

SOURCE-BASED INCREMENTAL LOADING

The purpose of incremental loading is to minimize the amount of data that needs to be loaded into the data warehouse, thus speeding up the loading process and minimizing any issues caused by putting a load on source systems. The key is to avoid transferring data that is already stored in the data warehouse. If your data contains one ore more fields that can define what records are new, you can utilize source-based incremental loading.

During the first deployment after incremental loading has been enabled, Jet Data Manager will create additional tables in the staging database and data warehouse that have an _ INCR or _I suffix. Jet Data Manager will then do a full load. During subsequent executions of the project, truncation is disabled so that the data from the full load is not removed. Using the field, or number of fields, that you have chosen, Jet Data Manager determines which records have been added to the data source since the last load and only transfers new records to the staging database and data warehouse.

Incremental load is enabled on the table level. Naturally, you will get the greatest increase in performance by enabling it on tables with a large amount of transactional data, such as those that contain large volumes of general ledger and inventory transactions. On the other hand, smaller tables with relatively few records, such as a Customer or Item table, usually do not take long to execute in the first place and enabling incremental load on them will only give you marginally better overall performance.

ENABLING SOURCE-BASED INCREMENTAL LOADING

To utilize source-based incremental loading on a table, the table must contain a field that represents new data. This could be an identifier field, an entry number, or possibly a date. I addition to that, the table must have a primary key defined. To define a field or fields to be used as the primary key for the table, right-click a field and click **Include in Primary Key**.

To enable source based incremental loading, follow the steps below.

1. Right click the table for which you wish to enable incremental loading and click **Table Settings**.
2. Click the **Data Extraction** tab and select **Enable Source Based Incremental Load**. If an error icon appears next to the setting, it means that another setting needs to be changed to enable source based incremental loading. Move you mouse over the error icon to see the error message.
3. Click **OK**. The table icon will now be overlaid with an "I" to make it easy for you to identify it as an incrementally loaded table.

4. If the table belongs to the staging database, right click the corrosponding source table under **Data Sources** in the project treeand click **Add Incremental Selection Rule.** - OR - If the table belongs to the data warehouse, right click the table and click **Add Incremental Selection Rule**.

5. The **Source Based Incremental Values** pane appears. Select the fields identifying which records have been added or changed since the last incremental load.
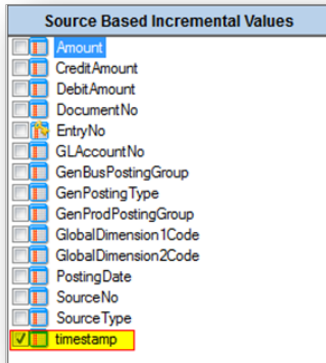
The fields you choose should ideally be fields that are generated by the system and incremented sequentially when new records are added. The table below lists recommended choices in some popular systems.

System	Recommended field
Dynamics NAV	Timestamp
Dynamics GP	DEX_ROW_ID
Dynamics AX	RECID

6. *Repeat the steps above for all tables you wish to enable incremental loading on.*
7. Right-click the **Business Unit**, click **Deploy and Execute** and click **Deploy and Execute Modified Tables and Views**. The **Deploy and Execute Modified Tables**
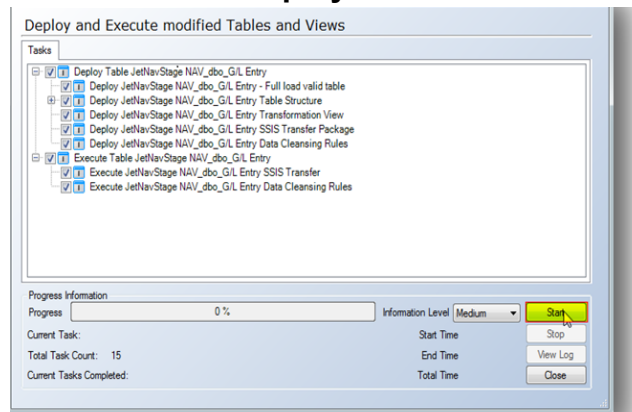
   and Views window opens.
8. Click **Start**to begin the first full load of the tables with source based incremental loading now enabled. The necessary incremental tables will be automatically added to the staging database and populated with the latest incremental values. The next time the table is executed, Jet Data Managerwill query these tables to determine the last record that was previously loaded and will only extract data from the data source was added after the last execution.

