



## Pentaho Metadata Editor User Guide



This document is copyright © 2011 Pentaho Corporation. No part may be reprinted without written permission from Pentaho Corporation. All trademarks are the property of their respective owners.

## Help and Support Resources

If you have questions that are not covered in this guide, or if you would like to report errors in the documentation, please contact your Pentaho technical support representative.

Support-related questions should be submitted through the Pentaho Customer Support Portal at <http://support.pentaho.com>.

For information about how to purchase support or enable an additional named support contact, please contact your sales representative, or send an email to [sales@pentaho.com](mailto:sales@pentaho.com).

For information about instructor-led training on the topics covered in this guide, visit <http://www.pentaho.com/training>.

## Limits of Liability and Disclaimer of Warranty

The author(s) of this document have used their best efforts in preparing the content and the programs contained in it. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, express or implied, with regard to these programs or the documentation contained in this book.

The author(s) and Pentaho shall not be liable in the event of incidental or consequential damages in connection with, or arising out of, the furnishing, performance, or use of the programs, associated instructions, and/or claims.

## Trademarks

Pentaho (TM) and the Pentaho logo are registered trademarks of Pentaho Corporation. All other trademarks are the property of their respective owners. Trademarked names may appear throughout this document. Rather than list the names and entities that own the trademarks or insert a trademark symbol with each mention of the trademarked name, Pentaho states that it is using the names for editorial purposes only and to the benefit of the trademark owner, with no intention of infringing upon that trademark.

## Company Information

Pentaho Corporation  
Citadel International, Suite 340  
5950 Hazeltine National Drive  
Orlando, FL 32822  
Phone: +1 407 812-OPEN (6736)  
Fax: +1 407 517-4575  
<http://www.pentaho.com>

E-mail: [communityconnection@pentaho.com](mailto:communityconnection@pentaho.com)

Sales Inquiries: [sales@pentaho.com](mailto:sales@pentaho.com)

Documentation Suggestions: [documentation@pentaho.com](mailto:documentation@pentaho.com)

Sign-up for our newsletter: <http://community.pentaho.com/newsletter/>

# Contents

About this Document.....	5
Knowledge Prerequisites.....	5
Introducing the Pentaho Metadata Editor.....	6
The Physical Layer.....	7
The Abstract Business Layer.....	7
The Business View.....	7
Incorporating Metadata.....	7
Installing the Pentaho Metadata Editor.....	9
Launching the Pentaho Metadata Editor.....	9
Adding a JDBC Driver.....	9
Creating a New Domain.....	11
Setting Up a Database Connection.....	11
An Example Connection: A JDBC Connection to a MySQL Database.....	13
An Example Connection: A JNDI Connection to a Hypersonic Database.....	13
Importing Physical Tables and Columns.....	13
Importing Tables Inside a Schema.....	14
Creating a Business Model.....	15
Creating Business Tables and Columns.....	15
Removing Unnecessary Business Columns.....	17
Creating Relationships between Business Tables.....	17
Creating Relationships Using the Editor Graph.....	19
Hadoop Hive-Specific SQL Limitations.....	20
Building a Business View.....	20
Building a Business View Using Manage Categories.....	20
Enriching Your Data — Understanding Metadata Concepts and Properties.....	23
Exceptions to the Rules.....	23
Required Properties.....	23
Applying Concepts and Properties in the Pentaho Metadata Editor.....	23
Applying Self Concepts: Managing Properties on the Model.....	24
The Properties List.....	24
The Properties Editor (Settings).....	24
Adding New Properties.....	25
Removing a Property.....	25
The Concept Editor.....	25
Building Concepts.....	26
Applying Concepts to Business Objects.....	26
Removing the Parent Concept.....	27
Adding Security to Metadata Business Objects.....	28
Defining Security Settings.....	28
Getting Security Settings Offline.....	29
Changing Security Constraints.....	29
Adding Row Level Security to a Pentaho Metadata Model.....	30
Applying Data Constraints.....	30
Importing and Exporting Domains.....	32
Importing a Domain.....	32
Exporting a Domain.....	32
Configuring and Managing Locales (Internationalization).....	33
Setting Up Locales.....	33
Domain Backup and Recovery.....	34
Publishing a Domain to a Pentaho BI Server.....	35
Before you Publish Your Domain.....	35
Publishing a Domain.....	35
Sample Metadata Model Walkthrough.....	37
Examining the Default Metadata Model.....	37

Creating a New Domain File and Connecting to the Database.....	38
Importing Physical Tables and Columns.....	39
Setting Security.....	40
Creating Your Business Model.....	41
Assigning Business Tables to Physical Tables.....	42
Defining Relationships between Business Tables.....	43
Creating Derived Columns (Part 1).....	44
Creating Derived Columns (Part 2).....	45
Creating the Business Model.....	46
Creating the Business View.....	46
<b>Managing Categories</b> .....	46
Testing Your Business Model with MQL Query Builder.....	47
Publishing Your Business Model to the Pentaho BI Server.....	48
Examining the Model in the Pentaho User Console.....	49
<b>Metadata Properties Reference</b> .....	52
Out-of-the-Box Properties.....	52
Custom Properties.....	53
Required Properties per Business Object.....	54
<b>Pentaho Metadata Formulas</b> .....	55
Supported Functions.....	57
Supported Operators.....	58
Supported Aggregate Functions.....	58
<b>Troubleshooting</b> .....	59
Hive Database Disappears From Database Connection Dialogue.....	59

# About this Document

---

This document provides you with information and instructions associated with using the Pentaho Metadata Editor, a tool that allows you to build metadata models and domains. Also included, is an exercise that walks you through creating a metadata model using a small sample database.

See [Sample Metadata Model Walkthrough](#) for more information.

## Knowledge Prerequisites

---

To use Pentaho Metadata Editor, you must have database administrator skills (such as knowing how to import tables, create relationships between tables, assign aggregations, add categories, assign security, etc.) and extensive knowledge of your database(s). You must also know what type of data interests your business users so you can determine whether a column should be a measure or a dimension, for example. These skills allow you to map a business user model (logical model) to a complex relational database. Ultimately, the metadata models you provide to business users allows them to create Pentaho reports without requiring DBA assistance.

See [Examining the Default Metadata Model](#) for information about how metadata models are implemented in the Pentaho BI server.



**Important:** The sample images in this document are of the Pentaho Metadata Editor version 3.0. If you are using a newer version of the product you may notice slight differences in the user interface; however, overall product functionality is similar in both versions.

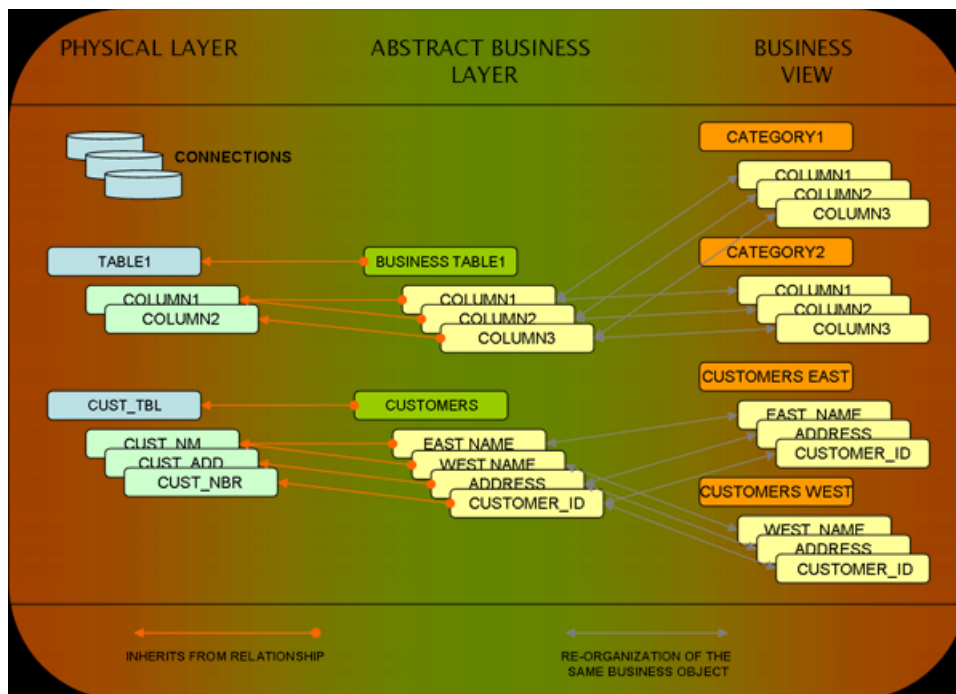
# Introducing the Pentaho Metadata Editor

Pentaho Metadata Editor (PME) is a tool that allows you to build Pentaho metadata domains and models. A Pentaho Metadata Model maps the physical structure of your database into a logical business model. These mappings are stored in a centralized metadata repository and allow administrators to:

- Create business-language definitions for complex or cryptic database tables
- Decrease the cost and impact associated with low level database changes
- Set security parameters limiting user's report access to data
- Drive formatting on text, date and numeric data improving report maintenance
- Localize the information to the user's regional settings

The goal of the Pentaho Metadata Editor is to simplify the experience of business users when they are create reports.

The metadata business model is actually one major component in a Pentaho Metadata Domain. The domain encapsulates both the physical descriptions of your database objects and the logical model (the business model), the abstract representation of the database. The purpose of this page is to give a Pentaho metadata consumer a good idea of the relationships involved throughout a Pentaho domain, in order to fully leverage the power of the business models. The following diagram depicts the business objects and relationships within the Pentaho metadata domain:



In this diagram, both the relationship arrows in the legend as well as the color of the business objects are significant. Each independent business object has its own color; thus, connections, physical tables, business tables, physical columns, business columns and categories are the main business objects in a Pentaho domain. The two types of relationships depicted in the diagram are key to understanding Pentaho metadata. The "inherits from" relationship is a relationship between two different business objects where one will inherit metadata from the other, with the ability to override any of the inherited metadata properties. An example of this is that a business table inherits metadata from its associated physical table. The second relationship, the "same object" relationship, is one where the two business objects are actually one and the same, just represented in a different organization or duplicated for usability. This is the relationship between business columns in the Abstract Business Layer, and business columns in the Business View.

**Important:** Although the Pentaho Metadata Editor displays a checkbox for each of the permissions available in the Pentaho BI Server, only **Execute** is enforced today. If a user has Execute permission, he or she will be able to view that resource; otherwise, the resource is hidden. Additional permissions may be enforced in a future release.



## The Physical Layer

---

The **Physical Layer** of a Pentaho domain encompasses connections, physical tables and physical columns. These objects represent the database(s) you are trying to model and enrich with metadata. The Physical Layer is not considered part of the business model, because not all connections defined in the Physical Layer will be used in every business model.

Multiple business models can be created in one domain, but those models must have one and only one connection reference. This means that you cannot mix and match physical tables from two different connections in the same model yet. We realize that this prevents the model from supporting multiple data sources, as well as severely limits the Pentaho Metadata Editor's ability to allow table changes across connections, a feature necessary for moving from dev to production databases, for example. Fortunately, you can get around the latter with a bit of hand-editing XML. Note that these features are important to us and are early inclusions in the metadata roadmap.

## The Abstract Business Layer

---

The **Abstract Business Layer** is the heavy lifter in the metadata business model. The business model encompasses the Abstract Business Layer and the Business View. In the Abstract Business Layer, you have business tables, business columns, and business relationships.

You can create business tables for any physical table you have defined in the Physical Layer. You can also create more than one business table to reference the same physical table. The same rules apply for business columns. This can be useful in a multitude of scenarios, filtering security or even data at this level being one example. The business table keeps a reference to the physical table that it models, and this allows a metadata inheritance relationship between physical tables and business tables. If you define metadata on a physical table, the business table will inherit that metadata, unless and until the business table itself has overridden the inherited metadata. This concept operates on a metadata property to property basis, meaning that each property can be overridden individually.

The same relationship exists between physical columns and business columns. If you define metadata on a physical column, the business column will inherit that metadata, unless and until the business column itself has overridden the inherited metadata. Then there are relationships. Don't confuse the term relationships here with the relationships described in the domain diagram. The relationships described here are not represented in that diagram. These relationships are mappings that you define to describe the relational (or other) bonds between your business tables. One example may be a one to many relationship between a customer table and an orders table. The strength in metadata relationships is that you can identify relationships between columns or tables in the Abstract Business Layer that may not be obvious in the Physical Layer. An example would be to compare budget, actual and sales numbers, using a formula-derived business column that is specific to your business rules.

## The Business View

---

The Business View is the part of the business model that applications will operate against, and end-users will see. The Business View is nothing more than "buckets" (called categories) for you to re-arrange and re-organize your business columns in a fashion that makes sense to the consumers of the data.

In the Business View, you can create any number of categories and arrange your business columns in those categories however it best suits your business, mixing and matching columns that derived from different business tables, even adding business columns more than once to different categories. The only restriction to the Business View is you cannot add the same business column more than once to a single category.

## Incorporating Metadata

---

Each business object in the domain can have metadata associated with it, with the exception of categories. In Pentaho terminology, a collection of metadata properties is called a concept.

Each business object can have three levels of **concepts**: its own (self or child concept), a parent concept and an inherited concept. All of the defined metadata properties from the three levels are available (to metadata-aware applications, end-users) on a business object. It is important to understand what the override hierarchy is, when more than one concept level has defined the same metadata property.

When a metadata property is set in the business object itself, and that same property is in play (either inherited or from a parent concept), the business object's self or child concept overrides all others. The next level in the hierarchy is the parent concept. So, for example, suppose the name property on a business object is *not* set on the object itself, but *is* inherited from its physical counterpart. In this instance, you set a parent concept on the business object with a name property defined in the parent concept. The parent concept's name property overrides the inherited name property. Ultimately, the inherited metadata properties are used if, and only if, the same property is not defined somewhere down the hierarchical chain, as part of a self concept or parent concept.



# Installing the Pentaho Metadata Editor

---

Installing the Pentaho Metadata Editor is as simple as extracting the files from the archive file that you downloaded from the [Pentaho download site](#). You may extract the files to any directory of your choice. There are a few additional requirements for the Pentaho Metadata Editor, listed below.

## JRE

You must have a JRE installed on your computer, version 1.5 or higher. The bin directory for the JRE should be registered on the system path.

## Pentaho BI Server

Having access to a Pentaho BI server is only a requirement if you want to set up security in your metadata model, or if you want to publish your metadata domain to a server. Otherwise, you can work with all other features in the Pentaho Metadata Editor without a Pentaho BI server available.

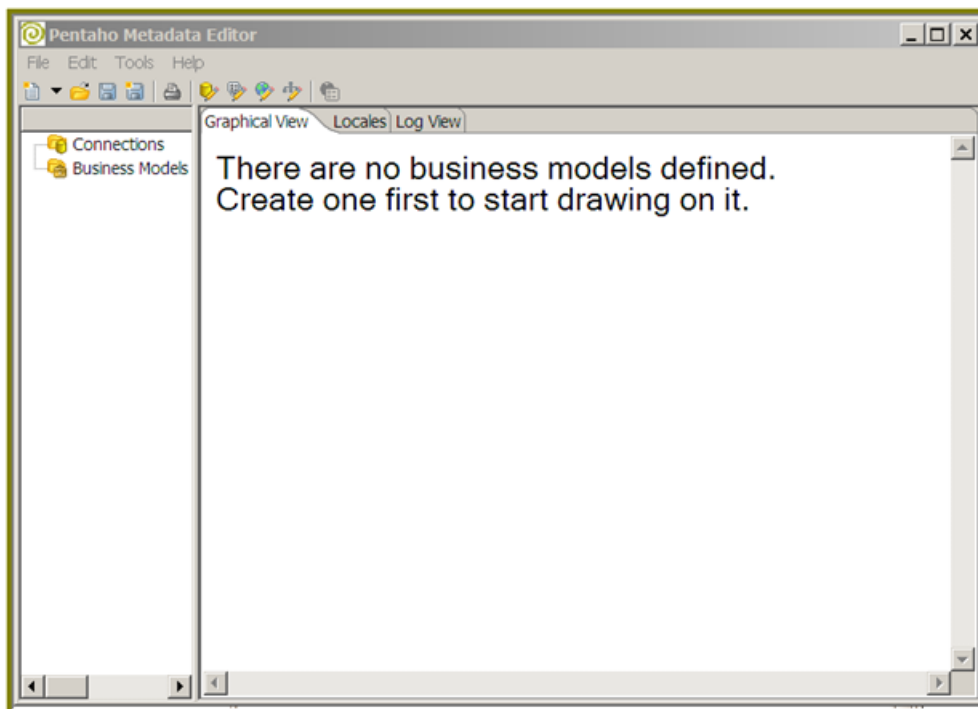
## Launching the Pentaho Metadata Editor

---

Once you have downloaded and extracted the files from the metadata archive, you should see two launch scripts in the root directory:

- **metadata-editor.bat** (for Windows)
- **metadata-editor.sh** (for all other platforms)


Execute the script that is appropriate for your platform, either by double-clicking on the file name in an explorer window, or as a command from a shell or command prompt. You should see the initial Pentaho Metadata Editor workspace, as shown below:



## Adding a JDBC Driver

---

Before you can connect to a data source in any Pentaho server or client tool, you must first install the appropriate database driver. Your database administrator, CIO, or IT manager should be able to provide you with the proper driver JAR. If not, you can download a JDBC driver JAR file from your database vendor or driver developer's Web site. Once you have the JAR, follow the instructions below to copy it to the driver directories for all of the BI Suite components that need to connect to this data source.


 **Note:** Microsoft SQL Server users frequently use an alternative, non-vendor-supported driver called JTDS. If you are adding an MSSQL data source, ensure that you are installing the correct driver.

## Backing up old drivers


You must also ensure that there are no other versions of the same vendor's JDBC driver installed in these directories. If there are, you may have to back them up and remove them to avoid confusion and potential class loading problems. This is of particular concern when you are installing a driver JAR for a data source that is the same database type as your Pentaho solution repository. If you have any doubts as to how to proceed, contact your Pentaho support representative for guidance.

## Installing JDBC drivers

Copy the driver JAR file to the following directories, depending on which servers and client tools you are using (Dashboard Designer, ad hoc reporting, and Analyzer are all part of the BI Server):

 **Note: For the BI Server:** before copying a new JDBC driver, ensure that there is not a different version of the same JAR in the destination directory. If there is, you must remove the old JAR to avoid version conflicts.

- **BI Server:** /pentaho/server/biserver-ee/tomcat/lib/
- **Enterprise Console:** /pentaho/server/enterprise-console/jdbc/
- **Data Integration Server:** /pentaho/server/data-integration-server/tomcat/webapps/pentaho-di/WEB-INF/lib/
- **Data Integration client:** /pentaho/design-tools/data-integration/libext/JDBC/
- **Report Designer:** /pentaho/design-tools/report-designer/lib/jdbc/
- **Schema Workbench:** /pentaho/design-tools/schema-workbench/drivers/
- **Aggregation Designer:** /pentaho/design-tools/agg-designer/drivers/
- **Metadata Editor:** /pentaho/design-tools/metadata-editor/libext/JDBC/

 **Note:** To establish a data source in the Pentaho Enterprise Console, you must install the driver in both the Enterprise Console and the BI Server or Data Integration Server. If you are just adding a data source through the Pentaho User Console, you do not need to install the driver to Enterprise Console.

