EspressReport® ES Cloud Version 7.2

ERES Cloud User's Guide

Quadbase

Quadbase Systems Inc.

ERES Cloud User's guide vi
Q. QuickStart 1
Q.1. Overview 1
Q.2. Start the Organizer 1
Q.3. Use Data Sources 1
Q.3.1. Open a Data Source Manager 1
Q.4. Report Designer 13
Q.4.1. Report Mapping 13
Q.4.2. Basic Report Formatting 27
Q.4.3. Advanced Reporting Features 36
Q.5. Chart Designer 49
Q.5.1. Chart Mapping 49
Q.5.2. Basic Chart Formatting 56
Q.5.3. Advanced Charting Features 61
Q.6. QuickDesigner Reports 64
Q.6.1. Create a Query 64
Q.6.2. Create a Report 69
Q.7. QuickDesigner Charts 74
Q.7.1. Create a Chart 74
Q.8. Online Maps 79
Q.8.1. Create Coordinates 79
Q.8.2. Create Online Map 83
Q.9. SVG Maps 90
Q.9.1. Create a Map 90
Q.9.2. Set SVG Map Thresholds 92
Q.9.3. Set SVG Map Drilldowns 95
Q.10. Publishing 96
Q.10.1. Dashboards 96
Q.10.2. URLs 108
Q.10.3. The Menu Page 110
Q.11. Alerts 110
Q.11.1. Alerts in Charts 110
Q.11.2. Alerts in Maps 112
Q.11.3. Alerts in Dashboards 113
Q.11.4. Alert Monitoring 114
1. Features and Release Notes 119
1.1. New Features List V7.0 119
1.1.1. New ERES Cloud Features 119
2. ERES Cloud Main Page 121
2.1. Upload Files 121
2.1.1. Excel Data Sources 122
2.1.2. Upload Other Files 124
3. Organizer 127
3.1. Using the Organizer 127
3.1.1. Starting the Organizer 127
3.1.2. The Organizer Interface 128
3.1.3. Projects and Folders 131
3.1.4. Files 131
3.1.5. Searching in Organizer 135
3.2. Scheduling & Archiving 136
3.2.1. Setting a Schedule 136
3.2.2. Setting an Archive 148
3.2.3. View Tasks 149
4. Data Sources 155
4.1. Data in Organizer 155
4.1.1. The Data Source Manager 155
4.1.2. Data from a Database 157
4.1.3. Data from XML and XBRL Files 188

4.1.4. Data from Text Files 193
4.1.5. Data from EJBs 194
4.1.6. Data from SOAP with WSDL support 196
4.1.7. Data from Salesforce 198
4.1.8. Data from Excel files 201
4.1.9. Using Data for Charts and Reports 203
4.1.10. Data Source Updating 206
4.1.11. CData JDBC drivers 208
4.2. Data in QuickDesigners and Maps 213
4.2.1. Select a Data Source 213
4.2.2. Queries 214
5. Designing Reports & Charts 225
5.1. Report Designer 225
5.1.1. Introduction to Report Designer 225
5.1.2. Report Types and Data Mapping 226
5.1.3. The Designer Interface 245
5.1.4. The Preview Window 299
5.1.5. Saving Reports 304
5.1.6. Using Formulas & the Formula Builder 305
5.1.7. Scripting 319
5.1.8. Drill-Down 330
5.1.9. Sub-Reports 337
5.1.10. Template Security 341
5.2. Chart Designer 345
5.2.1. Introduction to Chart Designer 345
5.2.2. Charting Basics 349
5.2.3. Chart Types and Data Mapping 354
5.2.4. The Chart Designer Interface 386
5.2.5. Chart Drill-Down 453
5.2.6. Saving Charts 458
5.3. QuickDesigner Reports 460
5.3.1. Start 460
5.3.2. Select a Data Source 460
5.3.3. Toolbar 461
5.3.4. Format Report Elements 462
5.3.5. Parameter Setting 469
5.3.6. Aggregation/Group 470
5.3.7. Data Formating 472
5.3.8. Sorting 475
5.3.9. Data Filtering 475
5.3.10. Save the Report 476
5.3.11. Export the Report 476
5.3.12. Open the Saved Report 477
5.3.13. Exit 477
5.4. QuickDesigner Charts 477
5.4.1. Start 478
5.4.2. Select a Data Source 478
5.4.3. Data Mapping and Ordering 479
5.4.4. Toolbar 481
5.4.5. Data Formatting 482
5.4.6. Save the Chart 502
5.4.7. Export the Chart 503
5.4.8. Open the Saved Chart 504
5.4.9. Exit 504
6. Designing Maps 505
6.1. Introduction to ERES Maps 505
6.2. Online Maps 505
6.2.1. Data Sources 505

6.2.2. Start Online Maps 506
6.2.3. Coordinates Editor 506
6.2.4. Coordinates Mapping 516
6.2.5. Create Online Maps 517
6.2.6. Save Map 529
6.2.7. Open Saved Map 530
6.2.8. Exit 531
6.3. SVG Maps 531
6.3.1. Introduction to SVG Maps 531
6.3.2. Area ID Mapping 531
6.3.3. SVG Maps Designing 531
6.3.4. Dynamic SVG Maps 542
6.3.5. Save Map 546
6.3.6. Open Saved Map 547
6.3.7. Exit 547
6.4. Map Viewer 547
6.4.1. Writing Map Viewer URL 548
6.5. Migrating Maps 549
6.A. List of SVG Map Images 549
7. Designing Dashboards 551
7.1. Introduction to Dashboards 551
7.2. Create Dashboard 552
7.2.1. Toolbar 553
7.2.2. Responsive Dashboard 554
7.2.3. Add Charts, Reports and Maps 555
7.2.4. Shared Parameters 558
7.2.5. Insert Labels 573
7.2.6. Add Background 575
7.2.7. Additional Options 576
7.2.8. Drilldown 578
7.2.9. Folders 579
7.2.10. Template Linkage 583
7.2.11. Dashboard Preview 584
7.3. Save Dashboard 584
7.4. Open Saved Dashboard 584
7.5. Exit 585
8. Publishing (Menu & URLs) 586
8.1. The Menu Page 586
8.1.1. Launching the Menu Page 586
8.1.2. Using the Menu Page 586
8.1.3. Mobile MenuPage 594
8.2. Image URLs 598
8.2.1. Generating URLs in Organizer 598
8.2.2. Writing Image URLs 600
8.3. Report URLs 617
8.3.1. Generating URLs in Organizer 617
8.3.2. Writing Report URLs 619
8.4. Dashboard URLs 631
8.4.1. Generating URLs in Organizer 631
8.4.2. Writing Dashboard URLs 633
8.5. Menu Page Listener 634
8.5.1. ERES Listener Manager 634
8.5.2. Using the Menu Page Listener 634
8.6. Report Viewer 636
8.6.1. The Report Viewer Parameters 637
8.6.2. Specifying Data for Report Viewer 638
8.6.3. Using Report Viewer 639
8.6.4. Connecting to the ERES Server 641

8.6.5. Swing Version 641
8.6.6. Exporting from Viewer 641
8.7. Page Viewer 641
8.7.1. Launching Page Viewer 642
8.7.2. The Page Viewer Parameters 643
8.7.3. Using Page Viewer 644
8.7.4. Buffer Time 645
8.7.5. Connecting to the ERES Server 645
8.7.6. Swing Version Available 645
8.7.7. Exporting from Viewer 645
8.7.8. Unmapped Drilldown Parameter 646
8.8. Chart Viewer 647
8.8.1. The Chart Viewer Parameters 648
8.8.2. Specifying the Data Source for Chart Viewer 650
8.8.3. Using Chart Viewer 652
8.8.4. Axis Rulers 654
8.8.5. Pop-up Menu 654
8.8.6. Swing Version Available 655
8.9. Dashboard Viewer 655
8.9.1. Preview Toolbar 655
8.9.2. Preview Options 656
9. Alerts 660
9.1. What is an Alert 660
9.2. Specifying alerts 660
9.2.1. Online Maps 660
9.2.2. SVG Maps 660
9.2.3. Charts 661
9.2.4. Reports 661
9.3. Dashboard alerts 662
9.3.1. Visual alert display 662
9.3.2. Alert messages 662
9.3.3. Sound alert 663
9.3.4. User interface 663
9.4. Monitoring 664
9.4.1. Alert notifications 664
9.4.2. Failed emails 664
9.4.3. User interface 664
9.5. Handling special situations 669
9.5.1. Change in script/control area/range 669
9.5.2. Missed monitoring jobs 669

ERES Cloud User's guide

Welcome to the ERES Cloud Online User's Guide.

Here you will find information about all the functions and features in EspressReport ES Service.



Tip

If you're reading the PDF version, enable *Bookmarks* toolbar in your PDF viewer.

Please select an option from the menu on the left to begin.

QuickStart

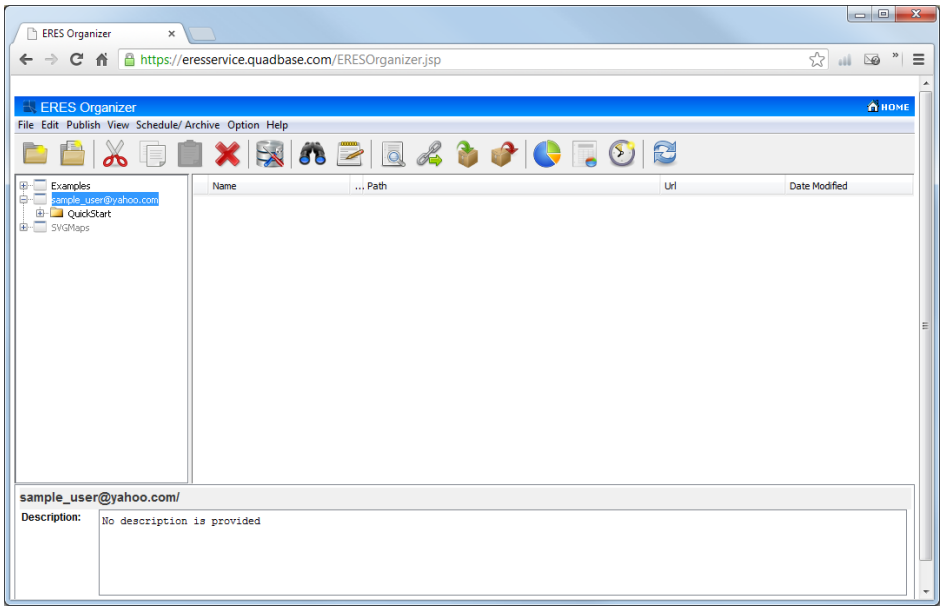
Q.1. Overview

Welcome to the ERES Cloud Quick Start Guide. This guide is the recommended starting point for users who are new to ERES and would like to get up to speed fairly quickly. This guide can also help users to quickly evaluate the product capabilities.

This guide provides step-by-step tutorials for report and chart creation and publishing features. Only a small portion of the features are described in this guide so for more information and detail, please refer to the other sections in the User's Guide.

Q.2. Start the Organizer

Click on the link labeled *Organizer* on the Start page. This will open the Organizer JSP page. Note that since the Organizer launches as an applet, it will take a few seconds to load.



Organizer Interface

The Organizer is a virtual file management system that allows you to create, manage, publish, and schedule reports and charts. The Organizer is used in conjunction with most of the other interfaces in ERES, and will be referenced extensively throughout this guide.

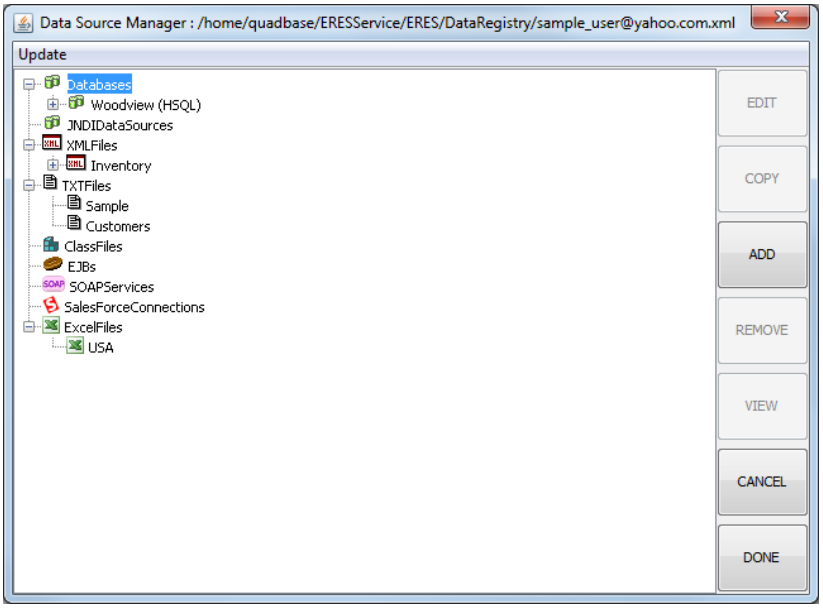
Reports and charts in the Organizer are grouped into projects and folders. There are three projects in the left side panel. The projects *Examples* and *SVGMaps* are read-only. Your project was created automatically with the name of your email address and it includes *QuickStart* folder with samples you will use in exercises.

Q.3. Use Data Sources

Data sources in ERES are maintained in XML data registry file. Your data registry file was created automatically with the name of your email address (your_email.xml). This file is managed from within the Organizer. Detailed documentation about the data sources can be found in Section 4.1 - Data in Organizer.

Q.3.1. Open a Data Source Manager

Open the Organizer and click the *Manage Data Sources* button  . The Data Source Manager window for your registry will open.

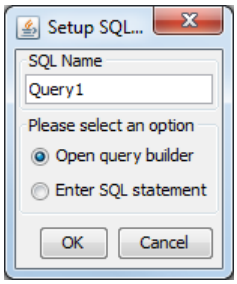


Data Source Manager Window

Q.3.1.1. Create a Query

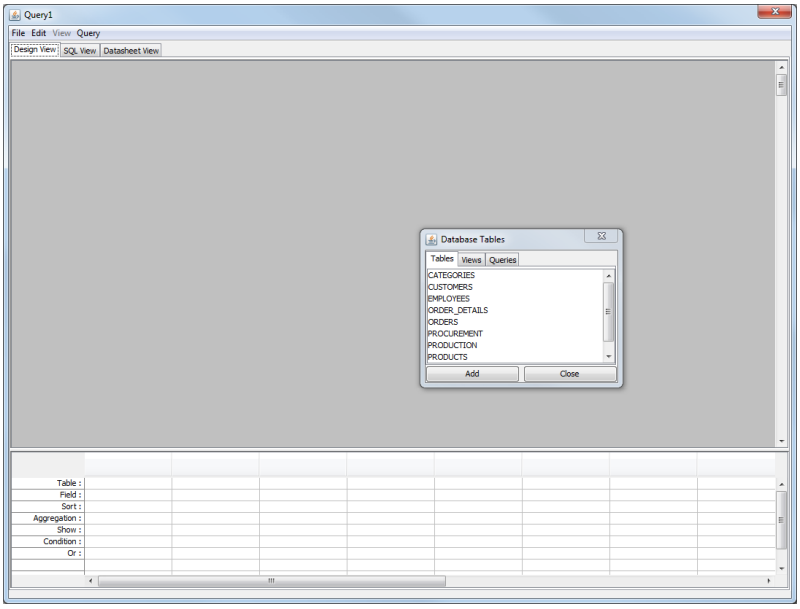
ERES provides a number of different interfaces to query a database to retrieve the report data. You can type a SQL statement, use the Query Builder, or use data views to create a query interface that insulates the end user from the database structure. In this example, we will use the Query Builder to create a query.

To create a new query, click to expand the *Woodview* node in the left-hand frame of the Data Source Manager. Two sub-nodes will appear, one called *Queries* and one called *Data Views*. Select the *Queries* node and click *Add*. A dialog will appear prompting you to specify a name for the query and to select whether to launch the Query Builder or to enter an SQL statement.



Query Name Dialog

Enter any name you would like, select *Open query builder*, and click on *OK*. The Query Builder will launch. You will see a separate window containing all of the tables for Woodview sitting over top of the main Query Builder window.

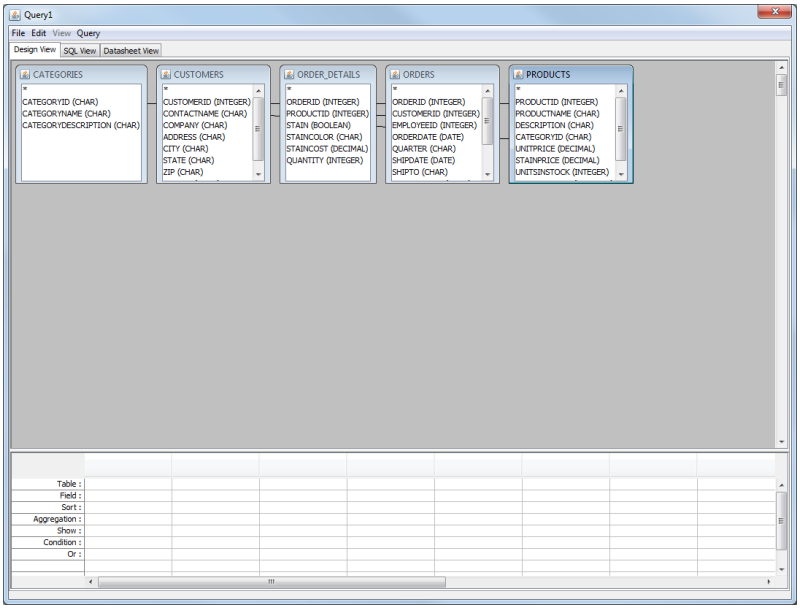


Query Builder Dialog

To add a table to the query, select a table in the Tables window and click the *Add* button. You can also double click on the table name. Using one of the two methods, add the following tables to the query:

```
CATEGORIES
CUSTOMERS
ORDER_DETAILS
ORDERS
PRODUCTS
```

The tables will appear in the top half of the Query Builder window. You will see the join lines connecting various fields in the tables.

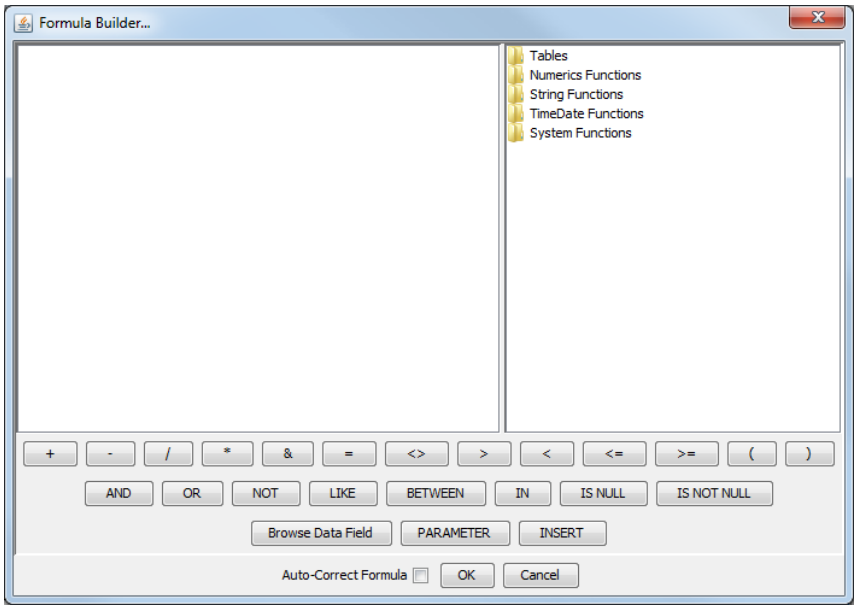


Query Builder with Tables

To add a field to the query, you can either double click on the field in the table window or double click on the *Table* and *Field* fields in the lower (QBE) portion of the Query Builder window and select the table and field from the drop-down menu. Using either method, add the following fields to the query.

```
ORDERID from ORDERS
COMPANY from CUSTOMERS
REGION from CUSTOMERS
CATEGORYNAME from CATEGORIES
PRODUCTNAME from PRODUCTS
UNITPRICE from PRODUCTS
QUANTITY from ORDER DETAILS
```

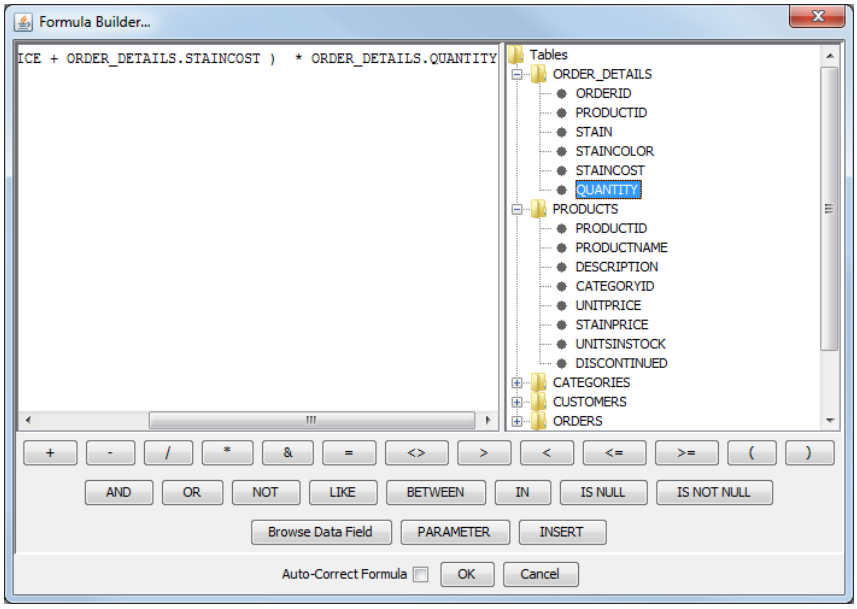
In the eighth column, which should be blank, right click in the*Field* field and select *Build* from the pop-up menu. This will open the Formula Builder interface allowing you to create a computed column.



Formula Builder Window

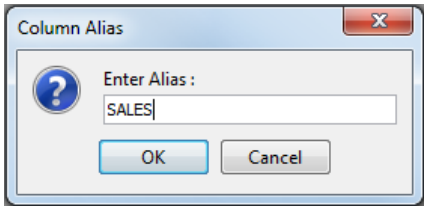
To build a column, first click the *left parenthesis* button. Then double click on the *Tables* folder. It will expand into five nodes, one for each table you selected for the query. Opening a table folder will list all of the column fields in that table. Open up the *PRODUCTS* folder, select *UNITPRICE* and click *Insert*. Then click the *add (+)* button. Next, insert *STAINCOST* from the *ORDER_DETAILS* table. Then click the *right parenthesis* button. Click the *multiply (*)* button and finally insert *QUANTITY* from the *ORDER_DETAILS* column. The finished formula should look like this:

```
(PRODUCTS.UNITPRICE + ORDER_DETAILS.STAINCOST) * ORDER_DETAILS.QUANTITY
```



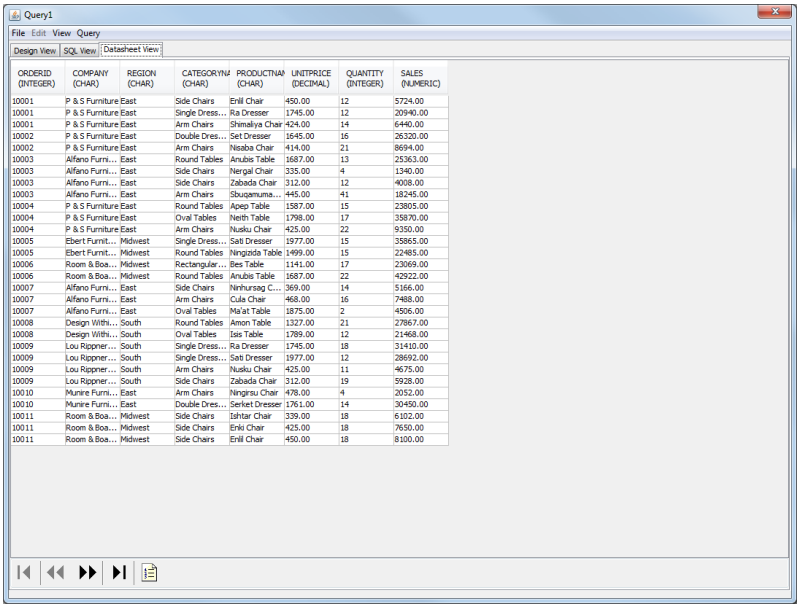
Formula Builder Window with Formula

Click the *OK* button and the built column will be added to the query. Next, we will give the column you built an alias. Right click on the column and select *Alias* from the pop-up menu. A window will appear asking you to enter a column alias. Enter "SALES" (without quotation marks) and click the *OK* button.



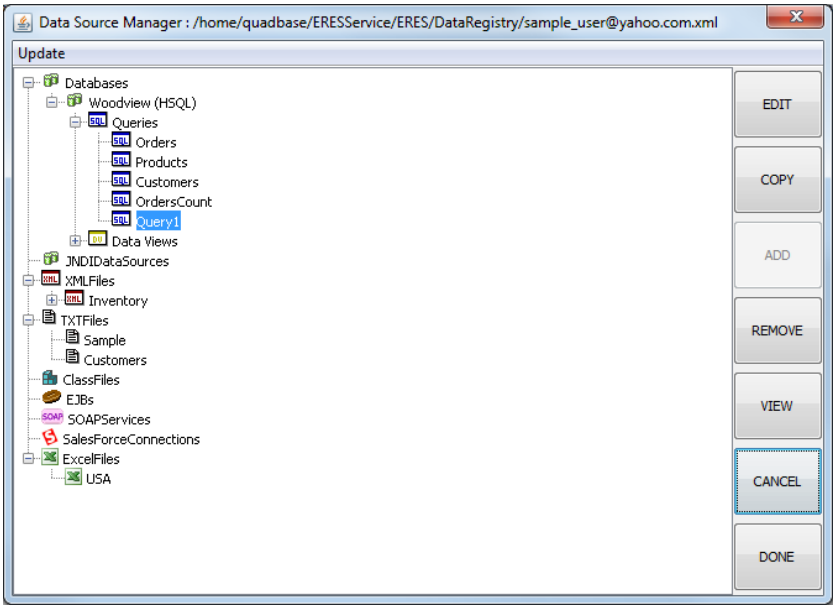
Column Alias Dialog

Click the *OK* button and you will see the column name change in the Query Builder. Now click on the *Datasheet View* tab in the Query Builder. Your query will run and you should see the first thirty records of the query results.



Query Builder Datasheet View

Now that you have finished designing the query, select *Done* from the File menu to save the changes. This will close the Query Builder window and return you to the Data Registry Manager window. There will now be a node under *Queries* for the query you have just designed.



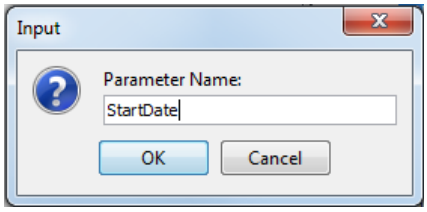
Data Source Manager With Query

Q.3.1.1.1. Add Query Parameters

ERES allows you to easily parameterize queries, allowing report and chart data to be dynamically filtered at run-time. In this tutorial we will add parameters to the query created in Section Q.3.1.1 - Create a Query.

To open the query that you created, select it, click the *Edit* button in the Data Source Manager, and click the *OK* button. Your query will re-open in the Query Builder. The Tables window will open on top of the Query Builder. Click *Close* to close the Tables window and scroll down in the lower (QBE) portion of the Query Builder window until you see the *Condition* field. Right click in the *Condition* field under the *ORDERID* column and select *Build* from the pop-up menu. This will bring up the Formula Builder allowing you to construct a condition for the query.

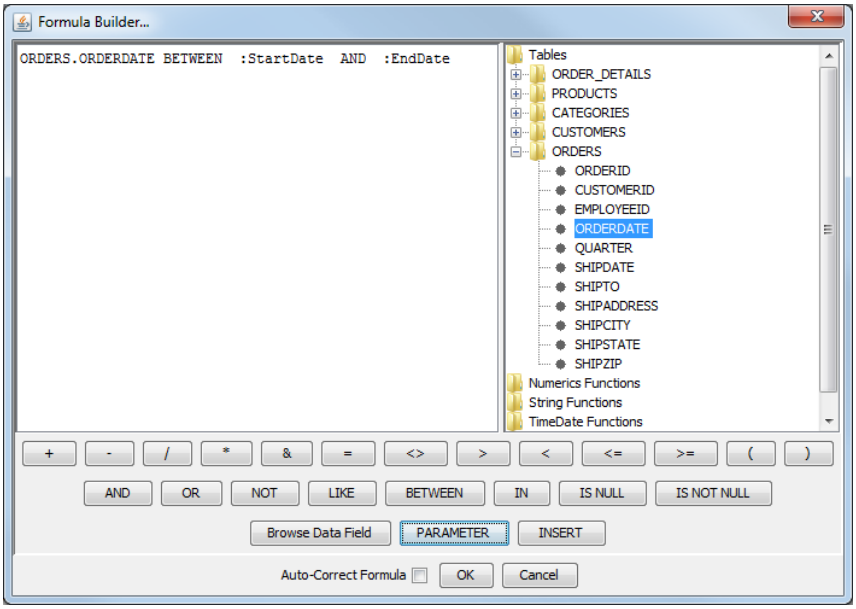
Double click on the *Tables* folder within the Formula Builder to expand it. Then expand the *ORDERS* node and double click on the *ORDERDATE* field. Next click the *Between* button and then click the *Parameter* button. This will bring up a dialog prompting you specify a name for the query parameter.



Parameter Name Dialog

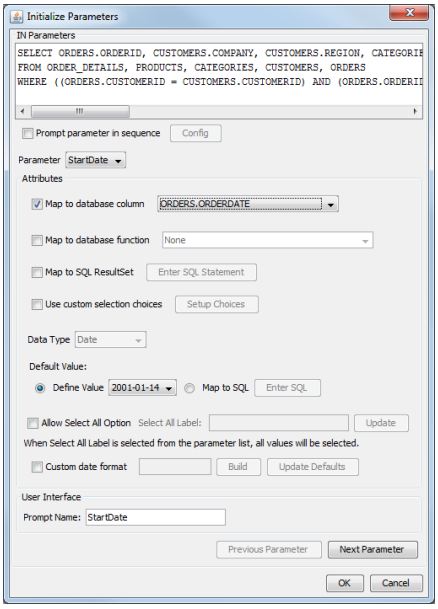
Enter **`StartDate`** as the parameter name and click the *OK* button. The parameter will be added to the query. Then click the *And* button. Click the *Parameter* button again. Enter **`EndDate`** as the second parameter name. The finished condition should look like this

```
Orders.OrderDate BETWEEN :StartDate AND :EndDate.
```



Formula Builder with Conditions

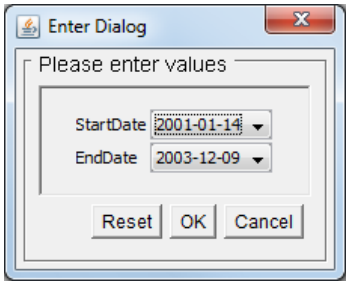
Click *OK* to close the Formula Builder and return to the Query Builder window. Now click on the *Datasheet View* tab. Because you have just added two parameters to the query, an initialization dialog will appear, asking you to specify a few properties for the query parameters.



Parameter Initialization Dialog

From this window, click on *Map to a database column* and select *ORDERS.ORDERDATE* from the drop-down menu. This will automatically fill the *Default Value* and *Data Type* options. Next, enter **`Start Date`** into the *Prompt Name*, then click the *Next Parameter* button and map the *EndDate* parameter to the same column. Click on the *Define Value* drop-down menu to select an end date. Select a date far enough from the start date that by default you will have more than a couple records to work with (this makes report design easier). Change its *Prompt Name* to **`End Date`**.

Click *OK* to close the initialization window once you have specified all the options. A new dialog will appear prompting you to select a date range by which to filter the result set.



Parameter Selection Dialog

Select the Start and End date that you would like and click *OK*. You will now see the filtered result in the datasheet window. Now, click *Done* from the File menu to save the changes you have made to the query.

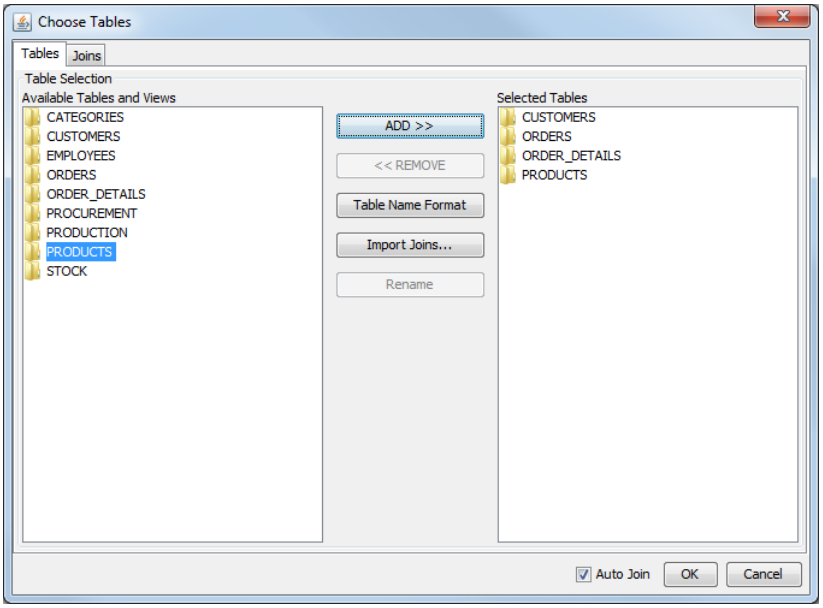
Q.3.1.2. Create a Data View

A unique feature in ERES is the ability to create data views. Data views are local schemas/views that allow an administrator to pre-configure a group of tables and fields, so that the end users only need to select fields and define simple conditions to create a query. Data views are also used with the ad hoc QuickDesigner. To create a data view, select the *Data Views* node under *Woodview* and click *Add*.

This will open a new dialog asking you to select database tables you would like to use. Select the following tables and add them to the *Selected Tables* panel by clicking on the *ADD* >> button:

```
CUSTOMERS
ORDERS
ORDER_DETAILS
```

```
PRODUCTS
```



Data View Tables Dialog

Next, click on the *Joins* tab. You will see a representation of the tables like in the Query Builder. You can see the auto-join lines between the tables. This window can be used to join the tables or modify the auto-joins if necessary (for more about joins please see Section 4.1.2.2.1.2 - Joins). Click *Ok* to finalize the table selection. The next window allows you to select and group fields for the view. At the top of the window you can specify a name for the view. Name it **`Invoices`**.

Next, double click on the *CUSTOMERS* folder to reveal the fields for that table. Add the following fields by double clicking them or selecting them and clicking on the *Add* >> button:

```
COMPANY
CONTACTNAME
ADDRESS
CITY
STATE
ZIP
```

Now add fields from the other tables as follows:

```
ORDERS:

ORDERDATE
SHIPDATE
SHIPTO
SHIPADDRESS
SHIPCITY
SHIPSTATE
SHIPZIP

ORDER_DETAILS:

ORDERID
STAIN
STAINCOLOR
QUANTITY

PRODUCTS:
```

```
PRODUCTNAME
UNITPRICE
STAINPRICE
```

Now click the *Add Heading* button. At the prompt, specify the name **`Customer Info`**. Add two more headings in the same way, one called **`Shipping Info`** and one called **`Order Info`**. Once they are created, select the following fields (Using **CTRL**+**Click** or **SHIFT**+**Click** for multiple selection):

```
COMPANY
CONTACTNAME
ADDRESS
CITY
STATE
ZIP
```

Once the fields are selected, click the *Group Fields* button and select *Customer Info* from the drop-down list. The fields will be moved under that heading. Next, select the following fields in the same way:

```
SHIPTO
SHIPADDRESS
SHIPCITY
SHIPSTATE
SHIPZIP
SHIPDATE
```

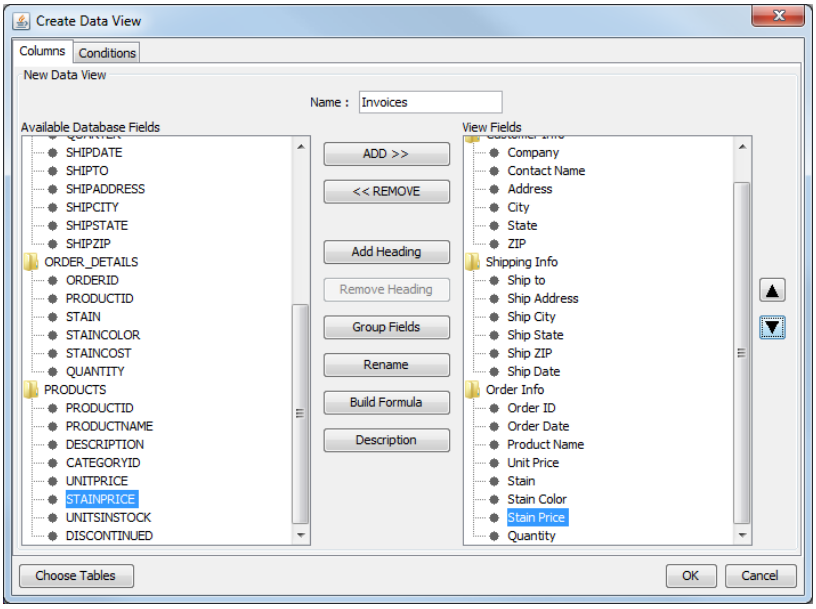
Add these fields to the *Shipping Info* group the same way as before. Next, select the following fields:

```
ORDERDATE
ORDERID
STAIN
STAINCOLOR
QUANTITY
PRODUCTNAME
UNITPRICE
STAINPRICE
```

Add these fields to the *Order Info* group. Next, select the *CONTACTNAME* field on the right side and click the *Rename* button. In the dialog, specify the name **`Contact Name`**. Repeat this for every field in order to give it proper names.

Next, select the *Order ID* field on the right side and click on the up arrow button to move the field to the top of the *Order Info* heading. Use the arrows to arrange the items in the *Order Info* heading in the following order:

```
Order ID
Order Date
Product Name
Unit Price
Stain
Stain Color
Stain Price
Quantity
```



Data View Fields Window

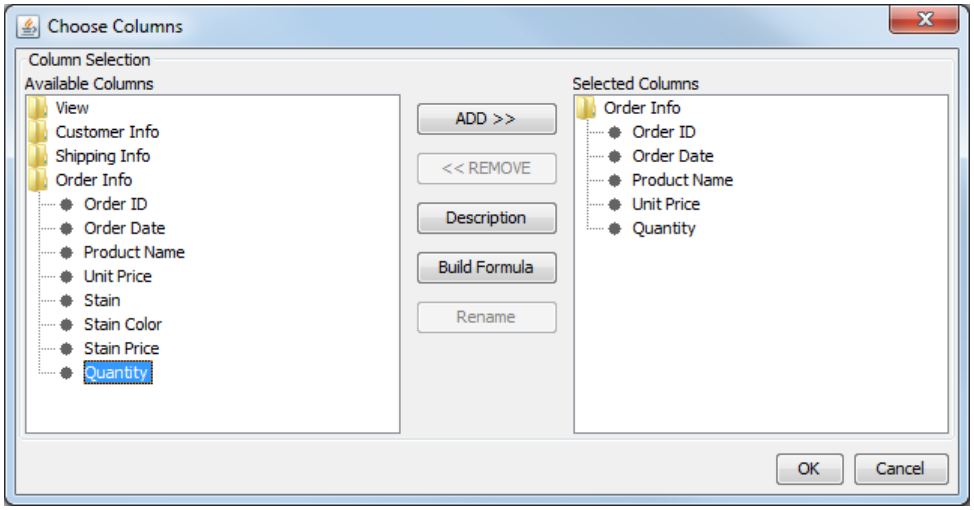
Now that the data view has been created, click the *OK* button in the fields window to save the view. It will be saved as a new node under *Data Views* in the Data Source Manager.

Q.3.1.2.1. Query a Data View

Now that the data view has been created, you can write queries against the view. This allows users to develop queries without knowing the underlying structure of the database. It also allows administrators to limit which database elements the user has access to. In this tutorial we will create a query for the data view you created in. Section Q.3.1.2 - Create a Data View.

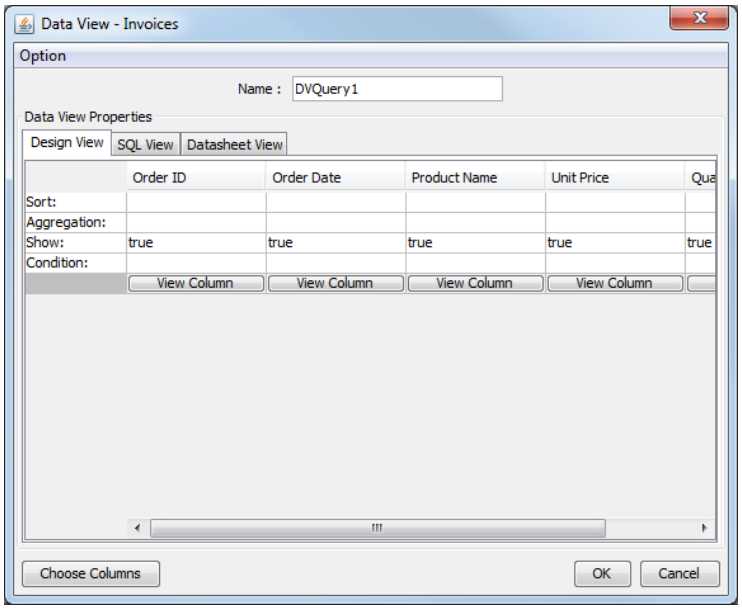
In the Data Source Manager, select the *Invoices* data view. Then click the *View* button. This will open a dialog, prompting you to select fields from the view. To select fields, first double click on a heading to expand it. Add the following fields to the query by double clicking or selecting them and clicking on the *ADD* button:

```
Order ID
Order Date
Product Name
Unit Price
Quantity
```



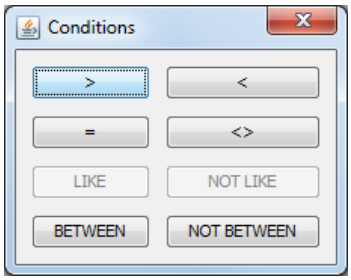
Data View Query Field Selection Dialog

Once you have finished adding the fields, click *OK*. This will bring up a new window allowing you to set conditions, grouping, and ordering for the query. Like the Query Builder this window also allows you to preview the query result with the *Datasheet View* tab.



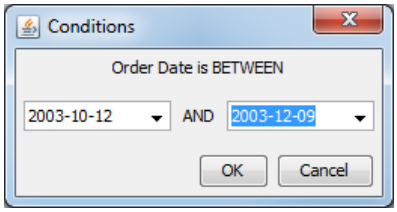
Data View Conditions Dialog

First specify a name for the query in the space provided at the top. Then double click on the *Condition* field for the *Order Date* column. This will bring up a dialog allowing you to specify a condition for the field.



Specify Condition Dialog

Click the *Between* button. A new dialog will appear prompting you to specify a start and end date with which to filter the results. Select **`2003-10-12`** as the first date and **`2003-12-09`** as the second.



Specify Condition Fields Dialog

Click *OK* to close the dialog and add the condition. You will be taken back to the conditions window. Now you can click on the *Datasheet View* tab to preview the query. Click *OK* in the main window and the query will be saved using the name you provided. You will then be taken to a screen showing the result of the data view query. Options in this screen allow you to use the results to create a report or chart. Click *Cancel* in this dialog to close the interface and return to the Data Source Manager. There will now be a new node for your query under the *Invoices* data view.

Q.4. Report Designer

After setting up data sources and creating queries, the next step to create reports is to take the results of the data source and map them to the report. Depending on the type of report being created, the mapping options will vary. This section only covers basic mapping for certain report layout. Once you will finish data mapping for your report, you will be taken to the Report Designer interface where you can customize many different features/properties of the report. Here in QuickStart, we will use only some of the most commonly used features in the Report Designer. For more information, please see Section 5.1 - Report Designer.

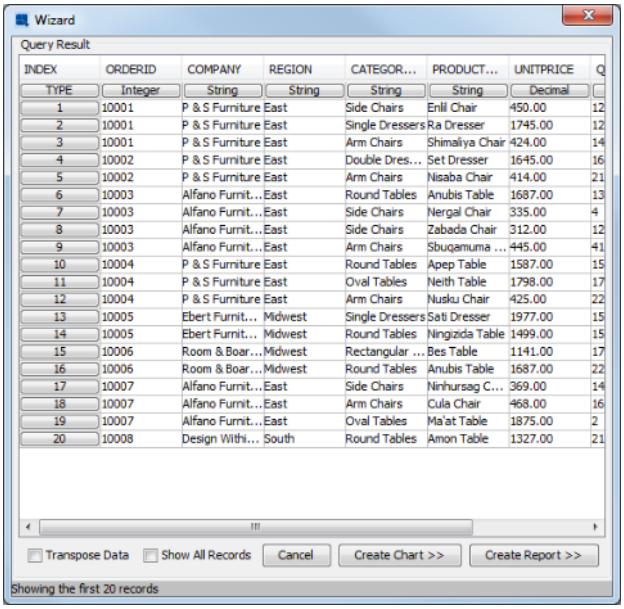
Q.4.1. Report Mapping

This section looks at several ways data can be mapped to the table structure of reports. This section will use the query created in Section Q.3.1.1 - Create a Query as the data source for the report.

Q.4.1.1. Simple Columnar Layout

The simple columnar layout is the most straight forward type of report mapping. Columns from the data source are drawn in a straight table in the report without any grouping and breaks.

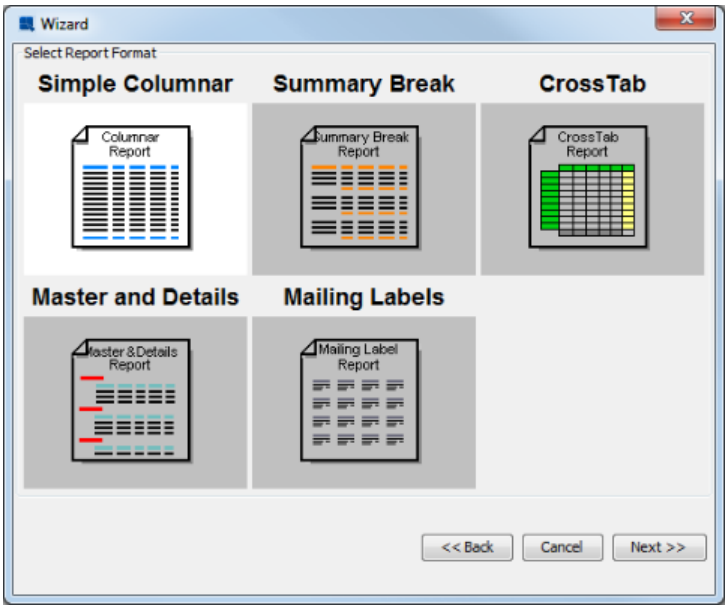
To begin creating a report, open your data registry and select the node for your query. Click the *VIEW* button. A new window will open, with a table containing the results of the query (first 20 records only). Notice that because your query contains a parameter, it initially runs with the default values that were specified when the parameters were initialized.



INDEX	ORDERID	COMPANY	REGION	CATEGOR...	PRODUCT...	UNITPRICE	Q
TYPE	Integer	String	String	String	String	Decimal	
1	10001	P & S Furniture	East	Side Chairs	Enli Chair	450.00	12
2	10001	P & S Furniture	East	Single Dressers	Ra Dresser	1745.00	12
3	10001	P & S Furniture	East	Arm Chairs	Shimaliya Chair	424.00	14
4	10002	P & S Furniture	East	Double Dres...	Set Dresser	1645.00	16
5	10002	P & S Furniture	East	Arm Chairs	Nisaba Chair	414.00	21
6	10003	Alfano Furnit...	East	Round Tables	Anubis Table	1687.00	13
7	10003	Alfano Furnit...	East	Side Chairs	Nergal Chair	335.00	4
8	10003	Alfano Furnit...	East	Side Chairs	Zabada Chair	312.00	12
9	10003	Alfano Furnit...	East	Arm Chairs	Sbuqamuma ...	445.00	41
10	10004	P & S Furniture	East	Round Tables	Apep Table	1587.00	15
11	10004	P & S Furniture	East	Oval Tables	Neith Table	1798.00	17
12	10004	P & S Furniture	East	Arm Chairs	Nusku Chair	425.00	22
13	10005	Ebert Furnit...	Midwest	Single Dressers	Sati Dresser	1977.00	15
14	10005	Ebert Furnit...	Midwest	Round Tables	Ningizida Table	1499.00	15
15	10006	Room & Boar...	Midwest	Rectangular ...	Bes Table	1141.00	17
16	10006	Room & Boar...	Midwest	Round Tables	Anubis Table	1687.00	22
17	10007	Alfano Furnit...	East	Side Chairs	Ninhursag C...	369.00	14
18	10007	Alfano Furnit...	East	Arm Chairs	Cula Chair	468.00	16
19	10007	Alfano Furnit...	East	Oval Tables	Ma'at Table	1875.00	2
20	10008	Design Withi...	South	Round Tables	Amon Table	1327.00	21

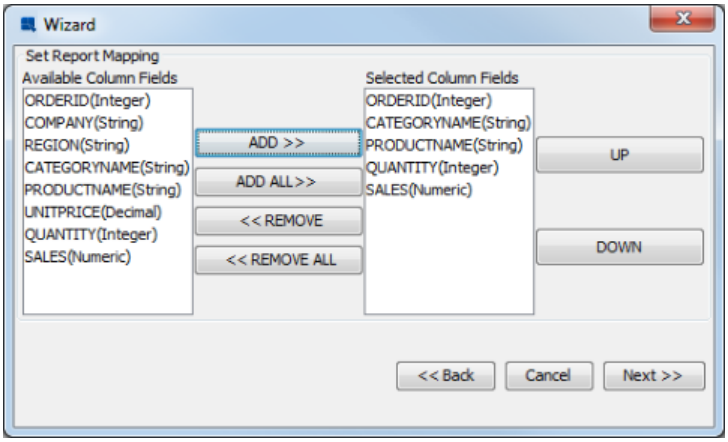
Query Result Screen

The bottom of the screen contains two buttons, *Create Chart* and *Create Report*, which allow you to use the query to design a chart or report. To continue designing the report, click *Create Report*. This will bring up a dialog asking you which report layout option you would like to use.



Select Report Layout Dialog

From this dialog, select *Simple Columnar* as the layout and click the *Next* button. The next step in the Report Wizard allows you to select which columns from the query you would like to use in the report, as well as re-arrange the column order. (Although the column order is not particularly important for the simple columnar layout, it can have a significant impact on other layouts depending on the mapping options selected).

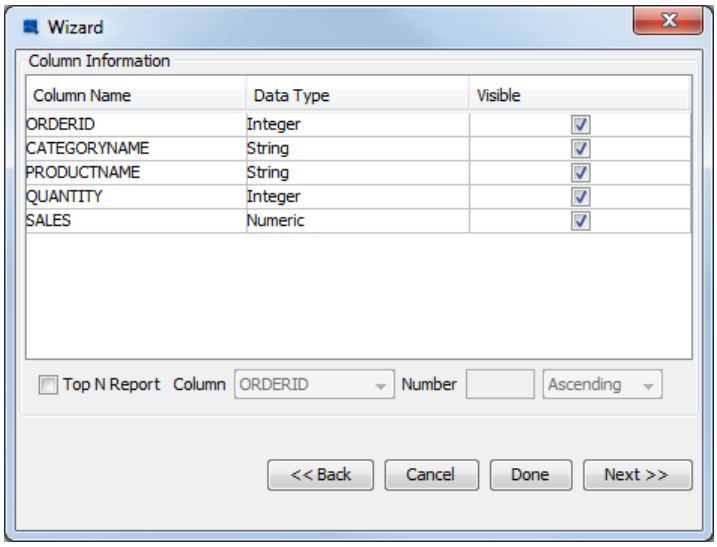


Column Selection/Ordering Dialog

In this dialog, select the following fields for the report:

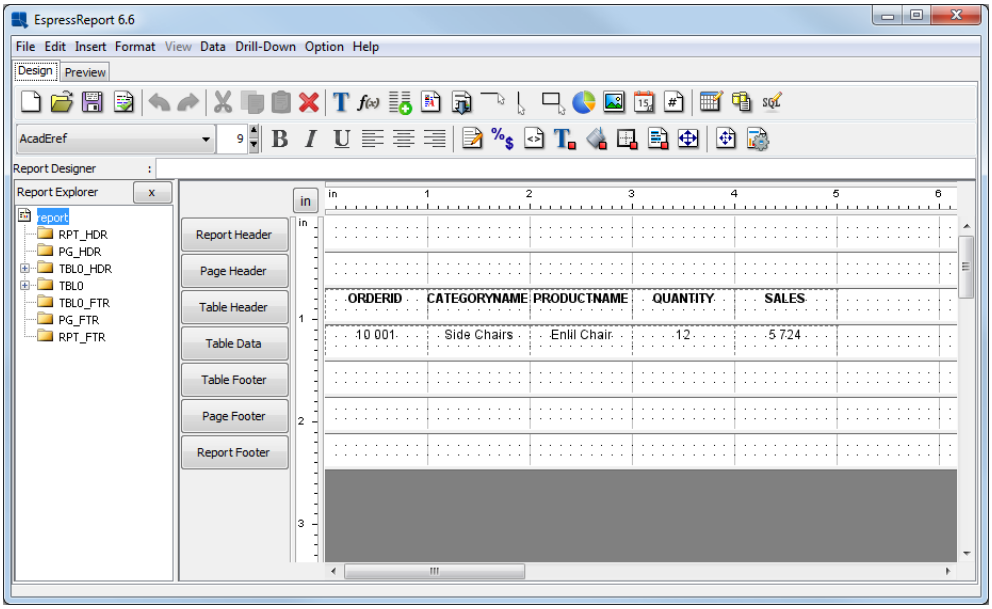
```
ORDERID
CATEGORYNAME
PRODUCTNAME
QUANTITY
SALES
```

Once you have selected the fields, click *Next* to continue on in the Wizard. The next dialog is the data mapping dialog. This is where you can select how to map the fields from the data source into the selected report layout. This is a simple columnar layout, so the only options are whether to set columns as visible or not and whether to generate a top N presentation.



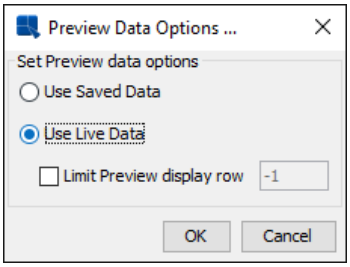
Data Mapping Dialog

In this dialog, select to leave all the columns visible and do not select the top N option. (For more about top N presentations, please see Section 5.1.2.1.1.1 - Top N Report). Next click the *Done* button. (There are some optional additional steps in the Wizard which will be explained later in this section). This will bring up the Report Designer interface with an unformatted version of your report.



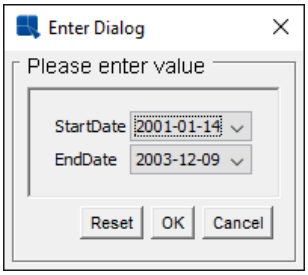
Simple Columnar Report in Design Window

Now click on the *Preview* Tab in the upper-left hand corner of the window. A dialog will open asking you whether you would like to preview the report with live or saved data.



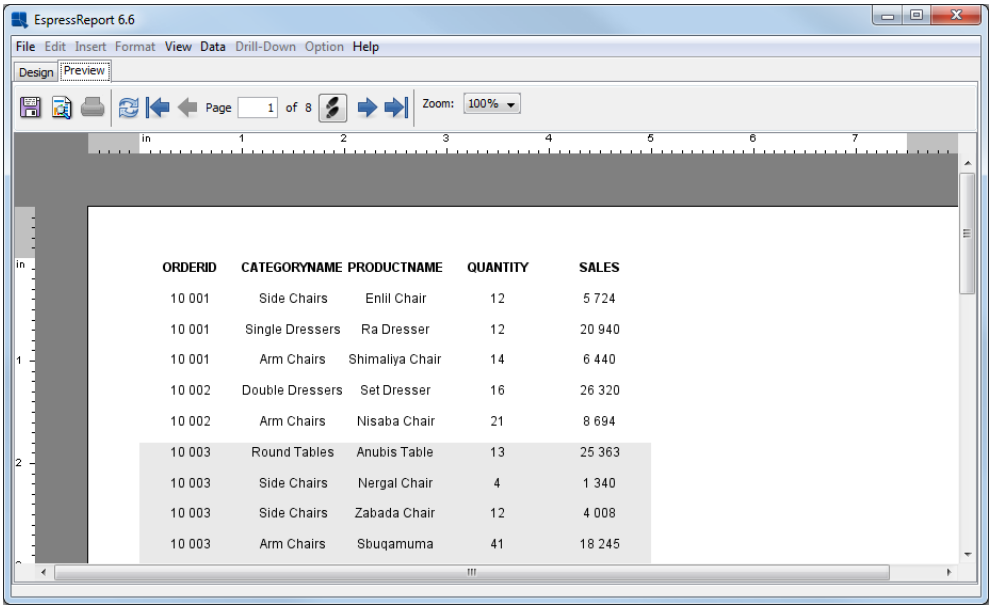
Preview Data Options

Select the *Live Data* option and click *OK*. A parameter selection dialog will appear prompting you to select a start and end date by which to filter the report.



Parameter Value Selection

Specify a range that is large enough to produce enough records and click *OK*. You will then see the report output. Notice how the simple columnar layout places the columns from the result set in the report directly, without any grouping, sorting, or summaries.



ORDERID	CATEGORYNAME	PRODUCTNAME	QUANTITY	SALES
10 001	Side Chairs	Enlil Chair	12	5 724
10 001	Single Dressers	Ra Dresser	12	20 940
10 001	Arm Chairs	Shimaliya Chair	14	6 440
10 002	Double Dressers	Set Dresser	16	26 320
10 002	Arm Chairs	Nisaba Chair	21	8 694
10 003	Round Tables	Anubis Table	13	25 363
10 003	Side Chairs	Nergal Chair	4	1 340
10 003	Side Chairs	Zabada Chair	12	4 008
10 003	Arm Chairs	Sbuqamuma	41	18 245

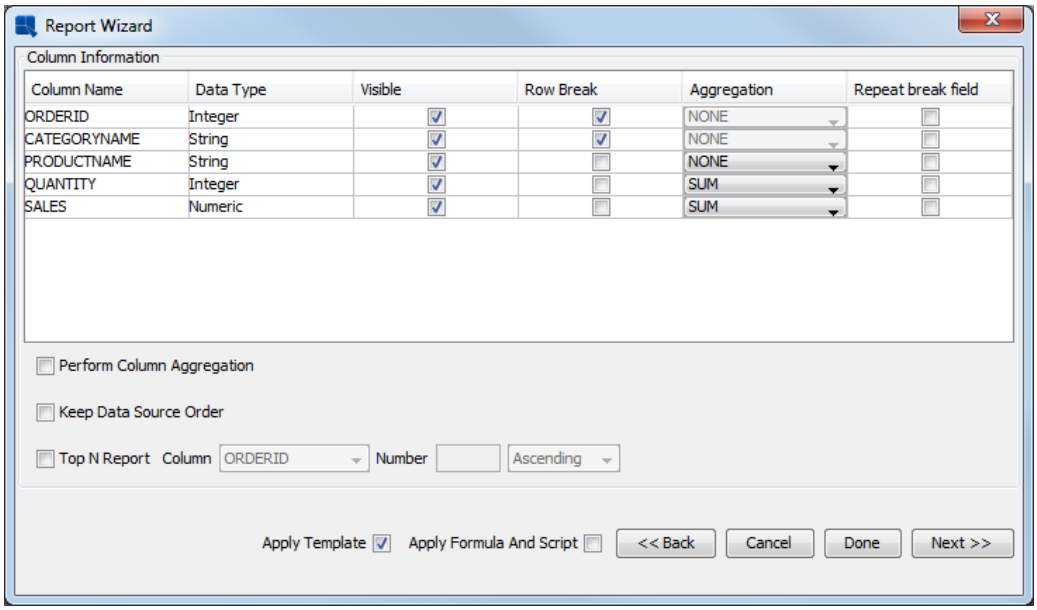
Simple Columnar Report Preview

Q.4.1.2. Summary Break Layout

The summary break layout is similar to the columnar layout, except it adds the ability to group and aggregate the report columns. A summary break report must be grouped by at least one column. Using the report created in Section Q.4.1.1 - Simple Columnar Layout, we will convert it into a summary break layout.

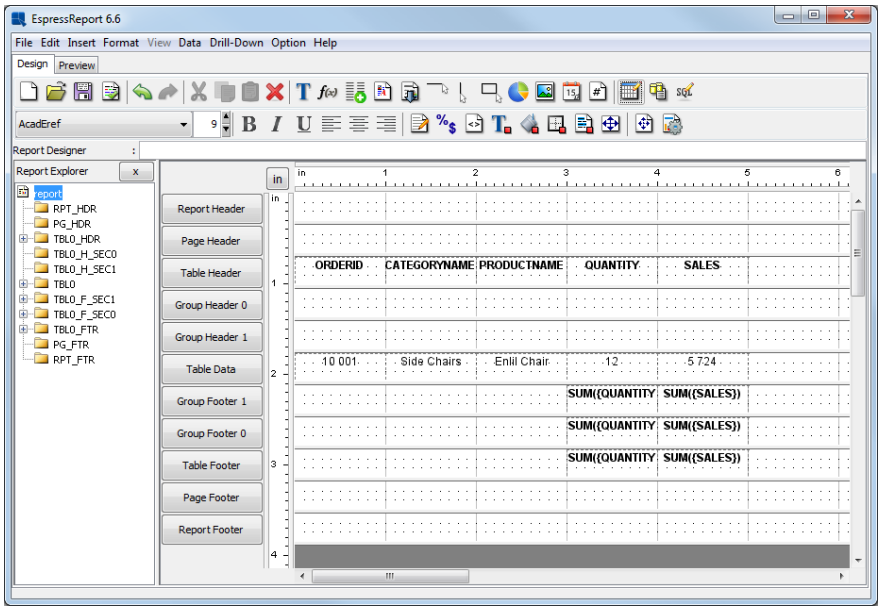
Go to the design tab and select the *Change Data Mapping* Icon on the toolbar  . This will return you to the Report Wizard, where we will select a different report layout. From the last data mapping dialog, hit the *Back* button twice until you get back to the *Select Report Format* window.

Change the report layout type to *Summary Break* and click *Next*. In the next screen, keep the column selection the same and click *Next* again to go to the data mapping window. You will notice that there are more options in this window than there were for the columnar layout.



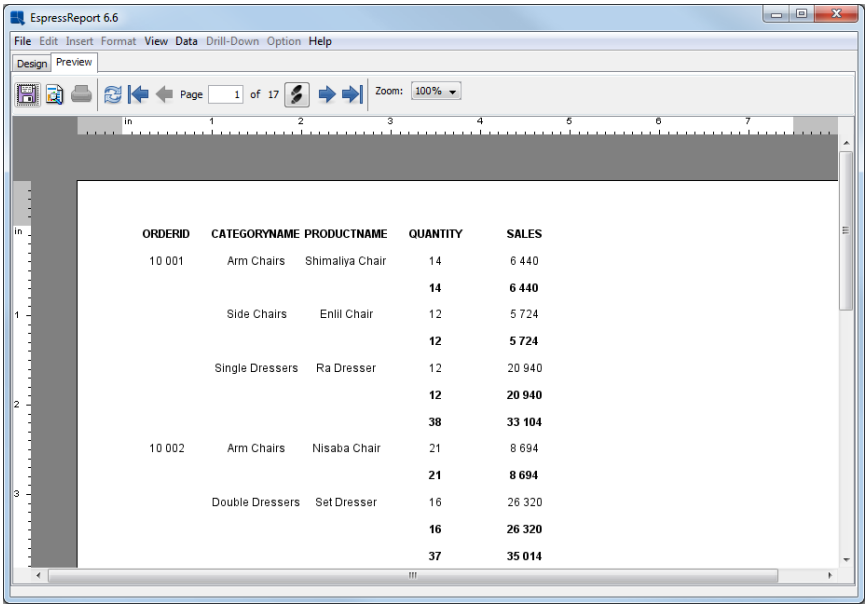
Data Mapping Screen for Summary Break Layout

In the data mapping dialog, check the option called *Row Break* for the first two columns. This will have the report group by those two columns. From the drop-down menus under *Aggregation*, select to *SUM* the *QUANTITY* and *SALES* columns. Also, uncheck the *Apply Template* option, as you do not need to carry over the formatting from the simple columnar layout. Once you have finished specifying options, again click the *Done* button. You will get a warning if you have not applied template. Click *Yes* to continue. You will be taken back to the Report Designer where the new mapping has taken effect.



Summary Break Report in Design Window

In the Designer you will notice that because there are two levels of nested grouping in the report there are now corresponding Group Header and Footer sections for each. For more about report sections and their behavior, please see Section 5.1.3.1 - Report Sections. Now click on the *Preview* tab to preview the report. Again, you will be prompted to specify parameter values. Once you see the report in the preview window, notice that the data is grouped by category name and order ID and that intermediate summaries are calculated for each group.



ORDERID	CATEGORYNAME	PRODUCTNAME	QUANTITY	SALES
10 001	Arm Chairs	Shimaliya Chair	14	6 440
			14	**6 440**
	Side Chairs	Enlil Chair	12	5 724
			12	**5 724**
	Single Dressers	Ra Dresser	12	20 940
			12	**20 940**
			38	**33 104**
10 002	Arm Chairs	Nisaba Chair	21	8 694
			21	**8 694**
	Double Dressers	Set Dresser	16	26 320
			16	**26 320**
			37	**35 014**

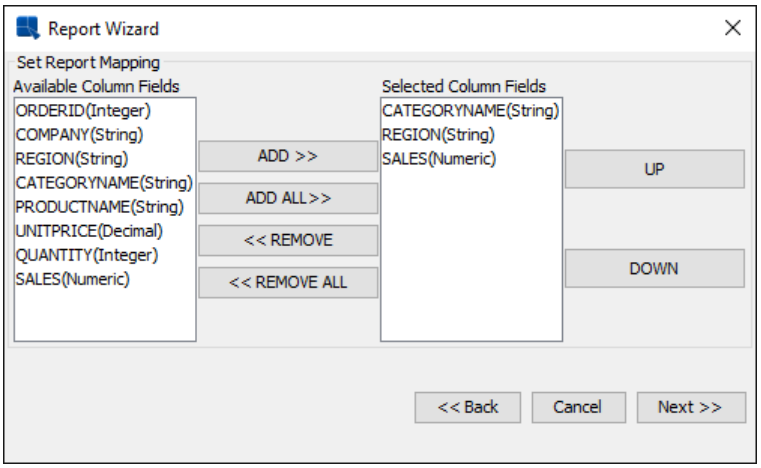
Summary Break Report Preview

Q.4.1.3. Crosstab Layout

A crosstab report shows data in a matrix-like form, allowing multi-dimensional data to be displayed in a two-dimensional layout. In the current example, we will use the crosstab layout to breakdown the sales column by product category and region.

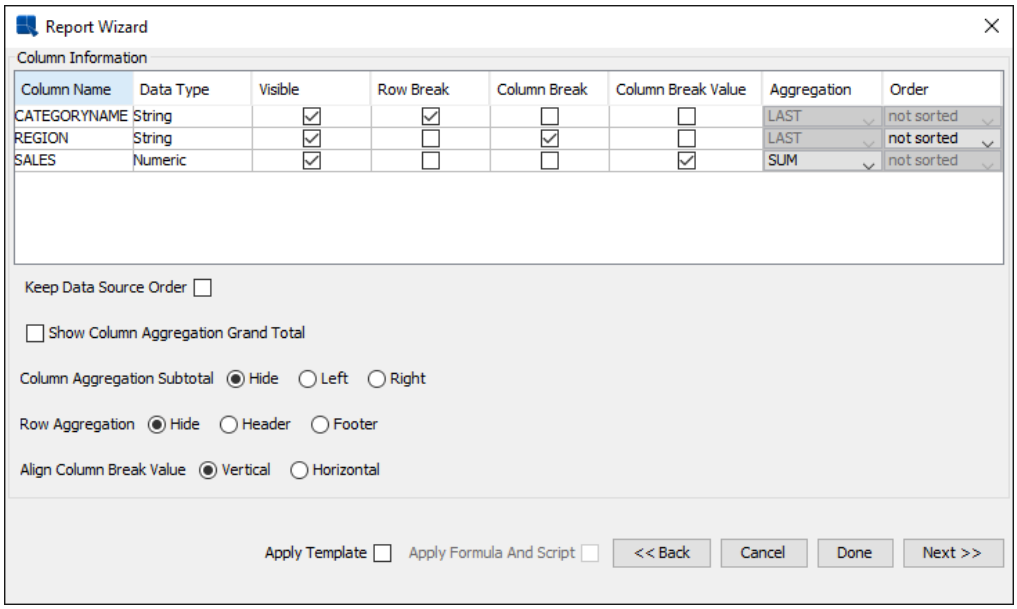
Go to the design window and select *Change Data Mapping*  . When the Report wizard re-opens, click the *Back* button until you get back to the layout selection screen. From this screen, select to use a *CrossTab* layout and click *Next*. At the column selection/ordering screen (next in the Wizard), change the selection to be the following:

```
CATEGORYNAME
REGION
SALES
```



Column Selection for Crosstab Report

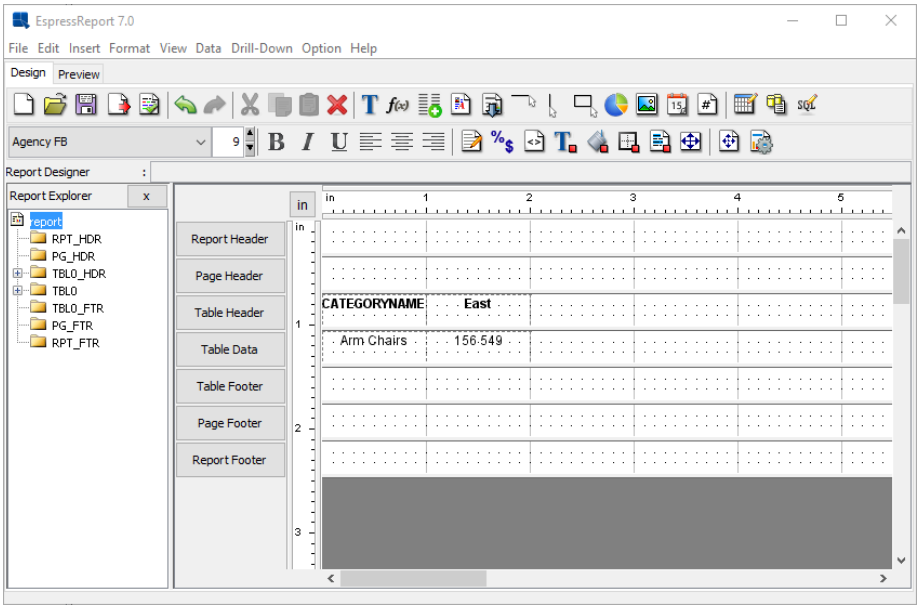
Once you have specified the columns in the correct order, click the *Next* button to bring up the data mapping option for the crosstab layout.



Data Mapping Screen for Crosstab Layout

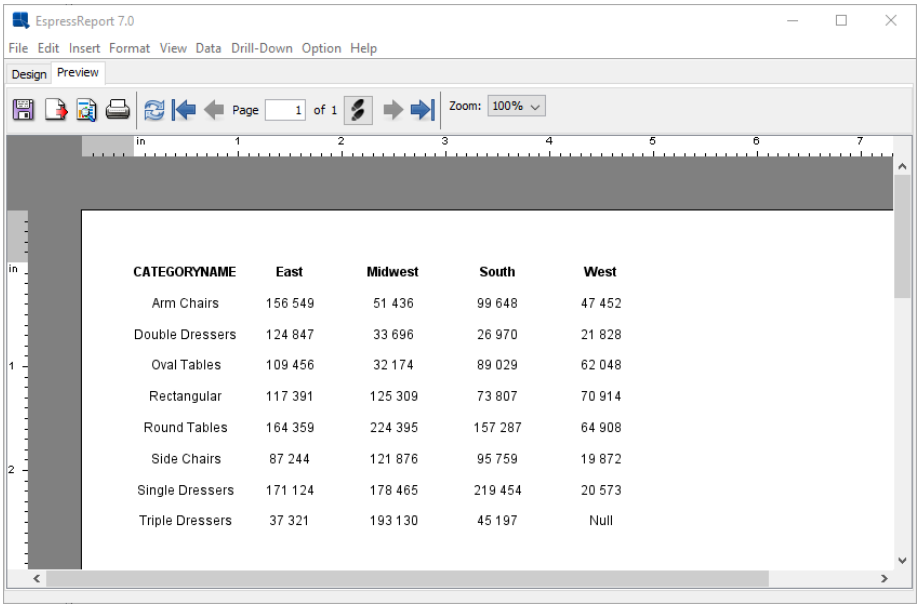
For the `CATEGORYNAME` column, check the option marked *Row Break*. This will create a row in the report data for each distinct category name. Next select the *Column Break* option for the `REGION` column. This will create a column in the report data for each distinct region. Leave the *Order* option as *not sorted*. Finally, select the *Column Break Value* option for the `SALES` column and click on the *Aggregation* menu to select *SUM*. This will give you the total sales for each category and region in the report.

Once you have finished specifying the options, click the *Done* button to go to the Report Designer. You will get a warning if no template was applied. Click *Yes* to continue.



Crosstab Report in Design Window

Now click on the *Preview* tab to preview the crosstab layout. As you can see a report column has been generated for each region, and has been automatically totaled both vertically (columns) and horizontally (rows).



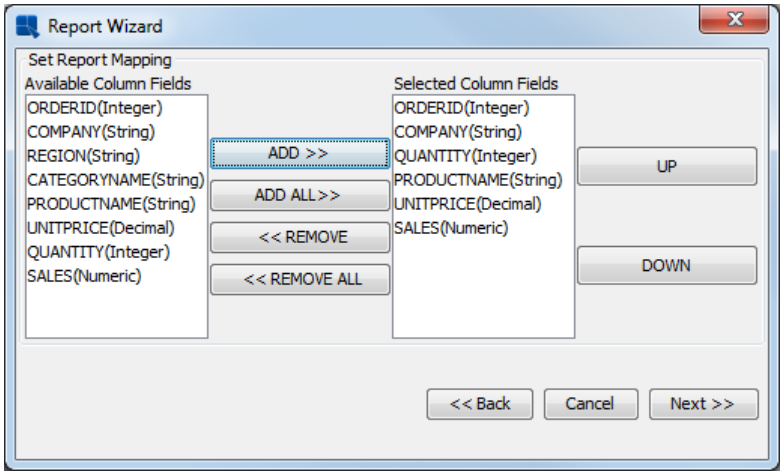
Crosstab Report Preview

Q.4.1.4. Master & Details Layout

Like the summary break layout, the master & details layout also allows you to group the data. It also allows you to automatically add column fields to the Group Header section, creating a one to many layout that can also be configured in a side-by-side layout.

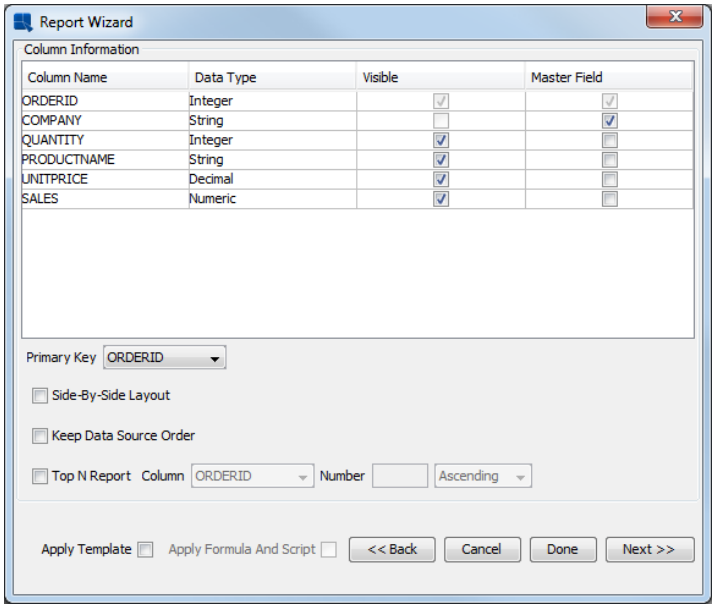
Go back to the design window and select *Change Data Mapping* from the Data menu. Again, navigate back to the layout selection screen. This time select the *Master and Details* layout. Click *Next* to get to the column selection screen. In the column selection screen, change the selection to include the following columns:

```
ORDERID
COMPANY
QUANTITY
PRODUCTNAME
UNITPRICE
SALES
```



Column Selection for Master & Details Report

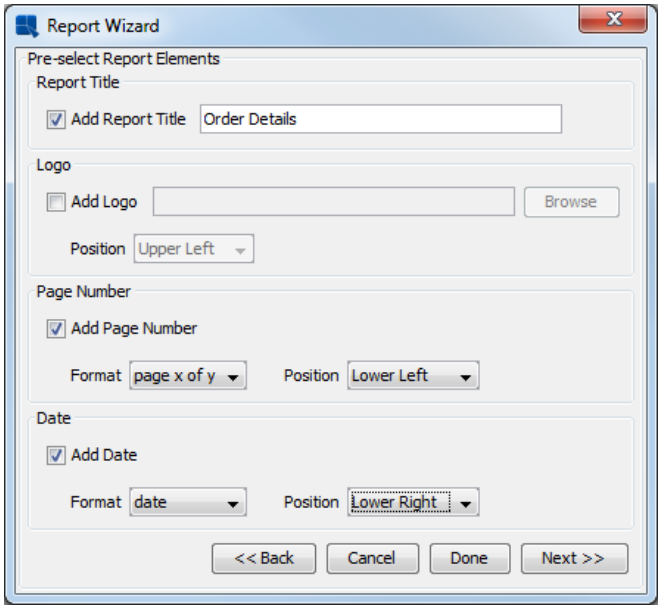
Once you have specified the columns in the correct order, click the *Next* button to bring up the data mapping dialog for the master & details layout.



Data Mapping Screen for Master & Details Layouts

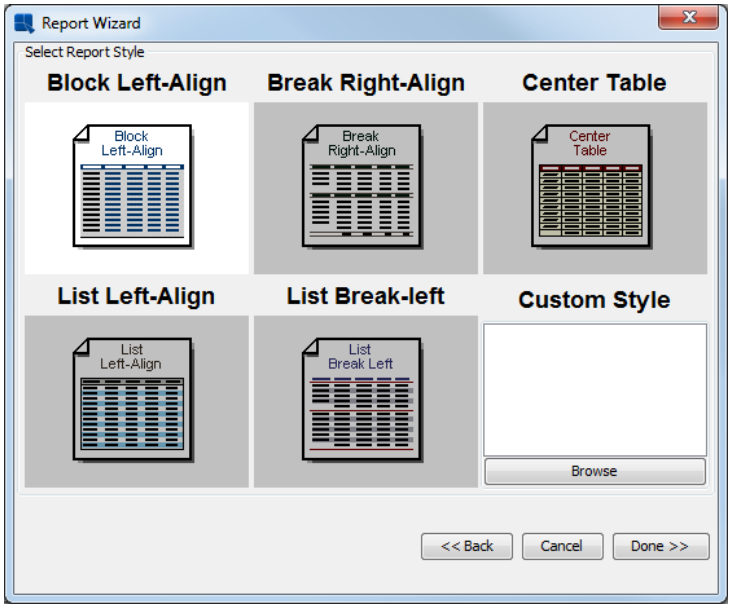
Select the `ORDERID` field as the *Primary Key* from the drop-down menu in the lower-left hand portion of the screen. This will group the data by the `ORDERID` field. Next, check the *Master Field* option for the `COMPANY` column. This will place the `COMPANY` field in the Group Header section of the report.

Instead of clicking *Done* to get to the Report Designer from this screen, click the *Next* button to invoke the additional pre-formatting options. You will get warning that all the formating will be lost. Click *Yes* to continue to the dialog that allows you to add several elements to the report.



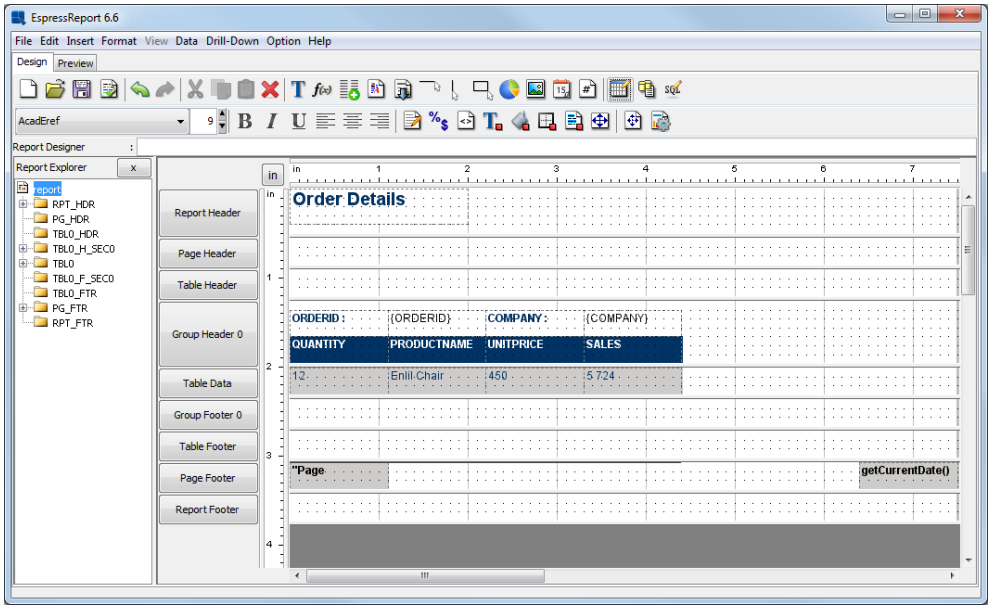
Add Report Elements Dialog

Check the boxes to add a *Report Title*, *Page Number* and *Date*. Enter the text you would like as the report title and specify the format and location of the date and time using the drop-down boxes for each option. Once you have finished setting the options, click *Next*. A new dialog will open allowing you to specify a style for the new report.



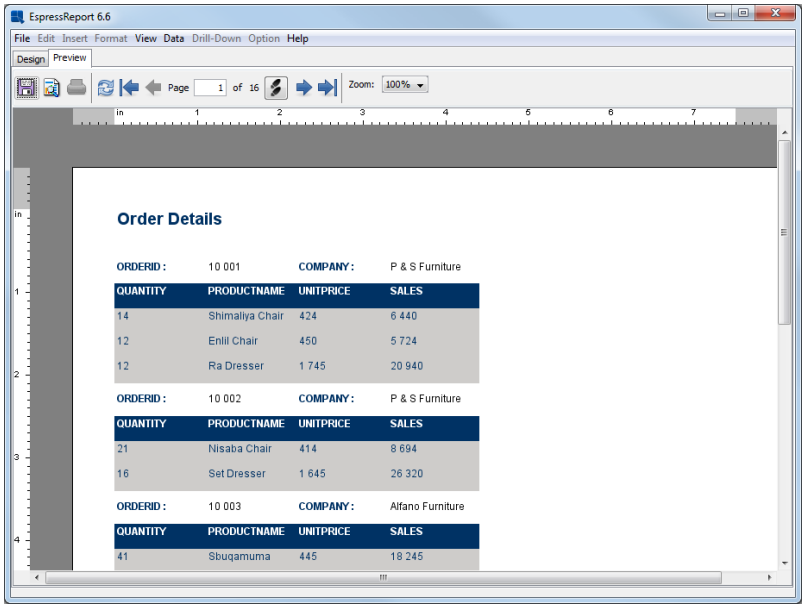
Report Style Selection

Select the *Block Left-Align* style and click the *Done* button. You will be taken to the Report Designer window, where the elements have been added. There is now some default formatting applied to the report.



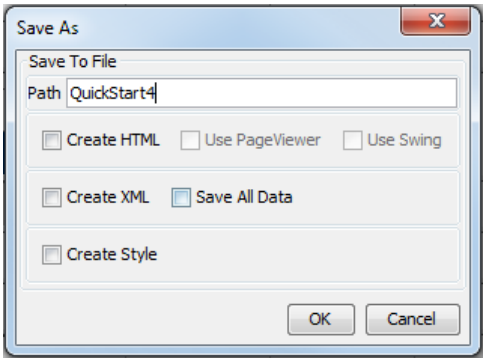
Master & Details Report in Design Window (with pre-formatting)

As you can see the `ORDERID` and `COMPANY` fields have been generated in the Group Header section. Now preview the report, and you will see a group with each order, as well as the page number and date in the section (header or footer) in which they were placed.



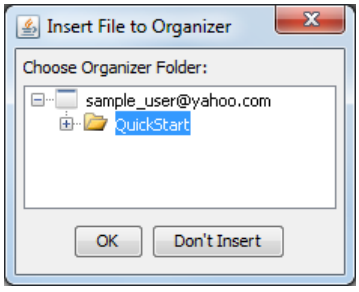
Master & Details Report in Preview Window

Now that you have finished creating a report, click the *Save* button on the toolbar to save the report  . This will bring up a new dialog, allowing you to specify save options for the report.



Save As Dialog

Enter a name for the report and click *OK*. The report will be saved in the `/your_email/Templates/` directory. Then, you will be prompted to add the file to the Organizer.



Insert File Dialog

Choose *QuickStart* folder. Click *OK* to add the file to the folder.

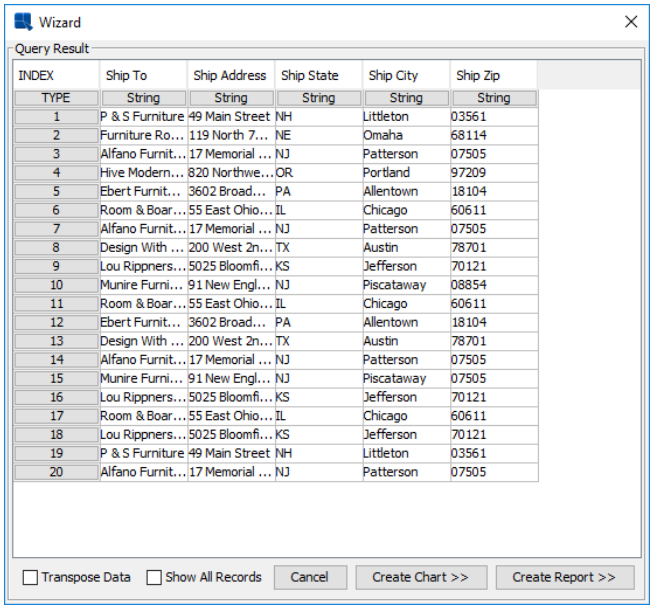
Q.4.1.5. Mailing Labels Layout

A mailing labels layout is similar to a simple columnar layout; however, it prearranges data in a manner that allows you to create mailing labels.

Close the design window and go back to the data registry to start designing a report. Navigate to *Invoicing* data view. Click the *View* button to create a new dataview query. Select the following fields from *shipping info* to the query:

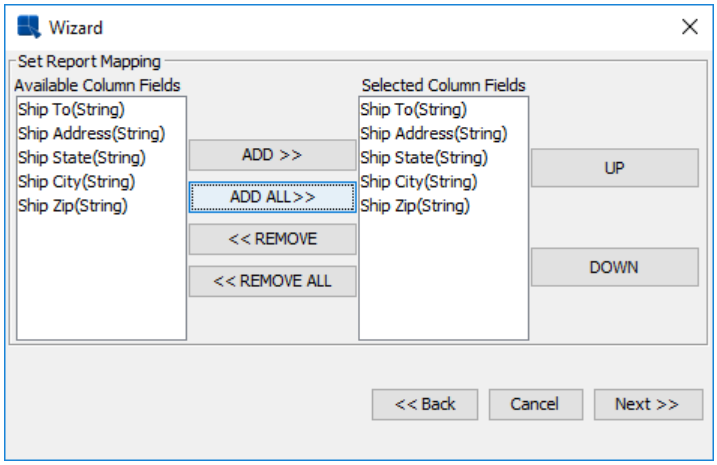
```
Ship To
Ship Address
Ship City
Ship State
Ship Zip
```

Specify a name for the query, then click *OK*. The query result show up:



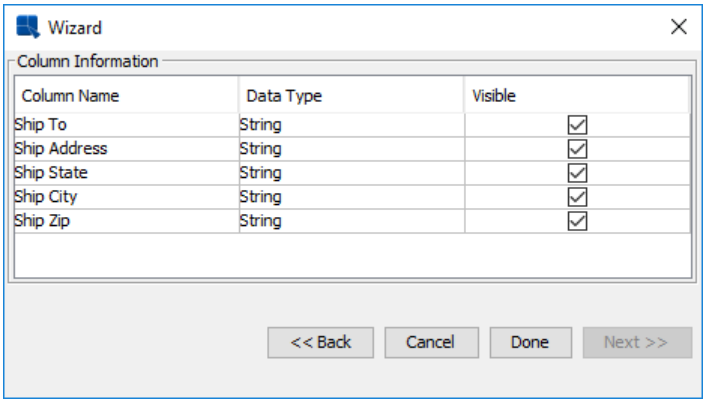
DataView Query Result for Creating Mailing Labels Report

click the *Create Report* button and select the *Mailing Labels* layout. The next screen is to select columns to the report, click *ADD ALL*:



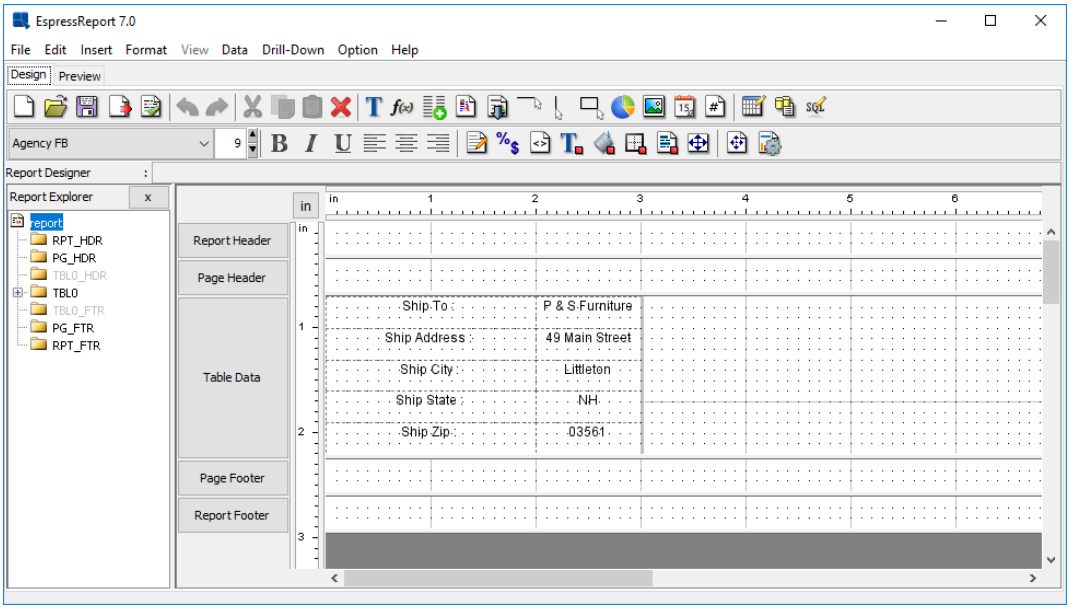
Column Selection for Mailing Labels Report

Once you have specified the columns in the correct order, click the *Next* button to bring up the data mapping dialog for the maling labels layout.



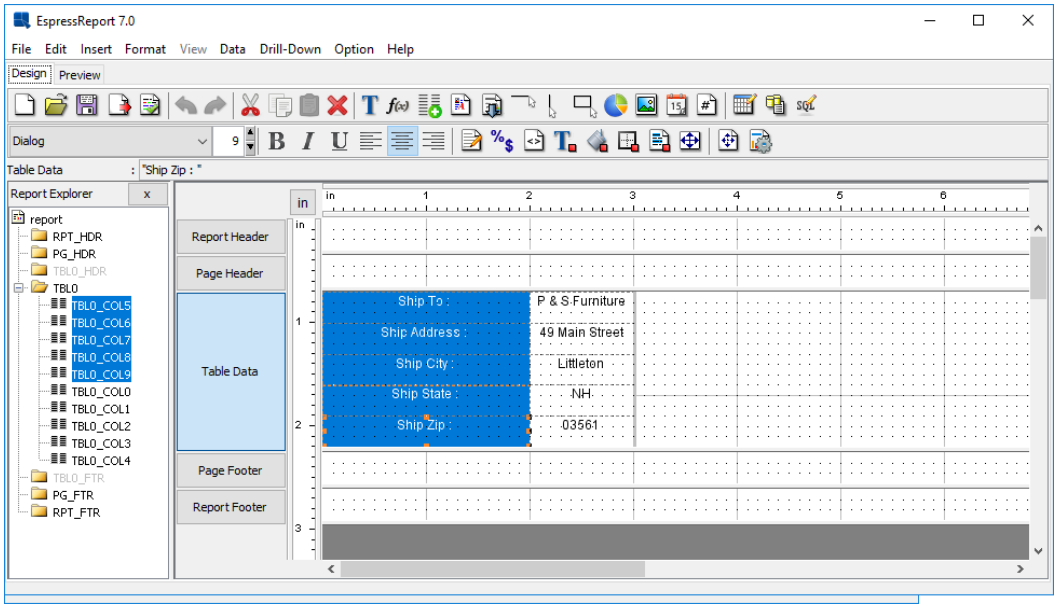
Data Mapping Screen for Mailing Labels Layouts

Click the*Done*, You will be taken to the Report Designer.



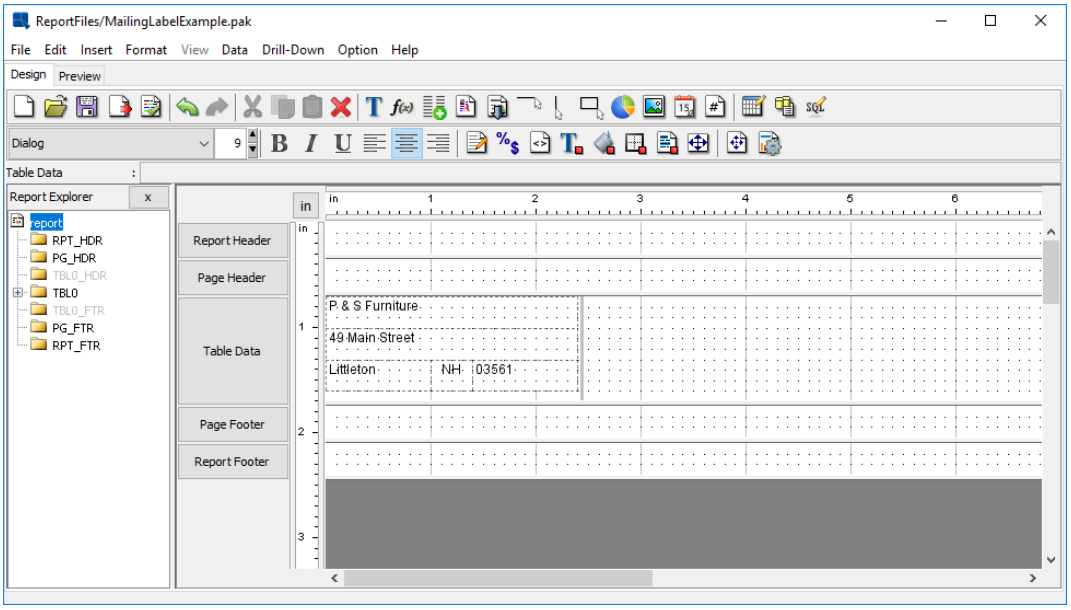
Mailing Labels Report in Design Window

We don't need the left columns in Mailing Label Report. We select all the left columns:



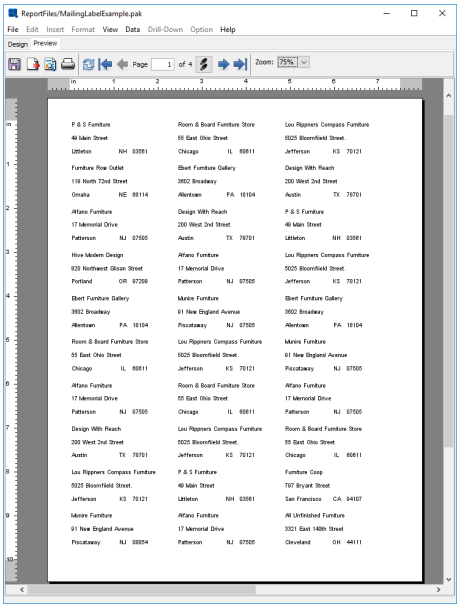
Mailing Labels Report Edit in Design Window

And remove them. We continue to edit the report to get general mailing label appearance:



Mailing Labels Report Appearance in Design Window

Now click on the *Preview* tab to preview the Mailing Labels layout.



Mailing Labels Report in Preview Window

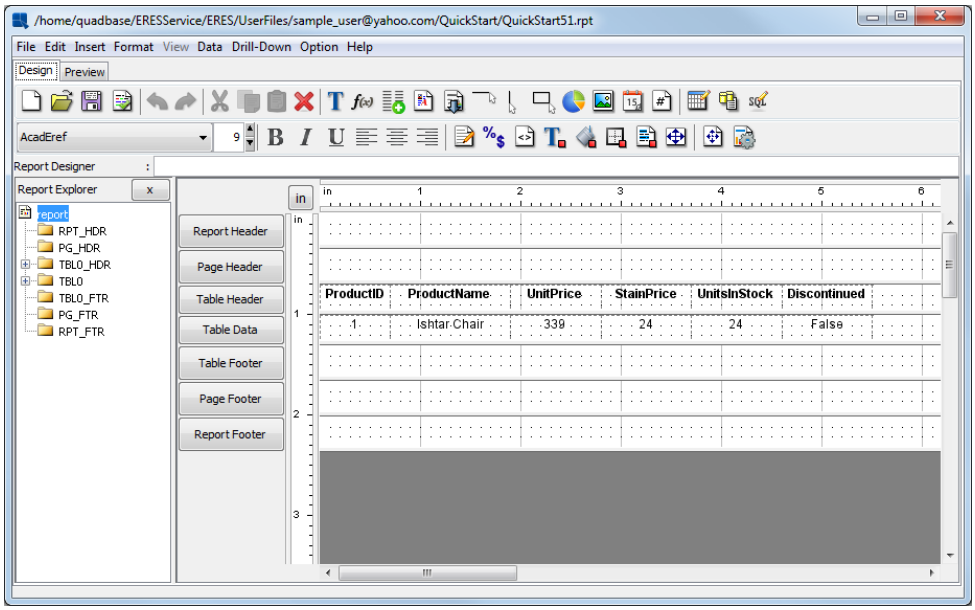
Now go to the Report Designer window, and select *Exit* from the *File* menu to close the interface. We want to save the report, so select *Yes* in the prompt and specify the name for the new-created Mailing Label Report.

Q.4.2. Basic Report Formatting

In this section, we will open an unformatted template and use some of the basic formatting features in EspressReport to create a polished presentation.

Q.4.2.1. Open a Report Template

If you do not have the Organizer open, open it. Then select QuickStart node under your project in the left-hand side. Double click the `QuickStart51.rpt` from the right-hand side to open it.

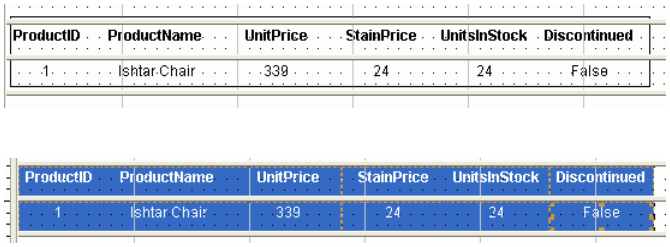


`QuickStart51.rpt` in the Report Designer

Q.4.2.2. Move and Align Report Elements

Report elements can be moved one by one, by click and dragging the cell. Report elements can also be moved as groups. To move a group of elements you must first select the group using a selection box. To activate the selection

box, left-click on the report and drag the mouse to draw a box around the report columns. They should become highlighted when you release the mouse button. To add some more elements to your current selection, press the **CTRL** key and click again, or draw another selection box while holding the **CTRL** key.

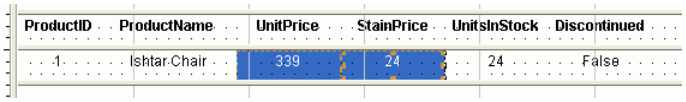


Drawing a Selection Box in Report Designer

Once the fields in the report have been selected, click and drag to indent them about half an inch. Next, click the *Left Alignment* button on the toolbar  . This will align all of the cell text to the left edge of the cell.

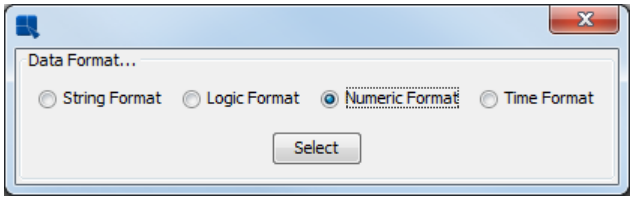
Q.4.2.3. Data Formatting

There are a number of options available that allow you to control how data will be displayed (date format, decimal places, rounding, etc). Using the selection box again, select the `UnitPrice`, and `StainPrice` columns (just the column fields not the headers).



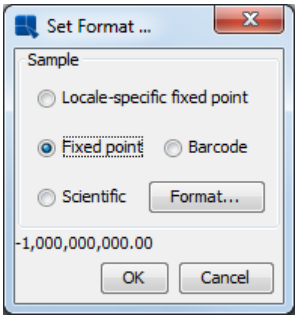
Two Column Selection

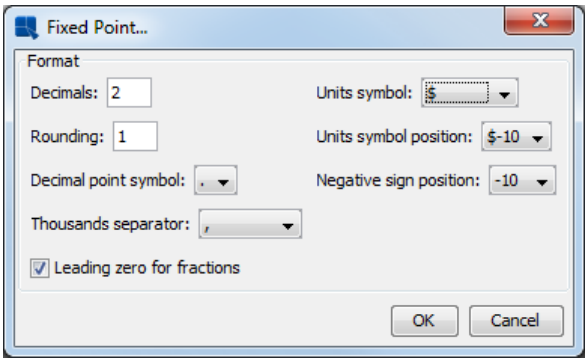
Once the columns have been selected, select the *Data Format* option from the *Format* menu. This will open a dialog prompting you to select which data type you would like to set the format for. Select *Numeric Format* and click *Select*.



Data Type for Formatting Dialog

This will open a new dialog. From this dialog select *Fixed Point* and click the *Format* button. At the next dialog select to have two decimal points and select the dollar sign as the *Units symbol*.





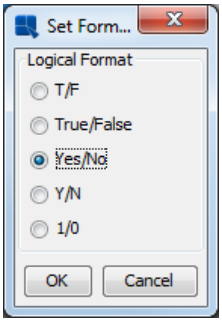
Numeric Format Dialogs

Click *OK* and *OK* again at the previous dialog. The selected fields will be converted to currency format.



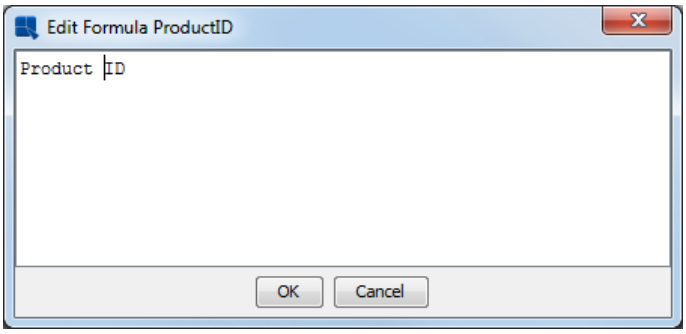
Numeric Data in Currency Format

Next, select the `Discontinued` column. The border outline will appear by clicking it. Again, click the *Data Format* button  . This time a dialog will appear allowing you to select a format for the Boolean column. Select *Yes/No* and click *OK*. The data in the column will now change.



Data Format Dialog for Boolean Data

Now we will edit the label text. By default column headers will display the column names from the database. However, you can override these headers with a custom one. Double click on the `ProductID` cell and a dialog will appear allowing you to modify the column header. Insert a space between `Product` and `ID` and click *OK*. The change will appear in the design window.



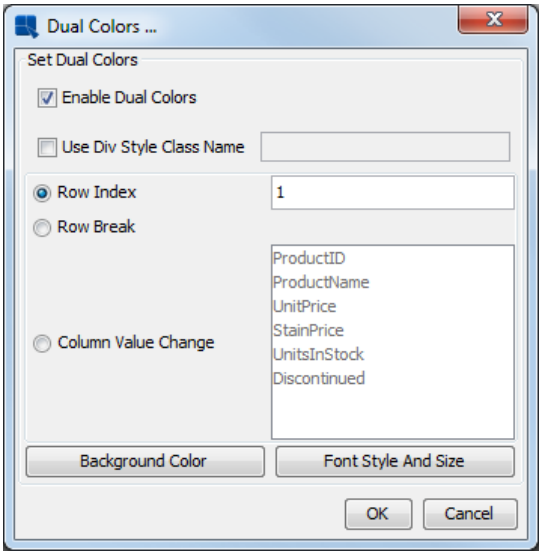
Edit Column Header Dialog

Repeat this for each column label except `Discontinued` so that you have inserted appropriate spaces in all of them.

Q.4.2.4. Set Dual Colors

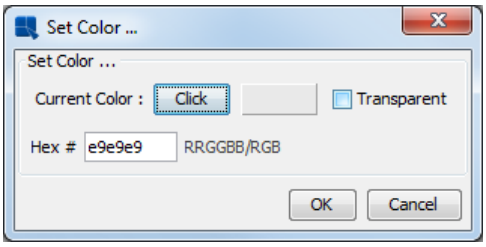
The dual colors feature in EspressReport allows users to differentiate different rows or groups of data by changing the background color, and/or font. Use the group selection tool to turn on dual colors and select all of the columns in the report (not headers).

Next, select the *Dual Colors* button on the toolbar  . A dialog will appear prompting you to set dual colors for the columns. Click the checkbox labeled *Enable Dual Colors*. Then select the *Row Index* radio button to indicate that you would like to change color based on a row value. Enter **`1`** for the row index value.



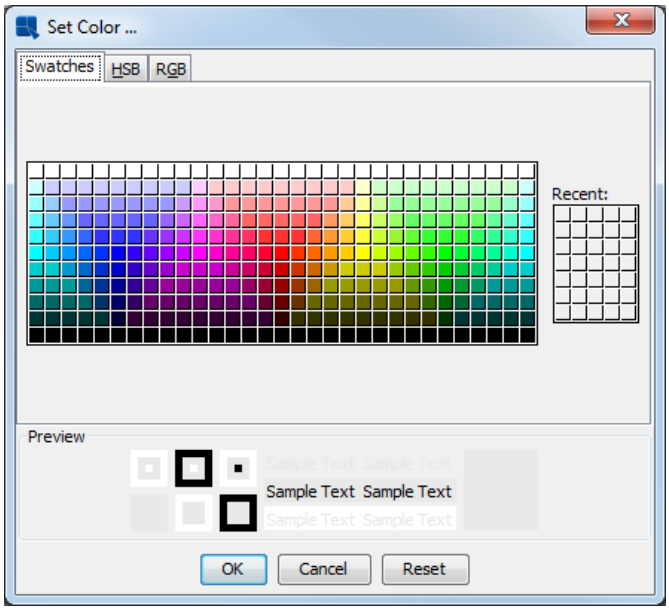
Dual Colors Dialog

Next, click the *Background Color* button. A dialog will appear giving you the option to set the background transparent and showing the current background color.



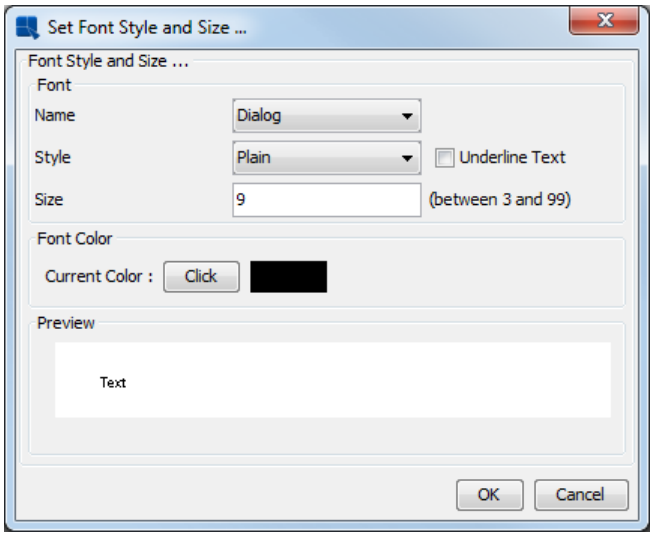
Background Transparency Dialog

Click the button labeled *Click* and a new dialog will appear with color swatches allowing you to select a new background color.



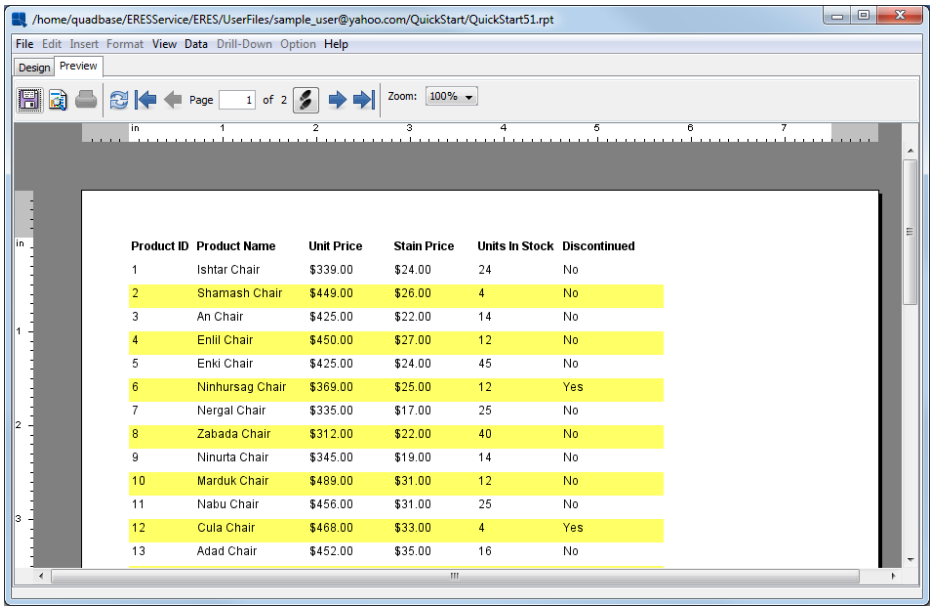
Choose Color Dialog

Select the new background color that you would like to use and click *OK*. You will return to the first dialog where the color selection will be reflected. Click the *OK* button again and you will return to the Dual Colors dialog. Once you are back at the Dual Colors dialog, click the *Font Style and Size* button. This will bring up a dialog allowing you to specify the font, font size, and font style for the alternating rows.



Font Style and Size Dialog

From this dialog, set the *Name* to **`Dialog`**, the *Style* to **`Plain`** and the *Size* to **`9`**. This will match the fonts for the alternating rows. Click *OK* to go back to the dual colors dialog and once more to return to the Report Designer. Now when you preview the report you will see alternating bands of color for each row.



Product ID	Product Name	Unit Price	Stain Price	Units In Stock	Discontinued
1	Ishtar Chair	$339.00	$24.00	24	No
2	Shamash Chair	$449.00	$26.00	4	No
3	An Chair	$425.00	$22.00	14	No
4	Enlil Chair	$450.00	$27.00	12	No
5	Enki Chair	$425.00	$24.00	45	No
6	Ninhursag Chair	$369.00	$25.00	12	Yes
7	Nergal Chair	$335.00	$17.00	25	No
8	Zabada Chair	$312.00	$22.00	40	No
9	Ninurta Chair	$345.00	$19.00	14	No
10	Marduk Chair	$489.00	$31.00	12	No
11	Nabu Chair	$456.00	$31.00	25	No
12	Cula Chair	$468.00	$33.00	4	Yes
13	Adad Chair	$452.00	$35.00	16	No

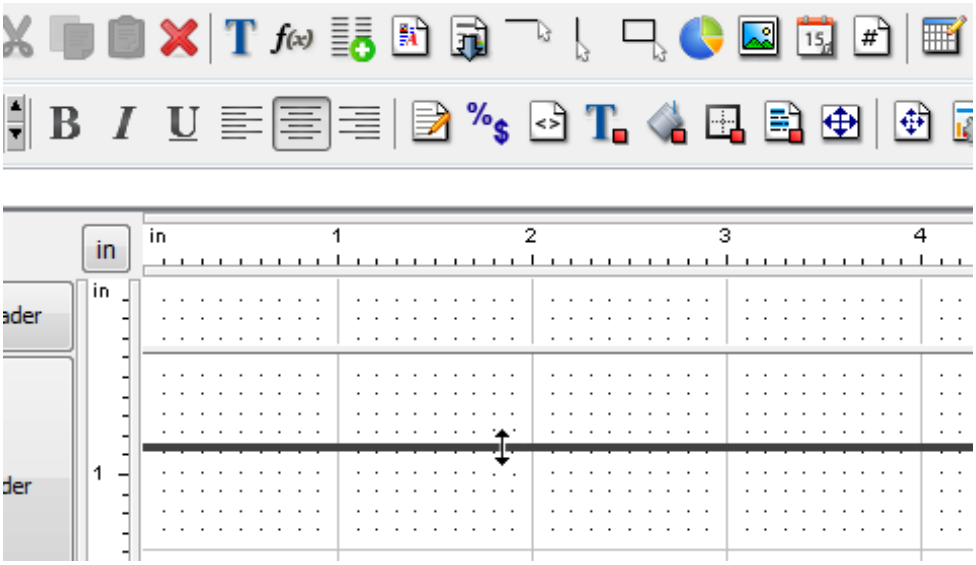
Dual Colors in Preview

Q.4.2.5. Inserting Elements

To further customize reports, many different types of elements can be added to a report template.

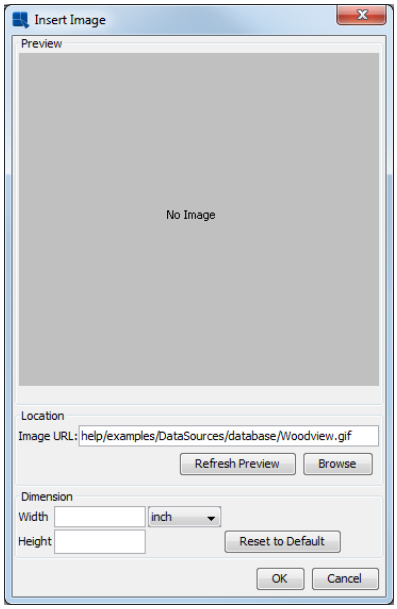
Q.4.2.5.1. Insert an Image

To insert an image into the report header, first you will want to resize the section to fit the image. In the design window, place the mouse over the lower section divider of the `Report Header` section, click and drag down making the section about an inch taller.



Re-Sizing a Report Section

Once the section is resized, click the *Image* button on the toolbar  . A small rectangle will follow your mouse pointer around the design window. Position the rectangle in the upper left-hand corner of the Report Header section and click. A dialog will appear prompting you to select an image to insert. Type in this Image URL: **`help/examples/DataSources/database/Woodview.gif`**.

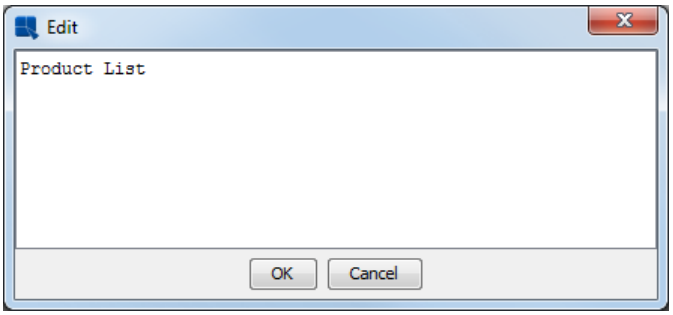


Insert Image Dialog

Click *OK* and the image will be inserted into the report. The image will be represented by a gray rectangle in the design view. You can see the image when you preview the report.

Q.4.2.5.2. Insert a Title

To insert a title, click the *Insert Label* button on the toolbar  . A small rectangle will now follow your mouse pointer around the design window. Position the rectangle next to the image you are inserted in and click. A dialog will appear prompting you to enter the label text. Type in the text of your desired title.



Insert Label Dialog

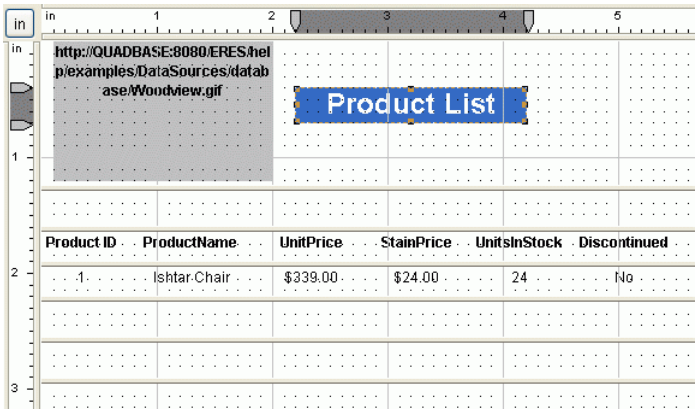
Click *OK* and the title will be added to the report. As you can see, the text is fairly small by default. To change this, change the font size dialog on the toolbar to be 18pt font.



Toolbar Font Options

When you do this you will notice that some of the text in the title cell has now disappeared. This is because the text is now larger than the space defined. To resize the cell, right click and drag until you can see the report title again.

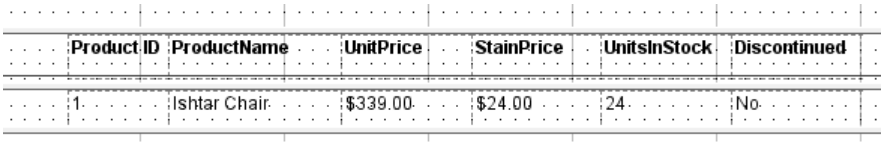
Then click the *Left Alignment* button on the toolbar to set the text alignment to the left as you did for the other report elements. Move and position the title cell so it is next to the inserted image.



Report Header with Image and Title

Q.4.2.5.3. Insert a Line

Next, we will add a horizontal line below the column headers. To do this, first resize the Table Header section slightly to provide some more space below the column headers. Next, select the *Insert Horizontal Line* button on the toolbar  . Your cursor will then change into a cross. Click below the `ProductID` header and drag across to the last column to draw the line.

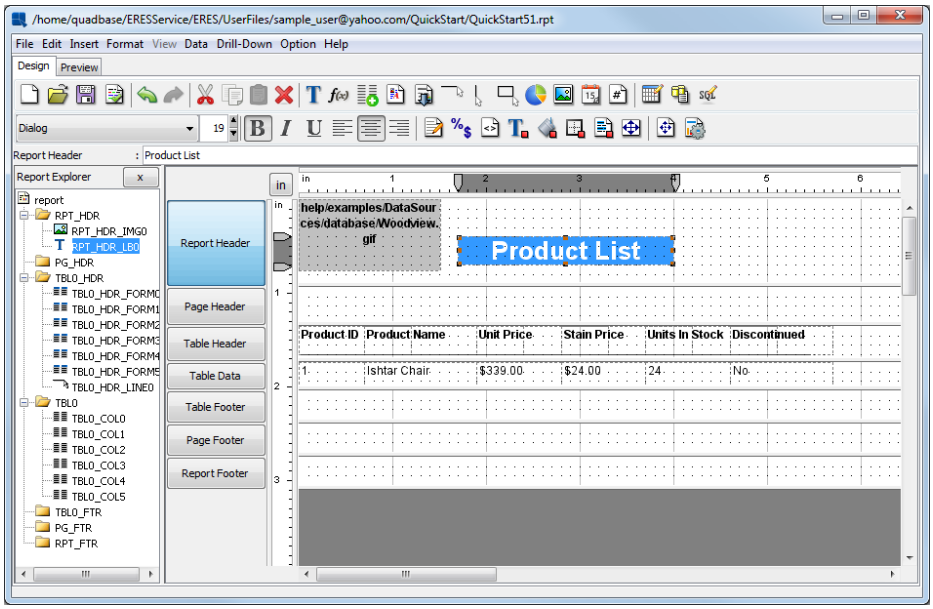


Drawing a Line in Report Designer

Q.4.2.6. View Report Elements in Report Explorer

Report Explorer from the *Option* menu is turned on by default. It is the panel on the left-hand side of the Report Designer, showing the report elements in a tree format. Click to expand some of the sections and you will see all

of the elements in the report represented in the tree. If you select one of the elements in the tree, the corresponding element in the report will be selected.



Designer with Explorer Open

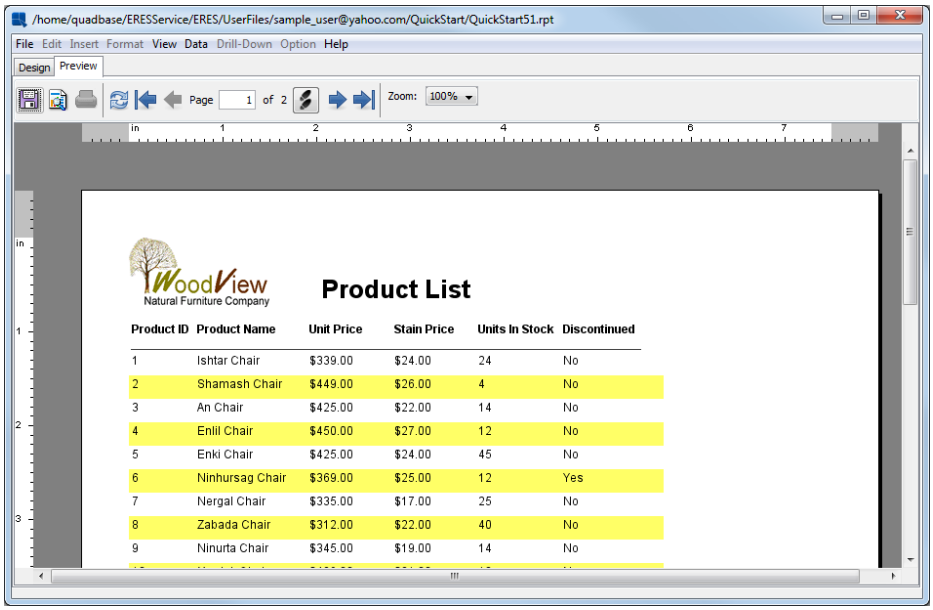
You can close the report explorer by selecting *Report Explorer* again from the *Option* menu.

Q.4.2.7. Set Section Options

In EspressReport, each of the report sections has a number of configurable options allowing you to display the data in the sections in any number of different ways. To invoke the options menu for a section, click the button for that section on the left-hand side of the design window. In this example, bring up the options menu for the *Table Header* section by clicking on the corresponding button.

From the pop-up menu select *Repeat On Every Page*. This will cause the Table Header to be drawn on each report page instead of only once (which is the default). For more about report section options, see Section 5.1.3.3 - Section Options.

Now that you have finished formatting the report, preview it to see the results.



Finished Report

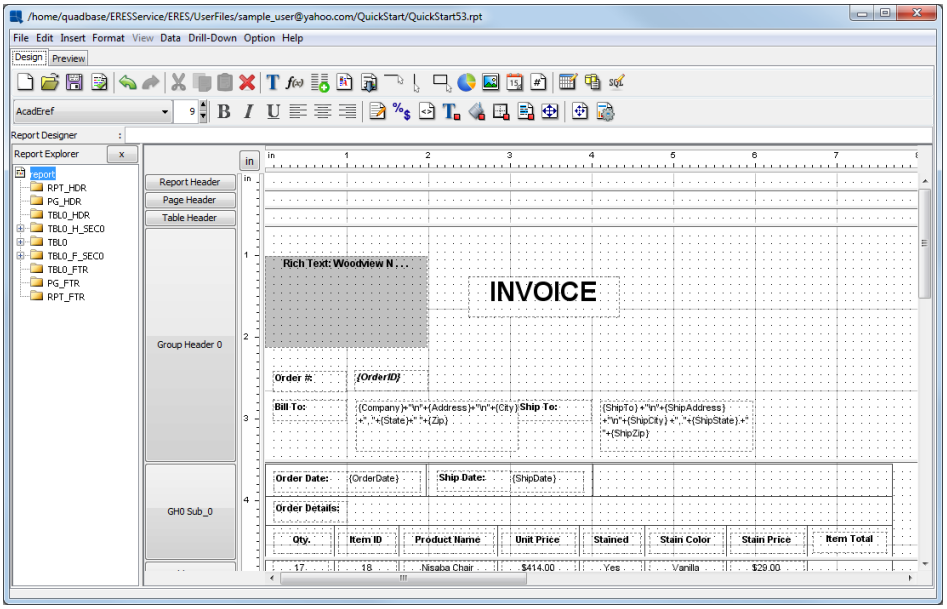
Now save the changes you have made to the report and exit from the Report Designer.

Q.4.3. Advanced Reporting Features

Q.4.3.1. Formulas & Scripting

EspressReport ES provides a large built-in formula and scripting library, giving you many ways to manipulate and analyze report data. In the following section we will take a template and use formulas and scripts to calculate/add some values to the report.

Select QuickStart node under your project in the left-hand side of the Organizer. Double click the `QuickStart53.rpt` from the right-hand side. This will open the report in Report Designer.

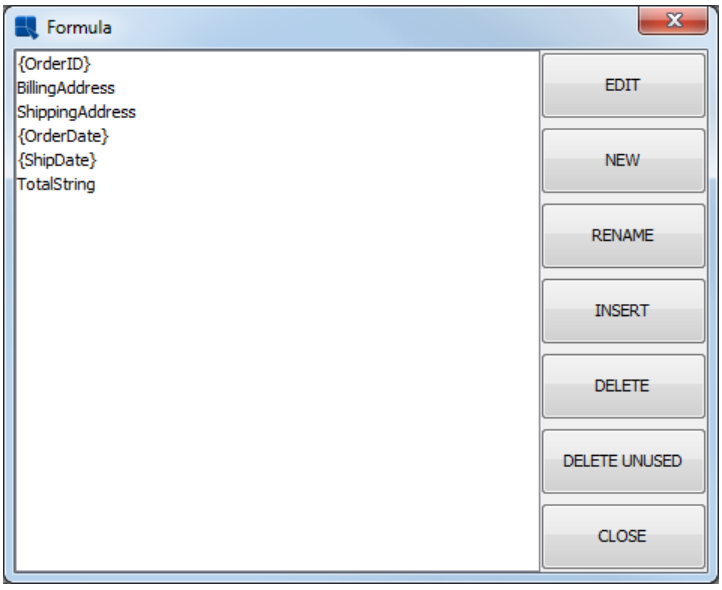


`QuickStart53.rpt` in Report Designer

The report is an invoice created using the Master & Details layout. Notice that the `Item Total`, as well as the `Sub-total` are blank. We will add in the formulas to calculate these values.

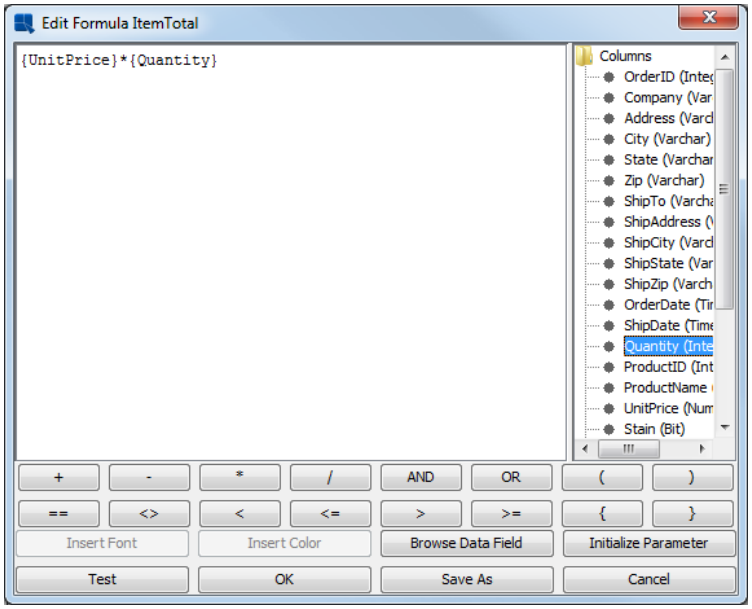
Q.4.3.2. Add a Formula

To insert a formula, click the *Insert Formula* button on the toolbar  . This will bring up a dialog containing all the formulas within the report. Notice that the template has several existing formulas. Click *New* to create a new formula. At the prompt, specify the name **`ItemTotal`** for the formula.



Report Formula List

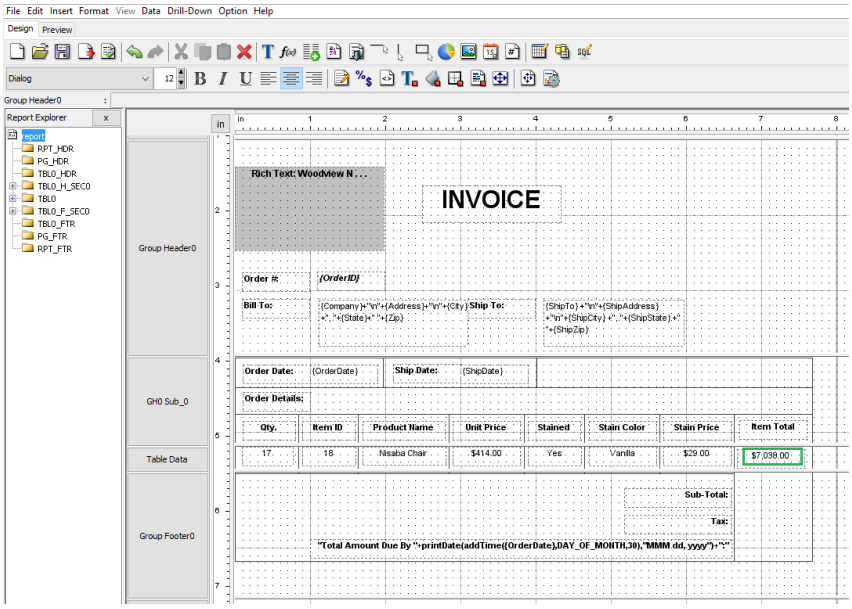
The Formula Builder window will then open. Double click on the *Columns* folder in the right-hand side to expand it. Then double click on the *UnitPrice* column to add it to the formula. Next, click the *multiply ("*")* button. Next, double click on the *Quantity* column to add it. The finished formula should look like this: `{UnitPrice}*{Quantity}`.



Formula Builder Window

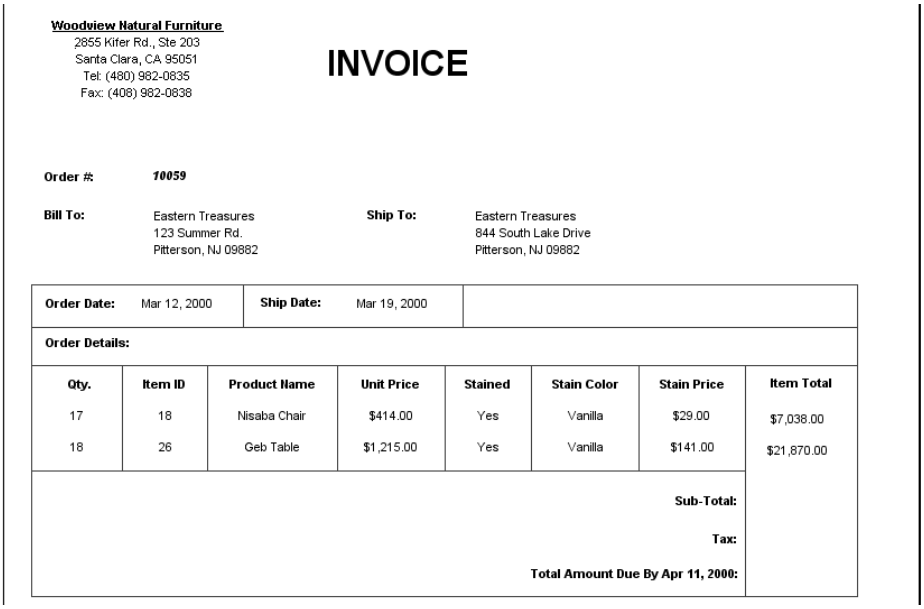
Click the *Test* button to ensure that the formula is entered correctly. Then click *OK*. You will be taken back to the formula list where your new formula has been added. From the formula list, select the *ItemTotal* formula that you just created and click the *Insert* button. The dialog will close and a small dotted rectangle will follow your pointer around the design window. Position the formula below the *Item Total* label and between the lines in the Table Data section and click. The formula will then be added to the report.

Next set the data format for the formula to be currency, like in Section Q.4.2.3 - Data Formatting.



Formula is Inserted

Now preview the report, notice that because the formula was added to the Table data section, it now computes for each row of data.



Woodview Natural Furniture
2855 Kifer Rd., Ste 203
Santa Clara, CA 95051
Tel: (480) 982-0835
Fax: (408) 982-0838

INVOICE

Order #: *10059*

Bill To: Eastern Treasures
123 Summer Rd.
Pitterson, NJ 09882

Ship To: Eastern Treasures
844 South Lake Drive
Pitterson, NJ 09882

Order Date: Mar 12, 2000 **Ship Date:** Mar 19, 2000

Order Details:

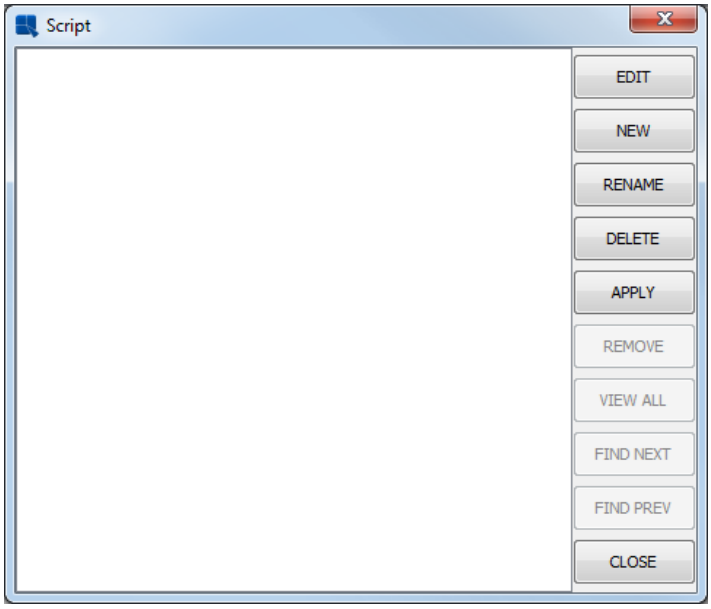
Qty.	Item ID	Product Name	Unit Price	Stained	Stain Color	Stain Price	Item Total
17	18	Nisaba Chair	$414.00	Yes	Vanilla	$29.00	$7,038.00
18	26	Geb Table	$1,215.00	Yes	Vanilla	$141.00	$21,870.00
						Sub-Total:	
						Tax:	
						Total Amount Due By Apr 11, 2000:	

Report with Formula

Q.4.3.3. Add a Script

Notice that the formula you added in the previous section does not correctly calculate the line total for the invoice. By only multiplying the unit price and the quantity the formula is ignoring whether an item was stained (which incurs an additional cost). To take care of this, we will use cell scripting.

To add a script, select the *ItemTotal* formula that you created in Section Q.4.3.2 - Add a Formula and click the *Scripting* button on the toolbar  . This will bring up a dialog containing all the scripts in the report. Since no scripts have been previously added the dialog will be blank.



Script List Window

Click *New* to create a new script. In the prompt, type the name **`StainCheck`** for the script. Click *OK* and the Formula Builder will open allowing you to add the script. Enter the following script:

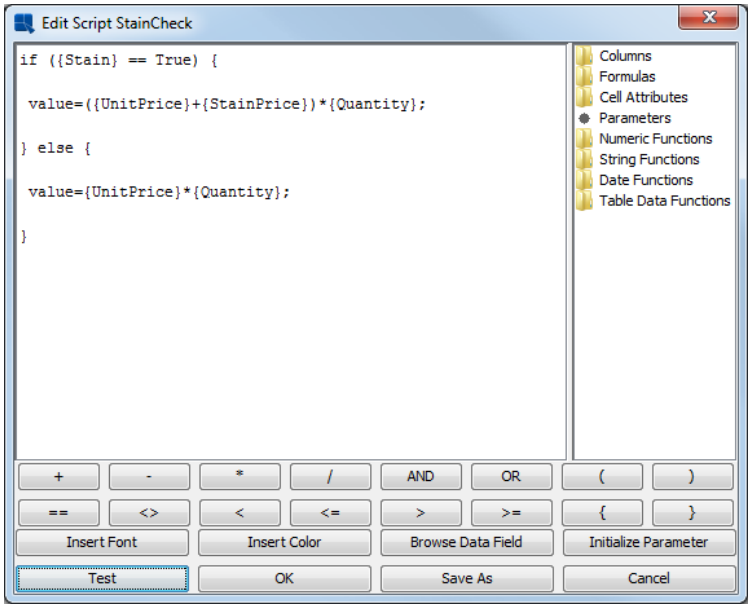
```
if ({Stain} == True) {

value=({UnitPrice}+{StainPrice})*{Quantity};

} else {

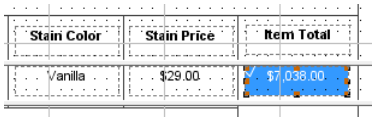
value={UnitPrice}*{Quantity}; }
```

The script dynamically modifies the `ItemTotal` price depending on whether the item has been stained or not. If the item has been stained(`{Stain} == True`) then the `ItemTotal` value includes the stain price. If the item has not been stained, the price is calculated in the same manner as before.



Formula Builder for Cell Scripts

Click *Test* to ensure that the script has been entered correctly. Then click *Ok*. You will return to the script list where the new script has been added. In the script list, select the script that you just created and click the *Apply* button to apply the script to the column. Notice that a check mark appears in the upper-left hand corner of the cell to which the script has been applied. Then click the *Close* button to close the dialog.

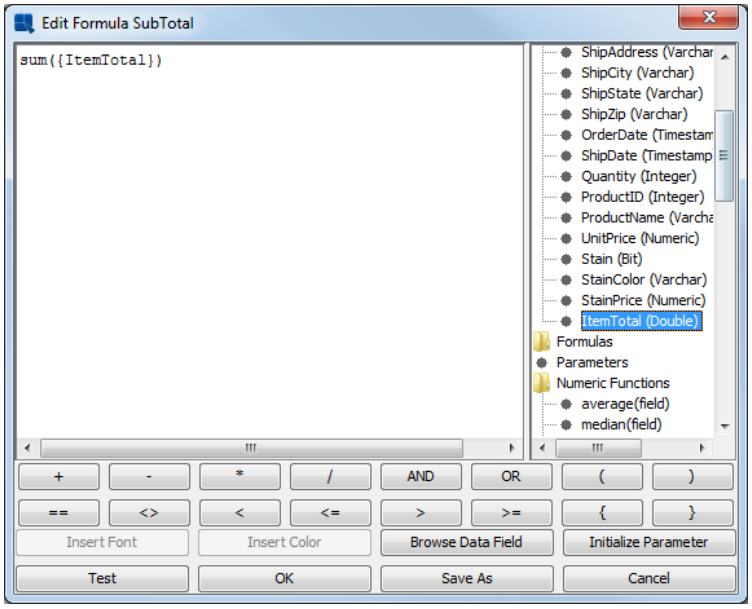


Cell with Applied Script in Design Window

Q.4.3.3.1. Add an Aggregation

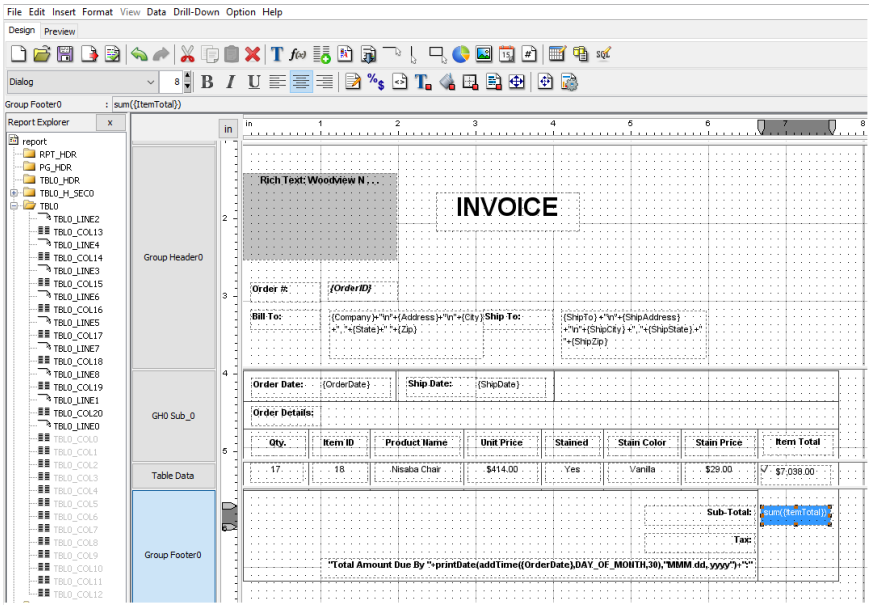
In ERES, formulas can also be used to aggregate report columns. Adding an aggregation is the same as adding a formula like in Section Q.4.3.2 - Add a Formula. Click on the *Insert Formula* icon and then click on the *New* icon to create a new formula. Name this formula **`SubTotal`**.

In the Formula Builder, double click on the *Numeric Functions* folder to expand it. Double click on the *sum(field)* function to insert it into the formula. Next use the cursor to highlight the `field` portion of the sum function. Double click on the *Columns* folder to expand it. At the end of the list of columns is one called `ItemTotal` that you created in Section Q.4.3.2 - Add a Formula. Double click it to add it to the formula. The finished formula should look like this: `sum({ItemTotal})`



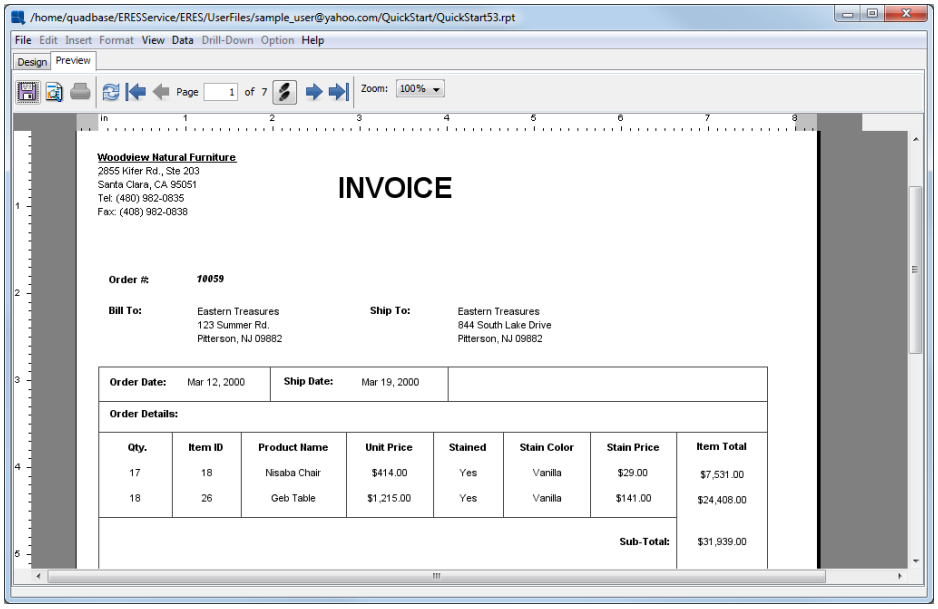
Formula Builder with Aggregation Formula

Click *Test* to ensure that the formula is entered correctly. Then click *OK* to return to the formula list. In the formula list, select the formula that you have created and click the *Insert* button. The dialog will close and a small dotted rectangle will follow your pointer around the design window. Position the formula in the *Group Footer* section of the report, below the `Item Total` column, and click to add it.



Formula Added

Because the formula is in the *Group Footer* section, it will not calculate until the report is run, and it only displays the text of the formula. Format the formula as currency like in Section Q.4.2.3 - Data Formatting. Now preview the report. Notice that the aggregation reflects the values that are modified by the cell script.



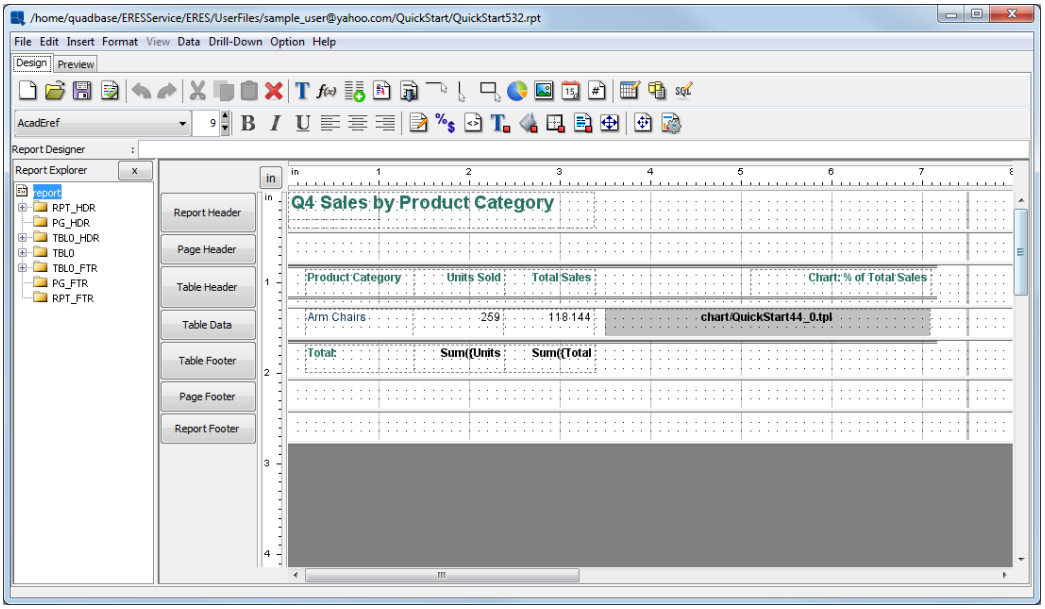
Report With Formulas

For additional tutorial, add two more formulas to the *Group Footer* section: one to calculate the sales tax and one to calculate a grand total for the order. Once you have finished, click the *Save* button to save the changes that you have made and close the Report Designer.

Q.4.3.4. Drill-Down

A unique feature of ERES is the ability to perform drill-down on reports automatically. Using drill-down, users can easily present summarized data in a top-layer report and can click through to view more detailed information. Using this feature, only one template has to be designed for each level of drill-down. For more information on this feature, see Section 5.1.8 - Drill-Down.

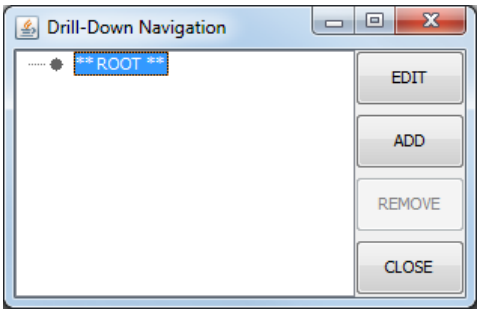
Select QuickStart node under your project in the left-hand side of the Organizer. Double click the `QuickStart532.rpt` from the right-hand side. This will open the report in Report Designer.



`QuickStart532.rpt` in the Design Window

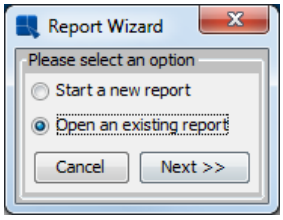
The report shows aggregated sales data grouped by product category. Next, we will link this template to another parameterized template, so that users can drill into each category and look at the sales figures for products in each

category. To add a layer of drill-down, select *Navigate* from the *Drill-Down* menu in the navbar. This will bring up a dialog showing the hierarchy of the any drill-down layers for the report. Only the root will show because there are not any layers defined.



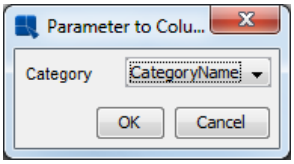
Drill-Down Navigation Dialog

To add a new layer of drill-down, click the *ADD* button. This will bring up a dialog asking you if you would like to use an existing template or create a new one. If you select to create a new template, you will go back to the data registry where you can start designing a report. You can also use an existing report. In either case, the report you use must have a parameterized query or class as the data source. For more on this see Section 4.1.2.2.2 - Parameterized Queries. In this instance, select to use an existing report and click *Next*.



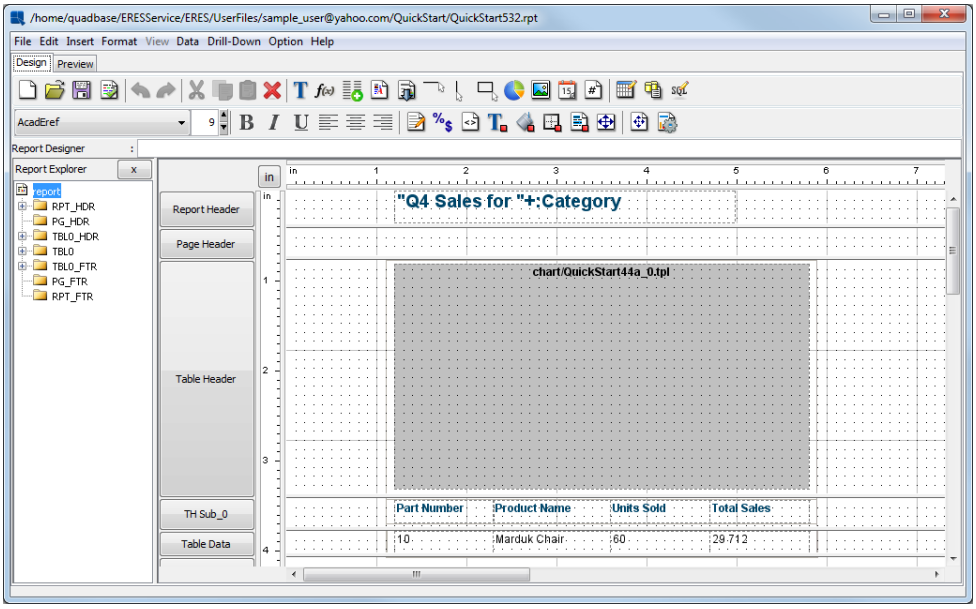
Report Options Dialog

At the prompt browse to the `/QuickStart/`, and select the `QuickStart532a.rpt` file. A dialog will then open prompting you to select a column from the primary report to map to the drill-down layer.



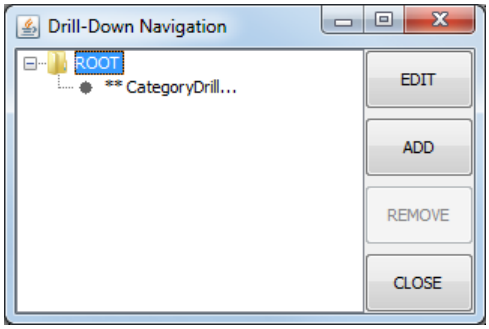
Column - Parameter Mapping Dialog

Select to map the `CategoryName` column to the parameter, and click *OK*. A dialog will open prompting you to specify a display name for the report. Enter a name and click *OK*. The drill-down layer will now open in the design window.



First Drill-Down Level in Report Designer

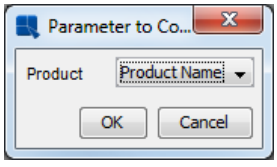
If you preview this level, you can see how the report is parameterized based on the Category Name. Back in the design window, again select *Navigate* from the *Drill-Down* menu. You will now see a new node for your added level under the ROOT node. The stars "**" indicate which level is currently open in the designer.



Drill-Down Navigation Dialog with Additional Layer

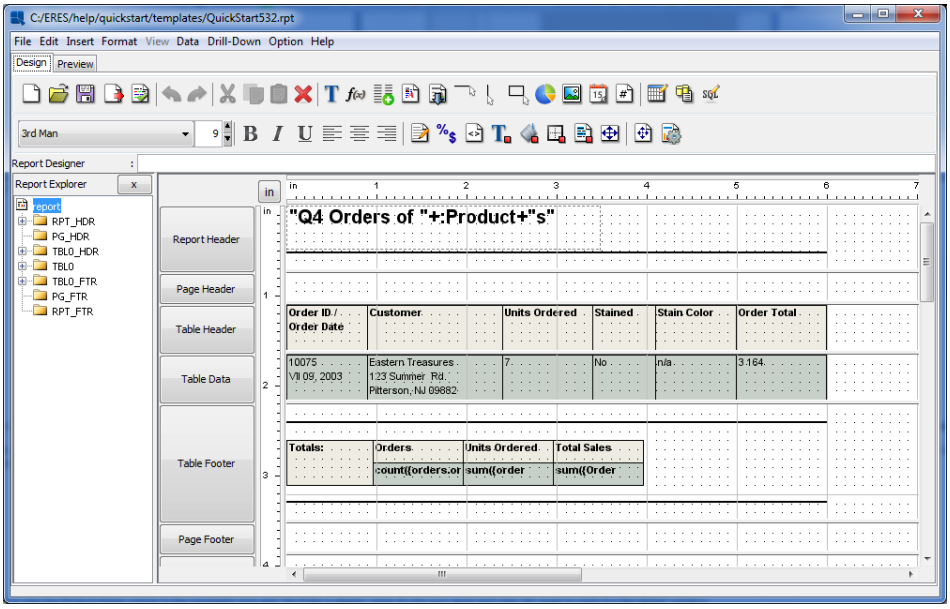
Now we will add one more layer of drill-down to this report, that allows users to drill through the product sales, to see the records of each order for a given product. To do this, select the node for the level you created, and click *Add*.

Again, select to use an existing template for the new drill-down layer. At the dialog, browse to the `/QuickStart/` and select the `QuickStart532b.rpt` file. This will bring up the parameter to column mapping dialog.



Second Parameter Mapping Dialog

Select to map the `Product Name` column to the parameter and click *OK* Enter a display name for the new layer and click *OK* again to open it in the design window.

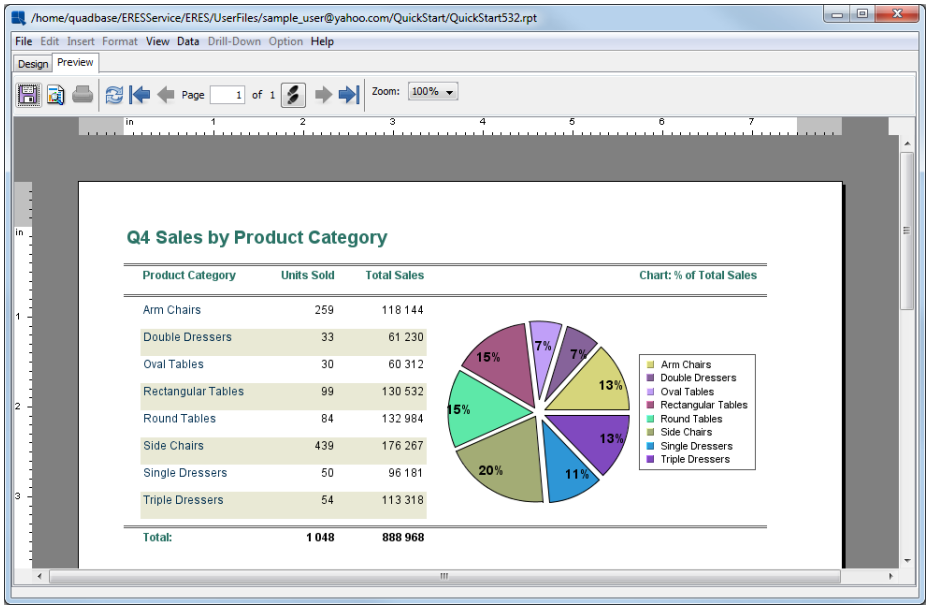


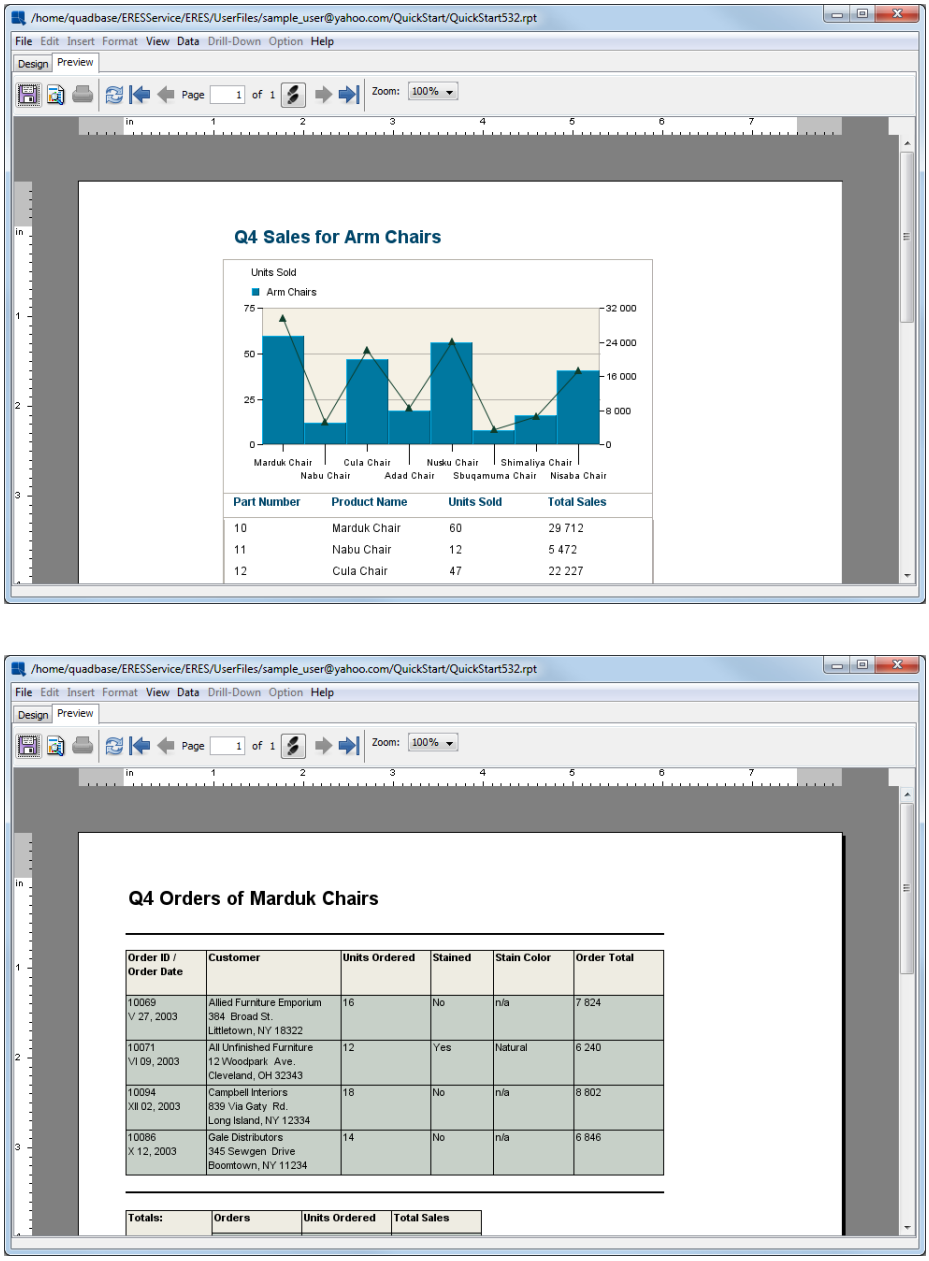
Second Drill-Down Level in Report Designer

Now, select *Navigate* from the *Drill-Down* menu again. Notice that there is a new node for the layer you just added. Select the root report, and click the *Edit* button to open it in the design window. You will see the first report open in the designer. Then click *Close* to dismiss the navigation dialog.

Now preview the report. Notice that the cursor changes when you mouse over a field in the Category Name column. Click on one and you will be taken to the next level that shows sales for each product in that category.

Within the product report you can click on a field in the `Product Name` column to go to the third level report showing the orders for that particular product.





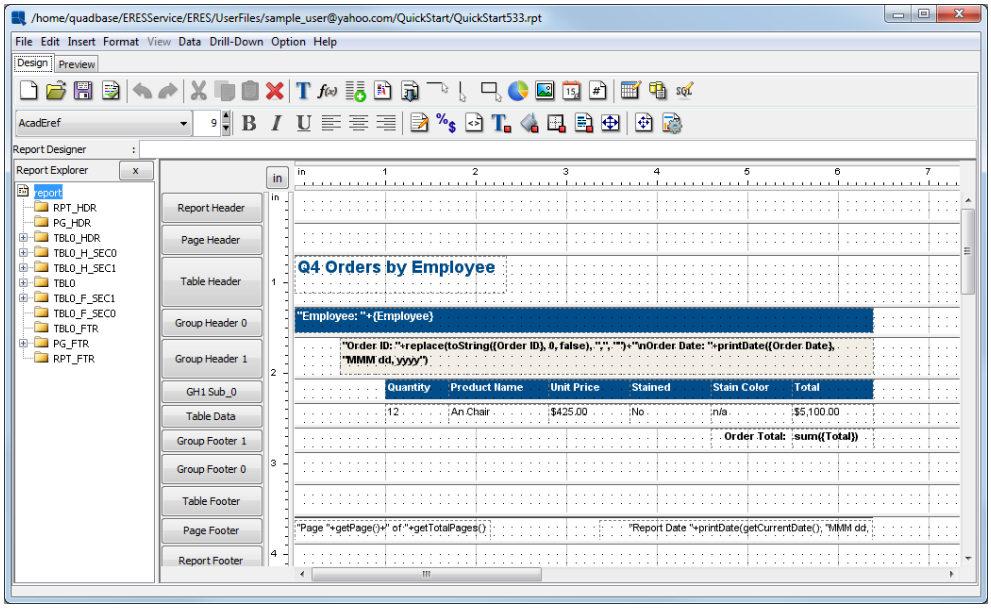
Drill-Down Reports in Preview Window

To navigate back to higher layers, right click and select *Back* from the pop-up menu. Now click the *Save* button to save the changes that you have made to the report.

Q.4.3.5. Sub-Reports

Another powerful feature in ERES is the ability to use sub-reports to create more complex report layouts and combine data from multiple sources in a report. For detailed information about sub-reports, seeSection 5.1.9 - Sub-Reports.

Select QuickStart node under your project in the left-hand side of the Organizer. Double click the `Quick-Start533.rpt` from the right-hand side. This will open the report in Report Designer.

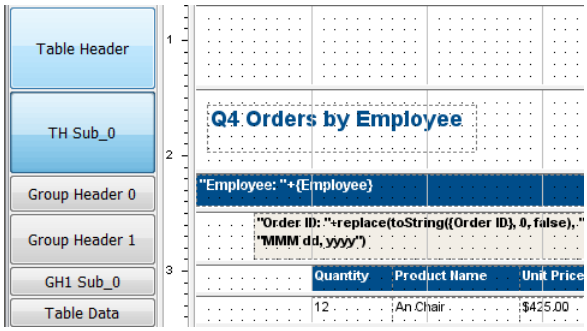


QuickStart533.rpt in Report Designer

The report uses two levels of nested grouping to show sales for each employee. We will now add a sub-report to the header that shows aggregated sales by category and employee (in a crosstab layout). Before adding the sub-report, we will create a new report section in which to place the sub-report. Although a sub-report can be placed anywhere in a report, it often makes sense to give a sub-report its own section, especially if the size of the sub-report can vary.

To insert a nested section, click the *Table Header* button on the left-hand side to bring up the section options menu. Select *Insert Section*. This will spawn a nested section for the table header. For more about nested sections, see Section 5.1.3.1.1 - Nested Sections.

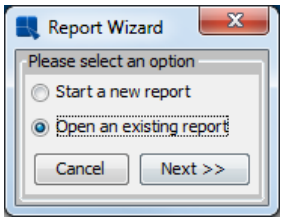
Next we will move the title into the new nested section. To do this, select the cell containing the report title and hit **CTRL**+**X** to cut it. Place the cursor in the new section and hit **CTRL**+**V** to paste. The cursor will turn to a cross. Position it where you would like the field and click to add it. You may want to resize the new section to fit the cell.



New Report Section

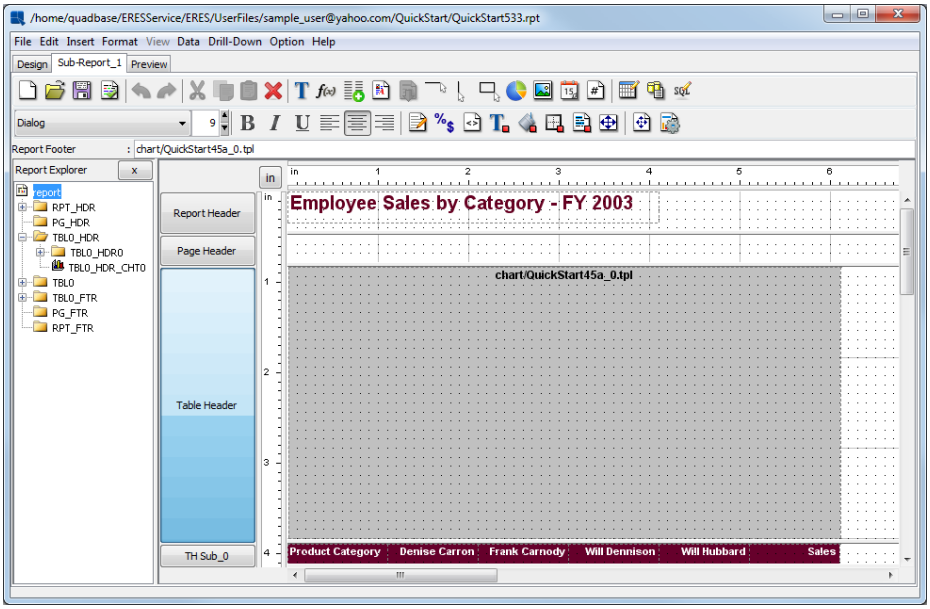
Next, click the *Insert Sub-Report* button on the toolbar  . Your mouse pointer will change to the shape of the plus sign. Press the left mouse button in the top-left corner of the `Table Header` section to insert the SubReport.

As with drill-down, you have the option of creating a new report for the sub-report or using an existing template. Unlike drill-down, however, the template does not have to have a parameterized query as the data source (unless you want to create linked sub-reports. For more on this, please see Section 5.1.9.4 - Linked Sub-Reports). Select to use an existing template and click *Next*.



Report Options Dialog

At the prompt, browse to the `help/quickstart/templates/`, select the `QuickStart533a.rpt` file, and click *OK*. The sub-report will then open in a new tab called `Sub-Report_1` in the Report Designer.

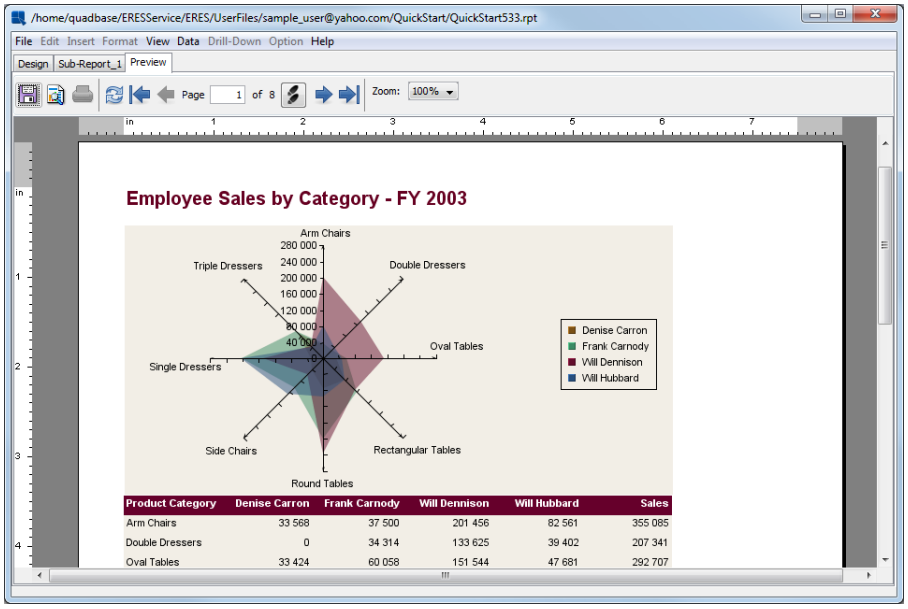


Sub-Report in Designer

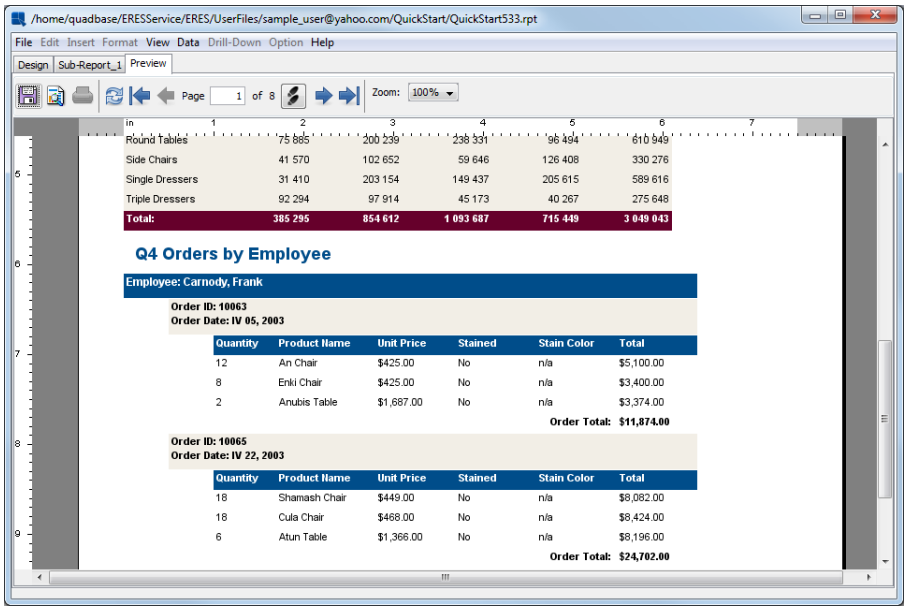
You can preview just the sub-report by toggling between the *Preview* and *Sub-Report* tab. Now go back to the main report by clicking the *Design* tab. The sub-report will appear as a small gray rectangle. Move the rectangle to the upper left-hand corner of the section and click and drag on the horizontal ruler to increase the width of the sub report to be about seven inches.

Next, we will set the sub-report to resize dynamically. Right click on the sub-report and select *Resize To Fit Content* from the pop-up menu. This will turn the dynamic sizing on or off.

Now preview the report. You can see that the entire sub-report runs before the main report.



Sub-Report with Main Report in Preview



Sub-Report with Main Report in Preview

Once you have finished modifying the report, click the *Save* button on the toolbar to save the changes you have made, and close the Report Designer.

Q.5. Chart Designer

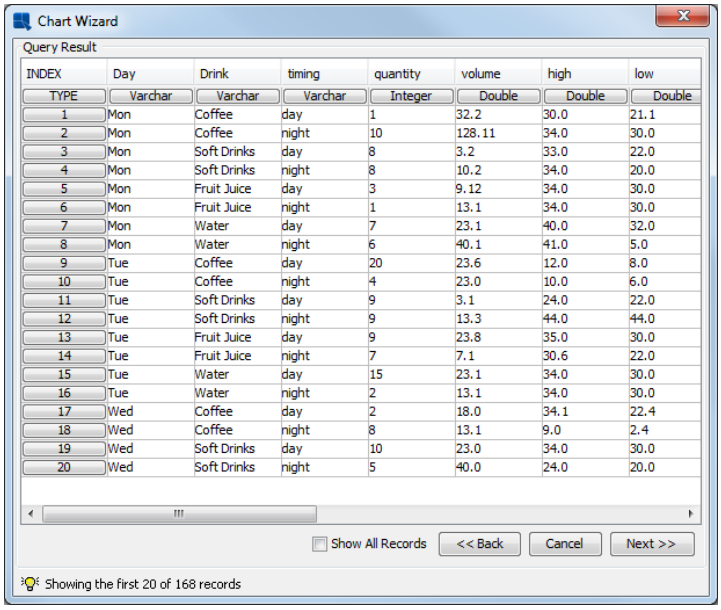
Chart Designer allows users to create and customize charts. This section shows data mapping, and how to use some of the most commonly used features in the Chart Designer. For more information about charts in Chart Designer, please refer to Section 5.2 - Chart Designer.

Q.5.1. Chart Mapping

This section will look at several ways that data can be mapped to charts.

Q.5.1.1. Column Chart

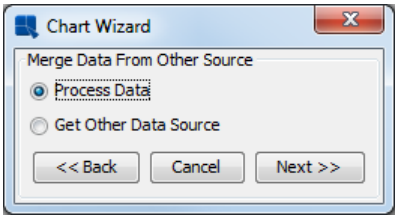
Column charts are a good starting point as the mapping for column charts is very similar to that of bar, area, and line charts. To begin creating a chart, click the *Chart Designer* button in the Organizer toolbar  . This will launch the Chart Designer interface and open the Data Source Manager. In the Data Source Manager, select the node for the text data source *TXTFiles/Sample* and click the *Next* button. A dialog will open a table with the contents of the text file (first 20 records).



INDEX	Day	Drink	timing	quantity	volume	high	low
TYPE	Varchar	Varchar	Varchar	Integer	Double	Double	Double
1	Mon	Coffee	day	1	32.2	30.0	21.1
2	Mon	Coffee	night	10	128.11	34.0	30.0
3	Mon	Soft Drinks	day	8	3.2	33.0	22.0
4	Mon	Soft Drinks	night	8	10.2	34.0	20.0
5	Mon	Fruit Juice	day	3	9.12	34.0	30.0
6	Mon	Fruit Juice	night	1	13.1	34.0	30.0
7	Mon	Water	day	7	23.1	40.0	32.0
8	Mon	Water	night	6	40.1	41.0	5.0
9	Tue	Coffee	day	20	23.6	12.0	8.0
10	Tue	Coffee	night	4	23.0	10.0	6.0
11	Tue	Soft Drinks	day	9	3.1	24.0	22.0
12	Tue	Soft Drinks	night	9	13.3	44.0	44.0
13	Tue	Fruit Juice	day	9	23.8	35.0	30.0
14	Tue	Fruit Juice	night	7	7.1	30.6	22.0
15	Tue	Water	day	15	23.1	34.0	30.0
16	Tue	Water	night	2	13.1	34.0	30.0
17	Wed	Coffee	day	2	18.0	34.1	22.4
18	Wed	Coffee	night	8	13.1	9.0	2.4
19	Wed	Soft Drinks	day	10	23.0	34.0	30.0
20	Wed	Soft Drinks	night	5	40.0	24.0	20.0

Contents of Text File

At the bottom of the dialog, click the *Next* button. A dialog will open asking you if you would like to select an additional data source for the chart.



Add Data Source Dialog

Select the *Process Data* option and click *Next*. You will then be taken to a dialog prompting you to select which type of chart you would like to create.

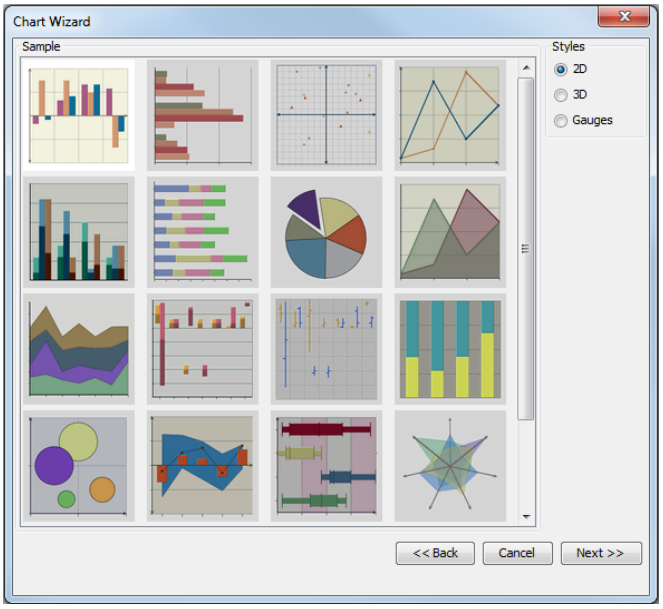


Chart Types Dialog

In this dialog, leave selected a *Column Chart* (the first image) and a 2D option (two-dimensional chart) as the data type and click *Next*. You will then be taken to the data mapping dialog which allows you to map columns from the data source to chart elements.

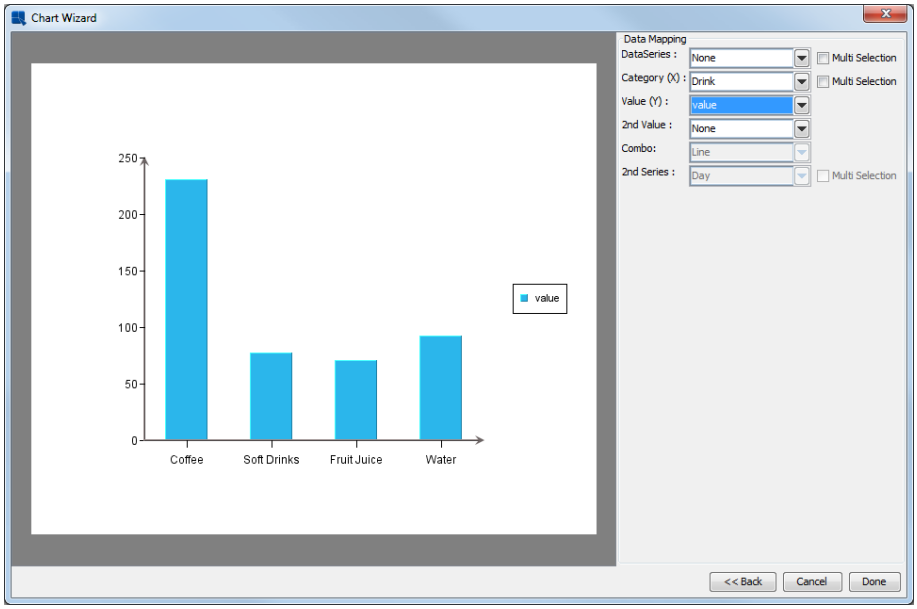


Chart Data Mapping Dialog

Set the *Data Series* to **`None`**, the *Category* to **`Drink`** and the *Value* to **`value`**. Then click *Done* to finish the Wizard and go to the Chart Designer interface.

If the generated chart does not fit in the viewport of the Chart Designer, select *Canvas* from the *Format* menu. This will bring up a dialog allowing you to set the size of the chart.

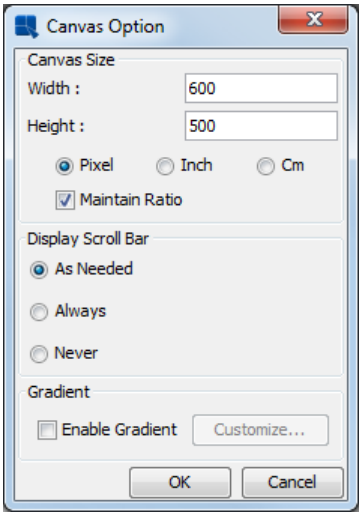


Chart Canvas Dialog

Enter a smaller size for the chart canvas and click *OK* to close the dialog. You should now see the whole chart in the viewport.

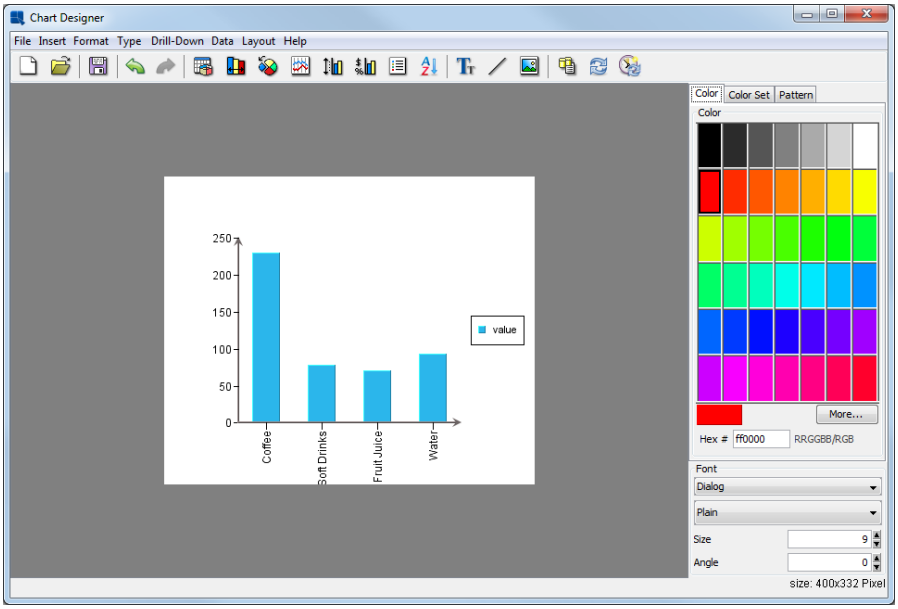


Chart Designer Showing Column Chart

Notice that the column chart contains a column for each distinct value in the Drink column and shows the corresponding value.

Q.5.1.1.1. Add a Data Series

At this point, the column chart only contains data points for the categories `Drink` column. However, ERES supports adding another dimension to this data by way of a series. To add a series to the column chart, click the *Data Mapping* button in the toolbar  . This will return you to the data mapping window.

In the mapping window, change the *Data Series* option from **`None`** to **`Day`** and click *Done* to return to the Chart Designer.

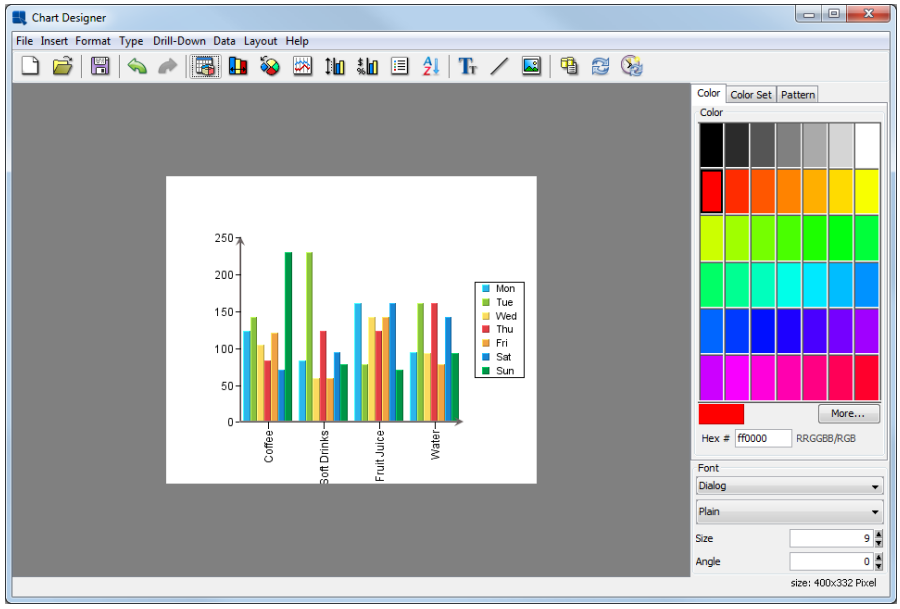


Chart Designer Showing Column Chart with Series

Notice now that instead of a single data point for each drink, the column for each drink is now comprised of seven small columns, one showing the data point for each day.

Q.5.1.1.2. 3D Column Chart

In a two-dimensional chart, the data series is represented in-line along the X axis. However, a three-dimensional representation provides another axis to work with.

In the Chart Designer select *3D Chart* from the *Type* menu to convert the column chart to a 3D representation.

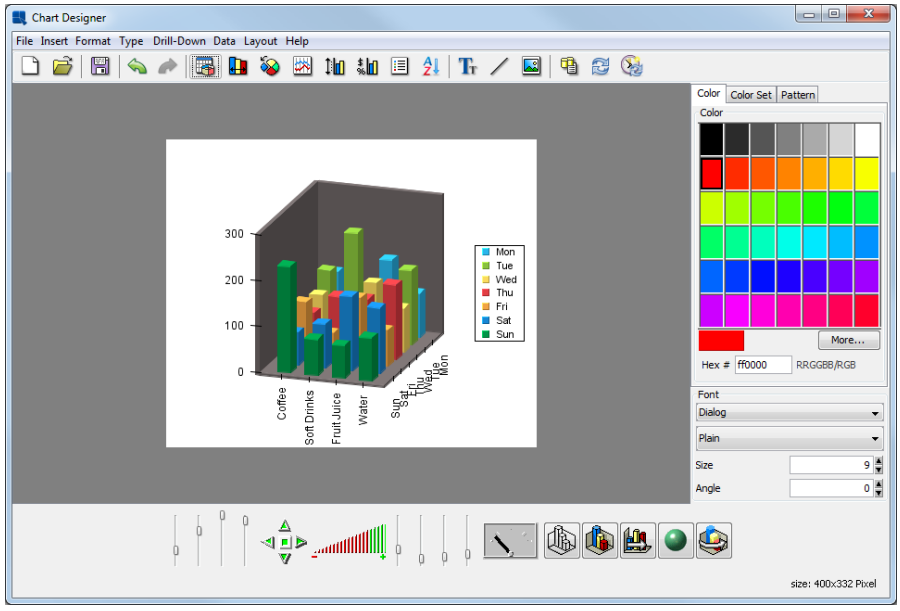


Chart Designer Showing 3D Column Chart

The 3D toolbar will automatically appear at the bottom of the Chart Designer when the chart is converted to 3D. The chart may appear squished when you convert it. You can use the sliders next to the zoom function to change the X, Y, and Z scaling for the 3D chart. You can also use the navigation buttons to position the chart in space. For more information about 3D features, see Section 5.2.4.4 - The Navigation Panel.

Notice that with the chart in 3D, the data series is now moved to the Z axis of the chart by default.

Q.5.1.2. Pie Chart

Pie charts are another commonly used chart type that shows values as a percentage of a whole. To convert your chart to a pie chart, first go to the *Type* menu and select *2D Chart* to convert your column chart back to two-dimensional form. Then select *Pie* from the *Type* menu. The chart will then be converted to a pie representation.

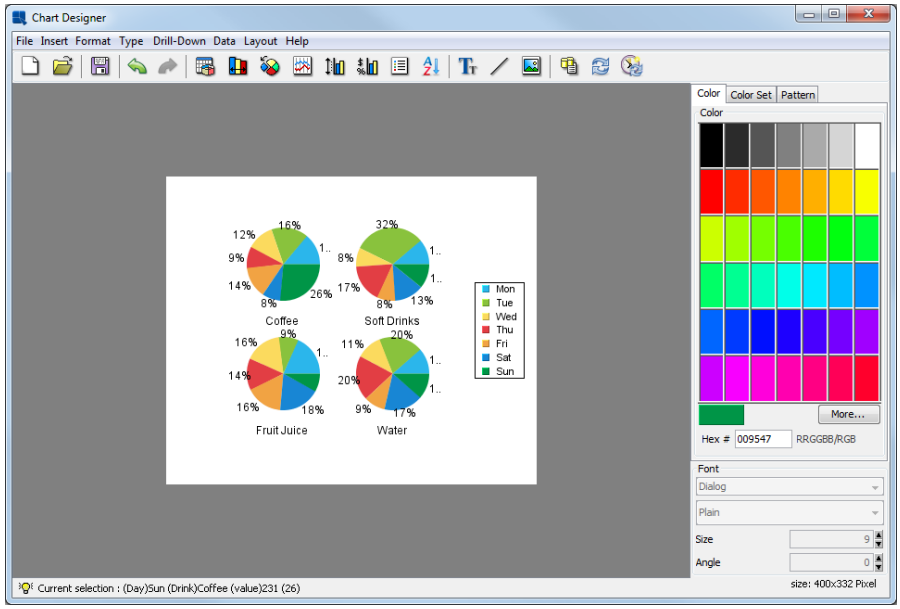
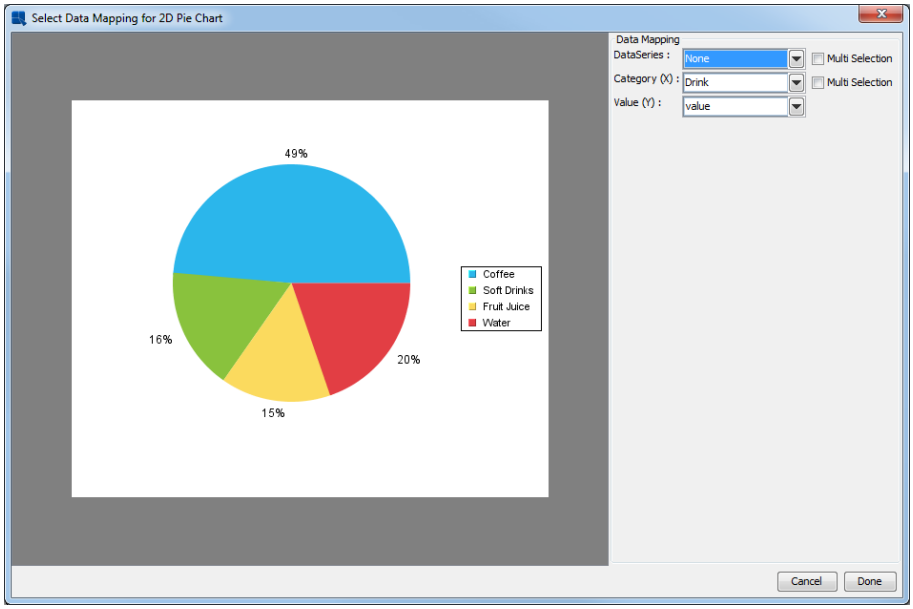


Chart Designer Showing Pie Chart with Series

Multiple pies are drawn when you have a data series; one for each category. Each category is broken down showing percentage contribution for the series elements.

In order to turn the chart into a single pie, we will remove the series. Click the *Change Data Mapping* button  on the toolbar to bring up the mapping options.



Data Mapping Options for Pie Charts

Change the *Data Series* option to **None** and click *Done*. Now in the Chart Designer you will see a single pie made up of your drink categories.

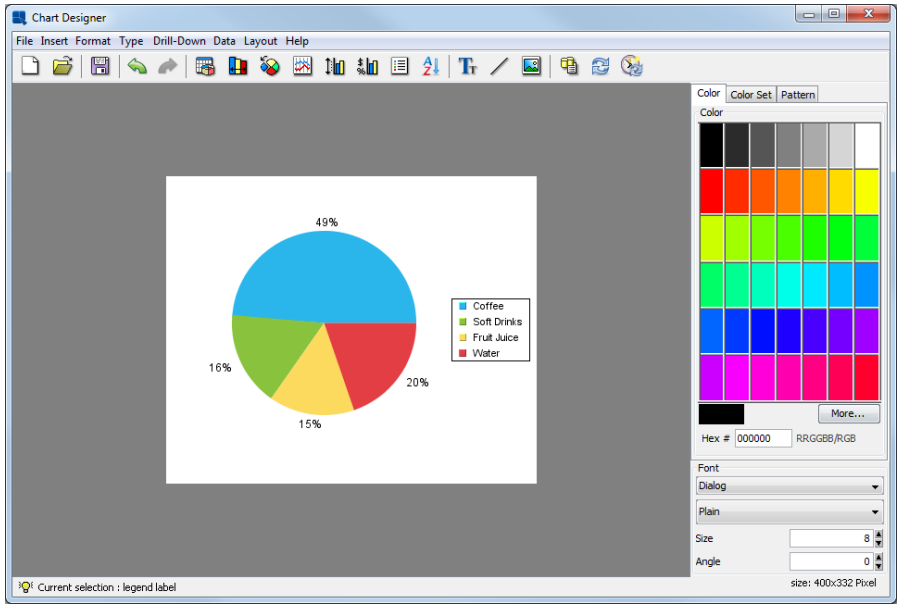


Chart Designer Showing Pie Chart

Q.5.1.3. Stack Column Chart

Another way to display multi-dimensional data is to use a stack type representation to show contributions to a total. To convert your pie chart into a stack column chart, select *Stack Column* from the *Type* menu. Chart Designer will ask whether you want to redo data mapping for the new chart type. Click *Yes* and confirm the change by clicking *Done* in the "Select Data Mapping for 2D Stack Column Chart" dialog. The chart will then be converted to a stack column layout.

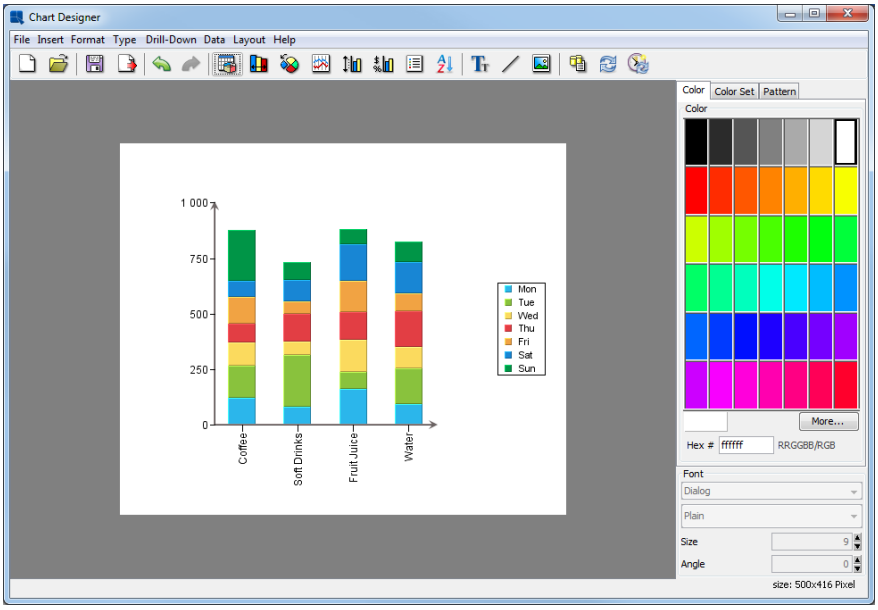
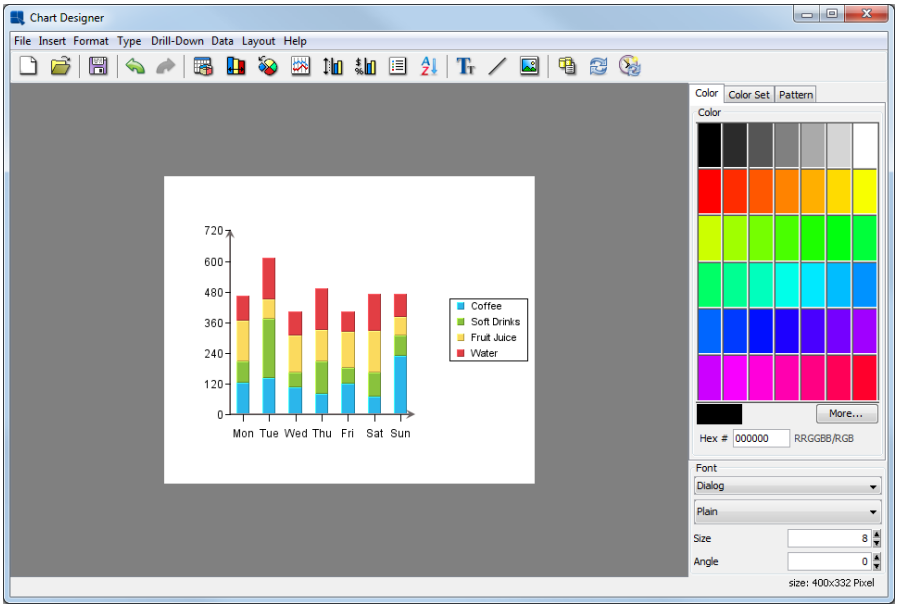


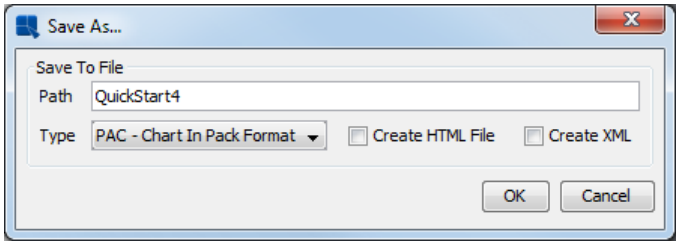
Chart Designer Showing Stack Column Chart

Notice that the chart now shows the columns as stacks made up of the values for each drink. Now click the *Data Mapping* button on the toolbar. In the data mapping window there is a new option called *Sum by*. This is set to the *Day* column. The *Sum by* column provides the individual stacks in these types of charts. Set *Sum by* to *Drink* and *Category* to *Day*. Click *Done* to return to the Chart Desinger. You should see the following chart.



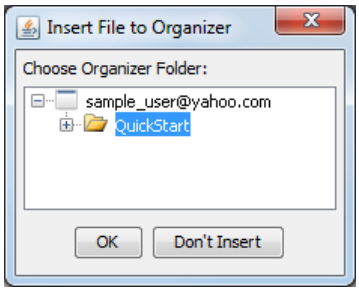
Stack Column Chart

Now that you have finished designing a chart, click the *Save* button on the Toolbar  . A dialog will open prompting you to specify a file name for the chart.



Save As Dialog

Enter a file name, and select to use either a PAC, CHT, or TPL format. For more about chart formats, see Section 5.2.6.2 - Saving Charts without Report Data. The chart will be saved in the `/user_email/Templates/` directory. You will then be prompted to add the file to the Organizer.



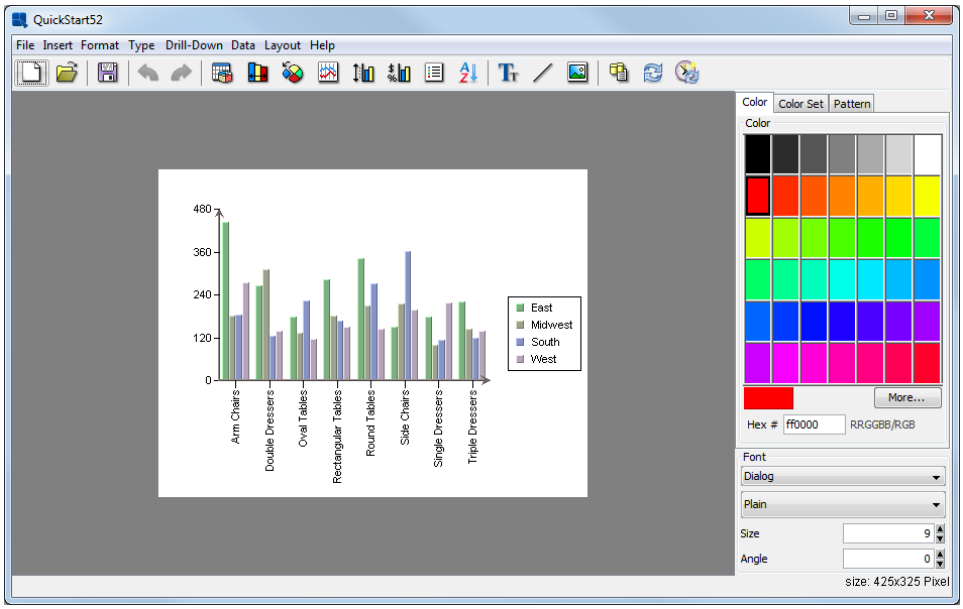
Isert File Dialog

Choose *QuickStart* folder. Click *OK* to add the file to the folder. Close the Chart Designer.

Q.5.2. Basic Chart Formatting

In this section, we will use some of the basic formatting features in the Chart Designer to change the look and feel of an unformatted chart.

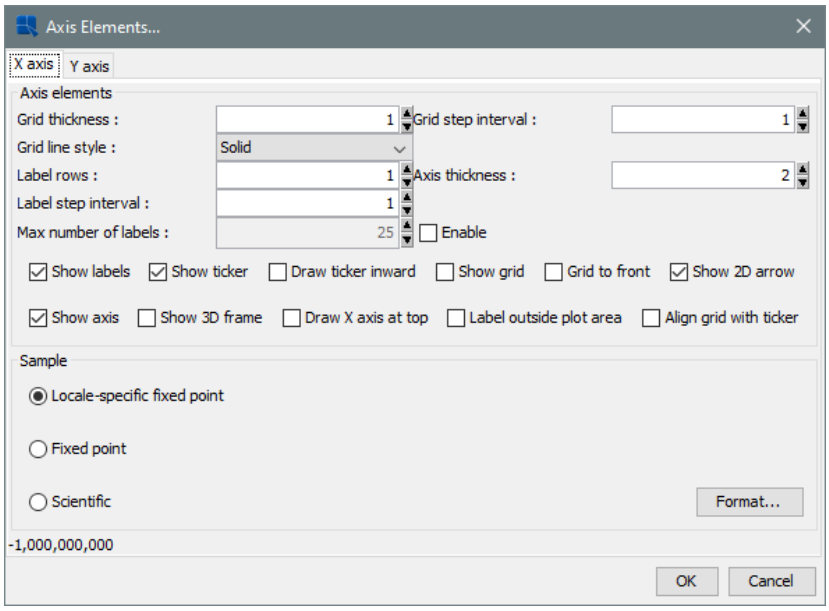
Select QuickStart node under your project in the left-hand side. Double click the QuickStart52.tpl from the right-hand side. This will open the chart in the Chart Designer. As you can see its relatively un-formatted.



Unformatted Chart in Designer

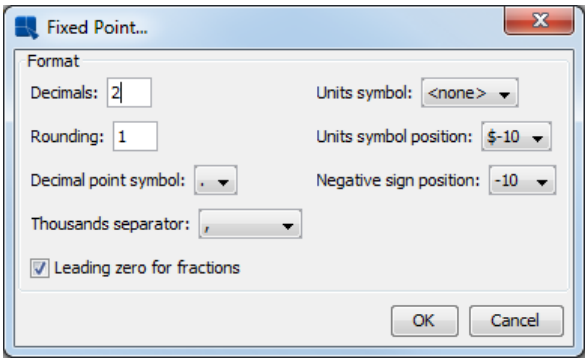
Q.5.2.1. Axis Options

Many chart appearance options are controlled using the Axis Elements dialog. To bring up this dialog, click the *Format Value Elements* button on the toolbar  . This will bring up a tabbed dialog allowing you to set different options for each chart axis. Click on the *Y Axis* tab to bring up options for the value axis.



Axis Elements Dialog

Check the box marked *Show grid* to add grid lines to the Y Axis. Select *Fixed point* for the data format and click the *Format* button. This will bring up an additional dialog allowing you to set format options for the numeric data. Select the number of decimals as **2** and click *OK*.



Numberic Format Dialog

Click *OK* again to dismiss the axis elements dialog and you will see the specified changes reflected in the chart.

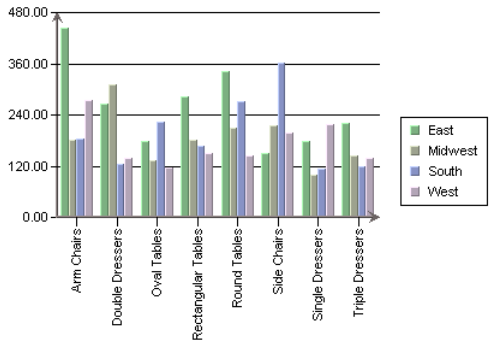
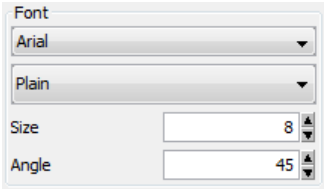


Chart with new Axis Formatting

Q.5.2.2. Modify Color and Font

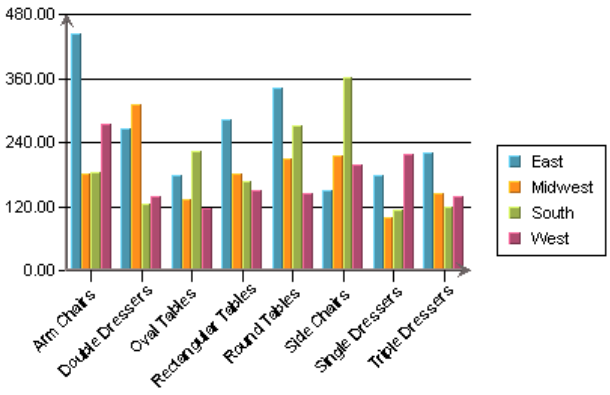
Chart colors can be selected individually or you can select predefined color set. In Chart Designer, click on the *Color set* tab on the right and select a color combination you like.

Now click on one of the X axis (category) labels. The lower left corner of the design window will reflect your selection. In the font dialog in the lower right corner, change the angle of the labels from 90 to 45 degrees and press Enter.



Font Panel

The chart labels will then rotate. You may need to adjust the position of the labels slightly. To do this, click and drag on a label. All the X axis labels will follow your cursor around the chart canvas.



Q.5.2.3. Add Titles

Next, we will add titles to the chart. To do this, select *Titles* from the *Insert* menu. This will bring up a dialog prompting you to enter a main title for the chart as well as titles for each of the axes.

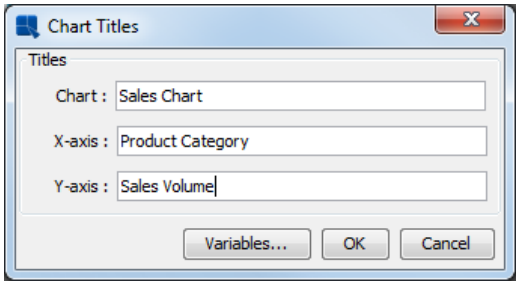


Chart Titles Dialog

Enter any titles that you would like for the various elements and click *OK*. The titles will then be added to the chart. Titles are placed automatically but you can manually adjust their positions by clicking and dragging the text on the chart canvas.

If you do not see chart titles, click and drag on the chart plot to move the chart. You should then be able to see the titles. You can adjust the font and font size of the chart titles by clicking them and changing the settings in the font panel.

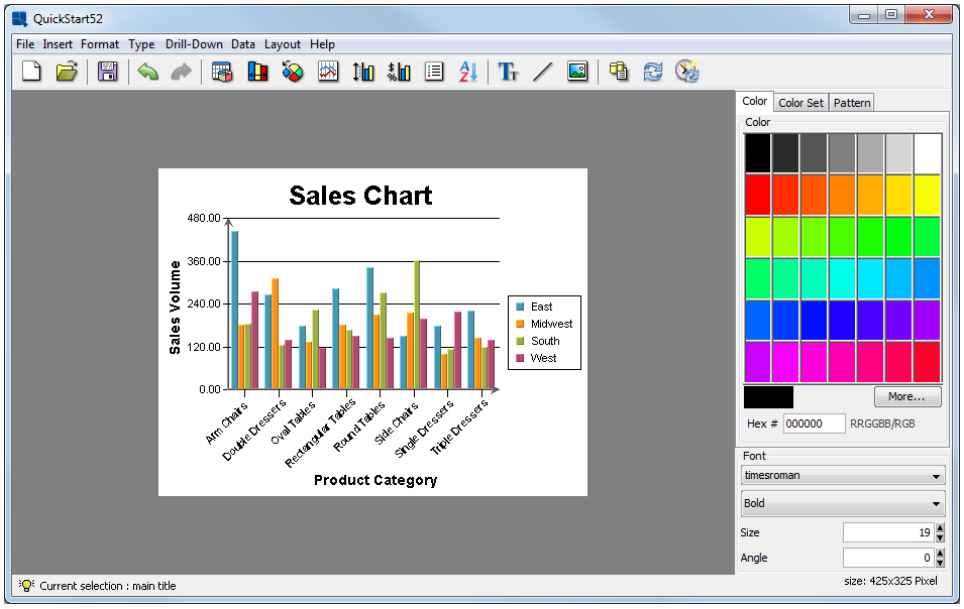
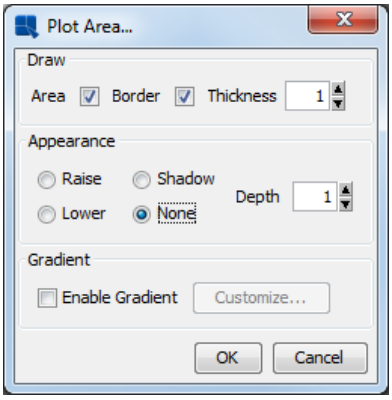


Chart With Titles

Q.5.2.4. Customize Plot Area

You can also customize the plot of two-dimensional charts, creating a separate background for the data points. To do this, select *Plot Area* from the *Format* menu. This will bring up a dialog allowing you to set display options for the chart plot.



Plot Area Dialog

Select to draw both the *Area* and the *Border* with a *Thickness* of **1**. Specify *None* for the *Appearance*. Once you have specified the options, click *OK* and the plot area for the chart will be modified. You can change the background color of the plot area by clicking to select it and then modifying the color in the color panel.

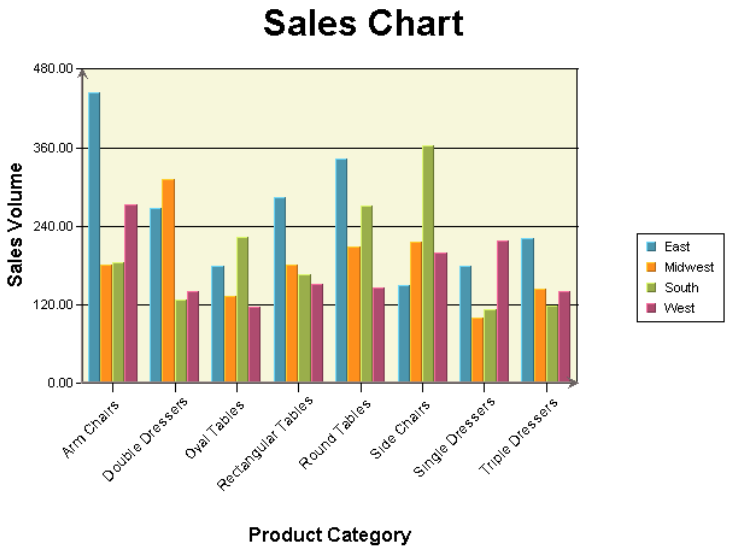


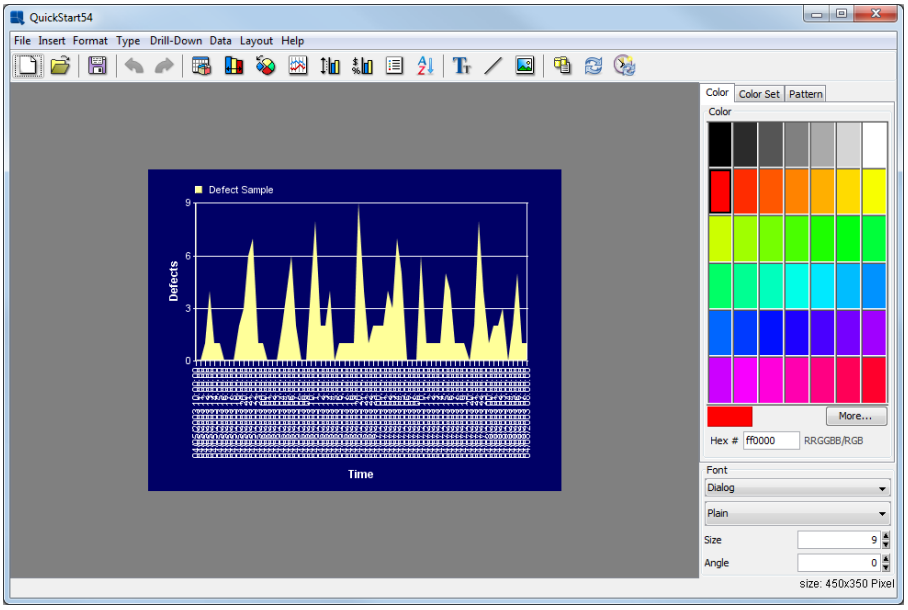
Chart With Colored Plot Area

Now save the changes you have made to the chart and exit the Chart Designer.

Q.5.3. Advanced Charting Features

In this section we will use some of the charting features in ERES to provide a chart that quickly conveys salient information to the viewer. For detailed information about the features discussed in this section and other charting features, please see Section 5.2.4 - The Chart Designer Interface.

Select QuickStart node under your project in the left-hand side of the Organizer. Double click the `Quick-Start54.tpl` from the right-hand side. This will open the chart in Chart Designer.



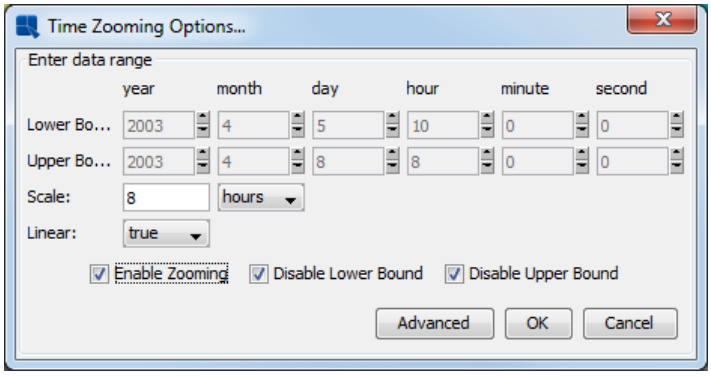
QuickStart54.tpl in Chart Designer

The chart shows hourly defect counts for a manufacturing process. Since the chart shows three days worth of data it is difficult to make out the individual points and the X axis labels are not legible because they are drawn on top of each other.

Q.5.3.1. Time-Series Zooming

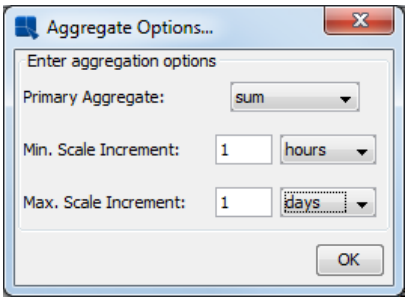
One way to improve the chart, is to use the time series zooming feature to aggregate the data into fewer data points. In this instance, lets look at the total number of defects for each eight hour shift, instead of plotting the data for each hour.

To turn on zooming, select *Time Zooming Options* from the *Format* menu. This will bring up the following dialog.



Zoom Options Dialog

In this dialog, set the *Scale* to **`8 hours`** and *Linear* option to **`true`**. Leave the *Disable Upper Bound* and *Disable Lower Bound* options checked. Check the *Enable Zooming* option. This will bring up a new dialog prompting you to select the aggregation for the zooming.



Aggregation Options Dialog

In this dialog, set the *Primary Aggregate* to be **`sum`**, the *Minimum Scale* to be **`1 hour`**, and the *Maximus Scale* to be **`1 day`**. Then click *OK* to return to the zoom options dialog, and *OK* again to return to apply the zoom setting. The chart now shows the total number of defects for each eight hour period.

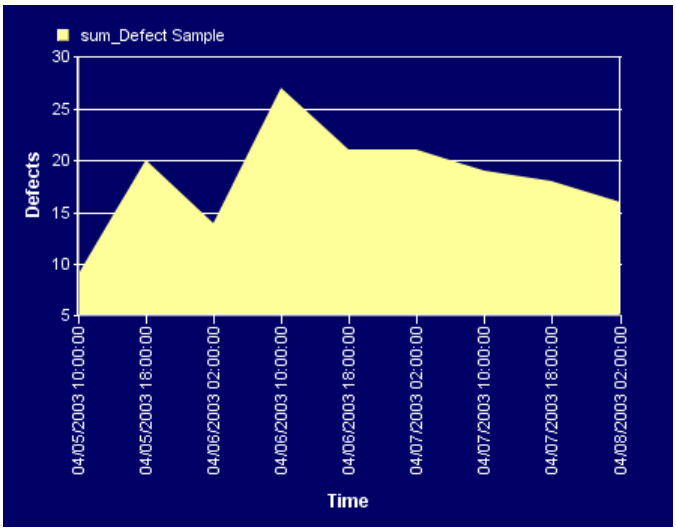
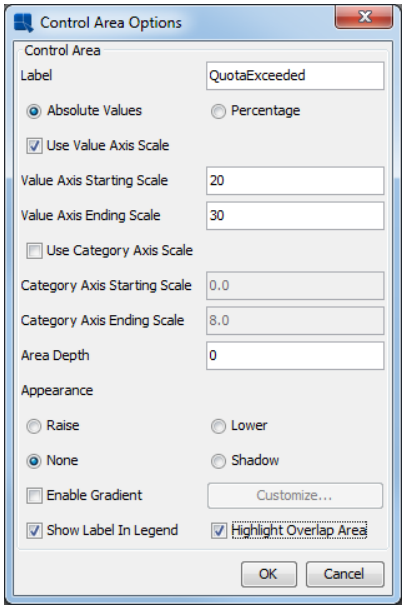


Chart With Zooming Applied

Q.5.3.2. Control Areas

Next, we will use the control area feature to highlight any shift where there were more than 20 defects.

To add a control area, select *Control Area* from the *Insert* menu and then click *Insert* in the control range list. This will bring up a dialog allowing you to select options for the new control area.



Control Area Dialog

In this dialog, enter a name for your control area. Then check the option marked *Use Value Axis Scale*, and enter **`20`** for the *Starting Scale*, and **`30`** for the *Ending Scale*. Then at the bottom of the dialog check the options marked *Show Label In Legend* and *Highlight Overlap Area*. Then click *OK* to apply the control area. Notice how only the areas where the number of defects is colored. You can set the color, by selecting the overlapped area, and selecting a new color from the palette on the right-hand side.

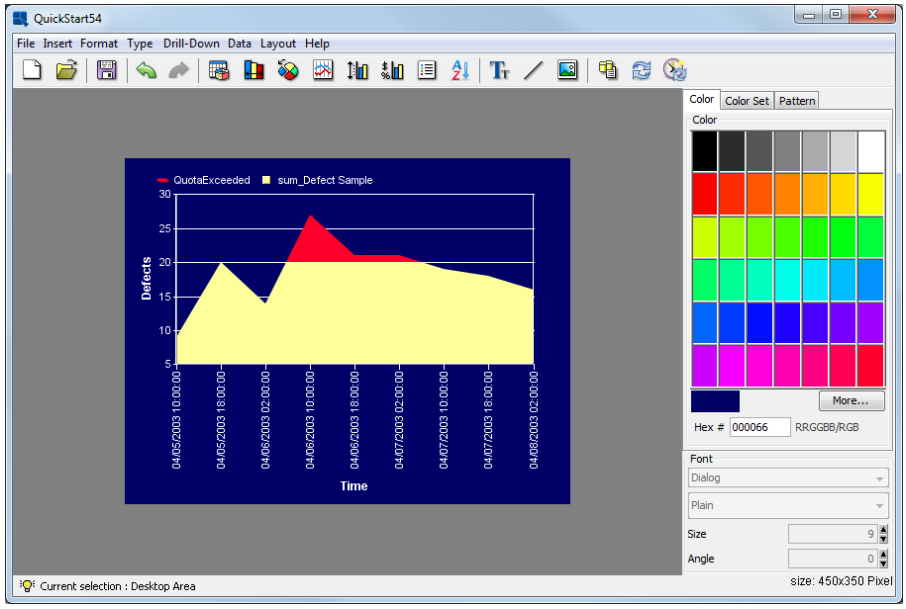


Chart With Control Areas Applied

Once you have finished modifying the chart, click the *Save* button on the toolbar to save the changes you have made and close the Chart Designer.

Q.6. QuickDesigner Reports

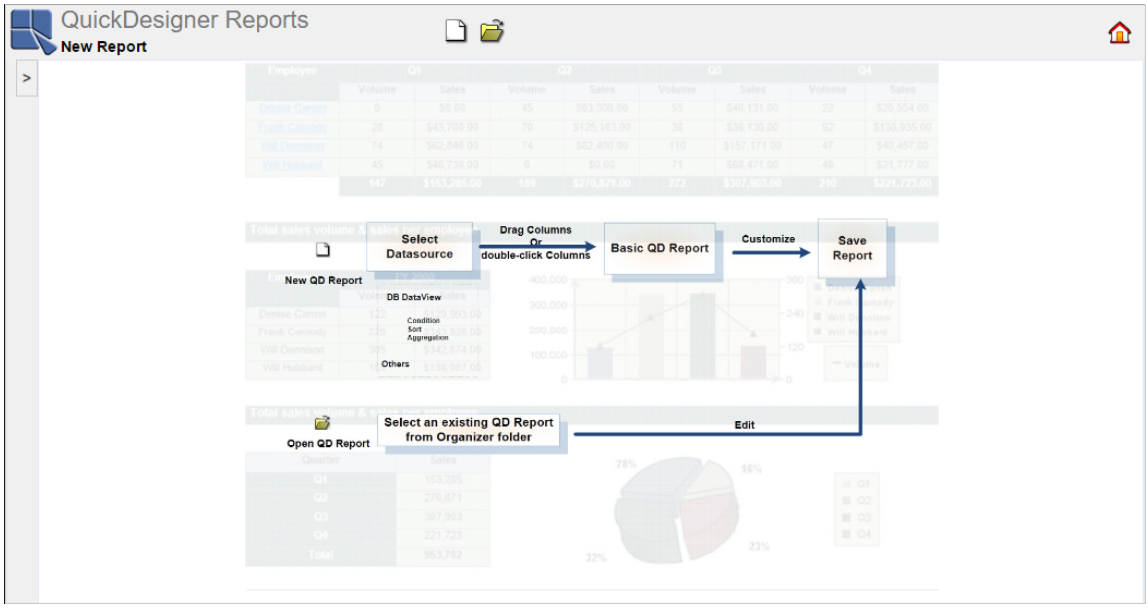
In the previous sections we looked at using Report Designer and Data Source Manager interfaces to query databases and create reports. ERES also provides the QuickDesigner Reports interface which provides a thin-client interface for creating ad hoc queries and reports. In this chapter, we will create a query and a report using QuickDesigner Reports.

To start QuickDesigner Reports, go to the ERES Start page and click the link labeled *QuickDesigner Reports*

Q.6.1. Create a Query

QuickDesigner Reports allows you to build a report using any data source that you can with Report Designer. However, if you want to modify a data source, you can only query DataViews or modify existing DataView queries using QuickDesigner Reports.

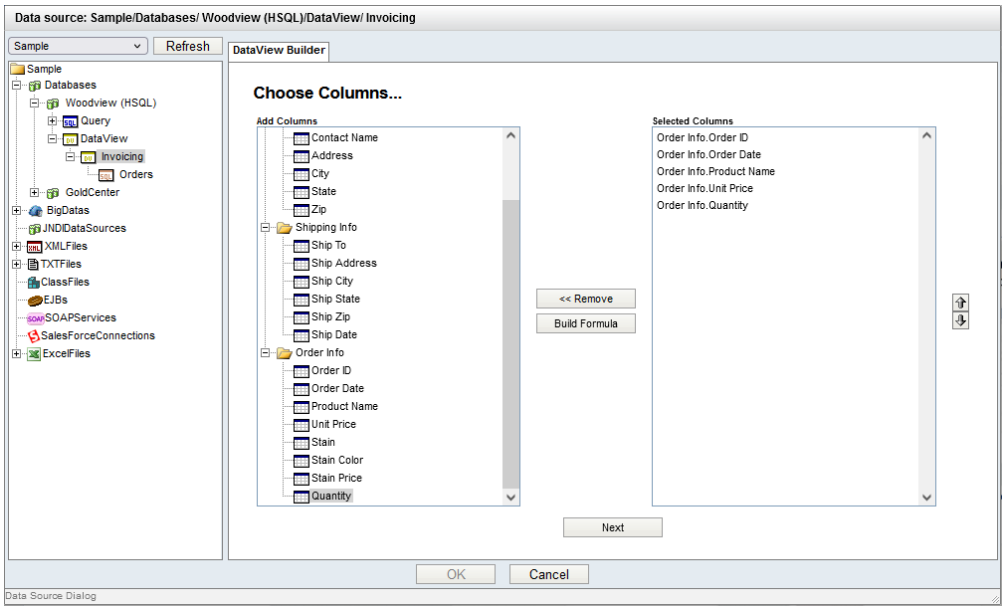
When QuickDesigner Reports opens, there are only two options accessible, to create a new report or open an existing one.



QuickDesigner Reports

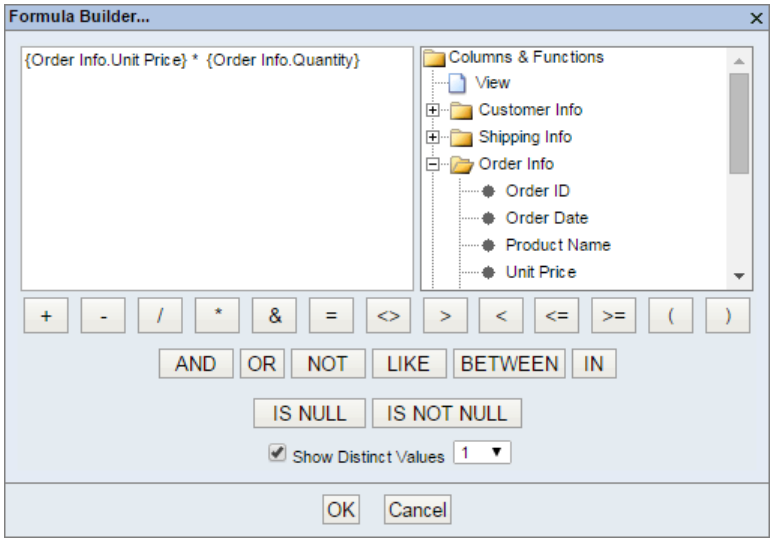
Click the  *New* button on the toolbar and the *Data Source Dialog* will open. Select the *Quickstart_sample* data registry from the left top drop-down list (the data registry you created in Section Q.3.1 - Open a Data Source Manager). The tab bellow shows all of the defined data sources in the registry. Expand the nodes and select the DataView `Invoicing` that you created in Section Q.3.1.2 - Create a Data View. You will see all headings and fields of the view in *DataView Builder* in the right pane of *Data Source Dialog*. Add the following fields from the `Order Info` heading by clicking on them in the left side panel.

```
Order ID
Order Date
Product Name
Unit Price
Quantity
```



Data Source Dialog

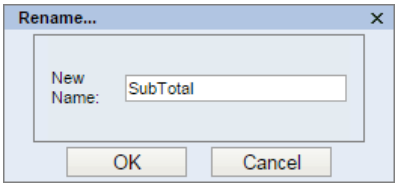
After adding the columns, click the *Build Formula* button to build a computed column. This will launch the Formula Builder in a new window.



Formula Builder

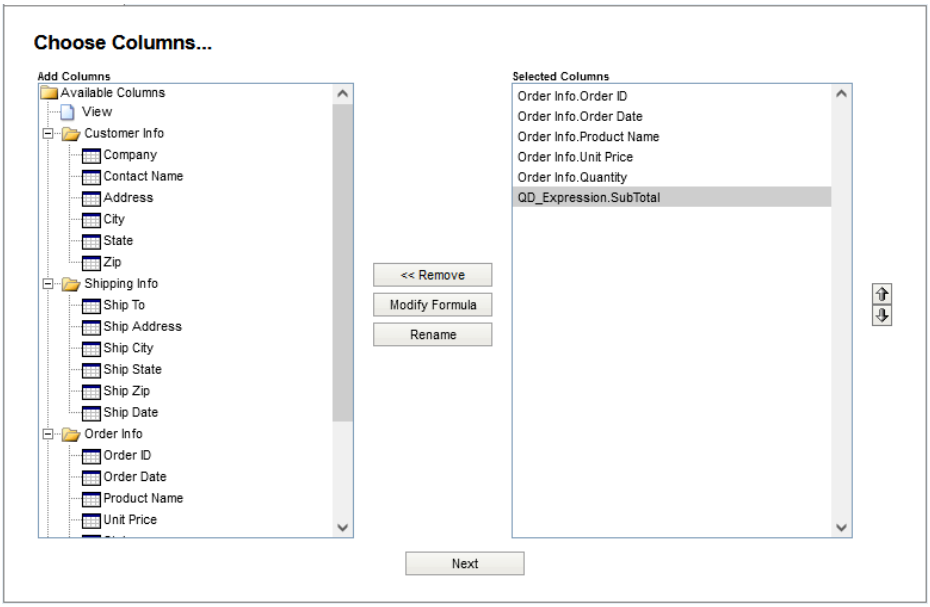
In the Formula Builder, expand the *View* node in the right-hand side. Then expand the *Order Info* node. Under this node click the *Unit Price* entry. The text will be added to the main formula window. Click the *multiply button (*)*. Finally, select the *Quantity* field. The finished formula should look like this: `{Order Info.Unit Price} * {Order Info.Quantity}`. Once you have finished, click *OK* to add the formula. A new column will be added in the column selection dialog.

To provide an alias for the expression column, select it and click the *Rename* button. This will bring up a dialog allowing you to specify the alias.



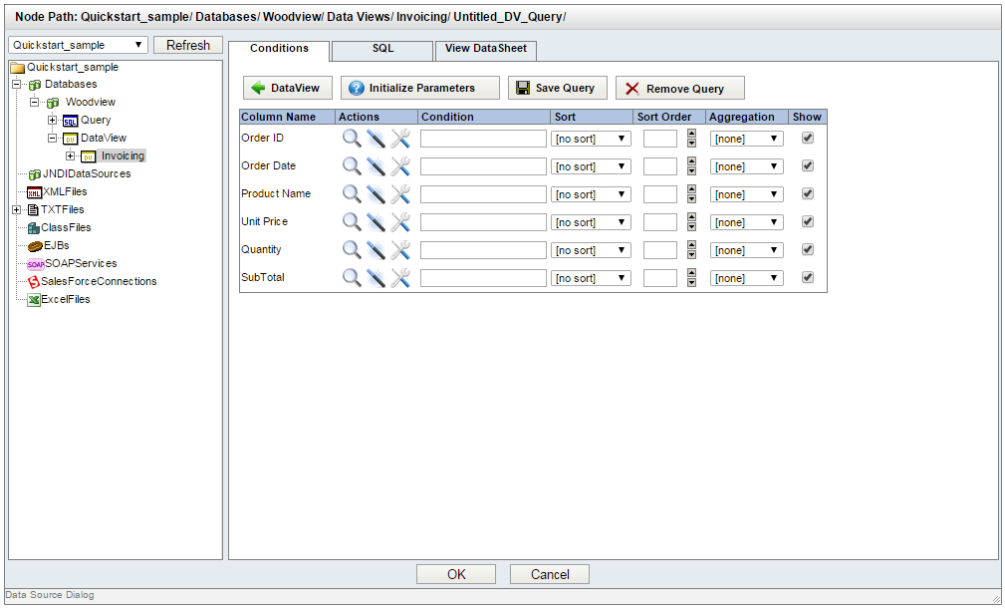
Rename Dialog

Enter **`SubTotal`** as an alias for the computed column and click *OK*. The computed column is renamed.



