



# 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.

## 2. What is Tableau?

- Overview . . . . . 2**
- What can I do with Tableau? . . . . . 2-3**
  - Visually Analyze Data Rapidly . . . . . 2-3
  - Create Insightful Reports and Views of Your Data . . . . . 2-4
  - Build Interactive Dashboards . . . . . 2-4
  - Share and Interact . . . . . 2-5
- What data can I analyze with Tableau? . . . . . 2-6**
- How Does Tableau Work? . . . . . 2-7**

## 3. The Tableau Environment

- Overview . . . . . 3-2**
- Opening and Closing the Application . . . . . 3-3**
- Start Page . . . . . 3-4**
- The Tableau Workspace . . . . . 3-6**
  - Data Window . . . . . 3-7
  - Toolbar . . . . . 3-8
  - Status Bar . . . . . 3-10
  - Cards and Shelves . . . . . 3-11
  - Reorganizing the Workspace . . . . . 3-15
- Workbooks and Sheets . . . . . 3-17**
  - Workbooks . . . . . 3-17
  - Sheets . . . . . 3-18
- Files and Folders . . . . . 3-23**
  - Changing the Repository Location . . . . . 3-25



## 4. Learning to Use Tableau

<b>Overview</b>	4-2
<b>Open Tableau</b>	4-3
<b>Connect to Data</b>	4-4
<b>Build Data Views (Manual)</b>	4-8
Example 1 - Yearly Profits	4-9
Example 2 - Quarterly Profits by Year	4-14
Example 3 - Quarterly Profits by Year and Market	4-16
Example 4 - Filter by Year and by Sales	4-18
Example 5 - Color Encode by Region	4-24
<b>Build Data Views (Automatic “Show Me!”)</b>	4-26
Yearly Profits	4-26
Sales versus Profit by Product and Customer	4-31
<b>Save Your Work</b>	4-34

## 5. Basic Connection

<b>Supported Data Sources</b>	5-3
<b>How to Connect to a Data Source</b>	5-4
Connect Using the Open Dialog Box	5-4
Connect Using the Connect to Data Dialog Box	5-5
<b>Examples – Connecting to Data Sources</b>	5-8
Microsoft Access Database	5-9
Microsoft Excel Workbook	5-11
OData	5-14
Tableau Data Extract File	5-16
Text File	5-17
Windows Azure Marketplace DataMarket	5-20
Firebird Database	5-22
Greenplum	5-24



IBM DB2 . . . . .	5-27
InterSystems Caché . . . . .	5-30
Microsoft Analysis Services Cube . . . . .	5-33
Microsoft PowerPivot . . . . .	5-36
Microsoft SQL Server Database . . . . .	5-38
MySQL Database . . . . .	5-41
Netezza . . . . .	5-43
Oracle Database . . . . .	5-46
Oracle Essbase Cube (aka IBM OLAP Server) . . . . .	5-49
ParAccel . . . . .	5-52
PostgreSQL Database . . . . .	5-54
Sybase IQ . . . . .	5-57
Teradata . . . . .	5-60
Vertica . . . . .	5-64
<b>Connecting to a Custom SQL Query . . . . .</b>	<b>5-68</b>
<b>Editing the Connection . . . . .</b>	<b>5-69</b>
<b>Renaming the Connection . . . . .</b>	<b>5-72</b>
<b>Duplicating a Connection . . . . .</b>	<b>5-73</b>
<b>Replacing Field References . . . . .</b>	<b>5-71</b>
<b>Refreshing the Data . . . . .</b>	<b>5-76</b>
<b>Closing the Connection . . . . .</b>	<b>5-77</b>

## 6. Clipboard Data Sources

## 7. Joining Tables

Overview . . . . .	7-2
Adding Tables . . . . .	7-3



---

Connecting to Multiple Tables . . . . .	7-3
Adding Tables to the Data window . . . . .	7-7
<b>Editing Tables . . . . .</b>	<b>7-8</b>

## 8. Working with Multiple Connections

<b>Understanding Data Integration . . . . .</b>	<b>8-3</b>
<b>Navigating Connections in the Data Window . . . . .</b>	<b>8-4</b>
<b>Adding a Secondary Connection . . . . .</b>	<b>8-6</b>
<b>Defining Relationships . . . . .</b>	<b>8-8</b>
<b>Example - Data Integration on a Worksheet . . . . .</b>	<b>8-12</b>
<b>Troubleshooting Data Integration . . . . .</b>	<b>8-17</b>

## 9. Extracting Data to the Data Engine

<b>Overview . . . . .</b>	<b>9-2</b>
<b>Creating an Extract . . . . .</b>	<b>9-3</b>
<b>Using Extracts . . . . .</b>	<b>9-7</b>

## 10. Managing Queries

<b>Overview . . . . .</b>	<b>10-2</b>
<b>Automatic Updates . . . . .</b>	<b>10-3</b>
<b>Cancel Query . . . . .</b>	<b>10-4</b>



**Abandoned Queries** . . . . . 10-5

**Precision Warnings** . . . . . 10-6

## 11. Understanding Data Fields

**Understanding the Data Window** . . . . . 11-3

Relational and Multidimensional Data . . . . . 11-6

Measure Values and Measure Names . . . . . 11-11

Number of Records . . . . . 11-11

Latitude and Longitude (generated) . . . . . 11-11

**Data Window Features and Functions** . . . . . 11-13

Organize the Data window . . . . . 11-13

Find Fields . . . . . 11-14

Rename Fields . . . . . 11-15

Hide or Unhide Fields . . . . . 11-16

Add Fields to the Data window . . . . . 11-17

**Editing Field Properties** . . . . . 11-18

Comments . . . . . 11-19

Aliases . . . . . 11-20

Colors . . . . . 11-27

Shapes . . . . . 11-27

Formats . . . . . 11-28

Sort . . . . . 11-29

Aggregation . . . . . 11-29

Changing Measure Names . . . . . 11-32

**Data Types & Roles** . . . . . 11-34

**Data Types & Roles** . . . . . 11-34

Data Types . . . . . 11-34

Data Roles . . . . . 11-37

Example: Continuous Dimensions and Discrete Measures . . . . . 11-42

## 12. Parts of the View



<b>Overview</b>	12-2
<b>Headers</b>	12-3
<b>Axes</b>	12-6
<b>Panes</b>	12-9
<b>Cells</b>	12-10
<b>Marks</b>	12-11
Mark Types	12-11
Stacking Marks	12-27
Changing Mark Size and Color	12-31
<b>Titles</b>	12-34
<b>Captions</b>	12-35
<b>Field Labels</b>	12-36
<b>Legends</b>	12-38

## 13. Building Views Manually

<b>Dragging Fields</b>	13-3
<b>Types of Shelves</b>	13-10
Columns and Rows Shelves	13-10
Pages Shelf	13-14
Filters Shelf	13-18
Level of Detail Shelf	13-21
Color Shelf	13-23
Size Shelf	13-31
Shape Shelf	13-38
Label Shelf	13-47
Path Shelf	13-49



<b>Multidimensional Hierarchies</b> . . . . .	<b>13-51</b>
Drilling Down and Up in a Hierarchy . . . . .	13-51
Building Views with Oracle Essbase . . . . .	13-54
Building Views with Microsoft Analysis Services Cubes . . . . .	13-58
Perfect Pivoting . . . . .	13-59
Defining Unique Values . . . . .	13-60
<b>Working with Large Views</b> . . . . .	<b>13-66</b>
<b>Example – Building Data Views Manually</b> . . . . .	<b>13-67</b>

## 14. Building Views Automatically

<b>Overview</b> . . . . .	<b>14-2</b>
<b>Show Me!</b> . . . . .	<b>14-3</b>
<b>Add to Sheet: Double-Click</b> . . . . .	<b>14-6</b>

## 15. Using Multiple Measures

<b>Individual Axes</b> . . . . .	<b>15-3</b>
<b>Blended Axes</b> . . . . .	<b>15-4</b>
<b>Dual Axes</b> . . . . .	<b>15-6</b>
<b>Combination Charts</b> . . . . .	<b>15-8</b>



## 16. Filtering

<b>Overview</b> . . . . .	<b>16-2</b>
<b>Adding Filters</b> . . . . .	<b>16-3</b>
Selecting Data to Filter . . . . .	16-3
Dragging Fields to the Filters Shelf . . . . .	16-7
Using Quick Filters . . . . .	16-38
<b>Global Filters</b> . . . . .	<b>16-49</b>
<b>Context Filters</b> . . . . .	<b>16-50</b>
Creating a Context Filter . . . . .	16-51
Example – Context Filters . . . . .	16-51
<b>Calculation Filters</b> . . . . .	<b>16-57</b>

## 17. Sorting, Grouping, and Sets

<b>Overview</b> . . . . .	<b>17-2</b>
<b>Sorting</b> . . . . .	<b>17-3</b>
Computed Sorting . . . . .	17-3
Manual Sorting . . . . .	17-14
<b>Groups</b> . . . . .	<b>17-18</b>
Creating Groups . . . . .	17-18
Editing an Existing Group . . . . .	17-22
Finding Members in the Groups Dialog Box . . . . .	17-25
<b>Sets</b> . . . . .	<b>17-26</b>
About Sets . . . . .	17-26
How to Create a Set . . . . .	17-28
Creating Sets Examples . . . . .	17-37



## 18. Dates and Times

- Overview . . . . . 18-2
- Changing Date Levels . . . . . 18-4
- Fiscal Dates . . . . . 18-8
- Perfect Pivoting with Dates . . . . . 18-10
- Continuous Dates . . . . . 18-12

## 22. Reference Lines and Bands

- Types of Reference Lines and Bands . . . . . 22-3
- Adding Reference Lines . . . . . 22-5
- Adding Reference Bands . . . . . 22-9
- Adding Reference Distrubtions (Bullet Graphs) . . . . . 22-12
  - Basic Reference Distributions . . . . . 22-12
  - Bullet Graphs . . . . . 22-15
- Editing Reference Lines and Bands . . . . . 22-19
- Removing Reference Lines and Bands . . . . . 22-20



---

## 19. Inspecting Data

- Overview . . . . . 19-2**
- Navigation Tools . . . . . 19-3**
  - Select . . . . . 19-3
  - Zoom . . . . . 19-5
  - Pan . . . . . 19-7
- Undo and Redo . . . . . 19-8**
- Drop Lines . . . . . 19-9**
- Summary Card . . . . . 19-11**
- View Data (Drill-Through) . . . . . 19-12**
- Describing the View . . . . . 19-16**



## 20. Actions

- Overview . . . . . 20-2**
- Filter Actions . . . . . 20-3**
- Highlight Actions . . . . . 20-8**
  - Selecting Marks to Highlight . . . . . 20-8
  - Color Legend Highlighting . . . . . 20-9
  - Highlight Toolbar Button . . . . . 20-11
  - Creating Advanced Highlight Actions . . . . . 20-13
- URL Actions . . . . . 20-16**
- Running Actions . . . . . 20-20**
- Actions and Dashboards . . . . . 20-22**
- Using Field Values in Actions . . . . . 20-24**
  - Using Field Values in URLs . . . . . 20-24
  - Using Field Values in Action Names . . . . . 20-25

## 21. Calculations

- Overview . . . . . 21-3**
- Aggregations . . . . . 21-4**
  - How Aggregation and Disaggregation Work . . . . . 21-8
  - Example – Aggregating and Disaggregating Data . . . . . 21-11
- Calculated Fields . . . . . 21-18**
  - How to Create a Calculated Field . . . . . 21-18
  - How to Create a Calculated Member . . . . . 21-21
  - Copying and Pasting Calculated Fields . . . . . 21-24
  - Example – Creating a Calculated Field . . . . . 21-26
  - Table Calculations . . . . . 21-40
  - Aggregate Calculations . . . . . 21-34



<b>Table Calculations</b> . . . . .	<b>21-40</b>
Quick Table Calculations . . . . .	21-49
Defining Basic Table Calculations . . . . .	21-50
Secondary Table Calculations . . . . .	21-67
Customizing Table Calculations . . . . .	21-68
<b>Binned Data</b> . . . . .	<b>21-72</b>
<b>Totals</b> . . . . .	<b>21-77</b>
Local vs. Server Computation . . . . .	21-77
Grand Totals . . . . .	21-78
Subtotals . . . . .	21-84
<b>Percentages</b> . . . . .	<b>21-87</b>
About Percentages . . . . .	21-87
Percentages and Aggregations . . . . .	21-88
Percentage Options . . . . .	21-91

## 22. Parameters

<b>Creating Parameters</b> . . . . .	<b>22-3</b>
<b>Editing Parameters</b> . . . . .	<b>22-9</b>
<b>Using Parameters in Calculations</b> . . . . .	<b>22-10</b>
<b>Parameter Controls</b> . . . . .	<b>22-12</b>
<b>Example - Parameters</b> . . . . .	<b>22-14</b>

## 23. Background Images

<b>Overview</b> . . . . .	<b>23-2</b>
<b>Adding Background Images</b> . . . . .	<b>23-3</b>



**Setting up the View** . . . . . 23-7

**Managing Background Images** . . . . . 23-8

Editing an Image . . . . . 23-8

Enabling/Disabling Images . . . . . 23-9

Adding Show/Hide Conditions . . . . . 23-10

Removing an Image . . . . . 23-13



---

## 24. Background Maps

<b>Overview</b> . . . . .	24-2
<b>Geographic Roles</b> . . . . .	24-3
<b>Building a Map View</b> . . . . .	24-6
<b>Map Options</b> . . . . .	24-7
Map Layers . . . . .	24-8
Data Layers . . . . .	24-10
Washout . . . . .	24-11
<b>Setting a Default Location</b> . . . . .	24-13
<b>Editing Locations</b> . . . . .	24-14
<b>Custom Geocoding</b> . . . . .	24-16
Creating an Import File . . . . .	24-16
Importing Custom Geocoding . . . . .	24-20
Saving Custom Geocoding . . . . .	24-22
<b>Background Map Sources</b> . . . . .	24-23
<b>Map Storing and Working Offline</b> . . . . .	24-26



## 25. Trend Lines & Statistics

- Overview . . . . . 25-3**
- Adding Trend Lines . . . . . 25-4**
  - Add Trend Lines to the View . . . . . 25-4
  - Why can't I add Trend Lines? . . . . . 25-6
  - Remove Trend Lines . . . . . 25-7
- The Trend Line Model . . . . . 25-8**
- Removing Factors from the Model . . . . . 25-9**
- Testing Significance . . . . . 25-12**
  - Entire Model Significance . . . . . 25-12
  - Significance of Specific Fields . . . . . 25-13
- Trend Lines Example . . . . . 25-14**
- Assumptions . . . . . 25-18**
- Trend Line Model Terms . . . . . 25-19**
- Commonly Asked Questions . . . . . 25-21**

## 26. Log Axes

---

+

## 27. Formatting

<b>Overview</b> . . . . .	27-3
<b>Worksheet Level Formatting</b> . . . . .	27-4
Opening the Format Window . . . . .	27-4
Fonts . . . . .	27-6
Alignment . . . . .	27-9
Shading . . . . .	27-10
Borders . . . . .	27-12
Lines . . . . .	27-12
<b>Formatting Specific Parts of the View</b> . . . . .	27-14
Fields . . . . .	27-14
Numbers . . . . .	27-16
Field Labels . . . . .	27-20
Legends . . . . .	27-21
Title and Caption . . . . .	27-23
Tooltips . . . . .	27-25
Null Values . . . . .	27-27
Reference Lines and Bands . . . . .	27-29
<b>Copying and Pasting Formatting</b> . . . . .	27-30
<b>Clearing Formatting</b> . . . . .	27-31
<b>Workbook Themes</b> . . . . .	27-32
<b>Resizing the Table</b> . . . . .	27-33
Resizing Rows and Columns . . . . .	27-33
Resizing the Entire Table . . . . .	27-35
Resizing Cells . . . . .	27-37
Reordering Rows and Columns . . . . .	27-39
Including and Excluding Rows and Columns . . . . .	27-39
Resize Keyboard Shortcuts and Commands . . . . .	27-40
Clearing Manual Sizing . . . . .	27-40
<b>Editing Axes</b> . . . . .	27-43
Changing the Axis Range . . . . .	27-43



Changing the Axis Appearance . . . . . 27-50  
Formatting Tick Marks . . . . . 27-52

## 28. Annotations and Mark Labels

**Overview** . . . . . 28-2

**Annotations** . . . . . 28-3  
Adding Annotations . . . . . 28-3  
Rearranging Annotations . . . . . 28-5  
Formatting Annotations . . . . . 28-11  
Removing Annotations . . . . . 28-15

**Mark Labels** . . . . . 28-16  
Showing and Hiding Mark Labels for the Worksheet . . . . . 28-16  
Showing and Hiding Individual Mark Labels . . . . . 28-17  
Moving Mark Labels . . . . . 28-18  
Editing Mark Labels with Aliases . . . . . 28-19

## 29. Dashboards

**Overview** . . . . . 29-2

**Creating Dashboards** . . . . . 29-3  
How to Create a Dashboard . . . . . 29-3  
Adding Views to a Dashboard . . . . . 29-3  
Adding Dashboard Objects . . . . . 29-6  
Removing Views and Objects from a Dashboard . . . . . 29-8

**Organizing Dashboards** . . . . . 29-10  
Showing and Hiding Parts of a Worksheet . . . . . 29-15  
Rearranging Dashboard Views and Objects . . . . . 29-15  
Setting the Dashboard Size . . . . . 29-18

**Understanding Dashboards and Worksheets** . . . . . 29-20



## 30. Publishing to Tableau Server

<b>Overview</b>	<b>30-2</b>
<b>Publishing Workbooks</b>	<b>30-3</b>
How to Publish Workbooks to the Server	30-3
Specifying Permissions	30-9
Showing and Hiding Worksheets	30-15
Embedding Passwords	30-16
Schedules	30-17
<b>Publishing Data Sources</b>	<b>30-19</b>
<b>Opening Workbooks from the Server</b>	<b>30-23</b>
<b>Importing Data Sources from the Server</b>	<b>30-26</b>
<b>User Filtering</b>	<b>30-29</b>
How to Create a User Filter	30-29
Previewing User Filters	30-34
Editing User Filters	30-35
Publishing with User Filters	30-37
Example - Setting User Filters	30-38
Filter Using a List of Users in your Data Source	30-45



# 31. Saving and Exporting

- Saving Your Work . . . . . 31-3**
- Workbooks . . . . . 31-3
- Bookmarks . . . . . 31-7
- Packaged Workbooks . . . . . 31-5
  
- Reverting Workbooks . . . . . 31-9**
  
- Exporting Your Work . . . . . 31-10**
- Export Data . . . . . 31-10
- Export as an Image . . . . . 31-21
- Exporting the Data Source Connection . . . . . 31-24**

# 32. Printing

- Overview . . . . . 32-2**
  
- Printing Results . . . . . 32-3**
- Page Setup . . . . . 32-3
- Printing . . . . . 32-6
  
- Publish as PDF . . . . . 32-8**
  
- Printing the Help . . . . . 32-10**
- Printing the Online Help . . . . . 32-10



## 33. Building Data Views Examples

<b>Bar Chart</b> . . . . .	<b>33-4</b>
Bar Chart–Basic . . . . .	33-5
Bar Chart–Color-Encoded . . . . .	33-7
Bar Chart–Color Encoded Nested Table . . . . .	33-9
Bar Chart–Filtered and Color-Encoded . . . . .	33-10
Bar Chart–Color Encoded and Sorted . . . . .	33-14
Bar Chart–Filtered, Color Encoded, and Level of Detail . . . . .	33-19
<b>Text Table</b> . . . . .	<b>33-22</b>
Text Table–Basic . . . . .	33-22
Text Table–Filter and Color Encode . . . . .	33-26
Duplicate as Crosstab . . . . .	33-28
<b>Line Chart</b> . . . . .	<b>33-30</b>
Line Chart–Basic . . . . .	33-31
Filter and Color Encode a Nested Table . . . . .	33-33
Line Chart–Filter and Color-Encode Multiple Measures . . . . .	33-35
Line Chart–Filter and Path Encode . . . . .	33-37
<b>Scatter Plot</b> . . . . .	<b>33-39</b>
Scatter Plot–Basic . . . . .	33-40
Scatter Plot–Filter and Color Encode . . . . .	33-42
Scatter Plot–Filter and Level of Detail . . . . .	33-45
Scatter Plot–Filter, Level of Detail, and Color Encode . . . . .	33-47
Scatter Plot–Filter, Level of Detail, Color Encode, and Shape Encode . . . . .	33-52
<b>Heat Map</b> . . . . .	<b>33-55</b>
Heat Map–Basic . . . . .	33-56
Heat Map–Size Encoded . . . . .	33-58
Heat Map–Shape and Size Encoded . . . . .	33-61
<b>Gantt Bar Chart</b> . . . . .	<b>33-63</b>
Gantt Bar Chart–Standard . . . . .	33-63
Gantt Bar Chart–Nonstandard . . . . .	33-66
<b>Pie Chart</b> . . . . .	<b>33-68</b>
Pie Chart– Maps . . . . .	33-68



# Appendix A: Functions, Operators, & Data Types

- Data Types** . . . . . **A-3**
  
- Formatting Literals** . . . . . **A-5**
  - Numeric Literals . . . . . A-5
  - String Literals . . . . . A-5
  - Date Literals . . . . . A-5
  - Boolean Literals . . . . . A-5
  - Null Literals . . . . . A-5
  
- Functions** . . . . . **A-6**
  - Number Functions . . . . . A-6
  - String Functions . . . . . A-10
  - Date Functions . . . . . A-13
  - Type Conversion . . . . . A-15
  - Logical Functions . . . . . A-17
  - Aggregate Functions . . . . . A-20
  - Pass Through Functions (RAWSQL) . . . . . A-21
  - User Functions . . . . . A-23
  
- Operators** . . . . . **A-37**
  - Precedence . . . . . A-39

## Appendix B: Performance Tips

## Appendix C: Tips and Tricks

## Appendix D: Keyboard Shortcuts



## **What is Tableau?**

---

---

## Overview

Tableau Software provides software applications for fast analytical and rapid fire business intelligence.

**Tableau Desktop** is a data visualization application that lets you analyze virtually any type of structured data and produce highly interactive, beautiful graphs, dashboards, and reports in just minutes. After a quick installation, you can connect to virtually any data source from spreadsheets to data warehouses and display information in multiple graphic perspectives. Designed to be easy to use, you'll be working faster than ever before.

**Tableau Server** is a business intelligence solution that provides browser-based visual analytics anyone can use at just a fraction of the cost of typical BI software. With just a few clicks, you can publish or embed live, interactive graphs, dashboards and reports with current data automatically customized to the needs of everyone across your organization. It deploys in minutes and users can produce thousands of reports without the need of IT services — all within your IT infrastructure.

**Tableau Reader** is a free viewing application that lets anyone read and interact with packaged workbooks created by Tableau Desktop.

The company is one of the 50 fastest growing software companies in the U.S. Our applications are being used by over 30,000 people worldwide. Customers include companies as diverse as Google, Cleveland Clinic, GM, Microsoft, Wells Fargo, the District of Columbia, Allstate, Cornell and Harvard.

**Tableau Public** is a free service that lets anyone publish interactive data to the web. Once on the web, anyone can interact with the data, download it, or create their own visualizations of it. No programming skills are required.

This section discusses the following topics:

- What can I do with Tableau?
- What data can I analyze with Tableau?
- How Does Tableau Work?



## What can I do with Tableau?

Imagine being able to answer virtually any business question by dragging-and-dropping your data into a free-form visual canvas. You create beautiful graphs, reports and dashboards. You then share those results in just a few clicks. Using Tableau Desktop, 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 faster and more efficiently than ever before. With Tableau Server and , you can share and embed your live, interactive views, reports and dashboards so that colleagues can interact, customize or monitor them.

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

### Visually Analyze Data Rapidly

#### See and Understand

People need effective views 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 ask rapid questions of your data by letting you iteratively create and modify live, interactive charts, reports and dashboards 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've never been able to before.

#### Filter and Identify Hidden Patterns

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 filter 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 step 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.

## **Create Insightful Reports and Views of Your Data**

### **Pivot and Crosstab**

With no training, users can conduct interactive “Q&A” sessions with data. They simply point Tableau to a database and start building drag & drop multi-dimensional crosstabs of their data.

A first report not only provides an insight (e.g., “Costs are increasing.”), but often points to additional questions (e.g., “Which products are responsible for the increase?”). You 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.

### **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.

## **Build Interactive Dashboards**

### **Build Dashboards People Can Understand**

Use Tableau to build dashboards that communicate clearly and directly. Each element of a dashboard presents information in the most effective way possible, based on the latest research in human perception. Tableau provides the display type



that best expresses the data—bar and line charts, maps, tables, scatter plots, and more. Tableau helps you build dashboards that inform and impress.

### **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. 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.

### **Interact and Drill-down**

Sometimes you need to answer additional questions within a dashboard. With Tableau, viewers can dynamically filter, highlight, drill-down and link across multiple views in one dashboard. This essentially creates an interactive visual analysis application on the fly.

## **Share and Interact**

### **Present**

Imagine pasting Tableau’s vivid multi-dimensional results into Microsoft Office applications and sharing them with others. Our users have a reputation for producing high-impact presentations that are easy to understand.

### **Publish and Embed**

Share your graphs, reports and dashboards by publishing them with Tableau Server. Anyone with proper data credentials can view and interact with those visualizations using just a browser. They can even save custom views, make comments or even tag favorites.. Don’t want people to visit a specific URL destination for their views? No problem – embed them in virtually any web application with just a few lines of code.

---

## What data can I analyze with Tableau?

Your data needs to be in a database, spreadsheet or structured text 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 depend on your purchase options. Refer to “Supported Data Sources” on page 5-3 for a complete list of supported data sources.

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).

---

## How Does Tableau Work?

While Tableau lets you analyze databases and spreadsheets like never before, 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 5-3 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 a Tableau worksheet. 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 4-1 to learn many of the basic steps involved when analyzing your data.

- 3 **Share** - You can share results with others either by sharing Tableau workbooks with other Tableau users, by pasting results into other applications such as Microsoft Office, by printing to a PDF or by using Tableau Server to publish or embed your views across your organization. Refer to “Publishing and Sharing” on page 32-1 for more information.







# The Tableau Environment

---

<b>Overview</b> . . . . .	<b>3-2</b>
<b>Opening and Closing the Application</b> . . . . .	<b>3-3</b>
<b>Start Page</b> . . . . .	<b>3-4</b>
<b>The Tableau Workspace</b> . . . . .	<b>3-6</b>
Data Window . . . . .	3-7
Toolbar . . . . .	3-8
Status Bar . . . . .	3-10
Cards and Shelves . . . . .	3-11
Reorganizing the Workspace . . . . .	3-15
<b>Workbooks and Sheets</b> . . . . .	<b>3-17</b>
Workbooks . . . . .	3-17
Sheets . . . . .	3-18
<b>Files and Folders</b> . . . . .	<b>3-23</b>
Changing the Repository Location . . . . .	3-25

---

## Overview

This section will introduce you to the Tableau environment including how to open and close the application, the workspace in general, and how your work is organized and stored. Specifically, this section discusses the following topics:

- Opening and Closing the Application
- Start Page
- The Tableau Workspace
- Workbooks and Sheets
- Files and Folders

You may also want to follow the step-by-step tutorial, [Learning to Use Tableau](#) to learn more about how to use the application.



# Opening and Closing the Application

## Open Tableau

There are many ways to open Tableau from your desktop computer.

**To open Tableau, do one of the following:**

- Double-click the Tableau icon on your desktop.
- Select **Start > All Programs > Tableau**.
- Double-click a Tableau workbook or bookmark file. Tableau files are typically stored in the My Tableau Repository folder of your My Documents folder. Refer to “Files and Folders” on page 3-23 to learn more about your repository.
- Drag a data source such as an Excel or Access file onto the Tableau icon or the application window. Tableau automatically makes a connection to the data source.

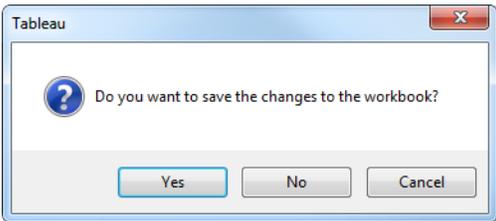
## Close Tableau

When you are done working in Tableau you should save your work and close the application.

**To close Tableau, do one of the following:**

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

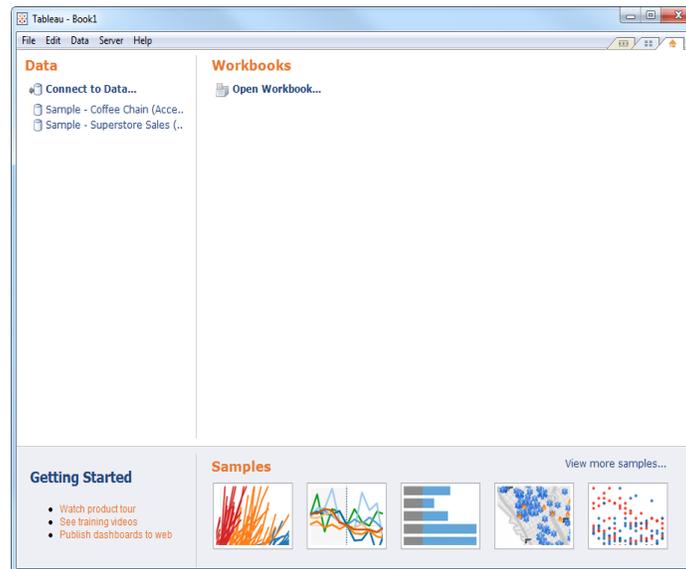
If your workbook has not been saved, you will be asked to whether you want to save with the dialog box shown below.



---

## Start Page

When you first open Tableau Visual Explorer, the Start Page opens to help you get started quickly. The Start Page contains many different resources that are useful whether you are first learning or very experienced.



**To open the start page, do one of the following:**

- Open Tableau by double-clicking the icon on your desktop.
- From an open workbook, click the third tab in the top right corner of the workbook.



The start page is split into four sections: Data, Workbooks, Getting Started, and Samples. Each section is described below.



## **Data**

The data section lets you either connect to a new data source or quickly open a saved connection. By default two sample data source connections are available. As you continue to create and save connections they'll be added to the list. The start page lists any data source connections that are saved in your repository.

## **Workbooks**

The workbooks section shows thumbnail images for workbooks you've recently opened. Click a thumbnail image to open the workbook. When you open Tableau for the first time, the workbooks section will be blank. As you create and save new workbooks, the nine most recently used ones are available on the start page.

## **Getting Started**

The getting started section contains links to resources that can help you learn how to use Tableau. You can watch a short flash video that introduces you to the workspace or browse more in-depth training videos on the Tableau web site. Finally, you can access our online Community Center where you can request support, participate in the forums, and download additional help.

## **Samples**

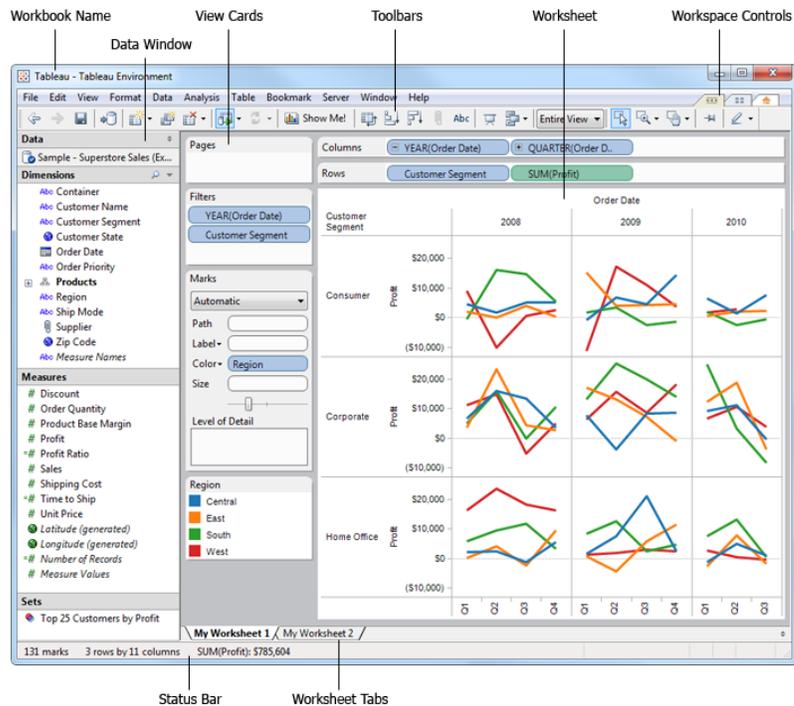
The samples section contains several sample workbooks that show off the kinds of analysis you can do with Tableau. Click on a thumbnail image in the samples areas to open a sample workbook.

## The Tableau Workspace

The Tableau workspace consists of menus, a toolbar, the Data window, cards that contain shelves and legends, and one or more sheets. Sheets can be worksheets or dashboards.

Worksheets contain shelves, which are where you drag data fields to build views. You can change the default layout of the shelves and cards to suit your needs, including resizing, moving, and hiding them.

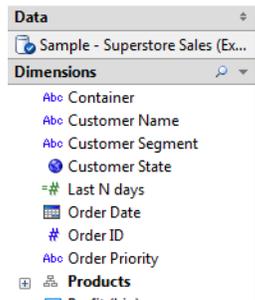
Dashboards contain views, legends, and quick filters. When you first create a dashboard, the Dashboard is empty and all of the worksheets in the workbook are shown in the Dashboard window.





## Data Window

Data fields appear on the left side of the workspace in the Data window. You can hide and show the Data window by selecting **View > Data Window**. 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 clicking the magnifying glass icon at the top of the window. 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 11-3 to learn more about the data window.



## Toolbar

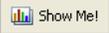
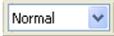
Tableau's toolbar contains commands such as Connect to Data, Show Me!, and Save. In addition, the toolbar contains analysis and navigation tools such as sort, group, and zoom. You can undock the toolbar by grabbing its left edge and then dragging it to a new location. You can hide or display a toolbar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



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. Use the drop-down list to control updates on the entire sheet or just quick filters.
	<b>Run Update:</b> runs a manual query of the data to update the view with changes when automatic updates are turned off. Use the drop-down list to update the entire sheet or just quick filters.



	<b>Show Me!:</b> displays up to twelve alternate views of the data, in addition to the best view according to best practices. 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>Show Mark Labels:</b> toggles between showing and hiding mark labels on the current view.
	<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>Fit Selector:</b> specifies a how the view should be sized within the application window. Select either a Normal fit, Fit Width, Fit Height, or Entire View.
	<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>Zoom:</b> displays greater detail of a specific area of the view. When you zoom in, the axes are locked to only show a portion of the data. Use the drop-down to select a level of magnification.

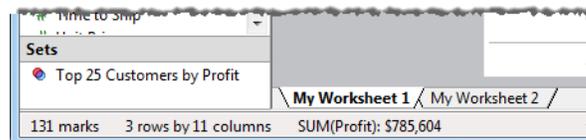


	<b>Pan:</b> moves your view of a table up and down as well as left and right. It is especially useful when looking at maps or when zoomed in on a view. Use the drop-down to incrementally pan.
	<b>Fix Axes:</b> toggles between locking the axes to a specific fixed range and showing all of the data in the view.
	<b>Highlight:</b> turns on highlighting for the selected sheet. Use the options on the menu to define how values will be highlighted. Refer to “Highlight Actions” on page 20-8 to learn more.

## Status Bar

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

An example of information about the view is shown below. In this particular view there are 131 marks shown in 3 rows and 11 columns. It also shows that the SUM(Profit) across the 131 marks is \$785,604.



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



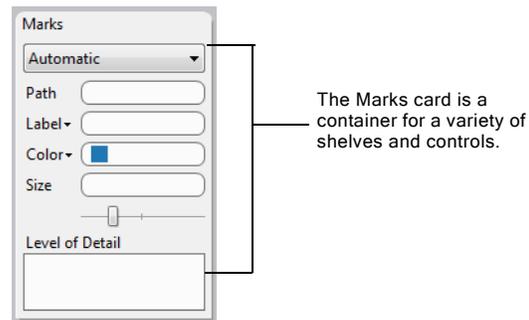
### Status Bar Warnings

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 10-5.</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 10-6.</p> <p><b>Geocoding Warning:</b> If Tableau cannot geocode some of your location values this warning will show. Geocoding warnings may happen if you have unknown location names or names that exist in multiple countries and states.</p>

### Cards and Shelves

Every worksheet contains a variety of different cards that you can show or hide. Cards are containers for shelves, legends, and other controls. For example, the Marks card contains the mark selector, the size slider, the mark transparency control, and the shape, text, color, size, angle, and level of detail shelves.



Cards can be shown and hidden as well as rearranged around the worksheet. Refer to “Reorganizing the Workspace” on page 3-15 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 13-10.

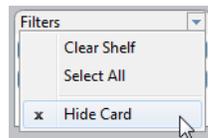
- **Columns Shelf** – contains the Columns shelf where you can drag fields to add columns to the view.
- **Rows Shelf** – contains the Rows shelf where you can drag fields to add rows to the view.
- **Pages Shelf**– contains the Pages shelf where you can create several different pages with respect to 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.
- **Map Legend** - contains the legend for the symbols and patterns on a map. The map legend is not available for all map providers.



- **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, Angle, 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.
- **Map Options** - allows you to modify the various labels and boundaries shown in the online maps. Also you can use this card to overlay metro statistical area information.
- **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.





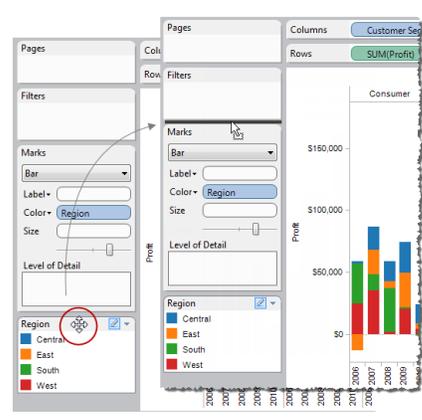
## Reorganizing the Workspace

You can rearrange and hide cards, toolbars, legends, shelves, and more.

### Rearranging Cards

A worksheet contains several cards that contain shelves, legends, and other controls (refer to “Cards and Shelves” on page 3-11). 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.



---

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

---

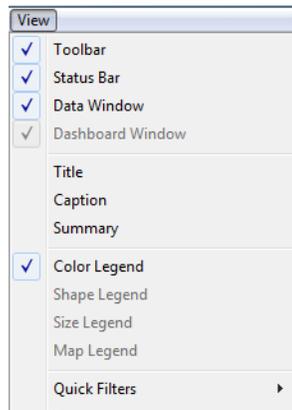
### Showing and hiding parts of the workspace

Just about everything in the workspace can be turned on and off so you don't have to clutter the worksheet with unnecessary cards, shelves, etc.

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



- Select **View** and then select what 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 card or shelf one at a time, you can switch into Presentation Mode. Presentation Mode hides everything on the sheet except for the view and its associated legends and quick filters.

**To toggle in and out of 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.
- Sheets – 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.

### 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.

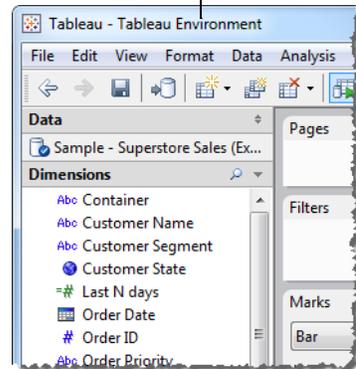
#### To open an existing workbook:

- Double-click the thumbnail image of the workbook on the start page. The start page shows workbooks that you've recently used.
- Select **File > Open** and navigate to the location of your workbook using the **Open** dialog box. Tableau workbooks have the **.twb** or **.twbx** file extensions.
- 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.



Workbook name



You can open multiple workbooks simultaneously. Each workbook is shown in its own window. Refer to “Saving and Exporting” on page 31-1 for information about saving a workbook.

## Sheets

Each workbook can contain worksheets and dashboards. A worksheet is where you build views of your data by dragging and dropping fields onto shelves. A dashboard is a combination of several views of data that you can arrange for presentation or to monitor. The sheets, whether worksheets or dashboards, display along the bottom of the workbook as tabs. In this section you’ll learn how to create, open, duplicate, hide, and delete sheets as well as how to organize sheets in a workbook. Specifically this section discusses the following topics:

- Creating New Sheets
- Undo, Redo, and Clearing Sheets
- Duplicating Sheets
- Hiding and Showing Worksheets
- Deleting Sheets
- Organizing Sheets



## Creating New 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.

## 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.

---

## 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 Data (refer to “View Data (Drill-Through)” on page 19-12). Finally, you can copy and paste the data into Excel. Refer to “Exporting Your Work” on page 31-10 to learn more.

### Hiding and Showing 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 29-20 to learn more about navigating between dashboards and worksheets.

You can show a hidden sheet by navigating to the dashboard that uses it. Select **Go to Sheet** on the dashboard view menu. The hidden sheet is shown temporarily next to the dashboard. When you navigate away from the hidden sheet, it is hidden again. Right-click the sheet tab of the hidden sheet and select **Unhide** to unhide it permanently.

### 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 and Showing Worksheets” on page 3-20 to learn more.

---

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

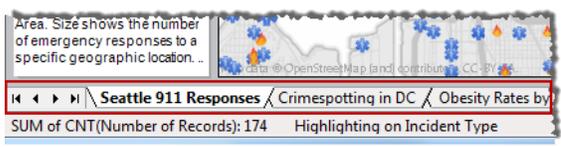
---

### Organizing Sheets

There are three ways to navigate and view the sheets in a workbook: the tabs at the bottom of the workbook, the filmstrip, and the sheet sorter. The tabs are useful for quickly navigating between a small number of sheets. If your workbook has a large number of sheets, you may find that the sheet sorter makes it easier to navigate them all.



**Sheet tabs.** Each sheet, whether worksheet or dashboard, is represented as a tab along the bottom of the workbook. Simply select the tab for the sheet you want to show in the workspace. On the left side of the tabs there are several controls that you can use to advance through each sheet or quickly jump to the first or last sheet in the workbook.

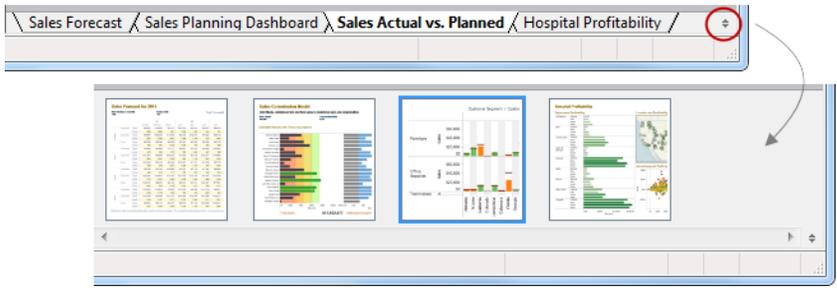


You can also navigate between sheets using the **Window** menu, or move through multiple worksheets by pressing **Ctrl+F6**.

You can also right-click these tabs to specify commands that apply to the entire selected sheet. For example you can create new sheets, duplicate sheets, copy formatting, and delete the sheet entirely. Finally, you can hold the control key when selecting to select and apply settings to multiple sheets all at once.

**Filmstrip.** Similar to the sheet tabs, the mini sheet sorter displays along the bottom of the workbook. However, instead of sheet names, the filmstrip shows a thumbnail image of each sheet. The filmstrip is useful when you are using Tableau to present your analysis and works well when you are working Presentation mode.

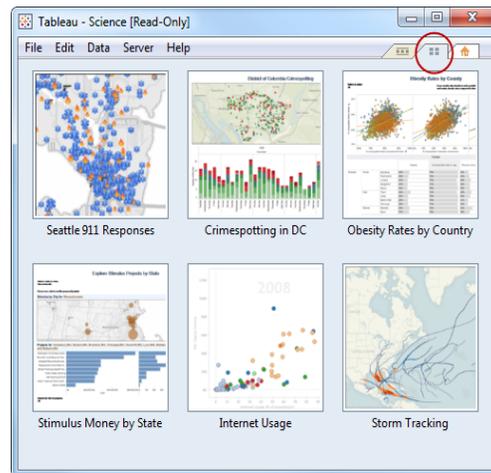
Open the filmstrip by clicking the arrows on the far right side of the sheet tabs at the bottom of the workbook.



Just like with the tabs, select the thumbnail image for the sheet you want to show in the workspace. You can right-click the images to specify commands that apply to each sheet.



**Sheet sorter.** The full sheet sorter shows all the sheets in a workbook as thumbnail images on a single page and is similar to the slide sorter in Microsoft Power Point. The sheet sorter is really useful when you have a large number of sheets in a workbook. Open the sheet sorter by clicking the sheet sorter tab in the upper right corner of the workbook.



From the sheet sorter you can drag and drop to reorder the sheets, create new sheets, and duplicate or delete existing sheets.



## 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 31-3 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 31-7 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 31-5 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 to the Data Engine” on page 9-1 to learn more about extracting data.

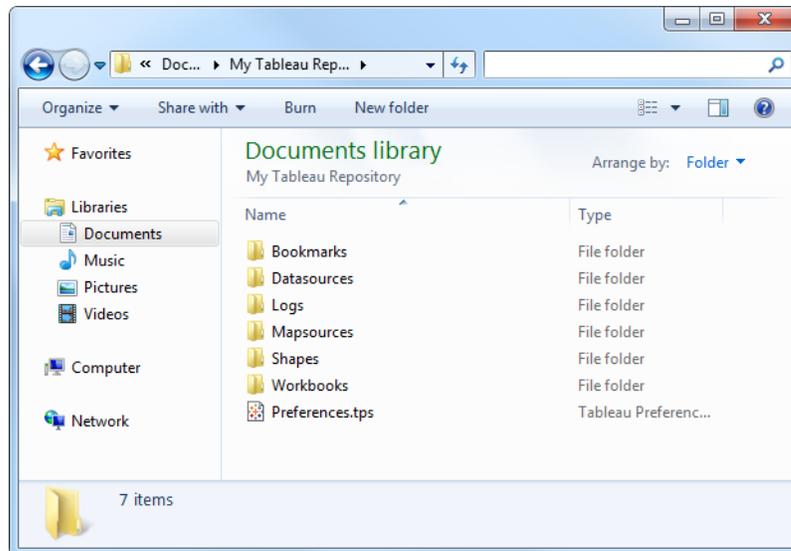
### Data Connection Files

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 “Replacing Field References” on page 5-71 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>4-2</b>
<b>Open Tableau</b> . . . . .	<b>4-3</b>
<b>Connect to Data</b> . . . . .	<b>4-4</b>
<b>Build Data Views (Manual)</b> . . . . .	<b>4-8</b>
Example 1 - Yearly Profits . . . . .	4-9
Example 2 - Quarterly Profits by Year . . . . .	4-14
Example 3 - Quarterly Profits by Year and Market . . . . .	4-16
Example 4 - Filter by Year and by Sales . . . . .	4-18
Example 5 - Color Encode by Region . . . . .	4-24
<b>Build Data Views (Automatic “Show Me!”)</b> . . . . .	<b>4-26</b>
Yearly Profits . . . . .	4-26
Sales versus Profit by Product and Customer . . . . .	4-31
<b>Save Your Work</b> . . . . .	<b>4-34</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.

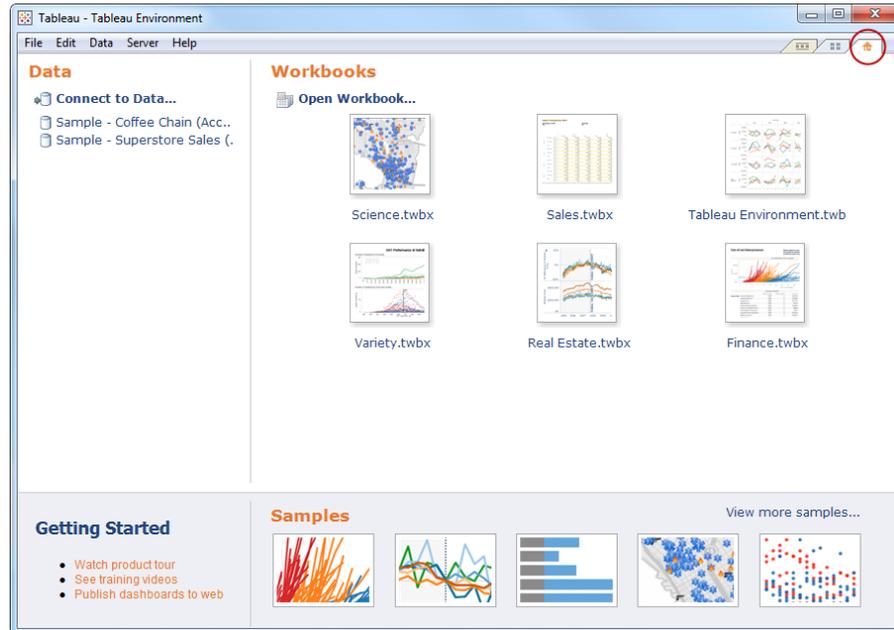
After reading this section, you will be able to build data views using your own relational or multidimensional data.

---

## Open Tableau

To open Tableau, on the Windows **Start** menu select **All Programs > Tableau Software > Tableau 6.0**, or double-click the desktop icon.

Tableau opens showing the start page. The start page contains recently used workbooks, saved data connections, sample workbooks, and some other getting started resources. Although the start page shows when you first open Tableau, you can always return to the start page after you start working by clicking the start page tab in the upper right corner of the workbook.



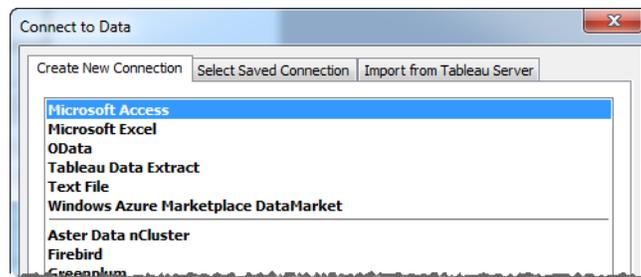
---

## Connect to Data

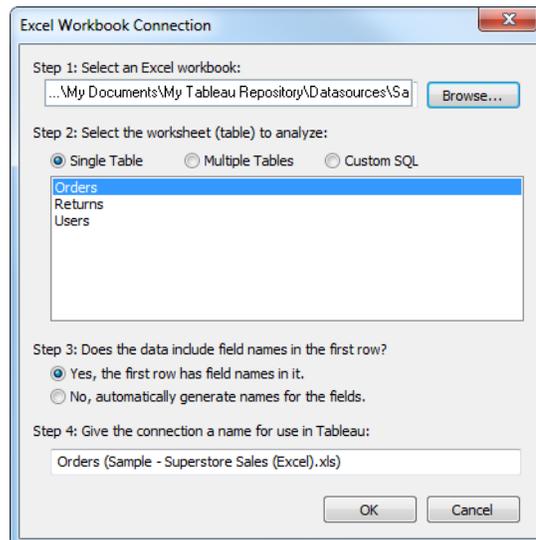
The first step to getting started with Tableau is to connect to the data you want to explore. For this example, you will connect to the Superstore Sales Excel data source that is installed with Tableau. There are many ways to connect to data, for example you can use the start page, the toolbar, or the Data menu. Follow the instructions below to connect to an Excel workbook from the start page.

### To connect to data from the start page:

- 1 Click **Connect to Data** in the Data section of the start page.
- 2 In the subsequent dialog box, select the connection type. For this example, select **Microsoft Excel** and then click **OK**.



- 3 Complete the **Excel Workbook Connection** dialog box. 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.



When you click **OK**, Tableau asks you if you want to do one of the following:

- Connect live – creates a live connection to your data source.
- Import all data – imports the entire data source into Tableau’s fast data engine as an extract. Importing your data can improve performance and add functionality.



- Import some data – allows you to specify a subset of data to import into Tableau’s fast data engine as an extract.



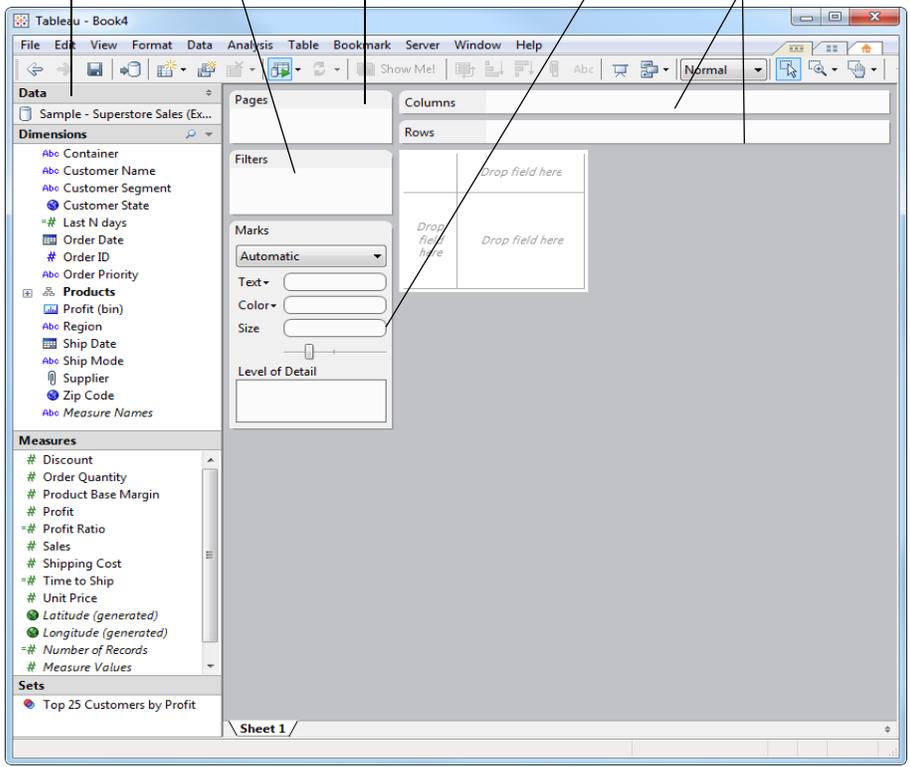


Once connected, the columns from the data source (e.g., Customer Name, Order Date, and Total Sales) are shown on the left side of the workbook in the Data window. Each column is shown as a separate field that you can drag and drop to start exploring your data.

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.

An important concept to understand about Tableau is that you can build views of your data by dragging fields from the data window to the shelves in the view. For example, below you can see the fields from the sample data source as well as the shelves you'll use 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. You will build the following views in this section:

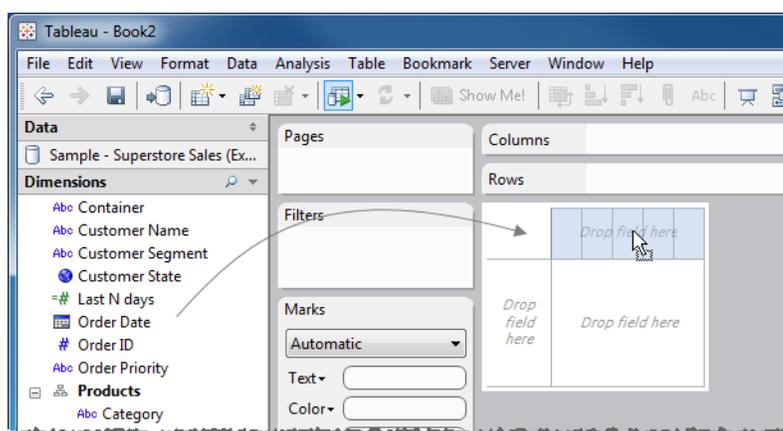
- Example 1 - Yearly Profits
- Example 2 - Quarterly Profits by Year
- Example 3 - Quarterly Profits by Year and Market
- Example 4 - Filter by Year and by Sales
- Example 5 - Color Encode by Region



## Example 1 - Yearly Profits

The first view you will build is a bar chart that displays 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 and drop it on the columns area of the workspace. As you drag fields over different areas in the view the transparent image changes into a visual representation of how the field will be added to the view if you drop it on the given area.



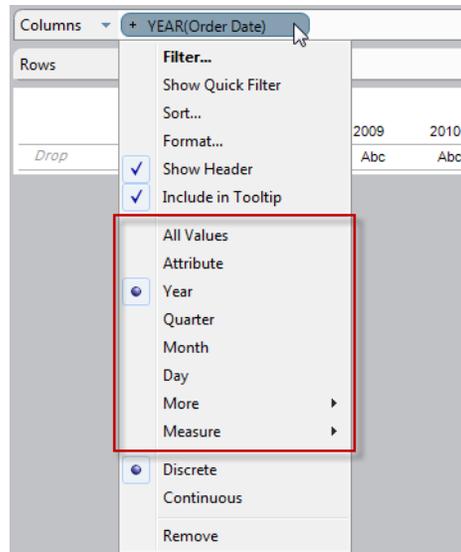
When you place a field in the columns area it is also added to the Columns shelf. You can also drag directly to these shelves. When you drag a field over a shelf, the  icon indicates that the shelf can accept the field.



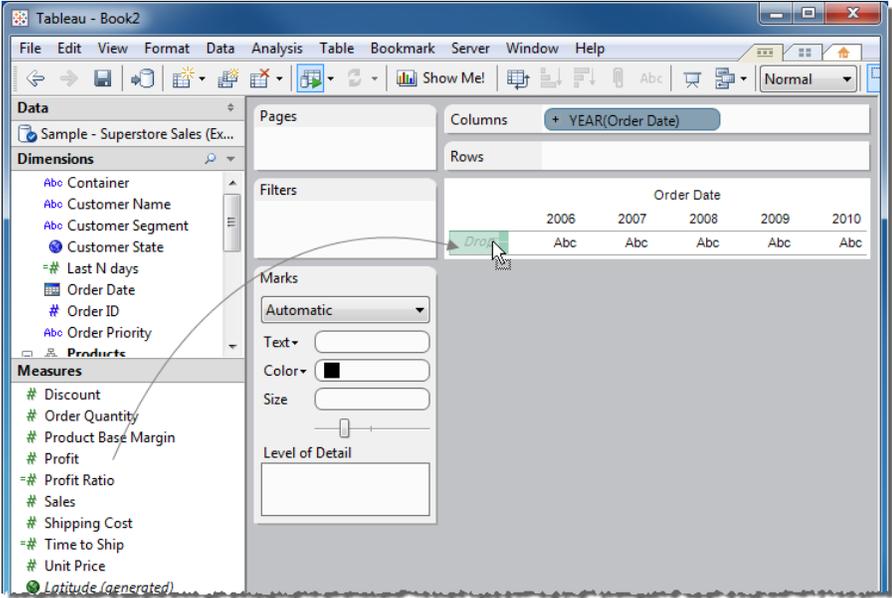
The resulting table has four columns and one row. Each column header represents a member of the Order Date field (2007, 2008, 2009, 2010). Each cell contains the label *Abc*, which indicates that the current mark type for this view is text.

Columns					
+ YEAR(Order Date)					
Rows					
Order Date					
	2006	2007	2008	2009	2010
<i>Drop</i>	Abc	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 menu.

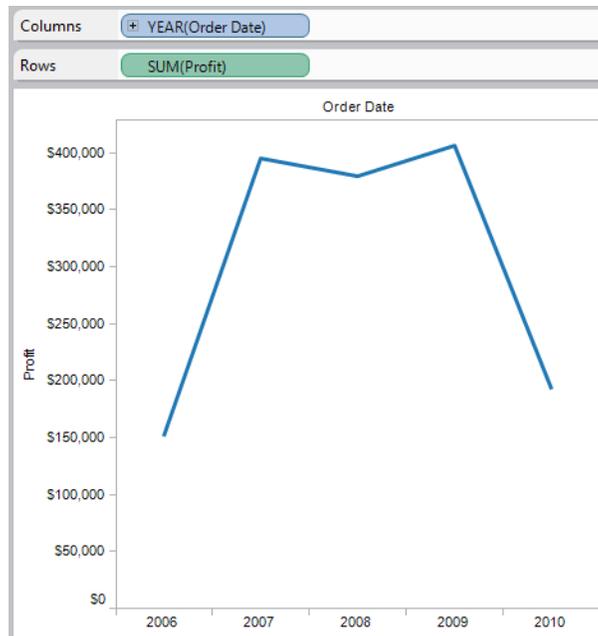


- 2 Click and drag the **Profit** field from the Data window and drop it on the rows area of the workspace.



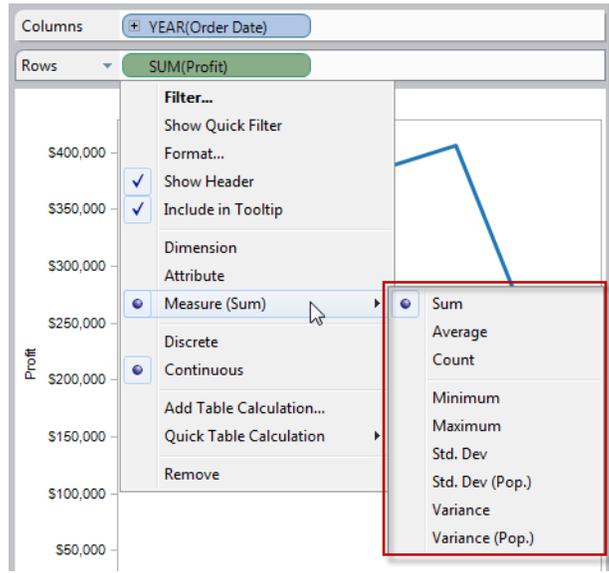


The table is automatically transformed into a line chart and a vertical axis is created for the measure.





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



The data are automatically displayed using lines because lines are a great way to compare data over time, and allow you to visually compare data and identify trends effectively.

Each bar represents the summation of the profit during the corresponding year. For example, Tableau identifies all the data rows associated with the year 2006 and then sums all the corresponding 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.



## Example 2 - Quarterly Profits by Year

In this example, you'll modify the view created in "Yearly Profits" on page 4-26 to include Quarter date level nested within the Years.

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

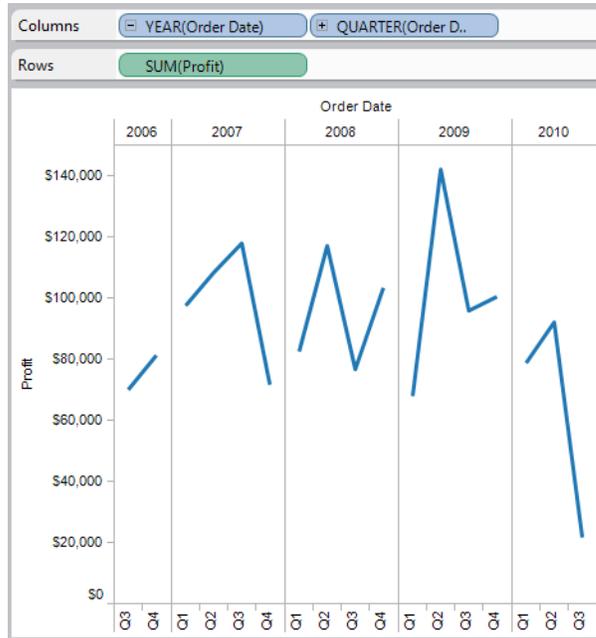


The other way is to drag the **Order Date** field to the **Columns** shelf again. Tableau automatically displays the dates using the next date level, which is quarters in this example.





The view is shown below.



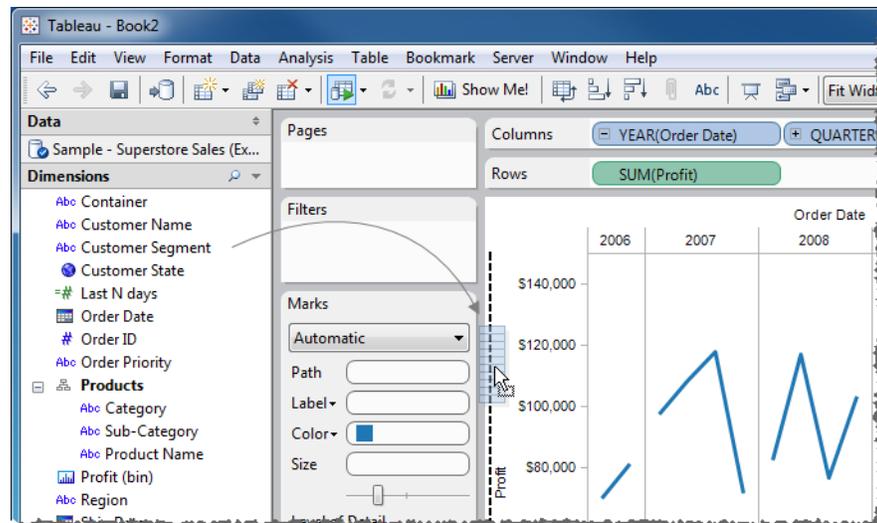
The new dimension divides the view into panes, where each pane contains the four quarters for a given year. This view is called a nested table because it displays multiple headers, with quarters nested within years.



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

In this example, you'll modify the view created in "Example 2 - Quarterly Profits by Year" on page 4-14 to include an additional dimension to break profit down by Customer Segment and Date.

- 1 Click and drag the **Customer Segment** dimension from the Data window and drop it to the left side of the Profit axis.



The field is added to the rows shelf and row headers are created. Each header represents a member of the Customer Segment field.

Tableau does not allow you to place a dimension to the right of a measure on either the **Rows** or **Columns** shelves.

The view is shown below.



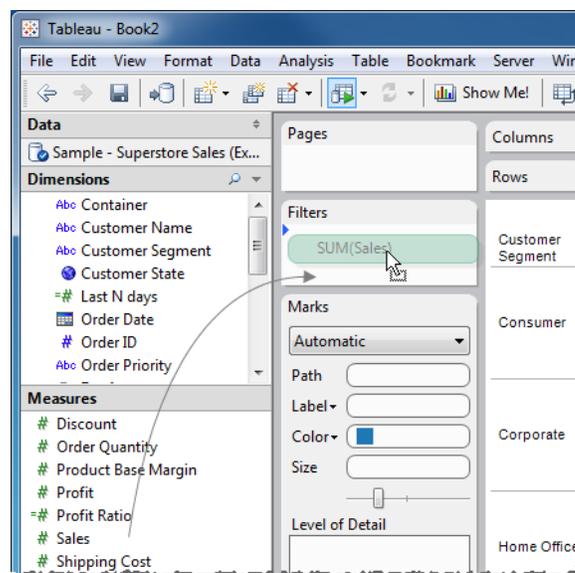
The new dimension divides the view into 20 panes: one for each area in the table where the year and customer segments intersect. This view is a more complex example of a nested table because it displays multiple row and column headers.



## Example 4 - Filter by Year and by Sales

In this example, you will modify the view created in “Example 3 - Quarterly Profits by Year and Market” on page 4-16 to only show profit for orders placed in 2009 or 2010 and have sales values greater than \$10,000.

- 1 Click and drag the **Sales** measure from the Data window and drop it 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 21-34 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).



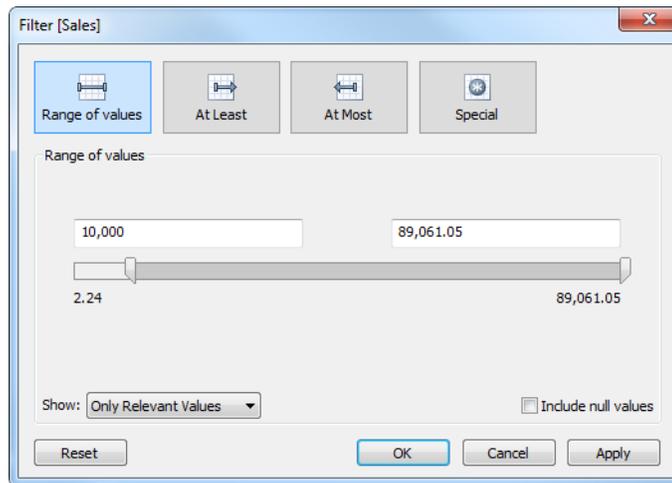
Filtering aggregated data means that the selected aggregation function (sum, average, and so on) is applied to the data and then it is filtered. Filtering disaggregated data means filtering individual data rows before any aggregation function is applied. In this example, filter **Sales** as a disaggregated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

- 2 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



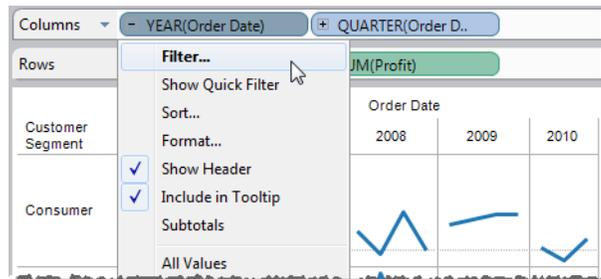
been filtering aggregated data, the limits would have been based on the selected aggregation.

Filter the data by typing a new value in the edit boxes or by moving the sliders. Type in 10,000 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 “Groups” on page 17-18 to learn more about filtering data.

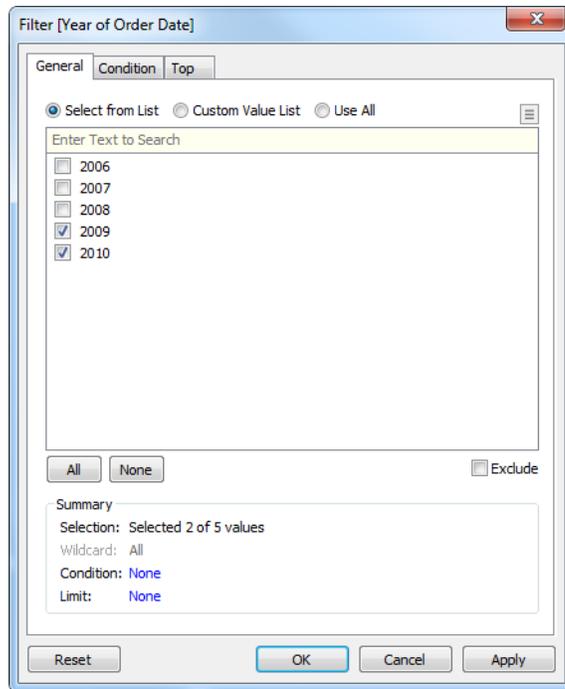
- 
- 3 Open the **Filter** dialog box by selecting **Filter** from the field menu for the Year(Order Date) field on the Columns shelf.. You can also open the dialog box by double-clicking on the field when it is on a shelf.



The dialog box below displays the dimension member names. By default, all members are selected. Deselect all years except 2009 and 2010 and click **OK**. Note that when a field is filtered, it is automatically placed on the **Filters** shelf. Additionally, the filter removes

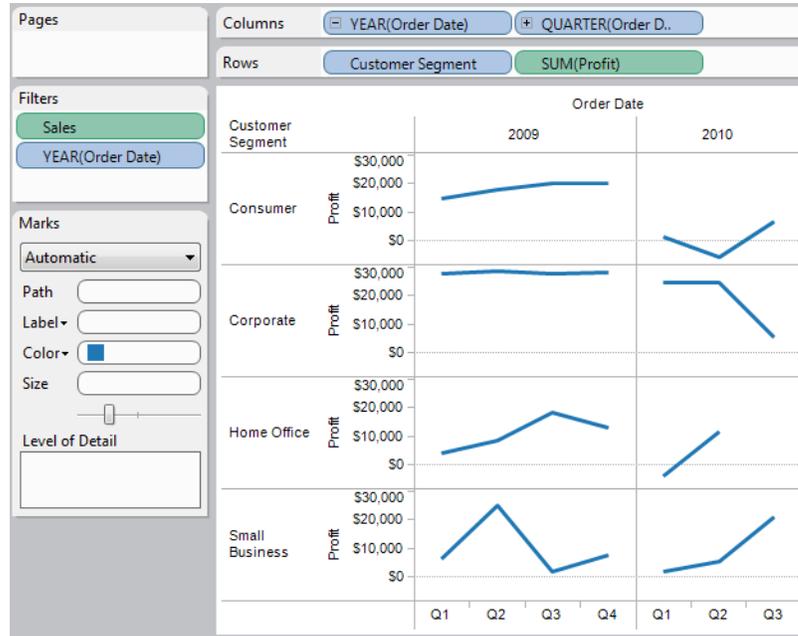


only the specified dimension members from the view. It does not modify the underlying data in any way.





The view is shown below.



The profit sum is calculated by including only the individual data rows that have sales greater than or equal to \$10,000. These data are then displayed by customer segment and by quarter for the years 2009 and 2010.

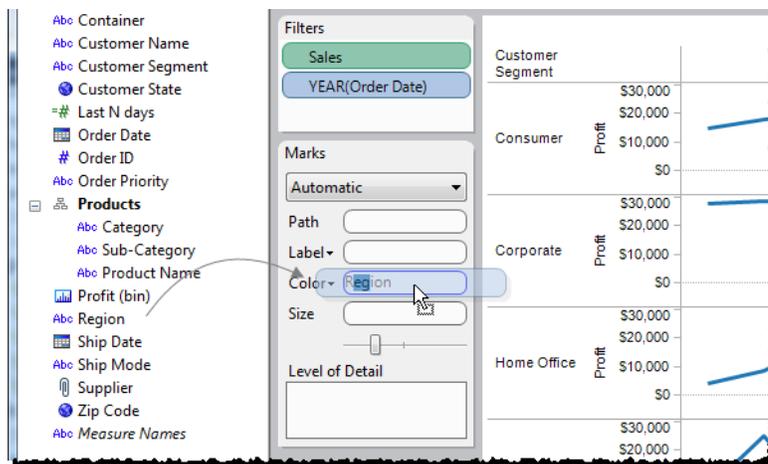
Note that you can quickly switch between the filtered and the unfiltered view by clicking the undo and redo buttons  on the toolbar. The undo/redo feature maintains a history of all the actions across the entire workbook, and provides an easy way to browse all the data views you created.



## Example 5 - Color Encode by Region

In this example you will modify the view created in “Example 4 - Filter by Year and by Sales” on page 4-18 to color the lines by Region.

- 1 Click and drag the **Region** dimension from the Data window and drop it on the Color shelf.

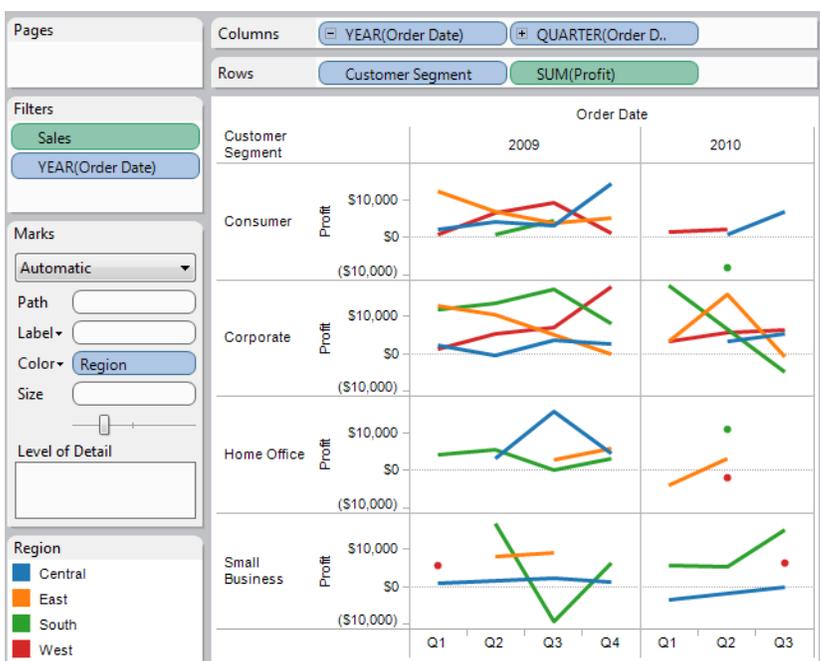


Placing a dimension on the **Color** shelf separates the marks according to the members in the dimension, and assigns a unique color to each member. The color legend displays each member name and its associated color.

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 pane now contains lines for each region and are colored accordingly.



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

In addition to creating views manually, you can use Show Me! to create views automatically. A complete description of Show Me! can be found in the section on “Building Views Automatically” on page 14-1.

Here are a few quick examples of building views automatically.

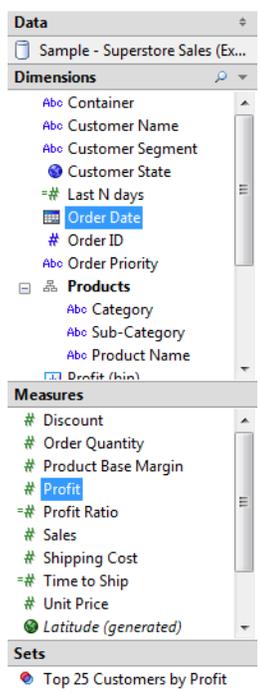
- Yearly Profits
- Sales vs. Profit by Product and Customer

### Yearly Profits

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



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



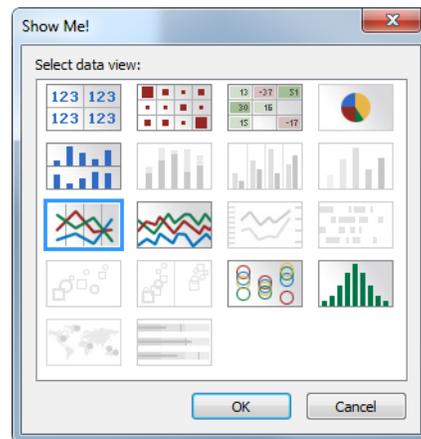
- 2 Click **Show Me!** on the toolbar.





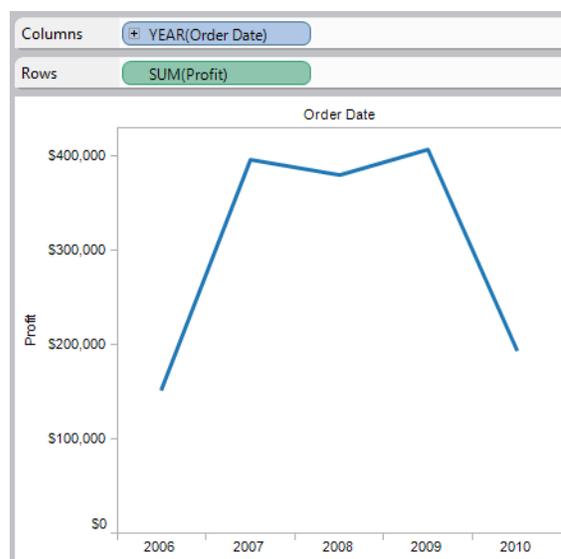
- 3 Select a type of view in the Show Me! dialog box.

Because you selected a date dimension and a measure, Tableau suggests you show the data as a discrete line view, which is generally the best way to look at measures over time.

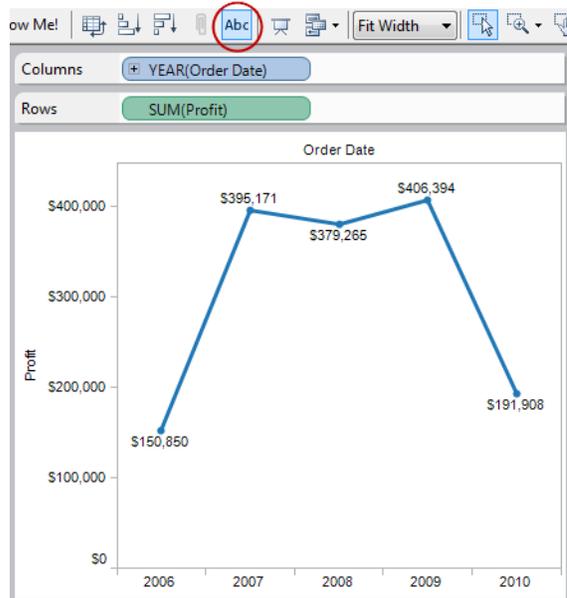




The final view is shown below.



You can also display the values along the line by turning on mark labels by clicking the **Mark Labels** button on the toolbar.

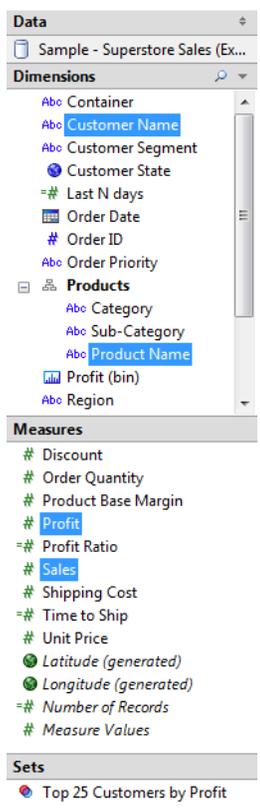




## Sales versus Profit by Product and Customer

For this example you'll create a scatter view using Show Me! that shows sales versus profit for each product and customer.

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

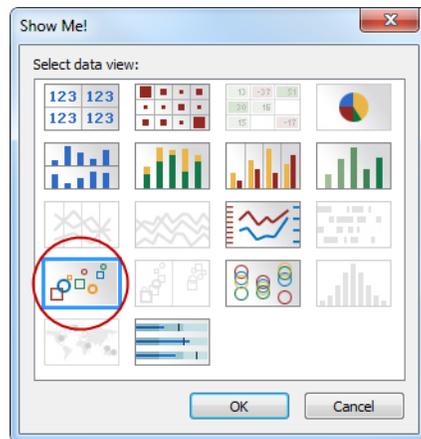


- 2 Click **Show Me!** on the toolbar.

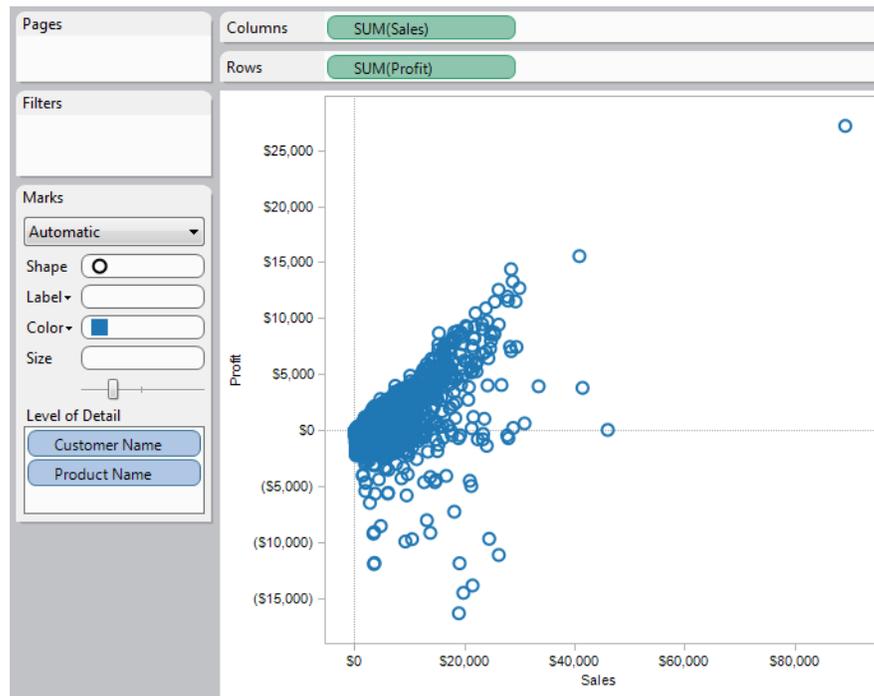




- 3 Select the scatter view type in the Show Me! dialog box.



The resulting view is shown below.



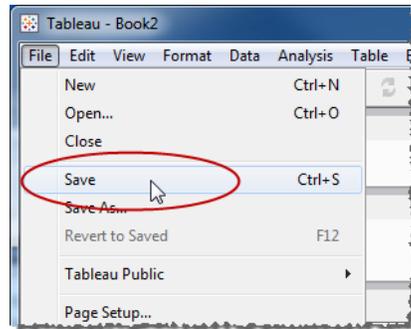
As you can see, Show Me! automatically creates a scatter view based on the fields you selected. You can now manually interact with this automatically-generated view refine it and answer questions.

---

## 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.

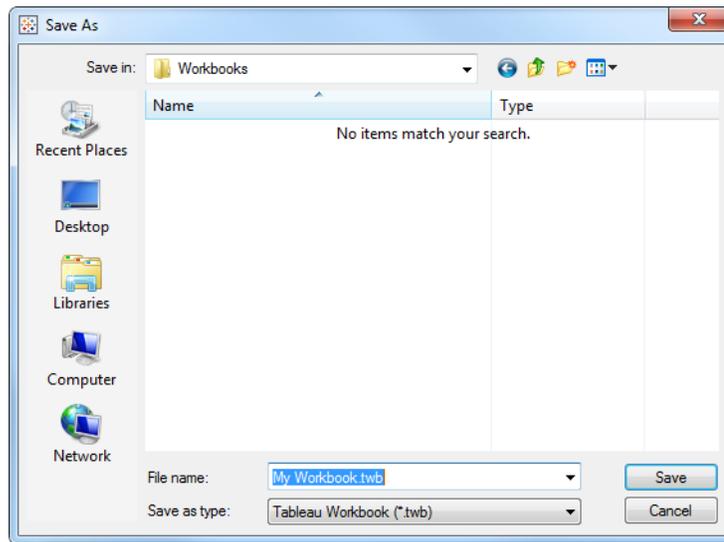
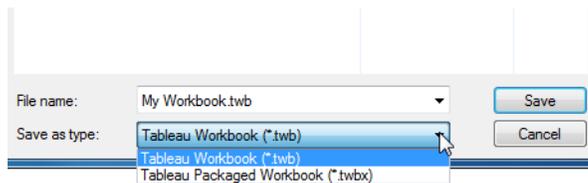


Tableau saves the workbook file with the .twb extension.



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 31, “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 31-5 to learn more.

---

## Basic Connection

---

<b>Supported Data Sources</b> . . . . .	<b>5-3</b>
<b>How to Connect to a Data Source</b> . . . . .	<b>5-4</b>
Connect Using the Open Dialog Box . . . . .	5-4
Connect Using the Connect to Data Dialog Box . . . . .	5-5
<b>Examples – Connecting to Data Sources</b> . . . . .	<b>5-8</b>
Microsoft Access Database . . . . .	5-9
Microsoft Excel Workbook . . . . .	5-11
OData . . . . .	5-14
Tableau Data Extract File . . . . .	5-16
Text File . . . . .	5-17
Windows Azure Marketplace DataMarket . . . . .	5-20
Firebird Database . . . . .	5-22
Greenplum . . . . .	5-24
IBM DB2 . . . . .	5-27
InterSystems Caché . . . . .	5-30
Microsoft Analysis Services Cube . . . . .	5-33
Microsoft PowerPivot . . . . .	5-36
Microsoft SQL Server Database . . . . .	5-38
MySQL Database . . . . .	5-41
Netezza . . . . .	5-43
Oracle Database . . . . .	5-46
Oracle Essbase Cube (aka IBM OLAP Server) . . . . .	5-49
ParAccel . . . . .	5-52
PostgreSQL Database . . . . .	5-54
Sybase IQ . . . . .	5-57
Teradata . . . . .	5-60
Vertica . . . . .	5-64
<b>Connecting to a Custom SQL Query</b> . . . . .	<b>5-68</b>
<b>Editing the Connection</b> . . . . .	<b>5-69</b>
<b>Renaming the Connection</b> . . . . .	<b>5-72</b>



## Overview

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
- Renaming the Connection
- Duplicating a Connection
- Replacing Field References
- Refreshing the Data
- Closing the Connection



## Supported Data Sources

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

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.

---

## How to Connect to a Data Source

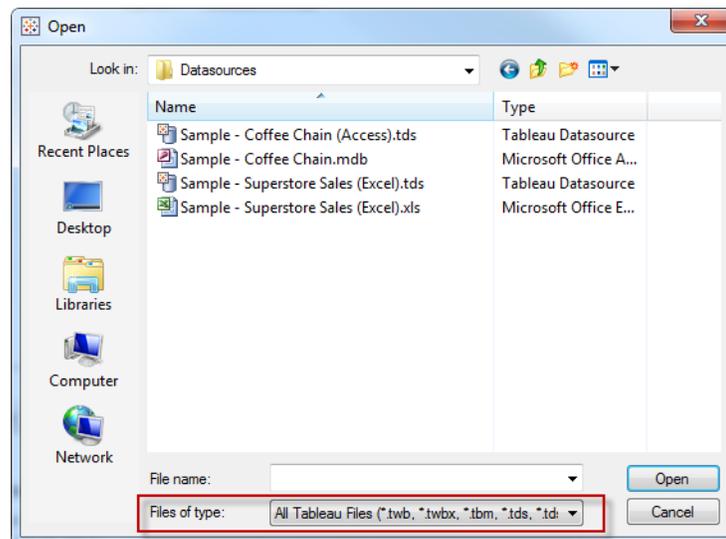
To build views of your data, you must first connect Tableau to a data source. Depending on the type of data you are connecting to, you can connect in one of two ways: with the **Open** dialog box or with the **Connect to Data Source** dialog box.