## Restarting


Once the driver JAR is in place, you must restart the server or client tool that you added it to.

## Connecting to a Microsoft SQL Server using Integrated or Windows Authentication

The JDBC driver supports Type 2 integrated authentication on Windows operating systems through the **integratedSecurity** connection string property. To use integrated authentication, copy the **sqljdbc\_auth.dll** file to all the directories to which you copied the JDBC files.

The **sqljdbc\_auth.dll** files are installed in the following location:

```
<installation directory>\sqljdbc_<version>\<language>\auth\
```

 **Note:** Use the **sqljdbc\_auth.dll** file, in the x86 folder, if you are running a 32-bit Java Virtual Machine (JVM) even if the operating system is version x64. Use the **sqljdbc\_auth.dll** file in the x64 folder, if you are running a 64-bit JVM on a x64 processor. Use the **sqljdbc\_auth.dll** file in the IA64 folder, you are running a 64-bit JVM on an Itanium processor.

## Creating a New Domain

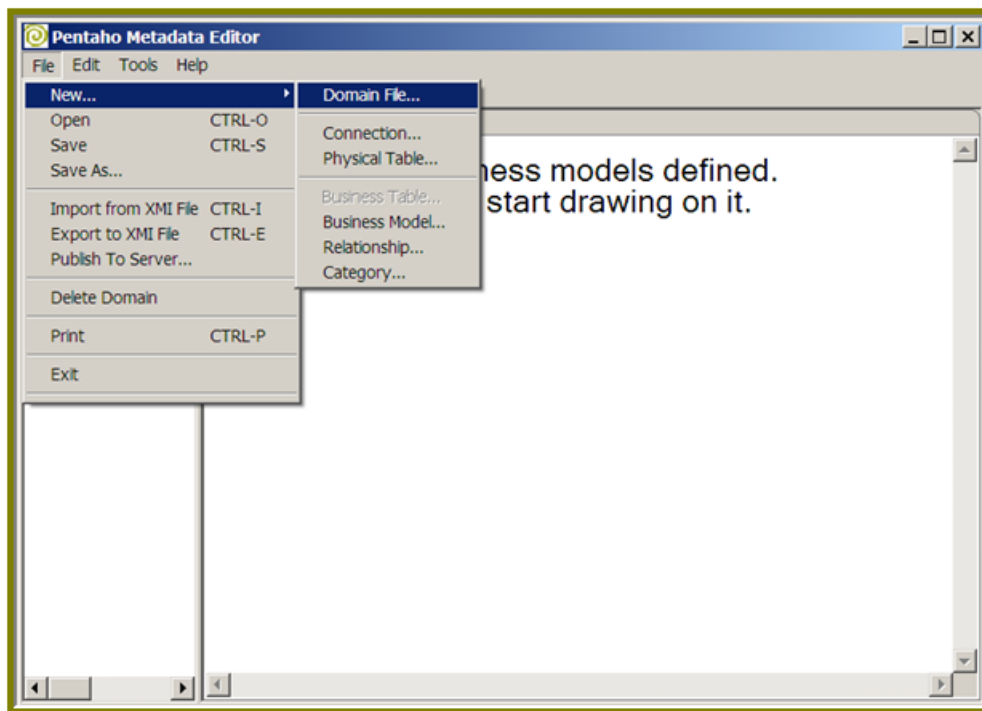
---

A metadata **domain** is a Pentaho term representing all of the business objects created, stored and used in the metadata layer. A domain may consist of one or more connections, one or more models, security information, business tables, business views, categories, columns and concepts. You can create and save multiple Pentaho metadata domains using the Pentaho Metadata Editor.

A metadata domain is accessed by the Pentaho BI server by exporting (or publishing) the domain to a file called, **metadata.xmi**, and placing the file in the Pentaho solutions repository at the root of a solution's directory structure.

In summary, a domain represents all of the associated modeled business entities. A domain can be viewed as metadata "document." Each solution is restricted to have at most one domain. Note that a solution repository can have multiple solutions. A domain must be published as a metadata.xmi file in the solution root. Pentaho Metadata Editor works with a domain at time only; for example, a "Sales" domain to define the relationships and entities used by sales team.

When you have first launch the Pentaho Metadata Editor, a new domain is automatically created for you. You can immediately begin adding connections, tables, columns, and more. If you want to start fresh with a new domain, select **File -> New -> Domain File** from the main menu.



There are several procedures associated with creating a new domain. These include:

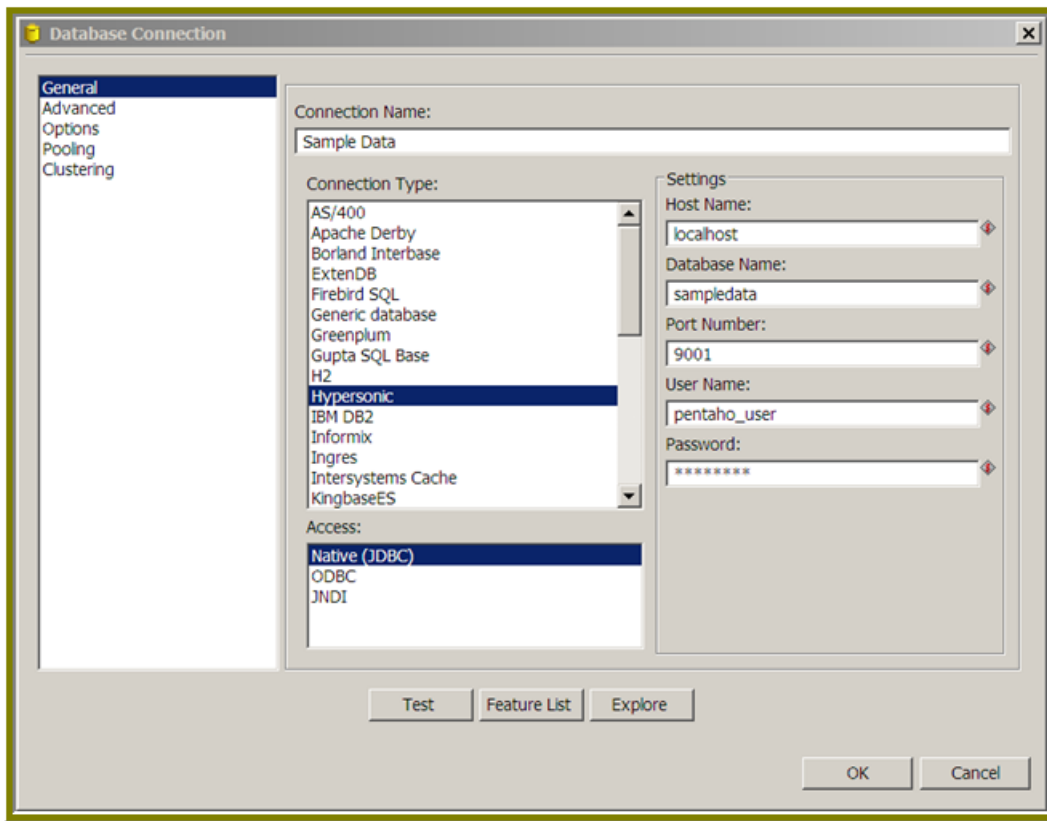
- Setting Up a Database Connection
- Importing Physical Tables and Columns
- Creating a Business Model
- Creating Business Tables and Columns
- Creating Relationships between Business Tables
- Building a Business View


## Setting Up a Database Connection

---

A connection represents connection information of a specific database, and acts as the parent in the hierarchy for all physical tables and physical columns that are defined for that database.

Pentaho metadata models can connect to most common relational databases using JDBC. The Pentaho Metadata Editor (and the Pentaho Metadata Architecture) supports a vast and rich set of data sources. Before you begin defining your business model, you must first describe the database or data source that you would like to model. You do this by defining one or more connections in the editor.



 **Note:** While the current implementation of Pentaho Metadata Editor supports multiple connections and multiple business models in the same domain, each business model must use physical objects (columns and tables) from a single connection only.

To create a new connection...

Step	Description
1	Right-click <b>Connections</b> navigation pane.
2	Select <b>New Connection</b> . The <b>Database Connection</b> dialog box appears.
3	Define your database connection as needed.
4	Click <b>Test</b> to ensure that your entries are correct.
5	Click <b>OK</b> to save your connection.

### Connection Types - Adding a New Database Type

First, note the **Connection Type** list in the dialog box. This is the list of database connections that the Pentaho Metadata Editor and metadata models can support. The list is quite extensive. If you do not see your database in the Connection Type list, you may be able to add it.

To add a new database type, you must copy the JDBC driver archive for your database into the PME install directory `...\metadata-editor\libext\JDBC`. Restart the Pentaho Metadata Editor, and you will see your database in the Connection Type list.

### Method of Access

Under Connection Type, you will see the **Access** list. Defining a JDBC or ODBC connection typically requires all of the remaining fields associated with the **General** tab to have the correct information; however, if you are into abstracting those details from your metadata domain, then use the JNDI method of access.

The JNDI access method keeps your server implementation cleaner as well; once you publish your domain to the server, as long as you have defined the JNDI connections with the same names, you still have a good implementation where your database information is only described to the JNDI layer. To take advantage of the JNDI method of access in the Pentaho Metadata Editor, you must define your database connection information in a properties file for the editor.

## An Example Connection: A JDBC Connection to a MySQL Database

This example is based on the Steels Wheels sample data. To connect to a MySQL database using JDBC

1. Name your connection, for example **MySQLSampleData**.
2. Select **MySQL** as the connection type.
3. Select **Native(JDBC)** as the access method.
4. Type **localhost** as the server host name. This assumes you have a local Pentaho BI Server running. If your server (or your sample database) is hosted elsewhere, type the name of that host instead.
5. Enter the name of the database, for example, **sampledata**.



**Note:** This example requires the default port for MySQL. Adjust the port as necessary for your setup.

6. Type the user name and password for the Steel Wheels sample data: **pentaho\_user**, **password**, respectively.
7. click **Test**. If all the information is entered properly, you should receive a "Test OK" message. If something is wrong, you should receive a message that is specific enough to help you determine the problem.
8. Click **OK** when you are done.

## An Example Connection: A JNDI Connection to a Hypersonic Database

Below are instructions for a sample JNDI connection to a Hypersonic database:

1. Name your connection, for example, **HySampleData**
2. Select **Hypersonic** as the connection type.
3. Select **JNDI**.



**Note:** Once you've selected JNDI, notice that the only remaining piece of information that you need to enter on the **General** tab is the database name.

4. Type **SampleData**.



**Note:** This name is actually the JNDI name that you will specify to map to the JNDI connection properties. In order for the Pentaho Metadata Editor to know what this name really means, you need to set up a simple JNDI properties file.

5. Navigate to [studio:PME install directory]/simple-jndi/jdbc.properties file. Open this file in your favorite text editor.
6. Type the following text into the jdbc.properties file:



**Note:**

```
SampleData/type=javax.sql.DataSource
SampleData/driver=org.hsqldb.jdbcDriver
SampleData/url=jdbc:hsqldb:hsqldb://localhost/sampledata
SampleData/user=pentaho_user SampleData/password=password
```

7. Save the file and close it.
8. Go back to the **Connection Information** dialog, and click **Test**.
9. Click **OK** when you are done.

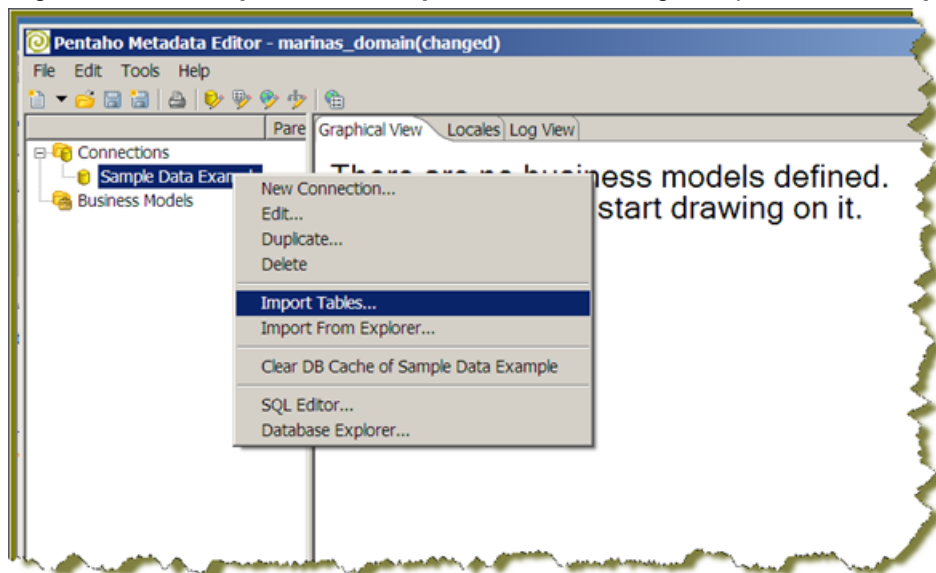
## Importing Physical Tables and Columns

After you define your database connections, the next step is to describe the physical database tables and columns that you want to include in your business model.

A **Business Model** contains all of the logical, abstract business objects and relationships that model your physical database objects in such a way that underlying changes to the physical database objects have minimal impact on your business, your business intelligence application, and your end users. There can be multiple business models in a single domain. A business model currently supports one database connection, and consists of business tables, relationships, and a multiple business views.


Fortunately, when you import a physical table, all of the table's columns come with it, so the import is a one step exercise instead of two. Later, you can remove the columns you do not want in the connection or the model. Below is an **example** of an import from HSQLDB (Hypersonic):

1. Right-click the **Sample Data Example** node in the navigation pane. Select **Import Tables** from the menu.



A dialog box that contains all physical tables available in the database appears.

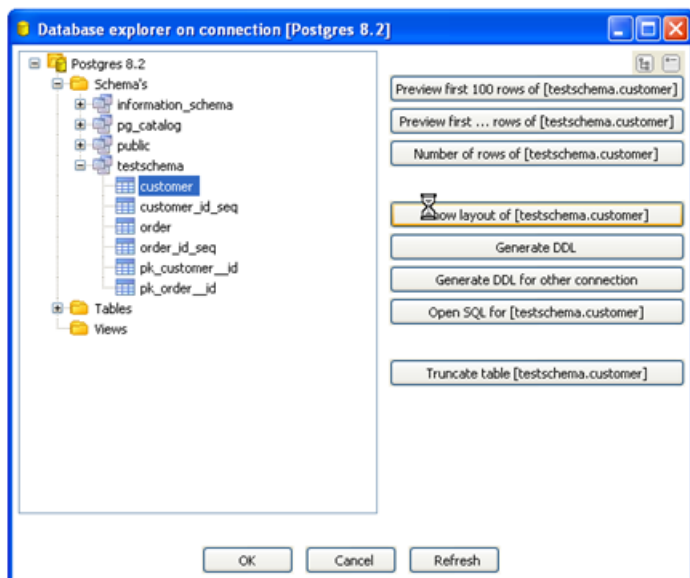
2. Select the tables you want to include.

 **Tip:** Press the <SHIFT> or <CTRL> key to select multiple tables.


3. Click **OK** when you are done.  
The tables you selected appear in a list in the navigation pane in the left side of the workspace.

## Importing Tables Inside a Schema

If you are accessing a table inside a schema, use the **Import from Explorer** command.



1. Right-click on the connection and select **Import from Explorer**.  
A dialog box that contains all physical tables available in the database appears.
2. Navigate under **Schema's** to locate the correct table(s) to import.

 **Important:** You must select the individual tables. Selecting the top level schema exclusively results in an error message.

3. Click **OK**.

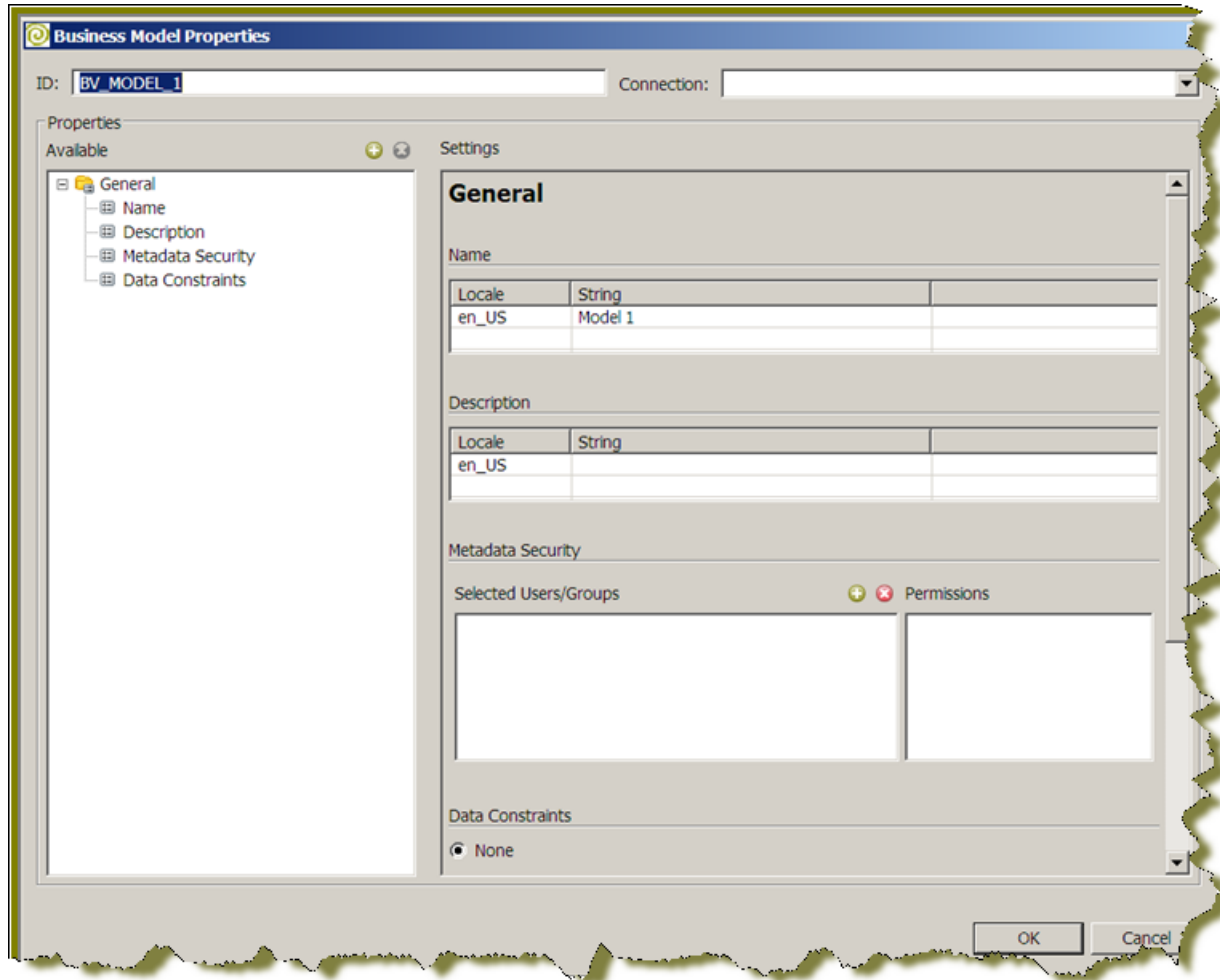
The tables you selected appear in a list in the navigation pane in the left side of the workspace.