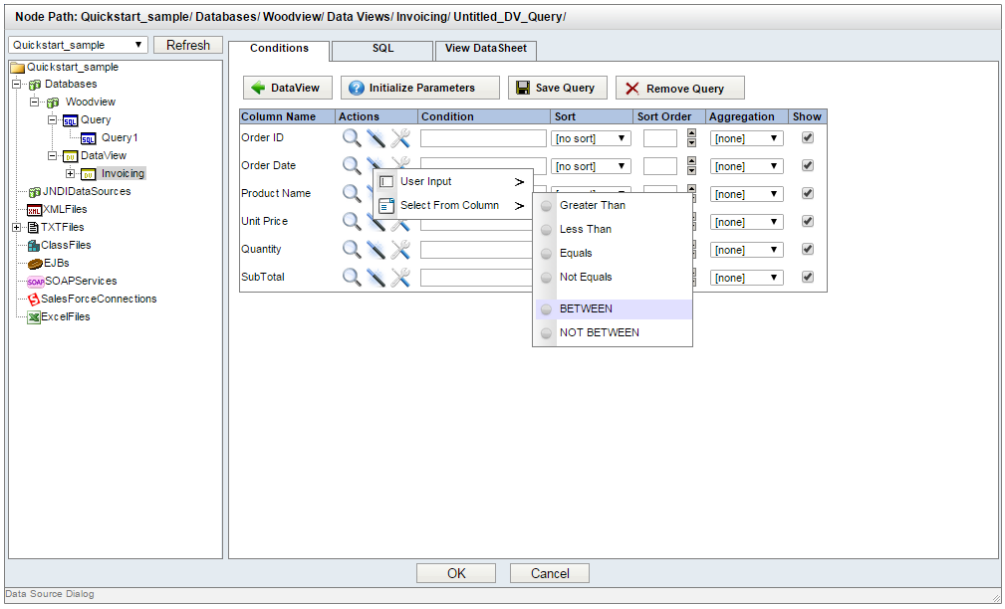
Renamed Column

Click the *Next* button to submit the column selection. The next dialog that opens allows you to set sorting, grouping, and conditions for the DataView query.



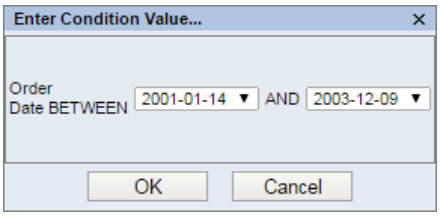
Conditions Dialog

Click the  *Condition Wizard* button for the `Order Date` field. Select the *Select From Column* option and then click the *Between* operator.



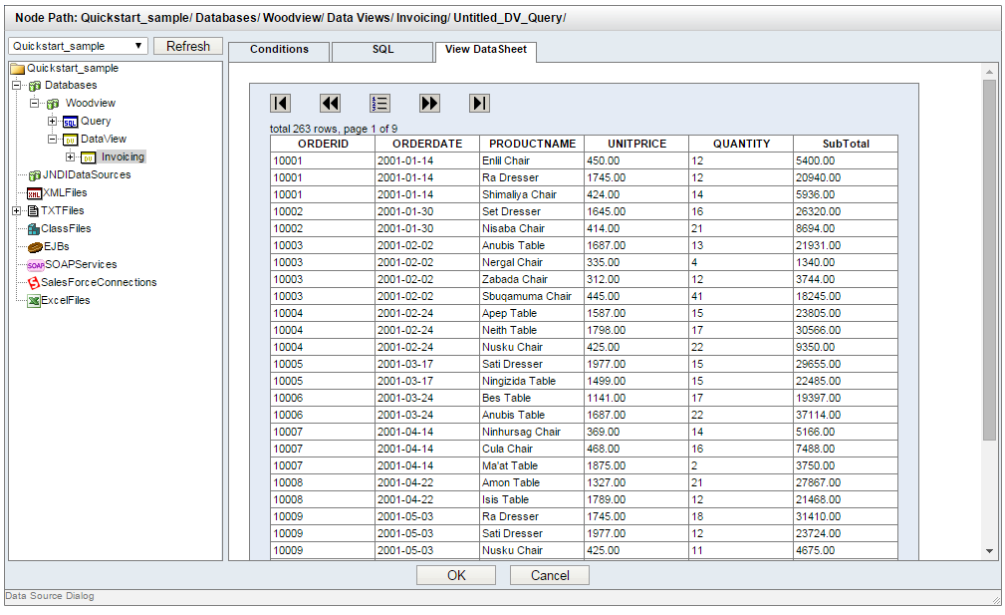
Condition Operators

This will open a dialog allowing you to select a date range by which to filter the query.



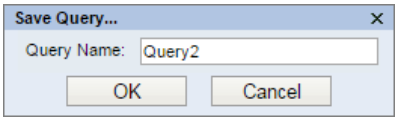
Select Condition Value Dialog

Select the date range that you would like and click *OK*. The condition will be added for the `Order Date` field. Now you can preview the finished query by clicking the *View DataSheet* tab.



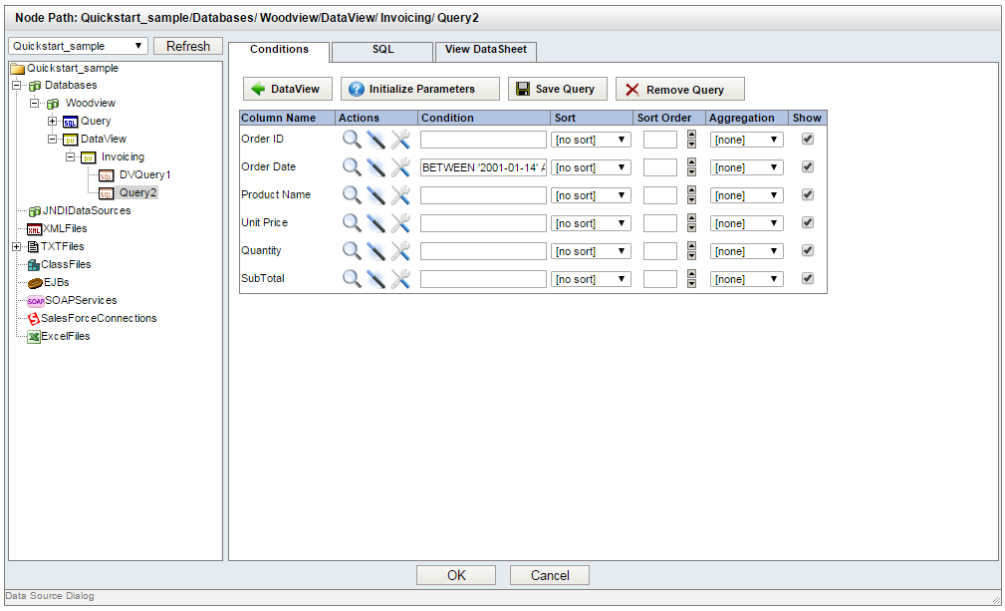
Data Preview

Next, go back to the *Conditions* tab and click the *Save Query* button to save this DataView query back to the registry. This will bring up a dialog prompting you to select a name for the query.



Save Query Dialog

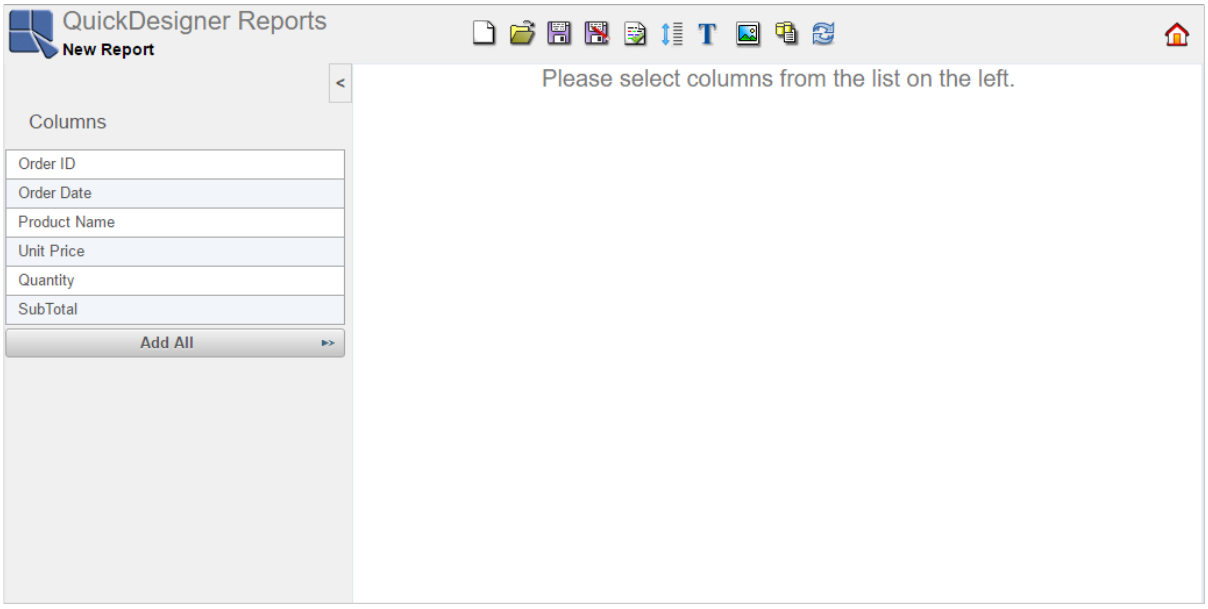
Enter the name you would like to use for the query and click *OK*. The new query appears under *DataView* node in the left pane of the *Data Source Dialog*.



New DataView Query

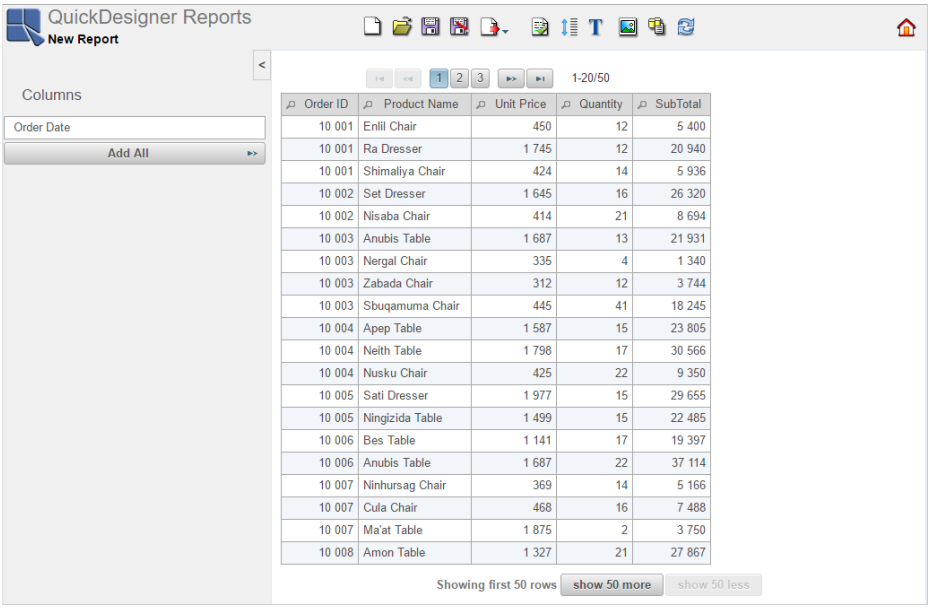
Q.6.2. Create a Report

Now that you have completed the query, click *OK* to close the *Data Source Dialog* and to open the main page of QuickDesigner Reports.



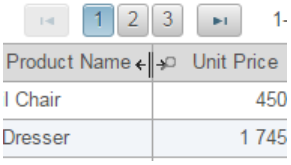
QuickDesigner Reports Main Page

You can insert any column to the report by using a double-click on the column or by dragging it to the report. So, insert columns *Order ID, Product Name, Unit Price, Quantity, SubTotal*.



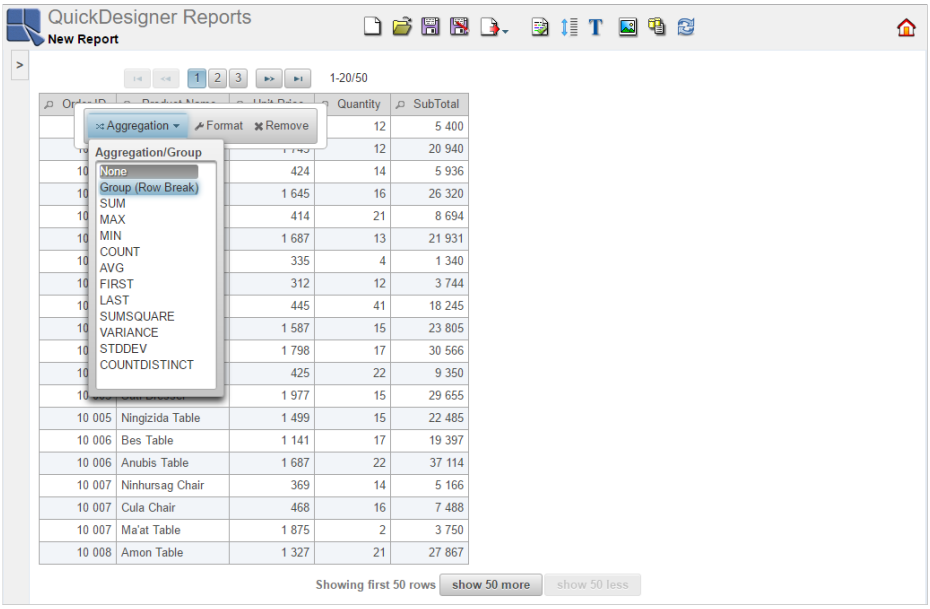
Columns Selecting

Collapse the left *Column* pane by clicking the *Collapse* button and adjust the width of columns. To do this, move your mouse over the right side of the column header. You will see a double arrow. Then left click and drag it.



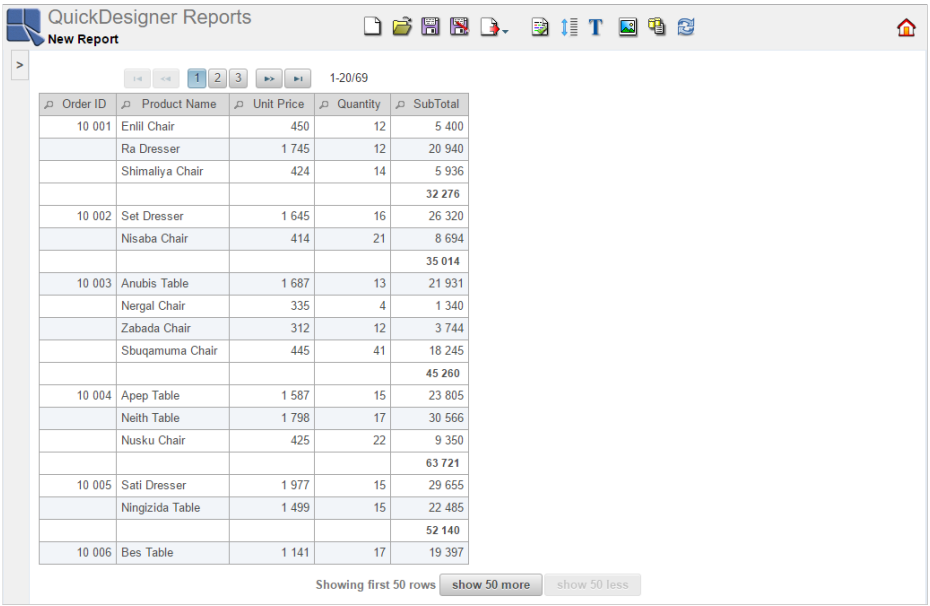
Column Width Setting

Right click on the *Order ID* header, select *Aggregation* and then *Group(Row Break)* option.



Aggregation Dialog

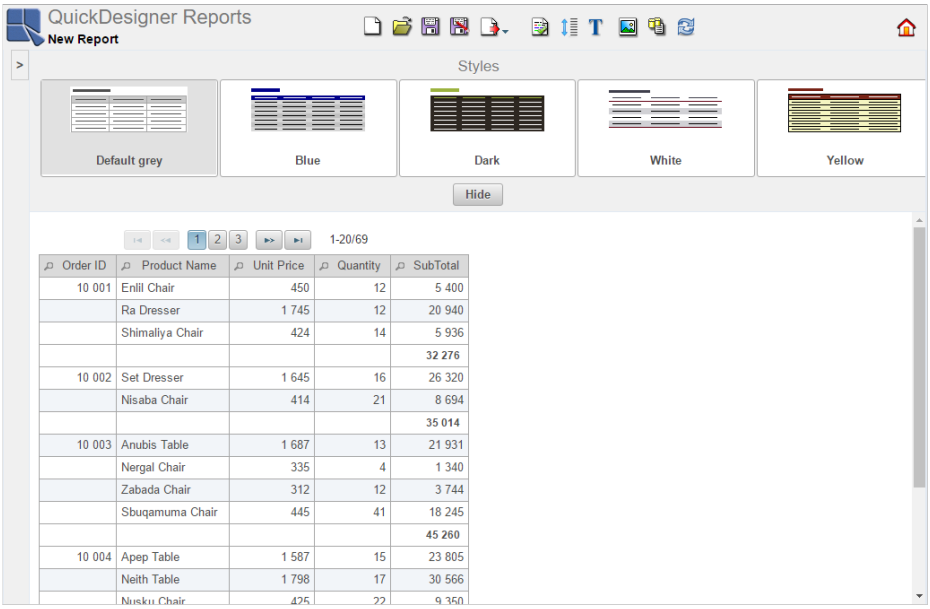
Right click on the *SubTotal* header, select *Aggregation* and then *SUM* option. Your report should look like this now:



Report with Aggregation

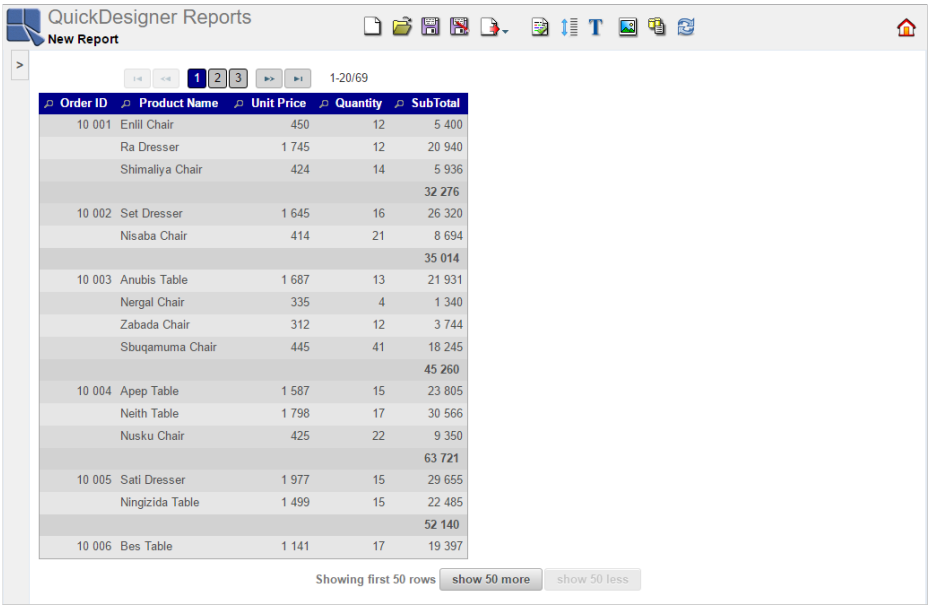
Q.6.2.1. Format Report Elements

QuickDesigner Reports allows you to change the style of the report by using several predefined styles. Click the  *Styles* button on the toolbar. The *Styles* dialog will open above the report.



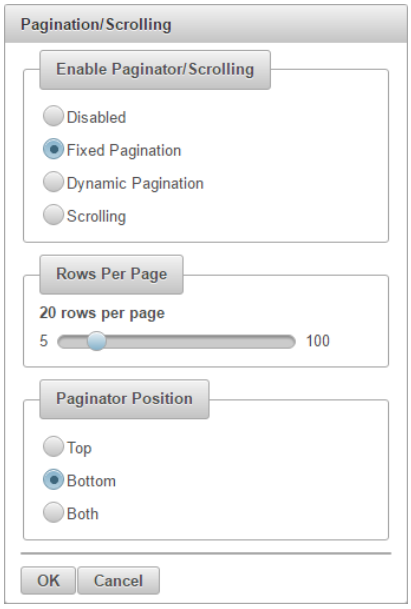
Styles Dialog

Select the style you like, e.g. the *Blue* style, and close the *Styles* dialog by clicking *Hide* button.



Style Applied

Next, place a pagination under the report. Click  *Pagination/Scrolling* button on the toolbar and *Pagination/Scrolling* dialog opens. Select *Bottom* option for *Paginator Position* and click *OK*.



Pagination/Scrolling Dialog

Now, add a report title. Click the  *Report Title* button on the toolbar and the *Report Title* dialog appears. Type **`Order Information`** as the title and click  *Apply* button.



Report Title Dialog

The report title appears on the left side of the report by default. Move it to the center of the report. Click on the title, hold the mouse button and move the mouse slightly. Dashed rectangles appear. Move the title to the middle rectangle and release the mouse button.



Title Placement

Your report should look like the report below. For more detail about all the formatting option available in Quick-Designer Reports, see Section 5.3.4 - Format Report Elements.

Order Information

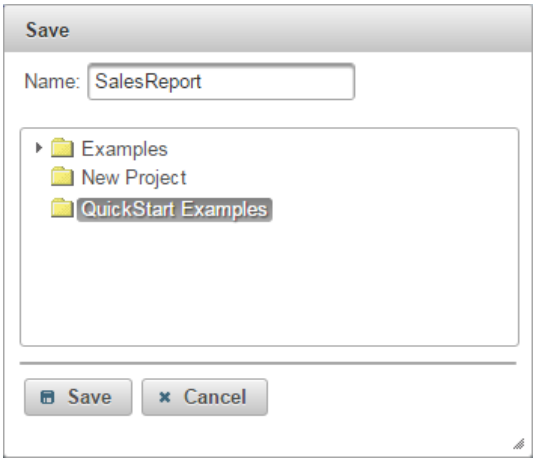
Order ID	Product Name	Unit Price	Quantity	SubTotal
10 001	Enlil Chair	450	12	5 400
	Ra Dresser	1 745	12	20 940
	Shimaliya Chair	424	14	5 936
				32 276
10 002	Set Dresser	1 645	16	26 320
	Nisaba Chair	414	21	8 694
				35 014
10 003	Anubis Table	1 687	13	21 931
	Nergal Chair	335	4	1 340
	Zabada Chair	312	12	3 744
	Sbuqamuma Chair	445	41	18 245
				45 260
10 004	Apep Table	1 587	15	23 805
	Neith Table	1 798	17	30 566
	Nusku Chair	425	22	9 350
				63 721
10 005	Sati Dresser	1 977	15	29 655
	Ningizida Table	1 499	15	22 485
				52 140
10 006	Bes Table	1 141	17	19 397

1 2 3 1-20/69

Report With Changes Applied

Q.6.2.2. Save the Report

Now that you have finished formatting, you can save the report you have created back to Organizer so that other users can view it. To save the report, click the *Save* button on the toolbar  . This will bring up a dialog allowing you to specify a name for the template.



Save Report Dialog

Enter a name for the report. Click the *Save* button to save the report.

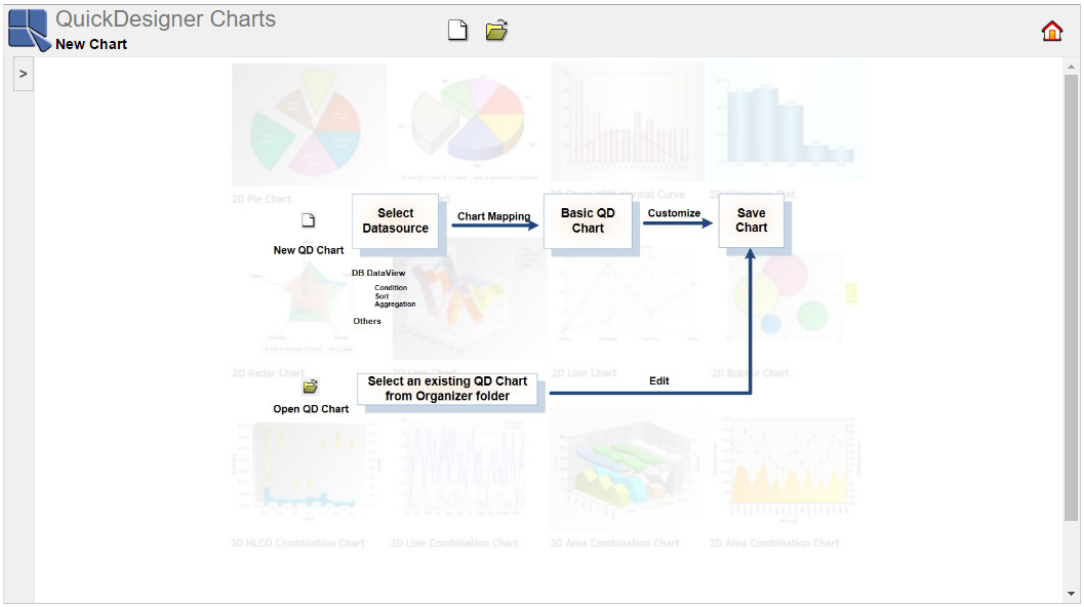
Q.7. QuickDesigner Charts

QuickDesigner Charts is a thin-client ad-hoc charting interface. It provides users a scaled-down design tool to create charts. With QuickDesigner Charts, end users can easily select, filter, and present data without mastering database structures, all with zero client download.

For more detailed information about all the chart options available in QuickDesigner Charts, see Section 5.4 - QuickDesigner Charts.

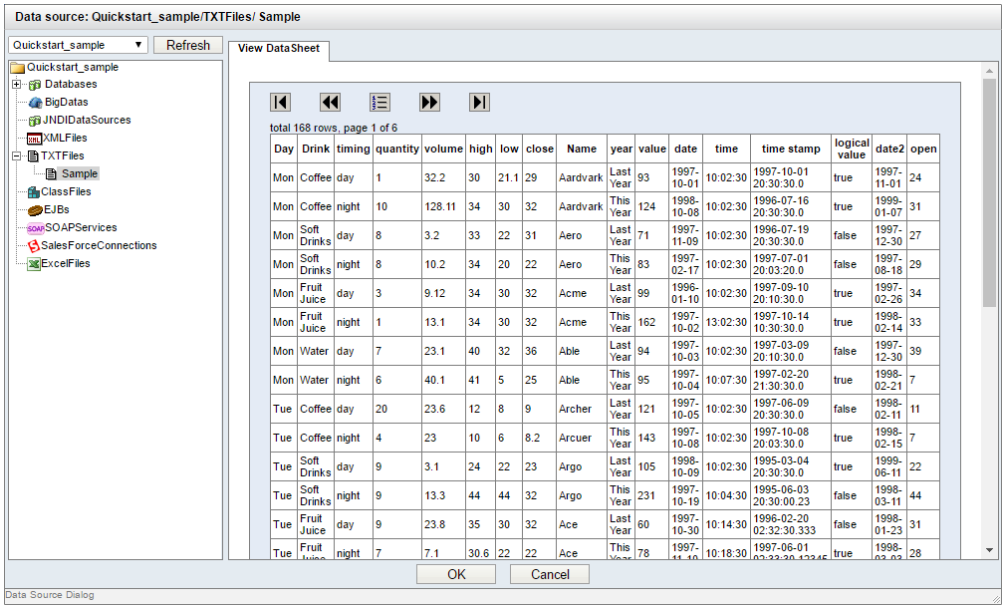
Q.7.1. Create a Chart

To start QuickDesigner Charts, go to the ERES Start page and click the link labeled *QuickDesigner Charts*. Quick-Designer Charts should open, and you can create a new chart or open an existing one.



QuickDesigner Charts

Click the *Create New Chart* button on the toolbar  . You will be taken to *Data Source Dialog*, where you can select a data source for the new chart. Select the `Sample.dat` text file data source. You can see records of the text file in the *View Data Sheet* window on the right side of the *Data Source Dialog*. Click the *OK* button to continue.



Data Source Dialog

You can see the *Chart Mapping* dialog that allows you to select the chart type and mapping options. Select to create a **`2D High Low`** chart using the first options in the dialog.

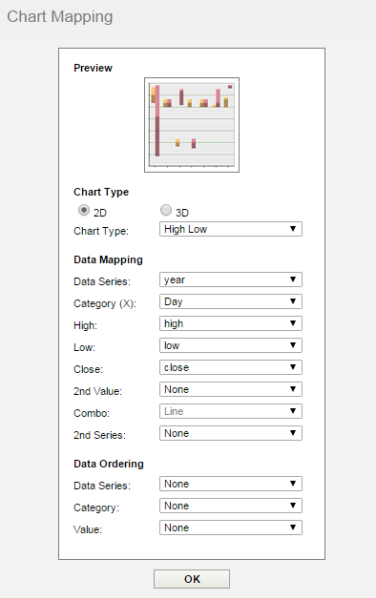
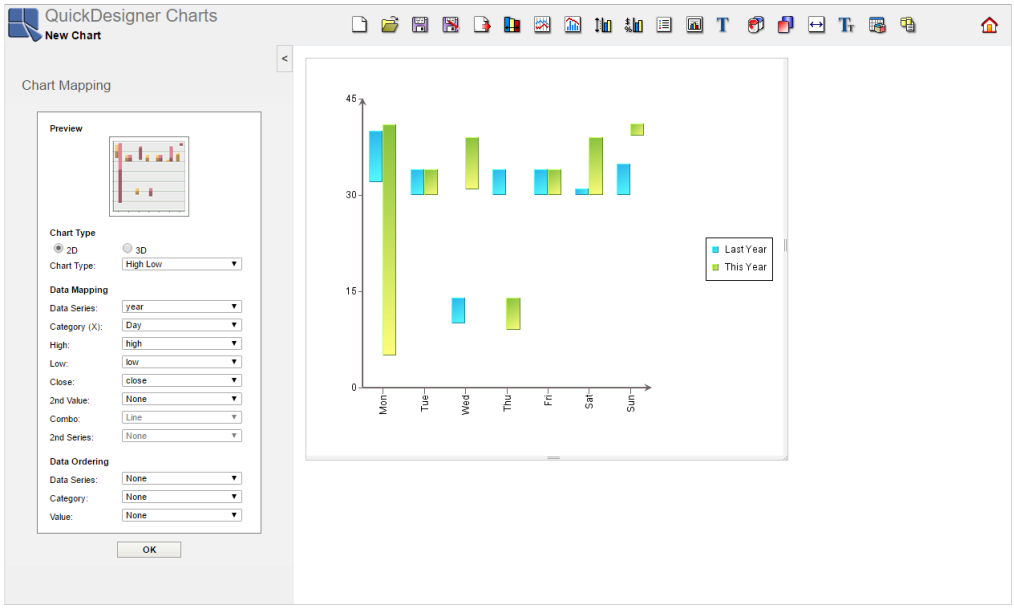


Chart Mapping Dialog

For the mapping options select **`year`** as the *Data Series*, **`Day`** as the *Category*, **`high`** as the *High*, **`low`** as the *Low*, and **`close`** as the *Close*. Click *OK* and the chart will open on the right side of QuickDesigner Charts.



QuickDesigner Charts Interface

Q.7.1.1. Format Chart Elements

Click the *Canvas Size* button on the toolbar  . This will bring up a dialog allowing you to set the size of the chart canvas. Set the size to be **`500`** by **`400`** pixels and click *OK*.

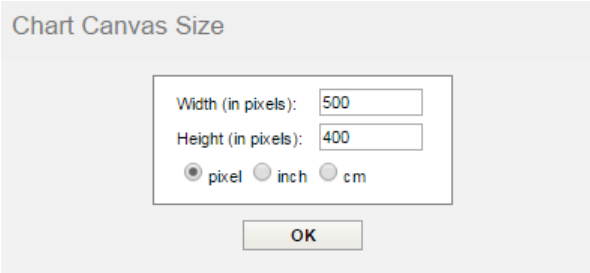


Chart Canvas Size Dialog

Next, click the *Font Settings* button on the toolbar  . This will bring up a dialog allowing you to set the font options for the different text elements in the chart. From the drop-down list in the top of the dialog select **`X Axis Labels`**. Then set the font to be **`Arial`**, **`8`** pt, and **`Plain`**.

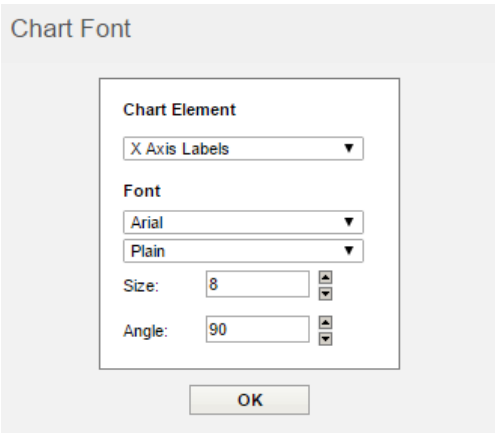
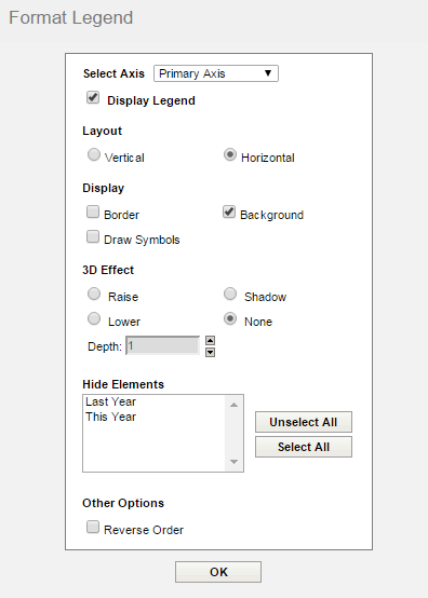


Chart Font Dialog

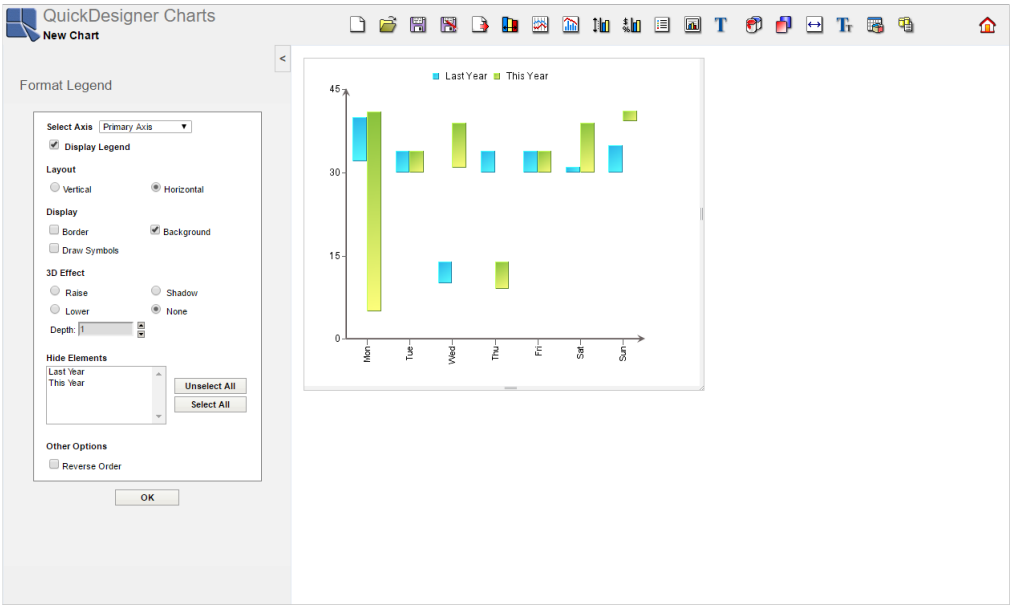
Now repeat this process for the Y axis labels and the legend text, setting them all to be **`Arial`**, **`8`**pt and **`Plain`**. Click *OK*.

Next, click the *Format Legend* button on the toolbar  . This will bring up a dialog allowing you to set display options for the chart legend. Change the *Layout* option to **`Horizontal`**, uncheck the *Border* option, check the *Background* option, change the *3D Effect* option to **`None`**, and then uncheck the *Background* option again. Click *OK* to apply the changes.



Format Legend Dialog

Now left click on the chart, hold the mouse button down and move the chart to keep the best position. Do the same with the legend. You can also set the plot area size by dragging the sides or dragging the corners of the chart. To do this, mouse over the side or the corner of the chart until you see the arrow. Then left click and drag.

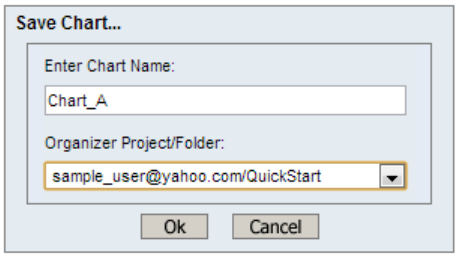


Finished Chart

For more detailed information about all the chart formatting options available in QuickDesigner Charts, see Section 5.4.5 - Data Formatting.

Q.7.1.2. Save the Chart

Now that you have finished formatting the chart, you can save it back to the Organizer. To save the chart, click the *Save* button on the toolbar  . This will bring up a dialog allowing you to specify a name for the chart.



Save Chart Dialog

Enter a name for the chart. Click *Ok* to save the chart. The window will then give you a message to indicate that the chart was saved successfully. Click *Ok* to close the dialog.

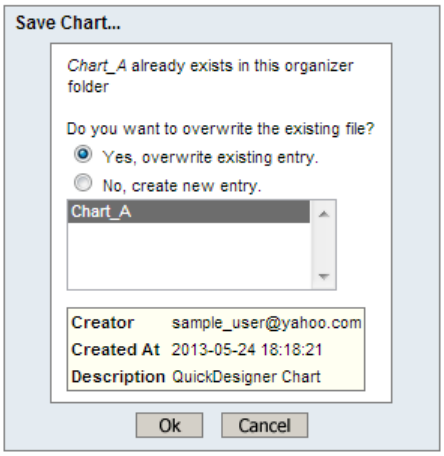


Chart Overwrite Dialog

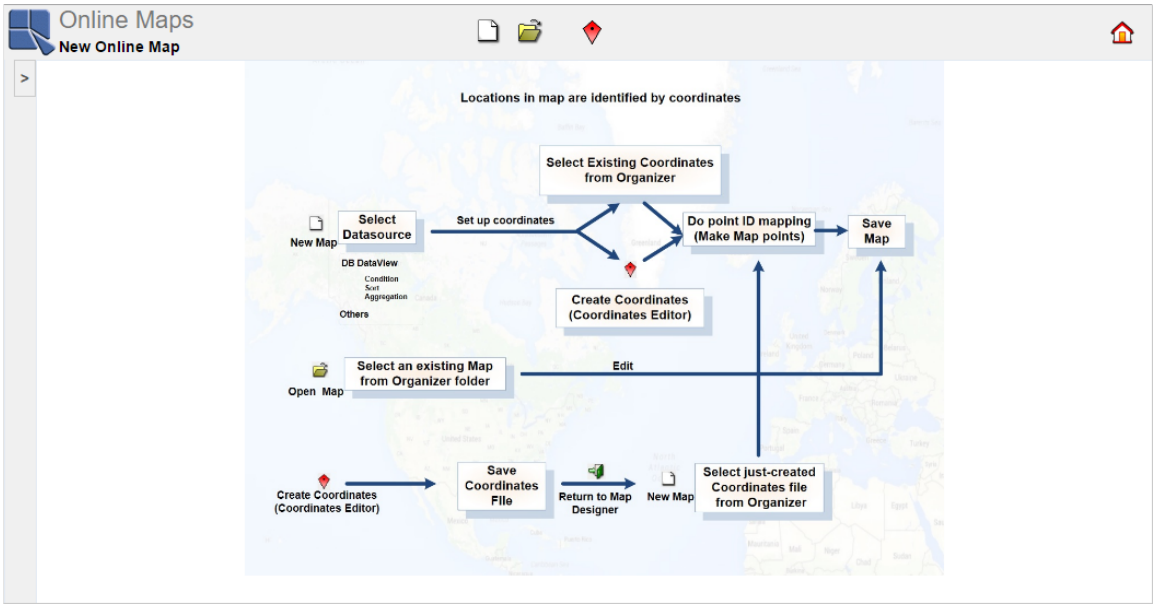
Q.8. Online Maps

An alternative way of presenting your data is by using the ERES maps feature. ERES maps are designed to report geographical data from data sources. They fetch geographical data from a data source and mark them on a map. There are two types of maps: Online Maps and SVG Maps. We will work with Online Maps at first and then with SVG Maps. (SVG Maps are described in Section Q.9 - SVG Maps.)

Q.8.1. Create Coordinates

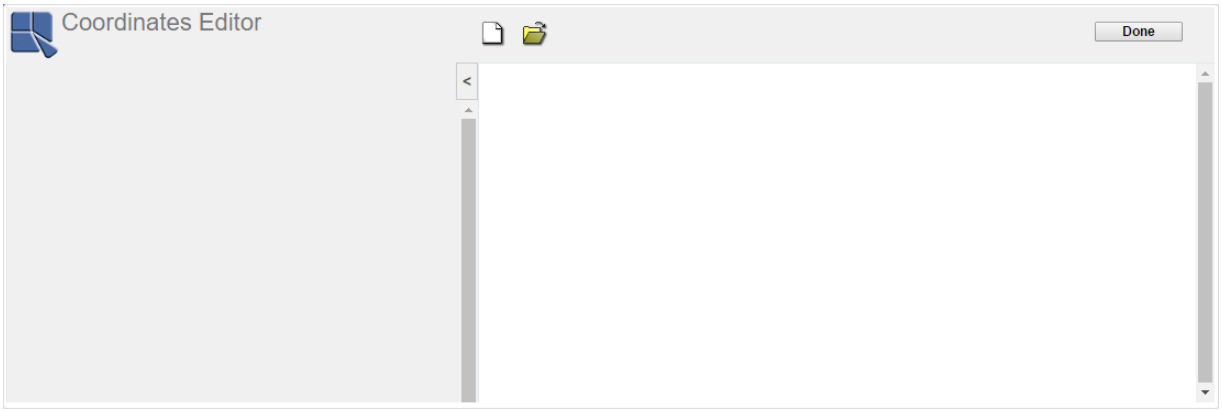
Before creating an Online Map, you will need to create a coordinates file. Coordinates files contain coordinates of places that you want to mark on the map. For more details about coordinates, see section Section 6.2.3.2 - Create Coordinates.

After launching Online Maps, you can create a new map, open some existing map, or create/edit coordinates.



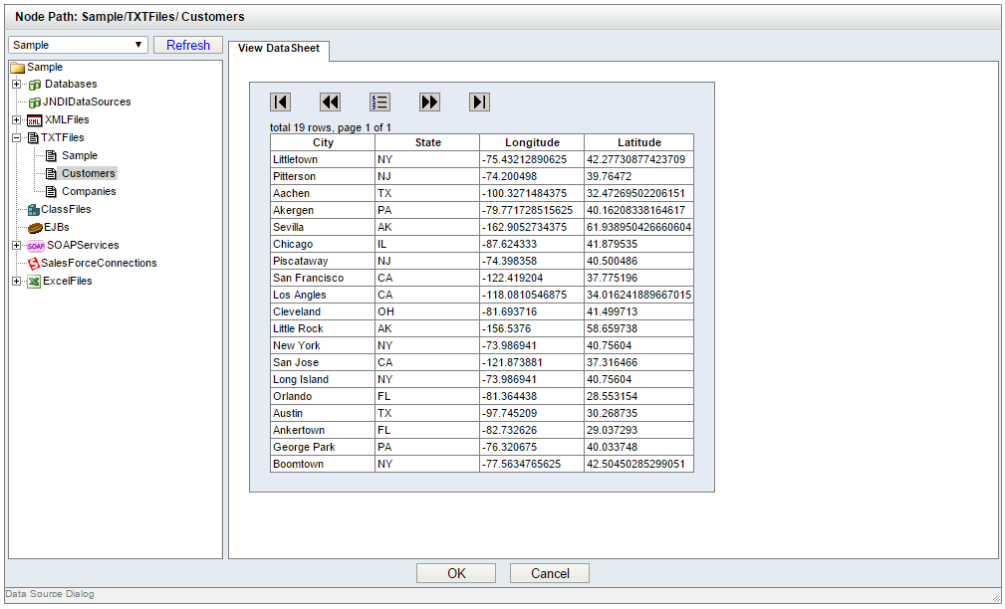
ERES Online Maps

Click the  *Coordinates Editor/Edit Coordinates* icon on the toolbar. The Coordinates Editor will open.



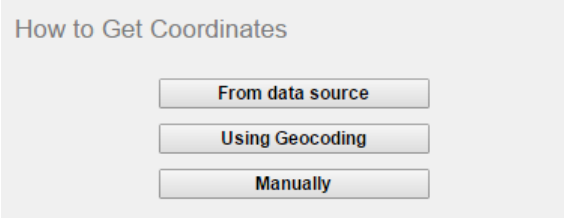
Coordinates Editor Start Options

Click the  *New Coordinates* icon. The *Data Source Dialog* will appear, allowing you to select a data source for the coordinates file. Click on the *TXTFiles* node to show all TXT data sources. Select the `Customers` data source. You can see a table displaying records of the selected data source in the right pane. It allows you to check whether you chose the right data source. Click *OK* to close the *Data Source Dialog*.



Selecting Data Source

You will be prompted to choose a method of obtaining coordinates from the data source. There are three methods for this purpose (to learn more about how to choose the right method, see section Section 6.2.3.2 - Create Coordinates). The best method for our `Customers` data source is to obtain the coordinates directly from the data source, so click the *From data source* option.



How To Get Coordinates

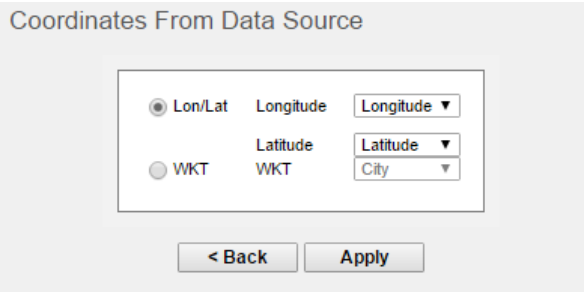
Next step is to select data source fields that contain the coordinates.



Note

This step would be different if you chose some other method of obtaining coordinates (all possibilities described in Section 6.2.3.2 - Create Coordinates).

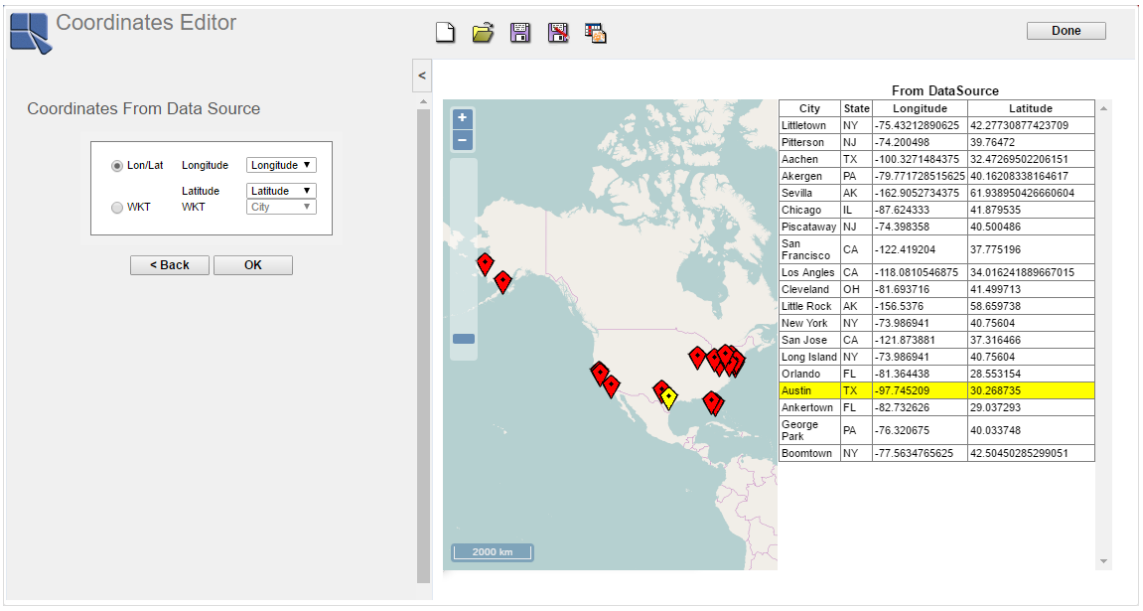
Each Coordinates data record consists of latitude and longitude (obtained automatically or manually) and a Point ID. The **Point ID** is an identifier of the place (e.g. city name, branch name, etc) and it is loaded from the data source you chose to create Coordinates from. The Point ID can consist of several **Point ID fields** - each data source column creates one field (except the fields, which contain longitude and latitude). Our data source has four columns `City`, `State`, `Longitude`, and `Latitude`. Therefore, the Point ID will consist of two fields - `City` and `State`. The upper option is labeled `Longitude` and the lower option is labeled as `Latitude`. Just set the upper field to data source field that contains information about longitude (in this case, it is the `Longitude` field) and the lower field to the data source field that contains latitude (`Latitude`). The following screenshot shows the correctly set dialog:



Select Fields With Coordinates

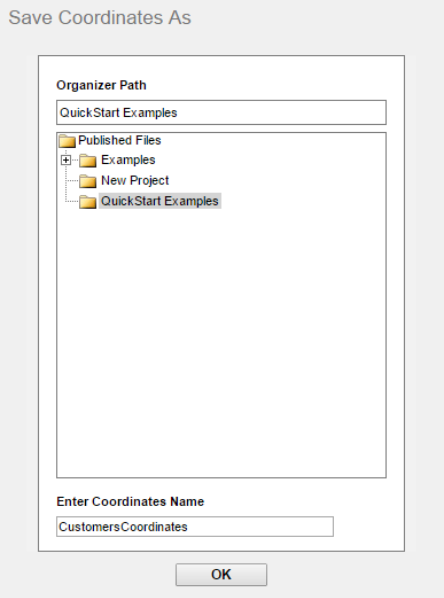
Your dialog should look like the one in the previous screenshot. If it does, click the *Apply* button.

You can see a map with coordinates table in the right pane of the Coordinates Editor. When you move the mouse arrow over one of the map markers or over a Point ID Table row, it will highlight the map marker and Point ID Table row that match together. (For more information see Section 6.2.3.3 - Coordinates Editor Interface).



Coordinates Editor

Now you may save the coordinates file. To save, click on the  *Save* icon on the toolbar.



Save Coordinates

Type **`CustomersCoordinates`** into the *Enter Coordinates Name* textfield and select **`your_email/QuickStart`** Organizer folder. Click *OK* to save the coordinates file.

Close the Coordinates Editor and go back to the Online Maps designer by clicking the *Done* button on the right side of the toolbar.

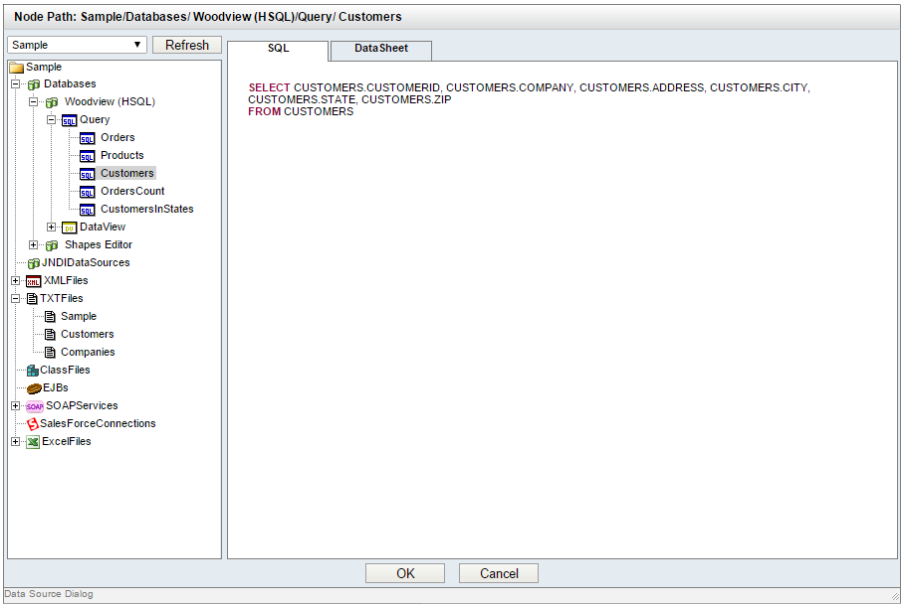
Q.8.2. Create Online Map



Note

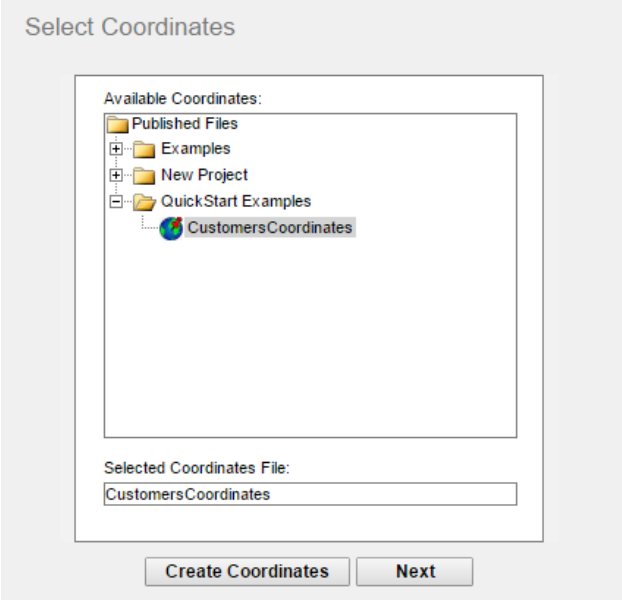
If you have not practiced creating coordinates, please do so. It may be extremely difficult to create an Online Map without proper coordinates.

Open the *Data Source Dialog* by clicking the  *New Map* icon on the Online Maps toolbar. We will use a different data source than the one we used for the coordinates in previous section. Open the *Databases* node, then the *Woodview (HSQL)* and the *Query* node and select the *Customers* query. Click *OK* to close the *Data Source Dialog*.



Selecting Data Source

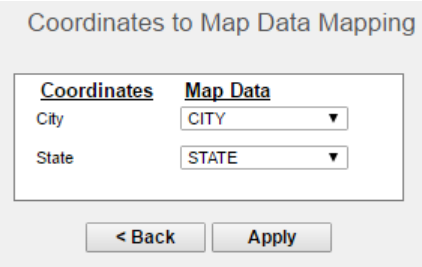
Now, you will have to select a coordinates file from the Organizer. You will be able to see all projects, folders, and coordinates that were inserted into ERES Organizer. To select the coordinates file we created in previous section, open the `QuickStart` folder under your project (click on it), then select the *CustomersCoordinates* coordinates and click *Next*.



Selecting Coordinates File

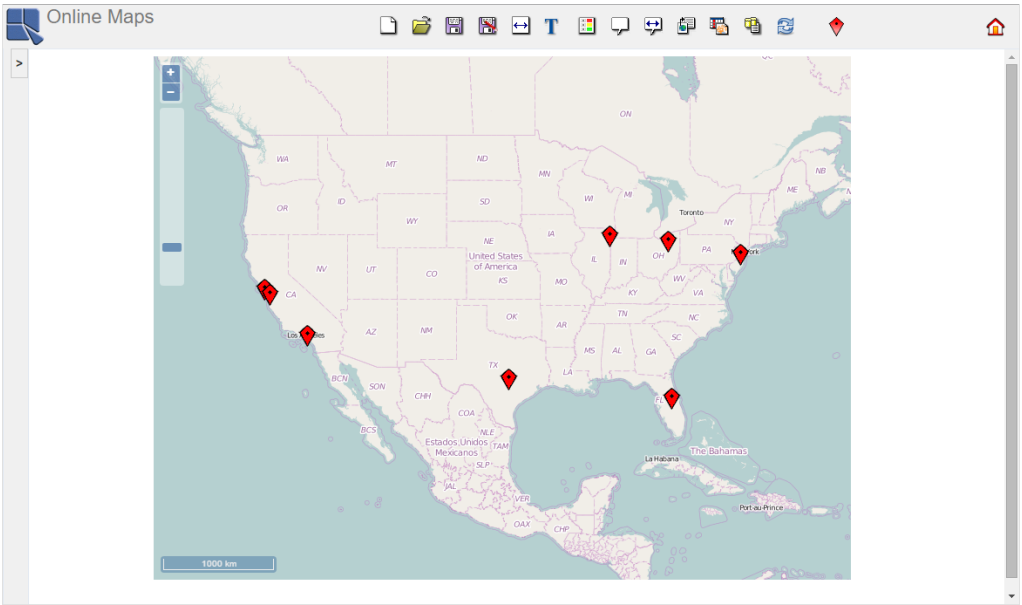
Next, you will need to set point mapping. To do this, you need to map at least one field from the ERES Online Maps data source to a Point ID field in the coordinates file data source. Both fields should contain the same type of information (for example, city name). But you can map more than one field if that is required to uniquely identify the location on the map (for example, city and state). When determining the map location of the ERES Online Maps data source record, the mapping function automatically searches for the same value of the mapped field(s) in ERES Online Maps data source and the mapped field(s) of the coordinates data source, and only the matched points will be visible on the map. In a nutshell, the coordinates data source contains information about location of data points on the map and ERES Online Maps data source contains the data that you want to report on the map. You may read more about point mapping in Section 6.2.4 - Coordinates Mapping.

In this tutorial, we will map the `City` field to the `CITY` field and the `State` field to the `STATE` field. Just map the fields and click *Apply*.



Set Point Mapping

Setting point mapping was the final step of the Maps Wizard. You should be able to see the Online Maps toolbar and the Online Map with some marker on it (the left pane is collapsed by clicking the < *Collapse* button on the screenshot).



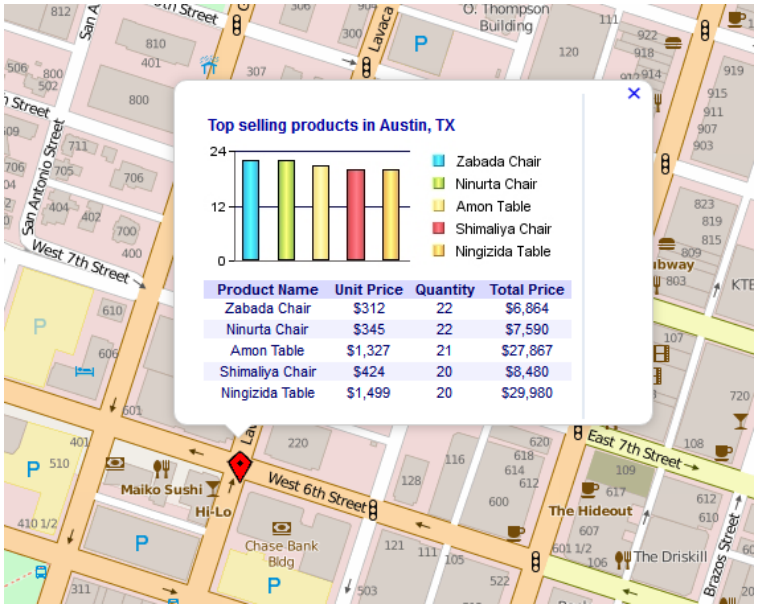
Online Maps Designer

The Online Maps toolbar allows you to configure the Online Map. You can immediately see results of all changes you make on the Online Map.

Save the map by clicking on the *Save* icon on the toolbar. Type **`CustomersOnlineMap`** into the *Enter Map Name* textfield. Select the **`your_email/QuickStart`** Organizer folder and click *OK* to save the map.

Q.8.2.1. Configure Tooltips

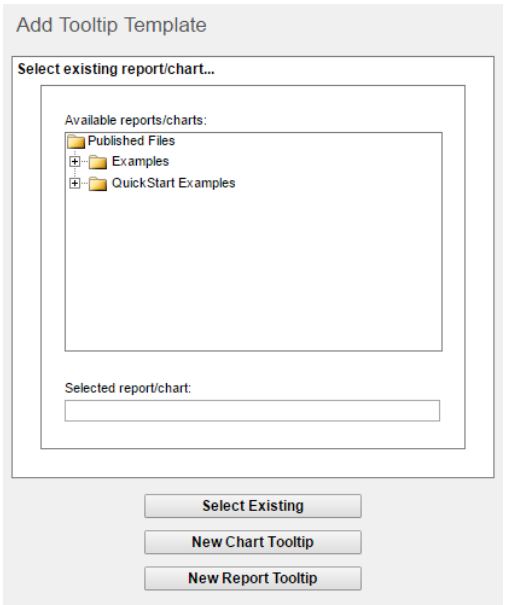
Tooltips show a brief report for a particular map marker. When the tooltips are enabled, each marker on the Online Map has its own tooltip. After setting the tooltip template (report or chart) for the Online Map, upon mouse over any map marker, the tooltip bubble will appear, displaying data relevant to the particular marker. Detailed information about the tooltips can be found in Section 6.2.5.7 - Tooltips.



Example: Google Map with Tooltip

Lets configure the tooltips for our CustomersOnlineMap that we created in the previous QuickStart section.

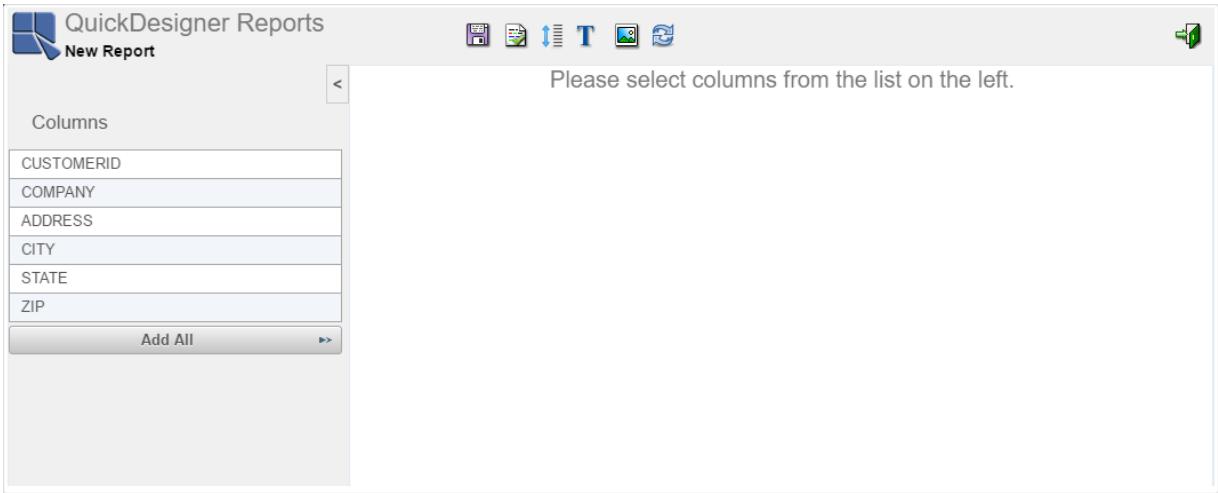
Click on the  *Tooltip template* icon on the toolbar. The *Add Tooltip Template* dialog will appear.



Add New Tooltip

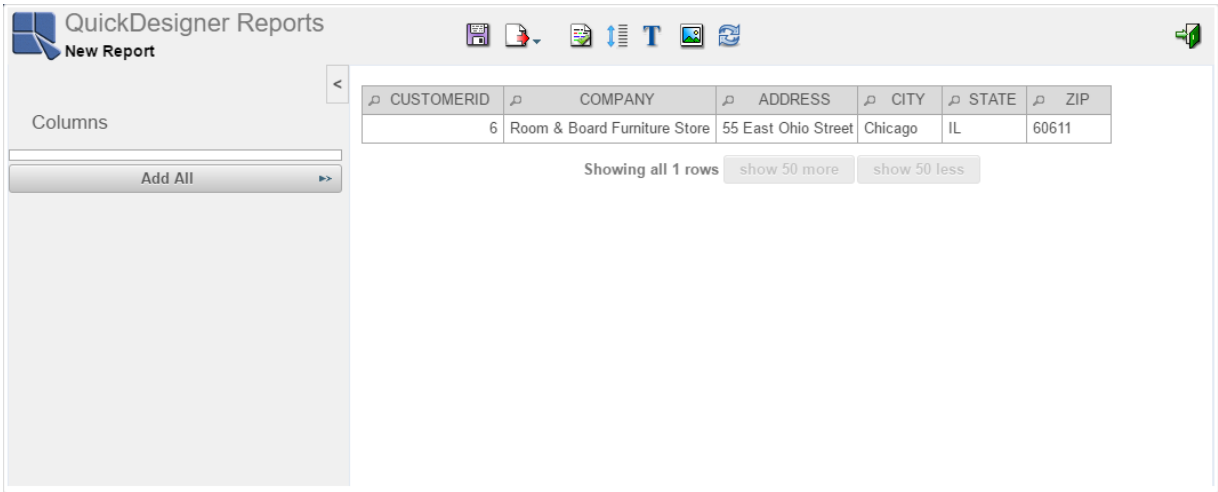
You can select an existing tooltip template or create a new chart or report tooltip template. When you choose to select new chart template, QuickDesigner Charts will launch. When you choose to select new report template, QuickDesigner Reports will launch. The launched QuickDesigner Report or QuickDesigner Charts will use the Map Data as a datasource to create a report or chart as a tooltip. In our example, Customer Query under Woodview(HSQL) Query node is the datasource used by Quick Designer to create tooltip. For now, we will just make a very simple report in QuickDesigner Reports. Visit Section Q.6 - QuickDesigner Reports, or the Section 5.3 - QuickDesigner Reports for information about how to use QuickDesigner Reports.

Click the *New Report Tooltip* option to open QuickDesigner Reports.



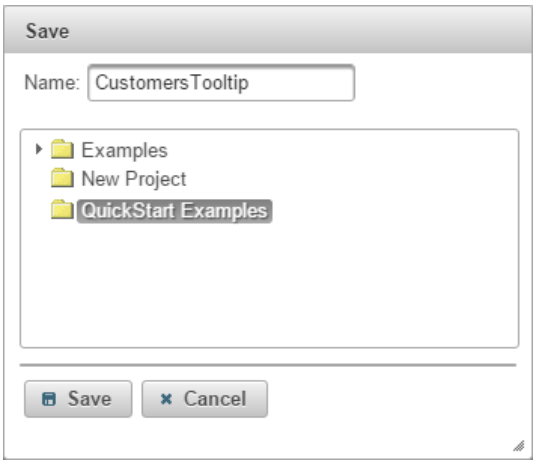
QuickDesigner Reports

Click *Add All* to add all columns to the report.



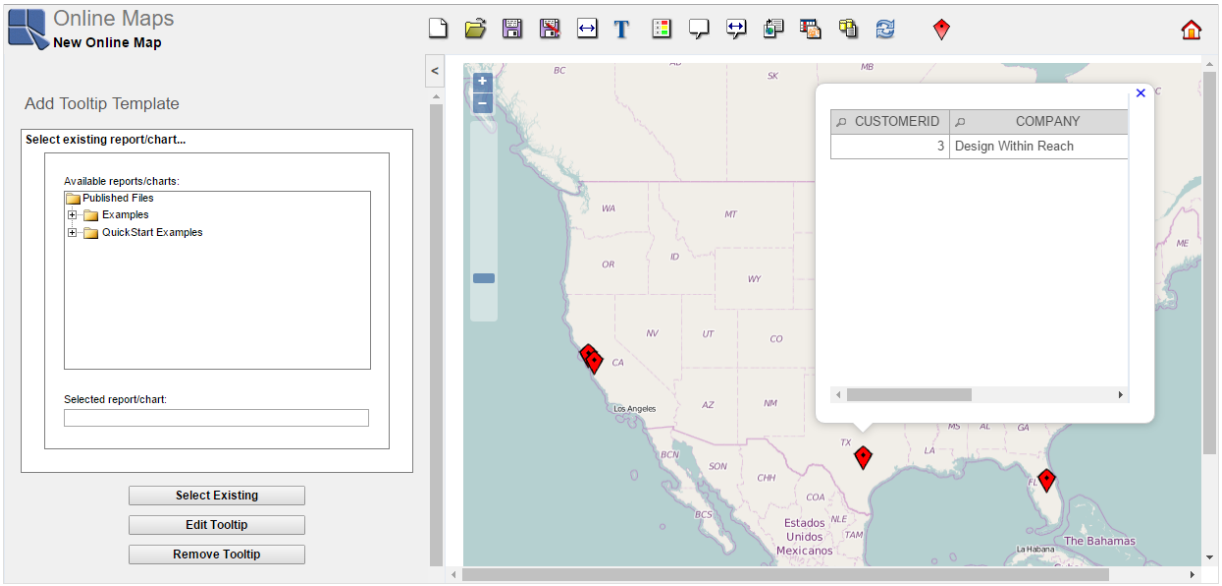
Tooltip Report

Click on the  *Save* icon on the QuickDesigner Reports toolbar. Enter **`CustomersTooltip`** as the *Name* and select the *QuickStart Examples* Organizer folder. Click *OK* to save the report.



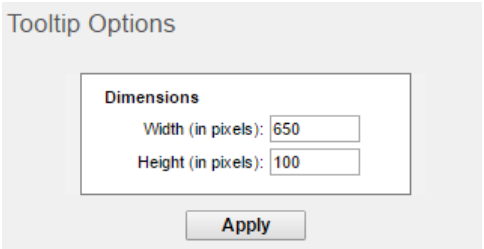
Save Report Dialog

Close QuickDesigner Reports by clicking the  *Return to Map Designer* icon on the toolbar.



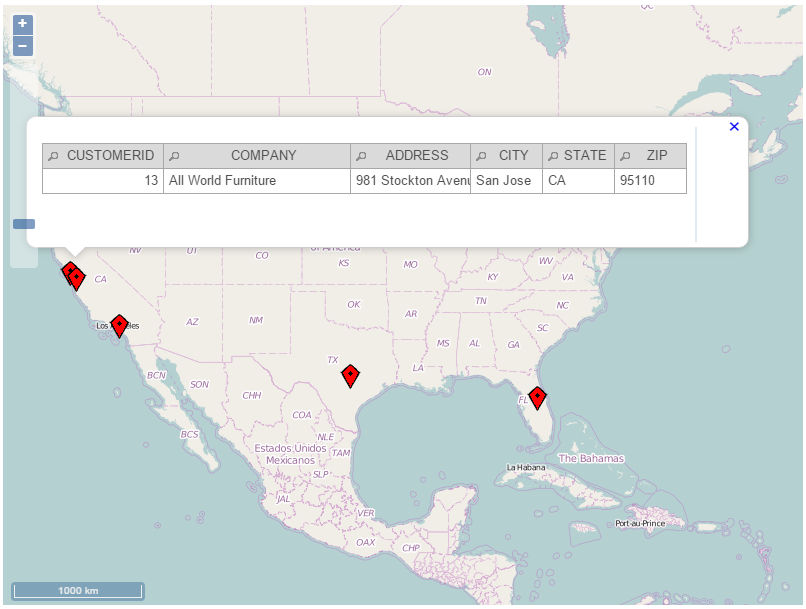
Adding Tooltip

Move the mouse arrow over some map marker and a tooltip bubble will appear. As you can see, the bubble has inconvenient dimensions for your report. To adjust the bubble dimensions, click on the  *Tooltip Options* icon on the toolbar. Set width to **650** and height to **100** and click *Apply*.



Tooltip Option Dialog

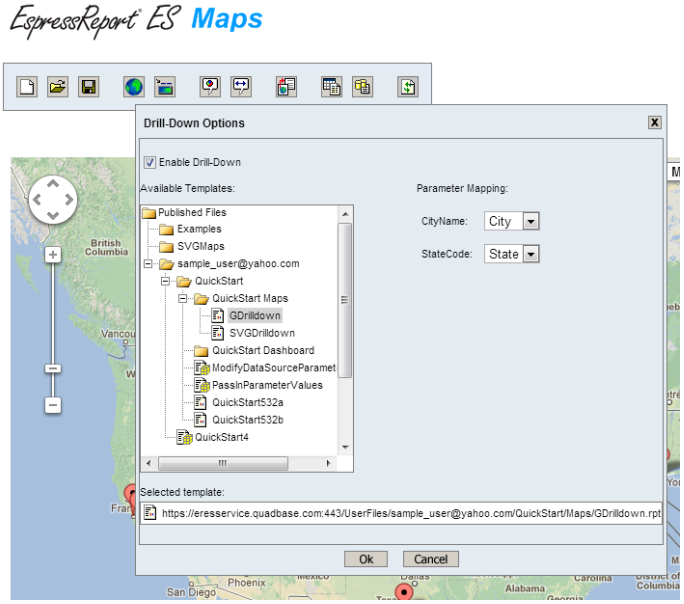
Close the currently open tooltip bubble (if any) and move mouse arrow over some marker. The tooltip should look better now.



Online Map with Tooltip

Q.8.2.2. Setting Drilldowns

Click on the  *Drilldown Options* icon on the toolbar. Select the *Enable Drilldown* option. The *Available Templates* treeview will open, showing all projects, folders and parameterized reports, charts and maps inserted in the Organizer. Open your project, then open the `QuickStart` and `QuickStart Maps` folder. Select the `GDrilldown` report. This report has two parameters: `CityName` and `StateCode`. You have to select the Online Map database fields containing corresponding data. In this case, map the `CityName` parameter to the `City` field and the `StateCode` parameter to the `State` field. Click *Apply*. The drilldowns are configured now. Click on a map marker, report will open in a new window. More info in Section 6.2.5.9 - Drilldowns.



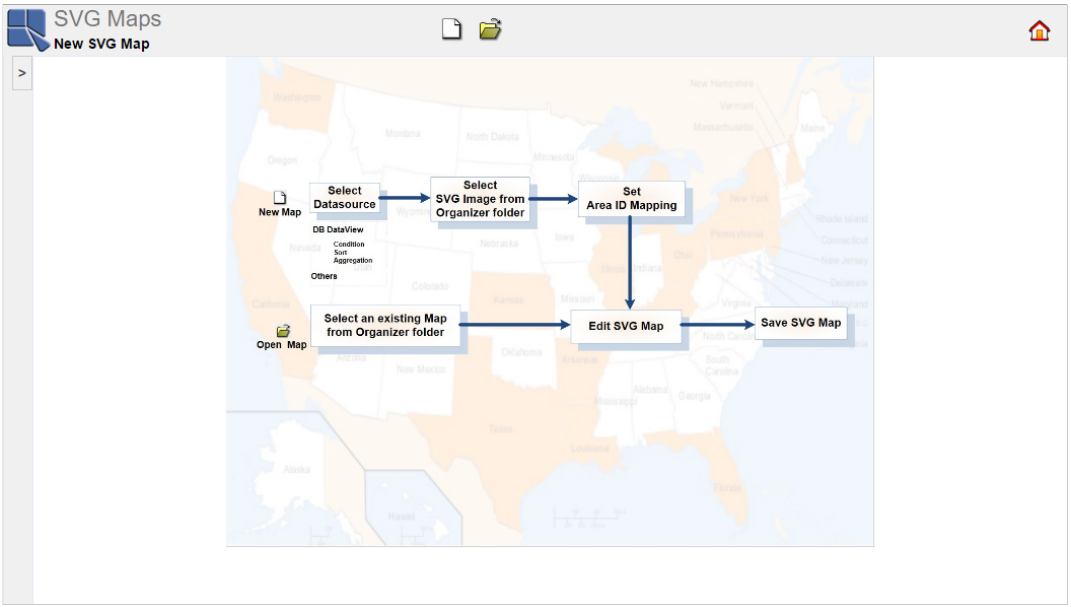
Drilldown Options

Q.9. SVG Maps

Unlike Online Maps, SVG Maps do not use map points. They use **map areas**. The map areas can be colored according to some values from the map data source. Coordinates of the map areas are defined in **SVG map image**. For more information see Section 6.3 - SVG Maps.

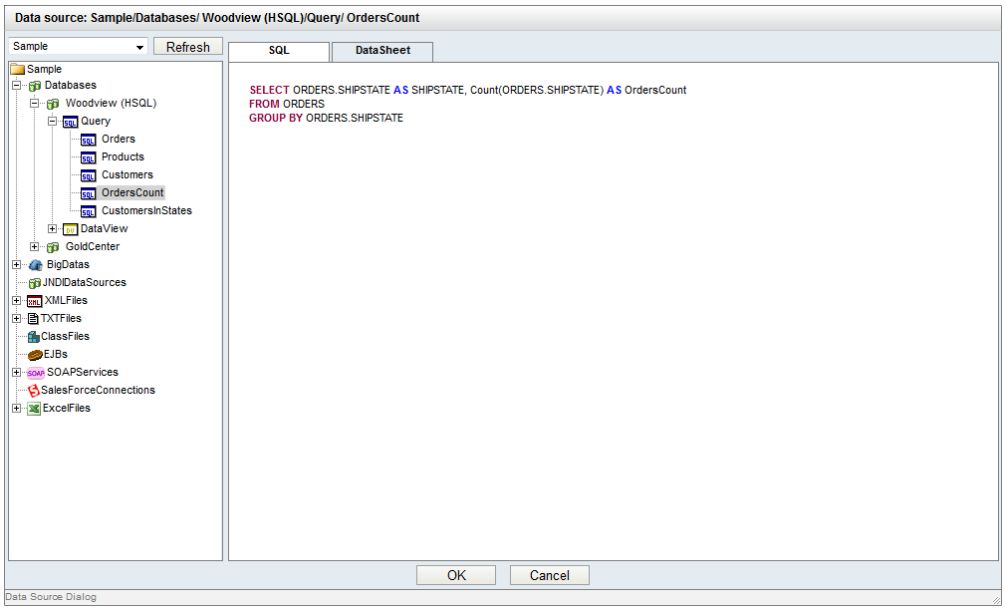
Q.9.1. Create a Map

To create a SVG Map, go to the ERES Start page and click the link labeled *SVG Maps*. You can create a new SVG map or open some existing SVG map.



New/Open Selection

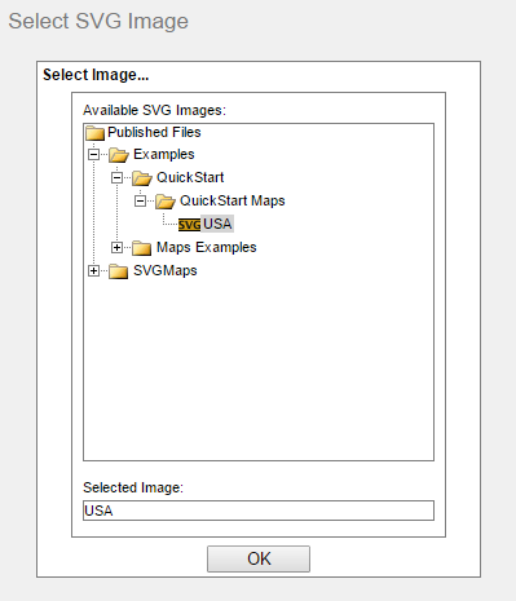
Click the  *Create New Map* button on the toolbar. You will be taken to the *Data Source Dialog*, where you can select a data registry and a data source for a new SVG Map. Select the *Sample* data registry (the left top corner of the window). Open the `Databases/Woodview (HSQL)/Query` node, and select the `OrdersCount` query.



Data Source Dialog

You can see records of the data source file in the *DataSheet* tab. Then click *OK* to close the *Data Source Dialog*.

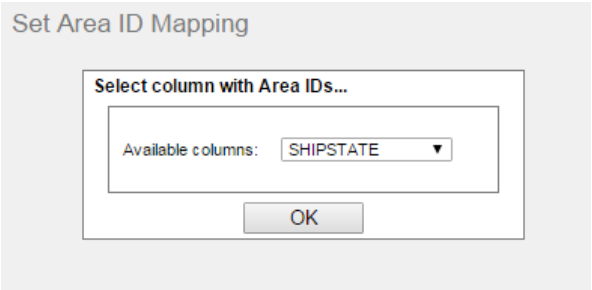
The next dialog allows you to select a Map Image. Map Images are SVG image files with data structures containing geographic data(see Section 6.3.2 - Area ID Mapping for more information about required image format), and must be inserted into the Organizer before they can be used.



Select SVG Image Dialog

The *Select SVG Image* dialog contains a tree showing all Projects, Folders and SVG images inserted in the Organizer. Open the `Examples` project, then the `QuickStart` and the `QuickStart Maps` folder. You should be able to see the SVG image named `USA` in the `QuickStart Maps` folder. Select it and click *OK*.

Now, you can set the connection between the data source and the SVG image. This is done by assigning a data source field to the SVG map Area IDs (see Section 6.3.2 - Area ID Mapping to learn more about area ID mapping). In this case, we have a map of the USA with state name abbreviations in its Area IDs. So we want to assign a field containing the same type of data, which is the *SHIPSTATE* field. Set the *Available columns* drop-down menu to *SHIPSTATE* and click *OK*.



Select Column with Area IDs

SVG map will open (*Set Area ID Mapping* dialog is colapsed on the screenshot).

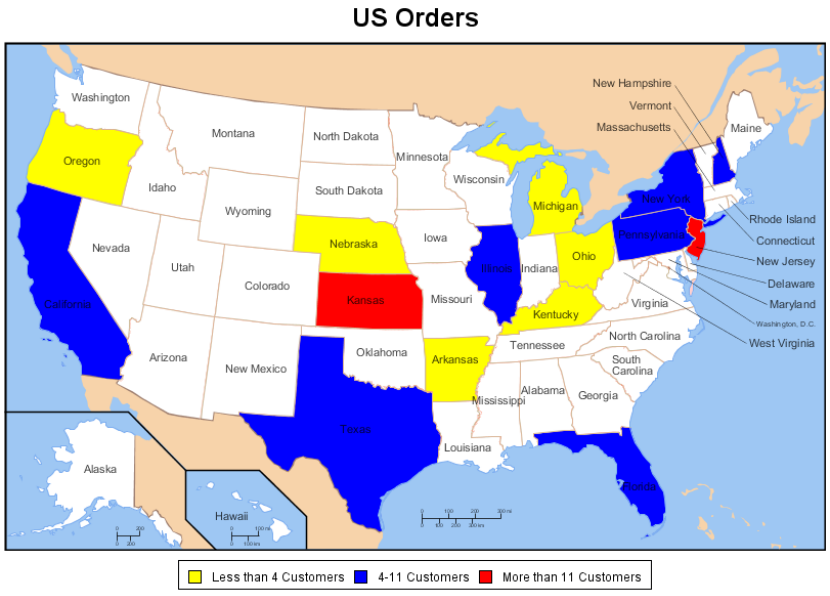


SVG Map Builder

As you can see, right now we have only the basic map with no data highlighted. There are two ways of adding data to the SVG Map: Thresholds and Drilldowns.

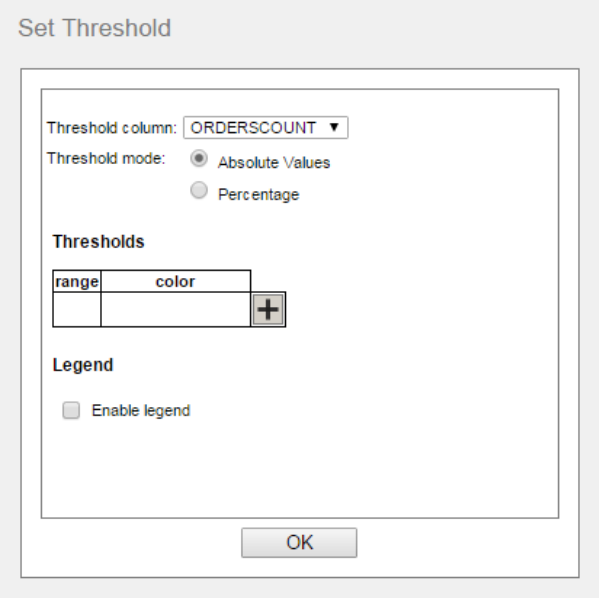
Q.9.2. Set SVG Map Thresholds

Thresholds allow you to configure rules for coloring map areas according to its value in a particular data source field (see Section 6.3.3.5 - Thresholds).



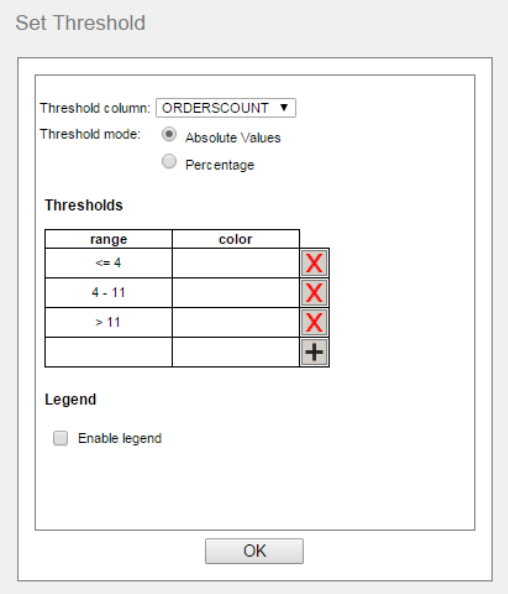
Example: SVG Map with Thresholds

To set thresholds, click on the  *Set Thresholds* icon on the toolbar. First of all, you have to select a data source column containing numerical data that you want to report. We will use a column labeled as **ORDERSCOUNT**, which holds information about number of orders made in a particular state.



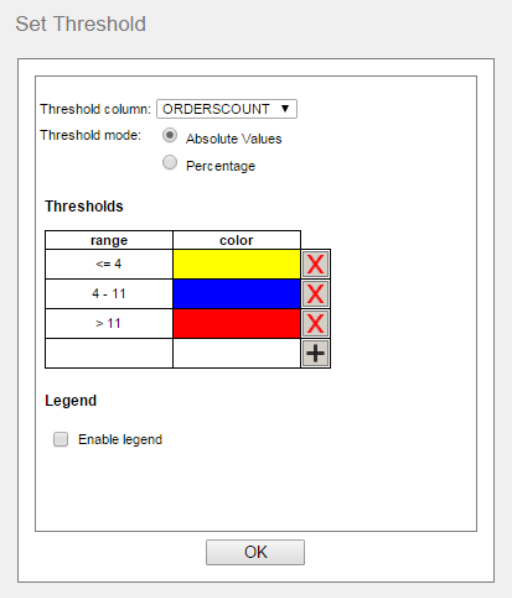
Set Threshold Dialog

Next, you can add some threshold value(s) in actual values, or percentage unit. We will use actual values for our example. Click on the *Plus* icon next to the blank field labeled as *color*. You will be prompted to enter threshold value, which can be done by two ways – you can type it manually into the text field or you can select the value from the drop-down menu. (There are three ways to enter threshold values for percentage units. The third way is to enter a number of intervals.) Select the value **4** from the drop-down menu and confirm the dialog by clicking *OK*. Click on the *Plus* icon once again, and choose the value **11** from the drop-down menu and click *OK*. ERES has automatically set the ranges between values you added.



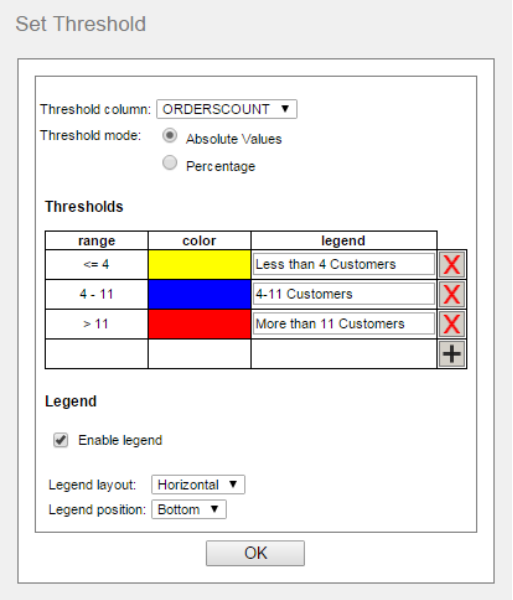
Setting Range

You can set color for each of the intervals. Click on the white color field next to the <=4 range. Click on some yellow field in the swatches color table and click *OK*. Color of the <=4 range has changed to yellow. Click on the color field next to the 4-11 range, choose a blue color and click *OK*. Click on the color field next to the >11 range, choose a red color and click *OK*. The Set Thresholds dialog should look like the one on the following screenshot:



Setting Color

Confirm the dialog by the *OK* button. Areas of some states will be filled with the colors you previously set. Without a legend, it can be unclear what the colors mean. So lets create one. Select the *Enable legend* option in the Set Threshold dialog. The *Thresholds* table has changed. It has one more column now labeled as *legend*. You should be able to write into the legend field. Fill in the column according to the following screenshot.

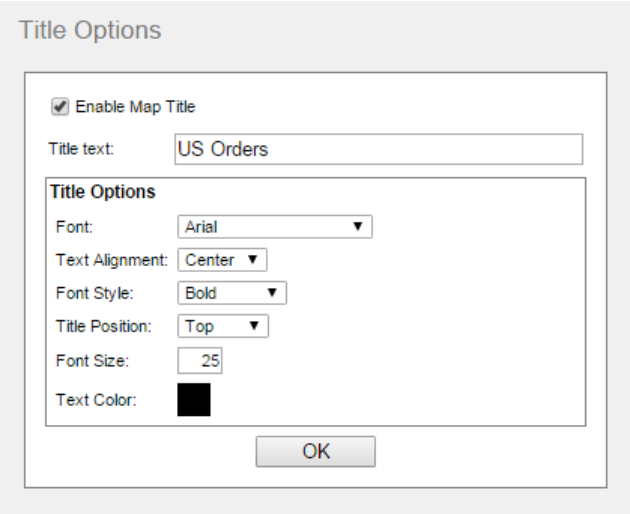


Set Thresholds- Legend

Set the *Legend layout* drop-down menu to **`Horizontal`** and set the *Legend position* to **`Bottom`**. Click *OK*. Then, without map title, you have the same SVG Map image shown in the begining of this part.

Q.9.2.1. Add a Map Title

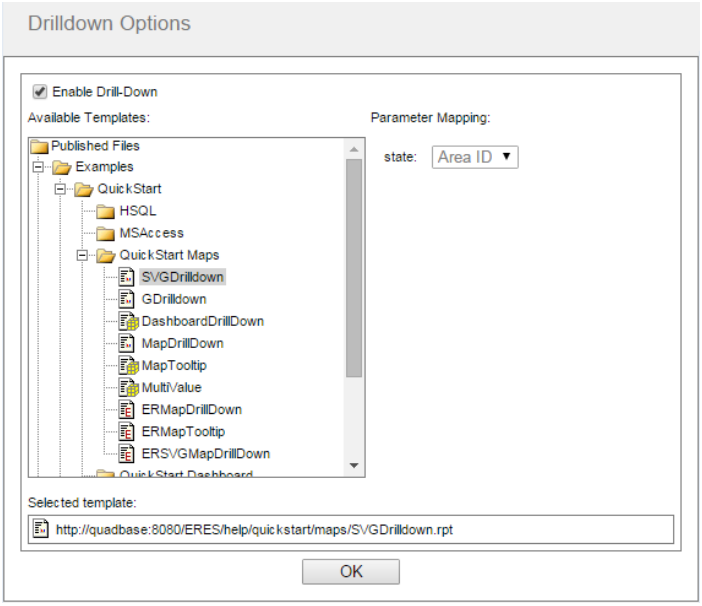
Click on the  *Map Title* icon on the toolbar. Select the *Enable Map Title* option. Type **`US Orders`** into the *Title text* field, set *Font Style* to **`Bold`** and confirm by the *OK* button.



Map Title

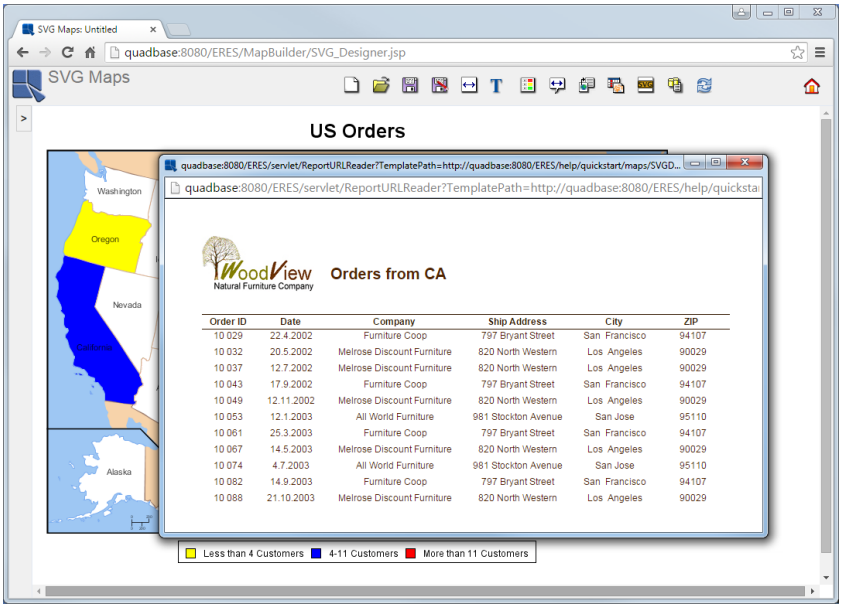
Q.9.3. Set SVG Map Drilldowns

Click on the  *DrillDown Options* icon on the toolbar. Select the *Enable Drill-Down* option. Organizer structure will show up in *Available Templates* treeview. You will be able to see only folders and parameterized reports, charts and maps inserted into Organizer. Open your project, then the `QuickStart/QuickStart Maps` folder and select the `SVGDrilldown` report. Click *OK*.



DrillDown Options

When you click on a colored state area, a drilldown report should open in a new window.



Map with Drilldown

To save the SVG Map, click on the  *Save* icon on the toolbar. Type **`OrdersSVGMap`** into the *Enter Map Name* text field and select the `your_email/QuickStart` folder from the Organizer treeview.

Q.10. Publishing

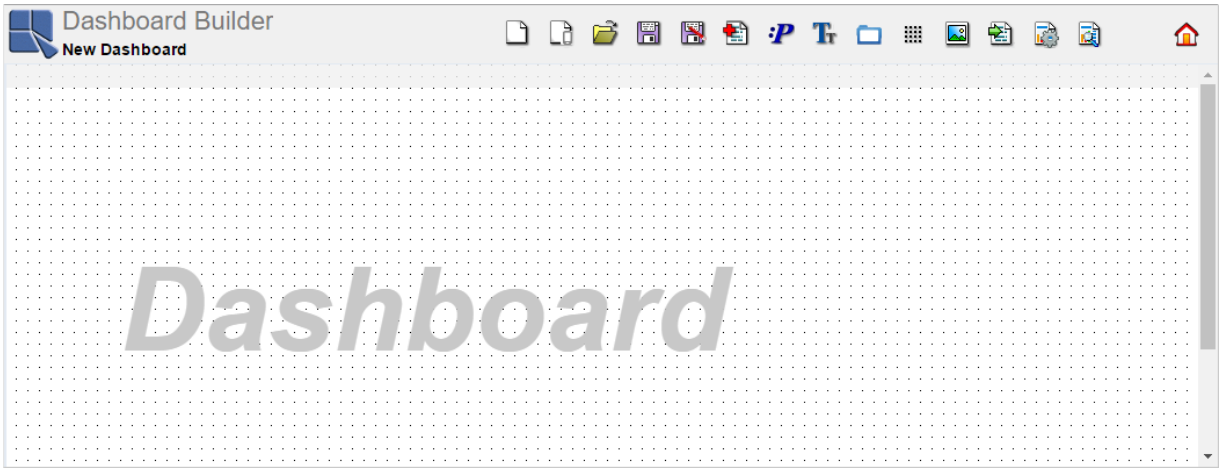
In the previous sections we looked at setting up ERES and creating reports, charts, and maps. This chapter discusses the automated publishing features in ERES.

Q.10.1. Dashboards

In addition to URLs and the menu page, ERES also allows you to publish reports and charts using a dashboard interface. Dashboards can place multiple charts and report tables into a single presentation page. Users can define common filters for the dashboard items and set up drill-down for individual dashboard items.

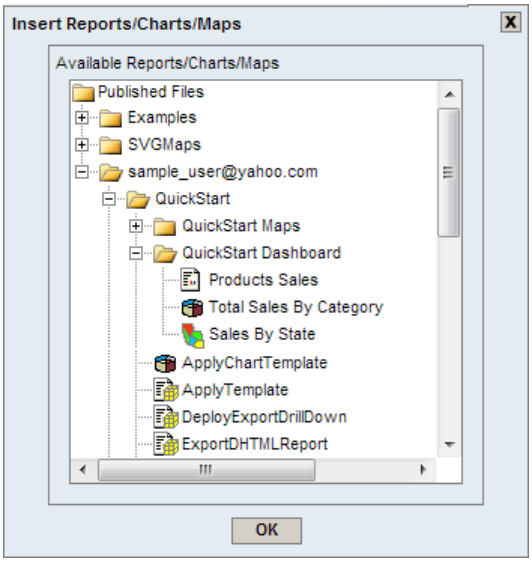
Q.10.1.1. Create a Dashboard

Dashboards are created in the thin-client Dashboard Builder interface. To launch the Dashboard Builder, click the *Dashboard Builder* link in the ERES start page. The interface will allow you to build a new dashboard.



Dashboard Builder Interface

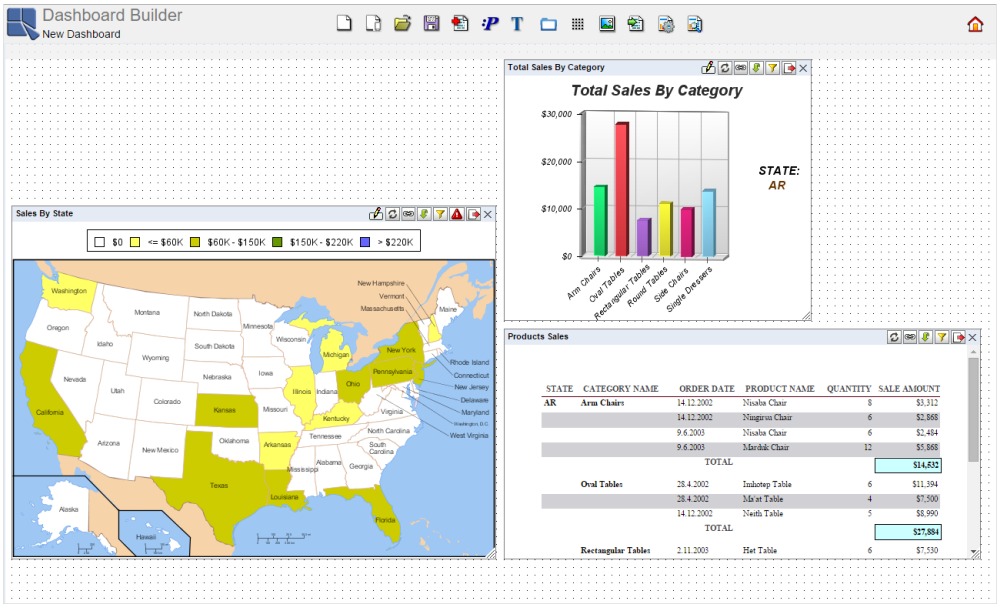
In this section we will guide you through creating a dashboard from report, chart, and map templates that come with ERES installation. At first we need to add some templates to the dashboard. To add reports, charts, or maps to the dashboard, click the *Add Report/Chart/Map* icon  on the dashboard toolbar. After you have clicked the button, the Insert Reports/Charts/Maps dialog will appear. The dialog contains a tree that mirrors the folder structure in the Organizer. Next, expand *your email* node and then the *QuickStart* and *QuickStart Dashboard* sub-nodes as shown on the image below.



Insert Reports/Charts/Maps Dialog

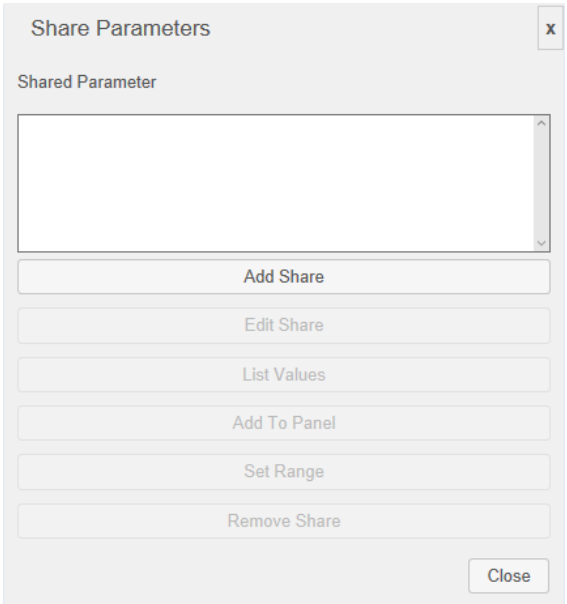
From the dialog click on all the templates (`Sales By State` map, `Products Sales` report, and `Total Sales By Category` chart) in the *QuickStart Dashboard* node. This will insert the templates into the dashboard. Once you have inserted the templates, close the dialog by clicking *Close* button.

Now move the *Sales By State* map template by clicking to the map header and drag it to the lower left position in the dashboard. Then select the *Total Sales By Category* chart and drag it to the upper right position in the dashboard. Finally, move *Products Sales* report and drag it below the chart. You can also resize the chart/report/map by clicking and dragging the sizing handle that appears in the lower right corner of the templates. After completing these steps, the dashboard should look like the following image:



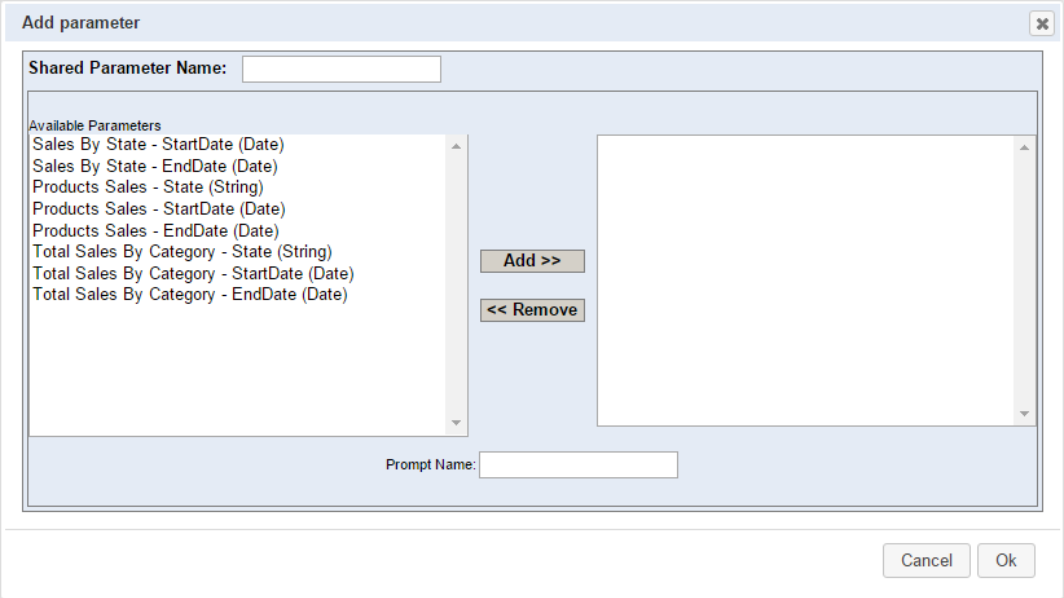
Dashboard with Added Templates

Next, we will add dashboard shared parameters that allow us to group common parameters from the chart, report, and map into a single parameter. First open the Shared Parameters dialog by clicking on the *Shared parameters* icon :P on the toolbar. This will open the Share Parameters dialog that allows you to add shared parameters to the dashboard.



Share Parameters Dialog

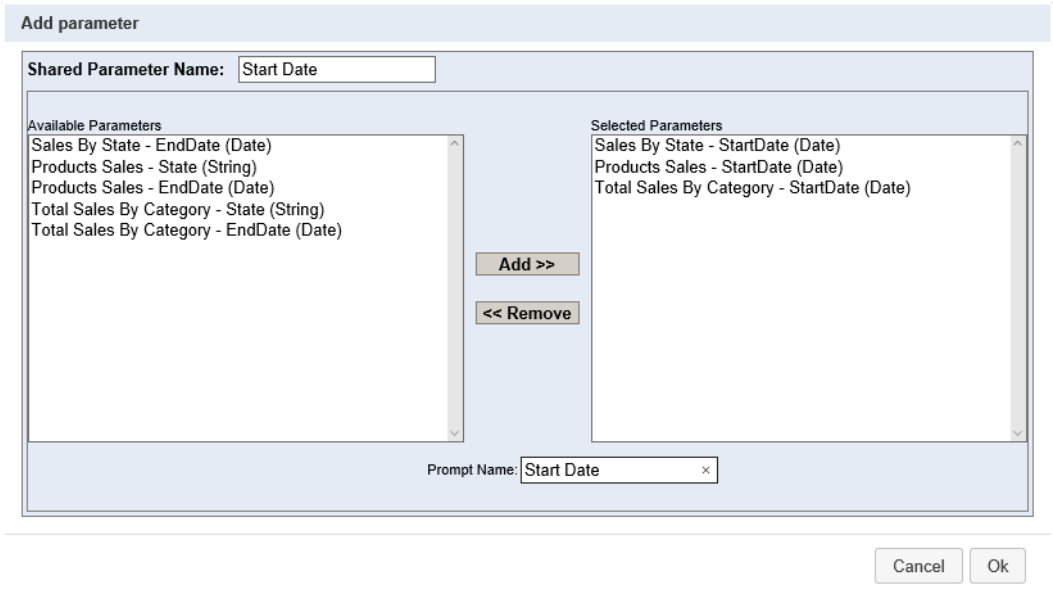
From the dialog, click *Add Share* button. This will open another dialog that allows you to specify a shared parameter.



Shared Parameter Dialog

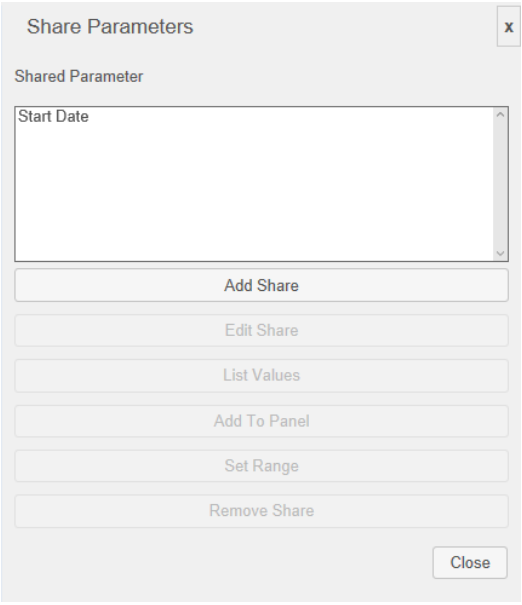
The left-hand side of the dialog shows all the available parameters from all the charts/reports/maps in the dashboard. You can select which parameters you would like to add to the selected parameters list by selecting them in the left-hand side and clicking the *Add* button. Note that all parameters in the selected parameters list must have the same data type.

In this example, we will create two shared parameters (`Start Date` and `End Date`) and put them into a date range panel. We will start with the first parameter (`Start Date`). First, enter the parameter name **`Start Date`** to the *Shared Parameter Name* text box at the dialog top. Enter the same text to the *Prompt Name* text field and then select all the `Start Date` parameters from the list on the left (you can select the parameters by **CTRL**+**Clicking** on the appropriate items in the list). After you have selected all the `Start Date` parameters from the list, click the *Add* button. This will add the selected parameters to the list on the right. At this moment the shared parameter dialog should look as follows:



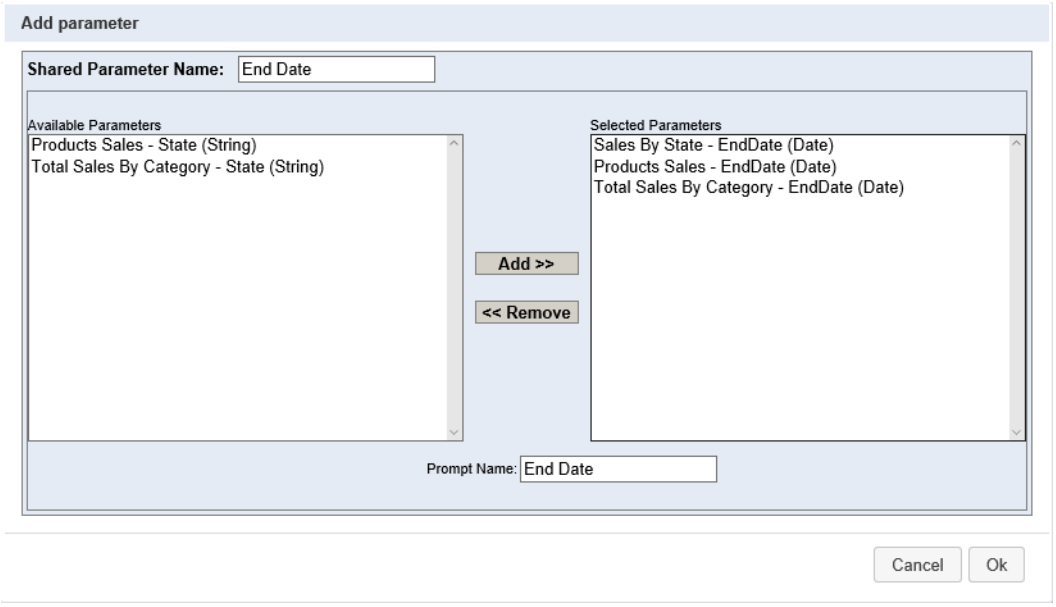
Shared Parameter Dialog - `Start Date` Parameter

Now click the *OK* button. This will put you back to the shared parameters dialog. As you can see the new shared parameter `Start Date` has been added to the shared parameters list.



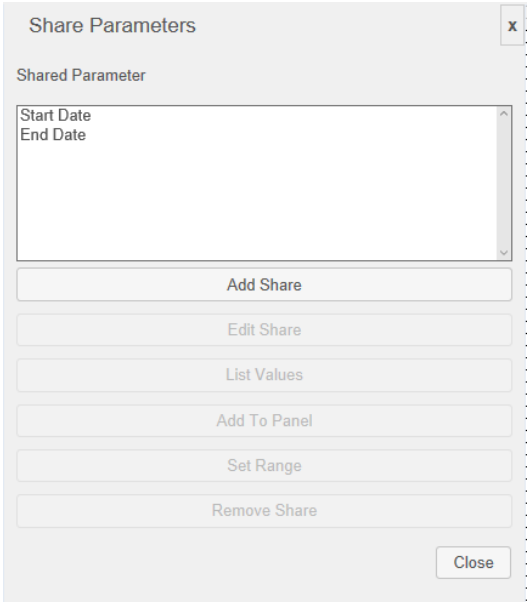
Share Parameters Dialog - `Start Date` parameter added

Next, we will add the second shared parameter `End Date` in the same way as mentioned above. From the dialog click the *Add Share* button again, which will open the shared parameter dialog. Enter the parameter name and the prompt name **`End Date`** and select all the `End Date` parameters from the list on the left. After you have selected all the parameters, click the *Add* button.



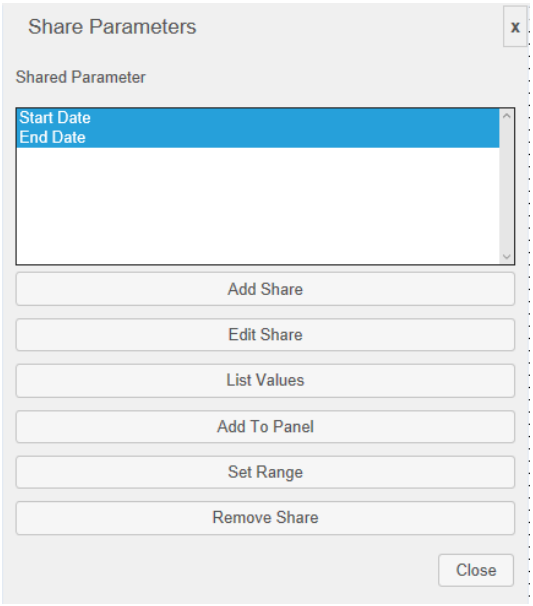
Shared Parameter Dialog - `End Date` Parameter

Click *OK* button. After adding both parameters you should see the same shared parameters dialog as follows:



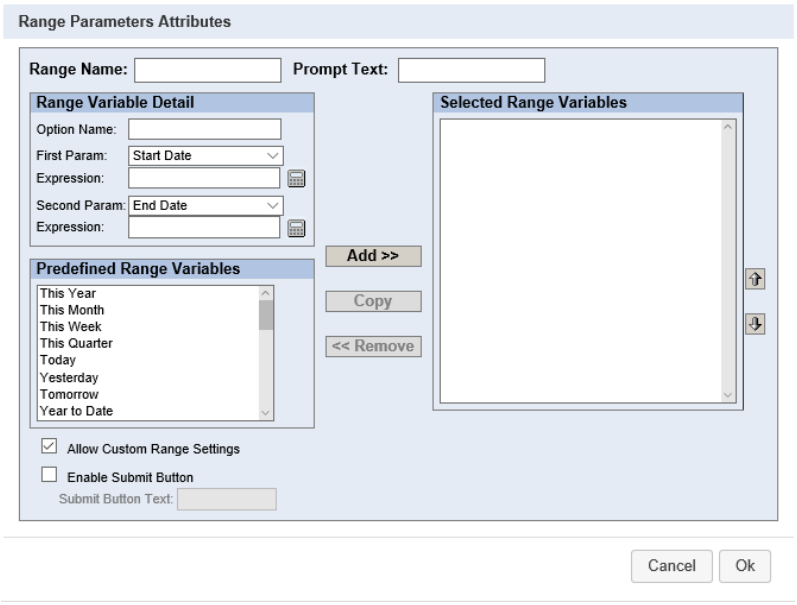
Shared Parameters Dialog - Added Parameters

In the dialog, select both parameters by **CTRL**+**Clicking** on the parameters as shown in the image below.



Shared Parameters Dialog - Selected Parameters

After selecting the parameters, click *Set Range* button. This will open the *Range Param Attributes* dialog that allows you to specify ranges for the parameters.



Set Date Range Dialog

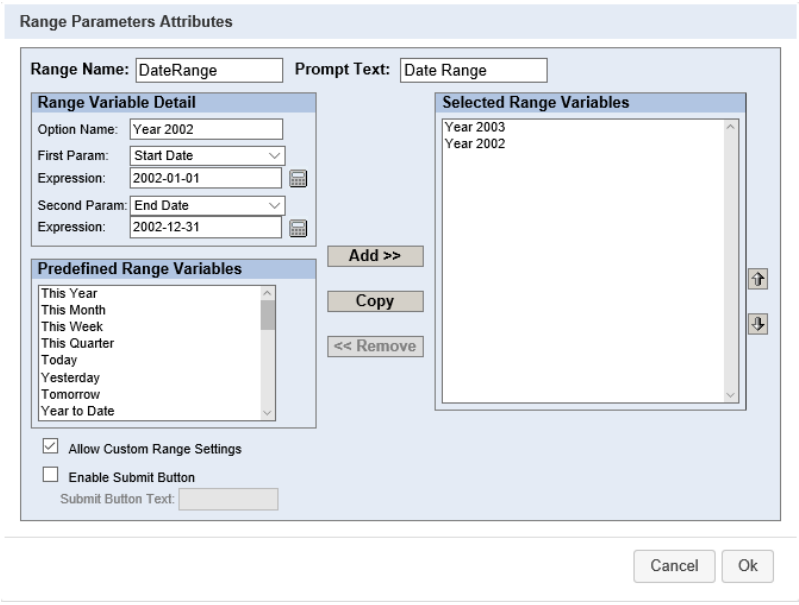
As you can see there are lot of options in the dialog. Please note that we will describe only the options necessary for the example. For more information about the rest of options, please see the Section 7.2.4.2 - Parameter Range.

First, specify the *Range Name* in the text box at the top of the dialog, e.g. **`DateRange`**. The name represents the title for the Date Range panel. Next, specify the *Prompt Text* to be **`Date Range`**. This text will be then displayed next to the selected date parameter drop-down menu in the Date Range panel. After entering the names, we will proceed with creating date ranges.

For simplicity we will add only two date ranges (`Year 2003` and `Year 2002`) in this example. The range `Year 2003` will include all the dates in year 2003 (dates from `2003-01-01` to `2003-12-31` in the Year-Month-Day format). Similarly, the range `Year 2002` will include all the dates in year 2002 (dates from `2002-01-01` to `2002-12-31`).

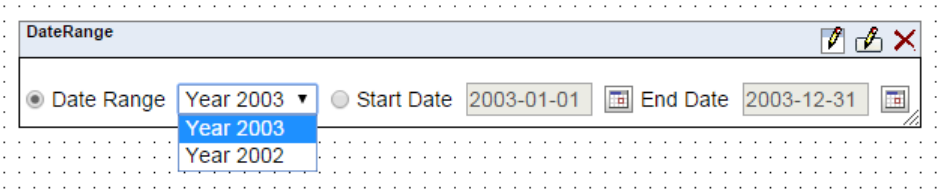
Now specify the name for the first range in the *Option Name* text box to be **`Year 2003`**. Next, enter the expressions for the `Start Date` and `End Date` parameters to be **`2003-01-01`** and **`2003-12-31`**. Once you have entered the *Range Name* as well as the expressions, click *Add* >> button. This will add the date range to the list of selected ranges on the right.

Similarly, add the second date range `Year 2002`. As mentioned above, first specify the range name to be **`Year 2002`** and enter the expressions for the `Start Date` and `End Date` parameters as **`2002-01-01`** and **`2002-12-31`**. After that, click the *Add* >> button to add the range to the list of selected ranges on the right.



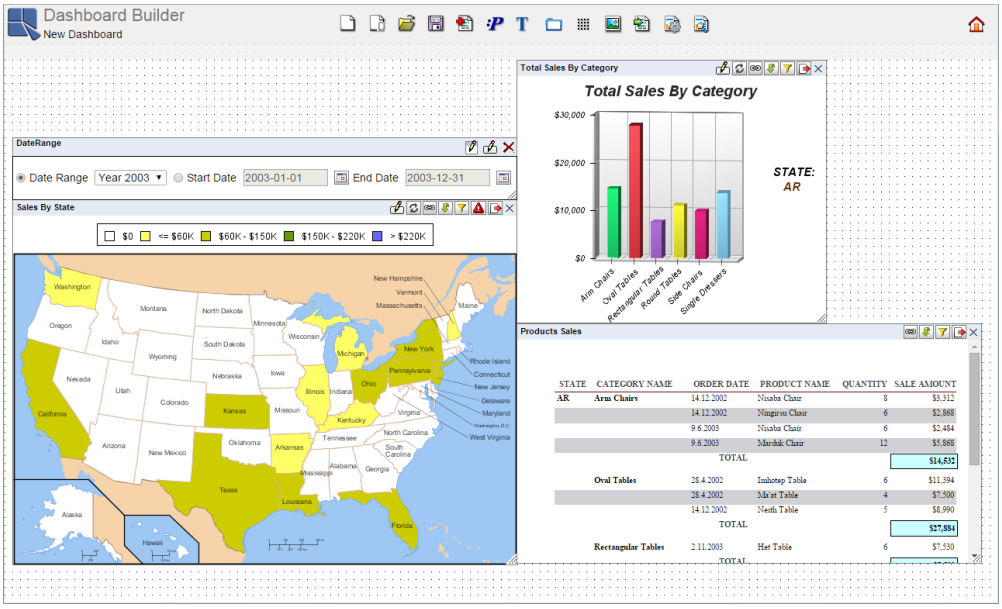
Set Date Range Dialog - Specified Ranges

Once you have added both date ranges into the selected date ranges list on the right, click *OK* button. After that click somewhere to the Dashboard Builder, and the date range panel appears.



Date Range Panel

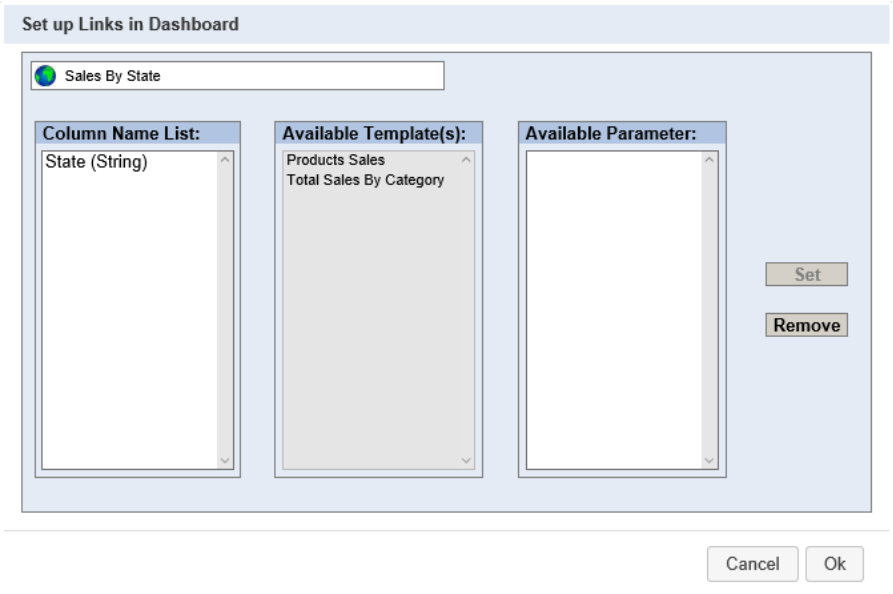
Now move the panel above the map as shown on the image below. To move the panel, just click on the panel header and drag it to the desired location. You can also resize the panel by clicking and dragging the lower right corner of the panel. The resizing rectangle will then appear, allowing you to adjust the size of the panel.



Dashboard With The Date Range Panel

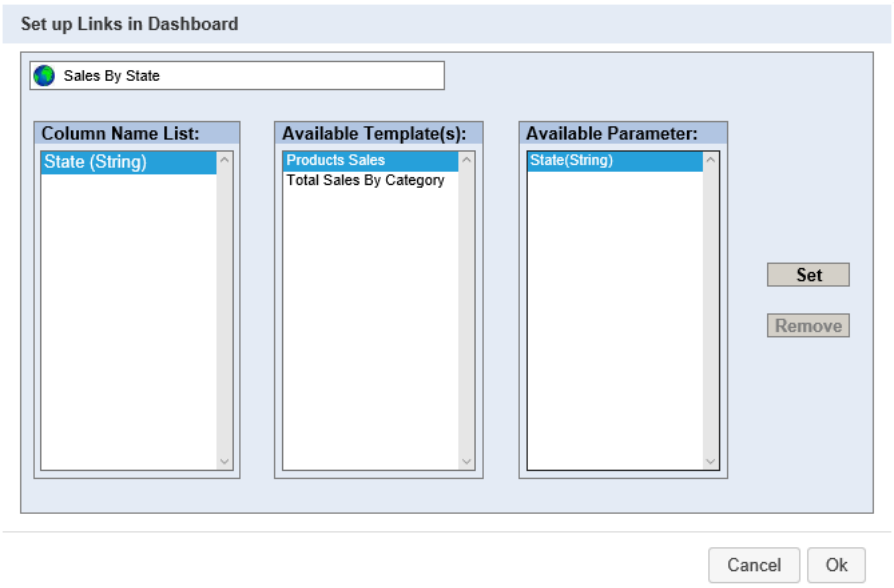
The last thing we will do in this example will be adding links from the map to the chart and report in the dashboard. This will enable sending the `State` parameter from the map data points to the chart and report.

To set up links from the map, click the *Add/Modify Link* icon from the map header. This will open the Set Up Link dialog below. There are three lists in the dialog. The first list contains all the available data columns and their data types in the source template (in our case ,`Sales By State` map). The second list shows all the available destination templates in the dashboard (in our case, `Product Sales` report and `Total Sales By Category` chart). Finally, the third column contains all the available parameters in the selected destinate template. Please note that the third column will show the parameters only after selecting the destination template in the second list.



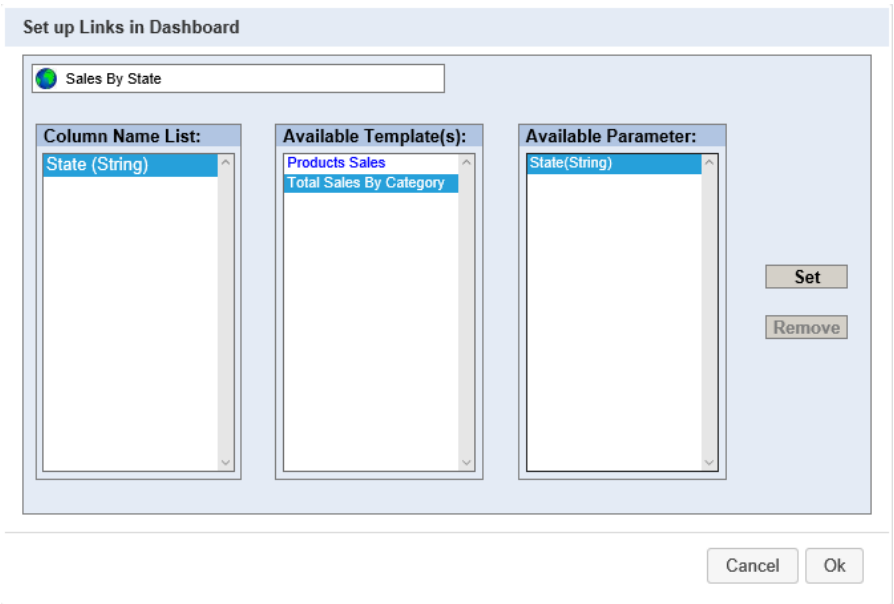
Set Up Links Dialog

To set up a link from the map to the `Product Sales` report, select *State (String)* item in the first list and the *Product Sales* item in the second list. This will cause showing the parameters from the report in the third list. After that select *State(String)* item from the third list. This should enable *Set* button as shown on the image below. Click the button to set up the first link.



Set Up Links Dialog - Link From The Map To The Report

Similarly, set up a link from the map to the `Total Sales By Category` chart. First, select *State (String)* item in the first list and the *Total Sales By Category* item in the second list. This will cause the parameters from the chart in the third list to be shown. From the third list select *State(String)* item. This should enable *Set* button as shown on the image below. Set up the second link by clicking the button.

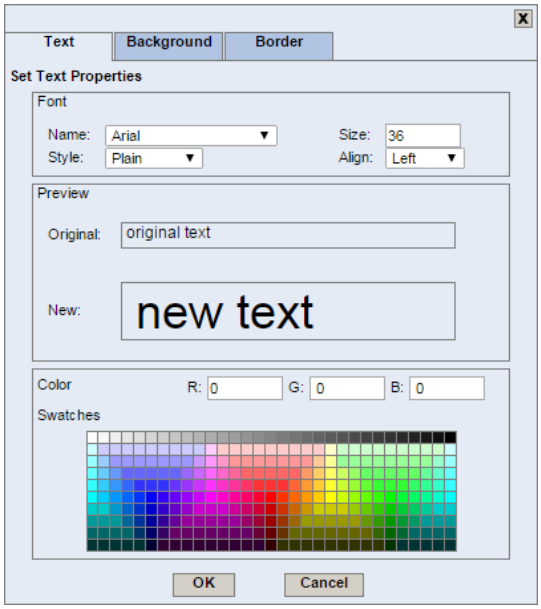


Set Up Links Dialog - Link From The Map To The Chart

Once you have set up both links, click *OK* button that will close the dialog. At this moment the links from the map should be set up. So if you now click on a data point in the map, the chart and report should refresh according to the data point parameter value.

To finish our dashboard, we will just add a dashboard title and choose a dashboard background color. To add the title, click the *Insert Label* icon  from the toolbar. A small rectangle will then follow your mouse cursor. Position the rectangle where you would like to insert the label and click. This should add the label panel into the dashboard. Next, double click to the label panel to edit the label text. Enter e.g. **SALES BY STATE DASHBOARD** text and click outside of the label panel. You may also want to resize the label because the label is now too small to be the

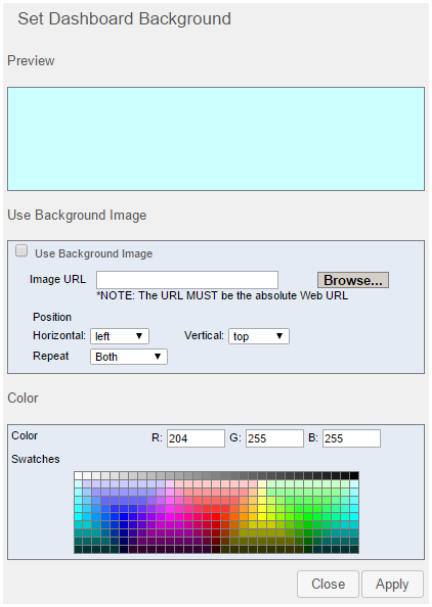
dashboard title. This can be done from the label properties dialog that opens after clicking on the *Edit* icon on the panel header.



Set Label Properties Dialog

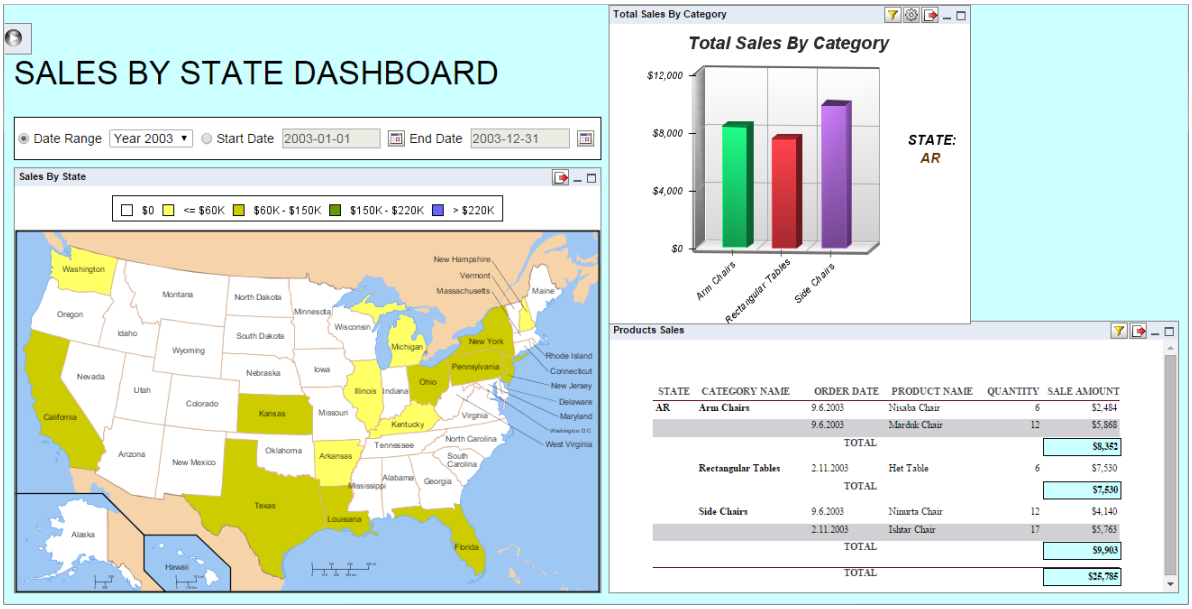
You can specify various properties for the label from this dialog. In the dashboard preview screen shot below, we just increased the label font size to **36**px in the *Text* tab, disabled the panel border in the *Border* tab (write **0** to the *Thickness* option), and set up the label background color to be transparent in the *Background* tab. Click *OK* to apply changes. Now, click and drag the bottom right corner of the label border to be the whole title in one line. For more information about the label properties, please see Section 7.2.5 - Insert Labels.

To change the dashboard background color, just click the *Add Dashboard Background* icon on the toolbar. This will open the *Set Dashboard Background* dialog from which you can choose the dashboard background color. Select the color and click *Apply*.



Set Dashboard Background Dialog

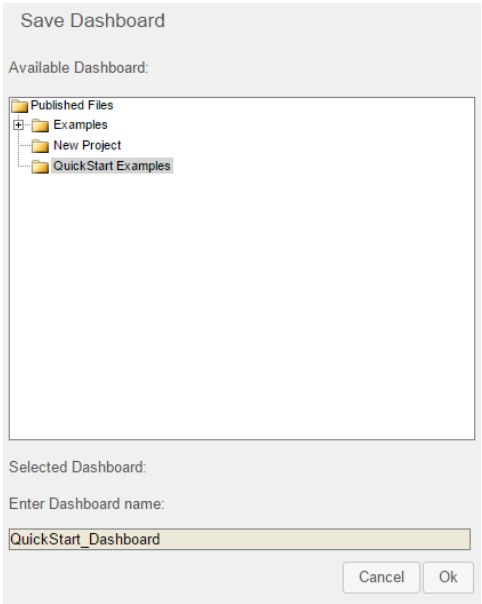
Once you have chosen the dashboard background color, you may preview the dashboard by clicking the *Preview* button  on the Dashboard Builder toolbar. The dashboard display will open in a new window. For more information about dashboard features, see Section 7.1 - Introduction to Dashboards.



Dashboard Preview

Q.10.1.2. Save the Dashboard

Close the dashboard preview window to return to the main Dashboard Builder interface. Click the *Save* button  on the toolbar. A dialog will open prompting you to specify a name for the dashboard.

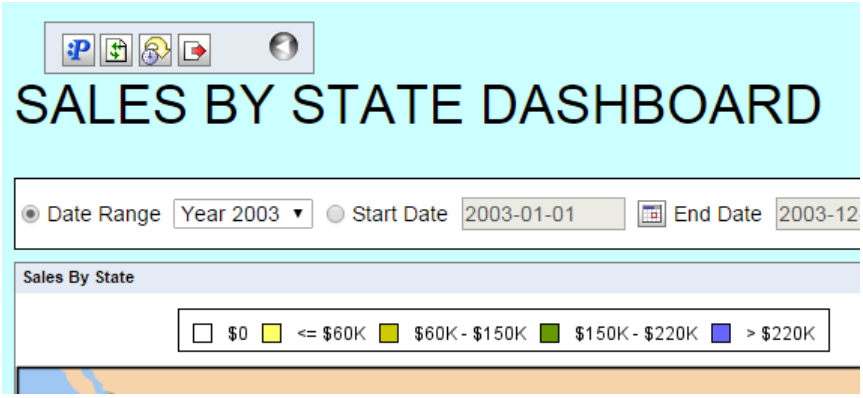


Save Dashboard Dialog

Enter a name for your dashboard. Click *Ok* to save the dashboard. The window will then give you a message that the dashboard was saved successfully. Click *Ok* to close the dialog.

Q.10.1.3. Export Dashboard to PDF

Click the *Options* button  on the Dashboard Builder toolbar and check the option *Show toolbar in preview* (in the *Other* section). Click *Apply* to apply the setting. Open the dashboard preview window by clicking the *Preview* button  . Then click the *Unpack* button  to open the preview toolbar. You can export dashboard to PDF file by clicking *Export Dashboard* button  from the preview toolbar.



Export Dashboard from Preview

All the reports, charts, SVG Map and other objects in the Dashboard will be exported to PDF file with the exception of Google Map due to licence restrictions.

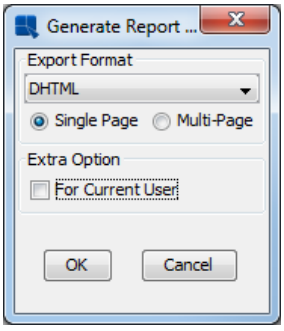
Alternatively, you may open the same .dsb file in Published Files (for more information please see Section 8.1 - The Menu Page), open the Dashboard toolbar and Export PDF there.

Note: There may be small differences in the appearance of elements in export (e.g. radio button shape can change from round to square).

Q.10.2. URLs

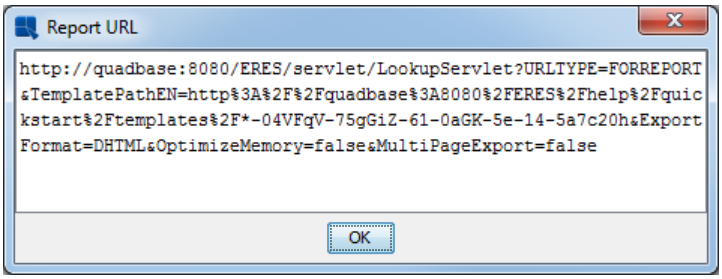
One automated deployment provided in ERES is the ability to run reports and charts via URL calls to the ERES server. Report and chart URLs are documented in detail in Section 8.2 - Image URLs, and Section 8.3 - Report URLs.

To generate a URL, go to the Organizer, select the `QuickStart532.rpt` file (drill-down example that can be created in Section Q.4.3.4 - Drill-Down), and click the *Generate URL* button on the toolbar  . A dialog will open prompting you to specify options for the generated URL.



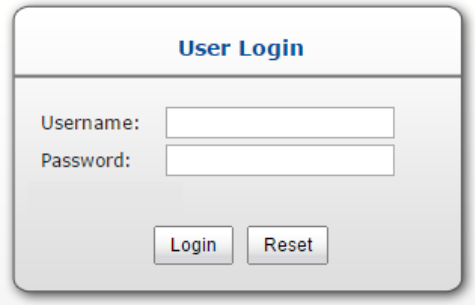
URL Options Dialog

In this dialog, select single page **`DHTML`** as the export format and disable the option marked *For Current User*. This will create the URL without encoding a username and password, prompting users to login when running the URL. Click *Ok* and the URL will be generated in the *Report URL* dialog.



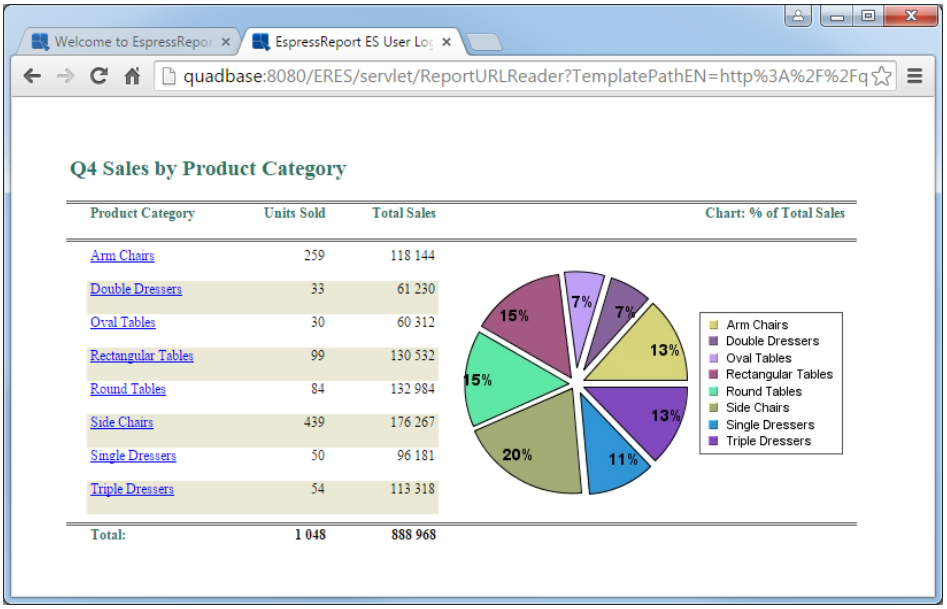
Generated URL

Now, select the generated URL and copy the text to the browser's location dialog. Hit enter to run the URL. You will first be directed to a login dialog because you selected not to pass in the username and password in the URL.



URL Login Dialog

Enter your *User Name* and *Password* in the dialog and click the *Login* button. The report will then load in the browser.



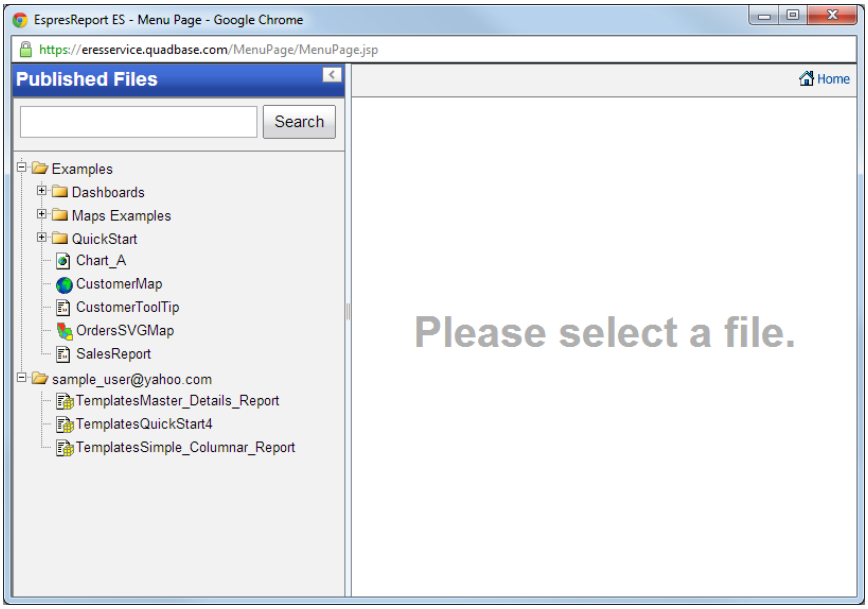
Generated Report

As you can see the drill-down column is generated as hyperlinks in the exported report. Click one of the links and you will go to the lower-level report in the same manner as you did in the preview window.

Q.10.3. The Menu Page

Another way that ERES automatically publishes reports, charts, maps, and dashboards is through the Menu Page. The Menu Page is a thin-client interface that allows users to run/view the reports, charts, maps, and dashboards to which they have access in the Organizer.

To launch the Menu Page, go to the ERES Start page, click the link labeled *Published Files* and the Menu Page will load.



Menu Page

The page contains a list of all the reports, charts, maps, and dashboards in your project and some examples. The only option available is to run the file because you do not have any active schedule or archive jobs. To open a file, expand respective project/folder nodes (in the left-hand tree-list) to locate the file and click on the file name. The file will load in the right-hand DHTML Viewer panel.

For more information about the options and functions available in the Menu Page, see Section 8.1 - The Menu Page. After you have finished viewing your report/chart/map/dashboard, close the window that contains it, and click the *Home* button in the menu interface to return to the start page.

Q.11. Alerts

In previous sections, you learned how to create reports, charts, and maps (let's call them “objects”), and how to publish them in dashboards. This section describes how to add and manage alerts in objects, how to “watch” alerts in dashboards and how to monitor alerts automatically.

An alert is a range of some indicator (for example, profits/losses, sales, up-time, etc) you want or do not want the indicator to reach. If the indicator's value reaches into the alert range (in other words: if an alert was triggered), you will be instantly informed about that, so you can fix issues in a timely manner. More detailed information about alerts can be found in Section 9.1 - What is an Alert.

First of all, you have to create the alert. That means that you have to select which indicator you want to watch and what values are critical for you. Each object (report, chart, or map) can have its own alerts.

Q.11.1. Alerts in Charts

In Section Q.5.3.2 - Control Areas, you learned how to add a “control area” into a chart. That is all you have to know to add alerts into charts – just open a chart in Chart Designer and add one or more control areas into it. Basically, you can think of a control area as an alert in a chart.

Launch ERES Organizer (as described in Section 3.1.1 - Starting the Organizer), select `QuickStart` folder under your project and double-click on the *QuickStart54.tpl* from the right-hand side to open it in the Chart Designer. The

chart shows defects counts for a manufacturing process for each 8 hours. We worked with this file in Section Q.5.3 - Advanced Charting Features and we set there the control area for more than 20 defects. We will add another alert that will indicate the defects that are not critical, although they are not optimal.

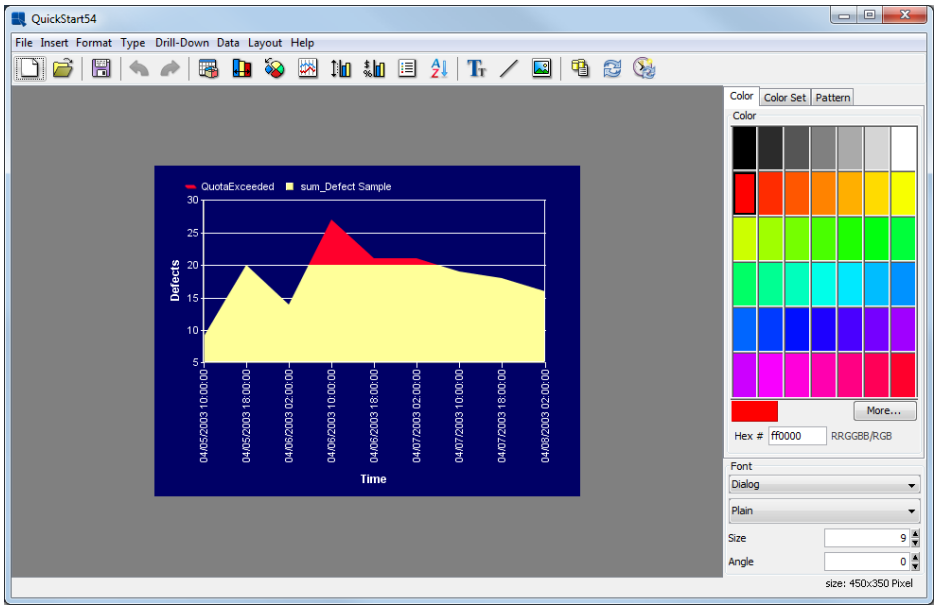


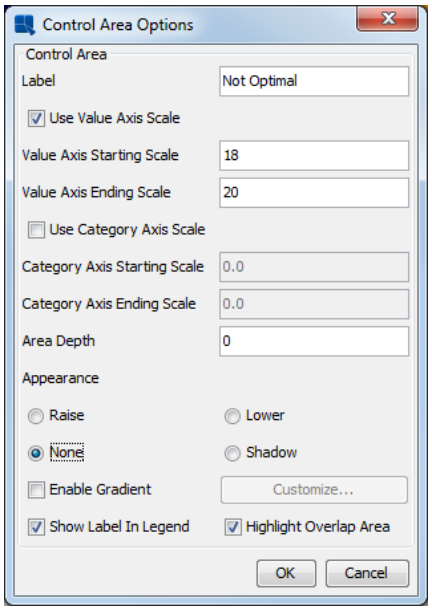
Chart With Contol Area

Open the *Insert* menu and choose the *Control Area* option. *Control Area Options* dialog has appeared.

To add a new control area, click the *Insert* button. Now, it is time to set up our new alert. The first option is called *Label*. The *Label* is set to **`New Range`** by default. Delete the default label and set it to **`Not Optimal`** instead.

Now we will set up the alert range. Set the *Value Axis Starting Scale* to **`18`** and the *Value Axis Ending Scale* to **`20`**.

To include the control area in the legend, choose the *Show Label In Legend* option. Choose also *Highlight Overlap Area*.



Control Area Options



Note

Complete description of this dialog can be found in Section 5.2.4.6.3 - Adding Control Areas

Click *OK* to close the dialog. As you can see, our new alert was created successfully. Now click the *OK* button again to go back to the Chart Designer. The new control area has appeared in the chart. It is highlighted by a random color. To change the color, click on the new area and then choose an indicative color (for example, light red) from the *Color* panel on the right side of the window. Save the chart and close the Chart Designer.

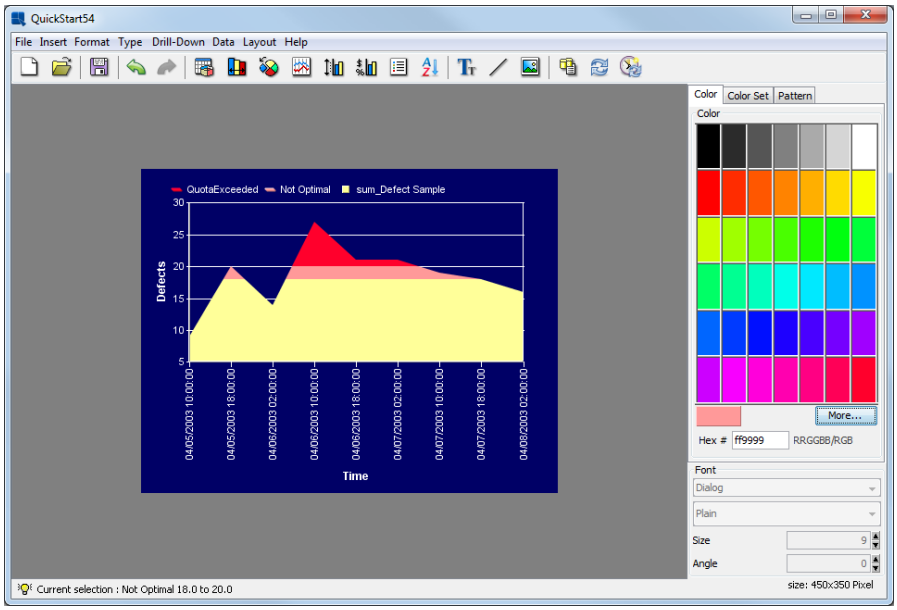


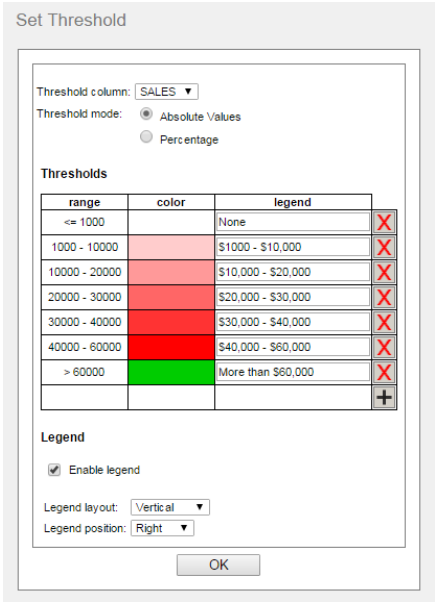
Chart With Two Control Areas

Q.11.2. Alerts in Maps

In the Section Q.9 - SVG Maps, you created a SVG map with thresholds. This means that you have also created a map with alerts because thresholds are equivalent to alerts in maps.

Let's add a threshold into an existing SVG map. Open ERES main page and click on the *SVG Maps* link. Click on the  *Open Map* icon on the toolbar. Expand `Examples/Dashboards/Alerts Components` nodes in the *Open Map* tree-list, select the `Sales By State` map and click the *OK* button. Click *Submit* to use default parameter values. To manage thresholds, click on the  *Set Thresholds* icon. To add a new threshold (alert), click on the  *Add Threshold* icon. In the *Enter threshold value* dialog, type **`60000`** and click *OK*. As new thresholds are added, some legend labels are not right. The legend labels are important for alerts because they also double as alert names. Set the *40000 – 60000* legend to **`$40,000 - $60,000`**. Now set the *> 60000* legend label to **`More than $60,000`**. Click on the *> 60000* color field. *Select Color* dialog should show up. Choose an appropriate color (for example, green) and click *OK*.

Now, the Set Thresholds dialog should look like this:



Set Thresholds Dialog

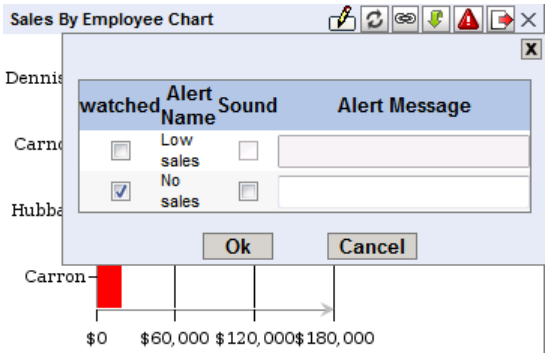
If it does, click *OK* to confirm the changes. The map has been changed since you added the threshold. Click on the *Save* icon. The copy of the read-only map is created. Write the map name, choose the folder under your project and click *OK*. Click icon in the top-right corner of the window to go back to the ERES main page.

Q.11.3. Alerts in Dashboards

In previous paragraphs, you added alerts into few objects. Now we want to make use of the alerts. There are two ways to utilize alerts – alert watching and alert monitoring. This paragraph describes how to set up alert watching. Alert watching is set up in dashboards. The basic principle is quite easy: insert some objects (reports, charts or maps) into a dashboard and pick the alerts you want to watch.

Log on to the ERES main page and launch *Dashboard Builder*. Click *Open* icon. Expand `Examples/Dashboards` nodes, select `Alerts Dashboard` dashboard, and click *Ok*. All dashboard objects have *Set Alert* icons in their header bars. That means that all dashboard objects have some alerts that can be watched.

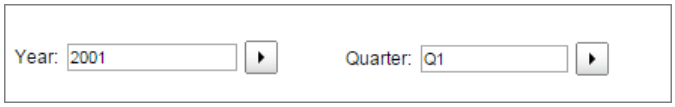
Locate the *Sales By Employee Chart* chart and click on its *Set Alert* icon. The following dialog should appear.



Sales By Employee - Watched Alerts

There is the list of all alerts. *No sales* alert is set to be watched. Leave this alert checked, click *Ok* to close the dialog and then click on the *Preview* icon in the main menu. Notice that the chart's border is flashing. That

means that the *No sales* alert was triggered. In other words: at least one *Sales By Employee Chart* chart bar has a zero value. Now locate following panel:



Parameter Panel

You can change the *Year* and *Quarter* parameters to see the situation in the required period.

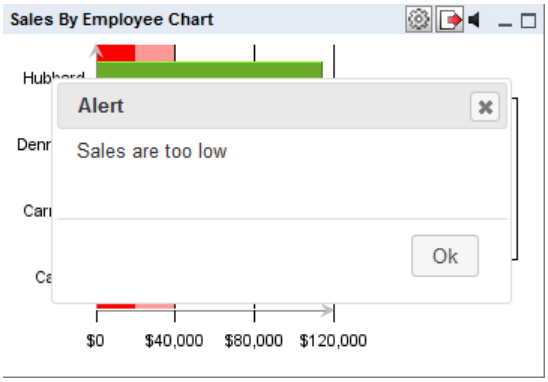
You may also notice that the `Sales Overview Report` header is flashing. That is because the report contains a watched alert which is triggered. You cannot see the report until you click on its tab header, but you can tell that an alert is triggered in the report.

Please note that alert highlighting, blinking and alert messages are visible only in *Preview* of the Dashboard Builder.

Q.11.3.1. Sound alerts and alert messages

Border flashing is not the only way to alert you. You can also set alert messages and sound alerts. Close the *Preview* window and click again on the  *Set Alert* icon for the *Sales By Employee Chart*. For our watched *Not Optimal* alert, check *Sound* and enter any message into *Alert Message* textbox. Click on *OK* and use *Preview* button to preview the dashboard.

Trigger the alert by changing the *Quarter* to *Q2* in the *Parameter Panel* and an alert message with your entered text will appear and you will also hear a beep sound. To mute the sound, click on the  *Mute* icon in upper right corner of the chart. Please note that each template has its own *Mute* button. To close the alert message, simply click on *OK* button.



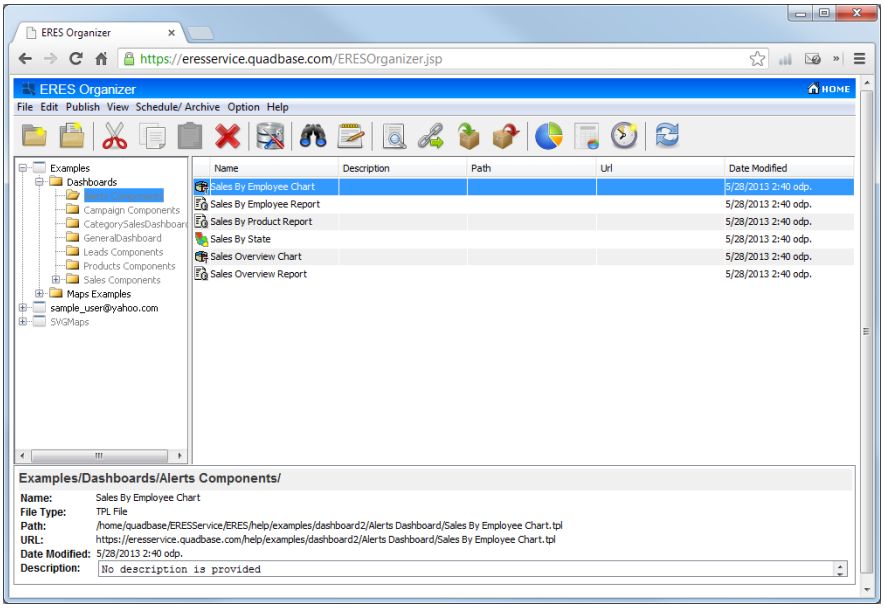
Triggered alert message

For more information about alerts, please refer to Section 9.3 - Dashboard alerts.

Q.11.4. Alert Monitoring

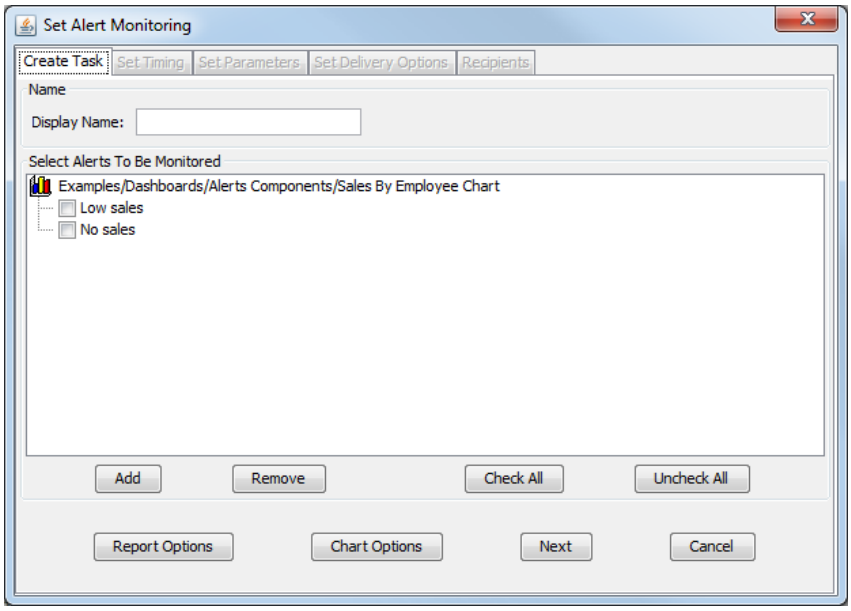
Another way of using alerts is alert monitoring. Unlike alert watching, alert monitoring checks alerts even when there is no dashboard running.

Lets set up a new alert monitoring job. Log on to the ERES main page and click the *Organizer* button. Expand `Examples/Dashboards/Alerts Components` nodes and select the `Sales By Employee Chart` file.



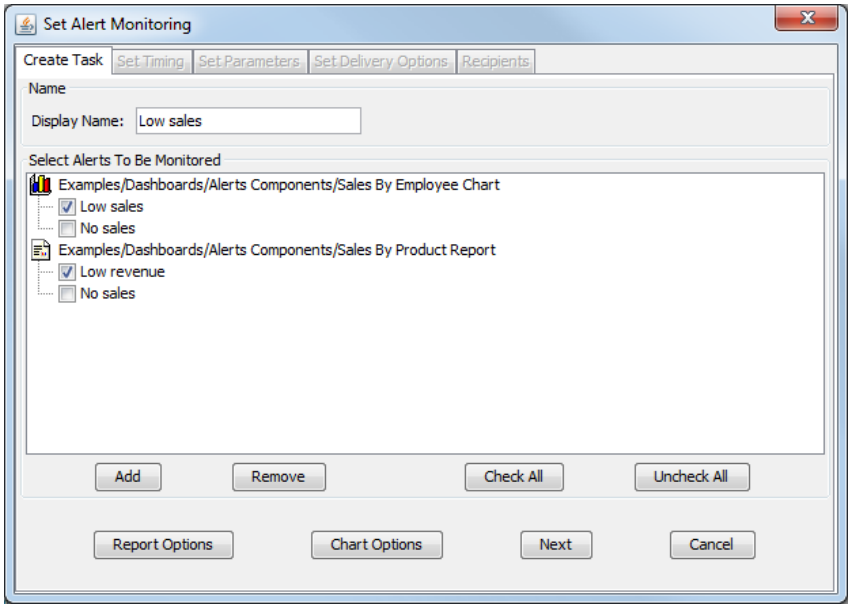
Sales By Employee Chart Location

Select *Set Alert Monitoring* from the *Schedule/Archive* menu. The following dialog should show up:



Setup Alert Monitoring Dialog

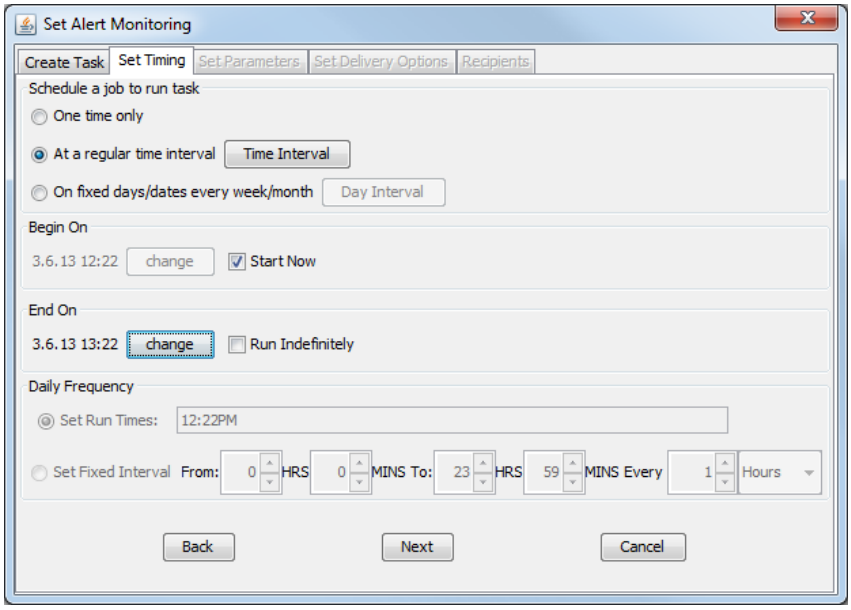
Click the *Add* button. Expand `Examples/Dashboards/Alerts Components` nodes, select the `Sales By Product Report` report and click *OK*.On the Set Alert Monitoring dialog select `Low sales` and `Low revenue` alerts. Type **`Low sales`** into the *Display Name* field. The dialog should look like this now:



Setup Alert Monitoring Dialog

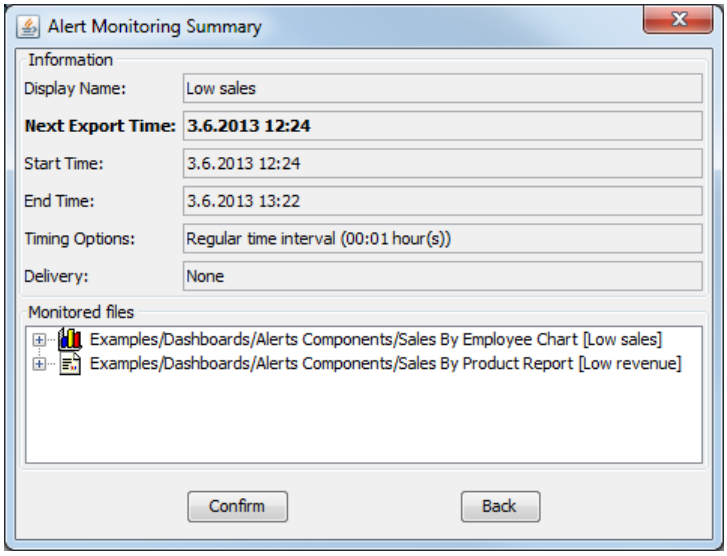
If it does, click *Next*.

You should be on the *Set Timing* tab now. By default, the alerts would be checked just once, few minutes after configuring it. Select the *At a regular time interval* option. Set Time Interval dialog should show up. Default time interval is **1** minute, which is fine for us, so just click *OK* to confirm it. Uncheck the *Run Indefinitely* checkbox and click the *change* button right next to it. Set the end time one hour to the future (so you do not have to delete the job manually after you are done with this QuickStart chapter) and click *Ok*. The dialog should look like this:



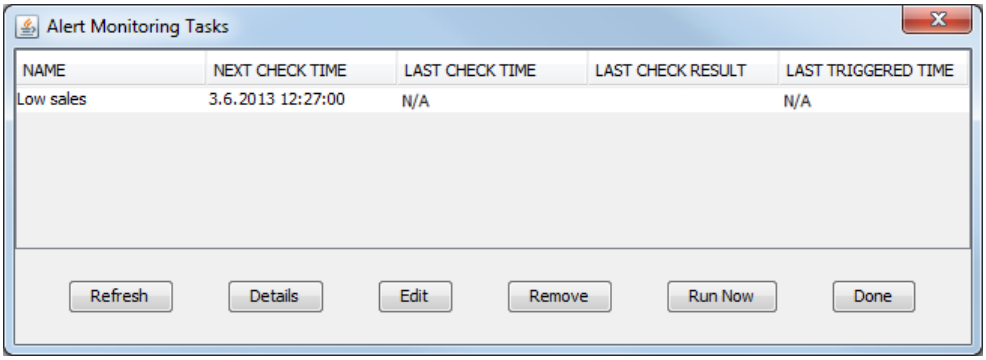
Set timing

Click *Next*. Click *Add* to add new parameter set to the *Sales By Employee Chart* chart. Click *OK* to use default parameters. Click on the *Sales By Product Report* tab header, click *Add* and *OK* again and then click *Next*. The next step is configuring delivery options. This requires further knowledge and we will skip this step right now. If you are interested in setting delivery options, see Section 9.4.3.1 - Create/Edit monitoring dialog. Click *Next*. The following dialog should appear allowing you to check your settings before saving the alert monitoring job. Click *Confirm* to save the alert monitoring job.



Alert Monitoring Summary

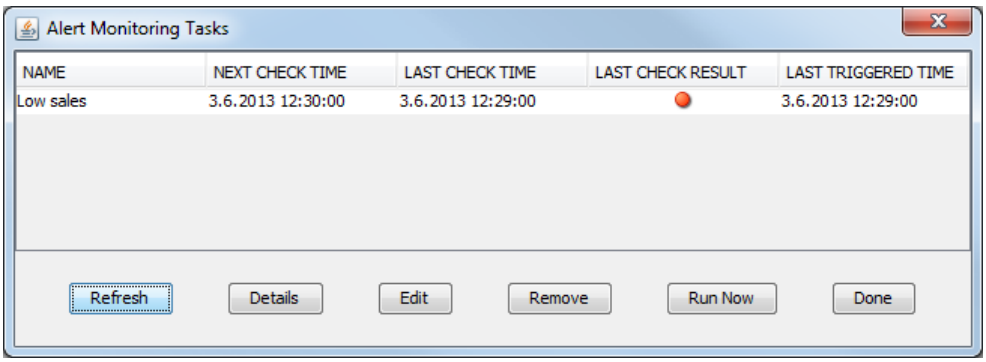
You should be back in Organizer now. Click on the *Schedule/Archive* pull-down option in the main Organizer menu and select *View Alert Monitoring Tasks* option. The following dialog should open:



Alert Monitoring Tasks Dialog

The `Low sales` task should be on the list. Check the *Last check time* value. If no check was performed until now, it will show `N/A`. You can either wait for next automatic check (look at the *Next check time* if you want to know when that's going to happen) and then click *Refresh*. You can also select the `Low sales` task and click *Run Now* and then click *Yes*.

After the task was checked, click *Refresh*. You should be able to see the results.

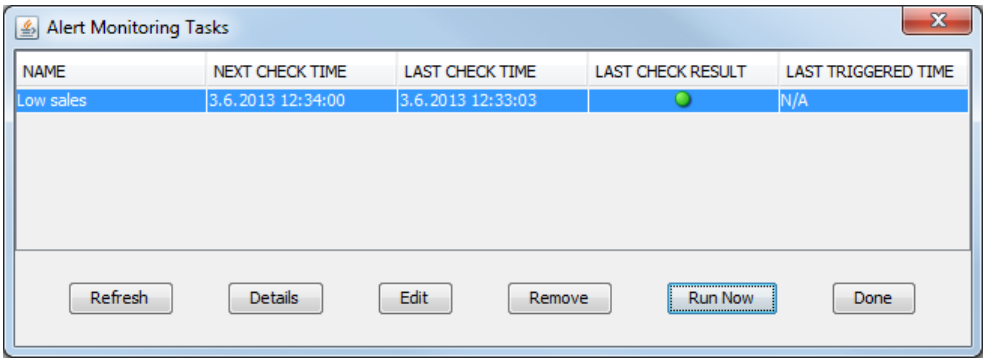


Alert Monitoring Tasks Dialog

The *Last check result* indicator is red. That means that at least one alert was triggered during last check. Also notice that the *Last triggered time* is now the same as *Last check time*.

Select the *Low sales* task and click *Edit*. Select the `Sales By Product Report` and then click the *Remove* button. Click *Done* and then click *Confirm* to save the changes.

Wait for the automatic check or click *Run Now* again. Click *Refresh*.



Alert Monitoring Tasks Dialog

Now, the *Last check result* indicator is green. That means that no alert was triggered during last check.



Note

Please, keep in mind that this is just a quick guide which should guide you through basic steps of configuring alert monitoring. In reality, data in the database would be changing constantly; therefore, you would not have to change alert monitoring job properties in order to trigger/not trigger an alert.

Chapter 1. Features and Release Notes

1.1. New Features List V7.0

1.1.1. New ERES Cloud Features

1. Unified data registry dialog in Quick Designers and Map Designers allows users to view query, query result, and test with different parameter values, so users can locate the desired query easily (Section 4.2 - Data in QuickDesigners and Maps).

2. Diagrams on start page of Quick Designers and Map Designers helps end-users to understand how to use the tools quickly (Section 6.2.2 - Start Online Maps).

1.1.1.1. Dashboard Builder (Chapter 7 - Designing Dashboards)

1. Being able to change chart type in dashboard preview gives end-users more control on dashboard contents (Section 8.9.2 - Preview Options).

2. Directly opening report/chart/map designer from dashboard builder makes it much more convenient to edit existing components or add new components (Section 7.2.3 - Add Charts, Reports and Maps).

3. Besides original static dashboard, responsive dashboard is added as a new dashboard type. Responsive Dashboard optimizes utilization of the display without exceeding the display width preset by user. Therefore, it's useful for mobiles (Section 7.2.2 - Responsive Dashboard).

4. Refresh Template icon on each report/chart/map makes it easier to test dashboard with real data (Section 7.2.3.3 - Chart/Report/Map Toolbar).

1.1.1.2. QuickDesigner Reports (Section 5.3 - QuickDesigner Reports)

1. Simplified UI includes only two steps: Select Datasource and Construct Report.

2. Constructing and formatting report in only one interface makes it intuitive and easy to see what you get and flexible to make changes (Section 5.3.4 - Format Report Elements).

3. Drag and drop to add report columns and order report column ease the work of report design (Section 5.3.4.1 - Add a Column, Section 5.3.4.5 - Column Order).

4. Right click on report column to do Aggregation/Format/Remove is straightforward (Section 5.3.6 - Aggregation/Group, Section 5.3.7 - Data Formating, Section 5.3.4.2 - Remove a Column).

5. 'Excel-like' search icon on each column header allows data filtering by column (Section 5.3.9 - Data Filtering).

6. Parameters pane on the right makes report testing easy (Section 5.3.5 - Parameter Setting).

7. Less items on Toolbar reduces confusion and saves time (Section 5.3.3 - Toolbar).

8. Collapse Sidebars allow more space for report design (Section 5.3.4.3 - Collapse Sidebars).

1.1.1.3. QuickDesigner Charts (Section 5.4 - QuickDesigner Charts)

1. Simplified UI includes only two steps: Select Datasource and Construct Chart.

2. Setting pane and chart preview pane are on the same window, so all design steps can be done on one window (Section 5.4.3 - Data Mapping and Ordering).

3. Improved toolbar shows Chart-Specific Options based on current chart type and combines settings to reduce icons (Section 5.4.5.5 - Chart-Specific Options, Section 5.4.4 - Toolbar).

4. Parameters setting on Chart Preview pane makes chart testing easy.

5. Collapse Sidebar allows more space for chart design.

1.1.1.4. Online Maps (Section 6.2 - Online Maps)

1. Online Map is a newly designed Google Map. It is more straightforward to create a new map (Section 6.2.2 - Start Online Maps).

2. Displaying setting on left pane, instead of popup dialog, makes UI more user-friendly.

3. Save dialog shows Organizer node tree. It makes it easier to locate the desired organizer folder to save map file (Section 6.2.6 - Save Map).

4. Heatmap feature provides visualization of data intensity (Section 6.2.5.6 - Heatmap).

5. Being able to set parameter values in Map Designer makes testing easier.

6. More map types satisfies different needs (Section 6.2.5.5 - Map Options).

7. Users can draw various shapes instead of simple map points. The shapes are drawn in Coordinates Editor (Section 6.2.3.2 - Create Coordinates, Section 6.2.3.2.3 - Manually).

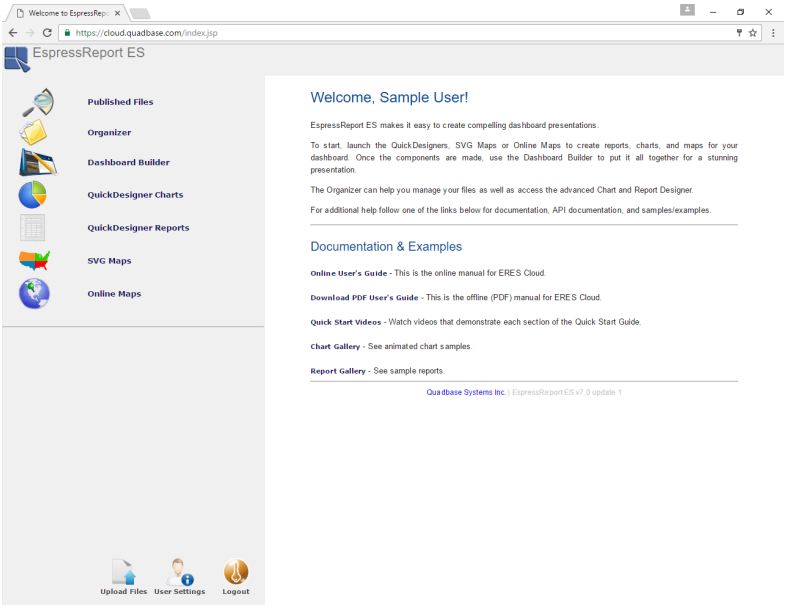
1.1.1.5. SVG Maps (Section 6.3 - SVG Maps)

1. Redesigned UI is more straightforward to create a new map.

2. Allow users to change SVG image after a map is created (Section 6.3.3.3.1 - Change SVG Image).

Chapter 2. ERES Cloud Main Page

After you log into the ERES Cloud, the main page will open. From this page, you can access all ERES Cloud features:



ERES Cloud Main Page

Published Files — Shows list of all charts, reports and dashboards you have created and also some example for you to try. For more details, see the Chapter 8 - Publishing (Menu & URLs) chapter.

Organizer — Organizer allows you to manage all files (charts, reports, dashboards, data sources etc.), organize the files into folders, manage your data sources and also launch the full version of the Chart Designer. Organizer documentation can be found in Chapter 3 - Organizer.

QuickDesigners — Create or edit charts and reports quickly and easily. To learn how to use the Quick-Designers, see Section 5.3 - QuickDesigner Reports and Section 5.4 - QuickDesigner Charts .

Map Designers — Allows you to visualize your data on maps. More on SVG Maps and Online Maps can be found in the Chapter 6 - Designing Maps chapter.

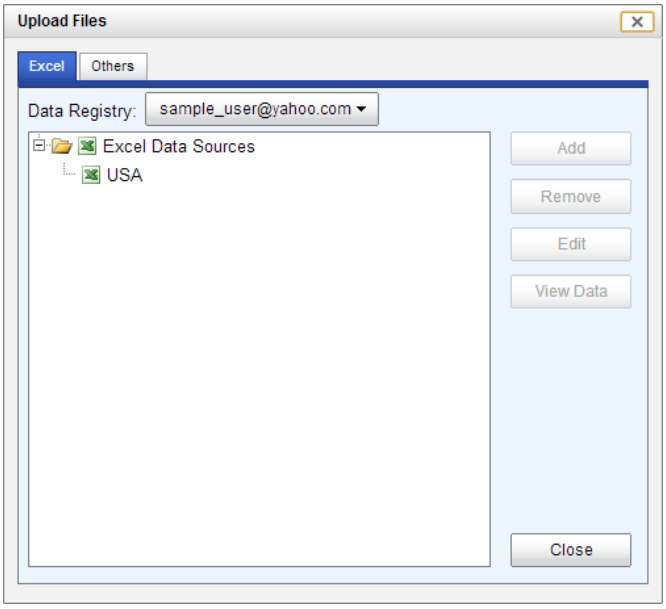
Dashboard Builder — After you have created some templates (i.e. reports, charts, and maps), you can publish them in well-arranged dashboards. See Chapter 7 - Designing Dashboards to learn more.

Upload Files — Upload Excel data spreadsheets, other data sources, images, etc. and use them for your charts, reports, and maps. To learn more, see Section 2.1 - Upload Files.

User Settings — Allows you to change your email address and password.

2.1. Upload Files

To manage data sources, click on the *Upload Files* button on the ERES Cloud main page. The following dialog will show up.



Upload Files

You can see two tabs in this dialog. *Excel* tab allows you to upload and manage your Excel data sources. *Others* tab allows you to upload and manage other files (e.g. images, text data sources).

2.1.1. Excel Data Sources

You can manage your own Excel data sources from *Excel* tab of the *Upload Files* dialog.

There are four buttons on the *Excel* tab:

ADD Allows you to upload a new excel file and create a data source from it.

REMOVE To delete a data source, select a data source node on the left and click on the *REMOVE* button.

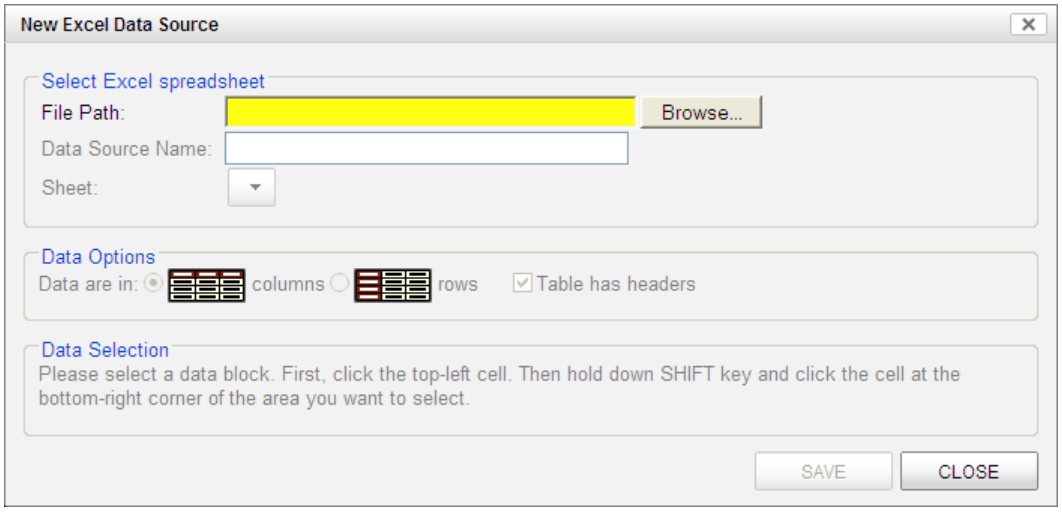
> **Warning**
>
> This action can't be undone!

EDIT Select an existing data source on the left and click on the *EDIT* button in order to modify the data source. This will open the same dialog that is described in the Section 2.1.1.1 - Add/Modify Excel Data Source chapter.

VIEW DATA To preview a data source, select the data source on the left and then click on this button. First 50 rows of the data will be displayed.

2.1.1.1. Add/Modify Excel Data Source

If you click on the *ADD* button on the *Excel Data Sources* dialog (as described in the previous paragraph), the following dialog will show up.



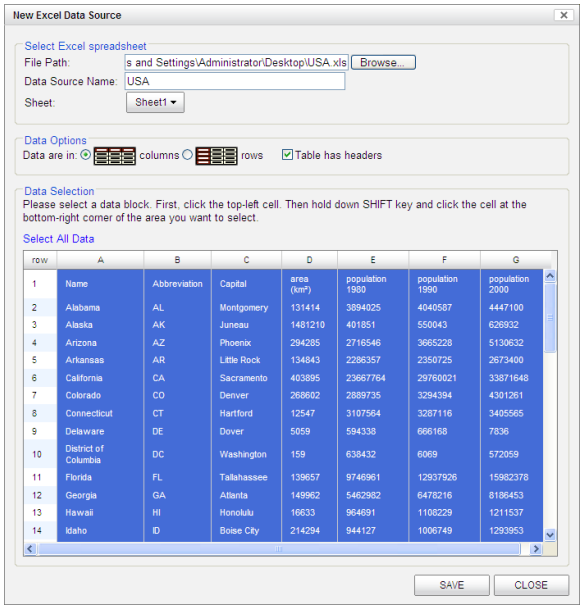
Add New Excel Data Source Dialog

If you want to create a new data source, first you have to upload an Excel file. Click on the *Browse...* button, select an excel file and click *Open*. Data from the file will be displayed in the *Data Selection* section.



Note

Only raw data are shown in the *Data Selection* table. No formatting (e.g. different fonts, colors, etc.) will be displayed.



Add/Modify excel Data Source Dialog

Data Source Name — This name will be used in ERES Cloud for this data source. By default, the excel file name will be used, but it can be easily changed to any other name.

Sheet — An excel file usually contains several sheets. Use this drop-down menu to switch between sheets. If you switch to another sheet, its data will be shown in the *Data Selection* table.

Data are in columns/rows — Allows you to set orientation of the data table. Data in columns mean, that each table column contains values of the same variable (for example: num-

ber of items in stock, product price etc.). Also, if you select the *Data are in columns* option, values in each column should have the same data type (except for header which is always converted to String). The same logic is applied to data in rows.

For example: take a look at the previous picture. Data in the picture are in columns, because each column contains data of a certain kind (for example: state name).

The *Table has headers* check-box allows you to use first field of each selected column/row as its header. If you want to use headers, you have to select the header row/column in the *Data Selection* table along with the data (as seen in the previous picture).



Note

If you choose the *Data are in rows* option and save the data source, the data orientation will be automatically transposed to vertical orientation (i.e. data will be in columns).

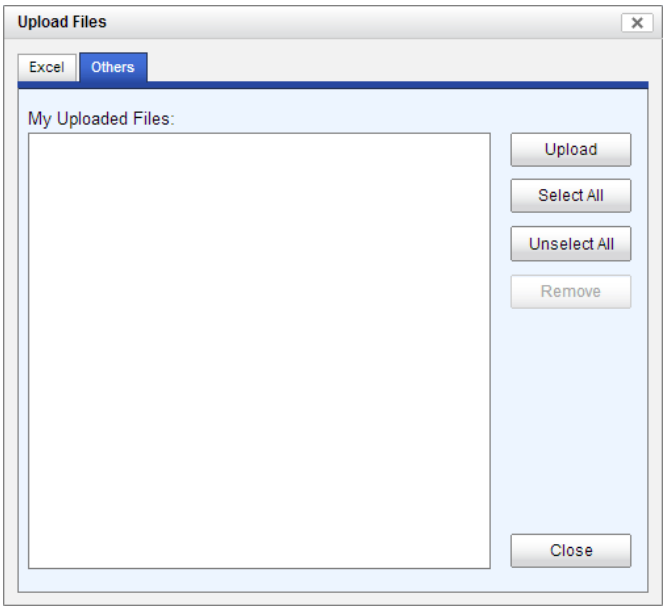
Data Selection

Allows you to select the data that will be used in the data source. By default, all non-empty cells are selected. If you want to exclude some columns/rows from the table or if an excel sheet contains several tables, you can narrow the selection in the *Data Selection* table. To do that, select a corner cell, then press the **SHIFT** key and then click on the opposite corner cell.

After you have finished configuring the data source, click on the *SAVE* button to insert the data source into your data registry.

2.1.2. Upload Other Files

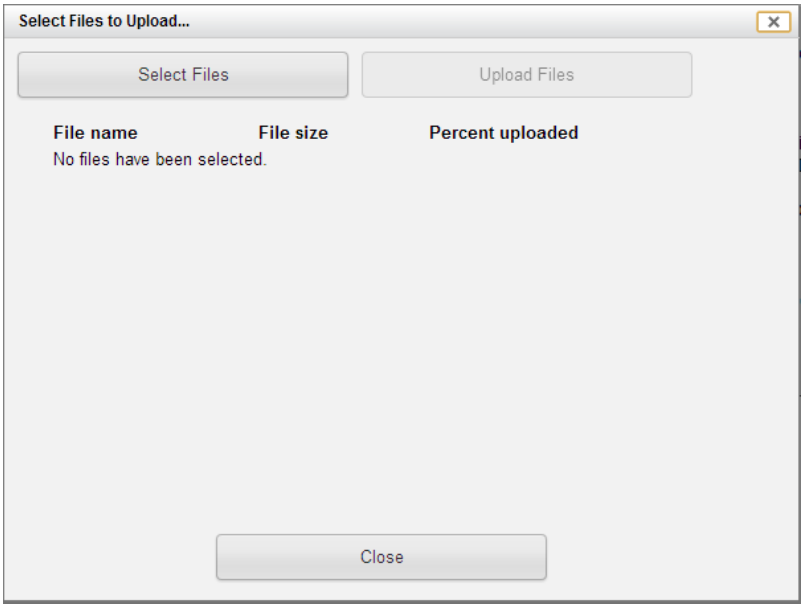
You can upload your own files from *Others* tab of the *Upload Files* dialog.



Upload Your Own Files

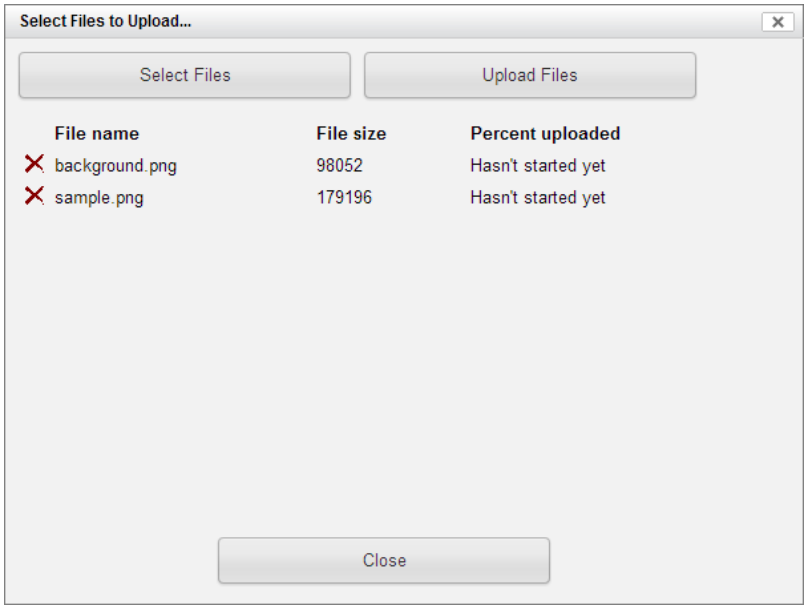
There are four buttons on the *Others* tab:

Upload Allows you to upload the files from your local computer. To do this click the *Upload* button. The next dialog will appear.



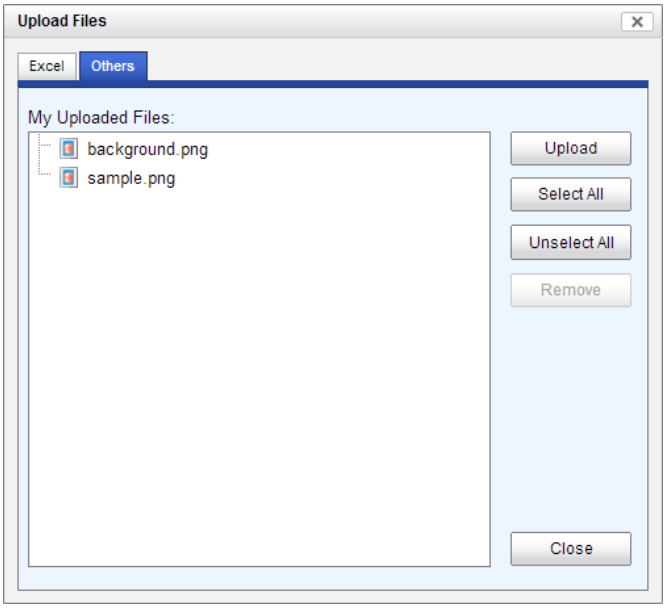
Select Files Dialog

Click the *Select Files* button. Your local computer directory will be opened to select the file. Select one or more files and confirm the selection. Your files are added to the *Select Files to Upload* dialog now.



Select Files To Upload

Click the *Upload Files* button. Your files are uploaded and you can see them in the *Upload Files* dialog. Your files were saved to the *Uploads* directory and you can use them by *Browse* function (e.g. from Report Designer).



Uploaded Files

Select All Allows you to select all uploaded files to remove them.

Unselect All You can unselect all selected files.

Remove Allows you to remove your uploaded files.

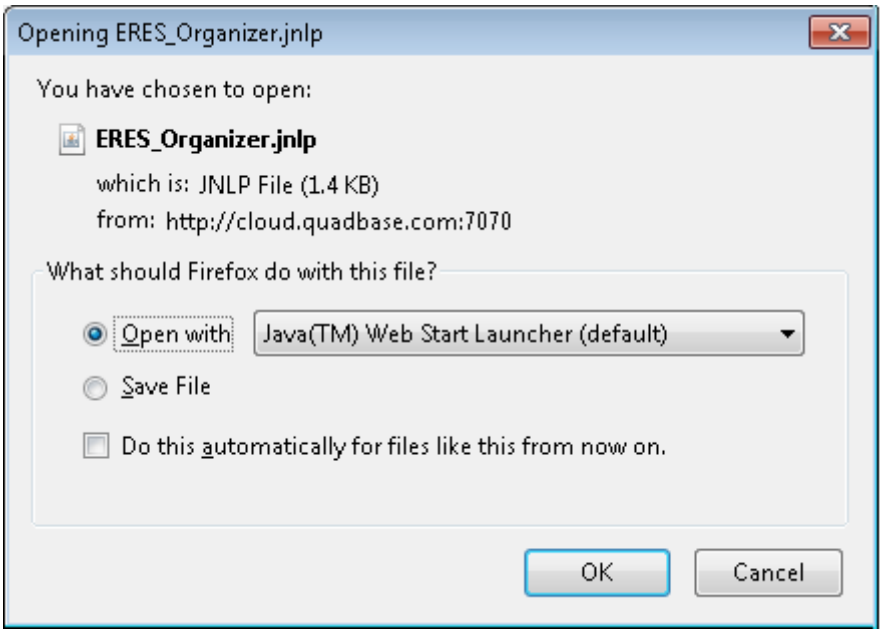
Chapter 3. Organizer

3.1. Using the Organizer

The ERES Cloud Organizer is a powerful graphical user interface that allows users to design charts and reports and organize them with other files in virtual folders. The Organizer contents are automatically published to the Menu page and users can automatically generate URLs to run reports and charts. In addition, schedule and archive jobs are created and maintained through the Organizer. The Organizer also includes the user permissions and security level setting options.

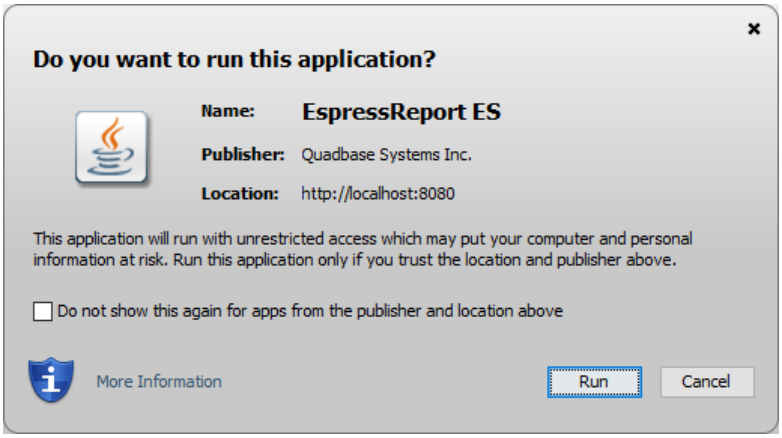
3.1.1. Starting the Organizer

To start the Organizer, click the link labeled *Organizer* in the Start page. This will open a pop-up window. Confirm `ERES_Organizer.jnlp` file open by Java Web Start Launcher. Then the Organizer will open.



Opening Organizer

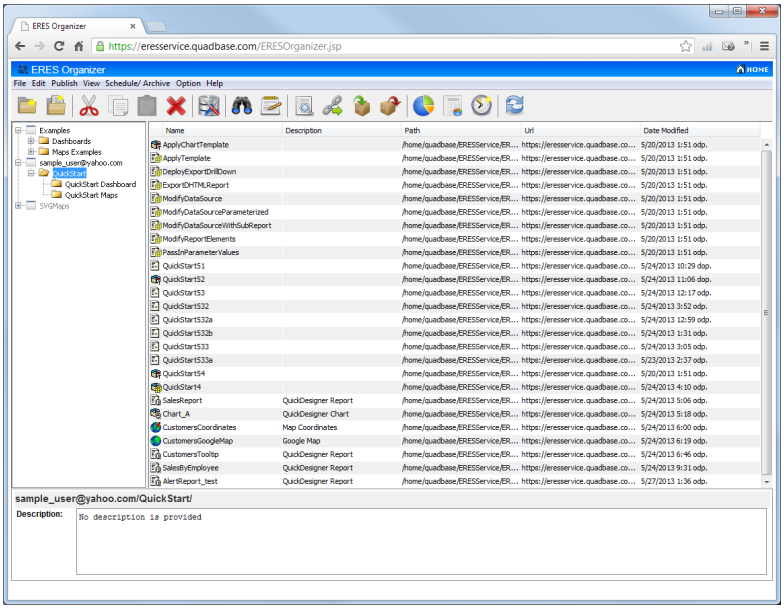
If you use Chrome as a browser, you have to save `ERES_Organizer.jnlp` at first, and then to open it from your download directory. You will see the next dialog. Click *Run* and Organizer will open.



Opening Organizer in Chrome

3.1.2. The Organizer Interface

Since the Organizer is at heart a file management system, it is no coincidence that the application window resembles a file browser. The left-hand side of the window contains a list of projects and their associated sub-folders, while the right-hand side lists the contents of a folder in the upper pane and the project/folder descriptions in the lower pane.



Organizer Interface

3.1.2.1. The Organizer Menus

Most of the Organizer functions can be controlled using the drop-down menus located at the top of the Application window.

3.1.2.1.1. File Menu

This menu performs most of the file operations. Options in this menu include:

New — If you select this option, another menu pops up.

- **Folder** — Creates a new folder in current project/folder.
- **Chart** — Launches ChartDesigner. (see Section 5.2.1 - Introduction to Chart Designer).
- **Report** — Launches ReportDesigner (see Section 5.1.1 - Introduction to Report Designer).
- **Dashboard** — Launches DashboardBuilder (see Chapter 7 - Designing Dashboards).

Manage Data Sources: — This option opens Data Source Manager that allows you to manage data files and database queries.

Open File: — This will open the currently selected file. If the file is a chart then the Chart Designer will open, allowing you to edit the chart. If the file is a report file then the Report Designer will open allowing you to edit the report.

Create PAK: — Pack a GXML (Online Map), SXML (SVG Map) or DSB (Dashboard) file and its components (Templates, Tooltips etc.) into a single file (see Section 3.1.4.5 - Dashboard and Map Packages).

Unpack File: — Unpack a GPAK (Online Map), SPAK (SVG Map) or DPAK (Dashboard) pack file into Organizer (see Section 3.1.4.5 - Dashboard and Map Packages).

Delete: This will delete the currently selected project, folder, or file.

Refresh: This will refresh your Organizer window, allowing it to reflect the latest changes/ modifications performed by other users.

Exit: This will close the Organizer.

3.1.2.1.2. Edit Menu

This menu allows you to edit, and cut/paste files and folders. Options in this menu include:

Insert File(s): This allows you to insert a report or chart into the selected project/folder.

Edit: This allows you to edit the currently selected project, folder, or file. For projects and folders, you can change the project/folder name as well as its description. For files, you can change the display name, the location path or URL, and the description.

Cut: This removes the selected project, folder, or file, and places it on the clipboard.

Copy: This makes a copy of the selected project, folder, or file, and places it on the clipboard.

Paste: This pastes the current clipboard project, folder, or file into the currently selected project or folder.

3.1.2.1.3. Publish Menu

This menu allows you to setup and view published reports and charts. Options in this menu include:

Preview Menu Page: This will load the ERES Menu page in a new browser window.

Generate Image URL: This option generates an Image URL for the currently selected chart.

Generate Report URL: This option generates a Report URL for the currently selected report.

Generate Dashboard URL: This option generates a Dashboard URL for the currently selected dashboard.

3.1.2.1.4. View Menu

This menu allows you to search for files or edit charts and reports. Options in this menu include:

Search: This allows you to search through projects, folders, and files.

Chart Designer: This will open the Chart Designer, allowing you to design or modify charts.

Report Designer: This will open the Report Designer, allowing you to design or modify reports.

3.1.2.1.5. Schedule/Archive Menu

The menu allows you to set and view options related to chart/report scheduling and archiving. Options in this menu include:

Set Schedule: This allows you to set a schedule for the currently selected chart, report, project, or folder.

Set Archive: This allows you to set an archive for the currently selected chart, report, project, or folder.

Set Alert Monitoring: This allows you to set an alert monitoring for the currently selected chart, report, dashboard, svg map, project, or folder.

View Schedule Tasks: This will open the log allowing you to view all the of the currently scheduled and archived jobs.

View Alert Monitoring Tasks: This will open the log allowing you to view all the current alert monitoring jobs.

3.1.2.1.6. Option Menu

Set Designer: This allows you to specify whether to open every chart or report in the same Chart Designer or Report Designer window (thus closing the current chart or report), or to open a new Chart Designer or Report Designer window for each chart or report opened.

3.1.2.1.7. Help Menu

This menu allows you to view a version of the program and to open a documentation. Options in this menu include:

About: This brings you an information about the version of the program.

Contents: This opens the ERES Cloud User's Guide.

3.1.2.2. The Organizer Toolbar

The toolbar at the top of the Organizer window offers easy access to some of the most commonly used features. The buttons perform the following functions:


Insert a new folder into the current project


Insert a new file into the current project or folder


Cut the currently selected project, folder, or file, and place on the clipboard


Copy the currently selected project, folder, or file, and place on the clipboard


Paste the current clipboard project, folder, or file into the currently selected project or folder


Delete the currently selected project, folder, or file


Manage Data Registries, files, queries, etc.


Search for projects, folders, or files


Edit the currently selected project, folder, or file


Preview the menu page


Generate a URL for the currently selected file


Pack a map or a dashboard and it's components into a single file (see Section 3.1.4.5 - Dashboard and Map Packages)


Unpack a GPAK (Online Map), SPAK (SVG Map) or DPAK (Dashboard) pack file into Organizer (see Section 3.1.4.5 - Dashboard and Map Packages).


Go to Chart Designer


Go to Report Designer


Set schedule for the currently selected chart, report, folder, or project


Refresh the Organizer window

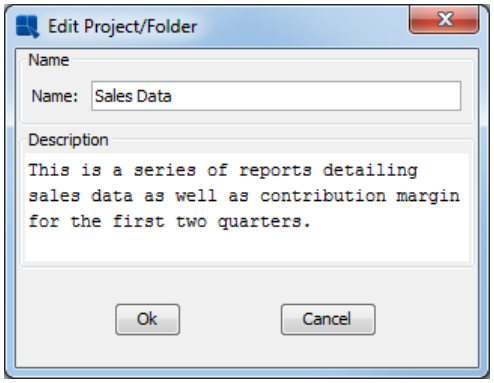
3.1.3. Projects and Folders

Projects and their associated folders are the central part of the Organizer interface. They provide the primary means of organizing charts and reports. Projects are the top level of organization. You can see read-only projects with examples and your project created automatically with the name of your email address. You can insert many layers of sub-folders into your project, creating just about any organizational structure.

It is important to note that all folders in the Organizer are virtual folders, meaning that they do not necessarily have a physical location on the server machine or mirror an actual directory structure. Projects and folders in the Organizer are pointers to real directories and files that may be stored in a variety of locations.

3.1.3.1. Creating and Modifying Folders

To insert a sub-folder into your project or folder, first select the project or folder in which you would like to place the new sub-folder. Then you can either click the *New Folder* button on the toolbar, select *New Folder* from the File menu, right-click in the left-hand pane and select *New Folder* from the pop-up menu, or press **Ctrl**+**F**. A node labeled *new folder* will appear under the currently selected project or folder.



Edit Project/Folder Dialog

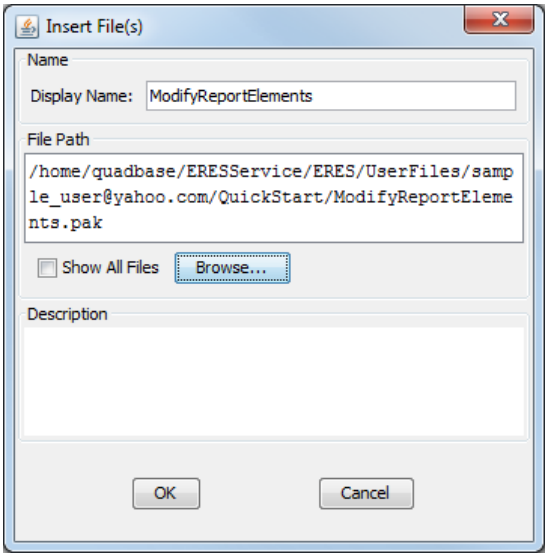
To edit the folder name and description, first select it, and then you can either click the *Edit* button on the toolbar, select *Edit* from the Edit menu, press *F2* key, or right-click, and select *Edit* from the pop-up menu. A dialog will then appear prompting you to change the folder name and description. Once you have entered your desired changes, click *Ok*, and the new folder name and description will be displayed in the lower right-hand pane.

3.1.4. Files

3.1.4.1. Adding and Modifying Files

To add a file, first select the folder or your project in which you would like to place the file. Then you can either click the *Insert File* button on the toolbar, select *Insert File* from the Edit menu, or right click in the upper right-hand pane, and select *Insert File* from the pop-up menu. A dialog box will appear, prompting you to enter a display name for the file (by default the file name is used), the location of the file (a direct path) and a description of the file. (If you select a path that is not the server root or a mapped virtual directory, you will be prompted to map one.) Click *Ok* and the file will be inserted into the current directory or folder.

To insert more than one file at a time, click the *Browse* button on the insert file dialog, and use **CTRL**+**Click** to select multiple files in the browse dialog. Then when you click *OK* all of the files will be added to the Organizer.



Insert File Dialog

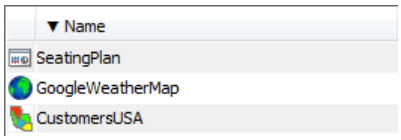
The list of files in the currently selected folder will be displayed in the upper right-hand pane of the Organizer. Each entry shows the file's display name, the file type, the file description, path, and URL.

To edit the file's display name, location, or description, first select the file that you would like to edit. Once you have selected a file, you can either click the *Edit* button on the toolbar, select *Edit* from the Edit menu, press *F2* key, or right click, and select *Edit* from the pop-up menu. A dialog box will appear allowing you to make changes to the display name, location, and description. Click *Ok* and the changes will be reflected in the file list.

To open a file, double click on it, or first select it, and then you can select *Open File* from the *File* menu, or right click and select *Open File* from the pop-up menu. If the file is a chart (PAC, CHT, TPL, or QCH) format, it will be opened in the Chart Designer. If the file is a report (PAK, RPT, QPT, or XML) format, it will be opened in the Report Designer. If the file is a dashboard, the Dashboard Builder interface will open in a new window. If the file is a map, Online Maps interface will open in a new window. All other file types, such as SVG, will be opened in a web browser window.

3.1.4.2. Arranging Files

To arrange files within folder or project, open the folder (click on it in the left-hand side tree-list) and then click on a column header in the right-hand side. Files in the folder/project will be arranged by the content of the selected column. To switch between ascending and descending order, click on the column header again.



A folder arranged by the file names in descending order.

3.1.4.3. File Identification Icons

Each file type can be easily identified by its icon. You can see these icons next to file names in Organizer and all dialogs in ERES designers where you need to select some file.



Report file created in Report Designer that contents supplementary files associated with it (`.pak`). For more information see Section 5.1.5 - Saving Reports.



Report template created in Report Designer (`.rpt`) in older versions of ERES (6.6 and older).


Chart file created in Chart Designer (`.cht`). For more information see Section 5.2.2.3.1 - Saving Chart Definitions.


Chart template created in Chart Designer (`.tpl`). For more information see Section 5.2.2.3.1 - Saving Chart Definitions.


Chart file created in Chart Designer that contents supplementary files associated with it (`.pac`). For more information see Section 5.2.2.3.1 - Saving Chart Definitions.


Report file created in QuickDesigner Reports (`.qdr`). For more information see Section 5.3 - QuickDesigner Reports.


Chart file created in QuickDesigner Charts (`.qch`). For more information see Section 5.4 - QuickDesigner Charts.


Map file created in Online Maps (`.gxml`). For more information see Section 6.2 - Online Maps.


Packed map file created in Online Maps and packed in Organizer (`.gpak`). For more information see Section 3.1.4.5 - Dashboard and Map Packages.


Coordinates file for Online Maps (`.cxml`). For more information see Section 6.2.1 - Data Sources.


SVG map file created in SVG Maps (`.sxml`). For more information see Section 6.3 - SVG Maps.


Packed SVG map file created in SVG Maps and packed in Organizer (`.spak`). For more information see Section 3.1.4.5 - Dashboard and Map Packages.


SVG map image (`.svg`). For more information see Section 6.3.1 - Introduction to SVG Maps.

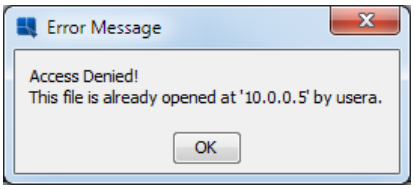

Dashboard file created in Dashboard Builder (`.dsb`). For more information see Section 7.1 - Introduction to Dashboards.


Packed dashboard file created in Dashboard Builder and packed in Organizer (`.dpak`). For more information see Section 3.1.4.5 - Dashboard and Map Packages.

3.1.4.4. Report Locking

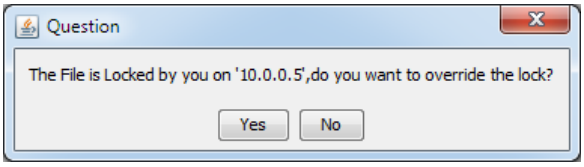
The ERES Organizer has an automatic report locking system designed to block multiple users from editing the same report in Report Designer at the same time. The system utilizes your IP address, your username, and a special token to determine whether you are allowed to open a specific file. Here is an example illustrating how the advanced file locking system operates.

Suppose you are logged in as userA and you are editing a report `Sales.pak`. One of your colleagues logs in as userB and attempts to edit the same `Sales.pak`. She would see the following error message, alerting her that the file is currently locked.



File is Locked

Suppose after opening the file, your browser crashed or you moved to a different computer. Now you log back in as userA and open the same file again. You will see the following message allowing you to override your own lock.



Override Your Lock

If you moved to another computer and the report is still opened on the previous computer, you will not be able to save on the previous computer. You must close the report and re-open it in order to have write privilege again.



Note

The admin has the ability to overwrite any lock.

3.1.4.5. Dashboard and Map Packages

A single map or a single dashboard usually consists of several files (for example, a dashboard needs several chart/report/map templates, a svg map needs a svg image, drill-down template, etc) which could make moving maps or dashboards to another machine very difficult, because you would have to move all the linked files as well. Fortunately, you do not have to move each file manually. A dashboard or a map and can be packed into a single file along with all the linked files (dashboard templates, drill-down templates, etc) without even knowing which exact files are needed for the dashboard or map to work. The packed files can be moved or archived quickly and easily.

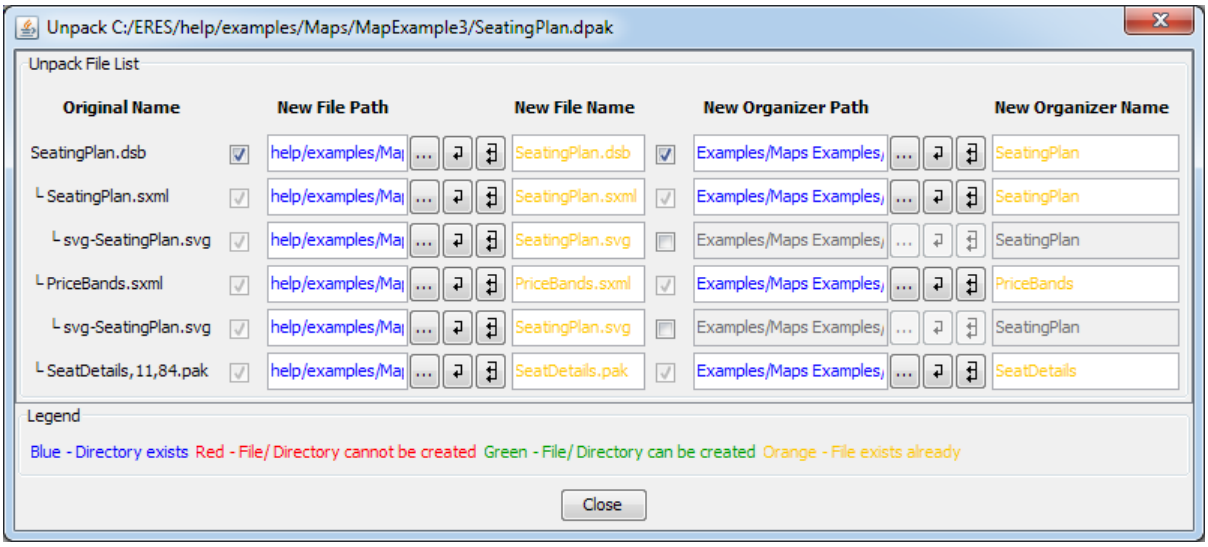
To pack a dashboard or a map (both SVG and Online maps can be packed), select the dashboard (`.dsb`) or map (`.sxml` or `.gxml`) file in Organizer and then click on the  *Create PAK* icon on the main Organizer toolbar. A dialog will pop up prompting you to select the pack file name. Enter a file name or use the default one and click *OK*. Another dialog should show up confirming that the package was created successfully and showing you the full file path. The packed file will be also automatically inserted into the current Organizer folder/project.

Packed files can be viewed in the Menu page without being unpacked, but they can not be modified in the SVG and Online Maps designers or in the Dashboard Builder. In order to modify a packed dashboard or map, you need to unpack the gpak, spak, or dpak file first.

To un-pack a dpak, spak, or gpak file that has not been inserted into Organizer, click on the  *Unpack File* icon on the main Organizer toolbar. A dialog will pop up, allowing you to choose the file on your hard drive and to enter a name which will be used for the main unpacked file (`.dsb, .sxml` or `.gxml` file).

If you want to unpack a file that has been inserted into Organizer, select the file in Organizer and then click on the  *Unpack File*. The same dialog will pop up but the file path and the file name fields will be filled automatically.

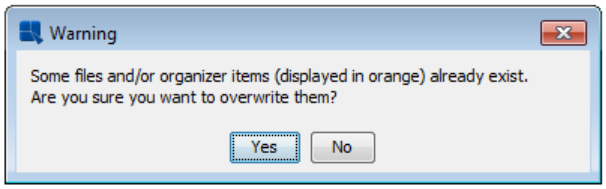
To view the list of all files that are stored in the package, click on the *Details* button. A new window will open.



Unpack File List Dialog

In this window, you can edit both physical and Organizer file path of all files that can be un-packed from the package (which also allows you to resolve potencial file name conflicts). After you are done editing the file paths, click the *Close* and then click *OK* to unpack the package.

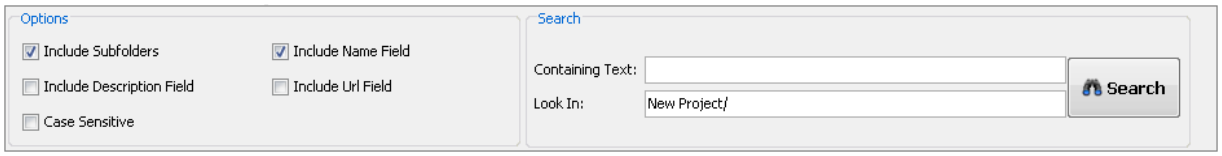
The ERES will now try to unpack all files from the packed file to their original locations. If there are any filename conflicts (i.e. there already is a file with the same name in the same path on your hard drive or in Organizer), the following dialog will show up:



If you click *Yes*, all existing files will be overwritten by the files from the package without no further questions. If you click *No*, the *Unpack File List* dialog will open again allowing you to edit the filename conflicts. Then click *Unpack* to unpack the files into their locations and to insert them into Organizer.

3.1.5. Searching in Organizer

Often, if you have many files in many folders in the Organizer, it can be hard to find a specific file. To locate files or folders, you can use the search function. To search, you can either click the  *Search* button on the toolbar, select *Search* from the View menu, press **Ctrl**+**F**, or right click in the upper right-hand panel and select *Search* from the pop-up menu. The search dialog will then appear.



Search Dialog

The search dialog prompts you to enter a search string in the project or folder in which you would like to search. By default, the currently selected project or folder is listed. If you leave the *Look In* field blank, it will search all projects.

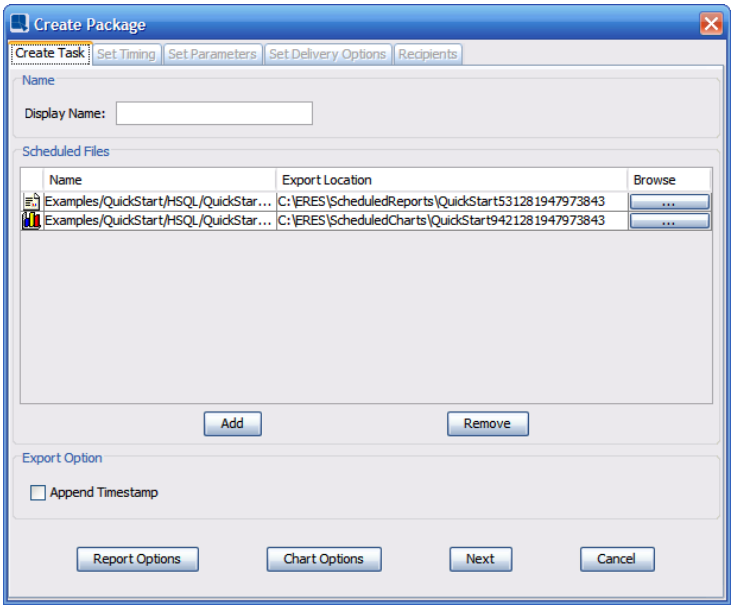
Other options allow you to include sub-folders in your search, and to match the search string to filenames, descriptions, and URLs. You can also match case in your search by selecting the *Case Sensitive* option.

3.2. Scheduling & Archiving

ERES is capable of scheduling/archiving any chart (`.pac, .cht` or `.tpl`) or report (`.pak, .rpt` or `.xml`) file that is placed in the Organizer. A schedule will generate an exported file for the report or chart at a regular interval specified by the user. An archive will generate a report or chart template with the current data stored in the template. The periodicity of schedule/archive can be set for specific time intervals, or for fixed days of the week or dates of the month.

3.2.1. Setting a Schedule

To schedule a chart or report, first, select the file in the Organizer that you would like to schedule. You can also select multiple files at once. Then click the *Set Schedule* button on the toolbar or right click on the selected file and select *Set Schedule* from the pop-up menu. The first dialog of the scheduling process will open.



Create task dialog

On this Wizard step, enter a *Display Name*.

In the *Scheduled files* list, you can see what files are being scheduled. You can also remove a file from the scheduled task or add a file.

The dialog has following objects:

Display name: scheduled task name.

Scheduled files:

- **Icon -** indicates the type of scheduled file.

   - Chart

   - Report

- **Name -** name of the file
- **Export location -** directory path. The file will be exported to the location at the time when the task is run. You can double click this field to alter the path manually.

  You can access the exported file from MenuPage by clicking on the  *Schedule* icon (as described in Section 8.1.2.1 - Viewing Reports and Charts).

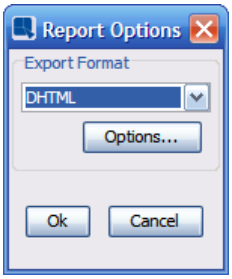
- **Browse -** clicking the "**...**" field will bring up a browse dialog allowing you to change file export location. This has the same effect as double clicking the *Export location* field, only this option allows you to browse file structure.

Add - this button allows you to add more files into the scheduled task.

Remove - removes selected file from the scheduled task.

Append timestamp - appends current time (i.e. time when the task is run) to the exported file's name. Also, if timestamp is appended to filename, the filename is always unique thus every time when the task is run, a new file is added into the export folder. If you disable the *Append timestamp* option, only one exported file will be kept in the export folder and it will be overwritten every time the task is run.

Report options - shows dialog allowing you to set report export format. For reports you can select DHTML, PDF, CSV, Excel (XLS), Excel 2007 (XLSX), text, rich text, or XML. Additional options allow you to select single page or multi-page exporting for DHTML and XML. The *Options...* button allows you to set format-specific options for the exported file. For DHTML you can set single or multi-page exporting. For DHTML you can also specify which CSS options to use. For PDF you can set encryption and for text you can specify the delimiter.



Report Options Dialog

Chart options - shows dialog allowing you to set chart export format. For charts you can select JPG, GIF, PNG, PDF, SVG, FLASH, or EXCEL. Other options allow you to specify PNG compression, set JPEG quality, and specify background transparency for GIFs. You can also enable and configure Flash animation if the image depiction should be animated. You can also specify the size of the exported image.

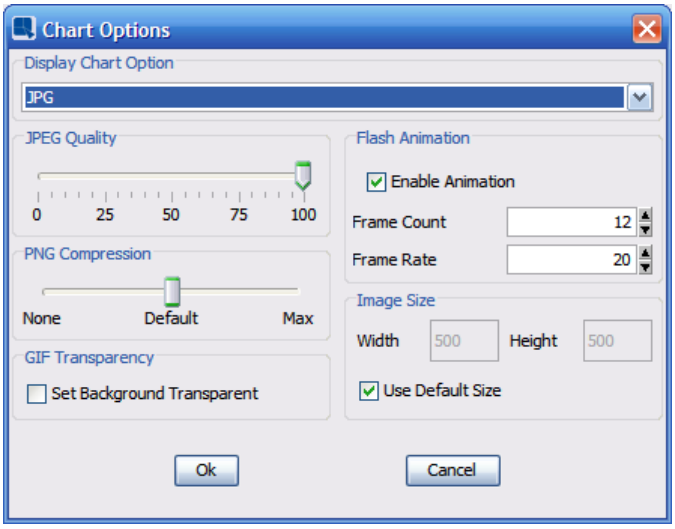
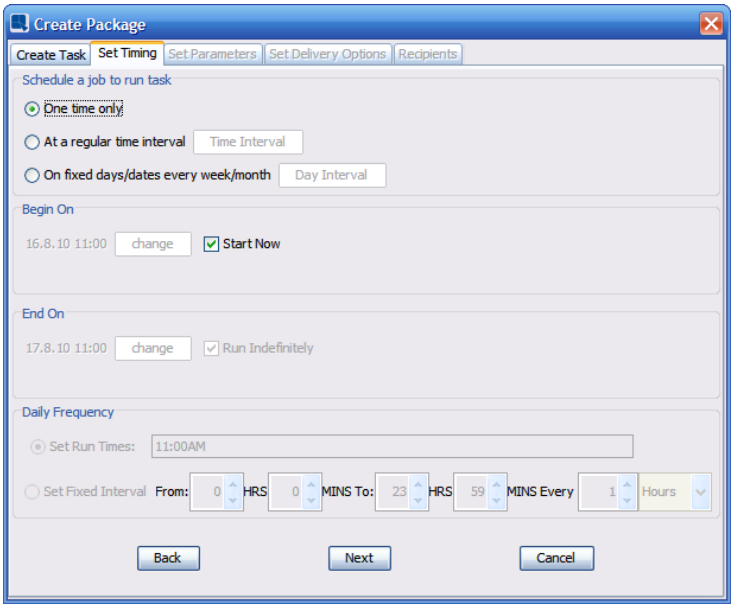


Chart Options Dialog

After you are done setting the first tab options, you can click *Next* button and proceed to the *Set timing* tab.



Set Timing Dialog

Set timing tab allows you to set a time pattern of the scheduled task.

You have following timing options:

One time only: task is run one time only at pre-defined time. You can choose when the task will be run using the *Begin on* option (will be described later).

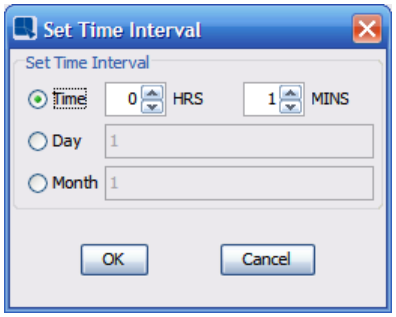
At regular time interval: scheduled task will be run at fixed intervals for a pre-defined period of time. For example, you can set the task to be run every X hours (or minutes, days, months...) from *Begin on* to *End on* time and date. When you select this option, the *Set time interval* dialog pops up. This dialog will be described later.

On fixed days/dates every week/ month: If you need to run your task irregularly, this option should do the trick. It allows you to run the task on certain week day(s) (for example, every Monday and Friday), or certain days of the month (for example, every the first of each month). This is set in the *Set fixed days/dates Interval* dialog which pops up when you select the option. Also, when you select this option, the *Daily frequency* section is enabled. These two dialogs will be described later.

You also need to set when the task should start running and when it should stop. You can do this using the following options:

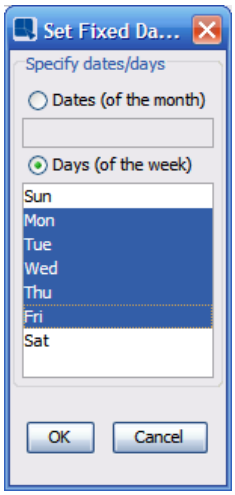
Begin on: allows you to set when the task will start running. By default, the *Start now* option is enabled. It means that the task start time and date will be set two minutes ahead of the time when the creation of the schedule job is finished. If you do not want the task to start immediately, uncheck this option to pick the start time and date from a calendar. You can access the calendar by clicking the *Change* button (which is available only when the *Start now* option is disabled).

End on: controls when the task will end. This option is very similar to the *Begin on* option. The only difference is that the *Start now* option is replaced by *Run indefinitely* option. If you choose this option, the task will never stop unless you stop it manually. You can also unselect this option and pick the *End on* time and date manually from a calendar.



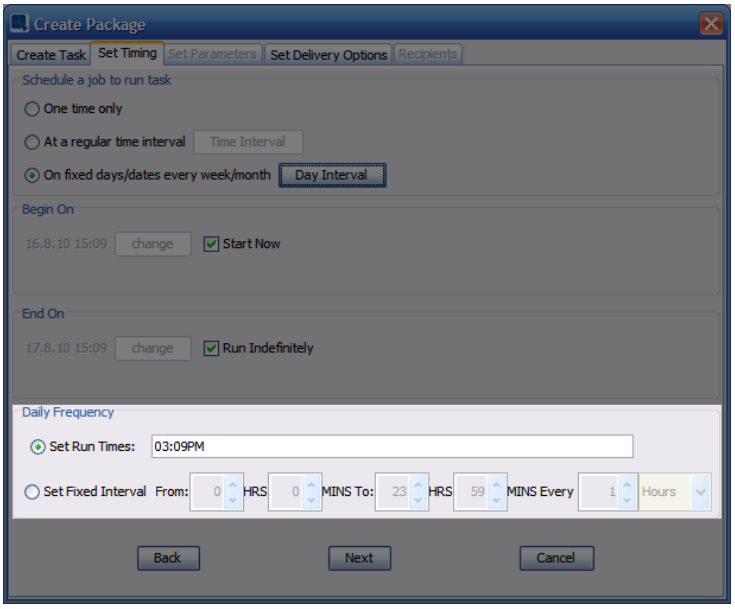
Set Time Interval Dialog

This dialog is enabled only if the *At regular time interval* option is selected. First of all, choose a time unit (namely, hours and minutes, days or months) that suits your needs. Next, enter a number of the time units that should be waited for between each task runs.



Set Fixed Days/Dates Dialog

This dialog is enabled only if the *On fixed days/dates every week/month* option is selected. First, choose whether you want to run the task on a weekly basis or on a monthly basis. Next, select days of the week or month task should be run on. After you are done setting this dialog, click *OK* and proceed to the *Daily frequency* settings.



Daily Frequency Settings

After you have set the days of the month/week you want the task to be run on, you have to set when exactly should it be run on those days. That can be done in the *Daily frequency* section. You have two basic options:

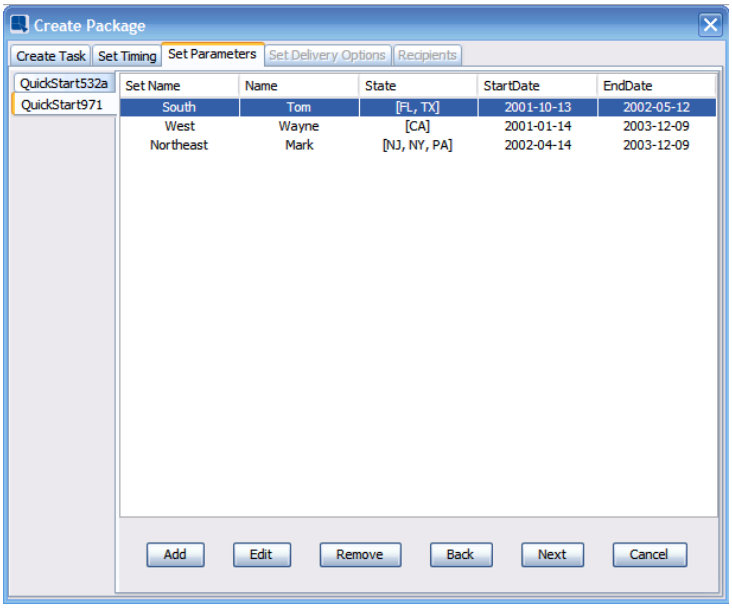
Set run times - allows you to manually enter the run times. You can add as many values as you want. You can use 12-hour or 24-hour format. Individual values are comma separated. **For example:** 08:05AM, 11:15PM, 22:35

Set fixed interval: use this option to run the task periodically from the *From:* time to the *To:* time every X hours/minutes. To set the times, you can use up and down arrows or you can enter the value manually (click in the field and type a number). If you enter an incorrect time (for example, 11:65), the value will be marked with red color. Correct values are marked with green color.

Once you have finished specifying the periodicity, click *Next* to continue with the Schedule Wizard.

3.2.1.1. Scheduling Parameterized Reports/Charts

If you have selected to schedule a report or chart that contains parameters, the next tab will appear allowing you to set the parameter values that you would like to use.



Set parameters dialog

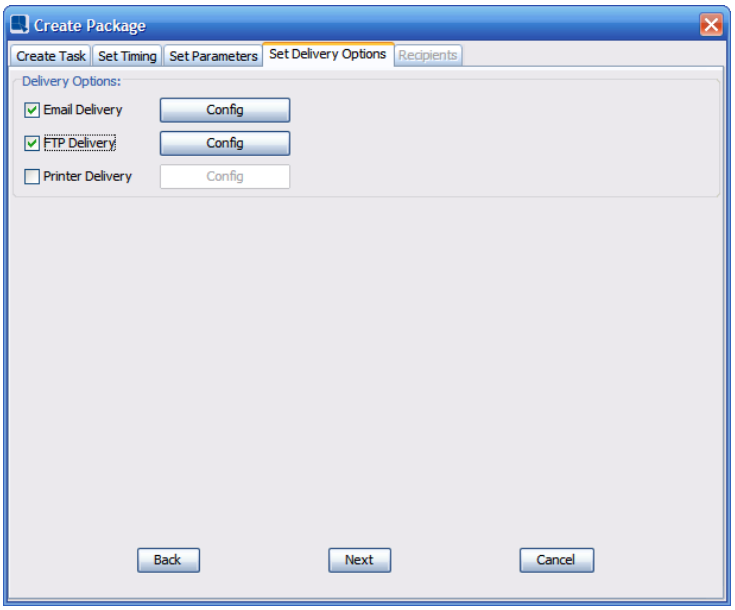
To add a set of parameters, click the *Add* button. This will bring up the parameter prompt dialog allowing you to select the set parameter values that you would like to use. Once you have selected a parameter set, your choices will appear in the dialog. You can add as many different combinations of parameter sets as you want. A separate file will be generated for each parameter set.

You can also specify a name for each parameter set by double clicking on the first column. The name specified here will be used in later dialogs to help you organize your recipients.

Once you have added all parameter sets, click the *Next* button to continue.

3.2.1.2. Additional Delivery Options

After you have finished setting up the schedule, you will be prompted to select delivery methods. These options allow you to email the generated files, print the exports, and upload the exported files to a FTP server. You do not have to select any additional delivery options.

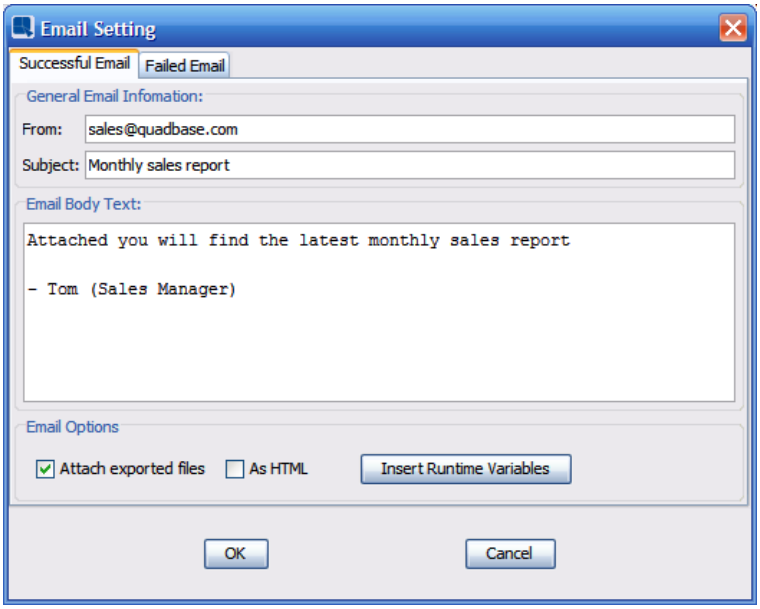


Delivery Options Dialog

If you select a delivery option, click the *Config* button to configure the delivery method.

