



# PDF Help

This PDF Help contains the complete documentation for Tableau. View this document electronically to take advantage of search and bookmarks. You can also print this document for hard copy reference. Please download the updated PDF Help regularly.

# Getting Started

---

- What is Tableau? . . . . . 1-1**
  - What Can Tableau Do? . . . . . 1-3
  - What Data Can Tableau Analyze? . . . . . 1-7
  - How Does Tableau Work? . . . . . 1-9
- The Tableau Environment . . . . . 2-1**
  - Opening and Closing Tableau . . . . . 2-3
  - The Tableau Interface . . . . . 2-5
  - Workbooks and Sheets . . . . . 2-16
  - Files and Folders . . . . . 2-21
- Learning to Use Tableau . . . . . 3-1**
  - Open Tableau . . . . . 3-3
  - Connect to Data . . . . . 3-5
  - Build Data Views (Manual) . . . . . 3-9
  - Build Data Views (Automatic “Show Me!”) . . . . . 3-29
  - Save Your Work . . . . . 3-37

---

# Connect to Data

---

<b>Basic Connection</b>	<b>4-1</b>
Supported Data Sources	4-3
How to Connect to a Data Source	4-6
Examples – Connecting to Data Sources	4-11
Connecting to a Custom SQL Query	4-65
Editing the Connection	4-66
Exporting the Connection	4-68
Refreshing the Data	4-70
Closing the Connection	4-71
 <b>Joining Tables</b>	 <b>5-1</b>
Adding Tables	5-2
Editing Tables	5-2
 <b>Extracting Data</b>	 <b>6-1</b>
Creating an Extract	6-3
Using Extracts	6-4
 <b>Managing Queries</b>	 <b>7-1</b>
Automatic Updates	7-3
Cancel Query	7-4
Abandoned Queries	7-5
Precision Warnings	7-6
 <b>Understanding Data Fields</b>	 <b>8-1</b>
Understanding the Data Window	8-3
Data Window Features and Functions	8-8
Renaming Fields and Dimension Members	8-11
Data Types & Roles	8-20

# Building Data Views

<b>Parts of the View</b>	<b>9-1</b>
Headers	9-3
Axes	9-6
Panels	9-9
Cells	9-10
Marks	9-11
Titles	9-34
Captions	9-35
Field Labels	9-36
Legends	9-38
 <b>Building Views Manually</b>	 <b>10-1</b>
Placing Fields on Shelves	10-3
Placing Fields on the Worksheet	10-6
Types of Shelves	10-10
Multidimensional Hierarchies	10-33
Example – Building Data Views Manually	10-44
 <b>Building Views Automatically</b>	 <b>11-1</b>
Automatic Views: Show Me!	11-3
Add to Sheet: Double-Click	11-8
 <b>Sorting, Filtering, and Grouping</b>	 <b>12-1</b>
Sorting	12-3
Filtering	12-19
Ad-Hoc Groups	12-62
Sets	12-70
 <b>Dates and Times</b>	 <b>13-1</b>
Changing Date Levels	13-4
Fiscal Dates	13-8
Perfect Pivoting with Dates	13-10
Continuous Dates	13-12



<b>Inspecting Data . . . . .</b>	<b>14-1</b>
Navigation Tools . . . . .	14-3
Undo and Redo . . . . .	14-8
Color Legend Highlighting . . . . .	14-9
Drop Lines . . . . .	14-12
Reference Lines . . . . .	14-14
Summary Card . . . . .	14-20
Viewing Underlying Data (Drill-Through) . . . . .	14-21
Describing the View . . . . .	14-24
 <b>Hyperlinks . . . . .</b>	 <b>15-1</b>
Adding URL Links . . . . .	15-3
Adding Sheet Links . . . . .	15-6
Launching Links . . . . .	15-9
Using Field Values in Hyperlinks . . . . .	15-10
Managing Links . . . . .	15-14
Example: Adding Hyperlinks . . . . .	15-16

## Advanced Analysis

<b>Calculations . . . . .</b>	<b>16-1</b>
Aggregations . . . . .	16-4
Calculated Fields . . . . .	16-18
Table Calculations . . . . .	16-40
Binned Data . . . . .	16-59
Totals . . . . .	16-63
Percentages . . . . .	16-74
 <b>Background Images . . . . .</b>	 <b>17-1</b>
Adding Background Images . . . . .	17-3
Setting up the View . . . . .	17-6
Managing Background Images . . . . .	17-8
 <b>Trend Lines and Statistics . . . . .</b>	 <b>18-1</b>
Adding Trend Lines . . . . .	18-4
The Trend Line Model . . . . .	18-8

Removing Factors from the Model . . . . .	18-9
Testing Significance . . . . .	18-13
Trend Lines Example . . . . .	18-16
Assumptions . . . . .	18-20
Trend Line Model Terms . . . . .	18-21
Commonly Asked Questions . . . . .	18-23

<b>Log Axes . . . . .</b>	<b>19-1</b>
---------------------------	-------------

## Publishing and Sharing

<b>Formatting . . . . .</b>	<b>20-1</b>
Worksheet Level Formatting . . . . .	20-4
Formatting Specific Parts of the View . . . . .	20-13
Copying and Pasting Formatting . . . . .	20-24
Clearing Formatting . . . . .	20-25
Workbook Themes . . . . .	20-26
Resizing the Table . . . . .	20-27
Miscellaneous Table Options . . . . .	20-35
Editing Axes . . . . .	20-37
Formatting Mark Colors . . . . .	20-49

<b>Annotations and Data Labels . . . . .</b>	<b>21-1</b>
Annotations . . . . .	21-3
Data Labels . . . . .	21-17

<b>Dashboards . . . . .</b>	<b>22-1</b>
Creating Dashboards . . . . .	22-3
Organizing Dashboards . . . . .	22-8
Understanding Dashboards and Worksheets . . . . .	22-13
Tips for Creating Useful Dashboards . . . . .	22-15

<b>Publishing to Tableau Server . . . . .</b>	<b>23-1</b>
Publishing Workbooks . . . . .	23-3
Publishing Data Sources . . . . .	23-14
Opening Workbooks from the Server . . . . .	23-18
Importing Data Sources from the Server . . . . .	23-21

User Filtering . . . . . 23-24

**Saving and Exporting . . . . . 24-1**

Saving Your Work . . . . . 24-3

Reverting Workbooks . . . . . 24-9

Exporting Your Work . . . . . 24-10

Exporting the Data Source Connection . . . . . 24-24

**Printing . . . . . 25-1**

Printing Results . . . . . 25-3

Publish as PDF . . . . . 25-7

Printing the Help . . . . . 25-8

## Reference

**Functions, Operators, & Data Types . . . . . 26-1**

Data Types . . . . . 26-3

Formatting Literals . . . . . 26-5

Functions . . . . . 26-6

Operators . . . . . 26-25

**Performance Tips . . . . . 27-1**

Performance Tips: All Data Sources . . . . . 27-3

Performance Tips: Relational Data Sources . . . . . 27-5

Performance Tips: Multidimensional Data Sources . . . . . 27-7

Speeding up Context Filters . . . . . 27-8

Extracting Large Text and Excel Files . . . . . 27-9

**Tips and Tricks . . . . . 28-1**

**Keyboard Shortcuts . . . . . 29-1**

**Glossary . . . . . 30-1**

**Index . . . . . 30-1**

## What is Tableau?

---

---

## Overview

Tableau is a visual analysis and reporting solution that allows you to explore and analyze databases with simple drag and drop operations. This includes small databases such as Excel and Access as well as large databases such as SQL Server and Oracle.

Tableau has brought visual analysis to leading companies from nearly every industry, from Fortune 500 multinationals to the US government. Our products are in use today by business intelligence professionals, marketers, analysts, finance departments, strategic planners, sales departments, supply-chain specialists, national security specialists, engineers, scientists, social scientists, and executives around the world.

This section discusses the following topics:

- What can Tableau do? – Discover the powerful functions of Tableau and how they can benefit you.
- What Data can Tableau Analyze? – Learn about the kinds of data sources supported by Tableau.
- How does Tableau work? – A simple overview of the steps it takes to build views in Tableau.

## What Can Tableau Do?


Imagine being able to explore any database using a free-form visual canvas that helps you answer questions. Using Tableau, you can build and interact with views of data. These views allow you to query, display, analyze, filter, sort, group, drill down, drill up, calculate, organize, summarize, and present data efficiently.

The various ways that Tableau can help you explore your data are discussed in more detail below.

### See and Understand

People need effective layouts of data to understand results, discover relationships, find patterns, locate outliers, uncover structure and summarize findings. How well can you see what is going on in your business?

Tableau lets you flip through hundreds of views of your databases in minutes. These views are fundamentally more useful for analysis than those provided by pre-canned reports and traditional dashboards. Tableau gives you interactive visual tables, picture-perfect data displays, side-by-side comparisons and graphic encodings using color, size and shape. Without any programming or training, users can see and understand databases like they have never been able to before.



### Pivot and Refine

Tableau is a system in which problem solving occurs in a tightly knit cycle of visual analysis. With no training, users can conduct fluid "Q&A" sessions with data. They simply point Tableau to a database and start building drag & drop views of data.

A first view not only provides an insight (for example, "Costs are increasing."), but also suggests additional questions (for example, "Which products are responsible for the increase?"). Users can immediately pivot the first view to reveal the answer to the second question. The true value of Tableau is that it helps people navigate a series of questions and answers using an insightful interface.



## Isolate and Extract

What if you could dive into databases and unearth hidden gems in seconds? The exceptions? The hidden clusters? Tableau users quickly uncover the "ah-ha" findings behind summaries and trends, because they can drill down, identify useful groups of information, and extract exactly the groups they want (e.g., a group of highly profitable customers; or bottlenecks in a process). Interesting groups of information stand out clearly when you adopt visual analysis. Even better, once you locate items of interest, you can select the data and export it to other applications with a few mouse clicks.



## Browse and Explore

Tableau is the world's leading exploratory browser for databases. A key productivity pattern in the analysis process is the ability to start with "big picture" summaries of data and then quickly focus on detailed areas of interest.

To conduct effective analysis, it is crucial for people to quickly change what data they are viewing and how it is being viewed. Tableau's flexible interface enables this free form exploration. Exploratory analysis is further supported with unlimited undo and redo, allowing people to surf their databases much like they surf the Web.



## Query and Summarize

Tableau accomplishes so much because it's an interactive query and analysis front-end for databases. Does Tableau connect to relational databases and files? Yes it does. Does Tableau connect to large OLAP data warehouses? Yes it does. With no knowledge of query languages like SQL or MDX, business people can produce analyses and reports that tap the value of their data.



## Monitor and Measure

Use Tableau to build analytical dashboards that compare information and track performance against goals. These dashboards can be based on multiple data sources. And they are fully interactive, allowing you to drill into and explore information directly from the dashboard. You can also apply common filters to all of the worksheets, allowing you to change the filter and watch an array of visual displays update simultaneously.



## Publish and Present

Imagine pasting Tableau's vivid multi-dimensional results into Web pages or Microsoft Office applications and sharing them with others. Our users have a reputation for producing



high-impact presentations that are easy to understand. If you believe in presenting fantastic visual summaries of data, Tableau is a crucial addition for your toolset.



## What Data Can Tableau Analyze?

Your data needs to be in a database format before you can analyze it with Tableau. Databases include relational databases and multidimensional OLAP databases. The specific databases your copy of Tableau can connect to depends on your purchase options.

- The Personal Edition of Tableau can connect to Access, Excel, text files, and Tableau Data Extract files.
- The Professional Edition can also connect to Firebird, Microsoft SQL Server, Microsoft Analysis Services, MySQL, PostgreSQL, IBM DB2, Netezza, Oracle, Hyperion Essbase, and IBM OLAP Server.

Refer to the Products page of the Tableau Web site for more information on specific data source compatibility requirements.

To see which data sources your copy of Tableau can connect to, select **Data > Connect to Data**. Any data source that is not supported by your version of Tableau is greyed out (contact Tableau to upgrade your database accessibility options). Refer to “Supported Data Sources” on page 4-3 to learn more.

### Relational Databases

Relational databases store data in tables. Tables consist of rows and columns. For example, a relational database stored in a Microsoft Excel worksheet named **Sales** is shown below.

Column

	A	B	C	D	E	F
1	Order Date	Customer Name	Region	Sales Total	Order Priority	Discount
2	11/19/97	Speros Goranitis	West	27176.81	4-NOT SPECIFIED	0.03
3	11/19/97	Speros Goranitis	West	1363.14	4-NOT SPECIFIED	0.04
4	11/19/97	Speros Goranitis	East	139.87	4-NOT SPECIFIED	0.07
5	11/19/97	Rick Reed	North	12633.4	2-HIGH	0.01
6	11/19/97	Rick Reed	West	12543.76	2-HIGH	0
7	11/19/97	Rick Reed	South	1414.12	2-HIGH	0.02
8	11/19/97	Rick Reed	South	5316.62	2-HIGH	0.08
9	11/19/97	Matthew Fleese	West	5367.26	1-URGENT	0.07
10	11/19/97	Matthew Fleese	North	18532.99	1-URGENT	0.07
11	11/19/97	Matthew Fleese	East	12.67	1-URGENT	0.03
12	11/19/97	Matthew Fleese	West	146.1	1-URGENT	0.06
13	11/19/97	Rick Reed	West	241.05	3-MEDIUM	0.08
14	11/19/97	Rick Reed	East	880.12	3-MEDIUM	0.08
15	11/19/97	Rick Reed	North	263.72	3-MEDIUM	0.1
16	11/19/97	Rick Reed	West	492.92	3-MEDIUM	0.02
17	11/19/97	Rick Reed	South	293	3-MEDIUM	0.06
18	11/21/97	Jim Sink	South	1002.58	2-HIGH	0.1
19	11/21/97	Jim Sink	West	1356.56	2-HIGH	0.04
20	11/21/97	Jim Sink	North	42.45	2-HIGH	0.03
21	11/21/97	Jim Sink	East	60.16	2-HIGH	0.07
22	11/21/97	Dario Medina	West	182.68	3-MEDIUM	0.05

Row

Each row represents a sales record. The columns hold specific information about every customer, such as the order date, customer name, geographic region, and so on. You can connect Tableau to any relational table like this and begin analyzing the data.

If your data is stored in multiple tables, you must relate (join) the data before analyzing it with Tableau. The easiest way to do this is to select the Multiple Tables option in the Connect Dialog box. From there you can define a join to as many tables as necessary. Refer to “Connecting to Multiple Tables” on page 5-2 to learn more. All major relational databases support views.

## **Multidimensional OLAP Database**

OLAP databases like Microsoft Analysis Services or Hyperion Essbase store data in multidimensional data cubes. Tableau can connect to these cubes directly.

## How Does Tableau Work?

Tableau is a visual analysis and reporting tool for databases. But you don't need to know anything about databases to use Tableau. In fact, Tableau is designed to allow business people with no technical training to analyze their data efficiently.

Tableau is based on three simple concepts:

- 1 **Connect** – Connect Tableau to any database you want to analyze. Note that Tableau does not import the data. Instead, it queries to the database directly. Refer to “Supported Data Sources” on page 4-3 to learn more.
- 2 **Analyze** – Analyzing data means viewing it, filtering it, sorting it, performing calculations on it, reorganizing it, summarizing it, and so on.

Using Tableau, you can do all of these things by simply arranging the fields of your data source on Tableau worksheets. When you drop a field on a worksheet, Tableau queries the data using standard drivers and query languages (like SQL and MDX) and presents a visual analysis of the data. Refer to “Learning to Use Tableau” on page 3-1 to learn many of the basic steps involved when analyzing your data. You can find more detailed information in “Building Views Manually” on page 10-1 and “Calculations” on page 16-1.

- 3 **Share** – You can share results with others either by sharing Tableau workbooks with other Tableau users, or by simply pasting results into other applications such as Microsoft Office. In addition, you can export the views as an image or publish to a PDF. Refer to “Saving and Exporting” on page 24-1 for more information.

# The Tableau Environment

---

<b>Overview</b>	<b>2-2</b>
<b>Opening and Closing Tableau</b>	<b>2-3</b>
<b>The Tableau Interface</b>	<b>2-5</b>
Data Window	2-6
Toolbars	2-8
Status Bar	2-10
Cards and Shelves	2-11
Rearranging Interface Components	2-14
<b>Workbooks and Sheets</b>	<b>2-16</b>
Workbooks	2-16
Worksheets and Dashboards	2-17
<b>Files and Folders</b>	<b>2-21</b>
Changing the Repository Location	2-23

---

## Overview

The purpose of this section is to provide you with an overview of the Tableau environment. The topics discussed are:

- Opening and Closing Tableau – Open and quit the application.
- The Tableau Interface – Understand the main components of the Tableau interface.
- Workbooks and Worksheets – Workbooks hold one or more worksheets. Worksheets hold views you create with your data.
- Files and Folders – Learn about the Tableau repository.

Refer to Chapter 3, “Learning to Use Tableau” if you want to learn how to use Tableau by following a simple example.

## Opening and Closing Tableau

You can open Tableau in one of four ways:

- Double-click the Tableau icon on your desktop.



Tableau 3.5

- Select **Tableau** on the Windows **Start** menu.
- Open a Tableau file.

Double-click on a Tableau workbook or bookmark file to open the file in Tableau. Tableau files are typically stored in the My Tableau Repository folder of your My Documents folder. Refer to “Files and Folders” on page 2-21 for more information.

- Drag a supported data file onto the Tableau icon or the application window.  
Drag an Excel or Access file onto the Tableau icon or application window to open Tableau and create a connection to the data. For other supported data sources, you must open Tableau first and then create a connection. Refer to “Supported Data Sources” on page 4-3 for a complete list of supported data sources.

If you start Tableau by double-clicking the desktop icon or the Windows **start** menu, the application window contains an empty workbook. Otherwise, the Data window shows the data fields in the data source. Refer to “The Tableau Interface” on page 2-5 for information about Tableau’s interface features.

---


**Note** Opening a file directly from a Web site or other remote source can cause the data source to open with an empty Data window. To avoid this behavior, copy the file to your local hard drive.

---

You can close Tableau one of these two ways:

- Click the close icon  located on right corner of the application title bar.
- Select the **File > Exit** menu item.

If your workbook has not been saved, the following dialog box opens.



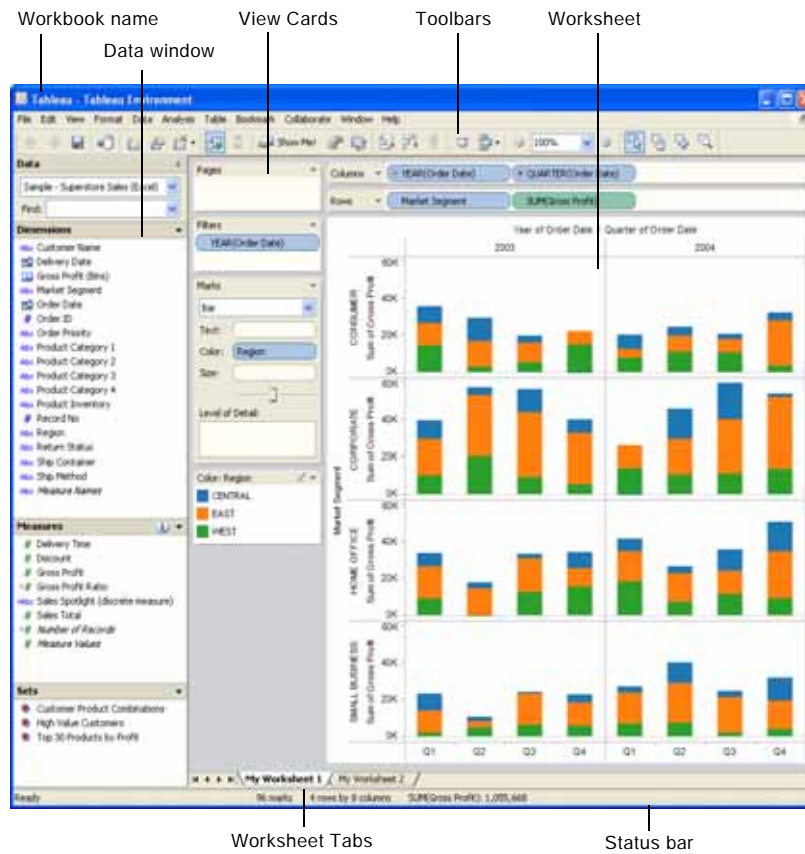


## The Tableau Interface


The Tableau interface consists of menus, a toolbar, the Data window, view cards that can hold one or more shelves, a workbook, worksheets, and dashboards as shown in the figure below.


Worksheets contain shelves such as **Rows** and **Columns**, which are where you drag data fields to build views. When you start Tableau, the interface displays with a default layout. You can change the layout to suit your needs, including resizing, moving, and closing interface components.

Dashboards contain views, legends, and quick filters. When you first create a dashboard, the Dashboard is empty and all of the worksheets in the workbook display in the Dashboard window.



## Data Window

Data fields appear on the left side of the interface in the Data window. You can hide and show the data window by selecting the **View > Data Window** menu item. You can also click the minimize button  in the upper right corner of the Data window.



You can search for fields in the Data window by typing in the box labeled **Find** (located at the top of the Data window) or by pressing **Ctrl + F** on your keyboard. Refer to “Find Fields” on page 8-8 to learn more about finding fields in the Data window.

You can right-click the fields in the Data window to access important commands. For example, hide, rename, and group fields.

Refer to “Understanding the Data Window” on page 8-3 for more information about the Data window.

## Toolbars













Tableau's toolbar contains commands such as Connect to Data, Show Me!, and Save. In addition, the toolbar contains analysis and navigation tools such as sort, group, and zoom. You can undock the toolbar by grabbing its left edge and then dragging it to a new location. You can hide or display a toolbar by selecting **View > Toolbar**.

### Tableau Toolbar



The toolbar helps you quickly access common tools and actions. The table below explains the functions of each of the toolbar buttons.


	<b>Undo:</b> undoes the last task you completed.
	<b>Redo:</b> repeats the last task you canceled with the Undo button.
	<b>Save:</b> saves changes made to the workbook.
	<b>Connect to Data:</b> opens a dialog box where you can create a new connection or select one from your repository.
	<b>New Sheet:</b> creates a new blank worksheet.
	<b>Duplicate Sheet:</b> creates a new worksheet containing the exact same view as the current sheet.
	<b>Clear:</b> clears the current worksheet. Use the drop-down list to clear specific parts of the view such as filters, formatting, and sizing.
	<b>Automatic Updates:</b> controls whether Tableau automatically updates the view when changes are made.
	<b>Run Update:</b> runs a manual query of the data to update the view with changes when automatic updates are turned off.
	<b>Show Me!:</b> automatically chooses the best view of the selected data fields.

	<b>Show Me! Alternatives:</b> displays up to fourteen alternate views of the data, in addition to the best view selected by Show Me! The options presented when this button is clicked depend on the data fields that have been selected.
	<b>Swap:</b> moves the fields on the Row shelf to the Column shelf and vice versa. The Hide Empty Rows and Hide Empty Columns settings are also swapped with this button.
	<b>Sort Ascending:</b> applies a manual sort in ascending order of a selected field based on the measures in the view.
	<b>Sort Descending:</b> applies a manual sort in descending order of a selected field based on the measures in the view.
	<b>Group Members:</b> creates an ad-hoc group by combining selected values.
	<b>Presentation Mode:</b> toggles between showing and hiding everything but the view.
	<b>View Cards:</b> shows and hides the specified cards in a worksheet. Select the cards you want to hide or show from a drop-down list.
	<b>Zoom Selector:</b> increases the size of the entire view by a specified percentage. Use the plus and minus buttons to zoom incrementally or select a percentage from the drop-down list.
	<b>Select:</b> selects data in a view. Click on a specific mark in the view or drag the select tool around several marks.
	<b>Pan:</b> moves your view of a table up and down as well as left and right.
	<b>Focus:</b> focuses on a specific part of a view. Drag the focus tool around the area you want to focus on. The focus tool is limited to one pane at a time.
	<b>Zoom:</b> increases the size of the view so you can examine of specific areas of the view up close. Hold down the CTRL key to zoom out.


## Status Bar

The status bar is located at the bottom of the Tableau workbook. It displays descriptions of menu items as well as information about the current view.

An example of a menu item description is shown below. This message displays when you select **File > Open**.





An example of information about the view is shown below. In this particular view there are 1023 rows and 3 columns.



You can hide the status bar by selecting **View > Status Bar**.

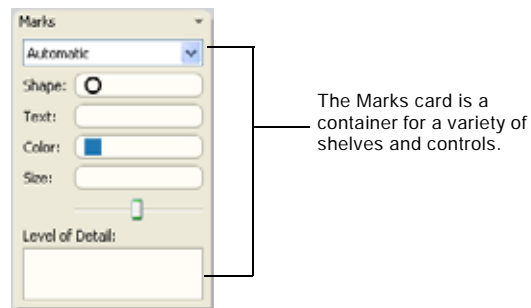
## Status Bar Warnings

Occasionally, Tableau will display warning icons in the bottom right corner of the status bar to indicate errors that have or may occur. Below are the possible warning icons and what they mean.

	<p><b>Cancel Query Indicator:</b> When you cancel multiple queries, an indicator appears to show you how many queries are still running on the database and using resources. For more information about this warning refer to “Abandoned Queries” on page 7-5.</p>
	<p><b>Precision Warning:</b> Some fields are more precise in the database than Tableau can model. When you add these fields to the view a precision warning is displayed in the status bar. For more information about this warning refer to “Precision Warnings” on page 7-6.</p>

## Cards and Shelves

Every worksheet contains a variety of different cards that can be turned on and off. Cards are containers for shelves, legends, and other controls. For example, the Marks card contains the mark selector, the size slider, and the shape, text, color, size, and level of detail shelves.



Cards can be shown and hidden as well as rearranged around the worksheet. Refer to “Rearranging Interface Components” on page 2-14 to learn more about moving the cards.


The following list describes each card and its contents. For more information about each shelf refer to “Types of Shelves” on page 10-10.

- **Columns Shelf** – contains the Columns shelf where you can drag fields to add columns to the view.
- **Rows Shelf** – contains the Rows shelf where you can drag fields to add rows to the view.
- **Pages Shelf** – contains the Pages shelf where you can create several different pages with respect to the members in a dimension or the values in a measure.
- **Filters Shelf** – contains the Filters shelf; use this shelf to specify the values to include in the view.
- **Measure Names/Values Shelf** – contains the Measure Names shelf; use this shelf to use multiple measures along a single axis.
- **Color Legend** – contains the legend for the color encodings in the view and is only available when there is a field on the Color shelf.
- **Shape Legend** – contains the legend for the shape encodings in the view and is only available when there is a field on the Shape shelf.
- **Size Legend** – contains the legend for the size encodings in the view and is only available when there is a field on the Size shelf.

- **Quick Filters** – a separate quick filter card is available for every field in the view. Use these cards to easily include and exclude values from the view without having to open the Filter dialog box.
- **Marks** – contains a mark selector where you can specify the mark type as well as the Path, Shape, Text, Color, Size, and Level of Detail shelves. The availability of these shelves are dependent on the fields in the view.
- **Title** – contains the title for the view. Double-click this card to modify the title.
- **Caption** – contains a caption that describes the view. Double-click this card to modify the caption.
- **Summary** – contains summary of each of the measures in the view including the Min, Max, Sum, and Average.
- **Current Page** – contains the playback controls for the Pages shelf and indicates the current page that is displayed. This card is only available when there is a field on the Pages shelf.



Each card has a menu that contains common controls that apply to the contents of the card. For example you can use the card menu to show and hide the card. Access the card menu by clicking on the arrow in the upper right corner of the card.




## Rearranging Interface Components

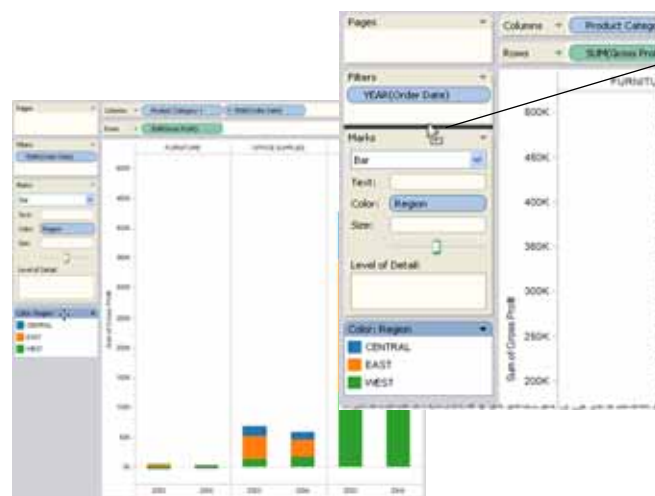
You can rearrange and hide various interface elements including cards and other worksheet elements. You can rearrange and hide the following:

- Cards
- Worksheet elements

### Rearranging Cards

A worksheet contains several cards that are containers for shelves, legends, and other controls (refer to “Cards and Shelves” on page 2-11). Each card can be rearranged to create a custom workspace.

To move a card, point the cursor at the title area of the card you want to move. When the cursor becomes the move symbol , click and drag the card to a new position. As you drag the card around the worksheet, the possible positions for it are highlighted with a black bar.



The black line indicates a possible place to drop the card.


**Note** You can restore the worksheet windows to their default state by selecting **View > Reset Cards**.

## Showing and hiding worksheet elements


Just about every element in the worksheet can be turned on and off so you don't have to clutter the worksheet with unnecessary elements.

### To show or hide the Data window, toolbar, and the status bar:

- Select **View** and then select the element you want to hide.



### To show or hide a card:


- Click **View Cards**  on the toolbar and then select the card you want to show or hide from the drop-down list. You can also show or hide cards using the **View** menu.

**Note** You can restore the cards to their default state by selecting **Reset Cards** on either the **View** menu or the **View Cards** drop-down menu on the toolbar.

## Presentation Mode

Sometimes you may want to use Tableau for presenting your findings. Rather than hiding each individual worksheet element one at a time, you can switch into Presentation Mode. Presentation Mode hides all worksheet elements except for the view and its associated legends and quick filters.

### To turn on Presentation Mode:

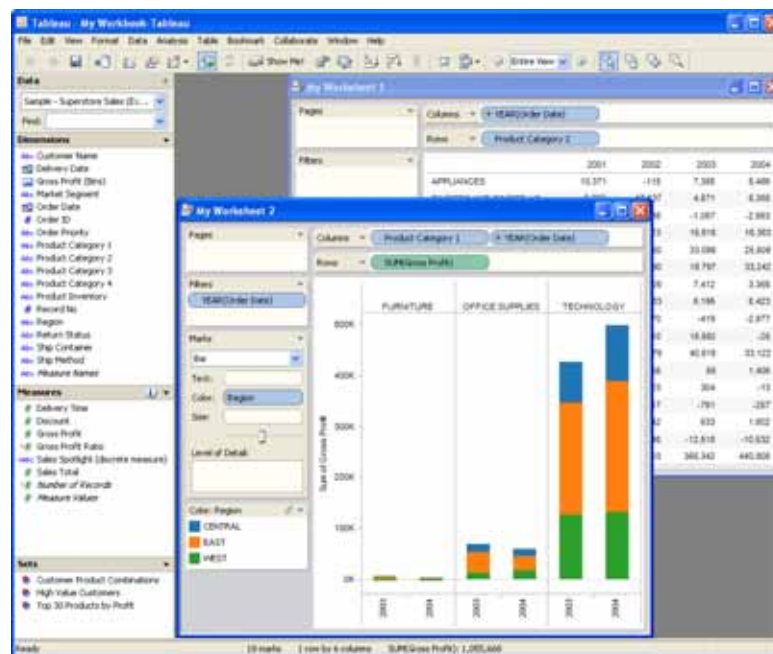
- Click the **Presentation Mode**  button in the toolbar or select **Window > Presentation Mode**.

## Workbooks and Sheets

Tableau uses a workbook and worksheet file structure, much like Microsoft Excel.

- **Workbooks** – hold the work you create and consists of one or more worksheets and dashboards. You can open multiple workbooks simultaneously.
- **Worksheets and Dashboards** – contain a particular view or dashboard of your data. A workbook can contain multiple worksheets, and each worksheet can be connected to a different data source.

As shown below, the workbook titled **MyWorkbook** consists of two worksheets offering two different views of the same data.



## Workbooks

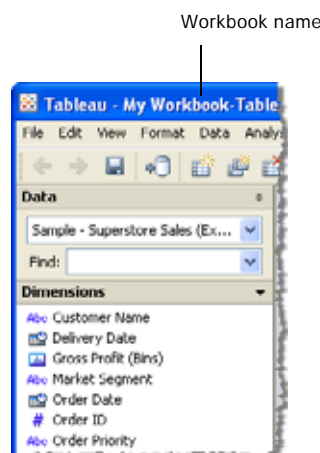
Tableau workbook files are much like Microsoft Excel workbooks. They contain one or more worksheets or dashboards and hold all of your work. They allow you to organize, save, and share your results.

When you open Tableau, a blank workbook is automatically created. You can also create a new workbook by selecting **File > New** or pressing **Ctrl+N**. You can open an existing workbook several different ways.

**To open an existing workbook:**

- Select **File > Open** and navigate to the location of your workbook using the **Open** dialog box. Tableau workbooks have the .twb file extension.
- Double-click on any workbook file.
- Drag any workbook file onto the Tableau desktop icon or onto the running Tableau application.

The workbook name is displayed in Tableau's title bar.



You can open multiple workbooks simultaneously. Each workbook is shown in its own window.

Refer to “Saving and Exporting” on page 24-1 for information about saving a workbook.

## Worksheets and Dashboards



Each workbook can contain worksheets and dashboards. Build views of data by placing fields on the shelves in a worksheet then combine several worksheet views into a dashboard where you can present and monitor data.


Each worksheet in a given workbook can be connected to a different data source and combined into a single dashboard. For example, you can connect to an Excel workbook, an Access database, and a SQL Server database in a single workbook by creating a new worksheet for each source. Then you can combine views from each of these worksheets into a single dashboard sheet.

This section discusses the following topics:

- Undo, Redo, and Clearing Sheets
- Creating and Opening Sheets
- Duplicating Sheets
- Hiding Worksheets
- Deleting Sheets

### Undo, Redo, and Clearing Sheets

Every Tableau workbook contains a history of the steps you have performed on the worksheets or dashboards. To move backward through the history, click **Undo**  on the toolbar or press **Ctrl+Z**. Similarly, move forward through the history by clicking **Redo**  on the toolbar or pressing **Ctrl+Y**.

You can remove all fields, formatting, sizing, axis ranges, filters, and sorts in the sheet by clicking **Clear**  on the toolbar. You can also use the Clear drop-down list on the toolbar to clear specific aspects of the view such as clear all formatting, sizing, filters, or sorts.

---

**Note** Using the clear commands on the toolbar does not clear the history. If you decide that you didn't want to clear the sheet, click the **Undo** button.

---

### Creating and Opening Sheets

To create a new worksheet, select **Edit > New Worksheet** or press **Ctrl+M**. Create a new dashboard by selecting **Edit > New Dashboard**. Tableau inserts the new sheet into the current workbook.

Tableau automatically generates sheet names. The first worksheet is named Sheet 1, the second worksheet is named Sheet 2, and so on. You can rename a sheet by selecting **Edit > Rename Sheet**. Alternatively, double click the name of the sheet on the sheet tab and type a new name.

You can navigate between different worksheets by using the controls in the lower left corner of the workbook and by selecting an appropriate sheet tabs.



You can also navigate between sheets using the **Window** menu, or move through multiple worksheets by pressing **Ctrl+F6**.

### Duplicating Sheets

Duplicating a sheet allows you to easily make a copy of a worksheet or dashboard. You can then modify the view without losing the original version. To duplicate the active sheet, select **Edit > Duplicate Sheet**.

**Duplicate As Cross-tab.** A cross-tab (sometimes referred to as a Pivot Table) is a table that summarizes data in rows and columns of text. It is a convenient way to display the numbers associated with a data view.

In Tableau, you can quickly create a cross-tab from a worksheet by selecting **Edit > Duplicate as Cross-tab**. This command inserts a new worksheet into your workbook and populates the sheet with a cross-tab view of the data from the original worksheet. Dashboard sheets cannot be duplicated as cross tabs.

There are other ways to see the numbers behind data views. For example, you can mouse-over any mark to display the associated numbers in a tooltip. You can also right-click a selection of marks and select View Underlying Data (refer to “Viewing Underlying Data (Drill-Through)” on page 14-21). Finally, you can copy and paste the data into Excel. Refer to “Exporting Your Work” on page 24-10 to learn more.

### Hiding Worksheets

A worksheet that is used in a dashboard cannot be deleted, but it can be hidden. You may want to hide a worksheet if you are sharing the dashboard with others and don’t want to clutter the workbook with all of the supporting worksheets.

You can hide the worksheets that are used in a dashboard by right-clicking the worksheet tab and selecting **Hide Sheet**. Keep in mind that someone viewing the dashboard can still access the hidden worksheet. Refer to “Understanding Dashboards and Worksheets” on page 22-13 to learn more about navigating between dashboards and worksheets.

### Deleting Sheets

Deleting a sheet removes it from the workbook. You can delete the active worksheet by selecting **Edit > Delete Sheet**. Alternatively, right-click the worksheet or dashboard tab along the bottom of the workbook and select **Delete Sheet**. Worksheets used in a dashboard cannot be deleted, rather you can hide the worksheet. Refer to “Hiding Worksheets” on page 2-19 to learn more.

---

**Note** There must always be at least one worksheet or dashboard in a workbook.


---




## Files and Folders

You can save your work using several different Tableau specific file types: workbooks, bookmarks, packaged data files, data extracts, and data connection files. Each of these file types are described below.


### Workbooks

Tableau workbook files have the .twb file extension and are marked with the workbook icon . Workbooks hold one or more worksheets and dashboards. Refer to “Workbooks” on page 24-3 to learn more about workbooks.


### Bookmarks

Tableau bookmark files have the .tbm file extension and are marked with the bookmark icon . Bookmarks contain a single worksheet and are an easy way to quickly share your work. Refer to “Bookmarks” on page 24-7 to learn more about bookmarks.


### Packaged Workbooks

Tableau packaged workbooks have the .twbx file extension and are marked with the packaged workbook icon . Packaged workbooks contain a workbook along with any supporting local file data sources and background images. This format is the best way to package your work for sharing with others who don’t have access to the data. Refer to “Packaged Workbooks” on page 24-5 to learn more.

### Data Extract Files

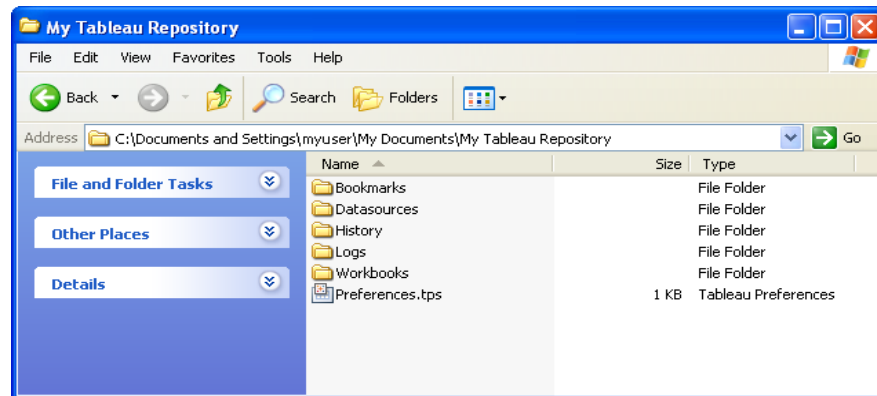
Tableau data extract files have the .tde file extension and are marked with the extract icon . Extract files are a local copy of a subset or entire data source that you can use to share data, work offline, and improve database performance. Refer to “Extracting Data” on page 6-1 to learn more about extracting data.

### Data Connection Files

Tableau data connection files have the .tds file extension and are marked with the data connection icon . Data connection files are shortcuts for quickly connecting to data sources that you use often. Refer to “Exporting the Connection” on page 4-68 to learn more.

These files can be saved in the associated folders in the My Tableau Repository directory, which is automatically created in your My Documents folder when you install Tableau.

Your work files can also be saved in other locations, such as your desktop or a network directory.



## Changing the Repository Location

You can specify a new location for the Tableau repository if you are not using the default location in your Documents folder. For instance, if you are required to have your data on a network server instead of on your local machine, you can point Tableau at the remote repository.

To specify the repository location, follow these three steps:

- 1 Select **File > Repository Location**.
- 2 Select a new folder that will act as the new repository location in the **Select a Repository** dialog box.
- 3 Restart Tableau so that it uses the new repository.

---

**Note** Changing the repository location does not move the files contained in the original repository, rather it creates a new repository where you can store your files.

---

# Learning to Use Tableau

---

<b>Overview</b>	<b>3-2</b>
<b>Open Tableau</b>	<b>3-3</b>
<b>Connect to Data</b>	<b>3-5</b>
<b>Build Data Views (Manual)</b>	<b>3-9</b>
Example 1 - Yearly Profits	3-10
Example 2 - Quarterly Profits by Year	3-16
Example 3 - Quarterly Profits by Year and Market	3-18
Example 4 - Filter by Year and by Sales	3-20
Example 5 - Color Encode by Region	3-26
<b>Build Data Views (Automatic “Show Me!”)</b>	<b>3-29</b>
Yearly Profits	3-29
Sales versus Profit by Product and Customer	3-33
<b>Save Your Work</b>	<b>3-37</b>

---

## Overview

The purpose of this section is to get you started with Tableau by presenting a simple example. The example takes you through all the basic steps you would use for your own work. These steps are:

- 1 Open Tableau – Open the application.
- 2 Connect to Data – Connect to the Superstore Sales Excel data, which is included with Tableau.
- 3 Build Data Views (Manual) – Build data views by placing fields on shelves.
- 4 Build Data Views (Automatic “Show Me!”) – Build data views using Tableau’s automatic building tools. Simply select fields in the Data window and let Tableau intelligently arrange them in a relevant view.
- 5 Save Your Work – When you have finished building data views, you should save your work as a Tableau workbook.

After reading this section, you will be able to build data views using your own relational or multidimensional data.

## Open Tableau

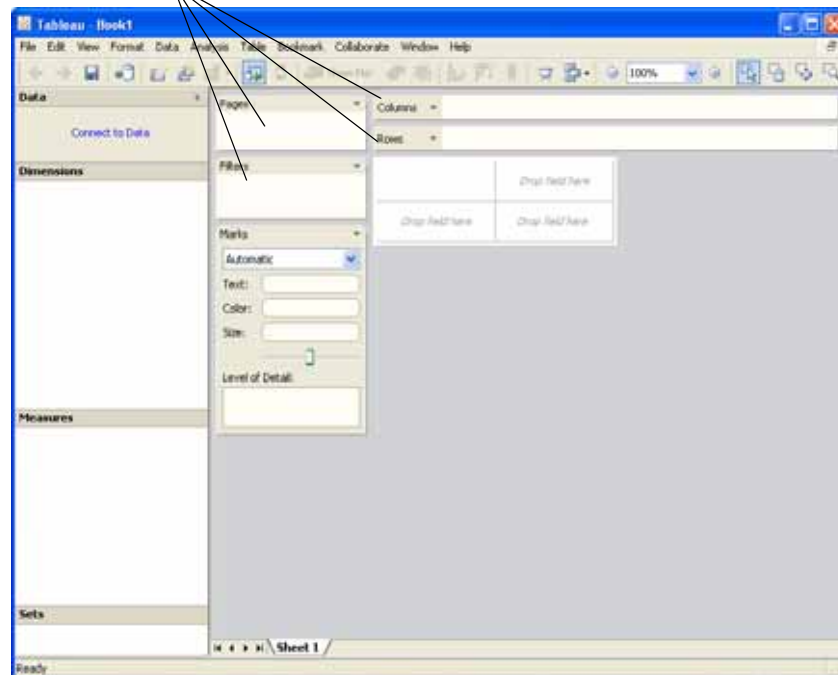
To open Tableau, on the Windows **Start** menu select **All Programs > Tableau Software > Tableau 3.5**, or double-click the desktop icon.



Tableau 3.5

Tableau opens with an empty workbook and worksheet. As shown below, the worksheet contains several shelves.

Shelves are part of every Tableau worksheet.



Shelves are where you place the fields that make up your data view. They allow you to create the rows and columns of a table, to exclude data from the view by filtering, to display additional levels of detail, and to encode the data by color, size, shape.

When you place the fields on the shelves you are generating a query of the database. For more information about queries refer to “Managing Queries” on page 7-1.

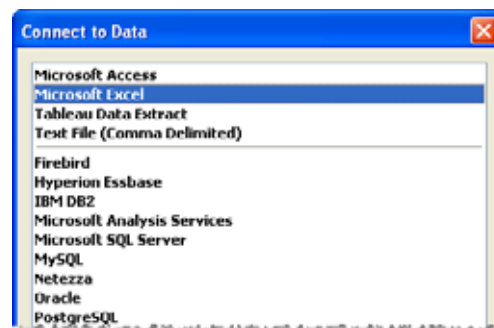
Each workbook can contain multiple worksheets. You can have multiple workbooks open at the same time. Each workbook displays in its own window.

## Connect to Data

To build views, you must first connect Tableau to data. In this example, you will connect to the Superstore Sales Excel worksheet that is installed with Tableau.

**To connect to data:**

- 1 Select **Data > Connect to Data**.
- 2 Select the connection type. For this example, select **Microsoft Excel** and then click **OK**.

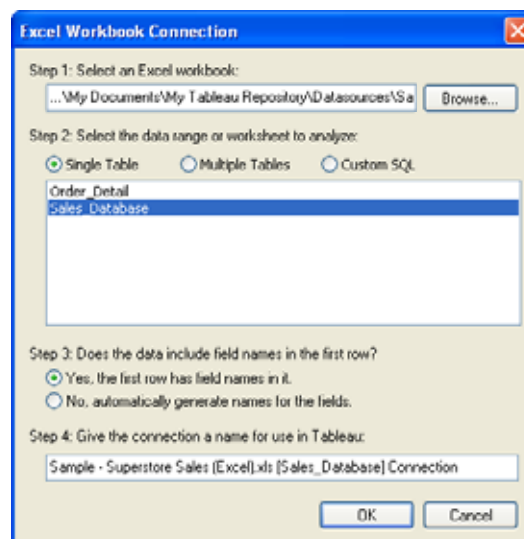





### 3 Complete the connection dialog.

Complete the **Excel Workbook Connection** dialog box as shown below. Note that the Excel data is located in the **Datasources** directory of the your Tableau Repository. By default, the Tableau Repository is created in your My Documents folder when Tableau is installed.

Click **Browse** to navigate to your Tableau repository and select Sample-Superstore Sales (Excel).xls.



When you click **OK**, Tableau asks you if you want to save a copy of the connection in the repository.



A copy of the connection is simply a shortcut to the data. The shortcut displays in the **Connect to Data** dialog the next time you open it. By selecting the shortcut, you can bypass the connection configuration in the future. Note that a shortcut already exists for this particular data.

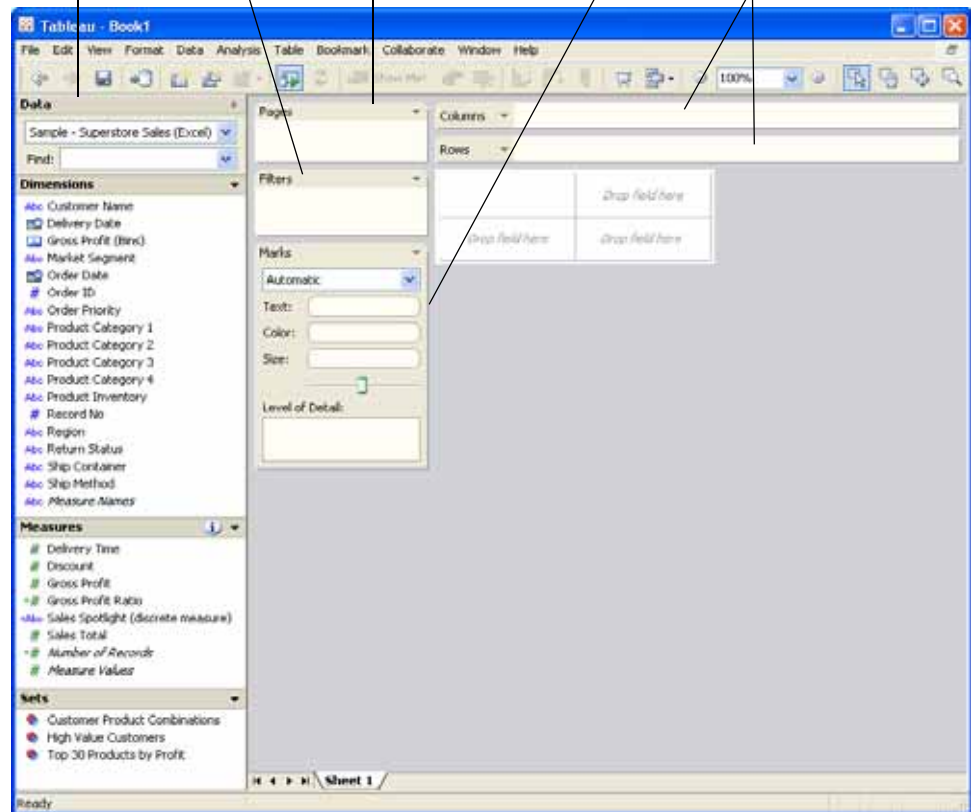
When the connection is established, the data fields show in the Data window and you can begin to build data views.

For the Excel workbook data, the fields are created from the columns of the Sales\_Database worksheet. Each column contains a unique attribute of the data such as customer name, sales total, product type, and so on.

The fields are organized into two sections: dimensions and measures. Dimensions typically hold categorical data such as product types and dates, while measures hold numeric data such as sales and profit.

The figure below shows the Data window as well as the shelves used in this example.

The Data window displays data fields. Exclude data by filtering. Separate your view into pages. Encode the marks using color, size, and shape. Create the columns and rows of the data view.



## Build Data Views (Manual)

You can build data views by dragging fields from the Data window and dropping them onto the shelves that are a part of every Tableau worksheet.

This section presents five data views using the Sample–Superstore Sales data that comes with Tableau. Each view builds upon its predecessor, and emphasizes one or two main features. The final view is a nested table that displays filtered data using stacked colored bars. You will build the following views in this section:

- Example 1 - Yearly Profits
- Example 2 - Quarterly Profits by Year
- Example 3 - Quarterly Profits by Year and Market
- Example 4 - Filter by Year and by Sales
- Example 5 - Color Encode by Region

After reading this section, you will have a basic understanding of important Tableau features and terminology such as placing fields on shelves, headers and axes, mark types, aggregation, filtering measures and dimensions, color-encoding a dimension, tooltips, and undo/redo.

All examples use the sample “Superstore” Excel data that comes with Tableau.

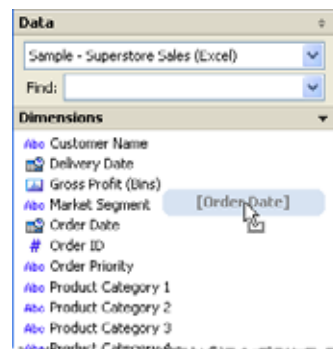
## Example 1 - Yearly Profits

The first view you will build is a bar chart that displays gross profit as a function of year. To create the view, follow the steps below:

- 1 Click and drag the **Order Date** field from the Data window to the view.

The **Order Date** field is located in the **Dimensions** area of the Data window because it contains date information, which are discrete values.

When you drag the field a transparent image follows the mouse cursor.



- 2 Place **Order Date** on the **Columns** shelf.

When you drag a field over a shelf, the ▼ icon indicates that the shelf can accept the field.



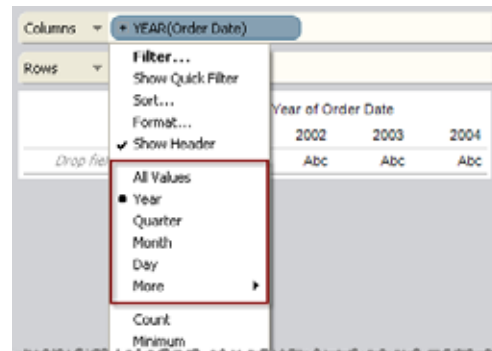
The resulting table has four columns and one row. Each column header represents a member of the Order Date field (2001, 2002, 2003, 2004). Each cell contains the label *Abc* , which indicates that the current mark type for this view is text.

The screenshot shows a data view interface. At the top, there is a 'Columns' dropdown menu with a blue button labeled 'YEAR(Order Date)'. Below it is a 'Rows' dropdown menu. The main area displays a table with the following structure:

	Year of Order Date			
	2001	2002	2003	2004
Drop field here	Abc	Abc	Abc	Abc


Notice that the field is colored blue indicating it is a dimension. Also, the field name changed to **YEAR(Order Date)** because year is the default date level for this field. The default date level is determined by the highest level that contains more than one distinct value (e.g., multiple years, multiple months, etc.). That means that if **Order Date**

contained data for only one year but had multiple months, the default level would be month. You can change the date level using the field's context menu.




- 3 Click and drag the **Gross Profit** field from the Data window.


The **Gross Profit** field is located in the **Measures** area of the Data window because it contains numeric data.



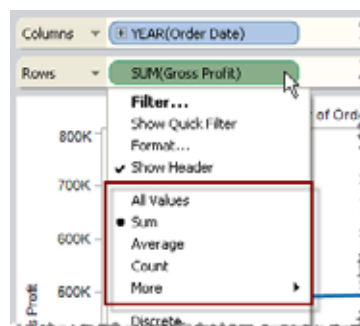
- 4 Place **Gross Profit** on the **Rows** shelf.



The table is automatically transformed into a line chart and a vertical axis is created for the measure.



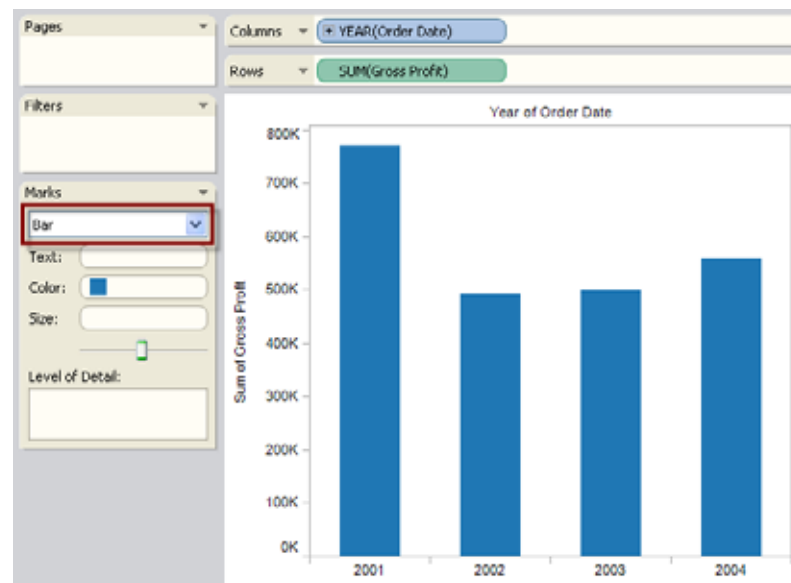
Notice that the Gross Profit field is colored green, indicating that it is a measure. Also, the field name changed to **SUM(Gross Profit)**. That's because by default, all measures for relational data are aggregated as a summation. You can change the aggregation using the field's context menu.






The data are automatically displayed using lines. Line charts are a great way to compare data over time, and allow you to visually compare data more effectively than just looking at the raw numbers.

In Tableau, lines are a type of mark. Line marks are used because the Mark Type setting is set to **Automatic** and the data view was built using one measure and one date dimension. You can easily select a different mark type on the **Mark** menu. For this example, change the mark type to **Bar**.



Each bar represents the summation of the gross profit for each year. For example, Tableau identifies all the data rows associated with the year 2001 and then sums all the corresponding gross profit values. Then it repeats this process for the remaining three years. You can display the value for each point in the line using tooltips.


To use tooltips, place the cursor over the mark of interest. For example, the sum of the gross profit for the year 2001 is \$769,935, which is significantly higher than the gross profit for the other years. You can also display the value by turning on data labels with **Format > Automatic Data Labels**.




## Example 2 - Quarterly Profits by Year

The next view we'll build is a nested bar chart that displays gross profits by quarter for several years. It builds upon the view created in "Yearly Profits" on page 3-29.


There are two ways to display quarters with years. One way is to drill down into **Order Date** by clicking the **+** control.



The other way is to drag the **Order Date** field to the **Columns** shelf again. Tableau automatically displays the dates using the next date level, which is quarters in this example.



The view is shown below.




The new dimension divides the view into four panes, where each pane contains the four quarters for a given year. This view is called a nested table because it displays multiple headers, with quarters nested within years.

You now have a view of the gross profit for every quarter of every year. As you can easily see, the most profitable quarter is Q4 of 2001 when the gross profit sum is \$233,618.


### Example 3 - Quarterly Profits by Year and Market

For this example, we'll again build a nested bar chart, but this time we'll add even another dimension so that we can analyze gross profit by market segment throughout each quarter and year. This example builds upon the view created in "Example 2 - Quarterly Profits by Year" on page 3-16.

- 1 Click and drag the **Market Segment** dimension from the Data window.




- 2 Place **Market Segment** on the **Rows** shelf.



Because you placed the field on the rows shelf, row headers are created. Each header represents a member of the Market Segment field. The new headers are located outside of (to the left of) the gross profit axes because **Market Segment** is to the left of **Gross Profit** on the **Rows** shelf.

Tableau does not allow you to place a dimension to the right of a measure on either the **Rows** or **Columns** shelves. That is, measures are always located inside dimensions in the data view.

The view is shown below.




The new dimension divides the view into 16 panes: one for each area in the table where the year and market segments intersect. This view is a more complex example of a nested table because it displays multiple row and column headers.

In the previous view, you saw that Q4 of 2001 is the most profitable quarter. In this view, you can see how each market performs in that quarter. You can see that the gross profit for the Corporate market in Q4 for the year 2001 is \$68,381.


## Example 4 - Filter by Year and by Sales

Now, we'll take the same view you built in "Example 2 - Quarterly Profits by Year" on page 3-16, but this time filter the data to only include sales from 2001 and 2004 and that are greater than \$2,000.

- 1 Click and drag the **Sales Total** measure from the Data window.

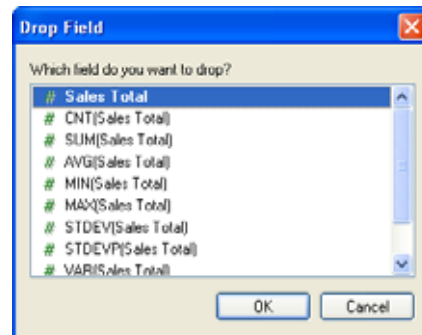


- 2 Place **Sales Total** on the **Filters** shelf.



When you place a measure on the Filter shelf, the **Drop Field** dialog box shown below automatically opens. You can choose to filter aggregated data using one of the standard aggregation functions, or you can choose to filter disaggregated data, which is the default choice. Refer to "Aggregate Calculations" on page 16-30 to learn more about

aggregating data. Note that aggregations do not apply to multidimensional data (because the data has been already aggregated on the server).




Filtering aggregated data means that the selected aggregation function (sum, average, and so on) is applied to the data and then it is filtered. Filtering disaggregated data means filtering individual data rows before any aggregation function is applied. In this example, filter **Sales Total** as a disaggregated measure.

- 3 After clicking **OK**, the **Filter** dialog box automatically opens. The dialog box displays the upper and lower limits of the measure. Because you are filtering disaggregated data, the limits reflect the minimum and maximum individual rows in the data. If you had



been filtering aggregated data, the limits would have been based on the selected aggregation.


Filter the data by typing a new value in the edit boxes or by moving the sliders. Type in 2000 for the lower limit and click the **OK** button.



Note that new axes are not created for the data view because the field does not add new rows or columns to the table. Only fields placed on the **Rows** or **Columns** shelves can add rows and columns. However, the data view can change to reflect the data rows excluded by the filter. Refer to “Filtering” on page 12-19 to learn more about filtering data.

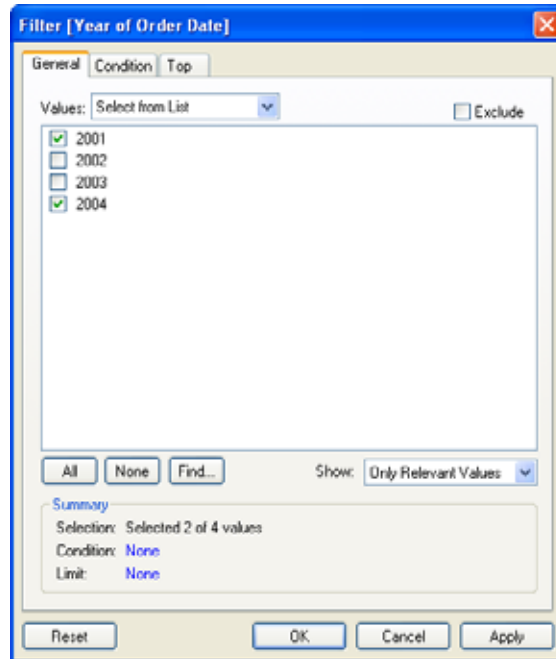
#### 4 Filter by **Order Date**.

Open the **Filter** dialog box by selecting **Filter** from the field’s context menu. You can also open the dialog box by double-clicking on the field when it is on a shelf.

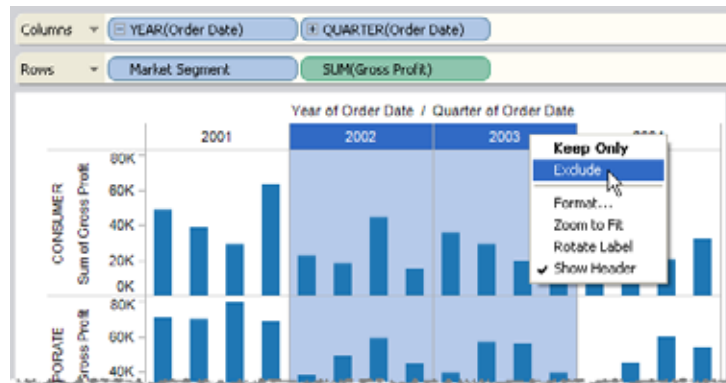


The dialog box below displays the dimension member names. By default, all members are selected. Deselect the years 2002 and 2003 and click **OK**. Note that when a field is filtered, it is automatically placed on the **Filters** shelf. Additionally, the filter removes

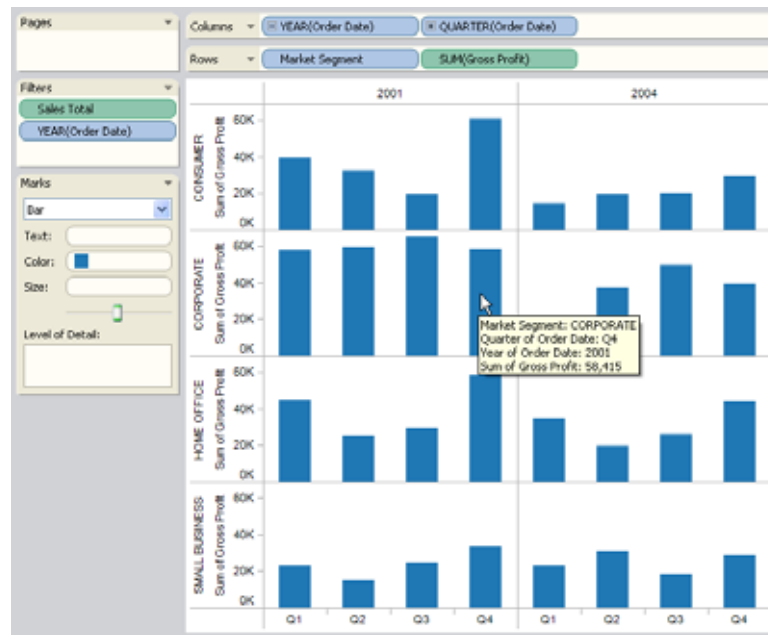
only the specified dimension members from the view. It does not modify the underlying data in any way.





You can also filter dimension members by right-clicking on one or more headers, and selecting **Exclude** from the context menu, as shown below. To select multiple headers, hold the **Ctrl** key while clicking.



The view is shown below.



The gross profit sum is calculated by including only the individual data rows that have sales greater than or equal to \$2,000. These data are then displayed by market segment and by quarter for the years 2001 and 2004. For example, the tooltip shows the gross profit sum for the Corporate market in Q4 for the year 2001 is \$58,415.

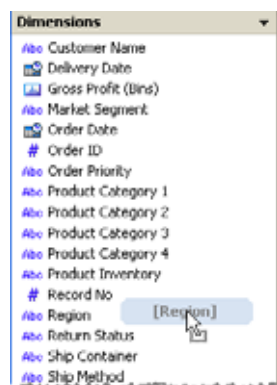
Note that you can quickly switch between the filtered and the unfiltered view by clicking the undo and redo buttons   on the toolbar. The undo/redo feature maintains a history of all the actions across the entire workbook, and provides an easy way to browse all the data views you created.

## Example 5 - Color Encode by Region


This view is a nested bar chart that displays gross profit by market segment and by quarter. The data are filtered so that sales less than \$2,000 are excluded, and the years 2002 and 2003 are excluded. The data are also color-encoded by region. It builds upon the view created in “Example 4 - Filter by Year and by Sales” on page 3-20.

To create the view, follow the steps below:


- 1 Click and drag the **Region** dimension from the Data window.



- 2 Place **Region** on the **Color** shelf.



Placing a dimension on the **Color** shelf separates the marks according to the members in the dimension, and assigns a unique color to each member. The color legend displays each member name and its associated color.



Notice that new headers are not created for the data view because the field does not add a new row or column to the table.

The final view is shown below.



Each column in the view is now separated into individual segments with unique colors. That is, the data are color-encoded. For the bar mark type, the segments are automatically stacked on top of each other.

Notice that some marks do not contain segments for all three regions. This is because that data did not make it through the sales filter. For example the Small Business market shows profits only in the East and Central regions in Q3 of 2004. The West region does not show profits because no sales figures reached \$2,000.

The tooltip shows the gross profit sum for the Corporate market in the East (orange) region in Q4 for the year 2001 is \$29,517.

## Build Data Views (Automatic “Show Me!”)

In addition to creating views manually, you can use Tableau’s “Show Me!” (**Ctrl+1**) functionality to create views automatically. A complete description of Show Me! can be found in the section on “Building Views Automatically” on page 11-1.

Here are a few quick examples of building views automatically.

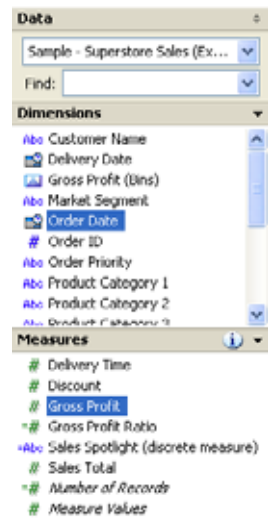
- Yearly Profits
- Sales vs. Profit by Product and Customer


### Yearly Profits

The first view is a line chart that displays gross profit as a function of year. To create the view using Show Me!, follow the steps below.

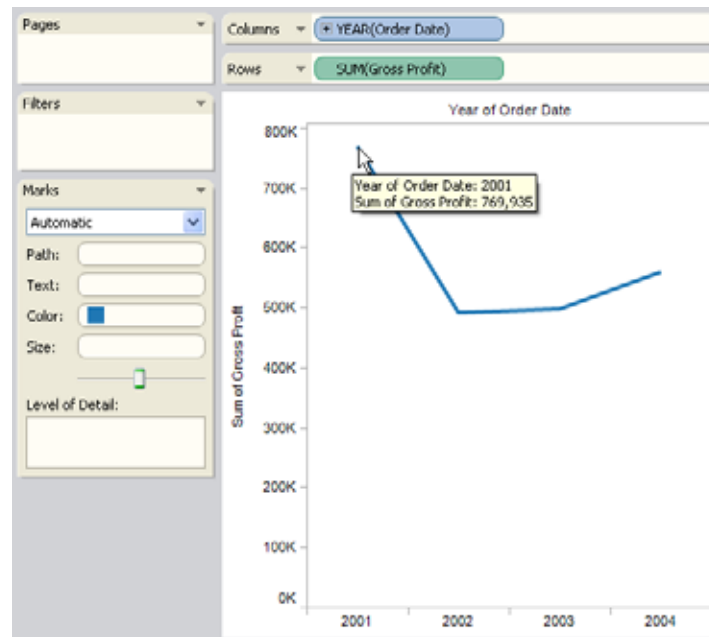


- 1 Select the **Order Date** and **Gross Profit** fields in the Data window. Press the Control (Ctrl) key to select multiple fields.




- 2 Click **Show Me!** on the toolbar .

The final view is shown below.



The line represents the summation of the gross profit for each year. For example, Tableau identifies all the data rows associated with the year 2001 and then sums all the corresponding gross profit values. It then repeats this process for the remaining three years. You can display the value for each point on the line using tooltips. To use tooltips, place the cursor over the mark of interest. For example, the sum of the gross profit for 2001 is \$769,935, which is significantly higher than the gross profit for the other years.


You can also display the value by turning on data labels by selecting **Format > Data Labels > On**.




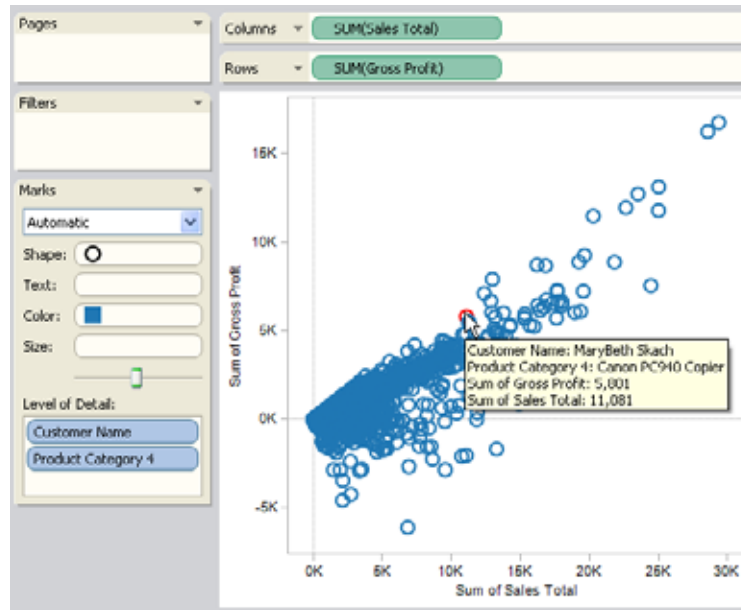
## **Sales versus Profit by Product and Customer**

To create a scatterplot view using Show Me! follow the steps below:


- 1 Select the fields labeled **Sales Total**, **Gross Profit**, **Product Category 4**, and **Customer Name** in the Data window. Hold the Control (Ctrl) key to select multiple fields.



- 2 Click **Show Me!** on the toolbar  **Show Me!** .



As you can see, the Show Me! command automatically creates a view of data, based on the fields you selected to analyze. This example shows two measures and two dimensions. The result is a scatterplot demonstrating the relationship between those two measures for every combination of the two dimensions. For instance, if you look at the point selected above, it represents the SUM(Sales) and the SUM(Gross Profit) for all “Canon PC940 Copiers” purchased by “MaryBeth Skach.” You can now manually interact with this automatically-generated view to create exactly the desired analysis.

- 3 Now what if you want to see this information presented in bars instead of a scatterplot? You can change to bars by clicking **Show Me! Alternatives**  on the toolbar. The

Show Me! Alternatives dialog box opens so you can pick the style of display for your data.




- 4 Select the graphic style you want to use and click **OK**.

**Note** Each view type requires a specific combination of dimensions and measures. If the view does not meet the requirements, the icon is grey in the dialog box. If you're not sure which type of view is best, hover the mouse over each icon in the dialog box to read a description of the view type and its requirements.

## Save Your Work

After you have created all the desired views of your data, you should save the results in a Tableau workbook. Saving a Tableau workbook allows you to save all of your worksheets for later use. It also allows you to share your results using a convenient file. To save a Tableau workbook, follow the steps below:

- 1 Select the **File > Save** menu item or type **Ctrl+S**.



- 2 Specify the workbook file name.

If this is the first time you are saving the workbook, the **Save As** dialog box opens. Type a name for the new workbook.




Tableau saves the workbook file with the .twb extension. The default location is the workbooks folder of the Tableau repository. However, you can save Tableau workbooks to any location you choose.



You can open the saved workbook file and Tableau will automatically connect to the data and restore your data views. Refer to Chapter 24, “Saving and Exporting” to learn more about saving your work.

---

**Note** When you share workbooks that are connected to file type data sources (e.g., Excel, Access, text, and data extracts), you also have to share the data source so others can connect to the data. An easy way to share your work and the data sources is to save a packaged workbook. Refer to “Packaged Workbooks” on page 24-5 to learn more.

---

---

<b>Overview . . . . .</b>	<b>4-2</b>
<b>Supported Data Sources . . . . .</b>	<b>4-3</b>
<b>How to Connect to a Data Source . . . . .</b>	<b>4-6</b>
Connect Using the Open Dialog Box . . . . .	4-6
Connect Using the Connect to Data Dialog Box . . . . .	4-7
<b>Examples – Connecting to Data Sources . . . . .</b>	<b>4-11</b>
Microsoft Access Database . . . . .	4-11
Microsoft Excel Workbook . . . . .	4-14
Tableau Data Extract database . . . . .	4-18
Text File . . . . .	4-20
Firebird Database . . . . .	4-25
Hyperion Essbase Cube (aka IBM OLAP Server) . . . . .	4-29
IBM DB2 . . . . .	4-33
Microsoft Analysis Services Cube . . . . .	4-41
Microsoft SQL Server Database . . . . .	4-45
MySQL Database . . . . .	4-49
Netezza . . . . .	4-53
Oracle Database . . . . .	4-57
PostgreSQL Database . . . . .	4-61
IBM OLAP Server . . . . .	4-64
<b>Connecting to a Custom SQL Query . . . . .</b>	<b>4-65</b>
<b>Editing the Connection . . . . .</b>	<b>4-66</b>
<b>Exporting the Connection . . . . .</b>	<b>4-68</b>
<b>Refreshing the Data . . . . .</b>	<b>4-70</b>
<b>Closing the Connection . . . . .</b>	<b>4-71</b>

---

To begin analyzing your data, first connect Tableau to one or more data sources. A data source can be as simple as an Excel workbook, or as elaborate as a SQL Server or Oracle data warehouse. After connecting, the data fields become available in the Data window on the left side of the workbook. This section describes the types of data supported and how to create and maintain a basic connection in the following sections:

- Supported Data Sources
- How to Connect to a Data Source
- Examples – Connecting to Data Sources
- Connecting to a Custom SQL Query
- Editing the Connection
- Exporting the Connection
- Refreshing the Data
- Closing the Connection

Tableau supports a wide variety of data sources, including Microsoft Office files, SQL databases, comma delimited text files, and multi-dimensional OLAP (cube) databases.

As shown below, these data sources are divided into two types: relational and multidimensional (OLAP). The data sources supported by your copy of Tableau are determined by the version purchased. Refer to the Products page of the Tableau Web site for more information on data source compatibility requirements.

Type	Data Source	Supported Versions
Relational	Microsoft Excel workbook	Excel 2000 or later (including Excel 2007)
	Microsoft Access database	MS Access 2000 or later (including Access 2007)
	Microsoft SQL Server database	SQL Server 2000 or later (including SQL Server 2005)
	PostgreSQL	Version 7.0 or later
	Firebird	Version 2.0 or later
	IBM DB2	Version 8.0 or later
	InterSystems Caché	Version 5
	MySQL database	Version 3.23 or later
	Netezza	Version 3.1.4 or later
	Oracle database	Oracle 9i or later
	Text file	N/A
	Tableau Data Extract	N/A

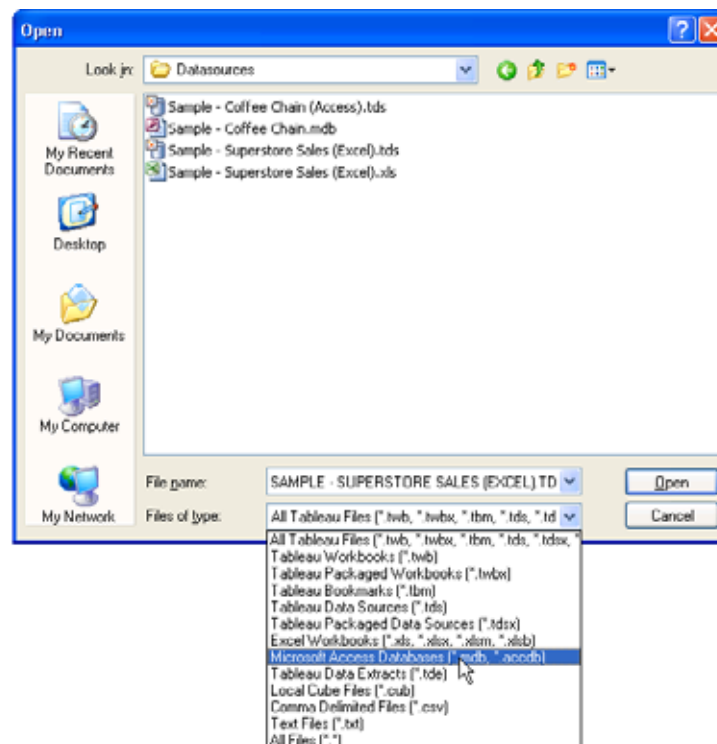
Type	Data Source	Supported Versions
Multidimensional (OLAP)	Microsoft Analysis Services (SSAS) cube	SSAS 2000 or later (including SQL Server 2005)
	Hyperion Essbase cube	Version 7.1, 7.x or later.
	IBM DB2 OLAP Server	Version 8.1 or later

To build views of your data, you must first connect Tableau to a data source. Depending on the type of data you are connecting to, you can connect in one of two ways: with the **Open** dialog box or with the **Connect to Data Source** dialog box.

You can select Tableau workbooks, Tableau packaged workbooks, Tableau bookmarks, Tableau data source connections, Excel workbooks, Access databases, Tableau data extract files, local Cube (.cub) files, comma delimited files (.csv), and text files using the **Open** dialog box.

**To connect using the open dialog box:**

- 1 Select **File > Open**.



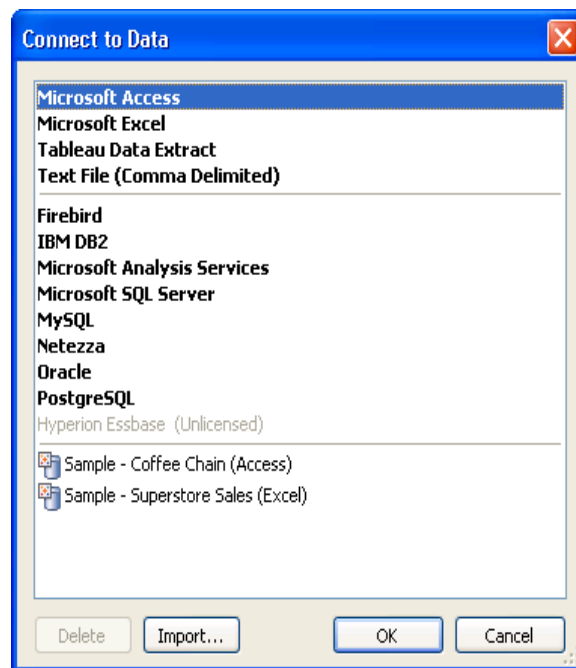
- 2 Locate and select the data file and click **Open**.
- 3 If you select a data source file, Tableau opens the appropriate connection dialog box. Complete the connection dialog box and click **OK**. Refer to “Examples – Connecting to Data Sources” on page 4-11 for information about connecting to specific types of data.

You can connect to any supported data source with the Connect to Data dialog box.

**To connect using the Connect to Data dialog box:**




- 1 Select **Data > Connect to Data** or press **Ctrl + D** on your keyboard.
- 2 In the Connect to Data dialog box, select the type of data you want to connect to. You can also select a saved data connection (TDS files) from the Favorites list. Refer to “Exporting the Connection” on page 4-68 to learn how to save a data connection.



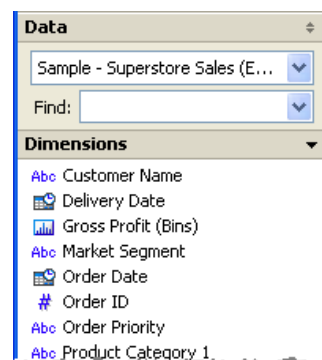
- 3 A data source-specific dialog box opens that allows you to complete the connection process. Refer to “Examples – Connecting to Data Sources” on page 4-11 for more information.

**Note** Another way to connect to data is to import from a workbook. A workbook can contain multiple worksheets, each of which can be connected to a different data source. To import a connection from a workbook, click the **Import from Workbook** link at the bottom of the Connect to Data dialog box.


- 4 After completing the connection, Tableau prompts you to specify whether you want to remember the connection. Select **Yes**, to create a shortcut to the data source, called a TDS file, so that in the future you can bypass configuring the connection.



After the connection is established, the data source fields display on the left side of the workbook in the Data window. Refer to “Understanding the Data Window” on page 8-3 to learn more.



You can display information about the connection by selecting **Data > Connection > Properties**. The properties of an example data source are shown below



## Examples – Connecting to Data Sources

This section contains examples that show you how to connect to the following Relational and Multidimensional Data supported by Tableau:

- Microsoft Access Database
- Microsoft Excel Workbook
- Tableau Data Extract database
- Text File
- Firebird Database
- Hyperion Essbase Cube (aka IBM OLAP Server)
- IBM DB2
- InterSystems Caché
- Microsoft Analysis Services Cube
- Microsoft SQL Server Database
- MySQL Database
- Netezza
- Oracle Database
- PostgreSQL Database
- IBM OLAP Server

### Microsoft Access Database

This example discusses how to connect Tableau to a Microsoft Access database.


---

**Note** Tableau supports all Access data types except OLE Object and Hyperlink.

---

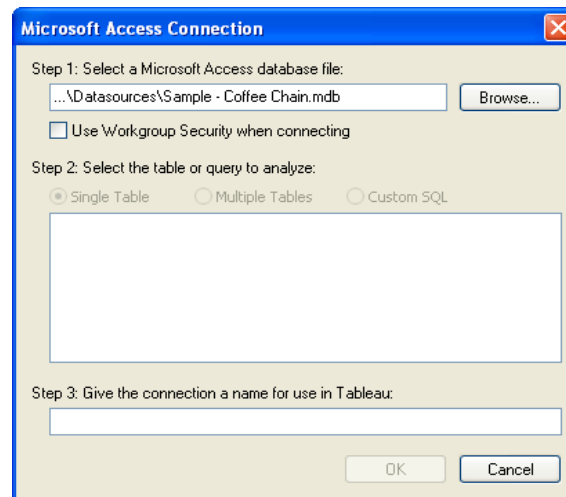
**To connect to an Access database:**

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Microsoft Access**.




- 3 Complete the connection.

Follow the steps on the **Microsoft Access Connection** dialog box to complete the connection. These steps are described below.



**Step 1: Select a Microsoft Access database file**

Select the database by typing its name or by navigating to the file. Access databases have the .mdb file extension.




Click **Browse** to navigate to the file using the Open dialog box.



You can select whether to use workbook security so you can connect to both password protected files and Access files protected by workgroup security. When you select the workbook security option, a logon dialog box opens where you can enter your password or select **Use workgroup security**. If the file is protected by workgroup security, type the System database, user name, and password into the corresponding text fields.

**Step 2: Select the table or query to analyze**

You can connect to a table or query. You can also connect to a set of relational tables that are related by join conditions. To do so, select the Multiple Tables or Custom SQL options. You can also add joins later. Refer to “Joining Tables” on page 5-1.

**Step 3: Give the connection a name for use in Tableau**

Specify a unique name for the connection. Note that a default name is automatically generated based on the table or query name.



If you just created the Access data source and it is still open on your computer, you cannot directly connect to it through Tableau. An error message appears informing you that there was a problem connecting to the data source. Simply close the Access data source on your computer and try to connect again. After you are connected you can reopen your file in Access.

---


**Note** If the Access file contains columns that are more than 254 characters wide, Tableau will not be able to use these fields. Either remove the columns from the table or modify them to fit within 254 characters prior to connecting in Tableau.

---

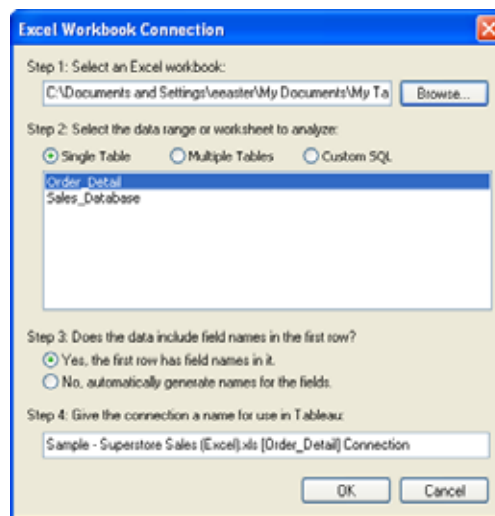
## Microsoft Excel Workbook

**To connect Tableau to an Excel workbook:**

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Microsoft Excel**.



- 3 Follow the steps on the **Excel Workbook Connection** dialog box to complete the connection. These steps are described below. Note that you cannot connect to password-protected Excel files.





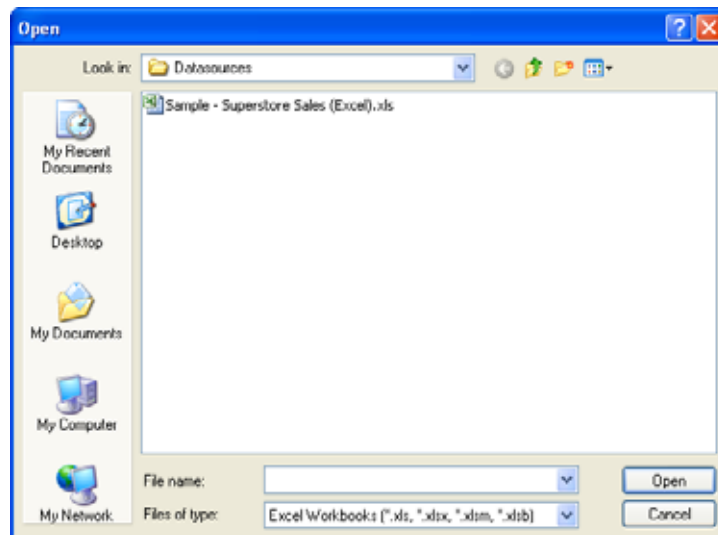
**Step 1: Select an Excel workbook**

Select the workbook by typing its name or by navigating to the file. Excel files have the .xls file extension.

Step 1: Select an Excel workbook:

Browse...

Click **Browse** to navigate to the file using the **Open** dialog box.

**Step 2: Select the data range or worksheet to analyze**

You can connect to a worksheet or a named range. Named ranges allow you to connect to just a specific portion of an Excel worksheet. You can create a named range in Excel by highlighting a range of cells and then selecting **Insert > Name > Define**. Then give the range of cells a name. You can now connect to this named range in Tableau in the same way you can connect to a worksheet.

You can also connect to a set of relational tables that are related by join conditions. To do so, select the Multiple Tables option. You can also add joins later. Refer to “Joining Tables” on page 5-1.

Step 2: Select the data range or worksheet to analyze:

☒ Single Table
 ☐ Multiple Tables
 ☐ Custom SQL

Order\_Detail  
Sales\_Database

### Step 3: Does the data include field names in the first row

Specify whether the workbook includes field names in the first row. If field names are not included, they can be automatically generated by Tableau with names such as F1, F2, F3, and so on.

Step 3: Does the data include field names in the first row?

☒ Yes, the first row has field names in it.
 ☐ No, Automatically generate names for the fields.

The Sales Database worksheet includes field names such as Record No, Order Priority, Sales Total, and so on.

	A	B	C	D	
1	Record No	Order Priority	Sales Total	Discount	Field names
2	1	5-LOW	2302.42	0.00	RE
3	2	5-LOW	996.66	0.10	EX
4	3	5-LOW	2664.66	0.07	RE
5	4	5-LOW	220.52	0.12	RE
6	5	3-MEDIUM	159.99	0.10	RE
7	6	3-MEDIUM	2678.68	0.06	RE
8	7	3-MEDIUM	888.31	0.07	EX
9	8	3-MEDIUM	54.95	0.04	RE
10	9	3-MEDIUM	90.42	0.05	RE

### Step 4: Give the connection a name for use in Tableau

Specify a unique name for the connection. Note that a default name is automatically generated based on the workbook name. You can connect to multiple Excel worksheets at the same time provided each worksheet has a unique connection name.

Step 4: Give the connection a name for use in Tableau:

Sample - Superstore Sales (Excel).xls [Sales\_Database] Connection

OK Cancel


**Note** If the Excel file contains columns that are more than 254 characters wide, Tableau will not be able to use these fields. Either remove the columns from the table or modify them to fit within 254 characters prior to connecting in Tableau.

## Tableau Data Extract database

This example discusses how to connect Tableau to a Tableau Data Extract file. To learn about how to create extracts refer to “Extracting Data” on page 6-1.

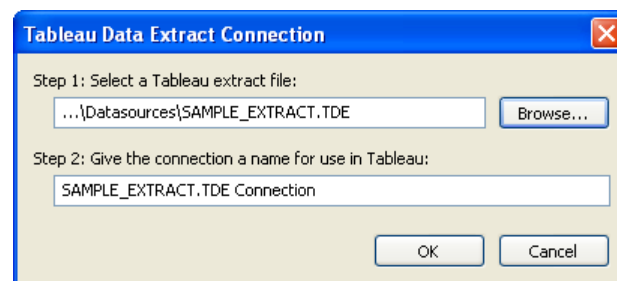
**To connect to Hyperion Essbase cube:**

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Tableau Data Extract**.



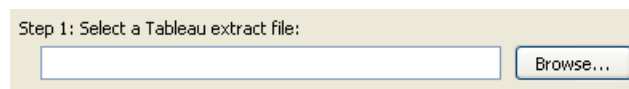
- 3 Complete the connection.

Follow the steps on the **Tableau Data Extract Connection** dialog box to complete the connection. These steps are described below.



**Step 1: Select a Tableau extract file.**

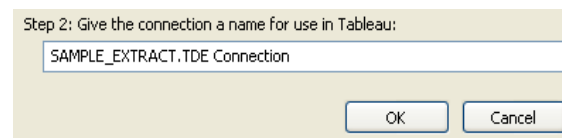
Type the name of a Tableau extract file or click Browse to navigate to the file on your computer.

A screenshot of a dialog box titled "Step 1: Select a Tableau extract file:". It features a text input field and a "Browse..." button to its right.

Step 1: Select a Tableau extract file:

**Step 2: Give the connection a name for use in Tableau**

Specify a unique name for the connection. Note that a default name is automatically generated based on the file name.

A screenshot of a dialog box titled "Step 2: Give the connection a name for use in Tableau:". It contains a text input field with the text "SAMPLE\_EXTRACT.TDE Connection" and "OK" and "Cancel" buttons at the bottom right.


Step 2: Give the connection a name for use in Tableau:

## Text File

This example discusses how to connect Tableau to a text file. Tableau connects to comma separated text (CSV) files. However, you can also connect to tab-delimited text files with a few extra steps. Refer to “How to Connect to Tab-Delimited Text Files” on page 4-23.

**To connect to a text file, follow these steps:**

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Text File (Comma Delimited)**.




- 3 Complete the connection.

Follow the steps on the **Text File Connection** dialog box to complete the connection. These steps are described below.




### Step 1: Select a comma-delimited text file

Select the database by typing its name or by navigating to the file.




Click **Browse** to navigate to the file using the Open dialog box.



Large text files often perform poorly as a data source, because the queries are slow. If Tableau determines the file is too big to perform well, you will be prompted to create an extract. Refer to “Extracting Large Text and Excel Files” on page 27-9.

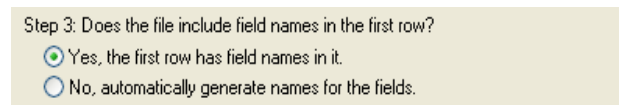
### Step 2: Select the data range to analyze

You can connect to a single file or a set of relational tables that are related by join conditions. To do so, select the Multiple Tables option. You can also add joins later. Refer to “Joining Tables” on page 5-1.




### Step 3: Does the data include field names in the first row

Specify whether the file includes field names in the first row. If field names are not included, they can be automatically generated by Tableau with names such as F1, F2, F3, and so on.




If your text file has field names, it will have descriptive entries as shown below.



**Step 4: Give the connection a name for use in Tableau**

Specify a unique name for the connection. Note that a default name is automatically generated based on the file name.

You can connect to multiple text files at the same time, but each one requires a unique connection name in Tableau.

**How to Connect to Tab-Delimited Text Files**

If you are using Tableau v2.1 or higher, you can also connect to tab-delimited text files with a few extra steps.

**To connect to tab-delimited text files:**

- 1 If you are using Tableau v2.1 rename your data file to have a .txt file extension. If you are using v3.0 or later there is no need to rename the file.
- 2 Create a new text file (using Notepad or another text editor) and type the following:

**[Your\_Datasource.txt]**

**Format =TabDelimited**

Substitute the name of your data file for "Your\_Datasource.txt."

- 3 Save the file in the same directory as the data file and call it **schema.ini**.
- 4 Return to Tableau and connect to the data file using the Text file instructions outlined above (refer to "Text File" on page 4-20).

You can connect to pipe delimited files by typing **Format = Delimited (|)** into the schema.ini file instead of TabDelimited as shown above.



---

**Note** If the text file contains columns that are more than 254 characters wide, Tableau will not be able to use these fields. Either remove the columns or modify them to fit within 254 characters prior to connecting in Tableau.

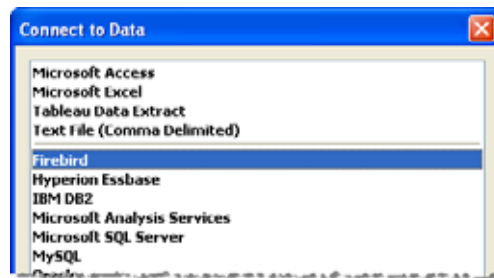
---

## Firebird Database

This example discusses how to connect Tableau to a Firebird database. Connecting to this data source requires the Professional Edition of Tableau. You can see what version you have by selecting **Help > About Tableau**.

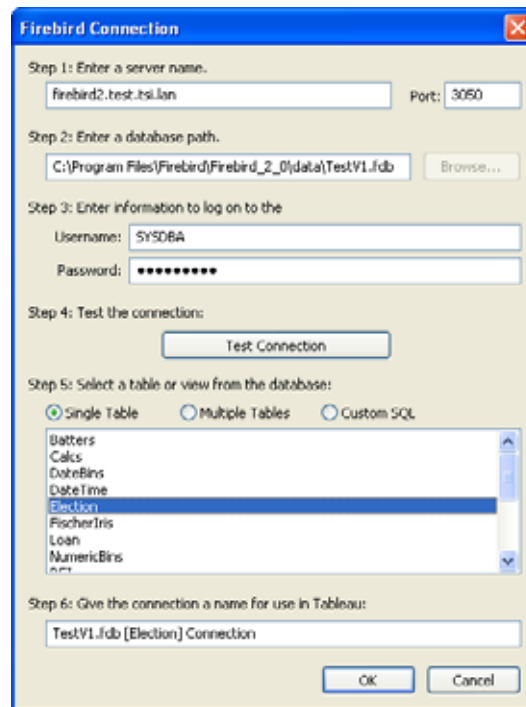
**To connect to a Firebird database:**

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Firebird**.



- 3 Complete the connection.

Follow the steps on the **Firebird Connection** dialog box to complete the connection. These steps are described below.



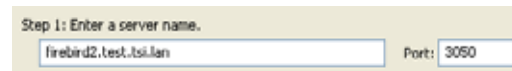
The image shows a 'Firebird Connection' dialog box with a blue title bar and a close button. It contains six steps for configuring a connection:

- Step 1: Enter a server name.** A text box contains 'firebird2.test.tsl.lan' and a 'Port' field contains '3050'.
- Step 2: Enter a database path.** A text box contains 'C:\Program Files\Firebird\Firebird\_2\_0\data\TestV1.fdb' and a 'Browse...' button.
- Step 3: Enter information to log on to the**
  - Username:** 'SYSDBA'
  - Password:** masked with nine dots.
- Step 4: Test the connection:** A 'Test Connection' button.
- Step 5: Select a table or view from the database:**
  - Radio buttons for 'Single Table' (selected), 'Multiple Tables', and 'Custom SQL'.
  - A list box showing database objects: 'Batters', 'Calcs', 'DateBins', 'DateTime', 'Election' (highlighted), 'FischerTris', 'Loan', 'NumericBins', and 'ACT'.
- Step 6: Give the connection a name for use in Tableau:** A text box containing 'TestV1.fdb [Election] Connection'.

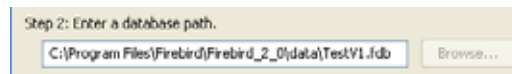
At the bottom right are 'OK' and 'Cancel' buttons.

**Step 1: Enter a server name**

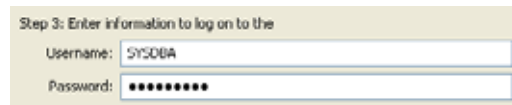
Type in the name of a specific server.

A screenshot of a dialog box titled "Step 1: Enter a server name." It contains a text input field with the value "firebird2.test.tsl.lan" and a port input field with the value "3050".**Step 2: Enter a database path**


Specify the location of the database.

A screenshot of a dialog box titled "Step 2: Enter a database path." It contains a text input field with the value "C:\Program Files\Firebird\Firebird\_2\_0\data\TestV1.fdb" and a "Browse..." button.**Step 3: Enter information to log on to the server**

Specify the username and password.


A screenshot of a dialog box titled "Step 3: Enter information to log on to the server". It contains two input fields: "Username:" with the value "SYSDBA" and "Password:" with a masked password represented by eight asterisks.**Step 4: Test the connection**

If the connection is successful, Tableau will alert you. If the connection is unsuccessful, verify that your user name and password are correct. If the connection continues to fail, your computer is having trouble locating the server. Contact your network administrator or DBA.

A screenshot of a dialog box titled "Step 4: Test the connection:". It contains a single button labeled "Test Connection".**Step 5: Enter a table or view from the database**

Firebird databases can contain multiple tables and views. Specify which table or view within the database you want to connect to. You can also connect to multiple tables or a specific query by selecting the Multiple Tables and Custom SQL options.

Refer to “Joining Tables” on page 5-1 to learn more about connecting to a set of relational tables that are related by join conditions.



#### Step 6: Give the connection a name for use in Tableau

Specify a unique name for the connection. Note that a default name is automatically generated based on the database name.




**Note** In order to take full advantage of the calculation language in Tableau, you need to load some extra functions into your Firebird database. These functions are located in three .dll files in C:\Program Files\Tableau\Tableau 3.0\bin\udf. Load the files using DEFINE FUNCTION commands. You can use the .sql script files located with the UDF dlls along with iSQL to load the function definitions into your server. Keep in mind that the scripts load all of the functions in the corresponding udf libraries, so if you have existing functions you may need to cut and paste.

## Hyperion Essbase Cube (aka IBM OLAP Server)

This example discusses how to connect Tableau to an Hyperion Essbase cube (also known as IBM OLAP Server). Connecting to this data source requires the Professional Edition of Tableau. You can see what version you have by selecting **Help > About Tableau**.

**To connect to Hyperion Essbase cube:**

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Hyperion Essbase**.



- 3 Complete the connection.

Follow the steps on the **Hyperion Essbase Connection** dialog box to complete the connection. These steps are described below.

The image shows a 'Hyperion Essbase Connection' dialog box with the following steps:

- Step 1: Enter a server name:** A text box containing 'blackcomb'.
- Step 2: Enter information to log on to the server:** A 'Username:' text box with 'test' and a 'Password:' text box with eight dots.
- Step 3: Test the connection:** A 'Test Connection' button.
- Step 4: Select an application on the server:** A dropdown menu showing 'Sample'.
- Step 5: Select a database from the application:** A list box with 'Basic' selected.
- Step 6: Give the connection a name for use in Tableau:** A text box containing '(Sample-Basic) Connection'.

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

### Step 1: Enter a server name

Type the name of the server.

This close-up shows Step 1: Enter a server name: with a text box containing 'blackcomb'.

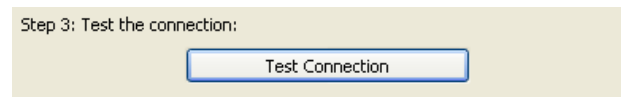
### Step 2: Enter information to log on to the server

Specify the username and password.

This close-up shows Step 2: Enter information to log on to the server: with a 'Username:' text box containing 'test' and a 'Password:' text box containing eight dots.

**Step 3: Test the connection**

Click **Test Connection**. If the connection is unsuccessful, verify that your user name and password are correct. If the connection continues to fail, your computer may be having trouble locating the server. Contact your network administrator or database administrator.

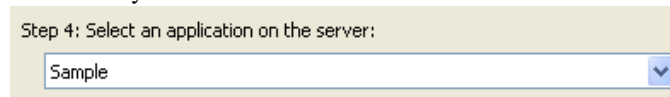
A dialog box with a light beige background. At the top, it says "Step 3: Test the connection:". Below this text is a single button labeled "Test Connection".

Step 3: Test the connection:

Test Connection

**Step 4: Select an application on the server**

Servers often contain multiple applications. Select the specific application on the server that you want to connect to.

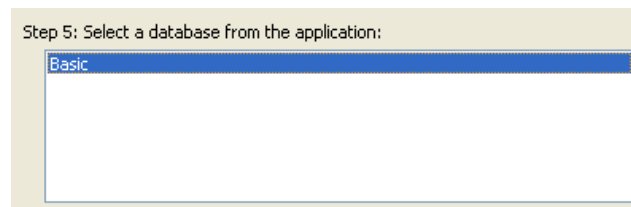
A dialog box with a light beige background. At the top, it says "Step 4: Select an application on the server:". Below this text is a dropdown menu showing "Sample" with a downward arrow on the right.

Step 4: Select an application on the server:

Sample

**Step 5: Select a database from the application**

Applications often contain multiple databases. Select the specific database that you want to connect to.

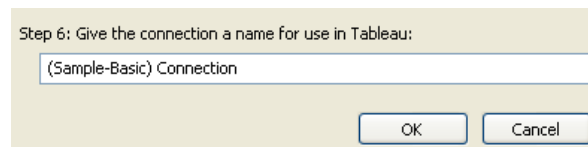
A dialog box with a light beige background. At the top, it says "Step 5: Select a database from the application:". Below this text is a list box with "Basic" selected and highlighted in blue.

Step 5: Select a database from the application:

Basic

**Step 6: Give the connection a name for use in Tableau**

Specify a unique name for the connection. Note that a default name is automatically generated based on the application name.

A dialog box with a light beige background. At the top, it says "Step 6: Give the connection a name for use in Tableau:". Below this text is a text input field containing "(Sample-Basic) Connection". At the bottom right are two buttons: "OK" and "Cancel".

Step 6: Give the connection a name for use in Tableau:

(Sample-Basic) Connection

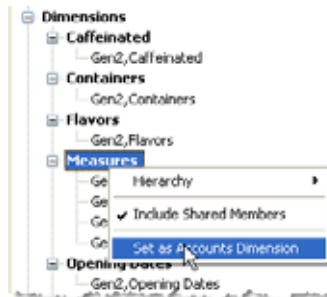
OK Cancel



**Step 7: Set an Accounts Dimension (Optional)**

In some case, the accounts dimension for your data source can appear in the **Dimensions** area of the Data window. This might occur if there is an error in the cube and another field is identified as the accounts dimension or there is no accounts dimension set at all. The accounts dimension defines the fields that are included as measures.

For example, as shown in the figure below, **Measures** is the accounts dimension for the data source but appears as a dimension in the Data window. To correct this error, right-click the field and then select **Set as Accounts Dimension** from the context menu.



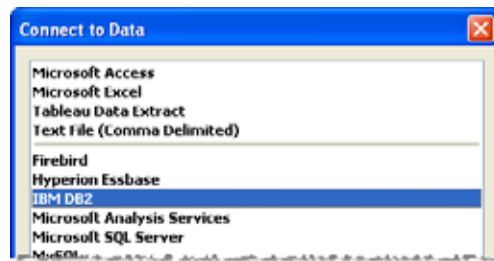
If there is no account dimension set at all, Tableau will prompt you to define one.

## IBM DB2

This example discusses how to connect Tableau to an IBM DB2 database. Connecting to this data source requires the Professional Edition of Tableau. You can see what version you have by selecting **Help > About Tableau**.

**To connect to an IBM DB2 database:**

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **IBM DB2** and click **OK**.



- 3 Complete the connection.

Follow the steps on the **IBM DB2 Connection** dialog box to complete the connection. These steps are described below.

The screenshot shows a 'DB7 Connection' dialog box with the following steps:

- Step 1: Enter a server name:** The text box contains 'db2-81.test.itallan' and the Port is set to '50000'.
- Step 2: Enter a database on the server:** The text box contains 'TESTV1'.
- Step 3: Enter information to log on to the database:** The Username is 'test' and the Password is masked with dots.
- Step 4: Test the connection:** A 'Test Connection' button is present.
- Step 5: Select a table or view from the database:** Radio buttons for 'Single Table' (selected), 'Multiple Tables', and 'Custom SQL' are shown. A list of tables is displayed, with 'Election (TEST.Election)' selected.
- Step 6: Give the connection a name for use in Tableau:** The text box contains 'TESTV1 [Election (TEST.Election)] Connection'.

At the bottom are 'OK' and 'Cancel' buttons.

### Step 1: Enter a server name

Type in the name of a specific server and the port to which you want to connect.

This close-up shows Step 1 of the connection process. The text box for the server name contains 'My Server' and the Port is set to '50000'.

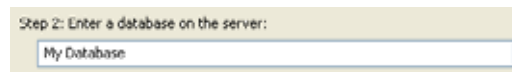
The port is dependent on the type of server you are connecting to and whether you are connecting to an encrypted port. The table below describes some common port

settings, however, it is possible that your server is configured to use a non-standard port. Contact an administrator if you don't know which port to connect to.

Type of Server	Security	Port
Linux, Windows, Unix	Non-Encrypted	50000
	Encrypted	60000

#### Step 2: Enter a database on the server

Servers often contain multiple databases. Type the name of a specific database on the server that you want to connect to.

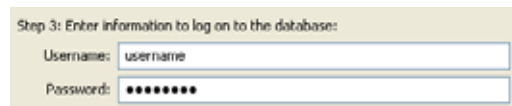


Step 2: Enter a database on the server:

My Database

#### Step 3: Enter information to log on to the server

Specify the username and password.



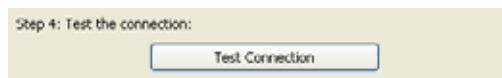
Step 3: Enter information to log on to the databases:

Username: username

Password: \*\*\*\*\*

#### Step 4: Test the connection

If the connection is unsuccessful, verify that your user name and password are correct. If the connection continues to fail contact your network administrator or DBA.




Step 4: Test the connection:

Test Connection

#### Step 5: Select a table or view from the database

IBM DB2 databases can contain multiple tables and views. Specify which table or view within the database you want to connect to. You can also connect to multiple tables or a specific query by selecting the Multiple Tables and Custom SQL options.

Refer to “Joining Tables” on page 5-1 to learn more about connecting to a set of relational tables that are related by join conditions.



#### Step 6: Give the connection a name for use in Tableau

Specify a unique name for the connection. Note that a default name is automatically generated based on the database name.



## InterSystems Caché

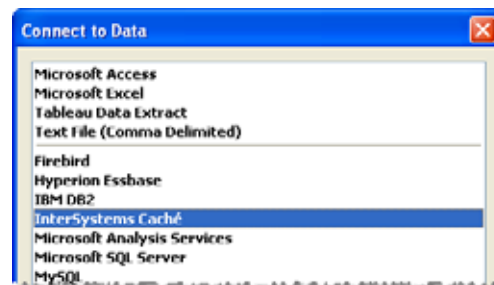
This example discusses how to connect Tableau to an InterSystems Caché database. Connecting to this data source requires a special license for Tableau that has Caché enabled. Additionally, you or a database administrator must load a stored procedure into your Caché 5 database before you connect using Tableau.

### To load the stored procedure into your Caché database:

- 1 Copy the the XML file on the following web page to the Caché server:  
[www.tableausoftware.com/community/support/cache](http://www.tableausoftware.com/community/support/cache)
- 2 Start Caché Studio at the server from the Cube in the system tray.
- 3 Within Studio, select **Tools > Import Local**.
- 4 In the subsequent dialog box select the XML file.
- 5 In the Import dialog box, make sure both the class name **TableauFuncs** and the **Compile Imported** option are selected.
- 6 Click **OK**.

### To connect to an InterSystems Caché database in Tableau:

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **InterSystems Caché** and click **OK**.



### 3 Complete the connection.

Follow the steps on the **InterSystem Caché Connection** dialog box to complete the connection. These steps are described below.

The dialog box titled "InterSystems Caché Connection" contains the following steps:

- Step 1: Enter a server name:** The text box contains "cacheserver.test.lan" and the Port is set to "1972".
- Step 2: Enter a database on the server:** The text box contains "testdb".
- Step 3: Enter information to log on to the database:** The Username text box contains "username" and the Password text box contains seven dots.
- Step 4: Test the connection:** A button labeled "Test Connection" is present.
- Step 5: Select a table or view from the database:**
  - Radio buttons for "Single Table" (selected) and "Multiple Tables".
  - A list box containing the following items: Election (TESTV1.Election), FischerIris (TESTV1.FischerIris), Loan (TESTV1.Loan), NumericBins (TESTV1.NumericBins), **REI (TESTV1.REI)** (highlighted), SeattleCrime (TESTV1.SeattleCrime), Securities (TESTV1.Securities), and SpecialData (TESTV1.SpecialData).
- Step 6: Give the connection a name for use in Tableau:** The text box contains "REI Connection".

At the bottom right are "OK" and "Cancel" buttons.

#### Step 1: Enter a server name

Type in the name of a specific server and the port to which you want to connect.

This close-up shows Step 1 of the dialog box. The text box for the server name contains "cacheserver.test.lan" and the Port text box contains "1972".

**Step 2: Enter a database on the server**

Servers often contain multiple databases. Type the name of a specific database on the server that you want to connect to. Database names are case sensitive.

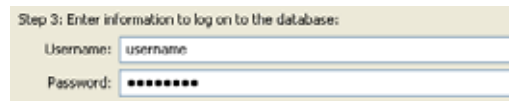


Step 2: Enter a database on the server:

testdb

**Step 3: Enter information to log on to the server**

Specify the username and password.



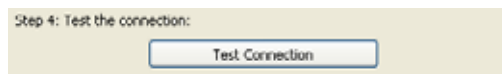
Step 3: Enter information to log on to the database:

Username: username

Password: .....

**Step 4: Test the connection**

If the connection is unsuccessful, verify that your user name and password are correct. If the connection continues to fail contact your network administrator or DBA.

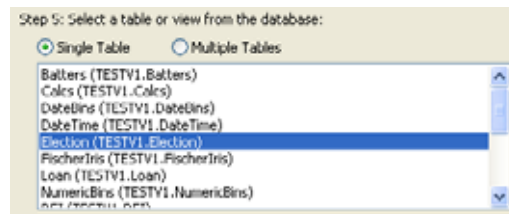


Step 4: Test the connection:

Test Connection

**Step 5: Select a table or view from the database**

InterSystems Caché databases can contain multiple tables and views. Specify which table or view within the database you want to connect to. You can also connect to multiple tables by selecting the Multiple Tables option. Refer to “Joining Tables” on page 5-1 to learn more about connecting to a set of relational tables that are related by join conditions.



Step 5: Select a table or view from the database:


☒ Single Table ☐ Multiple Tables

Batters (TESTV1.Batters)  
Coles (TESTV1.Coles)  
DateBins (TESTV1.DateBins)  
DateTime (TESTV1.DateTime)  
**Election (TESTV1.Election)**  
FischerIris (TESTV1.FischerIris)  
Loan (TESTV1.Loan)  
NumericBins (TESTV1.NumericBins)  
Test (TESTV1.Test)



**Step 6: Give the connection a name for use in Tableau**

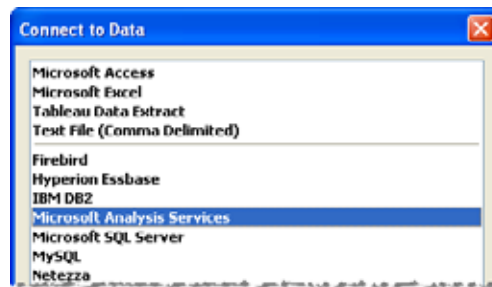
Specify a unique name for the connection. Note that a default name is automatically generated based on the database name.



This example discusses how to connect Tableau to a Microsoft Analysis Services (SSAS) cube. Connecting to this data source requires the Professional Edition of Tableau. You can see what version you have by selecting **Help > About Tableau**.

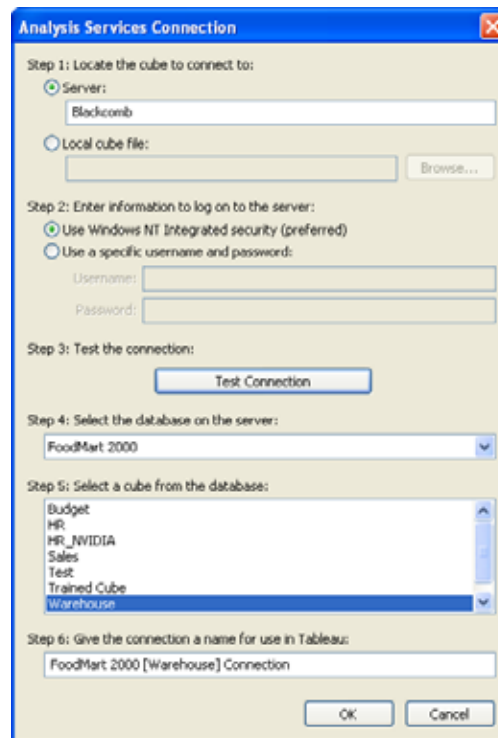
**To connect Tableau to a Microsoft Analysis Services Cube:**

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Microsoft Analysis Services**.



- 3 Complete the connection.

Follow the steps on the **Analysis Services Connection** dialog box to complete the connection. These steps are described below.



The image shows a Windows-style dialog box titled "Analysis Services Connection". It is divided into six steps for configuring a connection to an Analysis Services cube.

**Step 1: Locate the cube to connect to:**  
This step has two radio buttons. The first, "Server:", is selected and followed by a text box containing "Blackcomb". The second, "Local cube file:", is unselected and followed by a text box and a "Browse..." button.

**Step 2: Enter information to log on to the server:**  
This step has two radio buttons. The first, "Use Windows NT Integrated security (preferred)", is selected. The second, "Use a specific username and password:", is unselected and followed by "Username:" and "Password:" text boxes.

**Step 3: Test the connection:**  
This step contains a single button labeled "Test Connection".

**Step 4: Select the database on the server:**  
This step contains a dropdown menu currently showing "FoodMart 2000".

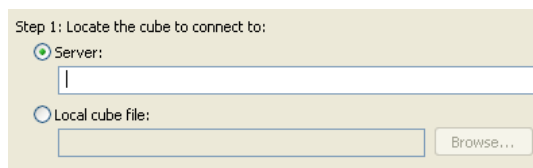
**Step 5: Select a cube from the database:**  
This step contains a list box with the following items: Budget, HR, HR\_NVIDIA, Sales, Test, Trained Cube, and Warehouse. The "Warehouse" item is currently selected and highlighted.

**Step 6: Give the connection a name for use in Tableau:**  
This step contains a text box with the text "FoodMart 2000 [Warehouse] Connection".

At the bottom right of the dialog are "OK" and "Cancel" buttons.