## Creating a Business Model

The business model is where the logical mapping of business objects to the physical tables and columns you set up takes place. The model is also the place where you define the relationships between your business tables, and organize the business view.

To create a business model...

1. Right-click **Business Models** in the navigation pane and select **New Business Model**.



The Business Model Properties dialog box appears as shown above.

2. At the top of the dialog box, there is an ID text field, pre-populated with a value. Pentaho recommends you accept the pre-populated value as this value *must* be unique across all models that you define.
3. To name your new model, enter a new value under the **Name** property text box on the right.
4. Click **OK** to save your changes and exit the dialog box.

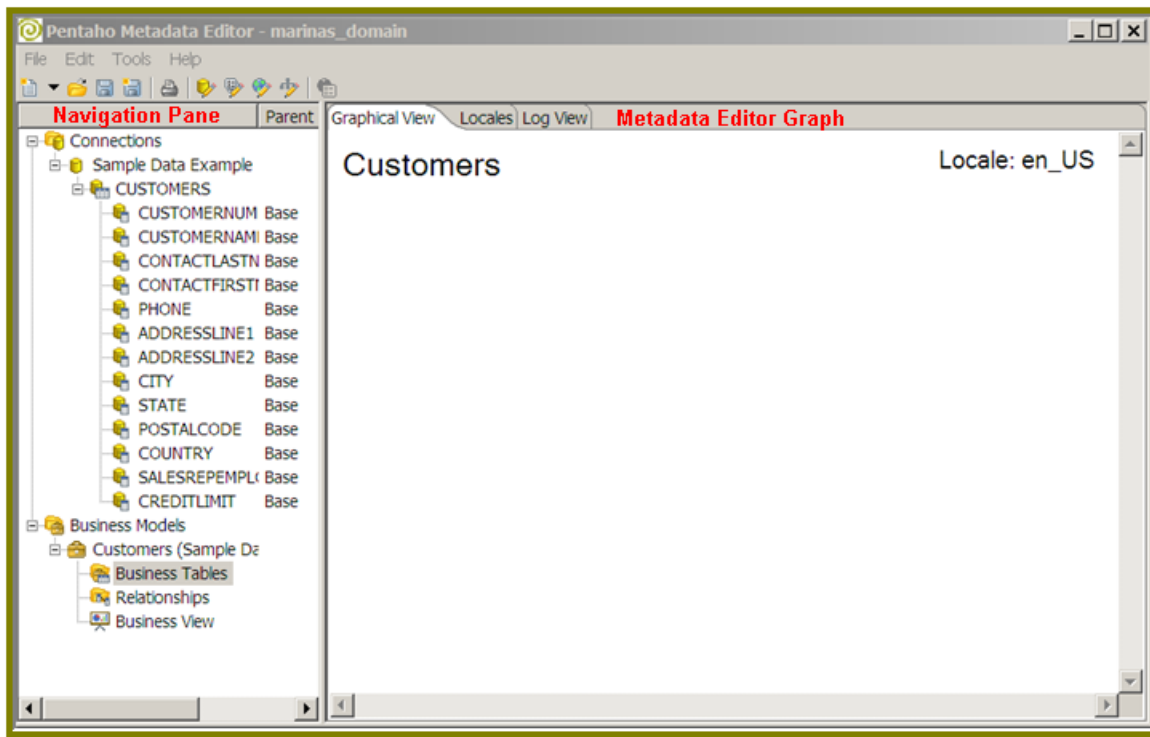
## Creating Business Tables and Columns

After you create a business model you must add the business tables and business columns; then, you create the relationships between your business tables.

Up to this point, you have used the navigation pane to build the business objects. The workspace on the right side of the page is called the **Metadata Editor Graph**. Use the Graph to lay out your business tables and show the




relationships between them. While you can accomplish the rest of your tasks by manipulating the objects in the navigation pane, you can also use the Graph.



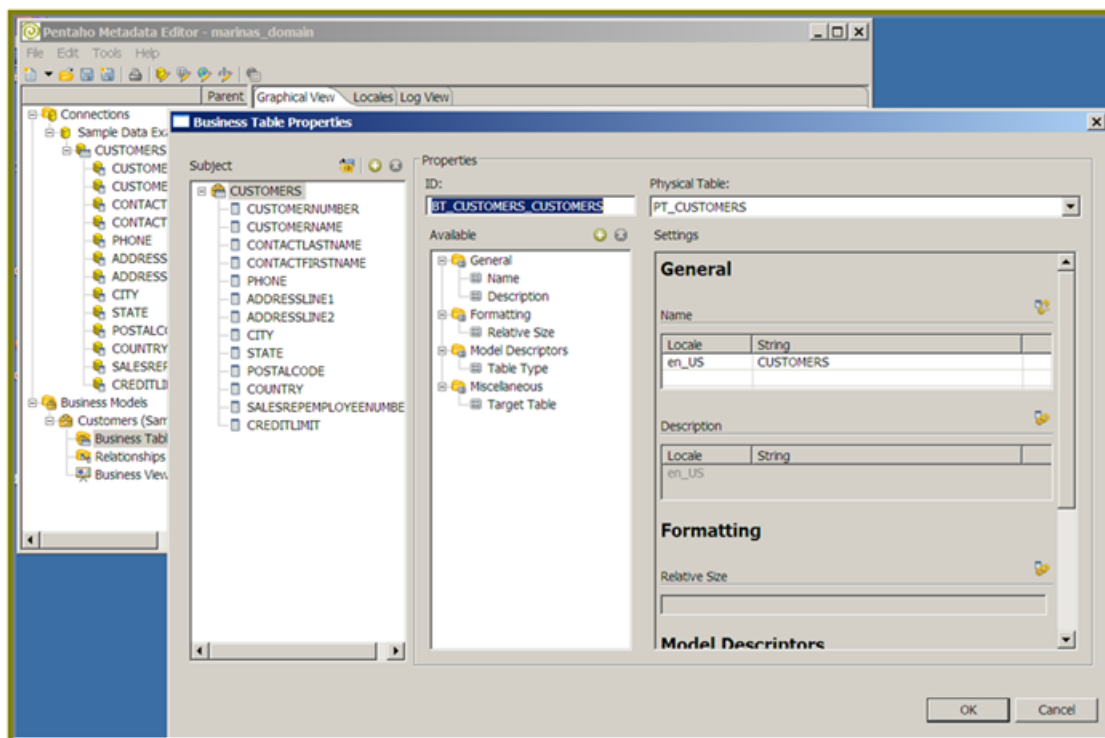
To create business tables and columns...

1. Right-click **Business Tables** in the navigation pane. Alternatively, you can right-click inside the **Metadata Editor Graph** workspace and select **Business Tables**.

 **Tip:** Drag the physical table from the navigation pane to the Graph workspace to create a new business table based on that physical table. Another shortcut is to drag the physical table in the navigation pane to the **Business Tables** node, which will also create a new business table based on that physical table.

A list of physical tables, each prefixed by the connection name to which they belong appears.

2. Select the physical table you want associated with the new business table.
3. Click **OK** to display the **Business Table Properties** dialog box.





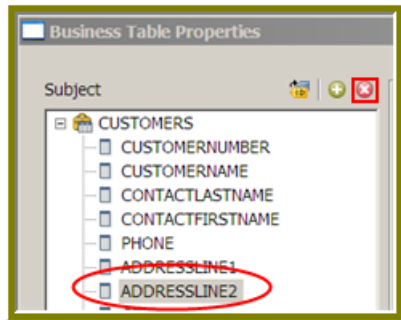
In the **Business Table Properties** dialog box, the **ID** field is pre-populated with a value. This ID identifies a specific business table, and must be unique. Pentaho recommends that you accept the default value for this field.


The table you chose in Step 3 is listed next, then below the fields is another navigation pane, with the name of your business table and all the columns inherited from the physical table included as business columns. The business columns are created for you when the new business table is created, so no need for the extra steps of creating them separately.

## Removing Unnecessary Business Columns

To delete unnecessary business columns...

1. In the **Business Table Properties** dialog box navigation pane, select the column you want to remove.



2. Click  (Delete Column).  
The column you selected is removed from the list.
3. Repeat Steps 1 and 2 with the remaining columns you want to delete.
4. Click **OK** when you are done.

## Creating Relationships between Business Tables

Once you have all of your business tables created, you must create relationships between the tables so that the query generators and SQL generators that work with Pentaho Metadata can create the data queries correctly. This is similar to drawing a relational diagram to show primary and foreign key relationships; however, relational links are not the only relationships that can be modeled. You can create a relationship between any two tables, link any two columns between them and dictate what the relationship is (one to many, many to many, and so on). The important pieces of information to know before you try to create a relationship are:

- What two business tables would you like to associate with this relationship?
- What columns in the business tables identify the relationship?
- What type of relationship is it — one to one, one to many, many to one, and so on?

To create a new relationship between business tables in the navigation pane, first make sure that the model you want to add this relationship to is selected, and that the **Relationships** node is visible.

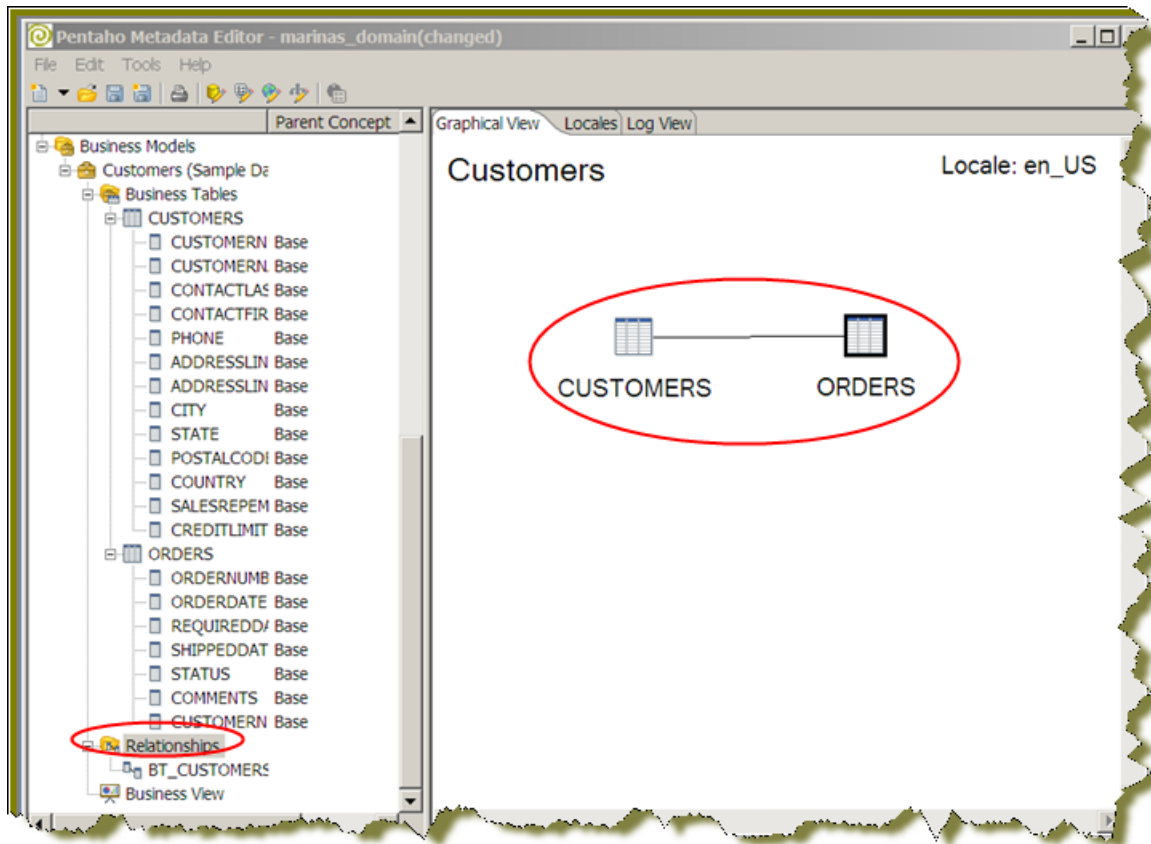
1. Right-click **Relationships** in the navigation pane.
2. Select **New Relationship**  
The **Relationship Properties** dialog box appears.
3. Select a business table from the **From Table/Field** list. This is the first relationship
4. Select a business table from the **To Table/Field** list. This steps sets up a relationship between two tables.
5. Now, specify the business columns (from the adjacent lists) from each business table that identify this relationship. If the business column names are similar, click **Guess Matching Fields** and let Pentaho Metadata Editor attempt to determine the columns for you.
6. Define the relationship from the **Relationship** drop down list. The relationships that can be chosen are defined with examples:

Relationship	Description
1:N	A <b>one-to-many</b> mandatory relationship is the most common relationship in databases. The primary key table

Relationship	Description
	contains only one record that relates to none, one, or many records in the related table. This relationship is similar to the one between you and one of your parents. You have one mother, but your mother may have several children.
<b>N:1</b>	A <b>many-to-one</b> is opposite of one to many (1:N) relationship.
<b>1:1</b>	In a <b>one-to-one</b> relationship, both tables are limited to one record only on either side of the relationship. Each primary key value relates to a single record, or no record, in the associated table. They are like spouses — you may be married, or not; however, if you are married, both you and your spouse can have only one partner. Most one-to-one relationships are forced by business rules. If you do not have a business rule, you can, in most cases, combine both tables into one table without breaking normalization rules.
<b>0:N</b>	A <b>zero to many</b> optional relationship indicates that a person may have no phone, one phone, or many phones, and that the phone may not be "owned," but can only be owned by a maximum of one person.
<b>N:0</b>	Opposite of a zero to many relationship
<b>0:1</b>	A <b>zero to one</b> relationship might indicate that a person may be a programmer, but a programmer must be a person. It is assumed that the mandatory side of the relationship is the dominant.
<b>1:0</b>	Opposite of a zero to one relationship
<b>N:N</b>	In a many to many relationship each record in both tables can relate to an unlimited number of records (or no records) in the other table. For example, if you have many siblings, your siblings also have many siblings. Many-to-many relationships must have a third table, referred to as an associate or linking table, because relational systems cannot accommodate the relationship directly.
<b>0:0</b>	A <b>zero to zero</b> optional relationship indicates that a person may occupy one parking space, but that a person is not necessary to have a space and a space does not need to have a person.

7. Click **OK** when you are done.

You should see a new relationship line drawn between the two tables on the Editor Graph, and the relationship represented in the navigation pane.

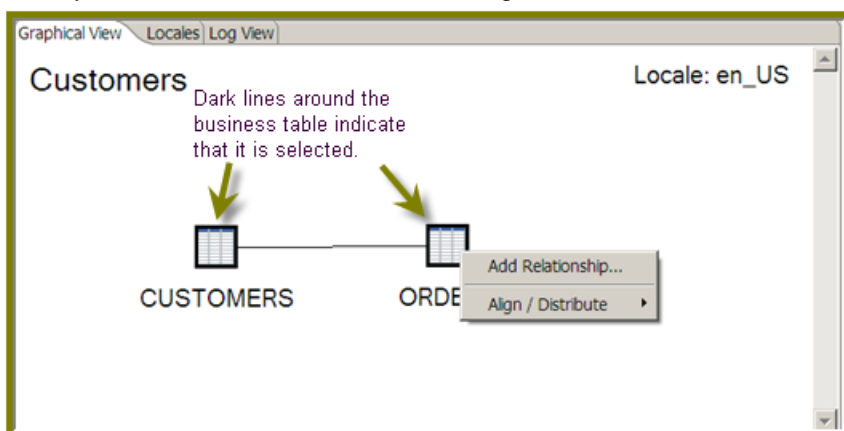


**Important:** Note that complex joins appear in the WHERE clause of the SQL statement, so currently any joining that takes place in the FROM clause of the SQL statement is not supported. An example of a complex join might be `TABLE_A.COL_A=TABLE_B.COL_A AND TABLE_A.COL_B=TABLE_B.COL_B`. This represents a join of two tables based on two key columns versus a single join column. Also note, the complex join expression provided must use the names of the physical tables and physical columns, not business tables and business column names.

## Creating Relationships Using the Editor Graph

In the Editor Graph, creating a new relationship is somewhat simplified because you select the two business tables on the canvas, and the **Relationship Properties** dialog box is pre-populated with your selections. Before you start, make sure that the model you want to add a relationship to is selected, and that the business tables are displayed in the Editor Graph.

1. Select the two business tables you want to include in the new relationship, either by clicking and dragging a marquee around the tables or by holding the **<SHIFT+CTRL>** keys, then clicking on the tables.
2. Once your business tables are selected, right-click on the selection. Click **Add Relationship...** in the popup menu.



The **Relationship Properties** dialog box appears.

3. Continue following steps 3 through 7 in [Creating Relationships between Business Tables](#) to finish creating the relationship.

## Hadoop Hive-Specific SQL Limitations

There are a few key limitations in Hive that prevent some regular Metadata Editor features from working as intended, and will limit the structure of your SQL queries in Report Designer:

- **Outer joins are not supported.**
- **Column aliases do not work.** This means you are forced to use the column ID nomenclature.
- **Each column can only be used once in a SELECT clause.** Duplicate columns in SELECT statements will cause errors.
- **ORDER BY can only work with column IDs, not table names.** When specifying an ORDER BY statement, you can only list the column ID. This means that if you have two tables with identical column names, there will be a name collision; therefore you should try to form your Hive structure such that column IDs are not shared. Secondly, table names are omitted from the syntax entirely: **ORDER BY [columnID]**; is the Hive-supported way, whereas **ORDER BY [table.columnID]**; is the traditional SQL way.
- **Conditional joins can only use the = conditional unless you use a WHERE clause.** Any non-equal conditional in a FROM statement will force Metadata Editor to use a cartesian join and a WHERE clause conditional to limit it. This is not much of a limitation, but it will seem unusual to experienced Metadata Editor users who are accustomed to working with SQL databases.

## Building a Business View

---

The last step in creating a business domain is building your business view. A **Business View** is a collection of business categories that represents the "view" of your model, typically consumed by your end users. Each model can have one and only one business view. Business views are made up of logically (logically relevant to your organization or end-users) organized business categories and business columns.