3.2.1.2.1. Email Delivery Options

When you click the 'Config' button for email delivery the following dialog will open allowing you to set an email message:



Email Settings Dialog

The first tab contains information for the email sent when the report or chart has successfully exported. The *General Email Information* portion of this tab allows you to specify the from addresses for the email as well as the subject.

The first tab is called *Successful email* and it allows you to set an email message which will be sent in case the scheduled task finished with no problems.

There are the following elements:

From: enter an email address into this field. This address will be used as an email sender's address. Please make sure that your SMTP server will not reject the address (depends on your SMTP server security settings).

Subject: email subject. Enter any convenient text in here.

Email Body Text: enter a message that will be sent on every task run. You can enter any text and also, you can use *Runtime variables* (will be described later).

Attach exported files: allows you to attach all exported reports and charts to the email.

As HTML: if this option is enabled, the email body text will be replaced by the scheduled report or chart in HTML format. This option is available only if the *Attach exported files* is enabled and there is only one scheduled report/chart in the scheduled task. Also, PDF, SVG and Flash charts cannot be added into HTML emails and only DHTML reports can be added into HTML emails. You can set the export format on the *Create task* tab, as described in Section 3.2.1 - Setting a Schedule.

Insert Runtime Variables - allows you to insert certain variables into the email body text. Variables are replaced by real values at the task run time. All variables also have a “default value”. If no convenient data can be inserted into the variable, the default value is used (for example, when user has not provided his/her name). You can choose between following variables:

<USERNAME> - recipient's user name. For example: you can use this variable in a salutation like “Hello <USERNAME>”. This variable works only for email recipients that were chosen from the ERES database (will be described later). If you add an email address to recipients list manually, no username is displayed because no username was provided.

<FULLNAME> - recipient's full name. Just like the <USERNAME> variable, this variable also does not work for manually entered recipients email addresses.

<EMAIL> - recipient's email address.

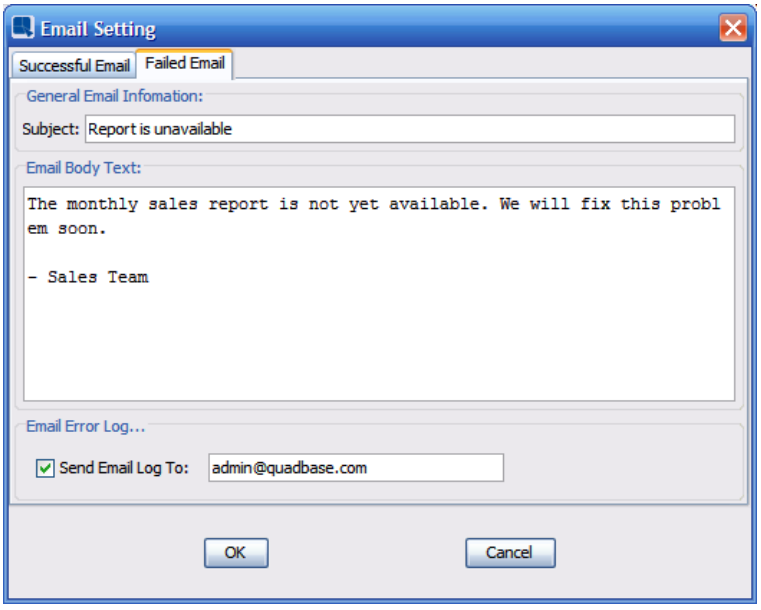
<TIMESTAMP> - current time and date. Shows when the email was sent.

<PARAMETERS> - list of all parameter sets and parameter values.

<SAVED_LINKS> - links of all scheduled files. Users can access the exported files through these links.

<LIVE_LINKS> - links of all scheduled files. Unlike the <SAVED_LINKS>, these links lead to the DHTML Viewer which allows users to access current version of the scheduled files.

The *Failed Email* tab allows you to send a different email (subject and body) to your recipients if an error has occurred during the task running (for example: one of the scheduled reports has been removed).

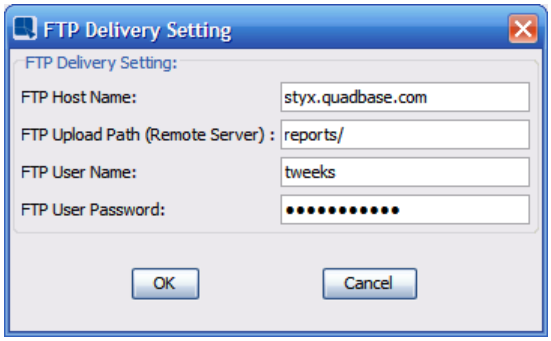


Failed email settings dialog

This tab also allows you to send the error logs to a specific user. Typically, this would be the admin or technical user who is able to diagnose and troubleshoot the problem.

3.2.1.2.2. FTP Delivery Options

When you click the *Config* button for FTP delivery, the following dialog will open, allowing you to set options for FTP schedule delivery.



FTP Delivery Options Dialog

The following information is required for FTP schedule delivery:

FTP Host Name: This is the host name of the FTP server to which you would like to upload the reports. If you have a Secure FTP Server (SSH FTP), enter **`sftp://`** in front of the hostname.

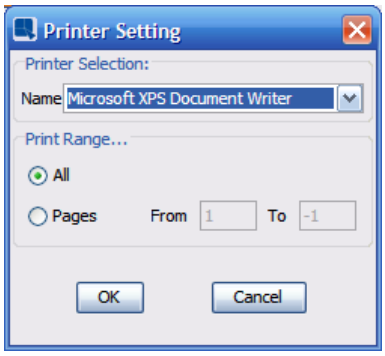
FTP Upload Path: This is the path or directory (relative to the FTP root) where you would like to write the exported report files.

FTP User Name: This is the user name for the FTP server.

FTP Password: This is the password for the FTP server.

3.2.1.2.3. Printer Delivery Options

When you click the *Config* button for printer delivery, the following dialog will open allowing you to set options for the print schedule delivery.

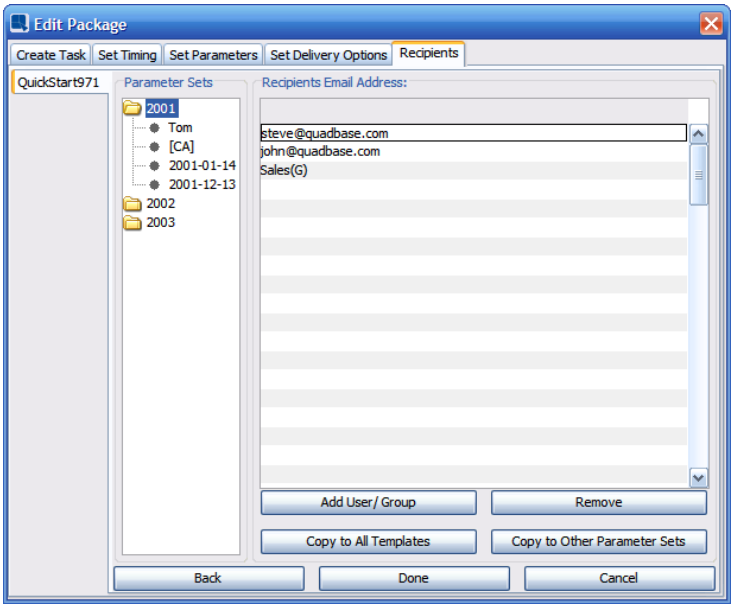


Printer Delivery Options Dialog

The first drop-down dialog contains a list of all the available system and network printers. You can select which one you would like to use. The second option allows you to select a page range for printing. You can print all of the pages or set a specific range. **`-1`** represents the last page. **Example,** if you want to print all of the pages except for the first one, set the range to *From* **`2`** , *To* **`-1`**.

3.2.1.3. Specifying Email Recipients

If you selected the email delivery, the next tab allows you to specify the recipient list. The recipient list is set up differently depending on the type of report or chart you are scheduling. If the report or chart contains parameters, you will see the following dialog. This allows you to send different parameter sets to different users or groups.



Set Recipients Dialog for Parameterized Reports

Entering recipients into the list is simple, click on the *Add User/Group* icon to select users and groups from the ERES database. You can also add regular email addresses by double clicking any row in the recipients list and typing it in manually. You may also click on *Add User/Group* and view the *Prev. List* tab for the previously listed recipients.

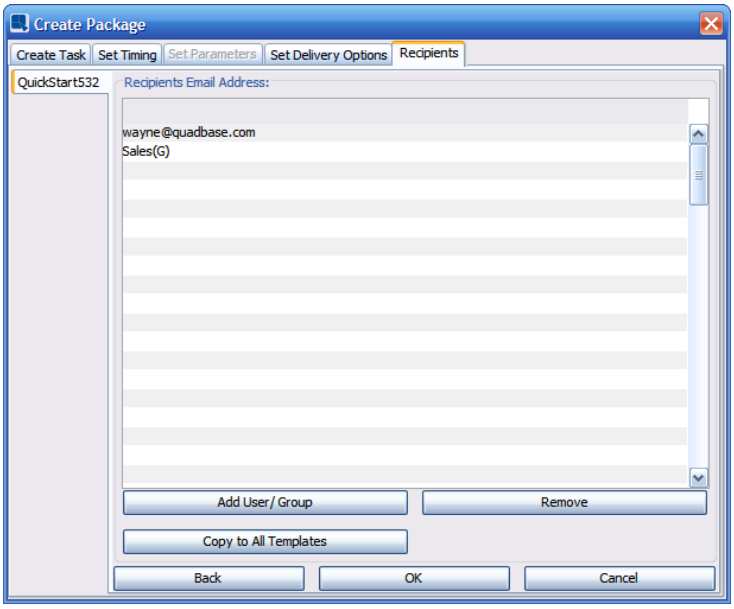
If you want to send all parameter sets or all templates to the same recipients, you do not have to add the recipients manually into all parameter sets and templates. Just add the recipients into one recipients list and then use one of the following buttons:

Copy to all templates: click this button to copy current recipient list into all scheduled templates (reports, charts) recipient lists.

Copy to other parameter sets: allows you to copy current recipient list to the rest of the parameter sets.

If you are scheduling a grouped report (such as a summary break report) and the report is using a database as a datasource, you will see the option for report bursting as well. Screenshots and detailed information for report bursting can be found in the Section 3.2.1.3.1 - Report Bursting.

For all other cases, you will see the following dialog containing only the recipients list.

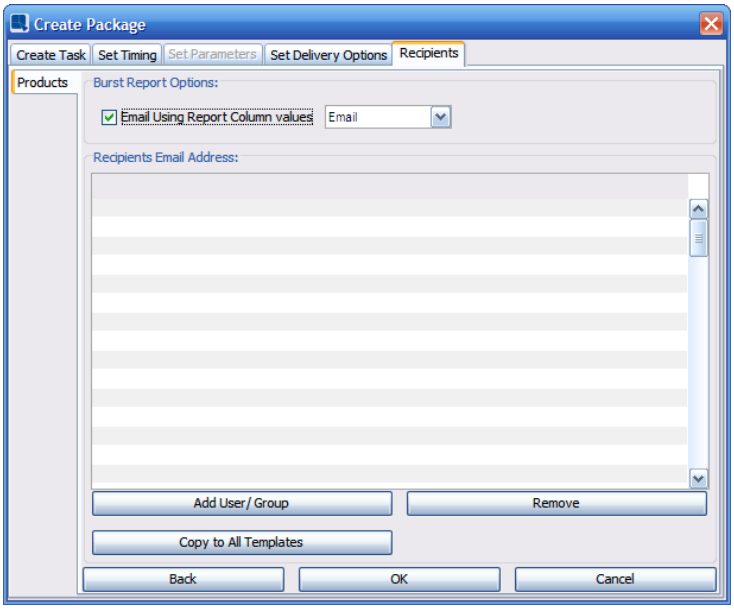


Set Recipients Dialog

3.2.1.3.1. Report Bursting

If the scheduled report contains grouped data (like a summary break report), you can enable “report bursting”. When a report is burst, the main report is divided into smaller reports and each group report is then sent to a different recipient. The recipients are taken from a database column so you do not have to enter them manually for each group of data. This feature offers significant performance and security improvements over running many small reports. The report burst option is available only for database data sources (MySQL, MS SQL, Oracle etc...).

If the report contains grouped data, the recipients dialog will contain a bursting option at the top.



Report Bursting

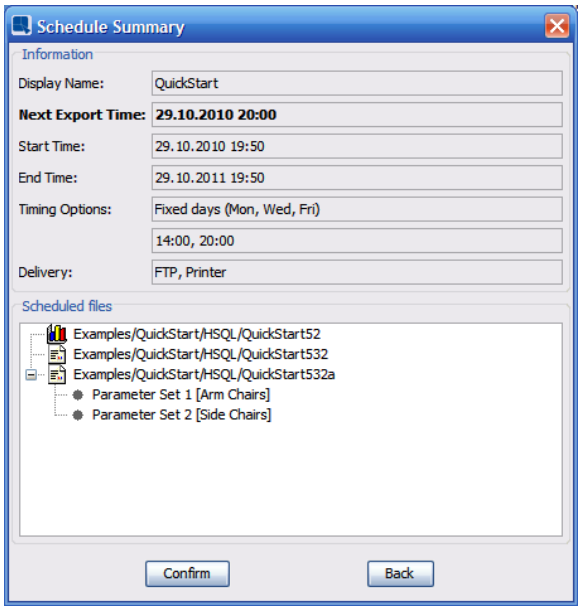
To set up the burst report feature, the report has to have a column that contains email addresses that correctly correspond to the group by (row break) column. For example, if the report is an account statement grouped by a customer id column, the report must also contain a column that contains the email address for each customer. Note that the email column does not need to be visible in the report in order to use the bursting feature.

If the report meets the aforementioned requirements, you can enable report bursting by selecting the option *Email Using Report Column Values* at the top of the dialog. This will run a single query for the report, break the report up by groups, export each group, and deliver the exports to the email address listed for each group.

3.2.1.4. Schedule Summary

After you confirm the last step of the *Create Package* Wizard, the *Schedule summary* dialog pops up. It offers you a quick overview of the scheduled task you have just finished configuring. If you are satisfied with the task's settings, click *Confirm*. If not, click *Back* to go back to the *Create Package* wizard.

It is recommended to always check the *Next Export Time*. If you are not satisfied with the next export time, it indicates that the timing has been set wrong. You can click *Back* and revise the timing settings.

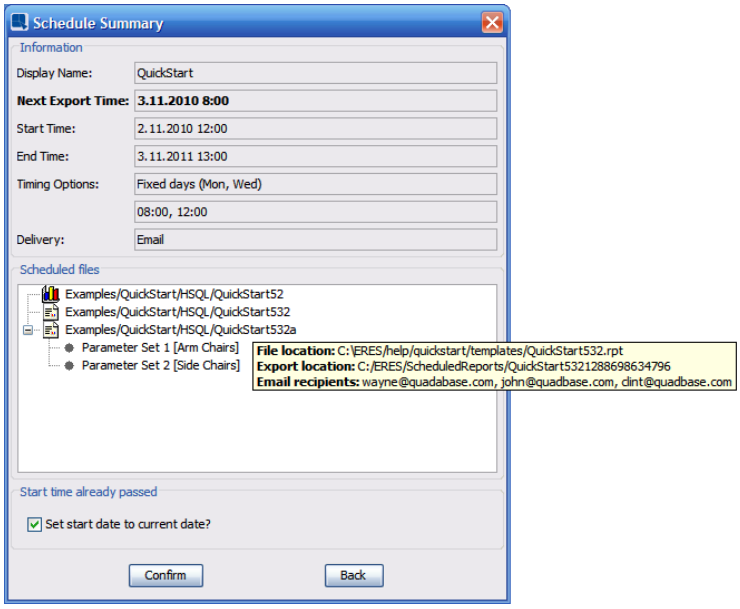


Schedule Summary

The *Schedule Summary* dialog has only one setting - *Set start date to current date*. This setting appears only when the start time has already passed. If you select this option, start time will be automatically set to current date/time.

Move the mouse over a scheduled object (report, chart, etc) in the *Scheduled files* list to display file and export locations.

If the *email delivery* is enabled, the tooltip also displays email recipients for the scheduled object. If the object has some parameters, there is a tooltip for each parameter set displaying email recipients.

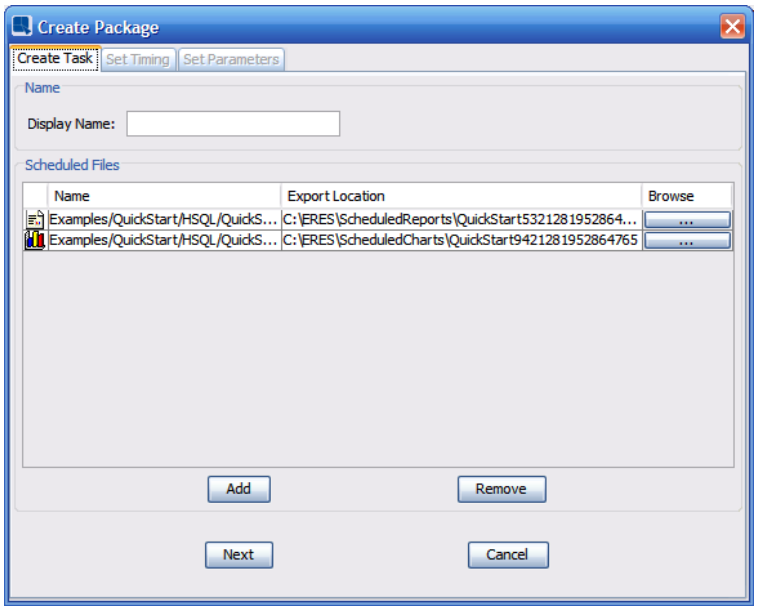


Schedule Summary

3.2.2. Setting an Archive

In addition to scheduling, which periodically exports the chart or report to a static file, ERES also supports archiving. An archive is similar to a schedule job except that instead of exporting a static copy of the report or chart, a version of the chart or report template is saved. The archived template contains all the data for that report or chart, thus saving a historical snapshot of the data that can be run/viewed at a later date.

To setup an archive, select the report or chart that you would like to use in the Organizer, and select *Set Archive* from the *Schedule/Archive* menu, or right click and select *Set Archive* from the pop-up menu. The following dialog will appear allowing you to set the name and location of the archive:



Set Archive Dialog

From this screen you are prompted to specify a *Display Name* for this archive. You can also select to specify the location and file name (without extensions) for the saved archive or allow the server to assign a location and file name. Note that each time the archive runs a new file will be created (the file name is appended with a timestamp)

creating a distinct file for each version. Archived files can also double as backup that can be used to restore old report/chart template version.

Set archive wizard is just a simpler version of the *Create scheduled task* wizard that was described in the previous chapters.

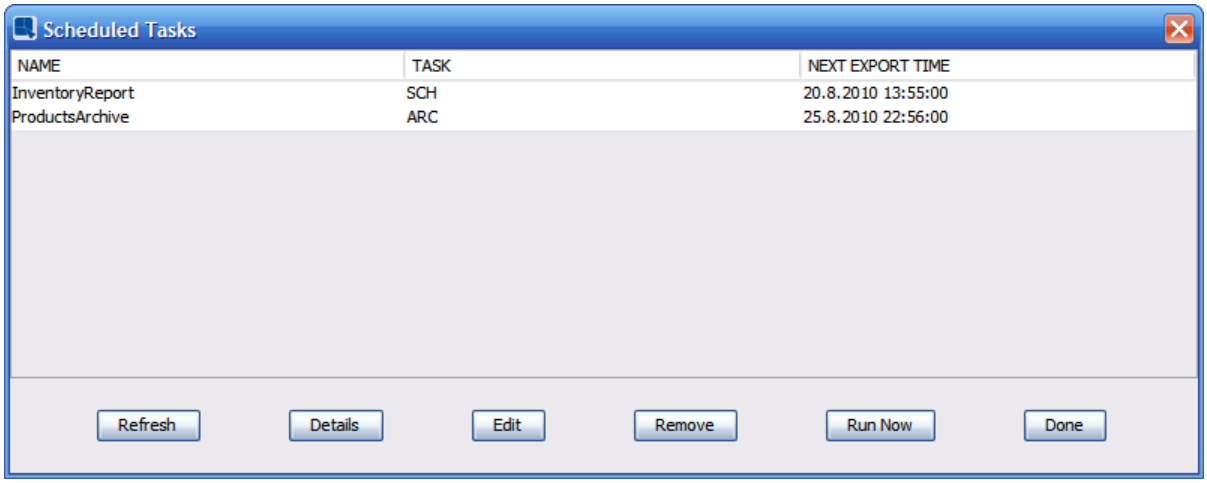
Once you have set the initial options for the archive, the next two tabs allow you to set the periodicity for the archive, and supply parameter sets (if any). These operations are exactly the same as for setting a schedule discussed in Section 3.2.1 - Setting a Schedule.

Archives cannot be delivered since they are in templates rather than exported formats. They can be accessed from the menu page, which will show a special icon  for files that are archived (as described in Section 8.1.2.1 - Viewing Reports and Charts).

3.2.3. View Tasks

You can see a list schedules or archives, by selecting *View schedule tasks* from the *Schedule/Archive* menu. A window will open with a list of all of the currently active schedule and archive jobs. The list displays the schedule name, task type (*SCH* for scheduled tasks, and *ARC* for archive tasks) and the next export time. This window also allows you to edit, remove, and immediately run the schedule or archive at the moment.

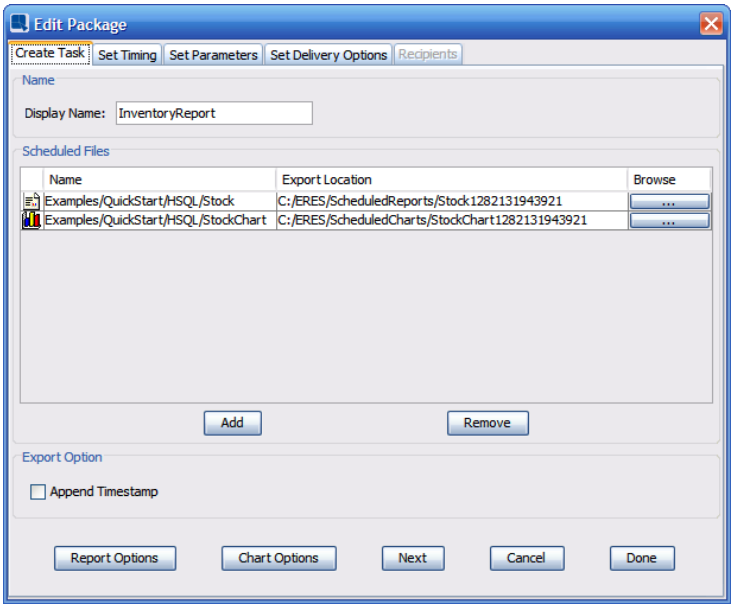
The *Schedule/Archive* menu also contains a very similar option called *View alert monitoring tasks*. This is useful when you are working with alerts monitoring. To learn more about alerts monitoring, please visit Section 9.4 - Monitoring.



View Scheduled Tasks

3.2.3.1. Editing Schedules and Archives

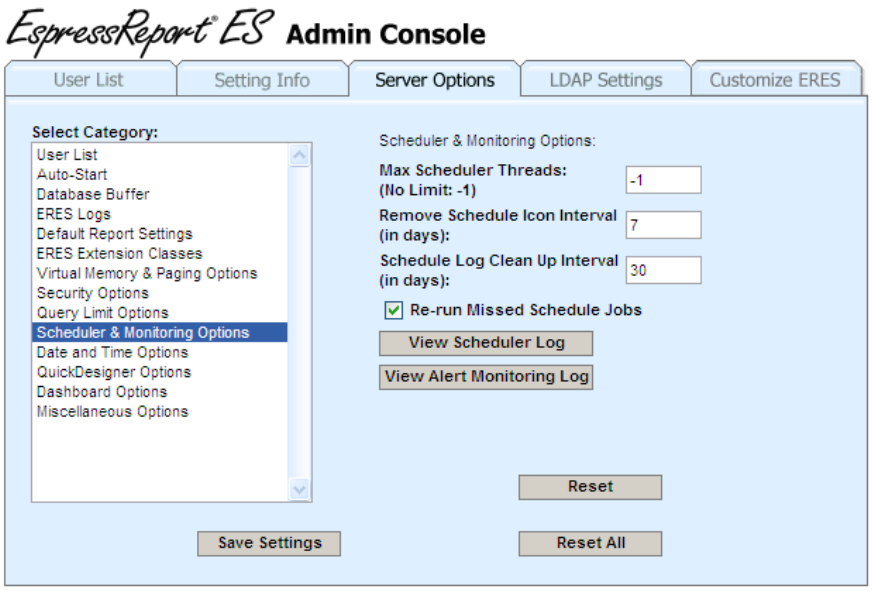
The *Edit* option in the *View Tasks* window allows you to modify the existing schedule or archive. The edit dialog looks very similar to the *Create scheduled task* dialog. However, rather than having to go through the steps in sequence, you can skip to any section by clicking on the tab names. There is also a *Done* button, which allows you to exit the dialog whenever you are satisfied with your changes.



Edit Scheduled Task Dialog

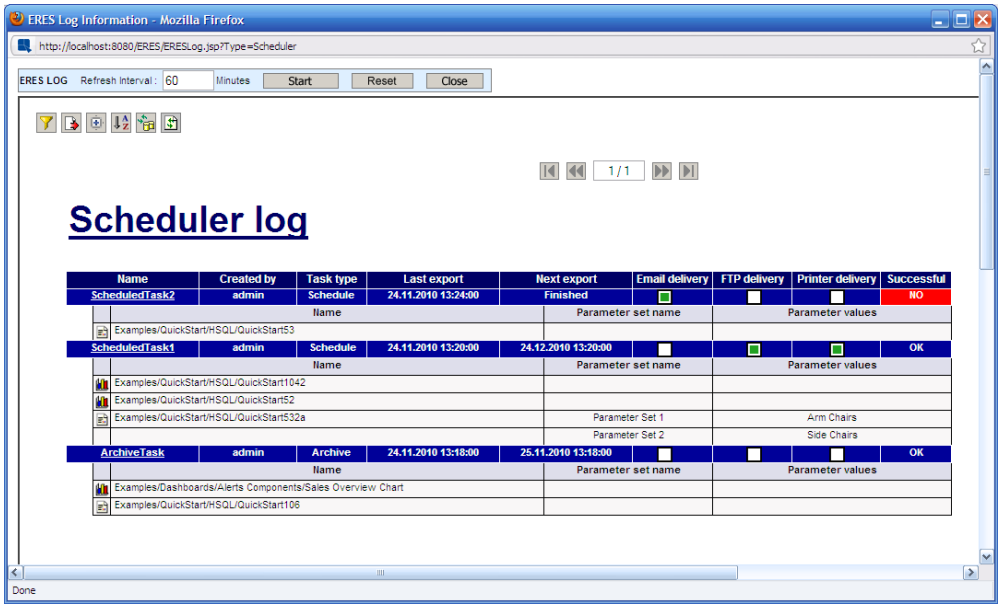
3.2.3.2. Schedule and Archive Logs

The admin also has the ability to view a log report containing detailed information about previously run schedules and archives. To view the detailed Scheduler Log, log in as the **`admin`**, and go into the administration console. In the *Server Options* tab, under the *Scheduler & Monitoring Options* section, click on the *View Scheduler Log* icon.



Admin Console

This will pop up a window containing the Scheduler Log. The root report displays information regarding the schedule list. This not only includes the currently active schedules and archives, it also includes any completed schedules and archives.



Scheduler Log

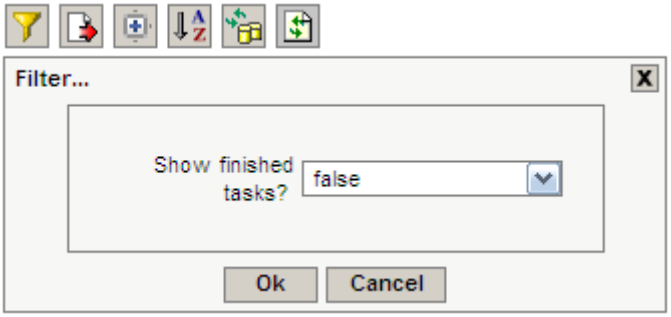


Note

The list contains only tasks that have been run at least once in the past. If there are no such tasks, *Scheduler log* will not be visible. A“No tasks found!” warning message will be displayed instead.

All tasks in the Scheduler Log have *Task type* indicator which tells whether the task is a schedule or an archive.

If you do not want the completed schedules and archives to be on the list, you can filter them out. Click on the  *Filter* icon, set the *Show finished tasks?* parameter to *false* and click *Ok*.



Scheduler Log Filter

The *successful* field indicates whether the task was run successfully during the last export time. If one of scheduled templates could not be exported or one of delivery methods could not be performed, the *successful* field will look like this:  . If that happens, you can get more detailed information in the task's drill-down. This will be discussed later.

Under each task's header, there is a list of its files. If there are several parameter sets added to a single template, the file will take several rows in the list, but its name will be displayed only once (as displayed in the following picture).

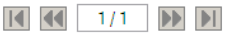
Name	Created by	Task type	Last export	Next export	Email delivery	FTP delivery	Printer delivery	Successful
ScheduledTask3	admin	Schedule	24.11.2010 13:38:00	Finished	☐	☐	☐	OK

Name	Parameter set name	Parameter values
Examples/QuickStart/HSQL/QuickStart532b	Parameter Set 1	Adad Chair
	Parameter Set 2	Hapi Table
	Parameter Set 3	Marduk Chair

Parameterized Report

The list contains one drill-down. If you click on a scheduled task name, another report will open, providing detailed information about the particular task.





21.11.2010 14:26:26

ScheduledTask3 log

Created by:	admin	
Next export time:	This task has finished	
Delivery settings:	Email delivery:	Yes
	FTP delivery	No
	Printer delivery:	No

Export time: 21.11.2010 14:20:00			
Name:	Examples/QuickStart/HSQL/QuickStart532	Export success:	OK
Template name:	C:\ERES/help/quickstart/templates/QuickStart532.rpt	Email success:	NO
Export location:	C:/ERES/ScheduledReports/QuickStart5321290345447281		
Parameter set name:			
Parameter value:			

Scheduler Log - Second Level

All tables between two consecutive blue headers show details for files that were run at the same run time.

Each scheduled file and parameter set has its own table. In the left part of the table, you will find file details, namely: *name* (also doubles as organizer path), *template name* (also doubles as disk path), *export location*, and *parameter settings*.

In the right part of the table, there are export and delivery methods results. All files have the *Export success* indicator. It indicates whether the file was successfully exported. This is the most important indicator, because if the file wasn not exported successfully, there is no file to be delivered by the delivery methods. If the export fails, be sure to check whether the template can be accessed by the ERES server (it could have been removed or renamed), if the export location can be accessed (there is enough free disk space and the ERES server has write permission to the directory) and if the file's data source can be accessed.

Each delivery method has its own indicator. If some delivery method was not enabled for the scheduled task, its indicator will not be visible in this table. If a delivery indicator is visible, you can click on it to drill-down.

Export time: 21.11.2010 14:20:00			
Name:	Examples/QuickStart/HSQL/QuickStart532	Export success:	OK
Template name:	C:\ERES/help/quickstart/templates/QuickStart532.rpt	Email success:	NO
Export location:	C:/ERES/ScheduledReports/QuickStart5321290345447281	FTP success:	OK
Parameter set name:		Print success:	OK
Parameter value:			

Export and Delivery Success

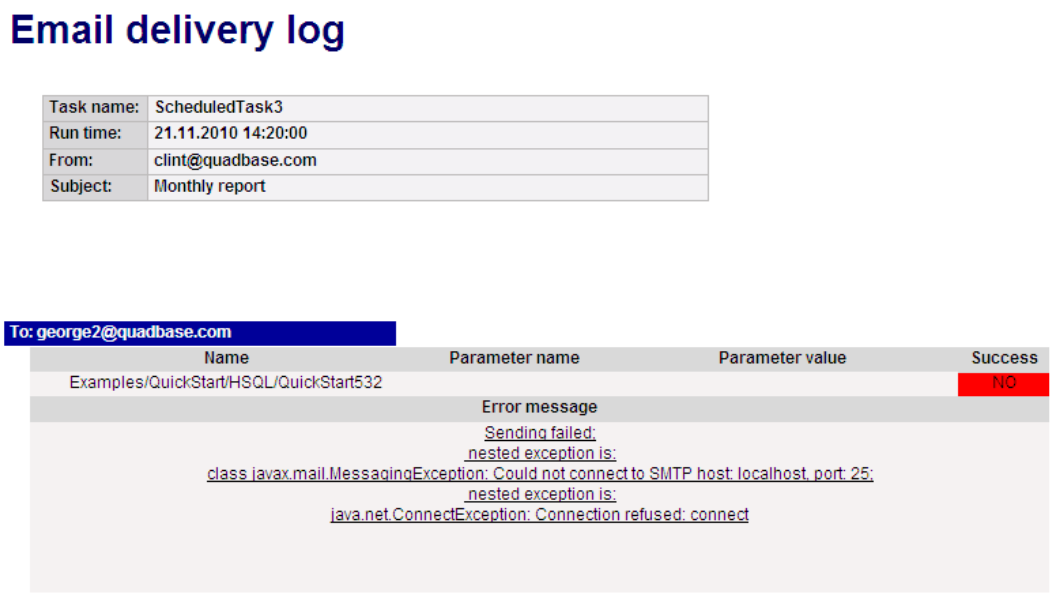
This log level can be also filtered. The log can either contain all exports or just failed exports. To enable the filter, click on the  *Filter* icon and set the *Failed exports only* parameter to *true*.

To return back to the root level, click on the  *Back* icon.

Email Delivery Log

If you click on an *Email delivery* indicator in the second Scheduler Log level, email delivery log will open. It contains details for one file and one run time only. Each email recipient has one table which shows whether the email was sent successfully to the recipient's email address. If the delivery fails, an error message will be displayed.

For example:



Failed Email Delivery Details

In the previous picture, there is a failed email delivery. The log shows that an email from `clint@quadbase.com` to `george2@quadbase.com` (not actual email addresses) couldn't be sent because the ERES server wasn't able to connect to SMTP host at localhost:25. It could be caused by a network problem or by incorrect SMTP settings. This is probably the most common problem with email delivery.

FTP delivery

FTP delivery log is much simpler – there is only one “recipient” so there is only one table. The table contains information about the scheduled file, ftp settings, run time, and *success* indicator and in case the delivery has failed – an error message. You can see such situation in following picture. The problem is that the ERES server was not able to connect to gamma.quadbase.com [https://data.quadbase.com/Docs72/ftp://gamma.quadbase.com] FTP host (not an actual FTP host).



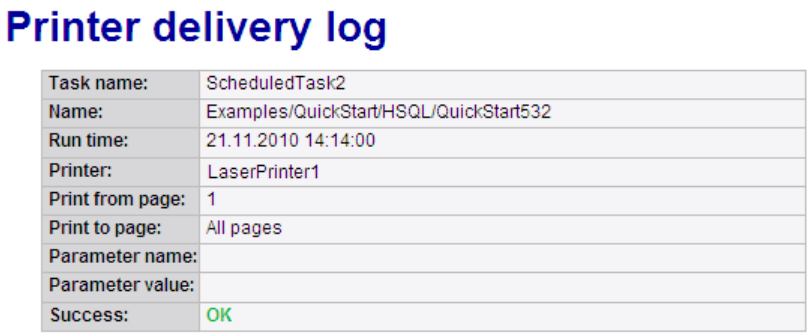
FTP Delivery log

Name:	Examples/QuickStart/HSQL/QuickStart532a
Run time:	21.11.2010 13:11:00
FTP Host:	gamma.quadbase.com
FTP Path:	/uploads
FTP User name:	jamie
Parameter name:	Parameter Set 2
Parameter values:	Single Dressers
Success:	NO
Error message	
gamma.quadbase.com	

Failed FTP Delivery

Printer Delivery Log

Printer delivery log is very similar to the FTP delivery log. It shows file and printer details and an error message (only in case the delivery has failed).



Printer delivery log

Task name:	ScheduledTask2
Name:	Examples/QuickStart/HSQL/QuickStart532
Run time:	21.11.2010 14:14:00
Printer:	LaserPrinter1
Print from page:	1
Print to page:	All pages
Parameter name:	
Parameter value:	
Success:	OK

Successful Printer Delivery

To go back to the second log level, click on the  *Back* button.

Chapter 4. Data Sources

4.1. Data in Organizer

ERES Cloud can draw data from JDBC compliant databases, text files, XML files, EJBs, and even bring in object/array data through class files. Prior to creating reports or charts, users must first set up data sources they want to use. Data source information, including database connection information, database queries, text file location, XML file/DTD/Schema location, Java class location, and EJB connection information is stored in XML registry files that are set up and defined within the Organizer interface.

4.1.1. The Data Source Manager

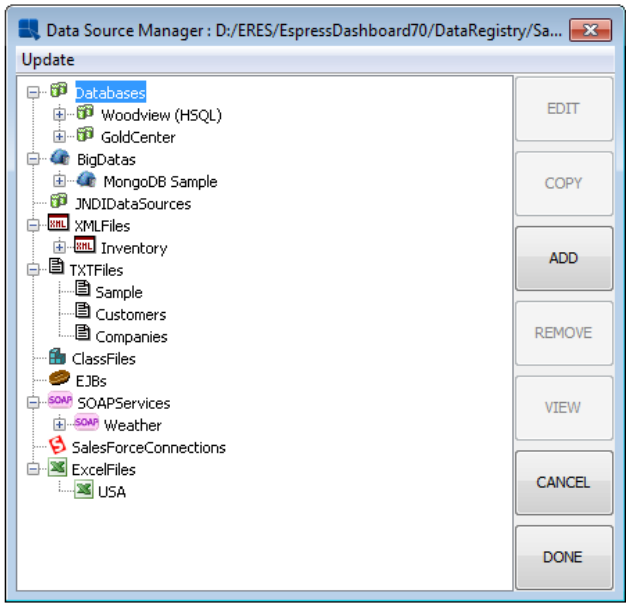
The Data Source Manager window allows you to add and edit data sources. The Data Source Manager is a visual representation of the registry XML file.



Note

The XML data source repository only stores location and connection information, not the actual data. It is not an XML file that contains data to be used in a report or chart.

4.1.1.1. Using Data Source Manager



Data Source Manager

The left side of the window contains a tree listing all of the data sources in the registry file. Grouped under *Databases* are individual databases and their associated queries and data views. Grouped under *JNDIDataSources* are database sources that use JNDI (Java Naming and Directory Interface) name to connect instead of JDBC. Grouped under *XMLFiles* are all XML files and their associated queries. Grouped under *TXTFiles* are all specified Text files. Grouped under *ClassFiles* are all specified class files, and grouped under *EJBs* are all specified EJB connections.

Right side of the window contains a series of buttons controlling all of the functions of the Data Source manager. Each button performs the following function:

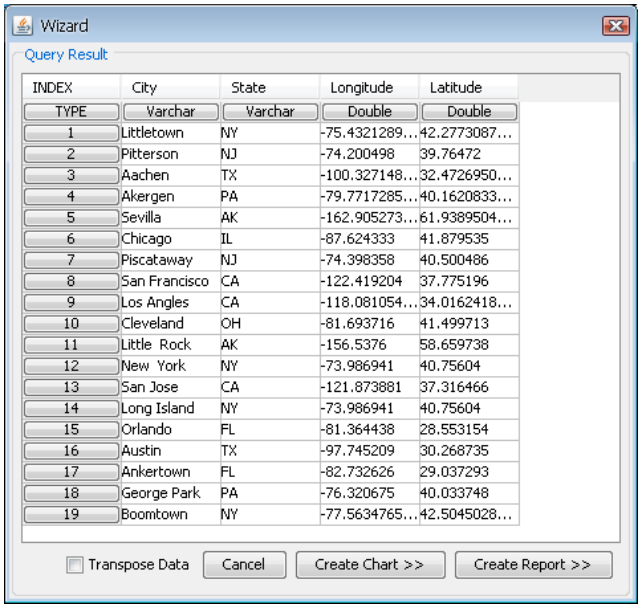
Edit: This option allows you to modify attributes of a data source. For a database, it allows you to change the connection information and modify queries/data views. For XML files, it allows you to change the file and its location, and modify XML queries. For text files, it allows you to change the display name and file location. For class files, it allows you to change the display name and modify the location. And for EJBs, it allows you to change the display name as well as parameter values.

Copy: This option is not only available for database nodes and bigdata nodes but also for the specified query or data view.

Add: This option allows you to add a data source. It will create a new source depending on which node is selected on the left side of the window. Hence, if you select *TXTFiles* and click the *Add* button, you will be prompted to add a new text file data source.

Remove: This option will remove the selected data source.

View: This will open a new window showing the tabular data from the selected data source, either query, XML file, text file, class file or EJB. For data views, the data view query interface will start. From this window, you can select to create a new report or a new chart by selecting the appropriate button.



View Data Source Dialog

Cancel: This will cancel the wizard process.

Done: This will save any changes you've made in the registry and return you to the registry list.

4.1.1.2. Data Source Node Locking

The data registry has an automatic node locking system designed to block multiple users from editing related nodes at the same time. The system utilizes your IP address, username, and a special security token to determine whether you are allowed to add, remove, or edit a specific node. If you are currently editing a node, other users will not be able to edit it. However, if you disconnect from the server or move to another computer and try to open the same node from the same username, an option box will appear allowing you to override your previous lock.

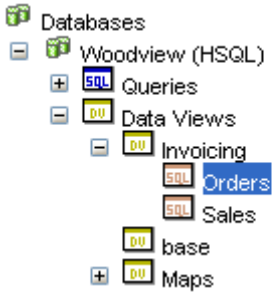
Please note that administrator is capable of unlocking any node, even if they did not open it in the first place.

The data registry is a tree where data type is the root and queries and files are the leaves. Changing any node within a particular branch will lock all parent nodes and all child nodes in that branch. Here are some examples of how the node locking system works:



Locking Query

In the above image, you can see that the database *Woodview (HSQL)* and query *Orders* are part of the same branch. So if you edit *Orders* query, it will lock both *Orders* and database *Woodview (HSQL)*. If you edit *Woodview (HSQL)*, the database will be locked and all child nodes connected to this database will be locked as well, which includes *Orders*, *Products*, *State*, and *Region*.



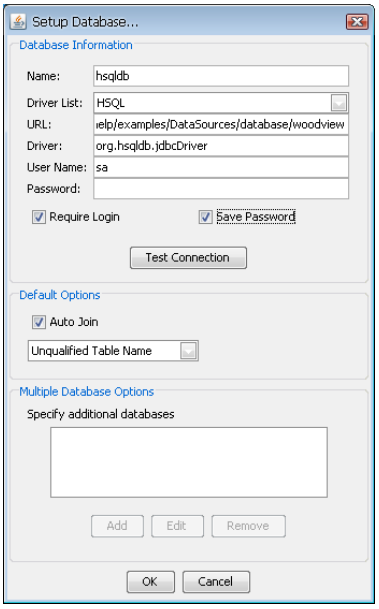
Locking Data Views

In the case of data views, there are three levels in the tree. In the above image, editing *Orders* data view query will lock the query, data view *Invoicing*, and database *Woodview (HSQL)*. If the data view *Invoicing* is being edited, both queries *Orders* and *Sales* will be locked, as well as database *Woodview (HSQL)*. However, in both of these scenarios, other data views and queries can still be edited. If the database *Woodview (HSQL)* is being edited, all data views as well as their queries will be locked.

4.1.2. Data from a Database

ERES can draw data from any JDBC compliant database. In order to connect to a database via a 3rd party driver (other than the JDBC bridge), you will need to add the classes for that driver to the classpath of the application server/servlet runner where you have deployed the ERES server. If you installed ERES with Tomcat, then you need to modify `setclasspath.bat/.sh` file in `/bin/` directory of your Tomcat installation to add the driver classes, or copy them to `<ERESInstallDir>/WEB-INF/lib` directory. Note that JDBC drivers for MS SQL Server, MySQL, Oracle, Informix and PostgreSQL databases are included as a convenience. Other database JDBC jar files are not included due to licensing, multiple drivers, and/or other concerns, although support for these databases exist and the jar files can be explicitly added.

The first step in using a database as the data source is to set up a database in the registry and specify the connection information. To add a database, click on the *Databases* node and click the *Add* button. This will bring up a window prompting you to specify the connection information for that database. You can choose a database to connect to from the Driver List or specify the information manually. Fields to enter are database name, URL, and driver. You can also select whether the database requires a login and whether you want to save username and password information. If you select to save login and password information, you can then enter these informations in *User Name* and *Password* textboxes. Then click the *Ok* button and the new database will be added to the Data Source Manager window.



Add Database Dialog

In order for ERES to create a connection to the database, the following information must be provided:

URL: This JDBC URL specifies the location of the database to be used. A standard JDBC URL has three parts, which are separated by colons:

jdbc:<subprotocol>:<subname>

The three parts of a JDBC URL are broken down as follows:

1. `jdbc` - the protocol. The protocol in a JDBC URL is always `jdbc`.

2. `<subprotocol>` - the name of the driver or the name of a database connectivity mechanism, which may be supported by one or more drivers. A prominent example of a subprotocol name is `odbc`, which has been reserved for URLs that specify ODBC data source names. For example, to access a database through a JDBC-ODBC bridge, you have to use a following URL:

   jdbc:odbc:Northwind

   In this example, the subprotocol is `odbc` and the subname `Northwind` is a local ODBC data source, i.e. `Northwind` is specified as a system DSN under ODBC.

3. `<subname>` - a way to identify the database. The subname can vary, depending on the subprotocol, and it can have a subsubname with any internal syntax the driver writer chooses. The function of a subname is to give enough information to locate the database. In the previous example, `Northwind` is enough because ODBC provides the remainder of the information.

Databases on a remote machine require additional information to be connected to. For example, if a database is to be accessed over your company Intranet, the network address should be included in the JDBC URL as part of the subname and should follow the standard URL naming convention of

//hostname:port/subsubname

Assuming you use a protocol called `VPN` for connecting to a machine on your company Intranet, the JDBC URL you need to use can look like this:

jdbc:vpn://dbserver:791/sales (similar to `jdbc:dbvendorname://machineName/SchemaName`)

It is important to remember that JDBC connects to a database driver, not the database itself. The JDBC URL that identifies the particular driver is determined by the database driver vendor. Usually, your database vendor also provides you with the appropriate drivers. It is highly recommend-

ed that users contact their database driver vendor for the correct JDBC URL that is needed to connect to the database driver.

Driver: This is the appropriate JDBC driver that is required to connect to the database. If you are using the JVM included within the installation (or Oracle's J2SE), use the following driver specification to connect to an ODBC data source:

sun.jdbc.odbc.JdbcOdbcDriver

You can also specify a JDBC driver name specific to your database if you are NOT using the JDBC-ODBC bridge. For example, the Oracle database engine will require `oracle.jdbc.driver.OracleDriver`.

User Name: This is the login used for the database.

Password: Password for the above user.

Once you specify the connection information, you can test the database connection by clicking the *Test Connection* button. This will test the connection using the information you've provided and report any problems.

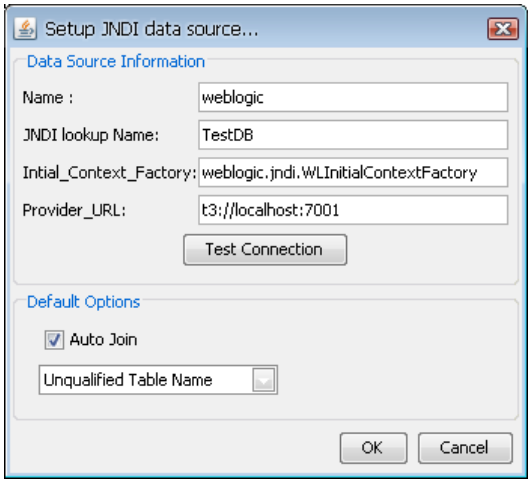
The *Default Options* portion of the dialog allows you to specify some properties for queries generated through the Query Builder interface or data views. You can specify whether to auto-join selected tables. Auto-join will attempt to join primary and foreign keys defined in the database. You can specify the table name format that should be used for queries either unqualified (only table name), or 2-part or 3-part qualified. Properties specified here will become the default setting for new queries and data views. They can also be modified for individual queries.

The *Multiple Database Options* portion of the dialog allows you to specify additional databases (i.e. additional database URLs) to obtain data from within the query. This option is only available when the database (original and any additional database) is MS SQL Server and 3-Part Qualified Table Name option is chosen. Note that the same login details as well as the same driver (as defined in the original connection) are used to connect to the specified additional databases as well. The query can obtain data by referencing a column in the additional database using a 3-Part table nomenclature.

4.1.2.1. JNDI Data Sources

In addition to connecting to databases via JDBC, ERES lets you use the JNDI (Java Naming and Directory Interface) to connect to data sources. In ERES, JNDI data sources are treated just like database data sources and support the same functionality (queries, parameters, data views, etc.). The advantage of using a JNDI data source is that it potentially makes it easier to migrate reports between environments. If data sources in both environments are set up with the same lookup name, reports can be migrated without any changes.

To connect to a JNDI data source in ERES, you must have a data source deployed in the same Web application environment as you are running the ERES server. To set up a JNDI data source, select the *JNDIDataSources* node in the Data Source Manager and click the *Add* button. This will bring up a dialog allowing you to specify the connection information.



JNDI Setup Dialog

The first option allows you to specify a display name for the data source. The second option allows you to specify the JNDI lookup name for the data source. The third option allows you to specify the initial context factory for the data source, and the last option allows you to specify the provider URL. This information will vary depending on the application server you're using as different vendors implement JNDI data sources differently. You can test the connection by clicking the *Test Connection* button.

4.1.2.2. Queries

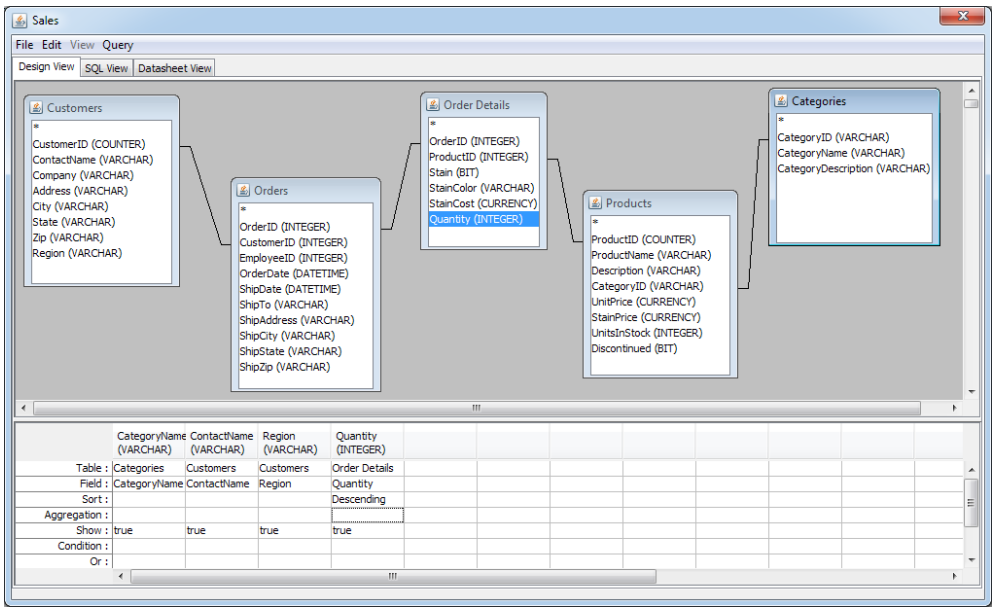
Once you add a database, a new node for your database will appear in the Data Source Manager window. When you expand the node, you will see two more nodes, one called *Queries* and one called *Data Views*. These are the two ways to retrieve data from your database. To create a new query, select the *Queries* node and click the *Add* button. A dialog will come up prompting you to specify a query name and select whether you would like to enter SQL statement or launch the Query Builder.

If you select to enter an SQL statement, a dialog box will come up allowing you to type in your SQL statement. From this dialog, you can also load a QRY or text file containing SQL text, or execute a stored procedure. If you select to launch the Query Builder, the Query Builder will open in a new window, allowing you to construct the query visually. After you finish building or entering the query, you will return to the Data Source Manager window and the query will appear as a new entry under the *Queries* node for your database.

4.1.2.2.1. Using Query Builder

The Query Builder is an integrated utility that allows you to construct queries against relational databases in a visual environment. To launch the query builder, add a new query within the Data Source Manager and select the *Open Query Builder* option. The Query Builder will then open in a new window. You can also launch the query builder to modify an existing query by double clicking the query name in the Data Source Manager.

The main Query Builder window consists of two parts. The top half of the window contains all the database tables selected for the queries and their associated columns. The top window also shows what joins have been set up between column fields. The lower half of the main window or QBE (query by example) window contains columns that have been selected or built for the query and their associated conditions.



Query Builder Window

There are three tabs at the top of the Query Builder window. These allow you to toggle between different views. The *Design View* tab is the main designer window described above, the *SQL View* tab shows the SQL statement that is generated by the current query and the *Datasheet View* tab shows the query result.

When you finish constructing the query, select *Done* from the File menu to return to the Data Source Manager.

4.1.2.2.1.1. Tables

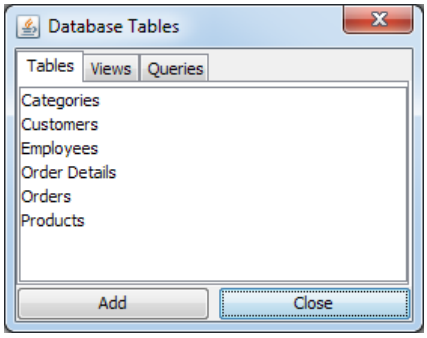
When the Query Builder launches for the first time, a tabbed window will appear containing a list of all the tables within the database. A second tab contains a list of all the views in the database, and the third tab contains a list of

other queries you have designed for the database under a heading called *Queries*. From this window, you can select the tables/views/queries from which you would like to build the query. You can also load a previously designed query as a table. To add a table, select it and click the *Add* button or double click on the table name. When a table is added, it will appear in the main Query Builder window and will display all columns within that table. To remove a table, right click within the table and select *Delete* from the pop-up menu. You can also specify a table alias and sort the fields alphabetically from this menu. You can close the tables window by clicking the *Close* button. To re-open it, select *Show Tables* from the *Query* menu.



Note

By default, the tables will appear using the name format you specified when setting up the database connection. You can change the naming by selecting *Table Name Format* from the Query menu.



Query Builder Tables Window

4.1.2.2.1.2. Joins

When you select database tables for the query, the Query Builder can auto-detect joins between column fields based on primary key-foreign key relationships in the database. Auto-joins will be added depending on which option you selected when setting up the database connection. Auto-joins will create a standard join between tables. A join is represented by a line drawn between two fields in the top half of the design window. To remove a join or edit its properties, right click on the line and select your choice from the pop-up menu. To add a join, click and drag one column field to another in a different table. A join will then appear. You can change the auto-join settings by selecting *Auto Join* from the Query menu.

Join Properties: Selecting *Join Properties* from the pop-up menu will bring up three options allowing you to select the type of join used between the column fields. Query Builder only supports equi-joins. Inequality joins can be easily accomplished using the *conditions* field. You can specify inner joins, left outer joins, and right outer joins. See the examples below for an explanation of the different join types.

Suppose you have the following two tables: Customers and Orders

CustomerID	CustomerName
1	Bob
2	Ivan
3	Sarah
4	Randy
5	Jennifer

OrderID	CustomerID	Sales
1	4	$2,224
2	3	$1,224
3	4	$3,115

OrderID	CustomerID	Sales
4	2	$1,221

An inner join on **CustomerID** on the two tables will result in combining rows from the **Customers** and **Orders** tables in such a way that each row from the **Customers** table will be "joined" with all the rows in the **Orders** table with the matching **CustomerID** value. Rows from the **Customers** table with no matching **CustomerID** fields from the **Orders** table will not be included in the query result set.

Now suppose you create a query by selecting the **OrderID**, **CustomerName**, and **Sales** fields with an inner join on the **CustomerID** field. The select statement generated by the Query Builder would look like this:

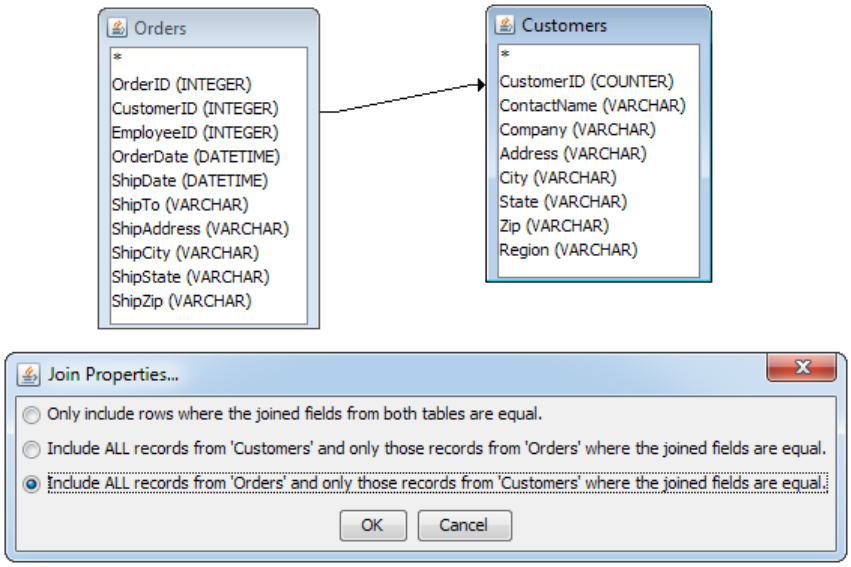
```
Select Orders.OrderID, Customers.CustomerName, Orders.Sales
From Customers, Orders
Where Customers.CustomerID = Orders.CustomerID
Order by Orders.OrderID;
```

The result of the query is shown below:

OrderID	CustomerName	Sales
1	Randy	$2,224
2	Sarah	$1,224
3	Randy	$3,115
4	Ivan	$1,221

As you can see, the **CustomerName** entries "Bob" and "Jennifer" do not appear in the result set. This is because neither customer has placed an order. There are situations where you may want to include all the records (in this example customer names) regardless whether matching records exist in the related tables(s) (in this case the **Orders** table). You can achieve this result using outer joins.

The Query Builder gives you the option of either right or left outer joins. The keywords "right" and "left" are not significant. It is determined by the order that the tables are selected in the Query Builder. If the outer table (the one that will have all records included regardless of matching join condition) is selected first, then Query Builder will use a right outer join. If the outer table is selected after the other join table, a left outer join is used. In our example, the **Customers** table has been selected before the **Orders** table, hence to select all of the records from the **CustomerName** field, the Query Builder will use a right outer join on the **CustomerID** fields.



Join Properties Dialog

Now, using the previous example, suppose you create the same query as before, except this time you specify to include all records from the **Customers** table. The select statement generated by the Query Builder would look like this:

```
Select Orders.OrderID, Customers.CustomerName, Orders.Sales
From Orders right outer join Customers on Orders.CustomerID =
 Customers.CustomerID
Order by Orders.OrderID;
```

The result of the new query is shown below:

OrderID	CustomerName	Sales
	Jennifer	
	Bob	
1	Randy	$2,224
2	Sarah	$1,224
3	Randy	$3,115
4	Ivan	$1,221

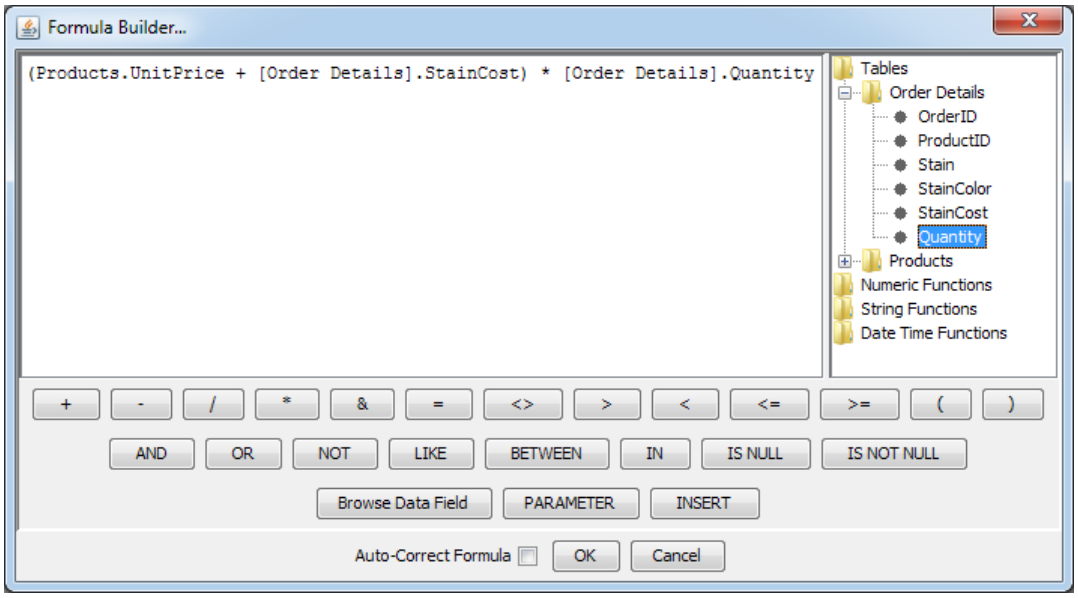
As you can see, all of the customer names have now been selected and null values have been inserted into the result set where there are no corresponding records. If you specify an outer join, the join line connecting the two tables in the Query Builder will become an arrow in the direction of the join.

4.1.2.2.1.3. Columns

The QBE window contains information on column fields selected for the query, as well as any conditions for the selection.

Selecting Column Fields: You can add column fields to the query from any table that has been selected in one of two ways. You can double-click on a field name within a table to add it to the query, or you can double-click on the *Table* or *Field* fields to bring up a drop-down menu with field choices. You can remove a column from the query by right clicking in the lower window and selecting *Delete Column* from the pop-up menu, or by selecting *Delete Column* from the Edit menu. Once you select a column field, you can specify how you want to sort the column, in ascending or descending order, by double clicking on the *Sort* field. You can also specify group by or column aggregation by double clicking on the *Aggregation* field. Aggregation options include: *Group By, Sum, Average, Min, Max, Count, Standard Deviation, Variance, First, Last,* and *Where*. If you select group by for one column, then you are required to specify group by (or aggregation) for all of the other columns. To specify a column alias, right click on the column and select *Alias* from the pop-up menu. You can perform a `SELECT DISTINCT` operation by selecting the *Select Distinct* option from the query menu.

Building Columns: To build your own column, right click on a blank column in the QBE window and select *Build* from the pop-up menu. This will launch the Formula Builder. The Formula Builder allows you to construct columns in a visual environment using the tables that you have selected and the formula library for the database that you are using. You can click the *Browse Data Field* button to see the first few records of data for any field in your query.



Formula Builder Window

Conditions: You can place conditions on the query selection by entering them in the *Condition* or *Or* fields. A condition placed in the *Condition* field creates an `AND` clause within the generated SQL. A condition placed in the *Or* field creates an `OR` clause within the SQL. Right clicking in either field and selecting *Build* from the pop-up menu will bring up the Formula Builder. In the Formula Builder, you can specify standard conditions, `=, <, >, BETWEEN, LIKE, NOT,` etc., as well as construct formulas to filter the query. You can also specify a query parameter here.



Note

ERES can auto-correct items entered as query conditions by appropriately appending the field name and encasing string arguments in quotes. For example, if you enter `= ARC`, ERES will change the query condition to `Categories.CategoryName='ARC'`. If you're using complex functions (i.e. database functions that take multiple string arguments), ERES may not be able to properly parse the function. You can turn off the auto-correct feature by unchecking the box at the bottom of the formula builder window.

4.1.2.2.1.4. Using Database Functions

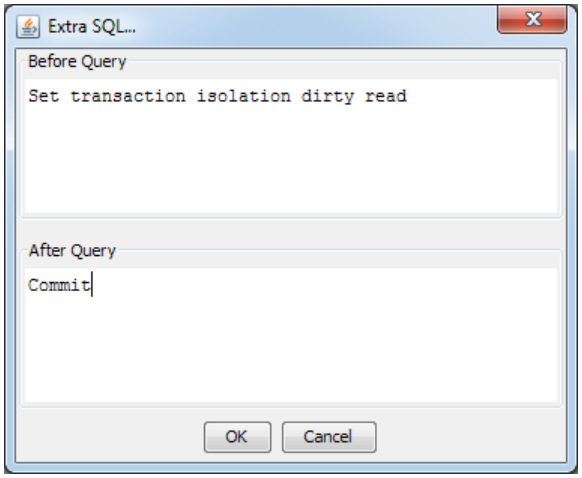
The formula builder component in the query builder allows you to use database specific functions when building a column or condition for the query. You can use the functions that are supplied or add your own to the interface.

ERES comes with the function libraries for Oracle, Access, MS SQL, and DB2 pre-loaded. They are stored in XML format in `DatabaseFunctions.xml` file in `userdb` directory. For databases with functions not stored in XML, ERES will use default ones. You can specify different database functions by editing the XML file or creating a new one based on `DatabaseFunctions.dtd` file in `userdb` directory. A sample database functions file might look like this:

```
<DatabaseFunctions>
      <Database ProductName="ACCESS">
            <FunctionSet Name="Numeric Functions">
                  <Function>Abs(number)</Function>
                  <Function>Atn(number)</Function>
            </FunctionSet>
      </Database>
</DatabaseFunctions>
```

4.1.2.2.1.5. Adding Extra SQL

Sometimes it is necessary to add extra SQL statements to run before or after a query. For example, you may need to set a transaction level or call a stored procedure before executing a query and/or commit a transaction or drop a temporary table after executing a query. The query builder allows you to specify these extra SQL statements by selecting *Extra SQL* from the Query menu. This will bring up a window allowing you to write statements to execute before and/or after a query.



Extra SQL Dialog

You can enter any SQL statements you would like to run before and/or after the query in the appropriate boxes. Once you finish, click the *Ok* button and the statements will be added to the query.

4.1.2.2.1.6. Query Output

The *SQL View* and *Datasheet View* tabs allow you to see two different views of the query.

SQL View: The *SQL View* tab shows you the SQL statement generated by the query in the design view. It allows you to see how the Query Builder is translating different operations into SQL. You can edit the generated SQL if you want, however, if you change the SQL and then return to the *Design View*, all changes you made will be lost. If you save a query after changing the SQL, the query will re-open in the *SQL View* tab if you select to edit it.

Datasheet View: The *Datasheet View* tab shows you the query result in data table form (this tab is also available in the Enter SQL dialog). The datasheet view will show you all the data that is drawn as a result of executing the query. Going to the datasheet view will also test the query to check for design errors. You can navigate the query result by using the toolbar at the bottom of the window.

 Go to the first page of the data table

 Go to the previous page of the data table

 Go to a specific row of data

 Go to the next page of the data table


Go to the last page of the data table


Set number of rows to display per page (default is 30)

Exporting Queries: You can export queries in one of two ways. You can output the SQL statement as a text or you can output the query result as CSV file. To export a query, select *Export* from the File menu. A second menu will appear giving you the option to *Generate SQL* or *Generate CSV*. Select the desired option and a dialog box will appear prompting you to specify the filename and location.



Note

To save the query and exit the Query Builder, select *Done* from the File menu.

4.1.2.2.2. Parameterized Queries

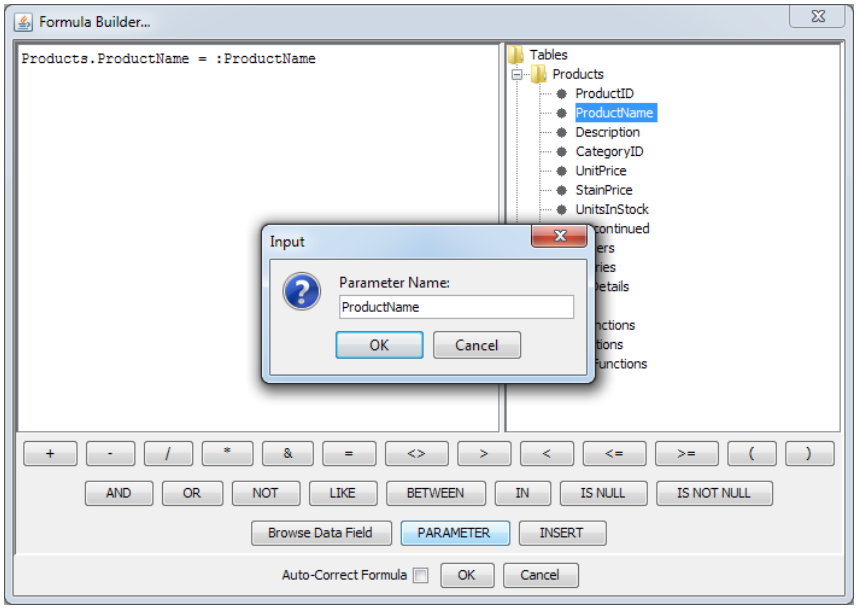
You can also use the Query Builder to design parameterized queries. This feature allows you to filter the data at run-time.

Query parameters can be defined when typing an SQL statement or using the Query Builder. They can also be defined when running data views (this is covered in the next section). A parameter is specified within an SQL statement by the `":"` character. Generally, the parameter is placed in the `WHERE` clause of an SQL Select statement. For example, the following SQL statement

```
Select * From Products Where ProductName = :Name
```

specifies a parameter called `Name`. You can then enter a product name at run-time and only retrieve data for that product.

Within the Query Builder, you can specify a query parameter by right clicking on the *Condition* field and selecting *Build* from the pop-up menu. The Formula Builder will open, allowing you to place a condition on the column.



Specifying a Parameter in the Formula Builder

You can insert a parameter by clicking the *PARAMETER* button. A second dialog will appear prompting you to specify a name for the parameter. Type the parameter name, click *OK* and then click *OK* again to close the formula builder. You can specify as many different parameters for query as you like.

4.1.2.2.2.1. Multi-Value Parameters

ERES supports a special kind of parameter that takes an array of values as input rather than a single value. Multi-value parameters are useful when you want to filter the result set based on an unknown number of values. For example, say a report is run to return a list of customers for a specific state/province. Users could select as many different states/provinces as they wanted and return the relevant information.

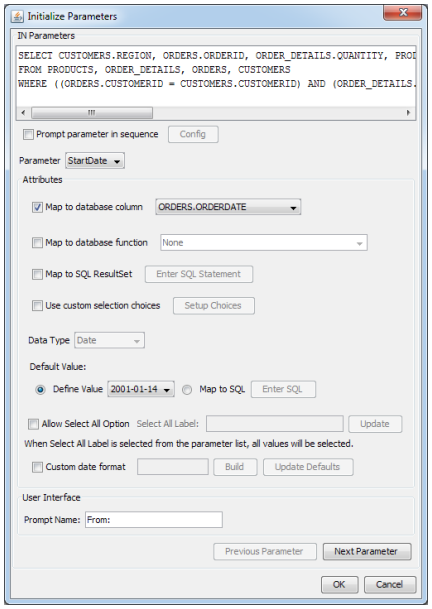
To create a multi-value parameter, place a parameter within an `IN` clause in an SQL statement. For example, the following query

```
Select Customers.Company, Customers.Address, Customers.City,
 Customers.State, Customers.Zip
From Customers
Where Customers.State IN (:State);
```

will create a multi-value parameter named `State`. Multi-value parameters will only be created if there is only one parameter in the `IN` clause. If you place more than one parameter in the `IN` clause, i.e. `Customers.State IN (:State1, :State2, :State3)`, it will create three single value parameters instead.

4.1.2.2.2.2. Initializing Query Parameters

When you attempt to save (by selecting *Done* from the File menu) or preview (by clicking the *Datasheet View* tab) a parameterized query, you will first be prompted to initialize the parameter. You can also initialize it by selecting *Initialize Parameters...* from the Query menu or by clicking the *Initialize Parameters* button in the Enter SQL Dialog.



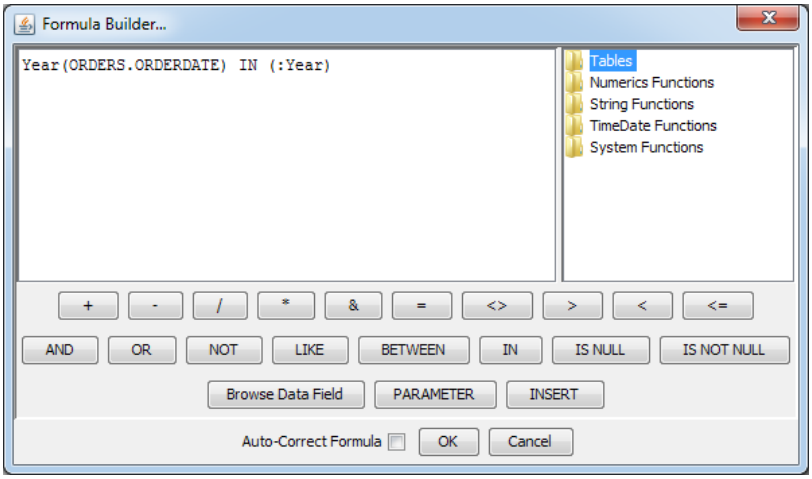
Initialize Parameter Dialog

From this dialog you can specify the following options:

Map to database column: This option allows you to specify a column from the database whose values will be used for the parameter input. Selecting this option modifies the parameter prompt that the end user will get when previewing or running the report in the Report Viewer. If you map the parameter to a database column, then the user will be prompted with a drop-down list of distinct values from which to select a parameter value. If you do not map, the user will have to type in a specific parameter value.

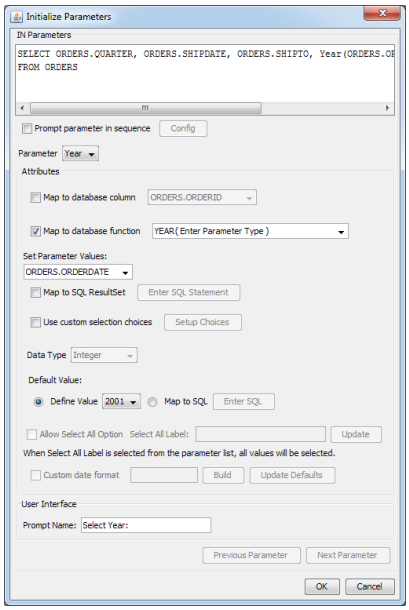
Map to database function:

The *map to database column* feature is very handy if you want to enter a valid value for a parameter from a list box, but sometimes you rather want a computed value or a derived value from a database column. For example, you want to find all orders from year 2007. However, OrderDate is a date. What you want is to apply the *Year* function to the OrderDate column. This is the impetus behind this feature. Mapping a parameter to a database function is very similar to mapping to a column. In the formula builder, enter a condition comparing a function result to a parameter as shown below:



Condition for Mapping to Database Function

In the initialize parameter dialog, check the *Map to database* function checkbox and the values will be automatically filled in.



Map Parameter to Database Function

The list of custom functions is extracted from `DatabaseFunctions.xml` file located in `<ERES Install>/userdb/` directory. Modify the `.xml` file if you wish to add a new database or custom functions. The new functions will appear in this list after you restart the program.

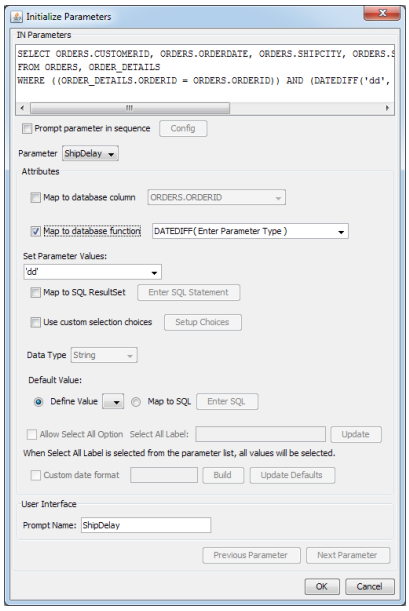
If your database is not listed in the `.xml` file, the function list will be populated with functions listed in the JDBC driver. However, the function para-

meters are not provided. For example, the HSQL database will not be listed in the `.xml` file.

An interesting example using the HSQL database is as follows. Suppose you would like to create a report for orders that were delayed. You can utilize the HSQL DateDiff function to find the number of days for the order to ship.

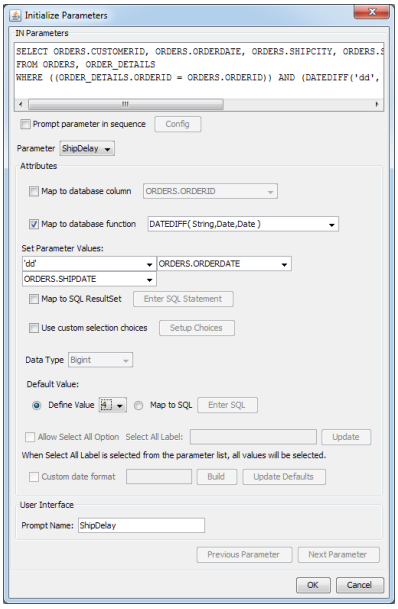
```
DATEDIFF('dd', ORDERS.ORDERDATE, ORDERS.SHIPDATE)
 >= :ShipDelay
```

This function finds the difference between the order date and the ship date and displays the result in terms of days. If you initialize the parameter and check the map to database function, the following window will be displayed:



No Parameter Types for HSQL Function

The `DateDiff` function requires a string and two date values for the parameters. Enter these parameter types in the parentheses. This will bring up three set parameter value lists. Enter **dd** (day) for the first parameter, select `Orders.OrderDate` from the list for the second parameter, and select `Orders.ShipDate` from the list for the third parameter. The default values will be updated with the function results.



Map Parameter to HSQL Function

Map to SQL ResultSet:

A parameter mapped to a database column will give you a list of distinct values in a drop-down list box for the user to choose from when running the report. However, to produce the list of values, a select distinct on the column with the joins and conditions from the query will be run. In some cases, this can be a time-consuming process. To obviate this problem, and in fact gain complete control as to what and how to populate the drop-down list box, you can write your own select statement to populate the drop-down list. An added bonus is that parameters that are in the query can be included in this query. With proper joins and parameters included, you can use this feature to facilitate cascading parameters (See Section 4.1.2.2.2.3 - Cascading Parameters). Below is an example:

Suppose you have two parameters in the query. So, your query is as follows:

```
SELECT CATEGORIES.CATEGORYID,
 PRODUCTS.PRODUCTNAME, PRODUCTS.UNITPRICE,
 PRODUCTS.UNITSINSTOCK
FROM PRODUCTS, CATEGORIES
WHERE ((PRODUCTS.CATEGORYID =
 CATEGORIES.CATEGORYID))
AND (((CATEGORIES.CATEGORYID =:category) AND
 (PRODUCTS.PRODUCTNAME =:product)))
```

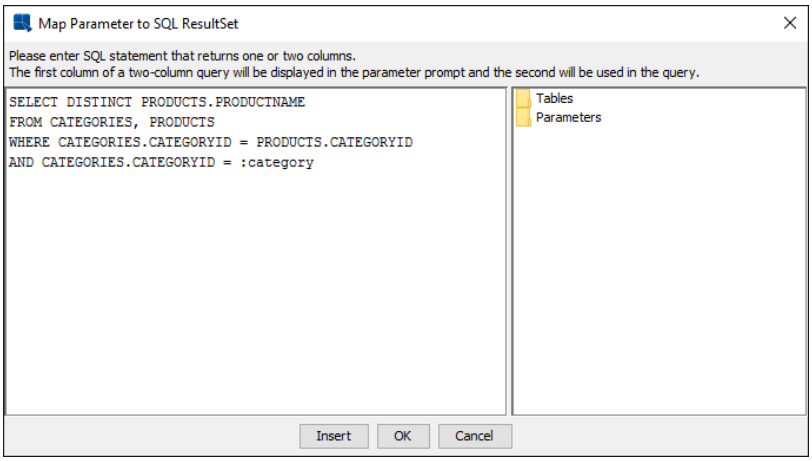
In the `config` prompt in `initialize parameter`, set the order for parameter prompting to `category` first, then `product`.

The select statement for parameter `category` can be the following:

```
SELECT DISTINCT CATEGORIES.CATEGORYID
FROM CATEGORIES
```

The select statement for parameter `product` will be as shown below:

```
SELECT DISTINCT PRODUCTS.PRODUCTNAME
FROM CATEGORIES, PRODUCTS
WHERE CATEGORIES.CATEGORYID = PRODUCTS.CATEGORYID
AND CATEGORIES.CATEGORYID = :category
```

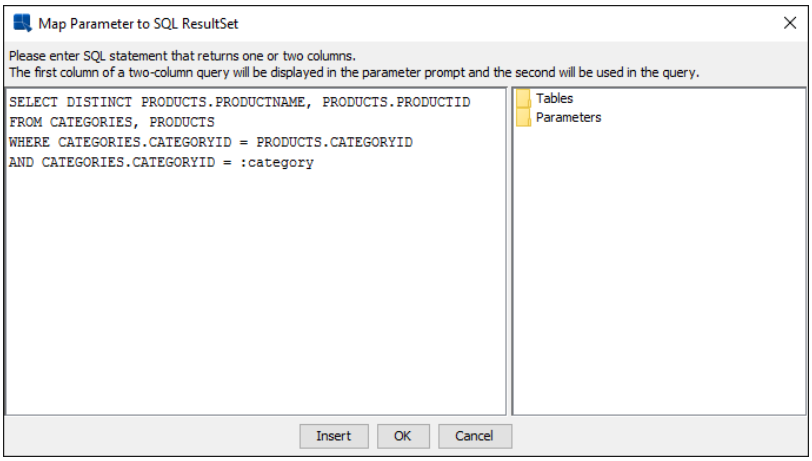


Select Statement for Product

When the user runs the template, `category` will be prompted first. Then the value of `category` chosen will be used to filter for `product`.

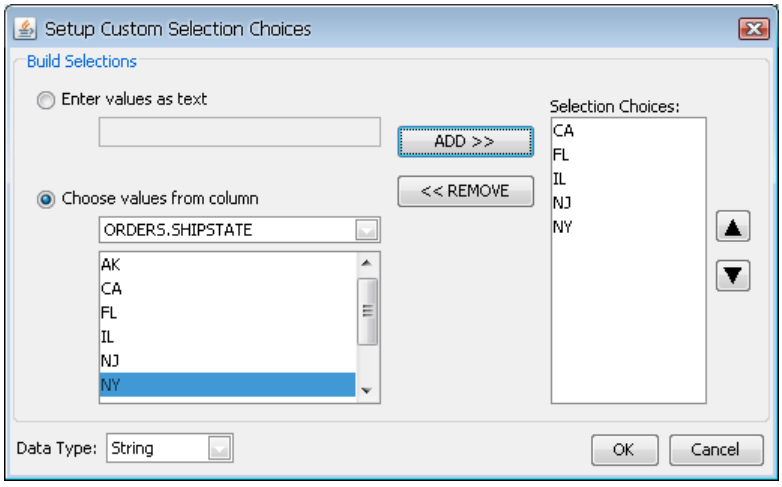
The select statement mapped to a parameter can have either one or two columns in the select list. It is clear that if one column is in the select list, it must be the column that supplies list of distinct values for the parameter. Another useful feature provided here is that you can actually select two columns in the select list such that the first one of the columns will supply values for the drop-down list while the second column will be the actual parameter value for the filter condition. Consider the following example.

Suppose your database has a table with product ID as the primary key. When your end user wants to search for products from the database, they want to use the *product name* as parameter (therefore it is listed in the query as first) since a *product ID* could be just a cryptic code (therefore it is listed in the query as second). Using this feature, you can choose product name for the values in the drop-down list and product ID as the actual value filter condition.



Select Statement with Two Columns

Use custom selection choices: Rather than having a drop-down menu with all the distinct column values, you can also create a custom list of parameter values. To create this list, select this option and click the *Setup Choices* button. This will launch a new dialog allowing you to create a list of choices.



Custom Parameter List Dialog

In this dialog, you can either enter custom values or select values from the distinct values of a column in the database. Once you finish specifying the values for the list, click the *OK* button and the choices will be saved.

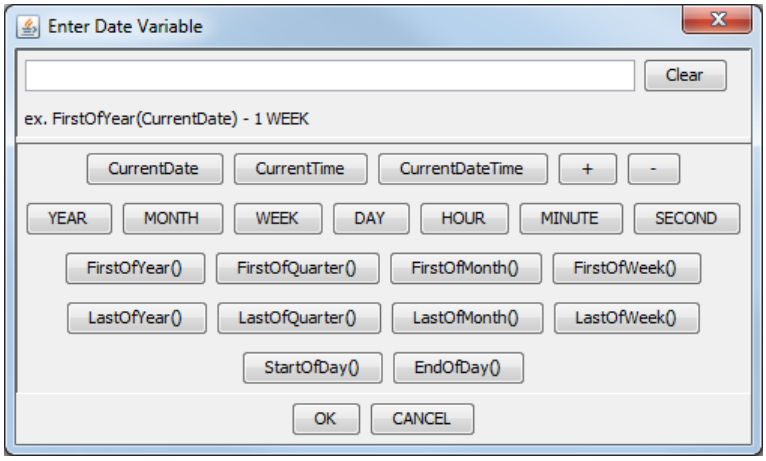
Default Value: This allows you to specify a default value for the parameter. Although you don't have to specify a default value, it is recommended that you do so. If you do not supply a default value, you cannot open or manipulate the report template without the data source present.

You can either select a single value manually (either choose it from a list or type it manually, it depends on the mapping method you chose) or map the default value to a SQL query.

For multi-value parameters (see Section 4.1.2.2.2.1 - Multi-Value Parameters), the SQL query can return more than one value. In such case, several values will be chosen as default parameter values.

Date Variable: This option is only available when the parameter is not mapped to a database column or function, or it's mapped to a SQL resultset and not set to a custom

selection choice. This option is only intended for parameters with variable type date/time. When you click this button, the following panel will pop up, listing all the supported keywords.



Enter Date Variable Dialog

This dialog allows you to select one of the three keywords: `CurrentDate`, `CurrentTime`, and `CurrentDateTime`. You can add or subtract units of time from the current date/time, allowing you to have a dynamic date range. For example, a report can have the following default values:

```
StartDate: CurrentDate - 1 WEEK
EndDate: CurrentDate
```

This would indicate that every time the report is run, the default prompt should be one week ago to the current date. Other supported time units are `YEAR`, `MONTH`, `DAY`, `HOUR`, `MINUTE`, and `SECOND`. This feature only supports a single addition or subtraction, it does not support multi-value parameters.

You can also use functions to define the parameter value:

FirstOfYear() Argument format: `CurrentDate`, `CurrentDateTime`, e.g. `FirstOfYear(CurrentDate)`

This function returns a date of the first day of the year from the argument. For example, when the argument evaluates to `2012-08-14`, the function returns `2012-01-01`.

LastOfYear() Argument format: `CurrentDate`, `CurrentDateTime`, e.g. `LastOfYear(CurrentDate)`

This function returns a date of the last day of the year from the argument. For example, when the argument evaluates to `2012-08-14`, the function returns `2012-12-31`.

FirstOfQuarter() Argument format: `CurrentDate`, `CurrentDateTime`, e.g. `FirstOfQuarter(CurrentDate)`

	This function returns a date of the first day of the quarter which includes the date from the argument. For example, when the argument evaluates to `2012-08-14`, the function returns `2012-07-01`.
LastOfQuarter()	Argument format: `CurrentDate`, `CurrentDateTime`, e.g. `LastOfQuarter(CurrentDate)` This function returns a date of the last day of the quarter which includes the date from the argument. For example, when the argument evaluates to `2012-08-14`, the function returns `2012-09-30`.
FirstOfMonth()	Argument format: `CurrentDate`, `CurrentDateTime`, e.g. `FirstOfMonth(CurrentDate)` This function returns a date of the first day of the month from the argument. For example, when the argument evaluates to `2012-08-14`, the function returns `2012-08-01`.
LastOfMonth()	Argument format: `CurrentDate`, `CurrentDateTime`, e.g. `LastOfMonth(CurrentDate)` This function returns a date of the last day of the month from the argument. For example, when the argument evaluates to `2012-08-14`, the function returns `2012-08-31`.
FirstOfWeek()	Argument format: `CurrentDate`, `CurrentDateTime`, e.g. `FirstOfWeek(CurrentDate)` This function returns a date of the first day of the week which includes the date from the argument. For example, when the argument evaluates to `2012-08-14`, the function returns `2012-08-12` (Sunday is considered as the beginning of the week).
LastOfWeek()	Argument format: `CurrentDate`, `CurrentDateTime`, e.g. `LastOfWeek(CurrentDate)` This function returns a date of the last day of the week which includes the date from the argument. For example, when the argument evaluates to `2012-08-14`, the function returns `2012-08-18` (Saturday is considered as the end of the week).
StartOfDay()	Argument format: `CurrentTime`, `CurrentDateTime`, e.g. `StartOfDay(CurrentDateTime)` This function returns a time of the start of the day from the argument. For example, when the argument evaluates to `2012-08-14`

`12:15:03`, the function returns `2012-08-14 00:00:00.0`.

EndOfDay() Argument format: `CurrentTime`, `CurrentDateTime`, e.g. `EndOfDay(CurrentDateTime)`

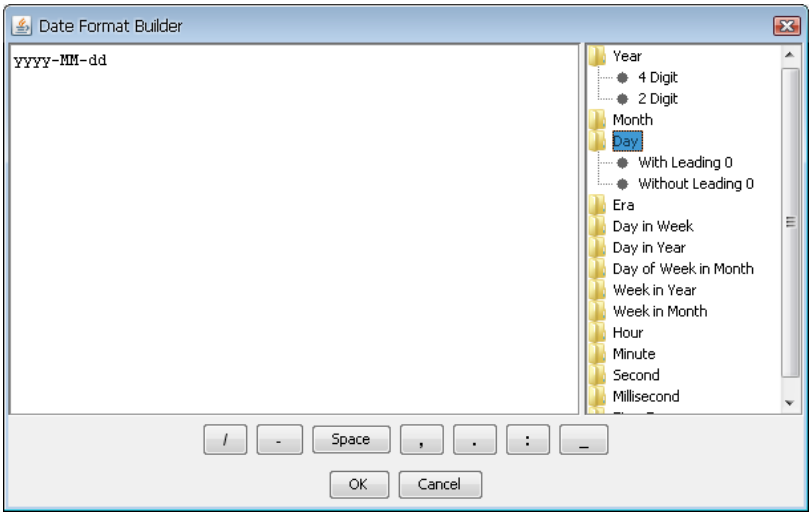
This function returns a time of the end of the day from the argument. For example, when the argument evaluates to `2012-08-14 12:15:03`, the function returns `2012-08-14 23:59:59:999`.

Data Type: This allows you to specify data type for the parameter value(s). If you mapped the parameter to a column, the data type is set automatically.

Allow Select All Option: Use this to add an option to the parameter prompt dialog that will allow users to select all parameter values even for single-value parameters. See the Section 4.1.2.2.2.4 - All Parameters for more details.

Custom Date Format: This allows you to set the format in which the date parameter should be entered. This option is only available when the data type is either date, time, or timestamp and the parameter is either mapped to column or not mapped at all.

When you check this option, the default custom date format is shown. The date format is built using a combination of characters that represent date/time elements. You can build the format easily using the date format builder by clicking the *Build* button.



Date/Time Format Builder

The builder contains a list of elements available on the right. You can mouse over the elements to see an example of each presentation. The bottom section contains a set of separators available for use.

You can also type in the format by hand either in the builder or directly in the Custom Date Format input box. For more information about characters and their formatting, please refer to the documentation for the **printDate()** report function in Section 5.1.6.2.8.3 - Date/Time Functions. Formatting for this option is the same as for the format argument of the function.

Prompt Name: This allows you to specify the prompt that is given to the user in the parameter dialog.

If you map the parameter, the user will see either a drop down box (single value parameter) or a list box (multi-value parameter) containing various options. If you choose not to map the parameter, the user will see a textbox to enter their own value. In case of a multi-value parameter, it is recommended to let the user inform in the parameter prompt that this parameter accepts multiple values. Users can separate multiple values using a comma (e.g. `ARC, DOD, TRD`). If the text requires the use of comma, the user can use quotes to include the comma within the filter string (e.g. `"Doe, John", "Smith, Mike"`).

Clicking the *Previous Parameter* and the *Next Parameter* buttons allows you to initialize each of the parameters that have been defined in the query.

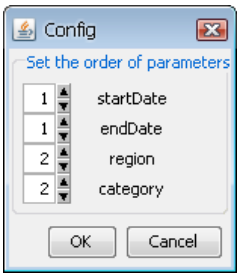
When you select to use a parameterized query to design a report, or open a report that uses a parameterized query, the report will load/start with the default values. You will be prompted to provide parameter values when you preview the report.

4.1.2.2.2.3. Cascading Parameters

By default, the user is prompted to enter all report parameters at once in the prompt dialog. This configuration, however, may not be the best approach if some parameters are mapped to database columns with a significant number of distinct values. It can be difficult to select from a very large list and depending on the parameter combination, users may be able to select parameters that do not return any data.

To assist with these problems, ERES provides a feature that allows the user to configure the order in which the parameters should be entered. With this feature enabled, the user enters parameters in the dialog in a pre-defined order. As such, each selection will be applied as a filter to the next parameter prompt(s). Using cascading parameters can limit the number of distinct values presented to the user and can prevent the user from selecting invalid parameter combinations.

To enable cascading parameters, check the option marked *Prompt parameter in sequence* in the parameter initialization dialog. Then click the *Config* button to set the order of the parameter prompts. A dialog will open showing all the parameters defined in the query.



Parameter Sequence Dialog

Using the spin boxes, you can set the sequence for the query parameters. User will be prompted to start with the lowest numbered parameter and work his/her way up to the highest one. If two or more parameters share the same number, the user will be prompted to enter those parameters at the same time (in the same dialog).

By default, the parameter values for the next level are generated by rerunning the entire query with the previously prompted parameters filled in. If the original query is slow to execute, you can improve performance by mapping higher order parameters to SQL Queries. You can even include previously selected parameter values in the mapped query. For more information, please see Section 4.1.2.2.2.2 - Initializing Query Parameters.

Please note that this feature can only be used with reports.

4.1.2.2.2.4. All Parameters

Sometimes you want to select all parameter values at once. The *All Parameters* feature allows you to do so.

Single-value parameters

It is possible to select all parameter values at once, even for parameters that do not allow multi-value selection.



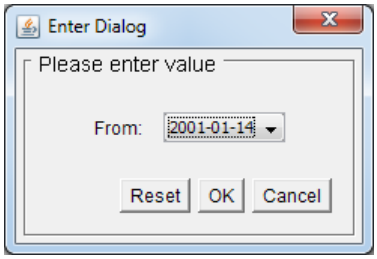
Note

There is a difference between multi-value selection and all-value selection. See the Inner Workings chapter to learn more.

For example: Let's assume you have a condition like this:

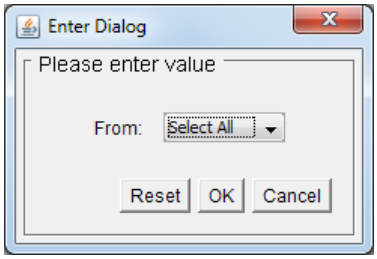
```
WHERE column = :Parameter
```

In such case, the parameter prompt dialog will not allow you to select more than one value.



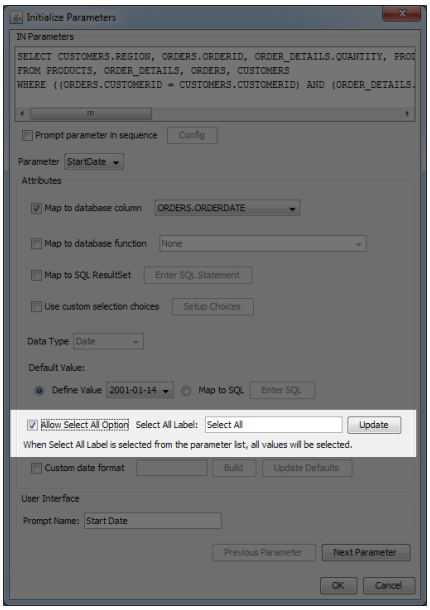
Typical single-value parameter

But you can use the *Select All Values* feature to add an option to the parameter value list that will allow viewers to select all parameter values at once (even if the parameter does not allow multi-value selection).



Single-value parameter with the Select All functions enabled

The *Select all* feature can be enabled in the *Initialize Parameters* dialog (see Section 4.1.2.2.2.2 - Initializing Query Parameters for more details) by selecting the *Allow Select All Option* checkbox.



This option is only available for parameters that meet the following requirements:

1. The parameter uses one of the following operators. If there are multiple occurrences of this parameter in the query, all parameter comparison operators have to be one of these:

   < less than

   <= less than or equal to

   > greater than

   >= greater than or equal to

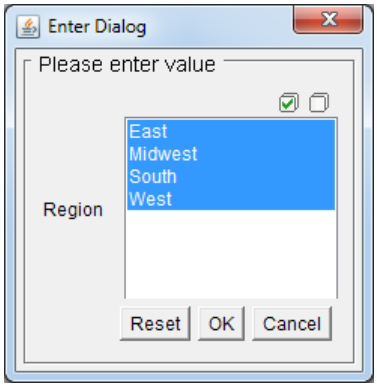
   = equal to

2. The parameter is mapped to the same column as the column from the parameter condition **or** the parameter isn't mapped to anything.

After you select the *Allow Select All* option, the *Select All Label* field activates, allowing you to enter a text that will be used for selecting all data from the query. For parameters that are mapped to a column, this text will be displayed in the parameter value list in 1st place. For parameters that aren't mapped to anything, entering this text as the parameter value will result in selecting all data.

Multi-value parameters

Unlike single-value parameters, the *Select All* feature is enabled for all multi-value parameters by default (in fact, it can't be disabled, because disabling it for multi-value parameters would make no sense). All you have to do to use this feature is to click on the *Select All* icon in the parameter prompt.



However, multi-value parameters can work in two modes:

1. If the parameter meets the conditions from the previous paragraph and the parameter is on the first cascading level (i.e. parameter cascading is disabled or the parameter is on the first cascading level), it is parsed by the SQL parser and the parameter condition is nullified. Nullifying the parameter optimizes the query and prevents it from causing performace issues or even errors. See the Inner Workings to learn more about how it works.

2. If the parameter doesn't meet the conditions from the previous paragraph or if it's not on the first cascading level, selected values will be injected to the query as a list of values separated by comma. If there is a large amount of values injected to the query as a list, the query can become quite long. Long queries can cause performance issues or even errors, so it is **not recommended** to use this option for parameters with many values.

Inner Workings

If a report/chart/map viewer chooses to select all parameter values for a single-value parameter or for a multi-value parameter that meets the conditions for parameter disabling, the query is then automatically parsed and a special condition is added to the parameter which basically disables the parameter.

For example: The following query

```
select *
from table
where column > parameter_value
```

Would be parsed and passed to the database as:

```
select *
from table
where ((column > parameter_value) OR (1 = 1))
```

This example also demonstrates another important thing: selecting all values for the < (less than) or > (greater than) operators returns all values from the table (if there are no other conditions) rather than returning no data at all (because condition like `WHERE <all data from the Date column> > Date` would return no data...).

Because ERES allows you to use many database systems, parsing may fail for certain complex queries in certain databases. In such case a warning dialog will be displayed.

In such situations, you have the following three options:

1. Try to modify the query so it can be parsed by our parser.
2. Add your own *Select all parameters* condition to the query.

   For example:

   ```
   WHERE ((column = :Parameter) OR (:Parameter
    LIKE 'selectall'))
   ```

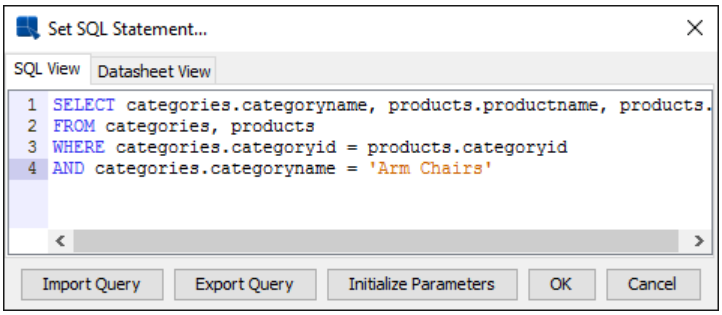
   Note

   If you embed all parameters directly to the query, leave the *Allow Select All Option* option **disabled**.

3. Contact Quadbase support [https://www.quadbase.com/support-overview/].

4.1.2.2.3. Entering SQL Statements

Typically, the Query Builder is recommended for creating queries. However, there are cases when it is necessary to enter SQL statements directly: for example if the query is already created in a QRY file, if the query is built into a stored procedure/function, or if the query requires commands not supported by the Query Builder. In these situations, select *Enter SQL statement* to open the Set SQL Statement window. Here you can enter SQL statements directly into the text area as shown below or you can load an existing QRY File.



Enter SQL Statement Dialog

To preview the result set, click on the *Datasheet View* tab.

4.1.2.2.3.1. Calling Oracle Stored Procedures

Compared to other database systems, Oracle uses a different approach when it comes to stored procedures and functions. For example, on MS SQL Server, using the `EXEC` command will return a result set. However, Oracle requires the use of an `OUT` parameter with a `REF CURSOR` type to return the result set. In addition, Oracle will not accept multiple statements from a single query. Therefore, it is necessary to store the query within a stored function and use special syntax to access the existing Oracle stored procedures.

To access your Oracle stored procedures, the first step is to define a weakly typed `REF CURSOR` using the following PL/SQL statement:

```
CREATE OR REPLACE PACKAGE types
AS

      TYPE ref_cursor IS REF CURSOR;

END;
```

This `ref_cursor` type will be used to store the query result set and return as an `OUT` parameter. The next step is to create a function which calls your stored procedure and executes your query. The following skeleton code will return a simple query using the `ref_cursor` type.

```
CREATE OR REPLACE FUNCTION my_function()

      RETURN types.ref_cursor

AS

      result_cursor types.ref_cursor;

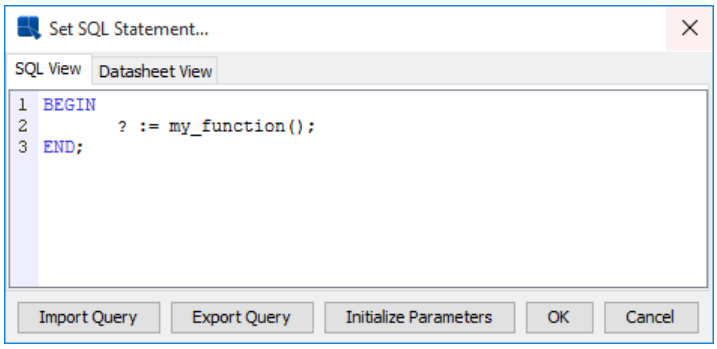
BEGIN

      do_stored_procedure();
      OPEN result_cursor FOR
            SELECT * FROM Categories

      RETURN result_cursor;

END;
```

Now that the Oracle stored function is set up, it can be easily called from ReportDesigner using a special PL/SQL like syntax. In the Set SQL Statement window, enter the following syntax to call the Oracle stored function:



Calling simple Oracle stored function

The `BEGIN ... END;` syntax alerts the system that the user is trying to access an Oracle stored function. And the "?" notifies the ReportDesigner that a variable is reserved for the `OUT` parameter. The JDBC syntax for calling Oracle stored procedures is as follows:

```
( call ? := my_function() )
```

However, ERES does not support this format. Preview the results by clicking the *Datasheet View* tab.

Here is a more practical example to illustrate how stored procedures can be used with ERES to develop useful solutions. Suppose you have a table called *employee_table* that stores an organization's location hierarchy such as the one shown below:

ID	NAME	PARENT	EMPLOYEE
1	All	NULL	0
2	America	1	0
3	Europe	1	0
4	New York	2	20
5	Santa Clara	2	30
6	Dallas	2	12
7	London	3	14
8	Paris	3	11

The table lists the various corporate locations in a tree structure. The numbers of employees are stored in the leaf nodes (e.g. New York, London, etc.) and each node contains information about its immediate parent. Suppose you want to create a report that displays the number of employees in a certain region and information about the separate branches within that region. For example, if the user inputs `ID = 2` (America), you want the report to display the total number of employees in America along with the branch locations. Using Oracle's `CONNECT BY` and `START WITH` clauses, the problem can be solved with two simple Oracle Stored Functions:

```
CREATE OR REPLACE FUNCTION sum_employees(locID IN NUMBER)
      RETURN NUMBER
AS
      sum_emp NUMBER;
```

```
BEGIN

      SELECT sum(employee) INTO sum_emp
      FROM employee_table
      CONNECT BY PRIOR id = parent
      START WITH id = locID;

      RETURN sum_emp;

END;

CREATE OR REPLACE FUNCTION regional_employees (locID IN NUMBER)

      RETURN types.ref_cursor

AS

      result_cursor types.ref_cursor;

BEGIN

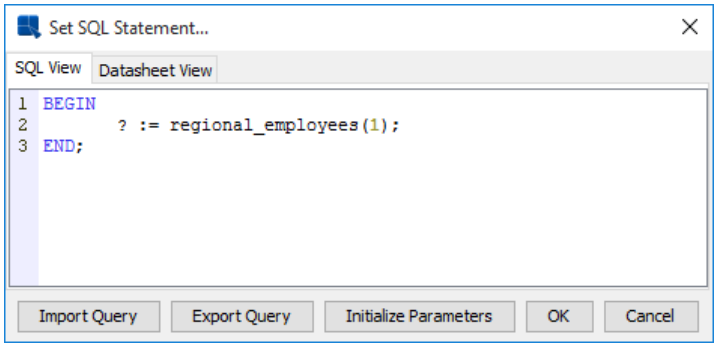
      OPEN result_cursor FOR
            SELECT id, name, sumEmployees(id) AS Employees
            FROM employee_table
            CONNECT BY prior id = parent
            START WITH id = locID;

      RETURN result_cursor;

END;
```

The function `sum_employees` takes the starting node as an argument and finds the sum of all leaf nodes that are descendents of that node. For example, `sum_employees(3)` returns 25 because there are 25 employees in Europe (14 in London, 11 in Paris). The second function, `regional_employees`, traverses through the tree structure starting with the `locID` and builds a result set from the `ID`, `Name` and the result from the `sum_employees` function. The result set is then returned through a `REF CURSOR`.

To call a stored function that requires an argument, enter the following statements in the Set SQL Statement window:



Calling regional_employees function

Preview the results by clicking the *Datasheet View tab*.

SQL View | Datasheet View

ID	NAME	EMPLOYEES
1	All	87
2	America	62
4	New York	20
5	Santa Clara	30
6	Dallas	12
3	Europe	25
7	London	14
8	Paris	11

Result set from regional_employees

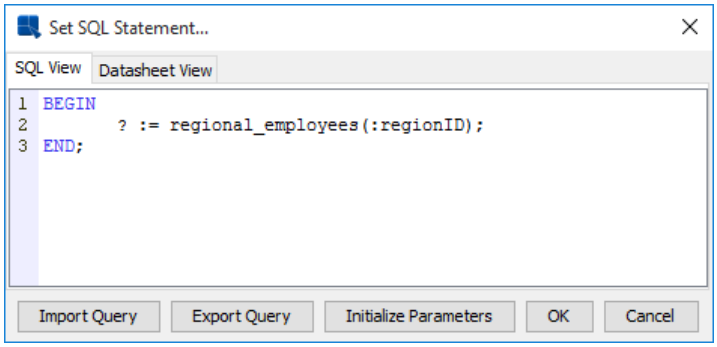
As seen in the results, the `CONNECT BY` clause traverses the tree recursively listing the American nodes together before listing the European nodes. If the user is only interested in European locations, they can enter **3** for the parameter and the following result set would return this:

SQL View | Datasheet View

ID	NAME	EMPLOYEES
3	Europe	25
7	London	14
8	Paris	11

Result set from regional_employees in Europe

To create a parameterized report, use the `:param_name` syntax. The SQL parser in ERES will be able to differentiate between the colon used for parameters and the one used for the assignment operator (`:=`). Here is an example of using the parameters:



Calling Oracle Stored Function using Parameter

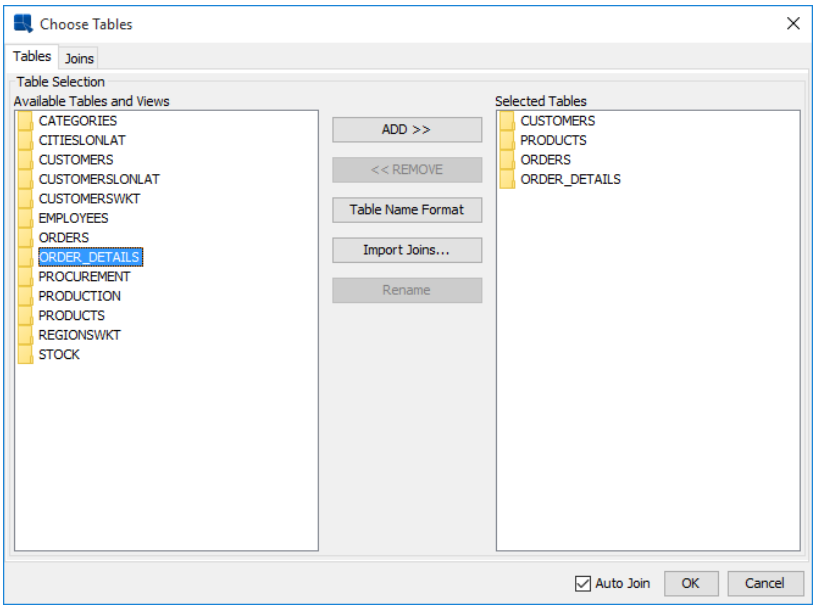
When using `IN` parameters, it is necessary to initialize the parameters prior to executing the query. It is especially important to set the correct default data type for executing stored procedures because the parameters cannot be mapped to existing columns. You can find more information about initializing parameters in Section 4.1.2.2.2 - Parameterized Queries.

4.1.2.3. Data Views

In addition to the query interfaces, ERES provides another way of retrieving database data - data views. Data views provide a simplified view of the database in which users can design queries by simply selecting fields without using the Query Builder or having any knowledge of the underlying database structure. Using data views, administrators can predefine tables, joins, and fields, creating a local schema for the user to select from.

For example, an administrator could set up a data view for the sales department. The appropriate database tables and fields are pre-selected and grouped in a manner congruent with business users' logic. For example a group called *invoices* would have the appropriate customer and order fields. End users would then select this data view, pick the pertinent fields, specify a date range, and then begin designing a report or chart.

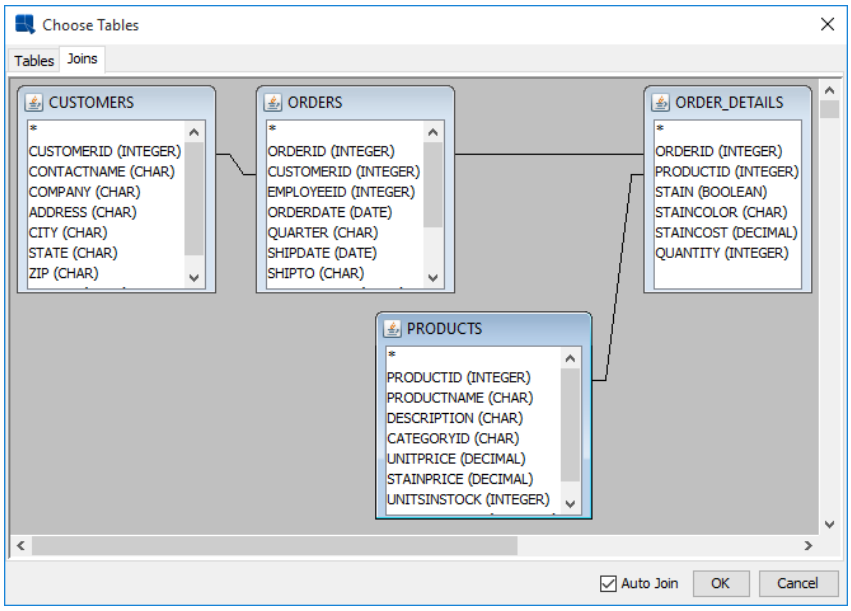
To create a data view, select the *Data Views* node in the Data Source Manager window and click the *Add* button. A new window will open allowing you to select the database tables you want to use for the data view.



Data View Choose Tables Dialog

Left window contains all available database tables and views. You can add a table by selecting it in the left window and clicking the *ADD>>* button. By default, the data view will use the name format you specified when setting up the database connection. You can change the naming by clicking the *Table Name Format* button, or specify a table alias by clicking the *Rename* button. You can also import selected tables and joins from another data view by clicking the *Import Joins…* button.

The *Joins* tab of this window allows you to specify the joins between the selected tables.

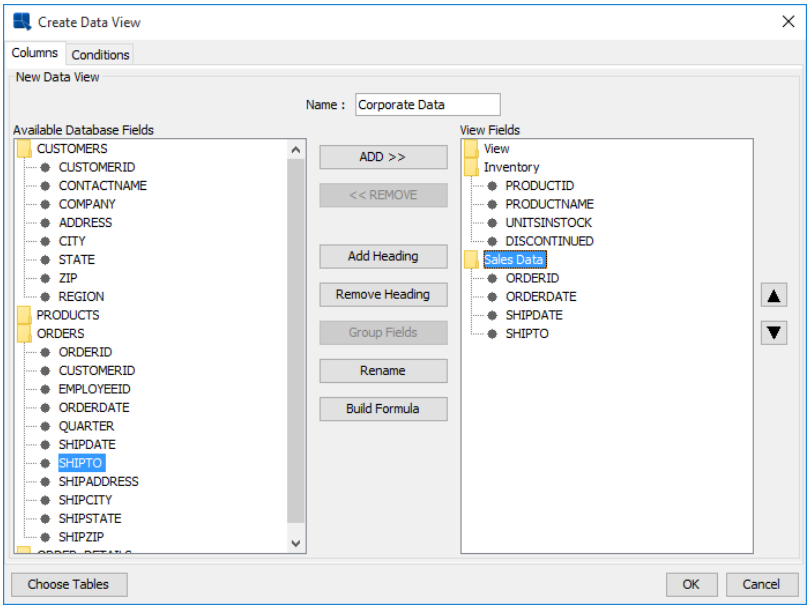


Data View Joins Dialog

The *Joins* tab shows all selected tables and their associated fields. The tables will be auto-joined depending on which option you selected when setting up the database connection. These auto-joins create a standard join between tables represented by a line. To remove a join or edit join properties, right click on the line and select your choice from the pop-up menu. To add a join, click and drag one column field to another in a different table. A join will then

appear. Data views use the same join properties as the Query Builder. For more information about join properties, please see Section 4.1.2.2.1.2 - Joins.

After you finish selecting and joining tables, click the *OK* button and a new window will open allowing you to construct the data view.



Create Data View Dialog

The left window contains a list of tables you have selected and their associated fields. Each folder represents a table and can be opened and closed by double clicking. The right window contains fields that have been selected for the data view. To add a field to a data view, select it in the left window and click the *ADD>>* button. Fields can be removed from the data view in the same way by selecting a field in the right window and clicking the *<<REMOVE* button. You can create a calculated column by clicking the *Build Formula* button. This will open the formula builder, allowing you to build the column. You can also define an alias by selecting any of the view fields in the right window and clicking the *Rename* button.

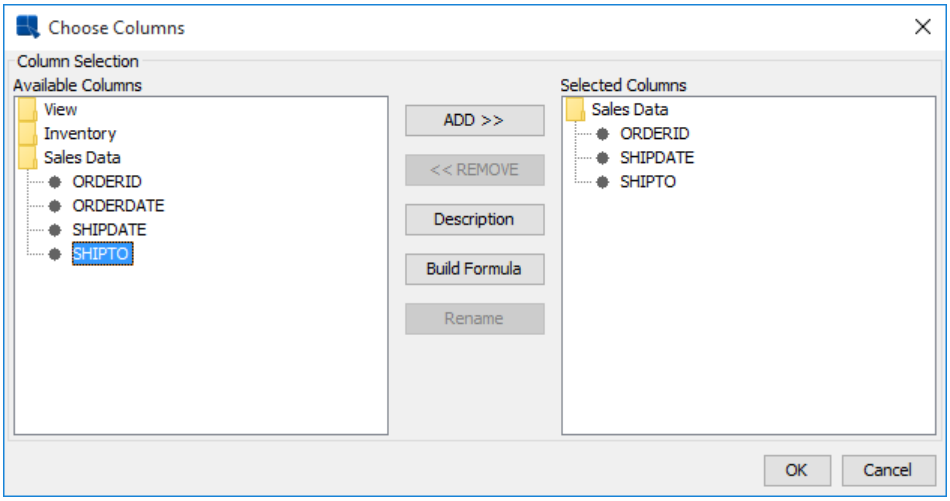
You can also group fields within the data view by adding headings. This allows you to create your own organizational structure of *virtual tables* that group data from different database tables under one heading. To create a heading, click the *Add Heading* button. You will then be prompted to specify a name for the heading. The new heading will then appear as a folder in the right window. To add fields under a heading, first select the fields you want to add from the right window and click the *Group Fields* button. You will then be presented with a drop-down menu, allowing you to select the heading under which you would like to add the fields.

The *Conditions* tab contains a formula builder window that allows you to specify certain filtering criteria for end users. Anything added in this window will be added to the *Where* clause of the generated SQL. For more information about using the formula builder, please see Section 4.1.2.2.1.3 - Columns.

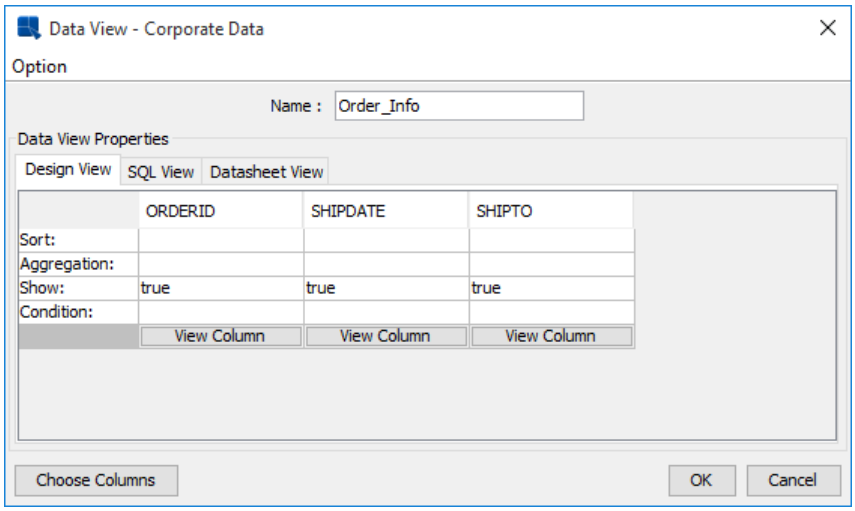
Once you finish creating the data view, click the *OK* button and the data view will be added to the Data Source Manager. Users can now use this view to construct ad-hoc queries.

When you design a report or chart using a data view as the data source (by selecting the data view and clicking the *Create Report* or *Create Chart* button), a window will open allowing you to select which fields in the view you want to use for the report. From this dialog, you can also build computed fields based on the available view columns.

After you select the fields, click the *OK* button and a new window will open allowing you to specify sorting, aggregation, and filtering conditions for the data view.



Data View Choose Fields Dialog



Data View Conditions Window

You can specify sorting, aggregation, and conditions for every field in the data view by double clicking on the respective field. Sorting and aggregation can be selected from drop-down menus. Double clicking on the *Conditions* field brings up a new window that allows you to specify simple selection criteria like `>`, `<`, `=`, and `between`. Users can build more advanced filtering criteria by right clicking on the *Conditions* field and selecting *Build* from the pop-up menu. This will open the Formula Builder window allowing you to build a condition. You can also display all of the unique values in the column by double clicking on the *View Column* button.

The Option menu in the upper left corner of the conditions window allows you to select a vertical/horizontal view for the conditions window, initialize any parameters in the data view, or save the query.

The selection set and conditions that you specify will be saved as a data view query with the name that you specify in the name field. Data view queries are saved under the node for the data view. A report created from the data view will refer to the data view query for updating/modification.

4.1.2.3.1. Data View Parameters

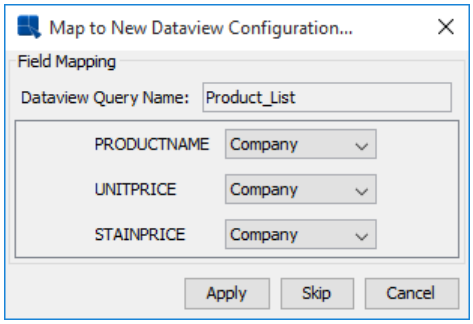
As with Query Builder, users can specify query parameters in Data Views. To add a parameter to a data view, select a data view in the Data Source Manager and click *View* to run the data view. After you select fields for the data view and you are in the conditions window, right-click in the *Condition* field for a column and select *Build* from the pop-up menu. This will bring up the formula builder, allowing you to specify a parameter in the same way as in Query Builder. For more information about this, please see Section 4.1.2.2.2 - Parameterized Queries.

Once you enter the parameter, you will be prompted to initialize it if you go to the *Datasheet View* tab and then click *Ok* to continue with the report wizard, or if you save the selections as a query. You can also initialize the parameter by selecting *Initialize Parameters* from the Option menu.

4.1.2.3.2. Updating Data View Queries

Sometimes you may need to make changes to the structure/make-up of the data view as your data model or requirements change. Changes could include adding/removing fields or re-naming them. You can propagate changes from the data view to its associated queries by selecting it in the data source manager and selecting *Data View Queries* from the Update menu.

All of the queries associated with the view will be scanned and any inconsistencies in fields or field names will be presented for you to update.



Update Query Fields Dialog

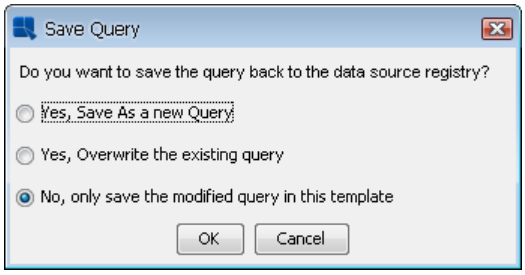
For each query, you will be prompted to change any fields that no longer match the data view structure. For each field, you can select a field from the data view to map it to, or remove the field from the query. If you do not want to change anything in the query, you can click the *Skip* button. The query will continue to run, but it will refer to the old data view structure. Click the *Apply* button to save the changes to the data view query.

4.1.2.4. Editing Queries

If you have selected to build a report using database data either by designing a query in the Query Builder, writing an SQL statement, or running a data view, you can modify the query directly from the Report Designer or the Chart Designer without having to go back to the Data Source Manager.

To modify a report or chart's query, select *Modify Query* from the Data menu in Report Designer or Chart Designer. If you have designed a query in the Query Builder, then the Query Builder interface will re-open allowing you to modify the query. If you have entered an SQL statement, a text box will open allowing you to modify the SQL. If you have used a data view, the data view conditions window will re-open allowing you to change the filters or pick additional fields.

Once you specify all the changes, you will be given an option to modify the query in the data registry, save a new query in the data registry, or modify only the query in the template.



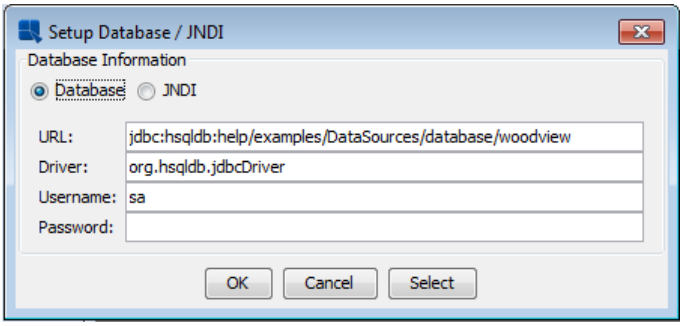
Saving Query Options

Once you specify the save options, the modified query will be applied to the report or chart. Note that if you made significant changes to the query, you may need to perform the data mapping again.

4.1.2.5. Editing Database Connections

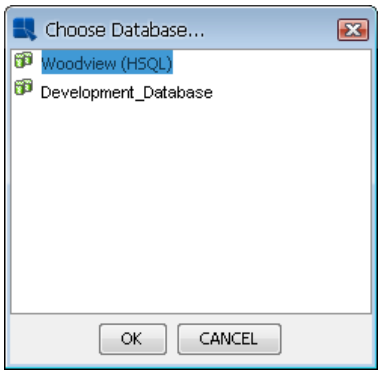
If you have selected to build a report or chart using database data, you can also directly edit a template's database connection information from within the Report Designer or the Chart Designer. To modify the connection informa-

tion in the report or chart, select *Modify Database* from the Data menu. This will bring up a dialog allowing you to specify a different URL, Driver, Username, and/or password for the report to use when connecting to the database.



Change Database Connection Dialog

In addition to manually entering the database information, you can retrieve the database connection information from a data registry. To do this, click the *Select* button on the Database Connection dialog. This will allow you to browse to XML registry file from which you want to pull the database connection. When you select a registry file, you will be presented with a list of databases defined in the registry.



Select Database from Registry

Select the database that you want to use and click the *Ok* button. The connection information for that database will be automatically applied to the connection dialog. After you set the connection information for the template, click the *Ok* button to apply the changes. A dialog will open asking you if you would like to verify the new connection information.



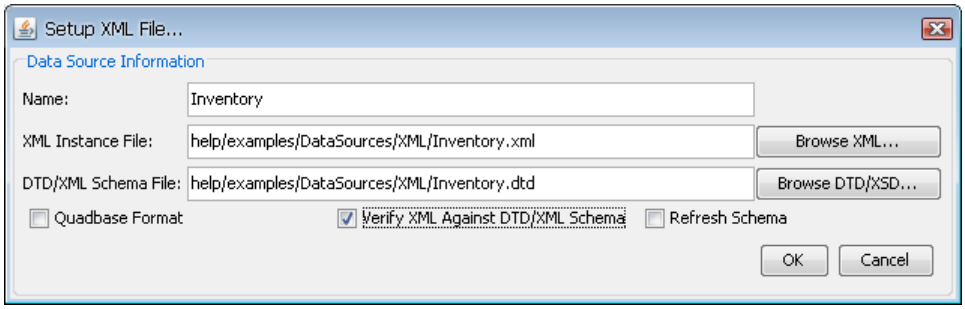
Verify Connection Dialog

Unless you know that data source isn't present, it's generally a good idea to check that the supplied connection information is correct. If you're changing the database connection information in a report, all the sub-reports, charts, and drill-down layers in the report that use the same connection information will be automatically updated as well.

Unlike the modify query feature, changes to a template's database connection will only saved be to the template. It will not be saved to the data registry.

4.1.3. Data from XML and XBRL Files

In addition to relational databases, ERES allows you to retrieve data and query XML files. XML data can be in virtually any format, but you need to specify a DTD file or an XML Schema (XSD) along with the XML data. To set up an XML data source, select *XMLFiles* node in the Data Source Manager and click the *Add* button. A dialog will open prompting you to specify options for the new XML source.



Setup XML Data Source Dialog

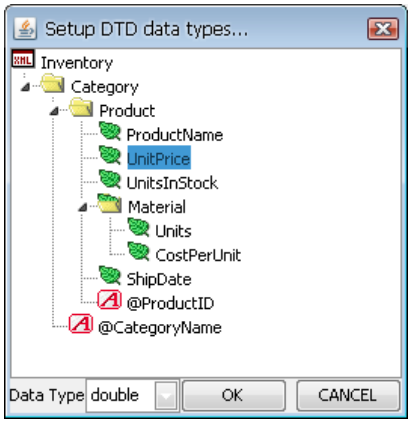
The first option allows you to specify a display name for the XML data source. The second option allows you to specify the location of the XML file from which you want to retrieve data. Note that you can also specify an XBRL file in this field. You can set up a data source that retrieves XML data from an HTTP server here as well, by adding the appropriate URL as the file location. The third option allows you to specify the location of a valid DTD or XML schema file for the XML file.

The *Quadbase Format* checkbox allows you to indicate whether the XML file is in the form of an XML export from ERES. For example, if you choose to export a report's data in XML format, you can read it back in using this format. When you use a file in this format, you do not have to specify a DTD or XML Schema.

The *Verify XML against DTD/XML Schema* checkbox will make sure that the supplied XML file/source complies with the layout specified in the DTD or XML schema file. Because queries are designed based on the structure of the DTD/XML Schema file, a non-conforming XML source could produce unexpected results. If the XML does not conform to the DTD or XML schema, you will be given a warning. You can, however, continue setting up the data source.

The *Refresh Schema* checkbox will only appear if you choose to edit an existing XML data source. Checking this option will reload the schema or DTD definition to incorporate any changes to the structure in the XML data source in the registry. This option is only necessary if the DTD or schema definition has changed since the data source was first created.

Once you finish setting up the XML file and the DTD/XML Schema, a new dialog will open allowing you to specify the data type for all the selectable elements in the XML data source. The dialog will be different depending on whether you are using a DTD file or XML Schema.



DTD Data Type Selection Dialog

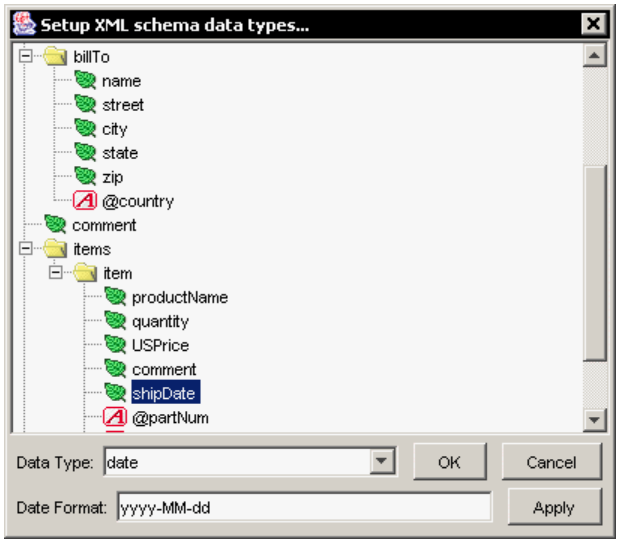
Because DTD files do not specify a correct data type for elements, all elements are considered to be Strings by default. To change the data type of an element, you have to select the element and pick a data type from the drop-down window at the bottom of the dialog. To ensure proper results when you query the XML file, you should set the data type for all selectable elements. This includes leaf nodes, parent nodes that contain data, and attribute elements. The following data types are supported:

- String

- Integer
- Double
- Date (If you specify date as the data type, you will also be required to specify the date format.)
- Boolean

Once you finish specifying the data types, click the *OK* button and the XML source will be added to the Data Source Manager.

If you're using an XML schema, a different data types dialog will open.

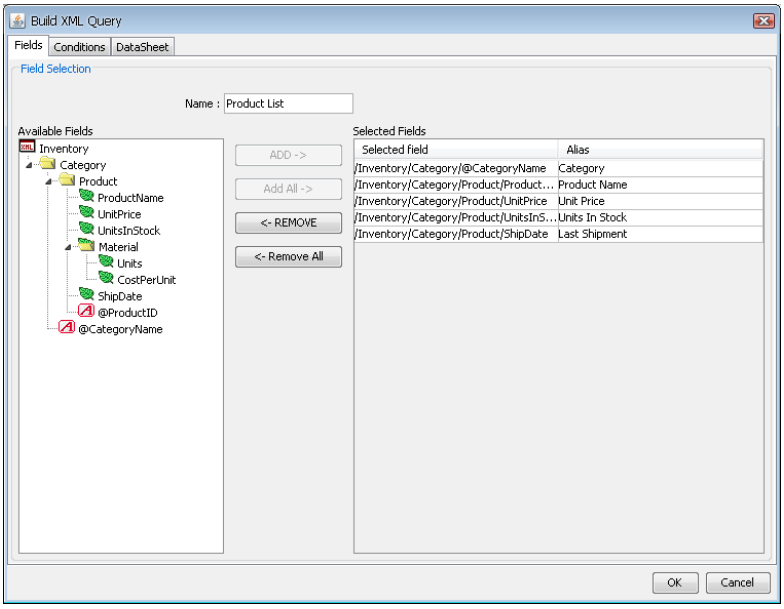


XML Schema Data Type Selection Dialog

Generally, the data types should already be defined in the XML schema file, but you can make any changes in this dialog. Once you finish specifying the data types, click the *OK* button and the XML source will be added to the Data Source Manager.

4.1.3.1. XMLQueries

Once you set up an XML data source, you can then create queries to select nodes, specify filtering conditions, and transform the tree structure into the tabular form used by ERES. To add a query, select the node for your XML source and click the *Add* button. This will launch the XML query builder interface that allows you to construct a query.



XML Query Field Selection Tab

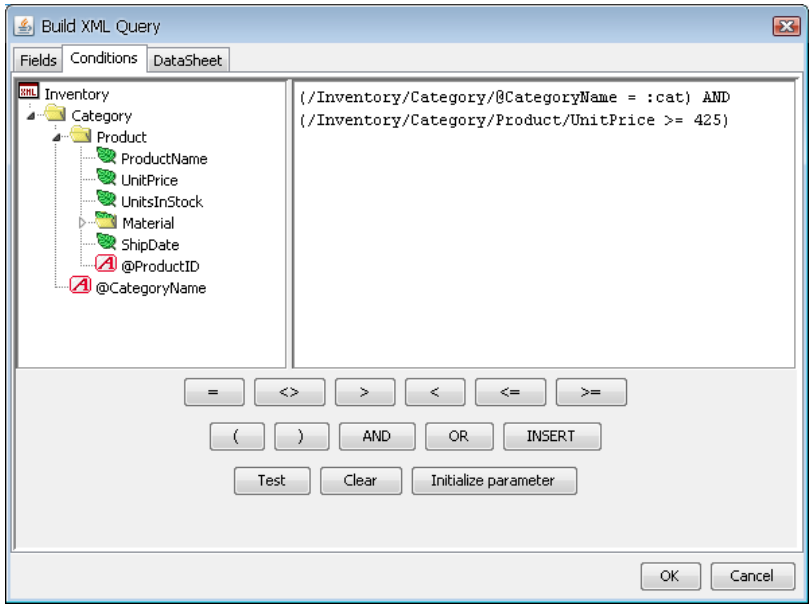
The first tab in the XML query builder allows you to select the fields/nodes from the XML file that you would like to use in the report. The left side of the window contains the tree structure from the DTD or XML schema file. You can pick any selectable elements and add them to the query by clicking the *Add* button. Selected fields will appear in the right side of the window. You can specify an alias for any field by double clicking the *Alias* field for a column and typing the new column alias.



Note

Each selected element will become a column in the report or chart data. For results where a one-to-one relationship cannot be determined, the tabular structure is built using all available permutations in the data (similar to a cross-join in SQL). For best results, it is recommended that you select fields for a query where a clear hierarchical relationship is present.

The *Conditions* tab of the XML query builder allows you to specify some basic filtering criteria for your selection.



XML Query Conditions Tab

You can specify an equal, not equal, greater than, less than, less or equal to, or greater or equal to condition for any selectable element in the XML file. You can also use the AND and OR operators to build compound conditions. Fields are specified using a direct path down the XML tree. Currently, only direct path is supported. You cannot use more complex XPath expressions. To add a field, you can double click on it in the left side, or you can select it and click the *Insert* button.

After you finish writing the conditions, click the *Test* button to verify that the syntax is correct.

The *DataSheet* tab allows you to preview the query result and see how the XML data is converted to tabular form. You can navigate through the result set the same way as in Query Builder ().

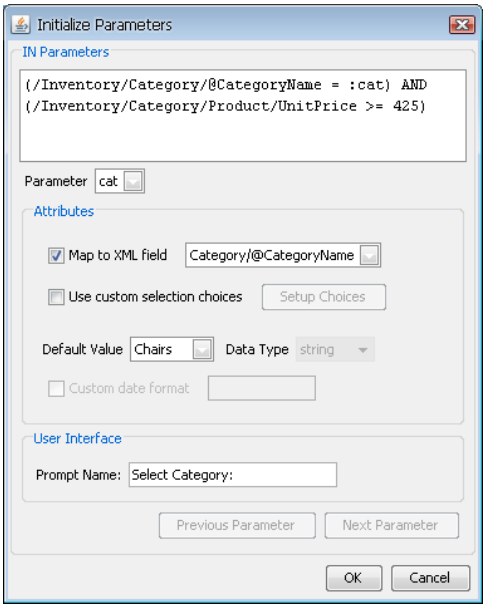
Once you select the fields and specify the appropriate conditions, click the *OK* button. The query will then be added as a new node under your XML source in the Data Source Manager and can now be used to create a report or chart.

4.1.3.1.1. XML Parameters

As with database queries, you can also specify parameters for XML queries. The same syntax ":" is used to denote a parameter in the XML condition as it is in a query condition. So the following XML condition:

```
/Inventory/Category/@CategoryName = :category
```

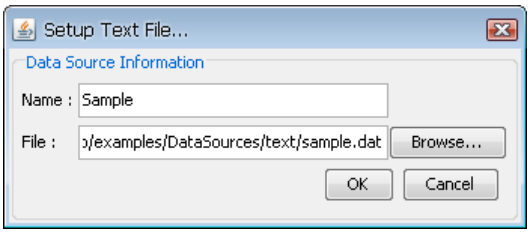
would place a dynamic filter on the query for the *CategoryName* attribute. XML parameters are initialized in the same way as query parameters. The initialization dialog will appear if you try to preview or close the query, or you can trigger it by clicking the *Initialize Parameters* button. The only difference is that instead of mapping to a database field, the parameter prompt can be mapped to a node in the XML file.



XML Parameter Initialization Dialog

4.1.4. Data from Text Files

ERES also allows you to retrieve data from flat text files. To add a text file as a data source, select the *TXTFiles* node in the Data Source Manager and click the *Add* button. A dialog will open prompting you to specify a display name and the location of the text file you want to use.



Add Text Data Source Dialog



Tip

The text file can be also retrieved from URL. To do so, enter the URL to the *File* text field. You have to enter a full URL with protocol etc. (e.g. http://www.quadbase.com/textfile.txt).

After you specify the information, click the *OK* button and the text source will appear under the *TXTFiles* node of the Data Source Manager window.

4.1.4.1. Formatting Requirements for Text Files

There are certain formatting requirements for the data within a text file in order make it readable by ERES. Generally, data is expected to be in a form similar to the following:

```
String,date,decimal
Name,day,Volume
"John","1997-10-3",32.3
"John","1997-4-3",20.2
"Mary","1997-9-3",10.2
"Mary","1997-10-04",18.6
```

The above data file is a plain text file. The first row specifies the data types and the second row specifies the field names. The third row and so on are the records. Every text file must consist of these three parts. There are four records, with three fields each in the example data file. The delimiter between the fields may be one of the following characters: `","`, `";"`, or `" "` (that is comma, semi-colon, or space). Each field can be put in quotes (single or double).

4.1.4.2. Data Types and Format for Text Files

In text data files, the data type is specified using a keyword. The following is a list of recognized keywords and their corresponding JDBC type and Java type.

Data File Keywords (Not Case Sensitive)	JDBC Type	Java Type in ERES
Boolean, logical, bit	BIT	Boolean
tinyint	TINYINT	byte
smallint, short	SMALLINT	short
int, integer	INTEGER	int
long, bigint	BIGINT	long
float	FLOAT	double
real	REAL	float
double	DOUBLE	double
numeric	NUMERIC	java.math.BigDecimal
decimal	DECIMAL	java.math.BigDecimal
date	DATE	java.sql.Date
time	TIME	java.sql.Time
timestamp	TIMESTAMP	java.sql.Timestamp
string	CHAR	String
varchar	VARCHAR	String
longvarchar	LONGVARCHAR	String

For certain data types, the data in a text file must be presented in a specific format. The following is a list of the data types that require specific formatting:

Data Type	Format	Example
Date	yyyy-mm-dd or yyyy-mm	2001-06-12 or 2000-06
Time	hh:mm:ss	12:17:34
Timestamp	yyyy-mm-dd hh:mm:ss	2001-06-12 12:17:34
Boolean	true/false, t/f, 1/0 (case insensitive)	true

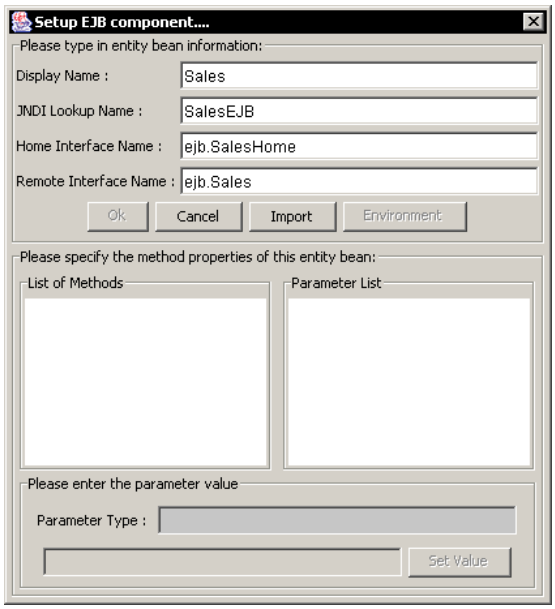
4.1.5. Data from EJBs

Using Enterprise JavaBeans™ technology, developers can simplify the development of large enterprise applications. With EJB technologies, developers can rely on building business logic and allow the application server (EJB container) to manage all system level services.

When working in the Java EE™ environment, persistent data interfaces are provided by entity beans. Although the underlying storage mechanism might be a relational database, the application data model is the EJB and it may not be desirable to have a reporting tool making redundant database connections. For this situation, ERES allows users to query data directly from an entity bean.

To add an EJB as a data source, the EJB must first be deployed in the application server and the client JAR file containing the appropriate stub classes must be added to your application server/servlet runner's classpath. Select

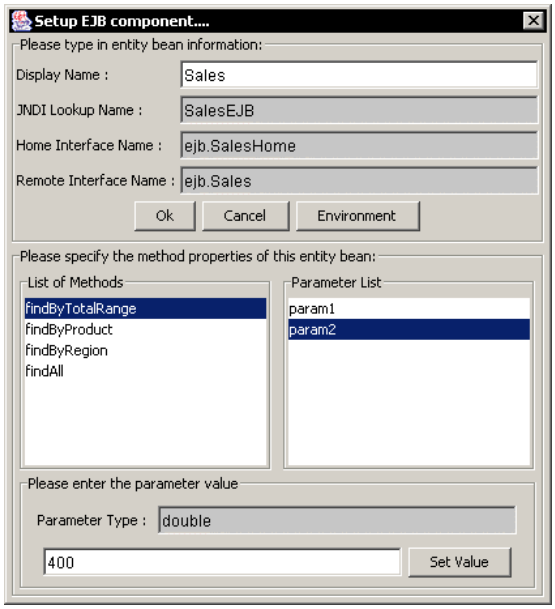
the *EJBs* node in the Data Source Manager and click the *Add* button. This will bring up a dialog allowing you to specify a display name and the connection information for the bean.



Add EJB Dialog

To connect to an EJB data source, you must provide the JNDI lookup name for the bean (this is specified when you deploy the bean). For EJB 1.1 users, specify the name for the home interface and the remote interface. For EJB 2.0 users, specify the local home interface and the local interface. Once you specify all the information, click the *Import* button, which will analyze the home interface and retrieve all of the finder methods. Any methods found will be populated in the *List of Methods* section in the dialog.

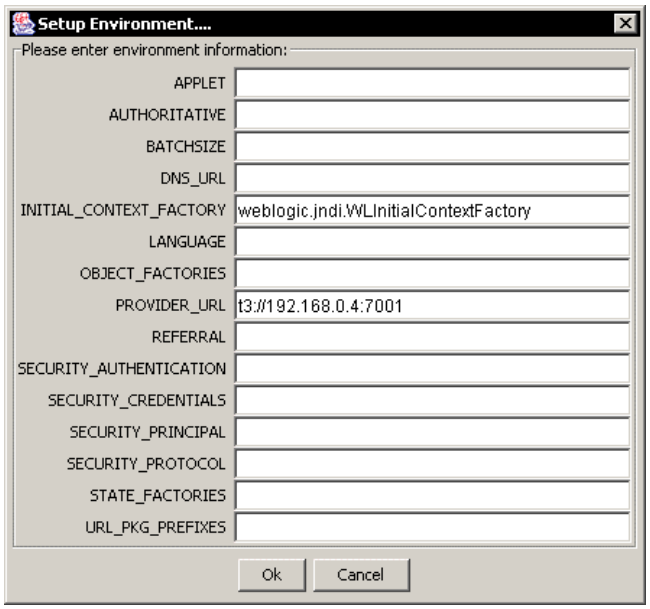
The same dialog can be used to filter the data being retrieved based on parameters that are present in the finder methods. When you select a method in the left dialog, any parameters present will appear in the *Parameter List* section. You can then click on a parameter to set its value.



Specifying EJB Parameter Values Dialog

When you select a parameter, the data type of the parameter will appear in the lower portion of the window. Below that you will be able to specify a value for the parameter. Be sure to enter a correct value for the data type, and then click the *Set Value* button. This will fix the parameter values. Once you finish setting up all parameter values, click

the *Environment* button. This will bring up a new dialog allowing you to specify environment properties for your application server. This information is necessary for the ERES Server to connect to the EJB.



EJB Environment Setup

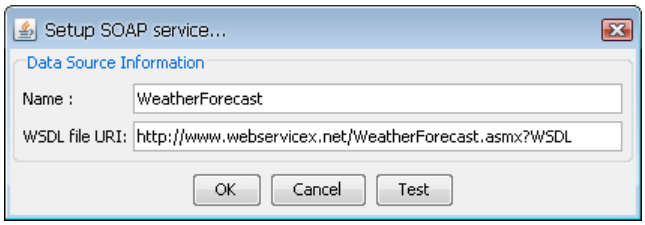
The fields here are the available environment properties for the JNDI context interface. You don't have to specify all values, only the information necessary for your environment (application server). Once you finish specifying the environment variables, click the *OK* button. You will be returned to the EJB setup window. Click *OK* again to finish setting up the EJB data source. A new node will appear under *EJBs* with your EJB.

You can modify the parameter values by selecting your EJB source, and clicking the *Edit* button.

4.1.6. Data from SOAP with WSDL support

ERES also allows you to retrieve data using SOAP (Service Oriented Architecture Protocol). To connect to a SOAP data source using WSDL, you don't need to know any URLs for the services, SOAP actions, operation names or parameters. All you need to know is a location of WSDL file, which contains all the necessary information.

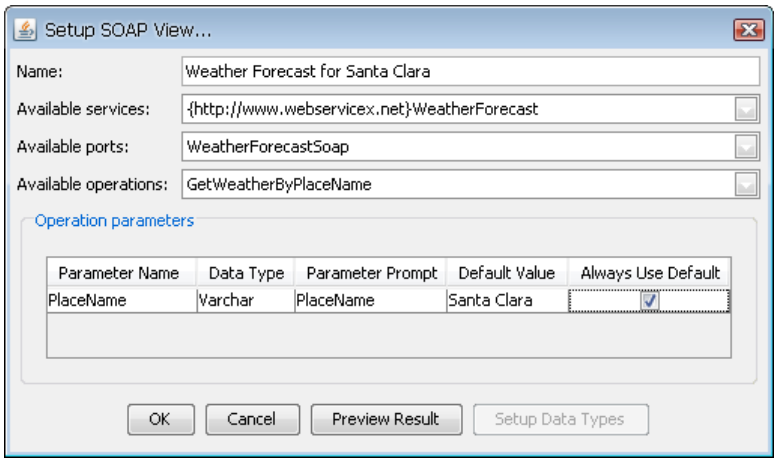
To set up a SOAP data source, select the *SOAPServices* node in the Data Source Manager and click the *Add* button. A dialog will open prompting you to specify options for the new SOAP data source.



SOAP data source setup

The first option allows you to specify a display name for the data source. The second option allows you to specify an URL of the WSDL file. Once you specify the connection information, you can test the connection to the WSDL file by clicking the *Test* button. This will test retrieving the WSDL file from the URI you've provided, check the file for any supported SOAP operations and report any problems. After clicking *OK*, a new SOAP data source node will be added to the data registry.

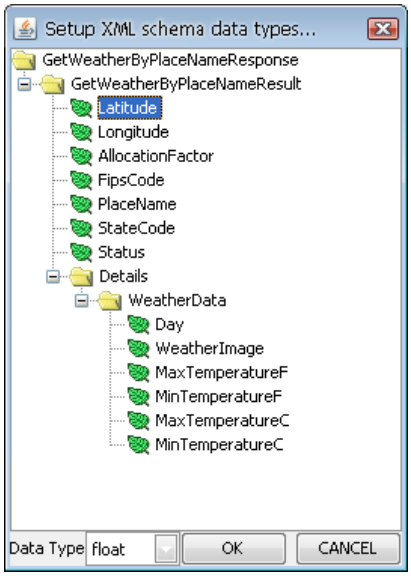
To add a new SOAP View, select an existing SOAP data source and click the *Add* button. A dialog will open prompting you for all the parameters necessary to make a SOAP query.



Setup SOAP View dialog

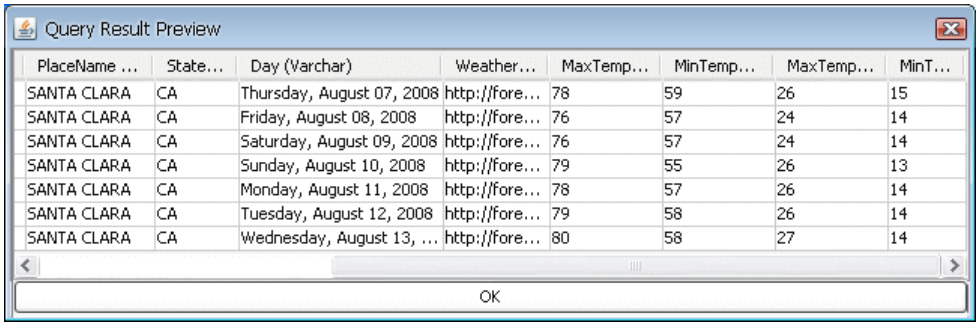
The first option in the dialog allows you to specify a display name in the Data Source Manager. Next, there are three drop-down menus in the dialog. The first drop-down menu contains all the SOAP services described in the WSDL file. Once you specify a service, the second drop-down list will be populated with all the ports of this service. Selecting a port will populate the last drop-down list with all the operations of this port. If you move the mouse over any drop-down list, a hint will appear with documentation for the service/port/operation (if the documentation is provided in the WSDL file). After specifying the service, port and operation, all the parameters of the operation will be read. If there are some parameters, they will be displayed in a table. The first two columns of the table are not editable. They are read from the WSDL file. The next 2 columns are editable and allow you to specify parameter prompts and default values. The last column contains a checkbox which allows you to choose whether the default parameter value will always be used or not. This means that this parameter value will be fixed and you will not be prompted for it. All the parameters that don't have this checkbox checked will be used as report/chart parameters.

The *Setup Data Types* button is only available when editing the SOAP View from the Data Source Manager and allows you to adjust data types when necessary. In order to verify result from the SOAP response, click the *Preview Result* button. All the default values will then be tested to see if they have proper data type or not. In case they do not match, you will be prompted to adjust them. After that, you will get the setup data types dialog (the same as for XML data source). If the data source is parameterized, you will get the parameter prompt dialog before specifying data types. Please note that in order to generate the XML schema properly, you have to specify existing parameter values.



Setup XML Schema Data Types dialog

You can setup data types from this dialog. The behavior is exactly the same as for XML data source with DTD schema (see Section 4.1.3 - Data from XML and XBRL Files). Once you finis specifying data types, click the *OK* button. A dialog will open showing you the result preview.



Query Result Preview dialog

Clicking the *OK* button in this dialog will take you back to the Setup SOAP View dialog. Once you finish specifying all necessary information, click the *OK* button in the dialog. A new node will be added under your SOAP data source in the Data Source Manager and can now be used to create a report or chart.

4.1.7. Data from Salesforce

The Salesforce data source is designed for existing Salesforce users who want to display their Salesforce data in ERES. The connection to the Salesforce server is established via SOAP using Salesforce Partner WSDL (version 13.0). Users communicate with Salesforce server by SOQL (Salesforce Object Query Language) queries. Please note that users must have valid Salesforce accounts with username and password to work with this data source. Moreover, users who use the ERES Salesforce data source must have access to Salesforce account from trusted networks. To add your IP address to the trusted IP list, you have to activate your computer as described below.

For more information about SOQL queries and activating Salesforce user's accounts from trusted networks, please visit the following Salesforce sites:

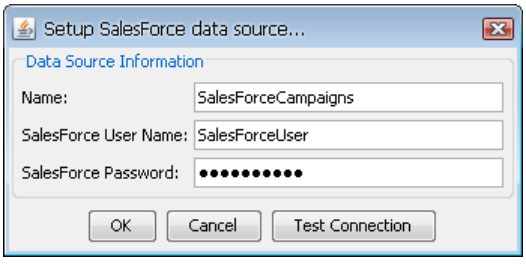
SOQL queries

https://developer.salesforce.com/docs/atlas.en-us.soql_sosl.meta/soql_sosl/sforce_api_calls_soql_sosl_intro.htm

Activating Salesforce user's accounts

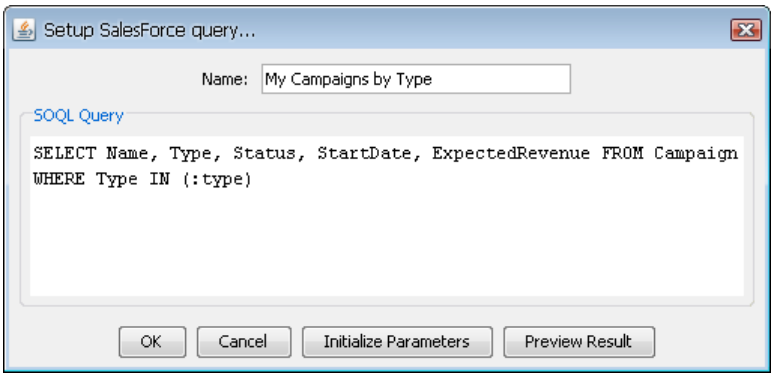
https://help.salesforce.com/s/articleView?id=sf.security_networkaccess.htm

To set up a Salesforce data source, select the *SalesForce* node in the Data Source Manager and click the *Add* button. A dialog will open prompting you to specify a display name for the data source, user name and password to your Salesforce account. Once you specify the connection information, you can test the connection to your Salesforce account by clicking the *Test Connection* button. This will test the connection using the information you've provided and report any problems.



Setup SalesForce Data Source Dialog

Once you add a Salesforce data source, a new node will appear in the Data Source Manager window. To add a new Salesforce query, click the *Add* button. A new dialog will open prompting you to specify a query name and SOQL query.



Setup SalesForce Query Dialog

Please note that only child-to-parent relationship queries are supported in the current ERES version. You cannot use parent-to-child queries (using nested SOQL queries). For more information about Salesforce relationship queries and their syntaxes, please visit the following Salesforce site: https://developer.salesforce.com/docs/atlas.en-us.soql_sosl.meta/soql_sosl/sforce_api_calls_soql_relationships.htm

Moreover, this dialog allows you to initialize query parameters in case that your query contains single value or multi value parameters. A parameter is specified within an SOQL statement using the ":" character. Generally the parameter is placed in WHERE clause of an SOQL Select statement. For example, the following SOQL statement

```
Select Name, Type, Status, ActualCost From Campaign Where Name
 = :CampaignName
```

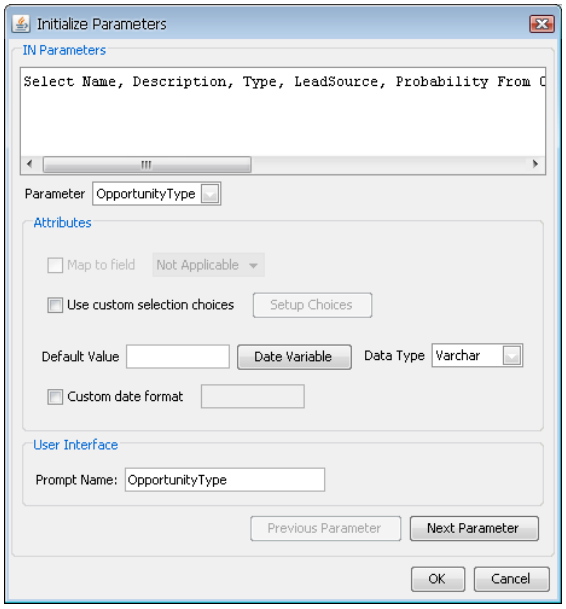
specifies a single value parameter called *CampaignName*. You would then be able to enter a campaign name at run-time and only retrieve data for that campaign.

Another example of SOQL statement shows using of multi value parameters that take an array of values as the input rather than single values.

```
Select Name, Description, Type, LeadSource, Probability From Opportunity
 Where Type IN (:OpportunityType) And LeadSource IN (:OppLeadSource)
```

The statement specifies two multi value parameters called *OpportunityType* and *OppLeadSource*. You would then be able to specify opportunity types and lead sources at run-time and you will only retrieve data according to specified parameters values.

In order to initialize SOQL query parameters, click the *Initialize Parameters* button. The initialize parameters dialog will then appear allowing you to specify parameters mapping.



Initialize Parameters Dialog

From this dialog, you can specify the following options:

Map to field: This allows you to specify a field from the Salesforce data source whose values will be used for the parameter input. Selecting this option modifies the parameter prompt that you will get when previewing or running the report/chart. If you map the parameter to a Salesforce field, you will be prompted with a drop-down list of distinct values from which you can select a parameter value. If you do not map, you will have to type in specific parameter value.

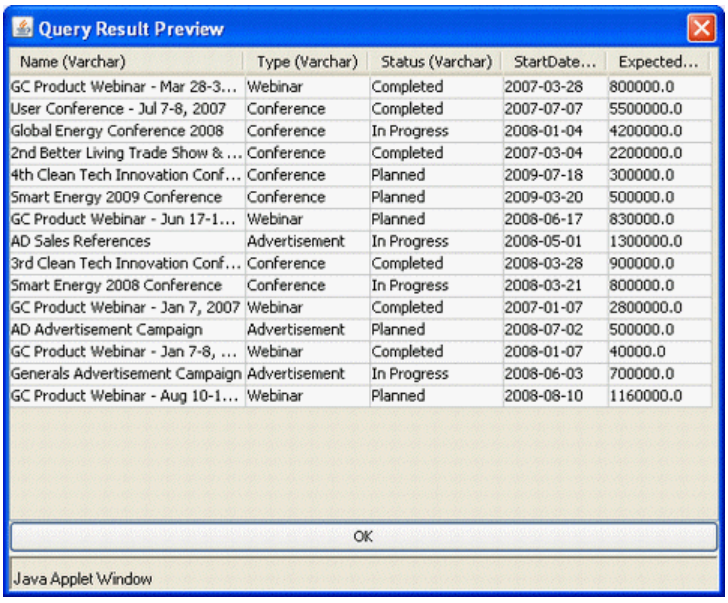
Use custom selection choices: Rather than having a drop-down menu with all the distinct column values, you can build a custom list of parameter values. To set up the list, select this option and click the *Setup Choices* button. This will launch a new dialog allowing you to create a list of choices.

The rest of the options are basically same as for database query parameters. For further information about initializing database query parameters, see Section 4.1.2.2.2.2 - Initializing Query Parameters. Once you specify mapping for all available parameters, click the *OK* button and you will be taken back to the Setup Salesforce Query dialog.

From the Setup SalesForce Query dialog, you can also preview the query result using the *Preview Result* button to verify output from your query. In case you have a parameterized query, the parameter prompt dialog will appear prompting you to specify parameter values. Once you specify the parameter values, click the *OK* button and the query result preview dialog will appear.



Parameter Prompt



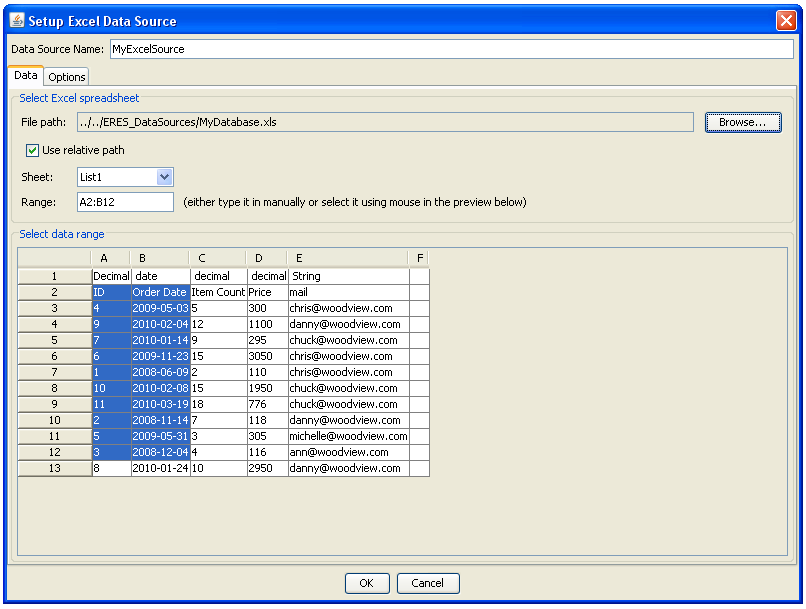
Query Result Preview Dialog

From this dialog, you can verify the query output. Clicking the *OK* button will take you back to the Setup SalesForce Query Dialog.

Once you specify the query, click the *OK* button. The query will then be added as a new node under your Salesforce data source in the Data Source Manager and can now be used to create a report or chart.

4.1.8. Data from Excel files

ERES also allows you to design reports and charts using data retrieved from Excel files. To add an Excel file as a data source, select the ExcelFiles node in the Data Source Manager and click the *Add* button. A dialog will open prompting you to specify the data source name and to select the Excel file from which the data should be imported.



Setup Excel Data Source dialog - Data

After selecting the Excel file, the imported data will be previewed in the dialog. If the checkbox *Use Relative Path* is checked, the file path to the selected Excel file will be set as relative to the ERES installation directory. Otherwise the full path will be used. In case the Excel file is stored on a different disk drive from the one which ERES is installed on, this option is not available. You can select the sheet (if there is more than one in the file) and the cells which are relevant for the data source being designed using your mouse or by specifying the range in the *Range* box (for instance, you can specify that your source will use the data from the columns A and B and from the rows 2 to 12 by typing **`A2:B12`** in the *Range* box; this is the format which is also used for ranges in MS Excel). You can also select the data by clicking the row header or the column header. Hold **Ctrl** or **Shift** to select more rows or columns. Click the top left corner to select all cells. Again, this behavior is similar to MS Excel.

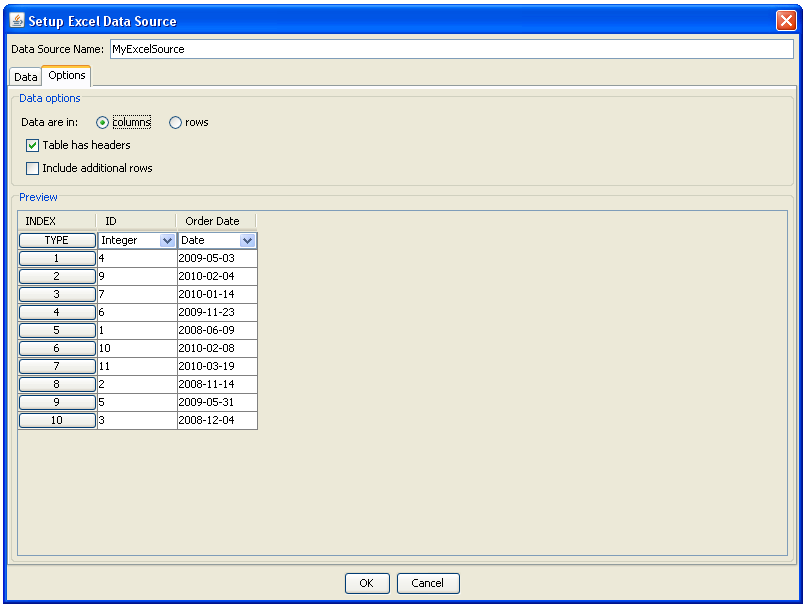
Please note that ERES can process both `*.xls` files and `*.xlsx` files. The `*.xlsx` files are used in Microsoft Excel 2007 and newer and they are based on Open XML.

ERES can also process basic Excel formulas. If it is not able to process the entered formula, an error message will be displayed.

Empty rows (in case the data are in columns) or columns (in case the data are in rows) at the end of the sheet are automatically removed from the data source, even in case they have been selected.

Click on the *Options* tab to specify whether the data are in columns or in rows to get the data structure correctly. You can also specify whether table headers are included in your data selection. The option *Include additional rows* or *Include additional columns* (depending on whether the data are in columns or rows) allows you to automatically include data rows or columns into the Excel file after the data source has been created without the need to change the data source in the Data Source Manager manually.

The bottom part of the dialog on the *Options* tab shows you the data source content following the current configuration. You can change the data type for each column of the data source there, if desired. The data types are detected automatically, therefore changing the type should not be necessary in most cases.



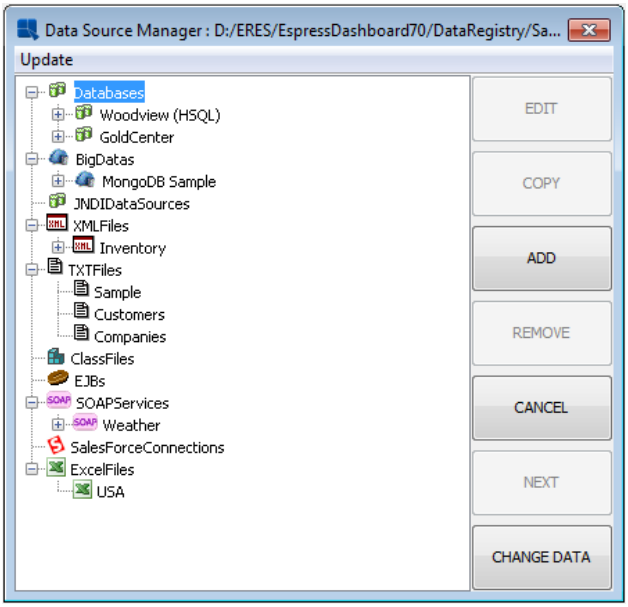
Setup Excel Data Source dialog - Options

Click *OK* to save the data source when the configuration is finished.

4.1.9. Using Data for Charts and Reports

Once you set up a data registry, the sources defined in the registry can be used for charts and reports. The first time you select to start a new report or chart, you will be prompted to select a registry you want to use. If there are currently no available registries (meaning that you haven't created any, or you don't have privileges to view any), you will be prompted to go the data registry manager to create one.

Once you select a registry, a data source manager window will open allowing you to select, add, or modify a data source that you want to use for the chart or report.



Data Source Manager Window for Chart Designer

After you first select a registry, the next time you design a report or chart, it will automatically connect to that registry and open the data source manager. You can change registries by clicking the *Change Data* button.

4.1.9.1. Using Multiple Data Sources

Once you select the data source you want to use for a report or chart and click the *Next* button, the next screen will display first twenty records from the data source (With the exception of data views, which will require you to select fields and set conditions first). You can display all of the records by checking the *Show All Records* box.

ERES allows you to construct reports and charts from multiple data sources. Once you have selected your first source, and click *Next* from the data table window, you will be presented with a dialog asking if you would like to process the current data or select another data source.

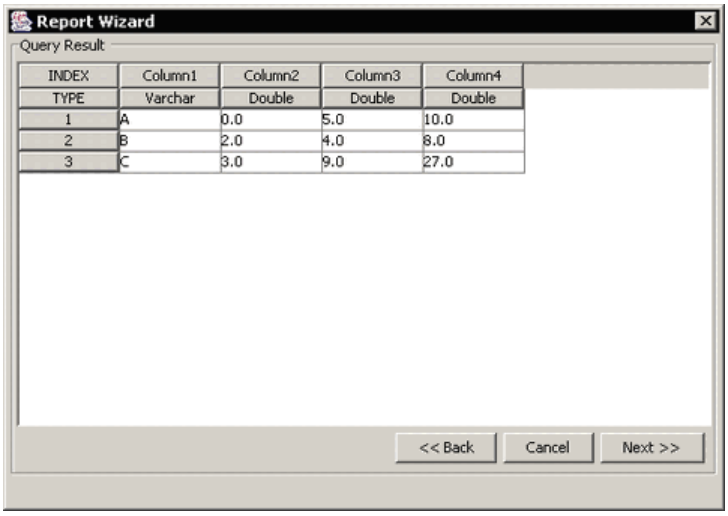


Additional Data Source Dialog

If you select *Process Data* and click the *Next* button, you will continue to the next step in the report or chart wizard. If you select *Get Other Data Source*, you will return to the Data Source Manager to select another data source for the report. You can repeat this process to select as many sources as you want.

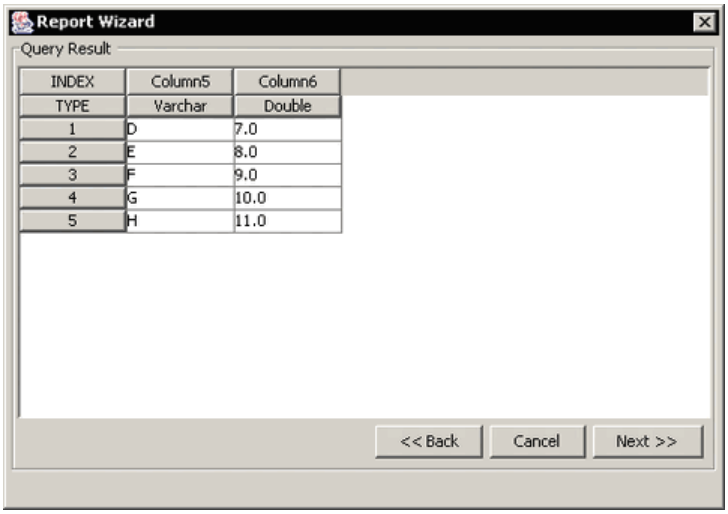
Multiple data sources are combined sideways in the data table. This means that the columns from the second (or third, fourth, etc.) data sources are placed to the right next to the columns from the first data source. If the columns from one data source have more rows than the other ones, null values will be placed in extra rows.

For example, assuming you want to use two data sources to create a report. The data table generated by the first source will look like this:



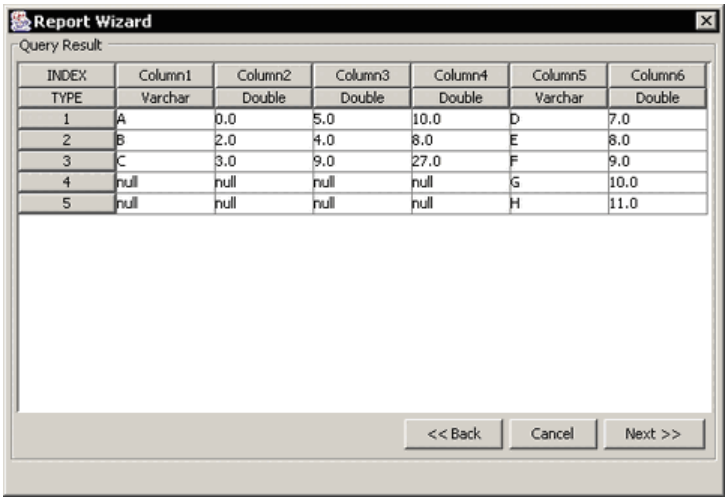
Data From the First Data Source

The data from the second source will look like this:



Data From the Second Data Source

When you combine two data sources, you will get the following data table:



Data From Combined Sources

As you can see, the two data sources have been placed side by side. Since the second source has more rows of data than the first, additional rows of null data were added.



Note

You cannot use parameterized data sources to create a multiple data source.

You can also use multiple data sources in reports by adding sub-reports. For more information about sub-reports, please see Section 5.1.9 - Sub-Reports.

4.1.9.2. Change Data Source

At any point during report or chart design, you can select to change the template's data source. Since report and chart templates are saved with their data source information, you must use the option within Report Designer or Chart Designer to change the template's data source. Simply altering a data source in the registry will not affect the template unless you use the data source updating feature (For more information about this, see Section 4.1.9 - Using Data for Charts and Reports). To change a template's data source, select *Change Data Source* from the Data menu, or click the *Change Data Source* button on the toolbar. This will bring up the Data Source Manager allowing you to select a new data source or modify an existing one.

If the new data source is significantly different from the previous data source in number of columns or data types, you might need to re-map the data to the report or chart. This is necessary because some layouts or computations may no longer work correctly with major changes to the structure of the data. For more information about data mapping, please see Section 5.1.2 - Report Types and Data Mapping for reports and Section 5.2.3 - Chart Types and Data Mapping for charts.

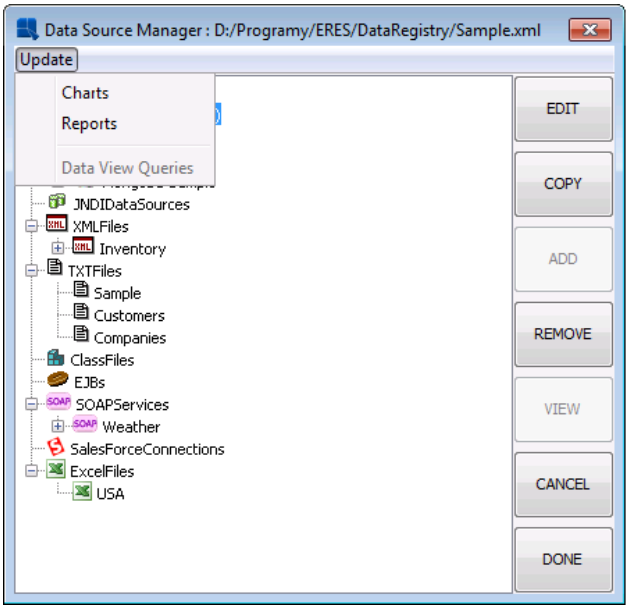
4.1.10. Data Source Updating

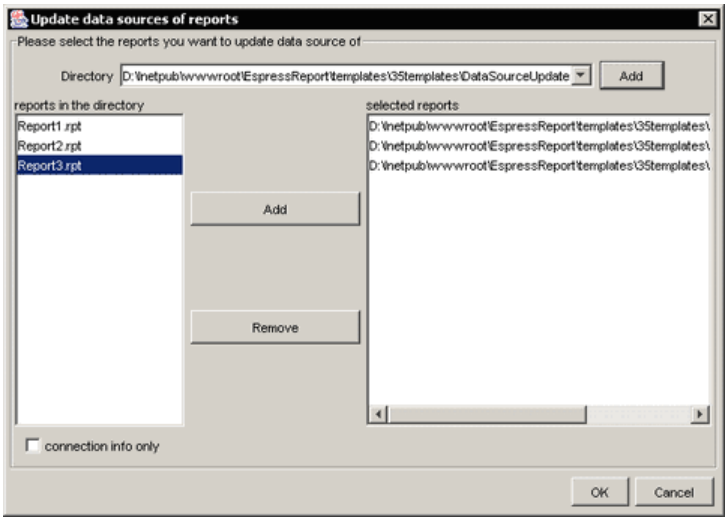
There are many circumstances where you will want to move a group of templates or a complete installation of ERES from one location to another. For example an application can move from a development to a production environment. In each environment, the location and connection information for data sources may be different. In this scenario, updating report templates one by one as detailed in the previous section isn't a feasible way to change the connection information for a large number of templates. Instead, ERES allows you to quickly update a group of templates based on information in the data registry.

Although report and chart templates maintain an internal copy of the data source information (allowing them to be deployed independently), they also maintain information (location & source) about the data registry from which they were created. Hence, when you modify a report or chart's query (as detailed in Section 4.1.2.4 - Editing Queries), you have the option to save the changes back to the data registry. In order to use this feature, you will have to keep your data registry up to date. This means that query changes should be saved back to the registry and changes in data view structure should be propagated to data view queries.

To use this feature, first make any modifications to the data registry that you want to propagate to the templates. These modifications can include database connection information, file locations for text, XML data files, and even changes to data views or queries that you want to pass to the templates.

From the data source manager, select *Reports* or *Charts* from the *Update* menu depending on whether you want to update charts or reports. This will bring up a dialog allowing you to select which files you want to update.

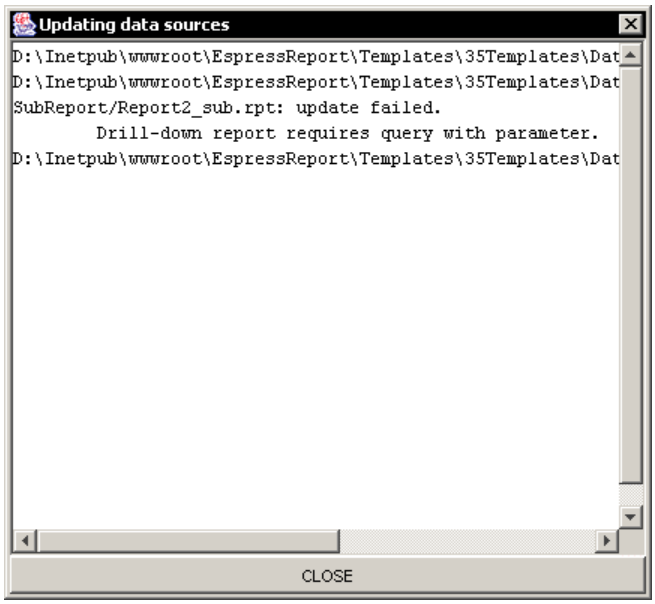




Select Reports for Data Source Updating

To select files to update, first browse to the directory that contains the reports/charts. After you select a directory, any report/chart templates will appear in the left side of the dialog You can select any templates you want like to update and click the *Add* button. You can navigate to as many different directories as you want to select templates. Selecting the *Connection Info Only* option will only update the connection information (database URL, driver, username, password, and locations for XML files, text files and Java classes). Queries and data view information will not be updated in the templates.

Once you finish selecting the templates you want to update, click *OK* and the updating process will begin. A dialog will display the current progress and any errors.



Updating Data Sources Progress/Log Dialog

A log file named `UpdateDataSources` will also be written in the root directory of the installation with the contents of the progress screen. Reports that fail the updating for various reasons can be updated manually using the option in Report Designer or Chart Designer.



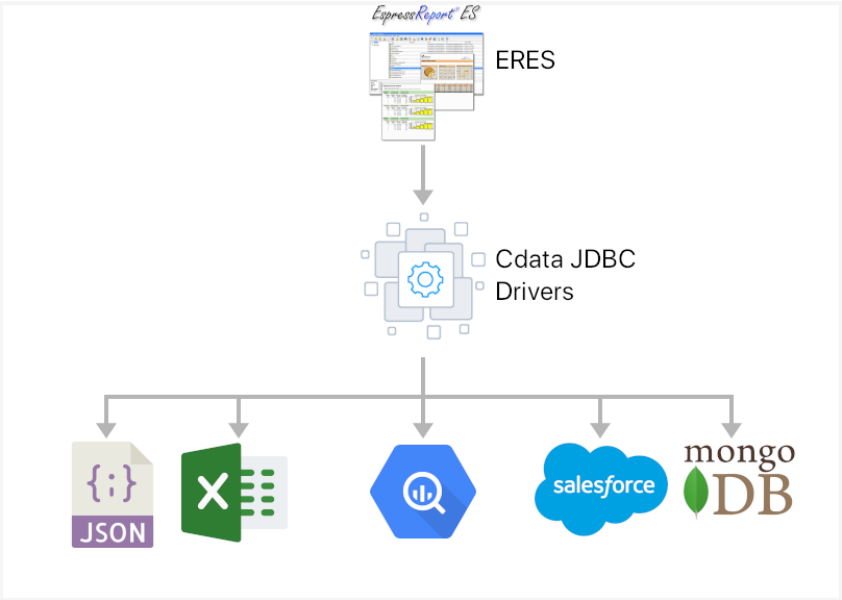
Note

Only reports or charts for the current data registry will be modified. If you select templates that do not retrieve their data from a source in the current registry, they will be ignored.

4.1.11. CData JDBC drivers

If you want to connect to a data source driver that doesn't ship with ERES Cloud, CData JDBC drivers are an interesting option.

CData JDBC drivers can be also used to increase capabilities of non-database data sources like Excel or JSON.





Tip

Although CData JDBC drivers are a 3rd party commercial product, you can install a free trial version to test the product before purchasing.



Note

CData JDBC drivers are a 3rd party product. If you want to use the drivers, you have to purchase them at https://www.cdata.com

4.1.11.1. Supported CData Drivers

Currently, we support the following CData JDBC drivers:

- Salesforce
- BigQuery
- Excel
- JSON
- MongoDB
- Kintone

Other CData JDBC drivers might work too but we can not guarantee full functionality for unsupported drivers. If you require a connection to a driver that is not listed here, please contact us at <support@quadbase.com>

4.1.11.1.1. Excel

Download: https://www.cdata.com/drivers/excel/jdbc

Official CData Documentation (for the JDBC driver only): https://cdn.cdata.com/help/RXF/jdbc

After downloading the driver, see the following chapters: Section 4.1.11.2 - CData JDBC driver installation, Section 4.1.11.3 - Deploying the CData JDBC Driver in ERES Cloud and Section 4.1.11.4 - Using the CData JDBC drivers in DataSource Manager

The CData Excel driver allows you to run SQL queries on top of an Excel file just like it was a database. This allows you to use the QueryBuilder, build Data Views in a graphical user interface as well as write SQL queries manually (in Query Builder).



Tip

You can also write queries with parameters and multi-value parameters using Excel files as the data source.

Since Excel does not have data types set for columns as a normal database has, determining the data type for the columns can be a bit tricky at times.By default, we added the following parameters to the CData Excel driver connection URL:

```
TypeDetectionScheme=RowScan;
RowScanDepth=10;
```

This makes the CData Excel driver scan the first ten rows when loading the selected Excel file and detect the data source for each column automatically based on the data in the first ten rows.

However, there are multiple options of how to determine the data type for each column.

For more options, read the CData JDBC Excel driver documentation about the TypeDetectionScheme parameter: https://cdn.cdata.com/help/RXH/jdbc/RSBExcel_p_TypeDetectionScheme.htm

You can change the TypeDetectionScheme value in the Setup Database… dialog in Data Source Manager in the URL field (the field where you entered the Excel file path).

4.1.11.1.2. JSON

Download: https://www.cdata.com/drivers/json/jdbc

Official CData Documentation (for the JDBC driver only): https://cdn.cdata.com/help/DJF/jdbc

After downloading the driver, see the following chapters: Section 4.1.11.2 - CData JDBC driver installation, Section 4.1.11.3 - Deploying the CData JDBC Driver in ERES Cloud and Section 4.1.11.4 - Using the CData JDBC drivers in DataSource Manager

The CData json driver allows you to run SQL queries on top of a json file just like it was a database. This allows you to use the QueryBuilder, build Data Views in a graphical user interface as well as write SQL queries manually (in Query Builder).



Tip

You can also write queries with parameters and multi-value parameters using JSON files as the data source.

4.1.11.1.3. Salesforce

To connect to Salesforce:

Download: https://www.cdata.com/drivers/salesforce/jdbc

Official CData Documentation (for the JDBC driver only): https://cdn.cdata.com/help/RFF/jdbc

After downloading the driver, see the following chapters: Section 4.1.11.2 - CData JDBC driver installation, Section 4.1.11.3 - Deploying the CData JDBC Driver in ERES Cloud and Section 4.1.11.4 - Using the CData JDBC drivers in DataSource Manager

4.1.11.1.4. MongoDB

To connect to MongoDB:

Download: https://www.cdata.com/drivers/mongodb/jdbc

A bug fix was added to newer MongoDB driver from https://cdatabuilds.s3.amazonaws.com/support/DGR-JV_8598.exe

Official CData Documentation (for the JDBC driver only): https://cdn.cdata.com/help/DGF/jdbc

After downloading the driver, see the following chapters: Section 4.1.11.2 - CData JDBC driver installation, Section 4.1.11.3 - Deploying the CData JDBC Driver in ERES Cloud and Section 4.1.11.4 - Using the CData JDBC drivers in DataSource Manager

4.1.11.1.5. BigQuery

To connect to BigQuery:

Download: https://www.cdata.com/drivers/bigquery/jdbc

Official CData Documentation (for the JDBC driver only): https://cdn.cdata.com/help/DBF/jdbc

After downloading the driver, see the following chapters: Section 4.1.11.2 - CData JDBC driver installation, Section 4.1.11.3 - Deploying the CData JDBC Driver in ERES Cloud and Section 4.1.11.4 - Using the CData JDBC drivers in DataSource Manager

4.1.11.1.6. Kintone

To connect to Kintone:

Download: https://www.cdata.com/drivers/kintone/jdbc

Official CData Documentation (for the JDBC driver only): https://cdn.cdata.com/help/DBF/jdbc

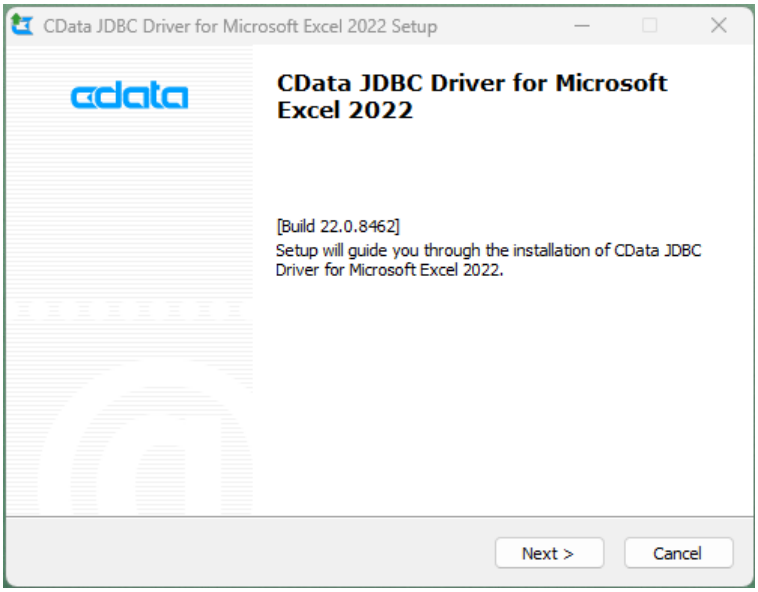
After downloading the driver, see the following chapters: Section 4.1.11.2 - CData JDBC driver installation, Section 4.1.11.3 - Deploying the CData JDBC Driver in ERES Cloud and Section 4.1.11.4 - Using the CData JDBC drivers in DataSource Manager

4.1.11.2. CData JDBC driver installation

The installation process of all CData JDBC drivers is basically the same for all the supported data sources. We'll guide you through the process of installing the CData Excel JDBC driver for example.

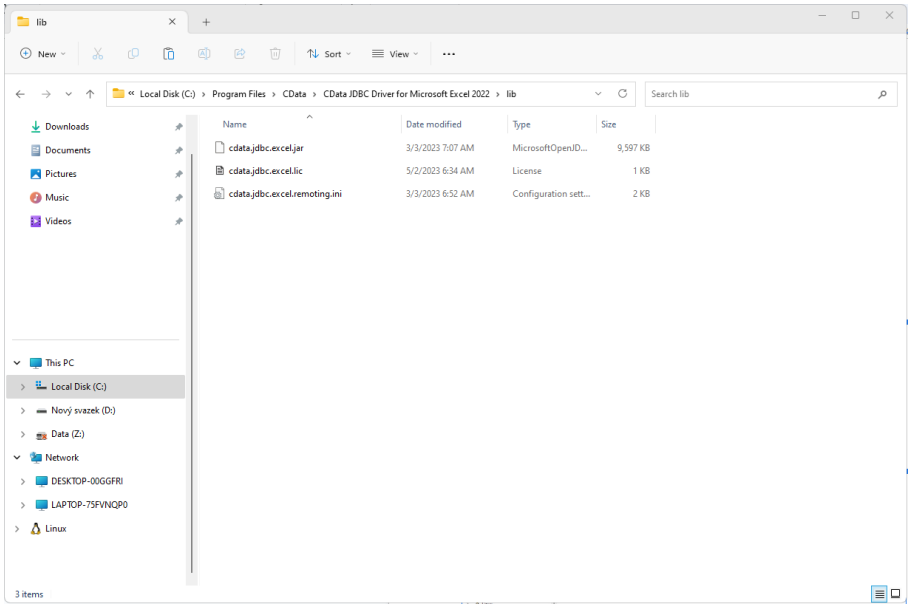
First, install the CData JDBC Driver of your choice. You will find the links to each supported CData JDBC driver in the chapters below.

After downloading the installer, proceed with the setup according to the dialogs.



In the following dialog, you can choose to either install a paid version of the product or a free 30-day trial version. In this example, we'll use the trial version.

When the setup is complete, JDBC Driver will be created in "CData JDBC Driver for Microsoft Excel 2022\lib".





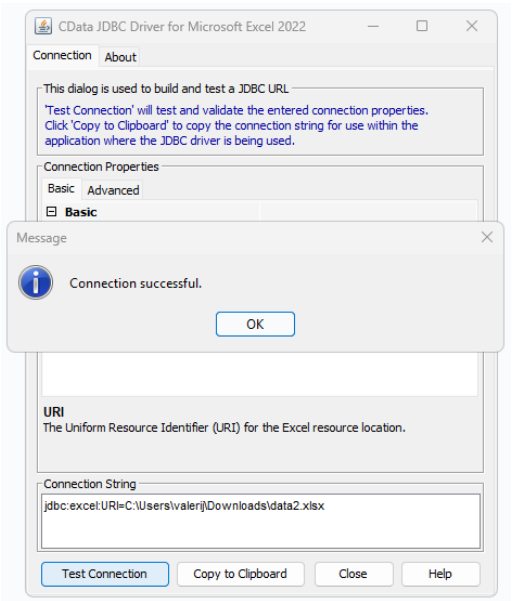
Tip

Executing "cdata.jdbc.excel.jar" will launch the connection test tool. You can use this tool to test or troubleshoot the CData driver.



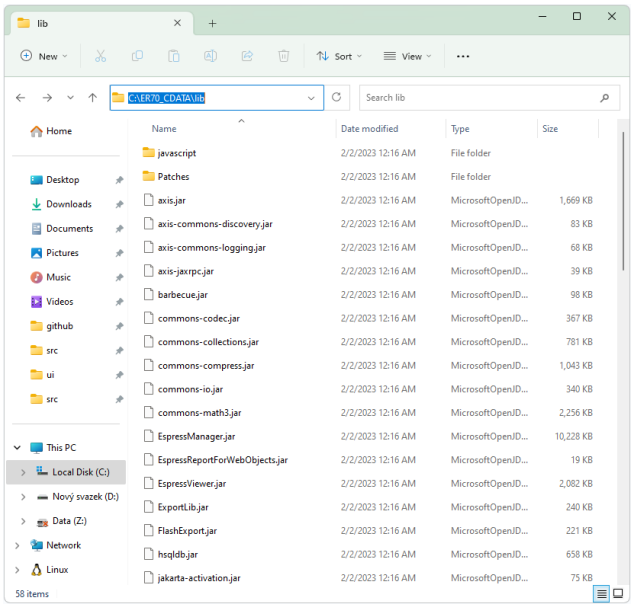
Tip

The connection test tool also displays you the connection URL that can be used in ERES Cloud Data Source Manager.



4.1.11.3. Deploying the CData JDBC Driver in ERES Cloud

Locate the "CData Installation Directory"\CData JDBC Driver for Microsoft Excel 2023\lib (for example: C:\Program Files\CData\CData JDBC Driver for Microsoft Excel 2023\lib) on your hard drive.The directory should contain three files: cdata.jdbc.excel.jar, cdata.jdbc.excel.lic, cdata.jdbc.excel.remoting

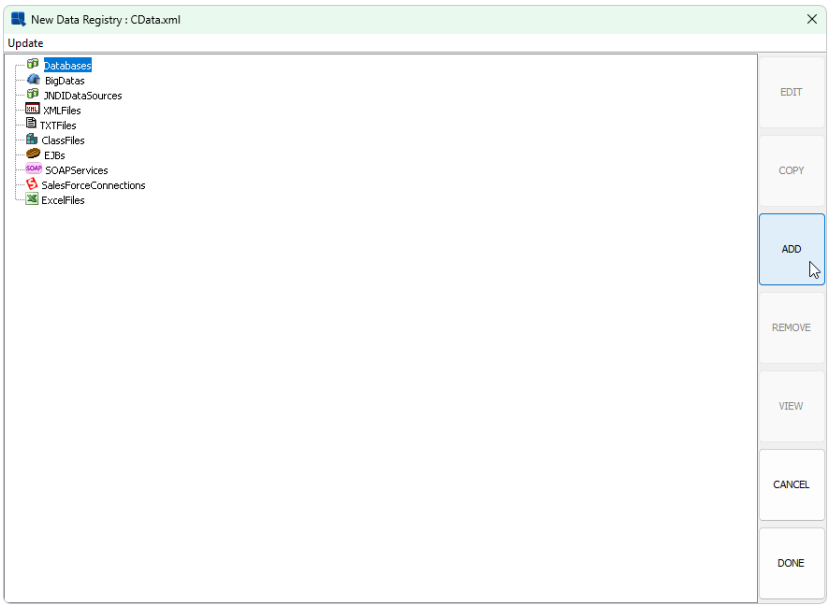


After you've copied the files, restart The EspressManager (if it is running).

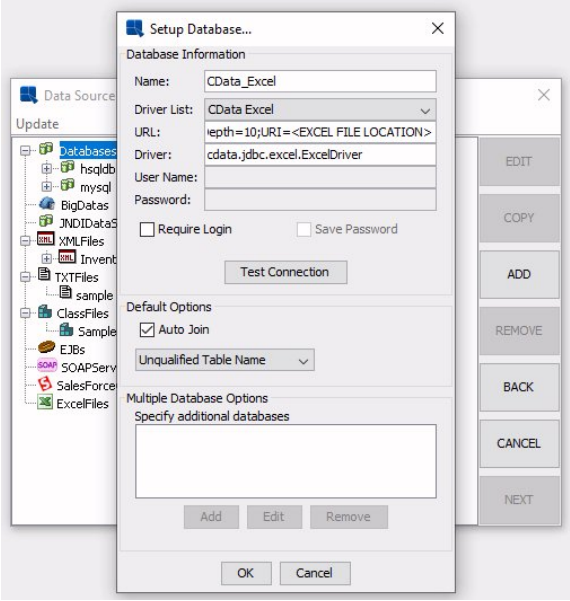
4.1.11.4. Using the CData JDBC drivers in DataSource Manager

Launch the DataSource Manager and open a data registry in it (existing one or a new one).

Select the “Databases” option in the tree-list and press the “ADD” Button.



In the “Driver List:” drop-down menu, select “CData Excel” (or any other CData JDBC driver you might be installing).



In the “URL:” text field, replace the placeholders (like "EXCEL FILE LOCATION") with real values.



Tip

Alternatively, you can replace the text in the “URL:” text field with the connection string obtained by the CData Test Connection tool described in the previous chapter.

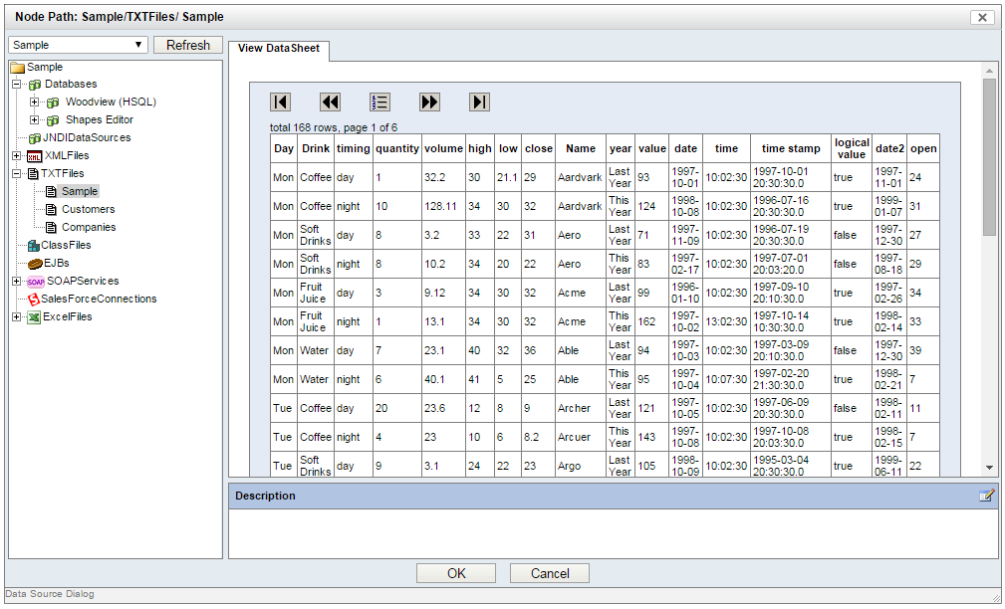
Click OK. You’re done. You can start using Excel files as if they were a database.

4.2. Data in QuickDesigners and Maps

4.2.1. Select a Data Source

QuickDesigner Reports, QuickDesigner Charts, Online Maps and SVG Maps use the same interface to manage data - *Data Source Dialog* . The *Data Source Dialog* allows end users to select, filter, and present data without mastering database structures, and all with zero client download. For full control of data sources, please see Section 4.1 - Data in Organizer.

The *Data Source Dialog* prompts you to select the data source that you would like to use.



Data Source Dialog

Once a data source is selected, you can see records of the source in the right-hand side *View DataSheet* pane. If you select a DataView as the data source, you will be able to create a new DataView query in the *DataView Builder* in the right pane. You can also select an existing DataView query and modify it in the right pane.

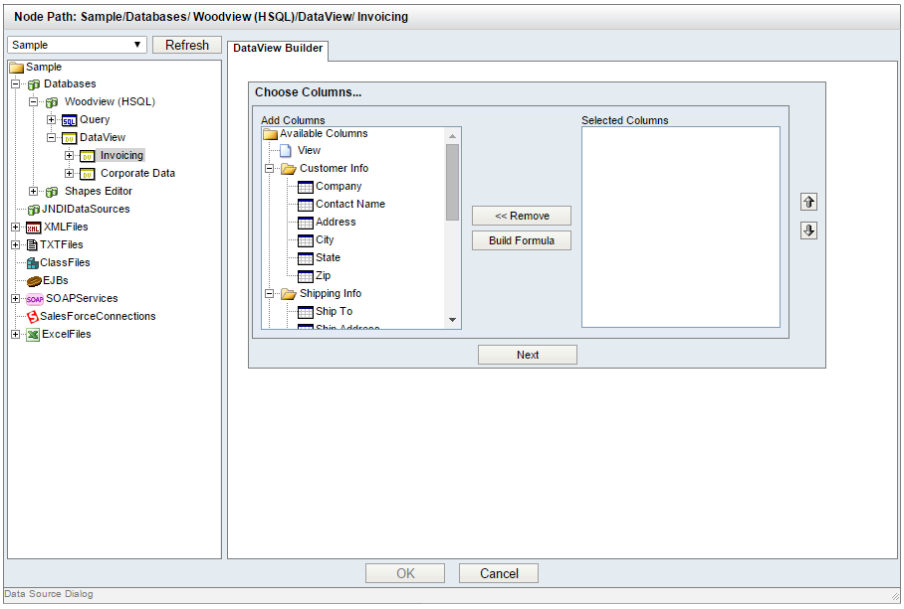
4.2.2. Queries

Queries work in conjunction with Data Views. Data Views allow an administrator to pre-select tables and fields from a database and create a local view from which end users can run queries. For more information about data views and setup, please see Section 4.1.2.3 - Data Views.

If you select either a data view or a data view query as the data source, you will be taken to *Data Source Dialog* query interface in the right-hand pane where you can create or modify a query based on the data view.

4.2.2.1. Select Fields

If you select a data view as the data source, the *DataView Builder* will open in the right pane allowing you to select fields from the view for your query.

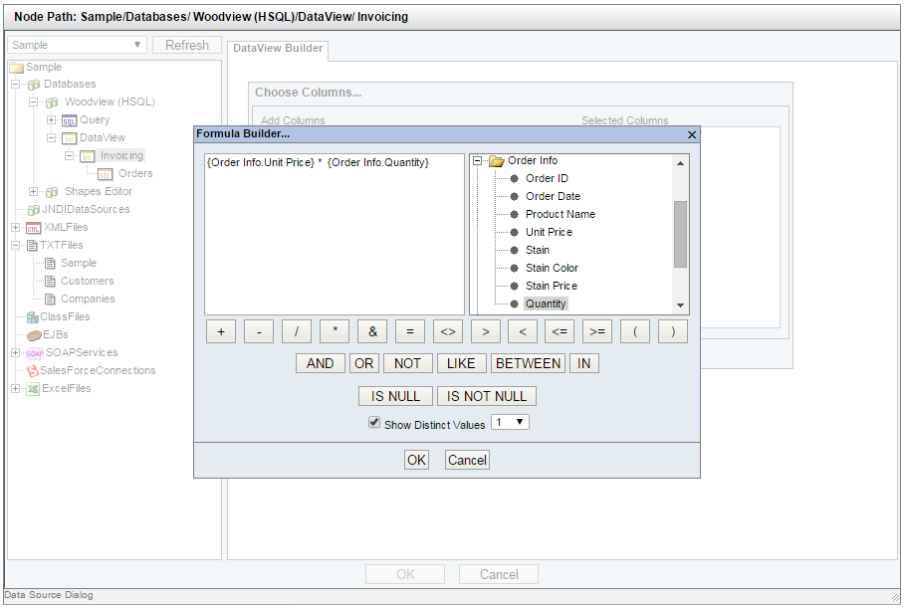


Select Columns Dialog

A select list for each heading in the view is generated on the left-hand side of the *DataView Builder*. To select fields for the query, click on a field on the left-hand side. The selected column will be added to the dialog on the right-hand side. To remove fields, you can select the columns from the right-hand side and click the *Remove* button.

4.2.2.1.1. Build Columns

You can also create computed columns from the *Choose Columns* dialog. To create a computed column, click the *Build Formula* button. This will launch the Formula Builder interface in a new window.

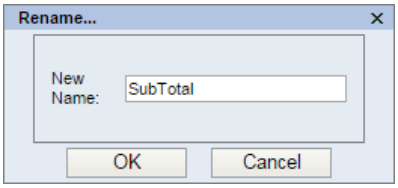


Formula Builder Window

The right-hand side of the Formula Builder contains a list of fields in the view, as well as the database functions. The buttons at the bottom are available operators and conditionals. The Boolean and conditional operators are more commonly used for building conditions which is covered in Section 4.2.2.2 - Set Conditions.

Once you have finished constructing the field, click *OK* to close the Formula Builder and return to the *Choose Columns* window where the computed field is added to the *Selected Columns* list. You can edit the formula by selecting it in the list and clicking the *Modify Formula* button.

You can also create an alias for any expression column by selecting it and clicking the *Rename* button. This will bring up a new window for you to specify the alias.

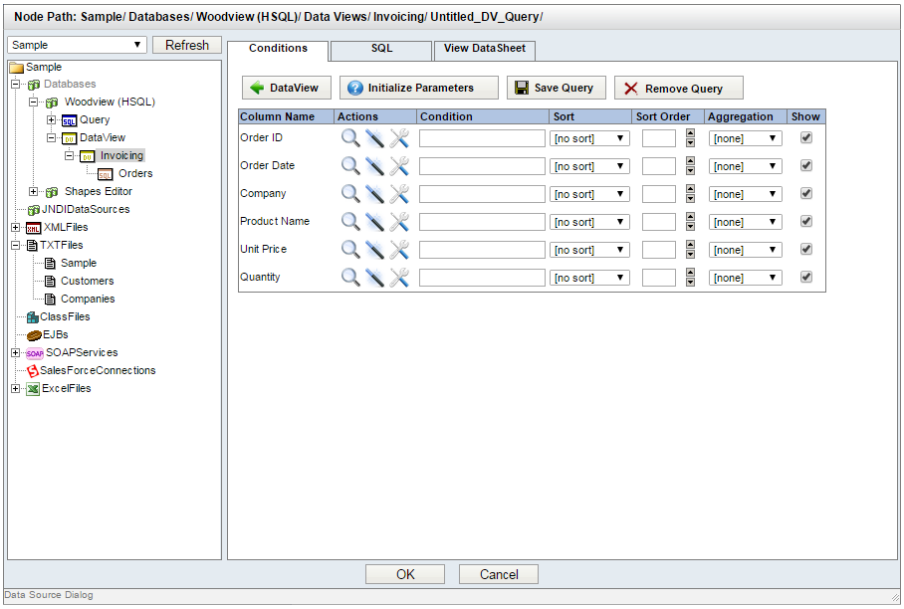


Rename Dialog

Once you have finished selecting the columns you would like to use, click the *Next* button to set conditions.

4.2.2.2. Set Conditions

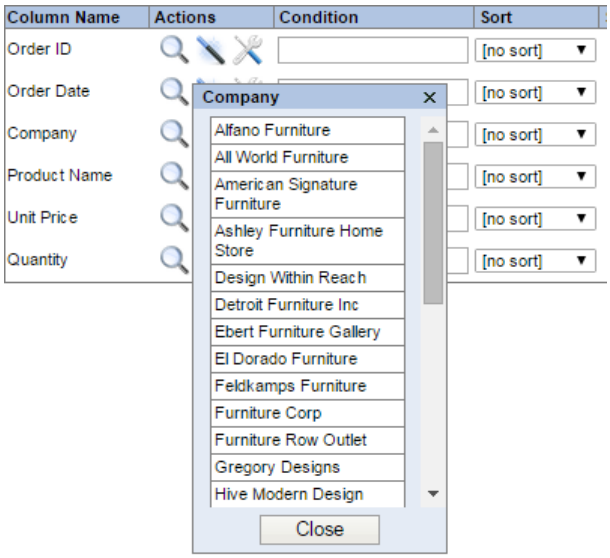
After selecting fields, or if you select a data view query as the data source, the conditions window opens allowing you to set aggregation, grouping, and filtering for the selected columns.



Conditions Dialog

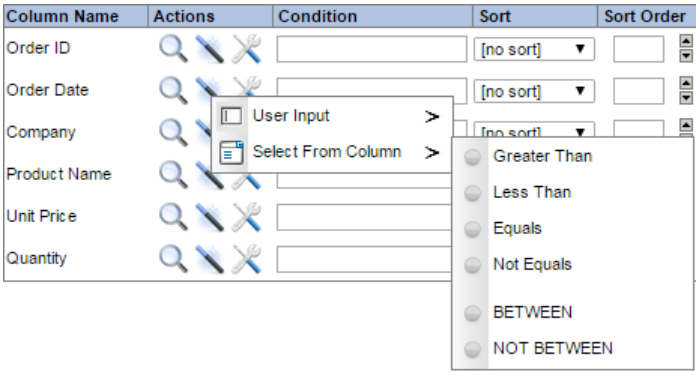
Each selected column is listed along with several options:

Actions: **Show Column Data:** This will bring up a new window with column data.



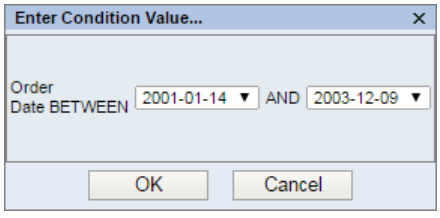
Show Column Data

Condition Wizard: This will bring up a dialog allowing you to set basic conditions from a field. The first dialog allows you to select whether you would like to enter values or select them from a list.



Condition Wizard

Once you have selected the condition, a new dialog opens allowing you to set the values for the condition.



Condition Value Dialog

Select the values that you would like and click *OK*. The condition will automatically be added to the *Condition* field.

Build Condition: This will bring back up the Formula Builder, allowing you to construct an expression for the condition.

Condition: This allows you to type a condition for the column.

Sort: This allows you to indicate if the column should be sorted. You can select whether to sort in ascending or descending order.

Sort Order: This allows you to specify the priority (order) for sorting columns.

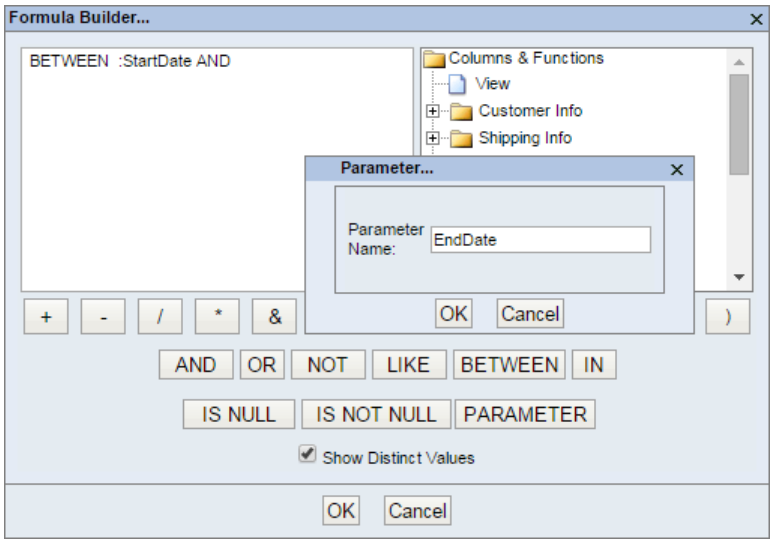
Aggregation: This allows you to specify any grouping or aggregation in the query.

Show: This option indicates whether the column should be visible or not.

4.2.2.2.1. Parameterized Queries

Data Source Dialog also allows you to define query parameters. Parameters are defined in the same way as with other query interfaces - using the `":"` character. You can define parameters directly in the conditions field, using the convention of `:ParameterName` to define a parameter.

Parameters can also be defined in the Formula Builder when building conditions. To define a parameter this way, click the  *Build Condition* button to open the Formula Builder, click the *Parameter* button to insert a parameter. This will bring up a new dialog allowing you to enter the parameter name.

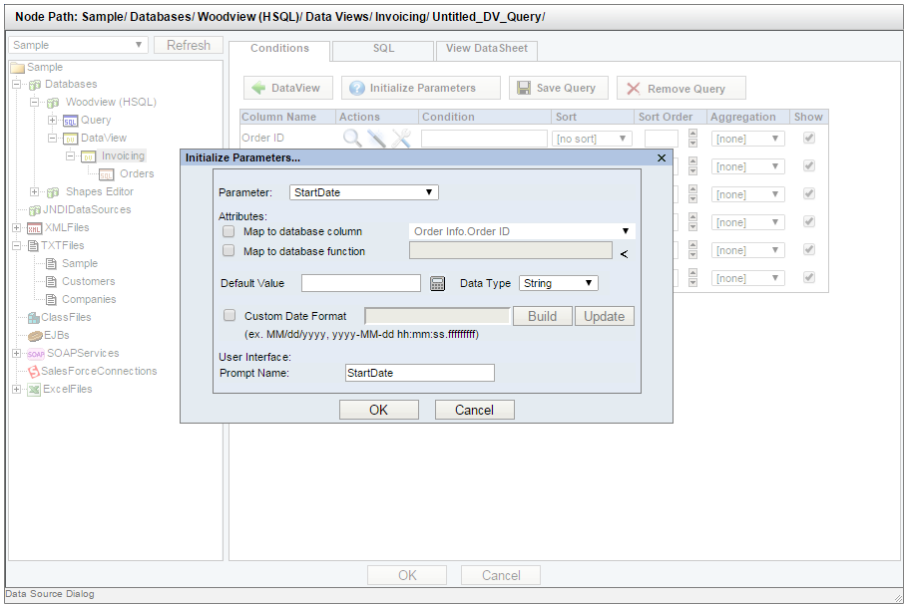


Specifying a Parameter in Formula Builder

For more information about query parameters, please see Section 4.1.2.2.2 - Parameterized Queries.

4.2.2.2.1.1. Initialize Query Parameters

Once you have specified parameters in the query, they will need to be initialized. You can initialize parameters by clicking the *Initialize Parameters* button. The initialization dialog will also open if you preview the query by clicking on the *View DataSheet* tab or if you click on the *OK* button closing *Data Source Dialog*.

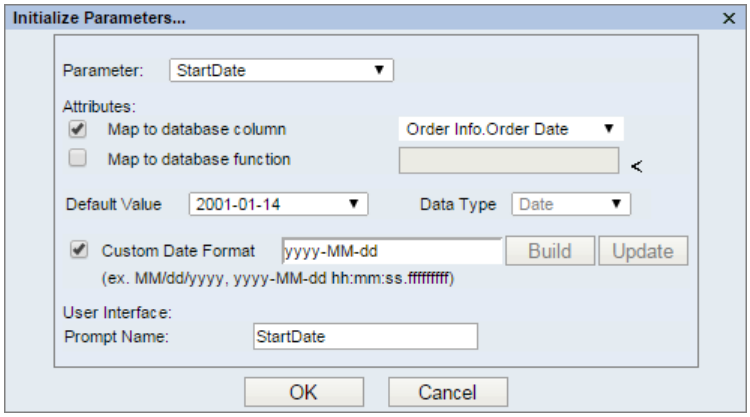


Initialize Parameter Dialog

The drop-down list at the top of the dialog indicates which parameter is currently being initialized. You can access and set the options for each parameter by selecting it from this list.

The following options are available for each parameter:

Map to database column: This allows you to specify a column from the database whose values will be used for the parameter input. Selecting this option modifies the parameter prompt that you will get when running the report or chart. If you map the parameter to a database column, you will be prompted with a drop-down list of distinct values from which to select a parameter value. If you do not map, you will have to type in the specific parameter value.

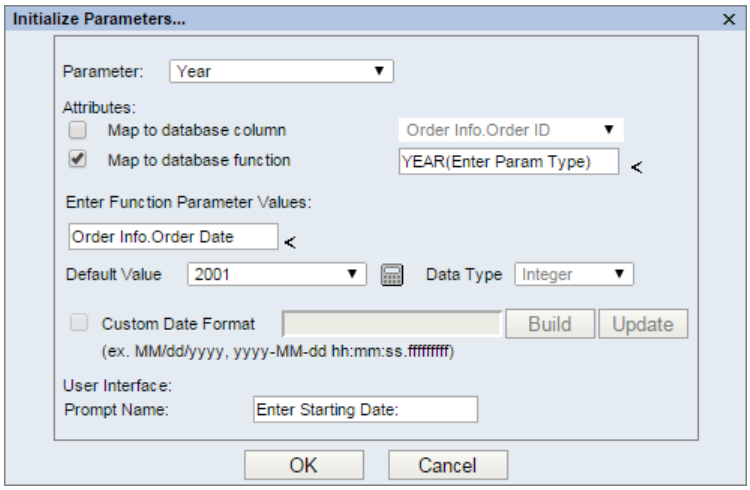


Initialize Parameter Dialog

Map to database function: Mapping a parameter to a database function is very similar to mapping to a column. This feature allows you to map to the results of any function from your database. For example, suppose you have the following condition in your query:

```
Year({Orders.OrderDate}) IN (:OrderYear)
```

In the initialize parameter dialog, check the *Map to database function* box, select the function from the list and enter the parameter value for your function. The default value and data type will be automatically updated



Initialize Parameter Dialog

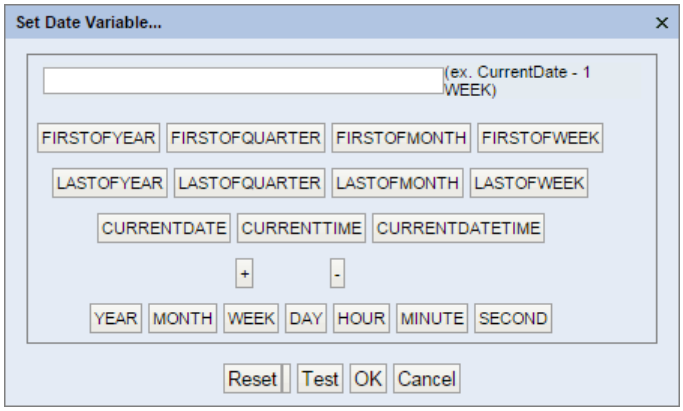
For certain databases, you may notice that the functions available do not specify the parameter types for functions. In this case, there will only be one function parameter field and it will be longer than normal. Selecting a column from the drop down list will be appended at the end of the field rather than replacing the existing text. You must enter all parameter values for the selected function, separating values with commas. Once you click away from this field, the default value field will be filled with function results and the data type will be automatically determined.

Default Value:

This allows you to specify a default value for the parameter. Although you do not have to specify a default value, it is recommended that you do so.

Date Variable:

This option is only available when the parameter is not mapped to a database column or function, and is only intended for parameters with variable type date/time. When you click this button, the following panel will pop up, listing all the supported keywords.



Enter Date Variable Dialog

This dialog allows you to select one of the three keywords: *CurrentDate*, *CurrentTime*, and *CurrentDateTime*. You may add or subtract units of time from the current date/time, allowing you to have a dynamic date range. For example, a report may have the following default values:

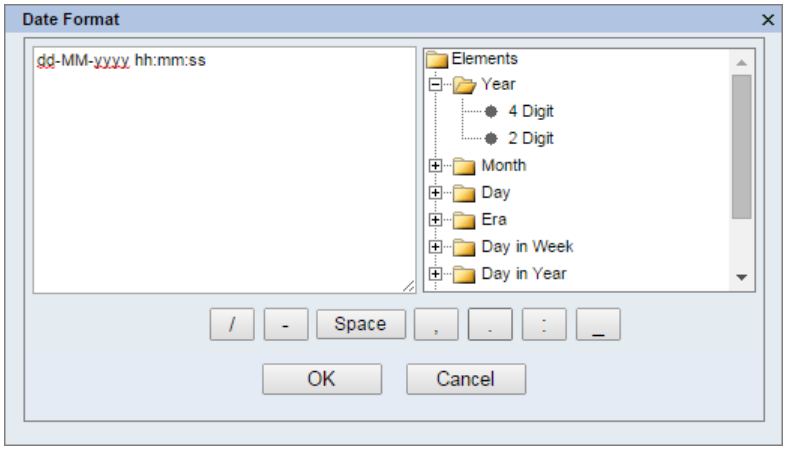
```
StartDate: CurrentDate - 1 WEEK
EndDate: CurrentDate
```

This would indicate that every time the report is run, the default prompt should be one week ago to the current date. Other supported measures are *YEAR*, *MONTH*, *DAY*, *HOUR*, *MINUTE*, and *SECOND*. This feature only supports a single addition or subtraction. This feature does not support multi-value parameters. Note: QuickDesigner Reports date values (in the right-hand panel) correspond to the variable values entered here, not the keywords themselves as in Report Designer.

You can also use functions to define the parameter value.

Data Type: This allows you to specify the data type for the parameter value(s). If you have mapped the parameter to a column,the data type is set automatically.

Custom Date Format: This allows you to set the format in which the date parameter should be entered. This option is only available when you have mapped the parameter to a date/time column or when you not mapped the parameter at all and the parameter data type is `date`, `time`, or `timestamp`. When you check this option you can specify a date format by either using the builder interface or by entering the date format in a combination of characters that represent time elements.



Date Format Dialog

If you click the *Build* button, the above dialog will appear. The date/time representations are listed on the right and optional spacers and symbols are shown as a collection of buttons on the bottom. Once you have finished creating the format, click *OK* to save.

The date and/or time format is a series of characters and delimiters. Letters are used to represent different elements of date/time data. For more on the characters and their formatting, please refer to the documentation for the printDate() report function in Section 5.1.6.2.8.3 - Date/Time Functions.

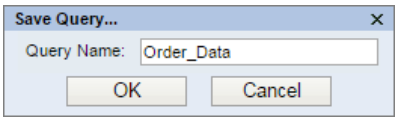
If the parameter is mapped to a column, once you have finished designing the format, click on the *update* button to refresh the parameter prompt with the new format.

Prompt Name: This allows you to specify the prompt that is given to the user in the parameter dialog.

Once you have finished setting values for all the parameters, click *Ok* to close the dialog and save the settings.

4.2.2.2.2. Save Queries

You can save the created query, or save the changes to the query if you selected a data view query to begin with by clicking the *Save Query* option in the conditions window. This will bring up a dialog prompting you to specify a name for the query.

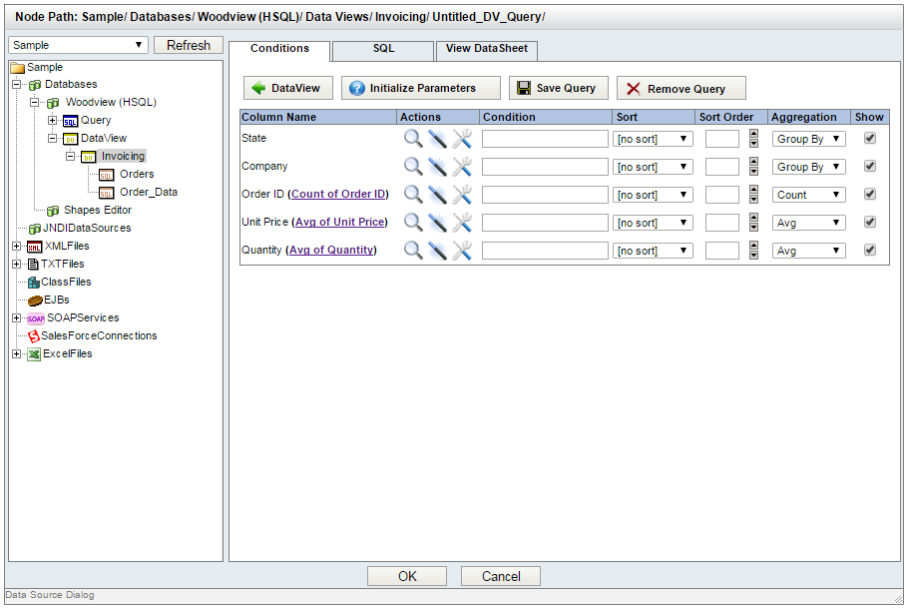


Save Query Dialog

Once you have specified the name, click *OK*. The query will be saved in the data registry as a new node under the data view.

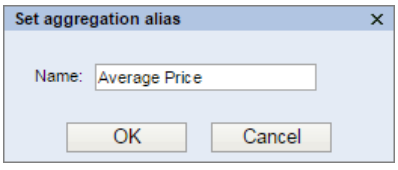
4.2.2.2.3. Aggregation Alias

If you add aggregation and group bys to the query, you are also able to specify an alias for the aggregated field.



Set Alias

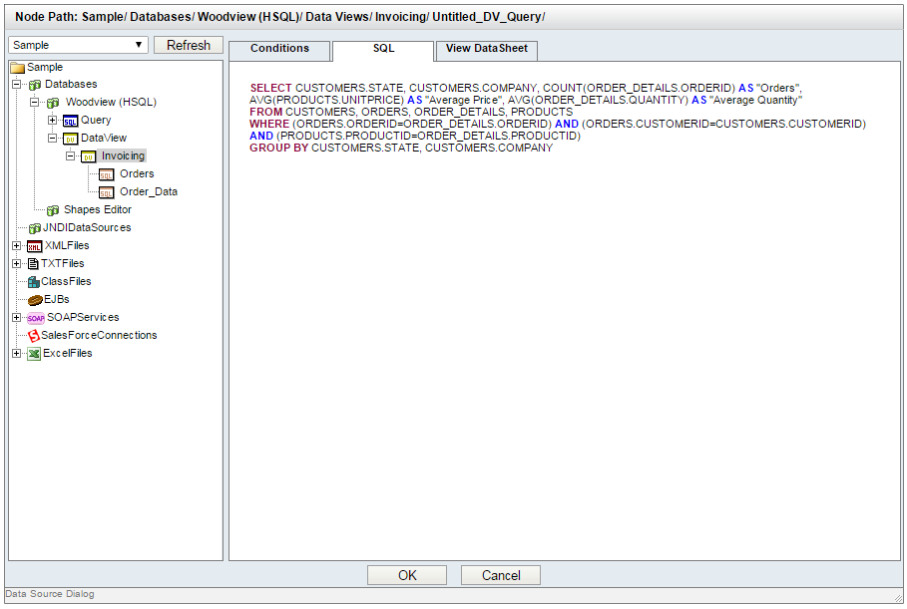
Clicking on the underlined option next to the column name will pop up a dialog allowing you to specify a name for the field. The name you specify here will replace the actual field name.



Alias Dialog

4.2.2.3. View SQL Query

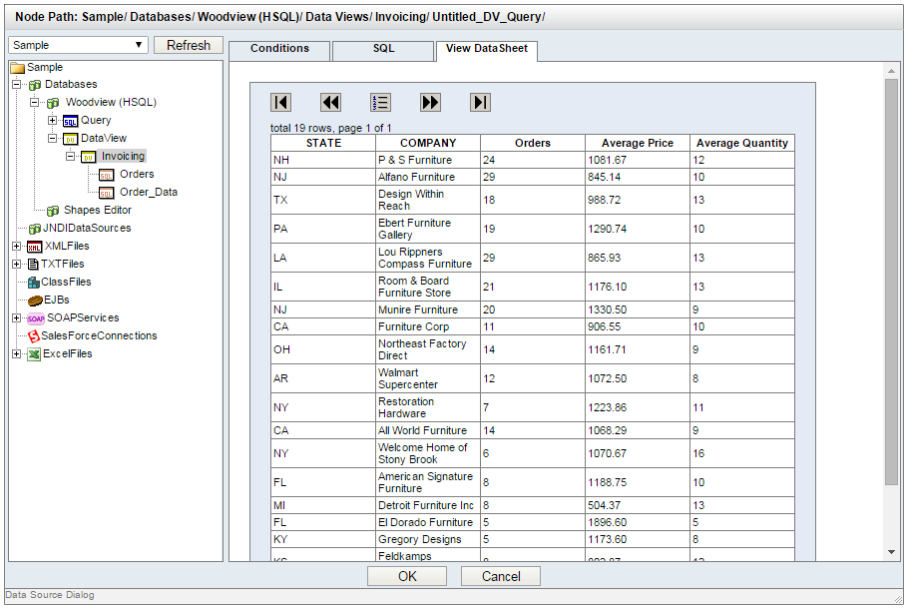
You can view the constructed query in the *SQL* tab of the *Data Source Dialog*.



SQL View

4.2.2.4. Preview Query Results

You can preview the query at anytime by clicking the *View DataSheet* tab in the *Data Source Dialog*.



Query Results

You can navigate through the result set using the icons at the top of the window.



Go to the first page of the data table.



Go to the previous page of the data table.



Set the number of records to display per page. This will open a new dialog prompting you to specify the record number.


Go to the next page of the data table.


Go to the last page of the data table.

Chapter 5. Designing Reports & Charts

5.1. Report Designer

5.1.1. Introduction to Report Designer

The Report Designer is a graphical user interface launched within the Organizer that allows users to create and customize reports. The simple drag and drop style interface and extensive editing/formatting capabilities makes the report design quick and easy.

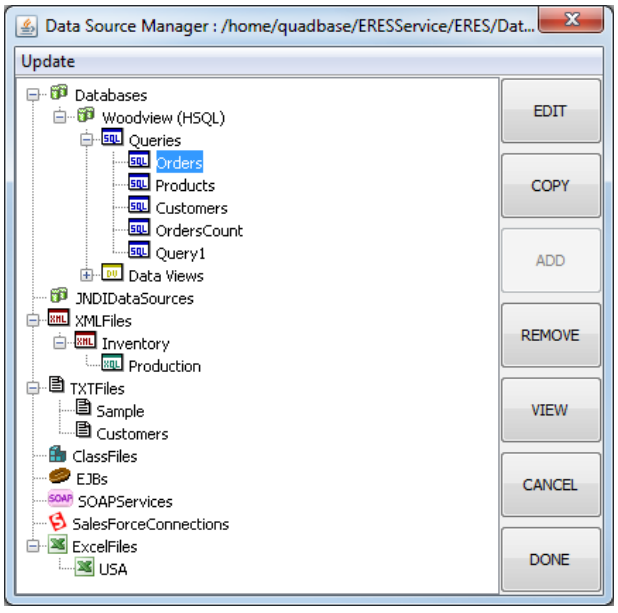
5.1.1.1. Starting Report Designer

The Report Designer interface is always loaded from within the Organizer. To start a new report, you can click the *Report Designer* button on the toolbar or select *Report Designer* from the *View* menu. You can also start Report Designer when creating or editing data sources in the registry. The *View* button allows you to preview a data source and select to build a chart or report. In addition, you can open the Report Designer to edit a report template file in the Organizer. To do this, first select the file that you would like to open, then select *Open File* from the *File* menu, or press **Ctrl**+**O**. You can also right click on the file and then select *Open File* from the pop-up menu, or simply double click the file.

5.1.1.2. Selecting a Data Source

The first step in designing a report is to select the data source from which the report is to be drawn. The first time you select to start a new report you will be prompted to select the data registry that you would like to use. If there are not currently any available registries (meaning that you have not created any or you do not have privileges to view any), you will be prompted to go the Data Registry Manager to create one. For more on data sources, please see Section 4.1 - Data in Organizer.

Once you have selected a registry, the Data Source Manager window will open allowing you to select, add, or modify a data source that you would like to use for the chart or report. Note that the registry will not open if you are building a report from the registry in the *Modify Data Sources* dialogs.



Data Source Manager Window for Report Designer

To use a particular data source, select the query, XML query, text file, data view, or data view query that you would like to use, and click the *NEXT* button. With the exception of data views, which will require you to select fields and set conditions first, the next screen will present the first twenty records from the data source. From this screen you can see all of the records returned by the selected source, by checking the *Show All Records* box.

5.1.2. Report Types and Data Mapping

ReportDesigner supports five basic report types: simple columnar, summary break, crosstab, master & details, and mailing label. Each report type has slightly different mapping options and can be used to generate a different style of report.

5.1.2.1. Simple Columnar Report

The simple columnar report is the most basic report type supported by ReportDesigner. It displays columnar data in a single table without any grouping or breaks.