### Connect Using the Open 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, and delimited text files using the **Open** dialog box.

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

- 1 Select **File > Open**.



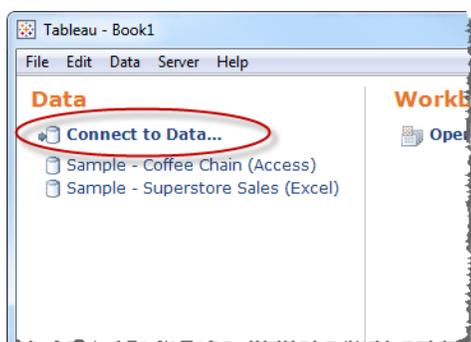
- 2 Locate and select the data file and click **Open**.
- 3 If you select a data source file, Tableau opens the appropriate connection dialog box. Complete the connection dialog box and click **OK**. Refer to “Examples – Connecting to Data Sources” on page 5-8 for information about connecting to specific types of data.

## Connect Using the Connect to Data Dialog Box

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

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

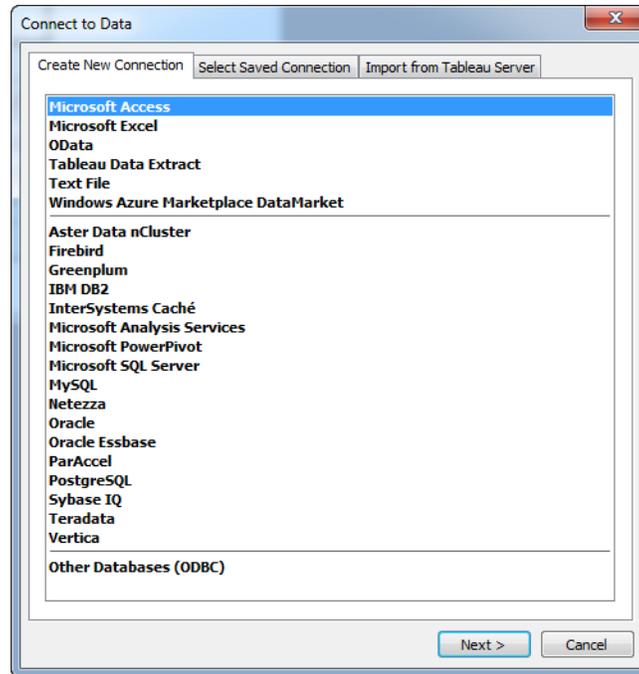
- 1 Select **Data > Connect to Data** or press **Ctrl + D** on your keyboard. You can also select the **Connect to Data** option on the start page.



- 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) or import from Tableau Server. Refer



to “In the Replace References dialog box, select a field from the new data source that corresponds to the invalid field.” on page 5-71 to learn how to save a data connection.



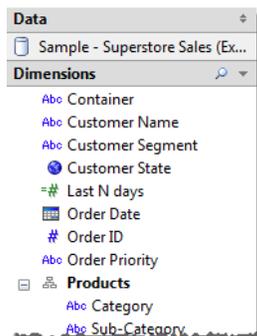
- 3 A data source-specific dialog box opens that allows you to complete the connection process. Refer to “Examples – Connecting to Data Sources” on page 5-8 for more information.

---

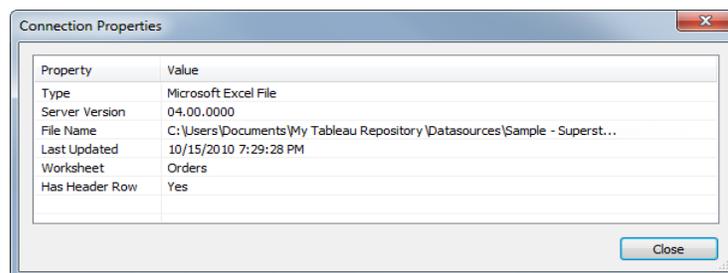
**Note** Another way to connect to data is to import from a workbook. A workbook can contain multiple worksheets with connections to a different data sources. To import a connection from a workbook click the **Import from Workbook** button at the bottom of the **Select Saved Connection** tab in the Connect to Data dialog box.

---

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 11-3 to learn more.



You can display information about the connection by selecting **Data > 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 sources supported by Tableau:

- Microsoft Access Database
- Microsoft Excel Workbook
- OData
- Tableau Data Extract File
- Text File
- Windows Azure Marketplace DataMarket
- Firebird Database
- Greenplum
- IBM DB2
- InterSystems Caché
- Microsoft Analysis Services Cube
- Microsoft PowerPivot
- Microsoft SQL Server Database
- MySQL Database
- Netezza
- Oracle Database
- Oracle Essbase Cube (aka IBM OLAP Server)
- ParAccel
- PostgreSQL Database
- Sybase IQ
- Teradata
- Vertica
- Other Databases (ODBC)

## Microsoft Access Database

This example discusses how to connect Tableau to a Microsoft Access database. Tableau supports all Access data types except OLE Object and Hyperlink.

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Microsoft Access**.
- 3 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

Type the file path navigate to the file by clicking Browse.

Select **Use workgroup security when connecting** when you are connecting to a password protected Access file or a file that is protected by workgroup security. When you select this 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, and Password into the corresponding text fields.

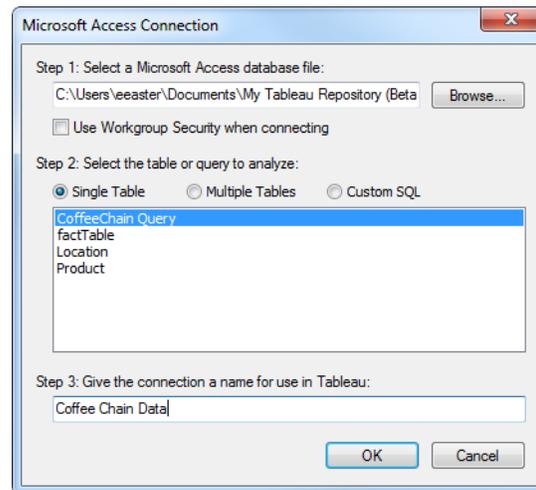
### Step 2: Select the table or query to analyze.

You can connect to a Single Table or query. Alternatively, you can connect to a set of relational tables that are related by join conditions. Select either Multiple Tables or Custom SQL when you are connecting to multiple tables. You can also add joins later. Refer to “Joining Tables” on page 7-1.

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

Specify a unique name for the connection. A default name is automatically generated.

A completed Microsoft Access Connection dialog box is shown below.



---

**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

This example discusses how to connect Tableau to a Microsoft Excel workbook.

- 4 Select **Data** > **Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 5 Select **Microsoft Excel**.
- 6 Follow the steps on the **Excel Workbook Connection** dialog box to complete the connection. These steps are described below. Tableau does not support connecting 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. 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 single 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 **Define Name** on the **Formulas** tab. 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 7-1.

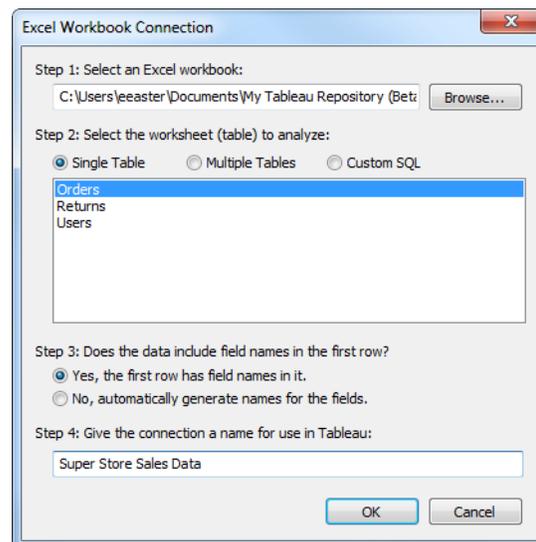
### Step 3: Does the data include field names in the first row

Specify whether the workbook includes column names in the first row. These names will become the names of the fields in Tableau. If column names are not included, they will be automatically generated by Tableau. You can rename the fields later.

### 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.

A completed Microsoft Excel Connection dialog box is shown below.





---

**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.

---



## **OData**

This example discusses how to connect Tableau to an OData data source.

- 7 Select **Data** > **Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 8 Select **OData**.
- 9 Follow the steps on the **OData Connection** dialog box to complete the connection. These steps are described below. Tableau does not support connecting to password protected Excel files.

### **Step 1: Select or enter a URL**

Enter the URL for the OData data source. Click the link to find an OData data source on Windows Azure Marketplace DataMarket.



### Step 2: Enter authentication information

If necessary, enter authentication information. You can authenticate using your Windows Azure Marketplace DataMarket account key or a username and password.

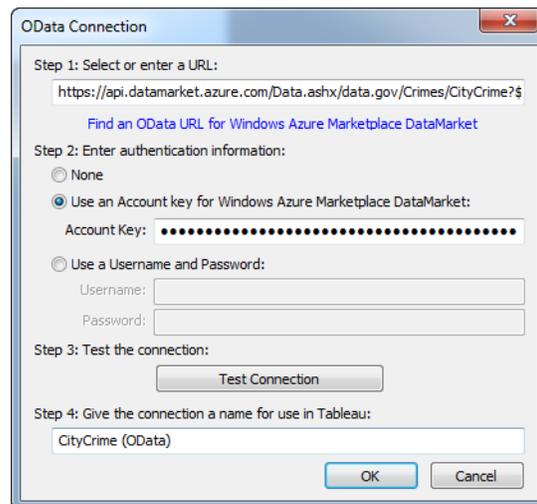
### Step 3: Test the Connection

If the connection is successful, Tableau will alert you. If the connection is unsuccessful, verify that the URL and authentication information are correct. If the connection continues to fail, your computer is having trouble locating the server. Contact your network administrator or database administrator.

### 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.

A completed OData Connection dialog box is shown below.





## Tableau Data Extract File

This example discusses how to connect Tableau to a Tableau Data Extract file. To learn about how to create extracts refer to “Extracting Data to the Data Engine” on page 9-1.

- 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.

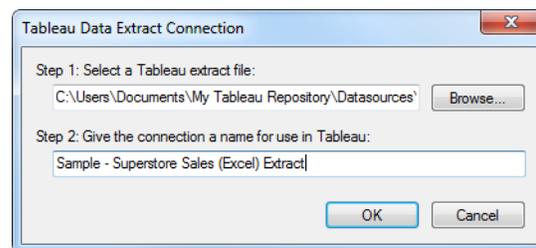
### Step 1: Select a Tableau extract file.

Type the name of a Tableau extract file or click **Browse** to navigate to the file on your computer. Extract files use the **.tde** file extension.

### 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 completed Tableau Data Extract Connection dialog box is shown below.



## Text File

This example discusses how to connect Tableau to a text file. Tableau connects to delimited text files.

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Text File**.
- 3 Follow the steps on the **Text File Connection** dialog box to complete the connection.

### Step 1: Select a delimited text file

Select the text file by typing the file path or by clicking **Browse** to navigate to the file.

### Step 2: Set options for the text file.

There are three options that you can specify to define the type of text file you are connecting to. Specify the following options:

- Select whether the first row contains column names. This option is selected by default. Alternatively you can have Tableau generate names when you connect. These names can be changed later.
- Select the character that is used to separate the columns. Columns can be separated by a Comma, Tab, Space, Colon, Vertical Bar (also known as the pipe character), or another character as defined in a schema.ini file. Refer to “Using Text Files with Alternate Delimiters” on page 5-19 to learn more about using other characters.
- Select a Character Set that describes the text file encoding. You can select ANSI, OEM, UTF-8, UTF-16, or Other. When you select Other you must specify the character set in the provided text field. This value will be verified when the connection is attempted.

### Step 3: Select the data range to analyze

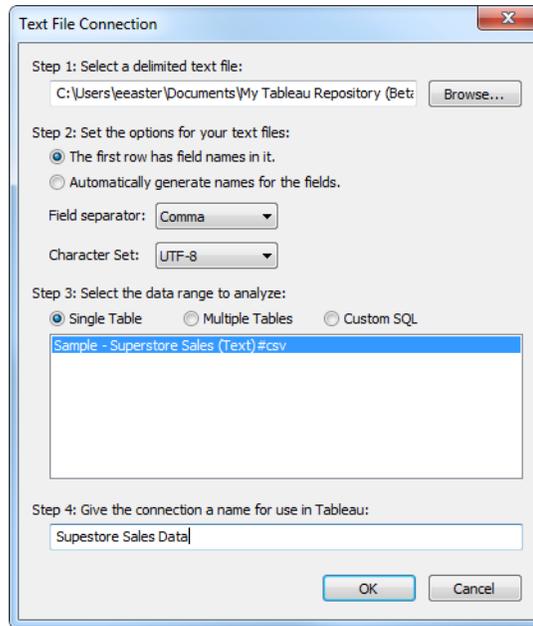
You can connect to a single file or a set of relational tables that are related by join conditions. To do so, select the **Multiple Tables** option. You can also add joins later. Refer to “Joining Tables” on page 7-1.

### 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.



A completed Text File Connection dialog box is shown below.



---

**Note** 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 35-9.

---



---

### Using Text Files with Alternate Delimiters

You can connect to text files that are delimited by commas, tabs, spaces, colons, or vertical bars using the Text File Connection dialog box. If your text file is delimited by an alternate character you must first create a schema.ini file that defines the delimiter. Then you can use the Other option when connecting in Tableau.

#### To create a schema.ini file:

- 1 Create a new text file (using Notepad or another text editor) and type the following:

```
[Your_Datasource.txt]
Format=Delimited(delimiter character)
```

Substitute the name of your data file for “Your\_Datasource.txt.” and the character used to separate the columns for “delimiter character.” For example, if you wanted to use the asterisk character as a delimiter in a file called Sales-Data.txt you would type the following into the schema.ini file.

```
[Sales-Data.txt]
Format=Delimited(*)
```

- 2 Save the file in the same directory as the data file and call it **schema.ini**.
- 3 Return to Tableau and connect to the data file using the Text file instructions outlined above (refer to “Text File” on page 5-17). Select **Other** as the Field Separator.

---

**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.

---



## Windows Azure Marketplace DataMarket

This example discusses how to connect Tableau to a Windows Azure Marketplace DataMarket data source.

- 4 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 5 Select **Windows Azure Marketplace DataMarket**.
- 6 Follow the steps on the **Windows Azure Marketplace DataMarket Connection** dialog box to complete the connection. These steps are described below. Tableau does not support connecting to password protected Excel files.

### Step 1: Select or enter a URL

Enter the URL for the Windows Azure Marketplace DataMarket data source. Click the link to find an OData data source in Windows Azure Marketplace DataMarket.



### Step 2: Enter authentication information

If necessary, enter authentication information. You can authenticate using your Windows Azure Marketplace DataMarket account key or a username and password.

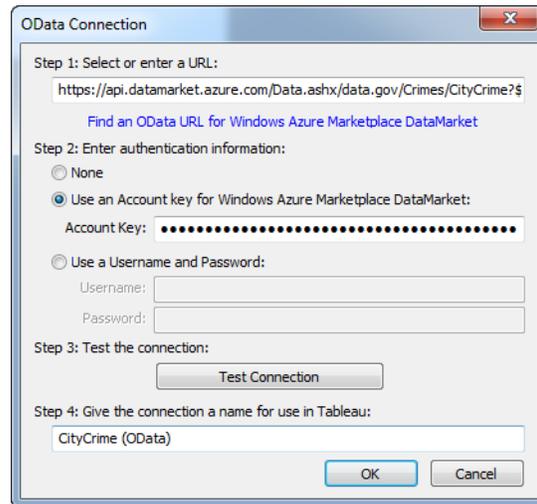
### Step 3: Test the Connection

If the connection is successful, Tableau will alert you. If the connection is unsuccessful, verify that the URL and authentication information are correct. If the connection continues to fail, your computer is having trouble locating the server. Contact your network administrator or database administrator.

### 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.

A completed Windows Azure Marketplace DataMarket Connection dialog box is shown below.





## Firebird Database

This example discusses how to connect Tableau to a Firebird database.

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Firebird**.
- 3 Follow the steps on the **Firebird Connection** dialog box to complete the connection. These steps are described below.

**Step 1: Type in the name of a specific server.**

**Step 2: Specify the location of the database.**

**Step 3: Specify your username and password to log onto the server.**

**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 database administrator.

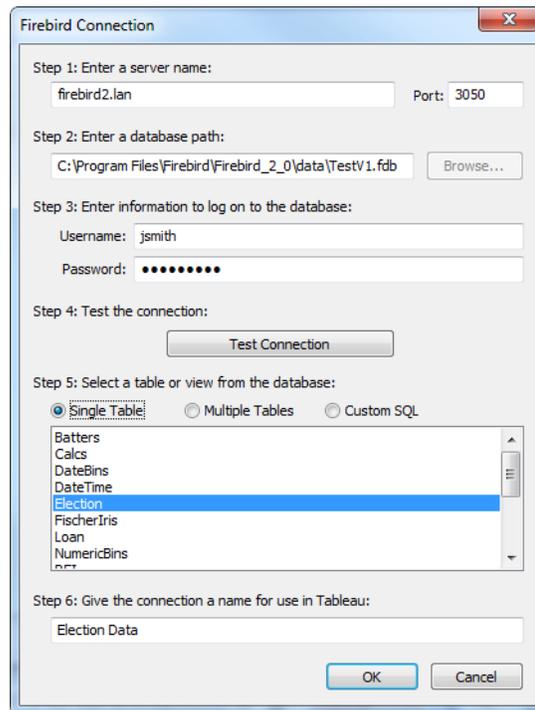
**Step 5: Enter a table or view from the database**

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 7-1 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 completed Firebird Connection dialog box is shown below.



**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 5.0\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.



## Greenplum

This example discusses how to connect Tableau to a Greenplum database.

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Greenplum** and click **OK**.
- 3 Follow the steps on the **Greenplum Connection** dialog box to complete the connection. These steps are described below.



**Step 1: Type in the name of a specific server and the port to which you want to connect.**

**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 3: Enter your username and password to log on to the server**

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

Greenplum 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 7-1 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 completed Greenplum Connection dialog box is shown below.



Greenplum Connection

Step 1: Enter a server name:  
greenplum33.lan Port: 5432

Step 2: Enter a database on the server:  
TestDB

Step 3: Enter information to log on to the database:  
Username: jsmith  
Password: ●●●●●●

Step 4: Test the connection:  
Test Connection

Step 5: Select a table or view from the database:  
 Single Table  Multiple Tables  Custom SQL

- Batters (public.Batters)
- Calcs (public.Calcs)
- DateBins (public.DateBins)
- DateTime (public.DateTime)
- Election (public.Election)
- FischerIris (public.FischerIris)
- Loan (public.Loan)
- NumericBins (public.NumericBins)
- ...

Step 6: Give the connection a name for use in Tableau:  
Election (public.Election) (TestV1)

OK Cancel



## **IBM DB2**

This example discusses how to connect Tableau to an IBM DB2 database. Tableau connects to IBM DB2 databases running on Linux, Windows, and Unix servers. iSeries (AS400 and z/OS) type servers are not supported.

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **IBM DB2** and click **OK**.
- 3 Follow the steps on the **IBM DB2 Connection** dialog box to complete the connection. These steps are described below.



### Step 1: Enter a server name

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

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 3: Enter information to log on to the server

Specify the username and 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 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 7-1 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 completed IBM DB2 Connection dialog box is shown below.



DB2 Connection

Step 1: Enter a server name:  
db2-81.test.tsi.lan Port: 50000

Step 2: Enter a database on the server:  
testv1

Step 3: Enter information to log on to the database:  
Username: test  
Password: ●●●●●●

Step 4: Test the connection:  
Test Connection

Step 5: Select a table or view from the database:  
 Single Table  Multiple Tables  Custom SQL

*#Tableau_2_Connect*
#Tableau_3_test
Batters
Calcs
DateBins
DateTime
Election
FischerIris
...

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

OK Cancel



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

#### Step 1: Enter a server name

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



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

Specify the username and 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 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 7-1 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 completed Intersystems Caché Connection dialog box is shown below.



**InterSystems Caché Connection**

Step 1: Enter a server name:  
cacheserver.test.lan Port: 1972

Step 2: Enter a database on the server:  
testdb

Step 3: Enter information to log on to the database:  
Username: username  
Password: ●●●●●●

Step 4: Test the connection:  
Test Connection

Step 5: Select a table or view from the database:  
 Single Table  Multiple Tables

Election (TESTV1.Election)
FischerIris (TESTV1.FischerIris)
Loan (TESTV1.Loan)
NumericBins (TESTV1.NumericBins)
<b>REI (TESTV1.REI)</b>
SeattleCrime (TESTV1.SeattleCrime)
Securities (TESTV1.Securities)
SpecialData (TESTV1.SpecialData)
SpecialData (TESTV1.SpecialData)

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

OK Cancel



## Microsoft Analysis Services Cube

This example discusses how to connect Tableau to a Microsoft Analysis Services (SSAS) 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.



---

### **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 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 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 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.

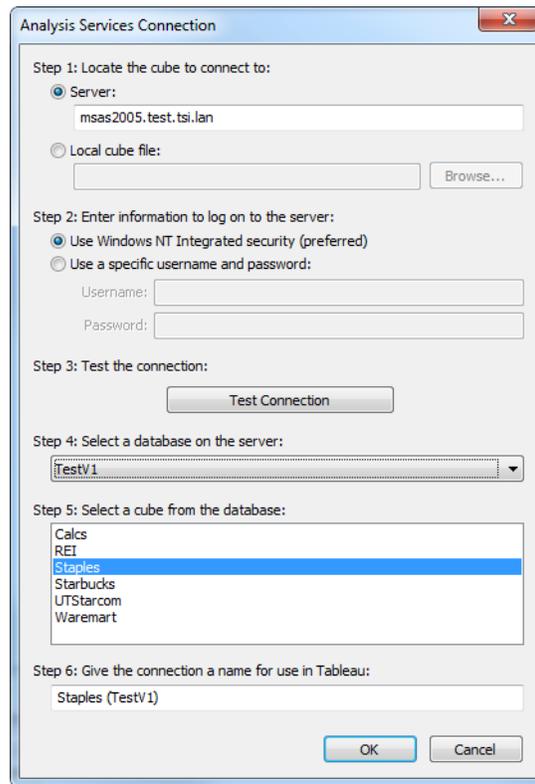
### **Step 5: Select a cube from the database**

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

### **Step 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.

A completed Microsoft Analysis Services Connection dialog box is shown below.



---

**Note** If you are connecting to a Microsoft Analysis Services 2000 data source, you are given the option to enable NonEmptyCrossjoin Optimization. When selected this option ensures that virtual dimensions are shown.

---



## Microsoft PowerPivot

This example discusses how to connect Tableau to PowerPivot Excel files. You can connect to workbooks on your local computer or browse to files stored on SharePoint.

To connect to PowerPivot data:

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



3 Complete the connection.

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

**Step 1: Choose a file on SharePoint or your Local machine.**

If you are connecting to a PowerPivot file on SharePoint you can type the URL or click Browse to navigate to the PowerPivot workbook. If you are connecting to a local file, click Browse to navigate to the local workbook.

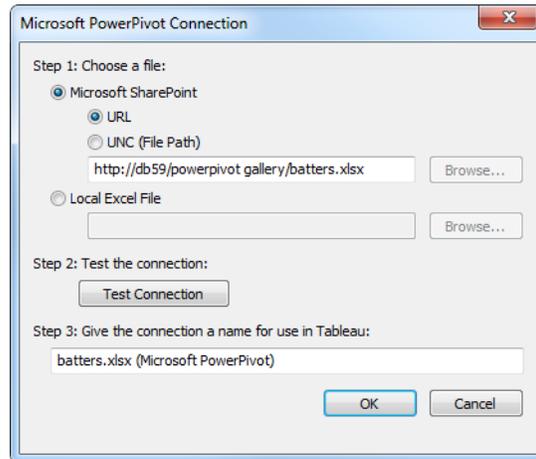
**Step 2: Test the connection**

Click **Test Connection**. If the connection is unsuccessful, verify that you have all the necessary drivers installed and that your SharePoint server is available. If the connection continues to fail, your computer may be having trouble locating the server. Contact your network administrator or database administrator.

**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 database and table/view name.

A completed Microsoft PowerPivot dialog box is shown below.





## Microsoft SQL Server Database

This example discusses how to connect Tableau to a Microsoft SQL Server database. Tableau supports all SQL Server data types except the following: binary objects, varbinary, image, SQL variant, Timestamp, and Unique Identifier.

- 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.



---

### **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.

### **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.

### **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 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.

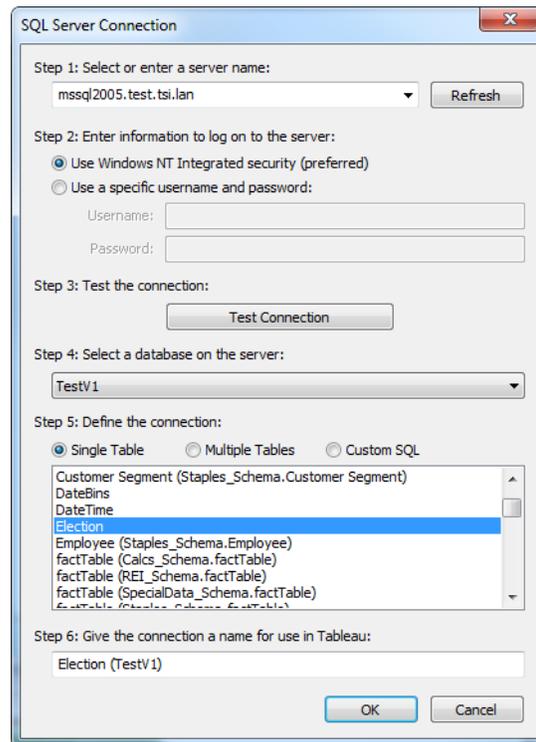
### **Step 5: Select a table or view from the database**

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 7-1 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.

A completed Microsoft SQL Server Connection dialog box is shown below.



## MySQL Database

This example discusses how to connect Tableau to a MySQL database.

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **MySQL**.
- 3 Follow the steps on the **MySQL Connection** dialog box to complete the connection. These steps are described below.



---

**Step 1: Enter a server name**

Type in the name of a specific server.

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

Specify the username and password.

**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.

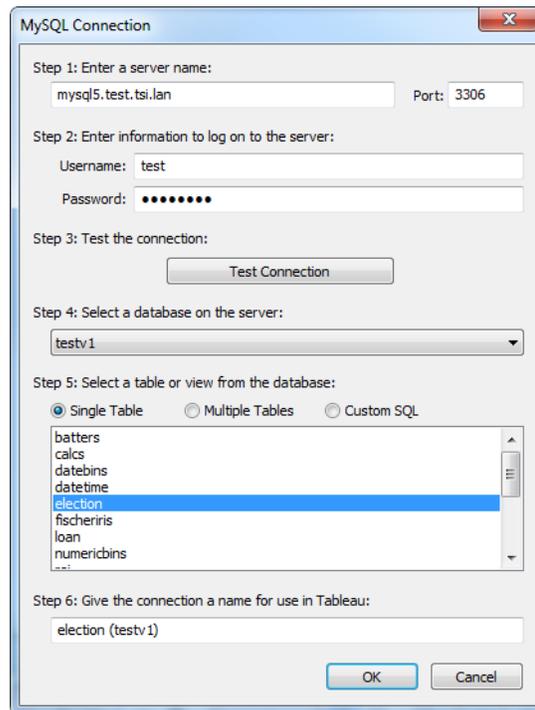
**Step 5: Enter a table or view from the database**

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 7-1 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 completed MySQL Connection dialog box is shown below.



## Netezza

This example discusses how to connect Tableau to a Netezza database.

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

#### 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.



---

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

Specify the username and 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 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 7-1 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 completed Netezza Connection dialog box is shown below.



**Netezza Connection**

Step 1: Enter a server name:  
netezza.server.lan Port: 5480

Step 2: Enter a database on the server:  
SALES

Step 3: Enter information to log on to the database:  
Username: username  
Password: ●●●●●●

Step 4: Test the connection:  
Test Connection

Step 5: Select a table or view from the database:  
 Single Table  Multiple Tables  Custom SQL

Batters
Calcs
DateBins
DateTime
<b>Election</b>
FischerIris
Loan
NumericBins
...

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

OK Cancel



## Oracle Database

This example discusses how to connect Tableau to an Oracle relational database. Tableau supports all Oracle data types except CLOB, NCLOB, binary objects, and database objects.

- 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.

---

**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.

### 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 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 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 7-1 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 or view name.

A completed Oracle Connection dialog box is shown below.



Oracle Connection

Step 1: Enter the Oracle connection name:  
oracle92.test.tsi.lan/ord9i.tsi Advanced...

Step 2: Enter information to log on to the server:  
 Use Windows NT Integrated security  
 Use a specific username and password:  
Username: test  
Password: ●●●●●●

Step 3: Test the connection:  
Test Connection

Step 4: Select a schema on the server:  
TestV1

Step 5: Select a table or view from the schema:  
 Single Table  Multiple Tables  Custom SQL

- Batters (TestV1.Batters)
- Calcs (TestV1.Calcs)
- DateBins (TestV1.DateBins)
- DateTime (TestV1.DateTime)
- Election (TestV1.Election)**
- FischerIris (TestV1.FischerIris)
- Loan (TestV1.Loan)
- NumericBins (TestV1.NumericBins)

Step 6: Give the connection a name for use in Tableau:  
Election (TestV1.Election) (TestV1)

OK Cancel

## Oracle Essbase Cube (aka IBM OLAP Server)

This example discusses how to connect Tableau to an Oracle Essbase cube (also known as IBM OLAP Server).

- 1 Select **Data > Connect to Data**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Oracle Essbase**.
- 3 Follow the steps on the **Oracle Essbase Connection** dialog box to complete the connection. These steps are described below.

### Step 1: Enter a server name

Type the name of the server.

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

Specify the 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.

### 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.

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

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

### 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.

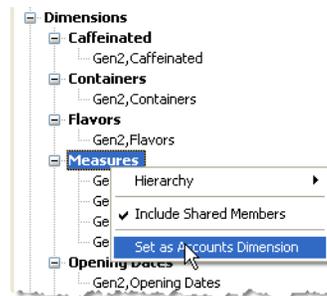
### Step 7: Set an Accounts Dimension (Optional)

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.

A completed Oracle Essbase Connection dialog box is shown below.



Oracle Essbase Connection

Step 1: Enter a server name:  
essbase901.test.tsi.lan

Step 2: Enter information to log on to the server:  
Username: admin  
Password: ●●●●●●

Step 3: Test the connection:  
Test Connection

Step 4: Select an application on the server:  
Sample

Step 5: Select a database from the application:  
Basic

Step 6: Give the connection a name for use in Tableau:  
Basic (Sample)

OK Cancel



## ParAccel

This example discusses how to connect Tableau to a ParAccel database.

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **ParAccel** and click **OK**.
- 3 Follow the steps on the **Greenplum Connection** dialog box to complete the connection. These steps are described below.



**Step 1: Enter a server name**

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

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

Specify the username and 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.

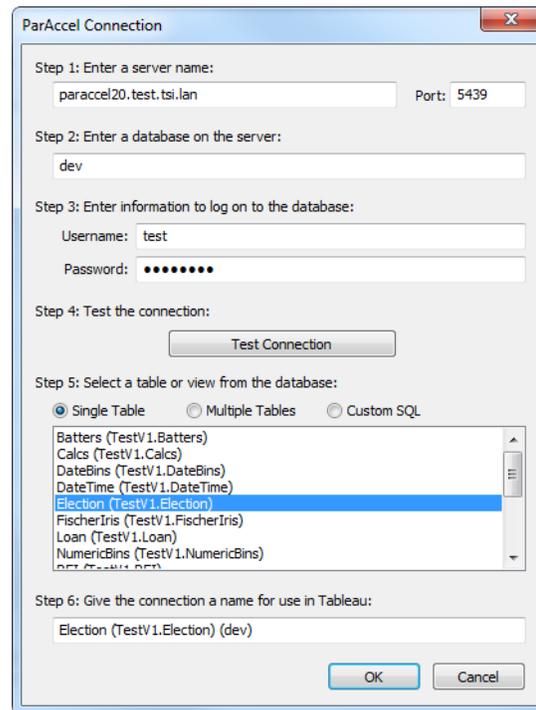
**Step 5: Select a table or view from the database**

ParAccel 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 7-1 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 completed ParAccel Connection dialog box is shown below.



## PostgreSQL Database

This example discusses how to connect Tableau to a PostgreSQL database.

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **PostgreSQL**.
- 3 Follow the steps on the **PostgreSQL Connection** dialog box to complete the connection. These steps are described below.

### Step 1: Enter a server name

Type in the name of a specific server.



---

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

Specify the username and 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 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 7-1 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 completed PostgreSQL Connection dialog box is shown below.



PostgreSQL Connection

Step 1: Enter a server name:  
postgres8.test.tsi.lan Port: 5432

Step 2: Enter a database on the server:  
TestV1

Step 3: Enter information to log on to the database:  
Username: test  
Password: ●●●●●●

Step 4: Test the connection:  
Test Connection

Step 5: Select a table or view from the database:  
 Single Table  Multiple Tables  Custom SQL

Batters
Calcs
DateBins
DateTime
Election
FischerIris
Loan
NumericBins
...

Step 6: Give the connection a name for use in Tableau:  
Election (TestV1)

OK Cancel



## **Sybase IQ**

This example discusses how to connect Tableau to a Sybase IQ database.

- 1 Select **Data > Connect to Data Source**, or press **Ctrl + D**, to open the dialog box.
- 2 Select **Sybase IQ** and click **OK**.
- 3 Follow the steps on the **Sybase IQ Connection** dialog box to complete the connection. These steps are described below.



---

**Step 1: Enter a server name**

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

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

Specify the username and 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 5: Select a table or view from the database**

Sybase IQ 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 7-1 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 completed Sybase IQ Connection dialog box is shown below.



Sybase IQ Connection

Step 1: Enter a server name:  
sybaseiq151.test.tsi.lan Port: 2639

Step 2: Enter a database on the server:  
TestV1\_151

Step 3: Enter information to log on to the database:  
 Use Windows NT Integrated security (preferred)  
 Use a specific username and password:  
Username: test  
Password: ●●●●●●

Step 4: Test the connection:  
Test Connection

Step 5: Select a database on the server:  
testv1

Step 6: Select a table or view from the database:  
 Single Table  Multiple Tables  Custom SQL

batters
calcs
datebins
datetime
<b>election</b>
fischeriris
loan
numericbins
...

Step 7: Give the connection a name for use in Tableau:  
election (testv1)

OK Cancel



## **Teradata**

This example discusses how to connect Tableau to a Teradata database.

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



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

**Step 1: Enter a server name**

Type in the name of a specific server.

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

Specify a username and password.

**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 or database administrator.

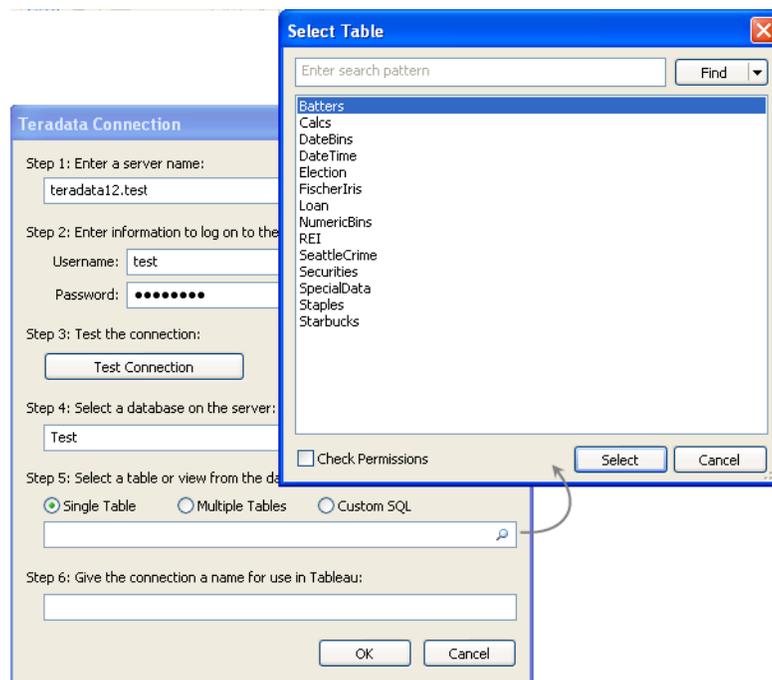
**Step 4: Enter a database on the server**

Servers often contain multiple databases. Type the name of the database you want to connect to or click the Search button to browse and select the database.



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

Teradata databases can contain multiple tables and views. Type the name of the table you want to connect to or click the Search button to browse and select the table. 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 7-1 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 completed Teradata Connection dialog box is shown below.



**Teradata Connection** ✕

Step 1: Enter a server name:

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

Step 3: Test the connection:

Step 4: Select a database on the server:  
 🔍

Step 5: Select a table or view from the database:  
 Single Table    Multiple Tables    Custom SQL  
 🔍

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



## **Vertica**

This example discusses how to connect Tableau to a Vertica database.

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



- 3 Follow the steps on the **Vertica Connection** dialog box to complete the connection.

**Step 1: Enter a server name**

Type in the name of a specific server and a port number.

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

Specify the username and 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 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 7-1 to learn more about connecting to a set of relational tables that are related by join conditions.

**Optionally, click the Initial SQL button to define a query that will execute before any other SQL is issued to set up the session information that may be needed for further queries to provide meaningful results. For example, you may want to set up some temporary tables for the session or set database environment variables such as time zone or locale.**



**After you enter the Initial SQL and click OK, the connection will be tested using this SQL and the list of database tables will be regenerated.**

**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 completed Vertical Connection dialog box is shown below.



Vertica Connection

Step 1: Enter a server name:  
vertica.test.tsi.lan Port: 5433

Step 2: Enter a database on the server:  
db

Step 3: Enter information to log on to the database:  
Username: test  
Password: ●●●●●●

Step 4: Test the connection:  
Test Connection

Step 5: Select a table or view from the database:  
 Single Table  Multiple Tables  Custom SQL

Batters
Calcs
DateBins
DateTime
Election
FischerIris
Loan
NumericBins

Step 6: Give the connection a name for use in Tableau:  
Election (TestV1)

OK Cancel

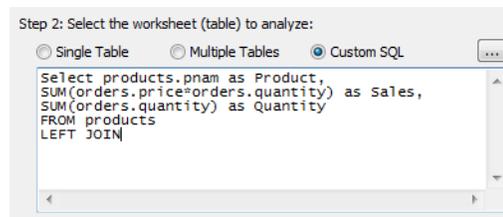
---

## 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

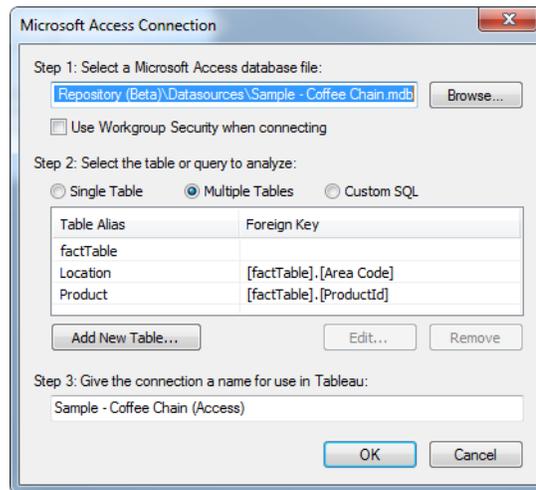
When a Tableau workbook connected to a data source, you can edit the connection. You might want to edit the data source connection to:

- 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 > 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 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.

---

## Renaming the Connection

When you connect to a data source you are given the option to give it a name for use in Tableau. You can change the name you specified by selecting **Data > Data Connection > Rename**. Naming a connection is useful when you have a single workbook connected to many data sources. The name you assign can help you keep track of the specifics of the connection. You can also review the connection properties by selecting **Data > Data Connection > Properties**.



## Duplicating a Connection

Sometimes you'll want to make changes to a data connection such as add more tables, hide and show fields, set field defaults, and so on. When you make these changes it affects all sheets that use the data connection. You can duplicate a connection so that you can make the changes without affecting the existing sheets. To duplicate a connection select **Data > Data Connection > Duplicate**. When you duplicate a connection the name has "(copy)" appended to the end. Refer to "Renaming the Connection" on page 5-72 to learn how to change this name.



## 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 5-5. If you often 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 or if you’ve 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.

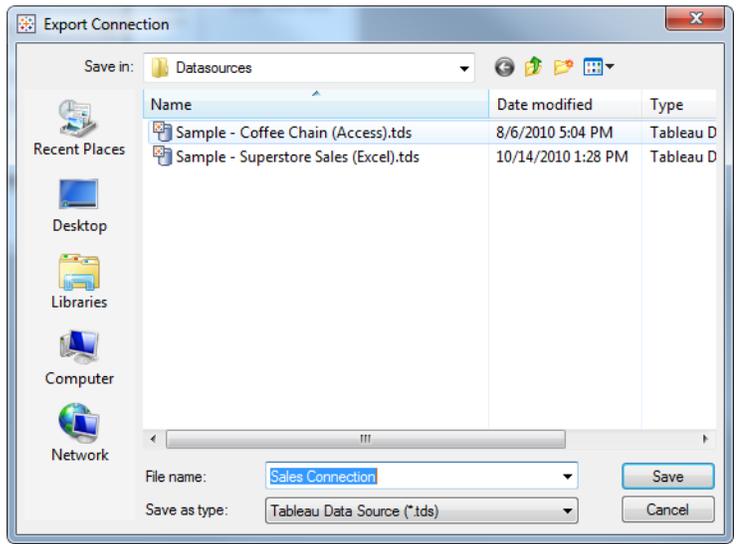
TDS files are stored in the Datasources folder of the Tableau Repository. They have the .tds extension and can be identified by the data source icon . If you move the TDS file to another location, you cannot access the file with the **Connect to Data** dialog box. However, you can access the file by 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, MySQL, and so on) as the original.

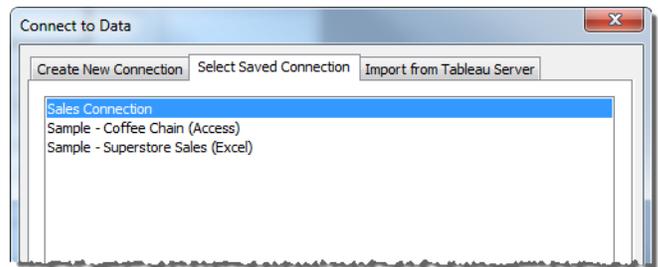


**To export the connection:**

- 1 Select **Data > 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.



---

## 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.

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.

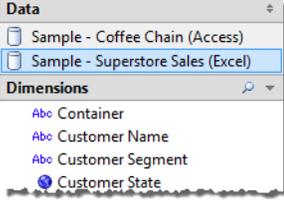


## Closing the Connection

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.

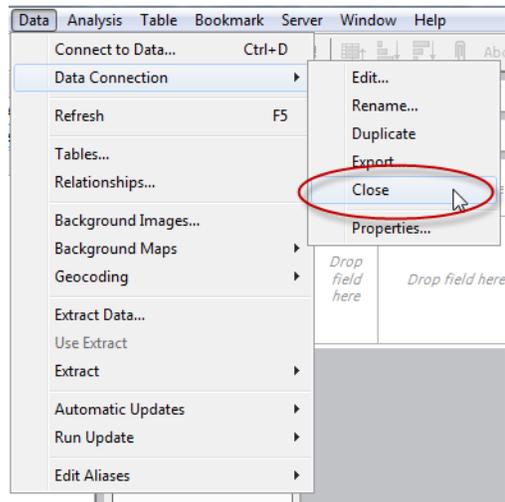
**To close a data connection:**

- 1 Select the data source in the Data window.





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



## Clipboard Data Sources

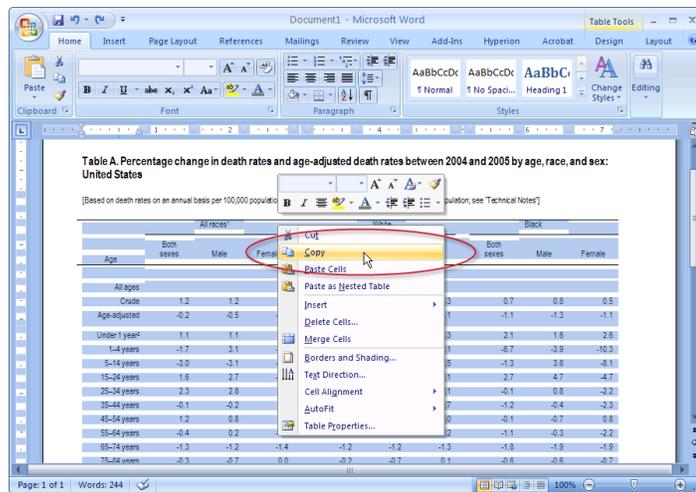
---

Sometimes you want to pull in data from an outside source for some quick analysis. Rather than create a whole data source and then connect in Tableau, you can copy and paste the data directly into the application. Tableau automatically creates a data source that you can begin analyzing. When you save the workbook, the data source is saved as a tab delimited text file into your Tableau Repository.

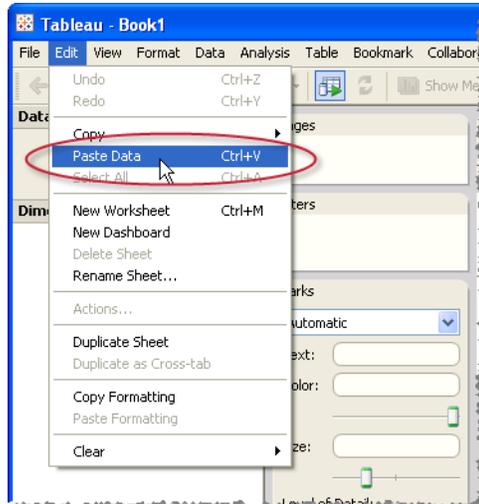
You can copy and paste data from a variety of office applications including Microsoft Excel and Word. You can also copy and paste html tables from webpages. Tables that are copied as comma separated values or tab delimited can be pasted into Tableau. Please be aware that not all applications use these formats when copying.

**To copy data into Tableau:**

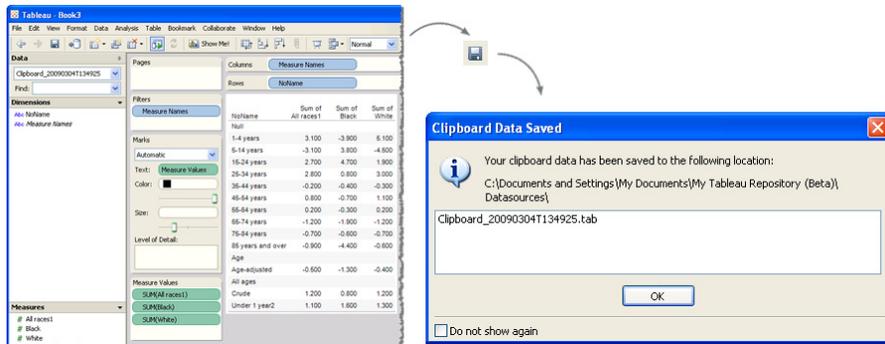
- 1 Select the data you want and copy it to the clipboard.



- 2 Open Tableau Desktop and select **Edit > Paste Data**. A new data source is automatically created.



- 3 Select **File > Save** to save the data source. When you save the workbook the data source is automatically put into your repository. If you save as a packaged workbook the data sources is saved with the workbook instead.





# Joining Tables

---

<b>Overview</b> . . . . .	<b>7-2</b>
<b>Adding Tables</b> . . . . .	<b>7-3</b>
Connecting to Multiple Tables . . . . .	7-3
Adding Tables to the Data window . . . . .	7-7
<b>Editing Tables</b> . . . . .	<b>7-8</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



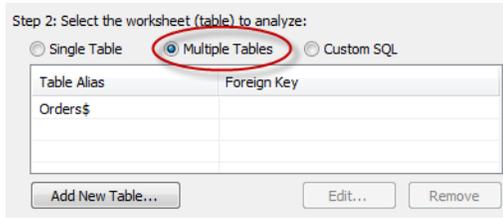
## Adding Tables

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.

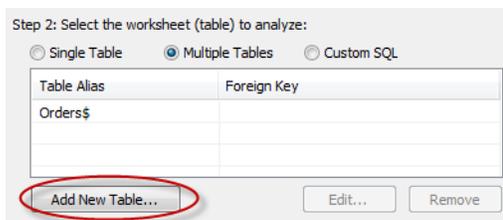
- Connecting to Multiple Tables
- Adding Tables to the Data window

### 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 5-8.
- 2 Select the table or view you want to start with (typically the fact table) and then select the **Multiple Tables** option.

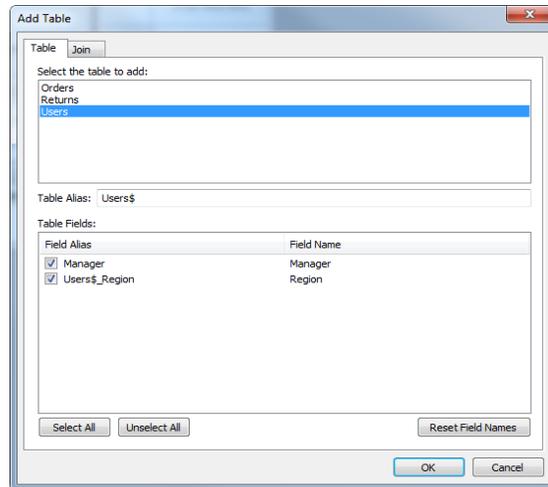


- 3 Click the **Add New Table** button at the bottom of the dialog box.





- 4 In the Add Table dialog box select a table to add to the Data window.



- 5 Optionally change Field aliases.

In the bottom half of the Add Table dialog box, there is a list of fields with their aliases. Double-click the field alias to change how the field will be displayed in the Data



window. This is often useful when you have duplicate field names across tables or your field names are not very understandable.

6 Add a join by selecting the **Join** tab.

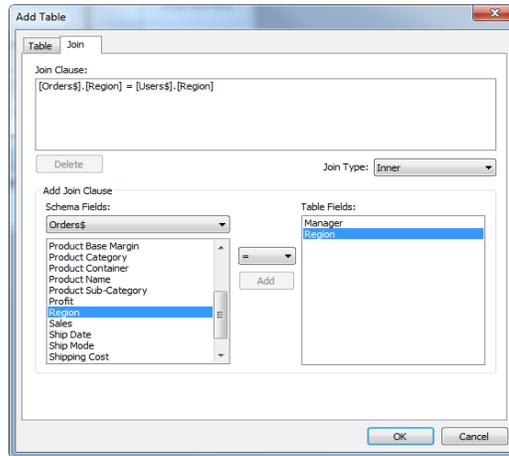
Tableau will automatically create a join for you based on the structure of your data. If an automatic join has been made, it is listed at the bottom of the dialog box. Select the Joins tab to inspect the join clause to make sure it is how you want to connect the tables.



You can delete unwanted join clauses by selecting it in the list of join clauses and clicking **Delete**.

7 In the Join pane, add one or more join clauses by selecting a field from the original table, a field from the added table, and an operator. Then click **Add** to add it to the list of Join Clauses.

For example, in a data source that has a table of customer information and another for product information, you could join the two tables based on an Product ID field that exists in both tables. Select Product ID in both the lists of fields, select the equal sign as the operator, and click **Add**.



- 8 Select the type of join from the **Join Type** drop-down list. You can select Inner, Left, Right, or Full.

Please note, you cannot nest Inner joins within Left or Right joins. These joins will cause a join expression not supported error.

- 9 When finished, click **OK**.

The tables are listed in the Connection dialog box along with the foreign keys.

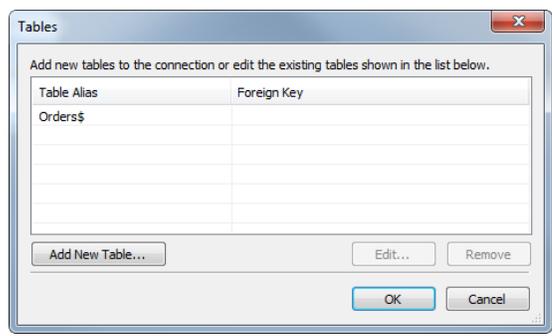
- 10 Complete the connection by giving it a name and clicking OK.

When you add joined tables, the Data window is automatically organized to use the **Group by Table** command. You can turn this feature off or change how the Data window is sorted using the Data window menu. Refer to “Organize the Data window” on page 11-13 to learn more.



## Adding Tables to the Data window

- 1 Select **Data > Tables**.



- 2 In the Tables dialog box click **Add New Table**.
- 3 In the Add Tables dialog box, select a table to add. Complete the dialog box by defining at least one join clause and optionally changing field aliases. Refer to “Connecting to Multiple Tables” on page 7-3, beginning at step 4, to learn more about the Add Table dialog box.
- 4 When finished, click **OK** and then **OK** again to close the dialog boxes.

---

**Note** When you connect to multiple tables you are essentially connecting to a denormalized view of the data source. This means that all queries are run against all tables and it is possible for some measures to be over counted. For example, suppose you have an employees table and an orders table. However you keep the employee salary measure in the orders table. The salary will be counted for each order the employee made. Use the MIN aggregation to remove the double counting.

---

---

## Editing Tables

You can modify the joined tables using the **Tables** command. You may want to edit a table to add or remove fields from the Data window, modify the join clause, or add more join clauses to further define how the table is connected to the original table.

**To edit a table:**

- 1 Select **Data > Tables**.
- 2 In the Tables dialog box, select the joined table you want to modify and click **Edit**.
- 3 In the subsequent dialog box, you can change the table and field aliases as well as add and remove fields from the Data window. Select the **Join** tab to edit the join clauses.
- 4 When finished click **OK** twice to close the Tables dialog boxes.

# Working with Multiple Connections

---

<b>Understanding Data Integration . . . . .</b>	<b>8-3</b>
<b>Navigating Connections in the Data Window . . . . .</b>	<b>8-4</b>
<b>Adding a Secondary Connection . . . . .</b>	<b>8-6</b>
<b>Defining Relationships . . . . .</b>	<b>8-8</b>
<b>Example - Data Integration on a Worksheet . . . . .</b>	<b>8-12</b>
<b>Troubleshooting Data Integration . . . . .</b>	<b>8-17</b>



## Overview

A workbook can contain multiple connections to multiple data sources. Each connection is listed at the top of the Data window. Each worksheet has a primary connection and can optionally have several secondary connections using data integration. The primary connection and the secondary connections are linked by specified relationships. Adding a secondary connection to a sheet can be useful when you have data in multiple data sources that you want to integrate into a single analysis. This section discusses primary and secondary connections and specifying relationships in the following sections:

- Understanding Data Integration
- Navigating Connections in the Data Window
- Adding a Secondary Connection
- Defining Relationships
- Example - Data Integration on a Worksheet
- Troubleshooting Data Integration



## Understanding Data Integration

Data integration is when you blend data from multiple data sources on a single worksheet. The data is joined on common dimensions. Data integration does not create row level joins and is not a way to add new dimensions or rows to your data. Refer to “Joining Tables” on page 7-1 to learn how to create those types of joins. Instead, data integration should be used when you have related data in multiple data sources that you want to analyze together in a single view. For example, you may have Sales data collected in an Oracle database and Sales Goal data in an Excel spreadsheet. To compare actual sales to target sales, you can blend the data based on common dimensions to get access to the Sales Goal measure.

To integrate data, you must first add one of the common dimensions from the primary data source to the view. For example, when blending Actual and Target sales data, the two data sources may have a Date field in common. The Date field must be used on the sheet. Then when you switch to the secondary data source in the Data window, Tableau automatically links fields that have the same name. If they don't have the same name, you can define a custom relationship that creates the correct mapping between fields.

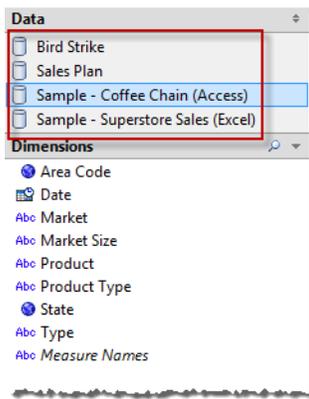
For each data source that is used on the sheet, a query is sent to the database and the results are processed. Then all the results are left joined on the common dimensions. The join is done on the member aliases of the common dimensions so if the underlying values aren't an exact match, you can fix it up in Tableau.

In general, a good test to see whether data can be integrated smoothly is to drag the dimensions from the primary data source into a text table on one sheet. Then on another sheet, drag the same fields from the secondary data source into a text table. If the two tables match up then the data is most likely going to blend correctly.

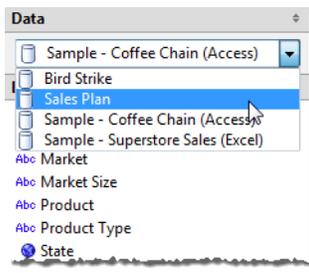


## Navigating Connections in the Data Window

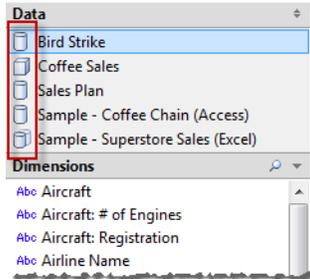
The top of the Data window lists all of the connections in a given workbook. Simply select the data source you want to use and the Data window updates to show the corresponding fields.



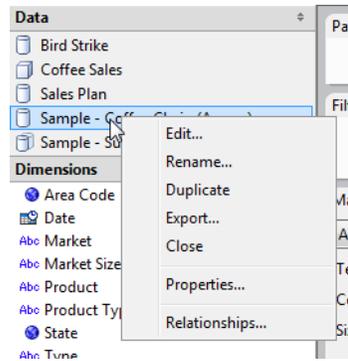
You can resize the connection list area in the Data window to save space. When you resize to a limited vertical height, the list is converted to a drop-down list.



Each connection has an icon to indicate the type of connection. For example, the icon can indicate whether the data source is relational, a cube, or a data extract.



You can right-click the connections to access the commands that are on the **Data > Connection** menu. For example, you can right-click a connection and rename, export, or close it.

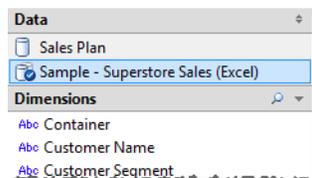




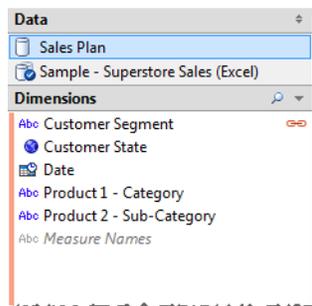
## Adding a Secondary Connection

Sometimes you may have data in two separate data sources that you want to analyze together on a single worksheet. While you can analyze several data sources together on the same worksheet, each worksheet must have a primary data source.

The primary data source is the connection that you first use in the view. After you drag fields to the view, the primary connection is marked with a blue check mark.



If you switch to another connection, you'll notice that the Data window is marked with an orange color to indicate that if you use fields from this connection, it will become the secondary connection.



### To add a secondary connection:

- 1 Connect to the primary data source and build a view.

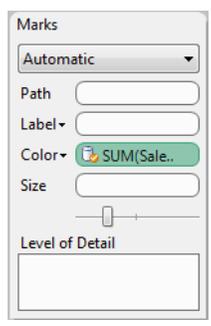


- 1 When you need additional data from a separate data source, select **Data > Connect to Data**.
- 2 Follow the steps in “How to Connect to a Data Source” on page 5-4 to connect to the secondary data source.
- 3 On the worksheet where you need the secondary data, select the new connection. Notice that the Data window is colored orange to indicate that it is a secondary data source.

If your view uses any fields that exist in the secondary data source, you'll see that Tableau automatically defined a relationship to link those two columns. Linked fields are marked with a link icon  in the Data window. You can also define your own custom relationships to handle columns that don't have matching names. You must have a linked field in the view to use data from the secondary connection. Refer to “Defining Relationships” on page 8-8 to learn how to link fields.

- 4 Drag the fields from the secondary data source into the view.

The fields in the view that are from the secondary data source are marked with an orange check mark to indicate that they are from the secondary data source.



**To remove a secondary connection:**

- 1 Remove any secondary fields that are used in the view.
- 2 Right-click the secondary connection at the top of the Data window and select **Close**.



## Defining Relationships

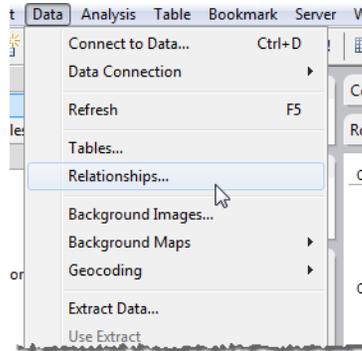
Tableau automatically recognizes when a field you are using in the view exists in a secondary data source. These fields are marked with a link icon in the Data window. You must have a linked field in the view in order to use data from the secondary data source. For example, the workbook shown below has two connections: Superstore Sales and Sales Plan. These two connections have related information including the columns for Customer Segment and Customer State. The data from Sales Plan (the secondary data source) cannot be used until one of those common fields has been added to the view.

You can modify the automatic relationships or create new custom relationships by selecting **Data > Relationships**.

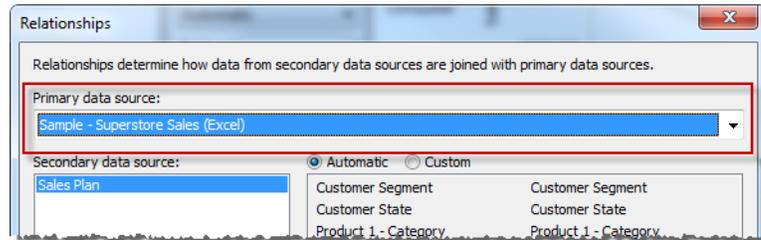
The Relationships dialog box lets you select a primary data source using the drop-down list at the top. Then you can select a secondary data source in a list on the left side of the dialog box. The right side of the dialog box lists any relationships that have been added.

**To add and edit relationships:**

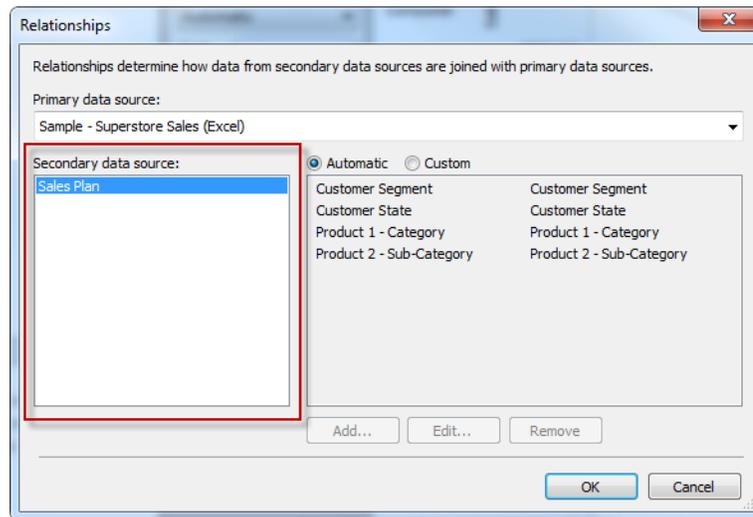
- 1 Select **Data > Relationships** to open the Relationships dialog box.



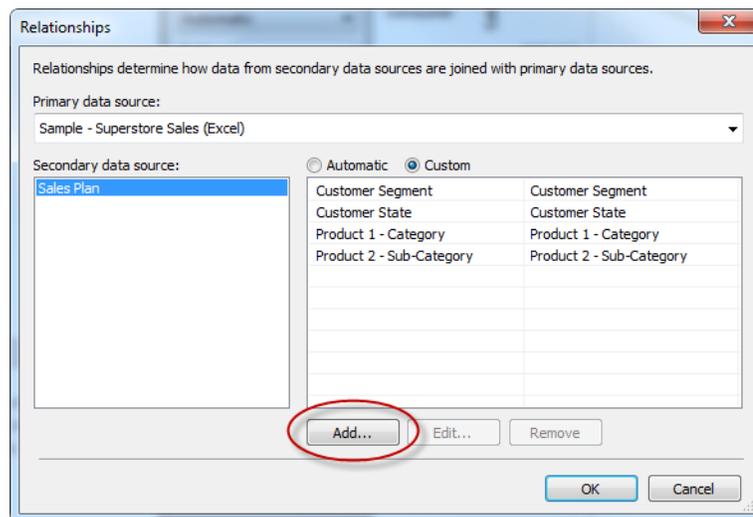
- 2 Make sure that the primary data source is selected from the down list.



- 3 Select a secondary data source from the list on the left.

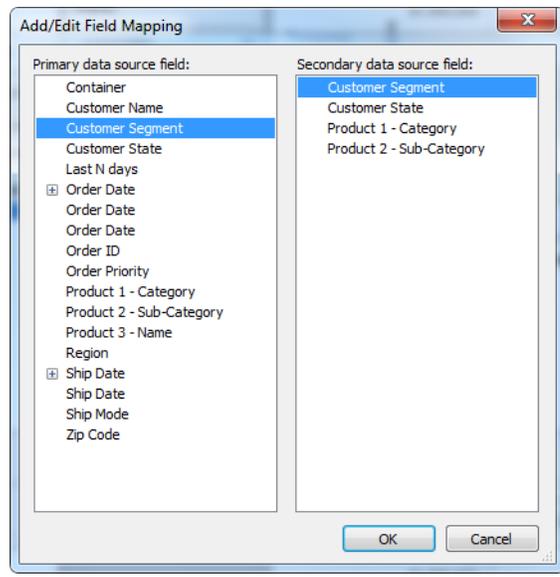


- 4 Select **Custom** at the top of the relationships list.
- 5 Click the **Add** button or select an existing relationship and click the **Edit** button.





- 6 Select a field in the primary data source and map it to matching field in the secondary data source.



- 7 When finished, click **OK**.
- 8 Add and Remove as many relationships as necessary, and when finished, click **OK** again.

The linked fields are marked with a link icon  in the Data window to show that they are related to a field in the primary data source.

The relationship matches values based on the member aliases. You can fix up fields that don't match by editing the aliases. For example, when mapping a State Name field in the primary data source to a State Abbreviation field in the secondary data source, "AK" will not map correctly to "Alaska". You'll have to modify the aliases in one of the data sources.

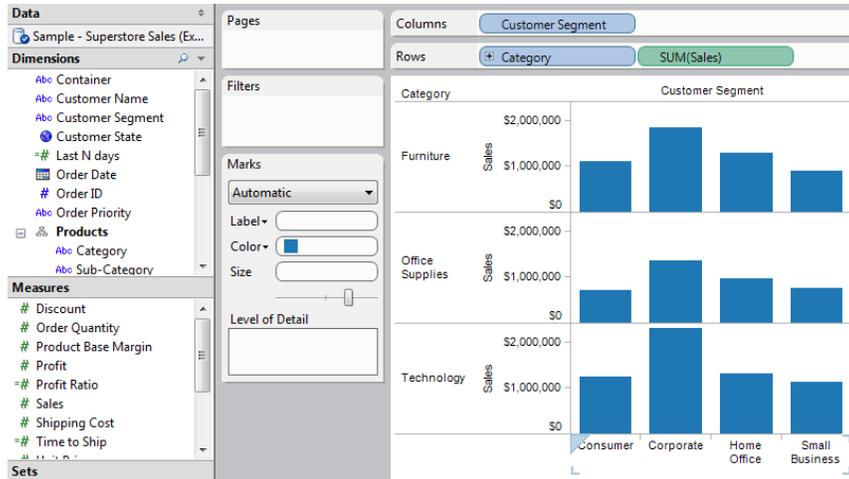


## **Example - Data Integration on a Worksheet**

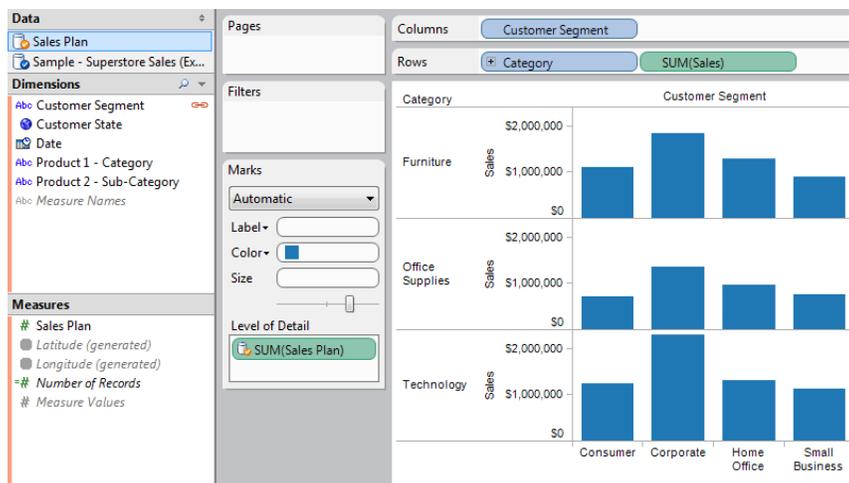
This example uses two data sources: Sample - Superstore Sales (Excel) that comes with the application and an auxiliary Excel file that contains forecasted sales information. An example of each of their columns are shown below:

Follow the steps below to use data from both data sources on a single worksheet.

- 1 Connect to Sample - Superstore Sales (Excel) and build a view that shows Sales by Customer Segment and Product Category.



- 2 Select **Data > Connect to Data** and connect to the Sales Plan spreadsheet.
- 3 Drag the **Sales Plan** measure to the Level of Detail shelf.

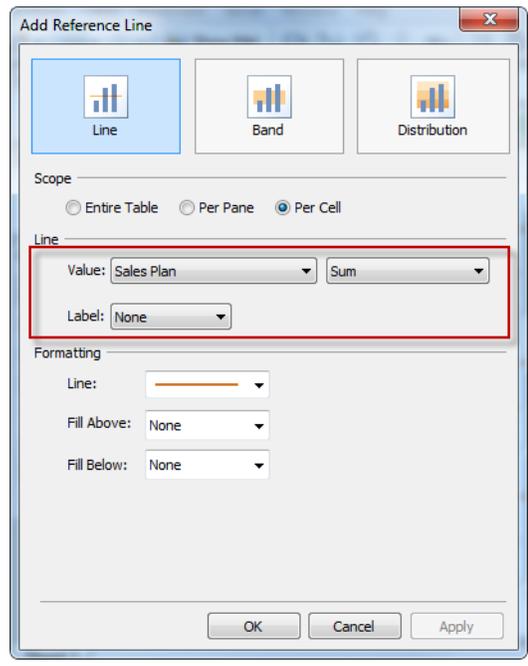




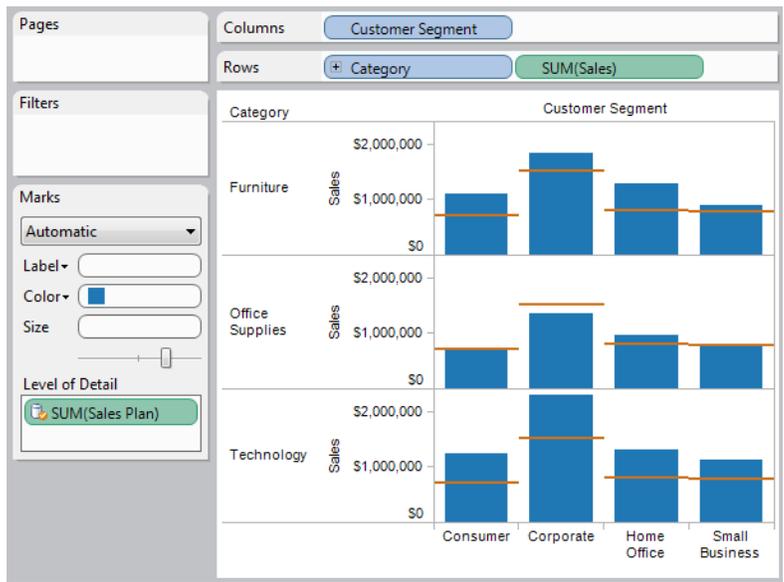
4 Right-click the Sales axis and select **Add Reference Line**.



5 In the Reference Line dialog box, add a reference line that shows Sales Plan per cell. When finished, click **OK**.



- 6 The Worksheet is now pulling data from the secondary data source (Sales Plan) to show how actual sales compared to the forecasted sales.



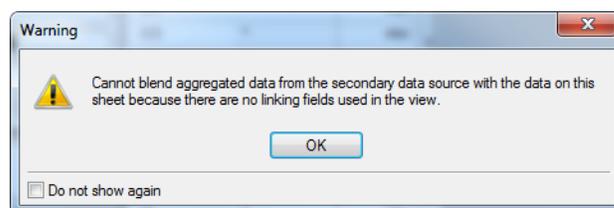
## Troubleshooting Data Integration

When you integrate data from multiple data sources, you may run into some of these common issues:

### Warning: Cannot Blend Because No Linking Fields

When you drag a field from a secondary data source to the view, you may see a warning that says:

Cannot blend aggregated data from the secondary data source with the data on this sheet because there are no linking fields used in the view.



This warning occurs when you have not used a linking field in the view. For example, if you have two data sources with the related dimensions State and Date, you must be using one of those primary dimensions in the view before you can blend data from the secondary data source.

The secondary data source may not have any relationships to the primary data source. Look in the Data window for the link icon. Tableau automatically links fields that have the same name. If your fields do not have the same name you'll have to create a custom relationship. Refer to “Defining Relationships” on page 8-8.

### Asterisks Show in the View

When you relate secondary data, make sure that there is only one matching member in the secondary dimension for each mark. If there are multiple matching members you will see an asterisk in the view. For example, say you have two data sources. The primary data source has a state field that contains state abbreviations. The secondary data source also contains a state field along with a customer segment field. Each state may have multiple customer segments (e.g., CA has Consumer and Corporate). When you relate the two data sources on state, you've created a relationship where state can have multiple customer segment values. When that happens you will see an asterisk in the view.

Primary Data source

DestState	Drop
AK	Abc
AL	Abc
AR	Abc
AZ	Abc
CA	Abc
CO	Abc
CT	Abc
DC	Abc
FL	Abc
GA	Abc

Secondary Data source

Customer State	Customer Segment	Drop
AL	Corporate	Abc
	Home Office	Abc
	Small Business	Abc
AZ	Consumer	Abc
	Corporate	Abc
	Home Office	Abc
	Small Business	Abc
CA	Consumer	Abc
	Corporate	Abc
	Home Office	Abc
	Small Business	Abc
CO	Consumer	Abc

Blended Data

Customer Segment is an attribute of State. Because there are multiple segments per state, an asterisk displays.

DestState	Customer Segment	Drop
AK	Null	Abc
AL	*	Abc
AR	Null	Abc
AZ	*	Abc
CA	*	Abc
CO	*	Abc
CT	*	Abc
DC	*	Abc

All secondary fields are aggregated. Dimensions are aggregated as Attributes (ATTR), which means that if there's only one member it will show the member value but if there are multiple members it will show an asterisk.

# Extracting Data to the Data Engine

---

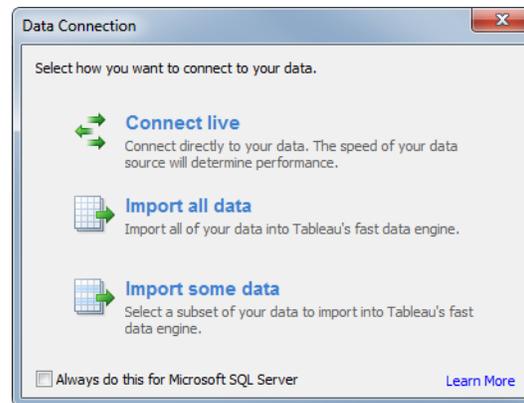
Overview . . . . .	9-2
Creating an Extract . . . . .	9-3
Using Extracts . . . . .	9-7
Upgrading Legacy Extracts . . . . .	9-8
Optimizing Extracts . . . . .	9-9

---

## Overview

When you import all or some of your data into the Tableau's data engine, you create a data extract. Extracts can improve performance, upgrade your data to allow for more advanced capabilities, and enable offline analysis. Extracts allow you to analyze data stored in large data sources that would otherwise take too long to query over and over.

When you connect to data, you are given the following options:



- Connect live – creates a live connection to your data source. No extract is created.
- Import all data – imports all of the data into the data engine.
- Import some data – allows you specify filters to import some of the data into the data engine.

This section discusses the following topics:

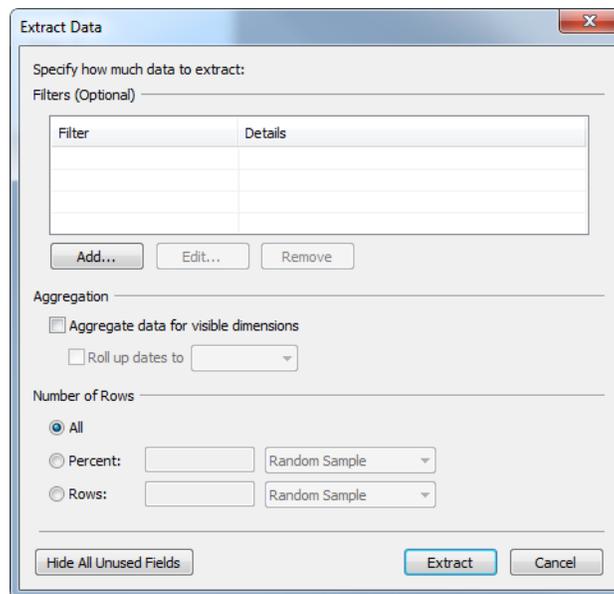
- Creating an Extract
- Using Extracts
- Upgrading Legacy Extracts
- Optimizing Extracts

---

## Creating an Extract

If you don't create a connection when you first open the data, you can create an extract later using the Data menu. Follow the steps below to create an extract.

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



- 2 Optionally define filters to limit the data that will be extracted. Any fields that are hidden in the Data window will be automatically excluded from the extract.

To add filters, click the **Add** button under the Filters list. Refer to “Dragging Fields to the Filters Shelf” on page 16-7 for more information about defining filters.



- 3 Specify whether to **Aggregate data for visible dimensions**. When you select this option the measures are aggregated using their default aggregation. Aggregating the data can minimize the size of the extract file and increase performance.

When you choose to aggregate the data you can also choose to **Roll up dates** to a specified date level such as Year, Month, etc.

Aggregation

Aggregate data for visible dimensions

Roll up dates to Month

The examples below show how the data will be extracted for each aggregation option.

Original Data		Table			Each record is shown as a separate row. There are 7 rows in the data source.
		A	B	C	
1		Date	Region	Sales	
2		1/1/2009	South	\$500	
3		1/1/2009	West	\$200	
4		1/1/2009	West	\$100	
5		1/1/2009	East	\$300	
6		1/2/2009	South	\$600	
7		1/2/2009	South	\$400	
8		1/2/2009	East	\$100	
9					



<b>Aggregate Data (no roll up)</b>	<table border="1"><thead><tr><th></th><th>A</th><th>B</th><th>C</th></tr></thead><tbody><tr><td>1</td><td>Date</td><td>Region</td><td>Sales</td></tr><tr><td>2</td><td>1/1/2009</td><td>East</td><td>\$300</td></tr><tr><td>3</td><td>1/1/2009</td><td>South</td><td>\$500</td></tr><tr><td>4</td><td>1/1/2009</td><td>West</td><td>\$300</td></tr><tr><td>5</td><td>1/2/2009</td><td>East</td><td>\$100</td></tr><tr><td>6</td><td>1/2/2009</td><td>South</td><td>\$1,000</td></tr></tbody></table>		A	B	C	1	Date	Region	Sales	2	1/1/2009	East	\$300	3	1/1/2009	South	\$500	4	1/1/2009	West	\$300	5	1/2/2009	East	\$100	6	1/2/2009	South	\$1,000	<p>Records with the same date and region have been aggregated into a single row.</p> <p>There are 5 rows in the data source.</p>		
	A	B	C																													
1	Date	Region	Sales																													
2	1/1/2009	East	\$300																													
3	1/1/2009	South	\$500																													
4	1/1/2009	West	\$300																													
5	1/2/2009	East	\$100																													
6	1/2/2009	South	\$1,000																													
<b>Aggregate Data (roll up dates to Month)</b>	<table border="1"><thead><tr><th></th><th>A</th><th>B</th><th>C</th><th>D</th></tr></thead><tbody><tr><td>1</td><td>Date</td><td>Region</td><td>Sales</td><td></td></tr><tr><td>2</td><td>1/1/2009</td><td>East</td><td>\$400</td><td></td></tr><tr><td>3</td><td>1/1/2009</td><td>South</td><td>\$1,500</td><td></td></tr><tr><td>4</td><td>1/1/2009</td><td>West</td><td>\$300</td><td></td></tr><tr><td>5</td><td></td><td></td><td></td><td></td></tr></tbody></table>		A	B	C	D	1	Date	Region	Sales		2	1/1/2009	East	\$400		3	1/1/2009	South	\$1,500		4	1/1/2009	West	\$300		5					<p>Dates have been rolled up to the Month level and records with the same region have been aggregated into a single row.</p> <p>There are 3 rows in the data source.</p>
	A	B	C	D																												
1	Date	Region	Sales																													
2	1/1/2009	East	\$400																													
3	1/1/2009	South	\$1,500																													
4	1/1/2009	West	\$300																													
5																																

4 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 aggregation and then extracts the number of rows from the filtered and aggregated results.

If you specify a percentage, keep in mind that the percentage sample is calculated differently depending on the type of data source you are using. The percentage value you specify indicates the probability of each row being included in the extract.

5 Click the **Hide All Unused Fields** button to quickly remove them from the extract. Refer to “Hide or Unhide Fields” on page 11-16 to learn how to hide fields.

6 When finished, click **Extract**.

7 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.

---



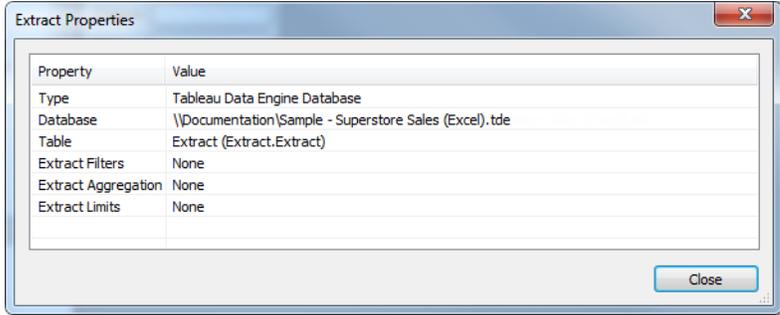
## Using Extracts

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 > Extract > Remove**. 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.

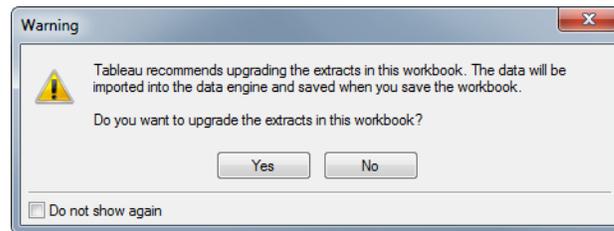
If the the underlying data has changed you can update the extract file with the new data by selecting **Data > Extract > Refresh**. You can see information about the extract by selecting **Data > Extract > Properties**.



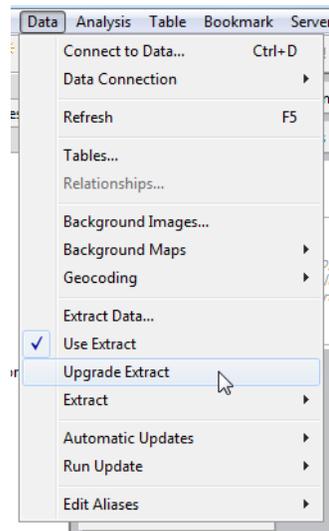
---

## Upgrading Legacy Extracts

If you have data extracts that were created before version 6.0, you should upgrade the extracts to use the data engine. When you open the workbook, you are given the option to upgrade the extracts.



You can also upgrade the extracts by selecting **Data > Upgrade Extract**.





## Optimizing Extracts

To improve performance when working with extracts you can optimize the extract. Optimizing moves any calculated fields you've created into the extract so that they don't have to be computed locally every time you use it.

Optimize the extract by selecting **Data > Extract > Optimize**.

When you modify the calculated field, the modified version will be used until you select **Data > Extract > Optimize** again.

Each time you optimize the extract, any deleted calculations will be removed from the extract, new ones will be added, and modified ones will be updated.



# Managing Queries

---

<b>Overview . . . . .</b>	<b>10-2</b>
<b>Automatic Updates . . . . .</b>	<b>10-3</b>
<b>Cancel Query . . . . .</b>	<b>10-4</b>
<b>Abandoned Queries . . . . .</b>	<b>10-5</b>
<b>Precision Warnings . . . . .</b>	<b>10-6</b>

---

## 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.

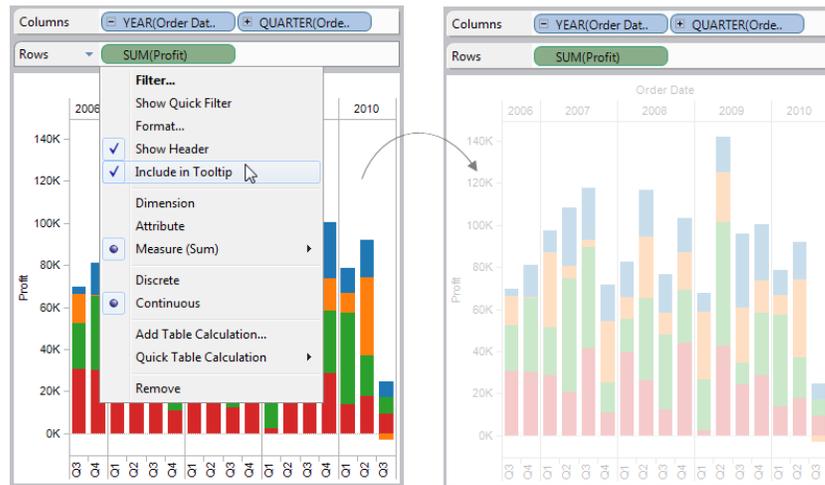
## Automatic Updates

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.