A business category is like a bucket where you group and re-group your business columns. Business categories can mimic your business table names or be named after your favorite rock stars. Categories do not have metadata associated with them, have no tie back to any business table (although the Editor Graph gives you the impression this relationship exists; don't be fooled), and have the simple purpose of allowing you to bucket the business columns in your model as intuitively as possibly for your data consumers.

Building a business view consists of creating your categories, then moving your business columns from the business tables into the categories. You can move columns from different business tables into the same category, and even duplicate the same business column into two different categories.

The Editor Graph represents the business tables portion of the business model only, so you'll use the navigation pane and the Category Editor to create a business view.

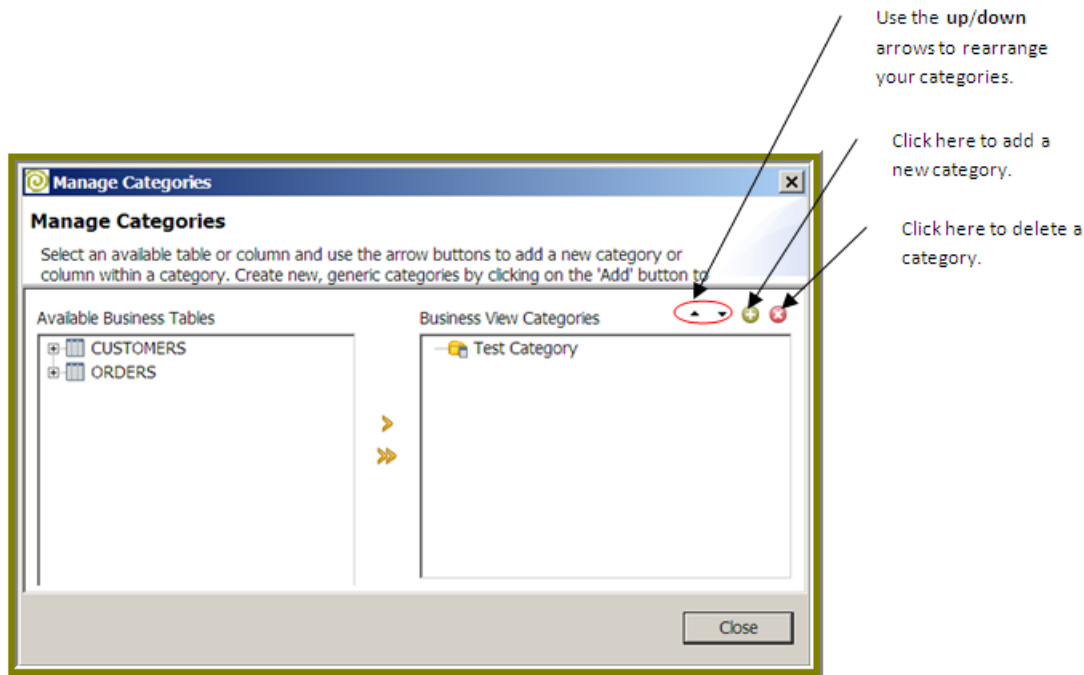
To create a new category using the navigation pane, first make sure that the model you want to add this category to is selected, and the **Business View** node is visible.

1. In the navigation pane, right-click **Business View**.
2. Select **New Category**.  
The **Business Category Properties** dialog box appears.
3. Type a name for the category in the **ID** field.
4. Click **OK**.

## Building a Business View Using Manage Categories

Follow the instructions below to build a business view using the **Manage Categories** dialog box. Before you begin, make sure that the model whose business view you want to create is selected, and the Business View node is visible in the navigation pane.

1. Double-click on Business View in the navigation pane. Alternatively, go to **Tools -> Manage Categories**.  
The Manage Categories dialog box appears.
2. Create categories directly from your business tables by moving columns into categories, add new categories, and remove categories and columns.



3. Click **Close** when you are done.

### Create a New Category from a Business Table

The navigation tree in the Manage Categories dialog box lists all of the business tables in your view. Follow the instructions below to create a new category from a business table in the Manage Categories dialog box:

1. First select the table in the navigation tree on the left.
2. Click > (Add).

Notice that your category has the same name as your business table, and all of your business columns were added to the new category by default.

3. Click **Close** if you are done.

### Moving Columns into Categories

Follow the instructions below to move columns into categories in the Manage Categories dialog box:

1. Select the table from the list of available business tables. To view the columns under a business table, click the plus sign (+) to the left of each table name.
2. Select the destination category from the list on the right.
3. Click > (Add).
4. Click **Close** if you are done.

### Adding a New Category


Follow the instructions below to add a new category in the Manage Categories dialog box.

1. Click + (Add) in the upper right corner of the Business View Categories list. The Business Category Properties dialog box appears.
2. Enter an ID (name) for your category, and click **OK**. The new category name appears in the **Manage Categories** dialog box.
3. Click **Close** if you are done.

### Remove Categories and Columns

Follow the instructions below to remove categories and/or columns in the Manage Categories dialog box:

1. Select the items you want to remove in the **Business View Categories** list on the right.

2. Click  (Delete).
3. Click **Close** when you are done.

# Enriching Your Data — Understanding Metadata Concepts and Properties

---

Presumably the end goal when creating a metadata domain is to have a model that is available to end-users and applications that allows them to consistently and richly enhance the data with which they are working. Whether that be through formatting, security metadata, or some other custom metadata, it all should result in a more efficient, consistent, clearer and smarter view of the data.

So far, you have walked through modeling your database in a manner that is more intuitive than its physical representation. The second half of the metadata work is defining the metadata. The Pentaho metadata paradigm uses the term **concept** to mean a collection of metadata properties that can be applied to a given business object (business table or column, for example).

A "property" has an identifier (a key into a map actually) and a value. This collection of properties is the map of attributes that you want to apply to a particular business object, such as a business column or business table. The concept IS the metadata you define for your business objects. Each business object (physical table, business table, columns, and so on) has its own concept whose properties override all other InheritedConcept or ParentConcept concepts.

Concepts can also be defined independent of any business object and can be structured in an inheritance hierarchy for better organization and management of your metadata. These independent concepts can then be applied to one or more business objects as a ParentConcept.

Refer to [Introducing the Pentaho Metadata Editor](#) for details regarding concepts and how they are used.

## Exceptions to the Rules

---

Like all exceptions, the exceptions to the metadata inheritance logic is complex, but well-grounded. There are two sets of properties that change how the typical inheritance is accomplished — **required properties** and the **default concept**.

## Required Properties

---

All physical and some business objects in the metadata model have a set of required properties. These properties are set automatically on creation of the object, and are not removable (although you can change their value). The purpose for required properties is disallow users from getting into a predicament where they have removed a property that is integral to the SQL generation process. For example, if a physical table did not have a Target Table property set, the SQL generator will cause errors, because it cannot access the correct physical table to query.

Using the previous example, when you set a parent concept on the physical table, and the parent concept has a Target Table property, the physical table does not recognize the parent concept's value for Target Table. This is because the physical table already has a Target Table property as part of its self concept, and the self concept always overrides the parent concept. And since you cannot remove a required property, the parent concept's Target Table will never be recognized at the physical level.

So how do you override a physical object's required property? You set a parent concept at the business model or business view level. Since the inherited properties are override-able, the parent concept at the business level wins.

All physical metadata objects, the business categories and the business model have required properties. You can see what's required for each by referring to [Required Properties per Business Object](#) on page 54.

## Applying Concepts and Properties in the Pentaho Metadata Editor

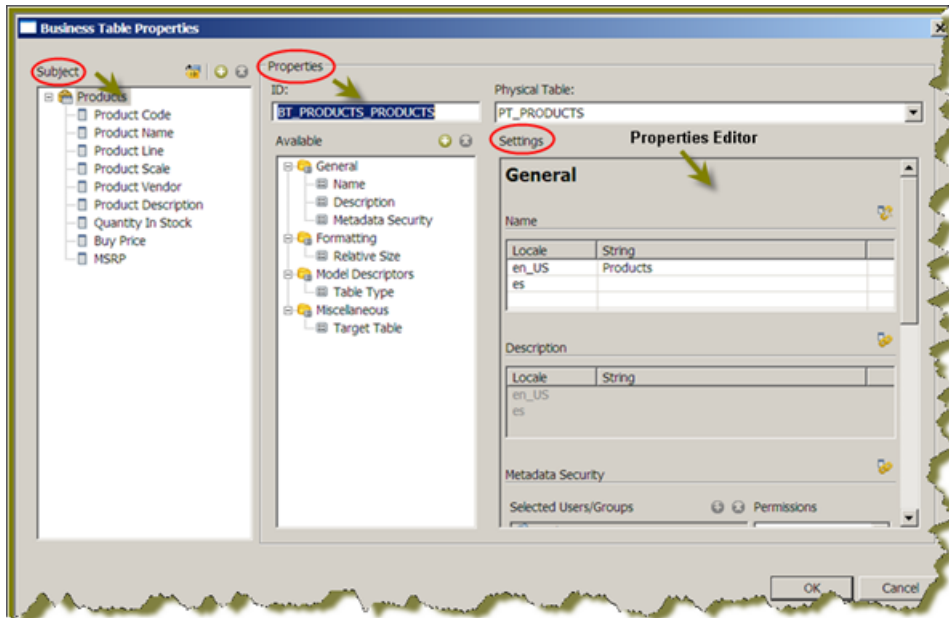
---

As explained previously, there are three levels of concepts for the business objects in the business model, which provides great power and flexibility in how you define your metadata. Since one of the levels is an inherited level, there are really only two concept application processes you need to learn to work with — applying **self concepts**, and applying **parent concepts**.





## Applying Self Concepts: Managing Properties on the Model

You may have noticed when you were creating your tables and columns, that the properties dialog box for every business object has two additional lists to the right of the **Subject** list. These lists are related to **concept editing**. The two additional lists are the **Properties** list, and the **Property Editor (Settings)**, and they operate in the same way regardless of the object to which you are applying metadata.



## The Properties List

The middle panel in the **Properties** dialog box is the **Properties** list. This is the list of currently applied metadata properties for the object that is selected in the **Subject** list. The properties are sorted into their appropriate categories for organization. At the top of the Properties list, there are two icons. The  (add icon) is for adding new metadata properties to this business table's self concept; the  (remove icon) is for removing properties from the self concept. Note that the remove button is disabled until a property is selected in the Property list.

### Helpful Icon Color Coding

There are four ways a property can show up in the Property List:

1. The property is a default property for the associated business object, and thus cannot be removed.
2. The property was inherited from the business object's physical ancestor.
3. The property has been set as part of the self concept.
4. The property has been set as part of the object's parent concept.

It may be difficult to set properties without understanding the override hierarchy, so to help make the hierarchy a bit more apparent in the Properties list, the color code alerts you to what concept level is in play.

- **Yellow Icon** = This property is inherited from the object's physical ancestor.
- **Blue Icon** = Pentaho set the property on the object's own concept. This could be overriding an inherited property.
- **Orange Icon** = This property is set as a result of the parent concept applied to the object.
- **Purple Icon** = This is a special icon reserved for security properties.


## The Properties Editor (Settings)

The **Properties Editor (Settings)** shows the property name and the associated value for that property in a scrolling list. The properties are displayed in the Properties Editor in the same order as they appear in the Properties list. If you click



on a property in the Property List, the Property Editor scrolls to locate the property in the editor. In the Properties Editor, you can set or modify the values of the properties applied to the business object.


## Override

There are icons in the upper right corner next to any property whose value can be overridden. If you want to modify a property by overriding its inherited value, you must first click  (the override button). To cancel an override, click the override button again.

## Adding New Properties

---


Before you add a new property to a business object, make sure that the object you want to apply the new property to is selected in the **Subject** list.

1. In the Properties dialog box, click  (Add) next to **Available**.  
The **Add New Property** dialog box appears. You are prompted with a list of *property choices* to apply.
2. Select a property.
3. Click **OK**.  
The property is now available for you to modify in the Property Editor.

## Removing a Property

---

Follow the instructions below to remove a property:

1. In the Properties list, select the property you want to remove.
2. Click  (Remove)



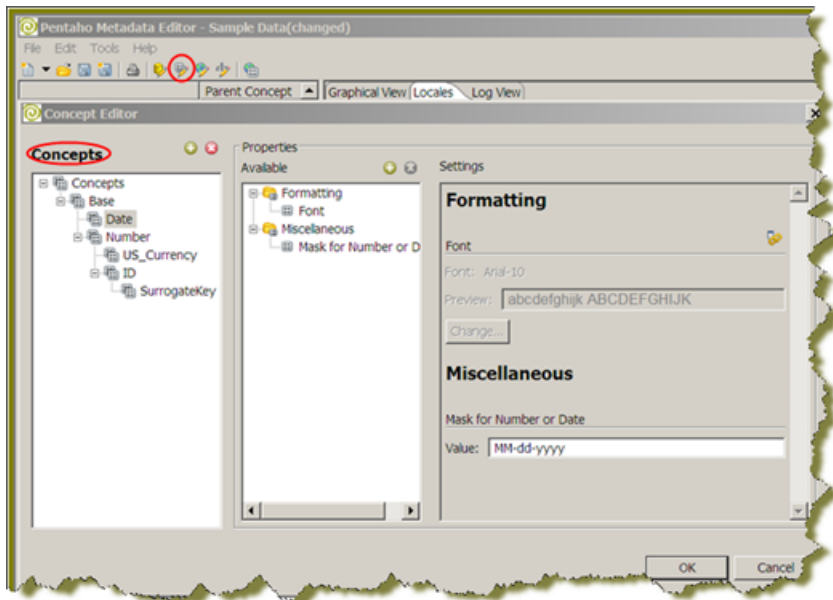
**Note:** If you want to remove inherited or parent concept properties, you must edit the inherited business object or the parent concept.

## The Concept Editor

---

Parent concepts are independent hierarchies of concepts that can be assigned to one or more business objects through the navigation tree. Before you can assign a parent concept you must first learn how to create one.

The **Concept** editor allows you to build concepts to be used as parent concepts. The ability to isolate the concepts, name them, then associate the named concept with one or more business objects gives you flexibility and good concept management. Another feature of the Concept editor is that you can define concepts that build upon other concepts. By nesting concepts this way, you minimize the number of properties you must repeatedly define and create a good inheritance hierarchy.







## The Default Concept

The default Concept is the default set of properties assigned to objects, mainly used for formatting. For example, the mask for all currencies could be set to `$$,##0.00;($,##0.00)` by default. Every model contains a default concept called **Base**. This concept is applied as the parent concept to all physical columns created under a connection. The default concept includes metadata properties that are commonly defined and sets a base line set of metadata for all physical columns in the metadata model. The purpose of the default concept is to provide legitimate values for common metadata properties that may not be set in other places in your hierarchy. You can update, add, and delete metadata properties on the default concept. You can also remove the default concept as the parent concept for any and all physical columns. You cannot, however, delete the default concept from the list of concepts in the Concept editor.

## Building Concepts

The set of instructions below is an example of how to build a concept.

1. On the main page of the Pentaho Metadata Editor toolbar, click  (Concept Editor).
2. With the **Base** concept selected in the **Concepts** list, click  (Add nested concept. Name this concept, **Number**).
3. Make sure **Number** is selected in the **Concepts** list. Note that the properties from Base has been inherited by Number.
4. Add a **Mask** property to **Number**, with a value of `###,##0.00`.
5. With **Number** selected in the **Concept** list, add another nested concept called **ID**.
6. Select **ID** in the **Concepts** list, find the inherited **Mask** property (under **Settings**).
7. Click  (Override) and enter an override value of **0** (zero) for the mask.

 **Note:** You now have three defined concepts that can be used anywhere in your business model. Each concept serves a different type of business object; **Base** for default, generic columns, **Number** for those columns you know contain numeric financial data, and **ID** for ID columns.

8. Click **OK** to save the concepts.

## Applying Concepts to Business Objects

Now that you have some defined concepts, you can designate some of them as parent concepts in our business model.

1. Select any physical table in the navigation tree. In this sample walkthrough exercises the **Customers** physical table is used.
2. Expand the columns under the table in navigation tree. Apply the **ID** concept to the column **Customernumber**.

3. Right-click (or <CTRL> + click) on the column, and select **Set Parent Concept** from the menu.
4. Choose the **ID** concept and click **OK**.  
You should see the concept in the right hand column of the navigation tree on all objects to which it was applied.

## Removing the Parent Concept

---

To remove a parent concept from an object in the business model, you clear the parent concept. 1. 2. 3. This removes the parent concept and its metadata properties from the object.

1. Select any physical table in the navigation tree.
2. Expand the columns under the table in navigation tree and select the column that contains the concept you want to remove.
3. Right-click (or <CTRL> + Click) on the column and select **Clear Parent Concept** from the menu.  
The parent concept and its metadata properties are removed from the object.

# Adding Security to Metadata Business Objects

Pentaho metadata provides a **Security Information** property that allows you to define table or column level security used by the Pentaho BI server. Before you can use this property, you must define settings in the Pentaho Metadata Editor that allow the Pentaho Metadata Editor to communicate with the BI server. The Pentaho Metadata Editor must be able to retrieve the list of users, roles and access control lists (ACLs) needed.

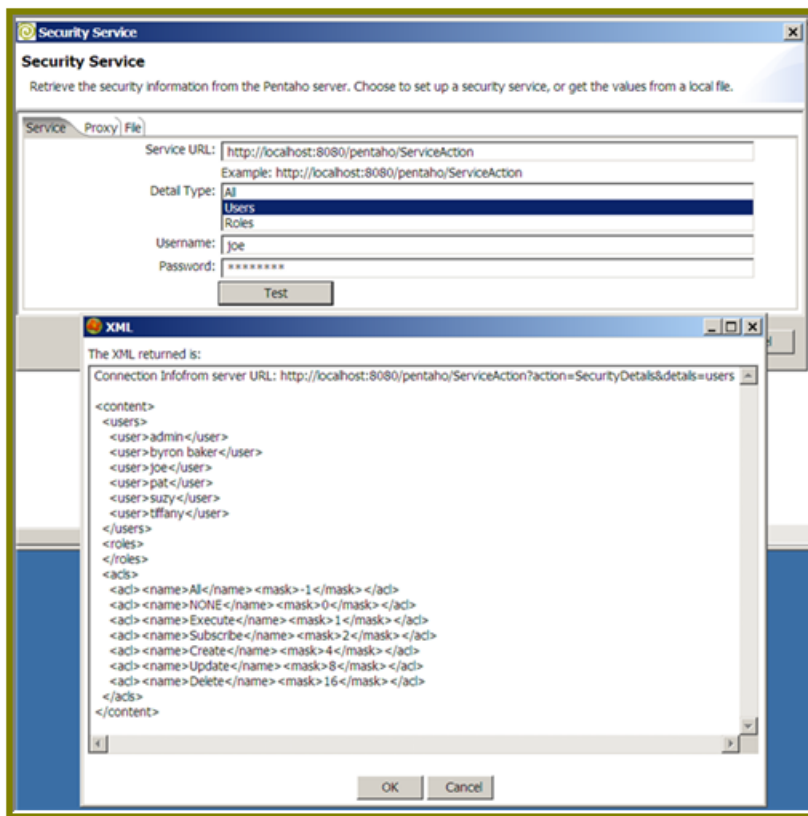
Because the security information used in the business model must be retrieved from a Pentaho BI server, you must have the following information available before you configure security:

- The base URL of the Pentaho BI server and the name of the service that executes security information retrieval. In the demo server, the base URL is **http://localhost:18080/pentaho**, (or **http://localhost:8080/pentaho**), and the default name for the service is **ServiceAction**.
- The logon credentials for the BI server.