**Step 1: Locate the cube to connect to**

Select whether to connect to a cube file remotely on a server or locally. To connect to a remote cube, type the name of a specific server into the text box. If you are connecting to the server via HTTP you can enter the URL as the server name. To connect to a local cube file, select **Local cube file** and click **Browse** to navigate to the cube file on your computer.



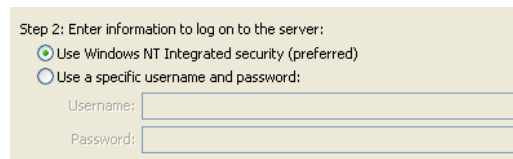
Step 1: Locate the cube to connect to:

☒ Server:

☐ Local cube file:

**Step 2: Provide login information to the server**

Specify whether to use Windows NT Integrated security or a specific username and password. If the cube is password protected, you must type your username and password.



Step 2: Enter information to log on to the server:

☒ Use Windows NT Integrated security (preferred)

☐ Use a specific username and password:

Username:

Password:

**Step 3: Test the connection**

Click **Test Connection**. If the connection is unsuccessful, verify that your user name and password are correct. If the connection continues to fail, your computer may be having trouble locating the server. Contact your network administrator or database administrator.



Step 3: Test the connection:

**Step 4: Select the database on the server**

Servers often contain multiple databases. Select the specific database on the server that you want to connect to.

Step 4: Select the database on the server:

FoodMart 2000

**Step 5: Select a cube from the database**

SSAS databases often contains multiple cubes. Specify which cube within the database you want to connect to.

Step 5: Select a cube from the database:

Budget  
HR  
HR\_NVIDIA  
Sales  
Test  
Trained Cube  
Warehouse

**Step 6: Give the connection a name for use in Tableau**

Specify a unique name for the connection. Note that a default name is automatically generated based on the database and table/view name.

Step 6: Give the connection a name for use in Tableau:

FoodMart 2000 [Warehouse] Connection

OK Cancel

This example discusses how to connect Tableau to a Microsoft SQL Server database. Connecting to this data source requires the Professional Edition of Tableau. You can see what version you have by selecting **Help > About Tableau**.

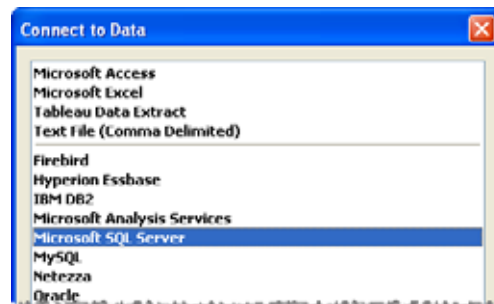
---

**Note** Tableau supports all SQL Server data types except the following: binary objects, varbinary, image, SQL variant, Timestamp, and Unique Identifier.

---

**To connect to an SQL Server database:**

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Microsoft SQL Server**.



### 3 Complete the connection.

Follow the steps on the **SQL Server Connection** dialog box to complete the connection. These steps are described below.

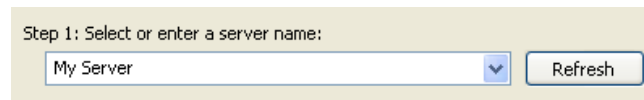
The screenshot shows the 'SQL Server Connection' dialog box with the following steps and settings:

- Step 1: Select or enter a server name:** The server name 'blackcomb' is entered in the text box. A 'Refresh' button is next to it.
- Step 2: Enter information to log on to the server:** The 'Use Windows NT Integrated security (preferred)' radio button is selected. Below it are empty text boxes for 'Username:' and 'Password:'.
- Step 3: Test the connection:** A 'Test Connection' button is present.
- Step 4: Select a database on the server:** The database 'pubs' is selected in the dropdown menu.
- Step 5: Define the connection:** The 'Single Table' radio button is selected. Below it is a list of tables: authors, discounts, employee, jobs, pub\_info, publishers, roysched, sales, and titles. The 'employee' table is highlighted.
- Step 6: Give the connection a name for use in Tableau:** The text 'pubs [employee] Connection' is entered in the text box.

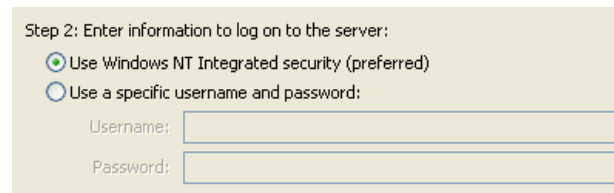
At the bottom right of the dialog box are 'OK' and 'Cancel' buttons.

**Step 1: Select or enter a server name**

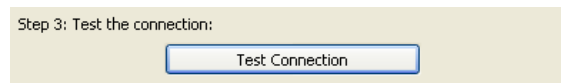
You can select a server from the drop-down menu. This menu will contain all the SQL Server databases on your computer and on the network to which you have access. Alternatively, you can type in the name of a specific server.

A screenshot of the first step in a wizard. It features a label 'Step 1: Select or enter a server name:' followed by a text box containing 'My Server' and a downward arrow icon. To the right of the text box is a button labeled 'Refresh'.**Step 2: Enter information to log on to the server**

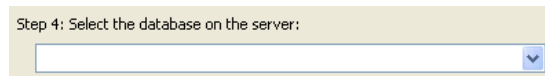
Specify whether the database is password protected. If it is password protected, you must type your Username and Password.

A screenshot of the second step in a wizard. It features a label 'Step 2: Enter information to log on to the server:'. Below the label are two radio button options: 'Use Windows NT Integrated security (preferred)' (which is selected) and 'Use a specific username and password:'. Below the second option are two text boxes labeled 'Username:' and 'Password:'.**Step 3: Test the connection**

Click **Test Connection**. If the connection is unsuccessful, verify that your user name and password are correct. If the connection continues to fail, your computer may be having trouble locating the server. Contact your network administrator or database administrator.

A screenshot of the third step in a wizard. It features a label 'Step 3: Test the connection:' followed by a button labeled 'Test Connection'.**Step 4: Select the database on the server**


Servers often contain multiple databases. Select the specific database on the server that you want to connect to.

A screenshot of the fourth step in a wizard. It features a label 'Step 4: Select the database on the server:' followed by a text box with a downward arrow icon.




**Step 5: Select a table or view from the database**

SQL Server databases often contain multiple tables and views. Specify a single table or view within the database you want to connect to. You can also connect to multiple tables or a specific query by selecting the Multiple Tables and Custom SQL options. Refer to “Joining Tables” on page 5-1 to learn more about connecting to a set of relational tables that are related by join conditions.

**Step 6: Give the connection a name for use in Tableau**

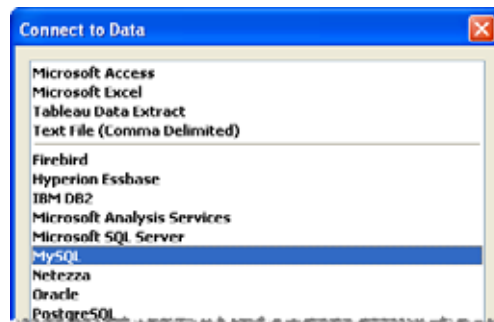
Specify a unique name for the connection. Note that a default name is automatically generated based on the database and table/view name.



This example discusses how to connect Tableau to a MySQL database. Connecting to this data source requires the Professional Edition of Tableau. You can see what version you have by selecting **Help > About Tableau**.

**To connect to a MySQL database:**

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **MySQL**.



- 3 Complete the connection.

Follow the steps on the **MySQL Connection** dialog box to complete the connection. These steps are described below.

The screenshot shows the 'MySQL Connection' dialog box with the following steps:

- Step 1: Enter a server name.** The text 'blackcomb' is entered in the server name field, and '3306' is in the port field.
- Step 2: Enter information to log on to the server:** The username 'test' is entered, and the password field is empty.
- Step 3: Test the connection:** A 'Test Connection' button is visible.
- Step 4: Select a database on the server:** The dropdown menu shows 'test' selected.
- Step 5: Select a table or view from the database:** The 'Single Table' radio button is selected. A list of tables is shown, with 'Coffee Chain' highlighted.
- Step 6: Give the connection a name for use in Tableau:** The text 'test [Coffee Chain] Connection' is entered in the name field.

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

### Step 1: Enter a server name

Type in the name of a specific server.

This close-up shows Step 1 of the dialog. The text 'Step 1: Enter a server name.' is at the top. Below it, the text 'blackcomb' is entered in the server name field, and '3306' is in the port field.


### Step 2: Enter information to log on to the server

Specify the username and password.


This close-up shows Step 2 of the dialog. The text 'Step 2: Enter information to log on to the server:' is at the top. Below it, the username 'test' is entered, and the password field is filled with dots.

**Step 3: Test the connection**


If the connection is successful, Tableau will alert you. If the connection is unsuccessful, verify that your user name and password are correct. If the connection continues to fail, your computer is having trouble locating the server. Contact your network administrator or DBA.

**Step 4: Enter a database on the server**

Servers often contain multiple databases. Select the specific database on the server that you want to connect to.

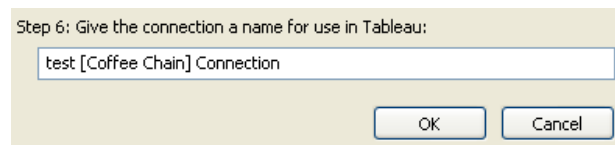
**Step 5: Enter a table or view from the database**

MySQL databases often contain multiple tables and views. Specify which table or view within the database you want to connect to. You can also connect to multiple tables or a specific query by selecting the Multiple Tables and Custom SQL options. Refer to “Joining Tables” on page 5-1 to learn more about connecting to a set of relational tables that are related by join conditions.



**Step 6: Give the connection a name for use in Tableau**

Specify a unique name for the connection. Note that a default name is automatically generated based on the database name.



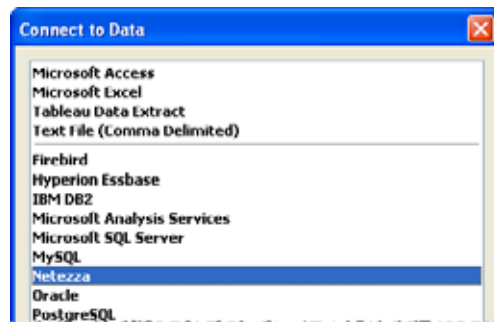
A screenshot of a dialog box titled "Step 6: Give the connection a name for use in Tableau:". The dialog box has a light beige background. Inside, there is a text input field containing the text "test [Coffee Chain] Connection". Below the input field, there are two buttons: "OK" and "Cancel".

## Netezza

This example discusses how to connect Tableau to a Netezza database. Connecting to this data source requires the Professional Edition of Tableau. You can see what version you have by selecting **Help > About Tableau**.

**To connect to a Netezza database:**

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Netezza** and click **OK**.



- 3 Complete the connection.

Follow the steps on the **Netezza Connection** dialog box to complete the connection. These steps are described below.

The screenshot shows the 'Netezza Connection' dialog box with the following steps:

- Step 1: Enter a server name:** The text box contains 'netezza.server.lan' and the Port is set to '5480'.
- Step 2: Enter a database on the server:** The text box contains 'SALES'.
- Step 3: Enter information to log on to the database:** The Username text box contains 'username' and the Password text box contains seven dots.
- Step 4: Test the connections:** A 'Test Connection' button is present.
- Step 5: Select a table or view from the database:** Radio buttons for 'Single Table' (selected), 'Multiple Tables', and 'Custom SQL' are shown. Below is a list box containing: Batters, Cals, DateBins, DateTime, Election (highlighted), FischerIris, Loan, NumericBins, and nrt.
- Step 6: Give the connection a name for use in Tableau:** The text box contains 'Election Connection'.

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

### Step 1: Enter a server name

Type in the name of a specific server and the port to which you want to connect. In most cases, the default value for Port will be the correct port number.

This close-up shows Step 1 of the dialog. The text box for the server name contains 'netezza.server.lan' and the Port text box contains '5480'.

**Step 2: Enter a database on the server**

Servers often contain multiple databases. Type the name of a specific database on the server that you want to connect to.

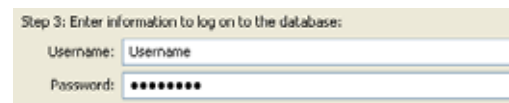


Step 2: Enter a database on the server:

SALES

**Step 3: Enter information to log on to the server**

Specify the username and password.



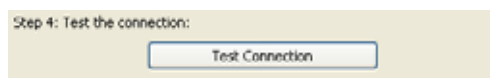
Step 3: Enter information to log on to the databases:

Username: Username

Password: .....

**Step 4: Test the connection**

If the connection is unsuccessful, verify that your user name and password are correct and you have all of the necessary drivers installed. If the connection continues to fail contact your network or database administrator.

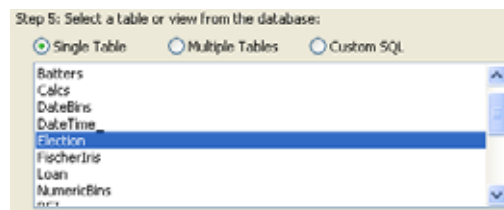


Step 4: Test the connection:

Test Connection

**Step 5: Select a table or view from the database**

Netezza databases can contain multiple tables and views. Specify which table or view within the database you want to connect to. You can also connect to multiple tables or a specific query by selecting the Multiple Tables and Custom SQL options. Refer to “Joining Tables” on page 5-1 to learn more about connecting to a set of relational tables that are related by join conditions.



Step 5: Select a table or view from the databases:


☒ Single Table ☐ Multiple Tables ☐ Custom SQL

Batters  
Cals  
DateBins  
DateTime  
Election  
FischerIns  
Loan  
NumericBins  
mvt



**Step 6: Give the connection a name for use in Tableau**

Specify a unique name for the connection. Note that a default name is automatically generated based on the database name.



## Oracle Database

This example discusses how to connect Tableau to an Oracle relational database. Connecting to this data source requires the Professional Edition of Tableau. You can see what version you have by selecting **Help > About Tableau**.

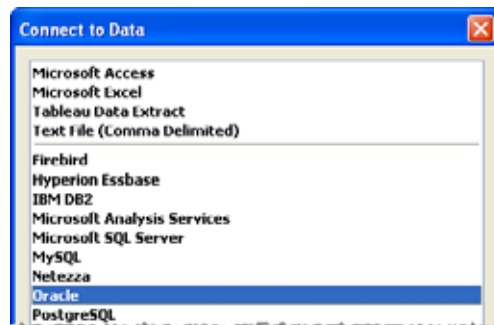
---

**Note** Tableau supports all Oracle data types except CLOB, NCLOB, binary objects, and database objects.

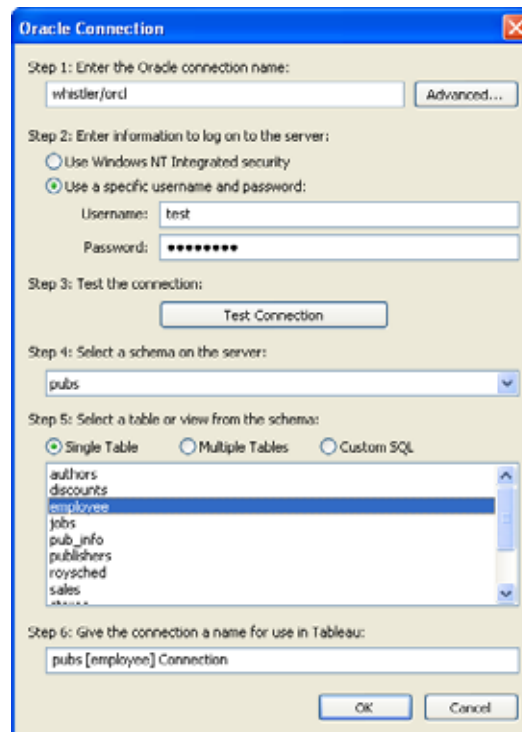
---

### To connect to an Oracle database:

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Oracle**.




- 3 Complete the connection. The information in this dialog box is specific to your server and database.



Follow the steps on the **Oracle Connection** dialog box to complete the connection. These steps are described below.

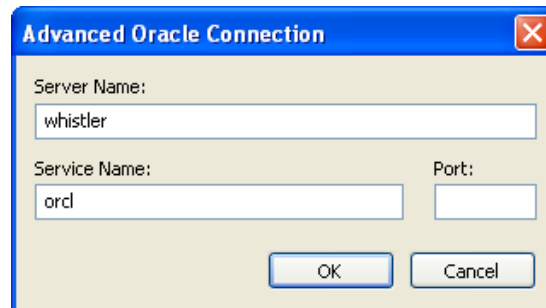
### Step 1: Enter the Oracle connection name

Type the Oracle connection name into the text box.



If you do not know the exact connection string to type, click the advanced button. In the Advance Oracle Connection dialog box, type the server name, service name, and

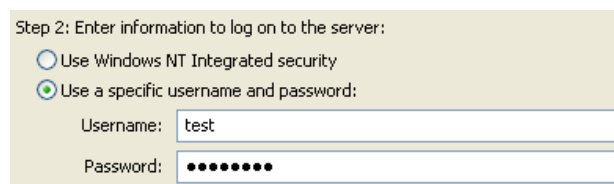
optionally specify the port. Then click **OK**. The connection name will be created based on these variables.

A screenshot of the 'Advanced Oracle Connection' dialog box. It has a blue title bar with the text 'Advanced Oracle Connection' and a red close button. The dialog contains three text input fields: 'Server Name:' with the value 'whistler', 'Service Name:' with the value 'orcl', and 'Port:' which is empty. At the bottom right, there are two buttons: 'OK' and 'Cancel'.

**Note** In order to use your net services definitions in Tableau, you must set either TNS\_ADMIN or ORACLE\_HOME as an environment variable. To set TNS\_ADMIN as the environment variable use the full path of the directory that contains the tnsnames.ora file. To set ORACLE\_Home as an environment variable use the path to the main Oracle directory.

### Step 2: Enter information to log on to the server

Specify whether you are using Windows NT integrated security or a specific username and password. If you are using a specific username and password, type them into the appropriate text boxes.

A screenshot of the 'Step 2: Enter information to log on to the server' dialog box. It has a title bar with the text 'Step 2: Enter information to log on to the server:'. Below the title bar, there are two radio buttons. The first is 'Use Windows NT Integrated security' and the second is 'Use a specific username and password:'. The second radio button is selected. Below the radio buttons, there are two text input fields: 'Username:' with the value 'test' and 'Password:' with a masked password represented by eight dots.

### Step 3: Test the connection

Click **Test Connection**. If the connection is unsuccessful, verify that your user name and password are correct. If the connection continues to fail, your computer may be

having trouble locating the server. Contact your network administrator or database administrator.

Step 3: Test the connection:

Test Connection

#### Step 4: Select a schema on the server

Servers often contain multiple schemas. Select the specific schema on the server that you want to connect to.

Step 4: Select a schema on the server:

pubs

#### Step 5: Select a table or view from the schema.

Oracle schemas often contain multiple tables. Specify which table within the schema that you want to connect to. You can also connect to multiple tables or a specific query by selecting the Multiple Tables and Custom SQL options. Refer to “Joining Tables” on page 5-1 to learn more about connecting to a set of relational tables that are related by join conditions.

Step 5: Select a table or view from the schema:

☒ Single Table ☐ Multiple Tables ☐ Custom SQL

authors  
discounts  
employee  
jobs  
pub\_info  
publishers  
roysched  
sales  
stamps

#### Step 6: Give the connection a name for use in Tableau

Specify a unique name for the connection. Note that a default name is automatically generated based on the database and table or view name.

Step 6: Give the connection a name for use in Tableau:

pubs [discounts] Connection

OK

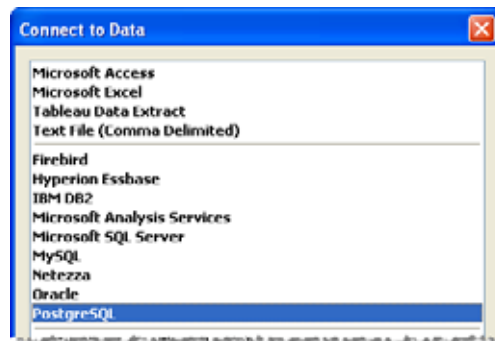
Cancel

## PostgreSQL Database

This example discusses how to connect Tableau to a PostgreSQL database. Connecting to this data source requires the Professional Edition of Tableau. You can see what version you have by selecting **Help > About Tableau**.

**To connect to a PostgreSQL database:**

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **PostgreSQL**.



- 3 Complete the connection.

Follow the steps on the **PostgreSQL Connection** dialog box to complete the connection. These steps are described below.

The screenshot shows a 'PostgreSQL Connection' dialog box with the following steps:

- Step 1: Enter a server name.** The text box contains 'postgres8.test.tsl.lan' and the Port is '5432'.
- Step 2: Enter a database on the server:** The text box contains 'TestV1'.
- Step 3: Enter information to log on to the**
  - Username: 'test'
  - Password: '\*\*\*\*\*'
- Step 4: Test the connection:** A 'Test Connection' button is present.
- Step 5: Select a table or view from the database:**
  - Radio buttons: ☒ Single Table, ☐ Multiple Tables, ☐ Custom SQL.
  - Table list: Batters, Calcs, DateBins, DateTime, Election (selected), FischerTris, Loan, NumericBins, etc.
- Step 6: Give the connection a name for use in Tableau:** The text box contains 'Election Connection'.

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

### Step 1: Enter a server name

Type in the name of a specific server.

This close-up shows Step 1 of the dialog box. The text box contains 'postgres8.test.tsl.lan' and the Port is '5432'.

**Step 2: Enter a database on the server**

Servers often contain multiple databases. Type the specific database on the server that you want to connect to.

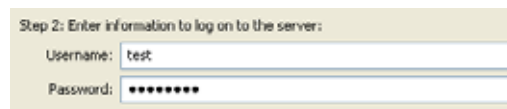


Step 2: Enter a database on the server:

TestV1

**Step 3: Enter information to log on to the server**

Specify the username and password.



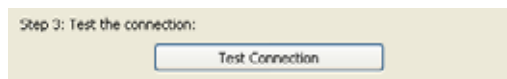
Step 3: Enter information to log on to the server:

Username: test

Password: \*\*\*\*\*

**Step 4: Test the connection**

If the connection is successful, Tableau will alert you. If the connection is unsuccessful, verify that your user name and password are correct. If the connection continues to fail, your computer is having trouble locating the server. Contact your network administrator or DBA.

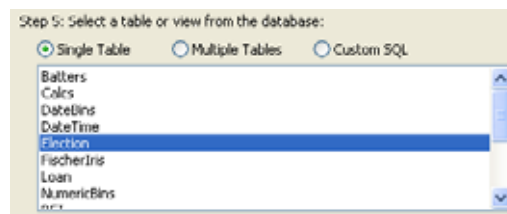


Step 4: Test the connection:

Test Connection

**Step 5: Enter a table or view from the database**

PostgreSQL databases can contain multiple tables and views. Specify which table or view within the database you want to connect to. You can also connect to multiple tables or a specific query by selecting the Multiple Tables and Custom SQL options. Refer to “Joining Tables” on page 5-1 to learn more about connecting to a set of relational tables that are related by join conditions.



Step 5: Select a table or view from the database:

☒ Single Table ☐ Multiple Tables ☐ Custom SQL

Balters  
Colis  
DateBins  
DateTime  
Election  
FischerIns  
Loan  
NumericBins  
Loan



**Step 6: Give the connection a name for use in Tableau**

Specify a unique name for the connection. Note that a default name is automatically generated based on the database name.




**IBM OLAP Server**


To connect to an IBM DB2 OLAP Server, follow the instructions for Hyperion Essbase as described in “Hyperion Essbase Cube (aka IBM OLAP Server)” on page 4-29.

## Connecting to a Custom SQL Query

For most relational data sources you can connect to a specific query rather than the entire data source. Often this can be useful when you know exactly the information you need and you understand how to write SQL queries.

**To connect to a custom SQL query:**

- 1 Select **Custom SQL** in the connection dialog box.
- 2 Type or paste the query into the text box. The  button in the upper right corner of the text box opens a larger editing window for more complex queries.



When you finish the connection, only the relevant fields display in the Tableau Data window.

If your SQL query references duplicate columns, you may get errors when trying to use one of the columns in Tableau. This will happen even if the query is valid. For example, consider the following query:

```
SELECT * from authors, titleauthor where authors.au_id = titleauthor.au_id
```

The query is valid, but the **au\_id** field is ambiguous because it exists in both the “authors” table and the “titleauthor” table. Tableau will connect to the query but you will get an error anytime you try to use the **au\_id** field. That’s because Tableau doesn’t know which table you are referring to.

## Editing the Connection

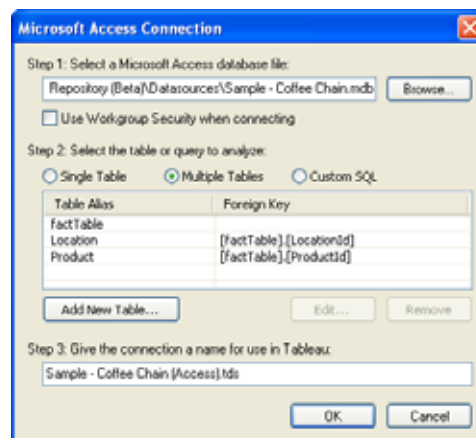
When a Tableau workbook is connected to a data source, you can edit the connection. You might want to edit the data source connection to:

- Specify a new location for the data source.  
Suppose the name or location of a data source you were using has changed and is no longer available using the previous connection information. In this case, you can direct the workbook to the correct location without losing your work.
- Apply analyses created using one data source to another data source.


Suppose you create a workbook containing several views involving markets, products, sales, and profits and you want to apply the analyses to a new data source. Instead of recreating each view from scratch, you can edit the original data connection and specify a new data source.

### To edit a data source connection:

- 1 Select **Data > Connection > Edit**.
- 2 Complete the data source-specific connection dialog box. For example, the **Microsoft Access Connection** dialog box is shown below. You can specify a new file, or you can select a different table to analyze.



## Replacing Field References

When you successfully connect to a new data source, all worksheets that previously referred to the original data source now refer to the new data source. If the new data source does not have the same field names that are used in the original workbook, the fields become invalid and are marked with an exclamation mark . You can quickly resolve the problem by replacing the field's references.

For example, say you have a workbook connected to a data source that contains a Product Category field. Then you edit the connection to point to a new data source that has all the same data but instead of Product Category, the field name has been changed to Product Type. The Product Category field remains in the Data window but is marked as invalid. To make the field valid, you can replace the references, which means you can map the invalid field to a valid field in the new data source (e.g., Product Category corresponds to Product Type).

### To replace the references for a field:

- 1 Right-click the invalid field in the Data window and select **Replace References**.
- 2 In the Replace References dialog box, select a field from the new data source that corresponds to the invalid field.


## Exporting the Connection

When you connect to a data source for the first time, Tableau asks you if you want to remember the connection by saving a copy in your repository. This copy acts as a shortcut allowing you quickly connect without opening the connection dialog box. Shortcuts appear in the **Connect to Data** dialog box as shown in “Connect Using the Connect to Data Dialog Box” on page 4-7. If you connect to the same data source multiple times, you should opt to remember the connection so that you can quickly reconnect.

At any time while connected to a data source, you can export the connection information. You might want to do this if:

- You did not choose to remember the connection when you first connected to the data source.
- You added custom fields to the Data window such as groups, sets, calculated fields, and binned fields or added joined tables. For subsequent connections to that source, use the shortcut so that you don't have to recreate the custom fields. Note that you can also save custom fields by saving the workbook or by creating a bookmark.

The connection is saved as a TDS (Tableau data source) file. Each TDS file references one data source, and stores information about the data source type and location, and any custom fields you have created. TDS files do not store data, workbooks, or worksheets.

TDS files are stored in the Datasources folder of the Tableau Repository. They have the .tds extension and can be identified by the  icon. If you move the TDS file to another location, you cannot access the file with the **Connect to Data** dialog box. However, you can access the file by selecting **File > Open** and navigating to the file. You can also connect by dragging the TDS file onto Tableau's desktop icon or onto the running application.


If you move the data source referenced by the TDS file, you will not be able to complete the connection. In this case, Tableau will ask you to replace the original data source with another one. Note that the replacement data source must be of the same type (Excel, Access, and so on) as the original.

**To export the connection:**

- 1 Select **Data > Connection > Export**.
- 2 Complete the **Export Data Source** dialog box by specifying a file name.



As shown below, the new TDS file is displayed in the **Connect to Data** dialog box.



The new shortcut appears in the dialog box.

## Refreshing the Data

If you make modifications to a data source such as adding new fields or rows, changing data values or field names, or deleting data, Tableau will reflect those changes the next time you connect to the data source. However, because Tableau queries a data source and does not import the data, you can immediately update Tableau to reflect the data source modifications without disconnecting.

---

**Note** You cannot modify an Excel file while it is connected to Tableau.

---

If you are connected to a data source that has been modified, you can immediately update Tableau with the changes by selecting **Data > Refresh**.

If you remove a field from a data source that is used in a Tableau worksheet, and then refresh the data source, a warning message displays indicating that the field will be removed from the view and the worksheet will not display correctly because of the missing field.


## Closing the Connection

You can close a data connection at any time. Doing so does not modify the data source. Instead, it disconnects Tableau from the data so that you can no longer query it. Additionally, the connection is cleared from the Data window and all open worksheets associated with the data are cleared. If you accidentally close a connection, use the Undo button to reconnect.


### To close a data connection:

- 1 Select the data source in the Data window menu.

The connection you want to close must be displayed in the Data window.



- 2 Select **Data > Connection > Close**.





# Joining Tables

---

<b>Overview</b>	<b>5-2</b>
<b>Adding Tables</b>	<b>5-2</b>
Connecting to Multiple Tables	5-2
Adding Tables to the Data window	5-2
<b>Editing Tables</b>	<b>5-2</b>

## Overview

Many relational data sources are made up of a collection of tables that are related by specific fields. For example, a data source for a publisher may have a table for authors that contains the first name, last name, phone number, etc. of clients. In addition, there may be another table for titles that contains the price, royalty, and title of published books. In order to analyze these two tables together, to answer questions like, how much was paid in royalties last year for a particular author, you would join the two tables using a common field such as Author ID. That way you can view and use the fields from both tables in your analysis. This section discusses the following topics:

- Adding Tables
- Editing Tables


## Adding Tables

You can connect to multiple tables when you first connect to a data source using the connection dialog box. There you can add tables, specify joins, and modify the field aliases in the case you have similarly named fields in each of the tables. You can also add tables after you have already connected to the data source. This section describes how to connect to multiple tables as well as adding tables to the Data window.

- Connecting to Multiple Tables
- Adding Tables to the Data window

## Connecting to Multiple Tables

- 1 In the Connection dialog box, complete the connection information according to the examples described in “Examples – Connecting to Data Sources” on page 4-11.
- 2 Select the table or view you want to start with (typically the fact table) and then select the **Multiple Tables** option.



# Extracting Data

---

Overview . . . . .	6-2
Creating an Extract . . . . .	6-3
Using Extracts . . . . .	6-4

---

## Overview

Extracts are saved subsets of the whole data source that you can use to improve performance, upgrade your data to allow for more advanced capabilities, and perform offline analysis. You can create an extract by defining filters and limits that include the data you want in the extract. Extracts allow you to analyze data stored in large data sources that would otherwise take too long to query over and over.

Extracts are useful in the following situations:

- If you are using a file based data source such as Excel or Access, extracts can improve performance even when you extract the entire data source.
- Extracts add functionality to file based data sources, such as the count distinct aggregation.
- If you are using a large data source, you can improve performance and limit the load on the server by using extracts to look at a subset of your data.
- If you are travelling and need to access your data offline, you can extract the relevant data to a local data source.

This section discusses the following topics:

- Creating an Extract
- Using Extracts

## Creating an Extract

- 1 Select **Data** > **Extract** to open the Extract Data dialog box.



- 2 Optionally define the extract if you are not extracting the entire data source:
  - Click **Add** to define a filter that includes the data you want to extract from the data source. Refer to “Configuring the Filter Dialog Box” on page 12-21 for more information about defining filters.
  - Select the number of rows you want to extract. You can extract All, a Percentage, or a specific number of rows from the data source. Note that Tableau first applies any filters and then extracts the number of rows from the filtered results.
- 3 When finished, click **Extract**.
- 4 In the subsequent dialog box, select a location to save the extract into and give the file a name. Then click **Save**.

---

**Note** Depending on the size of your data source, extracting data can take a long time. However, after you have extracted the data and saved it to your hard drive, performance will improve.

---

## Using Extracts

After you create an extract, the current workbook begins using the extract. However, the extract connection is not saved with the workbook until the next time you save. That means, if you close the workbook without saving first, the workbook will connect to the original data source the next time you open it.

You can toggle between using the extract and using the entire data source by selecting **Data > Use Extract**. Often you will want to create an extract with a sample of the data so you can set up the view and then switch to the whole data source. That way you won't have to perform long queries every time you place a field on the shelf.

You can remove an extract at anytime by selecting **Data > Remove Extract**. When you remove an extract you can choose to **Remove the extract** or **Remove and delete the extract file**. If you choose the latter, the extract file will be deleted from your hard drive.

---

**Note** When extracting data, make sure you save the file to your local hard drive. Tableau does not support extracting data to a network drive.

---

# Managing Queries

---

Overview . . . . .	7-2
Automatic Updates . . . . .	7-3
Cancel Query . . . . .	7-4
Abandoned Queries . . . . .	7-5
Precision Warnings . . . . .	7-6

---

## Overview



Queries are automatically generated every time you add a field to a shelf and interact with the view. Tableau offers several ways you can manage these queries once they are sent to the underlying data source. The following are some topics to understand when managing queries:


- Automatic Updates - switch between updating the view automatically and running updates manually.
- Cancel Query - sometimes a query may be taking longer than you expected, you can cancel the query so you can begin working in Tableau again.
- Abandoned Queries - Even though you may cancel a query in Tableau the query still executes on the data source. These are called abandoned queries.
- Precision Warnings - sometimes your fields contain more precise numbers than Tableau can display. In these cases, a precision warning is given.



## Automatic Updates

When you place a field on a shelf, Tableau generates the view by querying the data source. If you are creating a dense data view that involves many fields, the queries might be time consuming and significantly degrade system performance. In this case, you can instruct Tableau to turn off automatic updates.

By default, automatic updates are turned on and the toolbar button is highlighted . However, it is sometimes more efficient for Tableau to execute the queries you need for your final view, rather than for every intermediate step required to compose that view. You can turn off updates by pressing **F10** or the **Automatic Updates** toolbar button .

While automatic updates are turned off, you can still update the view at any time by clicking **F9** or the **Run Update**  on the toolbar. This way, you can update your data view at an intermediate step. It is possible to enter an invalid state when automatic updates are turned off. When this happens, the view is desaturated and invalid commands are disabled. The view and commands become available again when you click Run Update on the toolbar.

For example, the view below has automatic updates turned off. When the aggregation for Inventory is changed from a summation to an average, the view is desaturated to let you know that you have made a change to the view that has made the current view invalid.




## Cancel Query

This command is used any time you want to stop a query that is in process. You may want to cancel a query that is taking a long time to complete due to the size of the data source. When a query is taking a long time to complete, a progress dialog box opens. You can cancel a requested query by clicking the **Cancel** button on the Processing Request dialog box.

**To cancel a query:**

- 1 Click **Cancel** in the Processing Request dialog box.




- 2 After canceling a query the view becomes invalid because it is in an in-between state. The result is a blank view although all your fields are still on the shelves. To resume working with Tableau, alter the view in anyway and allow the query to complete.

---

**Note** Canceling a large number of queries can result in performance degradation in the underlying database. Although the query has been abandoned by Tableau it is still executing on the database. Refer to “Abandoned Queries” on page 7-5 for information on how to monitor cancelled queries.

---

## Abandoned Queries


When you cancel a query in Tableau, the database is told to stop processing the query. However, some databases do not support cancel (MS Excel, MS Access, Essbase, SSAS 2000). If you cancel a query using one of these types of data sources, the query is abandoned by Tableau but is still running in the background and using resources. When you have abandoned queries, an indicator appears in the bottom right corner of the workbook showing the number of queries still running . As queries in the background complete, the number will go down. It is important to monitor the number of queries running and not let the number get too high, otherwise you will see performance degradation of both Tableau and the underlying database.

---

**Note** Text, Microsoft Excel, and Microsoft Access data sources may be temporarily unavailable after cancelling a query because of a lock performed internally. You may have to wait until the abandoned query has completed before re-connecting.

---

## Precision Warnings

When you add a field to a view that contains values with more precision than Tableau can model, a warning icon  is displayed in the bottom right corner of the status bar. For example, a value in the database may have 22 decimal places but Tableau only supports up to 15. When you add that field to the view, you get a precision warning. If you click on the warning, you can read more details including the number of decimal places that have been truncated in the view.

Remember that the precision of the data displayed in Tableau will always first be dependent on the data in your database. If the values in your database exceed 15 decimal places, when you add them to the view, the value is truncated and a precision warning appears.

# Understanding Data Fields

<b>Understanding the Data Window . . . . .</b>	<b>8-3</b>
Relational and Multidimensional Data . . . . .	8-5
Measure Values and Measure Names . . . . .	8-7
Number of Records . . . . .	8-7
 <b>Data Window Features and Functions . . . . .</b>	 <b>8-8</b>
Organize the Data window . . . . .	8-8
Find Fields . . . . .	8-8
Rename Fields . . . . .	8-9
Hide or Unhide Fields . . . . .	8-9
Add Fields to the Data window . . . . .	8-10
 <b>Renaming Fields and Dimension Members . . . . .</b>	 <b>8-11</b>
Renaming Fields . . . . .	8-11
Member Aliases . . . . .	8-11
Example - Editing Member Aliases . . . . .	8-13
Changing Measure Names . . . . .	8-17
 <b>Data Types &amp; Roles . . . . .</b>	 <b>8-20</b>
Data Types . . . . .	8-20
Data Roles . . . . .	8-22
Example: Continuous Dimensions and Discrete Measures . . . . .	8-28

---

## Overview

The data in all data sources are categorized into fields such as Customer, Sales, Profit, Temperature, etc. When you connect to a data source with Tableau, the fields are displayed in along the left side of the workbook in the Data window. The fields are what you will use to build views of your data. Each field is automatically assigned a data type (such as integer, string, and date) and a pair of data roles. This section discusses the following topics:

- Understanding the Data Window
- Data Window Features and Functions
- Renaming Fields and Dimension Members
- Data Types & Roles

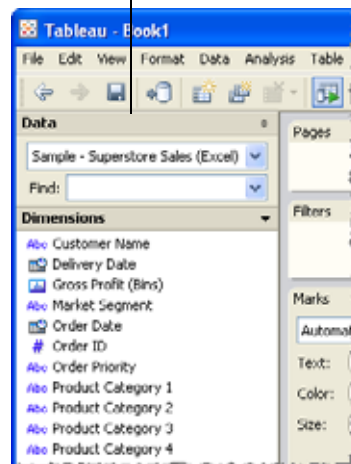
## Understanding the Data Window

All data sources contain fields. In Tableau, these fields appear in the Data window. For multidimensional data sources, the fields are given by the dimensions and measures of a cube. For relational data sources, the fields are given by the columns of a table or view. Each field contains a unique attribute of the data such as customer name, sales total, product type, and so on. For example, some of the fields of an Excel worksheet are shown below.

	A	E	C	Q	
1	Record No	Order Priority	Sales Total	Order Date	Fields
2	1	5-LOW	2302.42	1/12/2001	
3	2	5-LOW	896.68	1/12/2001	
4	3	5-LOW	2564.85	1/12/2001	
5	4	5-LOW	226.52	1/21/2001	
6	5	3-MEDIUM	158.88	1/21/2001	
7	6	3-MEDIUM	2679.61	1/21/2001	
8	7	3-MEDIUM	999.31	1/21/2001	
9	8	3-MEDIUM	56.95	1/21/2001	
10	9	3-MEDIUM	10.42	1/21/2001	
11	10	3-MEDIUM	4129.53	1/21/2001	

After you connect to a data source with Tableau, the data source fields appear on the left side of the workbook in the Data window.

Data window



The Data window organizes the fields into three areas:

- **Dimensions** – Fields that typically hold discrete qualitative data. Examples of dimensions include dates, customer names, and market segments (refer to “Data Types & Roles” on page 8-20).
- **Measures** – Fields that typically hold numerical data that can be aggregated. Examples of measures include sales, profit, number of employees, temperature, frequency, and pressure (refer to “Data Types & Roles” on page 8-20).
- **Sets**– An additional area that stores custom fields based on existing dimensions and criteria that you specify. Named sets from an SSAS Server also appear in Tableau in this area of the Data window. You can interact with these named sets in the same way you interact with any other custom set in Tableau. Refer to “Sets” on page 12-70.

For multidimensional data sources, fields are explicitly defined as dimensions or measures when the database is created. For relational data sources, Tableau automatically organizes the fields. By default, fields containing text, date or boolean values are dimensions, while fields containing numerical values are measures. To learn how to change the default organization refer to “Converting Measures to Dimensions” on page 8-24.

The Data window for an Excel worksheet (a relational database) is shown below. The **Discount** and **Gross Profit** fields contain numbers and appear as measures in the Data window. The **Order Priority** field contains text and the **Order Date** field contains dates. These fields appear as dimensions in the Data window.





**Note** By default, the field names defined in the data source are displayed in the Data window. You can rename fields as well as member names (refer to “Renaming Fields and Dimension Members” on page 8-11).

Refer to the topics below to learn more about the Data window:

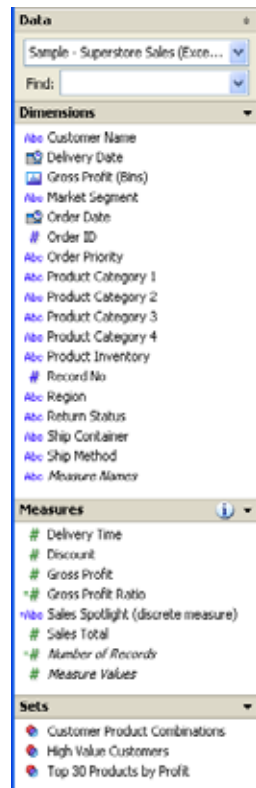
- Relational and Multidimensional Data
- Measure Values and Measure Names
- Number of Records

## Relational and Multidimensional Data

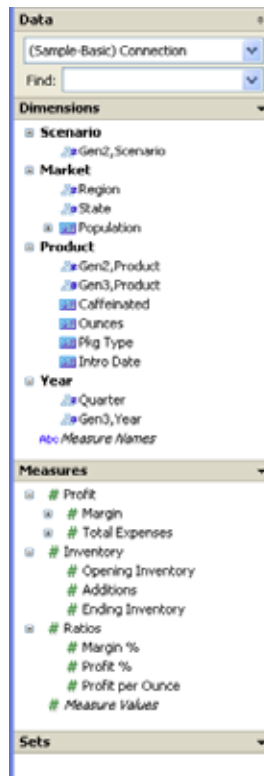
The Data window for a relational and a multidimensional data source are shown below. Note that the windows look essentially the same for both data sources in that the fields are organized into dimensions and measures. However, the multidimensional data source

contains *hierarchies* for the dimensions. For example, notice that the **Product** dimension in the multidimensional Data window contains hierarchical members such as Product Family, Product Department, and so on.

Relational Data



Multidimensional Data




You can expand or collapse the various areas or hierarchies in a multidimensional Data window by clicking the **+** button. You can hide the Data window altogether by selecting **View > Data Window**.

## Measure Values and Measure Names

The Data window contains a few fields that are not part of your data source, two of which are: **Measure Names** and **Measure Values**. The **Measure Values** field always appears at the bottom of the **Measures** area of the Data window and contains all the measures of your data source collected into one field. The **Measure Names** field always appears at the bottom of the **Dimensions** area of the Data window and contains all the names of the measures collected into a single dimension.

Tableau automatically creates these fields so that you can build certain types of data views that involve multiple measures. In particular, use these fields if you want to display multiple measures in the same pane simultaneously. As shown below, creating a view with **Measure Names** and **Measure Values** is one way to display all the data in your data source.



For an example of how to use the measure values and measure names fields refer to “Line Chart-Filter and Color-Encode Multiple Measures” on page 43-34.


## Number of Records

In addition the Measure Names and Measure Values fields, the Data window contains a Number of Records field that is also not part of the underlying data source. This field, represents the number of rows in the data source. It is useful when you are working with a data source that is primarily categorical resulting in very few measures.

## Data Window Features and Functions

### Organize the Data window

You can reorganize the Data window from its default layout by selecting from a variety of sorting options. These **Sort by** options are located in the Data window menu.



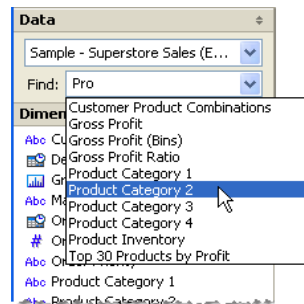
Select one of the following options to sort by:

- **Name** - lists the dimensions and measures in alphabetical order according to their field aliases.
- **Data source order** - lists the dimensions and measures in the order they are listed in the underlying data source.

You can also select to **Group by Table**, which is a command that toggles on and off. When you select this command, the dimensions and measures are grouped according to the table they belong to. This is especially useful when you have several joined tables.

### Find Fields

You can search for fields in the Data window. If there are many fields in your data source, it can be difficult to find a specific one like “Date” or “Customer” or “Gross Profit”. This is especially true for multi-dimensional (OLAP) databases that have large hierarchies. To search for a field, click the Find box at the top of the Data window (Ctrl + F) and type the name of the field you want to search for. Valid field names that fit the description appear in a drop-down list. Select the field you want and press enter on your keyboard to highlight the field in the Data window.

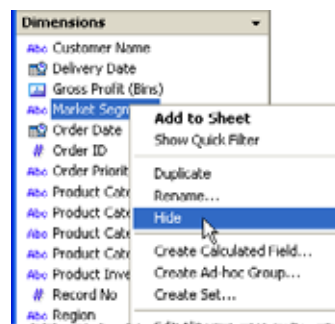


## Rename Fields


You can rename the fields displayed in the Data window by right-clicking the field and selecting **Rename**. The changed field name is saved with the workbook as well as when you export the connection. Renaming a field in Tableau does not modify the underlying data source. Refer to “Renaming Fields and Dimension Members” on page 8-11 for more information about renaming fields and dimension members.

## Hide or Unhide Fields

You can selectively hide or show fields in the Data window. To hide a field, right-click the field you want to hide and select **Hide**.



Notice that when fields are hidden, an information icon displays in the Data window.



Rest the mouse over the information icon to see a tooltip showing the number of fields that are hidden. When you want to change your fields from hidden to visible, select **Unhide Fields** on the Data window menu.



A dialog box opens containing all of your hidden fields. You can select as many or as few hidden fields as you want to show, and then click **OK**.

## Add Fields to the Data window

You can create calculated fields that appear in the Data window. These new computed fields can be used like any other field. Select **Create Calculated Field** on any of the Data window title menus. Alternatively, select **Analysis > Create Calculated Field**. Refer to “Calculated Fields” on page 16-18 for more information.

## Renaming Fields and Dimension Members

Tableau allows you modify both the field and dimension member names that display in the Data window and the view. When you modify the member names of a field the new names are called aliases. An alias is an alternate name assigned to data in a database, usually for the purpose of making the data more understandable. This section discusses the following topics:

- Renaming Fields
- Member Aliases
- Example - Editing Member Aliases
- Changing Measure Names

### Renaming Fields

You can assign an alternate name for a field that displays in the Data window as well as in the view. For instance, a field called Market Segment in the data source could be aliased to appear as Business Segment in Tableau. You can rename both dimensions and measures. Renaming a field does not change the name of the field in the underlying data source, rather it is given a special name that only appears in Tableau workbooks.

#### To rename a field:

- 1 Right-click the field name in the Data window you want to rename and select **Rename**.
- 2 Type the new name in the subsequent dialog box and click **OK**.

The field displays with the new name in the Data window.

#### To revert to the default field name:

- 1 Right-click the field and select **Rename**.
- 2 In the Rename dialog box, type a new name and click **OK**.

### Member Aliases

Member aliases are alternate names for specific members of dimensions. For example, you may want to assign aliases for the members of the “Market Segment” dimension. Perhaps you want the “Consumer” members of this field to display as “Home Consumer” in all views.

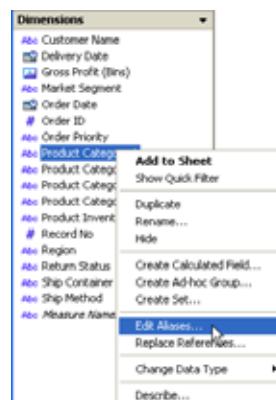
Member aliases can be created for the members of most dimensions in the Data window. You cannot, however, define member aliases for continuous dimensions and dates. Member aliases do not apply to measures. The precise way to use member aliases in Tableau depends on whether you are connected to a relational or a multi-dimensional data source.

The method for creating member aliases depends on the type of data source you are using:

- Relational Database Member Aliases
- Multi-dimensional (OLAP) Database Member Aliases

### Relational Database Member Aliases

To create a member alias for a relational data source, right-click a field name and select **Edit Aliases** from the context menu.



A dialog box opens allowing you to define specific aliases for each member. You can reset the member names back to their original names by clicking the **Clear Aliases** button in the bottom right corner of the Edit Aliases dialog box.


### Multi-dimensional (OLAP) Database Member Aliases

Member aliases for Hyperion Essbase multi-dimensional databases are created on the server by the server administrator and can be activated in Tableau using the **Alias File** option on the **Data** menu. Please talk to your Essbase database administrator. Member aliases are not supported by Microsoft Analysis Services databases.




By default, the alias for every member of every dimension is initially defined to be the original member name. For example, the figure below shows a bar chart built from an Essbase database. By default, the original members names are displayed (example on the left). As you can see, these names are not very intuitive. By selecting **Data > Alias File**, and selecting an appropriate alias file set up by the Essbase administrator, meaningful names are displayed in the headers.

Headers using default member names.



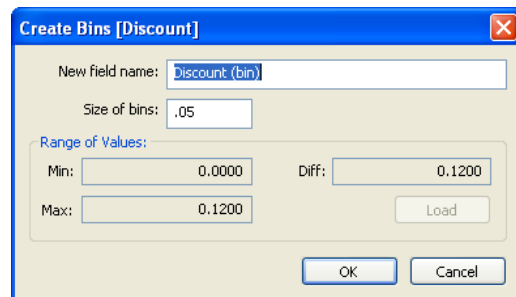
Headers using aliases from file.




## Example - Editing Member Aliases

The Superstore Sales Excel sample data source contains a measure called **Discount**, which contains discount values from 0 to 0.12. Suppose you want to analyze these data by categories: low discount, medium discount, and high discount.

To create the categories, you could first bin the measure so that when added to the view it create discrete headers instead of a continuous axis. In this example, define a bin size equal to 0.05, which produces three bins. The first bin contains the values 0 to 0.05, the second bin contains the values 0.05 to 0.10, and the third bin contains the values 0.10 to 0.15. The **Create Bins** dialog box for this field is shown below. Refer to “Binned Data” on page 16-59 for a complete description of how to bin data.



The new binned field is named **Discount (bin)** and appears in the **Dimensions** area of the Data window.



When you place **Discount (bin)** on the **Rows** or **Columns** shelf, the default aliases for the bins are given by the lower limit of the bin's numerical range



To improve the readability of the bins when they are displayed in Tableau, you can define aliases such as “Low discount”, “Medium discount”, and “High discount.”

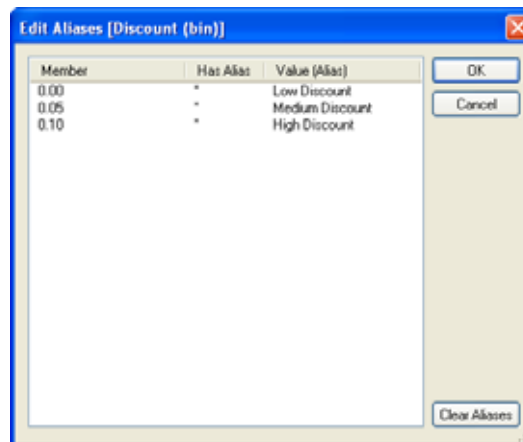
- 1 Open the Edit Aliases dialog box.

Right-click the name of the dimension in the Data window and select **Edit Aliases** on the context menu.




- 2 Define Aliases.

Assign an alias to every member of **Discount (bin)**. For example, the member originally labeled as 0.00 is now labeled as “Low Discount”.



You can change aliases at any time using the Edit Aliases dialog box. To do so, click on the alias you want to change and specify the new name. Use the Tab key to advance from one field to the next. To restore the original aliases, click the **Clear Aliases** button in the bottom right corner of the dialog box. You can also sort the members or their aliases by clicking the appropriate column header.

After completing the Edit Aliases dialog box, Tableau automatically turns on aliases if and displays the aliases in the view.



## Changing Measure Names

There are times that you will want to show multiple measures in a view and so you will use the **Measure Values** and the **Measure Names** fields. When you use Measure Names all of the measure names appear as row or column headers in the view. However, the headers include both the measure name and the aggregation label. So if you are showing the summation of gross profit the header displays as SUM(Gross Profit). You can change the names so that they do not include the aggregation label by editing the member aliases of the **Measure Names** field. This feature becomes particularly useful when you are working with a text table that shows multiple measures. For example, suppose you have a text table containing the aggregated gross profit of each product category by region.

Product Category 1

Region	FURNITURE	OFFICE SUPPLIES	TECHNOLOGY
CENTRAL	18,088	79,411	395,626
EAST	23,612	140,985	1,029,260
WEST	186	66,998	562,005

Now suppose you want to show both the Gross Profit and the Sales Total for each product category and region. When you add the Sales Total measure to the text table, the measures are combined and the **Measure Values** field is placed on the Text shelf. Additionally, the **Measure Names** field is added to the **Rows** shelf.

Product Category 1

Region	Measure Names	FURNITURE	OFFICE SUPPLIES	TECHNOLOGY
CENTRAL	Sum of Gross Profit	18,088	79,411	395,626
	Sum of Sales Total	230,885	395,366	1,071,019
EAST	Sum of Gross Profit	23,612	140,985	1,029,260
	Sum of Sales Total	762,068	893,321	2,814,419
WEST	Sum of Gross Profit	186	66,998	562,005
	Sum of Sales Total	337,957	440,755	1,538,908

Notice how the header names include the aggregation label. Those headers can be annoying if you are putting this view into a presentation. To change the measure names, right-click

the Measure Names field on the Rows shelf and select **Edit Aliases**. Make the changes and click **OK**.

Region	Measure Names	Product Category 1		
		FURNITURE	OFFICE SUPPLIES	TECHNOLOGY
CENTRAL	Gross Profit	18,088	79,411	396,826
	Total Sales	230,005	395,366	1,071,019
EAST	Gross Profit	23,612	140,985	1,029,260
	Total Sales	762,068	893,321	2,814,419
WEST	Gross Profit	186	66,998	562,005
	Total Sales	337,967	440,766	1,638,908

The Measure Names now display without the aggregation.

For more information about using **Measure Values** and **Measure Names** refer to “Measure Values and Measure Names” on page 8-7.

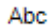




## Data Types & Roles

In Tableau, there are several data types that are supported. For example, you may have text values, date values, numerical values, and more. Each of the data types can take on different roles that dictate their behavior in the view. This section discusses the following:

- Data Types
- Data Roles
- Example: Continuous Dimensions and Discrete Measures

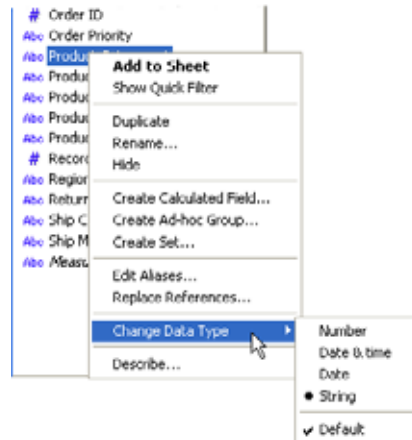
### Data Types

All fields in a data source have a data type. The data type reflects the kind of information stored in that field, for example integers (410), dates (1/23/2005) and strings (“Wisconsin”). The data type of a field is identified in the Data window by one of the icons shown below.

Icon	Description
	Text values
	Date values
	Date & Time values
	Numerical values
	Boolean values (relational only)

Sometimes Tableau may identify a field with a data type that is incorrect. For example, a field that contains dates may be identified as an integer rather than a date. You can change the data type in Tableau by right-clicking the field in the Data window, selecting Change Data Type, and then selecting the appropriate data type.





**Note** Sometimes the data in your database is more precise than Tableau can model. When you add these values to the view a precision warning will appear in the right corner of the status bar. To learn more about precision warnings refer to “Precision Warnings” on page 7-6.

### Mixed Data Types for Excel and CSV Files

Most columns in an Excel or CSV (comma separated value) file contain values of the same data type (dates, numbers, text). When you connect to the file, Tableau creates a field in the appropriate area of the Data window for each column. Dates and text values are dimensions, and numbers are measures.

However, a column might have a mixture of data types such as numbers and text, or numbers and dates. When you connect to the file, the mixed-value column is mapped to a field with a single data type in Tableau. Therefore, a column that contains numbers and dates might be mapped as a measure or it might be mapped as a date dimension. The mapping is determined by the data types of the first 16 rows in the data source. For example, if most of the first 16 rows are text values, then the entire column is mapped as text.

---

**Note** Empty cells also create mixed-value columns because their formatting is different from text, dates, or numbers.

---

Depending on the data type Tableau determines for each field, the field might contain Null values for the other (non matching) records as described in the table below.

Mapped Data Type	Treatment of Other Data Types in the Field
Text	Dates and numbers are treated as text. Nulls are not created.
Dates	Text is treated as Null. A number is treated as the day in numeric order from 1/1/1900. You can identify these values by creating row or column headers with the field.
Numbers	Text is treated as Null. A date is treated as the number of days since 1/1/1900. You can identify these values by converting the measure to a dimension, and then creating row or column headers with the field.

If using fields that are based on mixed-value columns introduces difficulties when analyzing your data, you can:

- Format empty cells in your underlying data source so they match the data type of the column.
- Create a new column in Excel that does not contain mixed values.

## Data Roles

In addition to a data type, every field in Tableau is characterized by two important additional settings that determine the role and behavior of the field when it is placed on a shelf.

To expose the full functionality of Tableau it is useful to control whether a field is a dimension or measure, and continuous or discrete. This section discusses the following topics:

- Data Roles: Dimension vs. Measure
- Converting Measures to Dimensions

- Data Roles: Continuous vs. Discrete
- Converting Discrete to Continuous Quantities

---

**Note** On a multidimensional data source, changing data roles is limited. You can change some measures from continuous to discrete, but in general, data roles on this type of data source cannot be changed.

---

### Data Roles: Dimension vs. Measure

**Dimensions.** Dimensions typically produce headers when added to the rows or columns shelves in the view. By default, Tableau treats any field containing qualitative, categorical information as a dimension. This includes, for instance, any field with text or dates values. However, in relational data sources, the actual definition of a dimension is slightly more complex. A dimension is a field that can be considered an independent variable.

This means that a measure can be aggregated for each value of the dimension. For instance, you might calculate the Sum of “Sales” for every “State”. In this case the State field is acting as a dimension because you want to aggregate sales for each state. The values of Sales is dependant on the State, so State is an independent field and Sales is a dependent field.

Such aggregation could also be computed for numeric fields that are treated as dimensions. For instance, you might want to calculate the SUM of Sales for each “Discount Rate” offered to customers. In this case the Discount Rate field acts as an independent field and the Sales field is dependent even though both fields are numeric. You can use a numeric field as the independent field by first converting the Discount Rate measure to a dimension.

**Measures.** Measures typically produce axes when added to the rows or columns shelves. By default, Tableau treats any field containing numeric (quantitative) information as a measure. However, in relational data sources, the actual definition of a measure is slightly more complex. A measure is a field that is a *dependent* variable; that is, its value is a function of one or more dimensions.

This means that a measure is a function of other dimensions placed on the worksheet. For instance, you might calculate the Sum of “Sales” for every “State”. In this case, the Sales field is acting as a measure because you want to aggregate the field for each state. But measures could also result in a non-numeric result. For instance, you might create a calculated measure called “Sales Rating” that results in the word “Good” if sales are good and “Bad” otherwise. In this case the “Sales Rating” field acts as a measure even though it

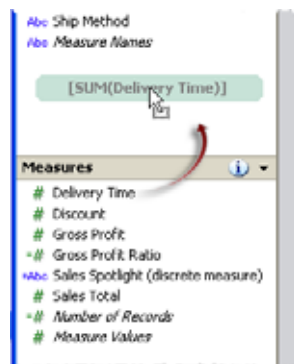
produces a non-numeric result. It is considered a measure because it is a function of the dimensions in the view.

### Converting Measures to Dimensions

By default, Tableau treats all relational fields containing numbers as measures. However, you might decide that some of these fields should be treated as dimensions. For example, a field containing ages may be categorized as a measure by default in Tableau because it contains numeric data. However, if you want to look at each individual age rather than an axis you can convert the Age field to a dimension.

**To convert a measure to a dimension do one of the following:**

- Click and drag the field from the measures area of the Data window and drop it into the dimensions area.




- Right-click the measure in the Data window and select **Convert to Dimension**.



**Note** You can also duplicate the measure before converting it. Duplicate the measure by right-clicking the field in the Data window and selecting Duplicate.

The **Delivery Time** field is now displayed in the **Dimensions** area of the Data window and is a discrete quantity as indicated by the # icon.



If you place the converted field on a shelf, it produces headers instead of an axis.

Columns ▾ Delivery Time						
Rows ▾ Product Category 1						
Delivery Time						
Product Category 1	1	2	3	4	5	6
FURNITURE	Abc	Abc	Abc	Abc	Abc	Abc
OFFICE SUPPLIES	Abc	Abc	Abc	Abc	Abc	Abc
TECHNOLOGY	Abc	Abc	Abc	Abc	Abc	Abc

### Data Roles: Continuous vs. Discrete


In addition to dimensions and measures, each field is categorized as either discrete or continuous. Below are example graphs illustrating the difference between these two data roles. Both examples show the Sum of Margin as a function of Inventory level. It is the same information presented in two different ways.

The Inventory dimension is discrete in this example.



Each inventory value is drawn as a header. Inventory appears **blue** on the Column shelf.

The Inventory dimension is continuous in this example.




Each inventory value is drawn along a continuous axis. Inventory appears **green** on the Columns shelf.


Whether a field is continuous or discrete is reflected in the color of the field's data type icon. In the Data window, blue icons indicate discrete and green icons indicate continuous fields.


Discrete fields always result in headers being drawn whenever they are placed on the row or columns shelves. Continuous fields always result in axes when you add them to the view. These roles are important because you may want to display your data continuously or discretely depending on what you are trying to see and the data itself. You can switch between continuous and discrete data roles. Refer to "Data Roles" on page 8-22.

### Converting Discrete to Continuous Quantities

When you are using a relational data source you can convert any numeric or date field into a continuous field. For example, in "Converting Measures to Dimensions" on page 8-24, the **Delivery Time** field was converted to a dimension. You can now convert this field to a continuous quantity by selecting **Convert to Continuous** from the field's right-click context menu.



The **Delivery Time** field is still displayed in the **Dimensions** area of the Data window, but now uses the  icon, which indicates it is a continuous quantity.



Placing the field on a shelf produces an axis. However, the field is not a measure because you cannot aggregate it using the usual set of aggregation functions such as sum and average. Refer to “Aggregations” on page 16-4 for more information.

---

**Note** You can also convert fields to continuous or discrete directly when they are on a shelf using the field’s menu. This converts the field while it is on the shelf but does not change its role in the Data window. That way you can use the field as continuous for a specific analysis and discrete elsewhere in the view.




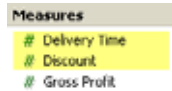
---

### **Example: Continuous Dimensions and Discrete Measures**

If a field in Tableau can have both a data type and have various data roles, you might be wondering about the full range of expressive possibilities. How do the dimension/measure settings interact with the discrete/continuous settings?


The answer is summarized in the following table. In essence, you can create discrete dimensions, continuous dimensions, discrete measures or continuous measures.



	Discrete	Continuous
<b>Dimensions</b>	<p>Discrete dimensions appear with a blue icon under dimensions.</p> <p>Example: Customer Name, Order Date.</p>  <p>Produces headers when placed on a Row or Column shelf. Discrete Dimension Example</p>	<p>Continuous dimensions appear with a green icon under dimensions.</p> <p>Example: Discount</p>  <p>Produces an axis when placed on a Row or Column shelf. Continuous Dimension Example</p>
<b>Measures</b>	<p>Discrete measures appear with a blue icon under measures.</p> <p>Example: Spotlight</p>  <p>Produces headers when placed on a Row or Column shelf. Discrete Measure Example</p>	<p>Continuous measures appear with a green icon under measures.</p> <p>Example: Discount</p>  <p>Produces an axis when placed on a Row or Column shelf. Continuous Measure Example</p>

Here are some examples of these various combinations in use using the sample Superstore Sales-Excel data that comes with Tableau.

### Discrete Dimension



Region is a discrete dimension.

SUM(Sales Total) is a continuous measure.

## Discrete Measure

The screenshot shows a Tableau worksheet with the following configuration:

- Columns:** + YEAR(Order Date)
- Rows:** Product Category 2
- Marks:**
  - Text: SUM(Sales To...)
  - Color: AGG(Sales Sp...)
  - Size: (empty)
  - Level of Detail: (empty)
  - Color: AGG(Sales Sp...)
  - Legend: BAD (blue square), GOOD (orange square)

The data table is as follows:


	2001	2002	2003	2004
APPLIANCES	62,220	18,348	24,961	36,360
BINDERS AND BINDER AC..	48,623	45,457	29,296	34,831
BOOKCASES	26,831	72,820	20,016	26,629
CHAIRS & CHAIRMATS	255,343	116,387	113,460	125,410
COMPUTER PERIPHERALS	156,535	85,714	101,517	73,869
COPIERS AND FAX	187,733	31,841	41,439	68,249
ENVELOPES	21,092	29,070	17,071	7,560
LABELS	19,261	20,729	17,269	17,146
OFFICE FURNISHINGS	63,146	16,941	26,272	20,728
OFFICE MACHINES	234,510	1,661	33,791	200
PAPER	190,735	167,783	171,055	175,713
PENS & ART SUPPLIES	20,410	29,794	14,874	17,971
RUBBER BANDS	4,266	3,672	2,767	3,674
SCISSORS, RULERS AND ..	7,167	8,341	2,443	6,282
STORAGE & ORGANIZATI..	164,866	111,297	79,331	106,096
TABLES	245,185	81,213	70,992	61,626
TELEPHONES AND COMM..	1,116,386	1,064,627	1,002,680	1,224,694

SUM(Sales Total) is a continuous measure.

AGG (Sales Spotlight) is a discrete measure made by a calculation with the formula:

IIF (SUM([Sales Total])>25000, "Good", "Bad") END


## Continuous Dimension



Discount is a continuous dimension.

Sales is a continuous measure.

## Continuous Measure



## Parts of the View

---

<b>Overview</b>	<b>9-2</b>
<b>Headers</b>	<b>9-3</b>
<b>Axes</b>	<b>9-6</b>
<b>Panes</b>	<b>9-9</b>
<b>Cells</b>	<b>9-10</b>
<b>Marks</b>	<b>9-11</b>
Mark Types	9-11
Stacking Marks	9-26
Changing Mark Size and Color	9-31
<b>Titles</b>	<b>9-34</b>
<b>Captions</b>	<b>9-35</b>
<b>Field Labels</b>	<b>9-36</b>
<b>Legends</b>	<b>9-38</b>


# Overview

This section describes the basic components of the views you can create in Tableau. The parts of a view can be categorized as either table components, which are part of every view, or optional components, which can be turned on or off.

## Table Components

Data views are displayed in a table on every worksheet. A table is a collection of rows and columns, and consists of the following components: Headers, Axes, Panes, Cells, and Marks.

In addition to these, you can optionally show Titles, Captions, Field Labels, and Legends.



## Headers

Headers are created when you place a dimension on the **Rows** shelf or the **Columns** shelf. The headers show the member names of each field on the shelves. For example, in the view below the column headers show the members of an Order Date field and the row headers show the members of a Product Category field.

Row Headers		Column Headers			
Columns		[1] YEAR(Order Date)			
Rows		Product Category 2			
		2001	2002	2003	2004
APPLIANCES		10,371	-115	7,385	5,485
BOOKCASES		3	-6,836	-1,067	-2,883
CHAIRS & CHAIRMATS		41,777	6,823	15,615	16,363
COMPUTER PERIPHERALS		45,775	26,060	33,096	25,508
COPIERS AND FAX		96,988	16,490	19,797	33,242
ENVELOPES		7,916	12,926	7,412	3,356
LABELS		9,019	10,283	8,195	8,423



You can show and hide row and column headers at anytime.

**To hide headers:**

- Right-click the headers in the view and select **Show Header**.

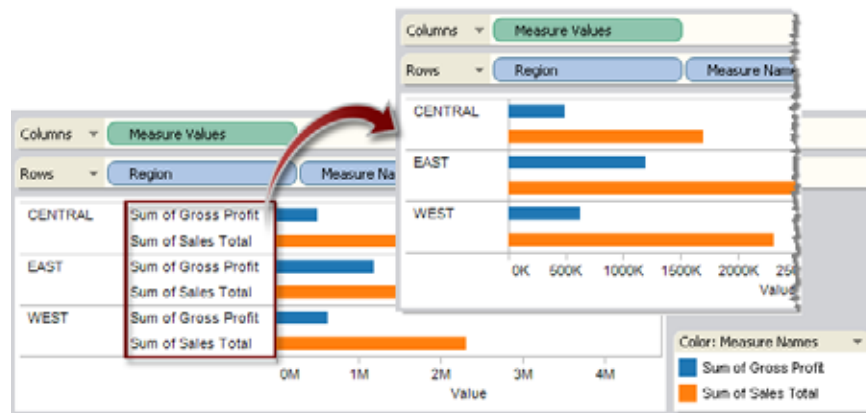


### To show headers:

- Select the field in the view whose headers you want to show and select **Show Header** on the field menu




Hiding headers can be really useful when you are working with multiple measures. For example, the view below shows both the sales and profit for each region along a single axis. You can see the view looks cluttered with the Measure Names headers showing. Because Measure Names is also indicated by the mark color, you can hide the excess headers to clean up the view.



## Axes


Axes are created when you place a measure on the **Rows** or **Columns** shelf. By default, the values of the measure field are displayed along a continuous axis.



You can show and hide axes at anytime.

**To hide axes:**

- Right-click the axis in the view and select **Show Header**.




### To show axes:

- Select the measure in the view whose axis you want to show and select **Show Header** on the field menu




Hiding axes can be really useful when you are working with maps, which often require latitude and longitude axes. For example, the view below shows a map houses sold in a particular neighborhood. By default, the latitude and longitude axes are shown although you may not find them useful. When these axes are hidden, the map view becomes a lot cleaner.



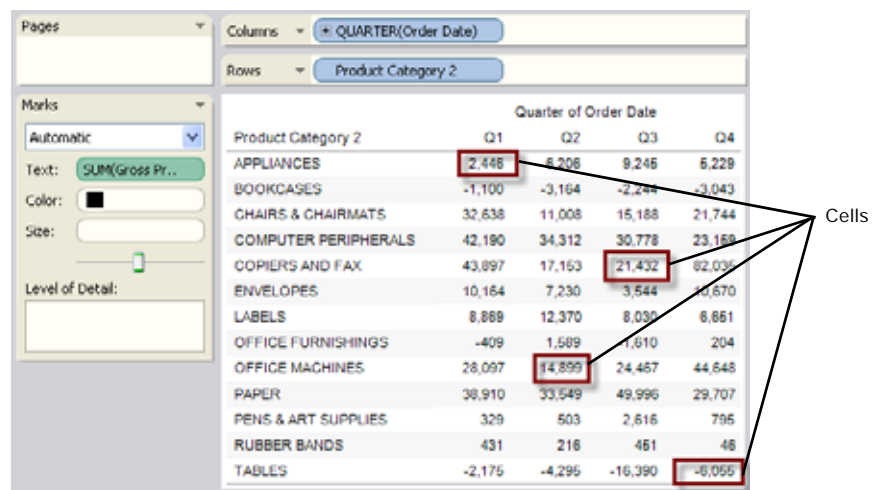
## Panes

Panes are created by the intersection of the rows and columns in a table. Depending on the table type, panes might be created by the intersection of an axis with headers, an axis with an axis, or headers with headers. Panes are identified by lines within the table.



## Cells

Cells are the basic components of any table you can create in Tableau. For a text table, the cell is the intersection of a row and a column, and is where the text is displayed. For other view types such as bar charts and scatter plots, identifying the cell is not always possible or useful.



The screenshot shows the Tableau interface with a table view. The Columns shelf contains 'Quarter of Order Date' and the Rows shelf contains 'Product Category 2'. The Marks card is set to 'Automatic'. The table displays sales data by quarter (Q1, Q2, Q3, Q4) for various product categories. A callout labeled 'Cells' points to several cells in the table, including '2,448' for Appliances in Q1, '21,432' for Copiers and Fax in Q3, '14,899' for Office Machines in Q2, and '-6,055' for Tables in Q4.

Product Category 2	Q1	Q2	Q3	Q4
APPLIANCES	2,448	8,206	9,246	5,229
BOOKCASES	-1,100	-3,164	-2,244	-3,043
CHAIRS & CHAIRMATS	32,638	11,008	15,188	21,744
COMPUTER PERIPHERALS	42,190	34,312	30,778	23,168
COPIERS AND FAX	43,897	17,153	21,432	62,030
ENVELOPES	10,164	7,230	3,544	16,670
LABELS	8,888	12,370	8,030	6,661
OFFICE FURNISHINGS	-409	1,589	1,610	204
OFFICE MACHINES	28,097	14,899	24,467	44,648
PAPER	38,910	33,549	49,996	29,707
PENS & ART SUPPLIES	329	503	2,616	795
RUBBER BANDS	431	218	461	48
TABLES	-2,175	-4,295	-16,390	-6,055

## Marks


Tableau does not use chart types to build data views. Instead, data are displayed with marks, where every mark corresponds to a row (or a group of rows) in your data source.

You can build views of your data by placing fields on shelves and by selecting the appropriate mark type (or by accepting the default mark type). Marks are discussed in the following sections:

- Mark Types
- Stacking Marks
- Changing Mark Size and Color

### Mark Types

Mark types are available from the **Mark** menu. All mark types can be modified by color-encoding and by size-encoding (except polygon) the data



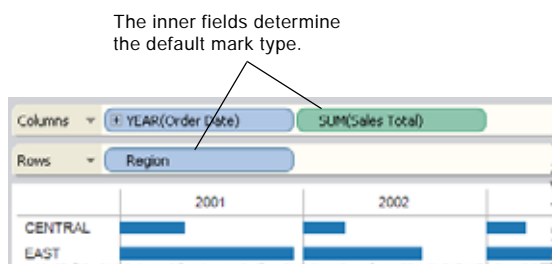
This section discusses the following supported mark types:

- Automatic Mark
- Text Mark
- Bar Mark
- Square Mark
- Circle Mark
- Shape Mark
- Line Mark
- Polygon Mark
- Gantt Bar Mark



## Automatic Mark

When the **Mark** menu is set to Automatic, Tableau automatically selects the best mark type for your data view. This mark type is determined by the inner fields on the **Rows** and **Columns** shelves.



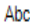
For example, if you create a view with a dimension as the inner field on both the **Rows** shelf and the **Columns** shelf, the text mark is automatically selected. If you create a view that has measures on both the **Rows** shelf and the **Columns** shelf, the shape mark is automatically selected. If you create a view with a dimension as the inner field on the **Rows** shelf and a measure on the **Columns** shelf (or vice versa), the bar mark is automatically selected. Note that Tableau automatically places measures inside dimensions when they share a shelf.

You can override the default selection and use any mark type that provides insight into your data. However, you should exercise some caution when manually selecting a mark type because the resulting view might hide important information about your data.

## Text Mark

The text mark type is useful when you want to display the numbers associated with one or more dimension members. This type of view is often called a text table, a cross-tab, or a Pivot Table. Tableau displays your data using text when:

- The **Mark** menu is set to Automatic, and you place one or more dimensions as the inner fields on both the **Rows** and the **Columns** shelves.
- You select **Text** from the **Mark** menu.

Initially, the data are displayed using the  icon.

	2001	2002	2003
APPLIANCES	Abc	Abc	Abc
BINDERS AND BINDER AG..	Abc	Abc	Abc
BOOKCASES	Abc	Abc	Abc
CHAIRS & CHAIRMATS	Abc	Abc	Abc
COMPUTER PERIPHERALS	Abc	Abc	Abc


To complete the view, you must place a field (typically a measure) on the **Text** shelf. As shown below, the **Sales Total** measure, which is aggregated as a summation, is used to complete the table.

	2001	2002	2003	2004
APPLIANCES	52,220	18,348	24,961	35,350
BINDERS AND BINDER..	48,623	45,457	29,296	34,831
BOOKCASES	25,831	72,820	20,015	26,529
CHAIRS & CHAIRMATS	255,343	116,387	113,460	125,410
COMPUTER PERIPHER..	156,535	85,714	101,517	73,869
COPIERS AND FAX	187,733	31,841	41,439	68,249
ENVELOPES	21,082	29,878	17,871	7,560
LABELS	19,261	20,729	17,269	17,146
OFFICE FURNISHINGS	53,146	16,941	25,272	20,728
OFFICE MACHINES	234,510	1,661	33,791	200
PAPER	190,735	167,783	171,055	175,713
PENS & ART SUPPLIES	20,410	29,794	14,874	17,971
RUBBER BANDS	4,286	3,672	2,767	3,874
SCISSORS, RULERS A..	7,167	8,341	2,443	6,282
STORAGE & ORGANIZ..	164,866	111,297	79,331	106,895
TABLES	245,185	81,213	70,992	61,626
TELEPHONES AND CO..	1,115,386	1,064,627	1,002,580	1,224,694

**Note** You can create a cross-tab of any data view by selecting the **Edit > Duplicate as Cross-tab** menu item. Refer to “Text Table” on page 43-22 for examples that show you how to build data views using text.

Because of the flexibility of Tableau, you might create a view that contains overlapping text. In this case, the following warning dialog box appears. If you do not want to display

this dialog box in the future, select the check box in the lower left. To display the dialog box again, select the **Help > Show Messages Again** menu item.



Overlapping text occurs when multiple data source values contribute to a single text table cell. There are three common cases to consider.

- **Level of detail** – If you place a dimension on the **Level of Detail**, **Color**, **Shape**, **Size**, or **Text** shelf, overlapping text occurs if multiple dimension members (levels of detail) contribute to a text table cell. To avoid overlapping text in this case, you might consider placing the dimension on the **Rows** or the **Columns** shelf.
- **Disaggregated data** – If you disaggregate a measure placed on the **Text** shelf, overlapping text occurs if multiple data source rows contribute to a text table cell. If you want to display disaggregated data, a text table is probably not the best choice. Instead, consider displaying the data in a scatter plot.
- **Non-unique names** – When connected to a multidimensional data source, a dimension can include non unique names. For example, an employee hierarchy for a company might contain two "Jane Smiths."

In the example below, overlapping text occurs when you disaggregate the **Sales Total** measure. As shown below, the cells contain overlapping sales data. This is because more than one data source row has a sale record for a given year and product. Note that Office Machines in 2004 indicates that there is only one sales record. However, this cell can still contain overlapping text if there are multiple data source rows with the same value. In this case, the overlapping text warning dialog box would still appear.

Columns: YEAR(Order Date)				
Rows: Product Category 2				
	2001	2002	2003	2004
APPLIANCES	8,000	5,000	8,000	8,000
BINDERS AND BINDER..	1,000	2,000	1,000	1,000
BOOKCASES	3,000	8,000	3,000	8,000
CHAIRS & CHAIRMATS	18,000	8,000	18,000	18,000
COMPUTER PERIPHER..	8,000	8,000	8,000	3,000
COPIERS AND FAX	28,000	12,000	18,000	18,000
ENVELOPES	2,000	8,000	3,000	3,000
LABELS	800	1,000	800	800
OFFICE FURNISHINGS	8,000	1,000	1,000	1,000
OFFICE MACHINES	18,000	800	18,200	200
PAPER	8,000	8,000	8,000	8,000
PENS & ART SUPPLIES	800	800	800	800
RUBBER BANDS	200	100	200	100
SCISSORS, RULERS A..	1,000	800	200	800
STORAGE & ORGANIZ..	18,000	18,000	18,000	18,000
TABLES	18,000	18,000	18,000	8,000
TELEPHONES AND CO..	18,000	18,000	18,000	18,000

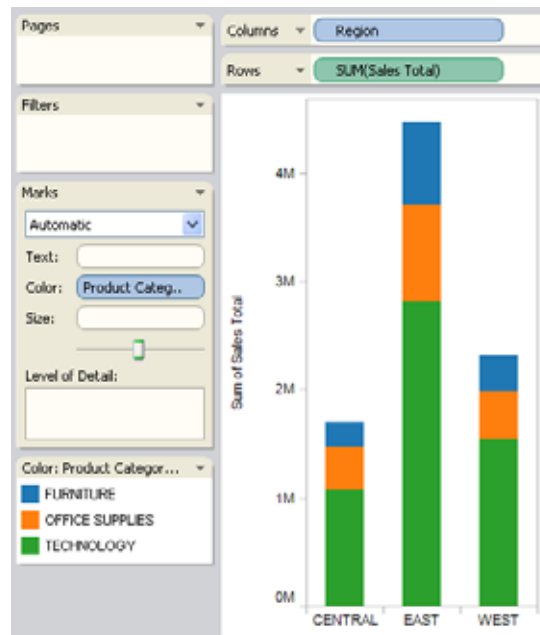
## Bar Mark

The bar mark type is useful when you want to compare measures across categories, or when you want to break data down into stacked bars. Tableau displays your data using bars when:

- The **Mark** menu is set to Automatic, and you place a dimension and a measure as the inner fields on the **Rows** and **Columns** shelves (or vice versa). If the dimension is a date dimension, the Line mark is chosen instead.
- You select **Bar** from the **Mark** menu.

Note that the marks are automatically stacked as described in “Stacking Marks” on page 9-26.

The data view shown below displays a dimension and a measure and is color-encoded by a dimension. Because the **Mark** menu is set to Automatic, the data are displayed using bars.

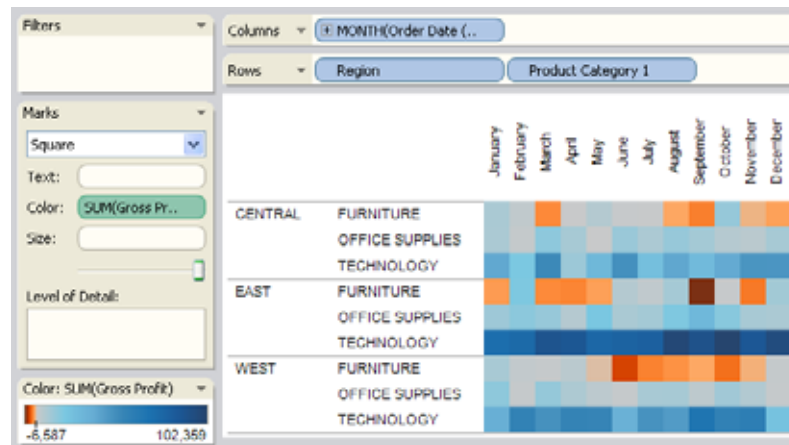


Refer to “Bar Chart” on page 43-4 for examples that show you how to build data views using bars.

### Square Mark

The square mark type is useful when you want to clearly see individual data points. When you select **Square** from the **Mark** menu, Tableau displays your data using squares.

The data view shown below displays several dimensions in both the rows and columns of a table. If the **Mark** menu was set to Automatic, the data would be displayed using text. By manually selecting Square, a completely different view is created. In particular, by placing a measure on the **Color** shelf, square marks can be used to create a heat map.

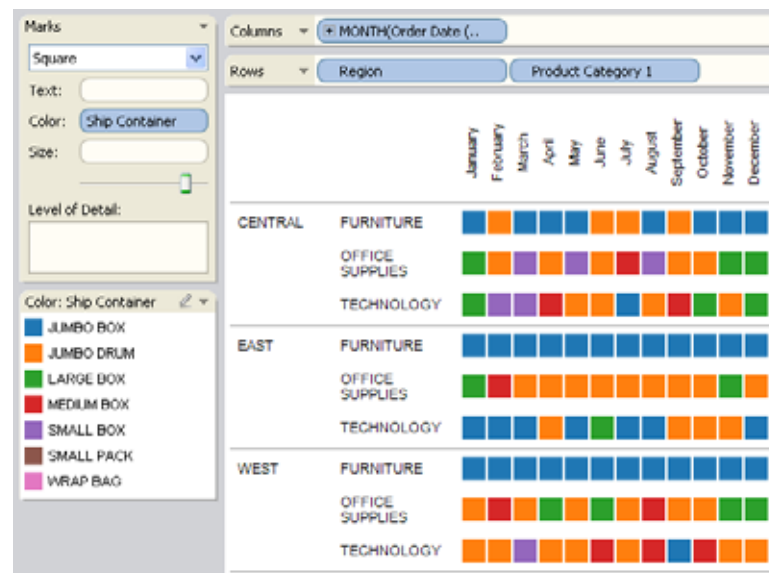


To reproduce this view, select the **Format > Cell Size > Square Cell** menu item and then adjust the size of the squares using the **Size** slider.


Refer to “Heat Map” on page 43-54 for detailed examples that show you how to build data views using squares.

Because of the flexibility of Tableau, you might create a view that contains overlapping data and is difficult to interpret. One way to do this is to place a dimension on the **Color** shelf. A view with overlapping data can be deceptive because only one of the marks for each cell is visible.

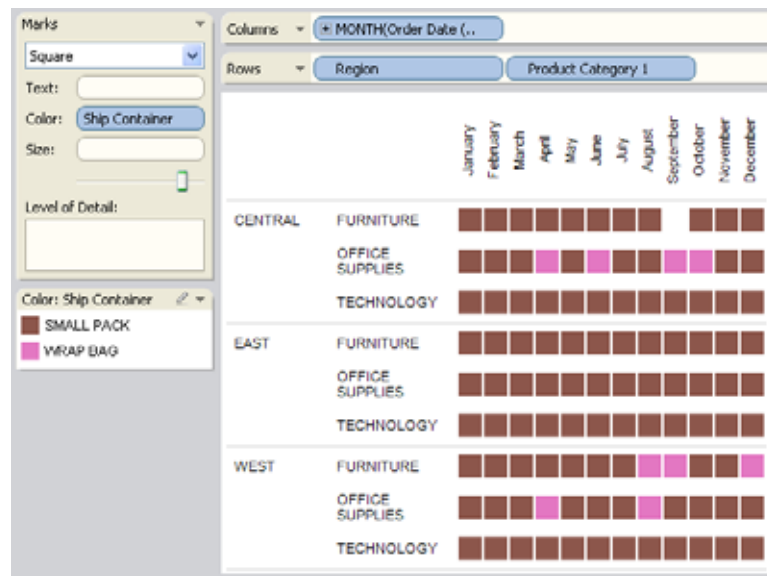
For example, suppose you replace the **Gross Profit** measure in the example above with the **Ship Container** dimension. As shown below, the squares indicate that there aren't any products shipped by Small Pack (brown) or Wrap Bag (pink).



Filter **Ship Container** to only include Small Pack and Wrap Bag.



As you can see, all of the squares have changed color showing that the marks overlapped.

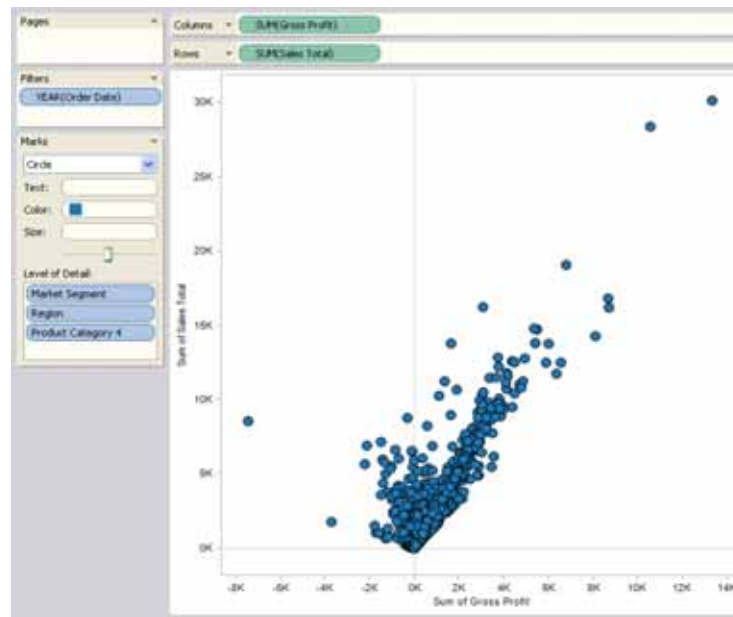


### Circle Mark

When you select **Circle** from the **Mark** menu, Tableau displays your data using circles.

As shown below, the data are displayed using circles. If the mark type was set to Automatic, Tableau would display the data using a shape (an open circle).






Refer to “Scatter Plot” on page 43-38 for examples that show you how to build a data view using circles.

### Shape Mark


The shape mark type is useful when you want to clearly see individual data points while also viewing categories associated with those points. Tableau displays your data using a shape when:

- The **Mark** menu is set to Automatic, and you place one or more measures on both the **Rows** and the **Columns** shelves.
- You select **Shape** from the **Mark** menu.

The view shown below displays the data from two measures. Because the **Mark** menu is set to Automatic, the data are displayed using a shape.

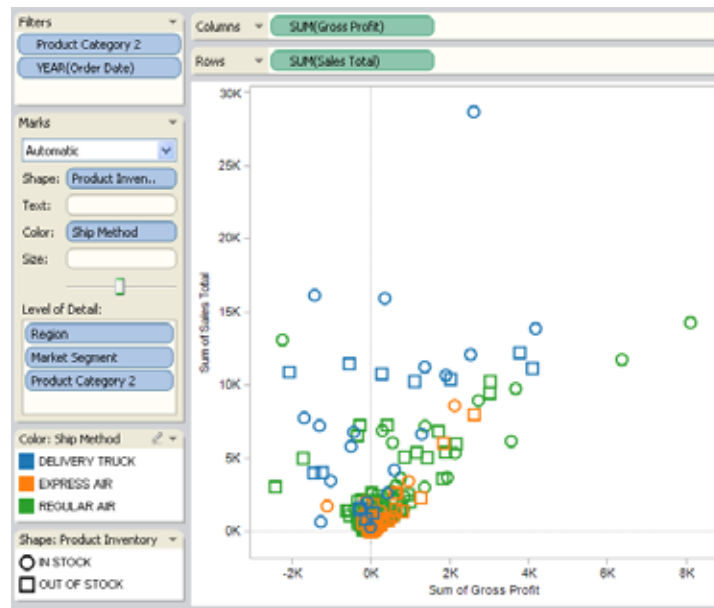


By default, the shape used is an open circle. You can select a different shape by double-clicking on the shape legend. As shown below, ten unique shapes are available.



To enhance the data view, you can place a dimension on the **Shape** shelf. Tableau separates the marks according to the members in the dimension, and assigns a unique shape to each member. The shape legend displays each member name and its associated shape.

As shown below, the **Product Inventory** dimension is used to shape-encode the view.



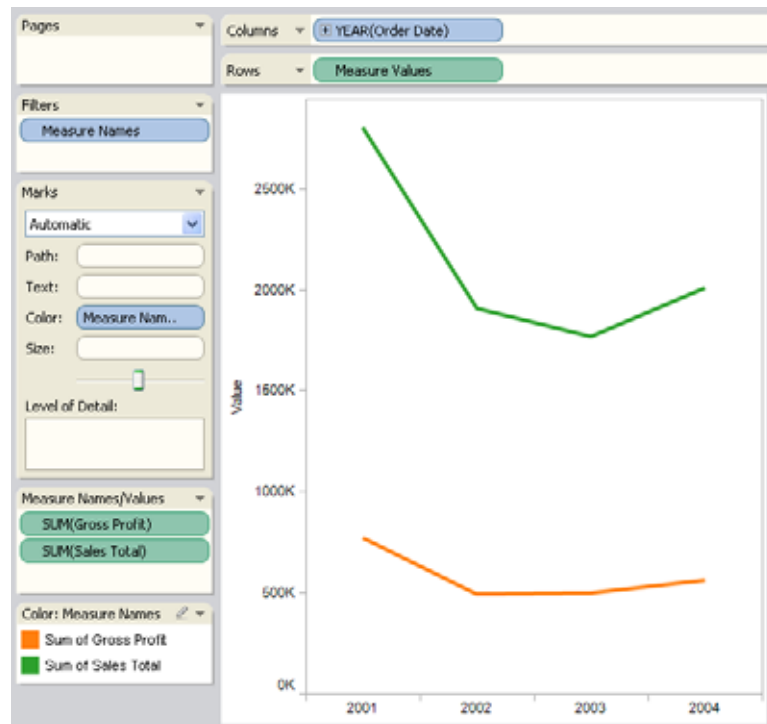
Refer to “Scatter Plot” on page 43-38 for examples that show you how to build data views using shapes.

### Line Mark

The line mark type is useful when you want to see trends in data over time, your data are ordered, or interpolation makes sense. Tableau displays data using lines when:

- The **Mark** menu is set to Automatic, and you place one or more measures on either the **Columns** shelf or the **Rows** shelf, and then plot the measures against a date dimension or a continuous dimension.
- You select **Line** from the **Mark** menu.

The data view shown below displays a dimension in the column of a table and several measures as the rows of the table.



With the line mark type, you can specify the drawing order of the line by placing a field on the **Path** shelf. Refer to “Line Chart-Filter and Path Encode” on page 43-36 for an example that uses path-encoded lines.

As the density of data increases, trends are often easier to see when using lines. This view shows 90 data points.



### Polygon Mark

Polygons are points connected by lines. The polygon mark type is useful when you want to connect points to create data areas. Tableau displays data using polygons when you select **Polygon** from the **Mark** menu.

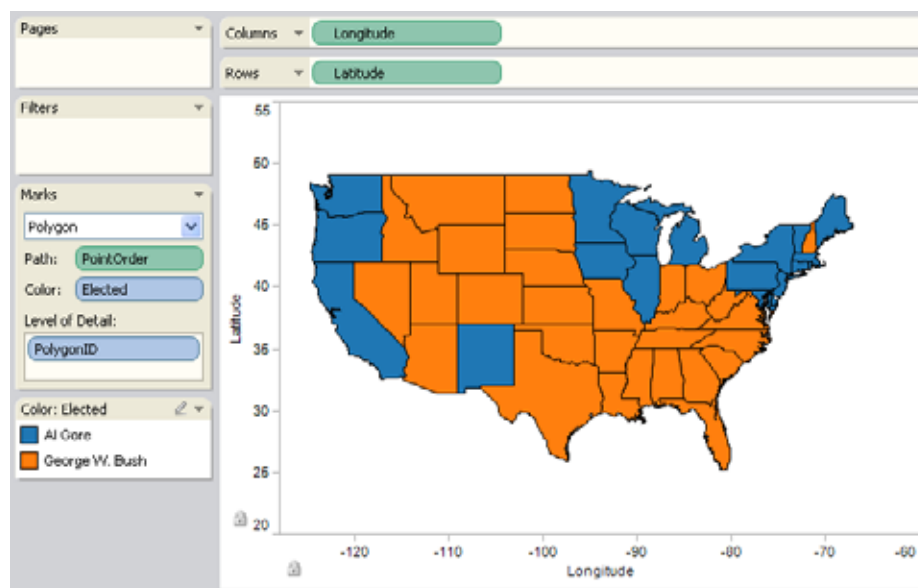
---

**Note** The polygon mark is not commonly used and often requires a specially constructed data source.

---

The view shown below comes from a specially constructed data source that holds geographic and election data. It displays the 48 contiguous US states as a function of latitude and longitude and color-encodes each state by the 2000 presidential election results.

If **Mark** is set to Automatic, the data will be displayed using a shape. By manually selecting Polygon, and adding additional fields to the view, a different view is created.



Every state is considered to be a polygon in the data source. The **PolygonID** field on the **Level of Detail** shelf is distinct for each US state. You can remove states from the view by filtering this field.

Additionally, you can specify the drawing order of the lines that constitute each polygon by placing a field on the **Path** shelf. In this example, the **PointOrder** measure is used to draw each state.

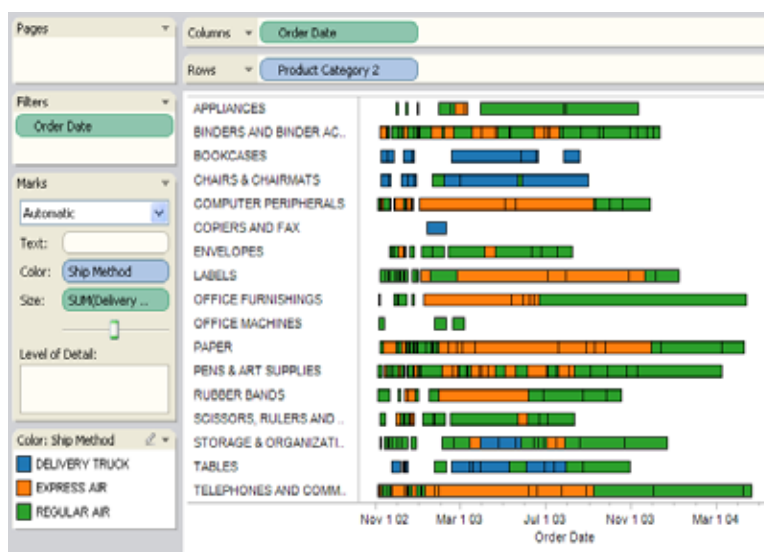
### Gantt Bar Mark

The Gantt bar mark type is useful when you want to view dates, project plans, or the relationships between different quantitative variables. Tableau displays your data using Gantt bars when:

- The **Mark** menu is set to Automatic and you place one or more dimensions on either the **Columns** shelf or the **Rows** shelf, and then plot the dimensions against a continuous quantity.
- You select Gantt Bar from the **Mark** menu.

The distinguishing characteristic of Gantt bars is that the length of every mark is proportional to the measure placed on the **Size** shelf.

The data view shown below displays a dimension as a function of a continuous date. If the **Mark** menu is set to Automatic, the data would be displayed using bars. By manually selecting Gantt Bar and adding additional fields to the view, a different view is created.




In particular, by placing the **Delivery Time** measure on the **Size** shelf, every bar in the view has been drawn with a length that indicates the delivery time of an order. Additionally, by placing the **Ship Method** dimension on the **Color** shelf, each bar is color-encoded by the shipping method.

## Stacking Marks

Stacking marks is relevant when your data view includes numeric axes. That is, at least one measure has been placed on the **Rows** or **Columns** shelf. When marks are stacked, they are drawn cumulatively along an axis. When marks are not stacked, they are drawn independently along an axis. That is, they are overlapping.

Stacking marks is particularly useful for bar charts which is why Tableau automatically stacks bars. You might find that stacking marks is useful for other marks such as lines as well. You can control whether marks are stacked or overlapping in any given view by selecting the **Analysis > Stack Marks** menu item. You can either allow Tableau to automatically select whether the marks are stacked or you can specify on or off. The default

mode is automatic. When you are in automatic mode, the Stack Marks menu indicates whether stacked marks is on or off.



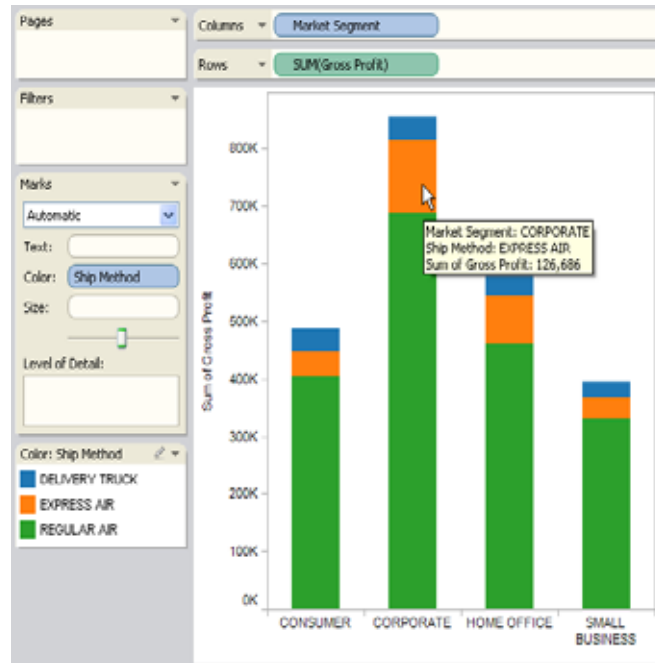
If you select **On** or **Off** on the **Stack Marks** menu, you are switched to manual mode. Your selection remains throughout any changes you make to the view.

The following examples illustrate stacking marks.

### Example – Stacking Bars

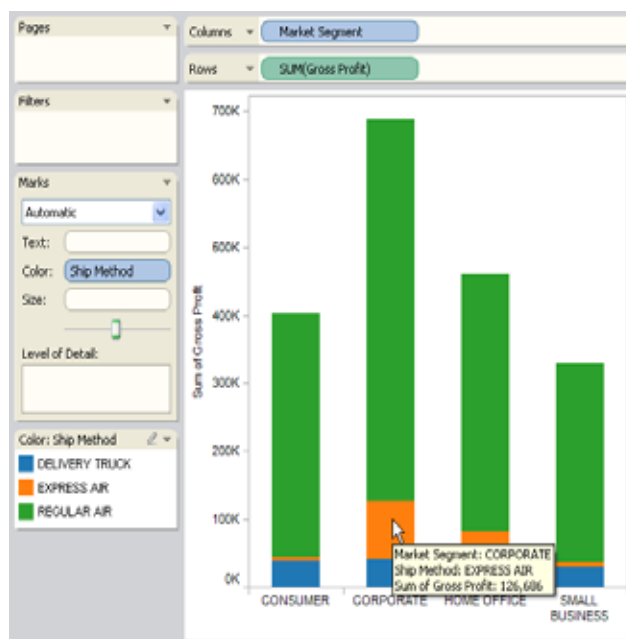
Consider the stacked bars view shown below. It was created by placing a dimension on Columns shelf, placing a measure on the Rows shelf, and color-encoding the data by a dimension.





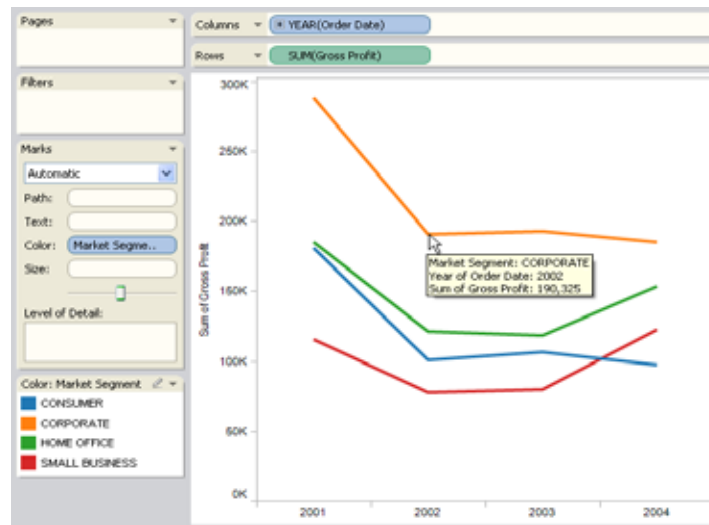
Because the mark type is a bar, Tableau automatically stacks the marks. This means that the marks are drawn cumulatively and the height of each stacked segment within each bar represents the value for that segment. For example, the sum of the gross profit for products shipped by Express Air (orange bar segment) in the Corporate market is \$126,686.

If you un-stack the marks, they all start from the horizontal axis. As shown below, you can still view the individual bar segments. Be aware, however, because un-stacked marks overlap, it is possible to create a view where bar segments are not visible.



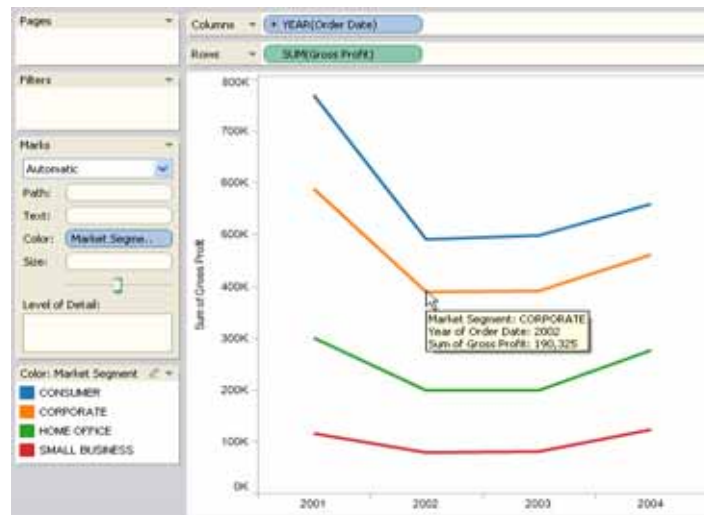
### Example – Stacking Lines

Consider the data view shown below. It was created by placing a date dimension on the **Columns** shelf, placing a measure on the **Rows** shelf, and color-encoding the data by a dimension. Because the mark type is a line, the marks are not automatically stacked. Instead, they are drawn independently from the horizontal axis.



Interpret any data point by reading the associated values from the horizontal and vertical axes. For example, in the year 2002, the Corporate (orange) sales totaled \$190,325. That is, the space between that data point and the horizontal axis is equal to the sum of the sales total for the Corporate market.

Now, stack the marks by selecting the **Analysis > Stack Marks > On** menu item. The stacked lines view is shown below.



In this view, the lines are no longer independent of each other. Instead, they are drawn cumulatively. The stacking order is given by the order of the dimension members in the data source. This order is reflected in the color legend, from bottom to top.

Therefore, the stacked Small Business (red) line is the same as its un-stacked version because it's at the bottom of the stacking list. The stacked Home Office (green) line is derived by adding its un-stacked values to the un-stacked Small Business values. The stacked Corporate (orange) line is derived by adding its un-stacked values to the stacked Home Office data. The stacked Consumer (blue) line is derived by adding its un-stacked values to the stacked Corporate data.

The vertical axis gives the new scale for the stacked marks. Interpret the space between consecutive lines as the sum of the sales total. The lines are no longer all compared to the horizontal axis.

For example, notice that the tooltip for the 2002 Corporate data still shows the sales total as \$190,325. The interpretation is that the space between the Corporate data and the Home Office data yields the sum of the sales total for the Corporate market.

## Changing Mark Size and Color

You can format marks by changing the mark size and color. This allows you to highlight specific data, to distinguish between marks effectively, and to create optimal presentations.

You can also display or remove mark borders. Refer to “Mark Borders” on page 20-18 for more information. This section discusses the following topics:

- Changing Mark Size
- Changing Mark Color

### Changing Mark Size

Each mark is displayed with a default mark size. You can change the size of marks at any time by moving the **Size** slider.



If you move the slider to the right, marks get larger. If you move the slider to the left, marks get smaller. The **Size** slider affects different marks in different ways, as described in the following table.

Mark Type	Description
Circle, Square, Shape, Text	Moving the slider makes the mark bigger or smaller.
Bar, Gantt Bar	Moving the slider makes bars wider or narrower.
Line	Moving the slider makes lines thicker or thinner.
Polygon	You cannot change the size of a polygon.

The size of your data view is not modified when you change marks using the **Size** slider. However, if you change the view size, the mark size might change to accommodate the new formatting. For example, if you make the table bigger, the marks might become bigger as well. Refer to “Formatting” on page 20-1 for more information.


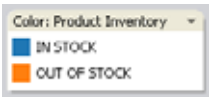

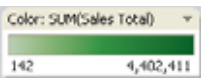
**Note** Changing the mark size is not the same as size-encoding the data using the **Size** shelf. Refer to “Size Shelf” on page 10-23 for more information.

## Changing Mark Color

Each mark is displayed with a color, which is presented in a color legend on the Tableau interface.

By default, all marks use the same color. However, you can display more than one color by placing a dimension or a measure on the **Color** shelf (**Ctrl+Alt+O**). Placing a dimension on the **Color** shelf separates the marks according to the dimension members and assigns a unique color to each member. Placing a measure on the **Color** shelf creates a continuous range of colors.

Depending on your data view, Tableau will use one of the four color legends described in the following table.

Legend Type	Description
	This is the default color. It is used when the <b>Color</b> shelf is not populated with a field. To edit the default color, double-click the blue box.
	This legend appears when the <b>Color</b> shelf is populated with a dimension. To edit a color, double-click anywhere in the legend.
	This is a diverging color legend and appears when the <b>Color</b> shelf is populated with a measure that contains both positive and negative numbers. To edit the colors, click any part of the color spectrum.
	This legend appears when the <b>Color</b> shelf is populated with a measure that contains only positive or only negative numbers. To edit the colors, click any part of the color spectrum.

Refer to “Color Shelf” on page 10-21 for more information about color encoding the view.


## Titles

You can add a title to any worksheet or dashboard. The title is displayed on the Title card.

### To show and hide titles:

- Select **View > Title** or click **View Cards**  on the toolbar and then select the **Title** card.

### Worksheet Title



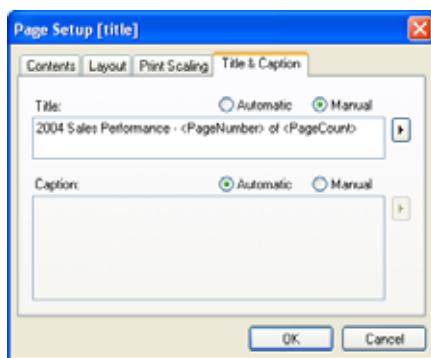
### Dashboard Title



By default, the title is the name of the sheet, but you can use a custom title and even include automatic text such as page number and sheet name.


### To edit titles:

- 1 Right-click on the title and select **Edit Title**.
- 2 In the Page Setup dialog box, type a new title into the **Title** text box. Use the arrow to the right of the text box to add automatic text such as page number, sheet name, page count, and more.



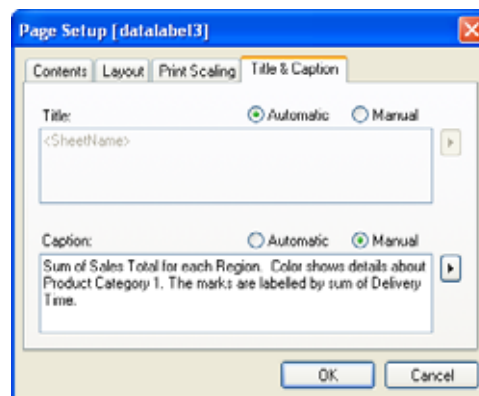
You can format the font, alignment, shading, and border of titles. Refer to “Title and Caption” on page 20-22 to learn more.

## Captions

All views can have a caption that is either automatically generated or manually created. The caption is displayed on the Caption card. To show the caption, select it on the **View Cards** toolbar menu  or select **View > Caption**.



The caption is automatically generated by default, however, you can edit the caption by double clicking the Caption card and selecting **Manual** in the subsequent dialog box.



The caption is part of the Page Setup settings and can optionally be printed and published with the view. Additionally, when you export the view as an image to another application like Microsoft PowerPoint, you can select to include the caption.

You can format the font, alignment, shading, and border of captions. Refer to “Title and Caption” on page 20-22 to learn more.



## Field Labels

Placing discrete fields on the rows and column shelves creates headers in the view that display the members of the field. For example, if you place a field containing products on the rows shelf, each product name is shown as row headers. In addition to showing these headers, you can show field labels, which are labels for the headers. In this example, the rows are labeled as Product Category, thus indicating that the list of products are members of the Product Category field.

Field labels indicate the data fields that are shown in the table.

Region / Year of Order Date								
Product Category	EAST				WEST			
	2001	2002	2003	2004	2001	2002	2003	2004
APPLIANCES	8,101	1,611	5,120	1,758	449	-1,836	49	3,871
BOOKCASES	139	-1,798	-822	-1,129	-136	-3,136	-136	-177
CHAIRS & CHAIRMATS	24,300	-296	13,780	10,146	6,846	3,998	-2,332	-737
COMPUTER PERIPHERALS	29,657	11,583	21,630	13,273	7,427	4,608	4,526	8,200
COPIERS AND FAX	25,176	16,480	4,439	26,276	48,827		16,368	
ENVELOPES	5,661	1,766	5,109	1,476	844	3,019	844	20

Field labels apply only to discrete fields. When you add continuous fields to the view, an axis is created. The axis is labeled with a header. Refer to “Axes” on page 9-6 to learn more about axes.

By default, field labels are shown. You can hide or show field labels at anytime.

**To show and hide field labels:**

- Select **Table > Field Labels for Columns** or **Field Labels for Rows**.



	2001	2002	2003	2004	Grand Total
APPLIANCES	62,220	18,348	24,961	36,360	130,879
BOOKCASES	25,831	72,620	20,516	26,629	145,196
CHAIRS & CHARMATS	266,343	116,367	113,480	126,410	615,601
COMPUTER PERIPHERALS	166,636	66,714	101,617	73,968	417,936
COPERS AND FAX	167,733	31,841	41,436		
ENVELOPES	21,062	29,878	17,871		
LABELS	19,261	20,729	17,206		
OFFICE FURNISHINGS	63,146	16,941	26,273		
OFFICE MACHINES	234,610	1,661	33,791	220	270,162
PAPER	160,736	167,762	171,066	176,713	708,288
PENS & ART SUPPLIES	20,410	26,794	14,674	17,971	69,849
RUBBER BANDS	4,386	3,672	2,767	3,974	14,800
TABLES	246,186	61,213	72,662	61,626	489,519
<b>Grand Total</b>	<b>1,466,276</b>	<b>676,791</b>	<b>666,294</b>	<b>634,227</b>	<b>3,432,670</b>

You can format the fonts, alignment, shading, and separators for field labels. Refer to “Field Labels” on page 20-21 to learn more.

## Legends

When you add fields to any of the encoding shelves such as the Color, Shape, and Size shelves, a legend appears to indicate how the view is encoded with relation to your data.



Not only do legends help you understand encodings, you can also use legends to sort, filter, and highlight specific sets of data.

Refer to “Sort by Drag and Drop” on page 12-17 to learn how to sort using legends.

Refer to “Filtering Discrete Fields” on page 12-20 to learn how to filter using legends.

Refer to “Color Legend Highlighting” on page 14-9 to learn how to highlight select items in the view.