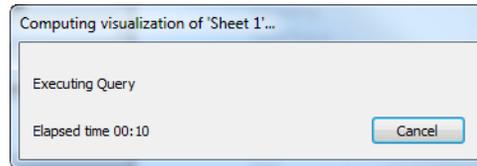
---

## 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 10-5 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>11-3</b>
Relational and Multidimensional Data . . . . .	11-6
Measure Values and Measure Names . . . . .	11-11
Number of Records . . . . .	11-11
Latitude and Longitude (generated) . . . . .	11-11
<b>Data Window Features and Functions</b> . . . . .	<b>11-13</b>
Organize the Data window . . . . .	11-13
Find Fields . . . . .	11-14
Rename Fields . . . . .	11-15
Hide or Unhide Fields . . . . .	11-16
Add Fields to the Data window . . . . .	11-17
<b>Editing Field Properties</b> . . . . .	<b>11-18</b>
Comments . . . . .	11-19
Aliases . . . . .	11-20
Colors . . . . .	11-27
Shapes . . . . .	11-27
Formats . . . . .	11-28
Sort . . . . .	11-29
Aggregation . . . . .	11-29
Changing Measure Names . . . . .	11-32
<b>Data Types &amp; Roles</b> . . . . .	<b>11-34</b>
<b>Data Types &amp; Roles</b> . . . . .	<b>11-34</b>
Data Types . . . . .	11-34
Data Roles . . . . .	11-37
Example: Continuous Dimensions and Discrete Measures . . . . .	11-42

---

## Overview

The data in all data sources are categorized into fields such as Customer, Sales, Profit, Temperature, etc. These fields are made from the columns in your data source. When you connect to a data source with Tableau, the fields are displayed 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:

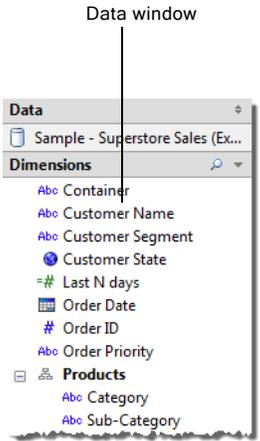
- Understanding the Data Window
- Data Window Features and Functions
- Editing Field Properties
- Data Types & Roles
- Data Types & Roles

# 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 determined by the dimensions and measures of a cube. For relational data sources, the fields are determined 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	B	C	D
1	Record No	Order Priority	Sales Total	Order Date
2	1	5-LOW	2302.42	1/1/2001
3	2	5-LOW	996.66	1/1/2001
4	3	5-LOW	2564.85	1/1/2001
5	4	5-LOW	226.52	1/2/2001
6	5	3-MEDIUM	159.98	1/2/2001
7	6	3-MEDIUM	2679.68	1/2/2001
8	7	3-MEDIUM	898.31	1/2/2001
9	8	3-MEDIUM	56.95	1/2/2001
10	9	3-MEDIUM	10.42	1/2/2001
11	10	3-MEDIUM	4129.53	1/2/2001

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.



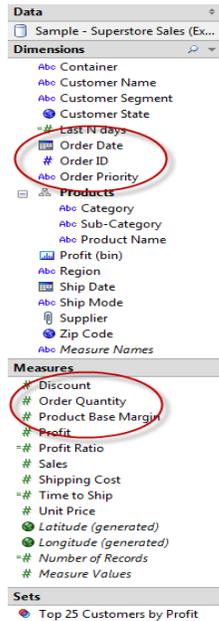


The Data window organizes the fields into three areas:

- **Dimensions** – Fields that typically hold discrete qualitative data. Examples of dimensions include dates, customer names, and customer segments (refer to “Data Types & Roles” on page 11-34).
- **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 11-34).
- **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 17-26.

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 11-38.

The Data window for an Excel worksheet (a relational database) is shown below. The **Discount** and **Order Quantity** 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 “Editing Field Properties” on page 11-18).

---

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

- Relational and Multidimensional Data
- Hierarchies (For Relational Databases)
- Measure Values and Measure Names
- Number of Records
- Latitude and Longitude (generated)



## 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.

The image displays two side-by-side screenshots of data windows, comparing a Relational Data Window and a Multidimensional Data Window.

**Relational Data Window:**

- Data:** Sample - Superstore Sales (Ex...)
- Dimensions:**
  - Abc Container
  - Abc Customer Name
  - Abc Customer Segment
  - Customer State
  - Last N days
  - Order Date
  - Order ID
  - Order Priority
  - Products**
    - Abc Category
    - Abc Sub-Category
    - Abc Product Name
  - Profit (bin)
  - Abc Region
  - Ship Date
  - Abc Ship Mode
  - Supplier
  - Zip Code
  - Abc Measure Names
- Measures:**
  - # Discount
  - # Order Quantity
  - # Product Base Margin
  - # Profit
  - # Profit Ratio
  - # Sales
  - # Shipping Cost
  - # Time to Ship
  - # Unit Price
  - Latitude (generated)
  - Longitude (generated)
  - # Number of Records
  - # Measure Values
- Sets:**
  - Top 25 Customers by Profit

**Multidimensional Data Window:**

- Data:** Basic (Sample)
- Dimensions:**
  - Scenario
    - Gen2, Scenario
  - Market**
    - Region
      - State
        - Major Market
        - New Market
        - Small Market
    - Population
  - Product**
    - Gen2, Product
    - Gen3, Product
      - Caffeinated
      - Ounces
      - Pkg Type
      - Intro Date
  - Year**
    - Quarter
    - Gen3, Year
  - Abc Measure Names
- Measures:**
  - # Profit
    - # Margin
    - # Total Expenses
  - # Inventory
    - # Opening Inventory
    - # Additions
    - # Ending Inventory
  - # Ratios
    - # Margin %
    - # Profit %
    - # Profit per Ounce
  - Latitude (generated)
  - Longitude (generated)
  - # Measure Values



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**.



## Hierarchies (For Relational Databases)

Unlike multidimensional data sources, relational data sources don't have built in hierarchies. However, often relational data sources have related dimensions that have an inherent hierarchy. For example, a data source may have fields for Country, State, and City. These fields could be grouped into a hierarchy called Location.

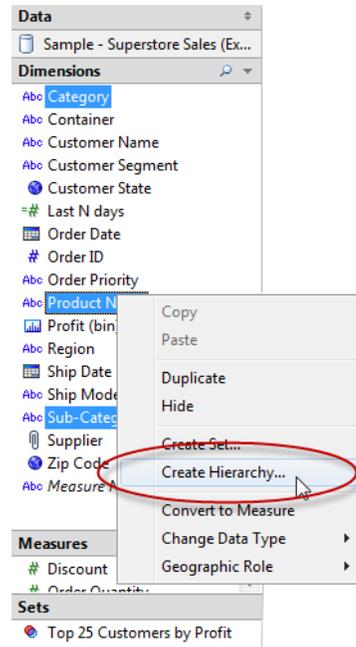
You can assemble dimensions into these hierarchies in the Data window. Hierarchies support single click navigation up and down the levels. When you use the fields in the view, a plus button  displays on the field so you can drill down and up in the hierarchy. Refer to "Drilling Down and Up in a Hierarchy" on page 13-51 to learn more about drilling through a hierarchy.

### To create a hierarchy:

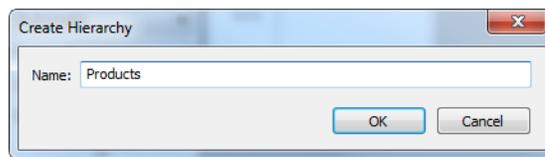
- 1 Do one of the following:
  - Right-click one or more selected fields in the Data window and select **Create Hierarchy**.



- In the Data window, drag and drop a field directly on top of another field.



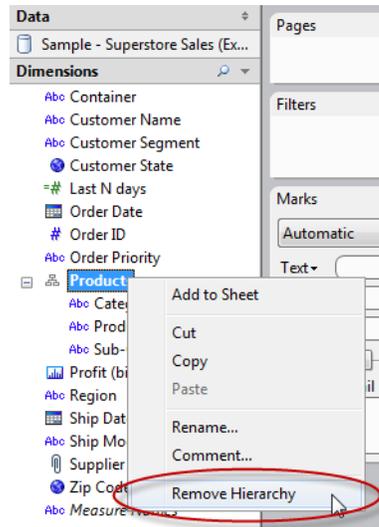
- 2 In the subsequent dialog box, specify a name for the hierarchy.



The hierarchy can be expanded and collapsed in the Data window. You can also drag and drop to add and, remove, and reorder the fields in the hierarchy.

**To remove a hierarchy:**

- Right-click the name of the hierarchy in the Data window and select **Remove Hierarchy**.

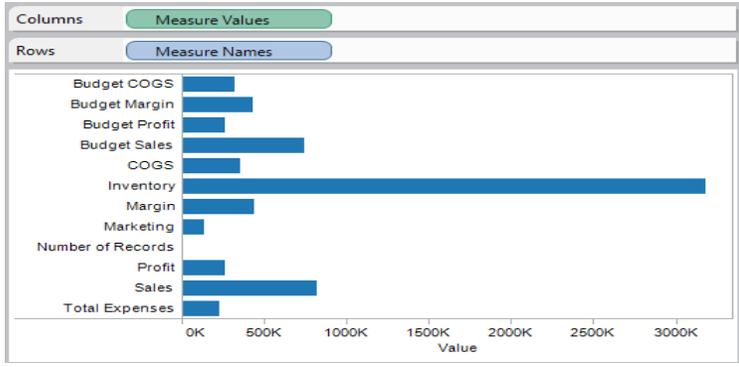




## Measure Values and Measure Names

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 33-35.

## 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.

## Latitude and Longitude (generated)

If you have defined any fields to be geographic fields, that is they can be used with maps, the Tableau automatically geocodes your data and includes Latitude (generated) and



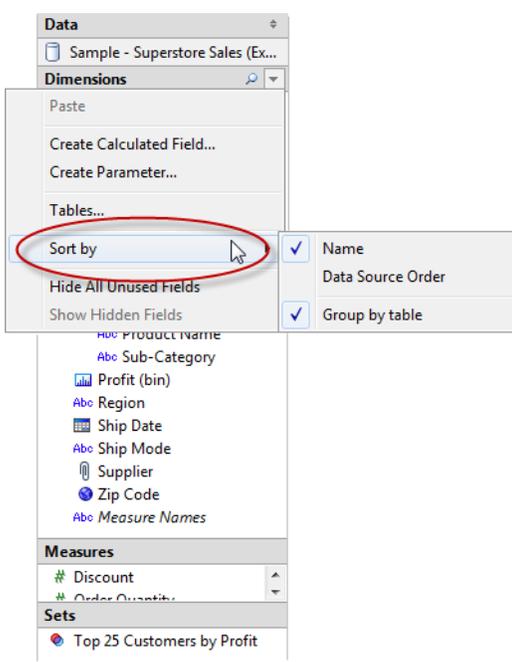
Longitude (generated) fields. You can use these fields to overlay your data on live maps. Refer to “Background Maps” on page 24-1 to learn more.



# 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:

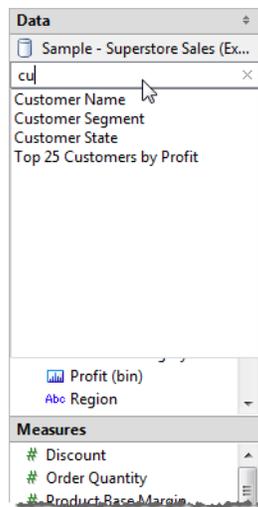
- **Name** - lists the dimensions and measures in alphabetical order according to their field aliases.
- **Data source order** - lists the dimensions and measures in the order they are listed in the underlying data source.

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 “Profit”. To search for a field, click the Find Field icon  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 assign an alternate name for a field that displays in the Data window as well as in the view. For instance, a field called Customer 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. The changed field name is saved with the workbook as well as when you export the connection.

### 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.

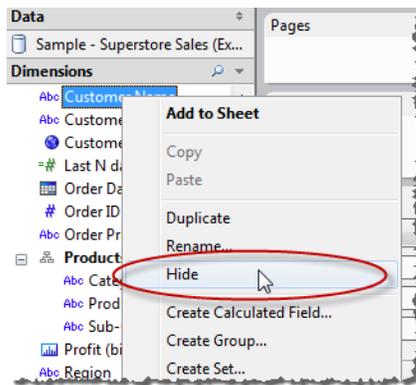
### To revert to the default field name:

- 1 Right-click the field and select **Rename**.
- 2 In the Rename dialog box, click **Reset** then click **OK**.

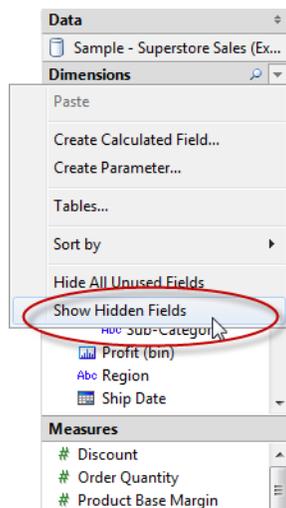


## Hide or Unhide Fields

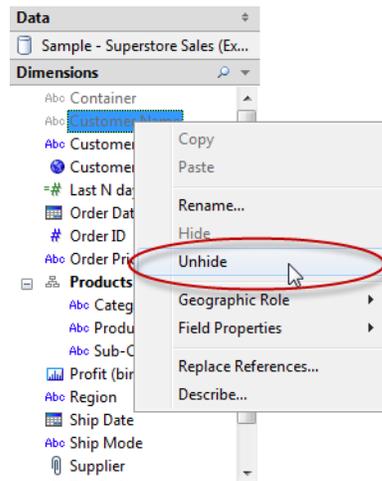
You can selectively hide or show fields in the Data window. To hide a field, right-click the field you want to hide and select **Hide**.



When you want to change your fields from hidden to visible, select **Show Hidden Fields** on the Data window menu.



The hidden fields are shown in gray in the Data window. You can then select one or more hidden fields, right-click and select **Unhide**.



Select **Hide All Unused Fields** on the Data window menu to quickly hide all of the fields that are not being used in the workbook.

**Note** Hiding fields can be a good way to decrease the size of a data extract file because hidden fields are automatically excluded from the extract.

## 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 the Data window menu. Alternatively, select **Analysis > Create Calculated Field**. Refer to “Calculated Fields” on page 21-18 for more information.

---

## Editing Field Properties

When you drag fields to shelves, the data is represented as marks in the view. You can specify settings for how the marks from each field will be displayed by setting mark properties. For example, when you place a dimension on the color shelf the marks will be colored by the values within that dimension. You can set the Color property so that anytime you use that dimension on the color shelf your chosen colors are used. Using field properties you can set the aliases, colors, and shapes, default aggregation, and so on. The field properties are discussed in the following sections.

- Comments
- Aliases
- Colors
- Shapes
- Formats
- Sort
- Aggregation
- Changing Measure Names

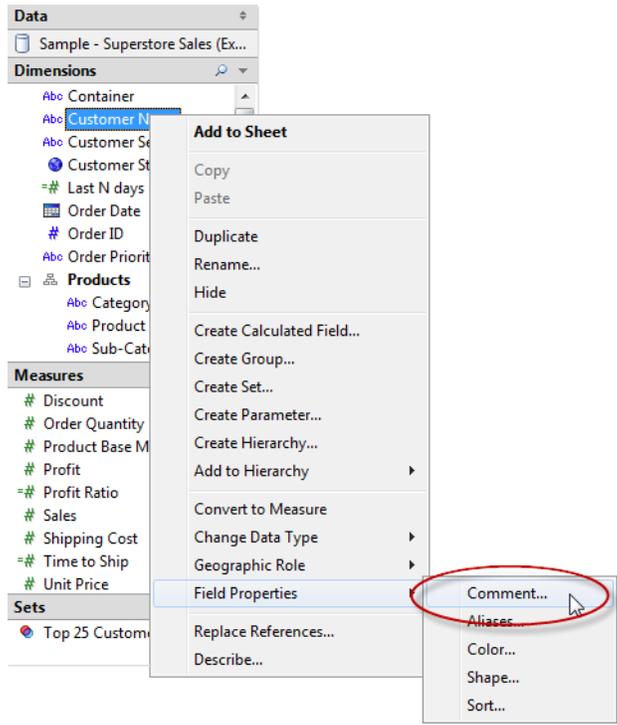


## Comments

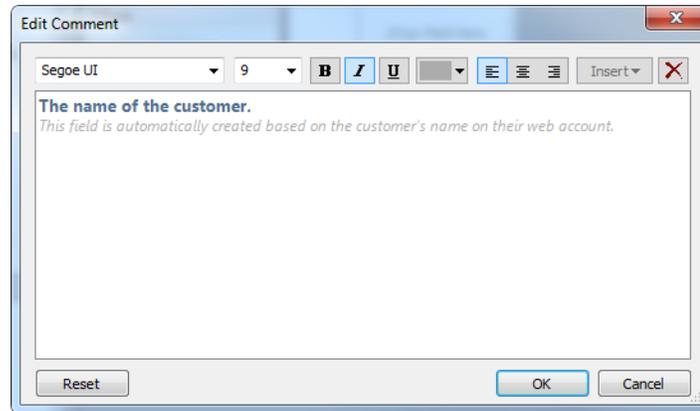
Fields can have comments that describe them. The comments display in a tooltip in the Data window and in the Calculated fields dialog box. Field comments are a good way to give more context to the data in your data source. Comments are especially useful when you are building a workbook for others to use.

### To add a to a field:

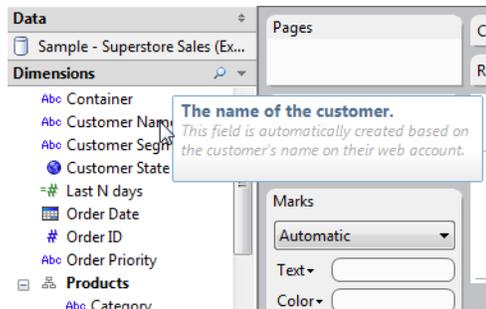
- 1 Right-click a field in the Data window and select **Field Properties > Comment**.



- 
- 2 Write a comment in the subsequent dialog box. Comments support rich text formatting that will be represented in the tooltip.



- 3 When finished, click **OK**.



## Aliases

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

Aliases can be created for the members of most dimensions in the Data window. You cannot, however, define aliases for continuous dimensions and dates and they do not apply

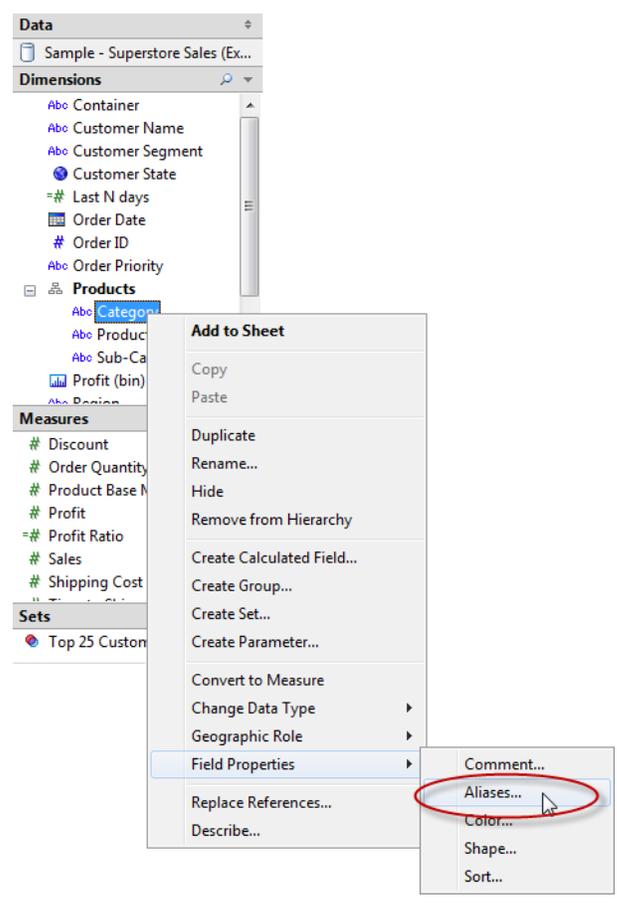


to measures. The method for creating aliases depends on the type of data source you are using:

- Aliases with a Relational Data Source
- Aliases with a Multidimensional Data Source

### Aliases with a Relational Data Source

To create an alias for a relational data source, right-click a field name and select **Field Properties > Aliases**.





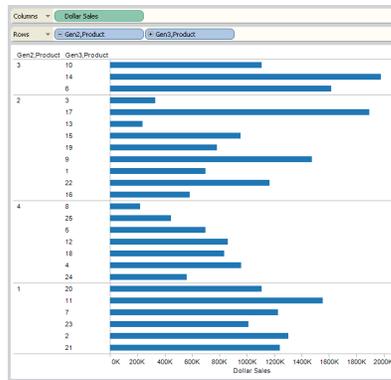
A dialog box opens allowing you to define aliases for each value within the selected dimension. You can reset the member names back to their original names by clicking the **Clear Aliases** button in the bottom right corner of the Edit Aliases dialog box.

### Aliases with a Multidimensional Data Source

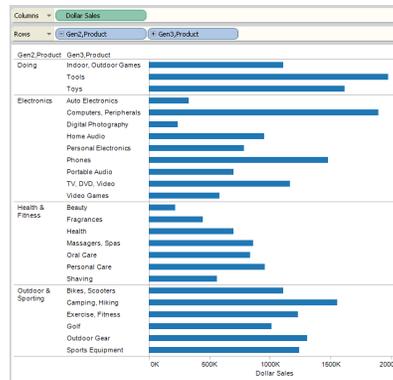
Aliases for Oracle Essbase multi-dimensional databases are created on the server by the server administrator and can be activated in Tableau using the **Alias File** option on the **Data** menu. Please talk to your Essbase database administrator to find out whether your database has aliases available. 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.



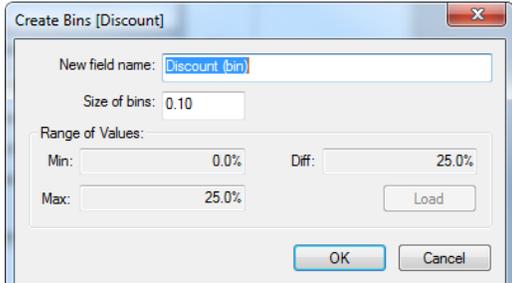
### Example - Editing Aliases

The Superstore Sales Excel sample data source contains a measure called **Discount**, which contains discount values from 0 to 0.25 or 25%. Suppose you want to analyze these data by categories: low discount, medium discount, and high discount.

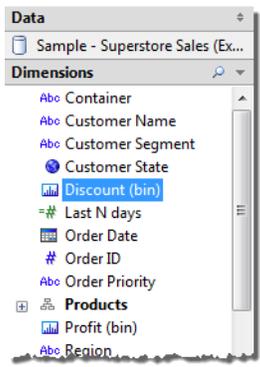
To create the categories, you could first bin the measure so that when added to the view it create discrete headers instead of a continuous axis. In this example, define a bin size equal



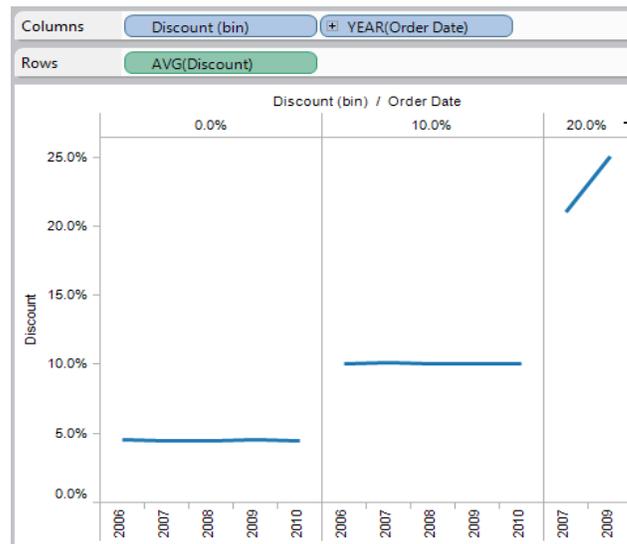
to 0.10, which produces three bins. The first bin contains the values 0 to 9%, the second bin contains the values 10% to 19%, and the third bin contains the values 20% and greater. The **Create Bins** dialog box for this field is shown below. Refer to “Binned Data” on page 21-72 for a complete description 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

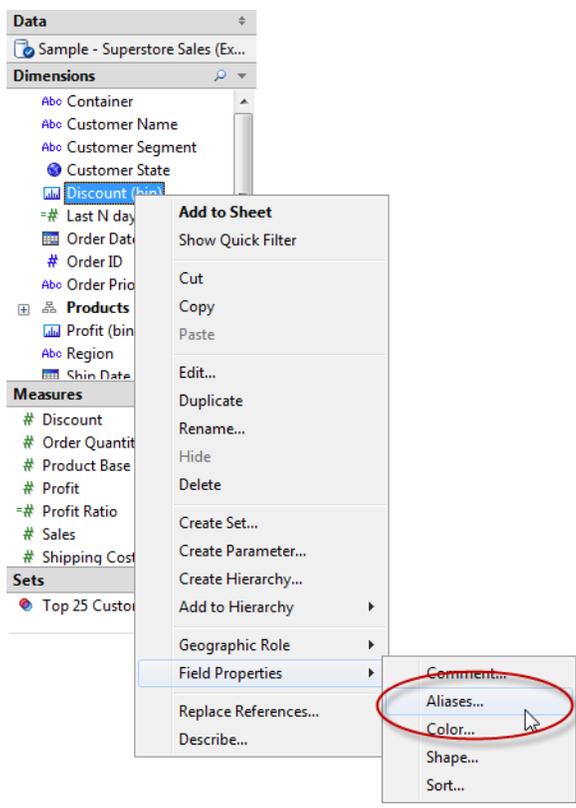


Original bin names.

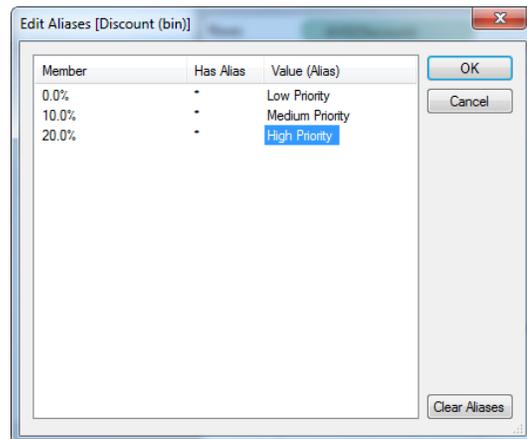
To improve the readability of the bins when they are displayed in Tableau, you can define aliases such as “Low discount”, “Medium discount”, and “High discount.”



- 1 Right-click the name of the dimension in the Data window and select **Field Properties** > **Aliases**.



- 2 Assign an alias to every member of the **Discount (bin)** field. For example, the member originally labeled as 0.0% is now labeled as “Low Discount”.

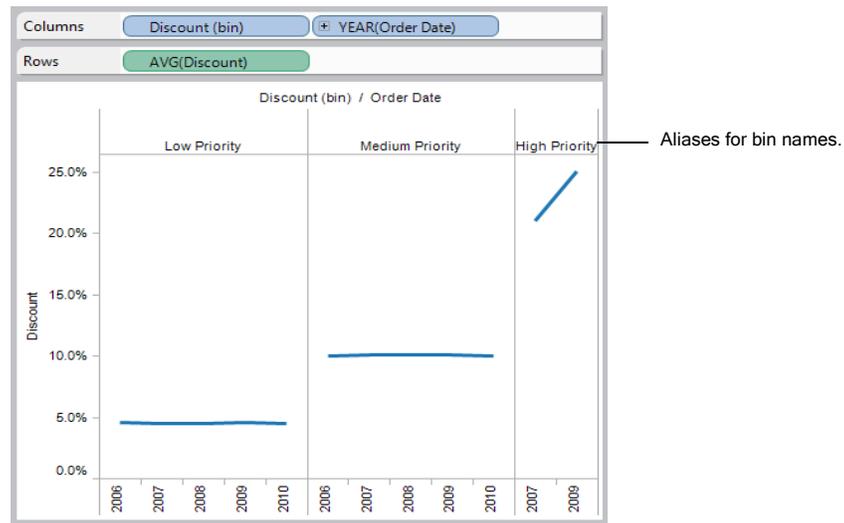


You can change aliases at any time using the Edit Aliases dialog box. To do so, click on the alias you want to change and specify the new name. Use the Tab key to advance from one value to the next. To restore the original aliases, click the **Clear Aliases** button in



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 displays the aliases in the view.



## Colors

When you use a dimension to color encode the view, default colors are assigned to the field's values. Color encodings are shared across multiple worksheets that use the same data source to help you create consistent displays of your data. For example, if you define the Western region to be green, it will automatically be green in all other views in the workbook. You can set the default color encodings for a field by right-clicking the field in the Data window and selecting **Field Properties > Color**. Refer to "Color Shelf" on page 13-23 to learn more about changing mark colors.

## Shapes

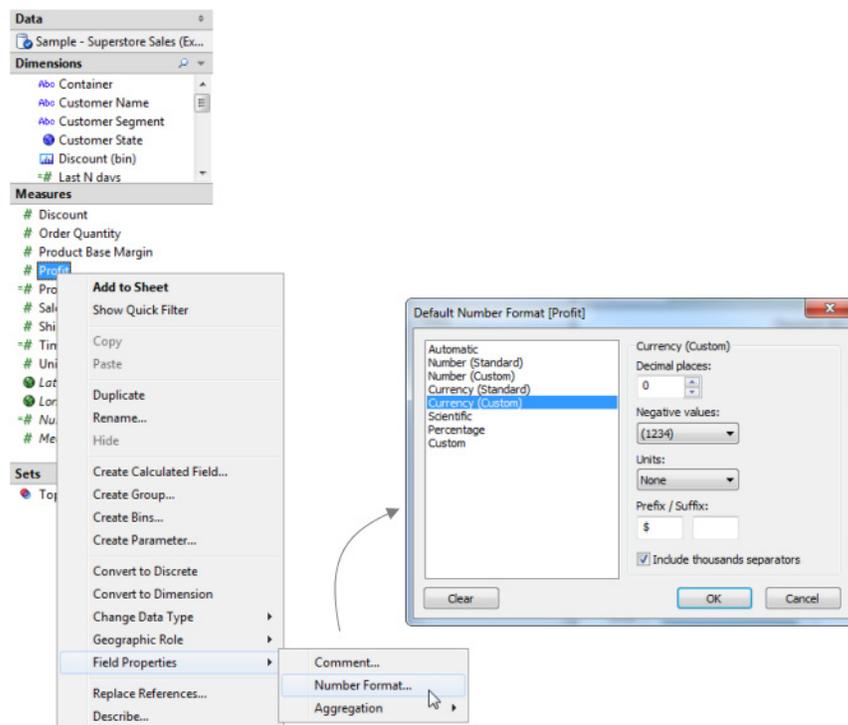
When you use a dimension to shape encode the view, default shapes are assigned to the field's values. Shape encodings are shared across multiple worksheets that use the same data source to help you create consistent displays of your data. For example, if you define the Furniture products are represented with a square mark, it will automatically be changed



to a square mark in all other views in the workbook. You can set the default shape encodings for a field by right-clicking the field in the Data window and selecting **Field Properties > Shape**. Refer to “Editing Shapes” on page 13-40 to learn more about changing mark shapes.

## Formats

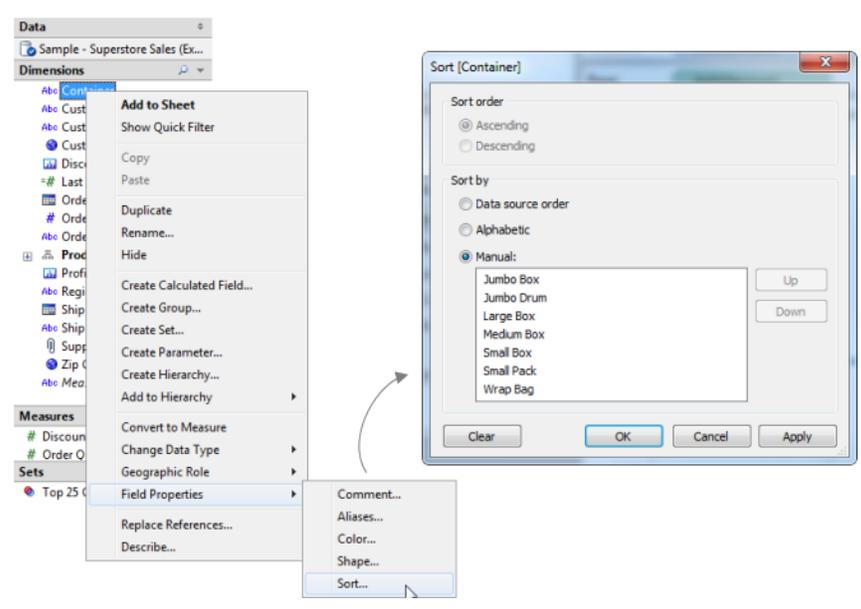
You can set the default text format for date and number fields. For example, you may want to always show the Sales values as currency using the U.S. dollar sign and two decimal points. On the other hand, you may want to always show Discount as a percentage. You can set the default formats by right-clicking a date or numeric field and selecting an option on the **Field Properties** menu. A dialog box opens where you can specify a default format.





## Sort

You can set a default sort order for the values within a categorical field so that every time you use the field in the view, they values will be sorted correctly. For example, let's say you have an Order Priority field that contains the values High, Medium, and Low. When you place these in the view, by default they will be listed as High, Low, Medium because they are shown in alphabetical order. You can set a default sort so that these values are always listed correctly. To set the default sort order right-click a dimension and select **Field Properties > Sort**. Then use the sort dialog box to specify a sort order. Refer to “Sorting” on page 17-3 to learn more about defining a sort order.



---

**Note** The default sort order also controls how field values are listed in a quick filter.

---

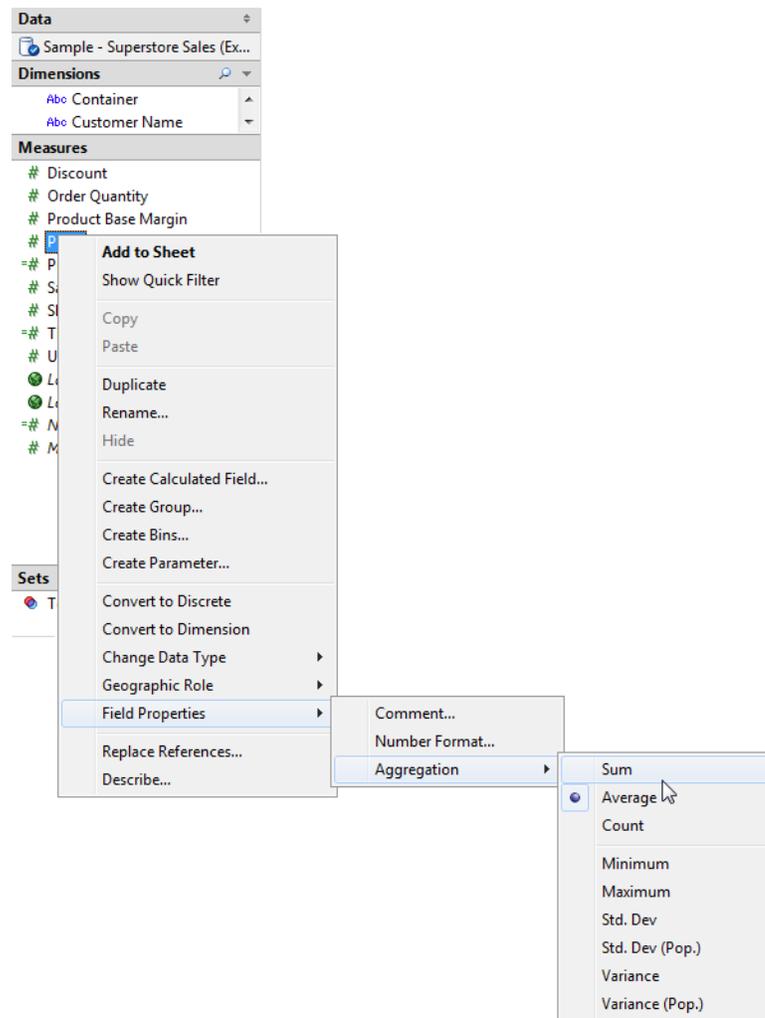
## Aggregation

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 **Field Properties > Aggregation**.
- 2 On the Aggregation 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	For SSAS data sources, this option computes the subtotal on the server.  For Essbase data sources, this option computes the total using the default aggregation determined by the data type (typically SUM).
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.



## 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 profit the header displays as SUM(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 profit of each product category by region.

The screenshot shows the Tableau interface with a text table. The Columns shelf contains 'Product 1 - Category' and the Rows shelf contains 'Region'. The Marks shelf is set to 'Automatic' and displays 'SUM(Profit)'. The table shows profit values for four regions across three product categories: Furniture, Office Supplies, and Technology.

Product 1 - Category			
Region	Furniture	Office Supplies	Technology
Central	44,510	121,977	126,900
East	-12,721	107,471	203,053
South	62,902	140,659	302,118
West	57,685	145,669	223,366

Now suppose you want to show both the Profit and the Sales for each product category and region. When you add the Sales 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.



The screenshot shows the Tableau interface with the following configuration:

- Columns:** Product 1 - Category
- Rows:** Region, Measure Names
- Marks:** Automatic, Text: Measure Values
- Measure Names/Values:** SUM(Profit), SUM(Sales)

Product 1 - Category				
Region		Furniture	Office Supplies	Technology
Central	Sum of Profit	44,510	121,977	126,900
	Sum of Sales	1,321,948	935,971	1,394,018
East	Sum of Profit	-12,721	107,471	203,053
	Sum of Sales	1,130,830	776,050	1,276,109
South	Sum of Profit	62,902	140,659	302,118
	Sum of Sales	1,354,835	990,105	1,787,040
West	Sum of Profit	57,685	145,669	223,366
	Sum of Sales	1,311,641	1,098,101	1,530,404

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**.

Product 1 - Category				
Region		Furniture	Office Supplies	Technology
Central	Profit	44,510	121,977	126,900
	Sales	1,321,948	935,971	1,394,018
East	Profit	-12,721	107,471	203,053
	Sales	1,130,830	776,050	1,276,109
South	Profit	62,902	140,659	302,118
	Sales	1,354,835	990,105	1,787,040
West	Profit	57,685	145,669	223,366
	Sales	1,311,641	1,098,101	1,530,404

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 11-11.

---

## 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:

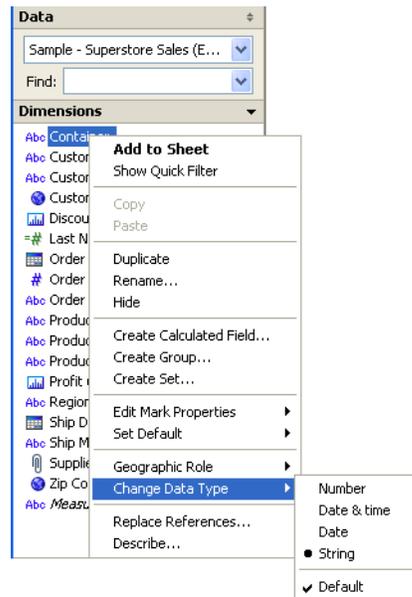
- Data Types
- Data Roles
- Example: Continuous Dimensions and Discrete Measures

### 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)
	Geographic values (used with maps)

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 10-6.

---

### 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.

<b>Mapped Data Type</b>	<b>Treatment of Other Data Types in the Field</b>
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 are dependent 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.

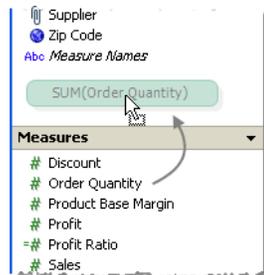
This means that a measure is a function of other dimensions placed on the worksheet. For instance, you might calculate the Sum of “Sales” for every “State”. In this case, the Sales field is acting as a measure because you want to aggregate the field for each state. But measures could also result in a 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**.

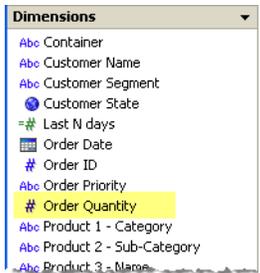


---

**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 **Order Quantity** 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.

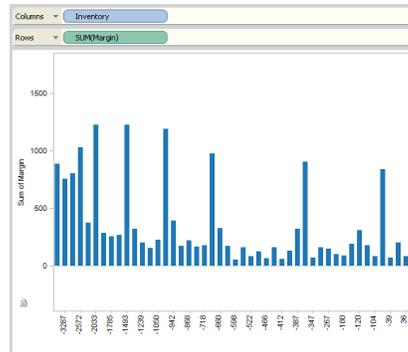
Product 1 - Category			
Order Quantity	Furniture	Office Supplies	Technology
1	Abc	Abc	Abc
2	Abc	Abc	Abc
3	Abc	Abc	Abc
4	Abc	Abc	Abc
5	Abc	Abc	Abc
6	Abc	Abc	Abc
7	Abc	Abc	Abc
8	Abc	Abc	Abc
9	Abc	Abc	Abc
10	Abc	Abc	Abc
11	Abc	Abc	Abc
12	Abc	Abc	Abc
13	Abc	Abc	Abc
14	Abc	Abc	Abc

**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.

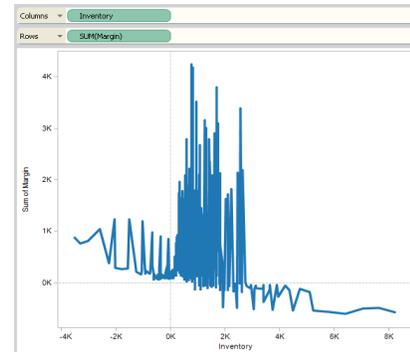


The Inventory dimension is discrete in this example.



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

The Inventory dimension is continuous in this example.



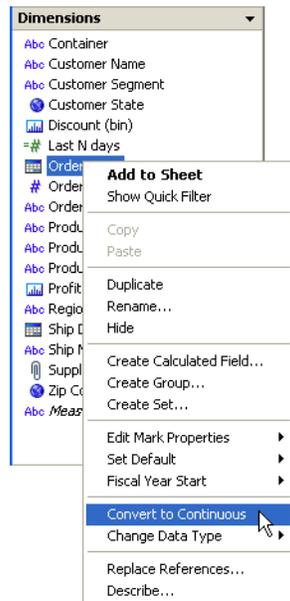
Each inventory value is drawn along a  
continuous axis. Inventory appears **green** on  
the Columns shelf.

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 11-37.

### 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 11-38, 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 **Order Date** field is still displayed in the **Dimensions** area of the Data window, but now uses a green 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 21-4 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.

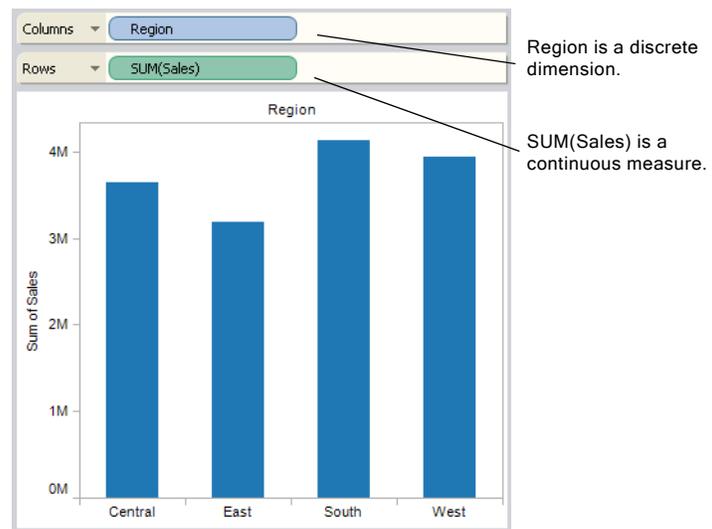


	<b>Discrete</b>	<b>Continuous</b>
<b>Dimensions</b>	<p>Discrete dimensions appear with a blue icon under dimensions.</p> <p>Example: Customer Name, Order Date.</p>  <p>Produces headers when placed on a Row or Column shelf.</p> <p>Discrete Dimension Example</p>	<p>Continuous dimensions appear with a green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



## Discrete Measure

Product Sub-Category	2006	2007	2008	2009
Appliances	13,808	5,499	7,524	21,850
Binders and Binder Accesso..	97,007	23,414	26,887	69,402
Bookcases	42,268	38,774	6,741	24,737
Chairs & Chairmats	32,756	43,074	39,120	21,443
Computer Peripherals	23,198	13,797	33,297	2,721
Copiers and Fax	35,348	64,681		34,671
Envelopes	2,535	5,748	2,179	6,697
Labels	346	593	318	1,310
Office Furnishings	24,666	23,893	20,868	16,974
Office Machines	65,983	43,101	49,812	72,448
Paper	12,412	13,974	12,261	8,976
Pens & Art Supplies	7,739	6,635	3,460	3,366
Rubber Bands	409	821	524	660
Scissors, Rulers and Trimme..	709	398		477
Storage & Organization	51,964	40,548	12,161	20,827
Tables	52,304	62,627	40,178	75,566
Telephones and Communica..	36,209	40,630	27,233	62,298

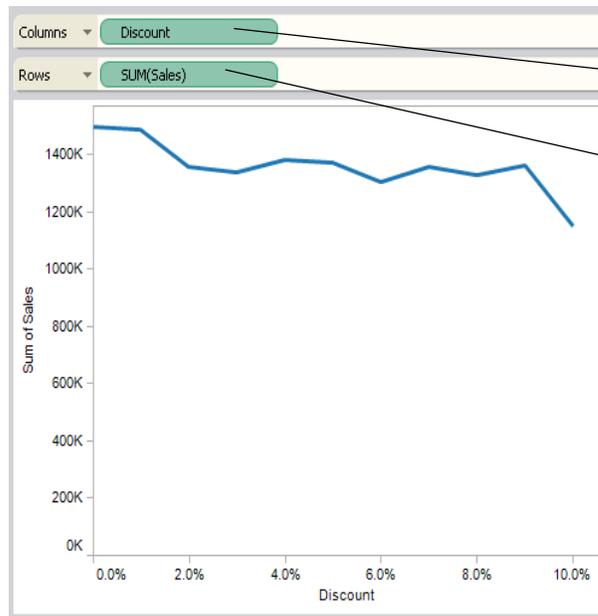
SUM(Sales) is a continuous measure.

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

IIF (Sum({Sales}) > 25000, "Good", "Bad")



## Continuous Dimension

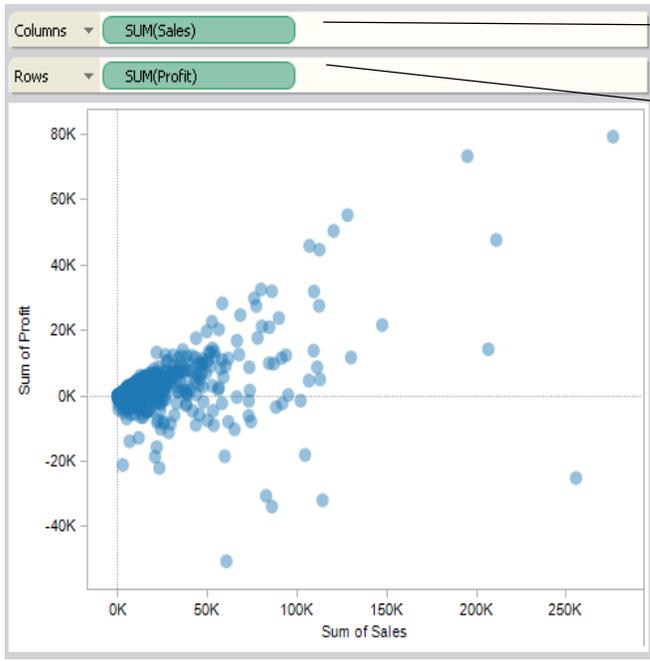


Discount is a continuous dimension.

Sales is a continuous measure.



## Continuous Measure



SUM(Sales) is a continuous measure.

SUM(Profit) is a continuous measure.



## Parts of the View

---

<b>Overview</b> . . . . .	<b>12-2</b>
<b>Headers</b> . . . . .	<b>12-3</b>
<b>Axes</b> . . . . .	<b>12-6</b>
<b>Panes</b> . . . . .	<b>12-9</b>
<b>Cells</b> . . . . .	<b>12-10</b>
<b>Marks</b> . . . . .	<b>12-11</b>
Mark Types . . . . .	12-11
Stacking Marks . . . . .	12-27
Changing Mark Size and Color . . . . .	12-31
<b>Titles</b> . . . . .	<b>12-34</b>
<b>Captions</b> . . . . .	<b>12-35</b>
<b>Field Labels</b> . . . . .	<b>12-36</b>
<b>Legends</b> . . . . .	<b>12-38</b>

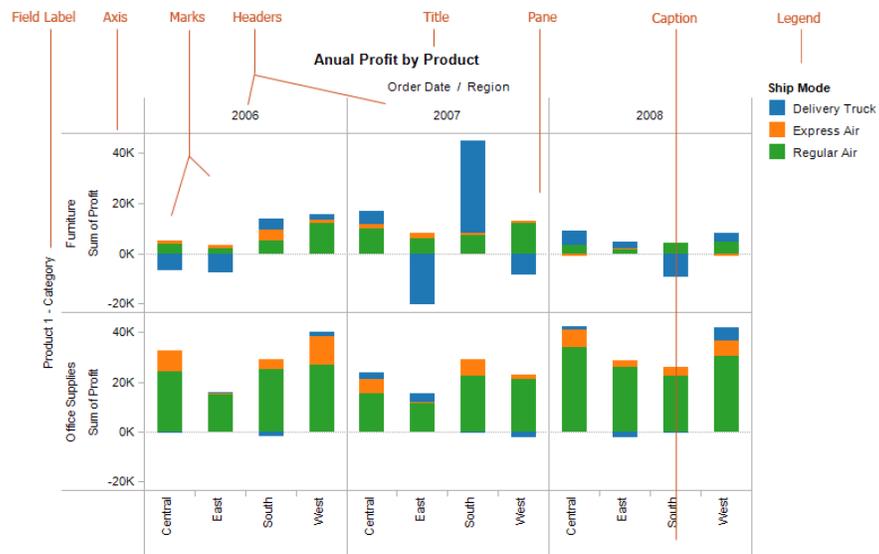
## Overview

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.



Sum of Profit for each Region broken down by Order Date Year vs. Product 1 - Category. Color shows details about Ship Mode. The view is filtered on Order Date Year and Product 1 - Category. The Order Date Year filter keeps 2006, 2007 and 2008. The Product 1 - Category filter keeps Furniture and Office Supplies.



# 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	+ YEAR(Order Date)			
Rows	Product 2 - Sub-Cate..			
	Order Date			
	2006	2007	2008	2009
Appliances	19,849	33,085	21,887	23,202
Bookcases	-3,738	-542	-4,944	-23,947
Chairs & Chairmats	55,817	19,310	40,187	34,625
Computer Peripherals	21,280	15,686	29,252	28,317
Copiers and Fax	33,389	50,521	58,896	24,555
Envelopes	7,451	13,914	12,697	14,649
Labels	2,623	2,361	5,704	3,001
Office Furnishings	24,035	28,456	26,216	18,171



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**.

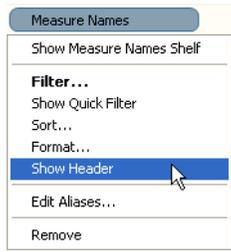
The screenshot shows a PivotTable with a context menu open over the 'Show Header' option. The PivotTable has 'Region' as the row field and 'Sum of Sales' as the value field. The context menu options are: Keep Only, Exclude, Format..., Rotate Label, Show Header (checked), and Edit Alias... The PivotTable data is as follows:

Region	Sum of Sales
Central	Sum of Sales
East	Sum of Sales
South	Sum of Sales
West	Sum of Sales
	Sum of Sales

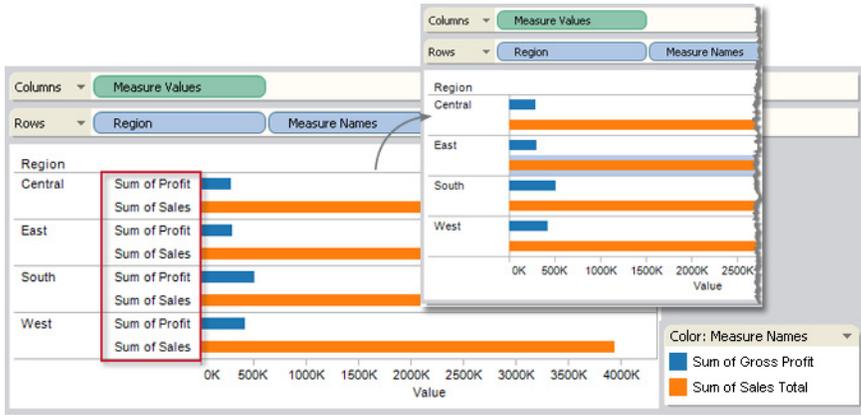


**To show headers:**

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



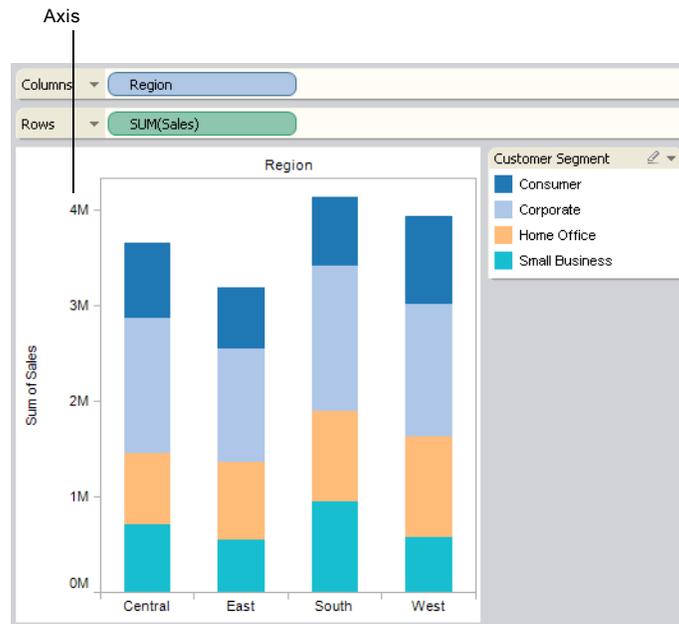
Hiding headers can be really useful when you are working with multiple measures. For example, the view below shows both the sales and profit for each region along a single axis. You can see the view looks cluttered with the Measure Names headers showing. Because Measure Names is also indicated by the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

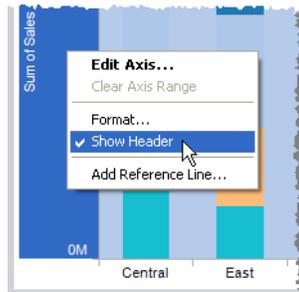




You can show and hide axes at anytime.

**To hide axes:**

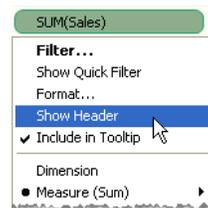
- Right-click the axis in the view and select **Show Header**.





**To show axes:**

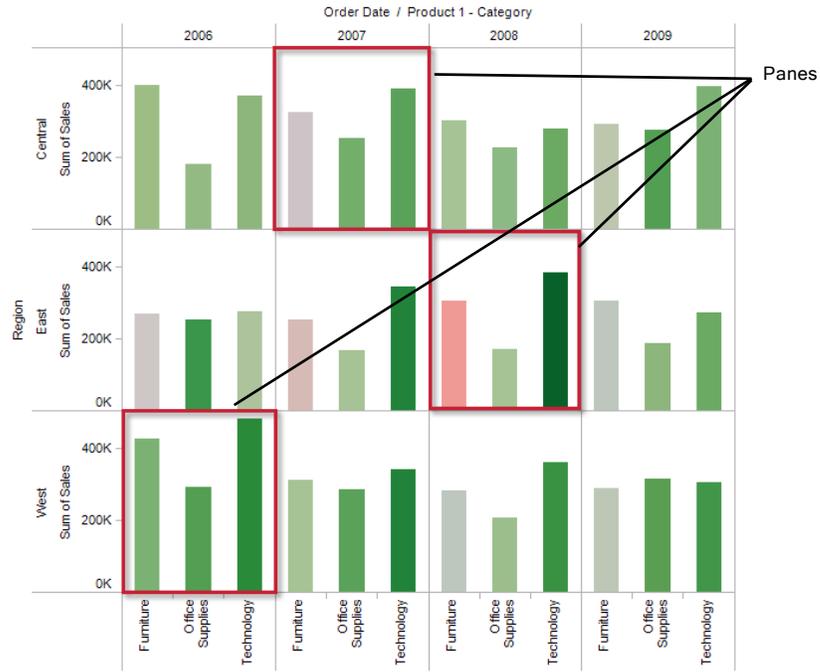
- Select the measure in the view whose axis you want to show and select **Show Header** on the field menu





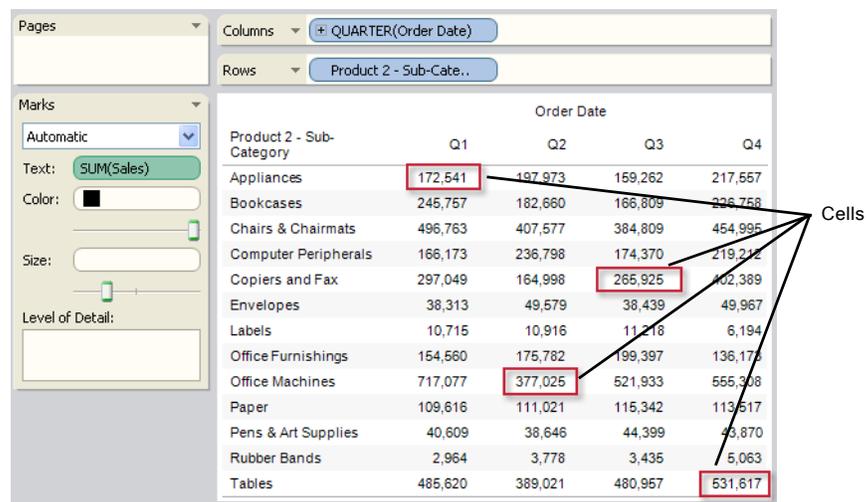
# Panes

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(Order Date)' and the Rows shelf contains 'Product 2 - Sub-Cate...'. The Marks card is set to 'Automatic' with 'SUM(Sales)' as the text. The table displays sales data for various product categories across four quarters (Q1, Q2, Q3, Q4). Several cells are highlighted with red boxes, and a label 'Cells' with arrows points to these specific cells.

Product 2 - Sub-Category	Order Date			
	Q1	Q2	Q3	Q4
Appliances	172,541	197,973	159,262	217,557
Bookcases	245,757	182,660	166,809	226,758
Chairs & Chairmats	496,763	407,577	384,809	454,995
Computer Peripherals	166,173	236,798	174,370	219,242
Copiers and Fax	297,049	164,998	265,925	102,389
Envelopes	38,313	49,579	38,439	49,967
Labels	10,715	10,916	11,218	6,194
Office Furnishings	154,560	175,782	199,397	136,173
Office Machines	717,077	377,025	521,933	555,308
Paper	109,616	111,021	115,342	113,517
Pens & Art Supplies	40,609	38,646	44,399	43,870
Rubber Bands	2,964	3,778	3,435	5,063
Tables	485,620	389,021	480,957	531,617

---

## Marks

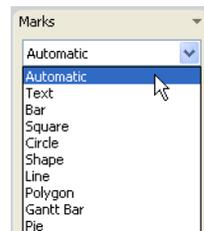
Tableau does not use chart types to build data views. Instead, data are displayed with marks, where every mark corresponds to a row (or a group of rows) in your data source.

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:

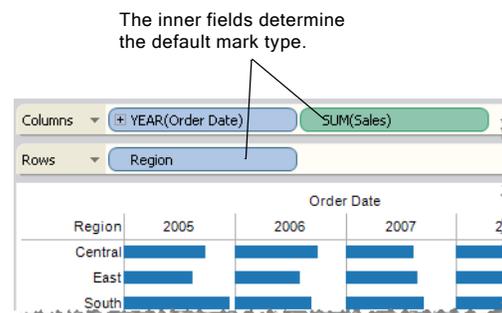
- Automatic Mark
- Text Mark
- Bar Mark
- Square Mark
- Circle Mark
- Shape Mark
- Line Mark
- Polygon Mark
- Gantt Bar Mark



- Pie 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.



Product 2 - Sub-Category	Order Date		
	2005	2006	2007
Appliances	Abc	Abc	Abc
Binders and Binder Accesso..	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** measure, which is aggregated as a summation, is used to complete the table.

Product 2 - Sub-Category	Order Date			
	2006	2007	2008	2009
Appliances	170,657	216,660	168,657	191,359
Binders and Binder Accesso..	360,210	178,639	185,123	300,550
Bookcases	253,942	263,566	140,664	163,810
Chairs & Chairmats	518,011	380,704	451,249	394,181
Computer Peripherals	215,565	175,327	191,042	214,620
Copiers and Fax	273,287	339,713	280,821	236,541
Envelopes	31,376	51,093	39,308	54,521
Labels	10,364	7,393	12,937	8,350
Office Furnishings	194,770	162,197	164,979	143,966
Office Machines	732,505	446,780	428,749	563,308
Paper	118,093	119,516	100,490	111,397
Pens & Art Supplies	42,340	46,728	42,671	35,786
Rubber Bands	3,142	3,754	4,375	3,969
Scissors, Rulers and Trimme..	16,653	6,877	10,645	46,821
Storage & Organization	297,775	289,979	243,495	268,525
Tables	478,981	429,657	500,322	478,255
Telephones and Communica..	472,252	442,048	471,010	504,005

**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 33-22 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** 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.

Product 2 - Sub-Category	Order Date			
	2006	2007	2008	2009
Appliances	████	████	████	████
Binders and Binder Accesso.	████	████	████	████
Bookcases	████	████	████	████
Chairs & Chairmats	████	████	████	████
Computer Peripherals	████	████	████	████
Copiers and Fax	████	████	████	████
Envelopes	████	████	████	████
Labels	████	████	████	████
Office Furnishings	████	████	████	████
Office Machines	████	████	████	████
Paper	████	████	████	████
Pens & Art Supplies	████	████	████	████
Rubber Bands	████	████	████	████
Scissors, Rulers and Trimme.	████	████	████	████
Storage & Organization	████	████	████	████
Tables	████	████	████	████
Telephones and Communica.	████	████	████	████



## 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 12-27.

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.



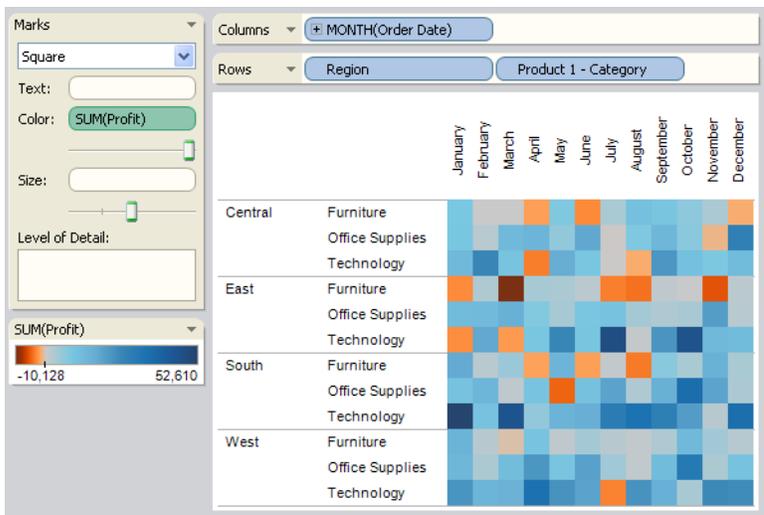
Refer to “Bar Chart” on page 33-4 for examples that show you how to build data views using bars.



## 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.



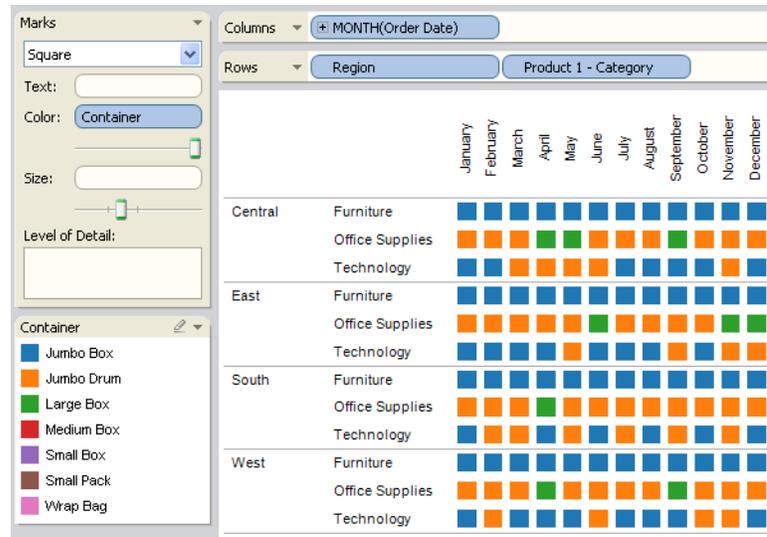
To reproduce this view, select the **Format > Cell Size > Square Cell** menu item and then adjust the size of the squares using the **Size** slider.

Refer to “Heat Map” on page 33-55 for detailed examples that show you how to build data views using squares.

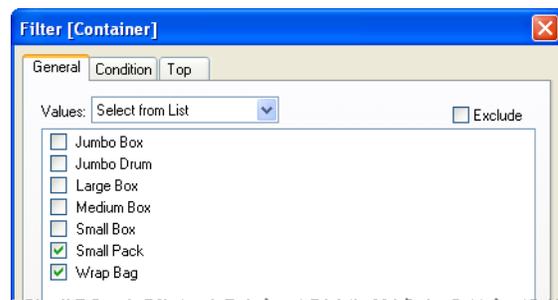


Because of the flexibility of Tableau, you might create a view that contains overlapping data and is difficult to interpret. One way to do this is to place a dimension on the **Color** shelf. A view with overlapping data can be deceptive because only one of the marks for each cell is visible.

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

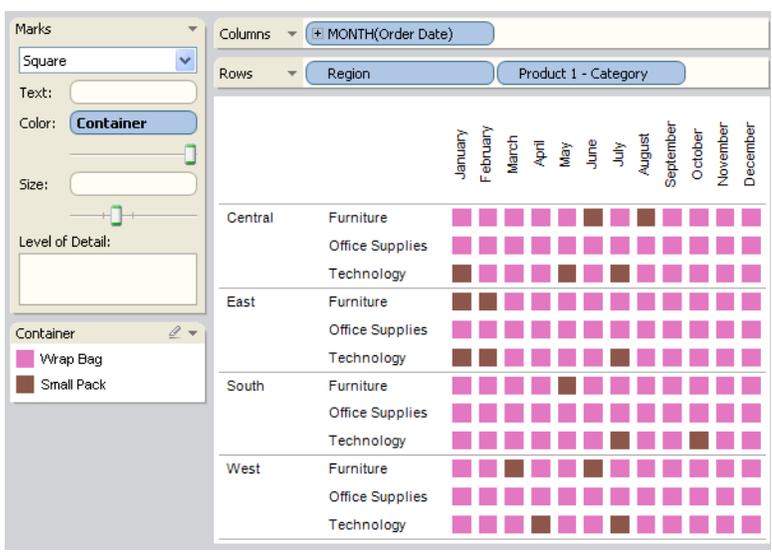


Filter **Container** to only include Small Pack and Wrap Bag.





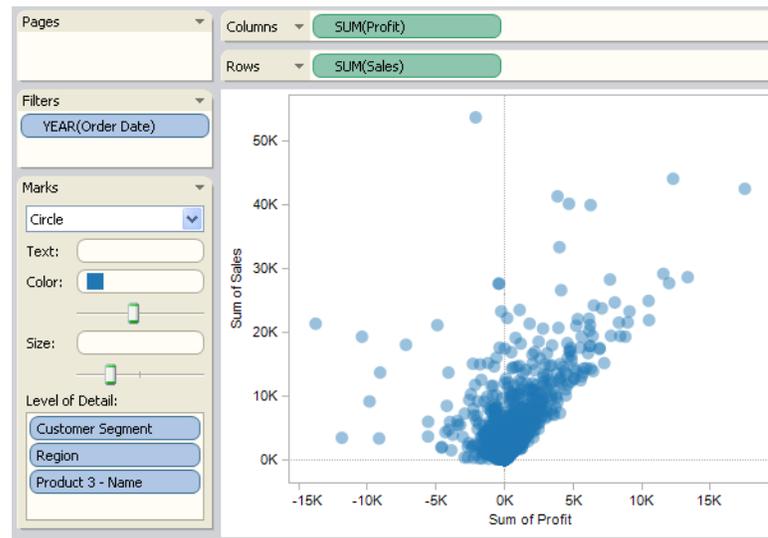
As you can see, all of the squares have changed color showing that the marks overlapped.



### 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).



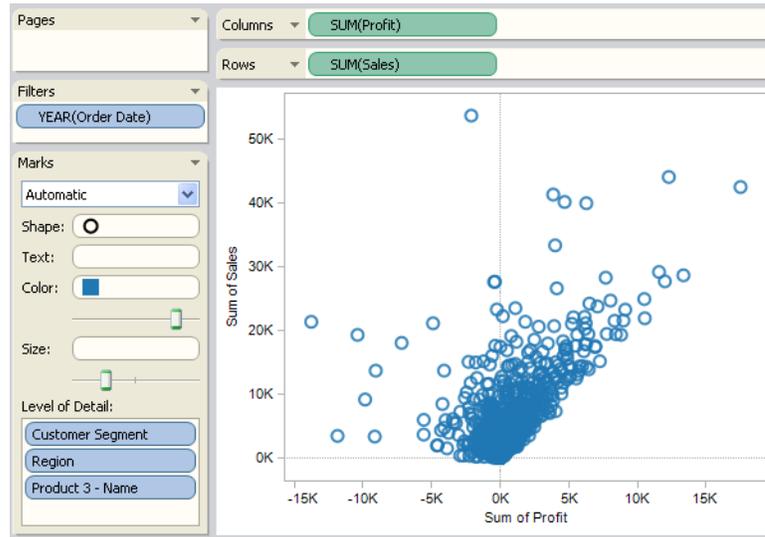
Refer to “Scatter Plot” on page 33-39 for examples that show you how to build a data view using circles.

### Shape Mark

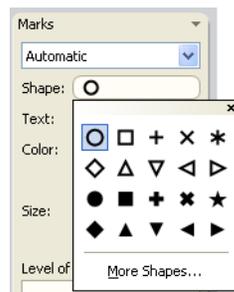
The shape mark type is useful when you want to clearly see individual data points while also viewing categories associated with those points. Tableau displays your data using a shape when:

- 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.

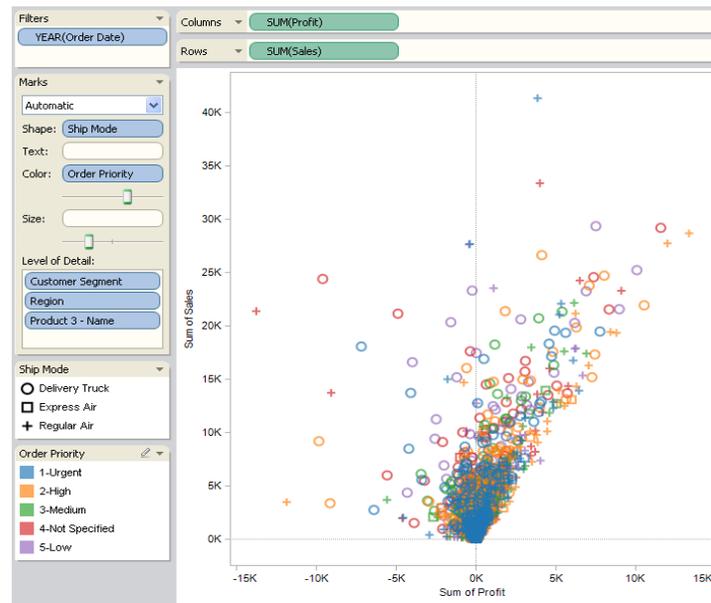


By default, the shape used is an open circle. You can select a different shape by clicking on the shape legend. As shown below, twenty unique shapes are available.



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 **Ship Mode** dimension is used to shape-encode the view.



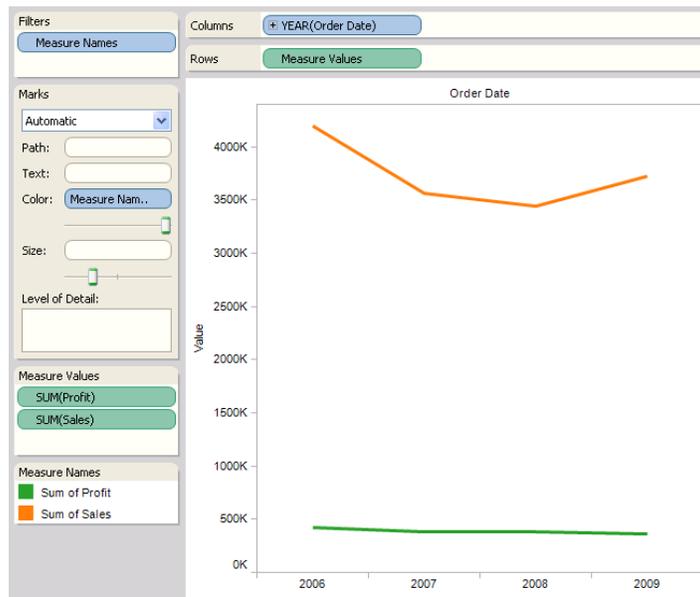
Refer to “Scatter Plot” on page 33-39 for examples that show you how to build data views using shapes.

### 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:

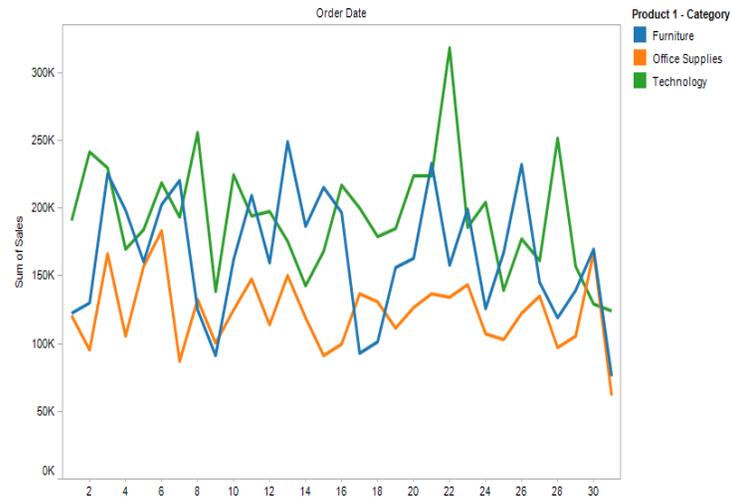
- The **Mark** menu is set to Automatic, and you place one or more measures on either the **Columns** shelf or the **Rows** shelf, and then plot the measures against a date dimension or a continuous dimension.
- You select **Line** from the **Mark** menu.

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 33-37 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.

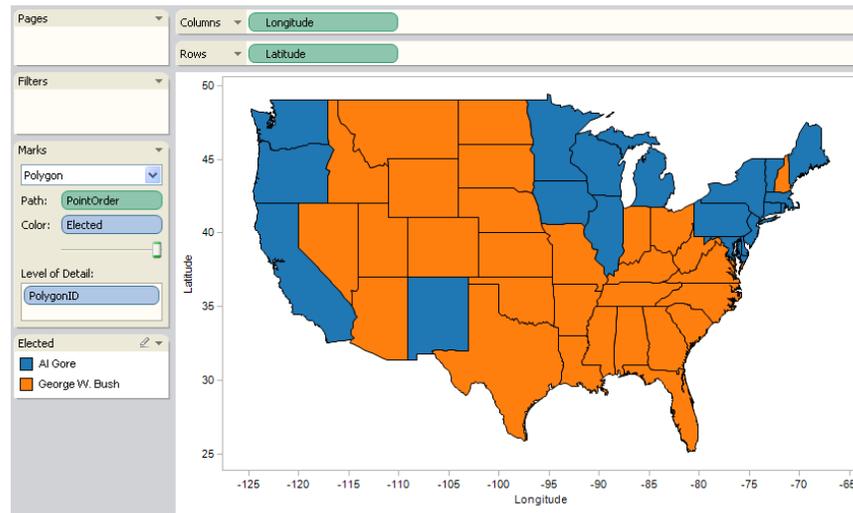
---

**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 **Time to Ship** 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 Mode** dimension on the **Color** shelf, each bar is color-encoded by the ship mode.

### Pie Mark

The pie mark can be used to show proportions. Although generally this type of information can be better shown using stacked bar charts (refer to “Stacking Marks” on page 12-27), there are cases where using pie marks can be very effective. For example, pie marks are very useful when trying to convey the percentage allocation of marketing expenses by state where the spending of geographically close states are very relevant.

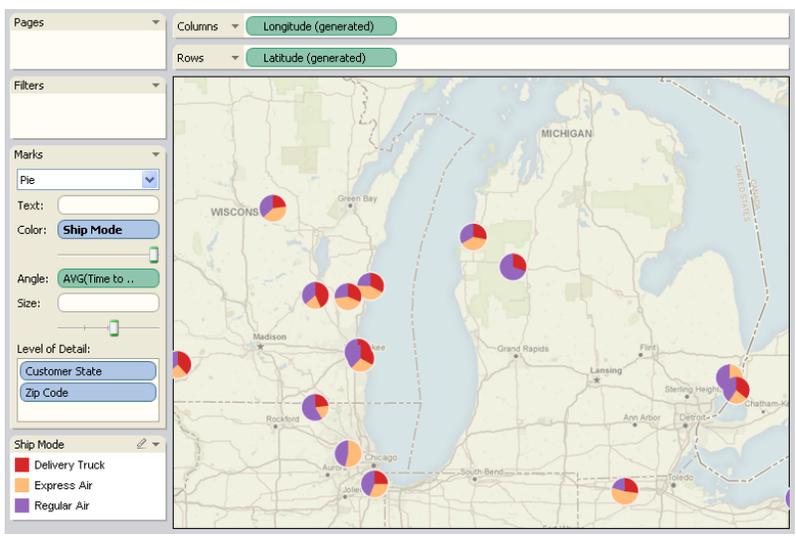
Tableau will never use the pie mark as an automatic mark type, but you can select **Pie** on the **Mark** menu.

When you select the Pie mark type, an additional shelf is available for angle. The Angle shelf determines the angular measure of the pie wedges. For example if you place a measure



such as Sales on the angle shelf, the total 360 degrees of the pie corresponds to the total sum of sales and each wedge is divided for the values of the field on the Color shelf.

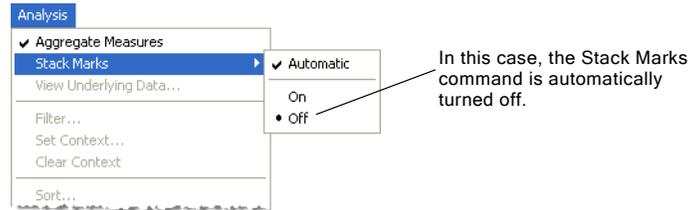
The view below shows the time it took to ship products by various ship modes. The data overlays a map and shows the information by zip code. We can quickly see that Regular Air takes the longest to ship in this particular region except in the south part of Michigan where they seem to have optimized for that ship mode.



## 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.

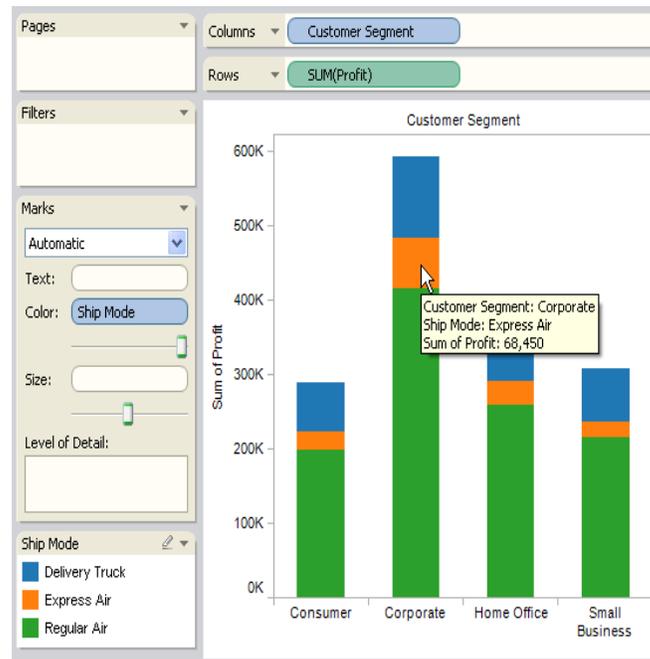


If you select **On** or **Off** on the **Stack Marks** menu, you are switched to manual mode. Your selection remains throughout any changes you make to the view.

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.





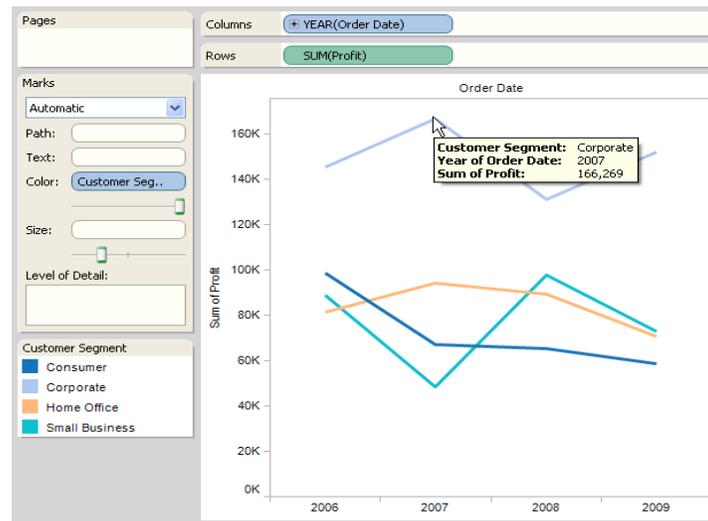
Because the mark type is a bar, Tableau automatically stacks the marks. This means that the marks are drawn cumulatively and the height of each stacked segment within each bar represents the value for that segment. For example, the sum of the profit for products shipped by Express Air (orange bar segment) in the Corporate market is \$68,450.

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.



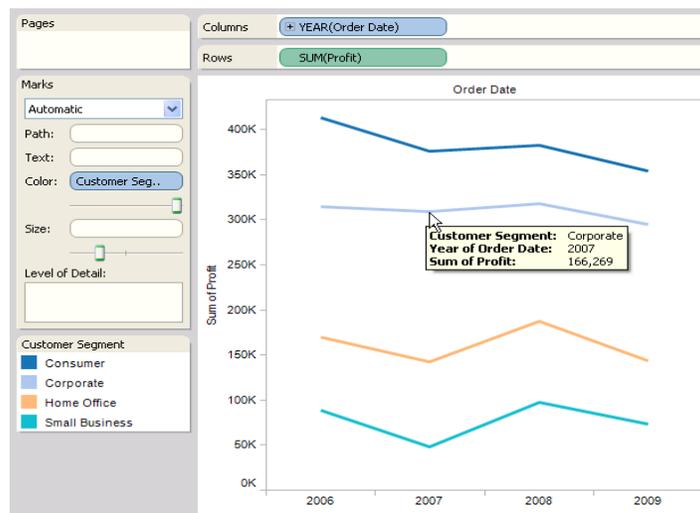
### Example – Stacking Lines

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 2007, the Corporate (light blue) sales totaled \$166,269. That is, the space between that data point and the horizontal axis is equal to the sum of the sales 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 (teal) line is the same as its un-stacked version because it's at the bottom of the stacking list. The stacked Home Office (peach) line is derived by adding its un-stacked values to the un-stacked Small Business values. The stacked Corporate (light blue) 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 profit. The lines are no longer all compared to the horizontal axis.

For example, notice that the tooltip for the 2007 Corporate data still shows the profit as \$166,269. The interpretation is that the space between the Corporate data and the Home Office data yields the sum of the profit 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.

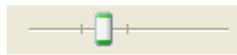


You can also display or remove mark borders. Refer to “Mark Borders” on page 13-28 for more information. This section discusses the following topics:

- 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.
Pie	Moving the slider makes the overall size of the pie bigger and smaller.

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 27-1 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 13-31 for more information.

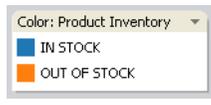
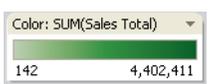


## Changing Mark Color

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, select <b>Format &gt; Marks</b> and modify the color in the Format window.
	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 13-23 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



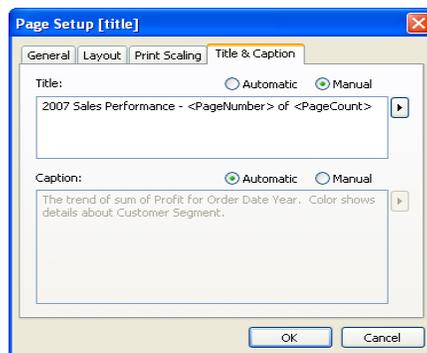
A screenshot of a dashboard title card. The title '2004 Sales Performance' is displayed in a red-bordered box above a data table.

	Quarter of Order Date					2001	2002	2003	2004
Product Catego...	Q1	Q2	Q3	Q4	APPLIANCES	10,371	-115	7,385	5,485
APPLIANCES	2,445	6,206	9,245	6,229					
BOOKCASES	-1,190	-3,164	-2,244	-3,043	BOOKCASES	-5,535	-957		2,663

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 27-23 to learn more.

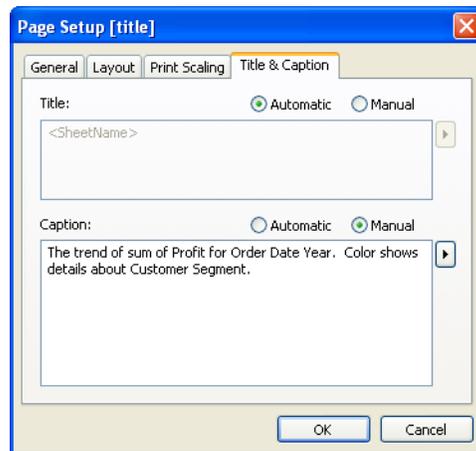
---

## 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.



Use the arrow to the right of the text box to add automatic text such as page number, sheet name, page count, and more.

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 27-23 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.

	East				West			
Product 1 - Category	2006	2007	2008	2009	2006	2007	2008	2009
Furniture	-1,025	-4,281	-11,996	4,580	30,134	15,679	4,687	7,184
Office Suppli..	49,760	16,089	15,547	26,075	42,990	40,164	20,710	41,805
Technology	13,707	62,681	91,234	35,431	56,813	63,373	51,240	48,087

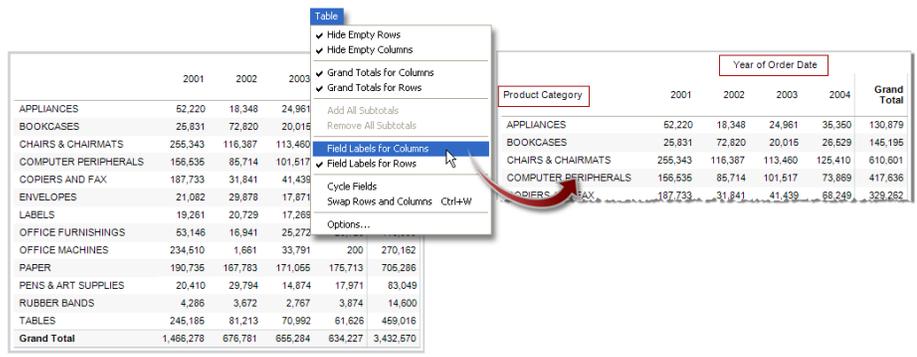
Field labels apply only to discrete fields. When you add continuous fields to the view, an axis is created. The axis is labeled with a header. Refer to “Axes” on page 12-6 to learn more about axes.



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**.



You can format the fonts, alignment, shading, and separators for field labels. Refer to “Field Labels” on page 27-20 to learn how.

---

## 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 17-16 to learn how to sort using legends.

Refer to “Filtering Dimensions” on page 16-7 to learn how to filter using legends.

Refer to “Color Legend Highlighting” on page 20-9 to learn how to highlight select items in the view.