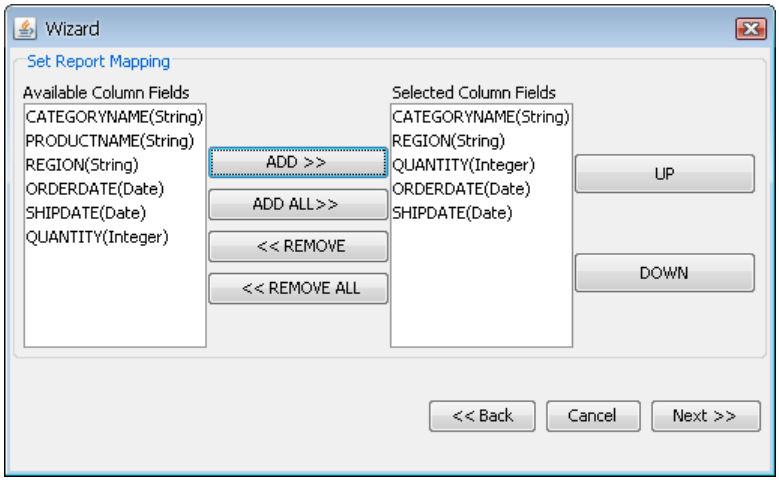


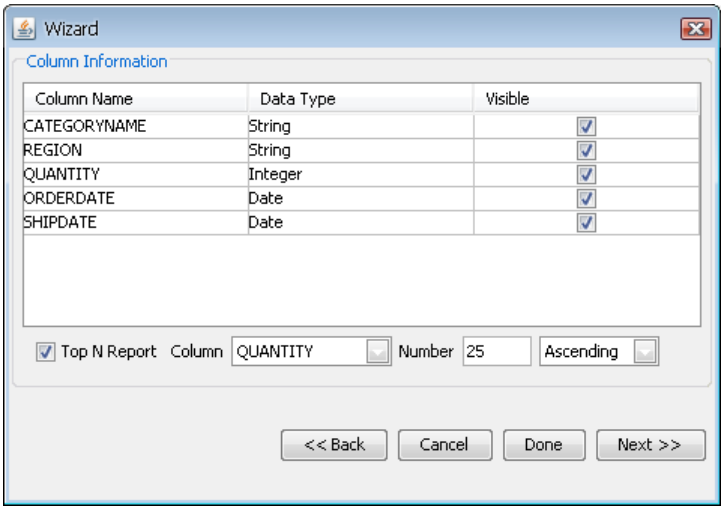
Simple Columnar Report

5.1.2.1.1. Data Mapping

The first step in data mapping for this type is to select the columns you want to include in your report from the *Set Report Mapping* window in the Report Wizard. The left side of the window lists all available columns from your query or data file and the right side lists columns that will be added to the report. Click on a column name to select it (**Shift**+**Click** for multiple selections) and press the *ADD* or *REMOVE* buttons to add it or to remove it from the report. The order in which you select columns in this window is the order in which they will appear in your report. You can change the order of the selected columns by highlighting the column you want to move and clicking the *UP* or *DOWN* buttons. Fields can also be moved by clicking and dragging the selected field.

The second step for report mapping is to set the column options. For simple columnar reports there is only one available option. The window displays a table showing all the columns you have selected for the report and all available options. The first field displays column names, the second one displays data type, and the third one displays a check box called *Visible*. Checking or un-checking this box will make the column visible or invisible within the body of the report (All columns are visible by default).





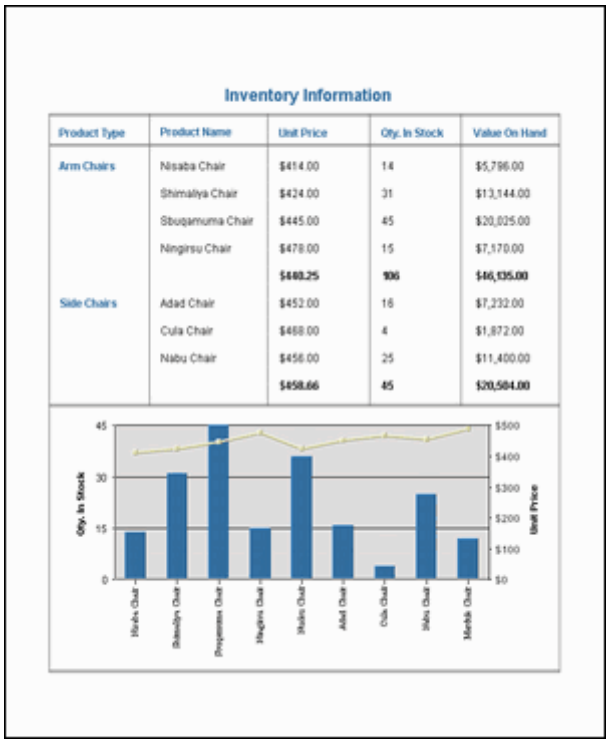
Simple Columnar Report Mapping

5.1.2.1.1.1. Top N Report

You can also create a Top N report using the columnar format. A Top N report will order and show you the highest set of values based on a particular column in the report. For example, you want to show top five customers based on total sales. To do this, check the *Top N Report* checkbox at the bottom of the data mapping window. You can then specify which column you want to use as the measure (for the above example, it would be total sales column), the number you want to retrieve (i.e. 10 or 20, etc), and whether you want to show the columns in ascending or descending order. The columnar report will then return the number of records specified sorted by the highest value for the specified column.

5.1.2.2. Summary Break Report

Like the simple columnar report, the summary break report type takes columnar data and presents it in a tabular form. However, unlike the simple columnar report, it allows you to break data into sections and insert summary fields.



Summary Break Report

5.1.2.2.1. Data Mapping

The first step in data mapping for this type is to select the columns you want to include in your report from the `Set Mapping` window in the Report Wizard. The left side of the window lists all available columns from your query or data file and the right side lists the columns that will be displayed in the report. Click on a column name to select it (**SHIFT**+**Click** for multiple selections) and press the *ADD* or *REMOVE* buttons to add it or remove it from the report. The order in which you select columns in this window is the order in which they will appear in your report. You can change the order of the selected columns by highlighting the column you want to move and clicking the *UP* or *DOWN* buttons. Fields can also be moved by clicking and dragging the selected field.

The second step in data mapping is to set column options. There are five column options available for summary break reports. The window displays a table showing all of the columns that you have selected for your report and all available options. The first field displays column names, the second one displays data type, and the third one displays a check box marked `Visible`. Checking or un-checking this box will make the column visible or invisible within the body of the report (All columns are visible by default). The fourth field is the row break field. Checking this box indicates that the report will break and insert column summaries every time the selected column field changes. By default, the first column selected for the report is the break field. The fifth field allows you to select aggregation. If the *Perform Column Aggregation* checkbox at the bottom of the window is checked, the selected aggregation will be performed on the entire column and the report will only contain summarized data. If the box is not checked, the report will display all data and only insert aggregations as summaries after each row break.

The drop-down menu allows you to select the aggregation operation to be performed for that column. The aggregation operations that can be performed are: `none`, `sum`, `maximum`, `minimum`, `count`, `average`, `first`, `last`, `sum of squares`, `variance`, `standard deviation`, and `count distinct`. If you select to perform column aggregation, you must select an aggregation option for each column. Columns selected as `row break` fields cannot be aggregated.

For example, suppose we have the following data table:

Order#	Product	Quantity
12	Chair	2
12	Table	3

Order#	Product	Quantity
14	Cabinet	2
14	Table	5

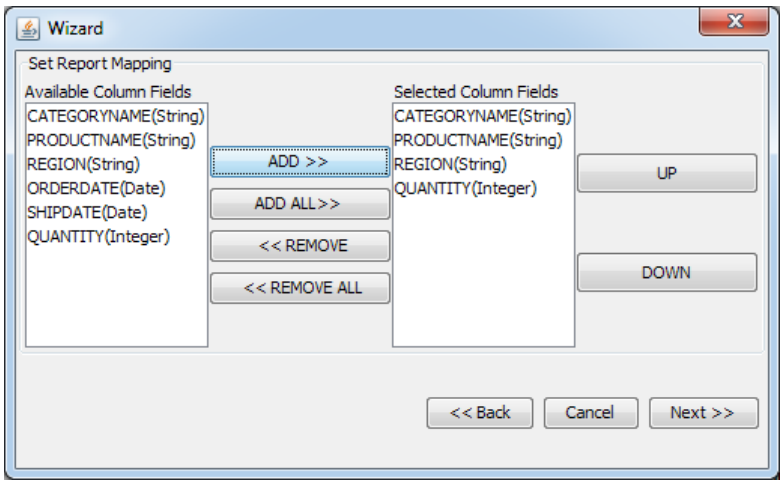
Setting `Order #` as the row break field and the aggregation on `Quantity` to `Sum` without checking *Perform Column Aggregation* will produce the following report:

Order#	Product	Quantity
12	Chair	2
	Table	3
		5
14	Cabinet	2
	Table	5
		7

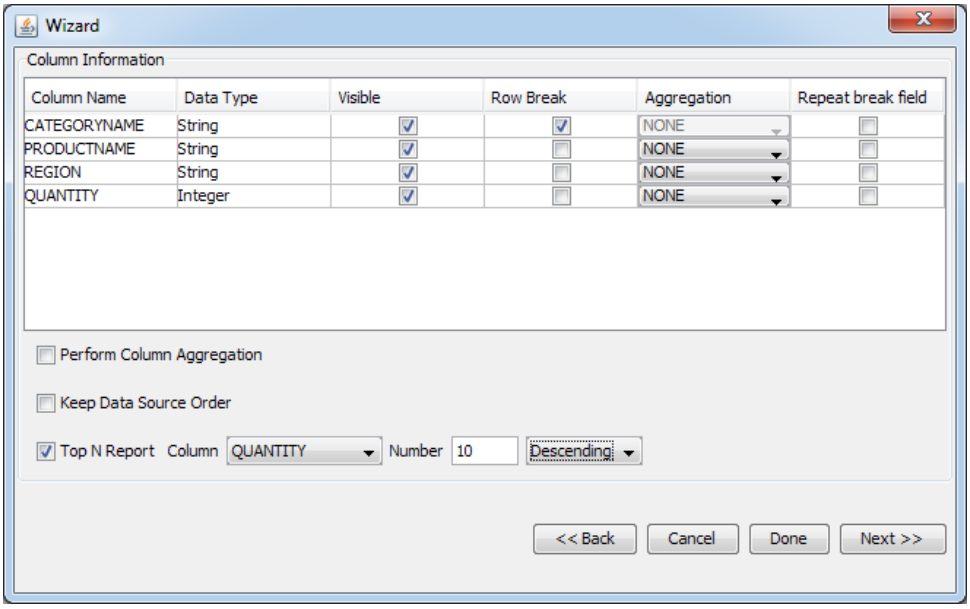
However, selecting *Perform Column Aggregation* and changing the `Product` aggregation to `Count` will produce the following report:

Order#	Product	Quantity
12	2	5
14	2	7
	2	**5**

The sixth field allows you to select whether to repeat a break field or not. By default, a row break column will only print each distinct value in the column once (i.e. once for each group). Selecting this option will cause the break fields to repeat for each row of data in the group.



Summary Break Report Mapping



Summary Break Report Mapping

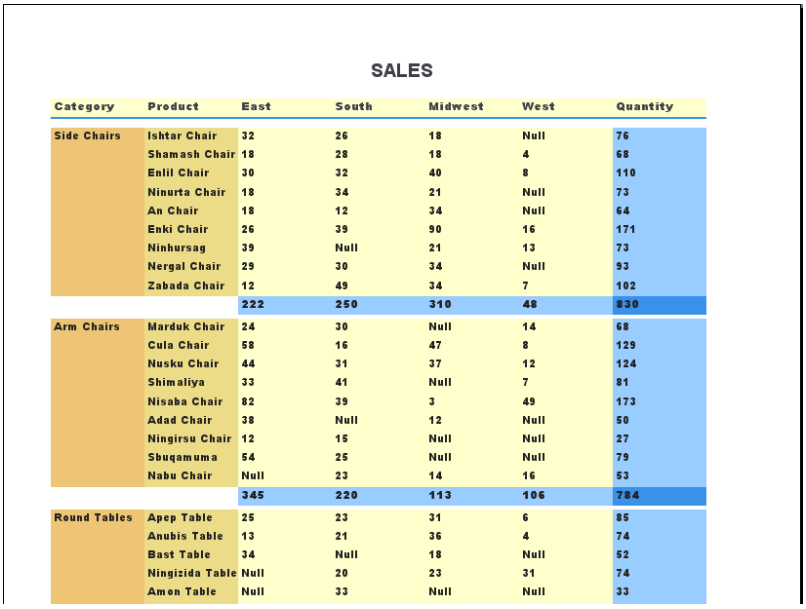
If you want to keep the data ordering specified in the data source, select the *Keep Data Source Order* option. To learn more about this feature, see Section 5.1.2.9 - Keep Data Source Order.

5.1.2.2.1.1. Top N Report

You can also create a Top N report using the summary break format. A Top N report will order and show you to highest set of values based on a particular column in the report. For summary break reports, you can also specify to show the highest values for each group. For example, you may want to show the top five customers in each region based on total sales. To do this, check the *Top N Report* checkbox at the bottom of the data mapping window. You can then specify which column you want to use as the measure (for the above example, it would be total sales column), the number that you want to retrieve (i.e. 10 or 20, etc), and whether you want to display the columns in ascending or descending order. The summary break report will then return the number of records specified (or number of records specified for each row break) sorted by the highest value for the specified column.

5.1.2.3. Crosstab Report

A crosstab report is a report format that shows and summarizes columnar data in a matrix-like form. Crosstab reports often resemble spreadsheets. Both rows and columns are summarized, allowing multi-dimensional data to be displayed in a two-dimensional format.



SALES

Category	Product	East	South	Midwest	West	Quantity
Side Chairs	Ishtar Chair	32	26	18	Null	76
	Shamash Chair	18	28	18	4	68
	Enlil Chair	30	32	40	8	110
	Ninurta Chair	18	34	21	Null	73
	An Chair	18	12	34	Null	64
	Enki Chair	26	39	90	16	171
	Ninhursag	39	Null	21	13	73
	Nergal Chair	29	30	34	Null	93
	Zabada Chair	12	49	34	7	102
		222	250	310	48	830
Arm Chairs	Marduk Chair	24	30	Null	14	68
	Cula Chair	58	16	47	8	129
	Nusku Chair	44	31	37	12	124
	Shimaliya	33	41	Null	7	81
	Nisaba Chair	82	39	3	49	173
	Adad Chair	38	Null	12	Null	50
	Ningirsu Chair	12	15	Null	Null	27
	Shuqamuma	54	25	Null	Null	79
	Nabu Chair	Null	23	14	16	53
		345	220	113	106	784
Round Tables	Apep Table	25	23	31	6	85
	Anubis Table	13	21	36	4	74
	Bast Table	34	Null	18	Null	52
	Ningizida Table	Null	20	23	31	74
	Amon Table	Null	33	Null	Null	33

Crosstab Report

The actual report is only constructed during running or preview. Since the crosstab table is essentially constructed from scratch every time the report is run, it is easier to create a smoothly collapsing and expanding crosstab table.

5.1.2.3.1. Data Mapping

Report mapping for crosstab reports is more complicated than other report types and may require some planning to be able to execute it correctly.

The first step in data mapping for this type is to select the columns you want to include in your report from the *Set Report Mapping* window in the Wizard. The left side of the window lists all available columns from your query or data file and the right side lists the columns that will be displayed in the report. Click on a column name to select it (**Ctrl**+**click** for multiple selections) and press the *ADD* or *REMOVE* buttons to add it or remove it from the report. The order in which you select columns in this window is the order in which they will appear in your report. Once you select the columns, click the *Next* button.

The second step in data mapping is to set column options.

The top of the column options window displays a table showing all of the columns that you have selected for your report and all available options.

- The first field displays the *Column Name*.
- The second field displays the *Data Type*.
- The third field displays a checkbox marked *Visible*. Checking or un-checking this box will make the column visible or invisible within the body of the report.
- The fourth field is the *Row Break* field. Selecting a column as the row break field will cause the report to insert a new crosstab row for each unique entry in the selected column.
- The fifth field is the *Column Break* field. Selecting a column as a column break field will cause the report to create a new column for each unique entry in the database.
- The sixth field is the *Column Break Value* field. Columns that you select as column break values become the fields that are summarized in the report.
- The seventh field allows you to select column *Aggregation*. A drop-down menu allows you to select the aggregation to be performed for that column. The aggregation operations that can be performed are: *sum*, *maximum*,

minimum, *count*, *average*, *first*, *last*, *sum of squares*, *variance*, *standard deviation*, and *count distinct*. Columns that have been selected as *Row Break* or *Column Break* fields cannot be aggregated.

- The eighth field *Order* allows you to specify ordering for the column break column. This order will determine how the crosstab columns are drawn from left to right. Ordering options are *not sorted*, *ascending*, and *descending*.



Note

The aggregation that will be performed depends on the aggregation of the *Column Break Value* column(s). For example: if you use the *MAX* aggregation, the greatest value from the row will be displayed in the *Grand Total* column.

Keep Data Source Order: If you want to keep the data ordering specified in the data source, select this option. To learn more about this feature, see Section 5.1.2.9 - Keep Data Source Order.

Reorder Columns: This option is selected by default. When enabled, columns are automatically rearranged so that all crosstab elements (such as *Row Break*, *Column Break*, and *Column Break Value*) appear on the left side of the report. If you want to place other columns before or between the *Row Break* and *Column Break* fields, you can uncheck the *Reorder Columns* option. This will preserve the column order as defined in the Report Wizard.

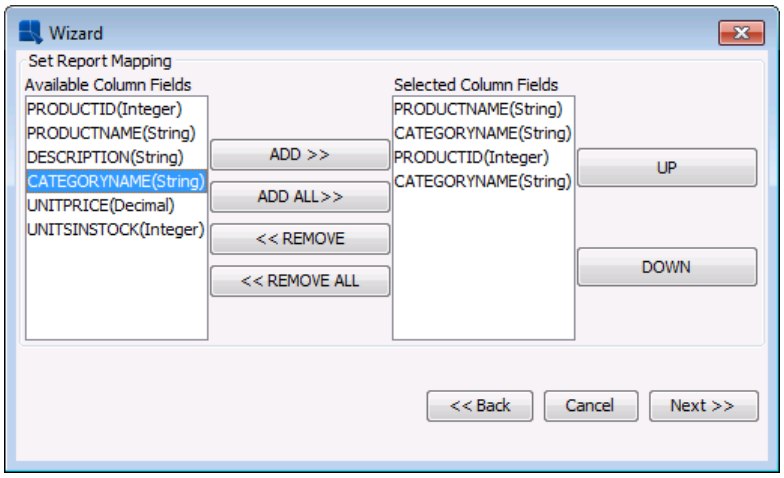
Show Column Aggregation Grand Total: If this option is enabled, a column will be added into the report. Each row of the column will aggregate all values from the corresponding row into a single value.

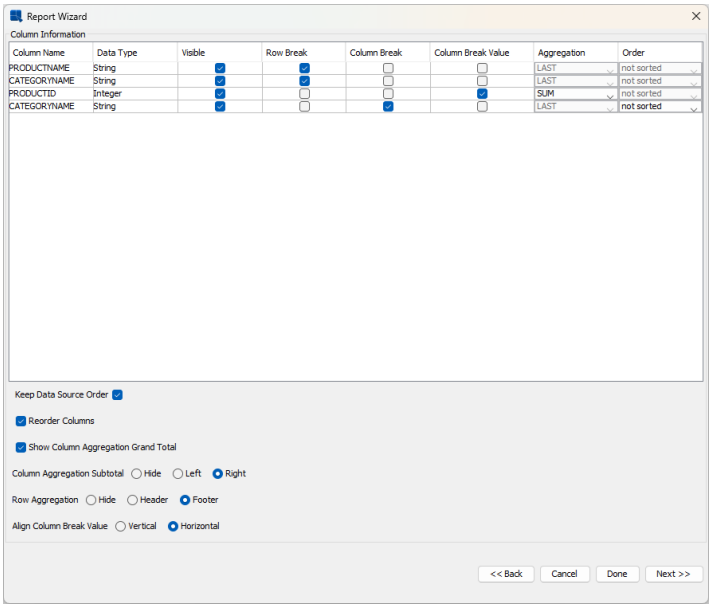
Column Aggregation Subtotal: Subtotal column is very similar to the *Grand Total* column, except it doesn't aggregate the whole row, but each one of the *Column Break* group. This option allows you to disable the subtotal column or set its position (right or left to the *Column Break* group).

Row Aggregation: This is also very similar to the *Grand Total* column. The main difference is that this option doesn't aggregate rows, but columns. Also, it doesn't add a column to the report, but it adds a row instead. Use this options to disable the *Row Aggregation* row, or to adjust it's position (on top of the table or below the table).

Align Column Break Value: If you select more than one *Column Break Value* columns, use this option to choose whether the values will be aligned in rows or columns.

For example, assume we have following crosstab mapping:





Crosstab Fixed-Field Mapping

Setting `CATEGORYNAME` and `PRODUCTNAME` as **`Row Break`** fields, `PRODUCTID` as a **`Column Break Value`** field with an aggregation of **`SUM`**, and `CATEGORYNAME` as a **`Column Break`** field with the *Show Column Aggregation Grand Total* option enabled will produce a report from the beginning of the *Crosstab Report* section.

You can easily change the crosstab settings using the *Change Data Mapping* option under the *Data* menu. This will open the crosstab report mapping dialog again. To learn more about this feature, see Section 5.1.2.8 - Change Data Mapping.

The following table shows various formatting in crosstab reports:

	Column Aggregation Subtotal	Row Aggregation	Align Column Break Value	Template
Report A [https://data.quadbase.com/Docs72/help/manual/code/export/Fixed_field_Crosstab_A.pdf]	Right	Footer	Horizontal	Template A [https://data.quadbase.com/Docs72/help/manual/code/templates/Fixed_field_Crosstab_A.pak]
Report B [https://data.quadbase.com/Docs72/help/manual/code/export/Fixed_field_Crosstab_B.pdf]	Right	Header	Horizontal	Template B [https://data.quadbase.com/Docs72/help/manual/code/templates/Fixed_field_Crosstab_B.pak]
Report C [https://data.quadbase.com/Docs72/help/manual/code/ex-	Left	Footer	Horizontal	Template C [https://data.quadbase.com/Docs72/help/manual/code/tem-

	Column Aggregation Subtotal	Row Aggregation	Align Column Break Value	Template
port/Fixed_field-_Crosstab_C.pdf]				plates/Fixed_-field-_Crosstab_C.-pak]
Report D [https://data.quadbase.com/Docs72/help/manual/code/export/Fixed_field-_Crosstab_D.pdf]	Left	Header	Vertical	Template D [https://data.quadbase.com/Docs72/help/manual/code/templates/Fixed_-field-_Crosstab_D.-pak]
Report E [https://data.quadbase.com/Docs72/help/manual/code/export/Fixed_field-_Crosstab_E.pdf]	Left	None	Vertical	Template E [https://data.quadbase.com/Docs72/help/manual/code/templates/Fixed_-field-_Crosstab_E.-pak]

5.1.2.3.2. Transposing Data

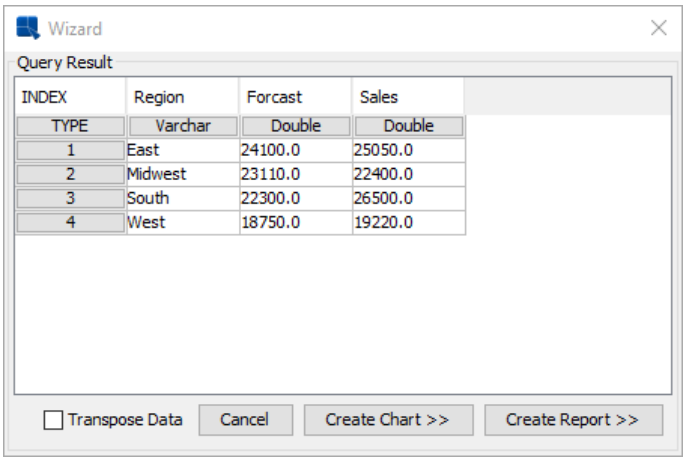
The crosstab report also provides a unique feature that allows users to create a transposed report. Data transposition allows users to create a report presentation where the input data is essentially rotated by 90 degrees. Columns become rows and rows become columns. For example, you have a following set of data:

Region	Forcast	Sales
East	24100	25050
Midwest	23110	22400
South	22300	26500
West	18750	19220

The data transposition feature can be used to create a report that will look like this:

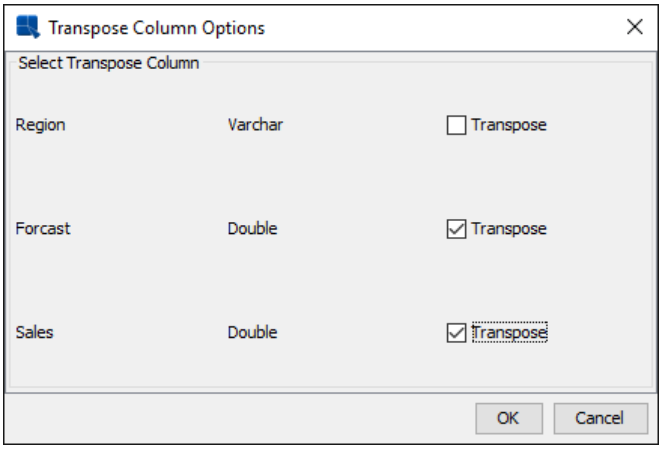
	East	Midwest	South	West
Forcast	24100	23110	22300	18750
Sales	25050	22400	26500	19220

To transpose the data, click the *Transpose Data* option at the bottom of the result screen that first appears when you select a data source for the report. The following dialog shows the initial result set from the sample data in the example above.



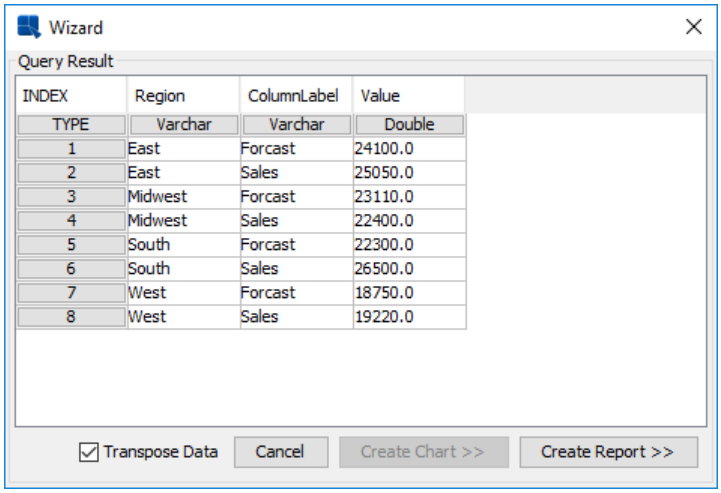
Data Table Dialog

When you check the option at the bottom, a new dialog will appear allowing you to select the columns that you want to transpose. In this example, you would select the `Forecast` and `Sales` columns. Note that in order to perform transposition, the selected columns must have the same data type.



Select Transpose Dialog

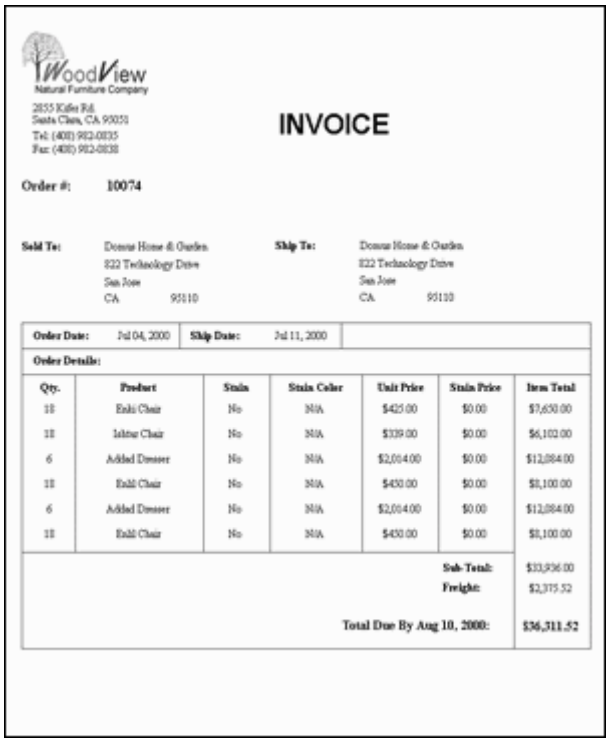
After making your selections, click the *OK* button to apply the changes. You will see the transposition in the data table dialog. The Sales and Forecast headers are transposed into a column called `ColumnLabel` and their values are merged into a column called `Value`. Now you can map this result set using a crosstab report where `ColumnLabel` is the Row Break, `Region` is the Column Break, and `Value` is the Column Break Value to create the final transposed report layout.



Data After Transposition

5.1.2.4. Master & Details Report

A Master & Details report is a set of tabular data that is grouped according to a master field. This report type is most commonly used when you have fields in your data table that have a one to many relationship. A good example of this is an invoice. For each order number in a database there will be customer information and several items with pricing information. A Master & Details report would be used to group the information according to order number.



Master & Details Report

5.1.2.4.1. Data Mapping

The first step in data mapping for this type is to select the columns you want to include in your report from the *Set Mapping Window* in the Report Wizard. The left side of the window lists all available columns from your query or data file and the right side lists columns that will be displayed in the report. Click on a column name to select it (**Ctrl**+**click** for multiple selections) and press the *ADD* or *REMOVE* buttons to add it or remove it from the report. The order in which you select columns in this window is the order in which they will appear in your report.

The second step in data mapping is to set the column options. There are three column options available for Master & Details reports. There is a drop-down menu at the bottom of the column options window. It is labeled `primary key`. The drop-down menu contains all of the columns that you have selected for the report. Selecting a column as the `primary key` will group the report according to that column. A new group will be created every time the value in the selected column changes.

The top of the column options window displays a table showing all of the columns that you have selected for your report and the available options. The first field displays column names, the second one displays data type, and the third one displays a checkbox marked *Visible*. Checking or un-checking this box will make the column visible or invisible within the body of the report. The fourth field is the `Master field` field. Selecting a column as a Master field will place the column value in the column header instead of the data section of the report.

A checkbox at the bottom of the data mapping window allows you to change the layout of the report. Checking the *Side-By-Side Layout* box will arrange the report so that the master section is displayed next to the details section rather than above it.

For example, suppose we have the following data table:

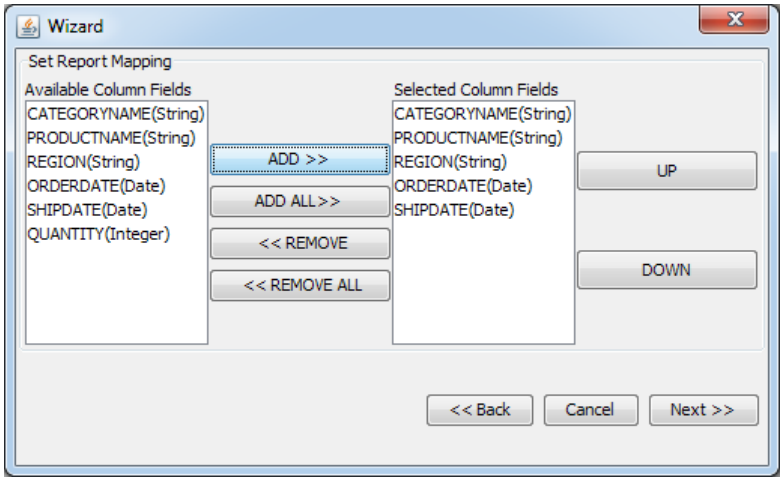
Order #	Customer Name	Product	Unit Price	Quantity
12	Paul Campbell	Chair	$24.95	4
12	Paul Campbell	Table	$127.50	1
14	Sally Hayes	Cabinet	$227.25	2
14	Sally Hayes	Chair	$24.95	2
14	Sally Hayes	Table	$127.50	1

Setting `Order #` as the **`primary key`** and `Customer Name` as a **`Master field`** without using side-by-side layout will result in the following report:

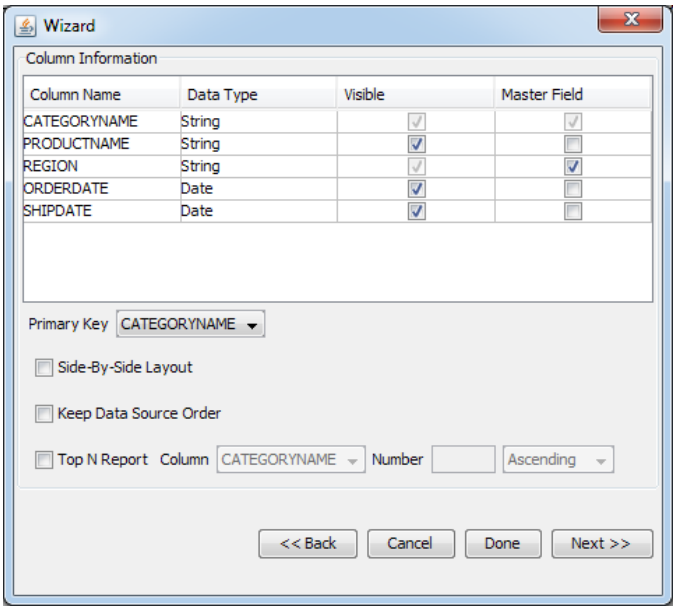
Order #	**12**	
Customer Name	**Paul Campbell**	
Product	**Unit Price**	**Quantity**
Chair	$24.95	4
Table	$127.50	1
Order #	**14**	
Customer Name	**Sally Hayes**	
Product	**Unit Price**	**Quantity**
Cabinet	$227.25	2
Chair	$24.95	2
Table	$127.50	1

However, the same report with side-by-side layout would look like this:

		Product	UnitPrice	Quantity
Order #:	**12**	Chair	$24.95	4
Customer Name:	**Paul Campbell**	Table	$127.50	1
Order#:	**14**	Cabinet	$227.25	2
Customer Name:	**Sally Hayes**	Chair	$24.95	2
		Table	$127.50	1



Master & Details Report Mapping



Master & Details Report Mapping

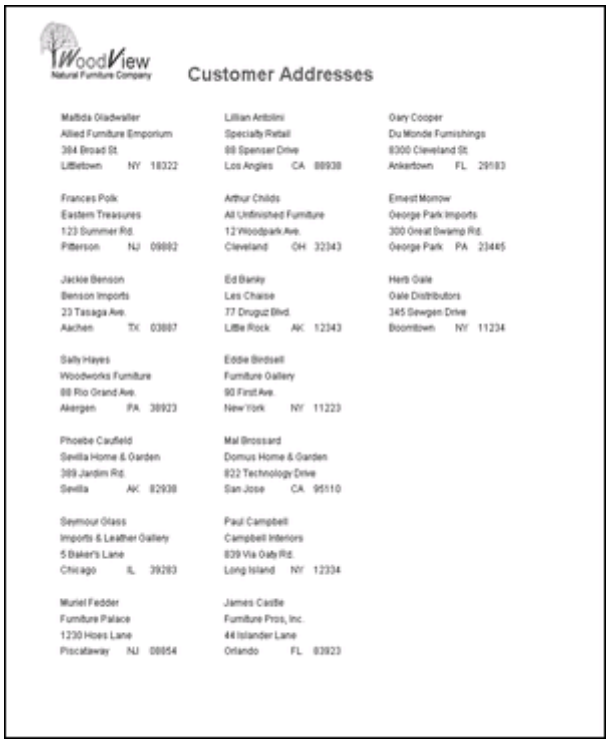
If you want to keep the data ordering specified in the data source, select the *Keep Data Source Order* option. To learn more about this feature, see Section 5.1.2.9 - Keep Data Source Order.

5.1.2.4.1.1. Top N Report

You can also create a Top N report using the Master & Details format. A Top N report will order and show you the highest set of values based on a particular column in the report. For Master & Details reports you can also specify to show the highest values for each group. For example, you may want to show the top five customers in each region based on total sales. To do this, check the *Top N Report* checkbox at the bottom of the data mapping window. You can then specify which column you want to use as the measure (for the above example it would be total sales column), the number that you want to retrieve (i.e. 10 or 20, etc), and whether you want to display the columns in ascending or descending order. The master & details report will then return the number of records specified (or number of records specified for each row break) sorted by the highest value for the specified column.

5.1.2.5. Mailing Label Report

A mailing label report is similar to a simple columnar report; however, it pre-arranges data in a way that allows you to create mailing labels. Data is arranged within the data table vertically and is wrapped automatically, making it easy to create a report showing an address list for mailing labels.

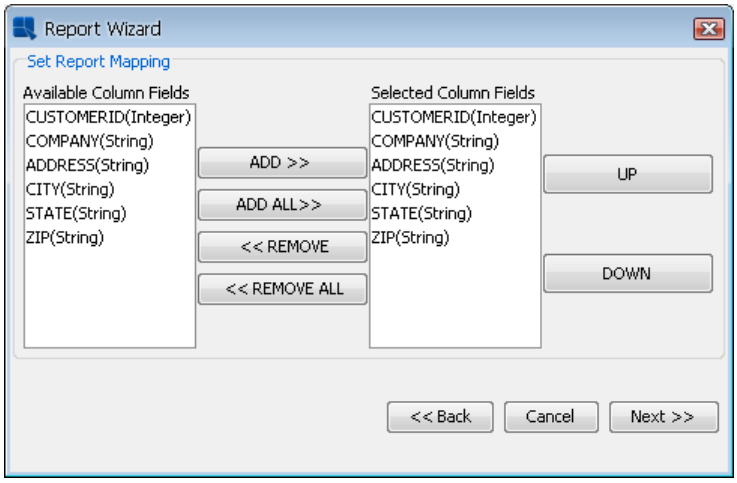


Mailing Label Report

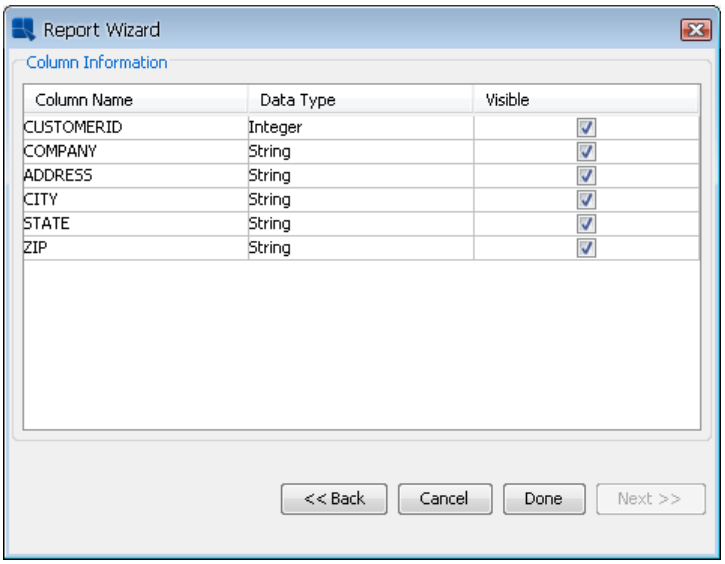
5.1.2.5.1. Data Mapping

Data Mapping for this type is exactly the same as for simple columnar report. The first step is to select the columns you want to include in your report from the *Set Report Mapping* window in the Report Wizard. The left side of the window lists all available columns from your query or data file and the right side lists columns that will be displayed in the report. Click on a column name to select it (**SHIFT**+**Click** for multiple selections) and press the *ADD* or *REMOVE* buttons to add it or remove it from the report. The order in which you select columns in this window is the order in which they will appear in your report.

The second step for data mapping is to set the column options. For mailing label reports there is only one available option. The window displays a table showing all of the columns you have selected for the report and all available options. The first field displays column names, the second one displays data type, and the third one displays a checkbox marked *visible*. Checking or un-checking this box will make the column visible or invisible within the body of the report (all columns are visible by default).



Mailing Label Report Mapping

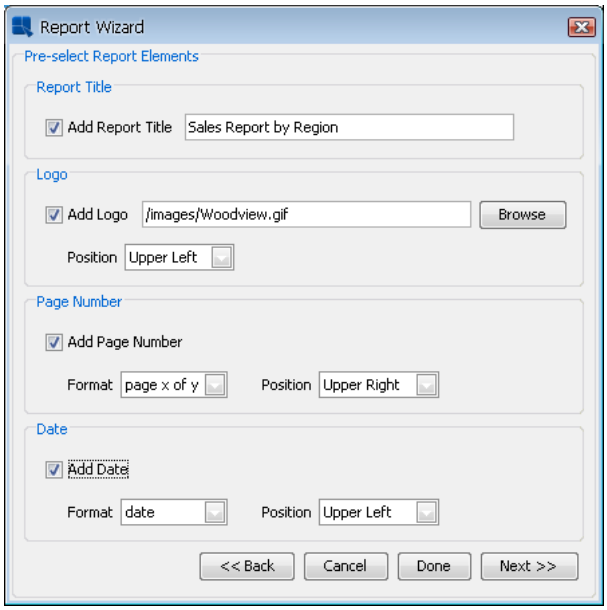


Mailing Label Report Mapping

5.1.2.6. Additional Formatting Options

After you finish specifying mapping options, you can dismiss the Wizard and begin editing/modifying the report. To exit the Wizard, click the *Done* button in the last data mapping window. This will take you back to the main designer window where a blank (unformatted) report will be generated based on your mapping specifications.

Instead of generating an unformatted report, ReportDesigner also provides several additional options in the Wizard that allows you to automatically place some elements in the report, as well as provide a default style for the report elements. To continue on in the Wizard, click the *Next* button in the last data mapping window. This will bring up a dialog prompting you to add several elements to the report.

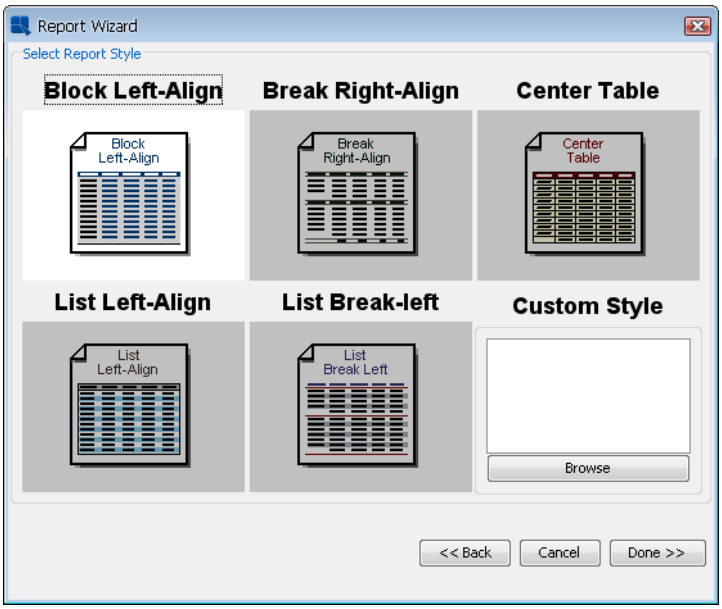


Add Elements Dialog

This dialog allows you to specify a report title and a logo for the report header. If you select to add a report title, you will be able to type in the title text. If you specify to add a logo, you can specify the location of the image file either using a file or URL path. You can also specify whether to place the logo in the upper right or upper left corner of the page.

You can also specify to add page numbers and date to the page headers and footers. You can select format and position of the elements either in the page header or footer and align them to the right, left, or center of the page.

After you finish specifying the elements, you can again click the *Done* button to dismiss the Wizard and begin editing/modifying the report with the added elements. If you click *Next*, you will be taken back to the last dialog in the Report Wizard. This dialog allows you to select a style for the report.



Report Style Dialog

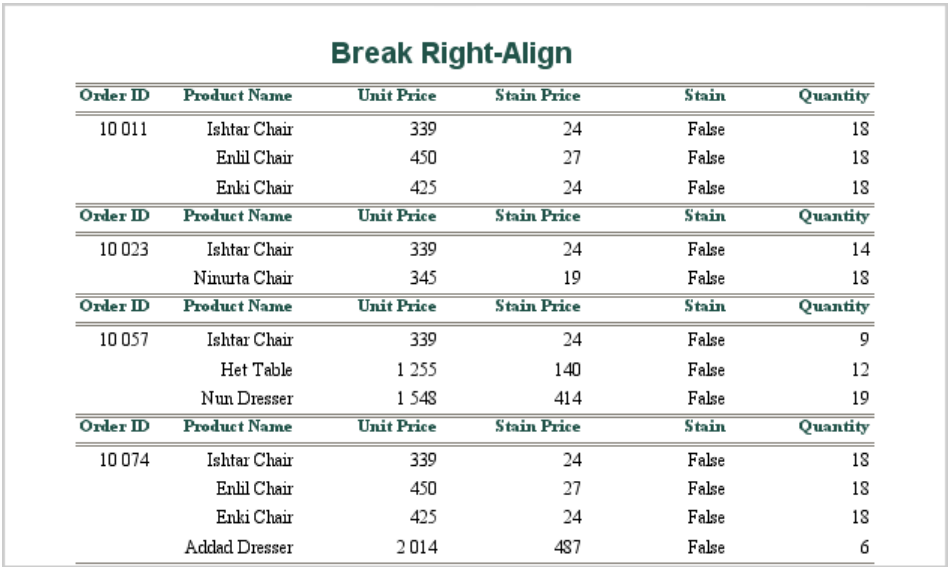
From here you can pick a pre-defined style or a custom style you want to use for the report. Selecting a style will apply formats to the report elements, as well as change the default attributes of new report elements. Below are examples of each of the pre-defined report styles when applied to the same summary break report:

Block Left-Align

Order ID	Product Name	Unit Price	Stain Price	Stain	Quantity
10 011	Ishtar Chair	339	24	False	18
	Enlil Chair	450	27	False	18
	Enki Chair	425	24	False	18
10 023	Ishtar Chair	339	24	False	14
	Ninurta Chair	345	19	False	18
10 057	Ishtar Chair	339	24	False	9
	Het Table	1 255	140	False	12
	Nun Dresser	1 548	414	False	19
10 074	Ishtar Chair	339	24	False	18
	Enlil Chair	450	27	False	18
	Enki Chair	425	24	False	18
	Addad Dresser	2 014	487	False	6

Block Left-Align

This style sets text alignment to the left and draws the report on a gray background with blue headers.

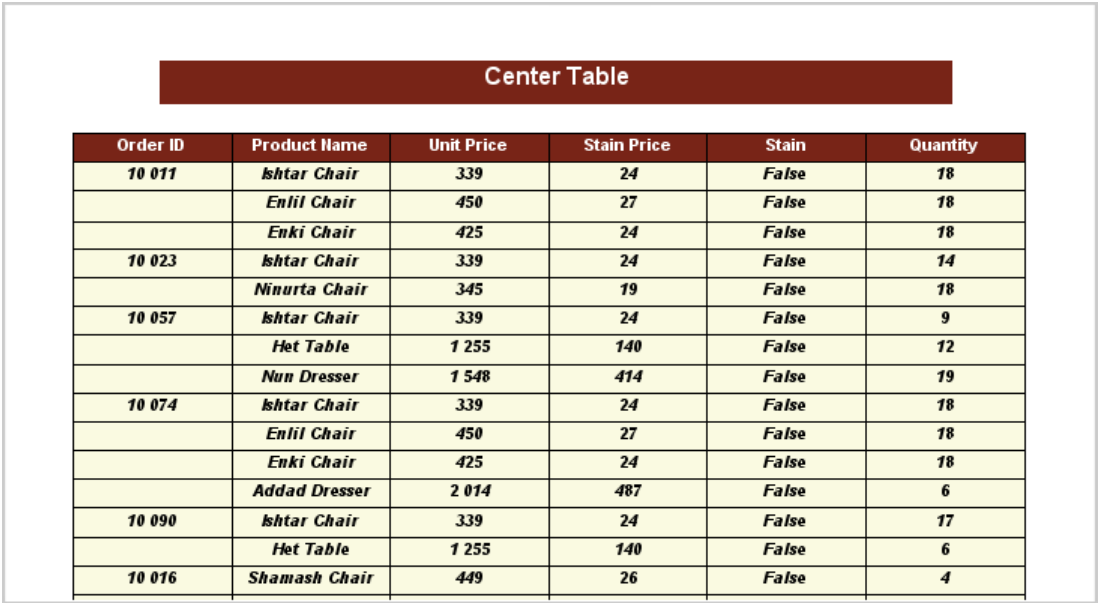


Break Right-Align

Order ID	Product Name	Unit Price	Stain Price	Stain	Quantity
10 011	Ishtar Chair	339	24	False	18
	Enlil Chair	450	27	False	18
	Enki Chair	425	24	False	18
Order ID	**Product Name**	**Unit Price**	**Stain Price**	**Stain**	**Quantity**
10 023	Ishtar Chair	339	24	False	14
	Ninurta Chair	345	19	False	18
Order ID	**Product Name**	**Unit Price**	**Stain Price**	**Stain**	**Quantity**
10 057	Ishtar Chair	339	24	False	9
	Het Table	1 255	140	False	12
	Nun Dresser	1 548	414	False	19
Order ID	**Product Name**	**Unit Price**	**Stain Price**	**Stain**	**Quantity**
10 074	Ishtar Chair	339	24	False	18
	Enlil Chair	450	27	False	18
	Enki Chair	425	24	False	18
	Addad Dresser	2 014	487	False	6

Break Right-Align

This style sets text alignment to the right and draws lines to demarcate the headers.

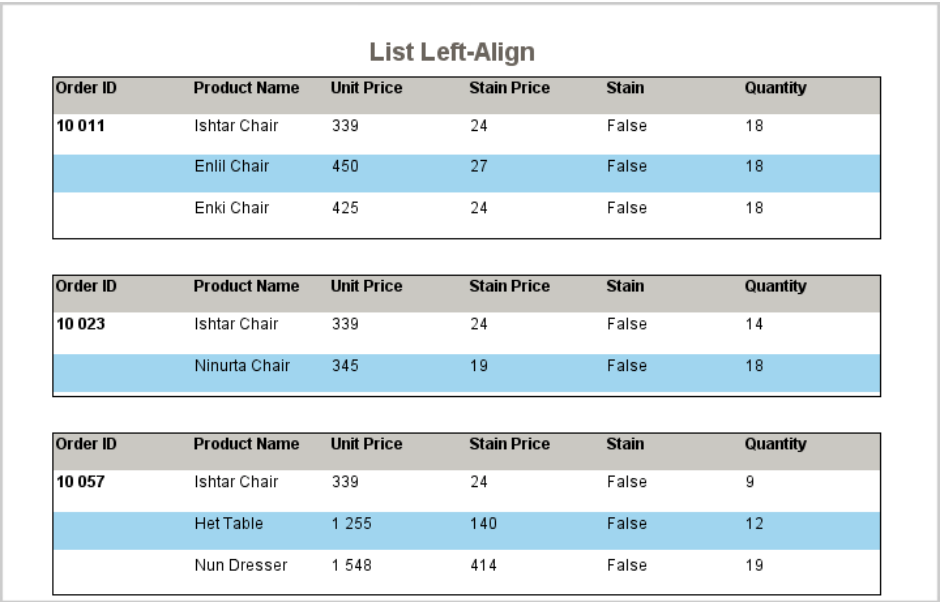


Center Table

Order ID	Product Name	Unit Price	Stain Price	Stain	Quantity
10 011	*Ishtar Chair*	*339*	*24*	*False*	*18*
	Enlil Chair	*450*	*27*	*False*	*18*
	Enki Chair	*425*	*24*	*False*	*18*
10 023	*Ishtar Chair*	*339*	*24*	*False*	*14*
	Ninurta Chair	*345*	*19*	*False*	*18*
10 057	*Ishtar Chair*	*339*	*24*	*False*	*9*
	Het Table	*1 255*	*140*	*False*	*12*
	Nun Dresser	*1 548*	*414*	*False*	*19*
10 074	*Ishtar Chair*	*339*	*24*	*False*	*18*
	Enlil Chair	*450*	*27*	*False*	*18*
	Enki Chair	*425*	*24*	*False*	*18*
	Addad Dresser	*2 014*	*487*	*False*	*6*
10 090	*Ishtar Chair*	*339*	*24*	*False*	*17*
	Het Table	*1 255*	*140*	*False*	*6*
10 016	*Shamash Chair*	*449*	*26*	*False*	*4*

Center Table

This style centers text and draws table borders around the report cells.



List Left-Align

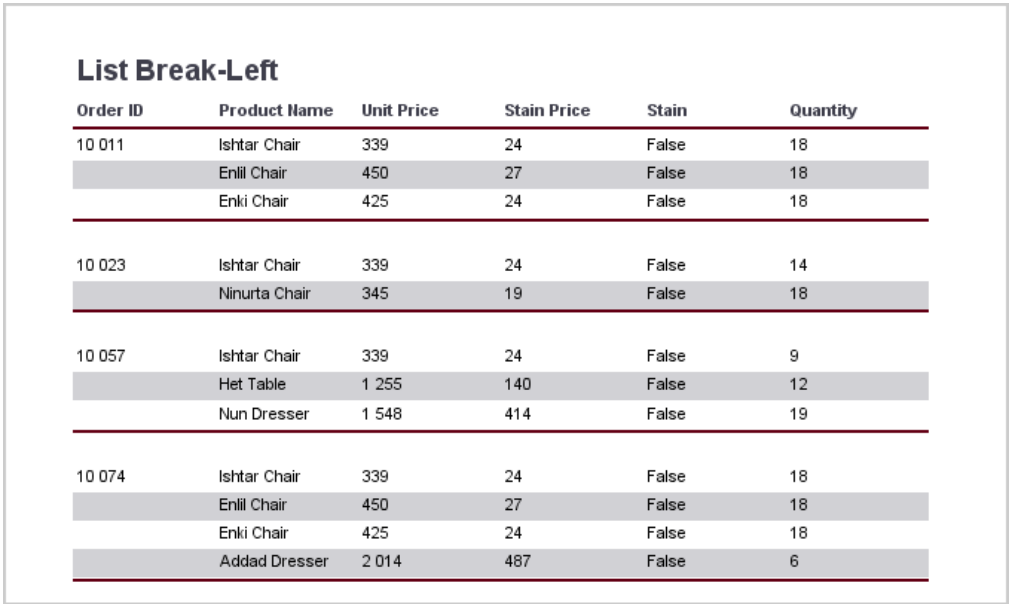
Order ID	Product Name	Unit Price	Stain Price	Stain	Quantity
10 011	Ishtar Chair	339	24	False	18
	Enlil Chair	450	27	False	18
	Enki Chair	425	24	False	18

Order ID	Product Name	Unit Price	Stain Price	Stain	Quantity
10 023	Ishtar Chair	339	24	False	14
	Ninurta Chair	345	19	False	18

Order ID	Product Name	Unit Price	Stain Price	Stain	Quantity
10 057	Ishtar Chair	339	24	False	9
	Het Table	1 255	140	False	12
	Nun Dresser	1 548	414	False	19

List Left-Align

This style sets text alignment to the left and draws alternating color for each row.



List Break-Left

Order ID	Product Name	Unit Price	Stain Price	Stain	Quantity
10 011	Ishtar Chair	339	24	False	18
	Enlil Chair	450	27	False	18
	Enki Chair	425	24	False	18
10 023	Ishtar Chair	339	24	False	14
	Ninurta Chair	345	19	False	18
10 057	Ishtar Chair	339	24	False	9
	Het Table	1 255	140	False	12
	Nun Dresser	1 548	414	False	19
10 074	Ishtar Chair	339	24	False	18
	Enlil Chair	450	27	False	18
	Enki Chair	425	24	False	18
	Addad Dresser	2 014	487	False	6

List Break-Left

This style sets text alignment to the left, adds lines to demarcate column headers, and draws alternating color for each row.

Report styles will only effect the appearance properties of the report elements and will not change the report type or data mapping options that were selected earlier. Once you specify a style, click the *Done* button to dismiss the Wizard and continue to the main designer window.

5.1.2.6.1. Custom Styles

In addition to the five pre-defined styles included with ReportDesigner, users can create their own style definitions. Style definitions are special versions of report templates that are saved with a `.stl` extension. Users can create a style definition using any report template.

5.1.2.6.1.1. Creating a Custom Style

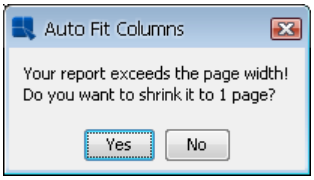
The first step in creating a custom style is to set the global formats. This allows you to control the default look and feel for all elements the user inserts into the report. For more information about global formats, see Section 5.1.3.9.1 - Global Formatting.

The second step in creating a custom style is to select the look and feel for the elements in each section of the report. To do this, select the report element whose formats you want to apply to all elements in the section when the style is applied and select *Set Attributes as Section Style* from the pop-up menu.

Once you finish setting the attributes, save the report template as a custom style by checking the *Create Style* option in the save as dialog.

5.1.2.7. Fit Columns to Page Width

Often when you finish with the Report Wizard, your created report will be wider than the default page width (8.5"). When this happens, you will be presented with a warning.

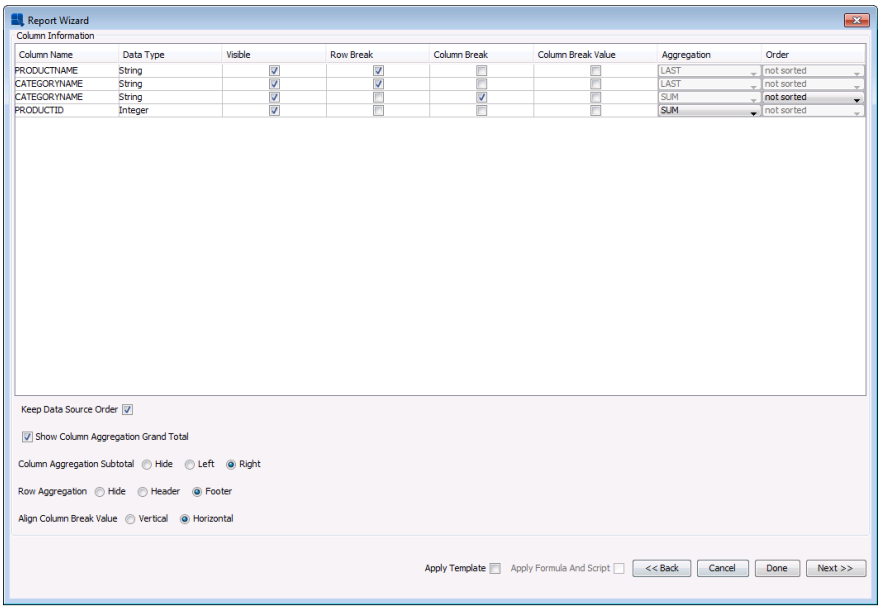


Auto-Fit Columns Dialog

The dialog gives you an option to shrink the report to fit within the page. If you select *Yes*, all columns will automatically shrink to fit the width of the page. The page boundary is indicated by a vertical bar within the design window. If you select to keep the columns overlapping the page boundary, the overflow will be printed on a new page when you preview the report. You can auto-fit the columns anytime you are editing the report by selecting *Auto Fit Columns* from the Format menu. The page width can also be adjusted by selecting *Page Setup* from the Option menu.

5.1.2.8. Change Data Mapping

Once you complete the data mapping, a rough version of the report will be generated. If it is not correct, you can change the data mapping by selecting *Change Data Mapping* from the *Data* menu, or clicking the *Change Data Mapping* button on the toolbar. This will take you back to the Report Wizard (starting with data mapping), allowing you to go back to change the report type and data mapping.

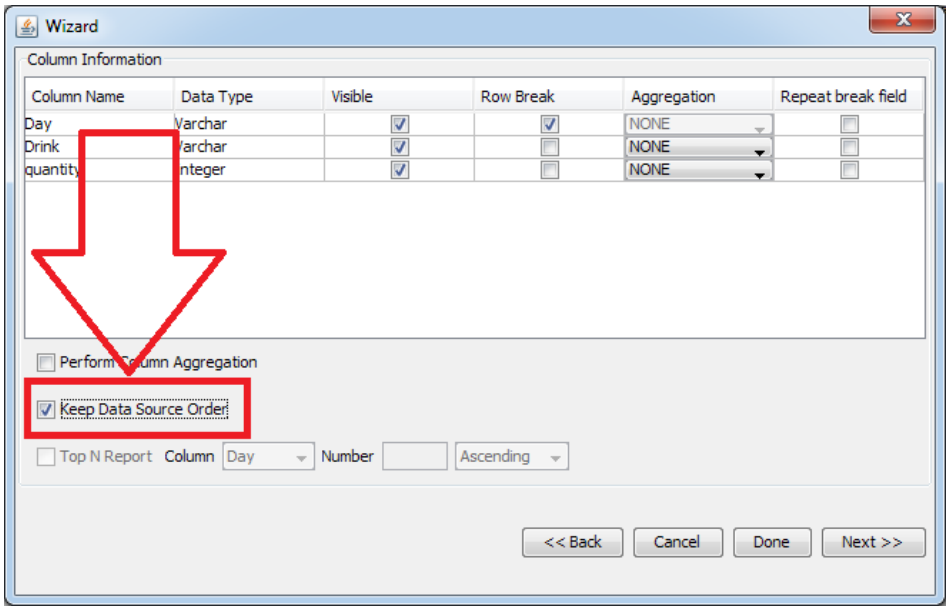


Change Data Mapping window

When you change the data mapping, there will be two checkboxes at the bottom of the last screen in the Report Wizard. One is marked *Apply Template* and the other *Apply Formula and Script*. Checking the first box will retain the formatting of your report; however, the labels and formulas may not be correct. Un-checking this box will create a blank report with the new data mapping. If you did not format lot of the report objects, it is recommended that you un-check this box. This will ensure that the new data mapping is completely accurate. The second box allows you to include functions and scripts placed in the data table section to be applied (formulas and scripts in other sections are always applied), when you select the *Apply Template* option. Please note that if you have changed the number of columns or the data types of certain columns, the functions and/or scripts may no longer work.

5.1.2.9. Keep Data Source Order

Summary Break, Cross-tab and Master&Details reports use memory-optimized query processing which can unfortunately change the order of the data. In other words: the data in the report may be displayed in different order than the data in the data source. To prevent this from happening, choose the *Keep Data Source Order* option in the report mapping dialog. This will disable the memory-optimized query processing which means that the data will be passed to the report without any optimalization (i.e. the original data ordering will be preserved).



Keep Data Source Order option

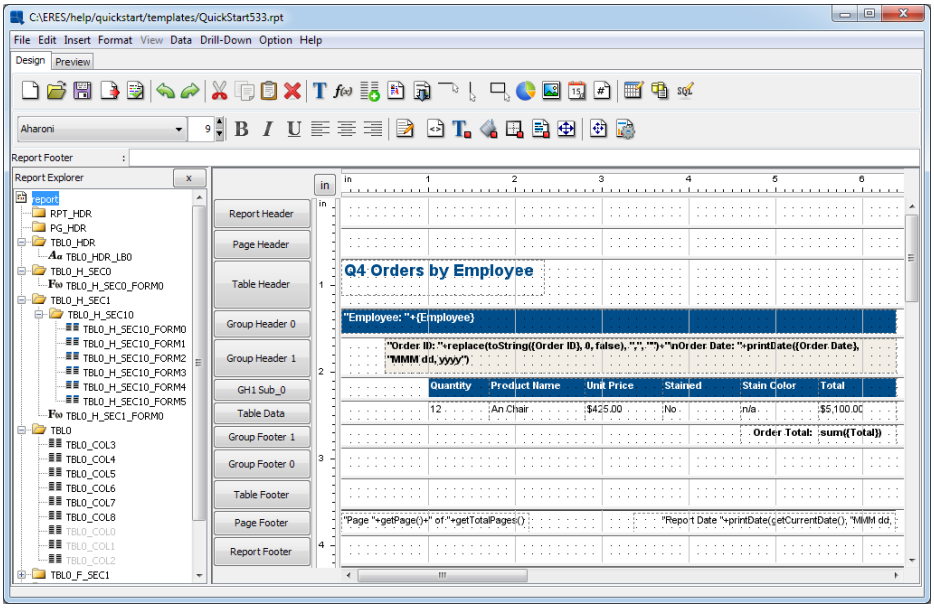


Note

If you choose the *Keep Data Source Order* option, the *Top N Report* option for Master&Details and Summary Break reports will be disabled.

5.1.3. The Designer Interface

Users who have some familiarity with other report writers will recognize the banded-style interface of the Report Designer. Report elements can be added and modified within the Designer window and then the resulting report can be previewed in the preview window.



Report Designer Interface

5.1.3.1. Report Sections

Each group within the designer interface represents a different report section. The section names are listed on the buttons to the left of each report group.

Report Header: This section is like the title of the report. It appears only once at the top of the report.

Page Header: Any element placed in this section will appear at the top of every page of the report.

Table Header: This section serves as a header to the detail or data section of the report. By default, it only appears once in the report.

Group Header: This section appears in reports with grouped data (i.e. master & details report), or data with row breaks inserted (i.e. summary break report). It repeats at the top of each grouping within the report.

Table Data: This is the main section of the report that contains most of the data. Data columns that have been selected for the report are placed in this section and are repeated for each entry in the column.

Group Footer: This section appears on reports with grouped data (i.e. Master & Details report) or data with row breaks inserted (i.e. summary break report). It repeats at the bottom of each grouping within the report.

Table Footer: This section serves as the footer or data section of the report. By default, it appears only once in the report.

Page Footer: Any element placed in this section will appear at the bottom of every page of the report.

Report Footer: This is the last summary of footer section of the report. It only appears once at the end of the report.

5.1.3.1.1. Nested Sections

There can also be nested sections in addition to the Report Header, Table Header, Group Header, Group Footer, Table Footer, and Report Footer sections. Nested sections are useful for placing sub-reports and rich text fields with variable lengths. This allows these type of elements to resize without overlapping any elements below. Nested sections inherit most of the section options from their parent sections, including page-breaking and repeating options.

To add a nested section, select *Insert Section* from the section options pop-up menu for the parent section. For more information about section options, options see Section 5.1.3.3 - Section Options. Note that if you select to insert a section from a nested section, it will inherit the same parent section (i.e. nested sections cannot be a parent).

5.1.3.2. Resizing Sections

To resize section height, place your mouse over the section divider within the design window. The pointer will change to a resize pointer. Click and drag the mouse to expand or decrease the section height.

5.1.3.3. Section Options

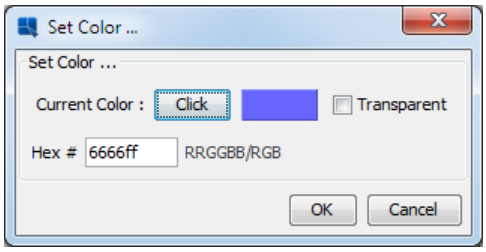
There are several formatting and display options available for each report section. You can access the section options menu by clicking the button with the section name at the left side of the Designer or by right clicking on a blank portion of the section field.

Background Color: This option is available for every section of the report. It allows you to set background color for the entire section. Selecting this option will bring up a dialog box prompting you to specify whether or not to set the background transparent. If you do not want a transparent background, uncheck the check-box and click the button. This will bring up a dialog prompting you to specify the background color. You can select font color from swatches or enter HSB/RGB values. You can also directly enter a HEX color value to the *Hex #* field.



Note

If there are no elements placed in a section, that section will not display. Hence, changing a section background color will have no effect on the finished report if the section is blank.



Background Color Dialog

Section Height: This option is available for every section of the report. It allows you to set the section height. It has the same effect as dragging the section's bottom edge, but unlike the drag&drop method, this option allows you to set the section height precisely in inches.

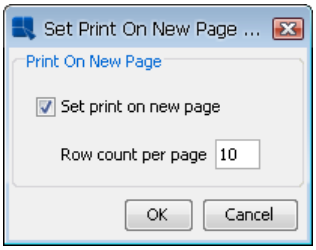
Insert Section: This option is available for Report Header, Table Header, Group Header, Group Footer, Table Footer, and Report Footer sections. Selecting this option will add a nested section to the currently selected report section.

Remove Section: This option is only available for nested sections that you have added to the report (regular sections can only be rendered invisible). This option will remove the nested section from the report and delete its contents.

Repeat On Every Page: This option is available for Table Header and Group Header sections. Selecting this option causes the section to repeat again after every page break (group headers will only repeat if the data for that group exceeds one page).

Print On a New Page: This option is available for Table Header, Group Header, Table Data, Group Footer, Table Footer, and Report Footer sections. For the Table Data section,

a new dialog will pop-up allowing you to select the number of rows to display on each page. This number will be set as a new maximum of table data rows for the page. After this amount of rows a page break will be inserted. This count does not include any formulas in the footer sections. For example, if the report contained row breaks and an aggregation calculated for each group, this element will not count towards the number of rows.



Print on New Page Dialog

Selecting this option for any other section inserts a page break each time the section appears.

New Excel Sheet: This option is only available for Group Header section(s). Selecting this option causes each group to be printed on a separate worksheet when the report is exported to Excel (XLS) or Excel 2007 (XLSX) format.

Reset Page Number: This option is only available for Group Header section(s). Selecting this option causes the report pagination to reset each time the section is run. Hence, each time the Group Header section runs, the current page number goes to 1. Note that setting this option will automatically enable the print on a new page option.

Fit Group On Page: This option is only available for Group Header section(s). It allows you to control some of the pagination for grouped data. Selecting this option will cause ReportDesigner to try to fit the entire group on a page. If the entire group cannot print on a page, it will print on a new page. If the group cannot fit on any page, (the fields are longer than the page height), it will be broken up into multiple pages.

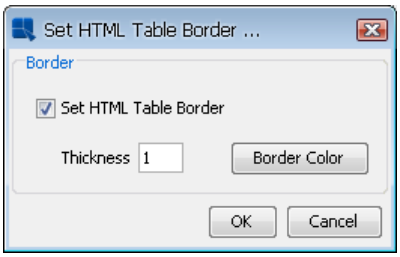
Skip First Value: This option is available for Table Header, Group Header, Group Footer, and Table Footer sections. This option is a subset of the print on a new page option. Selecting this option will ignore the page break for the very first time a section is run. For example, setting this option for the Group Header section will cause the report not to insert a page break for the first time the section is run. The page break will only occur when the report runs the Group Header for a second time.

Skip First Group Value: This option is only available for Group Header section. This option is a subset of the print on a new page option. Selecting this option will ignore the page break for the first value of each group of data. This option is used for nested grouping in a report where more than one group is set to print on a new page. To prevent a blank page for records where more than one grouped column value changes, the 'Skip First Group Value' can be set for the inner Group Header section.

Resize Cells Proportionally: This option only applies if you have set cells within the section to resize to fit content. If this option is selected, all cells in the section will resize with the variable height cells.

Invisible: This option is available for every section of the report. Selecting this option hides the entire section. To make the section visible again, click on the *Format* menu at the top of the designer and select *Make Section Visible* option.

HTML Border: This option is available for every section and only takes effect if the report is exported to HTML. Since tabular HTML (not DHTML CSS formatting) does not allow selective cell borders within a table, users can select to turn on or off a border for all cells in a section when the report is rendered in HTML. If you select this option, the following dialog will appear allowing you to set options for the HTML border:



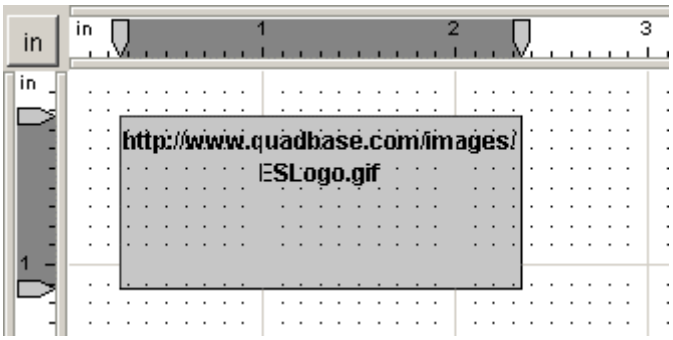
HTML Border Options

In this dialog, you can set thickness and border color by clicking the *Border Color* button. This will open the color palette dialog allowing you to pick a color for the HTML border.

Scripting: This option will bring up scripting interfaces allowing you to write, modify, or apply a script to dynamically modify certain section attributes. For more information about this feature, please see Section 5.1.7.2.7 - Section Scripts.

5.1.3.4. Rulers

The Report Designer has two rulers in the top left corner of the Design window. The rulers measure how the report will look in terms of the physical page dimensions. Please note that the dimensions do not include page margins. The rulers helps you to get a good idea of how the printed report will look like. You can also use the rulers to adjust the bounds of elements in your report. Selected element bounds will appear as shaded areas within the rulers and can be adjusted by clicking on and dragging the markers.



Rulers and in/cm button

5.1.3.4.1. Toggling Between Inches and Centimeters

The button in the upper left corner of the Design window allows you to toggle between inches and centimeters. By default, ReportDesigner uses inches for measurement. Depressing the button will change the ruler measures to centimeters. If centimeter is selected, all of the manually entered values like element bounds and page margins will also be entered and displayed in centimeters.

5.1.3.5. Designer Menus

Most of the Report Designer's features can be controlled using the drop-down menus at the top of the designer.

5.1.3.5.1. File Menu

This menu performs most of the file operations.

New: Creates a new report.

Open: Opens an existing report (`.pak/.rpt/.xml/.qrp/.stl` format).

Close: Closes the current report.

Apply Template: Imports and applies a template (`.pak/.rpt/.qrp` or `.xml`) file to the current report. All of the report attributes are changed to those of the template.

Save: Saves the current report.

Save As: Allows you to name and save the current report as a template (`.pak`) file. A checkbox within the *Save As* window allows you to generate an applet page with the Report Viewer embedded, as well as generate an XML template or a report style (`.stl`).

Export: This option exports the report to a selected format. Options include DHTML, PDF, CSV (data file), Excel (XLS), Excel 2007 (XLSX), RTF, TXT, and XML.

Print: Prints the current report (This option is only executable from Preview tab.)

Exit: Closes the application.

5.1.3.5.2. Edit Menu

This menu allows you to edit and cut/paste report elements.

Undo: Cancels the last operation performed in the Designer and reverts back to the previous state. The Designer will remember last 10 actions made.

Redo: Reverses the action of the undo command. For example, if you change the font color from black to red, you can undo this command to change it back to black and then redo this command to have it changed back to red again.

Edit: This option is available for labels, formulas, images, and charts. For labels, it will prompt you to change the label text. For formulas, it will prompt you to change the formula. For images, it will prompt you to change the image file URL. For charts, it will bring up Chart Designer to edit your chart.

Copy: This option makes a copy of the selected element and places it on the clipboard.

Cut: This option removes the selected element and places it on the clipboard.

Paste: This option pastes the current clipboard element into the report.

Delete: This will delete the selected element, except for elements in the Table Data section. It will not delete these, but render them invisible.

Remove Sub-Report: Removes a sub-report from the current report.

5.1.3.5.3. Insert Menu

This menu allows you to insert elements into the current report.

Insert Label: Inserts a label element into the current report.

Insert Formula: Inserts a formula element into the current report.

Insert Column Field: Inserts a column field into the report from the available fields (Available fields are those that have been selected for the report. See data mapping: Section 5.1.2 - Report Types and Data Mapping).

Insert Database Field: This option is only available if the template's data source is a database. It allows you to insert a field from a database that is not part of the original query.

Insert Column Header: Inserts a column header for one of the report fields.

Insert Parameter Value: Inserts a formula that returns the supplied value for a report parameter.

Insert Chart: Inserts a new chart into the current report.

Insert Image: Inserts an image from a file into the report. Images can be in JPEG, GIF, or PNG format.

Insert Rich Text Field: Inserts a new rich text into the current report.

Insert Table of Contents: Inserts a table of contents into the report.

Insert Line: Inserts a vertical or horizontal line into the report.

Insert Rectangle: Inserts a grid rectangle.

Insert Guideline: Inserts a guideline into the report. Guidelines can help you move and position elements in the report. For more information about guidelines, please see Section 5.1.3.7.14 - Guidelines.

Today's Date: Inserts the current date into the report in one of three ways: `Date`, `Date & Time`, or `Time`. The date function is actually a formula and the format can be adjusted after it is inserted. For more information about formula formatting, please see Section 5.1.3.7.3 - Data Formatting for Formulas and Column Fields.

Page Number: Inserts the page number into the report in one of two ways: `Page Number` or `Page Number of Total Pages (i.e. page X of Y)`. The page number function is actually a formula and it can be formatted after it is inserted. For more information about formula formatting, please see Section 5.1.3.7.3 - Data Formatting for Formulas and Column Fields.

Insert Sub-Report: Inserts a new or pre-existing sub-report into the current report. For more information about sub-reports, please see Section 5.1.9 - Sub-Reports.

5.1.3.5.4. Format Menu

The format menu contains formatting options for the currently selected element within the report. Some menu options will not work depending on the type of element which is currently selected.

Data Format: This option is available for formulas and column fields. It will bring up the data format editor for the data type of the selected element, either `string`, `numeric`, `date/time`, or `Boolean`.

Chart Export Format: This option is only available for charts. When reports are exported to DHTML format, charts are converted to static images. This option allows you to specify the image properties for the chart upon export.

Edit Attributes: This option is available for every report element. It will bring up a dialog box that allows you to adjust attributes depending on the selected element (e.g. bounds, background color, alignment, font style and size, rotation, border/round corners, data format).

Font Style and Size: This option is available for labels, formulas, and column fields. It brings up a dialog box prompting you to change the font appearance and the size of the selected element.

Background Color: This option is available for labels, formulas, and column fields. It allows you to adjust the background color of the selected element.

Dual Colors: This option is only available for elements in the Table Data section of the report. It allows you to specify alternating background colors, font colors, and font styles for the selected column, as well as the number of rows between changes or on which column field to base the change.

Line Color & Thickness: This option is only available for lines. It allows you to set the color and the thickness for the selected line(s).

Border: This option is available for every report element. It allows you to specify thickness (in pixels), and color and shape of corners (sharp or round) of the border drawn around the selected element. The border thickness/corners can be set individually for the each side/corner of the element.

HyperLink: This option is available for all elements except column fields. It allows you to hyperlink an element to a web page or another report.

Scripting: This option is available for all labels, formulas, and column fields. It allows you to create a script and apply conditional formatting to the selected element.

Rotation: This option is available for all labels, formulas, and column fields. It allows you to rotate the element 90 degrees clockwise or counter-clockwise.

Z-Index: This option allows you to set the z index for the selected report element. The z index controls how report elements behave (i.e. which one appears first) when placed on top of each other.

Bounds: This option is available for every report element. This feature allows you to manually enter the boundaries of an element. The measurements will be in inches or centimeters depending on which unit is selected with the toggle button in the upper left corner of the Designer window where the rulers meet.

Alignment: This option is available for labels, formulas, and column fields. It allows you to align the selected element horizontally to the left, center,right, as well as vertically to the top, middle,or bottom.

Wordwrap: This option is available for labels, formulas, and column fields. It allows users to enable or disable word wrapping within report elements.

Resize To Fit Content: This option is available for labels, formulas, and column fields. It causes the height of the selected element to adjust dynamically to fit its content.

Make Column Visible: This option is only available if you deleted an element in the Table Data section, or if you selected an invisible data column from the column options screen of the Report Wizard. It allows you to render any invisible column visible.

Make Section Visible: This option is only available if you rendered a report section invisible (See Section 5.1.3.3 - Section Options). It will make the section visible again.

Invisible: This option is available for elements in the Table Data section or for report sections. It will make the selected element invisible.

Column Wrap: This option allows column fields to be wrapped, so rather than breaking to the next page at the end of the page, they will continue next to the original column field within the current page.

Auto Fit Columns: This option will shrink all elements within the report to fit the report within the page width.

Swap Columns: This option allows the position of two elements to be swapped in the report. Elements must first be selected using **CTRL**+**Click**.

5.1.3.5.5. View Menu

The view menu contains options allowing you to navigate through the report in the Preview window and sort the data. It is only active when you are previewing the report.

First Page: This option will navigate to the first page of the report within the Preview window.

Previous Page: This option will scroll back one page within the Preview window.

Next Page: This option will scroll forward one page within the Preview window.

Last Page: This option will navigate to the last page of the report within the Preview window.

Go To Page...: This option will prompt you to enter a page number and then navigate to that page of the report within the Preview window.

Sort by (ascend): This option opens up a second menu allowing you to select one of the report columns to sort the report in ascending order (either numeric or alphabetical).

Sort by (descend): This option opens up a second menu allowing you to select one of the report columns to sort the report in descending order (either numeric or alphabetical).

Sort by... This opens a dialog which allows you to select nested sorting in a report (by more than one column). The direction of the sorting can be set for each column.

Launch Page Viewer: This will load the entire report in the page viewer window. For more information about this feature, see Section 5.1.4.2.4 - Using Page Viewer.

5.1.3.5.6. Data Menu

The data menu contains options that allow you to see and manipulate the data used in the current report and its properties.

Refresh: This option will re-query the data source used for the current report and then update the report.

Change Data Mapping: This will re-open the Report Wizard, allowing you to change data mapping for the current report.

Change Data Source: This will re-launch the Data Source Manager, allowing you to modify or change the report's data source.

Modify Database: If the report uses a database as the data source, this option allows you to modify the connection information that the report uses to connect to the database. For more information about this feature, please see Section 4.1.2.5 - Editing Database Connections.

Modify Query: If the report uses a database as the data source, this option allows you to modify the query used to retrieve the report data. For more information about this feature, please see Section 4.1.2.4 - Editing Queries.

View Column Mapping: This option will bring up a window that displays column mapping for the current report. The window displays the column index, name, data type, options, and whether it is currently visible or invisible.

View Table: This option will bring up a window containing the data table from which the current report is generated. The table will initially display only the first 20 records. Clicking on the *Show All Records* checkbox will display all of the records in the data table.

View Data Source Info: This option will bring up a window containing information about the data source that was used to create the template. The data source type, location, and the data registry location will be displayed.

Sub-Report Parameter Sharing: This allows you to link parameters defined within sub-reports to parameters in the main report or to the parameters defined in other sub-report. With this feature enabled, reports can share a user supplied parameter value automatically without having to pass the value to each sub-report. For more information about sub-report, see Section 5.1.9 - Sub-Reports.

Sub-Report Parameter Mapping: This allows you to link parameterized sub-reports to column fields in the primary report. For more information about sub-reports, see Section 5.1.9 - Sub-Reports.

Chart Parameter Mapping: This allows you to link parameterized charts (with an independent data source) to column fields in the report. For more information about this features, see Section 5.2.1.3.1 - Chart Parameter Linking.

Select Multiple Drill-Down Values: This option is only available when you preview a multi-value drill-down report. It allows you to select values to use when drilling to the next level. For more information about drill-down reports, please see Section 5.1.8 - Drill-Down.

Security: This will bring up the security settings dialog for the currently selected cell. For more information about the template security, please see Section 5.1.10 - Template Security.

Cell Properties: This option is available for all elements in the report. It will bring up a window displaying the properties of the selected element including its ID value, as well as column index, name, and data type for column fields.

Set Preview Security Level: This option allows you to specify which defined security level you want to use when previewing the template. For more information about the template security, please see Section 5.1.10 - Template Security.

Set Preview Data Options: This option allows you to turn on/off live data for the preview window as well as set the maximum number of records that should be returned when previewing the report.

Secured Query Parameters: This will bring up the parameter security dialog, allowing you to assign parameter values to security levels. For more information about the template security, please see Section 5.1.10 - Template Security.

Preview Parameter Prompt: This option allows you to enable/disable parameter prompting when you preview the report.

5.1.3.5.7. Drill-Down Menu

The drill-down menu contains options that allow you to add, remove, and navigate layers of drill-down within the report. For more information about drill-downs, please see Section 5.1.8 - Drill-Down.

Navigate: This option brings up the drill-down navigation window allowing you to edit, add, and remove layers of drill-down.

Drill-Down Link: This option allows you to link the currently selected element to a drill-down layer.

5.1.3.5.8. Option Menu

This menu contains display and printing options

Report Explorer: This option turns on/off the report explorer display. For more information about the explorer interface, see Section 5.1.3.8 - The Report Explorer.

Page Setup: This option allows you to customize the page size and margins. The measurements will be in inches or centimeters depending on which unit is selected with the toggle button in upper left corner of the designer window where the rulers meet.

HTML Page Title: This option allows you to specify the page title that is generated when the report is exported to DHTML format.

Background Color: This option allows you to set the page background color for the report. Specifying a color here will color the entire page in Report Viewer and when exporting to DHTML or PDF. Any non-transparent section color or report element background color will draw over the background color.

Background Image: This option allows you to specify a background image for the report.

Snap to grid: This option allows you to control the settings for the snap to grid feature. You can turn the feature on/off, as well as control the step size.

Font Mapping: This allows you to map system (true type) font files for the PDF export.

Viewer Font Setup: This option allows the font size for text in the Report Designer and Report Viewer to be relative to the screen resolution. Turned on by default, this feature allows for more precise conversions of reports to various export formats and between installations.

Show Cell Outline: This option draws the outline for all the elements in the report. It allows you to see the boundaries of all elements, as well as see elements where the text may not be visible.

Shift On/Off: This option allows you to turn the cell shifting mode on or off. If the shift mode is on, resizing and moving report elements will cause the surrounding elements to shift: accommodating the resized or repositioned element.

Show Formula Name: This option allows you to enable/disable displaying the function name, rather than the text of the function in the Design window. For more information about working with formulas, see Section 5.1.6 - Using Formulas & the Formula Builder.

Enable Table of Contents: This option allows you to show/hide a table of contents (if it is inserted).

Global Format: This option allows you to set formatting for each of the different type of report elements. It will apply changes to elements currently in the report, as well as change the default properties.

Null Data Handler: This option allows you to specify the treatment of null data within the report.

5.1.3.5.9. Help Menu

This menu allows you to view program version and open the User's Guide.

About: This option shows you information about the program version.

Contents: This option opens the ERES Cloud User's Guide.

5.1.3.6. Designer Toolbar

The toolbar at the top of the designer offers easy access to ReportDesigner's most commonly used features. The first row of buttons perform the following functions:


Start a new report


Open an existing report


Save the current report


Export the current report


Apply a template to the current report


Undo the last change


Redo the last undone change

 Cut the selected element and place it on the clipboard

 Copy the selected element

 Paste the current clipboard object into the report

 Delete the selected element

 Insert a label

 Insert a formula

 Insert a column field

 Insert a rich text field

 Insert a sub-report

 Insert a horizontal line

 Insert a vertical line

 Insert a grid rectangle

 Insert a chart

 Insert an image

 Insert the current date

 Insert page number

 Change Data Mapping

 Change Data Source

 Edit Query

The first two boxes in the second row of the toolbar allows you to select font face and size for the currently selected element. Buttons in the second row perform the following functions:

 Set the font style to bold

 Set the font style to italic

 Underline the text

 Set the horizontal alignment to left

 Set the horizontal alignment to center

 Set the horizontal alignment to right

 Edit the selected element

 Set data format for the current element

 Apply a script to the current element

 Set font style, size and color for the current element

 Set background color for the current element

 Set border thickness, color and corner rounding for the current element

 Set dual colors for the current element

 Set bounds for the current element

 Page setup

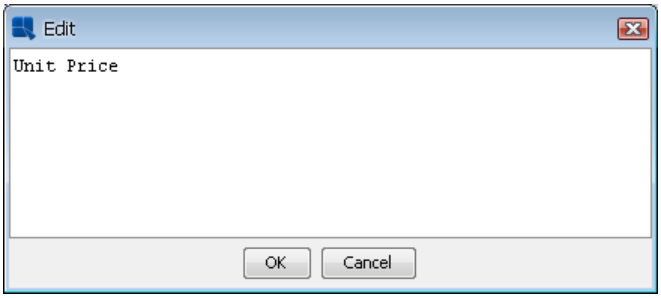
 Limit preview display rows

5.1.3.7. Inserting and Manipulating Report Elements

When you finish with the Report Wizard, a rough version of your report will be generated based on the mapping and it will display the options you selected. To polish the report, you may want to insert new report elements, as well as format the existing ones.

5.1.3.7.1. Inserting Elements

Labels: You can insert a label in one of two ways: by selecting *Insert Label* from the *Insert* menu or by clicking the  *Insert Label* button on the toolbar. After you select the *Insert Label* option, a small box will follow your mouse pointer around the Design window. Position the box where you want to insert the label and click. A dialog box will appear prompting you to enter the label text. Click the *OK* button and the label will appear in your report.



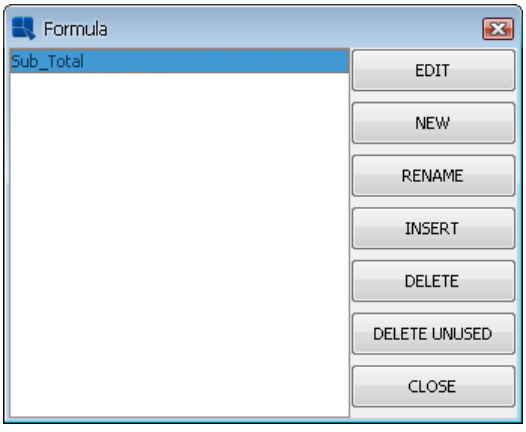
Insert Label Dialog

Formulas:

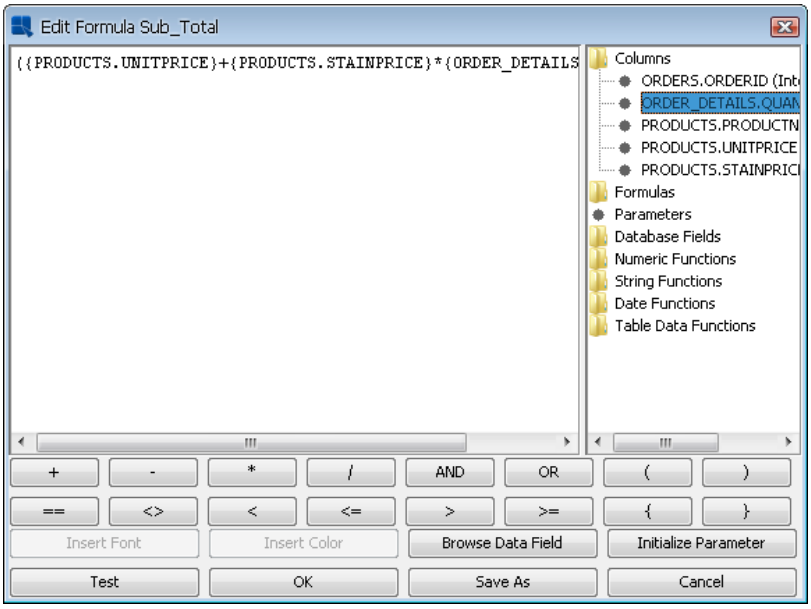
You can insert a formula in one of two ways: by selecting *Insert Formula* from the *Insert menu* or by clicking the *Insert Formula* button on the toolbar. After you select the *Insert Formula* option, a dialog box will open allowing to select a formula to insert. To insert a formula, select a formula from the list and click the *Insert* button. A small box will follow your mouse pointer around the Design window. Position the box where you want to insert the formula and click. The formula will be added. You can use an existing formula or create a new one. To create a new formula, click the *New* button in the formula list dialog. After entering a name for the new formula, the Formula Builder window will open allowing you to construct the formula (See Section 5.1.6 - Using Formulas & the Formula Builder for more information about using the formula builder and formula syntax). Any formula is automatically anchored to the report section in which it is inserted. This means that the formula will reset each time the section repeats.

Note

Only the text of the formula will appear in the design window (unless the formula is inserted in the Table Data section). You can choose to display the formula names instead of the text by selecting *Show Formula Name* from the Option menu.



Formula List



Formula Editor

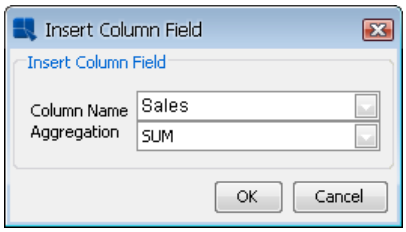
Column Fields:

You can insert a column field in one of two ways: by selecting *Insert Column Field* from the *Insert* menu or by clicking the  *Insert Column Field* button on the toolbar. After you select the *Insert Column Field* option, a dialog box will appear with two drop-down menus. The first one allows you to select the column field you want to insert into the report. The second one allows you to select aggregation (if any) for the column field. After you select the desired column field, click the *OK* button and a small box will follow your mouse pointer around the design window. Position the box where you want to place the column field and click. The column field will appear in the report. The column field is automatically anchored to the report section in which it is inserted. This means that each time the section repeats, the column field will display the next row in the data column.



Note

The column value will not appear in the Design window (unless it is placed in the Data Table section). All that will show is the column name `{name}`. You can see the column field value in the preview window.

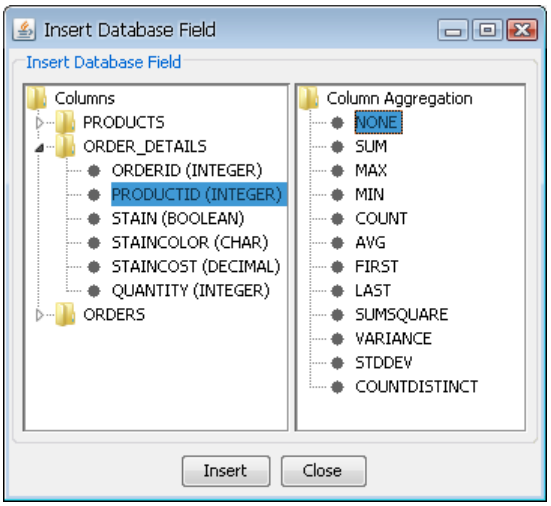


Insert Column Field

Database Field:

This option is only available if the current report uses a database as the data source. It allows you to add a field from the database into the report that may not have been initially selected by the query or as part of the data mapping. To insert a database field, select *Insert Database Field* from the *Insert* menu. This will open a dialog with all of the tables you have included in your query and their respective fields, allowing you to

select the field you want to add. Another option allows you to select the aggregation (if any) that you want to perform on the field.



Insert Database Field Dialog

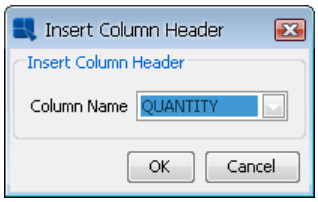


Note

If the template's query contains aggregation or a `Group by` clause, a third option will appear in the dialog prompting you to select `Database Aggregation` for the field. This can either be group by or any aggregation supported by the database.

Column Header:

Column headers are a unique report element that will dynamically display column header from the data source. You can add a column header by selecting *Insert Column Header* from the *Insert* menu. This will bring up a dialog prompting you to select the column for which you want to retrieve the header.

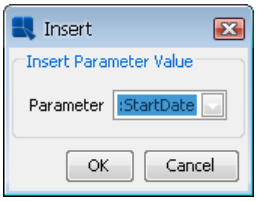


Insert Column Header Dialog

Select a desired column and click the *Ok* button. A box will then allow you to position the header wherever you want in the report. Click to add the column header.

Parameter Values:

You can insert user supplied values to query or formula parameters directly into the report. To insert a parameter value, select *Insert Parameter Value* from the *Insert* menu. This will bring up a dialog prompting you to select which parameter you want to display the value for.



Insert Parameter Value Dialog

Parameter values are also accessible through formulas. For more information about formulas, please see Section 5.1.6 - Using Formulas & the Formula Builder.

Images:

You can insert an image in one of two ways: by selecting *Insert Image* from the *Insert* menu or by clicking the  *Insert Image* button on the toolbar. After you select this option, a small box will follow your mouse pointer around the Design window. Position the box where you would like to place the image and click. A new window will appear prompting you to enter the URL of the image you want to insert.

There are two ways of inserting images: either locate the image on your hard drive or retrieve it from an URL.

To locate an image on your hard drive, click the *Browse* button. If you want to use an image from a different location, insert its URL.

To retrieve image from an URL, enter the URL in the *Image URL* text field. After that, click the *Refresh Preview* button to verify the URL. If the image from the URL appears in the *Preview* section, the URL is correct.



Note

You have to insert complete URL with protocol (e.g. `http://`).

If you save the report as PAK, the image will be stored in the PAK file along with the report.

If you want to re-open a RPT file (from an older ERES version) with a background image, please note that the image itself is not stored within the report. Only the path or URL is saved. If you move a RPT report, you need to be sure that it can still access the image using the specified path.

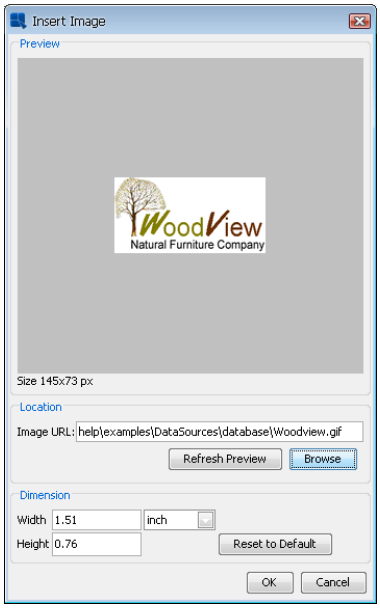
You can manually set the image dimensions in pixels, inches, or millimeters in the *Dimension* section.



Note

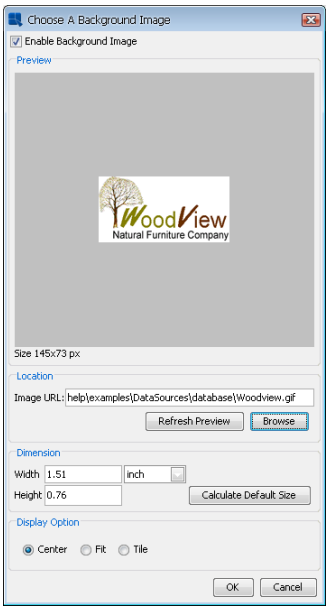
If the image dimension exceeds the page dimension, the image will not be shown (although space will be allocated for that image).

ReportDesigner also supports database images. You can add images from a database (BLOB format) by selecting the BLOB field as part of the query for the report. The images will be added as a column in the report. You can also retrieve images in the Table Data section via URL instead of using a BLOB field. For more information about this, please see Section 5.1.3.7.3 - Data Formatting for Formulas and Column Fields.



Insert Image Dialog

Background Images: ReportDesigner allows you to add background images to a report. Background images will underlay all other report elements and will repeat on every page. To add a background image, select *Background Image* from the *Option* menu. This will bring up a dialog prompting you to specify location of the image, as well as several display options.



Background Image Dialog

To insert a new image, select the *Enable Background Image* option. If you want to remove an existing background image and use simple background color instead, unselect this option.

This dialog also allows you to choose one of the *Display options* (*Center*, *Fit* and *Tile*).

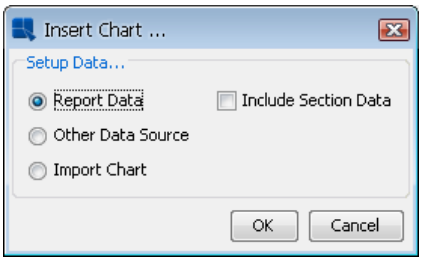
The rest of this dialog is the same as the *Insert image* dialog described in the previous section.

Lines/Rectangles: To insert a line or a grid rectangle, select either *Insert Line* or *Insert Rectangle* from the *Insert* Menu. If you select a line, you can further select whether you want to insert a vertical or horizontal line. You can also select the corresponding icons on the toolbar. After you select a line or rectangle, a cross or a box will follow the mouse pointer around the Design window. Position the cross/box where you would like to insert the line/rectangle and click, then drag the cross to draw the line or rectangle. A line or grid will then be inserted into the report.

> **Note**
>
> The line or grid will appear thicker than it actually is in the Design window.

5.1.3.7.1.1. Inserting Charts

You can insert a chart in one of two ways: by selecting *Insert Chart* from the *Insert* menu or by clicking the  *Insert Chart* button on the toolbar. After you select this option, a small box will follow your mouse pointer around the design window. Position the box where you want to place the chart and click. You will see the following dialog box:



Insert Chart Dialog

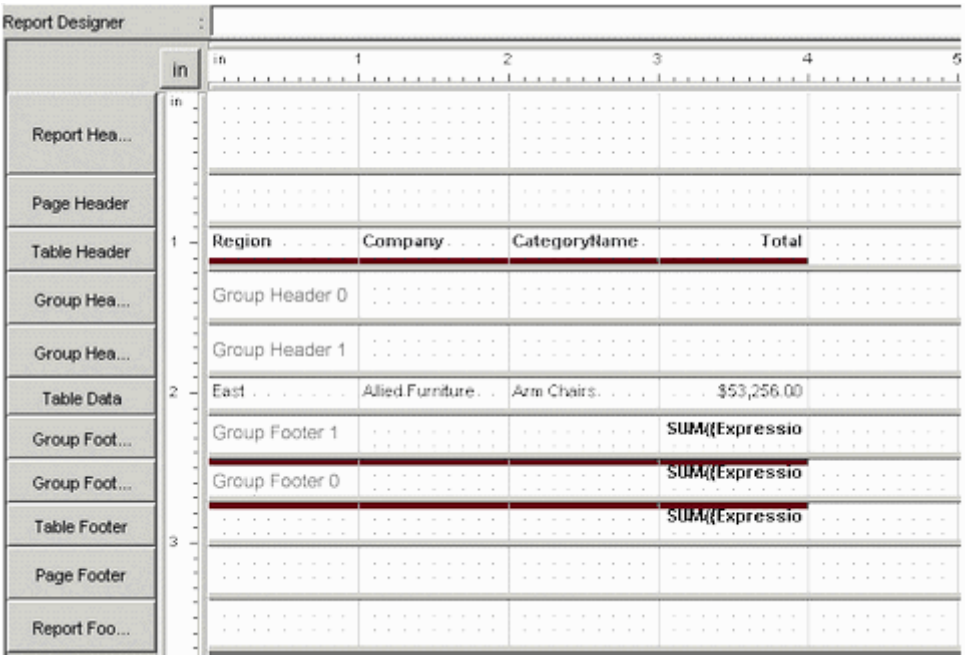
You have an option to use data from the current report or use an external data source. If you choose to use the data from the report, you will be shown a preview of your dataset and you will be taken directly to the chart selection window in the Chart Designer. Otherwise, you will be given an option to either select a different data source for your chart or import a chart file created before.

Once you select your options, the Chart Designer will open in a new window, allowing you to design and customize your chart. Please see the for more information about using the Chart Designer.

> **Note**
>
> The actual chart will not appear in the Design window. Instead, you will see a gray rectangle with the chart URL. You can see the actual chart in the preview window.

If your report contains grouping (Summary Break, CrossTab, or Master & Detail Reports), you will also have the option to *Include Section Data*. If you select this option, you will be able to use group aggregations that are located within any section that the chart encompasses. This means that you can create charts using data from any aggregation located in the same section as the chart or in any inner sections.

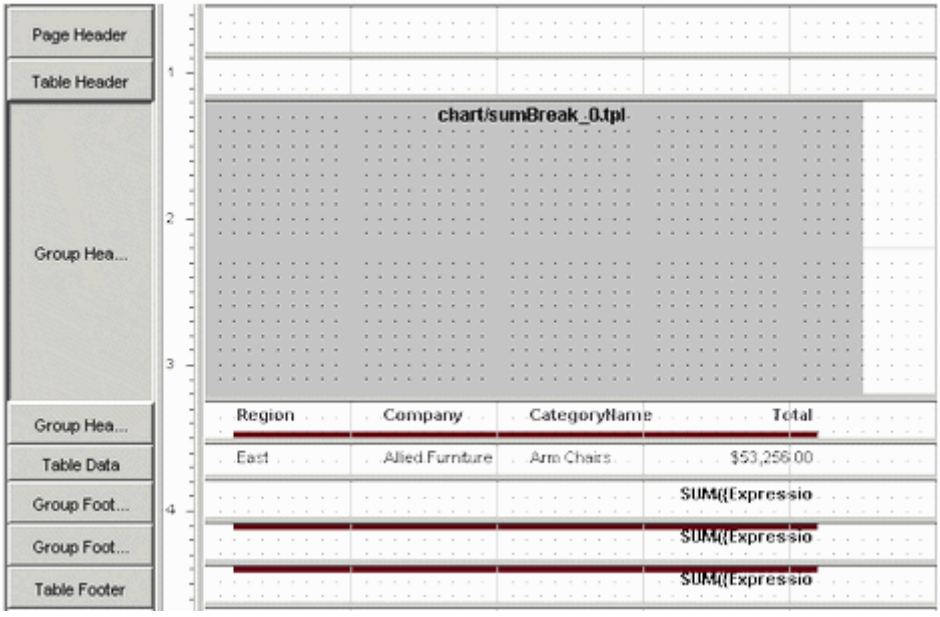
For example, if you were to create a Summary Break Report with two row break columns, this will result in two group header and two group footer sections. Assuming we have a group aggregation for a column in your report, the report might look like this:



Summary Break Report with Two Row Break Columns

If you position the chart in the inner group header (Group Header 1) or inner group footer (group footer 1), you will only be able to use the group aggregation for the inner section in your chart (the aggregation in Group Footer 1). If you positioned the chart in the outer group header (Group Header 0) or outer group footer (Group Footer 0), you will be able to use the aggregation for the outer section as well as the inner section in your chart because the outer section encompasses the inner section. Positioning the chart in the table header or table footer sections will allow you to use any group aggregation in the report.

Suppose you moved the column headers to Group Header 1 and positioned the chart in Group Header 0 as shown below.



Insert Chart in Outer Group Header

When you see the data mapping window in Chart Designer, you will be able to use the summary data from both Group Footer 0 (TBL0_F_SEC0_FORM0) and Group Footer 1 (TBL0_F_SEC1_FORM0).

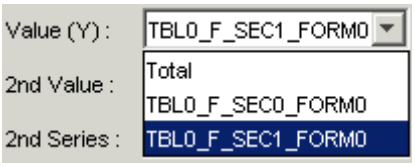
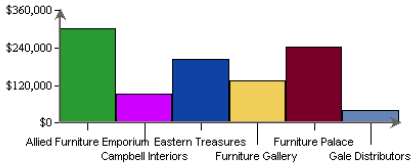


Chart Mapping with Section Data

To make the aggregation easier to recognize, you can use Custom IDs. For more information about using Custom IDs, see Section 5.1.6.2.1 - Using Column Field Data.

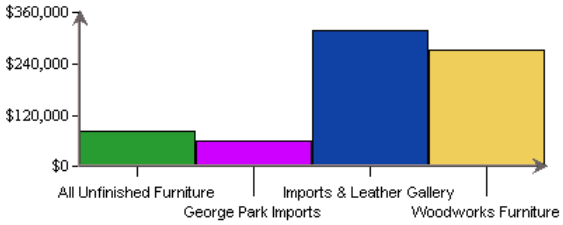
The resulting report will look like the image below. Notice that the chart plots the aggregation for each company. If you used the aggregation from Group Footer 0, the chart would have shown the aggregation for each region instead.



Region	Company	CategoryName	Total
East	Allied Furniture	Arm Chairs	$53,256.00
		Double Dressers	$58,144.00
		Oval Tables	$52,142.00
		Round Tables	$23,805.00
		Side Chairs	$23,540.00
		Single Dressers	$92,039.00
			$302,926.00
	Company	**CategoryName**	**Total**
	Campbell Interiors	Arm Chairs	$15,609.00
		Oval Tables	$11,394.00
		Rectangular Tables	$7,290.00
		Round Tables	$42,916.00
		Triple Dressers	$15,936.00
			$93,145.00
	Company	**CategoryName**	**Total**
	Eastern Treasures	Arm Chairs	$73,505.00
		Oval Tables	$3,750.00
		Rectangular Tables	$41,511.00

Report with Chart Using Section Data (East)

Since the chart is placed in Group Header 0, a new chart will be displayed for every region. The above image shows the chart for the East region only, but the report will contain charts corresponding to each of the other regions as well. Here is the chart for the Midwest region:



Region	Company	CategoryName	Total
Midwest	All Unfinished	Arm Chairs	$14,532.00
		Oval Tables	$27,884.00
		Rectangular Tables	$7,530.00
		Round Tables	$11,109.00
		Side Chairs	$9,903.00
		Single Dressers	$13,839.00
			$84,797.00
	Company	**CategoryName**	**Total**
	George Park	Arm Chairs	$5,424.00
		Oval Tables	$13,424.00
		Rectangular Tables	$5,416.00

Report with Chart Using Section Data (Midwest)

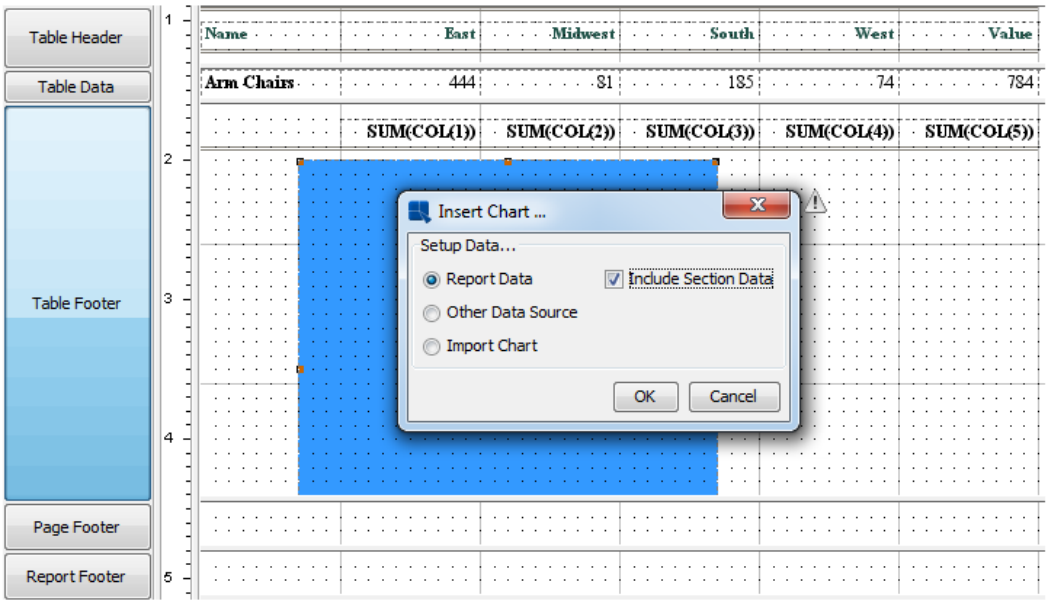
Keep in mind that if you select the *Include Section Data* option, the data available for the chart will vary between sections, so moving the chart between sections is not recommended. If you need to move the chart after it is made, make sure that the data mapped to the chart is also available from the location where you are moving the chart to.

Here is another example using *Include Section Data*. Suppose you had the following CrossTab report and you wanted to use the aggregation for each column as data points in your chart.

Name	East	Midwest	South	West	Value
Arm Chairs	444	81	185	74	784
Double Dressers	68	12	27	0	107
Oval Tables	79	33	24	16	152
Rectangular	85	82	67	52	286
Round Tables	143	109	72	46	370
Side Chairs	150	216	377	100	843
Single Dressers	79	100	114	18	311
Triple Dressers	21	45	19	41	126
	1 069	**678**	**885**	**347**	**2 979**

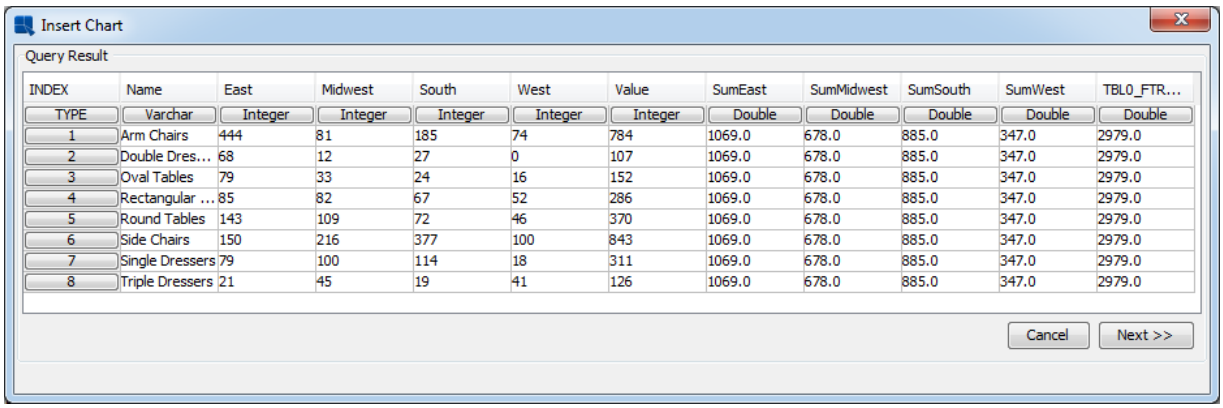
CrossTab Report with Group Aggregation

To make the aggregations easier for later use, give each group aggregation a Custom ID that matches their column header. For example, give `SUM(COL(1))` the ID **`SumEast`**, give `SUM(COL(2))` **`SumMidwest`** and so on. For more information about Custom IDs, see Section 5.1.6.2.1 - Using Column Field Data. Then insert the chart in the Table Footer Section using Report Data and Include Section Data.



Insert Chart in Table Footer

The formulas with custom IDs will be visible in the data source preview (if you selected the *Include section data* option).



Data source preview with section data

Since the data points are in different aggregations, you will need to use the transpose feature. To learn more about data transposition in charts, see the Section 5.2.3.1.1 - Data Transposition chapter.

In the data mapping window, select the *Multi Selection* option for the Category (X). This option allows you to select multiple columns for the category. Select all the formulas with custom IDs. The Value (Y) should be automatically set to *Value*.

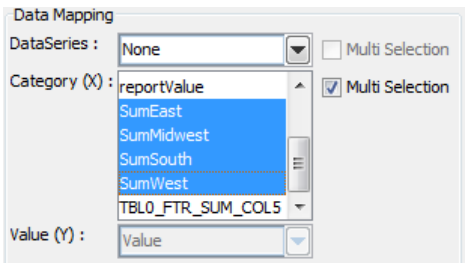
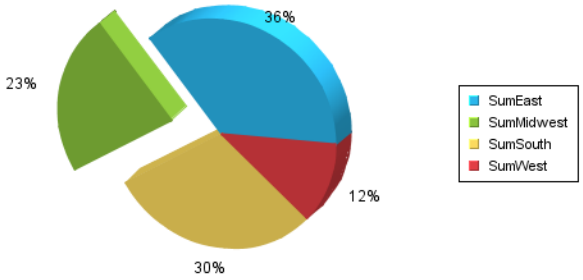


Chart Mapping for Transposed Data

Depending on the type of chart you select, the finished report might look like this:

Name	East	Midwest	South	West	Value
Arm Chairs	444	81	185	74	784
Double Dressers	68	12	27	0	107
Oval Tables	79	33	24	16	152
Rectangular	85	82	67	52	286
Round Tables	143	109	72	46	370
Side Chairs	150	216	377	100	843
Single Dressers	79	100	114	18	311
Triple Dressers	21	45	19	41	126
	1 069	**678**	**885**	**347**	**2 979**

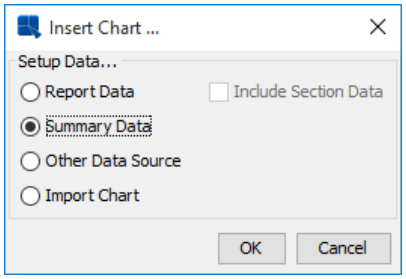


Report with Chart Using Section Data

5.1.3.7.1.1.1. Inserting Summary Charts into Crosstab Reports

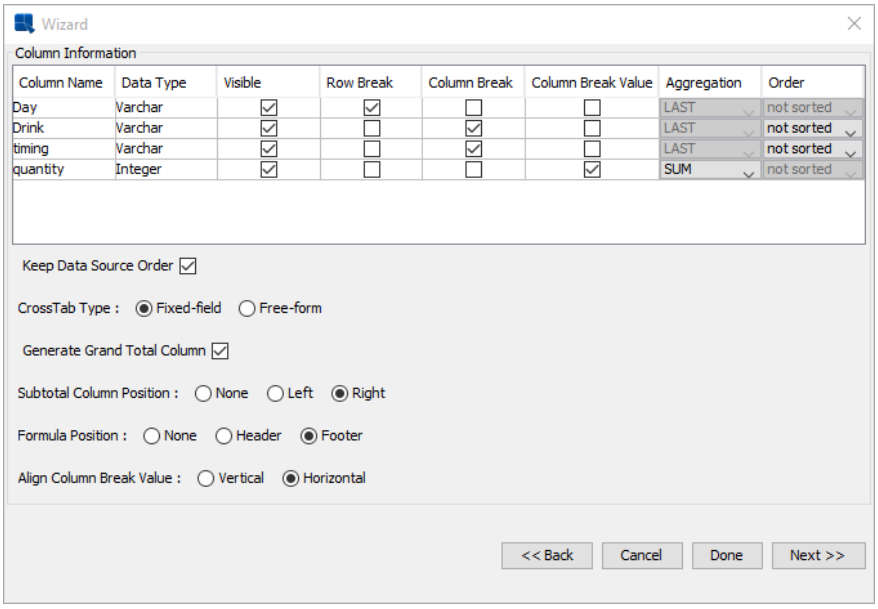
You can also insert summary charts into crosstab reports. A summary chart is a type of chart that is created from summarized data of a fixed-field crosstab report. For more information about crosstab reports, please see Section 5.1.2.3 - Crosstab Report.

When you are going to insert a chart into a crosstab report, you will still have the option to use data from the current report or use an external data source. Otherwise, you will be given an option to select a summary data source for your chart or import a chart file created before. Note that a summary chart can only be inserted into the Table Header/Footer sections, depending on the Formula Position. For example, if you created a fixed-field crosstab report with the formula position set to be drawn in the Header section, you can only insert a summary chart into the Table Header section. If you try to insert a chart to other report sections, the *Summary Data* option will not be available in the Insert Chart dialog.



Insert Chart Dialog

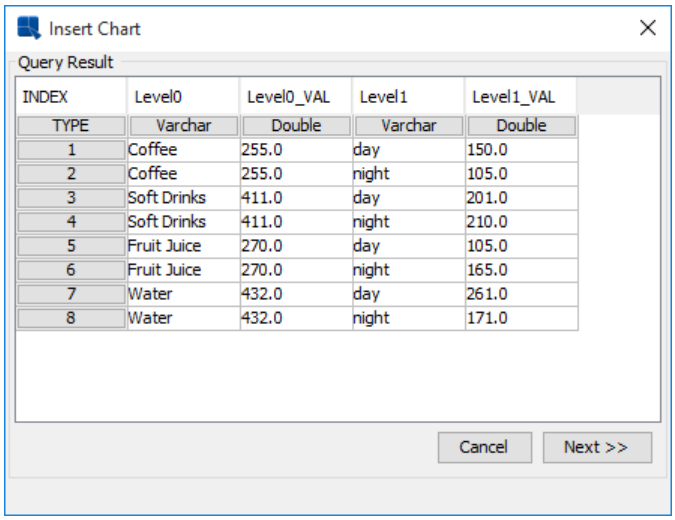
For example, say we have the following data mapping for a crosstab report:



Data Mapping

We selected `Day` as a *Row Break* field, `Drink` and `Timing` as *Column Break* fields, `Quantity` as a *Column Break Value* field with an aggregation of **`Sum`**. We also chose the *Formula Position* to be drawn in the Table Footer section of the report, which means we have to insert the summary chart into the Table Footer section in which data will be summarized.

Inserting a summary chart into the table footer section and selecting the **`Summary Data`** as datasource in the *Insert Chart* dialog will show the following Query Result dialog. The dialog shows data that will be used for a summary chart.



Query Result

Depending on the type of chart and formatting you select, the finished report might look like this:

	Coffee			Soft Drinks		
Day	day	night	SUM (Coffee)	day	night	SUM (Soft Drinks)
Mon	3	30	33	24	24	48
Tue	60	12	72	27	27	54
Wed	6	24	30	30	15	45
Thu	18	3	21	24	36	60
Fri	54	33	87	60	36	96
Sat	6	0	6	27	21	48
Sun	3	3	6	9	51	60
	150	**105**	255	201	210	**411**

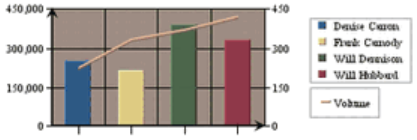


Report Preview

5.1.3.7.1.2. Inserting Elements into the Table Data Section

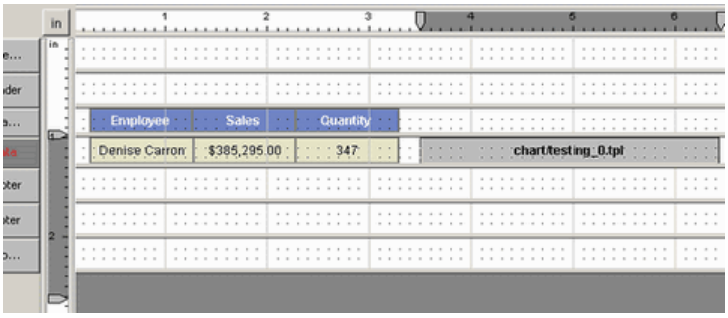
Generally, when elements are inserted into a report section, they will repeat each time that the report section repeats. However, this does not apply when certain elements are inserted into the Table Data section. Charts, images, and sub-reports do not repeat for each row of data (labels, column fields, and formulas do) when inserted into the Table Data section. This allows you to design a report in a side-by-side configuration where you might have a chart or in sub-report placed next to a data table as pictured below.

Employee	FY 2000	
	Sales	Volume
Denise Carron	$255,302.00	225
Frank Carnody	$217,890.00	333
Will Dennison	$390,813.00	371
Will Hubbard	$332,996.00	419



Side-by-Side Data Table & Chart

Normally, when you insert a chart, image, or sub-report into a report section, you will want to resize the section so that the entire element fits within the section (otherwise it is truncated). This is not the case for charts, images, or sub-reports inserted into the Table Data section. In this situation, you will want to keep the section the same height as your column fields regardless of the size of the element you're placing in the section.



Side-By-Side Data Table & Chart (Designer)

You can resize an element in the Table Data section using the rulers or by temporarily resizing the section to adjust the element size.

5.1.3.7.2. Editing Elements

You can edit elements in one of four ways: by selecting *Edit* from the *Edit* menu, clicking on the  *Edit* button in the toolbar, right clicking on the element, and then selecting *Edit* from the pop-up menu, or by double clicking on the element.

Editing Labels: When you select the *Edit* option for a label a dialog box will appear prompting you to change the label text. Click on *OK* and the label will be changed.

Editing Formulas: When you select the *Edit* option for a formula, the Formula Builder will appear prompting you to change the formula. Make any changes you want, then click the *OK* button and the formula will be changed.



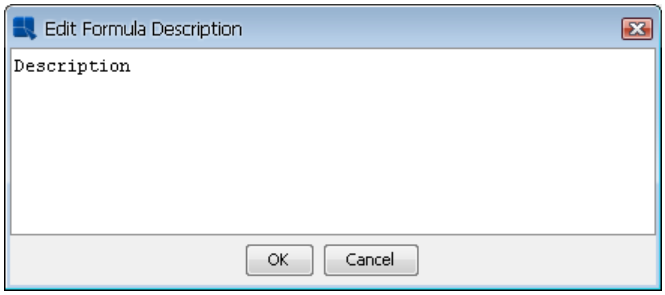
Note

If you use the same formula in more than one place in the report, modifying the formula will change it everywhere. You can avoid this by clicking the *Save As* button in the Formula Builder. This will allow you to specify a new name for the modified formula and modify it only for the particular element.

Editing Charts: When you select the *Edit* option for a chart, the Chart Designer will open, allowing you to edit and format the chart.

Editing Images: When you select the *Edit* option for an image, a new window appears prompting you to select a new image or directory URL. After selecting the new image, click the *OK* button and the new image will replace the old one.

Editing Column Headers: Column header is a unique type of report element that is generated when you first create the report. By default, column headers are dynamic and will display the name of the column even when the column changes (by changing data mapping, source, or applying a template). When you select to edit a column header, a dialog box will appear prompting you to change the text of the header.



Edit Column Header Dialog

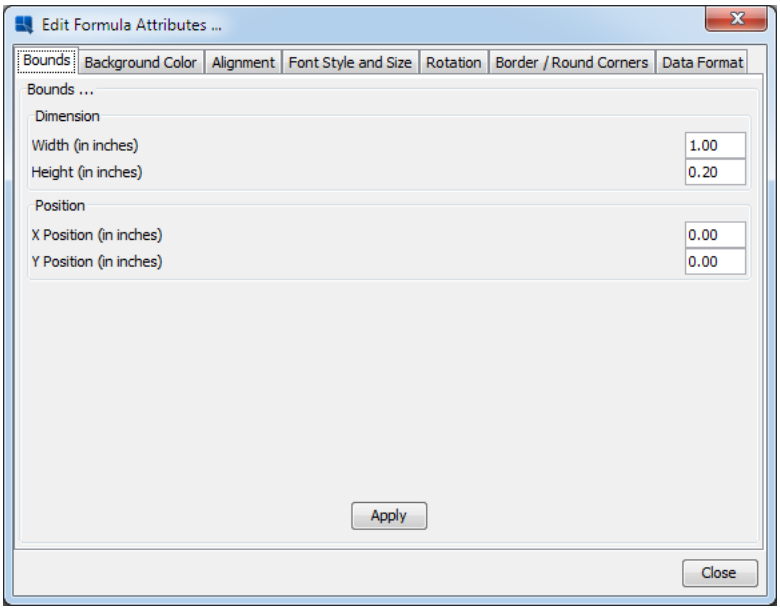
When you edit the header, a new text will be displayed. However, the header will no longer be dynamic. You can return the header to display the (dynamic) column name by right clicking on it and selecting *Original Column Header* from the pop-up menu.

You can also copy and paste cell attributes between report elements. This will apply all attributes including data formatting (assuming the data types are the same), alignment, font, color, and border attributes, as well as the bounds of a cell.

To copy the attributes of an element, right click on the element and select *Element Appearance* from the pop-up menu. This will expand into two additional choices. Select *Copy* from the second menu. To apply the copied attributes to another cell, repeat the same steps as before and select *Paste* from the secondary menu.

5.1.3.7.2.1. Editing Element Attributes

In addition to directly editing an element, you can directly access all of the element's attributes by selecting *Edit Attributes* option from the *Format* menu, or by right clicking on an element and selecting *Edit Attributes* from the pop-up menu. This will bring up a tabbed dialog that allows you to set a number of properties for the element at once.

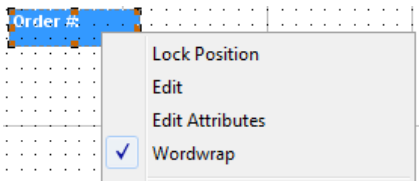


Multi-Attribute Editing Dialog

The options will vary depending on what type of element you select. For each tab in the dialog, you can set the option you want and then click the *Apply* button to set the changes for that attribute.

5.1.3.7.2.2. Word Wrapping

By default, any text within report elements, including labels, formulas, and columns will wrap to the next line if the cell boundaries are not wide enough to fit the text. However, this behavior can be disabled. To set text wrapping, select an element and select Format → Wordwrap, or right click on an element and select the option from the pop-up menu. This will disable/enable text wrapping.

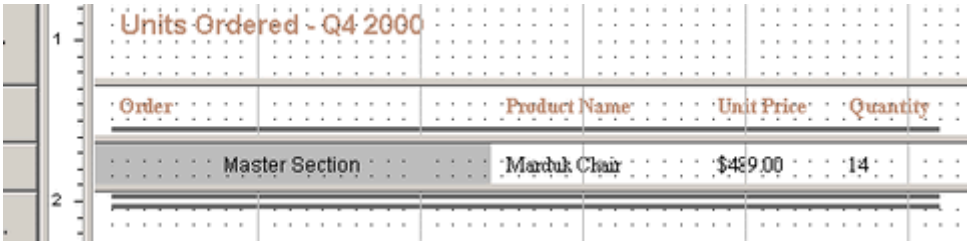


Word Wrapping Dialog

To disable word wrapping, un-check the option. Note that when word wrapping is disabled, the entire contents of the cell will be printed regardless of the cell boundaries. This can overlap and obscure other report elements. Also, word wrapping cannot be disabled for rich text fields as the rich text and Excel-based exports will still display the wrapped text.

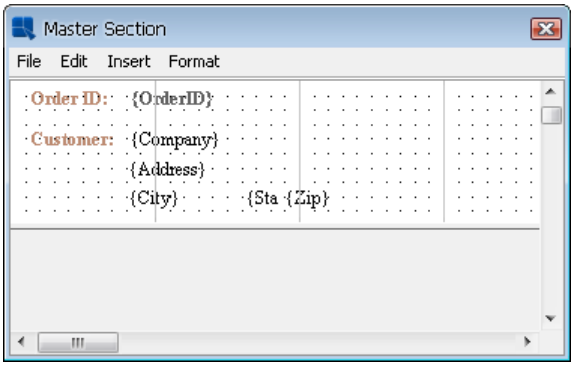
5.1.3.7.2.3. Editing Side-By-Side Master & Details Reports

As noted in Section 5.1.2.4 - Master & Details Report of this guide, you can specify a side-by-side layout for master & details reports. This will print the master field (group header) next to the details section (table data), rather than its traditional position above the details section. When you select this report format, the master section will appear as a gray rectangle next to the column fields within the table data section of the report.



Side-by-Side Master & Details

The master section can be moved in free-form like any other report element. You can edit the contents of the master field in one of three ways: selecting it and clicking the *Edit* button on the toolbar, right clicking it and selecting *Edit* from the pop-up menu, or double clicking it. The master section will open in a new window, allowing you to modify its contents.



Edit Master Section Dialog

The elements in this section can be moved, resized, and edited in the same way as other elements in the report. The master section acts the same as a group header section. It repeats for each unique value in the primary key column. After you finish editing the contents of the master section, it can be closed by selecting *Close* from the *File* menu.



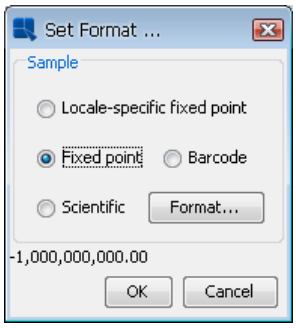
Note

After you finish editing the master section, it will automatically resize to fit its contents.

5.1.3.7.3. Data Formatting for Formulas and Column Fields

You can change the output of formulas and column fields in one of three ways: clicking the  *Data Format* icon in the toolbar, selecting *Data Format* from the *Format* menu, or right clicking on a formula or column field and selecting *Edit Attributes* from the pop-up menu. When you select the format option, a dialog box will appear (or the format tab of the multi-attribute editing dialog). The box that appears depends on what type of data is present in the selected element: numeric, string, date/time, or logical/boolean.

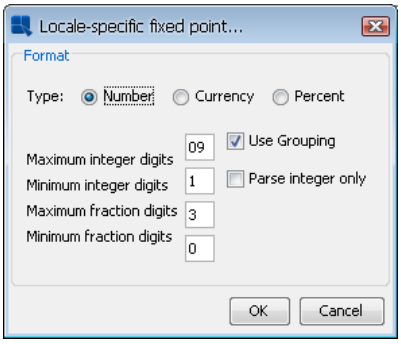
Formatting Numeric Data: The dialog box for numeric data contains four primary options for the data: locale-specific fixed point, fixed point, bar code, and scientific. You can select the option you want and then click on format for additional options.



Numeric Format Dialog

Locale-Specific Fixed Point:

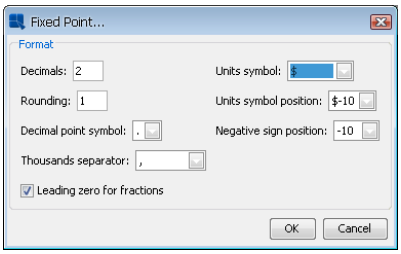
This will change the data format depending on the locale in which it is being viewed. Additional formatting for this option allows you to specify whether the data should be displayed as a number, currency, or percentage. Additionally, you can set the maximum and minimum number of integer digits and fraction digits. Other display attributes will vary depending on locale.



Locale-Specific Formatting

Fixed Point:

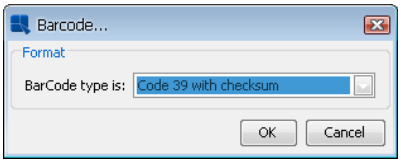
This will keep the data format consistent, regardless of locale. Additional formatting for this option allows you to set the number of decimals, rounding for digit number, unit symbols, negative sign positions, decimals and thousands separator, and specify leading zeroes for fractions.



Fixed Point Formatting

Bar Code:

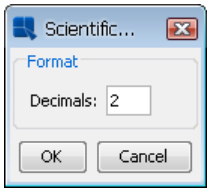
This will convert the data into a bar code. Additional formatting for this option allows you to select the symbology to use for the bar code. Supported symbologies are Code 39 with and without checksum, Interleave 2 of 5, UPC A, EAN 13, EAN 128, Standard 2 of 5 with and without check digit, Code 128, Code 128A, Code 128B, Code 128C, USD 3 with/without checksum, Code 3 of 9 with and without checksum, Global Trade Item Number, Random Weight UPCA, SCC 14 Shipping Code, Shipment Identification Number, SSCC 18, and US Postal Service.



Bar Code Formatting

Scientific:

This will display the data in scientific notation. Additional formatting for this option allows you to set the number of decimals.

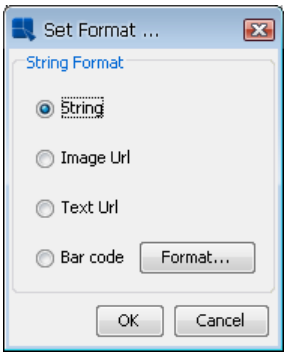


Scientific Formatting

After you finish selecting additional options, click the *OK* button to return to the main dialog box. Click the *OK* button and the data format will be changed.

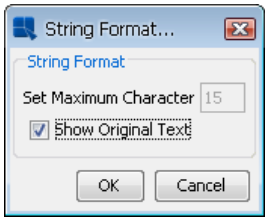
Formatting String Data:

The dialog box for string data contains four primary options: string, image URL, text URL, and bar code. You can select the option you want and in the case of string or bar code, click *Format* for additional options.



String Data Format

String: This dialog allows you to format the appearance of the string. The checkbox labeled *Show Original Text* controls whether the complete string for each data entry will be displayed. If you un-check it, you can then specify the maximum number of characters to be displayed. Click the *OK* button to return to the previous dialog.

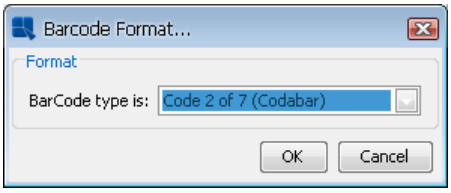


String Formatting

Image URL: This will convert the string into an image. This is used in situations where instead of storing an image as an object in a database (BLOB), you have stored URLs that point to image files on a server. Selecting this option will read the URL and retrieve the images to be placed in the report.

Text URL: This will convert the string into a large text object. This is used in situations where instead of storing a large text file as an object in a database (CLOB), you have stored URLs that point to files on a server. Selecting this option will read the URL and retrieve the text files to be placed in the report.

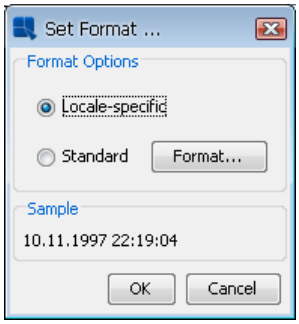
Bar Code: This will convert the data into a bar code. Additional formatting for this option allows you to select the symbology to use for the bar code. Available symbologies for string data are Code 39 with/without checksum, Code 2 of 7 (Codabar), Codabar, EAN 128, Code 128, Code 128A, Code 128B, Code 128C, Global Trade Item Number, Monarch, NW 7, SCC 14 Shipping Code, Shipment Identification Number, SSCC 18, US Postal Service, and USD 4.



Bar Code Formatting

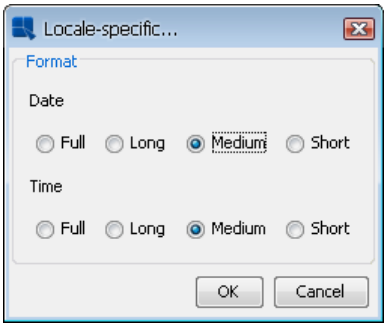
> **Note**
>
> Codabar will not accept start/stop characters in the data. If the input data is incorrect, the bar code may not be readable.

Formatting Date/Time Data: The dialog box for date/time data contains two primary options for the data: locale-specific and standard. You can select the option you want and click the *Format* button for additional options. The available additional options will vary depending on the nature of your data. Date, time, and timestamp data will bring up date, time, and date & time options respectively.



Date/Time Data Format

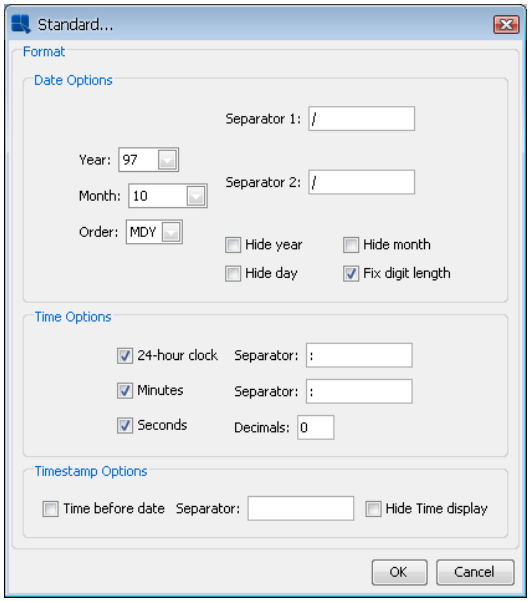
Locale Specific: This will change the format of the data depending on the locale in which it is being viewed. Additional formatting for this option allows you to select full, long, medium, or short notations for date and time information. Other display attributes will vary depending on locale.



Locale-Specific Formatting

Standard: This will keep the data format consistent, regardless of locale. Additional formatting for this option allows you to select year and month displays,

as well as the order in which month, day, and year information is presented. You can also select the characters to be used as separators. Time options allow you to display hours, minutes, and/or seconds, and select the separators between them. For timestamp data, you can select to display the time before or after the date and the separator to be used between them.



Standard Formatting

After you finish selecting additional options, click the *OK* button to return to the main dialog box. Click the *OK* button and the data format will be changed.

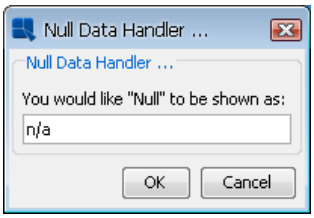
Formatting Logical Data: The dialog box for Logical or Boolean data contains five options for displaying the data. They are: T/F, True/False, Yes/No, Y/N, and 1/0. Select the format you want to use and then click the *OK* button to change the data format.



Boolean Data Format

5.1.3.7.3.1. Formatting Null Data

By default, null data will display in a report as `Null`. You can change the appearance of null data by selecting *Null Data Handler* from the *Option* menu. This will bring up a dialog, prompting you to specify a string to be displayed for nulls.



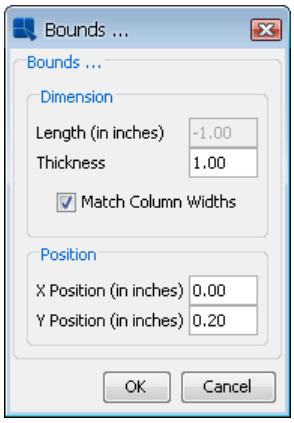
Null Data Handler

Enter a value you want and click the *OK* button. All the null values in the report will be displayed as the value you specified.

5.1.3.7.4. Line/Rectangle Format

You can format the appearance of lines and rectangles in the report by right clicking on a line or rectangle element. The pop-up menu will contain a list of available options.

Lines: There are three formatting options for lines. *Line Color & Thickness* allows you to set the color and thickness of the line. *Line Style* allows you to select the style of line - either solid, dotted, or double. *Set Bounds* allows you to specify the length and thickness of the line, as well as the X and Y coordinates of its origin. In addition, you can set the width of horizontal lines to match the aggregate width of columns in the report. This feature is useful for crosstab reports where the number of columns can vary and also for user-defined report styles. For vertical lines, you can set the length to match the section height.



Bounds Dialog for Horizontal Lines

Rectangles: There are four formatting options for grid rectangles. *Border* allows you to set the border thickness and color for the grid rectangle. *Background* allows you to set a background color for the fill area of the rectangle. *Set Bounds* allows you to specify the width and height of the grid rectangle, the X and Y coordinates of its origin, as well as the degree of rounding for the corners. *Border Style* allows you to select the line style of the border - either solid, dotted, or double.

5.1.3.7.5. Chart Export Format

This option allows you to specify attributes for charts exported to DHTML. You can format the chart export properties by selecting *Chart Export Format* from the *Format* menu or by right clicking on a chart and selecting *Export Format* from the pop-up menu. When you select this option, a dialog box will appear. You can select GIF, PNG, JPEG, or FLASH for the image type and set some options for each type. For GIF images, you can set the background to be transparent. For PNG images, you can set the compression. For JPEG images, you can set the quality (Image quality is directly proportional to file size). For FLASH images, you can enable animation and set the frame count and frame rate.

You can also specify to create an image map when the report is exported. If you create an image map, the hyperlinks defined in the chart will be active when the report is exported. If no hyperlinks are defined, the map will contain

display pop-up labels for the chart containing data point information. Note that FLASH images do not support image maps.

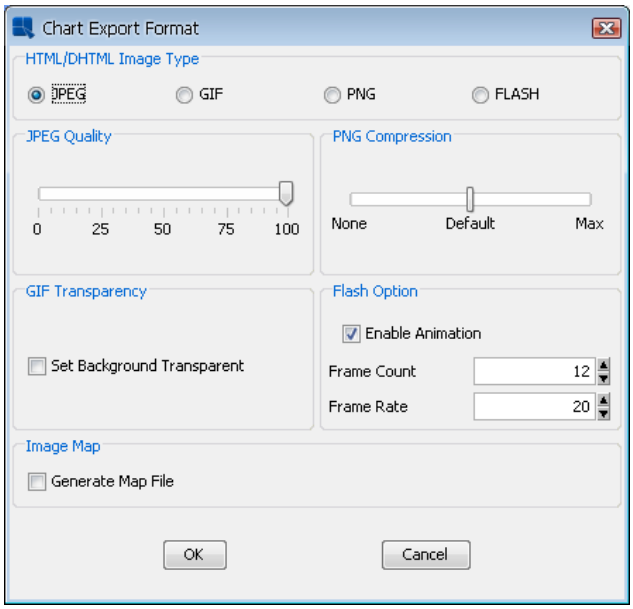
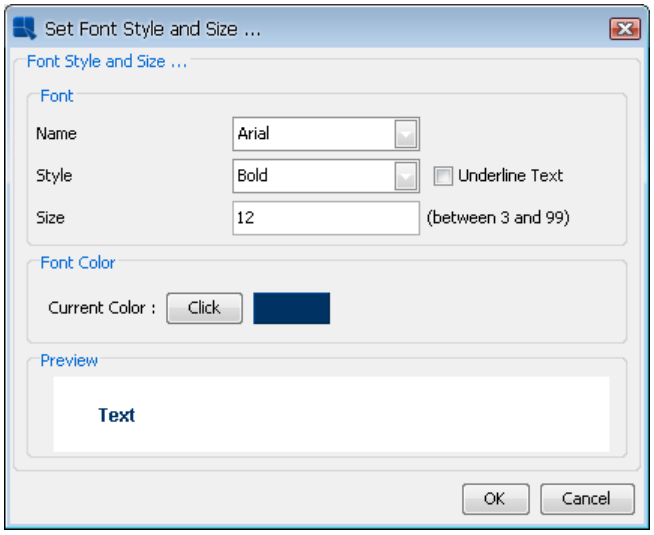


Chart Export Format

5.1.3.7.6. Font, Color, and Border Options

You can change the appearance of report elements by changing the font, color, and border properties.

Font Style and Size: The font style and size can be adjusted for labels, formulas, and column fields in one of three ways. You can directly modify the font style and size using the options in the Report Designer toolbar. In addition, you can also select *Font Style and Size* from the *Format* menu or right click on a report element and select *Edit Attributes* from the pop-up menu. The latter two options will invoke a separate dialog prompting you to select the font style and font size or in the attribute editing dialog, where font options can be set in the 'Font Style and Size' tab.



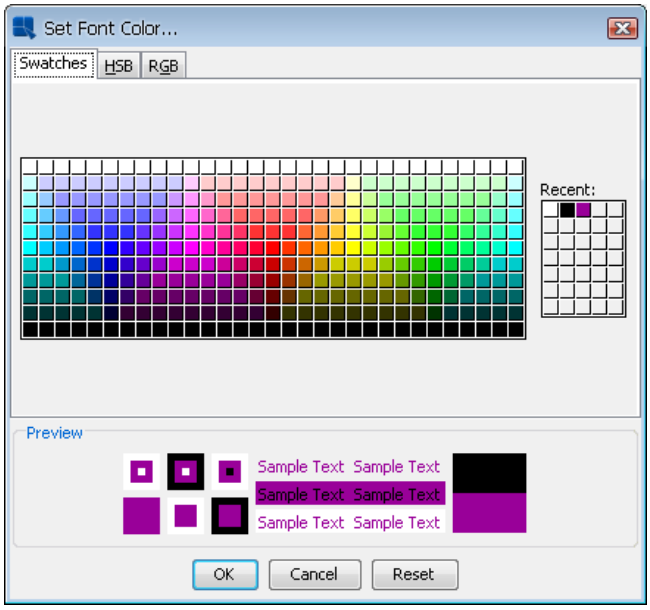
Font Style and Size Dialog

Alignment: The alignment for labels, formulas and column fields can be adjusted in one of three ways. You can directly adjust the horizontal alignment for an element by selecting the option(s) from the toolbar. You can also select *Alignment* from the

Format menu, or right click on the element and select *Edit Attributes* from the pop-up menu. The latter two options also allow you to select vertical as well as horizontal alignment for the element text.

Font Color:

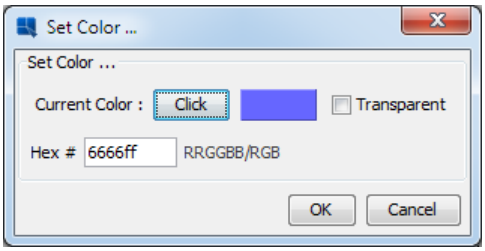
The font color can be adjusted for labels, formulas, and column fields in one of three ways: selecting *Font Style And Size* from the *Format* menu, clicking the  *Font Style And Size* button on the toolbar, or right clicking on the element and selecting *Edit Attributes* from the pop-up menu. Selecting this option will bring up a dialog box (or the multi-attribute editing dialog where options can be set in the *Font Style And Size* tab). To change the font color, click the *Click* button next to the *Current Color:* field. The Set Font Color dialog will appear allowing you to select font color from swatches. You can also enter HSB, HEX, or RGB values.



Font Color Dialog

Background Color:

The background color can be adjusted for labels, formulas, and column fields in one of three ways: selecting *Background Color* from the *Format* menu, clicking the  *Background Color* button on the toolbar, or right clicking on the element and selecting *Edit Attributes* from the pop-up menu. Selecting this option will bring up a dialog box (or the multi-attribute editing dialog where options can be set in the *Background Color* tab) prompting you to specify whether or not to set the background transparent. If you do not want a transparent background, uncheck the checkbox and click the button. This will bring up a dialog prompting you to specify the background color. You can select font color from swatches, or enter HSB or RGB values. You can also directly enter HEX color value to the *Hex #* field.

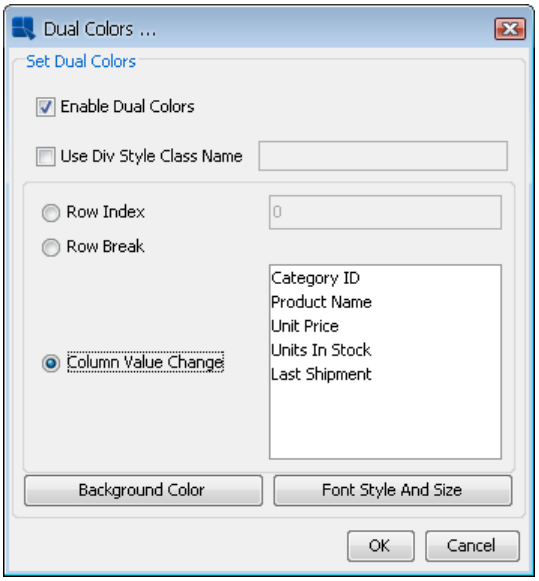


Background Color Dialog

Dual Colors:

You can specify alternating color for labels, formulas, or column fields that are in the data table section of the report in one of three ways: selecting *Dual Colors* from the *Format* menu, clicking the  *Dual Color* button on the toolbar, or right clicking on the element and selecting *Edit Attributes* from the pop-up menu. Selecting this option will bring up a dialog box (or the multi-attribute editing dialog where options can be set in the *Dual Colors* tab), allowing you to specify the number of rows between alternating colors. Instead of specifying alternate row numbers, you can also set dual colors to change on the row break (for summary break or crosstab reports) or as a particular column value changes. Using this feature, you can set the cell attributes to change when you reach a new group in a summary break report.

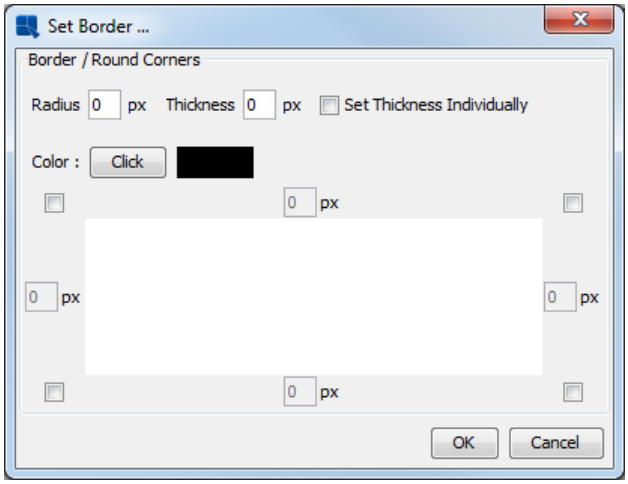
For alternate rows, you can set the background color, the font color, as well as the font style and size. Set the primary attributes as you would for any other report element.



Dual Colors Dialog

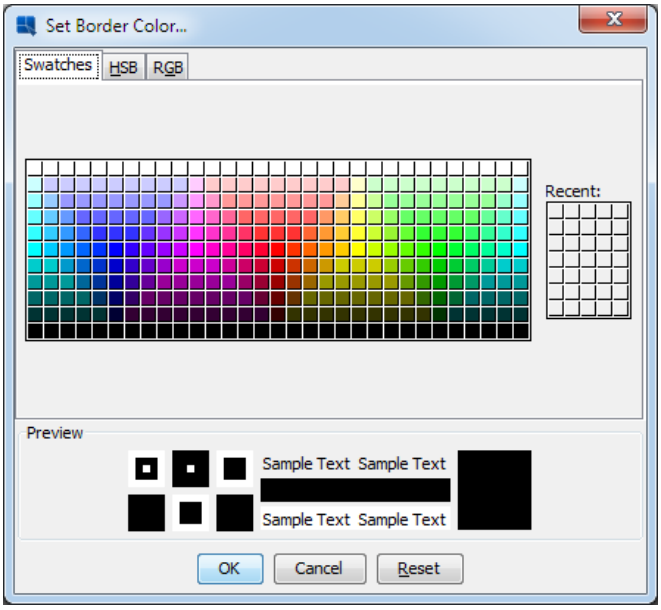
Border:

You can specify the thickness and color of the border to be drawn around any report element in one of three ways: selecting *Borders* from the *Format* menu, selecting the  *Border* button on the toolbar, or right clicking on the element and selecting *Edit Attributes* from the pop-up menu. Selecting this option will bring up a dialog box (or the multi-attribute editing dialog where options can be set in the *Border/Round Corners* tab) prompting you to enter the border thickness in pixels and to choose a color. Entering **`0`** as the border thickness value will remove the border.



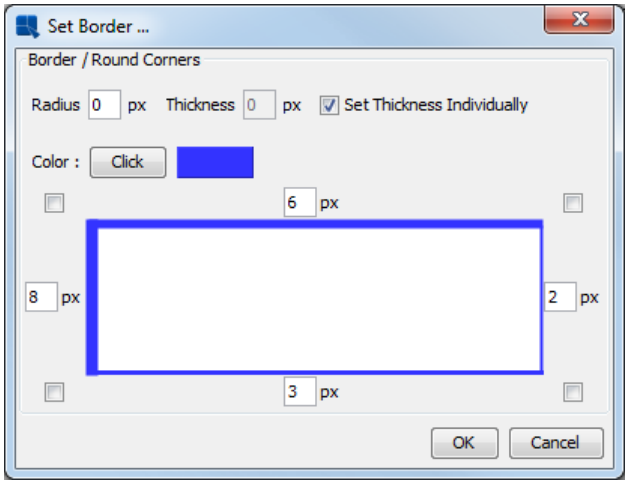
Border Dialog

To change the border color, click the *Click* button next to the *Color:* field. The *Set Border Color* dialog will appear allowing you to select border color from swatches. You can also enter HSB, HEX, or RGB values.



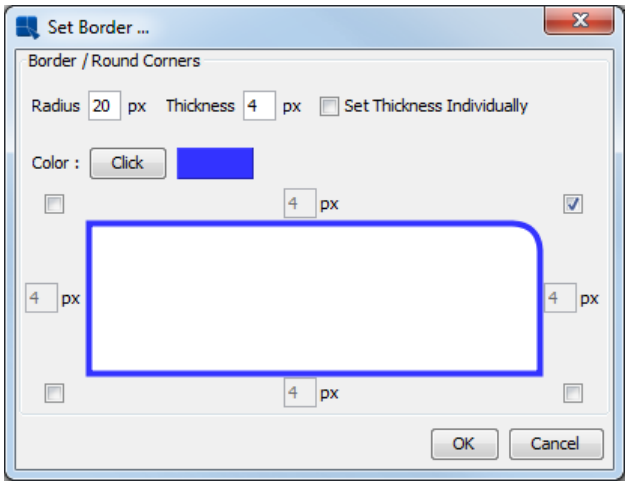
Border Color Dialog

To set the border thickness of the each side of the report element, check the *Set Thickness Individually* option. Then type the thickness (in pixels) in the text fields near the sides of the rectangle representing a selected element.



Border / Individual Thickness

To set round corners to an element, the elements need to have either non-transparent background color or a visible border, otherwise you will not see any changes. There are four checkboxes near the corners and each checkbox represents an element corner. To make a corner round, check its checkbox (you can also select several corners at once), then type a radius (in pixels) into the *Radius* text field.



Border / Round Corners Dialog



Note

If you plan on exporting reports with round corners, please note that round corners are available for DHTML and PDF export formats only.

5.1.3.7.6.1. Show Cell Outline

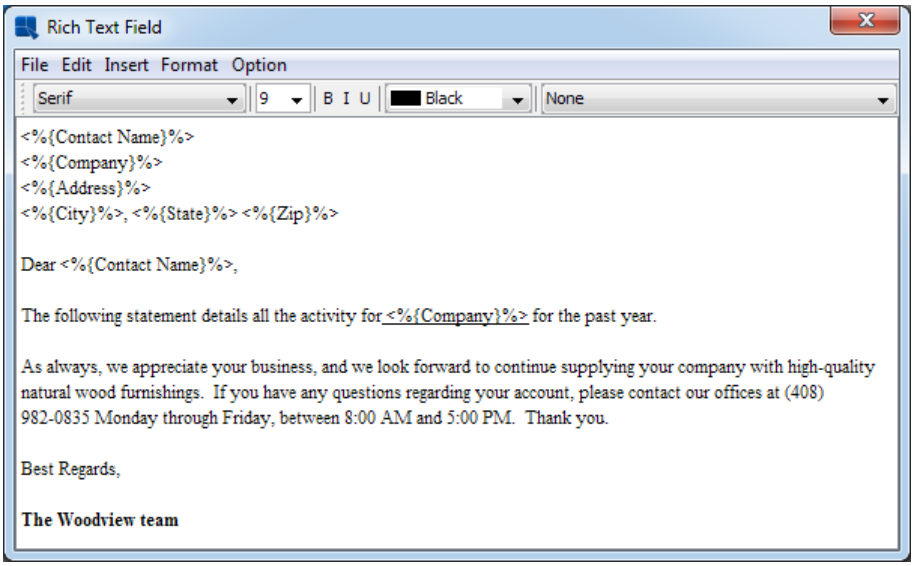
Sometimes you may want to show the boundaries or borders of report elements in the Design window, without actually turning on the borders. This is especially important if you have blank place-holder cells or columns where the first record is empty. To turn on the cell boundaries, select *Show Cell Outline* from the Option menu. This will draw a gray dotted border around each report element in the Design window.

5.1.3.7.7. Rich Text Fields

One of the common reporting needs is the ability to create form letters or other blocks of formatted text to a report and merge in data fields and functions. This functionality is available using rich text fields. Rich text fields, unlike

labels or string function fields, allow for complex paragraph, font, and color formatting within the cell. They also allow the ability to embed column fields and functions directly in the text flow.

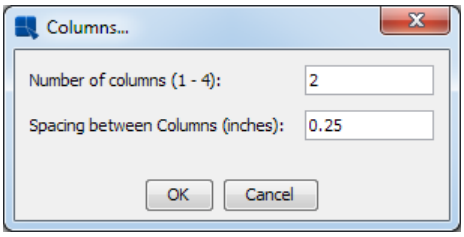
To add a rich text field to a report, select *Insert Rich Text Field* from the Insert menu, or click the  *Rich Text Field* button on the toolbar. After you select this option, a small box will follow your mouse pointer around the design window. Position the box where you want to insert the field and click. A new window will open, allowing you to enter the rich text field.



Rich Text Editor

The rich text editor works like a small word processor. The first drop-down box allows you to select the font and the font size. The three buttons allows you to specify bold, italic, or underlined text. The next drop-down-box allows you to specify font color. The last one allows you to select function or column field formats.

You can specify a columnar layout for the text by selecting *Columns* from the *Option* menu. This will bring up a dialog allowing you to specify the number of columns for the rich text field, as well as the spacing between the columns.



Rich Text Column Options Dialog

You can add in-line images to the text flow in rich text fields as well. To add an image, select *Insert Image* from the *File* menu. This will bring up a dialog prompting you to specify the image file you want to import. After you specify it, a new dialog will open, allowing you to specify the size for the in-line image in pixels. After you set the size, click the *Ok* button and the image will be added to the rich text field.

You can import any rich text file (.rtf format) into a rich text field. To do this, select *Import* from the *File* menu in the editor. This will bring up a dialog prompting you to specify the file you want to import.



Note

Certain paragraph settings will be lost when you import a rich text file.

5.1.3.7.7.1. Adding Formulas

You can add any formula directly into the text flow in a rich text field. It is also possible to add column fields and parameter values. To insert a formula, column field or a parameter value, click on the *Insert* menu, select a category (column, formula or parameter), and choose an item from the list of all available fields/parameters/formulas. The following syntax will be added into the text: *<% formula name/parameter name/column field name %>*.

The Insert → Formula list contains all formulas from the report. You can also create a custom formula directly in the rich text. To do so, simply type the *<%Formula Syntax%>* syntax (replace the *Formula Syntax* by the actual formula syntax). For more information about formula syntax, please see Section 5.1.6 - Using Formulas & the Formula Builder. This approach is recommended for simple formulas only. If you want to insert a more complex formula, you can create it in the report using the Formula Builder and then add it to the rich text field by selecting it from the Insert → Formula menu. See Section 5.1.6 - Using Formulas & the Formula Builder chapter to learn how to create formulas in the Formula Builder.

To format formulas and column fields, you must first pre-define a format. To do this, select *New* from the editor *Format* menu and select the data type for which you want to create a format. This will bring up the formatting dialog for that data type, allowing you to define the format. After you close the formatting dialog, you will be prompted to specify a name for the formatting. For example, you could define a format called *Currency* that uses locale-specific currency formatting for numbers.

To apply a pre-defined format of a formula, select the formula and then select the format that you want to apply from the drop-down box on the right side of the rich text editor.

5.1.3.7.8. Hyperlinks

You can apply a hyperlink to any report element except column fields. To apply a hyperlink to an element, select it and then select *HyperLink* from the *Format* menu or right click on the element and select *HyperLink* from the pop-up menu. A dialog box will then pop-up asking you to enter the link, hint, and target. You can enter any URL as the link or another `.pak` file. The hint will display text in a hint box on mouse over when the report is viewed.

If you are running ReportDesigner in stand-alone mode, then hyperlinks to URLs will not work from the Preview window. If you export to DHTML or PDF, links to `.pak` files will no longer work.



Hyperlink Dialog

You can also apply dynamic hyperlinks to report elements, including column fields, by using cell scripting. For more information about cell scripting, please see Section 5.1.7 - Scripting.

5.1.3.7.9. Column Wrapping

This feature is useful for reports that have few or narrow columns of data that are longer than one page. Rather than have the data take up only a portion of the page width and extend over multiple pages in length, you can 'wrap' the columns so that they will continue to the right of the original columns on the first page of the report.

For example, say you have the following report listing product names, and the number of units ordered:

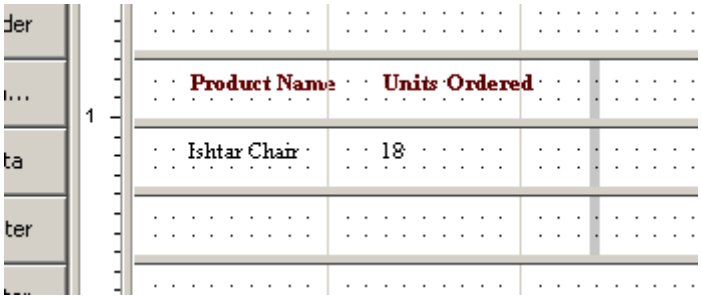
Product Name	Units Ordered
Adad Chair	4
Addad Dresser	5
Amon Table	2
An Chair	5

Product Name	Units Ordered
Anahita Dresser	3
Anubis Table	7
Apep Table	10
Apsu Dresser	3
Asherat Dresser	2
Atun Table	4
Bast Table	4
Bes Table	3
Cula Chair	10
Enki Chair	12
Enlil Chair	7
...	...
...	...

Assuming this report is long enough to encompass multiple pages, it is ideal to use column wrapping because it uses fewer pages and fills the blank space in the page width. The report will look like this with column wrapping:

Product Name	Units Ordered	Product Name	Units Ordered
Adad Chair	4	Bast Table	4
Addad Dresser	5	Bes Table	3
Amon Table	2	Cula Chair	10
An Chair	5	Enki Chair	12
Anahita Dresser	3	Enlil Chair	7
Anubis Table	7	...	...
Apep Table	10	...	...
Apsu Dresser	3		
Asherat Dresser	2		
Atun Table	4		

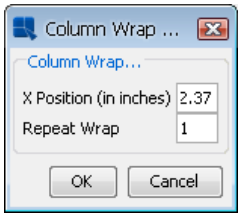
Column wrapping is available for simple columnar and summary break report types. To implement column wrapping, select *Column Wrap* from the *Format* menu. A gray vertical bar will be drawn across the middle of the Design window indicating the column wrapping placement.



Column Wrapping in Design Window

Column wrapping begins immediately to the right of the bar, so be sure to allow adequate space. Also, anything to the right of the column wrap bar will be truncated.

You can edit the column wrap properties by first selecting the column wrap bar and then selecting *Column Wrap* from the Format menu. You can also right click on the column wrap bar and then select *Column Wrap* from the pop-up menu. A dialog box will then appear allowing you to edit the properties of the column wrapping.



Column Wrapping Options

From this dialog, you can specify the X position of the column wrapping, as well as the number of times you want the columns to wrap in the page. The X position is the distance (in inches or centimeters) from the right side of the page where column wrapping is to begin. By default, column wrapping occurs once. Setting *Repeat Wrap* to **-1** will cause the wrapping to occur as many times as can fit within the page width.



Note

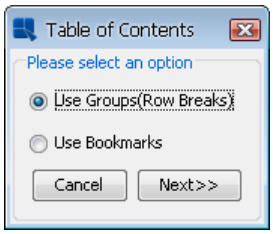
You can not specify column wrapping to occur more times than can fit within the page width. For example, if you specify three times, and the resultant report would be wider than the page, it will only repeat the wrapping twice.

5.1.3.7.10. Table of Contents

ERES has an ability to generate a table of contents for your report. The table of contents can either show the defined groups in the report, or a list of user-defined bookmarks. The TOC will display in Report Viewer, Page Viewer, DHTML, and PDF exports.

5.1.3.7.10.1. Adding a Table of Contents

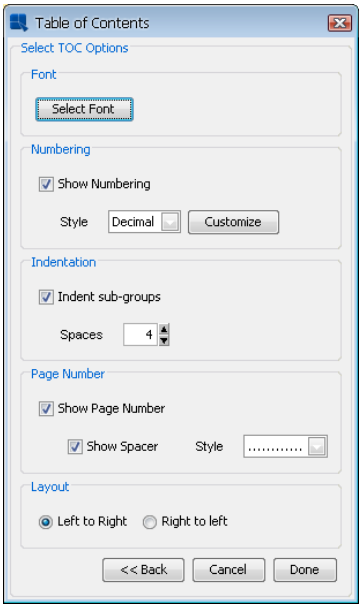
To add a table of contents to a report, select *Insert Table of Contents* from the *Insert* menu. Your cursor will turn into a cross. Click to place the table of contents in either the Report Header or Table Header sections. You can not add it to any other section. A dialog will open asking you if you would like to use groups or bookmarks for the table of contents.



TOC Type Dialog

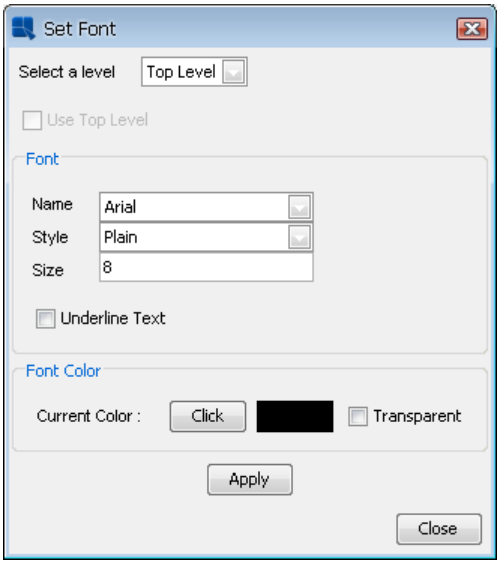
If you select to use groups, the table of contents will generate an entry for each group (and sub-group) in the report. In order to use this option, your report must contain grouped data (i.e. summary break, crosstab, master & details reports) with detail records. A crosstab report with one row break or a summary break report with column aggregation may be grouped, but because there are no detail records, you cannot create a table of contents for this presentation. If you select to use bookmarks, a one-level table of contents will be generated for each bookmark that is defined in the report. Bookmarks are defined using cell scripts. For more information about setting bookmarks, please see Section 5.1.7.2.1 - Formatting Actions.

The next dialog that appears allows you to specify options for the table of contents.



TOC Options Dialog

The first option in this dialog allows you to set the font for each level of the table of contents. When you click the *Select Font* button, a new dialog will open that allows you to specify font options.

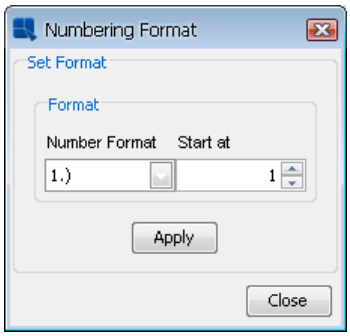


TOC Font Dialog

For each level you can set the font face, style, size, and color. Note that a level corresponds with a level of nested grouping in the report. If you select to create a report with bookmarks, only the top level will be available. Sub-levels can inherit the font from the parent level. After you make changes to the font for a particular level, click the *Apply* button to save the changes.

The next option allows you to select whether to show numbering for the table of contents and which style to use. You can customize each style by clicking the *Customize* button. The following styles are supported:

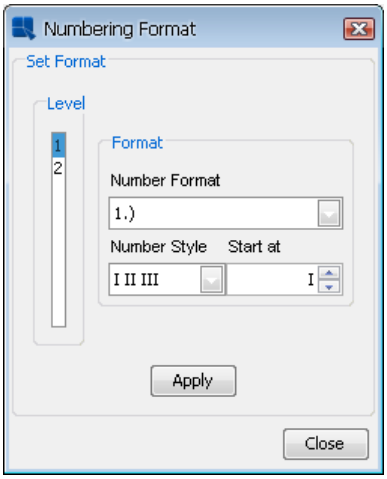
Decimal: This will use a numbered configuration for the table of contents entries with whole numbers for the outermost group and decimals for each sub-group. If you click the *Customize* button for this option, the following dialog will appear:



Customization Options for TOC Decimal Format

The first drop-down list allows you to select how the number should be formatted. The second option allows you to select the starting number. After you make your choices, click the *Apply* button to apply the changes.

Outline: This option will use an outline format allowing you to select numbers, letters, or Roman numerals for each level of the table of contents. If you click the *Customize* button for this option, the following dialog will appear:



Customization options for TOC Outline Format

The list on the left side allows you to select the level for which you would like to apply formatting. The first drop-down allows you to select the number format (this will apply whether you select numbers letters or Roman numerals for the level). The second drop-down allows you to pick the format for that level - either numbers, capital letters, lower-case letters, capital Roman numerals, or lower-case Roman numerals. The last option allows you to select the starting point for the selected style. After you make your choices, click the *Apply* button to apply the changes.

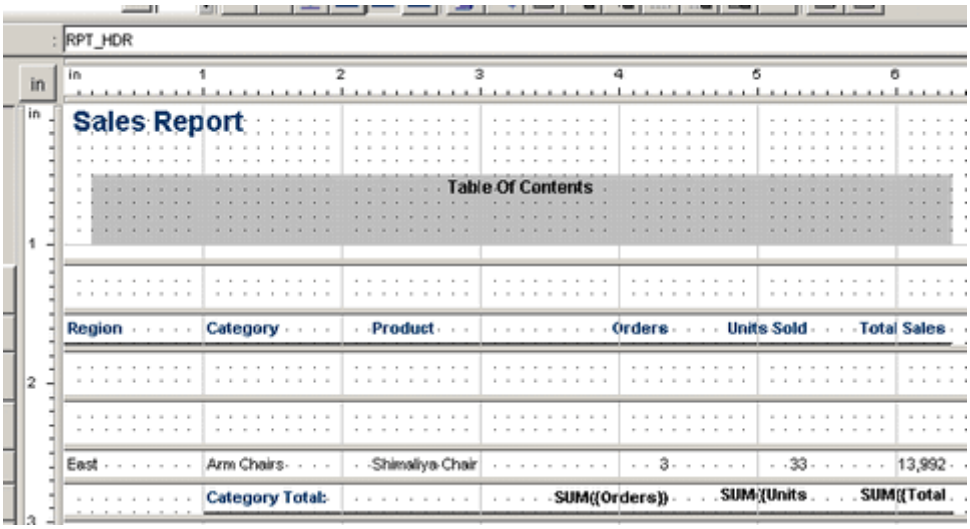
Bulleted: This option will use bullets for each entry. There are now additional customization options for this format.

The next option allows you to specify indentation for sub-groups. You can turn on/off the indentation and specify the number of spaces to use.

The next option allows you to specify whether to show the page number in the table of contents and whether to draw a spacer to the page number. You can draw a dotted line or a solid line for the spacer.

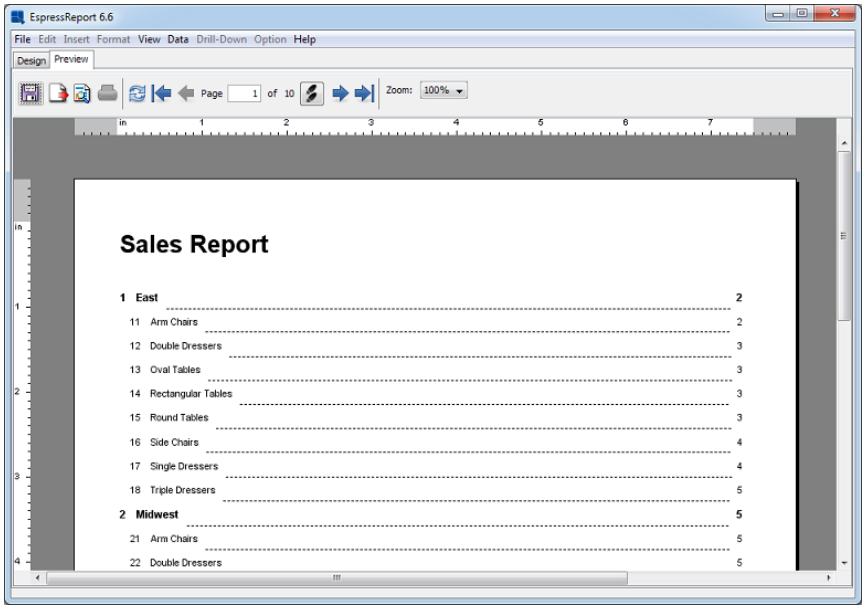
The last option in the dialog allows you to set the layout for the table of contents, either right to left, or left to right. After you finish setting options for the table of contents, click the *Done* button and it will be added to the report.

In the report, the table of contents will appear as a small grey rectangle. By default, the resize to fit content option will be turned on, but the width of the table of contents will be based on whatever size you allot for it in the section. Generally, you want to make this at least as wide as your report.



TOC in Design Window

You can see the complete table of contents when you preview the report. The table of contents is only supported for the DHTML and PDF exports. It will not appear if you export the report to other formats.



TOC in Preview Window

5.1.3.7.11. Moving and Resizing Report Elements

There are several ways to move and resize report elements. You can use rulers, mouse, or you can manually enter the element bounds.

Rulers:

You can resize and move elements using the rulers located in upper left corner of the Design window. When you select a report element by clicking on it with your mouse, a shaded area will appear on each ruler, marking the element's horizontal and vertical bounds.

On the ruler at the top of the Design window, clicking and dragging the left-hand marker will move the element horizontally. Clicking and dragging the right-hand marker will stretch the element horizontally.

On the ruler on the left side of the Design window, clicking and dragging the top marker will move the element vertically. Clicking and dragging the bottom marker will stretch the element vertically.

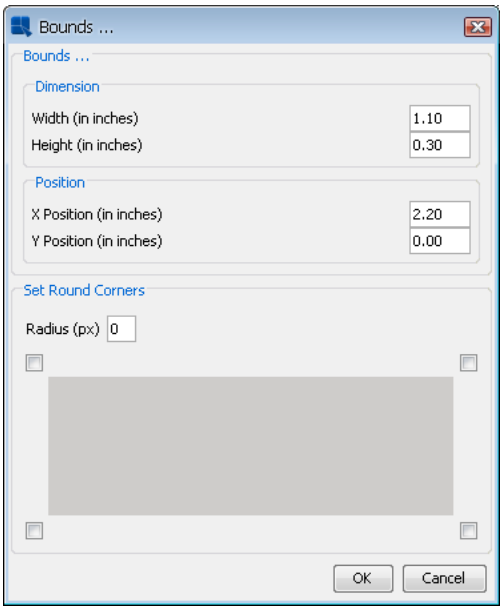
Mouse:

You can resize and move elements using your mouse. To move an element, simply click and drag it. To resize an element, click and drag on any of the sizing handles that appear around the edge of an object when it is selected. You can also right click and drag within the cell to resize it.

Manually Set Bounds:

You can manually set the bounds for any report element by clicking the  *Bounds* icon on the toolbar, or by selecting *Bounds* from the *Format* menu (You can also set bounds via the *Edit Attributes* dialog that you can open from the *Format* menu or by right-clicking an object and selecting it from the pop-up menu). A dialog box will appear prompting you to enter the new element bounds. The measurements will be in inches or centimeters, depending on which unit is selected with the toggle button in the upper left corner of the designer window where the rulers meet. The X and Y coordinates are for each specific report section. The point that marks an element's X and Y position is in the upper left corner of the element.

From the same dialog, you can also set round corners for the selected element. To learn more about round corners, please see Section 5.1.3.7.6 - Font, Color, and Border Options.

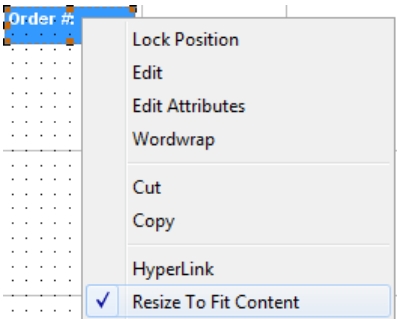


Set Bounds Dialog

Resizing To Fit Contents:

Often, you may have a column field that contains data of varying length. In this case the data may be truncated, or there may be a great deal of blank space within the cell when the field is short. To alleviate this problem, you can set the element to resize its height dynamically to encompass its contents.

To do this, first select the cell you want to resize and then select Format → Resize To Fit Content. The element will dynamically resize to fit its content. Note that only the height of the cell will resize. The width will remain the same.



Dynamic Resize Dialog

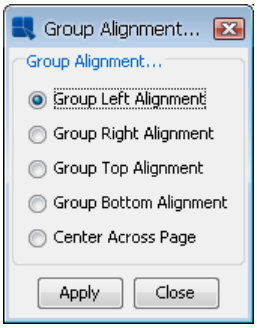
You can also invoke the section option called *Resize Cells Proportionally* to cause all of the other cells in the section to adjust their height with the resized cell. For more information about section options, please see Section 5.1.3.3 - Section Options.

Group Move/Resize:

You can move or resize a group of elements by selecting the elements first. Multiple elements in a report can be selected by drawing a selection box around them or by selecting them using **CTRL**+**Click**. You can draw a selection box by clicking and dragging on empty space in a report section. Once the elements are selected, click and drag on any element in the group to move that group. Click and drag on one of the re-sizing handles to adjust the size of elements in the group.

Group Alignment:

You can align a group of elements to the left side, right side, top, or bottom of a group. Select a group of elements and then right click on it. The *Group Edit* dialog box will appear giving you the option to align the group. Select *Group Left Alignment*, *Group Right Alignment*, *Group Top Alignment*, or *Group Bottom Alignment* and click the *Apply* button. All of the elements in the group will move to the left, top, or bottom edge of the group. You can also select the *Center Across Page* option. This will take the selected group of cells and center them in the page.



Group Options/Alignment Pop-up Dialog

Moving Elements Between Sections:

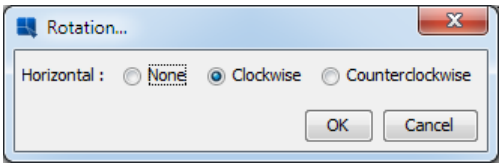
You can move a single element or group of elements between report sections by dragging them over the section boundaries. Note that elements cannot be dragged into the Table Data section.

Column Swap:

You can swap the position of two elements within a report section. This is primarily used to rearrange the position of column fields in the report. To swap two elements, select the elements using **CTRL**+**Click** or the selection box and then select *Swap Columns* from the Format menu. The positions of the two elements will be switched and the surrounding elements will be moved to accommodate the swap.

Rotate:

Text elements like labels, formulas, column headers, and column fields can be rotated 90 degrees clockwise or counter-clockwise. There are two ways to rotate elements. The first way to rotate an element is to right click on it and then move the mouse pointer over the *Rotate* option. A sub-menu will pop-up containing three options - *None*, *Clockwise* and *CounterClockwise*. The second way to rotate an element is to select it and then choose *Rotation* option from *Format* menu.



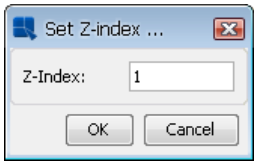
Rotation dialog

Note

Text rotation is available for DHTML, PDF, XSLX, and XLS export formats only.

5.1.3.7.11.1. Controlling Element Overlapping

In ReportDesigner, each element in a report has an associated Z index. This Z index number determines how elements behave when placed in the same space (on top of each other). When elements are placed in the same space, the element with the highest Z index will appear on top and all subsequent elements with lower Z indexes will be drawn successively below each other. To set the Z index for a report element, right click on it and select *Set Z-Index* from the pop-up menu. This will bring up a dialog allowing you to specify a number for the index of that element.



Z Index Dialog

Note

Certain export formats that require a tabular layout, like HTML and Excel, may not appear correctly if the report has elements that overlap each other.

5.1.3.7.11.2. Locking Element Position

In Report Designer, you can lock a cell position by right clicking on the element and selecting *Lock Position* from the pop-up menu. When this feature is enabled, the report element cannot be moved with mouse, rulers, or any of the group formatting or alignment features. It can only be moved when the lock position option is turned off.

5.1.3.7.12. Snap to Grid

By default, Report Designer operates in *Snap to Grid* mode. This forces report elements to move only in set increments. The grid layout is represented by small dots drawn in the Design window. By default, the grid step size is 0.1 in or 0.25 cm (depending on which measurement is selected).

To modify snap to grid options, select *Snap To Grid* from the Option menu. A dialog will open allowing you to set grid properties.



Snap To Grid Options

In this dialog, you can change the grid step size or disable the snap to grid feature. If you disable it, the report elements can be moved in free-form around the Design window.

5.1.3.7.13. Shift Mode

Shift mode is toggled on and off by selecting *Shift On/Off* from the *Option* menu. When shift mode is enabled, report elements will move to accommodate changes in other report elements. Hence, if you resize the width of one report element from 1" to 2", the elements next to the current element will shift to the right. If shift mode is not selected, other report elements will not move, even if the resized element ends up overlapping another.

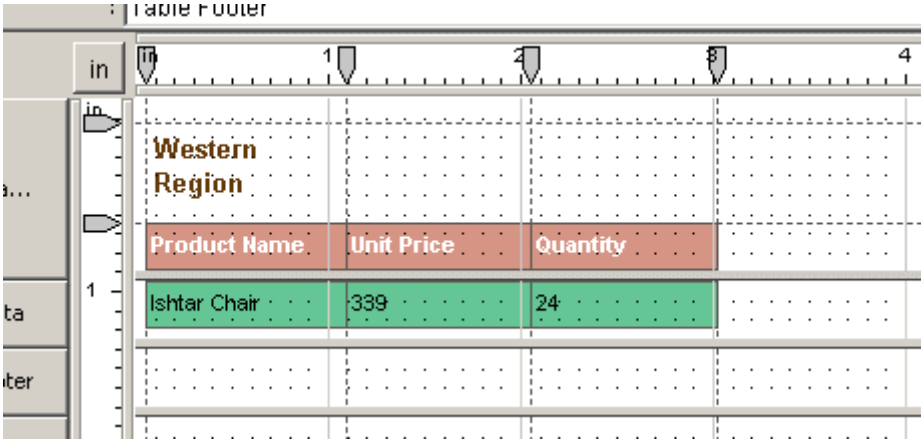


Note

Shift mode works best if you only resize elements or move them by a small amount. If you move elements a pronounced distance with shift mode enabled, it will displace other report elements by a pronounced distance as well, giving the report an odd appearance.

5.1.3.7.14. Guidelines

Another way to position and move groups of elements is to use guidelines. Guidelines allow you to place arbitrary position lines within the Design window and snap report elements to those lines. Guidelines allow you to precisely line up and position report elements without worrying about performing precise movements.



Elements Positioned with Guidelines

5.1.3.7.14.1. Inserting Guidelines

You can insert both horizontal and vertical guidelines by selecting *Insert Guideline* from the *Insert* menu. This option will open a sub-menu, prompting you to select whether you want a vertical or horizontal guideline. After you select your desired option, click in the design panel where you would like the guideline to be drawn. The guideline will appear as a dotted line with a marker in the upper or left ruler. You can move the guideline by clicking and dragging the marker in the ruler.

You can remove any guideline from a report by right clicking on the guideline (or the guideline marker in the ruler) and selecting *Delete* from the pop-up menu.



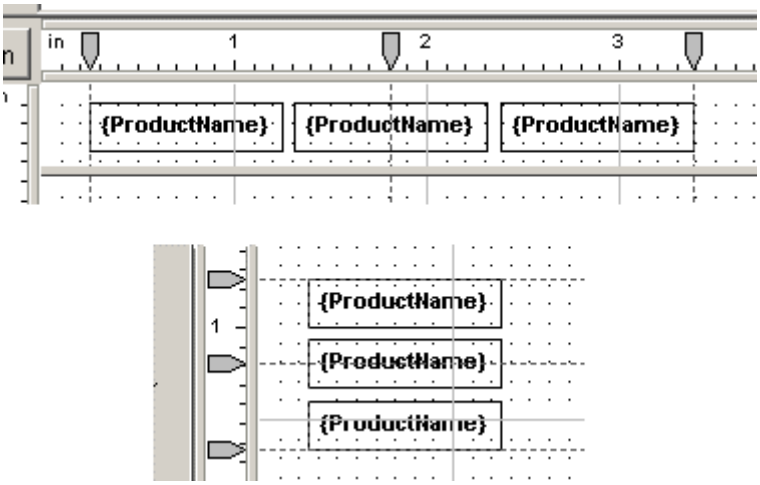
Note

Horizontal guidelines cannot be moved between report sections. If you wish to move your guideline to a new report section, you will need to insert a new guideline.

5.1.3.7.14.2. Positioning Elements with Guidelines

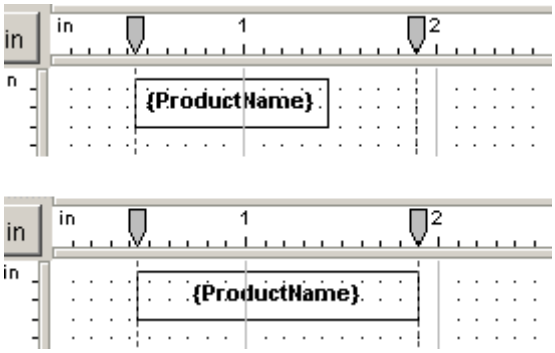
There are two ways in which objects can be snapped to guidelines. The first option is to snap all the elements that are near the guideline. To do this, right click on the guideline (or the guideline marker in the ruler) and select *Snap Elements in Range*. You can then further select to snap the cells left edge, right edge, or center. This will snap all the cells that are up to 0.1 inches away from the guideline.

The other option is to selectively snap elements to guidelines. To do this, first select the elements you want to snap to the guideline using **CTRL**+**Click** (you can select only one element, but you must use **CTRL**+**Click** to select it), then right click on the guideline (or the guideline marker in the ruler) to which you want to snap it. You can specify to snap the cell's left edge, right edge, or center to the guideline for both horizontal and vertical guidelines. Select the option you want from the pop-up menu and the cell(s) will snap to the guideline.



Guideline Positioning Options

Once an element is attached to a guideline, moving the guideline will cause the element to move as well. You can also snap an element to multiple guidelines. Snapping an element to two guidelines will cause it to stretch to fit.



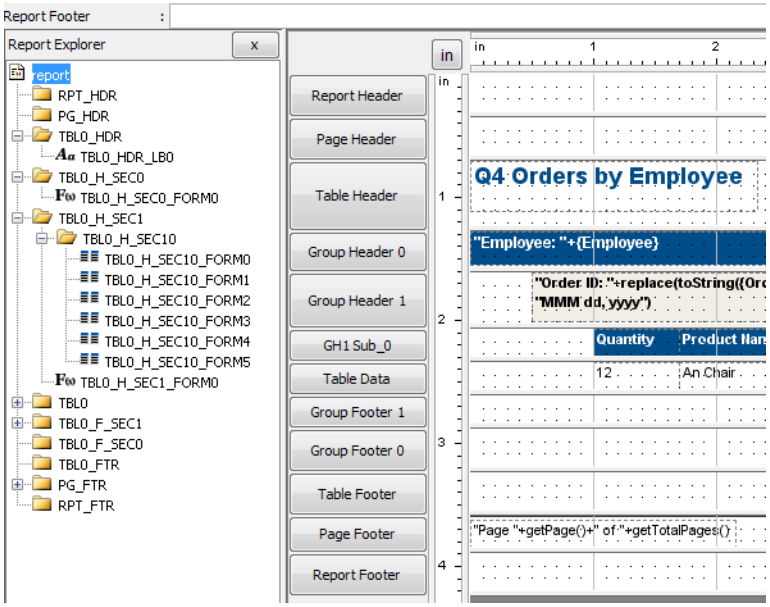
Element Snapped to Two Guidelines

When you have an element attached to two guidelines, moving either guideline will cause the element to resize (either stretch or shrink).

To release an element from a guideline, right click on it and select *Separate From Guidelines* from the pop-up menu. This will release the element from any guidelines to which it is attached. You can also release all of the elements attached to a guideline by right clicking on the guideline (or the guideline marker in the ruler) and selecting *Separate All* from the pop-up menu.

5.1.3.8. The Report Explorer

In addition to the main Design window, ReportDesigner provides another interface to select and edit the elements in a report - the Report Explorer. The Report Explorer displays all elements in the report as a tree structure on the left side of the Designer. It can be turned on and off by selecting *Report Explorer* from the *Option* menu.



Report Explorer Window

Each parent node in the tree represents a section in the report. Inside each section, an icon indicates element type and element ID. If you specified a custom ID for the element, it will display instead. For more information about custom IDs, see Section 5.1.6.2.1 - Using Column Field Data.

You can also edit report elements from within the Explorer. When you select an element in the tree, the element will also be selected in the Design window and the Design window will scroll in order to display that element on your screen. You can then use toolbars, menus, or pop-up menu to set any of the element properties. The pop-up menu can also be activated by right clicking on any element in the tree.

5.1.3.9. Global Formatting, Group Formatting, and Templates

To this point, you have only seen how to edit and format the properties of individual report elements. However, sometimes you may want to format multiple elements at once. To do this, you can use global formatting, group formatting, and templates.

5.1.3.9.1. Global Formatting

The global formatting feature allows you to give all the elements in your report a consistent look and feel. To set the global formats, select *Global Format* from *Option* menu. This will bring up a second menu with each of the different report element types listed: chart, column, formula, image, label, line, rectangle, column header, and title. Selecting one of these will bring up a tabbed dialog box prompting you to set the formats for those elements. Each tab contains formatting options for an element attribute. The formatting options are as follows:

Chart: border color, border thickness, set bounds.

Column: data format, font style and size, font color, background color, border color, border thickness, dual colors, set bounds, alignment, rotation, script.

Formula: data format, font color, font style and size, background color, border color, border thickness, set bounds, alignment, rotation, script.

Image: border thickness, border color, set bounds, script.

Label: font color, font style and size, background color, border color, border thickness, set bounds, alignment, rotation, script.

Line: line color, set bounds, script.

Rectangle: background color, border color, border thickness, set bounds, script.

Column Header: font color, font style and size, background color, border color, border thickness, set bounds, alignment, rotation, script.

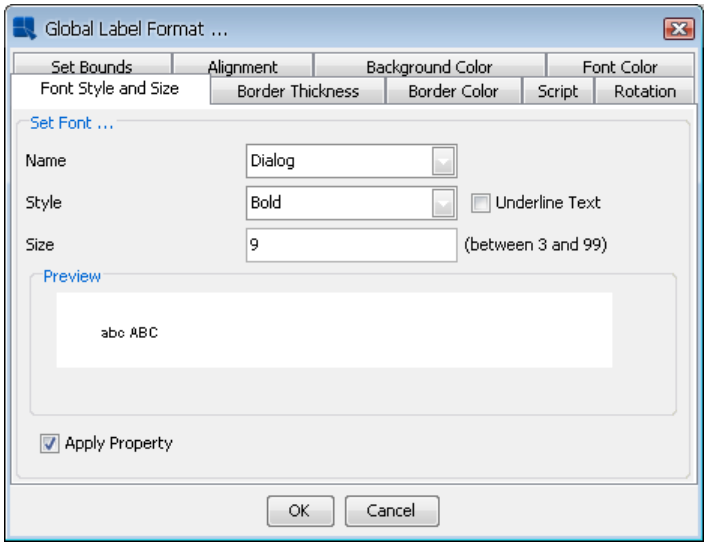
Title: font color, font style and size, background color, border color, border thickness, set bounds, alignment.

At the bottom of each tab is a checkbox marked *Apply Property*. Checking this box will apply the property to the global formats when you click *OK*. The box automatically becomes checked when you change a property. After you make all the property changes, clicking the *OK* button will change the properties of all elements of that type in the entire report. It will also become the default attribute for any additional elements placed in the report.



Note

You have an option to insert images in the original size (default attribute). This can be done using the checkbox labeled *Default Size* in the *Set Bounds* tab of the image global format dialog.



Global Format Dialog

5.1.3.9.1.1. Global Format Import/Export

You can pass global formats from one report to another using the Import/Export feature. You can export global formats by selecting *Export* from the *Global Format* sub-menu. This will bring up a dialog box prompting you to specify a filename. The global formats will then be saved as an XML file. You can load a global format XML file by selecting the *Import* option from the *Global Format* sub-menu. This will bring up a dialog box prompting you to specify the XML file you want to import. Click the *OK* button and the formats stored in the XML file will be applied to the current report.

5.1.3.9.1.2. Global Formatting & Crosstab Reports

Global formatting is very important when dealing with crosstab reports. Due to the nature of the report, the number of columns can increase or decrease depending on changes in the data, data source, or filtering criteria. When new columns are generated in a crosstab report, they will inherit the default formatting. In order to control the appearance of new columns in a crosstab report, you will need to set the default formatting using global formats.

Global formatting is also important because it controls the intervals in which new columns appear in the report. New columns will appear in intervals defined by the column width in the global formats. Therefore, if you set the default width to be inch, new columns in the crosstab report will appear at one inch intervals. Because of this, it is recommended that you maintain consistent widths for the data columns in a crosstab report. Setting different widths can result in unexpected behavior when new columns are added to the report.

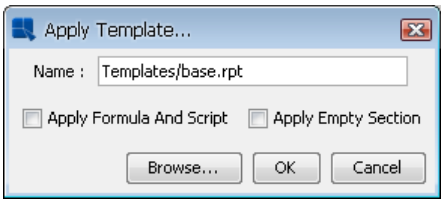
5.1.3.9.2. Group Formatting

To apply formatting to a group of elements, first select a group of elements either by drawing a selection box around them or using **CTRL**+**Click**. You can also select a row or column of cells by right clicking on a cell and selecting the *Select Column* or *Select Row* option from the pop-up menu.

Once you select the group, format the properties as you would for a single element. Formatting will only take effect on elements where it is applicable. For example, if you select four elements, two labels, a formula, and a chart, and then change the font size, it will have no effect on the chart.

5.1.3.9.3. Applying a Template

The features and elements of any report (`.pak, .qrp, .rpt` or `.xml`) file can be applied to another report. This is accomplished using the apply template feature. To apply a template to the current report, select *Apply Template* from the *File* menu or click on the *Apply Template Button* on the toolbar. A dialog box will appear, prompting you to select the file you want to apply. When you apply a template, all of the elements and their respective formatting will be applied to the current report. Only the data source information is not carried over (this includes functions and scripts in the table data section of the report).



Apply Template Dialog

There are two additional options available in the dialog.

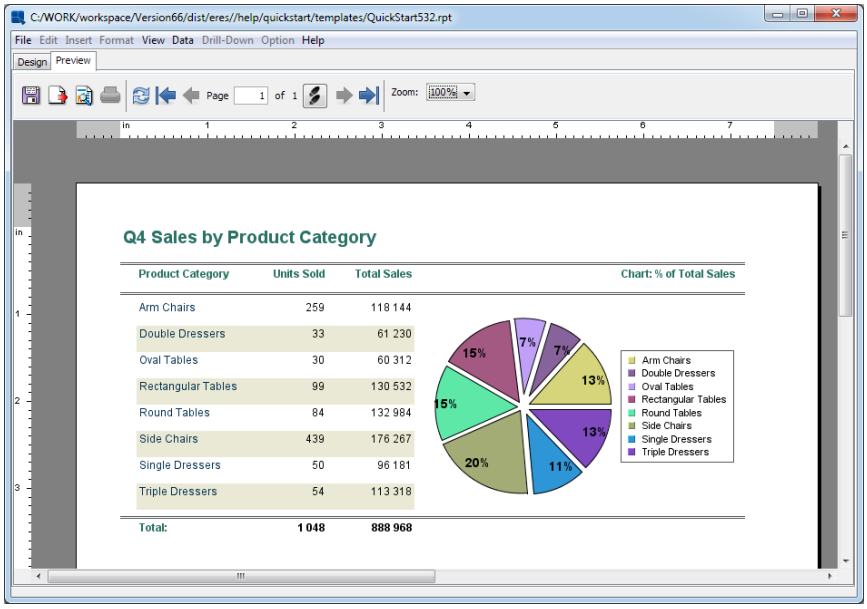
Apply Formula and Scripts: This option will apply the formulaic columns and cell scripts from the template onto the report. By default, they do not apply. Note that formulas and scripts from the template may not work correctly if the data types of the columns in the new report are different from those in the template.

Apply Empty Section: This option indicates whether or not to apply the blank sections from the template to the report. If you apply blank sections, they will overwrite sections in your report, essentially removing any elements you may have in the section. This option is generally used if you want to completely replicate the template that you are applying with the new report. If you do not apply blank sections, those sections in your report will not be overwritten. Only sections in the template with elements will be applied to your new report. This feature can be used if you have a certain default layout (headers/footers) that you want to pass among a group of reports. You can define a template where only the page headers/footers have defined elements and apply those headers/footers to other reports by selecting not to apply empty sections.

If you would like to have a report that has the same look and feel as another one without replacing the labels, formulas, etc, you can accomplish this by using the global format import/export feature rather than applying an entire template.

5.1.4. The Preview Window

After inserting or manipulating report objects, you can view the results in the Preview window. You can switch back and forth between the Design and Preview windows by using the tabs in the upper left corner of the Report Designer. The Preview window gives you an accurate picture of what the report will look like if it is printed or exported.



Preview Window

The Preview shows the page dimensions in inches or centimeters (depending on which unit is selected with the toggle button at the upper left corner of the designer window where the rulers meet). The drop-down menu on the right side of the toolbar allows you to set zooming for the preview.

5.1.4.1. Navigating the Report

You can navigate around the report using the toolbar or the *View* menu. The toolbar buttons perform the following functions:

 Go to the first page of the report

 Go to the previous page of the report

 Go to a specific page (that you enter in the *Page* text box)

 Go to the next page of the report

 Go to the last page of the report

These navigation functions can also be performed by selecting *First Page*, *Previous Page*, *Next Page*, *Last Page*, or *Go To Page...* from the *View* menu.

5.1.4.2. Other Preview Window Options

Other options in the Preview window toolbar perform the following functions:

 Save the file

 Export the file

 Load the entire report in the page viewer window. For more information about this, please see Section 5.1.4.2.4 - Using Page Viewer.

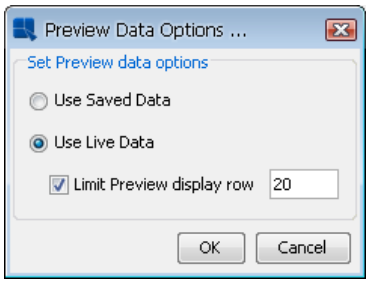
 Print the report (this option is not available if the Report Designer is run as applet)

Refresh Data

In addition, the *File*, and *Data* menus remain active, allowing you to manipulate the file and data.

5.1.4.2.1. Preview Data Options

You will be presented with the data options dialog the first time you preview a report. This dialog allows you to set several options related to how the Preview window displays report data.



Preview Data Options Dialog

The first option allows you to select whether or not to use live data when previewing the report. If you select live data, the report will connect to the data source and retrieve data every time the report is previewed. Note that if the data cannot be obtained (for example, if the database connection is down), the preview will show an error. You will then need to reload the report (once the connection has been reestablished) or go back to the *Preview Data Options* and select *Use Saved Data*. If you select to use saved data, the report will show whatever data it has when you preview the report. This could be either two records of back-up data that are stored in the template, the first twenty records that are retrieved when a report is first created, or the full dataset from the previous preview (if the option had been set to use live data).

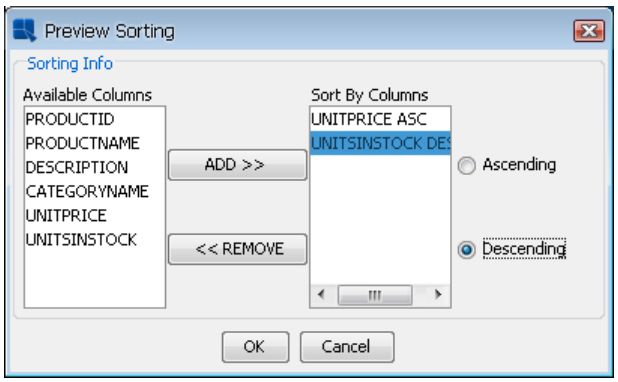
If you select to use live data to preview the report, a second option can be set that allows you to limit the number of records that should be retrieved from the data source. If your report uses a large dataset, this option allows you to see the report with data while limiting the processing time and preserving client memory.

Once you set this option the first time a report is previewed, you will not be prompted again when you preview the report. If you want to change these settings, select *Set Preview Data Options* from *Data* menu (this option is active only in *Design* tab). This will bring up the dialog again, allowing you to change the preview settings.

5.1.4.2.2. Sorting Data

You can sort the data within the Preview window based on any column in the report in ascending or descending order. To do this, select *'Sort by (ascend)'* or *'Sort by (descend)'* from the *View* menu. Each option will bring up a secondary menu containing all of the report column fields. Selecting a column field will cause the report to sort by that column in either ascending or descending order. This feature is also available in the Report Viewer applet.

To sort by multiple columns, select *Sort By...* from the *View* menu. This will bring up a dialog allowing you to select which columns to sort by and the direction.



Preview Sorting Dialog

To sort by a column, select it in the left side dialog and click the *Add* button. You can set the sorting direction by selecting the column on the right side and clicking either the *Ascending* or *Descending* button.



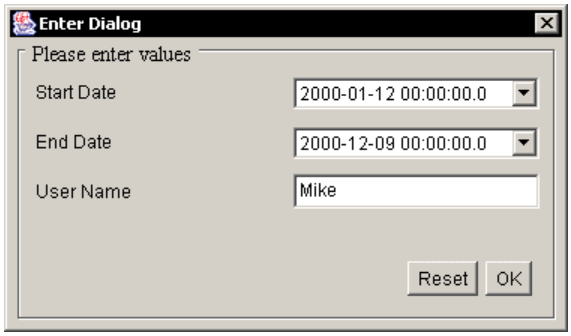
Note

Any sorting performed will not be saved with the template. When the template is re-opened (run), the data will use the original order. However, you can export the report after sorting and the exported file will reflect the re-ordered data.

It is not recommended that you use this feature with reports containing a large amount of data, as this will use a large amount of memory and can compromise performance.

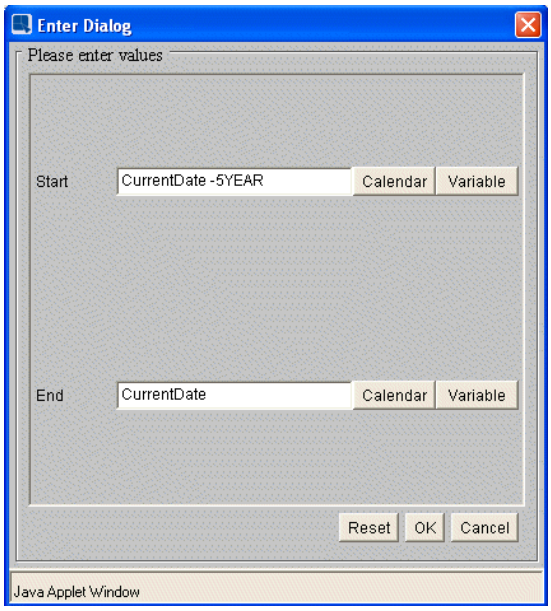
5.1.4.2.3. Parameter Prompts

If your report uses a parameterized query or if you have defined any formula parameters in the report, you will be prompted to select parameter values when you preview the report.



Parameter Prompt Dialog

You are prompted to type in or select parameter values depending on how the parameters are mapped. Clicking the *Ok* button will generate the report using the specified values. You can disable parameter prompting by toggling the *Preview Parameter Prompt* option in the Data menu. If date variables are not mapped to a database column, the parameter prompt will appear as below, with the options of entering a date from the calendar or a date variable.



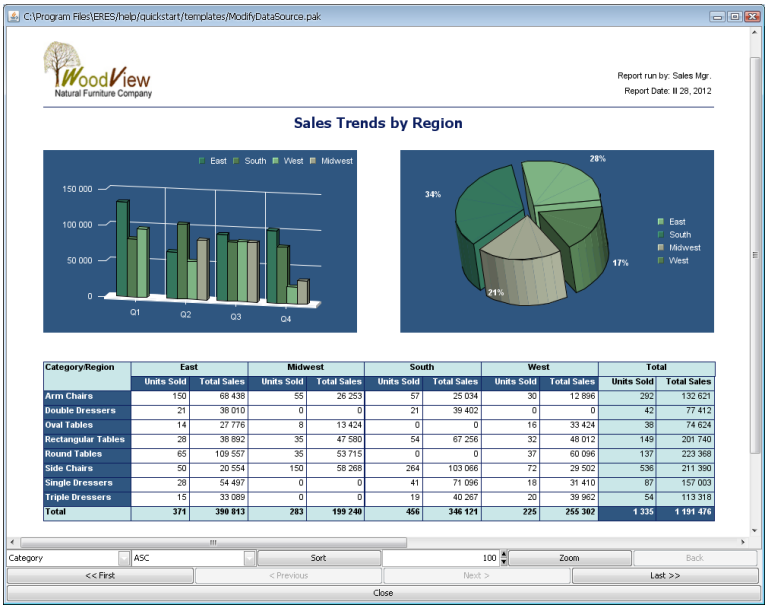
Date Parameter Prompt Dialog

For more information about query parameters, see Section 4.1.2.2.2 - Parameterized Queries. For more information about formula parameters, see Section 5.1.6.2.6 - Formula Parameters.

5.1.4.2.4. Using Page Viewer

Generally when you preview the report, the entire report with its data is kept in memory. This allows you to quickly preview and navigate through the generated report. However, for large reports, previewing the entire report could significantly affect system performance as the entire report is loaded into memory. To prevent this problem, you can limit the number of records that are fetched during preview as detailed in Section 5.1.4.2.1 - Preview Data Options.

On the other hand, you may want to preview the entire report without loading it into memory when creating an ad-hoc report. To do this, ReportDesigner allows you to load the report in the Page Viewer. To do this, select *Launch Page Viewer* from the *View* menu or click the  *Launch Page Viewer* icon on the toolbar. This will open a Page Viewer window containing the report.

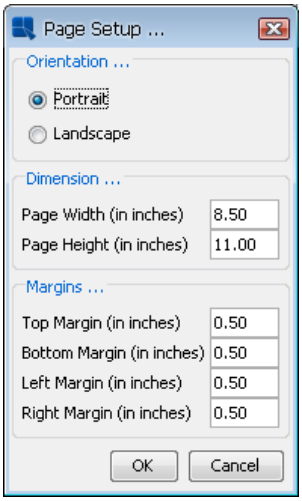


Report Shown in Page Viewer Window

The Page Viewer window will show you the entire report. You can navigate around the different pages by clicking the buttons on the toolbar at the bottom of the window or by right clicking and selecting to change pages from the pop-up menu. Because the Page Viewer is a static format, the window will not automatically reflect any subsequent changes made to the report. You will need to close the window and re-launch the Page Viewer. For more information about the Page Viewer, please see .

5.1.4.3. Setting Page Properties

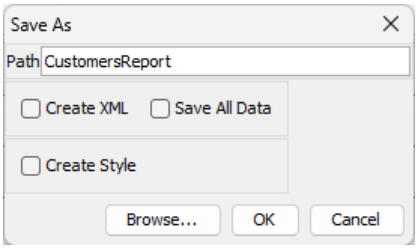
The output you see in the Preview window is based on the page dimensions and margins that are set in the Design window. The page dimensions are marked by the rulers, with the shaded areas indicating the page margins. To adjust these properties, select *Page Setup* from the *Option* menu. This will bring up a dialog box prompting you to set the page orientation (landscape or portrait), height, width, and margins. Measurements are in inches or centimeters, depending on which metric system is currently selected. Changing the orientation will also change the default print properties.



Page Properties Dialog

5.1.5. Saving Reports

You can save the current report by selecting *Save* or *Save As* from the *File* menu, or by clicking the  *Save* button on the toolbar. All reports are saved as templates either in `.pak` (binary) or XML format. Selecting the save option will simply overwrite the existing file, unless you are working on a new report that has yet to be saved. Selecting *Save As* option will bring up a dialog box prompting you to create a file name. The *Save As* dialog also has several checkboxes.



Save As Dialog

Create XML: This checkbox specifies whether or not to save the report in XML format. By default, reports are saved in binary (`.pak`) format. When this checkbox is checked, ReportDesigner will generate the template in XML format. The XML file can be saved and re-opened by the Designer. It can also be applied as a template to other reports.

Create Style: This checkbox specifies whether or not to save the report as a custom style. Reports saved in this format can be applied as a default look and feel to new reports. However, they cannot be used as a report any more. For more information about this featue, see Section 5.1.2.6.1 - Custom Styles.

Save All Data: This will save all of the report data in the template. By default, only two records are saved. This feature is useful if you would like to preserve the structure of a crosstab when changing data sources or editing the report, without the datasource. It also allows users to share a complete report when the data source is not available. Reports are always opened with back-up data in the Report Designer. Users can get the latest data by selecting the "Live Data" option when previewing the report.

5.1.6. Using Formulas & the Formula Builder

ReportDesigner supports a wide variety of formulas, giving you an important mechanism for manipulating and displaying report data. Using formulas, users can add summaries and aggregations, perform basic calculations, and build complex expressions using one of over 70 built-in functions.

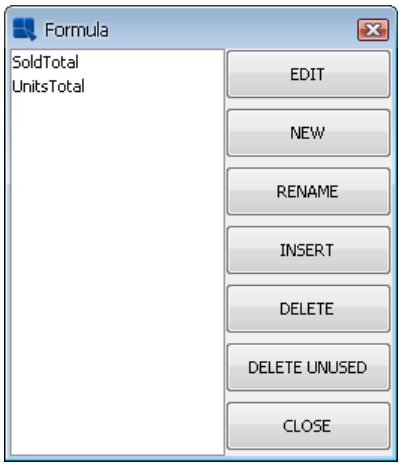


Note

Some of the syntax for ReportDesigner formulas has changed greatly between v2.51 and v3.0. In most cases, the deprecated syntax will continue to work correctly. However, it is recommended that you check older templates to make sure the functions still return correctly if you have upgraded to v3.0 or higher from v2.51 or lower.

5.1.6.1. Creating a Formula

When you select to insert a new formula into the report (either by selecting *Insert Formula* from the *Insert* menu or clicking the  *Insert Formula* button on the toolbar), a list of all the defined formulas in the report will appear.



Formula List Dialog

To insert a formula into the report, select a formula from the list and click the *Insert* button. The formula list dialog will then close and you will be allowed to place the formula.

The report section in which you place the formula is extremely important, since the formula will reset each time that section repeats in the report. For example, if you place a formula in the Group Footer of a summary break report, the formula will recalculate for each group in the report. If the formula references a column value, it will only use data within each break group. If you place the same formula in the Report Footer section, it would only calculate once and it would use all of the data in the column field, regardless of breaks and grouping.

Formulas and labels that are placed in the Table Data section of the report become computed columns and can be treated as column fields as well as formulas.

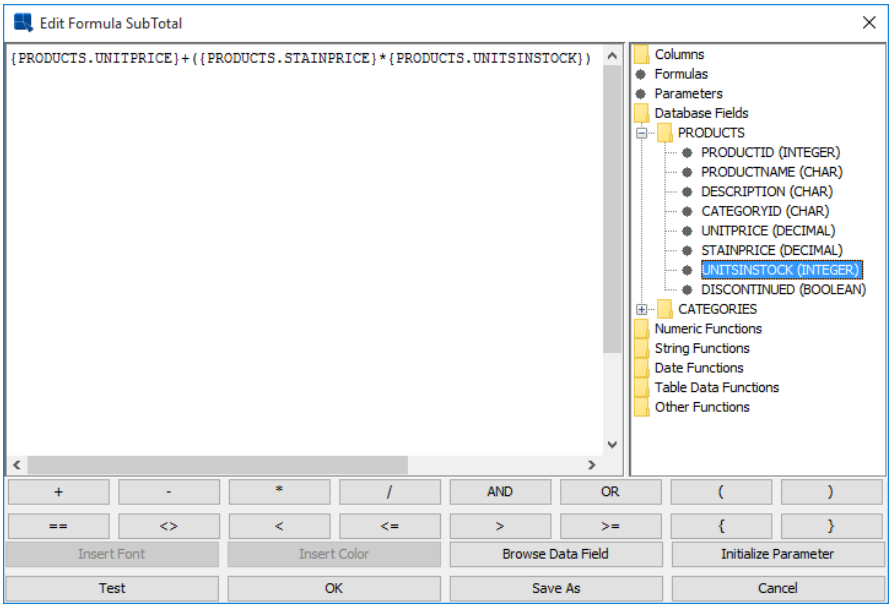


Note

Labels are automatically treated as Strings and can be edited by double clicking on them (i.e. to use a number as integer in a formula, open the formula builder and remove double quotes from the number).

5.1.6.1.1. The Formula Builder

Formulas are constructed using the Formula Builder interface. To launch the Formula Builder, you can select to create a new formula by selecting *New* from the formula list dialog. This will prompt you to specify a name for the new formula. Once you specify a name, the Formula Builder will open, allowing you to construct a formula. You can also use the Formula Builder to edit an existing formula. To edit a formula, select it from the list and click the *Edit* button or double-click on a formula cell in the Design window.



Formula Builder Window

The main window of the Formula Builder contains the text of the formula. The folders on the right side contain various elements that can be added to the formula, including column fields, other formulas, parameter values, database fields, and built-in functions. The first two rows of buttons contain the most common arithmetic and Boolean operators (operators and functions are discussed in the next section). The last two rows of buttons serve as command buttons and perform the following functions:

Insert Font: This option is not available for formulas, it is only available for scripting.

Insert Color: This option is not available for formulas, it is only available for scripting.

Browse Data Field: This option is for database fields. It will query the database and return the first 20 values for the column.

Initialize Parameter: This option will bring up the parameter initialization dialog for any parameters defined in the formula. Parameters are discussed in Section 5.1.6.2.6 - Formula Parameters.

Test: This option will check the current formula and see if it is correct.

Ok: This opton will close the formula builder and return to the formula list dialog, allowing you to insert the formula into the report.

Save As: This option allows you to save the current formula under a different name. Because you can re-use a formula in different places in the report, using the save as option allows you to modify a formula in one place, without changing it for every place where the formula is referenced.

Cancel: This option will close the Formula Builder without saving the formula/changes.

5.1.6.2. Formula Syntax

The following section details various operators and functions available in ReportDesigner and how to use them.

5.1.6.2.1. Using Column Field Data

You will often want to use data from column fields for your formulas. This can be easily accomplished in formulas using the following syntax: `{<Field Name>}`. The braces delimit the field and the field name is the name of the column field you're referencing. So in a basic example: `{UnitPrice} * {Quantity}` would multiply two column fields together.

In most cases, the column name is the name specified in the data source (i.e. the database column name or alias). If you are uncertain about the correct name, you can simply select the column field you want to use from the *Columns* node in the right panel in the formula builder.

There are two other functions you can use to retrieve column data:

id() This option will allow you to retrieve the value of any cell in the report, including column fields. The syntax is `id( <Object ID>)`. You can retrieve an object's ID by right clicking on the cell and selecting *Properties* from the pop-up menu. You can also retrieve the ID for column fields by selecting *View Column Mapping* from the Data menu.

You can also assign a custom ID to elements in the report, instead of using the default ID. To assign a custom ID to an element, right click on it in the Report Designer and select Custom ID from the pop-up menu. This will bring up a dialog allowing you to enter an ID for that element. The argument for the `id()` function can either be the original ID for the element or the defined custom ID.

Element IDs can also be used to get a handle to an element in the Report API. You can pass in either the original ID or the custom ID into the `getData()` method to get a handle to an element.

The `id()` function also serves another important purpose. It can be used to delay the calculation of computed columns. Normally, formulas in the Table Data section of the report will compute prior to any aggregations in the group or table footers. However, sometimes you may want to calculate the aggregations prior to the column fields. For example, you want to create a column that calculates each row's percentage of a total. In order to get the correct result, the total has to be calculated before the column. To do this, simply refer to the aggregation in the Table/Group Footer using the `id()` function. The finished function in the Table Data section would look like this: `{Quantity} / id(TBL0_FTR_FORM0)`.

col() This option is a deprecated function that is used in earlier version of ReportDesigner to retrieve data from column fields. Although this function is still valid, it is recommended that you use the new field notation. The syntax for this function is `col(Column Index)`. To view the column index, select *View Column Mapping* from the Data menu. This will display the assigned index number for each column field. Index values are assigned based on the order in which the columns are selected for the report. The first column has an index of 0.

5.1.6.2.1.1. Column Aggregation

You can aggregate any column field (including computed columns) using one of the aggregation functions provided with ReportDesigner. The following aggregation options are available:

average() This option returns the mean value of the specified column. The syntax is `average(<Column Field>)`. To return the average value of a column field named `UnitPrice`, you can use the following formula: `average({UnitPrice})`.

median() This option returns the median value for the specified column. The syntax is `median(<Column Field>)`. If the median of the column is two values (for an even number of rows), the aggregation will return the average of the two values.

count() This option returns a count of the values in the specified column. The syntax is `count(<Column Field>)`.

countdistinct() This option returns a count of all the distinct values in the specified column. The syntax is `countdistinct(<Column Field>)`.

max() This option returns the maximum value for the specified column. The syntax is `max(<Column Field>)`.

min() This option returns the minimum value for the specified column. The syntax is `min(<ColumnField>)`.

stddev() This option returns the standard deviation for the specified column. The syntax is `stddev(<Column Field>)`.

sum() This option returns the sum of the specified column. The syntax is `sum(<Column Field>)`.

sumsquare() This option returns the sum of squares for the specified column. The syntax is `sumsquare(<Column Field>)`.

variance() This option returns the variance for the specified column. The syntax is `variance(<Column Field>)`.

It is important to note that aggregation functions can only take column fields as an argument. You cannot use an expression or another function as an argument. Therefore, `sum({UnitPrice}*{Quantity})` is not valid. To make this calculation, you would first need to add `{UnitPrice}*{Quantity}` as a formula in the Table Data section of the report. Then you can sum the computed column created by the first formula using `sum(@formulaname)`.

5.1.6.2.2. Using Formulas

ReportDesigner allows you to easily re-use report formulas by plugging them directly into new formulas. You can refer to a formula using the following syntax: `@<FormulaName>`. The formula name is the name that was assigned to the formula the first time it was created. So for example if you created a formula `{UnitPrice}*{Quantity}` called `Total`, you could retrieve the result of this formula by typing `@Total`.

If you are uncertain about the correct name, you can simply select the formula you want to use from the *Formulas* node in the right panel of the Formula Builder.

You can also use this notation to aggregate computed columns if the formula has been placed in the Data Table section of the report.

5.1.6.2.3. Using Parameter Values

In addition to column field and formula values, you can also access parameter values as part of a formula. Parameter values are the user supplied values to either query or formula parameters. For more information about query parameters, see Section 4.1.2.2.2 - Parameterized Queries. For more information about formula parameters, see Section 5.1.6.2.6 - Formula Parameters.

To refer a parameter value, use the following syntax `:<ParameterName>` for query parameters and `?<ParameterName>` for formula parameters. For example, if the report has a query parameter named `StartDate` that prompts the user to supply a date, `:StartDate` will return the date that the user has supplied when the report is run.

If you are uncertain about the correct name, you can simply select the parameter you want to use from the *Parameters* node in the right panel of the Formula Builder. Also, if you type a formula parameter name wrong, the Formula Builder will assume that you are trying to define a new parameter.

5.1.6.2.4. Using Database Fields

If your report uses a database as the data source, you can also use database fields that were not selected for the report in a formula. Database fields are referenced in a similar manner to column fields and use the syntax `{<Table Name>.<Field Name>}`. If your database requires three part names, the fields should be referenced accordingly.

You can select database fields to add from the *Database Fields* node in the right panel of the Formula Builder. In addition, the Formula Builder allows you to preview database fields. To do this, first select a field in the right panel and then click the *Browse Data Field* button. This will bring up a new dialog showing the first few rows from the selected column.

5.1.6.2.5. Constants

When using constants in formulas, there are certain formats that are required for each data type.

Numeric Data: To represent a numeric constant, you simply have to type the number. Both integers and decimals can be accepted; however, you cannot type in thousands separators. So to specify the number 12.28 as a constant in a formula, type `12.28`.

String Data: To represent a string constant, type the string delimited with double quotes. For example, to specify a string Hello, you would type `"Hello"`.

Boolean Data: To represent a Boolean constant, type either true or false without any delimiters. To specify a false Boolean value, you would type `false`. Note that these are not case sensitive.

Date Data: To represent a date constant, type the date in the following format without delimiters: MM/dd/yyyy. To specify the date August 5th, 2002, you would type `08/05/2002`.

Time Data: To represent a time constant, type the time in the following format without delimiters: hh:mm:ss. To specify the time 2:30 PM, you would type `14:30:00`. Note that time must be in 24-hour format.

Timestamp Data: To represent a timestamp constant, type the date and time in the following format without delimiters: MM/dd/yyyy hh:mm:ss. To specify the time stamp August 5th, 2002 at 2:30 PM, you would type `08/05/2002 14:30:00`.

5.1.6.2.6. Formula Parameters

Rather than adding constants into a formula, you can specify a parameter. Using a parameter, you can collect constant values from the user (or elsewhere) at run-time and dynamically compute the formula based on those values. To specify a parameter in the formula, use a question mark "?" followed by the parameter name.

For example, the following formula `{Quantity}+?Value` would add a user supplied value to the `Quantity` field in a report.

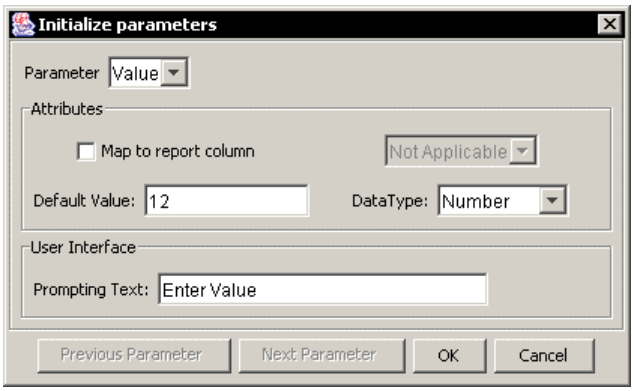


Note

If you want to define a new parameter within a formula, you have to use a unique name. If you specify a parameter name of a parameter that is already defined in another formula within the report, the new formula will return the value from the previous parameter.

5.1.6.2.6.1. Initializing Formula Parameters

Like query parameters (detailed in Section 4.1.2.2.2.2 - Initializing Query Parameters), you must initialize a formula parameter before a formula can be used in a report. To initialize any parameters defined in the current formula, click the *Initialize Parameter* button in the Formula Builder. This will bring up a dialog prompting you to specify options for the parameter.



Initialize Formula Parameter Dialog

From this dialog, you can specify a default value for the parameter, map the parameter to a report column (this will give the user a drop-down list of distinct values, rather than having them type in the parameter), specify the data type for the supplied parameter, as well as specify the prompt for the end user. Clicking the *Previous Parameter* and *Next Parameter* buttons allows you to initialize each of the parameters that have been defined in the formula. Click *OK* when you finish setting up the options for all parameters.



Note

For formula parameters, you must specify a default value for each parameter. Also, if the parameter is mapped to a report column, the first time the report is previewed, the dialog will only contain the first twenty values of that column. This is because the full dataset is not retrieved when the report is first created. At run-time, the dialog will contain all the available values in the report column.

5.1.6.2.7. Operators

ReportDesigner provides several arithmetic and Boolean operators that allow you to create expressions in formulas. Operators use in-fix notation that places the operator between arguments in the expression, i.e. `<argument1> <operator> <argument2>`.

You can automatically insert operators into a formula by clicking one of the operator buttons in the formula builder.

5.1.6.2.7.1. Arithmetic Operators

Four basic arithmetic operators are supported: `+`, `-`, `*`, and, `/`. These will add, subtract, multiply, and divide two objects respectively. For example `1 + 2` will return `3`.

Generally, arithmetic operators will take numbers as arguments and will return numbers with two exceptions. `+` can be used to concatenate strings, for example `Hello + "World"` would return `Hello World`. Also, `-` can be used to subtract two date objects and return the difference in days, for example `8/2/2002 - 4/7/2002` would return `117`.

5.1.6.2.7.2. Boolean Operators

Eight Boolean operators are supported: `AND`, `OR`, `==`, `<`, `<=`, `>`, `>=`, and `<>`. `AND` and `OR` take Boolean arguments and return Boolean values. For example `true AND true` would return `true` or `1 < 2 OR 4 < 3` would return `true`.

`==`, `<`, `<=`, `>`, `>=`, and `<>` are comparison operators that signify equal to, less than, less than equal to, greater than, greater than equal to, and not equal respectively. Comparison operators compare two objects of the same data type and will return a Boolean value, for example `3 < 4` would return `true` and `Yes == "No"` would return `false`.

Boolean operators are more commonly used in cell scripting rather than formulas. Cell scripting is covered in Section 5.1.7 - Scripting.

5.1.6.2.8. Functions

In addition to the basic operators, ReportDesigner provides a number of built-in functions. Function syntax generally takes the form of a function name outside of a set of parenthesis enclosing the comma-separated arguments `function(<argument>, <argument>)`.

You can insert functions into the formula by selecting the function you want under the *Numeric Functions*, *String Functions*, or *Date Functions* nodes in the Formula Builder. Functions are inserted with hints indicating the type and number of arguments it should take. In order for the function to return properly, these hints have to be removed and replaced with valid arguments.

5.1.6.2.8.1. Numeric Functions

The following functions take numeric arguments.

abs() This option will return the absolute value of a number. The syntax is `abs(<Number>)`. For example `abs(-12)` would return `12`.

acos() This option will return the arc cosine of an angle, in the range of 0 through pi. Syntax for this is `acos(<radians>)`. For example, `acos(-sqrt(2)/2)` would return `3*pi/4`.

asin() This option will return the arc sine of an angle in the range of -pi/2 through pi/2. The syntax is `asin(<radians>)`. For example, `asin(sqrt(2)/2)` would return `pi/4`.

atan() This option will return the arc tangent of an angle in the range of -pi/2 through pi/2. The syntax is `atan(<radians>)`. For example, `atan(1)` would return `pi/4`.

atan2() This option will convert rectangular coordinates (b, a) into polar coordinates (r, theta). This function returns theta by taking the arc tangent of b/a in the range of -pi to pi. The syntax is `atan2(<y_coordinate>, <x_coordinate>)`. For example `atan2(1,0)` would return `pi/2`.

ceil() This option will return the lowest integer value that is greater than the specified number. The syntax is `ceil(<Number>)`. For example, `ceil(15.3)` would return `16`.

cos()	This option will return the cosine of an angle in radians. The syntax is `cos(<radians>)`. For example `cos(3*pi()/4)` would return `-sqrt(2)/2`.
e()	This option will return the value e, the natural logarithm base. The syntax is `e()` without any arguments.
exp()	This option will return e raised to the nth power. The syntax is `exp(<Number>)`. For example `exp(3)` would return `e^3` or `20.086`.
factorial()	This option will return the factorial of the specified object. The syntax is `factorial(<Number>)`. For example `factorial(3)` would return `6`.
floor()	This option will return the highest integer value that is less than the specified argument. The syntax is `floor(<Number>)`. For example `floor(5.52)` would return `5`.
getColumnCount()	This option will return an integer indicating the total number of columns in the report. The syntax is `getColumnCount()` without any arguments. This function returns the total number of columns whether they are visible or not.
getRowIndex()	This option will return an integer indicating the index value of the current row (of Table Data) of the report. The syntax is `getRowIndex()` without any arguments.
getTotalRowIndex()	This option will return a count of all the rows (in Table Data) in the report. The syntax is `getTotalRowIndex()` without any arguments.
getRowIndexOfCurrentTable()	This option will return the current row in the current group (table) in a report. This function is generally only relevant in reports that have grouped data like the summary break and master & details layout. For example, if the current row is the third row in a group of data, `getRowIndexOfCurrentTable()` will always return 2 regardless of what the row index is for the total report (which you could retrieve using the `getRowIndex()` function). The syntax is `getRowIndexOfCurrentTable()` without any arguments.
getTotalRowIndexOfCurrentTable()	This option will return a count of all the rows in the current group (table) in a report. This function is generally only relevant in reports that have grouped data like the summary break and master & details layout. The syntax is `getTotalRowIndexOfCurrentTable()` without any arguments.
getSiblingCount()	This option will return a count of the number of sibling groups within a nesting level. This function is generally only relevant in reports that have grouped data like the summary break and master & details layout. For example, you have a summary break report with one level of grouping. There are three distinct values in your row break column creating three groups at the level. Within the group (i.e. Group Header, Table Data, or Group Footer sections) `getSiblingCount()` would return 3. The syntax is `getSiblingCount()` without any arguments.
getChildCount()	This option will return a count of the number of child groups inside a nesting level. This function is generally only relevant in reports that have grouped data like the summary break and master & details layout. The function will return the number of child group in outer most group when running in a table header or a table footer; It will return the number of child in the related group when running in a group footer/header. For example, you have the same example as described for the previous function. Within the group, the `getChildCount()` function would return 0 as there is only one level of grouping. However, outside the group (i.e. Table Header or Table Footer

sections) `getChildCount()` would return 3. The syntax is `getChildCount()` without any arguments.

getGroupIndex() This option will return an integer indication the index value of the current group. This function is only relevant in reports that have grouped data like the summary break and master & details layout. This function will only return the index of the inner most group for reports with nested groups.

log() This option will return the natural logarithm of the specified number. The syntax is `log(<Number>)`. For example `log(10)` would return `2.303`.

mod() This option will return the remainder after dividing two numbers. The first argument is the numerator and the second is the denominator. The syntax is `mod(<Number>, <Number>)`. For example `mod(12, 7)` would return `5`.

pi() This option will return the value pi. The syntax is `pi()` without any arguments.

pow() This option will return the value of the first argument raised to the power of a second argument. The syntax is `pow(<Number>, <Number>)`. For example `pow(2, 3)` would return `8`.

random() This option will return a random value greater than or equal to 0 and less than 1. The syntax is `random()` without any arguments.

rint() This option will round the specified argument to the nearest integer. The syntax is `rint(<Number>)`. For example `rint(5.6)` would return `6`.

sin() This option will return the sine of an angle in radians. The syntax is `sin(<radians>)`. For example `sin(pi()/2)` would return `1`.

sqrt() This option will return the square root of the specified number. The syntax is `sqrt(<Number>)`. For example `sqrt(64)` would return `8`.

tan() This option will return the tangent of an angle in radians. The syntax is `tan(<radians>)`. For example, `tan(pi()/4)` would return `1`.

toDegrees() This option will convert an angle measured in radians to degrees. The syntax is `toDegrees(<Number>)`. For example, `toDegrees(pi())` would return `180`.

toRadians() This option will convert an angle measured in degrees to radians. The syntax is `toRadians(<Number>)`. For example, `toRadians(180)` would return `3.142`.

toString() This option will convert a number into a string. The syntax is `toString(<Number>, <Number of Decimals>, <Round Up(True/False)>)`. For example `toString(12.216, 2, True)` would return "`12.22`" as a string.