# Building Views Manually

---

<b>Placing Fields on Shelves . . . . .</b>	<b>10-3</b>
<b>Placing Fields on the Worksheet . . . . .</b>	<b>10-6</b>
Dragging Discrete Fields to the View . . . . .	10-6
Dragging Continuous Fields to the View . . . . .	10-7
<b>Types of Shelves . . . . .</b>	<b>10-10</b>
Columns and Rows Shelves . . . . .	10-10
Pages Shelf . . . . .	10-12
Filters Shelf . . . . .	10-16
Level of Detail Shelf . . . . .	10-19
Color Shelf . . . . .	10-21
Size Shelf . . . . .	10-23
Shape Shelf . . . . .	10-28
Text Shelf . . . . .	10-29
Path Shelf . . . . .	10-31
<b>Multidimensional Hierarchies . . . . .</b>	<b>10-33</b>
Drilling Down and Up in a Hierarchy . . . . .	10-33
Building Views with Hyperion Essbase . . . . .	10-35
Building Views with Microsoft Analysis Services Cubes . . . . .	10-39
Perfect Pivoting . . . . .	10-40
Defining Unique Values . . . . .	10-41
<b>Example – Building Data Views Manually . . . . .</b>	<b>10-44</b>

---

## Overview

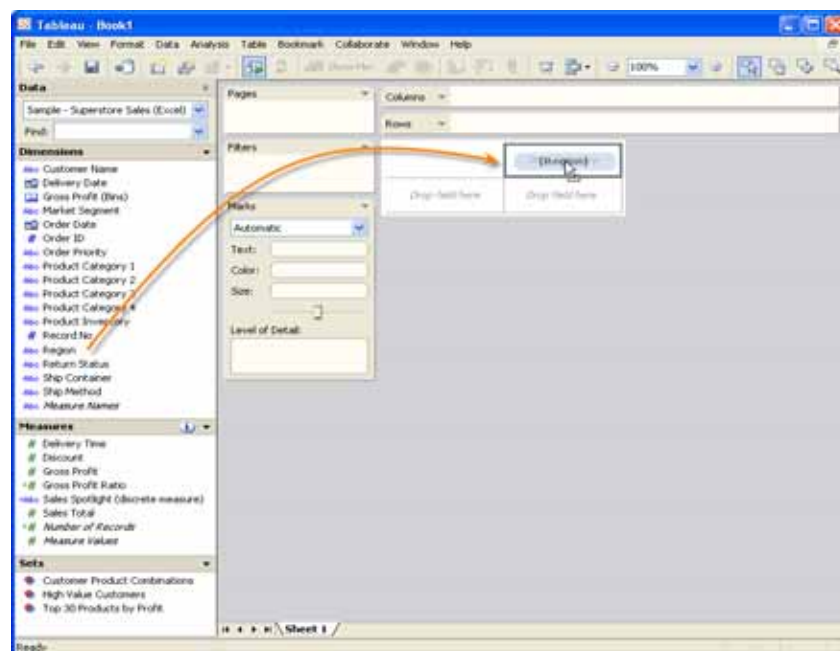
Building views in Tableau can be really easy if you understand some basic concepts of how it all work. This section discusses the following topics:

- Placing Fields on Shelves
- Placing Fields on the Worksheet
- Types of Shelves
- Multidimensional Hierarchies
- Example – Building Data Views Manually

## Placing Fields on Shelves

You can build views of your data by placing fields onto the shelves that are a part of every worksheet, or onto one of the components of a table (pane, header, axis, or legend). To manually place a field on a worksheet, click and drag the field from the Data window to the desired location in the view.

When you begin creating a new data view on a blank worksheet, you can place fields onto either the shelves or the row or column locations shown below.



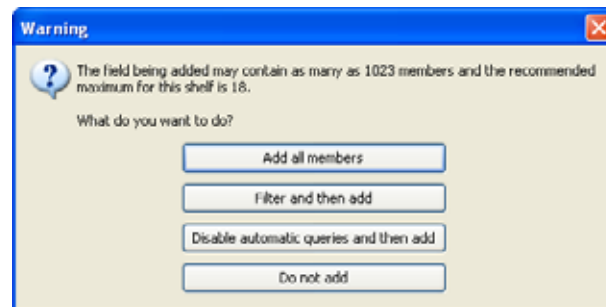
The color of the field indicates what behavior it will have when dropped. Blue is used to represent discrete fields, which will create headers when dropped. Green is used to represent continuous fields, which will create axes when dropped. Generally, dimensions are colored blue while measures are colored green. For a more advanced discussion of what the blue and green colors mean, refer to “Data Roles” on page 8-22.

You can drag fields from the Data window onto any shelf. You can also drag fields from one shelf to another shelf, and from one worksheet to another worksheet provided the


worksheets are connected to the same data source. Refer to “Understanding the Data Window” on page 8-3 for more information about the Data window.

The number of fields that you can place on the **Columns**, **Rows**, **Level of Detail**, **Filters**, and **Pages** shelves is unlimited. However, the **Color**, **Size**, **Shape**, **Text**, and **Path** shelves can hold only one field at a time.

If you place a dimension with a large number of members on a shelf, Tableau will present you with the following dialog box to help you avoid a performance degradation.



If you are building a data view that involves a large amount of data, it is generally more efficient to follow this procedure:

- 1 Turn off automatic updated by clicking the Pause Automatic Updates button  on the toolbar.
- 2 Place all desired fields on shelves.
- 3 Specify filters to restrict the data to the members of interest (refer to “Filtering” on page 12-19).
- 4 Turn on automatic updates by clicking the Resume Automatic Updates button on the toolbar.

At any time a query is taking too long, you can cancel the query by clicking **Cancel** in the progress dialog box. For more information about canceling queries refer to “Cancel Query” on page 7-4.

Replace a field with a second field on any shelf by dragging the second field directly on top of the first field. Refer to “Types of Shelves” on page 10-10 for more information about each of these shelves.

To remove a field from a shelf, drag it off the worksheet or select **Remove** on the field's context menu. To quickly remove multiple fields from a shelf, right-click the shelf and select **Clear Shelf** on the context menu.



## Placing Fields on the Worksheet

In addition to dragging fields to shelves, you can drag fields inside the table. This is an easy way to create a cross-tab or to work with multiple measures. The options available to you depend on the table type (text table, bar chart, scatter plot, etc.), and the table components (panes, headers, or axes). This section discusses the following topics:

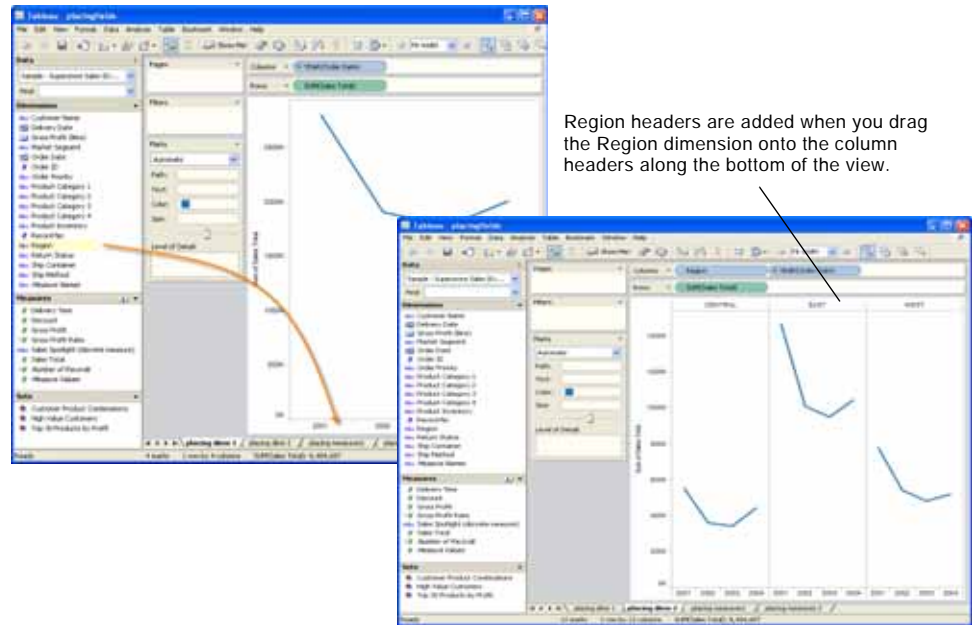
- Dragging Discrete Fields to the View
- Dragging Continuous Fields to the View

### Dragging Discrete Fields to the View

You can drag discrete fields to both headers and axes in the view. In addition, you can drag these fields to the pane itself. The result depends on where you drag the fields to. The following table describes what happens when you place a discrete field on each part of the view.

Location in View	Result
Row and Column headers	Tableau automatically creates row or column headers with the new discrete field members.
Quantitative axes	Tableau automatically creates row or column headers with the new discrete field members.
Pane	Tableau automatically color encodes the view using the data in the new discrete field. If the view is already color encoded, the new field is added to the view in an intelligent way.

For example, suppose you have a simple line graph with a single pane. When you drag a dimension to the column headers, additional column headers are created using the new dimension members. The same thing happens when you drag a dimension onto a quantitative axis.

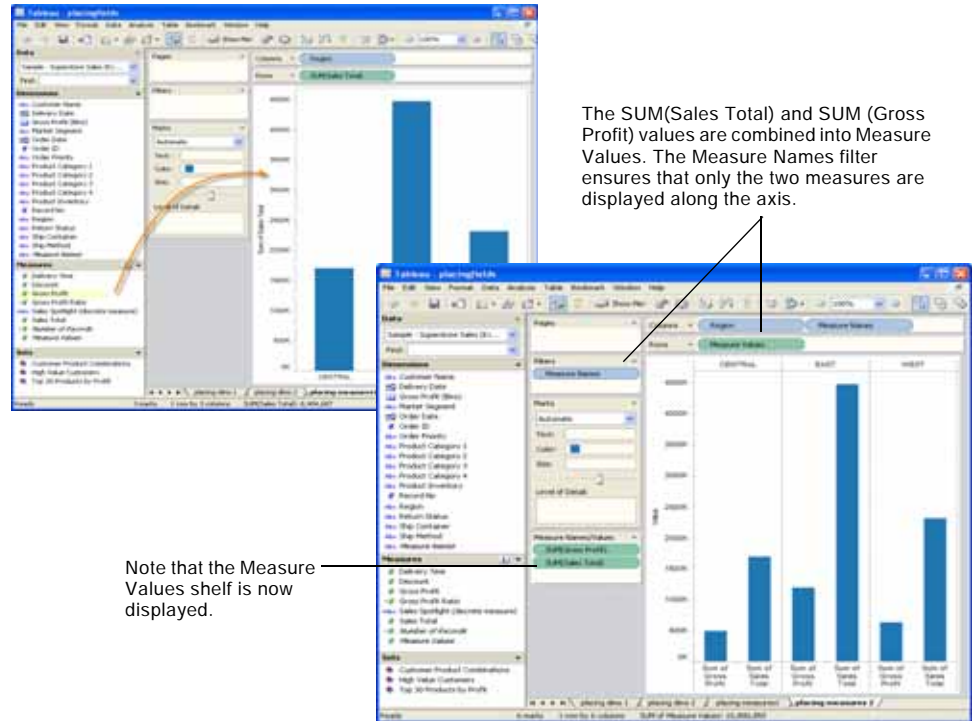


## Dragging Continuous Fields to the View

Similar to dragging discrete fields to the view, you can drag continuous fields onto the row and column headers, quantitative axes, and the pane itself. The table below explains the result of dragging a continuous field to various locations in the view.

Location in View	Result
Row and Column headers (nominative axes)	Tableau automatically creates a quantitative axis on either the Rows or Columns shelf.
Quantitative Axes	Tableau combines the new field with the existing data on the axis. The <b>Measure Values</b> field is placed onto the axis and filtered by <b>Measure Names</b> to only display the existing measures plus the new continuous field. <b>Measure Names</b> is added to either the rows or columns shelf. You cannot place continuous dimensions onto a quantitative axis. Refer to “Types of Shelves” on page 10-10.
Pane	Tableau automatically color encodes the view using the data in the new continuous field. If the view is already color encoded, the new field is placed on the <b>Level of Detail</b> shelf.

Suppose you have a simple bar chart with a quantitative axis. When you drag a measure onto the axis, the existing measure is combined with the new measure using the **Measure Values** field and a filter by **Measure Names** is placed on the view.



## Types of Shelves

Every worksheet in Tableau contains *shelves*. By placing fields on shelves, you can create the rows and columns of a data view, exclude data from the view, show additional levels of detail, and encode the data in various ways.

Shelves are described in the following sections:

- Columns and Rows Shelves
- Pages Shelf
- Filters Shelf
- Level of Detail Shelf
- Color Shelf
- Size Shelf
- Shape Shelf
- Text Shelf
- Path Shelf

Each section contains examples that illustrate how a simple data view is modified by placing a dimension or a measure on the shelf. For an example that utilizes many different shelves, refer to “Example – Building Data Views Manually” on page 10-44.

Note that some shelves are available only when certain mark types are used. For example, the **Shape** shelf appears when the shape mark type is used. This mark is selected automatically when you place measures on the both the **Rows** and **Columns** shelves and the mark type is Automatic. If you select a different mark type, the **Shape** shelf is hidden. Additionally, some shelves are not particularly useful with certain mark types. For example, size-encoding your data when using the line mark type is generally not useful. Refer to “Mark Types” on page 9-11 for more information about marks.

You should experiment with various combinations of shelves, fields, and mark types to find the optimal view for your data.

### Columns and Rows Shelves

The **Columns** shelf creates the columns of a table, while the **Rows** shelf creates the rows of a table. You can place an unlimited number of fields on these shelves.

When you place a dimension on the **Rows** or **Columns** shelf, headers for the members of that dimension are created. When you place a measure on the **Rows** or **Columns** shelf,

quantitative axes for that measure are created. As you build up your data view with more fields, additional headers and axes are included in the table and you get an increasingly detailed picture of your data.

In the view shown below, the members of the **Market Segment** dimension are displayed as column headers, while the **Gross Profit** measure is displayed as a vertical quantitative axis.

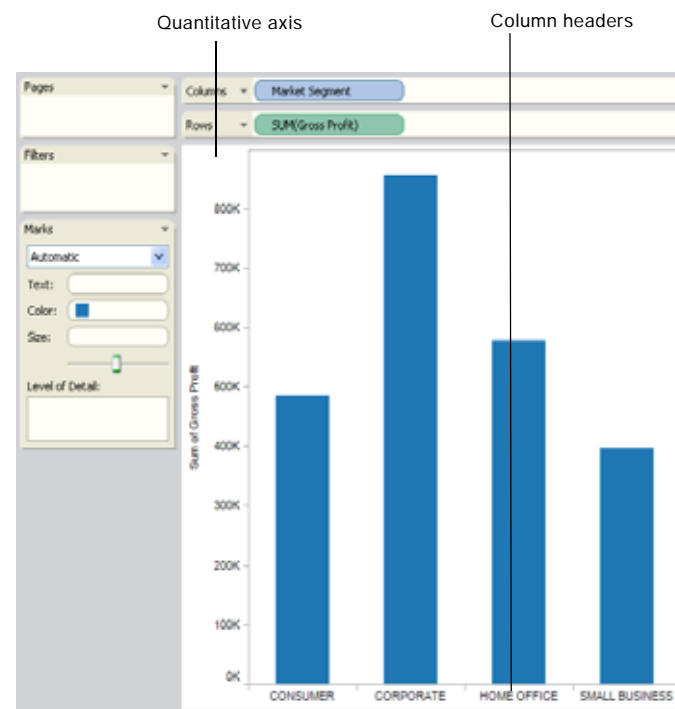
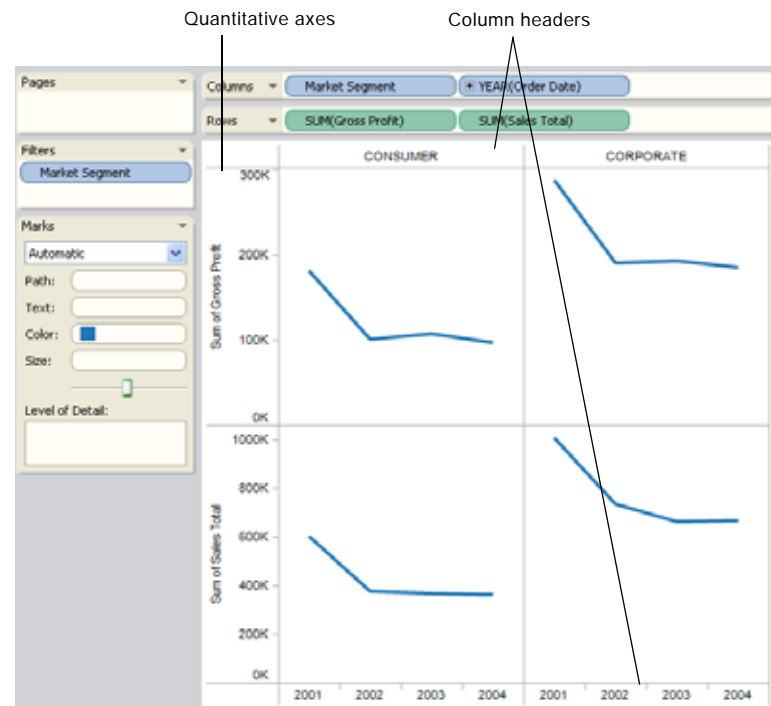


Tableau displays data using marks, where every mark corresponds to a row (or a group of rows) in your data source. The inner fields on the **Rows** and **Columns** shelves determine the default mark type. For example, if the inner fields are a measure and a dimension, the default mark type is a bar. You can manually select a different mark type using the **Mark** menu. Refer to “Mark Types” on page 9-11 for more information.

Adding more fields to the **Rows** and **Columns** shelves adds more rows, columns, and panes to the table.



## Pages Shelf

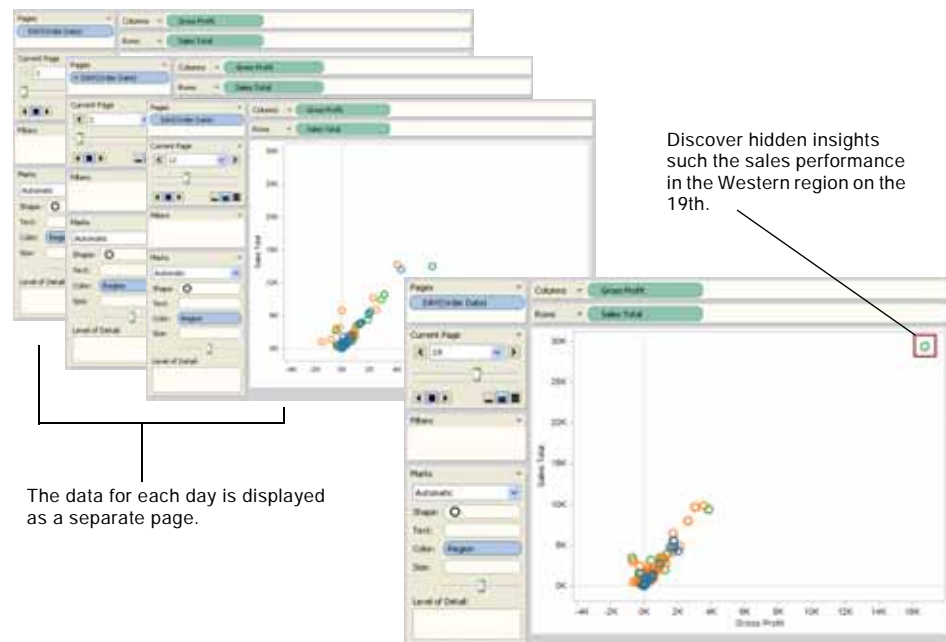
The **Pages** shelf lets you break a view into a series of pages so you can better analyze how a specific field affects the rest of the data in a view. When you place a dimension on the Page shelf you are basically adding a new row for each member in the dimension. When you place a measure on the Pages shelf, the measure is converted into a discrete measure.

The page shelf creates a view on a different page for each new row so you can easily flip through each view and compare them on a common axis. For example, the view below shows the Profit vs. Sales by Region for each day throughout the month.



You can see that it is difficult to see how these two measures have interacted from day to day. However, when you move the **Day** field to the **Pages** shelf and flip through the pages (one for each day) you can quickly discover hidden insights. In this example, it is interesting that the 19th is an especially big day in terms of sales and profit in the Western region.






The data for each day is displayed as a separate page.

When you add a field to the page shelf the **Current Page** card displays. Use this card to navigate through the pages.



There are three ways to navigate through the pages in a view.

**Jump to a specific page.** Select the member or value you want to view from the drop-down list on the Current Page card to display a specific page rather than scrolling through the entire sequence.




**Manually Advance through the pages.** You can manually advance through the sequence of pages by doing one of the following:

- Use the forward and back buttons on either side of the drop-down list to navigate through the pages one at a time.
- Use the Page Slider to quickly scroll forward and backward in the sequence of pages.
- Use the keyboard shortcuts below to scroll forward and backward in the sequence of pages.

F4	Starts and stops forward playback
SHIFT + F4	Starts and stops backward playback
CTRL + .	Skip forward one page
CTRL + ,	Skip backward one page

**Automatically Advance through the pages.** Use the playback controls to watch a slide show of the pages in the view. You can play forward, play backward, and stop.



---

**Note** You can control the speed of playback with the speed controls in the bottom right corner of the card. The smallest bar indicates the slowest playback speed.

---

## Filters Shelf

The **Filters** shelf allows you to specify which data to include and exclude. For example, you might want to analyze the gross profit for each market segment, but only for certain shipping containers and delivery times. By placing fields on the **Filters** shelf, you can create such a view.

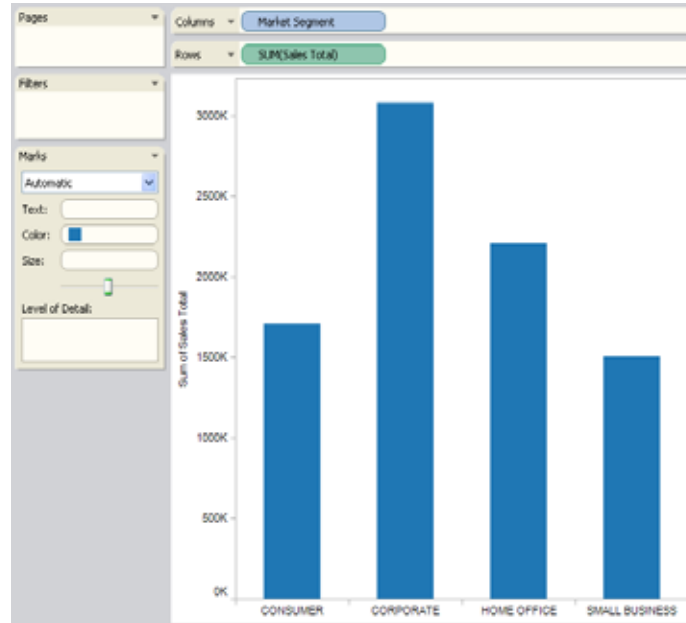
---

**Note** This section presents a brief overview of filtering. Refer to “Filtering” on page 12-19 for a complete description.


---

You can filter data using measures, dimensions, or both at the same time. Additionally, you can filter data based on the fields that make up the columns and rows of the table. This is called an *internal filter*. You can also filter data using fields that don’t contribute headers or axes to the table. This is called an *external filter*. All filtered fields display on the **Filters** shelf.

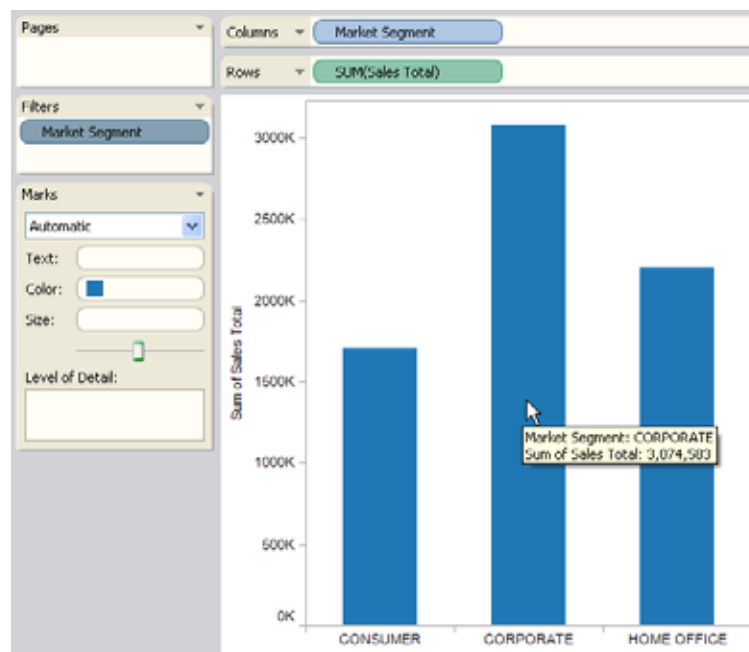
To illustrate the basic concepts of filtering, consider the following view.



Suppose you are not interested in the Small Business data. You can remove this column from the view by filtering the **Market Segment** dimension. To do so, select **Filter** from the field's context menu. The **Filter** dialog box opens. By default all members are selected. Un-check **Small Business** to exclude it from the view. All selected members will be included.




As shown below, **Market Segment** is automatically placed on the **Filters** shelf, and the view now contains three columns instead of the previous four.

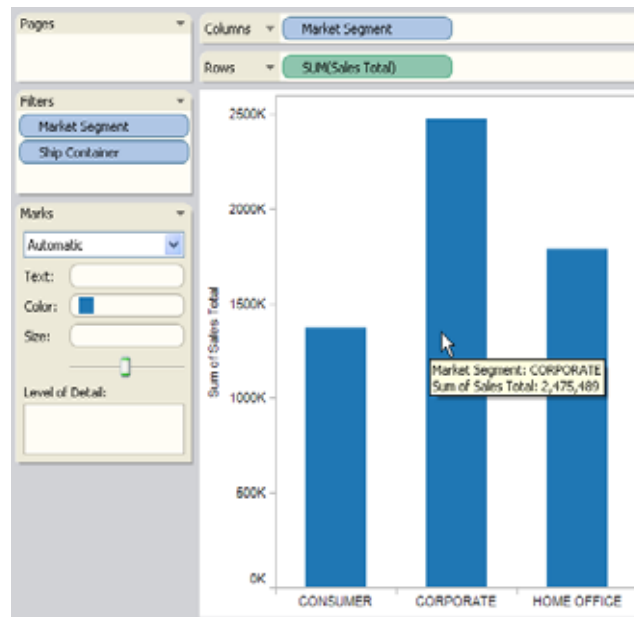


Suppose you want to only view sales for products that were shipped in boxes. To do this, place the **Ship Container** dimension directly on the **Filters** shelf. This is an example of an external filter because **Ship Container** is not part of the view. That is, it does not contribute row or column headers.

The **Filter** dialog box shown below automatically opens. By default, none of the members are selected. Select the members you want to keep as part of the view. All deselected members are excluded.



The modified data view is shown below. The tooltip shows that the sum of the sales for the Corporate market has decreased to \$2,475,489. This number is derived by summing all the rows in the data source that are associated with the Corporate market and that use a box as a shipping container.



The order of fields placed on the **Filters** shelf does not affect the data view because the filters are independent. The result of filtering by market segment, and then by shipping container is the same as filtering by shipping container and then by market segment.

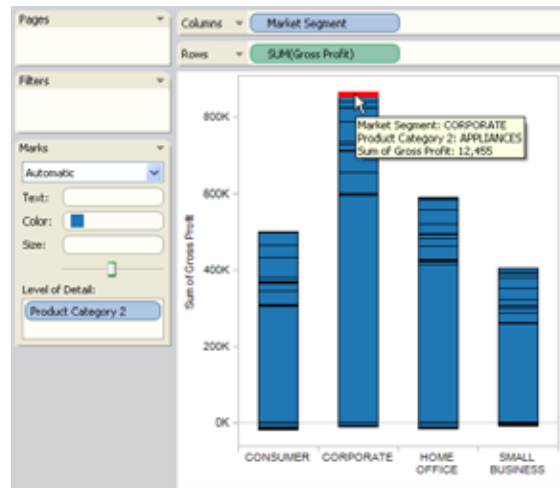
## Level of Detail Shelf

Whenever you place a dimension on the **Rows** or **Columns** shelf, the categorical members create table headers. The headers represent levels of detail because they separate the data source rows into specific categories. You can identify each category by the member name. For example, the **Market Segment** dimension separates the data source rows into four levels of detail: Consumer, Corporate, Home Office, and Small Business.

The **Level of Detail** shelf also allows you to separate the marks in a data view according to the members (levels of detail) of a dimension. However, unlike using the **Rows** and **Columns** shelf, using this shelf is a way to show more data without changing the table structure.

As shown below, the bars are separated into segments according to the members of the **Product Category 2** dimension. The size of each segment reflects the contribution to the

gross profit for a particular member. For example, view below shows that Appliances category in the Corporate market has a gross profit of \$12,455.



You can place any number of dimensions on the **Level of Detail** shelf. In fact, placing all dimensions on this shelf is one way to display all the rows of your data source.

---

**Note** The **Level of Detail** shelf works only if the measures that contribute axes to the table are aggregated. If the measures are disaggregated, then it isn't possible to separate the marks into additional levels of detail because all levels of detail are already shown.

---

Also, placing a measure on the **Level of Detail** shelf has no effect because measures do not contain members. However, you can place measures on this shelf if you want to export their values to Microsoft Access, copy their values to the Windows Clipboard, or view them in the tooltips.

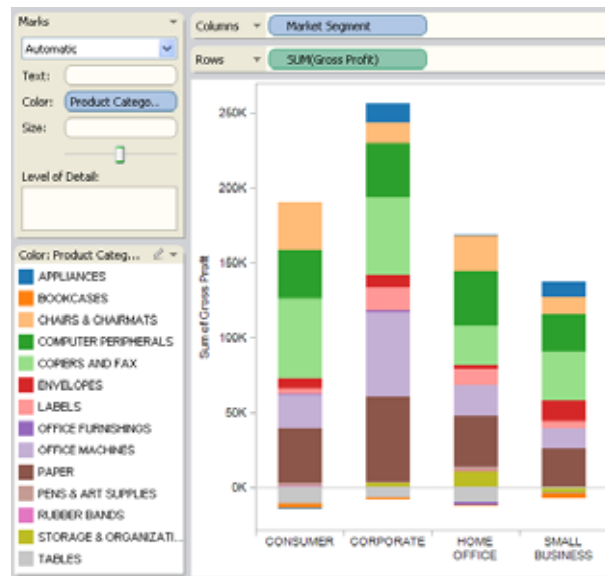
## Color Shelf

The **Color** shelf allows you to encode data by assigning different colors to the marks in a data view. The effect of color-encoding your data view depends on whether you use a discrete or continuous field.

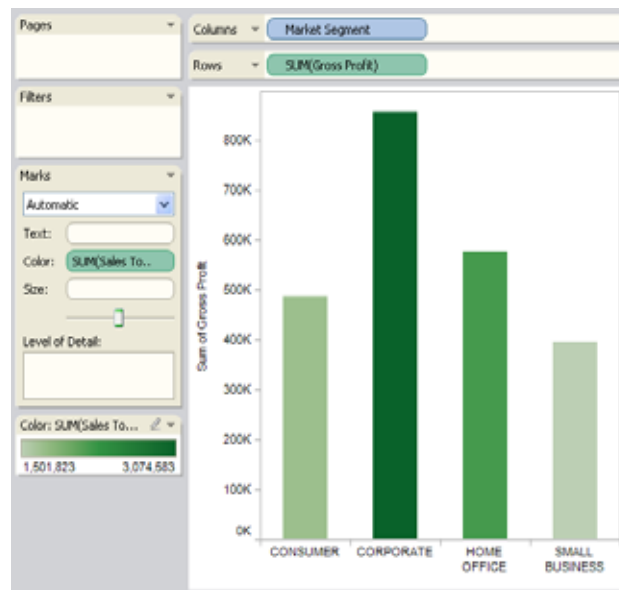
- **Discrete** – When you place a discrete field on the **Color** shelf, Tableau separates the marks according to the members in the discrete field and assigns a unique color to each member. The color legend displays each member name and its associated color. The default color palette contains 20 unique colors to encode discrete fields. If you have more than 20 members, the colors repeat. However, you can select from a variety of different color palettes that contain more unique colors. For information about selecting color palettes refer to “Color Shelf” on page 10-21.
- **Continuous** – When you place a continuous field on the **Color** shelf, Tableau draws each mark with a different color using a continuous range. A legend appears, showing the continuous color range.

As shown below, the bars are separated into colored segments according to the members of the **Product Category 2** discrete dimension. The dimension members and associated colors are displayed in the color legend.





In the view below, the bars are color-encoded by the **Sales Total** continuous measure. Note that each mark is encoded with only one color because the measure is aggregated. The continuous range of colors and numerical limits are displayed in the color legend.



To change the colors used for either dimensions or measures, double-click anywhere on the legend to open the Edit Colors dialog box. Refer to “Formatting Mark Colors” on page 20-49.

## Size Shelf

The **Size** shelf allows you to encode data by assigning different sizes to the marks in a data view. Depending on whether you use a discrete or continuous field you will add either categorical or quantitative sizes encoding. This section discusses the following topics:

- Categorical Sizes
- Quantitative Sizes

### Categorical Sizes

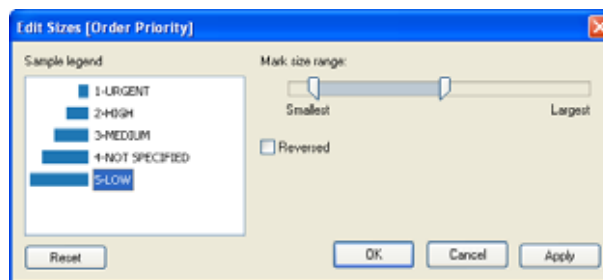
When you place a discrete field on the **Size** shelf, Tableau separates the marks according to the members in the dimension, and assigns a unique size to each member. Because size has

an inherent order to it (small to big), categorical sizes work best for ordered data like years or quarters.

Note that size-encoding data with a discrete field separates the marks in the same way as the **Level of Detail** shelf does, and then provides additional information (a shape) for each mark. When you add categorical size encoding to the view, a legend displays showing the sizes assigned to each member in the field placed on the size shelf. You can modify how these sizes are distributed in the Edit Sizes dialog box.

#### To edit categorical size encodings in a view:

- 1 Double-click on the legend or select **Edit Size** from the legend's menu to open the Edit Sizes dialog box.



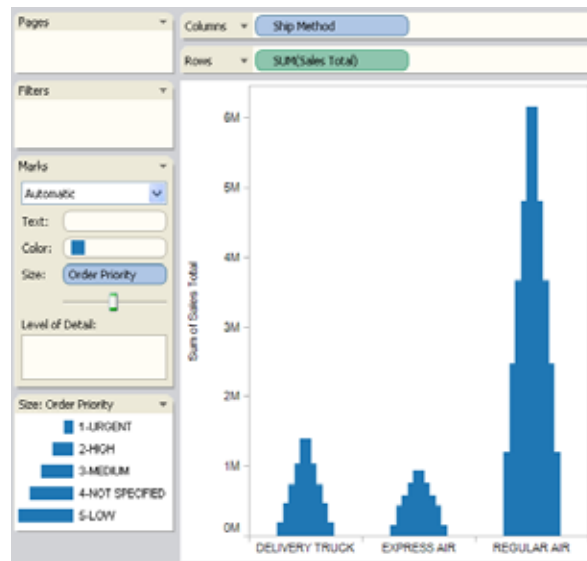
- 2 In the Edit Sizes dialog box, the sizes are displayed on the left and a size range slider is shown on the right. The sizes assigned to each member are distributed across the specified range. Use the slider to adjust the sizes assigned to each member.

You can also select **Reversed** to assign the largest mark to the smallest value and the smallest mark to the largest value.

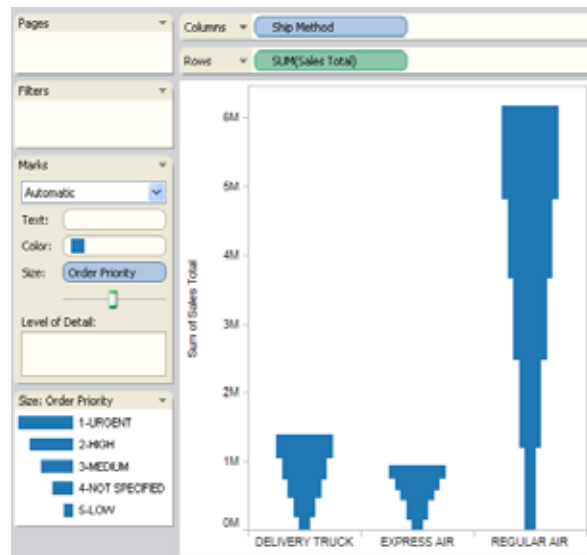
- 3 When finished click **OK**.

#### Example – Categorical Sizes

The view below shows the sales of a superstore broken down by shipping method. The order priority is indicated by the size of the bar.



In this case, the highest priority orders are shown with the smallest mark, which doesn't make sense. Use the Edit Sizes dialog box to Reverse the range so that the highest priority orders have the largest mark.



### Quantitative Sizes

When you place a continuous field on the **Size** shelf, Tableau draws each mark with a different size using a continuous range. The smallest value is assigned the smallest sized mark and similarly the largest value is represented by the largest mark.

When you add quantitative size encoding to the view, a legend displays showing the range of values over which sizes are assigned. You can modify how these sizes are distributed in the Edit Sizes dialog box.

**To edit quantitative size encodings:**

- 1 Double-click on the size legend or select **Edit Size** from the legend's menu to open the Edit Sizes dialog box.



- 2 In the Edit Sizes dialog box, select one of the following ways to map the sizes:
  - **Automatically** - selects the mapping that best fits your data. If the data is numeric and does not cross zero (e.g. all positive or all negative), the 'From zero' mapping is used. Otherwise, the 'By range' mapping is used.
  - **By range** - Uses the minimum and maximum values in the data to determine the distribution of sizes. For example, if a field has values from 14 to 25, the sizes will be distributed across this range.
  - **From zero** - Sizes are interpolated from zero making the maximum mark size assigned to the absolute value of the data value that is farthest from zero.
- 3 Use the range slider to adjust the distribution of sizes. When the From zero mapping is selected, the lower slider is disabled because it is always set to zero.
- 4 You can optionally select **Reversed** to assign the largest mark to the smallest value and the smallest mark to the largest value. This option is not available if you have selected to map the sizes from zero because the smallest mark is always assigned to zero.
- 5 Finally, you can select the **Start** and **End** checkboxes and manually type in a beginning and end value for the range of values to modify the distribution of sizes.
- 6 When finished, click **OK**.

### Example – Quantitative Sizes

The view below was created using a database that contains information about 1000s of TV advertisements for toys and video games for a one-month period. This example shows the amount of money each brand spends advertising on three television networks during the

weekends. The size of each mark indicates the amount of money spent on the advertisement while the color and shape represent the network and duration.



**Note** You can also change the size of the marks using the **Size** slider. For the line and bar mark types, the size slider controls the width of the mark. For the Gantt bar mark type, the size slider controls the length of the bar. For the other supported mark types, the size slider controls the area of the mark.

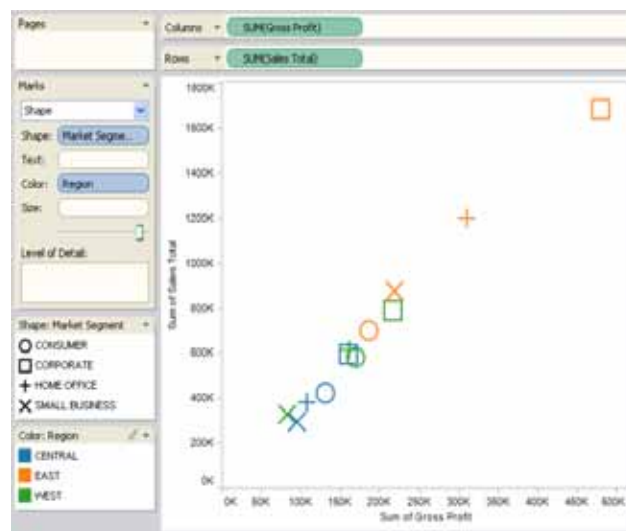
## Shape Shelf

The **Shape** shelf (**Ctrl+Alt+S**) allows you to encode data by assigning different shapes to the marks in a data view. When you place a dimension on the **Shape** shelf, Tableau separates the marks according to the members in the dimension, and assigns a unique shape to each member. The shape legend displays each member name and its associated shape. Ten unique shapes are used to encode dimensions. If you have more than 10 members, the shapes repeat. When you place a measure on the **Shape** shelf the measure is converted to a discrete measure.

Note that shape-encoding data separates the marks in the same way as the **Level of Detail** shelf does, and then provides additional information (a shape) for each mark. The **Shape**

shelf is available when you select the shape mark type from the **Mark** menu. It is the default mark type when measures are the inner fields for both the **Rows** shelf and the **Columns** shelf.

As shown below, the marks are separated into different shapes according to the members of the **Market Segment** dimension. Each shape reflects the each market segment's contribution to the gross profit and sales.



## Text Shelf

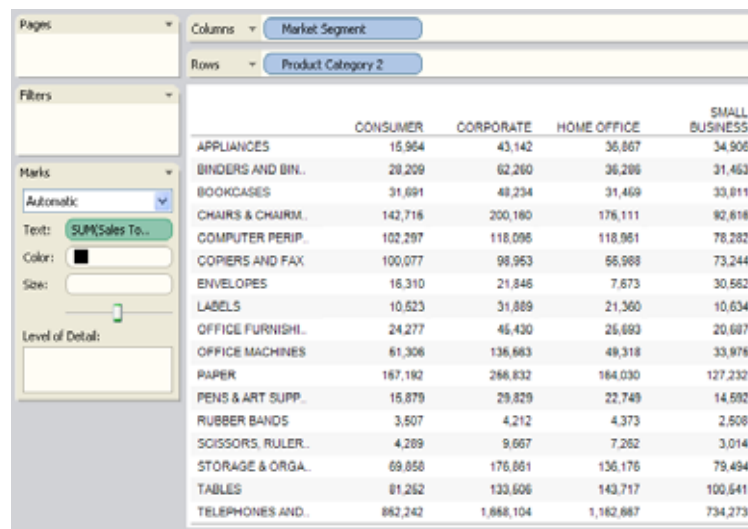
The **Text** shelf allows you to view the numbers associated with a data view, and to encode data by assigning text labels to the marks. The effect of text-encoding your data view depends on whether you use a dimension or a measure.

- **Dimension** – When you place a dimension on the **Text** shelf, Tableau separates the marks according to the members in the dimension. The text labels are given by the dimension member names.
- **Measure** – When you place a measure on the **Text** shelf, the text labels are given by the measure values. The measure can be either aggregated or disaggregated. However, disaggregating the measure is generally not useful because it often results in overlapping text.



Text is the default mark type when dimensions are the inner fields for both the **Rows** shelf and the **Columns** shelf. Refer to “Mark Types” on page 9-11 for more information. The text shelf is also available when you select the **Format > Automatic Data Labels** menu item.

The most common view using the **Text** shelf is a text table, which is also referred to as cross-tab or a PivotTable.

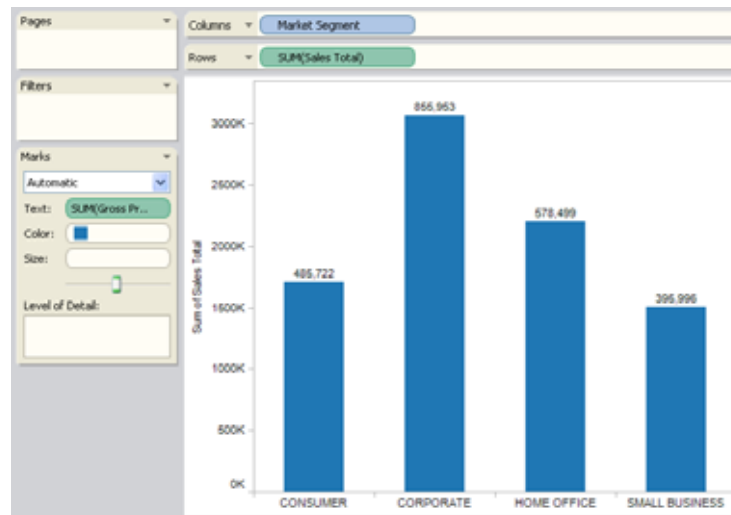


	CONSUMER	CORPORATE	HOME OFFICE	SMALL BUSINESS
APPLIANCES	15,964	43,142	36,667	34,906
BINDERS AND BIN.	29,209	62,260	36,206	31,463
BOOKCASES	31,691	49,234	31,469	33,811
CHAIRS & CHAIRM.	142,716	200,180	176,111	92,616
COMPUTER PERIP.	102,297	118,096	118,961	78,282
COPERS AND FAX	100,077	99,963	66,968	73,244
ENVELOPES	16,310	21,846	7,673	30,662
LABELS	10,623	31,889	21,360	10,634
OFFICE FURNISH.	24,277	46,430	26,693	20,687
OFFICE MACHINES	61,306	136,683	49,318	33,976
PAPER	167,192	266,832	164,030	127,232
PENS & ART SUPP.	15,879	29,829	22,749	14,592
RUBBER BANDS	3,507	4,212	4,373	2,608
SCISSORS, RULER.	4,269	9,667	7,262	3,014
STORAGE & ORGA.	69,858	176,861	136,176	79,494
TABLES	81,262	133,606	143,717	100,641
TELEPHONES AND.	862,242	1,668,104	1,162,667	734,273

**Note** You can display text labels with other mark types by selecting the **Format > Automatic Data Labels** menu item. Refer to “Data Labels” on page 21-17 to learn more about showing and hiding data labels.

If you place a dimension on the **Text** shelf, the marks are separated and labelled according to the dimension member names. If you place a measure on the **Text** shelf, the marks are labelled by the values contained by the measure.

As shown below, the heights of the bars are given by the **Sales Total** measure and the labels are given by the sum of the **Gross Profit** measure.



## Path Shelf

The **Path** shelf allows you to encode data by connecting marks using a particular drawing order. You can path-encode your data using either a dimension or a measure.

- **Dimension** – When you place a dimension on the **Path** shelf, Tableau connects the marks according to the members in the dimension. If the dimension is a date, the drawing order is given by the date order. If the dimension holds words such as customer names or product types, the drawing order is given by the order of the members in the data source. You can change the order by which data points are connected by changing the sort order of the members. Refer to “Sorting” on page 12-3.
- **Measure** – When you place a measure on the **Path** shelf, Tableau connects the marks according to the values of the measure. The measure can be aggregated or disaggregated.

The **Path** shelf is available only when you select the line or polygon mark type from the **Mark** menu. Refer to “Mark Types” on page 9-11 for more information.

To create a useful path-encoded view, your data table should contain at least one measure. This is because you cannot create a path that connects only categorical data (dimensions).

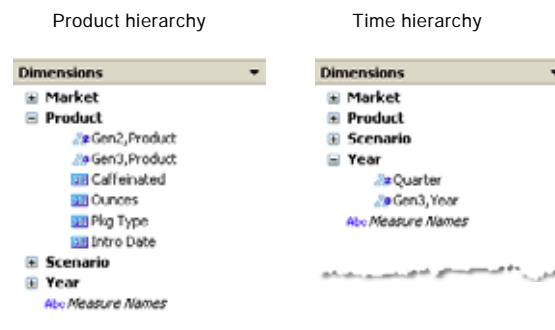
The view below was created using storm data from the Atlantic basin in 2005. The view uses line marks with the path determined by the date of the storm. In this example, it lets you see the path of the storm.



By placing the continuous date on the **Path** shelf, the lines are drawn in chronological order.

## Multidimensional Hierarchies

Multidimensional data sources contain hierarchies. For example, your database might contain a **Product** dimension that includes members such as product family, product department, and so on organized into a hierarchy, or you might have a **Time** dimension that includes years, quarters, and months.



When you are building a view using a multidimensional hierarchy, there are some important concepts to understand. This section discusses the following topics:

- Drilling Down and Up in a Hierarchy
- Building Views with Hyperion Essbase
- Building Views with Microsoft Analysis Services Cubes
- Perfect Pivoting
- Defining Unique Values


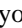
### Drilling Down and Up in a Hierarchy

One of the most useful ways to navigate hierarchies is to *drill down* or *drill up*. For example, if you are examining the sales totals for various years, you can then drill down and view sales for all of the months within each year. Alternatively, if you are examining sales totals for all months, you can then drill up and view the sales for each year.

You can drill down and drill up in Tableau by clicking on fields placed on shelves, or by selecting a hierarchy header in the table. These two methods are described below.

## Using Fields on Shelves

You can drill down or drill up by clicking on a dimension that is placed on any shelf. If the dimension is on the **Rows** or **Columns** shelf, drilling down shows more data (more headers) in the table, while drilling up shows less data in the table.

You can click on the plus/minus control that appears on any hierarchical dimension on any shelf. If a dimension member shows the plus sign , then its children are not already showing and you can drill down at least one level. If a dimension member shows the minus sign , then its children are already showing and you can drill up.

The following figure demonstrates drilling down one level in the hierarchy for the Store Country dimension to expose **States**.



## Using Headers

You can drill down and drill up for individual dimension members in a hierarchy by right-clicking a table header and selecting **Drill Down** or **Drill Up** from the context menu. This is often referred to as nonuniform drill down because you expose only the members of interest instead of exposing all the members of a given level.

For example, the following figure illustrates drilling down into the **Beverages** member. Note that new row headers are displayed in the table and that **Product Category**, which is the next member in the hierarchy, is automatically displayed on the **Rows** shelf.

The figure shows two Tableau views. The left view shows a table with 'Store Country' and 'Store State' on the Columns shelf and 'Product Department' on the Rows shelf. The 'Beverages' row is selected, and a context menu is open with 'Drill Down' highlighted. The right view shows the result of drilling down: 'Product Category' is added to the Rows shelf, and the table displays detailed data for Beverages, including Carbonated Beverages, Drinks, Hot Beverages, and Pure Juice Beverages.

		USA		
		CA	OR	WA
Alcoholic Beverages		\$4,011.37	\$3,460.32	\$6,667.39
<b>Beverages</b>		\$6,947.13	\$12,791.84	\$11,217.84
Dairy		\$2,185.34	\$1,729.84	\$3,143.42
Baked Goods		\$4,619.81	\$4,088.96	\$7,746.67
Breakfast Foods		\$1,797.19	\$3,143.42	\$1,729.84
Canned Foods		\$9,956.47	\$18,224.88	\$9,956.47
Canned Products		\$838.20	\$1,676.40	\$838.20
Deli		\$7,132.96	\$6,348.98	\$11,217.84

		USA		
		CA	OR	WA
Alcoholic Beverages		\$4,011.37	\$3,460.32	\$6,667.39
Beverages	Carbonated Beverages	\$1,096.84	\$1,511.01	\$3,028.50
	Drinks	\$1,710.08	\$1,478.12	\$2,456.09
	Hot Beverages	\$2,626.93	\$2,378.87	\$4,256.94
	Pure Juice Beverages	\$1,972.68	\$1,581.13	\$3,064.34
Dairy		\$2,185.34	\$1,729.84	\$3,143.42
Baked Goods		\$4,619.81	\$4,088.96	\$7,746.67
Breakfast Foods		\$1,797.19	\$3,143.42	\$1,729.84
Canned Foods		\$9,956.47	\$18,224.88	\$9,956.47
Canned Products		\$838.20	\$1,676.40	\$838.20
Deli		\$7,132.96	\$6,348.98	\$11,217.84

One reason to use nonuniform drill down is if your data source has a ragged hierarchy (asymmetric layout). You also might want to view the children for just the member of interest.

**Note** Drilling down and drilling up results in filtering the data. Refer to “Filtering” on page 12-19 for more information.

## Building Views with Hyperion Essbase

When Tableau is connected to an Hyperion Essbase data source, there are three important features that you should know about:

- **Generations and Levels** – For any given dimension, you can navigate using either levels or generations.
- **Shared Members** – You can choose to include or exclude dimension members that appear in more than one place in a hierarchy.
- **Setting the Default Member** – change the default member to avoid repeatedly setting filters.

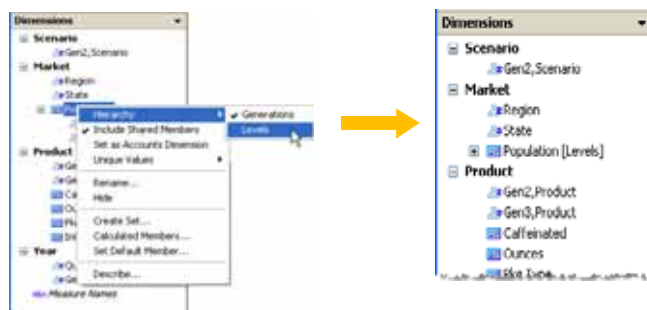
### Generations and Levels

In Tableau, you can work with either the generations or the levels of a dimension. The generations of a dimension are all members that are an equal distance from the root of the

dimension. The levels are all members that are an equal distance from the leaves of the dimension. For balanced dimensions, you'll typically want to work with generations. However, if your dimension is ragged, then it may make more sense to navigate using levels.

By default, the generations of each dimension are listed in the Data window. When you drag a dimension to a shelf, all generations that are ancestors of the selected generation (all generations that are above it in the hierarchy) are automatically included in the placement.

If you would rather navigate using the levels of a dimension, right-click the name of the dimension and then select **Hierarchy > Levels**.



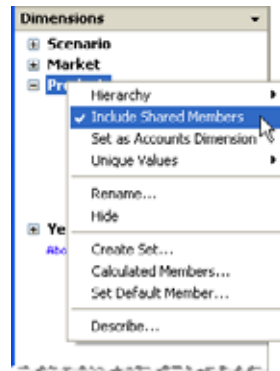
**Note** If you are using the same dimension in multiple worksheets, you can use levels in one worksheet and generations in another worksheet simultaneously. Furthermore, you can mix generations and levels from different dimensions in the same worksheet.

### Shared Members

Shared members are dimension members that appear in more than one place in a hierarchy. For example, Diet Coke might be part of the product generation. But it might be shared by both the diet colas branch and the colas branch of the product hierarchy above it. In the database, however, the data about Diet Coke is stored just once.

By default, Tableau includes shared members in all generations (or levels) of a dimension. This means that a shared member might appear multiple times in a table. If you choose to exclude shared members, they will appear only once in a table. By default, shared members are included for all dimensions. You exclude shared members for a given dimension

hierarchy by right-clicking the dimension name in the Data window and selecting **Include Shared Members** from the menu.





The figure below shows part of a data view where shared members are included (left) and excluded (right). Notice that diet drinks are shared members.

Shared members included.

Colas	Cola	
	Diet Cola	
	Caffeine Free Cola	
Root Beer	Old Fashioned	
	Diet Root Beer	
	Sasparilla	
	Birch Beer	
Cream Soda	Dark Cream	
	Vanilla Cream	
	Diet Cream	
Fruit Soda	Grape	
	Orange	
	Strawberry	
Diet Drinks	Diet Cola	
	Diet Root Beer	
	Diet Cream	

Shared members.

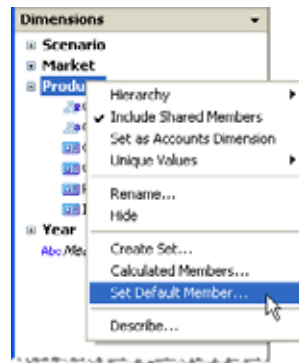
Shared members excluded.

Colas	Cola	
	Diet Cola	
	Caffeine Free Cola	
Root Beer	Old Fashioned	
	Diet Root Beer	
	Sasparilla	
	Birch Beer	
Cream Soda	Dark Cream	
	Vanilla Cream	
	Diet Cream	
Fruit Soda	Grape	
	Orange	
	Strawberry	

### Setting the Default Member

All multidimensional data sources have a default members that are set when the data source is first built. If you find that you are creating filters all the time to change look at the same specific data, you may find it useful to change the default member. For example, if you are the regional manager for the Western region in a company and you only want to look at your region's numbers, you can set the default member to the Western region.

You can change the default member in Tableau by right-clicking a dimension hierarchy and selecting **Set Default Member**.



In the subsequent dialog box, select from the following options:


- Default member defined on cube – uses the default member that was defined when the cube was built. This is the default setting in Tableau.
- (All) member for the hierarchy – uses the ALL member for the selected hierarchy as the default member.
- Selected member – uses the member that you select in the bottom half of the dialog box as the default member.

The default member determines how you view the cube and so is much more powerful than applying filters. All fields will be calculated based on the default member you select. In addition, these default member settings are saved with the connection.

## Building Views with Microsoft Analysis Services Cubes

When you build views in Tableau using a Microsoft Analysis Services Cube it is possible to have measures and dimensions that don't make a lot of sense when placed in the view together. For example, you may have a measure for Sales Quota. It won't make sense to place that measure against a dimension containing products if products don't have sales quotas. Tableau helps you figure out the dimensions and measure that can be used together in meaningful ways by highlighting unrelated dimensions and measures in gray. So in the last example, when we place Sales Quota onto a shelf, the products dimensions are highlighted in gray. **Highlighted dimensions are not disabled and can still be added to**

**the view.** When you add an incompatible measure to the view, the measure is marked with a caution symbol.



## Perfect Pivoting

In Tableau, *perfect pivoting* refers to working with hierarchies in these ways:


- Using varying levels of detail including skipping levels (for example, Country and City, but not State).
- Using varying levels of detail on different worksheet shelves simultaneously (for example, Product Family on the **Columns** shelf and Product Department on the **Color** shelf).
- Using varying levels of detail out of order (for example, Quarter before Year).

For example, the data source used to create the view below contains the **Store** hierarchy, which includes four levels. The hierarchy is broken and the **Store State** level is placed on the **Columns** shelf, while **Store City** level is placed on the **Color** shelf.

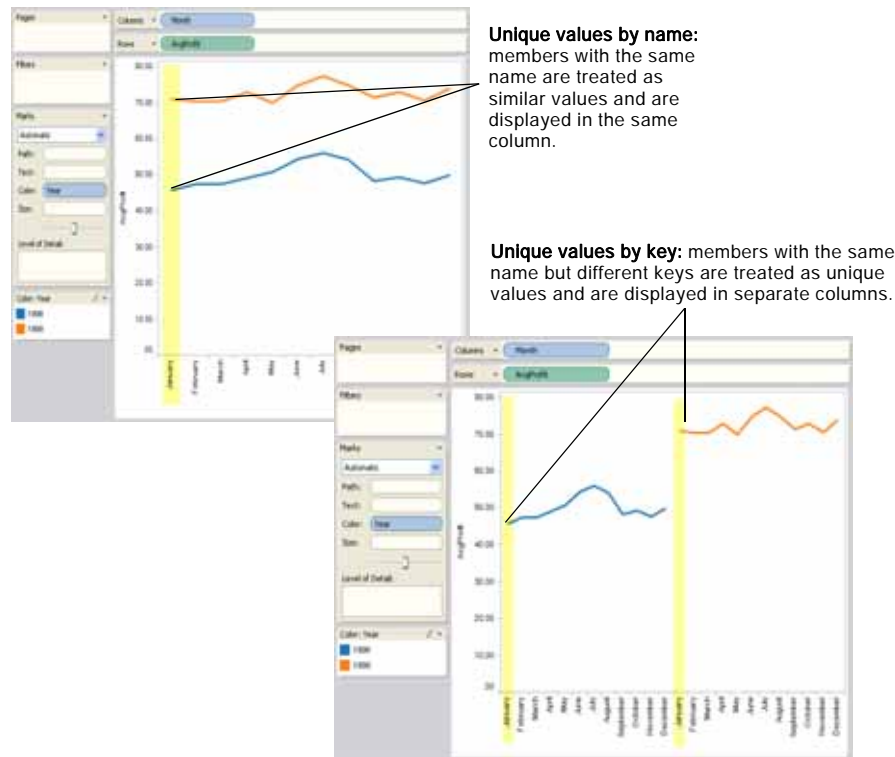


## Defining Unique Values

Sometimes, when you are building views in Tableau, a field will have multiple members with the same name. For example, you may have a view showing the average profit by month over several years. The month January appears multiple times (once for each year).



While the *name*, January, is repeated, each instance of January can either be considered similar or unique. If you consider them similar, they will appear in the same column if you decided to move the Year field to the color shelf. However, if you consider them unique, they will be treated as two different values.



It is generally okay to consider repeated names within date and time fields (like in the previous example) similar but if there are repeated names in the Customer Name field, you won't want to consider the two customers as the same person.

You can define how you want Tableau to determine whether repeated values are unique by right-clicking on the dimension and selecting one of the following on the **Unique Values** context menu:

- **By Key:** each member is considered unique based on the key given it by the system administrator when the database is set up. Members with the same name but different keys are treated as unique values.
- **By Name:** each member is considered unique based on the member name. Members with the same name (regardless of their keys) are treated as if they are the same.

By default, unique date and time values are determined by name and all other values are determined by key.

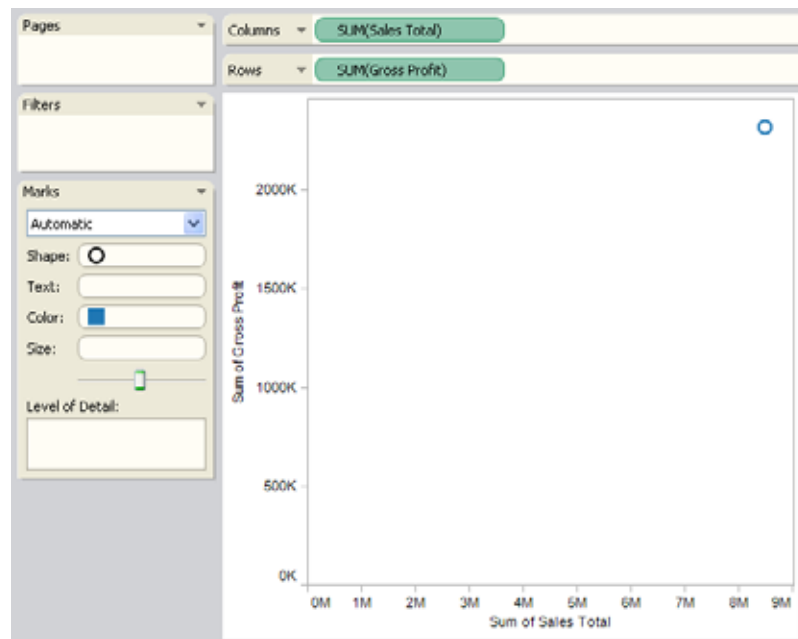
## Example – Building Data Views Manually

In this example, you will use the Sample - Superstore Sales (Excel) data source to create a view that contains two aggregated measures displayed as a scatter plot. The data are color-encoded and shape-encoded, and an additional level of detail is included. The data are also filtered.

To create the view, do the following:

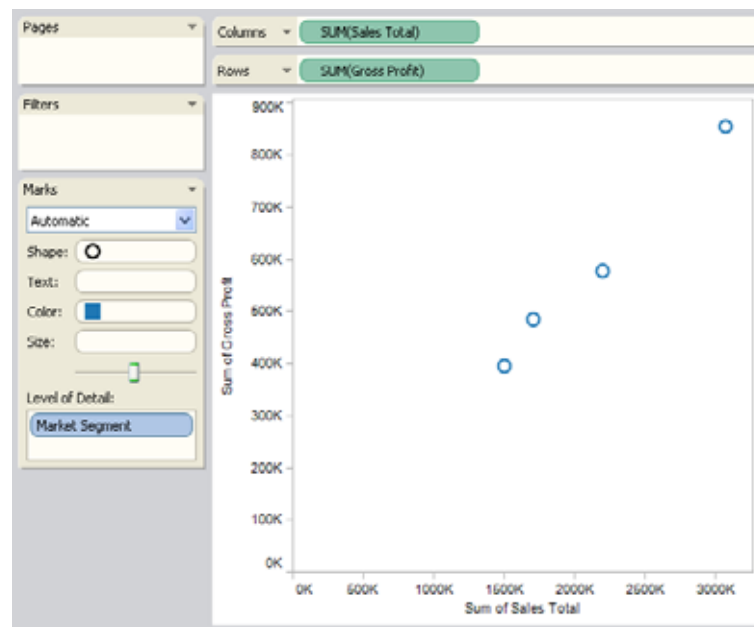
- 1 Place the **Sales Total** measure on the **Columns** shelf and the **Gross Profit** measure on the **Rows** shelf.

The measures are automatically aggregated and result in one data point. The data point is displayed using the shape mark type. Note that you are not displaying any levels of detail because dimension members are not included in the view.



- 2 Place the **Market Segment** dimension on the **Level of Detail** shelf.

The original mark is now separated into four marks, where each new mark is associated with a member (level of detail) of the **Market Segment** dimension.





3 Place the **Region** dimension on the **Color** shelf.


Each of the four marks is now separated into three new marks for a total of twelve. Each new mark is associated with a member of the **Region** dimension, and is encoded with a unique color. The color legend displays each member name and its associated color.




4 Place the **Product Category 1** dimension on the **Shape** shelf and filter the dimension to exclude Technology products.

Each of the twelve marks is now separated into three new marks for a total of 36. Each new mark is associated with a member of the **Product Category 1** dimension, and is encoded with a unique shape. The shape legend displays each member name and its associated shape.

To filter the data, select **Filter** from the **Product Category 1** context menu. The **Filter** dialog box opens. Deselect **Technology** to exclude it from the view.

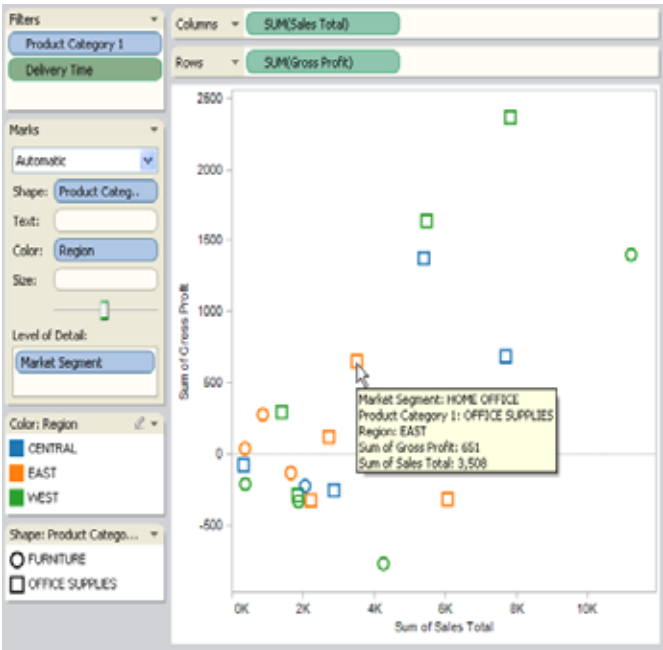


- 5 Place the **Delivery Time** measure on the **Filters** shelf. Select **Delivery Time** (without an aggregation) in the Filter Field dialog box and then filter the data to exclude delivery times less than 96 days.



The filter acts on every data source row before the profit and sales summations are performed. This results in the exclusion of additional data from the view.

The final view is shown below



# Building Views Automatically

<b>Overview</b> . . . . .	<b>11-2</b>
<b>Automatic Views: Show Me!</b> . . . . .	<b>11-3</b>
Show Me! . . . . .	11-4
Show Me! Alternatives . . . . .	11-6
<b>Add to Sheet: Double-Click</b> . . . . .	<b>11-8</b>

---

## Overview

Tableau contains a suite of tools designed to help you quickly create useful views. Two situations in which you would want to create views automatically are when you want:

- Better Insight – People often have difficulty mapping data to views that address their analytical or presentation needs. Tableau contains built-in rules that are used to examine data and suggest ways of looking at it. In this way the software acts as a tour guide for producing useful views of data.
- Time Savings – Building data views manually can sometimes be time consuming. Using Tableau's automatic features can help you work faster by giving you a starting view that you can further refine manually.

The specific methods for automatically generating views of data fall into two categories:

- Automatic Views: Show Me!
- Add to Sheet: Double-Click

## **Automatic Views: Show Me!**

Show Me! functionality is available in two different options:

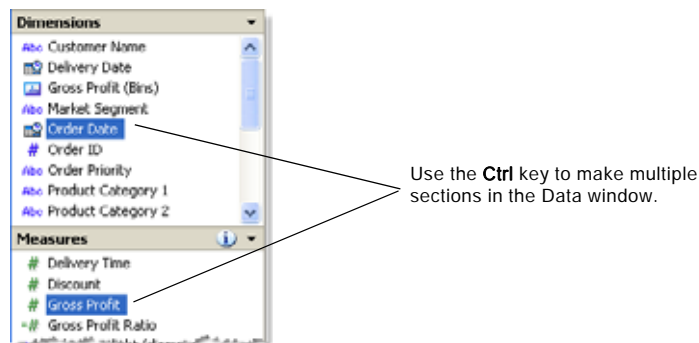
- Show Me!
- Show Me! Alternatives

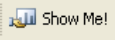
## Show Me!

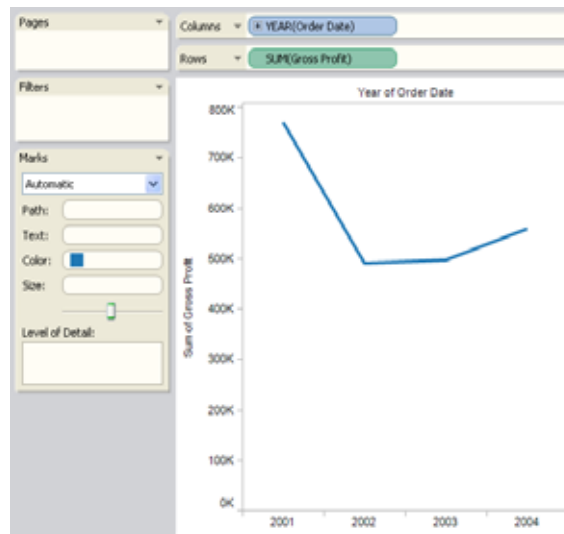
The first Show Me! method is the most straightforward in that it generates a single view of your data. Select fields you want to analyze in the Data window and press the Show Me! button on the toolbar. Tableau automatically generates a view of the data in those fields.

### 1 Select Input.

Select fields in the Data window that you want to analyze. Use the **Ctrl** key to make multiple selections.



- 2 Click **Show Me!**  on the toolbar.
- 3 View the Result. Tableau automatically creates a view of the data.



How does Tableau know what view to create? It examines information about the fields you selected in the Data window and makes a “best guess” suggestion based on that information. For instance, in the example shown above, a date field was selected as well as a continuous measure. Usually the best way to view a continuous measure over time is with a line.

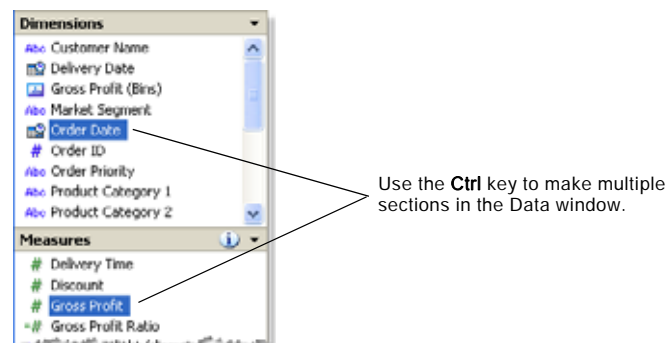


## Show Me! Alternatives

This choice adds a dialog box to the Show Me! method which permits alternative views of your data.


### 1 Select Input.

Select fields in the Data window that you want to analyze. Use the **Ctrl** key to make multiple selections.



### 2 Click **Show Me! Alternatives** on the toolbar.

### 3 Select an alternative in the dialog box:



Any alternative that is not grey will generate a view of your data. The default alternative is the one that is used by the Show Me! command. ToolTips describe the minimum requirements for each alternative.

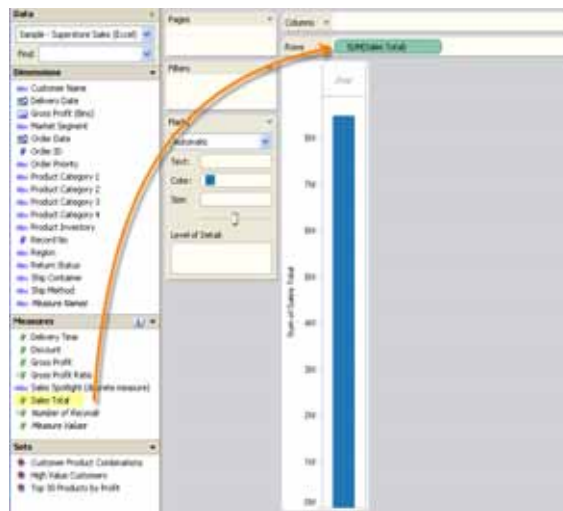
**4** View the Result.

## Add to Sheet: Double-Click

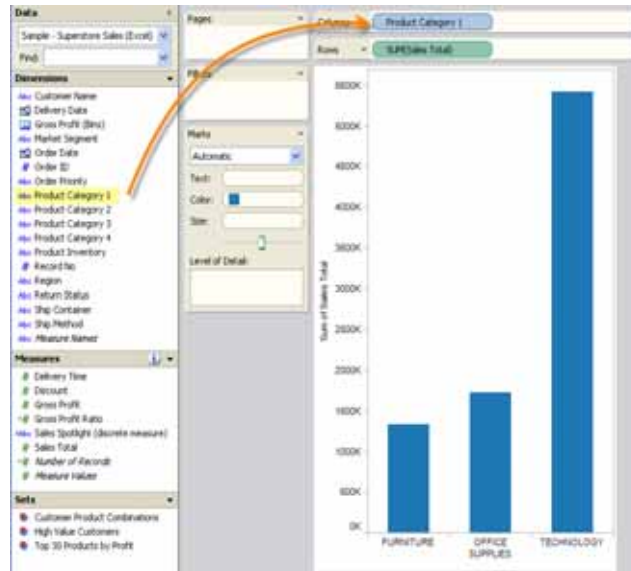
Tableau supports an additional method for automatically generating views of data called Automatic Double-Click. To use this method, double-click fields in the Data window you are interested in. Tableau automatically adds each field to the view. That is, each double-click results in an additional field added to a shelf in an intelligent way. Like the Show Me! feature, this function leverages Tableau's ability to make an intelligent "best guess" of where the field should be displayed.

**Here's how it works:**

- 1 Double-clicking the **Sales** measure in the Data window automatically adds that field to the view in an intelligent fashion.

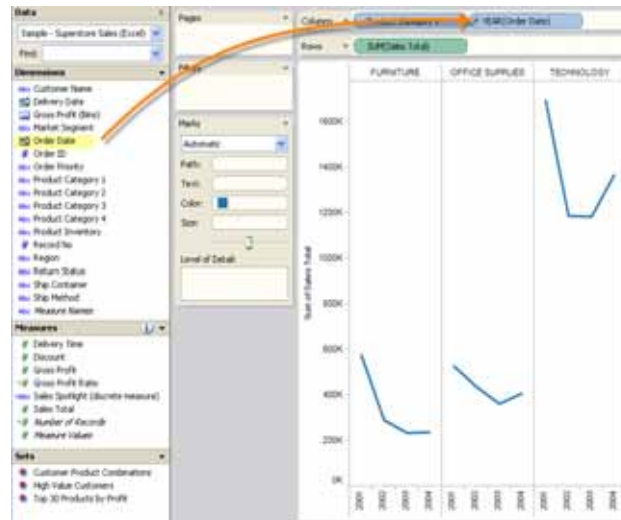


- 2 Double-clicking the **Product Category 1** dimension in the Data window automatically adds that field to the view in an intelligent fashion.



- 3 Double-clicking the **Order Date** dimension in the Data window automatically adds that field to the view in an intelligent fashion. As you double-click fields they are

successively added to the view. The order in which you click fields determines the type of view created.



The following table describes some of the rules used in creating automatic views via the Double Click method.

Text Table	Clicking a Dimension first produces a text table (or cross-tab). All subsequent clicks on fields result in refinement of the text table.
Bars	Clicking a Measure first and then a Dimension produces a bar view. All subsequent clicks result in refinement of the bar view, unless a Date Dimension is clicked, at which time the view is changed to a line.
Discrete Line	Clicking a Measure and then a Date Dimension produces a discrete line view. All subsequent clicks result in refinement of the discrete line view.

Continuous Line	Clicking a continuous dimension and then a Measure produces a continuous line view. All subsequent clicks on dimensions result in refinement of the continuous line view. All subsequent clicks on measures add quantitative axes to the view.
Scatter	Clicking a Measure and then another Measure produces a scatter view. Subsequent dimensions result in refinement to the scatter view. Subsequent measures will create a Scatter Matrix.

For a full discussion of dimensions, measures, data types and data roles, refer to “Data Types & Roles” on page 8-20.

# Sorting, Filtering, and Grouping

<b>Overview</b>	<b>12-2</b>
<b>Sorting</b>	<b>12-3</b>
Computed Sorting	12-3
Manual Sorting	12-15
<b>Filtering</b>	<b>12-19</b>
How Filtering Works	12-19
Filtering Discrete Fields	12-20
Filtering Continuous Fields	12-36
Quick Filters	12-46
External Filters	12-50
Context Filters	12-55
<b>Ad-Hoc Groups</b>	<b>12-62</b>
Creating Ad-Hoc Groups	12-62
Editing an Existing Group	12-66
Finding Members in the Ad-Hoc Group Dialog Box	12-68
<b>Sets</b>	<b>12-70</b>
About Sets	12-70
How to Create a Set	12-72
Creating Sets Examples	12-81

---

## Overview

After you understand the basics of building data views, use sorting, filtering, groups, and sets to further refine your views and extract exactly the information you are looking for. This section discusses how to re-order and sort the data in a view, filter out unnecessary rows and columns, and group dimension members into higher level categories, and create a set using multiple dimensions to create richer encodings.

- **Sorting** – Display your data in ascending or descending order based on other fields or custom formulas using computed sorts. Or you can manually sort your data to display in whatever order you choose.
- **Filtering** – Include only the relevant data and exclude the rest by selecting specific dimension members to remove, defining formulas that compute which rows and columns to include, and specifying a numeric range of values to include.
- **Ad-Hoc Groups** – Combine dimension members into higher level categories.
- **Sets** – Create a custom field based on existing dimensions that can be used to encode the view with multiple dimension members across varying dimension levels.



## Sorting

In Tableau, sorting a data view means arranging dimension members in a specified order. Tableau supports two different types of sorting:

- Computed Sorting
- Manual Sorting

### Computed Sorting

You might want to sort customers by alphabetical order, or sort a product line from lowest sales total to highest sales total. Both of these sorts are “computed sorts” because they use programmatic rules that you define to sort the field. This section discusses the following topics:

- About Computed Sorting
- How to Sort Data (Computed Sorts)
- Example – Sorting a Text Table
- Example – Sorting a Multidimensional Hierarchy

#### About Computed Sorting

Sorting dimensions in a computed manner follows these rules:

- You can sort any discrete field after it has been placed on a shelf (except the **Filters** shelf).
- Each dimension that appears on a worksheet can be sorted independently of any other dimension.
- The shelf location of the dimension determines the component of the data view that’s sorted. For example, if the dimension resides on the Columns shelf, the columns of the data view are sorted for that field. If the dimension resides on the **Color** shelf, the color encodings are sorted.
- Sorts are computed based on the values of the filters and sets in the view. Refer to “Filtering” on page 12-19 for more information.
- Sorted fields are identified with bold names.

Continuous fields are automatically sorted from lowest number to highest number (as indicated by the axes) and you cannot manually change the sort. However, you can reverse

the order of an axis using field specific formatting. Refer to “Editing Axes” on page 20-37 for more information.

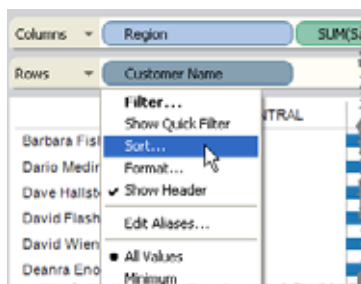
### How to Sort Data (Computed Sorts)

Use the sort dialog box to apply computed sorts to fields in the view.

#### To apply computed sorts:

- 1 Open the Sort dialog box.

Right-click on the field that you want to sort and select **Sort** from the its context menu.



- 2 Specify the sorting options.

Complete the **Sort** dialog box by specifying the following criteria:

- Sort order – Display the sort results in ascending or descending order.
- Sort by – Sort by one of these three options:

**Data source order** - the order that the data source naturally orders the data. Generally for relational data sources, this tends to be in alphabetical order. If you are using a cube, this order is the defined hierarchial order of the members within a dimension.

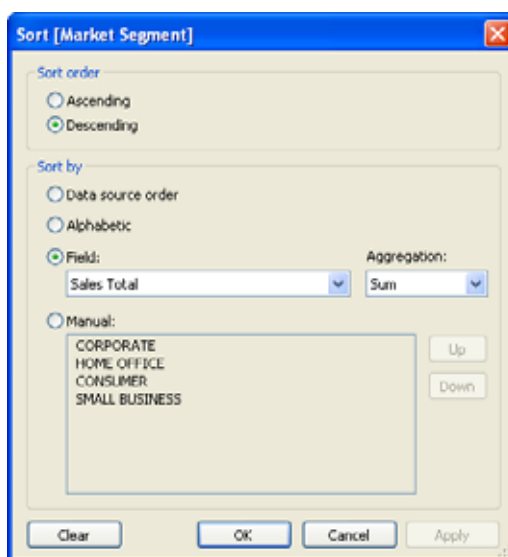
**Alphabetic** - the order of the letters in the alphabet.

**Field** - order the data based on the associated values of another field. For example, you could order several products by their total sales values.

When working with a relational data source and sorting by another field, you must also specify the aggregation function to use. This option is not available for

multidimensional data sources because aggregations are defined when the cube is created and cannot be modified in Tableau.

A typical scenario is to sort one or more discrete fields by a continuous field using the same aggregation. For example, the **Sort** dialog box shown below is configured to sort the members of the **Market Segment** field in descending order and by the sum of the **Sales Total** measure. The results will be displayed so that the member with the largest sales total is displayed first, the member with the second largest sales total is displayed second, and so on.



You should keep the following rules in mind when interpreting the sort results:

- Tableau computes the sort across the entire table using the specified criteria. Refer to “Example – Sorting a Text Table” on page 12-7 for an example.
- Sorts do not break the dimension hierarchy. Sorted fields are always displayed within the ordered context already set forth by the fields on the **Rows** and **Columns** shelves. This means that Tableau will not rearrange any of the headers of the fields that appear before

(to the left of) the sorted field. Refer to “Example – Sorting a Multidimensional Hierarchy” on page 12-10 for an example.

If you want to break the dimension hierarchy when sorting a multidimensional data source, place only the hierarchy level that you want sorted on the **Rows** or **Columns** shelf.

### Example – Sorting a Text Table

Using the Sample - Superstore Sales (Excel) data source, this example sorts the rows and columns of a text table to determine which products and years have the highest average delivery time. To create the view, follow the steps below:

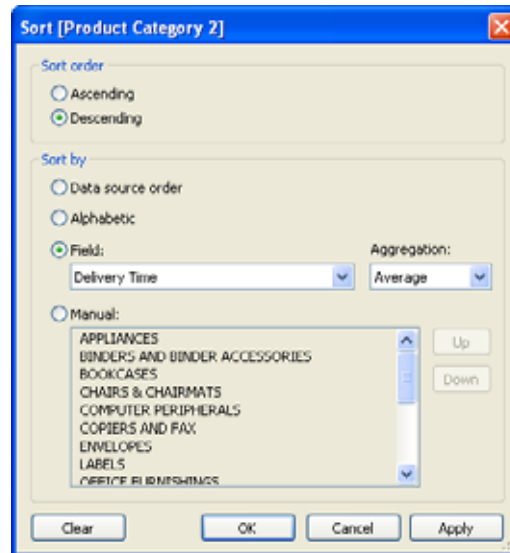
- 1 Place the **Order Date** dimension on the **Columns** shelf and the **Product Category 2** dimension on the **Rows** shelf.

Complete the text table by placing **Delivery Time** on the **Text** shelf and aggregating the measure as an average (select **Average** from the field's context menu). By default, the table is sorted in alphabetical order.

Columns ▾ + YEAR(Order Date)			
Rows ▾ Product Category 2			
	2001	2002	2003
APPLIANCES	3.88	4.95	25.80
BINDERS AND BINDER AC..	3.99	5.45	31.52
BOOKCASES	3.82	9.21	38.80
CHAIRS & CHAIRMATS	3.31	8.65	35.52
COMPUTER PERIPHERALS	3.69	5.62	28.40
COPIERS AND FAX	3.33	5.25	23.25
ENVELOPES	4.33	5.43	38.05
LABELS	3.84	5.54	31.91
OFFICE FURNISHINGS	3.27	5.17	26.99
OFFICE MACHINES	4.00	6.00	15.83

- 2 Sort the fields.

Right-click on **Order Date** field and select **Sort**. In the Sort dialog box select **Descending** as the Sort Order and sort by **Delivery Time** aggregated as an **Average**. When finished click **OK**. Then apply the same sort to **Product Category 2**.



The view is shown below. **Bookcases** is the top row in the table because it has the largest average delivery time across all years, while **Office Machines** is the bottom row in the table because it has the smallest average delivery time across all years. Similarly, **2003** is the left most column because it has the largest average delivery time for all products, while **2001** is the right most column because it has the smallest average delivery time for all products.

	2003	2002	2004	2001
BOOKCASES	38.80	9.21	9.10	3.82
ENVELOPES	38.06	6.43	6.36	4.33
CHAIRS & CHAIRML.	36.52	8.66	7.41	3.31
RUBBER BANDS	35.35	6.00	5.67	4.19
TABLES	34.09	7.22	8.69	3.64
LABELS	31.91	5.54	5.29	3.84
TELEPHONES AND..	30.88	5.62	5.44	3.86
PAPER	29.82	6.69	6.27	3.86
STORAGE & ORGA..	36.88	5.64	6.21	4.07
BINDERS AND BIN..	31.52	5.45	5.75	3.99
COMPUTER PERIP..	28.40	6.82	6.27	3.89
SCISSORS, RULER..	32.15	4.85	6.06	3.76
PENS & ART SUPP..	26.51	5.69	5.45	3.54
OFFICE FURNISHI..	26.99	5.17	5.46	3.27
APPLIANCES	25.00	4.95	4.94	3.88
COPIERS AND FAX	23.25	5.25	4.00	3.33
OFFICE MACHINES	16.83	6.00	6.00	4.00

At first glance, it's not clear if the data have been correctly sorted. That's because Tableau computes the sort across the entire table using the specified criteria. By turning grand totals on for both columns and rows, using the **Table** menu, you can see that the sort was performed correctly.

Pages

Filters

Marks

Automatic

Text: AVG(Delivery ..

Color:

Size:

Level of Detail:

Columns

+ YEAR(Order Date)

Rows

Product Category 2

	2003	2002	2004	2001	Grand Total
BOOKCASES	38.80	9.21	9.10	3.02	13.13
ENVELOPES	38.05	5.43	5.35	4.33	11.72
CHAIRS & CHAIRMATS	35.52	8.65	7.41	3.31	11.40
RUBBER BANDS	35.35	6.00	5.67	4.19	11.18
TABLES	34.09	7.22	8.59	3.54	11.03
LABELS	31.91	5.64	5.29	3.84	10.97
TELEPHONES AND CO..	30.88	5.62	5.44	3.06	10.79
PAPER	29.62	5.59	5.27	3.85	10.49
STORAGE & ORGANIZ..	35.88	5.64	6.21	4.07	10.32
BINDERS AND BINDER..	31.52	5.45	5.75	3.99	10.22
COMPUTER PERIPHER..	28.40	5.62	5.27	3.69	9.91
SCISSORS, RULERS A..	32.15	4.85	6.05	3.76	9.76
PENS & ART SUPPLIES	26.51	5.69	5.45	3.54	9.58
OFFICE FURNISHINGS	28.99	5.17	5.45	3.27	9.28
APPLIANCES	25.80	4.95	4.94	3.88	8.02
COPERS AND FAX	23.25	5.25	4.00	3.33	6.57
OFFICE MACHINES	15.83	8.00	6.00	4.00	5.22
Grand Total	30.34	5.73	5.52	3.78	10.35

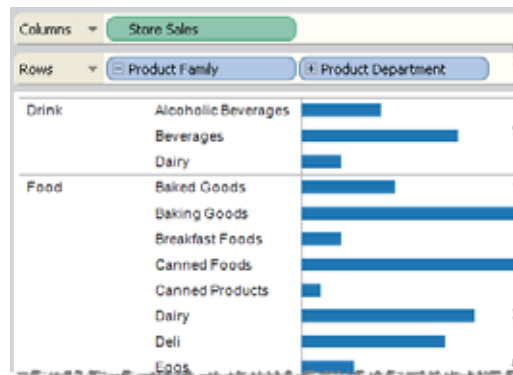
### Example – Sorting a Multidimensional Hierarchy

This example sorts the rows of a bar chart, created using a multidimensional data source, to determine which products have the highest store sales. To create the view, follow the steps below.



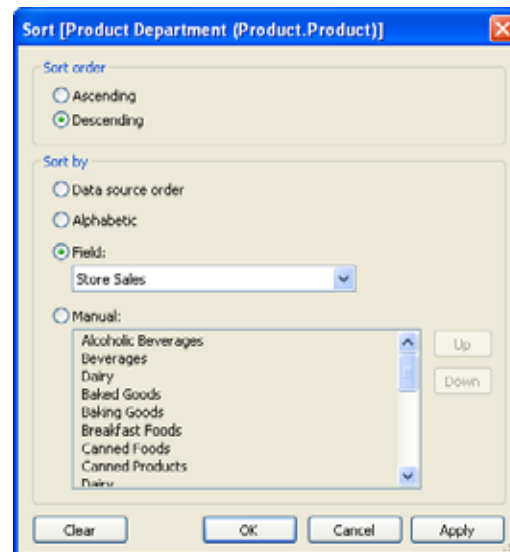
- 1 Place the **Store Sales** measure on the **Columns** shelf and the **Product Family** dimension on the **Rows** shelf.

Drill down one level in the hierarchy to display **Product Department**. By default, the table is sorted in based on the order in the data source.



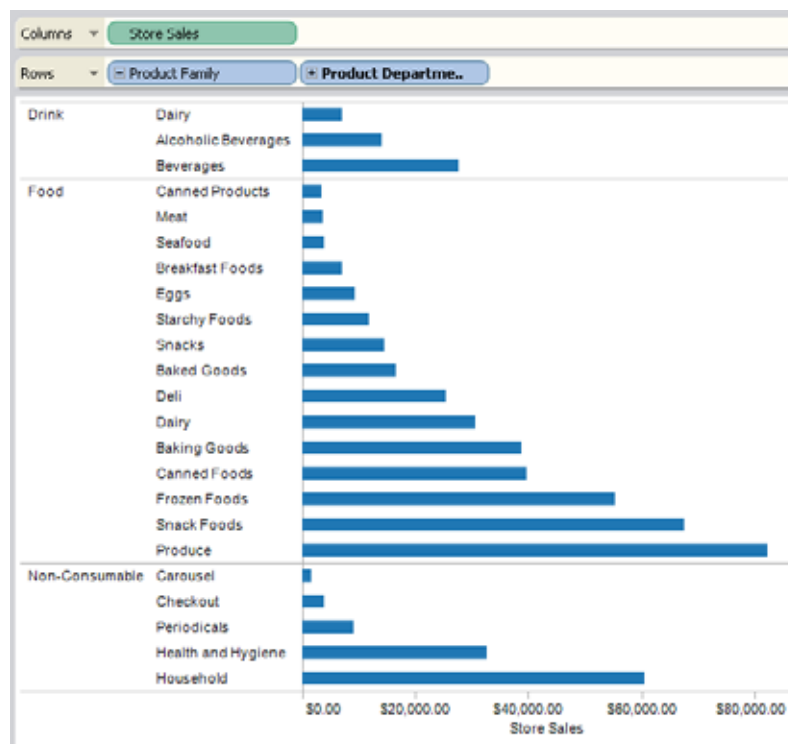
- 2 Sort **Product Department** in ascending order by the **Store Sales** measure.

Right-click on the Product Department field and select **Sort** on the field's context menu. In the Sort dialog box select **Ascending** as the Sort order and sort by **Store Sales**. When finished click **OK**.



The view is shown below. Notice that the **Product Department** members are sorted within each parent member. For example, **Dairy**, **Alcoholic Beverages**, and **Beverages**

are sorted only within the **Drinks** level. Tableau does not rearrange headers that appear before the sorted field.



- 3 If you want to order dimension members without regard to its parent, you should remove the **Product Family** field from the **Rows** shelf.

The sorted data are shown below.



Notice in that the member, **Dairy**, is displayed twice: once for the dairy sold within the Drink category and another for dairy sold within the Food category. The duplication is because the **Product** dimension is set to define unique values by key. All members in a dimension have an individual key that is unique. In this case the member name is Dairy but the keys are different so they are treated as two unique values. For more information about changing how Tableau determines unique values, refer to “Defining Unique Values” on page 10-41.


## Manual Sorting

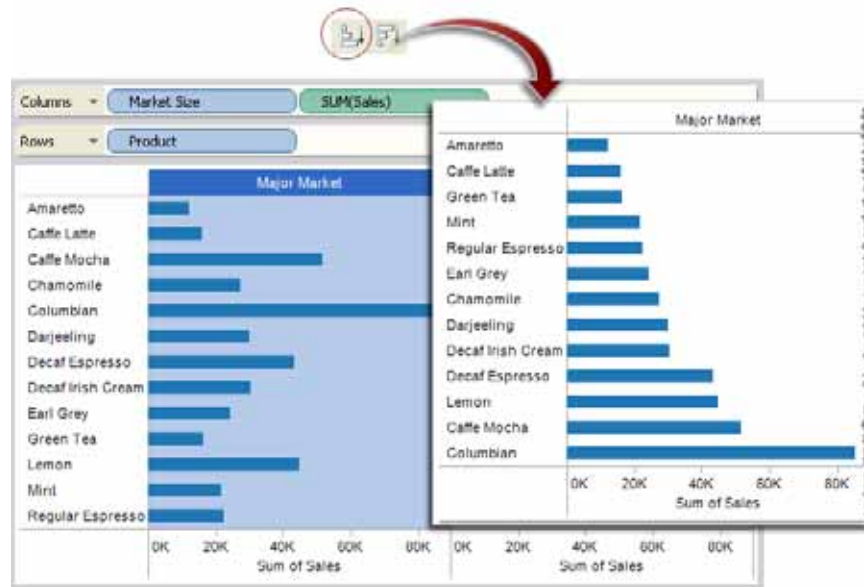
Manual sorting allows you to rearrange the order of dimension members in the table by dragging them in an ad-hoc fashion, giving precise control over how items appear next to one another in tables and in legends. It also gives you control over the order in which data is drawn on the screen. This control is useful when comparing specific pieces of data or interpreting overlapping data. Manual sorts can only be applied to discrete fields including a discrete measure. For information about continuous vs. discrete fields refer to “Data Roles” on page 8-22.

There are two ways to manually sort the data in a view. You can either select items in the view and use the Sort toolbar buttons or you can drag and drop headers in the view. This section discusses the following topics:


- Sorting using the Toolbar
- Sort by Drag and Drop

### Sorting using the Toolbar

The two sort buttons on the toolbar  manually sort a selection either in ascending or descending order based on the other fields in the view. For example, the view below shows sales by product and market size. When you select the Major Market column, thus selecting all of the products, the quick sort buttons sorts the product field by SUM(Sales), which is the measure in the view.




An easy way to anticipate how a selection will be sorted is by using the tool tips. Make a selection in the view and hover over the ascending or descending quick sort toolbar buttons to see a description of how the selection will be sorted.



Using the quick sort buttons creates a manual sort which you can always modify using the sort dialog box. Right-click a sorted field (indicated with bold text) and select **Sort** to open the Sort dialog box.

### Sort by Drag and Drop


- 1 Select the dimension member you want to move. This can be any dimension member that appears in a row or column header of a table, or in a legend like the color legend.
- 2 Drag the member to the desired location within that row, column or legend.



### Example- Manually Sorting Drawing Order

Changing the drawing order of a field allows you to see obscured data in your views in cases where data of one color or shape obscure data of another color or shape. For instance, if you can't see red marks in a scatter plot because they are obscured by green marks, you can change the drawing order so that the red points are drawn on top of the green points (and vice versa).

Change the drawing order of a field by re-arranging the order of dimension members in a legend. For instance, if you want to place red items in front of green items in a view, select the red legend entry and move it higher on the list of items shown in the legend. The marks are drawn in the view according to the order in the legend, from bottom to top. Also you can toggle back and forth between layered field items by dragging any one of the fields from top to bottom or from bottom to top.



Sorting the drawing order is not restricted to color legends. You can reorder shape legends as well. If you have multiple valid legends, the drawing order is defined first by shape, then by color. For example, suppose you have both a shape legend and a color legend. If you have a red circle on top of a green square, moving the green above the red in the color

legend will not necessarily move the green square on top of the red circle. It depends on the order in the shape legend first. If circles are above squares in the shape legend, no amount of reordering the color legend will get that square on top of the circle. Instead, move the square shape above the circle shape first and then reorder the color legend.



## Filtering

Filtering data means specifying specific data to include in the view thus excluding the rest. In Tableau you filtering differs depending on whether you are filtering discrete fields or continuous fields. In addition, there are different types of filters. All fields that are filtered appear on the **Filters** shelf. You can examine this shelf to quickly determine the data that has been removed from the view through a filter.

This section discusses the following topics:

- How Filtering Works
- Filtering Discrete Fields
- Filtering Continuous Fields
- Quick Filters
- External Filters
- Context Filters
- Global Filters


### How Filtering Works

To filter data, specify the dimension members or the range of values that you want to include or exclude from a view. You can filter a field that has already been placed on a shelf, or you can place a field directly on the **Filters** shelf. A field that appears only on the **Filters** shelf is called an external filter. Refer to “External Filters” on page 12-50 for more information. There is no limit for the number of dimensions and measures you can place on the **Filters** shelf.


By default, all filters are treated independently by Tableau. That means that the results generated by independent filters do not depend on the results generated by any other filters. Therefore the order of the filters on the shelf (and the order in which you define the filters) doesn't matter. However, you can create filters that are dependent on one another using context filters. You can think of a context filter as being an independent filter. Any other filters that you set are defined as dependent filters because they process only the data that passes through the context filters.

The following figure shows two dimensions and one measure on the **Filters** shelf. The filters are applied to the view according to the order of the fields placed on the shelf, from top to bottom. Because the filters are independent, changing the order of the fields does not change the results returned by Tableau.

Original order of independent filters.



Changing the order has no effect on results.



Suppose the **Ship Method** filter includes only the Regular Air member, the **Order Priority** filter includes only the Low and Medium members, and the **Sales Total** filter includes only values greater than or equal to 2,500. The data that passes through this set of filters must satisfy all the specified criteria.

The figure below shows two of the records passed by the filter. Notice that the data satisfy all of the filter criteria. Additionally, the order in which the criteria are applied to the data does not matter. Tableau can first filter all records with Regular Air, and then filter all records with Low or Medium order priority, and then filter all records with sales greater than or equal to 2,500 to produce the requested results. If the order in which the filters are applied is changed, Tableau will produce exactly the same result.

	A	B	C	D	E	F
1	Record No	Order Priority	Sales Total	Discount	Ship Method	De
2	1	5 LOW	2302.42	0	REGULAR AIR	
3	2	5 LOW	986.64	0.1	EXPRESS AIR	
4	3	5 LOW	2564.85	0.07	REGULAR AIR	
5	4	5 LOW	226.52	0.12	REGULAR AIR	
6	5	3 MEDIUM	189.98	0.1	REGULAR AIR	
7	6	3 MEDIUM	2679.60	0.06	REGULAR AIR	
8	7	3 MEDIUM	899.31	0.07	EXPRESS AIR	
9	8	3 MEDIUM	56.36	0.04	REGULAR AIR	
10	9	3 MEDIUM	30.42	0.05	REGULAR AIR	

These records pass through the filters.

For information about creating filters that are dependent on one or more other filters refer to “Context Filters” on page 12-55.

## Filtering Discrete Fields

You can filter any discrete field by selecting only the members that you want to include. All other members are excluded. You can specify a filter using one of the following methods:

- Configuring the Filter Dialog Box
- Filter by Selecting Table Headers

---

**Note** When filtering data using any of these methods, you are typically modifying the data view by removing row headers, column headers, encodings, and so on. You are not modifying an existing calculation or performing a new calculation on the underlying data. The exception is when you define an external filter, which can perform a calculation on the underlying data. Refer to “External Filters” on page 12-50 for more information.

---

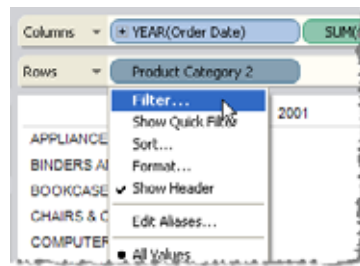
### **Configuring the Filter Dialog Box**

Use the Filter dialog box to specify the records you want to include and exclude from the view. The Filter dialog box contains three tabs: General, Condition, and Top. Use these tabs to create complex conditional rules for the filter.

**To filter dimensions using the Filter dialog box:**

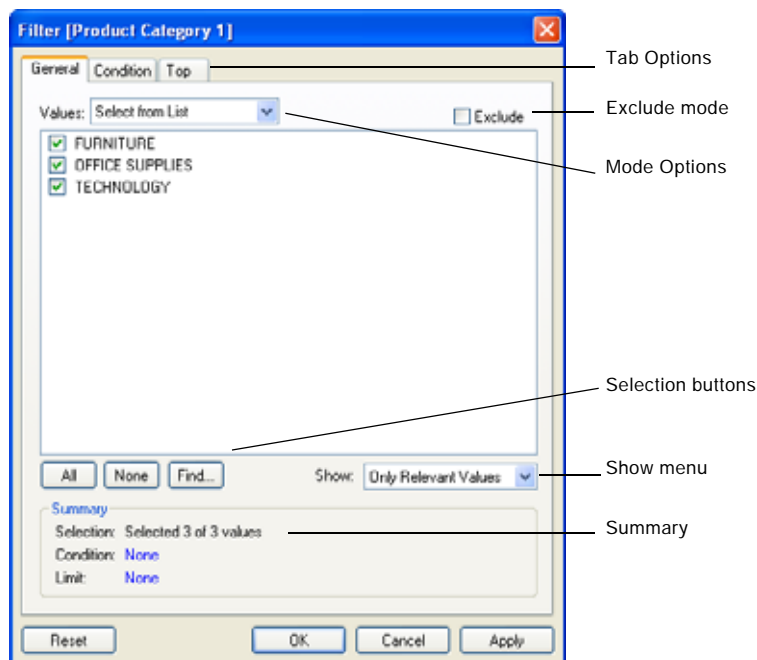
- 1 Open the **Filter** dialog box.

Double-click any discrete field on any shelf or select **Filter** from the field's right-click context menu. You can also place a dimension directly on the **Filters** shelf. The **Filter** dialog box opens automatically.



- 2 Complete the **Filter** dialog box.

The filter dialog box contains three tabs: General, Condition, and Top. Use these three tabs in combination to define formulas and rules that govern the data that is included in your view.



To learn more about each tab in the filter dialog box, refer to the topics listed below:

- Filter Dialog: General Tab
- Filter Dialog: Condition Tab
- Filter Dialog: Top Tab

### **Filter Dialog: General Tab**

Use the **General** tab in the **Filter** dialog box to select the members that you want to include. All deselected members are excluded. Completing the General tab differs depending on whether you are using a Relational or Multidimensional data source. For information on the General tab for a Relational data source refer to “Completing the General Tab: relational data source.” on page 12-24. For information on the General tab for Multidimensional data source refer to “Completing the General Tab: multidimensional data source.” on page 12-24.

**Completing the General Tab: relational data source.** To complete the **General** tab of the **Filter** dialog box select one of the following modes on the Values drop down list:

- **Select from List:** Displays a list of members in the data source that you can select from to include in the filter.

Selected mode is probably the mode you will use most often. By default selected dimension members are included in the filter, while deselected members are excluded. However, you can switch to Exclude mode by selecting **Exclude** in the upper right corner of the dialog box. When in Exclude mode, selected members are excluded while deselected members are included.

- **Type In:** Type explicit dimension member names into a text box to define a filter without querying the database.

Use Type In mode when you are using a large data source and queries are slow. If you know the dimension members you are interested in, you can type them into the text box or copy and paste them from another application. Make sure that each member is on its own line in the text box.


If you are using another mode and the query is taking a long time, you can cancel the query and return to the Filter dialog box in Type In mode. In addition, if you have automatic queries turned off, the filter dialog box will automatically open with Type In mode selected. For more information about automatic queries refer to “Automatic Updates” on page 7-3.

- **Use All:** Includes all of the members in the data source. Sometimes you will want to define a condition or limit filter that is based on all the data, even if that data changes over time, rather than selecting specific members to include or exclude from the filter. The Use All mode always includes every member in the database as the input to the condition or limit.

When you select **All** the rest of the options in the General tab are disabled. Refer to “Filter Dialog: Condition Tab” on page 12-26 and “Filter Dialog: Top Tab” on page 12-27 for information on setting conditions and limits for the filter.


**Completing the General Tab: multidimensional data source.** To complete the **General** tab of the Filter dialog box, select the fields you want to include in the view. The controls at the top of the dialog box are color coded to show the selected levels: the default color means

none are selected, blue means all members on that level are selected, and gray means some members are selected.

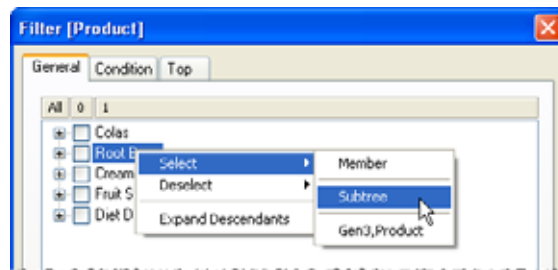


Because hierarchies can be very deep and wide, Tableau provides tools to help you find and select the members of interest. You can select hierarchy members in these four ways:

- **Select individual members** – Navigate through the hierarchy and click on the specific member names you want to include.
- **Use the Filter dialog box buttons** – Use **Expand All**, **Show Selected**, or **Find** to facilitate the navigation and selection process.
- **Select entire levels** – Use the hierarchy level selection controls at the top of the dialog box to include or exclude entire levels.

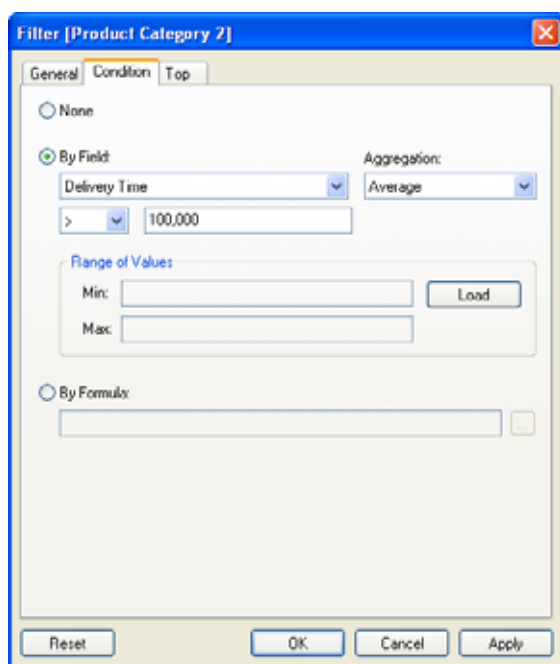


**Use the right-click context menus** – The context menu allows you to select, deselect, expand the selected hierarchy. For example, let's say you wanted to select the sub-tree of the Root Beer, right-click on Root Beer and choose **Subtree** from the **Select** menu.



### **Filter Dialog: Condition Tab**

Use the **Condition** Tab in the **Filter** dialog box to define rules to filter by. For example, In a view showing the average Delivery Time for a collection of products, you may want to only show the Products that have a Delivery Time that is greater than 100,000. In addition to creating field based rules to filter by, you can also write a custom formula



To complete the **Condition** tab select from the following:


- **None:** select this option if you do not want to add a condition to the filter. This is the default setting.
- **By Field:** select this option to specify a condition based on existing fields in the data source. Use the first two drop-down lists to select the field and aggregation you want to base the condition on. Then select a condition operator such as greater than, equal to, etc. Finally, type a criteria value into the text box. For example, to create the condition described above, select Total Sales and SUM from the first two drop-down lists. Then select Greater ( > ) from the operator list and type 100,000 into the text box.



---

**Note** You can use the **Range of Values** box to load the entire range of values for the selected field in the data source. The example above would not make sense if all the records in the data source for Total Sales were greater than 100,000 to begin with. Using the Range of Values box helps you decide a value that makes sense to the records in your data source. Click **Load** to view the range of values for the selected field.

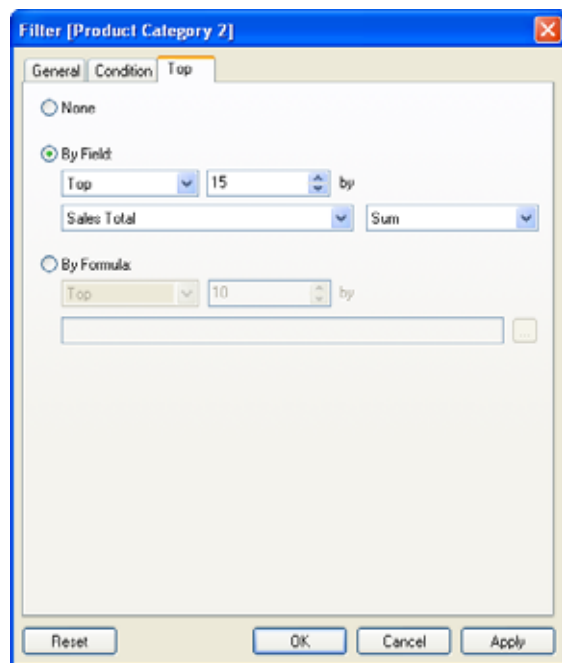
---

- **By Formula:** select this option for more advanced filter conditions. You can type a custom formula into the text box or open the formula editing dialog box by clicking the  button to the right of the text box.


For more information on writing formulas and functions refer to Chapter 26, “Functions, Operators, & Data Types”.

#### **Filter Dialog: Top Tab**

Use the **Top** tab in the **Filter** dialog box to define a formula that computes the data that will be included in the view. For example, in the same view discussed above that shows the average Delivery Time for a collection of products, you can decide to only show the Top 15 Products in terms of Sales. Rather than have to define a specific range for Sales (greater than 100,000) you can define a limit that is relative to the other members in the field. The formula defined on the Top tab is evaluated on the results of the formula on the Condition tab.




To complete the **Top** tab select from the following options:

- **None:** select this option if you do not want to add a limit to the filter. This is the default setting.
- **By Field:** select this option to add a simple limit based on an existing field in the data source. First select the limit range using the first two drop-down lists. For example you can select Top 10 or Bottom 20. Finally select the field and aggregation to base the limit on. So if you wanted to filter based on the Top 10 Sales, select Top and 10 from the first two drop-down lists and then select Sales and SUM from remaining lists.
- **By Formula:** select this option for more advanced filter limits. Select the limit range using the first two drop-down lists (e.g. Top 10 or Bottom 20). Then you can type a custom formula into the text box or open the formula editing dialog box by clicking the  button to the right of the text box. For more information on writing formulas and functions refer to Chapter 26, "Functions, Operators, & Data Types".

### Showing and Hiding Values in the Filter Dialog Box

The contents of the Filter dialog box is affected by the filters that are already set in the view. For example, if you create a filter on the **Market** category that only includes the western region, when you open the **States** filter dialog box, you typically only want to see the relevant values (the western states). If you do want to see all the values in the data source including the ones that don't pass the current filters, you can select **All values in Database** option in the **Show** list. The **Show** list includes and excludes data from displaying in the Filter dialog box so you can find what you are looking for quicker.



You can also choose to view **All values in Context** which means you will see the values that pass the current context filters. That option is only available if there are context filters set in the view. Refer to “How to Create a Context Filter” on page 12-56 for more information about context filters.

Using the **Show** menu in the filter dialog box you can select from the options described in the table below.

Menu Selection	Definition
<b>Only Relevant Values</b>	Only displays members (or the range of values) that would pass all currently set filters.
<b>All Values in Context</b>	Only displays members (or the range of values) that would pass all current set context filters and would ignore other filters.
<b>All Values in Database</b>	Ignores all filters and shows every value that occurs in the database

### Filter by Selecting Table Headers

You can filter discrete fields by selecting one or more headers in a table, and then selecting **Keep Only** or **Exclude** from the right-click context menu. You can also double-click on a single header to keep it and exclude all others.

When you select a table header that is part of a hierarchy, all of the next level headers are also selected automatically. For example, the view shown below consists of two unrelated dimensions placed on the **Columns** shelf, and two levels of the same hierarchy placed on the **Rows** shelf.

The selected row headers include the CA member of the **Store State** dimension, and the Portland and Seattle members of the **Store City** dimension. Note that when CA is selected, all members from the next (inner) level in the hierarchy are automatically selected.

The selected column headers include the Q1 and Q2 members of the **Quarter** dimension. Note that when Q1 and Q2 are selected, the inner dimension members (Drink and Food) are not automatically selected. This is because Drink and Food are members of the **Product Family** dimension, which is unrelated to the **Quarter** dimension.

Select headers from different hierarchy levels.

Select headers from the same outer dimension.

Columns		Rows							
+ Quarter		+ Product Family							
+ State/State		+ Store/City							

Selecting **Keep Only** keeps all selected headers as shown below. The view displays data for California, Portland, and Seattle for Drinks and Foods in Q1 and Q2. The view displays data for the states in the East and Southern region

Columns	Quarter	Product Family
Rows	Store State	Store City

Selecting **Exclude** excludes all selected headers. For example, suppose you want to keep or exclude data based on product family. To do this, select any single header of interest (Drink or Food). For example, the following view has the Drink header selected under the Q1 column. Selecting **Keep Only** keeps all headers labeled Drink and all outer headers.

Select headers from different hierarchy levels.

Select headers from the same inner dimension.

Columns		Rows			
Quarter		Store State			
Product Family		Store City			
		Q1		Q2	
		Drink	Food	Drink	Food
CA	Beverly Hills				
	Los Angeles	302.71		333.75	2,836.52
	San Diego			288.98	1,456.45
	San Francisco	45.54			
OR	Portland	350.78		375.57	2,251.08
	Salem	280.70	2,346.82	292.48	2,331.66
WA	Bellingham				
	Bremerton				
	Seattle	305.92	2,600.79	248.19	2,614.11
	Spokane				
	Tacoma	172.77	2,512.58	195.99	2,748.27
	Walla Walla	81.66	822.89		
	Yakima	256.42	1,682.85	203.27	1,404.17

### Example – Filtering Dimensions

This example filters headers and color encodings in a bar chart using the **Filter** dialog box. To filter the data, follow the steps below.

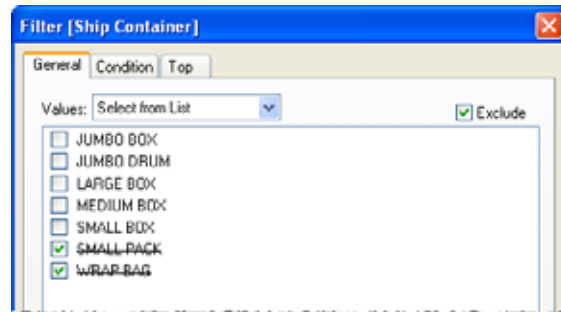
## 1 Create the data view.

The initial data view is shown below. It was created using the Superstore Sales Excel data source. The view shows the average regional delivery time for each product based on the ship container and ship method.

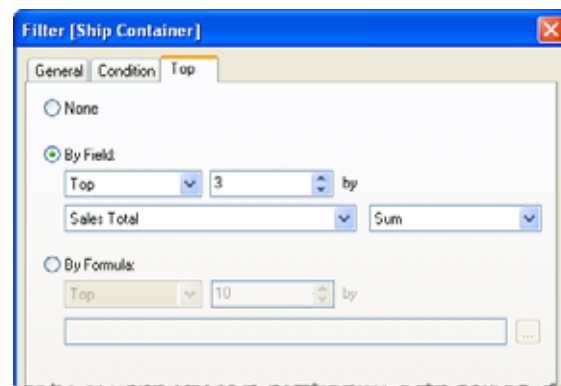


## 2 Filter the Ship Container dimension to exclude Small Pack and Wrap Bag and include the Top 3 remaining containers based on SUM(Sales).

Right-click the **Ship Container** dimension and select **Filter**. In the Filter dialog box, select the **Exclude** option in the upper right corner. Select **Small Pack** and **Wrap Bag** in the General tab as the members to exclude.



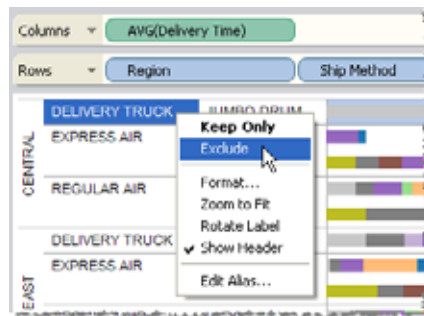
Switch to the Top tab and select **By Field**. Select Top 3 from first two drop-down lists. Then select **Sales Total** and **SUM** from the remaining drop-down lists. When finished click **OK**.