## Defining Security Settings

Follow the instructions below to define your security settings:

1. In the Pentaho Metadata Editor main window, go to **Tools -> Security**. The **Security Service** dialog box appears.
2. Type the correct URL in the **Service URL** text box.
3. Select **Users** from the **Detail Type** list.
4. Type **Joe** in the **User Name** text field. Joe is the sample admin user for the Pentaho User Console.
5. Type **password** in the **Password** text field. This password is associated with the sample admin user, Joe.
6. Click **Test**.



If the information you entered is correct, a listing of users and roles appear.

7. Click **OK** to exit.


## Getting Security Settings Offline

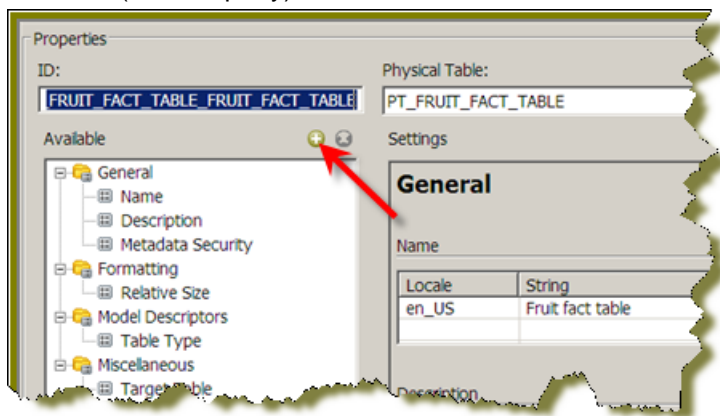
In instances when you want to work on your model and may not have access to the Pentaho BI Server, you can save your security information in a file. The Pentaho Metadata Editor will retrieve your settings from the file instead of accessing the server every time you open your domain.

1. Follow the instructions in .
2. After you click **Test**, all of the XML between the <content></content> tags, including content the tags themselves.
3. Paste the XML code into your favorite text editor, and save the file as **metadata\_security.xml**, in a location of your choice.
4. Click the **File** tab in the **Security Service** dialog box.
5. Browse to the file that you just saved.
6. Click **OK** to exit the dialog box.

## Changing Security Constraints

Follow the instructions below to change security constraints on a specific business table or column.

1. Open the **Properties** dialog box that contains the table or column you want to change.
2. Click  (Add Property).



The **Add New Property** dialog box appears.


3. Select the **Security Information** property and click **OK**.  
The Properties Editor (Settings) scrolls down to the **Metadata Security** section.
4. Add the individual role or user permissions to the business model, table, or column.  
These permissions be enforced in the Pentaho BI Platform after publishing the new metadata model.

# Adding Row Level Security to a Pentaho Metadata Model

Row Level Security allows you to control the results that are returned in a query based on a user's security level. You can specify which rows of data each User Role or User ID is allowed to retrieve from the database, based on some column of data, or combination of columns of data.

## Applying Data Constraints

In the Pentaho Metadata Editor, select the model to which you want to add Row Level Security, right click on the Model, and select **Edit**.



Row Level Security is only in effect at the Model level. Any data constraints defined below the Model Level, such as in a Business Table or Business Column, is ignored and not used. In the Model Properties dialog box, select the **General** -> **Data Constraints** .

### Global Constraint

If you are using the **Global Constraint**, a single MQL Formula is used to define security for all users. In addition to the standard MQL Functions available, there are also two additional functions:

- USER() - returns the name of the current user
- ROLES() - returns a list of roles associated with the current user

The example below defines an MQL formula that allows administrators full access; all other users have no access,


```
IN("Admin"; ROLES())
```

## Role-Based Constraints

If you are using **Role-Based Constraints**, the Metadata engine determines which MQL constraints are appropriate for the current user and applies them to the current query. Constraints may be added for each Role and User in a system. If zero constraints match a user and his or her roles, no data is returned by the MQL query. If more than one constraint applies to a user, the constraints are OR'ed together to determine row visibility.

This example below defines an MQL Formula for three different roles. The Admin role has full row visibility, the Sales and Engineering roles can access data that joins to rows associated with their specific department only.

Role	Constraint
Admin	TRUE()
Sales	[BC_DEPARTMENT]="Sales"
Engineering	[BC_DEPARTMENT]="Engineering"

 **Important:** Row Level Security Constraints are applied at the MQL Layer. The Business Columns referenced in the MQL Security Constraints will be resolved down to SQL Table Columns. The Tables which contain column references included in security constraints will be joined to your query, based on the relationships defined in the Business Model. It is recommended that you do not use outer joined business columns for the purposes of security constraints.

# Importing and Exporting Domains

---

When you save a domain, the domain is stored in a metadata repository. The Pentaho BI server does not use the metadata repository, instead it accesses an XML file exported from the Pentaho Metadata Editor. Exporting your domain is a good way to ensure safe backups of your domains.

## Importing a Domain

---

Follow the instructions below to import a domain. 4. You will be prompted to enter a name for this domain.

1. In the Pentaho Metadata Editor, go to **File -> Import From XMI File**.



**Note:** You may be prompted to save the currently active model if you have any unsaved changes pending. When you import a new domain, that domain becomes the active domain in the Pentaho Metadata Editor.

2. In the file browser, navigate to your domain file, select it, and click **OK**.
3. In the Save Model dialog box, type a name for the domain.



**Important:** If you enter the name of an existing domain, that domain is overwritten by the import.

## Exporting a Domain

---

Follow the instructions below to export a domain:

1. In the Pentaho Metadata Editor, go to **File -> Export to XMI File**.



**Note:** You may be prompted to save the currently active model if you have any unsaved changes pending.

2. Type a file name and select a location to save your file.



**Important:** The default extension for a Pentaho Metadata domain XML file is **.xmi**.

3. Click **Save**.

Once you have entered a name for your export file, the domain is exported to that file. You can inspect the export file using a text editor if you are curious about the underlying XML code.



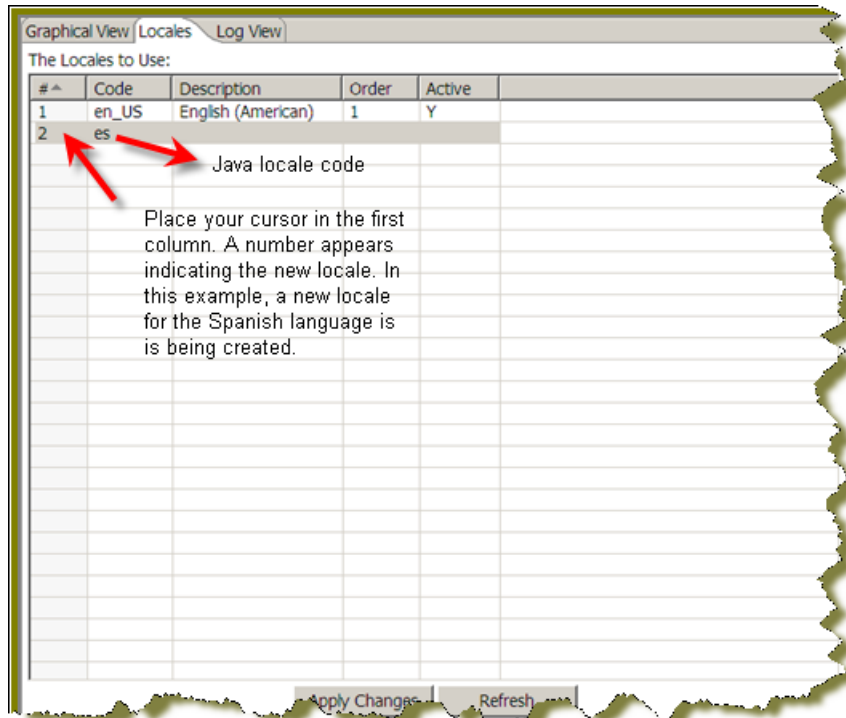
# Configuring and Managing Locales (Internationalization)

Pentaho Metadata Editor supports multiple locale entries for any text or string-based metadata property. But first you must specify the locales that you want to include in your model. To do this, use the Pentaho Metadata Editor **Locales Editor**.

## Setting Up Locales

Follow the instructions below to set up locales:

1. In the Pentaho Metadata Editor, go to **Tools -> Locales Editor**.
2. Place your cursor in the first column of the first empty row in the Locales Editor main table



3. In the next column and type the *Java Locale code* of the language you want to add.
4. Specify the order you want this locale to be used. In the example screen capture, **English (American)** is designated as the first language (Order = 1).
5. Type a **Y** in the last column to activate the locale, or an **N** to de-activate it.
6. Click **Apply Changes** when you are done.

## Domain Backup and Recovery

---

Each domain can be saved to the Common Warehouse Metadata (CWM) repository with any name you like. The Save and Save As options are available from both the Pentaho Metadata Editor File menu and the toolbar. Each time a domain is saved to the repository, a recovery export file of the domain is saved to the file system, under the .pentaho-meta directory. This directory is typically located in the your home directory. The recovery file contains the last successfully saved state of the domain. The files are named **recovery\_[studio:domainname].xmi**.

You can restore a domain to the last saved state by *importing the recovery file* importing the recovery file from the file system. Domain restoration may be necessary because you want to revert to the last known good state of your domain in the event of repository errors or corruption.

# Publishing a Domain to a Pentaho BI Server

---

You must share an XML representation of your domain with a Pentaho solution that your BI server recognizes so that the server can access your metadata domain. The implementation of metadata in the Pentaho BI Server requires that you follow some basic rules:

- There can only be one metadata domain per Pentaho solution.
- The metadata domain is associated with the solution by placing the XML format of the domain in the root directory of the solution.
- The domain XML file must be **metadata.xmi**.

You can accomplish this by [exporting your domain file](#) and copying it manually to the solution directory of your choice. Or, you can use the Pentaho Metadata Editor publish feature if you know the credentials to your server.

## Before you Publish Your Domain

---

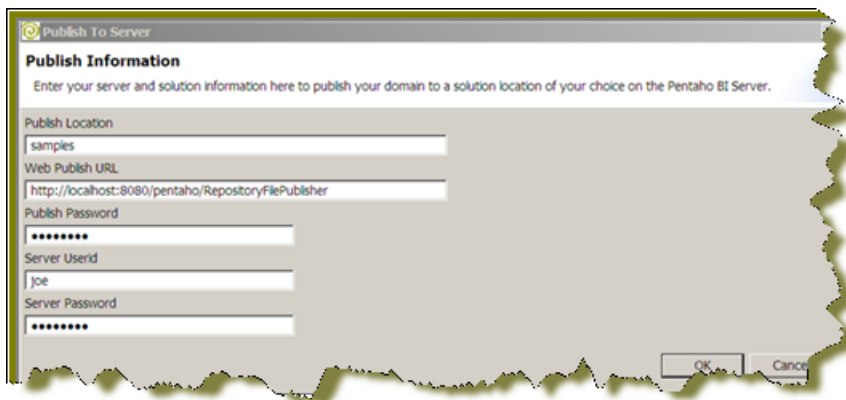
Before you publish a domain make sure you have the following information available:

- Has your Pentaho BI Server's publish password been set in `... \pentaho-solutions\system \publisher_config.xml`? Ask your server administrator if you are not sure. This password must be set so that the server publish utilities can be used.
- Do you have a user name and password to the Pentaho BI Server that provides administrative privileges?
- Do you know the base URL for your Pentaho BI server? This consists of all URL information up to and including the Web application context. For example, the base URL of the pentaho demo server is **http://localhost:18080/pentaho**.
- For publishing, you must append the base URL the to name of the publish service. The default is **RepositoryFilePublisher**. This is the value you must use, unless your server admin has changed the service name.
- You must also know the name of the solution to which you want to publish the domain. For example, the demo server ships with two different solutions, the **Samples** solution and the **Steel Wheels** solution. No slashes are required.

## Publishing a Domain

---

Follow the instructions below to publish your domain to the server.



1. In the Pentaho Metadata Editor go to **File -> Publish to Server**. The **Publish to Server** dialog box appears.
2. Type the **Publish Location**; this is your solution name.
3. Enter the **Web Publish URL**, which is the Pentaho BI Server base URL followed by the publish service name. **RepositoryFilePublisher** is the default.
4. Type the **Publish Password**.
5. Type the **Server Userid** (administrator user name).
6. Type the **Server Password** (administrator password).

7. Click **OK** to save your entries and exit the dialog box.



**Note:** If the metadata.xmi file already exists on the server, you are prompted to overwrite it.

### Troubleshooting

If publishing to the server fails, below are a few things you may want to verify to be sure your information and configuration is correct:

- Make sure that the publish password is set on the server. You cannot use the BI server publish utility if the publish password is not set. The publish password is defined in **publisher\_config.xml** in **pentaho-solutions\system** directory.
- Verify that your server credentials have admin level responsibility.
- Do not use a trailing or leading slashes in the server URL, solution name, or domain file name.

# Sample Metadata Model Walkthrough

---

This series of exercises walks you through the steps of examining a sample metadata model. Before you begin the exercise, you must restore the [current\\_fruit\\_data\\_20081107\\_1424.sql](#) file to a MySQL database and [launch the Pentaho Metadata Editor](#). This file contains a small sample database.

To watch the recorded demo on which this walkthrough is based, go to the [Pentaho Demo Page](#), scroll down to the topic, *The Fundamentals of Pentaho Metadata Editor 3.0 with a bit of Report Designer 3.0*, and click the **Streaming Flash** link. The walkthrough in this document does not include the Report Designer instructions.

## Examining the Default Metadata Model

---

- The Pentaho Metadata Editor must be running.
- The Pentaho BI server must be running.

Follow the instructions below to see an example of the default metadata model that is deployed with the Pentaho BI Server.

1. In the main window of the Pentaho Metadata Editor, go to **File -> Import from XMI File**.
2. Navigate to your installation of the BI Server (for example, C:\Program Files\pentaho\server\biserver-ee\pentaho-solutions\, if you used the installer).

There can be several root pentaho-solution folders and each of them may contain one metadata model repository.

3. Open the **steel-wheels** folder and click **metadata.xmi**.

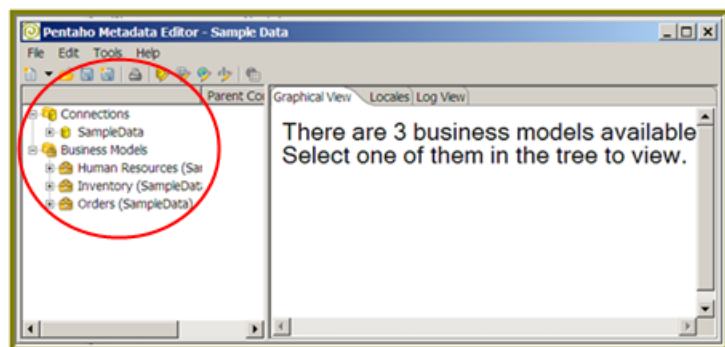
Metadata.xmi is the repository for Pentaho-related metadata and business views.

4. Type **Sample Data** in the **Save Model** dialog box.

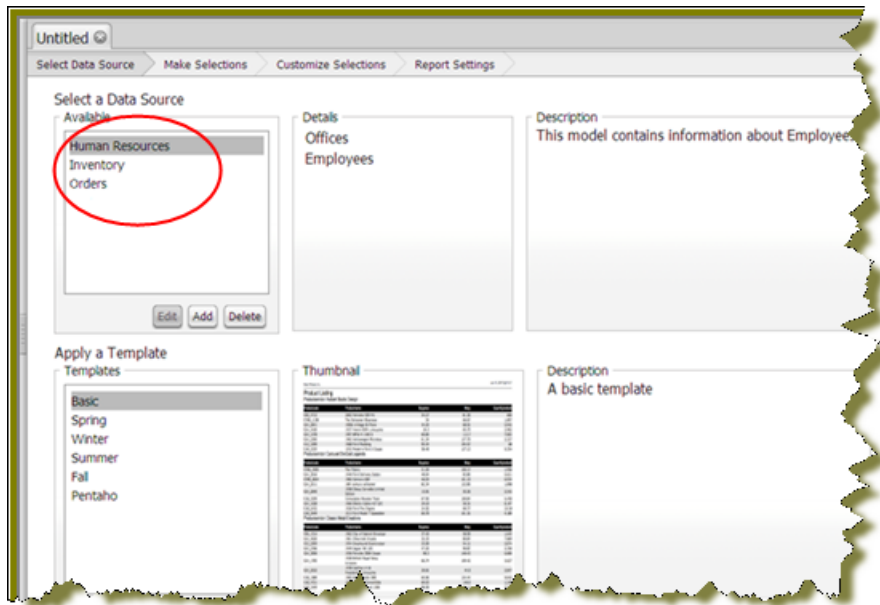


**Note:** If you see the message, "This model already exists...," click **Yes** to continue.

This step processes the file and displays the "structure" of the repository (connections and business models) on the navigation pane on the left.



5. Log into the Pentaho User Console, and click **New Report**.



Notice that the business models (Human Resources, Inventory, and Orders) are displayed. (To make it easier for end-users, the term "Business Model" has been replaced with "Data Source" in the Pentaho User Console in version 3.5).