5.1.6.2.8.2. String Functions

The following functions take string arguments:

getAllRowData() This option returns whole row of data in a format like `(COLUMN1_NAME) VALUE1 (COLUMN2_NAME) VALUE2 ...` Does not require any arguments.

getHeader() This option will return the column name for a column field. The syntax is `getHeader(<Column Field>)`. For example, if you have a column named `UnitPrice`, `getHeader({UnitPrice})` would return `UnitPrice` as a string.

getPage() This option will return the current page number. The syntax is `getPage()` without any arguments.

getTotalPages() This option will return the total number of pages for the current report. The syntax is `getTotalPages()` without any arguments.

getTotalSections() This option will return the total number of sections (i.e. the number of pages needed horizontally) for the current report. The syntax is `getTotalSections()` without any objects.

indexOf() This option will return the first index value where a specified pattern within a string occurs. There are two syntaxes for this function.

1. `indexOf(<String>, <Pattern>)`. For example `indexOf("Banana", "an")` would return `1`.
2. `indexOf(<String>, <Pattern>, <Starting Index>)`. For example, `indexOf("Banana", "an", 2)` would return `3`.

insert() This option allows you to insert new characters into a string. The syntax is `insert(<String>, < Character Number>, <New Characters>)`. For example `insert("Wood Natural Furniture", 4, "View")` would return `WoodView Natural Furniture`.

lastIndexOf() This option will return the last index value where a specified pattern within a string occurs. There are two syntaxes for this function:

1. `lastIndexOf(<String>, <Pattern>)`. For example `lastIndexOf("Banana", "an")` would return `3`.
2. `lastIndexOf(<String>, <Pattern>, <Starting Index>)`. For example, `lastIndexOf("Banana", "an", 2)` would return `1`.

replace() This option allows you to replace one set of characters in a string with another. There are two syntaxes for this function.

1. `replace(String, Character Start Number, Character End Number, New Characters)`. For example, `replace("Today is a rainy day", 11, 16, "sunny")` would return `Today is a sunny day`.
2. `replace(<String>, <Old Characters>, <New Characters>)`. For example, `replace("Today is a rainy day", "rainy", "sunny")` would return `Today is a sunny day`.

setMaxLength() This option allows you to set the maximum length of a string. The syntax is `setMaxLength(<String>, <Maximum Number of Characters>)`. For example, `setMaxLength("You're a firefighter",13)` would return `You're a fire`.

strcmp() This option compares two strings and returns the lexicographical difference between them. If the first argument is larger than the second one, the result is positive. If the first argument is smaller than the second one, the result is negative. If the strings are same, the result is zero. The syntax is `strcmp(<String>, <String>)`. For example, `strcmp("a", "c")` would return `-2`.

strcmpIgnoreCase() This option compares two strings and returns the lexicographical difference between them while ignoring case. The syntax is `strcmpIgnoreCase(<String>, <String>)`. For example, `strcmpIgnoreCase("a", "A")` would return `0`.

strcat() This option concatenates multiple strings together. The syntax is `strcat(<String>, <String>, <String>...)`. For example `strcat("Wood", "View")` would return `WoodView`. You can also use the `"+"` operator to concatenate strings.

strlen() This option will return an integer indicating the number of characters in a specified string. The syntax is `strlen(<String>)`. For example `strlen("ReportDesigner")` would return `14`.

substring() This option will return a portion of a string as specified by an argument. There are two syntaxes for this function:

1. `substring(<String>, <Character Start Number>)`. For example, `substring("unhappy", 2)` would return `happy`.
2. `substring(<String>, <Character Start Number>, <Character End Number>)`. For example `substring("smiles", 1, 5)` would return `mile`.

trim() This option will remove any leading or trailing spaces from the specified string. The syntax is `trim(<String>)`. For example, `trim(" Hello ")` would return `Hello`.

toLowerCase() This option will render any uppercase letters within the specified string to lowercase. The syntax is `toLowerCase(<String>)`. For example `toLowerCase("ABCdef")` would return `abcdef`.

toNumeric() This option will turn a string into a double. The syntax is `toNumeric(<String>)`. For example `toNumeric("425.52")` would return `425.52`. For this formula to work correctly, the string argument must contain numeric characters.

toUpperCase() This option will render any lowercase letters within the specified string to uppercase. The syntax is `toUpperCase(<String>)`. For example `toUpperCase("ABCdef")` would return `ABCDEF`.

5.1.6.2.8.3. Date/Time Functions

The following functions use date/time arguments. Some of the date/time functions use a special argument that indicates a calendar field. To specify a calendar field argument, type one of the following without any delimiters.

ERA	DAY_OF_MONTH	HOUR
YEAR	DAY_OF_YEAR	HOUR_OF_DAY
MONTH	DAY_OF_WEEK	MINUTE
WEEK_OF_YEAR	DAY_OF_WEEK_IN_MONTH	SECOND
WEEK_OF_MONTH	AM_PM	MILLISECOND

addTime() This option adds a specified amount of time to a given date/time. The syntax is `addTime(<Date/Time>, <Calendar Field>, <Number>)`. The number specifies the amount that the calendar field should be added to the date/time object. For example, `addTime(12/5/1998, MONTH, 5)` would return `May 5, 1999`, and `addTime(12/5/1998, DAY_OF_MONTH, -25)` would return `Nov, 10 1998`.

getAmPm() This option will return AM or PM as a string for a given time. The syntax is `getAmPm(<Time>)`. For example, `getAmPm(13:24:00)` would return `PM`.

getCurrentDate() This option will return the current date from the system. The syntax is `getCurrentDate()` without any arguments.

getCurrentDateTime() This option will return the current date and time from the system. The syntax is `getCurrentDateTime()` without any arguments.

getCurrentTime()

This option will return the current time from the system. The syntax is `getCurrentTime()` without any arguments.

getDateTime()

This option will return the value of a specified calendar field for a given date/time. The syntax is `getDateTime(<Date/Time>, <Calendar Field>)`. For example, `getDateTime(12:24:00, MINUTE)` would return `24`, and `getDateTime(10/10/2001, DAY_OF_WEEK)` would return `4` (meaning Wednesday).

getDayDifference()

This option will return the difference in days as an integer between two dates. The syntax is `getDayDifference(<Date>, <Date>)`. For example, `getDayDifference(5/1/2001, 3/1/2001)` would return `61`. You can also use the "-" operator to return the difference between dates.

getDayOfWeek()

This option will return the day of the week for the specified date. The syntax is `getDayOfWeek(<Date>)`. For example, `getDayOfWeek(10/10/2001)` will return `Wednesday` as a string.

getWeekOfYear()

This option will return the week of the year for the specified date. The syntax is `getWeekOfYear(<Date>)`. For example, `getWeekOfYear(10/10/2001)` will return `41`.

getEra()

This option will return the era for the specified date. The syntax is `getEra(<Date>)`. For example, `getEra(10/10/2000)` would return `AD` as a string.

getMonth()

This option will return the month for the specified date. The syntax is `getMonth(<Date>)`. For example, `getMonth(10/10/2000)` would return `October` as a string.

rollTime()

This option is a time rolling function. The syntax is `rollTime(<Date/Time>, <Calendar Field>, <Number>)`. Unlike the `addTime()` function, this function will only adjust the specified calendar field by the specified amount and will have no effect on the other fields. For example `rollTime(12/5/1998, MONTH, 5)` would return `May 5, 1998`, and `rollTime(12/5/1998, DAY_OF_MONTH, -25)` would return `Dec 11, 1998`. For the latter example, when the count reaches the beginning of the month, it starts over at the end without changing the month field.

printDate(), printDateTime(), printTime()

These three functions allow you to set the way in which date/time information is displayed beyond what is capable using the data formatting options. These functions will return date and/or time information as a string. Their syntax is as follows:

```
printDate(Date, Date Format)
printDateTime(Date, Date & Time Format)
printTime(Time, Time Format)
```

The date and/or time format is entered as a series of characters and delimiters. Letters are used to represent different elements of date/time data. The characters and what each of them represent are listed below:

Character	Represents	Output (text/number)	Example
G	era	text	AD
y	year	number	1996, 96
M	month in year	text or number (dependson length)	July, Jul, 07
d	day in month	number	10

Character	Represents	Output (text/number)	Example
h	hour am/ pm (1-12)	number	1
H	hour 24 hr. (0-23)	number	18
m	minute in hour	number	30
s	second in minute	number	55
S	millisecond	number	978
E	day in week	text	Tuesday, Tue
D	day in year	number	189
F	day of week in month	number	2 (as in 2 nd Wed. in July)
w	week in year	number	27
W	week in month	number	2
a	am/pm marker	text	AM, PM
k	hour 24 hr (1-24)	number	24
K	hour am/ pm (0-11)	number	0
z	time zone	text	Pacific Standard Time, PST

You can piece together almost any combination of these characters to produce a date expression in any format you want. The count of groups of characters determines the form that the element will take. For text elements with four or more characters in a group, the full form of the element will be used. If less than four characters are used, the short form will be used, if it exists. For example `EEEE` would return `Monday` and `EE` would return `Mon`. For month M which can be displayed as either text or a number, four or more in a group will display the full version, three will display the abbreviation, and two or less will display the number form.

For numeric elements, the count of characters is the minimum number of digits that the element will take. Shorter numbers will implement leading zeros. For example if the day of the date is 2, `dd` would return `02` and `d` would return `2`.

Any character that is not a-z or A-Z like ";", ":", "@", etc can be inserted anywhere within the string expression and will display as they are entered. You can also insert words and expressions by enclosing them within single quotes (type two single quotes to insert an apostrophe as text). Several examples are listed below:

`printDateTime(10/10/2001 21:15:12, "MMMM dd, yyyy 'at' hh:mm a z")` would return `October 10, 2001 at 9:15 PM PDT`

`printDate(10/10/2001, "EEE MMM dd'th', yyyy")` would return `Wed Oct 10th, 2001`

`printTime(13:22:12, "h 'o''clock' m 'minutes and' s 'seconds'")` would return `1 o'clock 22 minutes and 12 seconds`

`printDateTime(10/10/2001 13:22:12, "MM/dd/yy HH:mm:ss")` would return `10/10/01 13:22:12`

If you do not enter any information for the date/time format, the date will display as yyyy-mm-dd and the time will default to HH:mm:ss.S. So `printDateTime(10/10/2001 13:22:12)` would return `2001-10-10 13:22:12.0`.

toDate() This option will convert a long integer into a timestamp. The syntax is `toDate(<Number>)`. For example, `toDate(1025039526306)` would return `Jun 25, 2002 2:12:06 PM`.

5.1.6.2.8.4. Table Data Functions

The following functions take numeric arguments:

.row() This option will return a value of a column field according to the row index. To do this, you can append a column field with the following syntax: `.row(Row Index)`. The row index is the index value associated with each row in a report starting with 0. For example `{ProductName}.row(3)` would return the fourth value of the ProductName field.

.getScriptValue() This option will return value of a column field according to the row index after applying script. If there is no script applied on the column, it returns the same value as `.row()`. See Section 5.1.7.2.9 - Accessing Values Computed by Other Scripts for more details.

5.1.6.2.8.5. Other Functions

The following functions take string arguments:

first() This option will return the first value of a column field. The syntax is `first(<Column Field>)`. For example `first({UnitPrice})` would return the first value of the UnitPrice field.

last() This option will return the last value of a column field. The syntax is `last(<Column Field>)`. For example `last({UnitPrice})` would return the last value of the UnitPrice field.

5.1.6.2.8.6. Custom Functions

In addition to the functions provided with ReportDesigner, it is also possible to include your own custom functions in the Formula Builder when launching the Report Designer via API. Every Java function is supported.

5.1.6.3. Subreport Formula Access

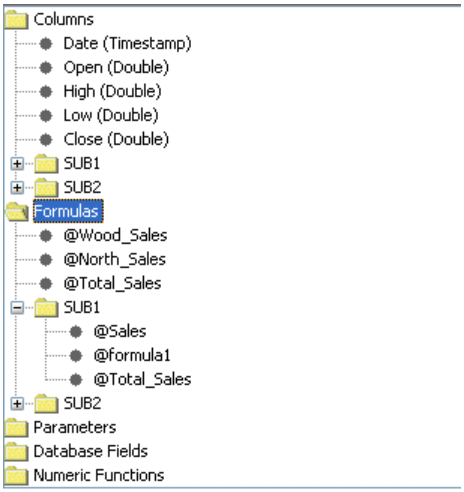
A main report can access its sub-reports' columns and formulas and use them as its own by using the prefix `SUB#` where "#" is the number of the subreport, beginning with 1. For example, the syntax to access the column `UnitPrice` from sub-report 1 is: `SUB1.{UnitPrice}` or `SUB1.COL(i)` where "i" is the column number of `UnitPrice`. Syntax to access the formula `Sales` from sub-report 1 is: `SUB1.@Sales`.



Note

Accessing a subreport column (i.e. `SUB1.{UnitPrice}`) returns a single value - the value from the first row of the column. You can get the value at a particular row by appending `.row(index)` at the end where "index" is the row number beginning with 0, but it will always only return a single value.

In the formula builder's help panel, under both Columns and Formulas, there is a folder for each subreport:



Sub-report columns and formulas can be used to build more complex expressions or scripts (i.e. `SUB1.@Total + SUB2.@Total`).

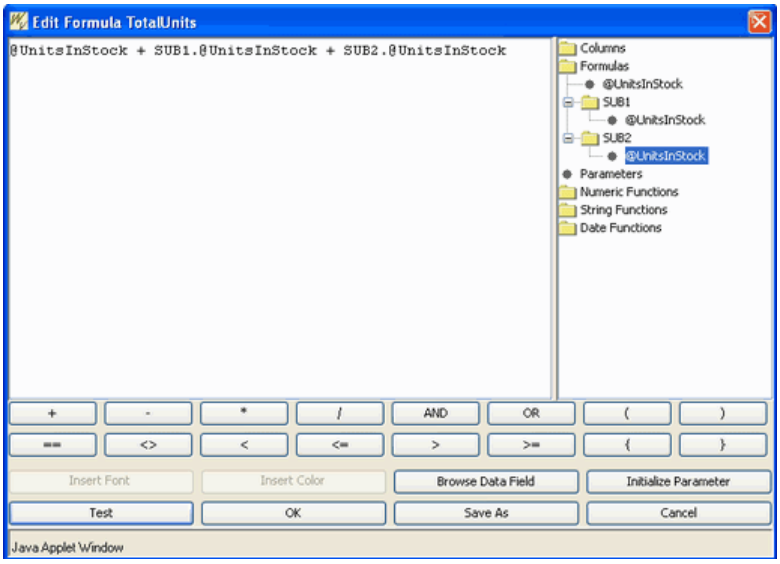
Although sub-report columns and formulas can be used in complex expressions, sub-report columns and formulas cannot be used in aggregation funtions (SUM, COUNT, MAX, AVERAGE, etc). For example, `SUM({UnitPrice})` is a valid formula, but `SUM(SUB1.{UnitPrice})` is not. To accomplish this, instead create a formula `SUB1.{UnitPrice}` and add it as a column into the main report. Then create a second formula which applies aggregation to the first formula (i.e. `SUM(@sub_total)`).

Although the main report can access any of the sub-report's columns or folders, sub-reports cannot access the main report's or any other sub-report's columns or formulas.

If a sub-report formula or column that is used by the main report is deleted, the formula in the main report will display a NULL value when running the report and script will not be applied.

For example, if product information is spread across several databases, but you would like to sum up all the units in stock across the three databases, you could create one main report, add several sub-reports and then create a formula to insert to the report footer adding up all the units in stock. To demonstrate this feature, open a report created from the Inventory XML file with the fields Category Name, Product ID, Product Name, Unit Price, and Units In Stock. Proceed to add subreports to it, one from the Woodview Access database, and one from the Woodview HSQL database, containing the same fields.

Within each subreport, as well as within the main report itself, create a formula called `UnitsInStock`, `SUM({UnitsInStock})`. Then you can create an ultimate formula within the main report using these three functions:



Creating a Formula Using Sub-report Formulas

This formula will take the formula values from each of the three reports and add them together to calculate the final value.

This feature is also applicable to linked sub-reports, although the application of it is different. For example, if you create a summary break report containing customer addresses with a sub-report in the group footer detailing all purchases made by that particular customer, you can sum up all the sales info within the main report.

First, create a formula within the sub-report entitled `Sum_Sales` which sums up all the sales in the sub-report. Return to the main report and create a formula `sub_total`, which is simply `SUB1.@Sum_Sales`. Insert it into the Table Data Section. It can be made invisible if you want to, but it must be in the Table Data Section. Then create a second formula in the main report, `total`, `SUM(@sub_total)`. This formula is added to the report footer and will sum up all the sales columns from all of the sub-reports.

5.1.7. Scripting

Cell scripting or conditional formatting allows you to dynamically modify certain object properties when specific conditions are met. Scripting allows you to highlight certain data values and to present a report that is easier to read. A simple example of this would be a report showing financial results or cash flows. Scripting could be applied to a column field so that negative numbers are shown in red, dynamically highlighting areas of revenue or cash shortages.

In ReportDesigner, scripts are element specific, meaning that in order to dynamically change the properties of a report cell, a script must be applied to that specific cell. Scripts can be applied to labels, column fields, formulas, or report sections.

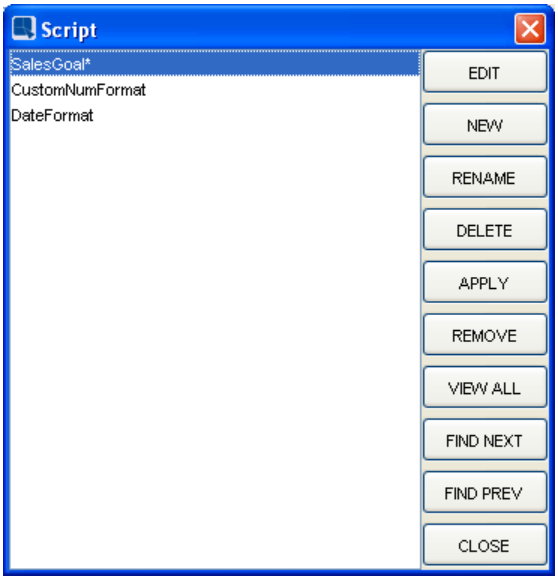


Note

Some of the syntax for ReportDesigner cell scripts has changed greatly between v2.51 and v3.0. In most cases, the deprecated syntax will continue to work correctly. However, it is recommended that you check older templates to make sure the scripts still function correctly if you have upgraded to v3.0 or higher from v2.51 or lower.

5.1.7.1. Creating a Script

To create a script, first select a cell you want to apply the script to and then select *Scripting* from the *Format* menu or click the  *Scripting* button on the toolbar. This will bring up a list of all the scripts that have been defined within the report.

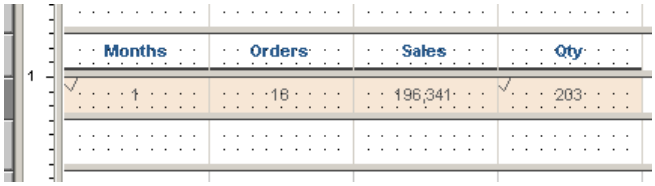


Script List Dialog

Like formulas, scripts are maintained on a report wide basis and can be re-used in multiple places in the report. However, unlike formulas, scripts cannot be inserted into the report, they can only be applied to existing report elements. A star next to a script name in the list indicates that the script is applied to the current cell.

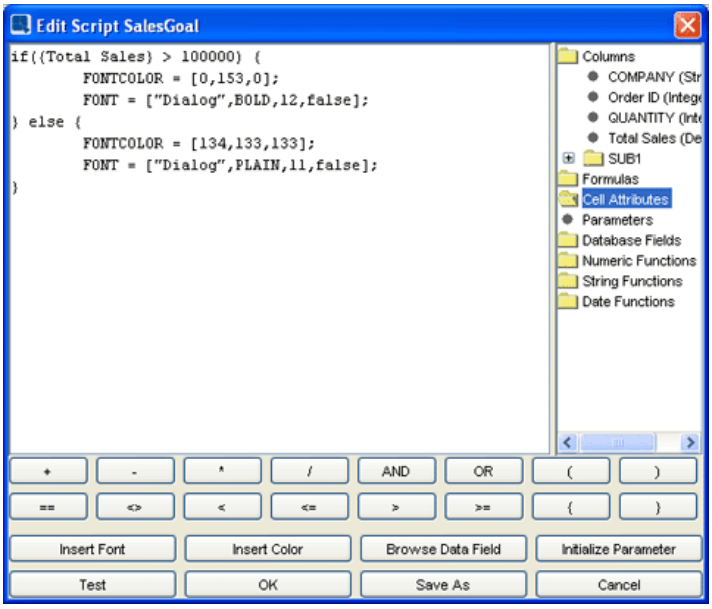
From this dialog, you can edit an existing script, create a new script, rename a script, delete a script, apply a script to the current element, and remove a script from the current element. If a script is being used, it can still be deleted (a message will appear asking if it is ok to proceed) and the script is first removed from any cell before being deleted. If the script is used by a security level, the script has to be removed manually first before it can be deleted. To apply a script to the currently selected report element, select a script from the list and click the *Apply* button. The script list will close and the script will be applied.

Any report element that has a script applied will have a small check mark in the upper left corner of the cell. Scripts will not execute until you preview or export the report.



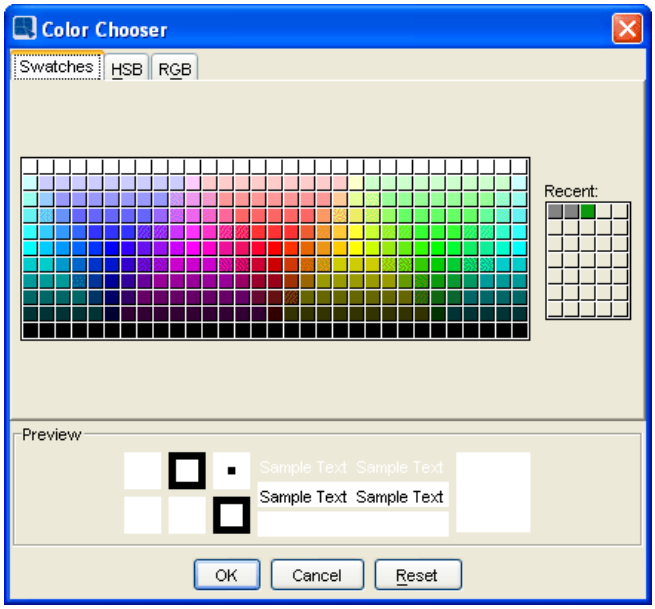
Script Indicators in Design View

If you select to create a new script by clicking the *New* button from the script list dialog, you will be prompted to specify a name for the new script. Once you specify a name, the Formula Builder will open, allowing you to construct a script.



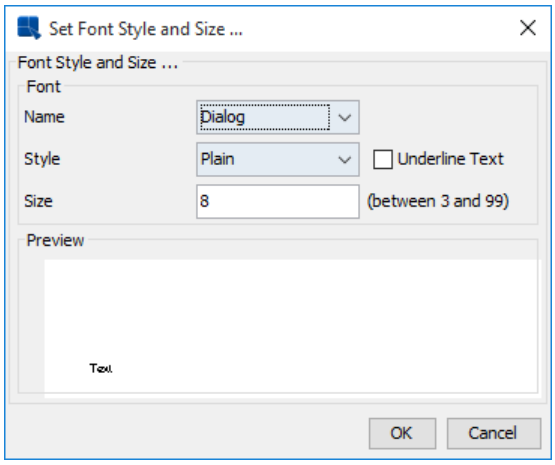
Formula Builder (Scripting)

Scripts are developed in the same formula builder interface as formulas. Although scripts use a different syntax, all of the formula fields, operators, and functions are also accessible in scripts. For more information about formulas and the Formula Builder, please see Section 5.1.6 - Using Formulas & the Formula Builder. In addition to the options available for formulas, there are two additional options available when creating scripts. The *Cell Attributes* folder on the right panel contains a list of the cell attributes (in addition to the value) that can be modified with a script. Attributes and their use are discussed in Section 5.1.7.2.1 - Formatting Actions. The *Insert Color* button will bring up a color selection dialog that allows you to visually pick a color for a cell attribute. Selecting a color will automatically insert an array of RGB values into the script.



Script Color Dialog

The *Insert Font* button will bring up a set font dialog for you to pick from the available fonts and allow you to set styles such as bold and underline.



Script Font Dialog

5.1.7.2. Script Syntax

For most scripts the syntax is fairly simple and consists of two basic components: condition and action.

```
if (Condition) {Action(s);}
```

The condition is generally a Boolean expression derived using one of the Boolean operators discussed in Section 5.1.6.2.7.2 - Boolean Operators For example, `({Quantity} < 5)` would return `true` when the value of the Quantity field is lower than 5. The AND and OR operators can be used to check multiple conditions. For example, `(({Quantity} < 5) AND ({InStock} == "Yes"))` would return `true` when the value of the Quantity field is lower than 5 and the InStock field equals `"Yes"`.

So a simple script that changes the font color to red when negative values are present in the report field would look something like this:

```
if ({SubTotal} < 0) {
    FONTCOLOR=[255,0,0];
}
```

More complex scripts can be created by specifying multiple conditions for multiple actions, using if else syntax.

```
if (Condition 1) {Action 1;}
else if (Condition 2) {<i>Action 2</i>;}
else {<i>Action 3</i>;}
```

There is no limit to the number of nesting levels (if else) that can be specified. Multiple actions for a single condition can also be specified.

```
if (Condition 1) {
    Action 1;
    Action 2;
    Action 3;
}
```

This will perform multiple actions when the condition is met. Note that each line of the script must end with a semicolon.

5.1.7.2.1. Formatting Actions

The following is a list of element attributes that can be modified using scripts.

VALUE: The action syntax that allows you to change the value of the element to which the script is applied is `VALUE=New Value;`. The new value does not necessarily have to be

of the same data type as the report element (for example you can replace number with string). However, changing string to number or number to date can cause formatting irregularities.

BGCOLOR: The action syntax that allows you to change the background color of the element is `BGCOLOR=[R,G,B]`; where R is a number for the red value of the new color, G is a number for the green value of the new color and B is a number for the blue value of the new color. You can automatically generate the RGB array using the *Insert Color* button in the Formula Builder.

FONTCOLOR: The action syntax that allows you to change the font color of the element is `FONTCOLOR=[R,G,B]`; where R is a number for the red value of the new color, G is a number for the green value of the new color and B is a number for the blue value of the new color. You can automatically generate the RGB array using the *Insert Color* button in the Formula Builder.

FONT: The action syntax that allows you to change the font, font style, and font size of the element is `FONT=[Font Name,Font Style,Font Size];`. The font name is the name of the font you want to use (i.e. “Dialog”, “Serif”, etc,). The options for font style are `BOLD`, `ITALIC`, or `BOLD+ITALIC`. The font size is a numeric value indication of the font size. For example, `FONT=[Dialog,BOLD,12];` would change the font to 12 point bold Dialog.

ALIGN: This action allows you to change the horizontal alignment of the data within the cell. The syntax is `ALIGN=Position;`. The position options are `left`, `right`, and `center`.

BORDERCOLOR: The action syntax that allows you to change the border color of the element is `BORDERCOLOR=[R,G,B]`; where R is a number for the red value of the new color, G is a number for the green value of the new color, and B is a number for the blue value of the new color. You can automatically generate the RGB array using the *Insert Color* button in the Formula Builder.

BORDERTHICKNESS: The action syntax that allows you to change the thickness of the element's border is `BORDERTHICKNESS=Thickness;`. The thickness is an integer indicating the number of pixels. For example, `BORDERTHICKNESS=2;` would set the border thickness to two pixels.

HYPERLINK: The action syntax that allows you to set the hyperlink properties of the element is `HYPERLINK=[URL,"Hint","Target"];`. The URL is the location you want the link to point to. The hint is the mouse over hint that will appear and the target is the specified target for the link. If you do not want to specify the hint or target, specify an empty string `""`. Note that you will not be able to click on the hyperlinks in the preview window if you launched the ERES Organizer using `bat` file, but the links will still work if you export the report to DHTML or PDF format.

BOOKMARK: This action allows you to specify a named location within the report. When you export to DHTML, it will become an anchor tag. When you export to PDF, it will become a bookmark within the generated PDF file. The syntax is `BOOKMARK=Name`. The bookmark can be any string value.

XPOSITION: The action syntax that allows you to change the X position of a cell within a section is `XPOSITION=Position;` (in inches or centimeters depending on which measurement is selected). For example, `XPOSITION=1.2;` would set the left edge of the report element 1.2 inches/centimeters away from the left edge of the section.

YPOSITION: The action syntax that allows you to change the Y position of a cell within a section is `YPOSITION=Position;` (in inches or centimeters depending on which measurement is selected). For example, `YPOSITION=0.3;` would set the top edge of the report element 0.3 inches/centimeters away from the top of the section.

WIDTH: The action syntax that allows you to set the width of a report element is `WIDTH=Size;` (in inches or centimeters depending on which measurement is selected). For example, `WIDTH=1.5;` would set the width of the element to 1.5 inches/centimeters.

HEIGHT: The action syntax that allows you to set the height of a report element is `HEIGHT=Size;` (in inches or centimeters depending on which measurement is selected). For example, `HEIGHT=0.25;` would set the height of the element to 0.25 inches/centimeters.

ROTATION: The action syntax that allows you to set the rotation of a report element is `ROTATION=Angle;` (in degrees, where negative value brings counterclockwise rotation, applicable values are `-90, 0` and `90`). For example, `ROTATION=-90;` would set the angle of the element to 90 degrees counterclockwise.

VISIBLE: This action allows you to control whether a cell is visible or not. The syntax is `VISIBLE=True/False;`. Setting visible to false will make the report element invisible.

There is also a special action called **ALERT**. To learn more about alerts in report scripts, please visit Section 9.2.4 - Reports.

In addition to modifying these properties for the current cell, you can also retrieve these properties from other report elements. You can retrieve the value from a report element in the same manner that you do in formulas either `{Column Field}`, `@Formula`, `:Parameter`, etc, or using the `id()` function for other cells or labels. For more information about this, please see Section 5.1.7.2.3 - Using Formulas.

You can retrieve all other attributes using the following syntax `ATTRIBUTE(Report Element)`. For example, `FONTCOLOR(id(RPT_HDR_LB0))` would return an integer array `[R,G,B]` representing the font color of a label in the Report Header section. You can use this to run comparisons based on attributes or to match the attributes of the current cell with those of another report element. For example, `BGCOLOR=BGCOLOR({ProductName});` would change the current cell's background color to match that of the ProductName column.

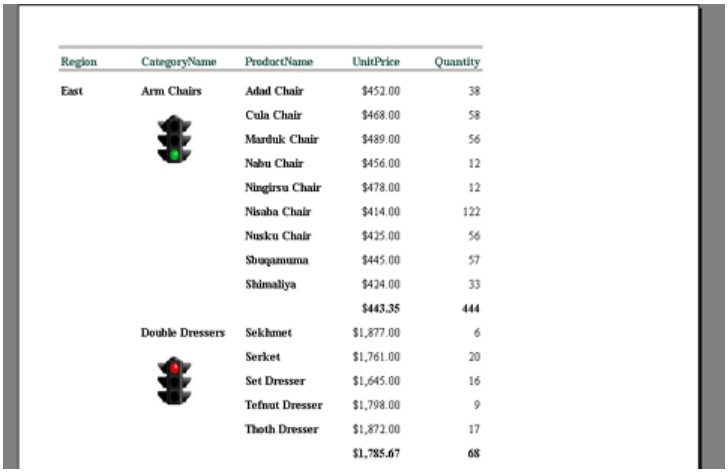
Depending on which type of element the script is being applied to, certain formatting actions may have no effect. For example, you can set scripts on lines and rectangles, but you can only modify their bounds and colors.

5.1.7.2.1.1. Scripting Image URLs

Utilizing the if else conditionals and the Image URL formatting option, you can create image URL scripts that change images dynamically based on data in your report. Suppose you want to include two simple images to help readers determine if the sales for a particular category has achieved the desired goal. Let's say that the sales goal for each category is 100 units sold. Using this information, you can write the following script to display the success image when the goal is met and the failed image when the goal is not met.

```
if(@SUM_Quantity > 100) {
    value = "http://www.quadbase.com/images/green.jpg";
} else {
    value = "http://www.quadbase.com/images/red.jpg";
}
```

Create a new formula cell that has a string data type (you can simply use `""` for the content) and apply this script to the cell. Then set the data formatting for this cell to be `"Image Url"`. More information about data formatting can be found in Section 5.1.3.7.3 - Data Formatting for Formulas and Column Fields. Here is how the report will look like using this script:



Region	CategoryName	ProductName	UnitPrice	Quantity
East	Arm Chairs	Adad Chair	$452.00	38
		Cula Chair	$468.00	58
		Marduk Chair	$489.00	56
		Nabu Chair	$456.00	12
		Ningirsu Chair	$478.00	12
		Nisaba Chair	$414.00	122
		Nusku Chair	$425.00	56
		Shuqamuna	$445.00	57
		Shimaliya	$424.00	33
			$443.35	444
	Double Dressers	Sekhmet	$1,877.00	6
		Serket	$1,761.00	20
		Set Dresser	$1,645.00	16
		Tefnut Dresser	$1,798.00	9
		Thoth Dresser	$1,872.00	17
			$1,785.67	68

Scripted Image URL

5.1.7.2.2. Variables & Arrays

In addition to the if else statements, you can also define variables and arrays in a script. Variables and arrays are defined at the beginning of the script (before any statements).

5.1.7.2.2.1. Variables

If you have a constant or expression that you intend to use in multiple places in a script, it can often be easier to define a variable rather than re-typing the constant or expression in every place where it is used. A variable can be declared at the beginning of a script using the following syntax: `DataType VariableName;`. The data type is the type for the variable value and can be either Number, Boolean, String, or Date.



Note

`DateVar` is used as the data type declaration for Date variables. This avoids confusion with the date() function. The variable name can be anything except an operator or function name.

The following example uses variables which contain an expression.

```
String Line1;
String Line2;
Line1 = {Address1} + "\n";
Line2 = {City} + ", " + {State} + " " + {Zip};

if ({Address2} == NULL) {

    VALUE = Line1 + Line2;
} else VALUE=Line1 + {Address2} + "\n" + Line2;
```

This script pieces together several report fields to create an address entry. The first two lines of the script declare the variables. The script then concatenates fields and strings depending on whether there is an Address2 entry or not. Using the variables prevents having to re-type the expressions. Note that “\n” is used to break the text to a new line.

You can also assign a value to a variable as part of an action in a script. You can use this to modify the variable value based on some conditions and then return the variable value to the cell. This is useful if you only want to write one value statement. The following example assigns a value to a variable:

```
if ({Flags} == "M") {
    LinkAppend = "Main";
} else if ({Flags} == "R") {
    LinkAppend = "Remote";
} HYPERLINK = ["StatPage.html#" + LinkAppend, "", ""];
```

In this example, a variable is defined without an initial value. The value of the variable is assigned based on a string flag in the report. At the end of the script a hyperlink is defined for the element that uses the variable value to specify a page location.

5.1.7.2.2.2. Arrays

In addition to variables, you can also declare arrays at the beginning of a script. Specific values from the array can be retrieved later. This is useful if you have a number of fixed values that you will be using in different places in the script like months or days of the week. An array declaration is similar to a variable declaration and uses the following syntax: `DataType[] ArrayName = [Value1, Value2, Value3, ...];`. Data types for arrays are the same as for variables and can be one of Number, Boolean, String, or Date. The variable name can be anything except an operator or function name. The initial value assignment is optional.

You can retrieve a value from an array using the following syntax: `ArrayName[Index].` The following example uses an array to replace the numeric values for a Month column (1-12) and replaces them with the proper month names.

```
String[] MonthNames = ["January","February","March","April",
"May","June","July","August","September","October",
"November","December"];

VALUE = MonthNames[({Month} - 1)];
```

Notice that the index specified is `{Month}-1`. This is because the first month value is 1 and the first index value for the array is 0.

Like variables, you can also assign or overwrite values in an array as part of an action. The syntax for this is `ArrayName[Index]=Value;`.

5.1.7.2.3. Using Formulas

You can use any of the ReportDesigner formulas from the main report or any sub-reports (for more on accessing sub-report formulas, see Section 5.1.6.3 - Subreport Formula Access) in both the condition and action components of the script, as well as for variable values. Most commonly formulas are used to retrieve values either from the current element or other report elements. You can retrieve the value from column fields using `{Column Field}`, the value from existing report formulas using `@FormulaName` and the user supplied values for query and formula parameters using `:ParamName` and `?ParamName` respectively.

In addition to these methods, there is one additional function that can be used in scripts.

this() - This function returns the value of the current element (i.e. the cell to which the script is applied). The syntax is `this()` without any arguments. The function can be useful if you're reusing scripts or if you want to ensure that the script is referring to the cell to which it is applied. For example `If (this() < 0) {FONTCOLOR=[255,0,0];}` would turn the font red, then the value of the current cell is less than zero.

You can also use the formulas to build expressions for comparison or to change an element's value. For more information about ReportDesigner formulas and their syntax, please see Section 5.1.6 - Using Formulas & the Formula Builder.

5.1.7.2.4. Checking for Nulls

Using cell scripts you can check for null values within report columns. To specify a null as a constant, simply use the word `null` without any delimiters. So a simple script that checks for nulls would look like this:

```
if (this() == null) {
    VALUE="yes";
}
```



Note

Even if you have set the null data handler in Report Designer, you can still use script in this manner. This is because the script will execute before the null handler has a chance to convert the content of the cell.

5.1.7.2.5. Row-Specific Options

Using scripts, you can retrieve and modify values for specific rows of column fields. Normally when you retrieve or replace the value of a column field you will get/replace the first value from the column or the current row if the script is applied to a field in the Table Data section. However, you also have the option of specifying which row you would like to get/replace. To do this, you can append a column field with the following syntax: `.row(Row Index)`. The row index is the index value associated with each row in a report starting with 0. For example `{ProductName}.row(3)` would return the fourth value of the ProductName field.

You can obtain the current row index as well as the total number of rows in the report using the `getRowIndex()` and `getTotalRowIndex()` functions respectively. The following example uses this feature to hide repeated values in a column field.

```
if (({Category}.row(getRowIndex()) == {Category}.row(getRowIndex()-1)) AND
 (getRowIndex() <> 0)) {
    FONTCOLOR=[255,255,255];
}
```

This example assumes that the background color for the report is white and hides the field by adjusting the font color accordingly.

5.1.7.2.6. Loops

In addition to conditional if else statements, ReportDesigner also allows you to add loops to scripts. Used in conjunction with variables and row-specific options/formatting, loops are useful if you need to perform running totals/calculations independent of the current row or section of the report. The basic syntax for loops is:

```
while (Condition) {Action(s);}
```

The specified actions will be performed as long as the condition (Boolean expression) is true.

For example, say you have the following report:

Product	Quantity	Price
Chair	6	$327.00
Chair	4	$218.00
Chair	12	$418.00
Chair	5	$221.00
Chair	18	$248.00
Table	4	$875.00
Table	2	$1,024.00
Table	8	$967.00
Table	7	$1,106.00

Adding a formula that calculates the total quantity to the Table Footer (`sum({Quantity})`) would return `66`. However, say you only wanted to show the total quantity for chairs in the Table Footer. To do this, you write a loop. Applying the following script to the aggregation.

```
Number i;
Number s;
i = 0;
s = 0;

while (i < getTotalRowIndex()) {
```

```
    if ({Product}.row(i) == "Chair") {

        s = s + {Quantity}.row(i);

    }

    i = i + 1;

}

VALUE = s;
```

will return `45` - the sum of the quantity column for `Chair`. The script loops through each row in the report and checks to see if the value for the Product column is “Chair”. If it is, the Quantity value for that row is added to the variable s.

5.1.7.2.7. Section Scripts

In addition to individual cells, scripts can also be applied to report sections. Based on conditions within the report data, individual section options can be turned on/off and section background colors can be set. To add a script to a report section, click on the button with the section name on the left side of the Designer. You can also right click on a blank portion of the section field and select *Scripting* from the pop-up menu.

Writing section scripts is exactly the same as writing scripts for report elements as they use the same syntax. The only difference is that instead of *Cell Attributes* folder in the Formula Builder, you will see *Section Attributes* folder containing various formatting actions that can be performed with section scripts.

5.1.7.2.7.1. Formatting Actions for Section Scripts

The following is a list of section attributes that can be modified using scripts (not all options are available for every report section).

BGCOLOR: The action syntax that allows you to change the background color of the element is `BGCOLOR=[R,G,B];` where R is a number for the red value of the new color, G is a number for the green value of the new color and B is a number for the blue value of the new color. You can automatically generate the RGB array using the *Insert Color* button in the Formula Builder.

PRINT_ON_NEW_-PAGE: This action allows you to turn on/off the print on new page section option. The syntax is `PRINT_ON_NEW_PAGE=True/False;` For more information about the print on new page feature, please see Section 5.1.3.3 - Section Options.

RESET_PAGE_NUM-BER: This action allows you to turn on/off the reset page number section option. The syntax is `RESET_PAGE_NUMBER=True/False;` For more information about the reset page number feature, please see Section 5.1.3.3 - Section Options.

VISIBLE: This action allows you to control whether a section is visible or not. The syntax is `VISIBLE=True/False;` Setting visible to false will make the section invisible. Using this attribute coupled with the nested section feature (Section 5.1.3.1.1 - Nested Sections) allows you to create different headers/footers to display depending on the data.

HEIGHT: This action allows you to set the height of a section. The syntax is `HEIGHT=Size;` (in inches or centimeters depending on which measurement is selected). For example `HEIGHT=0.25;` would set the height of the section to 0.25 inches/centimeters.



Note

Scripts that set or refer to section attributes can only be applied to sections. They cannot be applied to other report elements.

5.1.7.2.8. Commenting Scripts

You can also add comments to a script to explain its design and operation. A comment is preceded by two slashes `//`. Anything following this on the same line is treated as a comment. You can break comments into multiple lines, but each line must begin with `//`.

5.1.7.2.9. Accessing Values Computed by Other Scripts

If you have scripts applied to cells in your report, you can access these values using the syntax `.getScriptValue()`. For example, you may have this data:

```
String, int
Month, Revenue
January, 3567
"February", 6854
"March", 5421
"April", 8425
"May", 5611
"June", 2356
"July", 6543
"August", 2563
"September", 4432
"October", 8841
"November", 1023
"December", 2265
```

If you wish to take this data and create a year to date revenue column, you could use the following script, which would display the scripted value from the previous row added to the revenue from the current row. This script is written to be put on a formula consisting of the number 0 as a column in the table data section, refered to within this script as `{formula1}`.

```
if(getRowIndex()==0)

    VALUE={revenue};
else
    VALUE={revenue} + {formula1}.getScriptValue(getRowIndex()-1);
```

This would yield a report like this:

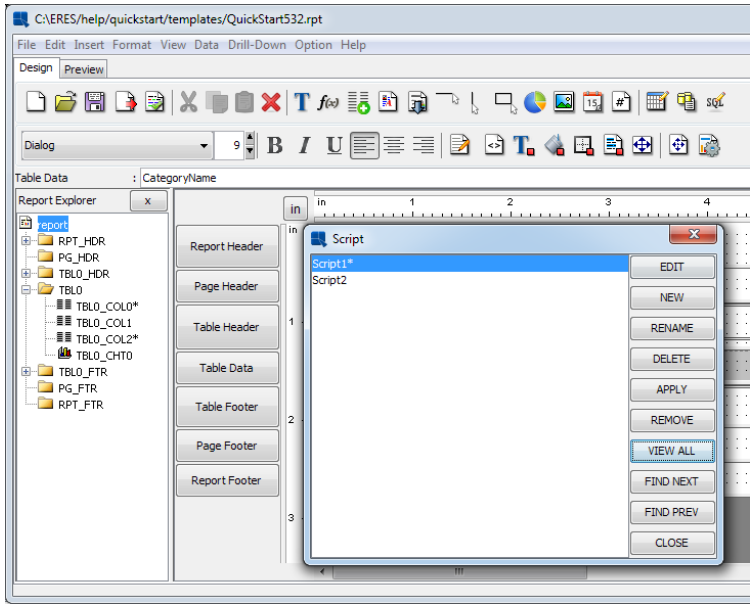
Month	Revenue	Year to Date
January	3,567	3,567
February	6,854	10,421
March	5,421	15,842
April	8,425	24,267
May	5,611	29,878
June	2,356	32,234
July	6,543	38,777
August	2,563	41,340
September	4,432	45,772
October	8,841	54,613
November	1,023	55,636
December	2,265	57,901

5.1.7.3. Managing Scripts

For reports that contain a number of scripts or scripts applied to a number of elements, there are several options available to help you manage and find the elements that the scripts are applied to. In the script dialog, you will notice that in addition to the standard options such as remove, delete, and rename, there are also buttons to view all, find next, and find previous elements. Make sure that you select a script before using these buttons.

VIEW ALL: This button will open the Report Explorer if it is not already open and highlight every element that the selected script has been applied to with an as-

terisk. The tree will automatically expand the sections that have highlighted cells. Clicking on one of these elements will center the designer view on that element allowing you to change or remove the applied script. The highlights will be changed if you select another script and click on VIEW ALL again. Closing the script dialog box will remove all the highlights and the explorer tree will collapse.



View All Result

FIND NEXT and FIND PREVIOUS

These buttons are used to cycle through elements that the selected script has been applied to. If you already have an element selected in the Designer when clicking on these buttons, the search will begin at the current element. If no elements are selected, the search will begin at the beginning of the report. If you select another script and click FIND NEXT, the asterisk will shift to the newly selected script and the first element that has the script applied will be highlighted in the designer. *FIND NEXT* and *FIND PREVIOUS* only cycle through visible elements. If you want to see invisible elements with the script, use *VIEW ALL* which shows all invisible elements as well as visible elements.

For more information about the Report Explorer, please see Section 5.1.3.8 - The Report Explorer

5.1.8. Drill-Down

Often reports are used to display large amounts of data. However, with large data sets (particularly if the data is fairly detailed) displaying everything in one report may not the best way to present the information. One solution to this problem is to create a top-level report that displays only summarized data and allow users to click through to see underlying data. This is the concept of drill-downs.

In ReportDesigner, users can easily display data in this way. Links can be placed on column fields in the primary (top-level) report to secondary (sub-level) reports. The sub-level reports use parameterized queries that accept the value from the primary report as the input. Hence, sub-level reports will only display data related to the particular value on which the user has clicked.

For example, say you are designing a report to show sales data for an entire year. Rather than creating a report that displays the entire year's worth of data, you could design a top-level report that shows aggregated data by month.

Month	Units Sold	Total Sales
January	5	$10,224.12
February	4	$12,286.14

Month	Units Sold	Total Sales
March	6	$14,415.16
April	5	$12,316.11
May	2	$9,482.64
June	3	$10,116.12
July	7	$17,624.18
August	5	$10,361.54
September	4	$10,116.82
October	3	$8,612.41
November	5	$10,552.24
December	11	$22,614.89

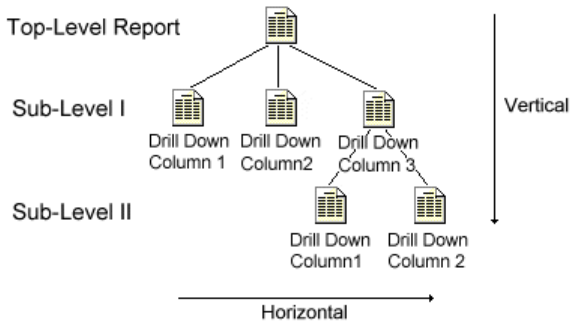
Users viewing the report could then drill-down to see more data for each month. In this case, say the user clicks on *March*. They would then be taken to an underlying report that offers detailed sales information for that month.

Sales Data For March 2001				
Order #	**Customer**	**Product**	**Quantity**	**Total**
10042	Allied Furniture Emporium	Chair	2	$6,425.16
10051	Furniture World	Table	1	$3,114.12
		Dresser	1	$2,116.83
10068	Domus Home & Garden	Chair	1	$800.65
		Dresser	1	$1,958.40
Total:			**6**	**$14,415.16**

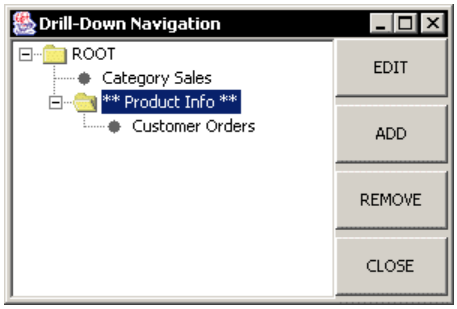
From this report, another layer of drill-down can be added, allowing users to see more detailed customer information. In this case, say the user clicks on *Allied Furniture Emporium*. They would be taken to a third report that shows more information about the customer.

Customer Information		
		Address:
Company:	Allied Furniture Emporium	384 Broad St. Littletown, NY 18322
Contact Name:	Matilda Gladwaller	
Order History:		
Year	**Orders**	**Total Sales**
1999	14	$52,614.18
2000	18	$68,115.89
2001	12	$38,812.12

Using drill-down in this manner allows large amounts of information to be presented in a way that is easy to understand and navigate. ReportDesigner allows horizontal and vertical layers of drill-down. This means that different columns in that same report can have drill-down (horizontal), or that sub-level reports can have subsequent drill-down reports (vertical). Users can also add drill-down layers to sub-reports.



Within the Report Designer, drill-down is controlled from the Drill-Down menu. Selecting the *Navigate* option from this menu brings up the Navigation window. From this window you can add and edit drill-down reports.



Drill-Down Navigation Dialog

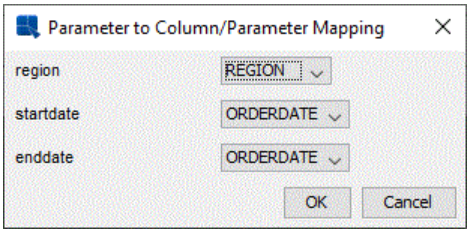
The left side of the Navigation window displays the hierarchy of drill-down reports (see the above diagram). The ROOT is the top-level report. The level that you are currently editing is marked with "**". To edit a different report, select it and click the *EDIT* button on the right side. That report will then open in the Designer, allowing it to be customized. Reports can also be removed by selecting the report and then clicking the *REMOVE* button.

5.1.8.1. Creating a Drill-Down Report

The first step in creating a drill-down report is to create the top-level report. Because drill-down relies on relationships between data sets, only reports that use parameterized queries, XML queries, or class files can be used for drill-down layers. When designing the top-level report, bear in mind that at least one of the columns or parameters will be passed as a parameter to any sub-level reports. After you finish designing the top-level report, go to the Drill-Down menu and select *Navigate*. The navigation window will open. Select the report under which you would like to add the drill-down layer (if the report has no drill-down, only the ROOT node will be visible), and click the *ADD* button. If you are adding drill-down to a sub-report, first go to the sub-report tab and then load the navigation dialog.

You will then be prompted to create a new report or use an existing report for the drill-down layer. You can use any existing report; however any report for a drill-down layer must have a parameterized query or class file as the data source. If the existing report contains drill-down levels of its own, all drill-down levels will be imported as well. If you select to create a new report, the Report Wizard will open, allowing you to select a data source, report type, and column mapping.

The next step is to map the columns or parameters in the top-level report, to the parameters in the sub-level report. You will be prompted to do this when you select an existing report for the drill-down layer, or when you select the data source for a new report for the drill-down layer. In either case, a dialog will appear prompting you to map the columns or parameters from the top-level report.



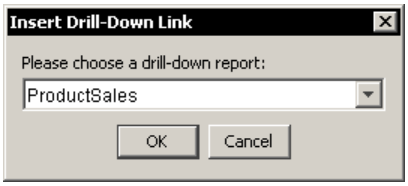
Parameter to Column/Parameter Mapping Window

The options that are available in the drop-down menu are based on data type. For example, if your parameterized query takes a string as a parameter, only columns containing string data can be mapped. If there are multiple parameters for the sub-level query, there will also be the option *None*. This will allow you to create unmapped drill-down reports. For more information, see Section 5.1.8.5 - Unmapped Drill-Down. Once you specify the drill-down mapping, you will be prompted to specify a name for the sub-level report. Now the sub-level report will open in the Design window, allowing you to format and customize it.

5.1.8.1.1. Drill-Down Links

By default, a link will be placed on the column in the top-level report which is mapped to the sub-level report. Instead of using the default link, you can place the link on a different column, or multiple columns using the *Drill-Down Link* option from the Drill-Down menu.

To add or remove a drill-down link from a report column, first make sure that you are currently editing the correct template (top or higher level). Then select a column field and select *Drill-Down Link* from the Drill-Down menu. This will bring-up a drop-down list allowing you to select which sub-level report you would like to link to. To remove a drill-down link, select the blank option at the top of the drop-down list.



Drill-Down Link Dialog

5.1.8.1.1.1. Linking from Charts

In addition to placing links on a column, you can also set drill-down links on a chart. With this feature, users can click on the data point in a chart and traverse to the next level of drill-down. To set a drill-down link on a chart, select the chart object in the Report Designer and select *Drill-Down Link* from the Drill-Down menu. Next, select the sub-level report just as you would for a column.

When you set a drill-down link in the chart depending on the elements in the chart and the number of parameters in the sub-level report, the category, data series, and/or sum by value for the data point that is selected are passed as the parameter value to the sub-level report. In order for this feature to function correctly, the chart will need to use the same data source as the report. The chart should also have the categories mapped to the column that was selected to match the parameter in the sub-level report. For more information about column mapping in charts, please see Section 5.2.2.2 - Basic Data Mapping and Section 5.2.3 - Chart Types and Data Mapping.

5.1.8.2. Multi-Value Drill-Down

In addition to drilling on one field at a time, ReportDesigner supports the concept of multi-value drill-down. Using this feature, you can drill into several different values from a top-level report at once.

For example say you have a report detailing sales information for 100 different customers. Using the multi-value drill-down, users could select the customers that most interest them and drill into their records at once instead of looking at each customer individually.

5.1.8.2.1. Creating a Multi-Value Drill-Down Report

Adding layers of multi-value drill-down to a report is almost exactly same as adding regular drill-down layers. The only difference is in the query in the lower-level report. To create a multi-value drill-down layer, the lower-level report must have a multi-value parameter in its query that is mapped to a column in the higher-level report.

Using the previous example, say the query for the lower level has a multi-value parameter for a customer id - something like `Where Customers.CustomerID IN (:CustID)`. If the `CustID` parameter is mapped to the `Customer ID` column in the top-level report, this will automatically create a multi-value drill-down layer. For more information about multi-value parameters, see Section 4.1.2.2.2.1 - Multi-Value Parameters. If you do not wish to create a multi-value drill-down layer, simply use a single value parameter in the lower-level report.



Note

You can have only one report per level when using multi-value drill-down. It does not support linking different columns from the same report to different reports.

You can place column links in multi-value drill-down reports; however, a dialog or form is used to allow users to select multiple column values and drill to the next level. For more information about this, please see Section 5.1.8.4 - Viewing Drill-Down Reports.

5.1.8.3. Crosstab Drill-Down

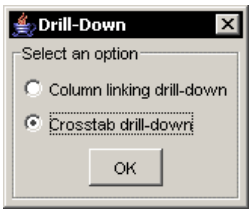
ReportDesigner provides a special implementation of the drill-down features when the top level report is a crosstab report. Using the crosstab drill-down features, users can click on a cell in the crosstab matrix and drill into information for that particular row or column. For example, assume you have the following report showing sales volume by product and region.

	East	South	Midwest	West
Chair	144	208	131	108
Cabinet	208	114	158	206
Dresser	100	101	112	109

Using standard drill-down configuration, if a user clicks on one of the hyperlinked cells, only the values from that row could be passed to the lower-level report. However, with crosstab drill-down, the column break and row break values for that cell are passed to the lower-level report. Therefore, if the user clicks on the cell value *144*, the values `"Chair"` and `"East"` would be passed to the parameters in the lower-level report.

5.1.8.3.1. Creating a Crosstab Drill-Down Report

When you select to add a drill-down layer to a crosstab report, you will be prompted whether you would like to create a standard (column linking) drill-down presentation or use crosstab drill-down.



Crosstab Drill-Down Options

If you select to use crosstab drill-down, the setup will continue just as it would for any other drill-down presentation. The only difference comes with parameter mapping. Instead of mapping parameters from the lower level report to columns in the main report, they are mapped to the row break and column break values for the top-level report.



Parameter-Column Mapping for Crosstab Drill-Down

Once you add the drill-down layer, the cells (column break values) in the crosstab matrix for the top level report will automatically be linked to the layer you have added.

5.1.8.4. Viewing Drill-Down Reports

You can directly navigate and view drill-down reports within the Report Viewer Applet (and the Preview window). When you open a report containing drill-down, you can click on any entry in a column with a drill-down link. You will be taken to the sub-level report, which will take the value for that row as the parameter.

For multi-value drill-down reports, you can select *Select Multiple Drill-Down Values* from the *Data* menu in the Preview window or from the Viewer pop-up menu. This will bring up a dialog containing all the values in the mapped column. You can then select the values you want and click the *Ok* button to go to the next level.



Multi-Value Drill-Down Selection Dialog



Note

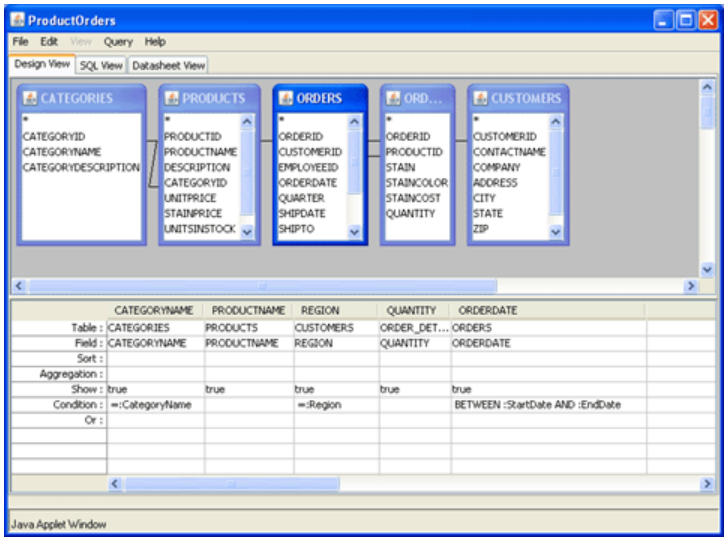
The templates for sub-level reports are saved in DrillDown directory. If you are viewing reports on another machine, the Report Viewer will look for the templates in that directory, so they have to be moved as well.

Drill-down reports can be exported to any format; however, only the DHTML and PDF formats can retain the drill-down functionality (multi-value drill-down reports only support DHTML). Any other format will only export the top-level report.

5.1.8.5. Unmapped Drill-Down

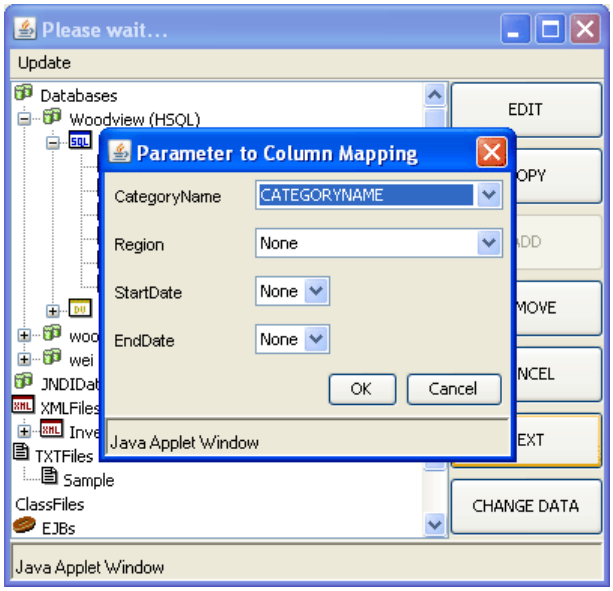
When you create a sub-level query with multiple parameters, you have the option of leaving some of these parameters unmapped, so that they can be prompted after clicking a drill-down link. To use this feature when selecting mapping for parameters, select **none** for the parameters you wish to remain unmapped. The sub-level report also contains a Data menu option *Preview Parameter Prompt* which determines if the parameter prompt will be displayed when you drill-down from the root level. If this level has never been previewed before, it will be turned off and use the default parameters. Otherwise, it will save the last used parameter values for future use.

Each sub-level drill-down report requires at least one mapped parameter. Thus, in order to create a report with unmapped parameters, the sub-level query must have at least two parameters. For example:



Drill-Down Level Query

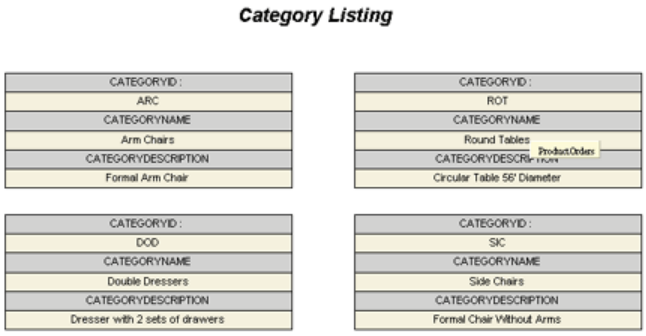
Map at least one parameter to a database column and select **none** for the other(s).



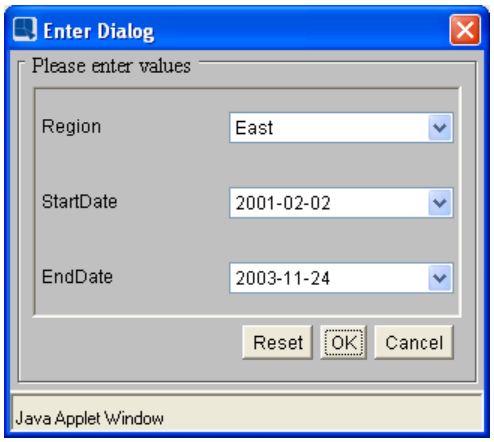
Creating Unmapped Parameters

The remainder of the creation of the drill-down report is like any other.

When you preview the report, the root report will be displayed as normal (with a parameter prompt if the root report contains parameters).



If you click on a link and the *Preview Parameter Prompt* menu option for the drill-down report is checked, you will see the following prompt. Notice that the `CategoryName` parameter is not shown:



Parameter Prompt for Unmapped Drill-down

The resulting report will look like this:

Product Orders

CATEGORYNAME	PRODUCTNAME	REGION	QUANTITY	ORDERDATE
Round Tables	Atun Table	East	21	Jul 11, 2003
	Apep Table	East	6	Nov 24, 2003
	Anubis Table	East	16	Feb 14, 2003
	Atun Table	East	6	Feb 21, 2003
	Ningizida Table	East	4	Jan 16, 2003
	Amon Table	East	12	Feb 14, 2003
	Ningizida Table	East	18	Jun 22, 2002
	Anubis Table	East	4	Jun 22, 2002
	Bast Table	East	16	Dec 4, 2001
	Bast Table	East	12	Mar 12, 2002
	Anubis Table	East	13	Feb 2, 2001
	Apep Table	East	15	Feb 24, 2001

From the DHTML Viewer or in an exported file, the default values will be used, but they can be changed from the filter options.

5.1.9. Sub-Reports

A sub-report is a report within a report. The sub-report is nested entirely within a section of a primary report. Each sub-report has its own data source and design. Using sub-reports you can create reports that use un-related data. You can also create a report that has more than one data table.

In ReportDesigner, sub-reports are generally unrelated to the primary report (however parameterized sub-reports can be linked to column fields in the primary report, see Section 5.1.9.4 - Linked Sub-Reports). Each sub-report uses its own data source and it can be designed independently from the primary report. It also runs entirely within the section in which it is placed in the primary report. You can either use an existing report (`.pak, .qrp, .rpt` or `.xml`) template or create a new one for the sub-report.

5.1.9.1. Adding a Sub-Report

To add a sub-report to an existing report, click the  *Insert Sub-Report* button on the toolbar or select *Insert Sub-Report* from the *Insert* Menu. A second menu will open giving you the option of adding a new sub-report or re-using an existing one. To insert a new sub-report, select *New...*.

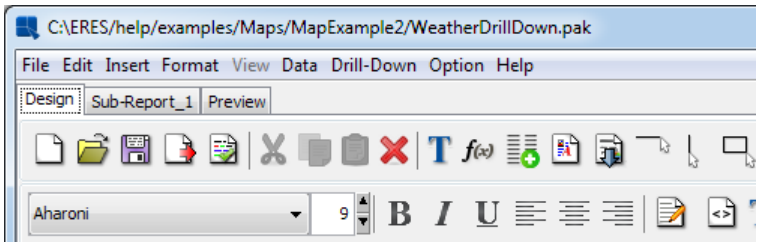
The cursor will turn to cross. Position it where you want to insert the sub-report and click. A dialog box will appear asking you whether you want to use an existing report or create a new one.

You can use any existing report in either `.pak, .rpt` or `.xml` format as the sub-report. If you select to use an existing report, you will be prompted to specify the location of the report file. After you specify the file, it will be inserted as a sub-report.

If you select to create a new report for the sub-report, the Report Wizard will launch, allowing you to go through the process of designing the report from data source selection to report mapping. Once you finish with report mapping, the new report will be inserted as a sub-report.

5.1.9.2. Editing a Sub-Report

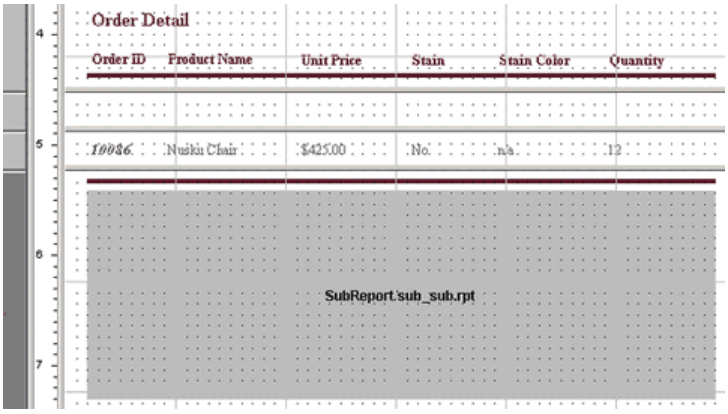
After a sub-report has been added, a new tab will appear at the top of the Report Designer window.



Sub-Report Tab

Selecting the *Sub-Report* tab will open a new design pane that allows you to modify the sub-report. You can modify the sub-report in the same way as you modify any report. For each sub-report you add, a new tab will appear. When you have a sub-report, the *Preview* tab becomes context sensitive. If you click on the *Preview* tab from the *Sub-Report* tab, the preview pane will only contain the sub-report. If you click on the *Preview* tab from the *Design* tab, you will preview the main report with the sub-report embedded.

Within the Design window, the sub-report will appear as a gray rectangle. It can be moved and resized like any object. You can also edit a sub-report by double-clicking on the corresponding rectangle in the Design window.



Sub-Report in Design Window

Sub-reports run entirely in the section in which they are placed and will generally repeat each time the report section repeats. This means that if the sub-report is placed in the table header section, it will only run once (unless the *Print on Every Page* option is enabled). However, if the sub-report is placed in the Group Header section, the sub-report will repeat for each group in the report.

The exception to this is the table data section. Sub-reports inserted into the table data section will not repeat for each row, but will rather run alongside the data in the primary report. This allows you to create a report with the effect of two data tables running next to each other.

5.1.9.2.1. Sub-Report sizing

The size of the gray rectangle in the main report Design window represents the size of the sub-report, so if the sub-report is larger than the space you allow in the Design window, portions of the sub-report may be truncated.

Fixing a sub-report's size within the main report may not be the best solution if the number of records in the sub-report vary. You can solve this by setting the sub-report to have variable height. Dynamic sizing for sub-reports is set in the same way as for other report cells using the resize to fit option. For more information about this feature, please see Section 5.1.3.7.11 - Moving and Resizing Report Elements.

You can also set the width of the sub-report to be dynamic. This feature is most useful for crosstab type reports where the number of columns in the report can vary as the data or filters change. This sets the sub-report to always draw wide enough to fit all the visible columns in the sub-report. To set dynamic width for the sub-report, right click on the sub-report in the Design window and select *Auto Resize Width* from the pop-up menu (To disable this feature, open the pop-up menu and click on this option again).



Note

If you set a sub-report to have variable height, any objects placed directly below the sub-report in the same section of the primary report will not shift to accommodate the resized sub-report. Hence, these objects may be overlapped by the sub-report.

To ensure that objects are not overlapped, it is recommended that you place sub-reports with dynamic height in their own report sections. If necessary, you can add nested sections to the report. For more information about this, see Section 5.1.3.1.1 - Nested Sections.

5.1.9.3. Removing a Sub-Report

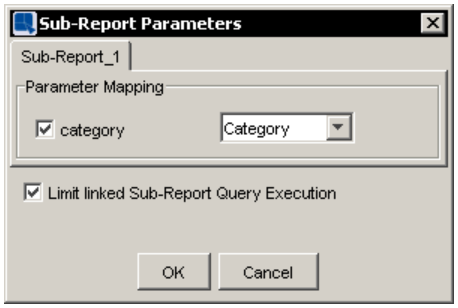
There are two ways to remove a sub-report from the primary report. You can simply select the sub-report object in the primary report and delete it. This will remove the sub-report from the primary report but it will keep it available. This means that the sub-report tab is still active and the sub-report can be re-inserted into the primary report.

The second method is to select *Remove Sub-Report* from the *Edit* menu. A second menu will open with a list of all the sub-reports embedded in the current report. Select the sub-report you want to remove. The sub-report tab will be removed and the sub-report will no longer be available.

5.1.9.4. Linked Sub-Reports

Sometimes you may want sub-reports to be coordinated with the data in the primary report. For example, if you have grouped data, you may want to filter a sub-report so that it shows data pertinent to each group. You can use linked sub-reports to do this.

Sub-reports are linked to primary reports using query parameters. So to create a linked sub-report, the sub-report must have a parameterized query. After you create the sub-report, you can link it by selecting *Sub-Report Parameter Mapping* from the *Data* menu. This will bring up a dialog allowing you to map column fields from the primary report to query parameters in the sub-report. This operation is similar to the one used in drill-down.



Sub-Report Parameter Mapping Dialog

The dialog contains a tab for each parameterized sub-report. Each parameter in the sub-report can be mapped to a different column field in the primary report.

Normally when linked sub-reports execute, the sub-report will issue its query each time it runs using the current value of the report column. However, if your report is configured such that the sub-report executes a large number of times, these repeated queries can have an impact on performance. If you check the last option in this dialog called *Limit Sub-report query execution*, ReportDesigner will try to merge the sub-report queries and limit the number of calls to the database when the sub-report runs. Note that this option will only take effect if the sub-report uses a database query to retrieve the data. This will not effect sub-reports that use parameterized classes or xml files as the data source.

When using linked subreports, it is recommended that you do not use the cascading parameters functionality, as this may potentially cause incorrect data to appear (depending on the cascading parameters and the order in which they are presented). The default value of the parameter will always be selected instead of the value passed by a column in the main report.

5.1.9.4.1. Using Linked Sub-Reports

Although linked sub-reports will run in any configuration and in any report section, it is important to be aware of the behavior of sub-reports when creating a linked presentation. As noted in Section 5.1.9.2 - Editing a Sub-Report, sub-reports will only run as many times as the section in which they are placed and will only run once when placed in the Table Data section.

Therefore, if you place a linked sub-report in a section that runs only once (Table Header, Report Header, Table Data, etc), the sub-report will only run once and it will use the first value from the column to which it is linked to filter the report. To generate a linked presentation where the sub-report runs for each value in the column to which it is linked, you should use a summary break layout. By setting the column as a row break and placing both the column and the sub-report in the Group Header section, you will get the desired result.



Note

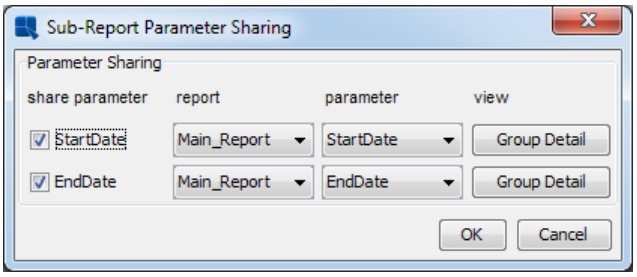
If a relationship between the column and the sub-report already exists in the database, you may want to generate a side-by-side master & details layout instead of linked sub-reports. Using the master & details layout will result in better performance. For more information about this, please see Section 5.1.2.4.1 - Data Mapping.

5.1.9.5. Parameter Linking

In addition to taking a value from a column in the main report, a sub-report can also share parameter values with the main report or with other sub-reports. In this arrangement, a user can enter values to filter the main report's data and have the same filters apply to the sub-report data.

Parameter linking is only available if parameters share the same data type. By default, if you create a sub-report that has parameters with the same name and data type as those in the main report, they will be linked. When parameter values are supplied to the main report, they are also automatically passed to the sub-report.

The parameter linking behavior can be modified by navigating to the sub-report that you want to modify and selecting *Sub-Report Parameter Sharing* from the Data menu (in the Sub-Report tab). This will bring up a dialog allowing you to modify the linking behavior for the sub-report parameters.



Parameter Sharing Dialog

For each parameter defined in the sub-report, you can enable/disable the linking. You can also set which parameter in the main report or in another sub-report you want to link to. The *Group Detail* button will launch a dialog showing the current relationships for that parameter.

Once you specify the settings you want for all parameters, click the *Ok* button to apply the changes.

For more about query parameters and initialization, please see Section 4.1.2.2.2 - Parameterized Queries of this guide.



Note

Parameter linking will be overridden if you select to map a column field from the main report to the sub-report parameter as described in Section 5.1.9.4 - Linked Sub-Reports.

5.1.10. Template Security

One of the most common needs with a reporting tool is to keep the number of different report templates at a minimum. In ERES, the administrator can grant access to specific reports for users and groups in the system. However, simply granting or denying access may not be enough to facilitate template re-use. To help with this, ERES provides a template security features that allow the administrator to configure different levels of access to the same report template for different users/groups in the system.

For example, say you have a group in ERES called *Sales*. The group all has access to a report that shows total sales grouped by sales region. Sales executives viewing the report need to see the totals as well as the contributions by each region, while individual account managers would only need to see the data for their region. Rather than creating a different report for each sub-group within the *Sales* group, the same template can have security settings applied. The base template for this scenario would look like this:

Sales Report				
Total Units Sold: 29 **Total Sales:** $24,449.69				
Eastern Region				
Order #	**Quantity**	**Product**	**Unit Price**	**Total Sales**
1001	2	Chair	$325.16	$650.32
	1	Table	$1,114.18	$1,114.18
			Order Total:	**$1,764.50**
1004	3	Table	$1,114.18	$3,342.54
	2	Dresser	$1,518.60	$3,037.20
			Order Total:	**$6,379.74**
Total for Eastern Region:				**$8,144.24**
Midwestern Region				
Order #	**Quantity**	**Product**	**Unit Price**	**Total Sales**
1003	4	Chair	$325.16	$1,300.25

Sales Report				
	2	Table	$1,114.18	$2,228.36
	1	Dresser	$1,518.60	$1,518.60
			Order Total:	**$5,047.21**
Total for Midwestern Region:				**$5,047.21**
Southern Region				
Order #	**Quantity**	**Product**	**Unit Price**	**Total Sales**
1005	2	Table	$1,114.18	$2,228.35
	1	Dresser	$1,518.60	$1,518.60
			Order Total:	**$3,746.95**
1007	3	Chair	$325.16	$975.48
	1	Table	$1,114.18	$1,114.18
			Order Total:	**$2,089.66**
Total for Southern Region:				**$5,836.61**
Western Region				
Order #	**Quantity**	**Product**	**Unit Price**	**Total Sales**
1006	2	Dresser	$1,518.60	$3,037.20
	4	Chair	$325.16	$1,300.25
	1	Table	$1,114.18	$1,114.18
			Order Total:	**$5,451.63**
Total for Western Region:				**$5,451.63**

Using the security features, this one template could be used to meet the requirements of all the members of the *Sales* group. In this example, assume that the report has a query parameter that filters based on region.

For executives, the view above is what they would want to see when viewing the report, so a security level called `Management` could be created that retrieves data for all regions. This is accomplished by mapping all the available values for the region parameter to the `Management` level. In this example the values for the Eastern, Midwestern, Southern, and Western regions would be assigned to the security level. Note that in ReportDesigner, in order to supply more than one value to a query parameter, you need to define a multi-value parameter. For more information about this, please see Section 4.1.2.2.2.1 - Multi-Value Parameters.

For the account managers, different security levels could be created for each region. The above example has four regions, so security levels `East`, `Midwest`, `South`, and `West` would be added. Since account managers should not see the aggregate data, cells containing the summaries in the header would be hidden for the region levels using the cell-level security features discussed in Section 5.1.10.2 - Cell-Level Security. To make sure that each level returns data for the appropriate regions the parameter value for the Eastern region would be mapped to the `"East"` security level, the parameter value for the Midwestern region would be mapped to the `Midwest` security level, the parameter value for the Southern region would be mapped to the `South` security level, and the parameter value for the Western regions would be mapped to the `West` security level. For more information about setting security levels to apply parameter values, see Section 5.1.10.2.1 - Security Parameters.

In the Organizer, the members of the *Sales* group can be associated with the defined security levels. Then when the report is run, the security level for the user who has requested the report is applied and the parameter values associated with that level are automatically supplied to the query without prompting the user. For example, here is the same report run by an account manager in the eastern region:

Sales Report
Eastern Region

Sales Report				
Order #	**Quantity**	**Product**	**Unit Price**	**Total Sales**
1001	2	Chair	$325.16	$650.32
	1	Table	$1,114.18	$1,114.18
			Order Total:	**$1,764.50**
1004	3	Table	$1,114.18	$3,342.54
	2	Dresser	$1,518.60	$3,037.20
			Order Total:	**$6,379.74**
Total for Eastern Region:				**$8,114.24**

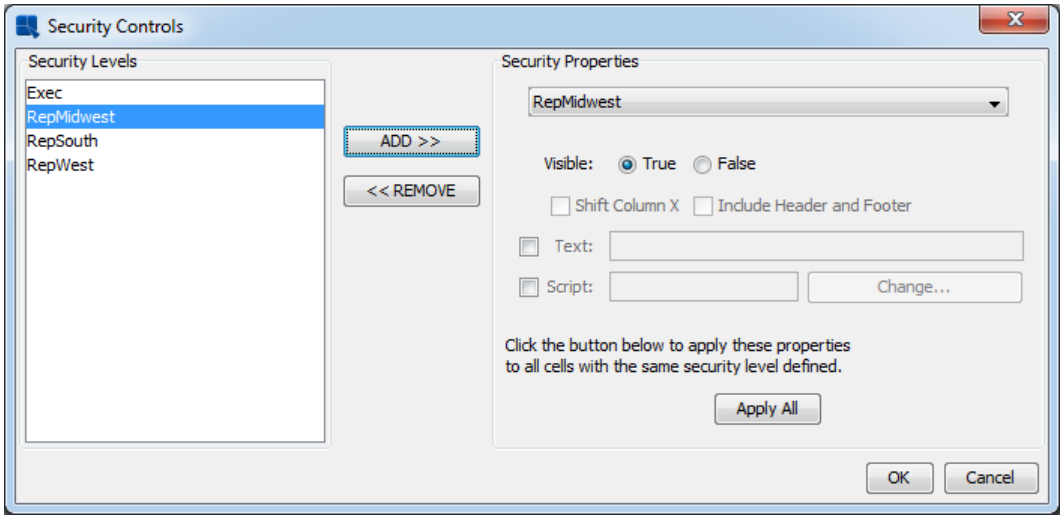
5.1.10.1. Security Levels

Security levels are unique groups of settings that can control cell and parameter behavior in reports. Levels are created by the administrator in the Organizer interface. The administrator can also associate users and groups with specific security levels.

Within a report, different cell behaviors and filtering can be associated with a specific security level. At run-time, based on the user login, the security level assigned to that user is applied to the report.

5.1.10.2. Cell-Level Security

For any cell/visible column in the report, you can set different behaviors for different security levels. To apply security settings to a cell, either select the cell and select *Security* from the *Data* menu, or right click on the cell and select *Security* from the pop-up menu. This will bring up a dialog allowing you to specify security settings for that cell.



Cell Security Settings Dialog

The dialog contains a list of all the defined security levels in the Organizer.

Security levels can be added/removed/modified in the Organizer.

To add security settings for a report cell, select the level for which you want changes to take effect and click the *Add* button. The security level will then appear in the drop-down list on the left side of the cell security dialog. Now you can modify the cell for that security level. There are three basic controls available for cells. You can render the cell invisible for a security level. If the cell is visible, you can either provide custom text for the cell or you can apply a cell script. For more information about the cell scripts, please see Section 5.1.7 - Scripting. If you select to make a column (cell in Table Data section) invisible, make sure to select the column header along with the data cell so they both contain the same security setting, then you can check the *Shift Column X* option to make the other columns shift to the left when the security level is applied.

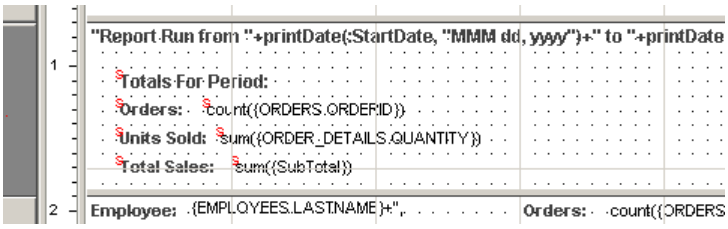


Note

When you specify to change the text of a cell, it will not work for default column headers. This is because column headers are unique elements that retrieve their value from the data source. If you wish to modify the text of a column header, you will need to convert it to a static label first. For more information about this, please see Section 5.1.3.7.2 - Editing Elements.

You can propagate this behavior to any cell that has controls for this security level by clicking the button on the right side. You can specify controls for as many different levels of the cell as you want.

Once you finish specifying the security settings for the cell, click the *Ok* button. You will go back to the Design window where a red *s* will appear in the upper left corner to indicate that the security settings have been applied to the cell.

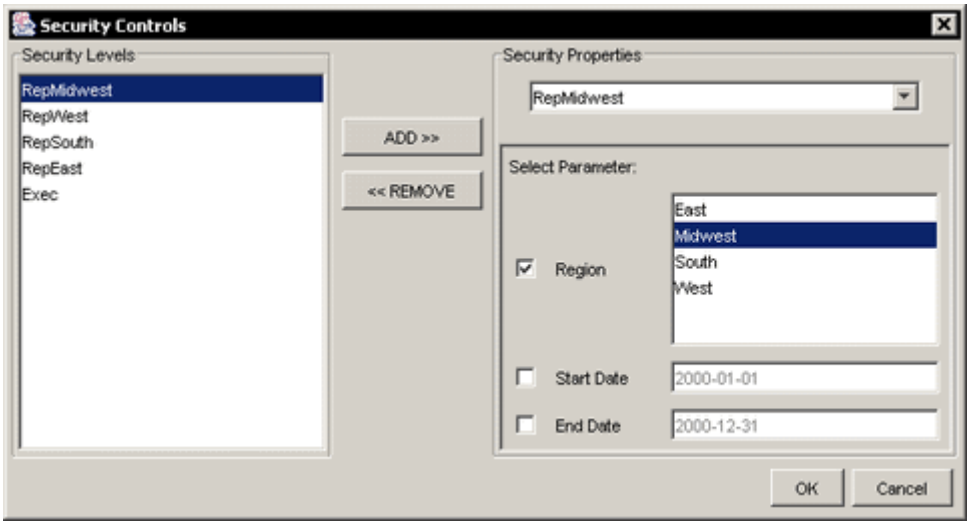


Report Cells with Security Settings

5.1.10.2.1. Security Parameters

In addition to controlling the specific behavior of cells in a report, you can also assign parameter values to a security level. For example, using the scenario described earlier where a sales report is grouped by region, you could add a query parameter that filters by sales region, and then assign the appropriate region value to a security level for that region. When an employee logs in, the report is run with the appropriate security level and the user will only see data for their sales region as the appropriate parameter value has been supplied to the report by the security level. To use security parameters, the report must use either a parameterized database query or a parameterized class as the data source.

To add security settings for parameters, select *Secured Query Parameters* from the *Data* menu. This will bring up a dialog allowing you to select query parameter values.



Security Parameters Dialog

Like the cell security dialog, the parameter security dialog contains a list of all the security levels defined in the Organizer on the left side.

To add security settings to report parameters, select the level you want to assign parameter values to and click the *Add* button. The security level will then appear in the drop-down list on the right side of the dialog. For each selected

security level, you can select which parameters you want to secure using the check box next to the parameter. For each secured parameter, you can select the value(s) that you want to be associated with this security level. Once you finish assigning parameter values, click the *Ok* button to return to the design window.

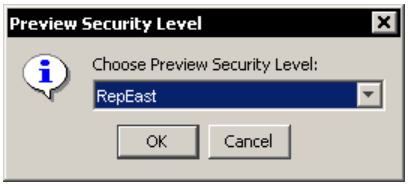
5.1.10.3. Security for Sub-Reports & Drill-Down Layers

Template security settings can be automatically applied to sub-reports as well as drill-down layers. In both cases, you need to edit the sub-report or drill-down layer (by clicking the sub-report tab, or navigating to the drill-down layer) to add security. In the sub-report or drill-down layer, you can add security settings just as you would for the main report. If you want to apply security to both the main report and the sub-report/drill-down layer, you will need to assign the security levels with the same name(s) as in your main report. When a report is run with a security level, the security level will be automatically applied to the sub-report(s) and/or drill-down layer(s).

For more information about drill-down and sub-reports, please see Section 5.1.8 - Drill-Down and Section 5.1.9 - Sub-Reports.

5.1.10.4. Viewing Secured Templates

You can preview the effects of your security settings in the Preview window of the Report Designer. To do this, select *Set Preview Security Level* from the *Data* menu. This will bring up a dialog prompting you to select the security level that you want to use to view the report.



Preview Security Level Dialog

Select the level you want and then preview the report. You will see the report with security settings applied for the level you selected. This level will remain selected until you change the settings. Selecting the blank entry at the top of the list in this dialog will set the report to run without applying any security settings.

5.1.10.4.1. Secured Report Designer

Only the administrator has access to the report security features. When other users access the Report Designer, it will run in secure mode. When the Report Designer is running in secure mode, users will only be able to preview reports at the level assigned to them. In addition, they will be unable to modify report cells with defined security settings or cell scripts that have been applied as security settings.

5.2. Chart Designer

5.2.1. Introduction to Chart Designer

The Chart Designer is a graphical user interface, launched within the Organizer that allows users to create and customize charts. The simple drop and drag style interface and extensive editing/formatting capabilities, makes chart design quick and easy.

5.2.1.1. Working with Charts

It is important to note that in ERES there are two ways that you can use charts. Charts can either be placed within reports or designed and run as stand-alone entities. A unique feature of charts placed in reports is that they can use the report data as their data source. In this instance, the chart cannot be deployed outside of the report.

Charts that use their own data sources can be deployed independently from reports, either in menu pages or using image URLs to embed them in web pages.

5.2.1.2. Starting Chart Designer

The Chart Designer interface is always loaded from within the Organizer. To start a new chart, you can click the *Chart Designer* button on the toolbar or select View → Chart Designer. You can also start Chart Designer when creating or editing data sources in the registry. The *View* button allows you to preview a data source and to build

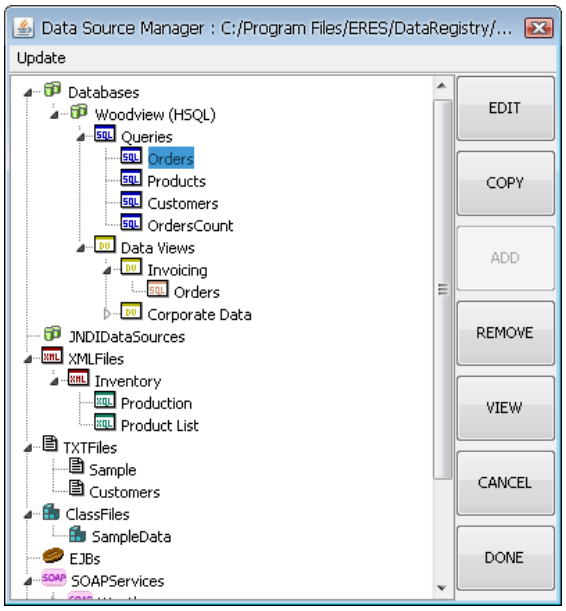
a chart or report. In addition, you can open the Chart Designer to edit a chart template file in the Organizer. To do so, first, select the file that you would like to open and then you can select File → Open File, right click on the file and select *Open File* from the pop-up menu, or double-click on the file.

If you are creating or adding an embedded chart within a report, the Chart Designer is launched from within Report Designer by either clicking on the *Insert Chart* button on the toolbar or by selecting Insert → Insert Chart.

5.2.1.3. Selecting a Data Source

The first step in designing a chart is to select the data source from which the chart is to be drawn. The first time you select to start a new chart you will be prompted to select the data registry that you would like to use. If there are not currently any available registries (meaning that you have not created any or you do not have privileges to view any), you will be prompted to go the Data Registry Manager to create one. For more on data sources, please see Section 4.1 - Data in Organizer

Once you have selected a registry, a Data Source Manager window will open allowing you to select, add, or modify a data source that you would like to use for the chart or report. Note that the registry will not open if you are building a chart from the registry in the *Modify Data Sources* dialogs.



Data Source Manager Window for Chart Designer

If you are launching Chart Designer from within Report Designer (using the *Insert Chart* option), the following dialog will be presented:

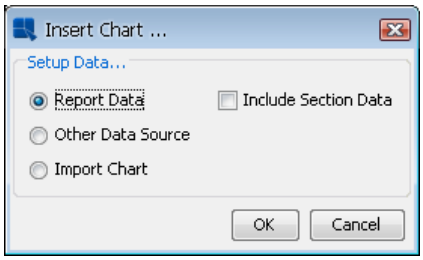


Chart Data Options

The following options are available:

Report Data: This will plot the chart using the data contained in the report. This data includes all of the report columns and computed columns that have been added using the Report functions. When using report data for the chart, you also have the option to include section data. This will take any formulas or labels from the report section in which

the chart is placed and make them available to the chart. Data for labels and formulas is imported as additional columns to the chart input data. The cell ID is used as the column name.

However, when you use report data, the chart cannot be deployed independently from the report, and certain charting features like time-based zooming and histograms are not available.

Other Data Source: This will allow you to select a different data source than the report data. If you select this option, you will go to the Data Source Manager window where you can create or select a data source for the chart.

When you use an independent data source the chart can later be deployed independently from the report. Also, all of the charting features are available. Note that some features like drill-down are only functional when the chart is deployed independently from the report.

Import Chart: Choose this option to import an existing chart to the report.

5.2.1.3.1. Chart Parameter Linking

If you are launching Chart Designer from within Report Designer (embedded chart) and you select to use an independent data source that is a parameterized query or class file, by default it will run with default values when the chart is run in the report. However, you can use the chart parameter linking feature for the chart to retrieve its parameter values based on column values in the report. This feature works in the same manner as the sub-report linking feature described in Section 5.1.9.4 - Linked Sub-Reports. To link a parameterized chart to a column, first, insert a parameterized chart into a report. Then, select *Chart Parameter Mapping* from the *Data menu* in Report Designer. This will bring up a dialog allowing you to map column fields from the report to query parameters in the chart.

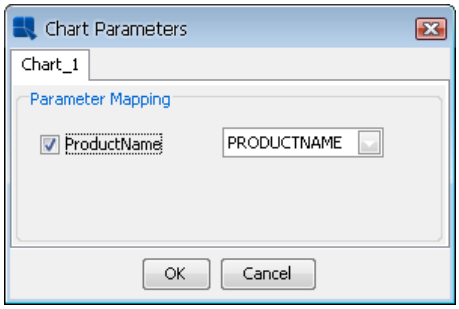


Chart Parameter Mapping Dialog

The dialog contains a tab for each parameterized chart. Each parameter in the chart can be mapped to a different column field in the report.

5.2.1.3.2. Transposing Data

If you preview a data source in the Data Source Manager (see Section 4.1.1.1 - Using Data Source Manager to learn more about the Data Source Manager), you can transpose the data source. This is useful when you want to see what the transposition will do to your data source before creating a chart. There is a better way to configure data transposition for charts (see Section 5.2.3.1.1 - Data Transposition); however, if you enable the transposition in the data source preview and choose to create a chart, the transposition settings will **not** be lost and the chart will use the transposed data (and you will be able to modify the transposition later in the *Data Mapping* dialog).



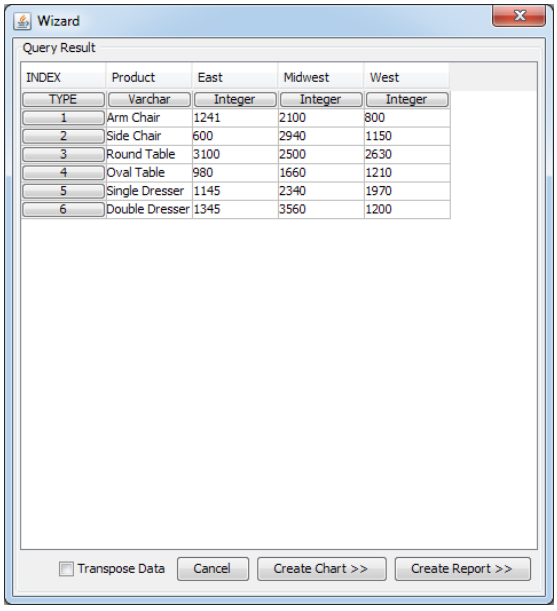
Note

This transposition method is available only in the data source preview dialog (that means that you have to open the Data Source Manager, select a data source and then click the *View* button). The other chart data transposition method is available in almost all charts (see Section 5.2.3.1.1 - Data Transposition for more details).

Once you preview your chart's data source, the next screen will present the first twenty records from the data source (With the exception of data views, which will require you to select fields and set conditions first). From this screen

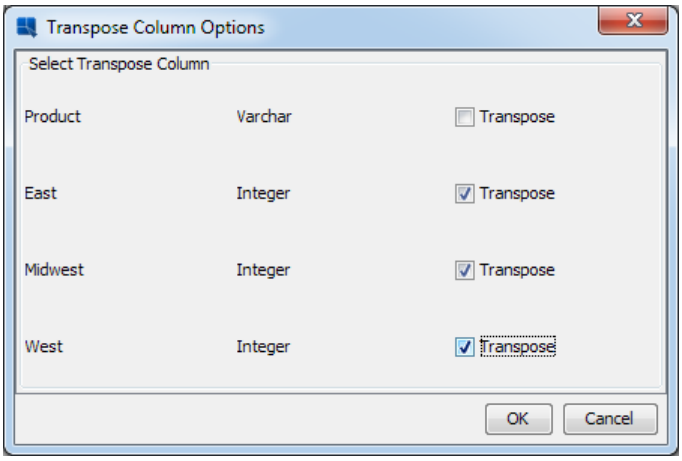
you can see all of the records returned by the selected source by checking the *Show All Records* box. Please note that the *Show All Records* button is hidden when the data source returns less than 20 rows of data.

Often, if you Are drawing data in spreadsheet (pivot table) format or from a crosstab report, you may want to use the column headers from the data source either as the categories or as a data series. For example, the following table is shown when a crosstab report is used as the data source.



Data Table Dialog

In order to be able to plot the regional column headers (East, Midwest, & West) in the chart, you will first need to transpose the data. To transpose the data, click the checkbox at the bottom of the data table window. A new window will open prompting you to select the columns that you would like to transpose.



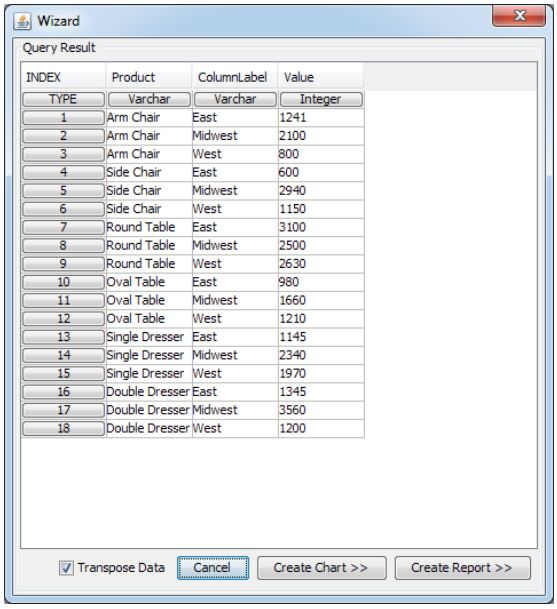
Select Transpose Dialog

Check the columns that you would like to transpose and click *Ok*. You will return to the data table where you will see the results of the data transposition. Columns that you have selected are removed from the data table and the transposed columns are placed at the end of the table as new columns.



Note

Columns that you select must have the same data type.



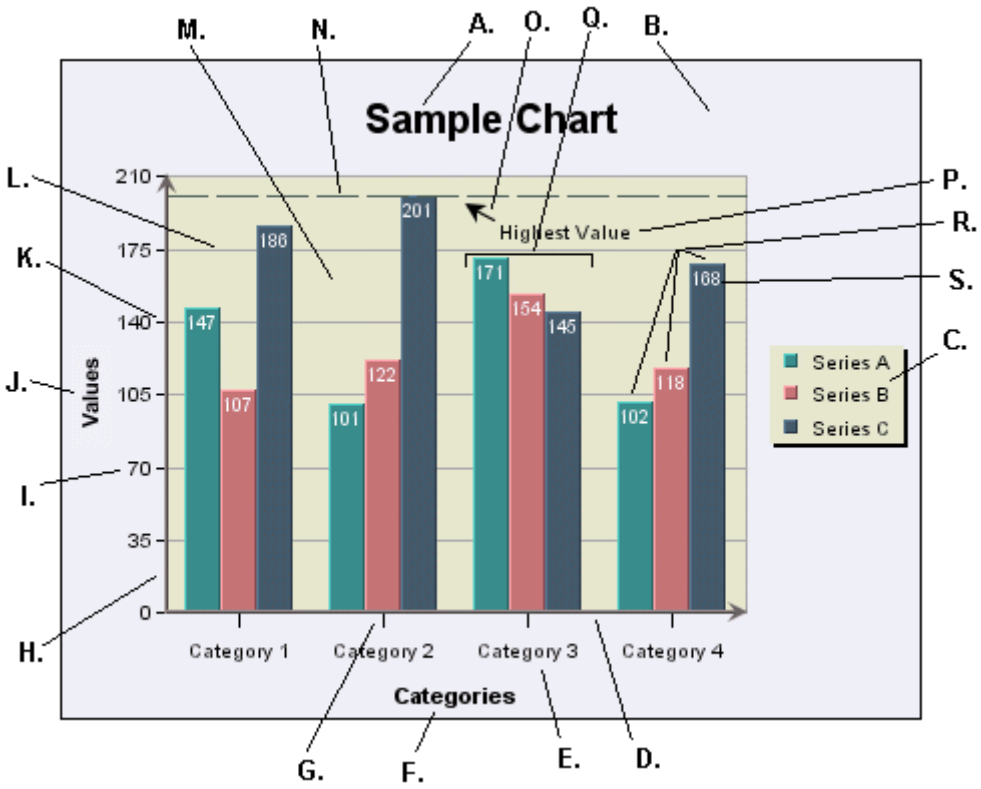
Data After Transposition

5.2.2. Charting Basics

This section covers some of the basics of charts including the different parts of a chart and the way in which tabular data either from the report or another data source is mapped to a chart.

5.2.2.1. What is a Chart

The following diagram illustrates the various components that make up a chart. Also, this diagram refers to various terms that will be used throughout this guide.



A. Main Title	This is the main title for the chart.
B. Chart Canvas	This is the background on which the chart is drawn. The canvas serves as the size/boundary for the chart and it is generally the same size as any exported image. You can modify the canvas color or place/add an image file as a background.
C. Legend	This is the chart legend. The legend shows your category or series names along with a color point. Secondary values, as well as added trend/control lines and control areas, can also be displayed in the legend.
D. Category (X) Axis	This is the X or category axis of the chart. Generally, the category axis plots the distinct entries from the dataset for which you want to plot values in the chart. (The values are generally plotted on the Y-axis.) Certain chart types such as bar and Gantt reverse this by drawing the categories on the Y-axis, and plotting the values on the X-axis. Other chart types like scatter and bubble plots, plot values on each axis to create a point in 2D or 3D space.
E. X-Axis Labels	These are the labels for the X-axis elements or categories.
F. X-Axis Title	This is the X-axis title.
G. X Ticker	These are the X-axis tickers. By default the tickers match each data point in the chart.
H. Value (Y) Axis	This is the Y or value axis of the chart. Generally, the Y-axis plots the values for each of the categories. By default the scale of the Y axis is generated to provide a best fit for the dataset; however, it can be manually adjusted. For combination charts, the second value axis is drawn at the right-hand side of the plot.
I. Y-Axis Labels	These are the labels for the Y-axis values.
J. Y-Axis Title	This is the Y-axis title.
K. Y-Axis Ticker	These are the Y-axis tickers.
L. Y Grid	These are grid lines drawn along each scale step in the Y-axis. Grid lines can also be drawn for the points on the X-axis (and Z-axis for 3D charts).
M. Plot Area	This is the area, bounded by the axes, where all the data points are plotted. You can fill the area with color and/or draw a border around the area, along with other options. The plot area can be moved and resized on the chart canvas.
N. Control Line	This is one of the special types of lines that can be added to a chart. In this instance, it is a control line that follows the highest value in a series. Control lines can also be drawn for averages and multiples of standard deviation. Users can also add a variety of trend lines to charts.
O. Floating Line	A floating line is an arbitrary line added to a chart. In this case it is being used as a pointer with an arrow. Floating lines move in relative position with the chart plot. They can also be used to create filled shapes.
P. Annotation Text	This is a piece of arbitrary text added to a chart (not labels or titles). You can place text anywhere in the chart canvas. Like floating lines, annotation text moves in relative position to the chart plot.
Q. Category Elements	This is the plot for a category element. There are three points plotted for each category because this chart has a data series.
R. Data Series Elements	These are the individual points that make up a category. A series allows groups of data to be plotted on a single chart. For more about information about categories and series, see Section 5.2.2.2 - Basic Data Mapping

S. Data Top Labels	These are labels placed at each data point in the chart that display the exact value for each point.

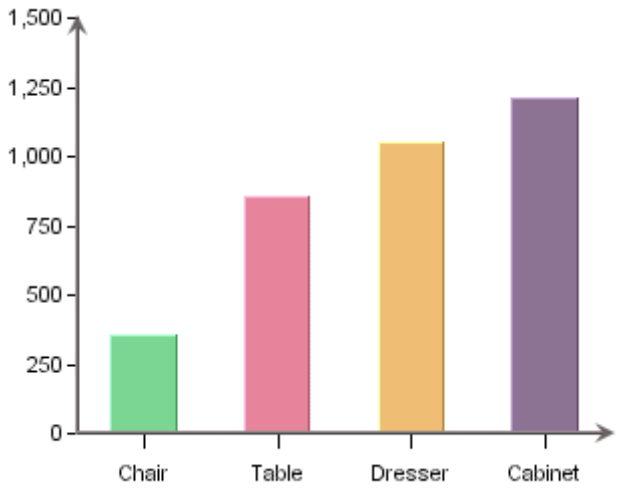
5.2.2.2. Basic Data Mapping

Data mapping is the way that raw data is rendered in the chart. Although data can be drawn from many sources, a chart looks for the basic structure of the data to be in the form of a table. Hence, data passed in as arguments, from the report or from XML files is converted to a table structure before mapping.

A basic set of data might look something like this:

Product	Sales
Chair	362
Table	862
Dresser	1052
Cabinet	1211

To plot this data in a chart, you would want to plot the **Sales** value for each entry in the **Product** column. Hence, the products are your category and the sales numbers are your values. In a chart you would map the **Product** column to the X (category) axis and the **Sales** column to the Y (value) axis. The resulting plot would look like this in a column chart:

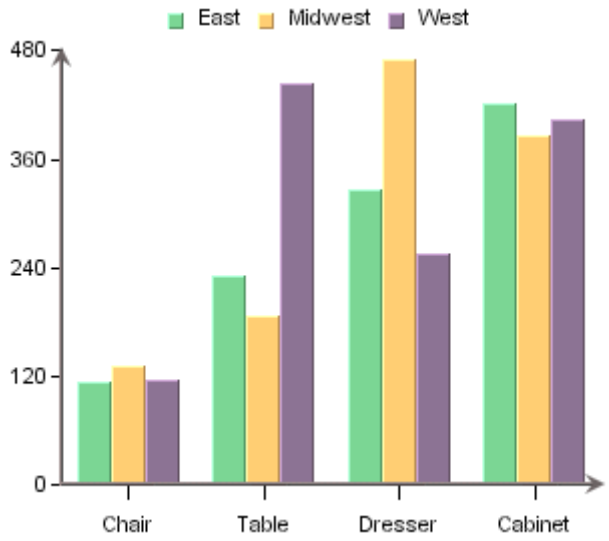


Here, a column is drawn to show the value for each distinct element in the category column. On top of the basic category values, additional information can be displayed in the form of a data series. For example, say that there is another element to our data that shows sales data not only for product, but also over a sales region. Our adjusted table would look like this:

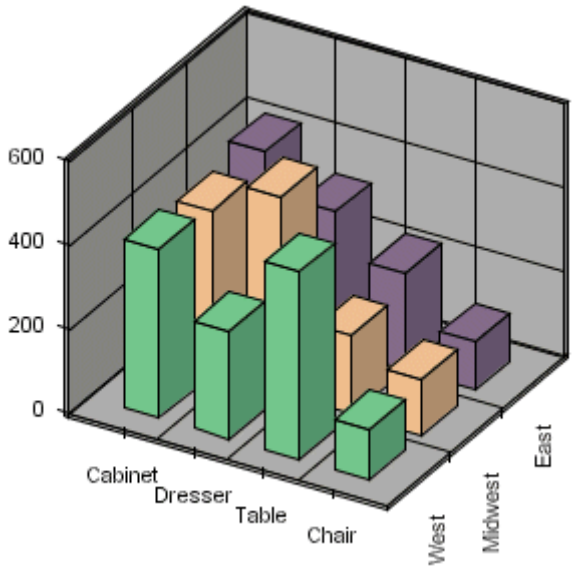
Product	Region	Sales
Chair	East	114
Chair	Midwest	131
Chair	West	117
Table	East	231
Table	Midwest	187
Table	West	444
Dresser	East	327

Product	Region	Sales
Dresser	Midwest	469
Dresser	West	256
Cabinet	East	422
Cabinet	Midwest	386
Cabinet	West	403

In order to show the value for each region per product we could add the **Region** column to the data mapping as a data series. Doing this gives us the following chart:



Now each category has three data points, one for each region. For two-dimensional charts the series is always displayed in-line. In three-dimensional charts, the series is drawn on the Z-axis by default, although it can be drawn in-line as well. Below, is the same chart show in 3D.



In this chart, the data series is drawn along the Z-axis. Note that the order of the categories has been changed to provide a better view of the data. This is the basic concept behind data mapping. Most chart types use this mapping technique or mapping options similar to this. Detailed data mapping instructions for each chart type are available in Section 5.2.3 - Chart Types and Data Mapping.

5.2.2.3. Saving and Exporting Charts

There are several options available both for saving chart definitions and exporting charts as image files. More information on how to save and export charts in the Chart Designer can be found in Section 5.2.6 - Saving Charts.

5.2.2.3.1. Saving Chart Definitions

There are two primary methods which you can use to store chart definitions created in ChartDesigner: either as chart or as template files.

Chart files Chart files save the chart in a binary file called `filename.cht`. A chart file stores both the definitions of the chart (type, dimension, etc.) as well as the data that was used to create the chart. Hence, chart files are portable. Any time you open a chart file, it will open with the original data that was used to create the chart. After opening the chart, you can refresh the data from the source or change the chart's data source entirely.

Template files Template files save the chart in a binary file called `filename.tpl`. A template file stores only the chart definitions and stores only 10 records of data with the chart. Hence, any time a template file is opened it will try to connect to the original data source to retrieve the data. Because of this, template files can be less portable than chart files. Template files can also be used to pass chart attributes from one chart to another. To do this, you can apply a template to a chart. This will carry over many of the attributes of the template to the current chart. For more information about applying templates, please see Section 5.2.6.2.1 - Working with Templates.

PAC files PAC files save the chart the way that it is ready for deployment. A `.PAC` file takes the chart and all the supplementary files associated with it - background images, dynamic drill-downs, and parametric drill-downs and places them in a single binary file. This configuration makes it easy to deploy charts and move them between ERES installations. If you reopen a `.PAC` file in Chart Designer, it will expand all the files back to their original location. Therefore, the target directories will need to exist for a `.PAC` file to be opened successfully. Refrain from renaming the `.PAC` file after its creation because this will lead to problems during expansion.

5.2.2.3.2. Generating Image Files

For image files, ERES Cloud can render charts in the following formats:

GIF ERES can generate GIF images using one of two compression methods - RLE or LZW. The LZW method is faster and produces smaller files; however, its use is protected by patent. You must obtain a license from Unisys in order to unlock the LZW compression. By default, GIF files are generated using RLE compression.

JPEG JPEG is another popular image format. It is a higher resolution image format than GIF and it is not patent protected. When generating a JPEG file you can specify the quality and compression of the file. The higher the quality, the larger the file.

PNG PNG is an image format that is less popular but it can be displayed in most browsers. It is a high-quality image with a smaller file size than JPEG. There are three different compression options available for this format.

SVG SVG (Scalable Vector Graphics) is a relatively new image format that saves the image as vectors in an XML-based text format. Generally, you will need a browser plug-in to view these images.

SWF SWF is an Adobe Flash file. The flash format is vector based and it allows the chart to be resized after export. Also, flash allows for high-resolution printing and produces a small file size.

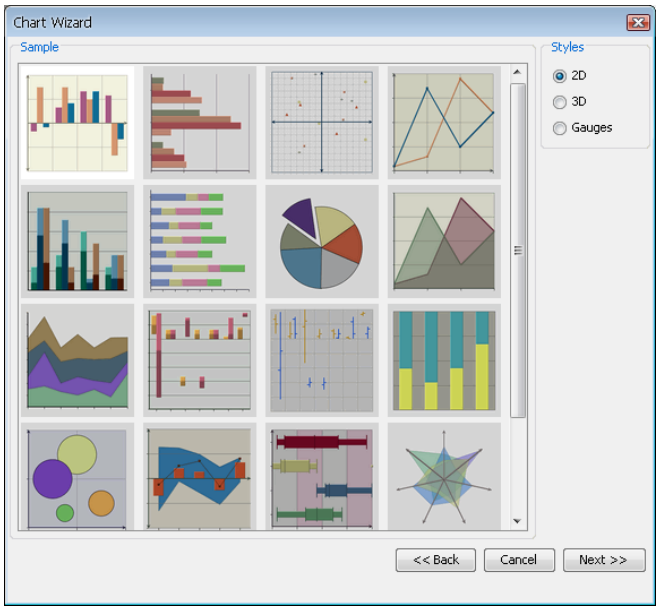
BMP This is a Windows bitmap format.

WMF WMF is the Windows Meta File format. This can be used for import/export into MS Office documents.

In addition to static images, ChartDesigner also allows the chart data to be exported as a text file, PDF, or an XML file.

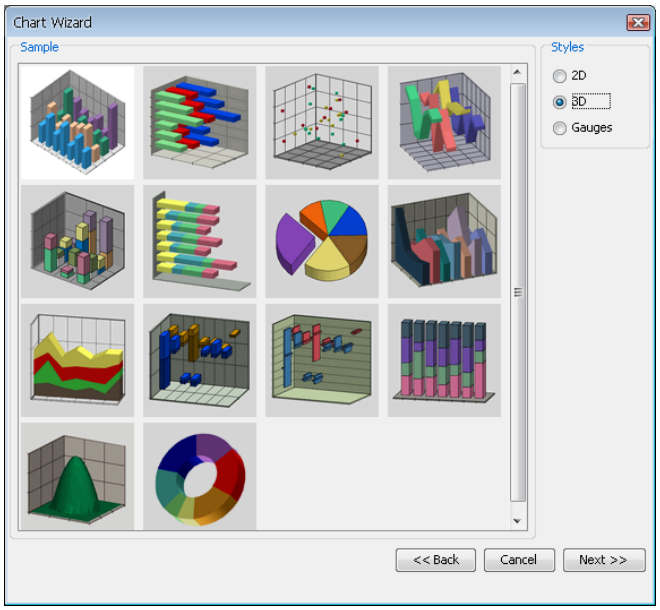
5.2.3. Chart Types and Data Mapping

Once you have selected the data that you would like to use for the chart, the next step in the Chart Wizard is to select the chart type that you would like to use. You will be presented with a dialog that allows you to specify the chart type.



Two-Dimensional Chart Types Selection Dialog

Each chart type represents a different way in which the data points are plotted to give a pertinent representation to all kinds of data. The different types of charts are broken down by dimension. In addition to basic chart types, users can create many different types of composite/combination charts by adding secondary values/series to the chart. You can toggle between the chart categories by selecting either *2D*, *3D* or *Gauges* in the right-hand side of the chart types dialog.



Three-Dimensional Chart Types Selection Dialog

This dialog enables you to select your chart type by either selecting the image and clicking *Next* or by double clicking on the chart image. You can select from one of the following chart types:

- Column

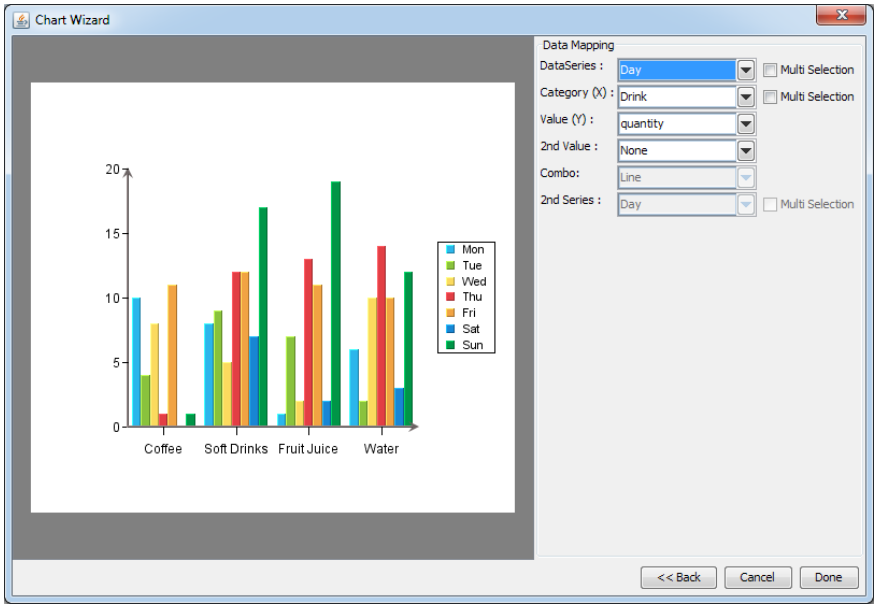
- Bar
- XY(Z) Scatter
- Line
- Stack Column
- Stack Bar
- Pie
- Area
- Stack Area
- High Low
- HLCO
- Percentage Column
- Doughnut
- Surface (Three-Dimensional Only)
- Bubble (Two-Dimensional Only)
- Overlay (Two-Dimensional Only)
- Box (Two-Dimensional Only)
- Radar (Two-Dimensional Only)
- Dial (Two-Dimensional Only)
- Gantt (Two-Dimensional Only)
- Polar (Two-Dimensional Only)
- Circular Gauge
- Square Gauge
- Semi Circular Gauge
- Square Gauge
- Quarter Circular Gauge
- Heatmap

Each of the chart types is described in detail later in this chapter, starting with Section 5.2.3.2 - Column Charts

For information on gauges, see Section 5.2.3.20.2 - Gauges

5.2.3.1. Data Mapping

After you have selected the chart type, the next step is to specify the data mapping for the chart. Data mapping is the way that the selected data source is rendered in the chart. The basics of data mapping are described in Section 5.2.2.2 - Basic Data Mapping. The data mapping screen in the Chart Wizard allows you to set the mapping options and also preview the results.



Data Mapping Dialog

The left-hand side of the dialog shows a preview of your chart. The right-hand side shows which columns from your data source have been selected to plot in the different chart elements (series, category, value, etc). By default, ERES will select the first available columns, based on the data type, to plot. Changing the data mapping is easy - click on the down-facing arrow next to a data field and select a different field. The chart preview in the left-hand part of the data mapping dialog will be updated immediately.

The specific mapping options vary for each chart type and are discussed later in this chapter starting with Section 5.2.3.2 - Column Charts.

5.2.3.1.1. Data Transposition

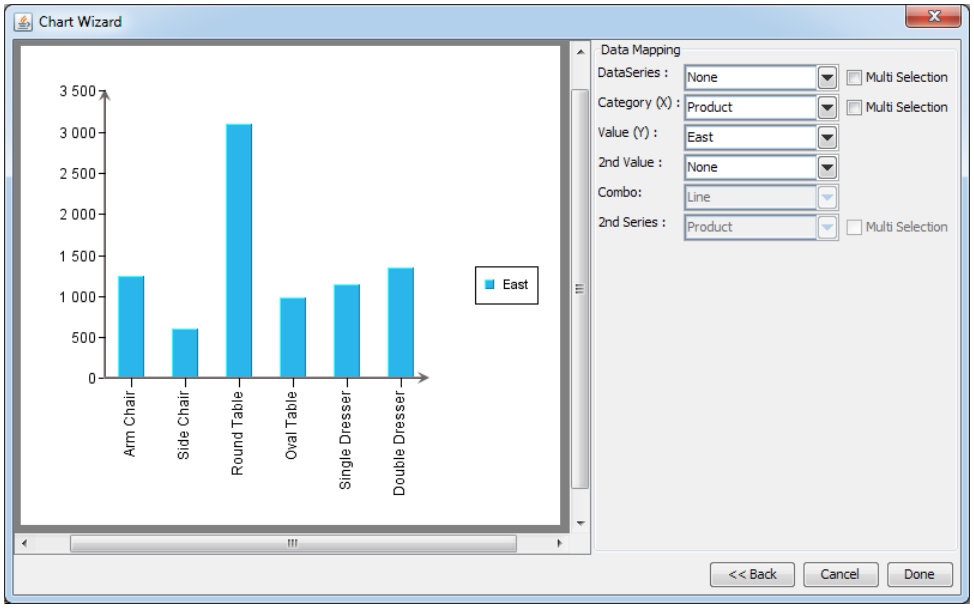
You can also set data transposing from the data mapping dialog. Transposing allows you to plot multiple data source columns in one chart without modifying the data source.

For example, we have a data source that looks like this:

INDEX	Product	East	Midwest	West
TYPE	Varchar	Integer	Integer	Integer
1	Arm Chair	1241	2100	800
2	Side Chair	600	2940	1150
3	Round Table	3100	2500	2630
4	Oval Table	980	1660	1210
5	Single Dresser	1145	2340	1970
6	Double Dresser	1345	3560	1200

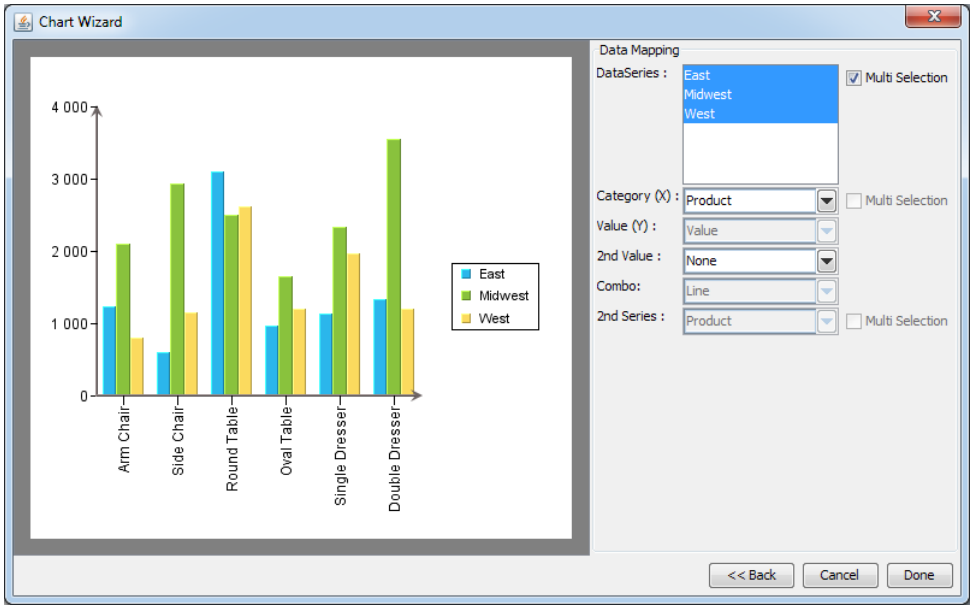
Example Data Source

We would like to plot data from all regions in one chart. The default mapping without data transposing does not allow us to do that. We would have to plot only one region in the chart or modify the data source, so we will have to transpose the data source.



Default Chart Data Mapping

To transpose the data, select the *Multi Selection* option next to the *Data Series* field. The field will change to a list allowing us to select several columns at the same time. We will just select all region columns and the chart will look like this:



Default Chart Data Mapping

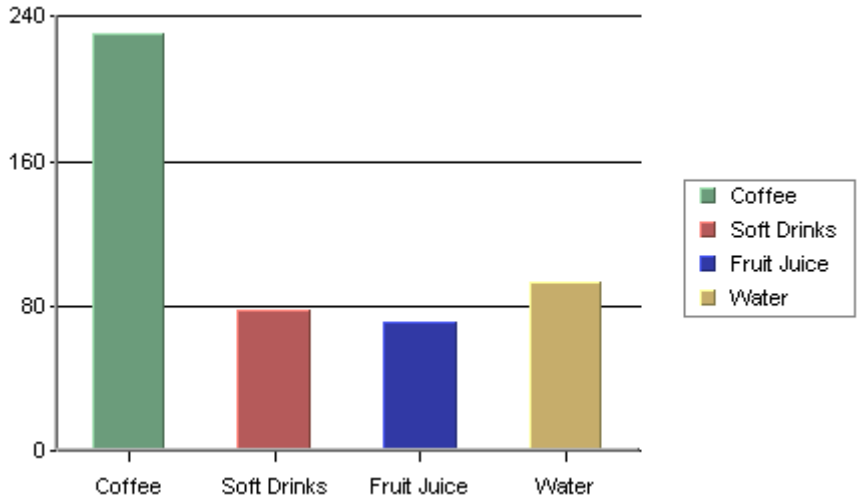


Note

Only one chart element (such as data series, category, value, or second series) can be transposed at a time. If you select the *Multi Selection* checkbox for one chart element, the *Multi Selection* checkboxes will be deactivated for other chart elements.

Transposing is not available for drill-down charts.

5.2.3.2. Column Charts

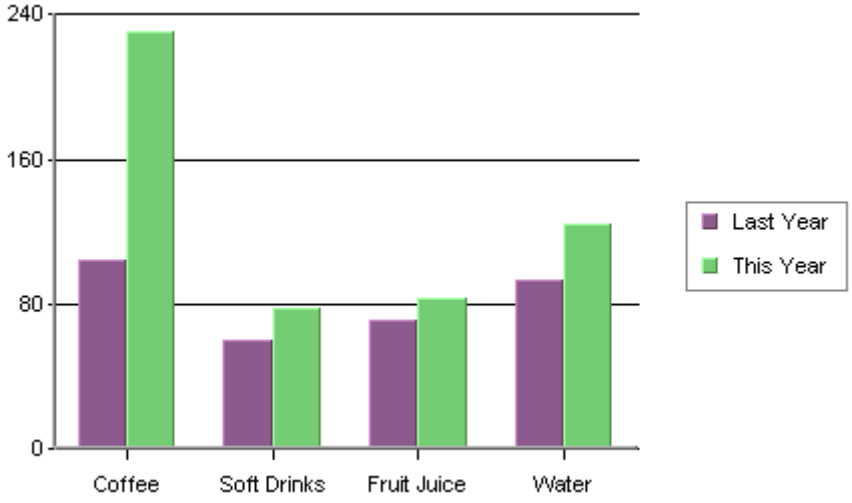


Column Chart (Without Data Series)

A column chart displays each row in the data table as a vertical bar (or column). The categories are listed along the X-axis and the values plotted along the Y-axis. Column charts are good for comparing discrete values for different groups. Each group is represented by a different color.

In a two-dimensional chart, if a data series has been selected then the entire series for a single category will be displayed in the XY plane. If a three-dimensional column chart is being used then all the vertical bars in a given data series are drawn using the same color along the Z-axis. If a data series is not present in the chart, the categories will be represented by the same color by default. Different colors for the categories are possible to set by clicking the  *Change chart options* icon on the toolbar (or selecting Format → Chart Options), and unchecking the *Single Color For All Categories* option.

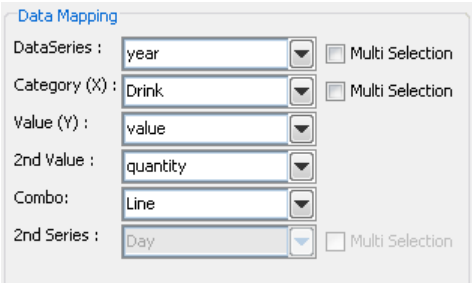
In the examples given, we have chosen Drink as the category variable and Value as the value variable (most examples use the `Sample.dat` file included in the ERES installation). Year has been selected as the data series variable for the chart below. Therefore, each name has two columns shown: one for *last year* and one for *this year* which are the only two values present in the data series column.



Column Chart with Data Series

5.2.3.2.1. Data Mapping

Data mapping for column charts is fairly straightforward. It is similar to the examples first presented in Section 5.2.2.2 - Basic Data Mapping. For column charts, the following options are available:



Mapping Options for Column Charts

Data Series: Choose a data column whose distinct values will determine the number of data series in the chart. Each element in a data series is drawn using the same set of colors and other attributes.

Category (X): Choose a data column whose distinct values will determine the categories.

Value (Y): Choose a data column to provide values for each category.

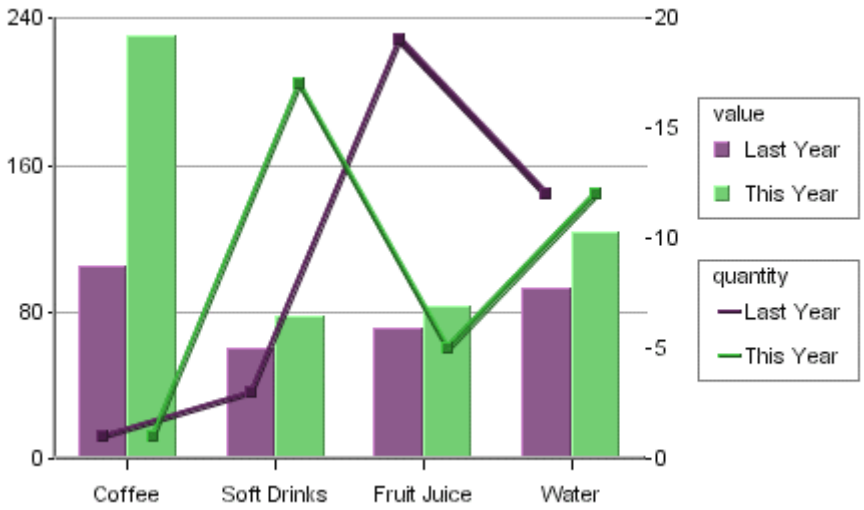
2nd value: Add a second value to create a combination chart.

2nd Series: Choose another column to be series for the secondary chart. This option is applicable only if the secondary chart is an overlay chart.

Combo: Choose the chart type for the secondary chart. For column charts the combo options are *Line* and *Overlay*.

The data mapping also allows you to transpose the data (in other words: to select several columns for a single category). To learn more about data transposition, please see Section 5.2.3.1.1 - Data Transposition.

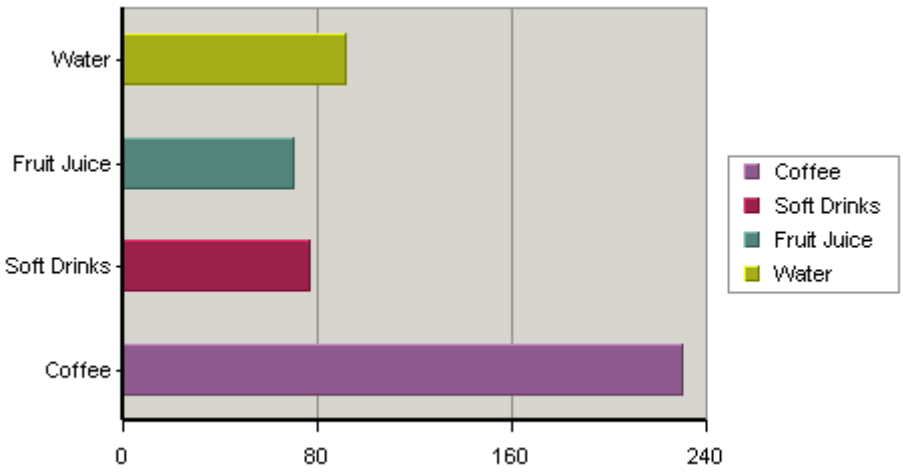
The last three data mapping options allow you to add a second value to the chart. ChartDesigner supports secondary values for all chart types except pie, radar, bubble, dial, surface, and scatter charts. In the example below, a second value of Quantity has been added to our column chart.



Column Chart With 2nd Value

As you can see, the second axis labels are drawn on the right-hand side of the plot area (you can choose to make this axis visible). Usually, the second value will share the same categories and series as the primary value. However, for two-dimensional column, stack column, stack area, high low, HLCO, and percentage column charts you can specify an overlay combination, which allows you to specify a second series. For more on overlay charts, please see Section 5.2.3.17 - Overlay Charts.

5.2.3.3. Bar Charts

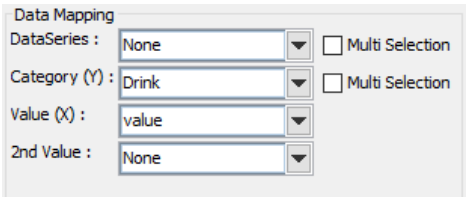


Bar Chart

A bar chart is essentially the same as a column chart, except that horizontal bars are drawn in the chart as opposed to the vertical bars which are used in a column chart. In a bar chart, the categories are plotted along the Y-axis and the values along the X-axis.

Just as for a column chart, if a data series has been selected then the entire series for a single category is displayed in the XY plane. If a three-dimensional bar chart is being used, all the horizontal bars in a data series are drawn using the same color. Each category is drawn using a different color. If the data series is not present in the chart, the categories will be represented by the same color by default. Different colors for the categories are possible to set by clicking the  *Change chart options* icon on the toolbar (or selecting Format → Chart Options), and unchecking the *Single Color For All Categories* option.

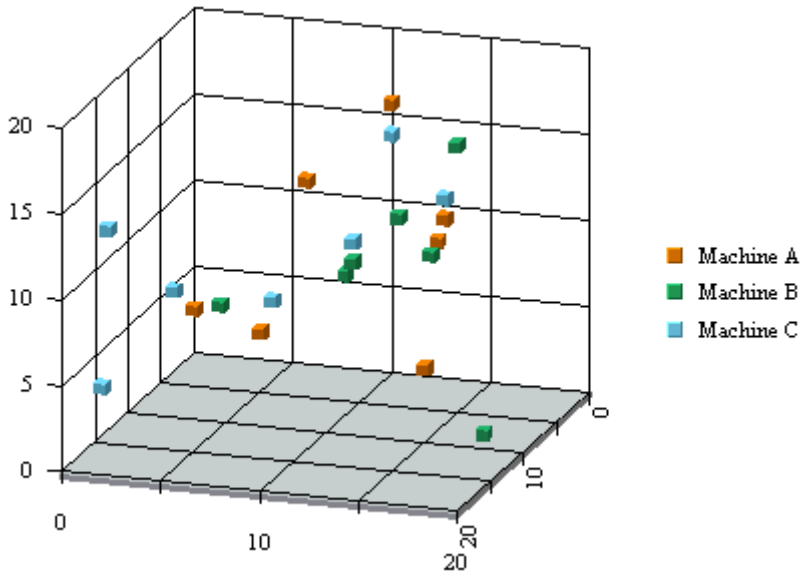
5.2.3.3.1. Data Mapping



Mapping Options for Bar Charts

The mapping for this chart type is similar to that of a column chart, except the category (X) and values (Y) under column chart become category (Y) and values (X) respectively. This is because values are represented vertically in a column chart, but horizontally in a bar chart. Please note that the *2nd Series* and *Combo* options are not available for bar charts. This is because the only combination available with bar charts is a line.

5.2.3.4. XY(Z) Scatter Charts

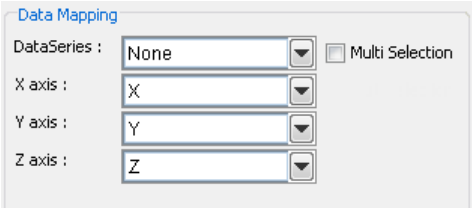


XY(Z) Scatter Chart

In an XY(Z) scatter chart, each selected row in the data table defines a point in two or three-dimensional space. Thus each column must contain either numeric or date/time values. A marker represents each point. The data columns that are in each row of the data table determine the spatial position of the marker.

Optionally, another data table column can be chosen to separate the markers into groups. Elements of each group have the same value on this column which is referred to as the data series column. Markers in the same group are drawn using the same drawing attributes; in other words, using the same shapes and colors. The X-axis scale of a scatter chart is linear. This means that unlike other chart types, the data points may or may not be evenly spaced along the X-axis.

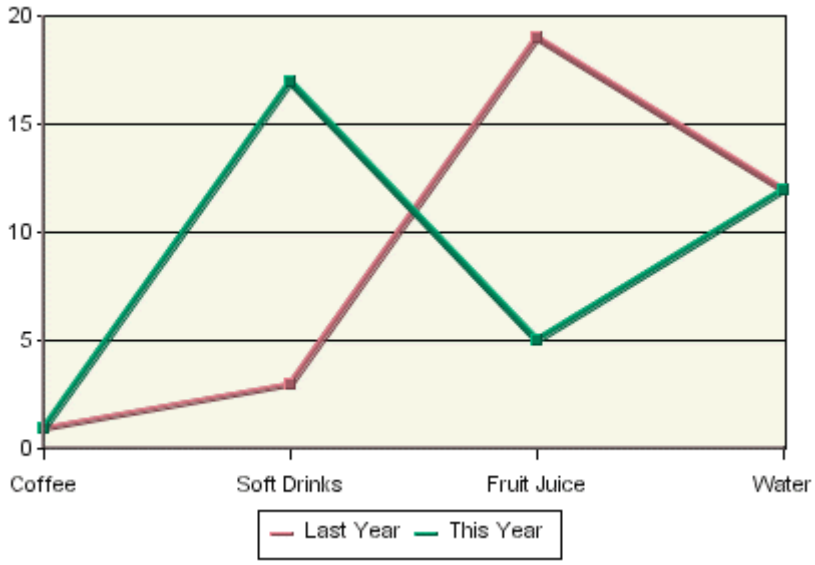
5.2.3.4.1. Data Mapping



Mapping Options for Scatter Charts

In a scatter chart, the X-axis, Y-axis, and Z-axis values determine the X, Y, and Z coordinates of a point respectively. The data series box allows you to choose a data column whose distinct values will determine the number of data series in the chart. Each element in a data series is drawn using the same set of drawing attributes(e.g., colors and markers).

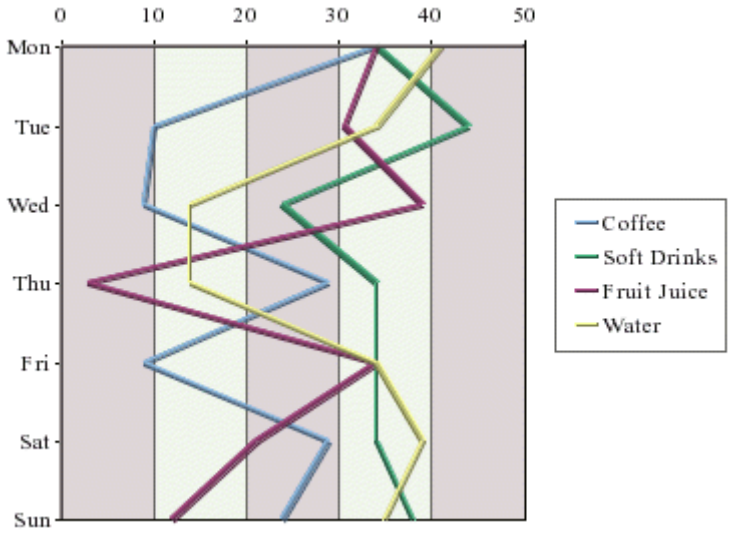
5.2.3.5. Line Charts



Line Chart

Line charts present data in a manner similar to column charts. For a two-dimensional line chart, a marker denotes a data point for each category. The value column, which must be numeric, determines the height of a marker. Each marker is joined with a line. If the chart has a data series, a separate line is drawn for each element in the series. Please note, the markers (points) are not shown by default. You can enable showing markers (points) in the *Line and Point* dialog. This dialog can be opened from the *Format* menu or by clicking the  *Format line and points* icon on the toolbar.

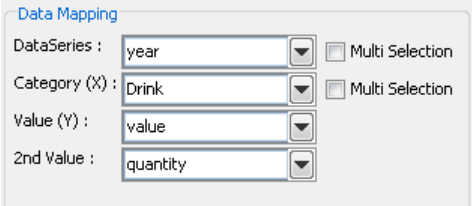
ChartDesigner allows 2D line charts to be displayed vertically in addition to the default horizontal setting. To use this feature in Chart Designer, create a line chart and then select Format → Chart Options. Next, select *Vertical*. The chart is rotated clockwise 90 degrees.



Vertical Line Chart

A three-dimensional line chart is an extension of its two-dimensional counterpart. It contains no additional information. The markers disappear (they are not available for 3D line charts) and thicker lines, spaced apart on the Z-axis, replace the thin lines.

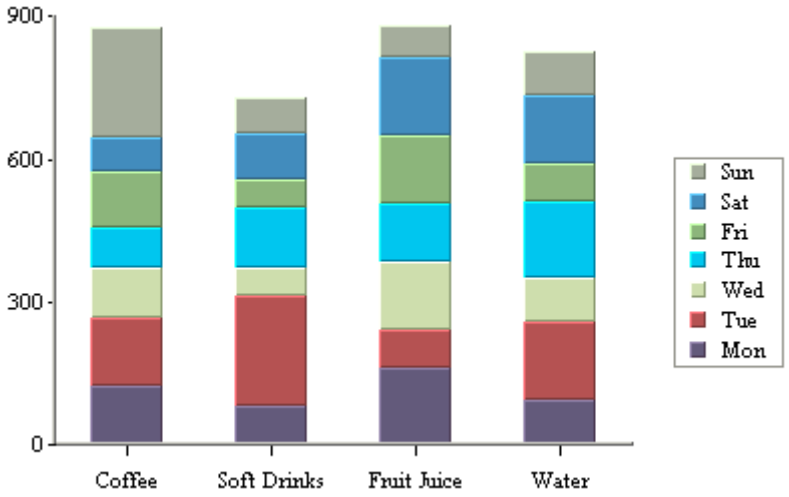
5.2.3.5.1. Data Mapping



Mapping Options for Line Charts

Data mapping for line charts is almost exactly the same as for column charts (discussed in Section 5.2.3.2.1 - Data Mapping). Line charts, however, do not have the *2nd Series* and *Combo* options. This is because the only combination available with line charts is a line. To create line combinations with other chart types (bar, column, stack column, etc) select the other chart type as the primary chart, and select *Line* as the *Combo* option.

5.2.3.6. Stack Column Charts



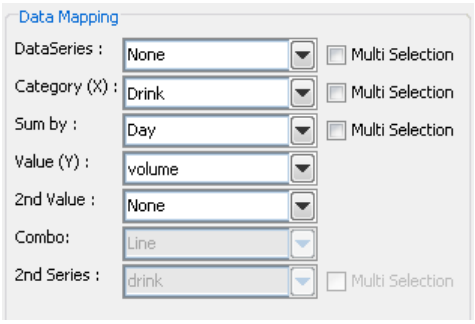
Stack Column Chart

A stack column chart is a derivative of a column chart in which each vertical column comprises a stack of bars. Each selected row in the data table is represented by a component of a chart column. For a two-dimensional stack column chart, the categories are plotted on the X-axis and the values on the Y-axis. The value column, which must be numeric, determines the length of each bar component. A third column, known as the sum-by column, represents a group of values for each category. A stack for a given category value is built up stacking the sum-by values for that category value. Each stack component of a chart column has a distinct sum-by value denoted by a distinct color. Each row in the data table must have a unique pair of values on both the category and sum-by columns. Components with negative values are separated from components with positive values and they are stacked up in the “negative” direction below the X-axis.

In our example above, we have chosen types of drink as category value and day representing the sum-by values. Another independent category (to be displayed on the Z-axis) may be optionally specified based on a fourth column as a data series. In such a case, each row in the data table must have a unique triplet of values on the category, sum-by and data series columns. In a two-dimensional chart, the entire data series for a single category is displayed side-by-side in the XY plane. In a three-dimensional stack column chart, the data series is displayed along the Z-axis.

5.2.3.6.1. Data Mapping

In stack column charts, each category is further subdivided into components. Hence, there is an additional sum-by option used to determine the sum-by values for the chart.



Mapping Options for Stack Column Charts

The data mapping options are as follows:

Data Series: Allows you to choose a data column whose distinct values will determine the number of data series in the chart. Each element in a data series is drawn using the same color.

Category (X): Allows you to choose a data column whose distinct values determine the categories. Values from this data column are used to calibrate the X-axis. Each category in a data series is drawn as a distinct column in the chart.

Sum-by: Allows you to choose a data column whose distinct values determine the components in each category. Each distinct value in this column determines a distinct stack component of a column.

Value (Y): Choose a data column to provide values for each category.

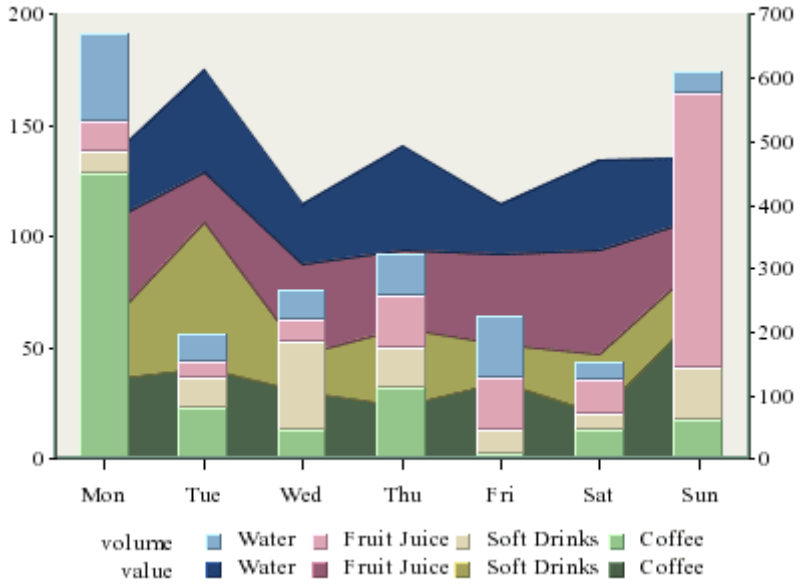
2nd value: Add a second value to create a combination chart.

2nd Series: Choose another column to be series for the secondary chart. This option is applicable only if the secondary chart is an overlay chart.

Combo: Choose the chart type for the secondary chart. For stack column charts the combo options are *Line*, *Stack Area*, and *Overlay*.

The data mapping also allows you to transpose the data (in other words: to select several columns for a single category). To learn more about data transposition, please see Section 5.2.3.1.1 - Data Transposition.

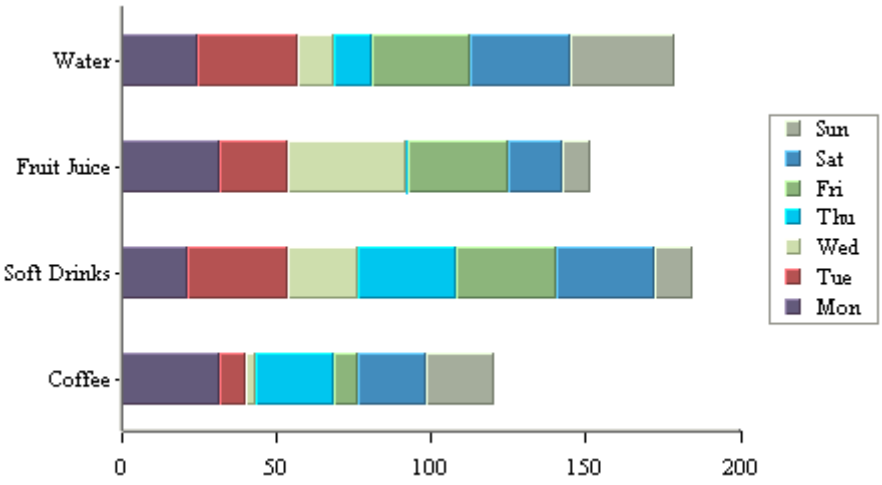
Stack column charts have a unique combination option which allows users to plot a second value as a stack area chart. Both values share the same category and sum-by values. This chart can compare component based categories for different values in the data.



Stack Column-Stack Area Combination Chart

With this combination chart you can specify to hide particular sum-by component in the chart to achieve a strip area effect.

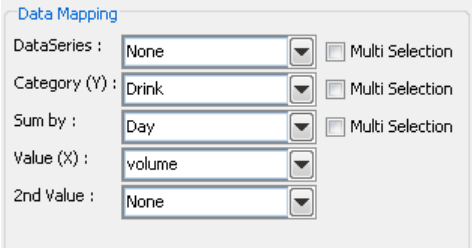
5.2.3.7. Stack Bar Charts



Stack Bar Chart

Like a stack column chart, the stack bar chart displays the chart categories as a sum of different stacked values - the sum-by column in the data source. The difference for a stack bar chart is that the categories are plotted on the Y-axis and the values are plotted along the X-axis.

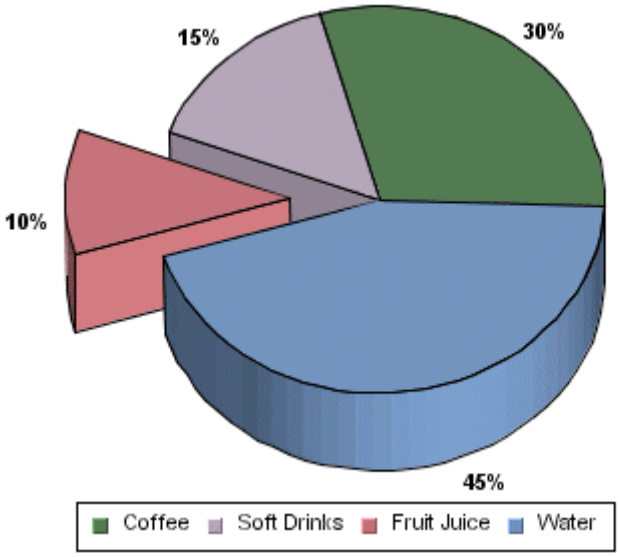
5.2.3.7.1. Data Mapping



Data Mapping Options for Stack Bar Charts

The mapping for this chart type is similar to that of a stack column chart, except the category (X) and values (Y) under column chart become category (Y) and values (X), respectively. This is because values are represented vertically in a column chart, but horizontally in a bar chart. Please note that the *2nd Series* and *Combo* options are not available for stack bar charts. This is because the only combination available with stack bar charts is a line.

5.2.3.8. Pie Charts

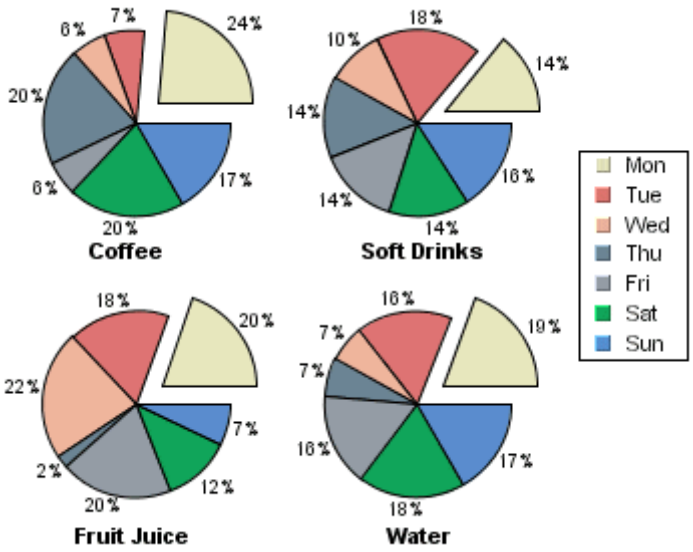


Pie Chart

In a pie chart, each row in the data is represented as a pie sector. The pie chart requires a category column and a value column. The value column must be numeric. The number of distinct values in the category column determines the number of pie sectors in the chart. All the values are summed up and each slice is assigned an angle according to its share in the total. Thus, the value column for each category determines the size of the slice representing that category.

A three-dimensional pie chart is simply an extension of its two-dimensional counterpart and contains no extra information.

Pie charts can also have a data series. If you specify a series in the data mapping, a separate pie will be drawn for each category element and it will be comprised of the series elements. A pie chart with a data series is displayed below:



Pie Chart With Data Series

When a series is present, all the separate pies are drawn on a single plot area and resized in proportion with the plot. You can choose to either stack the different pies or draw them in a line.

5.2.3.8.1. Data Mapping



Mapping Options for Pie Charts

For pie charts, the mapping is as follows:

Data Series: Allows you to choose a data column whose distinct values will determine the number of data series in the chart. Each element in a data series is drawn using the same color.

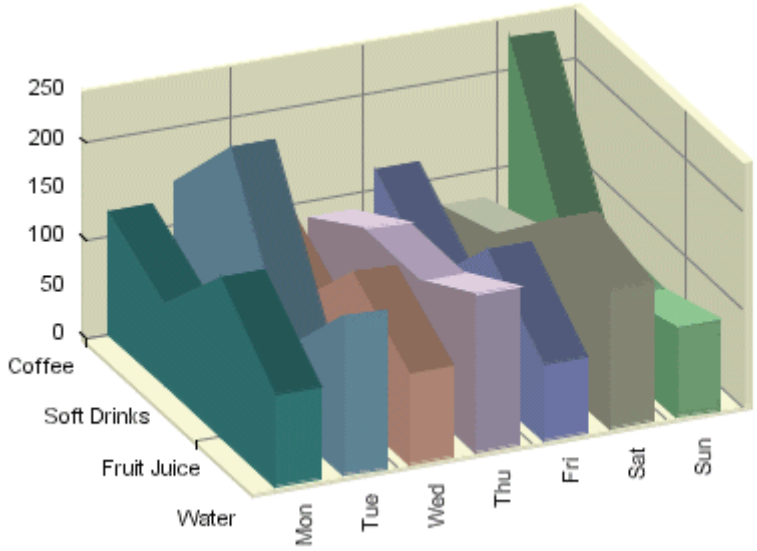
Category (X): Allows you to choose a data column whose distinct values determine the various categories.

Value (Y): Allows you to choose a data column to provide values for each category.

There is no *2nd Value* option, as you cannot make any combination charts with a pie chart.

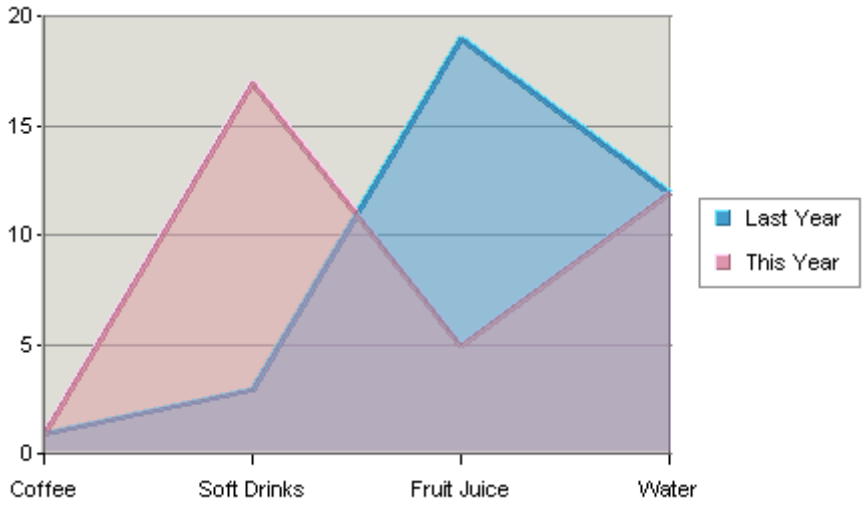
The data mapping also allows you to transpose the data (in other words: to select several columns for a single category). To learn more about data transposition, please see Section 5.2.3.1.1 - Data Transposition.

5.2.3.9. Area Charts



Three-Dimensional Area Chart

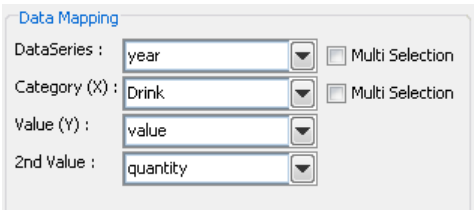
A three-dimensional area chart may be viewed as a derivative of a column chart. A three-dimensional area chart may be constructed from a three-dimensional column chart in the following manner. For each data series (Z-axis) value, the tops of all the columns are joined together by a thick line. The columns are then removed and the area (in the XY plane) under each line is filled with a distinct color.



Two-Dimensional Area Chart

A two-dimensional area chart is essentially a projection of its three-dimensional counterpart viewed along the Z-axis. As a result, a two-dimensional area chart may sometimes hide a data series value altogether. Therefore, it must be used with caution. The chart above illustrates this point: in a two-dimensional display some parts of each series are concealed by a series in front. To ameliorate this problem, you can change the ordering of the data series (see Section 5.2.4.8.3 - Data Ordering) or set the area translucent (see Section 5.2.4.1.3 - Format Menu) as in the example above.

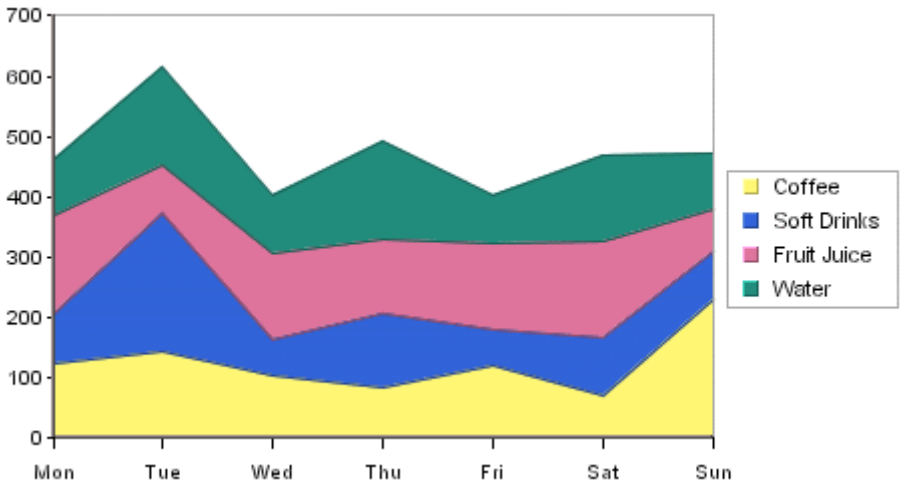
5.2.3.9.1. Data Mapping



Mapping Options for Area Charts

Data mapping for area charts is almost exactly the same as for column charts (discussed in Section 5.2.3.2.1 - Data Mapping). However, area charts do not have the *2nd Series* and *Combo* options. This is because the only combination available with area charts is a line.

5.2.3.10. Stack Area Charts



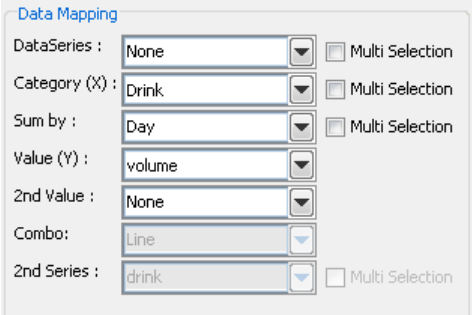
Stack Area Chart

A two-dimensional stack area chart may be viewed as a derivative of a two-dimensional stack column chart but without a data series. This chart may be constructed from a stack column chart as follows: The tops of each stack component that have the same color are joined together by lines. The stack columns are removed and the area between each set of lines is filled with a distinct color.

A three-dimensional stack area chart can have a data series and can also be derived from a three-dimensional stack column chart using the process stated above. In this case, the stack columns for each distinct data series value (a fixed Z-axis value) are joined separately, giving rise to multiple stacks.

Both two-dimensional and three-dimensional stack area charts may be constructed using a similar set of data as needed for their respective stack column chart counterparts. Note, as stated earlier, that a two-dimensional stack area chart cannot have a data series.

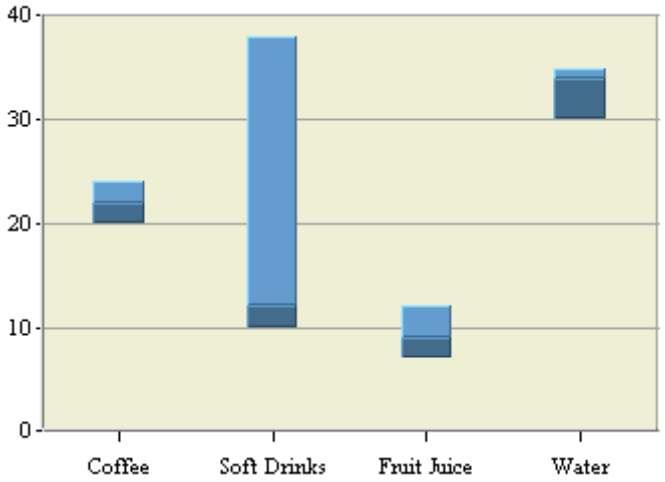
5.2.3.10.1. Data Mapping



Mapping Options for Stack Area Charts

Data Mapping for stack area charts is almost exactly the same as for stack column charts (covered in Section 5.2.3.6.1 - Data Mapping). However, two-dimensional stack charts cannot have a data series, and the combo options are *Line* and *Overlay*.

5.2.3.11. High-Low Charts

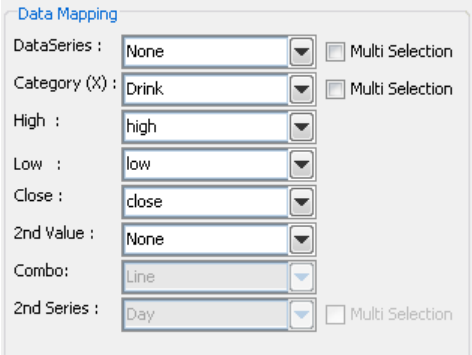


High-Low Chart

A high-low chart is also a derivative of a column chart, only that instead of one value column it uses two columns for high and low bounds of values. The data for the chart must contain at least three columns: category, high, and low. The category column can contain values of any type, while the high and low columns must be numeric.

Another option for a high-low chart is a data column called the *close* column. If a close column is also used, then the close value will lie somewhere between the high and low values. The portion of the bar between the low point and the lose point is rendered in a darker form of the same color as the portion between the close point and the high point. As for many other types of charts, a high-low chart may include a data series column.

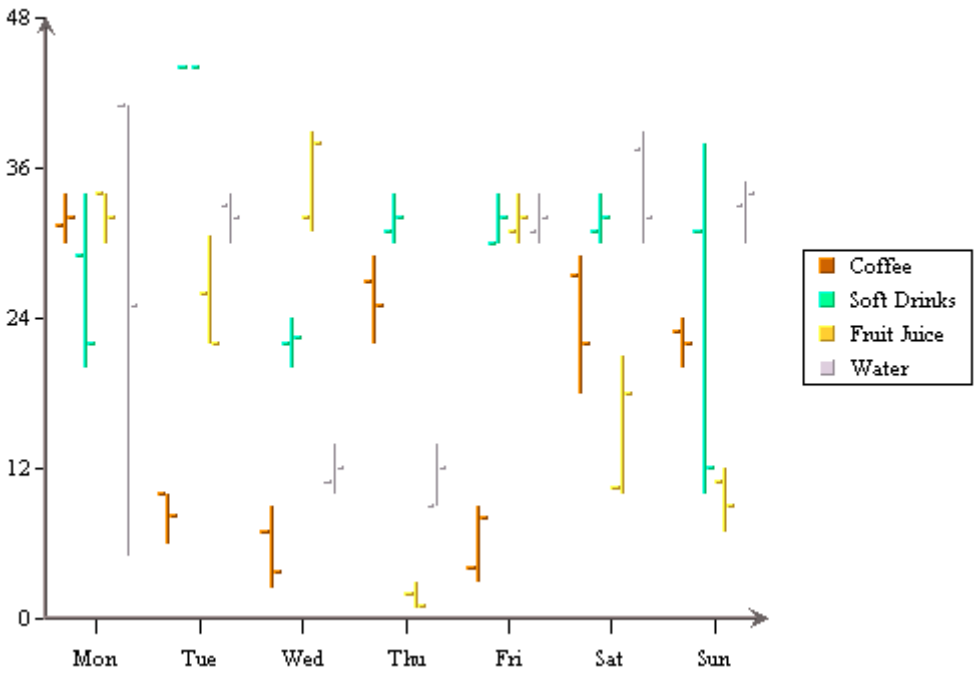
5.2.3.11.1. Data Mapping



Mapping Options for High-Low Charts

Data mapping for high-low charts is similar to column charts. You can select series and category columns in the same manner. The difference for high-low charts is that you need to specify at least two value columns - a high and a low value. These are the two points that are plotted for each category. A third value called *Close* can also be specified. Combo options for high-low charts are *Line*, *Column*, and *Overlay*.

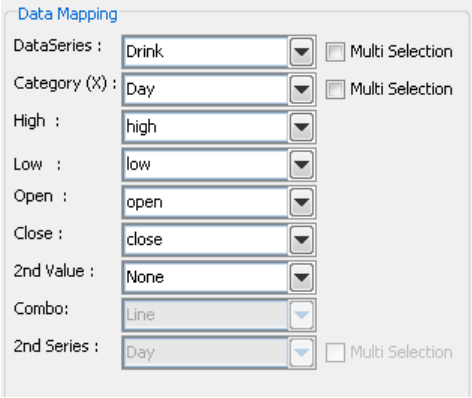
5.2.3.12. HLCO Charts



HLCO Chart

A HLCO (High Low Close Open) chart is similar to the high-low chart described above, except that it also contains an *open* column. It is useful in presenting data that fluctuates over discrete periods of time, such as stock prices or inter-day temperatures. The HLCO chart can also be displayed in a Candlestick format for both 2D and 3D charts. This feature is described in Section 5.2.4 - The Chart Designer Interface.

5.2.3.12.1. Data Mapping



Mapping Options for HLCO Charts

Data mapping for HLCO charts is similar to that for High-Low charts. The only difference is that instead of specifying two different value columns, you must specify four different columns - *High*, *Low*, *Open* and *Close*. Combo options are *Line*, *Column* and *Overlay*.

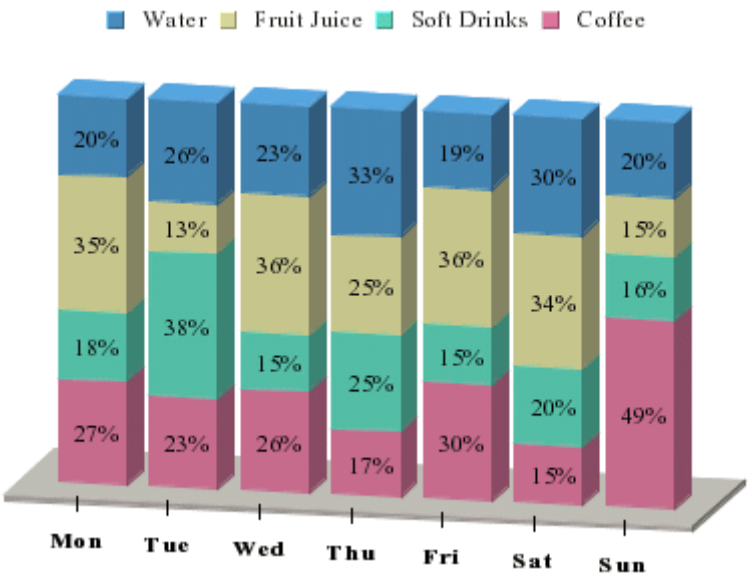
A common kind of chart is a one that plots both stock price and trading volume in the same chart. This can be accomplished using an HLCO-column combination in ChartDesigner.



HLCO-Column Combination Chart

This example plots stock price and volume data over a three month period.

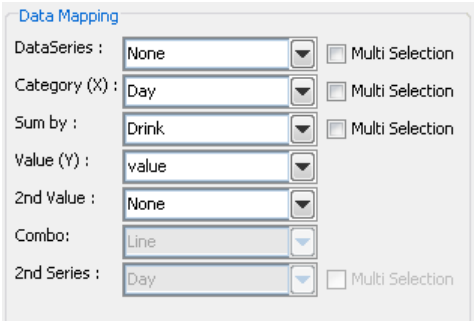
5.2.3.13. Percentage Column Charts



Percentage Column Chart

A percentage column chart may be viewed as a derivative of pie and column charts together. Each column in the chart corresponds to one pie. A percentage column chart has a category column corresponding to the X-axis value. A sum-by column represents the category in this case and the value column is the same as that of a pie chart.

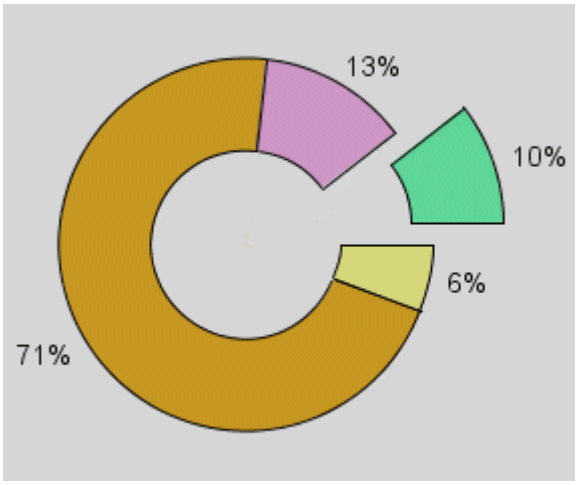
5.2.3.13.1. Data Mapping



Mapping Options for Percentage Column Charts

Data mapping for percentage column charts is the same as for stack column charts (covered in Section 5.2.3.6.1 - Data Mapping). The only difference is that the sum-by components are plotted as a percentage of the Value column instead of discrete values. Combo options for percentage column charts are *Line* and *Overlay*.

5.2.3.14. Doughnut Charts

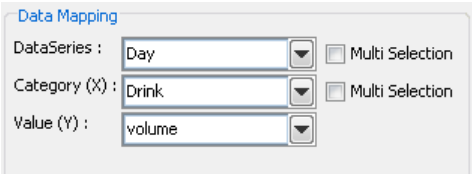


Doughnut Chart

A doughnut chart is the same as a pie chart, except for the “hole” in the center.

Just as for a pie chart, if a data series is selected, a separate doughnut will be drawn for each category element and each doughnut will be comprised of the series elements. When a series is present, all the separate doughnuts are drawn on a single plot area and resized in proportion with the plot. You have an option to either stack the different doughnuts or draw them in a line.

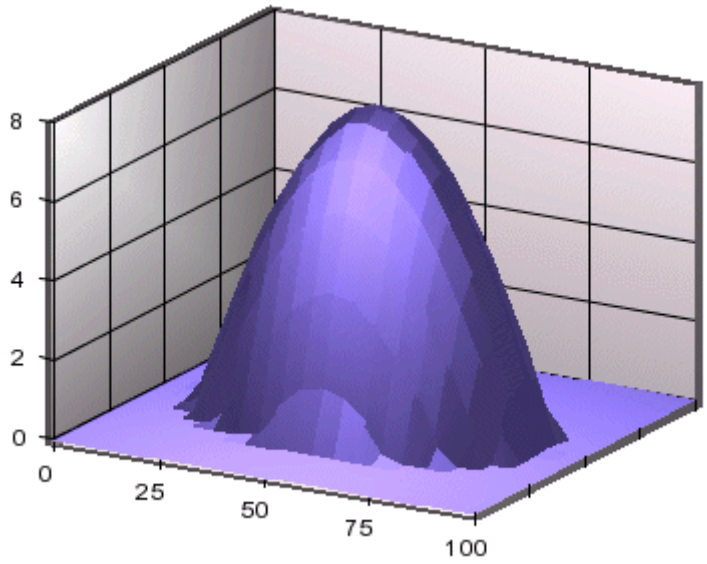
5.2.3.14.1. Data Mapping



Mapping Options for Doughnut Charts

Data mapping for doughnut charts is exactly the same as for pie charts (discussed in section Section 5.2.3.8.1 - Data Mapping).

5.2.3.15. Surface Charts



Surface Chart

A 3D surface chart is a scientific/business chart, which uses three sets of numeric data values to form a smooth surface in a three dimensional space. For each data point, two of the three values represent the coordinates on the horizontal plane and the third value is used to determine the vertical position of a data point. Input data can be in a form of a rectangular matrix (as shown below) or arranged in three columns where each column represents an axis. In the following sample data, the top row (column header) and the left column (row header) represent the coordinate values of the axes that form the plane on which the surface will be drawn. When the chart is viewed from above the plane, you can see a grid formed by joining every four adjacent points together. A given set of data cannot plot a surface chart if this grid cannot be drawn.

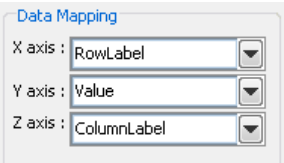
Sample surface chart data:

	0	20	40	60	70	80	100
20	0	0	0	0	0	0	0
30	0	10	10	10	10	10	10
40	0	10	25	25	25	2	10
80	0	10	25	30	30	30	25
90	0	10	25	30	30	30	25
100	0	10	10	25	25	25	10

The values of the column header and the row header are not required to be equally separated because a surface chart can properly scale the horizontal distances. This dataset can create a distorted inverted cone. A surface chart can be very impressive when animation is turned on.

A surface chart uses a set of X, Y, and Z coordinates to plot a 3D chart. The vertical axis is called Y-axis and each Y value is represented by a function of X and Z, f(X, Z). Data on XY plane (horizontal plane) simply looks like a square matrix. Surface charts do not support data series and cannot be converted to a 2D chart.

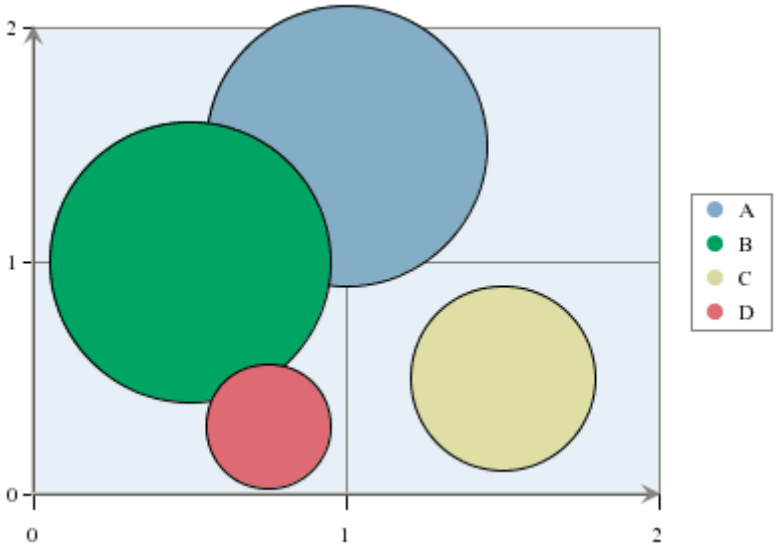
5.2.3.15.1. Data Mapping



Mapping Options for Surface Charts

The data mapping for a surface chart is similar to a three-dimensional scatter chart; the X-axis, Y-axis, and Z-axis values determine the X, Y, and Z coordinates of a point respectively. However, unlike a scatter chart, surface charts do not support data series and cannot be converted to a two-dimensional chart.

5.2.3.16. Bubble Charts



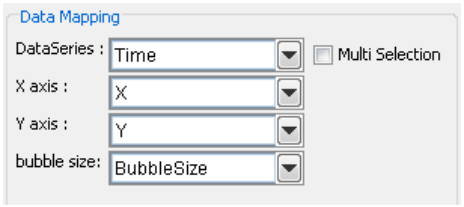
Bubble Chart

A bubble chart is used to represent data wherein the size of the bubble also provides information. A data point is represented by an XY coordinate and a third point, which is the radius of a circle or bubble.

This chart is available in a two-dimensional form only. For more information about the bubble charts and their display options, please see Section 5.2.4.9.1 - Bubble Charts.

5.2.3.16.1. Data Mapping

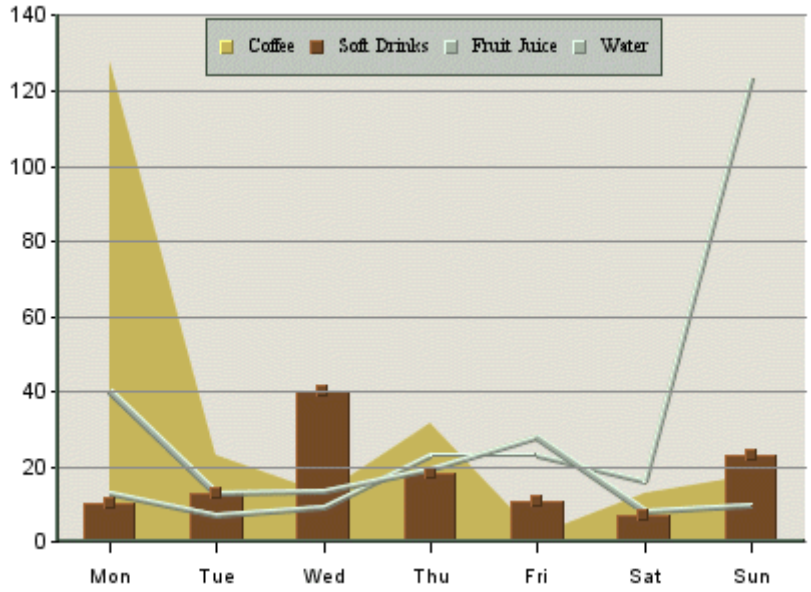
The data mapping for bubble charts is similar to three-dimensional scatter charts; however, instead of plotting a Z-axis position, the third value determines the size of the bubble (specifically the radius).



Mapping Options for Bubble Charts

Note that similarly to a scatter chart, the columns have to be numeric in order to be mapped successfully.

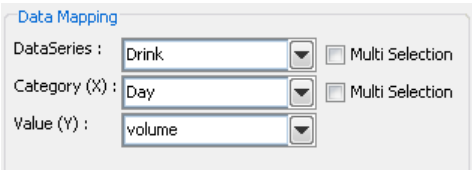
5.2.3.17. Overlay Charts



Overlay Chart

ChartDesigner supports a special chart type called overlay chart, which allows you to super-impose more than two charts with a common category axis. A different chart can be used to represent each element of a data series. This allows for more freedom while creating the chart and also allows more information to be represented. The chart types supported in an overlay chart are column, area, and line charts. Only two-dimensional charts can be included in an overlay chart.

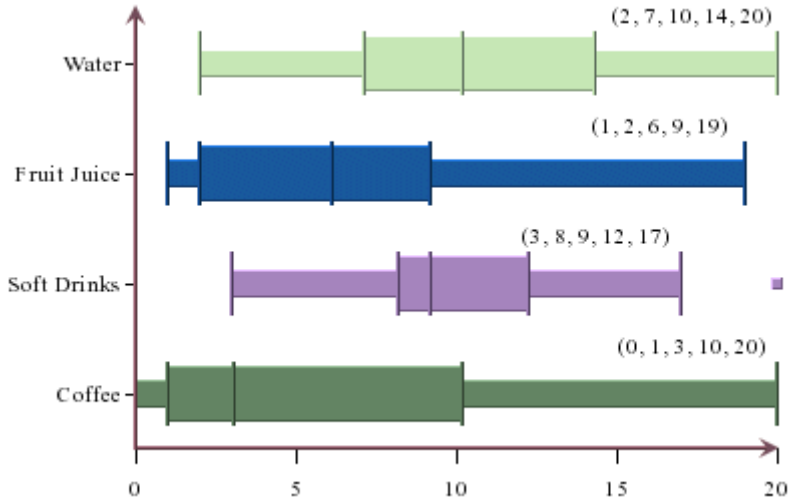
5.2.3.17.1. Data Mapping



Mapping Options for Overlay Charts

The data mapping for an overlay chart is similar to a column chart (covered in Section 5.2.3.2.1 - Data Mapping), except that the overlay chart plots each element in the data series as a separate chart. Please note that the *2nd Value*, *2nd Series* and *Combo* options do not apply to overlay charts. Overlay charts do not support secondary values. However, the different series elements can be plotted on different axes. For more about plotting options for overlay charts, see Section 5.2.4.9.3 - Overlay Charts.

5.2.3.18. Box Charts

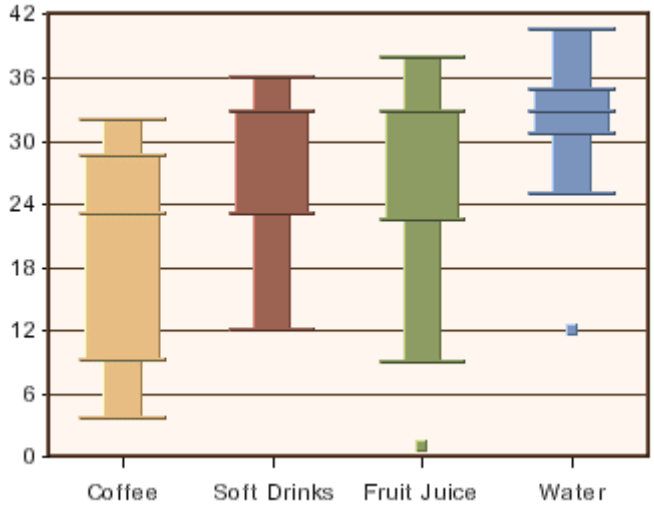


Box Chart

ChartDesigner provides a statistical chart type, the Box Chart, also known as Box and Whiskers Chart. Basically, the box chart provides a quick visual realization of summary statistics. A histogram shows the distributions of observed values so you can identify the peak(s), minimum, maximum values and clustering of data. A box chart, on the other hand, displays a summary of data distribution in quarterly percentiles (Minimum data point, 25th percentile, median, 75th percentile, and Maximum data point).

The middle line in the box represents the median. The other lines in each direction represent the 25th percentile increment (decrement). The minimum and maximum points are the observed minimum and maximum which are not outliers. The lines that extend from the edge of the box to the minimum/maximum are sometimes called whiskers. Hence, the name box and whisker chart. Outliers are any values that are 1.5 times (or more) the length of the box, that being the difference between the 75th and 25th percentiles. Outliers are calculated and shown as dots away from the box. The other nice feature about box plot is that you can stack up the boxes in one chart to compare summaries of different data sets.

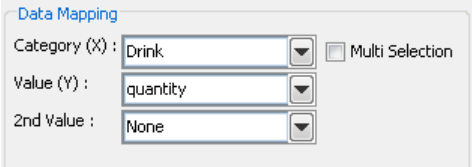
ChartDesigner allows box charts to be displayed vertically in addition to the default horizontal setting. To use this feature in Chart Designer, create a box chart and then select Format → Chart Options, and select *Vertical*.



Vertical Box Chart

This chart is available in two-dimensional form only.

5.2.3.18.1. Data Mapping



Mapping Options for Box Charts

For box charts the mapping is as follows:

Category (X): Allows you to choose a data column whose distinct values determine the various categories.

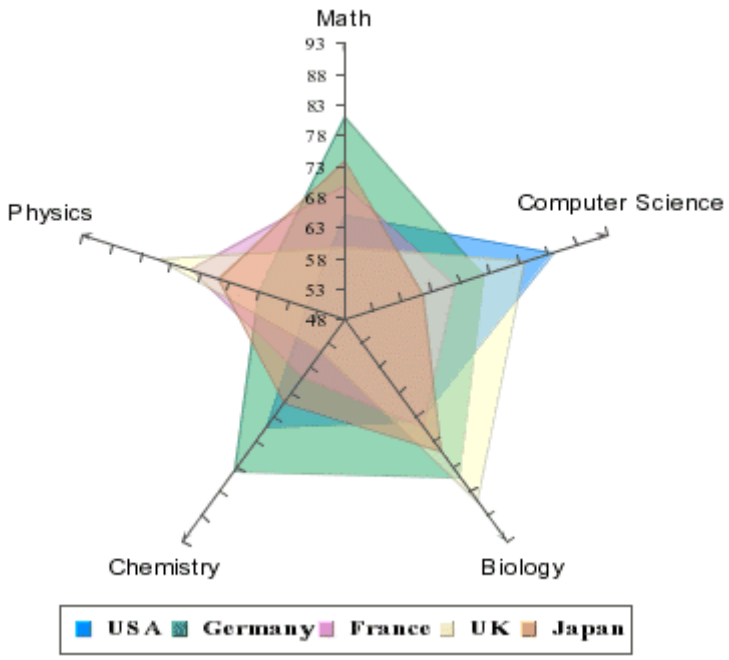
Value (Y): Allows you to choose a data column to provide values for each category. These values are used to calculate the minimum data point, the 25th percentile, the median, the 75th percentile, and the maximum data point.

2nd value: Add a second value to create a combination chart.

Please note that the *2nd Series* and *Combo* options are not available for box charts. This is because the only combination available with box charts is a line.

The data mapping also allows you to transpose the data (in other words: to select several columns for a single category). To learn more about data transposition, please see Section 5.2.3.1.1 - Data Transposition.

5.2.3.19. Radar Charts



Radar Chart

Radar charts are useful when you need to compare performance/measurement results, statistics, etc from different sources. For example, we can compare median test scores of children in the industrial nations on subjects such as math, computer science, and biology. A radar chart has multiple axes along which data can be plotted. Each axis is a category. The data is shown as points on the axis. The points belonging to one data series can be either joined or the enclosed area filled. A point close to the center on any axis indicates a low value while a point near the end represents a high value. In some cases, low values may be more desirable than high values, e.g. price/earning, price/sales, price/book of stocks.

This chart is available in a two-dimensional form only.

5.2.3.19.1. Data Mapping



Mapping Options for Radar Charts

For radar charts the mapping is as follows:

Data Series: Allows you to choose a data column whose distinct values will determine the number of data series in the chart. Each element in a data series is drawn using the same color along each axis.

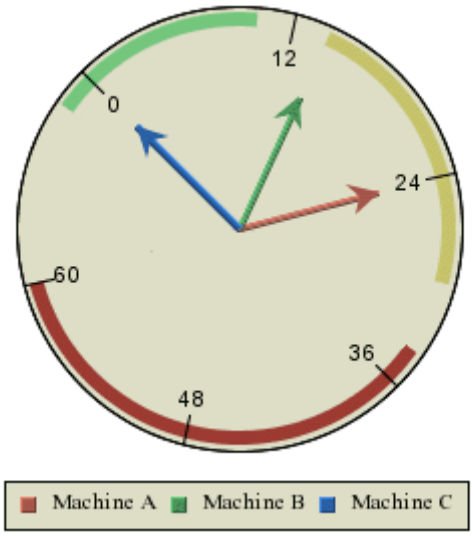
Category (X): Allows you to choose a data column whose distinct values determine the categories. The number of unique elements in this column determines the number of axes in the radar chart.

Value (Y): Allows you to choose a data column to provide the numeric value to plot against the category.

Please note that the *2nd Value*, *2nd Series*, and *Combo* options do not apply for radar charts. Radar charts do not support secondary values.

The data mapping also allows you to transpose the data (in other words: to select several columns for a single category). To learn more about data transposition, please see Section 5.2.3.1.1 - Data Transposition.

5.2.3.20. Dial Charts

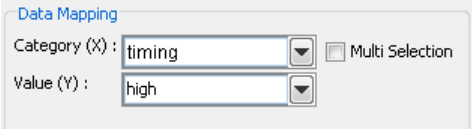


Dial Chart

Dial Charts are used to build circular charts (such as temperature gauges, speedometers, and clocks), create dashboards, or balanced scorecard applications. Here the data is represented as a “hand” on a “dial”. You can either use the whole dial (for instance for a clock) or just part of it (a semi- circular chart such as a gasoline gauge). Dial Charts support pop-up labels, drill-down, and user interactions such as rotating the hands or sectors of a dial.

This chart type is available in a two-dimensional form only.

5.2.3.20.1. Data Mapping

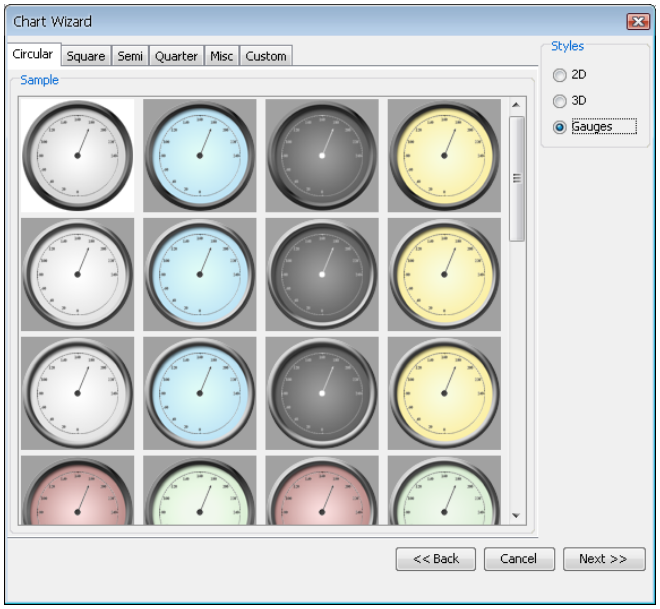


Mapping Options for Dial Charts

The data mapping of a dial chart is similar to a pie chart (detailed in Section 5.2.3.8.1 - Data Mapping). However, instead of pie wedges, the categories become dial hands. Also, dial charts do not support data series or secondary values.

5.2.3.20.2. Gauges

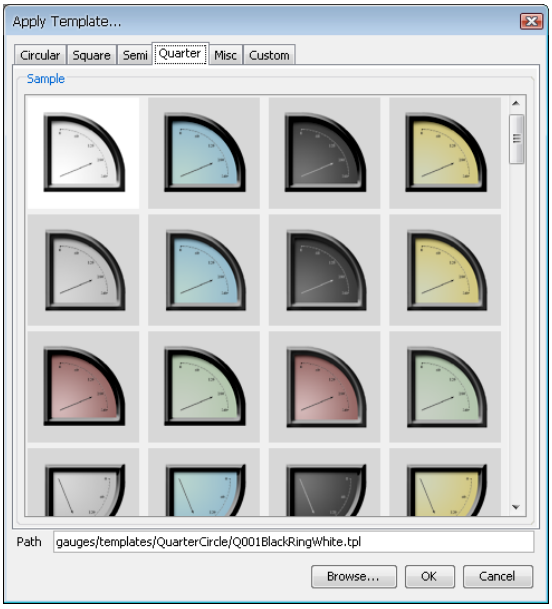
Gauges are special types of dial charts containing a plot background image and/or plot foreground image. To create a new gauge, select the *Gauges* radio button on the *Chart Type Selection* dialog.



Selecting a Gauge Template

There are several tabs of predefined gauge templates to help you quickly create a gauge.

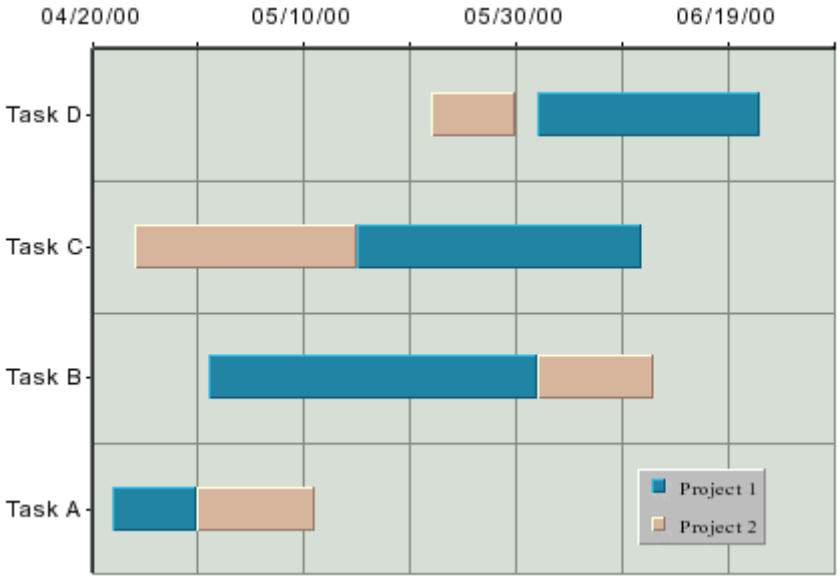
It is also possible to apply a gauge template onto an existing dial chart. Selecting apply template when the current chart is a dial chart will display the gauge tabs just like when creating a new chart.



Applying a Gauge Template

For more information regarding the dial background and foreground images, see Section 5.2.4.9.2 - Dial Charts.

5.2.3.21. Gantt Charts



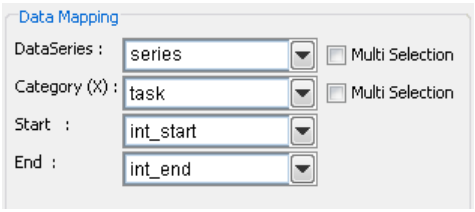
Gantt Chart

A Gantt or time chart is used to represent data where a time factor is being measured. This is often used in project or time management situations. A Gantt chart resembles a bar chart with the category axis on the Y-axis and the value axis on the X-axis. The value axis uses date, time, or timestamp data (numeric values can also be used). Each data point has a start and end time associated with it.

This chart type is available in a two-dimensional form only.

5.2.3.21.1. Data Mapping

The data mapping for Gantt charts is similar to high-low charts. However, instead of selecting a column for the high and low values, you select the start and end values.



Mapping Options for Gantt Charts

The mapping is as follows:

Data Series: Allows you to choose a data column whose distinct values will determine the number of data series in the chart.

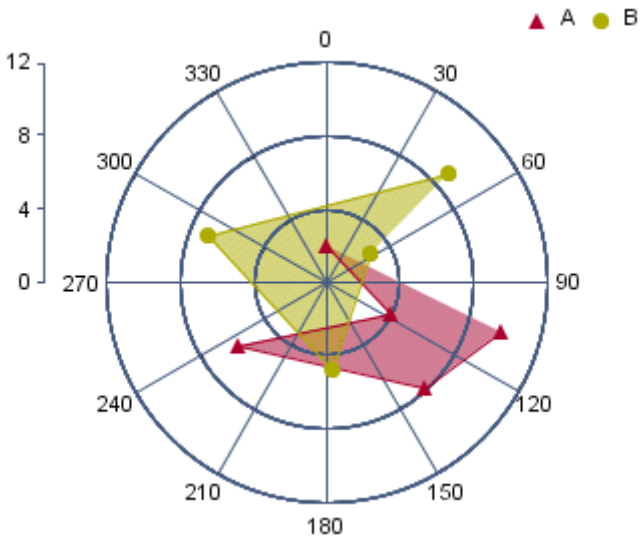
Category (X): Allows you to choose a data column whose distinct values will determine the categories.

Start: Allows you to choose a data column whose distinct values represent the starting time interval for the category.

End: Allows you to choose a data column whose distinct values represent the ending time interval for the category.

The data mapping also allows you to transpose the data (in other words: to select several columns for a single category). To learn more about data transposition, please see Section 5.2.3.1.1 - Data Transposition.

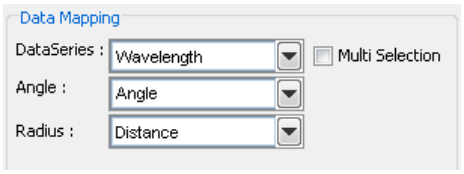
5.2.3.22. Polar Charts



Polar Chart

Like a two-dimensional scatter chart, a polar chart plots points on a plane. However, instead of rectangular coordinates, a polar chart plots points using polar coordinates (r, θ), where r represents the distance from the center of the circular plot and θ represents the angle of a ray emanating from the center of the plot and passing through the point. Like scatter charts, the data for polar chart coordinates must be numeric. The θ values can be supplied in degrees or radians. A third data series column can be used to separate the data points into groups.

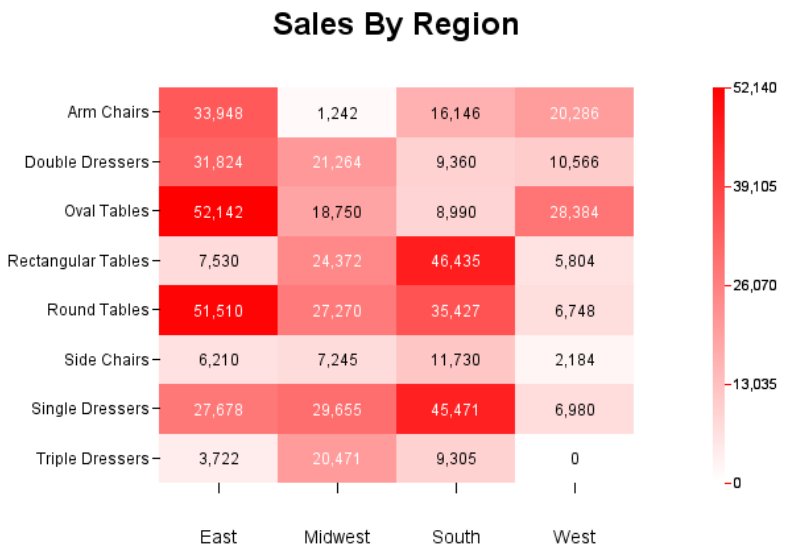
5.2.3.22.1. Data Mapping



Mapping Options for Polar Charts

The data mapping for polar charts is similar to scatter charts. The angle and radius options allow you to select the columns whose values will make up the polar coordinates. These must both be numeric. The data series box allows you to choose a data column whose distinct values will determine the number of data series in the chart. Each element in a data series is drawn using the same set of drawing attributes, e.g., colors and markers.

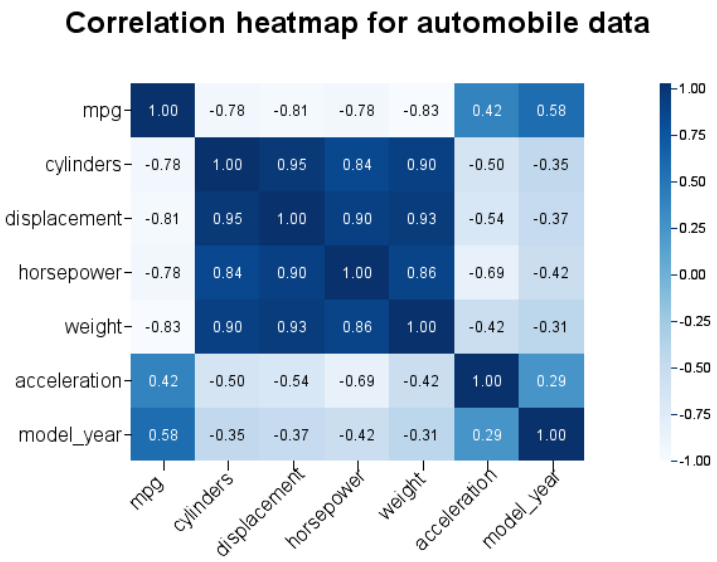
5.2.3.23. Heatmap Charts



Heatmap Chart

A heatmap is a data visualization tool that uses color to represent the values of a matrix. This method is particularly effective for displaying large amounts of data in a way that is easy to understand at a glance. The colors in a heatmap typically range from light (white) to dark (e.g.,red) along the color gradient, or from cool (e.g., blue) to warm (e.g., red), with each color representing a different value. This visual representation allows for quick identification of patterns, trends, and outliers within the data. A simple example that shows the revenue of certain product category in certain region is presented above. In this case, region is the X axis, product category is the Y axis and the plot is (X,Y, value).

In addition to just plotting (X,Y. value) straight from the input data, there is a correlation option that you can apply to the data. As such, you can easily visualize the correlation coefficients between pairs of columns from the data source, i.e. each entry in the matrix is the correlation coefficient between the corresponding columns on the X axis and the Y axis. Below is an example. In this example, the columns along the X axis are the same as those along the Y axis. But they do not have to be the same.



Correlation Heatmap Chart

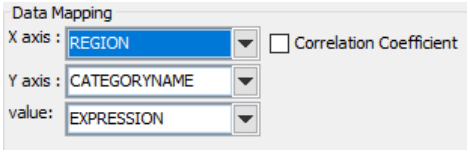


Note

Heatmap chart is available in a two-dimensional form only.

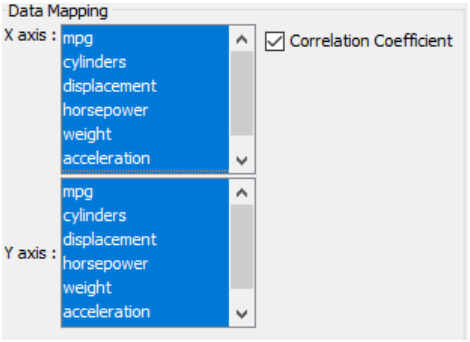
5.2.3.23.1. Data Mapping

Data mapping of heatmap without applying the correlation function is straightforward.



Heatmap data mapping

Data mapping of heatmap with correlation function applied.



Correlation heatmap data mapping

Note that the “Correlation Coefficient” checkbox is checked, and the X-axis and Y-axis show all the available numeric columns in the data source. You can select multiple columns in each of the axes.

In our earlier example, we selected all the columns for the X axis and all the columns for the Y axis. But you can choose different columns for each axis.

5.2.3.24. Changing Data Mapping or Data Source

Once you have finished setting the mapping options you will be taken to the main Chart Designer interface. Once you are in the Chart Designer you can return to the Data Mapping window by selecting Data → Modify Data Mapping, or by clicking the  *Change data mapping* icon. This will return you to the data mapping screen, allowing you to adjust the mapping for the chart.

You can also change a data source directly from the Chart Designer by selecting Data → Modify Data Source, or by clicking the  *Modify data source* icon on the toolbar. This will open the Data Source Manager.

5.2.4. The Chart Designer Interface

Once you have completed all the steps in the Wizard, the main Chart Designer interface will appear. Here you can customize and modify the appearance of the chart by changing properties and adding new elements.

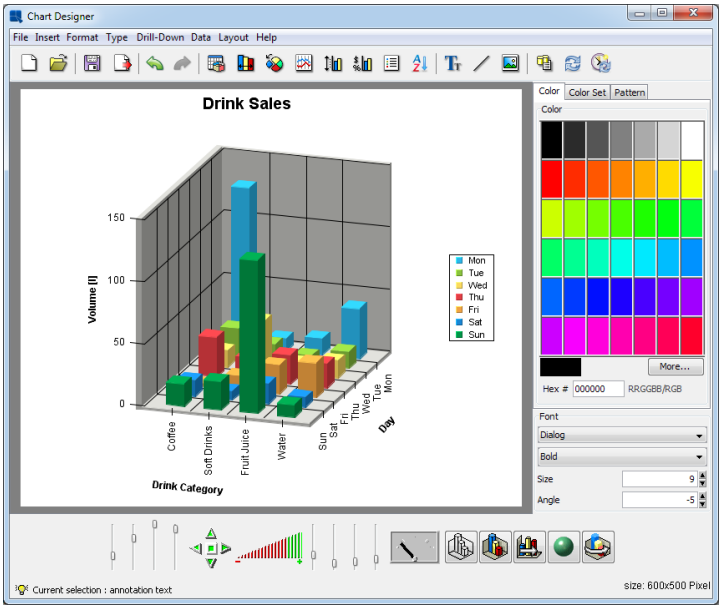


Chart Designer Interface

5.2.4.1. The Designer Menus

Most of Chart Designer's functions can be controlled from the menu bar at the top of the designer window. This section provides a brief overview of the available options. All of the features here are discussed later in this chapter.

5.2.4.1.1. File Menu

This menu performs basic file operations such as opening, closing, and saving files. Note that saving and opening options are not available when you use the report data as the data source for an embedded chart. For more information about file options, please see Section 5.2.6 - Saving Charts.

New: This will return you to the Data Source Manager to begin creating a new chart. If you have not saved your current chart, you will be prompted to do so.

Open: This allows you to open a saved chart definition. Files that can be loaded by Chart Designer include `.cht`, `.tpl`, `.pac`, and `.xml` chart definition files.

Close: This closes the current chart.

Apply Template: This option allows you to apply a template to the current chart. You can apply any chart template in either .tpl or .xml format. For more information about working with chart templates, please see Section 5.2.6.2.1 - Working with Templates.

Save: This saves the current chart.

Save As: This allows you to save the current chart. You can save the chart as a `.cht, .pac` or `.tpl` (a chart, a complete packed chart or a template file). The

`Create XML`

check-box allows you to create an XML chart definition file.

Export: This allows you to export the chart to a number of static image formats. The chart data can also be exported as XML file.

Exit: This closes the Chart Designer with the possibility to save the current file.

5.2.4.1.2. Insert Menu:

This menu allows you to add various elements to a chart.

Titles: This option allows you to automatically add titles to the chart. A main title can be specified as well as titles for each of the axes. Unlike annotation text, titles will size and position automatically with the chart.

Text: This allows you to add text or annotation to a chart. Text can be added anywhere on the chart and can have a number of different formatting properties. Variables can be added to the text for run-time substitution. For more on adding text to charts, see Section 5.2.4.6.1 - Adding Text

Background: This allows you to select an image to use as the chart background. Background images can be tiled, centered, or stretched to fit the entire canvas.

Dial Foreground: This option is only available for dial charts. This allows you to select an image to use as the dial chart foreground. The image can be stretched.

Dial Background: This option is only available for dial charts. This allows you to select an image to use as the dial chart background. The image can be stretched.

Link: This allows you to add a hyperlink to a data point or a collection of data points in a chart. To use hyperlinks in a chart, you will need to export a map file along with the image.

Line: This allows you to add an arbitrary floating line to a chart. These lines can be placed anywhere and can be used to draw enclosed shapes as well. Floating lines are often used as pointers and can be generated with arrowheads.

TrendLine: This allows you to add a trend line to the chart. ChartDesigner allows you to draw many different types of trend lines including: linear, a polynomial of any degree, a power, exponential, logarithmic, a moving average, exponential moving average, triangular moving average, cubic B-spline, and a normal distribution curve.

Horz/Vert Line: This allows you to add a fixed horizontal or vertical line to a chart. You have the choice of either adding a constant line (one that draws at a fixed value on the X or Y-axis) or a control line (draws lines based on a certain value range either average, minimum, maximum, or multiples of standard deviation).

Control Area: This allows you to draw a fixed area (either on the plot area of a 2D chart or on the face of a dial chart) to compare against the data values of the chart.

5.2.4.1.3. Format Menu

This menu allows you to edit and modify the properties of many different chart components.

Undo: Cancels the last operation performed in the designer and reverts back to the previous state. The designer will remember the last 10 actions made.

Redo: Reverses the action of the undo command. For example, if you change the font color from black to red, you can undo this command to change it back to black, and then redo this command to have it change to red again.

Data Properties: This option allows you to control several options for how the data is displayed. You can control the thickness of columns/bars, how null data is displayed, whether to draw data top labels or not, as well as enable and select a color for negative top labels.

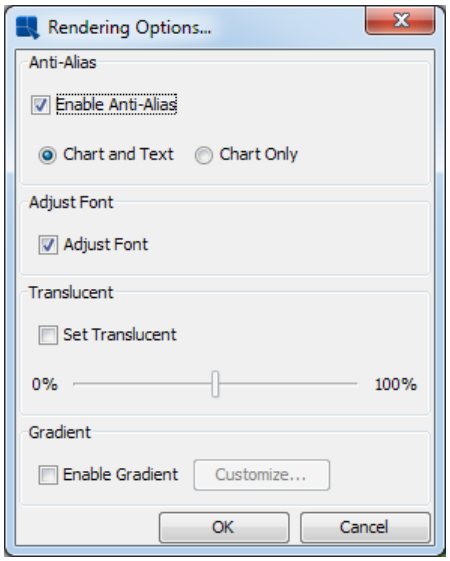
Histogram Options: This allows you to specify if you want to draw the chart as a histogram and display additional options to specify the frequency count.

Aggregation Options: This allows you to specify whether to aggregate the data before drawing the chart and display additional options to specify type of aggregation.

Zoom Options This allows you to enable/disable and set options for time based zooming. This option only applies if you have date/time data mapped to the category axis.

3D Display Options: This allows you to set several options that control the display of 3D charts. You can specify an inline series for 3D column (or similar) charts, as well as specify rendering approximation for 3D scatter and surface charts. (This improves performance for charts with a lot of data points.)

Rendering Options: This allows you to specify various rendering options for the chart for a better presentation.



Rendering Options Dialog

Among the options available are:

Enable Anti-Alias: This allows you to specify anti-aliasing for the chart. Anti-aliasing will smooth text and lines in the chart, creating a smoother appearance. This feature can be applied to either the entire file or just to the chart, which will leave the text unaffected.

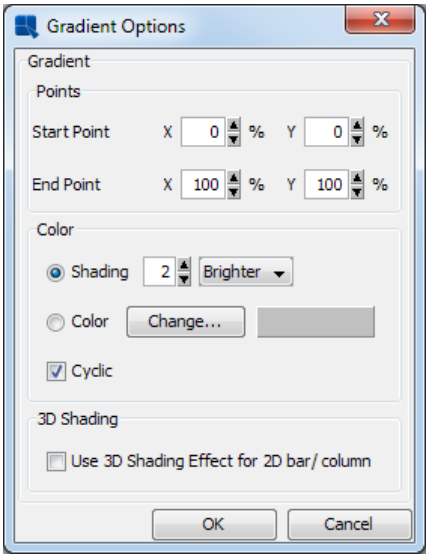
Adjust Font: This allows you to specify the font size for text in charts to be relative to the screen resolution. This feature allows for more precise conversions of charts to various export formats and between installations.

Translucent This allows you to specify whether the columns/bars/areas should be translucent. This

allows any columns/bars/areas “hidden” behind to show through. You can also specify the opacity by adjusting the slider, where 0% is completely opaque and 100% is completely transparent. This option is available for all 3D charts and for 2D Area, Radar and Gantt charts with data series.

Gradient:

This allows you to enable gradients. The gradient can be applied either across the entire canvas or to chart data points only. You can change the gradient options by clicking on Customize to show the following dialog:



Gradient Options Dialog

This dialog allows you to specify the start and end points for the gradient (based on the x-y plane), the gradient color scheme, and whether the gradient should be cyclic. The start and end points are represented as percentages with (0,0) being the top left corner of the canvas and (100, 100) being the bottom right corner of the canvas. You can also specify whether the gradient change is shading (becoming darker or brighter) or a different color. Lastly, you can specify the gradient to be cyclic i.e., toggle between the two opposites in the gradient. Note that the start and end points represent the first segment if the gradient is cyclic.

You can also enable 3D shading for certain 2D charts (Column, Bar, Stack Column, Stack Bar, 100% Column and Overlay) by selecting the *Use 3D Shading Effect for 2D bar/column* option (please note that this option is available for 2D bar/column charts only).

Font Mapping: This allows you to map system (true type) font files for the PDF export.

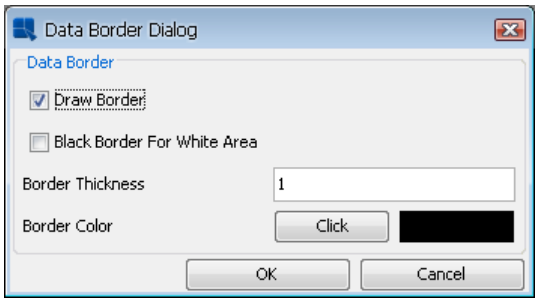
Chart Options: This brings up some specific options for the chart type that you are using. Options vary depending on chart type.

Axis Scale: This allows you to adjust the scale for any value axes. Automatic (best fit) scaling is used by default.

Axis Elements: This allows you to modify the appearance of the axes and axis labels. Options here include axis thickness, grid lines, label steps, and data formatting.

Canvas: This allows you to adjust the background canvas size. You can also specify whether to use scroll bars when the canvas size exceeds the view port.

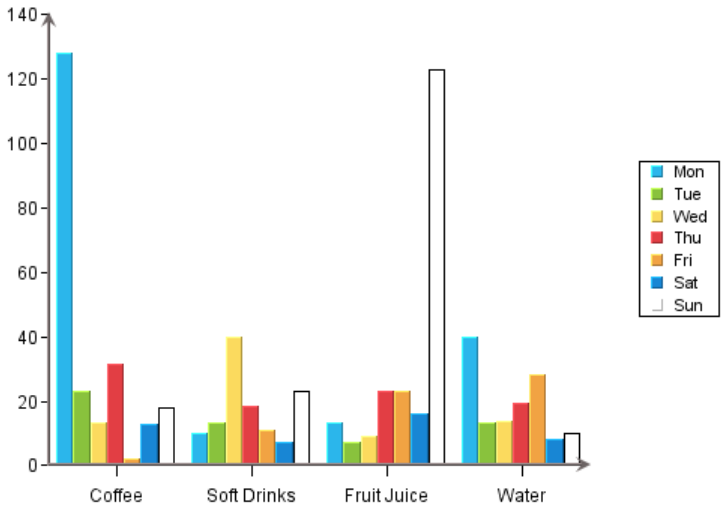
Data Border: Add a border to data elements (columns, bars, etc).



Data Border Dialog

If you select the *Draw border* option, the *Border thickness*, and the *Border color* fields will be enabled allowing you to configure the border properties.

The *Black border for white area* option can be selected even when the *Draw border* is disabled. If you do so, a simple black border will be added only to white data elements.



Column chart with the "Black border for white area" enabled



Note

This option is not available for Surface, Scatter, Line, Bubble, Box, Radar, Dial, and Polar charts.

Legend: This allows you to modify the display properties of the chart legend.

Lighting Model: This allows you to modify some of the lighting options for three-dimensional charts. You can modify both the light ambient color and the intensity.

Line and Point: This allows you to modify the display properties of any lines in a chart. For two-dimensional charts you can choose to display lines and points for any dataset as well as customize their appearance. You can also use this menu item to modify the display properties for any trend, floating, or horizontal/vertical lines.

Plot Area: This option allows you to customize appearance of the area bounded by the X and Y axes for two-dimensional charts.

No Data To Plot Message: This option allows you to set a message that appears if there is insufficient data to plot the chart. By default, the “No Data To Plot” message appears.

Flash Hintbox Customization: This option allows you to specify the font properties as well as the border and background color for the hint box in the flash export.

NULL Data Properties: This option allows you to show any category axis point that contains Null data and replace it with a different string. By default, any Null data category point is skipped.

Table: This allows you to add and configure a table displaying the chart data.

Text Properties: This option allows you to set a resize ratio for any text in the chart. From this option you can also specify to use Java 2D rotate text. This option gives rotated text a cleaner appearance. You can also specify any text replacement options.

Viewer Options: This menu item allows you to specify some configuration options for the Chart Viewer JNLP when the chart is viewed. You can control what options are available in the Viewer pop-up menu. These options can also be controlled through HTML parameters.

5.2.4.1.4. Type Menu

This menu allows you to change the current chart and its dimension. You can switch between 2D and 3D for chart types that support representations in each dimension. You can also change chart types. Note that as you switch between chart types, some formatting information will be lost. Also, you cannot switch between Gantt charts and other chart types.

5.2.4.1.5. Drill-Down Menu

This menu contains options allowing you to add and navigate drill-down layers in a chart. The drill-down features are explained in Section 5.2.5 - Chart Drill-Down. Note that drill-down is only available (functional) when charts are deployed independently from reports.

Add: This allows you to add a layer of data drill-down.

Remove This: This removes the current level of data drill-down.

Remove All: This removes all levels of data drill-down.

Previous This navigates to the previous layer of data drill-down.

Next: This navigates to the next layer of data drill-down.

Go To Top Level: This navigates to the top-level chart for data drill-down.

Dynamic: This allows you to enable dynamic data for drill-down.

Parameter Drill-Down: This brings up the parameter drill-down navigation window allowing you to edit, add, and remove layers of parameter drill-down.

5.2.4.1.6. Data Menu

This menu contains options that allows you to refresh, re-order or completely change the chart data.

Modify Data Mapping: This will bring back the Data Mapping Window, allowing you to change the data mapping for the chart.

Modify Data Source: This will take you back to the Data Source Manager, allowing you to select a new data source for the chart.

Modify Database: If the chart uses an independent data source that is from a database, this feature, like in Report Designer, will allow you to modify the database connection that the chart use. For more information about this feature, please see Section 4.1.2.5 - Editing Database Connections.

Modify Query: If the chart uses an independent data source that is from a database, this feature, like in Report Designer, will allow you to modify the query that retrieves the chart data. For more information about this feature, please see Section 4.1.2.4 - Editing Queries.

Query Parameters: This will allow you to re-initialize query parameters and change parameter values for the current chart. This option is only available if the chart uses a parameterized query as the data source.

Ordering: This option allows you to re-order the data points in a chart. You can arbitrarily change the order of any of the category elements or you can sort the categories by value.

Refresh: This will update the current chart with the latest data. The original data source must be available for this option to work.

Schedule Refresh: This allows you to set a schedule to refresh the data. This option is for deploying charts in the Chart Viewer JNLP. You can set a periodic refresh interval so that the chart will update itself with the latest data.

View Table: This option will bring up a window containing the data table from which the current chart is generated. The table will initially display only the first 20 records. Clicking on the *Show All Records* checkbox will display all of the records.

View Chart Data: Rather than viewing the entire data table you can just view the data points plotted in the current chart.

View Data Source Info: If the chart contains an independent data source, this option will bring up a dialog containing information about the data source that was used to create the chart. The data source type and location are displayed.

Go Back: This will go back to the original chart if you have traversed a link.

5.2.4.1.7. Help Menu

This menu allows you to view a version of the program and to open a documentation.

About: This shows you information about the version of the program.

Contents: This opens the ERES Cloud User's Guide.

5.2.4.1.8. Layout Menu

The options in this menu allow you to toggle various elements in the Chart Designer interface. From this menu you can turn on and off the font and color panel, the Chart Designer toolbar, and the navigation panel for 3D charts.

5.2.4.2. The Designer Toolbar

The toolbar at the top of the Chart Designer window offers easy access to the Chart Designer's most commonly used features and functions. The buttons perform the following functions:

 Start a new chart

 Open an existing chart

 Save the current chart

 Export the current chart

 Undo the last change

 Redo the last undone change

 Change data mapping for the current chart

 Modify data display properties

 Modify chart-specific options

 Modify line and point attributes

 Change axis scale

 Modify axis elements/display properties

 Modify legend display

 Change data ordering

 Insert annotation text

 Insert floating line

 Add/change background image

 Modify/change chart data source

 Refresh chart data

 Schedule periodic data refresh

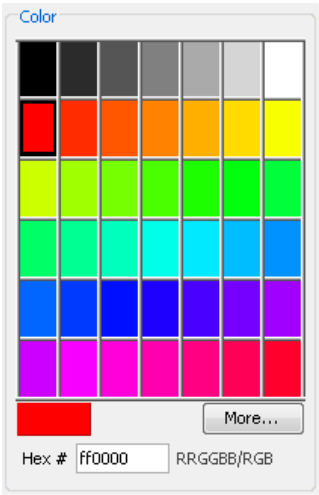
5.2.4.3. Color, Color set, Pattern, and Font Panels

The color, color set, and font panels on the right-hand side of the Chart Designer window allow you to modify the color of any chart object, as well as modify the font size and style for any text/labels in the chart. You can choose not to display these panels by toggling the Layout → Show font/color panel option.

For heatmap charts, you can use the *Color Panel* or *Color Set Panel* to set the colormap. The *Select Axis* option in the *Color Set Panel* is not applicable to heatmap charts. Additionally, the *Pattern Panel* does not respond to heatmap charts.

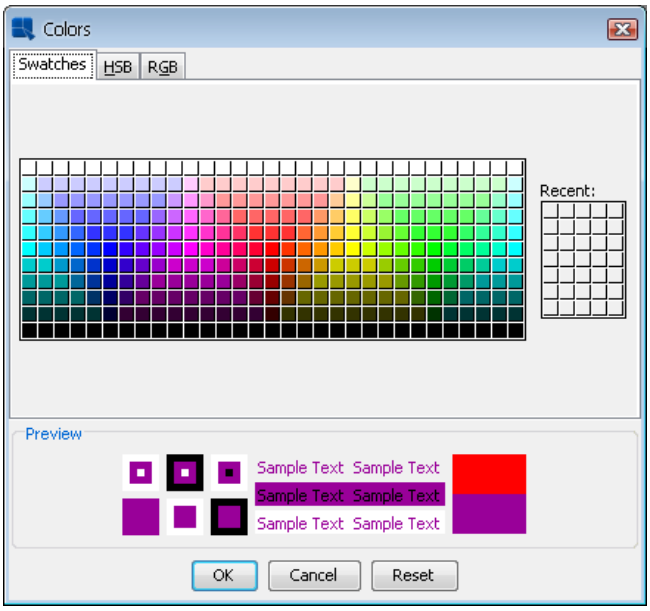
5.2.4.3.1. Color Panel

You can use the color panel to change the color of any element in the chart. To modify an element's color, first, click on it. The status bar at the bottom of the Chart Designer window will indicate which element has been currently selected. After you have selected the object, click on one of the panels in the color panel and the color of the object will change to reflect the color you selected.



Color Panel

To create a custom color for the object, first select it, then click the *More* button. This will bring up a new dialog allowing you to pick a color from a larger palette or to create a new color.



Additional Colors Dialog

From this dialog you can pick a new color from swatches or configure a custom color using HSB values or RGB values.

5.2.4.3.2. Color Set Panel

The color set panel allows you to choose a color scheme for your chart. It contains predefined color sets that can be applied to chart data points on the value or the second value axis.



Color Set Panel

In order to change a color set for the value or the second value axis, first select the *Color Set* tab. Then select the axis (*Value* or *2nd Value*) from the *Select Axis* select box and click the appropriate color set radio button. After that chart data points will get colors from the selected color set. Note that the *Select Axis* select box is only visible when the chart has the second value axis. By default, the value axis uses the first color set while the second value axis uses the seventh.

Please note that if you change a color of a data point manually, it will automatically unselect the selected color set. This is because of the color set that no longer corresponds to colors of data points in the chart. It is also important to note that if chart data points have more different colors than there are colors in the color set, it will automatically use colors from the beginning of the next color set, and so on. If there is no next color set available, it will use the first one instead.

5.2.4.3.2.1. Save Colors for Categories feature

Data points colors are closely related to the *Save Colors for Categories* feature that is available in the *Data Properties* dialog. The dialog can be opened by clicking the  *Change data properties* icon on the toolbar, or by using Format → Data Properties. (Please note that this feature is not available for Stack charts and it is disabled if the chart has Single Color for All Categories.) If the feature is turned on, colors of chart data points are assigned to names of categories (or series for charts with series). If such categories (or series) then appear in the chart, it will automatically use their assigned colors. If the feature is turned off (default), colors are assigned by data points order (i.e. the first data point will get the first available color from the color set, the second will get the second color, etc.). *Save Colors for Categories* setting is automatically saved in the `.pac` file. The following example shows scenarios with the feature on and off:

Assume that a chart has three categories (“A”, “B”, and “C”) with colors taken from the following color set (blue, green, yellow, red).

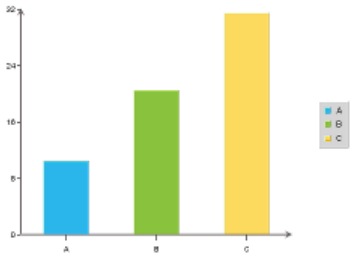


Image 1 - Example Chart

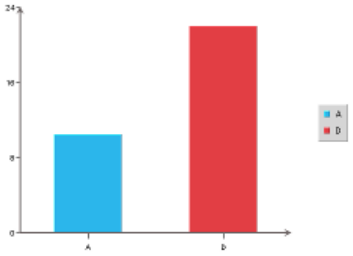


Image 2 - "Save Colors for Categories" Feature On

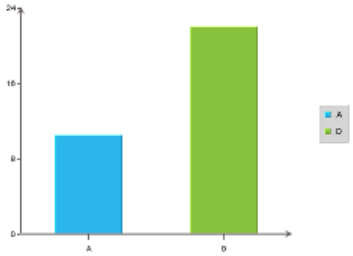


Image 3 - "Save Colors for Categories" Feature Off

Save Colors for Categories feature on

We saved the chart with the "Save Colors for Categories" feature enabled (see Image 1), so the following categories and colors were saved to the list of saved categories:

`category A` ... blue color

`category B` ... green color

`category C` ... yellow color

Now if you open the chart and data changes (e.g. there will be only categories `A` and `D` in the data - see Image 2), categories will have the following colors:

`category A` ... blue color (category `A` has blue color, because the category name `A` is present in the list of saved categories)

`category D` ... red color (category `D` has the next available color in the color set, because the category name `D` isn't present in the list of saved categories. In this case, the category will have red color, because blue, green, and yellow colors are already assigned to categories `A, B,` and `C`.

For this scenario, the color set will not be selected under the *Color Set* tab because colors of data points do not correspond to the color set.

Save Colors for Categories feature off

Here is the same situation, but assume that the chart has been saved with the "Save Colors for Categories" feature off (see Image 3). The list of saved categories will be empty this time.

After opening the chart, the categories will have the following colors:

`category A` ... blue color (category `A` has blue color, as it is the first available color from the color set)

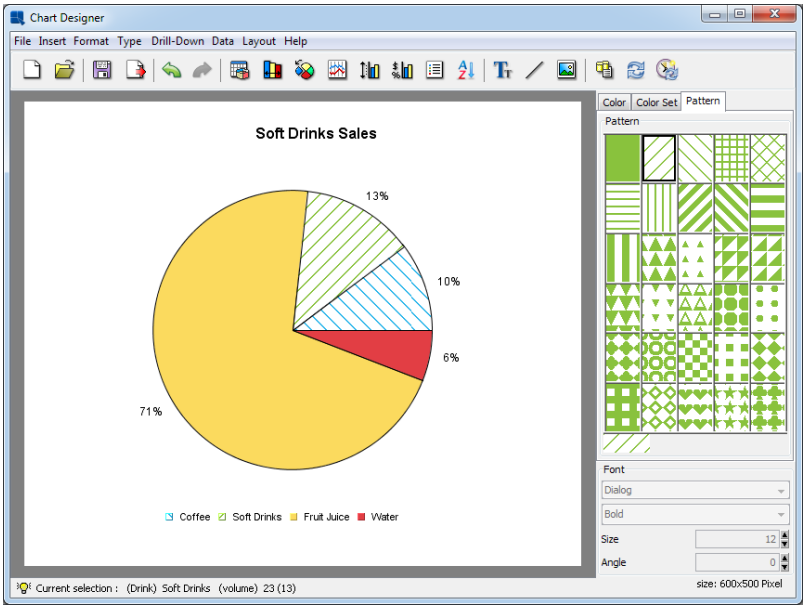
`category D` ... green color (category `D` has the second color from the color set, because the first is already assigned to the category `A`)

For this scenario, the color set will be selected under the *Color Set* tab because colors of data points correspond to the color set.

5.2.4.3.3. Pattern Panel

Unlike color and font panels, the pattern panel can only be applied to data points. There is a predefined pattern palette available for you to use. Similarly to how the color panel is used, you will need to first pick the data point you would like to change and then pick any pattern from the pattern palette. The pattern will be applied to the data point directly.

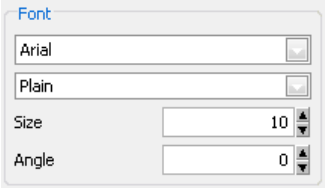
If a pattern has already been defined for a data point, the user can still change the color by selecting the color panel tab and picking a different color from the color palette. The color on the pattern will be changed to the new color immediately. The patterns shown on the pattern palette will change to the new color as well.



Pattern Example

5.2.4.3.4. Font Panel

You can use the font panel to modify the font, the font style, font size, and the font angle for labels, titles, or other text in a chart. To modify the font, you must first select the text object whose font you would like to change by clicking on it. The status bar at the bottom of the Chart Designer window will indicate which is the currently selected object.



Font Panel

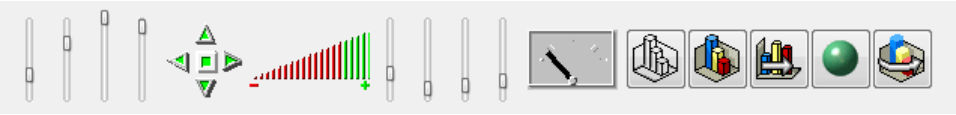
The first drop-down box allows you to select the font that you would like to use. The second drop-down box allows you to select the font style either plain, bold, italic, or bold and italic. The third box allows you to select the font size. The last box allows you to specify the angle of the text.

Certain groups of text (i.e. axis labels or data top labels) will all have the same properties. Hence, if you select one of text and modify the font properties it will apply to all of them.

ChartDesigner allows you to use the Java graphics libraries to give your text a cleaner appearance. For regular text you can use the chart anti-alias feature by selecting Format → Rendering Options. For rotated text (i.e. text not a 0 degrees), you can use the Java 2D rotate text feature by selecting Format → Text Properties. Note that these methods require Java 21 or higher.

5.2.4.4. The Navigation Panel

The navigation panel provides options for controlling a number of the properties specific to three-dimensional charts. It does not appear for two-dimensional charts and it can be hidden for three-dimensional charts by toggling the Layout → Show Navigation Panel option.



Navigation Panel

There are six controls and five buttons in the Navigation panel. Starting from the left, the six controls are step size, light position for X, Y and Z axes, navigation, zoom, scale and navigation speed. The buttons on the right control wire frame/solid mode, on/off border drawing, on/off inline series (for columnar and bar charts with series), on/off Gouraud shading, and on/off animation.

Step size: This slide bar allows you to set the incremental amount used by the navigation control to rotate or move the 3D chart. The higher the position of the horizontal bar, the larger is the increment for each step.

Light position for X, Y, and Z: User-defined light position is useful to control the position of the light and therefore the direction from which it will shine on each axis.

Navigation: This control allows you to rotate or translate the chart. The center button is a toggle switch. When it is in a depressed state (denoted by the red color of the button), clicking on any of the four triangular buttons moves the chart in the direction indicated by that button. When the center button is in an elevated state, clicking on any of the four triangular buttons rotates the chart in the direction indicated by that button. The speed of navigation may be controlled by the Navigation Speed control.

Zoom: This control allows you to effectively move the chart closer to or farther away from the viewer. To operate, you point the mouse to the control and click on the left mouse button. Then, while holding down the button, you drag the mouse to the left or right. When you drag the mouse to the right, the red area is extended to the right, and the chart will appear to be closer to you (zoom in to blow up the chart). When you drag the mouse to the left, the red area will move to the left, and the chart will appear to move farther away from you (zoom out to shrink the chart).

Scale: This is a set of four slide bars. The first three bars allow you to change the scale factors for the X, Y, and Z axes respectively. The fourth bar determines the thickness of a three-dimensional column, bar, line, or pie (depending upon the chart type). The higher the position of the horizontal bar, the larger the value. For two-dimensional charts, you can adjust the bar width for bar, column, stack bar, stack column, high-low, and HLCO charts by selecting Format → Data Properties.

Animation speed: This allows you to set the speed for chart navigation. This control is useful with the Navigation control and the Animation On/Off switch. It determines

how fast the chart moves or rotates. To operate, click and drag the indicator needle to the desired position. Moving the indicator needle to the right will speed up the translation or rotation and moving the needle to the left will slow it down.

In addition, the five buttons do the following:



Wire frame/ solid mode On/Off: This toggle switch allows you to view the three-dimensional chart as a wire frame when the switch is on or as a solid object when the switch is off.



Border Drawing On/Off: This toggle switch will draw a black outline on each edge of the chart when the switch is on.



Inline Series On/Off: This toggle switch will draw the series column in the same XY plane. This option is only available for columnar and bar charts with a data series and it does not appear on the navigation panel if the chart type is not applicable.



Gouraud Shading On/Off: Gouraud shading is a sophisticated and very realistic shading feature for three-dimensional charts. When the switch is on it will begin rendering the chart to shade each individual surface.



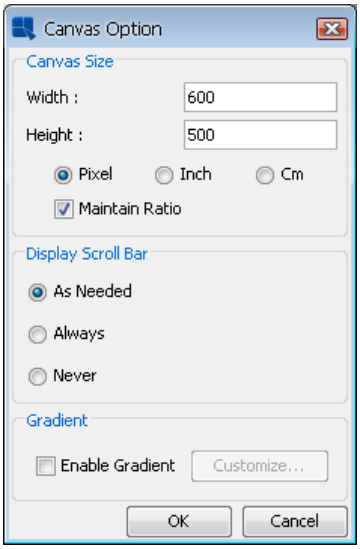
Animation On/Off: This toggle switch allows you to start/stop the animation of a three-dimensional chart. The speed of animation may be set using the Navigation Speed control. During animation, all panel controls except the animation speed control are disabled.

5.2.4.5. The Viewport

The viewport compromises the central portion of the Chart Designer window. Within the viewport you can select, move, and size all of the various chart elements on the canvas.

5.2.4.5.1. The Chart Canvas

The chart canvas is the background on which all of the chart elements are drawn. Its dimensions are the size of the finished chart. You can modify the size of the chart canvas by selection Format → Canvas. This will bring up a dialog prompting you to specify the new canvas dimensions.