- 3 Filter the Ship Method dimension to exclude Delivery Truck.

Right-click on the **Delivery Truck** row header and select **Exclude**. The Delivery Truck ship method is removed from each region in the view.






#### 4 Filter the color-encoding dimension.

Open the **Filter** dialog box for the **Product Category 2** dimension. Deselect the Computer Peripherals, Office Machines, and Telephones and Communication members from **Product Category 2**.

The final view is shown below. By examining the **Filters** shelf, you can easily see that the view includes three filters. Because **Ship Container** and **Ship Method** appear on both the **Rows** and **Filters** shelves, you know that row headers have been removed from the view. Because **Product Category 2** appears on both the **Color** and **Filters** shelves, you know that color-encodings have been removed. To determine which members have been excluded, open the **Filter** dialog box for each of these fields.



## Filtering Continuous Fields

You can filter continuous fields, such as numeric measures, by selecting a range of values that you want to include. All other values are excluded. Typically, the continuous fields you filter will be numeric measures, however, continuous date and time dimensions are another example of continuous fields. Use the Filter dialog box to define a filter for a continuous field.

This section discusses the following topics:

- Configuring the Filter dialog box: continuous fields
- Example – Filtering Continuous Fields
- Example- Filtering Continuous Dates


**Note** If you have a large data source, filtering continuous fields can lead to a significant degradation in performance. It is sometimes much more efficient to filter by creating a set containing the continuous fields and then applying a filter to the set. Refer to “How to Create a Set” on page 12-72. It is also efficient to filter a discrete dimension using a condition based on the continuous field. Refer to “Filtering Discrete Fields” on page 12-20.

### Configuring the Filter dialog box: continuous fields

To filter continuous fields using the Filter dialog box:

- 1 Open the **Filter** dialog box.


Select any continuous field on any shelf, and then select **Filter** from the context menu. You can also place a continuous field directly on the **Filters** shelf. The **Filter Field** dialog box opens.



**Note** If you place a field directly on the filter shelf, you need to specify an aggregation before the filter dialog box will open. Another dialog box opens first where you can select an aggregation for the field you are filtering by. Fields on the filter shelf can be either aggregated or disaggregated regardless of the other continuous fields in the worksheet.

## 2 Complete the **Filter** dialog box.

The **Filter** dialog box specifies the minimum and maximum values of the range to include in the view. Values inside the range are included in the view, while values outside the range are excluded (filtered). Follow the steps below to complete this dialog box.



### Step 1: Specify the values to include.

You can specify the kind of values you want to include when you are defining the filter using the drop-down list at the top of the dialog box. For example, you can choose to include all values in the specified range, non-null values only, and all values in the data source. Select from the following options:

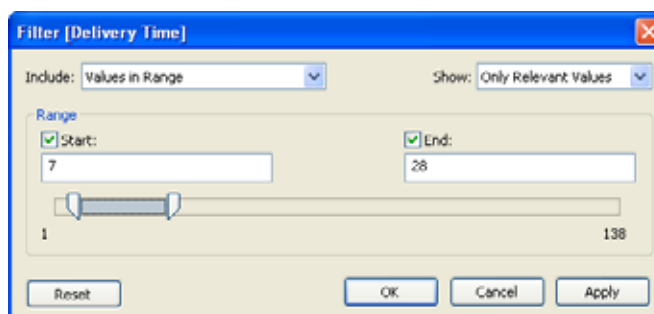
Values in Range	Defines a filter based on the values in the specified range. This is the most common case.
Values in Range and Null Values	Defines a filter based on the values in the specified range but also includes null values.
Null Values Only	Defines a filter that only includes the null values.

Non-Null Values Only	Defines a filter that includes all values that are not null.
All Values	Defines a filter that includes all values.

**Step2: Define a range.**

If you chose Values in Range or Values in Range and Null Values, you must now define the range. You can either use the text boxes or the slider to define this range. You can create a filter with an open-ended range, for example greater than 100,000, by deselecting the minimum or maximum check box. The text box is then disabled and the sliders adjust to only show the remaining slider handle.

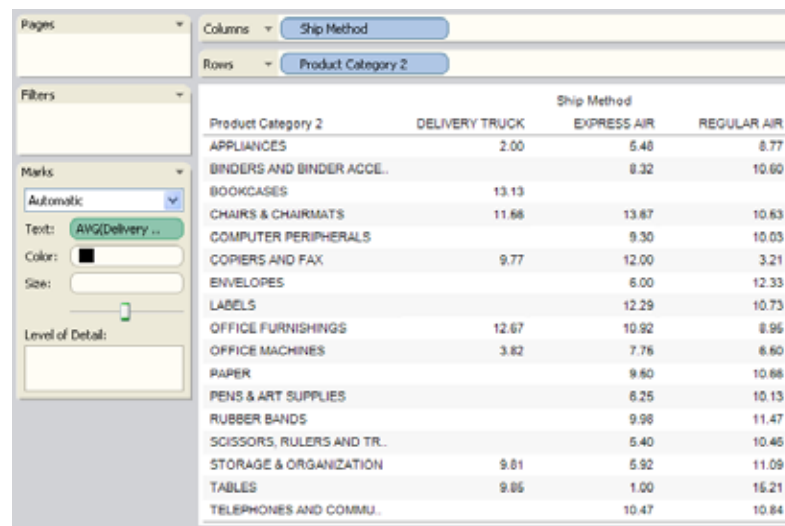
The following dialog box is configured to exclude the data that falls outside a specified range of delivery times. This dialog box is an example for numeric measures. A slightly different filter dialog box opens for continuous date and time fields.

**Example – Filtering Continuous Fields**

This example filters a text table using an aggregated measure, and then filters the table using the same measure in a unaggregated state. The steps are as follows:

## 1 Create the view.

The text table is shown below. It was created using the Superstore Sales Excel data source. The numbers are given by the **Delivery Time** measure aggregated as an average.



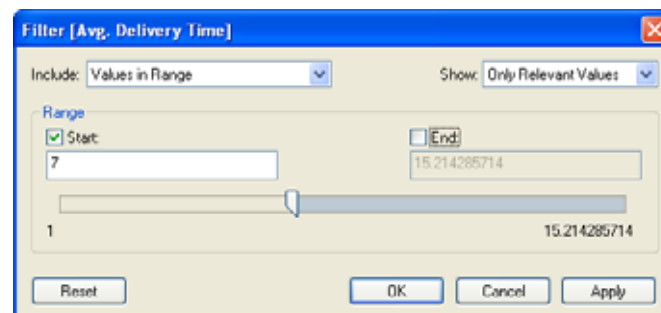
Product Category 2	Ship Method		
	DELIVERY TRUCK	EXPRESS AIR	REGULAR AIR
APPLIANCES	2.00	5.48	8.77
BINDERS AND BINDER ACCE..		8.32	10.60
BOOKCASES	13.13		
CHAIRS & CHAIRMATS	11.66	13.87	10.63
COMPUTER PERIPHERALS		9.30	10.03
COPIERS AND FAX	9.77	12.00	3.21
ENVELOPES		6.00	12.33
LABELS		12.29	10.73
OFFICE FURNISHINGS	12.67	10.92	8.95
OFFICE MACHINES	3.82	7.76	6.60
PAPER		9.60	10.68
PENS & ART SUPPLIES		6.25	10.13
RUBBER BANDS		9.98	11.47
SCISSORS, RULERS AND TR..		5.40	10.45
STORAGE & ORGANIZATION	9.81	5.92	11.09
TABLES	9.85	1.00	16.21
TELEPHONES AND COMMU..		10.47	10.84

## 2 Filter the data using an aggregated measure.

Suppose you want to filter by the average delivery time. To do this, open the **Filter** dialog box by selecting **Filter** from the field's context menu. You can also place the



**Delivery Time** measure directly on the **Filters** shelf and select the average from the **Filter Field** dialog box.

The **Filter** dialog box is shown below. It displays the limits of the aggregation for the selected measure. Specify a new lower limit of 7 and un-check the End checkbox so that the filter includes all values greater than 7.




The modified data view is shown below. Comparing this view with the original, unfiltered view is straightforward because the measure and the filter use the same aggregation. For example, Appliances shipped by Delivery Truck and Express Air are removed from the view because their average delivery time is less than 7 days, while

Appliances shipped by Regular Air remains in the view because its delivery time is greater than 7 days.

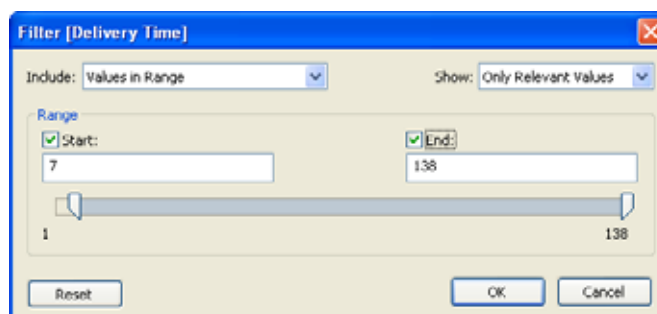
Pages	Columns	Ship Method		
	Rows	Product Category 2		
Filters		Ship Method		
AVG(Delivery Time)		Product Category 2	DELIVERY TRUCK	EXPRESS AIR
				REGULAR AIR
Marks		APPLIANCES		8.768
Automatic		BINDERS AND BINDER ACCE...	0.317	10.606
Text: AVG(Delivery ...		BOOKCASES	13.133	
Color: 		CHAIRS & CHAIRMATS	11.680	13.867
Size: 		COMPUTER PERIPHERALS		9.300
Level of Detail:		COPIERS AND FAX	9.769	12.000
		ENVELOPES		12.333
		LABELS		12.289
		OFFICE FURNISHINGS	12.667	10.922
		OFFICE MACHINES		7.760
		PAPER		9.504
		PENS & ART SUPPLIES		10.130
		RUBBER BANDS		9.976
		SCISSORS, RULERS AND TR...		10.457
		STORAGE & ORGANIZATION	9.806	
		TABLES	9.046	
		TELEPHONES AND COMMU...		10.472
				10.837

3 Filter the data using an unaggregated measure.

Now suppose you want to filter the view using the unaggregated **Delivery Time** measure. To do this, select **All Values** from the context menu of the AVG(Delivery Time) field on the Filters shelf.




The **Filter** dialog box is shown below. It displays the limits of the individual rows for the Delivery Time measure. Specify a new lower limit of 7.



The filtered data view is shown below. Notice that the numbers are very different from the original, unfiltered view. This is because Tableau excludes each row in the data



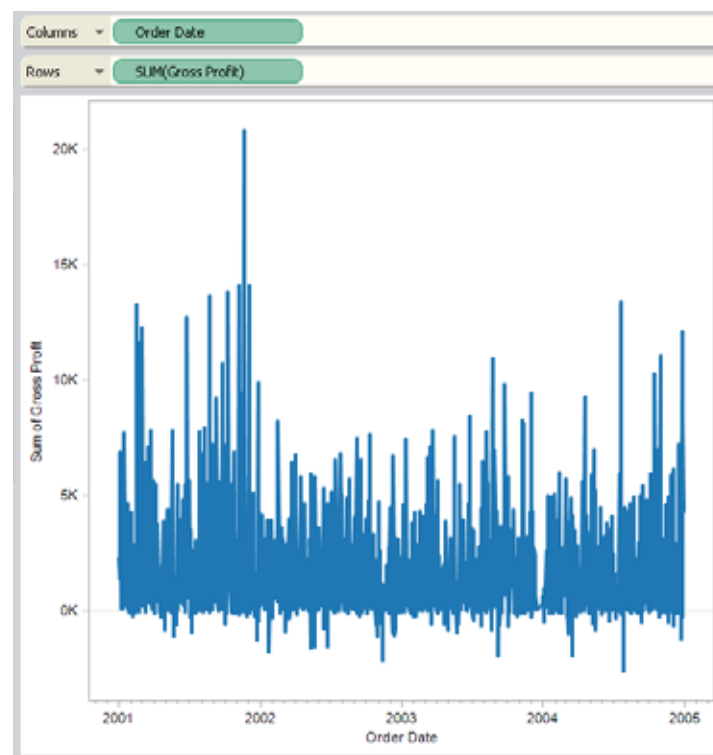
Pages	Columns	Ship Method			
	Rows	Product Category 2			
Filters			Ship Method		
Delivery Time		Product Category 2	DELIVERY TRUCK	EXPRESS AIR	REGULAR AIR
AVG(Delivery Time)		APPLIANCES		13.17	18.09
		BINDERS AND BINDER ACCE..		17.80	22.77
Marks		BOOKCASES	16.22		
Automatic		CHAIRS & CHAIRMATS	17.84	19.50	25.18
Text: AVG(Delivery ..		COMPUTER PERIPHERALS		20.80	20.72
Color: 		COPIERS AND FAX	16.14	12.00	16.00
Size: <input type="text"/>		ENVELOPES		12.00	26.48
		LABELS		29.24	22.40
Level of Detail:		OFFICE FURNISHINGS	18.00	26.11	19.48
		OFFICE MACHINES	12.00	13.50	13.09
		PAPER		22.07	21.71
		PENS & ART SUPPLIES		14.29	20.50
		RUBBER BANDS		20.00	23.40
		SCISSORS, RULERS AND TR..		12.75	24.47
		STORAGE & ORGANIZATION	14.11	14.50	24.48
		TABLES	17.20		29.54
		TELEPHONES AND COMMU..		24.50	21.50

### Example- Filtering Continuous Dates

This example filters a line graph, to show the gross profit over a specific range of time. The steps are as follows:

- 1 Create the view.

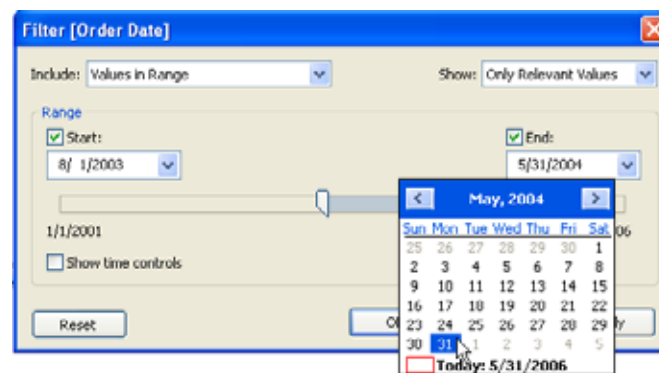
The line graph is shown below. It was created using the Superstore Sales Excel data source. Place **Order Date** on to the **Columns** shelf and select All Values as the aggregation. Then place **Gross Profit** onto the **Rows** shelf.



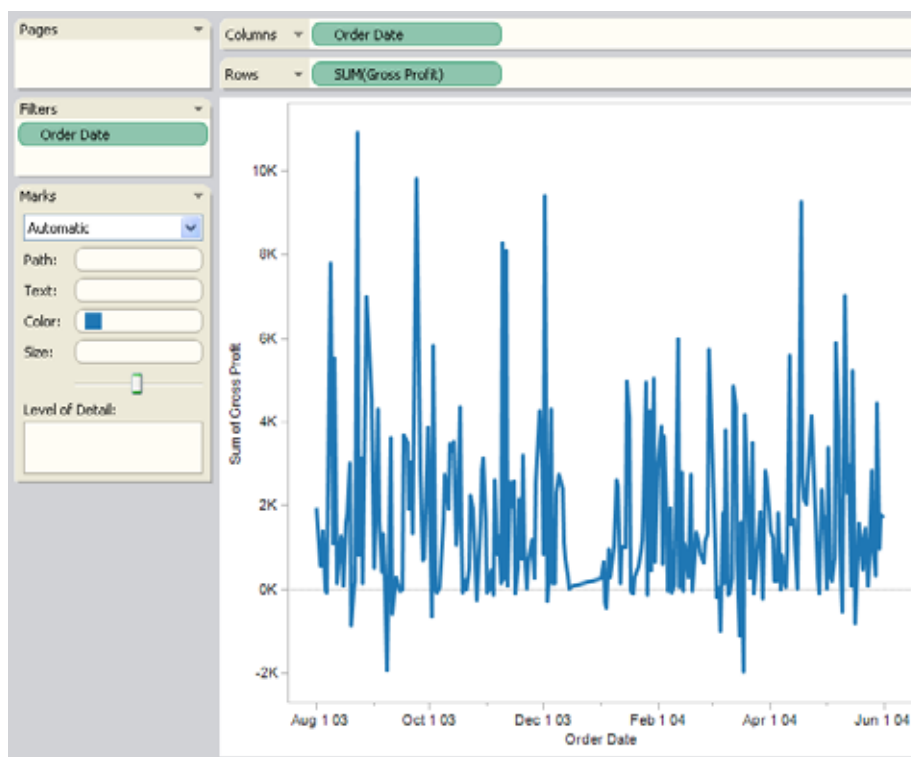
## 2 Filter the data to include a specific range of time.

Suppose you want to filter by the Order Date to show values from August 2, 2003 until May 1, 2004. To do this, open the **Filter** dialog box by selecting **Filter** from the **Order Date** context menu.

The **Filter** dialog box is shown below. It displays the Order Date limits. Use the drop-down calendars to specify a new lower limit of August 2, 2003 and an upper limit of May 1, 2004.

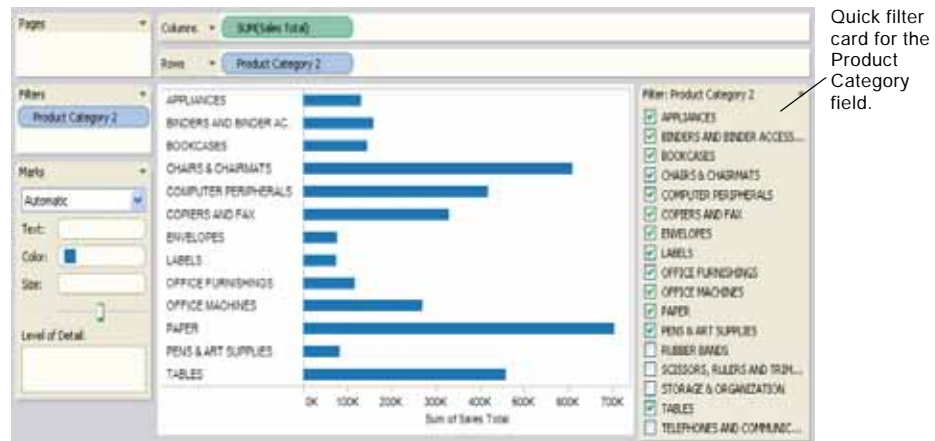


The modified data view is shown below.




## Quick Filters

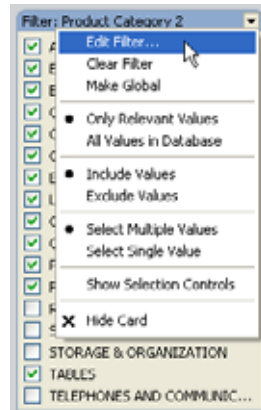
Tableau lets you quickly add and modify filters through its Quick Filters feature. When you turn on a Quick Filter, a smaller representation of the Filter dialog box opens as a new card. From there you can quickly decide what to include in the view.



A Quick Filter can be turned on for existing filters or for non-filtered fields. To show or hide a quick filter, select **Quick Filter** from the field's context menu.



From the Quick Filter card, you can open the Filter dialog box for more filtering options and to add conditions and limits by selecting **Edit Filter** on the card's drop down menu.




Also in the Quick Filter card's drop-down menu you can control the context of the filter. The data considered for each filter is affected by the existing filters in the view. Using the Quick Filter card's drop-down menu you can select the values in the view that the filter will affect. Select **Only Relevant Values**, **All Values in Context**, or **All Values in Database**. For more information about selecting a context for a filter refer to "Showing and Hiding Values in the Filter Dialog Box" on page 12-29.


To turn a quick filter off, select **Hide Card** on the Quick Filter card's drop-down menu or select **Show Quick Filter** on the field's context menu again.

You can show quick filters for both discrete fields and continuous fields. Use the checkboxes to include and exclude discrete dimension members and use the slider control the included values of a continuous measure.

Discrete Quick Filter




Continuous Quick Filter



### Quick Filters: Discrete Fields

For discrete fields the Quick Filter card shows a list of members that you can select or deselect to include or exclude from the view. There are two kinds of selection modes: Select Multiple Values or Select Single Values. Use the card menu to switch between the two different selection modes.

**Select Multiple Values Mode.** This is the default mode, where you can select and deselect multiple members in the field to include and exclude them from the view.



**Select Single Values Mode.** Use this mode if you want to select a single member at a time. This mode is useful when you want to see how each individual member affects the view.




---

**Note** If you are using a multidimensional data source, the Quick Filter card does not show the entire hierarchy. Rather, it shows a sample of the members in a dimension. Use the Filter Dialog box to view all the dimensions in a hierarchy.

---

### Quick Filters: Continuous Fields


For continuous fields the Quick Filter shows the filtered values as a pair of sliders that you can adjust to include or exclude more values. Click on the upper and lower limit readouts to enter the values manually.



Sometimes you may want to filter using an open-ended range as opposed to having fixed start and end values. For example you may want to include all values greater than 5000. To use an open-ended range, deselect the fixed value options on the card menu. These options are described below.

- Fixed start value – Deselect this option if you want to filter using an open-ended start value. This means that the filter will include all values less than the specified end value.
- Fixed end value – Deselect this option if you want to filter using an open-ended end value. This means that the filter will include all values greater than the specified start value.

The highlighted area is called the data bar. It indicates the range in which data points actually lie in the view. Use this indicator to determine a filter that makes sense for the data in your data source. For example, you may filter the Gross Profit field to only include values between \$100,000 and \$500,000 but your view only contains values between \$120,000 and \$150,000. The range of data you can see in the view is indicated by the data bar while the sliders show you the range of the filter.



## External Filters

The filters discussed so far apply only to the data that make up the view. With these types of filters, fields that appear on the **Filters** shelf also appear somewhere else in the view.

An External filter uses fields that do not contribute headers or axes to a view. For example, suppose your data view displays the total expenses as a function of product type, and you want to filter the data by sales and by state. You can do this by placing the **Sales** and **State** fields directly on the **Filters** shelf. These fields do not contribute to the rows or columns of



the view. Additionally, when a discrete field is specified as an external filter, Tableau automatically performs a calculation based on the members selected in the filter.

### How to Create an External Filter

Create an external filter by placing a dimension or a measure directly on the **Filters** shelf. You can also create an external filter using a set. For relational data sources, the set can contain members from any dimensions. For multidimensional data sources, the set cannot contain levels from different hierarchies.

---

**Note** Microsoft Analysis Services 2005 data sources generally do not support external filters created by levels and attributes of dimension that already exists in the view. The exception is if you limit the filter to a single member. For example, you cannot place the country level of a customer dimension on a shelf and then add an external filter for several specific customer cities. However, you can use an external filter if you select a single city rather than several.

---

#### To create an external filter with a discrete field:

- 3 Place a discrete field directly on the **Filters** shelf, the filter dialog box opens.
- 4 Using **Selected** mode in the **General** tab of the dialog box, select the members you want to include in the filter. If multiple members are selected, then a calculation is performed on the data. For relational data sources, the calculation can be any one of the aggregation functions. For multidimensional data sources, the calculation is always a summation. Refer to “Calculation Filter” on page 12-52 for more information.

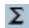
#### To create an external filter with a continuous field:

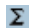
- 1 Place a measure directly on the **Filters** shelf, the **Filter Field** dialog box opens.
- 2 Select an aggregation for the measure or accept the default selection and disaggregate the data. For multidimensional data sources, this dialog box does not open because the aggregations cannot be changed in Tableau.
- 3 Use the Filter dialog box to specify a range of values to include.

## Calculation Filter

For discrete fields, an external filter is called a *calculation filter* because Tableau can perform a calculation on the selected dimension members. This occurs when:

- The dimension appears only on the **Filters** shelf.
- You select more than one member from the dimension.

The calculation icon  displays next to the field's name to indicate this operation.

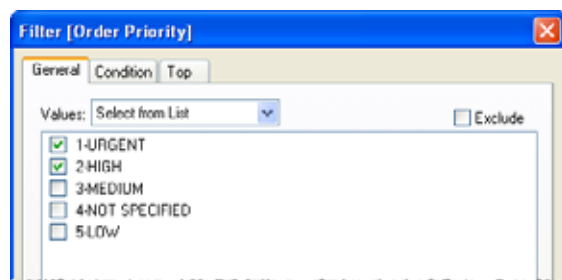
The calculation that is performed depends on the data source. For relational data sources, the calculation matches the aggregation for each measure used in the view. For multidimensional data sources, the calculation is always a summation and the  icon is used to indicate this operation.

This section discusses the following topics:

- Calculation filters: Relational Data Source
- Calculation Filters: Multidimensional Data Source

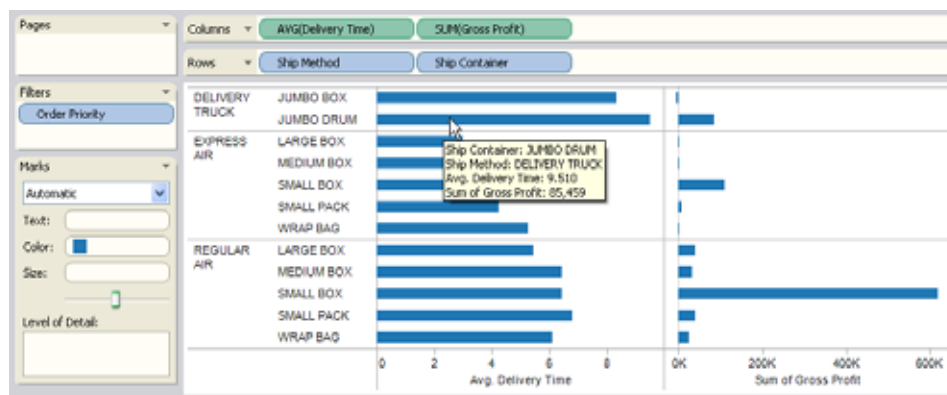
### Calculation filters: Relational Data Source

Consider the view shown below. It consists of the **Gross Profit** measure aggregated as a summation and the **Delivery Time** measure aggregated as an average. These measures are displayed with the **Ship Method** and **Ship Container** dimensions. An external filter that consists of two members of the **Order Priority** dimension is applied to the data.



Because the data source is relational, Tableau automatically applies the appropriate calculation to the members of the external filter based on the aggregation of each measure. Therefore, a summation is performed for **Gross Profit** and an average is performed for **Delivery Time**.

For example, the tooltip shows the data for Jumbo Drums delivered by truck. The average delivery time is 9.510 days. This number was calculated by averaging the delivery times for all the rows that have an Urgent or High order priority. Similarly, the sum of the gross profit is \$85,459. This number was calculated by summing the gross profit for all the rows that have an Urgent or High order priority.

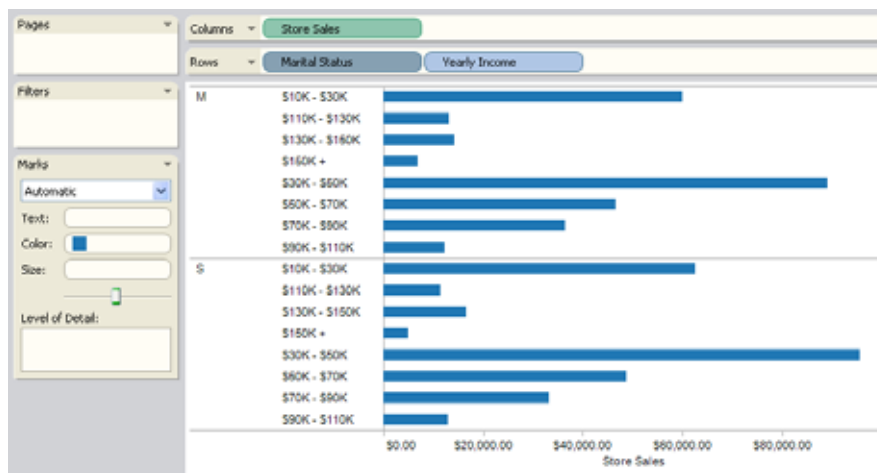


### Calculation Filters: Multidimensional Data Source


You can think of a calculation filter as “slicing” the cube’s data through the specified dimensions. Therefore, a calculation filter for a multidimensional data source is sometimes referred to as a slicer.

Because the aggregation for this arbitrary slice was not defined when the cube was created, Tableau automatically performs a summation. Fortunately, measures are usually aggregated as a summation. Therefore, applying a calculation filter produces a sum of a group of sums, which is a calculation that is useful and easy to interpret.

For example, consider the view shown below. It consists of the **Store Sales** measure displayed with the **Marital Status** and **Yearly Income** dimensions.

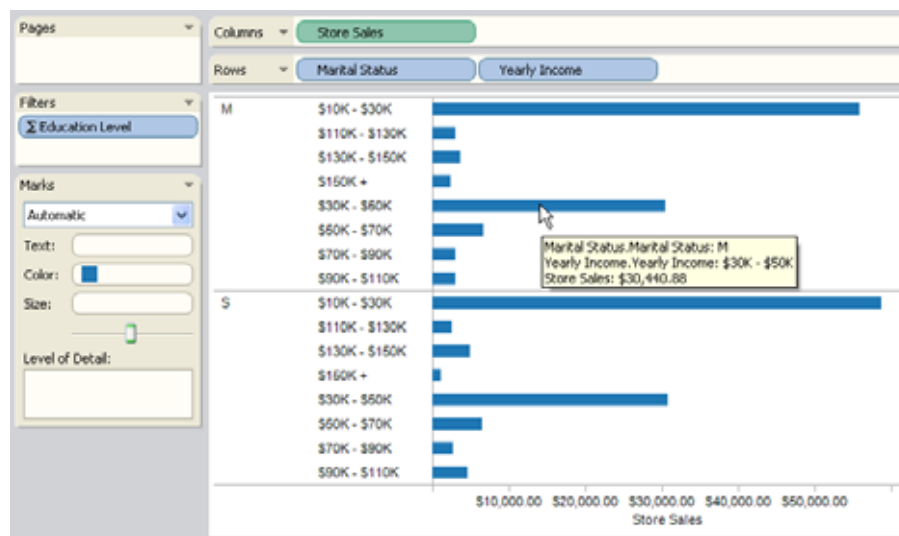


An external filter that consists of two members of the **Education Level** dimension is applied to the data.



Because the data source is multidimensional, Tableau automatically sums the members of the external filter. If **Store Sales** was aggregated as a summation when the cube was created, Tableau would sum a group of sums. If **Store Sales** was aggregated using a different function such as a percentage, Tableau would produce a summation of percentages, which is probably not a calculation that is useful to you.

For example, in this example store sales is \$30,441 for married customers with an annual income between \$30,000 and \$50,000. This number was calculated by summing **Store Sales** for the Partial College and Partial High School members.



## Context Filters

By default, all filters that you set in Tableau are computed independently. That is, each filter accesses all rows in your data source without regard to other filters. However, you can set one or more filters as *context filters* for the view. You can think of a context filter as being an independent filter. Any other filters that you set are defined as dependent filters because they process only the data that passes through the context filter.

You may create a context filter to:

- Improve performance – If you set a lot of filters or have a large data source, the queries can be slow. For relational data sources, you can set one or more context filters to improve performance. Refer to “How to Create a Context Filter” on page 12-56.
- Create a dependent numerical or top N filter – You can set a context filter to include only the data of interest, and then set a numerical or a top N filter and process only the data that passes through the context filter.

For example, suppose you’re in charge of breakfast products for a large grocery chain. Your task is to find the top 10 breakfast products by profitability for all stores. If the data source is very large, you can set a context filter to include only the products in your department. You can then set a top 10 filter by profit as a dependent filter, which would process only the data that passes through the context filter.

Context filters are particularly useful for relational data sources because a temporary table is created. This table is automatically generated by Tableau when you set the context, and acts as a separate (smaller) data source that results in increased performance when you build data views.

---


**Note** For Excel, Access, and text data sources, the temporary table is created as an Access table. For SQL Server, MySQL, and Oracle data sources, you must have permission to create a temporary table on the server.

---


For multidimensional data sources, temporary tables are not created and context filters only define which filters are independent and dependent.

### How to Create a Context Filter

To create a context filter, select **Add to Context** from the context menu of an existing categorical filter. Alternatively, you can select the **Analysis > Set Context** menu item. The context is computed once to generate the view. All other filters are then computed relative to the context. Context filters:

- Appear at the top of the **Filters** shelf.
- Are identified by the color grey and the  icon.
- Cannot be rearranged on the shelf.

As shown below, the **Product** dimension is set to be the context for a data view. The **Customers** filter is computed using only the data that passes through **Product**.



You can modify a context filter by:

- Removing the field from the **Filters** shelf – If other context filters remain on the shelf, a new context is computed.


- Editing the filter – A new context is computed each time you edit a context filter.
- Selecting **Remove from Context** – The filter remains on the shelf as a standard categorical filter. If other context filters remain on the shelf, a new context is computed.

### **Example – Context Filters**

This example filters a scatter plot using two filters configured with top N limits. One filter is set as a context filter and the other is a dependent filter. The steps are as follows:

## 1 Create the view.

A scatter plot is shown below. It was created using the Superstore Sales Excel data source, and displays the disaggregated **Sales Total** and **Gross Profit** measures color-encoded by the **Product Category 2** dimension.




## 2 Apply Top N filters.

The first filter is applied to the **Product Category 2** dimension. Place this dimension on the Filter shelf. In the filter dialog box, select the Top tab and specify a limit by field that includes the top 10 products by average delivery time. The second filter is applied to the **Customer Name** dimension. Place this dimension on the Filter shelf. In the filter dialog box, select the Top tab and specify a limit by field that includes the top 10 customers by average sales. For more information on the Top tab refer to “Filter Dialog: Top Tab” on page 12-27.

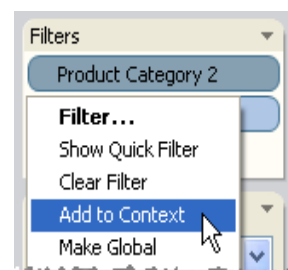


Both filters are independent filters and produce the view shown below.




### 3 Set the context filter.


Set **Product Category 2** as the context for **Customer Name** by selecting **Add to Context** from the **Product Category 2** field menu. The two filters are no longer independent of each other. Instead, **Customer Name** depends on the data that passes through **Product Category 2**.



Because the data source is relational, the progress bar indicates that a temporary Access table is being created based on the context.



The final view is shown below. The context filter removed some of the data points that did not pass through the filter.



## Global Filters

A global filter is a filter that applies to all worksheets in the workbook that are connected to the same data source. For example, you may have a filter that only includes a specific region or product of interest. Rather than adding this filter every time you create a new worksheet, you can simply create the filter once and then make it global.

### To make a global filter:

- Right-click an existing filter on the filter shelf and select **Make Global**.




The field is marked with a globe icon and the filter is applied to all worksheets in the workbook. Additionally, the filter is automatically added to any new worksheet you create. Any changes you make to the filter affects all of the worksheets.

At anytime you can make a global filter local again. When you make a filter local, the filter remains on all the worksheets, however, they are no longer tied together and can be deleted or modified on an individual basis.

### To make a global filter local:

- Right-click on the global filter on the filter shelf and select **Make Local**.



The globe icon is removed and the filter can once again be modified individually per worksheet.

## Ad-Hoc Groups

An ad-hoc group is a group of dimension members that have been combined into higher level categories. For example, if you are working with a view that shows average test scores by major, you may want to group certain majors together to create major categories. English and History may be combined into a group called Liberal Arts Majors while Biology and Physics may be grouped as Science Majors. This section discusses the following topics:

- Creating Ad-Hoc Groups
- Editing an Existing Group
- Finding Members in the Ad-Hoc Group Dialog Box

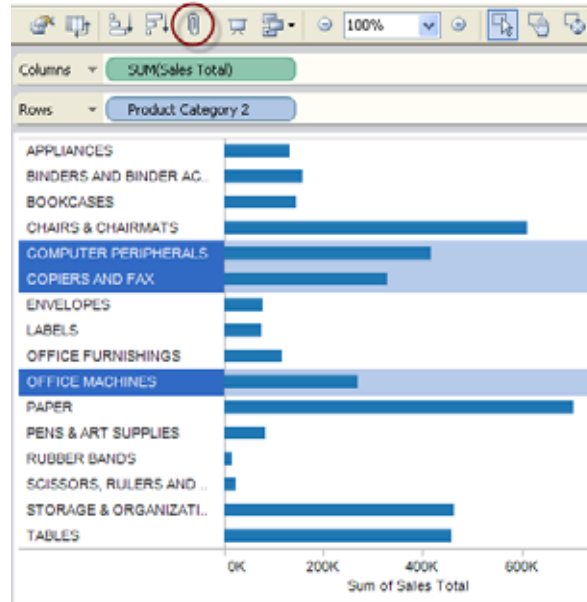
### Creating Ad-Hoc Groups

The most common way to create an ad-hoc group is through the group button on the toolbar. However, you can also create ad-hoc groups by right-clicking a dimension in the Data window and selecting Create Ad-Hoc Group.

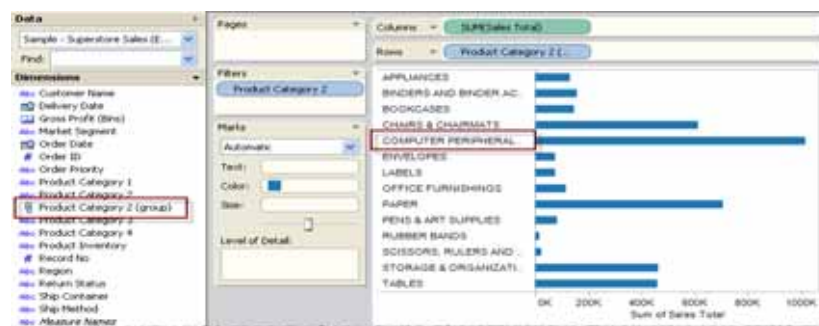
**To create an ad-hoc group using the toolbar:**

- 1 Hold the **CTRL** key on the keyboard to multi-select dimension members in the view.

- 2 Click the **Group** button  on the toolbar.



The selected members are combined into a single member and a new grouped field is added to the Data window. A default member name is automatically constructed using the combined member names.



You can use the grouped field just like any other field in the view, except the grouped

field is cannot be used to create calculated fields.

You can add to or remove members from a group by right-clicking the grouped field in the Data window and selecting **Edit**. In the Edit Ad-Hoc Group dialog box you can also change the default name of the group and combine fields into new groups. Refer to “Editing an Existing Group” on page 12-66 to learn more.

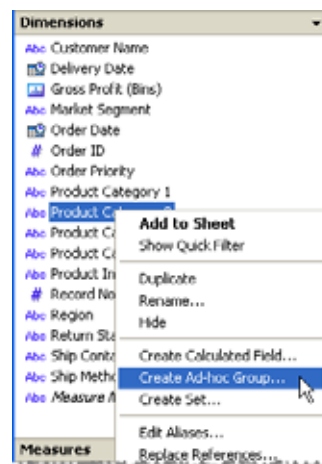
---

**Note** You can quickly un-group the dimension members by selecting the group in the view and clicking the Group button on the toolbar.

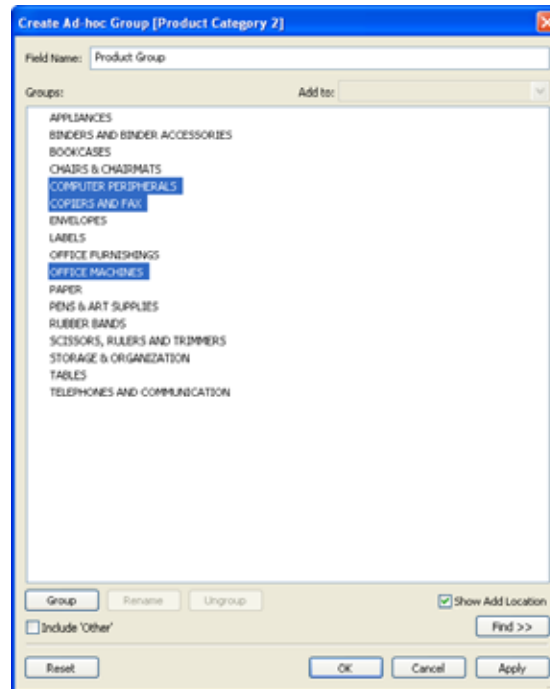
---

**To create ad-hoc groups from the Data window:**


- 1 Right-click a dimension in the Data window and select **Create Ad-Hoc Group**.



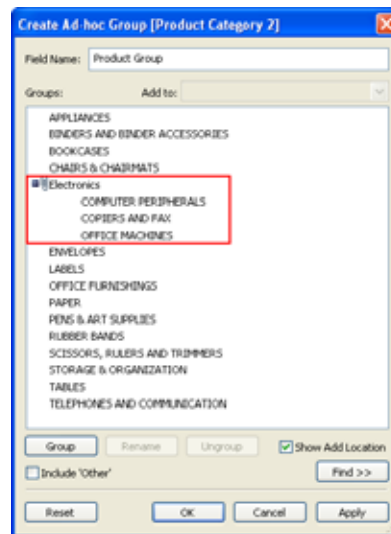
- 2 In the Create Ad-Hoc Group dialog box, select several members that you want to group. Hold the **CTRL** key on your keyboard to select multiple members.



- 3 Click the **Group** button at the bottom of the dialog box.



The selected members are combined into a single member. A default title is automatically constructed using the combined member names. Rename the group by selecting it in the list and clicking the **Rename** button at the bottom of the dialog box.

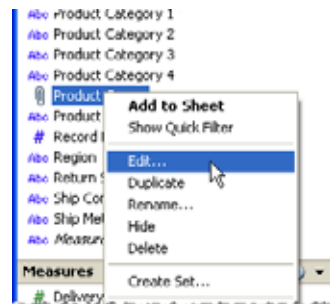


## Editing an Existing Group

After you have created an ad-hoc group either using the toolbar or from the Data window, you can add members to the group, change the default member names, as well as change the name of the grouped field using the Edit Ad-Hoc Group dialog box.

**To add members to an existing group:**

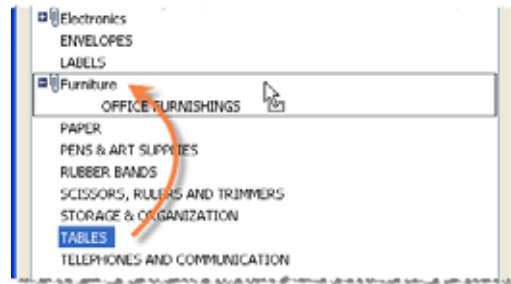
- 1 Right-click the grouped field in the Data window and select **Edit**.



- 2 In the Edit Ad-Hoc Group dialog box, do one of the following:




- Select one or more members and drag and drop them into the existing group. This method works best if you are working with a small dimension with few members.



- Select one or more members, right-click and select **Add To**. In the subsequent dialog box, select the group you want to add the selected members to and click OK.




- Select one or more members and select the group in the **Add to** drop down list at the top of the dialog box.



- 3 When finished, click **OK**.

**To rename a group:**

- 1 Right-click the grouped field in the Data window and select **Edit**.
- 2 In the Edit Ad-Hoc Group dialog box, select the grouped members and click the **Rename** button at the bottom of the dialog box.



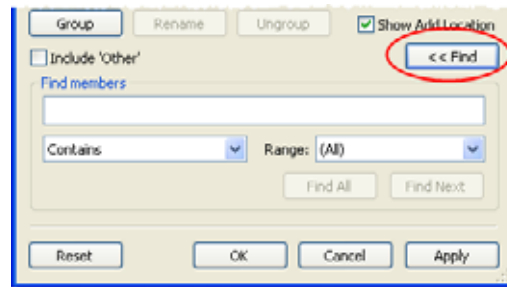
- 3 Type a new name and press **Enter** on your keyboard.
- 4 When finished, click **OK**.

### Finding Members in the Ad-Hoc Group Dialog Box

When you create groups from a large dimension with many members, use the Find option to quickly select the members you are looking for and add them to an existing group.

**To use the find options:**

- 1 Show the find options by clicking the Find button at the bottom of the dialog box.



- 2 Type all or part of the member name into the text box and select an appropriate result criteria from the drop down list. You can select whether to find members that start with, contain, or are an exact match to the search term.
- 3 Select a Range to search in. You can select to search all members, or within specific groups.
- 4 Click Find All to select all the matching members or select Find Next to manually navigate through each of the search results.
- 5 When you have found and selected the members of interest, you can quickly add them to an existing group by selecting the group from the **Add to** drop-down list at the top of the dialog box.

## Sets

Sets are custom fields you create that are based on existing dimensions, and that filter data using one or more criteria. You can create a set from any existing dimension. When you create a set for continuous dates associated with a relational data source, the set will be based on discrete values rather than a continuous range of values.

This section discusses the following topics:


- About Sets
- How to Create a Set
- Creating Sets Examples

### About Sets

The three main uses of a set are:

- Create a subset of the data – Select one or more dimension members that are of interest to you. For example, sort a field and select only cities on the west coast with populations greater than 500,000, or manually select outliers that appear in a scatter plot. Refer to “Example – A Set Containing a Subset” on page 12-81 for more information.
- Create unique encodings – Combine dimension members to create unique encodings. For example, create a set that combines market and product, and then color-encode a data view using the combined members. Refer to “Example – A Set Containing Unique Encodings” on page 12-85 for more information.
- Save filters for later use – once you have created a filter, you can save the filter as a set and use it in all of the worksheets in a workbook. This saves you from having to recreate the filter every time you want to use it.

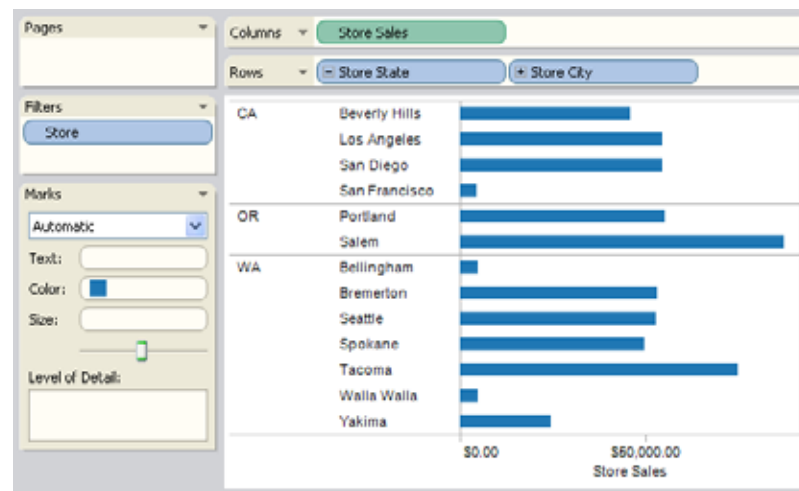
Tableau displays sets in the **Sets** area of the Data window and labels them with the  icon.



You can work with a set just as you would with any other dimension. For example, after placing a set on a shelf, you can filter the members, sort the members, and so on.

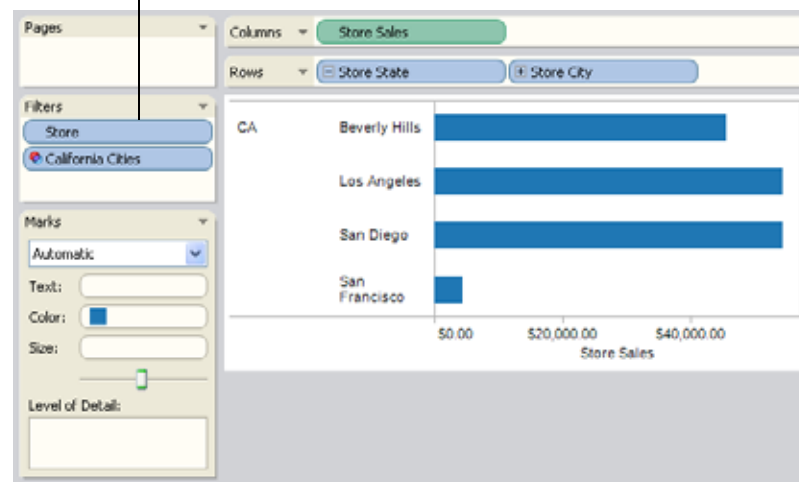
Additionally, sets are always treated as a filter. Therefore, when you place a set on a shelf, it is automatically placed on the **Filters** shelf as well.

Note that if you use a filter and a set that are based on the same dimension, the result is the intersection of the filter and the set or its descendants. For example, the following view filters the **Store** hierarchy to include only the states and the cities shown below.




If you create a set that includes only California, and then place the set on the **Filters** shelf, the resulting view will contain only the cities in California. That is, the view results from the intersection of the set and the **Store** filter.

A set and a filter based on the same dimension.



**Named Sets in Microsoft Analysis Services.** Named sets from your SSAS Server appear in Tableau in the form of Sets in the Data window. You can interact with these named sets in the same fashion you can interact with any custom sets in Tableau.



## How to Create a Set

You can create a set in one of the following ways:

- Create a Set by Selecting Marks
- Create a Set from a Field
- Create a Nest Set


The best method for you depends on your data characteristics, analysis needs, and so on. If you want to save the sets you create, you should save your work as a workbook or a bookmark. If you do not save any of your work and exit Tableau, your sets will be lost.

### **Create a Set by Selecting Marks**

Create a set by selecting marks if you want to create a subset of your data, and the data of interest can best be identified via the data view. For example, you might select outliers or the top few values from a field that's been sorted.

Create the set by manually selecting the desired marks in a data view, and then selecting **Create Set** from the view's right-click context menu.

For example, consider the scatter plot shown below. The view consists of two measures that are color-encoded by a dimension. A collection of data points deemed to be outliers are manually selected for a new set.

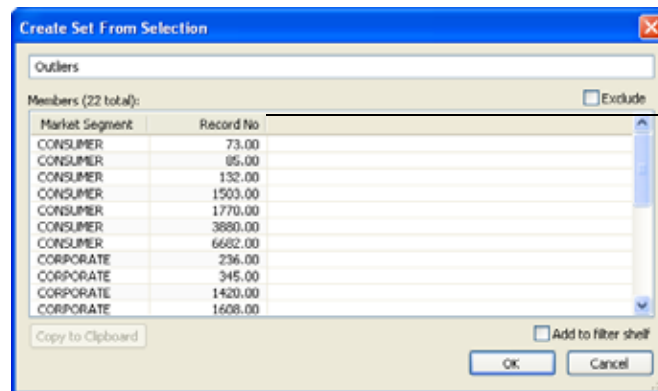


Selecting **Create Set** from the right-click context menu opens the **Create Set From Selection** dialog box. You can specify the set name, select one or more set members and copy them to the Windows Clipboard, click on a column header to sort the members, or right-click on a column header to remove the column or to restore the original sort order. Changing the sort order in the dialog box does not change the set definition. You should remove columns that aren't important to your analysis. This will make header labels easier to read and will improve performance.

Optionally select the **Exclude** checkbox in the upper right corner if you want the set to contain all members except the ones you selected.




**Note** You can optionally select to add the set to the filters shelf after you create it using the check box in the lower right corner of the dialog box.




Click on a header to sort the column.  
Right-click on a header to remove the column.

Tableau displays the new set in the **Sets** area of the Data window.



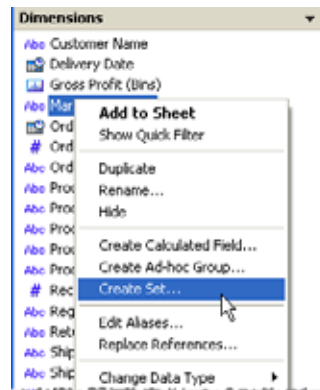
When you use the set in a data view, a header is created for each set member. As shown below, the header labels are given by the member names. For cubes, you can choose to display the fully qualified name by selecting the set in the Data window, and then selecting **Qualify Member Names** from the field's right-click context menu.



### Create a Set from a Field

Create a set from a field if you want to create a subset of a specific field.

Create the set by selecting right-clicking the field in the Data window and selecting **Create Set**.



The **Create Set** dialog box opens. Complete the dialog box by specifying the set name and selecting one or more dimension members. In addition, you can optionally define conditions and Top limits to further define the set. Refer to “Configuring the Filter Dialog Box” on page 12-21 to learn more about completing this dialog box.

The dialog box shown below is associated with a multidimensional data source. Notice that members are selected from different hierarchy levels.

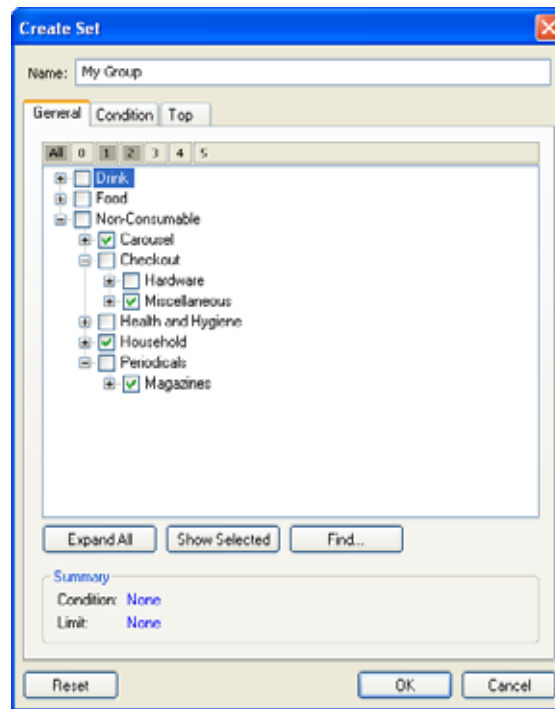


Tableau displays the new set in the **Sets** area of the Data window.



### Create a Nest Set

A nest set is a cross product of members from different dimensions. You would create a nest set if you want to encode a data view using multiple dimensions. Refer to “Example – A Set Containing Unique Encodings” on page 12-85 to learn more about this method.

---

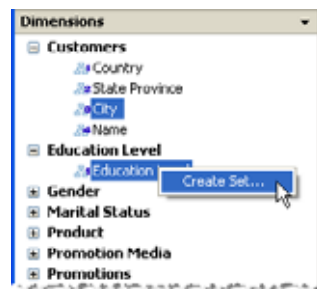
Create the next set by selecting multiple dimensions in the Data window and then selecting **Create Set** from the right-click context menu of a selected field.

---

**Note** For multidimensional data sources, you must select levels from different hierarchies.

---

For example, the selections shown below will produce a new set that consists of the **City** and **Education Level** dimensions.



The **Create Set From Selection** dialog box opens. You can specify the set name, select one or more set members and copy them to the Windows Clipboard, click on a column header to sort the members, or right-click on a column header to remove the column or to restore the original sort order.

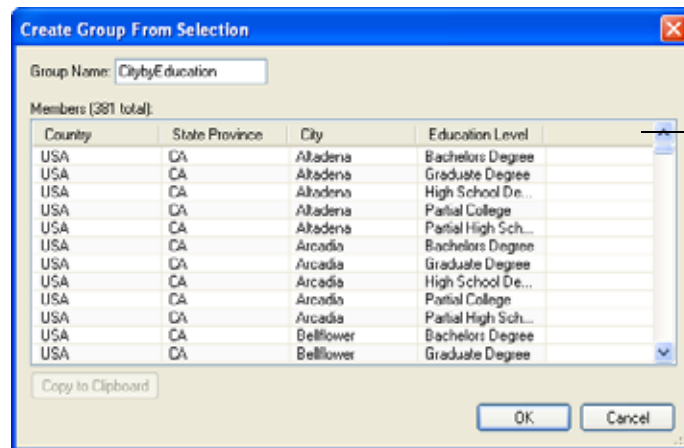


Tableau displays the new set in the **Sets** area of the Data window.



When you use the set in a data view, a header is created for each member. The header label is given by combining the original dimension names as shown below. For cubes, you can choose to display the fully qualified name.

Columns				
Product Family		Product Department		
Rows				
CitybyEducation				
	Drink			Baked Goods
	Alcoholic Beverages	Beverages	Dairy	
Altadena, Bachelors Degree	14.12	14.96	2.22	17.90
Altadena, Graduate Degree				8.84
Altadena, High School Degree	25.53	38.14	3.07	28.36
Altadena, Partial College	7.08	17.43	7.33	12.82
Altadena, Partial High School	17.07	67.62	14.70	44.36
Arcadia, Bachelors Degree	12.83	41.73	3.68	25.42
Arcadia, Graduate Degree	6.71	6.86		2.64
Arcadia, High School Degree	25.55	28.21	13.31	29.05
Arcadia, Partial College	1.95	14.18		
Arcadia, Partial High School	8.87	54.96	3.03	43.90
Stollflower, Bachelors Degree	12.54	67.04	16.97	39.97

## Creating Sets Examples

This section contains the following examples to help you understand how to create and use sets:

- Example – A Set Containing a Subset
- Example – A Set Containing Unique Encodings
- Example – Hierarchical Sets and their Descendents

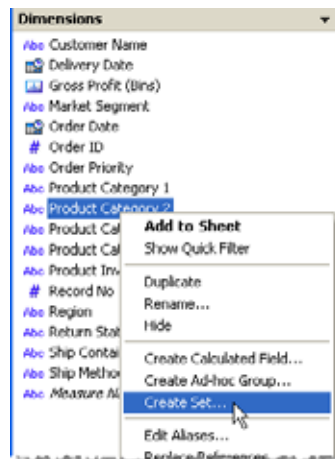
### Example – A Set Containing a Subset

One reason to create a set is so you can easily work with just the dimension members that are of interest to you. For example, you might want to work with specific geographic regions, high-value customers, or one product line in your organization. To create such a set, select the relevant dimension members using any of the methods described in “How to Create a Set” on page 12-72.

In this example, you will create a subset of the Sample Superstore data source using the **Create Set** dialog box. Follow the steps below:

- 1 Select the dimension that will form the set.

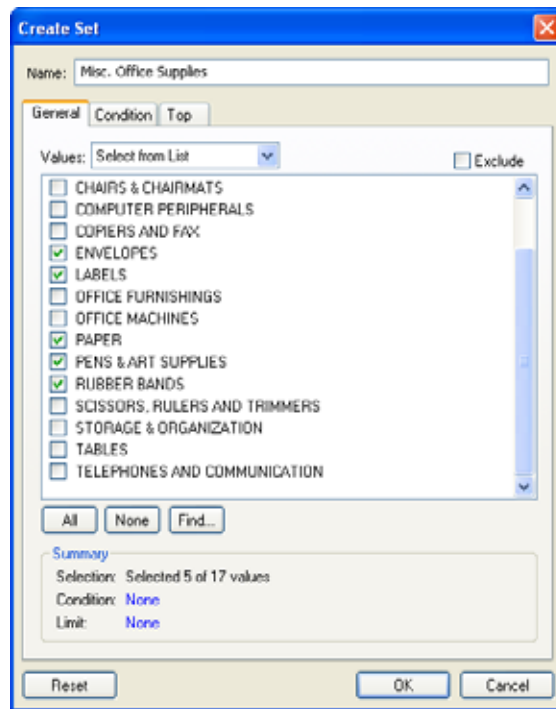
Right-click **Product Category 2** in the Data window, and select **Create Set**.



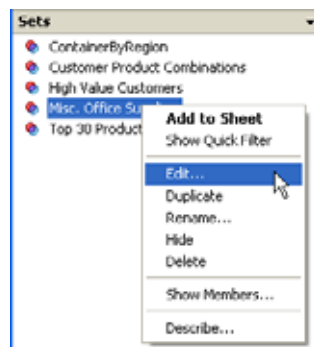
- 2 Complete the **Create Set** dialog box.

Specify the name of the set and select the dimension members that you want to include in the set. In this example, you are only interested in Envelops, Labels, Paper, Pens and Art Supplies, and Rubber Bands.

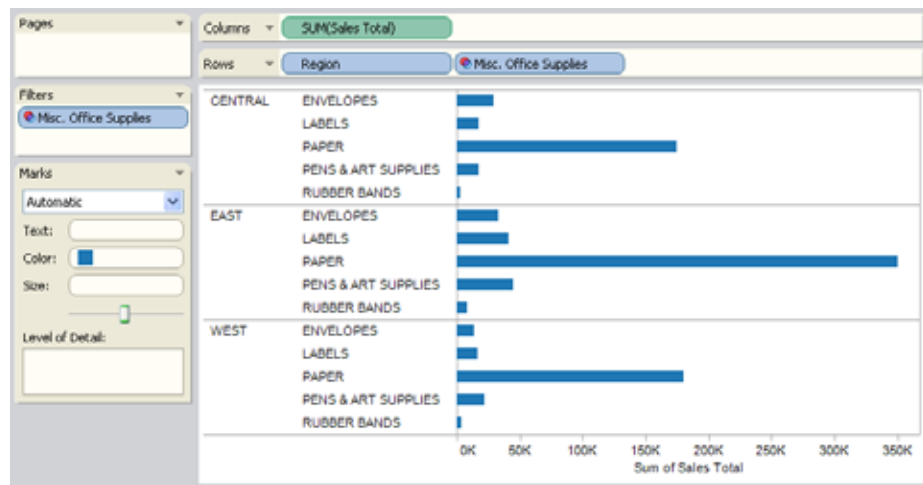




The new set displays in the **Sets** area of the Data window. You can edit the set, show set members, and so on using the right-click context menu.



You can use the set to create data views just like any other field. Note that for cubes, the members are displayed using qualified names. A qualified name has the form Food.Dairy or Drink.Beverages.Hot Beverages.Coffee, which provides information about a member's ancestors. You can turn off qualified names via the set's right-click context menu.



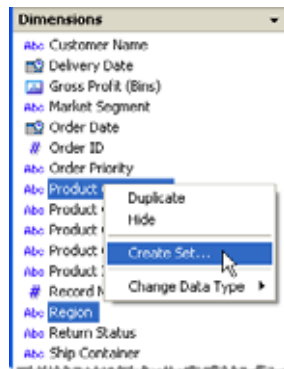
### Example – A Set Containing Unique Encodings

Encoding shelves such as **Color**, **Size**, and so on accept only one field at a time. Using the original data source fields, you are limited to encoding your data view with the members of only one dimension. By creating a set, you can encode the view with members from different dimensions.

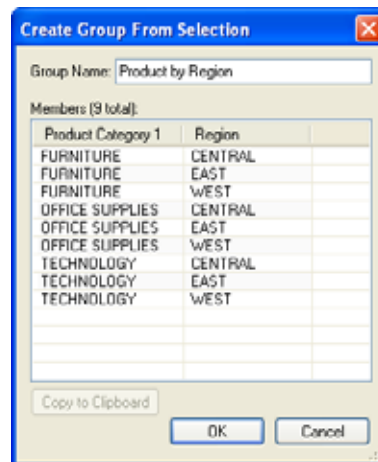
This example uses the Superstore Sales Excel data source to create a set that contains all the members from two different dimensions. The set is used to encode a data view by color, and is then filtered to include only the members of interest. The steps are as follows:

- 1 Create the set.

Create the set by selecting the **Region** and **Product Category 1** dimensions in the Data window, and then selecting **Create Set** from the context menu.



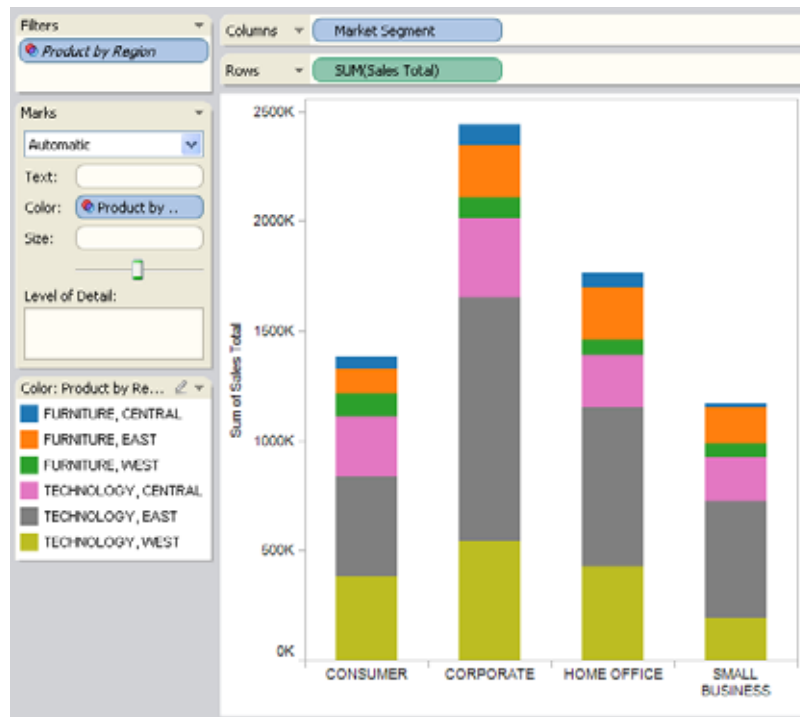
The **Create Set From Selection** dialog box opens. Call the new set **Product by Region**.



- 2 Encode the data view with the new set.


The data view shown below was created by placing the **Market Segment** dimension on **Columns** shelf, placing the **Sales Total** measure on the **Rows** shelf, and color-encoding the data using the new set.

When you place the set on the **Color** shelf, Tableau separates the marks according to the members in the set, and assigns a unique color to each member. The color legend displays each member name and its color.

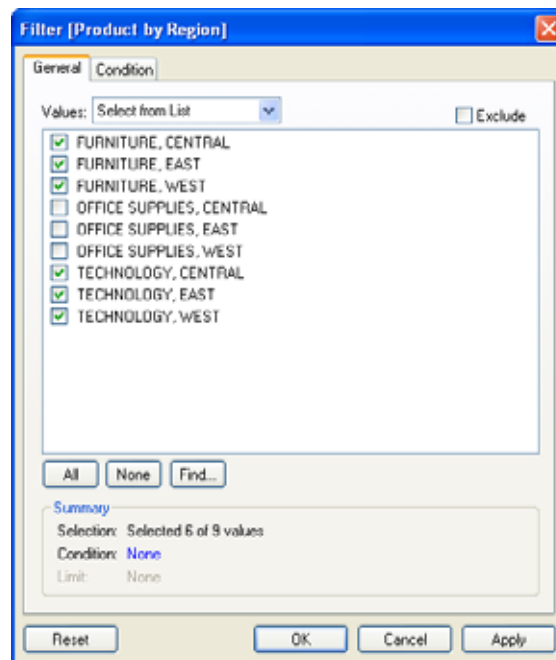


### 3 Filter the set.

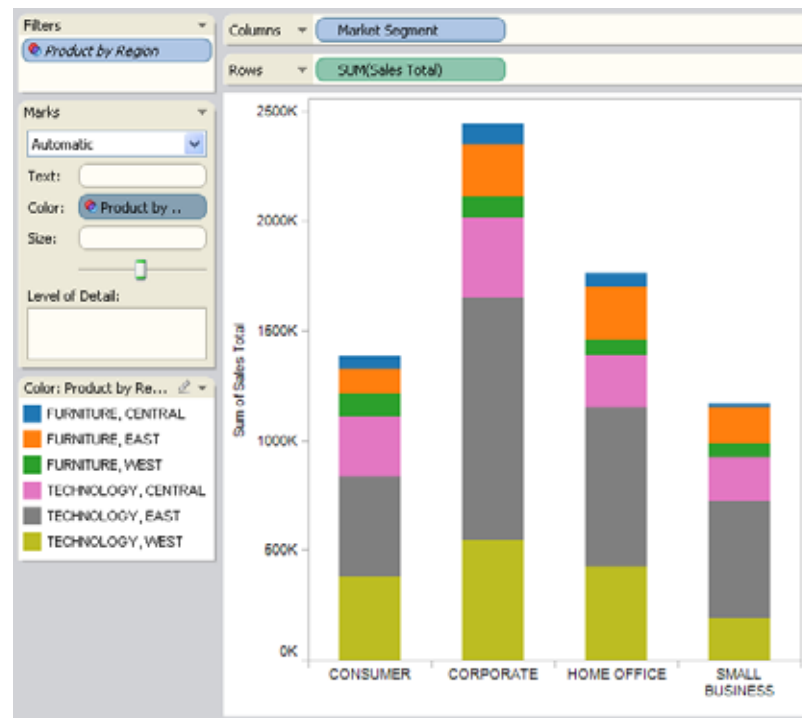
Filter the set to include only the dimension members of interest. You can open the **Filter** dialog box by selecting **Filter** on the set's field menu.



For this example, include only the Furniture and Technology products.

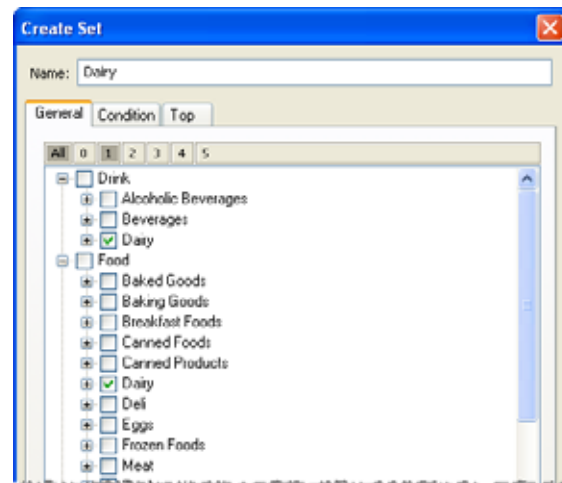


The final view is shown below. Note the name of the filtered set is italicized.



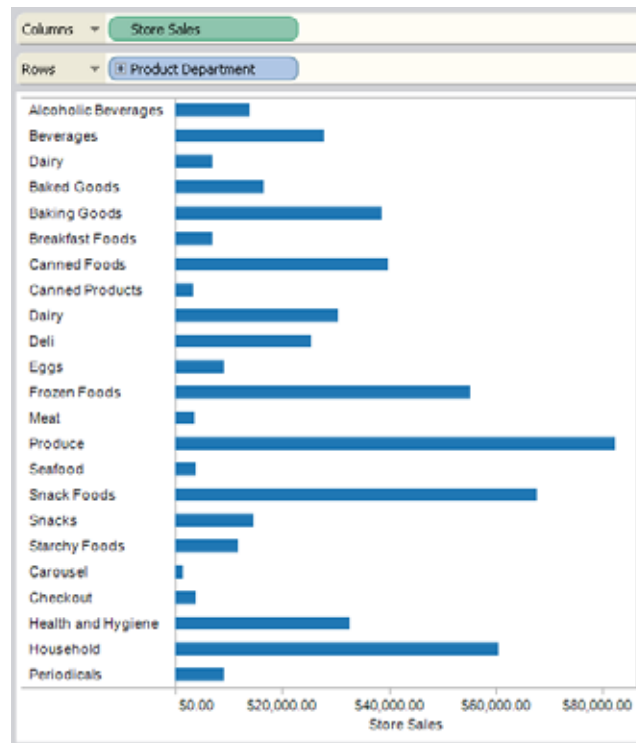
### Example – Hierarchical Sets and their Descendents

A hierarchical set filters data to the selected members and all of their descendents. For example, a set named **Dairy** is created from the **Product** hierarchy. As shown below, it includes only the Dairy product department.




Consider the following view. The **Product Category** dimension is placed on the **Rows** shelf and the **Store Sales** measure is placed on the **Columns** shelf.

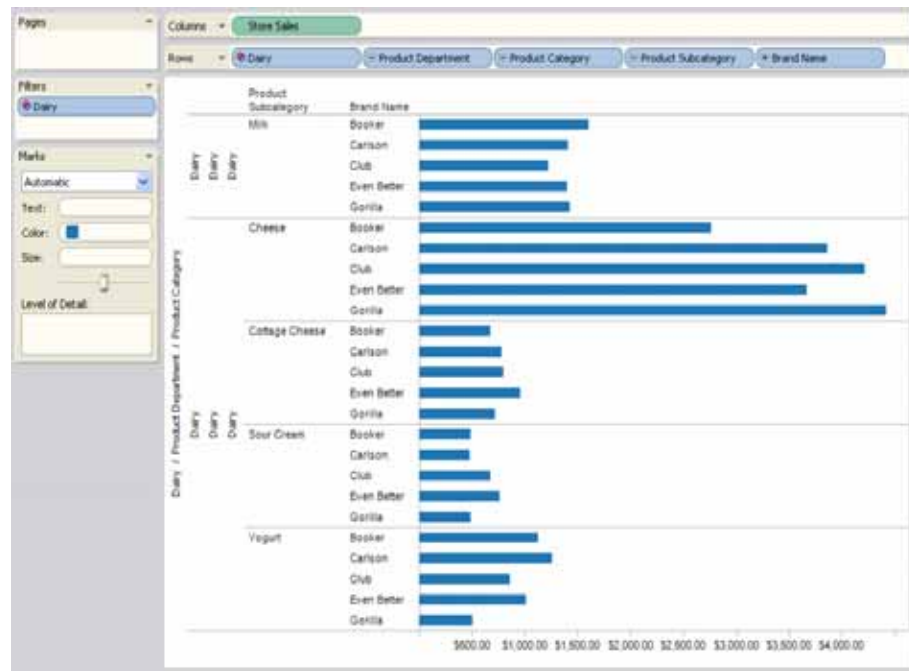




If you place the **Dairy** set on the **Filters** shelf, you can see that the view is filtered to include only the Dairy product categories.



As shown below, you can drill down into **Product Department** to reveal the **Product Category**, **Product Subcategory**, and **Brand Name** levels. As these descendents are revealed, row headers are added to the view. This is because a set filter allows you to view the levels of detail contained within the filtered members.



# Dates and Times

---

Overview . . . . .	13-2
Changing Date Levels . . . . .	13-4
Fiscal Dates . . . . .	13-8
Perfect Pivoting with Dates . . . . .	13-10
Continuous Dates . . . . .	13-12


---

## Overview

Working with dates in Tableau differs depending on whether you are using a Relational or Multidimensional Data Source. This section below discusses these differences.


### Multidimensional Data Sources


For multidimensional data sources, dates dimensions are usually organized into hierarchies that contain levels such as year, quarter, and month. These hierarchies are defined when the cube is created and you cannot modify them in Tableau. For example, the **Time** dimension from an SSAS data source is shown below.



When you place a multidimensional date on a shelf, the field is treated like any other dimension. For example, you can drill down, drill up, and so on.

### Relational Data Sources

For relational data sources, dates and times are automatically placed in the **Dimensions** area of the Data window and are identified by the  icon. For example, the **Delivery Date** and **Order Date** dimensions from an Excel data source are shown below.



When you place a relational date on a shelf, the field name is automatically modified to reflect the default date level. Tableau defines the default date level to be the level at which

---

there are multiple instances. For example, if the date field includes multiple years, the default level is year. However, if the date field contains data for just one year but includes multiple months, then the default level is month.

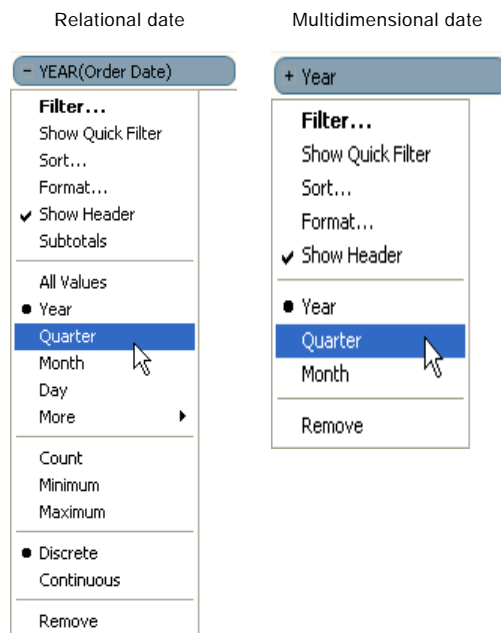
If you don't want Tableau to automatically select a date level and would rather have a date dimension be a continuous field, you can right-click the field in the Data window and select **Convert to Continuous**. The dimension turns green in the Data window and anytime you use the field it will be continuous. You can easily revert back by selecting **Convert to Discrete** from the field's context menu in the Data window. You can also convert a single field to continuous while it is on a shelf by selecting Continuous on its field menu. The field on the shelf turns green but the field in the Data window is still discrete.

To learn more about working with dates and times refer to the following topics:

- Changing Date Levels
- Fiscal Dates
- Perfect Pivoting with Dates
- Continuous Dates

## Changing Date Levels

For both relational and multidimensional data sources, you can change the date level using the field's context menu after dragging it to a shelf. For example, you can select the **Quarter** level as shown below.



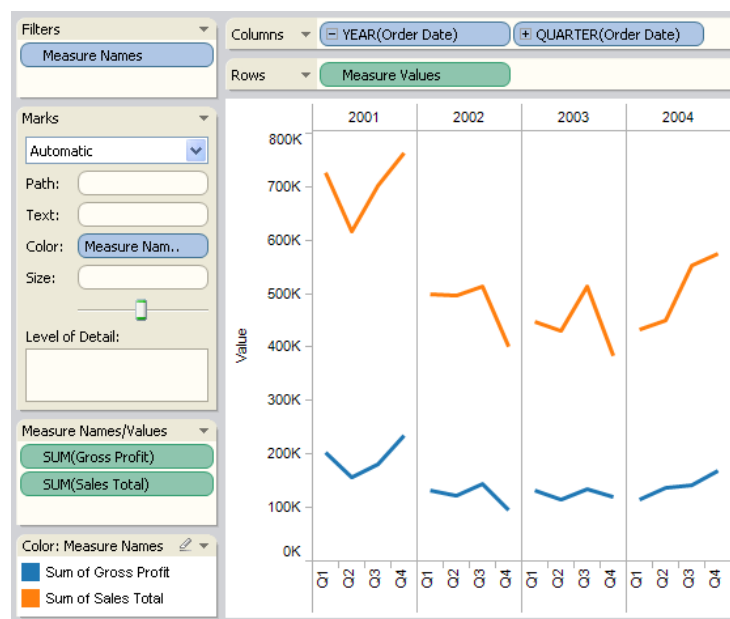
For multidimensional dates, the levels available in the context menu are given by the levels defined in the date hierarchy.

For relational dates, there are many different date levels available, as shown above. When you select a particular level, Tableau asks the data source to perform a computation on the date field. For example, suppose a particular row in your data source has a date entry of 01/23/03. The year is 2003, the quarter is 1 because January falls in the first quarter, and the week number is 4 because January 23rd falls in the fourth week. How the date level is computed depends on your data source because the computation is actually being done by the data source. Therefore, if your data source is configured to use a specific standard to compute week number, Tableau will use the same standard.

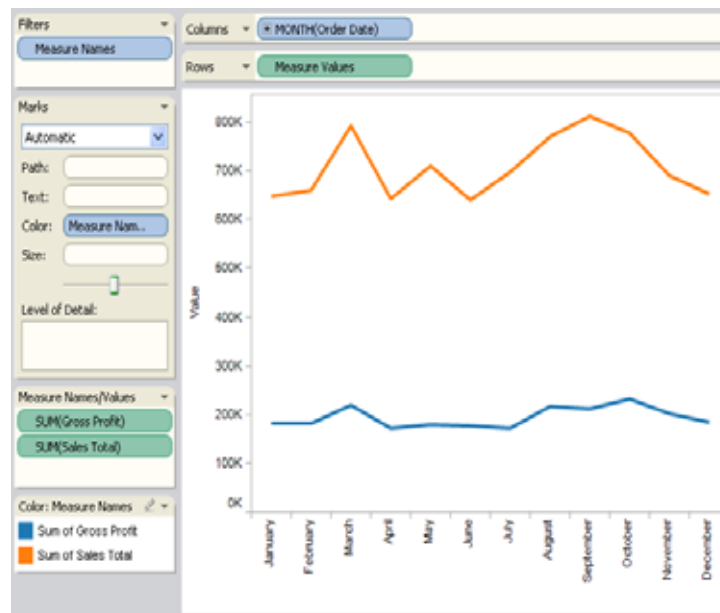
Note that some date levels might not make sense for your relational data source. For example, if the date format does not include time information such as hour, minute, or second, then selecting one of these options will not add any data to your view.

You can work with dates at varying levels of detail simultaneously. To do so, you can drill into dates by clicking the **+** control. You can also drag date fields to the **Rows** or **Columns** shelf multiple times in order to nest them and to drill down into them at varying levels of detail.

For example, the view shown below drills down into the year level to display the quarter level as well.

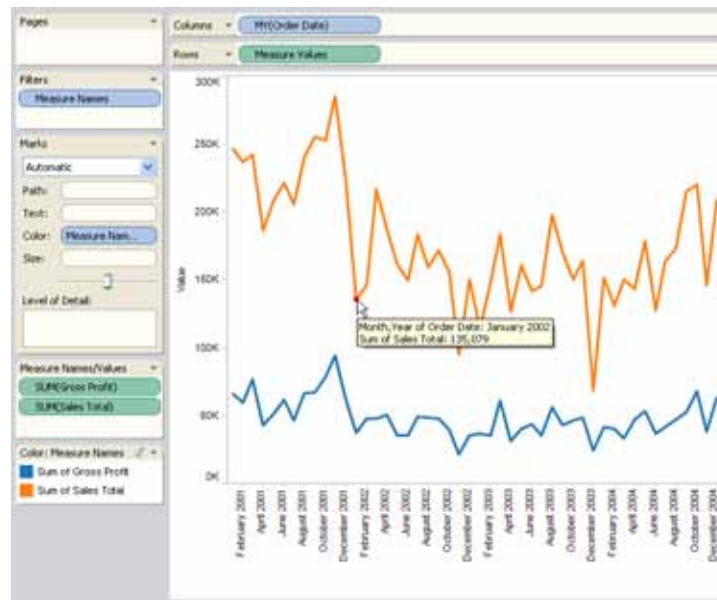


You can display the data by month by selecting **Month** from the date field's context menu. This displays the data for each month across all years.



To display finer granularity, you can select the **MMMM YYYY** level from the field menu. Tableau displays the dates using the month and the year.





## Fiscal Dates

Occasionally a date field needs to be expressed in terms of its fiscal date equivalent. For instance, calendar years always run from January 1st until December 31st. But an organization's fiscal year might start on a month other than January. For instance, a company's fiscal year might run from June 1st in one year through May 31st of the following year. In these cases, it's helpful to express the **Fiscal Year** and the **Fiscal Quarter** and the **Fiscal Week Number** rather than their calendar equivalents, when using the date field in a view.

To express date fields in fiscal terms, follow these steps:

- 1 Right-click the date dimension in the Data window and select **Fiscal Year Start**. This option is only available on fields that are classified as date dimensions.
- 2 Designate the start of the fiscal year by selecting a month from the subsequent context menu.

Whether a given level of a date dimension is affected by the conversion to a fiscal equivalent depends on the specific case. Consult the following table:

**Table 0-1:**

Date Level	When Converted to Fiscal
YEAR	The YEAR reflects the fiscal year. For instance, the year for the date June 1, 2004 would be shown as FY 2005.
QUARTER	The QUARTER reflects the fiscal quarter. For instance, the quarter for the date June 1, 2004 would be shown as Q1.
MONTH	No change in behavior. The calendar month is the same as the fiscal month.
DAY	No change in behavior. The calendar day is the same as the fiscal day.
HOURL	No change in behavior. The calendar hour is the same as the fiscal hour.

**Table 0-1:**

Date Level	When Converted to Fiscal
MINUTE	No change in behavior. The calendar minute is the same as the fiscal minute.
SECOND	No change in behavior. The calendar second is the same as the fiscal second.
WEEKNUMBER	The WEEKNUMBER reflects the fiscal week number. For instance, the week number for the date June 1, 2004 would be shown as 1.
WEEKDAY	No change in behavior. The calendar weekday is the same as the fiscal weekday.
MM/YYYY	No change in behavior. This date format always displays calendar dates, even when a fiscal year has been assigned.
M/D/Y	This date format always displays Calendar dates, even when a fiscal year has been assigned.

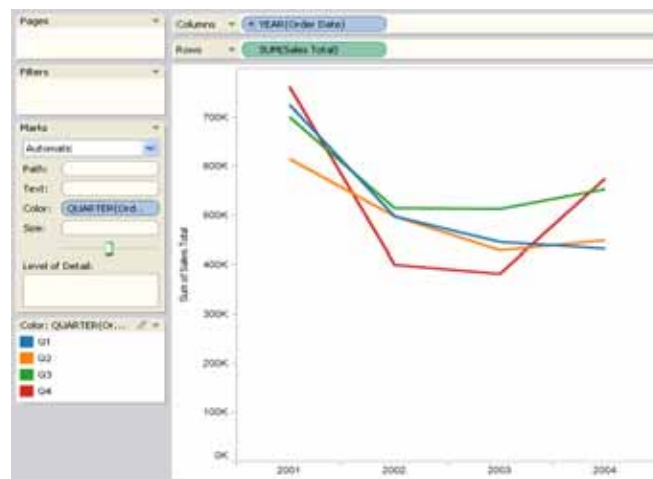
Notice that the only date level that expressly displays the conversion to a fiscal calendar is the YEAR level. Specifically, fiscal years are shown with the FY prefix. This is not true of fiscal quarters or week numbers, however, which are not shown with any special fiscal markings.

Fiscal year designations for any given date dimension are applied to all instances of the field in the Tableau workbook. Fiscal dates can only be applied to dimensions in a relational data source.

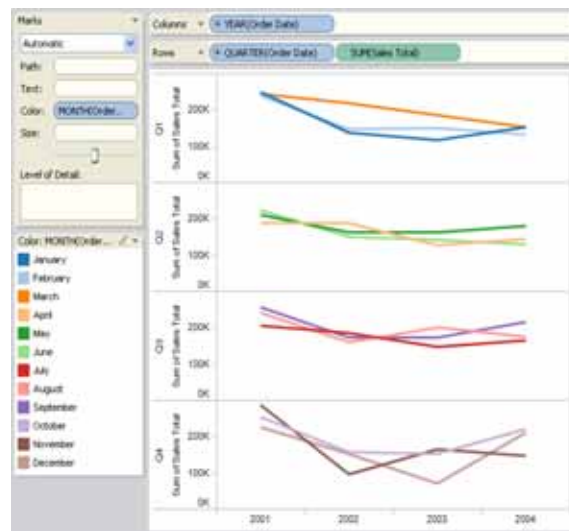
## Perfect Pivoting with Dates

You can perfect pivot dates by placing different date levels on different worksheet shelves simultaneously. Place the date field on a variety of shelves and then select the desired date level from the fields' context menus.

For example, the following line chart displays years as column headers and then color-encodes the marks by quarter.



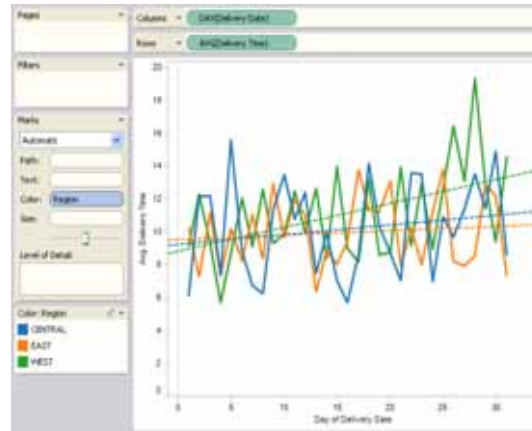
You can separate the marks by month and by quarter as shown below.



## Continuous Dates

You can treat a date as a continuous quantity after placing the field on a shelf. You do this by selecting **Continuous** from the field's context menu. This draws a quantitative axis for the date values. You can then change the displayed date range by double-clicking on the axis and specifying the desired range.

For example, the view below displays the average delivery time as a function of a continuous delivery date and color-encoded by region. As you can see, the color of the **Delivery Date** field changes from blue to green after it is converted to a continuous quantity.



Treating dates as a continuous quantity is particularly useful when you use Gantt bars or want to see trends using line charts as shown above.

By default, date dimensions are discrete fields for which Tableau automatically selects a date level when it is placed on a shelf. You can make a date dimension continuous by default by right-clicking the field in the Data window and selecting **Convert to Continuous**. The field turns green and is automatically converted to a continuous field when you drag it to a shelf. To revert to discrete again, right-click the field in the Data window and select **Convert to Discrete**.

# Inspecting Data

---

<b>Overview . . . . .</b>	<b>14-2</b>
<b>Navigation Tools . . . . .</b>	<b>14-3</b>
Select . . . . .	14-3
Pan . . . . .	14-4
Focus . . . . .	14-5
Zoom . . . . .	14-6
<b>Undo and Redo . . . . .</b>	<b>14-8</b>
<b>Color Legend Highlighting . . . . .</b>	<b>14-9</b>
<b>Drop Lines . . . . .</b>	<b>14-12</b>
<b>Reference Lines . . . . .</b>	<b>14-14</b>
Formatting Reference Lines . . . . .	14-17
<b>Summary Card . . . . .</b>	<b>14-20</b>
<b>Viewing Underlying Data (Drill-Through) . . . . .</b>	<b>14-21</b>
<b>Describing the View . . . . .</b>	<b>14-24</b>

---

## Overview


Once you have created your view, Tableau offers a selection of dynamic data inspection tools that help you isolate the data of interest and then continue to explore and analyze. For example, if you have a dense data view, you can focus on a particular region, select a group of outliers, view the underlying data source rows for each mark, and then view a summary of the selected marks include the average, minimum, and maximum values. This section discusses the following ways to inspect your data:

- Navigation Tools
- Undo and Redo
- Drop Lines
- Reference Lines
- Summary Card
- Viewing Underlying Data (Drill-Through)
- Describing the View



## Navigation Tools

The toolbar contains several navigation tools to help when you are inspecting the data in a view. When you select a navigation tool, the button is highlighted and the cursor changes to indicate which tool you are using.



The following navigation tools are available:

- Select
- Pan
- Focus
- Zoom

### Select

Selecting marks is useful when you want to visually identify a subset of the data view. Use the Selection tool by clicking on the Select icon in the toolbar. By default the selection tool is enabled when you start Tableau.

Enable selection mode.



You can select any individual mark by clicking on it. You can select multiple marks by holding down the **Ctrl** key. You can also drag the cursor to draw a box around the marks you want to select. Finally, you can combine these methods to quickly select all the marks of interest.

As shown below, after selecting the marks of interest, you can use the right-click context menu associated with any of the marks to copy the data, export the data, and so on.



## Pan

You can move your view of a table up and down as well as left and right with the pan tool. Panning is particularly useful when your data view contains many panes, and you want to move quickly from pane to pane.

Use the Pan tool by clicking on the pan icon on the toolbar.

Enable pan mode.



For example, the following view contains many panes. Panning allows you to easily move the view from Coffee products in Illinois to Tea in Washington.



## Focus


Focusing is useful when you want to see greater detail in a smaller area. Focusing also filters data.

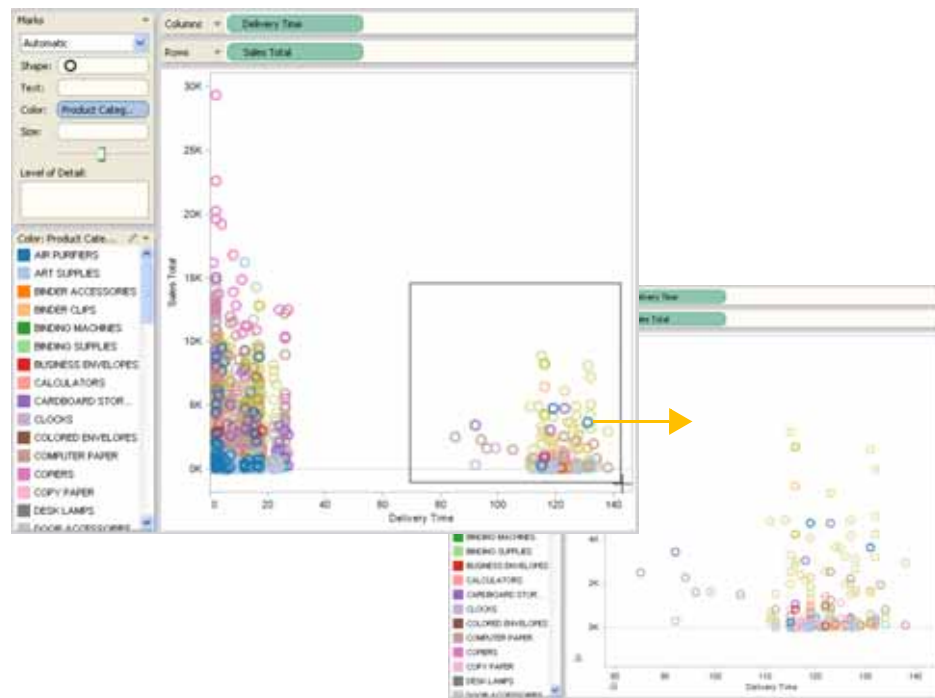
Use the focus tool by clicking on the Focus button on the toolbar.


Enable focus mode.



Note that while in focus mode, the focus selection is limited to a single pane. To focus, drag the cursor to draw a rectangle around that data you want to focus on. If the table contains axes, Tableau changes the axes limits based on the focus limits. If the table contains dimensions, Tableau filters the data by excluding dimension members that are outside the focused pane.

For example, the following view contains outliers. Focusing allows you to see more detail for these data. Notice that the  icon appears on each axis indicating that their limits have been modified and locked. To reset the limits, right-click an axis and select **Reset Axis Range** from the context menu.



If you want to reset the graph to its original view, right-click the view and select **Clear Manual Sizing** on the context menu. You can also use the **Undo** button  to undo each operation you performed on your view.


## Zoom

Zooming is useful when you have a lot of data in a view and you want to focus on a specific part of the view without excluding the rest. Unlike the Focus tool, which adds a filter to the view excluding all deselected marks, the Zoom tool simply increases the size of the view making it easier to read.

Use the zoom tool by clicking the **Zoom** button on the toolbar.

Enable zoom mode.




The cursor changes to a magnifying glass . To zoom out, click on the area of the view you want to zoom in on. To zoom out hold the **Ctrl** key on your keyboard and click on the view. You can click as many times as you want to continue zooming in or out. You can also zoom to a specific size using the zoom drop-down list on the toolbar. The following table describes the options in this list:

50%	Displays the view at one half the default size.
100%	Displays the view at the default size.
200%	Displays the view at twice the default size.
Entire View	Shrinks or Expands the view to fit the entire view within the active window.
View Width	Shrinks or Expands the view to fit the entire width within the active window regardless of the height.
View Height	Shrinks or Expands the view to fit the entire height within the active window regardless of the width.

## Undo and Redo

You can perform unlimited undo and redo of your actions. You can undo almost all actions in Tableau by pressing the **Undo** button on the toolbar. Likewise, you can redo almost all actions by pressing the **Redo** button on the toolbar.

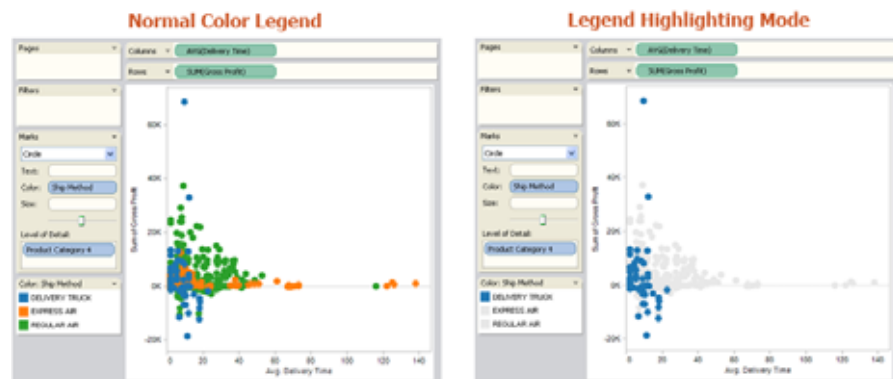


In this regard, every workbook behaves like a web browser. You can quickly return to a previous view. Or you can browse all the views of a data source that you have created. Tableau saves the undo/redo history across all worksheets until you exit. The history is not saved between sessions.

## Color Legend Highlighting


Color legend highlighting is a powerful analytical mode for the color legend that allows you to focus on select members in the view. When you turn on color legend highlighting the marks associated with the selected items in the color legend are colored while all other marks are gray.


For example, the views below show the relationship between delivery time and profit for several products. The view on the left uses the normal color legend, all marks are colored based on their shipping method. The view on the right uses legend highlighting to call out the products that were delivered via Delivery Truck.



You can easily switch between legend highlighting and normal modes using the color legend card menu. Then, if you like how a view is highlighted, you can assign the highlight colors to the color palette. The old colors are replaced with the highlight colors.

**To turn on color legend highlighting:**


- 1 Click the **Highlight** button  at the top of the color legend or select **Highlight Selected Items** on the color legend card menu.



- 2 Select an item in the color legend.

Once legend highlighting is turned on, you can quickly focus on specific data in the view by selecting different items in the color legend.

**To turn off color legend highlighting:**

- Click the **Highlight** button  at the top of the color legend or select **Highlight Selected Items** on the color legend card menu.




If you like how the view is highlighted and want to keep a specific member highlighted even when you turn off legend highlight mode, you can assign the highlight colors to the existing color palette. The original color legend is discarded and the highlight colors become the new color palette for the legend.




**To assign the highlight colors to the color palette:**

- Select **Assign Highlight Colors** to Palette on the color legend card menu.



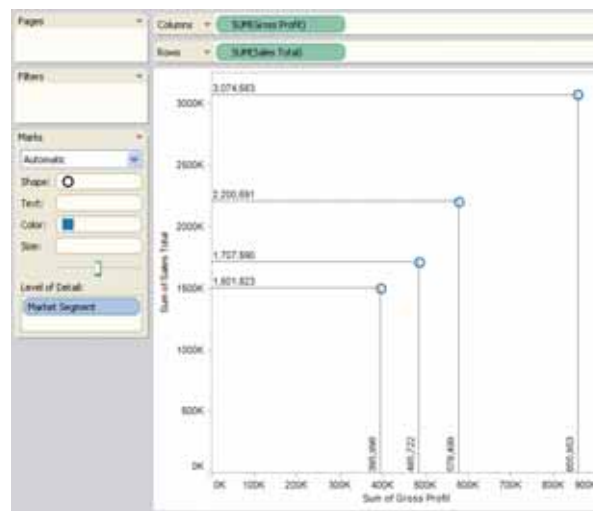
---

**Note** Color legend highlighting can also be used in dashboards. To turn on legend highlighting in a dashboard, simply select the color legend you want to use for highlighting and then click the highlight button  at the top.

---

## Drop Lines

Drop lines are most useful for distinguishing marks and calling out their position in the view. For example, in a view that is dense with scatter marks, you can turn on drop lines to show the position of a particular data point. When you add drop lines a line is extended from the marks to one of the axes. You can choose to show drop lines all the time or only when a mark is selected.




**To add drop lines to the view:**

- Right-click on the pane and select **Drop Lines**.

By default, drop lines are set to only show when the mark is selected. You can change this setting and specify other options in the Drop Lines dialog box.

**To edit drop lines:**

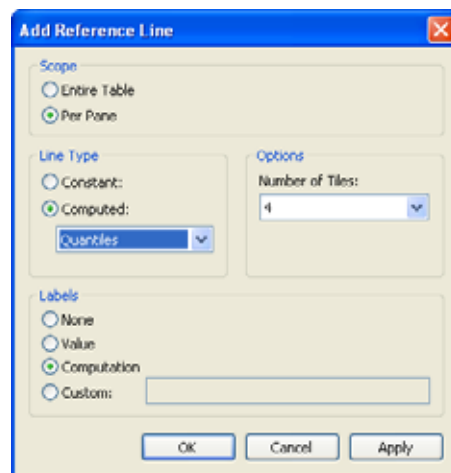
- 1 Right-click on the pane and select Edit Drop Lines to open the Drop Lines dialog box.



- 2 In the Drop Lines dialog box select an axis to draw the line to, whether to always show the drop lines, and whether to show labels.
- 3 When finished click **OK**.

## Reference Lines

A reference line is typically used to mark a specific value or region on an axis. For example, if you are analyzing the monthly sales for several products, you may want to include a reference line at the average sales mark so you can see how each product performed against the average. Tableau lets you add an unlimited number of constant and computed reference lines. Add reference lines using the Add Reference Line dialog box.




**To add a reference line:**

- 1 Right-click on a quantitative axis and select **Add Reference Line**.




- 2 In the Add Reference Line dialog box, select one of the following scopes:

- **Entire table** - adds a reference line to the entire table across all panes.



References lines are applied to the entire table.

- **Per pane** - adds a reference line on a per pane basis. Computed reference lines are recalculated for each pane in the view.



- 3 Select whether to use a constant or computed line type.

A reference line can either be placed at a specified constant value on the axis or at a computed value such as minimum or average. If you select a computed line type, you must also select the type of computation. Additionally, reference lines are computed using the data showing in the view, not on the underlying data. Select from the following:

- **Average** - places a line at the average value.
- **Confidence Interval**- places lines on the axis indicating the interval between which lie the specified percentage of values.
- **Maximum** - places a line at the maximum value.
- **Median**- places a line at the median value.
- **Minimum** - places a line at the minimum value.

- **Percentile** - places a line indicating a specified percentile. When you select this option, you must also select the percentage.
  - **Quantile** - breaks the view into a specified number of tiles using reference lines. When you select this computation, you must also select the number of tiles.
  - **Standard Deviation** - places lines at the specified number of standard deviations above and below the mean. When you select this option you must specify the factor, which is the number of standard deviations, whether the computation is on a sample or the population, whether to show both lines, just the upper line, or just the lower line.
  - **Sum** - places a line at the SUM of all the values in either the pane or the entire view.
- 4 Select one of the following label options:
- **None** –select this option to not include a label for the reference line.
  - **Value** – select this option to include a label that is the corresponding value along the axis.
  - **Computation** – select this option to display an automatic label. The label is based on the axis and the computation that is applied.
  - **Custom** – select this option to type a custom label into the text box.
- 5 When finished, click **OK**.

You can edit, remove, and format reference lines by right-clicking on the reference line in the view. To remove all reference lines from an axis, right-click on the axis and select **Remove all reference lines**. For more information about formatting reference lines refer to “Formatting Reference Lines” on page 14-17.

## Formatting Reference Lines

Tableau lets you format the reference lines you add to the view. You can change the style, width, and color of these lines as well as the font properties of the label. Specify format settings that apply at the worksheet level (refer to “Lines” on page 20-11) or format a specific reference line.

**To format a specific reference line:**

- 1 Right-click on the line and select **Format**.
- 2 In the Format window you can specify the line style, as well as the label font, alignment, number, and shading properties using the drop-downs.




- 3 When finished click **OK**.

The view below contains several reference lines, formatted in various ways.

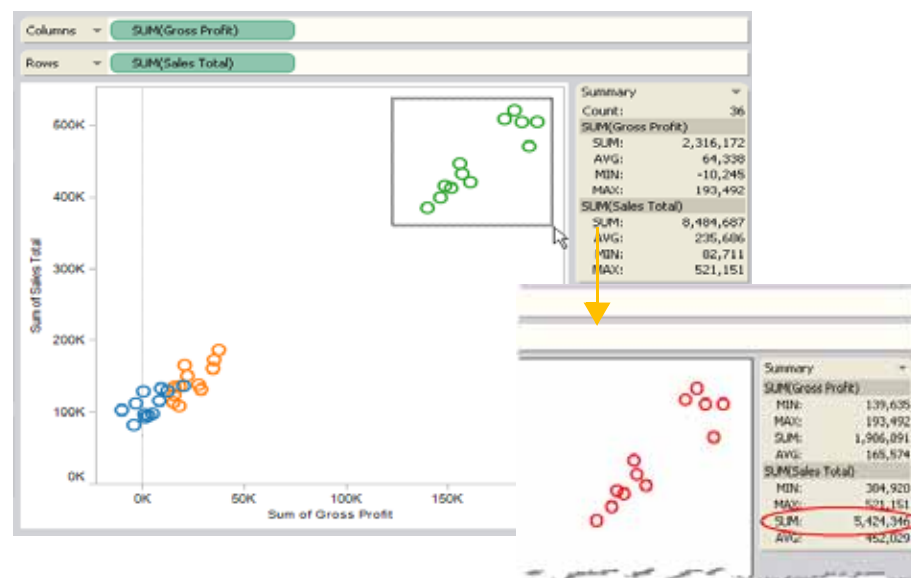




## Summary Card

The summary card is a really quick way to view information about a selection or the entire data source. The card shows the SUM, MIN, MAX, and Average for each measure in the view. You can hide or show the Summary Card by selecting it on the View Cards toolbar menu . You can also select **View > Cards > Summary**.

Consider this example, the view below is a scatter plot of profit vs. sales for three different product categories. You can see that the technology category contains high profit and high sales products (the green marks). When you select these marks, the summary card quickly shows you that these products account for \$5,424,346 in sales with a minimum sale of \$384,920.



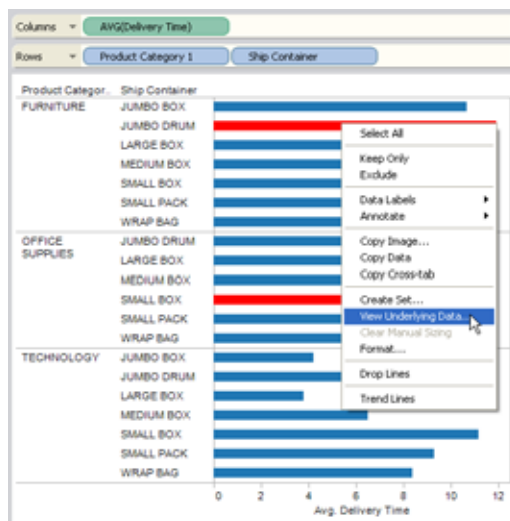
## Viewing Underlying Data (Drill-Through)

For some data sources, you can view the underlying disaggregated data associated with one or more marks in a data view by using the View Underlying Data (drill-through) function. You cannot view underlying data for views that are disaggregated.

In other words, you can display the values for each row in the data source that compose the marks. You might want to do this in order to verify the aggregated value associated with a mark, or to isolate and export the individual rows associated with data of interest such as outliers.

The View Underlying Data (drill-through) function works with all relational and multi-dimensional databases except Hyperion Essbase databases. While you can view underlying data with Microsoft Analysis Services multi-dimensional databases, the database must be drill-through enabled and there are some restrictions to the data you can view. For more information about drill-through with a Microsoft Analysis Services database refer to “Viewing underlying data (Microsoft Analysis Services)” on page 14-23.


The view shown below displays the average delivery time for two product dimensions as a bar chart. Suppose you want to view the underlying data for the largest marks in each pane. To do this, select the marks of interest, right-click in the table, and select **View Underlying Data** from the context menu. Alternatively, you can select the **Analysis > View Underlying Data** menu item.



For MySQL data sources, you might get an error message that states “Got a packet bigger than max\_allowed\_packet.” This happens when the data you are trying to access is larger than the maximum packet size of your MySQL server. To solve this problem, increase the value of the set-variable = max\_allowed\_packet entry in the my.cnf configuration file under [mysqld]. You must then restart mysqld as well as your servlet engine.

**Note** Viewing underlying data may not return any records if you are using a field that contains floating point values as a dimension. This is due to the precision of the data source and mainly occurs when you are connected to Microsoft Excel, Microsoft Access, or text files.

The underlying data for the selected marks are displayed in the **Underlying Data** dialog box. Notice that the number of rows that compose the underlying data is shown in the upper left of the dialog box. For this example, the number is 4,656.



You can sort the data by clicking one or more column headers. To restore the original sort order, right-click the sorted column header and select **Restore Original Sort Order**.

By default, the **Show all fields** check box is not selected. This means that only the fields placed on shelves (or fields referenced by a calculation placed on a shelf) for the current

worksheet are displayed. If you select this check box, Tableau will include all data source fields in the dialog box.


If you want to export one or more data source rows, select the data points of interest by clicking individual rows in the table, or by pressing the **Select All** button. You can then click **Copy to Clipboard** to copy the selected data to the Windows Clipboard and paste it into a another file.

---

**Note** You can view all data source rows by not selecting any marks, and then viewing the underlying data. This option is not available with a Microsoft Analysis Services database; you must select a single mark to view underlying data for.

---

**Viewing underlying data (Microsoft Analysis Services)** . Viewing underlying data with a Microsoft Analysis Services database works in almost the same way it does with relational data sources. The difference is that a Microsoft Analysis Services cube is generally set up and configured by an administrator who decides whether it is enabled for drill-through and the fields that a user is allowed to view. That means that when you try to view underlying data using a database that is not enabled, you may get an error message alerting you that the cube is not enabled for drill-through.



In addition, Microsoft Analysis Services databases limit viewing underlying data to a single mark at a time.

When you are viewing underlying data for a field, the **Show all fields** option is checked and disabled by default. With a Microsoft Analysis Services database, only the fields specified by the administrator are shown so you cannot choose to include all data source fields in the dialog box.

## Describing the View

Occasionally you may want to succinctly summarize an analysis you have completed on a worksheet. You might then want to remind yourself of what it shows (the filters that are applied, etc.), and finally, you may want to share a summary of the analysis with someone else.

When you choose **View > Describe Sheet**, you can view a description of the workbook, data source, fields and layout of the current worksheet. This summary includes the Caption in the first line, but expounds on other important summary information. This information can be copied and exported to other applications using the clipboard.

**Note** If you have Trend Lines turned on, the Describe Sheet dialog box includes information about the trend line model, including an anova table. Refer to “Trend Line Model Terms” on page 18-21 to learn more about the terms used to describe the model.



# Hyperlinks

---

<b>Overview . . . . .</b>	<b>15-2</b>
<b>Adding URL Links . . . . .</b>	<b>15-3</b>
<b>Adding Sheet Links . . . . .</b>	<b>15-6</b>
<b>Launching Links . . . . .</b>	<b>15-9</b>
<b>Using Field Values in Hyperlinks . . . . .</b>	<b>15-10</b>
Using Field Values in URLs . . . . .	15-10
Using Field Values in Link Names . . . . .	15-11
Valid and Invalid Links . . . . .	15-13
<b>Managing Links . . . . .</b>	<b>15-14</b>
Editing Links . . . . .	15-14
Enabling/Disabling Links . . . . .	15-15
Removing Links . . . . .	15-15
<b>Example: Adding Hyperlinks . . . . .</b>	<b>15-16</b>

---

## Overview

Tableau allows you to add context to your data using hyperlinks. Link to Web pages, files, and other Tableau worksheets directly from your analytical results. For example, in a view showing the home sales in a specific neighborhood, you could use hyperlinks to point to census data regarding the neighborhood, a satellite image of the house from a Web-based mapping service, and another worksheet showing similar houses in the area.

There are two kinds of hyperlinks in Tableau: URL links and sheet links. A URL link points to a web page, file, or other web resource, while a sheet link points to another sheet in the workbook.

This section discusses the following topics:

- Adding URL Links
- Adding Sheet Links
- Launching Links
- Using Field Values in Hyperlinks

---

**Note** For more information about hyperlinks and constructing URLs, please visit our Training & Support Center.

---




## Adding URL Links

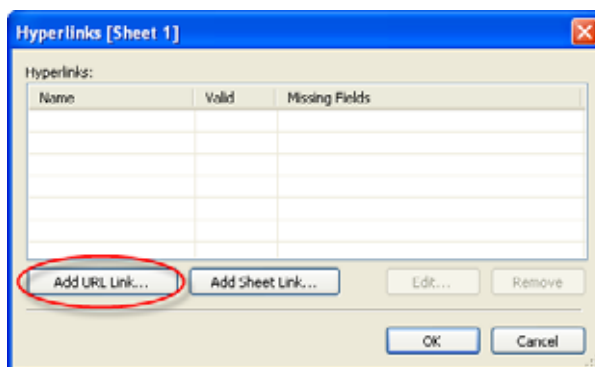
A URL link is a hyperlink that points to a Web page, file, or other web-based resource outside of Tableau. You can use Hyperlinks to link to more information about your data that may be hosted outside of your data source. To make the links relevant to your data, you can substitute field values of a selection into the URL as a parameters.

**To add a Hyperlink:**

- 1 Select **Data > Hyperlinks**.



- 2 In the Hyperlinks dialog box, click **Add URL Link**.



- 3 In the subsequent dialog box, specify a name for the link.

The name is what displays on the context menu in the view when you are launching the link (refer to “Launching Links” on page 15-9). In a real estate example, a link to a Google satellite image of a house could be called “Show satellite image.” Additionally, you can use variables in the name that will be filled in based on the values of the selection. Refer to “Using Field Values in Link Names” on page 15-11 to learn more.

- 4 Specify the URL.

You can use any URL that your browser can recognize including Web pages, ftp resources, and files.

Just as you can use variables in the *name* of the URL, you can also use field values as parameters in the URL. That means that you can send information about each selected mark to a given Web site. To learn more about using fields as parameters in your URL refer to “Using Field Values in URLs” on page 15-10.

- 5 Optionally select one or more of the following options:

- **URL Encode Data Values** - select this option if your data contains values that use characters that are not allowable in a URL. For example if one of your data values contains an ampersand, such as “Sales & Finance,” the ampersand must be translated into characters that your browser understands (URL encoded) if you want to include that value in the URL.

- **Enable Multi-Select** - select this option if you are linking to a Web page that can take lists of values as parameters in the link. For example, say you select several products in a view and you want to see each product's details hosted on a Web page. If the server can load multiple product details based on a list of identifiers (product ID or product name), you could use multi-select to send the list of identifiers as parameters.

When you enable multi-select you must also define the item delimiter, which is the character that separates each item in the list (often a comma). You must also define the Delimiter Escape, which is used if the delimiter character is used in a data value.

- 6 When finished, click **OK**.

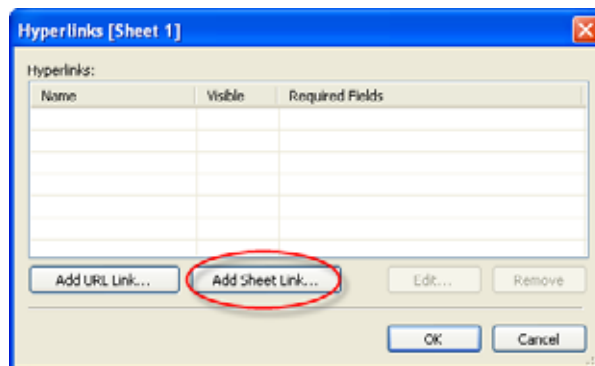
## Adding Sheet Links

Sheet links are a way to send information between worksheets. Typically a sheet link is used to send information from a selected mark to another sheet showing related information. For example, when looking at a view showing the sales price of houses, you may want to be able to select a particular house and show all comparable houses in a different view. You could define a sheet link to accomplish this task. First you need to decide what comparable means. In this case, say that comparable houses are houses with a similar sale price and square footage. A sheet link to show comparable houses can be defined by selecting a destination worksheet and defining filters on sales price and square footage.

Sheet links work by sending the data values of the relevant source fields as filters to the destination sheet. If you launch the sheet link described in this example from a house that sold for \$450,000, the destination sheet will have a filter to only show houses that sold for the same amount.

**To add a sheet link:**

- 1 Select **Data > Hyperlinks**.
- 2 In the Hyperlinks dialog box, click **Add Sheet Link**.



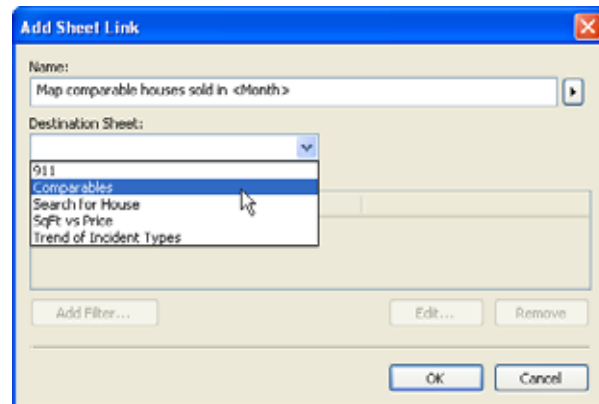
- 3 In the subsequent dialog box specify a name for the Sheet Link.

The name is what appears on the context menu in the view. For example, when sending housing information from one sheet to a map, the name could be “Map all comparable houses sold in February” You can use variables in the name that will be filled in based


on the values of the selected field. Refer to “Using Field Values in Link Names” on page 15-11.



- 4 Select the destination sheet using the drop-down list. The destination sheet must be in the current workbook.