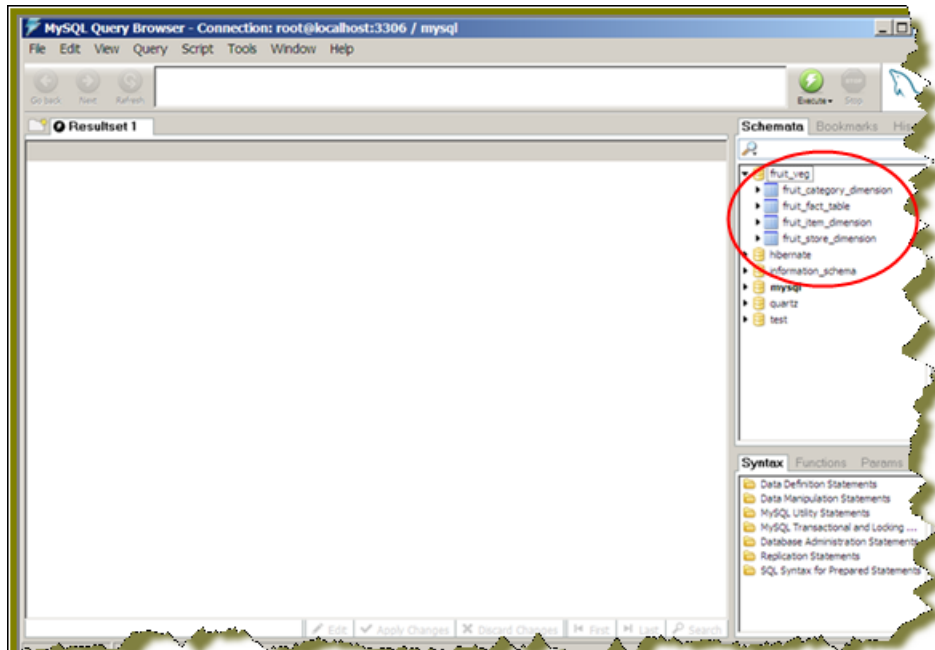
## Creating a New Domain File and Connecting to the Database

MySQL must be running.


Before you can start building your metadata model, you must create a new domain and connect to the MySQL, which contains the *fruit\_veg* database you restored.

1. In the Pentaho Metadata Editor main window, go to **File -> New Domain File**.
2. In the navigation pane, click **Connections -> New Connection**.  
The Database Connection dialog box appears.
3. In the Database Connection dialog box, type **Fruits and Vegetables** in the **Connection Name** text field.

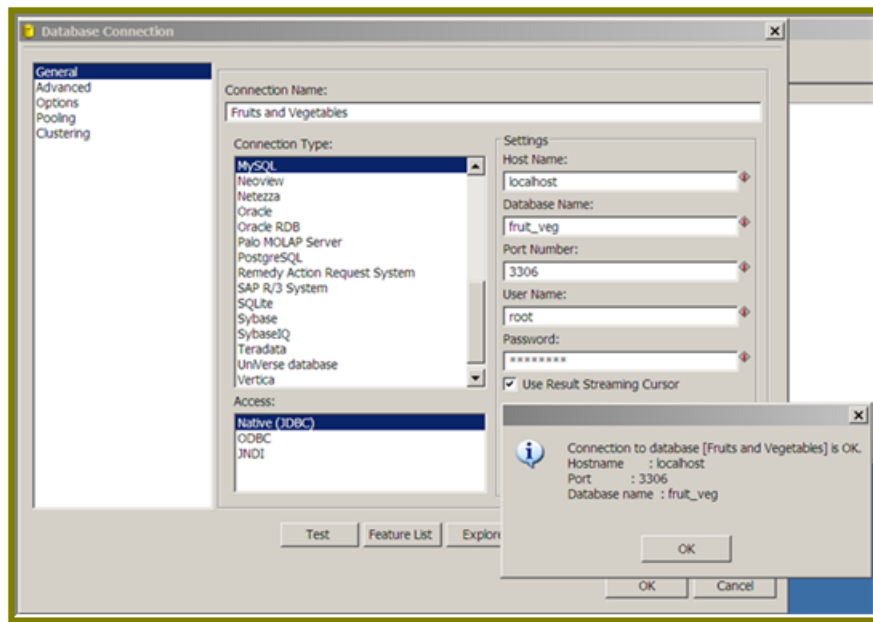
This is the database you restored from the **current\_fruit\_data 20081107 1424.sql** file.



4. Under **Connection Type**, select **MySQL**.
5. Under **Settings**, enter the information as shown in the table below:

 **Note:** It is assumed that you have an installation of the BI Server which includes a MySQL repository on your local device. This ensures that your settings are correct. The default **Port Number** for MySQL is entered automatically.

Setting	Entry
Host Name	localhost
Database Name	fruit_veg (This is MySQL schema name.)
User Name	pentaho_user (You may have to enter <b>root</b> as your user name if you have issues connecting to the database.)
Password	password

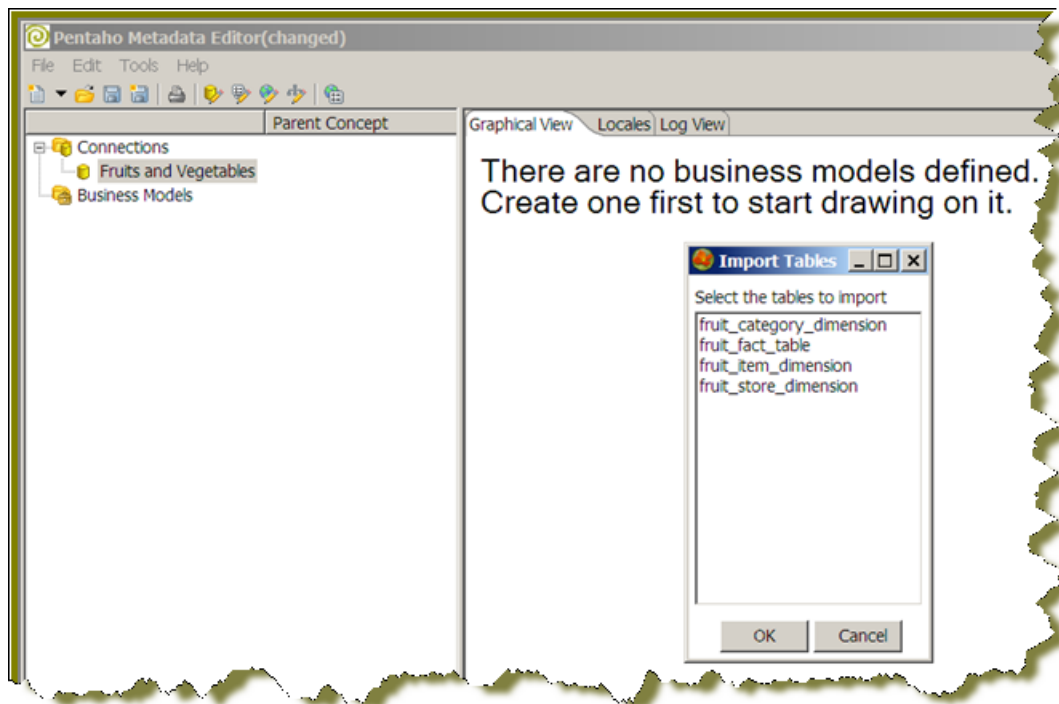


6. Click **Test** to make sure you are connected to the database. A success message appears when you are connected.

## Importing Physical Tables and Columns

After you define your database connections, the next step is to describe the physical database tables and columns that you want to include in your business model.

1. Right-click the **Fruit and Vegetables** node in the navigation pane. Select **Import Tables** from the menu.

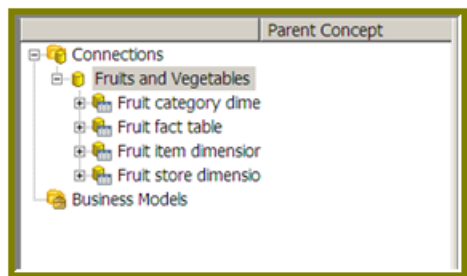


A dialog box that contains all physical tables available in the database appears.

2. Select all the tables.

 **Tip:** Press the <SHIFT> key to select multiple tables.

3. Click **OK** when you are done.



The tables you selected appear in a list in the navigation pane in the left side of the workspace.

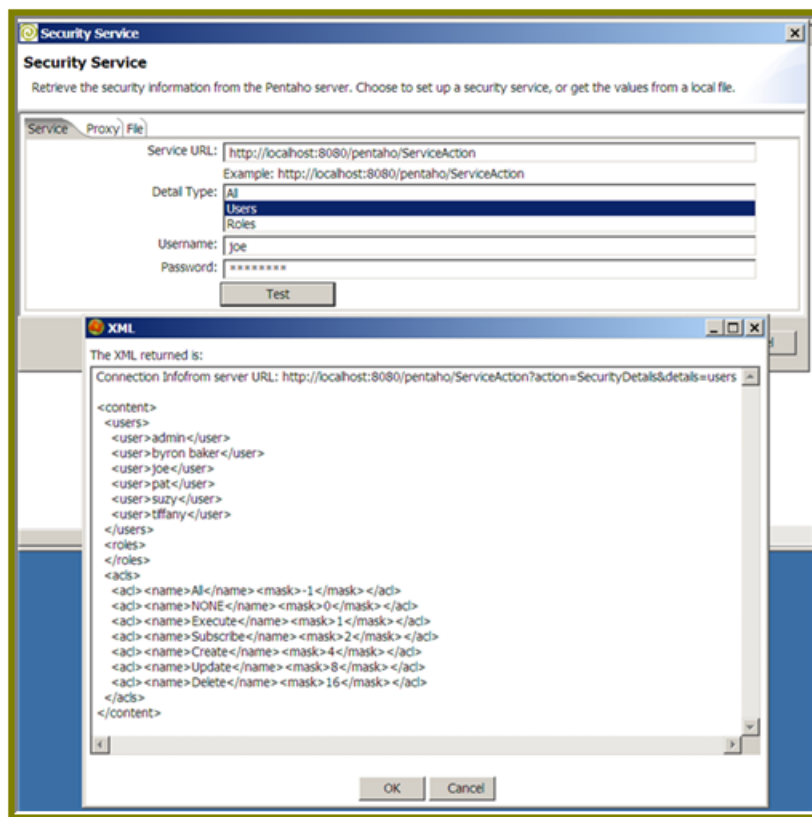
## Setting Security

- The Pentaho Metadata Editor must be running.
- The Pentaho BI server must be running.

Pentaho Metadata Editor provides you with a security service that allows you to define the table or column-level security used by the Pentaho BI Server. Before you can use this service, you must provide the URL to your Pentaho BI Server, so that the Pentaho Metadata Editor can retrieve the required list of users, roles and Access Control Lists (ACLs). By default, the security service URL is set to `http://localhost:8080/pentaho/ServiceAction`; however, if you used the installer to install the Pentaho BI Suite, the service must be set to `http://localhost:18080/pentaho/ServiceAction`. You must make sure that this URL is correct before you proceed.

1. In the Pentaho Metadata Editor main window, go to **Tools -> Security**. The **Security Service** dialog box appears.
2. Type the correct URL in the **Service URL** text box.
3. Select **Users** from the **Detail Type** list.
4. Type **Joe** in the **User Name** text field. Joe is the sample admin user for the Pentaho User Console.
5. Type **password** in the **Password** text field. This password is associated with the sample admin user, Joe.
6. Click **Test**.





If the information you entered is correct, a listing of users and roles appear.

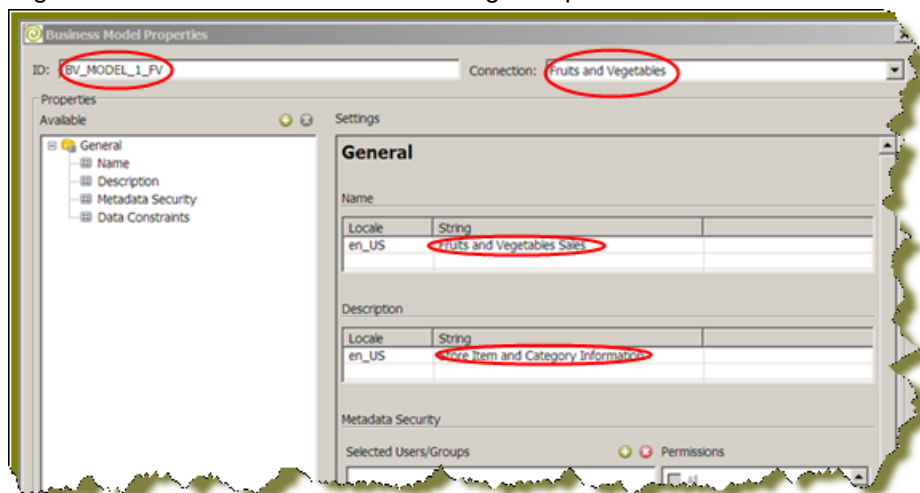
7. Click **OK** to exit.

## Creating Your Business Model

- The Pentaho Metadata Editor must be running.
- The Pentaho User Console must be running.


Follow the instructions below to create your business model:

1. Right-click **Business Models** in the navigation pane and select **New Business Model**.



The **Business Model Properties** dialog box appears as shown above.

2. In the ID text field, type **BV\_MODEL\_1\_FV**.
3. Select **Fruits and Vegetables** from the **Connection** list.
4. Type **Fruits and Vegetables Sales** in the **Name** field under **String**.  
This is the name of the metadata model that is visible to your users.
5. Type **Store Item and Category Information** in the **Description** field under **String**.

6. For the purpose of demonstration, click  (Add Users/Groups) under **Metadata Security**, to examine the **Add Permissions** dialog box.

This step allows you to set metadata security based on users and roles and to select the appropriate assigned permissions.

7. For the purposes of demonstration, scroll down to **Data Constraints** to examine the constraint options.

This step allows you to set global and role-based constraints associated with row-level security.

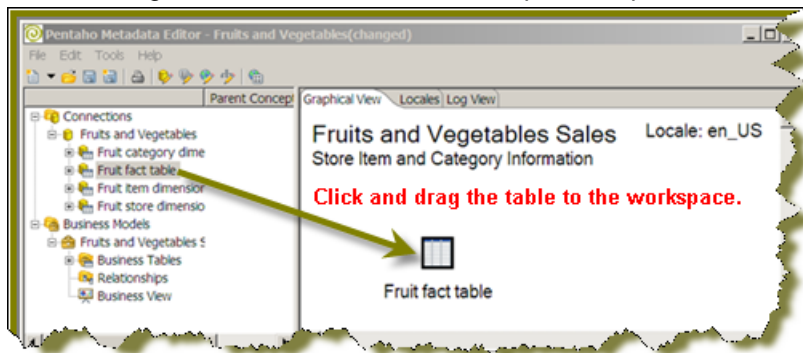
8. Click **Cancel** to exit the **Add Permissions** dialog box.
9. Scroll down to examine the options under **Data Constraints**.
10. Click **OK** to save your changes and exit the dialog box.

## Assigning Business Tables to Physical Tables

- The Pentaho Metadata Editor must be running.
- The Pentaho BI server must be running.

Now that you have your business model defined you must assign the business tables to the associated physical tables.

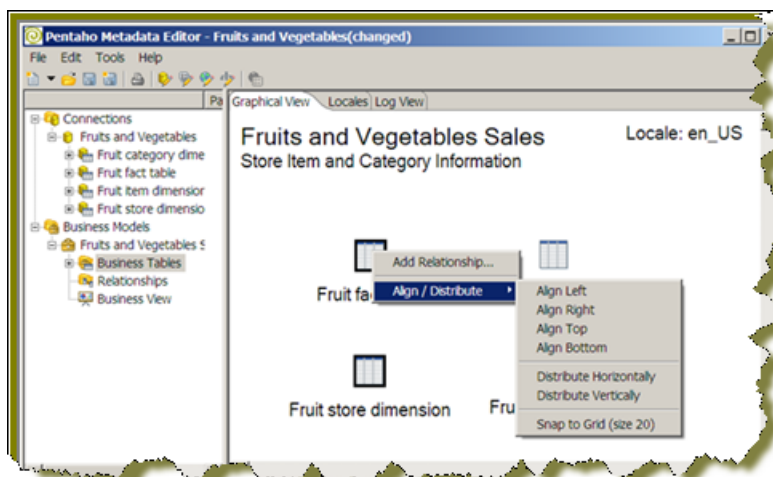
1. In the navigation pane, expand the **Fruits and Vegetables** node under **Connections**.
2. Click and drag the **Fruit fact table** to the Graph workspace.



The **Business Table Properties** dialog box appears. Each time you drag a table to the workspace, you can set additional parameters and properties in the Business Table Properties dialog box.

3. In the **Business Table Properties** dialog box, click **OK** to accept the defaults.
4. Click and drag the dimension tables (**Fruit category dimension**, **Fruit item dimension**, **Fruit store dimension**) to the workspace. Click **OK** to accept the defaults in the Business Properties dialog box for each dimension table.

When a table is selected, a dark border appears around it. To select more than one table, click and drag a marquee around the tables or press the **<SHIFT + CTRL>** keys to make your selections. Right-click and choose **Align/Distribute** to display the alignment options.

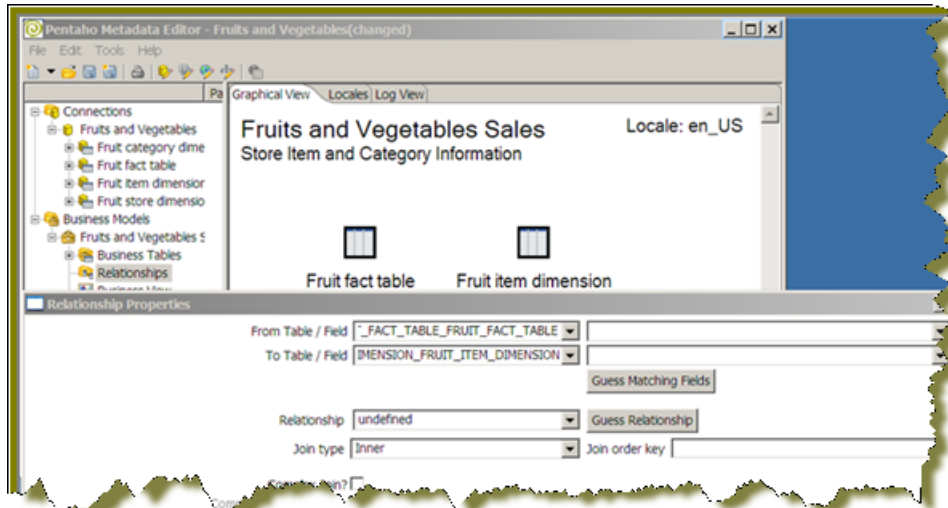


## Defining Relationships between Business Tables


- The Pentaho Metadata Editor must be running.
- The Pentaho BI server must be running.

After you assign your tables, you must define the relationships between them.

1. In the workspace, press <SHIFT + CTRL> keys to select the **Fruit fact table** and **Fruit item dimension** table.
2. Right-click and select **Add Relationship** to display the **Relationship Properties** dialog box.



3. Specify the business columns (from the adjacent lists) from each business table that identify this relationship. The **Fruit fact table** will connect to the **Fruit item dimension** table through the **ITEM\_ID**. The **Fruit item dimension** table will also connect through the **ITEM\_ID**.



### Note:

When the business column names are similar, click **Guess Matching Fields** and let Pentaho Metadata Editor attempt to determine the fields for you.


4. Select **N:1** (many-to-one) from the **Relationship** list and accept the **Join type** default (**Inner** join).



### Note:

Click **Guess Relationship** if you want Pentaho Metadata Editor to provide an initial guess at what type of relationship exists between the tables.