Canvas Formatting Dialog

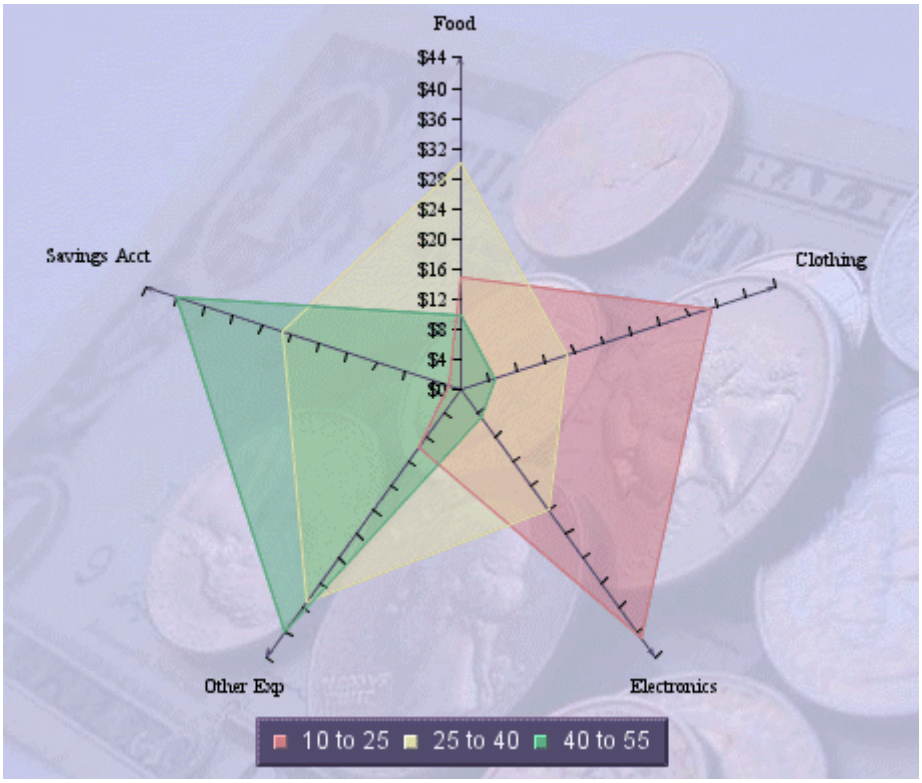
You can specify the canvas size in pixels, inches, or centimeters. From this dialog you can also specify when to use scroll bars in the viewport. By default the viewport will display scroll bars when the canvas is larger than the window. When the canvas is smaller than the viewport window, a dark gray area will appear around it.

If you're creating an embedded chart in a report, the canvas will resize to match the space that you have defined for the chart in the Report Designer. For example, if you set the space in the report to be 3 inches by 4 inches, the next time you edit the chart, the canvas will automatically size to 3 inches by 4 inches.

On this dialog, you can also set up gradient background for the canvas. The gradient settings are the same as in the *Rendering options* described in the Section 5.2.4.1.3 - Format Menu.

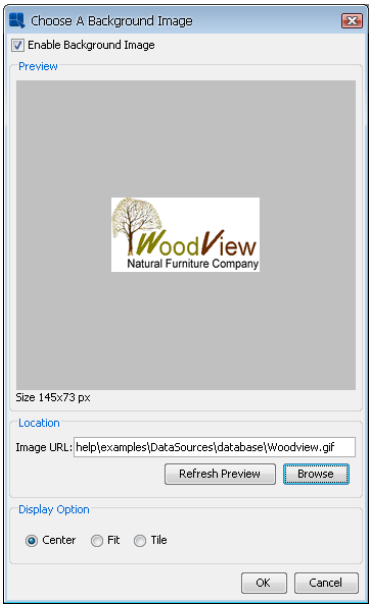
5.2.4.5.1.1. Background Images

You can add an image as the background of the chart instead of having a plain or a colored canvas.



Radar Chart with Background Image

You can add a background image by selecting Insert → Background or by clicking the  *Background* button on the toolbar. This will bring up a dialog allowing you to specify the background image for the chart.



Add Background Image Dialog

To insert a new image, select the *Enable Background Image* option. If you want to remove an existing background image and use simple background color instead, unselect this option.

There are two ways of inserting background images: either locate the image on the hard drive or retrieve it from an URL.

To locate an image on the hard drive, click on the *Browse* button.

To retrieve image from an URL, enter the URL in the *Image URL* text field. After that, click on the *Refresh Preview* button to verify the URL. If the image from the URL appears in the *Preview* section, the URL is correct.

If you add a background image and save the chart as TPL or CHT, the image itself is not stored with the chart. Only the path or URL is saved. If you move a TPL or CHT chart, you need to be sure that it can still access the image along the path specified. If you save the chart as PAC, the background image will be stored in the PAC file along with the chart.

5.2.4.5.2. Moving and Sizing Chart Elements

You can select any element in the chart by clicking on it. The status bar at the bottom of the Chart Designer window will indicate which object has been currently selected. Clicking and dragging on an object will move it around the chart canvas. Note that some objects like axis or data top labels move in tandem, while other objects like legends move independently.

You can move the entire chart plot by clicking in the plot area and dragging the mouse. This will move the entire chart around the canvas. To resize a chart click the plot area and 8 points will appear in the corners and on the sides of the plot area. Drag any of those points using the mouse to enlarge or shrink the plot area. You can also resize three-dimensional charts by using the zoom control in the navigation panel.

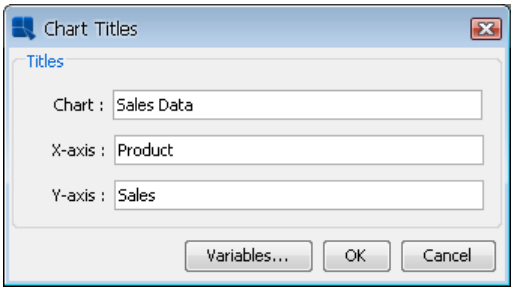
5.2.4.6. Adding Chart Elements

In addition to the default chart elements, ChartDesigner provides a number of additional elements that you can add to a chart.

5.2.4.6.1. Adding Text

There are two ways to add text to a chart: as titles or as plain text elements.

Adding Titles: To add titles to a chart, select Insert → Titles. This will bring up a dialog prompting you to enter titles for the chart.

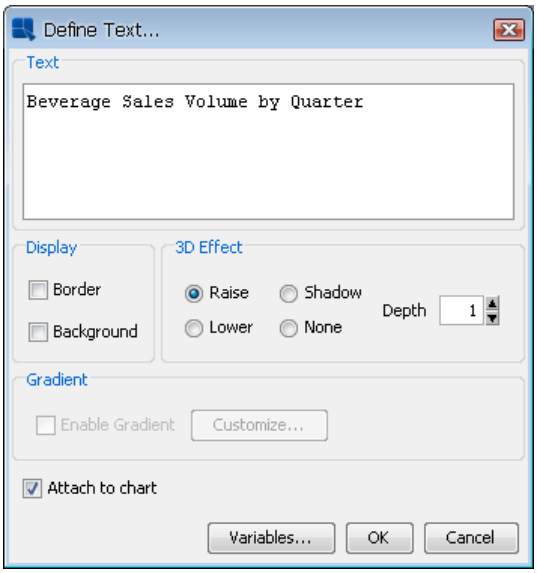


Add Titles Dialog

The dialog allows you to specify a main title for the chart and a title for each of the axes. Pie and dial charts, which do not have axes, prompt you for the chart title only. After you have finished specifying the titles click *OK* and they will be added to the chart. Titles are placed and sized automatically. However, they can be moved and the fonts can be changed.

Adding Text:

To add individual text fields to the chart, select Insert → Text or click the  *Add Text* button on the toolbar. This will bring up a dialog prompting you to add text to the chart.



Add Text Dialog

From this dialog you can specify the text and configure some display options. You can specify whether to draw a border around the text or around a background. You can also specify what effect you want to apply to the background.

On this dialog, you can also set up gradient background for the text label. The gradient settings are the same as in the *Rendering options* described in the Section 5.2.4.1.3 - Format Menu.

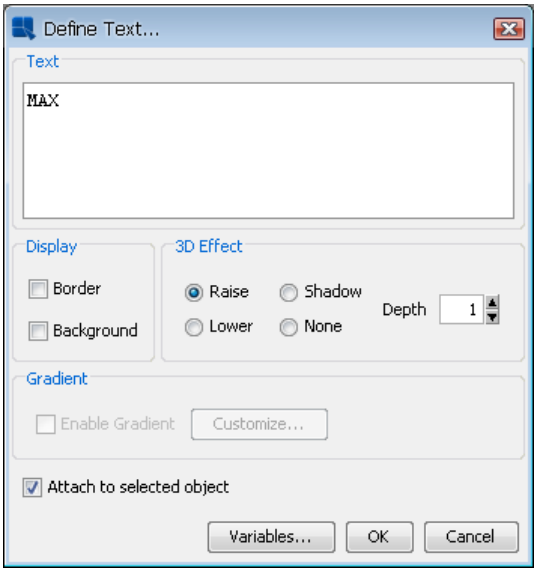
Once you have finished specifying the text, click *OK*. You will then be able to place the text in the chart. A small rectangle will follow your pointer around the chart canvas. Click the mouse where you would like to place the text.

Annotation Text:

ChartDesigner also supports annotation. Annotation allows you to attach labels or text fields to a particular element, such as a line or the chart plot. For example, you can insert a control line showing the maximum value to a chart and attach a text label called **MAX** to this control line. Each time the maximum value changes, the label will adjust its

position along with the control line. For more information about the control lines, please see Section 5.2.4.6.2.3 - Fixed Horizontal/Vertical Lines.

You can specify text to be annotation in two ways. To attach the text to the chart plot, select the *Attach to Chart* option when adding text. To attach the text to a specific object like a control line, first select the object, and then select Insert → Text or click the  *Add Text* button on the toolbar. The option for *Attach to selected object* will be automatically checked. Leave it checked and any text you add will be automatically attached to the object.

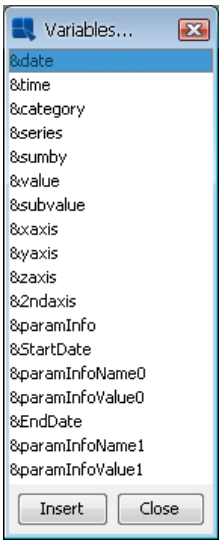


Add Text (Annotation) Dialog

5.2.4.6.1.1. TextVariables

ChartDesigner allows you to specify certain variables within text that allow for run time substitution based on certain values/objects in the chart. For example, if your chart uses a parameterized query as the data source, you could use the **`&paramInfo`** variable to display which parameter value(s) were selected at runtime.

Both the insert titles dialog and the add text dialog have a button marked *Variables* at the bottom. This will bring up a dialog with a list of variables you can use and it will allow you to select one to add to the title or to the text.



Variables Dialog

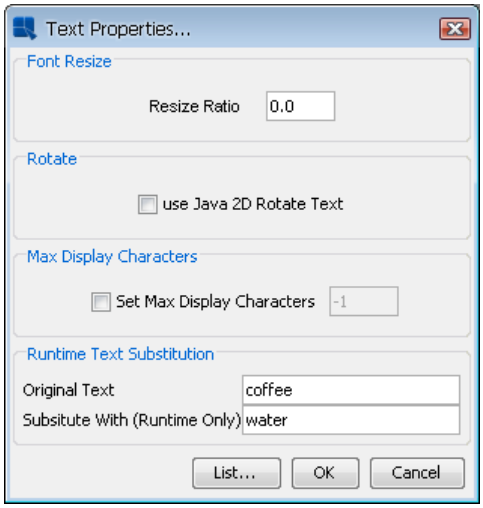
The following text variables are supported:

&drillInfo:	This displays which data point is being drilled on for drill-down charts. This variable does not work for parameter drill-down.
&paramInfo:	This displays the parameter value(s) that were selected. You can use this variable instead of &drillInfo for parameter drill-down charts.
&date:	This displays the date when the chart was last drawn/redrawn.
&time:	This displays the time when the charts were last drawn/redrawn.
&category:	This displays the name of the category column.
&series:	This displays the name of the data series column.
&sumby:	This displays the name of the sum-by column.
&value:	This displays the name of the value column
&subvalue:	This displays the name of the secondary value column
&xaxis:	This displays the name of the column that is mapped to the X-axis. This is for charts that map a value instead of a category to the X-axis like scatter or bubble charts.
&yaxis:	This displays the name of the column that is mapped to the Y-axis. This is for scatter, bubble, and surface charts.
&zaxis:	This displays the name of the column that is mapped to the Z-axis. This is for scatter, bubble, and surface charts.
&2ndaxis:	This displays the name of the column that is mapped to the 2nd-axis.
&paramInfoName<index>:	If the chart contains parameters, this displays the name of the parameter for the selected index.
&paramInfoValue<index>:	If the chart contains parameters, this displays the supplied parameter value for the selected index.
&<paramName>>:	If the chart contains a parameter with this name, this will display the value selected for that parameter.

5.2.4.6.1.2. Text Replacement

ChartDesigner allows you to overwrite a particular piece of text in a chart. This can be useful if the data source does not use particularly intuitive column names. Note that this feature will replace all instances of the text. For example, if you have a column chart without a series that displays a column name for both the X-axis label and the legend item, you cannot use text replacement to change only the label and not the legend item. The text replacement feature will also change only whole strings and not instances where there is a partial match.

To use text replacement, select Format → Text Properties. This will bring up a dialog allowing you to specify replaced text.



Text Properties Dialog

Please note that when making successive changes to the same text, the original text must be used. For example, assume you replaced the word “coffee” with “water”. Now if you want to change “water” to “soft drink” the text replacement should have the original text, which is “coffee” and then “soft drink” as the replacement. To remove any text replacement, simply replace the original string with itself. Hence, using the same example, you would replace “coffee” with “coffee”.

You can see a list of all the original text and the replacements by clicking on the *List* button. This will bring up a dialog listing all of the text replacement in the chart.

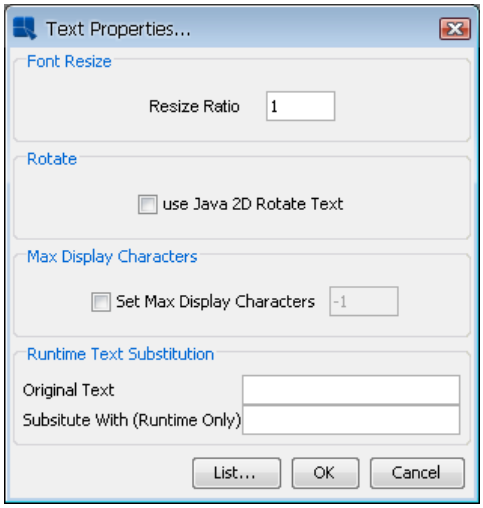


Replaced Text List

From this dialog you can select any of the replaced text and modify or undo the replacement by clicking the buttons at the bottom of the dialog.

5.2.4.6.1.3. Automatic Text Resizing

ChartDesigner has the ability to automatically adjust the font size of the text in a chart as it is resized. This is useful if you're using the same chart template to produce a number of charts in different sizes. You can specify a ratio for the font size to adjust based on changes in the chart canvas. To specify a resize ratio, select Format → Text Properties.

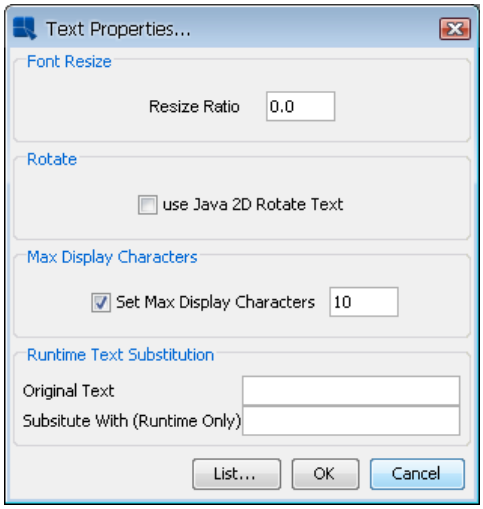


Text Properties Dialog

The ratio dictates the relative percent that the font should resize in regards to the canvas. For example, say you resize a chart from 500 x 500 pixels to 250 x 250. With a resize ratio of 1 then text with 12 point font would decrease to 6 point, decreasing by the same percent as the canvas. However, with a resize ratio of 0.5 the font would decrease half as much as the canvas so our 12 point font would only decrease to 9 point.

5.2.4.6.1.4. Text Cropping

Long labels or text in a chart can sometimes take up too much space in the chart plot, leaving little room for the actual chart. For situations like this, ChartDesigner offers a text cropping feature for chart text. Text that is longer than a user-supplied threshold will be truncated with "...". The hint box for the chart will show the whole label. To specify text cropping, select Format → Text Properties.



Text Properties Dialog

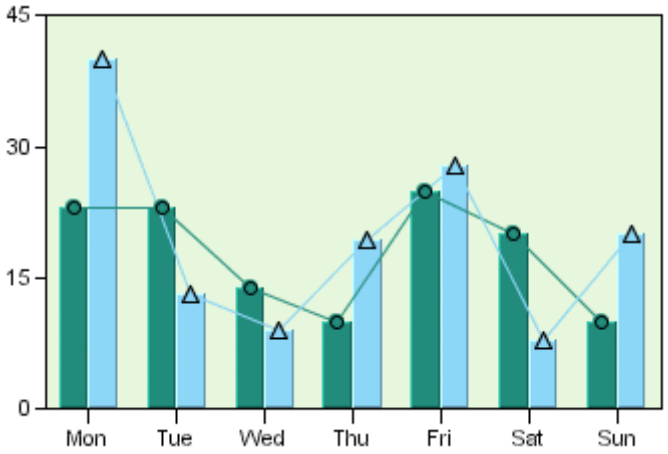
To enable text cropping, check the *Set Max Display Characters* option and specify the maximum character length in the dialog. Any text longer than the specified number of characters will be truncated.

5.2.4.6.2. Adding Lines

ChartDesigner allows you to add and format a number of different types of lines for charts.

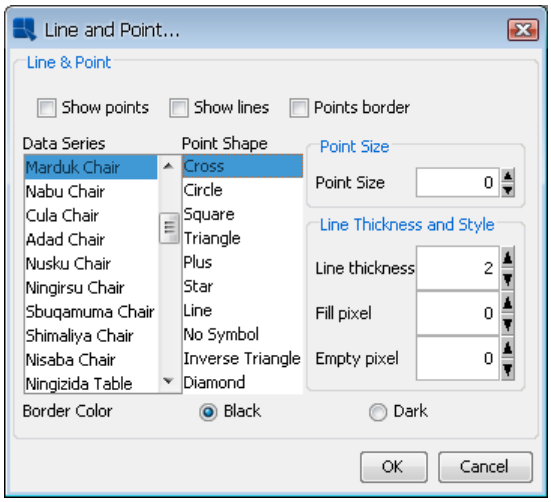
5.2.4.6.2.1. Line and Point Formatting

You can choose to display lines and points for all the data points in the chart for any two dimensional chart type. Note that some chart types already use this representation (i.e. line or scatter charts).



Column Chart with Lines and Points Defined

Line and point display is controlled by selecting Format → Line and Point, or click the  *Line and Point* button on the toolbar. This will bring up a dialog presenting several options.



Line and Point Dialog

The first three options allow you to specify whether you would like to show lines, points, and a points border for the chart. For radar charts you also have the option of showing areas. The remaining options allow you to customize the line and point displays for each element in the data series.

For each data series element, you can specify the point shape that you would like to use. You can also control the size of the points. The default point size is 0. You can specify point sizes of -1, -2, & -3 which represent sizes of 0.75, 0.5, and 0.25 respectively. At -3 (0.25), the point will be drawn as a dot regardless of the selected point shape.

For lines you can specify the line thickness, as well as customize a dash pattern. The dash pattern is created by specifying the number of filled pixels and the number of empty pixels (between 0 and 255). The line is then drawn by dividing into segments - the number of filled pixels followed by the number of empty pixels. Setting 0 for both will result in a solid line. Setting 255 for both will result in no line being drawn.

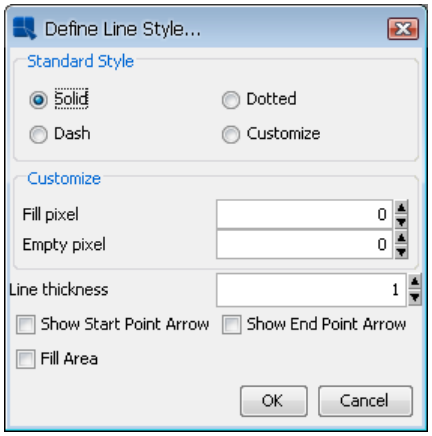
The last option allows you to change data point symbol border color to black or darker shade of symbol color.

5.2.4.6.2.2. Floating Lines

Floating lines are free-form lines that can be arbitrarily added to any place on the chart canvas. Often floating lines are used to point to a specific element in a chart. To add a floating line select Insert → Line, or click the  *Insert Line* button on the toolbar.

When you select this option, your mouse pointer will change to a cross. Click within the chart canvas where you would like the line to begin. Each additional click will add a point to the line, allowing you to add another segment. This way you can use floating lines to draw shapes as well. Once you have finished, right-click to stop drawing. The line will then be added.

Once a line has been placed on the canvas, it cannot be moved individually. It will move with the chart plot, like annotation text. To specify options for a floating line, first select it, and then select Format → Line and Point or click the *Line and Point* button on the toolbar. This will bring up a dialog allowing you to format various properties for the floating line.



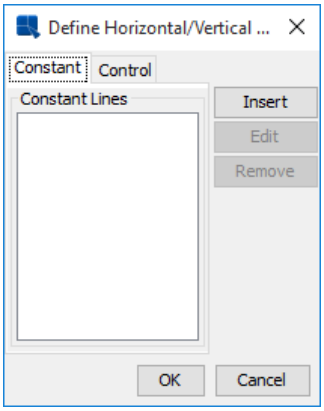
Line and Point Dialog for Floating Lines

The dialog allows you to specify a standard line style or to create a custom dash pattern in the same manner as line and point formatting. You can also specify the line thickness in pixels. The checkboxes at the bottom of the dialog allow you to place an arrowhead at the start and/or end of the line, as well as fill the area enclosed by the line to create a solid shape.

Floating lines can be deleted by selecting them and hitting the delete key.

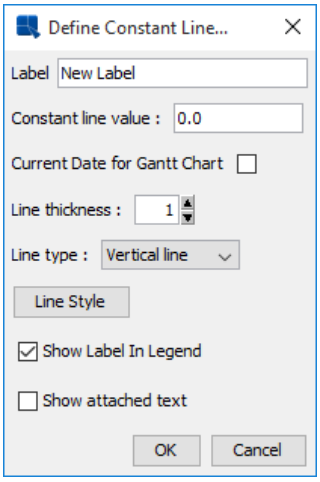
5.2.4.6.2.3. Fixed Horizontal/Vertical Lines

Fixed horizontal or vertical lines are lines that are drawn on one of the chart axes. These lines can also be drawn on three-dimensional charts where they appear as planes. There are two types of fixed lines: constant lines and control lines. Constant lines are lines that are fixed to a certain value in an axis. Control lines are drawn based on computed values that allow you to spot data points that are outside of certain value ranges. To add either type of line to a chart, select *Horz\Vert Line* from the *Insert* menu. This will bring up a dialog showing the list of existing horizontal/vertical lines, allowing you to edit the selected line, remove the selected line, and/or create a new line.



Define Horizontal/Vertical Lines Dialog

Clicking on *Insert* or clicking on *Edit* when an existing line is selected, respectively, will bring up a dialog allowing you to configure the selected line.



Constant Line Dialog

For constant lines you need to specify a label for the line, as well as the numeric value to use for the line. Note that for the category axis, the data points start with 0.5. You can also specify the line thickness and whether the line is horizontal or vertical. The last two options allow you to add an item to the chart legend for the line and whether to display any annotation for the line.

For Gantt charts, there is one extra option called *Current Date for Gantt Chart*. If you choose this option, the *Constant line value* field will deactivate and the current date will be used as the line value (the line position will be updated every time you run the chart).

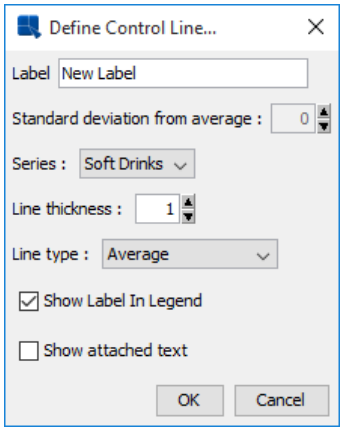
For radar charts, the horizontal/vertical option is disabled. Radar charts draw constant/control lines at the same point around all the chart axes (in a similar manner to the radar grid). In addition, for radar charts, an additional option named *Circular Style* is present. By default, lines in radar charts are drawn in a segmented fashion - straight lines connect the points on each axis. Selecting this option will draw the constant/control as a circle.



Note

If you have not specified any annotation for the line then none will appear if you select the last option. For more on adding annotation to a chart, please see Section 5.2.4.6.1 - Adding Text.

To add a control line, click on the *Control Line* tab in the dialog.

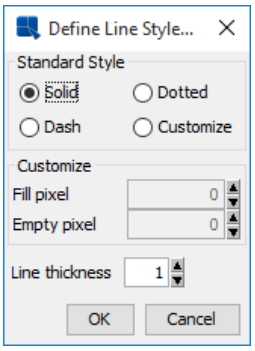


Control Line Dialog

For control lines you need to specify which series element you want to compute the value for (this option does not appear if no series is present) and how to compute the value. Options for control lines are average, minimum, maximum, and multiples of standard deviation.

After you have specified all of the options, the line will be added to the chart or the selected line will be changed, respectively. To edit any of the properties specified in the previous dialogs, you can select Insert → Horz/Vert Line again and select the line from the list. You can also double-click on the line that you want to modify.

You can change the appearance of fixed lines by first selecting the line and then selecting Format → Line and Point, or clicking the *Line and Point* icon on the toolbar. This will bring up a dialog allowing you to customize the line.



Line and Point Dialog for Fixed Lines

This dialog allows you to specify a standard line style or to create a custom dash pattern in the same manner as line and point formatting.

5.2.4.6.2.4. Trend Lines

A powerful feature of ChartDesigner is the ability to add trend lines to charts. Trend lines can help to show more details of a chart's data by exposing and highlighting certain trends.

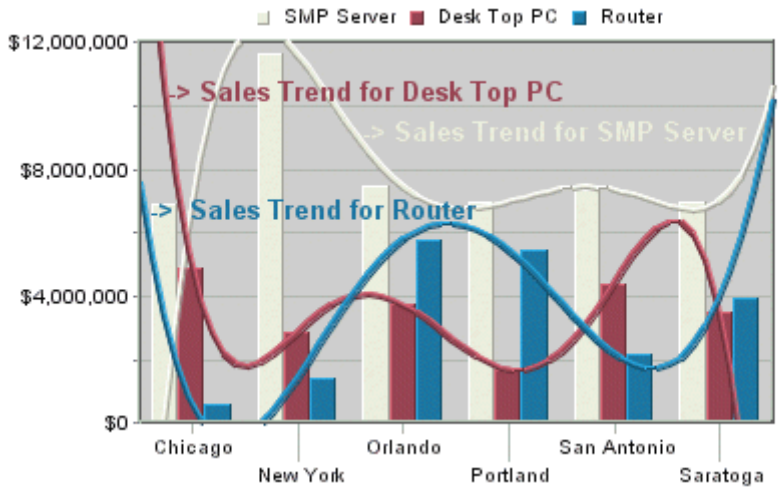
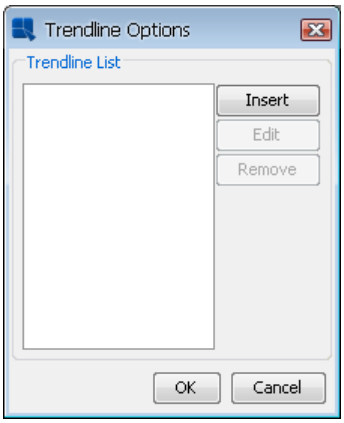


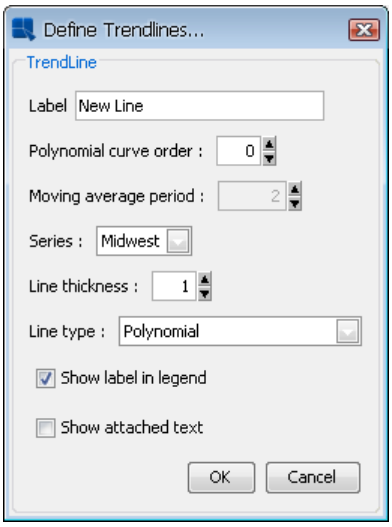
Chart with Trend Lines

To add a trend line to a chart, select Insert → Trendline. This will bring up a dialog showing the list of existing trend lines, allowing you to edit the selected trend line, remove the selected trend line, and/or create a new trend line.



Trend Line Options Dialog

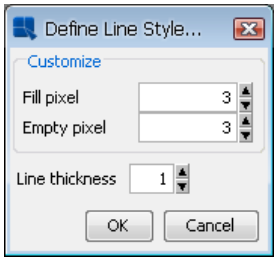
Clicking on *Insert* or clicking on *Edit* when an existing trend line is selected, respectively, will bring up a dialog allowing you to configure the selected trend line.



Define Trend Lines Dialog

In this dialog you can specify a label for the line, as well as what element of the data series to base the calculation on. The following types of trend lines are supported: a polynomial of any degree (please note that a linear trend line is a polynomial trend line of the 1st degree, i.e. the *Polynomial curve order* option has to be set to *1*), a power, exponential, logarithmic, a moving average, an exponential moving average, an triangular moving average, cubic B-spline, and a normal distribution curve. For moving averages you will need to specify the average period and for a polynomial you will need to specify the curve order. You can also specify the thickness of the line and configure whether a label in legend and the attached text should be shown. In case the chart has data series, you can configure the trendline for a specific series.

After you have specified all of the options, the trend line will be added to the chart or the selected trend line will be changed, respectively. To edit any of the properties specified in the previous dialog, you can select 'TrendLine' from the Insert menu again, and select the line from the list. You can change the appearance of the trend line by first selecting it, and then selecting Format → Line and Point, or clicking the *Line and Point* button on the toolbar. This will bring up a dialog allowing you to customize the lines.



Line and Point Dialog for Trend Lines

This dialog allows you to create a custom dash pattern in the same manner as line and point formatting.

5.2.4.6.2.4.1. Normal Distribution Curve

A special type of trend line that allows you to draw a normal distribution curve for the data in the chart. In order to plot a normal distribution curve, the chart must either be a two-dimensional column or bar chart, it cannot have a data series, and the category should be numeric. Assuming these conditions are met, you can specify a normal distribution curve as one of the trend line options.

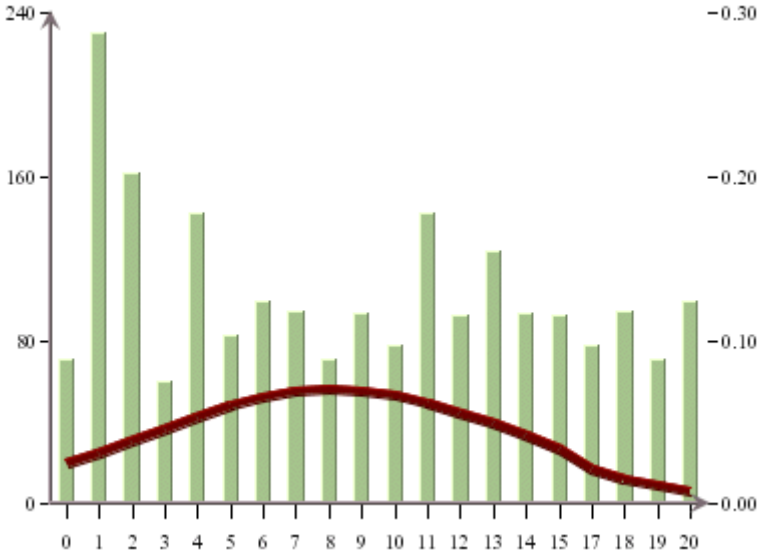
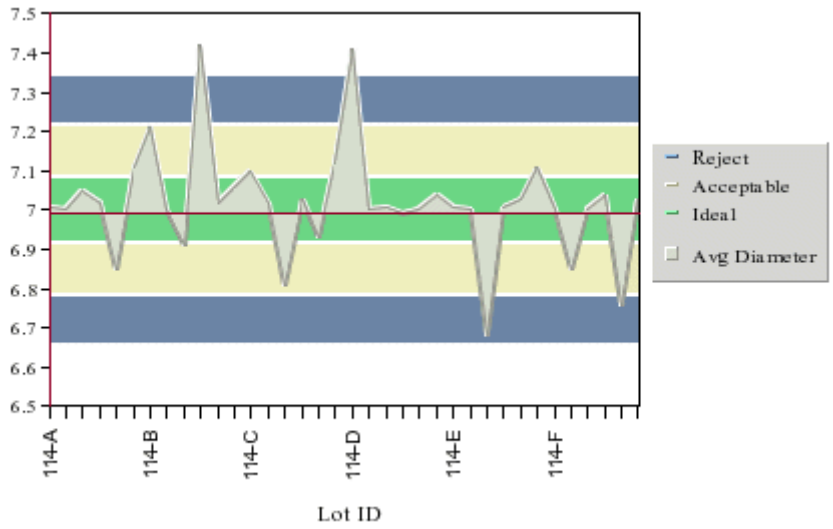


Chart with Normal Distribution Curve

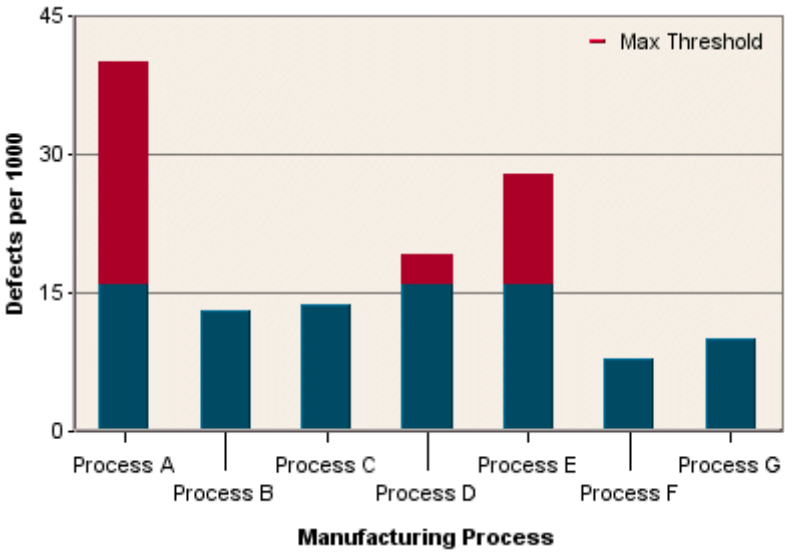
Since the scale for the curve is usually different than the scale for the value axis, the curve is shown on a secondary axis. You can modify the scale by changing the scale for the secondary axis.

5.2.4.6.3. Adding Control Areas

Control areas are useful for comparing the chart data against a certain range of data. For most two-dimensional charts, control areas are drawn as filled areas on the chart plot between a range of values on the chart's value axis and/or category axis. The data points are then drawn over top of the control areas giving you a quick visual reference to see which data points fall within the designated range. Instead of drawing a background area on the plot, the control areas can also be shown only where the data points intersect the control area. This feature gives users a clear visual reference when specific threshold values are reached.

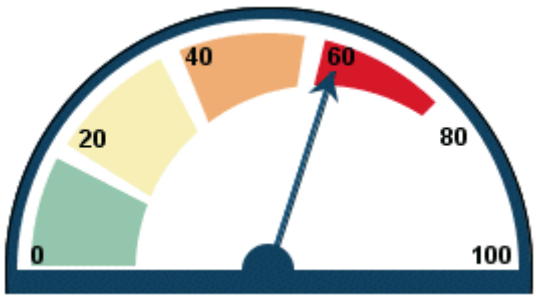


Two-Dimensional Area Chart with Control Areas



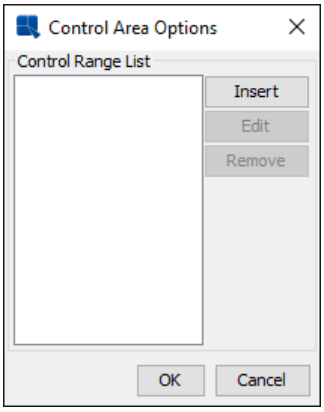
Column Chart with Control Area Drawn for Data Points

A special instance of control areas can be used for dial charts. For dial charts control areas are drawn as arcs on the face of the dial, allowing you to see if the dial hands (data points) fall within the range. Note that control areas are not available for radar and pie charts.



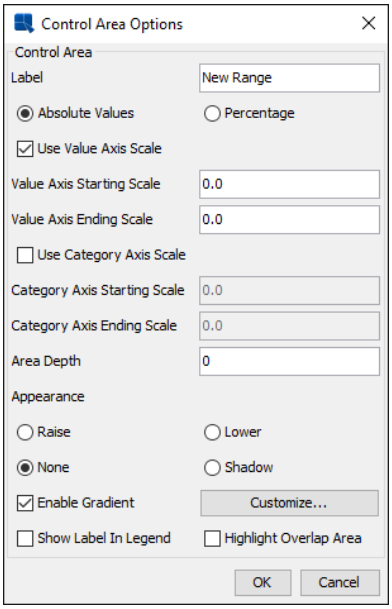
Dial Chart with Control Areas

To add a control area to a chart, select Insert → Control Area. The following dialog will appear showing the list of existing control areas, allowing you to edit the selected control area, remove the selected control area, and/or create a new control area.



Control Area List

Clicking on *Insert* or clicking on *Edit* when an existing control area is selected, respectively, will bring up a dialog allowing you to configure the selected control area. If your chart is not a dial chart, the following dialog will open.



Control Area Configuration Dialog

The following options are provided for control areas:

Label: This option allows you to specify a label for the control area.

Absolute Values: This option allows you to specify the scale in absolute values.

Percentage: This option allows you to specify the scale in percentage.

Use Value Axis Scale: This option allows you to specify whether the control area should be bounded by values on the value axis of the chart.

Value Axis Starting Scale: This is the lower bound for the control area on the value axis.

Value Axis Ending Scale: This is the upper bound for the control area on the value axis.

Use Category Axis: This option allows you to specify whether the control area should be bounded by values on the category axis of the chart.

Category Axis Starting Scale: This is the lower bound for the control area on the category axis.

CategoryAxis Ending Scale: This is the upper bound for the control area on the category axis.

Area Depth: This option specifies the depth for any of the appearance styles. If no style is selected, the depth will have no effect.

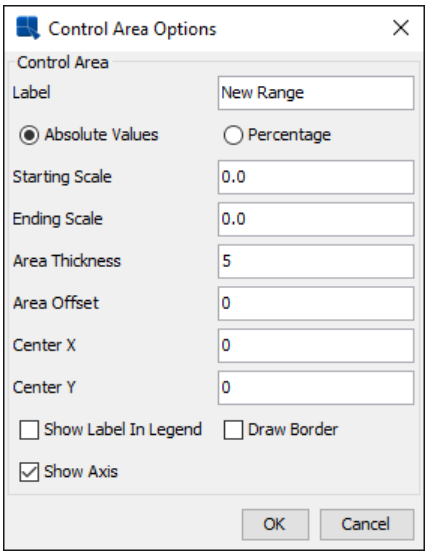
Appearance: This option allows you to specify a 3D or shadow effect for the control area. If the area depth is specified as zero, the appearance will not take effect.

Enable Gradient: Enable color gradient for the control. Gradient settings are described in Section 5.2.4.1.3 - Format Menu

Show Label In Legend: This option specifies whether or not to show the control area label in the chart legend.

Highlight Overlap Area: This option will only show the control area where the data points overlap the control area boundaries.

If your chart is a dial chart, then a different dialog will appear when you select Insert → Control Area and then click on the *Insert* or *Edit* button.



Control Area Dialog for Dial Charts:

The following options are provided for dial chart control areas:

Label: This option allows you to specify a label for the control area.

Absolute Values: This option allows you to specify the scale in absolute values.

Percentage: This option allows you to specify the scale in percentage.

Starting Scale: This is the value where the control area begins.

Ending Scale: This is the value where the control area ends.

Area Thickness: This option allows you to set the thickness for the control area

Area Offset: This option allows you to specify the offset in pixels from the edge of the dial chart

Center X: This sets the X coordinate for the center of the control area. 0 shares the same center point as the dial face. You can specify a new number (either negative or positive) pixels to specify an offset position from the center of the dial.

Center Y: This sets the Y coordinate for the center of the control area. It works in the same way as the previous option.

Show Label in Legend: Specifies whether to show the control area label in the chart legend.

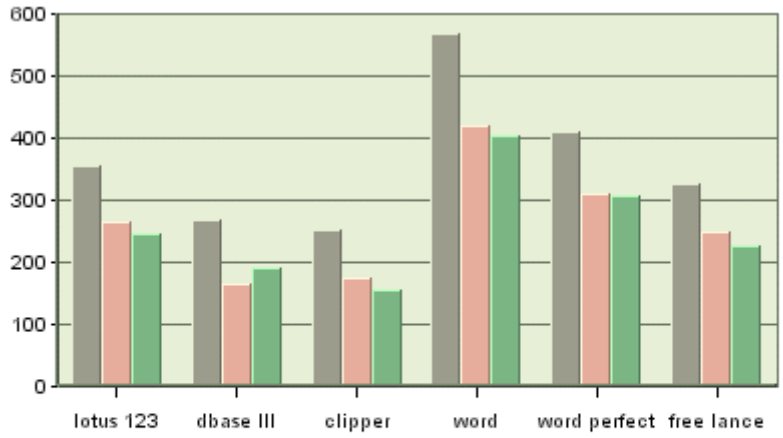
Draw Border: This option allows you to draw a border around the control area.

Show Axis: This option allows you to show or hide the axis for the remaining area not covered by the control range.

After you have specified all of the options, the control area will be added to the chart or the selected control area will be changed, respectively. To edit any of the properties specified in the previous dialog, you can select Insert → Control Area again, and select the control area from the list. You can also double-click on the control area that you want to modify.

5.2.4.6.4. Adding Tables

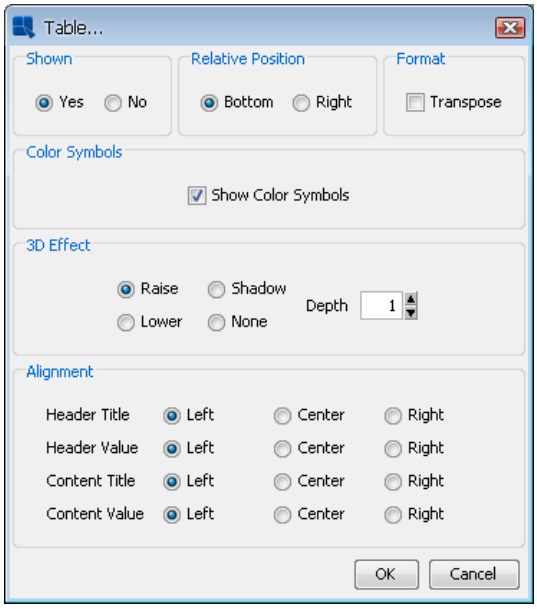
In addition to displaying charts, you can also display a table showing the data points displayed in the chart. The table can be placed below or to the right of the chart plot.



PERIOD_END	lotus 123	dbase III	clipper	word	word perfect	free lance
1987-03-31	355	270	253	569	410	327
1987-02-28	267	165	177	420	311	251
1987-01-31	248	193	155	405	307	229

Chart with Table

To add a table to a chart, or to modify the various display options for a table, select Format → Table. This will bring up a dialog allowing you to customize various options for the table display.



Format Table Dialog

From this dialog you can specify whether or not to display the table, as well as what relative position to give the table either to the bottom or right-hand side of the chart plot. You can also specify a 3D effect for the table and its depth.

The *Transpose* checkbox allows you to swap the columns and rows of the table. By default, the category elements are drawn as columns and the data series elements as rows.

The *Show Color Symbols* option allows you to show/hide color boxes for data series in the table.

quarter\drink	Coffee	Fruit Juice	Soft Drinks	Tea	Water
Q1	181	89	264	114	303
Q2	149	144	212	144	156
Q3	169	70	498	162	275
Q4	195	171	230	144	186

Chart table with color boxes

The *Alignment* options allows you to specify horizontal alignment of the text in the table cells, either left, center, or right. The alignment can be set for row headers, column headers, and inner table cells.

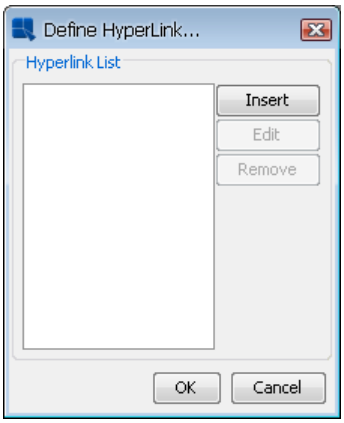


Note

If there isn't enough room in the chart canvas, not all data points will be displayed in the table. The table size adjusts with the canvas size and also with the font size in the table cells.

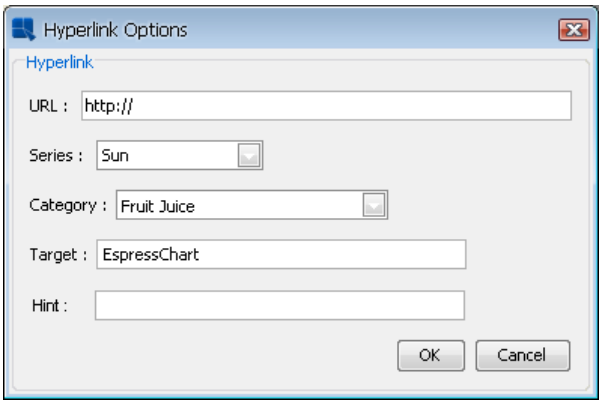
5.2.4.6.5. Adding Hyperlinks

ChartDesigner has an capability to add hyperlinks to any data point in a chart. Links can be specified for either single data points or multiple elements. Any added hyperlinks will be applied to both the data points on the chart and to their respective fields on the legend. To add a hyperlink to a chart, select Insert → Link or right-click on a data point and select *Insert Link* from the pop-up menu. This will bring up a dialog showing a list of existing hyperlinks which you can configure, remove, and/or create new ones.



Insert Link Dialog

Clicking on *Insert* or clicking on *Edit* when an existing hyperlink is selected will bring up a dialog allowing you to configure the selected hyperlink.



Define Link Dialog

The URL field allows you to specify a Web page that you want to link.

The *Series* and *Category* drop down menus allows you to select an element in the data series and category elements for the hyperlink. You can also link to all data series elements or all category elements.

You can specify the `Target` parameter recognized by HTML when specifying a hyperlink to be attached to a data point or data series. This can be used to determine whether the new HTML page should open in a new browser window, in the same browser window, or whether the new page should occupy the same portion of the page as the current page.

The *Hint* field allows you to enter text that will pop-up when you move your mouse cursor over a data point. If you want to create pop-up labels without hyperlinks, you can leave the URL field blank and only specify the hint.

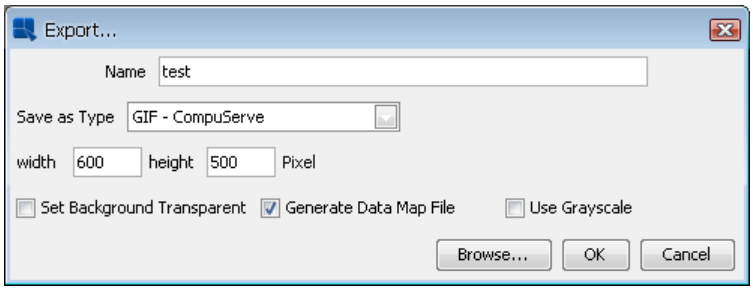
5.2.4.6.5.1. Viewing Hyperlinks

If you specify hyperlinks for charts, they will be active only when the chart is exported to Flash format*. For most image formats such as PNG, JPG, GIF, etc, an image map file containing information for the link will be automatically generated when you export the chart. You can insert the image and image map into an HTML file to view the image with clickable links.



Note

* For Flash export, when the user clicks on the hyperlink, there will likely be a warning message prompted by Flash saying that there was a potentially unsafe operation. You can turn this warning off by clicking on settings and adding the chart into the list of trusted locations.



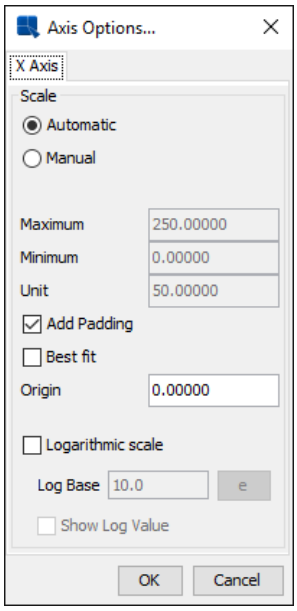
Export Dialog

5.2.4.7. Formatting Chart Axes

ChartDesigner provides a number of extensive formatting capabilities for the chart axes. Users can customize everything from the axis scale to the way how axis labels should be displayed.

5.2.4.7.1. Axis Scale

By default, the scale of any value axes in the chart is calculated to provide a 'best fit' for the data being plotted. This is often a useful feature if the data being displayed can change radically. However, you may often want to set the scale of the axes manually. To modify the axis scale, select Format → Axis Scale, or click the  *Axis Scale* button on the toolbar. This will bring up a dialog allowing you to format the scale for any value axes in the chart.



Axis Scale Dialog

The following options are provided:

Automatic: This turns on automatic scaling for the axis. This is the default option.

Manual: This turns on manual scaling, allowing you to set the axis scale to your preference.

Maximum: This is the highest value on the axis scale.

Minimum: This is the lowest value on the axis scale.

Unit: This is the step interval between successive labels.

Add Padding: This will raise the highest value of the axis to create a cushion between the max value of the data and the top of the chart.

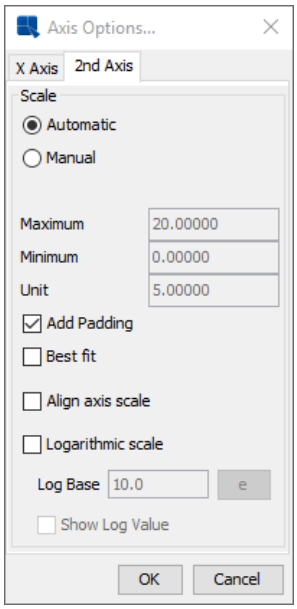
Best fit: This will automatically place the origin of the chart based on the minimum and maximum values of the data.

Origin: This allows you to specify where the X and Y axes should intersect. This is usually set to zero.

Logarithmic Scale: This option creates a logarithmic scale for the given axis. It's valid only if the axis in question contains positive values.

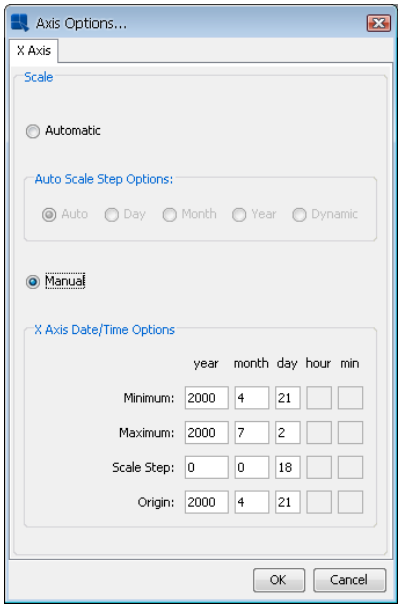
- **Log Base:** This allows you to specify the base for the log value.
- **Show Log Value:** This specifies whether to show log values in the axis labels or not.

The axis scale dialog will have a tab for each axis in the chart. There's a unique option available for secondary axes which allows you to align the axis scale. It will apply all options from the primary axis to the secondary axis, giving both axes the exact same scale.



Axis Scale Dialog for Secondary Axis

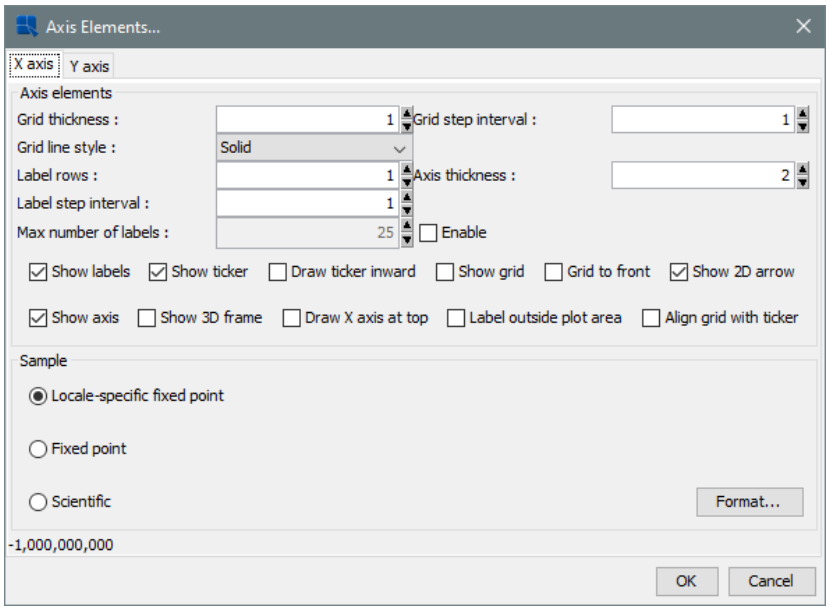
The axis scale dialog is different for a Gantt chart. You can still select the *Automatic* option to have the scale configured automatically or the *Manual* option to set the scale manually. The *Automatic* axis scale has a few Auto Scale Step Options: Auto, Day, Month, Year, and Dynamic. The *auto* option always finds a best fit. The *dynamic* option also finds a best fit, but unlike the *auto* option, it uses standard step intervals only (for example: 1 month, 1 year etc...). When the *Manual* axis scale is chosen, the *Maximum*, *Minimum*, *Unit* and *Origin* are replaced with *Maximum Date*, *Minimum Date*, *Scale Step* and *Origin Date* respectively and these new settings take in a Date/Time (represented by yyyy, MM, dd, hh, mm).



Axis Scale Dialog for Gantt Chart

5.2.4.7.2. Axis Elements

The appearance properties of the axes and the axis labels are controlled through the axis elements dialog. To invoke the axis elements dialog, select Format → Axis Elements or click the *Format Value Elements* button on the toolbar. This will bring up the following dialog, allowing you to customize elements in all chart axes. You can also customize the appearance of dial and pie chart labels from this dialog.



Axis Elements Dialog

A tab will appear in this dialog for each axis in the chart. The dialog allows you to perform the following options. Note that some options are only available for certain chart types, certain data types, and on certain axes.

Grid thickness: This allows you to specify the thickness of any grid lines along the axis.

Grid step interval: This option allows you to set the grid step interval for any grid lines along the axis.

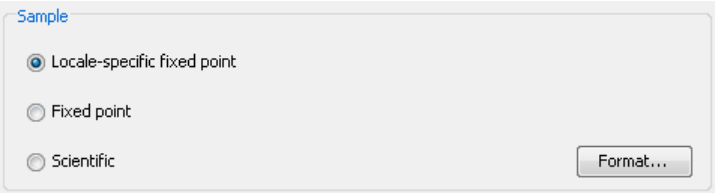
Grid line style:	This option allows you to select the grid line style (solid, dotted, dash).
Label rows:	This option allows labels to be displayed in alternating rows. This can prevent overlapping. This option is only available for X-axis.
Axis thickness:	This option allows you to set the thickness of the axis in pixels. Note that this setting is applied to all axes in the chart.
Label step interval:	This option allows you to set the label step interval for the data. For example, setting this to 2 will draw the label for every other data point in the chart.
Label interval unit:	This option allows you to select the unit to be used when sorting and representing time-based data (date, time, or timestamp). Selecting tickers will use the data as it is read by Chart Designer. You can also select Dynamic for time-based data and that will choose an appropriate scale (depending on number of data points and range of data).
Max number of labels:	This option allows you to select the maximum number of labels and tickers displayed on the axis. If the number of labels exceeds the max count, the label step will be re-calculated.
Ascending/Descending:	This option allows you to order and filter time-based data. You can sort the data in ascending or descending order, as well as specify the starting (or ending) point for the data.
Show labels:	This option allows you to remove or display the labels for each axis.
Show ticker:	This option allows you to remove or display the axis tickers.
Draw ticker inward:	This option will draw the axis tickers inside the plot area instead of outside (default).
Show sub-tickers:	This feature is only available for the value axis when the axis scale is set to logarithmic with a log base of 10. This feature will draw interval tickers (non-uniform) between the points on the value axis.
Show grid:	This option allows you to remove or display the grid for each axis.
Grid to front:	This option allows you to draw the grid lines on top of the data elements in the chart. By default the data points are drawn on top of the grid.
Show 2D arrow:	This option allows you to remove or display the arrowhead at the end of the axis. Note that this option applies to all chart axes and it is only available for two-dimensional charts.
Show axis:	This option allows you to remove or display the axis (for two-dimensional charts) or the wall (for three-dimensional charts).
Show 3D frame:	This option allows you to remove or display a frame around the chart. Note that this option applies to all chart axes and it is only available for three-dimensional charts.
Label outside plot area:	This option sets the labels to be placed outside of the plot area, irrespective of where the axis is. This feature can be useful for category axis labels if you're plotting data with both positive and negative values.
Align grid with ticker:	This option aligns the grid line with the ticker instead of placing it between tickers. This places the ticker and the corresponding grid line along the same line. This option only applies to the category axis of column-type charts.
Swap Y-axis position:	This option will swap the primary and secondary value axes. This option can only be found under the 2nd Axis tab.

Draw X-axis at top: This option allows the X-axis to be positioned at the top of the chart instead of the default bottom position. This option is available for two-dimensional column, bar, scatter, high-low, HLCO, bubble, and Gantt charts.

5.2.4.7.2.1. Axis Label Formatting

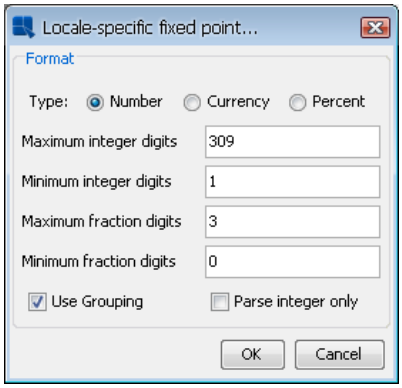
The axis elements dialog also allows you to format the appearance of the axis labels, depending on what type of data is plotted on the axis. The *Format Options* portion of the dialog contains the label formatting dialog.

Formatting Numeric Data For numeric data there are three primary options for display formatting: locale-specific fixed point, fixed point, and scientific. Additional options will be displayed if you click on the *Format* button.



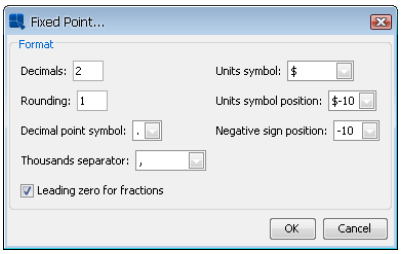
Numeric Data Format Options

Locale-Specific Fixed Point This will change the format of the data depending on the locale in which it is being viewed. Additional formatting for this option allows you to specify whether the data should be displayed as a number, currency, or percentage. In addition, you can set the maximum and minimum number of integer digits and fraction digits. Other display attributes will vary depending on locale.



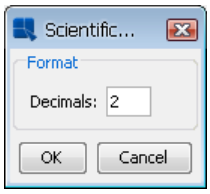
Locale-Specific Formatting Options

Fixed Point: This will keep the data format consistent, regardless of locale. Additional formatting for this option allows you to set the number of decimals, rounding for digit number, unit symbols, negative sign position, decimal and thousands separator, and enable leading zeros for fractions



Fixed Point Formatting Options

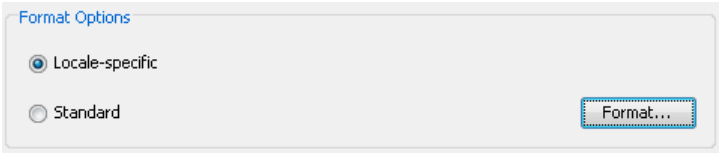
Scientific: This will display the data in scientific notation. Additional formatting for this option allows you to set the number of decimals.



Scientific Formatting Options

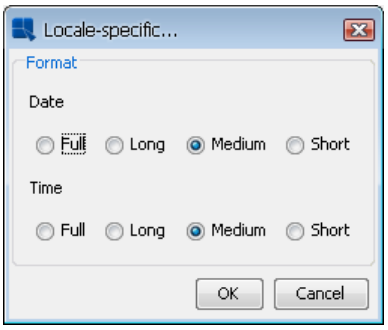
Formatting Date/Time Data

For date/time data there are two primary options for display formatting: locale specific and standard. Additional options will be displayed if you click on the *Format* button. The available options will vary depending on the nature of your data. Date, time, and timestamp data will bring up date, time, and date & time options respectively.



Date/Time Data Format Options

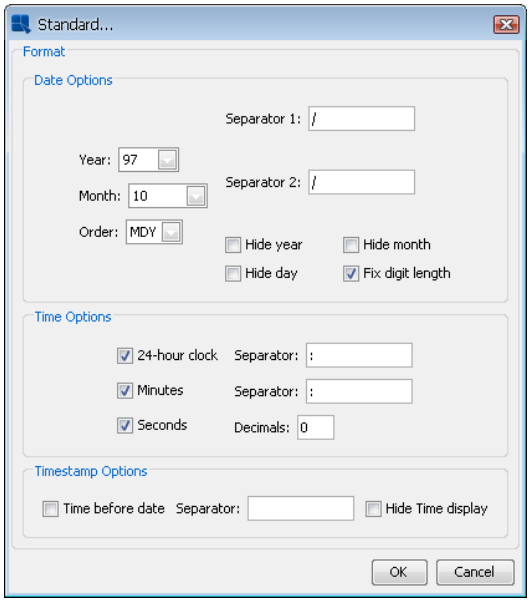
Locale-Specific: This will change the format of the data depending on the locale in which it is being viewed. Additional formatting for this option allows you to select full, long, medium, or short notations for date and time information. Other display attributes will vary depending on locale.



Locale-Specific Formatting Options

Standard:

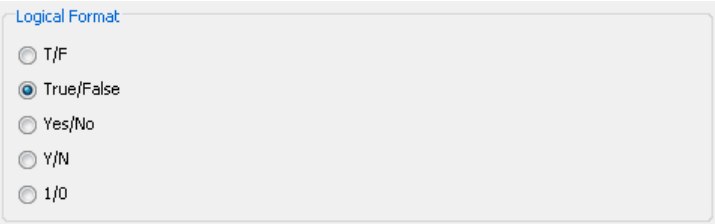
This will keep the data format consistent, regardless of locale. Additional formatting for this option allows you to select year and month displays, as well as the order in which month, day, and year information is presented. You can also select the characters that you want to be used as separators. Time options allow you to display hours, minutes, and/or seconds and also to select the separators between them. For timestamp data, you can select to display the time before or after the date, as well as the separator to be used between them.



Standard Formatting Options

Formatting Logical Data:

There are five options available for displaying logical or Boolean data: T/F, True/False, Yes/No, Y/N, and 1/0.



Logical Data Formatting Options

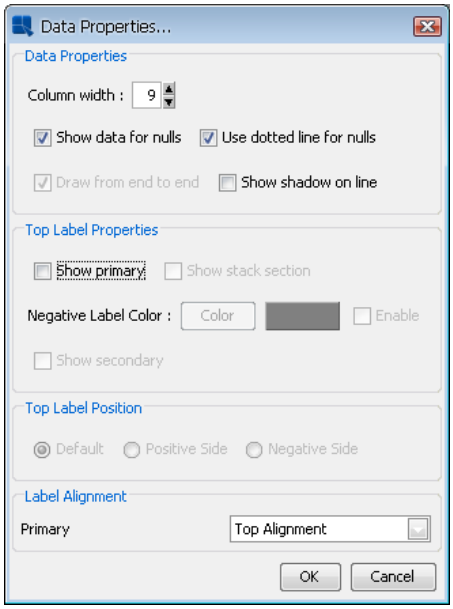
Any changes you make to the data formatting will take effect after you click on the *OK* button in the axis elements dialog. Note that there are no additional formatting options for string data.

5.2.4.8. Formatting Plot/Data Elements

ChartDesigner provides a number of ways to customize and configure the way data points are drawn and annotated on the chart, as well as the chart plot itself.

5.2.4.8.1. Data Properties

Many of the data display options are controlled through the data properties dialog. From this dialog you can control the size of bars/columns, set display options for null values, and specify options for data labels. To invoke the data properties dialog, select Format → Data Properties or click the  *Data Properties* button on the toolbar. This will bring up the following dialog:



Data Properties Dialog

This data properties dialog contains the following options:

Column width: This specifies the ratio of the bar/column width with respect to the gap between successive bars in the chart. Each unit represents $1/10_{th}$ of the space between data points. Therefore, entering **9** would leave 10% of the space between data points blank, while **10** would eliminate all space between bars/columns. This option only pertains to two-dimensional bar, column, stack bar, stack column, high-low, HLCO, and Gantt charts. To control the column thickness in three-dimensional charts, you can use the thickness of shape slider in the navigation panel.

Show data for nulls: This option will connect lines when null data is present. For example, if you have three points and the value of point 1 is 4, point 2 is null, and point 3

is 6, then a line will be drawn from 4 to 6. This option is only available for line charts and other two-dimensional charts with lines. All other chart types will not plot null data.

Use dotted lines for nulls: You can use this option to replace the full line with a dotted line. Like the show data for nulls option, this property is only available for lines.

Draw from end to end: This option allows you to draw two-dimensional line and area charts across the entire plot area, rather than offsetting to the first and last data points on the chart.

Show shadow on line: This option allows you to use shading on two-dimensional lines. However, the line must be thicker more than one pixel.

Show primary: This option will display data top labels for the primary values in the chart.

Show stack section: This option will display individual labels for each stack section for stack bar, stack column, and stack area charts.

Negative Label Color: This option will display the top labels with a value smaller than that of the origin in a different color that can be selected using the *Color* button after enabling the feature.

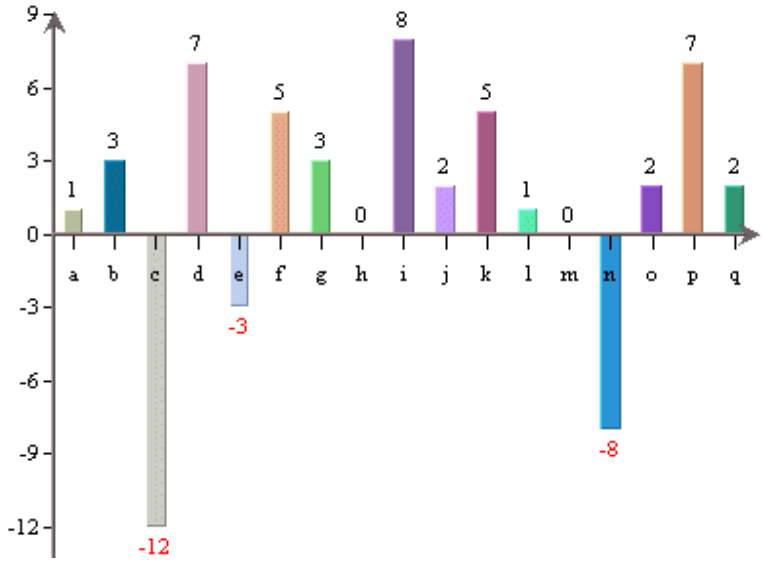


Chart with Colored Negative Top Labels

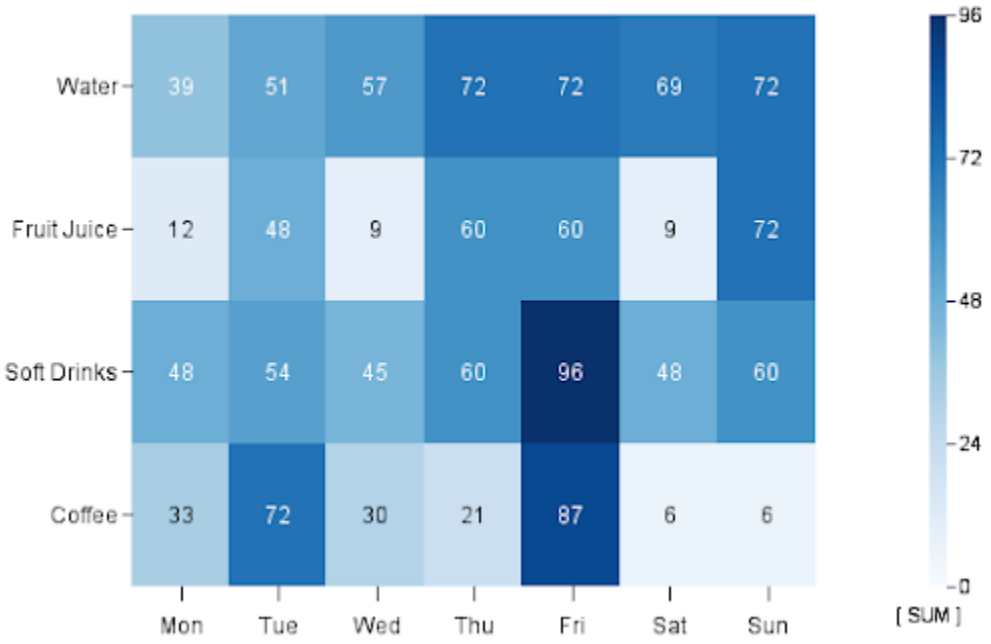
Show secondary: This option will display data top labels for the secondary values in the chart.

Top label position: This option allows you to specify where the data top labels should be drawn. By default, they are drawn above data points if they are positive and below data points if they are negative. Other options allows you to draw the labels to either positive or negative side.

Label Alignment: This option allows you to set the alignment for the data top labels. You can draw them at the top, bottom, or middle of the data points. In addition, you can select to draw the label inside the data point at the top or bottom. An additional option stack charts offers you to set the alignment for stack section labels.

5.2.4.8.1.1. Data Properties for Heatmap Chart

For heatmap charts, the Data Properties dialog allows users to configure the values displayed in the cells in the matrix.



Data Properties for Heatmap Chart

5.2.4.8.2. Date/Time Based Zooming

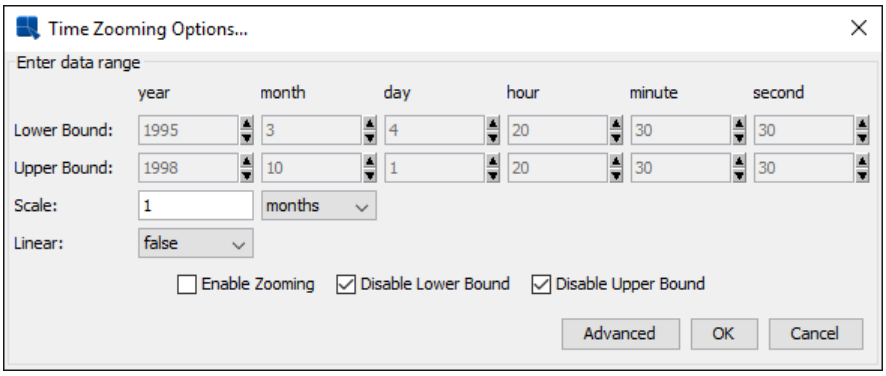
For charts displaying date or time data on the category axis, ChartDesigner provides a unique feature allowing users to perform date/time based zooming. Using this feature, you can group the category elements into user-defined intervals and aggregate the points in each group. You can also filter the data by specifying upper and lower bounds for the results.

For example, suppose your data contains daily sales volume for the past two years. Using zooming, you could aggregate the data to look at average volume per month, quarter, or year. Using the upper and lower bounds, you could narrow the range to look at weekly sales volume within a specific quarter.

Zooming is available for all chart types except scatter, surface, box, dial, polar, radar, bubble, and Gantt.

5.2.4.8.2.1. Adding Zooming

When you create a new chart with date, time, or timestamp data in the category axis, you can specify zooming options by selecting Format → Time Zooming Options.



Zoom Options Dialog

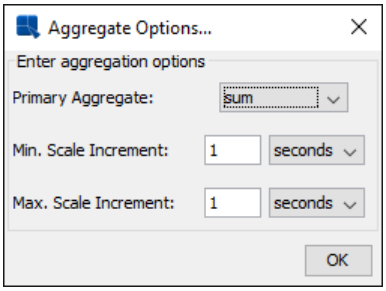
This dialog allows you to specify a lower and upper bound for the data, as well as the interval by which you would like to group the data. The scale specified here must be within the maximum and minimum scale specified in the aggregation options dialog.

This dialog also allows you to preserve a linear scale for the chart. By setting the *Linear* option to true, the chart will always display points for the grouped intervals, even if there is no data associated with a particular group. For

example, say again that you are measuring sales volume over a three month period - April, May, and June. If the input data has no records for May and you set the *Linear* option to true, a point will be drawn for May with a value of zero. If you set the *Linear* option to false, the data point for April will be immediately followed by June.

You can disable/enable zooming, as well as the lower and upper bound restrictions by using the checkboxes at the bottom of the dialog.

If you enable zooming (if you check *Enable Zooming* option), the dialog *Aggregate Options* will appear, prompting you to specify aggregation options for the grouped data points.



Aggregation Options Dialog

In this dialog, you can specify the Primary Aggregation, as well as the maximum and minimum scale increments that can be used when zooming the data. After you have specified your desired options, click on the *OK* button to return to zooming options.

Once you have finished specifying all the options, click on the *OK* button and the zooming will be applied to the chart.

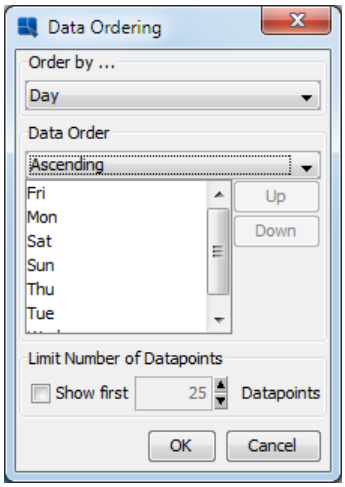
5.2.4.8.2.2. Zooming In Chart Viewer

When deploying charts using Chart Viewer, end users can perform dynamic zooming. To perform a time-series zoom in the Chart Viewer, **Ctrl**+**Click** on a point on the chart and drag it to another point in the chart. This will automatically zoom in based on the lower and upper bounds selected using the mouse. The aggregation is performed according to the options that were set during design time. You can undo the zoom by **Ctrl**+**Right-Click**.

The scale interval is chosen automatically, depending on the data and chosen bounds (as long as minimum 2 data points can be shown). The scale interval can also be changed in the Chart Viewer by pressing **Alt**+**Z**. This will bring up a dialog allowing the user to change the zoom settings.

5.2.4.8.3. Data Ordering

ChartDesigner allows you to change the order of the category and series elements. To modify the ordering, select Data → Ordering or click the  *Change Data Ordering* button on the toolbar. This will bring up the following dialog:



Data Ordering Dialog

There is an *Order By* list which contains category element, series element, and an option marked *VALUES*. You have the following options for the category and series elements:

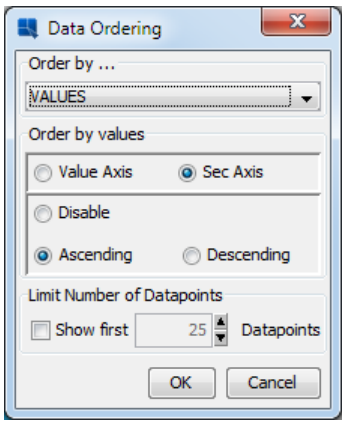
DataSource Order: This option turns the ordering off. The categories/series order will depend on the data source only and will not be altered by the ERES at all.

Ascending: This option will arrange the categories/series elements in ascending order. For example, if the category elements are strings, they will be arranged alphabetically.

Descending: This option will arrange the categories/series elements in descending order.

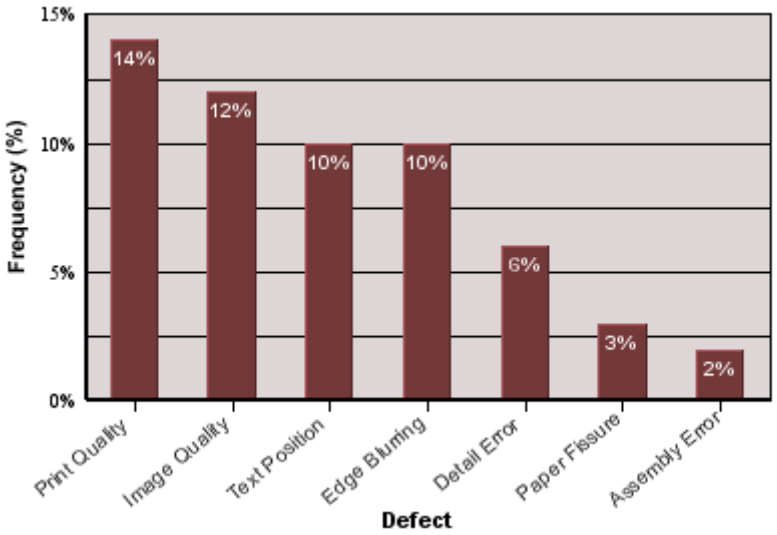
Customize: This option allows you to customize the categories/series order. To customize the order, select an item from the list of Categories/Series items and then move it upwards or downwards in the list by clicking on the *Up* or *Down* button (the buttons are inactive until you select the *Customize* option).

You can also sort the category elements based on their corresponding values. To do this, select the *VALUES* option in the data ordering dialog.



Value Data Ordering Dialog

If you choose the *VALUES* option, the entire dialog changes. From this dialog, you can specify to sort the category elements based on their corresponding values in the value, or secondary value axis. You can also specify whether to sort them in ascending or descending order. This type of ordering is called a *Pareto* chart and is often used in process control applications.



Pareto Chart

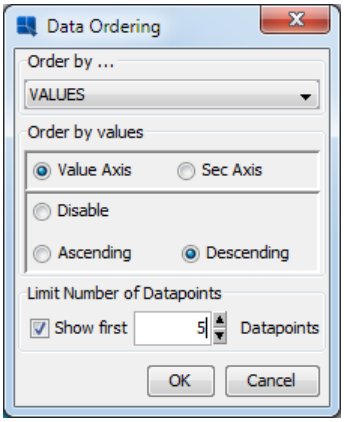
Please note that any sorting set will be re-applied if the chart is refreshed and/or if the data changes.

5.2.4.8.3.1. Top/Bottom N Charts

Sometimes you want to plot only a few highest or lowest values. To do that, you can use the *Top/Bottom N* function.

To enable this feature, choose Data → Ordering, or click the  *Change Data Ordering* button on the toolbar.

If the chart doesn't have any data series, the *Ordering* dialog will show the *Limit Number Of DataPoints* option.



This option can only be enabled if you sort the chart by categories or values, or by the values in ascending or descending order. If you have such chart, you can enable this function by selecting the *Show first* option. Then you can specify the maximum number of items that will be shown in the chart. If the data source returns more items than you specify in this option, excessive items will not be shown in the chart as if they didn't exist.

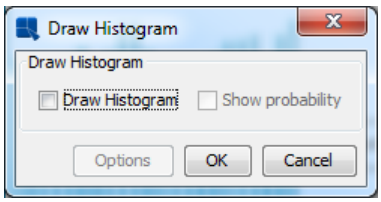
5.2.4.8.4. Histograms

Histogram is a useful analysis tool that allows you to track how often events occur and when and how a set of data falls into specific ranges. ChartDesigner allows you to plot histograms based on the category elements in a chart. You can plot histograms for all category data types except time-based data (date, time, or timestamp).

Histograms are calculated by counting data points or instances of each category element. For numeric categories, you can further specify upper and lower bounds, as well as the number of bins or bin width to create ranges for the frequency counts.

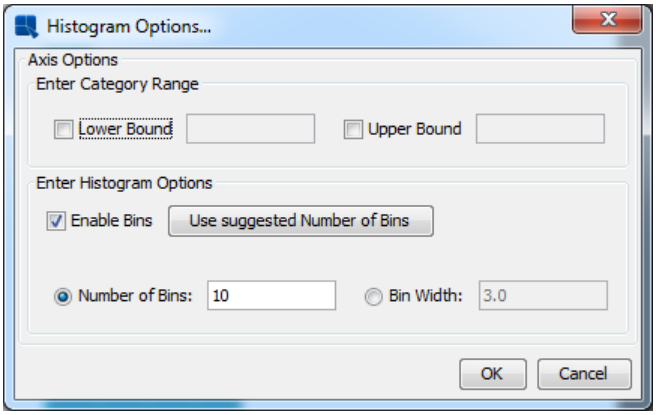
To create a histogram, you must start with a 2D column chart, bar chart, line chart or area chart. In the Data Mapping dialog, make sure DataSeries is set to *None* and the field you want to plot in the histogram is set in the Category axis.

Once you click *Done* in the Data Mapping dialog and the chart is shown on the canvas, select Format → Histogram Options. A dialog will appear allowing you to select a histogram plot. By default, the histogram is displayed as frequency count. You can choose to change it to display probability by checking the *Show probability* check box.



Select Histogram Dialog

If the values in the Category axis is numeric, you will see that the *Options* button is enabled. When you click Options button, another dialog will appear, allowing you to specify options for the histogram plot.



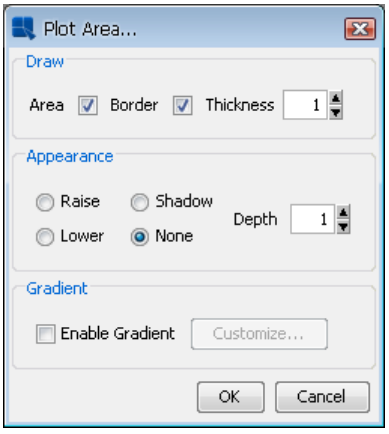
Histogram Options Dialog

From this dialog, you can set a lower or upper bound for that data being plotted. When you place bound restrictions, the histogram will not count data that falls outside of the range specified by the upper and lower bounds. If you select *Enable Bins*, you can specify the number of bins or set width of each bin. The default number of bins is 10. You can also click *Use suggested Number of Bins* if you wish to use a value calculated by the system. If *Enable Bins* is deselected, the frequency count will be performed on each Category value instead of a range of values for a bin.

If you enable scale, you can specify the number of bins or set width of each bin.

5.2.4.8.5. Formatting Plot Area

The plot area is the plane on which the data points are drawn for two-dimensional charts. You can customize the appearance of the plot area by selecting Format → Plot Area. Assuming the current chart is a two-dimensional chart, the following dialog will appear.



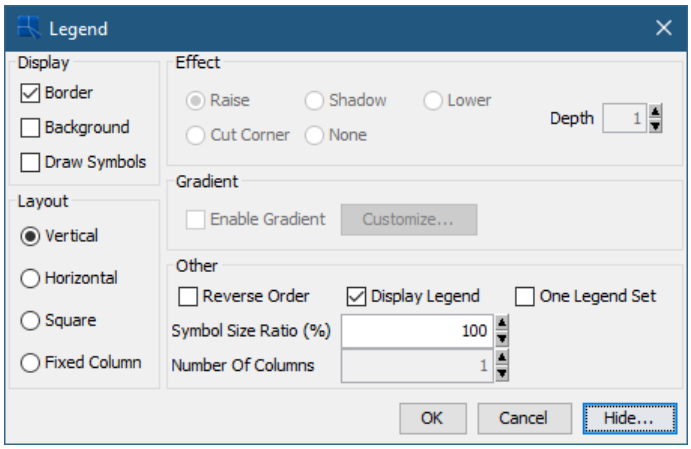
Plot Area Dialog

This dialog allows you to draw a border around the plot area, or fill it with a background color. If you fill the area, you can also specify certain 3D effects like raising, lowering or shadow.

On this dialog, you can also set up gradient background for the plot area. The gradient settings are the same as in the *Rendering options* described in the Section 5.2.4.1.3 - Format Menu.

5.2.4.8.6. Formatting Chart Legend

You can control and modify the display of the chart legend either by selecting Format → Legend, or by clicking on the  *Format Legend* button on the toolbar, or by selecting *Legend properties* from legend pop-up menu (which pops up when you right-click on the legend). This will bring up the following dialog, allowing you to customize the legend properties.



Format Legend Dialog

The dialog contains the following options:

Display: These options allow you to turn on or off the legend border and background. This also allows you to display the point symbols instead of lines or blocks in the legend.

Effect: This allows you to add a 3D effect to the legend. You can raise it, lower it, or draw a shadow. In addition to the 3D effects, you can also display the legend with cut corners.

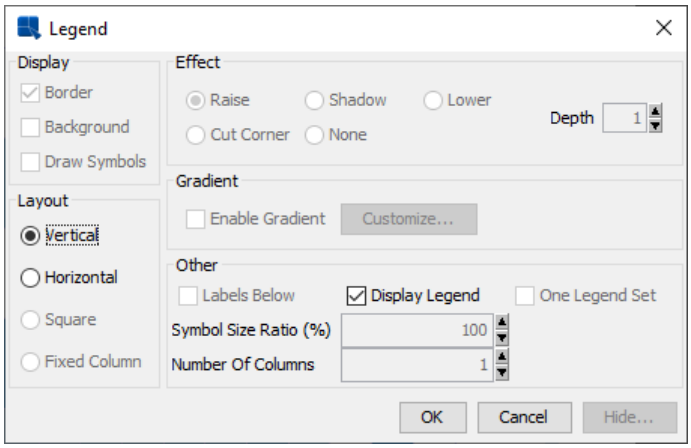
Layout: This allows you to change the legend from vertical, horizontal, square, or fixed column layout.

Gradient: Allows you to configure gradient for the legend background. The gradient settings are the same as in the *Rendering options* described in the Section 5.2.4.1.3 - Format Menu.

Other: This allows you to choose whether or not to display the legend, or to draw the legend in reverse order. You can set the fixed number of columns in legend in the *Number of columns* field. This field will be active only if you choose the *Fixed columns* layout option in the *Layout* section. You can set trend/control lines and chart data legend drawn as *One Legend Set*. You can also change the size of the symbols in the legend.

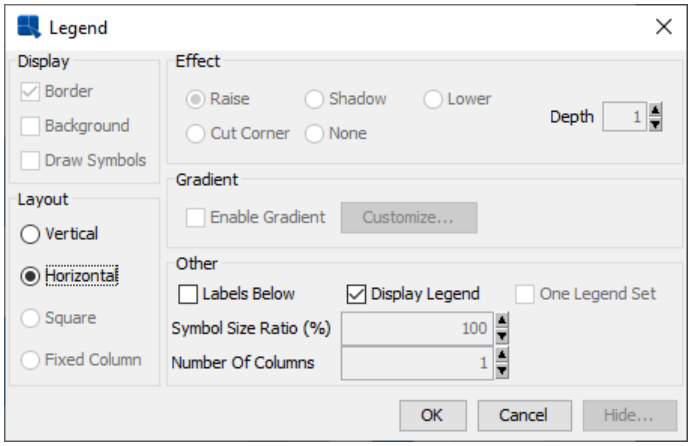
Additionally, you can remove specific category/series elements from the legend by clicking on the *Hide* button. This will bring up a list of the legend items, where you can select which elements you would like to hide.

The legend for heatmap charts differs from those of other charts, and some legend options in the *Legend* dialog are not applicable to heatmap charts. By default, the legend for heatmap charts is displayed vertically.



Heatmap Vertical Legend

In the horizontal layout, you can choose to display tickers and labels below the legend by selecting the *Labels Below* option.

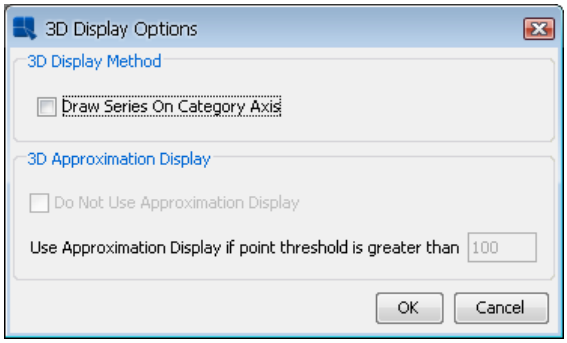


Heatmap Horizontal Legend

5.2.4.8.7. 3D Display Options

ChartDesigner renders three-dimensional charts in true 3D, allowing light source modification, panning, zooming, and rotation. However, 3D rendering can be very memory and CPU intensive. When charts have a lot of data points (like 3D scatter and surface charts), it's possible to run out of memory when generating the chart. To solve this problem, a rendering approximation feature is provided. Using this algorithm the chart is not rendered perfectly, but it's usually acceptable when a lot of points have to be shown.

By default, approximation is turned on at a threshold value of 100 points. This means that if a 3D chart has more than 100 data points, the approximation will be used. You can turn this feature off, or change the threshold value by selecting Format → 3D Display Options. This will bring up the following dialog.

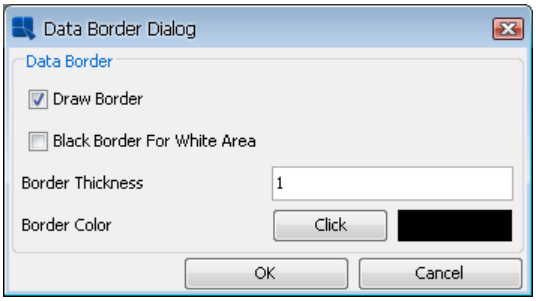


3D Display Options Dialog

The two options for 3D approximation allows you to turn on/off the approximation and set the threshold value. The other option in this dialog allows you to draw the series in-line (the same option is in navigation panel).

5.2.4.8.8. Data Border

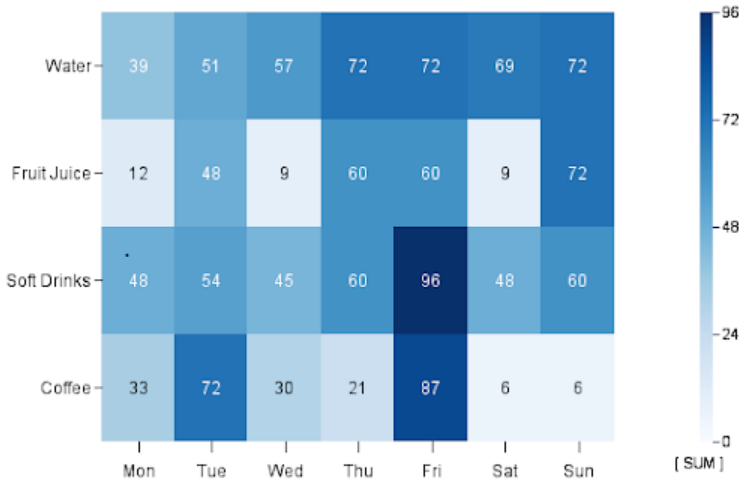
For column, bar, stack column, stack bar and HLCO charts, ChartDesigner allows you to configure a border around the columns. To set the border option, select Format → Data Border. This will bring up a dialog allowing you to set border options.



IData Border Dialog

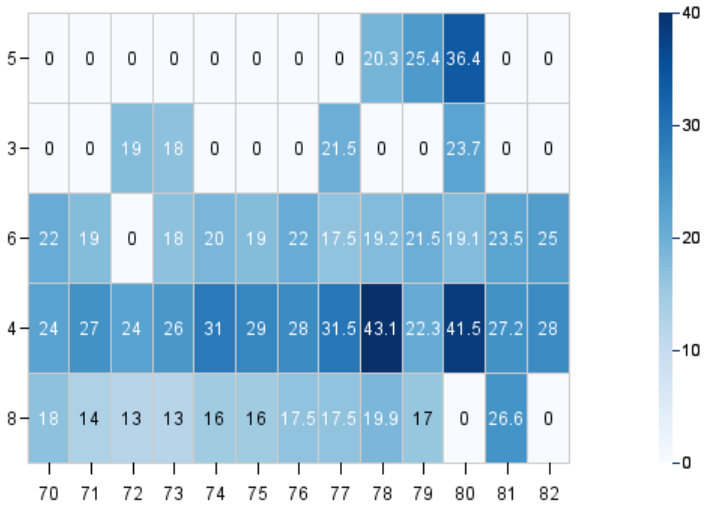
The first option allows you to turn on/off the data border. The second option allows you to set a black border for any white areas in the chart. Please note that the border is black only if the first option is unchecked and will only appear around white areas in the chart. The third and fourth options allows you to set border thickness and border color. If you click on the *Click* button, a new dialog will appear allowing you to select or enter a new color.

For heatmap charts, you can configure a border around the data rectangles. This is particularly useful when there are many light-colored cells.



Heatmap Data Border

The following is an example heatmap chart with border thickness set to one.

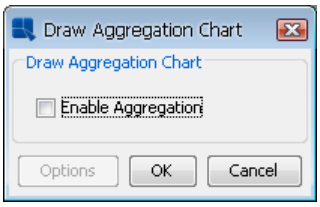


Heatmap chart with light gray colored border

5.2.4.8.9. Aggregation

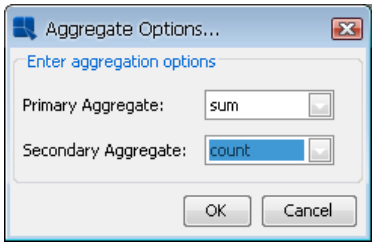
ChartDesigner allows you to aggregate data if there is more than one data point associated to a given category (and its series and/or stack, if a series and/or stack is present). This allows for a broader look at the data rather than just a single data point (out of many).

To aggregate the data, select Format → Aggregation Options. A dialog will appear allowing you to enable aggregation.



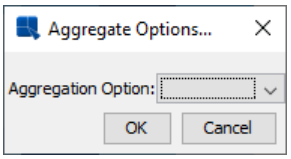
Select Aggregation Dialog

When you select *Enable Aggregation*, a second dialog box will appear, asking you which type of aggregation should be applied. You can choose from minimum, maximum, average, sum, count, first, last, sumsquare, variance, stddev, and countdistinct for the aggregates. You can specify a primary aggregate (aggregate applied to the column mapped to the primary axis) as well a secondary aggregate (aggregate applied to the column mapped to the secondary axis), if a secondary axis exists.



Aggregate Options Dialog

For heatmap charts, the aggregation option can be used when there are multiple data values at a single (x, y) point.



Aggregate Options

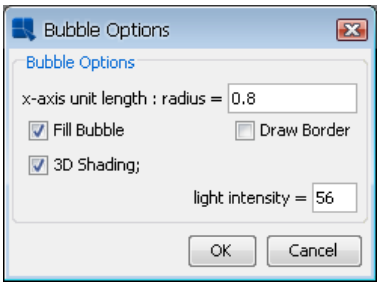
By default, no aggregation is applied to heatmap charts, and the first data value is always selected if there are multiple values at the same (x, y) point.

5.2.4.9. Chart-Specific Options

There are a number of formatting options that are unique to certain chart types. These options can be modified by selecting Format → Chart Options, or clicking the  *Chart Options* button on the toolbar. This will bring up a dialog that varies depending on the type of the current chart. Some chart types have no additional options.

5.2.4.9.1. Bubble Charts

For bubble charts, the following dialog is displayed:



Bubble Options Dialog

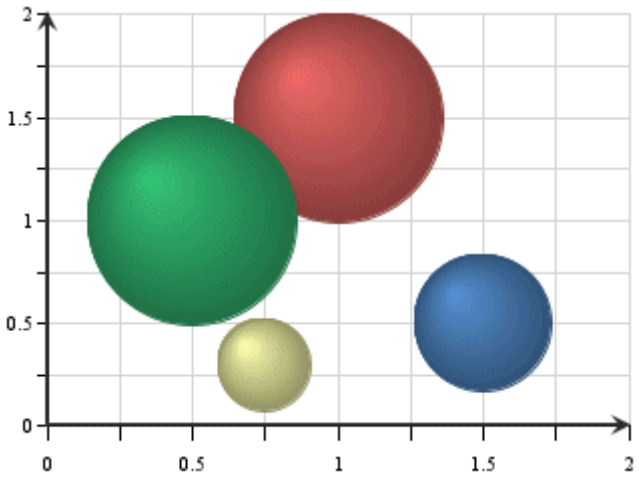
The following options are available for bubble charts:

x-axis unit length: radius — This option specifies the ratio of X-axis unit length to the radius of the bubble.

Fill Bubble: — This option allows you to specify whether or not to fill the bubble area.

Draw Border: — This option allows you to specify whether or not to draw a border around the bubbles.

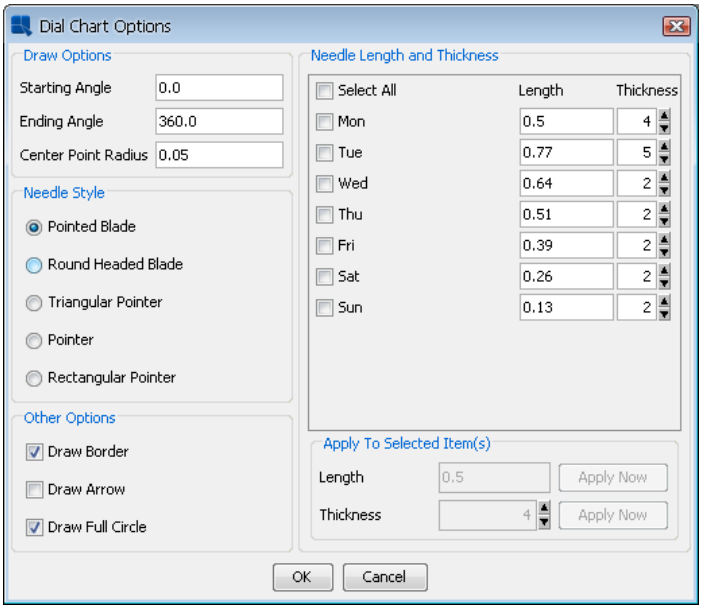
3D Shading: — This option allows you to add 3D shading to the bubbles. You can also specify the light intensity for the shading.



Bubble Chart with 3D Shading

5.2.4.9.2. Dial Charts

For dial charts, the following dialog is displayed.



Dial Options Dialog

The following options are available for dial charts:

Starting Angle: This option specifies the angle where the first axis label is to be set. This property also determines where the border and dial area will start if the *Draw Full Circle* option is unchecked. The angle is represented in degrees and is 0 by default. Assuming the dial chart is a clock face, 0 degrees is 12 o'clock.

Ending Angle: This option specifies the angle where the last axis label should be set. This property also determines where the border and dial area will end if the *Draw Full Circle* option is unchecked. The angle is represented in degrees and is 360 by default. Hence by default the labels (and data points) encompass the entire circumference of the dial.

Center Point Radius: This option specifies the radius for an inner circle, which starts from the center of the dial chart. The radius is specified as a ratio to the radius of the

dial plot. Hence, a value of 1 will make the inner circle encompass the entire dial. If the *Draw Full Circle* option is unchecked, only the portion of the center point that is within the starting and ending angles will be shown.



Dial Chart with Center Point Radius

Needle Style: This option specifies the type of needle to draw.

The Pointed Blade is a smooth line with a thick base. It becomes slightly thinner as it extends outwards, but finishes with a very sharp pointed tip.

The Round Headed Blade is similar to the Pointed Blade except for a round tip.

The Triangular Pointer is sharp and resembles a very thin triangle.



The Pointer is a step ladder pointer with three segments.

The Rectangular Pointer (as shown above on the picture) is a simple straight needle.

Default needle is Pointed Blade.

Draw Border: This option specifies whether or not to draw a border around the dial.

Draw Arrow: This option specifies whether or not to draw arrowheads at the end of the dial hands.

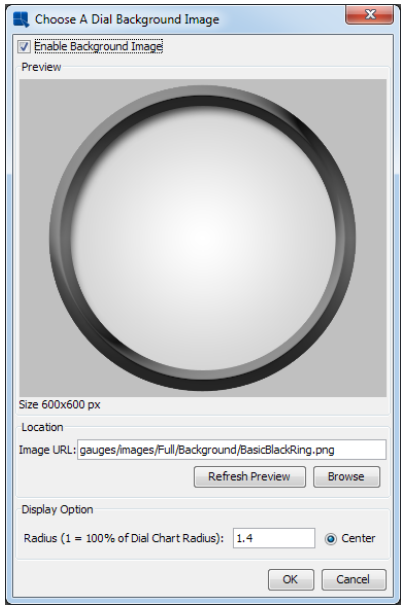
Draw Full Circle: This option specifies whether to draw the dial as a complete circle (360 degrees) or only draw the portion of the circle determined by the starting and ending angles.

Needle Length and Thickness: This option specifies the length and thickness of the needle. The needle length is measured from the center of the dial. The range is from 0 (center of the dial) to 1 (the end of the dial). The thickness determines the width of the needle, larger values results in a wider pointer. When creating a chart, the needle length is randomly generated. Default thickness is 2. Each category element is represented by a different needle. You can either change properties individually or change multiple categories at once. To change the property of each needle individually, directly change the values to the right of the category. To make changes to multiple needles, check each category or check the *Select All* option, set the properties in the lower right corner

and then click on the *Apply Now* button for each property changed. This will modify all checked categories to new values.

5.2.4.9.2.1. Gauge Images

Dial charts have an additional option to display a foreground or background image for the dial plot area. To add a foreground or background image, select Insert → Dial Foreground... or Insert → Dial Background.... These options are only enabled for dial charts.

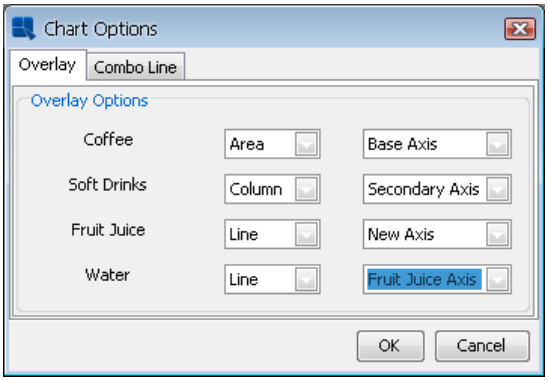


Dial Chart Background Image Dialog

Selecting an image works in the same way as the background image dialog, see Section 5.2.4.5.1.1 - Background Images. In the dial chart image dialog, there is also an option to specify the radius of the image. Specifying 1 for the radius will make the image the same width and height as the plot area. Increasing or decreasing this value will enlarge or shrink the images.

5.2.4.9.3. Overlay Charts

For overlay charts, the following dialog is displayed:

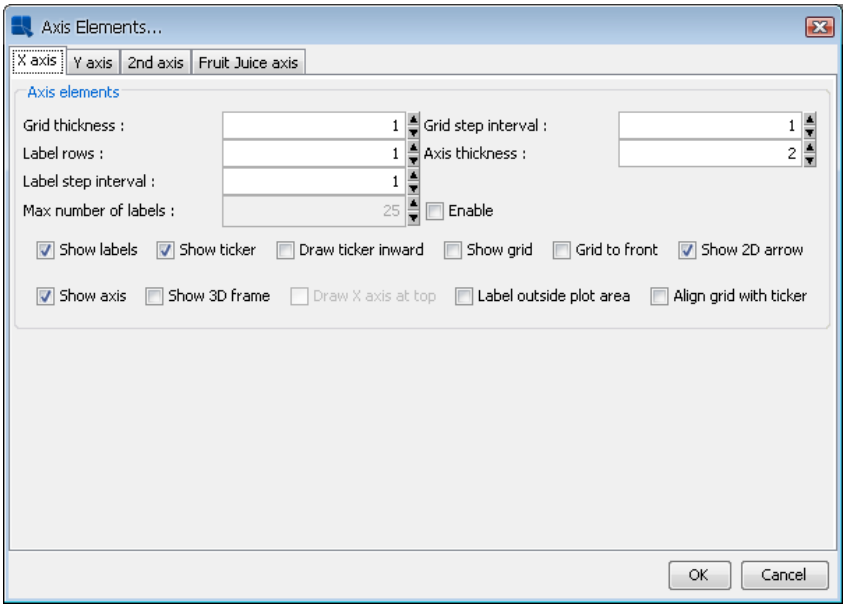


Overlay Options Dialog

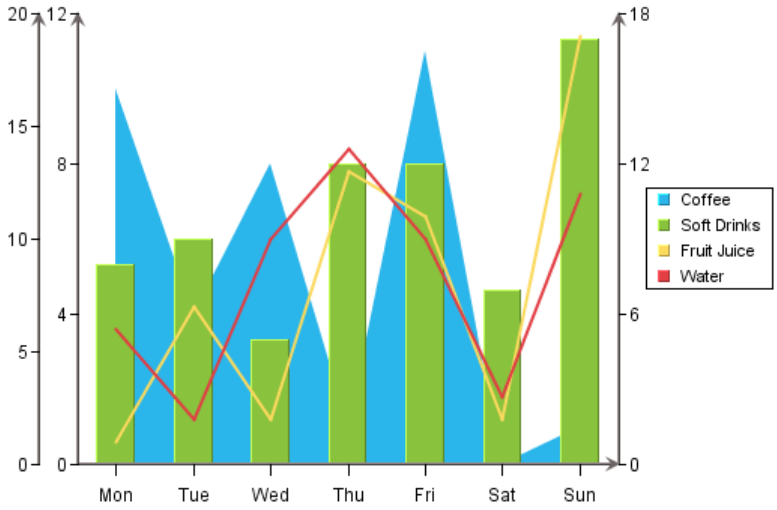
From this dialog, you can specify which chart type you would like to use for each element of the data series. Available chart types for the series elements in an overlay chart are column, area, and line. You can also choose not to display certain series elements.

From this dialog, you can also specify which axis you would like to use to plot a series element. You can place elements on the primary or secondary axes, or you can create new value axes for the series elements. To create a new value axis, select *New Axis* from the drop-down menu. Once you specify to use a new axis, a new option will be added to the drop-down menus for the other data series elements, allowing them to be drawn on the same axis

that you previously specified. Using a variety of axes allows you to precisely tune the scale that the different series elements use. Each of these axes will have its own tab in the Axis Elements window, where you can change the label step interval for each axis independently of the others.

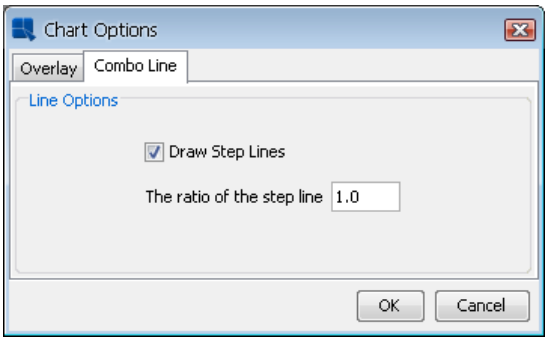


Axis Options Dialog for Multiple Value Axes



Overlay Chart with Multiple Value Axes

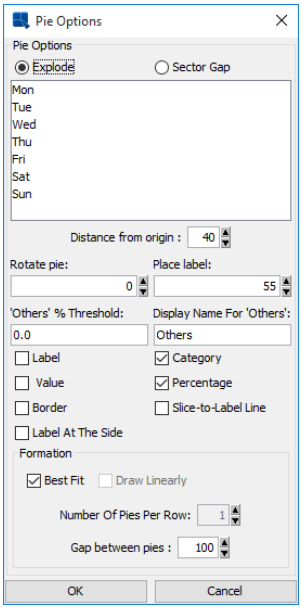
If one of the chart types used by the overlay chart is a line, you can specify whether or not to draw the line as a step line using the *Combo Line* tab. For more about step lines see Section 5.2.4.9.5 - Line Charts:



Combo Line Options for Overlay Charts

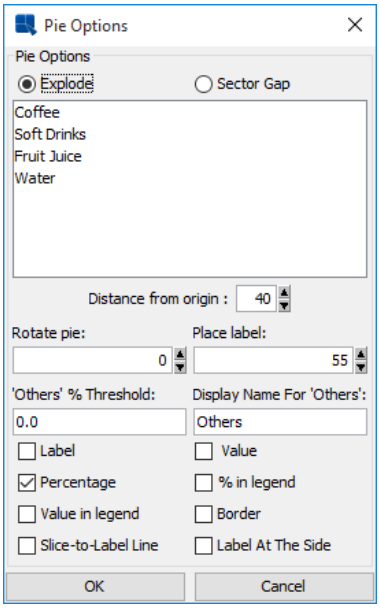
5.2.4.9.4. Pie Charts

For pie charts, the following dialog is displayed. (Note that different options will appear/disappear depending on whether the chart has a series, and if it's a 2D or 3D chart.) Here are the options available for a 2D chart with series:



Pie Options Dialog (With Series)

Here are the options available for a 2D chart without series. Notice that the formation options are removed and the *% in legend* and *Value in legend* options are added:



Pie Options Dialog (Without Series)

The following options are available for pie charts:

Explode: This option allows you to pick one or more category/series elements whose sections are to be drawn at a certain distance away from the center of the pie.

Sector Gap: This option allows you to pick one or more category/series elements whose sections are drawn at a certain distance away from the center of the pie and still maintain the same distance between pie slices and a circular boundary.

Distance from origin: This option allows you to specify how far the exploded/sector gap sections are to be drawn away from the center. This number, represented as a percentage of the radius, indicates the distance between the center and the tip of the pie slice to be exported.

Rotate pie: This option allows you to specify the number of degrees that the chart should be rotated in a clockwise direction. Available values are between 0 and 360

Place label: This option indicates the distance of the labels from the center of the pie. The position of an individual label can also be adjusted by dragging the text.

“Others” % Threshold: This feature is useful for pie charts that have a large number of small categories. Rather than draw a slice for each category, users can select a threshold value. Any category whose percentage of the value column is less than the threshold value will be lumped into an “Others” slice.

Display Name for “Others”: This option allows you to set the display name for the “Others” slice that is created for categories that fall below the supplied threshold value. This label will appear in the legend, and/or for the slice label.

Label: This option determines whether a category/series label should be drawn for each pie slice. By default, these only appear as legend items. Note that the label will not appear if the data for the slice is 0 or null.

Value: This option allows you to specify whether to display the actual value of each pie slice. Note that the value will not appear if the data for the slice is 0 or null.

Category: This option allows you to specify whether to display the category of each pie slice.

Percentage: This option allows you to display the percentage for each pie slice. The percentages are calculated by dividing the value of each section by the sum of all the values. Note that the percentage will not appear if the data for the slice is 0 or null.

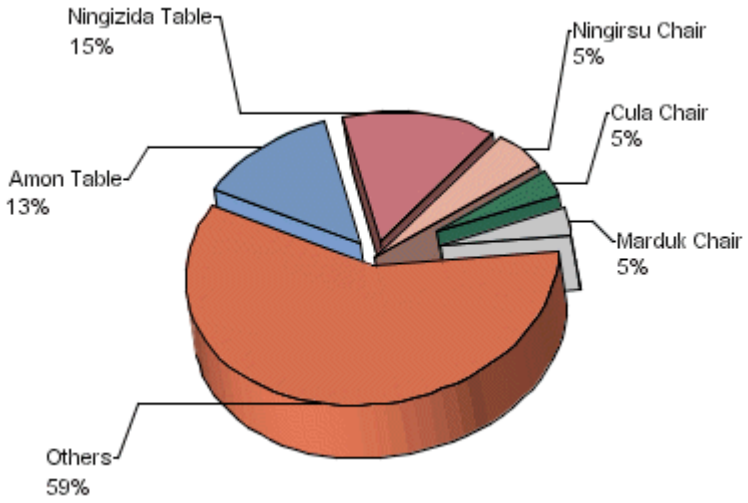
% in legend: This option allows you to display the percentage represented by each slice in the pie in the legend. This can be a preferable presentation if the pie slices become too thin. This option is only available if the pie chart does not have a data series.

Value in legend: This option allows you to display the value represented by each slice in the pie in the legend. This can be a preferable presentation if the pie slices become too thin. This option is only available if the pie chart does not have a data series.

Border: This option specifies whether to draw a border around each pie slice. This option is only available for two-dimensional pie charts. For three-dimensional pies, you can use the border drawing option on the navigation panel. Note that the border will not appear if the data for the slice is 0 or null.

Slice-to-Label Line: This option will draw a line from any label(s) to it's corresponding pie slice. Note that the slice-to-label line will not appear if the data for the slice is 0 or null.

Label at the Side: This option will place labels for the pie chart away from the plot around the outside of the chart. When used with the *Slick-to-Label Line* option, it gives users a way to display the pie labels for charts with many small categories without any text overlapping. Note that the label will not appear if the data for the slice is 0 or null.



Pie Chart with Side Labels and Lines

Best Fit: This option will arrange multiple pies in best configuration to fit the chart canvas. It's only available for pies with data series.

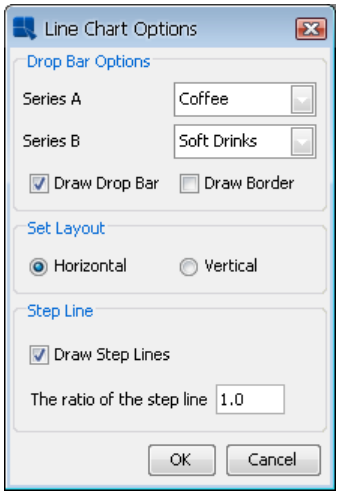
Draw Linearly: This option will arrange multiple pies in a straight horizontal line. It's only available for pies with data series.

Number of Pies Per Row: This option allows you to create a custom arrangement of multiple pies, by specifying the number of pies to draw in each row of the arrangement.

Gap between pies: This option allows you to specify the gap between the multiple pies. The number is a multiple of the pie radius, so the gap will adjust with the size of the chart plot.

5.2.4.9.5. Line Charts:

For two-dimensional line charts, one of two different dialogs will be displayed depending on whether the chart has a data series or not. If the chart has a data series then the following dialog is displayed.



Line Options Dialog (with series)

Line charts with data series have a specific option that allows you to draw drop bars between two series elements. The dialog options are as follows:

Series A: This option specifies the first series element for the drop bar.

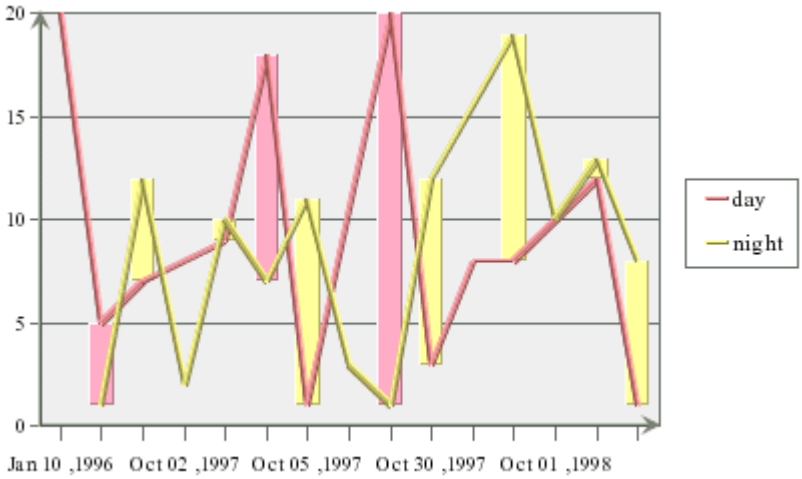
Series B: This option specifies the second series element for the drop bar.

Draw Drop Bar: This option specifies whether or not to draw drop bars.

Draw Border: This option specifies whether or not to draw a border around drop bars.

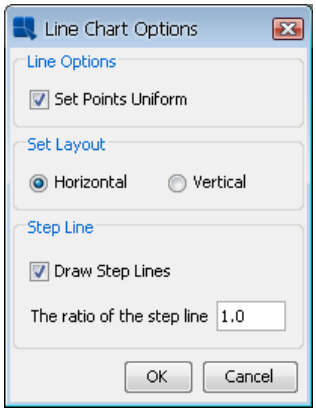
Set Layout: This option specifies whether to draw a line chart in vertical or horizontal orientation.

Step Line: This option allows you to draw line chart as a step line. You can also specify the step line ratio to use.



Line Chart with Drop Bars

Note that the color of the drop bar will vary depending on which series has the higher value for a given point. If the line chart does not have a series, then the following dialog will appear.



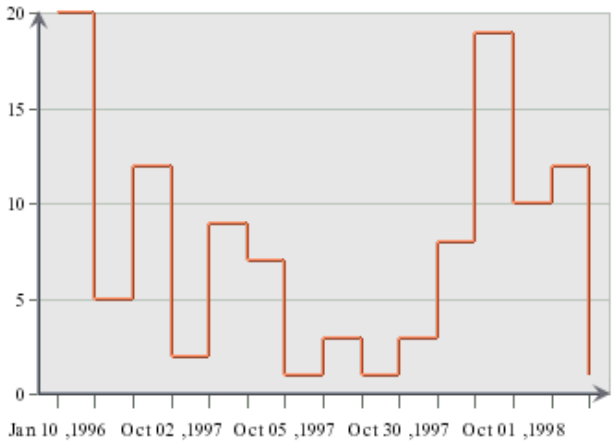
Line Options Dialog (without series)

The dialog options are as follows:

Set Points Uniform: This option specifies whether the point shapes and colors are uniform or not. Un-checking this option allows you to set multiple colors and point shapes for the data points. The points can be customized in the line and point dialog.

Set Layout: This option specifies whether to draw the line chart in vertical or horizontal orientation.

Step Line: This option allows you to draw the line chart as a step line. You can also specify the step line ratio to use.



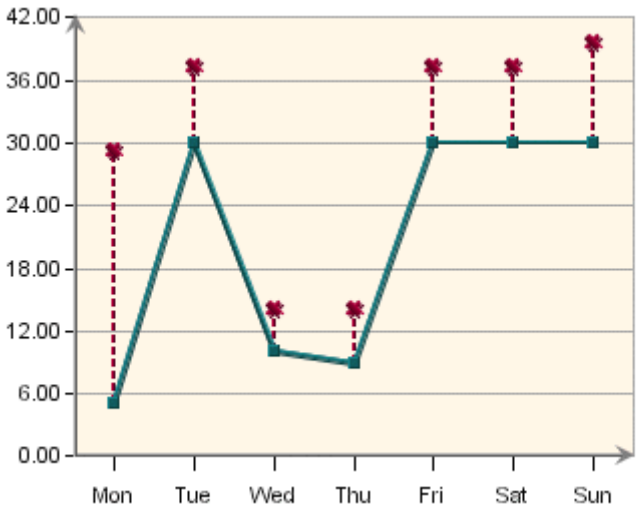
Line Chart with Step Lines

The step line ratio allows you to specify how far between points the horizontal portion of the step line should be drawn. A ratio of 1 draws the line horizontally to the next point on the chart and then draws vertically to that point, while a ratio of 0.5 draws the horizontal portion halfway between the two points. A ratio of 0 will result in the vertical portion of the line drawn first and then connect to the next point horizontally. Values for the ration are between 0 and 1.

There are no additional options for three-dimensional line charts.

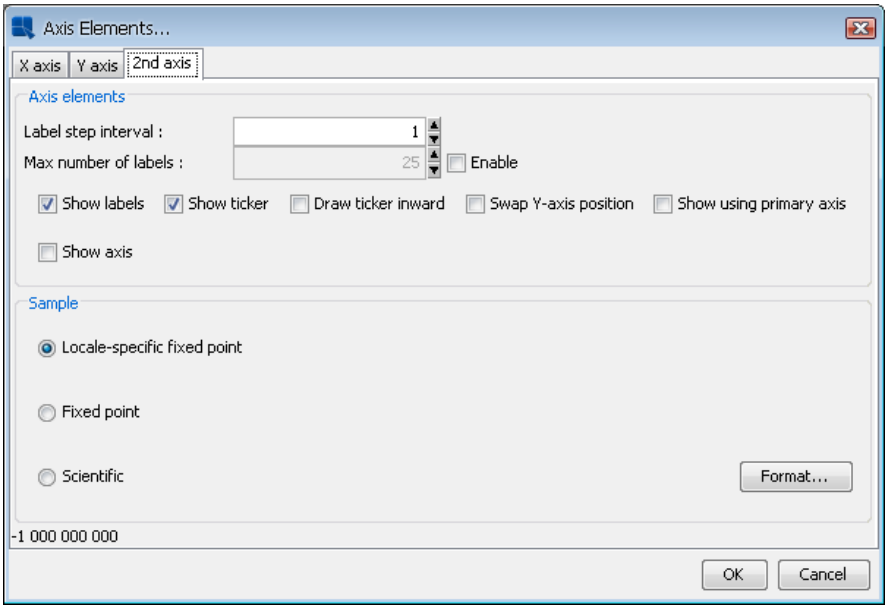
5.2.4.9.5.1. Double Value Line Charts

ChartDesigner contains a special option for line charts that allows you to have two values shown for the same line. Here the secondary axis is used to plot the second value (as in a line-line combination) and then combined with the line on the primary axis.



Double Value Line Chart

To create a double value line chart, design a line-line combination chart (a line chart with primary and secondary values). Then select Format → Axis Elements. This will bring up the axis elements dialog.

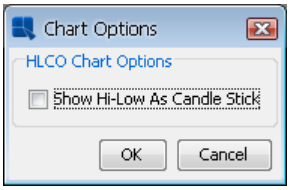


Axis Elements Dialog for Line-Line Combination Charts

Under the *2nd Axis* tab, there is a checkbox marked *Show using primary axis*. Check this box and click on the *OK* button. Your chart will now be drawn as a double value line chart.

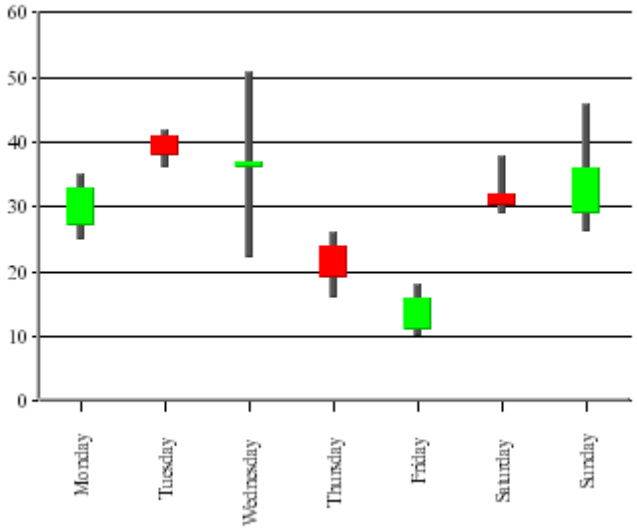
5.2.4.9.6. HLCO Charts

For HLCO charts, the following dialog is displayed:



HLCO Options Dialog

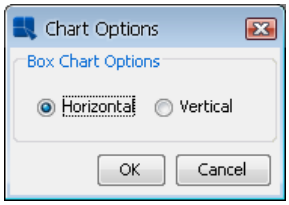
The *Show Hi-Low As Candle Stick* option will turn the HLCO chart into a candle representation. A candle HLCO chart blends high, low, close, and open data into a single object that resembles a candlestick.



HLCO Candlestick Chart

5.2.4.9.7. Box Charts

For box charts, the following dialog is displayed:

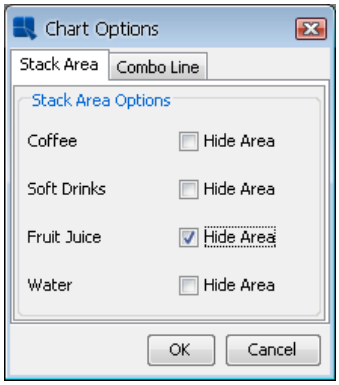


Box Options Dialog

This dialog allows you to specify whether to display the box chart in a horizontal or vertical orientation.

5.2.4.9.8. Stack Area Charts

For stack area charts, the following dialog is displayed:

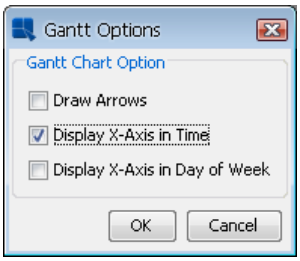


Stack Area Options Dialog

You can choose to hide any of the stacks in the chart by clicking on the corresponding check box. The *Combo Line* tab allows you to specify step lines if the chart is a line stack area combination. There are no additional options for three-dimensional stack area charts.

5.2.4.9.9. Gantt Charts

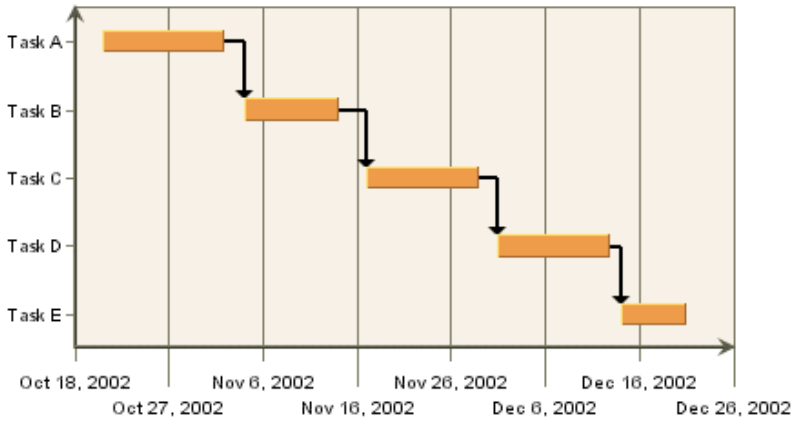
For Gantt charts, the following dialog is displayed:



Gantt Options Dialog

The following options are available for Gantt charts:

Draw Arrows: This option will draw connecting arrows between category elements of the Gantt chart. This allows you to illustrate a sequence between scheduled events. The arrows are drawn in the order the category elements appear in the data source.



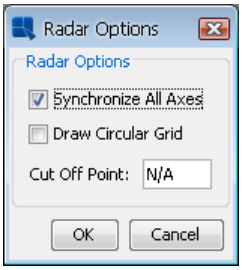
Gantt Chart with Arrows

Display X-Axis in Time: This option shows the ticker labels as time values instead of numeric values for the X-axis.

Display X-Axis in Day of Week: This option shows the ticker labels as days of the week with the date for each Sunday shown as well.

5.2.4.9.10. Radar Charts

For radar charts, the following dialog is displayed:



Radar Options Dialog

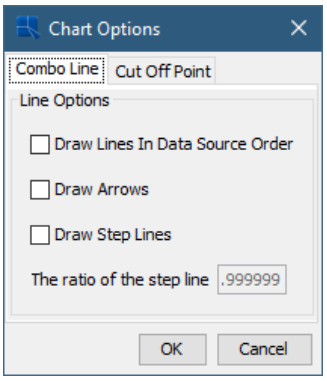
By default, the scale is same for all axes in the radar chart. Unchecking the *Synchronize All Axes* option will allow each axis in the radar chart to be scaled independently. You can select to use auto-scaling for each axis, or you can set the scales manually by invoking the axis scale dialog.

The second option allows you to set how the grid is drawn for the radar chart. By default, if the grid is enabled, it is drawn in straight lines that connect the tickers on each axis. Enabling the *Draw Circular Grid* option will draw the grid in a circle, similar to the polar chart grid.

The third option allows you to specify a cut-off point for the data points (areas) in the radar chart. You can enter a maximum value that should be shown in the chart. The areas bounded by the data points will not be drawn beyond the specified cut-off point.

5.2.4.9.11. Scatter Charts

For scatter charts, the following dialog is displayed:



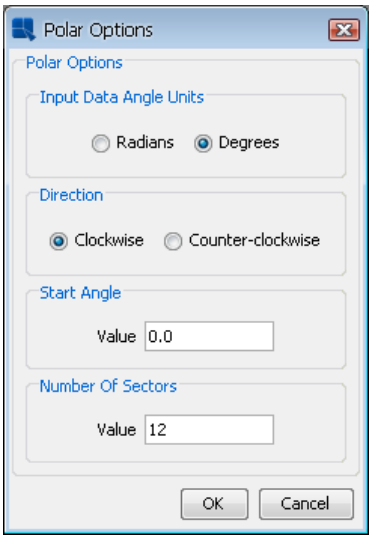
Scatter Options Dialog

The *Combo Line* tab allows you to set Draw Lines In Data Source Order, Draw Arrows on lines or Draw the connecting lines as a Step Lines, as well as specify the step lines ratio.

The *Cut Off Point* tab allows you to specify a maximum value for the Y point of the scatter coordinates. Any coordinates that fall beyond this threshold will not be plotted. Connecting lines will draw up to the edge of the threshold and continue to the next data point.

5.2.4.9.12. Polar Charts

For polar charts, the following dialog is displayed:



Polar Options Dialog

Scale: This option allows you to specify whether the input data for the angle (θ) portion of the data points is in radians or degrees. The chart will always display angles from 0 to 360. If the input data is in radians, it will be displayed as degrees.

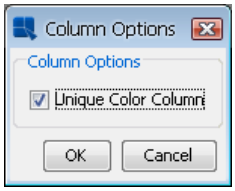
Direction This option allows you to specify whether the circular plot should be drawn clockwise or counter-clockwise.

Start Angle: By default, the top of the polar chart plot is 0 degrees. This option allows you to specify a different angle for the top of the plot. The argument for this angle is supplied in degrees or radians, depending on the scale you have chosen.

Number of Sectors: This option allows you to select number of sectors you'd like to show in the chart. Sectors are created by drawing additional polar axis lines at specified angle intervals. By default, four sectors are shown.

5.2.4.9.13. Column Charts with Series

For column charts with data series, the following dialog is displayed:

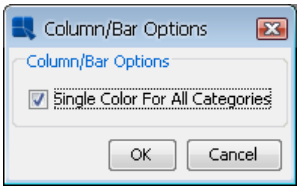


Column Options Dialog

Normally, when column charts have data series, each series has it's own color that is applied for every category in the chart. If you want to assign different colors to the columns in the chart regardless of the series, you can enable the *Unique Color Column* option in this dialog. When it's turned on, you can set color for each column in the chart individually.

5.2.4.9.14. Column/Bar Charts without Series

For column/bar charts without data series, the following dialog is displayed:

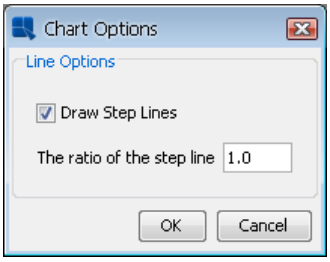


Column/Bar Options Dialog

Normally, when column/bar charts don't have data series, all categories in the chart have single color. If you want to assign different colors to the categories in the chart, uncheck the *Single Color For All Categories* option in this dialog.

5.2.4.9.15. Two-Dimensional Line Combination Charts

For any other two-dimensional line combination chart, the following dialog is displayed:

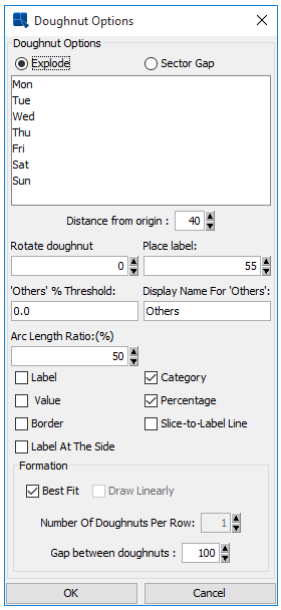


Line Combination Chart Options

This dialog allows you to specify whether to draw the combo line as a step line, as well as specify the step line ratio.

5.2.4.9.16. Doughnut Charts

The chart options for a doughnut chart are almost exactly the same as the options for a pie chart. You can refer to the options under Section 5.2.4.9.4 - Pie Charts for more details.



Doughnut Chart Options

The only option unique to doughnut charts is the Arc Length Ratio (%) which specifies the size of the hole at the center of the chart. A higher number results in a smaller hole.

5.2.5. Chart Drill-Down

Like with reports, it is also possible to add layers of drill-downs to charts. All of the drill-downs added will be applied to both data points on the plot area and their respective fields in the legend. Using chart drill-down, you can create a top-level chart that displays summarized data and allows users to click through specific data points to see underlying details.

The charting engine in ERES contains several different built-in drill-down mechanisms that allows you to create only one chart template for each level of drill-down. All of the drill-down options are available from the Drill-Down menu in Chart Designer.



Warning

It is important to note that you can only use drill-downs in charts when the chart is deployed independently from the report (in the Chart Viewer, its own servlet, etc). If a chart template that includes drill-down is placed in a report, only the top-level chart will display. Users will not be able to click through to the lower-level charts.

5.2.5.1. Data Drill-Down

Data drill-down allows you to group and display information that is based on a single data source. The advantage to this form of drill-down is that it works with any data source. However, it does not allow you to display loosely related information because all levels of drill-down will share the same value column from the input data.

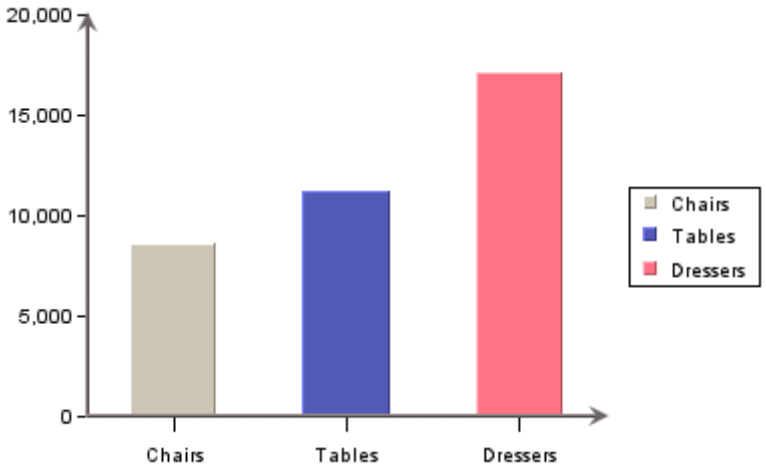
For example, assume you have the following table as the chart data:

Category	Product	Sales
Chairs	Elm Arm Chair	$8,216
Chairs	Pine Side Chair	$7,611
Chairs	Redwood Arm Chair	$8,625
Tables	Elm Round Table	$10,241
Tables	Pine Oval Table	$9,663
Tables	Oak Oval Table	$11,261
Dressers	Oak Single Dresser	$16,442
Dressers	Elm Double Dresser	$17,148

Using this data, you could create a top-level chart that displays total sales for each distinct product category and then create a lower-level chart that shows individual sales for each product in a category. The number of drill-down levels available depends on the number of groupings that are present in the input data.

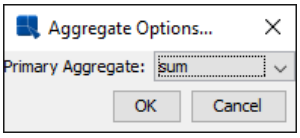
5.2.5.1.1. Adding Data Drill-Down

To add layers of data drill-down to a chart, you must first create a top-level chart. In the above example, we would create a chart with “Category” mapped to the category axis and “Sales” mapped to the value axis. In a column chart, it would look like this:



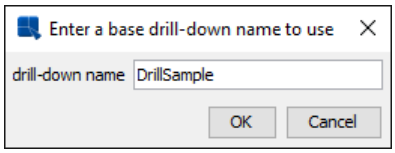
Top-Level Chart

To add a layer of drill-down, select Drill-Down → Add. This will bring up a dialog prompting you to specify the aggregation you want to use for the value axis. For example, selecting sum will display the total for each category in the top-level chart. Available aggregate functions are minimum, maximum, average, sum, count, first, last, sumsquare, variance, stddev, and countdistinct.



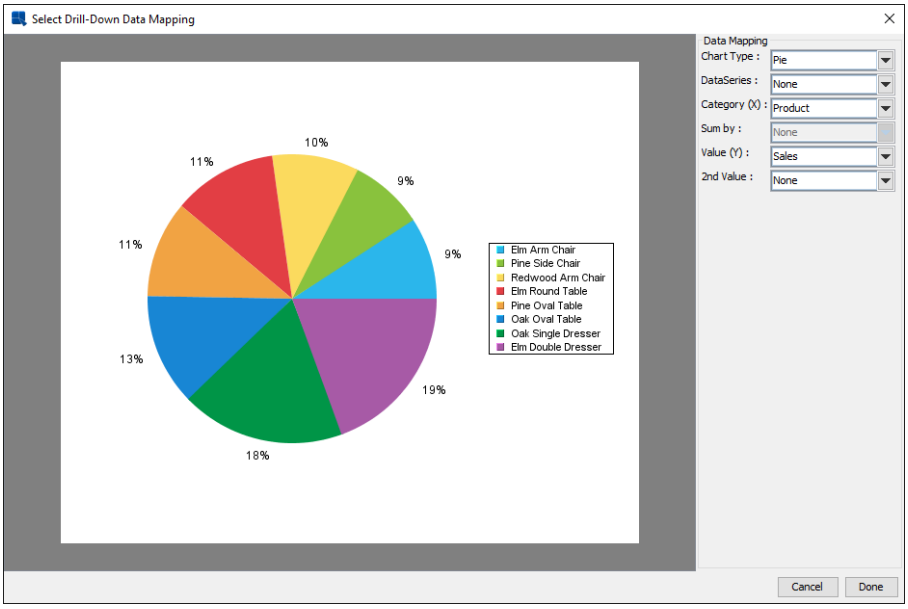
Aggregation Dialog

After you select the aggregation you want to use, click on the *OK* button. A new dialog will appear prompting you to specify a name for the drill-down chart. This name will be assigned to the chart templates that are created by the drill-down process. Note that all sub-level drill-down templates are saved in the `/drilltemplates/` directory.



Drill-Down Name Dialog

Once you have specified a name you want to use, click on the *OK* button. You will then be prompted to specify the chart type and data mapping for the sub-level chart.



Data Drill-Down Mapping Dialog

The options in the dialog are similar to those for normal data mapping. The major difference is that the first option allows you to select a chart type. Available chart types for data drill-down are column, bar, line, stack column, stack bar, pie, area, doughnut, overlay, radar, and dial. The data mapping options will change depending on the type of chart that you select.

Once you have finished specifying the mapping options, click on the *Done* button and you will go back to the Chart Designer where you can customize and modify your sub-level chart. You can add additional layers of drill-down by selecting Drill-Down → Add again.

You can navigate to a particular drill-down chart by selecting Drill-Down → Previous or Drill-Down → Next or by double clicking a data point. The *Previous* selection takes you up a level, while the *Next* selection goes down a level using the left-most category as the data to be drilled on. You can also go from any drill-down chart to the top-level chart by selecting Drill-Down → Go To Top Level. You can customize (i.e. assign colors, add title, axis labels, background image, ...etc.) the drill-down chart in the same way as the top-level chart. Everytime you change and leave a level, you will be prompted to save the chart. Make sure that you answer *yes* if you want the attributes of the drill-down level to be saved, otherwise you will you lose all the changes.

You can insert a drill-down chart between two other drill-down charts by going to the higher level drill-down chart and selecting Drill-Down → Add. The drill-down chart that you create will be inserted between the two previous drill-down charts.

You can remove a level of the drill-down by navigating to that level and selecting Drill-Down → Remove This.

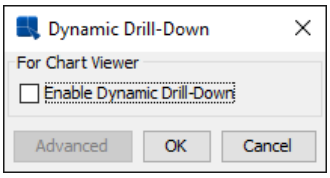
The entire drill-down chart can also be deleted by selecting Drill-Down → Remove All. Selecting Drill-Down → Previous brings the drill-down chart to a higher level, which is the same as right clicking and selecting *Back* from the pop-up menu. The *Next* selection brings the drill-down chart to a lower level, which has the same function as double clicking on an item in the chart.

Once you have finished creating and editing all the levels of the drill-down chart, you can save it by navigating to the top-level chart and selecting File → Save or File → Save as.

5.2.5.2. Dynamic Data Drill-Down

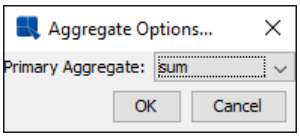
Usually, mapping and drill options for data drill-drown are fixed during design time. Dynamic drill-down is an additional option that allows you to select the mapping. The only thing specified during design time is the top-level chart and aggregation.

To create a chart with dynamic drill-down, first create a chart you want to use as the top-level chart (for example, you can create the same top-level chart as before) a then select Drill-Down → Dynamic. A dialog will appear allowing you to enable dynamic drill-down.



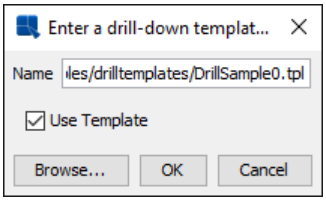
Enable Dynamic Drill-Down Dialog

Check the *Enable Dynamic Drill-Down* box and a new dialog will pop-up prompting you to select the aggregation.



Aggregation Dialog

Options for this dialog are the same as for regular data drill-down. Once you specify the aggregate you want to use, click on the *OK* button. Another dialog will pop-up, allowing you to specify a template you want to use for the sub-level charts.



Select Dynamic Drill-Down Template Dialog

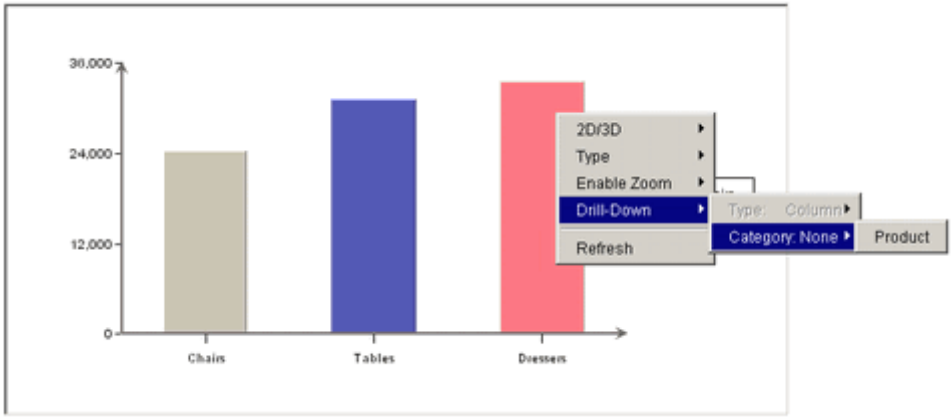
If you do not select to use a template, the sub-level charts will be generated using default appearance properties. After you finish selecting these options, your chart will change to display aggregated data on the value axis. You can change the option by selecting Drill-Down → Dynamic again and then clicking on the *Advanced* button.

Dynamic drill-down charts can only be viewed in the Chart Viewer JNLP. There you can configure setting of the next drill-down chart by right clicking on the desktop area which will bring up a pop-up menu. If the dynamic drill-down is enabled and the chart still has some unused fields (columns) to plot, the item *Drill-Down* will be displayed in the pop-up menu. Upon selecting *Drill-Down* in the pop-up menu, you can view the current setting for next drill-down. If *Category* is **None**, it means you have not configured the next drill-down chart yet. To configure the drill-down chart for next level, you must select *Category* (*Type*, *Series* and *SumBy* can be selected once *Category* has been set). After you create a drill-down chart, you can navigate to different levels in the Chart Viewer by either left clicking on a data point if you want to go down one level, or by right clicking and selecting *Back* from the pop-up menu if you want to go up one level, or by repeating the previous step to create another drill-down chart for next level.



Tip

After traversing to a lower-level of drill-down, right clicking on a data point will navigate you back to the top-level chart. To bring up the pop-up menu, right click on the chart canvas away from the chart data points.



Dynamic Data Drill-Down in Chart Viewer

5.2.5.3. Parameter Drill-Down

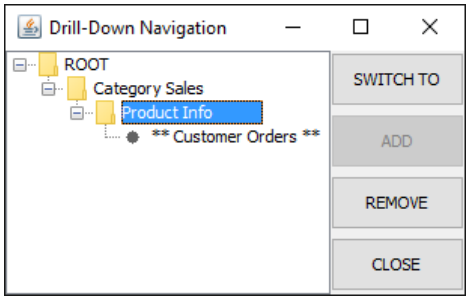
The third type of a drill-down you can use for charts is a parameter drill-down. Parameter drill-down is the most flexible implementation because you can relate the data between drill-down levels in any way you like. Specifically, you do not need to use the same value element for each level. Instead, parameter drill-down uses parameterized query feature to relate various chart levels and since it uses queries, the data source for the sub-level charts must be a database or parameterized class file.

For example, using our previous scenario, rather than always looking at sales, you want your top level chart to look at aggregated sales by category (same as before), then on the next level, you want to look at sales volume for each product, and from there, you want to look at inventory levels for each product by region. This can be accomplished with parameter drill-down.

With parameter drill-down, the category value of an element you click to drill on will be passed to the sub-level chart as the parameter value. Hence, if you click on the “Chairs” column, the value of “Chairs” would be passed to the query. Therefore, anything in the database that could be filtered by category name could be retrieved.

5.2.5.3.1. Adding Parameter Drill-Down

To add and edit various layers of a parameter drill-down, select Drill-Down → Parameter Drill-Down. This will bring up a navigation window that shows various levels of drill-down.



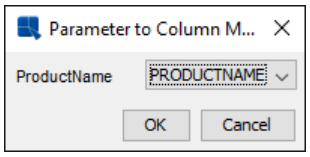
Parameter Drill-Down Navigation Window

Left-hand side of the navigation window displays the hierarchy of drill-down levels. The “ROOT” node is the top-level chart. The level you are currently editing is marked with **. To edit a different level, select it and click the *SWITCH TO* button on the right-hand side. The chart will then open in the Designer. Levels of drill-down can also be removed by selecting the node in the Navigation window and then clicking on the *REMOVE* button.

To add a new level of drill-down, select the chart under which you would like to add the layer (if you have only the top-level chart, then only the “ROOT” node will be visible), and click *ADD*. You will then be prompted to create a new chart or to use an existing chart for the drill-down layer. You can use any existing chart; however, any chart for a drill-down layer must have a parameterized query as the data source. If you select to create a new chart, the

Data Source Manager will reopen, allowing you to select a data source for the chart and to continue through the Chart Wizard steps.

The next step is to map the category and/or series columns from the top-level chart to the query parameter(s) in the sub-level chart. You will be prompted to do this when you select an existing chart for the drill-down layer or when you select the data source for a new chart for the drill-down layer. In either case, a dialog will appear prompting you to map the fields.



Parameter Mapping Window

Available options in the drop-down menus are based on data type. For example, if your parameterized query has string as the parameter, only fields containing string data can be mapped. Hence, it is important to consider what type of data will be passed from the top-level chart to the lower-level charts because if your category or series are not of the correct data type or if you do not have enough fields to pass to the lower-levels, the drill-down will not work correctly.

Once you correctly specify the parameter mapping, you will be prompted to specify a display name for the drill-down level. After you select a name, the sub-level chart will appear in the Designer, allowing you to customize it. You can navigate through the layers of drill-down using the Navigation window, or like data drill-down, you can double click on a data point to go down a level and right click and select *Back* from the pop-up menu to go up a level.

5.2.6. Saving Charts

After you finish designing a chart, you can save the definition as a chart or as a chart template, and provide XML definitions for both. You can also generate a number of static image exports of the chart or create a JSP/JNLP page with the Chart Viewer JNLP.

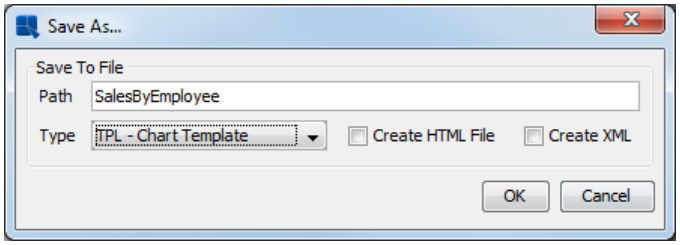
5.2.6.1. Saving Charts with Report Data

When you design a chart using report data, the chart cannot be run/viewed any place except within the report. Because of this, most of the file operations are disabled in the Chart Designer. When you opt to save a chart it will be automatically saved with the same name as the report in the `/chart/` directory of the installation. For example, if your report is named *SalesReport*, any charts in the report that use report data would be saved in the following convention `SalesReport_0.tpl, SalesReport_1.tpl,` etc

Although most of the file options are disabled when you use report data for the chart, you can apply a template to carry over the appearance properties of another chart. Chart templates are covered in Section 5.2.6.2.1 - Working with Templates.

5.2.6.2. Saving Charts without Report Data

If your chart uses an independent data source from the report, it can be saved in any place with any filename. To save the current chart, select File → Save or File → Save As, or click on the  *Save* button on the toolbar. Assuming you have not saved the chart before, or selected File → Save As, the following dialog will appear:



Save As Dialog

The first option allows you to specify a name and file path for the saved chart. You can browse to the appropriate file path by clicking on the *Browse* button at the bottom of the dialog. The second option allows you to specify which format you would like to use when saving the chart. As detailed in Section 5.2.2 - Charting Basics, there are two principle ways in which chart definitions can be saved.

Chart format: Chart files save the chart in a binary file called `filename.cht`. A chart file stores both the definitions of the chart (type, dimension, etc), as well as the data that was used to create the chart.

Template format: Template files save the chart in a binary file called `filename.tpl`. A template file only stores chart definitions. It doesn't store any chart data. Hence, any time a template file is opened, it will try to connect to the original data source to retrieve the data.

PAC format: PAC files save the chart in such a way that makes them ready for deployment. A `.PAC` file takes the chart and all the supplementary files associated with it, such as background images, dynamic drill-downs, or parametric drill-downs, and places them in a single binary file.

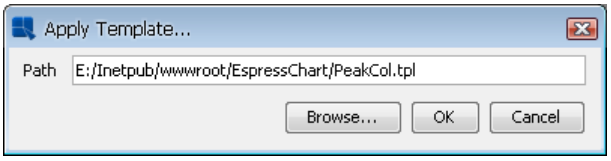
Once you select a format you want to use, click on the *OK* button and the chart will be saved.

5.2.6.2.1. Working with Templates

Chart templates are a special format in which chart definitions can be saved. Unlike chart files (`.cht` format) which saves both the chart definitions as well as the data for the chart, templates only save the definitions (i.e. the chart attributes and the layout of the individual components), and the data source information. This means that templates store the location/connection information for the data source that was used to create the chart, but they store none of the actual data in the file. (Templates do keep 10 records of back-up data, allowing them to be opened when the data source is not present).

Template files can be used in two ways. First, it can be opened, viewed, or exported, allowing you to see a chart with fresh data. This way data for the chart are retrieved based on the data source specified in the file. If the original data source is not accessible, this method cannot be used. The second way in which templates can be used is to apply its attributes to other charts. In this scenario, the chart to which the template is applied will inherit appearance properties of the template (including colors, fonts, size of chart components, position of legends, etc). This feature allows you to produce a consistent look and feel among charts.

You can apply a template in Chart Designer by selecting File → Apply Template. This will bring up a dialog prompting you to specify the template file you want to use.



Apply Template Dialog

The dialog allows you to specify the template file and its location. You can browse to a file by clicking on the *Browse* button.

Note that when the chart and the template being applied are different sizes (i.e. chart canvas size), the resulting chart may not display correctly. This is because the text size will not change between template and chart to which it is applied. While the other components will adjust to the size of the new canvas, the font will not. To keep a consistent look, the size of your template should be close to the size of the chart to which it will be applied.

Hyperlinks, floating lines, floating text, and axis scales defined in the template files carry over. You may have to redefine them in the chart, if necessary. Chart type and dimension are not modified by the template. For example, a three-dimensional chart will not be changed to a two-dimensional chart if the template is a two-dimensional chart. Likewise, a pie chart remains a pie chart when a template of a bar chart is applied. Although the chart type does not change when you apply a template, some appearance properties will not translate very well from a template with another chart type. For best results, try to apply templates of the same type and dimension to a chart.

5.2.6.2.2. Saving XML Templates

In addition to the two binary chart representations, you can also save the chart definitions in XML format. This option allows you to have a text-based chart template that can be used to modify chart properties outside of Chart Designer or the API.

To create an XML file when saving the chart, check the *Create XML File* box in the save as dialog when saving the chart. This will create an XML file for the chart. Note that the XML file is a representation of the `.tpl` format and not the `.cht` format.

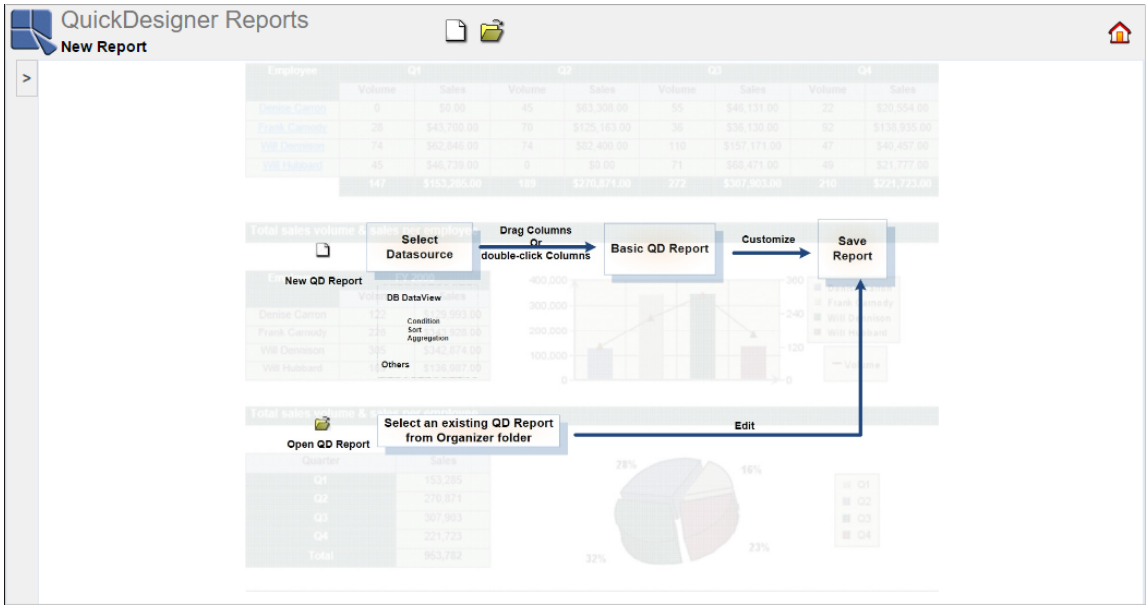
5.3. QuickDesigner Reports

QuickDesigner Reports is a thin-client ad-hoc reporting interface. It provides users a scaled-down design tool to create reports. With QuickDesigner Reports, end users can easily select, filter, and present data without mastering database structures, all with zero client download.

5.3.1. Start

QuickDesigner Reports can be started from the ERES Start page. If you have logged in as a user with design privileges, then you can follow QuickDesigner Reports link to begin using QuickDesigner Reports.

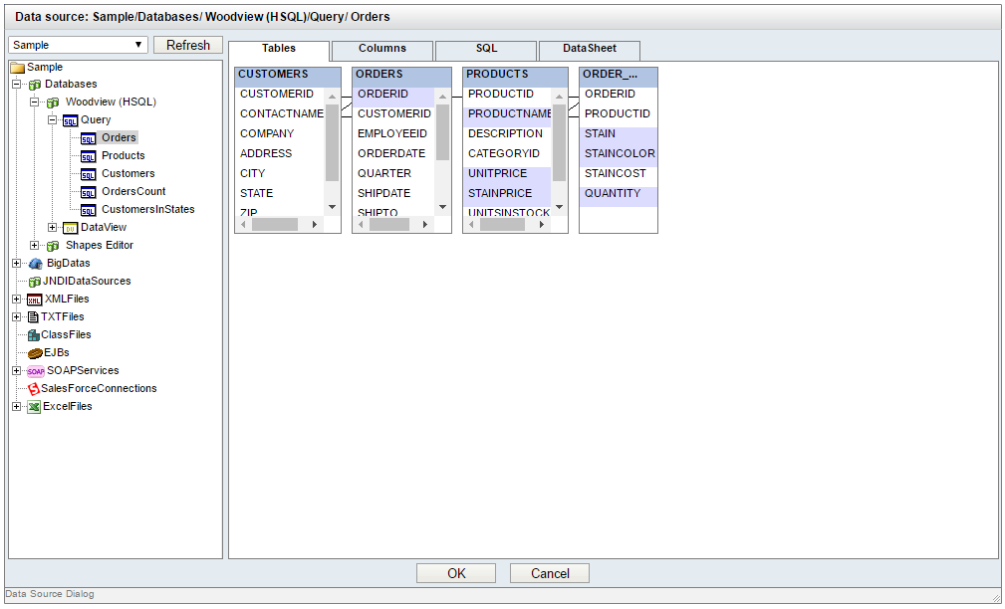
When the QuickDesigner Reports launches, the first page that appears prompts you to select between opening an existing report, or creating a new one.



QuickDesigner Reports Start Options

5.3.2. Select a Data Source

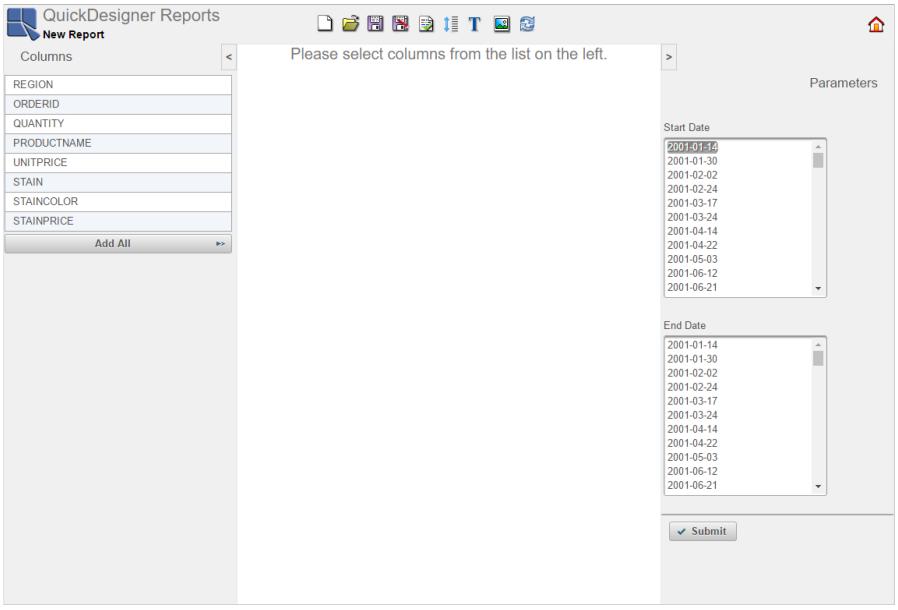
If you want to create a new report, click on the  *New* icon on the toolbar. The *Data Source Dialog* will appear and prompt you to select data source you want to use with QuickDesigner Reports.



Data Source Dialog

Unlike Report Designer or Chart Designer where the user can select to create new data sources or modify the existing ones, the only option is to select a data source for the report. The only exception to this is for Data Views or Data View Queries. If data sources of this type are selected, you will see the *DataView Builder* on the right side of the Data Source Dialog. The *DataView Builder* allows you to build or modify queries against the View. For more about managing data sources in QuickDesigner Reports, please see Section 4.2 - Data in QuickDesigners and Maps.

Once you finish selecting a data source, click on the *OK* button to close the *Data Source Dialog*.



QuickDesigner Reports

You can see columns of the data source on the left side, the central part for creating reports, and parameters pane on the right side (this pane is visible only if the data source has some parameter(s)).

5.3.3. Toolbar

A lot of formatting actions for reports in QuickDesigner Reports are accessed through the toolbar. The icons perform the following actions:

 Start a new report

 Open an existing report

 Save the current report

 Save report as

 Export the current report

 Set style for the report

 Pagination/Scrolling options

 Insert/Edit report title

 Insert report logo

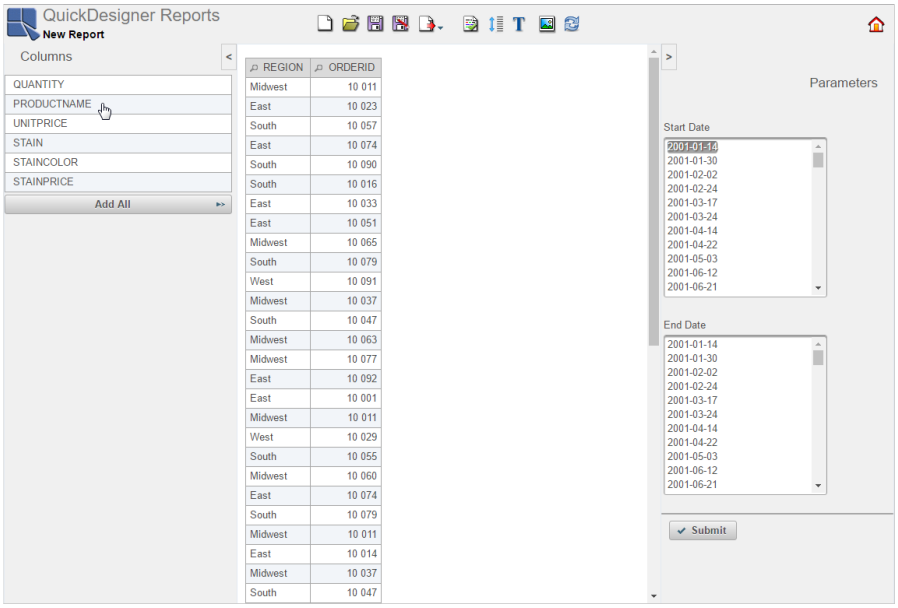
 Modify DataView Query - This option allows user to modify the DataView and DataView query used for the report. The icon is only visible in toolbar when a DV query is used as data source for the report (otherwise it's hidden).

 Refresh the report

5.3.4. Format Report Elements

5.3.4.1. Add a Column

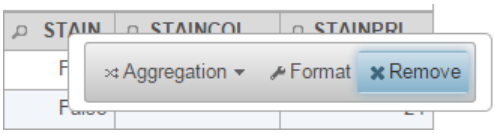
To insert a column to the report, double-click on the column in the left *Columns* pane, or click on the column and drag it to the report. You can also add all columns to the report in one step by clicking the *Add All* button.



Add a Column

5.3.4.2. Remove a Column

To remove a column from the report, right click on the column header and select *Remove* or drag the column back to the *Columns* pane. The column is removed from the report.

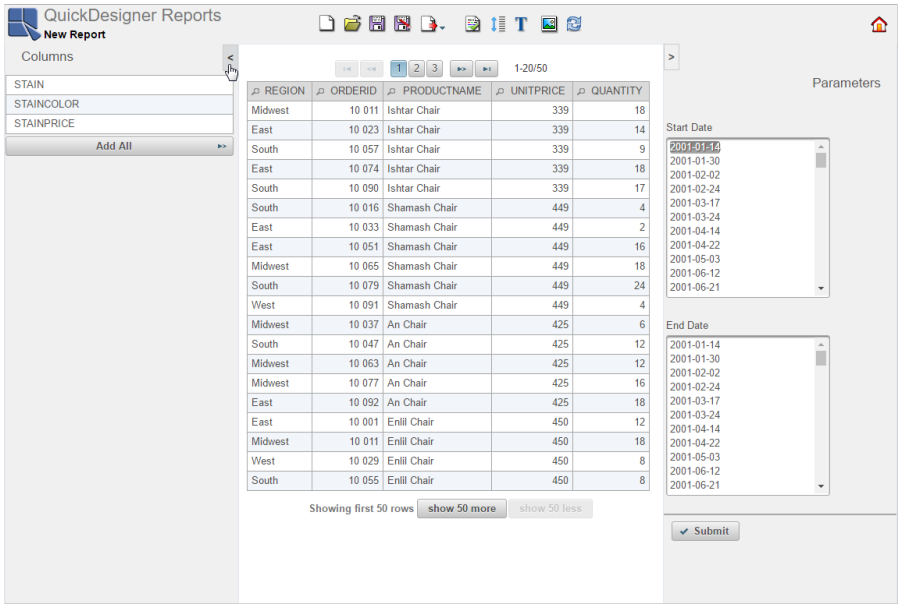


Remove a Column

5.3.4.3. Collapse Sidebars

When you add columns to the report, you can colapse the *Columns* pane to have more space for designing the report.

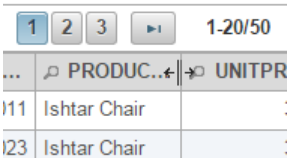
To do this, click on the < *Collapse* button at the top of the *Columns* pane. You can also collapse the *Parameters* pane on the right side (if the report has some parameters).



Collapse Sidebars

5.3.4.4. Column Width

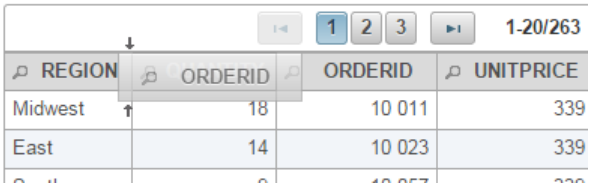
You can simply resize a column width to fit the content or to see the whole column header. Move the mouse over the right side of the column header to see double arrow, left click and drag to set the column width.



Column Width

5.3.4.5. Column Order

You can also reorder the columns. Left click on the column header and drag it. When you see the arrows, release the mouse button.



Column Order

5.3.4.6. Show More Rows

By default, only the first 50 rows will be displayed in the report. You can display more or less rows by clicking on the *show 50 more* or *show 50 less* buttons under the report.



Show More Rows

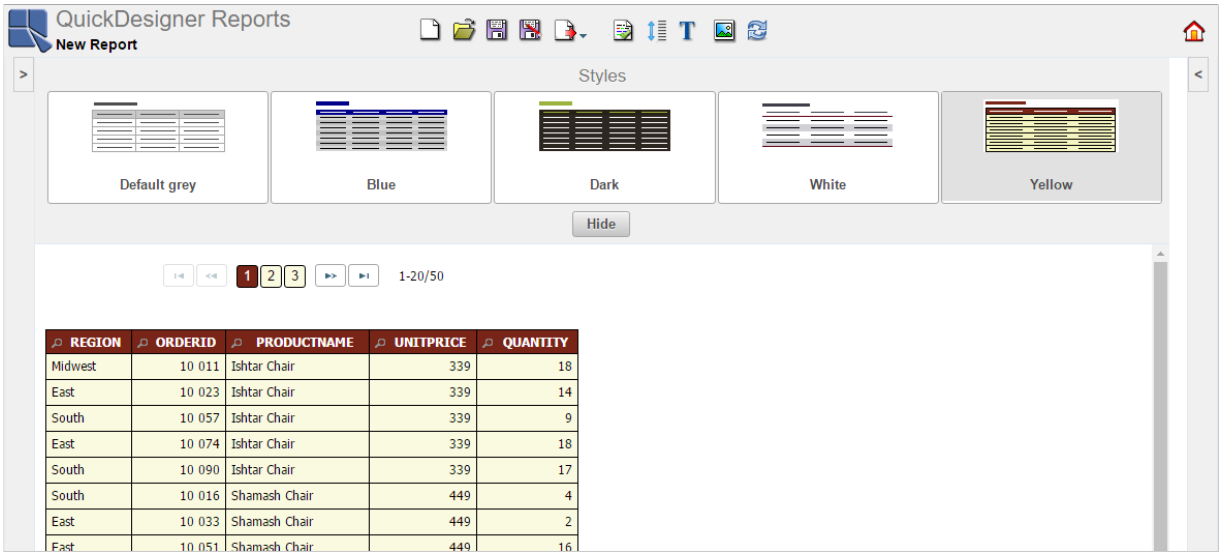
When all rows are displayed, you will also see the total number of rows.



Show Less Rows

5.3.4.7. Report Style

You can choose from five pre-defined report styles. Click on the  *Styles* icon on the toolbar to open the *Styles* dialog. Select the style you would like to use and close the *Styles* dialog by clicking on the *Hide* button.

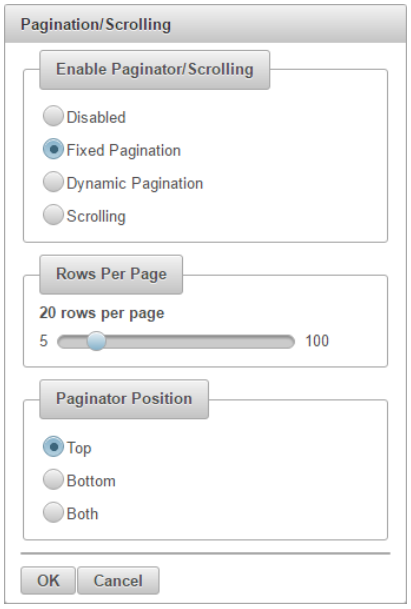


Styles Dialog

5.3.4.8. Pagination/Scrolling

Click the *Pagination/Scrolling* icon on the main toolbar  to open the *Pagination/Scrolling* dialog. You can choose fixed pagination, dynamic pagination that automatically keeps the height of your screen, scrolling, or disable this feature.

Fixed Pagination For fixed pagination, you can select a number of rows per page and a paginator position on the page.



Pagination Dialog

Paginator divides the report data into separate pages. Once you set pagination to be active, you can switch to any page of the report using buttons with page numbers. You can also switch to the first page by clicking on the  *First Page* button, to the last page by clicking on the  *Last Page* button, or to the previous and next page by clicking on the  *Previous Page* and the  *Next Page* buttons. Next to the buttons you can see which rows are displayed/total number of rows.



REGION	ORDERID	PRODUCTNAME	UNITPRICE	QUANTITY
Midwest	10 060	Enlil Chair	450	22
East	10 074	Enlil Chair	450	18
South	10 079	Enlil Chair	450	24

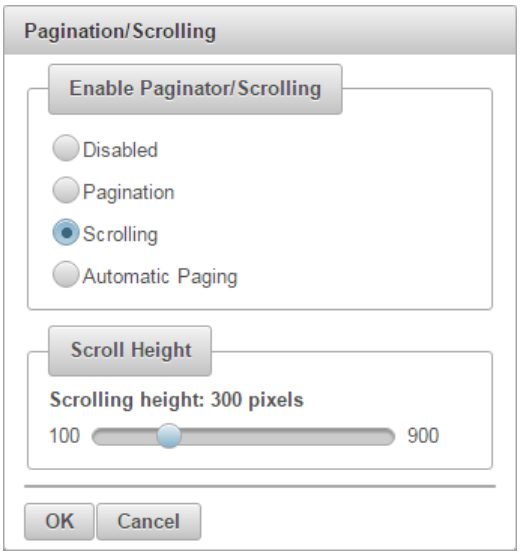
Report with Pagination Buttons

Dynamic Pagination

Dynamic pagination does not require any further setting. It will automatically adjust number of rows per page to the height of your screen. As a result, no scroll bar will appear in QuickDesigner Reports, Dashboard Builder, and Published Files. The number of rows per page may be different for different machines.

Scrolling

Scrolling allows you to display all the report data on a single page while not taking up too much space. All the displayed data is placed into a scrollable table with scrolling height of your choice. The scrolling height in pixels can be set in the *Scroll Height* section of the *Pagination/Scrolling* dialog.



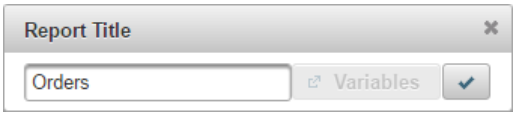
Scrolling Dialog

REGION	ORDERID	PRODUCTNAME	UNITPRICE	QUANTITY
Midwest	10 011	Ishtar Chair	339	18
East	10 023	Ishtar Chair	339	14
South	10 057	Ishtar Chair	339	9
East	10 074	Ishtar Chair	339	18
South	10 090	Ishtar Chair	339	17
South	10 016	Shamash Chair	449	4
East	10 033	Shamash Chair	449	2
East	10 051	Shamash Chair	449	16
Midwest	10 065	Shamash Chair	449	18

Report with Scrolling

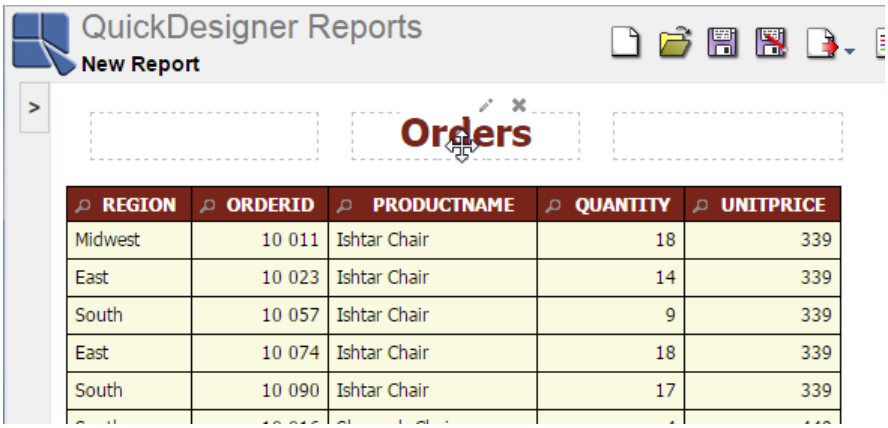
5.3.4.9. Report Title

You can add a report title. Click the  *Report Title* icon on the toolbar, and the *Report Title* dialog appears. Write a title of the report and click on the  *Apply* button or press *Enter* to insert the title.



Report Title Dialog

The report title is on the left by default. You can move it to the center or right side of the report. Click on the title, hold down the mouse button and move the mouse slightly. You can see three dashed rectangles that represent three positions of the title - left, center, right. Move the title to the position you want to use and release the mouse button.

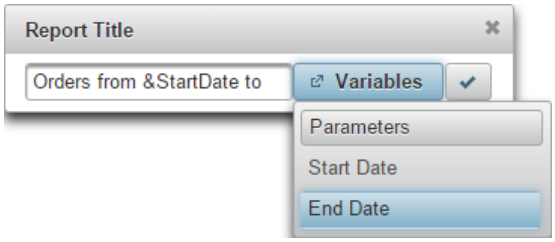


Report Title Moving

You can edit or remove the report title via the pop up *Edit* button when the cursor is pointed over the title, or via the *Report Title* icon on the toolbar. Both will open the *Report Title* dialog. There you can rewrite the title and apply it with the *Apply* button, or remove it by clicking on the *Remove* button. Then close the *Report Title* dialog by clicking on *X* in the upper right corner. You can also remove the report title directly via the pop up *X* button when the cursor is pointed over the title.

5.3.4.9.1. Report Title with Variables

If the data source is parameterized, you can add variables to the report title. Variables display the current parameter value in the report title and change dynamically according to which parameter value is selected.



Variables for Report Title

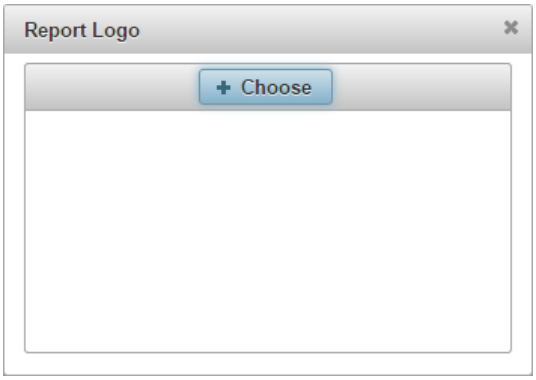
To add a variable, click on the *Variables* button in the *Report Title* dialog, select a parameter from the list, then it will be displayed in this format &<*parameterName*> in the textfield.



Report Title with Variables

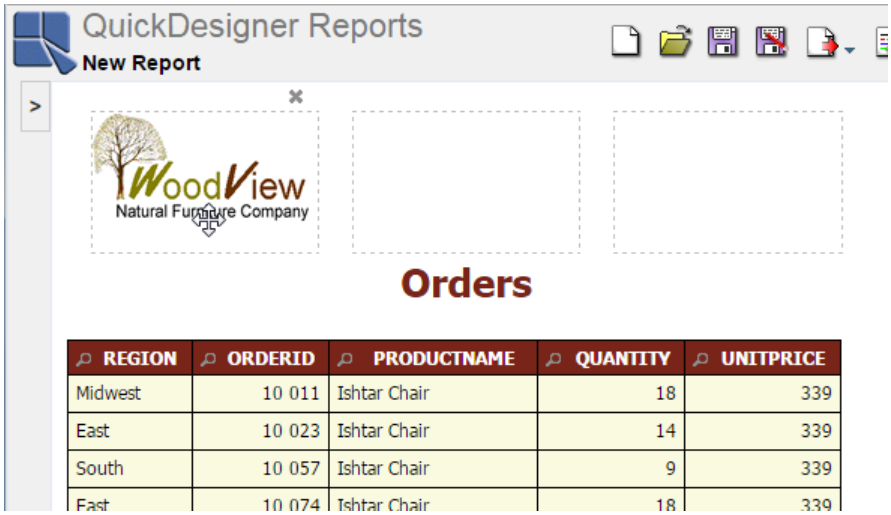
5.3.4.10. Report Logo

You can add a logo to your report. To do this, click on the *Report Logo* icon on the toolbar to open the *Report Logo* dialog. This dialog allows you to choose a logo from your local directory. Once a logo is selected, it is inserted to the top left.



Report Logo Choose

You can move the logo to the center or right side of the report. Click on the logo, hold down the mouse button and move the mouse slightly. You can see three dashed rectangles that represent three positions of the logo - left, center, right. Move the logo to the position you would like to use and release the mouse button.

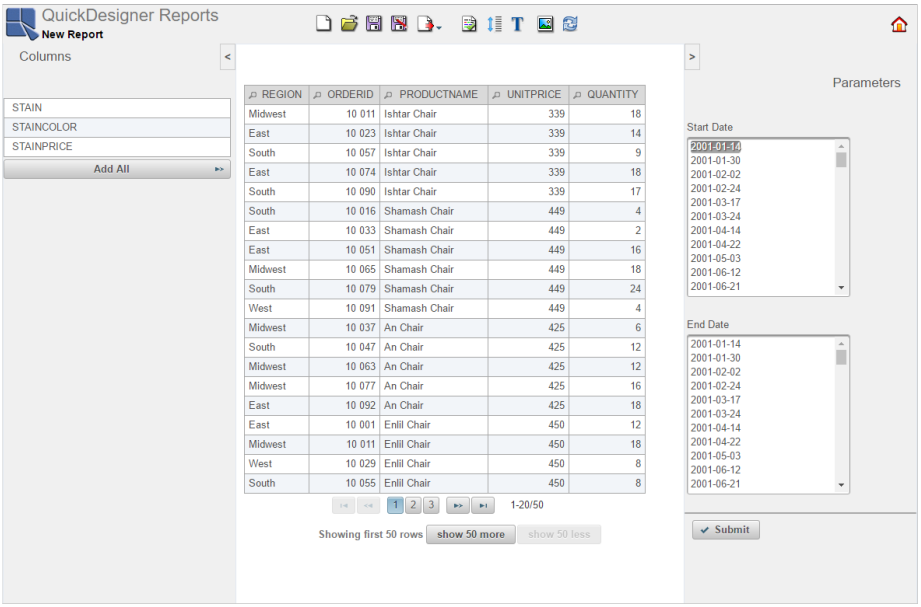


Report Logo Placement

You can remove the report logo via the pop up button when the cursor is pointed over the logo. If you want to change the logo, just click on the *Report Logo* icon to choose a new image. The new image will replace the old one.

5.3.5. Parameter Setting

If the report uses a parameterized data source, you can see the *Parameters* pane on the right side of the designer window. You can select parameter(s) value(s) and apply it (them) by clicking on the *Submit* button.

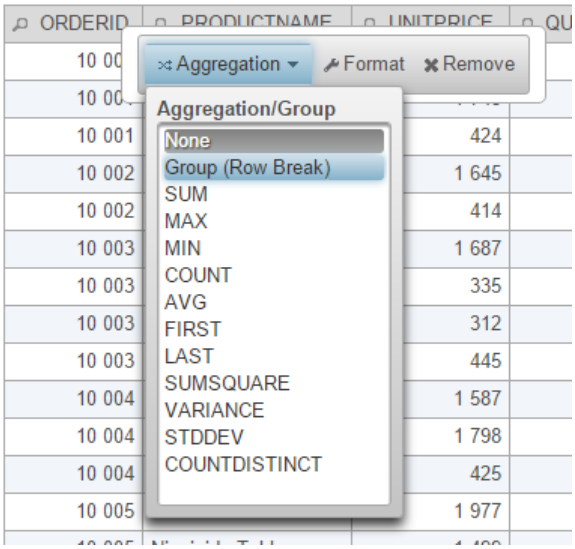


Parameters Pane

You can collapse or open this pane by clicking on the > or < *Collapse* buttons.

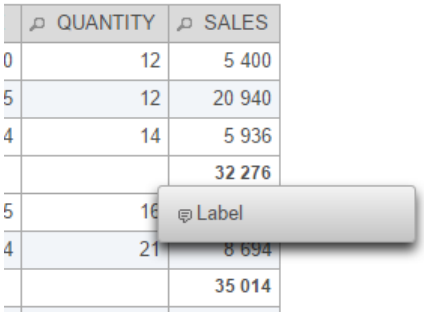
5.3.6. Aggregation/Group

You can add an aggregation or group to each column of the report. To do this, right click on the column header and move the mouse cursor over *Aggregation* button. The *Aggregation/Group* list will appear. Select an aggregation by clicking on it in the list. The aggregation will be assigned to the appropriate column. After you finish this setting, an extra line of the aggregation value will appear after each group (row break) and after the last aggregation value line, an extra line with a grand total value will appear.



Aggregation/Group Setting

Once the aggregations are set, you can set label for the row break and grand total. The grand total label is called Report Footer label. To set a label, right click on an empty cell next to the value you want to describe and then click on the pop up *Label* button.



Aggregation Label Button

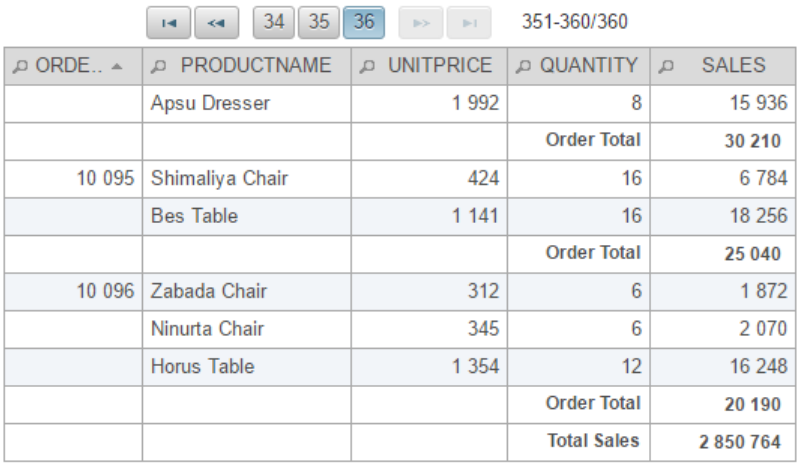
The *Footer Labels* dialog will appear. You can insert a text and decide how to align it in a cell - left or right. Then click on the *OK* button to apply the label.



Footer Labels Dialog

To edit or remove a label, right click on it and then click on the pop up *Label* button and edit the text or click on the *Remove* button in the *Footer Labels* dialog.

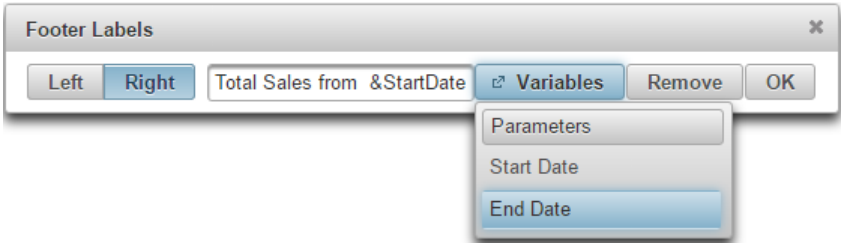
The next image shows a sample of the report with *Row Break*, *SUM* aggregation, group footer label, and report footer label.



ORDE..	PRODUCTNAME	UNITPRICE	QUANTITY	SALES
	Apsu Dresser	1 992	8	15 936
			Order Total	**30 210**
10 095	Shimaliya Chair	424	16	6 784
	Bes Table	1 141	16	18 256
			Order Total	**25 040**
10 096	Zabada Chair	312	6	1 872
	Ninurta Chair	345	6	2 070
	Horus Table	1 354	12	16 248
			Order Total	**20 190**
			Total Sales	**2 850 764**

Aggregation and Label

If the data source is parameterized, you can add variables to the footer labels. Variables display the current parameter value in the selected footer label(s) and change dynamically according to the selected parameter value.



Footer Labels Variables

To add a variable, click on the *Variables* button in the *Footer Labels* dialog and select a parameter from the list. It will be displayed as &<*parameterName*> in the textfield.

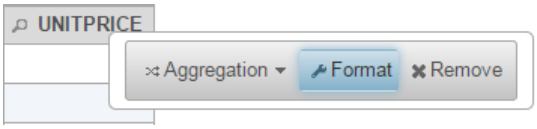
rus Table	1 354	12	16 248
		Order Total	20 190
	Total Sales from 2001-01-14 to 2003-12-09		2 850 764

Report Footer Label with Variables

If you add a variable to the label, the text may not fit the cell. In this case it is advisable to use text alignment.

5.3.7. Data Formating

Right click on the column header and select *Format* from the pop-up menu.

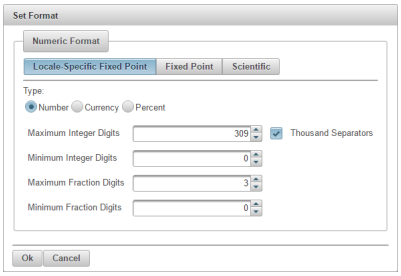


Format Button

When you select the *Format* option, the *Set Format* dialog will appear. The *Set Format* dialog is dependent on what type of data is present in the selected column: numeric, string, date/time, or logical/boolean.

Formatting Numeric Data: The dialog for numeric data contains three primary options for the data: *Locale-Specific Fixed Point*, *Fixed Point*, and *Scientific*. Each option opens a different tab.

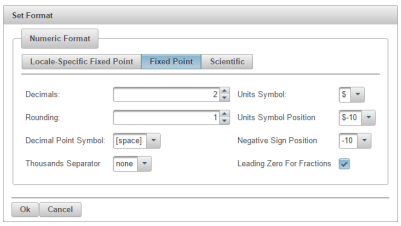
Locale-Specific Fixed Point: This will change the format of the data depending on the locale in which it is being viewed. Additional formatting for this option allows you to specify whether the data should be displayed as a number, currency, or percentage. In addition, you can set the maximum and minimum number of integer digits and fraction digits. Other display attributes will vary depending on locale.



Locale-Specific Formatting

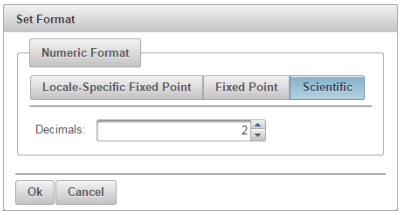
Fixed Point: This will keep the data format consistent, regardless of locale. Additional formatting for this option allows you to set the number of decimals, rounding for digit number, unit symbols, negative sign position, decimal

and thousands separator, and specify leading zeroes for fractions.



Fixed Point Formatting

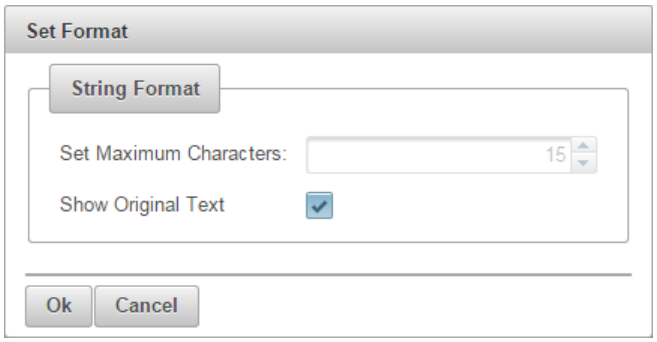
Scientific: This will display the data in scientific notation. Additional formatting for this option allows you to set the number of decimals.



Scientific Formatting

After you finish selecting options, click on the *Ok* button to apply formatting and close the *Set Format* dialog.

Formatting String Data: In this dialog, you can change maximum characters for the string format. To do this, uncheck the *Show Original Text* option and then set maximum characters.

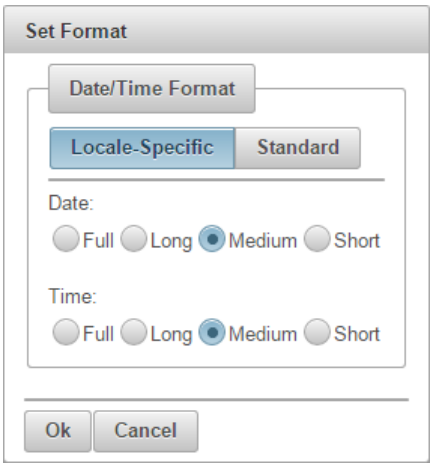


String Data Format

Click on the *Ok* button to apply changes and close the *Set Format* dialog.

Formatting Date/Time Data: The dialog for date/time data contains two tabs: *Locale-Specific* and *Standard*.

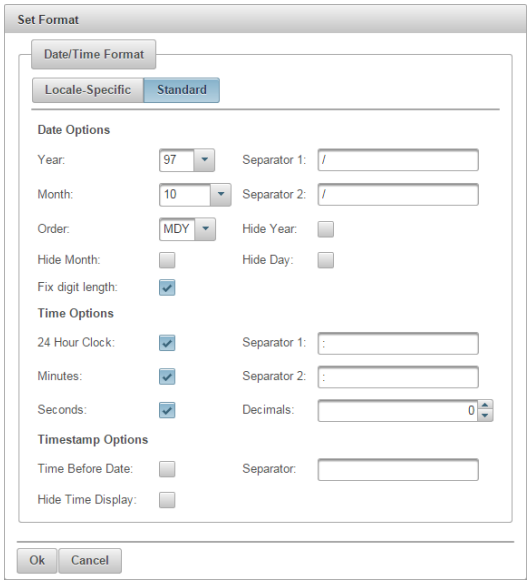
Locale Specific: This will change the format of the data depending on the locale in which it is being viewed. Additional formatting for this option allows you to select full, long, medium, or short notations for date and time information. Other display attributes will vary depending on locale.



Locale-Specific Formatting

Standard:

This will keep the data format consistent, regardless of locale. Additional formatting for this option allows you to select year and month displays, as well as the order in which month day and year information is presented. You can also select which characters should be used as separators. Time options allow you to display hours, minutes, and/or seconds, and select the separators between them. For timestamp data, you can select to display the time before or after the date and the separator to be used between them.

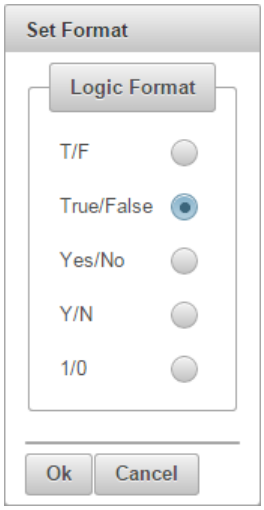


Standard Formatting

After you finish selecting options, click on the *Ok* button to apply formatting and close the *Set Format* dialog.

Formatting Logical Data:

The dialog for Logical or Boolean data contains five options for displaying the data: T/F, True/False, Yes/No, Y/N, and 1/0.

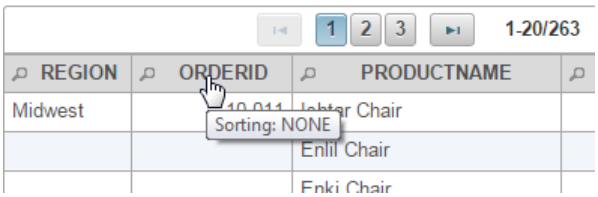


Logic Format

Select a format you want to use and then click on the *Ok* button to change the data format and close the *Set Format* dialog.

5.3.8. Sorting

When you hover your mouse cursor over the column header, you will see a label *Sorting: NONE*. Left click on the column header to apply the sorting. First click causes ascending sorting, second click causes descending sorting and third click cancels the sorting.



Sorting

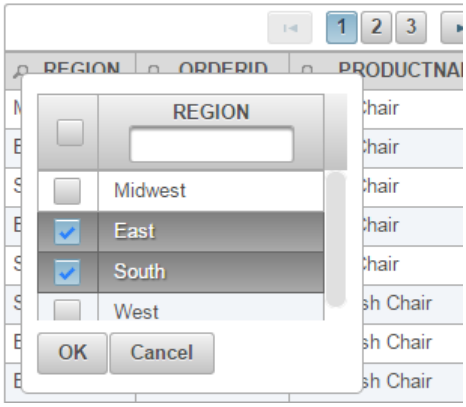
Sorted columns have arrow icons on the column header. The *Up Arrow* icon indicates ascending order and the *Down Arrow* icon indicates descending order.



Sorting Arrow

5.3.9. Data Filtering

You can filter data by clicking on the  *Filter Column* icon on the column header. A pop-up dialog will appear, which will allow you to select data for filtering. To select values which you want to be displayed in the report, check off the box next to the desired value(s). You can narrow down the values shown in the *Filter* dialog or search for particular values by typing them into the text field at the top of the dialog. To select all available values within the column, check the box in the upper left corner (next to the search field). Click on the *OK* button to apply the filtering and close the dialog.

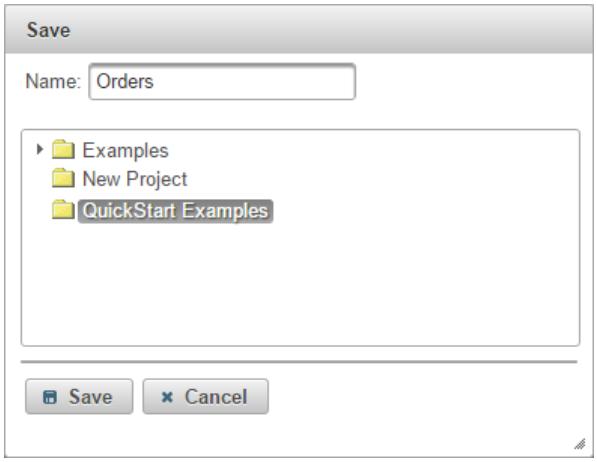


Filter Column

If there are several columns applying data filters, the report will display all filter conditions, i.e. different column filters has “AND” relationship.

5.3.10. Save the Report

You can save the report by clicking on the *Save* button on the toolbar  . This will bring up a dialog allowing you to specify a name for the report. Enter a name for the report. Click on the *Save* button to save the report.

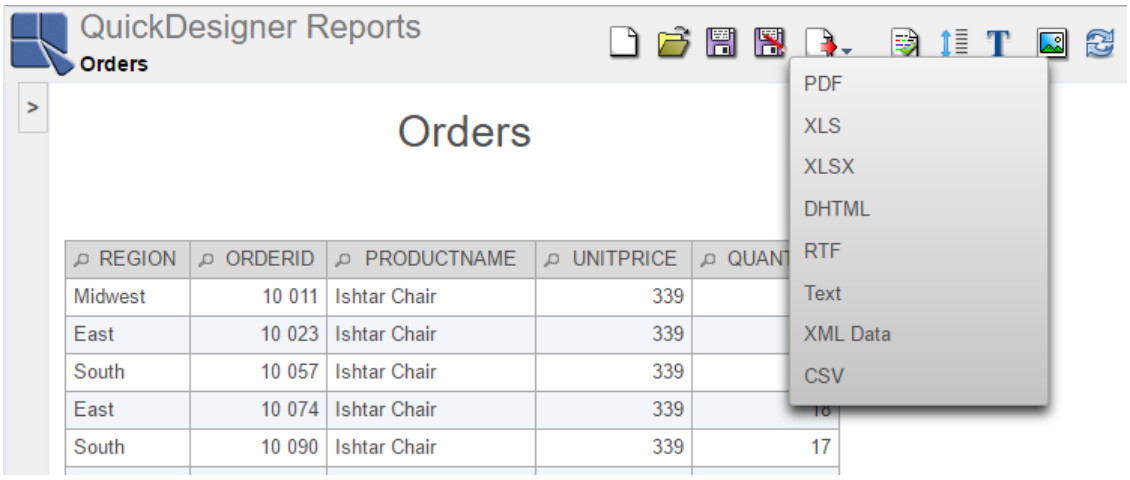


Save Dialog

After saving or opening existing report, the  *Save As* button allows you to save the existing or modified report under a different name.

5.3.11. Export the Report

You can also export the report in several file formats. To do so, click on the  *Export* icon on the toolbar. A pop-up dialog will appear, prompting you to specify the export format. Available options are PDF, XLS, XLSX, DHTML, RTF, Text, XML Data, and CSV.

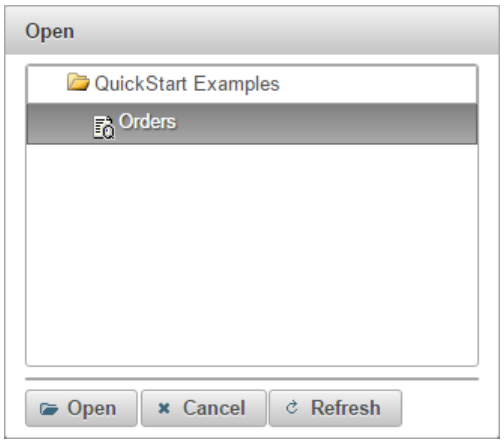


Export Options

Select the format and then you can save the generated file to your local system.

5.3.12. Open the Saved Report

You can open the saved report by clicking on the *Open* icon on the main toolbar. The *Open* dialog will appear.



Open the Report

Here you can see all reports from Organizer created in QuickDesigner Reports. Select a report and click on the *Open* button to open it.

5.3.13. Exit

You can exit QuickDesigner Reports either by clicking on the *Home* icon in the upper right corner, or by clicking on the *QuickDesigner Reports* title, or by clicking on the *Logo* icon in the upper left corner. Before closing, you will be prompted if you want to save an unsaved report.

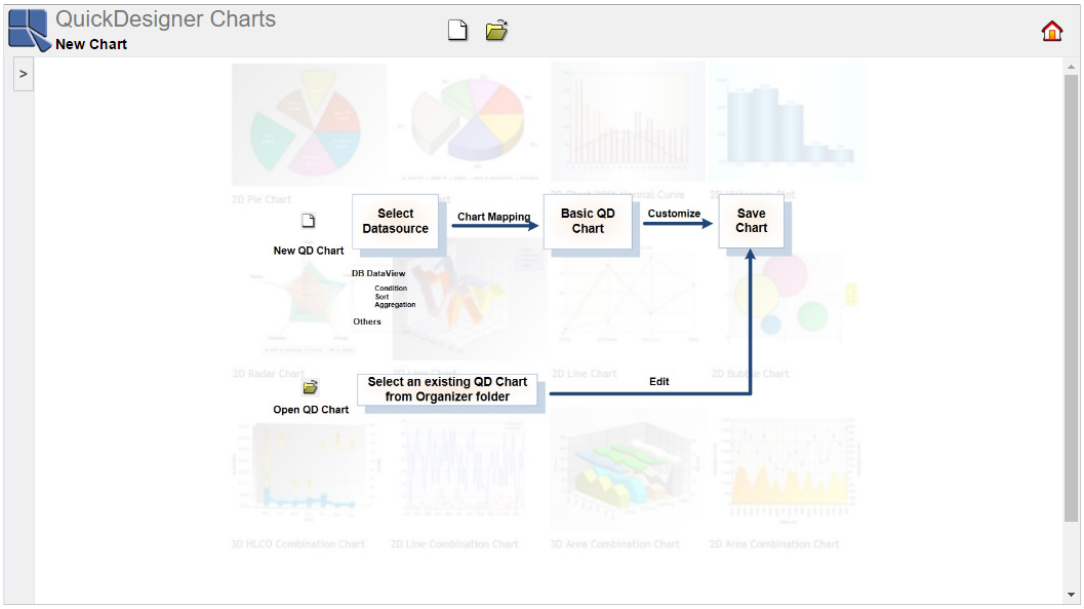
5.4. QuickDesigner Charts

QuickDesigner Charts is a thin-client ad-hoc charting interface. It provides users a scaled-down design tool to create charts. With QuickDesigner Charts, end users can easily select, filter, and present data without mastering database structures, all with zero client download.

5.4.1. Start

QuickDesigner Charts can be started from the ERES Start page. If you have logged in as a user with design privileges, you can follow QuickDesigner Charts link to start using QuickDesigner Charts.

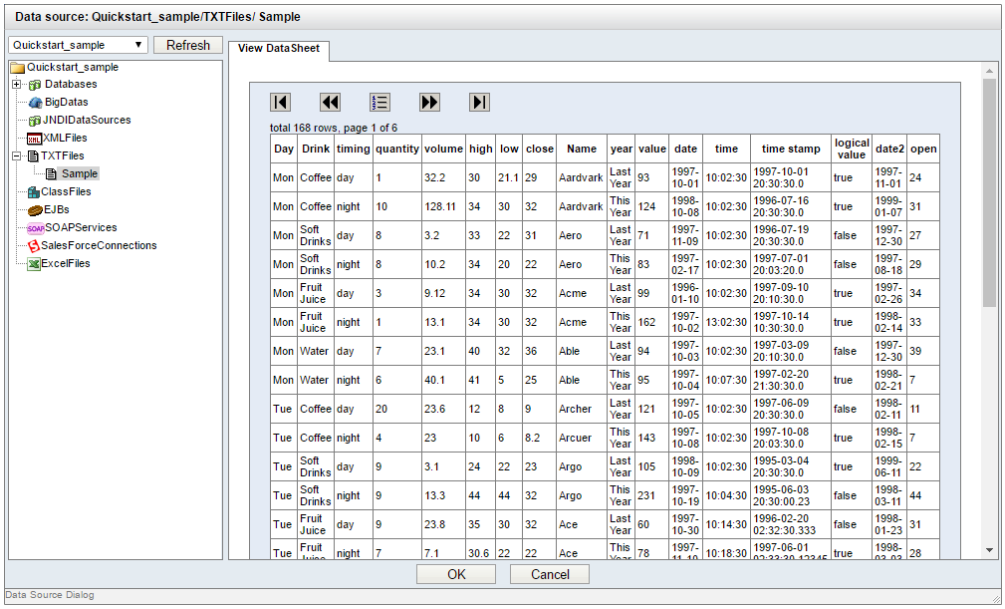
When QuickDesigner Charts launches, the first page that appears prompts you to select between opening an existing chart or creating a new one.



QuickDesigner Charts Start Options

5.4.2. Select a Data Source

If you want to create a new chart, click on the  *Create New Chart* icon. The *Data Source Dialog* appears and prompts you to select data source you want to use with QuickDesigner Charts.



Day	Drink	timing	quantity	volume	high	low	close	Name	year	value	date	time	time stamp	logical value	date2	open
Mon	Coffee	day	1	32.2	30	21.1	29	Aardvark	Last Year	93	1997-10-01	10:02:30	1997-10-01 20:30:30.0	true	1997-11-01	24
Mon	Coffee	night	10	128.11	34	30	32	Aardvark	This Year	124	1998-10-08	10:02:30	1996-07-16 20:30:30.0	true	1999-01-07	31
Mon	Soft Drinks	day	8	3.2	33	22	31	Aero	Last Year	71	1997-11-09	10:02:30	1996-07-19 20:30:30.0	false	1997-12-30	27
Mon	Soft Drinks	night	8	10.2	34	20	22	Aero	This Year	83	1997-02-17	10:02:30	1997-07-01 20:03:20.0	false	1997-08-18	29
Mon	Fruit Juice	day	3	9.12	34	30	32	Acme	Last Year	99	1998-01-10	10:02:30	1997-09-10 20:10:30.0	true	1997-02-26	34
Mon	Fruit Juice	night	1	13.1	34	30	32	Acme	This Year	162	1997-10-02	13:02:30	1997-10-14 10:30:30.0	true	1998-02-14	33
Mon	Water	day	7	23.1	40	32	36	Able	Last Year	94	1997-10-03	10:02:30	1997-03-09 20:10:30.0	false	1997-12-30	39
Mon	Water	night	6	40.1	41	5	25	Able	This Year	95	1997-10-04	10:07:30	1997-02-20 21:30:30.0	true	1998-02-21	7
Tue	Coffee	day	20	23.6	12	8	9	Archer	Last Year	121	1997-10-05	10:02:30	1997-06-09 20:30:30.0	false	1998-02-11	11
Tue	Coffee	night	4	23	10	6	8.2	Arcuer	This Year	143	1997-10-08	10:02:30	1997-10-08 20:03:30.0	true	1998-02-15	7
Tue	Soft Drinks	day	9	3.1	24	22	23	Argo	Last Year	105	1998-10-09	10:02:30	1995-03-04 20:30:30.0	true	1999-06-11	22
Tue	Soft Drinks	night	9	13.3	44	44	32	Argo	This Year	231	1997-10-19	10:04:30	1995-06-03 20:30:00.23	false	1998-03-11	44
Tue	Fruit Juice	day	9	23.8	35	30	32	Ace	Last Year	60	1997-10-30	10:14:30	1996-02-20 02:32:30.333	false	1998-01-23	31
Tue	Fruit Juice	night	7	7.1	30.6	22	22	Ace	This Year	78	1997-11-10	10:18:30	1997-06-01 02:32:30.12345	true	1998-03-03	28

Data Source Dialog

You can see the records of the data source in *View DataSheet* pane on the right. If you select DataView as the data source, you will be able to create a new DataView query in the *DataView Builder* in the right pane. You can

also select an existing DataView query and modify it in the right pane. For more about managing data sources in QuickDesigner Charts, please see Section 4.2 - Data in QuickDesigners and Maps.

Once you have finish selecting a data source, click on the *OK* button to close the *Data Source Dialog*.

5.4.2.1. Change a Data Source

After you return to the main QuickDesigner Charts window (after you created a new chart or opened an existing one), you can change the chart's data source by clicking on the  *Change Data Source* button on the toolbar. This will take you back to the *Data Source Dialog* and allow you to pick a new data source.

5.4.3. Data Mapping and Ordering

Once you select a data source or finish building a query and close the *Data Source Dialog* by clicking on the *OK* button, you will see *Chart Mapping* dialog in the left pane of QuickDesigner Charts.

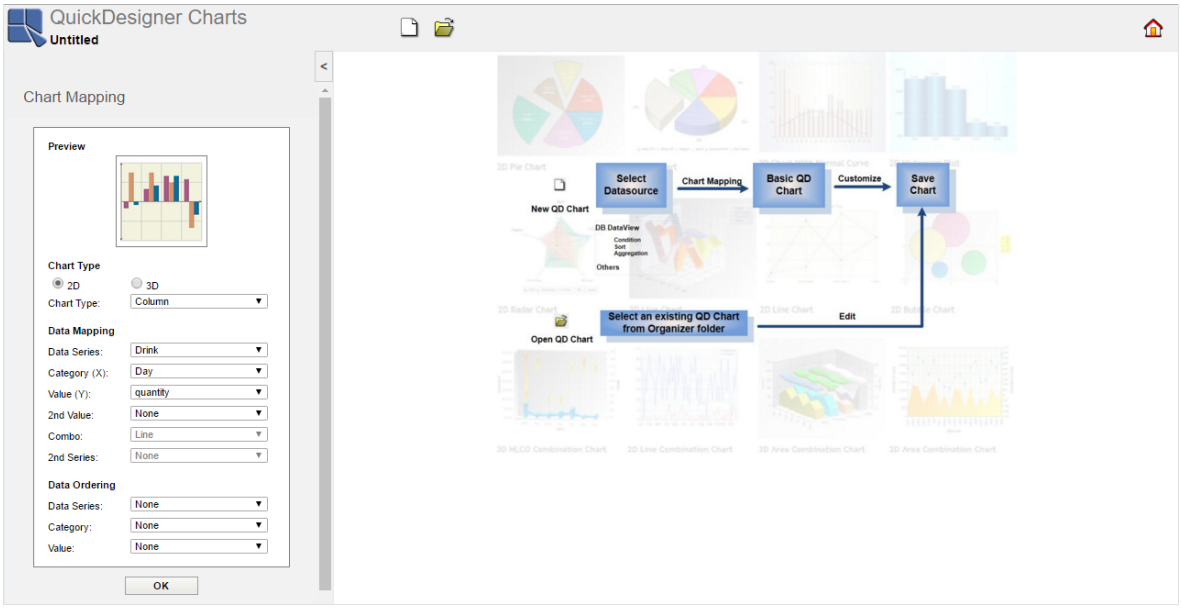


Chart Mapping Dialog

5.4.3.1. Mapping and Ordering Options

The *Chart Mapping* dialog allows you to select which chart type you want to create and which fields from the data source you want to map to the chart elements.

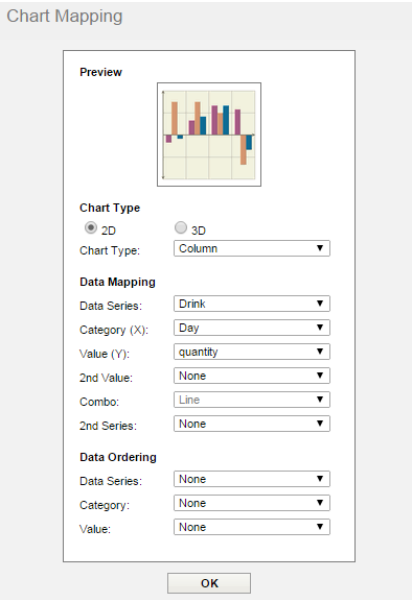
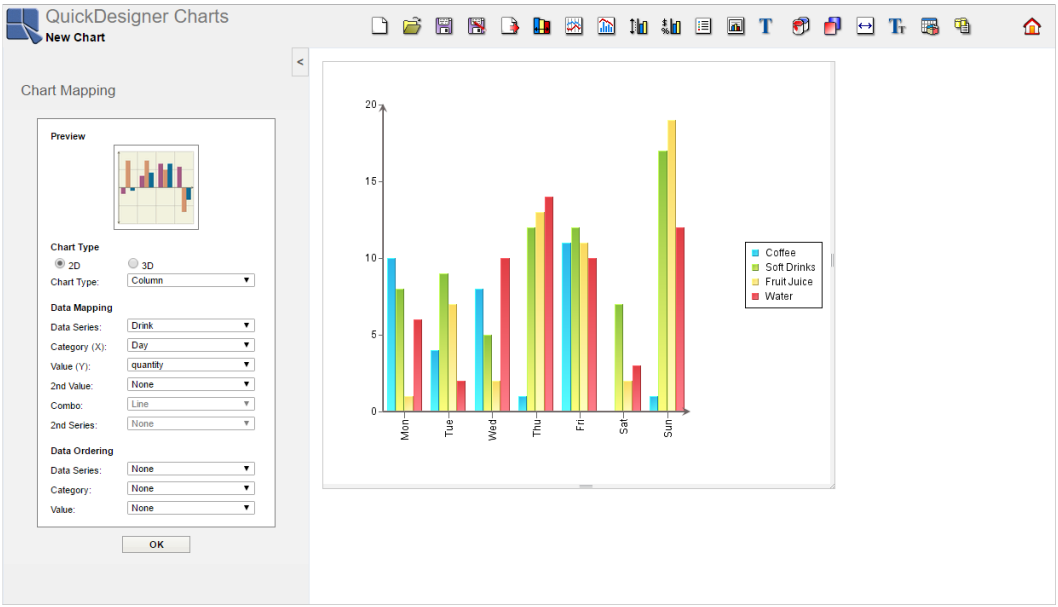


Chart Mapping Options

The *Chart Type* options allows you to select between two-dimensional and three-dimensional chart types, and the chart type you want to use. The *Data Mapping* options allows you to select data mapping for the chart. The *Data Ordering* options allows you to sort data in ascending or descending order. For a detailed description of all available charts and mapping options, see Section 5.2.3 - Chart Types and Data Mapping.

Once you select the type and mapping options, click on the *OK* button and the chart will appear in the right pane, as well as all icons.



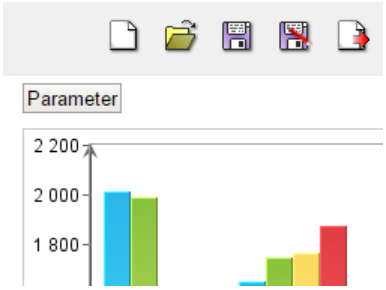
QuickDesigner Charts Main Page

5.4.3.2. Change Data Mapping

You can change the mapping options and the chart type by clicking on the  *Change Mapping/Chart Type* icon on the toolbar. This will open the *Chart Mapping* dialog in the left pane.

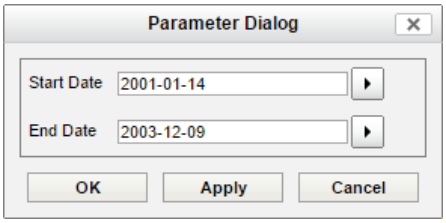
5.4.3.3. Parameter Setting

If the data source is parameterized, you can see *Parameter* button above the chart.



Parameter Button

After clicking on it, the *Parameter Dialog* will open allowing you to set parameter values. To apply a parameter value, click on the *Apply* button (the dialog stays open) or *OK* button (the dialog will be closed).



Parameter Dialog

5.4.4. Toolbar

All of the formatting actions for charts in QuickDesigner Charts are accessed through the toolbar. The icons perform the following actions:


Create a new chart


Open an existing chart


Save the chart


Save chart as


Export the chart


Set data properties for the chart


Chart-specific options (This option is only available for Pie, HLCO, Dial, Gantt, Polar, and Doughnut chart.)


Line and point settings


Trend line settings


Format axis scale

 Format axis elements

 Format chart legend

 Format plot area

 Add/Edit chart titles

 Set chart colors

 Set color gradient

 Set chart canvas size

 Set label/legend font

 Change chart type/mapping

 Change chart data source

5.4.5. Data Formatting

Once you finish specifying all the mapping options, you can format chart elements.

5.4.5.1. Position and Size

There are several options available in QuickDesigner Charts that allows you to control the size and position of the chart.

5.4.5.1.1. Canvas Size

The chart canvas size can be adjusted by clicking on the  *Canvas Size* button on the toolbar. This will bring up a dialog in the left pane allowing you to specify the size of the chart.

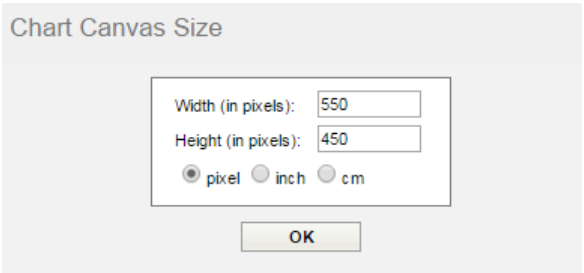
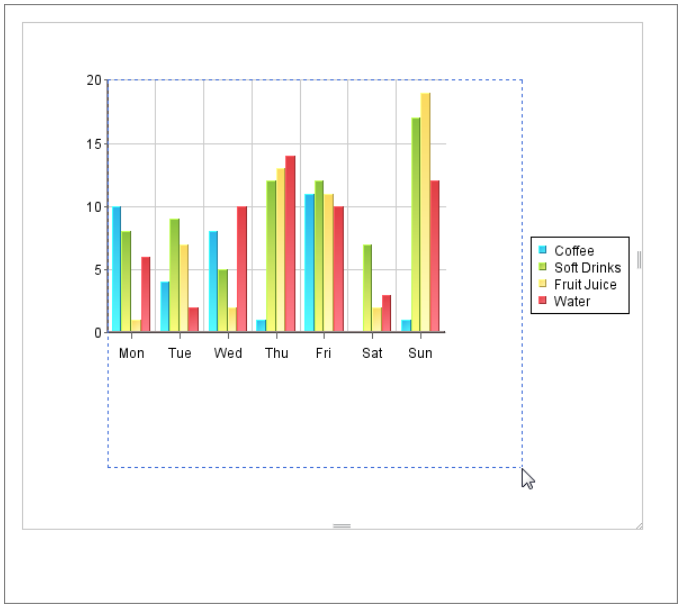


Chart Canvas Dialog

The dialog allows you to specify the width and height of the canvas in pixels, inches, or centimeters. Once you specify the new dimensions, click on the *OK* button and all changes will be applied. You can also change the canvas size freely using drag & drop. In this case, you can see the exact size in the *Chart Canvas Size* dialog.

5.4.5.1.2. Chart Position and Size

You can set the position of the chart within the canvas freely using drag & drop. You can set the canvas size or plot area size by dragging the sides or dragging the corner.



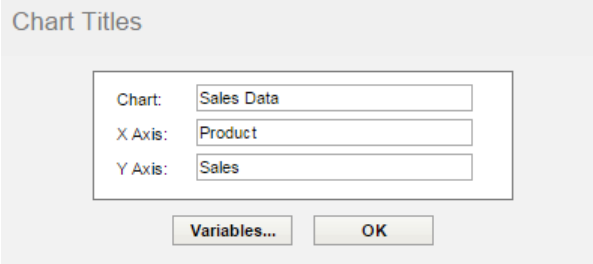
Setting Size

5.4.5.2. Font and Color Options

QuickDesigner Charts interface allows you to add/edit chart titles, set title/label font, and chart colors.

5.4.5.2.1. Chart Titles

To add or edit the chart titles, click on the T *Add/Edit Chart Titles* button on the toolbar. This will bring up a dialog in the left pane.

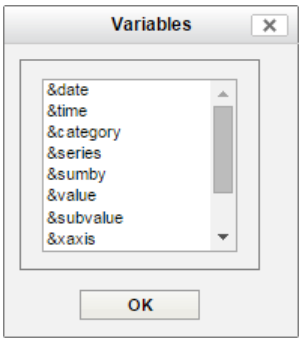


Insert/Edit Titles Dialog

You can add a main title to the chart, as well as one for each of the chart axes. Once you finish adding titles, click on the *Ok* button to apply the changes.

5.4.5.2.1.1. Text Variables

As with the Chart Designer, you can also add text variables that will provide run-time substitution based on certain values in the chart. To bring up a list of variables, click on the *Variables* button in the *Chart Titles* dialog. The variables list will open in a new window.



Text Variables Dialog

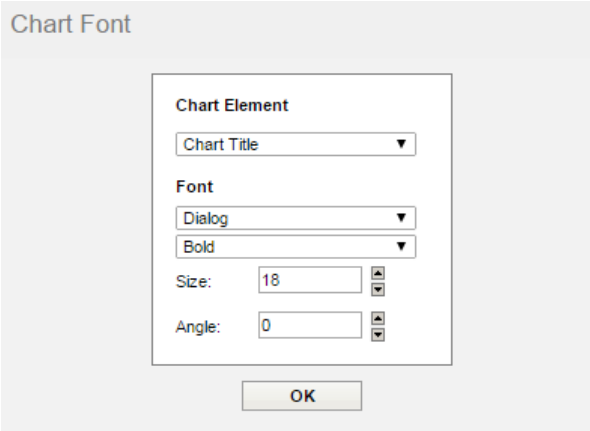
The following text variables are supported:

&date	This variable displays the date when the chart was last drawn/redrawn.
&time	This variable displays the time when the chart was last drawn/redrawn.
&category	This variable displays the name of the category column.
&series	This variable displays the name of the data series column.
&sumby	This variable displays the name of the sum-by column.
&value	This variable displays the name of the value column.
&subvalue	This variable displays the name of the secondary value column.
&xaxis	This variable displays the name of the column that is mapped to X-axis. This is for charts that map a value instead of a category to the X-axis like scatter or bubble charts.
&yaxis	This variable displays the name of the column that is mapped to Y-axis. This is for scatter, bubble, and surface charts.
&zaxis	This variable displays the name of the column that is mapped to Z-axis. This is for scatter, bubble, and surface charts.
&2ndaxis	This variable displays the name of the column that is mapped to 2nd-axis.
&paramInfoValue	If the chart contains a parameter, this displays the parameter value(s) that were selected.
&paramInfoName	If the chart contains a parameter with this name, this will display the name of that parameter.

To add a text variable to a title, select it from the list and click on the *OK* button.

5.4.5.2.2. Chart Fonts

To set the font for chart labels and titles, click on the *Font Settings* button on the toolbar. This will open a dialog in the left pane.



Font Settings Dialog

To set the font for a title or label, select it in the drop-down list at the top of the dialog. Then select font face, style, size, and angle using the options in the lower dialog. Once you finish setting up fonts, click on the *OK* button to apply the changes.

5.4.5.2.3. Chart Colors

You can set color for all chart elements by clicking on the  *Set Chart Colors* button on the toolbar. The *Chart Colors* dialog will open in the left pane.

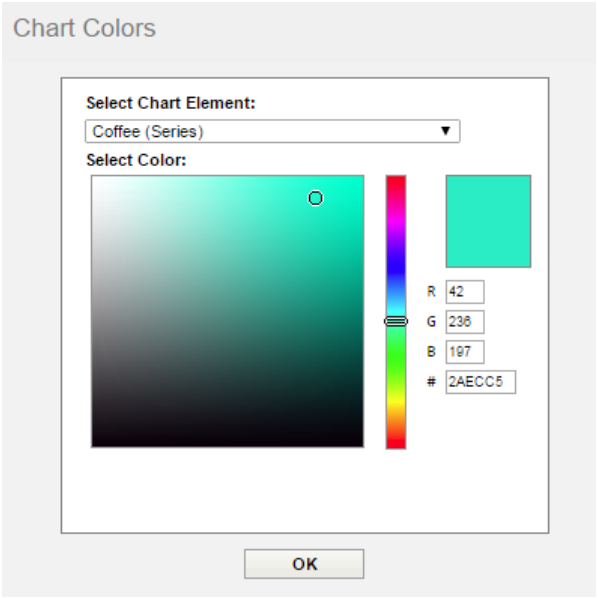


Chart Colors Dialog

To set the color for a chart element, first select it from the drop-down list at the top of the dialog. Then select a new color either by clicking on one of the swatches, entering RGB values for the color, or entering HEX code to the # field. Once you finish setting up the colors, click on the *OK* button to apply the changes.

Please note that some elements are not visible by default (e.g. data border). In this case, first make the element visible (see Section 5.4.5.4.1 - Data Properties) and then use *Chart Colors* dialog for setting a color of the element.

5.4.5.2.4. Color Gradient

You can set the color gradient for the chart by clicking on the  *Set Color Gradient* button on the toolbar. The *Color Gradient* dialog will open in the left pane.

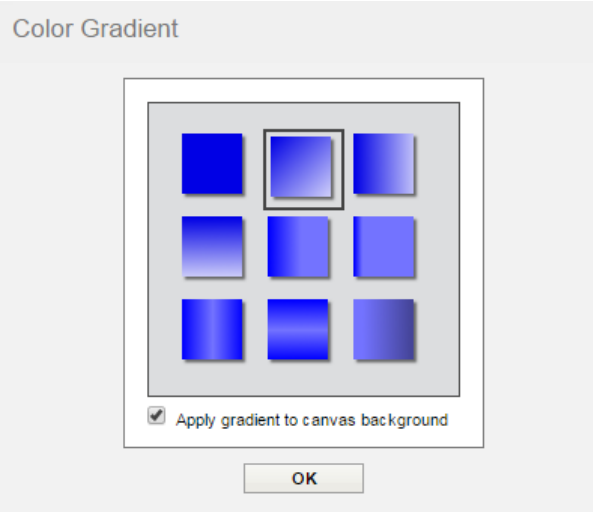


Chart Colors Dialog

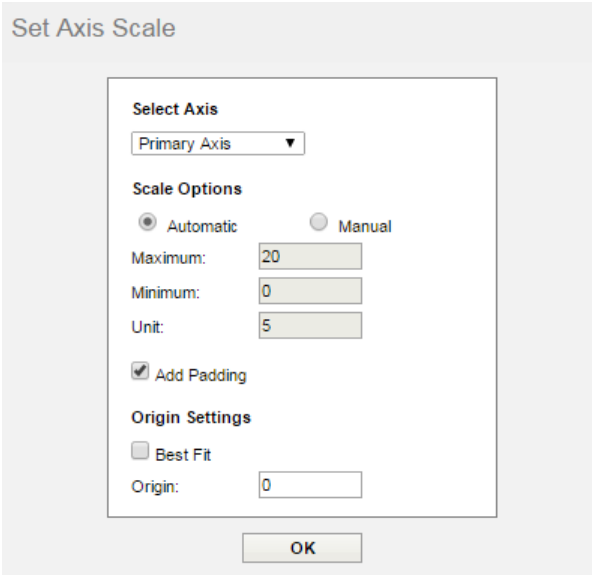
Select a gradient you want to use and click on the *OK* button to apply the changes. You can apply the same gradient to the canvas background.

5.4.5.3. Chart Axes

In QuickDesigner Charts, you have several formatting options for the chart axes. You can customize the axis scale as well as label/axis appearance options.

5.4.5.3.1. Axis Scale

By default, the scale of any value axes in the chart is calculated to provide the *best fit* for the plotted data. This is especially useful if the data change often. However, there are many times when you may want to set the scale of the axes manually. To modify the axis scale, click on the  *Axis Scale* button on the toolbar. A new dialog will open in the left pane.



Axis Scale Dialog

You can select the axis for which you would like to set the scale by selecting it in the drop-down list at the top of the dialog. The following scaling options are provided for each value axis:

Automatic: This option turns on automatic scaling for the axis. This is the default setting.

Manual: This option turns on manual scaling to set the axis scale to the options provided.

Maximum: This is the highest value on the axis scale.

Minimum: This is the lowest value on the axis scale.

Unit: This is the step size of the axis scale.

Add Padding: This option will add padding to the top of the automatic scale. If this option is disabled, the maximum scale will match the largest data point.

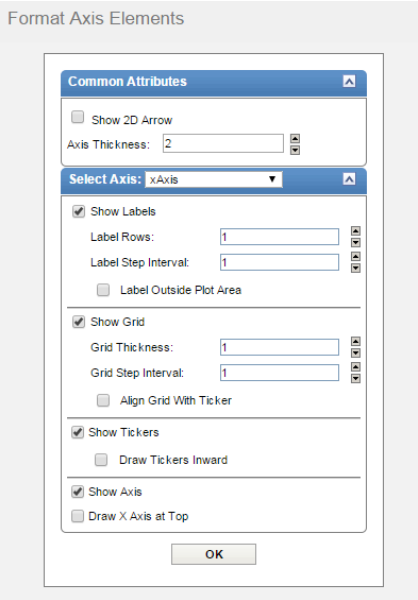
Best Fit: This option will automatically place the origin of the chart based on minimum and maximum values of the data.

Origin: This option allows you to specify where the X and Y axes should intersect. This is usually at zero.

Once you finish specifying the axis scale options, click on the *OK* button to apply the changes.

5.4.5.3.2. Axis Elements

Additional appearance properties for the axes and axis labels can be set in the *Format Axis Elements* dialog. To change these settings, click on the  *Axis Options* button on the toolbar. This will open the dialog in the left pane.



Axis Elements Dialog

To change the settings for an axis, first select it in the drop-down *Select Axis* list and then set the available options below. The dialog allows you to perform the following options. Note that some options are only available for certain chart types, certain data types, and on certain axes.

Show 2D Arrow: This option allows you to remove or display arrowhead at the end of the axis. Note that this option applies to all chart axes and is only available for two-dimensional charts.

Axis Thickness: This option allows you to set the thickness of the axis in pixels. Note that this setting is only available for two-dimensional charts and it is applied to all axes in the chart.

Show Labels: This option allows you to remove or display the labels for each axis.

Label Rows: This option allows labels to be displayed in alternating rows. This can prevent overlapping. This option is only available for the X-axis.

Label Step Interval: This option allows you to set the label step interval for the data. For example setting this to 2 would draw the label for every other data point in the chart.

Label Outside Plot Area: This option sets the labels to be placed outside of the plot area, regardless of where the axis is. This feature can be useful for category axis labels if you are plotting data with both positive and negative values.

Label Interval Unit: This option allows you to select the unit to be used when sorting and representing time-based data (date, time, or timestamp). Selecting tickers will use the data as it is read by Chart Designer.

Show Grid: This option allows you to remove or display the grid for each axis.

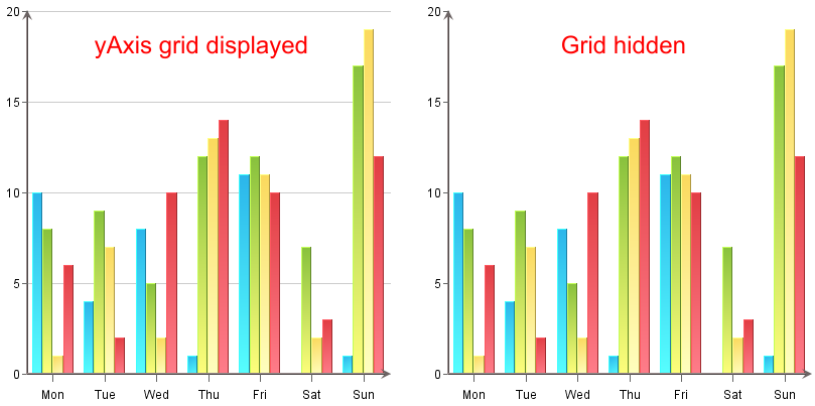


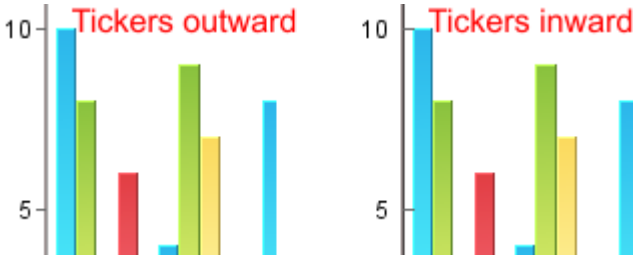
Chart With and Without Grid

Grid Thickness: This option allows you to specify the thickness of any grid lines along the axis.

Grid Step Interval: This option allows you to set the grid step interval for any grid lines along the axis.

Align Grid With Ticker: This option aligns the grid line with the ticker instead of placing it between tickers. This places the ticker and corresponding grid line along the same line. This option only applies to the category axis of column-type charts.

Show Tickers: This option allows you to remove or display the axis tickers and draw them inward or outward.



Tickers Drawn Outward and Inward

Draw Tickers Inward: This option will draw the axis tickers inside the plot area instead of outside (default).

Show Axis: This option allows you to remove or display the axis (for two-dimensional charts) or the wall (for three-dimensional charts).

Draw X Axis at Top: This option allows the X-axis to be positioned at the top of the chart instead of the default bottom position. This option is available for two-dimensional column, bar, scatter, high-low, HLCO, bubble, and Gantt charts.

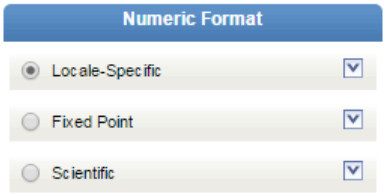
Swap Y Axis Position: This option will swap the primary and secondary value axes. This option can only be found under 2nd Axis tab.

5.4.5.3.2.1. Axis Label

The *Format Axis Elements* dialog also allows you to format the appearance of the axis labels, depending on what type of data is plotted on the axis.

- **Formatting Numeric Data:**

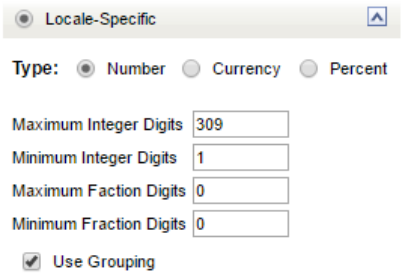
  For numeric data, there are three primary options for display formatting: locale-specific fixed point, fixed point, and scientific. Additional options will be displayed when you click on the *Format* button.



Numeric Data Format Options

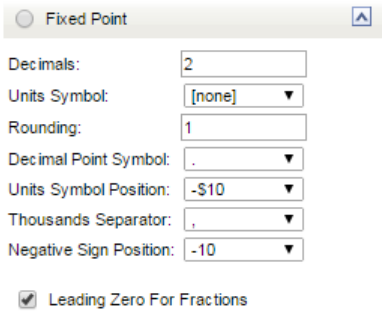
Locale-Specific: This will change the format of the data depending on the locale in which it is being viewed. Additional formatting for this option allows you to specify whether the data should be displayed as a number, currency, or percentage. In addition, you can set the maximum and minimum number of integer digits and fraction digits. *Use Grouping* brings space as the thousands separator. Other display attributes will vary depending on locale.

Please note that if you select *Currency* as *Type*, local-specific format will automatically take currency symbol from your system settings (e.g. it will add "$" to the label if you're in U.S.A.).