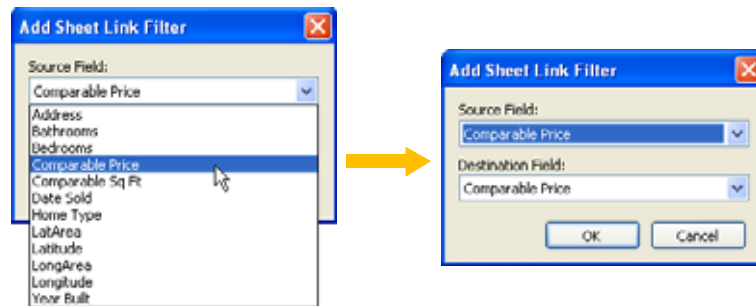


- 5 Click **Add Filter** to specify the data to show on the destination worksheet.



- 6 In the subsequent dialog box, select a source field and a destination field. When you launch the sheet link from a specific mark, the destination sheet will only show marks

whose destination fields match the source field values of that specific mark. In the comparable houses sheet link example, the Source Field is Comparable Price and the Destination Field is Comparable price. That means when you launch the sheet link for a house whose comparable price is greater than \$1.5 million, the destination worksheet will only show houses that also have a comparable price of greater than \$1.5 million.



7 Repeat this process of adding filters as many times as is necessary.

8 When finished, click **OK**.

If you are connected to a relational data source, you can add sheet links across data sources even if the field names are not exactly the same. One data source may have a field titled Latitude while another has a Lat field. Using the drop down lists in this dialog box, you can associate the Latitude field to the Lat field. When using a multidimensional data source, the destination sheet must use the same data source. Moreover, the source fields must match the destination fields.

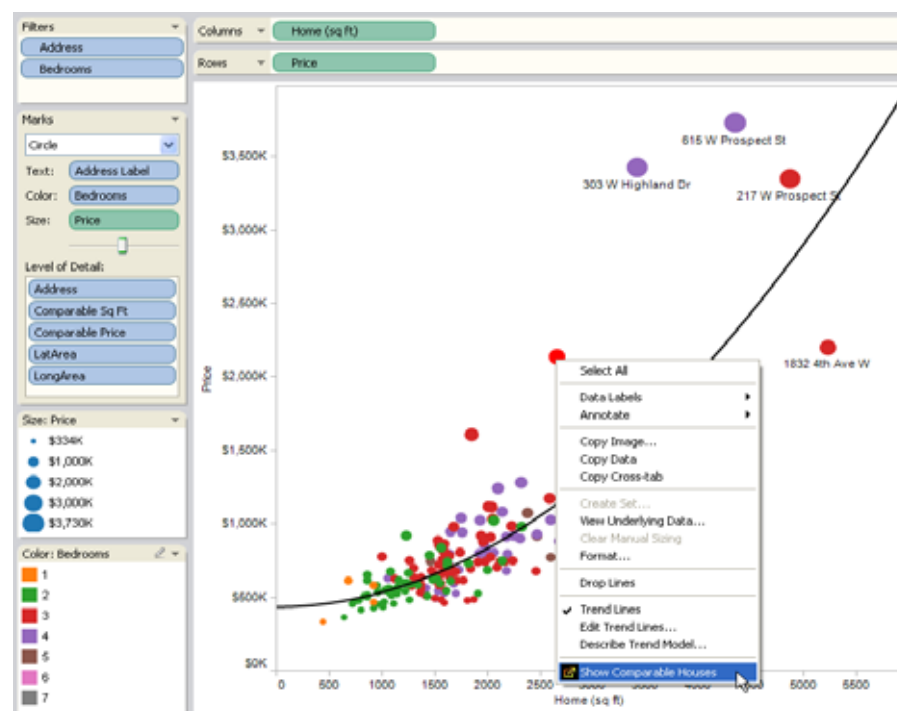
---

**Note** The fields available in the Destination Field drop-down list are dependent on what you selected as the Source Field. Only fields with the same data type as the source field can be selected as a destination field.

---

## Launching Links

Once you've added hyperlinks to the view, you can launch the links directly from the view by right-clicking on a data point or header and selecting the link on the context menu. If you have multi-select enabled, you can launch links from a selection of multiple marks in the view (refer to "Adding URL Links" on page 15-3 to learn more about multi-select).



However, links are not always visible for every worksheet and mark. Because links are mapped to specific fields in the data source, links will only be available for the worksheets that use the mapped fields. For example, if you add a hyperlink that uses both Latitude and Longitude as parameters in the link, the link will only be available to worksheets that use Latitude and Longitude in the view. Additionally, the link is only available on marks and headers that contain relevant values. Refer to "Valid and Invalid Links" on page 15-13 to learn more.

## Using Field Values in Hyperlinks

When you add a hyperlink in Tableau you often want to use values from your data as parameters in the name of the link as well as the link itself. Using fields as variables in the link name makes the menu item that launches the link specific to the selected mark. More commonly, using field values as parameters in the URL itself allows you to send information about a specific data point to the destination Web page. This section discusses the following topics:

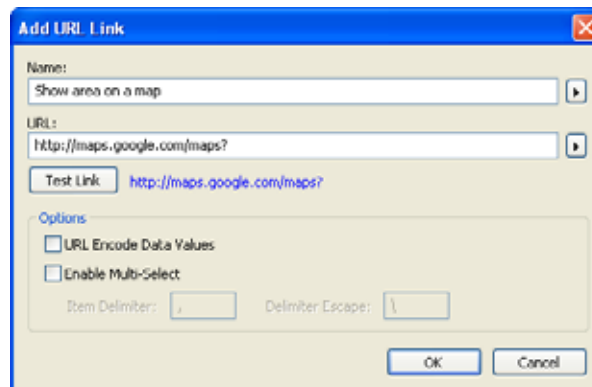
- Using Field Values in URLs
- Using Field Values in Link Names

### Using Field Values in URLs

Tableau lets you add fields as variables into the hyperlinks so when you follow the link the values of those fields are included. For example, when linking to an online mapping service, you can use an address field as a parameter so launching the link from a specific data point shows the address associated with that record on a map.

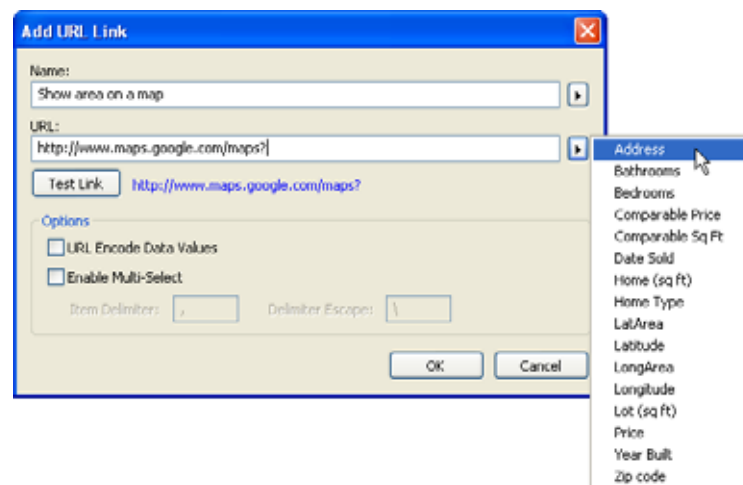
#### To add a field as a parameter in a URL:

- 1 In the Add URL Link dialog box, begin typing the URL for the link.





- 2 Place the cursor where you want to insert a field value.
- 3 Click the arrow to the right of the text box and select the field you want to add as a parameter. The field name is added to the URL between angle brackets.



You can continue adding field parameters as many times as you need to create the URL.

**Note** The list of available fields only includes non-aggregated fields and the first level in dates. To use aggregated field values or date bins as a parameter in a link, you must first create a calculated field and then use the name of that field in the link. Then, the calculated field must also be used in the view, for the link to be available. A good way to use these fields is by placing them on the Level of Detail shelf.

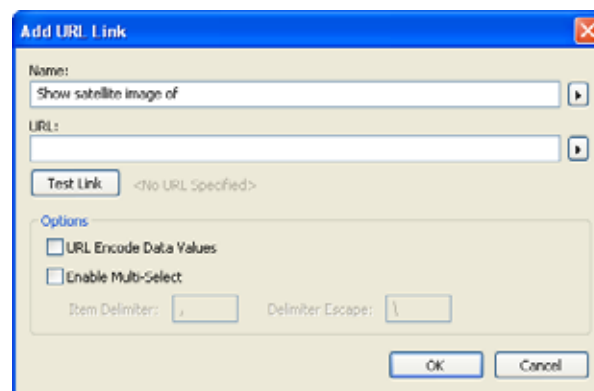
## Using Field Values in Link Names

In addition to using field values in the links themselves, you can use fields as variables in the link names. The name of the link displays on the context menu for the relevant marks allowing you to launch the link. Using field variables in the name is useful in making the link specific to the select mark. In a view showing real estate information, you could name a hyperlink to satellite images from an online mapping service, “Show satellite image of

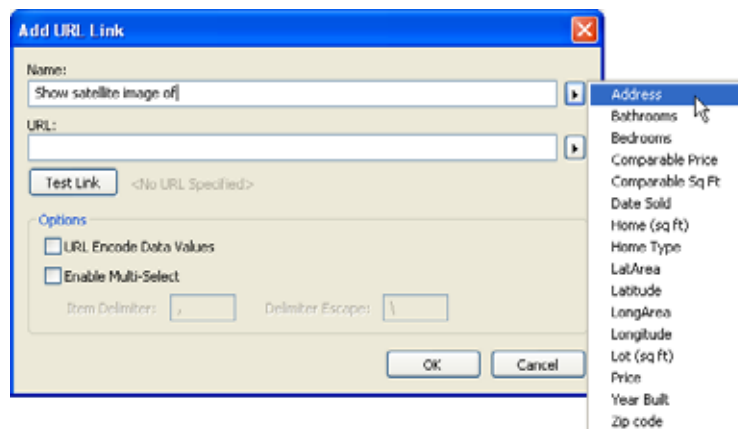
<Address>.” When you right-click on a specific mark, the <Address> tag is replaced with the location value associated with that mark.

**To add a field as a variable in a Name:**

- 1 In the Add URL Link dialog box, begin typing the name for the link.



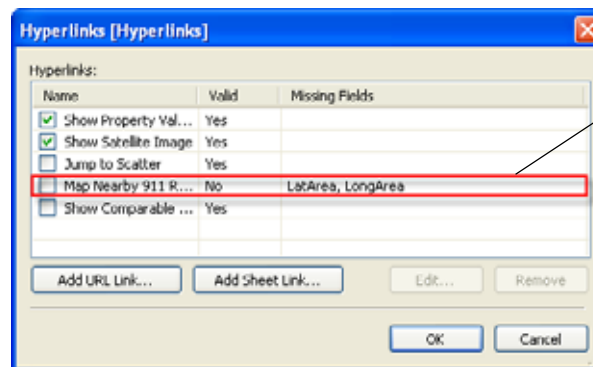
- 2 Place the cursor where you want to insert the field value.
- 3 Click the arrow to the right of the text box and select the field you want to add as a variable. The field name is added between angle brackets.



## Valid and Invalid Links

Sometimes when you add links to a workbook, the link is not valid for every worksheet. That's because the fields that you use as parameters in both the name and the URL must be used in the view. An easy way to add a field is to place it on the Level of Detail shelf. You can learn more about the valid and invalid links in the Hyperlinks dialog box.

The hyperlinks dialog box lists all the links added to the workbook, whether each link is valid for the active worksheet, and the missing fields if it is not valid.



This link is not valid for the active worksheet because the LatArea and LongArea fields are not used in the view. The link does not display on the right-click context menu.

## Managing Links

After you add a link of any kind, you can always go back and edit the definition or remove it entirely. In addition, you can enable and disable links depending for each worksheet. Links that are enabled display on the context menu for any relevant worksheet. This section discusses the following ways to manage your links:

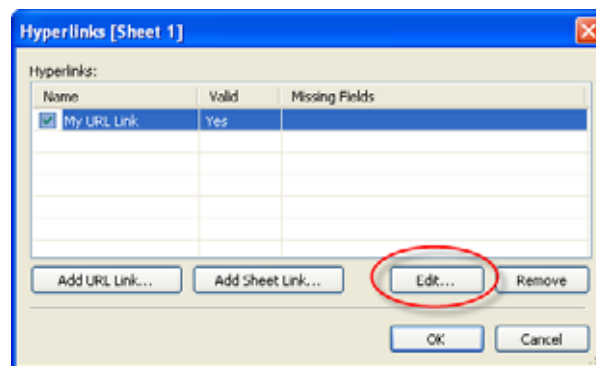
- Editing Links
- Enabling/Disabling Links
- Removing Links

### Editing Links

After creating a link, you can always go back and edit the name, URL, and change the URL options.

**To edit a link:**

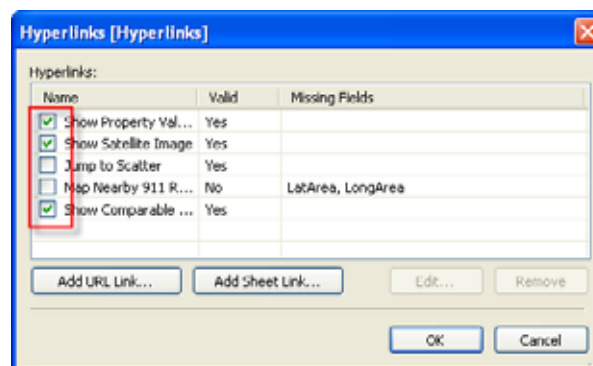
- 1 Select **Data > Hyperlinks**.
- 2 In the Hyperlinks dialog box, select the link you want to edit and click **Edit**. You can also double-click the link name.



- 3 Make the changes in the subsequent dialog box and click **OK**.
- 4 Click **OK** to apply the changes and close the Hyperlinks dialog box.

## Enabling/Disabling Links

When you add a link it displays on the context menu of all worksheets that use the fields that are used in the URL (refer to “Valid and Invalid Links” on page 15-13). However, often you will want to only show a subset of the links on any given worksheet. Use the checkboxes in the Hyperlinks dialog box to enable or disable hyperlinks for the active worksheet. Enabled links display in the context menu.



## Removing Links

After adding a link you can remove it using the Hyperlinks dialog box. When you remove a link, it is removed for the entire workbook. Refer to “Enabling/Disabling Links” on page 15-15 to learn about hiding links on specific worksheets.

**To remove a link:**

- 1 Select **Data > Hyperlinks**.
- 2 In the Hyperlinks dialog box, select the link you want to edit and click **Remove**.
- 3 Click **OK**.


## Example: Adding Hyperlinks

In this example you will build a view using the Sample - Superstore Sales data source that shows the Totals sales for several high priced products sold. Then you will add a hyperlink that opens a Google™ image search in your browser showing you pictures of a selected product in the view.


**Connect to the sample and build the initial view:**

- 1 Select **Data > Connect to Data Source**.
- 2 In the subsequent dialog box select the Sample - Superstore Sales (Excel) data source and click **OK**.


- 3 Drag the **Sales Total** measure to the Columns shelf and the **Product Category 4** dimension to the Rows shelf.




The resulting view shows the sales for many different kinds of products.



- 4 To make things a little simpler, filter to only show products with sales greater than \$5000. You can add a filter by right-clicking the Sales Total field on the Columns shelf and selecting **Show Quick Filter**.



- 5 On the Filter card, either use the sliders to set the lower limit to 5000 or click the text "8.35" and type 5000. Press Enter on your keyboard to apply the changes.



The resulting view shows the sales for only products with sales over \$5000.




### Add a hyperlink that goes to Google™ images search:

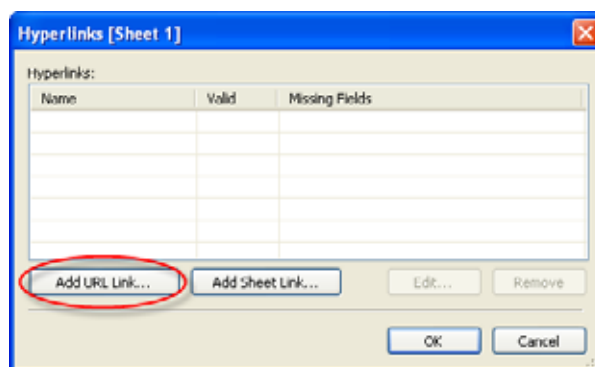
Now let's add a hyperlink from this view to a Google™ image search in your web browser that will show the selected product online. You can then launch the link by selecting a product to see online and selecting the launch from the right-click context menu.



- 1 Select **Data > Hyperlinks**.




- 2 In the Hyperlinks dialog box, click **Add URL Link**.




- 3 In the subsequent dialog box, first specify a name for the link. In this case type “Show Image on Google.” The name displays in the context menu when you right-click on a


mark or header in the view. Generally, you should use verbs that describe the link action in your hyperlink names.




- 4 To construct the URL start with the base web address, in this case type `http://www.google.com/` into the URL text box.



- 5 Google uses special words to determine what kind of search you are doing. For this example, you want to do an image search so type the word "images" after the base web address.

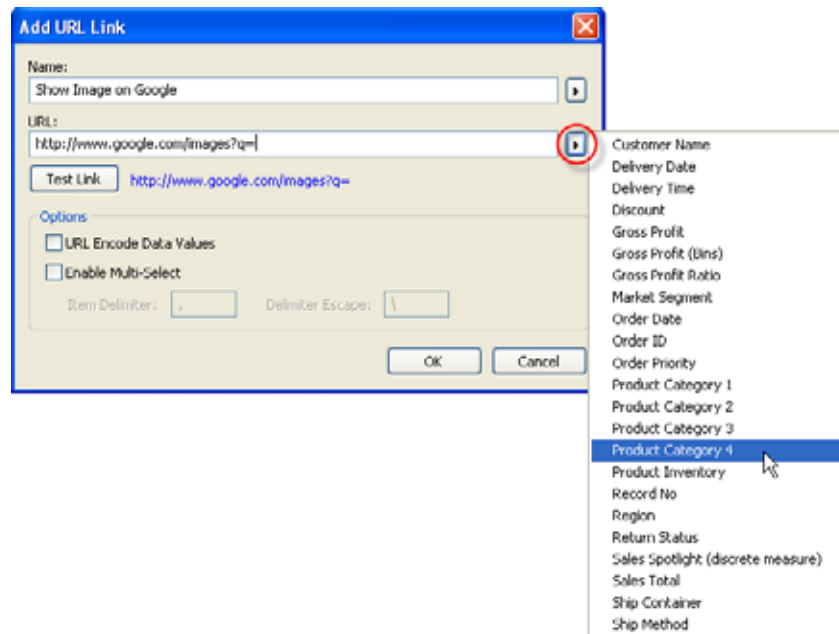


- 6 Next, to specify the search term type `"?q="` into the text box. These characters let the server know that the search term is coming.

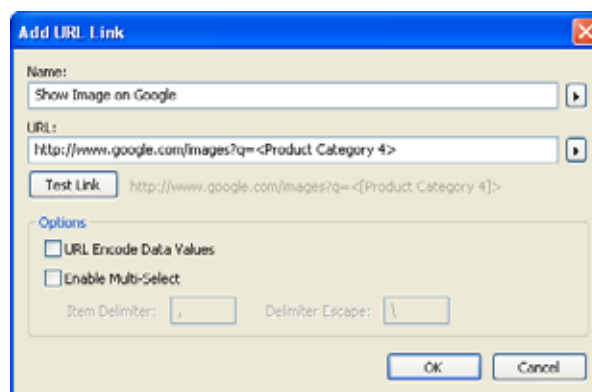


- 7 Finally, you need to specify the search term. You want the term to be the Product Category 4 value for any given mark. Rather than typing a value directly into the URL, you can insert a field name as a variable. When you launch the link from a selection, the

variable is replaced with the field's value. To insert a field as a variable, click the arrow on the right of the text box and select Product Category 4.



- 8 The final URL should be: "http://www.google.com/images?q=<Product Category 4>"



You can click the **Test Link** button, which will test the link substituting a random value from your data source for the Product Category 4 variable. For example, the URL the gets sent to the server could look like the following:

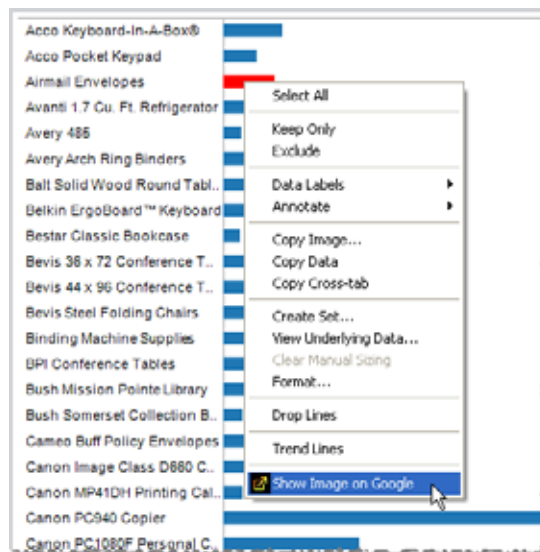
`http://www.google.com/images?q=12%20Colored%20Short%20Pencils`

Notice that the spaces in the product name are translated to the “%20” characters. That’s because spaces are not allowed in a URL, so they have to be encoded into a way the server can understand them.

- 9 Click **OK** in both the Hyperlinks dialog boxes.

#### Launch the link from a selection.

- 1 Right-click on a mark in the view.
- 2 Select the **Show Image on Google** item at the bottom of the context menu.



**Note** You can also right-click one of the table headers displaying the product categories.

# Calculations

---

<b>Overview</b>	<b>16-3</b>
<b>Aggregations</b>	<b>16-4</b>
How Aggregation and Disaggregation Work	16-8
Example – Aggregating and Disaggregating Data	16-10
<b>Calculated Fields</b>	<b>16-18</b>
How to Create a Calculated Field	16-18
How to Create a Calculated Member	16-20
Copying and Pasting Calculated Fields	16-22
Example – Creating a Calculated Field	16-24
Example-Spotlighting Using Calculations	16-37
Aggregate Calculations	16-30
<b>Table Calculations</b>	<b>16-40</b>
Quick Table Calculations	16-40
Defining Table Calculations	16-41
Advanced Table Calculations	16-58
<b>Binned Data</b>	<b>16-59</b>
<b>Totals</b>	<b>16-63</b>
Local vs. Server Computation	16-63
Default Aggregation	16-64
Grand Totals	16-65
Subtotals	16-71

<b>Percentages . . . . .</b>	<b>16-74</b>
About Percentages . . . . .	16-74
Percentages and Aggregations . . . . .	16-75
Percentage Options . . . . .	16-77

---

## Overview

To extract meaningful results from your data, you might want to perform one or more calculations. Some calculations are predefined in Tableau, while you can customize others to suit your specific needs. The following calculations are supported:

- Aggregations – View your data at different levels of detail. For example, you might want to view data in an aggregated state such as a summation or an average, or you might want to view the data in a disaggregated state and work with the individual rows of a data source.
- Calculated Fields – Create new fields that are based on existing data source fields, and common functions and operators. Use a standard dialog box that shows available functions and fields to author these custom fields. If you are using a multidimensional data source, you can also define MDX calculations.
- Table Calculations – Create calculations that are applied to the values in the entire table and are often dependent on the table structure itself, such as running totals and year to date growth.
- Binned Data – Create new fields that are based on binned measures.
- Subtotals – Add subtotals to the rows and columns of a table.
- Grand Totals – Add totals to the rows and columns of a table.
- Percentages – View data as percentages rather than as absolute numbers. The percentages can be based on rows, columns, panes, or the entire table.

---

**Note** You can use all of the above calculations if you are using a relational data source, however, multidimensional data sources do not support aggregations and binned data.

---

You can use all of these different types of calculations simultaneously. For example, you can create a new calculated field called **Profit** that is the difference between the **Sales** and **Cost** fields. You could then apply an aggregation (like a summation) to this new field in order to view total profit over time. You could then display the numbers as percentages and turn on grand totals to see how these percentages vary from category to category. Finally, you could bin the new field and display the data as a histogram.



## Aggregations

Sometimes it is useful to look at numerical data in an aggregated form such as a summation or an average. The mathematical functions that produce aggregated data are called *aggregation functions*. Aggregation functions perform a calculation on a set of values and result in a single value. For example, a measure that contains the values 1, 2, 3, 3, 4 aggregated as a sum results in a single value: 13.

---

**Note** You can aggregate measures using Tableau only for relational data sources. Multidimensional data sources contain aggregated data only.

---

For example, if you have 3,000 sales transactions from 50 products in your data source, you might want to view the sum of sales for each product, so that you can decide which products are the most important.

Tableau provides a set of predefined aggregations that are shown in the table below.

Aggregation	Description	Result for measure that contains 1, 2, 2, 3
All Values	Returns all unique values in a measure or dimension.	3 values (1, 2, 3)
Sum	Computes the sum of the numbers in a measure. Null values are ignored.	1 value (8)
Average	Computes the arithmetic mean of the numbers in a measure. Null values are ignored.	1 value (2)
Minimum	Computes the smallest number in a measure or continuous dimension. Null values are ignored.	1 value (1)
Maximum	Computes the largest number in a measure or a continuous dimension. Null values are ignored.	1 value (3)

Aggregation	Description	Result for measure that contains 1, 2, 2, 3
Standard Deviation	Computes the standard deviation of all values in the given expression based on a sample population. Null values are ignored. Returns a Null if there are fewer than 2 members in the sample that are not Null. Use this function if your data represents a sample of the population.	1 value (0.8165)
Standard Deviation Population	Computes the standard deviation of all values in the given expression based on a biased population. Assumes that its arguments consist of the entire population. Use this function for large sample sizes.	1 value (0.7071)
Variance	Computes the variance of all values in the given expression based on a sample. Null values are ignored. Returns a Null if there are fewer than 2 members in the sample that are not Null. Use this function if your data represents a sample of the population.	1 value (0.6667)
Variance Population	Computes the variance of all values in the given expression based on a biased population. Assumes that its arguments consist of the entire population. Use this function for large sample sizes.	1 value (0.5000)
Count	Counts the number of rows in a measure or a dimension. When applied to a dimension, Tableau creates a new temporary column that is a measure because the result of a COUNT is a number. You can count numbers, dates, booleans, and strings. Null values are ignored in all cases.	1 value (4)

Aggregation	Description	Result for measure that contains 1, 2, 2, 3
Count Distinct	Counts the number of unique values in a measure or dimension. When applied to a dimension, Tableau creates a new temporary column that is a measure because the result of a COUNT is a number. You can count numbers, dates, booleans and strings. Null values are ignored in all cases. This function is not supported for Microsoft Access, Microsoft Excel, and Text file data sources.	1 value (3)
Disaggregate	Returns all records in the underlying data source.	4 values (1, 2, 2, 3)

You can also define custom aggregations as described in “Aggregate Calculations” on page 16-30. Note that depending on the type of data view you create, Tableau will apply these aggregations at the appropriate level of detail. For example, Tableau will apply the aggregation to individual dimension members (the average delivery time in the East region), all members in a given dimension (the average delivery time in the East, West, and Central regions), or groups of dimensions (the sum of sales for all regions and for all markets).

You may specify a default aggregation for any measure that is not a user-defined aggregation. A default aggregation is a preferred calculation for summarizing a continuous or discrete field. The default aggregation is automatically used when a measure is first placed on a shelf. Right clicking any measure in the Data window displays the **Default Aggregation** option, as shown below. In this example, a user sets the default aggregation for the Budget Margin measure to be an Average.

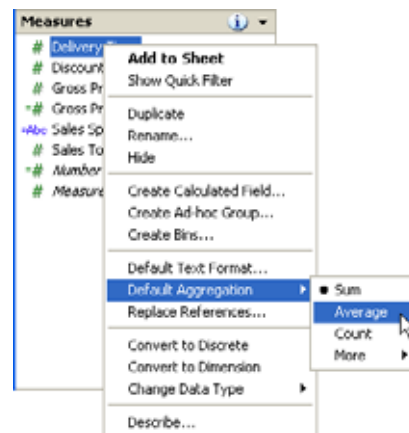


Tableau also allows you to view data in **disaggregated** form (relational databases only). This is an extremely powerful feature. When data are disaggregated, you can view all of the individual rows of your data source. For example, after discovering that the sum of sales for rubber bands is \$14,600, you might want to see the distribution of individual sales transactions. To answer this question, you need to create a view that shows individual rows of data. That is, you need to disaggregate the data (refer to “How Aggregation and Disaggregation Work” on page 16-8). Also, one way to look at disaggregated data is to view the underlying data that’s displayed in a table. Refer to “Describing the View” on page 14-24 for more information.

## How Aggregation and Disaggregation Work

When you place a measure on a shelf, Tableau automatically aggregates the data, usually by summing it. You can easily determine the aggregation applied to a field because the function always appears in front of the field's name when it is placed on a shelf. For example, **Gross Profit** becomes **SUM(Gross Profit)**.

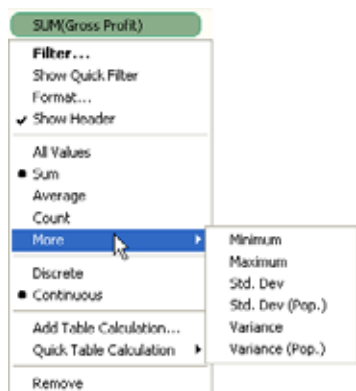
**Note** You can aggregate measures using Tableau only for relational data sources. Multidimensional data sources contain aggregated data only.

This section discusses the following topics:

- Aggregating Data
- Disaggregating Data
- Example – Aggregating and Disaggregating Data

### Aggregating Data

You can change the aggregation of a field by selecting a different function from the field's context menu. As shown below, all of the predefined aggregations are available from this menu.

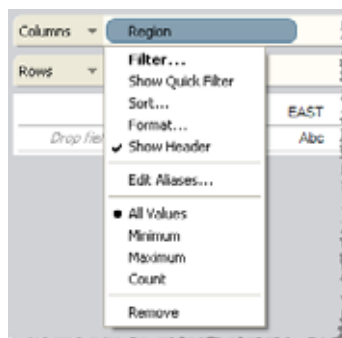


**Aggregating Measures.** You can assign a different aggregation to every measure you place on a shelf. For example, you can aggregate **Sales Total** as a summation, **Gross Profit** as a maximum, and **Discount** as an average.

You can change the aggregation state for all the measures on a worksheet by selecting the **Analysis > Aggregate Measures** menu item.



**Aggregating Dimensions.** When you aggregate dimensions, you create a new temporary measure column, so the dimension is now viewed as a measure. While you cannot apply any of the other predefined aggregations such as sum or average to a dimension, you can apply All Values, Minimum, Maximum, and Count.



## Disaggregating Data

Disaggregating your data allows you to view every row of the data source which can be useful when you are analyzing measures that you may want to use both independently and dependently in the view. For example, you may be analyzing the results from a product satisfaction survey with the Age of participants along one axis. You can aggregate the Age field to determine the average age of participants or disaggregate the data to determine at what age participants were most satisfied with the product.

---

**Note** If your data source is very large, disaggregating the data can result in a significant performance degradation. For tips on improving performance refer to Chapter 27, “Performance Tips”.

---

You can disaggregate all measures in the view by selecting **Analysis > Aggregate Measures**.

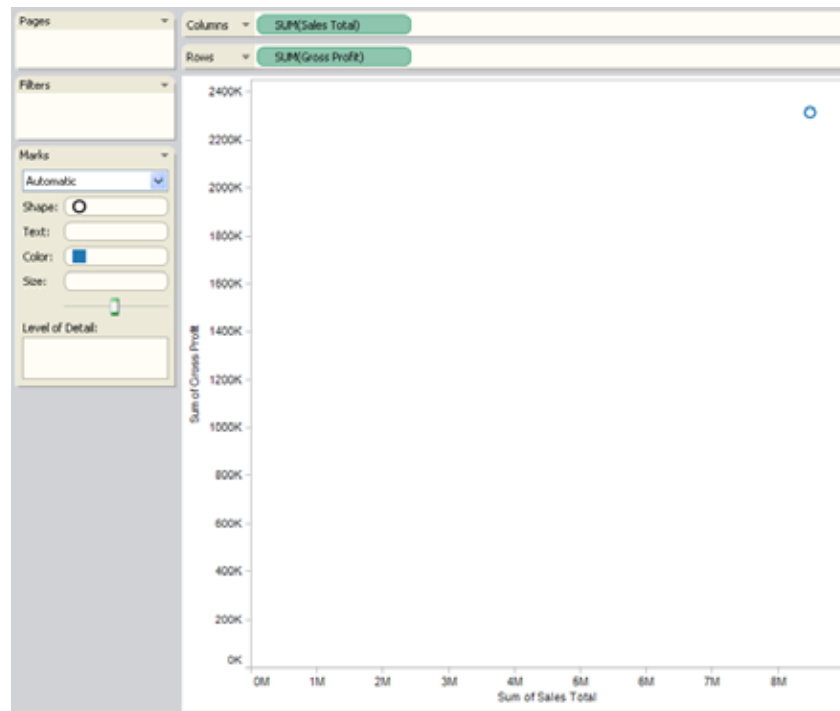


### Example – Aggregating and Disaggregating Data

This example includes several views of aggregated and disaggregated data created using the Sample – Superstore Sales data source. To create the views, follow these five steps:

- 1 Place the **Sales Total** measure on the **Columns** shelf and the **Gross Profit** measure on the **Rows** shelf.

The measures are automatically aggregated as sums. The aggregation is indicated by the field names and by the tooltip. The values shown in the tooltip are the sales total and the gross profit for the entire data source. That is, the summations are performed using every row in the data source.



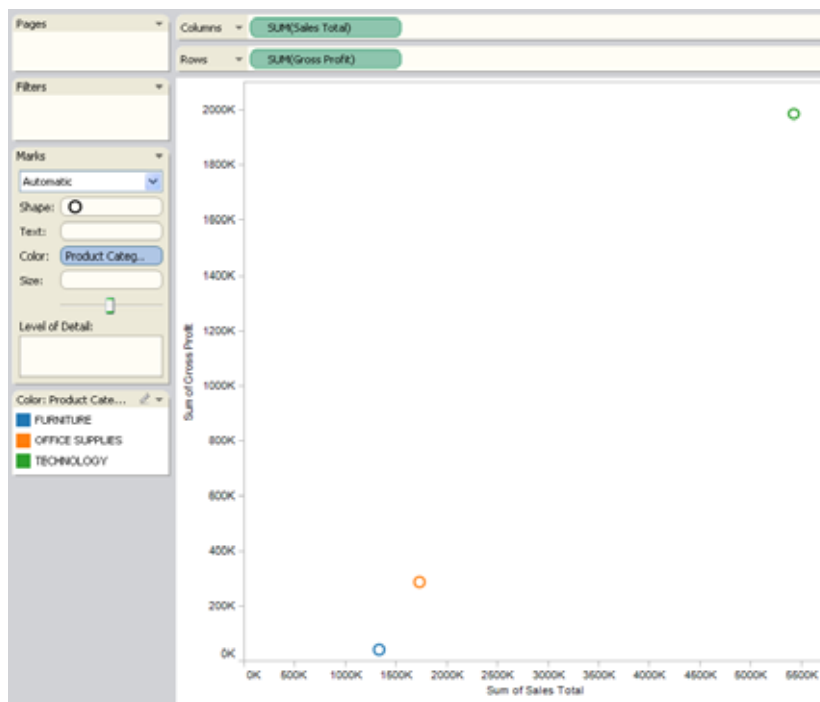
- 2 Place the **Product Category 1** dimension on the **Color** shelf.

One way to show more data in your view is to disaggregate the measures. Another way is to show additional levels of detail. For example, placing the **Product Category 1** dimension on the **Color** shelf separates the data into three marks—one for each dimension member—and then encodes the marks using color.

Although more marks are displayed, the measures are still aggregated. The single mark in the view indicates the sum of the sales and the sum of the gross profit for Office Supplies. If you were to sum the sales and gross profit values for the three marks, you

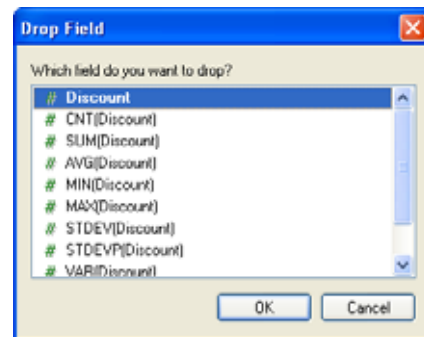


would produce the values for the entire data set as given in the previous step of this example.




- Place the **Discount** measure on the **Filters** shelf.

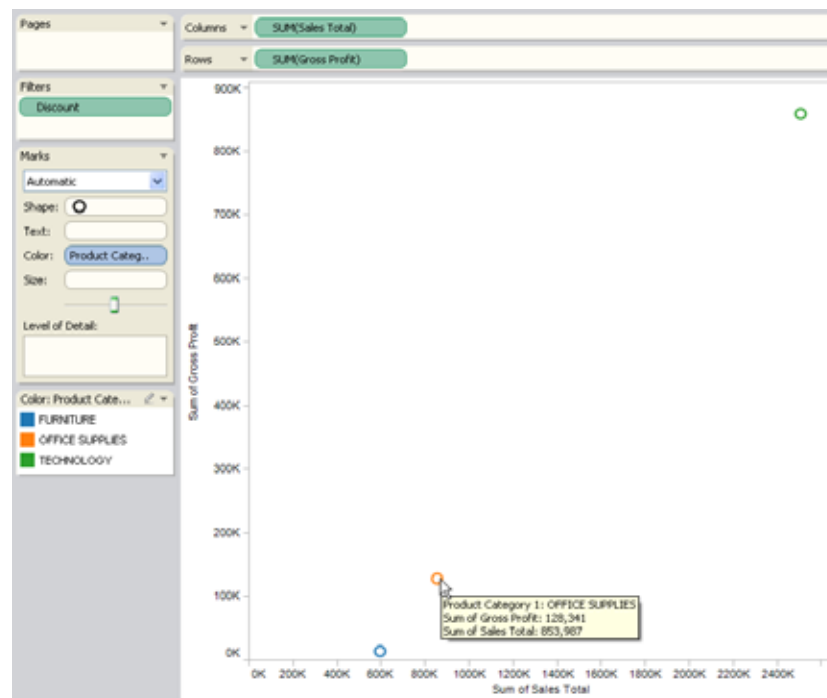
The **Filter Field** dialog box appears. Accept the default selection to filter the disaggregated measure.



Filter the data to only include discounts greater than 6% (0.06). Because **Discount** is disaggregated, Tableau applies the filter to each row in the data source before performing the aggregations for the **Sales Total** and **Gross Profit** measures.

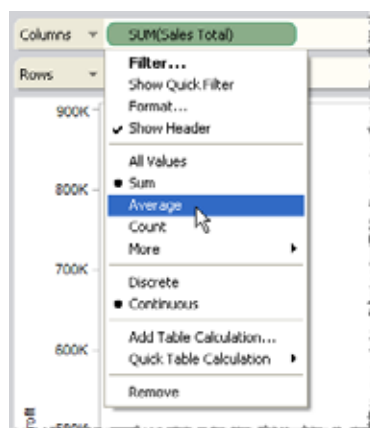


The view is shown below. The tooltip indicates that both the sales and the profit numbers are smaller than in the previous view. This is because data have been filtered out of the aggregation operation.

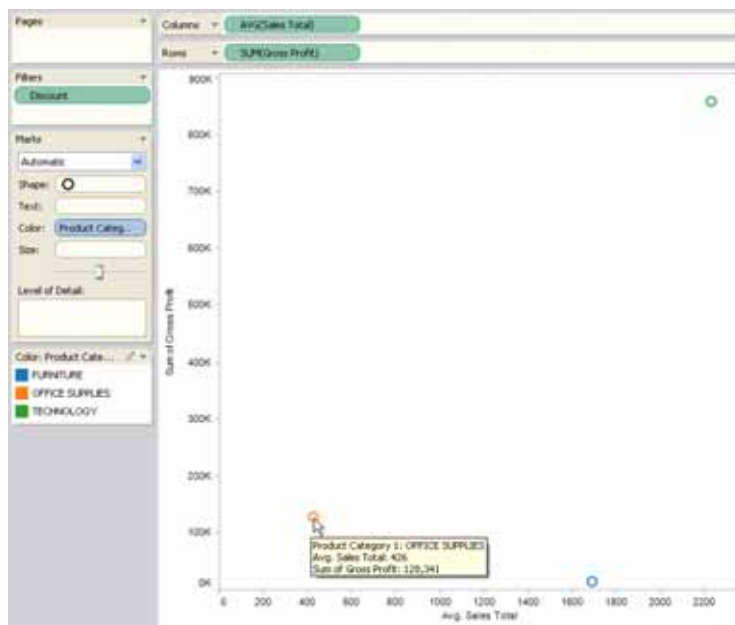


#### 4 Change the aggregation of **Sales Total** to an average.

The measures are not required to have the same aggregation. Change the aggregation by selecting **Average** from the field's context menu.




The view is shown below. The field name and tooltips indicate the new aggregation.



## 5 Disaggregate the data.

All measures—except those placed on the **Filters** shelf—must have the same aggregation state. That is, they must all be either aggregated or disaggregated.

You change the aggregation state by selecting the **Analysis > Aggregate Measures** menu item.



The view is shown below. Disaggregating the data displays every row in the data source that passes through the filter. The tooltip shows the profit and sales for one particular row.



## Calculated Fields

You might find that your data source doesn't include all of the fields needed to answer your questions. For example, you might want to create a new calculated field called **Profit** that is the difference between the **Sales** and the **Cost** fields, or you might want to create a conditional statement that divides the **Sales Budget** field into values that are under budget and values that are over budget.


Tableau allows you to create a new calculated field by defining a formula that is based on data source fields and other calculated fields, and that uses standard functions and operators.

This section discusses the following topics:

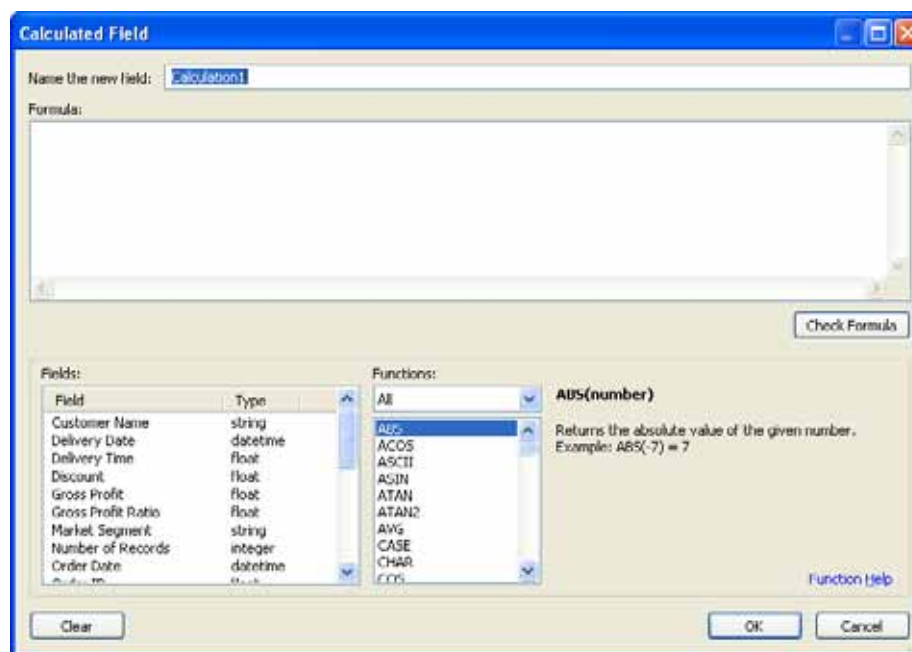
- How to Create a Calculated Field
- How to Create a Calculated Member
- Writing formulas in Tableau
- Example – Creating a Calculated Field
- Aggregate Calculations
- Example-Spotlighting Using Calculations

### How to Create a Calculated Field

To create a new calculated field, select **Analysis > Create Calculated Field**, or select **Create Calculated Field** on one of the Data window title menus. On a multidimensional data source, this option is only available from the Measures menu.



The **Calculated Field** dialog box opens.



To define the calculation do the following:

- 1 Specify a name for the new field.
- 2 Create a formula that defines the new field. Refer to “Writing formulas in Tableau” on page 16-22 for more information about how to define a formula.
- 3 Click **Check Formula** to verify that the formula is valid.
- 4 When finished, click **OK**.

The new calculated field displays in either the **Dimensions** area or the **Measures** area of the Data window depending on the data type returned by the calculation. Calculations that return a string or date are dimensions, while calculations that return a number are measures. In the latter case, you can convert the measure to a dimension if you want to treat the calculated values as discrete rather than continuous. Refer to “Converting Measures to Dimensions” on page 8-24 for more information about continuous and discrete fields.

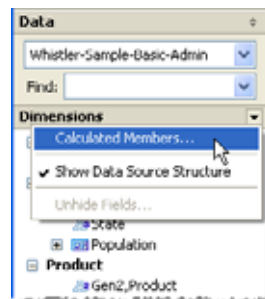


## How to Create a Calculated Member

If you are using a multidimensional data source, you have the option to create calculated members using MDX formulas instead of Tableau formulas. A calculated member can either be a calculated measure, which is a new field in the data source just like calculated fields, or a calculated dimension member, which is simply a new member within an existing hierarchy. For example, a Product dimension may have three members: Coke, Pepsi, and Coffee. You can define a new calculated member called “Colas” that sums the Coke and Pepsi members. Now when you place the Products dimension on the Rows shelf it displays four rows: Coke, Pepsi, Coffee, and Colas.


### Defining Calculated Members

You can define a calculated dimension member by selecting **Calculated Members** on one of the Data window title menus. The Calculated Members dialog box opens where you can create, delete, and edit any calculated members.

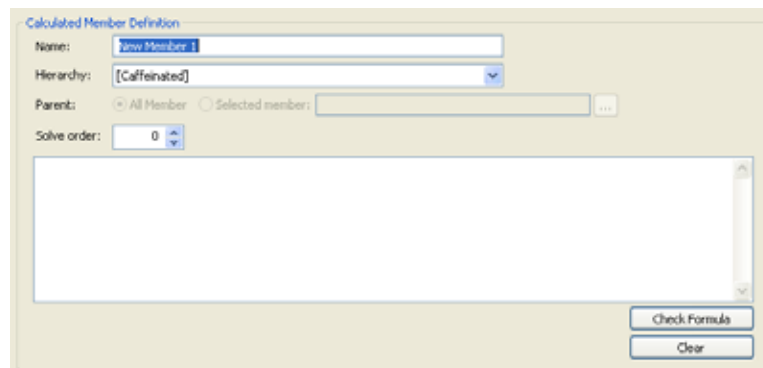


To create a new calculated member do the following:

- 1 Click **New** to add a new row to the list of calculated members at the top of the dialog box.



- 2 Type a **Name** for the new calculated member in the Member Definition area of the dialog box.



- 3 Select a hierarchy from the Hierarchy drop-down list. If you are creating a calculated measure, select **[Measures]** and define a result type using the **Result Type** drop-down list. Then skip to step five.
- 4 Specify the **Parent** member for the new calculated member. The **All Member** is selected by default, however, you can choose the **Selected Member** option to browse the hierarchy and select a specific parent member.
- 5 Give the new member a solve order.

Sometimes a single cell in your data source can be defined by two different formulas. The solve order defines the precedence given to each formula. Formulas with a lower

solve order are solved first. The default solve order is zero.

- 6 Type or paste an MDX expression into the **Formula** text box.
- 7 Click **Check Formula** to verify that the formula is valid.
- 8 When finished, click **OK**.

The new member displays in the Data window either as part of the Measures area, if you chose [Measures] as the parent member, or in the Dimensions area under the specified parent member. You can use the new member just like any other field in the view.

## Copying and Pasting Calculated Fields

Calculated fields are available to all sheets that use the same data source in a single workbook. In addition, you can copy and paste these custom fields between workbooks simply by right-clicking the field in the Data window and selecting **Copy**. Then in the new workbook, right-click the Data window and select **Paste**. You can copy and paste all custom defined fields such as calculated fields, ad-hoc groups, user filters, sets, and so on.

## Writing formulas in Tableau

When you are writing a calculation formula using Tableau formulas, make sure to include the following four parts:

## 1 Functions

The **Functions** area of the dialog box contains all the functions you can use to create a formula. The functions are organized into six categories, which are available from the drop-down menu. By default all functions are displayed.

You can display a brief description for each function by clicking its name in the list box. Double-click a function to include it in a formula. Refer to “Functions” on page 26-6 for a complete description of all functions.

## 2 Fields

All data source fields and calculated fields are listed in the **Fields** area of the dialog box. Binned fields and sets are not listed because they cannot be used in calculations.

The field names appear on the left and the field’s data type appears on the right. Refer to “Data Types” on page 8-20 to learn more about data types. You can click on the column heads to sort by name or by type.

Double-click a field name to include it in a formula. You can also just type the bare field name. However, if the field name includes special characters such as spaces, it must be delimited with square brackets as in SUM([Store Profit]). A right bracket (]) can be doubled to include it in the field name itself. For example, the field name “Store Profit]” would be written as [Store Profit]]] as in SUM([Store Profit]]]).

## 3 Operators

Operators are not available on the dialog box like functions and fields. Instead, you must manually type the operators into your formula. All standard operators such as addition (+), subtraction (–), multiplication (\*), and division (/) are supported. Refer to

“Operators” on page 26-25 for a complete description of the operators and the associated precedence rules.

#### **4 Comments (optional)**

You can insert custom comments for your calculations as a means of annotation for later review. To add a comment to a calculation type two forward slash characters into the formula pane.

For example:

Sales \* Profit //John’s calculation

In this example //John’s calculation is a comment.


A comment starts at the two forward slashes (//) and goes to the end of the line. A multiline comment can be written by starting each line with two forward slashes (//).

### **Example – Creating a Calculated Field**

In this example we will create a calculated field using Tableau formulas and use the new field in a data view. Then we’ll edit the field’s formula to create a new view, and finally delete the field from the Data window. This example uses the Sample - Superstore Sales (Excel) data source.

1 Create the view.

Select **New Calculated Field** on the Measures Data window title menu.

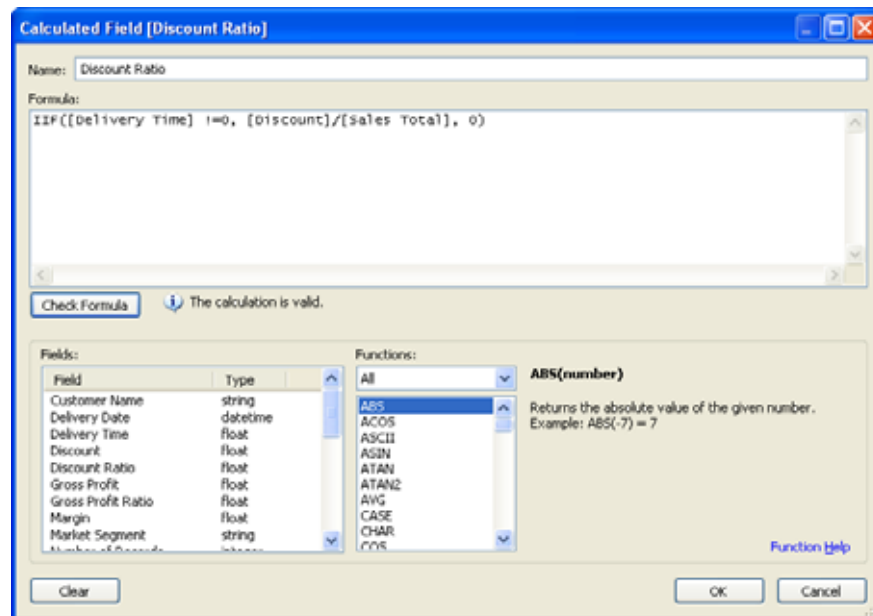


2 Complete the **Calculation** dialog box.


Name the new field **Discount Ratio** and enter the formula shown below.

`IIF([Delivery Time] !=0, [Disocunt]/[Sales Total],0)`

You can type the formula by double clicking the field names in the **Fields** list and functions in the **Functions** list. You must type the operators (!= and /) manually. Note that the IIF statement is used to avoid dividing by zero.



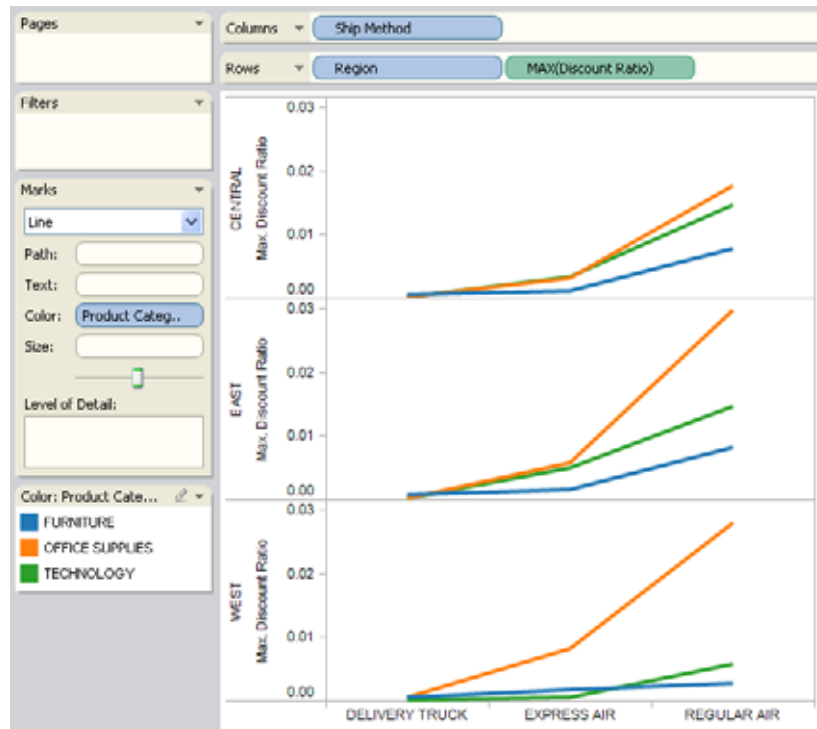
The new field displays in the **Measures** area of the Data window because the calculation returns a number. You can use this new field just like any other field.



### 3 Add the calculation to the view.

Place **Ship Method** on the columns shelf, **Region** on the Rows shelf, and **Product Category 1** on the Color shelf. Then place the new calculation, **Discount Ratio** onto the Rows shelf. Note that you can treat the new calculation just like any other measure. For example, you can apply an aggregation to it. Below, **Discount Ratio** is aggregated as a

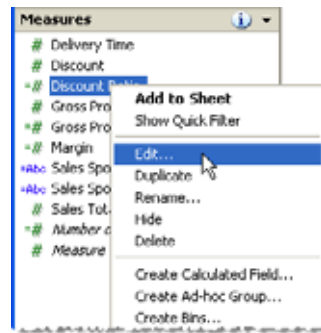
maximum. Make sure you select **Line** from the **Mark** menu to create a view like the one shown below.



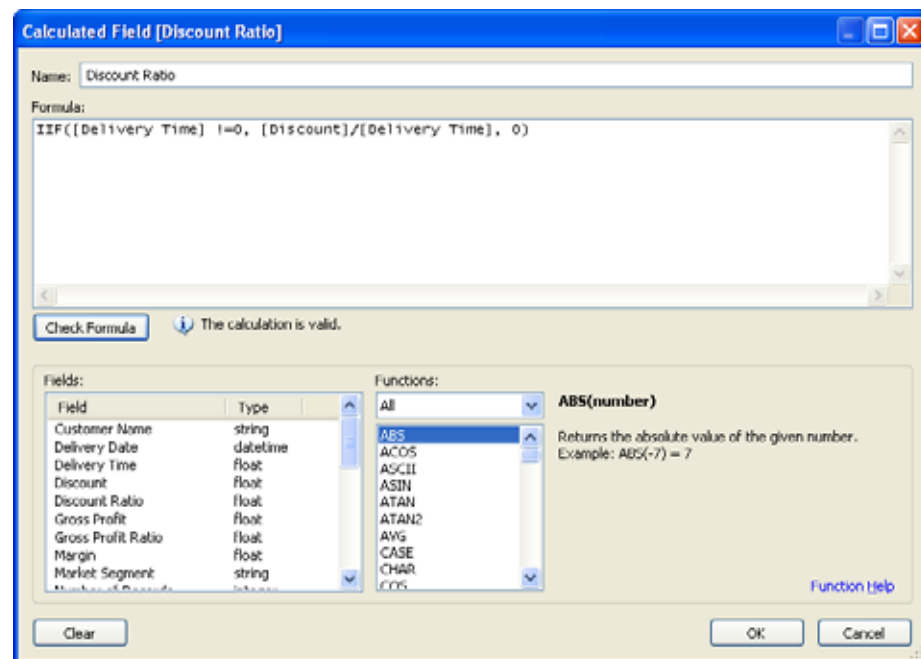
#### 4 Edit the calculation.

You can change the field's formula by right-clicking the field name in the Data window and selecting **Edit** or by selecting **Analysis > Edit Calculation**.

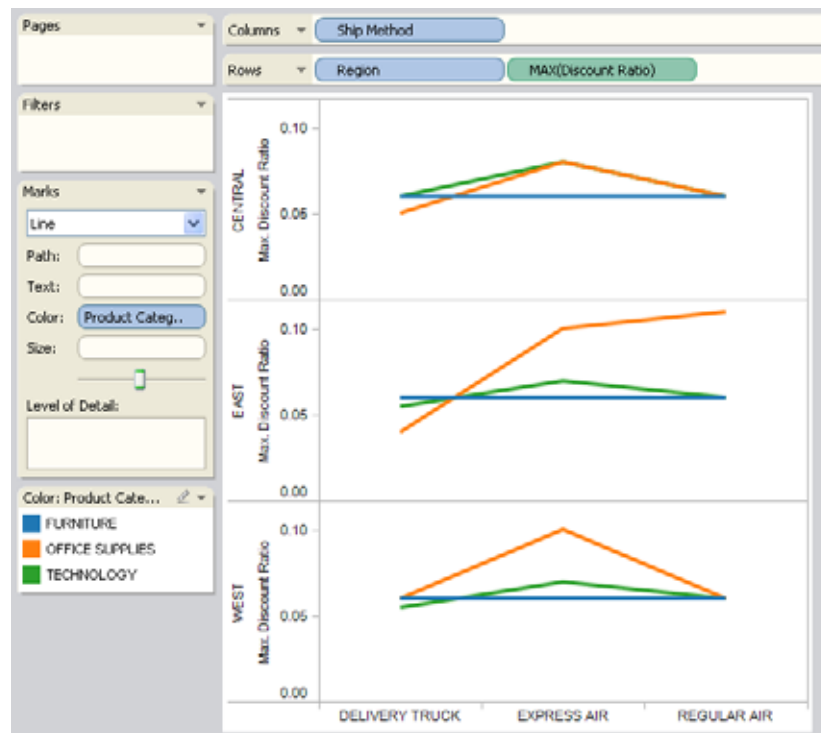




In the Calculated Field dialog box, change **Sales Total** to **Delivery Time**.



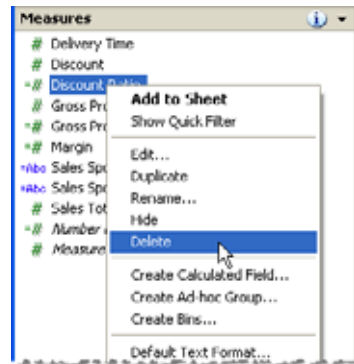
The view automatically updates after you click **OK** in the Calculated Field dialog box.



5 Delete the calculated field.

You can delete a calculated field by right-clicking the field name in the Data window and selecting **Delete** from the right-click context menu. Before deleting the field, you

might want to save your workbook. Refer to Chapter 24, “Saving and Exporting” for more information. If you do not save your work, the calculated fields will be lost.



## Aggregate Calculations

Aggregate functions allow you to summarize data. As described in “Aggregations” on page 16-4, Tableau includes a variety of predefined aggregations such as summation and variance. An *aggregate calculation* allows you to define aggregations other than these predefined choices.

This section discusses the following topics:

- About Aggregate Calculations
- How to Create an Aggregate Calculation
- Aggregate Calculations in a Disaggregated State
- Example – Aggregate Calculation

### About Aggregate Calculations

Suppose you want to analyze the overall gross margin for every product in your data source. One way to do this is to create a new calculated field called **Margin** that is equal to the gross profit divided by the sales total. Then you could place this measure on a shelf and use the predefined summation aggregation. In this scenario, **Margin** is defined as follows:

$$\text{Margin} = \text{SUM}(\text{Gross Profit} / \text{Sales Total})$$

This formula calculates the ratio of gross profit and sales for every row in the data source, and then sums the numbers. That is, the division is performed before the aggregation.

However, this is almost certainly not what you would have intended because summing ratios is generally not useful.

Instead, you probably want to know the sum of all gross profits divided by the sum of all sales. That formula is shown below.

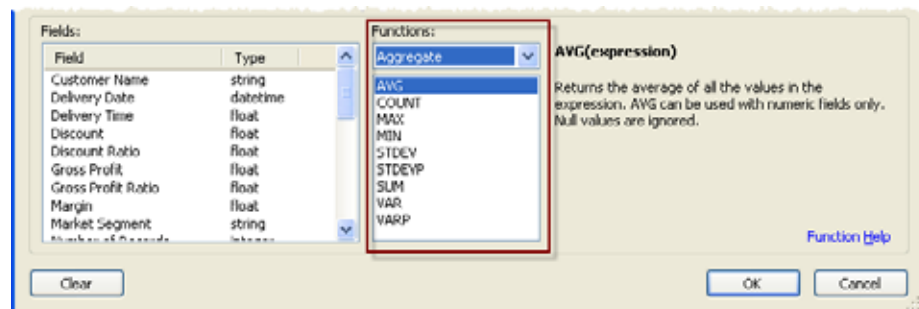
$$\text{Margin} = \text{SUM}(\text{Gross Profit}) / \text{SUM}(\text{Sales Total})$$

In this case, the division is performed after each measure is aggregated. An aggregate calculation allows you to create formulas like this.


### How to Create an Aggregate Calculation

When a calculation uses an aggregate function, it's called an aggregate calculation. You create an aggregate calculation by defining a new calculated field as described in “How to Create a Calculated Field” on page 16-18. The formula will contain one or more aggregate functions. You can easily pick an aggregate function from the **Calculation** dialog box by selecting **Aggregate** from the **Functions** menu as shown below.

These functions are identical to the predefined aggregate functions listed in “Aggregations” on page 16-4.



The aggregate calculation appears with the letters **AGG** in front of it when it is placed on a shelf.



When you create an aggregate calculation, no further aggregation of the calculation is possible. Therefore, the field's context menu does not offer any aggregation choices.

However, you can disaggregate the field. Refer to “Aggregate Calculations in a Disaggregated State” on page 16-32 for more information.

The rules that apply to aggregate calculations are:

- For any aggregate calculation, you cannot combine an aggregated value and a disaggregated value. For example,  $\text{SUM}(\text{Price}) * [\text{Items}]$  is not a valid expression because  $\text{SUM}(\text{Price})$  is aggregated and  $\text{Items}$  is not. However,  $\text{SUM}(\text{Price} * \text{Items})$  and  $\text{SUM}(\text{Price}) * \text{SUM}(\text{Items})$  are both valid.
- Constant terms in an expression act as aggregated or disaggregated values as appropriate. For example:  $\text{SUM}(\text{Price} * 7)$  and  $\text{SUM}(\text{Price}) * 7$  are both valid expressions.
- All of the functions can be evaluated on aggregated values. However, the arguments to any given function must either all be aggregated or all disaggregated. For example:  $\text{MAX}(\text{SUM}(\text{Sales}), \text{Profit})$  is not a valid expression because  $\text{Sales}$  is aggregated and  $\text{Profit}$  is not. However,  $\text{MAX}(\text{SUM}(\text{Sales}), \text{SUM}(\text{Profit}))$  is a valid expression.
- An aggregate calculation is always a measure.
- Like predefined aggregations, aggregate calculations are computed correctly for grand totals. Refer to “Grand Totals and Aggregations” on page 16-67 for more information.

### Aggregate Calculations in a Disaggregated State

If an aggregate calculation is disaggregated, the calculation is modified in a way that depends on the functions used. Every function has a disaggregated substitute, as shown below.

Aggregation Function	Disaggregated Substitute
AVG(data)	data
COUNT(data)	IIF(ISNULL(data),0,1)
COUNTD(data)	IIF(ISNULL(data),0,1)
MAX(data)	data
MIN(data)	data
STDEV(data)	Null
STDEVP (data)	IIF (ISNULL (data), Null, 0)
SUM(data)	data

Aggregation Function	Disaggregated Substitute
VAR(data)	Null
VARP (data)	IIF (ISNULL (data), Null, 0)


Note that STDEV and VAR are Null because those functions return Null if there are fewer than two elements in a group that are not Null, and each group has size 1 when it is disaggregated. Refer to “Aggregations” on page 16-4 for descriptions of the aggregation functions.

Therefore, if you define an aggregate calculation called Margin that is equal to SUM(Gross Profit)/SUM(Sales Total) and then disaggregate the data, it is interpreted as Gross Profit/Sales Total.

### Example – Aggregate Calculation

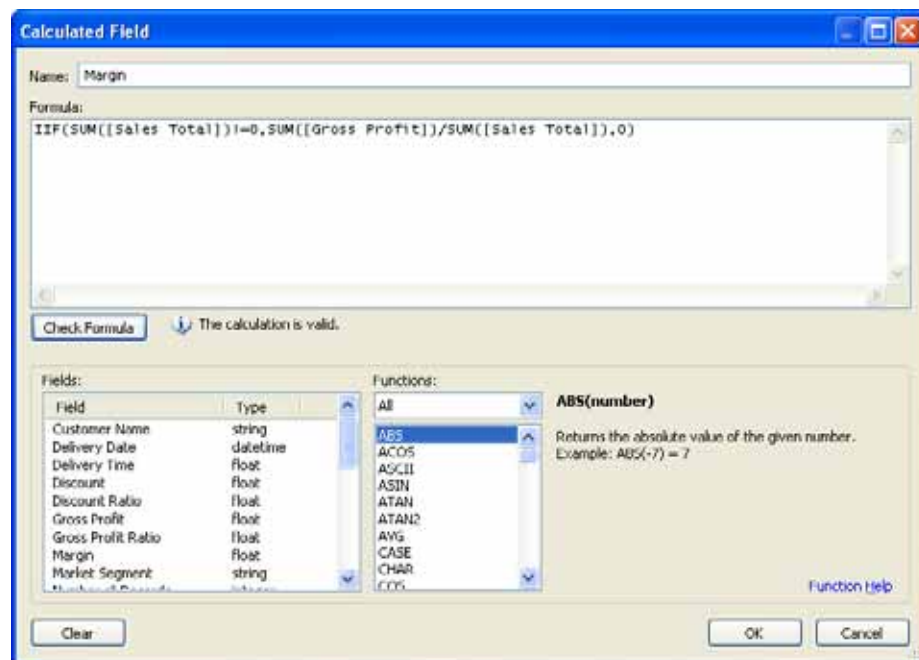
In this example you will use the Sample - Superstore Sales data source to create an aggregate calculation called Margin, and use the new field in a data view.

- 1 Select **New Calculated Field** on the Measures Data window title menu.




## 2 Define the calculation.

Name the new field **Margin** and enter the formula shown below. You enter the formula by selecting functions from the **Functions** area of the dialog box, and field names from the **Fields** area of the dialog box. You must type the operators (!= and /) manually. Note that the IIF statement is used to avoid dividing by zero.

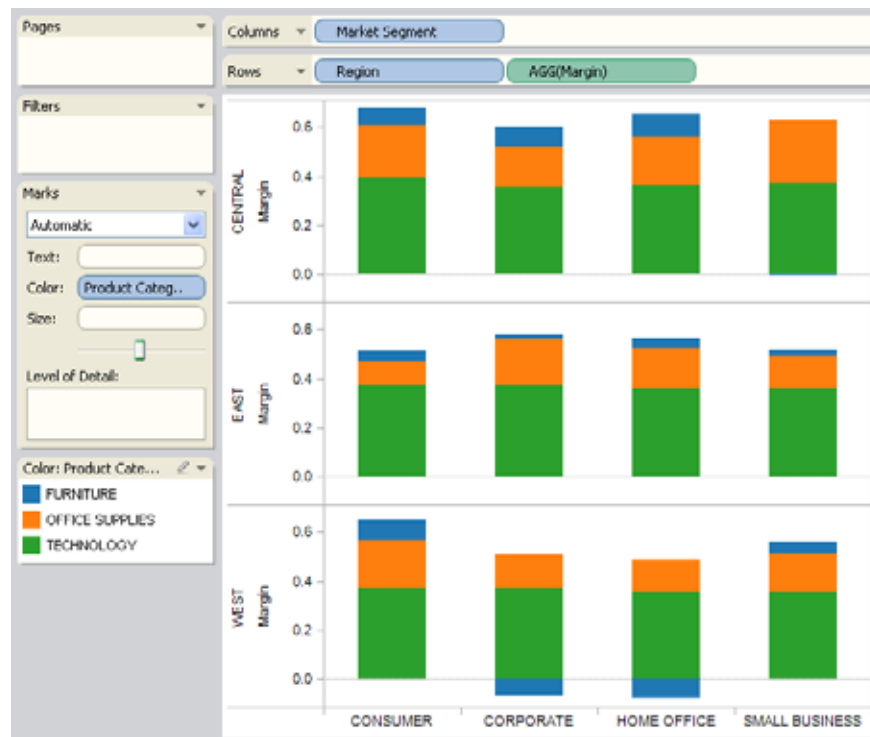


The new calculated field displays in the **Measures** area of the Data window where you can use it like any other measure.





A view using the new aggregate measure is shown below.



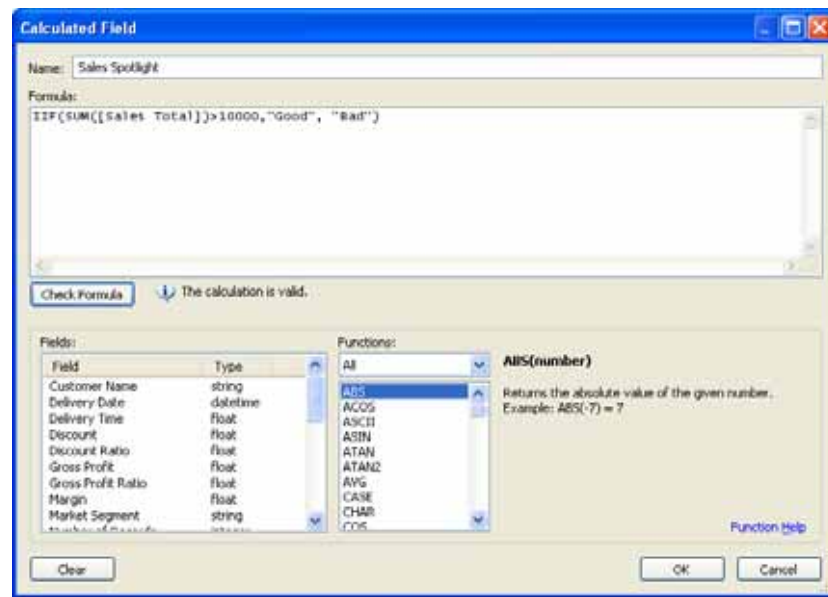
When **Margin** is placed on a shelf, its name is automatically changed to **AGG(Margin)**, which indicates it's an aggregate calculation. Additionally, the field's context menu does not include any aggregation choices because aggregating a field that's already aggregated is not possible.



### Example-Spotlighting Using Calculations

Spotlighting is a term that applies to a type of calculation. Specifically, a spotlighting calculation shows discrete thresholds based on the values of a measure. For instance, you might want to color-code sales totals so that those over 10,000 appear green and those below 10,000 appear red.

A spotlighting calculation is just a special case of a calculation that results in a discrete measure. A discrete measure is a calculation that is a dependant variable (and therefore a measure), but which results in a discrete result (as opposed to a continuous result). Thus the name discrete measure. Here is an example:



The formula in this example defines a discrete measure called “Sales Spotlight.” Discrete measures always appear with a blue “abc” icon in the Data window (green=continuous, blue = discrete). The example above is a measure because it is a function of another measure. It’s discrete because it produces discrete values (“Good” and “Bad”) as a result rather than continuous values like numbers. Here is an example of this categorical measure in use

The screenshot shows a Tableau interface with the following configuration:

- Columns:** Market Segment
- Rows:** Product Category 2
- Filters:** (Empty)
- Marks:**
  - Text: SUM(Sales To...)
  - Color: AGG(Sales Sp...)
  - Size: (Empty)
  - Level of Detail: (Empty)
  - Color: AGG(Sales Sp...)
  - Legend: Bad (Red), Good (Blue)

The resulting table is as follows:

	CONSUMER	CORPORATE	HOME OFFICE	SMALL BUSINESS
APPLIANCES	15,964	43,142	38,867	34,906
BINDERS AND BINDER AC..	28,209	62,260	36,206	31,453
BOOKCASES	31,691	48,234	31,459	33,811
CHAIRS & CHAIRMATS	142,716	200,180	175,111	82,616
COMPUTER PERIPHERALS	102,297	110,096	110,961	78,202
COPIERS AND FAX	100,077	98,953	56,968	73,244
ENVELOPES	18,310	21,848	7,873	30,662
LABELS	10,523	31,889	21,360	10,634
OFFICE FURNISHINGS	24,277	45,430	25,603	20,687
OFFICE MACHINES	51,306	136,683	49,318	33,975
PAPER	157,192	256,832	164,030	127,232
PENS & ART SUPPLIES	15,879	29,829	22,749	14,592
RUBBER BANDS	3,607	4,212	4,373	2,608
SCISSORS, RULERS AND ..	4,209	9,667	7,262	3,014
STORAGE & ORGANIZATI..	69,858	176,861	136,176	79,484
TABLES	81,262	133,606	143,717	100,641
TELEPHONES AND COMM..	852,242	1,658,104	1,162,667	734,273

Here the “sales spotlight” field is on the Color Shelf. It appears with the “AGG” prefix because it is an aggregate calculation. See the next section on “Aggregate Calculations” on page 16-30. Values above 10,000 and below 10,000 are assigned different colors. This type of discrete highlighting is often called spotlighting.

For a complete discussion of data types and data roles, including a discussion of discrete measures, see “Data Types & Roles” on page 8-20.

## Table Calculations

Another kind of calculation is a table calculation. Table Calculations are computations that are applied to the values in the entire table and are often dependent on the table structure itself. For example, in a sales environment, you can use table calculations to compute the running total of sales across a specified date range or to compare multiple months of sales and compute each month's contribution to the total sales.

All table calculations are computed locally using the values you see in the table. This means that computing a moving average for a measure that is aggregated as an average results in an average of the averages.

You can add table calculations to your view using either the predefined quick calculations or by specifying a custom definition. This section discusses the following topics:

- Quick Table Calculations
- Defining Table Calculations
- Advanced Table Calculations

### Quick Table Calculations

You can add common table calculations to your view using the Quick Table Calculations menu item on the field context menus. These quick calculations are predefined table calculations based on the most common scenarios.

#### To add a quick table calculation:

- 1 Right-click the measure you want to use in the table calculation and select Quick Table Calculation.
- 2 On the sub-menu select one of the following options:
  - Running Total
  - Difference
  - Percent Difference
  - Percent of Total
  - Moving Average
  - Year to Date (YTD) Total
  - Compound Growth Rate (CAGR)
  - Year over Year Growth

- Year to Date (YTD) Growth

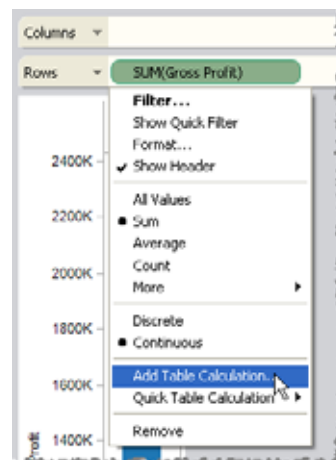
After adding the quick table calculation to the view, you can edit the definition by selecting Edit Table Calculation on the field's context menu.

## Defining Table Calculations

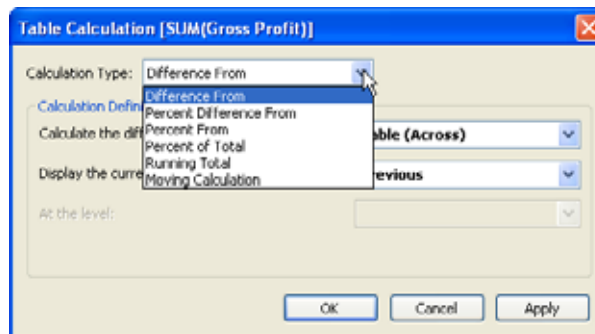
When you add a custom Table Calculation to the view, you need to specify the parameters that define the formula used in the computation. All of these parameters are set in the Table Calculation dialog box.

**To manually define a table calculation:**


- 1 Right-click the measure you want to use in the computation and select **Add Table Calculation**.



- In the Table Calculation dialog box, select one of the following types of calculations from the drop-down menu at the top:



- Difference From Calculation - compute the difference between two specified values.
  - Percent Difference From Calculation - compute the difference between two specified values as a percentage.
  - Percent From Calculation - compute each value as a percentage of another specified value.
  - Percent of Total Calculation - compute each value as a percentage of the total measure.
  - Running Total Calculation - compute cumulative totals along a specified dimension.
  - Moving Calculation - summarize a range of values using the specified aggregation to smooth short fluctuations and reveal long term trends.
- Define the formula using the drop-down lists in the bottom half of the dialog box. Learn more about how to define each type of calculation by selecting it in the list above.
  - When finished click OK. The measure is now marked as a table calculation and all the relevant values in the view are computed using the table calculation.



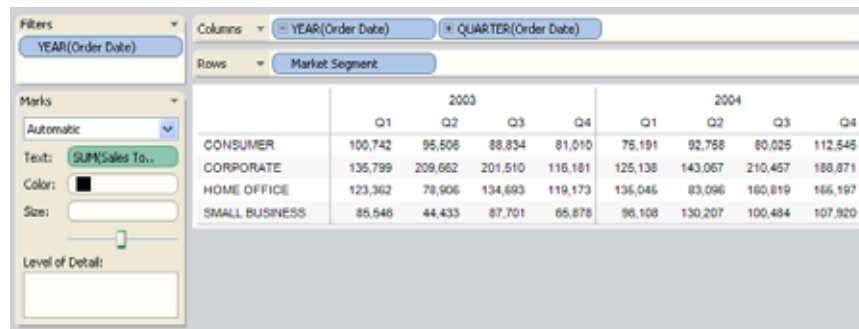
## Difference From Calculation

Use this type of calculation to compute the difference between two specified values in the table along a certain dimension. For example, compute the difference between 2003 and 2004 quarterly sales for four different market segments.

When defining a Difference From calculation, you need to specify a dimension or table structure to compute across, the dimension level to use in the computation (this is only required if you are computing across a dimension), and a value to compare the current value to. The following is an example of a Difference From calculation.

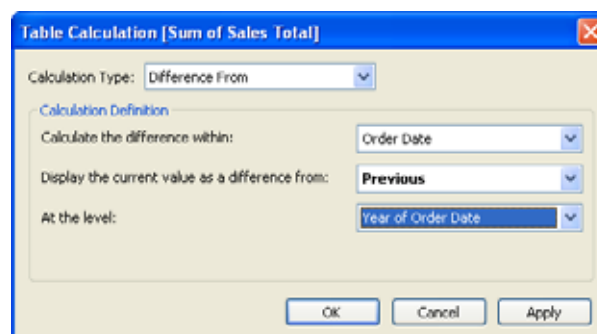
### Example: Difference From Calculation

The table below shows the 2003 and 2004 quarterly sales numbers for several different market segments of a superstore.



	2003				2004			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
CONSUMER	100,742	95,508	88,834	81,010	75,191	92,758	80,025	112,545
CORPORATE	135,799	209,662	201,510	116,181	125,138	143,067	210,457	168,871
HOME OFFICE	123,362	78,906	134,693	119,173	136,046	83,096	160,819	166,197
SMALL BUSINESS	85,546	44,433	87,701	65,878	96,108	130,207	100,484	107,920

To compute the difference between 2003 and 2004 sales, you can define a table calculation using the definition shown below.



**Table Calculation [Sum of Sales Total]**

Calculation Type: **Difference From**

**Calculation Definition**

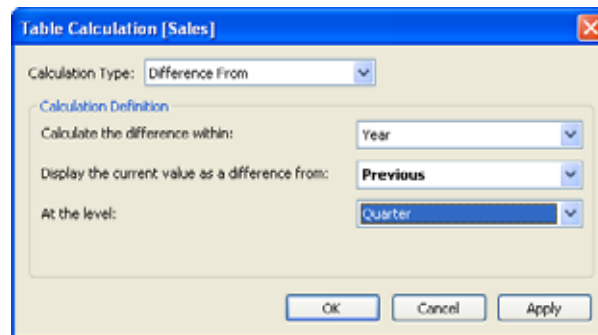
Calculate the difference within: **Order Date**

Display the current value as a difference from: **Previous**

At the level: **Year of Order Date**

OK Cancel Apply





The difference is calculated along the Order Date dimension at the level of year because we are comparing 2004 sales to 2003 sales. The table now displays the difference between each quarter in 2004 and the corresponding quarter in the previous year. Notice that there are no values for 2003. That's because there are no previous years to compute the difference from.

	2003				2004			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
CONSUMER					-25,551	-2,748	-8,808	31,536
CORPORATE					-10,662	-66,595	8,947	72,690
HOME OFFICE					11,683	4,191	26,125	45,024
SMALL BUSINESS					10,563	85,774	12,784	42,042

The view below may be more clear. It shows both the Difference From calculation and the Total Sales (before the computation). You can see that in the first quarter of 2004 the total sales was \$75,191 while in the same quarter in 2003 the total sales was \$100,742. The difference between these two values is -\$25,551. You can see that in

Compute the difference between the Q1 sales in 2003 and 2004.

The result is shown here.

		Q1		Q2		Q3		Q4	
		Sales Total	Diff from previous	Sales Total	Diff from previous	Sales Total	Diff from previous	Sales Total	Diff from previous
CONSUMER	2003	100,742		96,606		88,834		81,010	
	2004	75,191	-25,551	92,758	-2,748	80,025	-8,808	112,545	31,535
CORPORATE	2003	136,799		209,662		201,510		116,181	
	2004	125,139	-10,662	143,067	-66,595	210,457	8,947	188,871	72,690
HOME OFFICE	2003	123,362		78,906		134,693		119,173	
	2004	136,046	11,683	83,096	4,191	160,819	26,126	166,197	46,024
SMALL BUSINESS	2003	95,546		44,433		97,701		65,878	
	2004	96,108	10,563	130,207	85,774	100,484	12,784	107,920	42,042

**Note** You can add a Difference From calculation to your view quickly by right-clicking the measure you want to use in the computation and selecting **Quick Table Calculations > Difference**. This quick calculation computes the difference between values across rows where each difference is calculated against the previous value. Refer to “Quick Table Calculations” on page 16-40.

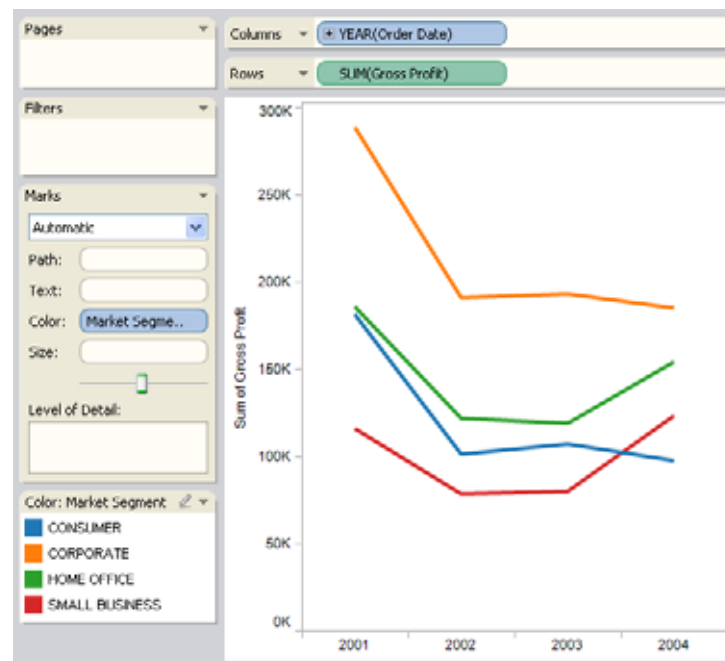
### Percent Difference From Calculation

Use this type of calculation to display the rate of change between two specified values in the table by computing the difference as a percentage. A common use of this type of calculation is to compute the percent gain year after year (CAGR).

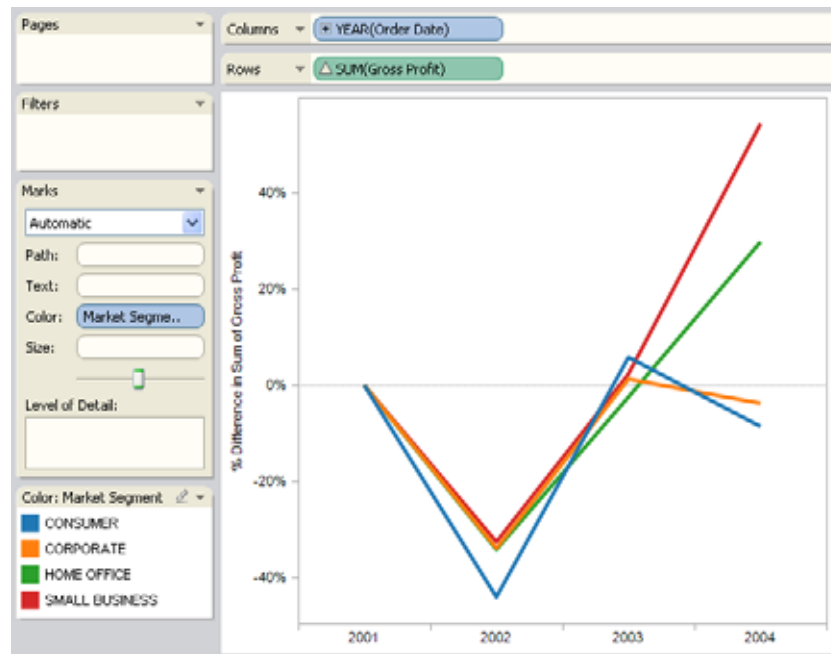
Similar to the Difference From calculation, to define a Percent Difference From calculation you need to specify a dimension or table structure to compute across, the dimension level to use in the computation (this is only required if you are computing across a dimension), and a value to compare the current value to. The following is an example of a Percent Difference calculation.

#### Example: Percent Difference From Calculation

The table below shows the gross profit for several markets over four years. Looking at the view, we can see that there was a drop in profit in all markets during 2002.



However, when we view this same view using a Percent Difference From calculation, it becomes clear that the consumer market segment had a much more drastic drop than the other markets.



You can define a Percent Difference From calculation like this using the table calculation definition shown below.

**Table Calculation [SUM(Gross Profit)]**

Calculation Type: Percent Difference From

Calculation Definition

Calculate the percent difference within: Order Date

Display the current value as a percentage of: Previous


At the level: YEAR(Order Date)

☐ Compute compounded rate

OK Cancel Apply

The difference is calculated along the Order Date dimension at the Year level because we are comparing year after year profit. Each value in the view is a difference of the previous year. The view below shows each year as a difference of 2001 Profit. You can see that the

Small Business Market is the only segment to recover from the drop in profit over the past three years.



**Note** Percent Difference From calculations are commonly used to calculate compound growth rates and year over year growth. You can quickly add these calculations by right-clicking the measure you want to use in the calculation and selecting the calculation on the Quick Table Calculation sub-menu. Refer to “Quick Table Calculations” on page 16-40.


### Percent From Calculation

A Percent From calculation is similar to the Percent Difference From calculation in that you can use it to compute the change between two values as a percentage. However, this type of calculation computes an absolute change. For example, use the Percent From calculation to compare the sales performance of several products.


When you define a Percent From calculation you need to specify a table structure or dimension to calculate the percentage from. If you select a dimension, you also need to select a level. Finally, you need to select a value that each value in the table will be displayed as a percentage of. The following is an example of a Percent From calculation.

### Example: Percent From Calculation

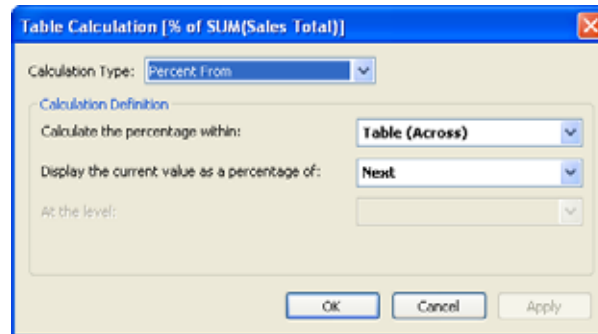
The view below shows the sales of the top thirty products by profit at a superstore. You can see that the top selling product is the Canon PC940 Copier.



While it is generally clear how each product is performing when compared to the next in terms of sales, you can use a Percent From calculation to compute just how much better each product is from the next. The view below has this type of calculation placed on the text shelf. Each mark is labeled with the percentage it is of the next product in the list. You can now see that the Canon Copier is 185% better in terms of sales than the Stacker Chairs.



This table calculation was computed using the definition shown in the dialog box below. The percentage is calculated across the rows and each product is displayed as a percentage of the next product.



### Percent of Total Calculation

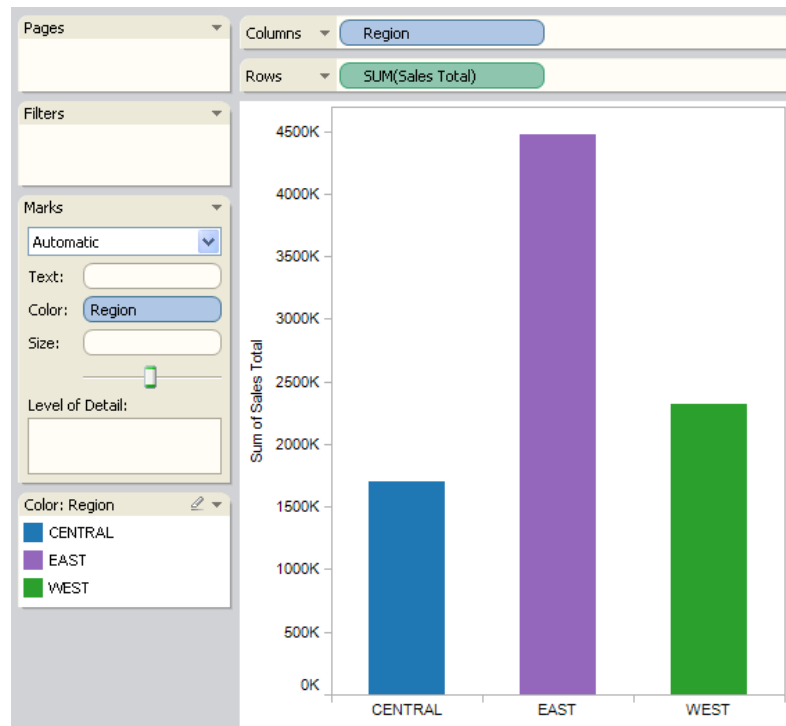
Use the Percent of Total calculation to compute the percentage a specified value contributes to the total. For example, you can use this type of calculation to view the contribution each member in a sales team makes to the total company sales each quarter.

For a Percent of Total calculation you must specify a dimension or table structure to compute across and the dimension level to use in the computation. The dimension level is only necessary if you have chosen to compute across a dimension rather than a table structure. The following is an example of a Percent of Total calculation.

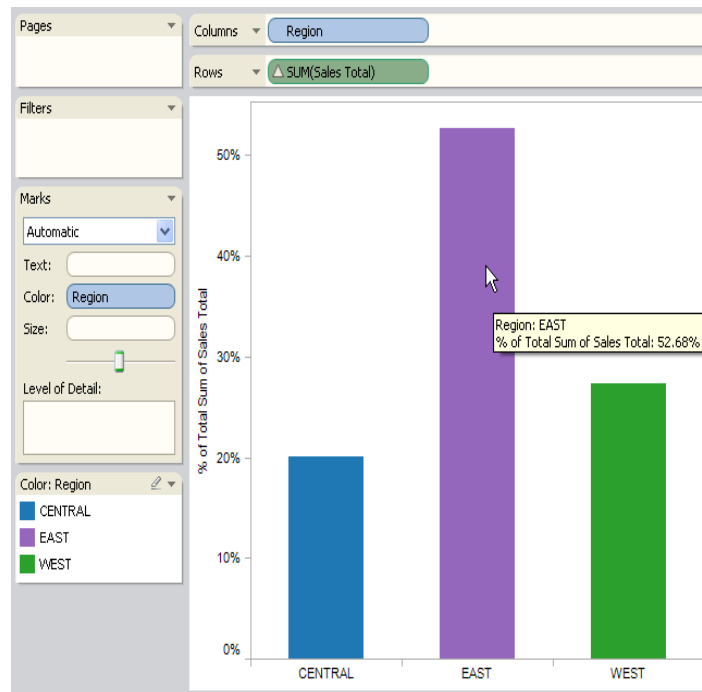
#### Example: Percent of Total Calculation

In this example, imagine that we are analyzing the performance of several stores. The view below shows the sales performance for three different regions. From the view we can see that the Eastern region (purple marks) has the highest sales.

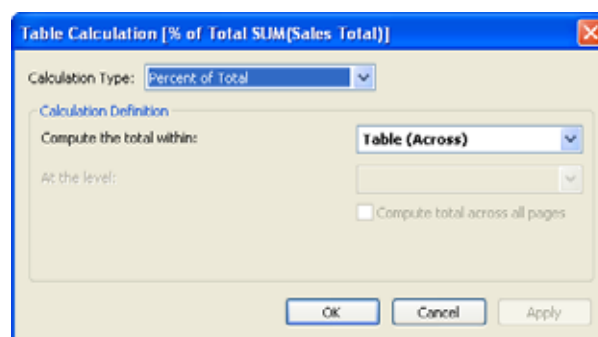




When we add the Percent of Total table calculation to the view, we can see that the East region accounts for just over 52% of the total sales.



The table calculation was computed using the definition shown below. The total was computed across each row in the table.



---

**Note** You can add a Percent of Total calculation quickly to your view by right-clicking the measure you want to use in the computation and selecting Percent Total on the Quick Table Calculations sub-menu. Refer to “Quick Table Calculations” on page 16-40.

---

### Running Total Calculation

Use the Running Total calculation to compute a cumulative total across a dimension or table structure. For example, you can use this type of calculation to calculate the cumulative sales totals for each quarter for several years.

When you define a running total calculation, you need to specify an aggregation to use when summarizing the values. For example, the most common aggregation will be sum so you can see the summation of values, but sometimes you may want to use average or another aggregation. You also need to specify the dimension to compute a running total across. This can either be an actual dimension in the data source or a table structure like rows or columns. Finally, you need to specify when to restart the at zero and begin totaling again. The following is an example of a Running Total calculation.

**Example - Running Total Calculation.** The view below shows the total quarterly sales from 2001 to 2004.

Columns: + QUARTER(Order Date)				
Rows: + YEAR(Order Date)				
	Q1	Q2	Q3	Q4
2001	725,234	615,250	699,985	761,852
2002	497,079	496,466	613,387	399,683
2003	446,449	428,507	512,737	382,242
2004	431,482	449,128	551,785	574,534

While it is useful to see each quarter's sales, you may also want to see the cumulative totals for each quarter in the year. To create this kind of view we can add a Running Total calculation. The view below shows the running totals for each quarter restarting at zero for each year. That means that the Quarter 4 shows the total sales for that year.

Columns		+ QUARTER(Order Date)			
Rows		+ YEAR(Order Date)			
		Measure Names			
		Q1	Q2	Q3	Q4
2001	Sales Total	725,234	616,260	699,906	781,882
	Running Total	725,234	1,340,493	2,040,400	2,802,320
2002	Sales Total	497,079	496,466	613,387	399,683
	Running Total	497,079	993,533	1,506,921	1,906,503
2003	Sales Total	445,449	428,507	512,737	382,242
	Running Total	445,449	873,956	1,386,693	1,768,935
2004	Sales Total	431,482	449,120	551,785	574,534
	Running Total	431,482	880,609	1,432,395	2,006,929

The SUM of Q1 and Q2 sales is show here.

This calculation was defined by the formula shown below. We are summarizing values as a sum along the Order Date dimension restarting at zero every Year.

**Table Calculation [Running Sum of SUM(Sales Total)]**

Calculation Type: **Running Total**

Calculation Definition

Summarize values using: **Sum**

Running within: **Order Date**

Restarting at zero for every: **YEAR(Order Date)**

☐ Perform a secondary calculation on the result

OK Cancel Apply

**Note** You can add a Running Total calculation to your view easily using the Quick Table Calculations menu. Right-click the measure you want to use in the calculation and select **Quick Table Calculations > Running Total**. Refer to “Quick Table Calculations” on page 16-40.


## Moving Calculation

A moving calculation is typically used to smooth short term fluctuations in your data so that you can see long term trends. A good example is when you are looking at securities data. There are so many fluctuations every day that it is hard to see the big picture through the


daily ups and downs. You can use a moving calculation to define a range of values to summarize using an aggregation of your choice.

When you define a moving calculation you must first specify the aggregation you want to use when summarizing that data. The most common aggregation for this type of calculation is an average. Next you need to specify the dimension to summarize across. You can select a table structure such as Rows or Columns or an actual field in your data source. Once you have selected a dimension, define the number of values before the current value and the number of values after the current value to include in the summary. You can also decide whether to include the current value using the checkbox on the right. The following is an example of a Moving Calculation.

**Example - Moving Calculation.** The view below shows the the discounts given at a superstore along a continuous date axis. As you can see, it is very difficult to see any kind of trend in this view.



However, if you add a Moving Average, the view becomes much more manageable.



This calculation was defined by the formula shown below. The values are summarized as an average along the rows in the view. Each value is an average of the seven days surrounding the current value (four days before and three days after). Note that we have opted to include the current value.

**Table Calculation [Moving Average of SUM(Discount)]**

Calculation Type: **Moving Calculation**

Calculation Definition

Summarize values using: **Average**

Moving within: **Table (Across)**

Previous Values: **4** Next Values: **3** ☒ Include current value

☐ Null if there are not enough values

☐ Perform a secondary calculation on the result

OK Cancel Apply

**Note** You can add a Moving Average to your view quickly using the Quick Table Calculations menu. Right-click the measure you want to use in the calculation and select **Quick Table Calculations > Moving Average**. By default this quick calculation will add a moving average across the rows in the view, summarizing the previous two values including the current value. Refer to “Quick Table Calculations” on page 16-40.

## Advanced Table Calculations

Table calculations can be very useful when you want to perform a calculation that applies to all of the data in the table. Most of the time you will only need to add a single calculation such as Difference From or Running Total. However, you may sometimes want to combine two calculations so that you perform one and then perform the next on the results. For example, when calculating the Year to Date Growth, you first need to calculate the cumulative totals and then calculate the percent difference each total is from the previous year. You can add a secondary calculation to Running Totals and Moving Calculations by selecting **Perform secondary calculation on the result** in the Table Calculation dialog box.

Table Calculation [SUM(Gross Profit)]

Calculation Type: Running Total

Calculation Definition

Summarize values using: Sum

Running within: Table (Across)

Restarting at zero for every:

☒ Perform a secondary calculation on the result

Secondary Type: Percent Difference From

Secondary Calculation Definition

Calculate the percent difference within: Table (Across)

Display the current value as a percentage of: Previous

At the level:

☐ Compute compounded rate

OK Cancel Apply

## Binned Data

Sometimes it's useful to organize the values of a measure into bins. For example, suppose you have a measure that holds the ages of customers ranging from 18 to 90. If you wanted to analyze how customer value breaks down by different age groups, you would bin the data. Also, to create a histogram you must first bin data.

In Tableau, bin data by highlighting a numeric dimension or measure in the Data window and selecting **Create Bins** from the context menu.

---

**Note** You can bin data only for relational data sources. This feature is not supported for multidimensional data sources.

---

When you bin a measure you create a new dimension. That's because you are creating discrete categories out of a continuous range of values. The following example walks you through creating a histogram using binned data.

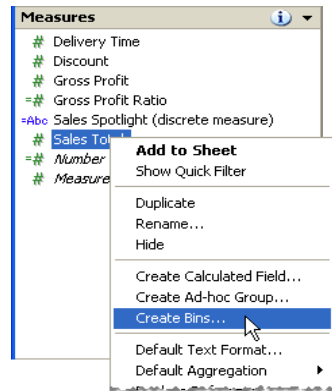
### Example – Creating a Histogram with Binned Data

Histograms are one way to display the distribution of values in a field. In Tableau, you can create a histogram by binning the values of a measure and then creating a view based on the measure and its binned values. This example uses the Sample - Superstore Sales data source.

**To create a histogram based on binned data:**



- 1 Select the Sales Total measure in the Data window and select Create Bins on the right-click context menu.



- 2 Complete the Create Bins dialog box.

When you bin a measure, you create a new field. The new field is a binned version of the original field. Specify the name of the new field and the size of each bin. To help you

determine the best bin size, press the **Load** button to display the range of values of the measure.

The binned field appears in the **Dimensions** area of the Data window because the bins are treated as discrete categories.



- 3 Place the **Sales Total** measure on the **Rows** shelf.

The measure is automatically aggregated as a summation, and an axis is created with a label given by the field name.

- 4 Place the **Sales Total (bin)** dimension on the **Columns** shelf.

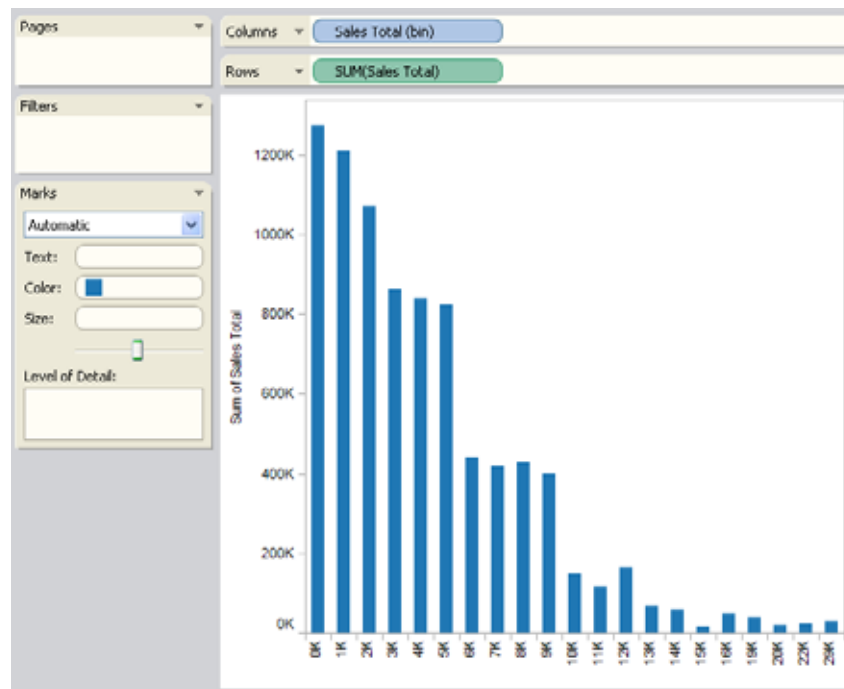
Row headers are created with labels given by the dimension member names.

---

**Note** Notice that all bins are of equal size. If you want to create variable sized bins, you can create a calculation using the using the CASE function. Refer to “Logical Functions” on page A-17.

---

The view is shown below.



Each bin acts as an equal-sized container that summarizes data for a specific range of values. Each bin label designates the lower limit of the range of numbers that is assigned to the bin. Note that the lower limit is inclusive. For example, the bin labelled **1K** contains numbers greater than or equal to 1,000, but less than 2,000.

**Note** This example shows how to build a histogram manually. You can also create a histogram automatically. Do this by (1) selecting a measure in the Data window; (2) clicking the “Show Me! Alternatives” button on the toolbar; (3) selecting the histogram option.

## Totals

You can automatically compute grand totals and subtotals for the data in a view. By default Tableau uses the underlying data to compute totals. However, if you are using a Multidimensional data source you can specify whether to do compute the total on the server using the underlying data or in locally using the data that you see in the table. The section discusses the following topics:

- Local vs. Server Computation
- Default Aggregation
- Grand Totals
- Subtotals

### Local vs. Server Computation

If you are using a multidimensional data source, you can specify whether to do the subtotal or total computation on the server using the underlying data in the data source or locally in Tableau using the data that you see in the table.

The default setting is to compute all totals on the server if you are connected to an SSAS data source and locally if you are connected to an Essbase data source using the aggregation specified in the cube. However, there are cases when it is not possible for the server to compute the expected subtotals due to filtering or perfect pivoting.

For example, let's say you have a view showing the sales of Amaretto, Columbian, and Decaf Irish Cream coffees. Then you filter the view to only show Amaretto and Columbian coffee sales. When you turn on subtotals for the Product Type field so you can see the total sales for all Coffees, one of the following will happen:

- If the total can be computed using the filter, the correct total will display. In this case Tableau would compute the total sales for all Coffees.
- If the total cannot be computed using the filter, the Totals cells in the view will be empty. In this case you would want to specify a local computation that only includes the values you see in Tableau.

#### To specify a local computation:

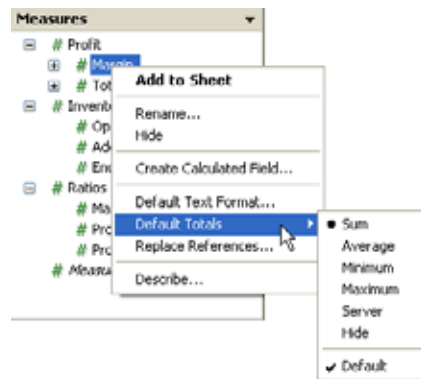
- 1 Select **Total Using** on the context menu of the measure you are using to calculate the subtotals.
- 2 Then select an aggregation to use.

## Default Aggregation

You can also specify a default aggregation for any measure. The default aggregation will be used automatically when the measure is first totaled in the view.

**To specify a default aggregation:**

- 1 Right-click any measure in the Data window and select **Default Totals**.
- 2 On the Default Totals list, select an aggregation.



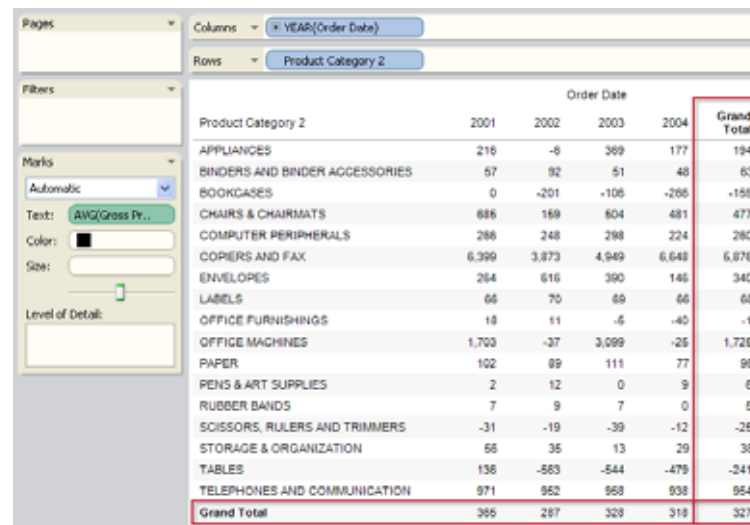
Whether you are specifying the aggregation for a field on a shelf or the default aggregation in the Data window, you can select from the following options:

Default	<p>For SSAS data sources, this option computes the subtotal on the server.</p> <p>For Essbase data sources, this option computes the total using the default aggregation determined by the data type (typically SUM).</p>
Sum	Displays the sum of all shown values.
Average	Displays the average of all the shown values.
Minimum	Displays the smallest shown value.
Maximum	Displays the largest shown value.

Server	Computes the subtotal on the server.
Hide	Hides all subtotals.

## Grand Totals

Any view in Tableau can include grand totals. For example, in a view showing the average gross profit for each product and year, you can turn on grand totals to also see the average gross profit for all products and all years.



Product Category 2	Order Date				Grand Total
	2001	2002	2003	2004	
APPLIANCES	216	-6	369	177	194
BINDERS AND BINDER ACCESSORIES	57	92	51	48	63
BOOKCASES	0	-201	-106	-266	-159
CHAIRS & CHAIRMATS	686	169	604	481	477
COMPUTER PERIPHERALS	286	248	298	224	280
COPERS AND FAX	6,399	3,873	4,949	6,643	6,876
ENVELOPES	264	616	390	146	340
LABELS	66	70	69	66	68
OFFICE FURNISHINGS	18	11	-6	-40	-1
OFFICE MACHINES	1,703	-37	3,099	-25	1,726
PAPER	102	89	111	77	95
PENS & ART SUPPLIES	2	12	0	9	6
RUBBER BANDS	7	9	7	0	5
SCISSORS, RULERS AND TRIMMERS	-31	-19	-39	-12	-25
STORAGE & ORGANIZATION	56	35	13	29	38
TABLES	136	-563	-544	-479	-241
TELEPHONES AND COMMUNICATION	971	952	958	938	954
<b>Grand Total</b>	<b>366</b>	<b>267</b>	<b>328</b>	<b>318</b>	<b>327</b>

## How to Turn on Grand Totals

You can calculate grand totals by selecting one of the **Grand Totals** options on the **Table** menu. The grand totals are added as an additional row or column to your table.



The following rules dictate whether you can turn on grand totals:

- The view must have at least one header – Headers are displayed whenever you place a dimension on the **Columns** shelf or the **Rows** shelf. If column headers are displayed, you can calculate grand totals for columns. If row headers are displayed, you can calculate grand totals for rows.
- Measures must be aggregated – The aggregation determines the values displayed for the totals. Refer to “Grand Totals and Aggregations” on page 16-67 for more information.
- Grand Totals cannot be applied to continuous dimensions.

The figure below is a text table with grand totals for both rows and columns.

Pages

Columns

Region

Rows

Product Category 2

Filters

Marks

Automatic

Text: SUM(Gross Pr..)

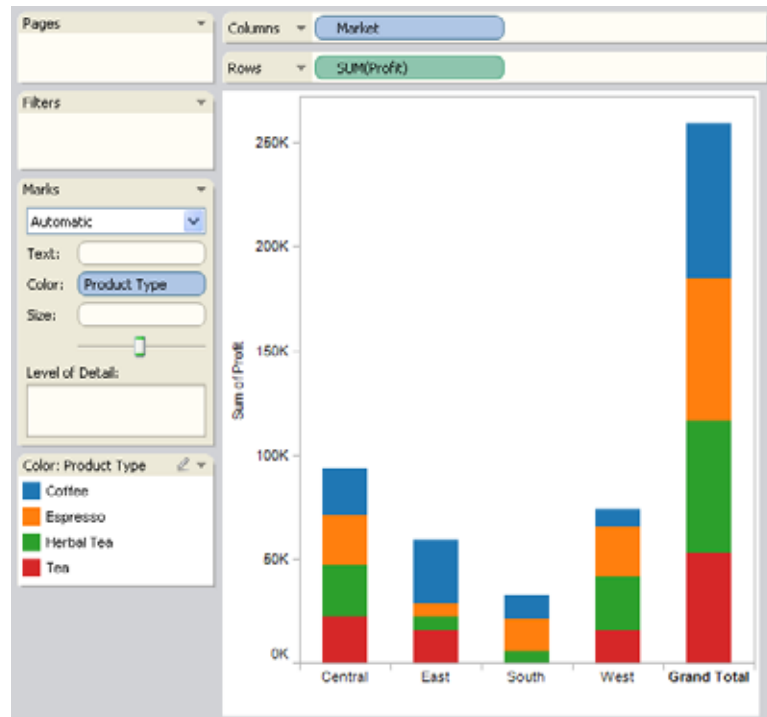
Color:

Size:

Level of Detail:

	CENTRAL	EAST	WEST	Grand Total
APPLIANCES	4,003	16,591	2,533	23,126
BINDERS AND BIN..	3,557	18,941	8,190	30,689
BOOKCASES	-2,260	-3,709	-3,583	-9,551
CHAIRS & CHAIRM..	24,876	47,929	7,774	80,578
COMPUTER PERIP..	29,534	76,143	24,762	130,439
COPIERS AND FAX	30,953	71,379	62,185	164,517
ENVELOPES	12,870	14,011	4,728	31,609
LABELS	8,207	20,216	7,497	35,920
OFFICE FURNISHI..	-759	2,696	-2,162	-225
OFFICE MACHINES	29,263	63,070	19,787	112,110
PAPER	44,109	73,709	34,344	152,162
PENS & ART SUPP..	982	2,691	570	4,243
RUBBER BANDS	180	482	493	1,144
SCISSORS, RULER..	-1,062	-1,690	42	-2,711
STORAGE & ORGA..	6,686	-3,976	8,601	11,211
TABLES	-3,769	-23,303	-1,043	-28,915
TELEPHONES AND..	305,885	818,669	455,271	1,579,825
Grand Total	493,126	1,193,867	629,189	2,316,172

You can also display totals for graphical views of data. In the figure below, only column totals are calculated because the table contains only column headers.



**Note** By default totals are computed on the server if you are connected to an SSAS data source and locally if you are connected to an Essbase data source using the aggregation specified in the cube. Refer to “Local vs. Server Computation” on page 16-63 to learn more.

### Grand Totals and Aggregations

Grand totals are computed using the aggregation of each measure. For example, if you are totaling the SUM(Gross Profit) for several products, the grand total will be the sum of the sums of profit. For aggregations such as SUM, you can easily verify the grand total because a summation of a group of sums is still a summation. However, be aware that your results



may be unexpected when using other aggregations, especially custom aggregations. For example, when looking at the average sales for several products, the grand total will be the average of the averages rather than the average of all sales. You can verify any calculation such as an aggregation or a grand total by viewing the underlying disaggregated data. Refer to “Viewing Underlying Data (Drill-Through)” on page 14-21 for more information.

The following table summarizes the standard aggregations and the grand totals that are calculated.

**Table 0-2: Grand Totals and Aggregations**

Aggregation	Calculation Description
Sum	The grand total using sum is the sum of the values shown in the row or column.
Average	The grand total using average is the averages shown in the row or column.
Minimum	The grand total using minimum is the minimum value shown in the row or column.
Maximum	The grand total using maximum is the maximum value shown in the row or column.
Standard Deviation	The grand total using standard deviation is the standard deviation values shown in the row or column.
Variance	The grand total variances are not the variances of the rows and columns in which they reside. Instead, the calculations are based on the underlying data behind the row or column.
Count & Count Distinct	The grand total counts are the counts of the rows and columns in which they reside.

### Example – Grand Totals and Aggregations

The figure shown below is a text table that displays the sales total aggregated as a sum. The grand total for the Jumbo Box shipped by Delivery Truck is \$431,837.55. Tableau calculates this number by summing all the rows in the data source that are associated with

the Jumbo Box and Delivery Truck fields. You can easily verify the number by summing the values for the Central, East, and West regions.

The screenshot shows a Tableau interface with a pivot table. The Columns shelf contains 'Region' and the Rows shelf contains 'Ship Method' and 'Ship Container'. The Marks shelf is set to 'Automatic' with the text 'SUM(Sales To...)'. The pivot table displays sales data for three regions: CENTR., EAST, and WEST. The rows are categorized by ship method (DELIVERY TRUCK, EXPRESS AIR, REGULAR AIR) and ship container (JUMBO BOX, JUMBO DRUM, LARGE BOX, MEDIUM BOX, SMALL BOX, SMALL PACK, WRAP BAG). The 'Grand Total' row at the bottom shows the sum of sales for each region and the overall total, which is 8,484,687.