241

TARGET-BASED INCREMENTAL LOADING

Target-based incremental loading is primarily used when there are no identifying fields that determine which records have been added since the last incremental update. With target based incremental loading, all of the data is transfered from the data source. Records are then compared against the existing records in the table and only new, updated, or deleted records are added to the staging database or data warehouse.

Target-based incremental loading is not as fast as source-based incremental loading, but is faster than a full load strategy.

ENABLING TARGET-BASED INCREMENTAL LOADING

To enable source-based incremental loading for at table, follow the steps below.

1. Right click the table for which you wish to enable incremental loading and click **Table Settings**.
2. Click the **Data Extraction** tab and select **Enable Target Based Incremental Load**. If an error icon appears next to the setting, it means that another setting needs to be changed to enable target based incremental loading. Move you mouse over the error icon to see the error message.
3. Click **OK**. The table icon will now be overlaid with an "T" to make it easy for you to identify it as an target-based incrementally loaded table.

In the project tree, a new item, **Incremental Settings**, is added below the list of fields.
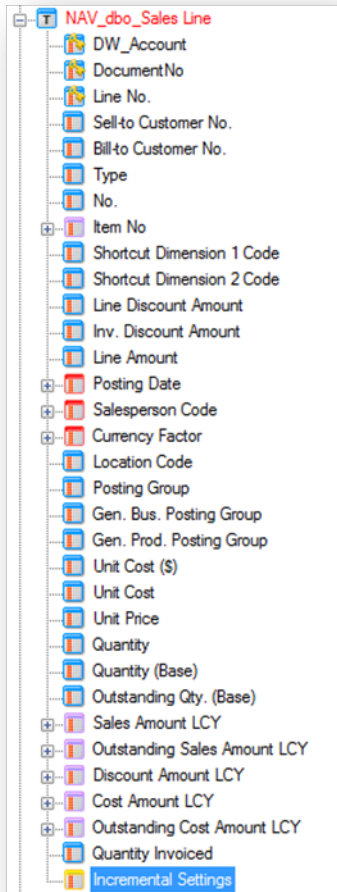
4. Click **Incremental Settings**. Three pane appear in the right-hand side of the application. In the first pane, **Target Based Incremental Keys**, the field or fields that represent the primary key for the table should be checked.
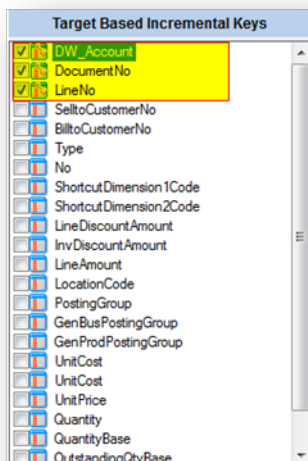
In the second pane, **Target Based Value Keys**, you should select the fields that would indicate a change in this table. Jet Data Manager will create a hash key field

based on the values of the fields you have selected and use this to determine if a record has been updated. In this example, all fields have been selected.
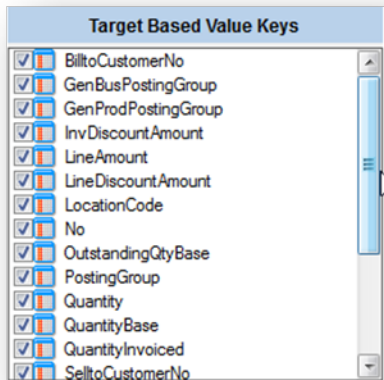
In the third pane, **Incremental Events**, you should select the events on the data source Jet Data Manager will take into consideration during execution.

- **Inserts** are any new records that have been added to the data source, based on the primary key on the table.
- **Updates** are any records that have modified or been changed since the last incremental load. This is determined based on the **Target Based Value Keys** selected.
- **Deletes** are any records that previously existed in the data source, but no longer exists based on the primary key. These will be removed from the staging database or data warehouse table.

5. Right-click the table and click **Deploy and Execute** to perform the initial load of the table.

The **Deploy and Execute** window opens.

244
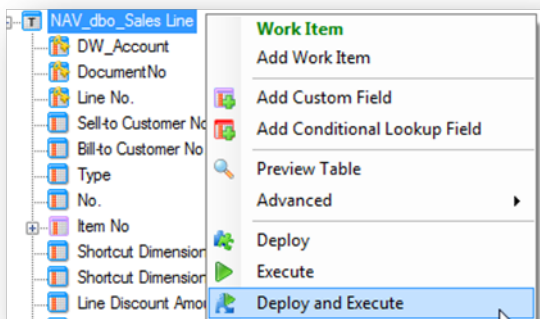
6. Click **Start** to begin the deployment and execution process. During this process, all of the data will be loaded, and the hash keys for the target-based incremental load will be generated by Jet Data Manager. Subsequent executions of the table will only load records that, depending on your settings, have been added, modified, or deleted from the data source since the last execution.