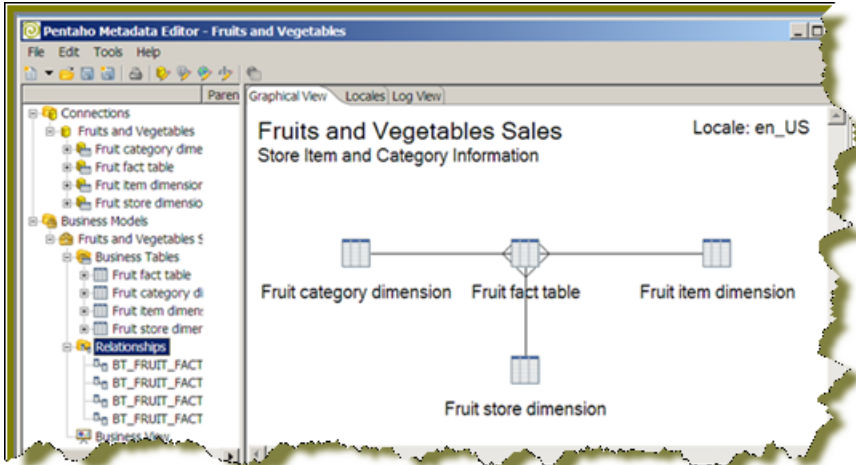
5. Click **OK** to save your changes and exit Relationship Properties dialog box. In the workspace, notice the relationship line connecting the two tables.
6. Repeat steps 1 through 5 to connect the **Fruit fact table** to the **Fruit store dimension** table. The **Fruit fact table** will connect to the **Fruit store dimension** table through the **STORE\_ID**. The **Fruit store dimension** table will also connect through the **STORE\_ID**. The relationship is will be **N:1** (one-to-many), and the Join type will be **Inner**.



7. Repeat steps 1 through 5 to connect the **Fruit fact table** to the **Fruit category dimension** table. The **Fruit fact table** will connect to the **Fruit store dimension** table through the **CAT\_ID** (category ID). The **Fruit store dimension**

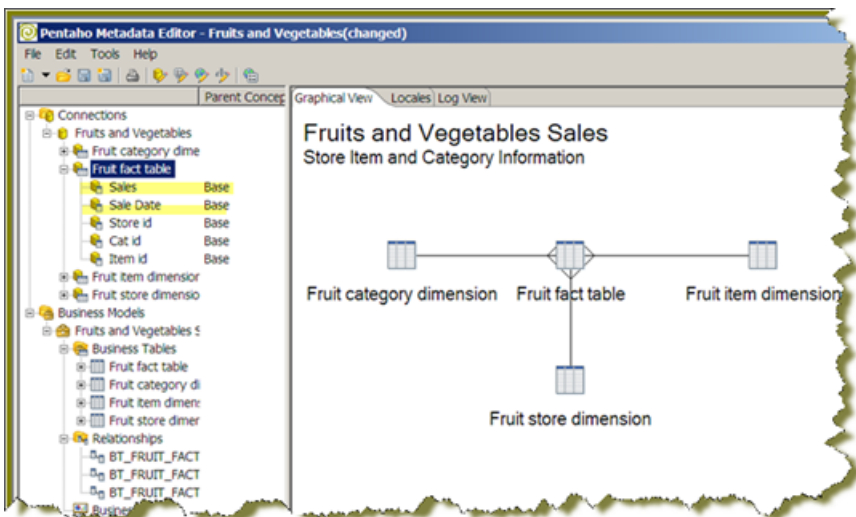
table will also connect through the **CAT\_ID**. The relationship is will be **N:1** (one-to-many), and the Join type will be **Inner**.

The relationships are defined.



## Creating Derived Columns (Part 1)

Now that you have defined your relationships, you are ready to create your business view; however, for the purpose of demonstration, you are going to create a business view that contains an aggregate of **Sales** and derived time information from the **Sales Date**, such as year, quarter, and month.



**Note:** Derived columns are sometimes referred to as "formula-based columns."

1. Under **Connections**, double-click on the **Fruit fact table** to open the **Physical Table Properties** dialog box.
2. Click **+** (Add New Column) to display the **New Column** dialog box.
3. Type **PC\_Total\_sales** in the text box.

This is the identifier for the column that will be an aggregate of total sales.


**Total sales** is added to the list of columns in the Fruit fact table.


4. In the **Physical Table Properties** dialog box, double-click the **Total Sales** column you just created and scroll down to **Aggregation Type** under **Model Descriptors**.
5. Select **Sum** from the **Aggregation Type** list.
6. Select **Numeric** from the **Data Type** list.
7. Type **-1** in the **Length** and **Position** text boxes.

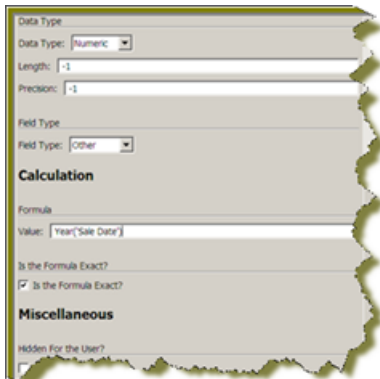
8. Scroll down to **Calculation** and type **Sales** in the **Value** text box.
9. Click **OK** to exit the dialog box.  
The **Total Sales** column appears in the navigation pane, under **Connections -> Fruit Fact table**.

## Creating Derived Columns (Part 2)

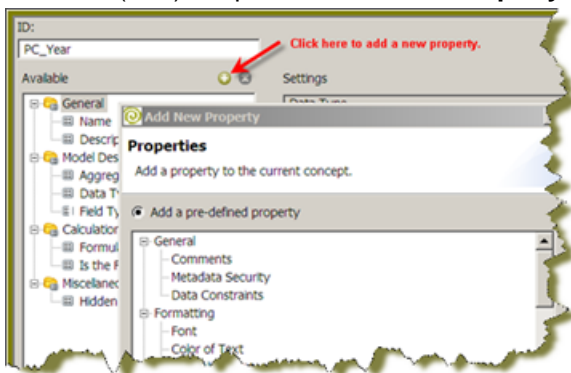
In this part of the exercise, you will create derived columns year, quarter, and month from **Sales Date**.

1. In the Pentaho Metadata Editor main window, right-click **Fruit Fact table** (under **Connections**) (in the navigation pane) and select **Edit** to display the **Physical Tables Properties** dialog box.
2. Click  (Add New Column) to display the **New Column** dialog box and type **PC\_Year** in the text box.  
A **Year** column is added under the Fruit Fact table.
3. Repeat Steps 1 and 2 to create the **PC\_Quarter** and **PC\_Month** columns.
4. Under **Fruit Fact table**, click **Year** to select it.
5. Scroll down to **Data Type** under **Model Descriptors** and select **Numeric** from the list.
6. Type **-1** in the **Length** and **Position** text boxes.
7. Type `Year('Sale Date')` in the **Value** test box under **Formula**.

 **Note:** This entry is a dialect-specific formula for MySQL. You are using the Year function.




8. Enable **Is the Formula Exact?**.
9. Click  (Add) to open the **Add New Property** dialog box and select **Text Alignment**.



The Text Alignment property appears in the Settings pane under **Formatting**.

10. Select **Left** from the **Alignment Type** list.

Numeric values default to right alignment.

11. Click  (Add) to open the **Add New Property** dialog box and select **Mask for Number or Date**.  
Format masks (for example, 'MMDDYYYY') are assigned to date and numeric fields; they are used to define how data returned for that field is displayed to users. Date fields must have format masks.

12. In the Settings pane, scroll down to **Miscellaneous** and type ##### in the **Value** text field under **Mask for Number or Date**.

This is your four-digit year.

### 13. Physical Table Properties


14. Now you must repeat Steps 4 through 11 for **Quarter** and **Month**. Refer to the tables below to complete this exercise.

Name	Data Type	Lgth. & Prec.	Formula	Align. Type	Mask
Quarter	Numeric	-1	Quarter('Sale Date')	Left	#
Month	Numeric	-1	Month('Sale Date')	Left	#

15. Click **OK** to exit the **Physical Table Properties** dialog box.

## Creating the Business Model

Now that you have created the derived columns, you are ready to add them to the business model.

1. In the navigation pane, under **Business Models**, expand **Business Tables** and double-click **Fruit fact table** to open the **Business Table Properties** dialog box.
2. Click  (Add) to open the **Add New Column** dialog box.
3. Press **<CTRL+Click>** to select **Total Sales**, **Year**, **Quarter**, and **Month** columns; click **OK** to exit the dialog box.


## Creating the Business View

Your business tables, relationships, and derived columns have been established. You are ready to create your business view.

1. In the navigation pane, right-click **Business View** to open the **Business Category Properties** dialog box.
2. In the **ID** text field, type **FV\_location\_id**.

This category will contain the Store ID.


3. Under **Name**, type **Location** in the **String** text field.
4. Under **Description**, type **Location** in the **String** text field.

 **Note:** In this dialog box you can also define options for metadata security. For the purpose of the exercise, you are accepting the default values.

5. Click **OK** to exit the dialog box.
6. Repeat steps 1 through 5 for the following categories: **FV\_Metrics\_id**, **FV\_Time**, and **FV\_items\_category\_id**.

Category	Name and Description
<b>FV_Metrics_id</b>	Measures
<b>FV_Time</b>	Time
<b>FV_items_category_id</b>	Items and Categories

Your categories are added to the business view.

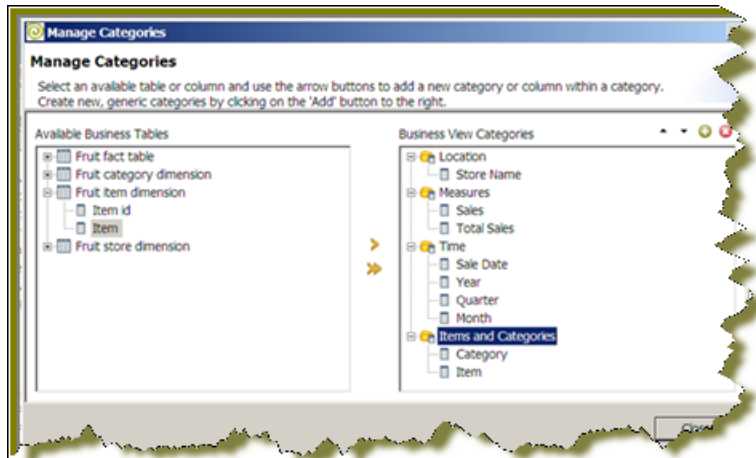


## Managing Categories

Follow the instructions below to select the columns and associate them your business categories.




1. Double-click on Business View in the navigation pane. Alternatively, right-click **Business Categories** and select **Manage Categories** or go to **Tools -> Manage Categories**.  
The **Manage Categories** dialog box appears.
2. Expand the **Fruit fact table** to display its associated columns.
3. Under **Business View Categories**, click **Measures** to select it.
4. Under **Available Business Tables**, click **Sales** to select it.
5. Click ➤ (Add) to move **Sales** under the **Measures** category.
6. Move **Sales Date, Year, Quarter** and **Month** under the **Time** category.
7. Expand the **Fruit store dimension** and move **Store Name** under the **Location** category.
8. Expand the **Fruit category dimension** and move **Category** under **Items and Categories**.
9. Expand the **Fruit item dimension** and move **Item** under **Items and Categories**.

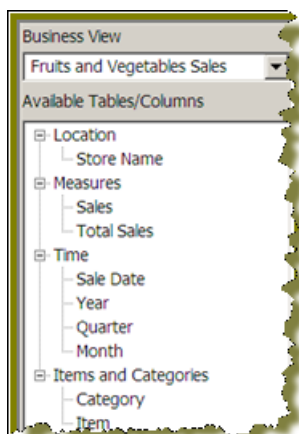


10. Click **Close** to exit the dialog box.
11. In the navigation pane, expand the **Business View** node to see the **Business View Categories**.

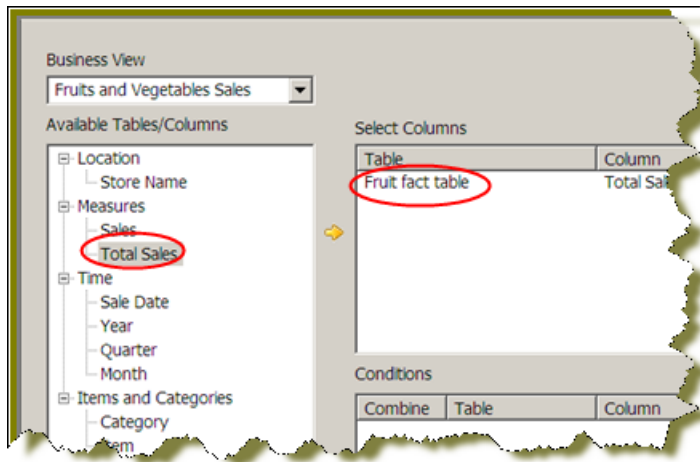
## Testing Your Business Model with MQL Query Builder



The MQL Query Builder allows you to select the categories and the fields that you populated in those categories thus giving you the ability to test your business model before you publish it to the BI Server.


1. In the Pentaho Metadata Editor toolbar, click  (Query Editor).  
The MQL Query Builder opens.
2. In the navigation pane of the MQL Query Builder, expand your categories.  
If you followed the instructions in [Managing Categories](#) on page 46 precisely, you should see results like the sample below:



3. Select **Total Sales** and click  (Add to Details) to move it under **Select Columns**.






4. Click  in the MQL Query Builder toolbar. The **Examine Preview Data** window opens. Notice that the value for **Total Sales** is listed.
5. Click **Close** to exit the **Examine Preview Data** window.  
Let's suppose you don't want to look at an aggregate value like Total Sales. Instead, you want to examine individual sales data by category, for example.
6. In the MQL Query Builder, select Total Sales (under Select Columns) and click  (Remove).






Total Sales is removed.

7. In the navigation page of the MQL Query Builder, move **Category** (Items and Categories) and **Sales** (Measures) and under **Select Columns**.

 **Note:** To determine how your data will be ordered in the preview window, select the items under Select Columns and use the yellow up/down arrows to rearrange the order as needed.

8. Click  to see the results in the **Examine Preview Data** window. Now, you can see sales results by the fruit and vegetable categories.
9. Let's suppose you want to look at aggregate sales by category. First, remove **Sales** from the items under Select Columns; (see Step 6) and place **Total Sales** under **Select Columns**.
10. Click  to see the aggregate values for Fruits and Vegetables.

 **Note:** You can examine the MQL and SQL queries associated with your choices by clicking the appropriate MQL and SQL icons,  , in the toolbar. MQL (Metadata Query Language) was developed by Pentaho to query the metadata layer. The query is built dynamically as you make your choices.

After you have tested your model and are happy with the results, you are ready to publish your model.

## Publishing Your Business Model to the Pentaho BI Server

Before you start the publish to server process, make sure that the BI Server is running. Refer to [Before you Publish Your Domain](#) on page 35.

1. Log into the Pentaho User Console using the credentials, **Joe/password**.
2. Right-click anywhere in the **Browse** pane of the Pentaho User Console and select **New Folder**. The **New Folder** dialog box appears.
3. Name the new solution folder, **demo**.
4. In the Pentaho Metadata Editor, go to **File -> Publish to Server**. You may get a message asking you to save your model.
5. In the Publish Information dialog box, type **demo** in the **Publish Location** field. This is the root solution folder to which the model will be published.

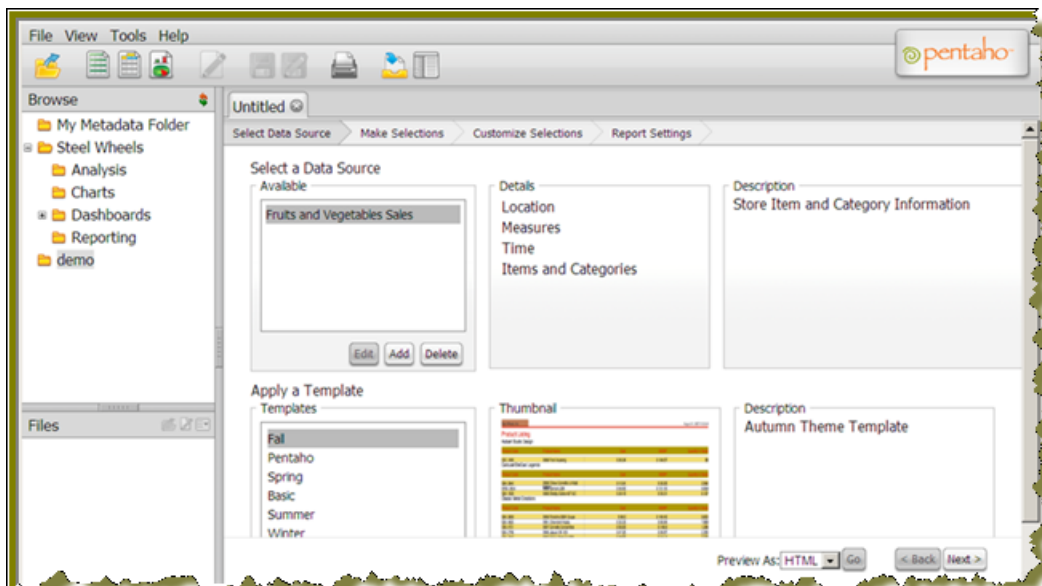



6. Type the **Web Publish URL**, which is the Pentaho BI Server base URL followed by the publish service name. **RepositoryFilePublisher** is the default.
7. Type the **Publish Password**. The default publish password is **password**.
8. Type the **Server User ID**; the default is **Joe**.
9. Type the **Server Password**; the default is **password**.
10. Click **OK** to exit the dialog box.  
A success message appears. You can verify that the metadata.xml file associated with your model was saved to the demo folder by navigating to `... \pentaho-solutions\demo` folder.
11. Log back into the Pentaho User Console and go to **Tools -> Refresh -> Reporting Metadata**.  
A success message appears.

You can now use the metadata model to create reports using the Report Designer, Design Studio, or Ad Hoc Query.

## Examining the Model in the Pentaho User Console

After you publish your model, it is available to users of the Pentaho User Console. The instructions that follow allow you to become familiar with how users take the data in a model to create an ad hoc report. This model is also available to them when they create a report in Report Designer.