		CENTR..	EAST	WEST	Grand Total
DELIVERY TRUCK	JUMBO BOX	59,928	253,710	118,200	431,838
	JUMBO DRUM	191,232	472,343	289,868	953,434
EXPRESS AIR	LARGE BOX	2,059	23,334	17,215	42,608
	MEDIUM BOX	2,819	9,213	1,881	13,913
	SMALL BOX	119,810	449,029	232,066	800,896
	SMALL PACK	12,900	21,732	9,485	44,117
	WRAP BAG	7,534	19,695	11,799	39,028
REGULAR AIR	LARGE BOX	86,208	293,764	135,033	514,996
	MEDIUM BOX	54,338	139,712	28,411	222,461
	SMALL BOX	1,067,001	2,445,364	1,316,226	4,828,591
	SMALL PACK	68,691	200,338	87,607	346,637
	WRAP BAG	34,848	141,574	69,849	246,271
Grand Total		1,697,270	4,469,797	2,317,619	8,484,687

Sum of sales for the entire data source.

The intersection of the grand total for columns and rows represents the grand total for the entire table. Tableau calculates this number by summing the sales totals for every row in the data source. Because the aggregation is a summation, you can verify this number by summing the grand totals for rows or for columns.

The figure shown below is a text table that displays the sales total aggregated as an average. The grand total for the Jumbo Box is \$3,129.26. Tableau calculates this number by averaging all the rows in the data source that are associated with the Jumbo Box and

Delivery Truck fields. You cannot verify this number by averaging the values for the Central, East, and West regions.

The image shows a Tableau interface with a pivot table. The Columns shelf contains 'Region' and the Rows shelf contains 'Ship Method' and 'Ship Container'. The Marks card is set to 'Automatic' and 'AVG(Sales Tot.)'. The table displays sales data for three regions: CENTRAL, EAST, and WEST, categorized by ship method (DELIVERY TRUCK, EXPRESS AIR, REGULAR AIR) and ship container (JUMBO BOX, JUMBO DRUM, LARGE BOX, MEDIUM BOX, SMALL BOX, SMALL PACK, WRAP BAG). The Grand Total for the entire data source is 1,197, which is highlighted with a red box.

		CENTRAL	EAST	WEST	Grand Total
DELIVERY TRUCK	JUMBO BOX	2,724	3,263	3,111	3,129
	JUMBO DRUM	4,664	3,969	4,141	4,145
EXPRESS AIR	LARGE BOX	686	1,944	3,443	2,130
	MEDIUM BOX	313	709	470	636
	SMALL BOX	990	1,227	1,240	1,192
	SMALL PACK	446	617	632	613
	WRAP BAG	204	207	303	228
REGULAR AIR	LARGE BOX	2,694	3,456	4,656	3,527
	MEDIUM BOX	1,647	1,726	838	1,503
	SMALL BOX	1,192	1,220	1,277	1,232
	SMALL PACK	528	661	684	639
	WRAP BAG	182	276	266	248
Grand Total		1,097	1,200	1,250	1,197

Average sales for the entire data source.

The intersection of the grand total for columns and rows represents the grand total for the entire table. Tableau calculates this number by averaging the sales totals for every row in the data source. Because the aggregation is an average, you cannot verify this number by averaging the grand totals for rows or for columns. To verify the grand total you need to average the rows in the data source, which are not visible in the table. You can display the rows by right-clicking in the table and selecting **View Underlying Data** on the context menu.

Region	Ship Container	Ship Method	Sales Total
EAST	SMALL BOX	REGULAR AIR	2,564.05
EAST	SMALL BOX	REGULAR AIR	2,302.42
WEST	SMALL BOX	REGULAR AIR	226.52
EAST	SMALL BOX	REGULAR AIR	159.98
EAST	SMALL BOX	REGULAR AIR	2,679.68
CENTRAL	SMALL BOX	REGULAR AIR	1,247.80
CENTRAL	SMALL BOX	REGULAR AIR	104.45
CENTRAL	SMALL BOX	REGULAR AIR	999.35
CENTRAL	SMALL BOX	REGULAR AIR	3,423.08
EAST	SMALL BOX	REGULAR AIR	4,638.68
EAST	SMALL BOX	REGULAR AIR	94.93
CENTRAL	SMALL BOX	REGULAR AIR	768.45
CENTRAL	SMALL BOX	REGULAR AIR	6,789.43
EAST	SMALL BOX	REGULAR AIR	1,895.32
EAST	SMALL BOX	REGULAR AIR	248.75
EAST	SMALL BOX	REGULAR AIR	3,715.50
WEST	SMALL BOX	REGULAR AIR	688.70
WEST	SMALL BOX	REGULAR AIR	3,221.68
WEST	SMALL BOX	REGULAR AIR	117.72
WEST	SMALL BOX	REGULAR AIR	1,690.74
WEST	SMALL BOX	REGULAR AIR	600.68
WEST	SMALL BOX	REGULAR AIR	2,783.93
CENTRAL	SMALL BOX	REGULAR AIR	1,054.00

## Subtotals

Any data view in Tableau can include subtotals. For example, you may have a view containing the total sales for two product types broken down by specific products. In addition to seeing the sales for each product you may want to see the total sales for each product type.

To calculate subtotals, identify the field you want to subtotal (Product Type in the above example) and select Subtotals from its context menu. The subtotal is the total of all the members in that field.

When you turn on subtotals for a specific field, the totals will change based on where that field is in the view. Consider the following example. The view below shows the sales for different product types sold across four different markets. Each product type is broken down by specific products. In addition, subtotals are turned on so that the view shows the total sales for each product type.

Columns		Market			
Rows		Product Type		Product	
		Central	East	South	West
Coffee	Amaretto	14,011	2,983		9,265
	Columbian	28,913	47,386	21,664	30,367
	Decaf Irish Cream	26,156	6,261	11,592	18,236
	<b>Total</b>	69,080	56,640	33,256	57,856
Espresso	Caffe Latte			15,442	20,458
	Caffe Mocha	36,218	16,646	14,163	18,676
	Decaf Espresso	24,485	7,722	15,384	30,578
	Regular Espresso		24,036		
	<b>Total</b>	59,703	48,405	44,989	69,911
Herbal Tea	Chamomile	38,570	2,194	11,186	25,632
	Lemon	21,978	27,176	14,497	32,274
	Mint	9,337	11,992		14,380
	<b>Total</b>	67,885	41,362	25,683	72,286
Tea	Darjeeling	30,209	14,096		28,769
	Earl Grey	32,881	6,505		27,387
	Green Tea	6,211	11,571		16,063
	<b>Total</b>	69,300	32,172		72,220

The subtotal is the sum of the sales for each product in the product type.

Now let's move the Product Type field from the Rows shelf to the Columns shelf. The view still shows the sales for four different product types; but now, the product types are broken down by market. Because subtotals were turned on for the Product Type, the subtotals are the sum of the sales completed in each market.

The subtotals are now the sum of the sales for each market.

Columns		Product Type		Market							
Rows		Product									
		Coffee					Espresso				
		Central	East	South	West	Total	Central	East	South	West	Total
Amaretto		14,011	2,993		9,285	26,289					
Caffe Latte									15,442	20,458	35,900
Caffe Mocha							35,218	16,646	14,163	18,876	84,904
Chamomile											
Columbian		28,913	47,386	21,664	30,367	128,319					
Darjeeling											
Decaf Espresso							24,485	7,722	15,384	30,578	78,168
Decaf Irish Cream		26,155	6,261	11,592	18,235	62,243					
Earl Grey											
Green Tea											
Lemon											
Mint											
Regular Espresso								24,036			24,036

**Note** By default subtotals are computed on the server if you are connected to an SSAS data source and locally if you are connected to an Essbase data source using the aggregation specified in the cube. That means if you are analyzing the average sales for each product, calculating the subtotals for each product type would result in the average sales of all products within that product type. This is not the case if you perform the computation locally instead of remotely. Refer to “Local vs. Server Computation” on page 16-63.

## Percentages

Any analysis in Tableau can be expressed in terms of percentages. For example, rather than viewing sales totals for every product, you might want to view each product's sales as a percentage of the total sales for all products.

You calculate percentages by selecting the **Analysis > Percentages Of** menu item. When you do this, all measures on the worksheet are displayed as a percentage based on all the table data.

---

**Note** The percentage options on the Analysis menu correspond to the percentage Table Calculations. When you select a percentage option, you are actually adding a Percent of Total table calculation. Refer to “Table Calculations” on page 16-40 to learn more.

---

This section discusses the following topics:

- About Percentages
- Percentages and Aggregations
- Percentage Options

### About Percentages

There are two factors that contribute to the percentage calculation:

- The aggregation – Percentages are calculated on the basis of the current aggregation for each measure. Refer to “Percentages and Aggregations” on page 16-75 for more information.
- The data to which you compare all percentage calculations – Percentages are a ratio of numbers. The numerator is the value of a given mark. The denominator depends on the type of percentage you want, and is the number to which you compare all your calculations. The comparison can be based on the entire table, a row, a pane, and so on. By default, Tableau uses the entire table. Other percentage calculations are available via the **Percentage of** menu item. Refer to “Percentage Options” on page 16-77 for more information.

The figure below is an example of a text table with percentages. The percentages are calculated with the **Sales Total** measure aggregated as a summation, and are based on the entire table.

Columns		9: YEAR(Order Date)				
Rows		Product Category 1		Region		
			2001	2002	2003	2004
FURNITURE	CENTRAL		1.159%	0.565%	0.393%	0.604%
	EAST		3.837%	1.873%	1.716%	1.566%
	WEST		1.834%	0.948%	0.600%	0.602%
OFFICE SUPPLIES	CENTRAL		1.383%	1.301%	0.948%	1.028%
	EAST		3.414%	2.237%	2.340%	2.538%
	WEST		1.434%	1.593%	0.953%	1.215%
TECHNOLOGY	CENTRAL		3.997%	2.376%	2.661%	3.589%
	EAST		10.065%	7.764%	7.128%	8.213%
	WEST		5.905%	3.813%	4.110%	4.310%

## Percentages and Aggregations

Percentages are computed on the basis of the aggregation for each measure. Standard aggregations include summation, average, and so on. Refer to “Aggregations” on page 16-4 for more information.

For example, if the aggregation applied to the **Sales Total** field is a summation, then the default percentage calculation (percent of table) means that each number displayed is the SUM(Sales Total) for that mark divided by the SUM(Sales Total) for the entire table.

In addition to using predefined aggregations, you can use custom aggregations when calculating percentages. You can define your own aggregations by creating a calculated field. Once the new field is created, you can use percentages on the field as you would any other field. Refer to “Aggregate Calculations” on page 16-30 for more information.

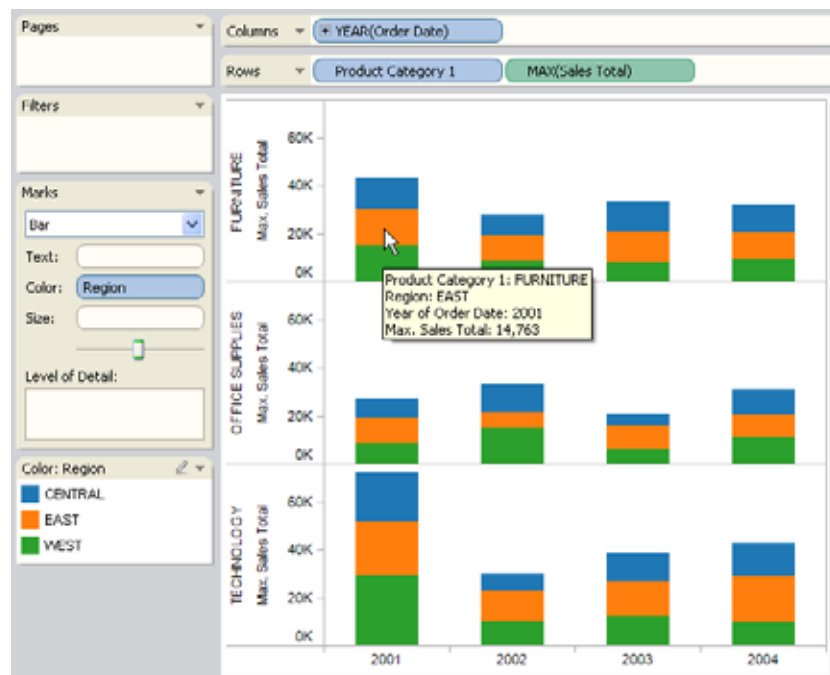
Percent calculations can also be applied to data that are disaggregated. In this case, all values are expressed as the percentage of a summation. You cannot choose an alternative aggregation.

### Example – Percentages and Aggregations

The view shown below is a nested bar chart created using two dimensions and a measure that is aggregated as a maximum. Additionally, the data are color-encoded by a dimension and the default percentage calculation has been applied. Notice that the axis labels are modified to reflect the percent calculation.




The tooltip reveals that the maximum sales total for furniture in the east in 2001 is 50.3% of the maximum for the entire table. What is the maximum for the table? You can see by inspection that the maximum occurs in the west in the year 2001. The tooltip for this bar segment would reveal a maximum sales total of 100%.



The next view displays two disaggregated measures as a scatter plot. Again, the default percentage calculation has been applied as reflected by the modified axis labels.

The tooltip shows that the selected data point has a gross profit of 0.0372% and a sales total of 0.1754%. Note that the percentage calculations are based on the entire data source.



## Percentage Options

Computing a percentage involves specifying a total on which the percentage is based. In Tableau, the default percentage calculation is based on the entire table. However, you can change the default by selecting a different percentage option from the **Analysis > Percentage of** menu. Select from the following options:

- Percent of Table
- Percent of Column
- Percent of Row
- Percent of Pane
- Percent of Row in Pane
- Percent of Column in Pane

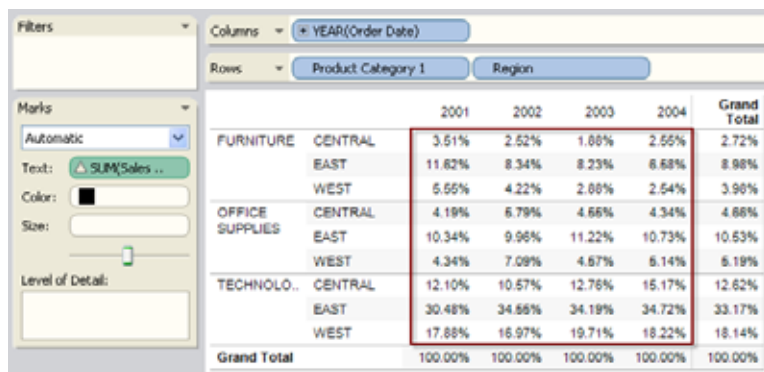
The option you choose is applied uniformly to all measures that appear on a worksheet. For instance, you cannot choose **Percent of Column** for one measure and **Percent of Row** for another.

If you are unsure what the current percentage calculation means, display the grand totals. This provides more information about each row and column. For example, if you select **Percent of Row** while displaying grand totals, you will see that the total for each row is exactly 100%. Refer to “Grand Totals” on page 16-65 for more information on grand totals.

The percent calculation options are described in the following sections. In each case, the grand totals are displayed as well.

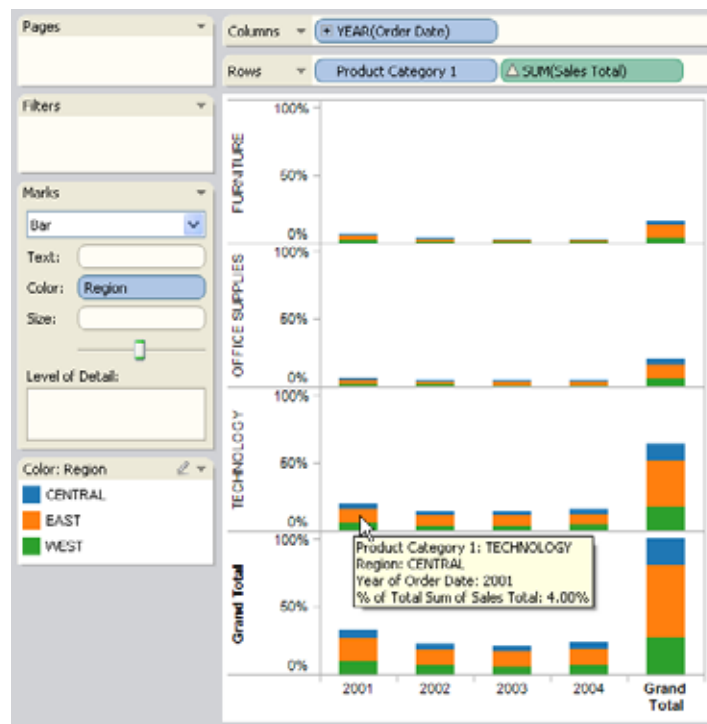
### Percent of Table

When you select **Percentage of Table**, each measure on the worksheet is expressed as a percentage of the total for the entire worksheet (table). For example, Technology in the East region accounts for 10.07% of total sales in 2001. The grand totals for rows show that 2001 accounts for 33.03% of the total sales, while Technology in the East accounts for 33.17%. Summing the grand totals for rows or for columns yields 100% of the total.



		2001	2002	2003	2004	Grand Total
FURNITURE	CENTRAL	3.51%	2.52%	1.88%	2.55%	2.72%
	EAST	11.82%	8.34%	8.23%	8.68%	8.98%
	WEST	5.55%	4.22%	2.88%	2.54%	3.98%
OFFICE SUPPLIES	CENTRAL	4.19%	5.79%	4.66%	4.34%	4.68%
	EAST	10.34%	9.96%	11.22%	10.73%	10.53%
	WEST	4.34%	7.09%	4.67%	5.14%	5.19%
TECHNOLO.	CENTRAL	12.10%	10.57%	12.76%	15.17%	12.62%
	EAST	30.48%	34.66%	34.19%	34.72%	33.17%
	WEST	17.88%	16.97%	19.71%	18.22%	18.14%
Grand Total		100.00%	100.00%	100.00%	100.00%	100.00%

An equivalent graphical table is shown below.

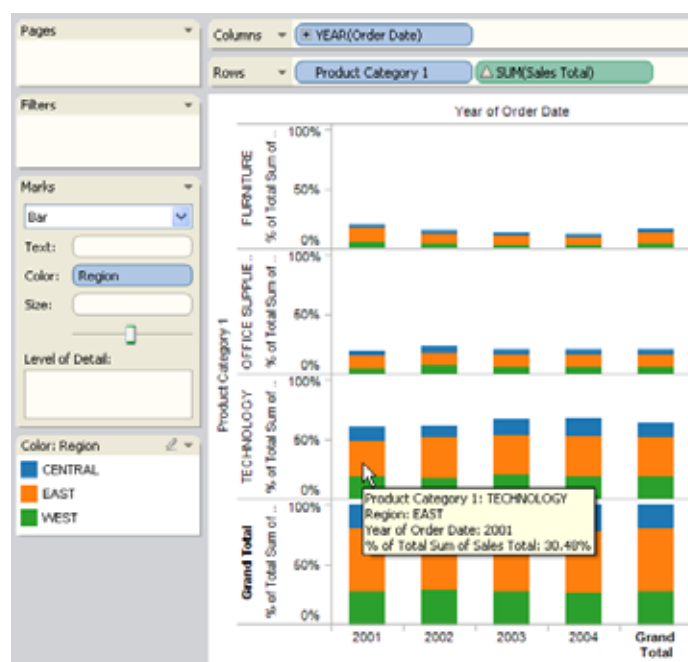


### Percent of Column

When you select **Percentage of Column**, each measure on the worksheet is expressed as a percentage of the total for the column. For example, Technology in the East region accounts for 30.48% of total sales in the 2001. 2001 is the column in this case. The grand total for columns yields the same results as “Percent of Table” on page 16-78 because all columns in the table are used in the calculation for a given region. For example, Technology in the East region accounts for 33.17% of the total sales. The grand total for rows yields 100% for each column (year) by definition.

Columns		+ YEAR(Order Date)				
Rows		Product Category 1    Region				
		Year of Order Date				
Product Category 1	Region	2001	2002	2003	2004	Grand Total
FURNITURE	CENTRAL	3.61%	2.82%	1.88%	2.66%	2.72%
	EAST	11.62%	8.34%	8.23%	6.58%	8.98%
	WEST	6.66%	4.22%	2.88%	2.64%	3.98%
OFFICE SUPPLIES	CENTRAL	4.19%	5.79%	4.55%	4.34%	4.66%
	EAST	10.34%	9.96%	11.22%	10.73%	10.63%
	WEST	4.34%	7.09%	4.57%	5.14%	5.19%
TECHNOLOGY	CENTRAL	12.10%	10.67%	12.76%	16.17%	12.82%
	EAST	30.48%	34.55%	34.19%	34.72%	33.17%
	WEST	17.88%	16.97%	19.71%	18.22%	18.14%
Grand Total		100.00%	100.00%	100.00%	100.00%	100.00%

An equivalent graphical table is shown below.

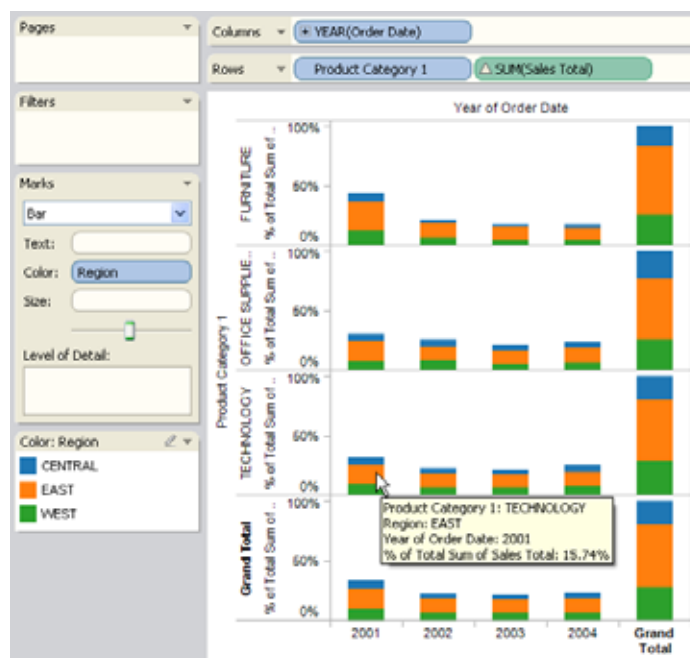


### Percent of Row

When you select **Percentage of Row**, each measure on the worksheet is expressed as a percentage of the total for the row. For example, Technology in the East region accounts for 30.34% of total sales in 2001. East is the row in this case. The grand total for rows are the same as “Percent of Table” on page 16-78. This is because all rows in the table are used in the calculation for a given year. For example, the year 2001 accounts for 33.03% of the total sales. The grand total for columns yields 100% for each row (region) by definition.

Columns		+ YEAR(Order Date)				
Rows		Product Category 1    Region				
		Year of Order Date				Grand Total
Product Category 1	Region	2001	2002	2003	2004	
FURNITURE	CENTRAL	42.81%	20.77%	14.43%	22.18%	100.00%
	EAST	42.72%	20.88%	19.10%	17.32%	100.00%
	WEST	48.03%	23.81%	18.06%	15.11%	100.00%
OFFICE SUPPLIES	CENTRAL	29.69%	27.91%	20.35%	22.06%	100.00%
	EAST	32.42%	21.26%	22.22%	24.10%	100.00%
	WEST	27.60%	30.66%	18.35%	23.39%	100.00%
TECHNOLOGY	CENTRAL	31.66%	18.82%	21.08%	29.43%	100.00%
	EAST	30.34%	23.41%	21.49%	24.76%	100.00%
	WEST	32.56%	21.02%	22.66%	23.76%	100.00%
Grand Total		33.03%	22.47%	20.85%	23.65%	100.00%

An equivalent graphical table is shown below. Note that the tooltip yields a different result than the text table because the row includes all regions.



### Percent of Pane

When you select **Percent of Pane**, each measure on the worksheet is expressed as a percentage of the total for the pane. This option is the same as **Percent of Table** when the table consists of only a single pane.

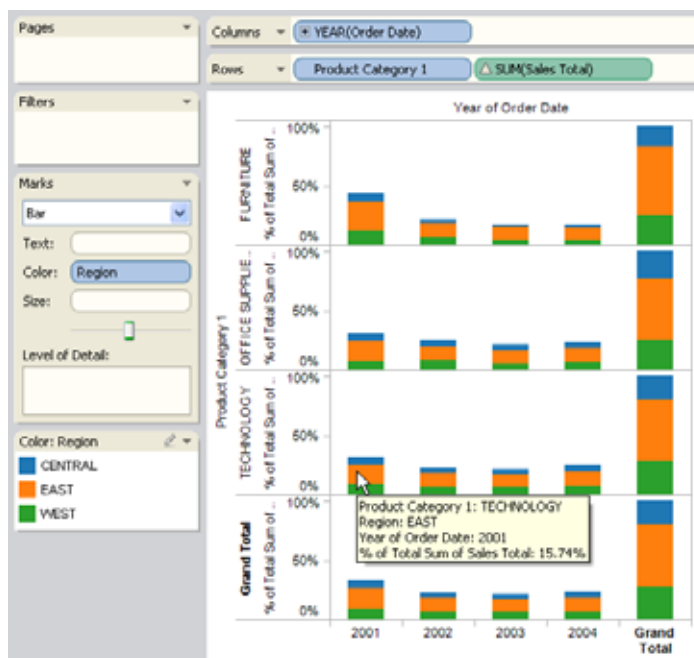
This table consists of three panes. Each pane encompasses one product category and all four years. Therefore, the table is one pane wide and three panes high. For more information about panes, refer to “Building Views Manually” on page 10-1.

For example, Technology in the East region accounts for 15.74% of total sales for the Technology pane, which is shown below. The grand total for rows yields the same results as “Percent of Table” on page 16-78. This is because all panes (Furniture, Office Supplies, and Technology) are used in the calculation. The grand total for columns represents the percentage of a given row for the pane. For example, the East region in the Technology pane accounts for 51.88% of the total for all rows in that pane. Summing the Central, East, and West regions yields  $19.74\% + 51.88\% + 28.37\% = 100\%$  (within the specified precision).

Columns		YEAR(Order Date)				
Rows		Product Category 1    Region				
		Year of Order Date				Grand Total
Product Category 1	Region	2001	2002	2003	2004	
FURNITURE	CENTRAL	7.39%	3.60%	2.50%	3.85%	17.35%
	EAST	24.46%	11.94%	10.94%	9.92%	67.26%
	WEST	11.89%	6.05%	3.82%	3.84%	25.39%
OFFICE SUPPLIES	CENTRAL	6.79%	6.38%	4.65%	5.04%	22.86%
	EAST	16.75%	10.98%	11.48%	12.45%	51.66%
	WEST	7.03%	7.81%	4.68%	5.96%	25.49%
TECHNOLOGY	CENTRAL	6.25%	3.72%	4.16%	5.61%	19.74%
	EAST	15.74%	12.14%	11.15%	12.85%	51.88%
	WEST	9.24%	6.96%	6.43%	6.74%	28.37%
Grand Total		33.03%	22.47%	20.85%	23.65%	100.00%



An equivalent graphical table is shown below.



### Percent of Row in Pane

When you select **Percentage of Row in Pane**, each measure on the worksheet is expressed as a percentage of the total for a given row within a pane. This option is the same as **Percent of Row** when the table is only a single pane wide.

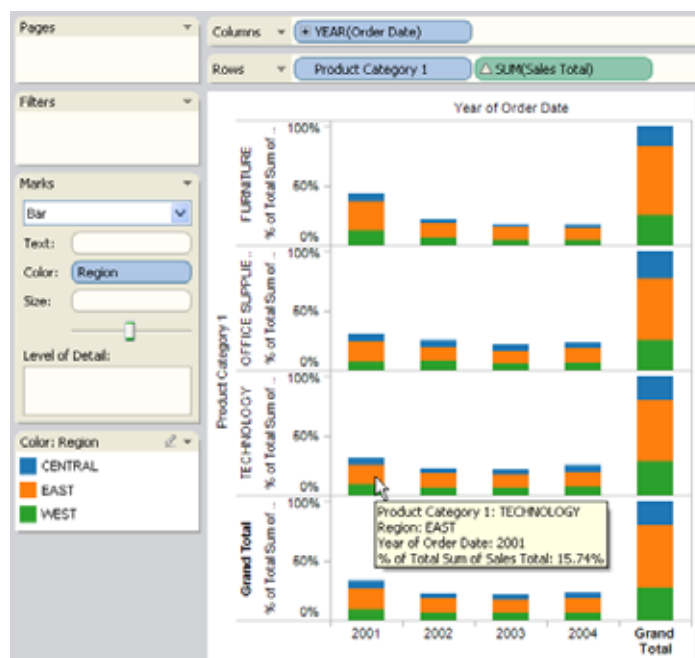
**Note** If you place **Measure Names** as the inner dimension on the **Columns** shelf, Tableau will return 100% for each mark because you cannot total up the values for multiple measure names. For example, you can't total up the values for SUM(Sales Total) and SUM(Gross Profit).

For example, Technology in the East region accounts for 30.34% of total sales in 2001. As shown below, the row is the East region and the pane is the Technology product category.

The grand total for columns is always 100%. This represents the sum of a given row within a pane. For example, the East region in the Technology pane sums to  $30.34\% + 23.41\% + 21.49\% + 24.76\% = 100\%$ . The grand total for rows is identical to the results given in “Percent of Row” on page 16-81 because the table is a single pane wide.

Columns		* YEAR(Order Date)				
Rows		Product Category 1    Region				
		Year of Order Date				
Product Category 1	Region	2001	2002	2003	2004	Grand Total
FURNITURE	CENTRAL	42.81%	20.77%	14.43%	22.18%	100.00%
	EAST	42.72%	20.86%	19.10%	17.32%	100.00%
	WEST	46.03%	23.81%	16.06%	15.11%	100.00%
OFFICE SUPPLIES	CENTRAL	29.69%	27.91%	20.35%	22.06%	100.00%
	EAST	32.42%	21.26%	22.22%	24.10%	100.00%
	WEST	27.60%	30.66%	18.35%	23.39%	100.00%
TECHNOLOGY	CENTRAL	31.66%	18.82%	21.08%	28.43%	100.00%
	EAST	30.34%	23.41%	21.49%	24.76%	100.00%
	WEST	32.66%	21.02%	22.66%	23.78%	100.00%
Grand Total		33.03%	22.47%	20.85%	23.65%	100.00%

An equivalent graphical table is shown below. Note that the tooltip yields a different result than the text table because the pane consists of different numbers of rows. In the text table, the pane has three rows. In the bar chart, the pane has only one row. In fact, the bar chart is identical to the view shown in “Percent of Pane” on page 16-82.



### Percent of Column in Pane

When you select **Percentage of Column in Pane**, each measure on the worksheet is expressed as a percentage of the total for a given column within a pane. This option is the same as **Percent of Column** when the table is only a single pane high.

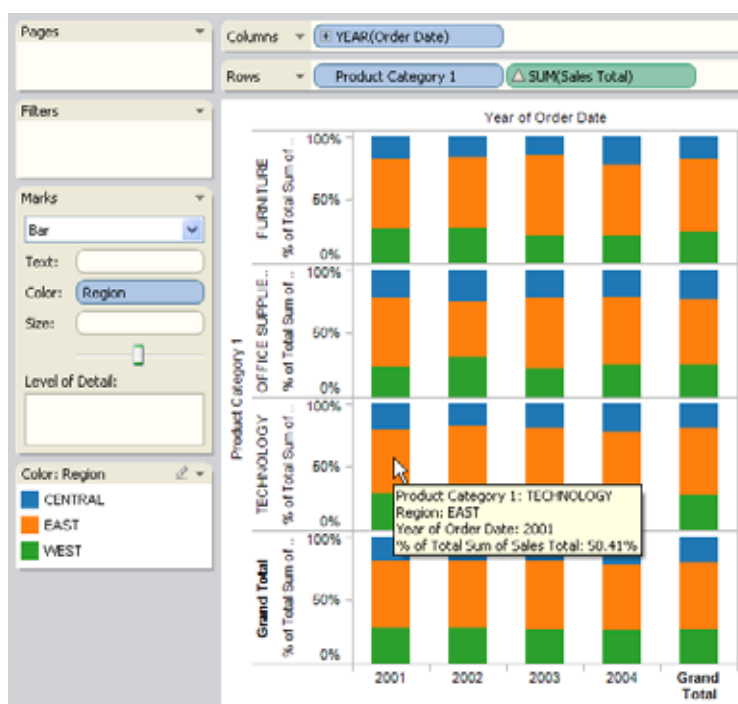
**Note** If you place **Measure Names** as the inner dimension on the **Rows** shelf, Tableau will return 100% for each mark because you cannot total up the values for multiple measure names. For example, you can't total up the values for SUM(Sales Total) and SUM(Gross Profit).

For example, Technology in the East region accounts for 50.41% of the total sales in 2001. As shown below, the column is the year 2001 and the pane is the Technology product category.

The grand total for rows is always 100%. This represents the sum of a given column within a pane. For example, in 2001 the Technology pane sums to 20.02% + 50.41% + 29.57% = 100%. The grand total for columns represents the percentage of a given row. For example, the East region in the Technology pane accounts for 51.88% of the total for all rows across all columns in that pane. Summing the Central, East, and West regions yields 19.74% + 51.88% + 28.37% = 100% (within the specified precision).

Columns		F YEAR(Order Date)				
Rows		Product Category 1 Region				
		Year of Order Date				Grand Total
Product Category 1	Region	2001	2002	2003	2004	
FURNITURE	CENTRAL	18.98%	18.89%	14.50%	21.88%	17.35%
	EAST	56.18%	55.31%	63.35%	56.35%	57.26%
	WEST	26.85%	28.00%	22.14%	21.79%	25.39%
OFFICE SUPPLIES	CENTRAL	22.20%	25.36%	22.35%	21.50%	22.86%
	EAST	54.79%	43.61%	55.17%	53.08%	51.65%
	WEST	23.01%	31.04%	22.48%	25.42%	25.49%
TECHNOLOGY	CENTRAL	20.02%	17.03%	19.14%	22.28%	19.74%
	EAST	50.41%	55.65%	51.29%	50.97%	51.88%
	WEST	29.57%	27.32%	29.57%	26.75%	28.37%
Grand Total		100.00%	100.00%	100.00%	100.00%	100.00%

An equivalent graphical table is shown below. Note that it uses equally-sized marks. This makes the table particularly easy to analyze because the grand totals do not dominate the overall scale as they do in other data views such as “Percent of Row in Pane” on page 16-84.



# Background Images

---

<b>Overview . . . . .</b>	<b>17-2</b>
<b>Adding Background Images . . . . .</b>	<b>17-3</b>
<b>Setting up the View . . . . .</b>	<b>17-6</b>
<b>Managing Background Images . . . . .</b>	<b>17-8</b>
Editing an Image . . . . .	17-8
Enabling/Disabling Images . . . . .	17-9
Adding Show/Hide Conditions . . . . .	17-9
Removing an Image . . . . .	17-11

---

## Overview

Background images are images that you display underneath your data in order to add more context to the marks in the view. A common use of background images is adding a map behind geographical information using latitude and longitude, zip code, or address values. For example, you may have a scatter plot that shows delivery times based on latitude and longitude. You can add a map as the background image so you can easily identify the regions that tend to have long delivery times.

Although maps are a common use of background images, you can use this feature for any data that relates to a two dimensional image. For example, use background images to show floor plans of a building, a model of the sea floor, or images of Web pages for analyzing Web logs.

This section discusses the following topics:

- Adding Background Images
- Setting up the View
- Managing Background Images

---

**Note** For more information about adding background images, map resources, and geocoding your data, please visit our Training & Support Center.


---

## Adding Background Images

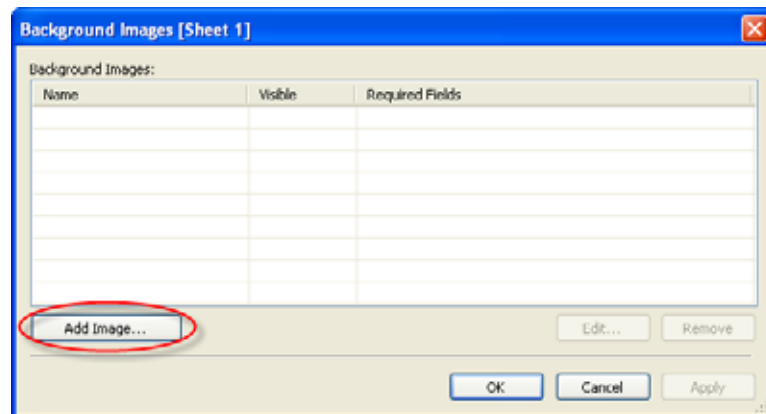
When you add a background image to the view, you need to specify a coordinate system by mapping both the x and y axes to the values of a fields in your data base. If you are adding a map, the x and y axes should be longitude and latitude expressed as a decimal. However, you can map the axes to any relevant fields based on your own coordinate system.

**To add a background image:**

- 1 Select **Data > Background Images**.



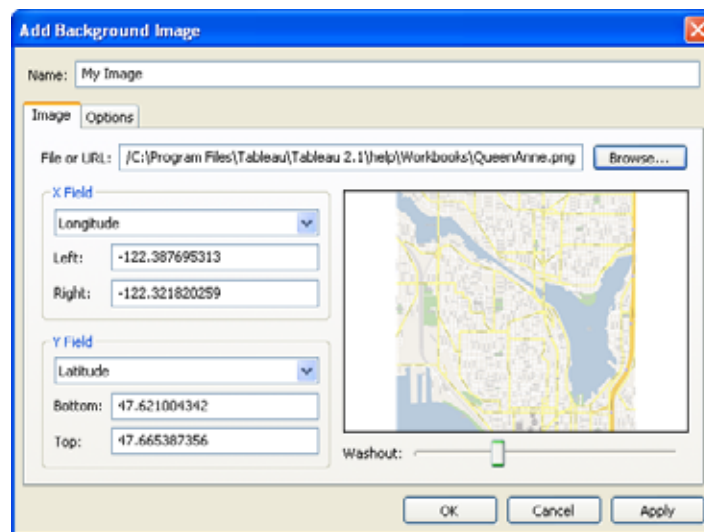
- 2 In the Background Images dialog box, click **Add Image**.



- 3 In the Add Background Image dialog box do the following:
  - Type a name for the image into the **Name** text box.

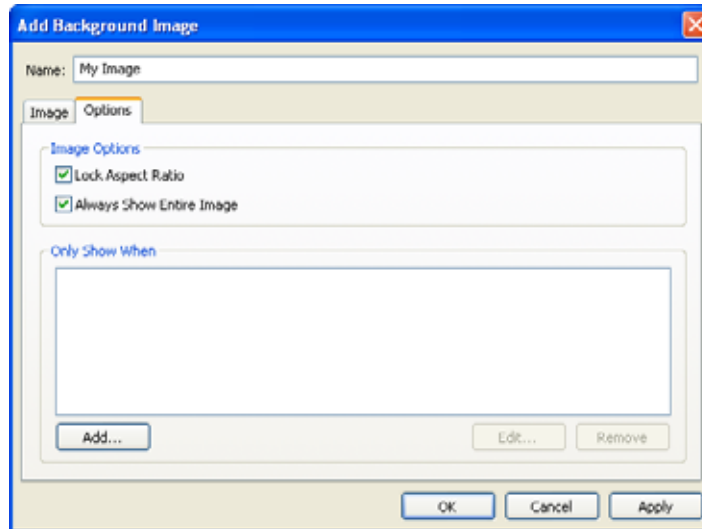


- Click **Browse** to navigate to and select the image you want to add to the background. You can also type a URL to link to an image hosted online.
- Select the field to map to the x-axis of the image and specify the left and right values. When adding a map, the longitude values should be mapped to the x-axis using decimal values (instead of degrees/minutes/seconds or N/S/E/W).
- Select the field to map to the y-axis of the image and specify the top and bottom values. When adding a map, the latitude values should be mapped to the y-axis using decimal values (instead of degrees/minutes/seconds or N/S/E/W).
- You can adjust the intensity of the image using the Washout slider. The farther the slider moves to the right, the more faded the image will appear behind your data.



- 4 You can specify the following options using the **Options** tab:
  - **Lock Aspect Ratio** - select this option to maintain the original dimensions of the image for any manipulations of the axes. Deselecting this option allows the image's shape to be distorted.
  - **Always Show Entire Image** - select this option to avoid cropping the image when the data encompasses only a portion of the image. If you lock both the axis in a view, this option may be negated.

- Add conditions for when to show the image. Refer to “Adding Show/Hide Conditions” on page 17-9 to learn more about defining conditions.



5 Click **OK**.

Refer to “Setting up the View” on page 17-6 to learn more about setting up the view so the image can display correctly.

---

**Note** In order to make the marks in a view more visible when placed on top of a background image, each mark is surrounded by a solid contrasting color called a halo. You can turn mark halos off by selecting **Format > Show Mark Halos**. Refer to “Mark Halos” on page 20-19 to learn how to change the color of the mark halos.

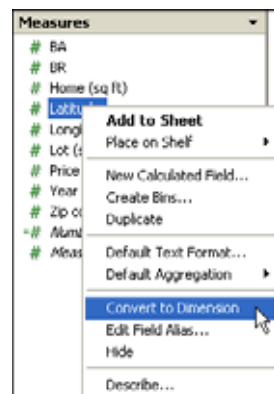
---

## Setting up the View

After you add a background image, you need to build the view in a way that matches the x and y mappings you specified for the image. That is, the fields you specified as x and y must be on the proper shelves and most often need to be continuous dimensions. Follow the steps below to set up the view correctly:

- 1 Convert the measures to dimensions.

If any of the fields you mapped to the x and y axes of the image are measures, you should convert them to dimensions. Drag the measure from the Measures area in the Data window to the Dimensions area. You can also right-click the measure and select **Convert to Dimension**.



For more information about converting fields refer to “Converting Measures to Dimensions” on page 8-24.

- 2 Place the field mapped to the x-axis on the **Columns** shelf.

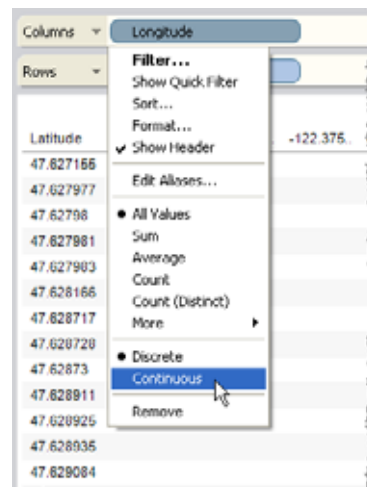
If you are working with maps, the longitude field should be on the columns shelf. It may seem backward at first, however, the fields on the columns shelf determine the values distributed across the x-axis.

- 3 Place the field mapped to the y-axis on the **Rows** shelf.

If you are working with maps, the latitude field should be on the rows shelf. It may seem backward at first, however, the fields on the rows shelf determine the values distributed across the y-axis.

- 4 Finally, convert the dimensions to continuous values.

In order to view the background image, you must use continuous dimensions. On the field menu for each dimension, select **Continuous**.



## Managing Background Images

You can add several background images to the workbook and then select the image or images to make active on each sheet. The Background Images dialog box lists all of the images, the required fields, and whether they are visible. The visibility is determined based on whether the required fields are used in the current view. This section discusses the following way to manage your images:

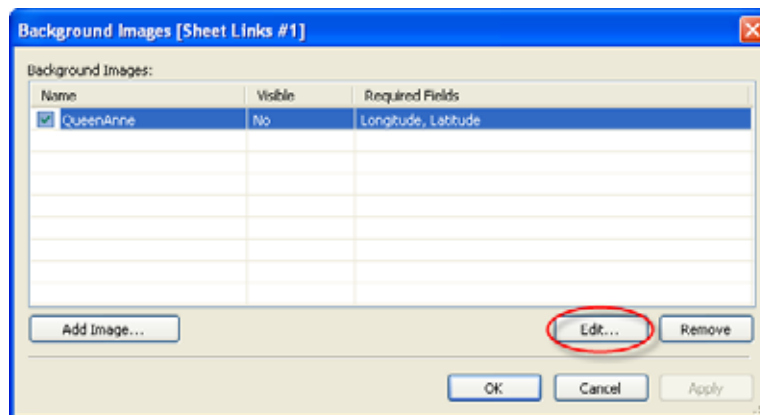
- Editing an Image
- Enabling/Disabling Images
- Adding Show/Hide Conditions
- Removing an Image

### Editing an Image

After adding a background image, you can always go back and edit the x and y field mappings as well as any of the options on the Options tab.

To edit an image:

- 1 Select **Data > Background Images**.
- 2 In the Background Images dialog box, select the image you want to edit and click **Edit** (you can also just double-click the image name).



- 3 In the Edit Background Image dialog box, make the changes to the image and click **OK**.

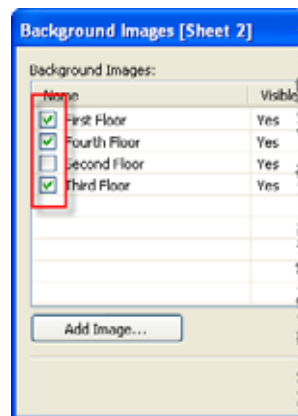
## Enabling/Disabling Images

Although you can add multiple images to a workbook, you may want to only use a subset of the images for a particular set of worksheets. For example, you may want to show a map of the entire United States of America on one view, and maps of individual states in other views.

Use the checkboxes in the Background Images dialog box to enable and disable the images for the current worksheet. You can show several images by enabling multiple images on a single worksheet. For example, you may have several images that you want to tile in the background to make a larger background image.

To enable or disable a background image:

- 1 Select **Data > Background Images**.
- 2 In the Background Images dialog box, select the checkboxes next to the images you want enabled.



- 3 Click **OK**.


## Adding Show/Hide Conditions

When you add a background image and enable it, the image will be shown automatically on any worksheet that has the required fields used in the view. To avoid showing an image


on all the worksheets, you can specify show/hide conditions. Show/Hide conditions are conditional statements that you define to specify when to show the image. For example, you may have a floor plan image for a multi-story building. While each image is associated with the same coordinates (the corners of the building), you do not want to show the first floor map when looking at the third floor information. In this case, you can specify a condition to only show the first floor image when the Floor field is equal to one.

**To specify show/hide conditions:**


- 1 Select **Data > Background Images**.
- 2 In the Background Images dialog box, select the image you want to add a condition to and click **Edit**.
- 3 In the subsequent dialog box, select the **Options** tab.
- 4 Click the **Add...** button at the bottom of the dialog box.



- 5 Select a field to base the condition on. In the example described above, the field is Floor.



- 6 Specify when you want to show the image by selecting a value of the field. For this example, one is selected.



- 7 Click **OK**.

A condition statement is added to the image. In the building floor plan example, the condition statement is Only show the image when Floor is equal to One.

- 8 Click **OK** twice to close the Background Image dialog boxes and apply the changes.

## Removing an Image

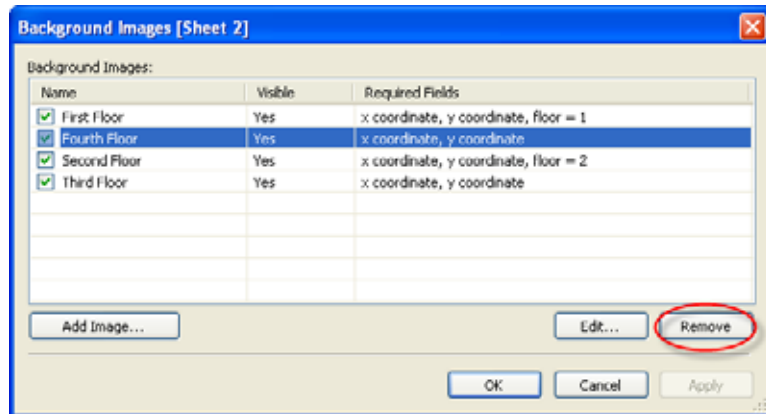
When you no longer want to use a background image you can either disable it (refer to “Enabling/Disabling Images” on page 17-9) or remove it, making it unavailable to all worksheets.

**To remove an image:**

- 1 Select **Data > Background Images**.



- 2 In the Background Images dialog box, select the image you want to remove and click **Remove**.



- 3 Click **OK**.

# Trend Lines and

<b>Overview</b>	<b>18-3</b>
<b>Adding Trend Lines</b>	<b>18-4</b>
Add Trend Lines to the View	18-4
Why can't I add Trend Lines?	18-6
Remove Trend Lines	18-7
<b>The Trend Line Model</b>	<b>18-8</b>
<b>Removing Factors from the Model</b>	<b>18-9</b>
<b>Testing Significance</b>	<b>18-13</b>
Entire Model Significance	18-13
Significance of Specific Fields	18-14pas
<b>Trend Lines Example</b>	<b>18-16</b>
<b>Assumptions</b>	<b>18-20</b>
<b>Trend Line Model Terms</b>	<b>18-21</b>
<b>Commonly Asked Questions</b>	<b>18-23</b>

---

# Statistics

---

---

## Overview

You can use Tableau's trend line feature to incrementally construct interactive models of behavior that you can use to make predictions about your data. For example, answer questions like whether profit is predicted by time, or whether average delays at an airport are significantly correlated with the month of the year. This section discusses the following topics:

- Adding Trend Lines
- The Trend Line Model
- Removing Factors from the Model
- Testing Significance
- Trend Lines Example
- Assumptions
- Trend Line Model Terms
- Commonly Asked Questions

## Adding Trend Lines

When you add trend lines to the view, you can specify several options about how you want them to look and behave. This section discusses:

- Add Trend Lines to the View
- Why can't I add Trend Lines?
- Remove Trend Lines


### Add Trend Lines to the View

- 1 Select **Analysis > Trend Lines** or right-click on the pane and select **Trend Lines**.


This command adds a linear trend line and will add a trend line for each color if there is a discrete field on the color shelf. You can optionally continue with the steps below to specify different trend line options.

- 2 Right-click on the pane and select **Trend Line Options** to open a dialog box where you can specify the following options:
  - Select either a **Linear**, **Logarithmic**, or **Polynomial** model type.


Linear Model



Logarithmic Model

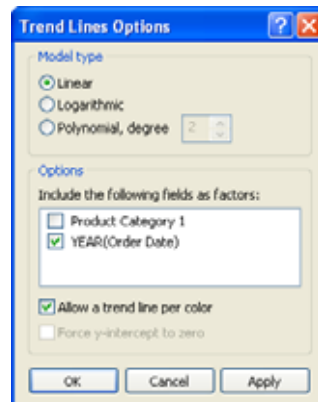


Polynomial Model



- Select fields you want to include as factors in the trend line model. For example, on a view of yearly sales for three different products, you may want to see the overall sales


trend across all products rather than a different line for each product. In this case, you would exclude the product field as a factor. As shown here:




For more information on excluding factors refer to “Removing Factors from the Model” on page 18-9.

- Select whether to exclude color. When you have color encodings in your view, you can use this option to add a single trend line that models all of the data ignoring the encoding.

Including Color



Excluding Color



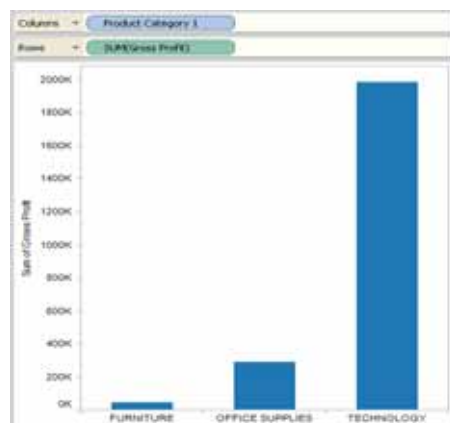
- Select whether to force the intercept to zero. This option is useful when you know that you want your trend line to begin at zero. For example, you may have an analysis of the number of products sold over time. You know that you started out with zero products sold, so you can force the trend line to begin there.

3 When finished click **OK**.

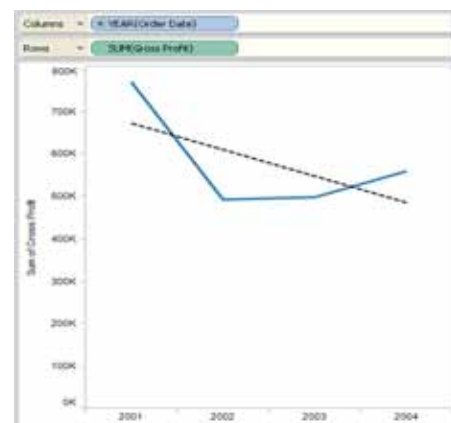
## Why can't I add Trend Lines?

To add trend lines to a view, both axis must contain a field that can be interpreted as a number. For example, you cannot add a trend line to a view that has a Product Category file, that contains strings, on the columns shelf and a gross profit measure on the rows shelf. However, you can add a trend line to a view of sales over time because both sales and time can be interpreted as numeric values.

Cannot add a trend line.



Can add a trend line.



On multidimensional data sources, the date hierarchies actually contain strings rather than numbers thus trend lines will not be allowed. Additionally, the 'm/d/yy' and 'mmmm yyyy' date formats on all data sources do not allow trend lines.

If you have trend lines turned on and you modify the view in a way where trend lines are not allowed, the trend lines will not show. When you change the view back to a state that allows trend lines, they will re-appear.

---

**Note** Tableau automatically stacks bar marks in many cases. However, trend lines cannot be turned on for stacked bars. You can turn off stacked marks by selecting **Off** on the **Analysis > Stacked Marks** menu.

---

## Remove Trend Lines

To remove trend lines from the view, simply select **Analysis > Trend Lines** or right-click the pane and select **Trend Lines**. The Trend Lines command toggles between on and off. If you have specified any trend line options, they will be used the next time you turn on trend lines. However, if you close the workbook with trend lines turned off, the trend line options will be reset to the default settings.



## The Trend Line Model

Any time you add a trend line to your view you are building a statistical model. You are answering the question of whether the factors in your view predict a specific value (measure). A simple example is to wonder whether profit is predicted by time in a view that shows the gross profit of a company over four years.


When you add a trend line in Tableau, the line you see is defined by a mathematical formula of the form:

$$Y = \text{factor 1} * \text{factor 2} * \dots * \text{factor N} * f(x) + \text{random error}$$


The term Y in the above expression is called the response variable and corresponds to the value you are trying to predict while the term X is the independent variable. The factors in the expression correspond to the categorical fields in the view (time in the profit example above). In addition, each factor is represented as a matrix. The \* is a particular kind of matrix multiplication operator that takes two matrices with the same number of rows and returns a new matrix with the same number of rows. That means that in the expression  $\text{factor 1} * \text{factor 2}$ , all combinations of the members of factor 1 and factor 2 will be introduced. For example, if factor 1 and factor 2 both had three members, then a total of nine variables will be introduced into the model formula by this operator.

## Removing Factors from the Model


You can remove factors from the model using the Trend Lines Options dialog box. Often you will want to remove factors because you want the trend line model to be based on the entire row in the table rather than broken up by the members or values of a field. Consider the following example. The view below shows the monthly sales for two different products, broken down by region.




When you select a trend line in the view, you can see that a separate model is created for each market segment.




Now remove Market as a factor in the model by deselecting it in the Trend Lines Options dialog box.



You can see that the trend line model is now based on all rows associated with the Product Type field. That is, the trend line depends only on the product type and is the same in each row of a given column.





## Testing Significance

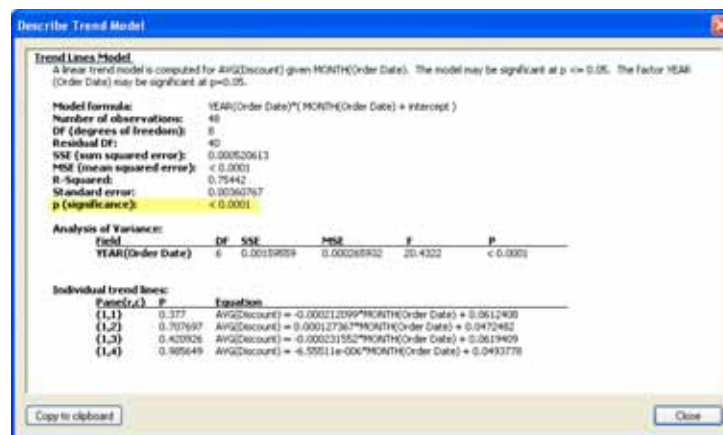
Once you've added a trend line to the view, you typically want to know the goodness of fit of the model, which is a measure of the certainty of the model's predictions. In addition, you may be interested in the significance of each factor. You can view these numbers using the Describe command.

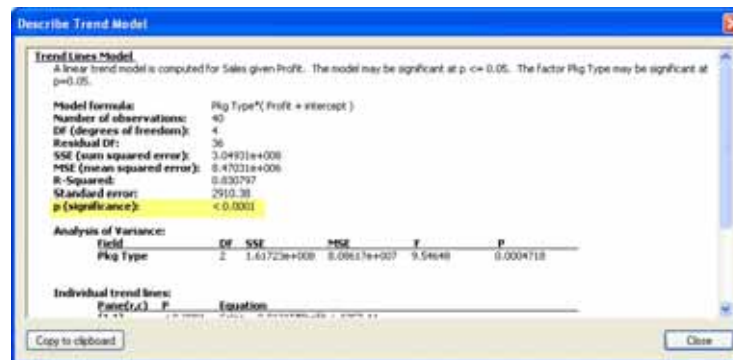
In the Sheet Description dialog box, there is a section of statistics about the Trend Line Model. When you are testing the significance, you are most concerned with the p-values. The smaller the p-value, the more significant the model or factor is. It is possible to have a model that has statistical significance but a factor that not really adding much to the model. The following topics are discussed in this section:

- Entire Model Significance
- Significance of Specific Fields
- Significance of Individual Trend Lines

### Entire Model Significance

To identify the significance of the entire trend line model, select **Analysis > Describe Trend Line Model**. In the Analysis of Variance table, the p-value is listed. The smaller the p-value, the more significant the model is.





This p-value is a comparison of the fit of the entire model to the fit of a model composed solely of the grand mean (the average of data in the data view). That is, this p-value assesses the explanatory power of the quantitative term  $f(x)$  in the model formula, which can be linear, polynomial or logarithmic with the factors fixed. It is common to test significance by applying the "95% confidence" rule. This equates to a p-value of 0.05 or less.

## Significance of Specific Fields

To identify the significance of specific fields in a trend line model, select **Analysis > Describe Trend Line Model**. In the Analysis of Variance table, which is often referred to as an ANOVA table, each of the fields that are factors in the model are listed. For each field, among other values, you can see the p-value. The p-value indicates how much that field adds to the significance of the entire model. The smaller the p-value the more that field improves the model. The values displayed for each field are derived by comparing the entire model to a model that does not include the field in question.

The following image shows the ANOVA table for a view of quarterly sales for the past two years of three different product categories.

Analysis of Variance:

Field	DF	SSE	MSE	F	P
Product Category 1	8	3.22326e+011	4.02908e+010	105.315	< 0.0001
YEAR(Order Date)	6	9.40812e+009	1.56802e+009	4.09862	0.0180582

As you can see, the p-value for Product Category 1 is smaller than the p-value for Order Date. That leads us to believe that the product category is more statistically significant in predicting the total sales than the year. However, both of these factors are statistically

significant in this model. (It is common to test significance by applying the “95% confidence” rule. This equates to a p-value of 0.05 or less.)

### Significance of Individual Trend Lines

You can view the p-value and formula for individual trend lines in two ways. Either right-click a specific trend line in the view and select **Describe Trend Line** or select **Analysis > Describe Trend Line Model** and see the Individual Trend Line table.

The Individual Trend Line table in the Describe Trend Line Model dialog box lists of all the trend lines in your view along with their p-value and the formula. This is an easy way to quickly determine which (if any) trend lines are statistically significant.




## Trend Lines Example

This example walks through the steps of incrementally testing the significance of various factors to solve a problem. This example uses the Sample - Superstore Sales data source.

We start with the question: what is causing high discount rates at a superstore? While statistics cannot answer this question, we can discover the variables that are correlated with high discounts. The view below shows the average monthly discount rates of all stores between 2002 and 2004. Refer to “Adding Trend Lines” on page 18-4 to learn how to add and remove trend lines.



Our first thought is that the discount rates may be high because items are out of stock and employees are offering discounts for the inconvenience. When we put the Product Inventory field on the rows shelf, we see that there was a large spike in discount rates for out of stock items in 2002 and other increases in 2004.




However, when looking at the ANOVA table for the trend line model, we see that the p-value for product inventory is 0.0671058. For most, this number is too high to be significant, indicating that we cannot predict discount rates based on inventory. (It is common to test significance by applying the “95% confidence” rule. This equates to a p-value of 0.05 or less.)

Analysis of Variance:

Field	DF	SSE	MSE	F	P
Product Inventory	8	0.000741799	< 0.0001	1.92971	0.0671058
YEAR(Order Date)	12	0.0041986	0.000349883	7.28148	< 0.0001

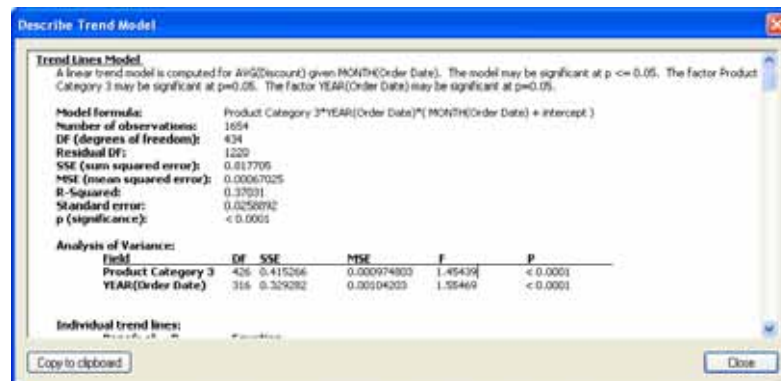
Still on the search to find out what is causing the high discount rates, we decide that it could be based on product categories. Certain categories may have higher discounts applied. It seems plausible so we put Product Category 3 on the rows shelf.



When we open the Describe Trend Line Model dialog box, and look at the ANOVA table this time, we see that the p-value for product category 0.0001 and the p-value for the entire model is less than 0.0001. With that, we can be confident that the changes in average discount rates over time are in fact quite different for different products. In other words we can be statistically confident that the different trend lines slopes for different products isn't just due to randomness, but rather there is a real correlation between a product and the slope of this line.

Analysis of Variances					
Field	DF	SS	MSE	F	P
Product Category 3	426	0.415266	0.000974003	1.45439	< 0.0001
YEAR(Order Date)	316	0.329282	0.00104203	1.55469	< 0.0001

In addition to product category 3, we notice that the year field offers a statistically significant improvement to the model. Refer to "Entire Model Significance" on page 18-13 for more information.



## Assumptions

The p-values reported in Tableau trend lines depend on some assumptions about the data. The first assumption is that, whenever a test is performed the model for the mean is (at least approximately) correct.

The second assumption is that the "random errors" referred to in the model formula (refer to "The Trend Line Model" on page 18-8) are independent across different observations and that they all have the same distribution. This would be violated if the response variable had much more variability around the true trend line in one category vs. another. For instance, in the example shown in this chapter, this would be violated if the discount rate had much more variability around the true trend line in one market vs. another.

## Trend Line Model Terms

When you describe the trend line model, there are several values listed. This section discusses what each of these values means.

### Model Formula

This is the formula for the full trend line model. The formula reflects whether you have specified to exclude factors from the model.

### Number of Observations

This is the number of rows used in the view.

### Residual DF (residual degrees of freedom)

For a fixed model, this value is defined as the number of observations minus the number of parameters estimated in the model.

### DF (degrees of freedom)

In the ANOVA table, this column refers to the degrees of freedom of the smaller model in the given row.

### SSE (sum squared error)

The term SSE generally refers to a “sum of squared errors.” The errors are the difference between the observed value and the value predicted by the model. In the ANOVA table, this column is actually the difference between the SSE of the simpler model in that particular row and the full model, which uses all the factors. This SSE also corresponds to the sum of the differences squared of the predicted values from the smaller model and the full model.

### MSE (mean squared error)

The term MSE refers to “mean squared error” which is the SSE quantity divided by its corresponding degrees of freedom.

### R-Squared

A measure of the percent of the variability explained by the full model as compared to a smaller model.

If  $SSE(F)$  is the sum of squared errors of the full model, and  $SSE(S)$  the sum of squared errors from the smaller model, then

$$R^2 = (SSE(S) - SSE(F)) / SSE(S)$$

measures the relative drop in variability of the response when using the full model as opposed to the smaller model. In the Describe Trend Line Model dialog box, this model is the model that has a different mean for each pane. This is different than the most common use of R-squared, which compares the model to a model with just a grand mean.

### **Standard error**

The square root of the MSE of the full model. It is an estimate of the standard deviation (variability) of the "random errors" in the model formula.

### **P (significance)**

The probability that an F random variable with the above degrees of freedom exceeds the observed F in this row of the ANOVA table.

### **Analysis of Variance**

This table, also known as the ANOVA table, lists information for each factor in the trend line model. The values are a comparison of the model without the factor in question to the entire model, which includes all factors.

### **Individual trend lines**

This table provides information about each trend line in the view. For each trend line, you can see its p-value as well as its equation. This is helpful when you have a lot of panes in the view, each with its own trend line. Looking at the list you can see which, if any, are the most statistically significant.

## Commonly Asked Questions

This section describes some commonly asked questions regarding trend lines in Tableau.

### **How do I change the confidence level used in the model?**

Tableau does not enforce a confidence level. It simply reports the significance of the whole model, or of a specific field, by showing the p-value. A p-value of 0.05, for instance, implies that you can be 95% sure of the model. A p-value of 0.10 means that you can be 90% confident of the model's prediction. So read the reported p-values, and then use whatever rules you wish to determine statistical significance.

### **What does it mean if the p-value for the model is significant but the p-value for the specific field in the ANOVA table is not significant?**

The p-value in the ANOVA table indicates whether the field adds or detracts from the significance of the entire model. The smaller the p-value the more the field improves the model. The values displayed for each field are derived by comparing the entire model to a model that does not include the field in question. So, this means that the model is statistically significant, but that you cannot be confident that the specific field in question adds anything to it. One thing to inspect in a case like this is whether you are better off removing the factor from the model.

### **What does it mean if the p-value for the model is not significant but the p-value for the specific field in the ANOVA table is significant?**

This could happen in a case when there is no "trend" within each pane. For example, the lines are flat, but the mean varies across a given factor.



## Log Axes

---

---

## Overview


Sometimes you will have a measure that uses a logarithmic scale as opposed to linear. For example, some well known logarithmic scales include the Richter magnitude scale to measure the strength of earthquakes, pH to measure acidity, and the stellar magnitude scale, which measures the brightness of stars. You can Edit the axis scale for any measure to be logarithmic using the Edit Axis dialog box. By default the tick marks are drawn at powers of ten, however, you can specify any base that is greater than 1.

### To change the scale of an axis:


- 1 Right-click the axis in the view and select **Edit Axis**.



- 2 In the Edit Axis dialog box, select **Logarithmic** for the axis scale.



- 3 Select the **Tick Marks** tab.



---

4 Select one of the following Major Tick mark options:

- **Automatic** - the major tick marks are drawn at powers of 10.
- **Fixed** - the major tick marks are drawn at a specified exponent. Type a number into the Powers of text box.
- **None** - major tick marks are not shown.

5 When finished, click **OK**.

---

**Note** You can also reverse the axis by selecting **Reverse** in the Scale area on the General tab of the Edit Axis dialog box.

---

# Formatting

---

<b>Overview . . . . .</b>	<b>20-3</b>
<b>Worksheet Level Formatting . . . . .</b>	<b>20-4</b>
Opening the Format Window . . . . .	20-4
Fonts . . . . .	20-5
Alignment . . . . .	20-8
Shading . . . . .	20-9
Borders . . . . .	20-11
Lines . . . . .	20-11
<b>Formatting Specific Parts of the View . . . . .</b>	<b>20-13</b>
Fields . . . . .	20-13
Numbers . . . . .	20-15
Marks . . . . .	20-17
Field Labels . . . . .	20-21
Legends . . . . .	20-22
Title and Caption . . . . .	20-22
<b>Copying and Pasting Formatting . . . . .</b>	<b>20-24</b>
<b>Clearing Formatting . . . . .</b>	<b>20-25</b>
<b>Workbook Themes . . . . .</b>	<b>20-26</b>
<b>Resizing the Table . . . . .</b>	<b>20-27</b>
Resizing Rows and Columns . . . . .	20-27
Resizing the Entire Table . . . . .	20-29
Resizing Cells . . . . .	20-31
Reordering Rows and Columns . . . . .	20-33
Including and Excluding Rows and Columns . . . . .	20-33
Resize Keyboard Shortcuts and Commands . . . . .	20-34
Clearing Manual Sizing . . . . .	20-34
<b>Editing Axes . . . . .</b>	<b>20-37</b>

Changing the Axis Range . . . . .	20-37
Changing the Axis Appearance . . . . .	20-45
Formatting Tick Marks . . . . .	20-47
<b>Formatting Mark Colors . . . . .</b>	<b>20-49</b>
Categorical Colors . . . . .	20-49
Quantitative Colors . . . . .	20-51

---

## Overview

Formatting is an important part of both your analysis and presentation. You can format almost everything you see on a worksheet including the fonts, shading, alignment, borders, and graph lines. For example, in a text table you may want to add banded shading to help you visually separate consecutive groups of rows or columns. In a scatter view with reference lines you may want to change the line thickness and color. All of these settings can be changed using the Format window.

Most often you will want to specify format settings for the entire worksheet, all rows, or all columns. However, Tableau also allows you to format individual parts of the view as well. For example, you can format specific fields, resize the cells and the table, and edit individual axes.

This section discusses the following topics:

- Worksheet Level Formatting
- Formatting Specific Parts of the View
- Copying and Pasting Formatting
- Clearing Formatting
- Resizing the Table
- Editing Axes
- Formatting Mark Colors

## Worksheet Level Formatting

Most commonly you will want to specify format settings at the worksheet level. For example, you may want to use a specific font for data labels, remove all the borders in a text table, or add shading to every other column in a view. These format settings can be specified using the Format window.

The Format window, when opened, replaces the Data window on the left side of the worksheet. There you can use a series of drop-downs to specify format settings for either the entire sheet, all rows, or all columns. This section discusses opening the Format window and the available settings in the following sections:


- Opening the Format Window
- Fonts
- Alignment
- Shading
- Borders
- Lines

### Opening the Format Window

Use the Format menu to open the Format window.

**To open the Format window:**


- Select Format and then select the part of the view you want to format.



The Format window opens on the left side of the workbook, replacing the Data window. At the top of the Format window, there is a toolbar where you can quickly switch between each of the types of format settings available.



The Format window also contains three tabs: Sheet, Rows, and Columns. Switch between these tabs to apply formatting to the entire sheet, just the rows, or just the columns.



## Fonts

In the Format window, select the Font command **A** on the toolbar to see the Font format settings. You can specify font properties for the entire worksheet, just the rows, or just the columns.

### Sheet Font Settings

For the entire worksheet, you can specify the font, style, size, and color for both the pane text and header text. For example, in the view below, the header text is set to be bold.



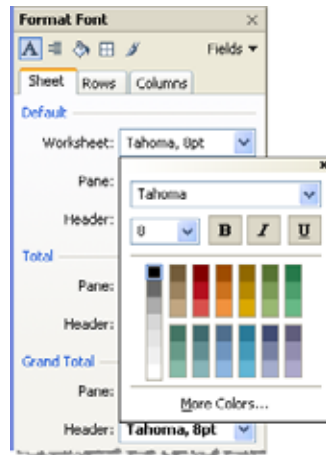


If you have totals or grand totals in the view, you can specify special font settings so that these values can stand out from the rest of the data. This is particularly useful when you are working with a text table. The view below shows a text table in which the grand totals are formatted to be bold and dark red.

The screenshot shows the 'Format Font' dialog box with the 'Grand Total' section highlighted. The text table displays 'Product Category 2' for four years (2001-2004) and a 'Grand Total' column. The 'Grand Total' row is formatted to be bold and dark red.

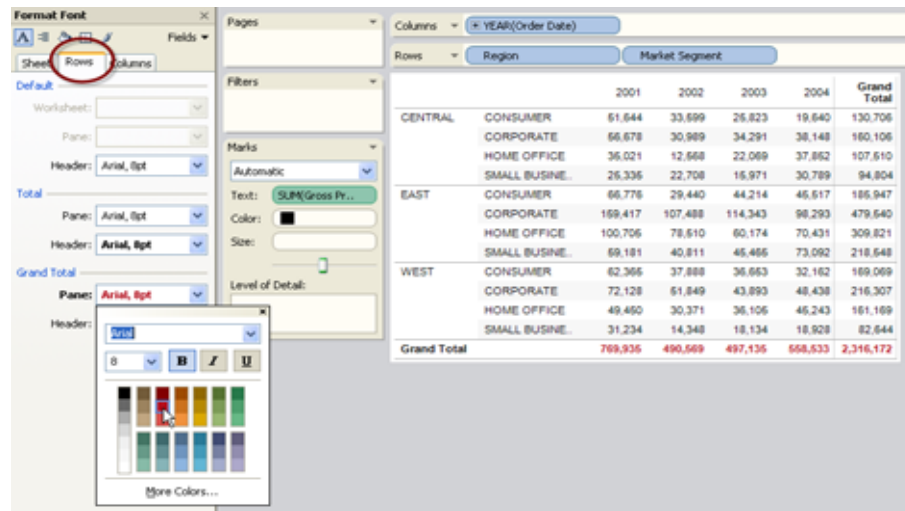
	2001	2002	2003	2004	Grand Total
APPLIANCES	10,371	-115	7,368	5,485	23,129
BINDERS AND BNL	8,293	12,137	4,871	6,388	30,689
BOOKCASES	3	-5,836	-1,067	-2,863	-9,561
CHAIRS & CHARM	41,777	8,823	18,519	18,363	80,679
COMPUTER PERIP	46,776	26,060	33,096	26,508	130,438
COPERS AND FAX	96,968	15,490	18,797	33,242	164,517
ENVELOPES	7,916	12,305	7,412	3,355	31,609
LABELS	9,519	10,263	8,196	8,423	36,400
OFFICE FURNISHI	2,600	870	-419	-2,977	-326
OFFICE MACHINES	93,664	-110	18,592	-26	112,110
PAPER	44,246	34,179	40,615	33,122	152,162
PENS & ART SUPP	429	2,365	66	1,406	4,243
RUBBER BANDS	431	423	304	-13	1,144
SCISSORS, RULER	-1,065	-817	-781	-267	-2,711
STORAGE & ORGA	6,534	2,242	833	1,902	11,211
TABLES	7,081	-12,946	-12,518	-10,632	-28,915
TELEPHONES AND	398,972	368,603	366,342	440,908	1,879,825
<b>Grand Total</b>	<b>788,936</b>	<b>490,969</b>	<b>687,138</b>	<b>868,533</b>	<b>2,315,172</b>

Finally, you can use the Worksheet drop-down to specify the properties of all text in the worksheet.




### Row and Column Font Settings

Switch to the Rows or the Columns tabs to specify font properties for just the rows or just the columns. Here you have the same options as you do for the Sheet in that you can modify the font, style, size, and color for both the pane text and the header text. In the view below the Grand Totals are formatted red for just the Rows. Notice that the Grand Totals for Columns are not affected by this setting.



However, sometimes the settings don't make sense for particular views and are greyed out. For example, in the view above, the default pane text is in both rows and columns. Specifying font properties on a row or column level does not make sense, to change the pane text, switch back to the Sheet tab.

## Alignment

In the Format window, select the Alignment command  on the toolbar to see the Alignment format settings. You can specify alignment properties for the entire worksheet, just the rows, or just the columns using the tabs across the top of the Format window.


Using the alignment drop-downs you can specify the text alignment for both the pane and header text. For each of these text areas you can specify the following alignment options:

- Horizontal - controls whether the text aligns on the left side, right side, or is centered.
- Vertical Alignment - controls whether the text aligns across the top, middle, or bottom.
- Direction - rotates the text so that it runs vertically or horizontally.

- Wrap - controls whether long headers and pane text wrap to the next line rather, than being abbreviated.



## Shading

In the Format window, select the Shading command  on the toolbar to see the Shading format settings. The shading settings control the background color of the pane and headers for normal areas, totals, and grand totals. In addition, you can add row or column banding.


### Row and Column Banding Settings

Sometimes, rather than selecting a single background color for of a pane or header, you may want to alternate the color from row to row or column to column at varying intervals. This kind of shading is called banding. Banding is particularly useful when you are working with a text table. Adding alternating shading can help you distinguish between consecutive rows or columns.

In the Format window, the banding settings allow you to select a color, the size, and level at which you want to apply the banding. Each of these properties are described below.

- Selecting a Color - select the color you want the bands to be using the drop down for either the pane or the header areas.

- Selecting a Band Size - the size refers to the frequency of bands. For example, rather than shade every other row, you may want to shade every three rows. Slide the size selector right and left to specify the band size.



		Year of Order Date		
		2001	2002	2003
FURNITURE	BOOKCASES	0.0001	-0.0001	-0.0028
	CHAIRS & CHARMATS	0.1836	0.0588	0.1378
	OFFICE FURNISHINGS	0.0470	0.0396	-0.0166
	TABLES	0.0289	-0.1894	-0.1783
OFFICE SUPPLIES	APPLIANCES	0.1988	-0.0063	0.2989
	BINDERS AND BINDER AC.	0.1706	0.2675	0.1883
	ENVELOPES	0.3786	0.4328	0.4149
	LABELS	0.4683	0.4961	0.4746
	PAPER	0.2320	0.2037	0.2374
	PENS & ART SUPPLIES	0.0210	0.0790	0.0037
	RUBBER BANDS	0.1008	0.1162	0.1087
	SCISSORS, RULERS AND	-0.1473	-0.0740	-0.3197

		Year of Order Date		
		2001	2002	2003
FURNITURE	BOOKCASES	0.0001	-0.0001	-0.0028
	CHAIRS & CHARMATS	0.1836	0.0588	0.1378
	OFFICE FURNISHINGS	0.0470	0.0396	-0.0166
	TABLES	0.0289	-0.1894	-0.1783
OFFICE SUPPLIES	APPLIANCES	0.1988	-0.0063	0.2989
	BINDERS AND BINDER AC.	0.1706	0.2675	0.1883
	ENVELOPES	0.3786	0.4328	0.4149
	LABELS	0.4683	0.4961	0.4746
	PAPER	0.2320	0.2037	0.2374
	PENS & ART SUPPLIES	0.0210	0.0790	0.0037
	RUBBER BANDS	0.1008	0.1162	0.1087
	SCISSORS, RULERS AND	-0.1473	-0.0740	-0.3197

- Selecting a Level - when you have nested tables where you have multiple dimensions on the rows and columns shelves, you may want to add banding at a particular level. For example, in the view below, both Year and Quarter are on the Columns shelf. When you set banding to the first level, the shading alternates between each member of the Quarter dimension. However, if you set banding to the second level, the shading alternates between each member of the Year dimension.


Slide the level selector right and left to specify the banding level.



		Year of Order Date : Quarter of Order Date							
		2000				2001			
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
APPLIANCES		194	1,800	8,080	1,341	2,847	1,77		
BINDERS AND B		1,043	1,893	2,126	20	2,192	1,12		
BOOKCASES		-201	-461	-603	128	-773	-10		
CHAIRS & CHAI		8,162	-224	8,771	2,916	2,678	4,46		
COMPUTER PR.		12,433	7,912	9,703	3,047	8,627	7,72		
COPIERS AND P		9,862		4,439	8,796	967	9,86		
ENVELOPES		3,369	1,898	-81	2,377	330			
LABELS		2,189	2,822	1,974	1,430	2,238	2,39		
OFFICE FURNIS.		148	242	-838	-170	-1,168	-10		
OFFICE MACH.		4,167		-13	14,447		-26		
PAPER		12,738	7,967	14,262	8,628	7,160	8,49		
PENS & ART SU.		48	-124	271	-138	-107	93		
RUBBER BANDS		142	116	47	0	38			
SCISSORS, RUL.		-47	-471	-17	-246	117	1		
STORAGE & CR.		3,408	-2,734	-170	38	435	-1,30		
TABLES		-2,729	-1,866	-8,898	-2,237	-9,017	-2,29		
TELEPHONES &		91,779	96,943	96,969	80,868	94,679	106,0		


		Year of Order Date : Quarter of Order Date							
		2000				2001			
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
APPLIANCES		194	1,800	8,080	1,341	2,847	1,769	980	777
BINDERS AND B		1,043	1,893	2,126	20	2,192	1,159	804	1,163
BOOKCASES		-201	-461	-603	128	-773	-102	-886	-1,632
CHAIRS & CHAI		8,162	-224	8,771	2,916	2,678	4,888	3,241	8,886
COMPUTER PR.		12,433	7,912	9,703	3,047	8,627	7,776	4,272	7,539
COPIERS AND P		9,862		4,439	8,796	967	9,869		23,401
ENVELOPES		3,369	1,898	-81	2,377	330	787	2,302	-6
LABELS		2,189	2,822	1,974	1,430	2,238	2,887	1,949	1,871
OFFICE FURNIS.		148	242	-838	-170	-1,168	-108	-1,088	-826
OFFICE MACH.		4,167		-13	14,447				-26
PAPER		12,738	7,967	14,262	8,628	7,160	8,409	8,188	10,988
PENS & ART SU.		48	-124	271	-138	-107	870	440	493
RUBBER BANDS		142	116	47	0	38	-86	87	-23
SCISSORS, RUL.		-47	-471	-17	-246	117	182	-282	-276
STORAGE & CR.		3,408	-2,734	-170	38	436	-4,916	4,869	-1,362
TABLES		-2,729	-1,866	-8,898	-2,237	-9,017	-2,782	-4,288	-897
TELEPHONES &		91,779	96,943	96,969	80,868	94,679	106,240	116,700	121,480

Borders


In the Format window, select the Borders command  on the toolbar to see the Borders format settings. Borders are the lines that surround the table, pane, cells, and headers in a view. You can specify the border style, width, and color for the cell, pane, and header areas. Additionally, you can format the row and column dividers.

Row and Column Divider Settings

Row and column dividers are a good way to visually break up a view and are most commonly used in nested text tables. You can modify the style, width, color, and level of the borders that divide each row or each column using the row and column divider drop-downs. The level refers to the header level you want to divide by. For example, in the view below, there are three dimensions on the rows shelf: Market Segment, Region, and Product Category 1. When adding a row divider you can switch between dividing every product, every region, or every market segment by sliding the level selector right and left.


Level: 

		Year of Order Date			
		2001	2002	2003	2004
CONSUMER	CENTRAL FURNITURE	2,816	3,406	-2,066	-120
	OFFICE SUPPLIES	5,537	7,789	4,244	2,423
	TECHNOLOGY	43,292	22,404	23,636	17,338
	EAST FURNITURE	4,906	-2,677	6,776	-2,663
	OFFICE SUPPLIES	8,411	-1,899	6,706	1,107
	TECHNOLOGY	53,459	34,017	32,734	47,063
	WEST FURNITURE	4,973	2,434	666	920
	OFFICE SUPPLIES	6,491	6,867	6,138	186
	TECHNOLOGY	50,902	29,598	29,861	31,057
	CENTRAL FURNITURE	6,243	-1,060	499	2,702
	OFFICE SUPPLIES	6,616	8,989	3,316	4,678
	TECHNOLOGY	44,820	23,049	30,476	30,768

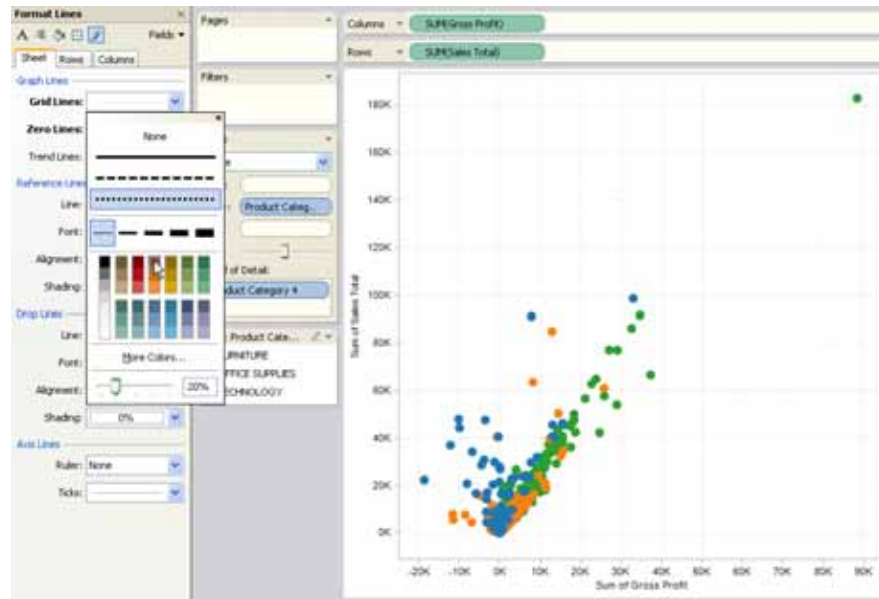
Level: 

		Year of Order Date			
		2001	2002	2003	2004
CONSUMER	CENTRAL FURNITURE	2,816	3,406	-2,066	-120
	OFFICE SUPPLIES	5,537	7,789	4,244	2,423
	TECHNOLOGY	43,292	22,404	23,636	17,338
	EAST FURNITURE	4,906	-2,677	6,776	-2,663
	OFFICE SUPPLIES	8,411	-1,899	6,706	1,107
	TECHNOLOGY	53,459	34,017	32,734	47,063
	WEST FURNITURE	4,973	2,434	666	920
	OFFICE SUPPLIES	6,491	6,867	6,138	186
	TECHNOLOGY	50,902	29,598	29,861	31,057
	CENTRAL FURNITURE	6,243	-1,060	499	2,702
	OFFICE SUPPLIES	6,616	8,989	3,316	4,678
	TECHNOLOGY	44,820	23,049	30,476	30,768

Lines

In the Format window, select the Lines command  on the toolbar to see the Lines format settings. The lines settings control the lines that are part of the graph such as grid lines and zero lines as well as lines that help you inspect data such as trend lines, reference lines, and drop lines. You can specify the style, width, and color for each of these lines. Additionally,

you can specify font, alignment, and shading settings for reference line and drop line labels. For example, in the view below, the grid lines are turned on and zero lines turned off.



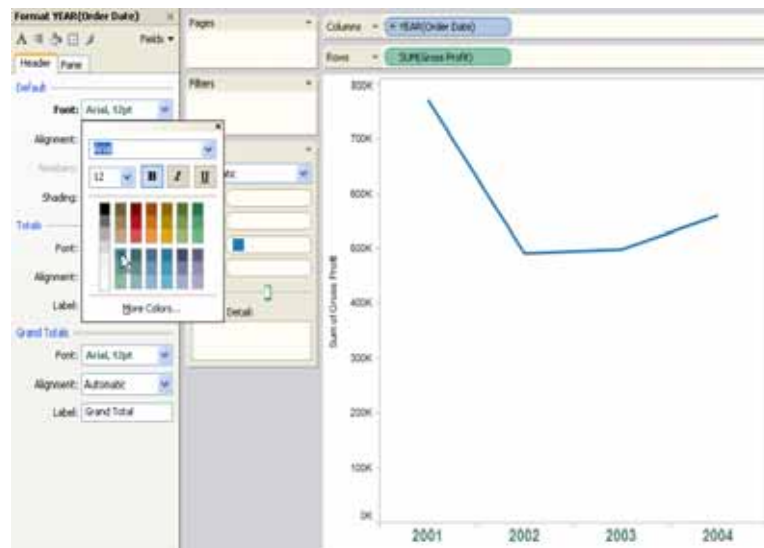
## Formatting Specific Parts of the View

Although you will more commonly want to apply format settings on a large scale such as the entire worksheet, all rows, or all columns, sometimes you may want to format specific parts of the view. You can specify individual format settings by right-clicking a specific part of the view and selecting Format. This section discusses how to format the following parts of the view:

- Fields
- Numbers
- Marks
- Legends
- Title and Caption

### Fields

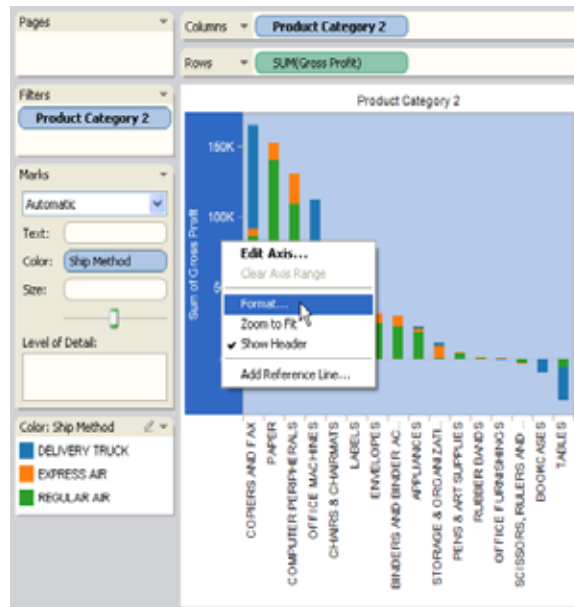
Rather than formatting all rows or all columns in a worksheet, you can specify format settings that only apply to a specific field in the view. For example, in the view below, the Year(Order Date) field has been formatted so that the headers are 12pt, bold, and blue. Notice that the header values along the Gross Profit axis are not affected.





### To format a specific field:

- 1 Right-click the part of the view you want to format and select **Format**.



The Format window opens containing settings relevant to the selected field.

- 2 Make changes in the Format window as necessary.

For discrete fields such as Region or Customer you can specify font and alignment properties for both header and pane areas. For continuous fields such as Profit or Sales you can format font properties for the pane and axis as well as number format and tick mark colors. Refer to “Editing Axes” on page 20-37 to learn more about other axis options. The view is updated as you make changes so you can quickly see the colors and formats that work with your view.

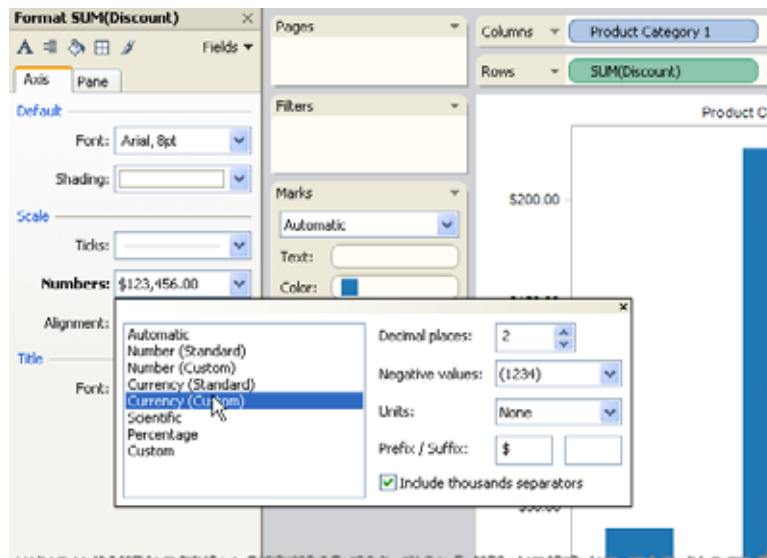
- 3 When finished, click the ‘x’ in the upper right corner of the Format window to return to the Data window.

## Numbers

When you format a measure, you can specify the number format for both the axis and the pane text. You can select from a set of standard formats, such as number, currency, scientific, and percentage; or you can define a custom number format using Microsoft Excel style format codes.

### To specify number format:

- 1 Right-click a measure or axis in the view and select **Format**.
- 2 In the Format window, click the **Numbers** drop-down.
- 3 Select a number format. Some formats require additional settings. For example, if you select Currency (Custom), you must also specify the number of decimal places, how to treat negative values, the units, a prefix or suffix, and whether to include thousands separators.



Below is a list of the number formats and associated options available in Tableau

Number Format	Format Options
<b>Automatic:</b> format is automatically selected based on either the format specified by the data source or the data contained in the field.	None
<b>Number (Standard):</b> format is based on locale selected.	<b>Locale:</b> number format changes based on the geographical location selected.
<b>Number (Custom):</b> format is customized to your liking.	<p><b>Decimal Places:</b> the number of decimal places to display.</p> <p><b>Negative Values:</b> how negative values will be displayed.</p> <p><b>Units:</b> the number will be displayed using the specified units. For example, if the number is 20,000 and the units are thousands, the number will be displayed as 20K.</p> <p><b>Prefix/Suffix:</b> characters that will precede and follow each displayed number.</p> <p><b>Include thousands separators:</b> whether the number will show separators every thousand (example: 100,000 vs. 100000).</p>
<b>Currency (Standard):</b> format and currency symbol is based on locale selected.	<b>Locale:</b> currency format based on the geographical location selected.

Number Format	Format Options
<b>Currency</b> <b>(Custom):</b> format and currency symbol is customized to your liking.	<b>Decimal Places:</b> the number of decimal places to display. <b>Negative Values:</b> how negative values will be displayed. <b>Units:</b> the number will be displayed using the specified units. For example, if the number is 20,000 and the units are thousands, the number will be displayed as 20K. <b>Prefix/Suffix:</b> characters that will precede and follow each displayed number. <b>Include thousands separators:</b> whether the number will show separators every thousand (example: 100,000 vs. 100000).
<b>Scientific:</b> numbers are displayed in scientific notation.	<b>Decimal:</b> the number of decimal places to display.
<b>Percentage:</b> numbers are displayed as a percentage with the percent symbol. The value of 1 is interpreted as 100% and 0 as 0%.	<b>Decimal:</b> the number of decimal places to display.
<b>Custom:</b> format is based entirely on what is specified in the format options.	<b>Custom:</b> type in the format you want to use. This format can be specified by an Excel style number code.

## Marks

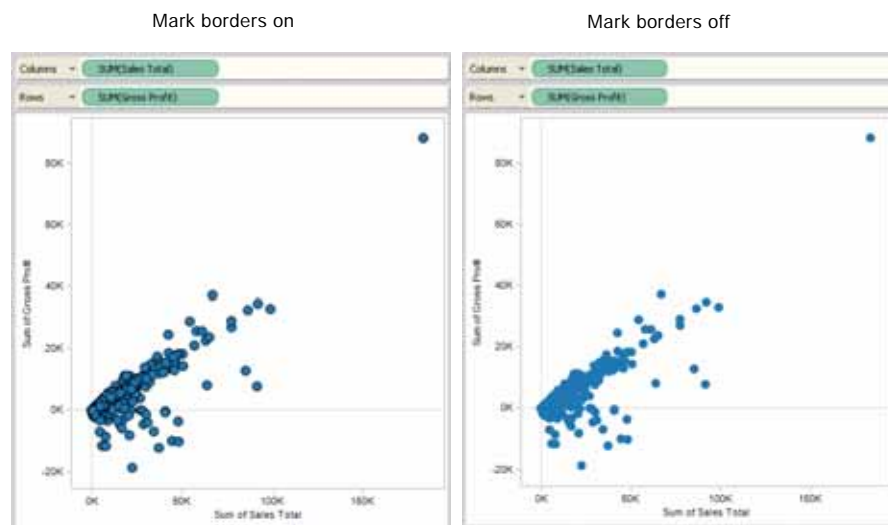
All views contain marks to represent the data visually. There are several different kinds of marks that you can select from depending on your data and how you want to view it. Once you have the view created, you can format the marks by turning the mark borders on or off, and specifying other mark options available for specific mark types. This section discusses the following topics:

- Mark Borders
- Mark Halos
- More Mark Formatting Options

## Mark Borders

By default, Tableau displays all marks without a border. You can turn on the mark borders at any time by selecting the **Show Mark Borders** option on the **Format** menu. All but the text, line, and shape mark types can have a border turned on.


Borders are often useful for distinguishing between closely spaced marks. For example, the view shown below has mark borders turned on (left) and turned off (right). As you can see, when borders are turned off, the marks become indistinguishable in the areas where they are tightly clustered.




Leaving mark borders off is particularly useful when you are viewing a large number of small marks that are color-encoded. It can be difficult to see the color encoding because the borders dominate the marks.

For example, the view shown below displays bars that are segmented by a large number of color-encoded dimension members. As you can see, when mark borders are turned on some marks are difficult to identify by color. When borders are turned off, the marks can easily be distinguished.


Mark borders on



Mark borders off



You can also turn mark borders on and off using the title menu on the Marks card.



## Mark Halos

In order to make the marks in a view more visible when placed on top of a background image or dark background, each mark is surrounded by a solid contrasting color called a halo. You can turn mark halos on and off as well as specify their color using the Format menu.

**To turn halos on or off:**

- Select **Format > Show Mark Halos**.

### To specify the halo color:

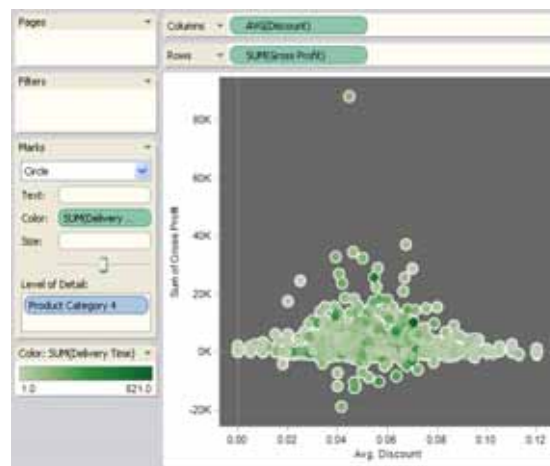
- 1 Select **Format > Halo Color > Set**.
- 2 In the Color dialog box, select a new color.
- 3 When finished, click **OK**.

---

**Note** Reset the halo color to the default color by selecting **Format > Halo Color > Reset**.

---

The view below has a dark shading with white mark halos turned on.



### More Mark Formatting Options

To further refine your view, you can modify the types of marks as well as the shape and size encoding. For more information about these options refer to “Marks” on page 9-11. You can also specify new color encodings. Refer to “Formatting Mark Colors” on page 20-49 to learn more.

## Field Labels

Field labels are row and column headings that indicate the data fields used to create the Table. By default, field labels are shown but you can choose to hide them. When field labels are showing they display in three different parts of the view: rows, columns, and the corner. The view below shows an example of each of these types of field labels.

Row Field Labels			Corner Field Labels			Column Field Labels					
Columns	Region		YEAR(Order Date)								
Rows	Market Segment		Product Family		Product Category						
			Region / Year of Order Date								
			EAST				WEST				
			2001	2002	2003	2004	2001	2002	2003	2004	
CONSUMER	OFFICE SUPPLIES	Product Family	Product Category								
		APPLIANCES	3,250	823	179	326		811		4,633	
		ENVELOPES	4,241	466	2,829	798	2,744		2,744	117	
		LABELS	992	2,461	752	838	309	291	309	678	
		PAPER	19,316	9,581	16,579	10,463	16,967	10,632	16,368	6,593	
	TECHNOLOGY	PENS & ART SUPPLIES	2,319	1,718	2,096	2,560	970	254	391	996	
		RUBBER BANDS	716	306	285	369	626	377	501	42	
		COMPUTER PERIPHERALS	16,729	9,693	9,327	6,402	8,175	10,443	7,646	4,906	
		COPERS AND FAX	9,256			16,193	41,867		12,602		
		OFFICE MACHINES	23,132	621	6,183		13,312				
Total		79,951	25,658	38,231	37,939	83,849	22,708	39,451	17,965		
CORPORATE	OFFICE SUPPLIES	APPLIANCES	10,137	2,345	9,115	2,945	1,721	2,439	179	9,504	
		ENVELOPES	7,371	448	7,150	378		952			
		LABELS	4,971	3,896	4,624	3,640	1,936	1,264	1,864	1,667	
		PAPER	35,376	33,032	30,770	44,360	13,775	12,467	12,326	16,421	
		PENS & ART SUPPLIES	4,628	7,076	2,690	3,483	1,884	3,197	1,462	1,026	
	TECHNOLOGY	RUBBER BANDS	851	489	407	803	407	175	233	47	
		COMPUTER PERIPHERALS	23,121	13,429	16,709	14,673	7,688	6,077	3,811	8,166	
		COPERS AND FAX	22,626	15,673		16,816	30,878				
		OFFICE MACHINES	64,549		14,279		16,374	670	290		
		Total		173,631	76,167	85,644	87,088	74,664	27,140	20,155	36,709

You can format the font, shading, alignment, and separators for each of these types of field labels.

### To format field labels:

- 1 Select **Format > Field Labels** or right-click a field label in the view and select **Format**.
- 2 In the Format window, specify setting the font, shading, and alignment of the field labels.



---

**Note** When you have multiple dimensions on the rows or columns shelves the field labels will be displayed adjacent to each other in the table. Each field label is separated from the others with a forward slash symbol. Specify a different separator in the Format window.

---

## Legends

When you encode the marks using the color and size shelves a legend card displays in the worksheet. You can format the legend font, shading, border, and alignment.

### To format legends:

- 1 Select **Format > Legends** or right-click the legend and select **Format**.
- 2 In the Format window, specify settings for the body and title of the legends.

---

**Note** The legend format settings apply to all legends, you cannot format individual legends separately.

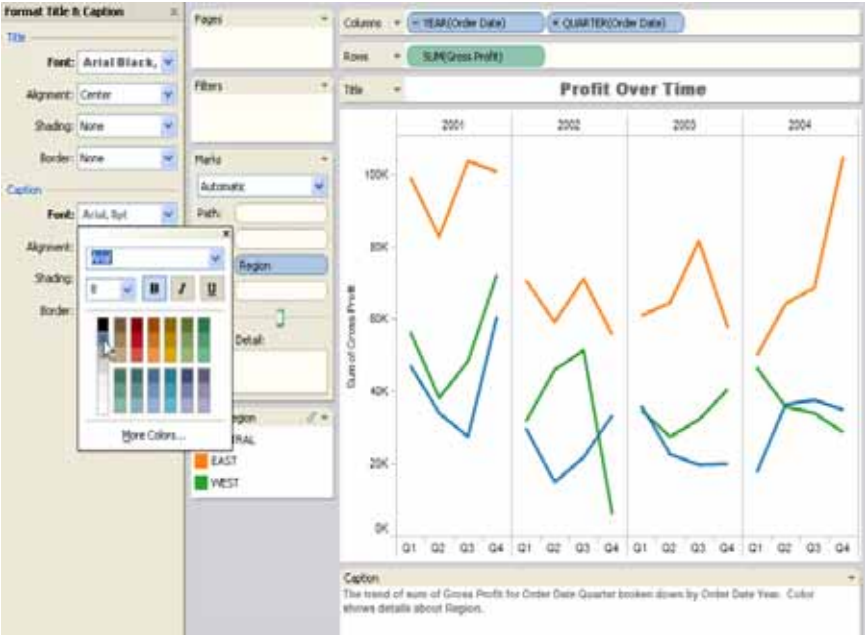
---

## Title and Caption

You can add titles and captions to any worksheet using the card menu on the toolbar. Refer to “Showing and hiding worksheet elements” on page 2-15 to learn more about turning these parts of the view on and off. After you add a title or caption, you can format its font, alignment, shading, and border using the Format window.

### To format titles and captions:

- 1 Select **Format > Title & Caption** or right-click the title or caption in the view and select **Format**.
- 2 Make the necessary changes in the Format window.



## Copying and Pasting Formatting

After you format a worksheet you can copy the format settings to one or more other worksheets. When you copy the formatting from a worksheet, it copies all of the format settings specified by the Format window. However, this command does not copy manual sizing, zoom settings, default label orientation, etc. Also formatting applied to individual reference lines and annotations is not copied with this command.

**To copy and paste formatting between worksheets:**

- 1 Select the worksheet from which you want to copy formatting.
- 2 Select **Edit > Copy Formatting**.
- 3 Select the worksheet you want to paste the formatting into.
- 4 Select **Edit > Paste Formatting**.

---

**Note** You can also copy and paste formatting using the worksheet tabs. Right-click the worksheet tab you want to copy the formatting from and select **Copy Formatting**. Then select one or more other worksheet tabs (hold the Ctrl key on your keyboard to select multiple tabs), right-click, and select **Paste Formatting**.

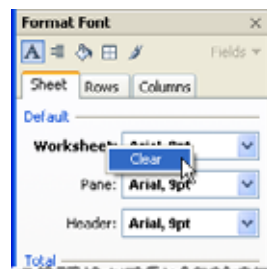
---

## Clearing Formatting

Any time you make changes to a setting in the Format window, the label of the setting is bolded to indicate that it is not the default.


### To clear individual settings in the Format window:

- In the Format window, right-click the label of the setting you want to clear and select Clear.



### To clear all custom formatting in the worksheet:

- On the toolbar, click the clear menu and select **Clear Formatting** on the drop-down.



**Note** You can also click the Clear button at the bottom of the Format window to clear all format settings currently showing in the Format window.



## Workbook Themes

The default formatting follows best practices for displaying information graphically and improves how the graphics look when exported to the web, PowerPoint, PDF, or Office documents. This default formatting is controlled by the workbook theme. All workbooks created using the latest version of Tableau Visual Explorer use the Default workbook theme. However, older Tableau Visual Explorer workbooks may use the classic theme. You can upgrade older workbooks to use the default theme.

### To upgrade workbook themes:

- 1 Open the workbook using Tableau Desktop v3.5 or later.
- 2 Select **Format > Workbook Theme > Default**.
- 3 Save the workbook.

Classic Theme

		2001	2002	2003	2004
FURNITURE	BOOKCASES	25,831	72,820	20,015	26,529
	CHAIRS & CHAIRMATS	255,343	116,387	113,460	125,410
	OFFICE FURNISHINGS	53,146	16,941	25,272	20,728
	TABLES	245,185	81,213	70,992	61,626
	<b>Total</b>	<b>579,505</b>	<b>287,361</b>	<b>229,740</b>	<b>234,294</b>
OFFICE SUPPLIES	APPLIANCES	52,220	18,348	24,961	35,350
	BINDERS AND BINDER ACCE..	48,623	45,457	29,296	34,831
	ENVELOPES	21,082	29,878	17,871	7,560
	LABELS	19,261	20,729	17,269	17,146
	PAPER	190,735	167,783	171,055	175,713
	PENS & ART SUPPLIES	20,410	29,794	14,874	17,971
	RUBBER BANDS	4,286	3,672	2,767	3,874
	SCISSORS, RULERS AND TRI..	7,167	8,341	2,443	6,282
	STORAGE & ORGANIZATION	164,866	111,297	79,331	106,895
	<b>Total</b>	<b>528,651</b>	<b>435,299</b>	<b>359,868</b>	<b>405,622</b>
TECHNOLOGY	COMPUTER PERIPHERALS	156,535	85,714	101,517	73,869
	COPIERS AND FAX	187,733	31,841	41,439	68,249
	OFFICE MACHINES	234,510	1,661	33,791	200
	TELEPHONES AND COMMUN..	1,115,386	1,064,627	1,002,580	1,224,694
	<b>Total</b>	<b>1,694,163</b>	<b>1,183,843</b>	<b>1,179,327</b>	<b>1,367,012</b>

Default Theme

		Year of Order Date			
Product Family	Product Category	2001	2002	2003	2004
FURNITURE	BOOKCASES	25,831	72,820	20,015	26,529
	CHAIRS & CHAIRMATS	255,343	116,387	113,460	125,410
	OFFICE FURNISHINGS	53,146	16,941	25,272	20,728
	TABLES	245,185	81,213	70,992	61,626
	<b>Total</b>	<b>579,505</b>	<b>287,361</b>	<b>229,740</b>	<b>234,294</b>
OFFICE SUPPLIES	APPLIANCES	52,220	18,348	24,961	35,350
	BINDERS AND BINDER AC..	48,623	45,457	29,296	34,831
	ENVELOPES	21,082	29,878	17,871	7,560
	LABELS	19,261	20,729	17,269	17,146
	PAPER	190,735	167,783	171,055	175,713
	PENS & ART SUPPLIES	20,410	29,794	14,874	17,971
	RUBBER BANDS	4,286	3,672	2,767	3,874
	SCISSORS, RULERS AND ..	7,167	8,341	2,443	6,282
	STORAGE & ORGANIZATI..	164,866	111,297	79,331	106,895
	<b>Total</b>	<b>528,651</b>	<b>435,299</b>	<b>359,868</b>	<b>405,622</b>
TECHNOLOGY	COMPUTER PERIPHERALS	156,535	85,714	101,517	73,869
	COPIERS AND FAX	187,733	31,841	41,439	68,249
	OFFICE MACHINES	234,510	1,661	33,791	200
	TELEPHONES AND COMM..	1,115,386	1,064,627	1,002,580	1,224,694
	<b>Total</b>	<b>1,694,163</b>	<b>1,183,843</b>	<b>1,179,327</b>	<b>1,367,012</b>

## Resizing the Table

Tableau allows you to change the size of the rows, columns, and cells that compose a table. The best way to resize your table depends on the view type and the table components you want to resize. This section discusses the following:

- Resizing Rows and Columns
- Resizing the Entire Table
- Resizing Cells
- Reordering Rows and Columns
- Including and Excluding Rows and Columns
- Resize Keyboard Shortcuts and Commands
- Clearing Manual Sizing

### Resizing Rows and Columns

Sometimes the rows and columns are not quite wide or tall enough. You can either resize rows and columns using the cell size commands on the Format menu or by manually dragging the header and axis borders in the view.

#### Using the Cell Size Commands

By selecting **Format > Cell Size** and then the **Taller**, **Shorter**, **Wider**, or **Narrower** menu items you can resize row and columns.

For example, suppose you want to increase the width of the columns and the height of the rows for the view shown below. You can use the **Taller** and **Wider** menu items or the keyboard shortcuts **Ctrl + up arrow** and **Ctrl + right arrow**. The views below use both these commands to make the view more readable.

Columns ▾

Rows ▾

		2003				2004	
		Q1	Q2	Q3	Q4		
CONSUMER	CENTRAL	9,495	12,816	3,631	-119	7	
	EAST	11,706	13,829	11,055	7,624	4	
	WEST	14,483	2,789	4,870	14,512	7	
CORPORATE	CENTRAL	10,158	4,319	12,800	7,214		
	EAST	19,336	32,801	34,932	27,274	11	
	WEST	8,679	20,109	8,867	5,237	15	
HOME OFFICE	CENTRAL	7,285	3,399	2,561	9,824	6	
	EAST	17,379	14,103	18,520	10,172		
	WEST	8,983	-145	12,235	18,031	18	
SMALL BUSINESS	CENTRAL	8,988	2,051	945	3,987	3	
	EAST	12,370	3,680	16,883	12,632	16	
	WEST	1,610	4,632	5,248	5,644	6	

Columns ▾

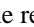
Rows ▾

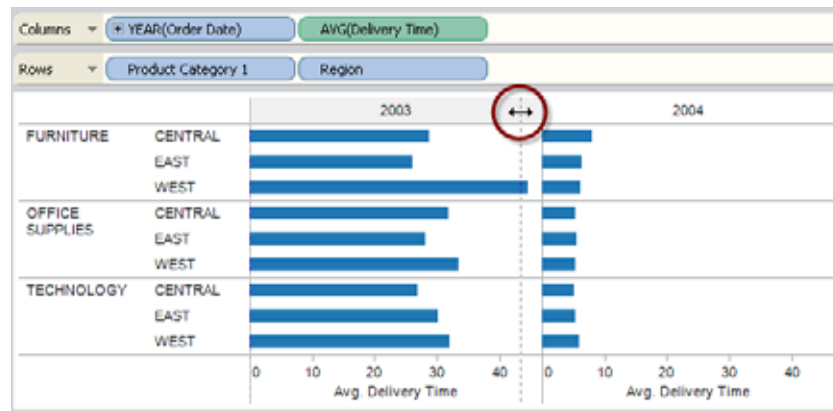
		2003				2004	
		Q1	Q2	Q3	Q4	Q1	Q2
CONSUMER	CENTRAL	8,486	12,816	3,631	-119	7,722	6,816
	EAST	11,706	13,829	11,055	7,624	4,849	8,906
	WEST	14,483	2,789	4,870	14,512	7,627	11,018
CORPORATE	CENTRAL	10,158	4,319	12,800	7,214	-133	16,187
	EAST	19,336	32,801	34,932	27,274	11,876	16,799
	WEST	8,679	20,109	8,867	5,237	13,948	10,680
HOME OFFICE	CENTRAL	7,285	3,399	2,561	9,824	6,696	3,948
	EAST	17,379	14,103	18,520	10,172	18,098	15,333
	WEST	8,983	-145	12,235	18,031	18,339	8,988
SMALL BUSINESS	CENTRAL	8,988	2,051	945	3,987	3,216	11,432
	EAST	12,370	3,680	16,883	12,632	16,900	21,434

**Note** For a given field, all members will have the same width and the same height. That is, you cannot resize individual field members.

## Manually Resizing Rows and Columns

To manually resize the widths or heights of row and column headers or axes:

- 1 Place your cursor over the vertical or horizontal border of a header or axis.
- 2 When you see the resize cursor , click and drag the border left and right or up and down.



## Resizing the Entire Table

You can increase or decrease the size of the entire table by selecting **Bigger** or **Smaller** on the **Format > Cell Size** menu. For example, to increase the width of the columns and the height of the rows for the view shown below, you can select **Format > Cell Size > Bigger** until the view is of the desired size. This option increases both the width and height of the panes in an intelligent way. Notice that the size of the row headers increase horizontally when you resize the table.

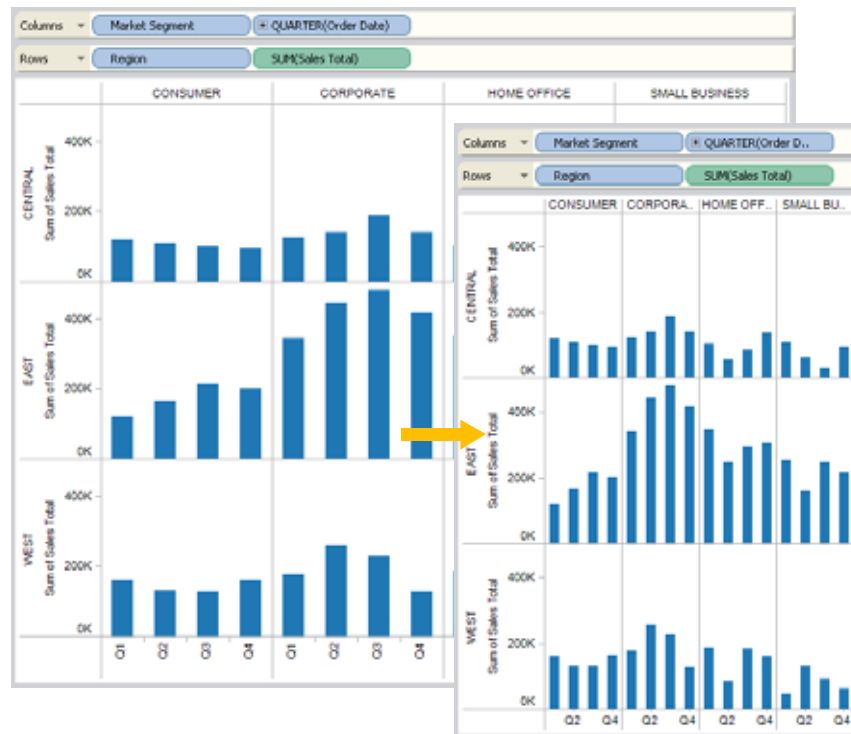


		2003				2004			
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
CONSUMER	CENTRAL	9,496	12,816	3,631	7	11,706	13,829	11,056	7
	EAST	11,706	13,829	11,056	7	14,483	2,789	4,870	14
	WEST	14,483	2,789	4,870	14	10,168	4,319	12,600	7
CORPORATE	CENTRAL	10,168	4,319	12,600	7	19,336	32,801	34,932	27
	EAST	19,336	32,801	34,932	27	9,679	20,109	8,867	6
	WEST	9,679	20,109	8,867	6	7,206	3,390	2,561	0
HOME OFFICE	CENTRAL	7,206	3,390	2,561	0	17,379	14,103	18,520	10
	EAST	17,379	14,103	18,520	10	8,983	-146	12,236	15
	WEST	8,983	-146	12,236	15	8,988	2,051	945	3
SMALL BUSINESS	CENTRAL	8,988	2,051	945	3	12,370	3,580	16,883	12
	EAST	12,370	3,580	16,883	12	1,610	4,632	6,248	6
	WEST	1,610	4,632	6,248	6				

		2003		
		Q1	Q2	Q3
CONSUMER	CENTRAL	9,496	12,816	3,631
	EAST	11,706	13,829	11,056
	WEST	14,483	2,789	4,870
CORPORATE	CENTRAL	10,168	4,319	12,600
	EAST	19,336	32,801	34,932
	WEST	9,679	20,109	8,867
HOME OFFICE	CENTRAL	7,206	3,390	2,561
	EAST	17,379	14,103	18,620
	WEST	8,983	-146	12,236
SMALL BUSINESS	CENTRAL	8,988	2,051	945
	EAST	12,370	3,580	16,883
	WEST	1,610	4,632	6,248

For the view shown below, you can select **Format > Cell Size > Smaller** to decrease the size of the table.