# Building Views Manually

---

<b>Dragging Fields</b> . . . . .	<b>13-3</b>
<b>Types of Shelves</b> . . . . .	<b>13-10</b>
Columns and Rows Shelves . . . . .	13-10
Pages Shelf . . . . .	13-14
Filters Shelf . . . . .	13-18
Level of Detail Shelf . . . . .	13-21
Color Shelf . . . . .	13-23
Size Shelf . . . . .	13-31
Shape Shelf . . . . .	13-38
Label Shelf . . . . .	13-47
Path Shelf . . . . .	13-49
<b>Multidimensional Hierarchies</b> . . . . .	<b>13-51</b>
Drilling Down and Up in a Hierarchy . . . . .	13-51
Building Views with Oracle Essbase . . . . .	13-54
Building Views with Microsoft Analysis Services Cubes . . . . .	13-58
Perfect Pivoting . . . . .	13-59
Defining Unique Values . . . . .	13-60
<b>Working with Large Views</b> . . . . .	<b>13-66</b>
<b>Example – Building Data Views Manually</b> . . . . .	<b>13-67</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:

- Dragging Fields
- Types of Shelves
- Multidimensional Hierarchies
- Working with Large Views
- Example – Building Data Views Manually

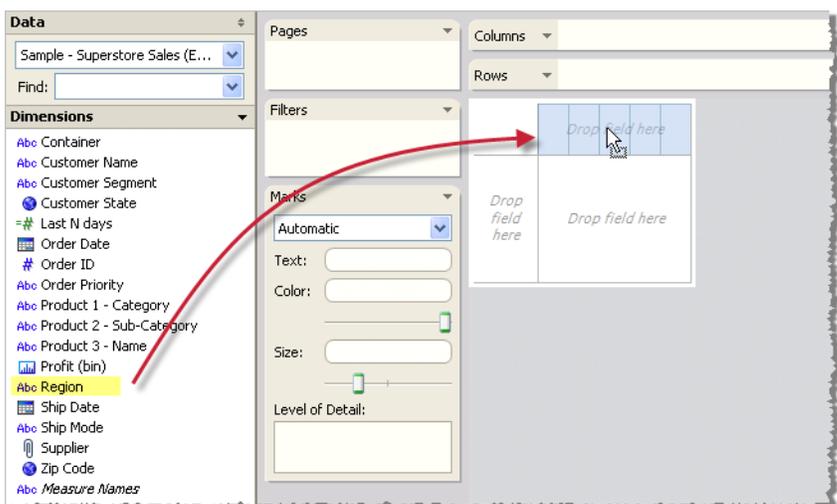


# Dragging Fields

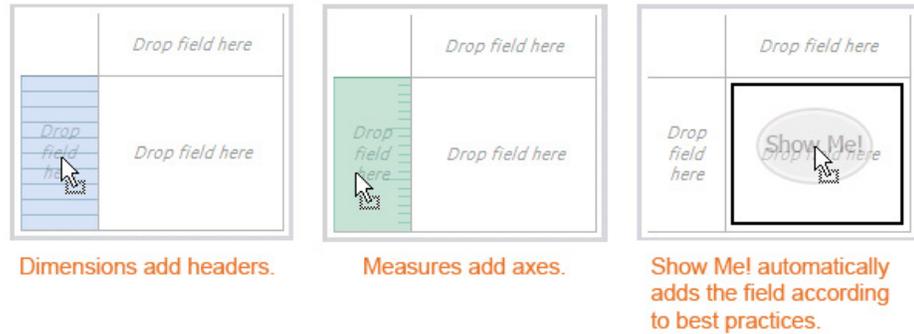
You can build views of your data by dragging fields from the Data window to the view. You can drag fields to a variety of active areas in the view or place them on the shelves that are part of every worksheet.

## The Basics

When you begin creating a new data view on a blank worksheet, drag a field from the Data window to and drop it in the view.

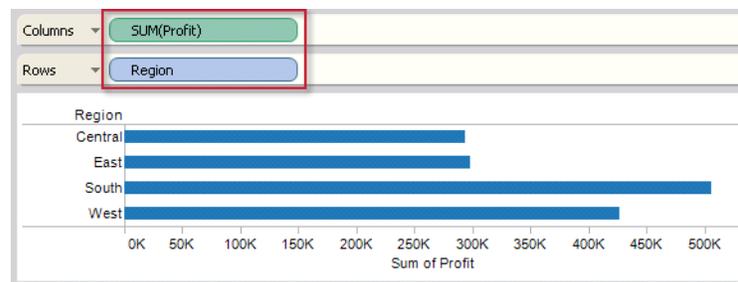


While dragging fields you can pause on the active areas in the view to see how the field will be added to the view. For example, in general dimensions will add row and column headers to the view while measures add continuous axes. Below are some examples of how fields can be added to the view.



For a more advanced discussion of dimensions and measures, refer to “Data Roles” on page 11-37.

When you drag a field to one of the active areas in the view, the field is added to the view and displays on one of the shelves. For example, in the view below the Regions are shown as Rows and Profit is shown as a Column with a continuous axis.



You can drag fields directly to the shelves instead of the active areas in the view. You can also drag fields from one shelf to another shelf. 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. Refer to “Types of Shelves” on page 13-10 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.



## Adding More Fields

You can add as many fields as necessary by dragging and dropping them on the different areas of the view. Once there are more fields in the view there are some extra active areas. For example you can add replace fields by dropping them on existing headers and axes in the view. Or instead of replacing the field you can blend multiple measures onto a single axis. Finally, you can rearrange the rows and columns in the view.

## Adding Headers Using Dimensions

You can add headers by dragging a dimension and dropping on either side of existing headers, or to the left of an axis. For example, in the view below you can add the Region field by dragging it and dropping it to the right of the product names.

The screenshot shows the Tableau interface with a pivot table. The 'Columns' shelf contains 'YEAR(Order Date)' and the 'Rows' shelf contains 'Product 2 - Sub-Category'. The 'Filters' shelf contains 'Region: East' and 'Product 2 - Sub-Category'. The 'Marks' shelf is set to 'Automatic' with the text 'SUM(Sales)'. The pivot table displays sales data for various product sub-categories across the years 2005, 2006, 2007, and 2008. A red circle highlights a dotted black line in the table, indicating an active area for adding headers. A red arrow points from the 'Region' dimension in the left pane to this active area.

Product 2 - Sub-Category	Order Date			
	2005	2006	2007	2008
Appliances	53,154	19,509	43,027	36,145
Binders and Binder Accesso...	99,622	28,031	30,620	71,247
Bookcases	72,633	45,440	35,139	59,688
Chairs & Chairs Mats	85,827	74,620	124,820	131,466
Computer Peripherals	32,840	35,923	24,530	51,962
Copiers and Fax	62,134	73,304	120,520	7,371
Envelopes	10,473	7,169	5,186	5,330
Labels	2,249	1,375	4,564	2,796
Office Furnishings	28,496	37,515	39,279	26,387
Office Machines	105,656	137,628	96,914	109,146
Paper	24,446	27,649	26,419	24,203
Pens & Art Supplies	4,524	12,132	5,236	10,515
Rubber Bands	462	1,817	955	550
Scissors, Rulers and Trimme...	1,383	1,322	1,870	3,066

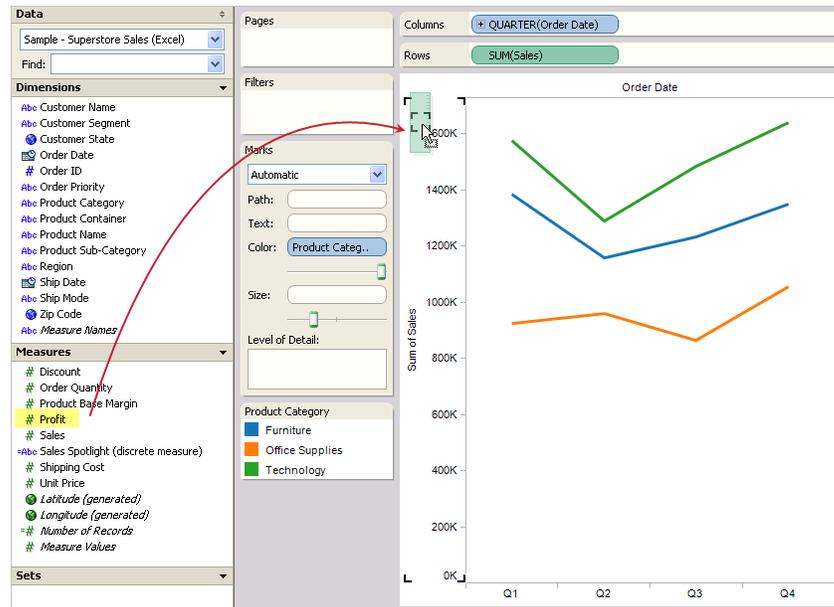
You can see that as you hover over the view, a dotted black line indicates active areas where you can add headers.

## Adding Axes Using Measures

You can add axes by dragging a measure and dropping it on an active area in the view. If an axis already exists in the view you can replace the existing axis, blend the new measure with the existing axis, or add a secondary axis.

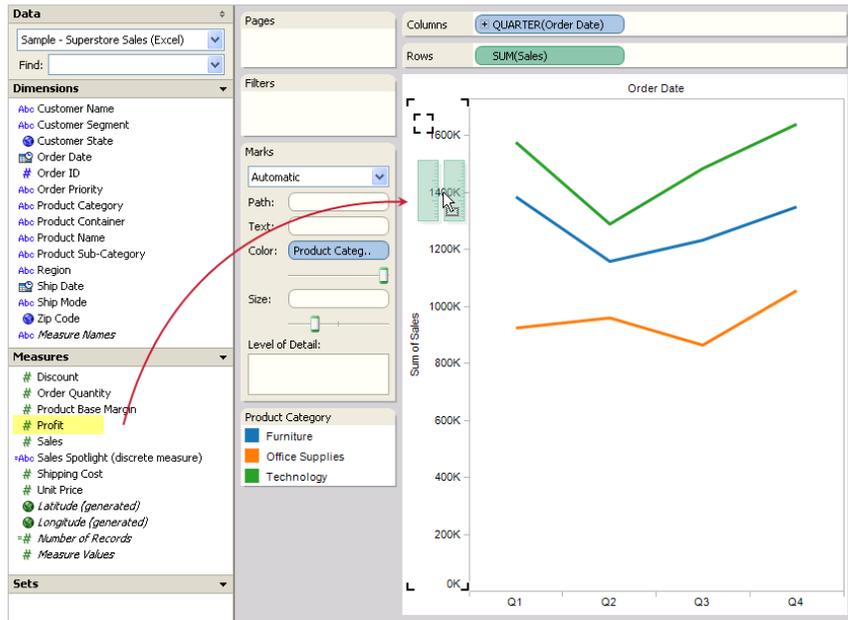


**Replace the Existing Axis.** Drag the new measure to the top left portion of the axis in the view. A small square drop zone appears and a single axis icon displays to indicate that a single axis will be left when you drop the measure.



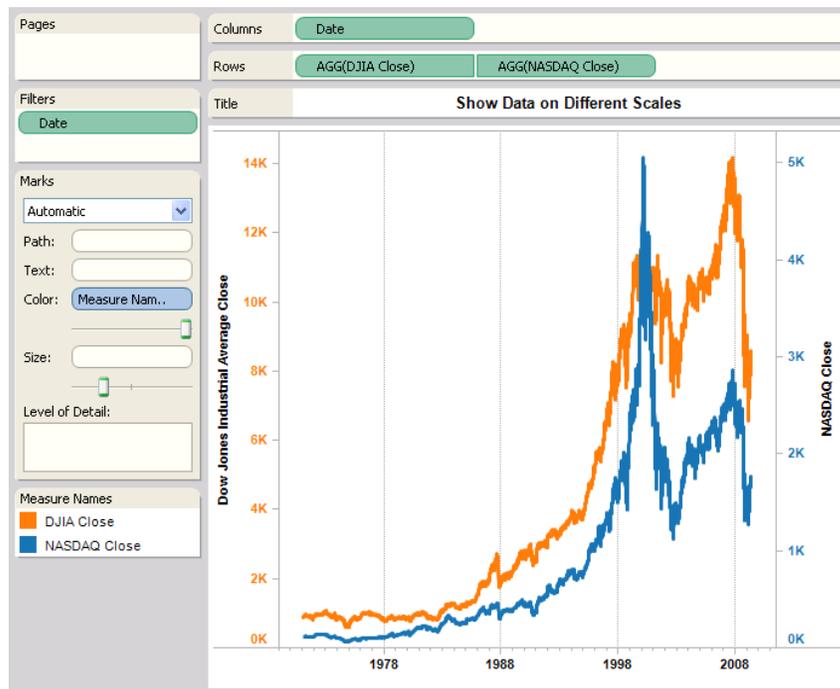


**Blend the Measures on Single Axis.** You can show multiple measures on a single axis by dragging the new measure directly on top of the existing axis. Blending measures uses the Measure Names and Measure Values fields. For a more details example of blending measures refer to “Line Chart–Filter and Color-Encode Multiple Measures” on page 33-35.





**Add a Secondary Axis.** Drag the field to the right side of the view to add the measure as a secondary axis. Secondary Axes are useful when you want to compare two fields that have different scales. In this case, blending these axes would distort the view. Instead you can add a secondary axis. You can add up to four axes to the view: two on the Columns shelf and two on the Rows shelf. Below is an example of a secondary axis view showing the Dow Jones Industrial Average and NASDAQ close values over time.





## Rearranging the Rows and Columns

Finally, you can rearrange the rows and columns in the view by dragging the selection border for headers or an axis.

Use the selection handle to rearrange the rows and columns in the view.

Region	Product 1 - Cate..	2005	2006	2007	2008
Central	Furniture	401,702	325,745	301,788	292,712
	Office Supplies	179,753	251,679	227,352	277,186
	Technology	372,012	388,864	277,817	355,325
East	Furniture	269,097	251,486	304,332	306,916
	Office Supplies	252,122	165,750	168,050	190,128
	Technology	277,221	344,862	371,579	282,447
South	Furniture	347,066	347,062	367,383	293,323
	Office Supplies	326,182	219,318	202,805	241,800
	Technology	563,114	328,583	350,882	544,462
West	Furniture	427,839	311,832	283,711	288,259
	Office Supplies	292,553	283,892	207,766	313,891
	Technology	481,263	341,558	359,687	347,895

---

## 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
- Label 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.

Some shelves are available only when certain mark types are used. For example, the **Shape** shelf appears when the shape mark type is used. Additionally, some shelves are not particularly useful with certain mark types. Refer to “Mark Types” on page 12-11 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. Tableau can also help you determine the best way to display your data using Show Me! Refer to “Building Views Automatically” on page 14-1 to learn more.

### 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 **Customer Segment** dimension are displayed as column headers, while the **Profit** measure is displayed as a vertical quantitative axis.

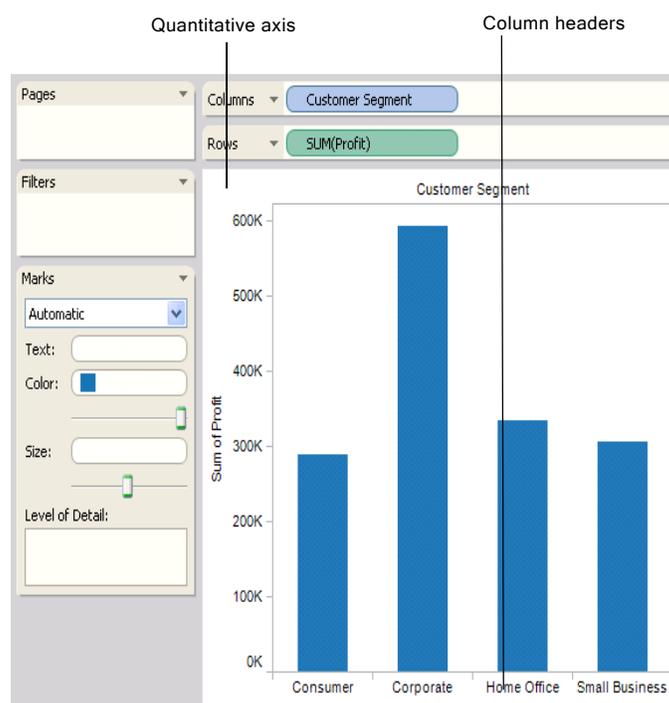
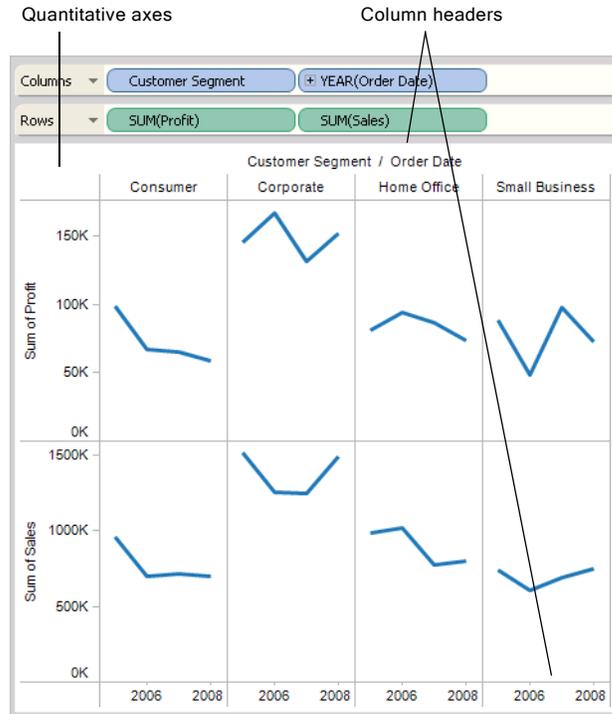


Tableau displays data using marks, where every mark corresponds to a row (or a group of rows) in your data source. The inner fields on the **Rows** and **Columns** shelves determine the default mark type. For example, if the inner fields are a measure and a dimension, the default mark type is a bar. You can manually select a different mark type using the **Mark** menu. Refer to “Mark Types” on page 12-11 for more information.



Adding more fields to the **Rows** and **Columns** shelves adds more rows, columns, and panes to the table.



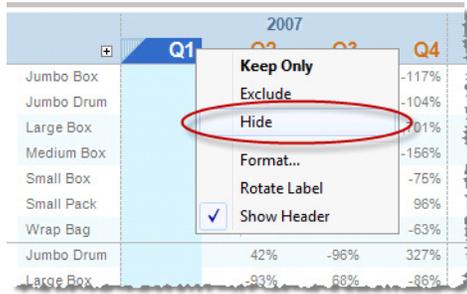


### Hide Rows and Columns

Generally you will add dimensions and measures to create the rows and columns of the table and you'll either include all data or add filters to only show a subset. However, when you filter data it is also excluded from calculations and other computations performed on the summarized data in the table. For example, Table Calculations depend on the data shown in the view for computations such as year over year growth and running totals. In these cases you can hide the rows and columns that you don't want to show without changing the calculation.

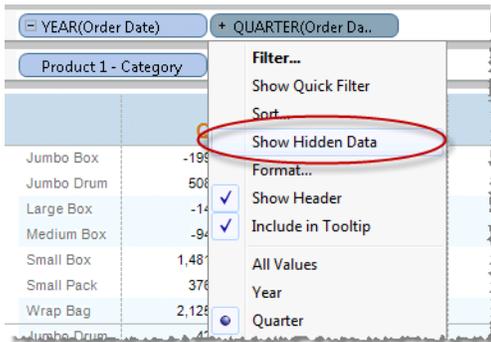
#### To hide a row or column:

- Right-click the row or column you want to hide and then select **Hide**.



#### To show hidden data:

- Open the field menu for a field that has hidden columns or rows and select **Show Hidden Data**.





Hiding columns is especially useful when using table calculations that compare to previous or next. In that case, there is always a row or column that doesn't show data because there is no data to compare to. You can simply hide the empty column without modifying the table calculation.

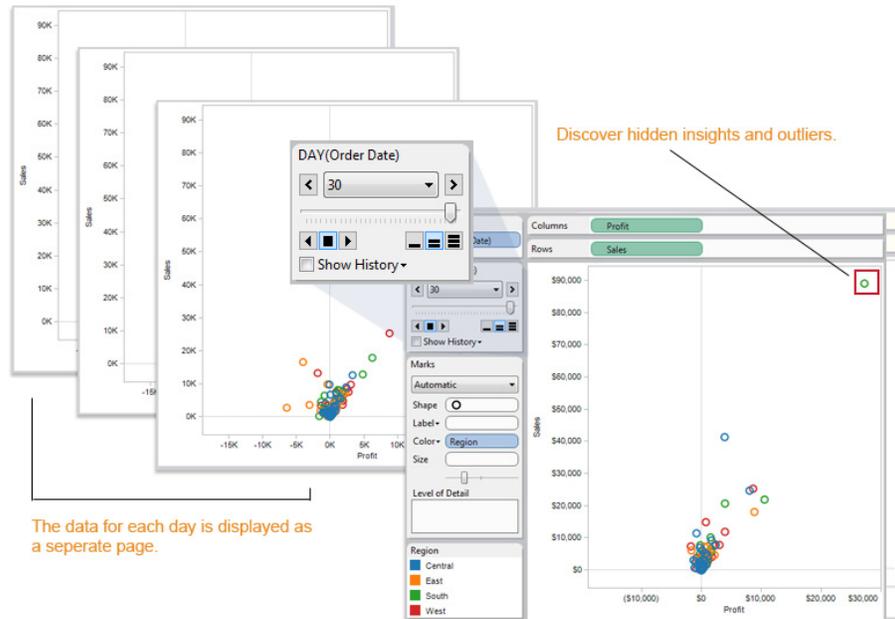
## Pages Shelf

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.



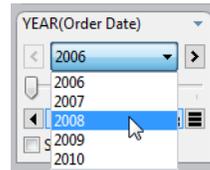
The data for each day is displayed as a separate page.

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.

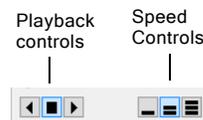


**Manually Advance through the pages.** You can manually advance through the sequence of pages by doing one of the following:

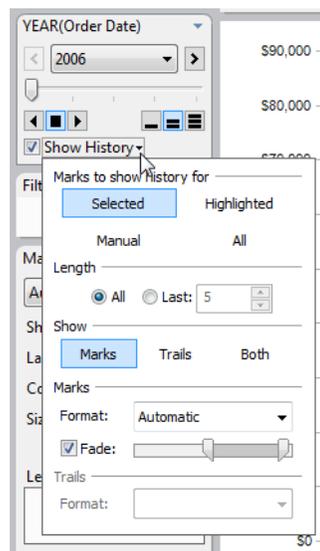
- Use the forward and back buttons on either side of the drop-down list to navigate through the pages one at a time.
- Use the Page Slider to quickly scroll forward and backward in the sequence of pages.
- Use the keyboard shortcuts below to scroll forward and backward in the sequence of pages.

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. 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.



**Page History.** Optionally show page history using the Show History checkbox. When you show history, marks from previous pages are shown in addition to the previous page. Open the drop-down control for history to specify what marks to show and when to show them.



The history drop-down control has the following options:

- Marks to show history for – select whether you want to show history for just selected marks, highlighted marks, marks that you’ve manually selected to show history for, or all marks. You manually show history for marks by right-clicking the mark in the view and selecting an option on the Page History menu.
- Length – select the number of pages to show in the history.
- Show – specify whether to show the historical marks, a line tracing through the previous values (trails), or both.
- Marks – format the historical marks including the color and how much to fade them. If the color is set to automatic, the marks will either use the default mark color or the color encoding on the Color shelf.
- Trails – format the lines that are drawn through the historical marks. This option is only available if Trails is selected in the Show options.

---

**Note** Page trails may not display if there are multiple marks per color on a page. Make sure that the level of detail for the view is less than or equal to the level of detail on the pages and color shelves.

---



## Filters Shelf

The **Filters** shelf allows you to specify which data to include and exclude. For example, you might want to analyze the profit for each customer segment, but only for certain shipping containers and delivery times. By placing fields on the **Filters** shelf, you can create such a view.

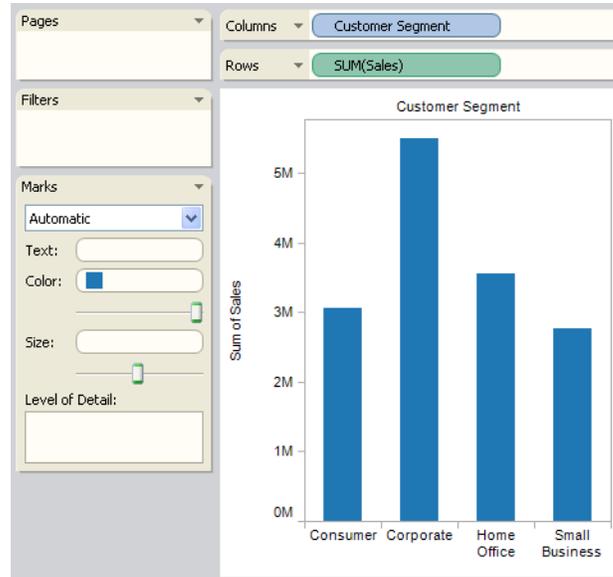
---

**Note** This section presents a brief overview of filtering. Refer to “Groups” on page 17-18 for a complete description.

---

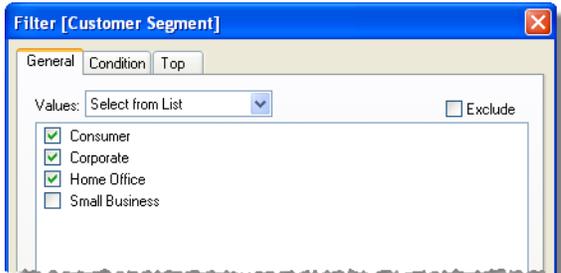
You can filter data using measures, dimensions, or both at the same time. Additionally, you can filter data based on the fields that make up the columns and rows of the table. This is called an *internal filter*. You can also filter data using fields that don't contribute headers or axes to the table. 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.

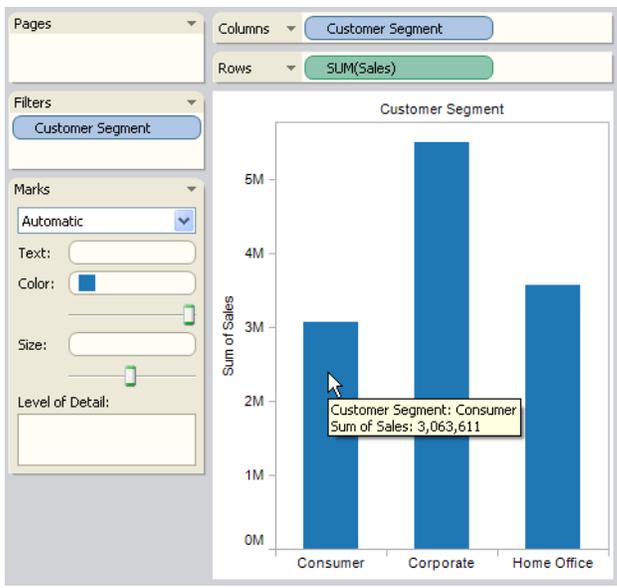




Suppose you are not interested in the Small Business data. You can remove this column from the view by filtering the **Customer Segment** dimension. To do so, select **Filter** from the field's context menu. The **Filter** dialog box opens. By default all members are selected. Un-check **Small Business** to exclude it from the view. All selected members will be included.



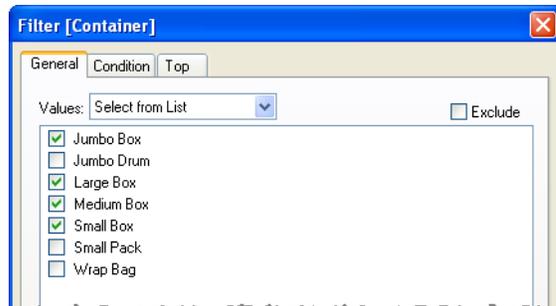
As shown below, **Customer Segment** is automatically placed on the **Filters** shelf, and the view now contains three columns instead of the previous four.



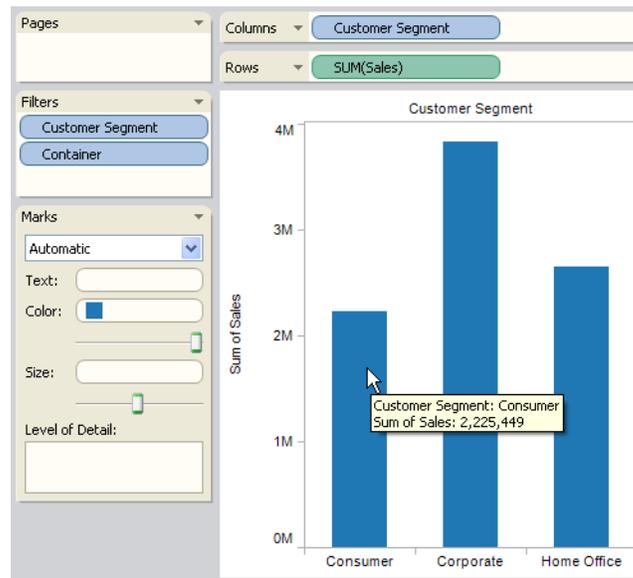


Suppose you want to only view sales for products that were shipped in boxes. To do this, place the **Container** dimension directly on the **Filters** shelf. This is an example of an external filter because **Container** is not part of the view. That is, it does not contribute row or column headers.

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 Consumer segment has decreased to \$2,225,449. 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.



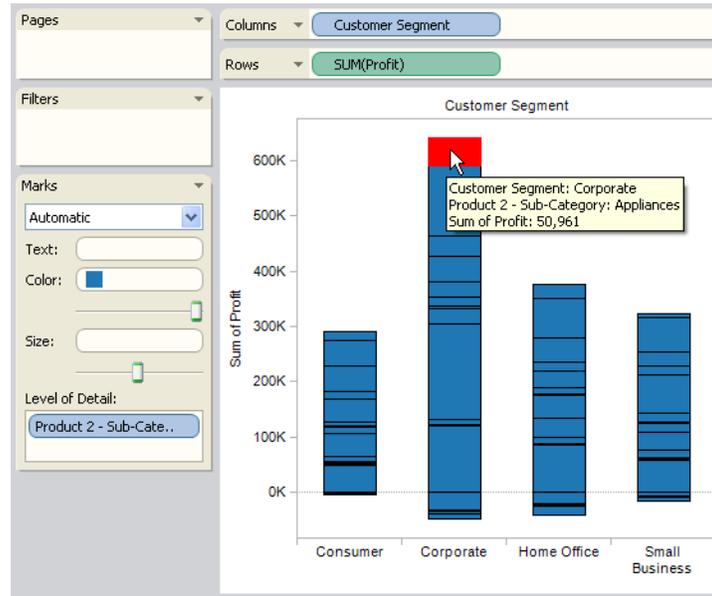
The order of fields placed on the **Filters** shelf does not affect the data view because the filters are independent. The result of filtering by customer segment, and then by container is the same as filtering by container and then by customer segment.

## 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 **Customer 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 2 - Sub-Category** dimension. The size of each segment reflects the contribution to the profit for a particular member. For example, the view below shows that Appliances category in the Corporate market has a profit of \$50,960.





---

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

All marks have a default color that is used when there are no fields on the color shelf. Most marks use a blue color while text marks are shown in black.

The **Color** shelf encodes data by assigning different colors to the marks in a data view based on the values of a field. The effect of color-encoding your data view depends on whether you use categorical or quantitative colors. You can also use the drop-down control to specify other color properties such as transparency, borders, and halos. The color shelf is discussed in the following topics:

- Categorical Colors
- Quantitative Colors
- Transparency
- Effects

---

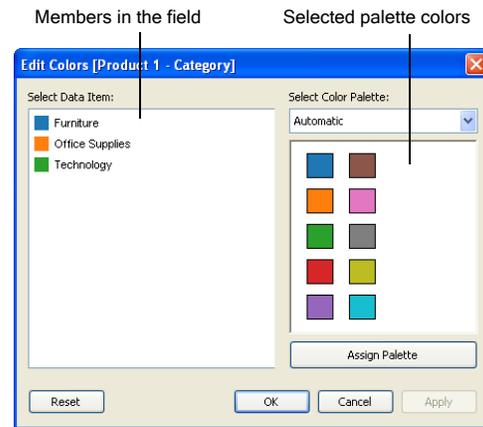
**Note** Color encodings are shared across multiple worksheets that use the same data source to help you create consistent displays of your data. For example, if you define the Western region to be green, it will automatically be green in all other views in the workbook. You can set the default color encodings for a field by right-clicking the field in the Data window and selecting **Edit encodings > Color**.

---



## 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.

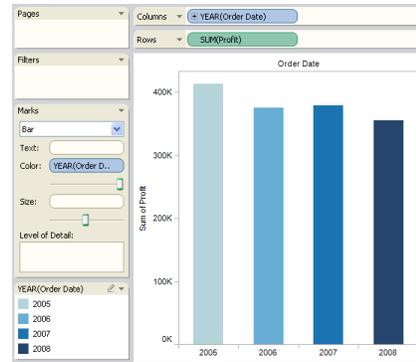
You can select a different color palette from the drop down list in the upper right of the Edit Color dialog box. Select from either categorical palettes or ordinal palettes. A categorical palette, such as **Tableau 20** contains several distinct colors that can be assigned to dimension members that 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: Product1 - Sub-Category has no natural order.



Ordinal color palette: Order Date has a chronological order.

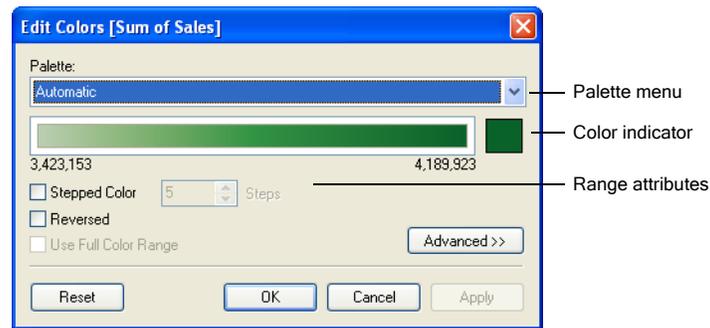


Once you select a palette, click **Assign Palette** to automatically assign the new palette colors to the members in the field. When finished, click **OK** to view the changes and close the dialog.

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

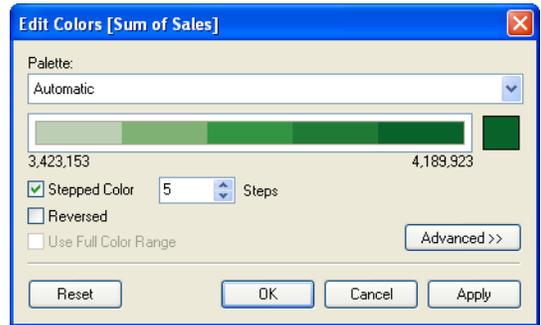
When you add a measure to the **Color** shelf a quantitative legend is added creating a continuous range of colors. You can modify the colors used in the range, the distribution of color, and other range attributes in the Edit Color dialog box. Right-click the legend and select **Edit Colors** or double-click 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**.

Each of the options for formatting quantitative colors are described below:

**Using Stepped Color.** You can 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**.



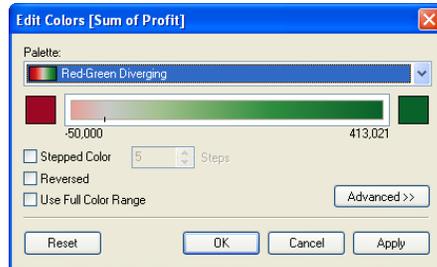
If a diverging color palette is selected, the center point is shown on the color ramp with a small black mark. When the number of steps is odd, the center mark is placed in the middle of the center step. When the number of steps is even, the center mark is placed at the boundary of the center-most two 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**.

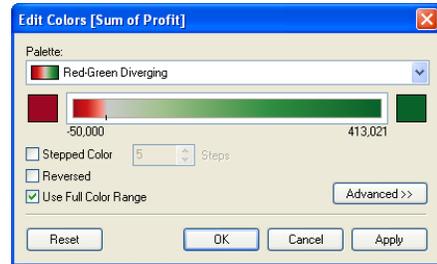
**Using the Full Color Range.** When you are using a diverging color palette you can select to **Use Full Color Range**. When you select this option, Tableau assigns the starting number a full intensity and the ending number a full intensity. If the range is from -10 to 100, the color representing negative 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**.

## Transparency

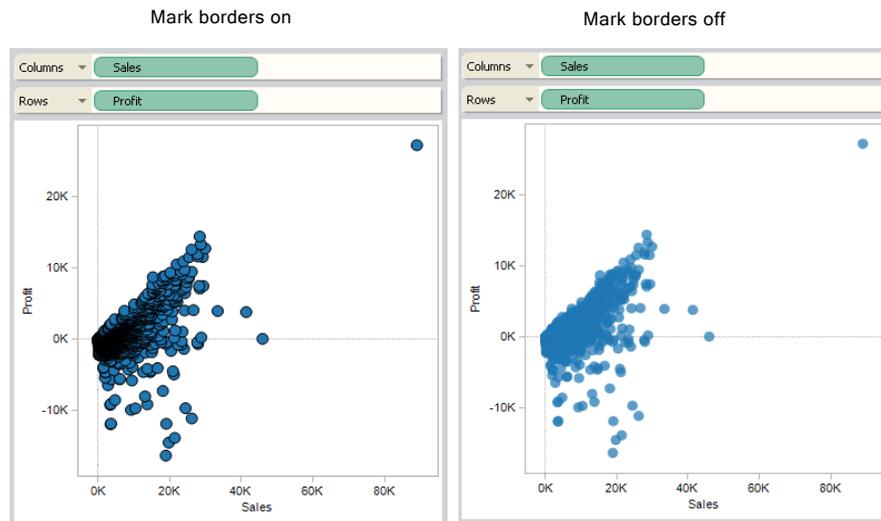
You can also modify the transparency of the marks drop-down control next to the Color shelf. This is especially useful in dense scatter plots or when you are looking at data overlaying a map or background image. As you slide the slider toward the left the marks become more transparent.

## Effects

Use the drop-down control next to the color shelf to modify other color properties. You can

**Mark Borders.** By default, Tableau displays all marks without a border. You can turn on the mark borders for all mark types except text, line, and shape. Turn on mark borders by selecting a color on the color shelf drop-down control.

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.



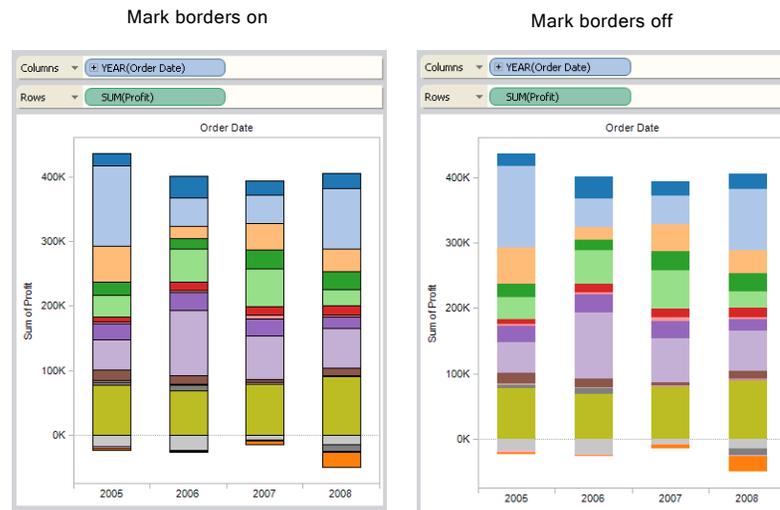
---

**Note** You can also use transparency to show the density of marks.

---

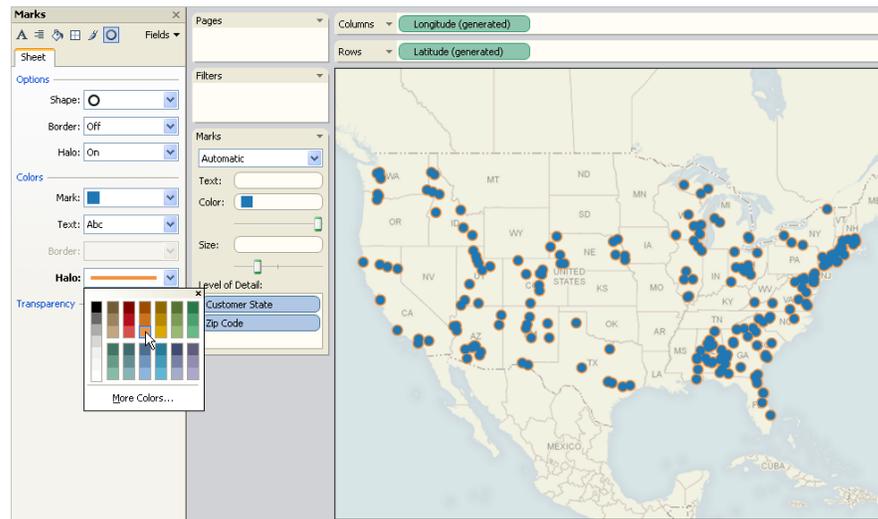
Leaving mark borders off is particularly useful when you are viewing a large number of small marks that are color-encoded. It can be difficult to see the color encoding because the borders dominate the marks.

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 Halos.** In order to make the marks in a view more visible when placed on top of a background image or map, each mark is surrounded by a solid contrasting color called a halo. Mark halos are available when you have a background image or background map. You can turn mark halos by selecting a color on the color shelf drop-down control.

The view below uses a map so the marks are surrounded by orange halos to make them stand out.



**Markers.** When you are using the Line mark type, you can add a marker effect to show and hide the points along the line. You can show selected points, all points, or no points. Select a marker effect on the color shelf drop-down control.

## 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 size encodings. 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.

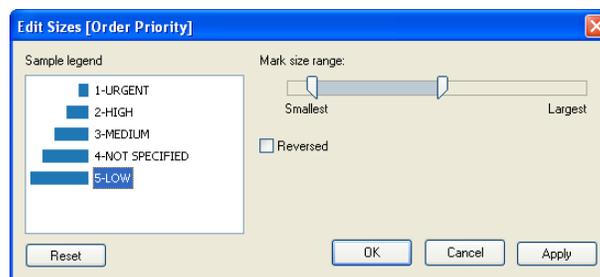
Note that size-encoding data with a discrete field separates the marks in the same way as the **Level of Detail** shelf does, and then provides additional information (a shape) for each mark. When you add categorical size encoding to the view, a legend displays showing the



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.



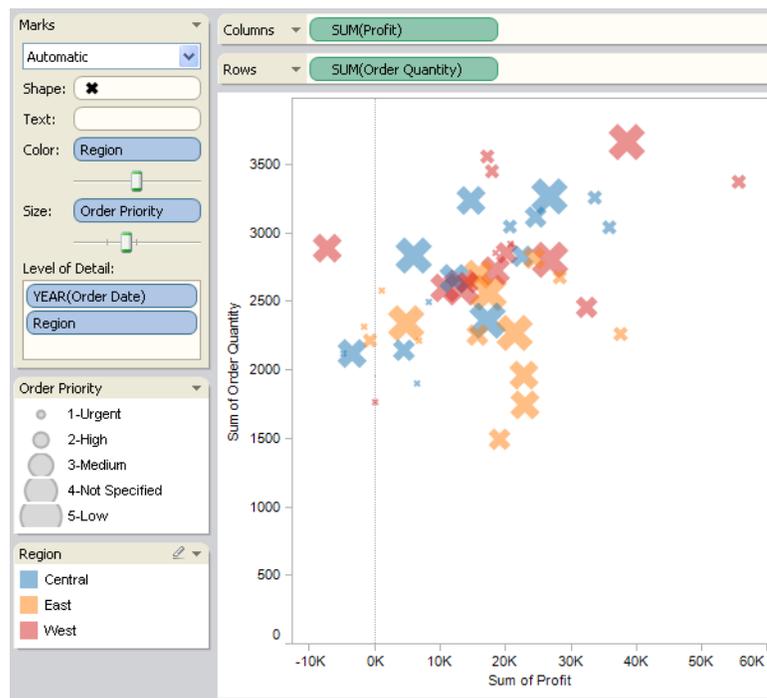
- 2 In the Edit Sizes dialog box, the sizes are displayed on the left and a size range slider is shown on the right. The sizes assigned to each member are distributed across the specified range. Use the slider to adjust the sizes assigned to each member.

You can also select **Reversed** to assign the largest mark to the smallest value and the smallest mark to the largest value.

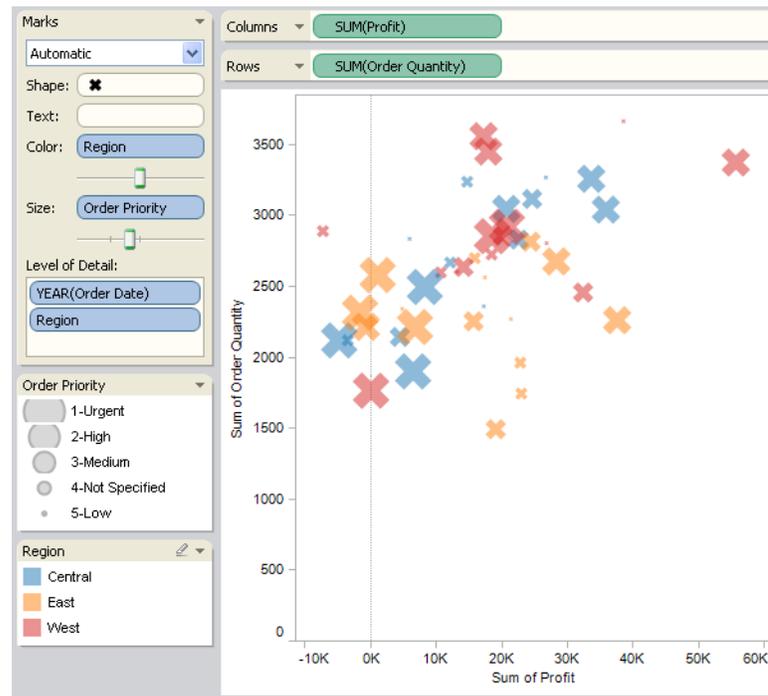
- 3 When finished click **OK**.

**Example – Categorical Sizes**

The view below shows the sales and profit of a superstore broken down by region and order date. The order priority is indicated by the size of the mark.



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

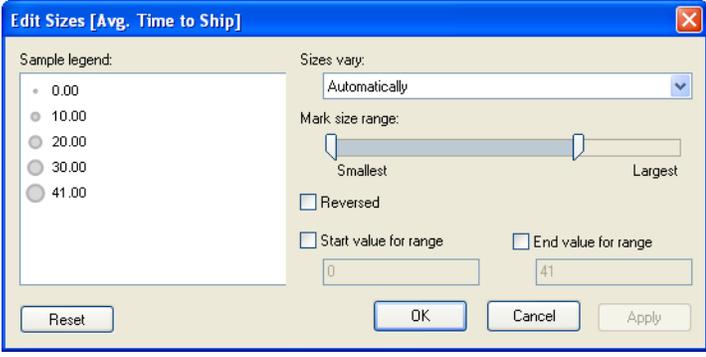
When you place a continuous field on the **Size** shelf, Tableau draws each mark with a different size using a continuous range. The smallest value is assigned the smallest sized mark and similarly the largest value is represented by the largest mark.

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.

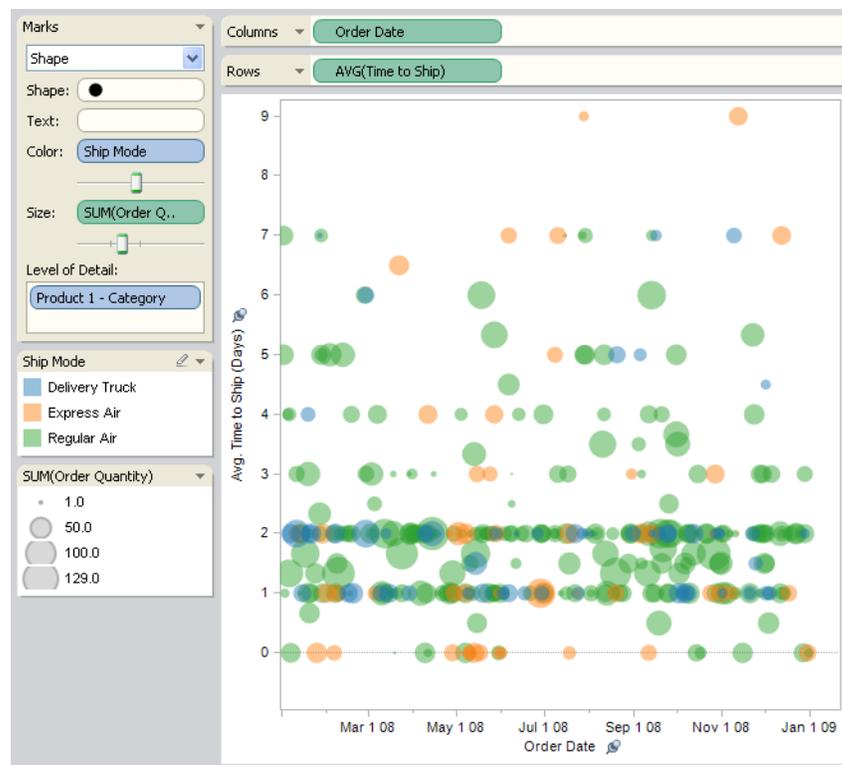


- 2 In the Edit Sizes dialog box, select one of the following ways to map the sizes:
  - **Automatically** - selects the mapping that best fits your data. If the data is numeric and does not cross zero (e.g. all positive or all negative), the 'From zero' mapping is used. Otherwise, the 'By range' mapping 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 analyzes the time it takes to ship products based on their ship mode, order date, and the size of the order. The size of each mark represents the order quantity while the color represents the Ship Mode. Looking at the view you can quickly see that most products ship within 1 and 2 days. However, larger orders shipped via Regular Air tend to take longer, especially during the second quarter. Curiously, there are a couple of smaller orders that were shipped via Express Air that took a long time to ship.



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 allows you to encode data by assigning different shapes to the marks in a data view. This section discusses the following topics:

- About Shapes
- Editing Shapes
- Custom Shapes
- Tips for Creating Custom Shapes

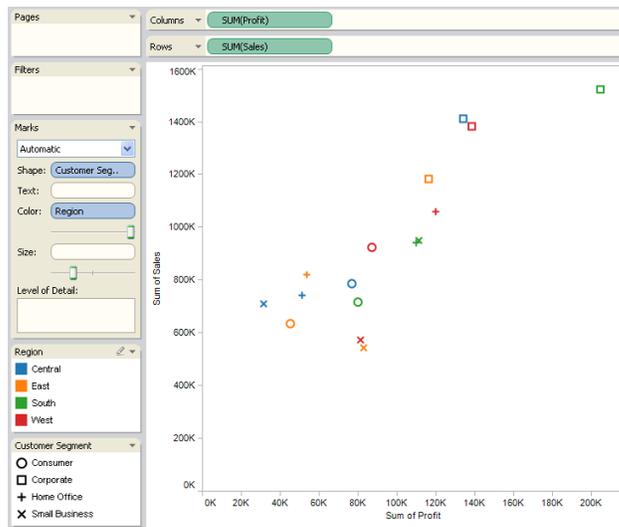


## About Shapes

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. 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 **Customer Segment** dimension. Each shape reflects the customer segment's contribution to the profit and sales.



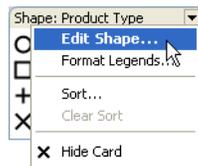


## Editing Shapes

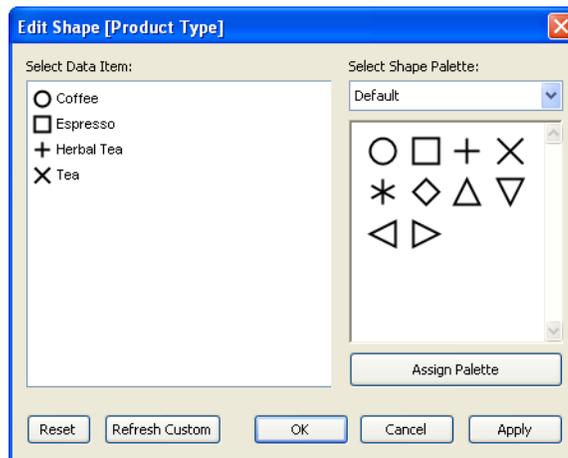
By default, ten unique shapes are used to encode dimensions. If you have more than 10 members, the shapes repeat. In addition to the default palette, you can choose from a variety of shape palettes such as filled shapes, arrows, and even weather symbols.

### To edit shapes:

- 1 Double-click the Shape Legend or select **Edit Shapes** on the legend's card menu. If there is no shape encoding, you can open the Edit Shapes dialog box by clicking the shape shelf itself and then selecting **More Shapes**.



- 2 In the Edit Shape dialog box, select a member on the left and then select the new shape in the palette on the right. You can also click the Assign Palette button to quickly assign the shapes to the members of the field.





---

Select a different shape palette using the drop-down list in the upper right of the Edit Shape dialog box.

---

**Note** Shape encodings are shared across multiple worksheets that use the same data source to help you create consistent displays of your data. For example, if you define Furniture products to be represented by a square, they will automatically be squares in all other views in the workbook. You can set the default shape encodings for a field by right-clicking the field in the Data window and selecting **Edit encodings > Shape**.

---

### Custom Shapes

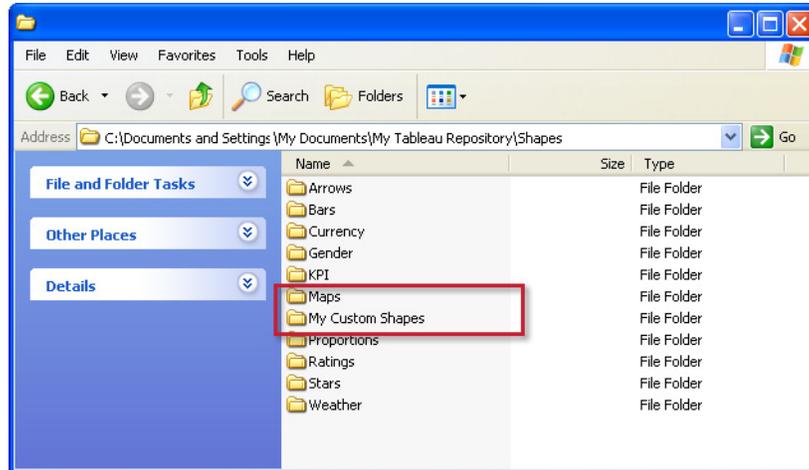
You can add custom shapes by adding the shape image files to the Shapes folder in your Tableau Repository located in your Documents folder. When you use custom shapes, they are saved with the workbook. That way the workbook can be shared with others.

#### To create custom shapes:

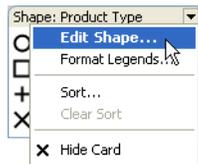
- 1 Create your shape image files. Each shape should be saved as its own file and can be in many image formats including bitmap (.bmp), portable network graphic (.png), JPEG, graphics interchange format (.gif), and so on. Refer to “Tips for Creating Custom Shapes” on page 13-46 for some tips on making useful shapes.
- 2 Place the shapes into the My Tableau Repository folder located in your Documents folder. The shapes should be put into a new folder inside the Shapes folder. The name



of the folder will be used as the name of the palette in Tableau. In the example below, two new palettes are created: Maps and My Custom Shapes.

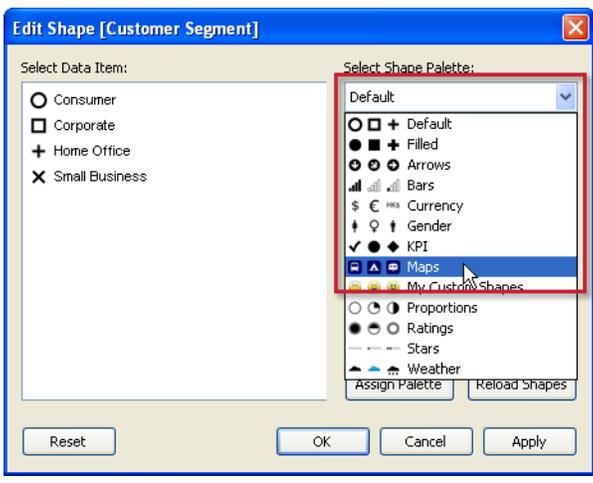


3 In Tableau, open the Edit Shape dialog box.

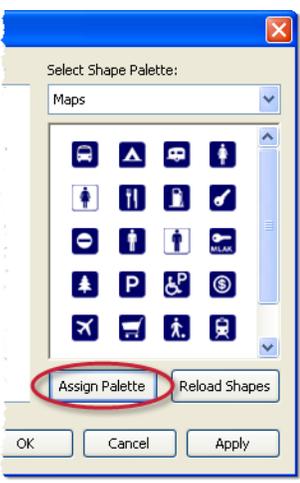




- 4 Choose the new custom palette in the drop-down list in the upper right of the dialog box. If you modified the shapes while Tableau was running, you may need to click the **Reload Shapes** button so the new shapes are available in the dialog box.



- 5 You can either assign members shapes one at a time, or click the Assign Palette button to automatically assign the shapes to the members.





You can return to the default palette at anytime by clicking the **Reset** button. If you open a workbook that uses custom shapes that you don't have, the workbook will show the custom shapes. However, you can click the **Reload Shapes** button in the Edit Shapes dialog box to use the ones in your repository instead.





## Tips for Creating Custom Shapes

When you create custom shapes there are a few things that you can do to improve how your shapes look and function in the view. Below are some tips to help you make good custom shapes. If you are creating your own shapes, we recommend following general guidelines for making icons or clip art.

- **Suggested Size** - unless you plan on using the Size shelf to make the shapes really large, you should try to make your original shape size close to 32 pixels by 32 pixels. However, the original size is dependent on the range of sizes you want available in Tableau. You can resize the shapes in Tableau using the Size shelf as well as the cell size options on the Format menu.
- **Adding Color Encoding** - if you plan to also use the Color shelf to encode the shapes with color, you should use a transparent background. Otherwise, the entire square of the image will be colored rather than just the symbol. GIF and PNG file formats both support transparency. GIF files support transparency for a single color that is 100% transparent, while .png supports alpha channels with a range of transparency levels available on every pixel in the image. When Tableau color encodes the symbol, the amount of transparency for each pixel will not be modified, so you can maintain smooth edges.
- **File Formats** - Tableau does not support symbols that are in the Enhanced Meta File format (.emf). The shape image files can be in one of the following formats: .png, .gif, .jpg, .bmp, and .tiff.



## Label Shelf

The **Label** 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 12-11 for more information.

The most common view using the **Text** shelf is a text table, which is also referred to as cross-tab or a PivotTable.

Product 2 - Sub-Category	Customer Segment			
	Consumer	Corporate	Home Office	Small Business
Appliances	112,337	306,807	187,463	140,726
Binders and Binder Accesso..	167,218	405,003	255,582	196,719
Bookcases	184,348	292,854	194,718	150,062
Chairs & Chairmats	373,903	660,781	403,420	306,040
Computer Peripherals	123,514	320,480	202,771	149,789
Copiers and Fax	233,017	364,176	248,344	284,825
Envelopes	30,692	57,963	35,613	52,030
Labels	6,215	15,152	7,806	9,869
Office Furnishings	113,940	226,389	219,624	105,959
Office Machines	538,623	891,427	409,378	331,914
Paper	85,450	154,710	113,350	95,986
Pens & Art Supplies	37,098	61,203	33,383	35,841
Rubber Bands	3,054	4,759	4,087	3,339
Scissors, Rulers and Trimme..	29,405	23,110	23,643	4,838
Storage & Organization	236,919	336,052	303,651	223,152
Tables	439,609	658,171	463,059	326,375
Telephones and Communica..	348,268	721,989	458,870	360,186



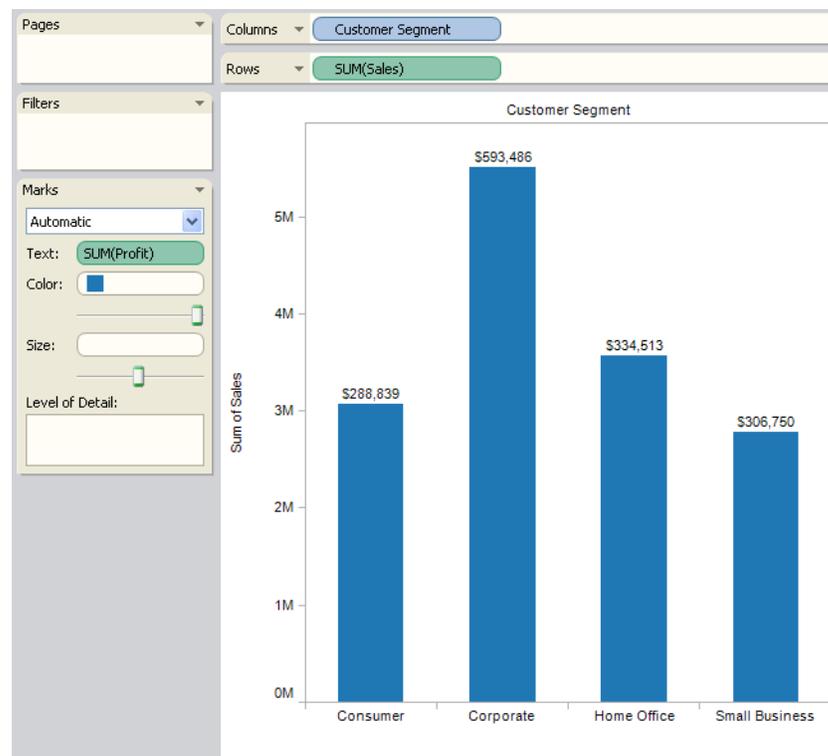
---

**Note** You can display text labels with other mark types by selecting clicking **Show Mark Labels** on the toolbar. Refer to “Mark Labels” on page 28-16 to learn more about showing and hiding mark labels.

---

If you place a dimension on the **Text** shelf, the marks are separated and labeled 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 measure and the labels are given by the sum of the Profit measure.



## Path Shelf

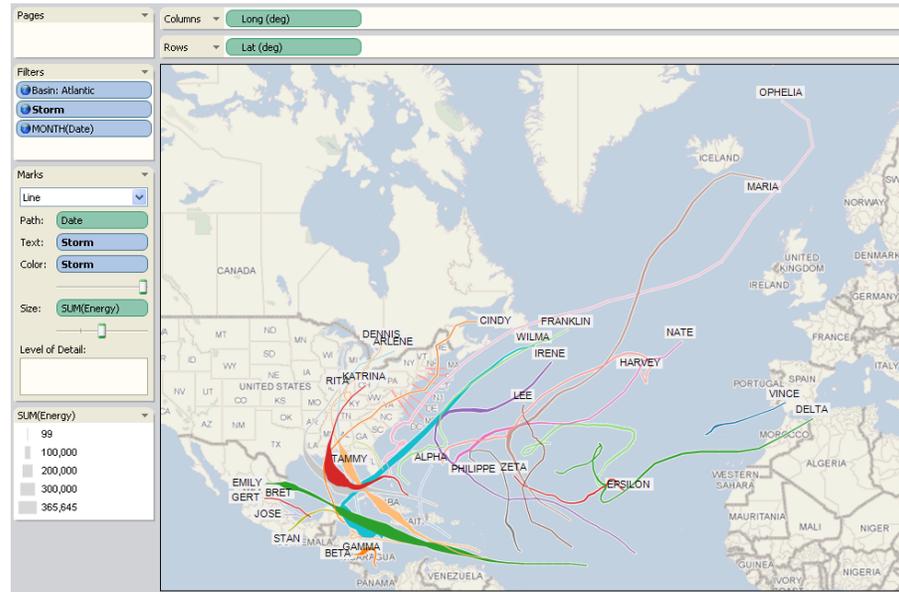
The **Path** shelf allows you to encode data by connecting marks using a particular drawing order. You can path-encode your data using either a dimension or a measure.

- **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 17-3.
- **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 12-11 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.

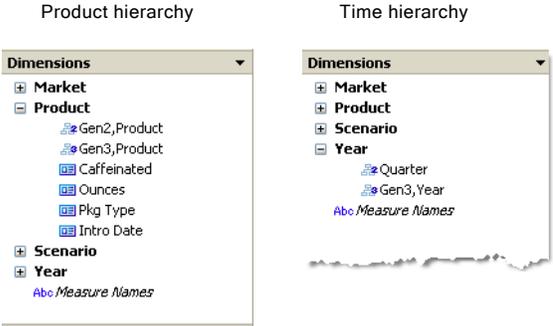


By placing the continuous date on the **Path** shelf, the lines are drawn in chronological order.



# 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 Oracle 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.



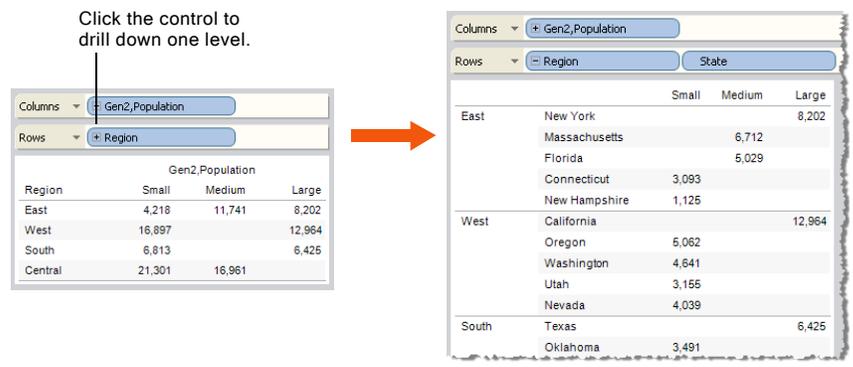
### Using Fields on Shelves

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.



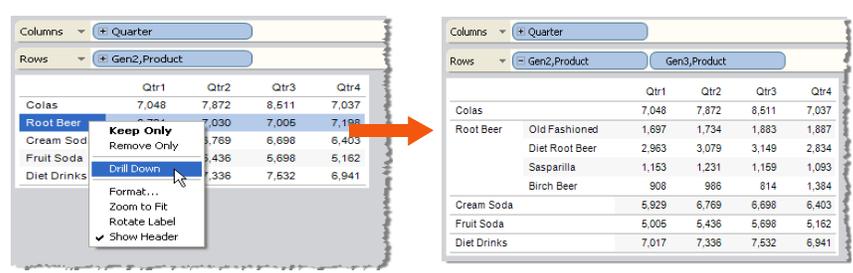
The following figure demonstrates drilling down one level in the hierarchy for the **Region** dimension to expose the **States**.



### 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 **Root Beer** member of the **Gen2,Product** dimension. Note that new row headers are displayed in the table and that **Gen3,Product**, which is the next generation in the hierarchy, is automatically displayed.



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 “Groups” on page 17-18 for more information.

---

## Building Views with Oracle Essbase

When Tableau is connected to an Oracle Essbase data source, there are three important features that you should know about:

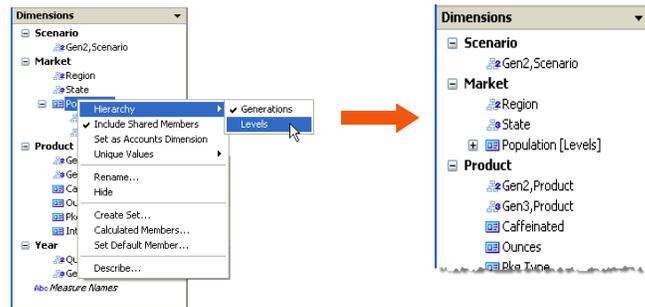
- Generations and Levels – For any given dimension, you can navigate using either levels or generations.
- Shared Members – You can choose to include or exclude dimension members that appear in more than one place in a hierarchy.
- Setting the Default Member – change the default member to avoid repeatedly 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**.



---

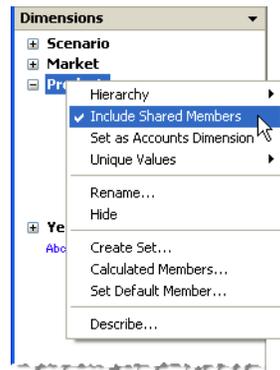
**Note** If you are using the same dimension in multiple worksheets, you can use levels in one worksheet and generations in another worksheet simultaneously. Furthermore, you can mix generations and levels from different dimensions in the same worksheet.

---

### 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.

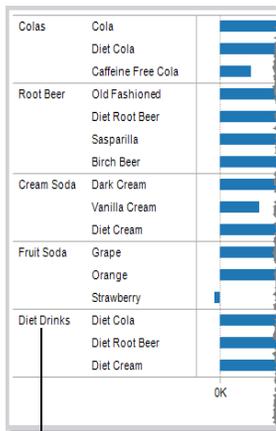
By default, Tableau includes shared members in all generations (or levels) of a dimension. This means that a shared member might appear multiple times in a table. If you choose to exclude shared members, they will appear only once in a table. By default, shared members are included for all dimensions. 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.





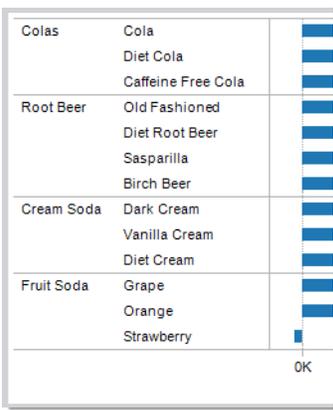
The figure below shows part of a data view where shared members are included (left) and excluded (right). Notice that diet drinks are shared members.

Shared members included.



Shared members.

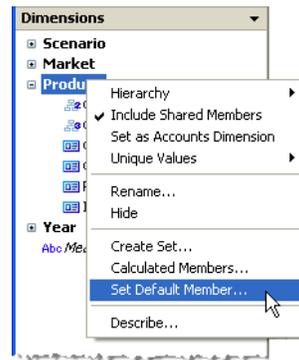
Shared members excluded.



### Setting the Default Member

All multidimensional data sources have default members that are set when the data source is first built. If you find that you are creating filters all the time to 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.

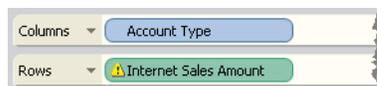
The default member determines how you view the cube and so is much more powerful than applying filters. All fields will be calculated based on the default member you select. In addition, these default member settings are saved with the connection.

## 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.

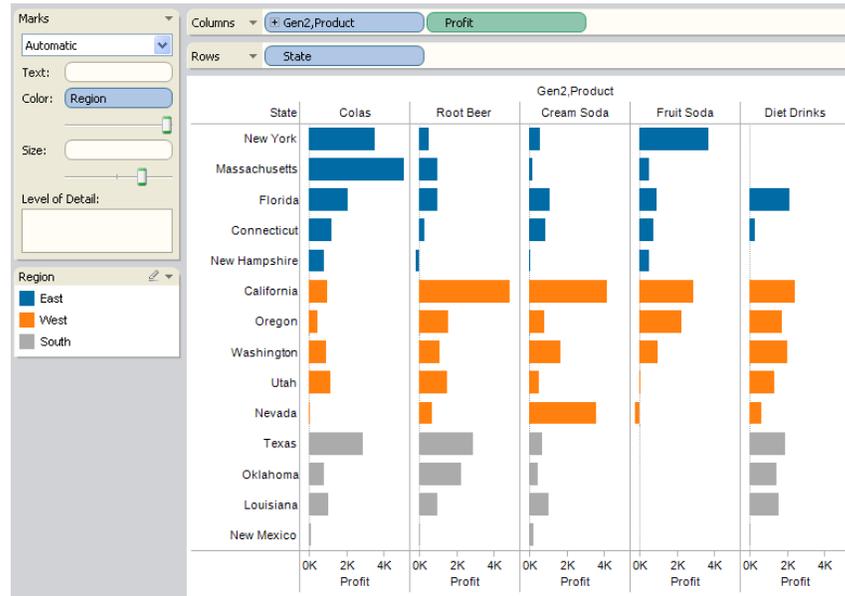


## Perfect Pivoting

In Tableau, *perfect pivoting* refers to working with hierarchies in these ways:

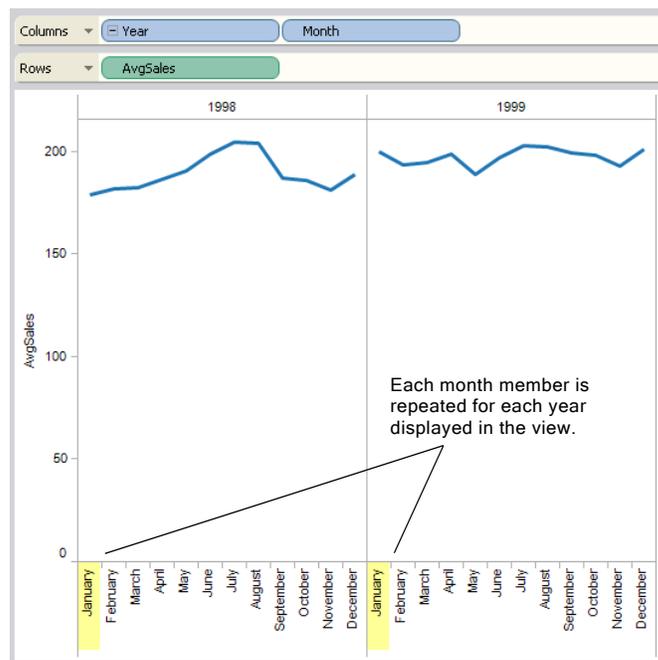
- Using varying levels of detail including skipping levels (for example, Country and City, but not State).
- Using varying levels of detail on different worksheet shelves simultaneously (for example, Product Family on the **Columns** shelf and Product Department on the **Color** shelf).
- Using varying levels of detail out of order (for example, Quarter before Year).

For example, in the following view the **Market** hierarchy is broken up to show the **Sate** level as **Rows** and the **Region** level as **Color**.

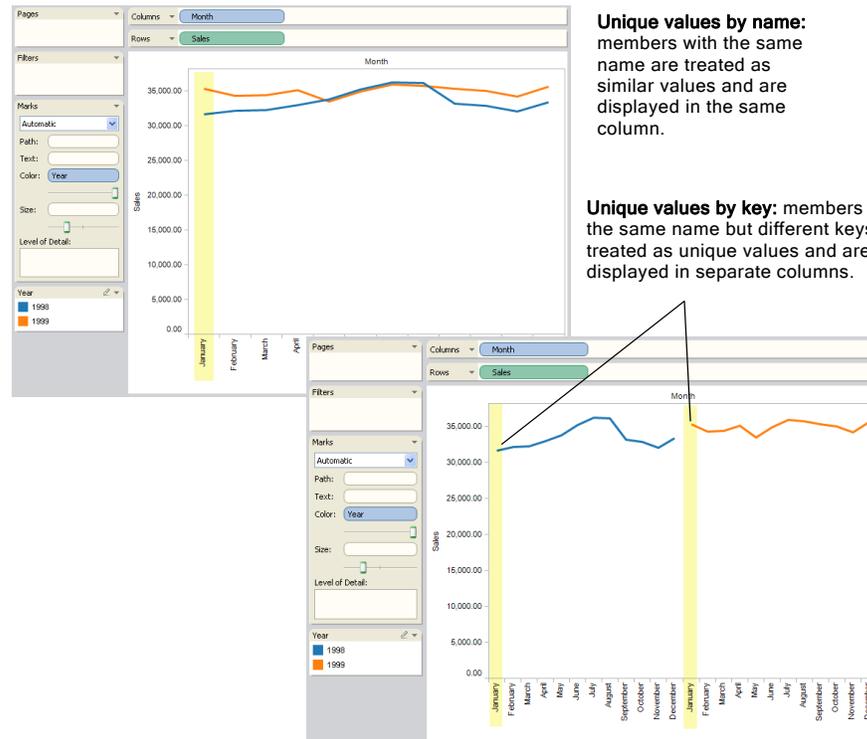


## Defining Unique Values

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.



**Unique values by name:** members with the same name are treated as similar values and are displayed in the same column.

**Unique values by key:** members with the same name but different keys are treated as unique values and are displayed in separate columns.

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:

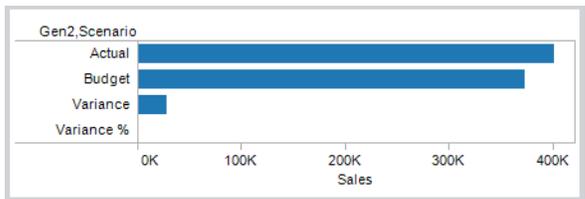
- **By Key:** each member is considered unique based on the key given it by the system administrator when the database is set up. Members with the same name but different keys are treated as unique values.
- **By Name:** each member is considered unique based on the member name. Members with the same name (regardless of their keys) are treated as if they are the same.



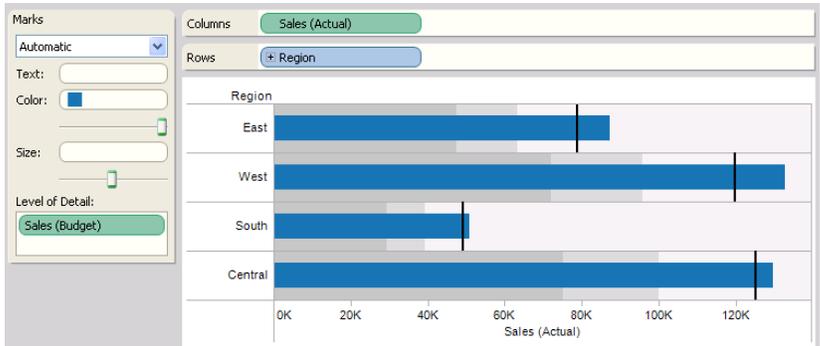
By default, unique date and time values are determined by name and all other values are determined by key.

### Utility Dimensions

Oracle Essbase databases sometimes have special dimensions used to model comparative values such as Actual vs. Budgeted or Current Year vs. Previous Years. These dimensions are the utility dimensions and are often set up as Scenario or Years. For example, the members of a Scenario dimension are shown below.



In the above view, you can see Actual Sales, Budgeted Sales, and so on. However, what if you wanted compare Actual Sales to Budgeted Sales in a bullet graph? In that case you need to set the Scenario dimension to be used as the utility dimension. When you set a dimension as the utility dimension you can then specify which member of the utility dimension to use for each measure in the view. For example, below is a bullet graph showing actual sales to budgeted sales by region.

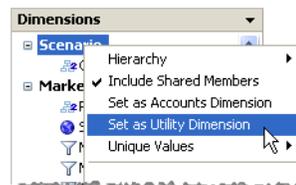


You can see that the Sales measure is used twice in the view: once to show actual and once to show budgeted.



**To use a dimension as the utility dimension:**

- 1 Right-click the dimension in the Data window and select **Set as Utility Dimension**.



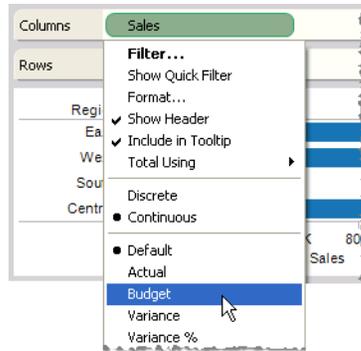
The dimension is hidden in the Data window and can no longer be used as a dimension field in the view. The Measures area of the Data window indicates that there is a utility dimension.



- 2 Drag a measure to the view.

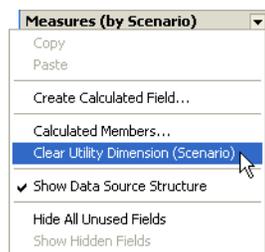


- Open the field menu for the measure in the view and select the member of the utility dimension you want to anchor the measure to.



**To remove a utility dimension:**

- Open the drop-down menu next to the Measures title in the Data window and select **Clear Utility Dimension**.

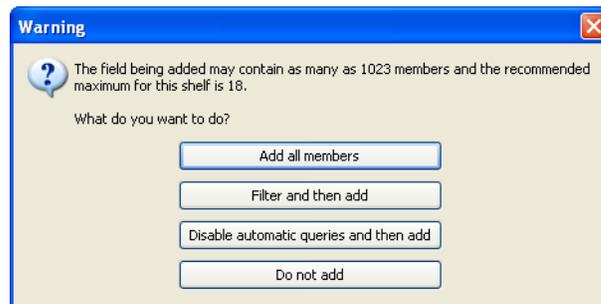


When you remove the utility dimension, measures that reference it in the view are no longer valid.

---

## Working with Large Views

Placing dimensions with a large number of dimensions on a shelf may take a long time and generally won't be very useful when they are added. Tableau will present you with the following dialog box with the options to make it more manageable.



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 updates 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 “Groups” on page 17-18).
- 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 10-4.



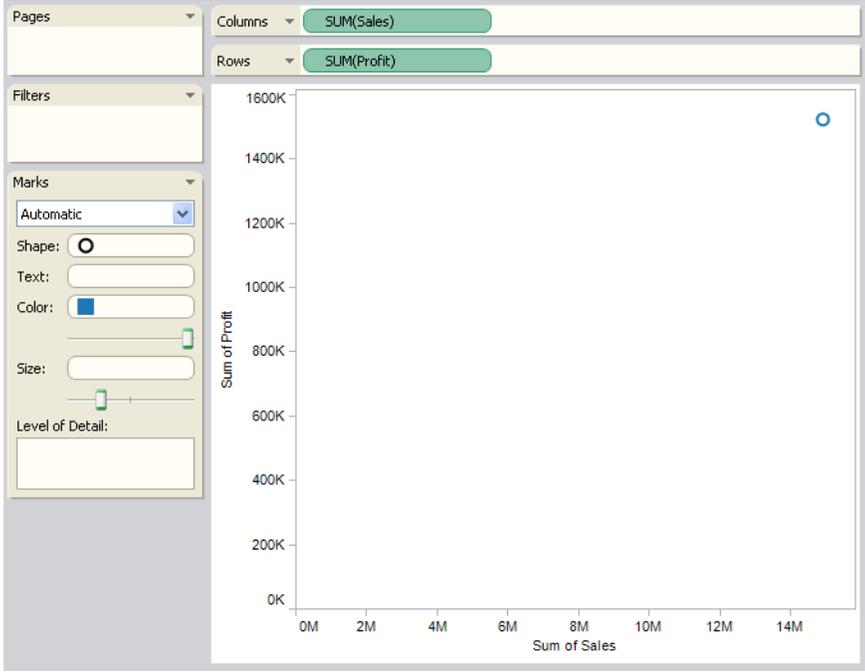
## Example – Building Data Views Manually

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** measure on the **Columns** shelf and the **Profit** measure on the **Rows** shelf.

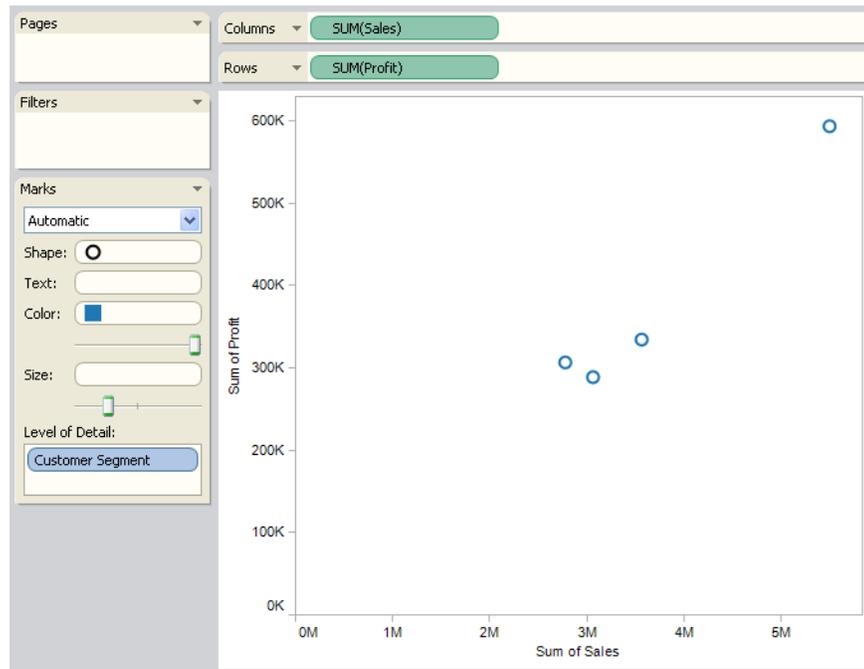
The measures are automatically aggregated and result in one data point. The data point is displayed using the shape mark type. Note that you are not displaying any levels of detail because dimension members are not included in the view.





- 2 Place the **Customer Segment** dimension on the **Level of Detail** shelf.

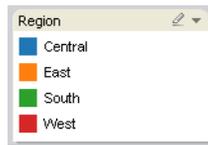
The original mark is now separated into four marks, where each new mark is associated with a member (level of detail) of the **Customer Segment** dimension.





3 Place the **Region** dimension on the **Color** shelf.

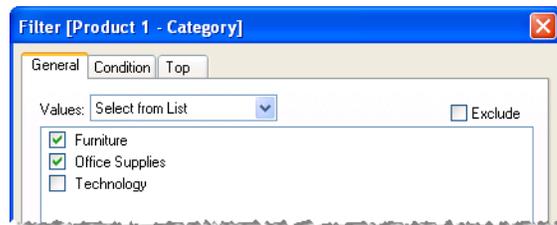
Each of the four marks are now separated into four new marks for a total of sixteen. Each new mark is associated with a member of the **Region** dimension, and is encoded with a unique color. The color legend displays each member name and its associated color.



4 Place the **Product 1 -Category** dimension on the **Shape** shelf and filter the dimension to exclude Technology products.

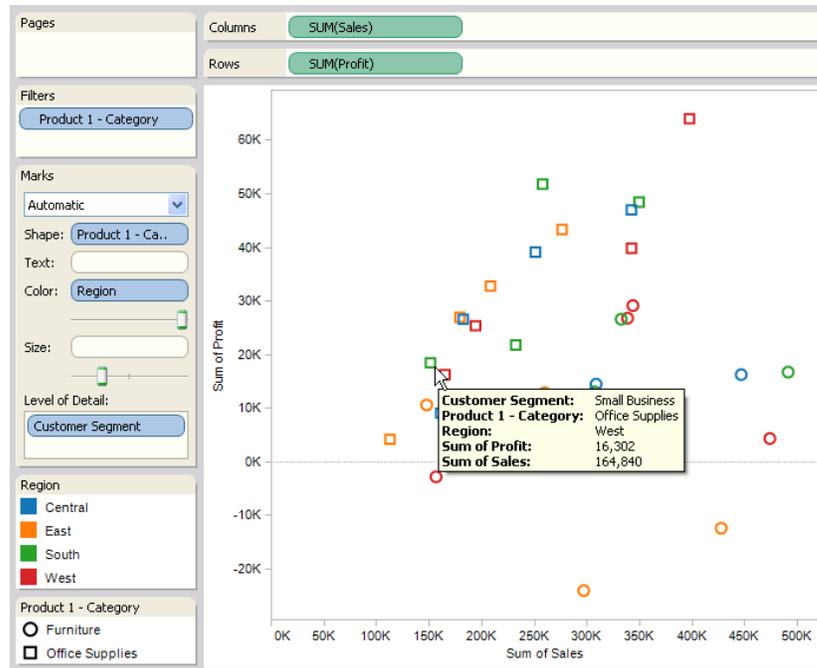
Each of the twelve marks are now separated into three new marks for a total of 48. Each new mark is associated with a member of the **Product 1 - Category** dimension, and is encoded with a unique shape. The shape legend displays each member name and its associated shape.

To filter the data, select **Filter** on the **Product 1 - Category** context menu. The **Filter** dialog box opens. Deselect **Technology** to exclude it from the view.





The final view is shown below



# Building Views Automatically

---

Overview . . . . .	14-2
Show Me! . . . . .	14-3
Add to Sheet: Double-Click . . . . .	14-6

---

## 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:

- Show Me!
- Add to Sheet: Double-Click

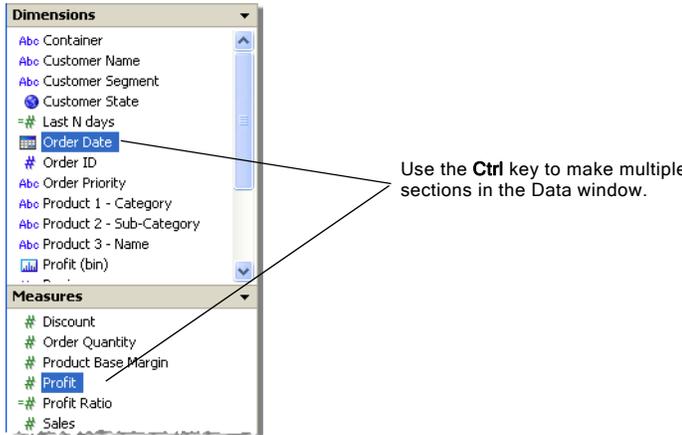


# Show Me!

Show Me! creates a view based on one or more selected fields. When you use Show Me! simply select fields you want to analyze in the Data window and press the Show Me! button on the toolbar. Tableau automatically evaluates the selected fields and gives you the option of several types of views that would be appropriate for those fields.