MULTIPLE ENVIRONMENTS

With multiple environments, you can have a dedicated development environment with automatic transfer of the latest version of your project to the production environment. This ensures that the production environment is always online and available for end-users.

A dedicated development environment enables you to work within non-production environments. This is useful when an organization needs to ensure that the production environment is always available for end-users. For example, the organization could have an environment called "Development" where changes are made, dimensions are updated, and measures are created. Once these modifications are tested, they can be transferred to the live production environment directly from Jet Data Manager.

PREREQUISITES

Before setting up development environments, ensure that the following prerequisites are met:

- All servers used in the development and production environments must have the same version of Jet Data Managerinstalled. This applies to bit version(32-bit or 64-bit) as well.

- Ensure that Jet Data Manager service is installed and started on the server(s) you wish to deploy on. Detailed instructions are provided below.

- Ensure that a project repository has been created on all the servers that are used in the production environment and development environment.

- The user account(s) that will be used to set up the multiple environments **may** need to have Read permissions to the Event Log. This can be set up in the Registry Editor under the "HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\eventlog" node. You can right-click the "eventlog" node, select **Permissions**, and add the users that will utilize Multiple Environment Deployment and assign them "Read" permissions.

SETTING UP A DEVELOPMENT ENVIRONMENT

The following example shows multiple environments using a single **Development** server and a single **Production** server.
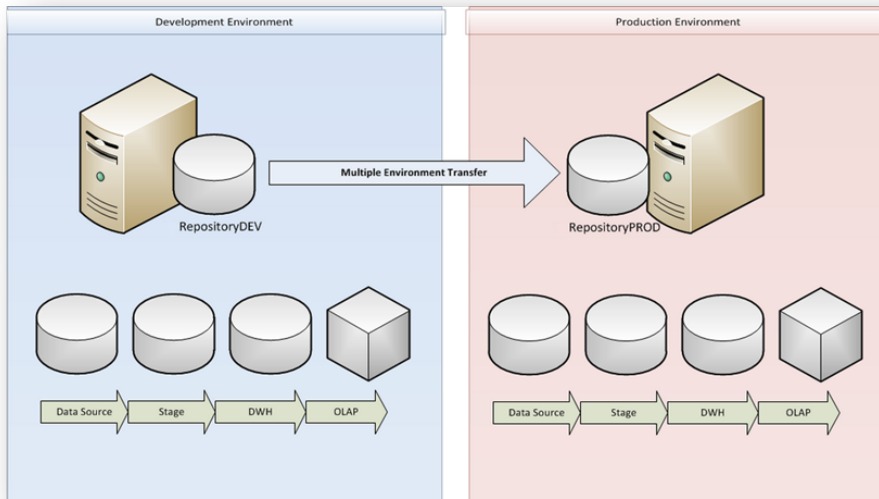
Note: You can set up as many environments as you need. Setup of additional environments follows the same steps as listed below.

SETTING UP THE PRODUCTION SERVER

The first step is to set up the **Production** environment on the production server.

1. Log on to the production server and open Jet Data Manager.
2. In the ribbon, on the **Tools** tab in the **Environment** group, click **Environment Properties**

3. This window will open. Click**New Environment**.

4. The **Add Environment** window will open:
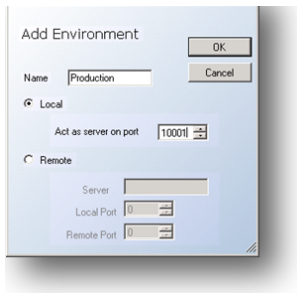
- Type a **Name**. In this example, the name is "Production".
- Select **Local**, as deployment will only be done into this environment and not from it.
- Enter a port to use in the **Act As Server On Port** box. Make sure the port is free to avoid any conflicts on the network.

5. Close Jet Data Manager.

6. Next you need to make sure that the Jet Data Manager erver service is correctly set up. Click/Right-click **Start**, click **Run**, type **Services.msc** and click **OK**. Locate the server service in the list. It will be named Jet Data Manager **Server [version]**.