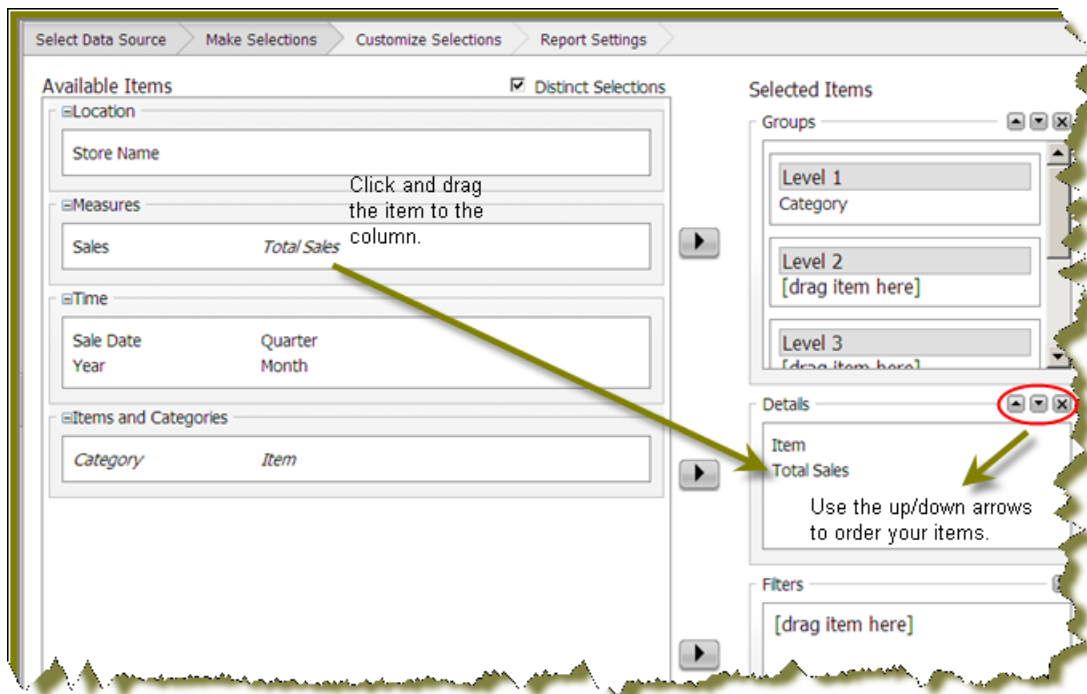



1. Log on to the Pentaho User Console using the default credentials, **joe/password**.
2. In the menu, go to **Tools -> Refresh -> Reporting Metadata**.  
This step ensures that the new model is loaded into the Pentaho User Console.
3. In the toolbar or quick launch bar, click  (New Report). Alternatively, in the menu go to **File -> New -> New Report**.

The Ad Hoc Reports interface opens. Ad Hoc Reports allow users to select the published business models and to view descriptive information associated with the model. It also allows them to select a template for the report they will create using the data in the model. Users can output their reports in HTML, PDF, Excel and CSV format.

In the sample screenshot, the Fruits and Vegetable model is already selected. Under **Details** (Location, Measures, Time, and so on) is the Business View you created. Users can see descriptive information can select a template.

4. Click **Next** (Make Selections) to start selecting items that are in the categories you specified in the Pentaho Metadata Editor.  
You are going to create a simple sales report.
5. Under **Available Items**, click **Total Sales** and drag it to the **Details** column; click **Category** and drag it to the **Groups** column. Move **Item** to the **Details** column.



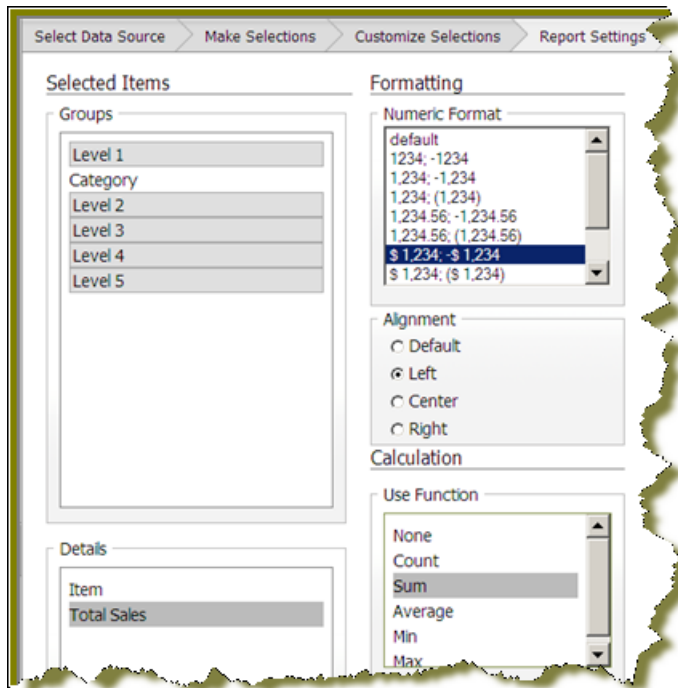
 **Tip:** You can use the up/down arrows in the columns boxes to adjust or reorder the position of your items so that they display properly in your report. You can also delete unwanted items.

6. Click **Go** to preview the report.

Your report appears. Click out of Preview mode to continue making adjustments to you report.

Category: Fruits	
Item	Total Sales
Apples	5,000
Bananas	7,150
Oranges	5,200
Category: Vegetables	
Item	Total Sales
Carrots	3,530
Lettuce	4,045
Peas	3,675

7. Click **Customize Selections** to apply additional formatting to the items you selected. In this example, click **Total Sales** and select the dollar sign for the **Numeric Format**, **Left** for the **Alignment**, and **Sum** for the **Use Function**. Click **Go** to preview the report.



The formatting is applied and you should see both the subtotal and grand total associated with your fruits and vegetables sales.

Category: Fruits	
Item	Total Sales
Apples	\$ 5,660
Bananas	\$ 7,150
Oranges	\$ 5,260
<b>Total Fruits</b>	<b>\$ 18,070</b>
Category: Vegetables	
Item	Total Sales
Carrots	\$ 3,530
Lettuce	\$ 4,045
Peas	\$ 3,675
<b>Total Vegetables</b>	<b>\$ 11,250</b>
<b>Grand Total</b>	<b>\$ 29,320</b>

# Metadata Properties Reference

This section contains the Metadata Properties Reference.

## Out-of-the-Box Properties

The table below shows the properties that are provided with Pentaho Metadata Editor.

\*Localized Properties

ID	Description	Category	Values
Name*	This property describes the display name for the business object	General	Alphanumeric
Description*	A descriptive text entry describing the business object	General	Alphanumeric
Comments*	Additional comments regarding the business object.	General	Alphanumeric
Security Information	Security rules for granting/restricting access to the business object	General	Determined by security widget; see Adding Security
Font	The font properties to apply to this business object	Formatting	Determined by selections in the font dialog box
Color of Text	The foreground or text color for the business object	Formatting	Determined by selections in the color dialog box
Text Alignment	The text alignment for the business object	Formatting	Left   Right   Centered   Justified
Background color	The background color for the business object	Formatting	Determined by selections in the color dialog box
Relative Size	This property is normally associated with business tables and is used to calculate join paths. The sum of all table relative sizes in a path are calculated when deciding on a multi-table join and the multi-table join with the smallest summed value is used for the join path.	Formatting	Numeric
Aggregation Rule	Determines the method of aggregating the data from this business object	Model Descriptors	None   Sum   Count   Distinct Count   Minimum   Maximum
Data Type	Data type for this business object	Model Descriptors	Unknown   String   Date   Boolean   Numeric   Binary   Image   URL + Length and Precision (Integers)
Field Type	The type or relationship purpose this field serves	Model Descriptors	Other   Dimension   Fact   Key   Attribute
Table Type	Table type is used to automatically determine relationship types with other tables. For instance, if a Fact table is joined with a Dimension table, this is normally an N to 1 relationship	Model Descriptors	Other   Dimension   Fact

ID	Description	Category	Values
Formula	This property allows you to create a calculation defining the business object. See <a href="#">Pentaho Metadata Formulas</a> Pentaho Metadata Formulas for more information.	Calculation	Alphanumeric
Is the Formula Exact?	Determines if the formula must be parsed as a Metadata Formula. If false, this field is treated as a database column name .	Calculation	Boolean
Column Width	The width of the column as represented for display	Miscellaneous	Pixels   Percent of Page Width   Inches   Centimeters   Points ) + Integer
Hidden For the User?	Hides the object from being displayed in the Business View of the model	Miscellaneous	Boolean
Mask for Number or Date	The format mask to use when the data for this object is displayed. For dates, the format follows the date/ time patterns for a Java <a href="#">SimpleDateFormat</a> format mask. For numbers, the pattern should follow the Java <a href="#">DecimalFormat</a> format mask.	Miscellaneous	Alphanumeric
Target Schema	Defines the database schema to use when querying the business object	Miscellaneous	Alphanumeric
Target Table	Defines the physical database table from which the business object is defined	Miscellaneous	Alphanumeric



## Custom Properties

You can define any number of custom properties. You must give your property an ID to set its type.

- String
- Date
- Numeric Value
- Color
- Font
- Type of Field
- Type of Aggregation
- Boolean
- Field Data Type
- Localized String
- Type of Table
- URL
- Security
- Text Alignmen
- Column Width

## Required Properties per Business Object

The table below contains the required properties for given business objects. Required properties cannot be deleted.

Required =  Not Required = 

ID	Physical Table	Physical Column	Business Category	Business Model
Name				
Description				
Security Information				
Table Type				
Relative Size				
Formula				
Field Type				
Data Type				
Aggregation Rule				
Is the Formula Exact?				
Hidden for the User?				
Font				
Mask for Number or Date				
Color of Text				
Color of Background				

# Pentaho Metadata Formulas

Formulas have multiple uses in Pentaho Metadata. The first use of formulas in Pentaho Metadata is in the constraint definition of a Metadata Query, also known as MQL. A constraint function references business table columns and uses various comparison operators to determine which subset of data the business user is interested in.

The second use is in the definition of Physical Table Columns. In addition to Physical table columns mapping directly to a database table column, physical table columns defined in Pentaho Metadata may also be defined as a formula. This allows for combining of multiple columns into a single column, and also for doing more advanced aggregate calculations within aggregate table definitions.

The third use is in the definition of complex joins within business model relationships. This allows for multiple key joins as well as other logic when joining tables.

The fourth use is row level security. Under the covers, Pentaho Metadata uses JFreeReport's libFormula package for interpreting formulas. The goal is to support OpenFormula syntax in the Metadata environment. Formulas are first interpreted by libFormula, and then in the Metadata system are converted to native SQL depending on the type of database used.

## First Use — MQL Constraints

Below is an example of an MQL Constraint formula:

```
OR ([BT_CUSTOMERS.BC_CUSTOMERS_CUSTOMERNAME] = "EuroCars";  
    (( [BT_CUSTOMERS.BC_CUSTOMERS_CREDITLIMIT] * 2) / 3 > 1000))
```


Note the `OR` function. This is a boolean function that has two parameters, separated by semicolons. These parameters are boolean expressions. The first boolean expression references a business column from the Metadata model. All references appear with brackets around them `[]`. This reference refers to the business table, and then to the business column. This boolean expression does some arithmetic and checks to see if the final value is larger than 1000. In the second expression, business column `BT_CUSTOMERS.BC_CUSTOMERS_CUSTOMERNAME` is compared to `EuroCars`. Notice that double quotes are used when referring to text; double quotes are required.

## Second Use — Physical Table Column Formulas

Below is an example of a Physical Table Column Formula:

```
[QUANTITYORDERED] * [PRICEEACH]
```

The references in this example specifically refer to the database column, not derived the physical column definitions. All operators and functions may be used in the definition of the physical table column. For this formula to be recognized, the `isExact` property of the physical table column must be set to `true`. The referenced physical column must be explicitly defined in the metadata model.

 **Note:** In earlier versions of Pentaho Metadata Editor, (prior to the 3.5 release), aggregation functions had to be specified explicitly and the aggregation rule had to be selected. This is no longer necessary; the query that is generated will use the selected aggregation rule during execution.

## Multi-table expressions — Formulas can use any business column in the model

Since the latest versions (after 2008/03/14) it is possible to define formulas that use business columns from anywhere in the business model. For example suppose there are two business tables:

- Orders (fact table), ID=BT\_ORDER\_FACT
- Product (dimension), ID=BT\_PRODUCT

Suppose you want to calculate the turnover based on:

- Orders (fact table), ID=BT\_ORDER\_FACT
- Product (dimension), ID=BT\_PRODUCT
- The number of products sold, from the Orders table, ID=BC\_FACT\_ORDER\_NRPRODUCTS
- The price of the product, from the Product table, ID=BC\_DIM\_PRODUCT\_PRICE

To get there, you must define a new business column, say in the Orders business table (although you could take Product too):

- Table: Orders (BT\_ORDER\_FACT)
- ID = BC\_FACT\_ORDER\_TURNOVER
- Name = Turnover
- Formula = [BT\_ORDER\_FACT.BC\_FACT\_ORDER\_NRPRODUCTS]
- Exact = Yes
- Aggregation Rule = SUM

The SQL generator is now going to replace the business columns by their respective SQL variants. As such, you must make sure that the business columns are resolving correctly. In this specific case, this means you want the two columns to be non-aggregated. If you now select the single business column BT\_FACT\_ORDER\_TURNOVER, below is the SQL that is generated:

```
SELECT
    SUM( BT_ORDER_FACT.NRPRODUCTS * BT_PRODUCT.PRICE ) AS COL0
FROM
    FACT_ORDER BT_ORDER_FACT ,DIM_PRODUCT BT_PRODUCT
WHERE
    ( BT_ORDER_FACT.PRODUCT_TK =
    BT_PRODUCT.PRODUCT_TK )
```

Now, suppose you want to generate the multiplication of the two sums (different use-case). You define the formula as "[BT\_ORDER\_FACT.BC\_FACT\_ORDER\_NRPRODUCTS] \* [BT\_PRODUCT.BC\_DIM\_PRODUCT\_PRICE]" (without the SUM) and specify an aggregation for the two business columns in use. The generated SQL will be as follows:

```
SELECT
    SUM( BT_ORDER_FACT.NRPRODUCTS ) * SUM( BT_PRODUCT.PRICE ) AS COL0
FROM
    FACT_ORDER BT_ORDER_FACT ,DIM_PRODUCT BT_PRODUCT
WHERE
    ( BT_ORDER_FACT.PRODUCT_TK =
    BT_PRODUCT.PRODUCT_TK )
```

It is possible to create two versions of the used business columns, one aggregated (exposed to the users) and one non-aggregated (hidden from the users) for example.

The SQL generator works recursively. That means that it is possible to create a formula that calculates 7% (taxes for example) of the turnover:

- ID = BC\_FACT\_ORDER\_TURNOVER\_TAXES
- Name = Turnover Taxes
- Formula = [BT\_ORDER\_FACT.BC\_FACT\_ORDER\_TURNOVER] \* 7 / 100
- Exact = Yes

If you add that column to the selection, you get one extra column as shown below:

```
( SUM( BT_ORDER_FACT.NRPRODUCTS * BT_PRODUCT.PRICE ) * 7 / 100 ) AS COL1
```

## Formula Syntax

*Function syntax:*

```
FUNCTION_NAME ( PARAM ; PARAM )
```

*Text (requires double quotes):*

```
"TEXT"
```

*Parenthesis are used for formula precedence:*

```
( 1 + 2 ) * 3
```

## Metadata References

*Business Column References:*

```
[<BUSINESS_TABLE_ID> . <BUSINESS_COLUMN_ID> ]
```

*Physical Column References (only used in physical column formula definitions):*

```
[ <PHYSICAL_COLUMN_NAME> ]
```



## Supported Functions

The table below contains a listing of supported functions. Examples are show below each supported function.

Function Name	Parameters	Description
OR	Two or more boolean expression parameters	Returns true if one or more parameters are true

```
OR( [BT_CUSTOMERS.BC_CUSTOMERS_CUSTOMERNAME] = "EuroCars";
    [BT_CUSTOMERS.BC_CUSTOMERS_CREDITLIMIT] > 1000 )
```

Function Name	Parameters	Description
AND	Two or more boolean expression parameters	Returns true if all parameters are true

```
AND( [BT_CUSTOMERS.BC_CUSTOMERS_CUSTOMERNAME] = "EuroCars";
    [BT_CUSTOMERS.BC_CUSTOMERS_CREDITLIMIT] > 1000 )
```

Function Name	Parameters	Description
LIKE	Two parameters	Compares a column to a regular expression, using "%" as wild cards

```
LIKE( [BT_CUSTOMERS.BC_CUSTOMERS_CUSTOMERNAME]; "%SMITH%" )
```

Function Name	Parameters	Description
IN	Two or more parameters	Checks to determine if the first parameter is in the following list of parameters

```
IN( [BT_CUSTOMERS.BC_CUSTOMERS_CUSTOMERNAME]; "Adam Smith"; "Brian Jones" )
```

Function Name	Parameters	Description
NOW	none	The current date

```
NOW( )
```

Function Name	Parameters	Description
DATE	Three numeric parameters, year, month, and day	A specified date

```
DATE( 2008; 4; 15 )
```

Function Name	Parameters	Description
DATEVALUE	One text parameter "year-month-day"	A specified date

```
DATEVALUE( "2008-04-15" )
```

Function Name	Parameters	Description
CASE	Two or more parameters	Evaluates the first, third, etc. parameter, and returns the second,

Function Name	Parameters	Description
		fourth, etc parameter value if there are an odd number of parameters, the last parameter is returned if no other parameter evaluates to true. Note that when using this function, the formula must be set on a new column as shown below.

```

CASE( [BT_CUSTOMERS.BC_CUSTOMERS_CUSTOMENAME] = "EuroCars"; "European
Cars";
      [BT_CUSTOMERS.BC_CUSTOMERS_CUSTOMENAME] = "AsiaCars";
      "Asian Cars"; "Unknown Cars"
)

```

Function Name	Parameters	Description
COALESCE	One or more parameters	Returns the first non null parameter

```

CASE( [BT_CUSTOMERS.BC_CUSTOMERS_CUSTOMENAME] = "EuroCars"; "European
Cars";
      [BT_CUSTOMERS.BC_CUSTOMERS_CUSTOMENAME] = "AsiaCars";
      "Asian Cars"; "Unknown Cars"
)

```

## Supported Operators

The table below contains a listing of supported operators.

Operator	Description
=	Returns true if two expressions are equal
>	Returns true if first expression is larger than the second
<	Returns true if first expression is smaller than the second
>=	Returns true if first expression is larger than or equal to the second
<=	Returns true if first expression is smaller than or equal to the second
<>	Returns true if two expressions are not equal
+	Adds two values
-	Subtracts two values
*	Multiplies two values
/	Divides two values

## Supported Aggregate Functions

The table below contains a listing of supported aggregate functions.

Function Name	Description
SUM	Sums a specific columns values determined by grouping
COUNT	Counts a specific columns values determined by grouping
AVG	Averages a specific columns values determined by grouping
MIN	Selects the minimum column value determined by grouping
MAX	Selects the maximum column value determined by grouping

# Troubleshooting

---

This section contains reported or expected problem descriptions and solutions.

## Hive Database Disappears From Database Connection Dialogue

---

If your Hive database type disappears from the list in the database connection dialogue in Report Designer or Metadata Editor, this probably means that your Pentaho BI Suite For Hadoop license has expired, or the license file has been moved, deleted, or become corrupt. You must install a valid license key in order to reintroduce this functionality in Report Designer.

If you do not update the licence key, all of your existing Hive-based reports will still work. However, you will not be able to alter the Hive connection details or create a new Hive connection.