Locale-Specific Formatting Options

Fixed Point: This will keep the data format consistent, regardless of locale. Additional formatting for this option allows you to set the number of decimals, rounding for digit number, unit symbols, negative sign position, decimal and thousands separator, and to use *Leading Zero For Fractions* which determines whether it should display 0 before decimals for numbers from -1 to 1 (e.g. 0.15/ .15).



Fixed Point Formatting Options

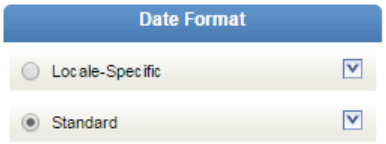
Scientific: This will display the data in scientific notation. Additional formatting for this option allows you to set the number of decimals.



Scientific Formatting Options

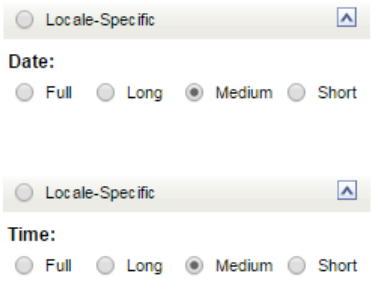
- **Formatting Date/Time Data:**

  For date/time data, there are two primary options for display formatting: locale specific and standard. Additional options will be displayed if you click on the *Format* button. The available options will vary depending on the nature of your data. Date, time, and timestamp data will bring up date, time, and date & time options.

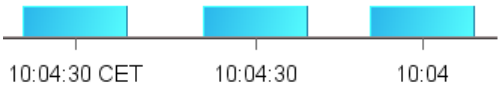


Date/Time Data Format Options

Locale-Specific: This option will change the format of the data depending on the locale in which it is being viewed. Additional formatting for this option allows you to select full, long, medium, or short notations for date, and time information. Other displayed attributes will vary depending on your locale.

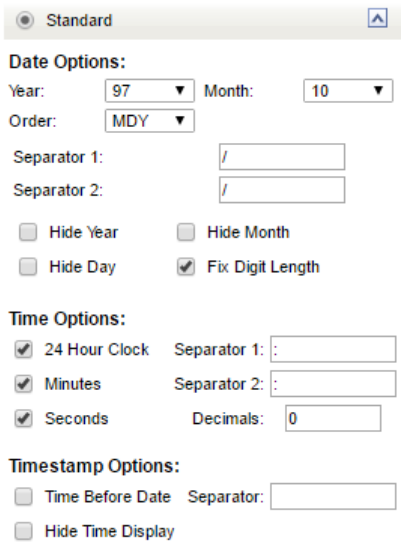


Locale-Specific Formatting Options



Example of Long/Medium/Short Time Format

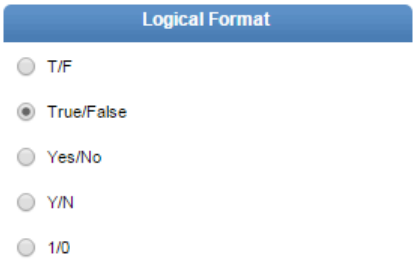
Standard: This option will keep the data format consistent, regardless of locale. Additional formatting for this option allows you to select year and month displays, as well as the order in which month, day, and year information is presented. You can also select the characters to be used as separators. Time options allow you to display hours, minutes, and/or seconds, and select the separators between them. For timestamp data, you can select to display the time before or after the date and the separator to be used between them. *Fix Digit Length* option adds zeros to days and months between 1 – 9 to keep the length of the date consistent (two digits for both day and month).



Standard Formatting Options

- **Formatting Logical Data:**

  There are five options available for displaying logical or Boolean data: T/F, True/False, Yes/No, Y/N, and 1/0.



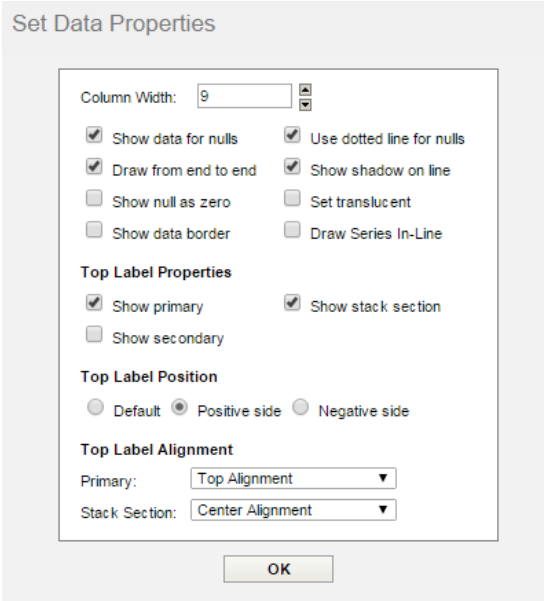
Logical Formatting Options

5.4.5.4. Plot/Data Elements

In QuickDesigner Charts, there are many ways to configure how data points and chart plot should be drawn.

5.4.5.4.1. Data Properties

Many of the data display options are controlled through the data properties dialog. From this dialog, you can control the size of bars/columns, set display options for null values, and specify options for data labels. To invoke this dialog, click on the  *Data Properties* button on the toolbar. The dialog will open in the left pane.

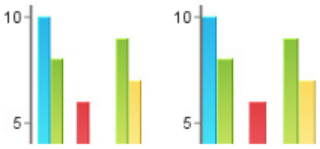


Data Properties Dialog

This data properties dialog contains the following options:

Column Width:

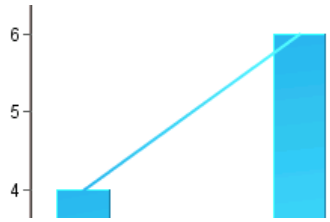
This option specifies the ratio of the bar/column width relatively to the gap between successive bars in the chart. Each unit represents 1/10th of the space between data points. Therefore, entering **9** would leave 10% of the space between data points blank, while **10** would eliminate all space between bars/columns. This option only pertains to two-dimensional bar, column, stack bar, stack column, high-low, HLCO, and Gantt charts.



Column width 9 vs 10

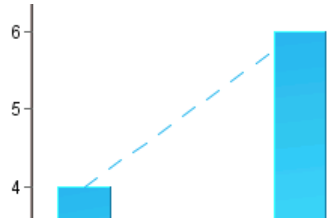
Show data for nulls:

This option will connect null data with lines. For example, if you have three points and the value of point 1 is 4, point 2 is null, and the value of point 3 is 6, then a line will be drawn from 4 to 6 for the three points. This option is only available for line charts or other two-dimensional charts with lines. All other chart types will not plot null data.



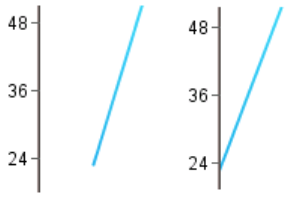
Connected line with null data used.

Use dotted lines for nulls: Instead of drawing a whole line through null values in the chart, you can use this option to draw a dotted line. Like the show data for nulls option, this option is only available for lines.



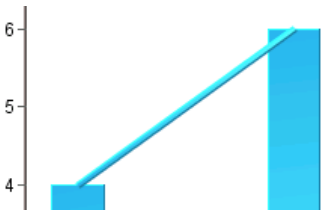
Dotted line with null data used.

Draw from end to end: This option allows you to draw 2D line and area charts across the entire plot area, rather than offsetting to the first and last data points on the chart.



End to end disabled/enabled.

Draw shadow for lines: This option specifies whether to draw shadow on 2D lines. In order for shading to apply, the line must be thicker than one pixel.



Connected line with shadow.

Show null as zero: This option will display *0* for null data.

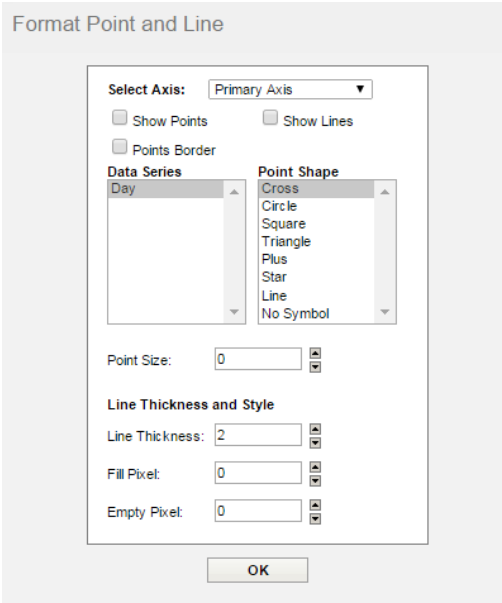
Set translucent: This option will draw all data points in translucent color, allowing overlapped elements to be visible. This is useful for radar charts and two-dimensional area charts with data series. This option is also available for bubble and Gantt charts. Note that because this feature requires Java 21 or higher, the translucent elements will not display if you are deploying charts in applets and the client is not using the Java plugin.

Show data border:	This option will draw a border around columns and areas.
Draw Series In-Line:	This option will display series in one line. It only pertains to three-dimensional column, stack column, percent column, and bar charts.
Show primary:	This option will display data top labels for the primary values in the chart.
Show secondary:	This option will display data top labels for the secondary values in the chart.
Show stack section:	This option will display individual labels for each stack section for stack bar, stack column, and stack area charts.
Top label position:	This option allows you to specify where the top labels should be drawn. By default, they are drawn above data points if they are positive and below data points if they are negative. Using these options, you can force the data points to always draw to the positive or negative side of the data.
Label Alignment:	This option allows you to set the alignment for the data top labels. You can draw them at the top, bottom, or middle of the data points. In addition, you can select to draw the label inside the data point at the top or bottom. An additional option for stack charts allows you to set the alignment for stack section labels.

Once you finish setting up the options, click on the *OK* button to apply the changes.

5.4.5.4.2. Line And Point

For any two dimensional chart type, you can choose to display lines and points for all data points in the chart. Note that some chart types already use this representation (i.e. line or scatter charts). Line and point settings can be adjusted by clicking on the  *Line and Point Settings* button on the toolbar. The dialog will open in the left pane.



Line and Point Dialog

The first three options allows you to specify whether you want to show lines, points, and points border for the chart. For radar charts, you also have the option of showing areas. The remaining options allows you to customize the line and point displays for each element in the data series.

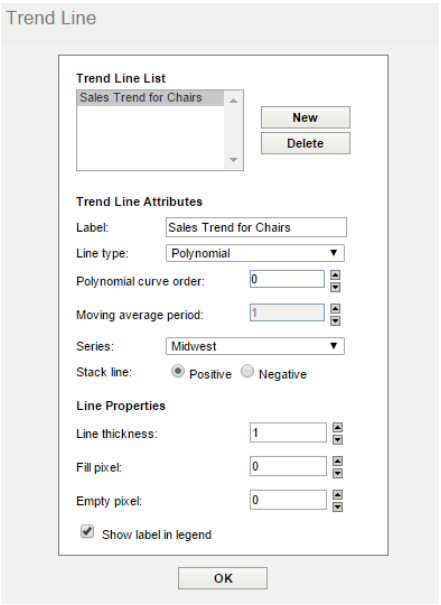
For each data series element, you can specify the point shape you want to use. You can also control the size of the points. The default point size is 0. You can specify point sizes of -1, -2, & -3 which represent sizes of 0.75, 0.5, and 0.25. At -3 (0.25), the point will be drawn as a dot regardless of the selected point shape.

For lines you can specify the line thickness as well as customize a dash pattern. The dash pattern is created by specifying the number of filled pixels and empty pixels (between 0 and 255). The line is then drawn by dividing into segments - the number of filled pixels followed by the number of empty pixels. Setting 0 for both will result in a solid line. Setting 255 for both will result in no line being drawn.

Once you finish setting up the options, click on the *OK* button to apply the changes.

5.4.5.4.3. Trend Line

A powerful feature in QuickDesigner Charts is the ability to add trend lines to charts. Trend lines can help you show more details of a chart's data by exposing and highlighting certain trends. Trend lines settings can be adjusted by clicking on the  *Trend Line* button on the toolbar. The dialog will open in the left pane.



Trend Line Dialog

There is a list of existing trend lines at the top of the dialog. You can select existing trend line and edit it, remove the selected trend line, and/or create a new trend line.

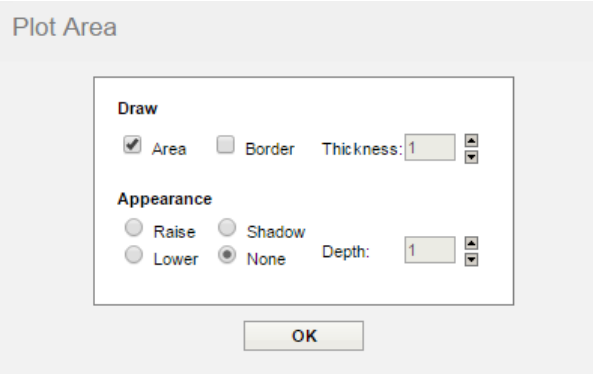
You can specify a label for the line as well as on which element of the data series the calculation should be based. The following types of trend lines are supported: polynomial of any degree, exponential, logarithmic, power, moving average, exponential moving average, triangular moving average, and cubic B-spline. For moving averages you will need to specify the average period and for a polynomial you will need to specify the curve order. You can also specify the thickness of the line and configure whether a label in legend and the attached text should be shown. In case the chart has data series, you can configure the trendline for a specific series.

Please note that there is a *Stack line* option in the *Trend Line* dialog. This option is only available for Stack Column, Stack Bar, and Stack Area chart. You can select whether the line will be drawn for positive or negative values.

Once you finish setting up the options, click on the *OK* button to apply the changes.

5.4.5.4.4. Plot Area

The plot area is the plane on which the data points are drawn for two-dimensional charts. You can customize the appearance of the plot area by clicking on the  *Format Plot Area* button on the toolbar. The dialog will open in the left pane.



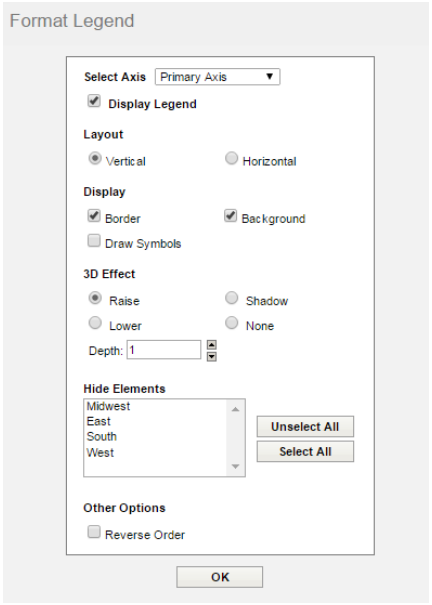
Plot Area Dialog

This dialog allows you to draw a border around the plot area or fill it to provide a background color. If you fill the area, you can also specify several 3D effects like raising, lowering, or drawing a shadow. You can emphasize the 3D effect by increasing the *Depth* option.

Once you finish setting up the options, click on the *OK* button to apply the changes.

5.4.5.4.5. Chart Legend

You can control and modify the display of the chart legend by clicking on the  *Format Legend* button on the toolbar. The dialog will open in the left pane.



Format Legend Dialog

The drop-down list at the top of the dialog allows you to select which legend (primary or secondary) you want to modify. The formatting options are as follows:

Display Legend: This option allows you to turn on or off the legend.

Layout: This option allows you to change the legend from vertical to horizontal layout.

Display: This option allows you to display border, background, and symbols for the legend.

3D Effect: This option allows you to add a 3D effect to the legend. You can raise it, lower it, or draw a shadow. The 3D effect can be emphasized by increasing the *Depth* option.

Hide Elements: This option allows you to select certain category/series elements to hide in the legend.

Other Options: This option allows you to choose whether or not to draw the legend in reverse order.

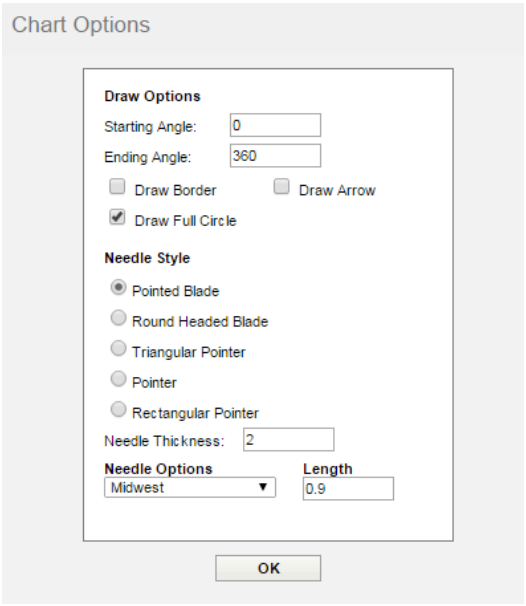
Once you have finish seting up the options, click on the *OK* button to apply the changes.

5.4.5.5. Chart-Specific Options

There are a number of formatting options that are unique to certain chart types - Pie Chart, Dial Chart, HLCO Chart, Gantt Chart, Polar Chart, Doughnut Chart. These options can be modified by clicking on the  *Chart Type Options* button on the toolbar in these types of charts. This will bring up a dialog in the left pane that varies depending on the type of the current chart. Besides the chart types listed here, other chart types have no additional options.

5.4.5.5.1. Dial Charts

For dial charts, the following dialog is opened:



Dial Chart Options Dialog

The following options are available for dial charts:

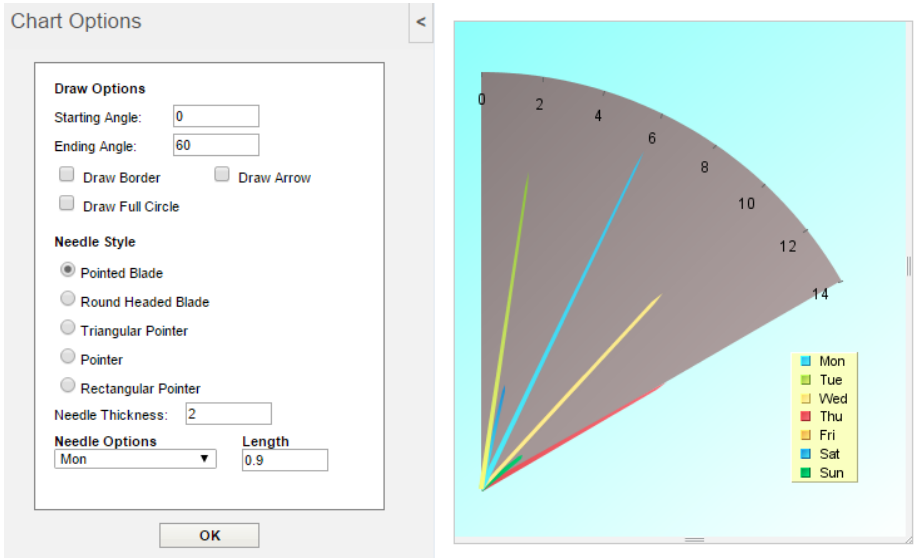
Starting Angle: This option specifies the angle of the first axis label. The angle is represented in degrees and is set to 0 by default. Assuming the dial chart is a clock face, 0 degrees is 12 o'clock.

Ending Angle: This option specifies the angle of the last axis label. The angle is represented in degrees and is set to 360 by default. By default, the labels (and data points) encompass the entire circumference of the dial.

Draw Border: This option specifies whether or not to draw a border around the dial.

Draw Arrow: This option specifies whether or not to draw arrowheads at the end of the dial hands.

Draw Full Circle: This option specifies whether to draw the dial as a complete circle (360 degrees) or only draw the portion of the circle determined by the starting and ending angles.



Circular Sector

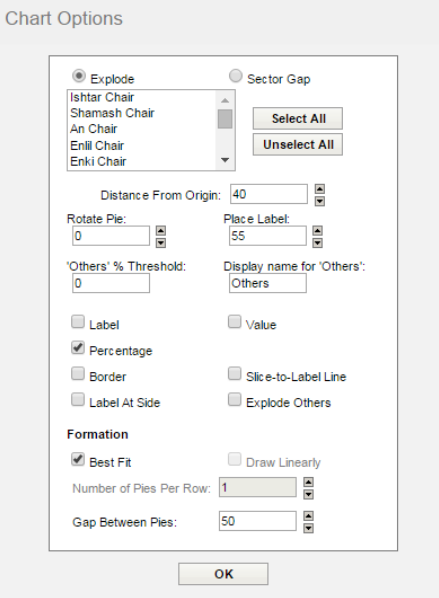
Needle Style: This option allows you to set style of the needle. The default one is *Pointed Blade*.

Needle Options: This option specifies the distance of the hand for each needle (if multiple needles exist) from the center of the dial. The range is from 0 (center of the dial) to 1 (end of the dial). You can adjust the needle length for each category by selecting the category from the drop-down list.

5.4.5.5.2. Pie and Doughnut Charts

Please note that doughnut chart specific options are exactly same as for pie charts, so even though the next section will only mention pie charts, the same options also apply for doughnout charts.

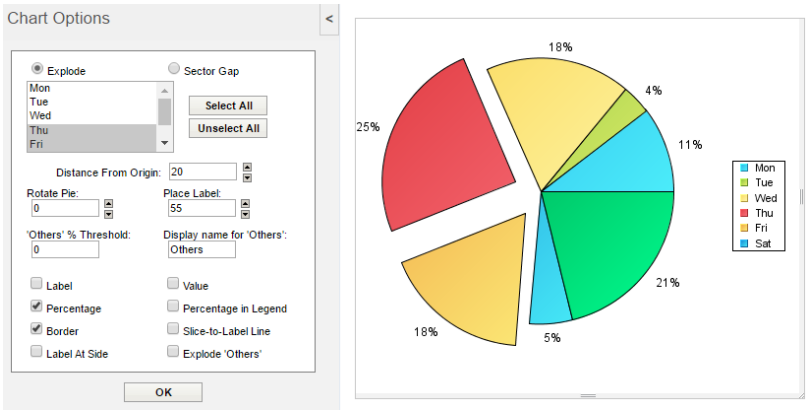
The following dialog is opened for pie charts:



Pie Chart with Series Options Dialog

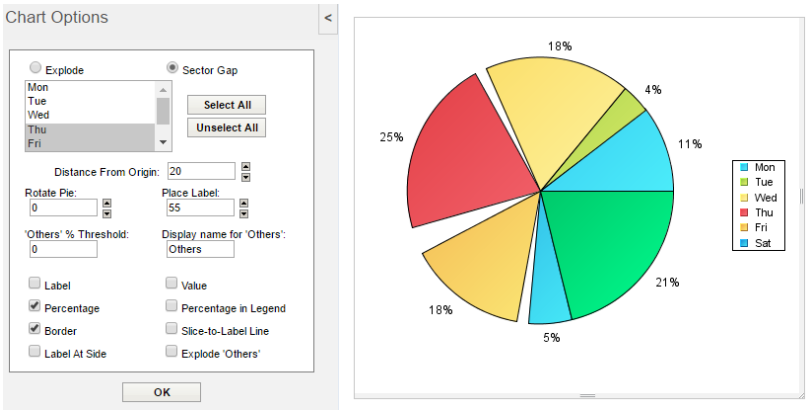
Different options will appear in this dialog depending on whether the pie chart is two-dimensional or three-dimensional and whether or not it has a series. The following options are available for pie charts:

Explode: This option allows you to pick one or more category/series elements whose sections are to be drawn at a certain distance away from the center of the pie.



Exploded Slices

Sector Gap: This option allows you to pick one or more category/series elements whose sections maintain the same distance from the center of the pie. This option creates a gap between the slices.



Sector Gap Applied on Slices

Distance From Origin: This option allows you to specify how far the exploded sections are to be drawn away from the center. This number, represented as a percentage of the radius, indicates the distance between the center and the tip of the pie slice to be exploded. For sector gap, this option is used to define the gap between the slices representing selected categories/series.

Rotate Pie: This option allows you to specify the number of degrees that the chart should be rotated. To rotate the pie in a clockwise direction, use negative values (i.e. -90). To rotate it in a counter-clockwise direction, use positive values (between 1 - 360).

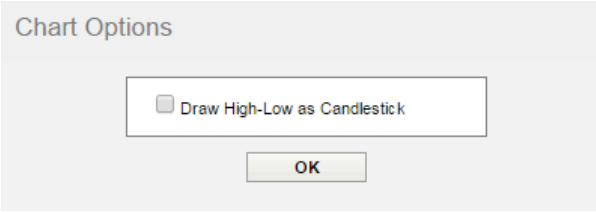
Place Label: This option indicates the distance of the labels from the center of the pie.

'Others' % Threshold: This feature is useful for pie charts that have a large number of small categories. Rather than drawing a slice for each category, you can select a threshold value. Any category with less than the threshold value will be moved into “Others” slice.

Display Name for 'Others':	This option allows you to set the display name for the “Others” slice that is created for categories that fall below the set threshold value. This label will appear in the legend and/or for the slice label.
Label:	This option determines whether a category/series label should be drawn for each pie slice. By default, it only appears as legend.
Value:	This option allows you to specify whether to display the actual value of each pie slice.
Percentage:	This option allows you to display the percentage for each pie slice. The percentages are calculated by dividing the value of each section by the sum of all the values.
Percentage in Legend:	This option allows you to display the percentage in the legend represented by each slice in the pie. In case a pie slice is too thin, this can be a preferable presentation. This option is only available if the pie chart does not have a data series.
Border:	This option specifies whether to draw a border around each pie slice.
Slice-to-Label Line:	This option will draw a line from any label(s) to its corresponding pie slice.
Label At Side:	This option will place the labels for a pie chart away from the plot around the outside of the chart. When used with the "Slice-to-Label Line" option, it gives you a way to display the pie labels for charts with many small categories without any text overlapping.
Explode Others:	This option allwos you to draw the “Others” segment at a certain distance from the center of the pie chart.
Best Fit:	This option will arrange multiple pies in best configuration to fit the chart canvas. It's only available for pie charts with data series.
Draw Linearly:	This option will arrange multiple pies in a straight horizontal line. It's only available for pie charts with data series.
Number of Pies Per Row:	This option allows you to create a custom arrangement of multiple pies by specifying the number of pies to be drawn in each row. It's only available for pie charts with data series.
Gap Between Pies:	This option allows you to specify a gap between multiple pies. The number represent multiple of the pie radius, so the gap will adjust with the size of the chart plot. It's only available for pie charts with data series.

5.4.5.5.3. HLCO Charts

The following dialog is opened for HLCO charts:

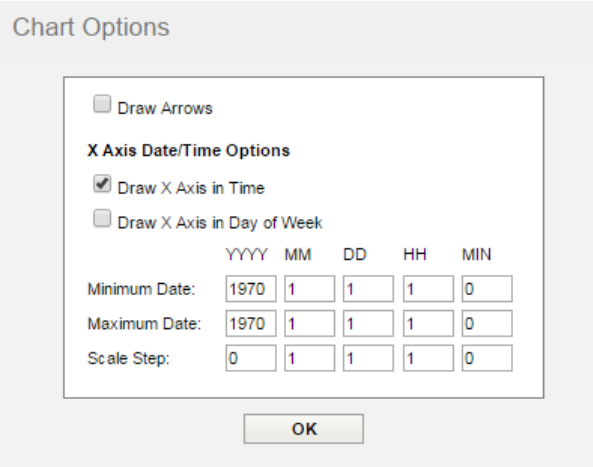


HLCO Chart Options Dialog

The *Draw Hi-Low As Candle Stick* option will turn the HLCO chart into a candle representation. A candle HLCO chart blends high, low, close, and open data into a single object that resembles a candlestick.

5.4.5.5.4. Gantt Charts

The following dialog is opened for Gantt charts:



Gantt Options Dialog

The following options are available for Gantt charts:

Draw Arrows: This option will draw connecting arrows between category elements of the Gantt chart. This allows you to illustrate a sequence between scheduled events. The arrows are drawn in the order the category elements appear in the data source.

Display X-Axis in Time: This option shows the ticker labels as time values instead of numeric values for the X-axis.

Display X-Axis in Day of Week: This option shows the ticker labels as days of the week with the date for each Sunday shown as well.

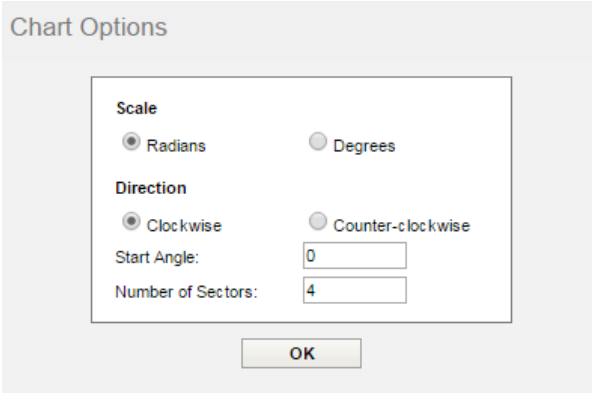
Minimum Date: This option specifies the beginning date for the X-axis. The format is year, month, day, hour, and minute.

Maximum Date: This option specifies the end date for the X-axis. The format is year, month, day, hour, and minute.

Scale Step: This option specifies the scale step for the X-axis. The format is year, month, day, hour, and minute.

5.4.5.5.5. Polar Charts

For polar charts, the following dialog is opened:



Polar Chart Options Dialog

The following options are available for polar charts:

Scale: This option allows you to specify whether the input data for the angle (θ) portion of the data points is in radians or degrees. The chart will always display angles from 0 to 360. If the input data is in radians, it will be displayed as degrees.

Direction: This option allows you to specify whether the circular plot should be drawn clockwise or counter-clockwise.

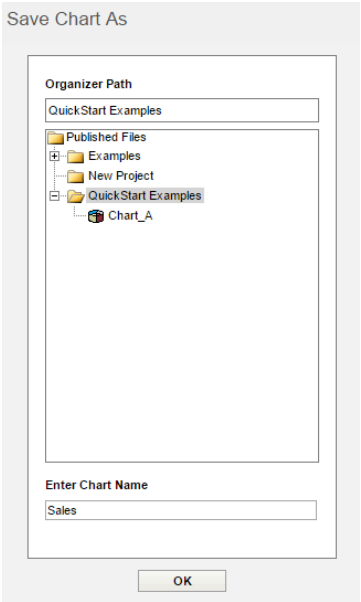
Start Angle: By default, the top of the polar chart plot is 0 degrees. This option allows you to specify a different angle for the top of the plot. It can be displayed either in degrees or radians, depending on the scale you chose.

Number of Sectors: This option allows you to select the number of sectors you want to show in the chart. Sectors are created by drawing additional polar axis lines at specified angle intervals. By default, four sectors are shown.

5.4.6. Save the Chart

Once you create a chart in QuickDesigner Charts, you can share it with other users by saving it on the server and inserting it into the Organizer.

To save a chart, click on the  *Save* (or  *Save Chart As* for saving the file under a different name that is already saved) button on the toolbar. A dialog will open in the left pane.



Save Chart Dialog

Here you can see all Organizer projects and folders you have access to. Select the project or folder you want to save the chart into. Specify a name for the chart at the bottom of the dialog. The chart will be saved with `.qch` extension.

Once you finish specifying the options, click on the *OK* button to save the chart. The new file will be saved in the ChartFiles folder in the ERES installation directory as `ChartNameInOrganizer_Username_TimeStamp.qch`. This should help to avoid unintentional overwriting of chart files.

If you try to save a file with the same name as another file, you will be asked to confirm overwriting it or using a different name. If you wish to save the changes in the currently opened chart, do not change the chart name and the Organizer folder in the Save Chart dialog. Click *Yes* in the dialog and the chart file will be updated.

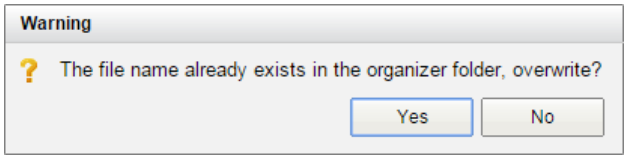
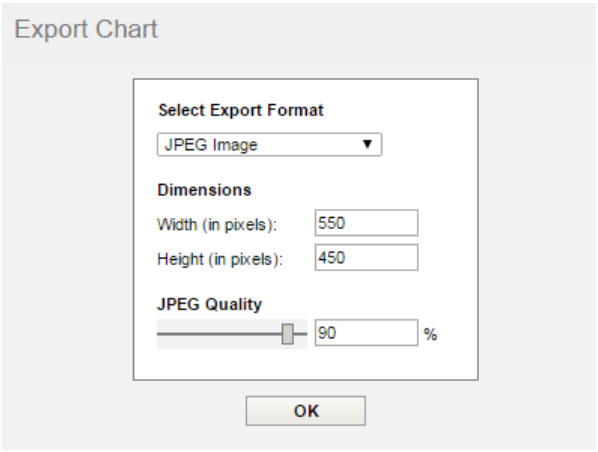


Chart Overwrite Dialog

5.4.7. Export the Chart

You can export a chart in a number of different formats. To do so, click on the  *Export* button on the toolbar and a dialog will open in a new window prompting you to specify the export options.



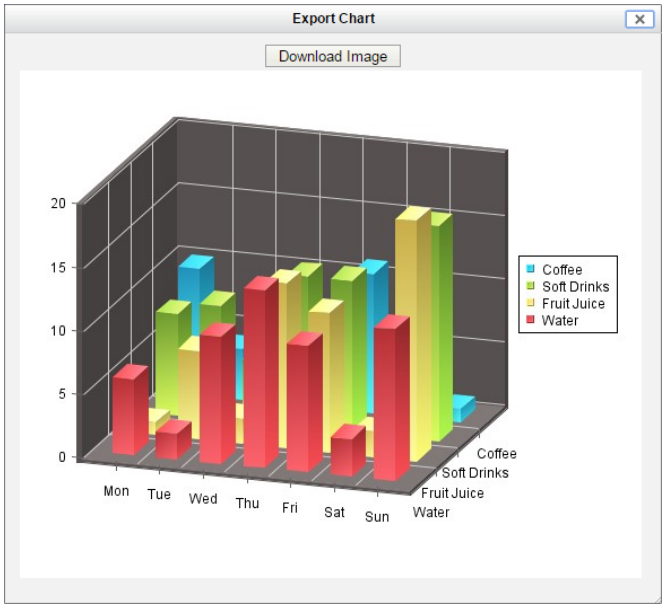
Export Chart Dialog

The first option allows you to specify the format in which to export the chart. The available formats are GIF, JPEG, PNG, BMP, PDF, SVG, SWF (Flash), Text, XML, and Chart Image Map.

The second option allows you to set the dimensions of the exported image. By default, these will match the canvas dimensions of your chart.

The last options allow you to set some image type-specific options. (The options appear only for appropriate format.) You can set the background transparent for GIF files (please note that the result is not good for more than 256 colors - e.g. with Color Gradient), set the quality for JPEG files, as well as specify compression for PNG files.

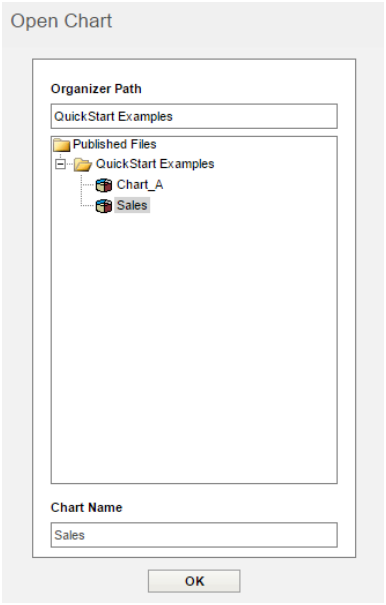
Once you finish specifying the options for the exported chart, click on the *OK* button and a new window will open containing the exported chart. You can save the generated file to your local system by clicking on the *Download Image* button.



Exported Chart

5.4.8. Open the Saved Chart

You can open a saved chart by clicking on the *Open* icon on the main toolbar. The *Open Chart* dialog will appear.



Open Chart Dialog

Here you can see all charts created in QuickDesigner Charts. Select a chart and then click on the *OK* button to open it.

5.4.9. Exit

You can exit QuickDesigner Charts either by clicking on the *Home* icon in the upper right corner, clicking on the *QuickDesigner Charts* title, or by clicking on the *Logo* icon in the upper left corner. Before closing, you will be asked if you want to save an unsaved chart.

Chapter 6. Designing Maps

6.1. Introduction to ERES Maps

ERES Maps are designed to report geographical data from data sources. They fetch geographical data from data source and mark them on a map. There are two types of maps: Online Maps and SVG Maps.

Online Maps The **Online Maps** feature in ERES displays geographical data using high quality on-line maps made by Google and Open Street Maps. Online Maps provides street maps and satellite maps, spanning over the entire world. By using high quality satellite imagery, a very wide zoom range is possible, making it possible to display every imaginable area of the world. Online Maps are useful for displaying exact spots on the maps. These spots are called **map points** and they are marked with **map markers**. A powerful tool provided by ERES Online Maps is **Geocoding**. Geocoding can locate buildings/cities on map by knowing only their addresses. Usage of Online Maps is limited by Online Maps Terms of Service (To see the current version, visit Google Maps API Terms of Service [https://developers.google.com/maps/terms.html], Open Street Maps API Terms of Service [https://www.openstreetmap.org/copyright/en]).

Online Maps can contain tooltips and drilldowns. **Tooltips** are small reports or charts which are displayed on mouse over a map marker. They display brief summary information related to the selected map point. **Drilldowns** are links to parameterized charts, reports or other maps that open when you click on a map marker. They are usually used to display more detailed information about the selected map point.

SVG Maps **SVG Maps** display data on SVG (Scalable Vector Graphics) maps. They differ from ERES Online Maps in two ways. First, SVG Maps work with geographical **areas** (e.g. a state is marked by the whole territory, not just by a point in the middle of the state). These areas can be colored according to set criteria. Second, the SVG Maps feature requires **SVG images** as the map source and do not use any online map source. The SVG Maps feature may be used to display spatial (non-geographic) and customized maps (e.g. maps of buildings, floor plans, mine maps, etc.).

6.2. Online Maps

ERES Online Maps use two types of files: Coordinates and Map Files. **Coordinates** contain location of individual map points (e.g. cities, company branches). Coordinates have to be created before creating a map. One set of Coordinates can be shared by several maps, so you do not have to store the same file several times. It is enough to add a map point in the Coordinates and it is added to all maps that use these Coordinates. For more information about Coordinates, see Section 6.2.3.2 - Create Coordinates. **Map File** displays map markers based on Coordinates and defines Drilldowns, Tooltips, and Heatmaps.

6.2.1. Data Sources

Each ERES Online Map uses data from two data sources - the first one is used for Coordinates and the second one is for the map itself. This chapter explains where and how these data sources are used.

The Coordinates data source is used when you are creating Coordinates. It is used for obtaining all the Point IDs. The **Point ID** is an identifier of a place (e.g. city name, branch name etc.) and it is loaded from the data source where you selected to create the Coordinates. The Point ID can consist of several **Point ID fields** - each data source column creates one field (except the fields which contain longitude, latitude, or WKT). For example, if your data source contains two columns - City and State, the Point ID will consist of two fields - City and State.

Each map point should have its unique Point ID, so any map point can be distinguished from the others. The data source should be chosen wisely to fulfill this condition (each row of data should be unique, there must not be duplicates).

It is recommended not to use parameterized data source for Coordinates, because there is no way to set parameter values for Coordinates. If you use a parameterized data source for Coordinates, the default parameter values are always used. This does not apply to multivalue parameters which use all the values (e.g. if you use `CustomersInStates` HSQL query (`Databases/Woodview/Query/CustomersInStates`) to create coordinates using geocoding - all param values (states) will be used).

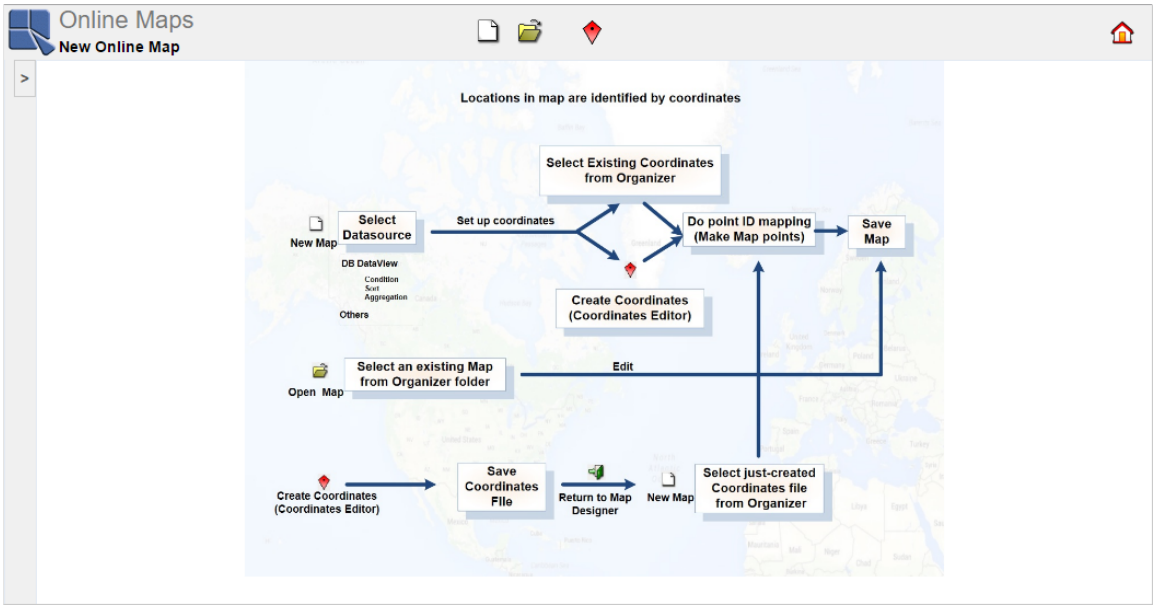
The Point ID is used for mapping map markers to the map data source in the Map File, so it should contain data which occurs in the map data source and which uniquely identifies map points. For example, if you want to display cities where your company has some customers, you should use City name and State name as the Point ID fields. The State name is useful in case when there are more cities with the same name in several different states. Optionally, you can add more Point ID fields for more detailed specification of city. For example if there are two cities with the same name in the same state, you should use more Point ID field(s) which will distinguish between these two cities, a ZIP code for instance.

The Coordinates data source can also contain coordinates (longitudes, latitudes, or WKT) of the map points. The longitudes and latitudes have to be stored in two separated columns of double data type, while WKT is a text field. These columns are not used as Point ID fields. If the coordinates are read from the data source, they are fetched every time the map is generated, so if there is a change in the Coordinates data source, it will be reflected in the map immediately. If the coordinates are not read from the data source, they are stored directly in the Coordinates file. In this case, the Coordinates data source is accessed only when you are creating or editing the Coordinates in the Coordinates Editor. When you are creating or displaying a map that is using this Coordinates, the data source is not accessed at all, because all the necessary information is in the Coordinates file, so it is not necessary for this Coordinates data source to be available when the map is created/displayed.

The map data source has two purposes. The first one is that it contains data that are used for Tooltips and Heatmaps. The second purpose is filtering. The map data source is mapped to the Coordinates as described in the Section 6.2.4 - Coordinates Mapping and only those map points from the Coordinates that have at least one associated row in the map data source are displayed. The map data source can be parameterized. In this case, the whole map is parameterized. If you create/open map that is using parameterized data source, there will be a filter icon in the map which opens parameters dialog that allows you to enter parameter values.

6.2.2. Start Online Maps

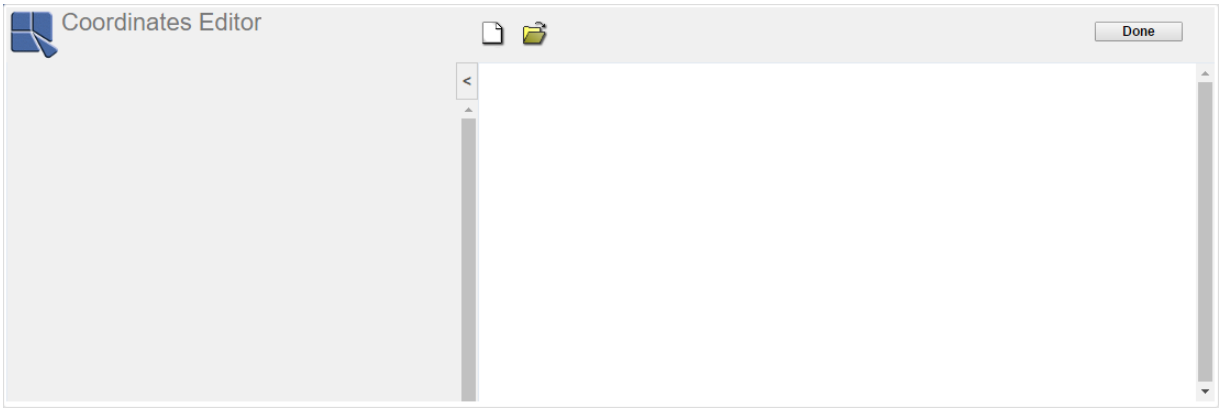
Online Maps can be started by clicking on the *Online Maps* link on the ERES Main page. After launching Online Maps, you will see a diagram that shows how to work with Online Maps. You can choose to create a new map, open an existing one, or open the Coordinates Editor where you can create/edit Coordinates.



Online Maps Start Options

6.2.3. Coordinates Editor

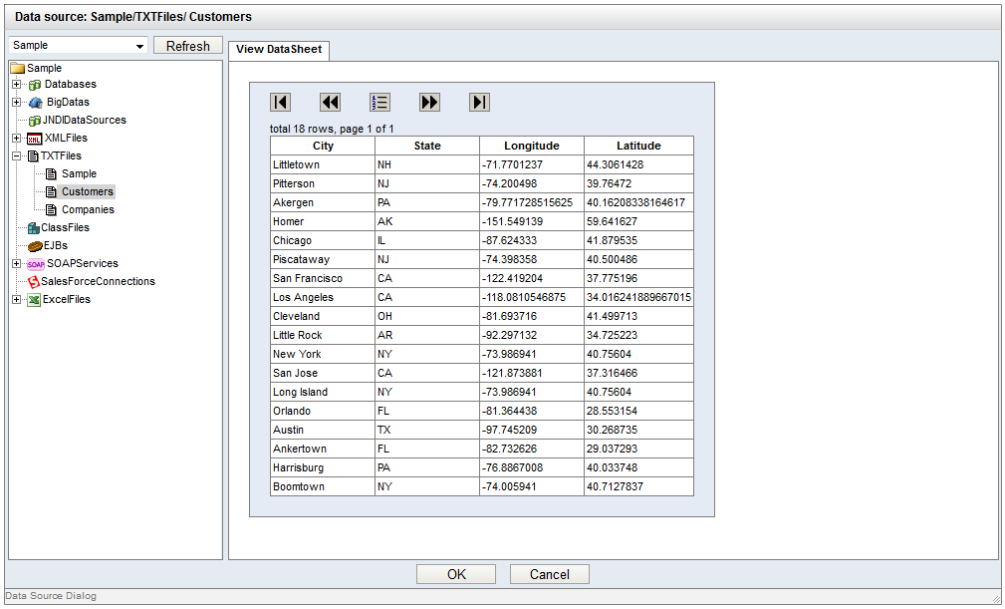
Click on the  *Coordinates Editor/Edit Coordinates* icon on the toolbar. The *Coordinates Editor* will open.



Coordinates Editor Start Options

6.2.3.1. Select Data Source

Open *Data Source Dialog* by clicking the  *New Coordinates* icon on the toolbar. Select the desired data source. A view of the data will appear in the right pane. It allows you to check whether you chose the right data source. Click *OK* to close the *Data Source Dialog*.

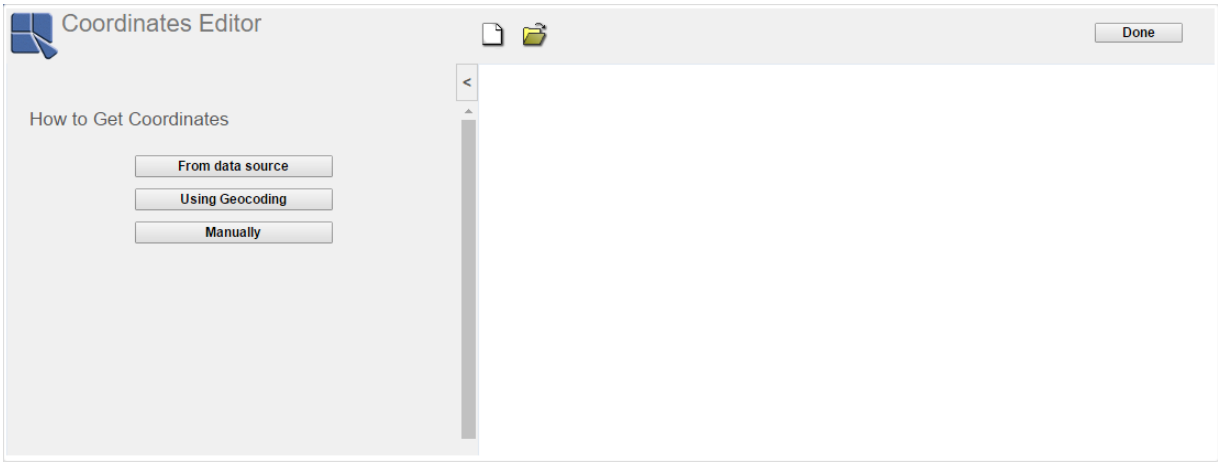


Data Source Dialog

6.2.3.2. Create Coordinates

The **Coordinates** file contains the coordinates of map points. Each Coordinates data record consists of WKT (obtained automatically or manually) and Point ID (for more information about Point ID, see Section 6.2.1 - Data Sources). Note: Existing coordinates files or data sources with longitude and latitude columns are still supported, but only WKT will be generated by geocoding or manual entry.

After selecting the data source, you will be prompted to select a method for obtaining Coordinates from the data source.



Get Coordinates

There are three ways how to get Coordinates, which are described below. Status of each pointID is determined by colors in the pointID table (For more information about colors, see Section 6.2.3.3 - Coordinates Editor Interface):

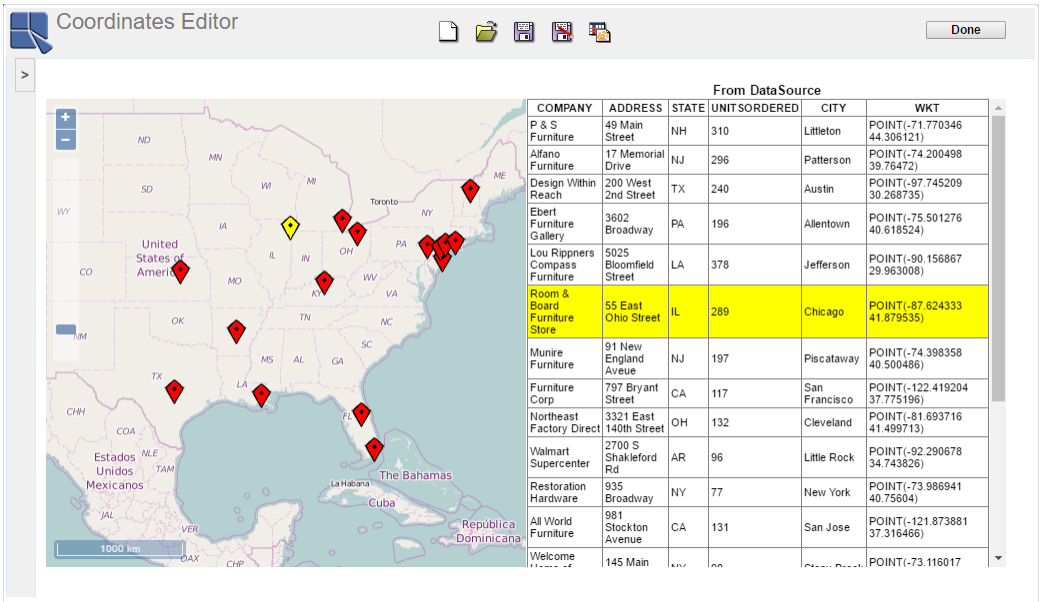
6.2.3.2.1. From Data Source

This option requires two numeral data source columns containing information about the latitude and longitude or WKT of particular map point. Latitude and longitude have to be in different columns. The data type of these columns should be double and they should contain latitude and longitude in decimal degree notation. Positive longitude means East longitude, negative means West. Positive latitude means North latitude, negative means South.

In this case, map point latitudes and longitudes are read from the data source every time a map, which is using this Coordinates, is open. If the data source changes, the change is reflected in the map (e.g. if a new map point is added, it will be displayed on the map). It is not necessary to edit the Coordinates, updating the data source is good enough.

If Geocoding or Manual insertion of the map point coordinates is used, it is necessary to edit the Coordinates every time you want to add or modify (move) a map point.

If the Coordinates are obtained from a data source, all map points will be placed automatically according to the data from the data source. These rows will have white color. In this case, the Coordinates are read-only - you can only check which map marker belongs to which Point ID by holding mouse over a map marker/Point ID. The associated Point ID/map marker will be highlighted with yellow color. If you want to change the position of a map point, you have to modify the data source directly (i.e. change the longitude and latitude or WKT in the appropriate data source column).



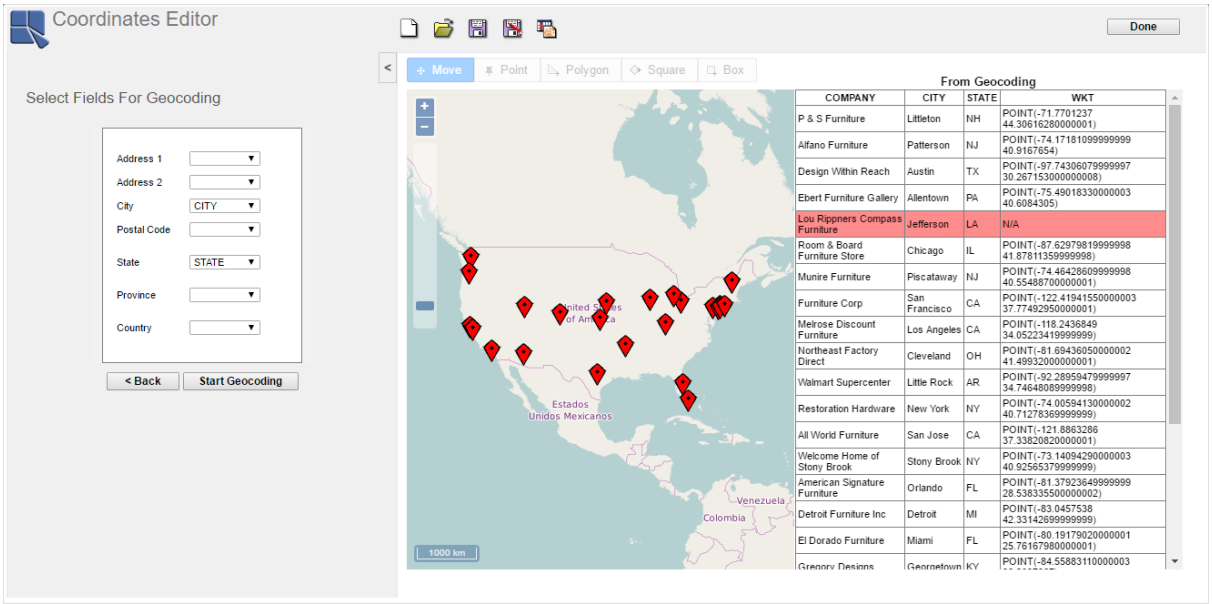
Coordinates from Data Source

6.2.3.2.2. Geocoding

Positions will be determined automatically using point addresses. The addresses can consist of several data source columns. You can select columns, which contain Address 1, Address 2, City, Postal Code, State, Province, and Country. Some of the fields may already be pre-selected when there is a convenient column in your data source (e.g. if the data source contains a column called *City*, it will be pre-selected in *City* field). You do not have to fill all the fields. If you do not fill any field, you will be automatically switched to manually option (see the next chapter). The name of the Geocoding fields are only illustrative, they do not have to be filled with exactly this information. The only rule you should keep in mind is to use the more specific information (e.g. street address) before the less specific (e.g. state). If your data source column contains more then just one field, use the first suitable field (e.g. if there is column with full address, you should use Address 1 field and leave the rest empty). Geocoding is very adaptive as it does not require any specific format of addresses and is able to find even inaccurate or incomplete addresses.

Please note that Geocoding may not be supported for all countries. To see the list of currently supported countries, visit Google Maps Coverage [https://developers.google.com/maps/coverage] and find your country in the list.

If you selected Geocoding, the Geocoding process is started automatically. It goes through all Point IDs and tries to determine positions of the corresponding points from their addresses. If it finds exactly one result, the map point is inserted automatically. The new point and the corresponding Point ID will turn yellow for a moment and then it will turn white, which indicates that the position is already determined. The toolbar is disabled during Geocoding, so you have to wait until it is finished before performing other actions. You can watch the Geocoding progress on the status bar. After the Geocoding is done, you should insert points that were not inserted automatically. These points had either no Geocoding results or more than one Geocoding result. The process is same as inserting the point manually and it is described below.



Coordinates by Geocoding

If you used Geocoding and you insert the rest of the points manually, the Point ID row may turn green instead of dark red after you click on it. This occurs when there is more than one Geocoding result. These results are displayed on the Online map with green markers so you can select the right one by clicking on it. If the desired map point is not among the Geocoding results, you can still add the map point manually on the correct place in the same way as manual inserting. If you hold mouse over a Geocoding result (green map marker), the full address of this map point is displayed to help you identify the correct place.



Selecting Geocoding Result

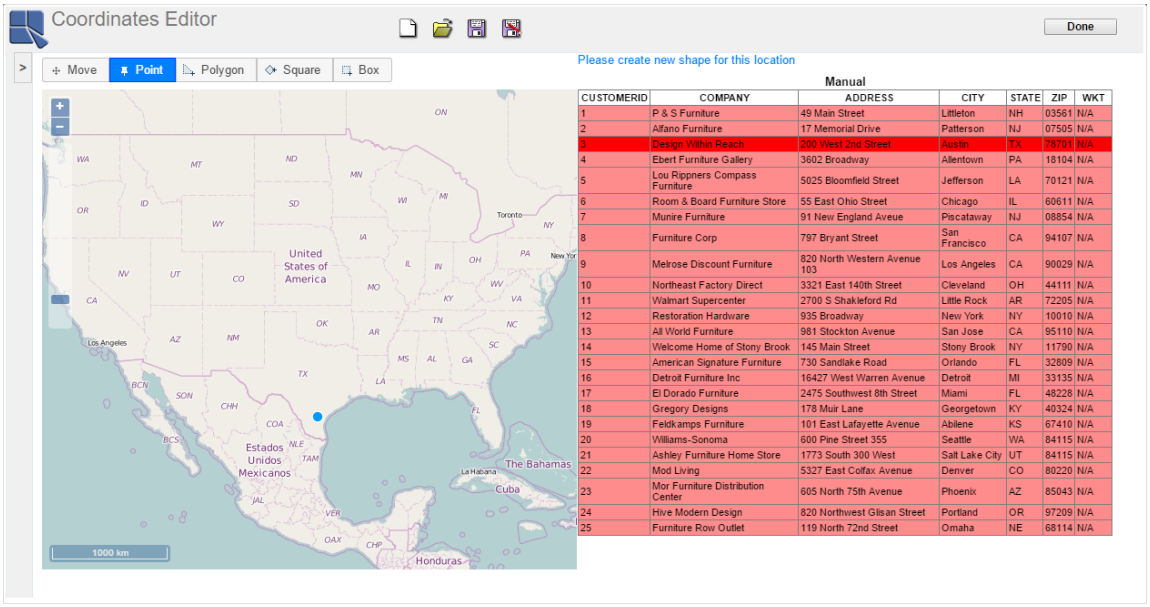
Please note that the Geocoding is not 100% accurate, so you should check all results that were inserted automatically.

6.2.3.2.3. Manually

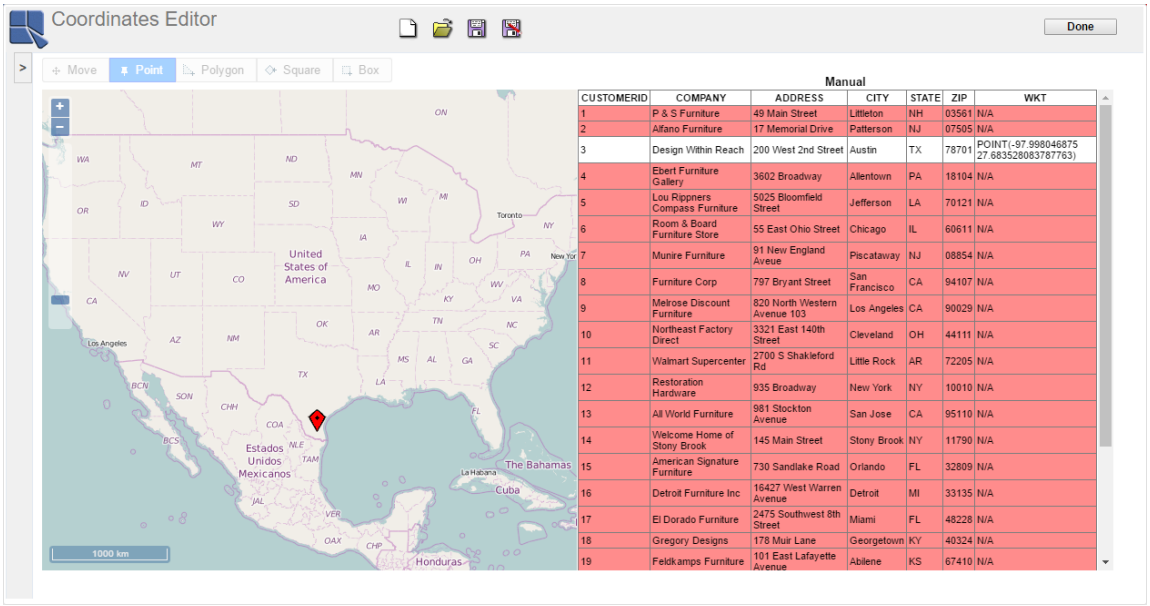
This option does not require any geographical data at all. Coordinates of map points are determined manually by clicking on the map.

In the *Manually* case, all the Point IDs are light red at the beginning, which means that their position is not determined. To place a map point, click on a Point ID (row in the table) first. The Point ID row will now turn dark red. *Point* option should be active as default on the map toolbar (*Point* is blue). Mouse over the map and you will see a small blue circle under your mouse pointer. Then click on the location where you want to place your map point.

A new map marker will be placed on the selected map point and the appropriate row will become white. Then you can place another map point in the same way. To freely navigate across the map, select *Move* option on the toolbar. Now you can move and zoom the map any way you need. You can also use another shapes to define a location on the map - *Polygon*, *Square*, and *Box* on the map toolbar. To use these options, select the shape, place the first point on the map by clicking the mouse, move the mouse and place another point(s).

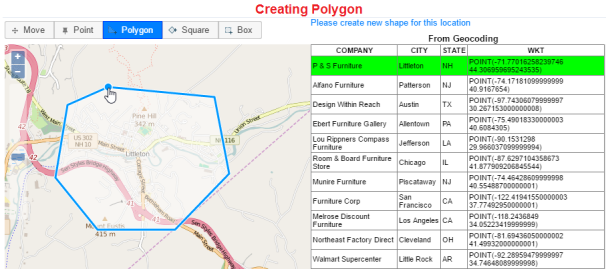


Selecting Location

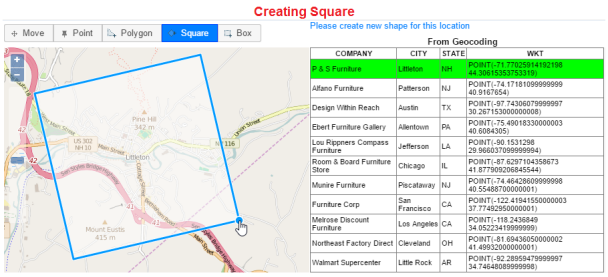


Selected Location

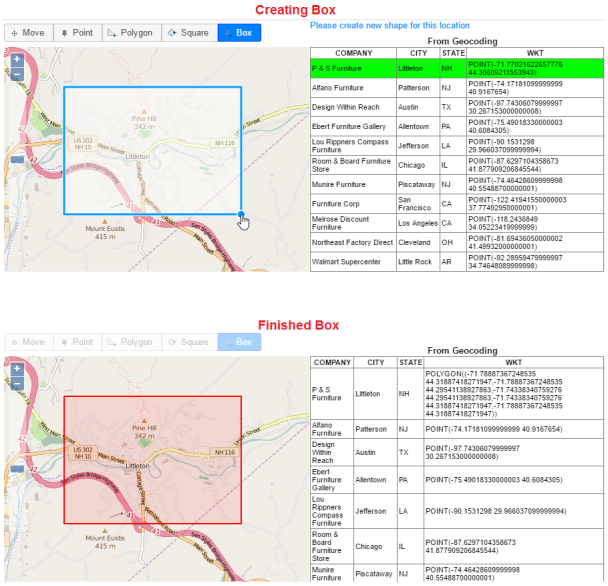
The following images show the formation of shapes (blue color) and their final form (red color).



Polygon



Square

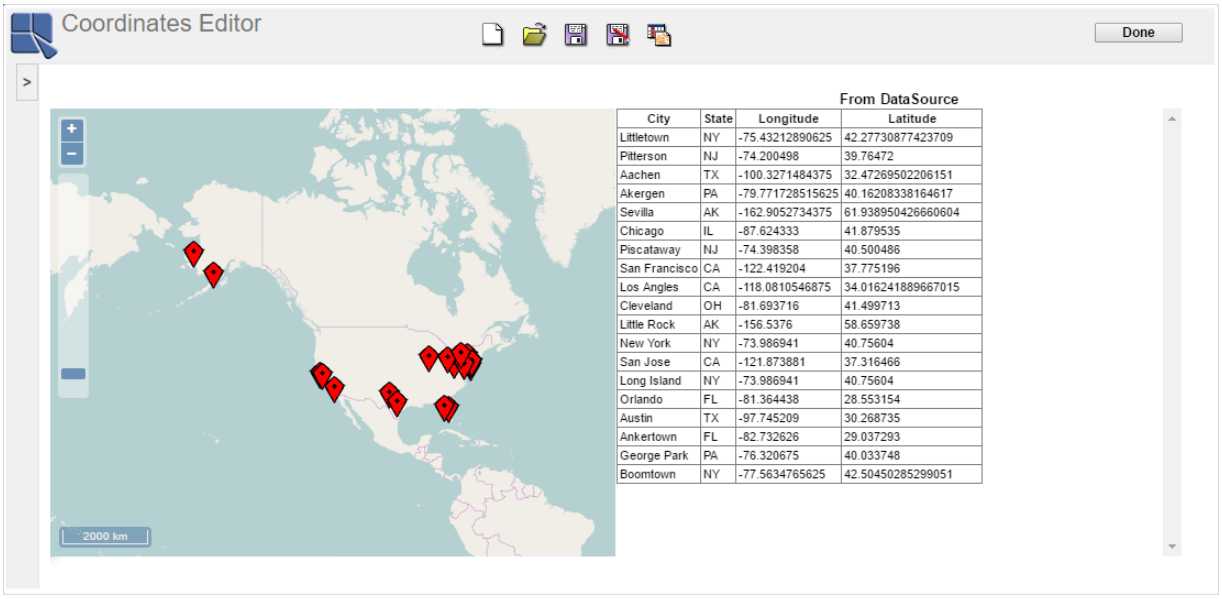


Box

You can move any of the already existing map points by dragging the appropriate map marker with mouse. You can also resize a shape by dragging the edge of a side to any direction.

6.2.3.3. Coordinates Editor Interface

After choosing a method of obtaining coordinates, you will get a map with a point ID table.



Coordinates Editor

The following buttons are on the toolbar:



New: Creates new Coordinates.



Open: Opens existing Coordinates.



Save: Saves current Coordinates to a file and Organizer. Note that all saved files will include userID and timestamp in addition to the name you provided. If an Organizer item with the supplied name already exists in the folder, you will be given a choice to override the existing item or create a new entry.



Save As: Allows you to save the existing or modified Coordinates to a different file or location.



Change Point Mapping: Allows you to change point mapping.

The left part of the window displays the Online Map, which displays individual map points using the map markers. The Online Map can be scrolled and zoomed to find a specific object.

Point ID table in the right part of the window shows all the Point IDs, which were read from the data source, and longitude and latitude of map points. Each row represents one Point ID and can have one of the following background colors:

White: Location of the corresponding map point has been successfully determined.

Light red: Location of the corresponding map point has not been determined yet. There are multiple locations for the point. This is not static status color.

Grey: Location of the corresponding map point has not been determined yet. There is no location for the point. When you click this grey-colored row, it turns into dark red. This is not static color.

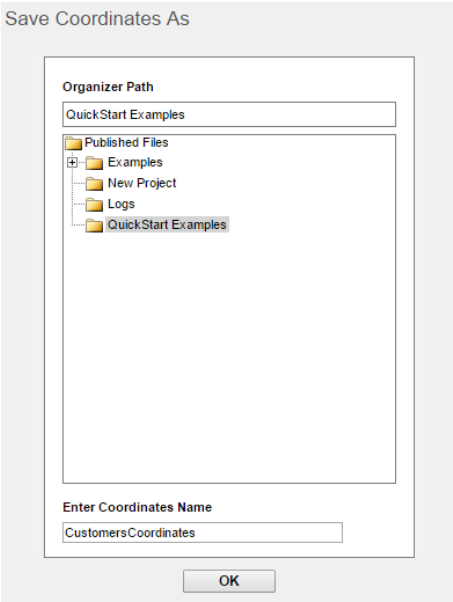
Dark red: Marks the point ID, which is currently under mouse cursor. If you click on it, it will either change into green (multiple locations), or grey (No location found), waiting for users to enter the location manually.

Yellow: Highlights the row with Point ID of the map point under mouse cursor on the map (this is used to determine which Point ID belongs to which map marker).

Green: Indicates that there are Geocoding matches for the selected Point ID.

6.2.3.4. Save Coordinates

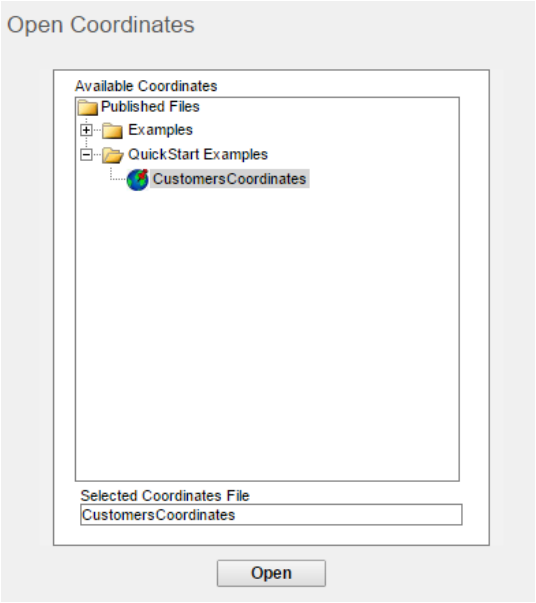
If you want to use the Coordinates for creating maps, be sure to save it to Organizer (click the  *Save* icon, enter a Coordinates name, and select any Organizer folder). Then click the *Done* button on the right side of the toolbar to close the Coordinates Editor and go back to Online Maps main window.



Save Dialog

6.2.3.5. Open, Edit Coordinates

To open a saved coordinates file, first open Coordinates Editor by clicking the  *Coordinates Editor/Edit Coordinates* icon on the main toolbar of the Online Maps designer. Then click the  *Open* icon on the main toolbar of Coordinates Editor.

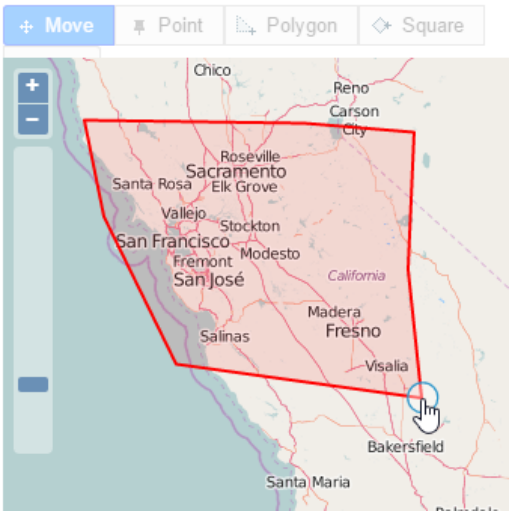


Open Coordinates

You can see all coordinates files from Organizer. Select a coordinates file and click the *Open* button.

Now you can edit coordinates obtained via geocoding or manual insertion. Coordinates obtained from data source are not possible to edit. You can change a point when you click the respective row in the pointID table - then you can either select again one of the geocoding results (applies only to geocoding coordinates) or replace a current point/shape with a new one in the map.

You can directly edit a map point using the drag and drop feature or edit a shape by stretching its sides.



Shape Editing

Please note that if the coordinates obtained via geocoding are not complete (not all rows are white), the geocoding process starts each time the file is opened.

Once you finish editing the coordinates, save the file and click the *Done* button to close the Coordinates Editor.

6.2.4. Coordinates Mapping

ERES Online Maps needs Coordinates to map the data source records to map points. When creating Online Maps, you have to map the Point ID fields from the Coordinates file to the map data source columns. According to this mapping, data source rows are associated with map points. There can be more data source rows associated with one map point but not vice versa (because Point IDs must be unique). The rows associated with a map point are used to generate Tooltips (see Section 6.2.5.7 - Tooltips) for the appropriate map marker.

Not all the Point ID fields have to be mapped. For example, if your Point ID has three fields - City, State, and Country, you can map only Country and State. But you have to keep in mind that any map point should be uniquely determined by the mapped Point ID fields. So if there are two Point IDs in your Coordinates, which differ only in City, but have equal Country and State, you have to use the City as well to distinguish between those two map points. You will get no error message if you use wrong mapping, but the data can be displayed incorrectly.

Some map points from the Coordinates may not have any associated rows. These map points are not displayed on the map using map markers. Only the map markers, for which some data was found in the data source, are displayed. So you can have Coordinates containing all the major cities in USA, but you can display only the one you are interested in. And you can use the same Coordinates for different map, which will show different set of cities, because the data source will contain data for different cities, so you do not have to maintain too many Coordinates files if you choose suitable Point ID fields.

Example:

Your data source contains the following data and you want to display data grouped by cities on a map:

Row #	City	State	Customer Name	Sales
1	San Fran-cisco	CA	Company A	$152,560
2	New York	NY	Company B	$240,468
3	Los Angeles	CA	Company C	$335,256
4	San Fran-cisco	CA	Company D	$80,381

Row #	City	State	Customer Name	Sales
5	Boston	MA	company E	$23,540
6	Chicago	IL	Company F	$124,532

You will have to create the Coordinates first. In this case, only the City name could be used as the Point ID because it is unique, but we will use two Point ID fields instead - City and State, in case that there will be more cities with the same name in different states in the future. So we will create a data source with two columns - City and State. This data source has to contain all the listed cities (and could contain some more). We can either write a suitable query to obtain these two columns from the original data source (if using a database for the data source) or we can, for example, create a simple text file data source with these two columns.

Next, we can start Coordinates Editor and create the Coordinates. Assuming we have a data source that does not contain data for longitudes and latitudes or WKT, so we cannot use coordinates from data source but we know part of an address, we will use Geocoding. Using Geocoding, we will assign WKT to every city, so the table in the Coordinates Editor will look like this (note that there is one extra city - Orlando, which is not in the map data source):

City	State	WKT
San Francisco	CA	POINT(-122.41941550000003 37.77492950000001)
New York	NY	POINT(-74.00594130000002 40.71278369999999)
Los Angeles	CA	POINT(-118.2436849 34.05223419999999)
Chicago	IL	POINT(-87.62979819999998 41.87811359999998)
Orlando	FL	POINT(-81.37923649999999 28.538335500000002)

Now you can create a map from the data source using these Coordinates. You map City Point ID field to the City data source column and State Point ID field to the State data source column. The Point IDs will be paired with the data source rows - San Francisco with the first and fourth, and Chicago with the sixth. Orlando does not have any matching row in the data source, so it will not be displayed at all. If you add a Tooltip, the associated rows will be used to generate it. For example, the Tooltip for San Francisco will be created using data from the first and fourth row of the data source.

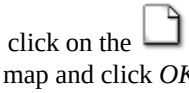
If you add a new city to the data source and you want to display it too, you have to edit the Coordinates data source first and add the city there. Next, you can open the Coordinates in the Coordinates Editor and there will be one new row in the table with this new city. This enables you to add the new map point for this city. You do not have to modify anything in your map. Therefore, if this Coordinates is shared among several maps, the map point is added automatically to all these maps.

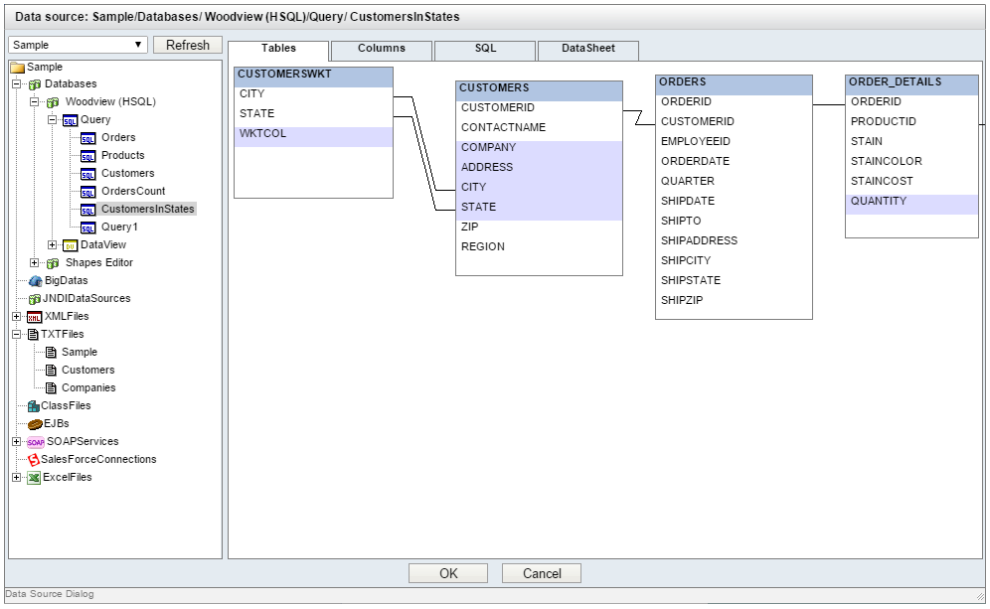
6.2.5. Create Online Maps

ERES Online map is a special kind of report that allows displaying geographical data on map. It points out the reported geographical data on a map using map markers. Markers can be used as drill-down links to parameterized reports, charts, or maps. Markers may also have tooltips. **Tooltips** look like bubbles, containing a report or chart. When tooltips are enabled, every map marker has its own tooltip bubble, which displays data associated to this map point. Only one tooltip bubble can be displayed at the time. A tooltip appears when you move your mouse cursor over a map marker. It closes automatically when you open a different tooltip, or it can be closed by clicking on the *X* sign in the upper right corner of the tooltip bubble.

To create a Online Map, launch Online Maps from the ERES main page.

6.2.5.1. Select Data Source

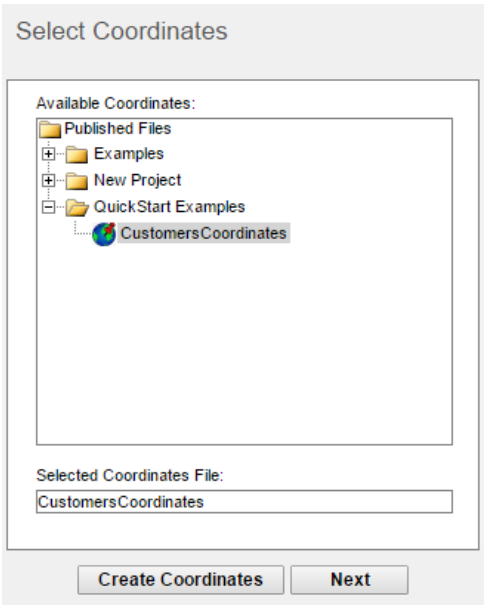
To create a new Online Map, click on the  *New Map* icon on the toolbar. The *Data Source Dialog* will appear. Select a data source for your map and click *OK*.



Data Source Dialog

6.2.5.2. Select Coordinates

Once you select a data source for your map, you will be prompted to select Coordinates file.



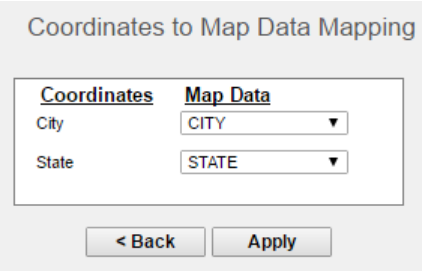
Select Coordinates

Select a Coordinates file and click the *Next* button. You can also launch the Coordinates Editor from this dialog by clicking the *Create Coordinates* button and creating a new coordinates file. Once the new coordinates file is saved, click the *Done* button in the upper right corner of the Coordinates Editor. You will be brought directly to *Coordinates Mapping* which is described in the next chapter.

6.2.5.3. Coordinates Mapping

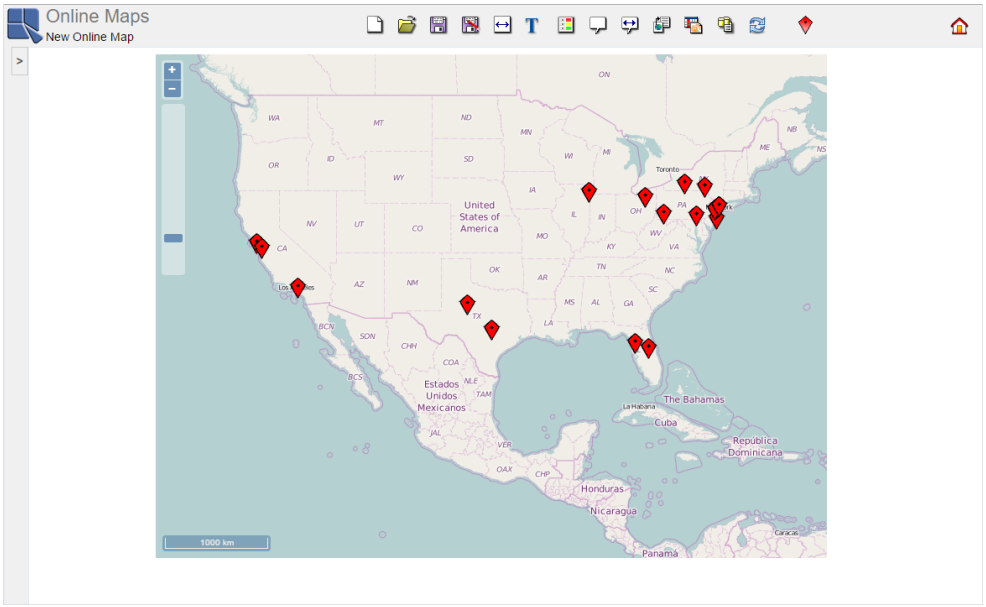
You will be prompted to select point mapping, i.e. you have to pair columns that will enable ERES Online Maps to assign records from maps record's data source to suitable Point ID fields from the Coordinates (e.g. if you want to show cities with your affiliates on map, your Coordinates may contain a column called *City* which is mapped to *Affiliates_City* from the map's data source). For more information, see Section 6.2.4 - Coordinates Mapping. Please note that Online Maps designer automatically recognizes which columns are the most suitable for Coordinates

mapping. Therefore, some of the fields may already be pre-selected when there is a convenient column in your data source (e.g. if the data source contains a column called *City*, it will be pre-selected in *City* field).



Coordinates Mapping

You have to map at least one column. Then click *Apply* and the Online Map will appear.



Online Maps Designer

You can change the Coordinates mapping by clicking the  *Change Point Mapping* icon on the toolbar.

6.2.5.4. Toolbar

These buttons are on the Online Maps toolbar:



New Map: Creates a new Map.



Open Map: Opens an existing Map.



Save Map: Saves the current Map to a file and Organizer. Note that all saved files will include userID and timestamp in addition to the name you provided. If an Organizer item with the supplied name exists in the folder, you will be given a choice to override the existing item or create a new entry.



Save Map As: Allows you to save the existing or modified Map to a different file or location.



Set Map Options: Allows you to change the way the Online Map is displayed.

Map Title Options: Allows you to set map title and its options.

Set Heatmap: Allows you to color-code map markers.

Tooltip Template: Adds/Modifies report/chart templates used as tooltips.

Tooltip Options: Allows you to change dimensions of tooltip bubbles.

Drilldown Options: Configures drill-downs.

Change Point Mapping: Changes the point mapping.

Change Data Source: Changes the data source of current map. Note that changing data source will result in loosing tooltip and drill-down settings.

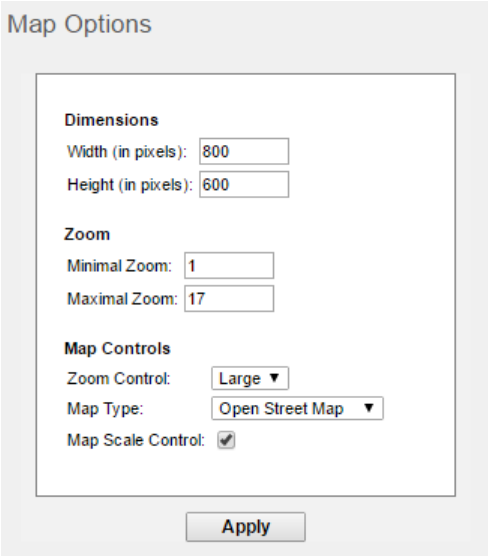
Refresh: Refreshes the map data (if the data source changed).

Coordinates Editor/Edit Coordinates: Opens Coordinates Editor that allows you to create/edit coordinates.

Home: Closes Online Maps designer. Before closing, you will be asked if you want to save an unsaved map.

6.2.5.5. Map Options

To set map options, click on the *Set Map Options* icon on the toolbar.



Map Options

Dimensions: Allows you to set width and height of the map (in pixels).

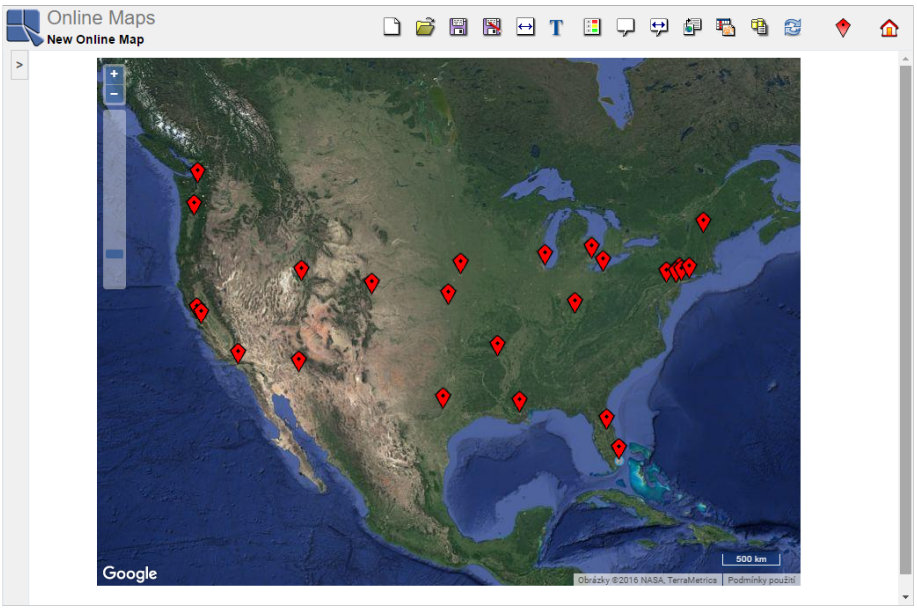
Zoom: Allows you to set minimal and maximal zoom of the map (in the range 1-17).

Map Controls: **Zoom Control** - shows/hides the zoom controls in the upper left corner (large, small, none).



Zoom Control Large

Map Type - allows you to set map type (Open Street Map, Google Street Map, Google Satellite).



Google Satellite

Map Scale Control - shows/hides the scale control in the lower right-hand corner.



Scale Control

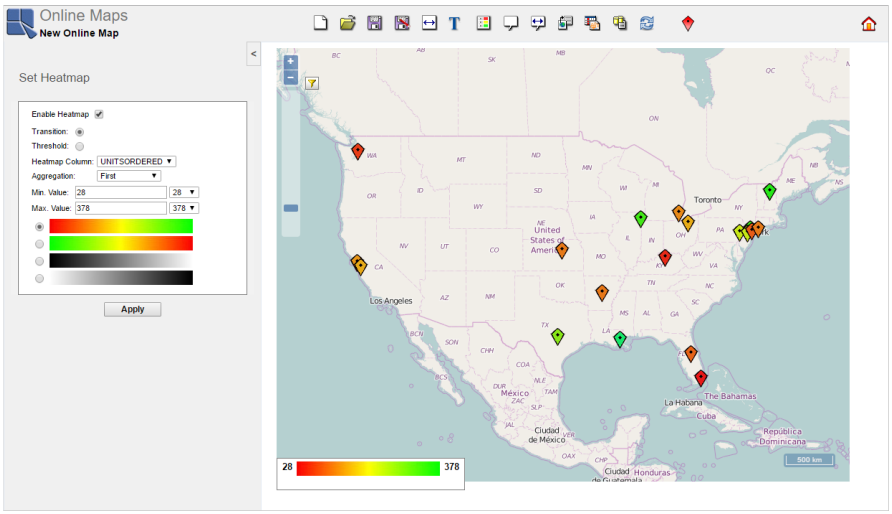
6.2.5.6. Heatmap

Heatmaps allows you to display map points in different colors based on values from a column in the data source. You can choose between two heatmap types - transition or thresholds.

Please note that some of the data sources have multiple records belonging to single map points in coordinates file. The *Aggregation* option allows users to choose a preferred way on how to aggregate values in the given heatmap column. The default aggregation value is set to *First* and this value is recommended only for maps where multiple records for one point in the map are not expected. In other cases, it is better to choose a more appropriate aggregation.

To set a heatmap, click on the  *Set Heatmap* icon on the toolbar. The *Set Heatmap* dialog appears in the left pane. At first check off the *Enable Heatmap* option and select a type of the heatmap - *Transition* or *Threshold*. The *Set Heatmap* dialog will be adjusted accordingly:

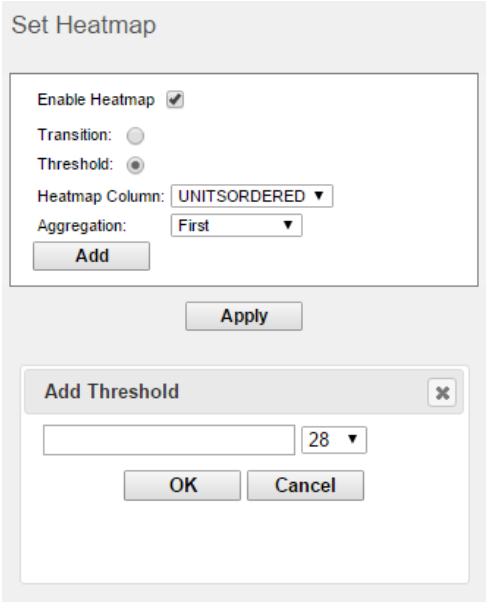
Transition Heatmap: Select a *Heatmap Column* from the drop-down list, select an *Aggregation*, and set minimal and maximal value of the selected column values. You can also choose a value from the drop down lists which contain the precise values from the selected column. Select a color gradient and click *Apply* to apply your settings on the map.



Set Transition Heatmap

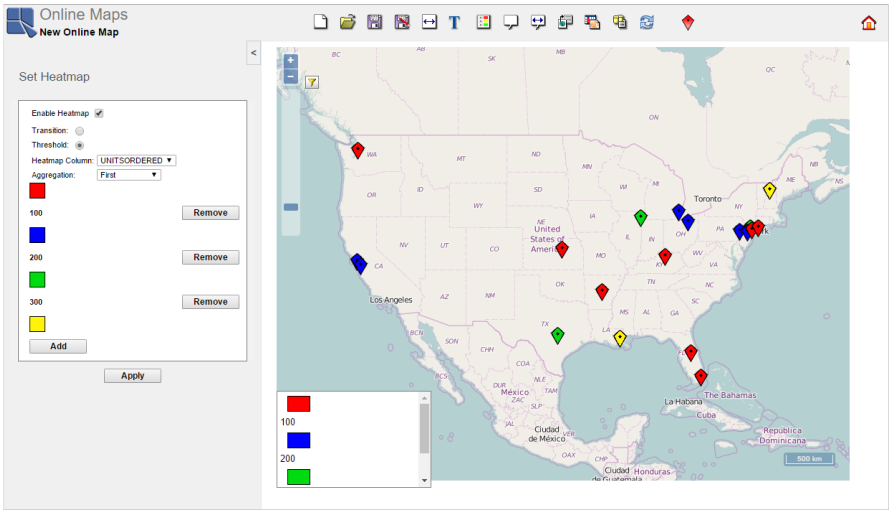
You can enable/disable a heatmap by checking/unchecking *Enable Heatmap* option.

Threshold Heatmap: Select a *Heatmap Column* from the drop-down list, select an *Aggregation*, and click the *Add* button to add a threshold. *Add Threshold* dialog will appear.



Set Threshold Heatmap

Type a threshold value or choose a value from the drop-down list which contains precise values from the selected column and click the *OK* button. Set another thresholds in the same way and set colors for each threshold by clicking on the color square and selecting the color. Then click *Apply* to apply the heatmap.

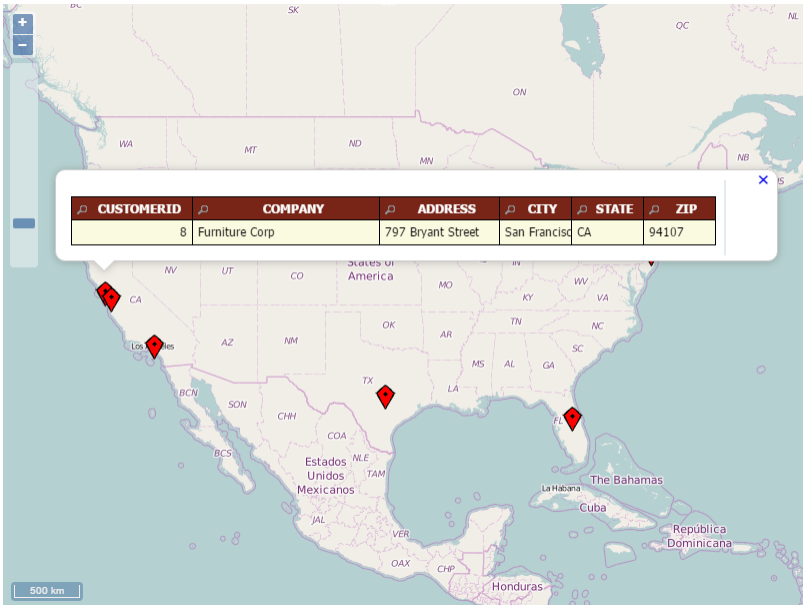


Set Threshold Heatmap

You can remove each threshold by clicking the *Remove* button next to the threshold. To enable/disable a heatmap, check/uncheck *Enable Heatmap* option.

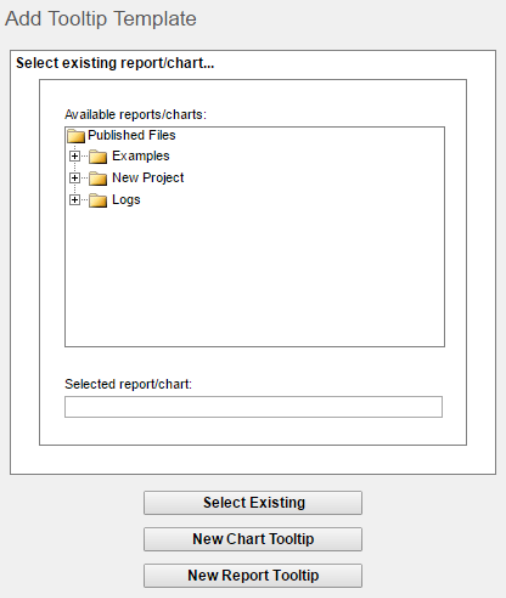
6.2.5.7. Tooltips

Tooltips may contain a report or chart that shows up when you move your mouse cursor over any map marker on the Online Map. Tooltips display single or multiple-row data that is related to the particular map point. As an example, a report tooltip will display data from the map datasource using parameter(s) from the *Map Data* of the corresponding map point. The *Map Data* is the same as the list in the *Change Point Mapping* dialog box. The report or chart file that is displayed in the tooltip bubble is called **Tooltip Template**.



Online Map with Tooltip

To enable tooltips for ERES Online Map, open the map in Online Maps and click on the  *Tooltip Template* icon. The *Add Tooltip Template* options will appear in the left pane.



Add Tooltip Template

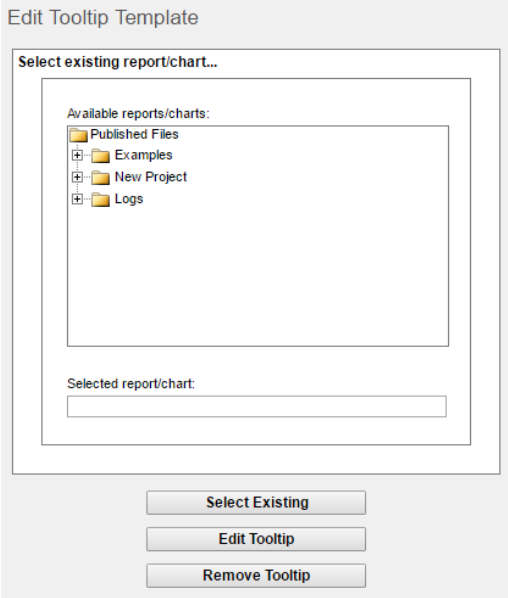
You can select an existing tooltip, create chart tooltip, or create report tooltip.

Select Existing: To select an existing Tooltip Template, select it from *Available reports/charts* list and click the *Select Existing* button. Using this option, you can also use the report/chart templates created in the Report Designer/Chart Designer, so you are not limited to use just the QuickDesigner Charts or QuickDesigner Reports to design Tooltip Templates. In this case, it is recommended to use exactly the same data source for creating this Tooltip Template report/chart as for the map, otherwise there might be a problem with data mapping.

New Chart Tooltip: Launches QuickDesigner Charts that will allow you to create a new chart file. Create a chart, save it, and leave QuickDesigner Charts by clicking the *Return to Map Designer* icon on the toolbar. (For more information about working in QuickDesigner Charts, see Section 5.4 - QuickDesigner Charts.)

New Report Tooltip: Launches QuickDesigner Reports that will allow you to create a new report file. Create a report, save it, and leave QuickDesigner Charts by clicking the *Return to Map Designer* icon on the toolbar. (For more information about working in QuickDesigner Reports, see Section 5.3 - QuickDesigner Reports.)

To edit, select, or remove the current tooltip, click on the *Tooltip Template* icon. The *Edit Tooltip Template* options will appear in the left pane. Please note that from this dialog, you can only directly edit files created in QuickDesigner Reports and QuickDesigner Charts (`.qdr` or `.qch`).



Edit Tooltip Template

Select Existing: Allows you to select other Tooltip Templates.

Edit Tooltip: Launches QuickDesigner Report or QuickDesigner Chart where you can edit the current tooltip template if it's a `.qdr` or `.qch` file (i.e. files from one of the QuickDesigners).

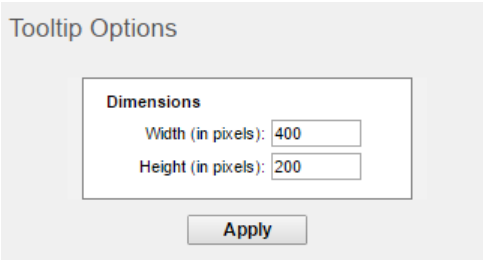
Remove Tooltip: Removes current Tooltip Template.



Note

To create a new chart/report tooltip you have to remove current tooltip at first and then will be *New Chart Tooltip* and *New Report Tooltip* options available.

To configure dimensions of the tooltip bubbles, click on the  *Tooltip Options* icon on the toolbar. Set new dimensions and click *Apply*.



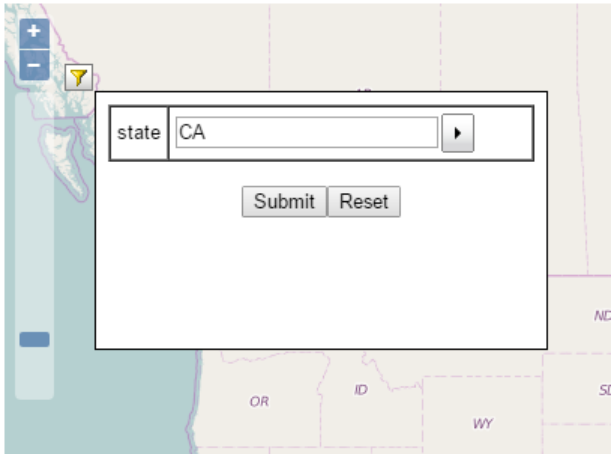
Tooltip Options

Changes will take effect after closing the current tooltip bubble (if a bubble is open, otherwise the changes will take effect immediately). It is not necessary to close the current tooltip bubble, you can just mouse over some other map point.

6.2.5.8. Parameter Setting

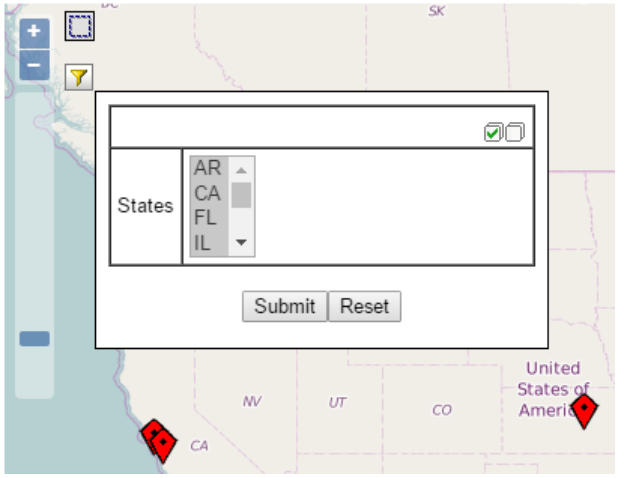
You can change parameters directly on the map by clicking the  *Set parameters* icon in the upper left corner of the map (The icon is displayed only if the data source is parameterized). A new dialog will open where you can select parameter value(s).

If the parameter is single-value, the next dialog will display a drop-down list. Select a parameter value and click the *Submit* button to apply it on the map. (You can reset your selection by clicking the *Reset* button.)



Parameter Setting

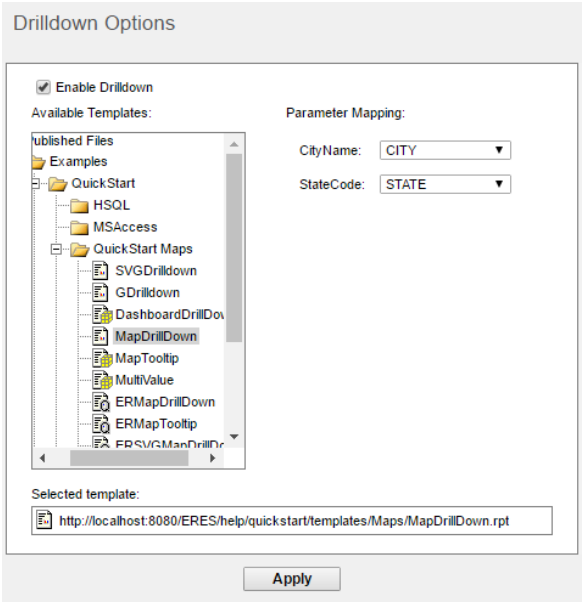
If the parameter is multi-value, the next dialog will display a list of parameter values. You can select parameter values by clicking on them (use **Ctrl+click** for multiple selection), select all parameter values by clicking the *Select All* icon, deselect all parameter values by clicking the *Select None* icon, and reset your selection by clicking the *Reset* button. Then apply your selection to the map by clicking the *Submit* button.



Parameter Setting

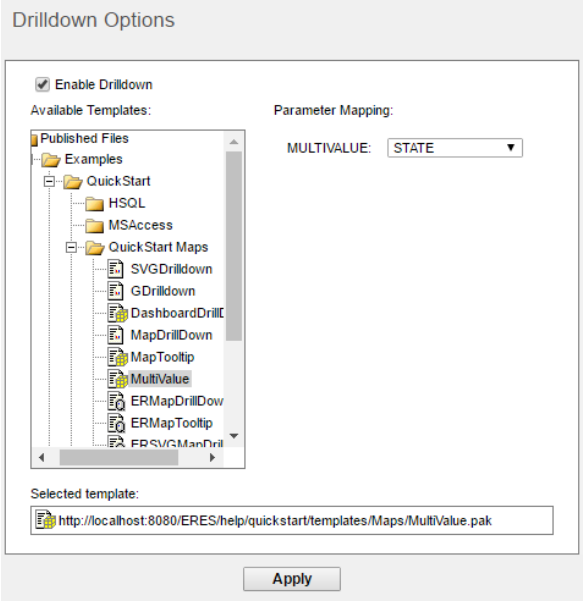
6.2.5.9. Drilldowns

To enable drilldowns in a map, click on the *DrillDown Options* icon. Check the *Enable Drilldown* checkbox which will show you all parameterized reports, charts and maps saved in Organizer. Click on a suitable parameterized report, chart, or map and select parameter mapping. All drilldown parameters have to be mapped to some Point ID field from the Coordinates. Mapping in this case means that the Point ID values of the selected map point will be used as drilldown parameter values (each map marker has unique Point ID). When you click on a map marker on the map, the drilldown report/chart/map will open in a new browser tab.



Drilldown Options

If all the drilldown parameters are multi-value (see Section 4.1.2.2.2.1 - Multi-Value Parameters), you will be allowed to select multiple map markers at once. For example, you can use `Examples/QuickStart/QuickStart Maps/MultiValue` data source, map *MULTIVALUE* parameter to *STATE* column, and then click *Apply*.



Multi-value Data Source

Selecting multiple map markers can be done by clicking the *multi-value Drilldown* button (see the image below; the button is displayed only if the map contains multi-value Drilldown).



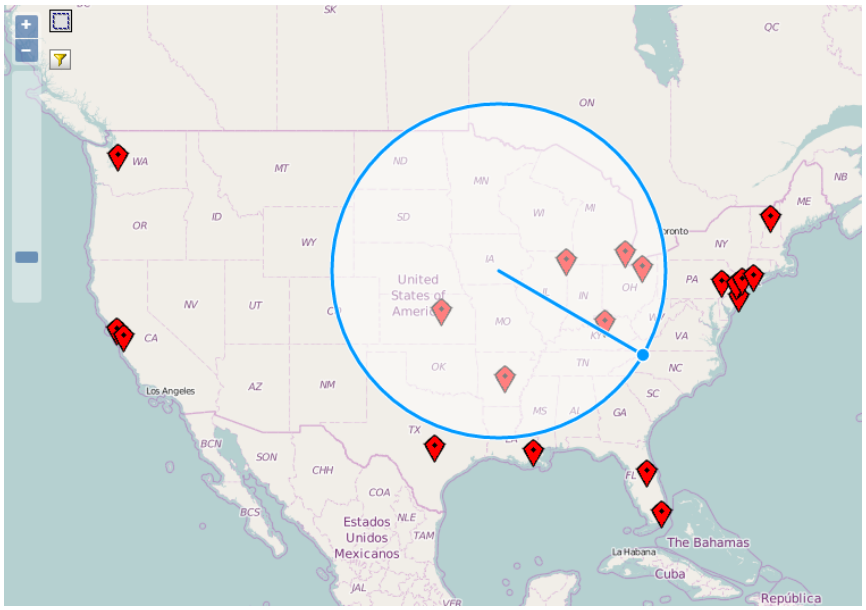
Multi-value Drilldown Button

After clicking the *Multi-value drilldown* button, the *Multivalue drilldown* dialog will appear.



Multi-value Drilldown Dialog

You can choose a shape of the bounding box (rectangle or circle). The circle can be fixed or drawn. Once you select a shape, click the *OK* button. Then you can select multiple map markers by dragging using bounding box. Left click on the map (this determines rectangle corner or circle center), hold mouse button down and drag it, and then release the mouse button to determine size of the bounding shape.



Bounding Box

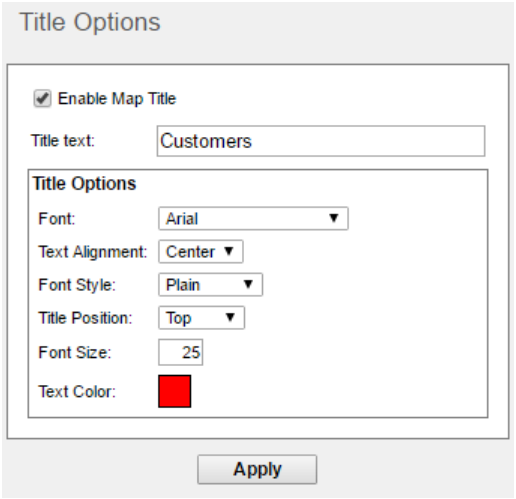
If you select fixed circle as a bounding box, you can set a diameter of the circle. Write a diameter value and click the *OK* button. You will see a circle with the diameter you just set. Then click to the map. Your click is a center of the circle bouding box. A new tab with a drilldown of your selection (markers within the circle) will appear.



Fixed Circle Bounding Box

6.2.5.10. Map Title

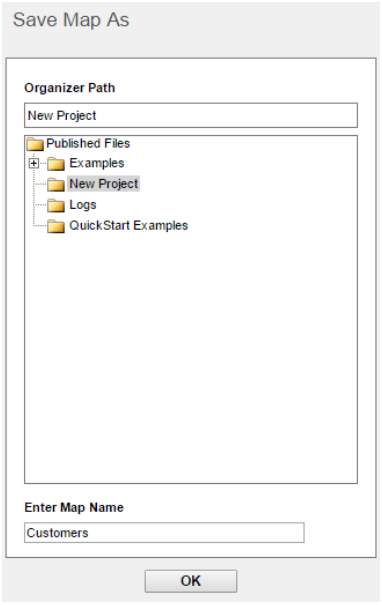
To enable map title, click on the  *Map Title Options* icon and check the *Enable Map Title* checkbox. Then you can insert the map title text and select various options - font, text alignment, font style, title position, font size (positive integers only), and text color. The color can either be chosen from swatches or inserted manually by entering red, green and blue components (all the values has to be integers from 0 to 255).



Map Title Options

6.2.6. Save Map

You can save the Online map by clicking the  *Save* button on the toolbar. This will open a dialog allowing you to specify a name for the map. Enter a name for the map and click the *OK* button to save it.

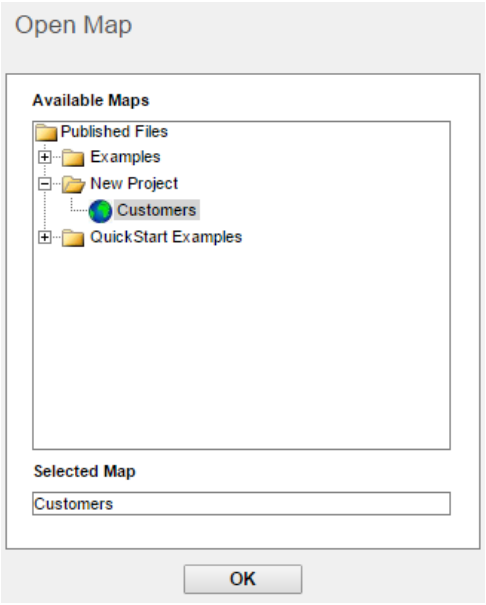


Save Dialog

After saving or opening an existing Online map, the  *Save As* button allows you to save the existing or modified map into a different name.

6.2.7. Open Saved Map

You can open a saved map by clicking the  *Open* icon on the main toolbar. The *Open Map* dialog will appear.



Open a Map

All Online maps created in Online Maps designer are visible in Organizer. Select a map and click the *OK* button to open it.

6.2.8. Exit

You can exit Online Maps designer either by clicking the  *Home* icon in the upper right corner, clicking the *Online Maps* title, or by clicking the  *Logo* icon in the upper left corner. Before closing, you will be asked if you want to save an unsaved map.

6.3. SVG Maps

6.3.1. Introduction to SVG Maps

Unlike Online Maps, SVG Maps do not use map points. They use **map areas**. The map areas can be colored based on values from the map data source. Coordinates of the map areas are defined in **SVG map image**. Virtually any SVG image can be used for SVG Maps, but it usually requires some modifications, which can be done in any graphical editor with SVG support. ERES does not contain any integrated graphical editor, so you have to use third party software. The recommended software is **Inkscape**, which can be downloaded at www.inkscape.org [https://inkscape.org/]. Inkscape is released under GPL license version 2 [http://www.gnu.org/licenses/old-licenses/gpl-2.0.html], so it is free for commercial and non-commercial use.

SVG (Scalable Vector Format) images consist of objects. These objects can be used as map area. To distinguish among individual map areas, ERES uses **Area ID**. Each map area must have an Area ID, but unlike Point ID, Area ID does not take up more than one field (it always consists of only one field). The Area ID has to be unique for every area and they are read from the SVG Image object ID attribute. This attribute can be set in most graphical editors. Please see the next chapter for a step-by-step guide for the recommended tool - Inkscape.

ERES is distributed with several SVG images that are already prepared to be used for creating SVG Maps without any modification. These images can be found in `<ERES_Installation_directory>/MapFiles/SVG`. They are also included in Organizer in SVGMaps project. You can find the list of all these maps in the Appendix 6.A - List of SVG Map Images.

6.3.2. Area ID Mapping

SVG Map needs an SVG image that defines the map areas. All SVG objects in the SVG image can be used as map areas. The only requirement is that they must have an Area ID assigned as described in the previous section. These Area IDs are used for mapping the map data source rows to map areas (SVG objects). Unlike the Point IDs, Area IDs cannot have multiple fields, so the mapping is much simpler. It is enough to select just one data source column, which will be mapped to the Area ID. A map area is associated with the data source row, which has the value of the mapped column equal to the Area ID (comparison is case sensitive).

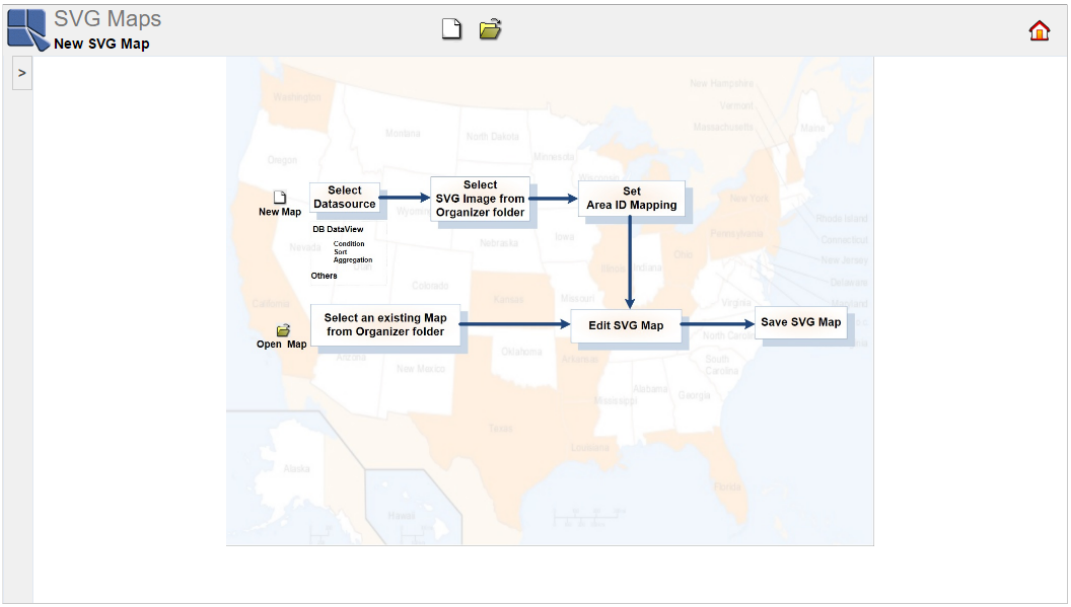
In the SVG Maps, the map data source is used for coloring areas according to numerical value of a data source column. It brings one more limitation for the map data source compared to Online Maps. There has to be at lest one matching data source row for one Area IDs. If there is no matching row, the area will have original (unchanged) color. If there is exactly one, the fill color of the matching map area will be changed according to Thresholds (see Section 6.3.3.5 - Thresholds). If there is more than one matching row, the behavior is undefined. There will be no error message, but the data displayed may not make sense.

6.3.3. SVG Maps Designing

6.3.3.1. Start

SVG Maps can be started from the ERES Start page. If you have logged in as a user with design privileges, then you can follow SVG Maps link to begin using SVG Maps.

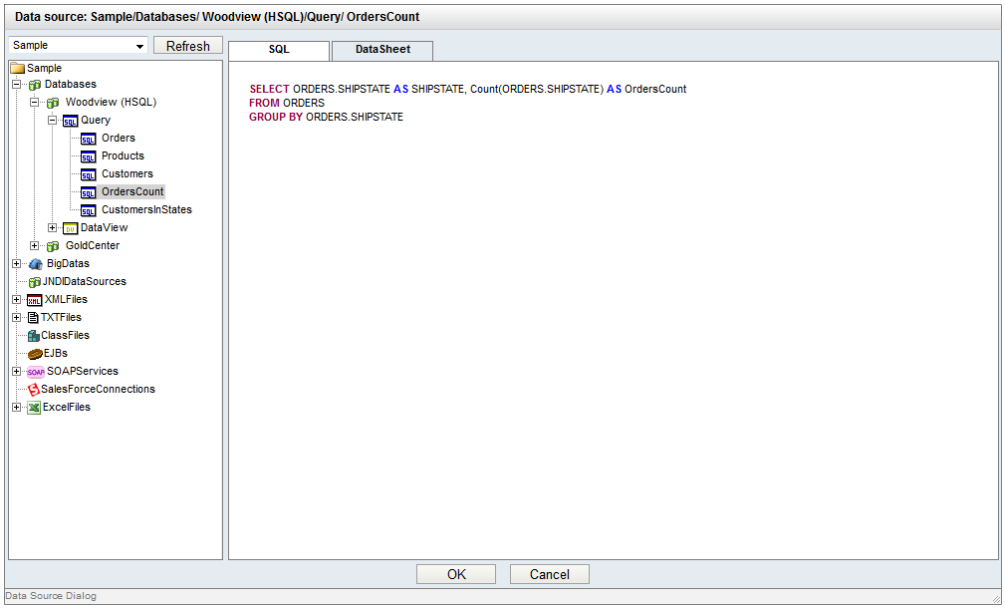
When the SVG Maps designer launches, the first page that appears prompts you to select between opening an existing map, or creating a new one.



SVG Maps Start Options

6.3.3.2. Select Data Source

If you want to create a new SVG map, click on the  *Create New Map* icon on the toolbar. The *Data Source Dialog* appears and prompts you to select data source you want to use with SVG Maps.



Data Source Dialog

Unlike Report Designer or Chart Designer where the user can select to create new data sources or modify existing ones, the only option is to select a data source for the map. The only exception to this is for Data Views or Data View Queries. If data sources of this type are selected, you will see the *DataView Builder* on the right side of the Data Source Dialog. The *DataView Builder* allows you to build or modify queries against the View. For more information about managing data sources, please see Section 4.2 - Data in QuickDesigners and Maps.

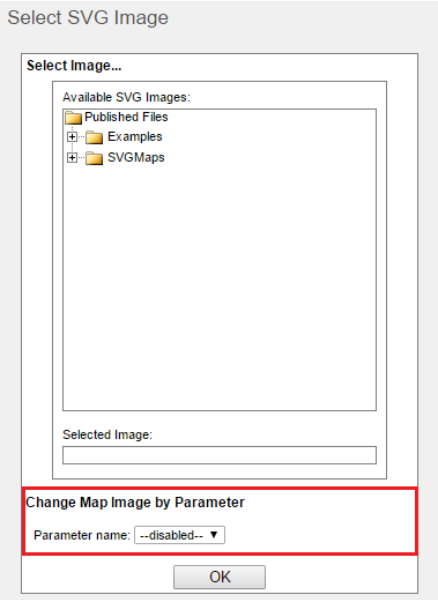
Once you finish selecting a data source, click the *OK* button to close the *Data Source Dialog*.

6.3.3.2.1. Change Data Source

To change a data source (for SVG maps that have already been created), click the  *Change Data Source* icon on the toolbar. The *Data Source Dialog* will open and you will see current data source. Select other data source from the tree list and then click the *OK* button to apply it.

6.3.3.2.2. Parameterized Data Source

You can use parameterized data source for SVG maps. For more information about parameterized data source in SVG maps, please see Section 6.3.4 - Dynamic SVG Maps.

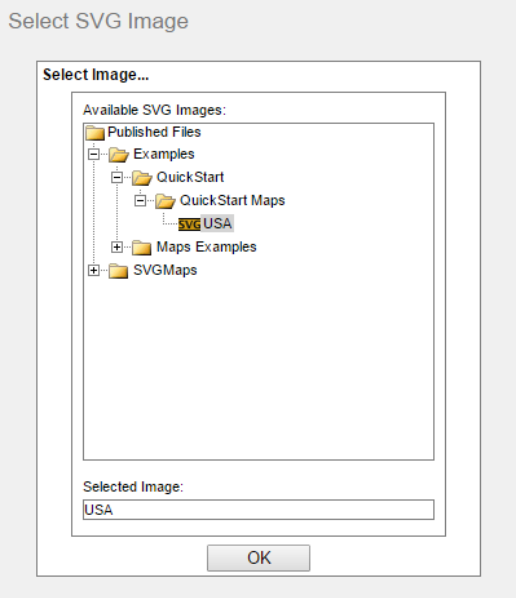


Selecting Parameterized Data Source

To change a parameter from the toolbar of the SVG Maps designer, click the  *Refresh/Set Parameters* icon on the toolbar. You will see available parameters. Set the parameter(s) and click *Submit* to apply this setting.

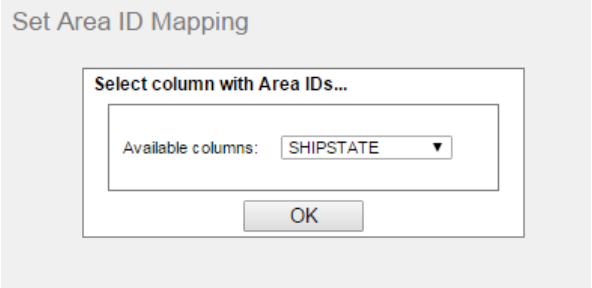
6.3.3.3. Select SVG Image

Now you can see the *Select SVG Image* dialog on the left side of the SVG Maps designer. This dialog prompts you to select a Map Image. Map Images are SVG image files with data structures containing geographic data (see Section 6.3.2 - Area ID Mapping for more information about required image format) and must be inserted into the Organizer before they can be used.



Select SVG Image Dialog

The *Select SVG Image* dialog contains a tree-list showing all projects, folders, and SVG images inserted in the Organizer. Select an SVG image you want to use as the map image and click the *OK* button. The *Set Area ID Mapping* dialog appears. Select a data source column that should be mapped to the Area IDs and then click the *OK* button.



Select Column with Area IDs

SVG map will open (*Set Area ID Mapping* dialog is collapsed on the screenshot).

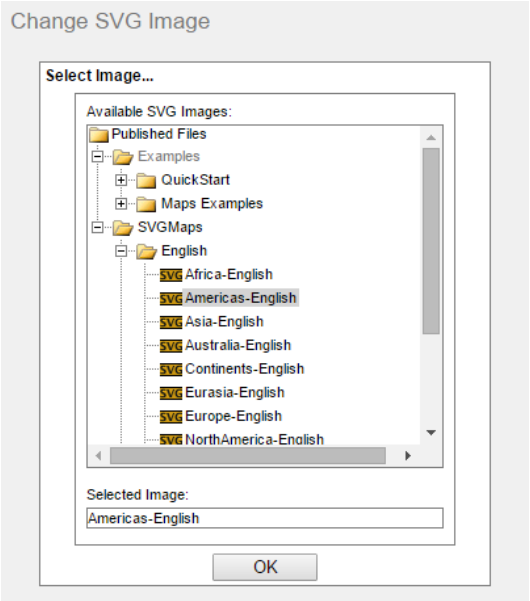


SVG Maps Interface

As you can see, right now we have only the basic map with no data highlighted. There are two ways of adding data to the SVG Map: Thresholds and Drilldowns.

6.3.3.3.1. Change SVG Image

To change an SVG image (for SVG maps that have already been created), click the  *Change SVG Image* icon on the toolbar. Select an image from the tree list and then click the *OK* button to apply the image.



Changing SVG Image

6.3.3.4. Toolbar

A lot of formatting actions for reports in SVG Maps are accessed through the toolbar. The icons perform the following actions:

 Start a new SVG map

 Open an existing SVG map

 Save the current SVG map

 Save the SVG map as

 Set a map size

 Insert/Edit a map title

 Set thresholds

 Insert/Edit a tooltip

 Insert a drilldown

 Change area ID mapping

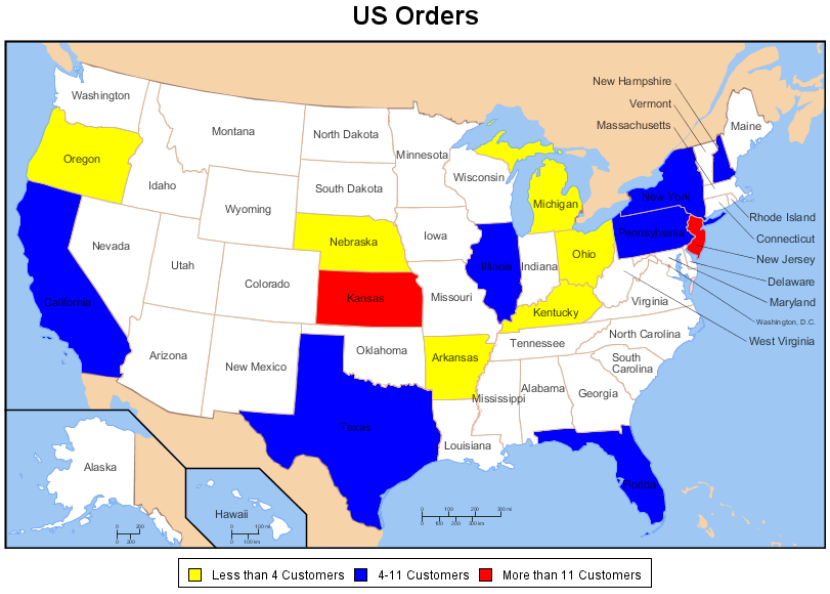
 Change SVG image

 Change data source

 Refresh/Set parameters

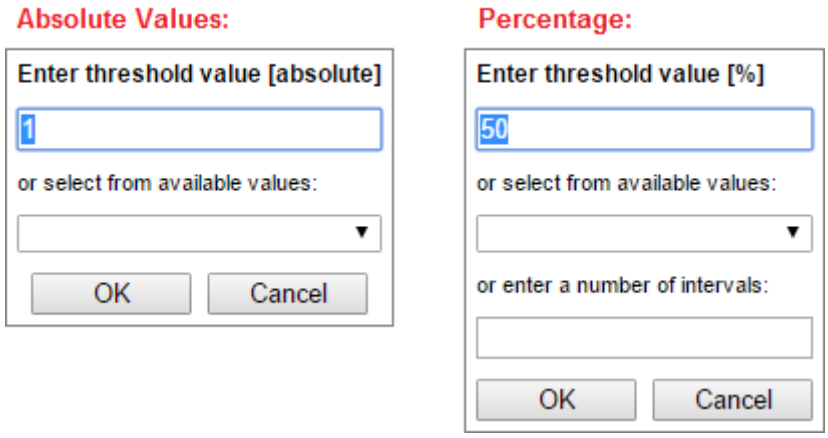
6.3.3.5. Thresholds

Thresholds allow you to differentiate areas by color based on value in a particular data source column. If a threshold is set, all the map areas (i.e. objects with a record in the map data source) are colored with color assigned to the appropriate value range.



Example: SVG Map with Thresholds

To set the thresholds, click on the *Set Thresholds* icon on the toolbar and select a threshold column (i.e. column that will be compared to the ranges defined below). Only numerical column can be selected. You can add a threshold value(s) in actual values or percentage unit. To add a threshold value, click the button next to the empty table field. A dialog box will open prompting you for the new threshold value.

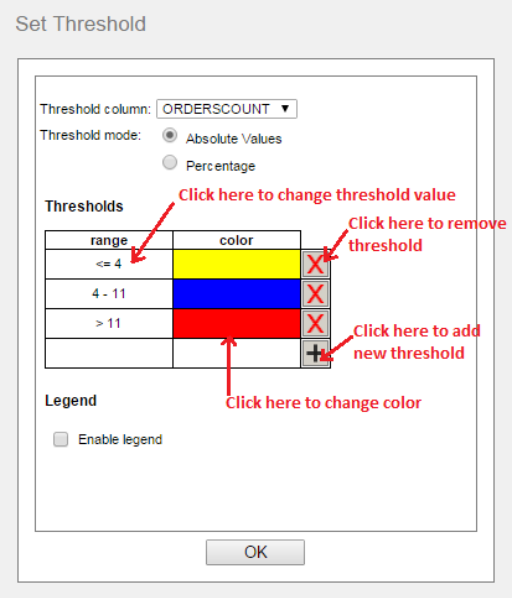


Entering Threshold Values

You can either enter it manually or select one of the suggested values from the drop-down list at the bottom of the dialog. This list contains all the values from the selected threshold column rounded to integers (There is a third way to enter threshold values for percentage units. It is an option to enter a number of intervals). After adding the first threshold value, two ranges are added - one for values below and one for values above the given threshold. Any threshold value added after that will create only one new range by splitting one of the old ranges in two. You can edit any of the existing thresholds by clicking on the appropriate range (the first column in the table).

To change the color assigned to a value range, click on the empty table field next to the value range you want to change the color of. You can then select one of the predefined colors or choose any color by entering its red, green and blue component values (each value is an integer within the 0-255 range). To remove a threshold value, click on the red button next to the value you want to remove. Changes will take effect after you confirm the *Set Threshold* dialog by clicking the *OK* button.

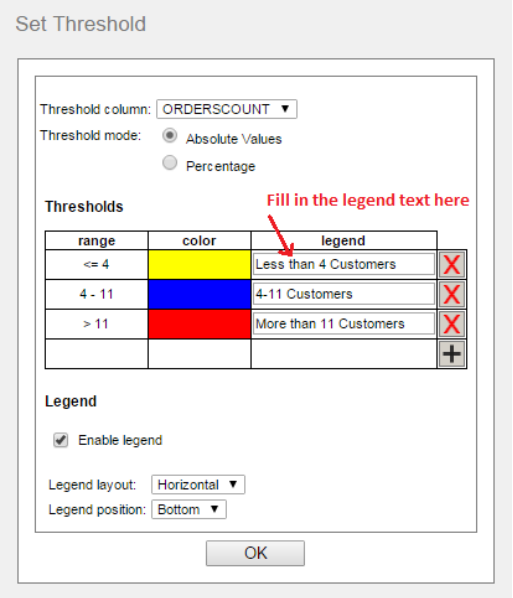
To remove all thresholds and disable coloring of areas, delete all threshold values or select an empty field for the threshold column.



Set Threshold Dialog

6.3.3.5.1. Legend

The *Set Threshold* dialog also allows you to add a legend to describe the value ranges. To enable it, check the *Enable Legend* checkbox. Then you can select a layout and position of the legend. If the legend is enabled, legend input fields next to the value ranges will appear. Insert a description of the ranges in these fields. If the legend is too large to fit one column (for vertical layout) or row (for horizontal layout), it will be split into several columns/rows. If it still does not fit the map canvas even after splitting, it will not be displayed at all. If this ever happens, you can try to resize the map or change the position or layout of the legend.

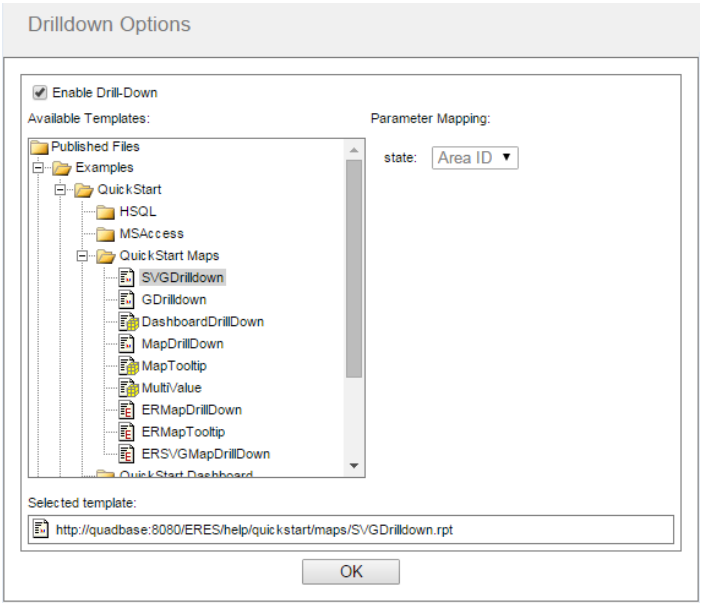


Setting Legend

6.3.3.6. Drilldowns

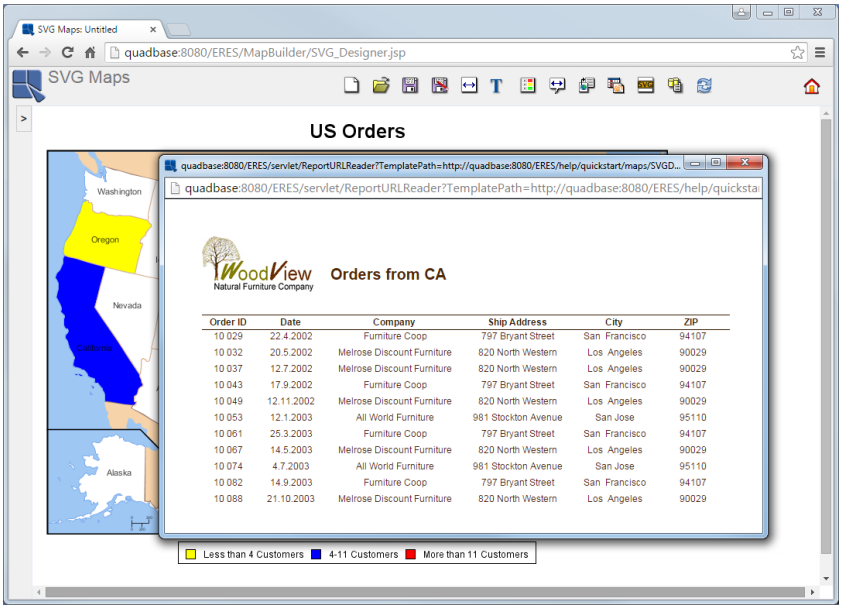
Click on the  *Drilldown Options* icon on the toolbar and select *Enable Drilldown* option. The Organizer structure will show up in the *Available Templates* treeview. Select a parameterized report, chart, or map. All parameters will be automatically mapped to the Area ID. This means that all parameters of the drilldown report, chart or map will be set to the Area ID of the selected area. It is not possible to change the mapping. SVG map also does not

support multi-value drilldowns (multi-value parameters are treated as single-value). Once you select a file, click the *OK* button to open it.



Drilldown Options

When you click in an area matching the ID mapping (e.g. a colored state), a drilldown report will open in a new window.

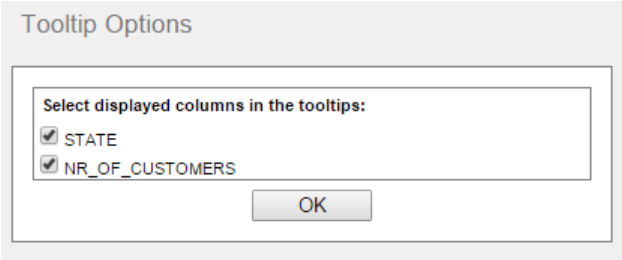


Map with Drilldown

6.3.3.7. Tooltips

Tooltips in SVG maps are different than tooltips in Online maps. In SVG maps, there are no reports in tooltips, just data from data source fields associated with particular map area.

To set up tooltips in a SVG map, click on the  *Tooltip Options* icon on the main toolbar. The following dialog will appear:



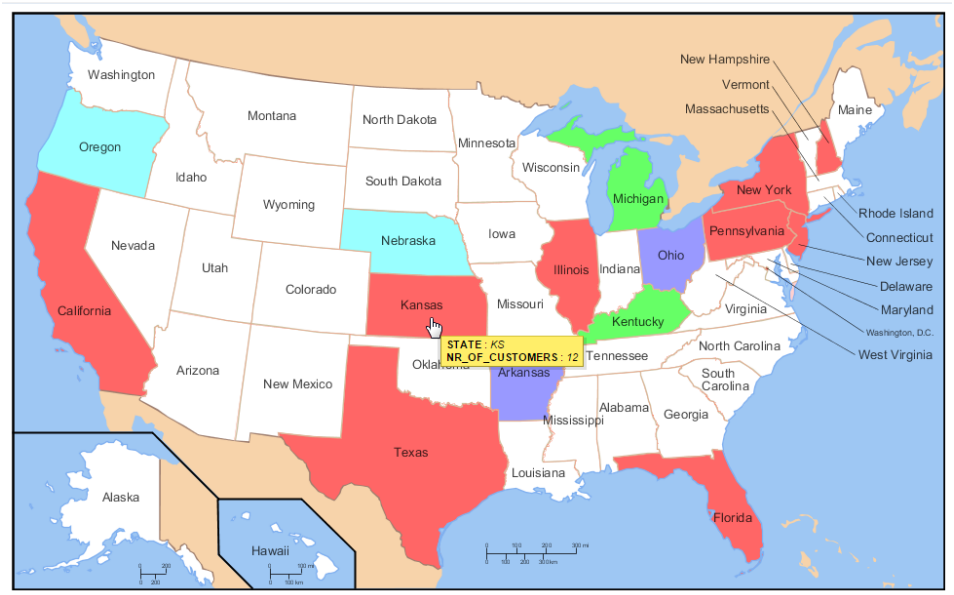
Tooltip options dialog

On this dialog, you can choose which data source fields will be displayed in tooltips, which will pop up in yellow background when you move your mouse cursor over a map area.



Note

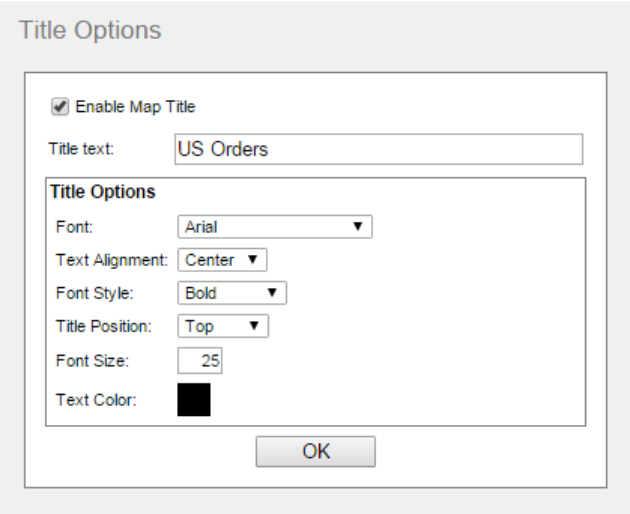
If you deselect all options on the *Tooltip options* dialog, default browser tooltips will be displayed (in grey background), showing SVG area IDs.



SVG Map with a Tooltip

6.3.3.8. Map Title

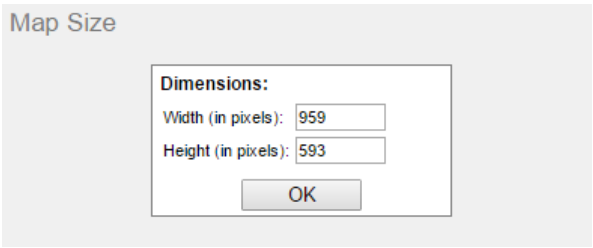
To insert a map title, click the *Map Title* icon on the toolbar. Check the *Enable Map Title* checkbox. Then you can insert the map title text and select various options - font, font style, font size (positive integers only), text alignment, title position, and text color. The color can be either choosed from swatches or inserted manually by entering red, green and blue components (all values has to be integers from 0 to 255). To apply your settings, click the *OK* button.



Map Title

6.3.3.9. Map Size

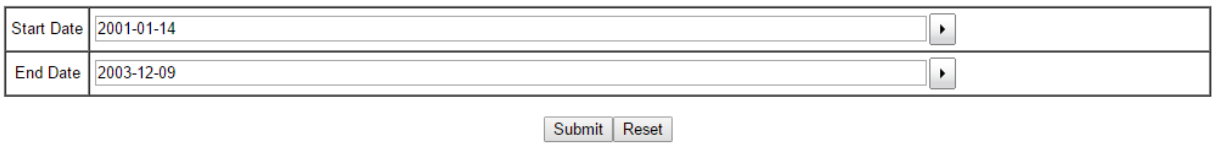
To set a map size, click the  *Set Map Size* icon on the toolbar. Enter dimensions and then click the *OK* button to apply your setting.



Map Size Setting

6.3.3.10. Change Parameter Values

To change parameter values (for maps with parameterized data source) or refresh the map, click the  *Refresh* icon on the toolbar and the Parameter setting dialog will appear. Select a parameter value(s) and then click the *Submit* button to apply parameter values on the map (You can reset your selection by clicking the *Reset* button).



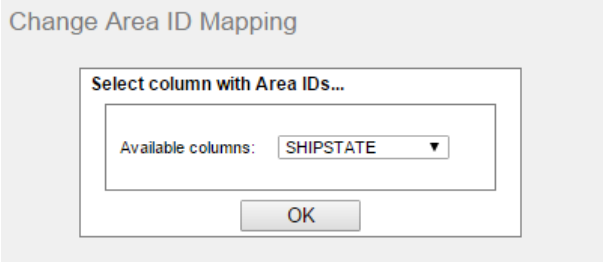
Parametr Setting Dialog for Single-value Parameters



Parameter Setting Dialog for Multi-value Parameters

6.3.3.11. Change Area ID Mapping

To change area ID mapping, click the  *Change Area ID Mapping* icon on the toolbar. Select a column from the drop down menu and then click the *OK* button to apply your setting (For more information about area ID mapping, please see Section 6.3.2 - Area ID Mapping).



Changing Area ID Mapping

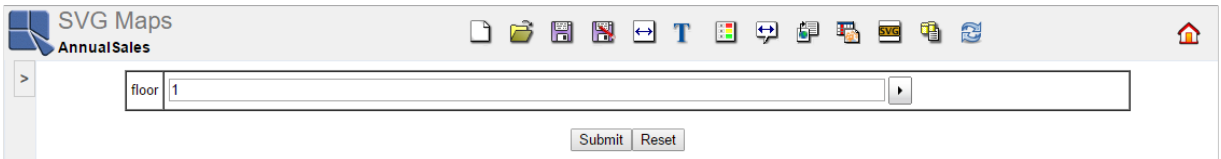
6.3.4. Dynamic SVG Maps

ERES Dynamic SVG Maps allow changing SVG image according to a parameter, i.e. you can use several SVG images in a single SVG map. Dynamic SVG map can be used also as a drilldown, so you can create useful SVG map with drilldown - for example continent to country, country to state, state to county, etc.

There are some rules and recommendations for Dynamic SVG Maps:

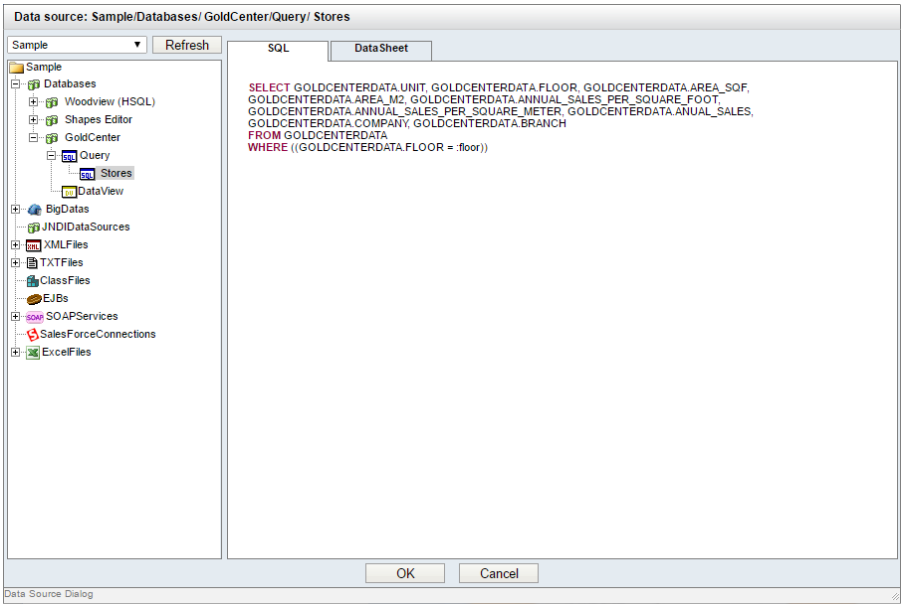
- SVG images have to be in format: `imageName_PARAMETER.svg`, where `imageName` is the same for all SVG images and `PARAMETER` is used as the parameter.(e.g. `worldmap_political, worldmap_physical, worldmap_topographic`,...)
- Data source has to be parameterized where one of parameters takes the same values at the end of the SVG images file names (after the underscore character).
- SVG images have to be saved inside a single folder.
- One of the SVG images has to be in Organizer. This image is used as default for creating a map and it is also used in case the SVG image is not available due to a wrong parameter, missing file, etc.
- It is recommended to have all SVG images in the same size.

It is possible to change SVG Images by clicking the  *Refresh/Set Parameters* icon on the toolbar. You will see a drop down menu for selecting a parameter. Select a parameter and click the *Submit* button. The corresponding image will be used for the map. All settings of the map (thresholds, legend, title, tooltips, drilldowns) will be applied to each SVG image.



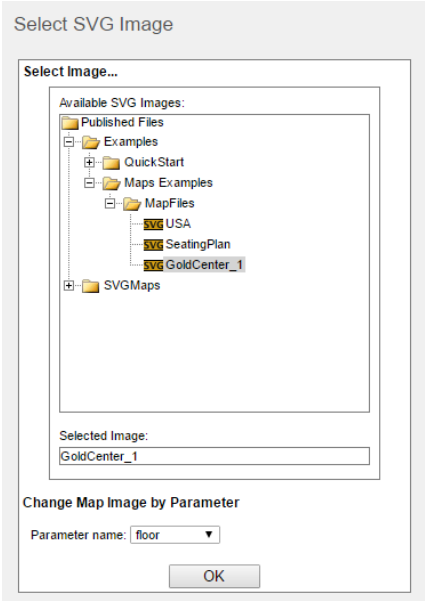
Parameter Selection

Let's take a look at the following example of a shopping center with three floors. We want to display the status of annual sales per square foot of stores. Open SVG Maps designer (from the ERES main page) and click the  *Create New Map* icon on the toolbar. *Data Source Dialog* will appear. Select *Sample* data registry and then select `Databases/GoldCenter/Query/Stores` query as a data source. Click the *OK* button to close the *Data Source Dialog*.



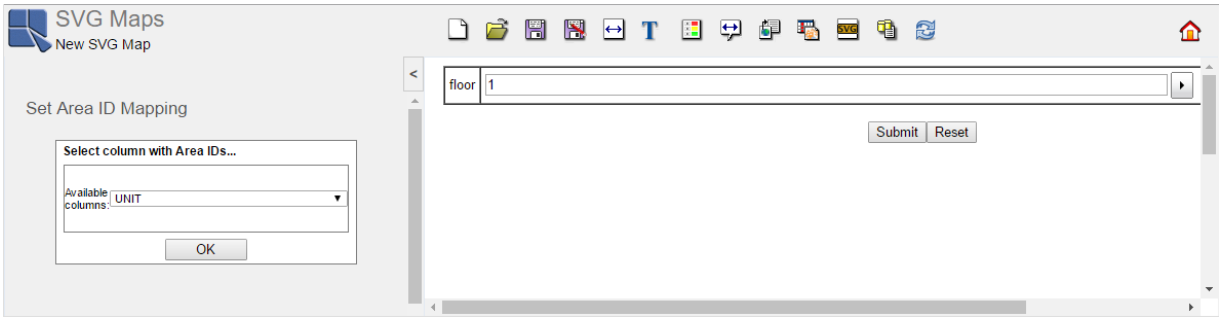
Select Data Source

Now you are prompted to select an SVG image. Select `Examples/Maps Examples/MapFiles/GoldCenter_1` SVG image. We have three SVG images for this example (`GoldCenter_1.svg`, `GoldCenter_2.svg`, and `GoldCenter_3.svg` in the *ERES/help/examples/Maps/MapExample4* folder), but you can see only one of them in the *Select SVG Image* dialog because only one was inserted into Organizer (it is not necessary to insert all SVG images to Organizer). Next select a parameter for changing SVG images in the *Change Map Image by Parameter* dropdown menu. Select the parameter *floor* and click the *OK* button.



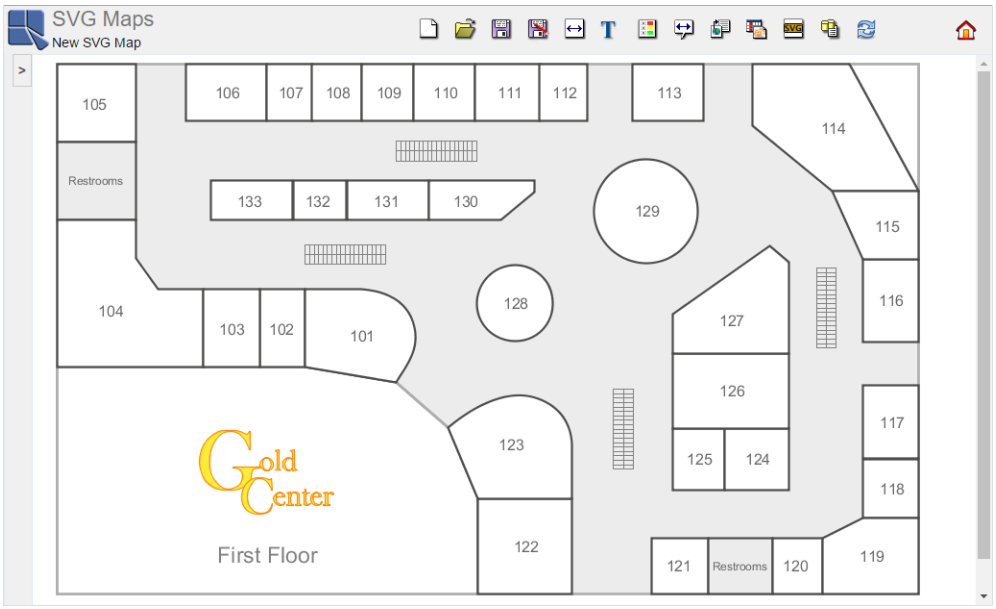
Select SVG Image

Next step is to map IDs to the data source column. Select *UNIT* and click the *OK* button. You can also select a parameter, a number of floor in our case, in this step. Leave the first floor selected and click the *Submit* button to open a map.



Set Mapping and Parameter

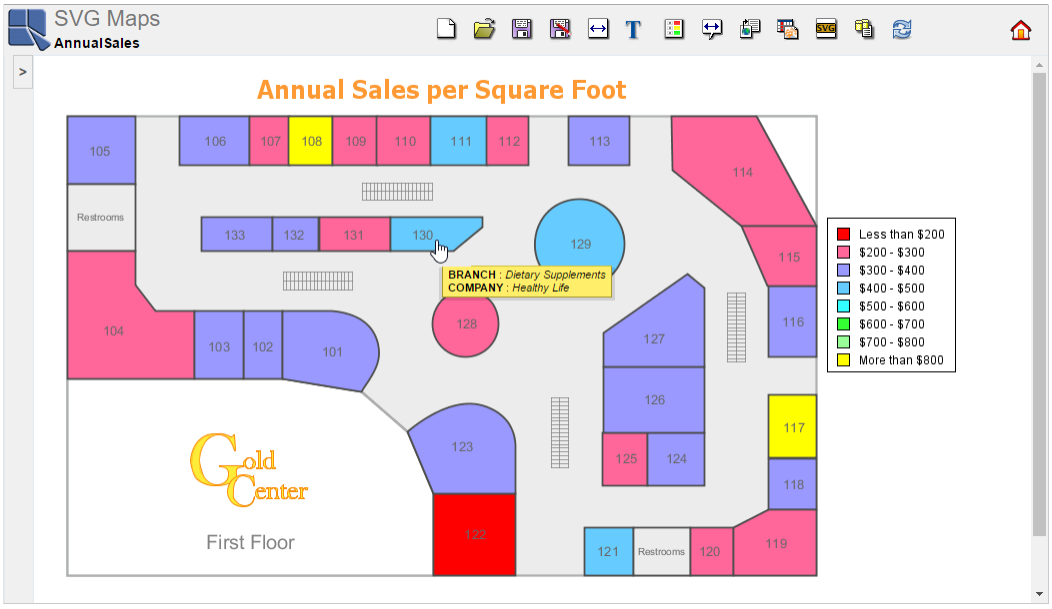
You can now see a map of the first floor (Left pane is collapsed in the next image).



Set Mapping and Parameter

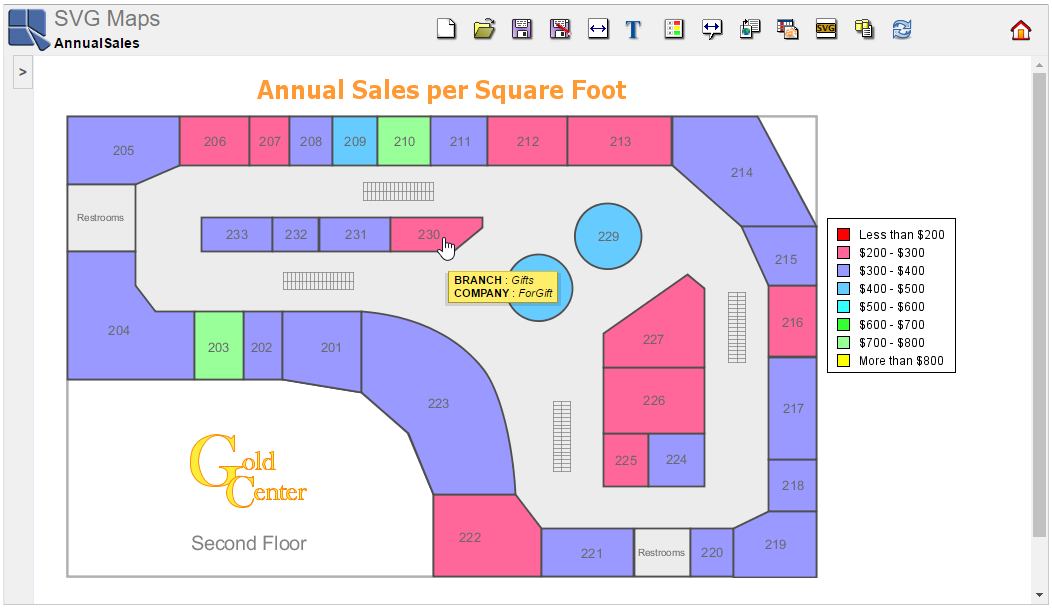
To change the SVG image (the other floor of the shopping centre in our case), you only need to change the parameter value by clicking the  *Refresh/Set Parameters* icon and selecting parameter value from the dropdown menu.

The next image shows the finished SVG map (you can find this SVG map in `Examples/Maps Examples/MapFiles/AnnualSales`). There are set thresholds (for more information, see Section 6.3.3.5 - Thresholds), a legend (for more information, see Section 6.3.3.5.1 - Legend, a title (for more information, see Section 6.3.3.8 - Map Title), and a tooltip (for more information, see Section 6.3.3.7 - Tooltips).



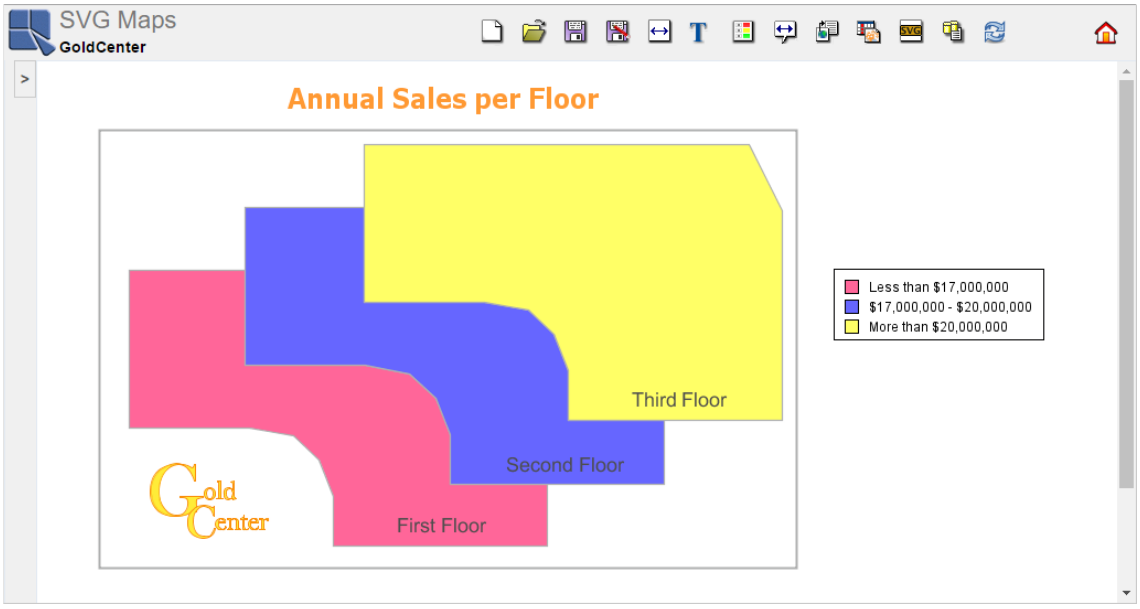
Finished Map

The next image shows the SVG map for the parameter value *2*, which is the second floor of the shopping center. All adjustments are also valid for the SVG image of the second floor.



Changed SVG Image

Dynamic SVG map can also be used as a drilldown. For example, open `GoldCenter` SVG map from `Examples/Maps Examples/MapFiles` directory. You should see this map:



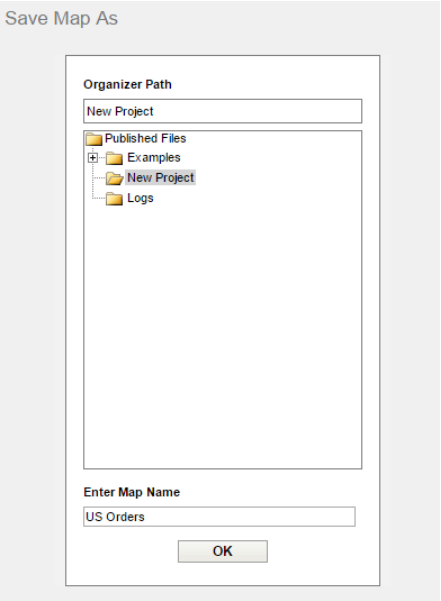
Basal SVG Map

Add a drilldown by clicking the  *Drilldown Options* button, select `Examples/Maps Examples/MapFiles/AnnualSales` file in the *Drilldown Options* dialog and click the *OK* button. Then you can click on the area of some floor to open the drilldown.

You can also open the finished SVG map with drilldown `GoldCenterWithDrilldown` under `Examples/Maps Examples/MapFiles` node.

6.3.5. Save Map

You can save the SVG map by clicking the  *Save* button on the toolbar. This will open a dialog allowing you to specify a name for the map. Enter a name for the map and click the *OK* button to save it.

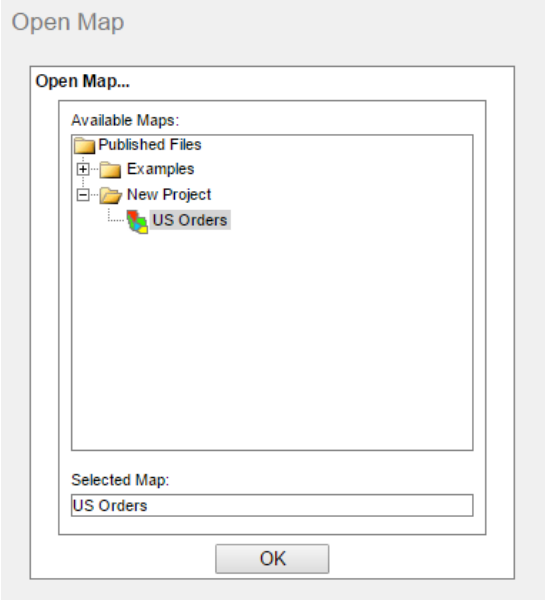


Save Dialog

After saving or opening an existing SVG map, the  *Save As* button allows you to save the existing or modified map into a different name.

6.3.6. Open Saved Map

You can open a saved map by clicking the  *Open* icon on the main toolbar. The *Open Map* dialog will appear.



Open the SVG Map

All SVG maps created in SVG Maps designer are visible in Organizer. Select a map and click the *OK* button to open it.

6.3.7. Exit

You can exit SVG Maps designer either by clicking the  *Home* icon in the upper right corner, clicking the *SVG Maps* title, or clicking the  *Logo* icon in the upper left corner. Before closing, you will be asked if you want to save an unsaved map.

6.4. Map Viewer

Map Viewer is a JSP page used for viewing maps. It is located in `http://machine:port/<ERES_CONTEXT>/MapBuilder/MapViewer.jsp` and it can be used for both Online Maps and SVG Maps. If you open a map from Menu Page, it is automatically opened in the Map Viewer (If you open a map from Organizer, it is opened in corresponding map designer). Another way to use the Map Viewer is to write your own map URL. This is described in the next chapter.



ERES Maps Viewer (showing SVG Maps with threshold)

6.4.1. Writing Map Viewer URL

You can write your own Map Viewer URL. The URL always begins with the location of the Map Viewer JSP: `http://<machine>:<port>/<ERES_CONTEXT>/MapBuilder/MapViewer.jsp?`. Following the question mark you need to specify Map Viewer parameters. Parameters are separated by the "&" character.

The most important Map Viewer parameter is **filePath**. It specifies the path to the map you want to display. The value can be either URL or file path on the server. The file path can be absolute or relative to the ERES installation directory. This parameter is mandatory and must always be specified.

The remaining Map Viewer parameters are used for setting map parameters. These parameters have no function in non-parameterized maps, so they are ignored if the map is not parameterized. The parameters are:

ForHTMLParamPage: This is a Boolean flag that indicates whether or not to display parameter page before displaying the map. If this argument is set to true then the URL will return the parameter page and ignore any parameter values passed into the URL. This argument is set to false by default.

QueryParamName: This parameter is used to specify parameter for parameterized map. QueryParamName specifies the name of the parameter you will be supplying a value for (Parameter names are specified when you create the query. For more information about this, see Section 4.1.2.2.2 - Parameterized Queries). It is always followed by QueryParamSize and QueryParamValue parameters.

QueryParamSize: This parameter is used to specify the number of values that will be passed into a particular map parameter. As described in Section 4.1.2.2.2.1 - Multi-Value Parameters, some queries can have multi-value parameters. This URL parameter allows you to specify if the parameter (indicated by the previous QueryParamName parameter) takes multiple values. This parameter is always followed by one or more QueryParamValue arguments for each of the parameter values for a particular parameter. If you do not specify this parameter, it will assume that the number of the values to be passed in is one.

QueryParamValue: This parameter is used to specify a parameter value you want to pass to the query parameter specified by the QueryParamName argument. For multi-value parameters, a separate QueryParamValue argument must be supplied for each distinct parameter value that you are passing in. For more information about creating parameterized queries, please see Section 4.1.2.2.2 - Parameterized Queries.

Example:

To open GoogleMapParam (one of the map examples - parameterized Online Map) with parameter `State` set to `CA`, use this URL:

```
http://machine:port/<ERES_CONTEXT>/MapBuilder/MapViewer.jsp?filePath=help/
examples/Maps/MapExam-
ple1/GoogleMapParam.gxml&QueryParamName=State&QueryParamSize=1&QueryParamValue=CA
```

Please replace `machine` with your machine name or IP, `port` with the port Tomcat is running on (8080 by default) and `<ERES_CONTEXT>` with your ERES context (ERES by default).

6.5. Migrating Maps

Often you may need to move maps from one location to another. For example, maps have to be moved between the server and the develop machine. On the develop machine the tooltips/drill-downs/svg images locations can change and may not be located from the data stored in the map `.gxml` or `.sxml` file. Therefore, ERES uses SPAK (for SVG Maps) and GPAK (for Online Maps) files to move or archive maps along with its components. To learn how to create a SPAK or a GPAK file, see Section 3.1.4.5 - Dashboard and Map Packages.



Note

The SPAK and GPAK files can be viewed in the MenuPage, but they cannot be opened in the SVG or Online map designers. To edit a map from a SPAK or a GPAK file, you have to unpack it first (to learn how to unpack DPAK or SPAK files, see Section 3.1.4.5 - Dashboard and Map Packages).

6.A. List of SVG Map Images

The following SVG map images are distributed with ERES. They are prepared to be used for creating SVG Maps. They all have uniform look - the land is gray and borders and oceans are white. You can change the colors in any SVG editor (e.g. Inkscape) or you can edit them directly in the SVG files (this requires at least basic knowledge of the SVG format).

The maps are split in two directories according to the Area IDs naming convention. The maps in the `<ERES_Installation_directory>/MapFiles/SVG/English` directory use English country (continent) names, while the maps in the `<ERES_Installation_directory>/MapFiles/SVG/ISO` directory use ISO 3166-1-alpha-2 [https://www.iso.org/iso-3166-country-codes.html] code. All Area IDs in these maps are in `UPPER CASE`. The English names are exactly the names from this table [https://www.iso.org/obp/ui/], only the spaces are replaced with underscores (_), because spaces are not supported in the SVG objects IDs. However, underscores are automatically converted to spaces before comparing them to the data source columns, so you should still use spaces in your data source.

These SVG map images are already inserted in the Organizer in the SVGMap project. Please note that this project is hidden in the Menu Page, but it can still be displayed and modified in Organizer and the files from this project can be used for creating maps.

This is the list of all SVG map images distributed with ERES with short description. The `<NAMING>` placeholder can be one of “English” or “ISO”, depending on the naming convention used.

Africa-<NAMING>.svg	Map of all countries in Africa
Americas-<NAMING>.svg	Map of all countries in North and South America
Asia-<NAMING>.svg	Map of all countries in Asia
Australia-<NAMING>.svg	Map of all countries in Australia and Oceania
Continents-English.svg	Map of the whole world divided by continents. The Area IDs are English names of the continents. This map is not available with ISO naming convention because there are no codes for the continents.

Eurasia-<NAMING>.svg	Map of all countries in Europe and Asia
Europe-<NAMING>.svg	Map of all countries in Europe
NorthAmerica-<NAMING>.svg	Map of all countries in North America
SouthAmerica-<NAMING>.svg	Map of all countries in South America
WorldCompact-<NAMING>.svg	Map of all countries in the world except Antarctica.
WorldCompact-Small-Countries-<NAMING>.svg	Same as WorldCompact-<NAMING>.svg, but very small countries are represented by circles, so they are more visible.
World-<NAMING>.svg	Map of all countries in the world including Antarctica.
World-SmallCountries-<NAMING>.svg	Same as World-<NAMING>.svg, but very small countries are represented by circles, so they are more visible.

Chapter 7. Designing Dashboards

7.1. Introduction to Dashboards

Another way the ERES allows users to publish reports, charts and maps is through dashboards. Dashboards allow users to easily consolidate summarized information in a single page or briefing. Dashboards in ERES require no special skill or administrative support. Simple thin-client interfaces allow all users to easily create customized information presentations.

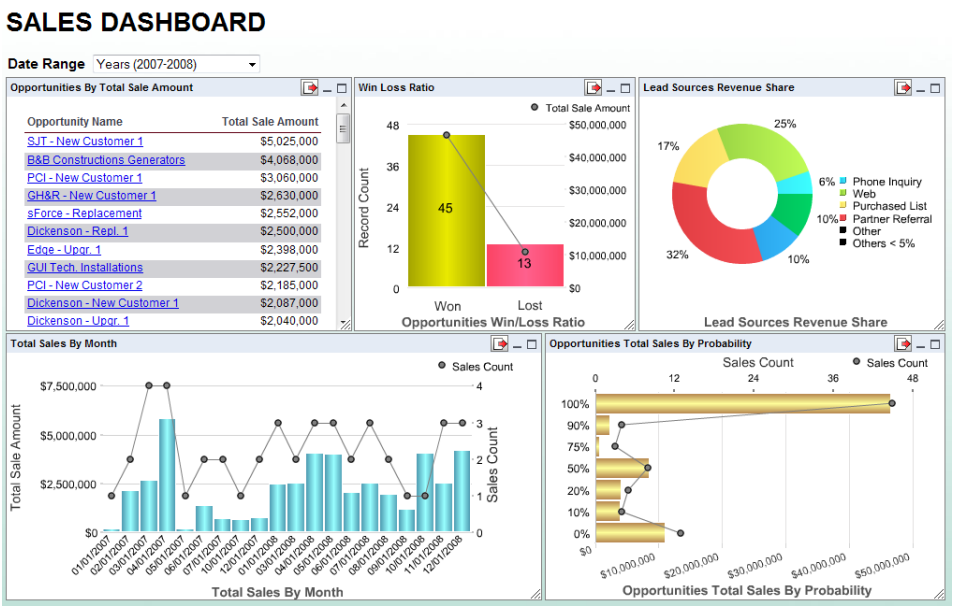
Here are some sample dashboard presentations.



Sales Dashboard



Summary Dashboard



Products Dashboard

The top-level or main presentation consists of a set of charts, reports or maps arranged on the page. You can add shared parameters in a panel, or list the values of parameters so that the dashboard users can apply filters to refresh the applicable charts/reports/maps.

Individual options for each chart or report allow users to apply additional filters as well as export the charts, reports and maps to various formats.

Users can also add layers of drilldown to each chart/report/map in the main presentation. The drilldown can point to a report, chart or map. This allows each piece of summarized information in the presentation to be immediately actionable.

7.2. Create Dashboard

Dashboards in ERES can be created using the Dashboard Builder interface. This easy-to-use thin-client interface allows you to easily create dashboard presentations using charts, reports and maps. You can simply insert dashboard components (charts, reports and maps) that are already deployed in the Organizer to the dashboard. Alternatively, you can create them from scratch directly in the Dashboard Builder. The newly created item will be automatically inserted in the dashboard. One convenient feature is that you can edit a chart, report or map in the Dashboard Builder interface without leaving it.

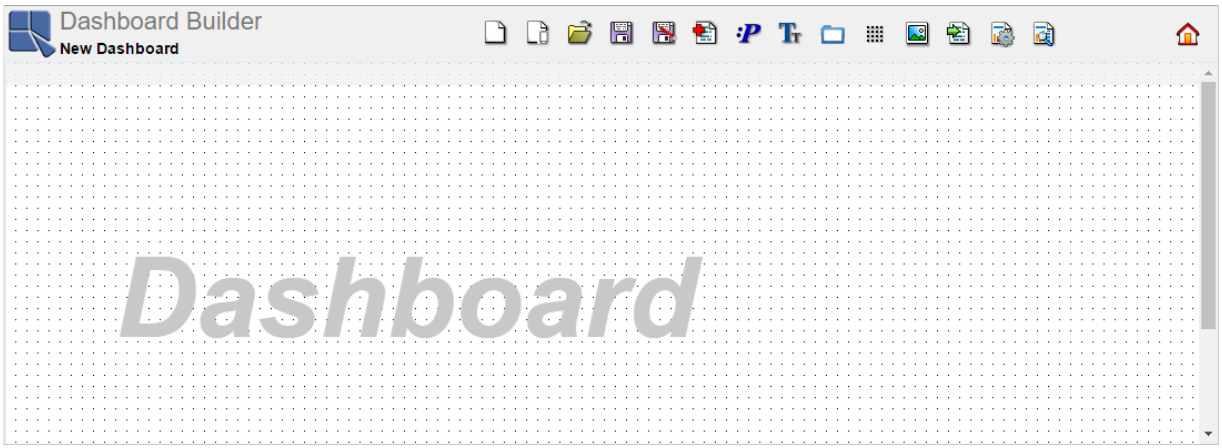
The dashboard layout can be either static or responsive. As you may know, responsive layout is great for displaying content on small devices such as mobile phones. In a static dashboard, the width is dependent on the size and layout of its components. If the dashboard is wider than the screen, all the components stay in place and a horizontal scroll-bar appears. Unlike the static dashboard, the responsive layout, using Responsive Dashboard feature, optimizes utilization of the display without exceeding the display width. Therefore the layout is automatically rearranged. The width of the responsive dashboard can also be set manually by dragging a limiter, a vertical line on the right side of the Dashboard Builder working space (in case you want the dashboard to be narrower than the screen).



Note

It is not possible to display static dashboards as responsive and it is not possible to change them to responsive and vice versa.

The Dashboard Builder can be launched from the ERES Start page. If you've logged in as a user with design privileges, you can follow the Dashboard Builder link to open the interface.



Dashboard Builder Interface

7.2.1. Toolbar

The top part of the Dashboard Builder interface contains a small toolbar that allows you to initiate the following actions:

Start a new static dashboard

Start a new responsive dashboard

Open an existing dashboard

Save the current dashboard

Add report/chart/map to the current dashboard

Add share parameter(s)

Insert a label into the current dashboard

Insert Folder

Hide/Show Grid

Add dashboard background

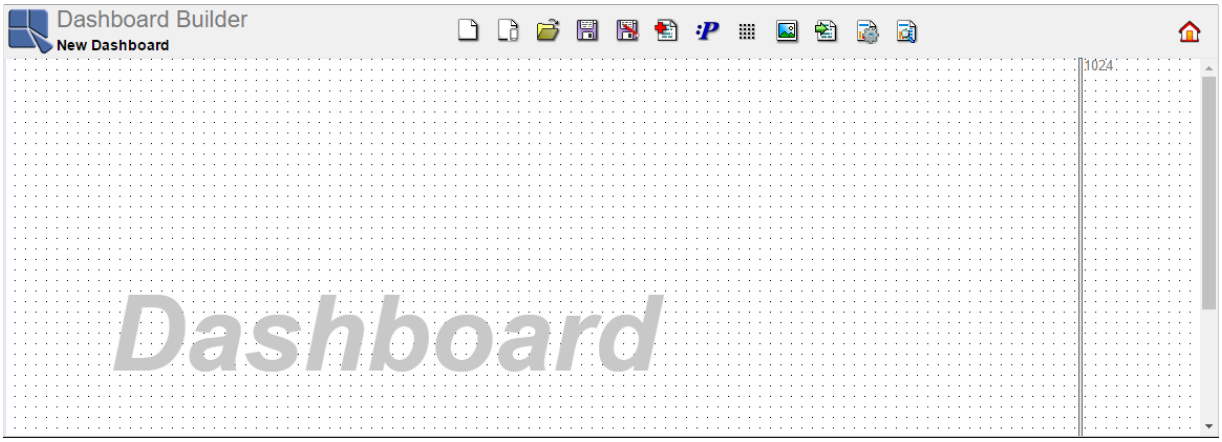
Migration

Dashboard additional options

Preview the current dashboard

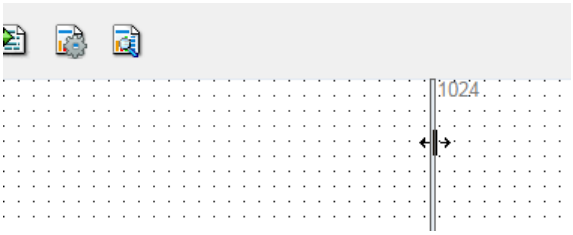
7.2.2. Responsive Dashboard

To create a new responsive dashboard, click the  *New Responsive Dashboard* icon on the toolbar.



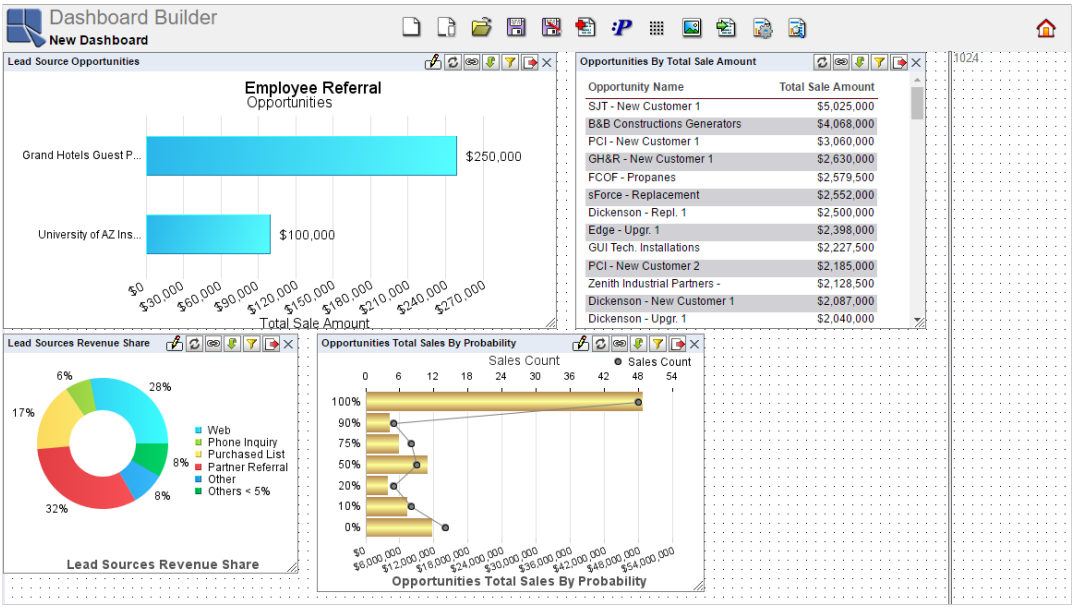
New Responsive Dashboard

You can click and drag the limiter to set a dashboard width.



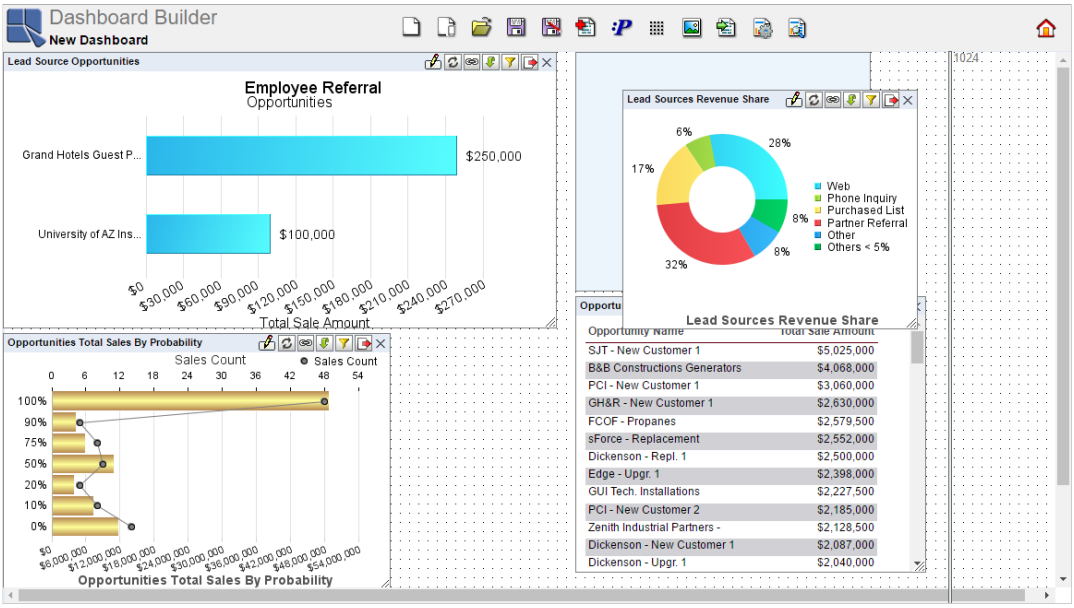
Dashboard Limiter Setting

Now you can add reports/charts/maps and resize them. The limiter (or the screen width) keeps the width of the dashboard. The layout of reports/charts/maps is arranged automatically.



Dashboard With Limiter

To change the layout manually, click on the header of the report/chart/map, hold the mouse button and move it. If there is enough space for the component, a blue rectangle will appear. Release the mouse button. The moved component is placed instead of the blue rectangle and the other components are arranged as best as possible.



Dashboard Layout Changing

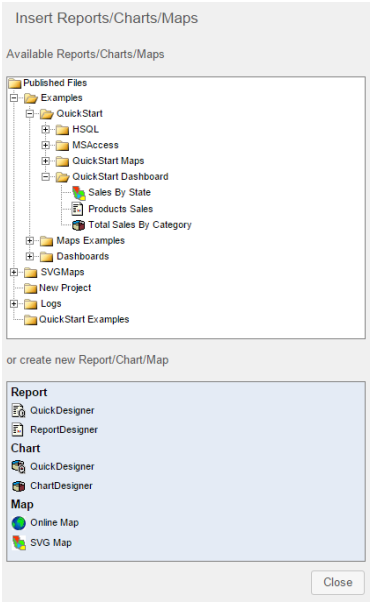


Note

It is not possible to insert labels and folders (see Section 7.2.9 - Folders) into responsive dashboards. Also, you cannot show component header bar (see Section 7.2.7 - Additional Options). Hence dynamic charts are not supported in preview and published dashboards.

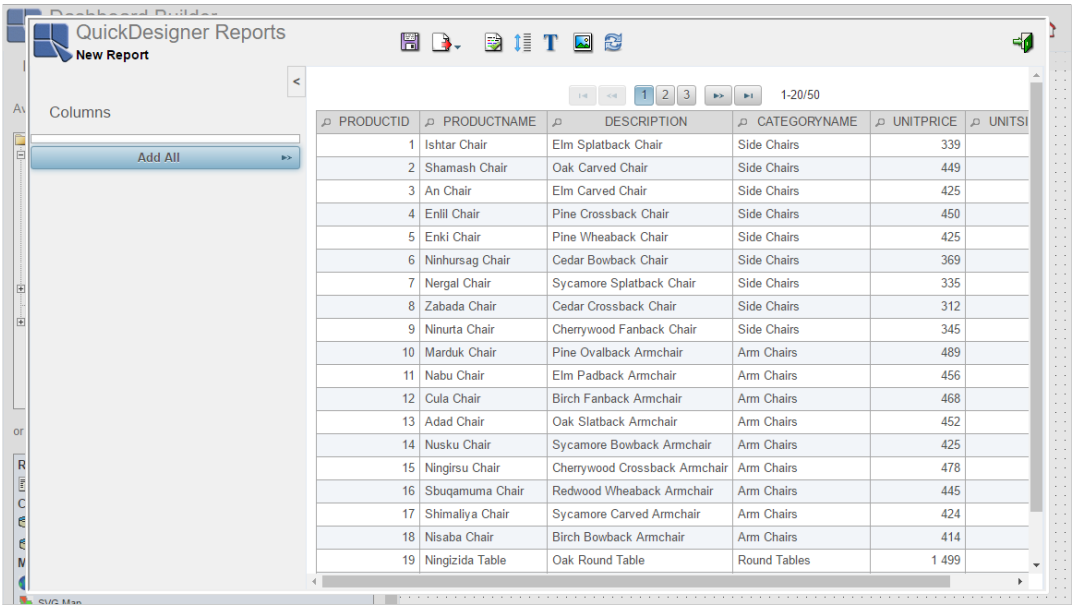
7.2.3. Add Charts, Reports and Maps

To add a chart, report or map to the dashboard, click the  *Add Report/Chart/Map* button on the toolbar. After you do so, the *Insert Reports/Charts/Maps* dialog will appear in the left pane. The dialog contains a tree that mirrors the folder structure in the Organizer. All charts, reports and maps to which you have access to are listed here. In order to insert a chart, report or map into the dashboard, simply click on the chart or report entry in the tree. The chart, report or map will then be added to the dashboard.



Insert Report/Chart/Map Dialog

You can also create a new report/chart/map directly from the Dashboard Builder. To do this, click the relevant link in the lower pane of the *Insert Reports/Charts/Maps* dialog (titled *create new Report/Chart/Map*) and you will be taken to the corresponding designer. Create a new report/chart/map in the designer and save it. Then close the designer or click the exit icon  to return to the Dashboard Builder. Your new component is added to the dashboard. For more information about working in designers, please see Section 5.1 - Report Designer, Section 5.2 - Chart Designer, Section 5.3 - QuickDesigner Reports, Section 5.4 - QuickDesigner Charts, Section 6.2 - Online Maps, Section 6.3 - SVG Maps.



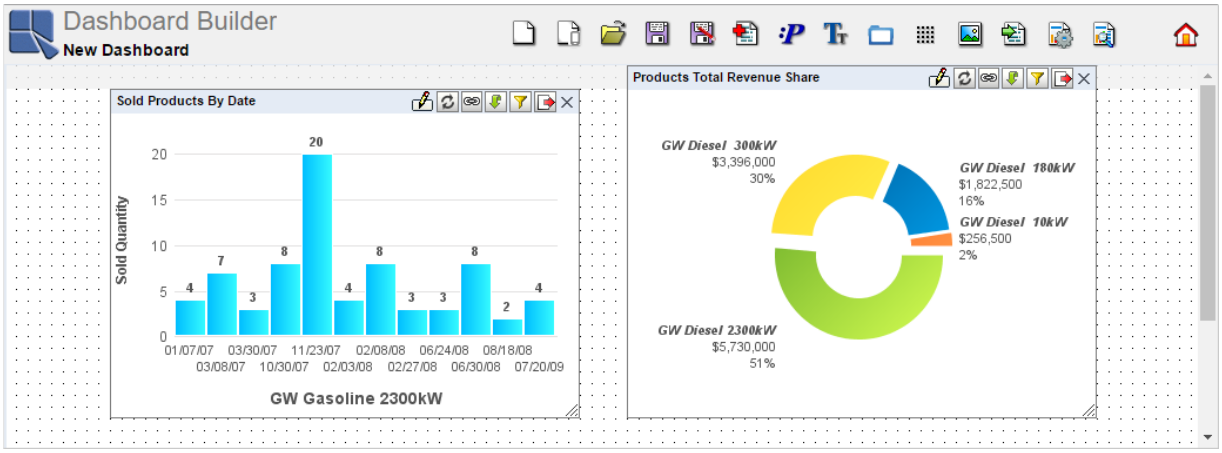
QuickDesigner Reports Opened from Dashboard Builder

7.2.3.1. Move and Resize Charts, Reports and Maps

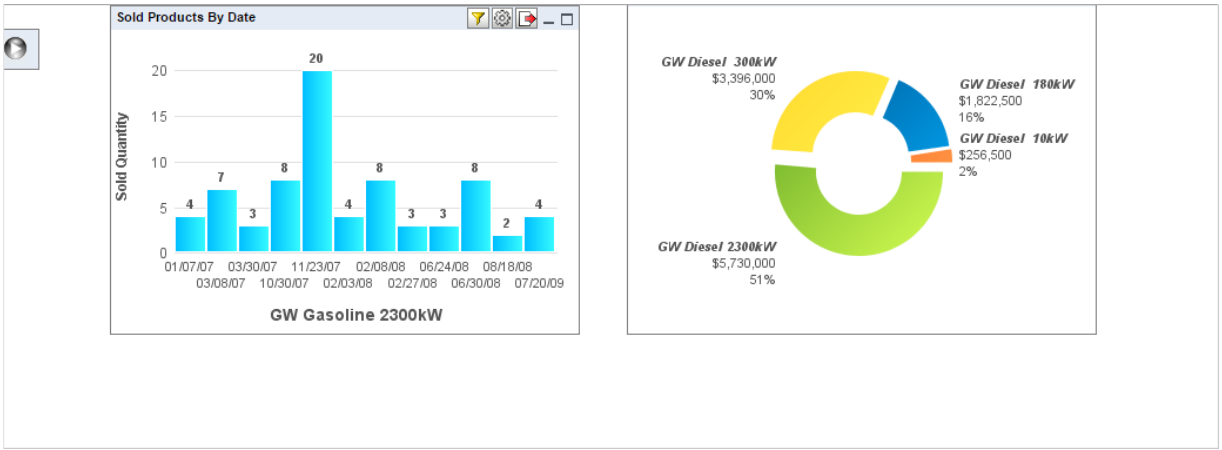
Charts, reports and maps in the dashboard can be moved or resized in free-form. You can resize and move charts, reports and maps using the mouse. To move a chart/report/map, simply click on the header bar of the chart/re-

port/map window and drag it. To resize a chart/report/map, click on the lower right corner of the chart/report/map window and drag on the sizing handle that appears in the lower right corner of the window.

You can see a gray stripe under the Dashboard Builder toolbar. This stripe allows you to place a report/chart/map so that its header bar isn't visible in the preview window. You can see the difference in two images below.



Report/Chart/Map Placement

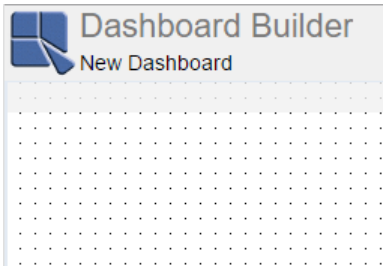


Report/Chart/Map in Preview

If a report, chart or map contains alerts, *Watch alerts* icon  appears in the header bar. To learn more about alerts in dashboards, please visit Section 9.3 - Dashboard alerts.

7.2.3.2. Snap to Grid

You can see a grid of dots in Dashboard Builder that helps you to align charts, reports and maps in the dashboard. The grid is visible by default and it snaps all dashboard elements (reports, charts, maps).



Dashboard Grid

You can hide/show the grid by clicking the *Hide/Show* icon on the toolbar. To position charts/reports/maps completely free, open the *Dashboard Options* dialog by clicking the *Options* icon and disable the option *Snap to Grid* under *Other* section.

7.2.3.3. Chart/Report/Map Toolbar

Each chart/report/map in the dashboard has its own toolbar that allows you to initiate the following actions:

Edit Template - This button opens appropriate designer that allows you to edit a template.

Refresh Template - This button refreshes a template.

Add/Modify Link - This button allows you to use data from one chart/report/map as a parameter in another. (see Section 7.2.10 - Template Linkage)

Add/Modify Drilldown - This button allows you to add/modify drilldowns. (see Section 7.2.8 - Drilldown)

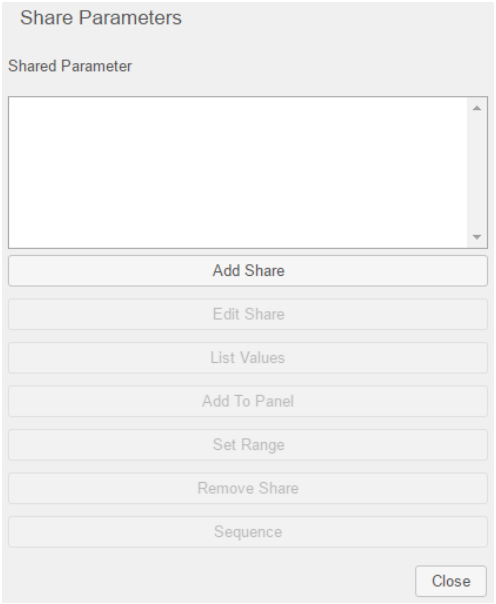
Filter Data - This button allows you to specify filters/parameters for a template. (see Section 8.9.2 - Preview Options)

Export - This button allows you to export chart/report/map. (see Section 8.9.2 - Preview Options)

7.2.4. Shared Parameters

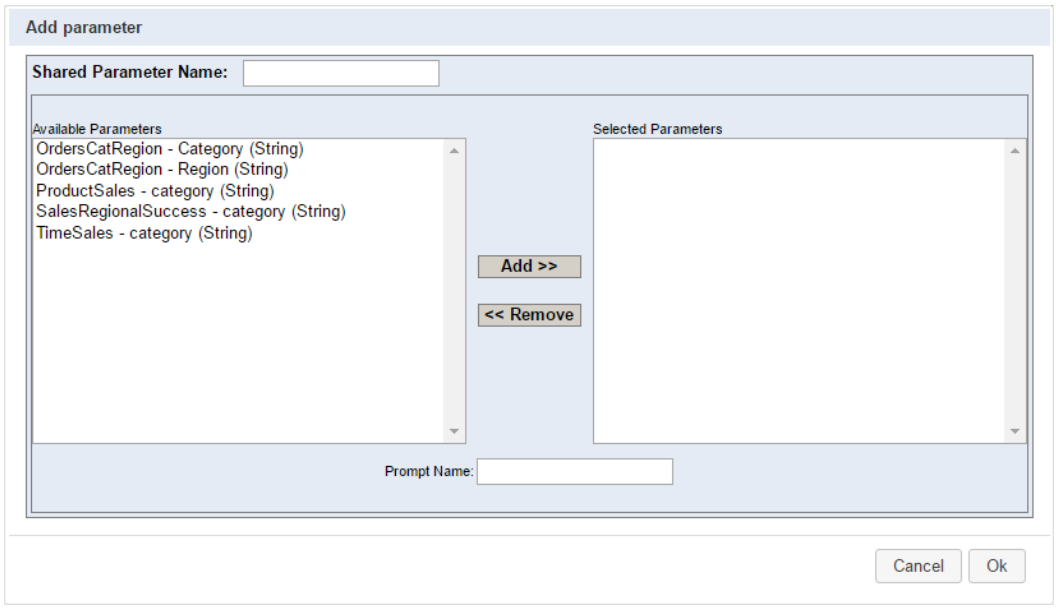
Shared parameters allows you to group common parameters from all charts/reports/maps into a single parameter. You can apply a common filter to some or all of the dashboard items at the same time. All of the dashboard's shared parameters are listed in the *Shared Parameter* list in the *Share Parameters* dialog that can be opened by clicking the *Share Parameters* icon on the Dashboard Builder toolbar.

For example, create a dashboard by adding all files from the `Examples/Dashboards/CategorySales-Dashboard` node (OrdersCatRegion, ProductSales, SalesRegionalSuccess, TimeSales). Then click the *Share Parameters* icon. You will now see *Share Parameters* dialog in the left pane.



Share Parameters Dialog

To create a shared parameter, click the *Add Share* button in the Share Parameters dialog. This will open a dialog that allows you to set options for the share.



Shared Parameter Dialog

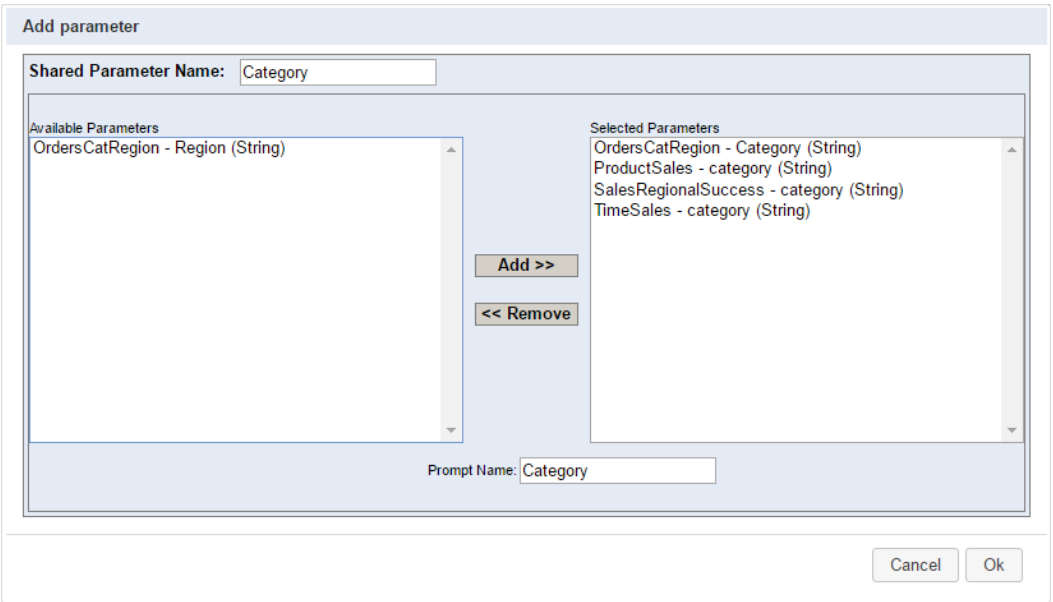
The left side of the dialog shows all parameters for all the selected charts/reports/maps in the dashboard. You can select which parameters you want to add to the share by selecting them in the left side window and clicking the *Add* button. Note that all parameters in the share must have the same data type.

The parameter mapping options (whether the parameter is mapped to a database column) is based on the first parameter selected for a share. For more information about parameter mapping, see Section 4.1.2.2.2.2 - Initializing Query Parameters.

You can enter a prompt name for each shared parameter. Once you finish entering all informations for the shared parameter, be sure to specify a name for it in the field at the top and then click the *Ok* button to save the changes. The parameter will then be added to the *Shared Parameters* list in the share parameters dialog.

You can edit or remove a shared parameter by selecting it in the list and clicking the *Edit Share* or *Remove Share* buttons.

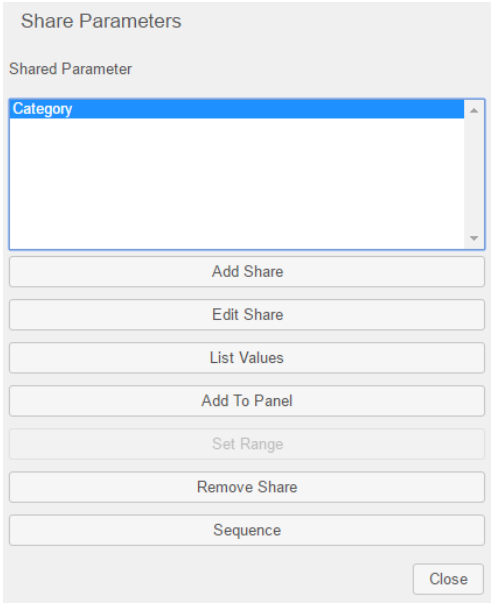
For our example, write *Category* as the *Shared Parameter Name* and the *Prompt Name* and select all available parameters *Category*. Click the *Ok* button to create a shared parameter.



Shared Parameter Dialog

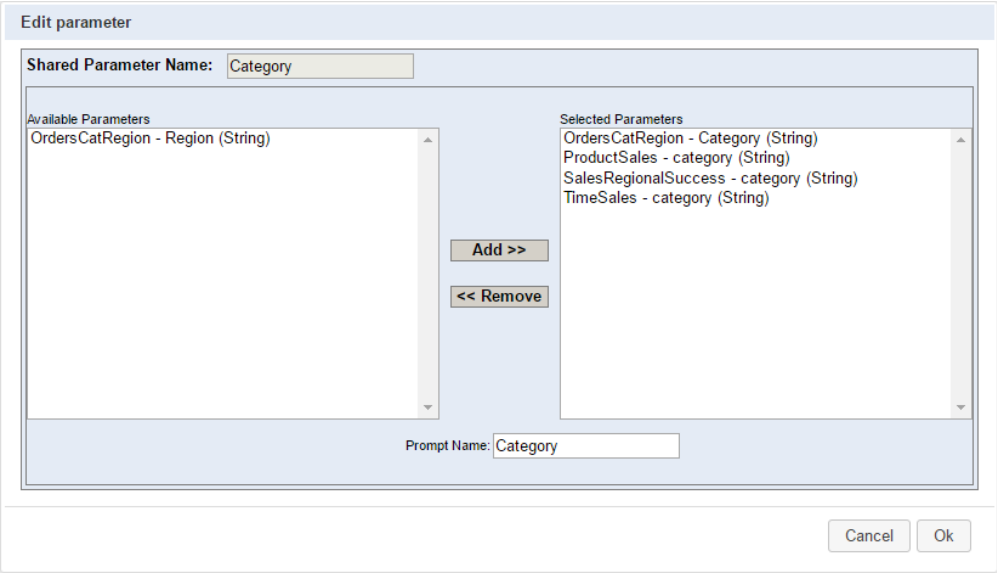
7.2.4.1. Arrange Shared Parameters

In addition to adding/editing/removing shared parameters, you can also list values of a shared parameter into a set of buttons in the dashboard or arrange a shared parameter to be displayed in the panel that allows you to enter parameter value right from the dashboard. Please note that in order to list the values of a parameter, the shared parameter can be either single or multi-valued and must have predefined selectable values. To add the shared parameters to the dashboard, simply click the *List Values* or *Add to Panel* button in the *Share Parameters* dialog.



Share Parameters Dialog

Edit Share Allows you to edit a shared parameter. To do this, select a shared parameter and click the *Edit Share* button. This will open *Edit parameter* dialog.



Edit Shared Parameter

List Values

If the parameter has a small number of values, you can choose to list the values as buttons or images. For example, let's say your dashboard shows sales data for your products and you have a relatively small number of products. You can select this option to make a list which makes it easier for the users to pick a value for the filter. To list the values for a shared parameter, first select it from the shared parameters list and click the *List Values* button. This will open the parameter value-list panel that contains the parameter values in buttons.



Parameter Value-list Panel

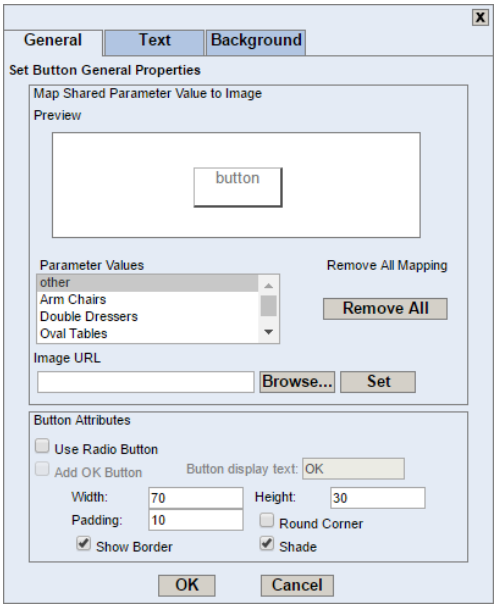
There are four buttons available in the panel header bar.



Modify Panel Attributes: This button will open a new dialog that allows you to modify panel attributes such as background color and panel border. For more information about the dialog, please see Section 7.2.5 - Insert Labels.



Modify Button Attributes: This button will open the following dialog that allows you to modify button attributes.

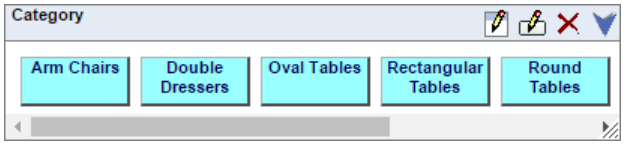


Set Button Properties Dialog

The first option in this tab allows you to map shared parameter value to an image. To map a shared parameter value, first select it either from the left panel, specify the image URL or browse for it using the *Browse* button and click the *Set* button. In order to remove image URL mapping, select appropriate parameter value, clear the image URL text box and press the *Set* button or use the *Remove All* button to remove all mapping. The other options in the dialog allows you to specify button width, height and padding. You can also choose whether to show the buttons border or whether the buttons should be shaded or not. Another option allows you to display the boxes with a radio button. This option is not available for multi-valued parameters, which are automatically given checkboxes.

The other tabs of the Set Button Properties dialog allows you to set the button text and background. The tabs are basically the same as for labels and they are described in the following Section 7.2.5 - Insert Labels.

The next screenshot shows our example with a modified view of buttons.



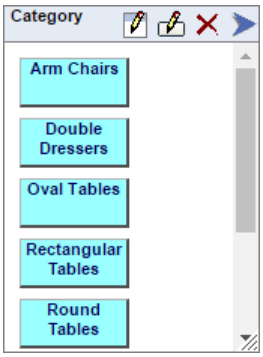
Modified Buttons



Delete: This button will delete the parameter value-list panel.



Set Buttons Alignment: These buttons allows you to set vertical or horizontal alignment of buttons. By default, buttons are aligned horizontally and you will only see the  *Vertical Align* button. When you click on it, the alignment will be changed to vertical and the  *Horizontal Align* button will appear.



Vertical Buttons Alignment

Add To Panel To arrange the layout of shared parameters by adding them to a panel, first select it from the shared parameters list and click the *Add to Panel* button. This will open a parameter panel with the parameter's list box added to the panel.



Parameter Panel

In the Dashboard Builder, you can move/resize the panel or move the small parameter box inside the panel. If you close the panel, you can still use the shared parameter with the preview toolbar (see Section 8.9.1 - Preview Toolbar). If you change the parameter value, the dashboard will be automatically refreshed.

There are three buttons available in the panel header bar:

Disable Auto Rearrangement: This button will disable/enable auto rearrangement.

Modify Detach Panel Attributes: This button will open a new dialog that allows you to modify the parameter panel attributes such as adding an *OK* button, changing panel layout, text label properties, background color and border for the panel. For more information about the dialog, please see Section 7.2.5 - Insert Labels.

Delete: This button will delete the parameter panel.

Set Range This option allows you to define a range across two shared parameters. For more information, please see Section 7.2.4.2 - Parameter Range.

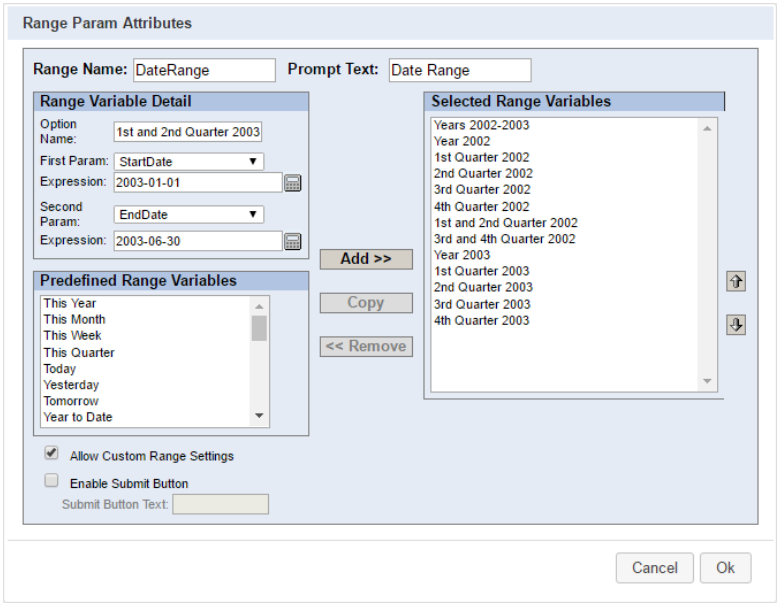
Remove Share This option allows you to remove shared parameter. To do this, select the shared parameter and click the *Remove Share* button.

Sequence This option allows you to set an order of shared parameters. For more information, please see Section 7.2.4.3 - Cascading Parameters.

7.2.4.2. Parameter Range

This function allows you to define a range across two shared parameters. To use this function, select two related shared parameters (such as start and end date, or integer objects representing a range) and then click the *Set Range* button. This will bring you to the *Range Param Attributes* dialog. To create a user defined range variable, simply fill out the *Range Variables Detail* dialog box. To see an example, please see the QuickStart Guide (Section Q.10.1.1 - Create a Dashboard). The next screenshot shows other setting of parameter range for the example created in QuickStart Guide.

In addition to creating your own range variables, you can also create range variables using *Predefined Range Variables* as described below.



Parameter Range Dialog

The dialog has these following options:

Range Name: This field allows you to enter a title for the parameter range window.

Prompt Text: This text is displayed next to the selected combo parameter dropdown menu.

Range Variable Detail: This window has several sub-fields:

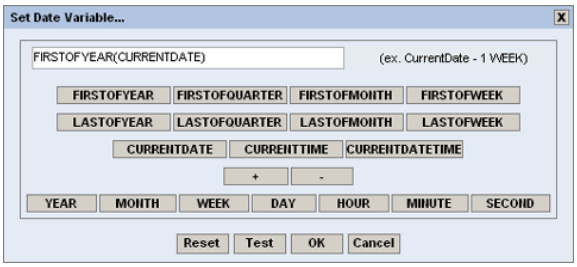
Option Name: This is the display name in the *Selected Range Variables* list.

First Param: This is the first parameter in the combined parameter, usually the starting value.

Second Param: This is the last parameter in the combined parameter, usually the ending value.

Expression: There are two expression fields, one for each parameter.

For date parameters, you can click on the *Date calculator* button  next to the expression field.



In this screen, you can enter one of the three keywords: `CurrentDate`, `CurrentTime`, and `Current-`

`DateTime`, or one of the function names, such as `FIRSTOFTHISYEAR`. Details about functions are listed in the list of functions.

You can add or subtract time units from the current date/time, allowing you to have a dynamic date range. For example, a report may have the following default values:

StartDate: CurrentDate - 1 WEEK

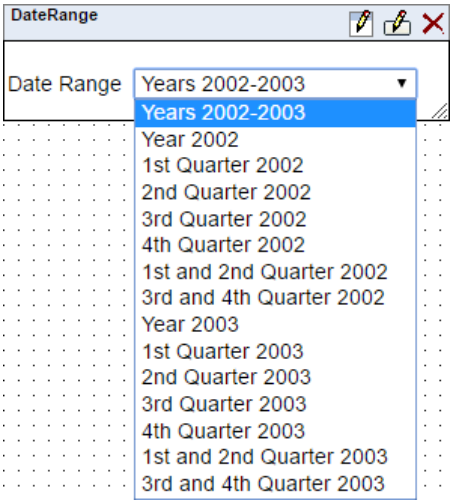
EndDate: CurrentDate

This would indicate that every time the report is run, the default prompt should be one week ago of the current date. Other supported options are `YEAR`, `MONTH`, `DAY`, `HOUR`, `MINUTE`, and `SECOND`. This feature only supports single addition or subtraction and it does not support multi-value parameters.

Selected Range Variables:

This window displays all range variables that have been defined. Clicking on any range variable in this list allows you to modify it, delete it or copy it. Once the variable is selected, you can see its detail in the *Range Variable Detail* window where you can modify it. You can move the selected range variable up or down in the list by clicking the  arrows right next to the window.

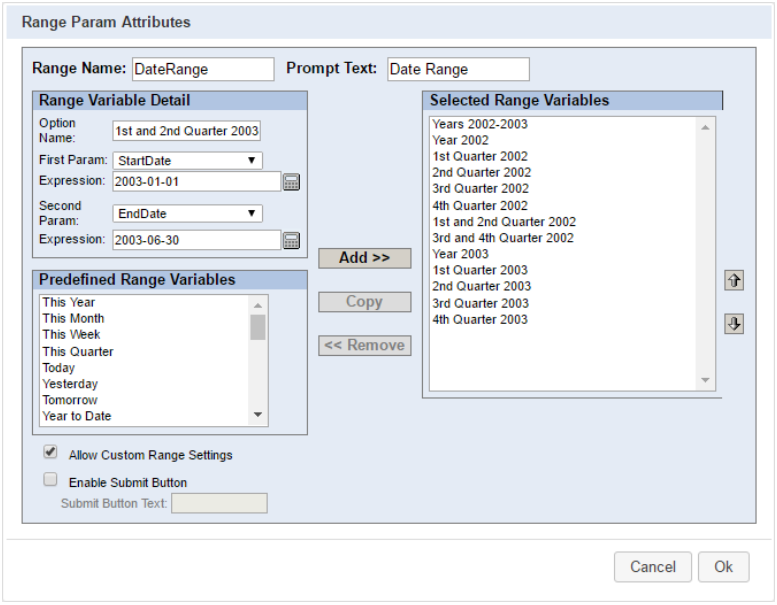
The next screenshot shows a dropdown menu in the dashboard that contains defined range variables (range variables from the *Selected Range Variables* window).



Range Variables in Dashboard

Add:

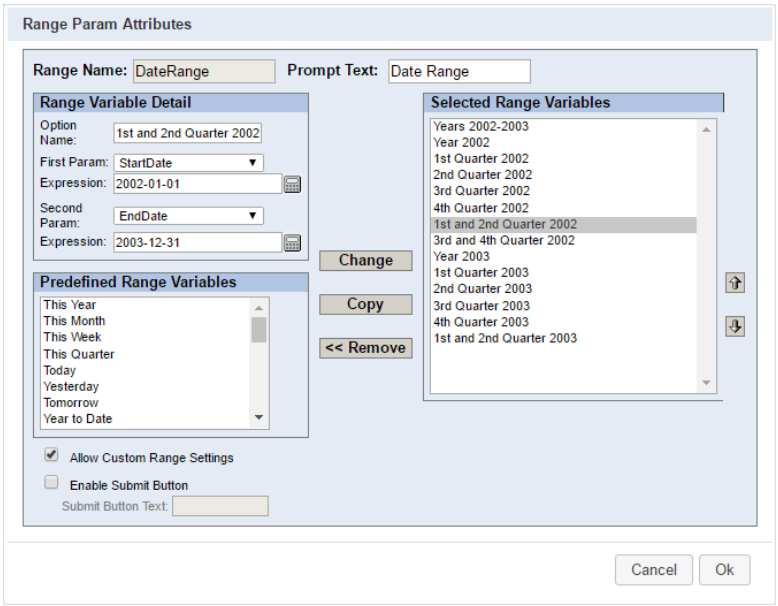
This button has two modes: *Change* or *Add*. When the text of the button reads *Add*, the parameter information in the *Range Variable Detail* window will be added into the *Selected Range Variables* list.



Range Variable Adding

Change:

This button has two modes: *Change* or *Add*. Clicking on an item in the *Selected Range Variables* list will change the text to *Change*. When in this mode, clicking the button will save any changes made in the *Range Variable Detail* window to the currently selected range variable.

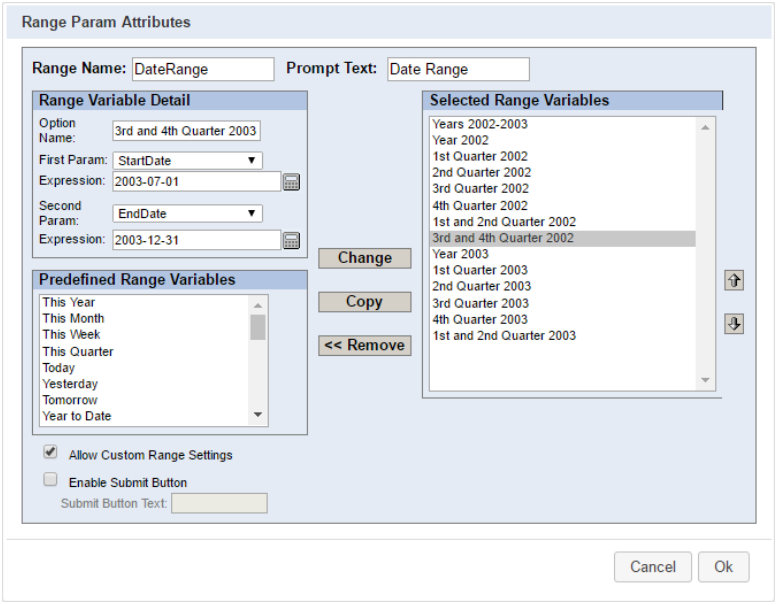


Range Variable Change

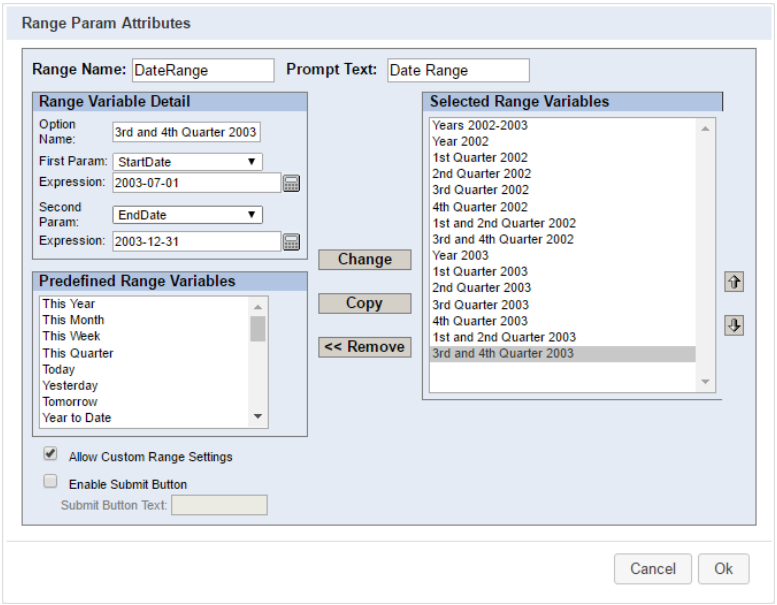
Copy:

This button will copy the selected item and add it to the *Selected Range Variables* list under a new name. The new name is created automatically by adding *_0,_1,_2,etc.* or you can change it in the *Range Variable Detail* window. To copy a variable under your name, select it, change a name in the *Range Variable Detail* window and then click the *Copy* button.

Copy can help you to create range variables more quickly. For example, you can select a range variable, change only the year and then click the *Copy* button.



Range Variable before Copy



Range Variable Copied

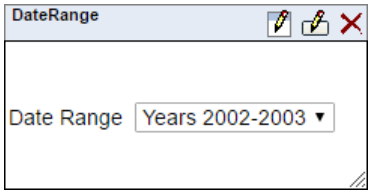
Remove: This button will remove the selected item from the *Selected Range Variables* list.

Predefined Range Variables: This window contains predefined date range variables for your convenience. It contains commonly used date range expressions such as `This Year, This Month, This Week, This Quarter,` and many more. To use a predefined variable, simply select the needed predefined range variable and click the *Add* button. You can add more than one expression and all of them will be listed under the *Selected Range Variables* window.

Allow Custom Range Settings: This option allows you to specify a parameter range by using calendar and parameter input boxes in the parameter range panel. The following screen-shots show the difference between with and without custom range.



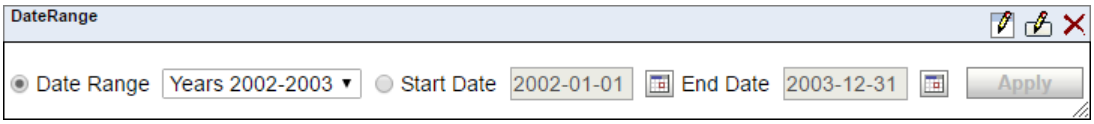
Enabled Custom Range Settings



Disabled Custom Range Settings

Enable Submit Button — If you allow *Custom Range Settings*, you can add *Submit Button*. This button allows you to apply custom range when both parameters are set. Without this button, each parameter is applied immediately when it is set. You can also determine a text of this button.

Once you finish setting up the parameter range in the *Range Param Attributes* dialog, click the *Ok* button to close the dialog. A rectangle will follow your mouse pointer around the dashboard design window. Position the rectangle and click. A parameter range panel will appear.



Parameter Range Panel with Submit Button

You can select range variables from the dropdown menu where are all variables from the *Selected Range Variables* window, or use customer text boxes (text boxes are available only when *Allow Custom Range Settings* option was checked). The text boxes represent the first and second parameters in the combined parameter. To use customer text boxes, check the radio button next to them and insert parameter values. Note that you can use the calendar to specify the date parameters. To open the calendar, simply click the *Calendar* button  next to the textbox.

There are three buttons available in the panel header bar:

 **Modify Panel Attributes:** This button will open a new dialog that allows you to modify panel attributes such as text label properties, background color and border for the panel. For more information about the dialog, please see Section 7.2.5 - Insert Labels.

 **Modify Range Param Attributes:** This button will open the *Range Param Attributes* dialog that allows you to edit range variables.

 **Delete:** This button will delete the range parameter panel.



Note

If you delete this panel, all range variables that were defined in the *Range Param Attributes* dialog will be lost.

7.2.4.3. Cascading Parameters

By default, you are prompted to enter all report parameters at once in the prompt dialog. However, this configuration may not be the best solution in case some parameters are mapped to database columns with a significant number

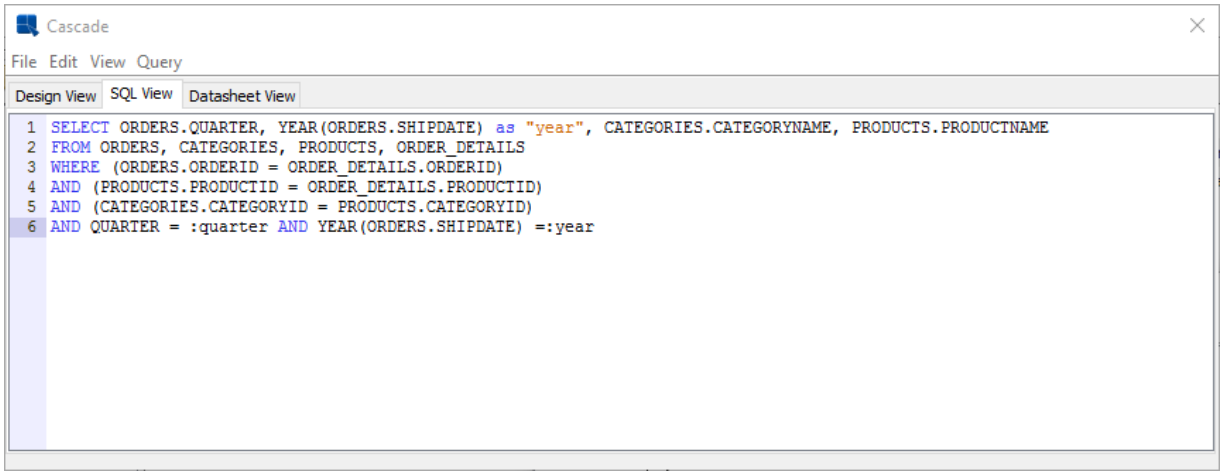
of distinct values. It can be difficult to go through a very large list and you may select values that don't return any data, depending on the parameter combination.

To assist with these problems, ERES provides a feature that allows you to configure the order in which the parameters should be entered. With this feature enabled, you can enter parameters in the dialog in a pre-defined order. As such each selection will be applied as a filter to the next parameter prompt(s). Using cascading parameters can limit the number of distinct values presented and can prevent selecting invalid parameter combinations.

First, create a query named e.g. `Cascade` with parameter sequence in Woodview HSQL database that is in ERES.

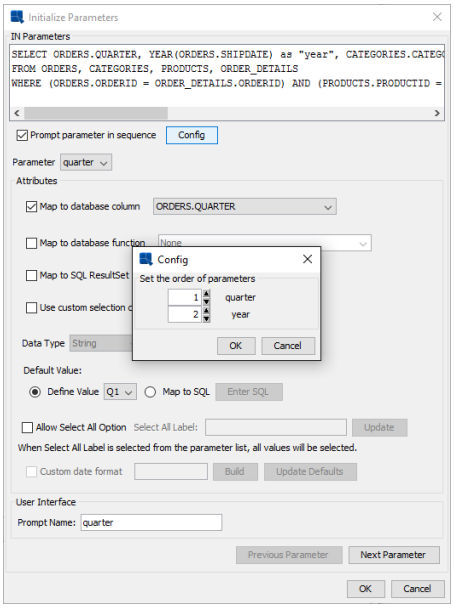
To do this, open Organizer, click the  *Manage Data Sources* icon to open the *Data Registry Manager* dialog, select a data registry and click the *Edit* button. Select `Queries` under `Databases/Woodview (HSQL)` node and click the *ADD* button to create a new query. Type *Cascade* as a name and for example, use the following query:

```
SELECT ORDERS.QUARTER, YEAR(ORDERS.SHIPDATE) as "year",
 CATEGORIES.CATEGORYNAME, PRODUCTS.PRODUCTNAME
FROM ORDERS, CATEGORIES, PRODUCTS, ORDER_DETAILS
WHERE (ORDERS.ORDERID = ORDER_DETAILS.ORDERID)
AND (PRODUCTS.PRODUCTID = ORDER_DETAILS.PRODUCTID)
AND (CATEGORIES.CATEGORYID = PRODUCTS.CATEGORYID)
AND QUARTER = :quarter AND YEAR(ORDERS.SHIPDATE) =:year
```



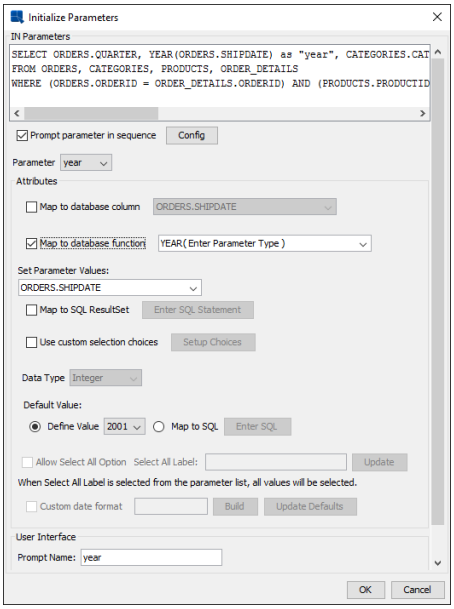
Cascading Parameters Query

Now initialize query parameters. Check *Prompt parameter in sequence* and click the *Config* button. Set the parameter sequence according to the image below, i.e. the *quarter* parameter will be in the first level and *year* in the second level. Then map the first parameter *quarter* to database column *ORDERS.QUARTER*.



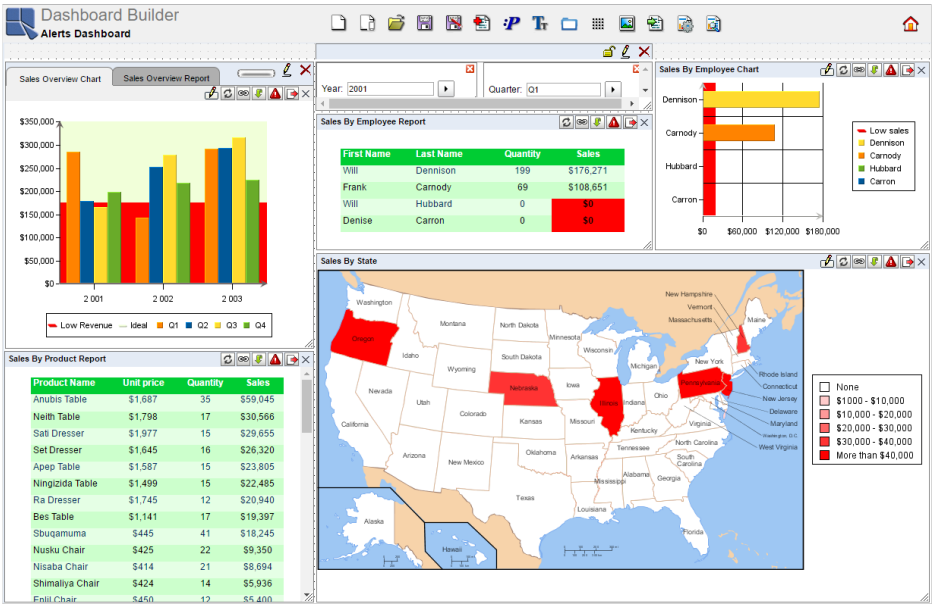
Initialize First Parameter

Map the second parameter *year* to database function *YEAR* and *Set Parameter Values* to *ORDERS.SHIPDATE*. Then click *OK* to save the query.



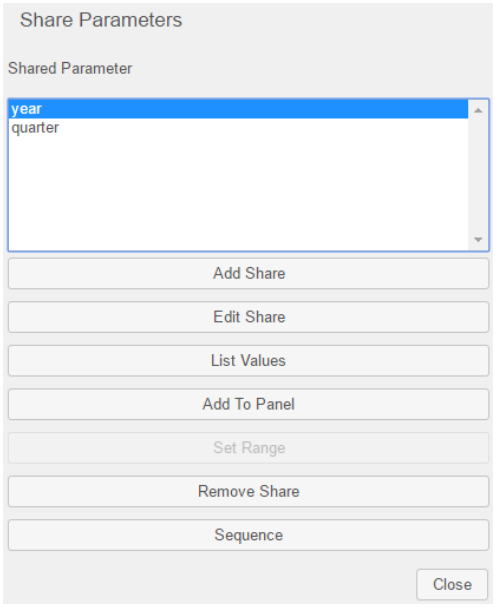
Initialize Second Parameter

Go to Dashboard Builder and open the *Alerts* dashboard (`Published Files/Examples/Dashboards/Alerts Dashboard`).



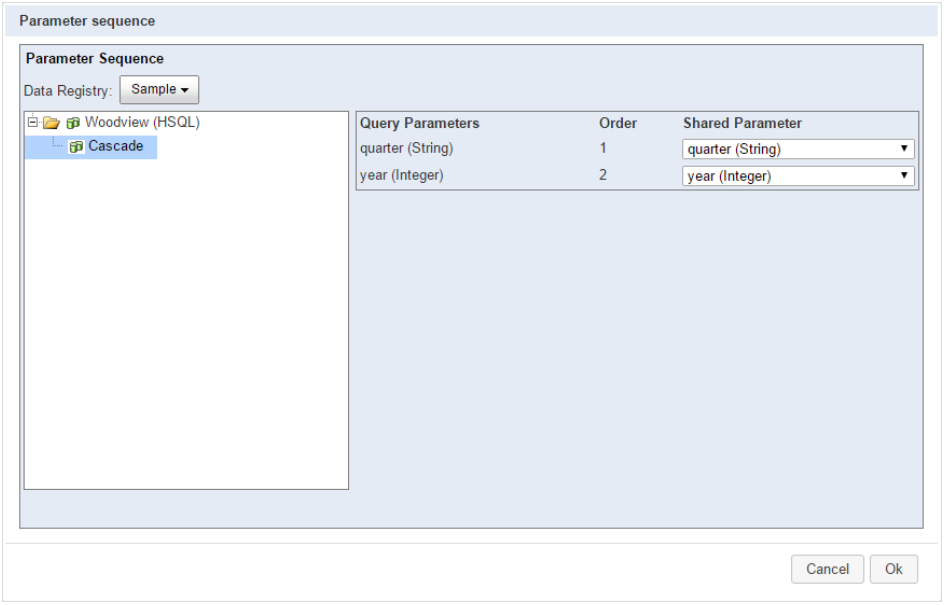
Alerts Dashboard

Then click the  *Share Parameters* icon on the Dashboard Builder toolbar to show the list of shared parameters in the dashboard (Notice that there are two shared parameters *quarter* and *year* already added in a dashboard. You can remove them to make room for the cascading parameters). From the dialog, select one of parameters to activate options and click the *Sequence* button at the bottom.