7. Right click the service and click **Properties**.

8. On the **General** tab, in the **Startup type** list, click **Automatic**.

9. On the **Log On** tab, click **This account**. In **This Account**, enter the account that was used for setting up the production environment, i.e. the account you are logged in with. Type the password for the account in the **Password** and **Confirm password**

boxes. Click **OK**.
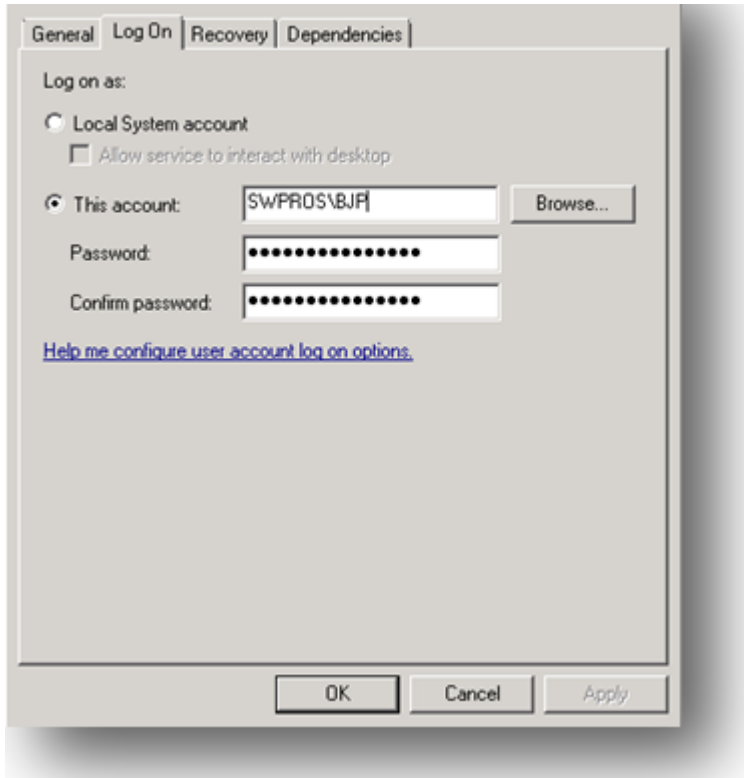
10. Right click the service in the list an click **Start**.
11. Log off from the production server.

SETTING UP THE DEVELOPMENT SERVER

The next step is to set up the development environment on the development server.

1. Log on to the development server and open Jet Data Manager.
2. In the ribbon, on the **Tools** tab in the **Environment** group, click **Environment Properties**

3. This window will open. Click **New Environment** to create the development environment.

4. The **Add Environment** window will open:
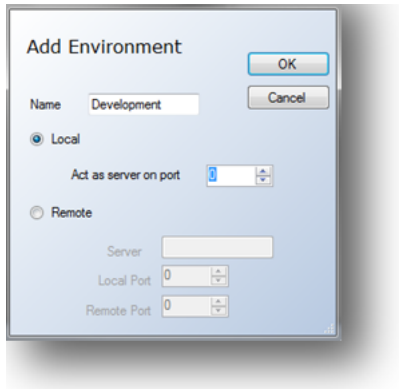
- Type a **Name**. In this example, the name is "Development".
- Select **Local**, leave **Act as Server on Port** set to 0 and click **OK**.

5. Click **New Environment** again to specify the production environment on the development server. This is so that the development environment knows where the production environment exists.
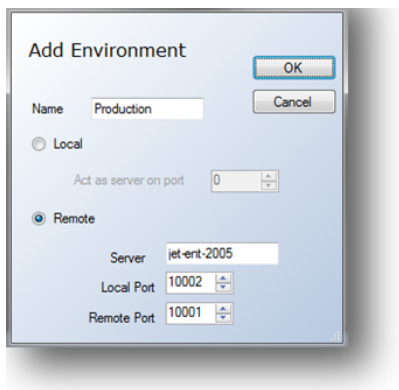
- Type a **Name**.
- Click **Remote** and fill in the remote server information:
    - **Server**: Type the server name or IP address of the production server.
    - **Local Port:** Enter any open port that is not being used by another application.
    - **Remote Port:**Enter the port you selected when setting up the production environment on the production server.

6. Click **OK**.

CREATING GLOBAL DATABASES

Global databases allow Jet Data Manager to know where the related databases reside for the Production and Development environments. For example, the location of the Staging Database for both the Production and Development environments will be specified.