### 1 Select Input.

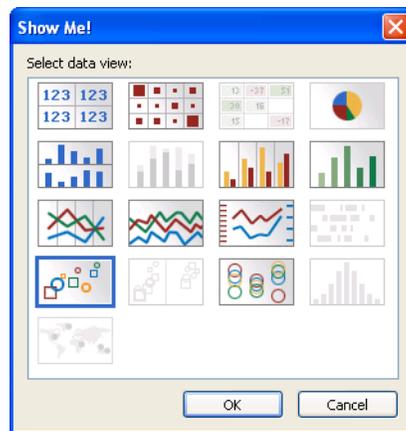
Select fields in the Data window that you want to analyze. Use the **Ctrl** key to make multiple selections.



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



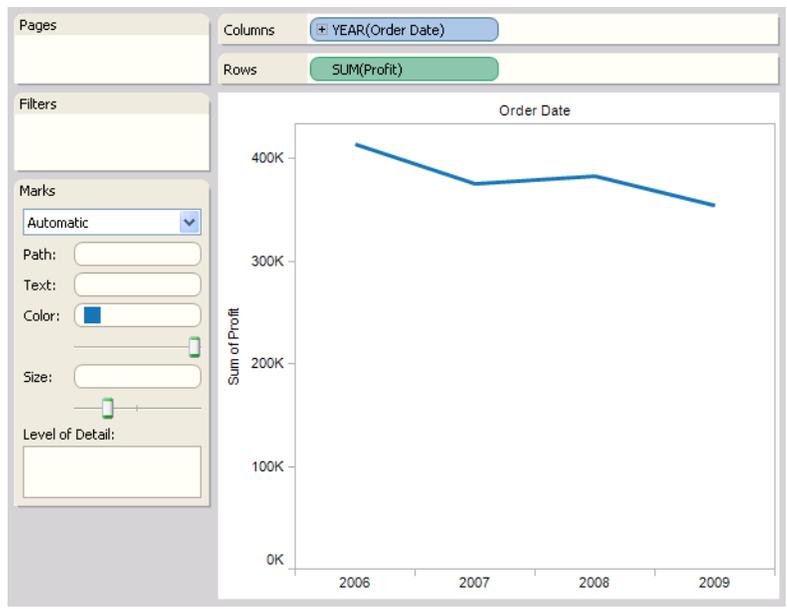
- 3 Select the type of view you want to create in the dialog box.



Any alternative that is not grey will generate a view of your data. Tooltips describe the minimum requirements for each alternative.



4 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 suggestion based on best practices for presenting data. 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.

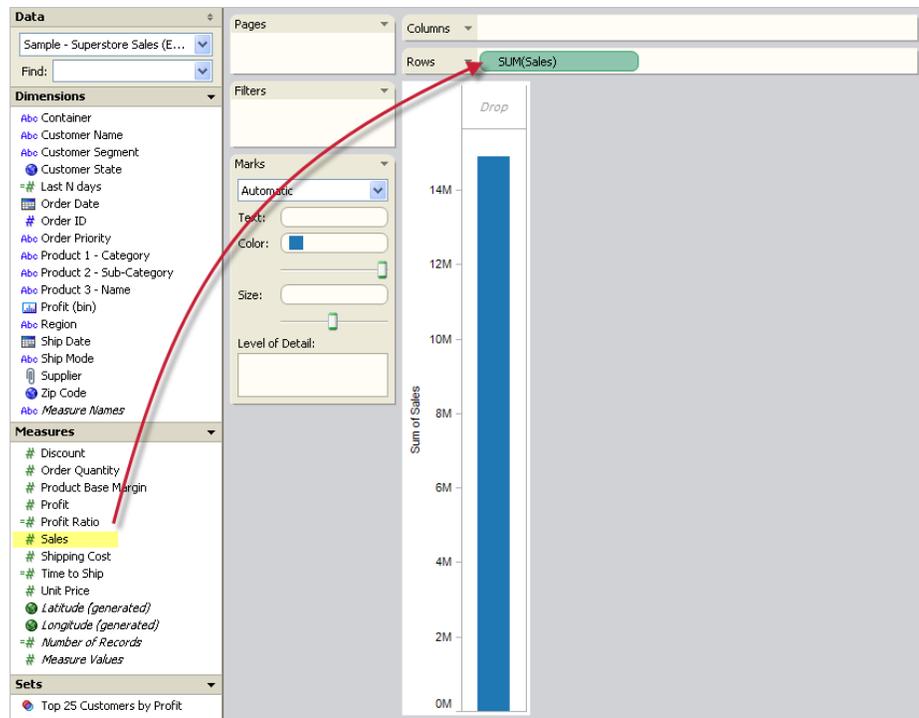
---

## Add to Sheet: Double-Click

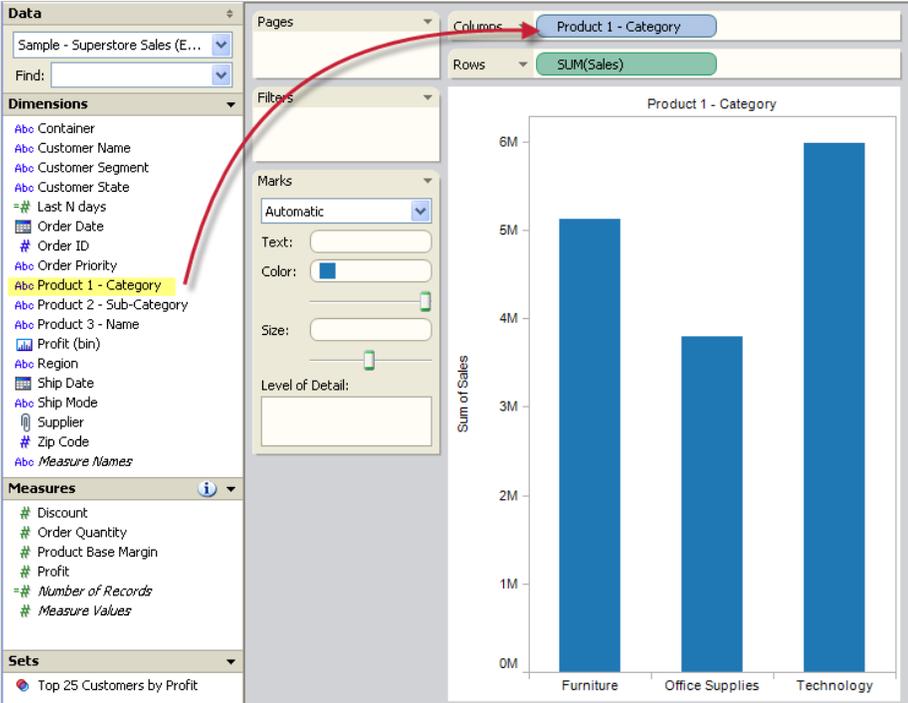
Tableau supports an additional method for automatically generating views of data called Automatic Double-Click. To use this method, double-click fields in the Data window you are interested in. Tableau automatically adds each field to the view. That is, each double-click results in an additional field added to a shelf in an intelligent way. Like Show Me!, this function leverages Tableau’s ability to make an intelligent “best guess” of how the data should be displayed.

**Here’s how it works:**

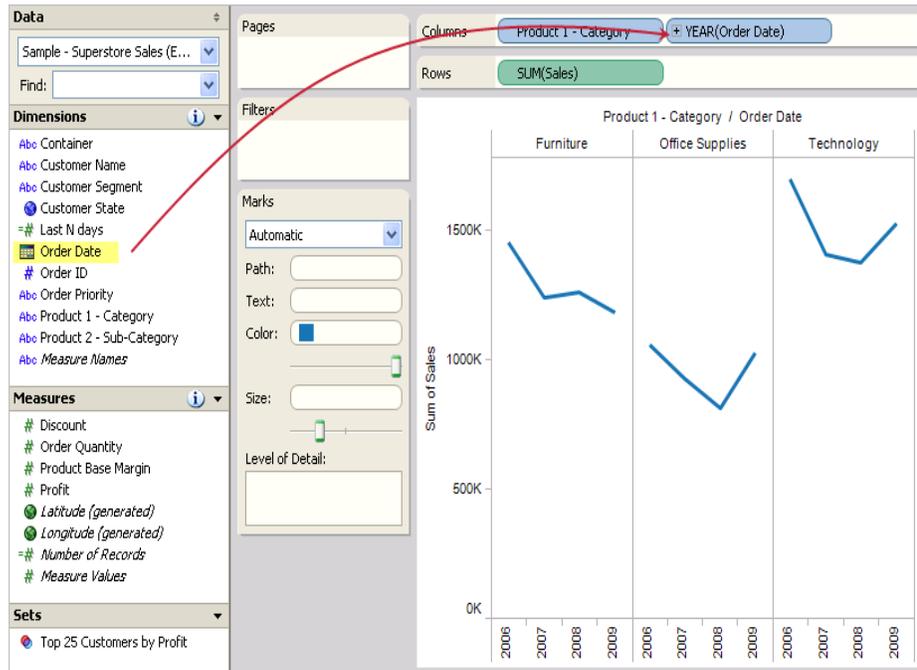
- 1 Double-clicking the **Sales** measure in the Data window automatically adds that field to the view in an intelligent way.



- 2 Double-clicking the **Product 1 - Category** 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 way. 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	Adding a dimension first produces a text table (or cross-tab). All subsequent clicks on fields result in refinement of the text table.
Bars	Adding 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 added, at which time the view is changed to a line.
Line	Adding a measure and then a date dimension produces a line view. All subsequent clicks result in refinement of the line view.
Continuous Line	Adding a continuous dimension and then a measure produces a continuous line view. Subsequent dimensions result in refinement of the continuous line view. Subsequent measures add quantitative axes to the view.
Scatter	Adding 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.
Maps	Adding a geographic field produces a map view with latitude and longitude as axes and the geographic field on the Level of Detail shelf. Subsequent dimensions add rows to the view while subsequent measures further refine the map by adding size and color encoding.



## Using Multiple Measures

---

<b>Overview</b> . . . . .	<b>15-2</b>
<b>Individual Axes</b> . . . . .	<b>15-3</b>
<b>Blended Axes</b> . . . . .	<b>15-4</b>
<b>Dual Axes</b> . . . . .	<b>15-6</b>
<b>Combination Charts</b> . . . . .	<b>15-8</b>

---

## Overview

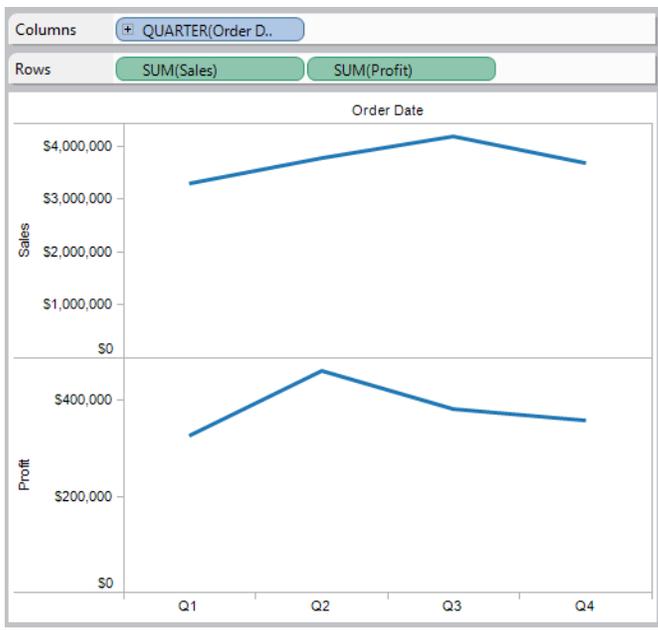
There are lots of different ways to compare multiple measures in a single view. For example, you can create individual axes for each measure or you can blend the two measures to share an axis and finally, you can add dual axes where there are two independent axes layered in the same pane. In any of these cases you can customize the marks for each axis to use multiple mark types and add different levels of detail. Views that have customized marks are called combination charts. This section discusses the following topics:

- Individual Axes
- Blended Axes
- Dual Axes
- Combination Charts



## Individual Axes

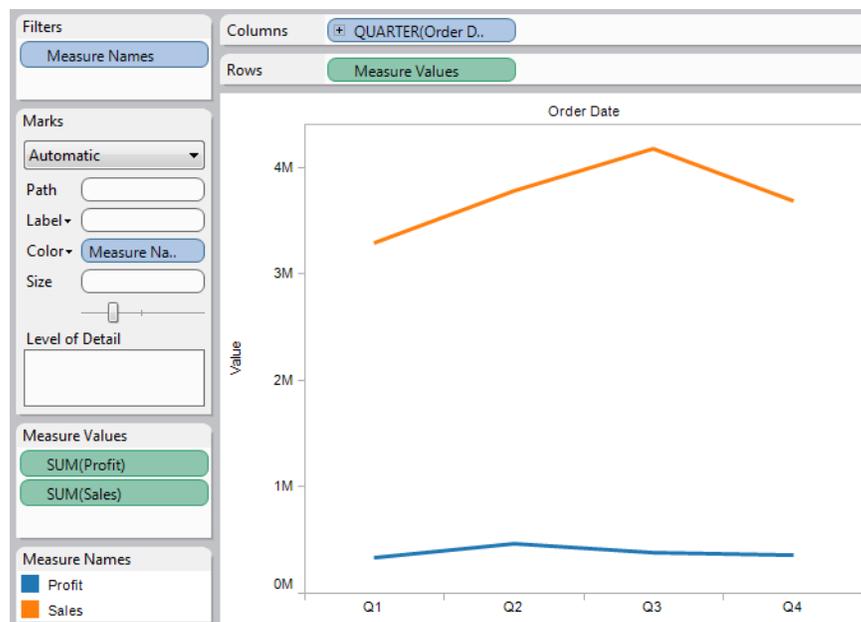
.Add individual axes for each measure by dragging measures to the Rows and Columns shelves. Each measure on the Rows shelf adds an additional axis to the rows of the table. Each measure on the Columns shelf adds an additional axis to the columns of the table. For example, the view below shows quarterly sales and profit. The Sales and Profit axes are individual rows in the table and have independent scales.



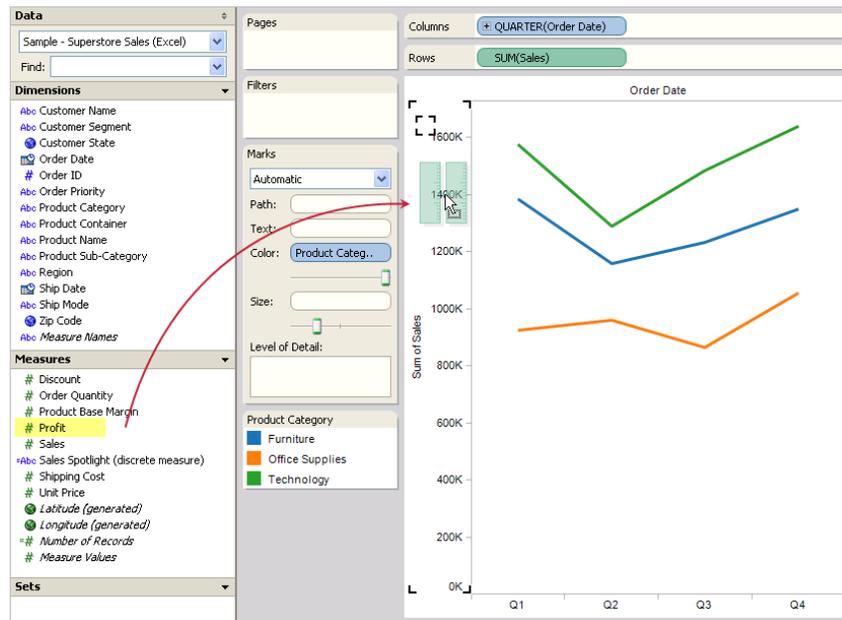
---

## Blended Axes

Measures can share a single axis so that all the marks are shown in a single pane. Instead of adding rows and columns to the view, when you blend measures there is a single row or column and all of the values for each measure is shown along one continuous axis. For example, the view below shows quarterly sales and profit on a shared axis.



To blend multiple measures, simply drag one measure or axis and drop it onto an existing axis.



Blending measures uses the Measure Names and Measure Values fields, which are generated fields that contain all of the measure names in your data source and all of the measure values. The shared axis is created using the Measure Values field. The Measure Names field is added to the Color shelf so that a line is drawn for each measure. Finally, the Measure Names field is filtered to only include the measures you want to blend.

For a detailed example of blending measures refer to “Line Chart–Filter and Color-Encode Multiple Measures” on page 33-35.

---

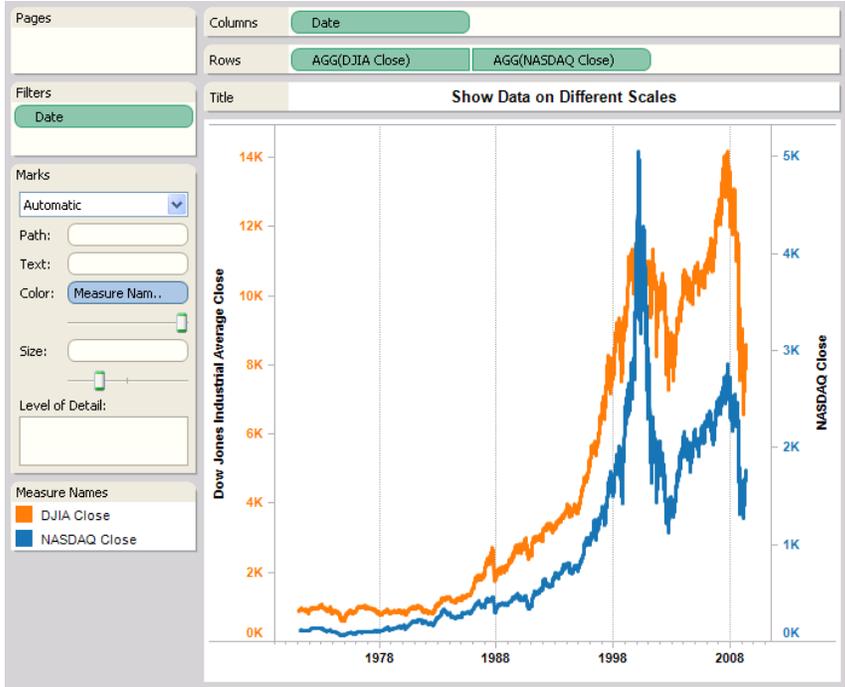
**Note** Blending axes is most appropriate when comparing measures that have a similar scale and units. If the scales of the two measures are drastically different, the trends may be distorted.

---

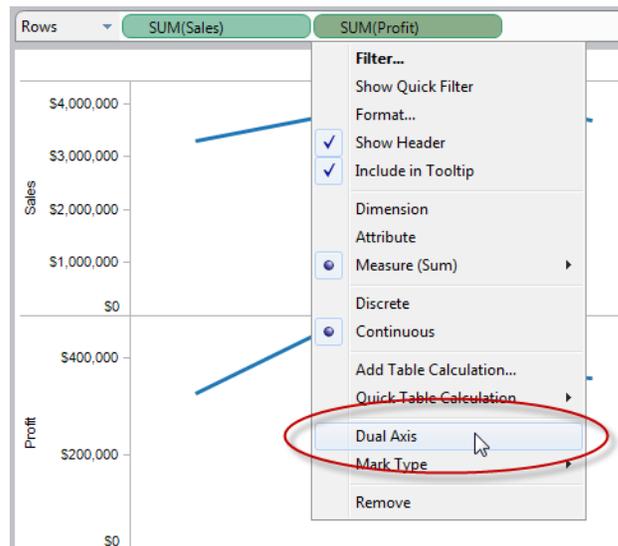


# Dual Axes

You can compare multiple measures using dual axes, which are two independent axes that are layered on top of each other. Dual axes are useful when you have two measures that have different scales. For example, the view below shows Dow Jones and NASDAQ close values over time. The two axes are independent scales but the marks are layered in the same pane.



To add the measure as dual axis drag the field to the right side of the view and drop it when you see a black dashed line. You can also select Dual Axis on the field menu for the measure.



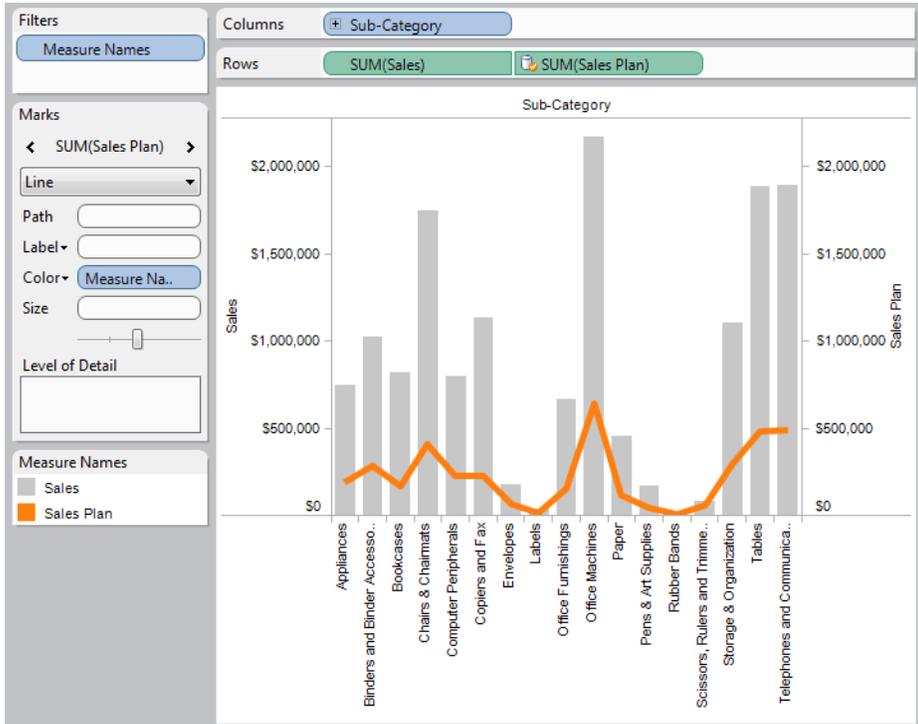
You can add up to four layered axes: two on the Columns shelf and two on the Rows shelf.

**Note** When you are using dual axes you should make sure that the two axes scales align with each other so you can make a correct comparison. You can easily line the two axes up by right-clicking the secondary axis and selecting **Synchronize Axis**.

# Combination Charts

When working with multiple measures in a view, you can customize the mark type for each distinct measure. For example, you can create a view with a line showing a target amount across several months and a bar chart showing the actual attainment for the months. These measures can be displayed as individual axes, blended axes, or dual axes.

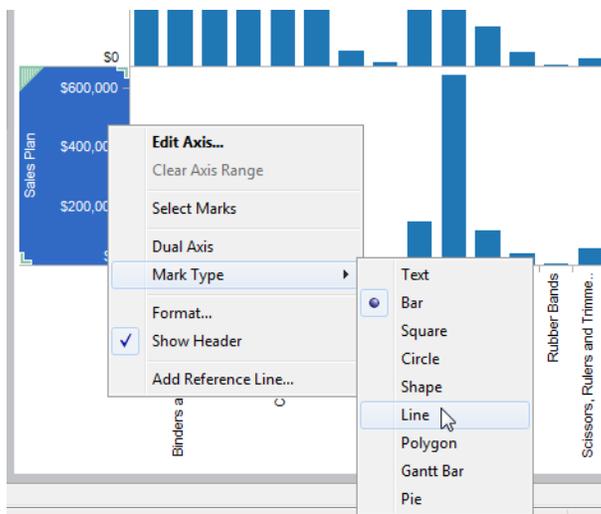
Because each measure can have customized marks, you can customize the level of detail, size, shape, and color encoding for each measure too.



To customize the marks for a measure:



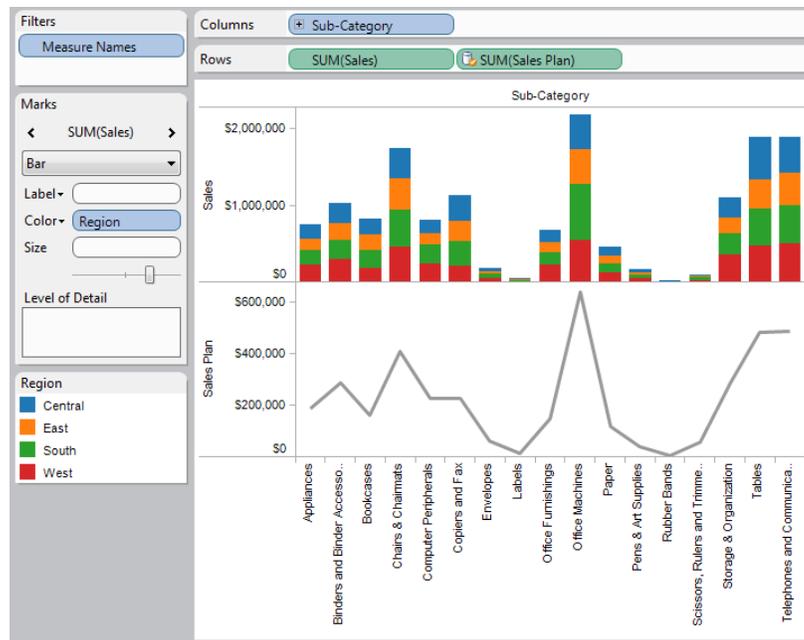
- 1 Right-click the axis for the measure you want to customize and select **Mark Type** and then select a custom mark type.



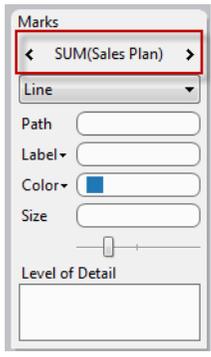
- 2 The Marks card switches into an advanced mode that shows the mark properties for the measure you customized. Any changes to the mark type, shape, size, color, and level of detail will be applied to the selected measure. For example, in the view below the marks card is showing the properties for the SUM(Sales) measure. When Region is placed on



the Color shelf, the encoding and level of detail is only applied to the SUM(Sales) layer. The SUM(Sales Plan) is now broken down by Region.

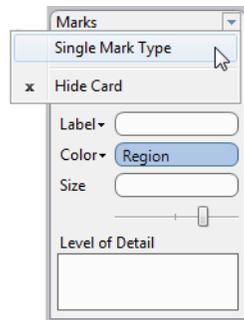


When the Marks card is in the Advanced mode you can switch between each of the measures in the view using control at the top of the Marks card. Select ALL to modify properties for all measures at once.





Select **Single Mark Type** on the Marks card menu to make all measures use the properties that are currently showing in the marks card.



# Filtering

---

<b>Overview</b> . . . . .	<b>16-2</b>
<b>Adding Filters</b> . . . . .	<b>16-3</b>
Selecting Data to Filter . . . . .	16-3
Dragging Fields to the Filters Shelf . . . . .	16-7
Using Quick Filters . . . . .	16-38
<b>Global Filters</b> . . . . .	<b>16-49</b>
<b>Context Filters</b> . . . . .	<b>16-50</b>
Creating a Context Filter . . . . .	16-51
Example – Context Filters . . . . .	16-51
<b>Calculation Filters</b> . . . . .	<b>16-57</b>

---

## Overview

Narrow the data shown in a view using Filter. Filters are defined by selecting specific dimension members or a range of measure values. For example a view showing product sales in four different regions may be filtered to only show three regions. In this case, the filter is created by selecting the specific regions to show. Another example may be to filter the same view to only show sales between \$5000 and \$20,000. For this filter a range of values is defined. In addition to these basic filters you can create complex computed filters to show the top 10 products based on sales, all products sold in the last 30 days, and so on.

All fields that are filtered display on the Filters shelf so you can quickly determine the data that has been removed from the view. Filters are applied to the view in the order they appear on the Filters shelf. However, by default filters are evaluated independently from each other, which means each filter is evaluated against the entire underlying data source and other filters are not taken into account. When working with independent filters, the order on the filters shelf does not change the results.

This topic discusses how to add basic filters and then refine them in the following topics:

- Adding Filters
- Global Filters
- Context Filters
- Calculation Filters



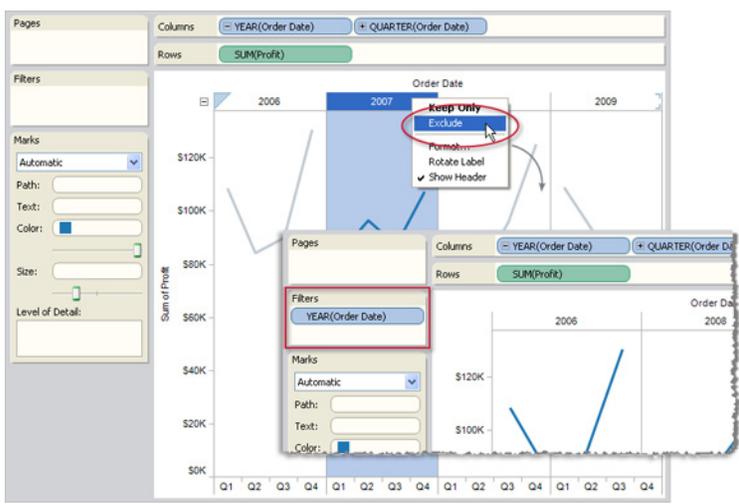
## Adding Filters

All fields that are filtered show on the Filters shelf. You can add a filter either by selecting data in the view, dragging a field to the Filters shelf, or turning on quick filters. Learn about each of these ways to add filters in the following topics:

- Selecting Data to Filter
- Dragging Fields to the Filters Shelf
- Using Quick Filters

### Selecting Data to Filter

You can filter data by selecting headers or marks in the view and then selecting **Keep Only** or **Exclude** on the right-click context menu. The dimension members are removed from the view and the filtered fields are added to the Filters shelf.



### Selecting Headers

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 East member of the **Region** dimension, and the Texas and Louisiana members of the **State** dimension. Note that when East is selected, all members from the next (inner) level in the hierarchy are automatically selected.

The selected column headers include the Colas and Root Beer members of the **Gen2, Product** dimension. Note that when these outer dimensions are selected, the inner dimension members from **Pkg Type** are not automatically selected. This is because the Gen2, Product and PkgType dimensions are unrelated.

Select headers from different hierarchy levels.

Select headers from the same outer dimension.

Columns		Gen2, Product		Pkg Type					
Rows		Region		State					
Region	State	Colas		Root Beer	Cream Soda		Fruit Soda	Diet Drinks	
		Bottle	Can	Bottle	Bottle	Can	Bottle	Bottle	Can
East	New York		\$8,940	\$7,939	\$9,305		\$8,514		
	Massachuse..		\$6,518	\$5,180	\$1,418		\$1,541		
	Florida		\$5,867	\$5,283	\$3,630	\$1,074	\$2,487		\$4,142
	Connecticut		\$3,378	\$3,090	\$3,849		\$2,094		\$1,384
	New Hamps..		\$1,467	\$1,570	\$2,180	\$965		\$1,109	
West	California	\$1,324	\$10,772	\$16,794	\$8,073	\$3,055	\$7,424	\$7,151	\$4,887
	Oregon	\$1,433	\$2,817	\$6,743	\$1,581	\$875	\$6,543	\$3,810	\$2,166
	Washington		\$4,937	\$4,704	\$1,243	\$3,378	\$4,774	\$2,068	\$6,268
	Utah	\$1,764	\$3,469	\$4,237	\$3,101	\$1,467	\$3,267	\$1,315	\$3,552
	Nevada		\$1,790	\$1,722	\$8,850	\$3,768	\$13,026	\$613	\$4,593
South	Texas		\$8,073	\$7,424	\$1,291	\$1,615		\$2,276	\$3,378
	Oklahoma		\$2,446	\$6,543	\$2,890	\$1,562		\$2,799	\$2,427
	Louisiana		\$2,660	\$4,774	\$2,085	\$1,797		\$1,371	\$3,214
	New Mexicc		\$3,101	\$3,267	\$825	\$593		\$1,087	\$2,124
Central	Illinois		\$8,850	\$13,026	\$4,918	\$3,183	\$4,241	\$5,875	\$6,834
	Ohio		\$1,515	\$2,724	\$5,228	\$865	\$676	\$5,867	\$1,418
	Wisconsin		\$1,562	\$5,283	\$3,142	\$1,417	\$1,697	\$3,077	\$1,074
	Missouri			\$3,849	\$2,440	\$1,531	\$1,368	\$2,864	\$1,125
	Iowa		\$593	\$1,331	\$1,129	\$3,651	\$7,151	\$12,943	\$516
	Colorado		\$3,183	\$4,918	\$4,241	\$2,948	\$3,810	\$4,459	\$1,965
Grand Total		\$12,841	\$93,293	\$109,086	\$64,436	\$36,969	\$84,230	\$38,240	\$67,438

Selecting **Keep Only** keeps all selected headers as shown below. The Product field is filtered to show Colas and Root Beer and the Market field is filtered to show the Eastern region as well as Texas and Louisiana in the Southern region.



Columns		Pkg Type		
Rows		State		
Region	State	Colas		Root Beer
		Bottle	Can	Bottle
<b>East</b>	New York		\$8,940	\$7,939
	Massachusetts		\$6,518	\$5,180
	Florida		\$5,867	\$5,283
	Connecticut		\$3,378	\$3,090
	New Hampshire	\$1,467	\$1,570	\$2,180
<b>South</b>	Texas		\$8,073	\$7,424
	Louisiana		\$2,660	\$4,774
<b>Grand Total</b>		<b>\$1,467</b>	<b>\$37,006</b>	<b>\$35,870</b>

Selecting **Exclude** excludes all selected headers as shown below. The Product field is filtered to show Cream Soda, Fruit Soda, and Diet Drinks. The Market field is filtered to show the Western and Central regions along with the remaining states in the Southern region.

Columns		Pkg Type				
Rows		State				
Region	State	Cream Soda		Fruit Soda	Diet Drinks	
		Bottle	Can	Bottle	Bottle	Can
<b>West</b>	California	\$8,073	\$3,055	\$7,424	\$7,151	\$4,887
	Oregon	\$1,581	\$875	\$6,543	\$3,810	\$2,166
	Washington	\$1,243	\$3,378	\$4,774	\$2,068	\$6,268
	Utah	\$3,101	\$1,467	\$3,267	\$1,315	\$3,552
	Nevada	\$8,850	\$3,768	\$13,026	\$613	\$4,593
<b>South</b>	Oklahoma	\$2,890	\$1,562		\$2,799	\$2,427
	New Mexico	\$825	\$593		\$1,087	\$2,124
<b>Central</b>	Illinois	\$4,918	\$3,183	\$4,241	\$5,875	\$6,834
	Ohio	\$865	\$676	\$5,867	\$1,418	\$2,109
	Wisconsin	\$1,417	\$1,697	\$3,077	\$1,074	\$4,090
	Missouri	\$1,531	\$1,368	\$2,864	\$1,125	\$3,132
	Iowa	\$3,651	\$7,151	\$12,943	\$516	\$7,657
Colorado	\$2,948	\$3,810	\$4,459	\$1,965	\$6,865	
<b>Grand Total</b>		<b>\$41,893</b>	<b>\$32,583</b>	<b>\$68,485</b>	<b>\$30,816</b>	<b>\$56,704</b>



---

**Note** These options are not available if a Wildcard Match filter is already specified for the field. Refer to “Filtering Dimensions” on page 16-7 to learn more about Wildcard Match filters.

---

### Selecting Marks

Instead of selecting headers to filter, you can filter individual marks in a view. This method is useful when you are looking at a scatter plot and you want to focus on a set of outliers or remove them so you can better focus on the rest of the data. Select individual marks or click and drag to select several marks. Then right-click and select **Keep Only** or **Exclude**.





## Dragging Fields to the Filters Shelf

Another way to create a filter is to drag a field directly to the filters shelf. When you add a field to the filters shelf, the Filter dialog box opens so you can define the filter. The Filter dialog box differs depending on whether you are filtering a dimension, measure, or date field. Learn how to create filters using the filters shelf in the following topics:

- Filtering Dimensions
- Filtering Measures
- Filtering Dates

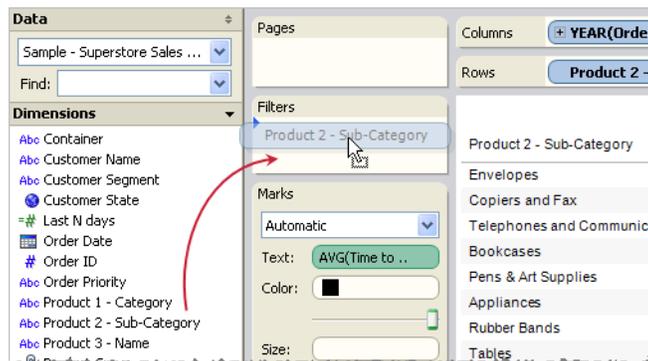
### Filtering Dimensions

Dimensions contain discrete categorical data so filtering this type of field generally involves selecting the values to include or exclude. You can create a basic categorical filter or you can define conditions and limits to create a more complex filter definition. Learn how to filter dimensions in the following topics:

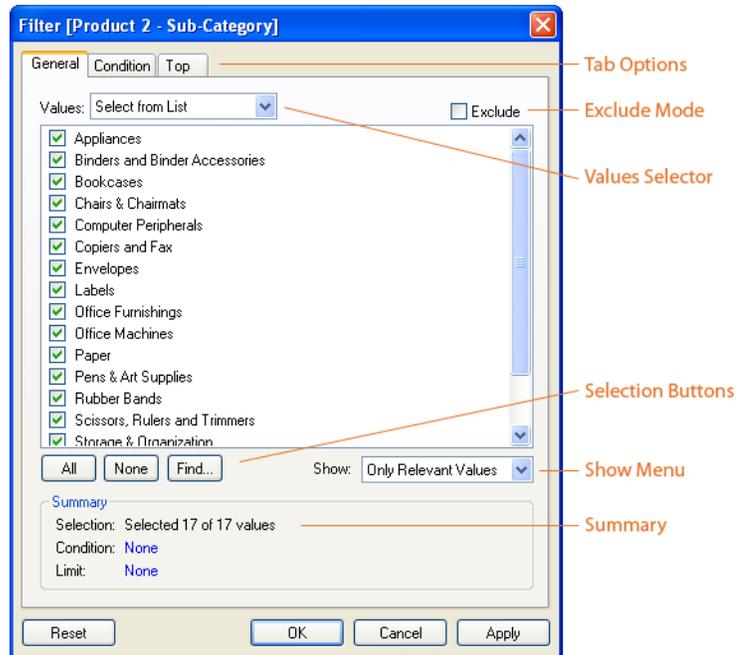
- Basic Categorical Filters
- Adding Conditions to Filters
- Adding Limits to Filters
- Example - Filtering Dimensions

### Basic Categorical Filters

- 1 Drag a field from the Data window to the Filters shelf. You can also right-click a field on any shelf and select **Filter**.



- 2 Use the General Tab of the Filter dialog box to select the values you want to include or exclude.



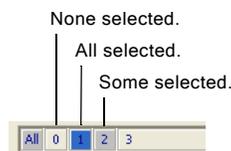
Each option on the General tab is described below:

- **Values Selector** - Use the Values drop-down list to choose a method of selecting values. Depending on the data source you are using and the type of dimension you are filtering, you can select from the following options:
  - **Select from List** - select from a list of the values (requires a database query to get the values)
  - **Wildcard Match** - type several characters to select all values that match the given pattern. You can use the asterisk character as a wildcard character. For example, type `ca*` to select all values that start with the letters "ca."
  - **Type In** - type explicit dimension member names into a text box to define a filter without querying the database. Use this option 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.

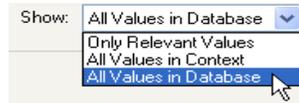
- **Use All** - select all of the members in the data source. Sometimes you will want to define a condition or limit filter that is based on all the data, even if that data changes over time. Rather than selecting specific members to include or exclude from the filter, the Use All option always includes every member in the database as the input to the condition or limit.
- **Exclude Mode** - By default, selected members when defining a filter will be included and deselected members will be excluded. However, sometimes it is easier to define what you don't want instead of all of the values you do want show. Select the **Exclude** option in the upper right corner of the dialog box to make your selections excluded from the filter instead of included.
- **Selection Controls** - These selection controls are available for multidimensional data sources and help you quickly select entire levels. Located at the top of the dialog box, the numbers indicate each level. The color shows what values are selected. The default color means no values are selected for that level, blue means all members on that level are selected, and gray means some members are selected.



- **Finding Values** - When you are working with a field that has a lot of members you may want to search the values and quickly select the ones you are looking for. Click the **Find** button to open the Find Values dialog box. Search results show in the bottom half of the dialog box. Select the values you want and click OK. You can choose to either **Replace** the selection with the values you selected in the Find dialog box or just **Add to Selection**.
- **Show Menu** - 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** menu. The Show menu includes



and excludes data from displaying in the Filter dialog box so you can find what you are looking for quicker.

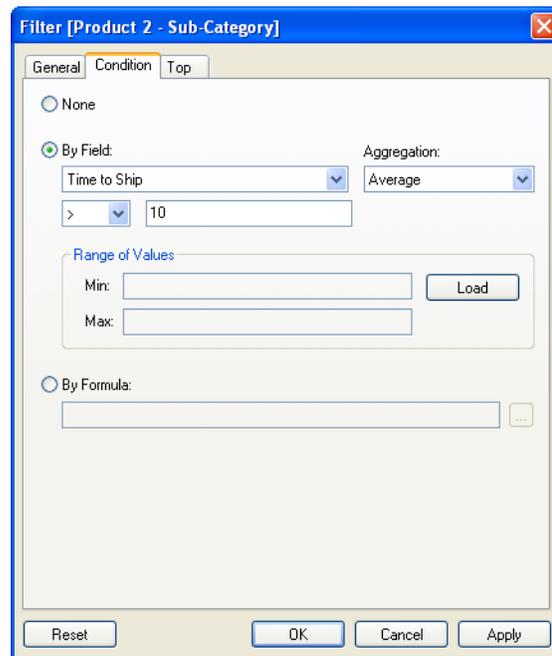


Using the **Show** menu in the filter dialog box you can select from the options described in the table below.

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

### Adding Conditions to Filters

Use the **Condition** Tab in the **Filter** dialog box to define rules to filter by. For example, in a view showing the average Time to Ship for a collection of products, you may want to only show the Products that have a Time to Ship that is greater than 10 days. You can use the built-in controls to write a condition or you can write a custom formula.



Each option on the Condition tab is described below:

- **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 Time to Ship and AVG from the first two drop-down lists. Then select Greater ( > ) from the operator list and type 10 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 Time to Ship were greater than 10 days 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.

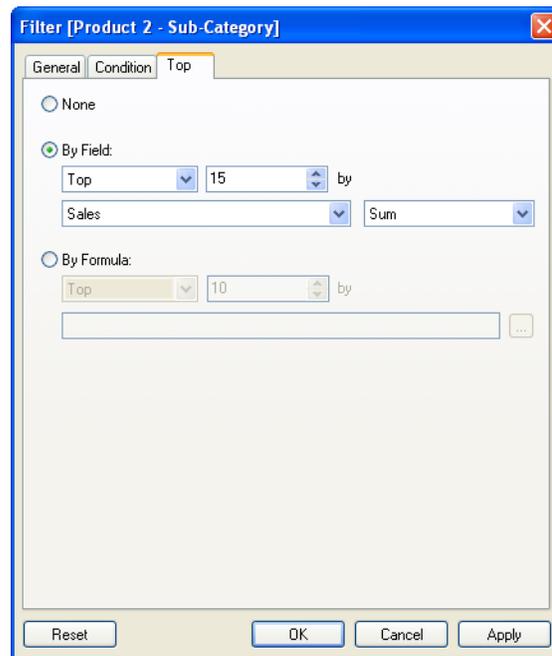
---

- **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 34, “Functions, Operators, & Data Types”.

### **Adding Limits to Filters**

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 Time to Ship 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 (e.g., 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.



Each option on the **Top** tab is described below:

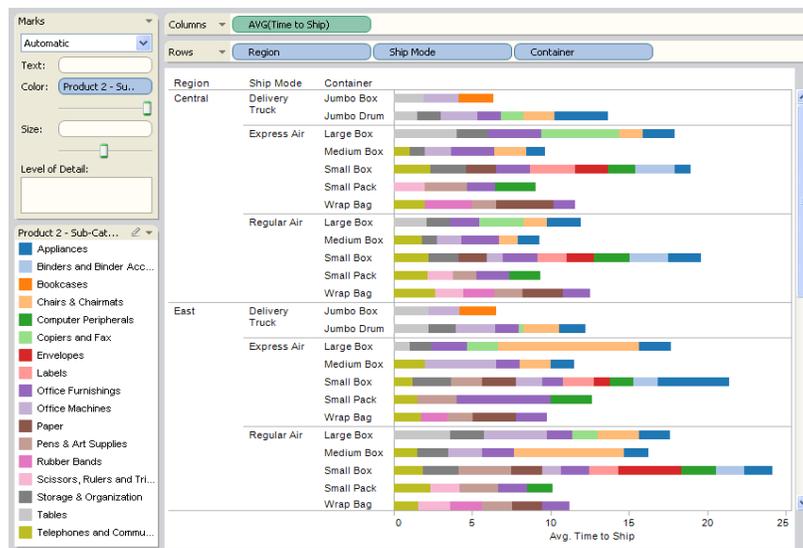
- **None:** select this option if you do not want to add a limit to the filter. This is the default setting.
- **By Field:** select this option to add a simple limit based on an existing field in the data source. First select the limit range using the first two drop-down lists. For example you can select 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 34, "Functions, Operators, & Data Types".



## 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 initial data view shown below. It was created using the Superstore Sales Excel data source. The view shows the average regional time to ship for each product based on the container and ship mode.

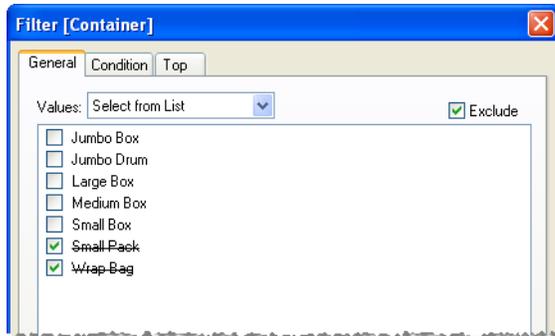


- 2 Create a basic filter on the Container dimension that excludes the Small Pack and Wrap Bag shipping containers.

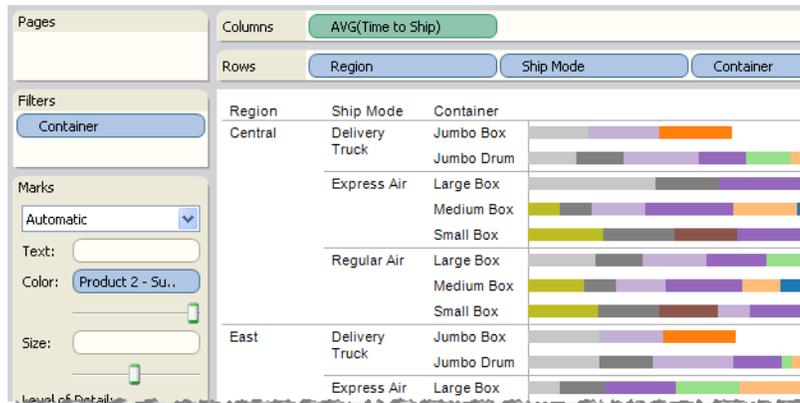
Drag the **Container** dimension to the **Filters** shelf to open the Filter dialog box. Click the **None** button at the bottom of the list to deselect all of the shipping containers. Then



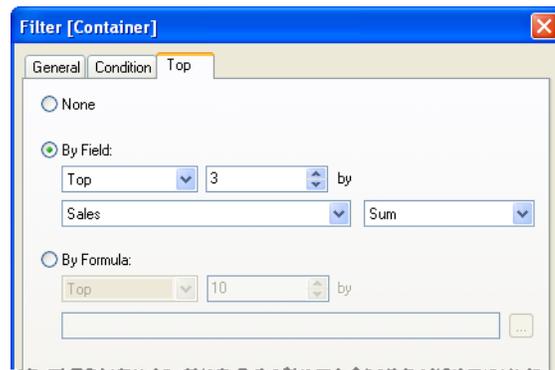
select the **Exclude** option in the upper right corner of the dialog box. Finally, select **Small Pack** and **Wrap Bag**. When finished click **OK**.



- 3 The view updates to only show orders that were not shipped in a Small Pack or Wrap Bag.



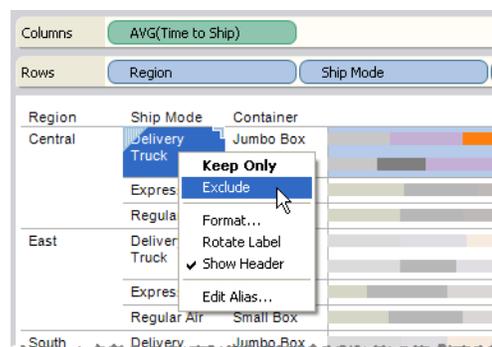
- 4 Now let's refine the filter on Container by adding a limit. Right-click the Container field on the Filters shelf and select Filter. The Filter dialog box opens. Leave the selections as they are.
- 5 Switch to the Top tab and select **By Field**. Select Top 3 from first two drop-down lists. Then select **Sales** and **SUM** from the remaining drop-down lists. When finished click **OK**.



The Top formula is computed after the selections on the General tab. So first Tableau computes all orders that were not shipped in Small Pack or Wrap Bag containers. Then the view shows just the top 3 of those orders in terms of sales.

- 6 Now let's add a new filter on Ship Mode to exclude orders that were shipped via Delivery Truck.

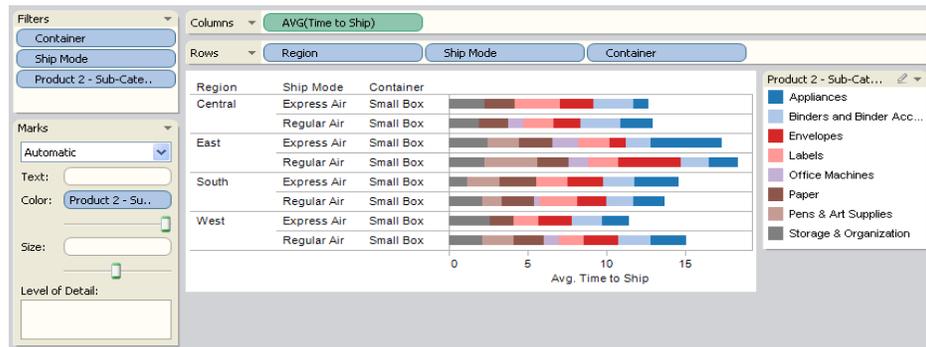
Right-click the **Delivery Truck** row header and select **Exclude**. The Delivery Truck ship mode is removed from each region in the view.



- 7 Finally, let's filter the Product 2 - Sub-Category dimension to minimize the number of colors being used in the view. Drag the **Product 2 - Sub-Category** dimension to the Filters shelf.

- 
- 8 In the Filter dialog box, deselect the Computer Peripherals, Office Machines, and Telephones and Communication values.

The final view is shown below. Take a look at the **Filters** shelf. You can easily see that the view is filtered on three separate fields. To determine which values have been excluded, open the **Filter** dialog box for each of these fields.



## Filtering Measures

Measures contain quantitative data so filtering this type of field generally involves selecting a range of values that you want to include. There are four types of quantitative filters: Range of Values, At Least, At Most, and Special. Learn about filtering measures in the following topics:

- Basic Quantitative Filters
- Showing and Hiding Values in the Filter Dialog Box
- Example – Filtering Measures

---

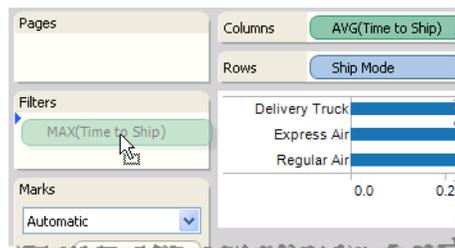
**Note** If you have a large data source, filtering measures can lead to a significant degradation in performance. It is sometimes much more efficient to filter by creating a set containing the measure and then applying a filter to the set. Refer to “How to Create a Set” on page 17-28.

---

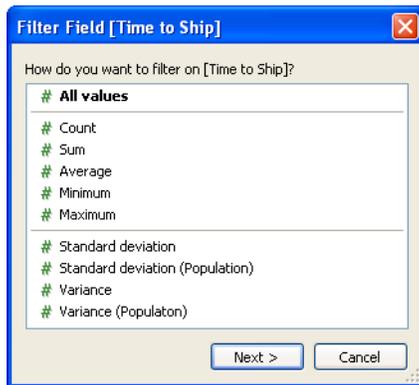


## Basic Quantitative Filters

- 1 Open the Filter dialog box dragging a measure on any shelf.



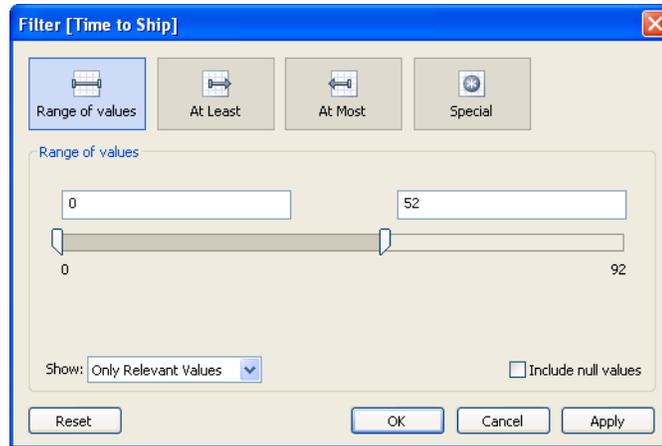
- 2 The Filter Field dialog box opens where you need to specify an aggregation. When finished, click **Next**.



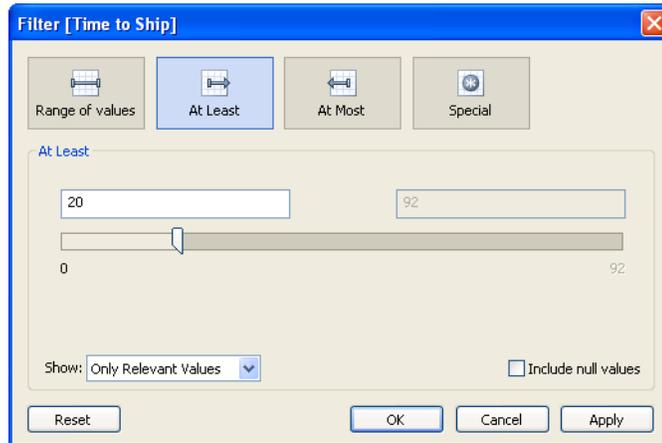
- 3 The Filter dialog box opens. There are four types of quantitative filters: Range of Values, At Least, At Most, and Special. Each of these types of filters are described below:



- **Range of Values** - Specify the minimum and maximum values of the range to include in the view. The values you specify are included in the range.

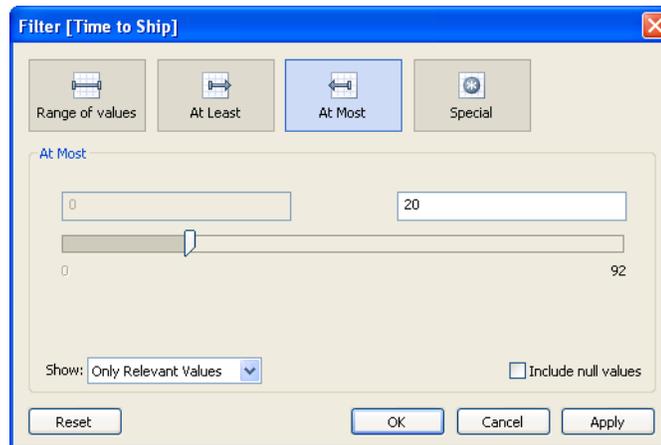


- **At Least** - Include all values that are greater than or equal to a specified minimum value. This type of filter is useful when the data changes often so specifying an upper limit may not be possible.

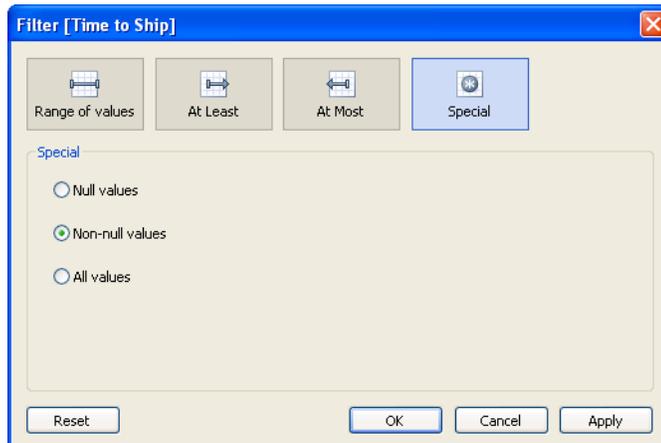




- **At Most** - Include all values that are less than or equal to a specified maximum value. This type of filter is useful when the data changes often so specifying a lower limit may be not be possible.



- **Special** - This special type of filter helps you filter on Null values. Include only **Null values**, **Non-null values**, or **All Values**.

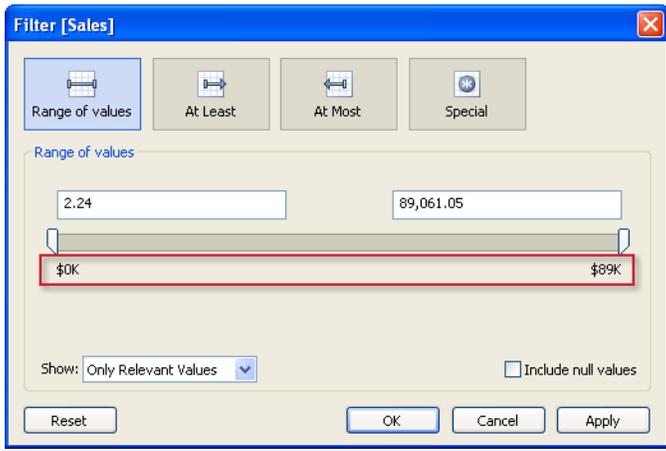


- 4 When finished defining the filter click **OK**.

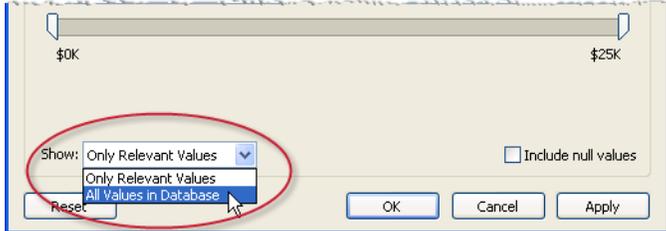


## Showing and Hiding Values in the Filter Dialog Box

The filter dialog box shows the minimum and maximum values for the field below the range slider. These numbers give you context when you are deciding the range of values to include in the filter.



These minimum and maximum values are affected by the other filters set on the view. For example, a database may include records with sales ranging from \$0 to \$89K. If you created a filter on the Sales field the minimum and maximum values shown in the filter dialog box would indicate this range. However, let's say you then filter the view to only show Office Supply products, which sell for between \$0 and \$25K. By default the filter dialog box will consider that filter and only show the office supplies range. You can use the **Show** menu in the bottom left corner of the dialog box to switch between **Only Relevant Values** and **All Values in the Database**. These options only affect the range that is shown in the filter dialog box and doesn't change how the filter will be applied to the view.





### Example – Filtering Measures

This example filters a text table using an aggregated measure, and then filters the table using the same measure in an unaggregated state.

- 1 Create the initial view using the Sample - Superstore Sales (Excel) data source. The text table is shown below.

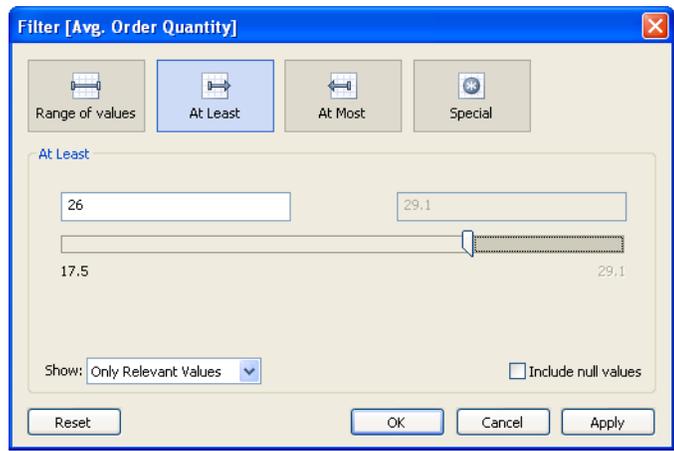
The screenshot shows the Tableau interface. On the left, the 'Marks' shelf is set to 'Automatic' with the text 'AVG(Order Qu...)'. The 'Columns' shelf contains 'Ship Mode' and the 'Rows' shelf contains 'Product 2 - Sub-Cate...'. The main view displays a text table with the following data:

Product 2 - Sub-Category	Delivery Truck	Express Air	Regular Air
Appliances	25.29	24.24	25.33
Binders and Binder Accesso..		25.12	25.10
Bookcases	26.00		
Chairs & Chairmats	24.77	26.57	25.00
Computer Peripherals		26.56	25.19
Copiers and Fax	28.20	25.33	23.36
Envelopes		27.56	26.56
Labels		27.06	25.00
Office Furnishings	26.00	25.97	26.31
Office Machines	23.72	17.60	23.08
Paper		24.98	25.24
Pens & Art Supplies		27.03	26.85
Rubber Bands		25.35	27.97
Scissors, Rulers and Trimme..		29.10	27.01
Storage & Organization	27.23	28.54	24.56
Tables	25.46	23.00	25.27
Telephones and Communica..		24.77	26.23



- 2 Filter the data to only show orders with an average quantity of 26 or more. You can create this type of filter by dragging the **Order Quantity** measure to the Filters shelf and select **Average** as the aggregation.

The Filter dialog box is shown below. This type of filter is an At Least filter with the minimum value set to 26.



- 3 When finished, click **OK**.

The modified 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, Copiers & Faxes shipped by Express Air and Regular Air are removed from the



view because the average order quantity is less than 26, while Copiers & Faxes shipped by Delivery Truck remains in the view because the average order quantity is greater than 26.

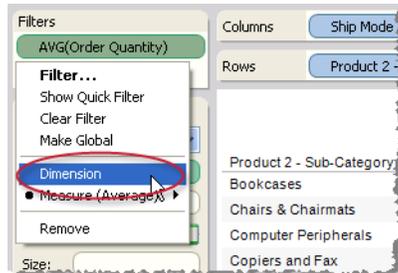
The screenshot shows a Tableau interface with the following settings:

- Filters:** AVG(Order Quantity)
- Columns:** Ship Mode
- Rows:** Product 2 - Sub-Cate..
- Marks:** Automatic
- Text:** AVG(Order Qu..)
- Color:** Black
- Size:** (slider)
- Level of Detail:** (empty)

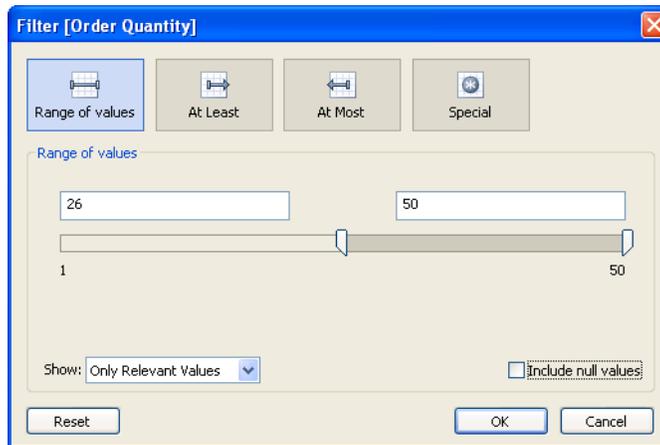
The resulting table is as follows:

Product 2 - Sub-Category	Delivery Truck	Express Air	Regular Air
Bookcases	26.000		
Chairs & Chairmats		26.571	
Computer Peripherals		26.561	
<b>Copiers and Fax</b>	<b>28.200</b>		
Envelopes		27.563	26.556
Labels		27.059	
Office Furnishings	26.000		26.311
Pens & Art Supplies		27.025	26.848
Rubber Bands			27.969
Scissors, Rulers and Trimme..		29.100	27.008
Storage & Organization	27.226	28.541	
Telephones and Communica..			26.229

- 4 Now let's filter the same view using a disaggregated measure. Suppose you want to filter the view using the disaggregated **Order Quantity** measure. To do this, select **Dimension** on the context menu of the AVG(Order Quantity) field on the Filters shelf.



The **Filter** dialog box is shown below. It displays the limits of the individual rows for the Order Quantity measure. Specify a new lower limit of 26.



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



source that has an order quantity that is less than 26, and then aggregates the remaining rows as an average.

The screenshot shows the Tableau interface with the following settings:

- Filters:** Order Quantity
- Columns:** Ship Mode
- Rows:** Product 2 - Sub-Cate..
- Marks:** Automatic, Text: AVG(Order Qu..), Color: Black, Size: [Slider]
- Level of Detail:** [Empty]

Product 2 - Sub-Category	Delivery Truck	Express Air	Regular Air
Appliances	38.182	37.885	38.513
Binders and Binder Accesso..		36.527	38.028
Bookcases	36.709		
Chairs & Chairmats	37.920	38.333	36.762
Computer Peripherals		37.632	37.205
Copiers and Fax	39.625	41.167	39.074
Envelopes		36.650	39.083
Labels		37.684	37.685
Office Furnishings	41.333	37.909	39.148
Office Machines	38.144	35.750	38.778
Paper		37.385	37.540
Pens & Art Supplies		39.475	38.165
Rubber Bands		37.000	38.681
Scissors, Rulers and Trimme..		38.417	38.343
Storage & Organization	36.667	38.349	37.127
Tables	38.692	40.000	40.636
Telephones and Communica..		38.542	38.425

## Filtering Dates

Date fields are a special kind of dimension that Tableau often handles differently than standard categorical data. This is especially true when you are creating date filters. Date filters are extremely common and fall into three categories: Relative Date Filters, which show a date range that is relative to a specific day; Range of Date Filters, which show a defined range of discrete dates; and Discrete Date Filters, which show individual dates that you've selected from a list. Learn about each of these types of filters in the following topics:

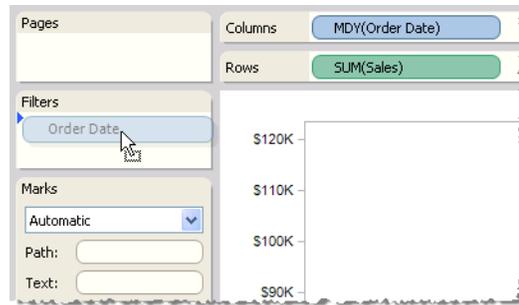
- Relative Date Filters
- Range of Dates
- Other types of Date Filters
- Discrete Date Filters
- Example - Filtering Dates



### **Relative Date Filters**

A relative date filter lets you define a range of dates that updates based on the date and time you open the view. For example, you may want to see Year to Date sales, all records from the past 30 days, or bugs closed last week. Relative date filters can also be relative to a specific anchor date rather than today. Follow the steps below to create a relative date filter.

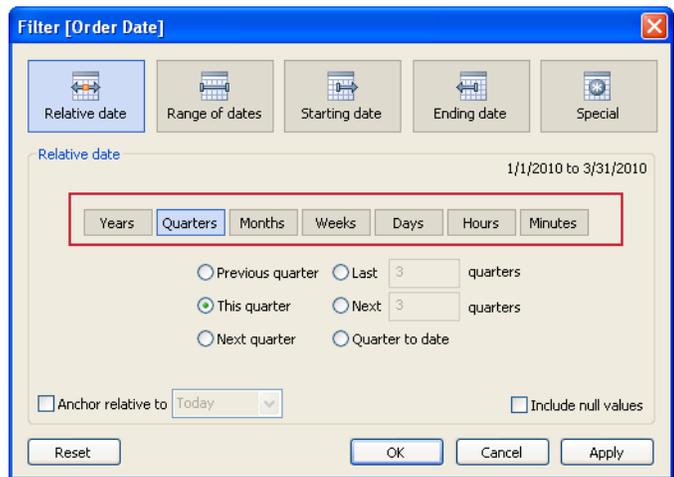
- 1 Drag a date field from the Data window and drop it on the Filters shelf.



- 2 In the Filter Field dialog box, select **Relative to Now** and then click **Next**.



- 3 The Filter dialog box opens showing the Relative to Now options. Select a unit of time to filter by. For example, to filter to show the last 2 quarters, select Quarter as the time unit.

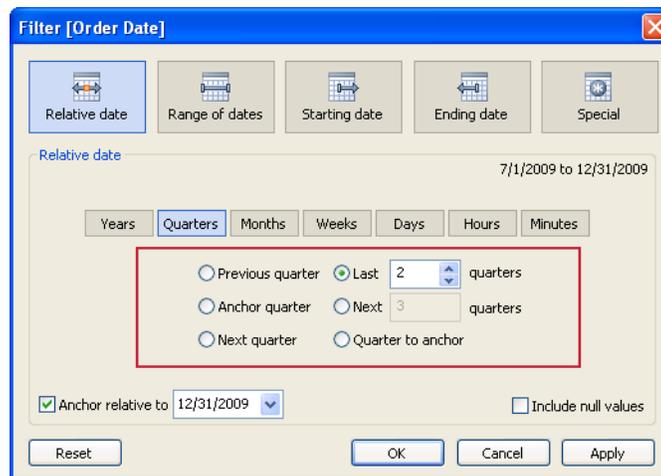


- 4 Use the rest of the controls to define the date filter. You can select from a variety of common options including current, previous, and next. By default, the filter is relative



to today. To make the filter relative to an alternate date select the **Anchor relative to** option in the bottom left corner the and select the date to anchor to.

The date period includes the current unit of time. For example, selecting Last 2 Quarters will include the current quarter and the previous quarter. Use the preview in the upper right corner to check your filter settings.



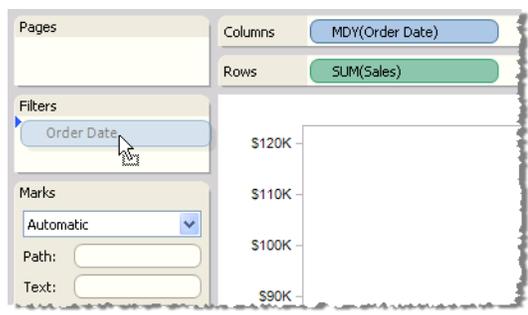
5 When finished, click **OK**.

### **Range of Dates**

Use this type of filter to define a fixed range of dates. For example, you may want to see all orders placed between March 1, 2009 and June 12, 2009. The Range of Dates filter is similar to the Range of Values option when creating Basic Quantitative Filters. Follow the steps below to create a Range of Dates filter.



- 1 Drag a date field from the Data window and drop it on the Filters shelf.



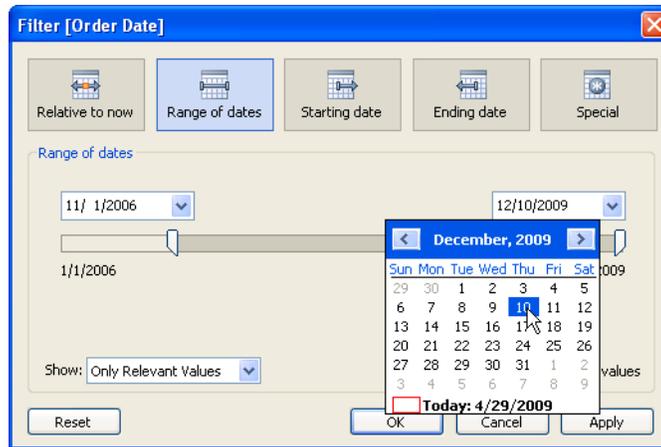
- 2 In the Filter Field dialog box, select **Range of Dates** and then click **Next**.



- 3 The Filter dialog box opens showing the Range of Dates options. Use the slider or the drop-down date controls to select minimum and maximum dates for the range you want



to include. The range is inclusive, which means that the minimum and maximum dates are included in the filter.



4 When finished, click **OK**.

---

**Note** If the field also includes Time you can select the Show Times option to further refine your filter range.

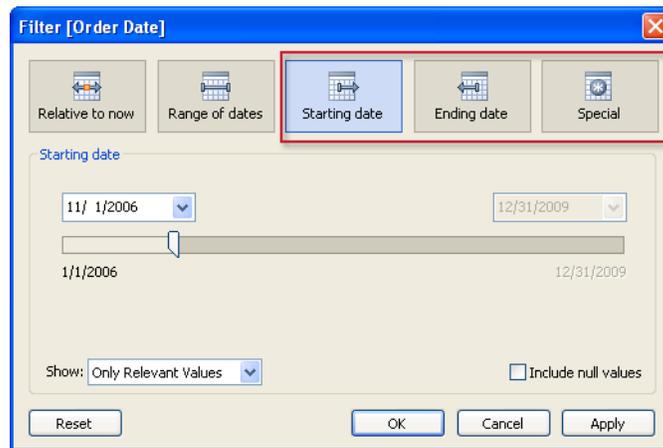
---

### **Other types of Date Filters**

Just like Basic Quantitative Filters, you can also filter dates by defining just a Starting Date or an Ending Date. These filters are useful when you want to define an open ended range.

In addition, you can create Special filters that include only Null dates, Non-null dates, or All dates.

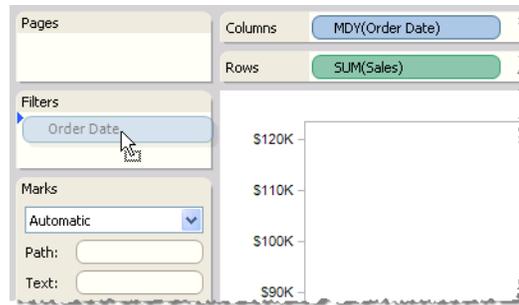
Use the options at the top of the Filter dialog box to define these types of filters.



### **Discrete Date Filters**

Sometimes you may want to filter to include specific individual dates or entire date levels. This type of filter is called a Discrete Date Filter because you are defining discrete values instead of a range. Follow the steps below to create a discrete date filter.

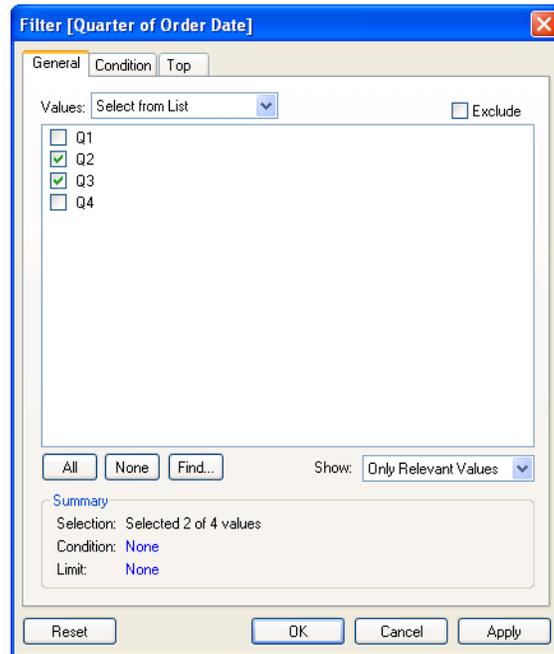
- 1 Drag a date field from the Data window and drop it on the Filters shelf.



- 2 In the Filter Field dialog box, select a date level or select **Individual dates** and then click **Next**.



- 3 In the Filter dialog box, select the dates you want to include. Refer to Basic Categorical Filters to learn more about specifying this type of filter.

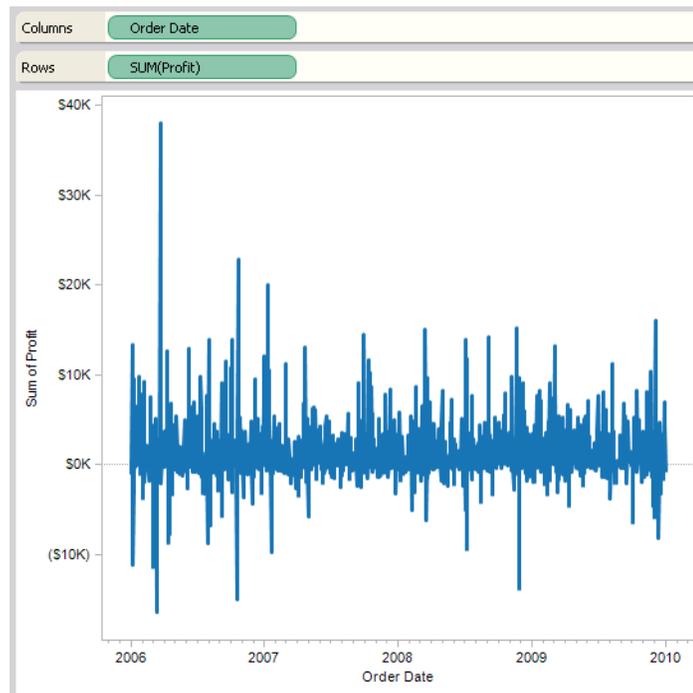


4 When finished, click **OK**.

### **Example - Filtering Dates**

This example filters a line graph, to show the profit over a specific range of time. The steps are as follows:

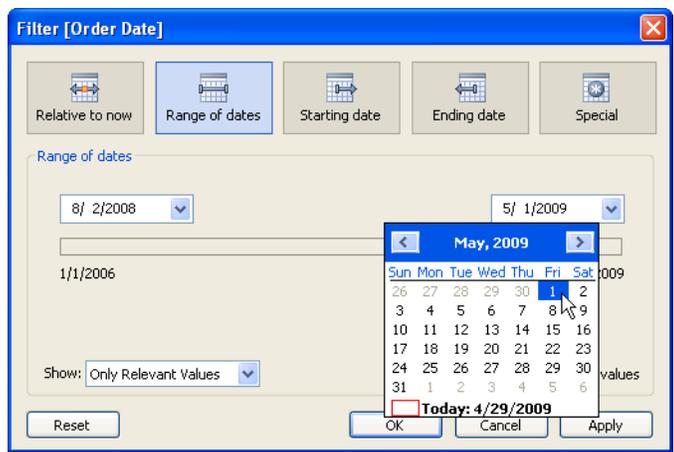
- 1 Create the initial view shown below. It was created using the Superstore Sales Excel data source. Place **Order Date** on to the **Columns** shelf and select All Values as the aggregation. Then place **Profit** onto the **Rows** shelf.





- 2 Now let's filter the view to include only orders that were place between August 2, 2008 and May 1, 2009. To create this filter drag the Order Date field to the Filters shelf and select Range of Dates in the Filter Field dialog box. Then click Next.

The **Filter** dialog box is shown below. It displays the Order Date limits. Use the drop-down date controls to specify a new lower limit of August 2, 2008 and an upper limit of May 1, 2009.



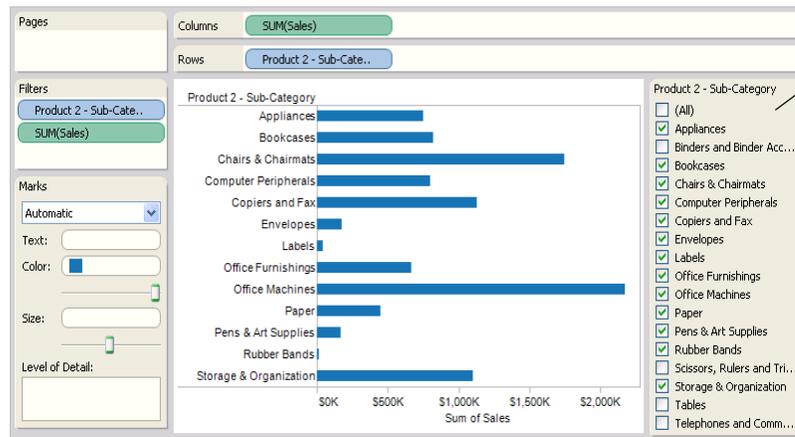
The filtered view is shown below.



## Using Quick Filters

Tableau lets you quickly add and modify filters using Quick Filters. 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. Learn more about quick filters in the following topics:

- Turning on Quick Filters
- Quick Filter Options
- Searching Quick Filters



Quick filter card for the Product Category field.

## Turning on Quick Filters

A Quick Filter can be turned on for existing filters or for non-filtered fields. To show or hide a quick filter, select **Show Quick Filter** from the field's context menu.





## Quick Filter Options

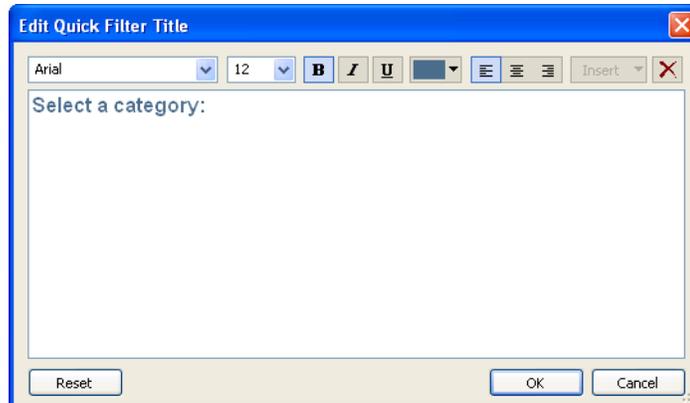
After you've turned on a quick filter there are many different options that let control how the filter works and its appearance. You can access these options using the card menu in the upper right corner of the quick filter card. Some options are available for all types of filters and others depend on whether you're filtering a Categorical field (dimensions) or a Quantitative field (measures). Finally, you can customize how quick filters display on the sheet, in dashboards, or when published to Tableau Server.

- General Quick Filter Options
- Categorical Quick Filter Options
- Quantitative Quick Filter Options
- Customizing Quick Filters

### General Quick Filter Options

- **Edit** - This option opens the main Filter dialog box so you can further refine the filter by adding conditions and limits.
- **Clear Filter** - Removes the filter from the Filters shelf and removes the quick filter.
- **Make Global** - Make the filter global, which means it applies to all sheets that use the same data source. Refer to Global Filters to learn more.
- **Only Relevant Values** - Specifies which values to show in the quick filter. When you select this option other filters are considered and only values that pass these filters are shown. For example, a quick filter on State will only show the Eastern states when a filter on Region is set. You can use the toggle at the top of the quick filter card to switch between this option and the All Values in Database option.
- **All Values in Database** - Specified which values to show in the quick filter. When you select this option all values in the database are shown regardless of the other filters on the view.

- 
- **Edit Title** - By default the title of the quick filter is the name of the field being filtered. Use this option to modify the title. Click **Reset** to return to the default title.

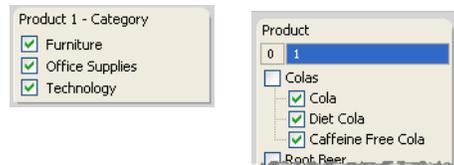


- **Hide Card** - Hides the quick filter card but does not remove the filter from the Filters shelf.

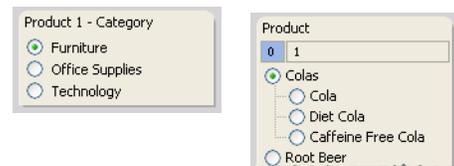


## Categorical Quick Filter Options

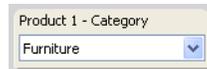
- **Include Values** - The items selected in the quick filter will be included in the view.
- **Exclude Values** - The items selected in the quick filter will be excluded from the view.
- **Multiple Values List** - Displays the values in the quick filter as a list of checkboxes where multiple values can be selected.



- **Single Value List** - Displays the values of the quick filter as a list of radio buttons where only a single value can be selected at a time. An “All” option can be added to the list to let you quickly select all values without switching to a multiple values list.



- **Compact List** - Displays the values of the quick filter in a drop-down list where only a single value can be selected at a time.

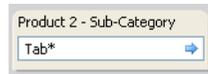




- **Slider** - Displays the values of the quick filter along the range of a slider. Only a single value can be selected at a time. This option is useful for dimensions that have an implicit order such as dates.



- **Wildcard Match** - Displays a text box where you can type a few characters. All values that match those characters are automatically selected. You can use the asterisk character as a wildcard character. For example, you can type "tab\*" to select all values that begin with the letters "tab". Pattern Match is not case sensitive. If you are using a multidimensional data source, this option is only available when filtering single level hierarchies and attributes.





## Quantitative Quick Filter Options

- **Range of Values/Dates** - 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.

The darker area inside the slider range 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 Sales field to only include values between \$200,000 and \$500,000 but your view only contains values between \$250,000 and \$320,000. 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.



- **At Least/Starting Date**- shows a single slider with a fixed minimum value. Use this option to create a filter using an open ended range.

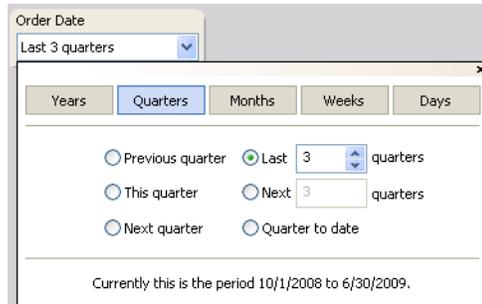


- **At Most/Ending Date** - shows a slider with a fixed maximum value. Use this option to create a filter using an open ended range.





- **Relative to Now** - shows a control where you can define a dynamic date range that updates based on when you open the view. The option is only available for filters on continuous date fields.

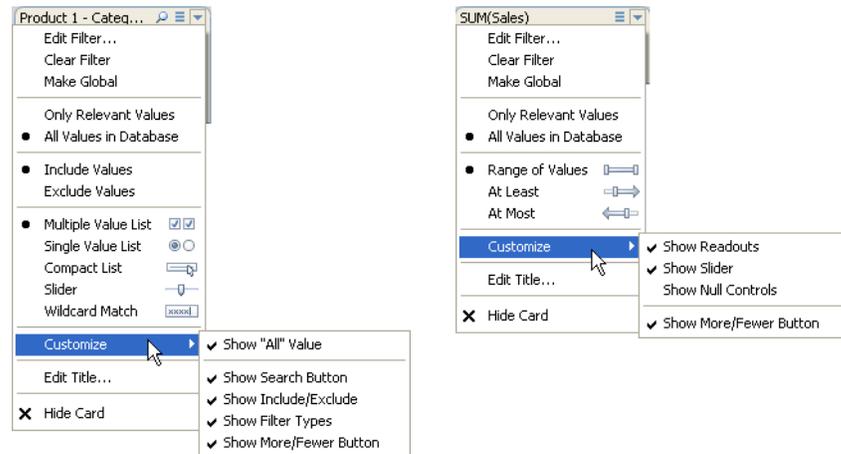


- **Browse Periods** - shows common date ranges such as past day, week, month, three months, one year, and five years. This option is only available for filters on continuous date fields.



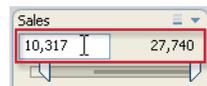
### Customizing Quick Filters

You can control how a quick filter control appears on the sheet, in dashboards, or when published to Tableau Server. Customize quick filters by selecting Customize on the quick filter card menu.



Then select from the following options:

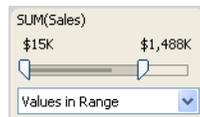
- **Show “All” Value** - toggles whether to show the “All” option that displays by default in multiple values and single value lists.
- **Show Search Button** - toggles whether to show the search button at the top of the quick filter.
- **Show Include/Exclude** - toggles whether to show the Include Values and Exclude Values commands on the quick filter card menu. When shown, users can switch the quick filter between include and exclude modes.
- **Show Filter Types** - toggles whether to let users change the type of quick filter is shown. For example, when shown, a user can change a multiple values list to a compact list.
- **Show More/Fewer Button** - toggles whether to show the More/Fewer button at the top of the quick filter.
- **Show Readouts** - controls whether the minimum and maximum values are displayed as text above a range of values. The readouts can be used to manually type a new value instead of using the sliders.