## Resizing Cells

Any table you can create in Tableau has the cell as its basic component. For a text table, the cell is what you would expect. It is the intersection of a row and a column, and is where the text is displayed.

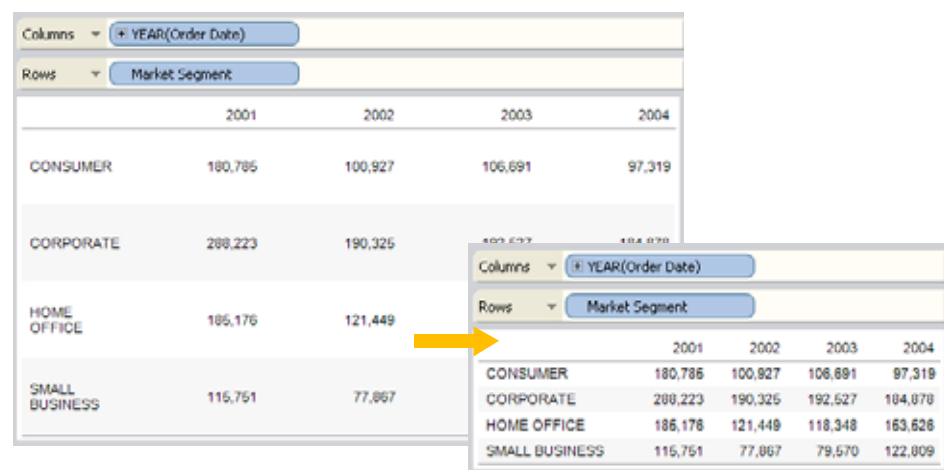
However, depending on the view you construct, identifying the cell is not always possible or useful like in the case of a scatter plot.

Manipulating cells to enhance your data view is useful when dimensions are the inner fields of both the Rows and Columns shelves. In this case, there are two shortcuts you can select on the **Format > Cell Size** menu:

- **Square Cell** – Adjusts the view so the cell has a 1:1 aspect ratio. This results in a square cell, which is particularly useful for heat maps.

- **Text Cell** – Adjusts the view so the cell has a 3:1 aspect ratio. This is particularly useful for text tables.

For example, the text table shown below is modified by selecting **Text Cell** from the **Format** menu. This enforces a cell aspect ratio of 3:1 and results in a compact table that is easy to read.



The image shows two versions of a table. The left version is a 'Text Cell' view where the cells are tall and narrow (3:1 aspect ratio). The right version is a 'Square Cell' view where the cells are more compact (1:1 aspect ratio). A yellow arrow points from the left table to the right table, indicating the transformation.

	2001	2002	2003	2004
CONSUMER	180,785	100,927	106,691	97,319
CORPORATE	288,223	190,325	192,527	184,878
HOME OFFICE	185,176	121,449	118,348	163,626
SMALL BUSINESS	115,751	77,867	79,570	122,809

The heat map shown below is modified by selecting **Square Cell** on the **Format > Cell Size** menu. This enforces a cell aspect ratio of 1:1 and results in a compact table that is easy to analyze. You can also use the **Size** slider to adjust the size of each mark.

The image shows two screenshots of a table interface. The top screenshot shows a table with columns for years (2001, 2002, 2003) and quarters (Q1, Q2, Q3, Q4), and rows for market segments (CONSUMER, CORPORATE, HOME OFFICE, SMALL BUSINESS). The bottom screenshot shows a zoomed-in section of the table with a yellow arrow pointing to it from the top view. The bottom view shows columns for years (2001, 2002, 2003, 2004) and quarters (Q1, Q2, Q3, Q4), and rows for market segments (CONSUMER, CORPORATE, HOME OFFICE, SMALL BUSINESS).

**Note** After changing the cell size, you can use **Ctrl+B** and **Ctrl+Shift+B** to decrease or increase the table size while maintaining the cell aspect ratio.

## Reordering Rows and Columns

The members in the view may not always be ordered exactly how you want it. You can reorder the rows and columns by clicking and dragging a header to a new position. Moving columns and rows around is equivalent to manually sorting. For more information about manually sorting your data refer to “Manual Sorting” on page 12-15.

## Including and Excluding Rows and Columns

Sometimes you will want to restrict certain members of a field from displaying. You can easily exclude a row or column by right-clicking the header and selecting **Exclude** on the context menu. Excluding a row or column creates a filter. For more information about filtering by selecting table headers refer to “Filter by Selecting Table Headers” on page 12-29.


## Resize Keyboard Shortcuts and Commands

Every view you build in Tableau will be different and require different sizing techniques. Using the keyboard shortcuts to resize rows and columns as well as the entire table makes building views more efficient. The table below defines some keyboard shortcuts and menu commands for common sizing actions.

Command	Keyboard Shortcut	Menu Command
Taller	CTRL + Up	Format > Cell Size > Taller
Shorter	CTRL + Down	Format > Cell Size > Shorter
Wider	CTRL + Right	Format > Cell Size > Wider
Narrower	CTRL + Left	Format > Cell Size > Narrower
Bigger	CTRL + SHIFT + B	Format > Cell Size > Bigger
Smaller	CTRL + B	Format > Cell Size > Smaller

## Clearing Manual Sizing

You can clear the custom sizing at any time using the Clear command on the toolbar. Select Clear Manual Sizing on the Clear drop-down.



You can also revert to the last saved state by selecting **File > Revert to Saved**. This option discards all unsaved changes including manual sizing. Refer to “Reverting Workbooks” on page 24-9.

## Miscellaneous Table Options

In addition to the standard formatting, there are some other settings that define the table structure. You can modify these settings using the Table Options dialog box. There you can specify the aspect ratio, the default number format, row and column attributes, and the default label orientation for labels along the bottom of the view. These settings apply to the whole view, however, some can be overridden by the changes made in the Format window.

### Setting Aspect Ratio

The aspect ratio refers to the ratio of the pane width to the pane height. You can choose to constrain the aspect ratio to a specified amount or not constrain it at all. An unconstrained axis range can be useful because it means that the axes don't have to be the same length. Anytime you manually resize a row or column, you are unconstraining the aspect ratio (refer to "Resizing Rows and Columns" on page 20-27). The aspect ratio setting only applies to views containing continuous axes on both the row and column shelves. Nominative axes are not affected by the aspect ratio settings.

### Setting Default Number Format

You can define the number of decimal places to display by default for numbers in the view. If you select **Automatic**, Tableau automatically decides the number of decimal places based on the data in the field. If you select **Manual**, you can decide to show up to 16 decimal places.

### Setting Row Attributes

Select from the following Row attributes:

- **Maximum levels of row labels:** determines the number of fields that can be added to the Rows shelf before the headers are combined on the same level.
- **Maximum levels of horizontal row labels:** determines the number of fields that can be placed on the Rows shelf before headers are automatically oriented vertically rather than horizontally.

### Setting Column Attributes

Select from the following column attributes:

- **Maximum levels of column labels:** determines the number of fields that can be placed on the **Columns** shelf before Tableau begins to combine the labels.

- **Show innermost level at bottom of view when there is a vertical axis:** displays the innermost level of column headers at the bottom of the view (as opposed to the top) when a vertical axis is added to the view.
- **Default orientation of labels at bottom of view:** determines whether labels at the bottom of the view are oriented horizontally or vertically by default. You can toggle between the horizontal and vertical options by pressing **Ctrl + L** on your keyboard.

## Editing Axes

When you add a measure to the Columns or Rows shelf, you add an axis to the view. For each axis you can specify the range, scale, tick mark properties, and more. Edit the axes to create a view that best matches your data and focuses on the relevant information. For example, you may have a view showing the Gross Profit over four years. The automatic axis may range from 0 to \$800,000 but your profits never went below \$400,000. You could adjust the Axis Range so that it starts at \$400,000 thus focusing on where the data points actually lie.



Axis formatting options are available in the Edit Axis dialog box. This section discusses the following topics:

- Changing the Axis Range
- Changing the Axis Appearance
- Formatting Tick Marks


### Changing the Axis Range

An axis shows data points that lie within a range of values. You can limit the axis range in order to focus the view to where the data points lie.

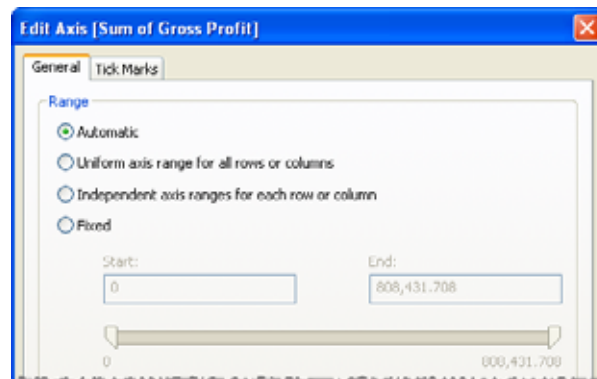
**To change the axis range:**



- 1 Right-click on the axis that you want to edit and select **Edit Axis**.



- 2 In the Edit Axis dialog box, select one of the following:
  - Automatic - The axis range is automatically decided based on the data used in the view.
  - Uniform axis range for all rows or columns - the axis range is the same across all panes in the view.
  - Independent axis ranges for each row or column - the axis ranges vary across each pane in the view.
  - Fixed - specify the start and end values for the axis. Fixed axes are applied across all panes in the view.



- 3 When finished, click **OK**.

Learn more with “Example – Changing the Axis Range” on page 20-39.

### Example – Changing the Axis Range

In this example you will build three views using the same data, however, each view will use a different axis range format. These views use the Sample-Superstore Sales spreadsheet to display the aggregated total sales for three product categories over the course of four years. The first view uses a uniform axis range for all rows in the view; the second view uses an independent axis range for each row in the view; and finally, the third view uses a custom defined fixed axis range.


#### View 1- Building a view with a uniform axis range

A uniform axis range means that the same range is applied to each row or column in your view. The range is automatically generated based on the underlying data values.

- 1 Place the **Order Date** dimension on the **Columns** shelf and the **Product Category 1** dimension on the **Rows** shelf.

## 2 Place the **Sales Total** measure on the **Rows** shelf.

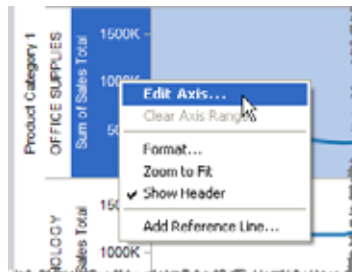
The measure is automatically aggregated as a summation and an axis is added to the view. By default the view uses a uniform axis range. Notice that the axis range is the same, from 0K to 1500K, for each product category.



## View 2-Building a View using independent axis ranges

When you use an independent axis range, each row or column will have its own axis range based on the underlying data values.

- 1 Right-click the SUM(Total Sales) axis in the view and select **Edit Axis**.




- 2 In the Edit Axis dialog box, select **Independent axis range for each row or column**.



- 3 When finished, click **OK**.

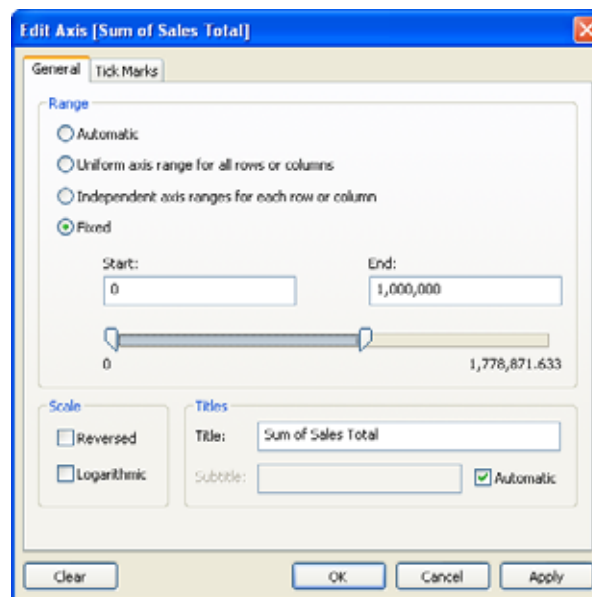
The axis range for each product category are now independent from each other. The Technology category still ranges from 0K to 1500K but the Office Supplies category ranges from 0K to a little over 500K and the Furniture category ranges from 0K to 600K.




### View 3- Building a view using a fixed axis range

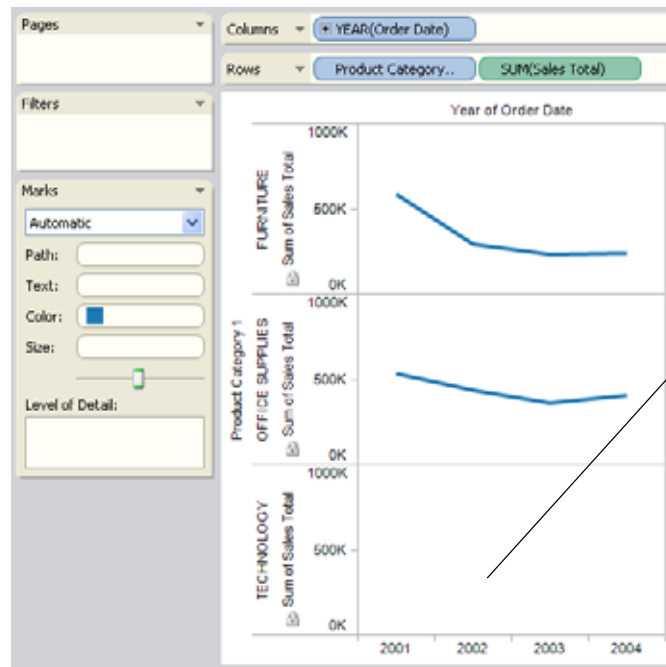
So far, Tableau has automatically generated the axis range based on the underlying data values. You can also define your own range that is applied uniformly across the rows or columns in the view.

- 1 Right-click on the SUM(Total Sales) axis in the view and select **Edit Axis**.
- 2 In the Edit Axis dialog box, select **Fixed**. Then define **Start** and **End** values either by typing into the text boxes or by dragging the sliders toward each other. For this example, type 0 as the **Start** and 1000000 as the **End**.




- 3 When finished, click **OK**.

Notice that the Technology category doesn't have any transactions with Sales below 1000K so nothing displays in its pane. The axes are marked with a lock symbol  indicating that you have limited the axis range and that some data may not be showing.




All three views are shown below.


View 1: Uniform range



View 2: Independent range



View 3: Fixed range




## Changing the Axis Appearance


Every axis has a title that is automatically generated based on the fields in the view. You can specify a custom axis title and add a subtitle using the Edit Axis dialog box. In addition, you can specify the scale of the axis such as whether to use a logarithmic scale and whether to reverse the axis.

### To change the axis title:

- 1 Right-click on the axis that you want to edit and select **Edit Axis**.
- 2 In the Edit Axis dialog box, type a new title into the text box.



- 3 Deselect the **Automatic** check box to add a custom subtitle.



- 4 When finished, click **OK**.

### To change the axis scale:

- 1 Right-click the axis that you want to edit and select **Edit Axis**.
- 2 In the bottom left of the Edit Axis dialog box, optionally select one of the following options:
  - Reversed- select this option to reverse the order of values on the axis.



- Logarithmic - select this option to use a logarithmic scale on the axis. Refer to “Log Axes” on page 1919-1 to learn more.



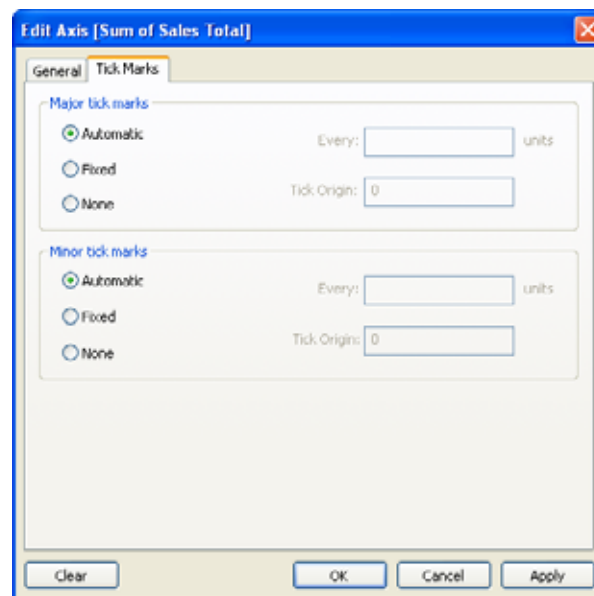
- 3 When finished, click **OK**.

## Formatting Tick Marks

You can specify how often the tick marks are displayed along the axis. Tableau allows you to modify both the Major and Minor tick marks. Major tick marks are accompanied by unit labels while Minor tick marks simply represent smaller increments between the major marks. You can choose to use automatic or fixed tick marks or have none at all.

### To format tick marks:

- 1 Right-click the axis you want to edit and select Edit Axis.
- 2 In the Edit Axis dialog box, select the Tick Marks tab.
- 3 For both Major and Minor tick marks, select from one of the following options:
  - Automatic - select this option to automatically show tick marks based on the data in the view.
  - Fixed - select this option to specify how often the tick mark should display and the starting value.
  - None - select this option to hide the tick marks completely.



- 4 When finished, click **OK**.

## Formatting Mark Colors

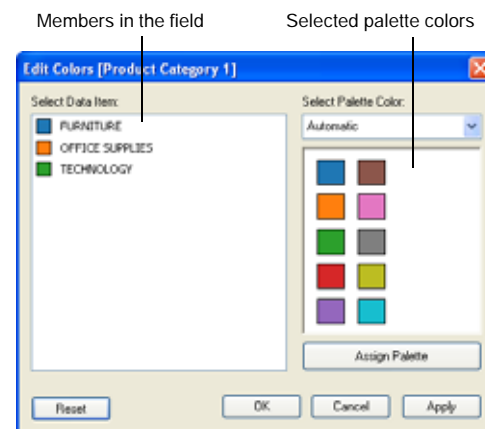
You can add colors to the view in several different ways. One useful way to use color is to add a field to the **Color** shelf, which color encodes the view with the data contained in the field. Placing a dimension on the **Color** shelf separates the marks according to the dimension members and assigns a unique color to each member. This is called categorical color encoding. Placing a measure on the **Color** shelf creates a continuous range of colors. This is called quantitative color encoding.

This section discusses the following topics:

- Categorical Colors
- Quantitative Colors

### Categorical Colors

When you add a dimension to the **Color** shelf a categorical legend is added based on the members in the dimension field. You can modify the colors used in the legend by right-clicking on the legend and selecting **Edit colors** or by double-clicking on the legend. The **Edit Colors** dialog box for a categorical legend is shown below.

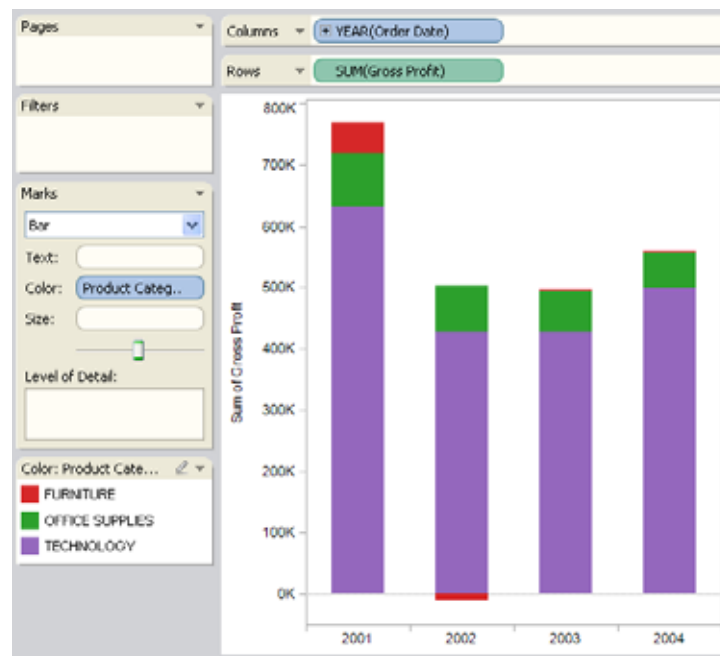


To change the color of a member, select the member on the left and then select the new color in the palette on the right. When finished, click **OK** to close the format dialog box.

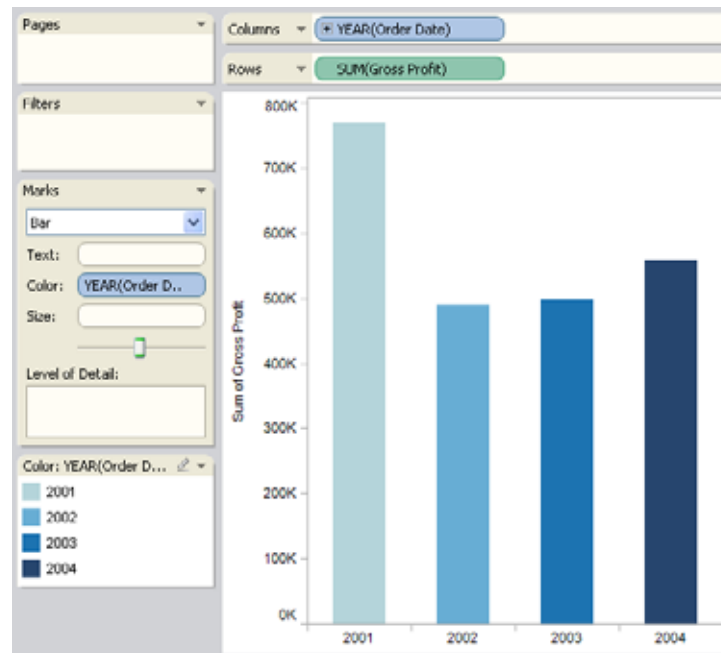
You can select a different color palette from the drop down list in the upper right of the **Edit Color** dialog box. Select from either categorical palettes or ordinal palettes. A categorical

palette, such as **Tableau 20** contains several distinct colors that can be assigned to dimension members that have no inherent order. Ordinal palettes contain a spectrum of related colors, which can be used for dimension members that have an associated order such as dates and numbers. The views below show a categorical palette versus an ordinal palette.

Categorical color palette: Product Category 1 has no natural order.



Ordinal color palette: Order Date has a chronological order.

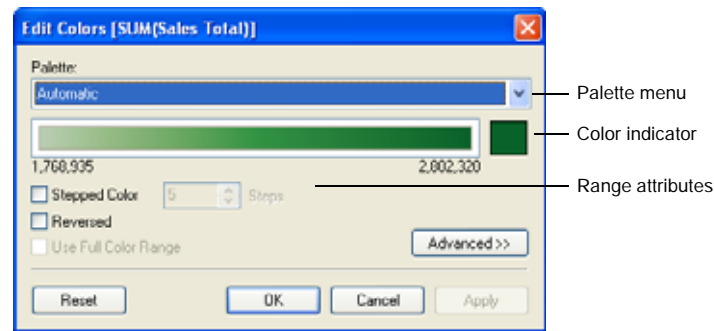


Once you select a palette, click **Assign Palette** to automatically assign the new palette colors to the members in the field. When finished, click **OK** to view the changes and close the dialog.

To return to the automatic color settings that Tableau provides by default click **Reset** in the Edit Colors dialog box and then click **OK**.

## Quantitative Colors

When you add a measure to the **Color** shelf a quantitative legend is added creating a continuous range of colors. You can modify the colors used in the range, the distribution of color, and other range attributes in the Edit Color dialog box. Right-click the legend and select **Edit Colors** or double-click on the legend. The **Edit Colors** dialog box for a quantitative legend is shown below.




To change the color used in the range, simply click on the color indicator to the right of the range and click on a new color in the spectrum. You can select a new palette from the **Palette** menu. You can choose between a sequential palette and a diverging palette. A sequential palette shows a simple range of values using color intensity to indicate one end of the range from the other. A diverging palette shows two ranges of values using color intensity to show the magnitude of the number and the actual color to show which range the number is from. Diverging palettes are most commonly used to show the difference between positive and negative numbers. When finished, click **Apply**.

The following topics discuss the options for formatting quantitative colors:

- Using Stepped Color
- Reversing the Color Palette
- Using the Full Color Range
- Limiting the Color Range
- Resetting the Color Range

### Using Stepped Color

You can also modify how the colors are distributed by selecting **Stepped Color**. The stepped color option groups the values into uniform bins each given a unique color. Use the text box to specify how many bins you want to use. For example, if you had a range of values from 0 to 100 and you select 5 steps, the color range would be broken up every 20 units. That means that all points between 0 and 20 would be colored the same, all points between 21 and 40 would be colored the same and so on. The dialog box below shows the color range broken up into five steps. When finished, click **Apply**.



If a diverging color palette is selected, the center point is shown on the color ramp with a small black mark. When the number of steps is odd, the center mark is placed in the middle of the center step. When the number of steps is even, the center mark is placed at the boundary of the center-most two steps.

### Reversing the Color Palette

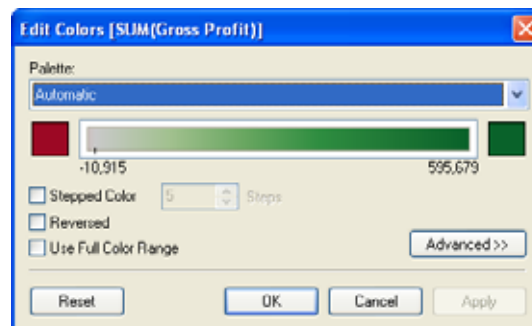
Select **Reversed** to switch the order of colors in the range. For example, if you want lower values to have a darker intensity in a sequential palette, reverse the palette. Alternatively, if you are using a diverging color palette with red representing -100 to 0 and blue representing 0 to 100, you can switch the colors using the reverse option to make blue represent the negative range and red represent the positive range. When finished, click **Apply**.

### Using the Full Color Range

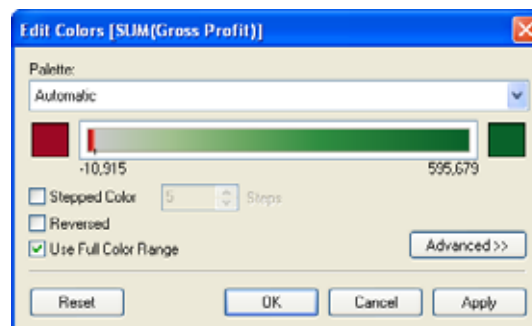
When you are using a diverging color palette you can select to **Use Full Color Range**. When you select this option, Tableau assigns the starting number a full intensity and the ending number a full intensity. If the range is from -10 to 100, the color representing negative numbers changes in shade much more quickly than the color representing positive numbers. If you do not select **Use Full Color Range**, Tableau assigns the color intensity as if the range was from -100 to 100 so that the change in shade is the same on both sides of zero. The example below shows a diverging color palette for values from -10 to 150. Without using the full color range, -10 is represented by a light red color. When the full color range is used, -10 is represented by a full red. When finished, click **Apply**.



Diverging color palette without full color range.



Diverging color palette using the full color range.



## Limiting the Color Range

You can limit the range that the colors are distributed across using the **Advanced** options. When you click **Advanced** in the Edit Colors dialog box, you can select to specify the start, end, and center values on the range by selecting the check box and typing a new value into the textbox. The **Start** value is the lower limit in the range, the **End** value is the upper limit, and the **Center** value is the where the neutral color is located on a diverging color palette. When finished, click **Apply**.

## Resetting the Color Range

To return to the automatic color settings that Tableau provides by default click **Reset** in the Edit Colors dialog box and then click **OK**.



# Annotations and Data Labels

<b>Overview</b>	<b>21-2</b>
<b>Annotations</b>	<b>21-3</b>
Adding Annotations	21-3
Rearranging Annotations	21-5
Formatting Annotations	21-11
Removing Annotations	21-15
<b>Data Labels</b>	<b>21-17</b>
Showing and Hiding Data Labels for the Worksheet	21-18
Showing and Hiding Individual Data Labels	21-18
Moving Data Labels	21-19
Formatting Data Labels	21-20
Editing Data Labels with Aliases	21-20
Example - Data Labels	21-21

---

## Overview

Annotations call attention to specific marks, points, or areas in a view. An annotation, sometimes called a call-out, is most commonly displayed as a text box with a line pointing to a specific point or mark. You can also add an area annotation, which calls out several marks or a region of the view. Additionally, you can use data labels to call out marks of interest or more commonly to label the view to make it more understandable. You can show data labels for all the marks in the view, or selectively show and hide individual labels. This section discusses the following topics:

- Annotations
- Data Labels

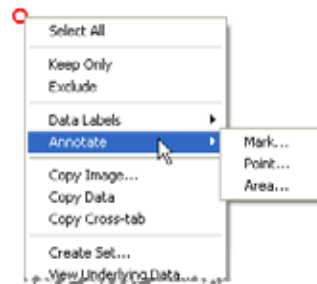
In Tableau there are three kinds of annotations: mark, point, and area. After you add an annotation, you can edit, re-position, format, and remove it. This section discusses the following topics:

- Adding Annotations
- Rearranging Annotations
- Formatting Annotations
- Removing Annotations

Annotations are an important part of publishing and sharing a view. Use annotations to call out a specific mark, a specific point such as a value on the axis or a reference line, or an area such as a cluster of scatter marks.

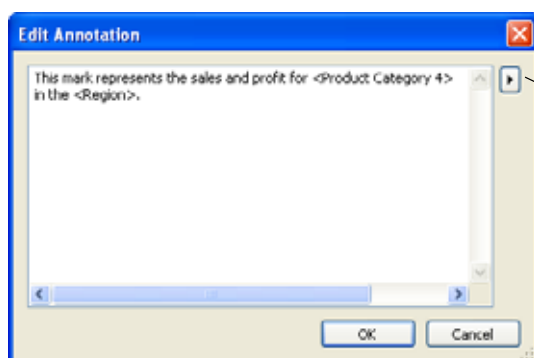
**To add an annotation:**

- 1 Right-click the view where you want to add an annotation and select **Annotate**.




- 2 On the sub-menu select one of the following types of annotations:
  - **Mark** - select this option to add an annotation that is associated with the selected mark. This option is only available if a mark is selected.
  - **Point** - select this option to annotate a specific point in the view.

- **Area** - select this option to annotate an area in the view such as a cluster of outliers or a targeted region of the view.
- 3 In the Edit Annotation dialog box, type the text you want to show in the annotation. Use the arrow button to insert dynamic parameters into the annotation text. For example, the annotation can display data values that update as the underlying data changes.



Use the arrow button to add data values as dynamic parameters in the annotation text.

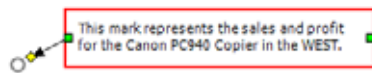
- Insert the \* parameter using the arrow button to show the data from the tooltip. As you add more detail to the view, this text is updated to show the live tooltip.
- 4 When finished, click **OK**.



After you add an annotation, you can move it around, resize it, adjust the line, and move the text. Each type of annotation can be rearranged and modified in different ways. This section discusses how to rearrange, resize, and modify each type of annotation.


### Mark Annotations

When you select a mark annotation the body and line are selected and several resize handles display. Use these handles to resize the body and line.



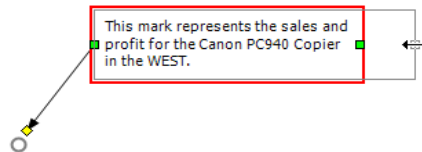
**To reposition the body:**

- Click and drag the body of the selected annotation to a new position.



#### To resize the body:

- Click and drag the body resize handle ■ left and right. The text and height are automatically adjusted to fit the width of the body.



#### To resize the line:


- Click and drag the line resize handle ◆.



### Point Annotations


A point annotation marks a specific point in the view such as a reference line or a value on an axis. Point annotations display as text with a line. When you select a point annotation, several resize handles display. Use these handles to reposition and resize the body and line.






### To reposition the body:

- Click and drag the body of the selected annotation to a new position. As you move the body, the line is automatically resized so that it continues to point at the specific point you selected.




**To resize the body:**


- Click and drag the side resize handles ■ left and right. The text and height are automatically adjusted to fit the width of the body.

**To move the line end point:**

- Click and drag the end point of the line ■ so that it points at a new location.


**Area Annotations**

An area annotation is a way to highlight or call out an area in the view. Area annotations are not associated with any particular mark, in fact, these annotations are commonly used to call out several marks. When you select an area annotation, several resize handles and two text handles display. Use these handles to reposition and resize the box and text.




**To reposition the box:**

- Click and drag the box of the selected annotation to a new position.




**To resize the box:**

- Click and drag one of the box resize handles ■.





**To reposition the text:**

- Click and drag the center text handle ◆ to a new position.



**To resize the text width:**


- Click and drag the right text handle  left and right. The text height is automatically adjusted to fit the width.




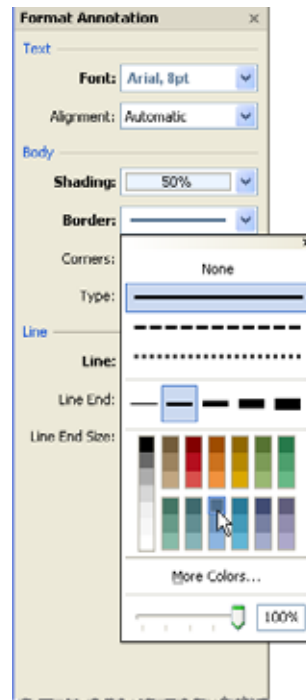
For each annotation you can modify the text, body, and line. For example, you can specify whether the body should be a box, a single edge, or not shown at all. Additionally, you can specify whether the lines on mark and point annotations end with an arrow, dot, or a simple line.

**To format annotations:**


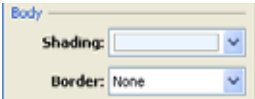

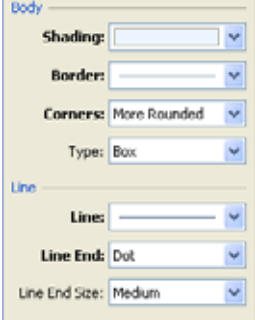

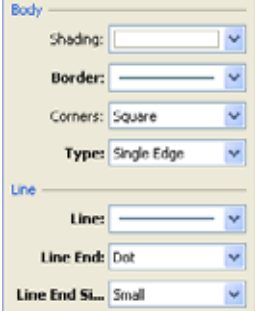
- 1 Select one or more annotations, right-click one of the selected annotations, and select **Format**. The Format window opens showing the relevant settings.





- 2 In the Format window, use the drop-downs to specify font properties, text alignment, line style, and shading.



Example Formatting Options

Annotation Style	Format Window Settings
	Default format settings for point and mark annotations.
	
	
	



Annotation Style	Format Window Settings
	
	
	

To learn more about the Format window refer to “Formatting” on page 20-1.

At any time you can remove one or more annotations.

**To remove an annotation:**

- 1 Select one or more annotations to remove.
- 2 Right-click one of the selected annotations and select **Remove** or click the **Delete** key on your keyboard.

Data labels are values shown next to each data point in a view. For example, in a view that shows product category sales over time as a line, you can turn on data labels so the sales values display next to each point along the lines. However, data labels don't have to be measure values. In the same example, you could turn on data labels that display the name of each product category next to each line.



You can show and hide data labels for the entire worksheet or selectively turn them on for specific marks. The data label values are dependent on the fields used in the view and are controlled specifically by the field on the Text shelf. This section discusses the following topics:

- Showing and Hiding Data Labels for the Worksheet
- Showing and Hiding Individual Data Labels
- Moving Data Labels
- Formatting Data Labels
- Editing Data Labels with Aliases
- Example - Data Labels

You can show and hide data labels for the whole worksheet by selecting an option on the **Format > Data Labels** menu. You can select from the following options:

- **On** - many data labels are shown but some are selectively hidden to avoid overlapping labels. This is useful when labeling dense scatter plots.
- **On (All)** - all data labels are shown in the view.
- **Off** - all data labels are hidden.

These selections override the settings for individual labels. For example, if you specifically turned on several data labels in the view and then selected **Data Labels > Off** on the **Format** menu, all data labels will be hidden.


When you turn on data labels, Tableau automatically selects a field to base the labels on. A label is then drawn next to each mark on the worksheet. However, in a dense view with many marks data labels that would otherwise overlap other labels are hidden. You can always turn on the hidden data labels by selecting **Data Labels > On (All)** or by turning on the label for specific marks, refer to “Showing and Hiding Individual Data Labels” on page 21-18 to learn how.

In the case where there’s a field on the text shelf, the data labels are based on that field. If you want to change the values displayed for the data labels, place another field on the **Text** shelf. This is one way to view additional information about your data without changing the basic layout. Tableau attempts to place data labels in the optimal locations. However, depending on your particular data view, the labels might overlap. Refer to “Showing and Hiding Individual Data Labels” on page 21-18 to learn how to selectively show and hide a specific data label.

Rather than showing all data labels, especially in a dense view, you may want to show labels for a selection of individual marks. You can use data labels to call out the values of specific marks of interest as well as hide overlapping data labels. Depending on the view, you may want to first turn on all labels and then hide specific labels or you may want to start with all data labels hidden and selectively show individual labels. You can show and hide individual data labels using the right-click context menus in the view.

**To show or hide individual data labels:**


- 1 Right-click the mark you want to show or hide a data label for.
- 2 On the context menu select **Data Labels** and then one of the following:
  - **Automatic** - select this option to allow Tableau to turn the label on and off depending on the view and the selection on the Format > Data Labels menu.
  - **On** - select this option to show the data label even when it would otherwise be hidden (based on the selection on the Format > Data Labels menu).
  - **Off** - select this option to hide the data label even when it would otherwise be shown (based on the selection on the Format > Data Labels menu).



After you show a data label in a view, you can reposition it to best fit your view and presentation. For example, in a stacked bar chart, the data labels are automatically placed in the center of each bar. However, you may want to stagger the labels so that the longer ones don't overlap.

**To move a data label:**

- 1 Select the mark whose data label you want to move.
- 2 Click and drag the yellow move handle to a new location.



Data labels are included as part of the pane format settings. You can change the font, color, alignment, and number format of all data labels using the Format window. Refer to “Formatting” on page 20-1 to learn more about using the Format window.

---

**Note** You can change the number format of the data labels by formatting the specific field that the labels are based on. Refer to “Numbers” on page 20-15 to learn how to format specific fields.


---

Another way to modify data labels is to edit the aliases of a field. An *alias* is an alternative name assigned to a dimension member, or to a field name. Tableau gives you the ability to display and edit aliases for data sources that support this feature. When you edit the aliases you can change the names of the members in a field, thus modifying the data labels displayed in a the view. For more information about Aliases and how to edit them refer to “Member Aliases” on page 8-11.

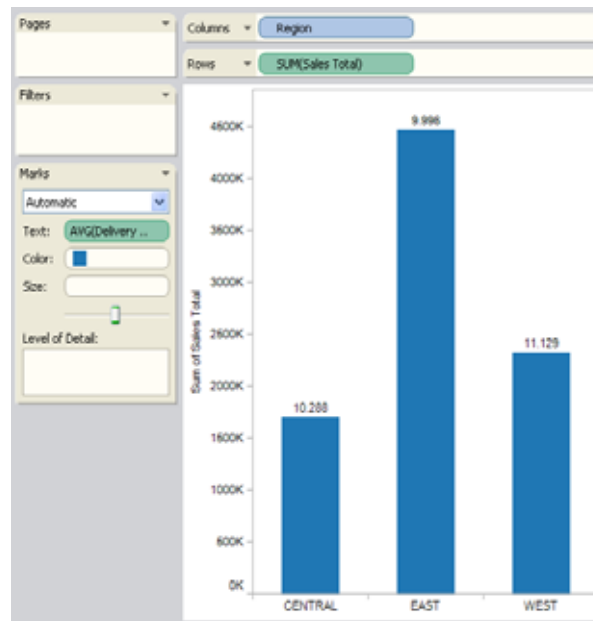
The data view shown below has data labels set to **On**.



By default, Tableau displays the measure values as the label.

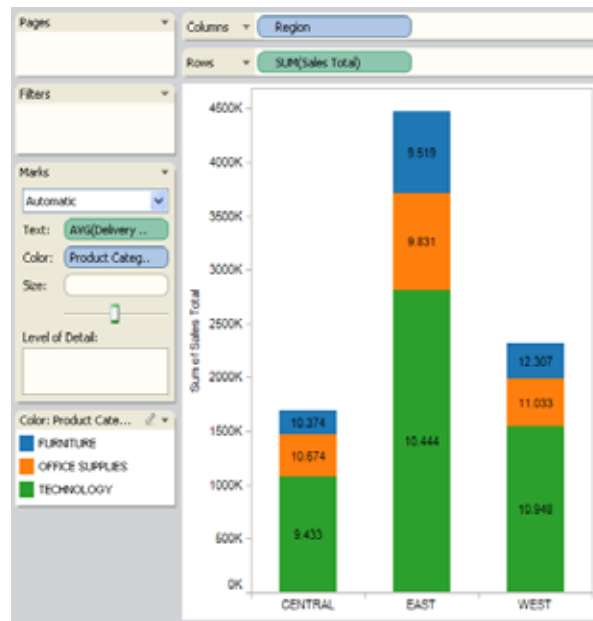


If you place a field on the Text shelf, Tableau replaces the default label with the values of that field. In the view shown below, the average delivery time is displayed as the data label for each mark. Adding the delivery time field adds additional information to the view: the labels show one value (the average delivery time), while the bar size shows another value (the sum of the sales total).




If you place a dimension on the Color shelf, you increase the level of detail and the number of data labels that are displayed.






Adding the color made the data labels more difficult to see. To make the labels more readable you can format the pane text so that the data labels are a different color (refer to “Fonts” on page 20-5) or try changing the mark type to a line by selecting **Line** on the Mark drop-down list.

If you place a dimension on the Text shelf, the data labels are defined by the dimension member names.




Notice in the view above that the Office Supplies data label is overlapping the Furniture mark. Click and drag the label to a more appropriate location using the yellow move handle.




The view below shows the profit versus sales for several products by region. By default, the data labels are hidden because there isn't a field on the text shelf. You can turn on data labels by selecting **Format > Data Labels > On (All)**. However, in a dense view such as this one, you probably want to select **On** instead so that labels don't overlap. All three views for each of these settings are shown below.

Data Labels &gt; Off



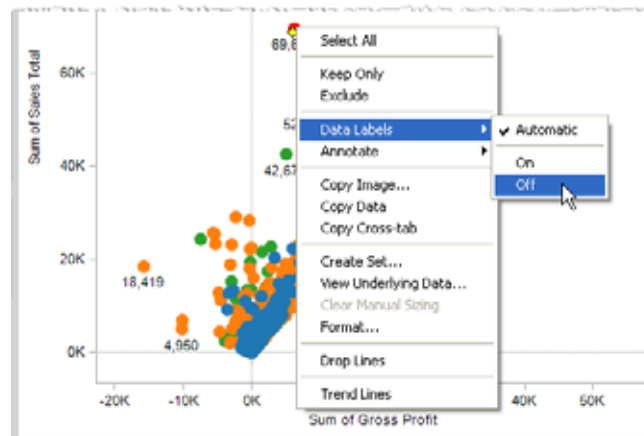
Data Labels &gt; On All



Data Labels &gt; On



Even still, you may not like how Tableau has selected to show or hide labels. For emphasis on some specific data points, you can right-click on the marks you want to show or hide labels for and select **Data Labels > On** or **Data Labels > Off**.



---

<b>Overview</b>	<b>22-2</b>
<b>Creating Dashboards</b>	<b>22-3</b>
How to Create a Dashboard	22-3
Adding Views to a Dashboard	22-3
Removing Views from a Dashboard	22-5
<b>Organizing Dashboards</b>	<b>22-8</b>
Showing and Hiding Parts of the Worksheet	22-8
Rearranging Dashboard Views	22-9
Setting the Dashboard Size	22-11
<b>Understanding Dashboards and Worksheets</b>	<b>22-13</b>
<b>Tips for Creating Useful Dashboards</b>	<b>22-15</b>

---

A dashboard is a collection of several worksheets shown in a single place, so you can compare and monitor a variety of data simultaneously. For example, you may have a set of views that you review every day. Rather than flipping through each worksheet, you can create a dashboard that displays all the views at once.

Similar to worksheets, dashboards are shown as tabs at the bottom of the workbook and update with the most recent data from the data source. When you create a dashboard, you can add views from any worksheet in the workbook. From the dashboard, you can format, annotate, drill-down, edit axes, and more.

Each view you add to the dashboard is connected to its corresponding worksheet. That means when you modify the worksheet, the dashboard is updated and when you modify the view in the dashboard, the worksheet is updated.

This section discusses the following topics:

- Creating Dashboards
- Organizing Dashboards
- Understanding Dashboards and Worksheets
- Tips for Creating Useful Dashboards


Create a dashboard just like any other worksheet, then add and remove views. This section discusses the following topics:

- How to Create a Dashboard
- Adding Views to a Dashboard
- Removing Views from a Dashboard

You can create a dashboard in much the same way you create a new worksheet. After you create a dashboard you can add and remove views.

**To create a dashboard:**

- Select **Edit > New Dashboard**.



A new tab for the dashboard is added along the bottom of the workbook. Switch to the new dashboard to add views.

---

**Note** You can also create a new dashboard by clicking the arrow next to the New Sheet button on the toolbar and selecting **New Dashboard**.

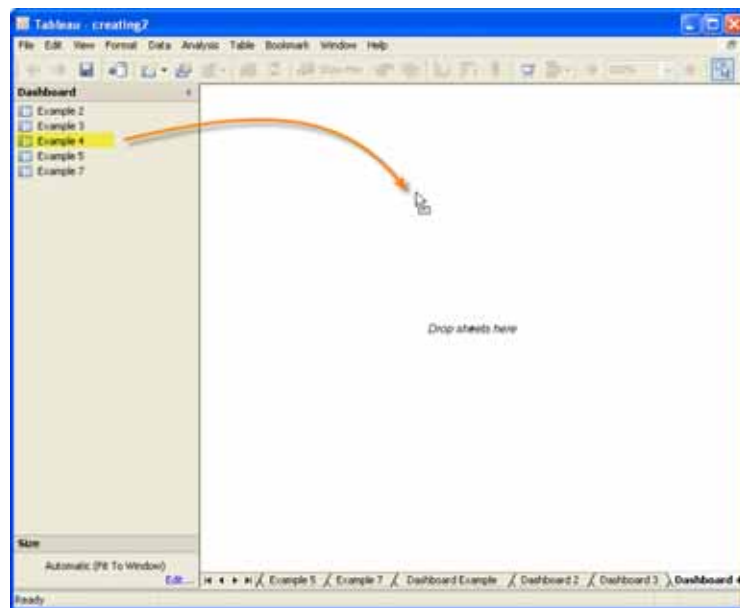
---

When you open a dashboard the Dashboard window replaces the Data window on the left side of the workbook. The Dashboard window lists the worksheets that are currently in the

workbook. As you create new worksheets, the Dashboard window updates so you always have all worksheets available when adding to a dashboard.

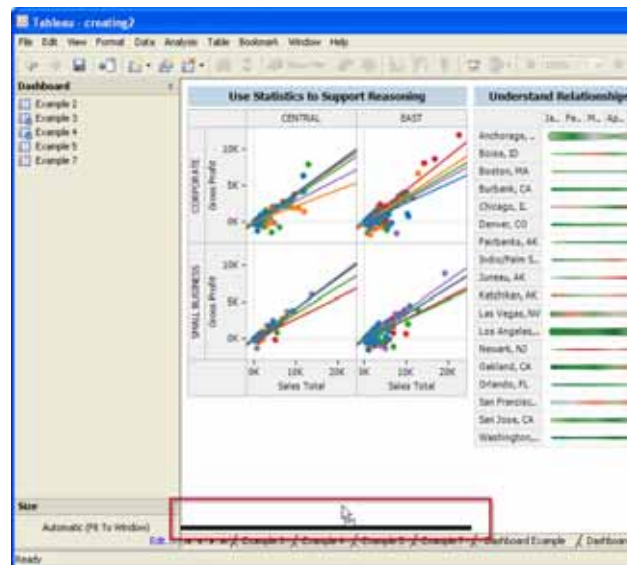
**To add a view to a dashboard:**

- Click and drag a worksheet from the Dashboard window to the dashboard on the right.



Continue to drag as many of the worksheets to the dashboard as you like. Notice as you drag worksheets around over the dashboard a black bar displays, indicating the various places you can drop it.



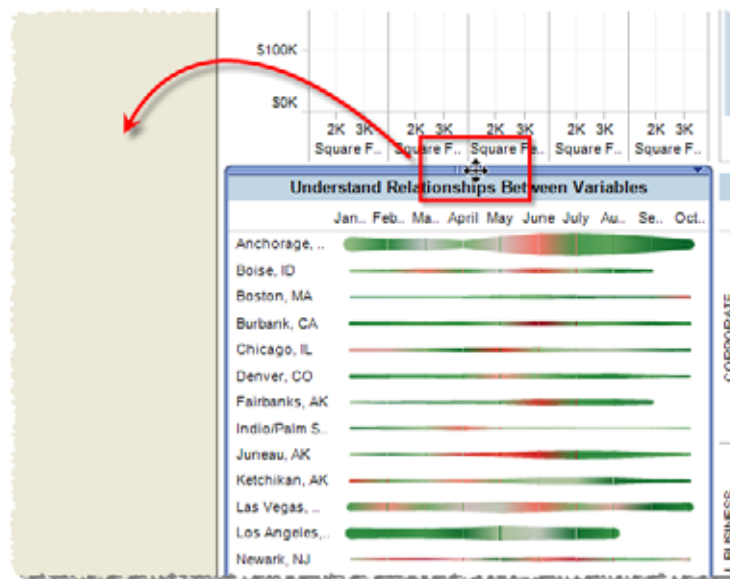


After a view is added to the dashboard, the worksheet is marked with a check mark in the Dashboard window.

After you add a worksheet to a dashboard, you can remove it in a number of different ways including dragging it out of the dashboard, using the context menus in the Dashboard window, or using the dashboard view menu.


**To remove a view by dragging:**

- 1 Select the view you want to remove from the view.
- 2 Click the move handle at the top of the view and drag it off the dashboard.



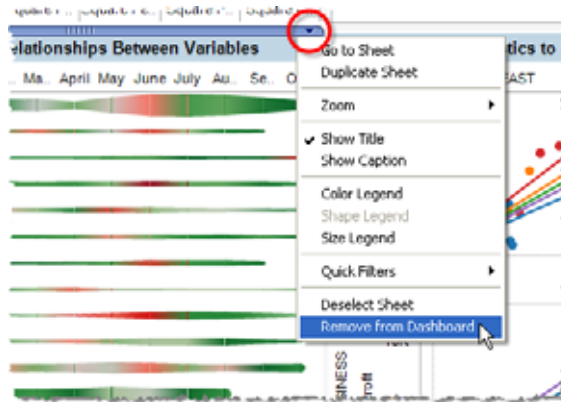
**To remove a worksheet using the Dashboard window:**

- Right-click the worksheet in the Workbook pane and select **Remove from Dashboard**.



**To remove a worksheet using the dashboard view menu:**

- 1 Select the view in the dashboard that you want to remove.
- 2 Select **Remove from Dashboard** on the dashboard view menu.



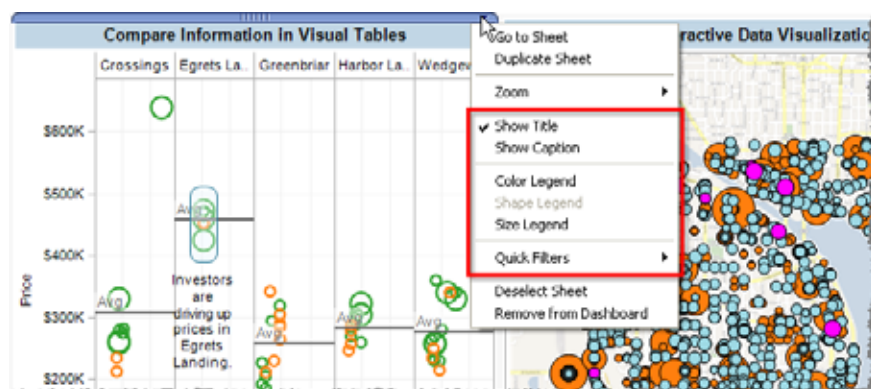
A dashboard can be made up of several views, legends, and quick filters. Each of these can be rearranged and hidden to help you highlight the most important information for your analysis. This section discusses the following topics:

- Showing and Hiding Parts of the Worksheet
- Rearranging Dashboard Views

As you drag worksheets to the dashboard, the view from the worksheet is automatically displayed. However, you may want to show other parts of the worksheet such as legends, titles, captions, and quick filters. You can show these parts of the worksheet using the dashboard view menu.

**To show and hide parts of the worksheet:**

- 1 Select a view in the dashboard.
- 2 On the dashboard view menu, you can select to show items that are available for the selected view. For example, you can show the title, caption, legends, and a variety of quick filters.



---


**Note** Quick filters are only available for the fields used in the original view.

---

Rearrange the views, legends, and quick filters in a dashboard in a way that best fits your analysis or presentation. You can rearrange the parts of a dashboard using the move handle at the top of a selected view, legend, or quick filter.

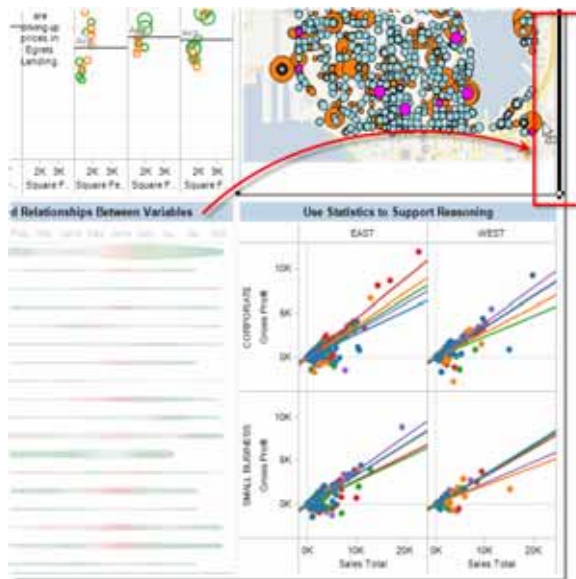
**To move a view, legend, or quick filter:**

- 1 Select the view, legend, or quick filter that you want to move.
- 2 Click and drag the move handle at the top of the selected item to a new location.




- 3 Drop the dashboard element in a new location.

As you drag the element around the dashboard area, a black bar highlights the available places you can drop it.

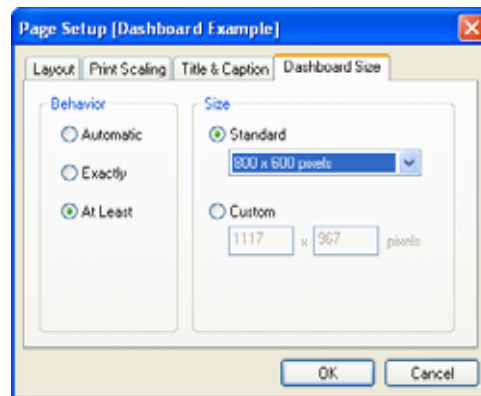


By default, the views in a dashboard automatically scale to fill the entire window. That means that the size of the dashboard can vary from machine to machine, depending on the screen resolution and the size of the application window. You can, however, specify a fixed size for the dashboard, which ensures a consistent layout when you share the dashboard with others. The size of the dashboard is displayed at the bottom of the Dashboard window.



#### To specify the size of a dashboard:

- 1 Click **Edit** in the Size area of the Dashboard window.
- 2 In the Page Setup dialog box, make sure the **Dashboard Size** tab is selected.



- 3 Select one of the following behaviors:
  - **Automatic** - the dashboard automatically resizes to fill the application window. This is the default behavior.

- **Exactly** - the dashboard always remains a fixed size. If the dashboard is larger than the window the dashboard becomes scrollable.
  - **At Least** - the dashboard automatically resizes to fill the application window, but will not scale smaller than a specified minimum size.
- 4 Select one of the following types of sizes:
- **Standard** - select from a list of common sizes, such as 800 x 600.
  - **Custom** - type a custom width and height into the text boxes. By default, these text boxes are filled with the dashboard's current width and height.
- 5 When finished, click **OK**.




The views in a dashboard are connected to the worksheets that they represent. That means when you make changes to the worksheet, the dashboard is updated and subsequently, any changes you make to the dashboard affect the worksheet. This interaction is important to remember when you are annotating, formatting, and resizing the views in your dashboard.

While dashboards are an easy way to summarize and monitor at a glance, you can go back edit the original view by jumping to a selected worksheet. Additionally, you can duplicate worksheets directly from the dashboard to perform in-depth analysis without affecting the dashboard. Finally, you can hide worksheets that are used in dashboards so that they are not shown as tabs along the bottom of the workbook.

**To jump to a worksheet from a dashboard view:**


- 1 Select the view you want to see full size.
- 2 Select **Go to Sheet** on the dashboard view menu.



**To duplicate a worksheet from a dashboard view:**


- 1 Select the view you want to duplicate.

- 2 Select **Duplicate Sheet** on the dashboard view menu



**To hide a worksheet:**

- Right-click the worksheet tab that you want to hide and select **Hide**.

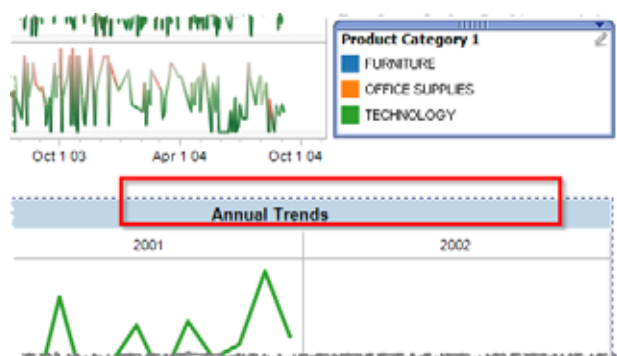


A well designed dashboard can be useful for analyzing, monitoring, reporting, and presenting your data. However, a poorly laid out dashboard can actually detract from the data you are trying to present. This section provides tips for how to create dashboards that don't waste space, facilitate your analysis, resize well, and are visually pleasing.

If you include a lot of quick filters and legends you may find having them distributed throughout the entire dashboard throws off the overall layout and distracts from the views. To avoid this, by default legends and quick filters are organized into a single column on the right side of the dashboard.



When you select a legend or quick filter in a dashboard, the associated view is highlighted with a dashed border.



To avoid over crowding your dashboard, limit your dashboard to include between three and five views that focus on a specific topic.

When you create a dashboard for sharing with others, consider who will be viewing the dashboard and their screen size. By default, dashboards automatically resize to fit within the application window, making it hard to ensure a consistent layout from screen to screen. You can, however, specify a fixed size to alleviate the problem. Refer to “Setting the Dashboard Size” on page 22-11 to learn more.

<b>Overview . . . . .</b>	<b>23-2</b>
<b>Publishing Workbooks . . . . .</b>	<b>23-3</b>
How to Publish Workbooks to the Server . . . . .	23-3
Specifying Permissions . . . . .	23-7
Showing and Hiding Worksheets . . . . .	23-12
Embedding Passwords . . . . .	23-13
<b>Publishing Data Sources . . . . .</b>	<b>23-14</b>
<b>Opening Workbooks from the Server . . . . .</b>	<b>23-18</b>
<b>Importing Data Sources from the Server . . . . .</b>	<b>23-21</b>
<b>User Filtering . . . . .</b>	<b>23-24</b>
How to Create a User Filter . . . . .	23-24
Previewing User Filters . . . . .	23-29
Editing User Filters . . . . .	23-30
Publishing with User Filters . . . . .	23-32
Example - Setting User Filters . . . . .	23-33
Filter Using a List of Users in your Data Source . . . . .	23-40

---

Tableau Server is where you and others can share views of your data within your company. You can publish workbooks and data sources to the server, as well as open workbooks and data sources from the server. Tableau Server is server software as defined in Tableau Software's End User Software License Agreement.

You must be added as a user on Tableau Server and be given publishing permission to publish views from Tableau Professional. Contact your server administrator to find out whether you have access to Tableau Server. You can also contact a Tableau sales representative to learn more about purchasing Tableau Server.

This section discusses how to publish to the server in the following topics:

- Publishing Workbooks
- Publishing Data Sources
- Opening Workbooks from the Server
- Importing Data Sources from the Server
- User Filtering

---

**Note** Refer to the online help accessible from Tableau Server to learn more about browsing the server as well as managing users, groups, and projects.

---


Publish views of your data to Tableau Server by publishing the workbook. Workbooks can be organized into projects and assigned tags, which are keywords that are used for search. When you publish a workbook you can add it to existing projects, hide individual sheets, add tags, specify permissions to regulate access to the workbook on the server, and choose to . This section discusses the following topics:

- How to Publish Workbooks to the Server
- Specifying Permissions
- Showing and Hiding Worksheets

After you create a workbook, you can publish to Tableau Server by following the steps below.


**To publish a workbook to the server:**

- 1 Select **Collaborate > Publish Workbook**.




- 2 Type the following into the Tableau Server Login dialog box:
  - **Server:** the server name (Example: sales\_server,
  - **User Name:** your user name. If Tableau Server is configured to use Active Directory, type your Windows user name (the domain is not required), otherwise, type your Tableau Server user name.

- **Password:** your password




- 3 In the Publish Workbook to Tableau Server dialog box, select a project to publish the workbook into.

A project is like a folder that can contain workbooks and data sources. Tableau Server comes with one project called Default. Leave the project set to **Default** to add the workbook to this pre-set project. All workbooks must be published into a project.



- 4 Type a name for the workbook into the **Name** text box.





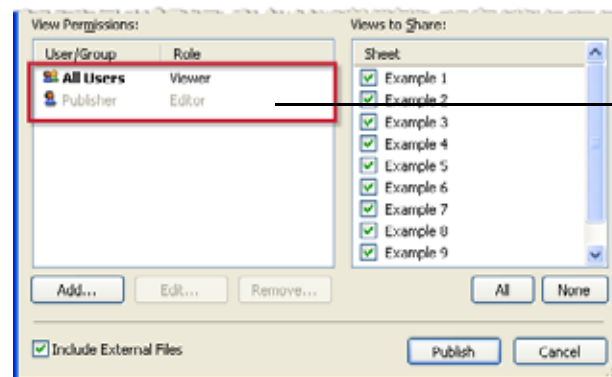
**Note** Use the drop-down list to select an existing workbook on the server. When you publish using an existing workbook name, the workbook on the server is overwritten with your workbook. You must be allowed the **Write** capability to overwrite workbooks on the server.

- 5 Optionally type one or more keywords that describe the workbook into the Tags text box. Tags help you and others find related workbooks when browsing the server.

Each tag should be separated by either a comma or a space. If the tag contains a space, type the tag surrounded by quote marks (e.g., "Sales Quotes").

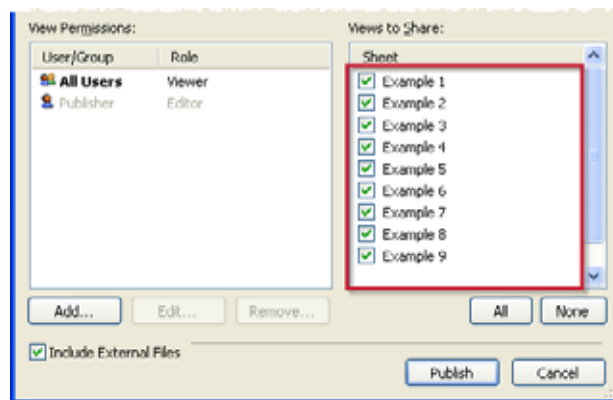
Tags: development finance examples

- 6 Optionally specify permissions to allow or deny access to the workbook on the server. By default all users can interact with the workbook and you, as the publisher, are allowed all capabilities. Refer to "Understanding Capabilities" on page 23-11 to learn more.



The default permissions allow all users to interact with the workbook and you, the publisher, are allowed all capabilities.

- 7 Select the sheets you want to share on Tableau Server. Any sheets that are not selected are hidden on the server. Refer to “Showing and Hiding Worksheets” on page 23-12 to learn more.



- 8 If the workbook contains one or more user filters you can specify what the thumbnail on the server will look like using the Generate Thumbnails as User: drop-down list. Refer to
- 9 Select whether to Include External Files. When you include external files, a copy of any referenced external file data sources or background images are published along with the workbook. External files include Excel, Access, Text, Data Extract, and image files. If you don't include these files, others may not be able to see the worksheets online.



- 10 Optionally decide whether to embed user names and passwords so server users don't have to have an account on the database to see the views. Refer to “Embedding Passwords” on page 23-13 to learn more.
- 11 When finished, click **Publish**.

---

**Note** If you are publishing a workbook that references data sources or images on a mapped drive, you should make sure to check the Include External Files option when publishing. If you do not want to publish the external files to the server, change the connection information so that the workbook references a full UNC path. For example rather than connecting to D:\datasource.xls you would connect to \\filesrv\datasource.xls.

---

When you publish a workbook, you have the option to specify permissions both for groups and specific users. Permissions allow or deny access to the workbook and its contained views on the server. By default all users are allowed to view the workbook and you, as the publisher, are allowed all capabilities. Tableau Server has three pre-defined sets of permissions called roles. Roles make it easy to assign common sets of permissions. This section discusses the following topics:


- Adding Permissions
- Editing and Removing Permissions
- Understanding Capabilities

### **Adding Permissions**

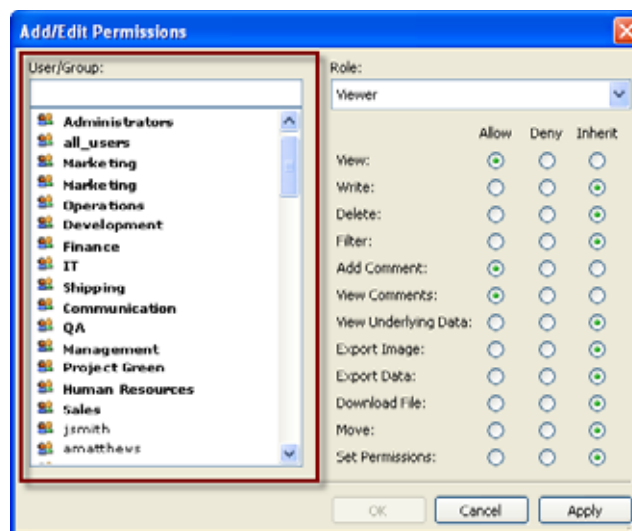
You can add permissions to a workbook in the Publish Workbook dialog box. After a workbook has been published, you can also modify and add permissions directly on Tableau Server. Refer to the Tableau Server online help to learn more.

#### **To add permissions:**

- 1 In the Publish Workbook to Tableau Server dialog box, click the Add button in the bottom left corner.




- 2 Select who you want authorize. You can select a group of users or a specific user in the User/Group list.



- 3 Select a role from the drop-down list to assign common sets of capabilities. Tableau Server has the following three pre-set roles:
  - **Viewer** - allows users to view the workbook on the server, as well as add and view comments.
  - **Interactor** - allows users to view the workbook on the server, apply filters, view the underlying data, export the image, and export the data. All other capabilities are inherited from the user's group and project permissions.

- **Editor** - allows all capabilities.



You can also specify custom sets of capabilities. Select whether to **Allow** or **Deny** each of the listed capabilities. If you leave the capability set to **Inherit**, the permission will be inherited from the user's group and project permissions. Refer to "Understanding Capabilities" on page 23-11 to learn more about each capability.

- 4 When finished, click **OK**. The new permission displays in the Publish dialog box.



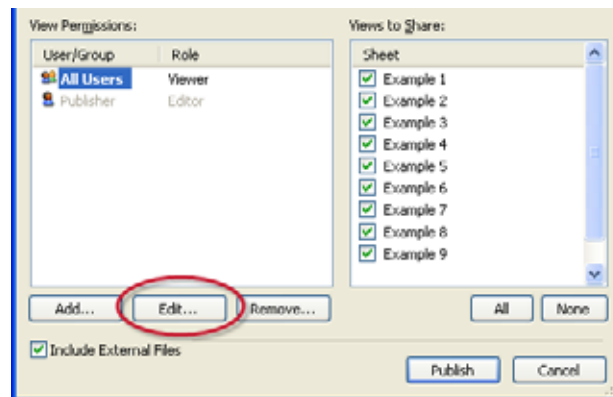
**Note** If you are adding permissions for several users and groups, click the **Apply** button when you are finished specifying permissions. The set of permissions is added and the Add/Edit Permissions dialog box remains open so you can continue to add more.

## Editing and Removing Permissions

In the Publish Workbook to Tableau Server the current permissions assigned to the workbook are shown in the bottom left. You can add, edit, and remove these permissions. Refer to "Adding Permissions" on page 23-7 to learn more about adding permissions.

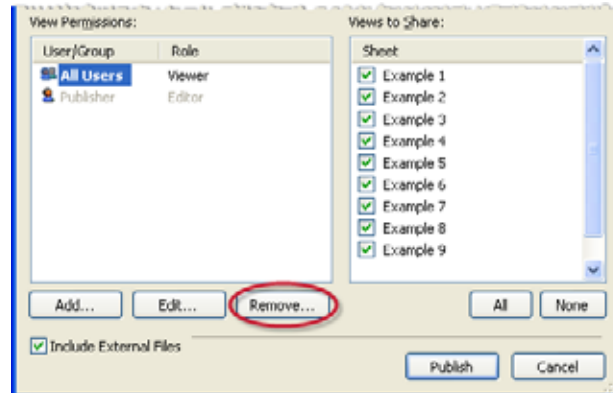
**To edit existing permissions:**

- Select the user or group whose permissions you want to modify and click **Edit**.



**To remove existing permissions:**

- Select the user or group whose permissions you want to delete and click **Remove**.



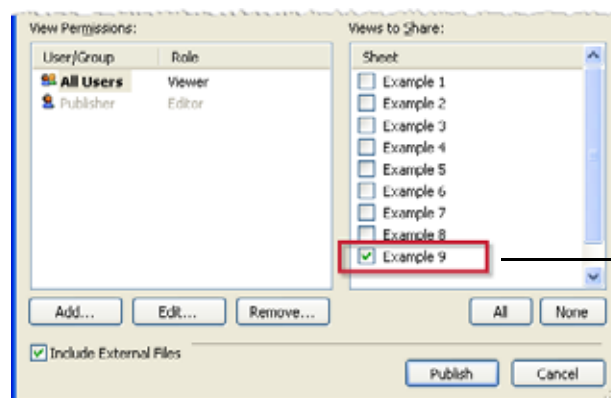
## Understanding Capabilities

A capability is an action or set of actions that can either be allowed or denied to users and groups when publishing a workbook. The following table describes each of the capabilities:

Capability	When allowed, users can...
View	View the worksheets on Tableau Server.
Write	Overwrite the workbook. If you allow someone this capability and she re-publishes the workbook, she becomes the publisher and is given all capabilities. Your access to the workbook will be controlled by your group permission and any permissions the new publisher sets.
Delete	remove the workbook and all of its contents from the server
Filter	Modify quick filters, keep only filters, and exclude
Add Comment	Add comments to the views in the workbook
View Comments	View the comments associated with the views in the workbook
View Underlying Data	View the underlying data for the worksheets in the workbook.
Export Image	Export each worksheet as an image.
Export Data	Export the data to a text file.
Download file	Open the workbook from the server using Tableau Professional
Move	Move workbooks between projects.
Set Permissions	Specify permissions for the workbook and all of the contained views.

By default, all sheets in a workbook are published to Tableau Server and can be viewed by all users. However, when you publish, you have the option to hide specific worksheets so they are not accessible on the server. Hidden sheets can still be accessed when the workbook is opened from the server using Tableau Professional. Users must be allowed the Download File capability to open the workbook from the server. Refer to “Adding Permissions” on page 23-7 to learn more about how to allow or deny this capability.

Showing and hiding worksheets is useful when you want to publish a complete dashboard without publishing the worksheets that make up the dashboard. For example, when you publish a workbook that has several worksheets and a summary dashboard, you can select to hide the individual worksheets and only show the dashboard. Only the dashboard will show on the server. Remember though, anyone allowed the Download File capability can open the workbook from the server and access the hidden worksheets.




In this example only the dashboard (Example 9) is shown. All other views will be hidden on the server.



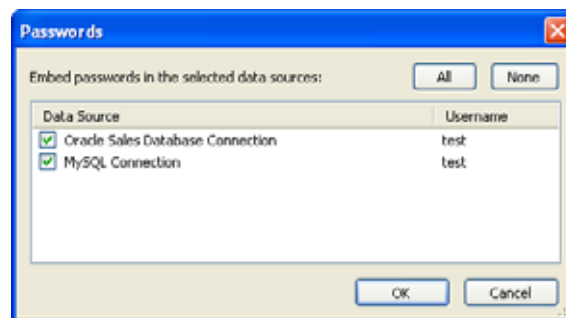
If your workbook connects to a data source that requires a user name and password, by default people viewing it on the server will also be prompted for a user name and password. However, as the author you can choose to specify a database user and password that will always be used when a view is opened on the server. The login information is not made public to the server users rather the server user will just automatically be logged in so they can see the view.

**To embed passwords in a published workbook:**

- 1 In the Publish Workbook dialog box, click the **Passwords** button in the bottom left corner.



- 2 The data source connections that require a user name and password are listed along with the user name you are currently logged in with. Select the connections that you want to embed passwords for and then click **OK**.



---

**Note** Whether you are allowed to embed passwords is controlled by the Tableau Server administrator. Administrators can allow authors to embed passwords with the Settings on the Maintenance page of the server.

---

In addition to publishing workbooks, you can also publish a data source. Publish data sources to share them with others who have access to Tableau Server. Data sources can be imported using Tableau Professional. This section discusses how to publish a data source to Tableau Server.

**To publish a data source:**

- 1 Select **Collaborate > Publish Data Source**.




- 2 Type the following into the Tableau Server Login dialog box:
  - **Server:** the server name
  - **User Name:** your user name. If Tableau Server is configured to use Active Directory, type your Windows user name, otherwise, type your Tableau Server user name.

- **Password:** your password




- 3 In the Publish Data Source dialog box, select a project to publish the data source into.

A project is like a folder that can contain workbooks and data sources. Tableau Server comes with one project called Default. Leave the project set to **Default** to add the data source to this pre-set project. All data sources must be published into a project.



- 4 Type a name for the data source into the **Name** text box.




---

**Note** Use the drop-down list to select an existing data source on the server. When you publish using an existing data source name, the data source on the server is overwritten. You must be allowed the **Write** permission to overwrite data sources on the server.

---

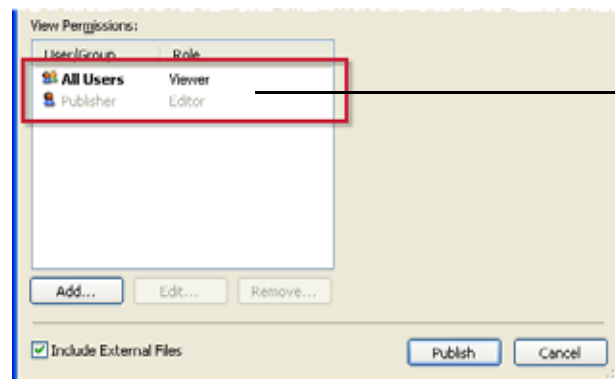
- 5 Optionally type one or more keywords that describe the data source into the Tags text box. Tags help you and others find related data sources when browsing the server.

Each tag should be separated by either a comma or a space. If the tag contains a space, type the tag surrounded by quote marks (e.g., “Sales Quotes”).



Tags: development finance examples

- 6 Optionally specify permissions to allow or deny access to the data source on the server. By default all users can view the data source and you, as the publisher, are allowed all capabilities. Refer to “Specifying Permissions” on page 23-7 to learn more.



The default permissions allow all users to import the data source and you, the publisher, are allowed all capabilities.

- 7 If your workbook contains active user filters you must specify how to generate the thumbnail images for the server. Refer to “Publishing with User Filters” on page 23-32 to learn more about how to do so.
- 8 Finally, if you are publishing an external file data source or a data source that is on a mapped drive select to Include External Files. When you include external files, a copy of the data source is published. External file data sources include Excel, Access, Text,

Data Extract, and image files. If you don't include these files, others may not be able to see the worksheets online.



- 9 When finished, click **Publish**.

If you have been allowed the Download File capability for a workbook, you can use Tableau Professional to open the workbook from the server. When you open a workbook from the server and make changes, you can either save it to your hard drive or, if you have been allowed the Write capability, you can republish the workbook to the server.

**To open a workbook from the server:**

- 1 Select **Collaborate > Open Workbook**.




- 2 Type the following into the Tableau Server Login dialog box:
  - **Server:** the server name
  - **User Name:** your user name. If Tableau Server is configured to use Active Directory, type your Windows user name, otherwise, type your Tableau Server user name.


- **Password:** your password



- 3 In the Open Workbook from Tableau Server dialog box, select the workbook you want to open and click **Open**.



You can find workbooks using the **Find** drop down lists. You can search all workbooks on the server or find by tags, publisher, project, or workbooks that you published.





If you have been allowed the Download File capability for a data source, you can use Tableau Professional to open the workbook from the server.

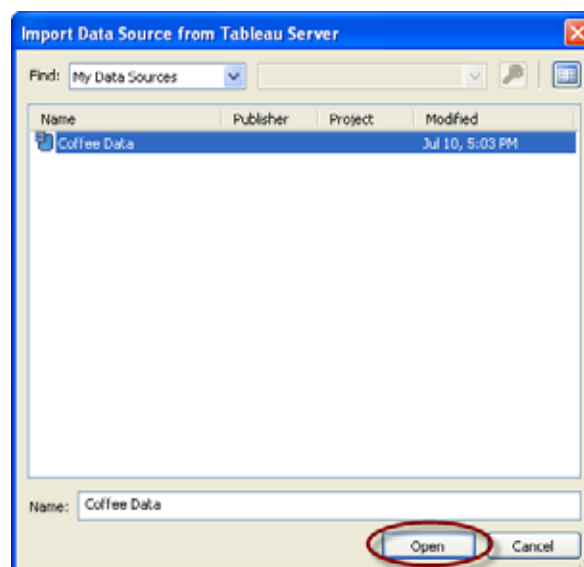
**To import a data source from the server:**

- 1 Select **Collaborate > Import Data Source**.
- 2 Type the following into the Tableau Server Login dialog box:
  - **Server:** the server name
  - **User Name:** your user name. If Tableau Server is configured to use Active Directory, type your Windows user name, otherwise, type your Tableau Server user name.


- **Password:** your password



- 3 In the Import Data Source from Tableau Server dialog box, select the data source you want to import and click **Open**.



You can find data sources using the **Find** drop down lists. You can search all data sources on the server or find by tags, publisher, project, or data sources that you published.



User filtering is a special kind of filter that allows you to limit the data any given person can see in a published view. For example, in a sales report that gets shared with regional managers, you may want to only allow the Western Regional Manager to see the western sales, the Eastern Regional Manager to see the eastern sales, and so on. Rather than create a separate view for each manger, you can define a user filter that allows each manager to see the data for a particular region.

A user filter is defined for an individual field and users or groups are given permission to see a subset of the members in that field. In the sales report example above, the user filter is defined for the Region field and each manager is given permission to see a corresponding region.

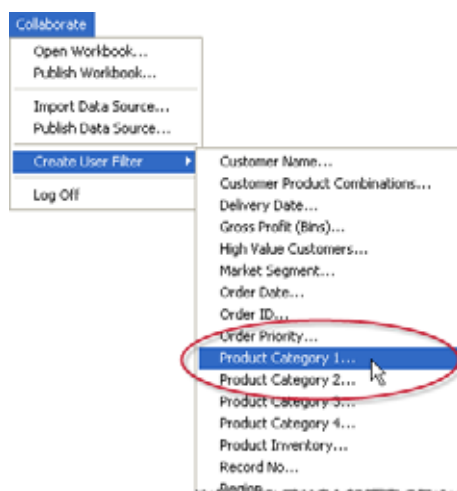
You can define a user filter for any dimension or multidimensional hierarchy. In addition you can define user filters for sets, binned fields, and ad-hoc groups that you've created. The user list comes from Tableau Server. When you publish to Tableau Server the the view is adjusted based on the who is logged in and looking at it.

This section discusses the following topics


- How to Create a User Filter
- Previewing User Filters
- Editing User Filters
- Example - Setting User Filters

You can create as many user filters as you like for a given view.


- 1 Select **Collaborate > Create User Filter** and then select a field that you want to filter on. For example, if you are limiting product data each person can access select the Product field.



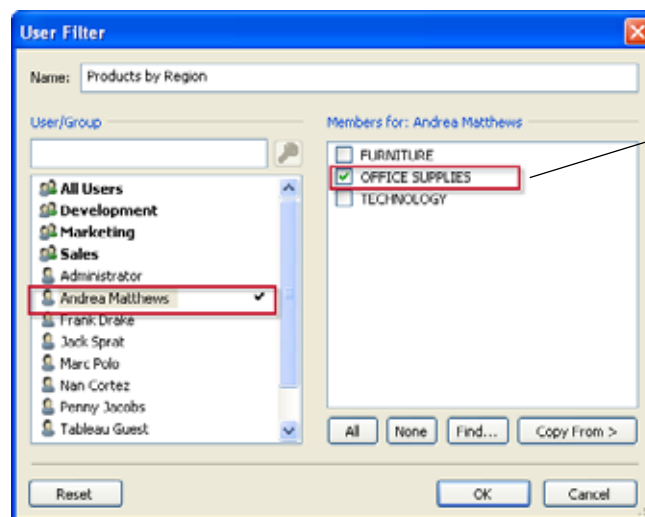
- 2 If necessary log in to Tableau Server. Refer to “How to Publish Workbooks to the Server” on page 23-3 to learn more about logging in.



- 3 In the User Filter dialog box, type a **Name** for the set of rules you are creating. For example, if you are filtering on product information you could name it Products by User.




- 4 Select a user or group in the list on the left; then on the right select the members of the field that the selected users are allowed to see. Repeat this process as necessary until everyone is assigned the correct set of members.

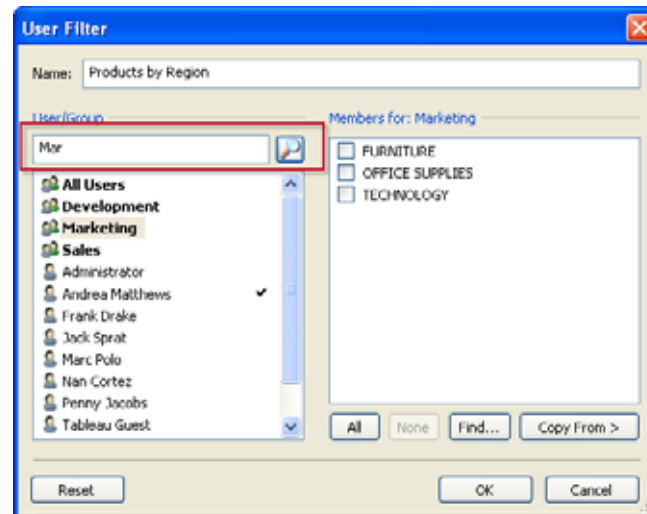


Here Andrea Matthews will only be able to see the Office Supplies data in the view.

- 5 When finished click **OK**.

### Finding Users in the User Filter Dialog Box

The User Filter dialog box makes it easy to find and select users or groups that you want to set a filter for. Simply begin typing the name of a user or group into the the text box at the top of the list of users. The first matching user or group is automatically selected. Click the **Search** button  to find the next matching user or group and so on.




### Finding Field Members in the User Filter Dialog Box

Some fields have a large number of members that are difficult to select one by one. You can find and select members easily using the Find dialog.

**To find members in the User Filter dialog box:**

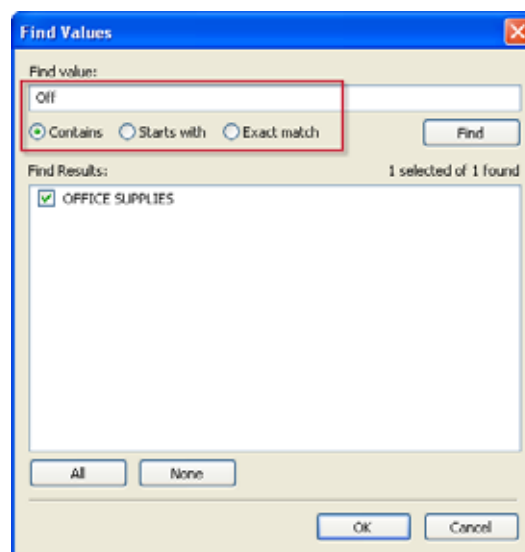
- 1 Click the **Find** button at the bottom of the list of members.



- 2 Type all or part of the member name into the **Find value** text box at the top of the Find Values dialog box and then click **Find**. You can change the search criteria by selecting

whether to return members that **Contain**, **Starts with**, or is an **Exact match** to the text you typed.

- 3 Select one or more members from the results shown in the bottom half of the dialog box. After you select the members of interest you can continue to search for other members until you have all necessary members selected.



- 4 When finished, click **OK** to return to the User Filter dialog box.

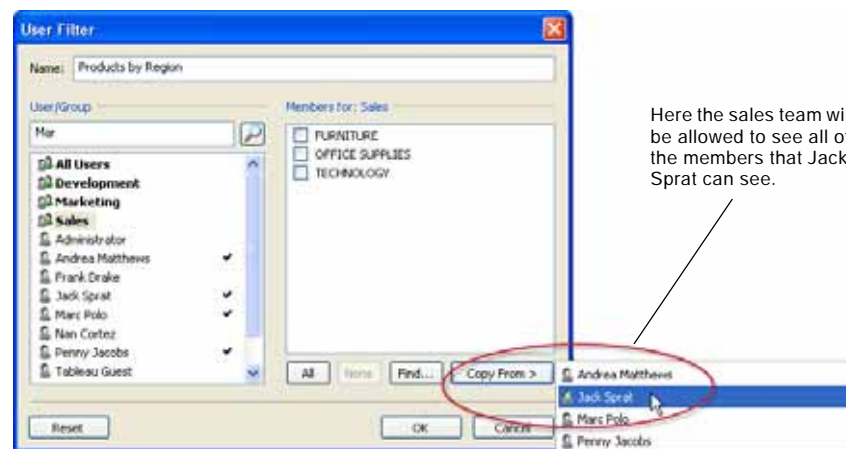
### Copying member selections from other users

As you specify the members each user or group can see you may want to just duplicate the member selections you already set for another user or group. For example, if you specify that the Product Manager can see a list of 50 products and then decide that you want to share the same products with everyone else on the her team; you can simply duplicate the member selection instead of having to select the 50 products for every member of the team.

**To copy a member selection from another user or group:**



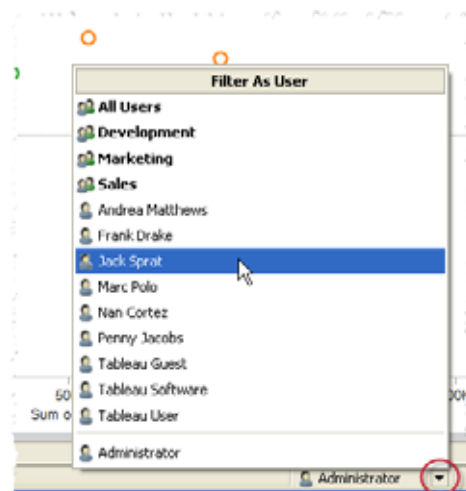
- 1 Select the user or group that you want to copy the member selection to.
- 2 Click the **Copy From** button at the bottom of the member list.
- 3 Select the user or group that you want to copy the member selection from.



If a workbook contains one or more user filters, the User Filter menu displays in the bottom right corner of the workbook window. The User Filter menu allows you to preview what each user or group will see when they look at the view on Tableau Server.

**To preview a user filter:**

- 1 Open the User Filter menu by clicking on the black arrow in the bottom right corner of the workbook window.
- 2 Select a user or group that you want to preview as.
- 3 The selected user or group is shown in blue at the bottom of the workbook window and the view updates to only show the appropriate data.




**Note** At any time you can return to viewing the workbook as yourself by selecting your user name on the User Filter menu. Your user name is listed at the bottom below a separator.

After you create a user filter you can go back and edit it just like you can edit other sets.

**To edit a user filter:**

- 1 Right-click the user filter in the Sets area of the Data window and select **Edit**.



- 2 If necessary log in to Tableau Server. Refer to “How to Publish Workbooks to the Server” on page 23-3 for details on logging into the server.
- 3 In the User Filter dialog box, make the necessary changes and click **OK**.

---


**Note** If you modify a user filter while not logged into the server, the set will only show users and groups that have a user filter specified. Log in to see all users and modify the user filter.

---

When you publish a workbook it and its contained sheets are represented with thumbnail images on the server. You can specify what these thumbnail images will look like by selecting to generate thumbnails as a specific user. For example, if you want the thumbnail image to show all three regions of a sales forecast, you can select to generate thumbnails as the manager who is allowed to see all regions.

**To specify how to generate the thumbnails:**

- In the publish dialog box, select a user or group in the **Generate Thumbnails as User** drop-down list.



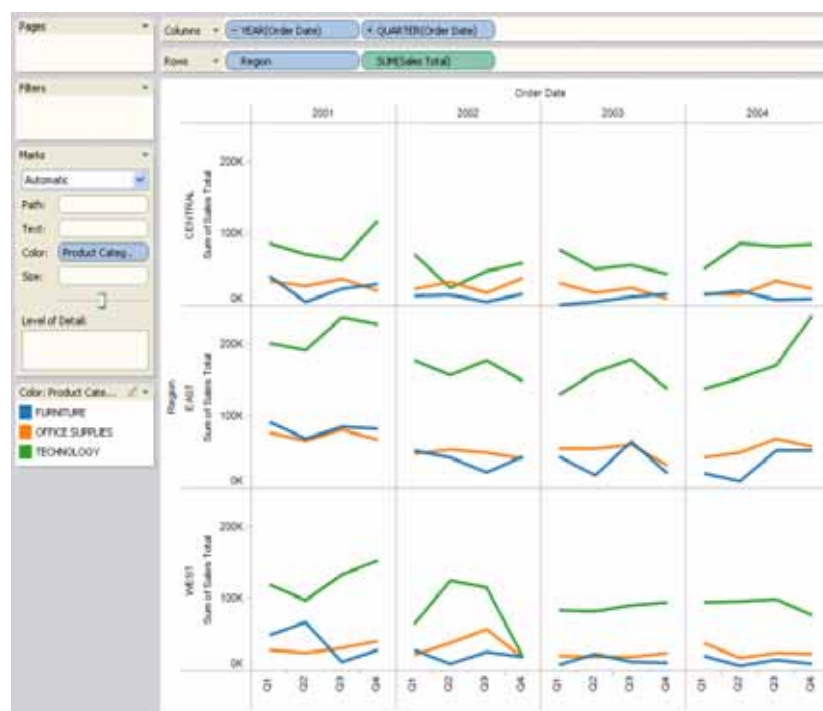
If the user you select cannot see any data, a “blank” thumbnail will be used. The blank thumbnail is shown below.



**Note** The Generate Thumbnails as User option in the publish dialog box is only available when the workbook contains one or more active user filters.

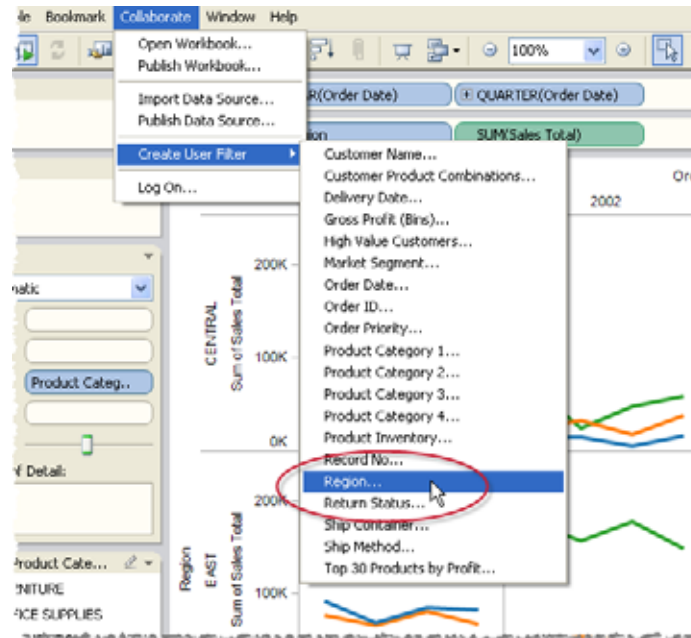
This is a simple example that explains how to set user filters on a sales report so that when it is published regional managers will only see data for their respective region. This example uses the Sample- Superstore Sales (Excel) data source that comes with Tableau.

The sales report is shown below.



As you can see the view shows the quarterly sales over several years for each region and product type. Now suppose you are going to publish this view to the server so each of the three regional managers can track their sales. However, you want to limit each manager to only see data relevant to their region. To do that you need to set up a user filter.

- 1 Select **Collaborate > Create User Filter**. Then select **Region** because that the field you want to use for filtering the view.




- 2 You may need to log into the server if are not already logged in. Log in by typing the Server name, your user name, and password. You user name and password is either specific for the server or is your Windows user name and password.



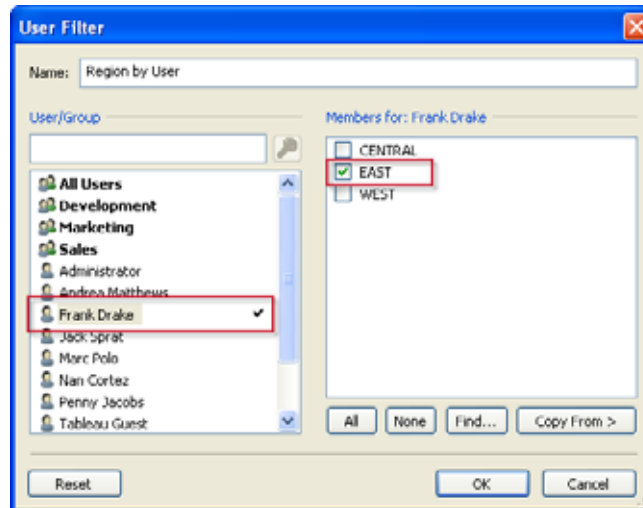
- 3 Type a name for the user filter. In this case we'll name it Region by User. After

specifying the filter, it will display in the Sets area of the Data window by this name.



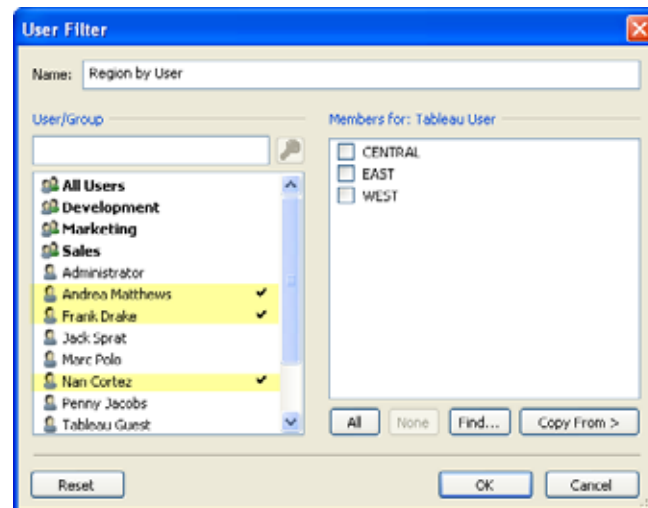
- 4 Then select a user or group on the left side for who you want to filter data. Then on the right side of the dialog box, select the members of the Region field that the selected user

or group is allowed to see. For example, below Frank Drake is the Eastern Regional Manager, so you select his name and then select East.



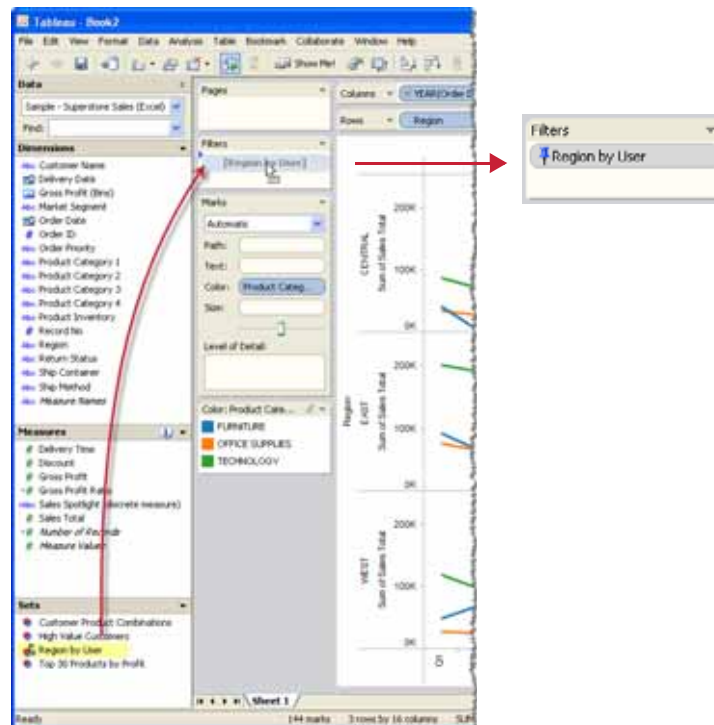
- 5 Repeat this process as many times as possible. Here we have three managers so Frank Drake gets to see the East, Andrea Matthews gets to see the Central region, and Nan Cortez gets to see the West region.





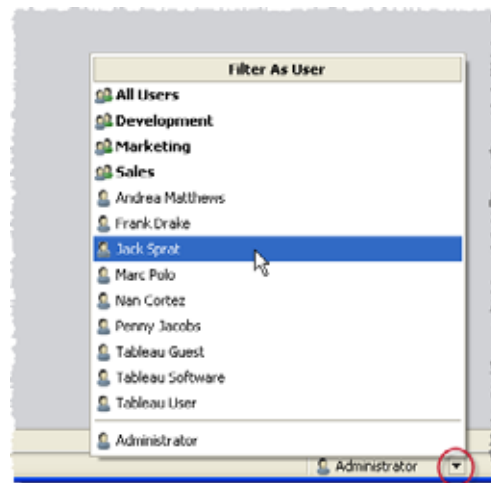
- 6 When finished click **OK**.
- 7 The user filter displays in the Sets area of the Data window. Now that it is defined you need to apply it to a view. Simply drag the Region by User user filter to the Filters shelf.

The filter becomes a context filter. Refer to “Context Filters” on page 12-55 to learn more.




**Note** When you apply a user filter to a view, it’s possible that you don’t see anything. For example, in this case you have not allowed yourself to see either of the three regions. You can edit a user filter by right-clicking it in the Data window and selecting **Edit**.

At anytime you can preview what the view will look like for each user using the User Filter menu in the bottom right corner of the workbook. This menu lists all users and groups. Selecting a user or group lets you preview what the selected user will see after the view is published.




For example, below are the three views each of the regional managers will see respectively.


Frank Drake - West



Andrea Matthews - Central



Nan Cortez - East



Now when you publish the view to the server each users will only see the data you have allowed them to see.


Sometimes you may have a field that contains your users in the data source. You can create user filters based on the names in the field rather than creating a set using the users on the server. For example, in a view showing the sales performance for several employees you could set a filter that only shows the data for the employee that is currently signed into the server. User filters that are based off of users that are part of your data source can be created using a calculated field and the Special functions.

#### Example of a user filter using users from the data source:

This view shows the Annual sales performance for a list of managers. When the view is published, you may only want each manager to see their own sales numbers. In order to do that you need to create a user filter that restricts the manager field to only include the user that is currently logged in.



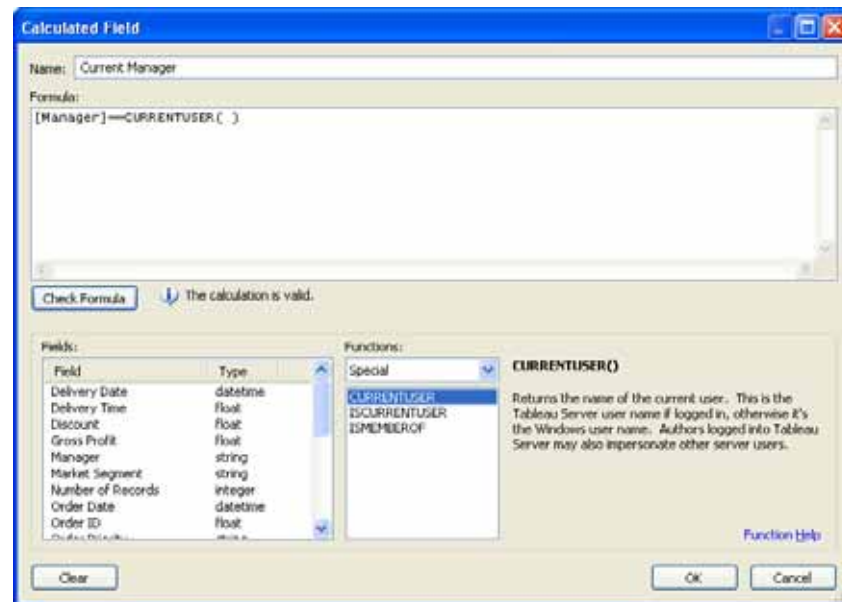
- 1 If you have not already logged in, select **Collaborate > Log In** to log into Tableau Server.
- 2 Create a calculated field by selecting **Create Calculated Field** on the Data window menu.




- 3 In the Calculated Field dialog box name the field. In this example we'll call it Current Manager.



- 4 Then use the Fields and Special functions in the bottom half of the dialog box to create calculation that looks like the one below.



- 5 When finished click **OK**.
- 6 The field displays in the Dimensions area of the Data window.



- 7 Place the new field on the Filters shelf. If you are not one of the users in the view (in this example one of the managers) the only option in the Filter dialog box will be False. In

that case you should select nothing and simply click **OK**. Then you can preview the view as another user and edit the filter to include True.

Otherwise, if you are one of the managers you can select **True** so you only include managers who match the current user.

You can preview what other users will see using the User Filter menu in the bottom right corner of the workbook. Refer to “Previewing User Filters” on page 23-29 to learn more.

The view below shows what Dave Hallstren would see if he was logged in.

Manager	2001	2002	2003	2004
Dave Hallstren	OK BK 10K Sum of Sales Total	OK BK 10K Sum of Sales Total	OK BK 10K Sum of Sales Total	OK BK 10K Sum of Sales Total

Refer to “Special Functions” on page 26-23 for more information about each special function available.

---

<b>Saving Your Work . . . . .</b>	<b>24-3</b>
Workbooks . . . . .	24-3
Bookmarks . . . . .	24-7
Packaged Workbooks . . . . .	24-5
 <b>Reverting Workbooks . . . . .</b>	 <b>24-9</b>
 <b>Exporting Your Work . . . . .</b>	 <b>24-10</b>
Export Data . . . . .	24-10
Export as an Image . . . . .	24-21
 <b>Exporting the Data Source Connection . . . . .</b>	 <b>24-24</b>



---

When you create views in Tableau, you will most likely want to save or share your work. You can export your work in a number of different formats to be used in different applications such as Microsoft PowerPoint and Microsoft Excel. This section discusses the following saving and exporting topics:

- Saving Your Work – how to save either the entire workbook or a single worksheet within a workbook.
- Exporting Your Work – export the data to another application such as PowerPoint presentations and other documents.
- Exporting the Data Source Connection – save changes you make to a data source such as added sets and calculations.

When you create useful views of your data, you should save the results. Tableau provides three ways for you to save your work:

- Workbooks – Saves all open worksheets.
- Packaged Workbooks – Saves the workbook along with all referenced local file data sources and images into a single file.
- Bookmarks – Saves the current worksheet.


You can share workbooks and bookmarks with your co-workers provided they can access the relevant data sources. If your co-workers do not have access to the data sources you can save a packaged data file.

Note that custom fields such as binned measures, calculated fields, groups, and sets are saved with the workbook or bookmark.

When you open Tableau, it automatically creates a new workbook. Workbooks hold the work you create and consist of one or more worksheets. Each worksheet contains a particular view of your data.

**To save a Tableau workbook:**

- 1 Select **File > Save** or type **Ctrl+S**.



- 2 Specify the workbook file name in the **Save As** dialog box.

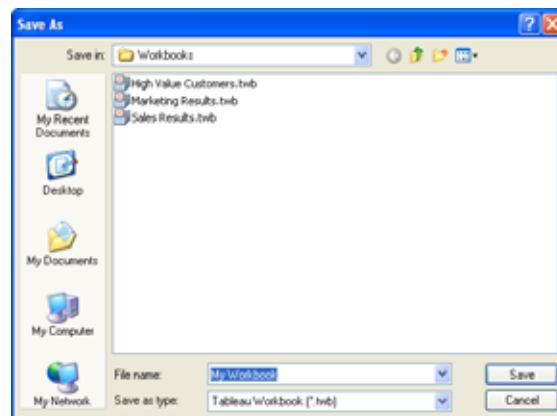


Tableau saves the file with the .twb extension. The default location is the **Workbooks** folder of the Tableau repository. However, you can save Tableau workbooks to any directory you choose.

---

**Note** Tableau file names cannot include any of the following characters: forward slash (/), backslash (\), greater-than sign (>), less-than sign (<), asterisk (\*), question mark (?), quotation mark ("), pipe symbol (|), colon (:), or semicolon (;).

---

To save an extra copy of a workbook that you already have open, select **File > Save As** and proceed by saving the file with a new name.

Workbooks often reference external resources. For example, workbooks might reference local file data sources such as Excel, Access, and Extract files and sometimes reference background images. When you save the workbook, it is linked to these resources. The next time you open the workbook the views are automatically updated with any changes made to the data and images. While in most cases you will want to save the workbook in this way, if you are sharing it with someone who does not have access to the referenced resources or publishing the workbook to Tableau Server, you can save a packaged workbook instead.


Packaged workbooks contain the workbook along with a copy of any local file data sources and background images. The workbook is no longer linked to the original data sources, rather it points to the copy that is included in the package. These workbooks are saved with the .twbx file extension. Others can open the packaged workbook using Tableau.

**To save a packaged workbook:**


- 1 Select **File > Save As**.



- 2 Specify a file name for the packaged workbook in the Save As dialog box.



- 3 Select **Tableau Packaged Workbooks** on the Save as type drop-down list.



- 4 Click **Save**.

The default location is the **Workbooks** folder of the Tableau repository. However, you can save packaged workbooks to any directory you choose.

The following types of files are included in packaged workbooks:

- Excel Files
- Access Files
- Text Files

- Tableau Data Extract Files
- Local Cube Files
- Background Image Files

Packaged Workbooks can be unpackaged at anytime in the Windows Explorer.

**To unpackage a workbook:**

- Right-click the packaged workbook file (.twbx) in Windows Explorer and select **Unpackage**.

When you unpackage a workbook you will see the regular workbook file (.twb) along with a folder that contain any data sources and images that were packaged with the workbook.

You can save a single worksheet as a Tableau bookmark. Bookmarks can be accessed from any workbook using the Bookmarks menu. Bookmarks are convenient when you have a variety of worksheets that you like to access frequently.

**To save a Tableau bookmark:**

- 1 Select **Create Bookmark** from the **Bookmark** menu.



- 2 Specify the bookmark file name and location in the Create Bookmark dialog box.

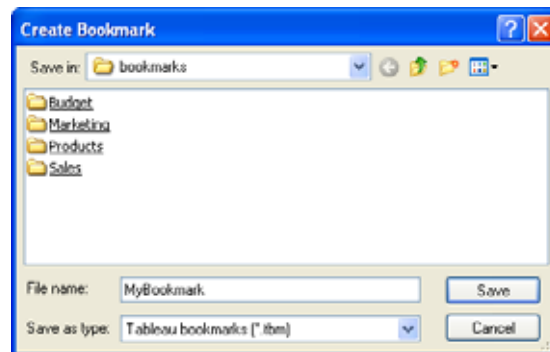


Tableau saves the file with the .tbm extension. The default location is the **Bookmarks** folder of the Tableau Repository. However, you can save bookmarks to any location you choose. Bookmarks that are not stored in the Tableau repository do not display in the **Bookmark** menu.

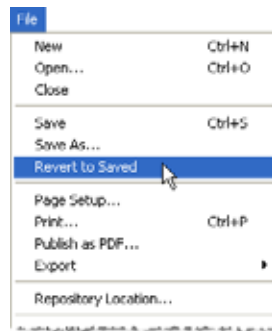
You can organize bookmarks into folders in the same way you organize files or documents. This is especially useful when you have a large number of bookmarks to manage. For example, you might organize bookmarks based on employee name, product types, sales results, and so on. You can organize bookmarks by creating a new folder, renaming an existing folder, renaming existing bookmark files, and so on.

Delete bookmarks in the same way you would delete any other file on your computer. After you delete a bookmark from the **Bookmarks** folder in the Tableau Repository, it is removed from the Bookmarks menu the next time you start the application.


Sometimes you will want to undo all of the changes you've made to a workbook and just start back to how it was when you opened it. Rather than clicking through your history using the Undo button, you can revert the workbook to the last saved state. When you revert to saved all unsaved changes are discarded.

**To revert to the last saved state of a workbook:**

- 1 Select **File > Revert to Saved**.



- 2 Click **Revert** in the warning dialog box.





After you have created several data views, you might want to export your results to other applications. Tableau provides several methods for you to export your work:

- Export Data – Copy the data from a view to an Excel worksheet or export as an Access database.
- Export as an Image – Copy images of your views into other applications such as Microsoft Office or PowerPoint. You can also include the images in web pages.

Exporting your results is a convenient way to share your work with coworkers who do not have access to Tableau, or to include your work as part of a presentation or document. To learn more about printing and publishing your work refer to Chapter 25, “Printing”.

---

**Note** You can also use Tableau to present your data. Refer to “Presentation Mode” on page 2-15 to learn more about presenting your work.

---

Suppose you want to export data from Tableau to another application, or create a new data source that contains a portion of the records in your original data source. There are several ways to complete these tasks in Tableau. You can

- Copy Records To Clipboard
- Copy Underlying Records to Clipboard
- Export Records To Microsoft Access
- Copy Cross-tab to Clipboard
- Export Cross-tab to Excel
- Extracting Data

All of these methods are discussed in this section.

---

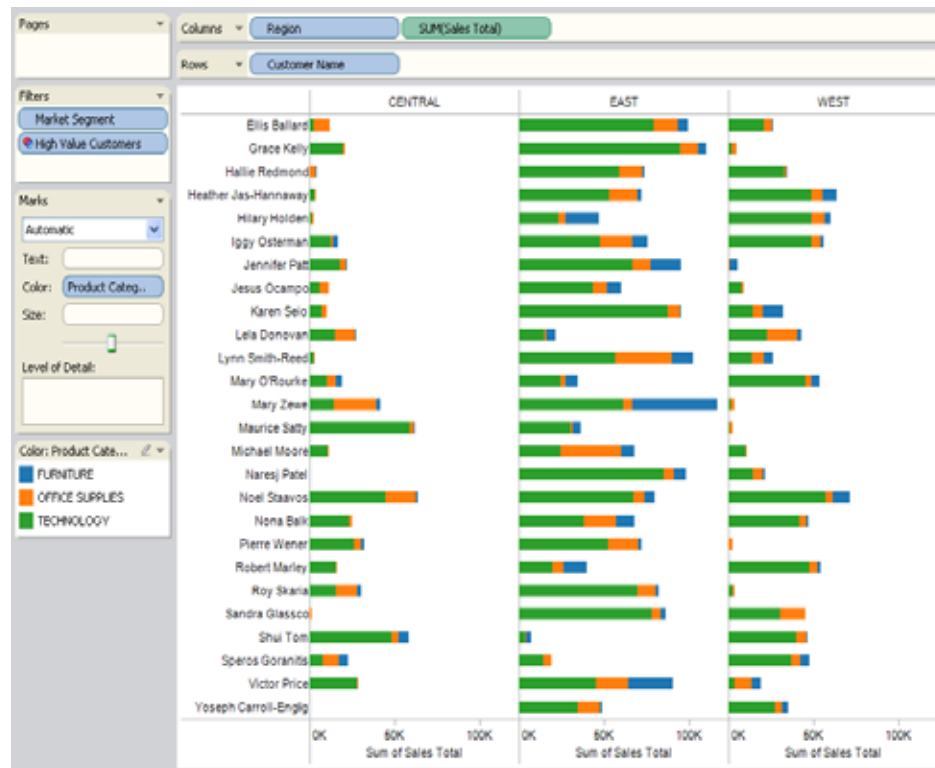
**Note** Excel worksheets hold 64,000 rows of data. If you copy more than this number of rows, Tableau pastes only the first 64,000 rows of the selected set into Excel. For datasets larger than 64,000 it is recommended that you export to Access instead.

---

When exporting data, you should keep these rules in mind:

- You can select any portion of a data view to export. Refer to Chapter 14, “Inspecting Data” to learn how to select data. If you want to export all data in a view, choose **Select All** on the **Edit** or right-click context menu. Copying and exporting to a cross tab always exports all the data in the view regardless of what you have selected.
- The fields that are exported to the new data source come from the fields placed on the worksheet shelves. The exception is fields that are external filters, and appear only on the **Filters** shelf.
- If you want to include other fields (either dimensions or measures) with the exported data without changing the basic view, you should place those fields on the **Level of Detail** shelf.

For example, you might create a view that contains data consisting of only high value customers, and then create a new data source containing only the data for those customers. The following data view shows high value customers displayed as a bar chart. The view is used in the following two sections, which describe how to export the data to an Excel worksheet, an Access database, and extract the data.

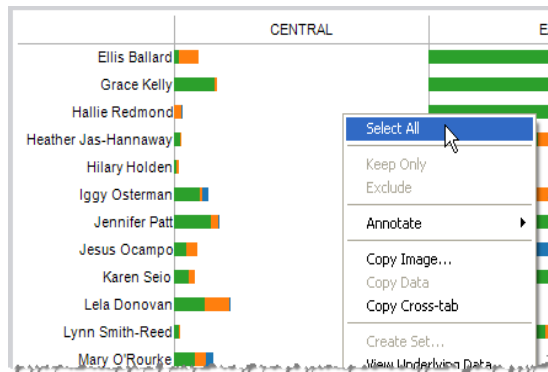


**Note** When exporting data to a Microsoft Office application, remember that Tableau only support Office 2000 or higher (including Office 2007).

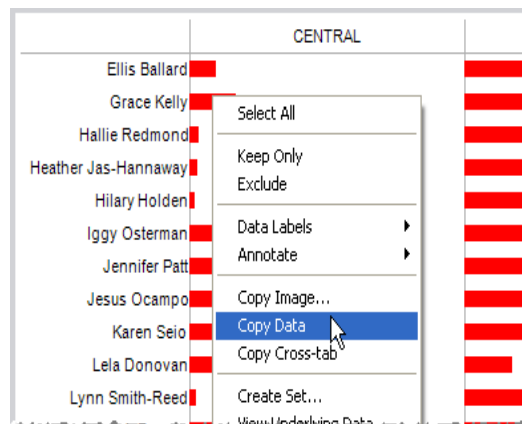
### Copy Records To Clipboard

Typically this function is used to copy records from Tableau into Microsoft Excel. To create an Excel spreadsheet from Tableau data, follow these three steps:

- 1 Select the desired data in Tableau. For this example, all the data are selected.