1. Click **New Global Database** from the Environment Dialog window.

2. The **Add Global Database** window will open.
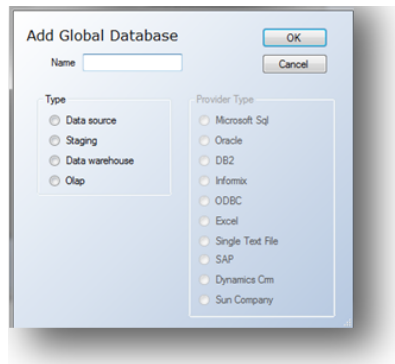
3. Within the Add Global Database Dialog, you will be creating a series of databases that will be used in your project. You will create the following Global Databases:
   - Data Source
   - Staging Database
   - Data Warehouse
   - OLAP
4. Assign a name to your data source, select **Data Source** in the **Type** section and select the relevant **Provider Type**. In this example, we will name our global database "GlobalDataSource" and use a provider type of Microsoft SQL.
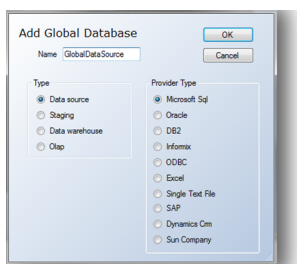
5. Repeat the previous step for all databases in the project (Staging, Data warehouse, and OLAP). Note that you do not need to select a Provider Type for staging, data warehouse, and OLAP.

   Your results should look similar to those shown below:
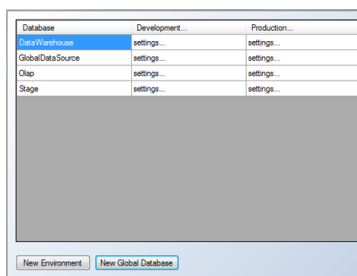
CONFIGURING GLOBAL DATABASES

Within the **EnvironmentProperty**window you should have a data source, data warehouse, OLAP, and staging database. Each environment, Development and Production, has a "**Settings..."** section for each database. You can also click the environment name to access additional settings.

CONFIGURING THE DATA SOURCE

1. Select the **"Settings…"** field from the data source row in the Development column from the **Environment Properties** window. This will display the Settings pane to the right.
2. Enter the following information:
   - Server: This will be the server address of the development server. Since this is currently on the development server, this can be localhost or the name of the server.
   - Name of the database from which data is extracted. This will be the name of the NAV, AX, GP, or other database. In our example, this is JetCorpDemo.

3. Next click "**Settings…"** on the data source row in the **Production** column. Enter the following configuration:
   - **Server:** This will be the name of the server on the Production Environment. In our example, the server name is jet-ent-2005.
   - **Catalog:** This will be the name of the database from which data is extracted, for instance your ERP system. In our example, this is "JetCorpDemo".

Note: If you are using your live ERP database for extracting data in both the development and production environments, then the server name and catalog in both the Development and Production columns will be the same.

CONFIGURING THE STAGING DATABASE

1. Next you will need to configure another Global Database for the staging. Click **"Settings…"** on the **Stage** row in the **Development** column to display the **Settings** pane to the right.
2. Enter the following configuration:
   - **Server**: This will be the server address of the development server. In our example, this is localhost.
   - **Catalog**: This will be the name associated with the staging database in the development environment. In our example, we use StageDev.

3. Next click **"Settings..."** on the **Stage** row of the **Production** column.
4. Enter the following configuration:
   - **Server:** This will be the name of the server for the Production Environment. In our example, the server name is "jet-ent-2005".
   - **Catalog:** This will be the name associated with your staging database in the production environment. In our example, we use StageProd.

CONFIGURING THE DATA WAREHOUSE

Next you will need to configure another Global Database for the data warehouse.

1. Click "Settings...: on the data warehouse row in the Development column to display the Settings pane to the right.
2. Enter the following configuration:
   - **Server:** This will be the name of the server for the development environment. In our example, this is localhost.
   - **Catalog:** This will be the name associated with your data warehouse in the development environment. In our example, we use DataWarehouseDev.

3. Next Click "Settings..." on the data warehouse row of the **Production** column.
4. Enter the following configuration:
   - **Server**: This will be the name of the server for the Production Environment. In our example the server name is "jet-ent-2005".
   - **Catalog:**This will be the name associated with your data warehouse in the production environment. In our example, we use "DataWarehouseDev".

CONFIGURING THE OLAP DATABASE

Next you will need to configure another Global Database for the OLAP cubes.