- **Show Null Controls** - shows a drop-down list that lets you control how the filter handles null values. You can select from from the following options:



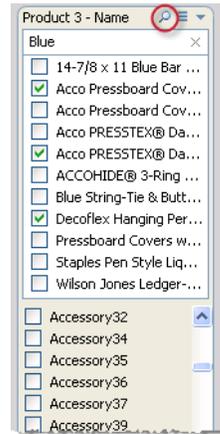
- Values in Range - the filter only includes values within the specified range.
- Values in Range and Null Values - the filter includes values within the specified range as well as null values.
- Null Values Only - the filter includes only null values.
- Non-Null Values Only - the filter includes only values that are not null.
- All Values - the filter includes all values. Use this option to quickly reset the selected range to include all values.



### Searching Quick Filters

Sometimes a categorical quick filter may contain a lot of values. You can use the Search option to quickly find and select the values you want. To open the search field, click the Search icon in the upper right corner of the quick filter card. Then start typing what you want to select. Matching values that contain the specified characters will show directly below the search field where you can select or deselect them as needed.

By default, search will return all values that contain the search term. You can use the asterisk character as a wildcard to restrict the results to values that begin with or end with the specified characters. For example, searching for “Bl\*” will find all values that start with the characters b and l. Search is not case sensitive.



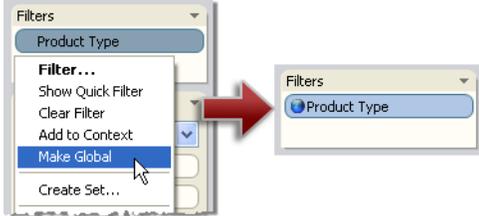


## Global Filters

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 sheet, 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.

---

## 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 categorical 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. You can set one or more context filters to improve performance.
- 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.

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 breakfast products. Then you can create 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.

Learn how to create context filters in the following topics:

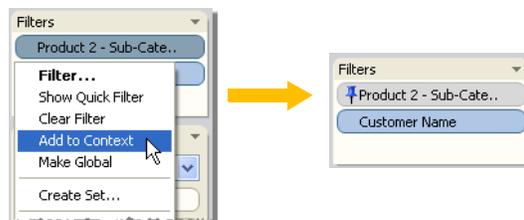
- Creating a Context Filter
- Example – Context Filters

## Creating 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 a grey color and the pushpin 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:

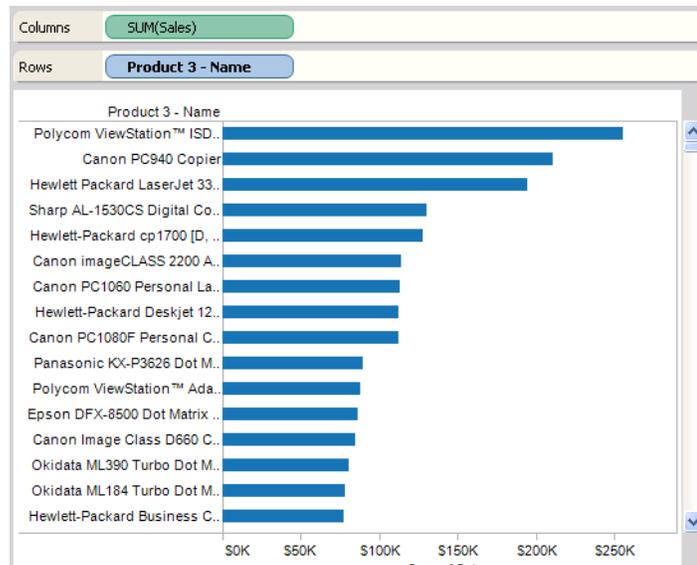
- Removing the field from the **Filters** shelf – If other context filters remain on the shelf, a new context is computed.
- Editing the filter – A new context is computed each time you edit a context filter.
- Selecting **Remove from Context** – The filter remains on the shelf as a standard categorical filter. If other context filters remain on the shelf, a new context is computed.

## Example – Context Filters

This example walks you through how to create a context filter. First you'll filter a view to show the top 10 products by sales. Then you'll create a context filter on product category so you can see the top 10 furniture products.



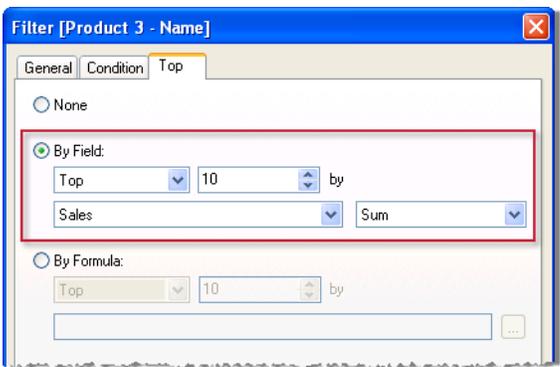
- 1 Use the Sample - Superstore Sales data source to create the initial view shown below. The view shows the sales for all products sorted with the highest sale at the top.



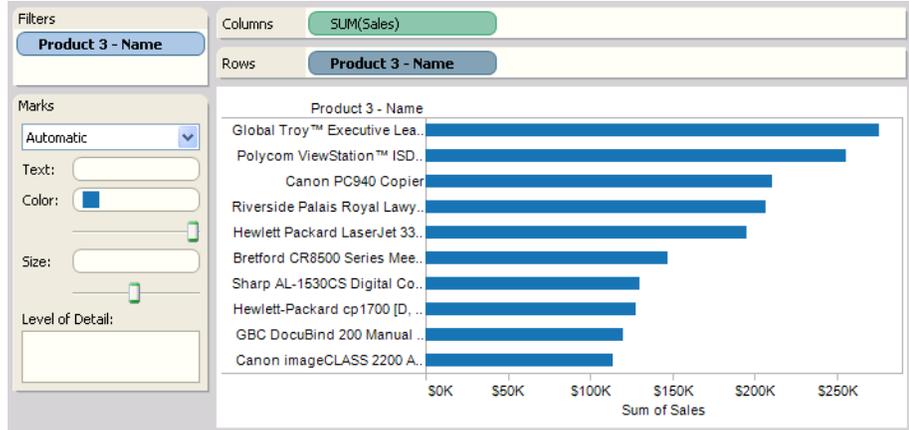
- 2 Now create a Top 10 filter to just show the top selling products. You can create this filter by dragging the **Product 3- Name** field to the Filters shelf. In the filter dialog box, switch



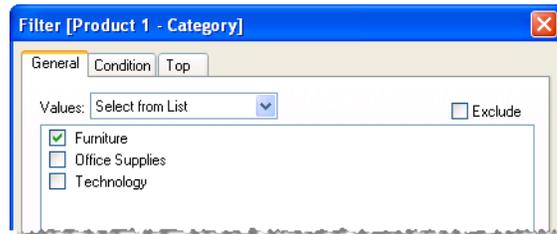
to the **Top** tab and define a filter that is Top 10 by Sum of Sales. Refer to “Adding Limits to Filters” on page 16-12 to learn more about defining a Top N filter.



- 3 When you click **OK**, you'll see that the view is filtered to show the top 10 products in terms of sales.



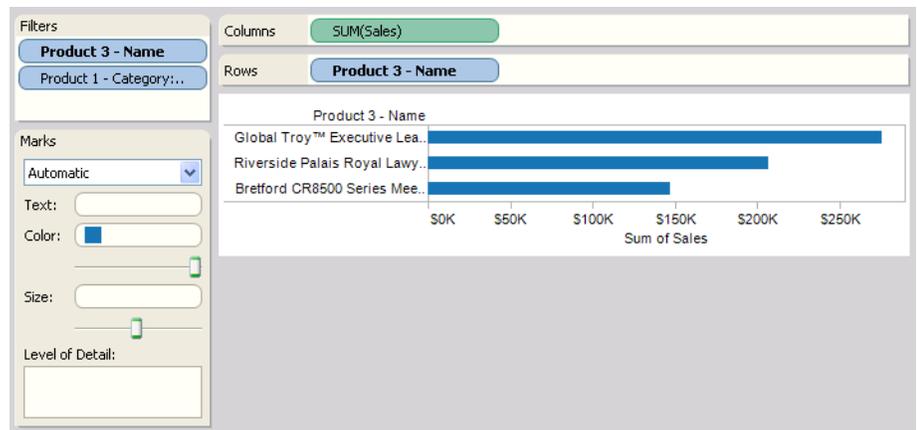
- 4 Now, let's add another filter to only show only furniture products. Drag the **Product 1 - Category** field to the Filters shelf and select **Furniture**. When finished, click **OK**.



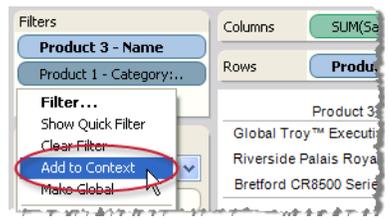
- 5 The view is filtered but instead of 10 products, it now only shows 3. The reason is because by default all filters are evaluated separately and the view shows the union of



the results. So this view shows that three of the top 10 overall products are furniture products.



- To find out what the top 10 furniture products are we need to make the Product 1 - Category filter a context filter. Right-click the field on the Filters shelf and select **Add to Context**.



- The filter is marked as a context filter and the view updates to show the top 10 furniture products. Tableau has first evaluated the data source and identified all of the furniture products. Then the Top 10 filter is evaluated on the results of that context.





# Calculation Filters

Filters on dimensions that are not used elsewhere in the view are called calculation filters. For these types of filters, Tableau performs a calculation on the selected dimension members. This occurs when:

- The dimension is only on the **Filters** shelf (not used on other shelves).
- You define the filter to include multiple values.

The calculation icon  $\Sigma$  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  $\Sigma$  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

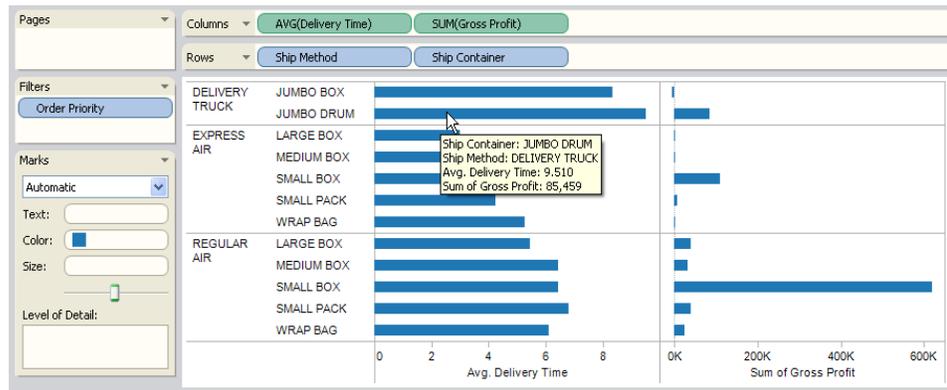
Consider the view shown below. It consists of the **Profit** measure aggregated as a summation and the **Order Quantity** measure aggregated as an average. These measures are displayed with the **Ship Mode** and **Container** dimensions. An external filter that consists of two members of the **Order 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 **Profit** and an average is performed for **Order Quantity**.



For example, the tooltip shows the data for Jumbo Drums delivered by truck. The average order quantity is 24.3. This number was calculated by averaging the order quantities for all the rows that have an Urgent or High order priority. Similarly, the sum of profit is \$114,363. This number was calculated by summing the profit for all the rows that have an Urgent or High order priority.

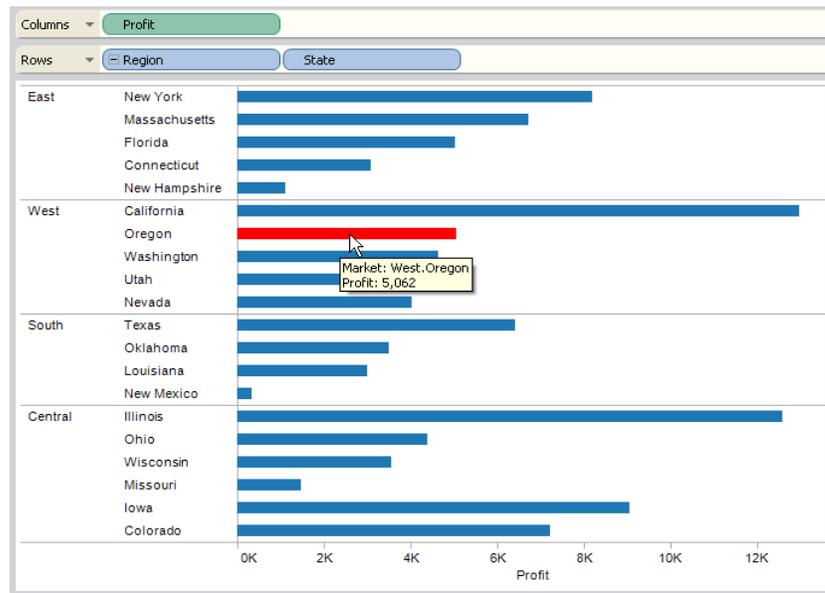


### Calculation Filters: Multidimensional Data Source

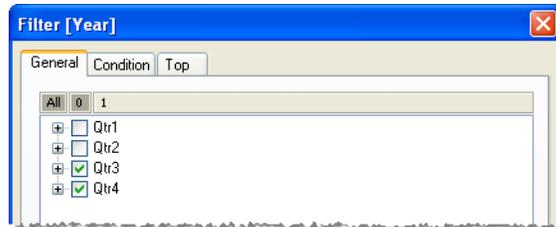
You can think of a calculation filter as “slicing” the cube’s data through the specified dimensions. Therefore, a calculation filter for a multidimensional data source is sometimes referred to as a slicer.

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.

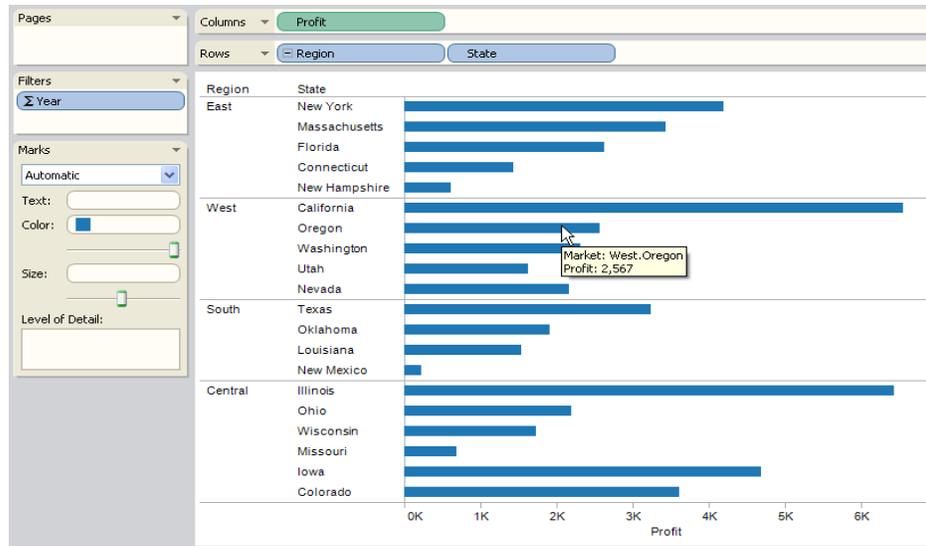
For example, consider the view shown below. The tool tip shows that the profit for Oregon is 5,062.



An external filter that consists of the Q3 and Q4 members of the **Quarter** dimension is applied to the data.



For example, in the view below the profit for Oregon is now \$2,567. This number was calculated by summing the data for Qtr3 and Qtr4.



# Sorting, Grouping, and Sets

---

<b>Overview</b> . . . . .	<b>17-2</b>
<b>Sorting</b> . . . . .	<b>17-3</b>
Computed Sorting . . . . .	17-3
Manual Sorting . . . . .	17-14
<b>Groups</b> . . . . .	<b>17-18</b>
Creating Groups . . . . .	17-18
Editing an Existing Group . . . . .	17-22
Finding Members in the Groups Dialog Box . . . . .	17-25
<b>Sets</b> . . . . .	<b>17-26</b>
About Sets . . . . .	17-26
How to Create a Set . . . . .	17-28
Creating Sets Examples . . . . .	17-37

---

## Overview

After you understand the basics of building data views, use sorting, 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, 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.
- **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 to highest sales. 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 “Groups” on page 17-18 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 27-43 for more information.

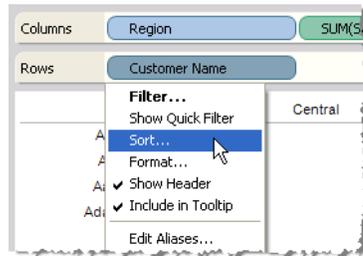
### How to Sort Data (Computed Sorts)

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 hierarchal order of the members within a dimension.

**Alphabetic** - the order of the letters in the alphabet.

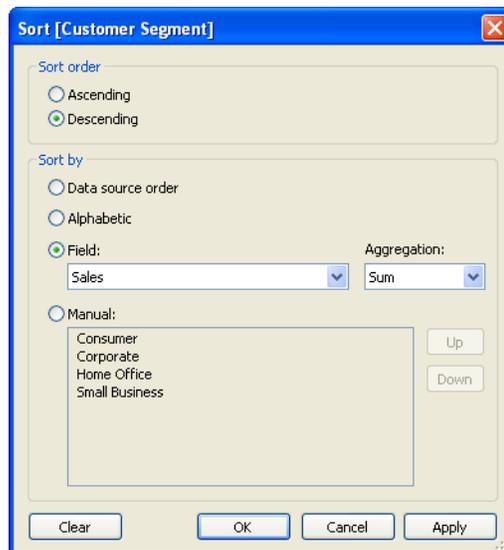
**Field** - order the data based on the associated values of another field. For example, you could order several products by their total sales values.

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 dimensions by a measure. For example, the **Sort** dialog box shown below is configured to sort the members of the **Customer Segment** field in descending order and by the sum of the **Sales** measure. The results will be displayed so that the member with the highest sales is displayed first, the member with the second highest sales 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 17-6 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 17-10 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 discounts. To create the view, follow the steps below:

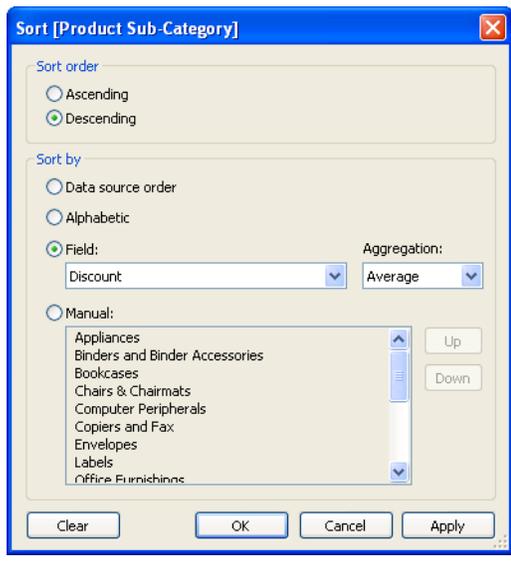
- 1 Place the **Order Date** dimension on the **Columns** shelf and the **Product Sub-Category** dimension on the **Rows** shelf.

Complete the text table by placing **Discount** on the **Text** shelf and aggregating the measure as an average (select **Measure > Average** from the field's context menu). By default, the table is sorted in alphabetical order.

	2006	2007	2008
Appliances	5.08%	5.03%	5.25%
Binders and Binder Accesso..	5.03%	4.63%	5.33%
Bookcases	4.20%	4.40%	5.17%
Chairs & Chairmats	4.93%	5.42%	4.90%
Computer Peripherals	4.94%	4.61%	5.08%
Copiers and Fax	5.05%	5.19%	5.55%
Envelopes	5.08%	5.44%	4.63%
Labels	5.62%	5.20%	4.60%
Office Furnishings	4.54%	5.28%	5.02%

- 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 **Discount** aggregated as an **Average**. When finished click **OK**. Then apply the same sort to **Product Sub-Category**.



The view is shown below. **Rubber Bands** is the top row in the table because it has the largest average discount across all years, while **Telephones and Communications** are at the bottom in the table because that category has the smallest average discount across all years. Similarly, **2008** is the left most column because it has the largest average discount for all products, while **2006** is the right most column because it has the smallest average discount for all products.



The screenshot shows the Tableau interface with the following configuration:

- Columns:** YEAR(Order Date)
- Rows:** Product Sub-Categ..
- Marks:** Automatic (dropdown), Text: AVG(Discount) (green pill), Color: black (color picker), Size: (slider), Level of Detail: (empty)

	2008	2009	2007	2006
Rubber Bands	5.14%	4.98%	4.92%	5.79%
Copiers and Fax	5.55%	4.80%	5.19%	5.05%
Storage & Organization	5.27%	5.26%	5.30%	4.68%
Pens & Art Supplies	5.14%	4.80%	5.11%	5.38%
Office Machines	5.54%	5.29%	4.71%	4.82%
Appliances	5.25%	4.90%	5.03%	5.08%
Binders and Binder Accesso..	5.33%	5.13%	4.63%	5.03%
Chairs & Chairmats	4.90%	4.87%	5.42%	4.93%
Tables	5.09%	4.95%	4.39%	5.49%
Paper	5.07%	5.12%	4.97%	4.61%
Labels	4.60%	4.35%	5.20%	5.62%
Scissors, Rulers and Trimme..	5.32%	5.13%	4.06%	5.03%
Office Furnishings	5.02%	4.77%	5.28%	4.54%
Envelopes	4.63%	4.24%	5.44%	5.08%
Computer Peripherals	5.08%	4.82%	4.61%	4.94%
Bookcases	5.17%	5.64%	4.40%	4.20%
Telephones and Communica..	4.69%	5.08%	4.84%	4.56%

At first glance, it's not clear if the data has 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.



Columns		+ YEAR(Order Date)				
Rows		Product Sub-Categ..				
	2008	2009	2007	2006	Grand Total	
Rubber Bands	5.14%	4.98%	4.92%	5.79%	5.18%	
Copiers and Fax	5.55%	4.80%	5.19%	5.05%	5.15%	
Storage & Organization	5.27%	5.26%	5.30%	4.68%	5.12%	
Pens & Art Supplies	5.14%	4.80%	5.11%	5.38%	5.11%	
Office Machines	5.54%	5.29%	4.71%	4.82%	5.06%	
Appliances	5.25%	4.90%	5.03%	5.08%	5.06%	
Binders and Binder Accesso..	5.33%	5.13%	4.63%	5.03%	5.04%	
Chairs & Chairmats	4.90%	4.87%	5.42%	4.93%	5.03%	
Tables	5.09%	4.95%	4.39%	5.49%	4.98%	
Paper	5.07%	5.12%	4.97%	4.61%	4.94%	
Labels	4.60%	4.35%	5.20%	5.62%	4.93%	
Scissors, Rulers and Trimme..	5.32%	5.13%	4.06%	5.03%	4.92%	
Office Furnishings	5.02%	4.77%	5.28%	4.54%	4.90%	
Envelopes	4.63%	4.24%	5.44%	5.08%	4.86%	
Computer Peripherals	5.08%	4.82%	4.61%	4.94%	4.85%	
Bookcases	5.17%	5.64%	4.40%	4.20%	4.80%	
Telephones and Communica..	4.69%	5.08%	4.84%	4.56%	4.80%	
<b>Grand Total</b>	<b>5.07%</b>	<b>4.97%</b>	<b>4.94%</b>	<b>4.89%</b>	<b>4.97%</b>	

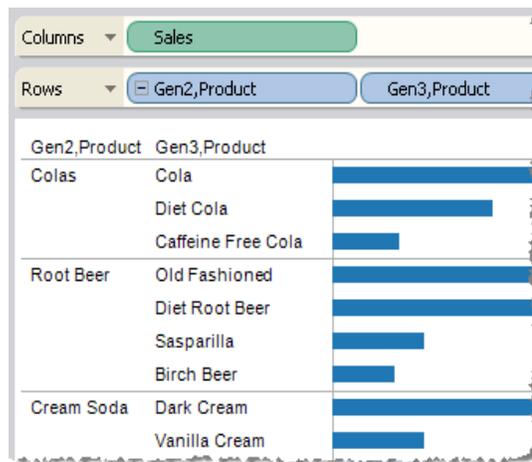


### Example – Sorting a Multidimensional Hierarchy

This example uses a multidimensional data source to sort the rows of a bar chart in order to determine which beverages have the highest sales. To create the view, follow the steps below.

- 1 Place the **Sales** measure on the **Columns** shelf and the **Gen2,Product** dimension on the **Rows** shelf.

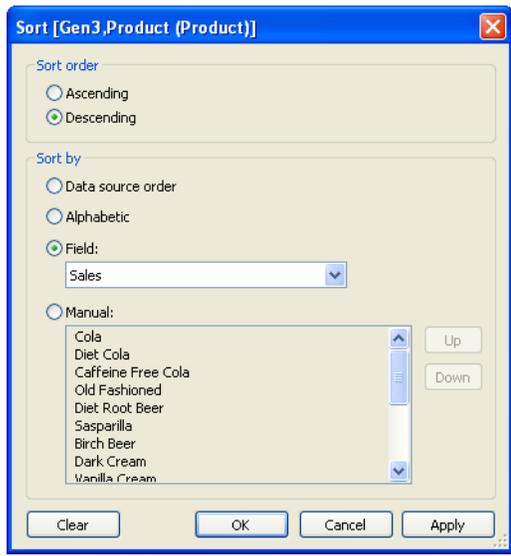
Drill down one level in the hierarchy to display **Gen3,Product**.





2 Sort **Gen3,Product** in ascending order by the **Sales** measure.

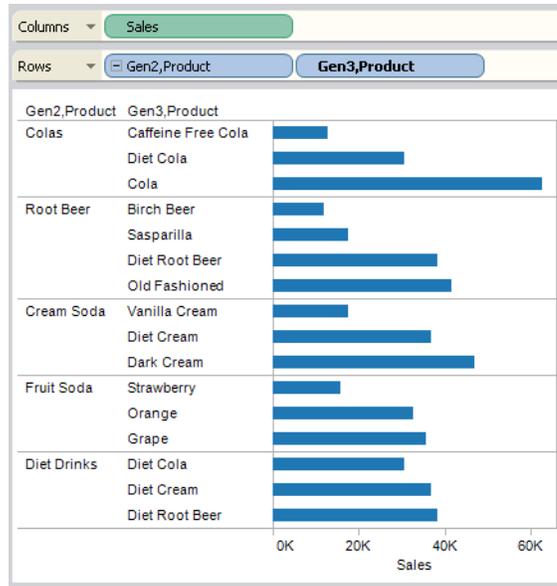
Right-click on **Gen3,Product** and select **Sort** from the field's context menu. In the Sort dialog box select **Ascending** as the Sort order and sort by the **Sales** field.



The view is shown below. Notice that the **Gen3,Product** members are sorted within each parent member. For example, Cola, Diet Cola, and Caffeine Free Cola are sorted only

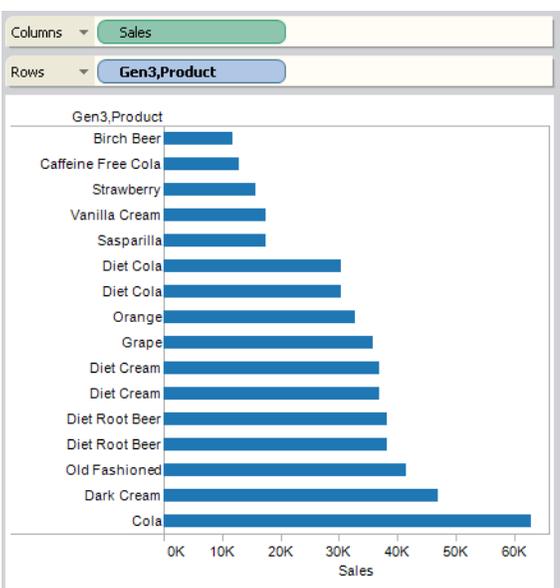


within the Colas level. Visual Explorer does not rearrange headers that appear before the sorted field.





- 3 If you want to order dimension members without regard to its parent, you should remove **Gen2,Product** from the **Rows** shelf. The sorted data are shown below.





## Manual Sorting

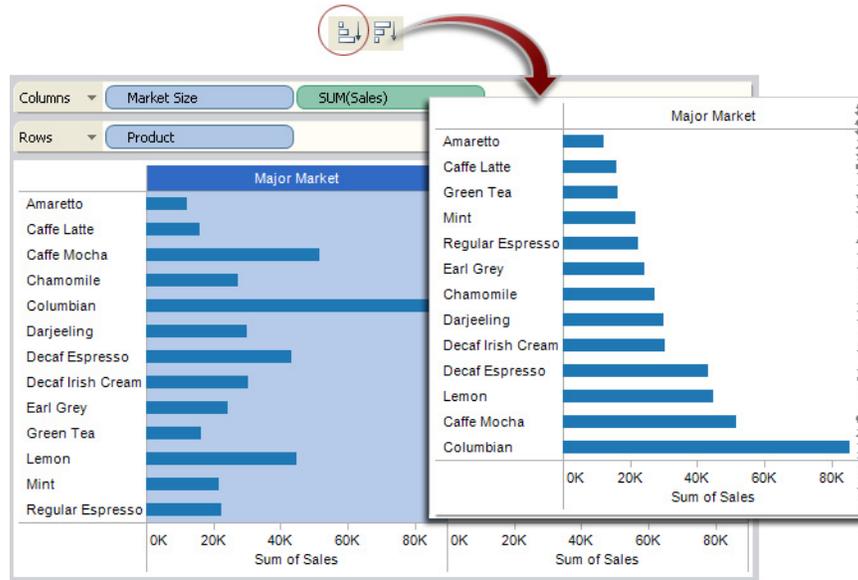
Manual sorting allows you to rearrange the order of dimension members in the table by dragging them in an ad-hoc fashion, giving precise control over how items appear next to one another in tables and in legends. It also gives you control over the order in which data is drawn on the screen. This control 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 11-37.

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.

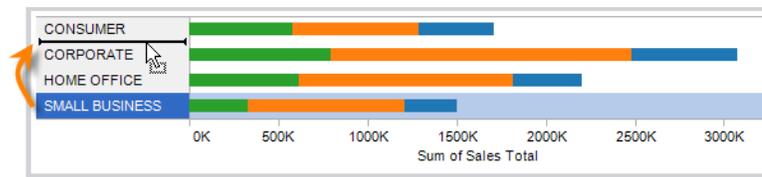


Using the quick sort buttons creates a manual sort which you can always modify using the sort dialog box. Right-click a sorted field (indicated with bold text) and select **Sort** to open the Sort dialog box.



## Sort by Drag and Drop

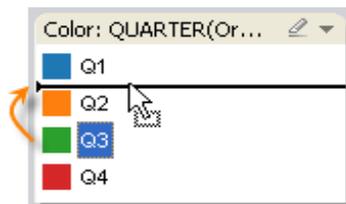
- 1 Select the dimension member you want to move. This can be any dimension member that appears in a row or column header of a table, or in a legend like the color legend.
- 2 Drag the member to the desired location within that row, column or legend.



## 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.

---

## Groups

A 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 Groups
- Editing an Existing Group
- Finding Members in the Groups Dialog Box

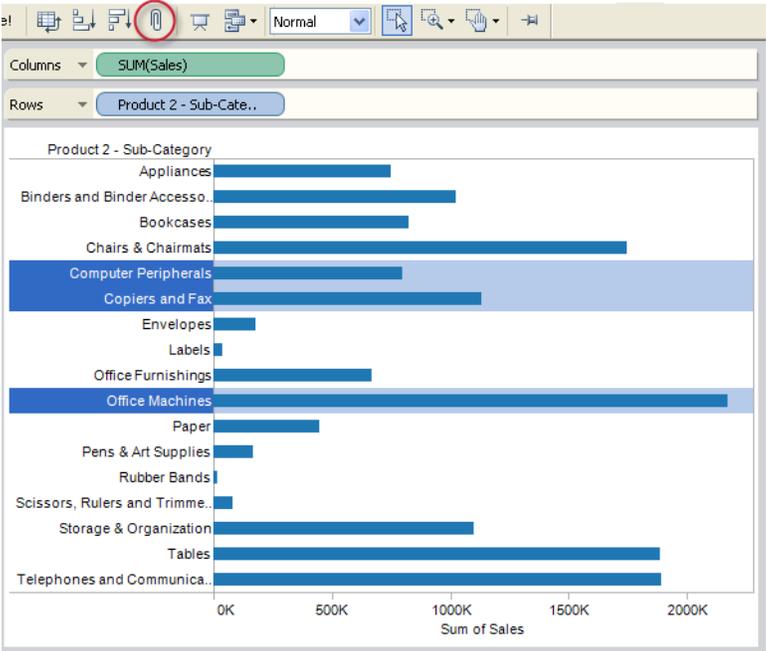
### Creating Groups

The most common way to create a group is through the group button on the toolbar. However, you can also create groups by right-clicking a dimension in the Data window and selecting **Create Group**.

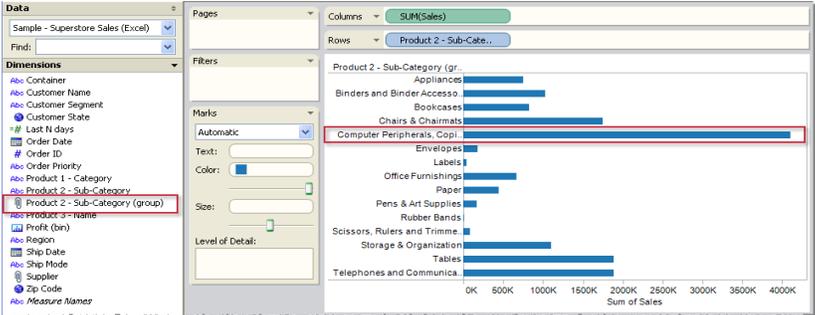
#### To create a group using the toolbar:

- 1 Hold the **CTRL** or **Shift** 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 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 17-22 to learn more.

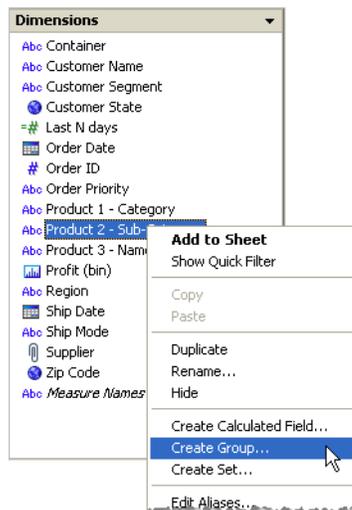
---

**Note** You can quickly un-group the dimension members by selecting the group in the view and clicking the Group button on the toolbar.

---

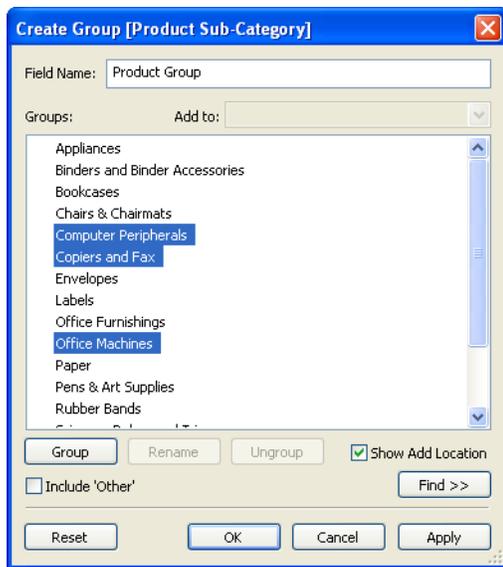
### To create groups from the Data window:

- 1 Right-click a dimension in the Data window and select **Create Group**.

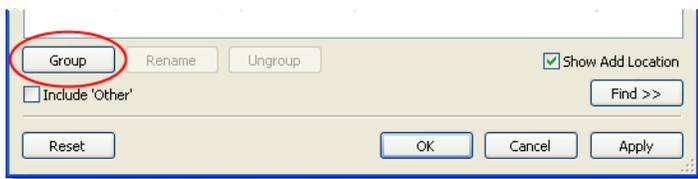




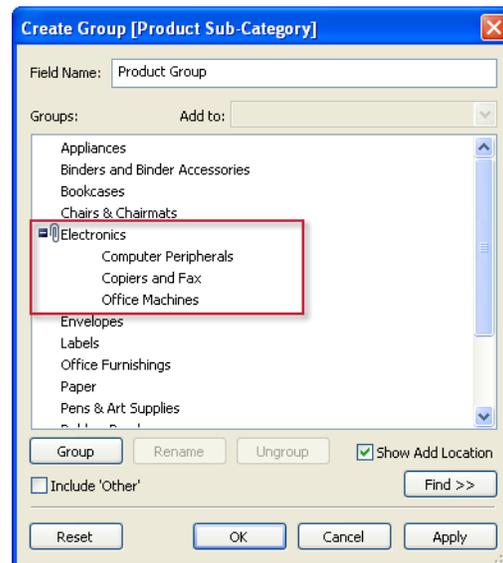
- 2 In the Create Group dialog box, select several members that you want to group. Hold the **CTRL** key on your keyboard to select multiple members.



- 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.



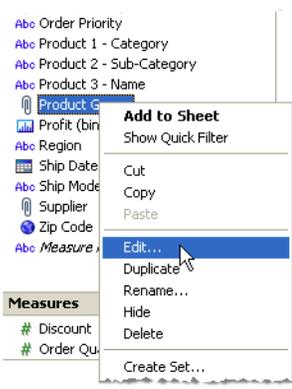
## Editing an Existing Group

After you have created a 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 Group dialog box.

**To add members to an existing group:**

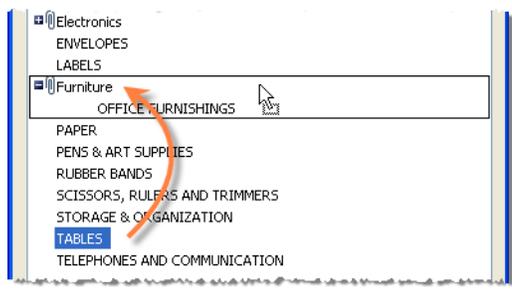


1 Right-click the grouped field in the Data window and select **Edit**.



2 In the Edit 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 dimension that has few members.

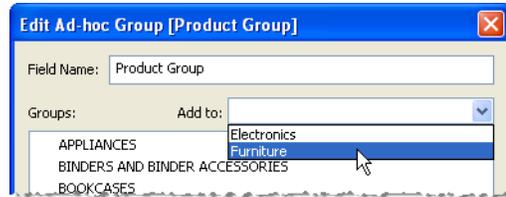


- Select one or more members, right-click and select **Add To**. In the subsequent dialog box, select the group you want to add the selected members to and click **OK**.





- 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 Group dialog box, select the grouped members and click the **Rename** button at the bottom of the dialog box.



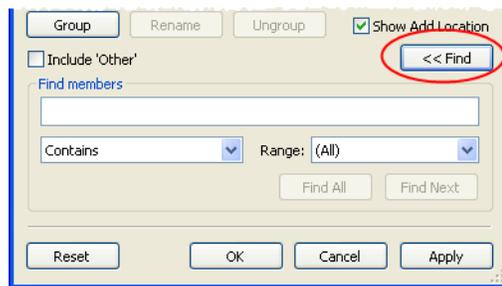
- 3 Type a new name and press **Enter** on your keyboard.
- 4 When finished, click **OK**.

## Finding Members in the Groups 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 17-38 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 17-42 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.

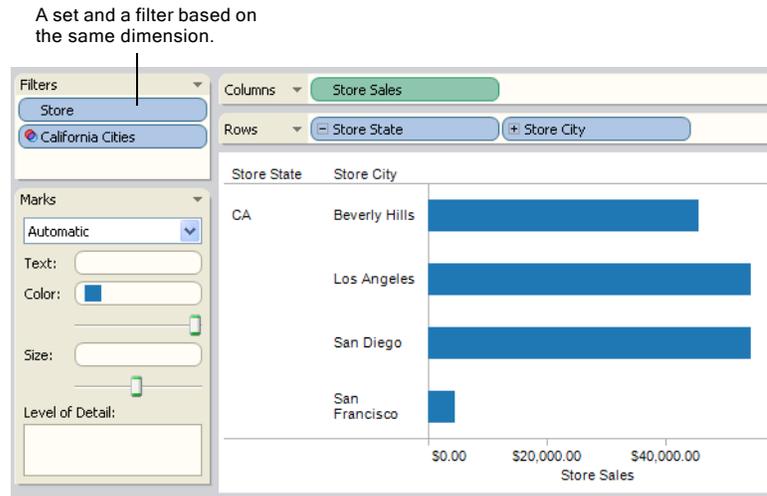


Additionally, sets are always treated as a filter. Therefore, when you place a set on a shelf, it is automatically placed on the **Filters** shelf as well.

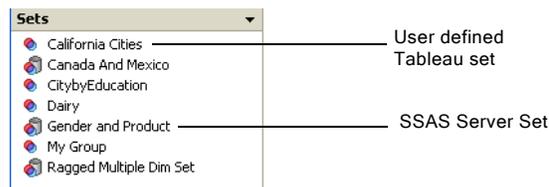
Note that if you use a filter and a set that are based on the same dimension, the result is the intersection of the filter and the set or its descendents. For example, the following view filters the **Store** hierarchy to include only the states and the cities shown below.



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.



**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.



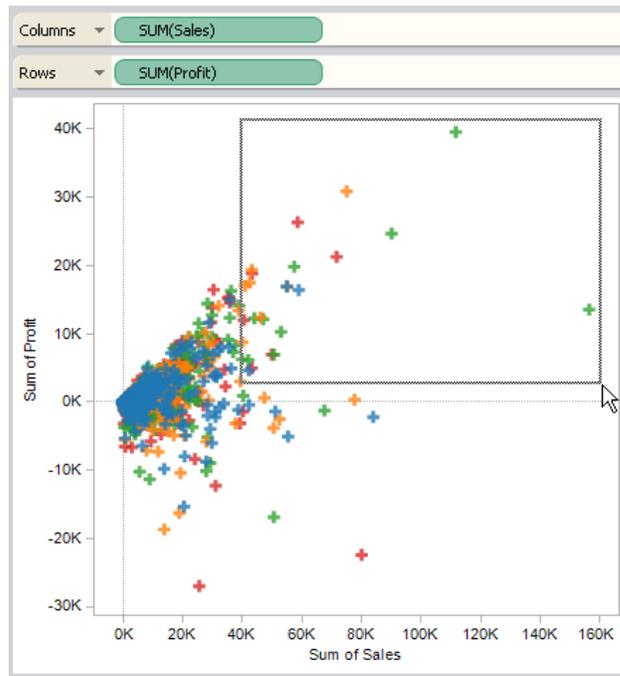
### **Create a Set by Selecting Marks**

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.

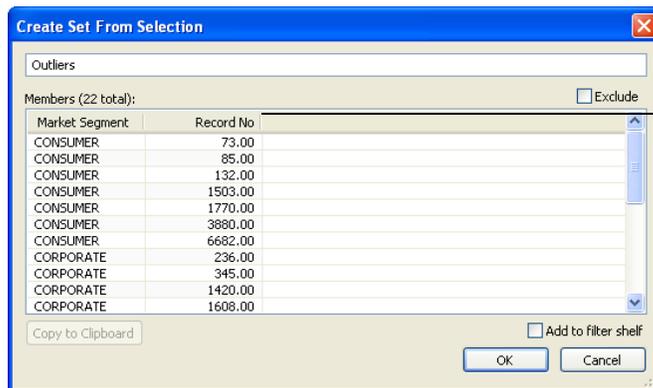
Optionally select the **Exclude** checkbox in the upper right corner if you want the set to contain all members except the ones you selected.



---

**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.

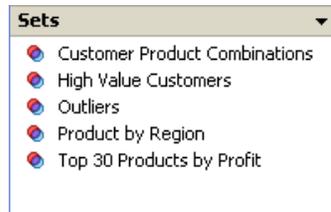
---



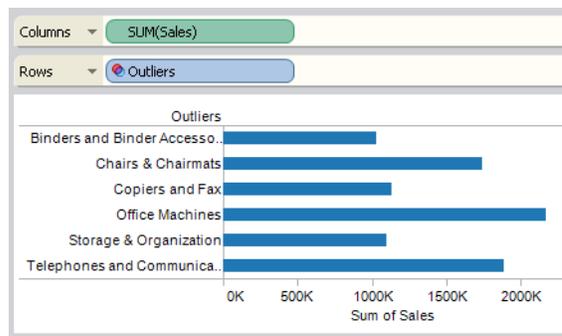
Click on a header to sort the column.  
Right-click on a header to remove the column.



Tableau displays the new set in the **Sets** area of the Data window.



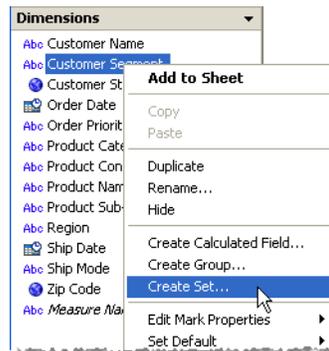
When you use the set in a data view, a header is created for each set member. As shown below, the header labels are given by the member names. For cubes, you can choose to display the fully qualified name by selecting the set in the Data window, and then selecting **Qualify Member Names** from the field's 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 “Filtering Dimensions” on page 16-7 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.

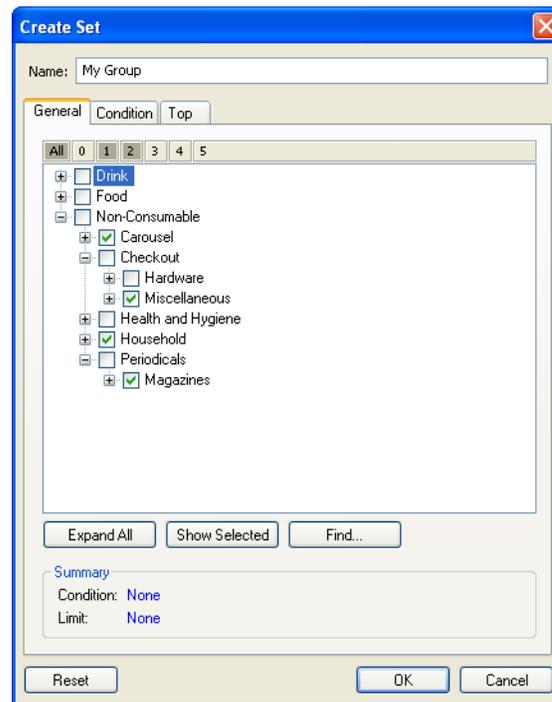


Tableau displays the new set in the **Sets** area of the Data window.





## 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 17-42 to learn more about this method.

Create the nest set by selecting multiple dimensions in the Data window and then selecting **Create Set** from the right-click context menu of a selected field.

---

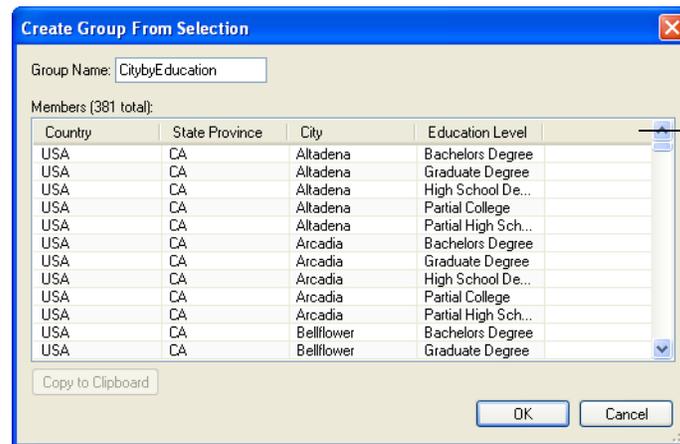
**Note** For multidimensional data sources, you must select levels from different hierarchies.

---

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.



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 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.



	Drink			
	Alcoholic Beverages	Beverages	Dairy	Baked Goods
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	28.35
Altadena, Partial College	7.06	17.43	7.33	12.92
Altadena, Partial High School	17.07	67.62	14.70	44.35
Arcadia, Bachelors Degree	12.83	41.73	3.68	25.42
Arcadia, Graduate Degree	6.71	5.85		2.54
Arcadia, High School Degree	25.55	28.21	13.31	29.05
Arcadia, Partial College	1.95	14.18		
Arcadia, Partial High School	8.67	54.95	3.03	43.90
Ballwin, Bachelors Degree	12.51	57.74	15.90	29.04

## 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**

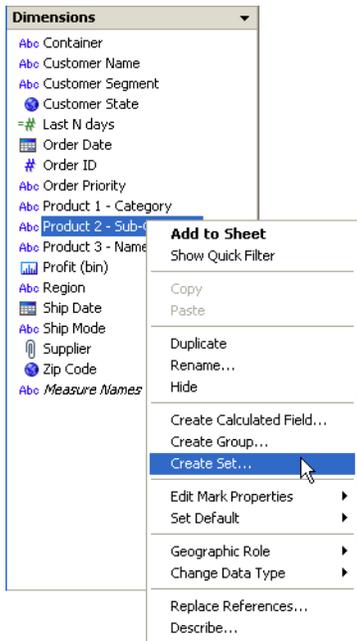
One reason to create a set is so you can easily work with just the dimension members that are of interest to you. For example, you might want to work with specific geographic regions, high-value customers, or one product line in your organization. To create such a set, select the relevant dimension members using any of the methods described in “How to Create a Set” on page 17-28.

In this example, you will create a subset of the Sample Superstore data source using the **Create Set** dialog box. Follow the steps below:

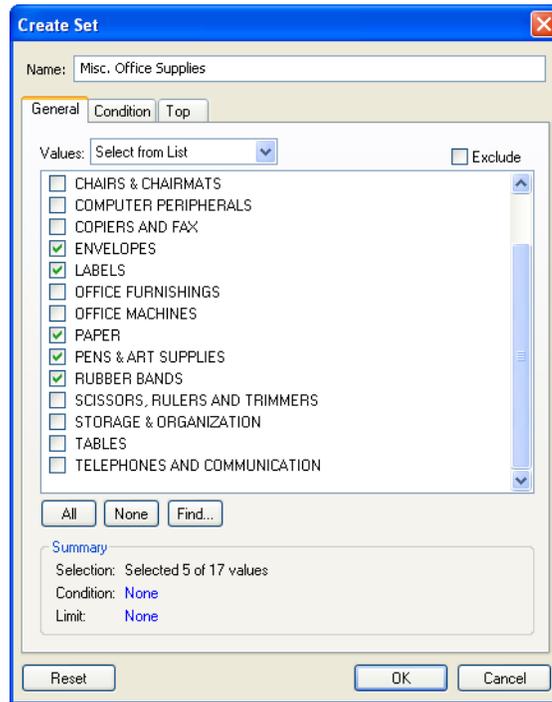


- 1 Select the dimension that will form the set.

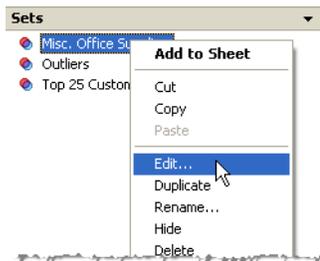
Right-click **Product 2 - Sub-Category** in the Data window, and select **Create Set**.



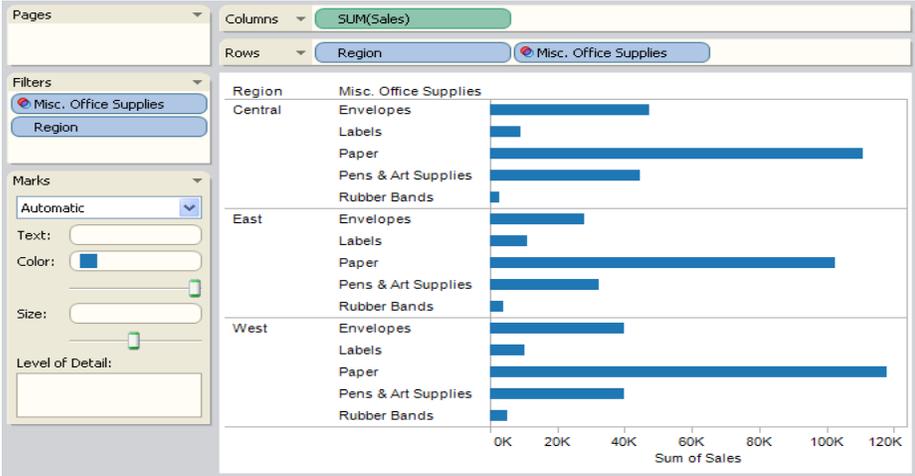
- 2 In the **Create Set** dialog box, specify the name of the set and select the dimension members that you want to include in the set. In this example, you are only interested in Envelopes, Labels, Paper, Pens and Art Supplies, and Rubber Bands.



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 1 - Category** 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 **Customer Segment** dimension on **Columns** shelf, placing the **Sales** measure on the **Rows** shelf, and color-encoding the data using the new set.

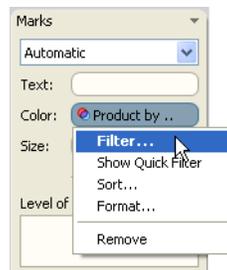


When you place the set on the **Color** shelf, Tableau separates the marks according to the members in the set, and assigns a unique color to each member. The color legend displays each member name and its color.



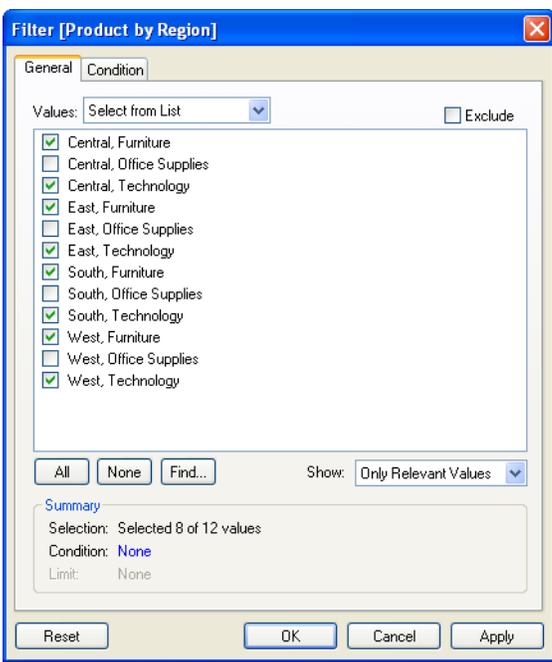
### 3 Filter the set.

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.



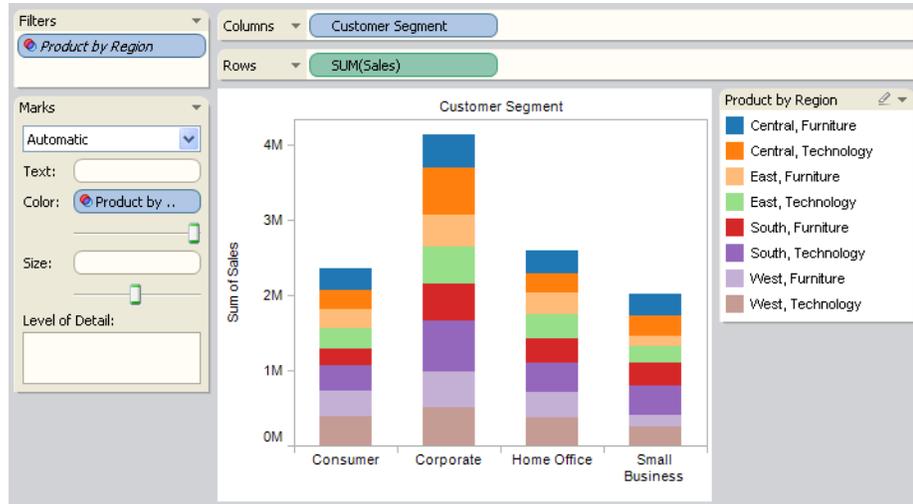


For this example, include only the Furniture and Technology products.



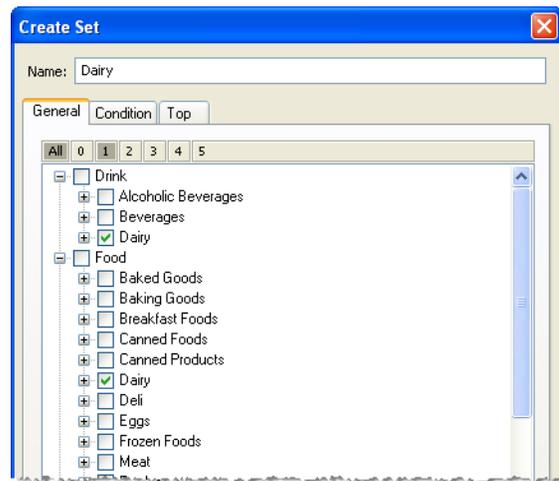


The final view is shown below. Note the name of the filtered set is italicized.

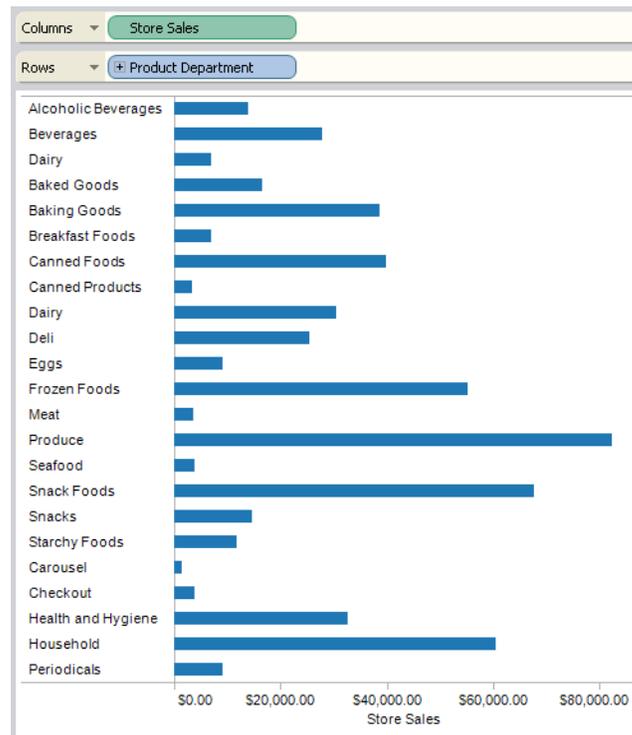


### Example – Hierarchical Sets and their Descendents

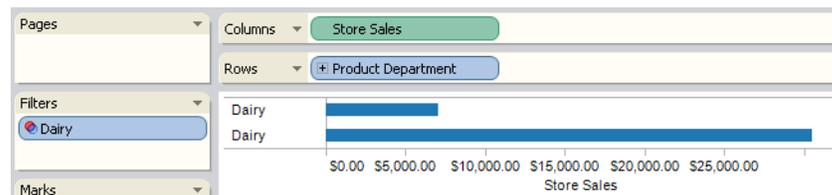
A hierarchical set filters data to the selected members and all of their descendents. For example, a set named **Dairy** is created from the **Product** hierarchy. As shown below, it includes only the Dairy product department.



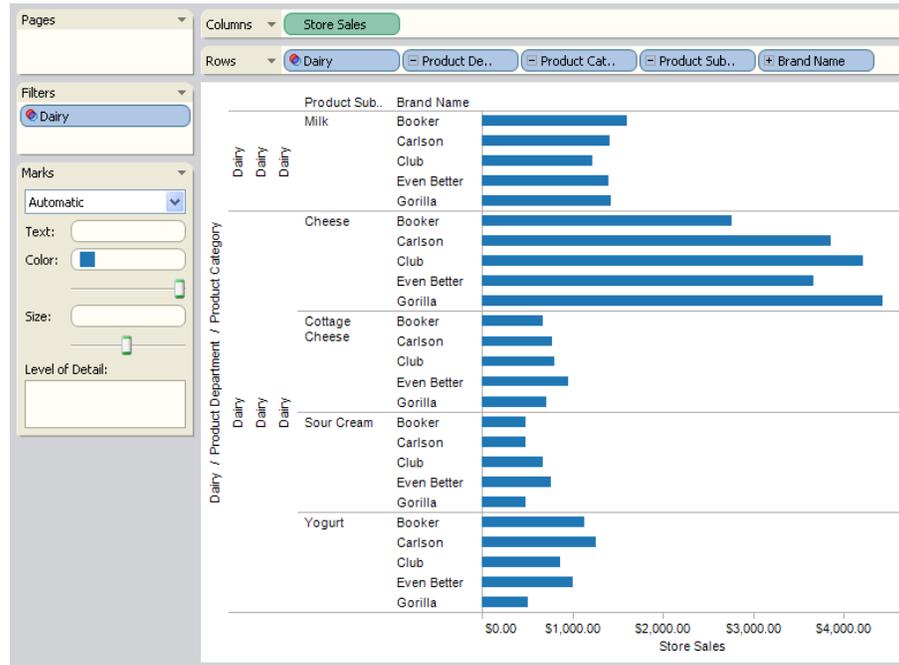
Consider the following view. The **Product Category** dimension is placed on the **Rows** shelf and the **Store Sales** measure is placed on the **Columns** shelf.



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

---

<b>Overview</b> . . . . .	<b>18-2</b>
<b>Changing Date Levels</b> . . . . .	<b>18-4</b>
<b>Fiscal Dates</b> . . . . .	<b>18-8</b>
<b>Perfect Pivoting with Dates</b> . . . . .	<b>18-10</b>
<b>Continuous Dates</b> . . . . .	<b>18-12</b>

---

## 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

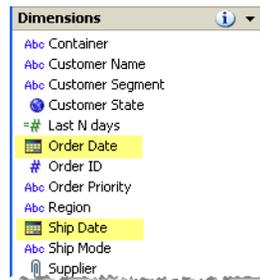
For multidimensional data sources, dates dimensions are usually organized into hierarchies that contain levels such as year, quarter, and month. In addition, some multidimensional data sources have time intelligence enabled, which enables different ways to look at these data levels such as Months by Year, Months by Quarter, Weekends, etc. These are represented as attributes of the hierarchy. These hierarchies and attributes 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 **Order Date** and **Ship 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.

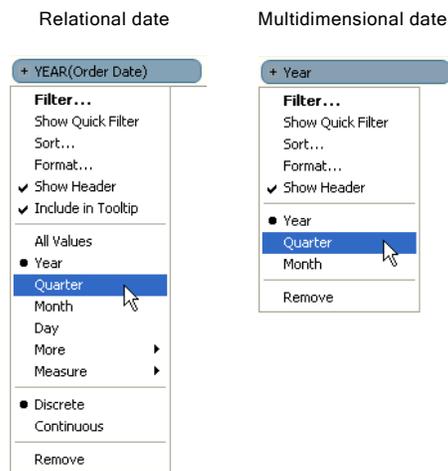
To learn more about working with dates and times refer to the following topics:

- 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 01/23/07. The year is 2007, 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**



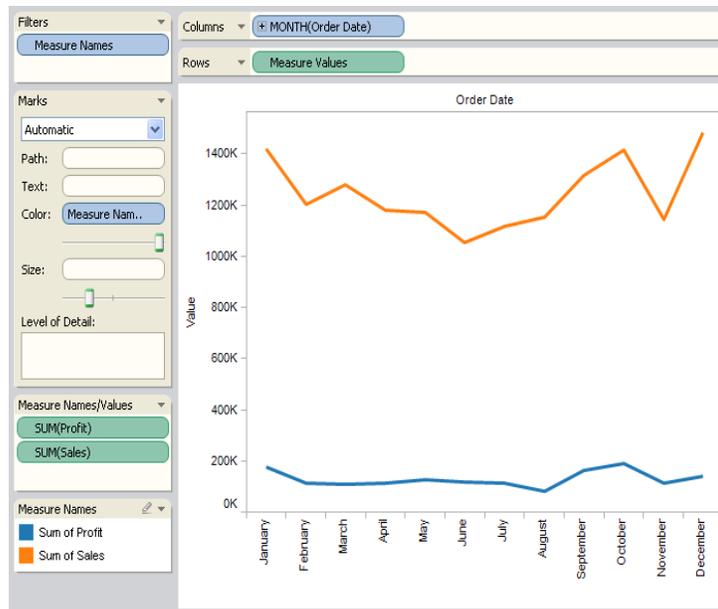
shelf multiple times in order to nest them and to drill down into them at varying levels of detail.

For example, the view shown below drills down into the year level to display the quarter level as well.

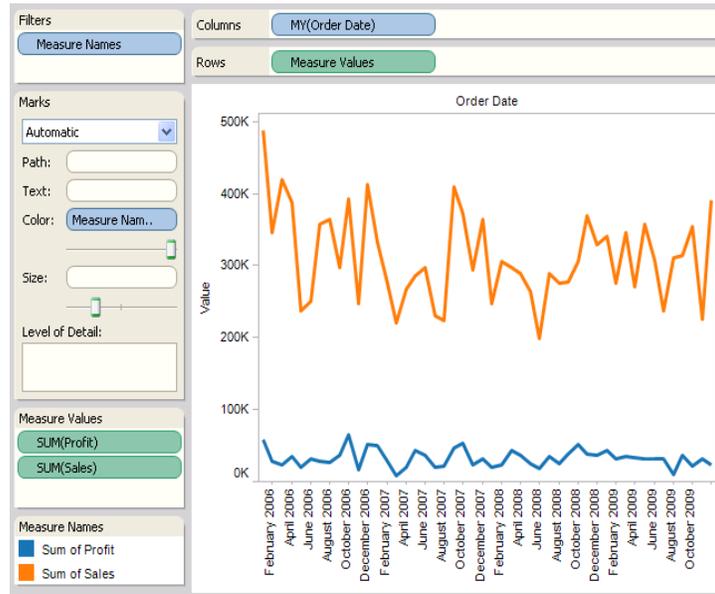




You can display the data by month by selecting **Month** from the date field's context menu. This displays the data for each month across all years.



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 15-1:**

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 15-1:**

<b>Date Level</b>	<b>When Converted to Fiscal</b>
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.

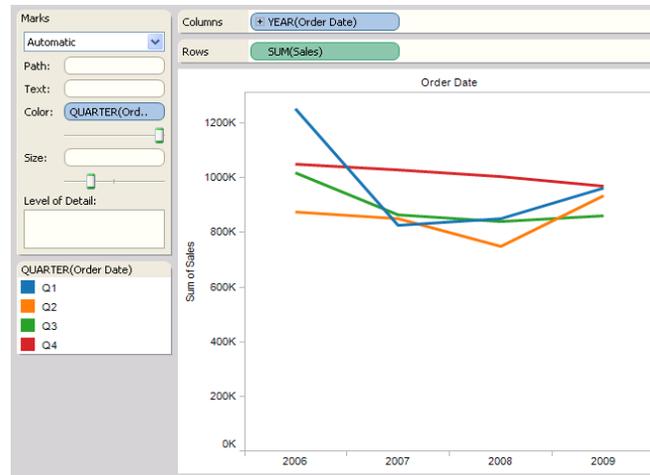
Fiscal year designations for any given date dimension are applied to all instances of the field in the Tableau workbook. Fiscal dates can only be applied to dimensions in a relational data source.

---

## Perfect Pivoting with Dates

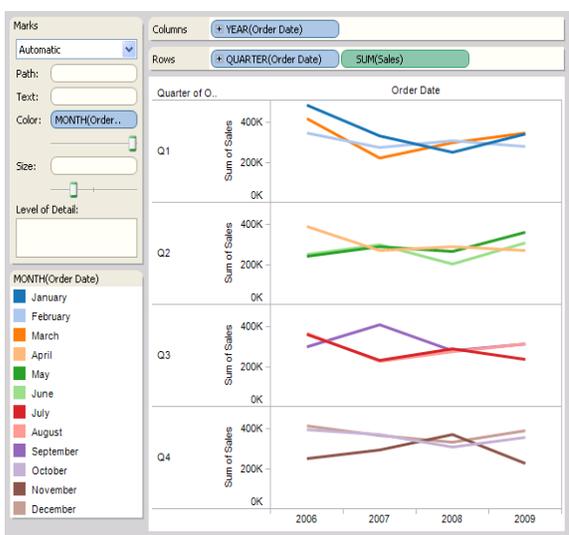
You can perfect pivot dates by placing different date levels on different worksheet shelves simultaneously. Place the date field on a variety of shelves and then select the desired date level from the fields' context menus.

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.

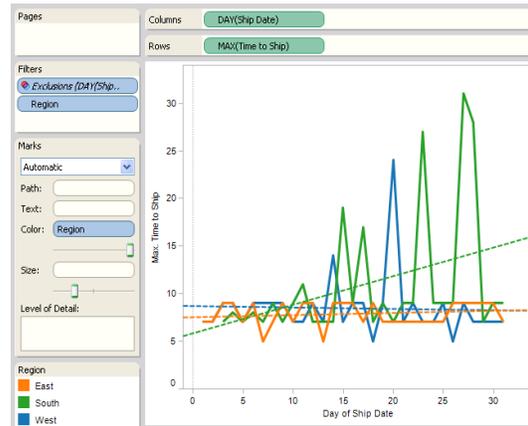


---

## 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 time to ship as a function of a continuous ship date and color-encoded by region. As you can see, the color of the **Ship 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>19-2</b>
<b>Navigation Tools</b> . . . . .	<b>19-3</b>
Select . . . . .	19-3
Zoom . . . . .	19-5
Pan . . . . .	19-7
<b>Undo and Redo</b> . . . . .	<b>19-8</b>
<b>Drop Lines</b> . . . . .	<b>19-9</b>
<b>Summary Card</b> . . . . .	<b>19-11</b>
<b>View Data (Drill-Through)</b> . . . . .	<b>19-12</b>
<b>Describing the View</b> . . . . .	<b>19-16</b>

---

## Overview

Once you have created a 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
- Summary Card
- View 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.



Each tool has a set of keyboard shortcuts that may make it easy to switch between tools and extend the functionality. Refer to “Navigation Tool Shortcuts” on page 37-4 to learn more.

The following navigation tools are available:

- Select
- Zoom
- Pan

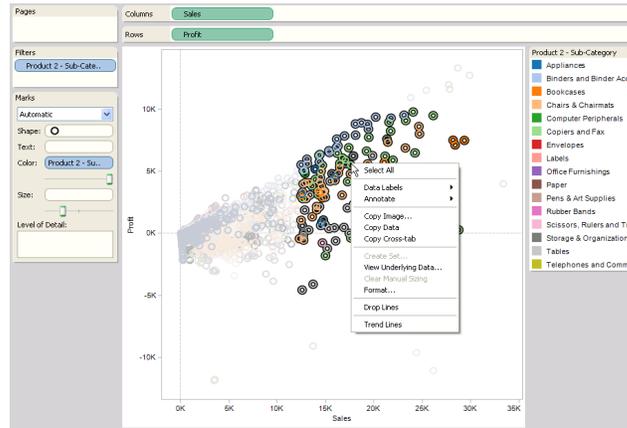
### Select

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.



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.





## 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. There are two ways to zoom: you can focus, which filters the view only show the marks of interest, or you can zoom, which increases the size of the view making it easier to read.

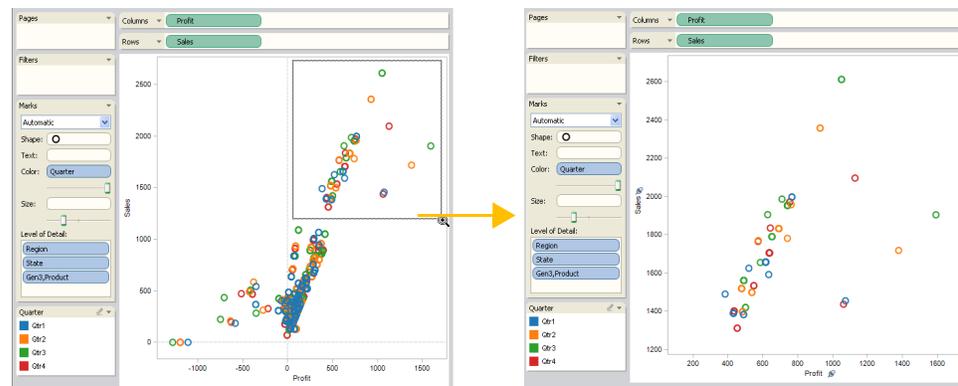
Use the zoom tool by clicking the **Zoom** button on the toolbar. The cursor changes to a magnifying glass .



### Focus Zoom

To focus zoom, click the area of the view you want to focus on. You can also click and drag to select an area to focus on. When you focus zoom, you can only select marks within a single pane. 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 **Clear Axis Range** from the context menu.

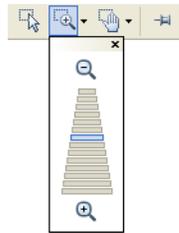


If you want to reset to the original view, use the **Undo** button  to undo each operation you performed on your view.



## Zoom

To zoom, hold the Alt key on the keyboard and click on the area of the view you want to zoom in on. You can also hold the Alt key and click and drag to zoom in on a specific area. To zoom out hold the Alt and Ctrl keys 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 incrementally zoom using the zoom drop-down list on the toolbar.



## Pan

You can move your view of a table up and down as well as left and right with the pan tool. There are two uses of panning. The first is when you have zoomed in on a view, particularly a map, and want to move the map around to see other marks of interest. The second is 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.

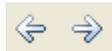


To pan around a view to see more marks of interest, simply click and drag the pan tool. If your view is really large with multiple panes, you can scroll from pane to pane by holding the Alt key on your keyboard while you click and drag the mouse.

---

## 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.

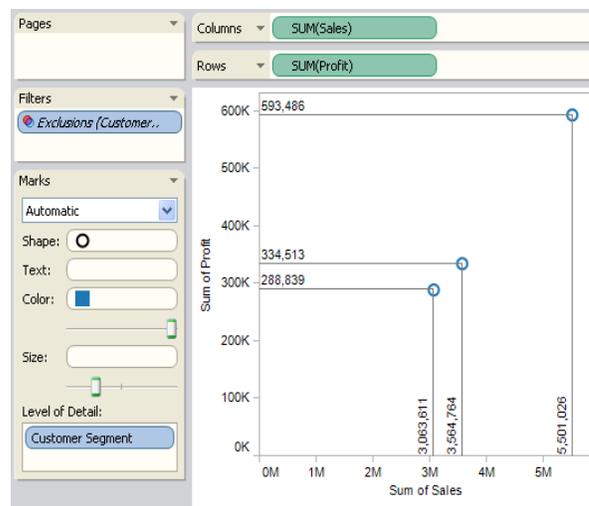


In this regard, every workbook behaves like a web browser. You can quickly return to a previous view. Or you can browse all the views of a data source that you have created. Tableau saves the undo/redo history across all worksheets until you exit. The history is not saved between sessions.

---

## 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.

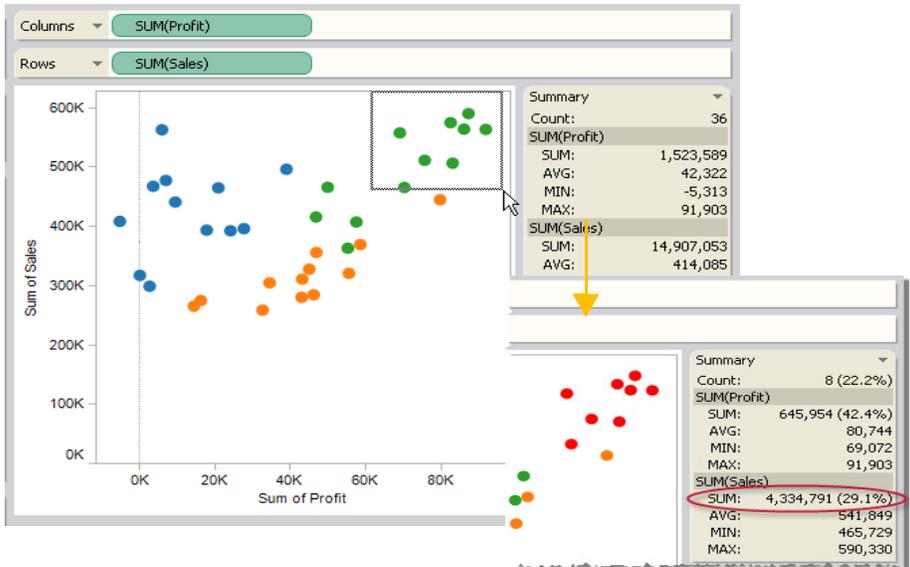


- 2 In the Drop Lines dialog box select an axis to draw the line to, whether to always show the drop lines, and whether to show labels.
- 3 When finished click **OK**.

# 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 \$4,334,791 in sales with a minimum sale of \$465,729.



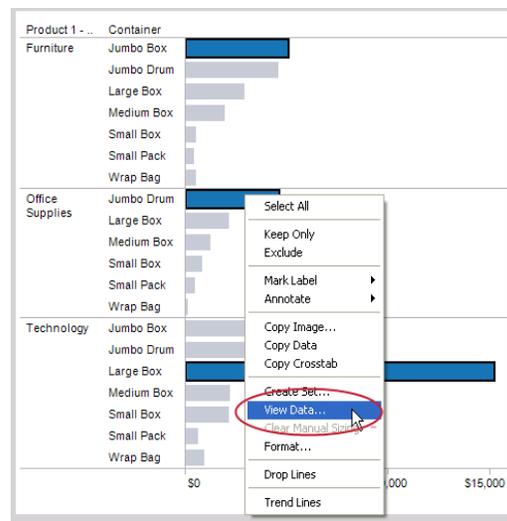
---

## View Data (Drill-Through)

The View Data command lets you display the values for each row in the data source that compose the marks. It also shows you the summary data based on the aggregations in the view. You might want to do view data 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 Data command works with all relational and multi-dimensional databases except Oracle Essbase databases. While you can view 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 see. For more information about drill-through with a Microsoft Analysis Services database refer to “View Data (Microsoft Analysis Services)” on page 19-14.

The view shown below shows the average order quantity for two product dimensions as a bar chart. Suppose you want to view the data for the largest marks in each pane. To do this, select the marks of interest, right-click in the table, and select **View Data** on the context menu. Alternatively, you can select the **Analysis > View 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 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.

## Underlying Data

The underlying data for the selected marks are displayed on the **Underlying** tab in the View Data dialog box. Notice that the number of rows that compose the underlying data is shown in the upper right of the dialog box.

Click a heading to sort the data.

Select to show all data source fields.

Customer Name	Customer Seg...	Customer State	Last N days	Order Date	Order ID	Order Priority
Liz Pelletier	Corporate	Ohio	567	7/15/2008	32.00	High
Emily Phan	Consumer	Maryland	967	6/11/2007	132.00	Medium
Liz Willingham	Home Office	North Carolina	786	12/29/2007	298.00	Critical
Betsy French	Consumer	Michigan	123	10/2/2009	293.00	High
Ricardo Block	Corporate	Alabama	866	9/20/2007	358.00	High
Carli Sayre	Corporate	Idaho	1141	12/19/2006	359.00	Medium
Sonia Cooley	Corporate	Connecticut	336	3/3/2009	384.00	Low
Dave Halsten	Corporate	Texas	63	12/1/2009	514.00	High
Tamara Chand	Consumer	Delaware	1106	1/23/2007	640.00	High
Sanjit Jacobs	Consumer	Idaho	375	6/3/2007	898.00	High
Mike Villorini	Small Business	Arizona	1368	5/6/2006	962.00	Low
Jas O'Carroll	Corporate	Connecticut	629	5/14/2008	965.00	Low
Maureen Gastline...	Consumer	California	226	6/21/2009	967.00	Medium
Halle Redmond	Home Office	South Carolina	244	6/3/2009	1027.00	Medium
Sylvia Foulston	Home Office	New Jersey	363	2/14/2009	1154.00	Critical
Don Weiss	Corporate	South Carolina	246	6/1/2009	1285.00	Critical
Don Weiss	Corporate	South Carolina	246	6/1/2009	1285.00	Critical
Ritza Hightower	Corporate	Connecticut	97	10/28/2009	1345.00	Low
Theone Picoeoner	Corporate	Mass	229	6/18/2009	1346.00	High

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 selected. Clear this options to only show the fields placed on shelves (or fields referenced by a calculation placed on a shelf) for the current worksheet.

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 clicking 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 selecting View Data. This option is not available with a Microsoft Analysis Services database; you must select a single mark to view data for.

## Summary Data

The summarized data is shown on the Summary tab. The summarized data is a text table of the aggregated data for only the fields shown in the view.

Product 1 - Cat...	Container	Avg. Sales
Technology	Large Box	\$15,212
Furniture	Jumbo Box	\$5,075
Office Supplies	Jumbo Drum	\$4,602

## View Data (Microsoft Analysis Services)

View Data with a Microsoft Analysis Services database works 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 see. That means that when you try to view 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 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

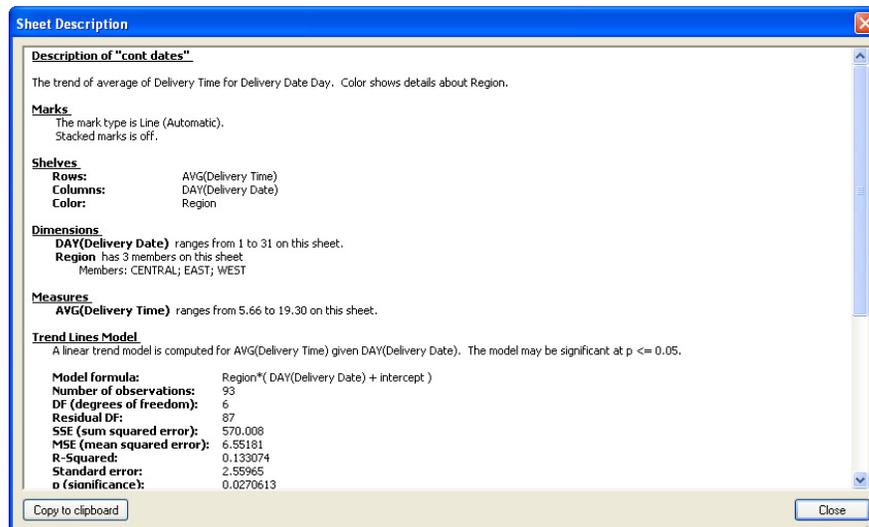
Occasionally you may want to succinctly summarize an analysis you have completed on a worksheet. You might then want to remind yourself of what it shows (the filters that are applied, etc.), and finally, you may want to share a summary of the analysis with someone else.

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 25-19 to learn more about the terms used to describe the model.

---



# Actions

---

<b>Overview</b> . . . . .	<b>20-2</b>
<b>Filter Actions</b> . . . . .	<b>20-3</b>
<b>Highlight Actions</b> . . . . .	<b>20-8</b>
Selecting Marks to Highlight . . . . .	20-8
Color Legend Highlighting . . . . .	20-9
Highlight Toolbar Button . . . . .	20-11
Creating Advanced Highlight Actions . . . . .	20-13
<b>URL Actions</b> . . . . .	<b>20-16</b>
<b>Running Actions</b> . . . . .	<b>20-20</b>
<b>Actions and Dashboards</b> . . . . .	<b>20-22</b>
<b>Using Field Values in Actions</b> . . . . .	<b>20-24</b>
Using Field Values in URLs . . . . .	20-24
Using Field Values in Action Names . . . . .	20-25

---

## Overview

Tableau allows you to add context and interactivity to your data using actions. Link to web pages, files, and other Tableau worksheets directly from your analytical results. Use the data in one view to filter data in another as you create guided analytical stories. Finally, call attention to specific results using highlighting.

For example, in a dashboard showing home sales by neighborhood you could use actions to help you quickly see relevant information for a selected neighborhood. Select a neighborhood in one view which then highlights the related houses in a map view, filters a list of the houses sold, and opens a webpage showing census data for the neighborhood.

There are three kinds of actions in Tableau: Filter, Highlight, and URL actions. This section discusses the following topics:

- Filter Actions
- Highlight Actions
- URL Actions
- Running Actions
- Actions and Dashboards
- Using Field Values in Actions

---

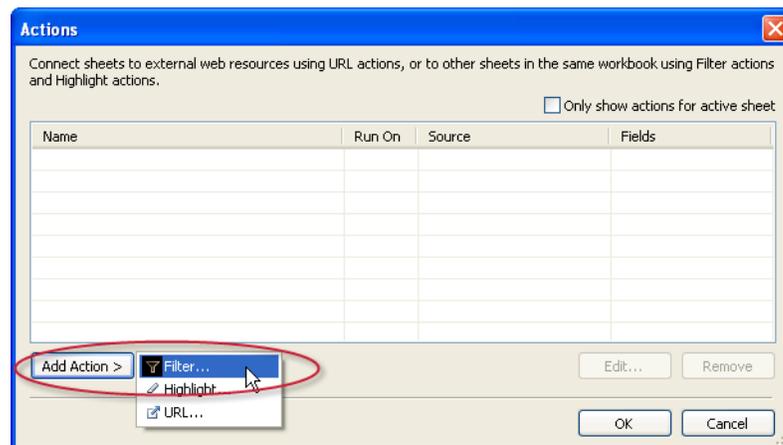
## Filter Actions

Filter actions are a way to send information between worksheets. Typically a filter action 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 filter action 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 filter action to show comparable houses can be defined by selecting a destination worksheet and defining filters on sales price and square footage.

Filter actions work by sending the data values of the relevant source fields as filters to the destination sheet. If you launch the filter action 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 create a filter action:

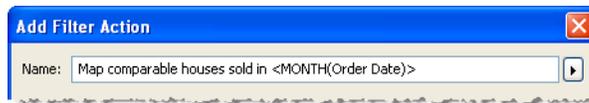
- 1 Select **Edit > Actions**.
- 2 In the Actions dialog box, click **Add Action** and then select **Filter**.



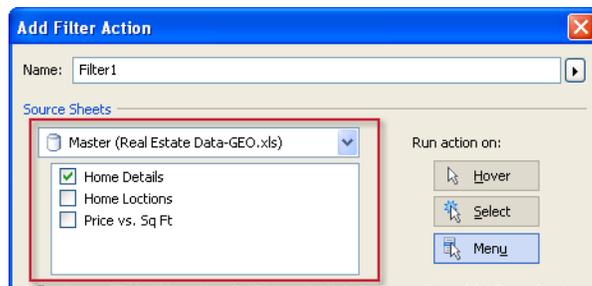


- 3 In the subsequent dialog box specify a name for the Action.

Use a name that defines the action. If you choose to run the action using the menu the name is the option that shows on the menu. 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 Action Names” on page 20-25.



- 4 Use the drop-down list to select a source sheet or data source. When you select a data source or dashboard sheet you can further refine by selecting the individual sheets you want to launch the action from.





- 5 Then select how you want to launch the action. Select one of the following options:
- **Hover** - rest the pointer over a mark in the view to run the action. This option works well for highlight and filter actions within a dashboard.
  - **Select** - click on a mark in the view to run the action. This option works well for all types of actions.
  - **Menu** - right-click a selected mark in the view and then select an option on a the context menu. This option works well for filter and URL actions.



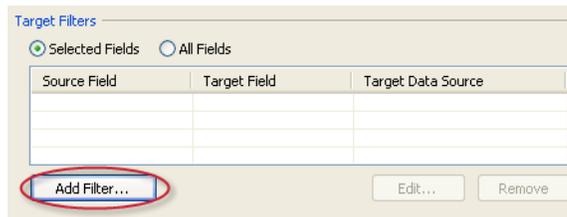
- 6 Use the second drop-down list to select a target sheet. When you select a dashboard sheet you can further refine the target by selecting one or more sheets within the dashboard.



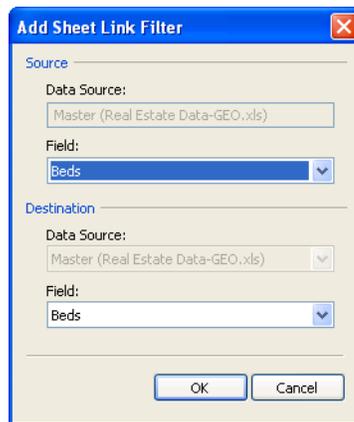
- 7 Specify what do do when the select is cleared in the view. You can select from the following options:
- **Leave the filter** - leaves the filter on the target sheets. The target views in the dashboard will show the filtered results.
  - **Show all values** - changes the filter to include all values.
  - **Exclude all values** - changes the filter to exclude all values. This option is useful when you are building dashboards that only show some sheets if a value in another sheet is selected.