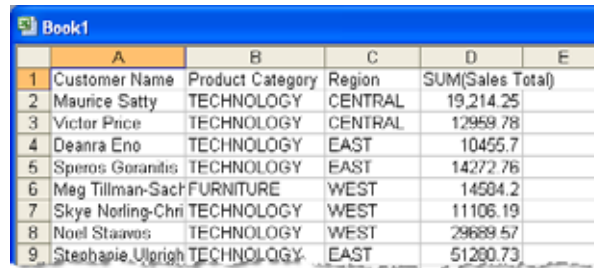


- 2 Select **Copy Data** from the **Edit** menu or right-click the view and select **Copy Data** from the context menu (as shown below).



- 3 Open an Excel worksheet and select **Paste** from Excel's **Edit** menu or press **Ctrl+V**. Notice that the fields placed on the **Rows**, **Columns**, and **Color** shelves are copied into

the worksheet. However, the **Market Segment** field is not copied because it is an external filter (it appears only on the **Filters** shelf).



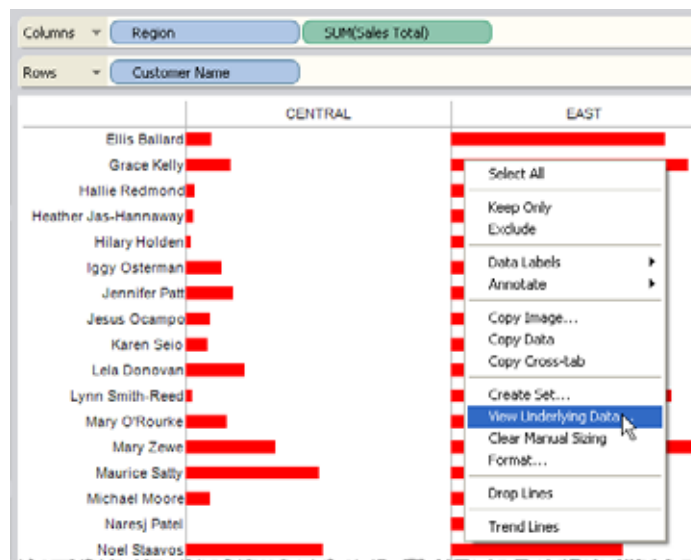
	A	B	C	D	E
1	Customer Name	Product Category	Region	SUM(Sales Total)	
2	Maurice Satty	TECHNOLOGY	CENTRAL	19,214.25	
3	Victor Price	TECHNOLOGY	CENTRAL	12,959.78	
4	Deanra Eno	TECHNOLOGY	EAST	10,455.7	
5	Speros Goranitis	TECHNOLOGY	EAST	14,272.76	
6	Meg Tillman-Sach	FURNITURE	WEST	14,584.2	
7	Skye Norling-Chri	TECHNOLOGY	WEST	11,106.19	
8	Noel Staavos	TECHNOLOGY	WEST	29,689.57	
9	Stephanie Ulgrigh	TECHNOLOGY	EAST	51,280.73	

### Copy Underlying Records to Clipboard

Copying underlying data can be used to copy the disaggregated data behind a view. For a full discussion of drill-through, refer to “Viewing Underlying Data (Drill-Through)” on page 14-21.

To copy underlying records:

- 1 Select the desired data in Tableau.
- 2 Right-click the selected records and select **View Underlying Data** on the context menu.



- 3 In the resulting dialog box, select the data you want to copy. You can select all the data by clicking **Select All**. Click **Copy to Clipboard**. Notice that you can add more columns

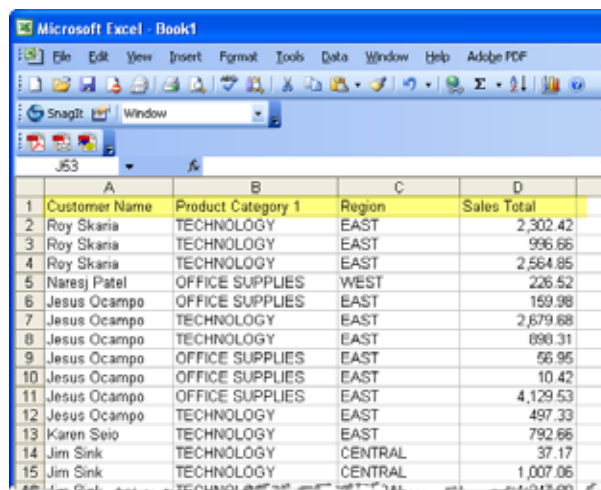
than those that are referenced by selecting **Show all fields** in the upper right hand corner of the dialog box.



Jesus Ocampo	OFFICE SUPPL	CENTRAL
Heather Jap-Han	OFFICE SUPPL	WEST
Heather Jap-Han	OFFICE SUPPL	WEST
Lynn Smith-Reed	TECHNOLOGY	EAST
Lynn Smith-Reed	FURNITURE	EAST
Lynn Smith-Reed	TECHNOLOGY	EAST

Select All Copy to Clipboard

- 4 Open an Excel worksheet and select **Paste** from Excel's Edit menu or press **Ctrl+V**. Notice that the fields placed on the Rows, Columns and Color shelves are copied onto the worksheet.

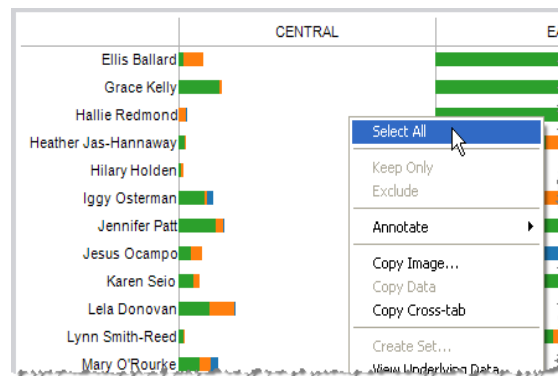


	A	B	C	D
1	Customer Name	Product Category 1	Region	Sales Total
2	Roy Skaria	TECHNOLOGY	EAST	2,302.42
3	Roy Skaria	TECHNOLOGY	EAST	966.66
4	Roy Skaria	TECHNOLOGY	EAST	2,564.85
5	Naresj Patel	OFFICE SUPPLIES	WEST	226.52
6	Jesus Ocampo	OFFICE SUPPLIES	EAST	159.98
7	Jesus Ocampo	TECHNOLOGY	EAST	2,679.68
8	Jesus Ocampo	TECHNOLOGY	EAST	890.31
9	Jesus Ocampo	OFFICE SUPPLIES	EAST	56.95
10	Jesus Ocampo	OFFICE SUPPLIES	EAST	10.42
11	Jesus Ocampo	OFFICE SUPPLIES	EAST	4,129.53
12	Jesus Ocampo	TECHNOLOGY	EAST	497.33
13	Karen Seio	TECHNOLOGY	EAST	792.66
14	Jim Sink	TECHNOLOGY	CENTRAL	37.17
15	Jim Sink	TECHNOLOGY	CENTRAL	1,007.06

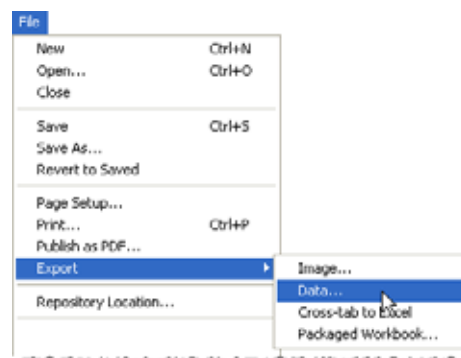
### Export Records To Microsoft Access

To create an Access database from Tableau data, follow these steps:

- 1 Select the desired data in Tableau. For this example, all the data are selected.

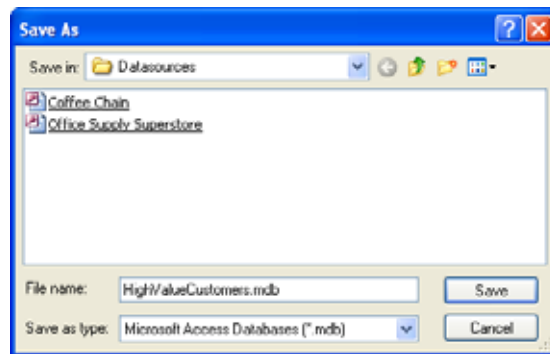


- 2 Select **File > Export > Data**.

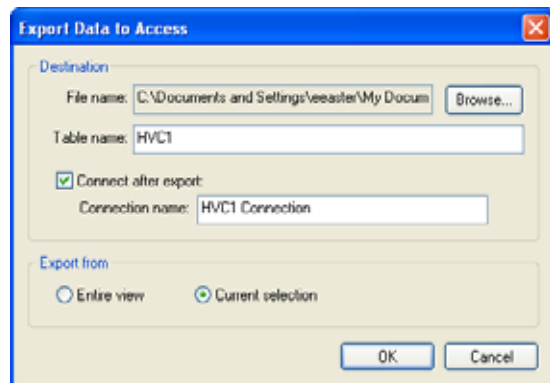


- 3 Select a location and name for your Access database by completing the **Save As** dialog box. Access databases end with the .mdb file extension.





- 4 Complete the Export Data to Access dialog box. For this example, name the table HVC1. The Connect after export option allows you to immediately connect to the new data source and continue working without interrupting your work flow.



### Copy Cross-tab to Clipboard

You can copy a cross-tab (text table) version of a view to the clipboard and transfer it to another application. For instance, you might want to transfer a cross-tab in Tableau to a cross-tab in Microsoft Excel. Or you may want to transfer the data behind a graphical view in Tableau to Excel in a cross-tab format. Copying a cross-tab to clipboard is restricted by the following general rules:

- This feature applies to all records in the view. It cannot be used on a subset of records.

- You can use this feature for aggregated views only. It cannot be used on disaggregated views of data, because a cross-tab is by definition an aggregated view of data. In other words, the **Aggregate Measures** option on the **Analysis** menu must be on in order for this function to work properly. For a full discussion of Aggregated and Disaggregated data refer to “Aggregations” on page 16-4

Other restrictions may apply depending on the data in your view. You cannot copy a cross-tab if the view contains continuous dimensions such as continuous dates and times.

To copy a view as a cross-tab to the clipboard:

- 1 Right-click any view in Tableau and select **Copy Cross-tab**. This copies all data in the current view to the clipboard in a cross-tab format.



- 2 Open an Excel worksheet and select **Paste** from Excel's Edit menu or press **Ctrl-V**.

The screenshot shows a Microsoft Excel worksheet with the following data:

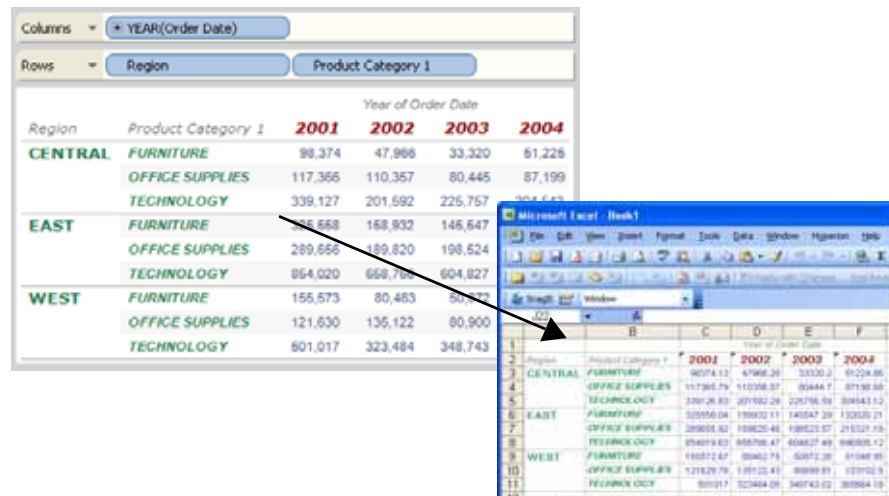
	A	B	C	D	E
1			CENTRAL	EAST	WEST
2	Barbara Fisher	FURNITURE		14,966	
3	Barbara Fisher	OFFICE SUPPLIES	4,211	11,473	1,645
4	Barbara Fisher	TECHNOLOGY	2,362	8,722	2,759
5	Dario Medina	FURNITURE	6,265	2,782	6,766
6	Dario Medina	OFFICE SUPPLIES	5,696	4,403	3,076
7	Dario Medina	TECHNOLOGY	8,782	40,056	7,560
8	Dave Hallsten	FURNITURE		5,063	
9	Dave Hallsten	OFFICE SUPPLIES	1,749	5,414	6,213
10	Dave Hallsten	TECHNOLOGY	3,795	19,642	4,681
11	Dave Flushing	FURNITURE	58	4,233	2,345
12	Dave Flushing	OFFICE SUPPLIES	9,860	49,800	15,322
			2,024	17,492	26,069

Notice that the pasted data always appears as a cross-tab in Excel even if the initial view of the data in Tableau was not in a cross-tab format.

### Export Cross-tab to Excel

There is a more direct way to transfer a cross-tab view of data to Microsoft Excel. Select **File > Export > Cross-tab to Excel**. Tableau automatically pastes a cross-tab version of

the current view into a new Excel workbook. This option automatically opens a new instance of the Excel application.



**Note** Although, copying a cross-tab to Excel is more direct, it can decrease performance because it is copying the formatting as well as the data. If the view you are exporting contains a lot of data, a dialog box opens asking whether you want to copy the formatting options. Disregarding the format saves time.

## Extracting Data

Another way to export all or subsets of your data to a new data source is to use Tableau's Extract feature. To learn more about extracts refer to "Extracting Data" on page 6-1.

Suppose you want to transfer your Tableau results into a presentation, report or web page. Tableau gives you several options:

- Copy to Another Windows Application
- Export to an Image File
- Publish as PDF

## Copy to Another Windows Application

To insert an image of your Tableau results into another Windows application such as PowerPoint, Word, or Excel, follow these steps:

- 1 Select **Copy > Image** on the **Edit** menu.
- 2 In the Copy Image dialog box, select the contents you want to include in the image and the legend layout if the view contains a legend.
- 3 When finished click **Copy**. When you do this, Tableau copies the current data view to the Windows Clipboard.
- 4 Open the target application and select one of the following:
  - **Paste** from the **Edit** menu or type **Ctrl+V**.
  - **Paste Special** from the **Edit** menu. If you select Paste Special you can enable enhanced visual and printer quality options.

In the Paste Special dialog box select how you want to paste the image. The Special dialog box in Microsoft PowerPoint Paste is shown below.



In most cases, paste the images as an Enhanced Meta File to get the best presentation quality.

## Export to an Image File

The export image command saves the current view as an image file. You can export to an image file with the following three steps.

- 1 Select **Export > Image** from the **File** menu.
- 2 In the Export Image dialog box, select the contents you want to include in the image and the legend layout if you are including a legend. When finished click **Save**.
- 3 In the Save Image dialog box, navigate to where you want to save the image file and type a file name into the text box. Select a file format from the **Save as type** drop-down menu. When finished click **Save**.

### **Publish as PDF**

You can publish one or more views to PDF by selecting **File > Publish as PDF**. For more information refer to “Publish as PDF” on page 25-7.

When you first connect to a data source you have the option to save the connection to your repository. Saving the connection creates a shortcut to the data source and lets you avoid having to create a new connection every time you want to use that source. If you decide not to save the data source upon connection, you can always export the data source at any time you are connected. Exporting the data source is useful if you didn't save the connection when you first connected but you want to later or if you have added custom fields such as ad-hoc groups and sets to the Data window. For more information about exporting the data source refer to "Exporting the Connection" on page 4-68.

---

<b>Overview . . . . .</b>	<b>25-2</b>
<b>Printing Results . . . . .</b>	<b>25-3</b>
Page Setup . . . . .	25-3
Printing . . . . .	25-5
<b>Publish as PDF . . . . .</b>	<b>25-7</b>
<b>Printing the Help . . . . .</b>	<b>25-8</b>
Printing the Online Help . . . . .	25-8



---

After creating a view or several views in Tableau you can print them. The first thing you should do before printing is specify how you want the printed page to look using the Page Setup settings. Then you can print to a printer or publish to a PDF. You can also print the Tableau Help directly from your web browser or by obtaining a PDF. This section discusses the following printing topics:

- **Printing Results** – Set up the page and print one or more worksheets within a workbook.
- **Publish as PDF** – Export one or more worksheets to a PDF document.
- **Printing the Help** – Print an individual help topic or obtain a PDF of the Tableau documentation.

Once you have a view or several views created in Tableau, you can print them or publish them as a PDF. This section walks you through the following topics:

- Page Setup – how to specify the page setup options
- Printing – print one or more worksheets to a hard copy.

Before printing there are several options you can set to specify how the worksheet will look when it is printed. For example, you can select which elements to include, printed page orientation, where you want to put the legend, margins, and more. These settings are specified in the Page Setup dialog box. You can set different page setup options for each worksheet in the workbook. That way you can have different titles, captions, legend settings, etc. for each worksheet you want to print. To open the page setup dialog box select **File > Page Setup**.

When the Page Setup dialog box opens, the following four categories of settings display:

- Contents – select which elements you want to show on the printed page as well as how to print the pages defined by the Pages shelf.
- Layout – specify margins, centering, and legend layout.
- Print Scaling – scale the view to fit on the specified number of pages.
- Title/Caption – add text that will be printed with each worksheet.

### Contents

Use the Contents tab to select the elements you want to show when you print. You can show or hide the title, view, caption, color legend, shape legend, and size legend.

If you have used the Pages Shelf to build your view, you can select whether to print the current page only or all pages. To learn more about the Pages Shelf refer to “Pages Shelf” on page 10-12.

### Layout

Use the Layout tab to specify the page margins, centering, and legend layout.

**Margins.** Specify top, bottom, left, and right margins by typing values into the text boxes.

**Centering.** Optionally, select whether to center the view horizontally or vertically on the page.

**Legend Layout.** If you include one or more legends, you can specify how you want the legends to appear on the printed page. Select a legend layout from the options at the bottom of the dialog box.

### Print Scaling

Use the Print Scaling tab to fit the view to a certain size and change the page orientation. These options only affect printed documents. The scaling options you specify here will not affect exported image or PDF publishing. For more information on publishing to a PDF refer to “Publish as PDF” on page 25-7.


**Print Scaling.** You can scale your view to fit within a single page or scale it across multiple pages. Select from the following options:

- Scale to – Scales the view to the specified percentage of its normal size.
- Fit to – Scales the view to fit within the specified area. Select the number of printed pages across and down. For example, if you have a really wide view that is not very tall, you can specify to fit it to three pages across by one page down.

**Page Orientation.** Use the page orientation settings to specify how you want the view oriented on the printed page. For example, if you have a view that is really wide but not that tall you should select the **Landscape** orientation. Select from the following page orientation options:

- Use Printer Setting – Use the page orientation that is already specified by the printer.
- Portrait – Rotates the view so that it is oriented vertically on the printed page.
- Landscape – Rotates the view so that it is oriented horizontally on the printed page.

The following diagram shows the difference between portrait and landscape page orientations.



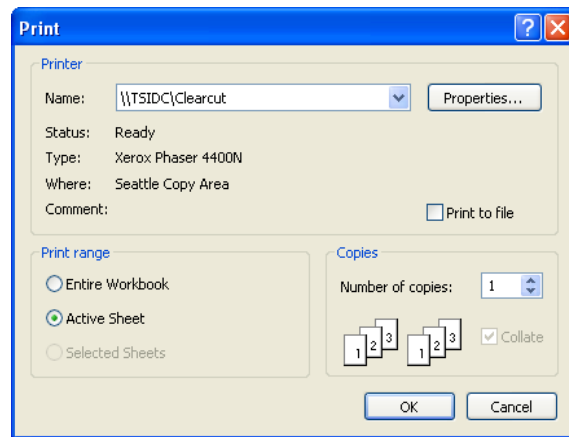
### Title/Caption

Use the Title/Caption tab to add custom text that is printed on the top or bottom of every page. An automatic title and caption are created for you based on your view. To modify the title and caption, select **Manual** and type the text you want into the respective text box.

Use the drop-down lists to insert automatic text such as the sheet name, page number, etc. The following automatic text options are supported:

- SheetName – displays the name of the worksheet.
- PageName – displays the name of the page based on the fields on the Pages shelf.
- PageNumber – displays the the page number based on the fields on the Pages shelf.
- PageCount – displays the total number of pages in a worksheet based on the fields on the Pages shelf.

After you have specified the Page Setup settings (refer to “Page Setup” on page 25-3), you can print by selecting **File > Print**. In the Print dialog, select a printer, specify a print range, and select the number of copies you want to print.



### Changing the Print Range

When you print from a workbook with multiple worksheets, each worksheet represents one or more printed pages, depending on the page set-up.

If you select **All** in the Print range area, all of the sheets will print on their own separate pages. You can print specific sheets by specifying a range of sheets.

Select from the following print ranges:

- Entire Workbook - Prints all the worksheets in the workbook.
- Active Sheet - Prints only the sheet currently displayed in the workbook.
- Selected Sheets - Prints the selected sheets.

---

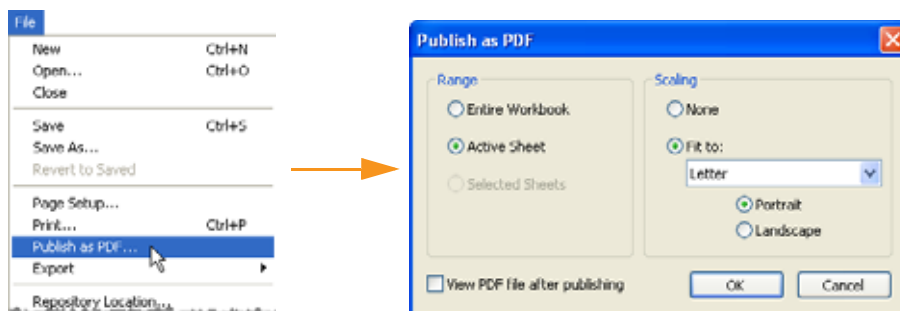
**Note** You can select multiple worksheets in a workbook by holding down the CTRL or Shift keys while clicking the worksheet tabs that you want to select.

---

You can publish views as PDF files rather than printing them as hard copies, using the Publish as PDF command. You do not need to have Adobe Acrobat installed on your computer.

**To publish as a PDF:**

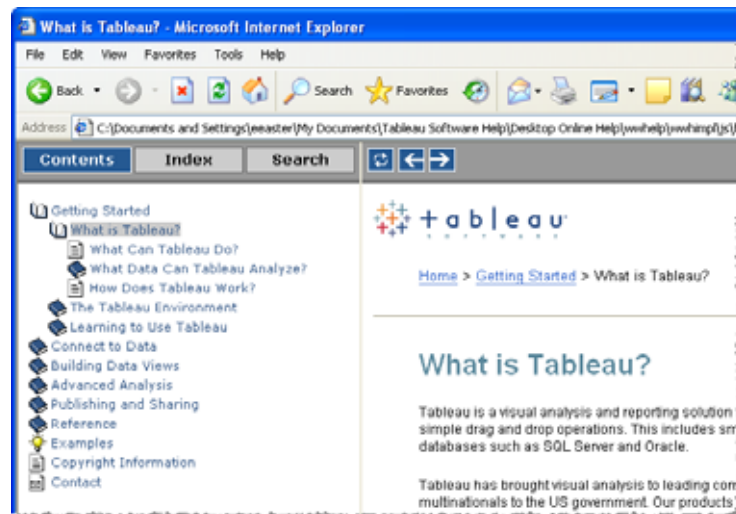
- 1 Specify page setup options for each sheet in your workbook (refer to “Page Setup” on page 25-3).
- 2 Select **File > Publish as PDF**.



- 3 In the Publish as PDF dialog box, select the print Range:
  - Entire Workbook - Publishes all the sheets in the workbook.
  - Active Sheet - Publishes only the sheet currently displayed in the workbook.
  - Selected Sheets - Publishes the selected sheets.
- 4 Select one of the following Scaling options:
  - None - Publishes each sheet at its actual size, no scaling occurs.
  - Fit to - Scales each sheet to fit onto the specified paper size and orientation.
- 5 Optionally select **View PDF file after publishing** if you want to automatically open the PDF when you are done publishing. This option is only available if you have Adobe Acrobat Reader or Adobe Acrobat installed on your computer.
- 6 Click **OK** and select where you want to save the PDF. Then click **Save**.

## Printing the Online Help

You can print an individual Help topic using the print options in your Web browser.



In addition to printing individual help topics, you can also download an offline help system and a printable PDF. Refer to “Accessing the Help” on page 39-1 to learn more.

# Functions, Operators, & Data Types

<b>Data Types</b> . . . . .	<b>A-3</b>
<b>Formatting Literals</b> . . . . .	<b>A-5</b>
Numeric Literals . . . . .	A-5
String Literals . . . . .	A-5
Date Literals . . . . .	A-5
Boolean Literals . . . . .	A-5
Null Literals . . . . .	A-5
<b>Functions</b> . . . . .	<b>A-6</b>
Number Functions . . . . .	A-6
String Functions . . . . .	A-10
Date Functions . . . . .	A-13
Type Conversion . . . . .	A-15
Logical Functions . . . . .	A-17
Aggregate Functions . . . . .	A-19
Pass Through Functions (RAWSQL) . . . . .	26-21
Special Functions . . . . .	26-23
<b>Operators</b> . . . . .	<b>A-25</b>
Precedence . . . . .	A-27



---

## Overview

This section explains how to use and combine the various data types supported by Tableau. In addition, this section discusses how to format and use the building blocks of formulas in Tableau. These parts include literal expressions, functions, and operators.

All of these features are important to understand when you create custom fields such as calculations.

- Data Types – The data types supported by Tableau and how to use and combine them.
- Formatting Literals – Using literal expressions to represent numbers, strings, dates, and more.
- Functions – How to use the functions available in Tableau to build formulas.
- Operators – How to use the operators available in Tableau to build formulas.

## Data Types

Tableau supports string, date/datetime, number, and boolean data types. These data types are automatically handled in the proper fashion. However, if you create calculated fields of your own, you need to be aware of how to use and combine the different data types in formulas. For example, you cannot add a string to a number. Also, many functions that are available to you when you define a calculation only work when they are applied to specific data types. For example, the DATEPART() function can accept only a date/datetime data type as an argument. So, you can write DATEPART('year',#April 15,2004#) and expect a valid result: 2004. You cannot write DATEPART('year',"Tom Sawyer") and expect a valid result. In fact, this example returns an error because "Tom Sawyer" is a string, not a date/datetime.

---

**Note** Although Tableau will attempt to fully validate all calculations, some data type errors cannot be found until the query is run against the database. These issues appear as error dialogs at the time of the query rather than in the calculation dialog box.

---

The data types supported by Tableau are described below. Refer to “Type Conversion” on page 26-15 to learn about converting from one data type to another.

### STRING

A sequence of zero or more characters. For example, "Wisconsin", "ID-44400", and "Tom Sawyer" are all strings. Strings are recognized by single or double quotes. The quote character itself can be included in a string by repeating it. For example, 'O"Hanrahan'.

### DATE/DATETIME

A date or a datetime. For example "January 23, 1972" or "January 23, 1972 12:32:00 AM". If you would like a date written in long-hand style to be interpreted as a date/datetime, place the # sign on either side of it. For instance, “January 23, 1972” is treated as a string data type but #January 23, 1972# is treated as a date/datetime data type.

### NUMBER

An integer or floating-point number. For example 10, 10.5, 3.999999, 1.23E+11, and 1.0E-6. Sometimes your data source may contain more precise values than Tableau can model. If you use these values in a view, Tableau displays a precision warning letting you know that the value has been truncated. For more information about precision warnings refer to “Precision Warnings” on page 7-6.

**BOOLEAN**

A field that contains the values TRUE or FALSE. An unknown value arises when the result of a comparison is unknown. For example, the expression `7 > Null` yields unknown. Unknown booleans are automatically converted to Null

## Formatting Literals

When you are using functions you will sometimes want to use literal expressions to represent numbers, strings, dates, and more. A literal expression signifies a constant value that is represented “as is.” For example, you may have a function where your input is a date. Rather than just type “May 1, 2005”, which would be interpreted as a string, you would type #May 1, 2005#, which is equivalent to using a date function to convert the argument from a string to a date (refer to “Date Functions” on page 26-13). You can use numeric, string, date, boolean, and Null literals. The way to format each of these literals is described below.

### Numeric Literals

A numeric literal is written exactly like you usually write numbers. If you want to input the number one as a numeric literal you would type 1. Subsequently, if you want to input the number 3.1415 as a numeric literal you would type 3.1415.

### String Literals

A string literal can be written either using single quotations or double quotations. If your string has a single or double quotation within it, simply type the symbol twice. For example, if you wanted to input the string “cat” as a string literal you could type ‘cat’ or “cat”. Additionally, if you want to type the string “She’s my friend.” as a string literal you could type ‘She’s my friend.’ or “She’s my friend.”

### Date Literals

Date literals are signified by the pound symbol (#). If you wanted to input the date “August 22, 2005” as a literal date you would type #August 22, 2005#.

### Boolean Literals

Boolean literals are written as either true or false. If you wanted to input “true” as a boolean literal you would type true.

### Null Literals

Null literals are written simply as Null. If you wanted to input “Null” as a Null literal you would type Null.

## Functions

The calculation functions are grouped into the following categories:

- Number Functions
- String Functions
- Date Functions
- Type Conversion
- Logical Functions
- Aggregate Functions
- Pass Through Functions (RAWSQL)
- Special Functions

These are the same categories used in the **Calculation** dialog box. The aggregate functions such as sum, average, and so on are described in “Aggregate Calculations” on page 16-30.

### Number Functions

#### **ABS(number)**

Returns the absolute value of the given number. For example,

$$\text{ABS}(-7) = 7$$

$\text{ABS}([\text{Budget Variance}])$  returns the absolute value for all the numbers contained in the Budget Variance field.

#### **ACOS(number)**

Returns the arc cosine of the given number. The result is in radians. For example,

$$\text{ACOS}(-1) = 3.14159265358979$$

#### **ASIN(number)**

Returns the arc sine of a given number. The result is in radians. For example,

$$\text{ASIN}(1) = 1.5707963267949$$

#### **ATAN(number)**

Returns the arc tangent of a given number. The result is in radians. For example,

$$\text{ATAN}(180) = 1.5652408283942$$

**ATAN2(y number, x number)**

Returns the arc tangent of two given numbers (x and y). The result is in radians. For example,

$$\text{ATAN2}(2, 1) = 1.10714871779409$$

**COS(number)**

Returns the cosine of a given number specified in radians. The number is in radians. For example,

$$\text{COS}(\text{PI}() / 4) = 0.707106781186548$$

**COT(number)**

Returns the cotangent of a given number specified in radians. The number is in radians. For example,

$$\text{COT}(\text{PI}() / 4) = 1$$

**DEGREES(number)**

Converts a given number in radians to degrees. For example,

$$\text{DEGREES}(\text{PI}() / 4) = 45.0$$

**EXP(number)**

Returns e raised to the power of the given number. For example.

$$\text{EXP}(2) = 7.389$$

$$\text{EXP}(-[\text{Growth Rate}] * [\text{Time}])$$

**LN(number)**

Returns the natural logarithm of a number. Returns Null if number is less than or equal to 0.

**LOG(number [, base])**

Returns the logarithm of a number for the given base. If the base value is omitted, base 10 is used.

**MAX(number, number)**

Returns the maximum of the two arguments, which must be of the same type. Returns Null if either argument is Null. MAX can also be applied to a single field in an aggregate calculation. For example,

MAX(4,7)

MAX(Sales,Profit)

MAX([First Name],[Last Name])

**MIN(number, number)**

Returns the minimum of the two arguments, which must be of the same type. Returns Null if either argument is Null. MIN can also be applied to a single field in an aggregate calculation. For example,

MIN(4,7)

MIN(Sales,Profit)

MIN([First Name],[Last Name])

**PI( )**

Returns the numeric constant pi: 3.14159.

**POWER(number, power)**

Raises the number to the specified power.

For example,

$\text{POWER}(5,2) = 5^2 = 25$

POWER(Temperature, 2)

You can also use the ^ symbol.

For example,

$5^2 = \text{POWER}(5,2) = 25$

**Radians (number)**

Converts the given number from degrees to radians. For example,

RADIANS(180) = 3.14159

**ROUND(number, [decimals])**

Rounds numbers to a specified number of digits. The decimals argument specifies how many decimal points of precision to include in the final result, and it is not required. If the decimals variable is omitted, number is rounded to the nearest integer. For example,

$\text{ROUND}(7.3) = 7$

$\text{ROUND}(-6.9) = -7$

$\text{ROUND}(123.47, 1) = 123.5$

$\text{ROUND}(\text{Sales})$  rounds every Sales value to an integer

Note that some databases such as MS SQL Server, allow specification of a negative length, where -1 rounds number to 10's, -2 rounds to 100's, and so on. This is not true of all databases to which you can connect. For example, it is not true of Excel or Access.

**SIGN(number)**

Returns the sign of a number: The possible return values are -1 if the number is negative, 0 if the number is zero, or 1 if the number is positive. For example, if the average of the profit field is negative, then

$\text{SIGN}(\text{AVG}(\text{Profit})) = -1$

**SIN(number)**

Returns the sine of a number specified in radians. The result is in radians. For example,

$\text{SIN}(0) = 1.0$

$\text{SIN}(\text{PI}() / 4) = 0.707106781186548$

**SQRT(number)**

Returns the square root of a number. For example,

$\text{SQRT}(25) = 5$

**SQUARE(number)**

Returns the square of a number. For example,

$\text{SQUARE}(5) = 25$

**TAN(number)**

Returns the tangent of a number specified in radians. The result is in radians. For example,



$\text{TAN}(\text{PI}() / 4) = 1.0$

## String Functions

### **ASCII(string)**

Return the ASCII code for the first character of string. For example,

$\text{ASCII}('A') = 65$

### **CHAR(number)**

Returns the character encoded by the ASCII code number. For example,

$\text{CHAR}(65) = 'A'$

### **FIND(string, substring, [start])**

Returns the index position of substring in string, or 0 if the substring isn't found. If the optional argument start is added, the function does the same thing, but ignores any instances of substring that appear before the index position start. The first character in the string is position 1. For example,

$\text{FIND}(\text{"Calculation"}, \text{"alcu"}) = 2$

$\text{FIND}(\text{"Calculation"}, \text{"Computer"}) = 0$

$\text{FIND}(\text{"Calculation"}, \text{"a"}, 3) = 7$

$\text{FIND}(\text{"Calculation"}, \text{"a"}, 2) = 2$

$\text{FIND}(\text{"Calculation"}, \text{"a"}, 8) = 0$

### **LEFT(string, number)**

Returns the left-most number of characters in the string. For example,

$\text{LEFT}(\text{"Matador"}, 4) = \text{"Mata"}$

### **LEN(string)**

Returns the length of the string. For example,

$\text{LEN}(\text{"Matador"}) = 7$

### **LOWER(string)**

Returns the lower case version of the string. For example,

`LOWER("ProductVersion") = "productversion"`

**LTRIM(string)**

Returns the string with any leading spaces removed. For example,

`LTRIM(" Matador ") = "Matador "`

**MAX(a, b)**

Usually applied to numbers, but also works on strings. Returns the maximum of a and b (a and b must be of the same type). With strings, MAX finds the value that is highest in the sort sequence defined by the database for that column. It returns Null if either argument is Null. For example,

`MAX ("Apple", "Banana") = "Banana"`

**MID(string, start, [length])**

Returns the string starting at index position start. The first character in the string is position 1. If the optional argument length is added, the returned string includes only that number of characters. For example,

`MID("Calculation", 2) = "alculation"`

`MID("Calculation", 2, 5) = "alcul"`

**MIN(a, b)**

Usually applied to numbers, but also works on strings. Returns the minimum of a and b (a and b must be of the same type). With strings, MIN finds the value that is lowest in the sort sequence. It returns Null if either argument is Null. For example,

`MIN ("Apple", "Banana") = "Apple"`

**RIGHT(string, number)**

Returns the right-most number of characters in string. For example,

`RIGHT("Calculation", 4) = "tion"`

**RTRIM(string)**

Returns the string with any trailing spaces removed. For example,

`RTRIM(" Calculation ") = " Calculation"`

**SPACE(number)**

Returns a string that is composed of the specified number of repeated spaces. For example,

SPACE(1) = " "

**TRIM(string)**

Returns the string with leading and trailing spaces removed. For example,

TRIM(" Calculation ") = "Calculation"

**UPPER(string)**

Returns the lower case version of the string. For example,

UPPER("Calculation") = "CALCULATION"

## Date Functions

The data functions are given below. Many of the examples use the # symbol with date expressions. Refer to “Date Literals” on page 26-5 for an explanation of this symbol. Additionally, many of these functions use date\_part, which is constant string argument. The valid date\_part values that you can use are given in the table below.

date_part	Values
'year'	Four digit year
'quarter'	1-4
'month'	1-12 or "January", "February", and so on
'dayofyear'	Day of the year; Jan 1 is 1, Feb 1 is 32, and so on
'day'	1-31
'weekday'	1-7 or "Sunday", "Monday", and so on
'week'	1-52
'hour'	0-23
'minute'	0-59
'second'	0-60

### DATEADD(date\_part, increment, date)

Returns an increment added to date. The type of increment is that specified in date\_part. For example,

```
DATEADD('month', 3, #April 15, 2004#) = #July 15, 2004#
```

This expression adds three months to the date #April 15, 2004#, and results in #July 15, 2004#.

### DATEDIFF(date\_part, date1, date2)

Returns the difference between date1 and date2 expressed in units of date\_part. For example,

```
DATEDIFF('month', #July 15, 2004#, #April 15, 2004#) = -3
```

This expression returns -3 because April is three months before July.

### **DATENAME(date\_part, date)**

Returns date\_part of date as a string. For example,

```
DATENAME('year', #April 15, 2004#) = "2004"
DATENAME('month', #April 15, 2004#) = "April"
```

### **DATEPART(date\_part, date)**

Returns date\_part of date as an integer. For example,

```
DATEPART('year', #April 15, 2004#) = 2004
DATEPART('month', #April 15, 2004#) = 4
```

### **DATETRUNC(date\_part, date)**

Truncates the specified date to the accuracy specified by the date\_part. This function returns a new date. For example, when you truncate a date that is in the middle of the month at the month level, this function returns the first day of the month.

```
DATETRUNC('quarter', #August 15, 2005#) = July 1, 2005
DATETRUNC('month', #April 15, 2007#) = April 1, 2007
```

### **DAY(date)**

Returns the day of the given date as an integer.

```
DAY(#April 12, 2005#) = 12
```

### **ISDATE(string)**

Returns true if a given string is a valid date. For example,

```
ISDATE("April 15, 2004") = true.
```

### **MAX(expression) or MAX(expr1, expr2)**

Usually applied to numbers but also works on dates. Returns the maximum of a and b (a and b must be of the same type). Returns Null if either argument is Null. For example,

```
MAX(#January 1, 2004#, #March 1, 2004#) = #March 1, 2004#
MAX([ShipDate1], [ShipDate2])
```

**MIN(expression) or MIN(expr1, expr2)**

Usually applied to numbers but also works on dates. Returns the minimum of a and b (a and b must be of the same type). Returns Null if either argument is Null. For example,

MIN(#January 1, 2004# ,#March 1, 2004#) = #January 1, 2004#

MIN([ShipDate1], [ShipDate2])

**MONTH(date)**

Returns the month of the given date as an integer. For example,

MONTH(#April 12, 2005#) = 4

**NOW( )**

Returns the current date and time. For example,

NOW( ) = “5/10/2006 1:08:21 PM”

**TODAY( )**

Returns the current date. For example,

TODAY( ) = “5/10/2006”

**YEAR (date)**

Returns the year of the given date as an integer. For example,

YEAR(#April 12, 2005#) = 2005

**Type Conversion**

The result of any expression in a calculation can be converted to a specific data type. The conversion functions are STR(), DATE(), INT(), and FLOAT(). For example, if you want to cast a floating point number like 3.14 as an integer, you could write INT(3.14). The result would be 3, which is an integer. The casting functions are described below.

---

**Note** A boolean can be cast to an integer, float, or string. It cannot be cast to a date. True is 1, 1.0, or “1”, while False is 0, 0.0 or “0”. Unknown maps to Null.

---

**DATE(expression)**

Returns a date given a number, string, or date expression. For example,

DATE([Employee Start Date])

DATE("April 15, 2004") = #April 15, 2004#

DATE("4/15/2004")

DATE(#2006-06-15 14:52#) = #2006-06-15#

Note that the quotation marks are required in the last two examples.

**DATETIME(expression)**

Returns a datetime given a number, string, or date expression. For example,

DATETIME("April 15, 2005 07:59:00") = April 15, 2005 07:59:00

**FLOAT(expression)**

Casts its argument as a floating point number. For example,

FLOAT(3) = 3.000

FLOAT([Age]) converts every value in the Age field to a floating point number.

**INT(expression)**

Casts its argument as an integer. For expressions, this function truncates the result to the closest integer. For example,

INT(1.0/3.0) = 0

INT(4.0/1.5) = 3

INT(0.50/1.0) = 0

INT(0.501/1.0) = 1

When a string is converted to an integer it is converted to a float first and then rounded.

**STR(expression)**

Casts its argument as a string. For example,

STR([Age]) takes all of the values in the measure called Age and converts them to strings.

## Logical Functions

**CASE *expression* WHEN *value1* THEN *return1* WHEN *value2* THEN *return2*...ELSE *default return* END**

The CASE statement is another method used to perform logical tests and return appropriate values. It is often easier to write than IIF or IF THEN ELSE statements. The CASE statement evaluates *expression* and compares it to a sequence of values, *value1*, *value2*, etc. and returns a result. The first value that matches expression will result in returning the corresponding return expression. If no match is found the default return expression will be used. If there is no *default return* and no values match, then Null is returned. For example,

```
CASE [Region] WHEN "West" THEN 1 WHEN "East" THEN 2 ELSE 3 END  
CASE LEFT(DATENAME('weekday',[Order Date]),3)  
WHEN "Sun" THEN 0  
WHEN "Mon" THEN 1  
WHEN "Tue" THEN 2  
WHEN "Wed" THEN 3  
WHEN "Thu" THEN 4  
WHEN "Fri" THEN 5  
WHEN "Sat" THEN 6  
END
```

While an IF statement is used to perform a sequence of arbitrary tests, a CASE statement is used to search for a match to an expression. A CASE statement can always be written as an IF statement, although the CASE statement will generally be more concise.

If you need to include numeric comparisons in your conditions, use a nested IF clause. The CASE function compares strings only. For instance, suppose you want to break the values of the Sales field into three custom categories: one for sales less than 200, one for sales between 200 and 300, and one for sales between 300 and 400. The formula would be:

```
IF  
[Sales] < 200 THEN "Low"  
ELSEIF [Sales] < 300 THEN "Medium"
```



```
ELSEIF [Sales] < 400 THEN "High"
ELSE "NULL"
END
```

---

**Note** Many times you can use an ad-hoc group to get the same results as a complicated case statement. Refer to “Ad-Hoc Groups” on page 12-62 to learn more.

---

### **IIF(test, then, else, [unknown])**

The IIF function is used to perform logical tests and return appropriate values. The first argument of an IIF function must be a boolean. A boolean can be a boolean field in the data source, or the result of a logical expression using operators (or a logical comparison of AND, OR, or NOT). If *test* evaluates to TRUE, then this function returns *then*. If *test* evaluates to FALSE, then this function returns *else*.

A boolean comparison may also yield the value UNKNOWN (neither TRUE nor FALSE), usually due to the presence of Null values in *test*. The final argument to IIF is returned in the event of an UNKNOWN result for the comparison. If this argument is left out, Null is returned. For example,

```
IIF(7>5, “Seven is greater than five”, “Seven is less than five”)
```

```
IIF([Cost]>[Budget Cost], “Over Budget”, “Under Budget”)
```

```
IIF([Budget Sales]!<0,[Sales]/[Budget Sales],0)
```

```
IIF(COGS>[Budget COGS], IIF(Sales>=[Budget Sales], “Over Cost Budget and  
Over Sales Budget”, “Over Cost Budget and Under Sales Budget”), “Under Cost  
Budget”)
```

### **IF test THEN value END / IF test THEN value ELSE else END**

The IF THEN ELSE function is used to perform logical tests and return appropriate values, but has a different format and slightly different semantics than the IIF statement. The IF THEN ELSE statement evaluates a sequence of *test* conditions and returns the *value* for the first condition that is true. If no condition is true, the *else* value is returned. Each test must be a boolean, which may either be a boolean field in the data source or the result of a logical expression. The final ELSE is optional but if it is not provided and there is no true *test* expression, then the function returns Null. All of the *value* expressions must be of the same type. For example,

```
IF [Cost]>[Budget Cost] THEN "Over Budget" ELSE "Under Budget" END
```

```
IF [Budget Sales]!=0 THEN [Sales]/[Budget Sales] END
```

### **IF *test1* THEN *value1* ELSEIF *test2* THEN *value2* ELSE *else* END**

There is no built-in limit to the number of ELSEIF *test* THEN *value*'s in an IF expression, however, individual databases may impose a limit on IF statement complexity. While an IF statement can be rewritten as a series of nested IIF statements, there are differences in how the expressions will be evaluated. In particular, an IIF statement distinguishes TRUE, FALSE and UNKNOWN, while an IF statement only worries about TRUE and not true (which includes both FALSE and UNKNOWN). For example,

```
IF [Region]="West" THEN 1 ELSEIF [Region]="East" THEN 2 ELSE 3 END
```

### **ISDATE(string)**

Returns TRUE if the string argument can be converted to a date. Otherwise it returns FALSE. For example,

```
ISDATE("January 1, 2003") = TRUE, ISDATE("Jan 1 2003") = TRUE
```

```
ISDATE("1/1/03") = TRUE
```

```
ISDATE("Janxx 1 2003") = FALSE
```

### **ISNULL(expression)**

Returns TRUE if the expression is Null. Otherwise, returns FALSE.

### **MIN(expression) or MIN(expression1,expression2)**

Returns the minimum of an expression across all records or the minimum of two expressions for each record.

## **Aggregate Functions**

### **AVG(expression)**

Returns the average of all the values in the expression. AVG can be used with numeric fields only. Null values are ignored.

### **COUNT(expression)**

Returns the number of items in a group. Null values are not counted.

**COUNTD(expression)**

Returns the number of distinct items in a group. Null values are not counted. This function is not available if you are connected to MS Excel, MS Access, or a text file. You can extract your data into an extract file to gain this functionality. Refer to “Extracting Data” on page 6-1.

**MAX(expression)**

Returns the maximum of an expression across all records. If the expression is a string value, this function returns the last value where last is defined by alphabetical order.

**MIN(expression)**

Returns the minimum of an expression across all records. If the expression is a string value, this function returns the first value where first is defined by alphabetical order.

**STDEV(expression)**

Returns the statistical standard deviation of all values in the given expression based on a sample of the population.

**STDEVP(expression)**

Returns the statistical standard deviation of all values in the given expression based on a biased population.

**SUM(expression)**

Returns the sum of all values in the expression. SUM can be used with numeric fields only. Null values are ignored.

**VAR(expression)**

Returns the statistical variance of all values in the given expression based on a sample of the population.

**VARP(expression)**

Returns the statistical variance of all values in the given expression on the entire population.

## Pass Through Functions (RAWSQL)

The pass-through functions can be used to send SQL expressions directly to the database without being interpreted by Tableau. If you have custom database functions that Tableau doesn't know about you can use the pass-through functions to call these custom functions.

Your database usually will not understand the field names that are shown in Tableau. Because Tableau does not interpret the SQL expressions you include in the pass-through functions, using the Tableau field names in your expression may cause errors. You can use a substitution syntax to insert the correct field name or expression for a Tableau calculation into pass through SQL. For example, imagine you have a function that computes the median of a set of values. You could call that function on the Tableau column [Sales] like this:

```
RAWSQLAGG_REAL("MEDIAN(%1)", [Sales])
```

In addition, because Tableau does not interpret the expression, you must define the aggregation. Use the RAWSQLAGG functions when you are using aggregated expressions.

### **RAWSQL\_BOOL("sql\_expr", [arg1], ...[argN])**

Returns a Boolean result from a given SQL expression. The SQL expression is passed directly to the underlying database. Use %n in the SQL expression as a substitution syntax for database values.

In the example, %1 is equal to [Sales] and %2 is equal to [Profit].

```
RAWSQL_BOOL("IIF( %1 > %2, True, False)", [Sales], [Profit])
```

### **RAWSQL\_DATE("sql\_expr", [arg1], ...[argN])**

Returns a Date result from a given SQL expression. The SQL expression is passed directly to the underlying database. Use %n in the SQL expression as a substitution syntax for database values. In this example, %1 is equal to [Order Date].

```
Example: RAWSQL_DATE("%1", [Order Date])
```

### **RAWSQL\_DATETIME("sql\_expr", [arg1], ...[argN])**

Returns a Date and Time result from a given SQL expression. The SQL expression is passed directly to the underlying database. Use %n in the SQL expression as a substitution syntax for database values. In this example, %1 is equal to [Delivery Date].

```
Example: RAWSQL_DATETIME("%1", [Delivery Date])
```

**RAWSQL\_INT(“sql\_expr”, [arg1], ...[argN])**

Returns an integer result from a given SQL expression. The SQL expression is passed directly to the underlying database. Use %n in the SQL expression as a substitution syntax for database values. In this example, %1 is equal to [Sales].

Example: RAWSQL\_INT(“500 + %1”, [Sales])

**RAWSQL\_REAL(“sql\_expr”, [arg1], ...[argN])**

Returns a numeric result from a given SQL expression that is passed directly to the underlying data source. Use %n in the SQL expression as a substitution syntax for database values. In this example, %1 is equal to [Sales]

Example: RAWSQL\_REAL(“-123.98 \* %1”, [Sales])

**RAWSQL\_STR(“sql\_expr”, [arg1], ...[argN])**

Returns a string from a given SQL expression that is passed directly to the underlying data source. Use %n in the SQL expression as a substitution syntax for database values. In this example, %1 is equal to [Customer Name].

Example: RAWSQL\_STR(“%1”, [Customer Name])

**RAWSQLAGG\_BOOL(“sql\_expr”, [arg1], ...[argN])**

Returns a Boolean result from a given aggregate SQL expression. The SQL expression is passed directly to the underlying database. Use %n in the SQL expression as a substitution syntax for database values.

In the example, %1 is equal to [Sales] and %2 is equal to [Profit].

Example: RAWSQLAGG\_BOOL(“SUM( %1 ) > SUM( %2 )”, [Sales], [Profit])

**RAWSQLAGG\_DATE(“sql\_expr”, [arg1], ...[argN])**

Returns a Date result from a given aggregate SQL expression. The SQL expression is passed directly to the underlying database. Use %n in the SQL expression as a substitution syntax for database values. In this example, %1 is equal to [Order Date].

Example: RAWSQLAGG\_DATE(“MAX(%1)”, [Order Date])

**RAWSQLAGG\_DATETIME(“sql\_expr”, [arg1], ...[argN])**

Returns a Date and Time result from a given aggregate SQL expression. The SQL expression is passed directly to the underlying database. Use %n in the SQL expression as a substitution syntax for database values. In this example, %1 is equal to [Delivery Date].

Example: `RAWSQLAGG_DATETIME("%1", [Delivery Date])`

### **RAWSQLAGG\_INT("sql\_expr", arg1, ...argN)**

Returns an integer result from a given aggregate SQL expression. The SQL expression is passed directly to the underlying database. Use %n in the SQL expression as a substitution syntax for database values. In this example, %1 is equal to [Sales].

Example: `RAWSQLAGG_INT("500 + SUM(%1)", [Sales])`

### **RAWSQLAGG\_REAL("sql\_expr", arg1, ...argN)**

Returns a numeric result from a given aggregate SQL expression that is passed directly to the underlying data source. Use %n in the SQL expression as a substitution syntax for database values. In this example, %1 is equal to [Sales]

Example: `RAWSQLAGG_REAL("SUM( %1)", [Sales])`

### **RAWSQLAGG\_STR("sql\_expr", arg1, ...argN)**

Returns a string from a given aggregate SQL expression that is passed directly to the underlying data source. Use %n in the SQL expression as a substitution syntax for database values. In this example, %1 is equal to [Customer Name].

Example: `RAWSQLAGG_STR("AVG(%1)", [Discount])`

## **Special Functions**

The special functions can be used to create user filters that are based on user lists in your data source. For example, say you have a view that shows the sales performance for each employee. When you publish that to the server you may want to only allow employees to see their own sales numbers. You can use the special function `CURRENTUSER` to create a field that returns True if the user name of the person logged into the server is the same as the employee name in the view. That way when you filter the view using this calculated field, only the data for the user that is currently logged in is shown.

### **CURRENTUSER()**

Returns the name of the person currently using Tableau. This is the Tableau Server user name if user is logged in otherwise this function returns the Windows user name. Use this function to create calculations that are dependent on the current user.

Example: `[Manager]==CURRENTUSER()`

If the manager Dave Hallsten was logged in, this function would only return true if the Manager field in the view is also equal to Dave Hallsten. When used as a filter this calculated field can be used to create a user filter that only shows data that is relevant to the person logged into the server.

**ISCURRENTUSER(string)**

Returns True if the username of the person currently using Tableau matches the given string. The username for the person currently using Tableau is the Tableau Server username if they are logged in, otherwise it's the Windows username.

Example: ISCURRENTUSER("Dave Hallsten")

This returns true if Dave Hallsten is the current user, otherwise it returns false.

**ISMEMBEROF(string)**

Returns true if the person currently using Tableau is a member of a group that matches the given string. If the person currently using Tableau is logged in, the group membership is determined by groups on Tableau Server. If the person is not logged in, this function returns false.

Example: IF ISMEMBEROF("Sales") THEN "Sales" ELSE "Other" END

## Operators

To create calculated fields and formulas, you need to understand the operators supported by Tableau. This section discusses the basic operators that are available and the order (precedence) of operations.

### + (addition)

This means addition when applied to numbers and concatenation when applied to strings. When applied to dates, it can be used to add a number of days to a date. For example,

`7 + 3`

`Profit + Sales`

`'abc' + 'def' = 'abcdef'`

`#April 15, 2004# + 15 = #April 30, 2004#`

### – (subtraction)

This means subtraction when applied to numbers and negation if applied to an expression. When applied to dates, it can be used to subtract a number of days from a date. Hence it can also be used to calculate the difference in days between two dates. For example,

`7 - 3`

`Profit - Sales`

`-(7+3) = -10`

`#April 16, 2004# - 15 = #April 1, 2004#`

`#April 15, 2004# - #April 8, 2004# = 7`

### \* (multiplication)

This means numeric multiplication. For example,  $5 * 4 = 20$ .

### / (division)

This means numeric division. For example,  $20 / 4 = 5$ .

### % (modulo)

This calculates a numeric remainder. For example,  $5 \% 4 = 1$ .



**=, =, >, <, >=, <=, !=, <> (comparisons)**

These are the basic comparison operators that can be used in expressions. Their meanings are as follows: = or =(equal to), > (greater than), < (less than), >= (greater than or equal to), <= (less than or equal to), != and <> (not equal to).

Each operator compares two numbers, dates, or strings and returns a boolean (TRUE or FALSE). Booleans themselves, however, cannot be compared using these operators. For example, TRUE=TRUE is not a valid expression. To compare booleans in this way, use the logical operators AND and OR. For example, TRUE AND TRUE is a valid expression.

**^ (power)**

This symbol is equivalent to the POWER function. It raises a number to the specified power.

For example:

$$6^3 = 216$$

**AND**

This is a logical operator. An expression or a boolean must appear on either side of it. For example,

IIF(Profit=100 AND Sales=1000, "High", "Low")

If both expressions are TRUE (i.e., not FALSE and not UNKNOWN), then the result is TRUE. If either expression is UNKNOWN, then the result is UNKNOWN. In all other cases, the result is FALSE.

If you create a calculation in which the result of an AND comparison is displayed on a worksheet, Tableau displays TRUE and FALSE. If you would like to change this, use the Format area in the format dialog.

**OR**

This is a logical operator. An expression or a boolean must appear on either side of it. For example,

IIF(Profit=100 OR Sales=1000, "High", "Low")

If either expression is TRUE, then the result is TRUE. If both expressions are FALSE, then the result is FALSE. If both expressions are UNKNOWN, then the result is UNKNOWN.

If you create a calculation in which the result of an OR comparison is displayed on a worksheet, Tableau displays TRUE and FALSE. If you would like to change this, use the

Format area in the format dialog. The OR operator employs "short circuit evaluation." This means that if the first expression is evaluated to be TRUE, then the second expression is not evaluated at all. This can be helpful if the second expression results in an error when the first expression is TRUE, because the second expression in this case is never evaluated.

## NOT

This is a logical operator. It can be used to negate another boolean or an expression. For example,

```
IIF(NOT(Sales = Profit),"Not Equal","Equal")
```

## Precedence

All operators are evaluated in a specific order. For example,  $2*1+2$  is equal to 4 and not equal to 6. The reason is that the  $*$  operator is always evaluated before the  $+$  operator.

The following table shows the order in which operators are evaluated. The first line has the highest precedence. Operators on the same line have the same precedence. If two operators have the same precedence they are evaluated from left to right in the formula.

Precedence	Operator
1	– (negate)
2	^ (power)
32	*, /, %
4	+, –
5	==, >, <, >=, <=, !=
6	NOT
7	AND
8	OR

Parentheses can be used as needed. Operators that appear within parentheses are evaluated before those outside of parentheses, starting from the innermost parentheses and moving outward. For example,  $(1 + (2*2+1)*(3*6/3)) = 31$ .

## Performance Tips

---

---

## Overview

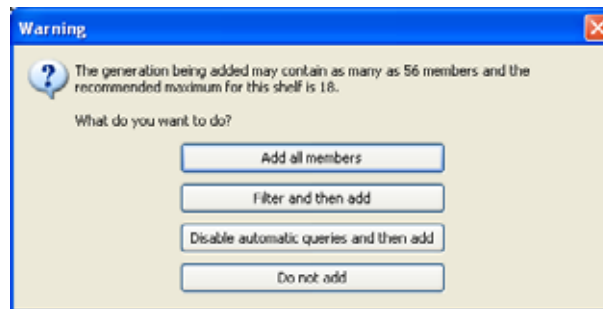
This section includes tips that will help improve Tableau’s performance when you use large data sources. One of the important things to understand is that Tableau is only as fast as your data source. So if your data source responds slowly to queries, then Tableau must wait for the data source before displaying results. As a result, a “best practice” suggestion is to use Tableau with databases that are suitable for real-time querying and analysis. This section discusses performance tips for the following categories:

- **All Data Sources**- general performance tips that can be applied to all data sources.
- **Relational Data Sources**- performance tips specific to relational data sources.
- **Multidimensional Data Sources**- performance tips specific to multidimensional data sources.
- **Speeding up Context Filters**- tips on increasing the performance of context filters.
- **Extracting Large Text Files**- tips for working with large text files.


## Performance Tips: All Data Sources

The following performance tips apply to all data sources supported by Tableau.

- **Turn off Automatic Updates**– When you place a field on a shelf, Tableau generates the view by automatically querying the data source. If you are creating a dense data view, the queries might be time consuming and significantly degrade system performance. In this case, you can instruct Tableau to turn off queries while you build the view. You can then turn queries back on when you are ready to see the result. Refer to “Automatic Updates” on page 7-3 for more information.
- **Look for Warnings**– Tableau displays a performance warning dialog box when you attempt to place a large dimension (with many members) on any shelf. The dialog box provides four choices as shown in the figure below. If you choose to add all members, then your might experience a significant degradation in performance.



- **Avoid Generating Too Many Panes** – If you attempt to create many panes in a table, Tableau will display the warning dialog box shown below. In most cases, you should not display more than the recommended number of panes because such a view is not very useful.



## Performance Tips: Relational Data Sources

The following performance tips apply to relational data sources.

- **Context filters** – If you are setting filters that significantly reduce the data set size, and that will be used for more than several data views, you should set those filters as context filters. Refer to “Context Filters” on page 12-55 for how to create context filters. For more information about performance improvement with context filters refer to “Speeding up Context Filters” on page 27-8.
- **Aggregate measures** – If the views you create are slow, make sure you are working with aggregated measures rather than disaggregated measures. When views are slow it usually means you are trying to view many rows of data at once. You can reduce the number of rows by aggregating the data. In other words, make sure the **Aggregate Measures** option on the **Analysis** menu is selected.
- **Sets** – If you want to filter a dimension to remove members based on a range of measure values, you should create a set rather than using a quantitative filter. For instance, you can create a set that only returns the Top 50 items in a dimension, rather than all of the items in a dimension.

When creating a group from a selection as described in “Create a Set by Selecting Marks” on page 12-73, make sure you've included only the columns of interest. Each additional column included in the set will result in decreased performance.

- **Extract Large Text and Excel Files** – If your data source is large text or Excel file, you should create a Tableau Extract to improve performance and gain new functionality. Note that if you connect Tableau to a large text file, you will be prompted to extract the data if the file is considered to be too large to perform well. Refer to “Extracting Large Text and Excel Files” on page 27-9.
- **Use a database server** – You should consider storing your data in a database server like Microsoft SQL Server. The Professional Edition of Tableau can connect to these larger database servers.
- **Create indexes for tables** – Index the tables in your relational database. To successfully index your data set, you should identify the fields that you frequently filter on and add them to the index. If you have a field that you use as a context filter often, consider setting it as your primary index. If you are working with Access tables that have more than 200,000 rows of data, consider setting indexes on the tables. You can learn how to do this by searching for “index” in the Access online help. Access allows you to store 2 GB of data (approximately 1-2 million rows) in a database, but it performs poorly well below this limit.

- **Break up your data** – If you have a lot of data, consider breaking it into smaller pieces. For example, you can create a cluster of Access tables that address specific subsets of your data.
- **Add filters first** - If you are working with a large data source and have automatic updates turned off, it is possible to create a really slow query when adding filters to the view. Rather than build the view and then specify filters, you should first specify the filters and then drag fields to the view. That way, when you run the update or turn on automatic updates, the filters will be evaluated first.



## Performance Tips: Multidimensional Data Sources

The following performance tips apply to multidimensional data sources.

- **Filtering** – If your cube has a single large dimension, you should set a filter directly on that dimension rather than setting a filter on another dimension or measure. For example, suppose you want to reduce the numbers of products being displayed in a data view. It is much more efficient to set the filter directly on **Products** or to create a computed set based on **Products** (such as Top 10) rather than filtering other fields such as **Location** or **Profit**.

Also, you should avoid selecting large numbers of members from a large dimension. When a dimension is large, it is best to keep the size of the filter to less than a thousand members.

- **Sets** – When creating a set from a selection as described in “Create a Set by Selecting Marks” on page 12-73, make sure you've included only the columns of interest in the Create Set from Selection dialog box. Each additional column included in the group will result in decreased performance. For instance, if you create a set that contains all regions with sales between 8000 and 15000 but you include a column that does not affect the members of the set, you may notice a performance decrease. You can remove extra columns by right-clicking the column and selecting **Remove This Column** from the context menu.
- **Sorting** – Avoid applying sorts to levels within a very large hierarchy.
- **Large Root Levels** – If you are working with a dimension whose root level is greater than 1000 but it is not too large (greater than 100,000) you should avoid using the filter dialog to filter the data. Instead, drag the dimension to a shelf and use the **Exclude** command in the headers context menus to limit the data that is displayed in the view. In this particular case, dragging a dimension with this size root level to the filter shelf may cause a long query.

## Speeding up Context Filters

To improve performance of context filters, especially on large data sources, follow these general rules.

- Using a single context filter that significantly reduces the size of the data set is much better than applying many context filters. In fact, if a filter does not reduce the size of the data set by one-tenth or more, it is actually worse to add it to the context because of the performance cost of computing the context.
- If you do have multiple filters to add to the context it is better to create all of the filters first and then create a context that includes them all. To create a context that includes them all, select **Analysis > Set Context** and then add the multiple filters to the context all at once. Using the standard **Add to Context** command in the context menus of each filter will force Tableau to compute the context once per filter which can degrade performance.
- Complete all of your data modeling **before** creating a context. Changes in the data model such as converting dimensions to measures, require recomputing the context.
- Set the necessary filters for the context and create the context **before** adding fields to other shelves. Doing this work first makes the queries that are run when you drop fields on other shelves much faster.
- If you want to set a context filter on a date you can use a continuous date. However, using date bins like `YEAR(date)` or context filters on discrete dates are very effective.

---

**Note** If your data set is heavily indexed, context filters may not provide performance improvement and may actually cause slower query performance.

---

## Extracting Large Text and Excel Files

When you connect to large text or Excel files in Tableau, you may experience poor performance because queries are slow. If your file has more than 20,000 rows you should extract the data using the Extract feature in Tableau. If Tableau determines your file to be too large (greater than 10 MB) you are prompted to extract the data. When you extract the data you will not only improve performance but also enable more functions such as count distinct.

### To extract a text or Excel file into an Tableau extract file:

- 1 Create a new connection and select the text or Excel file you want to connect to. For information about connecting to a text or Excel file refer to “Examples – Connecting to Data Sources” on page 4-11.
- 2 When you click **OK** you are prompted to create an extract file. You can also select Data > Extract. You can specify optional filters to define a subset of data or click **Extract** to include all the data. Refer to “Creating an Extract” on page 6-3 to learn more about defining an extract.
- 3 When you click **Extract**, the data from the text or Excel file is converted and stored with the name you specify.

## Tips and Tricks

---

---

## Overview

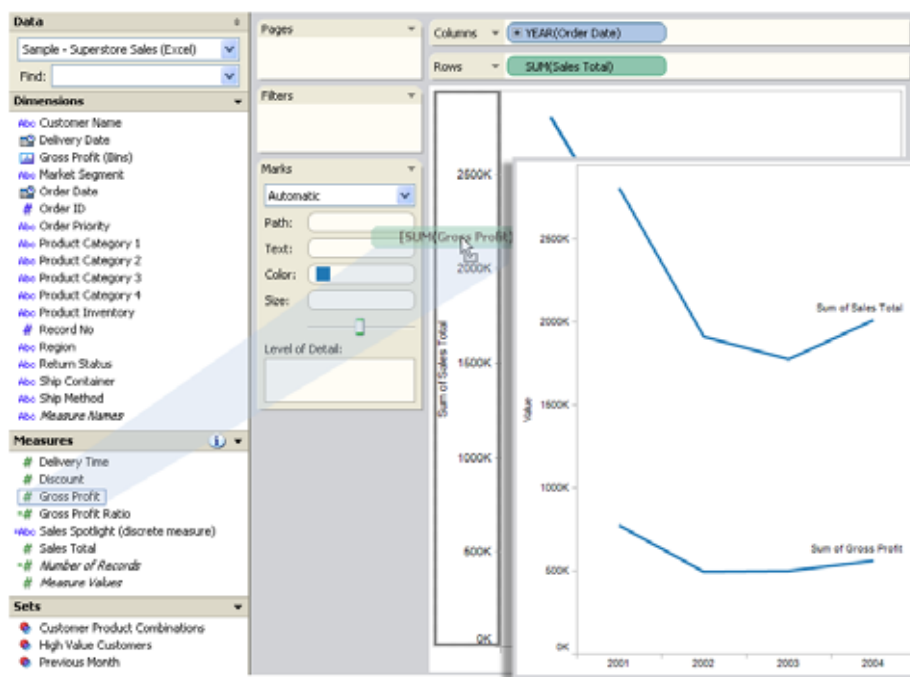
This section provides a collection of tips to help you use Tableau. You can learn more about all of these topics by searching the Tableau Help system.

- Show Multiple Measures on a Single Axis
- Convert Measures to Dimensions
- Convert Dimensions to Measures
- Duplicate Any View as a Cross-tab
- Manually Sort Items in the Color Legend
- Duplicate Fields in the View
- Create Moving Date Filters
- Resize Individual Views in a Dashboard
- Use the Pages Shelf to Split Large PDF Reports
- Specify Page Setup for Multiple Worksheets
- Visual Cues for Fields

## Show Multiple Measures on a Single Axis

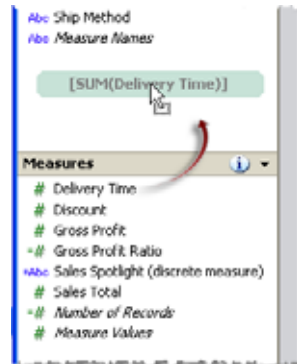
In addition to dragging measures onto shelves, you can drag measures onto the axis created by a previous measure. This allows multiple measures to share a single axis.

In the example below, both the Sales Total field and the Gross Profit field share a common axis. This view was created by dragging the Sales Total field to the Row Shelf and then dragging the Gross Profit field onto the axis created by the Sales Total field. The result is the creation of a “multiple measures” axis that is labeled ‘value.’ For a complete description refer to “Measure Values and Measure Names” on page 8-7.



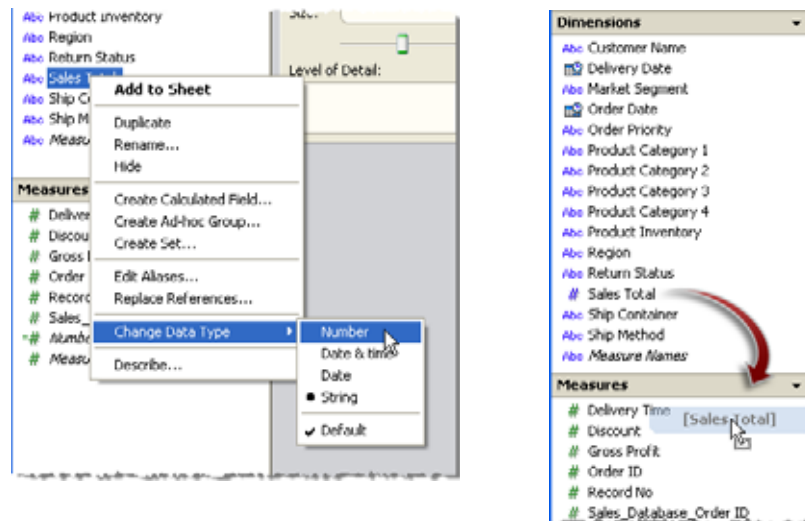
## Convert Measures to Dimensions

By default, Tableau treats all relational fields containing numbers as measures. However, you might decide that some of these fields should be treated as dimensions. For example, a field containing ages may be categorized as a measure by default in Tableau because it contains numeric data. However, if you want to look at each individual age rather than an axis you can convert the Age field to a dimension. Convert a measure to a dimension by dragging it from the Measures area to the Dimensions area in the Data window. For a complete description refer to “Converting Measures to Dimensions” on page 8-24.



## Convert Dimensions to Measures

By default, Tableau treats all relational fields containing text or date values as dimensions. However, you might decide that some of these fields should be treated as measures. For example, a field that contains sales figures with null values represented as “N/A” or “\*” may be categorized as a dimension by default in Tableau because it contains text values. You can convert the dimension to a measure first by changing the data type to a number and then by dragging it from the Dimensions area to the Measures area in the Data window.





## Duplicate Any View as a Cross-tab


You can duplicate any view in Tableau as a cross-tab by selecting **Edit > Duplicate as Cross-tab**. The result will contain only dimensions in the **Rows** and **Columns** shelves. If measure names are part of the original view, they will be displayed in the cross-tab using the **Measure Names** field.

You can also export a cross-tab directly to Microsoft Excel by selecting **File > Export > Cross-tab to Excel**. For more information about exporting as a cross-tab refer to “Export Cross-tab to Excel” on page 24-20.



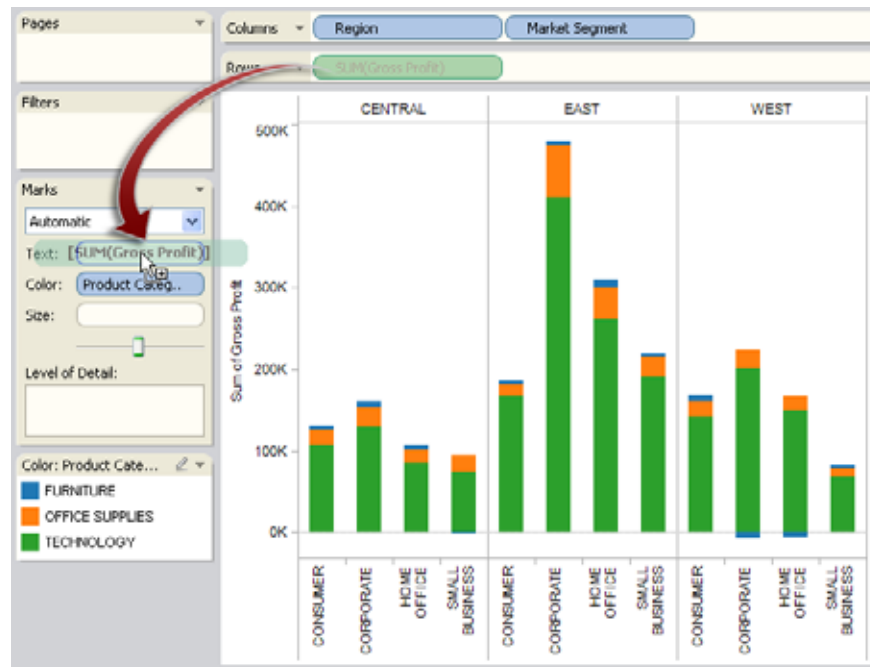
## Manually Sort Items in the Color Legend

Manually sorting the items in the color legend change the drawing order of a field so you can see obscured data in your views. For instance, if you can't see green marks in a scatter plot because they are obscured by orange marks, you can manually sort the items in the color legend so that the green points are drawn on top of the orange points. The marks are drawn in the view according to the order in the legend, from bottom to top. You can manually sort a color legend simply by dragging and dropping the items in the legend. Refer to “Example- Manually Sorting Drawing Order” on page 12-17 to learn more.



## Duplicate Fields in the View

Once you add fields to the view, you can duplicate them simply by hold down the Ctrl key on your keyboard while you drag it to another shelf. For example, in the view below, the SUM(Gross Profit) field is duplicated and placed on the Text shelf.



## Create Moving Date Filters


Tableau includes the `Today()` and `Now()` functions which help you create a date filter that updates over time. For example, you can create date filters such as Current Month, Previous Quarter, Past 7 Days, and more.

### To create a moving data filter:

1 Right-click the date field in the Data window and select **Create Set**.

2 Name the set and select **Use All** from the Values drop-down list.

The Use All selection will make is to that the filter is created using all values even as new dates are added to the data source.

3 Select the **Condition** tab in the Create Set dialog box, select **By Formula**, and click the  button to the right of the text box.

4 In the Calculated Field dialog box, type a formula that includes the dates you want to include in the view.

5 When finished click **OK** twice to save and close the dialog boxes.

Below is a table of formulas for common moving date filters. To learn more about date functions refer to “Date Functions” on page 26-13.

Date Filter	Formula*
Current Month	<code>datetrunc('month',[Date])=datetrunc('month',today())</code>
Current Quarter	<code>datetrunc('quarter',[Date])=datetrunc('quarter',today())</code>
Past 7 Days	<code>[Date]&lt; today( ) AND [Date]&gt;=today( )-7</code>
Previous Quarter	<code>datepart('quarter',[Date]) = datepart('quarter',today())-1</code>
Year to Date	<code>datepart('day',[Date])&lt;datepart('day',today()) AND datepart('year',[Date])= datepart('year',today())</code>

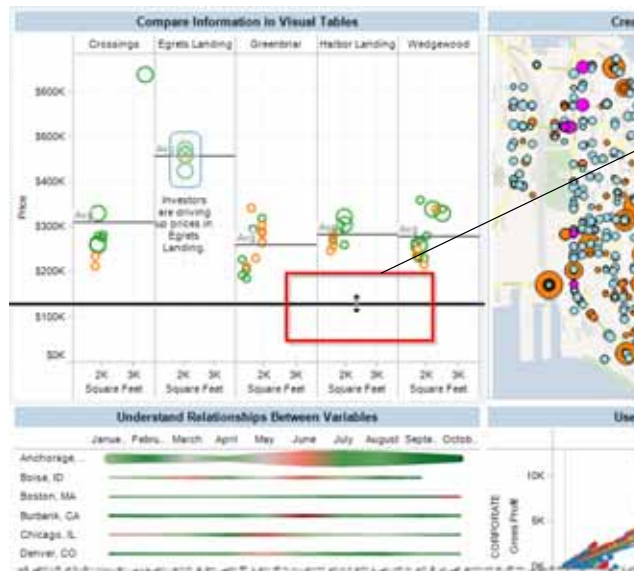
\*Replace the `[Date]` field in the formulas above with the name of your date field.

## Resize Individual Views in a Dashboard

When you have several views in a dashboard, you can click and drag the divider lines to resize each area of the dashboard. Typically, dragging a divider line resizes the entire row or column area, as shown below:




However, you can hold down the Shift key on your keyboard while dragging a divider line, to resize an individual view.




Hold down the Shift key while dragging the divider resizes an individual view. The views on the right are not resized.

## Use the Pages Shelf to Split Large PDF Reports

When you publish a large view to a PDF, the entire view is either scaled to fit on a single page or the page size is scaled to fit the view. Either way, the view is not split well across several printable pages. You can use the pages shelf to split the report into meaningful sections that are then split into a multiple page PDF report. For example, the view below show the total monthly expenses by Product and Market. When published to a PDF, the page is really long.



If you move the Market field to the Pages shelf, the published PDF fits on a standard letter sized paper and contains four pages — one for each market.




## Specify Page Setup for Multiple Worksheets

Before printing a single worksheet or several worksheets, you can specify how you want the document to look when printed using the Page Setup options. For example, you can specify whether to include the legends, titles, captions, margins, page orientation, and more. Each worksheet can have its own page setup. However, often you will want to specify global page setup options. Hold the Ctrl key to select one or more worksheets and then set the page setup options.

**To specify page setup for multiple worksheets:**

- 1 Hold the Ctrl key on your keyboard and select one or more worksheets that you want to specify the page setup for.



- 2 Select **File > Page Setup**.
- 3 Make the necessary changes in the Page setup dialog box and click **OK**. Refer to “Page Setup” on page 25-3 to learn more about each option in this dialog box.







## Visual Cues for Fields























Tableau provides many visual cues to help you evaluate the type of data that's displayed in the Data window, and the state of a data view. This section discusses the visual cues used for:





- Fields in the Data window
- Fields on Shelves

### Fields in the Data window

The following table explains each of the icons displayed in the Data window. Each icon in the table can be modified by one of four indicators.




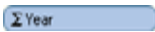



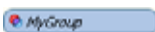



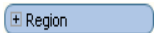

-  Blue icons indicate that the field is discrete.
-  Green icons indicate that the field is continuous.
-  Icons preceded by the equal sign (=) indicate that the field is a user-defined calculation or a copy of another field
-  Icons with an exclamation mark on them indicate that the field is invalid.

Visual Cue	Description
  	The field contains text values.
   	The field contains numeric values.
   	The field contains only date values.
   	The field contains both date and time values.
 	The field contains boolean (true or false) values.
 	The field is a user-defined set.
	The field is a server named set.
	The field is a numeric bin.
	The field is an ad-hoc group.

Visual Cue	Description
	The field is an attribute of a cube.
	The field is a slowly changing attribute of a cube.
	The field is a level in a hierarchy. Levels greater than five are shown without numbers.
	The field is a hierarchy in a dimension that contains multiple hierarchies.

## Fields on Shelves

Fields placed on shelves use a combination of icons, colors, and text styles as visual cues.

Visual Cue	Description
	A blue field on a shelf indicates a discrete field.
	A green field on a shelf indicates a continuous field.
	A bold name indicates a sorted field.
	The  icon indicates a calculation filter (slicer).
	The  icon indicates a set.
	An italicized name indicates a filtered set.
	A grey field with the  icon and indicates a context filter.
	The delta icon indicates that the field is a table calculation.
	The plus and minus controls appear when the field is part of a hierarchy that you can drill down and up in.
	

## Keyboard Shortcuts

---

---

Tableau keyboard shortcuts are listed below.

Keyboard Shortcut	Description
Ctrl+A	Select all data
Ctrl+B	Smaller table
Ctrl+Shift+B	Larger table
Ctrl+C	Copy selected data
Ctrl+Alt+C	Place selected field on columns shelf
Ctrl+D	Connect to data source
Ctrl+E	Describe Sheet
Ctrl+F	Makes the find command in the Data window active
Ctrl+Alt+F	Place selected field on filters shelf
Ctrl+H	Switch in and out of Presentation Mode
Ctrl+Atl+I	Place selected field on size shelf
Ctrl+L	Flip orientation of column labels
Ctrl+Alt+L	Place selected field on level of detail shelf
Ctrl+M	New worksheet
Ctrl+N	New workbook
Ctrl+O	Open file
Ctrl+Alt+O	Place selected field on color shelf
Ctrl+P	Print
Ctrl+Alt+P	Place selected field on pages shelf

---

Keyboard Shortcut	Description
Ctrl+Alt+R	Place selected field on rows shelf
Ctrl+S	Save file
Ctrl+Alt+S	Place selected field on shape shelf
Ctrl+Alt+T	Place selected field on text shelf
Ctrl+V	Paste clipboard
Ctrl+W	Swap rows and columns
Ctrl+X	Cut selection
Ctrl+Alt+X	Place selected field on rows shelf
Ctrl+Y	Redo
Ctrl+Alt+Y	Place selected field on columns shelf
Ctrl+Z, Backspace	Undo
Ctrl + Alt + Backspace	Clear the current worksheet.
Ctrl+(left arrow)	Make rows narrower
Ctrl+(right arrow)	Make rows wider
Ctrl+(down arrow)	Make columns shorter
Ctrl+(up arrow)	Make columns taller
Ctrl+1, Ctrl+Shift+1, Ctrl+!	Show Me!
ENTER	Add the selected field to the sheet
F1	Opens the Help
Ctrl+F4	Deletes the current worksheet or hides the worksheet if it is used in a dashboard
Alt+F4	Closes the current workbook

---

Keyboard Shortcut	Description
F4	Starts and stops forward playback on the pages shelf.
Shift + F4	Starts and stops backward playback on the pages shelf.
F5	Refreshes the data source.
Ctrl + .	Skip forward one page.
Ctrl + ,	Skip backward one page.
Ctrl+Tab, Ctrl+F6	Cycle backward through open worksheets
Ctrl+Shift+Tab, Ctrl+Shift+F6	Cycle forward through open worksheets
Shift+F6	Select mode
Shift+F7	Pan mode
Shift+F8	Focus mode
Shift+F9	Zoom mode
Ctrl + Click	Zooms out while in zoom mode
F9	Run Query
F10	Toggles Automatic Updates on and off
F12	Reverts workbook to last saved state.



# Glossary

---

---

## **Ad-hoc Groups**

Ad-hoc groups are dimension members that have been combined into higher level categories. For example, a dimension that contains states could be grouped into regions using Ad-hoc Groups. Ad-hoc groups are marked with a paper clip icon in the Data window.

## **Aggregation**

An aggregation results from a mathematical operation applied to a measure. Predefined aggregations include summation, average, and so on. Dimensions can be converted to measures by aggregating them as a count.

For relational data sources, all measures must be either aggregated or disaggregated (unless they appear on the **Filters** shelf). Tableau automatically aggregates measures, usually as a summation, when they are placed on a shelf. For multidimensional (OLAP) data sources, aggregations are defined when the cube is created and cannot be modified in Tableau.

## **Alias**

An alias is an alternative name assigned to a field or to a dimension member. Tableau supports both field aliases and member aliases.

## **Axis**

An axis is displayed in a table when you place a continuous field on the Rows or Columns shelf. The axis labels are given by the name of the measure.

## **Bookmarks**

Bookmarks contain the data view from a single worksheet. You can create and display bookmarks using the **Bookmark** menu.

Bookmarks behave like web browser bookmarks. They can be accessed without opening any other document and are a convenient way to quickly display different analyses. You should save bookmarks in the bookmarks folder of the Tableau repository.

## **Caption**

A description of the current view on the active worksheet. For example, “Sum of Sales for each Market”. Captions can be automatically generated or custom. Show and hide the caption by selecting **View > Caption**.



---

## **Cell**

Any table you create in Tableau has the cell as its basic element. Controlling cells to enhance your data view is useful for text tables and heat maps.

## **Color Legend**

The color legend displays the colors associated with a measure or dimension members. The default legend is modified when you place a dimension or a measure on the **Color** shelf.

## **Color Shelf**

The **Color** shelf allows you to encode data by assigning different colors to the marks in a data view. The shelf accepts a measures or a dimension.

When you place a dimension on the **Color** shelf, Tableau separates the marks according to the members in the dimension, and assigns a unique color to each member. When you place a measure on the **Color** shelf, Tableau draws each mark with a different color using a continuous range. In both cases, a legend describes the color encoding.

## **Columns Shelf**

The **Columns** shelf allows you to create the columns of a data table. The shelf accepts any number of dimensions and measures.

When you place a dimension on the **Columns** shelf, headers for the members of that dimension are created. When you place a measure on the **Columns** shelf, quantitative axes for that measure are created.

## **Cross-tab**

In Tableau, a cross-tab is another name for a text table. Text tables provide an easy way to display the numbers associated with dimension members.

## **Dashboard**

A dashboard is a collection of several worksheets shown in a single location where you can compare and monitor a variety of data simultaneously.

## **Data View / View**

A data view, also referred to simply as a view, is a representation of your data in a Tableau worksheet or dashboard. You can create data views by placing fields on shelves.

---

## **Data window**

The Data window displays the fields of the data sources to which Tableau is connected. The fields are divided into dimensions and measures. The Data window also displays custom fields such as calculations, binned fields, and groups.

You build views of your data by dragging fields from the Data window onto the various shelves that are a part of every worksheet.

## **Data Source**

To build data views, you must connect to a data source. Tableau supports many standard relational and multidimensional (OLAP) sources.

When connected, Tableau does not save a copy of your data. Instead, it saves information about where the data source is located. When you build data views, Tableau sends the appropriate queries to your data source. A convenient location to save data sources is the data sources folder of the Tableau repository.

## **Dimensions**

Dimensions are independent fields. Dimensions typically hold discrete data such as hierarchies and members that cannot be aggregated. Examples of dimensions include dates, customer names, and market segments. Tableau does, however, support continuous dimensions.

## **Encoding**

In Tableau, encoding refers to a particular visual representation of your data. You can encode your data by color, shape, size, and path using the associated worksheet shelves.

## **Fields**

Field is another name for a dimension or a measure. All databases contain fields. Fields contain data.

For relational data sources, fields are the columns of a table. For multidimensional (OLAP) data sources, fields are the dimensions of a cube. Each dimension or column contains a unique attribute of the data such as customer name, sales total, product type, and so on.

## **Field Label**

Field labels are titles that indicate the fields that are used in a view. For example, in a view that has rows for East, Central, and West might have a Region field label at the top of the column indicating that each row is a member of the Region field.

---

## **Filters Shelf**

The **Filters** shelf allows you to exclude data from a view. You can filter data using measures, dimensions, and both measures and dimensions at the same time.

You can filter data based on the fields that make up the columns and rows of the table. This is called an internal filter. You can also filter data based on fields that do not compose the table. This is called an external filter.

## **Format window**

The Format window is a pane that, when open, displays on the left side of the workbook. The Format window contains formatting settings that control the entire worksheet as well as individual fields in the view.

## **Headers**

Headers are displayed in a table when you place a dimension on the **Rows** or **Columns** shelf. The header labels are given by the dimension member names.

## **Level of Detail Shelf**

The **Level of Detail** shelf allows you to separate the marks in a data view according to the level of detail (members) of a dimension.

The shelf works only if your data are aggregated. If your data are disaggregated, then it isn't possible to separate the marks into additional levels of detail. Additionally, placing a measure on the shelf has no effect on the table structure because measures do not contain members.

## **Marks**

Marks visually represent one or more rows in a data source. Mark types can be a bar, line, square, and so on. You can control the type, color, and size of marks.

Tableau can automatically select a mark type, or you can manually select the mark type from the **Mark** menu.

## **Measures**

Measures are fields that are dependent variables. They are typically quantitative fields or calculated fields like sales, temperature or frequency. Discrete measures can also be created in Tableau.

---

## **Multidimensional Data Source**

In Tableau, a multidimensional (OLAP) data source can be a Microsoft Analysis Services cube or a Hyperion Essbase cube.

## **Pages Shelf**

The page shelf lets you split a view into a sequence of pages based on the members and values in a discrete or continuous field. Adding a field to the Page shelf is like adding a field to the Rows shelf except a new page is created for each new row.

## **Pane**

Tables consist of one or more panes. The number of panes in a view depends on the number and type of fields placed on the **Rows** and **Columns** shelves.

## **Path Shelf**

The **Path** shelf allows you to encode data by connecting marks using a particular drawing order. The shelf accepts measures and dimensions.

Dimensions connect the marks according to the members in the dimension. If the dimension is a date, the drawing order is given by the date order. If the dimension holds words, the line is drawn based on the order of the words in the data source. Measures connect the marks according to the values of the measure. The measure can be aggregated or disaggregated.

## **Query**

Tableau communicates with your databases with queries. Queries are questions that databases can understand and answer. Common query languages include SQL and MDX.

Every time you build a view of your data, Tableau translates your actions into queries and retrieves the requested information from the data source. If you are building a dense data view, you can turn queries off until all desired fields are placed on shelves.

## **Relational Data Source**

In Tableau, a relational data source can be an Excel workbook, an Access database, a comma delimited text file, a MySQL database, an Oracle database, a SQL Server database, a Firebird database, a PostgreSQL database, or a Tableau Data Extract file.

---

## Repository

The Tableau repository holds workbooks, bookmarks, and data sources. It is located in a folder called My Tableau Repository inside of your My Documents folder.

## Rows Shelf

The **Rows** shelf allows you to create the rows of a data table. The shelf accepts any number of dimensions and measures. When you place a dimension on the **Rows** shelf, headers for the members of that dimension are created. When you place a measure on the **Rows** shelf, quantitative axes for that measure are created.

## Set

A set is a custom field you create by filtering existing dimensions. They appear at the bottom of the Data window in the **Sets** area. The three main uses of a set are to create a subset of the data, apply a numerical or a top N filter, and to create unique encodings. You can use sets in data views just like any other dimension.

## Shape Legend

The shape legend displays the shapes associated with dimension members. The legend appears on worksheets that have a dimension placed on the **Shape** shelf.

## Shape Shelf

The **Shape** shelf allows you to encode data by assigning different shapes to the marks in a data view. The shelf accepts dimensions only.

When you place a dimension on the shelf, Tableau separates the marks according to the members of the dimension, and a legend describes the encoding. You cannot place a measure on the shelf because measures do not contain members.

## Shelves

You build views of your data by placing fields onto the shelves that are a part of every worksheet.

Some shelves are available only when certain mark types are selected. For example, the **Shape** shelf is available only with the Shape mark type.

## Size Shelf

The **Size** shelf allows you to encode data by assigning different sizes to the marks in a data view. The shelf accepts measures and dimensions.

---

When you place a dimension on the shelf, Tableau separates the marks according to the members in the dimension, and assigns a unique size to each member. When you place a measure on the shelf, Tableau assigns a different size to each mark using a continuous range.

### **Table**

The visual presentation of a data view is contained within a table. Tables consist of panes, headers, axes, and cells.

### **Text Shelf**

The **Text** shelf allows you to view the numbers associated with a data view, and to encode data by assigning text labels to the marks. The most common view using the **Text** shelf is a text table. The shelf accepts measures and dimensions.

### **Undo/Redo**

You can undo any action in Tableau by clicking **Undo** on the toolbar. Likewise, you can redo any action by clicking **Redo** on the toolbar.

Using **Undo** and **Redo**, you can quickly return to a previous view or you can browse all the views of a data source that you have created. The undo/redo history is not saved between Tableau sessions.

### **Workbooks / Packaged Workbooks**

Workbooks hold one or more worksheets and dashboards. By saving a workbook, you can save all open sheets in one file that can then be easily shared.

### **Worksheets**

Worksheets hold your data views. You can save individual worksheets as bookmarks.

Each worksheet can be connected to only one data source. However, different worksheets in a workbook can be connected to different data sources.

## A

- About Shelves 10-10
- Access database 4-11
- Accessing Data 4-6
- Accounts Dimension 4-32
- Add to sheet
  - double-click 11-8
  - drag to shelves 10-3
  - drag to worksheet 10-6
- Ad-hoc Groups 2-9, 12-62
  - creating 12-62
  - editing 12-66
  - finding members 12-68
- Aggregate Calculations 16-30
  - about 16-30
  - creating 16-31
  - disaggregated 16-32
- Aggregate Functions 26-19
  - AVG 26-19
  - COUNT 26-19
  - COUNTTD 26-20
  - MAX 26-20
  - MIN 26-20
  - STDEV 26-20
  - STDEVP 26-20
  - SUM 26-20
  - VAR 26-20
  - VARP 26-20
- Aggregation Labels
  - removing from measure names 8-17
- Aggregations 16-4, 16-8
  - all values 16-4
  - average 16-4
  - count 16-5
  - count distinct 16-6
  - grand totals 16-67
  - maximum 16-4
  - measures 16-8
  - minimum 16-4
  - percentages 16-75
  - standard deviation 16-5
  - standard deviation population 16-5
  - sum 16-4
  - variance 16-5
  - variance population 16-5
- Alerts
  - show/hide 9-13
- Aliases 8-11
  - changing measure names 8-17
  - editing data labels 21-20
  - fields 8-11
  - members 8-11
  - multidimensional 8-12
  - Relational 8-12
- Alignment
  - formatting 20-8
- All Values 16-4
- Analysis
  - aggregate measures 16-8
- Analysis Services 4-41
- Annotations 21-3
  - adding 21-3
  - formatting 21-11
  - rearranging 21-5
  - removing 21-15
- Aspect Ratio 20-35
- asymmetric layout 10-35
- Automatic mark type 9-12
- automatic updates 7-3
- Automatic Views 11-3
  - Double-Click 11-8
  - Show Me! 11-4
  - Show Me! Alternatives 11-6

Average 16-4

## Axes

- appearance 20-45
- creating in a table 10-10
- formatting 20-37
- range 20-37
- reverse 20-45
- scale 20-37
- tick marks 20-47
- title 20-45

## B

### Background Images 17-1

- adding 17-3
- aspect ratio 17-4
- can't see marks 20-19
- disabling 17-9
- editing 17-8
- enabling 17-9
- removing 17-11
- show/hide conditions 17-9

### Banding 20-9

## Bar

- equal size 16-88
- mark type 9-15

### Binned Data 16-59

### Bookmarks 24-7

- deleting 24-8
- removing 24-8

### Boolean data type 8-20

## Borders

- formatting 20-11
- marks 20-18

### Build Views Automatically

- double-click 11-8
- Show Me! 11-3

### Building Views

- automatically 11-4
- Hyperion Essbase 10-35
- Microsoft Analysis Services Cube 10-39

## C

### Calculated Fields

- creating 16-18

### Calculated Members

- creating 16-20

### Calculations

- aggregate calculations 16-30
- aggregations 16-4, 16-8
- binned data 16-59
- discrete measure 16-37
- grand totals 16-65
- local vs. server 16-63
- new fields 16-18
- percentages 16-74
- spotlighting 16-37
- subtotals 16-71
- user filters 23-40
- writing formulas 16-22

### Cancel Query

- abandoned queries 7-5
- reconnecting 7-5

### Cards 2-11

- caption 2-12
- card menu 2-13
- color legend 2-11
- columns 2-11
- current page 2-12
- filters 2-11
- marks 2-12
- measure names/measure values 2-11
- pages 2-11



- quick filters 2-12
- rows 2-11
- shape legend 2-11
- size legend 2-11
- summary 2-12, 14-20
- title 2-12
- Cell Size
  - clear 20-34
  - Narrower 20-27
  - shorter 20-27
  - square cell 20-31
  - taller 20-27
  - text cell 20-31
  - wider 20-27
- Chart Type
  - histogram 16-59
- Circle mark type 9-19
- Clear
  - manual sizing 20-34
- clear worksheet 2-18
- Close Connection 4-71
- closing sheets 2-20
- Collaborate
  - data sources 23-14
  - import data source 23-21
  - open workbooks 23-18
  - workbooks 23-3
- Color
  - encoding 20-49
  - full color range 20-53
  - limiting the range 20-54
  - resetting the range 20-54
  - reverse color palette 20-53
- Color Shelf
  - continuous fields 10-21
  - discrete fields 10-21
- Color shelf 10-21
- Colors
  - formatting 20-49
- Column Attributes 20-35
- Column Labels
  - formatting 20-35
  - maximum number 20-35
- Columns shelf 10-10
- Comments. See Annotations
- Connecting 4-6
  - Extract Data File 4-18
  - Firebird 4-25
  - Hyperion Essbase 4-29
  - IBM OLAP Server 4-64
  - MS Access 4-11
  - MS Analysis Services Cube 4-41
  - MS Excel 4-14
  - MySQL 4-49
  - Oracle 4-57
  - PostgreSQL 4-61
  - SQL Server 4-45
  - SSAS 4-41
  - tab-delimited text files 4-23
  - Tableau Data Extract 4-18
  - Text of CSV file 4-20
- Connection
  - adding tables 5-2
  - close 4-71
  - editing tables 5-2
  - exporting 4-68
  - saving 4-68
- connection
  - data source 4-6
- context filters 12-55
- Continuous
  - dates 13-12
- Continuous and Discrete 8-22
- Continuous Fields

- color shelf 10-21
- filtering 12-36
- Converting
  - discrete to continuous 8-26
  - measures to dimensions 8-24
- Converting data types 26-15
  - DATE 26-16
  - DATETIME 26-16
  - FLOAT 26-16
  - INT 26-16
  - STR 26-16
- Copy & Paste
  - ad-hoc groups 16-22
  - calculated fields 16-22
  - sets 16-22
  - user filters 16-22
- copy image 24-21
- Count 16-5
- Count Distinct 16-6
- Cross-Tab
  - duplicating worksheets 2-19
- CSV file 4-20
  - mixed data types 8-21
- Custom SQL 4-65
- D**
- Dashboards 2-17, 22-1
  - adding viewing 22-3
  - creating 22-3
  - organizing 22-8
  - rearranging views 22-9
  - removing views 22-5
  - show/hide 22-8
  - tips 22-15
- Data Labels
  - editing with aliases 21-20
- moving 21-19
- show/hide 21-18
- Data labels 21-17
- Data Roles 8-20, 8-22
  - continuous Dimensions 8-32
  - continuous dimensions 8-28
  - continuous measures 8-33
  - continuous vs. discrete 8-26
  - converting discrete to continuous 8-26
  - converting measures to dimensions 8-24
  - dimension vs. measure 8-23
  - Discrete Dimension 8-30
  - Discrete Measures 8-31
  - discrete measures 8-28
- Data Source
  - close 4-71
  - connection examples 4-11
  - export connection 4-68
- data source
  - connection 4-6
  - refresh 4-70
  - supported 4-3
- data source properties 4-9
- Data Types 8-20, 26-3
  - BOOLEAN 26-4
  - changing 8-20
  - DATE/TIME 26-3
  - mixed (Excel and CSV) 8-21
  - NUMBER 26-3
  - STRING 26-3
- data types
  - mixed 8-21
- Data window
  - features and functions 8-8
- data window 2-6
- Database
  - connecting 4-6

- Databases
  - oracle 4-57
- Date
  - continuous 13-12
  - See Dates
- date
  - perfect pivoting 13-10
- Date data type 8-20
- Date Functions 26-13
  - DATEADD 26-13
  - DATEDIIF 26-13
  - DATENAME 26-14
  - DATEPART 26-14
  - DATETRUNC 26-14
  - DAY 26-14
  - ISDATE 26-14
  - MAX 26-14
  - MIN 26-15
  - MONTH 26-15
  - NOW 26-15
  - TODAY 26-15
  - YEAR 26-15
- Dates 13-1
  - computation standards 13-4
  - levels 13-4
- Default Fields
  - Measure Names 8-7
  - Measure Values 8-7
  - number of records 8-7
- Default Member 10-35
- Deleting
  - annotations 21-15
  - bookmarks 24-8
- deleting sheets 2-20
- dialog box
  - Connect to Data Source 4-7
  - Open 4-6
- Dimensions
  - add to sheet 11-8
  - filtering 12-20
  - gray highlight 10-39
  - incompatible measures 10-39
- dimensions 8-4
  - Columns shelf 10-10
  - Filters shelf 10-16
  - Level of Detail shelf 10-19
  - Path shelf 10-31
  - Rows shelf 10-10
  - Shape shelf 10-28
  - Text shelf 10-29
- Dimensions and Measures 8-22
- Disaggregation 16-7
- Discrete Fields
  - color shelf 10-21
- Discrete fields
  - filtering 12-20
- discrete measure calculation 16-37
- Dividers
  - row and columns 20-11
- Double-Click
  - add to sheet 11-8
- drag and drop warning 10-4
- Drill down
  - nonuniform 10-34
- duplicate worksheet 2-19

**E**

- Equally-Sized Bars 16-88
- Essbase 4-29
- Excel data source 4-14
- Excel workbook
  - mixed data types 8-21
- Export
  - data source connection 4-68
  - to print 25-3
- export 24-10
  - Access 24-16
  - image 24-21
  - measures 10-21
- external filters 12-50
- Extract Data File
  - connecting 4-18
- Extracts
  - creating 6-3
  - defining 6-3
  - removing 6-4
  - using 6-4

**F**

- Field Aliases 8-11
- Field Names
  - unknown 4-67
- Fields
  - dropping on worksheets 10-1
  - formatting 20-13
- Files and Folders 2-21
- Filter
  - continuous fields 12-36
  - dimensions 12-20
  - discrete fields 12-20
  - measures 12-36
  - rows and columns 20-33

- user filters 23-24
- filter 12-19
  - context 12-55
  - external 12-50
  - summation icon 12-52
- Filters
  - sets 12-70
- Filters shelf 10-16
- Firebird
  - connect to 4-25
- Focus Tool 14-5
- Focusing 14-5
- Fonts
  - formatting 20-5
- Formatting
  - aspect ratio 20-35
  - axis appearance 20-45
  - axis range 20-37
  - borders 20-11
  - cell size 20-31
  - centering text 20-8
  - clear from worksheet 2-18
  - clear worksheet 2-18
  - colors 20-49
  - column labels 20-35
  - default number format 20-35
  - defaults 20-4
  - dividers 20-11
  - fields 20-13
  - fonts 20-5
  - format window 20-4
  - grid lines 20-11
  - headers 20-35
  - keyboard shortcuts 20-34
  - lines 20-11
  - mark borders 20-18
  - mark halos 20-19

- maximum number 20-35
- panes 20-35
- reference lines 14-17
- row and column banding 20-9
- row and column order 20-33
- row and column size 20-27
- row labels 20-35
- shading 20-9
- table layout 20-27
- table size 20-29
- text 20-5
- text alignment 20-8
- tick marks 20-47
- worksheet 20-4
- zero lines 20-11

## Functions 26-6

- aggregate 26-19
- converting data types 26-15
- dates 26-13
- logical 26-17
- pass through 26-21
- special 26-23
- strings 26-10
- type conversion 26-15

## G

### Gantt bar

- mark type 9-25

### generations 10-35

### Grand Totals 16-65

- aggregations 16-67

### grand totals

- sorting 12-9

### Group Members 2-9

### Grouping Measures 16-59

## H

### Headers

- creating in a table 10-10

### Heat Map

- overlapping data 9-18

### Heat Maps

- cell size

- Cross Tabs

- cell size 20-31

### Hide

- values in the filter dialog 12-29

### Hide All

- shelves 2-15

### hierarchy 8-5

### Histogram 16-59

### Histograms 16-59

### history 14-8

### HTML Help

- printing 25-8

### HVE Interface 2-5

### Hyperion Essbase 4-29

- building views 10-35

- set accounts dimension 4-32

### Hyperlinks 15-1

- editing 15-14

- example 15-16

- invalid 15-13

- launching links 15-9

- managing 15-14

- sending field values 15-10

- sheet links 15-6

- url links 15-3

- valid 15-13

**I**

IBM DB2 OLAP Server 4-64  
Images. See Background Images 17-1  
Import Data Source  
    from server 23-21  
import from workbook 4-9  
Invalid Fields 4-67  
ISO 8601 13-4

**J**

Joins  
    adding tables 5-2  
    editing 5-2

**K**

keyboard shortcut 29-1  
Keyboard shortcuts  
    table layout 20-34

**L**

Legends  
    filtering 12-20  
    highlighting 14-9  
    sorting 12-17  
Level of Detail  
    can't see mark borders 20-18  
Level of Detail shelf 10-19  
    exporting measures 10-21  
levels 10-35  
Line  
    mark type 9-22  
Lines  
    formatting 20-11  
Links. See also Hyperlinks

**Literal expressions**

    booleans 26-5  
    dates 26-5  
    Null 26-5  
    numeric 26-5  
    strings 26-5  
Logical Functions 26-17  
    CASE 26-17  
    IF THEN ELSE END 26-18  
    IF THEN ELSEIF 26-19  
    IIF 26-18  
    ISDATE 26-19  
    ISNULL 26-19  
    MIN 26-19

**M**

Maps. See Background Images  
Marks 9-11  
    automatic 9-12  
    bar 9-15  
    borders 20-18  
    changing size 9-32  
    circle 9-19  
    gantt bar 9-25  
    halos 20-19  
    line 9-22  
    polygon 9-24  
    shape 9-20  
    square 9-16  
    text 9-12  
marks  
    stacking 9-26  
Maximum 16-4  
Measure Names 8-7, 8-17  
    editing 8-17  
    editing aliases 8-17

- Measure names
    - removing 8-17
  - Measure Values 8-7, 8-17
  - Measures
    - add to sheet 11-8
    - caution symbol 10-39
    - filtering 12-36
    - gray highlight 10-39
    - incompatible dimensions 10-39
    - setting as accounts dimension 4-32
  - measures 8-4
    - Columns shelf 10-10
    - Filters shelf 10-16
    - Level of Detail shelf 10-21
    - Path shelf 10-31
    - Rows shelf 10-10
    - Shape shelf 10-28
    - Text shelf 10-29
  - Member Aliases 8-11
  - Menu Commands
    - table layout 20-34
  - Message
    - show/hide 9-13
  - Microsoft Access database 4-11
  - Microsoft Analysis Services Cube
    - building views 10-39
  - Microsoft Analysis Services cube 4-41
  - Microsoft Excel workbook 4-14
  - Microsoft SQL Server database 4-45
  - Minimum 16-4
  - mixed data types 8-21
  - moving the repository 2-23
  - MSAS 4-41
  - Multidimensional
    - schema 8-5
  - multidimensional
    - data sources 4-3
    - hierarchy 8-5
    - slicer 12-53
  - Multiple Measures 8-7, 8-17
  - MySQL database 4-49
    - max\_allowed\_packet error 14-22
- N**
- Navigation Toolbar
    - focus 2-9
    - pan 2-9
    - select 2-9
    - zoom 2-9
  - new workbook 2-17
  - nonuniform drill down 10-34
  - Number Format 20-35
  - Number functions
    - ABS 26-6
    - ACOS 26-6
    - ASIN 26-6
    - ATAN 26-6
    - ATAN2 26-7
    - COS 26-7
    - COT 26-7
    - DEGREES 26-7
    - EXP 26-7
    - LN 26-7
    - LOG 26-7
    - MAX 26-8
    - MIN 26-8
    - PI 26-8
    - POWER 26-8
    - RADIANS 26-8
    - ROUND 26-9
    - SIGN 26-9
    - SIN 26-9
    - SQRT 26-9

SQUARE 26-9  
TAN 26-9  
Number of Records 8-7  
Numbers  
    default format 20-35  
Numerical data type 8-20

## O

One-click Sort 2-9  
Open dialog box 4-6  
Open Workbook  
    from server 23-18  
Opening HVE 3-3  
Opening Tableau 3-3  
opening Tableau 2-3  
Operators 26-25  
    - (subtraction) 26-25  
    % (modulo) 26-25  
    \* (multiplication) 26-25  
    + (addition) 26-25  
    / (division) 26-25  
    AND 26-26  
    comparisons 26-26  
    NOT 26-27  
    OR 26-26  
Oracle Database  
    configuring 4-57  
Order of Operations  
    See Precedence  
overlapping text 9-13

## P

Pages Shelf 10-12  
panning 14-4  
Pass Through Functions 26-21

RAWSQL\_BOOL 26-21  
RAWSQL\_DATE 26-21  
RAWSQL\_DATETIME 26-21  
RAWSQL\_INT 26-22  
RAWSQL\_REAL 26-22  
RAWSQL\_STR 26-22  
RAWSQLAGG\_BOOL 26-22  
RAWSQLAGG\_DATE 26-22  
RAWSQLAGG\_DATETIME 26-22  
RAWSQLAGG\_INT 26-23  
RAWSQLAGG\_REAL 26-23  
RAWSQLAGG\_STR 26-23  
Path shelf 10-31  
PDF  
    print to 25-7  
    publish 25-7  
Percentages 16-74  
    about 16-74  
    aggregations 16-75  
    column 16-79  
    column in pane 16-87  
    options 16-77  
    pane 16-82  
    row 16-81  
    row in pane 16-84  
    table 16-78  
perfect pivoting  
    date 13-10  
Performance  
    all data sources 27-3  
    context filters 27-8  
    large text files 27-9  
    multidimensional data sources 27-7  
    relational data sources 27-5  
    tips 27-1  
Polygon mark type 9-24  
PostgreSQL



- connect to 4-61
- Precedence 26-27
- Precision
  - warning 7-6
- Presentation Mode 2-15
- Presenting
  - presentation mode 2-15
- Printing 25-3, 25-5
  - HTML Help 25-8
  - page orientation 25-4
  - to PDF 25-7
- Publish as PDF 25-7
- Publish to Server 23-1
- Publishing Data Sources 23-14
- Publishing Workbooks
  - hide sheets 23-12
  - passwords 23-13
  - permissions 23-7
  - server 23-3
  - user filters 23-24

## Q

- qualified member names 12-76
- Queries
  - abandoned 7-5
  - data source unavailable 7-5
  - precision warnings 7-6
- quitting Tableau 2-3

## R

- ragged hierarchy 10-35
- RAWSQL 26-21
- redo 14-8
- Reference Lines
  - formatting 14-17
- refresh data source 4-70
- Relational
  - schema 8-5
- relational
  - data sources 4-3
- Removing
  - bookmarks 24-8
- Rename Fields 8-11
- Rename Members 8-11
- rename sheet 2-18
- Re-Ordering
  - rows and columns 20-33
- Replace References 4-67
- Repository 2-21
  - moving 2-23
- Resizing
  - cells 20-31
  - entire table 20-29
  - rows and columns 20-27
- Row attributes 20-35
- Row Labels
  - maximum number 20-35
- Row Lables
  - formatting 20-35
- Rows and Columns
  - chaning order 20-33
  - include/exclude 20-33
  - re-ordering 20-33
  - show/hide 20-33
- Rows shelf 10-10
- Running Totals 16-40

**S**

save workbook 24-3

## Schema

add fields 8-10

find 8-8

hide/unhide fields 8-9

multidimensional 8-5

organize 8-8

relational 8-5

rename fields 8-9, 8-11

## schema

hiding 8-6

selecting marks 14-3

## Server 23-1

publishing workbooks 23-3

## Sets 12-70

about 12-70

creating 12-72

## Shading

formatting 20-9

shape mark type 9-20

Shape shelf 10-28

shared members 10-36

## Sheets

closing 2-20

deleting 2-20

renaming 2-18

## shelf

Color 10-21

Columns 10-10

Level of Detail 10-19

Path 10-31

Rows 10-10

Shape 10-28

Size 10-23

Text 10-29

## Shelves

pages 10-12

types 10-10

## shortcut

data source connection 4-68

keyboard 29-1

## Show

values in the filter dialog 12-29

## Show All

shelves. 2-15

## Show Me! 11-4

alternatives 11-6

## Show Me! Toolbar

Show Me! 2-8

Show Me! Alternatives 2-9

swap 2-9

## Show/Hide

data labels 21-18

parts of dashboard 22-8

rows and columns 20-33

## Size

marks 9-32

## Size shelf 10-23

## slicer 12-53

## Sort Ascending 2-9

## Sort Descending 2-9

## Sorting

manually 20-33

## sorting 12-3

grand totals 12-9

## Special Functions 26-23

CURRENTUSER 26-23

ISCURRENTUSER 26-24

ISMEMBEROF 26-24

## Spotlighting 16-37

## SQL

custom query 4-65

Server database 4-45

- typing in 4-65
- Square Cell 20-31
- Square mark type 9-16
- stacking marks 9-26
- Standard Deviation 16-5
  - population 16-5
- Standard Toolbar
  - automatic queries 2-8
  - clear sheet 2-8
  - connect to data source 2-8
  - duplicate sheet 2-8
  - new sheet 2-8
  - presentation mode 2-9
  - redo 2-8
  - run query 2-8
  - save 2-8
  - undo 2-8
  - worksheet windows 2-9
  - zoom 2-9
- Status Bar 2-10
  - precision warning 7-6
  - warnings 2-10
- String Functions 26-10
  - ASCII 26-10
  - CHAR 26-10
  - FIND 26-10
  - LEFT 26-10
  - LEN 26-10
  - LOWER 26-10
  - LTRIM 26-11
  - MAX 26-11
  - MID 26-11
  - MIN 26-11
  - RIGHT 26-11
  - RTRIM 26-11
  - SPACE 26-12
  - TRIM 26-12

- UPPER 26-12
- Subtotals 16-71
  - local vs. server 16-63
- Sum 16-4
- Summary Card 14-20
- summation icon 12-52
- supported data sources 4-3
- Symbols
  - caution 7-6

## T

- Tab Delimited Text File
  - connecting to 4-23
- Table
  - components 9-2
- table
  - axes creation 10-10
  - header creation 10-10
- Table Calculations 16-40
  - advanced 16-58
  - cumulative total 16-54
  - defining 16-41
  - difference from 16-43
  - moving average 16-40, 16-55
  - moving calculation 16-55
  - multiple calculations 16-58
  - percent difference from 16-45
  - percent from 16-48
  - percent of total 16-51
  - running total 16-40, 16-54
- Table Layout
  - formatting 20-27
  - keyboard shortcuts 20-34
  - menu commands 20-34
- Table Options 20-35
- Tableau Interface 2-5

- changing 2-14
- rearranging 2-14
- Tableau Server 23-1
- Tables
  - editing 5-2
- Text
  - formatting 20-5
- text
  - mark type 9-12
  - overlapping 9-13
- Text Cell 20-32
- Text data type 8-20
- Text file 4-20
- Text shelf 10-29
- Tick Marks 20-47
- Time 13-1
- Tips and Tricks 28-1
- Toolbar
  - Group 2-9
  - Sort Ascending 2-9
  - Sort Descending 2-9
- toolbars 2-8
  - navigation toolbar 2-8
  - standard toolbar 2-8
- ToolTips 3-14
- Totals
  - grand totals 16-65
  - local vs. server 16-63
  - subtotals 16-71
- Type Conversion 26-15
  - DATE 26-16
  - DATETIME 26-16
  - FLOAT 26-16
  - INT 26-16
  - STR 26-16

## U

- undo 14-8
- Updates
  - automatic 7-3
- User Filters 23-24
  - creating 23-24
  - editing 23-30
  - example 23-33
  - previewing as user 23-29
  - publishing 23-32
  - special functions 23-40
  - using data source 23-40

## V

- Variance 16-5
  - population 16-5
- View Cards
  - rearranging 2-14
- View Statistics 14-20
- viewing underlying data 14-21
- Views
  - automatic 11-3

## W

- Warnings
  - precision 7-6
  - show/hide 9-13
- Workbooks 2-16
- workbooks
  - importing from 4-9
  - saving 24-3
- Worksheets 2-17
  - clearing 2-18
  - dropping fields 10-1
  - duplicating 2-19

- duplicating as cross-tab 2-19
  - new 2-18
  - remove contents 2-18
- worksheets
  - clearing 2-18

**Z**

- Zoom mode 14-6
- Zoom tool 14-6

# Copyright Information

---

Copyright ©2003-2008 Tableau Software, Incorporated and its licensors. All rights reserved. Patents pending.

Portions of the code copyright ©1997-2008 The Board of Trustees of the Leland Stanford Junior University. All rights reserved.

Tableau's installation includes an unmodified executable version of the Firebird database. The source code for that database can be found at <http://www.firebirdsql.org>.

This product is Client Software as defined in Tableau Software's End User Software License Agreement.

---