1. Click "Settings..." on the OLAP row in the Development column to display the Settings pane to the right..
2. Enter the following configuration:
   - **Server:** This will be the server address of the development server. In our example, this is "localhost".
   - **Catalog:**This will be the name associated with the OLAP database in the development environment. In our example, we use "OlapDev".


3. Next, Click "Settings..." on the OLAP row of the **Production** column.

4. Enter the following configuration:
   - **Server:** This will be the name of the server for the Production Environment. In our example, the server name is "jet-ent-2005".

- **Catalog:** This will be the name associated with your OLAP database in the production environment. In our example, we use "OlapProd".

CREATING THE GLOBAL DATABASES

The final step in the configuration process is to test and create the global databases on SQL Server. This can be done from inside the Environmental Properties window. This needs to be done for both the development and production environments. Right-click "Settings..." and select Test Connection. If you get an error message, it generally means that the database has not been created yet. Right-click "Settings...", and select Create Database. Then retest the connection.

Perform this check on all Global Databases for both the Development and Production environments. Once this has been completed and all "Test Connection" responses return "Connection OK", click the OK button to close the Environment Properties window, and save all changes.

Note: The data source does not have the option to Create Database. This database represents the data source that Jet Data Manager is extracting from and will already exist in your infrastructure. An example of this will be your NAV, GP, or AX database.

CONFIGURE PROJECT CONNECTIONS

The environments have now been set up, and the global databases have been configured. The next step is to configure the connections in the project to utilize these Global Databases.

1. Open your project, and navigate to Data Sources at the bottom of the Data Tab. Right-click the adapter and select Edit Microsoft SQL Provider.

   Note: This will vary depending on your data source type.

2. Select **Use Global Database** for the data source, choose the Global Database that represents your data source, and click OK. There will generally be only one Global Database displayed in the drop-down list.


3. Navigate to your Staging Database, right-click the database, and select Edit Staging Database.

4. Select **Use Global Database** for the data source, choose the Global Database that represents your staging database, and click OK. There will generally be only one Global Database displayed in the drop-down list.


5. Navigate to your **Data warehouse**, right-click the database, and select Edit Data Warehouse.


6. Select **Use Global Database** for the data source, choose the Global Database that represents your data warehouse, and click There will generally be only one Global Database displayed in the drop-down list.

7. Navigate to the OLAP Database on the **Cubes** tab, right-click the OLAP Database, and select **Edit OLAP Server**.

8. Select **Use Global Database** for the data source, choose the Global Database that represents your OLAP database, and click OK. There will generally be only one Global Database displayed in the drop-down list.
9. The final step is to deploy and execute the project to ensure your project is properly configured and ready for transfer. From the **Project** tab, click **Manual Deployment and Execution**, and then click **Start**.


TRANSFER THE PROJECT FROM DEVELOPMENT TO PRODUCTION

You are now ready to transfer the project from the development to the production environment.

1. Log in to the **Development Environment** server, and open Jet Data Manager.
2. On the **Tools** tab, click **Multiple Environment Transfer**.

3. Click **Transfer** to migrate the project from the development server to the production server.

A dialog will appear asking you to confirm the transfer. Click **OK**.
4. Deploy the project on the production environment.
   - If you are transfering the project from development to production for the first time, or if you simply wish to deploy all objects, right click the **Production Environment** folder and click **Deploy**. When the **Deployed** line in the **Production** column changes to "Yes", the process has finished.
   - If only some objects were changed in the development environment and are in need of deployment, right click the **Production Environment** folder and click **Partial Deployment**. The **Remote Deployment Window** opens. A list of deployable objects is displayed. Select the objects you wish to deploy and click

Deploy. While deployment is under way, the **Partial Deploy** window displays the deployment status. When deployment has finished, a window opens with a list of the deployment tasks completed. Click **Close** to close the window

5. Click **Close** to close the **Multiple Environment Transfer** window.

EXECUTION PACKAGES

Execution packages will automatically update the staging database, data warehouse, and OLAP cubes on a scheduled basis. Since projects deployed from the Development Environment will replace packages in the Production Environment, it is recommended that the desired execution packages be set up in the Development Environment. This way, they are seamlessly transferred to the Production Environment with the package transfer. It may not be desirable to have automatic execution enabled in the development environment. This can be disabled by ensuring that the "Jet Data Manager Server Scheduler" service is disabled on the machine hosting the development environment. For more information regarding the configuration of Execution Packages, see Execution Packages.