- 8 Setup one or more filters to specify the data that you want to show on the target sheets. You can filter on **All Fields** or define filters on **Selected Fields**.
- 9 If you are defining filters for specific fields click **Add Filter**.



- 10 In the Add Filter dialog box select a source and target data sources and fields. When you run the action from a specific mark on the source sheet, a filter is added to the target sheet that only includes values for the target field that match the values of the source field. In the comparable houses sheet link example, the Source Field is Beds and the Target Field is Beds. That means when you launch the sheet link for a house that has 3 bedrooms, the destination worksheet will only show houses that also have 3 bedrooms.



- 11 When finished, click **OK** three times to close the dialog boxes and return to the view. 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 as the source sheet. Moreover, the source field names must match the destination field names.

---

**Note** The fields available in the Target 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.

---



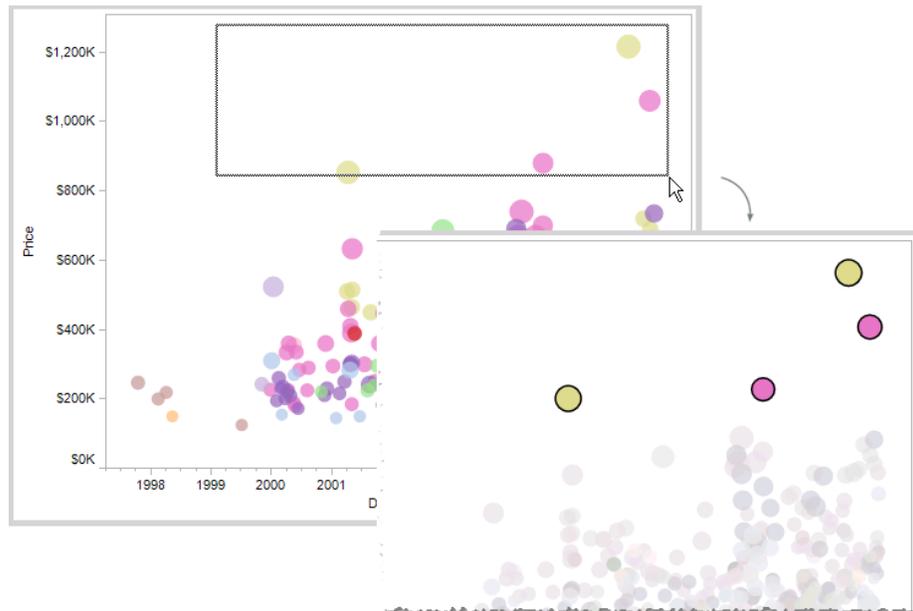
## Highlight Actions

Highlight actions allow you to call attention to marks of interest by coloring select marks and dimming all others. You can highlight marks in the view by selecting the marks you want to highlight, use the color legend to select related marks, or create an advanced highlight action.

### Selecting Marks to Highlight

When you select a mark in the view all other marks are dimmed to draw attention to the selection. Selection is saved with the workbook and can be included when publishing. The simplest way to add highlighting to a view is to select the marks you want to highlight.

You can select multiple marks by holding down the Ctrl key on your keyboard while you select each mark. You can also click and drag the pointer to select all marks in a specific area of the view.





## 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 order quantity and profit for several products. The view on the left uses the normal color legend, all marks are colored based on their shipping mode. The view on the right uses legend highlighting to call out the products that were delivered via Delivery Truck.

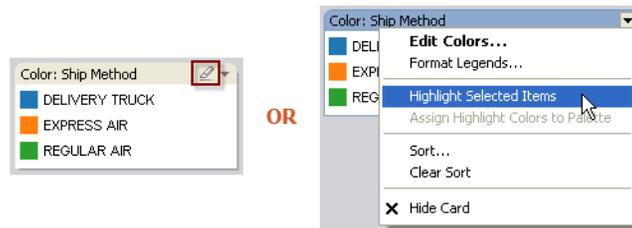


You can easily switch between legend highlighting and normal modes using the color legend card menu. Then, if you like how a view is highlighted, you can assign the highlight colors to the color palette. The old colors are replaced with the highlight colors.



### 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.

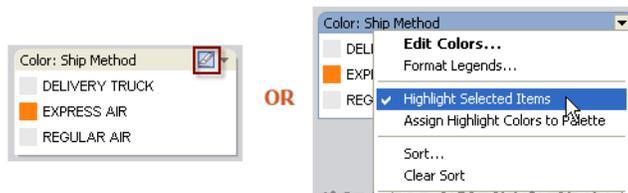


- 2 Select an item in the color legend.

Once legend highlighting is turned on, you can quickly focus on specific data in the view by selecting different items in the color legend. When color legend highlighting is turned on a Highlight Action is created and can be modified in the Actions dialog box.

### 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.



When you turn color legend highlighting off the action is removed from the Actions dialog box.



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.

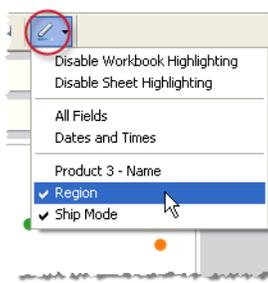
**To assign the highlight colors to the color palette:**

- Select **Assign Highlight Colors to Palette** on the color legend card menu.

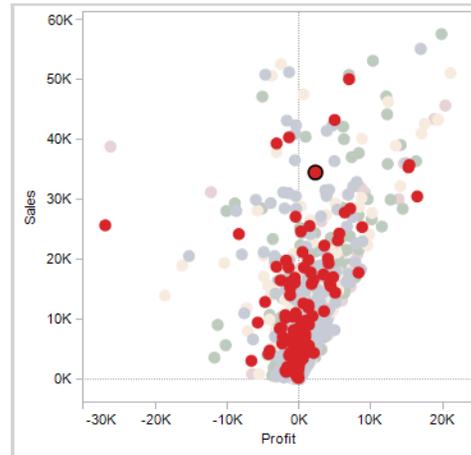


## Highlight Toolbar Button

Another way to add a highlight action is using the highlight button in the toolbar. Similar to the color legend highlighting, the toolbar button lets you highlight a collection of related marks in the view. To turn on highlighting, select the fields you want to use for highlighting on the toolbar menu. Then select a mark in the view to see the related data.



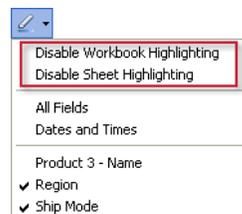
For example, the view below shows sales vs. profit by region. When a mark is selected, all other marks from that region that were shipped using the selected ship mode are highlighted. In this case you can quickly see all products from the Wester region that were shipped via Delivery truck.



The toolbar menu also lets you highlight on All fields or Dates & Times. All fields will consider all fields when determining matching records. Dates & Times considers all date and time fields.

When you use the Highlight toolbar button an action is created in the Actions dialog box. You can modify the action to create more advanced highlighting behavior.

Finally, you can use the toolbar button to disable highlighting across the entire workbook or for just the active sheet.



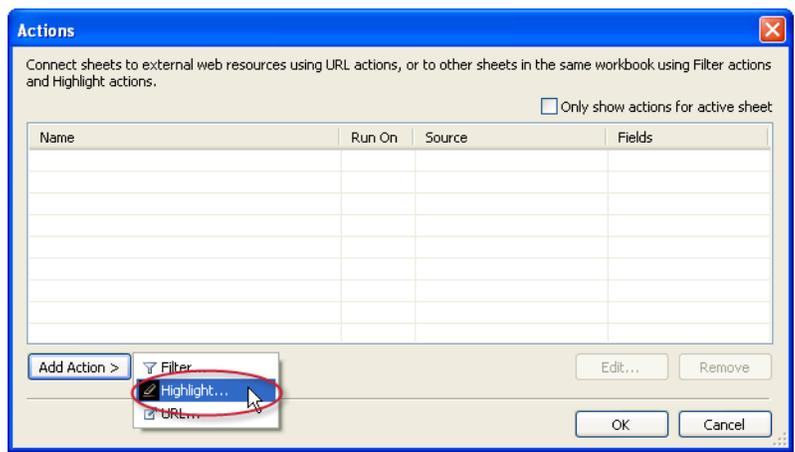


## Creating Advanced Highlight Actions

You can define more advanced highlight actions using the Actions dialog. There you can specify source and target sheets along and the fields you want to use for highlighting. Follow the steps below to create a Highlight Action.

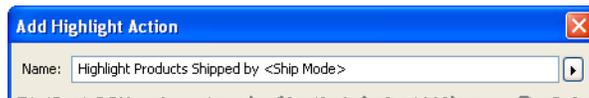
**To create a highlight action:**

- 1 Select **Edit > Actions**.
- 2 In the Actions dialog box click the **Add Action** button and then select **Highlight**.

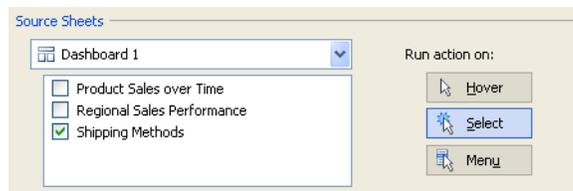




- 3 Give the action a name that will identify it in the Actions dialog. Try to make it descriptive. For example, Highlight Products Shipped by Delivery Truck. You can use variables in the name that will be filled in based on the values of the selected field. Refer to “Using Field Values in Action Names” on page 20-25.



- 4 Use the drop-down list to select the Source sheet or data source. If you select a data source or a dashboard sheet you can further select individual sheets within them.



- 5 Select how you want to launch the action. You can select from the following options:
  - **Hover** - rest the pointer over a mark in the view to run the action. This option works well for highlight and filter actions within a dashboard.
  - **Select** - click on a mark in the view to run the action. This option works well for all types of actions.
  - **Menu** - right-click a selected mark in the view and then select an option on a the context menu. This option works well for filter and URL actions.

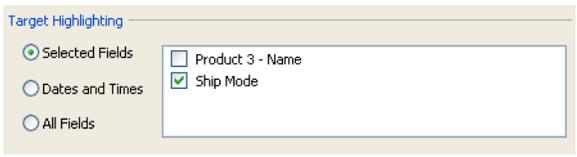




- 6 Select a Target sheet. If you select a dashboard you can further select individual sheets within the dashboard.



- 7 Select the fields you want to use for highlighting. Select from the following options:
  - **Selected Fields** - marks in the target sheet are highlighted based on select fields. For example, highlighting using the Ship Mode field will result in an action that highlights all marks in the target sheet that have the same ship mode as the selected mark in the source sheet.
  - **Dates and Times** - marks in the target sheet are highlighted when their date and time match those of the marks selected in the source sheet. All dates and time fields are considered when determining a match.
  - **All Fields** - marks in the target sheet are highlighted when they match the marks selected in the source sheet. All fields are considered when determining a match.



- 8 When finished, click **OK** twice to close the dialog boxes and return to the view.

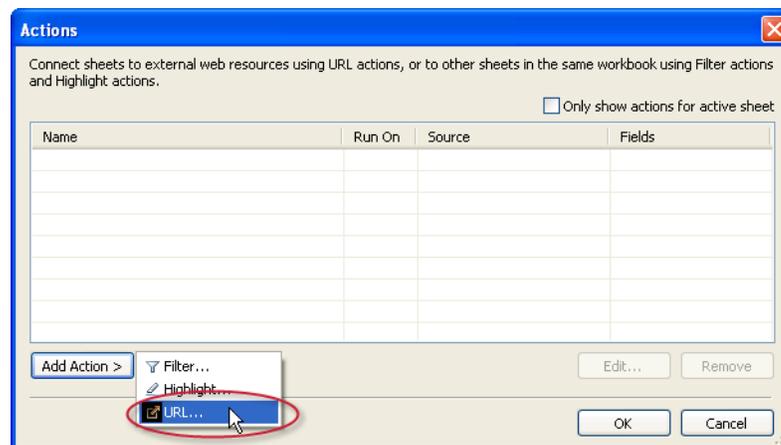
---

## URL Actions

A URL action is a hyperlink that points to a Web page, file, or other web-based resource outside of Tableau. You can use URL actions to link to more information about your data that may be hosted outside of your data source. To make the link relevant to your data, you can substitute field values of a selection into the URL as parameters.

### To add a Hyperlink:

- 1 Select **Edit > Actions**.
- 2 In the Actions dialog box, click **Add Action** and then select **URL**.

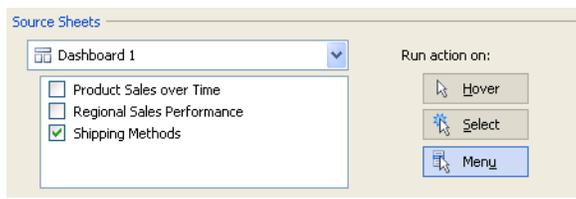




- 3 In the subsequent dialog box, specify a name for the link. Make the name descriptive of the action. If you choose to run the action using the menu the name is the option that shows on the menu. For example, when linking to more product details, the name could be “Show More Details for Binder Clips.” 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 Action Names” on page 20-25.



- 4 Use the drop-down list to select a source sheet or data source. If you select a data source or dashboard you can select individual sheets within it.



- 5 Select the fields you want to use for highlighting. Select from the following options:
  - **Hover** - rest the pointer over a mark in the view to run the action. This option works well for highlight and filter actions within a dashboard.
  - **Select** - click on a mark in the view to run the action. This option works well for all types of actions.
  - **Menu** - right-click a selected mark in the view and then select an option on a the context menu. This option works well for filter and URL actions.





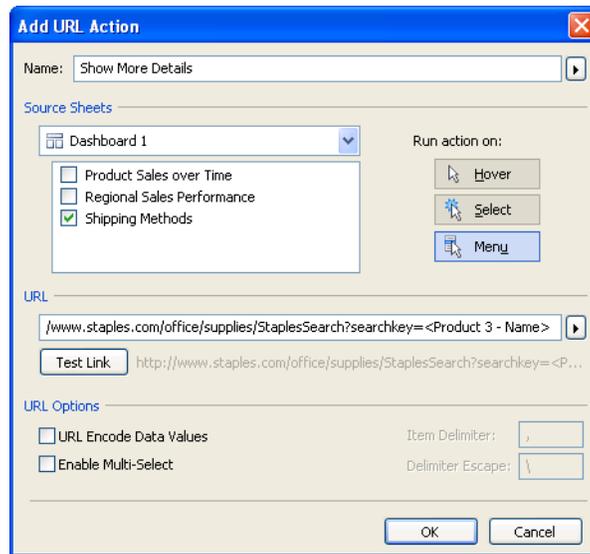
- 6 Specify the URL. You can use any URL that your browser can recognize including web pages, ftp resources, and files.

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 website. To learn more about using fields as parameters in your URL refer to “Using Field Values in URLs” on page 20-24.

- 7 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 webpage 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 webpage. 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.



8 When finished, click **OK** twice to close the dialog boxes and return to the view.

---

**Note** URL actions can also point to a web page object in a dashboard. Refer to “Actions and Dashboards” on page 20-22 to learn more about how actions work with dashboards.

---

---

## Running Actions

Depending on how the action is created you can run an action using one of the following three methods:

- **Hover** - rest the pointer over a mark in the view to run the action. This option works well for highlight and filter actions within a dashboard.
- **Select** - click on a mark in the view to run the action. This option works well for all types of actions.
- **Menu** - right-click a selected mark in the view and then select an option on a the context menu. This option works well for filter and URL actions.

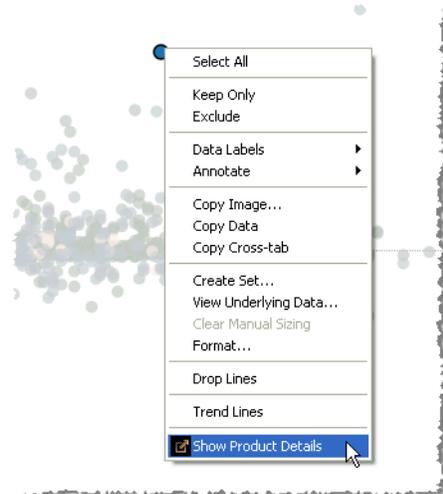
**Hover**



**Select**



**Menu**



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



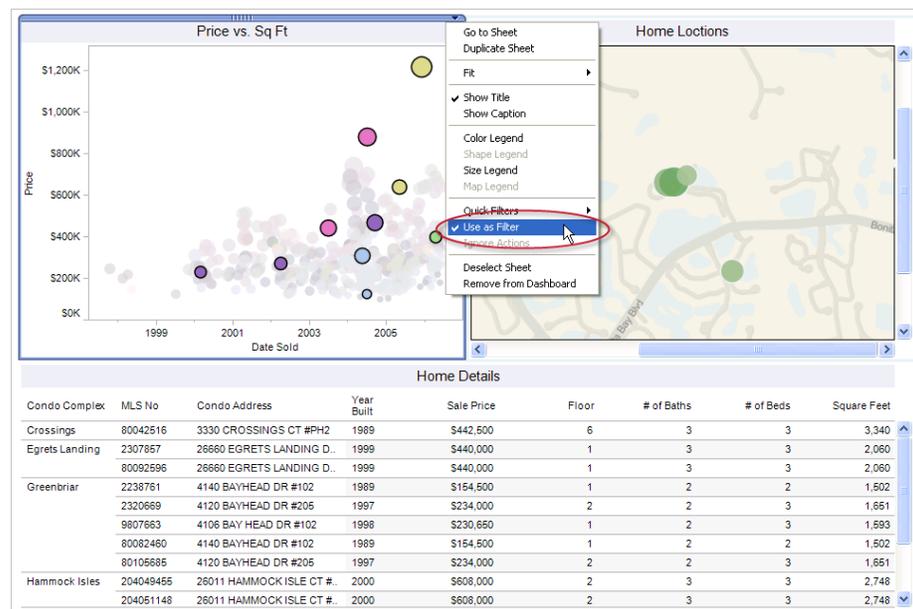
Longitude in the view. Additionally, the link is only available on marks and headers that contain relevant values.

## Actions and Dashboards

Actions often have special behavior when the source or destination is a dashboard. Filter and Highlight actions can affect other views in the dashboard and URL actions can update a webpage object so you don't have to open your web browser. Finally, you can create simple Filter and Highlight actions using special menu options so you don't have to open the Actions dialog box.

### Example: Filter Actions in a Dashboard

This example shows how to create a filter action in a dashboard. The example shows a Real Estate dashboard with three views. Using the **Use as Filter** option you can set one of the views to act as a filter on all the other views in the dashboard. In this case the scatter plot in the upper right is filtering the map view and the text table to show more details about the selected houses.



The **Use as Filter** command can only apply to one view at a time. A filter action is created that you can modify in the Actions dialog box.

## Example: URL Actions in a Dashboard

This example shows how a URL action works with a web page object in a dashboard. Below is a dashboard showing sales information by product for several stores in a coffee franchise. Included in the dashboard is a web page object that shows product details. The text table has a URL action that points at that web page. When you launch the action the web page automatically updates within the dashboard rather than opening a web browser.

Product Type	Product	Central		East		Market South		West		Grand Total	
		Sum of Sal...	Sum of Profit	Sum of Sal...	Sum of Profit	Sum of Sales	Sum of Profit	Sum of Sal...	Sum of Profit	Sum of Sal...	Sum of Profit
Coffee	Amaretto	14,012	5,104	2,994	1,010			9,263	-1,224	26,269	4,890
	Columbian	28,911	8,525	47,385	27,256	21,663	8,767	30,352	11,256	128,311	56,804
	Decaf Irish Cream	26,157	9,635	6,262	2,726	11,596	2,935	18,233	-1,307	62,248	13,989
Espresso	Caffe Latte					15,443	3,873	20,456	7,502	35,899	11,376
	Caffe Mocha			6,646	-6,232	14,186	5,202	18,874	4,066	84,904	17,678
	Decaf Espres			720	2,411	15,381	5,930	30,578	12,302	78,162	29,502
	Regular Espr			1,031	10,065					24,031	10,065
Herbal Tea	Chamomile			193	764	11,183	3,178	25,631	8,854	75,578	27,231
	Lemon			1,177	7,902	14,494	2,593	32,273	13,121	95,926	29,869
	Mint			991	-2,243			14,384	4,328	35,710	6,154
Tea	Darjeeling			694	6,500			28,773	11,784	73,151	29,053
	Earl Grey			3,404				27,382	10,426	66,772	24,164
	Green Tea			5,209	1,227	11,576	5,884	16,065	-7,112	32,850	-231
<b>Grand Total</b>		<b>265,045</b>	<b>93,852</b>	<b>178,576</b>	<b>59,217</b>	<b>103,926</b>	<b>32,478</b>	<b>272,264</b>	<b>73,996</b>	<b>819,811</b>	<b>259,543</b>

**Caffe Latte**

Rich, full-bodied Starbucks® espresso in steamed milk lightly topped with foam.

may we suggest?

Try something new

[Cappuccino](#)

nutrition facts table (customize)

---

## Using Field Values in Actions

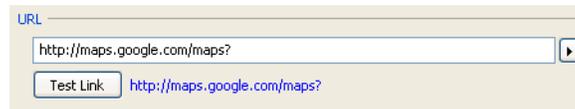
When you add an action in Tableau you often want to use values from your data as parameters in the name of the action as well as the action itself. Using fields as variables in the action name makes the menu item that launches the action specific to the selected mark. More commonly, using field values as parameters in the URL of a URL action allows you to send information about a specific data point to the destination webpage.

### Using Field Values in URLs

Tableau lets you add fields as variables into URL actions 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.

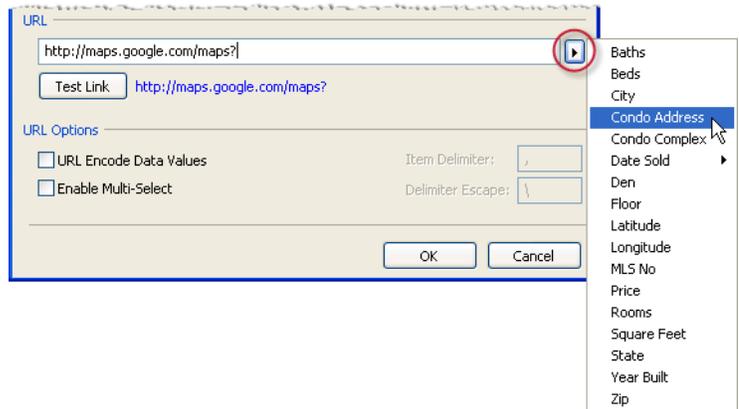
#### To add a field as a parameter to a URL action:

- 1 In the Add URL Action 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. To use aggregated field values as a parameter in a link, you must first create a calculated field and then use the name of that field in the link. The calculated field must also be used in the view in order for the link to be available. A good way to use these fields is by placing them on the Level of Detail shelf.

---

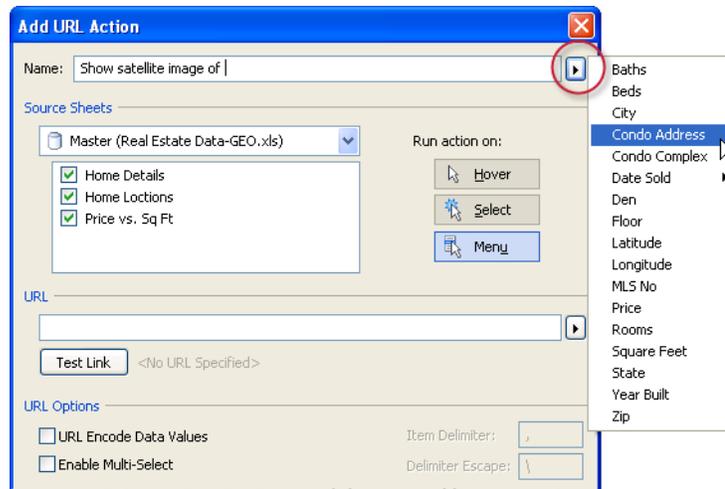
### Using Field Values in Action Names

In addition to using field values in URLs, you can use fields as variables in the action names. The name of the action displays on the context menu when an action is launched using the menu. Using field variables in the name is useful in making the action specific to the selected mark. In a view showing real estate information, you could name a URL action that points at 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 Action dialog box, begin typing the name for the action.
- 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.



# Calculations

---

<b>Overview</b> . . . . .	<b>21-3</b>
<b>Aggregations</b> . . . . .	<b>21-4</b>
How Aggregation and Disaggregation Work . . . . .	21-8
Example – Aggregating and Disaggregating Data . . . . .	21-11
<b>Calculated Fields</b> . . . . .	<b>21-18</b>
How to Create a Calculated Field . . . . .	21-18
How to Create a Calculated Member . . . . .	21-21
Copying and Pasting Calculated Fields . . . . .	21-24
Example – Creating a Calculated Field . . . . .	21-26
Table Calculations . . . . .	21-40
Aggregate Calculations . . . . .	21-34
<b>Table Calculations</b> . . . . .	<b>21-40</b>
Quick Table Calculations . . . . .	21-49
Defining Basic Table Calculations . . . . .	21-50
Secondary Table Calculations . . . . .	21-67
Customizing Table Calculations . . . . .	21-68
<b>Binned Data</b> . . . . .	<b>21-72</b>
<b>Totals</b> . . . . .	<b>21-77</b>
Local vs. Server Computation . . . . .	21-77
Grand Totals . . . . .	21-78
Subtotals . . . . .	21-84



<b>Percentages . . . . .</b>	<b>21-87</b>
About Percentages . . . . .	21-87
Percentages and Aggregations . . . . .	21-88
Percentage Options . . . . .	21-91



## 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
ATTR	Returns the value of the given expression if it only has a single value for all rows in the group, otherwise it displays an asterisk (*) character. Null values are ignored.	N/A
Dimension	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)



<b>Aggregation</b>	<b>Description</b>	<b>Result for measure that contains 1, 2, 2, 3</b>
Maximum	Computes the largest number in a measure or a continuous dimension. Null values are ignored.	1 value (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)



<b>Aggregation</b>	<b>Description</b>	<b>Result for measure that contains 1, 2, 2, 3</b>
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)
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 21-34. 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. Change the default aggregation by right-clicking a measure in the Data window and selecting **Field Properties > Aggregation**. Below the default aggregation for the Budget Margin measure is set to Average.

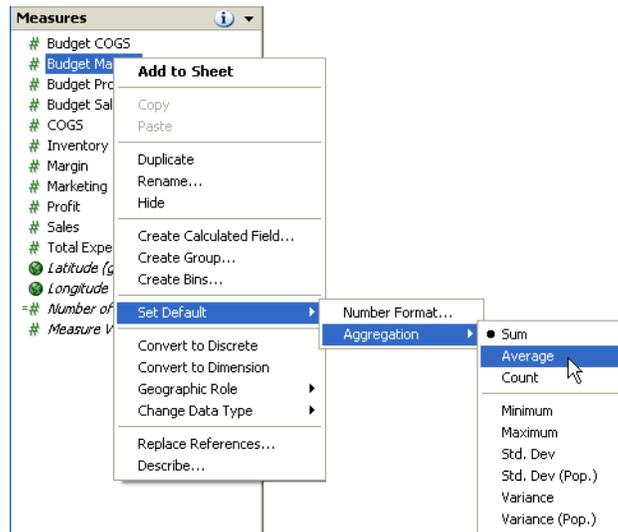


Tableau also allows you to view data in **disaggregated** form (relational databases only). This is an extremely powerful feature. When data are disaggregated, you can view all of the individual rows of your data source. For example, after discovering that the sum of sales for rubber bands is \$14,600, 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 21-8). 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 19-16 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, **Profit** becomes **SUM(Profit)**.

---

**Note** You can aggregate measures using Tableau only for relational data sources. Multidimensional data sources contain aggregated data only.

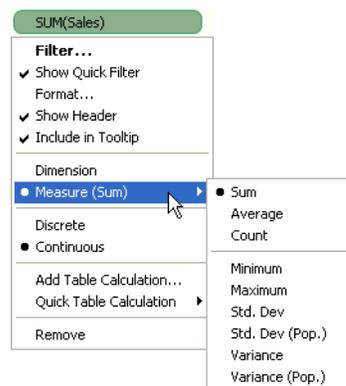
---

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** as a summation, **Profit** as a maximum, and **Discount** as an average.

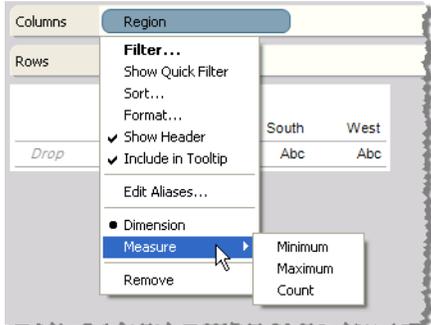


You can change the aggregation state for all the measures on a worksheet by selecting the **Analysis > Aggregate Measures** menu item.



When all measures are disaggregated you see a mark for each row in the view. You cannot select specific marks to Keep Only, Exclude, or create a Set when all measures are disaggregated.

**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 all of the other predefined aggregations to a dimension, you can apply Dimension, Minimum, Maximum, and Count.





## Disaggregating Data

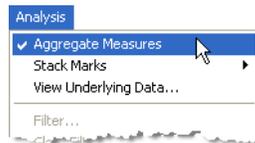
Disaggregating your data allows you to view every row of the data source which can be useful when you are analyzing measures that you may want to use both independently and dependently in the view. For example, you may be analyzing the results from a product satisfaction survey with the Age of participants 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 “Performance Tips” on page 35-1.

---

You can disaggregate all measures in the view by selecting **Analysis > Aggregate Measures**.



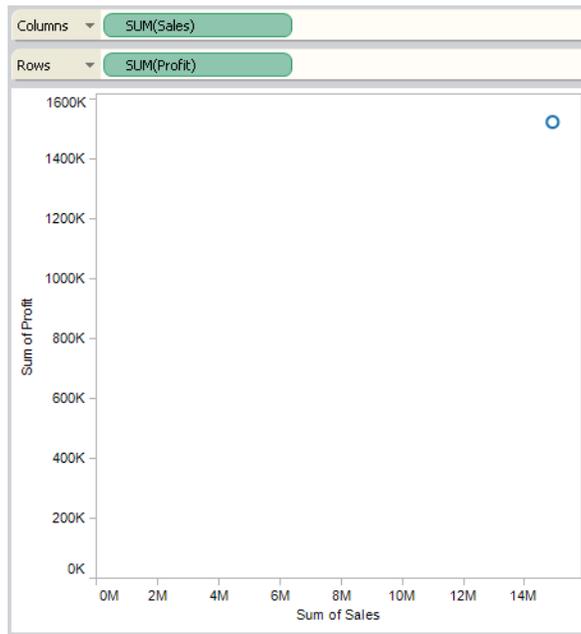


### Example – Aggregating and Disaggregating Data

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** measure on the **Columns** shelf and the **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 and the profit for the entire data source. That is, the summations are performed using every row in the data source.



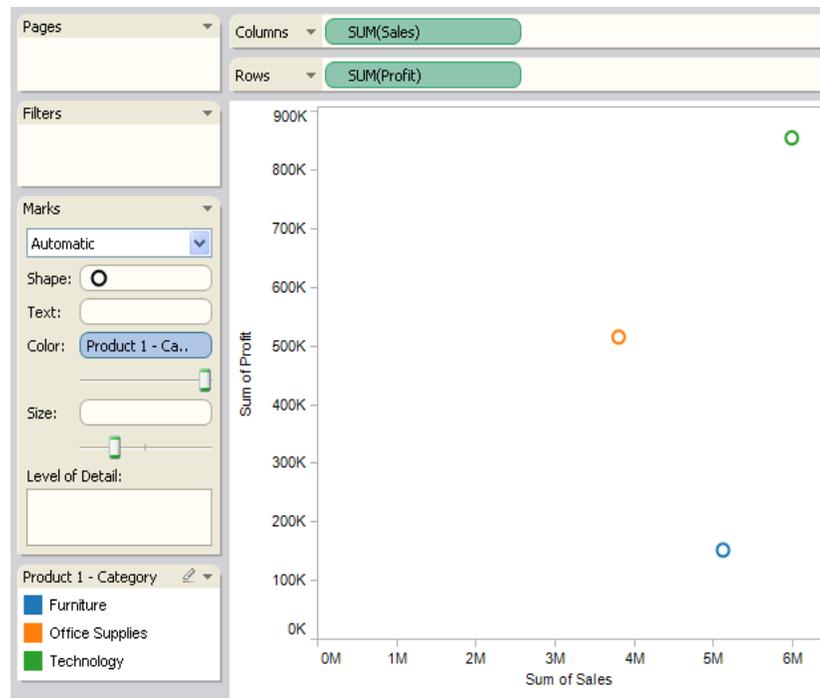
- 2 Place the **Product 1 - Category** dimension on the **Color** shelf.

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 1- Category**



dimension on the **Color** shelf separates the data into three marks—one for each dimension member—and then encodes the marks using color.

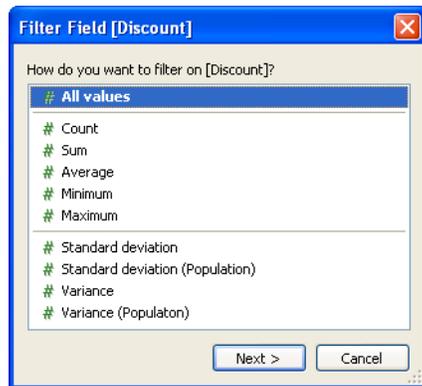
Although more marks are displayed, the measures are still aggregated. The single mark in the view indicates the sum of the sales and the sum of the profit for Office Supplies. If you were to sum the sales and profit values for the three marks, you would produce the values for the entire data set as given in the previous step of this example.





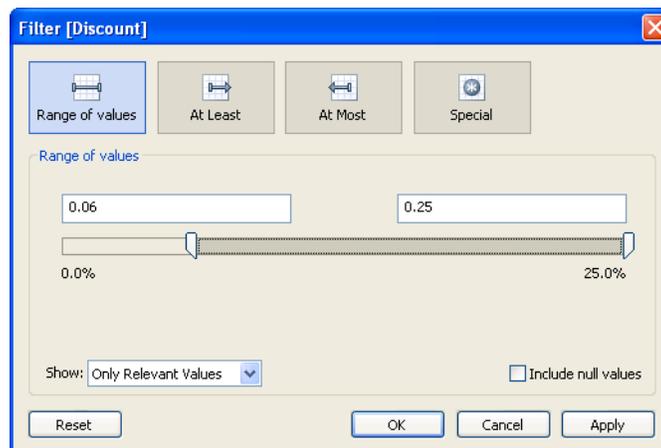
3 Place the **Discount** measure on the **Filters** shelf.

In the **Filter Field** dialog box select **All Values** to filter on the disaggregated measure.



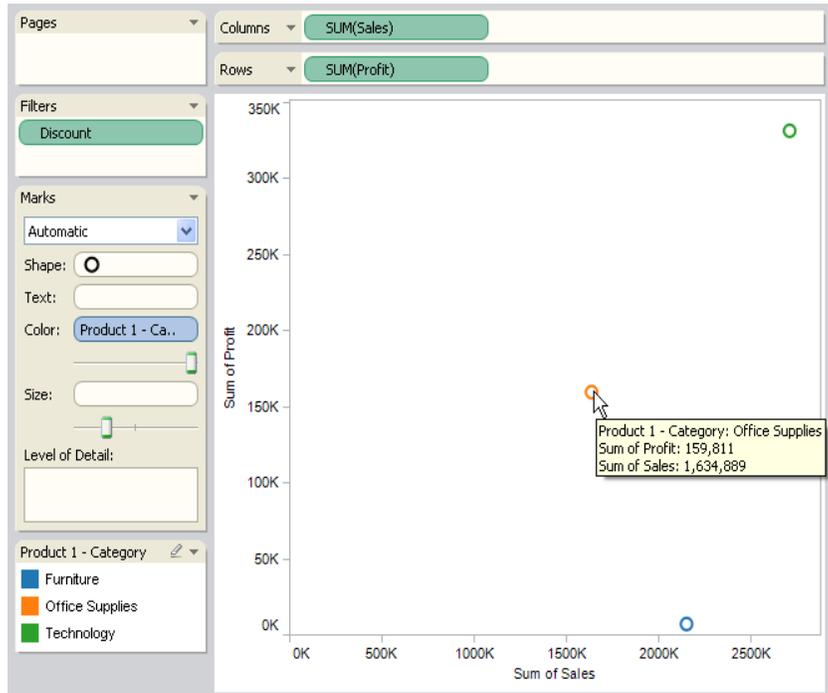


Filter the data to only include discounts greater than 6% (0.06). Because **Discount** is disaggregated, Tableau applies the filter to each row in the data source before performing the aggregations for the **Sales** and **Profit** measures.



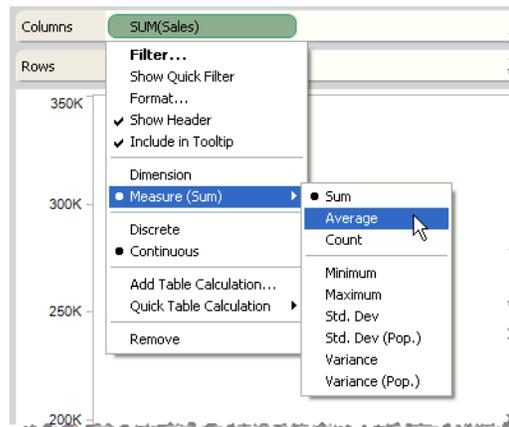


The view is shown below. The tooltip indicates that both the sales and the profit numbers are smaller than in the previous view. This is because data have been filtered out of the aggregation operation.

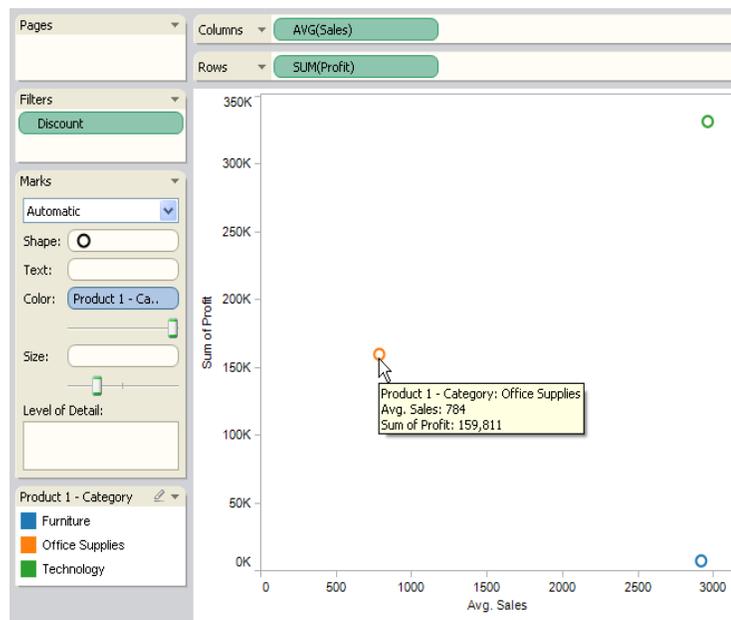


4 Change the aggregation of Sales to an average.

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.

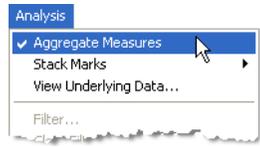




5 Disaggregate the data.

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.

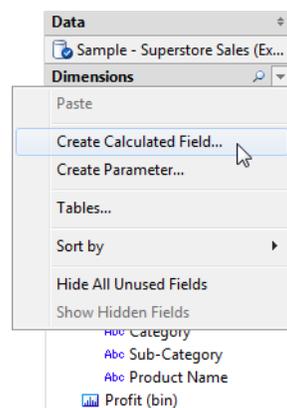
Tableau allows you to create a new calculated field by defining a formula that is based on data source fields and other calculated fields, and that uses standard functions and operators.

This section discusses the following topics:

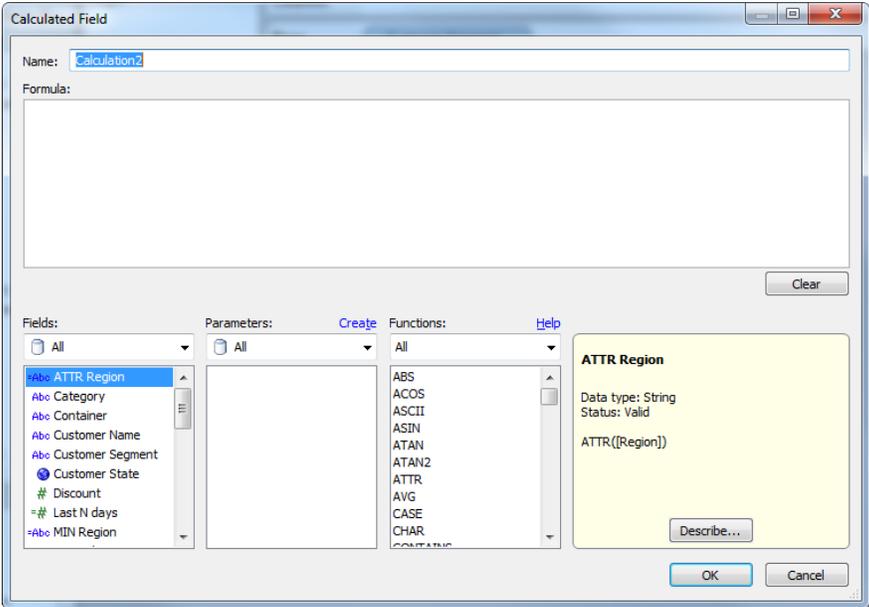
- How to Create a Calculated Field
- How to Create a Calculated Member
- Writing formulas in Tableau
- Example – Creating a Calculated Field
- Aggregate Calculations
- Table Calculations

### 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 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 21-24 for more information about how to define a formula.
- 3 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 11-38 for more information about continuous and discrete fields.

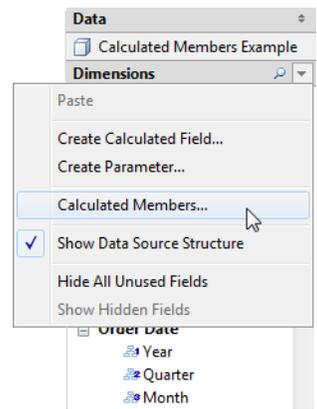


## How to Create a Calculated Member

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

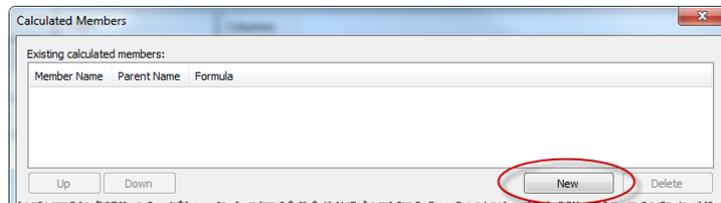
You can define a calculated dimension member by selecting **Calculated Members** on the Data window menu. The Calculated Members dialog box opens where you can create, delete, and edit any calculated members.



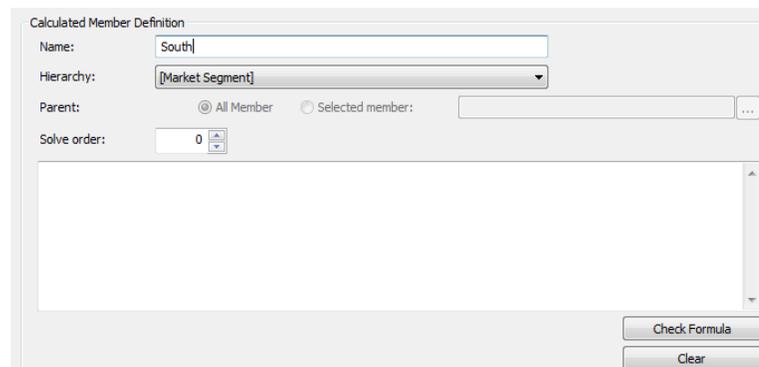
To create a new calculated member do the following:



- 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

The formula editor has built-in coloring and validation to help you avoid syntax errors. As you write the formula, syntax errors are underlined with a red squiggly line. Hover over the error to see directions for fixing it. Also any errors with the calculation are shown in a drop-down list. When the calculation is valid, a green check mark is displayed.

When you are writing formulas, any part that displays in bold indicates that it will be computed locally within Tableau on the aggregated results. Any normal weight text will be computed at the database level.

Formulas are made up of the following 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 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. Functions are colored black in the formula. Refer to “Functions” on page 34-6 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’s data type and the name display in the list. Refer to “Data Types” on page 11-34 to learn more about data types. Use the drop-down menu to select a secondary data source and see its fields.

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]]].

Fields are colored orange in the formula.

### 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 34-37 for a complete description of the operators and the associated precedence rules. Operators are colored black in the formula.

### 4 Parameters (optional)

Parameters are placeholder variables that can be inserted into calculations to replace constant values. When a parameter is used in a calculation, you can then use a parameter control to dynamically change the value. Refer to “Parameters” on page 22-1 to learn more. Parameters are colored purple in the formula.

### 5 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 (//). Comments are colored green in the formula.

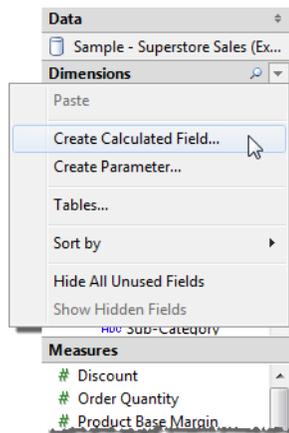


## Example – Creating a Calculated Field

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 Data window menu.

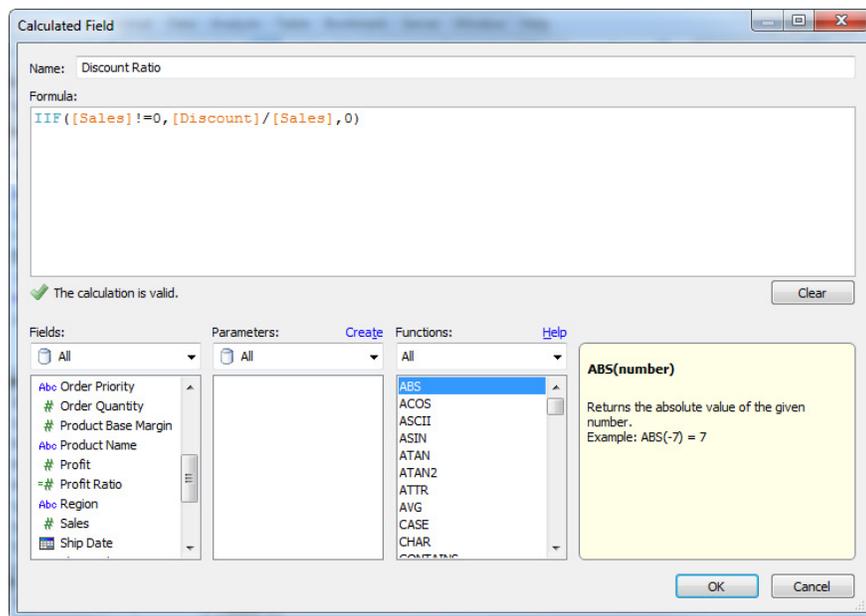


2 Complete the **Calculation** dialog box.

Name the new field **Discount Ratio** and enter the formula shown below.

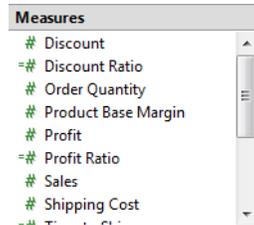
`IIF([Sales] !=0, [Discount]/[Sales],0)`

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.

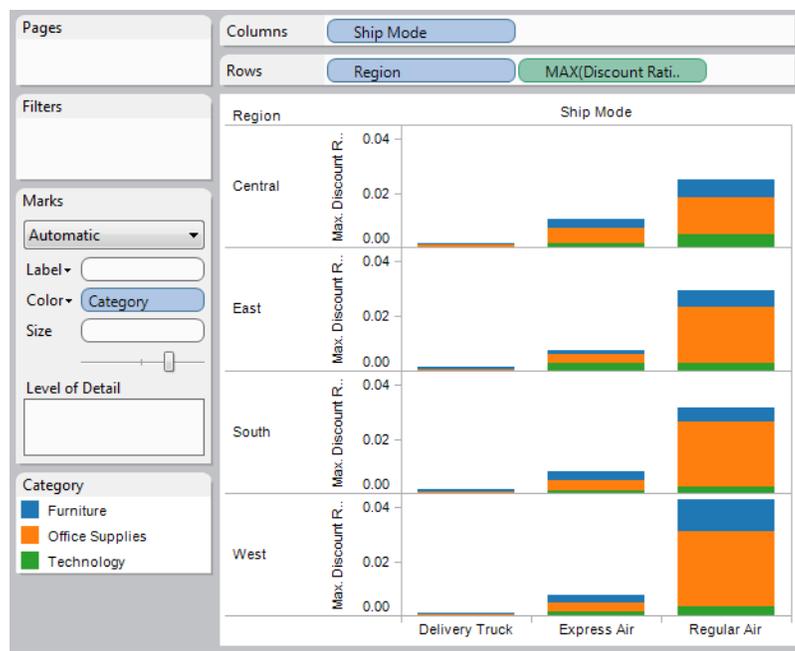


**3** Add the calculation to the view.

Place **Ship Mode** on the columns shelf, **Region** on the Rows shelf, and **Product 1-Category** 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

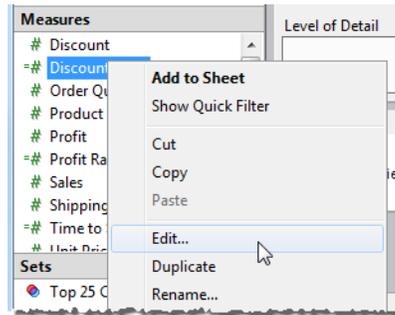


example, you can apply an aggregation to it. Below, **Discount Ratio** is aggregated as a maximum.

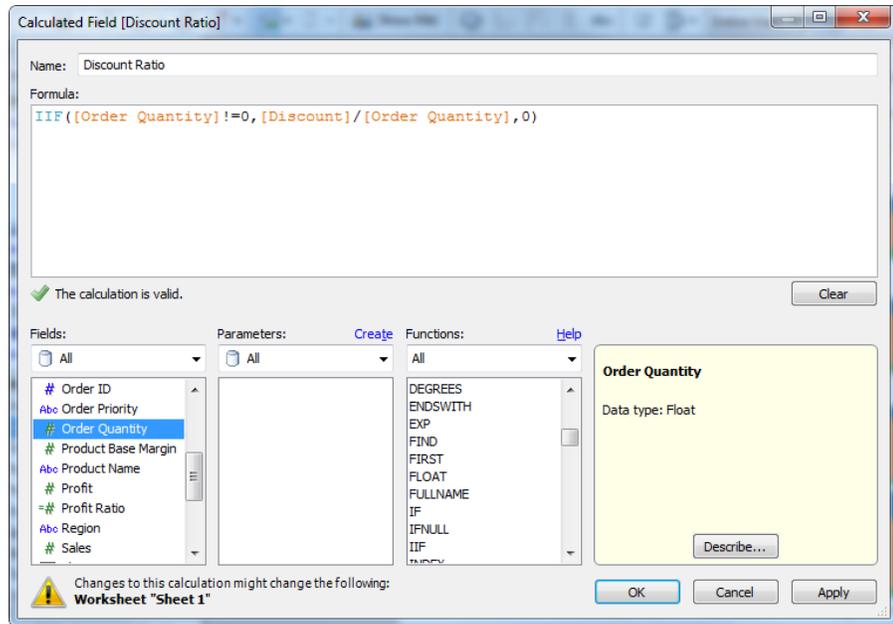


4 Edit the calculation.

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**.

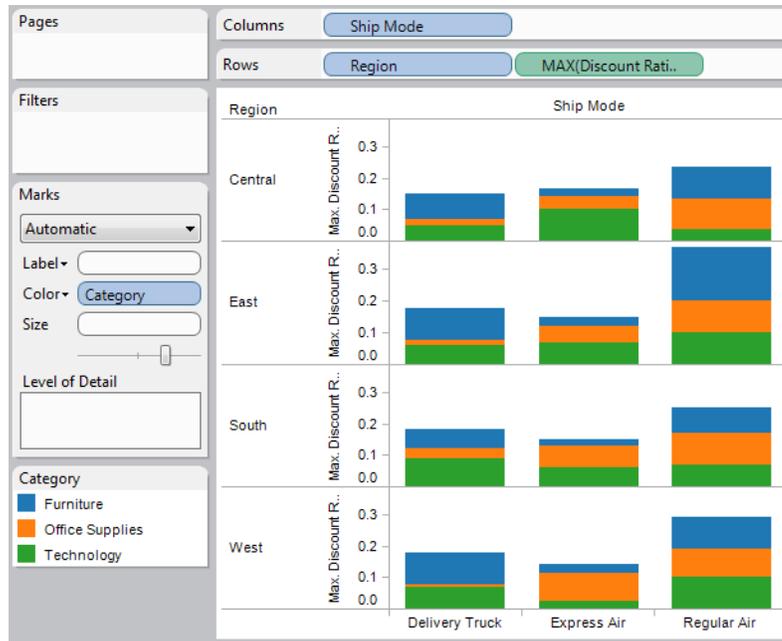


In the Calculated Field dialog box, change Sales to Order Quantity.





The view automatically updates after you click **OK** in the Calculated Field dialog box.

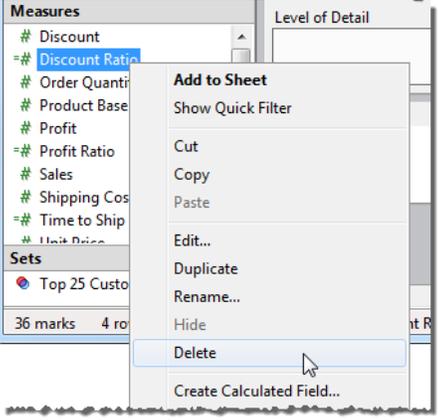


5 Delete the calculated field.

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 “Saving and Exporting” on page 31-1 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 21-4, 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 profit divided by the sales. Then you could place this measure on a shelf and use the predefined summation aggregation. In this scenario, **Margin** is defined as follows:

$$\text{Margin} = \text{SUM}([\text{Profit}] / [\text{Sales}])$$

This formula calculates the ratio of 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 profits divided by the sum of all sales. That formula is shown below.

$$\text{Margin} = \text{SUM}([\text{Profit}]) / \text{SUM}([\text{Sales}])$$

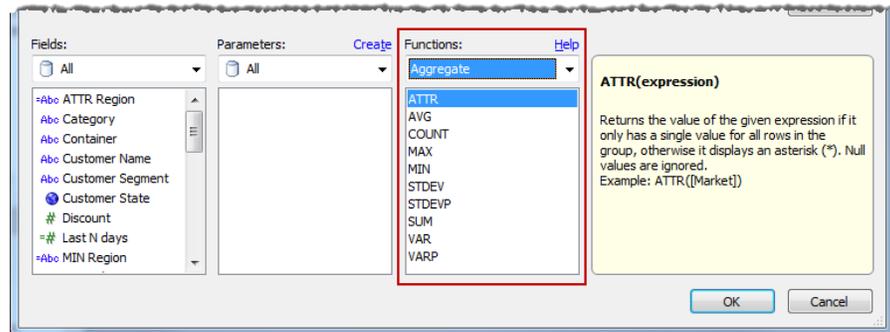
In this case, the division is performed after each measure is aggregated. An aggregate calculation allows you to create formulas like this.



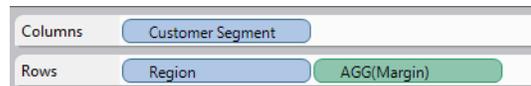
## How to Create an Aggregate Calculation

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 21-18. 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 21-4.



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 21-36 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,  $SUM(Price) * [Items]$  is not a valid expression because  $SUM(Price)$  is aggregated and  $Items$  is not. However,  $SUM(Price * Items)$  and  $SUM(Price) * SUM(Items)$  are both valid.



- Constant terms in an expression act as aggregated or disaggregated values as appropriate. For example:  $SUM(Price*7)$  and  $SUM(Price)*7$  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:  $MAX(SUM(Sales),Profit)$  is not a valid expression because Sales is aggregated and Profit is not. However,  $MAX(SUM(Sales),SUM(Profit))$  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 21-80 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
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 21-4 for descriptions of the aggregation functions.

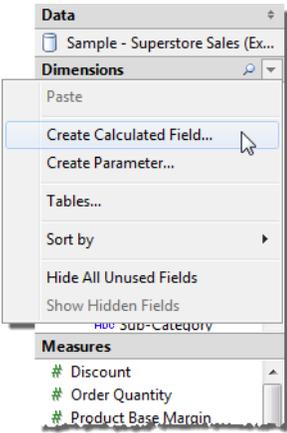


Therefore, if you define an aggregate calculation called Margin that is equal to  $SUM(\text{Profit})/SUM(\text{Sales})$  and then disaggregate the data, it is interpreted as Profit/Sales.

**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 Data window menu.

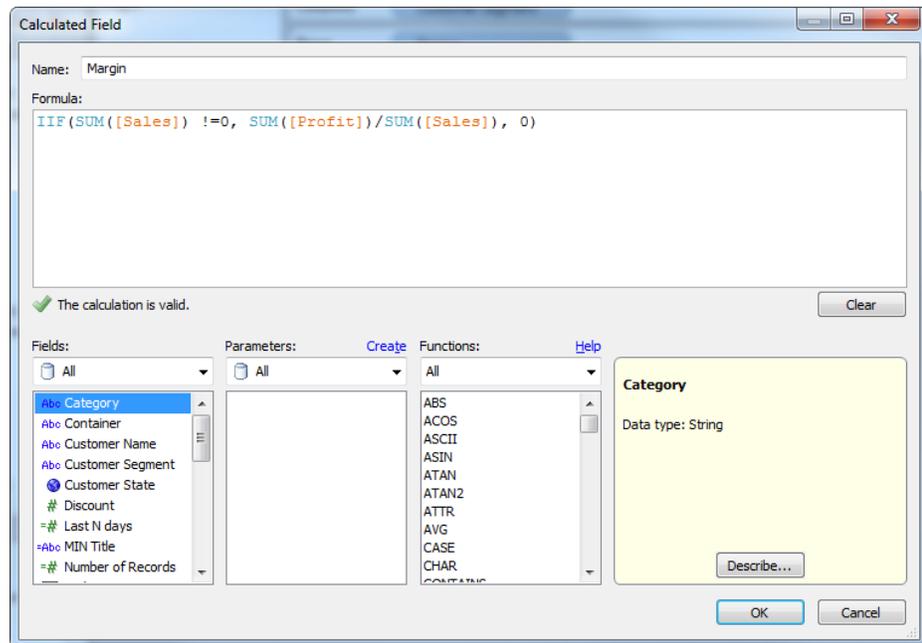


2 Define the calculation.

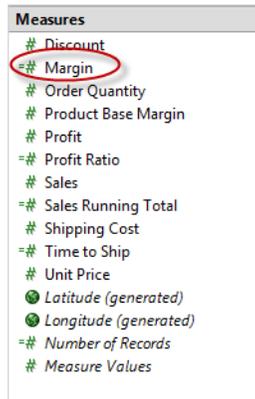
Name the new field **Margin** and enter the formula shown below.

`IIF(SUM([Sales]) !=0, SUM([Profit])/SUM([Sales]), 0)`

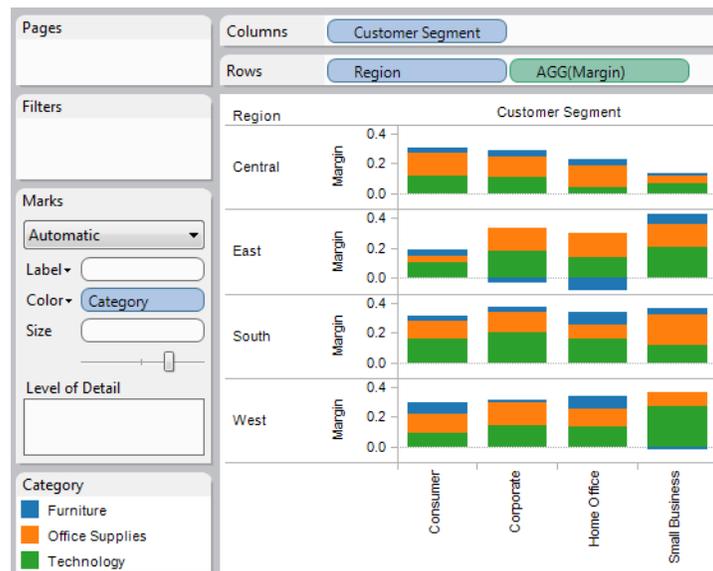
You enter the formula by selecting functions from the **Functions** area of the dialog box, and field names from the **Fields** area of the dialog box. You must type the operators (!= and /) manually. Note that the IIF statement is used to avoid dividing by zero.



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.



## 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.

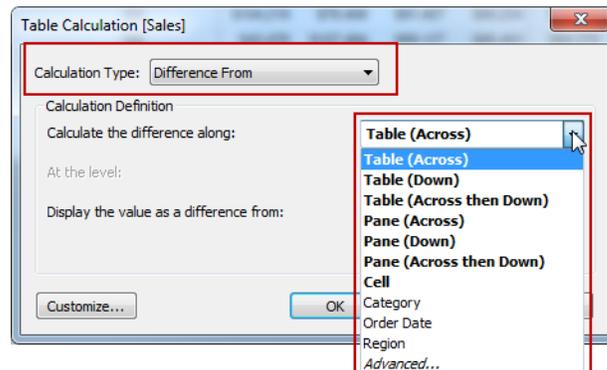
All table calculations are computed locally using the values you see in the table. This means that computing a moving average for a measure that is aggregated as an average results in an average of the averages.

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:

- Understanding Table Calculations
- Quick Table Calculations
- Defining Basic Table Calculations
- Secondary Table Calculations
- Customizing Table Calculations

### Understanding Table Calculations

Table calculations are computations that are applied to the values in the table. These computations are unique in that they use data from multiple rows in the database to calculate a value. To create a table calculation, you need to define both what you want to compute and what to compute along. For example, when defining a running sum calculation for several years you are computing a running sum along the Date field. These are defined in the calculated field dialog box using the Calculation Type and Calculate Along drop-down menus.



The definition of what to compute along has two parts: addressing fields and partitioning fields.

### Addressing and Partitioning

The addressing fields define what part of the table you are computing along. The partitioning fields define how to group the calculation. In the example of a running sum of product sales across several years, the addressing field is the Date field while the partitioning field is the product field. When you define the addressing for a table calculation, all the other fields are used for partitioning.

You can specify the addressing in the Table Calculation dialog box. The addressing can be relative to the table structure or a specific field. Each addressing option is described below.

**Table (Across).** This option sets the addressing to compute along the entire table moving horizontally through each partition. For example, the view below shows quarterly sales by region and product category. When a calculation addressing is set to Table Across, the fields that span horizontally across the table are the addressing fields (Category and Region). All the other fields (Year, Quarter) are partitioning. The addressing fields are shown in orange while partitioning fields are shown in blue.



Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$56,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

That means that each partition will be the combination of Year and Quarter.

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$56,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

**Partition**  
**1**  
**2**  
**...**

**Table (Down).** This option sets the addressing to compute along the entire table moving vertically through each partition. For example, the same view from above is shown below with the addressing set to compute along Table Down. The fields that span vertically (Year, Quarter) are now the addressing fields and the rest of the fields are partitioning (Category and Region). The addressing fields are shown in orange while partitioning fields are shown in blue.



Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$56,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

That means that each partition is the combination of Category and Region.

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$56,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

Partition 1 2 ...

**Table (Across then Down).** This option sets the addressing to compute across the entire table horizontally and then down the table vertically. This means that both the fields that span across the table and down the table are addressing fields.



Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$66,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$66,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$68,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,046	\$163,606	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

That means that the entire table is the partition. The computation will compute across, move to the next row and continue to compute across, and so on.

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$66,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$66,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$68,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,046	\$163,606	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

Partition 1

**Pane (Across).** This option sets to compute across the pane horizontally. The fields that span across the pane horizontally are the addressing fields. However, the fields that separate the panes are now partitioning fields. In the example below Category becomes a partitioning field along with Year and Quarter. Region is the addressing field.



Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$56,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

That means that the combination of Year, Quarter, and Category is the partition.

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$56,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

**Pane (Down).** This option sets the addressing to compute down the table within the pane. The fields that separate the pane (Category and Year) are partitioning fields. In addition, the Region field becomes a partitioning field while the Quarter field is the addressing field.

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$56,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042



That means that the combination of Year, Category, and Region is the partition.

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$56,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

**Pane (Across then Down).** This option sets the addressing to compute across within the pane, then move to the next row and continue to compute across. The addressing fields are both the fields that run across the table horizontally and down the table vertically (Region and Quarter). The partitioning fields are the fields that define the pane (Category and Year).

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$56,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

That means that the combination of Category and Year make up the partition.



Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$64,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$66,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$167,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

**Cell.** This option sets the addressing to the individual cells in the table. All fields become partitioning fields. This option is generally most useful when computing a percent of total calculation.

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$64,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$66,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$167,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

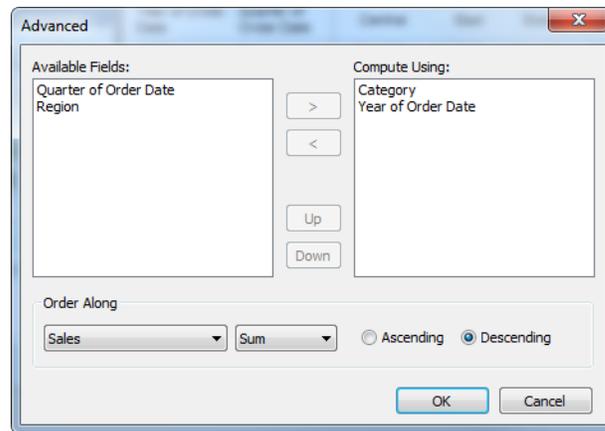
That means that the partition is the combination of Category, Region, Year, and Quarter.

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$64,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$66,705	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$167,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042



**Individual Fields.** This option sets the addressing to compute across the field you specify. The benefit of this option is that you get absolute control over how the calculation will be computed. Be careful though, because, addressing on an individual field means that when you rearrange the table, the calculation may no longer match the table structure.

**Advanced.** The advanced option lets you specify multiple fields to act as the addressing fields. When you select Advanced another dialog box opens where you can specify one or more fields to act as addressing fields. Then you can specify how to order those fields.



For example, in the view below the addressing fields are set to Category and Year. These are ordered by SUM(Sales) Descending. That means that the combination of Quarter and Region create the partition. Q1 Central exists four times in the table, and that is the partition.

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$98,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$66,705	\$103,741	\$103,680	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$104,219	\$78,127	\$69,127	\$65,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$43,478	\$107,494	\$69,127	\$65,453	\$64,378	\$44,255	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042



Because the order is set to SUM(Sales), the calculation is computed based on their SUM(Sales) values from highest to lowest.

Year of Order Date	Quarter of Order Date	Category / Region							
		Furniture				Office Supplies			
		Central	East	South	West	Central	East	South	West
2008	Q1	\$54,839	\$41,843	\$54,162	\$96,604	\$69,738	\$20,531	\$38,466	\$85,317
	Q2	\$71,341	\$56,706	\$103,741	\$103,660	\$43,070	\$38,547	\$65,251	\$108,219
	Q3	\$10,219	\$78,408	\$91,407	\$56,204	\$70,649	\$29,540	\$80,171	\$68,221
	Q4	\$4,478	\$107,494	\$89,127	\$85,453	\$64,378	\$44,265	\$23,944	\$52,238
2009	Q1	\$41,344	\$77,462	\$29,104	\$66,449	\$62,809	\$58,498	\$28,940	\$47,086
	Q2	\$78,894	\$64,045	\$163,505	\$69,963	\$46,951	\$21,775	\$66,439	\$50,694
	Q3	\$157,720	\$39,988	\$75,475	\$82,548	\$65,696	\$53,205	\$72,427	\$43,299
	Q4	\$74,210	\$77,246	\$49,308	\$44,382	\$87,564	\$65,831	\$34,448	\$87,042

### 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.

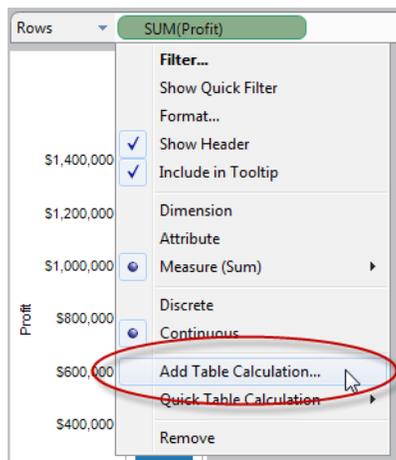


## Defining Basic Table Calculations

When you add a 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.

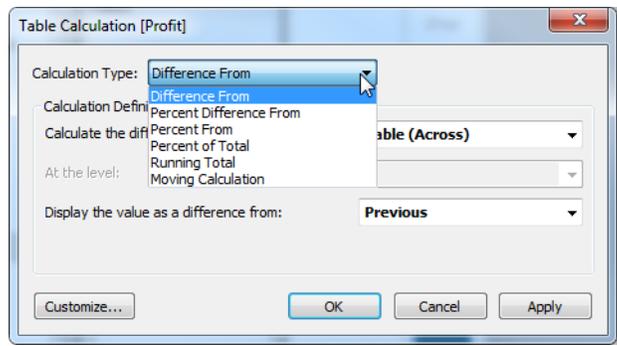
**To manually define a table calculation:**

- 1 Right-click the measure you want to use in the computation and select **Add Table Calculation**.

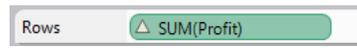




- 2 In the Table Calculation dialog box, select one of the following types of calculations from the drop-down menu at the top:



- 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.
- 3 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.
  - 4 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 2006 and 2007 quarterly sales for four different customer 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 2006 and 2007 quarterly sales numbers for several different customer segments of a superstore.

The screenshot shows the Tableau interface with the following configuration:

- Filters:** YEAR(Order Date)
- Columns:** YEAR(Order Date), QUARTER(Order Date)
- Rows:** Customer Segment
- Marks:** Automatic, SUM(Sales)

Customer S..	Order Date							
	2008				2009			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Consumer	191,307	136,044	187,366	200,602	203,761	186,186	166,203	161,240
Corporate	336,885	240,004	299,277	366,362	368,219	414,564	331,463	375,100
Home Office	177,731	225,400	178,564	206,091	220,707	203,956	171,392	188,252
Small Busin..	141,189	147,135	173,327	229,252	168,067	128,329	199,719	252,816

To compute the difference between 2008 and 2009 sales, you can define a table calculation using the definition shown below.

The dialog box 'Table Calculation [Sales]' is configured as follows:

- Calculation Type:** Difference From
- Calculation Definition:**
  - Calculate the difference along: Order Date
  - At the level: Year of Order Date
  - Display the value as a difference from: Previous
- Buttons:** Customize..., OK, Cancel, Apply

The difference is calculated along the Order Date dimension at the level of year because we are comparing 2008 sales to 2009 sales. The table now displays the difference between each quarter in 2009 and the corresponding quarter in the previous year. Notice that there are no values for 2008. That's because there are no previous years to compute the difference from. You can hide that column without affecting the calculation. Refer to "Hide Rows and Columns" on page 13-13 to learn more.

The screenshot shows a Tableau interface with the following configuration:

- Filters:** YEAR(Order Date)
- Columns:** YEAR(Order Date), QUARTER(Order Date)
- Rows:** Customer Segment
- Marks:** Automatic, SUM(Sales)

The resulting table displays sales data for four customer segments across four quarters for both 2008 and 2009. The 2008 columns are blank, and the 2009 columns show sales values and differences from 2008.

Customer S...	2008				2009			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Consumer					12,454	50,141	-31,163	-49,362
Corporate					31,334	174,560	32,176	8,738
Home Office					42,976	-21,444	-7,172	-17,839
Small Busin..					26,878	-18,807	26,392	23,564

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 2009 the total sales was \$203,761 while in the same quarter in 2008 the total sales was \$191,307. The difference between these two values is \$12,454.



Compute the difference between the Q1 sales in 2008 and 2009.

The result is shown here.

		Order Date							
		Q1		Q2		Q3		Q4	
Customer Segment	Year of Order Date	Sum of Sales	Difference in Sum of Sales	Sum of Sales	Difference in Sum of Sales	Sum of Sales	Difference in Sum of Sales	Sum of Sales	Difference in Sum of Sales
Consumer	2008	191,307		136,044		187,366		200,602	
	2009	203,761	12,454	186,186	50,141	156,203	-31,163	151,240	-49,362
Corporate	2008	336,885		240,004		299,277		366,362	
	2009	368,219	31,334	414,564	174,560	331,453	32,176	375,100	8,738
Home Office	2008	177,731		225,400		178,564		206,091	
	2009	220,707	42,976	203,966	-21,444	171,392	-7,172	188,252	-17,839
Small Business	2008	141,189		147,135		173,327		229,252	
	2009	168,067	26,878	128,329	-18,807	199,719	26,392	252,816	23,564

**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 21-49.

### Percent Difference From Calculation

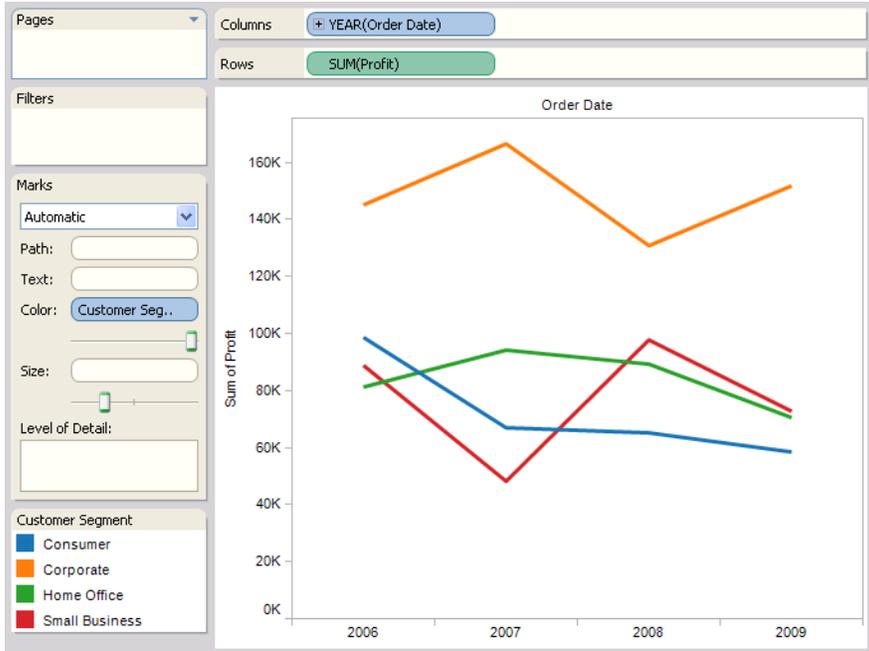
Use this type of calculation to display the rate of change between two specified values in the table by computing the difference as a percentage. A common use of this type of calculation is to compute the percent gain year after year (CAGR).

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.

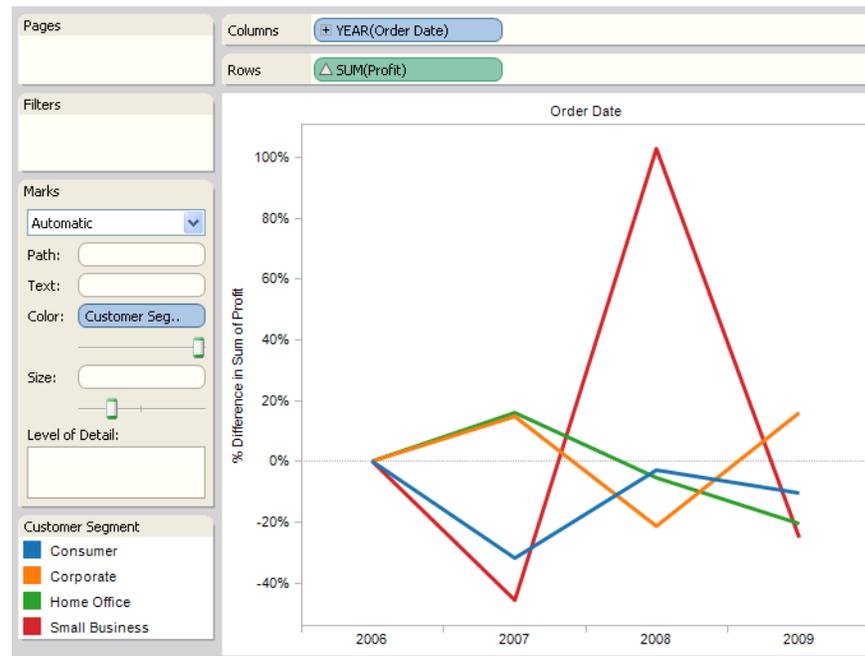
#### Example: Percent Difference From Calculation



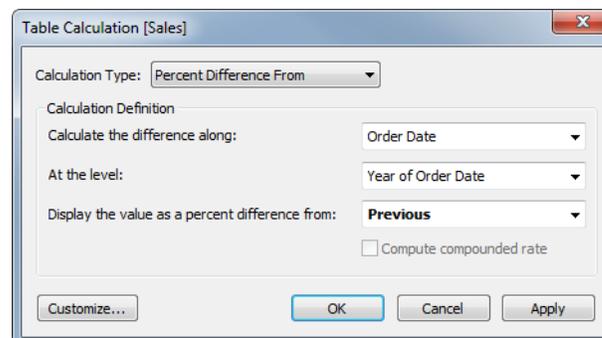
The table below shows the profit for several customer segments over four years. Looking at the view, we can see that there was a drop in profit in the Small Business and Consumer segments in 2007.



When we view this same view using a Percent Difference From calculation, it becomes clear that Small Business segment rebounded quite dramatically in 2008.



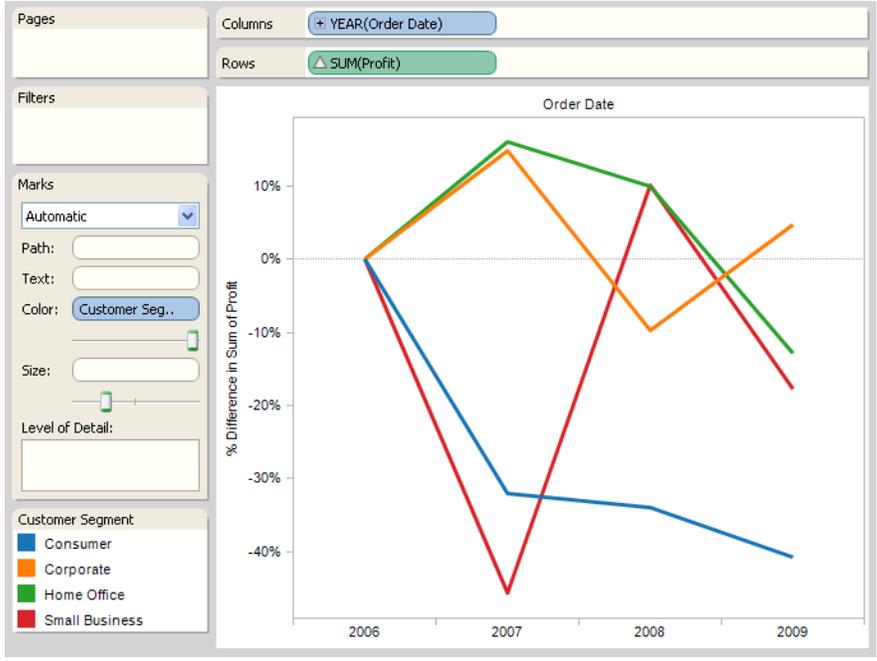
You can define a Percent Difference From calculation like this using the table calculation definition shown below.



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 2006 Profit. You can see that the Corporate segment is the only one to be above the 2006 profits in 2009.



**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 21-49.

### Percent From Calculation

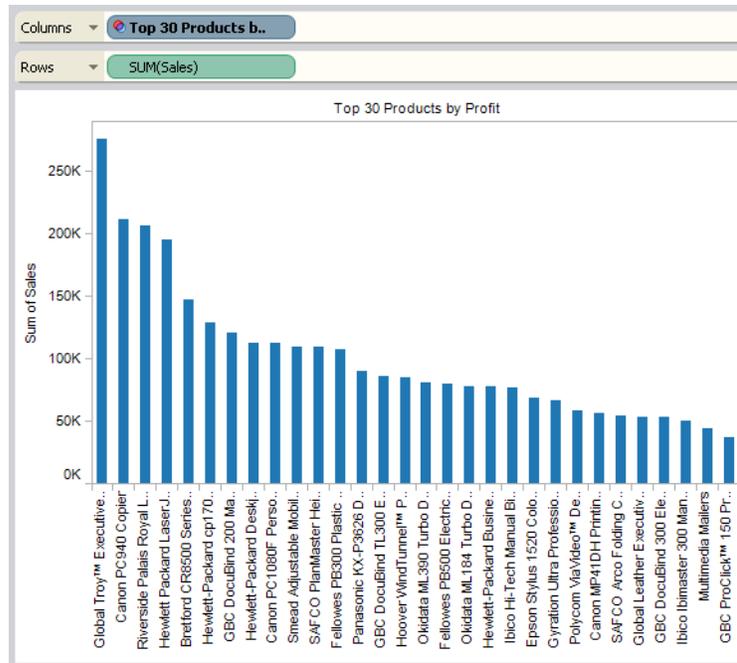
A Percent From calculation is similar to the Percent Difference From calculation in that you can use it to compute the change between two values as a percentage. However, this type of calculation computes an absolute change. For example, use the Percent From calculation to compare the sales performance of several products.



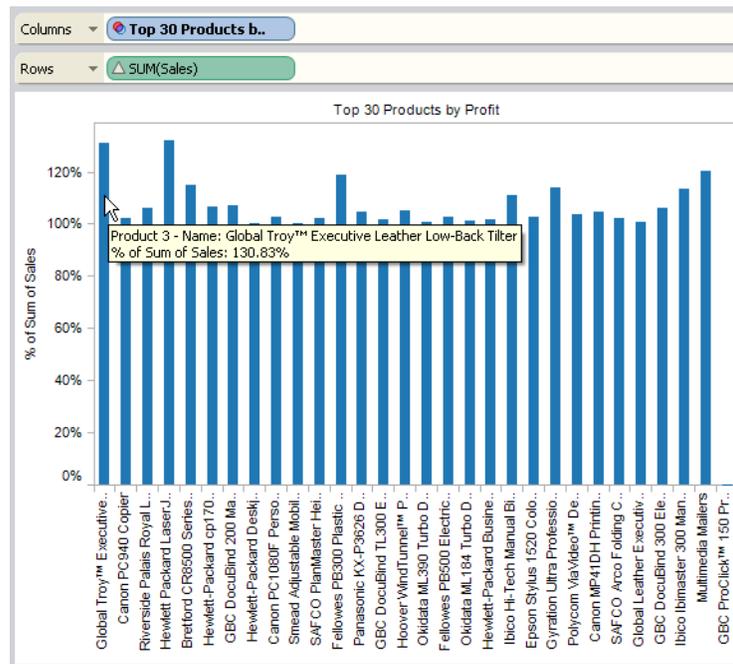
When you define a Percent From calculation you need to specify a table structure or dimension to calculate the percentage from. If you select a dimension, you also need to select a level. Finally, you need to select a value that each value in the table will be displayed as a percentage of. The following is an example of a Percent From calculation.

### Example: Percent From Calculation

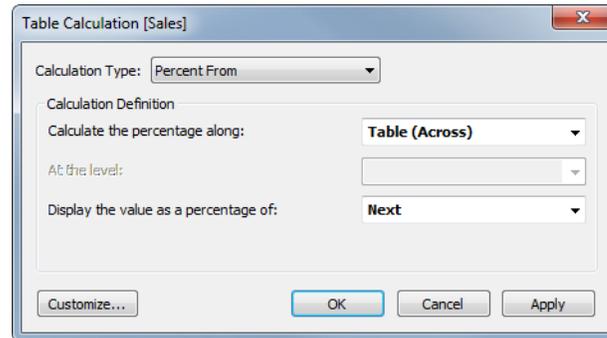
The view below shows the sales of the top thirty products by profit at a superstore. You can see that the top selling product is the Global Troy Executive Leather chair.



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 Leather Chair is 130% better in terms of sales than the Canon Copier.



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.



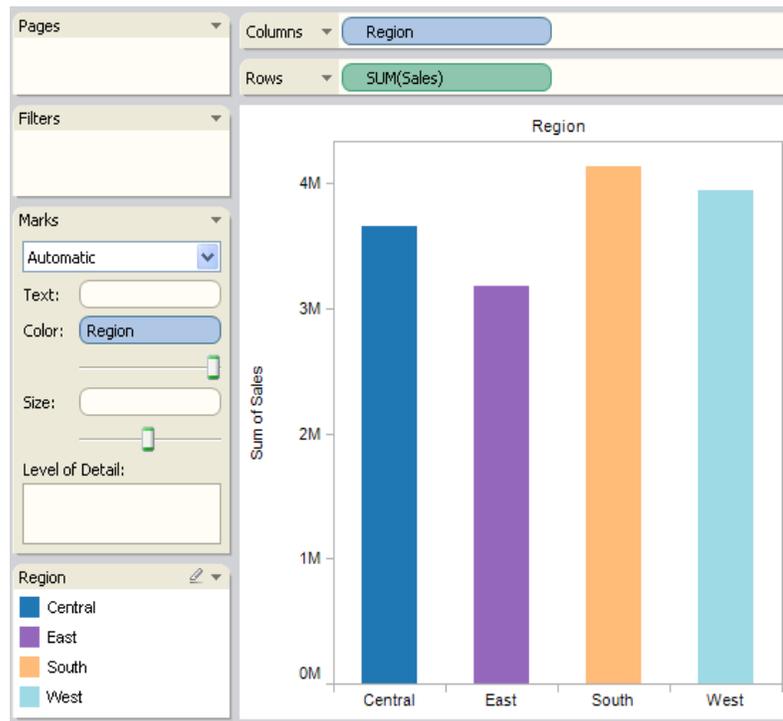
### Percent of Total Calculation

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.

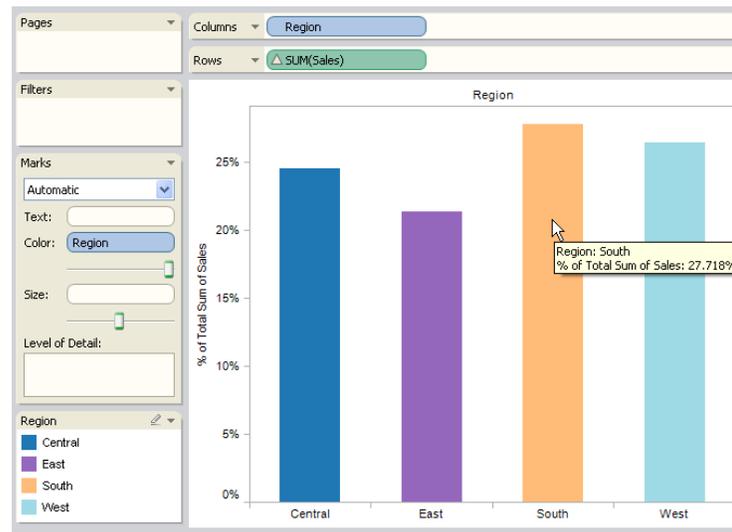
For a Percent of Total calculation you must specify a dimension or table structure to compute across and the dimension level to use in the computation. The dimension level is only necessary if you have chosen to compute across a dimension rather than a table structure. The following is an example of a Percent of Total calculation.