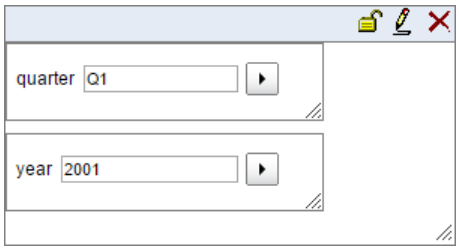
Shared Parameters Setting

The *Parameter Sequence* dialog will then appear. From the query tree on the left, select *Cascade* query you have created in previous steps. It should automatically map query parameters to shared parameters in the dashboard.



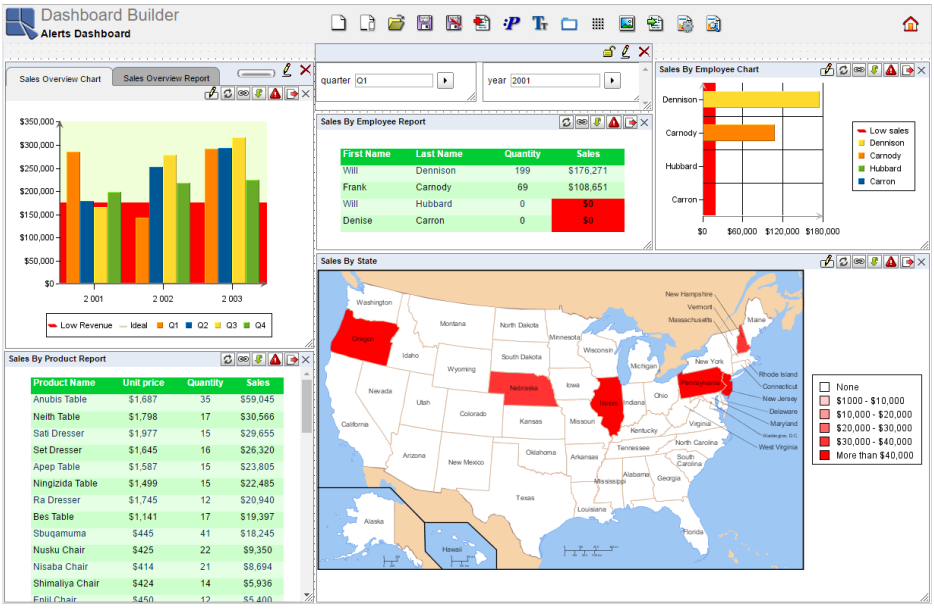
Parameter Sequence Dialog

Click the *OK* button and place the parameter panel into a dashboard.



Parameters Panel

The image below shows inserted cascade parameters panel in the dashboard. Now you can specify a parameter value for the first and second level. Once you specify the last parameter in the sequence, the dashboard templates should refresh according to the selected values.



Dashboard with Cascade Parameters Panel

7.2.5. Insert Labels

You can insert a label into a dashboard by clicking the *Insert Label* button on the dashboard builder toolbar (Please note that labels are not available for responsive dashboards). After you click the button, a small rectangle will follow your mouse pointer around the builder interface. Position the rectangle where you want to insert the label and click. The label will be inserted into the dashboard. You can edit the label text by double-clicking on the label body.



Inserting a Label



Editing a Label

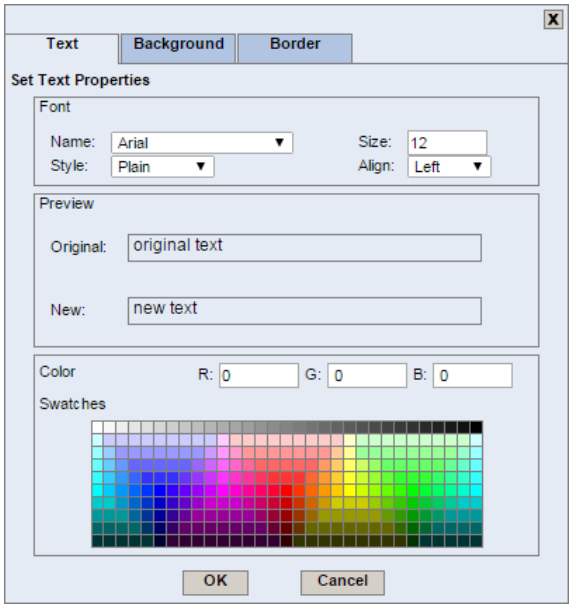
Labels can be moved or resized in the same way as charts, reports and maps. To move a label, simply click on the header bar of the label window and drag it. To resize a label, click on the lower right corner of the label window and drag on the sizing handle that appears around the lower right edge of the window.

There are three additional buttons available for a label in the label header bar.

Move to Back: This button will move the label to back.

Edit: This button will open a new dialog that allows you to specify text label properties, background color and border for the label.

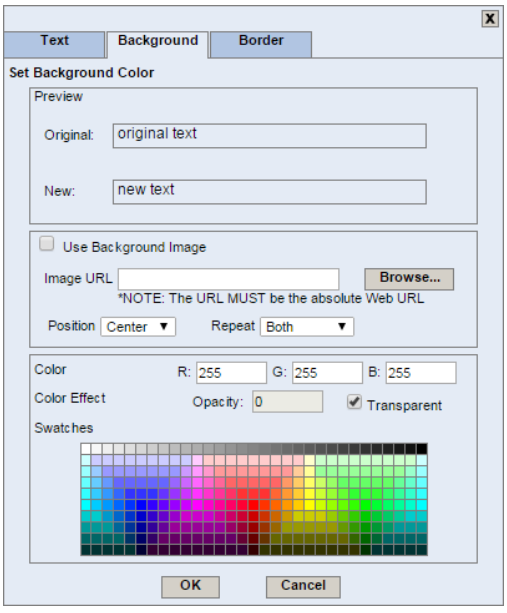
Text Label Properties: The first tab of the edit window allows you to adjust text label properties. You can adjust the font type, font style, font size and alignment of the label as well as changing font color using swatches or RGB values. In the preview section you can still see the original and new text.



Label Text Properties

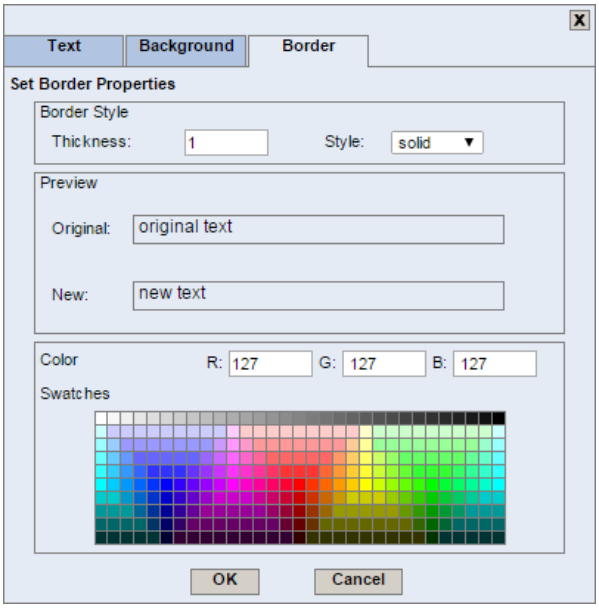
Label Background Color:

The second tab of the edit window allows you to set label background color or to use background image. The background color can be adjusted using swatches or RGB values. You can also specify the label opacity or choose whether or not to set the label background transparent. If you want to use a label background image, simply check the *Use Background Image* checkbox and enter the image URL or use the *Browse* button to browse to the appropriate background image. Please note that the URL must be absolute Web URL. In addition to adding background images, you can also specify the background image position (available positions are: center, left, right, top, bottom and fixed) or choose whether to set the dashboard background image to be repeated (available options are: both, horizontal, vertical and none).



Label Background Color

Border Properties: The third tab of the edit window allows you to set label border style and border color. The label border can be adjusted using swatches or RGB values. If you want to use border style, specify the border thickness and select appropriate border style (available options are: solid, dotted, dashed, double and groove). In the preview section you can still see the original and new text and border.



Label Border Properties

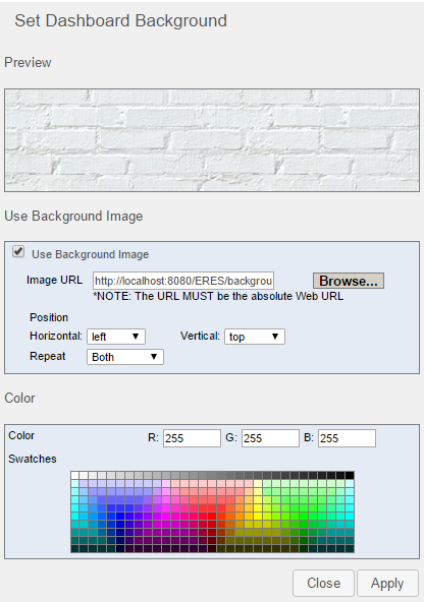
Delete: This button will delete the label.

Please note that the labels can also contain any HTML tags, which may override external settings from the dialog.

For example, let's say you have a label with the font type set to `Arial`. However, once you enter HTML tag `<font face="Garamond">Label content</font>` into the label, `Arial` font will be overriden with `Garamond` font for the `Label content` text.

7.2.6. Add Background

You can also add dashboard background by clicking the *Add Dashboard Background* button on the dashboard builder toolbar. This will open a dialog that allows you to set the dashboard background.



Dashboard Background Dialog

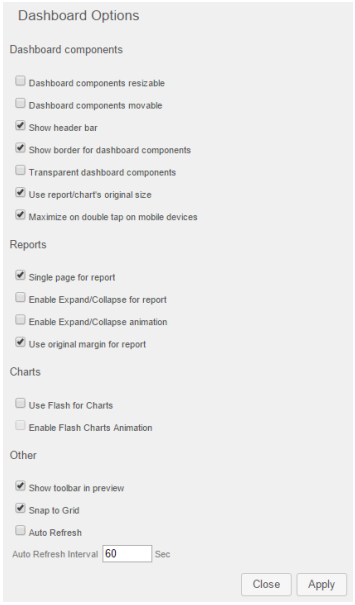
The first option in the dialog allows you to use a background image for a dashboard. If you want to use a background image for the dashboard, simply check the *Use Background Image* checkbox and enter the image URL or use the *Browse* button to browse to the appropriate background image. Please note the URL must be the absolute Web URL. In addition to adding background images, you can also specify the background image position (available positions are: center, left, right, top and bottom) or choose whether to set the dashboard background image to be repeated (available options are: both, horizontal, vertical and none). In the preview you can still see updated dashboard background.

The second option in the dialog allows you to set dashboard background color. The background color can be adjusted using swatches or RGB values.

Once you finish setting up the dashboard background options, click the *Apply* button to apply the changes. You can close this dialog by clicking the *Close* button or by clicking the *X* sign in the upper right corner.

7.2.7. Additional Options

You can also specify additional dashboard options by clicking the *Options* button  on the dashboard builder toolbar. This will open a dialog that allows you to set the dashboard components resizable/movable and other options.



Dashboard Options Dialog

Dashboard components resizable: This option allows you to resize the dashboard components. It doesn't affect the Dashboard Builder as it works in dashboard viewer (i.e. when you open a dashboard from MenuPage etc.) and dashboard preview only.

Dashboard components movable: This option allows you to move the dashboard components. Like the previous option, this one also affects deployed dashboards only.

Show header bar: Shows/hides header bars for templates in the dashboard. Affects deployed dashboards only. In Dashboard Builder, headers are always visible.

Show border for dashboard components: Shows/hides borders for templates in the dashboard. Affects deployed dashboards only. In Dashboard Builder, borders are always visible.

Transparent dashboard components: Enable transparent background of dashboard components (e.g. transparent background of chart/report/map title, legend).

Use report/chart's original size: This option only affects charts/reports at the moment you add them to the dashboard. If this option is enabled, charts/reports will be added to the dashboard using their original size settings (i.e. canvas size for charts and page size for reports). If this option is disabled, all reports/charts will be resized to standard size and their original size settings will be ignored. This option doesn't affect charts/reports after they have been added to the dashboard so you can resize them as you like.

Maximize on double tap on mobile devices: Enable this option to maximize a dashboard template by double tapping on mobile devices.

Single page for report: This option allows you to set whether reports will be displayed in single page or multi-page format

Enable Expand/Collapse for report? Enables Expand and Collapse function for reports.

Enable Expand/Collapse animation: Enable animation for Expand and Collapse function.

Use original margin for report: This option will set reports to be displayed with their original margins as defined in the Page Setup in the Report Designer. If this option is disabled, report margins are set to 0.2 inches.

Use Flash for Charts:	Use Flash for charts instead of the default PNG format.
Enable Flash Charts Animation:	Enable/disable animation for Flash charts.
Show toolbar in preview:	Hides/displays toolbar in deployed dashboards. For more information about the preview toolbar, please see Section 8.9.1 - Preview Toolbar.
Snap to Grid:	Snaps components to grid when resizing/moving.
Auto Refresh:	Enables auto refresh for deployed dashboards (this option doesn't affect the Dashboard Builder - it affects dashboard viewer and dashboard preview only). In already deployed dashboards, auto-refresh can be also enabled/modified from the preview toolbar.

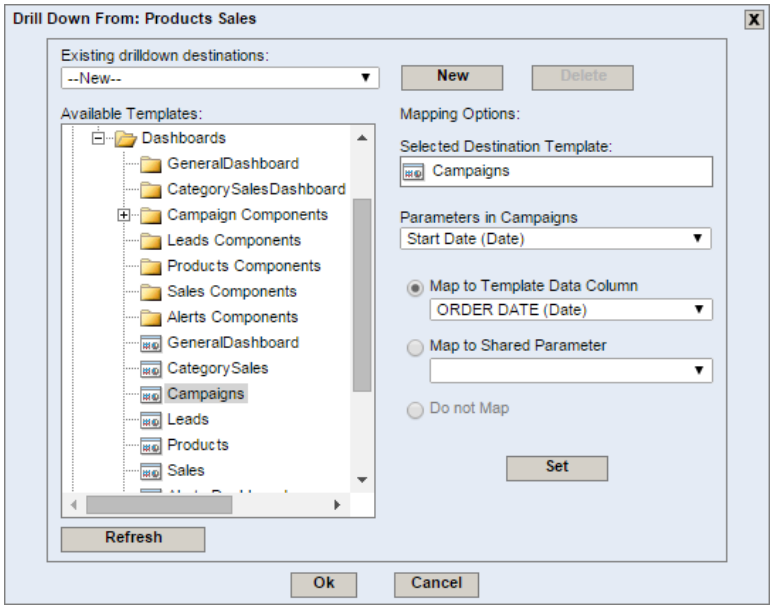
Once you finish setting up the dashboard options, click the *Apply* button to apply the settings. You can close this dialog by clicking the *Close* button or by clicking the *X* sign in the upper right corner.

7.2.8. Drilldown

You can add a drilldown to any chart/report/map in the dashboard. The drilldown feature allows you to map the chart data points, report columns, and/or shared parameters to the lower-level chart or report parameters. This creates a relationship where an end-user can either click on a data point in a dashboard chart, or row in a table and be taken to a lower-level report or chart that shows detailed information for the selected point.

You can specify more than one drilldown relationship within a template (chart/report/map) although you can only have a maximum of two drilldowns for charts and SVG maps. For a report, any column can be used as a drilldown anchor.

To add a layer of drilldown, click the *Add/Modify Drilldown* button in the chart/report/map header bar. A dialog will open that allows you to configure drilldown options for the chart, report or map.



Drilldown Options Dialog

The first option in the dialog, Existing drilldown destinations, is a dropdown selection box that will contain a list of the drilldowns that have been set in the template. The first element is always *--New--*. You can select a drilldown from the list and the component information will be filled automatically. If you choose *--New--*, all dialogs will be cleared and new entries into the dialogs can be entered. The items in the list are names of the drilldown targets with an icon that can help you to identify the type of the file (see Section 3.1.4.3 - File Identification Icons). The *Delete* button will only be enabled when you select an existing drilldown from the *Existing drilldown destinations*. When you hit this button, the current selected drilldown will be removed from the list and the whole dialog will be reset.

The field on the left side of the dialog contains a tree that mirrors the Organizer folder structure. All of the reports/charts/maps and dashboards to which you have access to will be listed. Please note that drilldowns to dashboards will be discussed in the next section.

To select a chart, report or map as the drilldown destination (lower-level), click on it in the list. The dialog will update to show the name of the currently selected drilldown report/chart/map. The first drop-down list below the name will be populated with all of the parameters in the selected report/chart/map.

Each parameter in the destination report/chart/map can be unmapped (i.e. it does not pass the parameter value and shows the drillable link no matter where you click on the report/chart) or mapped to either a column (data element) in the top-level chart/report/map, or one of the shared parameters defined in the dashboard. To set an unmapped link, simply select the *Do not Map To Column*. Note that this option is only available if the next drillable link is a dashboard. To set the parameter, first select the destination parameter in the drop-down list. Then select the radio button to indicate whether it should be mapped to the top-level chart/report/map or a shared parameter. Then select the element/field or shared parameter from the appropriate drop-down list.

Note that a column selected for the drilldown in the *Map to Template Data Column* can only be used once, i.e. you can have only one drilldown per column. A shared parameter in the *Map to Shared Parameter* can be used multiple times in the dashboard. This is because a shared parameter isn't tied to a specific template in the dashboard and it behaves like a parameter value provider for the drilldown destination template.

If all destination template parameters are mapped to shared parameters and if the parent template is a report, the first column of the report is used as the anchor for the drilldown link, although it won't pass the element value to the destination template. If the parent template is a chart or a SVG map, the data points become the anchor for the drilldown links.

You can click the *Set* button to save the drilldown settings that is currently being worked on. If the drilldown is new, a new entry in the *Existing Drilldown destinations* will be added, else the drilldown settings are updated in the browser memory. To make the saved settings permanent, click the *OK* button.

The *Refresh* button reloads the folder tree with latest organizer folder list. Note that all folders will be shown in a collapsed view.

Once you finish selecting the options, click the *Ok* button to save the changes. You will be returned back to the Dashboard Builder interface.

7.2.8.1. Drilldown to Dashboard

In addition to drilldown to any chart/report/map in the dashboard, you can also map the chart data points, report columns, and/or shared parameters to another dashboard. The dashboard builder supports both parameterized and non-parameterized dashboard as the destination of the drilldown. If the target dashboard has no parameter or you don't setup any parameter mapping, the template from which the drilldown link starts will become clickable and the result of the click will popup the target dashboard in a new window without any default parameter setup.

In case you want to set a parameterized dashboard to be a drilldown destination, you will need to map shared parameter(s) from the target (child) dashboard to chart data points/report columns or shared parameters in the mother dashboard. It will behave like setting up a drilldown link for a report, chart or map, as mentioned before. The corresponding field(s) will then be changed to a hyperlink and when you click the link, the target dashboard will pop up and the shared parameter of the target dashboard will be updated.

You can set up a drilldown link without any parameter mapping to another dashboard. If the template has both unmapped and mapped drilldowns and is a report or chart, the drilldown that appears depends on which point on the report/chart was clicked. If it is a datapoint that is used as a drilldown link for a mapped drill-down, a mapped drillable link will appear. If not, an unmapped drillable link will appear.

The other drilldown functionalities are basically the same as mentioned in the previous section. For more details about setting drilldowns and the drilldown options dialog, please navigate to the previous Section 7.2.8 - Drilldown.

7.2.9. Folders

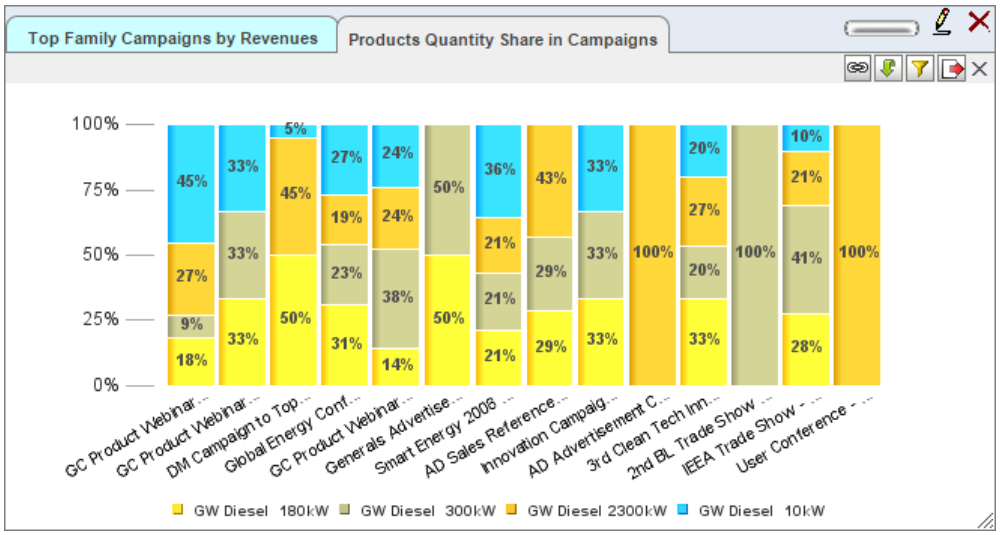
Folders allow you to place tabbed objects into your dashboard. Folders can contain charts, reports or maps, and use a tabbed interface that allows you to switch between current objects to display. You can insert a folder into

a dashboard by clicking the *Insert Folder* button on the dashboard builder toolbar. After you do so, a small rectangle will follow your mouse pointer around the builder interface. Position the rectangle where you want to insert the folder and click. The folder will be inserted into the dashboard. To move a folder, click and drag your mouse over the right side of the folder's header bar.



New Folder

After adding a folder to the dashboard, you can drag and drop reports, charts and maps into it. After you add an object to a folder, a new tab will be visible inside the folder, showing the names of the objects within it. To switch between objects, simply click on the tab of the object you wish to view. Objects can be removed from the folder by clicking and dragging on the object's title bar.

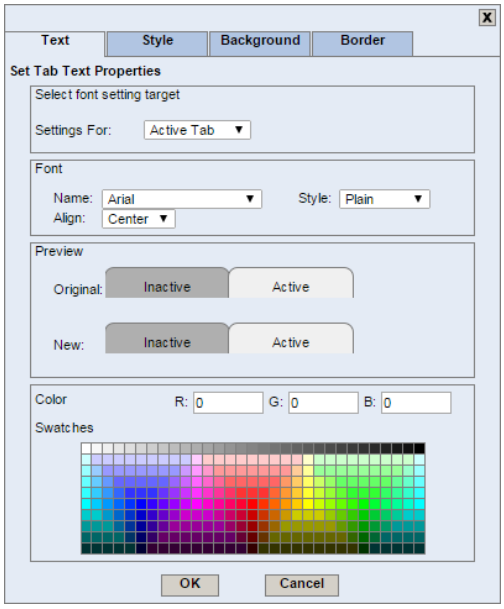


Folder with objects and style formatting

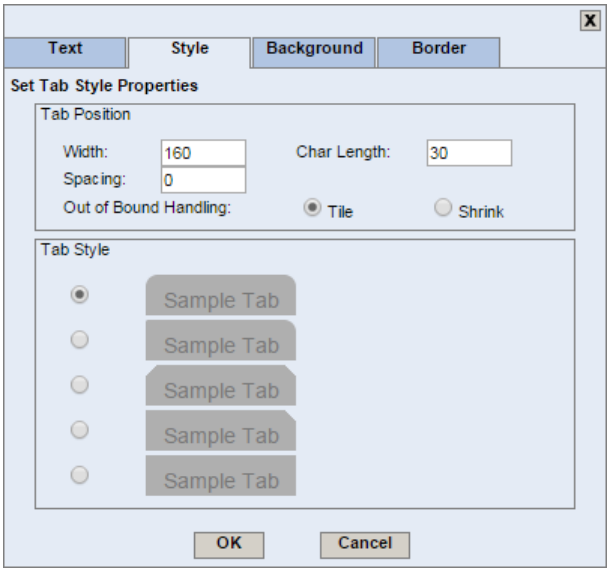
There are two additional buttons available for a folder in the folder header bar:

Edit: This button will open a new dialog that allows you to specify text label properties, style, background color and border for the folder tabs.

Text Properties: The first tab of the edit window allows you to adjust text label properties. You can adjust the font type, font style, font size and alignment of the label as well as changing font color using swatches or RGB values.



Style Properties: The second tab of the edit window allows you to set style properties. You can adjust tab shape and position.



Width: This option allows you to set a width of tabs in pixels.

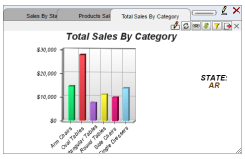
Char Length: This option allows you to set character length. Shortened file names will end with three dots (included in the length).

Spacing: This option allows you to set a space between tabs in pixels.

Out of Bound Handling:

Tile - The width of the tabs will be kept even if the

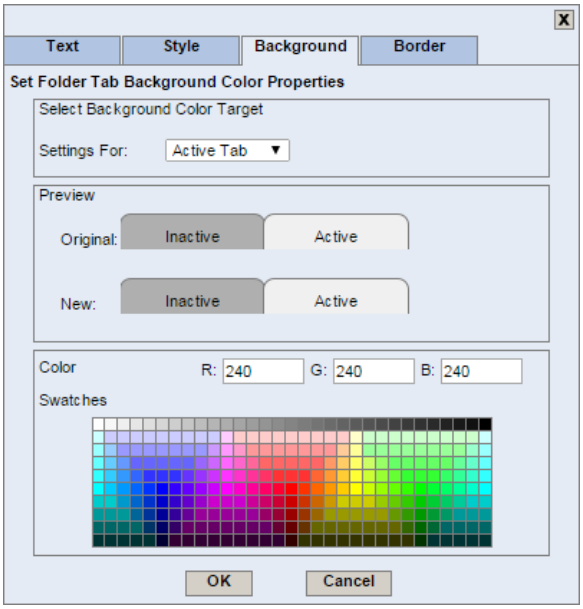
whole tabs are not visible.



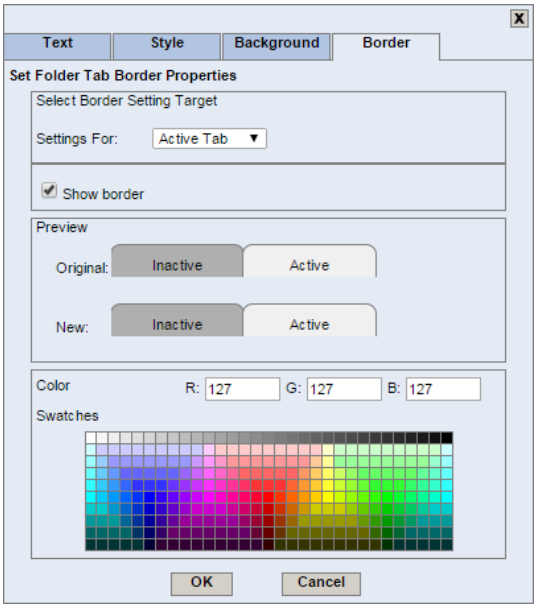
Shrink - The width of the tabs will adjust to make the whole tabs visible.



Background Properties: The third tab of the edit window allows you to set the background color for either active or inactive tabs. The background color can be adjusted using swatches or RGB values.



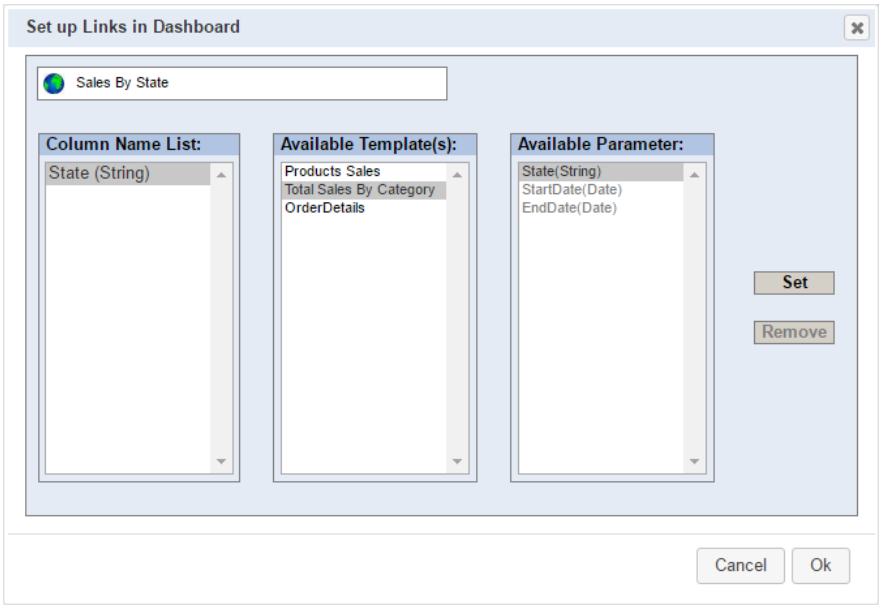
Border Properties: The fourth tab of the edit window allows you to set the tab border color. The tab border can be adjusted using swatches or RGB values.



Delete: This button will delete the folder.

7.2.10. Template Linkage

Template Linkage allows you to use data from one chart, report or map as a parameter in another. To use this feature, click the *Add/Modify link* button on the source chart, report or map to bring up the *Set up Links in Dashboard* screen.



On this screen, there are three columns:

Column Name List - This is a list of available columns and their data type in the source template. Click on a value from this column to use it as a parameter in another template.

Available Template(s) - This lists the other templates that are open in the current dashboard. Click on a template name to select it. While a template is selected, it will be highlighted in the dashboard.

Available Parameter - After selecting a column name and template, this list of available parameters will appear. Select a parameter you want to map to the selected column name in the first column.

When you are done selecting values in each column, click the *Set* button. You can then select new fields if you wish to add additional links. Once a link has been established, clicking on a value in the specified source column (or chart object) will cause the target templates to be refreshed using data from that column.

To remove a link, select it and then click the *Remove* button.

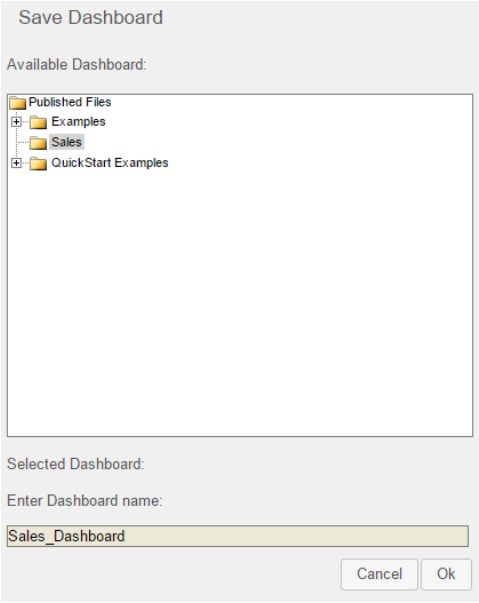
7.2.11. Dashboard Preview

At any point in the design process, the dashboard can be previewed by clicking the *Preview* button  on the dashboard builder toolbar. This will open a new window showing the dashboard. Dashboards that have been deployed to the Organizer can also be viewed in the Menu Page.

For more information about dashboard preview options, please see Section 8.9 - Dashboard Viewer.

7.3. Save Dashboard

Once you have finished setting up all the dashboard definitions, you can save the dashboard by clicking the *Save* button  on the toolbar. A new dialog will open prompting you to specify the name for the dashboard.



Save Dashboard Dialog

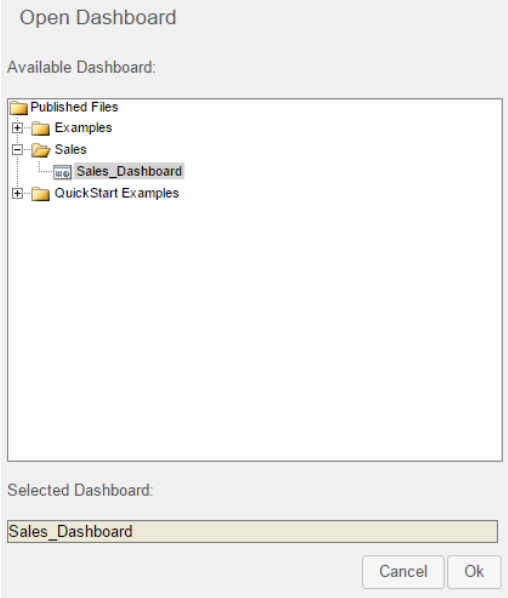
The first option allows you to specify a name for the dashboard. The dashboard will be saved with a `.dsb` extension in the `/DashboardFiles/` directory of your ERES installation.

The second option allows you to insert the dashboard into the Organizer. The drop-down list contains all of the projects and folders in Organizer to which you have access. Select the project or folder in Organizer where you'd like to place the dashboard.

Once you have finished specifying options, click the *Ok* button to save the file.

7.4. Open Saved Dashboard

You can open the saved dashboard by clicking the  *Open* icon on the main toolbar. *Open Dashboard* dialog will appear in the left pane.



Open the Dashboard

You can see all dashboards from Organizer created in Dashboard Builder. Select a dashboard and click *Ok* to open it.

7.5. Exit

You can exit Dashboard Builder by clicking the *Home* icon in the upper right corner, by clicking the *Dashboard Builder* title, or by clicking the *Logo* icon in the upper left corner. Before closing, you will be asked if you want to save an unsaved dashboard.

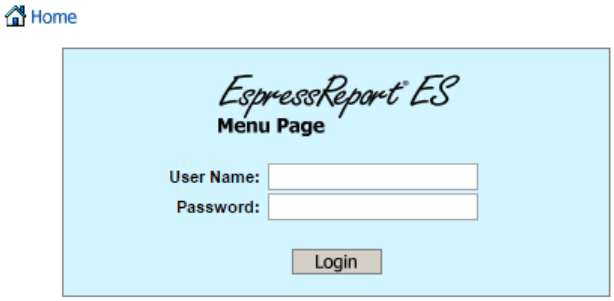
Chapter 8. Publishing (Menu & URLs)

8.1. The Menu Page

In ERES, all reports, charts, maps and dashboards created or added to the Organizer are automatically published. By logging into the menu page, users can access and run any files to which they have privileges in the Organizer. The menu page is included with ERES and deploys without any code.

8.1.1. Launching the Menu Page

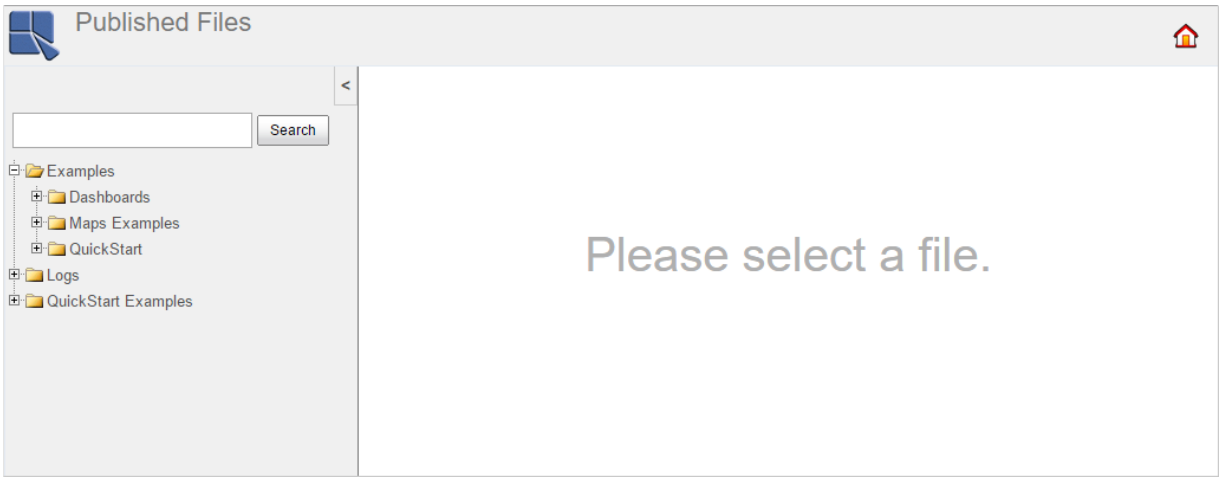
The menu page can be launched from the ERES start page. Once you've successfully logged in, click the *Published Files* link.



Direct Access to Menu Page

8.1.2. Using the Menu Page

When the menu page first loads, it will show a collapsed tree-list of all the reports, charts, maps, dashboards, and other files in the Organizer which you have privileges to see in a tabular format.



Menu Page

8.1.2.1. Viewing Reports and Charts

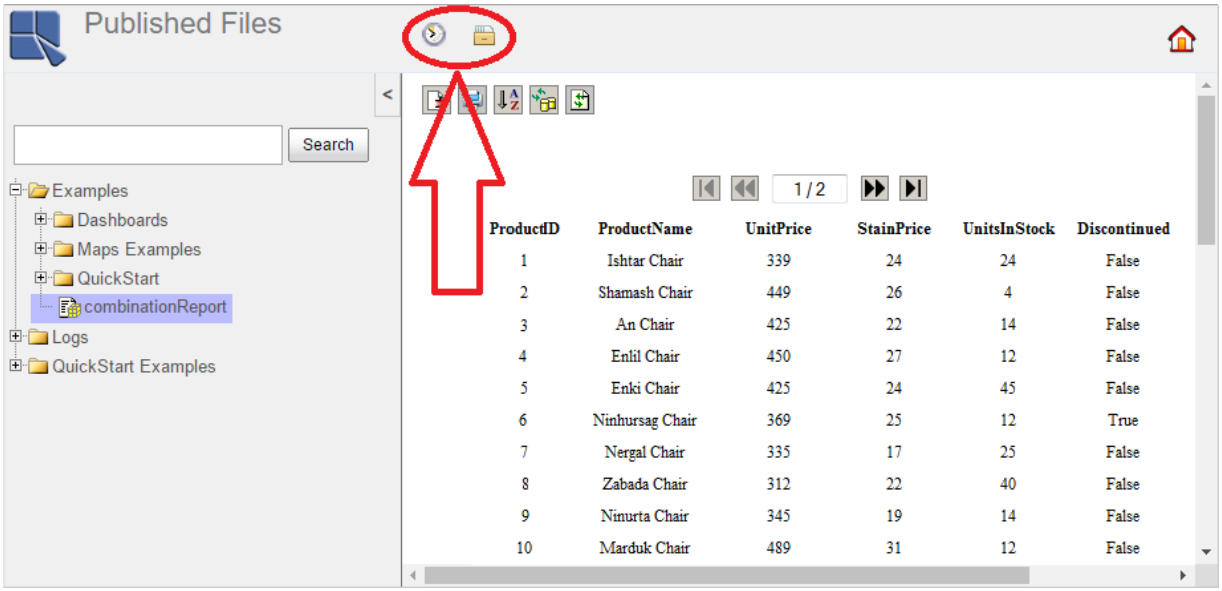
To open a file, expand respective project/folder nodes (in the tree-list on left side) to locate the file and click on the filename. The file will load in the right *DHTML Viewer* panel.

You can also search for specific files. To do so, enter search terms in the upper left dialog and click the *Search* button. The tree-list will reload, showing only relevant files. To cancel the search filter, click the *Close Search* icon next to the search field. Note that the *Close Search* icon will not be visible if there is no filter applied to the list at the moment.

You can hide the left panel by clicking on the < *Collapse* icon located in the top-right corner of the navigation panel.

To close the MenuPage and to return to the ERES start page, click the button in the top-right corner of the MenuPage.

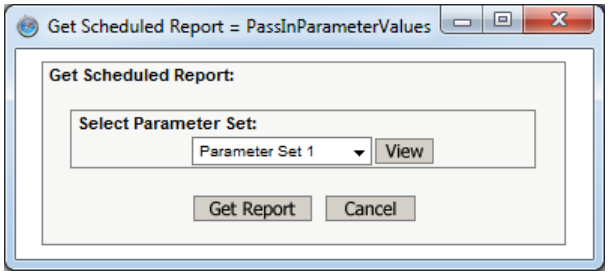
If you open a report or a chart that was scheduled or archived (i.e. the task was run at least one time, it doesn't matter whether the task has already finished or whether it is still active), the following options will be available.



Location of the View Scheduled/Archived version icons

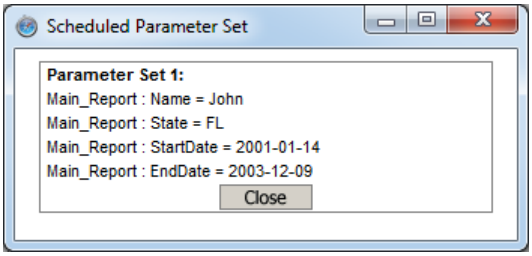
View Scheduled Report/Chart:

This option allow you to view the latest scheduled version for the report or chart with file extensions as `.pac, .cht, .pak, ,rpt, .tpl` and `.xml`. This dialog is only available if an active schedule job is defined for the report or chart. If the report contains parameters, clicking this will bring up a new dialog prompting you to select a parameter set to view for the report.



Select Parameter Set Dialog

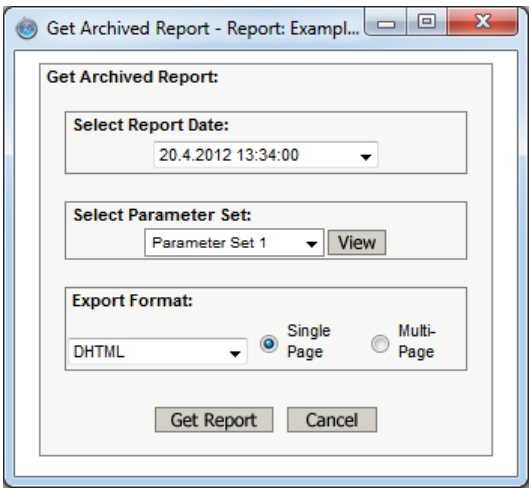
Here you select which parameter set you would like to use to see the latest scheduled version. Parameter sets are defined when the schedule job is created. For more about setting schedules see Section 3.2 - Scheduling & Archiving. You can see which values are included in a parameter set by clicking the *View* button. This will open a new window showing the parameter values.



Parameter Values Dialog

Once you've selected the parameter set that you would like to use, click *Get Report* (or *Get Chart*) and the scheduled report/chart will open in a new window. If the report or chart contains no parameters, the scheduled version will open directly when you click the *View Scheduled Report/Chart* button.

View Archived Report/Chart:
This option allows you to view an archived chart or report. This option is only available if an active archive is defined for the report or chart with file extensions as `.pac, .cht, .pak, ,rpt, .tpl` and `.xml`. Clicking this option will bring up a dialog in a new window allowing you to specify options for the generated report or chart.



View Archive Dialog

The first option allow you to select the version of the report or chart you would like to run. The drop-down list contains dates and times for each version. If the report or chart is parameterized, then the second option will allow you to select a parameter set, in the same manner as for scheduled reports and charts. The last option allow you to select the export format. These options are the same as when running a report or chart.

Once you have finished setting options for the report or chart, hit the *Get Report* or *Get Chart* button. The exported report or chart will open in a new window.

8.1.2.1.1. Viewing Reports with Encrypted Data

If your report contains encrypted data, you need to do two things in order to view the data.

1. You need to create an XML file that gives the database URL, database driver, name of column to be decrypted, and the function to be applied when the data is being retrieved.

2. You need to include the *ReplaceColumnInfoList* option in the command line, i.e. `servercommand.txt` before you start ERES server.

8.1.2.2. Using the DHTML Report Viewer

If you click on a report in the menu page tree-list, the report will open in the DHTML report viewer. This interface is a thin-client component that allows you to view and interact with the report. In the DHTML report viewer, the

first page will show as soon as it is ready, while the rest of the report continues to be processed on the server side. This means that for larger reports, navigating to a page that is not yet exported will result in delay. Before a report is finished processing, the toolbar options will not be available and clicking on any of them will result in a warning message. The main window of the interface provides HTML view of the report with a small toolbar in the top-left corner of the page. Each option on the toolbar will bring up a dialog allowing you to make changes to the report.

The main window also contains a page navigator that contains first page, previous page, next page, last page buttons and a textbox that allows you to jump to any specific page number. The textbox displays the current page number in the form <Current Page Number> / <Total Page Number>. To jump to a specific page, simply type the page number into the textbox and press **Enter**. You can also enter the value in the same fractional format that is used in the display. For example, to jump to page 7 of a 22 page report, you can enter **7** or **7 / 22**.

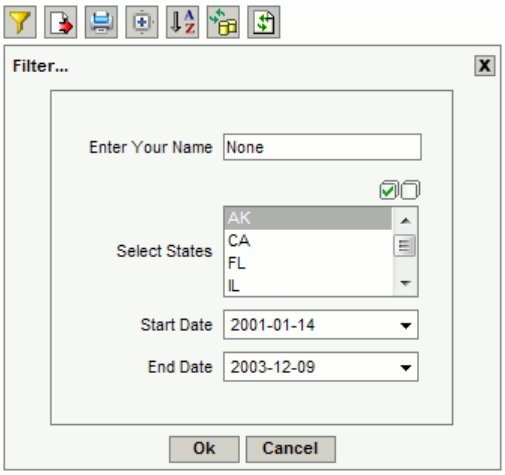


DHTML Viewer Interface

8.1.2.2.1. The DHTML Report Viewer Toolbar

The options in the DHTML report viewer toolbar are as follows:

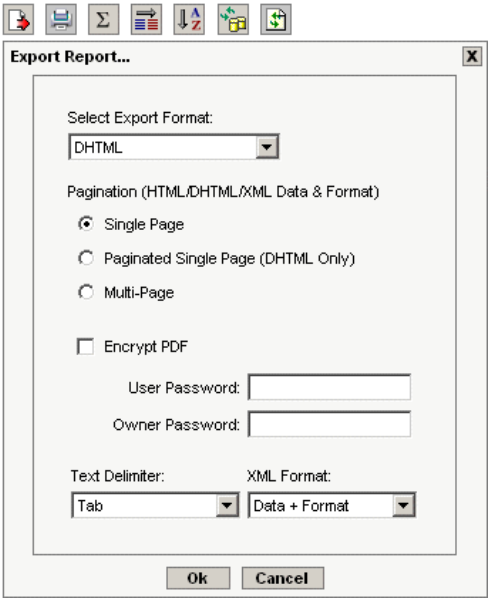
Filter Report: This option is only available if the report is parameterized. It will bring up a dialog containing all parameters defined in the report.



Filter Dialog

Select new values and click the *Ok* button to update the report. There is a validation system for unmapped parameters that will check entered values against expected datatype and alert you if it does not match.

Export Report: This option allow you to export the report in a variety of file formats. It will bring up a dialog prompting you to select export options.



Export Dialog

The first option in the dialog allows you to specify in which format you want to export the report. Available options are HTML, PDF, CSV, Excel (XLS), Excel 2007 (XLSX), DHTML, text, XML and rich text.

Other options allow you to set single or multi-page export for HTML, DHTML, and XML exports. PDF format allows you to enable encryption. For text export, you can select delimiter to be used, and for XML export, you can select whether to export only report data or XML description of the report and data.

Once you finish specifying options for the exported report, click the *Ok* button. A new window will open containing the exported report. The toolbar will be active on this window, allowing you to save the generated file to your local system or print the report locally.

Print Current Report: This option allows you to print current page of the report on user's local printer, as displayed inside the right-side panel.

1. By default, all browsers don't print background colors and images to save ink. You have to change print settings to allow that.

   How to set it in different browsers:

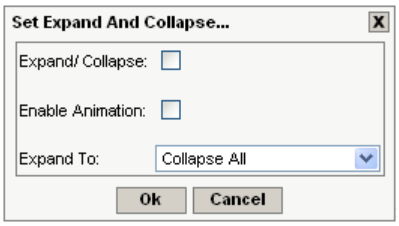
   Chrome: Check the *Background colors and images* option. It is in the *Print* dialog that opens automatically.

   IE 11: Click the gear icon in the upper right corner → Print → Page Setup → check *Print Background Colors and Images*

   Firefox: Print → Page Setup → check *Print Background (colors & images)*

2. Reports wider than page might be cut off. If you need to print wider reports, please export them to PDF and print the PDF.

Expand and Collapse View: This option allows you to add expand and collapse controls to any grouped report.

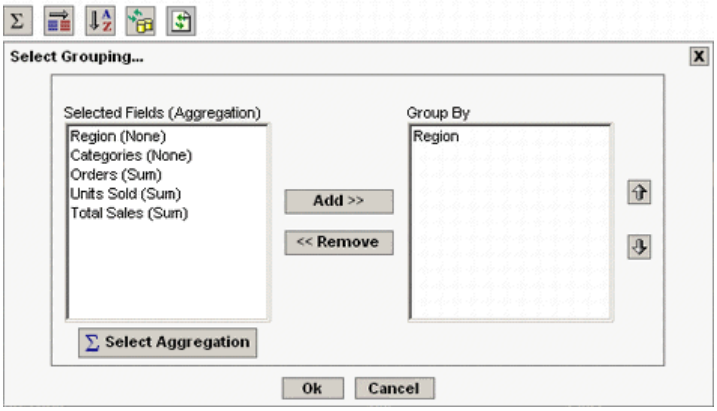


Expand and Collapse Dialog

The first option in the dialog allows you to turn the control for expanding and collapsing on. Enable Animation allows the grouping to fade in and out as it is expanded and collapsed. Expand To specifies the initial presentation of the report, i.e. collapsed completely, expanded completely or expanded to a specific group level.

Once you finish specifying the options, click the *Ok* button. The controls will now be added to the report.

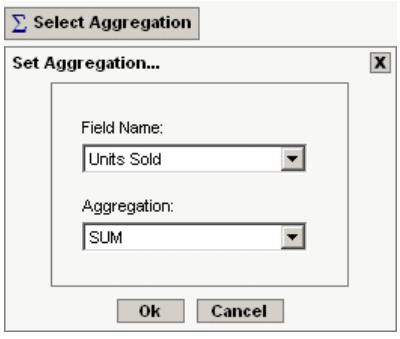
Change Grouping/Summaries: This option is only available if the report was created using QuickDesigner. It opens a dialog that allows you to edit the report grouping and aggregation.



Grouping Dialog

To add or remove grouping, select a column and click the *Add* or *Remove* button. If you select more than one level of grouping, you can change the nesting order using the arrow buttons on the right side of the dialog.

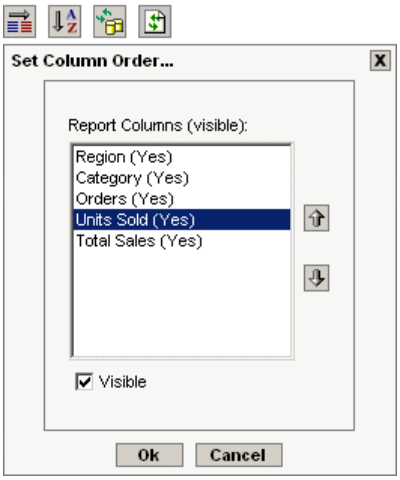
To set column aggregation, click the *Select Aggregation* button below the column list on the left side. This will open a new dialog allowing you to change the aggregation for the columns in the report.



Aggregation Dialog

The top drop-down list contains all selected columns for the report and the lower list contains all available aggregation options. To set aggregation for a column, select it in upper list and set the desired aggregation in lower list. Once you finish setting up the aggregation, click the *Ok* button and the dialog will close. The new aggregations will be reflected in the column list.

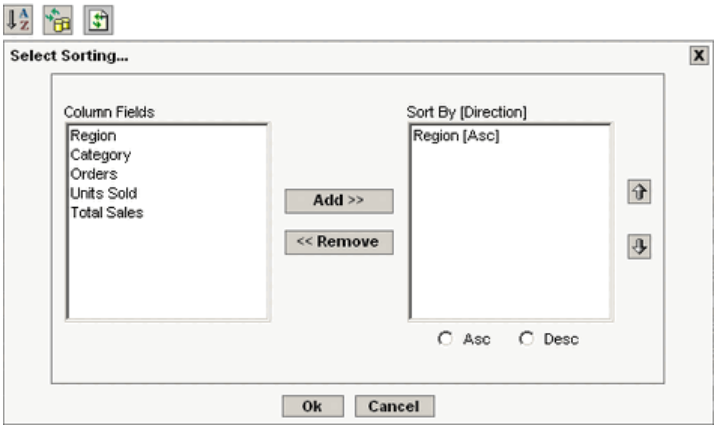
Change Column Order/Visibility: This option is only available if the report was created using Quick-Designer. It opens a dialog that allows you to change the column order in the report as well as show/hide columns.



Column Order Dialog

The dialog contains a list of all columns in the report. To change the column order, select a column in the list and click on the *up* or *down arrow* buttons to change the order. The order from top to bottom in the dialog represents the column order from left to right in the report. To show or hide a report column, check or un-check *Visible* option in the bottom of the dialog.

Change Sorting Options: This option allows you to change the sorting in the report. It will bring up a dialog prompting you to set sorting options.



Sorting Dialog

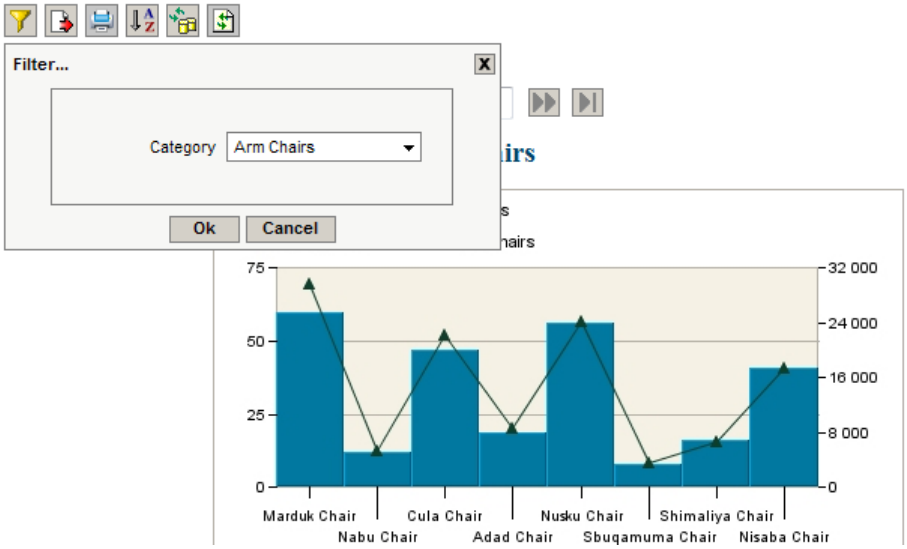
The left side of the dialog contains all columns in the report. To sort it by a column, select it in the left side and click the *Add* button. You can control the sort priority by selecting a sort-by column on the right side and clicking on the *up* or *down arrow* buttons to change its position. You can control the sort direction by selecting a sort-by column on the right side and clicking on the *Asc* or *Desc* radio buttons to indicate the direction.

Refresh Data: The DHTML Viewer caches the report and therefore when you close and reopen the same report, the data is not updated. This option will refresh the data in the report.

Reload Report: When you run a report in the DHTML Viewer, make some changes to this report and then try to open this report again in DHTML Viewer, you'll notice that the report does not reflect the new changes. This is because the DHTML Viewer caches the report to improve performance. To see the changes, use the *Reload Report* button to have the DHTML Viewer read the report structure again and also refresh the data.

8.1.2.2.2. Parameterized Reports and Reports with Unmapped Drilldown Parameters

When viewing a parameterized report in the DHTML Viewer, the parameters are initially set to their default values. You can then change the parameter value using the *Filter* icon described above. The DHTML Viewer can also be used to view drill-down reports with unmapped parameters. Similar to a regular parameterized report, the drilldown report will display using default values, but you can adjust them by using the *Filter* icon.



Drilldown with Unmapped Parameter

8.1.2.3. Using the DHTML Chart Viewer

If you click on a chart in the menu page tree-list, the chart will open in the DHTML chart viewer. This interface is a thin-client component that allows you to view and interact with the chart. The DHTML chart viewer displays the chart using default parameter value (if the chart is parameterized). The parameter value can be changed using the toolbar options. The main window of the interface provides you with HTML view of the chart with a small toolbar in the top-left corner of the page. Each option on the toolbar will bring up a dialog allowing you to modify the chart's look and feel, and navigate through the chart.

8.1.2.3.1. The DHTML Chart Viewer Toolbar

The options in the DHTML chart viewer toolbar are as follows:



Filter Parameter : This option is only available when the chart is parameterized. Clicking on the icon will bring up a dialog that allows you to specify parameter values for the chart. You can choose/select the values and then click the *OK* button to update the chart. For unmapped parameter, the value is validated against the expected data type and alerts you in case of a mismatch.



Export Chart : This option allows you to export the chart to a variety of formats such as PDF, PNG, GIF, JPEG and SVG. The exported chart is then displayed in a new window from where it can be saved to your local system or printed.



3D Display Option : This option is only available when the chart is hree dimensional. Three-dimensional charts have several additional options. These options can be accessed by clicking the *3D Display Options* button on the toolbar. This will open a dialog in a new window. Available options include:

- Arrow buttons at the top of the dialog allows you to adjust the viewing angle of the chart. The 3D chart can be rotated horizontally and vertically.
- The first three sliders allows you to control the chart scaling of the chart in the X, Y, and Z directions. The scaling features control the proportion of the chart in each direction. The final slider controls the thickness of the data points.

- The remaining options allow you to toggle wire frame mode for the chart, show/hide the border around the chart, and draw the series inline (instead of on the Z axis) for charts with data series. The last option also enables 3D approximation for the chart. 3D approximation is useful when dealing with charts with a large number of data points. The approximation can improve 3D rendering performance when a large amount of data points exists on the chart. This feature is automatically enabled when the chart contains more than 100 data points.

Once you finish setting the options, click the *Ok* button to save the changes.



Zoom Option : The Zoom option dialog is only available for a chart with zoom feature enabled. You can modify the zoom properties such as lower bound, upper bound and zoom scale, as well as change the period layout to linear. Then click the *Ok* button to apply the new zoom changes.



Refresh Data : You can use the refresh data option to fetch the latest data for the chart.



Reload Chart : Changes in the chart are not always reflected in the DHTML chart viewer. This is because the DHTML chart viewer caches the chart to improve performance. You can click on the *Reload* icon to display the changes and refresh the data.



Back Button : The *Back* button appears if you are in a sub-level chart. Clicking the *Back* button will take you back to the previous level.

The DHTML chart viewer also supports drill-down. You can drill-down to the next level chart (if the chart is a drill-down chart) by clicking on the desired data point.

The DHTML chart viewer also pops up a hintbox that displays the data information/hyperlink information for the specific data point. If a hyperlink is embedded, clicking on the data point will direct the user to the linked page.

8.1.3. Mobile MenuPage

If you access ERES main page from a mobile device, a special simplified ERES start page will be displayed.



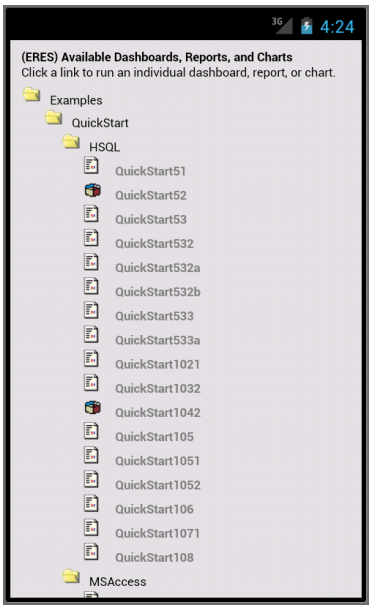
ERES mobile start page



Note

The mobile features don't work in so-called *PC view*, which can be enabled in some mobile web browsers. In *PC view*, the standard ERES start page and MenuPage will be displayed.

Enter your user name and password and tap on the *Login* button. You will be redirected directly to a special mobile MenuPage which allows you to view reports, charts, maps and dashboards.



Mobile MenuPage

The mobile MenuPage is easy to use - scroll/swipe up or down to locate a template or a dashboard and tap on its name to open it.

Mobile Report Viewer

If you open a report from the mobile MenuPage, standard DHTML Viewer will open. To learn more about the DHTML Viewer, see Section 8.1.2.2 - Using the DHTML Report Viewer.

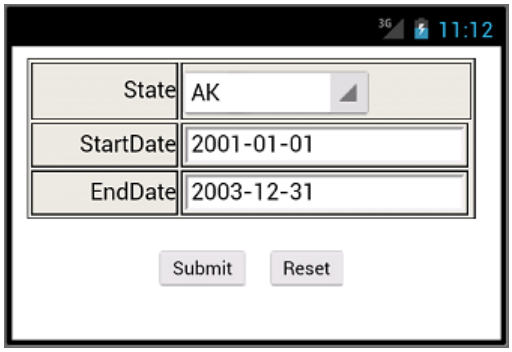


DHTML Viewer

Mobile Chart and Map Viewer

Chart/map viewer has no control elements.

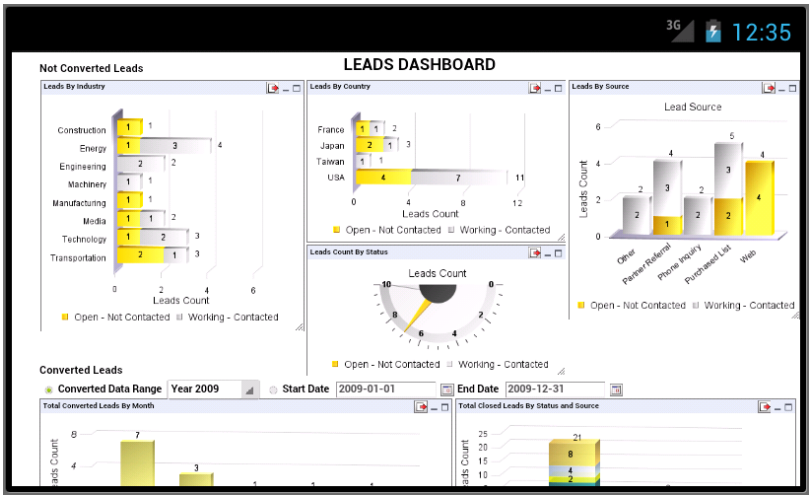
If the chart/map has any parameters, a parameter prompt will be displayed first. Enter any parameter values and tap on the *Submit* button to display the chart/map.



Parameter prompt

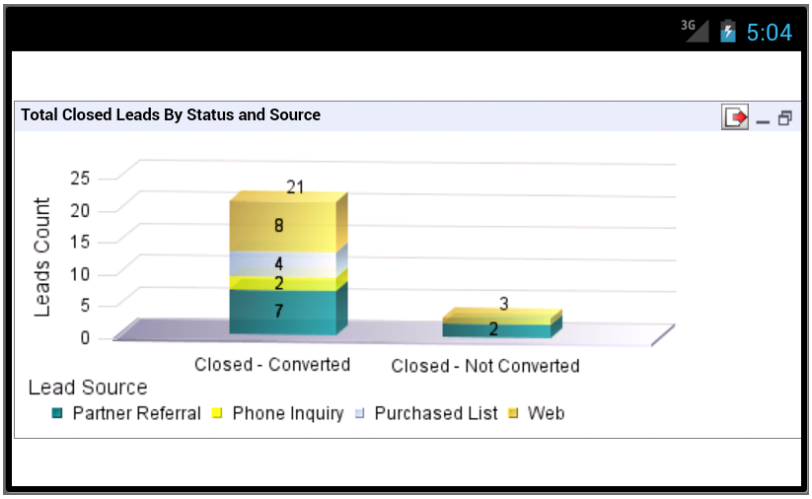
Dashboards

If you open a dashboard on a mobile device, you can swipe the dashboard to navigate in it.



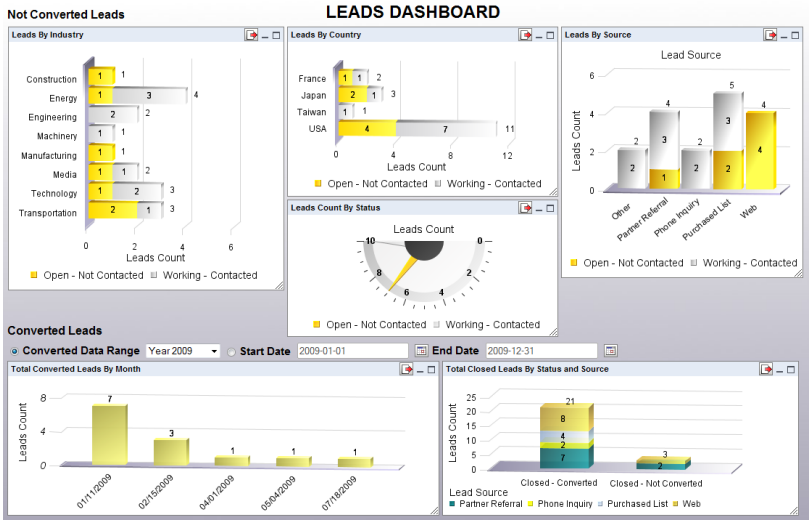
However, on smaller devices, navigating through large dashboards may require extensive zooming and scrolling. In such cases, you can maximize a dashboard template to display the dashboard in a more mobile-friendly way.

To maximize a template, click on its ☐ *maximize* icon, or double-tap/pinch on the template.

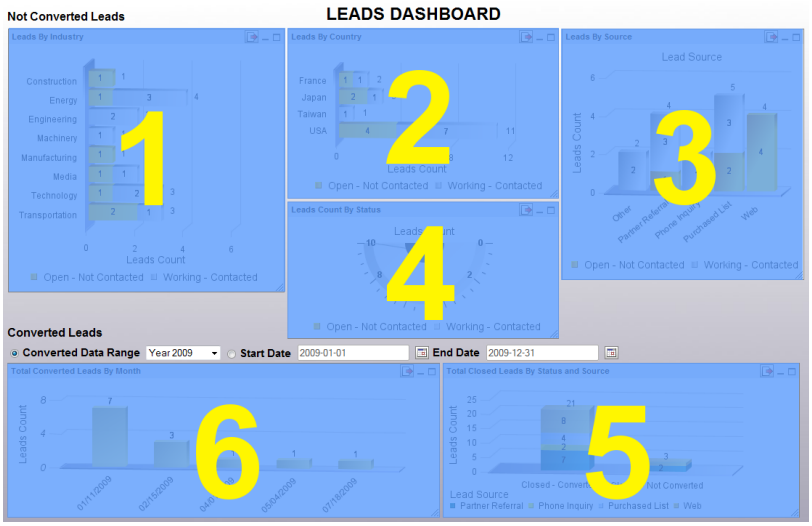


When navigating in maximized-mode, all dashboard templates will be arranged in a single loop, so you don't have to remember the dashboard lay-out at all. If you swipe on a maximized template to the left, the next template will be displayed.

For example: Templates from the following dashboard:



Would be displayed in the following order:



8.2. Image URLs

In addition to the menu and the API, EspressReport ES also allows you to deploy charts using image URLs (Reports can also be deployed using URLs as described in Section 8.3 - Report URLs). Image URLs are `http` calls to the server that can either return a chart as a complete `html` page or return an image to be embedded in an existing Web page in an `<img src>` or `<embed>` tag (SVG/PDF). You can use an URL to call a pre-defined chart (`.cht` file) or chart template (`.tpl` file), or to create a chart using default appearance properties. Parameters entered into the URL can control various chart properties.

8.2.1. Generating URLs in Organizer

The easiest way to generate image URLs for your Web pages is to have the Organizer do it for you. To generate an image URL in the Organizer, first select a chart file (`.cht` or `.tpl`) that you want to use. Then you can either select *Generate Image URL* from the Publish menu, click the *Generate URL* button on the toolbar, or right-click and select *Generate Image URL* option from the pop-up menu. A dialog will appear prompting you to specify attributes for the generated image.

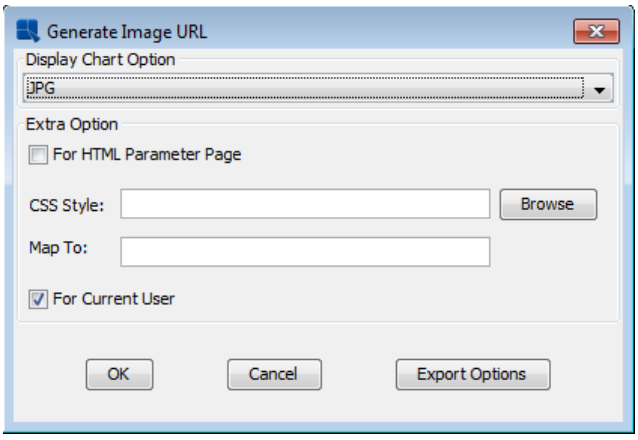
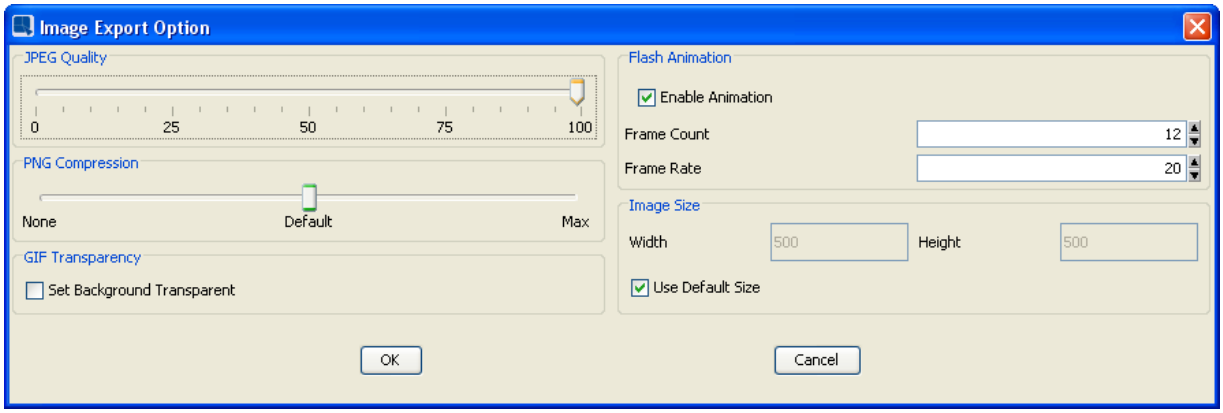


Image URL Options Dialog

The first option in the dialog allows you to select the export format for the chart - GIF, JPEG, PNG, PDF, SWF (Flash), or SVG. If the chart contains parameters, you will see options for the HTML Parameter Page (shown above). Unchecking the *For Current User* button will prompt viewers to login first in order to view the chart. You can set additional export options by clicking the *Export Options* button. This will bring up a new dialog with several options for the different chart formats.



Export Options Dialog

Options in this dialog allow you to set the JPEG quality. The higher the quality setting, the larger the generated image. You can also set PNG compression and whether to make the exported GIF image transparent. The export method allows you to select either frame or buffer method for generating the chart. Each method can result in faster performance depending on the system, the size of the generated image, and the number of data points in the chart. You can also enable and configure Flash animation if the image depiction should be animated. The last option allows you to specify the image size. You can specify the image size in pixels, or select default size. Clicking the *Use Default Size* button causes the image to be the same size as the chart canvas.

8.2.1.1. URLs for Parameterized Charts

If your chart is parameterized, then there are two ways in which the chart can be run. You can either pass parameters directly into the URL as arguments, or have the URL generate a parameter prompt. If the chart contains parameters, then several options will appear in the image URL options dialog as pictured above. To generate an HTML parameter prompt check the *For HTML Parameter Page* option.

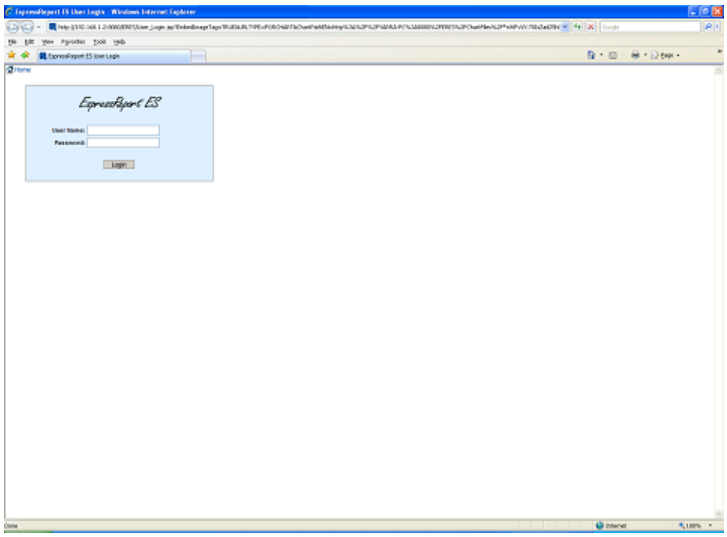
If you select to generate a parameter prompt, you can also specify a CSS file to use to format the prompt. The dialog has options allowing you to specify the file and http path to the CSS file (An example of a parameter page CSS file is available under `<ERESInstallDir>/help/examples/URL`).

If you select not to generate a parameter prompt, then a dialog will open asking you to select parameter values which will be added as arguments to the generated URL.

For information on running the chart URL, see Section 8.2.1.3 - Running Image URLs.

8.2.1.2. For Current User

Checking this option will embed the current user in the URL. If you uncheck this option, the user will be re-directed to a login page before viewing the chart.



User Login Dialog

By default, the URL will redirect to the login page under the ERES context (i.e. `http://machinename:port/ERES`. If you have deployed ERES in a different context, you can use the `JSPCONTEXT` parameter to pass in a different location.

8.2.1.3. Running Image URLs

Once you finish specifying options for the image URL, a new dialog will open containing the generated URL string. If you're running Organizer as an applet, the URL will load in a new browser window (allowing you to copy and paste).



Generated URL

Image URLs can be run in one of two ways. The image URL can either generate a whole (HTML) page containing the generated chart or it can be used to stream the chart image into an existing page via an image tag (i.e. `<img src="ImageURL">`). To generate the chart as a complete page, you need to add the following parameter to the URL: `EmbedImageTag=true`. Additionally, if the URL re-directs to the login page (no username and password supplied) or if you select to generate a parameter prompt for the chart, it will be returned in a complete page.

To embed the chart in an existing page, you must pass an username and password in with the URL or have it retrieve user information from the session (unless you're creating a new chart with the URL) and pass any parameter values into the chart in the URL.

8.2.2. Writing Image URLs

Rather than having the Organizer generate the image URL, you can write your own. You can use an existing chart or template file, or you can generate a chart on the fly using default properties. You can control various chart attributes using URL parameters. You can also modify an URL generated by the Organizer.

8.2.2.1. URL Syntax

The syntax for image URLs is fairly simple. Every URL begins with a call to the Server: `http://<machinename>:<port>/<context>/LookupServlet?`. The machinename is either the URL, machi-

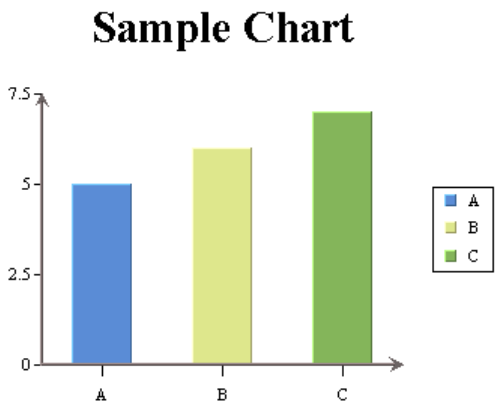
nename, or ip address of the server machine. The port is the port number for the application server, the default Tomcat for instance uses port 8080. The context is the context for the ERES servlets. By default, the context is `ERES/servlet/`, but this can be changed by the administrator. Following the question mark users need to specify whether the URL will return a chart or report using the `URLTYPE` parameter. Following the `URLTYPE` users can add parameters to the URL. Parameters are separated by the "`&`" character.

The following structure will generate a simple column chart in a Web page.

```
<img src="http://<machinename>:<port>/<context>/LookupServlet?
URLTYPE=FORCHART&MainTitle=Sample+Chart&ChartType=Column&CategoryName=A;B;C
&ChartData=series1;5;6;7&DisplayType=GIF&Width=400&Height=300">
```

For example, if the server was set up on www.quadbase.com [http://www.quadbase.com] with port 8080 and the context is `ERES/servlet/`, then the image tag should look like this.

```
<img src="http://www.quadbase.com:8080/ERES/servlet/LookupServlet?
URLTYPE=FORCHART&MainTitle=Sample+Chart&ChartType=Column&CategoryName=A;B;C
&ChartData=series1;5;6;7&DisplayType=GIF&Width=400&Height=300">
```



Generated Chart

As you can see, this URL generates a column chart with three points on the category (X) axis, and three data points in 1 series on the value (Y) axis.

In order for an URL to run, you must at minimum specify a chart or template file, or specify data, either directly within the URL or using a file or database. Also any input parameter string should use the "+" character in place of a space. An actual space character may disrupt the creation of the image file.

The same URL could be used to generate a Web page with the chart embedded. Simply add the following parameter to the end: `EmbedImageTag=true`. Then the URL can be used as a direct link, and doesn't need to be included in an existing Web page.

8.2.2.2. URL Parameters

URL parameters allow you to specify which chart or template file to use, as well as data or data source information, and chart appearance such as chart type and image dimensions.

8.2.2.2.1. Login Parameters

Login parameters allow you to pass username and password into image URLs. A login is required if the URL references a chart or template file (using the `ChartPath` or `TemplatePath` parameters). If login information is not specified in the URL, then the URL will re-direct to a login page before showing the chart.

USERNAME: This allows you to specify an username. If you generate an URL in Organizer and select to use an user, the username will be encoded in the generated URL. You can also specify the username as plain text.

```
&USERNAME=user
```

PASS: This allows you to specify the user password. If you generate an URL in Organizer and select to use an user, the password will be encoded in the generated URL. The password can also be specified using plain text.

```
&PASS=password
```

USESESSION: This allows you to have the server read the username and password from the session. If this option is enabled the information will be retrieved from session parameters `USERNAME` and `PASS`. This option is only available if you will be deploying the URL within a servlet or JSP application.

```
&USESESSION=true
```

JSPCONTEXT: This allows you to specify the location (context) of the ERES installation. This location is used when the URL re-directs to the login page. The default value is `/ERES`, and the URL will work without setting this parameter in most deployments. You only need to specify this parameter if ERES is deployed in a context of a different name.

```
&JSPCONTEXT=/MyReports
```

8.2.2.2.2. Data Input Parameters

TemplatePath: This allows you to specify the `.tpl` or template file that you would like to use to create the chart. Because the `.tpl` format is saved without data, you can specify different data when using templates. If you do not specify different data ERES will try to reload data from the template's original source.

```
&TemplatePath=http://machinename:port/ChartFiles/
templatefile.tpl
```

If *USERNAME* and *PASSWORD* is not provided, a login screen will appear prompting for one. If the user provided does not have permission for the template, a chart will not appear.



Note

You can use both absolute and relative paths for the template file location.

ChartPath: This allows you to specify the `.cht` or chart file that you would like to use to create the chart. Because the `.cht` file is saved with the chart data, you cannot specify different data when using chart files.

```
&ChartPath=http://machinename:port/ChartFiles/
chartfile.cht
```



Note

You can use both absolute and relative paths for the chart file location. If either a chart or template file is specified in the URL, then a login is required to access the file.

DataFilePath: This allows you to specify a data file that will be used to plot the chart, either text or XML. For more on data file specifications, please see Section 4.1.4 - Data from Text Files.

```
&DataFilePath=http://machinename:port/chartdata.txt
```



Note

You can use both absolute and relative paths (relative to ERES installation directory) for the data file location.

From XML files you can use only those, that are in Quadbase format (please see Section 4.1.3 - Data from XML and XBRL Files).

DataSourceClass: This allows you to specify a class file that will retrieve data to plot the chart.

```
&DataSourceClass=package.class
```

DBSourceInfo: This allows you to specify a database and SQL query to be used to retrieve chart data. The parameter takes five input values: `database URL, database driver, username, password,` and `query`. Each input value is separated by a semi-colon. For more on connecting to a database, please see Section 4.1.2 - Data from a Database.

```
&DBSourceInfo=url;driver;username;password;query
```

If you do not wish to enter an `username` or `password` use the space character "+" instead. Do not drop the input value. The parameter will not run with insufficient inputs. In your SQL statement be sure to substitute the "+" character for any spaces. Also, substitute the word `equal` or `equals` for the "=" sign.

QueryParamName: This parameter is used if the chart or template file you're using contains a parameterized query. QueryParamName specifies the name of the query parameter you will be supplying a value for. (Parameter names are assigned when you initialize a query parameter. For more on this, see Section 4.1.2.2.2 - Parameterized Queries) It is always followed by the `QueryParamSize`, and the `QueryParamValue` parameter.



Caution

You must specify all query parameters (or none of them), else your chart will not be displayed.

```
&QueryParamName=customerID&QueryParamSize=1&QueryParamValue=4
```

QueryParamSize: This parameter is used to specify the number of values that will be passed into a particular chart parameter. As detailed in Section 4.1.2.2.2.1 - Multi-Value Parameters, some

queries can have multi-value parameters. This URL parameter allows you to specify if the parameter (indicated by the previous `QueryParamName` parameter) takes multiple values. This parameter is always followed by one or more `QueryParamValue` arguments for each of the paramater values for a particular parameter. If you do not specify this parameter then it will assume that the number of values to be passed in is one.

QueryParamValue: This parameter is used if the chart or template file you're using contains a parameterized query. `QueryParamValue` specifies the value that you would like to pass to the query parameter specified by the `QueryParamName` argument. For multi-value parameters a separate `QueryParamValue` argument must be supplied for each distinct parameter value that you're passing in. For more on creating parameterized charts, please see Section 4.1.2.2.2 - Parameterized Queries.

ColumnMapping: If you have specified a data source for the generated chart, this argument allows you to map columns from your data file, or database query result set to the chart. Column mapping is specified using index values. Index values are assigned based on the order they appear in your data file from left to right, or the order in which they are selected in your SQL statement. Index values start with **`0`**, so the first column is **`0`** and the second is **`1`**, etc. To perform column mapping, specify the column index for the respective axes for the chart. If you do not want to map one of the axes use the value **`-1`**. Do not leave it blank. Mapping varies for different chart types.



Caution

You must specify this parameter along with `DataFilePath` parameter for proper chart generation.

For Scatter, Surface, and Bubble Charts:

```
&ColumnMapping=series;xvalue;yvalue;zvalue;subValue
```

For High-Low, and HLCO Charts:

```
&ColumnMapping=series;category;high;low;open;close;subValue
```

For Gantt Charts:

```
&ColumnMapping=series;category;end;start;-1;-1;subValue
```

For all other chart types:

```
&ColumnMapping=series;category;sumby;value;subValue
```

For more detailed information about column mapping for all chart types, please see Section 5.2.3.1 - Data Mapping.

TransposeData: This is a Boolean flag that denotes whether the data is to be transposed or not. It is followed by the `SelectTranspose` parameter which dictates which columns to transpose. By default this is false.

```
&TransposeData=true
```

SelectTranspose: This allows you to specify the columns that you would like to transpose. To select columns, specify the column indices separated by a semi-colon. If you do not specify any columns then all of them will be transposed.

```
&SelectTranspose=ColumnIndex;ColumnIndex…
```

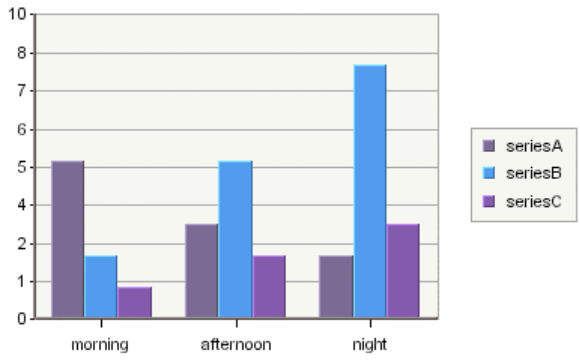
For an explanation of how data transposition works, please see Section 5.2.1.3.2 - Transposing Data.

ChartData: This parameter allows you to supply data for the chart within the URL string. The `ChartData` parameter specifies the data points and series for the chart. It is generally used in conjunction with the `CategoryName` parameter which specifies the individual category names for the data points, and the `DataType` parameter which specifies that data type for your chart data. Stack-type charts with a sum by value, can have the sum by names specified using the SumbyName parameter.

The structure of the `ChartData` parameter is a series name followed by data points. Each series is separated by the | character, and each data point is separated with a semi-colon. For chart types that require more than one value per data point, the sub-values are separated by a comma. The following examples illustrate how the `ChartData` parameter is used for various chart types.

For Column, Bar, Area, Line, Pie, and Dial charts, the structure resembles the following:

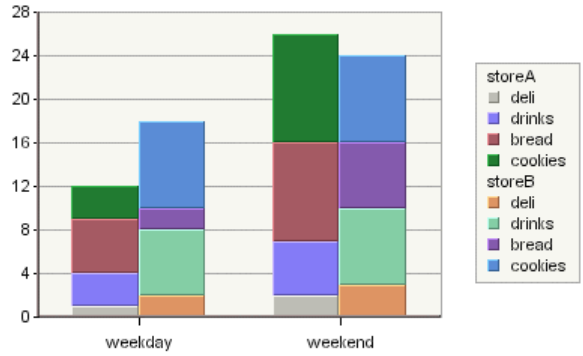
```
&ChartData=seriesA;5;3;2|seriesB;2;5;8|seriesC;1;2;3
&CategoryName=morning;afternoon;night&DataType=integer
```



This generates a chart that has three series, `seriesA`, `seriesB`, and `seriesC`. The chart also has three integer data points, `morning`, `afternoon`, and `night` within each of the series.

For Stack Column Stack Bar, Stack Area, and Percent Column charts the `SumbyName` parameter has to be specified in addition to the `CategoryName` and `DataType` parameters. The structure for these charts resembles the following:

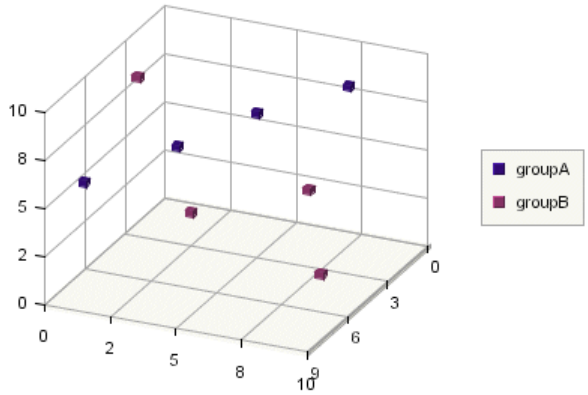
```
&ChartData=storeA;1,3,5,3;2,5,9,10|storeB;2,6,2,8;3,7,6,8
&CategoryName=weekday;weekend&SumbyName=deli;drinks;bread;cookies
&DataType=integer
```



This generates a chart with two series, `storeA`, and `storeB`, which have two respective points. The individual data points have four sum by segments separated by a comma.

For Bubble, Scatter and Surface Charts, the `ChartData` parameter value structure is similar to that of the Stack-type charts, however, rather than using the `CategoryName` parameter for the categories, it uses `CategoryLength`, an integer value that specifies the maximum length of the series. The structure resembles the following:

```
&ChartData=groupA;2,5,3;4,6,1;1,6,8;8,9,2|
groupB;1,9,4;7,4,3;9,2,6;2,1,2&CategoryLength=4&DataType=integer
```



For Surface charts, there are some additional restrictions. First, surface charts cannot contain more than one series. Second, each data point must be of the format `x,z,y,` the `x` and `z` values signify the position on the horizontal plane, while the `y` value determines the height. And finally, every unique `x` value must be paired with every unique `z` value, for example the following is an acceptable dataset

```
0,1,5;
0,2,6;
0,3,5;
1,1,5;
1,2,6;
1,3,5
```

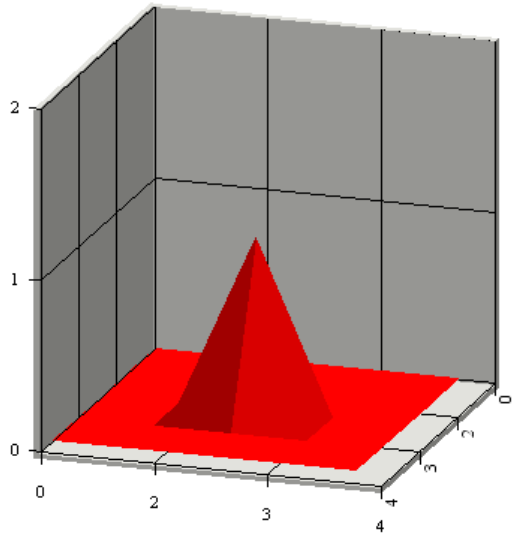
Every unique `x` value (`0` and `1`) were paired with every unique z value (`1, 2` and `3`). While the following dataset is not acceptable

```
0,1,2;
0,2,4;
0,3,2;
1,1,2;
1,3,2
```

because when `x = 1`, there is no corresponding `z` value which equals `2`. Also note that there is no restriction on the `y` value. It is not necessary to use the `CategoryLength` parameter since there can only be one series, for example:

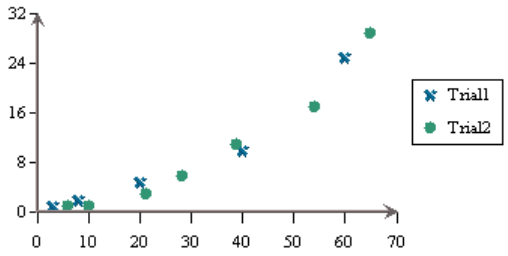
```
&ChartData=A;
0,0,0;0,1,0;0,2,0;0,3,0;0,4,0;
1,0,0;1,1,0;1,2,0;1,3,0;1,4,0;
2,0,0;2,1,0;2,2,1;2,3,0;2,4,0;
3,0,0;3,1,0;3,2,0;3,3,0;3,4,0;
4,0,0;4,1,0;4,2,0;4,3,0;4,4,0
```

Yields the following graph:



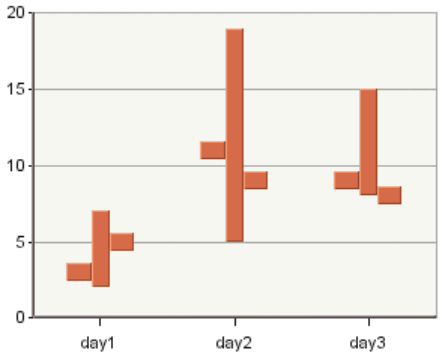
For two-dimensional Scatter Charts, each data point has two sub data points. All other charts have three sub data points separated by a comma. In addition, if the number of data points is not the same for all series, the maximum number of data points in a series must be specified. Use the `CategoryLength` parameter to specify the maximum number of data points, for example:

```
&ChartData=Trial1;3,1;8,2;20,5;40,10;60,25|
Trial2;6,1;10,1;21,3;28,6;39,11;54,17;65,29
&CategoryLength=7
```



For HLCO, High-Low, and Gantt Charts, the ChartData parameter specification is similar to that of the Stack-type charts. For HLCO Charts, 4 sub data points within each individual data point specify the `High, Low, Open,` and `Close` values. For High-Low Charts, 2 sub data points specify the `High`, and `Low` values. For Gantt Charts, 2 sub data points specify the `Start` and `End` values. The `ChartData` structure resembles the following:

```
&ChartData=priceA;7,2,3,5;19,5,11,9;15,8,9,8
&CategoryName=day1;day2;day3&DataType=integer
```



This data generates a HLCO chart with 1 series `priceA`, and three integer data points `day1, day2,` and `day3`. Each data point has a `High, Low, Open,` and `Close` value.

For more information about the various chart types, and data mapping, please see Section 5.2.3 - Chart Types and Data Mapping.

CategoryName: This parameter allows you to specify the individual category names of the data points. Each category name is separated by a semi-colon. The data points are specified within the `ChartData` parameter.

```
&CategoryName=day1;day2;day3…
```

CategoryLength: This parameter allows you to specify the maximum length of the series in scatter, and bubble charts. The argument is specified as a single integer. This needs to be specified when the individual series are not the same length. If `CategoryLength` is not specified, then the chart will be plotted assuming that the series length is never longer than the first series specified. Hence, if the second series is longer than the first, the additional data points will not be plotted unless the `CategoryLength` parameter is used.

```
&CategoryLength=4
```

SumbyName: This parameter allows you to specify the individual sumby names for the data points. This parameter is used when you're entering data for Stack-type charts.

```
&SumbyName=deli;drinks;bread;cookies…
```

DataType: This parameter allows you to specify the type of data values that are used by the chart. The default type is `float`. The arguments that can be specified for this parameter are `int, integer, float, double, boolean, date, time,` and `timestamp`. The values are not case sensitive. `Date, time,` and `timestamp` types can only be used for Gantt charts.

```
&DataType=integer
```



Note

`DataType` is only specified for the value axis. You can use different types of data for the category axis, however they are all treated as strings. To use numeric, or date data without having it converted to string, design the chart using a database or data file.

8.2.2.2.3. Appearance Parameters

MainTitle: This allows you to specify the main title of the chart. By default, the chart has no title. However, if you're using a template file the titles will carry over, unless you specify different titles using these parameters.

```
&MainTitle=Sample+Chart
```

XTitle: This allows you to specify the X-axis title.

```
&XTitle=Time+of+Day
```

YTitle: This allows you to specify the Y-axis title.

```
&YTitle=Sales
```

ZTitle: This allows you to specify the Z-axis title.

```
&ZTitle=seriesA
```

Width: This allows you to specify the width of the generated image in pixels. If you do not specify a width, then the image will have the same width as the chart canvas.

```
&Width=400
```

Height:

This allows you to specify the height of the generated image in pixels. If you do not specify a height, then the image will have the same height as the chart canvas.

```
&Height=444
```

FitOnCanvas:

This is a Boolean flag that will cause the chart (not including labels, legend, etc.) to be shrunk and/or repositioned to fit on the canvas. By default, this is `false`.

```
&FitOnCanvas=true
```

ChartRelativePosition:

This allows you to specify the X and Y position of the lower left-hand corner of the chart to the lower left-hand corner of the chart canvas or image. The two arguments, numbers between 0 and 1 separated by a semi-colon, specify the X and Y positions respectively. `0;0` places the chart at the lower left-hand edge of the image, and `1;1` places the chart at the upper right-hand edge. By default, the chart relative position is `0.2;0.2`

```
&ChartRelativePosition=0.23;0.18
```

ChartRelativeSize:

This allows you to specify the size of the chart (relative width and height) relative to the chart canvas or image. The two arguments, numbers between `0` and `1` separated by a semi-colon, specify the width and height respectively. `1;1` causes the chart to be the same size as the image, and `0;0` will have no chart. By default, the chart relative size is `0.6;0.6`.

```
&ChartRelativeSize=0.63;0.7
```

ChartDimension:

This allows you to specify whether the chart is two-dimensional or three-dimensional. By default the chart is two-dimensional. The argument for this parameter is either `2D` or `3D`.

```
&ChartDimension=3D
```

ChartType:

This parameter allows you to select which chart type to use. By default, a Column Chart type is used. The arguments for this parameter can be: `Column`, `Bar`, `Line`, `Pie`, `Area`, `Overlay`, `Box`, `Radar`, `Dial`, `StackColumn`, `StackBar`, `StackArea`, `PercentColumn`, `Scatter`, `Bubble`, `Surface`, `HLCO`, `HiLow`, and `Gantt`.

```
&ChartType=StackArea
```



Note

For Overlay, Box, and Radar charts, you can only open existing charts through the `ChartPath` parameter, you will not be able to specify `ChartData` for these chart types.

3DDrawMode: This allows you to specify the way in which three-dimensional charts should be drawn. There are five options that are specified with an integer argument from `1` to `5`. `1` draws the chart as a wire frame without filling in the surfaces. `2` draws a border around the chart frame. `3` draws a border and applies Gouraud shading. `4` draws the chart with Gouraud shading, but doesn't draw a border. `5` draws the 3D chart as standard (without borders or specialized shading). (This is also the default).

```
&3DDrawMode=2
```

For more information about three-dimensional rendering options, see Section 5.2.4.4 - The Navigation Panel.

8.2.2.2.4. Export Parameters

DisplayType: This parameter allows you to select the image format in which you would like to render the chart. Options are `GIF, JPG, PNG, PDF,` and `SVG`. The default type is `JPG`.

```
&DisplayType=GIF
```



Note

If you're generating a parameter page for the chart (see Section 8.2.2.2.5 - Parameter Page Parameters) you can specify this argument more than once in an URL to give the user a drop-down list from which to select the image format. If you specify more than one format without generating a parameter page, the chart will be generated using the format specified in the first argument.

ExportOnServer: This is a Boolean flag that allows you to specify whether the chart should be written as a file on the server-side, or streamed back to the client. If this argument is set to `true`, then the `ExportPath` argument should also be used to indicate where the file should be written.

```
&ExportOnServer=true&ExportPath=C:\ERES\ExportFiles
\chart.gif
```

ExportPath: This allows you to specify a file path to export the chart to. This argument should always be used with the `ExportOnServer` argument. When a file path is specified, the chart will be exported in the selected format to the server-side.



Caution

When using ExportPath parameter and the file with the same name already exists, it is replaced without any warnings.

UseBuffer: This is a Boolean flag that allows you to specify if the image buffer is used.

```
&UseBuffer=true
```

ExportMethod: This allows you to specify which of the two available image generation methods the server should use. You can specify either `B` for the buffer method or `F` for the frame method. Each method can result in faster performance depending on the system, the size of the generated image, and the number of data points in the chart.

```
&ExportMethod=B
```

GIFTransparent: This is a Boolean flag that allows you to set the background of generated GIF images transparent. This parameter only works if the display type is set to `GIF`.

```
&GIFTransparent=true
```

PNGCompression: This parameter allows you to specify the compression of PNG images. The lower the compression is set, the faster the file will be generated. Lower compressions produce much larger files. There are four compression settings that can be specified using an integer argument. `0` is default compression. This level is the mid-level compression. `1` is fast compression. This generates the PNG file with minimal compression, causing fast file generation with a fairly large file. `2` is maximum compression. This generates a fully compressed PNG file, causing the file to be very small, but taking more generation time. `3` is zero compression. This generates a PNG file without any compression. It is generated very quickly, but has a very large file size.

```
&PNGCompression=1
```

JPEGQuality: This parameter allows you to specify the quality of generated JPEG images. As quality increases, so does the generated file size. The argument for this parameter is an integer between `1` and `99`. `1` is minimum quality, and `99` is the maximum. By default the quality is set to `99`.

```
&JPEGQuality=99
```

8.2.2.2.5. Parameter Page Parameters

ForHTMLParamPage: This is a Boolean flag that indicates whether or not to return a parameter page instead of a chart. This option only works it the chart contains query parameters. If this argument is set to `true` then the URL will return the parameter page, and ignore any parameter values passed into the URL. By default this argument will be `false`.

```
&ForHTMLParamPage=true
```

ParamPageCssStyle: This allows you to specify a CSS file to format the generated parameter page.

```
&ParamPageCssStyle=http://machinename:port/files/
ParamPage.css
```



Note

You can use both absolute and relative paths to the CSS file.

ParamPageTitle: This parameter allows you to specify a title for the parameter page.

```
&ParamPageTitle=Select+Chart+Options
```

ParamPageTitleFontName: This parameter allows you to specify the font to use for the parameter page title.

```
&ParamPageTitleFontName=Arial
```

ParamPageTitleFontSize: This parameter allows you to specify the font size for the Parameter page title. The size is specified as HTML size (i.e. `<font size=3>`) rather than point size. It ranges from `1` to `7`, where `1` is the smallest and `7` is the largest.

```
&ParamPageTitleFontSize=5
```

ParamPageTitleFontStyle: This allows you to apply a font style to the text of the parameter page title. The style is specified as an integer from `0 - 3`. The numbers indicate the following styles: `0` - Plain, `1` - Bold, `2` - Italic, `3` - Bold + Italic.

```
&ParamPageTitleFontStyle=1
```

ParamPageTitleFontColor: This parameter allows you to specify the font color for the parameter page title. The color is specified with a six digit hexadecimal number like in HTML.

```
&ParamPageTitleFontColor=404040
```

ParamPageTextFontName: This parameter allows you to specify the font to use for the prompt text in the parameter page.

```
&ParamPageTextFontName=Arial
```

ParamPageTextFontSize: This parameter allows you to specify the font size for the Parameter page prompt text. The size is specified as HTML size (i.e. `<font size=3>`) rather than point size. It ranges from `1` to `7`, where `1` is the smallest and `7` is the largest.

```
&ParamPageTextFontSize=3
```

ParamPageTextFontStyle: This allows you to apply a font style to the parameter page prompt text. The style is specified as an integer from `0 - 3`. The numbers indicate the following styles: `0` - Plain, `1` - Bold, `2` - Italic, `3` - Bold + Italic.

```
&ParamPageTextFontStyle=0
```

ParamPageTextFontColor: This parameter allows you to specify the font color for the parameter page prompt text. The color is specified with a six digit hexadecimal number like in HTML.

```
&ParamPageTextFontColor=000000
```

ParamPageBorderThickness: This parameter allows you to specify the thickness of the border drawn around the prompts in the parameter page. The thickness is specified in pixels. Specifying `0` will draw no border around the prompts.

```
&ParamPageBorderThickness=1
```

ParamPageBorderColor: This parameter allows you to specify the color for the border drawn around the prompts in the parameter page. The color is specified with a six digit hexadecimal number like in HTML.

```
&ParamPageBorderColor=C6C6C6
```

ParamPageTableBackgroundColor: This parameter allows you to specify the background color for the table containing the prompts in the parameter page. The color is specified with a six digit hexadecimal number like in HTML.

```
&ParamPageTableBackgroundColor=FFFFFF
```

ParamPageAlignment: This parameter allows you to specify the alignment in the page for the table containing the prompts. Options are `LEFT`, `RIGHT`, and `CENTER`. The default alignment is `CENTER`.

```
&ParamPageAlignment=LEFT
```

ParamPageLayout: This parameter allows you to specify the layout of the parameter page. You can either align the parameter prompts vertically (default) or horizontally. The options for this parameter are `HORIZONTAL` and `VERTICAL`.

```
&ParamPageLayout=HORIZONTAL
```

ParamPageResetEnabled: This is a Boolean flag that allows you to show/hide the reset button for the parameter form. By default the button is shown.

```
&ParamPageResetEnabled=false
```

ParamPageOuterBorder: This is a Boolean flag that allows you to draw only the outer border for the parameter prompt table. This parameter will have no effect if the border thickness is set to zero.

```
&ParamPageOuterBorder=true
```

8.2.2.3. Examples

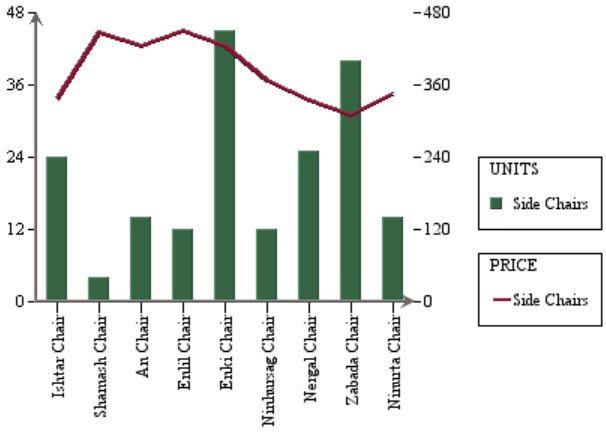
The following examples are sample image URLs that use one or more of the parameters listed in the preceding section. To use these, simply change `machinename` to your host name, `port` to the server port that you are using, and `context` to your servlet context. Note also that these examples all use the `EmbedImageTag` parameter to return a whole Web page containing the chart. If you want to use these examples in an existing Web page, set this parameter to `false`.

These URL examples are also included in a file named `ImageURL.txt` under `help/examples/URL`.

Example 1: Query data from a database, and plot a column chart

This example uses the Woodview HSQL database that is included with the ERES installation under `help/examples/DataSources/database`. To run this example you will need to have the HSQL JDBC driver in your classpath.

```
http://machinename:port/context/LookupServlet?URLTYPE=FORCHART
&DBSourceInfo=jdbc:hsqldb:help/examples/DataSources/database/woodview;
org.hsqldb.jdbcDriver;sa;+;select+c.categoryname+as+category,+
p.productname+as+product,+p.unitsinstock+as+units,+p.unitprice+
as+price+from+categories+c,+products+p+where+c.categoryid+equal+
p.categoryid+and+c.categoryid+equal+'SIC'&ColumnMapping=0;1;-1;2;3
&ChartRelativePosition=0.1;0.1&Width=420&Height=320
&EmbedImageTag=true&DisplayType=GIF
```

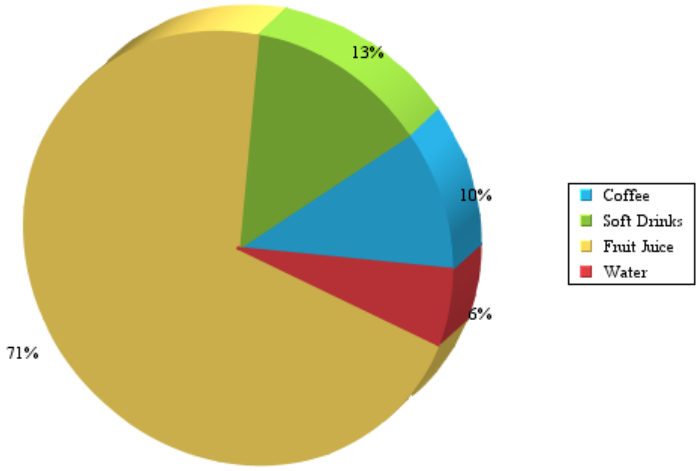


The query retrieves four columns from the database. The `UnitPrice` column is plotted as a sub-value which is the line. Note that the third argument in the `ColumnMapping` parameter is a `-1`. This is because column charts have no Sumby mapping.

Example 2: Draw data from a text file, and plot a 3D pie chart

This example draws data from the `sample.dat` file which is located in the `help/examples/DataSources/text` directory under the server root. You may need to modify the URL for the text file depending on how you have ERES deployed.

```
http://machinename:port/context/LookupServlet?URLTYPE=FORCHART
&DataFilePath=http://machinename:port/ERES/help/examples/DataSources
/text/sample.dat&ColumnMapping=-1;1;-1;4;-1&ChartType=Pie
&ChartDimension=3D&3DDrawMode=2&ChartRelativePosition=0.1;0.2;
&ChartRelativeSize=0.7;0.7&Width=420&Height=320
&EmbedImageTag=true&DisplayType=GIF
```

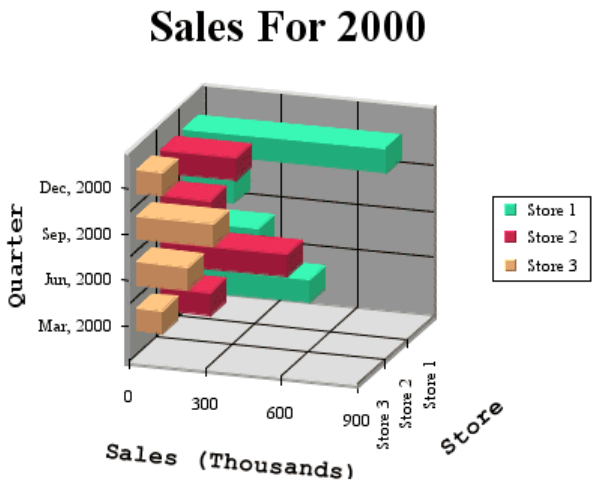


Since pie charts only have category, and value mappings, the column mapping uses `-1` for the series, sum by, and sub-value mappings. The `3DdrawMode` parameter is used to specify a border around the chart frame.

Example 3: Create a 3D Bar Chart, and add chart titles

This example uses the `ChartData` parameter to pass data within the URL string into the chart.

```
http://machinename:port/context/LookupServlet?URLTYPE=FORCHART
&ChartData=Store+1;500;300;200;800|Store+2;200;500;200;300|
Store+3;100;200;300;100&CategoryName=Mar,+2000;Jun,+2000;Sep,+2000;
Dec,+2000&DataType=integer&ChartType=Bar&ChartDimension=3D
&ChartRelativePosition=0.3;0.22&ChartRelativeSize=0.50;0.50&Width=420
&Height=320&MainTitle=Sales+For+2000&YTitle=Quarter
&XTitle=Sales+(Thousands)&ZTitle=Store&EmbedImageTag=true
&DisplayType=PNG
```



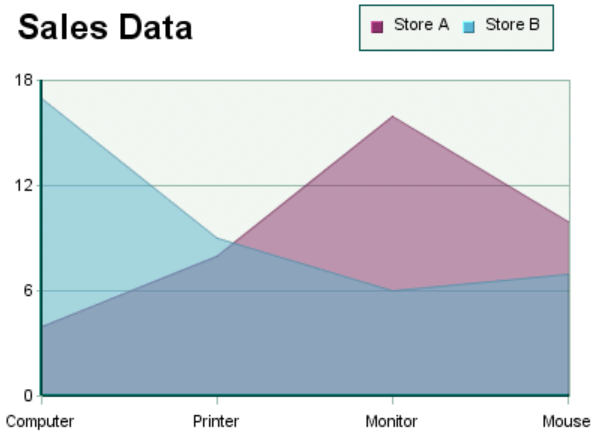
The chart is a 3D chart, which by default plots the series on the Z axis. X, Y, and Z axis labels have been added, as well as the chart title.

Example 4: Use a template file to apply appearance properties to a new dataset

This example uses the `Example4.tpl` file which is included in the `help/examples/URL` directory of the ERES installation. To run this example, you will need to add this file in the Organizer (see Section 3.1.4.1 - Adding and Modifying Files).

The URL includes the USERNAME and PASS parameters. You can either supply these parameters, or remove them. If they're removed, you will be re-directed to the login page before viewing the chart.

```
http://machinename:port/context/LookupServlet?URLTYPE=FORCHART
&USERNAME=username&PASS=password&TemplatePath=http://machinename:port
/ERES/help/examples/URL/Example4.tpl&ChartData=Store+A;4;8;16;10|
Store+B;17;9;6;7&DataType=Integer&CategoryName=Computer;Printer;
Monitor;Mouse&ChartType=Area&EmbedImageTag=true&DisplayType=PNG
```



As you can see, the template file has applied all of the appearance properties, including many that cannot be controlled using an Image URL. You main need to modify the URL for the template file depending on how you have ERES deployed.

Image URL parameters can only control some of the chart properties, and most of the properties are displayed using default values. For greater control over the appearance of the chart, you should create a chart or template file using the Chart Designer, and then call or apply that file with an image URL.

8.3. Report URLs

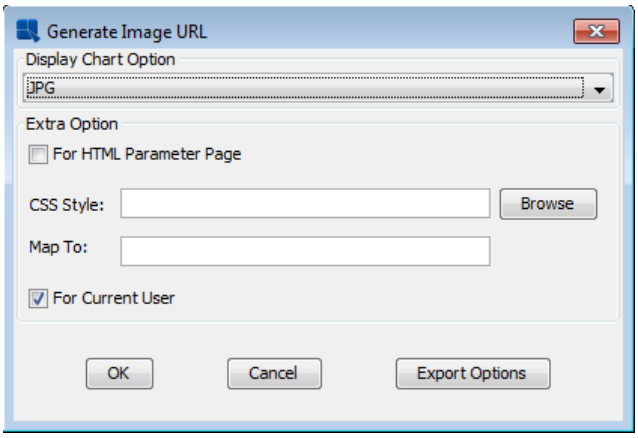
In addition to the menu and the API, EspressReport ES also allows you to deploy reports using URLs. Report URLs, like image URLs, allow reports to be called via an http call to the server. However unlike chart images generated by image URLs, report URLs cannot embed in an `<img src>` or other HTML tag. Report URLs are whole URLs that are run by pointing the web browser to the URL address, or by specifying a hyperlink to the report URL in another Web page.

Like Image URLs, report URLs can be used to generate reports from scratch, or to run existing report templates. However due to the number of available controls in reports, URL parameters are generally related to data input, and exporting.

8.3.1. Generating URLs in Organizer

The easiest way to generate report URLs is to have the Organizer do it for you. To generate a report URL in the Organizer, first select the report template (`.pak, .rpt` or `.xml`) that you would like to use. Then you can select *Generate Report URL* from the Publish menu, Click the *Generate URL* button on the toolbar, or right-click and

select *Generate Report URL* from the pop-up menu. A dialog will appear prompting you to specify some of the attributes for the generated report.



Report URL Options Dialog

The first option allows you to select whether to use the optimized memory exporting. This feature is useful for reports that use a large amount of data. Rather then fetching all of the data from the source at once. This feature will only export 5000 records at a time. This allows reports that use a lot of data to be run without running out of system memory. This feature is only applicable if the report draws its data from a database.

The second option allows you to select the format in which you would like to generate the report. Available formats include: HTLM, DHTML, PDF, Excel (XLS), Excel 2007 (XLSX), CSV, text, and rich text.

The third option allows you to select single page or multi page export. This option is only applicable to HTML or DHTML reports. Selecting single page will generate the entire report within a single html file. Selecting multi page will generate a new html file for each page of the report.

8.3.1.1. URLs for Parameterized Reports

If your report is parameterized, then there are two ways in which the report can be run. You can either pass parameters directly into the uRL as arguments, or have the URL generate a parameter prompt. If the report contains parameters, then several options will appear in the report URL options dialog as pictured above. To generate an HTML parameter prompt check the *For HTML Parameter Page* option.

If you select to generate a parameter prompt, you can also specify a CSS file to use to format the prompt. The dialog has options allowing you to specify the file and http path to the CSS file. (An example of a parameter page CSS file is available under `<ERESInstallDir/help/examples/URL>`).

If you select not to generate a parameter prompt, then a dialog will open asking you to select parameter values which will be added as arguments to the generated URL.

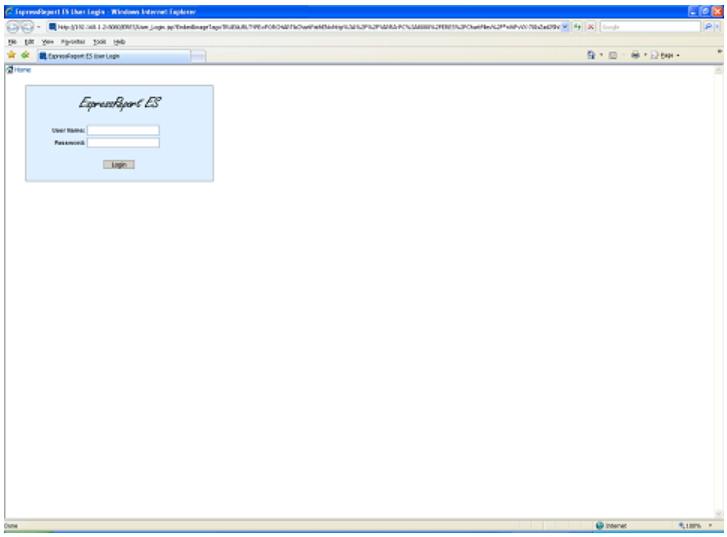
8.3.1.2. User Options

Report URLs can have an username and password passed in as arguments in the URL. If the administrator is generating the URL then an option will appear in the report URL options dialog as pictured above that allows you to select the user for which the URL is being generated. Clicking the *Select* button will bring up a list of defined users.

The URL can also be configured to read the username and password from the session of the URL is going to be run in a JSP or servlet environment. If this option is enabled the URL will try to read the username and password from session parameters `USERNAME` and `PASS`.

If an user other than the administrator generates the URL, then an option named *For Current User* will appear. Checking this option will embed the current user in the URL.

In either case if an user is selected for the URL the encoded username and password are added as arguments to the URL. If no username or password is supplied in the URL then when it is run, the user will be re-directed to a login page before viewing the report.

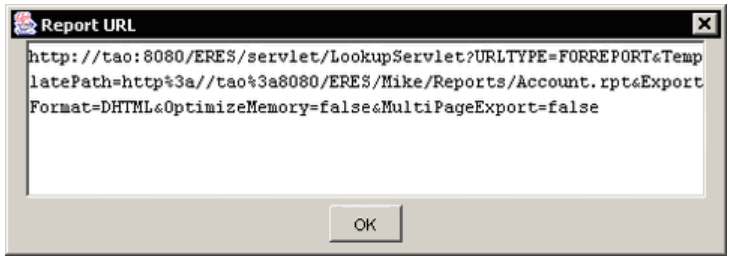


User Login Dialog

By default, the URL will redirect to the login page under the ERES context (i.e. `http://machinename:port/ERES`). If you have deployed ERES in a different context, you can use the `JSPCONTEXT` parameter to pass in a different location.

8.3.1.3. Running Report URLs

Once you have finished specifying options for the report URL a new dialog will open containing the generated URL string. If you're running Organizer as an applet, then the URL will load in a new browser window (allowing you to copy and paste).



Generated URL

When run, URLs will generate the report in the specified format and stream it back to the browser. For some formats like PDF or Excel, you will need to have the appropriate plug-in installed on the client in order for the report to display correctly.

8.3.1.4. Running Reports with Encrypted Data

To run a report contains encrypted data, you need to do two things in order to view the data.

1. You need to create an XML file that gives the database URL, database driver, name of column to be decrypted, and the function to be applied when the data is being retrieved.

2. You need to include the `ReplaceColumnInfoList` option in the command line, i.e. `servercommand.txt` before you start ERES server.

8.3.2. Writing Report URLs

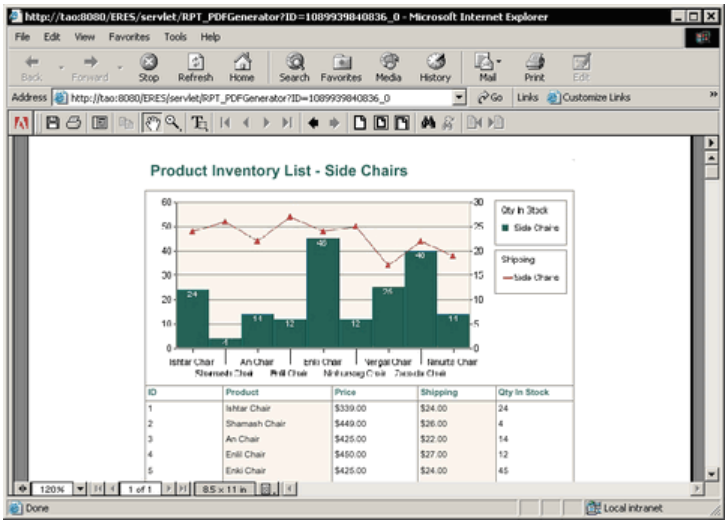
Rather than having the Organizer generate the report URL, you can write your own. You can run an existing report template, or you can pass a new data source to a report template, or even create a new report (un-formatted). You can control various report data and output properties using URL parameters. You can also modify an URL generated by the Organizer.

8.3.2.1. URL Syntax

The syntax for report URLs is fairly simple. Every URL begins with a call to the Server: `http://machinename:port/context/LookupServlet?`. Following the question mark users need to specify whether the URL will return a chart or report using the `URLTYPE` parameter. Following the `URLTYPE` users can add parameters to the URL. Parameters are separated by the "&" character.

For example the following URL will run the `ProductList.rpt` template in the `help/examples/URL` directory of the ERES installation and return it in PDF format.

```
http://machinename:port/context/LookupServlet?URLTYPE=FORREPORT
&USERNAME=username&PASS=password&TemplatePath=http://machinename:port
/ERES/help/examples/URL/ProductList.rpt&ExportFormat=PDF
```



Generated Report

As you can see the URL streams a PDF file back to the browser.

To run this example, you will need to add this file in the Organizer (see Section 3.1.4.1 - Adding and Modifying Files). The URL includes the `USERNAME` and `PASS` parameters. You can either supply these parameters, or remove them. If they're removed, you will be re-directed to the login page before viewing the report. You may also need to modify the URL for the template file depending on how you have deployed ERES.

In order for an URL to run, you must at minimum specify a report template, or specify data source - either a file or database. Also, any input parameter string should use the "+" character in place of a space. An actual space character may disrupt the creation of the report.

8.3.2.2. URL Parameters

URL parameters allow you to specify which report template to use, as well as data source information, and output format appearance such as PDF or HTML.

8.3.2.2.1. Login Parameters

Login parameters allow you to pass username and password into report URLs. A login is required if the URL references a template file (using the `TemplatePath` parameters). If login information is not specified in the URL then the URL will re-direct to a login page before showing the report.

USERNAME: This allows you to specify an username. If you generate an URL in Organizer and select to use an user, the username will be encoded in the generated URL. You can also specify the username as plain text.

```
&USERNAME=user
```

PASS: This allows you to specify the user password. If you generate an URL in Organizer and select to use an user, the password will be encoded in the generated URL. The password can also be specified using plain text.

```
&PASS=password
```

USESESSION: This allows you to have the server read the username and password from the session. If this option is enabled the information will be retrieved from session parameters `USERNAME` and `PASS`. This option is only available if you will be deploying the URL within a servlet or JSP application.

```
&USESESSION=true
```

JSPCONTEXT: This allows you to specify the location (context) of the ERES installation. This location is used when the URL re-directs to the login page. The default value is `/ERES`, and the URL will work without setting this parameter in most deployments. You only need to specify this parameter if ERES is deployed in a context of a different name

```
&JSPCONTEXT="/MyReports"
```

8.3.2.2.2. Data Input Parameters

TemplatePath: This allows you to specify the report template you would like to run. A report template is saved with its data source information, so if a different data source is not specified, it will retrieve data from the source with which it was designed.

```
&TemplatePath=http://machinename:port/ReportFiles/
reportfile.rpt
```



Note
You can use both absolute and relative paths for the template file location.

DataFilePath: This allows you to specify either a text or XML file from which you would like to retrieve data for the report. For more on data file specifications, please see Section 4.1.4 - Data from Text Files.

```
&DataFilePath=http://machinename:port/reportdata.txt
```



Note
You can use both absolute and relative paths for the data file location. Also, only XML files in the Quadbase format can be specified (The format generated when you export XML data from EspressReport). Other XML formats will not work in Report or Image URLs.

DataSourceClass: This allows you to specify a class file that will retrieve data to plot the chart.

```
&DataSourceClass=package.class
```

DBSourceInfo: This allows you to specify a database and SQL query to be used to retrieve report data. The parameter takes five input values: database URL, database driver, username, password, and query. Each input value is separated by a semi-colon. For more on connecting to a database, please see Section 4.1.2 - Data from a Database.

```
&DBSourceInfo=url;driver;username;password;query
```

If you do not wish to enter an username or password use the space character "+" instead. Do not drop the input value. The parameter will not run with insufficient inputs. In your SQL statement be sure to substitute the "+" character for any spaces. Also, substitute the word `equal` or `equals` for the "=" sign.

QueryParamName: This parameter is used if the report you're using contains a parameterized query. `QueryParamName` specifies the name of the query parameter you will be supplying a value for. (Parameter names are specified when you create the query. For more on this, see Section 4.1.2.2.2 - Parameterized Queries) It is always followed by the `QueryParamSize`, and the `QueryParamValue` parameter.

```
&QueryParamName=customerID&QueryParamSize=1&QueryParamValue=4
```

QueryParamSize: This parameter is used to specify the number of values that will be passed into a particular report parameter. As detailed in Section 4.1.2.2.2.1 - Multi-Value Parameters, some queries can have multi-value parameters. This URL parameter allows you to specify if the parameter (indicated by the previous `QueryParamName` parameter) takes multiple values. This parameter is always followed by one or more `QueryParamValue` arguments for each of the paramater values for a particular parameter. If you do not specify this parameter then it will assume that the number of values to be passed in is one.

QueryParamValue: This parameter is used if the report you're using contains a parameterized query. `QueryParamValue` specifies a value that you would like to pass to the query parameter specified by the `QueryParamName` argument. For multi-value parameters a separate `QueryParamValue` argument must be supplied for each distinct parameter value that you're passing in. For more on creating parameterized queries, please see Section 4.1.2.2.2 - Parameterized Queries.

FormulaParamName: This parameter is used if the report you're running contains a formula parameter. `FormulaParamName` specifies the name of the formula parameter for which you will be supplying a value. (Parameter names are specified when you write the formula, for more on this, see Section 5.1.6.2.6 - Formula Parameters). It is always followed by the `FormulaParamValue` parameter.

```
&FormulaParamName=UserName&FormulaParamValue=Jeff
```

FormulaParamValue: This parameter is used if the report you're running contains a formula parameter. `FormulaParamValue` specifies a value that you would like to pass to the formula parameter indicated by the preceeding `FormulaParamName` argument. For more on formula parameters, see Section 5.1.6.2.6 - Formula Parameters.

ColumnMapping: If you have specified a new data source for the report, this argument allows you to map columns from your data file, or database query result set to the report. Column mapping is specified using index values. Index values are assigned based on the order

they appear in your data file from left to right, or the order in which they are selected in your SQL statement. Index values start with `0`, so the first column is `0` and the second is `1`, etc. To perform column mapping, Specify the index values (separated by a semi-colon) from the data source that you would like to include in the report (in the order you want them to appear in the report). You can also provide names for the column in your data source with this parameter as well. This is optional, but setting names allows you to refer to columns by name for the other data mapping parameters. To add column names place them after the column index separated by a comma.

```
&ColumnMapping=3,Category;1,Product;4,Price;5,Sales;10,Employee
```

This example will generate a report that has the fourth, second, fifth, sixth, and eleventh columns from the data source as the first through fifth columns in the report respectively.



Note

Column order appears from left to right in the generated report.

RowBreak:

This argument allows you to set row break columns for a summary break or crosstab report. Columns are specified by index number separated by semi-colons. You can also refer to column names that you established in the `ColumnMapping` parameter. For more information about report mapping, please see Section 5.1.2 - Report Types and Data Mapping.

```
&RowBreak=0;1 or &RowBreak=Category;Product
```



Note

The numbers in this argument refer to the report columns and not the data source columns. Hence the above example sets the first two columns of the report to be row breaks. These columns would be the fourth and second columns from the data source (using the `ColumnMapping` example).

ColumnBreak:

This argument allows you to set column break columns for a crosstab report. Columns are specified by index number, and separated by semi-colons. You can also refer to column names that you established in the `ColumnMapping` parameter. For more information about report mapping, please see Section 5.1.2 - Report Types and Data Mapping.

```
&ColumnBreak=4 or &ColumnBreak=Employee
```



Note

The numbers in this argument refer to the report columns and not the data source columns.

ColumnBreakValue:

This argument allows you to set column break value columns for a crosstab report. Columns are specified by index number, and separated by semi-colons. You can also refer to column names that you established in the `ColumnMapping` parameter. For

more information about report mapping, please see Section 5.1.2 - Report Types and Data Mapping.

```
&ColumnBreakValue=3 or &ColumnBreakValue=Sales
```



Note

The numbers in this argument refer to the report columns and not the data source columns.

Aggregation:

This argument allows you to specify column aggregation for non-row break columns in summary break reports, and column break value columns in crosstab reports. Columns are specified by index number followed by a comma and then the aggregation type. You can also refer to column names that you established in the `ColumnMapping` parameter. Available aggregations are `SUM`, `COUNT`, `SUMSQUARE`, `AVG`, `VARIANCE`, `MAX`, `FIRST`, `STDDEV`, `MIN`, `LAST`, `COUNTDISTINCT`.

```
&Aggregation=2,AVG;3,SUM or
&Aggregation=Price,AVG;Sales,SUM
```

By default, non-row break columns in summary break reports have no aggregation, and column break value columns for crosstab reports use `SUM`. For more information about report mapping, please see Section 5.1.2 - Report Types and Data Mapping.



Note

The numbers in this argument refer to the report columns and not the data source columns.

PrimaryKey:

This argument allows you to set the primary key column for a master & details report. It only accepts one parameter - the column index of the primary key column. You can also refer to column names that you established in the `ColumnMapping` parameter.

```
&PrimaryKey=1 or &PrimaryKey=Product
```



Note

The number in this argument refer to the report columns and not the data source columns.

MasterField:

This argument allows you to specify which columns should be included in the master field for a master & details report. Columns are specified by index number, and separated by semi-colons. You can also refer to column names that you established in the `ColumnMapping` parameter.

```
&MasterField=0;2 or &MasterField=Product;Price
```

The numbers in this argument refer to the report columns and not the data source columns.

ReportType: This argument allows you to specify which report type to use for the generated report. Available types are "Columnar", "CrossTab", "MailingLabels", "MasterDetails", and "Summary". Report types are discussed in Section 5.1.2 - Report Types and Data Mapping. The default report type is "Columnar".

```
&ReportType=MasterDetails
```

8.3.2.2.3. Export Parameters

Export Format: This argument allows you to specify the format in which you would like to display the report. Available formats are `DHTML`, `PDF`, `TXT`, `CSV`, `XLS` and `XLSX`. The default format is `DHTML`.

```
&ExportFormat=PDF
```



Note

If you're generating a parameter page for the report (see Section 8.3.2.2.4 - Parameter Page Parameters) you can specify this argument more than once in an URL to give the user a drop-down list from which to select the export format. If you specify more than one format without generating a parameter page, the report will be generated using the format specified in the first argument.

ExportOnServer: This is a Boolean flag that allows you to specify whether the report should be written as a file on the server-side, or streamed back to the client. If this argument is set to `true`, then the `ExportPath` argument should also be used to indicate where the file should be written.

```
&ExportOnServer=true&ExportPath=C:\ERES\ExportFiles
\report.pdf
```

ExportPath: This allows you to specify a file path to export the report to. This argument should always be used with the `ExportOnServer` argument. When a file path is specified, the report will be exported in the selected format to the server-side.

MultiPageExport: This is a Boolean flag that allows you to specify whether or not to use multi page exporting for HTML and DHTML export formats. By default multi page exporting is not used.

```
&MultiPageExport=false
```



Note

This parameter will not work if you use a different data source than the one specified in the report template.

OptimizeMemory: This is a Boolean flag that allows you to specify whether or not to use memory optimized exporting when generating the report. This feature is explained in Section 8.3.1 - Generating URLs in Organizer. By default memory optimized exporting is not used.

```
&OptimizeMemory=true
```

8.3.2.2.4. Parameter Page Parameters

ForHTMLParamPage: This is a Boolean flag that indicates whether or not to return a parameter page instead of a report. This option only works it the Report contains query or formula parameters. If this argument is set to `true` then the URL will return the parameter page, and ignore any parameter values passed into the URL. By default this argument will be `false`.

```
&ForHTMLParamPage=true
```

ParamPageCssStyle: This allows you to specify a CSS file to format the generated parameter page.

```
&ParamPageCssStyle=http://machinename:port/files/
ParamPage.css
```



Note

You can use both absolute and relative paths to the CSS file.

ParamPageTitle: This parameter allows you to specify a title for the parameter page.

```
&ParamPageTitle=Select+Report+Options
```

ParamPageTitleFontName: This parameter allows you to specify the font to use for the parameter page title.

```
&ParamPageTitleFontName=Arial
```

ParamPageTitleFontSize: This parameter allows you to specify the font size for the Parameter page title. The size is specified as HTML size (i.e. `<font size=3>`) rather than point size.

```
&ParamPageTitleFontSize=5
```

ParamPageTitleFontStyle: This allows you to apply a font style to the text of the parameter page title. The style is specified as an integer from `0 - 3`. The numbers indicate the following styles: `0` - Plain, `1` - Bold, `2` - Italic, `3` - Bold + Italic.

```
&ParamPageTitleFontStyle=1
```

ParamPageTitleFontColor: This parameter allows you to specify the font color for the parameter page title. The color is specified with a six digit hexidecimal number like in HTML.

```
&ParamPageTitleFontColor=404040
```

ParamPageTextFontName: This parameter allows you to specify the font to use for the prompt text in the parameter page.

```
&ParamPageTextFontName=Dialog
```

ParamPageTextFontSize: This parameter allows you to specify the font size for the Parameter page prompt text. The size is specified as HTML size (i.e. `<font size=3>`) rather than point size.

```
&ParamPageTextFontSize=3
```

ParamPageTextFontStyle: This allows you to apply a font style to the parameter page prompt text. The style is specified as an integer from `0 - 3`. The numbers indicate the following styles: `0` - Plain, `1` - Bold, `2` - Italic, `3` - Bold + Italic.

```
&ParamPageTextFontStyle=0
```

ParamPageTextFontColor: This parameter allows you to specify the font color for the parameter page prompt text. The color is specified with a six digit hexidecimal number like in HTML.

```
&ParamPageTextFontColor=000000
```

ParamPageBorderThickness: This parameter allows you to specify the thickness of the border drawn around the prompts in the parameter page. The thickness is specified in pixels. Specifying `0` will draw no border around the prompts.

```
&ParamPageBorderThickness=1
```

ParamPageBorderColor: This parameter allows you to specify the color for the border drawn around the prompts in the parameter page. The color is specified with a six digit hexidecimal number like in HTML.

```
&ParamPageBorderColor=C6C6C6
```

ParamPageTableBackground-Color: This parameter allows you to specify the background color for the table containing the prompts in the parameter page. The color is specified with a six digit hexidecimal number linke in HTML.

```
&ParamPageTableBackgroundColor=FFFFFF
```

ParamPageAlignment: This parameter allows you to specify the alignment in the page for the table containing the prompts. Options are `LEFT, RIGHT,` and `CENTER`. The default alignment is `CENTER`.

```
&ParamPageAlignment=LEFT
```

ParamPageLayout: This parameter allows you to specify the layout of the parameter page. You can either align the parameter prompts vertically (default) or horizontally. The options for this parameter are `HORIZONTAL` and `VERTICAL`.

```
&ParamPageLayout=HORIZONTAL
```

ParamPageResetEnabled: This is a Boolean flag that allows you to show/hide the reset button for the parameter form. By default the button is shown.

```
&ParamPageResetEnabled=false
```

ParamPageOuterBorder: This is a Boolean flag that allows you to draw only the outer border for the parameter prompt table. This parameter will have no effect if the border thickness is set to zero.

```
&ParamPageOuterBorder=true
```

8.3.2.3. Examples

The following examples are sample report URLs that use one or more of the parameters listed in the preceding section. To use these, simply change `machinename` to your host name, `port` to the server port that you are using, and `context` to your servlet context. All of these examples use templates that are in the `help/examples/URL` directory, so you'll need to make sure the templates have been added to the Organizer and that you have permission to view them.

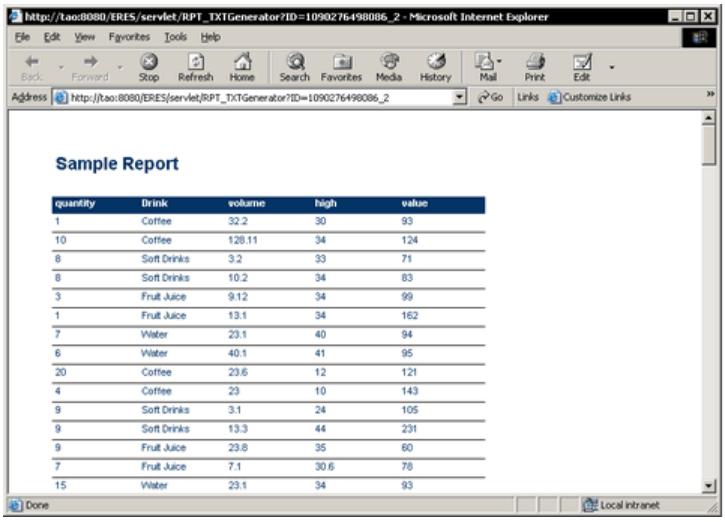
The URLs include the `USERNAME` and `PASS` parameters. You can either supply these parameters, or remove them. If they're removed, you will be re-directed to the login page before viewing the reports.

These URL examples are also included in a file named `ReportURL.txt` under `help/examples/URL`.

Example 1: Draw data from a text file

This example draws data from the `Sample.dat` file under the `help/examples/DataSources/text` directory, and creates a simple columnar layout. To run this example, you will need to add `help/examples/URL/Example1.rpt` to the Organizer (see Section 3.1.4.1 - Adding and Modifying Files).

```
http://machinename:port/context/LookupServlet?URLTYPE=FORREPORT
&USERNAME=username&PASS=password&DataFilePath=http://machinename:port
/ERES/help/examples/DataSources/text/sample.dat&ColumnMapping=3;1;4;
5;10&TemplatePath=http://machinename:port/ERES/help/examples/URL
/Example1.rpt&ExportFormat=DHTML
```



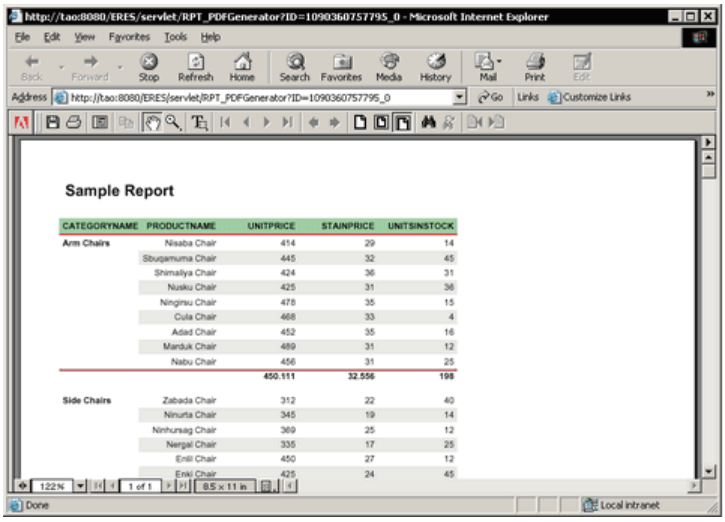
Generated Report

This report draws data from fourth, second, fifth, sixth, and eleventh columns of the `sample.dat` file. You may need to modify the URL for the data and template files depending on how you have ERES deployed.

Example 2: Draw data from a database and generate a summary break report

This example uses the Woodview HSQL database that is included with the ERES installation under `help/examples/DataSources/database` also names the columns in the data mapping parameters. To run this example you will need to have the HSQL JDBC driver in your classpath. To run this example, you will need to add `help/examples/URL/Example2.rpt` to the Organizer (see Section 3.1.4.1 - Adding and Modifying Files).

```
http://machinename:port/context/LookupServlet?URLTYPE=FORREPORT
&USERNAME=username&PASS=password&DBSourceInfo=jdbc:hsqldb:help
/examples/DataSources/database/woodview;org.hsqldb.jdbcDriver;sa;+;
select+c.categoryname,+p.productname,+p.unitprice,+p.stainprice,+
p.unitsinstock+from+categories+c,+products+p+where+c.categoryid+
equal+p.categoryid+and+(c.categoryname+equal+'Side+Chairs'+or+
c.categoryname+equal+'Arm+Chairs')&ColumnMapping=0,Category;1,Product;
2,Price;3,StainPrice;4,UnitsInStock&RowBreak=Category
&Aggregation=Product,NONE;Price,AVG;StainPrice,AVG;UnitsInStock,SUM
&ReportType=Summary&TemplatePath=http://machinename:port/ERES
/help/examples/URL/Example2.rpt&ExportFormat=PDF
```



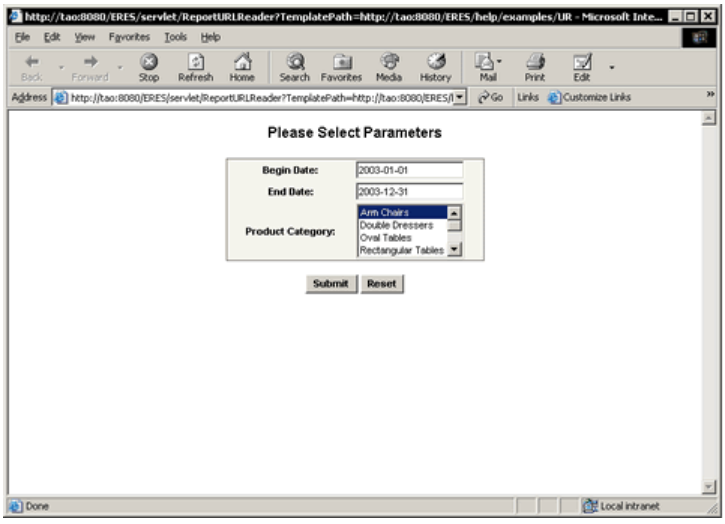
Generated Report

You may need to modify the URL for the template file depending on how you have ERES deployed.

Example 3: Open a parameterized report and generate an HTML parameter prompt

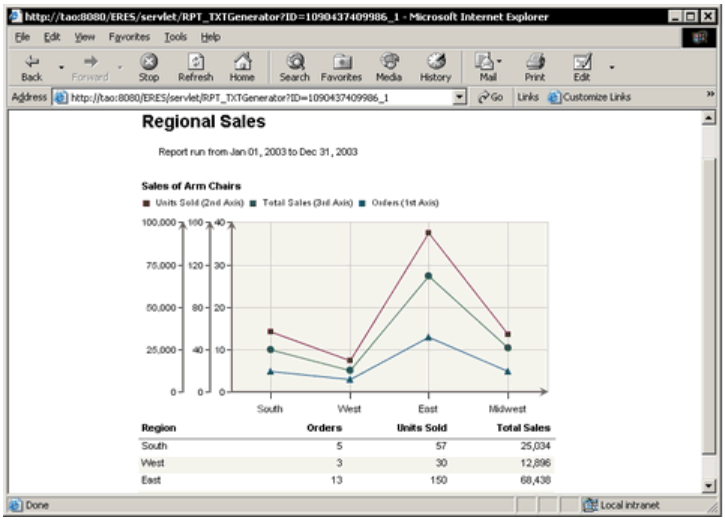
To run this example, you will need to add `help/examples/URL/RegionalSales.rpt` to the Organizer (see Section 3.1.4.1 - Adding and Modifying Files).

```
http://machinename:port/context/LookupServlet?URLTYPE=FORREPORT
&USERNAME=admin&PASS=admin&TemplatePath=http://machinename:port
/ERES/help/examples/URL/RegionalSales.rpt&ForHTMLParamPage=true
&ParamPageCssStyle=http://machinename:port/ERES/help/examples/URL
/ParamPage.css&ParamPageBorderThickness=1&ParamPageTitle=Please+
Select+Parameters&ParamPageOuterBorder=true
&ParamPageTableBackgroundColor=F7F7EF&ParamPageBorderColor=6C6464
&ExportFormat=DHTML
```



Generated Parameter Prompt

The URL first returns a parameter prompt based on the parameters specified in the URL and formatted with the `.css` file under `help/examples/URL`. Clicking *Submit* in this page returns the report with the specified parameters.



Generated Report

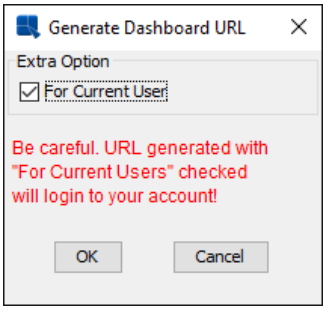
You may need to modify the URL for the template and CSS file depending on how you have ERES deployed.

8.4. Dashboard URLs

In addition to the menu and the API, ERES also allows you to deploy dashboards using URLs. Dashboard URLs allow dashboards to be called via http call to the server. Dashboard URLs are whole URLs that are run by pointing the web browser to the URL address, or by specifying a hyperlink to the dashboard URL in another Web page.

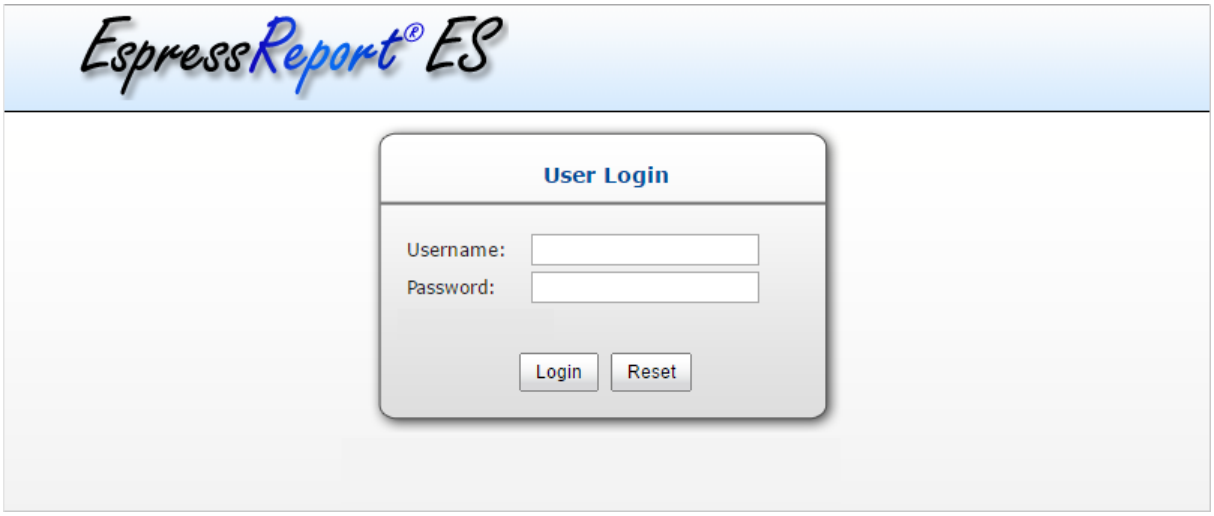
8.4.1. Generating URLs in Organizer

The easiest way to generate dashboard URLs is to have the Organizer do it for you. To generate a dashboard URL in the Organizer, first select the dashboard file (`.dsb`) that you want to use. Then you can select *Generate Dashboard URL* from the Publish menu, click the *Generate Dashboard URL* button on the toolbar, or right-click and select *Generate Dashboard URL* from the pop-up menu. A new dialog will appear allowing you to specify user option for the generated dashboard URL.



Generate Dashboard URL Dialog

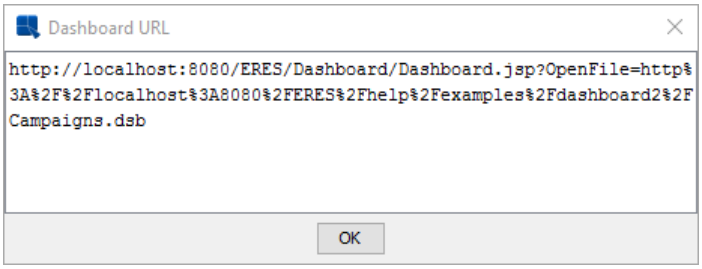
Dashboard URLs can have username and password passed in as arguments in the URL. If the *For Current User* option is checked, the current user will be embedded in the URL. If no username or password is supplied in the URL (*For Current User* option is unchecked) when it is run, the user will be re-directed to a login page before viewing the dashboard.



User Login Page

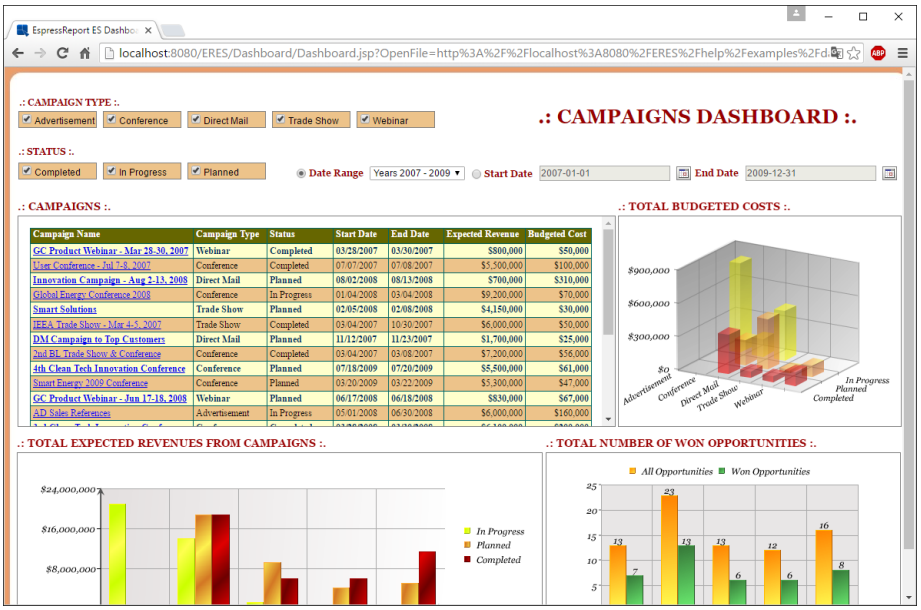
8.4.1.1. Running Dashboard URLs

Once you specify the *For Current User* option and click the *OK* button, a new window will open containing the generated URL string. From this window, you can copy the dashboard URL.



Generated Dashboard URL

In order to view a dashboard, just copy and paste the URL into the browser window.



View Dashboard URL

8.4.2. Writing Dashboard URLs

Rather than having the Organizer generate the dashboard URL, you can write your own by using the URL syntax below.

8.4.2.1. URL Syntax

The syntax for dashboard URLs is fairly simple. Every URL begins with a call to the Server: `http://<MACHINE_NAME>:<PORT>/<CONTEXT_ROOT>/Dashboard/Dashboard.jsp?`. The `MACHINE_NAME` is either the URL, machine name, or ip address of the server machine. The `PORT` is the port number for the application server, the default for Tomcat is 8080. The `CONTEXT_ROOT` is the context root for the ERES. By default, the context root is `ERES`, but this can be changed by administrator. Following the question mark, users need to specify either the `USERNAME` and `PASS` parameters, or the `USESESSION` parameter instead, as well as the dashboard file location.

The following URL will open the `Population.dsb` dashboard file for the specified user.

```
http://<MACHINE_NAME>:<port>/<CONTEXT_ROOT>//Dashboard/Dashboard.jsp?
USERNAME=qMFq-5dab27h&PASS=qMFq-5dab27h&OpenFile=http://
<MACHINE_NAME>:<port>/<CONTEXT_ROOT>/help/examples/Dashboard/Population.dsb
```

For example, if the server was set up on www.quadbase.com [http://www.quadbase.com] with port 8080 and the context root is `ERES`, then the dashboard URL will look like this.

```
http://www.quadbase.com:8080/ERES/Dashboard/Dashboard.jsp?
USERNAME=qMFq-5dab27h&PASS=qMFq-5dab27h&OpenFile=http://
www.quadbase.com:8080/ERES/help/examples/Dashboard/Population.dsb
```

8.4.2.2. URL Parameters

URL parameters allows you to specify the dashboard file you want to run as well as specify whether to use username and password that are stored in the browser session or not.

OpenFile: This allows you to specify the dashboard file that you want to run.

```
&OpenFile=http://machinename:port/ERES/DashboardFiles/
TestDashboard.dsb
```



Note

You can use both absolute and relative paths for the dashboard file location.

USERNAME: This option allows you to specify a username. If you generate an URL in Organizer and select to use an user, the username will be encoded in the generated URL. You can also specify the username as plain text.

```
&USERNAME=user
```

PASS: This option allows you to specify a user password. If you generate an URL in Organizer and select to use an user, the password will be encoded in the generated URL. The password can also be specified using plain text.

```
&PASS=password
```

USESESSION: This allows you to have the server read the username and password from the session. If this option is enabled, the information will be retrieved from session parameters `USERNAME` and `PASS`. This option is only available if you will be deploying the URL within a servlet or JSP application.

```
&USESESSION=true
```

8.5. Menu Page Listener

The menu page and URL publishing options provide a convenient deployment vehicle for reports and charts. Using these interfaces, users can easily deploy charts and reports to the Web without any coding. However, the interfaces themselves do not provide any significant run-time customization capability. Report and chart templates are run in a pretty much “as is” format.

For users that would like to provide some additional logic, or run-time customization to reports and charts, ERES provides the Menu Page Listener interface which is part of the ERES extension package. Using this method, users can implement listeners that will intercept the report or chart prior to export, and use the APIs to modify it.

8.5.1. ERES Listener Manager

The ERES Listener Manager is a user-implemented class that is used to manage several listeners that monitor events on the server, including the Menu Page Listener. In order to implement the Menu Page Listener, users must implement the Listener Manager. Below is a sample Listener Manager class:

```
package extensionClasses;

import quadbase.reportorganizer.ext.*;

public class MyEresListenerManager extends DefaultListenerManager {

	public MyEresListenerManager() {}

	public EresSchedulerListener getSchedulerListener() {

		return new MyEresSchedulerListener();

	}

	public MenuPageListener getMenuPageListener() {

		return new MyMenuPageListener();

	}

}
```

Users can impelement the `ERESListenerManager` interface or extend `DefaultListenerManager`. The above code implements two listeners, the Scheduler Listener and the Menu Page Listener. The Scheduler Listener takes effect when a schedule job executes.

8.5.2. Using the Menu Page Listener

The following code shows a sample implementation of the Menu Page Listener:

```
package extensionClasses;
```

```
import java.lang.*;
import java.awt.*;
import quadbase.reportorganizer.ext.*;
import quadbase.reportdesigner.ReportAPI.*;
import quadbase.reportdesigner.ReportElements.*;
import quadbase.reportdesigner.util.*;
import quadbase.ChartAPI.*;
import quadbase.reportdesigner.report.Formula;

public class MyMenuPageListener implements MenuPageListener {

      public QbReport modifyBeforeRun(QbReport report, String username) {

            System.out.println("Calling Menu Page Listener...");

            try {

                  // Create a new label to show user that is running the
 report
                  ReportCell userLabel = new ReportCell();
                  userLabel.setText("Report Run By: " + username);
                  userLabel.setFont(new Font("Arial", Font.PLAIN, 8));
                  userLabel.setAlign(IAlignConstants.ALIGN_RIGHT);
                  userLabel.setHeight(0.2);
                  userLabel.setWidth(1.6);
                  userLabel.setY(0);
                  userLabel.setX(5.90);

                  // Create a new formula to show the report run date/time
                  ReportCell runDate = new ReportCell();
                  Formula runDateFormula = new
 Formula("ReportRunDate", "\"Report Run Date: \" +
 printDateTime(getCurrentDateTime(), \"MMM dd, yyyy h:mm a\")");
                  report.addFormula(runDateFormula);
                  runDate.setFormulaObj(runDateFormula);
                  runDate.setFont(new Font("Arial", Font.PLAIN, 8));
                  runDate.setAlign(IAlignConstants.ALIGN_RIGHT);
                  runDate.setHeight(0.2);
                  runDate.setWidth(2.5);
                  runDate.setY(0.2);
                  runDate.setX(5.0);

                  // Add new cells to page header section
                  report.getPageHeader().addData(userLabel);
                  report.getPageHeader().addData(runDate);

                  // Adjust section height if necessary
                  if (report.getPageHeader().getHeight() < 0.4)
                        report.getPageHeader().setHeight(0.4);

            } catch (Exception ex) {

                  ex.printStackTrace();

            }

            return report;
```

```
	}

	public QbChart modifyBeforeRun(QbChart chart, String username) {

		System.out.println("Calling Menu Page Listener");

		try {

			// Define an array of chart colors
			Color gold = new Color(227,215,130);
			Color salmon = new Color(199,85,90);
			Color burgandy = new Color(125,18,66);
			Color slate = new Color(81,119,156);
			Color teal = new Color(130,203,217);
			Color blue = new Color(26,42,103);
			Color beige = new Color(188,204,177);
			Color[] chartColors = {gold, salmon, burgandy, slate,
 teal, blue, beige};

			// Apply new colors to the chart
			chart.gethDataPoints().setColors(chartColors);

		} catch (Exception ex) {

			ex.printStackTrace();

		}

		return chart;

	}

}
```

This example implements to `modifyBeforeRun` methods, one for charts and one for reports. With this code deployed, everytime a report is run, the username and the run time is added to the report header. Everytime a chart is run, the chart colors are modified to use the palette defined in the method. Note that this is only one example. Any report/chart API code can be used in these interfaces to modify the reports and charts prior to export.

8.6. Report Viewer

The Report Viewer is an applet that allows you to view a report dynamically through a Web browser. The Viewer can read report template files (`.rpt`, and `.xml`). The Report Viewer can be automatically embedded in an HTML page or manually placed in a Web page.

Inside Report Viewer, you go to different sections and pages of the report. There are built-in callback mechanisms that would allow users to click on a data element to jump to a related URL. Report Viewer also supports scheduled refresh (whereby a report's data is updated at regular intervals specified by the designer).

To create a stand-alone HTML page with the Viewer embedded, check the *Create HTML* option when saving a report in the Report Designer.

In addition you can embed the Viewer with your own Web page using the following syntax:

```
<applet codebase= ".." code =
 "quadbase.reportdesigner.ReportViewer.Viewer.class" width= 100% height=100%
 archive="ReportViewerWithChart.jar">
<PARAM name="filename" value="yourReport.pak">
```

```
</applet>
```

The parameter value of `filename` specifies the name of the file that contains the report. The filename value can also be specified by a URL; for example, you can prefix it by `http://` for accessing a remote data file. When viewed in the Report Viewer, the `.pak` file will connect to the data source specified, dynamically fetch the data, and generate the report.

8.6.1. The Report Viewer Parameters

Without using the Report Designer, you can also pass data via parameters to the Report Viewer applet. That is, you can use the Report Viewer to directly view a data file, or pass the data directly in the form of lines of data, in the HTML code. The following is a list of parameters:

Report Parameters

comm_protocol: The protocol to be used, to connect to the ERES Server

comm_url: The URL to connect to the ERES Server

servlet_context: The context in which the ERES Server is running

RefreshInterval: Specifies the interval in seconds to refresh the data (i.e., getting the data from the data source)

FileName: Specifies the name of the `.pak` file

ReportData: Specifies the report as a string

SourceDB: Specifies the database information

SourceFile: Specifies the data file information

SourceData: Specifies the data information

ReportType: Specifies the report type

Aggregation: Specifies the Aggregation Method

ColInfo: Sets the column mapping

RowBreak: Specifies the column(s) to be used as Row Breaks

ColumnBreak: Specifies the column to be used as Column Break

ColumnBreakValue: Specifies the column to be used as the values for the Column Break columns

PrimaryKey: Specifies the column to be used as Primary Key

MasterKey: Specifies the column(s) to be used as Master Key

RefreshData: If false, report cannot be refreshed. The default is `true`

ExportEnabled: If false, report cannot be exported. The default is `true`

Example: With the parameter `RefreshInterval` you can insert the following line:

```
<PARAM name="RefreshInterval" value="60">
```

In the above example, Report Viewer applet will fetch data from the specified data source (with the help of the ERES Server) and redraw the report every 60 seconds - all transparently. It is useful for accessing databases, which are updated frequently. Note that when specifying parameters that can have multiple values (or needs to be have multiple values), the separator is a space (“ ” without the double quotes). For example:

```
<PARAM name="RowBreak" value="0 1 2">
```

In the above example, the Row Break is set on Column `0`, Column `1` and Column `2` of the Column Mapping (i.e., the first three columns). The values for the parameter have a space (" " without the double quotes) separating them.

8.6.2. Specifying Data for Report Viewer

Using parameters, you can specify a data source for the Report Viewer in order to display different data with a report template.

8.6.2.1. Data From a Database

The following is a piece of HTML code that uses the Report Viewer to view a report using data extracted from a database.

```
<applet code="quadbase.reportdesigner.ReportViewer.Viewer.class" width=640
 height=480>

      <PARAM name="sourceDB" value="jdbc:odbc:DataSource,
 sun.jdbc.odbc.JdbcOdbcDriver, username, password, select * from products">
      <PARAM name="ColInfo" value="0 1 3">
      <PARAM name="ReportType" value="SummaryBreak">
      <PARAM name="RowBreak" value="0">
      <PARAM name="Aggregation" value="SUM">

</applet>
```

The arguments of **`ColInfo`** specify the mapping for the report.

The argument **`ReportType`** specifies the type of report to be displayed and it can be one of

- SimpleColumnar
- SummaryBreak
- CrossTab
- MasterDetails

The argument **`RowBreak`** specifies the column on which the row break is applied.

The argument **`Aggregation`** specifies the aggregation method applied and it can be one of:

- NONE
- SUM
- MAX
- MIN
- COUNT
- AVG
- FIRST
- LAST
- SUMSQUARE
- VARIANCE

- STDDEV
- COUNTDISTINCT

8.6.2.2. Data From a File

This HTML code generates a report using data from a data file.

```
<applet code = "quadbase.reportdesigner.ReportViewer.Viewer.class" width=640
 height=480>

      <PARAM name="sourceFile" value="http://.../test.dat">
      <PARAM name="ColInfo" value="1 0 2 3">
      <PARAM name="ReportType" value="SimpleColumnar">

</applet>
```

8.6.2.3. Data From an Argument

It is possible to have the Report Viewer read in data directly from the HTML file itself, rather than from a data file or database.

```
<applet code = "quadbase.reportdesigner.ReportViewer.Viewer.class" width=640
 height=480>

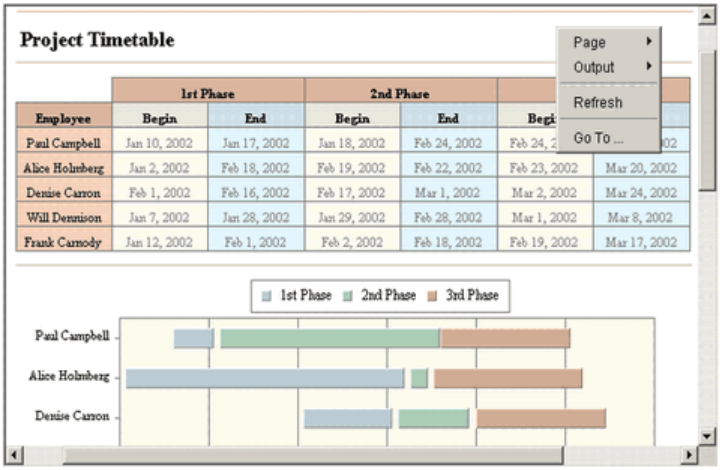
      <PARAM name="sourceData" value="int, string, int | value, name, vol |
 10, 'John', 20 | 3, 'Mary', 30 | 8, 'Kevin', 3 | 9, 'James', 22">
      <PARAM name="ColInfo" value="1 0 2">
      <PARAM name="ReportType" value="SimpleColumnar">

</applet>
```

The format for the parameter `SourceData` is the same as the data file format, except that each line is ended by a vertical bar "|".

8.6.3. Using Report Viewer

Using the Report Viewer, you can browse through the report using the Pop-Up Menu. Right-clicking on the report displays a Pop-Up Menu which can be used to traverse the report.



Report Viewer Window with Pop-Up Menu

The Pop-Up Menu can show the following options:

- **Back:** Go back to the previous report. This option is shown when a hyperlink is clicked on.
- **Section:** Contains various options to navigate the various sections

  First Section: Go to the first section

  Previous Section: Go to the previous section

  Next Section: Go to the next section

  Last Section: Go to the last section
- **Page:** Contains various options to navigate the various pages

  First Page: Go to the first page

  Previous Page: Go to the previous page

  Next Page: Go to the next page

  Last Page: Go to the last page
- **Output:** Contains various formats to output the report to

  Generate DHTML (Single Page): Export the report to DHTML onto a single file

  Generate DHTML (Multiple Pages): Export the report to DHTML onto multiple files, one file per page.

  Generate PDF: Export the report to PDF

  Generate CSV: Export the report to CSV

  Generate EXCEL (XLS): Export the report to Excel (`xls` file)

  Generate EXCEL 2007 (XLSX): Export the report to Excel 2007 (`xlsx` file)

  Generate TXT: Export the report to TXT

  Generate RTF: Export the report to RTF

  Generate XML (Data + Format): Export the complete report to XML

  Generate XML (Pure Data): Export the report data only to XML

  Print: Print the report
- **Refresh:** Refresh the data from the data source
- **Go To:** Go to the specified page and section
- **Sort by (ascend):** Sort the report based on a selected column's ascending values
- **Sort by (descend):** Sort the report based on a selected column's descending values
- **Show/Hide Report Toolbar:** This either shows or hides the navigation toolbar at the bottom of the viewer window.

In addition to the Pop-Up Menu, you can also

- Left single click to jump to the hyperlink associated with the cell. (Note that this creates a new browser window if the associated URL is **NOT** a `.pak` file. Users cannot use the *Back* button on the browser to return to the previous report)

- Run the mouse over a cell containing a hyperlink to see the Hint text associated with the link.

8.6.4. Connecting to the ERES Server

In order for the Report Viewer applet to connect to the ERES Server (to fetch the report data, and save exported files) you will need to specify the following parameters in your HTML code:

```
<PARAM name="comm_protocol" value="servlet">
<PARAM name="comm_url" value="http://machine:port">
<PARAM name="servlet_context" value="ERES/servlet">
```

Adding the above params to your HTML code will make Viewer connect to the server at `http://machine:port/ERES/servlet`.

8.6.5. Swing Version

A JFC/Swing version of the Report Viewer can also be used by referring to `SwingReportViewer.jar` instead of `ReportViewer.jar` in the `lib` directory under the ERES install directory. Call the following class when using the Swing viewer: `quadbase.reportdesigner.ReportViewer.swing.Viewer.class`. You can also specify to use the Swing Viewer when creating an applet page in the Report Designer.

8.6.6. Exporting from Viewer

This section pertains to exporting from the viewer in an Applet only. When using the Report Viewer API to create a `Component` that displays a report, the user can right-click on the `Component` to launch a pop-up menu with various options. One of the options (Output → Server → Generate...) is for exporting the report to different file formats. This will result in creating the exported file on the server side. However, there is an API feature that utilizes a server-side Java Servlet to stream back the exported content to the client's browser. The following describes how to use this feature.

In the code that retrieves the report `Component`, two lines of code need to be added:

```
Viewer viewer = new quadbase.reportdesigner.ReportViewer.Viewer();
Component comp = viewer.getComponent(report);
viewer.setExportServlet("http://host:port/ERES/servlet/
ViewerExportServlet");
viewer.setDynamicExport(true, "host", port, "ERES/servlet/");
```

The argument for the `setExportServlet` method is the url location of your deployed `ViewerExport-Servlet`. Please see the javadoc Specification for more details about the arguments of these two methods.

Then, when the user views the report Component, there will be new options under Output → Client → (Generate...). All of these options will result in a pop-up browser window that contains the streamed content of the report.

8.7. Page Viewer

The Report Viewer component offers users the capability of viewing and manipulating a report through a Web browser. The viewer loads the report template on the client and shows it to the user with the latest data. Report Viewer, however, is not the ideal deployment solution when the report contains a large amount of data. In these situations, loading the entire report on the client would cause the client to quickly run out of memory.

For situations like this, EspressReport provides another viewer called Page Viewer. Instead of loading the entire report on the client, Page Viewer employs a page serving technology that allows the client to load only one page of the report at a time. This allows the report to load faster, and conserve the client memory.

Page Viewer works by exporting the report into a series of page (`.page`) files that contain the drawing information for a set of pages in the report. The `.page` files are then loaded in the viewer on the client. Users running with

EspressManager, can directly pass in a report template as part of Page Viewer code. The report will be exported on the server with the latest data, and the client will load the `.page` files. Users can export the report in "view" format either from the API or Report Designer. These exported files can be loaded directly in Page Viewer.

8.7.1. Launching Page Viewer

The Page Viewer can be automatically embedded in a menu page or be manually embedded in an applet using a `.pak, .rpt` file or an exported view (`.view`) file. Page Viewer can also be used to load a report directly from the Preview window in Report Designer.

To embed the Page Viewer within a menu page simply select `RPT` as the export format when viewing the menu page. For more information about the menu page, please see Section 8.1 - The Menu Page.

To embed Page Viewer in an applet, simply use the following code:

```
<applet codebase=".." code="quadbase.reportdesigner.PageViewer.Viewer.class"
 width=100% height=100% archive="PageViewer.jar">

      <PARAM name="filename" value = "yourReport.pak">

</applet>
```

The parameter value of `filename` specifies the name of the file that contains the report. You can specify the name of a report template or an exported view file for this parameter. When using a view file with Page Viewer, you do not have to connect to the ERES Server. The filename value can also be specified by a URL; for example, you can prefix it by `http://` for accessing a remote data file.

You may also use Page Viewer in an application, simply use the following code to get a `java.awt.Component` Object.

```
quadbase.reportdesigner.PageViewer.Viewer.getComponent
(java.awt.Frame frame, java.lang.String filename, long bufferTimeInSec)
```

In addition to passing in the template file (or view file), you can also pass in a `QbReport` object. Note that instead of using the Viewer class, you would be using the ViewerAPI class.

```
quadbase.reportdesigner.PageViewer.ViewerAPI.getComponent(java.awt.Frame
 frame, QbReport report, long bufferTimeInSec, java.lang.String
 securityLevel);
```

The above methods are the simpliest of all methods for getting Page Viewer as a Component. For more information about the parameters for this constructor or information on other methods related to Page Viewer, please consult the Report API Specification.

Note that on the server side, you may need to copy the `PageViewer.jar` file from the `lib/` directory of the ERES installation to some directory of your web server that is accessible by the web client. Please also note that if you are having troubles with refreshing reports, see Section 8.7.4 - Buffer Time that deals with the `BufferTime` parameter.

8.7.1.1. Launching from Report Designer

Page Viewer can also be invoked from Report Designer. When previewing a report with a large amount of data, the report can be loaded in a Page Viewer window. This allows users to preview/view the entire report without loading all the data in memory. For more information about this feature, please see Section 5.1.4.2.4 - Using Page Viewer.

8.7.2. The Page Viewer Parameters

Without using Report Designer, you can also pass data via parameters to Page Viewer applet. That is, you can use Page Viewer to directly view a data file, or pass the data directly in the form of lines of data, in the HTML code. The following is a list of parameters:

Page Parameters

comm_protocol: The protocol to be used to connect to the ERES Server

comm_url: The URL to connect to the ERES Server

servlet_context: The context in which the ERES Server is running

Filename: Specifies the name of the `.pak` file

BufferTime: The time, in seconds, that controls how often to get new pages from the server. It looks at the local version of the `.page` file (if available), and if the time since it is last updated exceeds the `bufferTimeInSec` parameter, then new version of the page is requested again from the server. Otherwise it uses the local version of the `.page` file. If the local version of the `.page` file does not exist, it gets the page from the server

DataHintOffsetX: When the user moves the mouse over a data point (a bar on a bar graph or a point on a line graph) on a chart, a hint box that shows the details of the data point is shown. This parameter sets the *horizontal* offset position of the hint box relative to the data point

DataHintOffsetY: When the user moves the mouse over a data point (a bar on a bar graph or a point on a line graph) on a chart, a hint box that shows the details of the data point is shown. This parameter sets the *vertical* offset position of the hint box relative to the data point

DataHintBgColor: When the user moves the mouse over a data point (a bar on a bar graph or a point on a line graph) on a chart, a hint box that shows the details of the data point is shown. This parameter sets the *background* color of the hint box

DataHintFontColor: When the user moves the mouse over a data point (a bar on a bar graph or a point on a line graph) on a chart, a hint box that shows the details of the data point is shown. This parameter sets the *font* color of the hint box

DataHintFont: When the user moves the mouse over a data point (a bar on a bar graph or a point on a line graph) on a chart, a hint box that shows the details of the data point is shown. This parameter sets the *font type* of the hint box

Printing: Valid values: `true|false`. Sets whether the print option is shown in the pop-up menu

RefreshData: Valid values: `true|false`. If false, report cannot be refreshed. The default is `true`

PopupMenu: Valid values: `true|false`. Sets whether the pop-up menu is shown

EnableExport: Valid values: `true|false`. If false, report cannot be exported. The default is `true`

ShowDataHint: Valid values: `true|false`. Sets whether to show a data hint box when mouse over a data point in a chart

ShowLinkHint: Valid values: `true|false`. Sets whether to show a data hint box when mouse over a hyperlink

Example: With the parameter `BufferTime` you can insert the following line:

```
<PARAM name="BufferTime" value="3600">
```

In the above example, when Page Viewer applet opens, it will fetch data from the specified data source (with the help of the server) if more than 1 hour (3600 seconds) has passed since the local existing `.page` file has been updated. Otherwise Page Viewer will display the local existing `.page` file, if it exists. If it doesn't exist, it also fetches data from the specified data source.

Note that when specifying color parameters Use the Hexadecimal format for RGB colors that is standard in HTML tags. That means the first two hexadecimal digits are for red, the next two are for green, the last two are for blue. A total of six hexadecimal digits has to be specified for a color. For example:

```
<PARAM name="DataHintBgColor" value="ff00aa">
```

8.7.3. Using Page Viewer

Using Page Viewer, you can browse through the pages of the report using the Pop-Up Menu. Right clicking on the report displays a Pop-Up Menu which can be used to traverse the report. The Pop-Up Menu can show the following options:

- **Back:** Go back to the previous report. This option is shown when a hyperlink is clicked on.
- **Section:** Contains various options to navigate the various sections

  First Section: Go to the first section

  Previous Section: Go to the previous section

  Next Section: Go to the next section

  Last Section: Go to the last section
- **Page:** Contains various options to navigate the various pages

  First Page: Go to the first page

  Previous Page: Go to the previous page

  Next Page: Go to the next page

  Last Page: Go to the last page
- **Output:** Contains various formats to output the report to

  Generate DHTML (Single Page): Export the report to DHTML onto a single file

  Generate DHTML (Multiple Pages): Export the report to DHTML onto multiple files, one file per page.

  Generate PDF: Export the report to PDF

  Generate CSV: Export the report to CSV

  Generate EXCEL (XLS): Export the report to Excel (`xls` file)

  Generate EXCEL 2007 (XLSX): Export the report to Excel 2007 (`xlsx` file)

  Generate TXT: Export the report to TXT

  Generate RTF: Export the report to RTF

  Generate XML (Data + Format): Export the complete report to XML

  Generate XML (Pure Data): Export the report data only to XML

Print: Print the report

- **Refresh:** Refresh the data from the data source
- **Go To:** Go to the specified page and section
- **Sort by (ascend):** Sort the report based on a selected column's ascending values
- **Sort by (descend):** Sort the report based on a selected column's descending values
- **Show/Hide Report Toolbar:** This either shows or hides the navigation toolbar at the bottom of the viewer window.

In addition to the Pop-Up Menu, you can also

- Left single click to jump to the hyperlink associated with the cell. (Note that this creates a new browser window if the associated URL is **NOT** a `.pak` or a `.rpt` file. Users cannot use the Back button on the browser to return to the previous report)
- Run the mouse over a cell containing a hyperlink to see the Hint text associated with the link.

8.7.4. Buffer Time

A problem working with the view (`.view`) and page (`.page`) files associated with Page Viewer is that while saving memory and boosting performance by caching the report page by page on a file system, Page Viewer files' data might become obsolete in time. In other words, they might not reflect the latest report data. The solution to this is to have a mechanism for updating the Page Viewer files during certain time intervals.

Page Viewer provides the mechanism to adjust the time elapsed until a new view file needs to be fetched from the server. This time interval, also known as the **Buffer Time**, can be adjusted by setting the parameter `BufferTime` in Page Viewer applet. Simply pass in the buffer time in seconds as the value to the parameter `BufferTime`. The Page Viewer will then fetch new pages whenever the current version of the VIEW file becomes older than the specified buffer time.

Also, Page Viewer has a built-in mechanism that automatically eliminates old view and page files when they have not been updated for 3 days or more.

8.7.5. Connecting to the ERES Server

In order for the Page Viewer applet to connect to the ERES Server (to fetch the page files and save exported files) you will need to specify the following parameters in your HTML code:

```
<PARAM name="comm_protocol" value="servlet">
<PARAM name="comm_url" value="http://machine:port">
<PARAM name="servlet_context" value="ERES/servlet">
```

Adding the above parameters to your HTML code will make Viewer connect to the server at `http://machine:port/ERES/servlet`.

8.7.6. Swing Version Available

A JFC/Swing version of the Page Viewer can also be used to referring to `SwingPageViewer.jar` instead of `PageViewer.jar` in the `ERES/lib` directory.

8.7.7. Exporting from Viewer

This section pertains to exporting from the viewer in an Applet only. When using the Page Viewer API to create a `Component` that displays a report, the user can right-click on the `Component` to launch a pop-up menu with various options. One of the options (Output → Server → Generate...) is for exporting the report to different file

formats. This will result in creating the exported file on the server side. However, there is an API feature that utilizes a server-side Java Servlet to stream back the exported content to the client's browser. The following describes how to use this feature.

In the code that retrieves the report `Component`, two lines of code need to be added:

```
Viewer viewer = new quadbase.reportdesigner.PageViewer.Viewer();
Component comp = viewer.getComponent((Applet)null, reportTemplate, 0);
viewer.setExportServlet("http://host:port/ERES/servlet/
ViewerExportServlet");
viewer.setDynamicExport(true, "host", port, "ERES/servlet/");
```

where `reportTemplate` is a string containing the name and path of the report template file. You can also pass in a `QbReport` object using the following code:

```
ViewerAPI viewer = new quadbase.reportdesigner.PageViewer.ViewerAPI();
Component comp = viewer.getComponent((Applet)null, report, 0, null);
viewer.setExportServlet("http://host:port/ERES/servlet/
ViewerExportServlet");
viewer.setDynamicExport(true, "host", port, "ERES/servlet/");
```

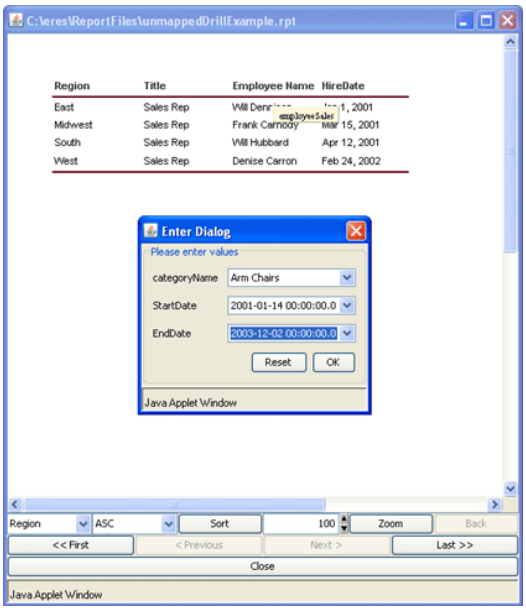
where report is a `QbReport` object.

The argument for the `setExportServlet` method is the url location of your deployed `ViewerExport-Servlet`. Please see the javadoc Specification for more details about the arguments of these two methods.

Then, when the user views the report`Component`, there will be new options under Output → Client → (Generate...). All of these options will result in a pop-up browser window that contains the streamed content of the report.

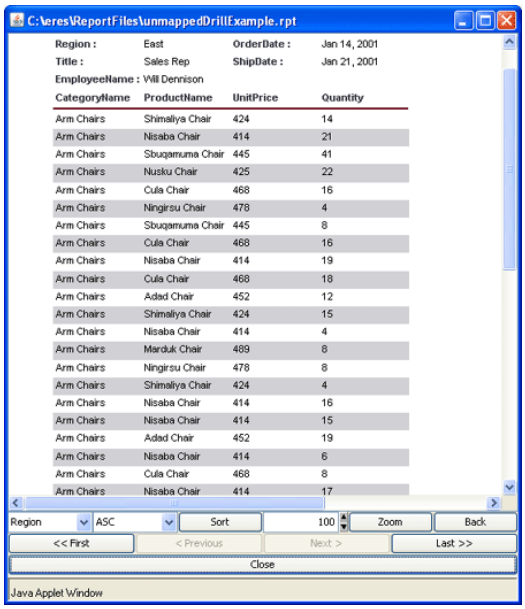
8.7.8. Unmapped Drilldown Parameter

When viewing a report with unmapped drilldown parameters, clicking on the links in the main report will cause a parameter selection dialog to pop up allowing you to fill in the unmapped parameters.



Parameter Selection Dialog

Once the values are entered, click *OK* and the drill report will be shown. To return to the previous level, click on the *back* button in the page viewer toolbar.



Drill-Down Level

8.8. Chart Viewer

All charts created by Chart Designer can be saved in a range of file formats that may be pasted into documents. These formats include BMP, JPG, PNG, PDF, SVG, SWF, WMF, and GIF. In addition, Chart Designer provides the option of saving a chart as a `.cht`, `.tpl`, or `.xml` file.

Chart Viewer is a JNLP that enables you to view and manipulate a chart dynamically through a web browser. Viewer reads the file (in .cht or .tpl format) as outputted by Chart Designer or the API, and then displays the chart. The small size of the data file makes it suitable to distribute the chart image over the web. The data is encrypted while being transferred from EspressManager to Chart Viewer and so lends a degree of security to sensitive information.

Files in the `.tpl` format can also be viewed using Chart Viewer. When a web page that contains a .tpl file is viewed using a browser, fresh data will be obtained automatically by Chart Viewer. Thus, a single chart template can be used to supply up-to-the-minute charts to users in real time.

Inside Chart Viewer, you can drag the chart, legend, title, or label to position the object. You can also resize the chart and drill-down on data points or a series of data. For three-dimensional charts, users can use the navigation panel to pan, zoom, rotate in each direction, and translate. Also, individual x, y, and z-axis scaling, thickness ratio adjustment, real time three-dimensional animation, etc can all be preformed easily. There are built-in callback mechanisms that let a user click on a data element to view the underlying data or to jump to a related URL. Chart Viewer is written in pure Java that runs on all platforms that support Java. Chart Viewer also supports scheduled refresh (where a chart's data is updated at regular intervals specified by the designer) and parameter serving where a chart's parameters are provided at load time.

Chart Viewer is an applet that allows you to view and manipulate a chart dynamically through a web browser. The Viewer can read chart files (`.cht`), template files (`.tpl`) and chart XML files. Chart Viewer can be automatically embedded in a menu page, in an HTML page, or manually placed in a web page.

Inside the applet, you can drag the chart, legend, title, or label to position the object. You can resize the chart and drill-down on data points or a series of data. For three-dimensional charts, users can use the navigation panel to pan, zoom, rotate in each direction, and translate. Also, individual x, y, and z-axis scaling, thickness ratio adjustment, real time three-dimensional animation, etc can be preformed easily. There are built-in callback mechanisms that let a user click on a data element to view the underlying data or to jump to a related URL. Chart Viewer also supports scheduled refresh (whereby a chart's data is updated at regular intervals specified by the designer).

To embed Chart Viewer within a menu page, simply select *CHT* as one of display options when viewing the menu page. For more information about the menu page, please see Section 8.1 - The Menu Page.

To create a stand-alone HTML page with the Viewer embedded, check the *Create HTML File* option when saving a chart in the Chart Designer.

In addition you can embed the Viewer with your own Web page using the following syntax:

```
<applet-desc
      name="Chart Viewer"
      main-class="quadbase.chartviewer.Viewer"
      width="800"
      height="600">
        <param name="filename" value="help/examples/ChartAPI/data/test.tpl"/
>
        <param name="preventSelfDestruct" value="false"/>

</applet-desc>
```

The parameter `filename` specifies the file name of the file that contains the chart data and you can prefix it by `http://` for accessing a remote data file. When viewed by Chart Viewer, a chart saved in the chart format (`.cht`) will use the data stored in that file for plotting.

A chart saved in the template format (`.tpl`) allows Chart Viewer to dynamically fetch the data from a database or a data file depending on where you specify the data source of chart to be when using Chart Designer to create the template (the database name, user name, password, etc are all stored in the `.tpl` file).

8.8.1. The Chart Viewer Parameters

You can also pass data and chart viewing control information via parameters to the Chart Viewer JNLP without using Chart Designer. That is, you can use the Chart Viewer to directly view a data file or pass the data directly in the form of lines of data, along with other control information, in the HTML code. By default, the parameter is true if the parameter is of a true/false type. The following is a list of parameters:

Chart Parameters (Common to Two and Three-Dimensional Charts):

mainTitle	the main title of the chart
xTitle	x-axis title
yTitle	y-axis title
zTitle	z-axis title
RefreshInterval	scheduled refresh interval in seconds
DragLegend	if false, the legend(s) can not be moved
DragChart	if false, the chart can not be repositioned or resized
ShowDataHint	if false, data information box will not be shown when left mouse click on chart data
ShowLinkHint	if false, hyperlink information box will not be shown when right mouse click on chart data
DataHintBgColor	set background color of data information box
LinkHintBgColor	set background color of hyperlink information box
DataHintFontColor	set font color of data information box
LinkHintFontColor	set font color of hyperlink information box

DataHintFont	set font of data information box
LinkHintFont	set font of link information box
DataHintOffsetX	set x offset of the data information box
DataHintOffsetY	set y offset of the data information box
LinkHintOffsetX	set x offset of the link information box
LinkHintOffsetY	set y offset of the link information box
Printing	if false, this will disable the ability to export the chart (using **Ctrl**+**P** and/or **Ctrl**+**J**) in a browser
filename	name of the template file to be applied to the chart
xAxisRuler	if true, show the x-axis ruler (for 2D charts only)
yAxisRuler	if true, show the y-axis ruler (for 2D charts only)
sAxisRuler	if true, show the secondary-axis ruler (for 2D charts only)
ResizeChart	if false, the chart cannot be resized
ResizeCanvas	if false, the canvas cannot be resized
comm_protocol	the protocol to be used, to connect to the ERES Server
comm_url	the URL to connect to the ERES Server
servlet_context:	the context in which the ERES Server is running
RefreshData	if false, the chart data cannot be refreshed
PopupMenu	if false, the pop-up menu will not be displayed
TypeMenu	if false, the type sub-menu will not be displayed in the pop-up menu
DimensionMenu	if false, the dimension sub-menu will not be displayed in the pop-up menu

For Three-Dimensional Charts only

Toggle3Dpanel	if false, the navigation panel can not be toggled to be visible or invisible
Drawmode	set different mode of drawing 3D chart. Available draw modes are **`Flat`** (default), **`WireFrame`**, **`Flat Border`** (which draws a black border around the flat shading model), **`Gouraud`**, and **`Gouraud Border`**
NavColor	set navigation panel color
navpanel	if false, the Navigation Panel is not displayed when a 3D chart is being viewed. The Navigation Panel is never displayed when a 2D chart is being viewed
GouraudButton	if false, the Gouraud shading button in the navigation panel is hidden
AnimateButton	if false, the animation speed control in the navigation panel is hidden
SpeedControlButton	if false, the speed control button in the navigation panel is hidden

Data Input Parameters

sourceDB	set the database information in order to generate the chart
sourceData	set the data information in order to generate the chart
sourceFile	set the datafile information in order to generate the chart
datamap	set the column mappings for the chart
TransposeData	set the data to be transposed before using it to generate the chart
chartType	set the chart type for the generated chart
ParameterServer	update the data in the chart dynamically
transposeData	if true, transpose the data

*Example:*the parameters `mainTitle`, `xTitle`, `yTitle`, and `zTitle` are used to specify the main title and axis title of the chart, they will override the ones defined in the template:

```
<PARAM name="mainTitle" value="This is the main Title">
<PARAM name="xTitle" value="x axis title">
<PARAM name="yTitle" value="y axis title">
<PARAM name="zTitle" value="z axis title">
```

*Example:*With the parameter `RefreshInterval` you can specify:

```
<PARAM name="RefreshInterval" value="60">
```

In the above example, the applet will fetch data from a database or data file (with the help of the server) and redraw the chart every 60 seconds - all transparently. It is useful for accessing databases in which the data changes frequently.

8.8.2. Specifying the Data Source for Chart Viewer

Using parameters, you can specify a data source for the Chart Viewer in order to display different data with a chart template or create a chart from scratch.

8.8.2.1. Data Read From a Database

This is some sample JSP/JNLP code that uses the Chart Viewer to view a chart drawn using data extracted from a database.

```
<?xml version="1.0" encoding="utf-8"?>
<jnlp spec="1.0+" href="EspressViewer.jnlp">
<information>
<title>Espress Viewer</title>
<vendor>Quadbase Systems Inc.</vendor>
<offline-allowed/>
</information>
<resources>
<j2se version="1.8+" max-heap-size="1024m"/>
<jar href="lib/EspressViewer.jar"/>
</resources>
<security>
<all-permissions/>
</security>
```

```
<applet-desc
name="Chart Viewer"
main-class="quadbase.chartviewer.Viewer"
width="640"
height="480">

<PARAM name="sourceDB" value="jdbc:odbc:DataSource,
 sun.jdbc.odbc.JdbcOdbcDriver, username, password, select * from products">
<param name="dataMap" value="0 1 -1 3">
<param name="chartType" value="3D Column">

</applet-desc>
<update check="always" policy="always"/>
</jnlp>
```

The arguments of `dataMap` specify how the chart utilizes different columns from the input data to plot the chart. In case of a scatter chart, they are series, x-value, y-value, and z-value. For a high low open close or high low chart the numbers are series, category, high, low, open, and close. For all other charts, the arguments are series, category, sumBy, and value. If you would like more information,please see the chapter on Column Mapping in the Chart API reference. The argument `chartType` specifies the type of chart to be displayed and it can be one of:

- 2D Column
- 3D Column
- 2D bar
- 3D bar
- 2D stack bar
- 3D stack bar
- 2D stack column
- 3D stack column
- 2D area
- 3D area
- 2D stack area
- 3D stack area
- 2D line
- 3D line
- 2D pie
- 3D pie
- 2D scatter
- 3D scatter
- 2D High Low
- 3D High Low

- 2D HLCO
- 3D HLCO
- 2D 100% Column
- 3D 100% Column
- 3D Surface
- 2D Bubble
- 2D Overlay
- 2D Box
- 2D Radar
- 2D Dial
- 2D Gantt
- 2D Polar

8.8.2.2. Data Read From a Data File

This HTML/jsp code draws a chart using data from a data file:

```
<applet-desc code = "quadbase.chartviewer.Viewer.class" width=640
 height=480>

      <param name="sourceFile" value="http://.../test.dat">
      <param name="dataMap" value="-1 0 -1 1">
      <param name="chartType" value="3D Pie">
</applet-desc>
```

8.8.2.3. Data Read From an Argument

It is possible to have Chart Viewer read in data directly from the HTML/jsp file itself rather than from a data file or from a database.

```
<applet-desc code = "quadbase.chartviewer.Viewer.class" width=640
 height=480>

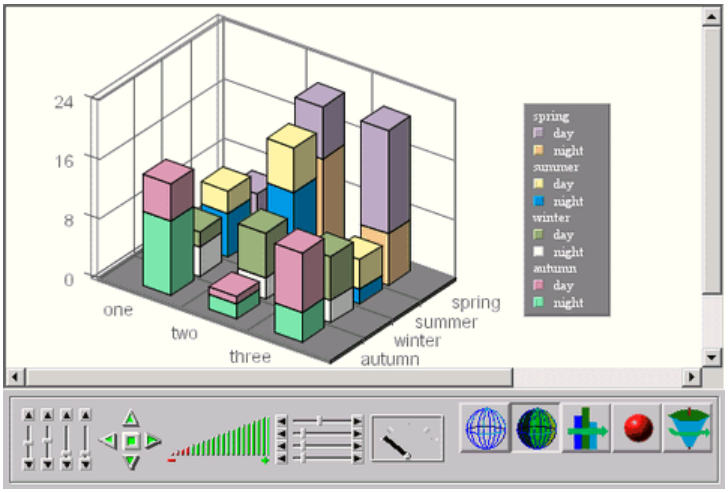
      <param name="sourceData" value="int, string, int | value, name, vol |
 10, 'John', 20 | 3, 'Mary', 30 | 8, 'Kevin', 3 | 9, 'James', 22">
      <param name="dataMap" value="-1 1 -1 2">
      <param name="chartType" value="3D Bar">

</applet-desc>
```

The format for the parameter sourceData is the same as the data file format, except that each line is ended by a vertical bar "|".

8.8.3. Using Chart Viewer

You can manipulate the chart appearance simply by using the mouse in Chart Viewer. For three-dimensional charts, Chart Viewer screen comprises of two sections: the *drawing panel* and the *navigation panel*. Only the *drawing panel* is visible for two-dimensional charts. Within the *drawing panel*, the chart itself is located on the *plot area*.



A Sample Chart Viewer Display

Below is a list, not by order or importance, of chart manipulations that are possible when using Chart Viewer:

- Using the navigation panel, you can:
  - Change light position
  - Rotate the chart
  - Translate the chart
  - Zoom in/out on the chart
  - Scale the chart in x, y, z dimension
  - Adjust the thickness ratio of bar/line/point/pie
  - Start the chart animation and control the speed of the animation
  - Toggle between wireframe and solid mode
  - Draw a black outline on all edges of the chart
  - Perform Gouraud shading
- Drag the left mouse button on the chart, main title, x, y, and z-labels, or legend to move the item
- Drag the right mouse button on plot area to resize the chart
- Press the **Alt** key and drag the right mouse button on plot area to resize the canvas
- To toggle the navigation panel (show and hide), **double click** the left mouse button on the drawing panel
- Using the mouse you can also query data points individually
- **Left single click** on a data point to view the data associated with that point
- A **right single click** on data provides the name of the hyperlink associated with the data point
- **Left double click** to jump to the hyperlink associated with the data point. (Note that this creates a new browser window. Users can not use the *Back* button on the browser to return to the previous chart)
- **Right double click** to jump back to the previous chart (if appropriate)
- Use **Alt**+**Z** to specify the zoom-in parameters

- Use **Ctrl**+**R** to refresh data manually
- Use **Ctrl**+**Left Click** to select the bounds to zoom in and **Ctrl**+**Right Click** to zoom out

8.8.4. Axis Rulers

There is now an option to show Axis Rulers in Chart Viewer (by setting a parameter `{axisRuler}` in the applet-desc tag) and Chart API (Please refer to the API Documentation quadbase.util.IAxisRuler [https://data.quadbase.com/Docs72/eres/help/apidocs/quadbase/util/IAxisRuler.html]). This provides a reference point when scrolling is enabled and the chart is moved such that the chart's own axes are not visible. This feature is for 2D charts only.

8.8.5. Pop-up Menu

If the pop-up menu option was selected during the chart creation/editing, the chart can be modified. The menu is opened by right clicking on the chart and the options there can be selected by highlighting and left clicking on it. Options available for viewing, when enabled, are dynamic data drill-down, changing chart types, axis zooming, and time-series data zooming.

8.8.5.1. Changing Chart Dimension and Type

The chart dimension and chart type can be changed by using the pop-up menu and selecting either the *2D/3D* option or the *Type* option.

Please note that *Overlay*, *Box* and *Dial* chart options do not appear if the chart is three-dimensional.

8.8.5.2. Axis Zooming

In the case of a large chart whose dimensions exceed that of the viewport (i.e. if a chart is large enough that it cannot be completely seen within window), an option exists to allow for axis zooming. This allows you to move all the axis and zoom in and out of the axis.

This option is only available for two-dimensional charts. Please note that the bubble, dial, pie, polar, radar, scatter, stack area, and vertical box charts do not support this feature.

You can enable this option by opening the pop-up menu and selecting *Enable Zoom*. You can choose from the zooming in on the x-axis, the y-axis or both.

To zoom in, left click on and drag out the desired area. If only the y-axis zooming is selected, you need not worry about the length or position in the x-axis.

Scrolling is enabled only for axes that have been zoomed. There will be scrollbars visible on the particular axis when axis zooming is enabled. You can left click and drag this scrollbar to shift the viewport of the applet.

Holding down the control key while left clicking will take the applet back to the previous zoom. Please note that only one previous zoom is stored in memory and thus you can only go back one step. To go all the way back to the default x and y scale, press the **Home** key on your keyboard.

8.8.5.3. Zooming

If the chart being shown is a zoom enabled chart, the zoom option will be available from the pop-up menu. Using the zooming option, you can group the category elements into user-defined intervals and aggregate the points in each group. For details on date/time based zooming, see Section 5.2.4.8.2 - Date/Time Based Zooming. Enabling zoom will allow you to input various zoom options such as lower bound, upper bound (you can also disable the bounds in which case it will go from the first data point to the last), the time scale, and whether you want the x-axis to be linear or not.

8.8.5.4. Dynamic Data Drill-Down

You can configure the next drill-down by selecting *drill-down* from the pop-up menu. This option is available only if dynamic data drill-down is enabled for the chart. After selecting the *Drill-Down* option the in the pop-up menu, you can view the current setting for the next drill-down chart. If the *Category* is *None*, the next level has not yet been configured. If there are unused columns in the data used to generate the chart, you can select a column for the

Category. After the *Category* has been selected, the *Type* of the chart in the next level can be set (along with *Series* and *SumBy*). The data points can now be left clicked to move down a level and right clicked to move up a level.

8.8.5.5. Query Parameter

When viewing a parameterized chart, you can enter in a different value for the parameter (or parameters depending on the number) by opening the pop-up menu and selecting the *Query Parameter* option. The chart is then redrawn with the new data based on the parameters selected.

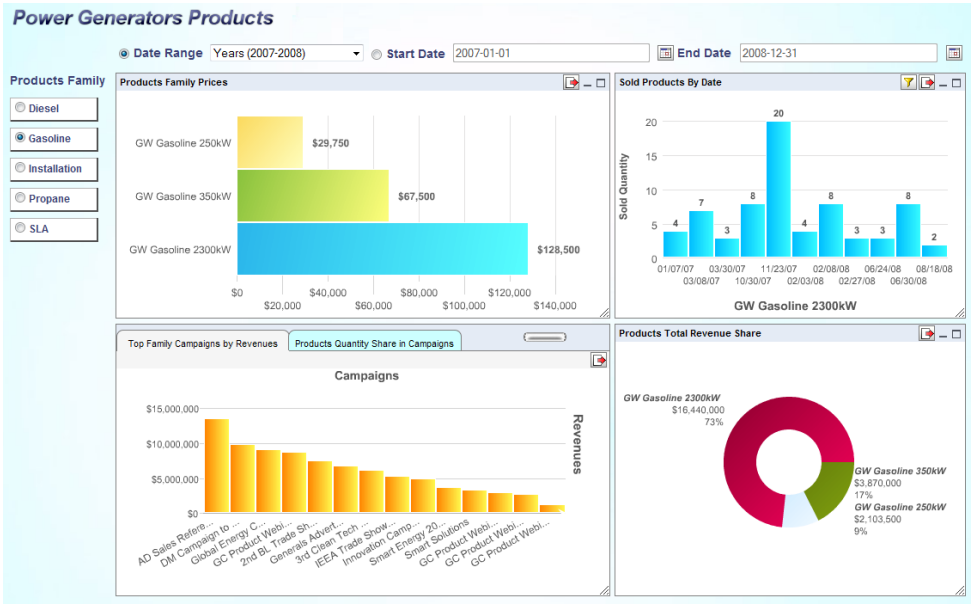
8.8.6. Swing Version Available

A JFC/Swing version of the Chart Viewer can also be used to referring to `SwingEspressViewer.jar` instead of `EspressViewer.jar` in the `ERES/lib` directory. Call the following class when using the Swing viewer: `quadbase.chartviewer.swing.Viewer.class`.

8.9. Dashboard Viewer

Dashboard Viewer is an applet that enables you to view and manipulate a dashboard dynamically through a web browser. Viewer reads the file (in .dsb format) as outputted by Dashboard Builder and then displays the dashboard.

The dashboard can be previewed by clicking the *Preview* button  on the Dashboard Builder toolbar. This will open a new window showing the dashboard.



Dashboard

8.9.1. Preview Toolbar

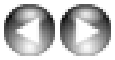
The top part of the Preview interface contains a small toolbar that allows you to initiate the following actions:

 Set shared parameter(s)

 Refresh a dashboard

 Set auto refresh

 Export Dashboard to PDF (Note: All objects in dashboard including Online maps are exported, however, if map type is set to Google street map or Google satellite, map imagery will be missing in the exported file due to its license restrictions.)

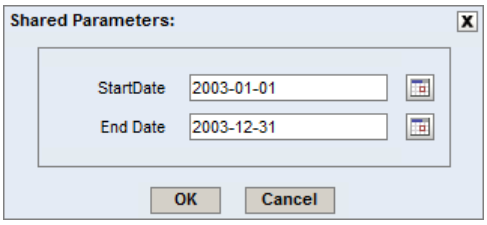
 Pack/Unpack the preview toolbar

8.9.2. Preview Options

This section gives you description of options available from the preview toolbar and additional options available from the chart/report/map header bar. You can hide/show the toolbar by clicking the  /  *Pack/Unpack* icons.

Set Shared Parameters:

To set shared parameters, click the *Set Shared Parameters* button  on the preview toolbar. After you click the button, the shared parameters dialog will then appear. The dialog allows you to specify values for shared parameters.



Set Shared Parameters Dialog

Note that you can use calendar for specifying date parameters. To open the calendar, simply click the *Calendar* button  in the shared parameters dialog.

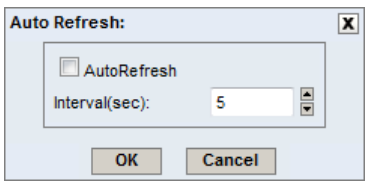


Calendar

Once you finish specifying the shared parameters, click the *Ok* button to save the changes. You will be taken back to the Preview.

Auto Refresh:

If you want your dashboard to be refreshed periodically, click the *Auto Refresh* button  on the preview toolbar. After you do so, the *Auto Refresh* dialog will then appear. You have to check off the *AutoRefresh* checkbox to enable the auto refresh feature. The dialog allows you to set the refresh interval in seconds. For example, if you specify the refresh interval for 10 seconds, the dashboard will refresh every 10 seconds.

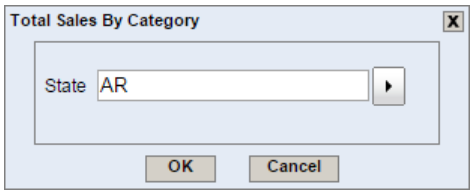


Auto Refresh Dialog

Once you finish setting up the auto refresh options, click the *Ok* button to save the changes. You will be taken back to the Preview.

In addition to the options on the preview toolbar, there are additional options available for a chart/report/map that can be triggered by clicking the small buttons in the chart/report/map header bars.

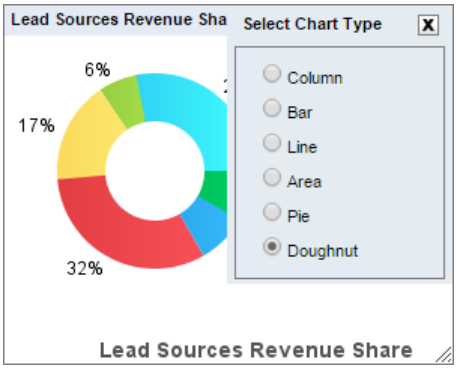
Filter: This option will open a new dialog that allows you to specify additional filters/parameters for the chart/report/map, if it has parameters that aren't associated with one of the dashboard's shared parameters.



Dashboard Filter Chart/Report/Map Dialog

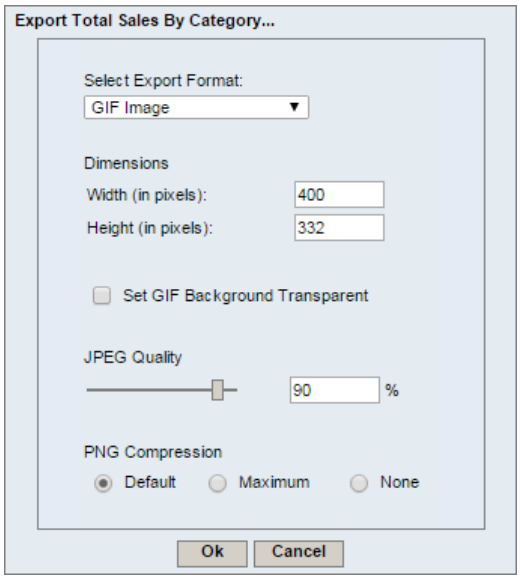
You can select parameter values in this dialog. Once you complete your selection, click the *Ok* button. The dashboard will refresh and the chart/report/map will reflect the new parameter values.

Change Chart Type: This option will open a new dialog that allows you to change a chart type. The chart type is changed only in the Preview. It is not possible to save the changed view of the chart, but you can export it.



Select Chart Type Dialog

Export: This option will open a new dialog that allows you to export the chart/report/map in a variety of formats. For charts the following dialog is displayed:



Dashboard Export Chart Dialog

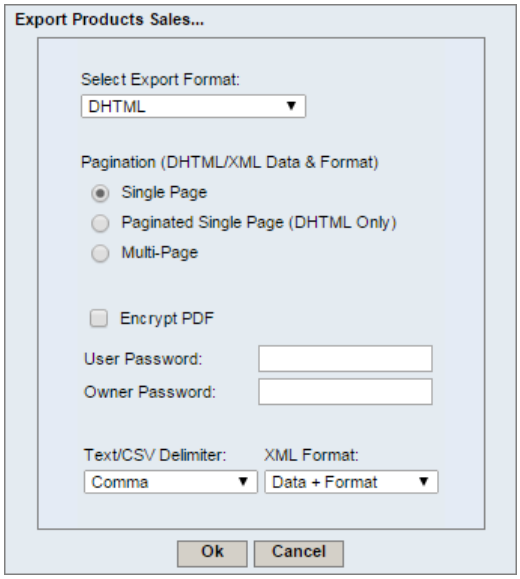
The first option allows you to specify the format in which to export the chart. The available formats are GIF, JPEG, PNG, PDF, SVG, Excel Image (XLS), Flash, Text Data File, XML Data File and MAP Data File (HTML Image Map).

The second option allows you to set dimensions of the exported image. By default, these will match the canvas dimensions of your chart.

The last set of options allows you to set some image type-specific options. You can set the background transparent for GIF files, quality for JPEG files, and specify encoding for PNG files.

Once you finish specifying the options for exporting the chart, click the *Ok* button. A new window will open containing the exported chart.

For reports the following dialog is displayed:



Dashboard Export Report Dialog

The first option in the dialog allows you to specify in which format you want to export the report. Available options are PDF, CSV, Excel (XLS), Excel 2007 (XLSX), DHTML, text, XML and rich text.

Other options allows you to set single or multi-page export for DHTML and XML exports. You can enable encryption for PDF format. For text export you can select to use the delimiter and for XML export you can select whether to export only report data, or XML description of the report and data.

Once you finish specifying the options for exporting the report, click the*Ok* button. A new window will open containing the exported report.

Online Maps and SVG Maps support only one export format, PDF, so no export format settings dialog is displayed for maps. Please note that Online Maps cannot be exported if Google street map/Google satellite is used as map type, due to license restrictions. In this case, map imagery will be missing in the exported file.



Maximize: This option maximizes the report/chart/map.

Minimize: This option minimizes the report/chart/map (only template header bar is visible).



Restore: This option resizes the report/chart/map to its original size.

Chapter 9. Alerts

9.1. What is an Alert

An alert is basically an exception handling mechanism that notifies user when certain data is out of a pre-defined range. Typically, the user is looking at some KPIs (key performance indicators) that are important to his/her organization. Some example KPIs on which a user may want to set alerts are inventory level, profit margin, sales growth rate, percentage of dropped calls in a mobile network, number of intrusion attempts on a company network, response time on customer support/service center etc..

ERES supports two types of alerts - dashboard alerts and monitoring alerts. Dashboard alerts are evaluated only on currently opened dashboards. They are not triggered when the dashboard is not open. Monitoring alerts are alerts in scheduled objects (reports, charts, maps, dashboards). This is similar to scheduling (described in Section 3.2 - Scheduling & Archiving), except that at the time when alert monitoring task is run, alert conditions are evaluated and the monitoring task actions (i.e. Send email, upload to FTP) are run only if an alert was triggered. If no alerts in the alert monitoring task are triggered, no action is performed.

9.2. Specifying alerts

Alerts are specified in a displayable object, namely maps, charts and reports.

9.2.1. Online Maps

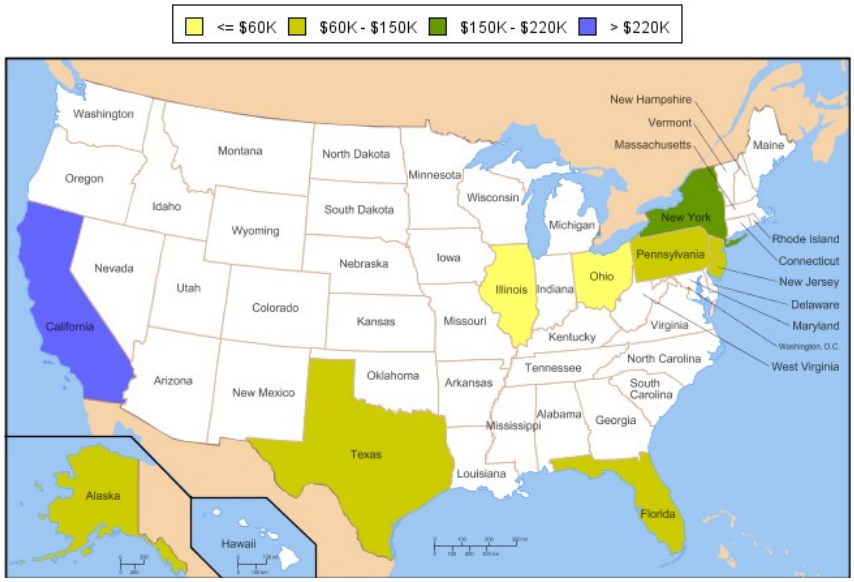
There is currently no support for alerts in Online Maps.

9.2.2. SVG Maps

For SVG Maps, alerts can be specified on all thresholds in the map. You can select ranges that are critical for you. If any of the areas on the map would be in that range, alert will trigger.

If the map doesn't have any thresholds, alerts will not be available for the map.

Example:



SVG map with thresholds

You can for example say that ranges <= $90000 and $90000 - $330000 are critical and you will get alert if there are any yellow or light green areas in the map.

9.2.3. Charts

For Charts, you can specify alerts on control areas. Similarly to SVG maps, you can choose the ranges that are critical for you. If there are no control ranges, alerts will not be available for the chart.

Example:

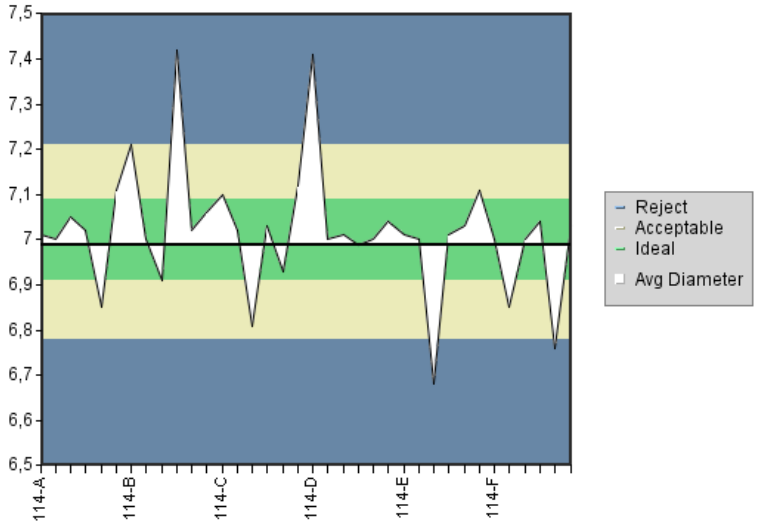


Chart with control areas

You can say that the Reject control area is critical and you will get alert whenever there is at least one data point in this (blue) control area.

Control area ranges are evaluated as left-closed, right-open intervals [start, end). In other words, the alert will trigger if the following condition is true: start <= ActualValue < end.

9.2.4. Reports

For reports, you can specify alerts in scripts. There is a special Formatting action called `ALERT`. You can assign any two strings to this action. First string will indicate name of the alert that should be triggered. Second string represents alert details. Alert details can be any sentence explaining the meaning of the alert or a convenient function result (for example: `getAllRowData()` - explained in Section 5.1.6.2.8.2 - String Functions). Several scripts can trigger single alert, or you can specify several different alerts in one report. For more information about scripting visit Section 5.1.7.1 - Creating a Script.

In the alert `UI`, you will get names of all possible alerts that can be triggered in the report and you can choose which ones you want to watch.

Example:

There are these two scripts in a report:

```
if ({Profit} < 0){
 FONTCOLOR = [255, 0, 0];
 ALERT = ["Negative Profit",getAllRowData()];
} else {
 FONTCOLOR = [0, 255, 0];
}
```

and

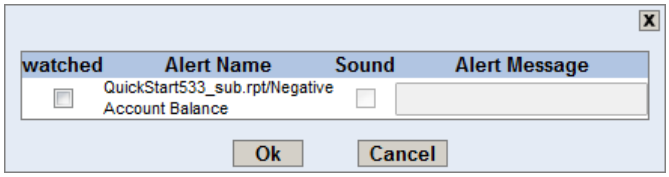
```
if ({AccountBalance} < 0){
 FONTCOLOR = [255, 0, 0];
 ALERT = ["Negative Account Balance","You're in debt!"];
}
```

If there is any row where `AccountBalance` is negative, the alert called `Negative Account Balance` will trigger. If there is any row where `Profit` is negative, the alert called `Negative Profit` will trigger. If both conditions are satisfied at once, both alerts will trigger.

The named alerts are analogical to ranges in maps and control areas in chart. They also have a name and you can select which ones you want to watch (you do not have to watch all of them). How to watch an alert will be described later.

9.2.4.1. Subreports

Subreports support alerts as well. The alert names are propagated to the main report. The alert names are prepended with subreport name to keep them unique in the main report (main report can have alerts with the same names and we don’t want to mix them).



Alert in subreport

9.2.4.2. Charts in reports

For charts in reports, names of the control ranges are propagated to the report. They will also be prepended with the chart name to avoid naming conflicts.

9.3. Dashboard alerts

Dashboard alerts are set up for individual templates (reports, charts, maps) in the dashboard using the user interface in the dashboard builder. Alert (exception) conditions are evaluated when the templates are exported, i.e. during initial loading, auto refresh, manual refresh of the dashboard, changing parameters and resizing the template. Auto refresh is enabled every time you preview or view a dashboard. Default auto refresh interval is 60 seconds.

9.3.1. Visual alert display

If an alert was triggered, color of template border will change and the border will start blinking. For templates in folders, tab header color will change and blink.

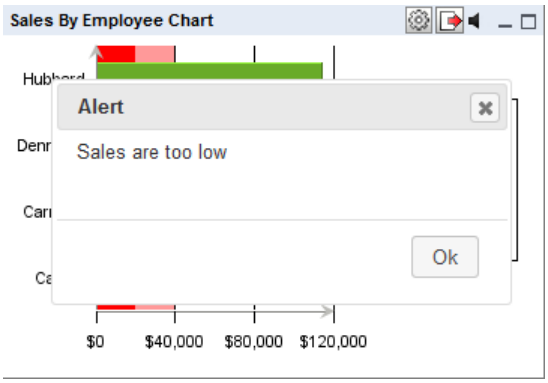
In addition, the template will have its own “native” highlighting of the values:

- coloring of areas in the SVG maps
- control ranges in charts
- for reports, visual changes that can be done by scripting to change text color, background color, border color, change image, change value, etc. (complete list of all formatting options can be found in Section 5.1.7.2.1 - Formatting Actions).

Please note that alert highlighting, blinking and alert messages are visible only in *Preview* of the Dashboard Builder.

9.3.2. Alert messages

If you also set up an alert message, a message with your text will pop-up.



Triggered alert message

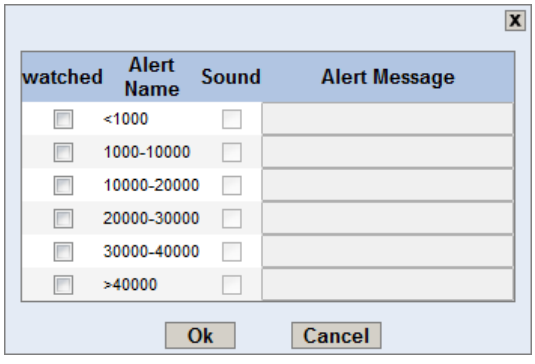
9.3.3. Sound alert

In some cases, sound alerts may be preferred over the visual alerts. When an alert is triggered, the template will continuously make beeping sound. If the alert is muted and you change a parameter in the template and the watched alert is triggered again, the alert sound will be active again. However, if the same watched alert is triggered after auto refresh, the alert sound will not reset, but if another watched alert is triggered, the alert sound will be active again. If the alert is disabled during auto refresh, it will reset back to active state.

9.3.4. User interface

Alerts dialog can be accessed using this button  in the header of every template in the dashboard. This button is only visible for templates that can have some alerts. It will not be displayed for templates that cannot have alerts (i.e. all Google maps, SVG maps without thresholds, charts without control ranges and reports without scripts with ALERT variable).

After clicking the button, you will get a dialog with all possible alerts that can be watched (i.e. names of all control ranges for charts, all ranges for SVG maps or all possible values of the ALERT variable for reports).



Watch alerts dialog

All the *watched* checkboxes will be unchecked by default, which means that no alert is watched. After checking a *watched* checkbox, you enable *Sound* checkbox and *Alert message* textbox, which will allow you to use sound alerts and alert messages. The text you fill in alert message textbox will then display as a pop-up message when the watched alert is triggered. If you leave it empty, no message will appear. You can then check the checkboxes and start watching the alerts. The alert name cannot be changed. It will always be read from the template.

All alerts will be joined using logical OR. If any of these alerts trigger, the dashboard alert will trigger. The AND operation only makes sense for reports and this kind of conditions can always be set directly in the script in the report.

9.4. Monitoring

Monitoring is supported for dashboards and also for individual reports, charts and maps (let's call them **monitored objects**). Monitoring is similar to scheduling. The monitored objects are checked in given time intervals and notifications are sent if alert occurs in some of the monitored objects. This is done in background, so the notifications are sent even if the monitored object is not viewed.

The user interface to set up monitoring is similar to that of scheduling (scheduling described in Section 3.2 - Scheduling & Archiving). If the monitored object has parameters, user can set up several parameter sets (just like for scheduling). Monitoring has the same options for testing intervals as scheduling (i.e. one time, time interval, fixed days, etc.). Monitoring is not supported for drilldowns.

Administrators and designers are the only users who can set up monitoring (only for objects they have read access to). They can also select recipients of notification emails. Viewers are not allowed to set up monitoring.

9.4.1. Alert notifications

Alert notifications are sent via emails. Exported templates can also be uploaded to FTP. Alert notification can be sent every time an alert occurs or it can be send only when an alert starts (i.e. if alert wasn't triggered during last check, but it is triggered during current test).

You can set email subject and body. They have to be the same for the entire monitoring (for all templates, alerts and parameter sets). Body of the notification email can contain some variables that will be substituted during runtime. List of these variables will be detailed later.

You can set different recipients for different templates and different parameter sets (just like in scheduling). It is not possible to set different recipients for different alerts in one template. If you want that, you have to create another alert monitoring task (you can add as many tasks as you want).

9.4.2. Failed emails

If some of the monitored objects are not available, or their data sources are not available, or if there are any other problems with exporting the monitored objects or alert evaluation (e.g. watched control area was removed), the user can be notified with failed emails.

The UI is exactly the same as the Failed emails dialog in scheduling (described in Section 3.2.1.2.1 - Email Delivery Options). You can set failed email subject and body. Recipients are the same as for the notifications. Error log (with stack traces) can be sent to a selected user (typically admin).

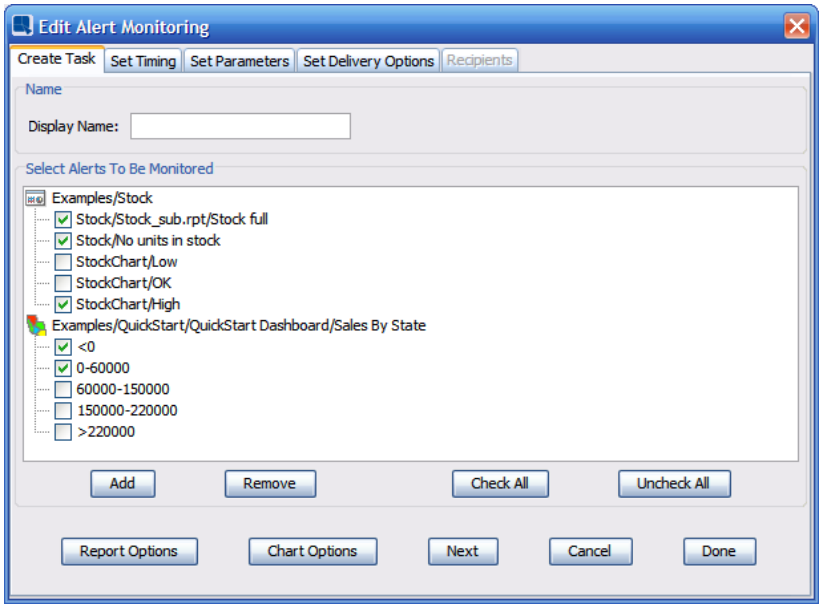
9.4.3. User interface

9.4.3.1. Create/Edit monitoring dialog

Monitoring is set up in the Organizer in a dialog similar to the Schedule dialog (see Section 3.2.1 - Setting a Schedule for more details).

To set up, please use the following steps.

Select alerts - This step is similar to the first step of the Schedule Package wizard. Select the templates that you want to monitor in this monitoring task. You can select any dashboard, stand-alone report, chart or map.



Create monitoring task dialog

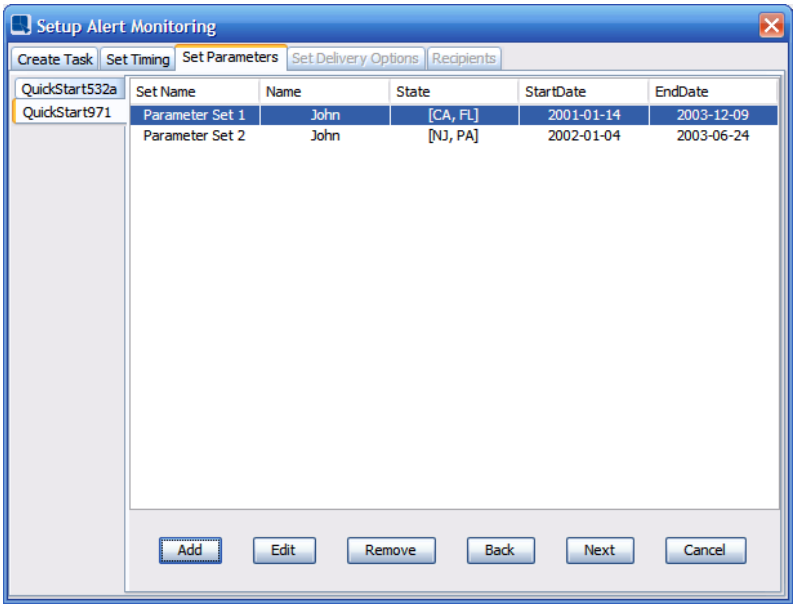
Unlike the list in schedule package, the monitored template list is a tree.

The first level of the tree is the selected templates (before the dialog was opened) or templates added using the *Add* button. The leaf nodes are available alerts (for their parent node). You can check alerts that should be monitored. You have to select at least one alert for each template. If you don't want to monitor any alert of a monitored object, remove the object by selecting it and clicking the *Remove* button.

If the selected template is dashboard that already have some dashboard alerts defined, the checkboxes are initialized according to the dashboard alerts (the alerts that are watched in the dashboard are checked). If the selected template is not dashboard, all the checkboxes are unchecked by default.

Set timing – this step is exactly the same as the Set timing in Scheduling (see Section 3.2.1 - Setting a Schedule for more details).

Set parameters – this step is skipped if there are no parameters in the selected templates. If you have selected to monitor a report, chart, map or dashboard that contains parameters, the next tab will appear allowing you to set the parameter values you want to use.



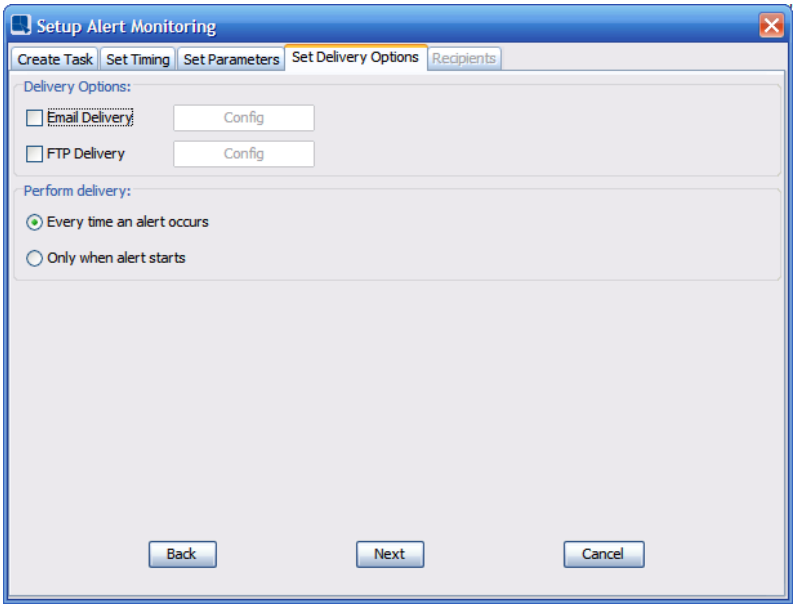
Set parameters dialog

To add a set of parameters, select the *Add* button. This will bring up the parameter prompt dialog allowing you to select the set of parameter values you want to use. Once you have selected a parameter set, your choices will appear in the dialog. You can add as many different combinations of parameter sets as you like. A separate file will be generated for each set of parameters you specify. Alerts will be evaluated separately for each parameter sets. Only the templates that triggered some alerts will be sent in the notification email or uploaded to FTP.

You can also specify a name for each parameter set by double clicking on the first column. The name specified here will be used in later dialogs to help you organize your recipients.

Once you specify all parameters, click the *Next* button to continue.

Set delivery options - This dialog allows you to configure the delivery options for the alert notifications. It is almost the same as Set Delivery Option tab in the Scheduling dialog (described in Section 3.2.1.2 - Additional Delivery Options).



Delivery options dialog

Printer delivery is not available for alert monitoring.

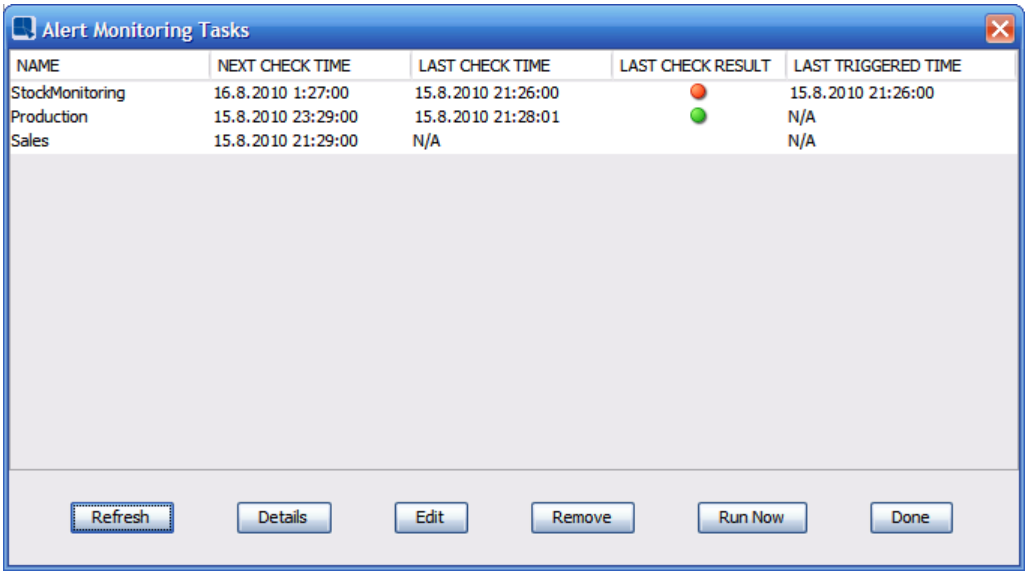
There is an extra option that is not available for scheduling. You can select if you want to perform the delivery (send emails or upload files to FTP) every time an alert occurs or if you only want to perform it when alert starts (i.e. if alert wasn't triggered during last check, but it is triggered during current test).

There are also two extra runtime variables that can be used in email body. It is *<ALERTS>* and it is replaced by names of all alerts that have been triggered during the export and *<ALERT_DETAILS>*. For reports, alert details are replaced by alert details specified in the *alert details* parameter (see Section 9.2.4 - Reports to learn more). For charts, it shows values of the tooltips that would be displayed on the data points that triggered alerts. For maps, it shows all values from the data source rows associated to the areas that triggered alert.

Email Recipients - This dialog allows you to set the recipients of the emails. It is almost the same as the Recipients dialog for scheduling (described in Section 3.2.1.3 - Specifying Email Recipients. The only difference is that alert monitoring doesn't support report bursting. Specifying Email Recipients). You can assign different recipients for individual templates and parameter sets.

9.4.3.2. Monitoring list

You can view monitoring list from pull-down menu Schedule/Archive → View Alert Monitoring Tasks. The dialog allows users to view, edit and remove current monitoring tasks.



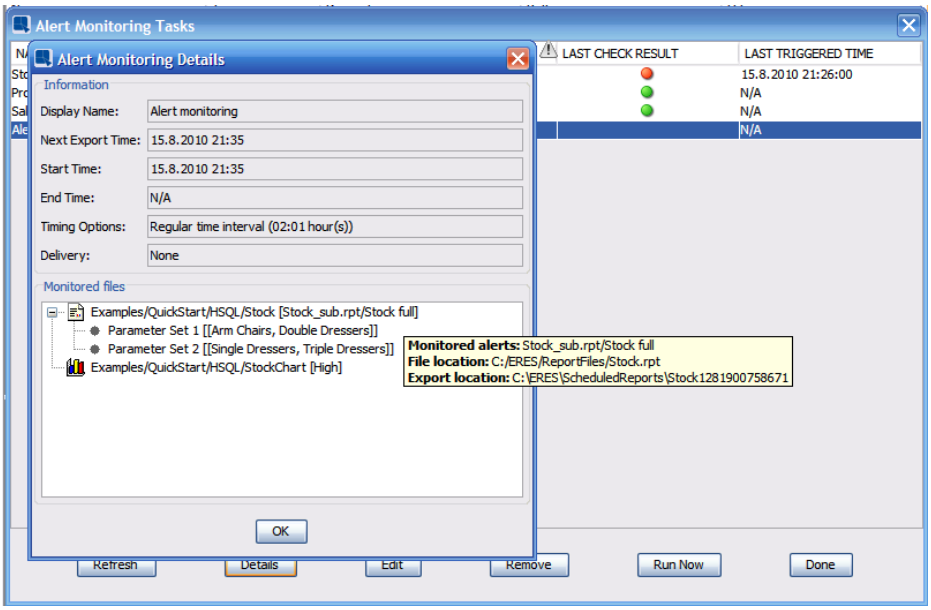
View alert monitoring tasks dialog

It also displays the following info about each monitoring task.

- name
- next check time
- last check time
- last check result (alert was triggered or not) – red light means that an alert was triggered, green light means no alerts
- last triggered time

The *Check Now* button is similar to the *Run Now* button in the Schedule list dialog. It checks selected alerts immediately (and sends notifications if necessary).

You can use the *Details* button to view more details about selected monitoring.



Alert monitoring details dialog

The Monitored files list contains list of all monitored files with all monitored alerts (in the square brackets). It also contains all parameter sets if the template is parameterized. If you hold cursor over a template or parameter set in the Monitored files list, you will get a tooltip with more details.

9.5. Handling special situations

9.5.1. Change in script/control area/range

If for some reason an alert is removed or renamed, the monitoring job or dashboard alert cannot find the missing alert. The following approaches will be taken to resolve the situation.

- For monitoring – An error email is sent. An explanation is included, e.g. alert `<ALERT_NAME>` not found in the template `<TEMPLATE_NAME>`.
- For dashboard alert – An error icon is displayed in the template header. If you mouse over the error icon, you will see error message explaining the problem. If you are using the Dashboard Builder at that time, you can fix the problem immediately.

9.5.2. Missed monitoring jobs

If a monitor check is missed,because server was not running, it will run immediately after the server is started. Each missed task will only run once, even when more checks were missed. After the missed monitoring jobs were run, checks will be performed as before according to their schedule.