#### Example: Percent of Total Calculation

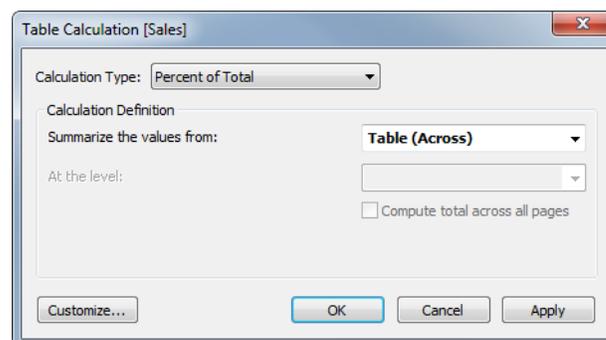
In this example, imagine that we are analyzing the performance of several stores. The view below shows the sales performance for three different regions. From the view we can see that the Southern region (orange marks) has the highest sales.



When we add the Percent of Total table calculation to the view, we can see that the South region accounts for just over 27% 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 21-49.

---



## 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 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 2006 to 2009.

Year of Orde..	Order Date			
	Q1	Q2	Q3	Q4
2006	1,249,502	874,001	1,017,316	1,049,104
2007	825,652	848,340	860,896	1,025,743
2008	847,112	748,583	838,533	1,002,307
2009	960,755	933,034	858,767	967,408

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.



Year of Orde..		Order Date			
		Q1	Q2	Q3	Q4
2006	Sales Total	1,249,502	874,001	1,017,316	1,049,104
	Running Total	1,249,502	2,123,503	3,140,819	4,189,923
2007	Sales Total	825,652	848,340	860,896	1,025,743
	Running Total	825,652	1,673,992	2,534,888	3,560,631
2008	Sales Total	847,112	748,583	838,533	1,002,307
	Running Total	847,112	1,595,695	2,434,228	3,436,536
2009	Sales Total	960,755	933,034	858,767	967,408
	Running Total	960,755	1,893,788	2,752,555	3,719,964

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 [Sales]

Calculation Type: Running Total

Calculation Definition

Summarize values using: Sum

Running along: Order Date

Restarting every: Year of Order Date

Perform a secondary calculation on the result

Buttons: Customize..., 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 21-49.

### Moving Calculation

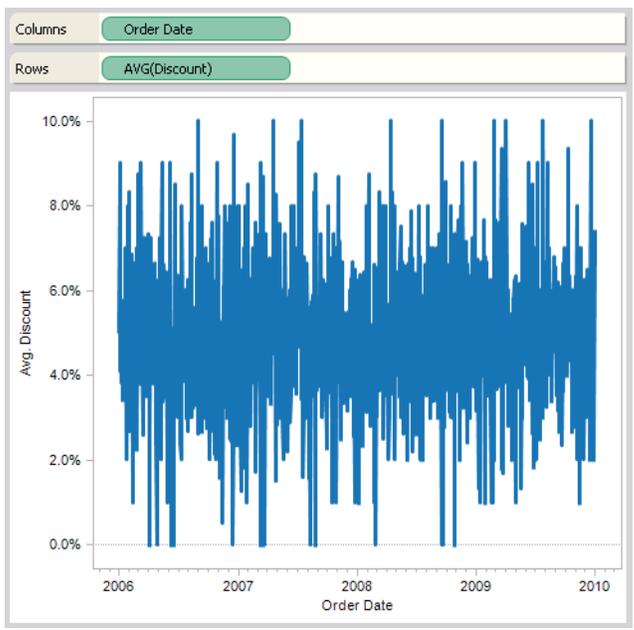
A moving calculation is typically used to smooth short term fluctuations in your data so that you can see long term trends. A good example is when you are looking at securities data. There are so many fluctuations every day that it is hard to see the big picture through the



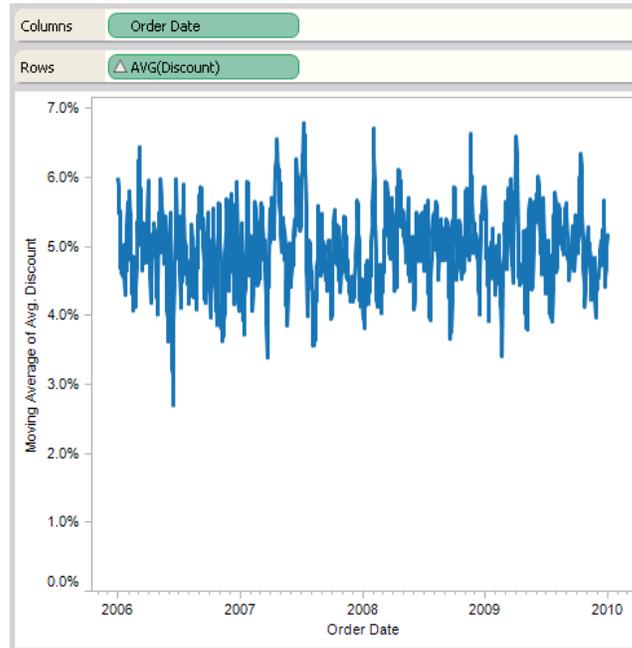
daily ups and downs. You can use a moving calculation to define a range of values to summarize using an aggregation of your choice.

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.

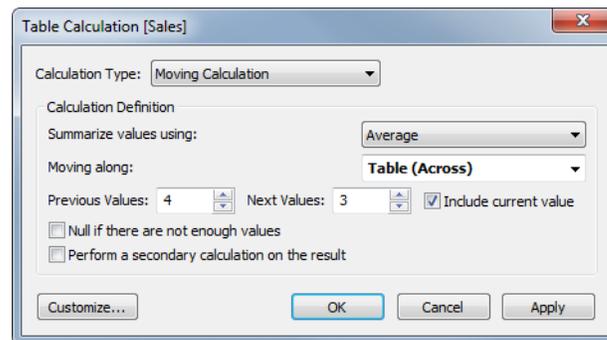
**Example - Moving Calculation.** The view below shows the the discounts given at a superstore along a continuous date axis. As you can see, it is very difficult to see any kind of trend in this view.



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.





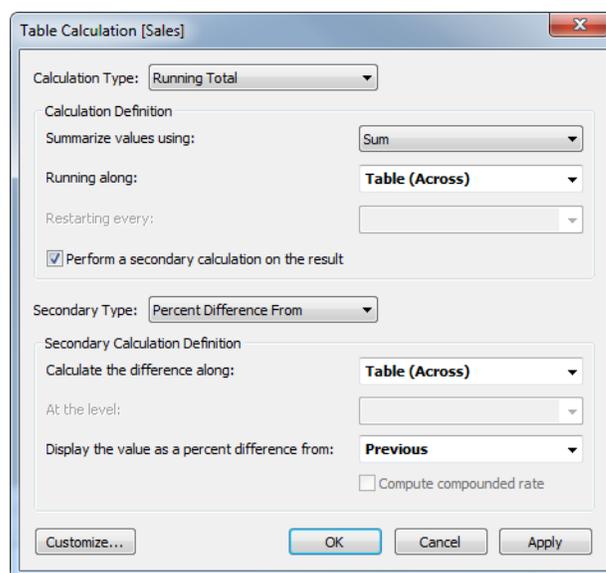
---

**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 21-49.

---

## Secondary 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.

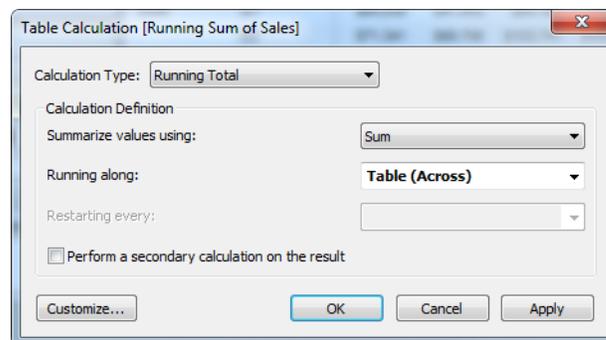




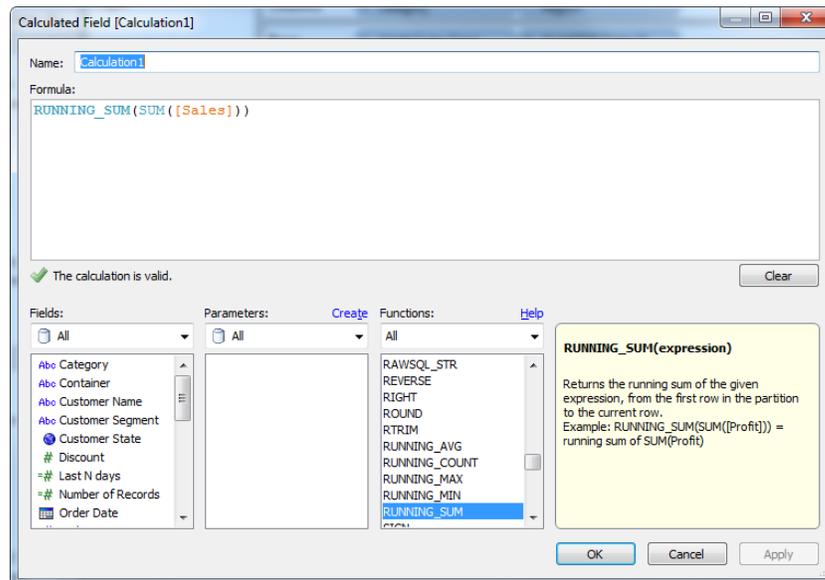
## Customizing Table Calculations

Table calculations are a special type of calculated field that computes on the local data in Tableau. While you can use the built-in table calculations such as Percent of Total, Difference From, Running Total, and so on; the functions required to define these calculations are also available for use in your own custom calculated fields. Customizing table calculations allows you to compute values such as the difference in number of orders this quarter versus an average quarter, total sales for regions that have above average margin, time since first click on a website, average temperature based on the last three days weighted at 10%, 40%, and 50%, and so much more.

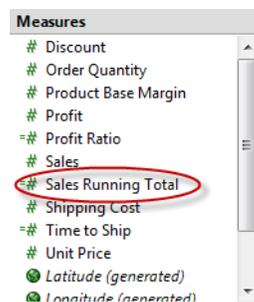
An easy way to become familiar with the Table Calculation functions is to add a basic table calculation and then click the Customize button in the lower left corner of the Table Calculation dialog box.



When you click Customize, the Calculated Field dialog box opens showing the formula for the calculation. You can see that it uses special functions. Refer to “Table Calculation Functions” on page 34-25 to learn more about each of these functions.



After you customize the calculated field, the changes are not saved until you click OK in the Calculated Field dialog box and in the Table Calculation dialog box. The new table calculation field is added to the Data window.



When you use that field in other views, it uses the default addressing and partitioning specified when the table calculation was created. You can change the addressing by right-clicking the field and select an option on the Compare To context menu.



The screenshot shows the Tableau interface with a context menu open over a 'Sales Running Total' pill. The menu includes options like 'Filter...', 'Show Quick Filter', 'Format...', 'Include in Tooltip', 'Discrete', 'Continuous', 'Compute using', 'Edit Table Calculation...', and 'Remove'. The 'Compute using' sub-menu is expanded, showing options like 'Table (Across)', 'Table (Down)', 'Table (Across then Down)', 'Table (Down then Across)', 'Pane (Across)', 'Pane (Down)', 'Pane (Across then Down)', 'Pane (Down then Across)', 'Cell', 'Category', 'Order Date', and 'Region'. In the background, a table shows quarterly sales data for 2009.

Year	Q1	Q2	Q3	Q4
2009	41,344	71,341	104,219	43,478
	118,807	128,046	182,627	150,971
	147,100	306,147	274,100	240,100
	74,210	151,456	197,708	167,720

Table calculation formulas much use aggregated data throughout the entire formula. When you are referring to a dimension in the formula you can use the `MIN([Dimension])` or `MAX([Dimension])` aggregations. However, this trick only works if the view is grouped by that dimension. That is, when there is a single dimension value for the row in question. For example, the view below shows several customer segments. Each segment corresponds to 4 regions. The `MIN([Region])` returns “Central” even though there are three other regions.

Customer Segment	Drop
Consumer	Central
Corporate	Central
Home Office	Central
Small Business	Central

You can instead use the `ATTR([Dimension])` aggregation. When you use `ATTR` the dimension value is used when you are grouping by the dimension. If there are multiple values it shows an asterisk. Nulls are ignored. The same view is shown below using `ATTR([Region])`.



Customer Segment	<i>Drop</i>
Consumer	*
Corporate	*
Home Office	*
Small Business	*

The ATTR aggregation is especially useful when you are working with multiple data sources on a single sheet.

---

## 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.

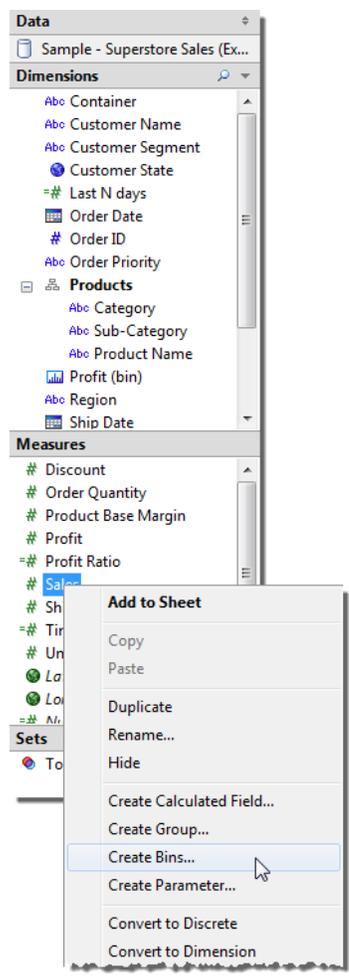
### Example – Creating a Histogram with Binned Data

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 measure in the Data window and select Create Bins on the right-click context menu.

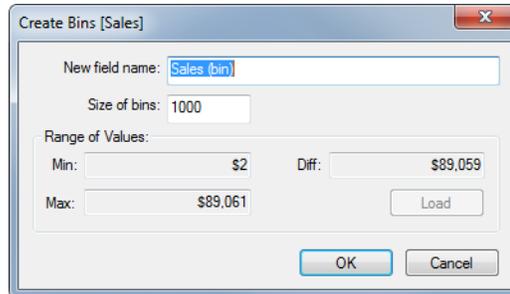


- 2 Complete the Create Bins dialog box.

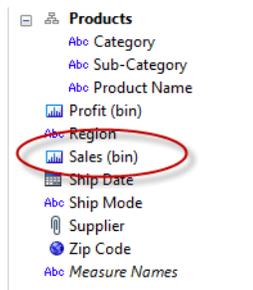
When you bin a measure, you create a new field. The new field is a binned version of the original field. Specify the name of the new field and the size of each bin. To help you



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** 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 (bin)** dimension on the **Columns** shelf.

Row headers are created with labels given by the dimension member names.



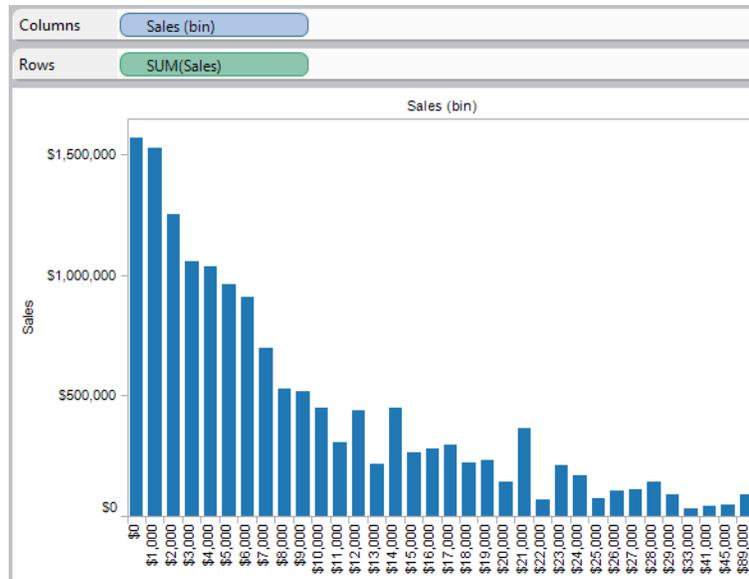
---

**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! 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
- 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.

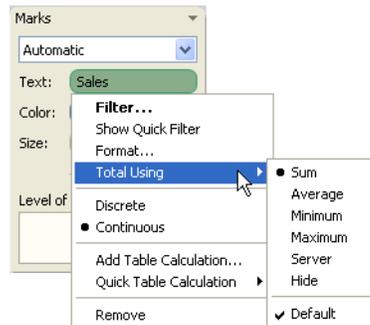
For example, let's say you have a view showing the sales of Amaretto, Columbian, and Decaf Irish Cream coffees. Then you filter the view to only show Amaretto and Columbian coffee sales. When you turn on subtotals for the Product Type field so you can see the total sales for all Coffees, one of the following 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.



Server computation is not always available and sometimes the totals will be blank for specific members in the view. When using server computation keep in mind the following information.

- Server computation is only available for ASO cubes.
- Server computation is not available for dynamic hierarchies. That means if the members in the view are part of a dimension or hierarchy that is tagged as dynamic, they cannot be included in the set of values you are using to compute the totals and will show up as blank in the view.
- If you are computing totals for a calculated field whose formula makes assumptions about other calculated members at different levels in the hierarchy, the totals will display as blank in the view.

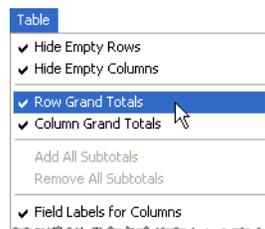
## Grand Totals

Any view in Tableau can include grand totals. For example, in a view showing the average profit for each product and year, you can turn on grand totals to also see the average profit for all products and all years.

Product 2 - Sub-Category	Order Date				Grand Total
	2006	2007	2008	2009	
Appliances	19,849	33,085	21,887	23,202	98,023
Binders and Binder Accesso..	123,814	44,426	44,769	94,165	307,174
Bookcases	-3,738	-542	-4,944	-23,947	-33,171
Chairs & Chairmats	55,817	19,310	40,187	34,625	149,938
Computer Peripherals	21,280	15,686	29,252	28,317	94,534
Copiers and Fax	33,389	50,521	58,896	24,556	167,361
Envelopes	7,451	13,914	12,697	14,649	48,711
Labels	2,623	2,361	5,704	3,001	13,689
Office Furnishings	24,035	28,456	26,216	18,171	96,878
Office Machines	47,143	100,290	71,300	57,858	276,590
Paper	17,135	12,930	4,280	11,643	45,987
Pens & Art Supplies	2,111	2,463	1,986	992	7,552
Rubber Bands	-2	-217	-121	263	-78
Scissors, Rulers and Trimme..	-2,906	-1,340	-1,822	-1,731	-7,799
Storage & Organization	5,143	8,045	-446	-10,226	2,517
Tables	-17,142	-23,350	-6,894	-13,686	-61,271
Telephones and Communica..	77,018	69,171	79,287	91,475	316,952
<b>Grand Total</b>	<b>413,021</b>	<b>375,209</b>	<b>382,234</b>	<b>353,125</b>	<b>1,523,589</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.



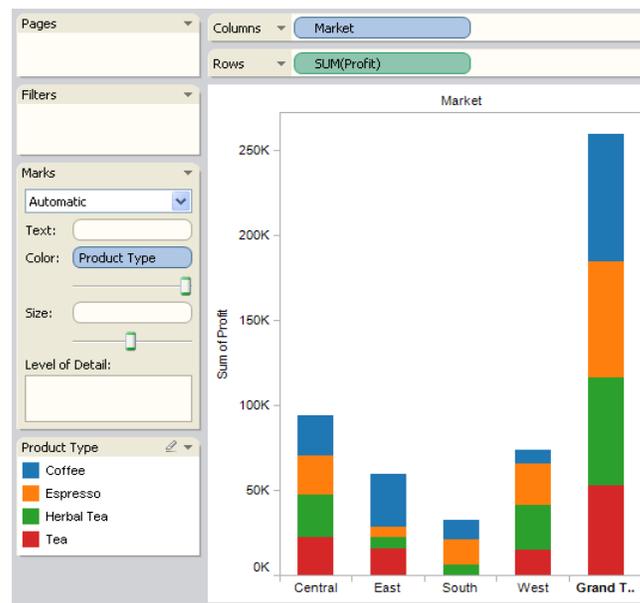
The following rules dictate whether you can turn on grand totals:

- 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 21-80 for more information.



- Grand Totals cannot be applied to continuous dimensions.

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 21-77 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(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 “View Data (Drill-Through)” on page 19-12 for more information.

The following table summarizes the standard aggregations and the grand totals that are calculated.

**Table 15-2: Grand Totals and Aggregations**

<b>Aggregation</b>	<b>Calculation Description</b>
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 aggregated as a sum. The grand total for the Jumbo Box shipped by Delivery Truck is \$2,814,305. 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, South, and West regions.



The screenshot shows the Tableau interface with the following configuration:

- Columns: Region
- Rows: Ship Mode, Container
- Marks: Text: SUM(Sales)

Ship Mode	Container	Region				Grand Total
		Central	East	South	West	
Delivery Truck	Jumbo Box	746,008	555,527	771,204	741,567	2,814,306
	Jumbo Drum	795,699	863,327	980,069	775,398	3,414,494
Express Air	Large Box	86,128	50,922	53,292	62,476	252,818
	Medium Box	35,354	15,596	37,047	29,401	117,398
	Small Box	180,677	106,604	213,880	205,835	706,996
	Small Pack	23,391	12,588	18,847	14,305	69,131
	Wrap Bag	5,103	9,800	10,051	12,726	37,680
Regular Air	Large Box	389,082	371,342	542,789	504,046	1,807,258
	Medium Box	149,191	136,713	151,470	139,519	576,893
	Small Box	1,070,918	905,200	1,163,709	1,244,812	4,384,638
	Small Pack	116,547	104,429	115,943	140,939	477,858
Wrap Bag	53,838	50,942	73,681	69,123	247,585	
<b>Grand Total</b>		3,651,938	3,182,989	4,131,981	3,940,146	<b>14,907,053</b>

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 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 aggregated as an average. The grand total for the Jumbo Box is \$5,350. 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.

		Region				Grand Total
Ship Mode	Container	Central	East	South	West	
Delivery Truck	Jumbo Box	\$5,695	\$5,192	\$5,211	\$5,297	\$5,350
	Jumbo Drum	\$5,604	\$5,794	\$5,537	\$5,101	\$5,607
Express Air	Large Box	\$5,383	\$3,917	\$4,099	\$3,905	\$4,359
	Medium Box	\$2,525	\$1,300	\$2,646	\$2,262	\$2,215
	Small Box	\$1,173	\$881	\$1,312	\$1,247	\$1,172
	Small Pack	\$585	\$672	\$673	\$433	\$562
	Wrap Bag	\$189	\$306	\$228	\$303	\$260
Regular Air	Large Box	\$5,380	\$5,305	\$5,539	\$4,820	\$5,254
	Medium Box	\$1,913	\$2,170	\$1,761	\$1,681	\$1,861
	Small Box	\$1,125	\$1,137	\$1,165	\$1,232	\$1,167
	Small Pack	\$574	\$561	\$537	\$621	\$574
	Wrap Bag	\$214	\$257	\$249	\$244	\$241
<b>Grand Total</b>		\$1,769	\$1,799	\$1,811	\$1,725	\$1,775

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 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 Data** on the context menu.



The screenshot shows a 'View Data' window with a table of sales data. The table has four columns: Container, Region, Ship Mode, and Sales. The data is as follows:

Container	Region	Ship Mode	Sales
Wrap Bag	East	Express Air	\$143
Wrap Bag	East	Express Air	\$146
Wrap Bag	East	Express Air	\$129
Wrap Bag	East	Express Air	\$814
Wrap Bag	East	Express Air	\$767
Wrap Bag	East	Express Air	\$151
Wrap Bag	East	Express Air	\$309
Wrap Bag	East	Express Air	\$21
Wrap Bag	East	Express Air	\$2,330
Wrap Bag	East	Express Air	\$410
Wrap Bag	East	Express Air	\$1,341
Wrap Bag	East	Express Air	\$76
Wrap Bag	East	Express Air	\$1,032
Wrap Bag	East	Express Air	\$435
Wrap Bag	East	Express Air	\$42
Wrap Bag	East	Express Air	\$103

## 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.



Product Type	Product	Market			
		Central	East	South	West
Coffee	Amaretto	14,012	2,994		9,263
	Columbian	28,911	47,385	21,663	30,352
	Decaf Irish Cream	26,157	6,262	11,596	18,233
	<b>Total</b>	69,080	56,641	33,259	57,848
Espresso	Caffe Latte			15,443	20,456
	Caffe Mocha	35,218	16,646	14,166	18,874
	Decaf Espresso	24,483	7,720	15,381	30,578
	Regular Espresso		24,031		
<b>Total</b>	59,701	48,397	44,990	69,908	
Herbal Tea	Chamomile	36,571	2,193	11,183	25,631
	Lemon	21,982	27,177	14,494	32,273
	Mint	9,335	11,991		14,384
	<b>Total</b>	67,888	41,361	25,677	72,288
Tea	Darjeeling	30,284	14,094		28,773
	Earl Grey	32,883	6,507		27,382
	Green Tea	5,209	11,576		16,065
	<b>Total</b>	68,376	32,177		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.

	Coffee					Espresso				
	Central	East	South	West	Total	Central	East	South	West	Total
Amaretto	14,011	2,993		9,265	26,269					
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	30,357	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 21-77.

## Percentages

Any analysis in Tableau can be expressed in terms of percentages. For example, rather than viewing sales 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 21-40 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 21-88 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 21-91 for more information.



The figure below is an example of a text table with percentages. The percentages are calculated with the **Sales** measure aggregated as a summation, and are based on the entire table.

Columns		+ YEAR(Order Date)			
Rows		Product 1 - Category		Region	
Product 1 - Category	Region	Order Date			
		2006	2007	2008	2009
Furniture	Central	2.695%	2.185%	2.024%	1.964%
	East	1.805%	1.687%	2.042%	2.052%
	South	2.328%	2.328%	2.464%	1.968%
	West	2.870%	2.092%	1.903%	1.934%
Office Supplies	Central	1.206%	1.688%	1.525%	1.859%
	East	1.691%	1.112%	1.137%	1.265%
	South	2.188%	1.471%	1.360%	1.622%
	West	1.963%	1.904%	1.395%	2.104%
Technology	Central	2.496%	2.609%	1.864%	2.661%
	East	1.860%	2.313%	2.570%	1.817%
	South	3.777%	2.204%	2.354%	3.652%
	West	3.228%	2.291%	2.413%	2.056%

## 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 21-4 for more information.

For example, if the aggregation applied to the **Sales** field is a summation, then the default percentage calculation (percent of table) means that each number displayed is the SUM(Sales) for that mark divided by the SUM(Sales) 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 21-34 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.

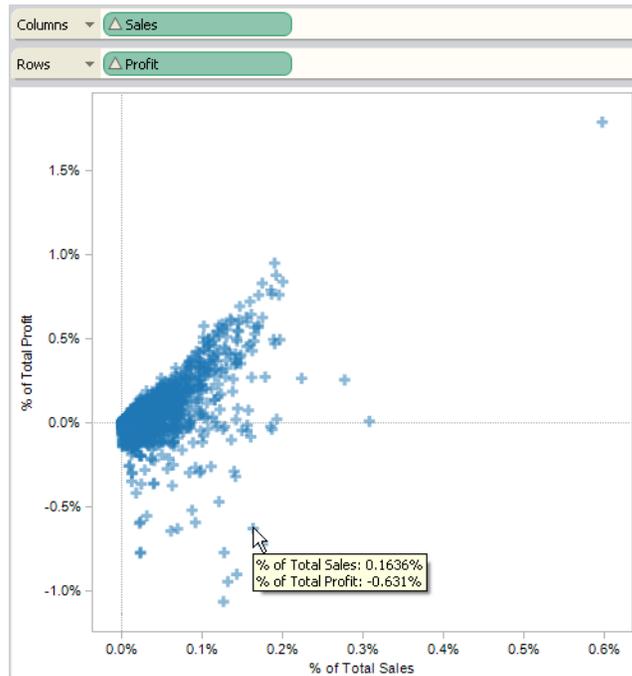
The tooltip reveals that the maximum sales for furniture in the east in 2005 is 31.6% of the maximum for the entire table. What is the maximum for the table? You can see by inspection that the maximum occurs in the South in the year 2005. The tooltip for this bar segment would reveal a maximum sales 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 profit of -0631% and a sales of 0.1636%. 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 21-78 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 2.50% of total sales in 2006. The grand totals for rows show that 2009 accounts for 24.95% of the total sales. Summing the grand totals for rows or for columns yields 100% of the total.

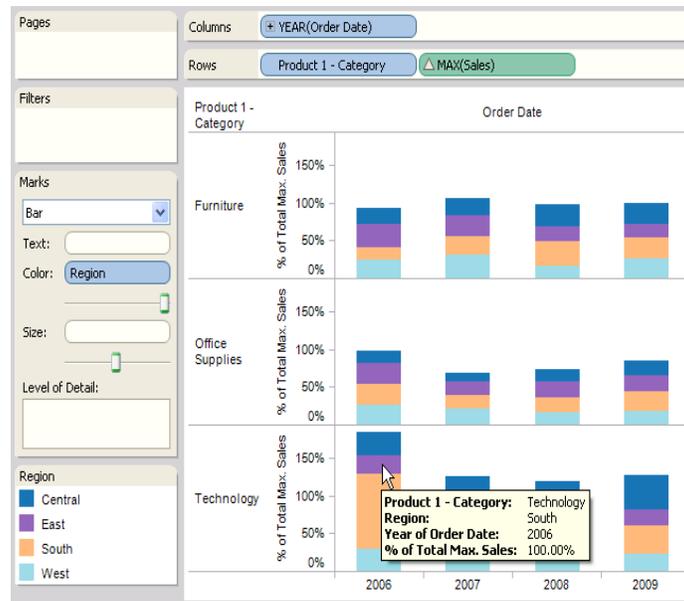


Columns: + YEAR(Order Date)

Rows: Product 1 - Category, Region

Product 1 - Category	Region	Order Date				Grand Total
		2006	2007	2008	2009	
Furniture	Central	2.69%	2.19%	2.02%	1.96%	8.87%
	East	1.81%	1.69%	2.04%	2.05%	7.59%
	South	2.33%	2.33%	2.46%	1.97%	9.09%
	West	2.87%	2.09%	1.90%	1.93%	8.80%
Office Supplies	Central	1.21%	1.69%	1.53%	1.86%	6.28%
	East	1.69%	1.11%	1.14%	1.27%	5.21%
	South	2.19%	1.47%	1.36%	1.62%	6.64%
	West	1.96%	1.90%	1.40%	2.10%	7.37%
Technology	Central	2.50%	2.61%	1.86%	2.66%	9.63%
	East	1.86%	2.31%	2.57%	1.82%	8.56%
	South	3.78%	2.20%	2.35%	3.65%	11.99%
	West	3.23%	2.29%	2.41%	2.06%	9.99%
<b>Grand Total</b>		28.11%	23.89%	23.05%	24.95%	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 6.62% of total sales in the 2006. 2006 is the column in this case.

Columns		+ YEAR(Order Date)				
Rows		Product 1 - Category	Region			
Product 1 - Category	Region	Order Date				Grand Total
		2006	2007	2008	2009	
Furniture	Central	9.59%	9.15%	8.78%	7.87%	8.87%
	East	6.42%	7.06%	8.86%	8.22%	7.59%
	South	8.28%	9.75%	10.69%	7.89%	9.09%
	West	10.21%	8.76%	8.26%	7.75%	8.80%
Office Supplies	Central	4.29%	7.07%	6.62%	7.45%	6.28%
	East	6.02%	4.66%	4.93%	5.07%	5.21%
	South	7.78%	6.16%	5.90%	6.50%	6.64%
	West	6.98%	7.97%	6.05%	8.43%	7.37%
Technology	Central	8.88%	10.92%	8.08%	10.66%	9.63%
	East	6.62%	9.69%	11.15%	7.28%	8.56%
	South	13.44%	9.23%	10.21%	14.63%	11.99%
	West	11.49%	9.59%	10.47%	8.24%	9.99%
<b>Grand Total</b>		100.00%	100.00%	100.00%	100.00%	100.00%

## 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 21.72% of technology sales in the east across all years. East is the row in this case.



Product 1 - Category	Region	Order Date				Grand Total
		2006	2007	2008	2009	
Furniture	Central	30.39%	24.64%	22.83%	22.14%	100.00%
	East	23.80%	22.24%	26.91%	27.05%	100.00%
	South	25.62%	25.62%	27.12%	21.65%	100.00%
	West	32.62%	23.77%	21.63%	21.98%	100.00%
Office Supplies	Central	19.21%	26.89%	24.29%	29.61%	100.00%
	East	32.49%	21.36%	21.85%	24.31%	100.00%
	South	32.94%	22.15%	20.48%	24.42%	100.00%
	West	26.64%	25.85%	18.94%	28.56%	100.00%
Technology	Central	25.92%	27.09%	19.36%	27.64%	100.00%
	East	21.72%	27.02%	30.03%	21.22%	100.00%
	South	31.51%	18.39%	19.64%	30.46%	100.00%
	West	32.32%	22.94%	24.16%	20.59%	100.00%
<b>Grand Total</b>		28.11%	23.89%	23.05%	24.95%	100.00%

## Percent of Pane

When you select **Percentage 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 several 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 13-1.



The screenshot shows the Tableau interface with the following configuration:

- Columns:** YEAR(Order Date)
- Rows:** Product 1 - Category, Region
- Marks:** Automatic, SUM(Sales)

Product 1 - Category	Region	Order Date				Grand Total
		2006	2007	2008	2009	
Furniture	Central	7.85%	6.36%	5.90%	5.72%	25.82%
	East	5.26%	4.91%	5.94%	5.98%	22.09%
	South	6.78%	6.78%	7.18%	5.73%	26.47%
	West	8.36%	6.09%	5.54%	5.63%	25.62%
Office Supplies	Central	4.73%	6.62%	5.98%	7.29%	24.63%
	East	6.63%	4.36%	4.46%	4.96%	20.42%
	South	8.58%	5.77%	5.34%	6.36%	26.05%
	West	7.70%	7.47%	5.47%	8.25%	28.90%
Technology	Central	6.21%	6.49%	4.64%	6.62%	23.97%
	East	4.63%	5.76%	6.40%	4.52%	21.31%
	South	9.40%	5.49%	5.86%	9.09%	29.85%
	West	8.04%	5.70%	6.01%	5.12%	24.87%
<b>Grand Total</b>		28.11%	23.89%	23.05%	24.95%	100.00%

### 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) and SUM(Profit).

The grand total for columns is always 100%. This represents the sum of a given row within a pane.



Columns		+ YEAR(Order Date)				
Rows		Product 1 - Category	Region			
Product 1 - Category	Region	Order Date				Grand Total
		2006	2007	2008	2009	
Furniture	Central	30.39%	24.64%	22.83%	22.14%	100.00%
	East	23.80%	22.24%	26.91%	27.05%	100.00%
	South	25.62%	25.62%	27.12%	21.65%	100.00%
	West	32.62%	23.77%	21.63%	21.98%	100.00%
Office Supplies	Central	19.21%	26.89%	24.29%	29.61%	100.00%
	East	32.49%	21.36%	21.85%	24.31%	100.00%
	South	32.94%	22.15%	20.48%	24.42%	100.00%
	West	26.64%	25.85%	18.94%	28.56%	100.00%
Technology	Central	25.92%	27.09%	19.36%	27.64%	100.00%
	East	21.72%	27.02%	30.03%	21.22%	100.00%
	South	31.51%	18.39%	19.64%	30.46%	100.00%
	West	32.32%	22.94%	24.16%	20.59%	100.00%
<b>Grand Total</b>		28.11%	23.89%	23.05%	24.95%	100.00%



### 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) and SUM(Profit).

---

The grand total for rows is always 100%. This represents the sum of a given column within a pane.

		Order Date				
Product 1 - Category	Region	2006	2007	2008	2009	Grand Total
Furniture	Central	27.79%	26.35%	24.00%	24.80%	25.82%
	East	18.61%	20.34%	24.21%	25.92%	22.09%
	South	24.01%	28.08%	29.22%	24.85%	26.47%
	West	29.59%	25.23%	22.57%	24.42%	25.62%
Office Supplies	Central	17.11%	27.34%	28.15%	27.14%	24.63%
	East	24.00%	18.00%	20.99%	18.47%	20.42%
	South	31.05%	23.82%	25.11%	23.68%	26.05%
	West	27.85%	30.84%	25.75%	30.71%	28.90%
Technology	Central	21.97%	27.70%	20.25%	26.12%	23.97%
	East	16.37%	24.57%	27.94%	17.84%	21.31%
	South	33.25%	23.41%	25.59%	35.85%	29.85%
	West	28.42%	24.33%	26.22%	20.19%	24.87%
<b>Grand Total</b>		100.00%	100.00%	100.00%	100.00%	100.00%



# Parameters

---

<b>Overview . . . . .</b>	<b>22-2</b>
<b>Creating Parameters . . . . .</b>	<b>22-3</b>
<b>Editing Parameters . . . . .</b>	<b>22-9</b>
<b>Using Parameters in Calculations . . . . .</b>	<b>22-10</b>
<b>Parameter Controls . . . . .</b>	<b>22-12</b>
<b>Example - Parameters . . . . .</b>	<b>22-14</b>



## Overview

Parameters are dynamic values that can replace constant values in calculations. For example, you may create a calculated field that returns true if Sales is greater than \$500,000 and otherwise return false. You can replace the constant value of “500000” in the formula with a parameter. Then using the parameter control you can dynamically change the threshold in your calculation. This section discusses creating parameters and using them in calculations in the following sections:

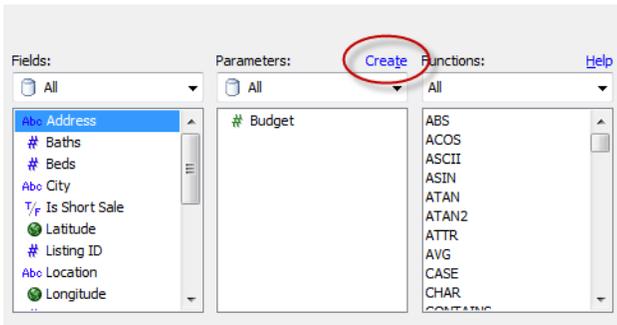
- Creating Parameters
- Editing Parameters
- Using Parameters in Calculations
- Parameter Controls
- Example - Parameters



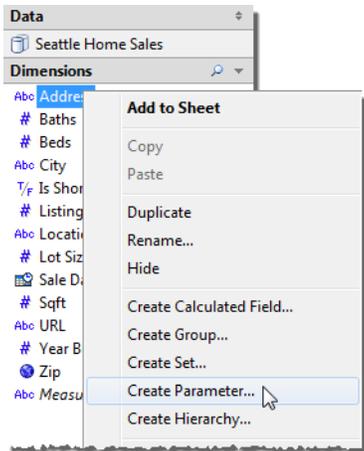
# Creating Parameters

You can create a new parameter based on a selected field, or you can create a new parameter from the Calculated Field dialog box. Follow the instructions below to create a new parameter.

- 1 Do one of the following:
  - In the Calculated field dialog box, click the **Create** link at the top of the list of parameters.

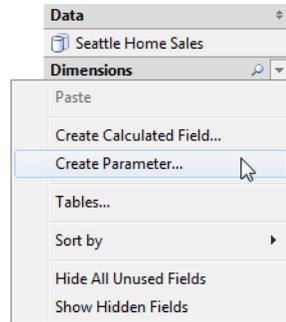


- In the Data window, right-click a field to base the parameter on and select **Create Parameter**.

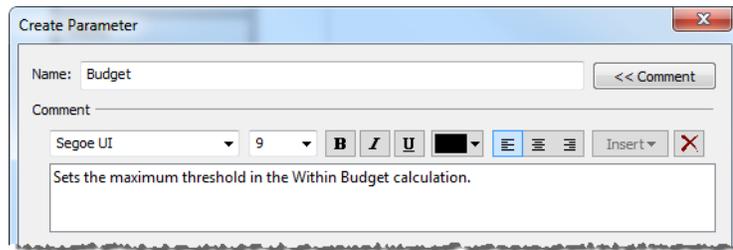




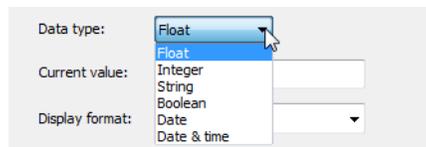
- Open the Data window menu using the drop-down arrow in the upper right corner and select **Create Parameter**.



- 2 In the Create Parameter dialog box, give the field a **Name** and optionally write a **Comment** to describe the parameter.



- 3 Specify the data type for the values it will accept. Parameters can be a float (non integer number), integer, string (text), boolean (true/false), date, or date & time.



- 4 Specify a current value. This is the default value for the parameter.



Current value: 10/27/2010

5 Specify the display format to use in the parameter control.

Display format: Automatic

Allowable values:

- Automatic
- Standard Long Date
- Standard Short Date
- 3/14/2001
- 3/14/01
- 03/14/01
- 03/14/2001
- 01/03/14
- 2001-03-14
- 14-Mar-01
- Wednesday, March 14, 2001
- March 14, 2001
- Wed, 14 March, 2001
- 14 March, 2001
- March, 2001
- 2001 Q1
- 2001 W11
- Custom

Automatic

Locale: Automatic

- 6 Specify how the parameter will accept values. You can select from the following options:
- All - the parameter control is a simple type in field.
  - List - the parameter control provides a list of possible values for you to select from.
  - Range - the parameter control lets you select values within a specified range.

The availability of these options is determined by the data type. For example, a string parameter can only accept all values or a list. It does not support a range.

If you select List, you must specify the list of values. Click in the left column to type a value. Each value can also have a display alias. You can copy and paste a list of values



by clicking Paste from Clipboard. Alternatively you can Add the members of a field as the list of values by selecting Add from Field.

Allowable values:  All  List  Range

List of values

Value	Display As
1	Apples
2	Bananas
3	Oranges
4	Pears
Add	

Add from Parameter ▶

Add from Field ▶

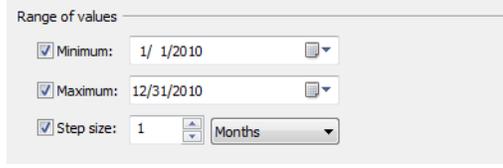
Paste from Clipboard

Clear All

If you select Range you must specify a minimum, maximum and step size. For example, you can define a date range between January 1, 2010 and December 31, 2010, with the

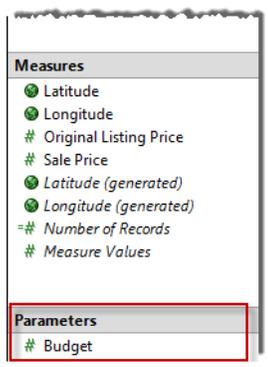


step size set to 1 month to create a parameter control that lets you select each month in 2010.



7 When finished, click **OK**.

The parameter is listed in the Parameters section at the bottom of the Data window.



It is also available in the Parameters list in the Calculated Field dialog box. Parameters are global across the workbook and can be used in any worksheet.



Fields: [All](#)

- # Address
- # Baths
- # Beds
- Abc City
- T/f Is Short Sale
- Latitude
- # Listing ID
- Abc Location
- Longitude

Parameters: [Create](#)

- # Budget

Functions: [Help](#)

- ABS
- ACOS
- ASCII
- ASIN
- ATAN
- ATAN2
- ATTR
- AVG
- CASE
- CHAR
- CONTAINS



## Editing Parameters

You can edit parameters from the Data window or the parameter control. Follow the instructions below to edit a parameter:

- 1 Do one of the following:
  - Right-click the parameter in the Data window and select **Edit**.
  - Select **Edit Parameter** on the parameter control card menu.
- 2 In the Edit Parameter dialog box, make the modifications as necessary.
- 3 When finished, click **OK**. The parameter is updated along with any calculations that use it.

---

**Note** You can delete a parameter by right-clicking it in the Data window and selecting **Delete**. Any calculated fields that use the deleted parameter will become invalid.

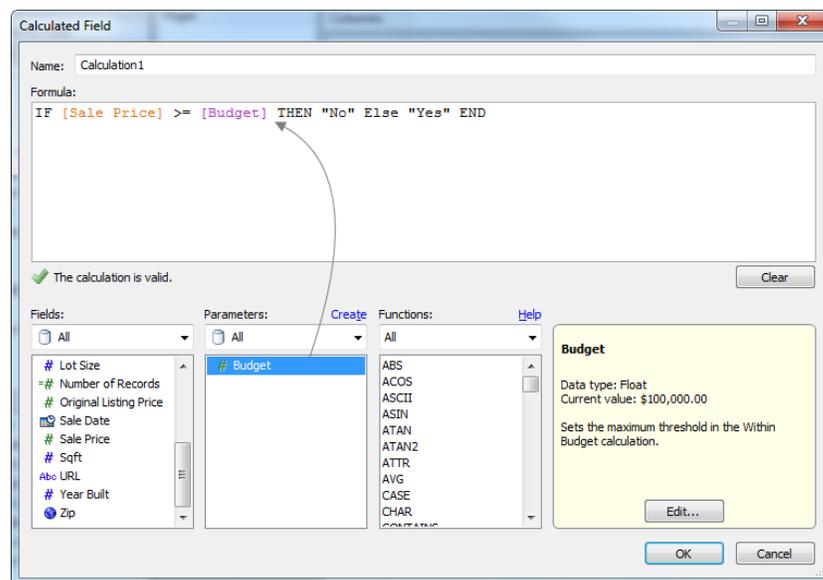
---

---

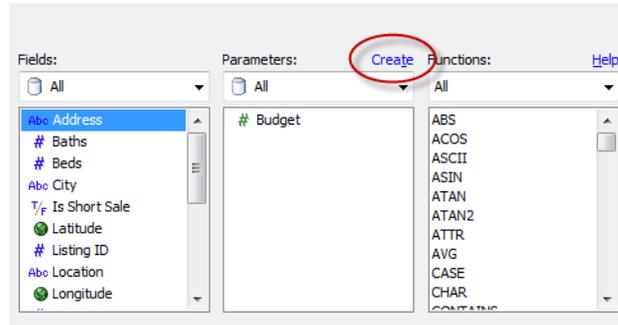
## Using Parameters in Calculations

Parameters give you a way to dynamically modify values in a calculation. Rather than manually editing the calculation (and all dependent calculations), you can use a parameter. Then when you want to change the value, you open the parameter control, change the value, and all of the calculations that use that parameter are updated.

A list of parameters are available in the Calculated Field dialog box. Simply double-click the parameter in the Parameters list to add it to your formula.



Because you often don't realize you need a parameter until you are authoring the calculation, you can create parameters directly from the Calculated Field dialog box. Click the **Create** link at the top of the Parameters list. Refer to "Creating Parameters" on page 22-3 for detailed instructions on creating parameters.



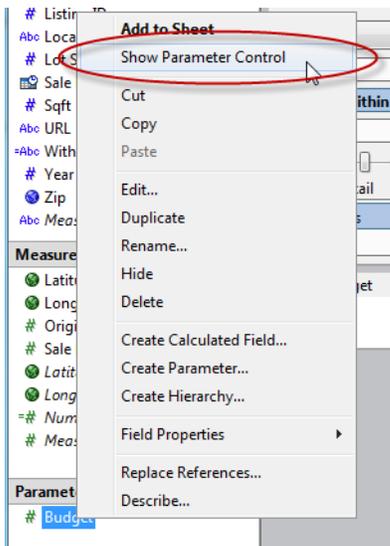
Similar to fields, parameters are written within square brackets in the formula. Parameters are shown as purple in the formula editor.



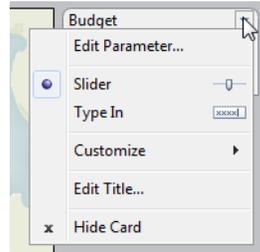
# Parameter Controls

The parameter control is a worksheet card that lets you modify the parameter value. Parameter controls are very similar to quick filter cards in that they contain controls that modify the view. You can open parameter controls on worksheets and dashboards and they are included when you save to the web or publish to Tableau Server.

Open the parameter control by right-clicking the parameter in the Data window and selecting Show Parameter Control.



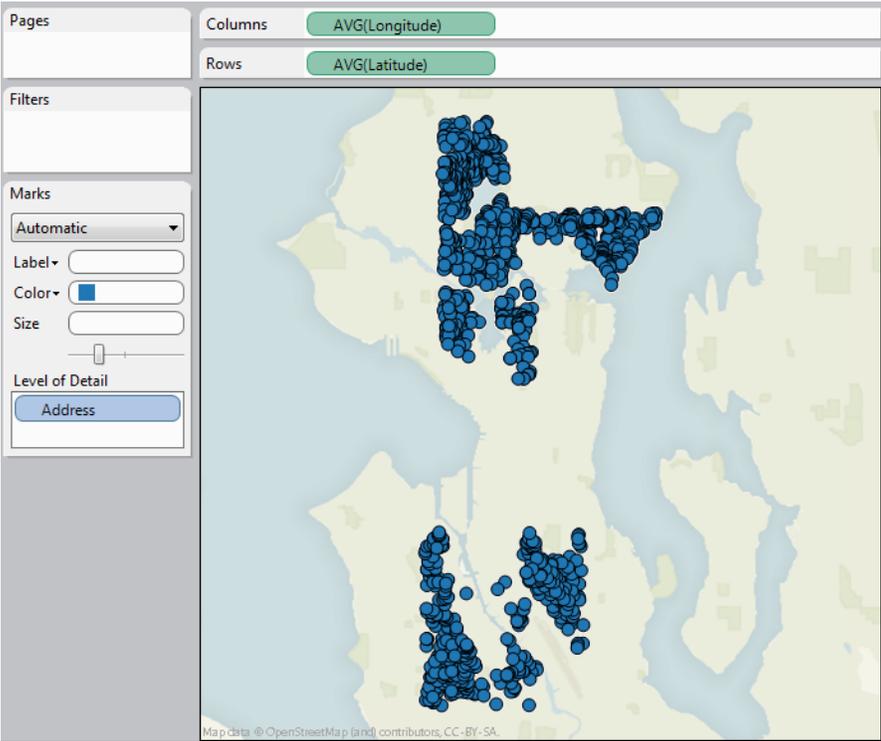
Like other cards, Parameter Controls have a menu that you can open using the drop-down arrow in the upper right corner of the card. Use this menu to customize the display of the control. For example, you can show a list of values as radio buttons, a compact list, a slider, or a type in field. The options available on this menu depend on the data type of the parameter as well as whether it accepts All, a List, or a Range of values



---

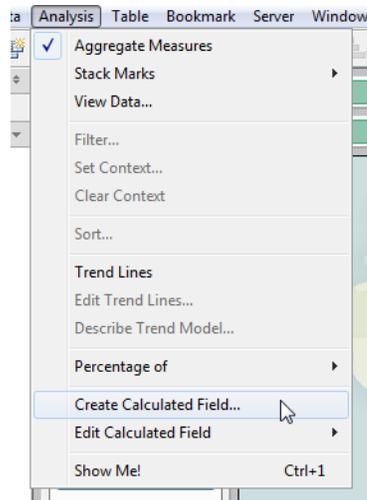
## Example - Parameters

This example uses a data source that contains home sale information. The example below shows a view of recently sold houses in the Seattle area.



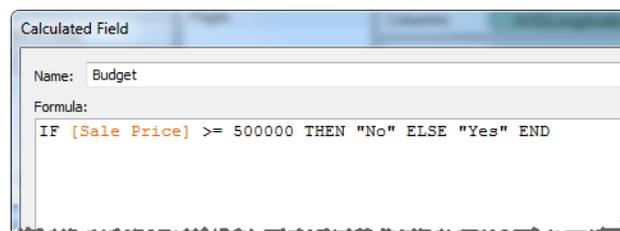
Follow the steps below to create a parameter that is used to color the houses based whether they fit into a specific budget.

- 1 Open the calculated field dialog box by selecting **Analysis > Create Calculated Field**.



- 2 Name the field "Within Budget."
- 3 In the formula type a boolean expression using the IF function. The formula is shown below:

```
IF [Sale Price] >= 500000 THEN "No" ELSE "Yes" END
```

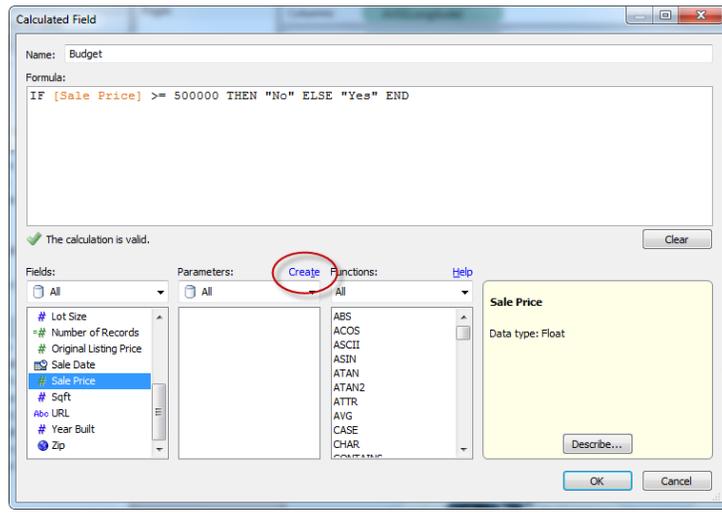


This formula will create a field that returns No if the sale price is over a budget of 500000. However, this threshold should be dynamic so that you can quickly change the

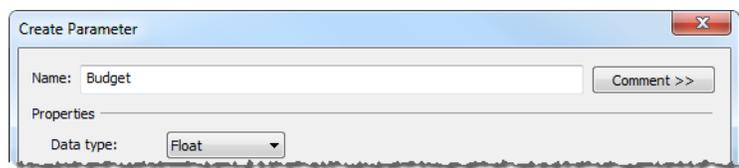


budget without editing the calculation. We need to replace the constant value with a parameter.

- 4 Click the **Create** link at the top of the Parameters list in the Calculated Field dialog box.



- 5 Name the parameter "Budget."



- 6 Select the Float data type because the values will be currency.



- 7 Set the **Current Value** to 500000. This is the default value for the parameter and just gives it a starting point.



Current value:

8 Change the display format to **Currency (Standard)**.

Create Parameter

Name: Budget Comment >>

Properties

Data type: Float

Current value: 500,000

Display format: \$500,000.00

Allowable values:

- Automatic
- Number (Standard)
- Number (Custom)
- Currency (Standard)
- Currency (Custom)
- Scientific
- Percentage
- Custom

Currency (Standard)

Locale: English (United States)

9 Allow a Range of values. Set the **Minimum** to 100,000 and the **Maximum** to 3,000,000. Set the **Step size** to 25,000.

This definition means that the budget can be between \$100K and \$3M in increments of \$25K.

Allowable values:  All  List  Range

Range of values

Minimum: 100,000

Maximum: 3,000,000

Step size: 25,000

- 
- 10 When finished, click **OK**.

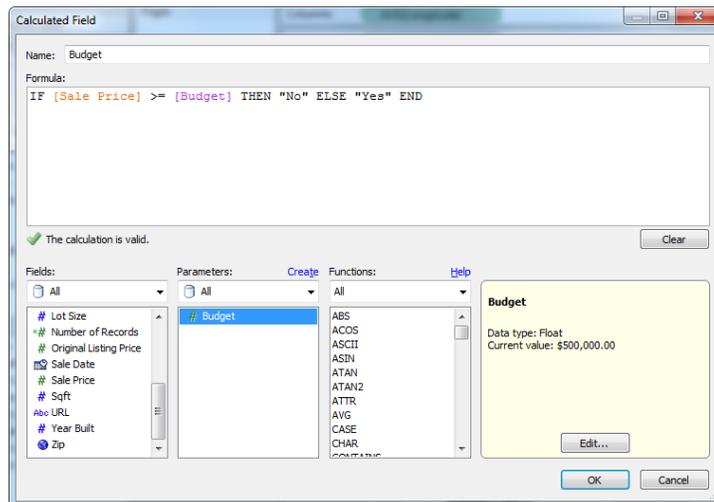
The screenshot shows the 'Create Parameter' dialog box with the following settings:

- Name: Budget
- Data type: Float
- Current value: 500,000
- Display format: \$500,000.00
- Allowable values: Range (selected)
- Range of values:
  - Minimum: 100,000
  - Maximum: 3,000,000
  - Step size: 25,000

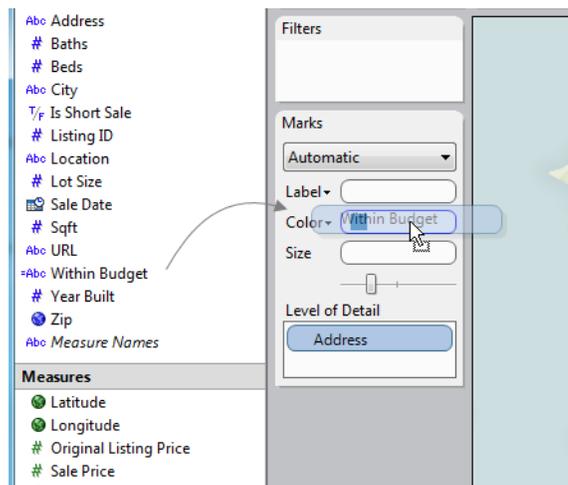
The 'OK' button is circled in red.

- 11 Back in Calculated Field dialog box, delete the constant value of “5000000” in the formula. Then double-click the Budget parameter in the Parameters list to add it to the formula. The formula now looks like this:

```
IF [Sale Price] >= [Budget] THEN "No" ELSE "Yes" END
```

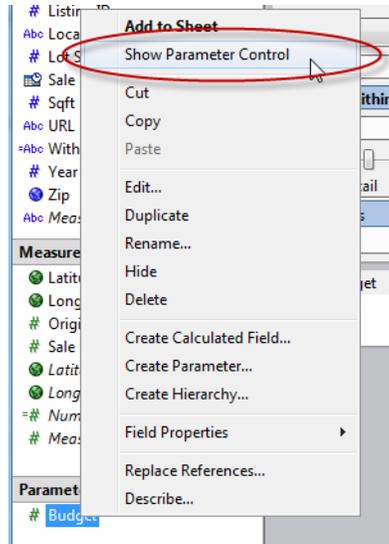


- 12 Make sure the calculation is valid and then click **OK**.
- 13 The new field is added to the Dimensions area of the Data window. Drag the **Within Budget** field to the **Color** shelf.

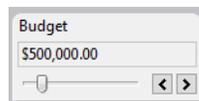




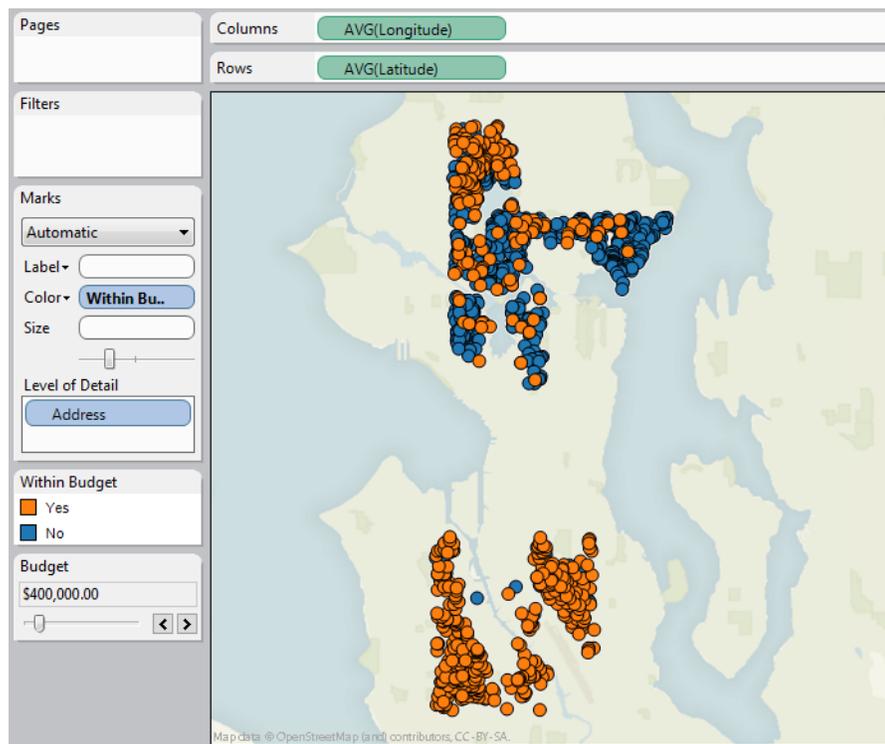
- 14 Right-click the Budget parameter at the bottom of the Data window and select Show Parameter control.



By default the control is shown as a slider. Clicking the right arrow in the control will increase the budget by \$25K.



The view below shows the budget set to \$400,000. Orange house are within budget and blue are not.



# Background Images

---

<b>Overview</b> . . . . .	<b>23-2</b>
<b>Adding Background Images</b> . . . . .	<b>23-3</b>
<b>Setting up the View</b> . . . . .	<b>23-7</b>
<b>Managing Background Images</b> . . . . .	<b>23-8</b>
Editing an Image . . . . .	23-8
Enabling/Disabling Images . . . . .	23-9
Adding Show/Hide Conditions . . . . .	23-10
Removing an Image . . . . .	23-13

---

## 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 custom map images that correspond to a coordinate system in your data. For example, you may have data that corresponds to several floors in a building. You can use background images to overlay that data on the actual floor plan of the building to give more context. Other examples of using background images include showing a model of the sea floor, images of web pages for analyzing web logs, and even levels from video games to visualize player statistics.

While Tableau allows you to load dynamic maps from the online and offline provider, background images allow you to use your own custom images whether they are special maps or any other image that corresponds to your data.

This section discusses the following topics:

- Adding Background Images
- Setting up the View
- Managing Background Images

Refer to “Background Maps” on page 24-1 to learn more about using dynamic maps.

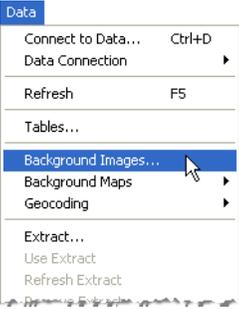


# Adding Background Images

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 fields in your database. 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.

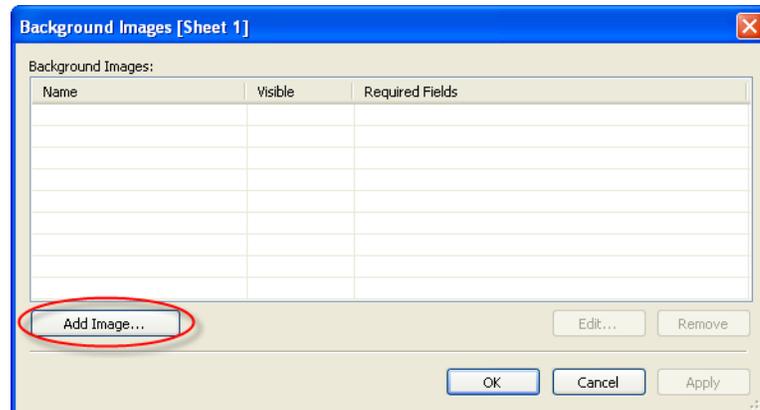
**To add a background image:**

- 1 Select **Data > Background Images**.





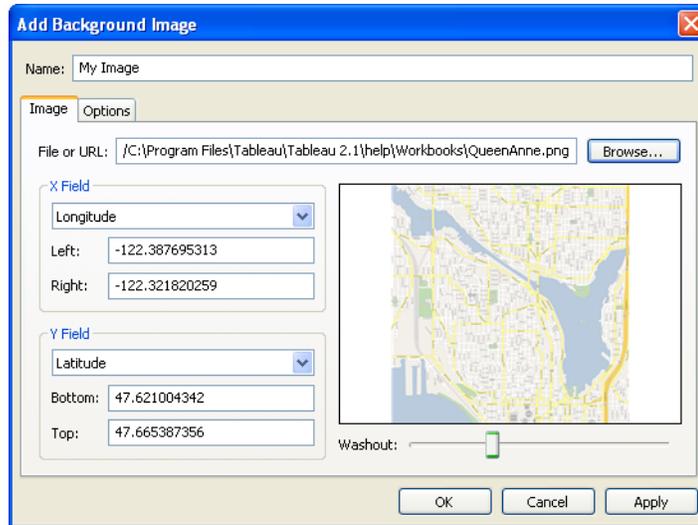
- 2 In the Background Images dialog box, click **Add Image**.



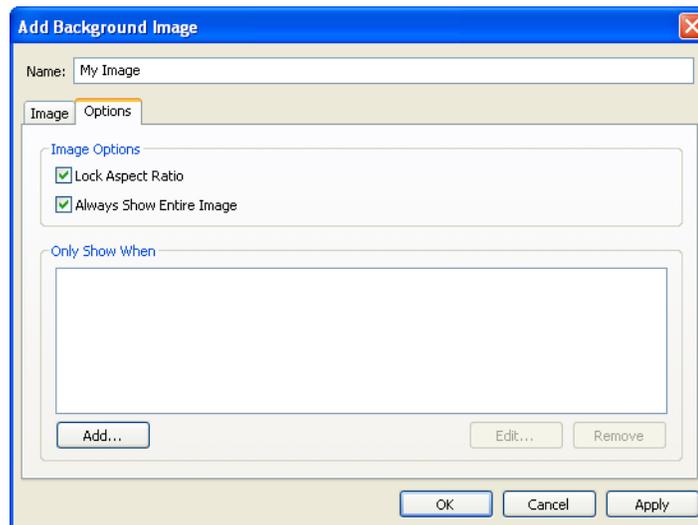
- 3 In the Add Background Image dialog box do the following:
  - Type a name for the image into the **Name** text box.
  - Click **Browse** to navigate to and select the image you want to add to the background. You can also type a URL to link to an image hosted online.
  - Select the field to map to 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 23-10 to learn more about defining conditions.



5 Click **OK**.

Refer to “Setting up the View” on page 23-7 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 13-30 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. Follow the steps below to set up the view correctly:

- 1 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.

- 2 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.

---

## 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:

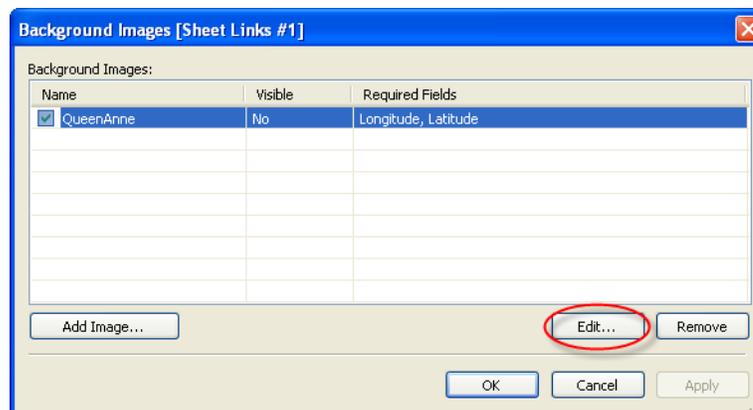
- Editing an Image
- Enabling/Disabling Images
- Adding Show/Hide Conditions
- Removing an Image

### 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**.

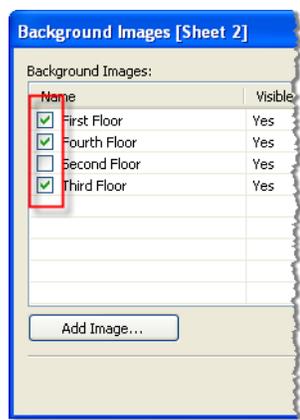
### Enabling/Disabling Images

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**.

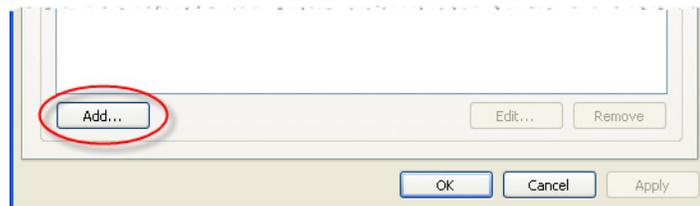


## Adding Show/Hide Conditions

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.

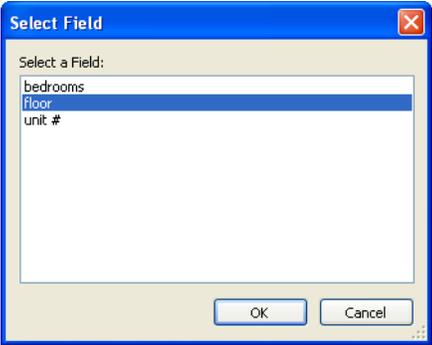
**To specify show/hide conditions:**

- 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.



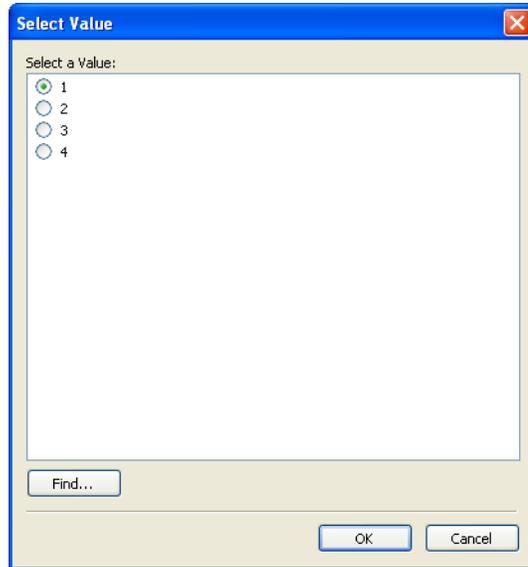


5 Select a field to base the condition on. In the example described above, the field is Floor.





- Specify when you want to show the image by selecting a value of the field. For this example, one is selected.



- 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.

- Click **OK** twice to close the Background Image dialog boxes and apply the changes.

---

**Note** When you add multiple conditions the background image will only show when all conditions are met. For example, if a background image has two conditions on Property Name and Floor, it will only show when Property is Greenwood Estates and Floor is 3.

---

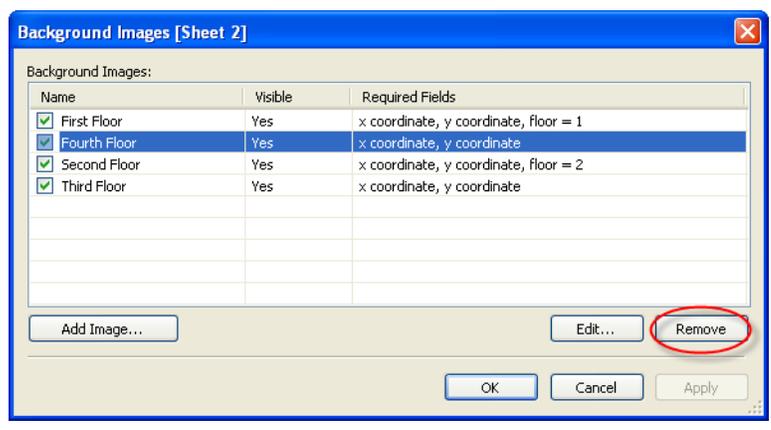


## Removing an Image

When you no longer want to use a background image you can either disable it (refer to “Enabling/Disabling Images” on page 23-9) 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**.



# Background Maps

---

<b>Overview</b> . . . . .	<b>24-2</b>
<b>Geographic Roles</b> . . . . .	<b>24-3</b>
<b>Building a Map View</b> . . . . .	<b>24-6</b>
<b>Map Options</b> . . . . .	<b>24-7</b>
Map Layers . . . . .	24-8
Data Layers . . . . .	24-10
Washout . . . . .	24-11
<b>Setting a Default Location</b> . . . . .	<b>24-13</b>
<b>Editing Locations</b> . . . . .	<b>24-14</b>
<b>Custom Geocoding</b> . . . . .	<b>24-16</b>
Creating an Import File . . . . .	24-16
Importing Custom Geocoding . . . . .	24-20
Saving Custom Geocoding . . . . .	24-22
<b>Background Map Sources</b> . . . . .	<b>24-23</b>
<b>Map Storing and Working Offline</b> . . . . .	<b>24-26</b>

---

## Overview

When you are connected to data that contains location information you can automatically show the data on an online or offline map. Once a map is loaded, you can zoom in and out, pan around the map, or focus on specific regions using the Navigation tools. This section discusses how to build a view using Tableau’s background maps feature in the following topics:

- Geographic Roles
- Building a Map View
- Map Options
- Setting a Default Location
- Editing Locations
- Custom Geocoding
- Background Map Sources
- Map Storing and Working Offline

If you would like to overlay your data on a static image instead you can use a background image. Refer to “Background Images” on page 23-1 to learn how.

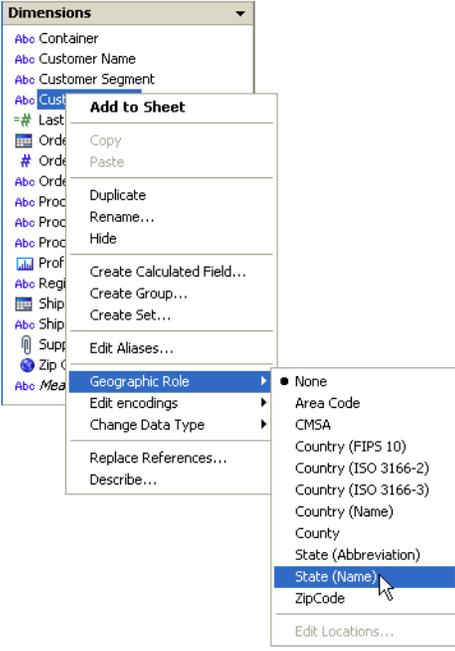


# Geographic Roles

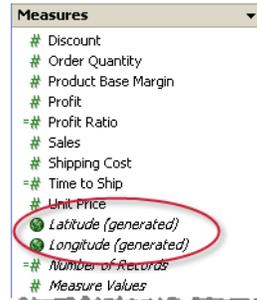
The first step to building a map view is to specify the fields that contain the location data. Tableau will automatically assign geographic roles to fields with common location names such as State, Country, and so on. You can manually assign geographic roles to fields that don't use common names and weren't automatically detected.

**To assign a geographic role:**

- 1 Right-click the field in the Data window that contains geographic data and select **Geographic Role**. Then select the type of data the field contains.



When you assign a field a geographic role it is marked in the Data window with a globe icon 🌐. What that means is that Tableau has automatically geocoded the information in that field and associated each value with a latitude and longitude value. Two automatic fields are added to the Measures area of the Data window: Latitude (generated) and Longitude (generated). These fields are available anytime you have use Tableau to geocode your data.



Below is a table describing the types of information that Tableau can geocode.

<b>Geographic Role</b>	<b>Description</b>
Area Code	U.S. Area Codes; numbers only. <i>Example: 206, 650, 415</i>
CMSA	U.S. Consolidated Metropolitan Statistical Area. <i>Example: Dallas-Fort Worth-Arlington</i>
Country (FIPS 10)	Worldwide countries using the U.S. Federal Information Processing Standard No. 10. <i>Example: AF, CD, IZ</i>
Country (ISO 3166-2)	Worldwide countries using the alpha 2 of the 3166 standard set by the International Organization for Standardization. <i>Example: AF, BH, UA</i>
Country (ISO 3166-3)	Worldwide countries using the alpha-3 of the 3166 standard set by the International Organization for Standardization. <i>Example: AFG, BHR, UKR</i>
Country (Name)	Worldwide country names in English. <i>Example: United States, Japan, Italy</i>
County	U.S. county names. <i>Example: King, Clark, Oneida</i>
State (Abbreviation)	Worldwide states and province abbreviations. <i>Example: WA, BC, AB</i>

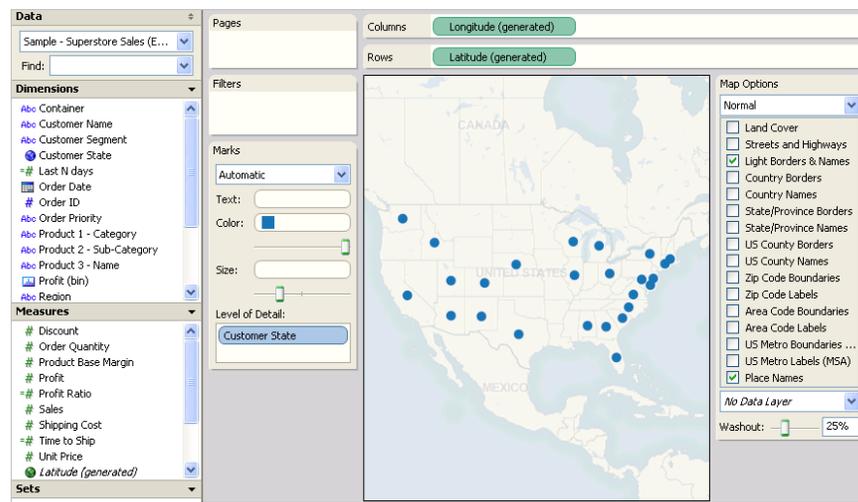


<b>Geographic Role</b>	<b>Description</b>
State (Name)	Worldwide state and province names. <i>Example: California, Alberta, Berlin</i>
ZipCode	U.S. 5 digit zip codes. <i>Example: 94404</i>

---

## Building a Map View

You can use the generated latitude and longitude fields to build map views simply by double-clicking a geographic field in the Data window. The Show Me! double-click rules will automatically add the generated Latitude and Longitude fields to the shelves and place the geographic field on the **Level of Detail** shelf.



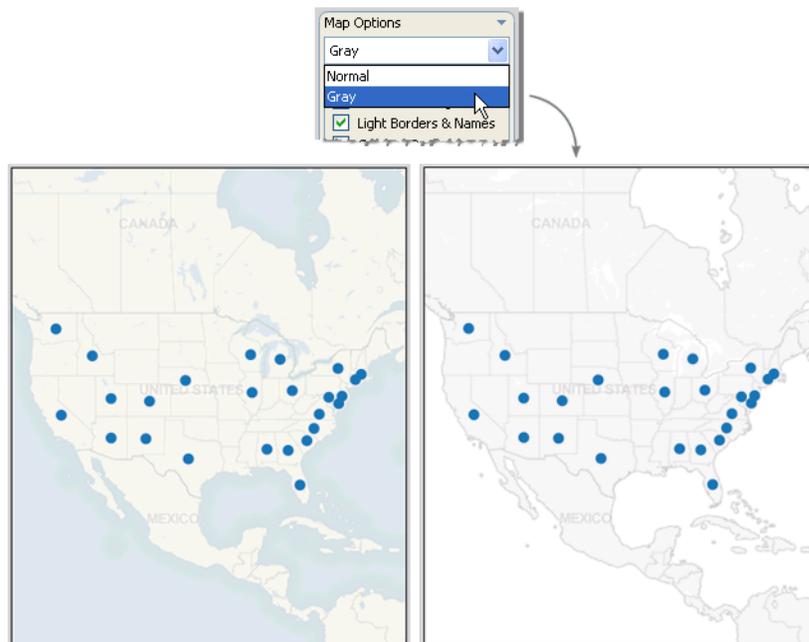
The marks in the view correspond to the field on the Level of Detail shelf. For example, above the State field is used in the view so there is a single mark for each state. As you add more geographic fields to the Level of Detail shelf, the marks in the view will be broken down by the members of those fields. Refer to “Level of Detail Shelf” on page 13-21 to learn more about how this shelf works.

When you first create a map view you have several options available including a variety of layers such as streets and highways, place names, and so on. Refer to “Map Options” on page 24-7 to learn more.

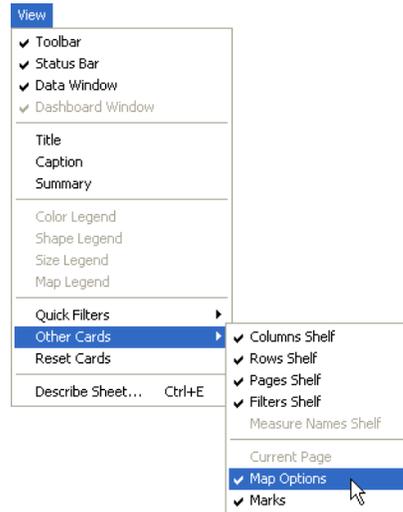
---

## Map Options

When you load a map there are several options to help you control the display of the map. For example, you can select between the normal map style or show it in grayscale using the drop-down menu at the top of the Map Options card.



You can turn on layers such as streets and highways, state and county boundaries, and place names. You can also apply data layers containing census information such as population and average household income. Finally you can adjust the washout to make the map fade into the background more. All of these options are set from the Map Options card, which displays by default when you create a map image. You can also turn it on by selecting **View > Other Cards > Map Options**.



## Map Layers

Tableau online maps provide several layers that can mark points of interest on the map. For example, overlay streets and highways, county boundaries, and more. These layers are displayed in the Map Options card.

Turn on layers by selecting them in the Map Options card.





Many layers are only visible at specific zoom levels. Refer to “Zoom” on page 19-5 to learn more about the navigation tools. Each of the map layers are described in the table below.

<b>Layer Name</b>	<b>Description</b>
Landcover	Shades wilderness areas and parks to give the map more depth.
Streets and Highways	Marks freeways and highways as well as small city streets. This layer includes the highway and street names as well.
Light Borders & Names	Shows country names, country borders, state names, state borders, and major roads as light gray.
Country Borders	Highlights country borders in a darker gray.
Country Names	Highlights country names in a darker gray.
State/Province Borders	Highlights state and province borders.
State/Province Names	Highlights state and province names.
US County Borders	Highlights U.S. county borders.
US County Names	Highlights U.S. county names.
Zip Code Boundaries	Marks U.S. zip code boundaries. You must zoom in to see this layer.
Zip Code Labels	Shows labels for U.S. zip codes. You must zoom in to see this layer.
Area Code Boundaries	Marks the U.S. area code boundaries. You must zoom in to see this layer.
Area Code Labels	Shows labels for the U.S. area codes. You must zoom in to see this layer.
US Metro Boundaries (MSA)	Marks U.S. Metropolitan Statistical Areas.



Layer Name	Description
US Metro Labels (MSA)	Shows labels for the U.S. Metropolitan Statistical Areas.
Place Names	Displays the names of places from country names and borders to city names, bodies of water, parks, universities, and more. This layer is very dependent on the zoom level.

If you find that you have a set of layers that you'd always like to show whenever you create new workbooks or sheets, you can set the current selections as the default.

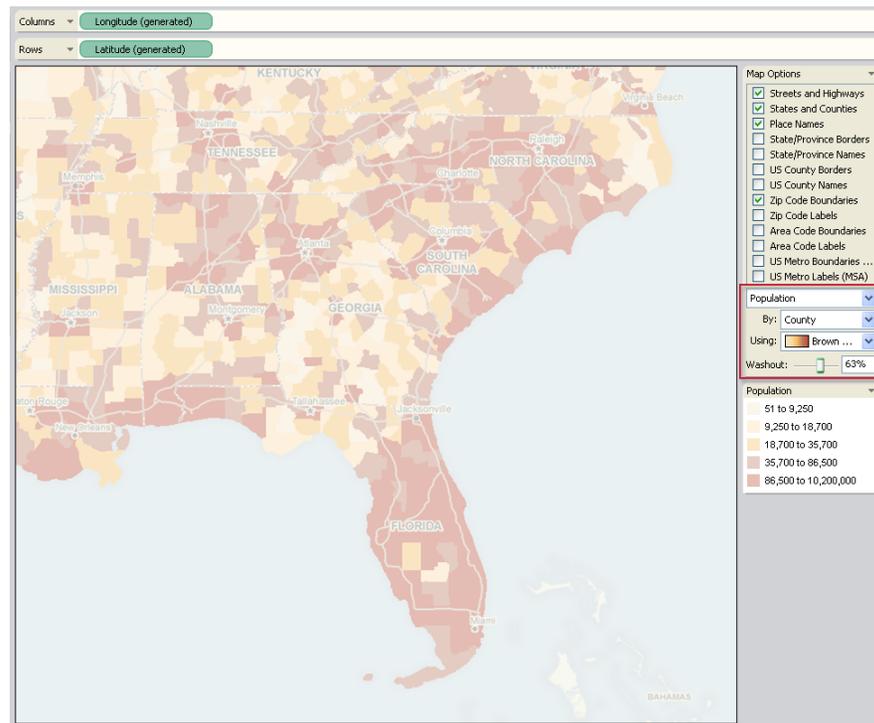
**To set the current map options as the default:**

- 1 Select the layers you want to show by default.
- 2 Select **Update Defaults** on the Map Options card menu.



## Data Layers

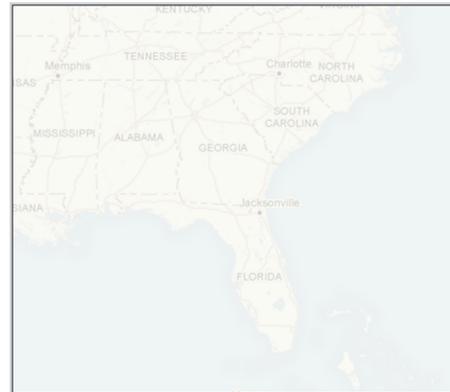
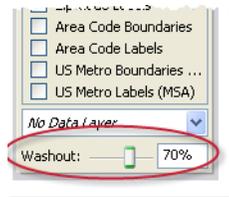
In addition to map layers you can turn on a variety of predefined data layers that show census information. The data layers are added as shading to the map and a legend is shown to explain the colors. You can change the level to show as well as the color legend using the drop-down menus on the map options card.



Hide and show the legend by selecting **View > Map Legend**.

## Washout

In addition to layers, you can use the Washout slider at the bottom of the Map Options card to control the intensity, or luminance, of the background map. The farther the slider moves to the right, the more faded the map will display behind your data.





## Setting a Default Location

Sometimes your location names may exist in multiple countries and states. For example, Clark county exists in both Washington and Nevada. In those cases, Tableau does not have enough context to geocode the location. You can set your default location to tell Tableau which country and state your data refers to.

You can set your default location by selecting **Data > Geocoding > Set Default Location**. Then select the country and/or state that your data refers to.



If you have not set a default location Tableau maps the location name to the first matching location. A warning message shows in the status bar to indicate when this has happened.

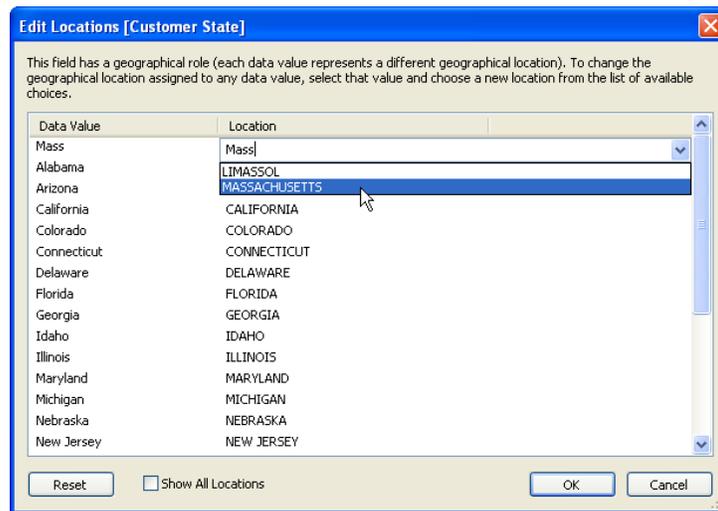
---

## Editing Locations

Sometimes Tableau will not recognize one or more of the location names in your data. When this happens those values are marked as unrecognized and mapped to the (0,0) latitude and longitude coordinates. Often this happens because your data values may be spelled incorrectly or use an abbreviation that Tableau does not understand. You can edit the unrecognized location names to map to known locations by right-clicking the geographic field in the Data window and selecting **Geographic Role > Edit Locations**.



In the Edit Locations dialog box, select the correct location for each of the unrecognized values. For example, below Mass is unrecognized. Select Massachusetts to map it to the correct location.



---

**Note** Only unused locations are shown in the list, however, you can select the **Show All Locations** option at the bottom of the dialog box to select from all locations.

---

---

## Custom Geocoding

By default Tableau recognizes a set of geographic roles that can be used to automatically geocode your data and create map views. For example, Tableau recognizes country names, state names, area codes, and so on. Refer to “Geographic Roles” on page 24-3 for the complete list of built-in geographic roles. If your data contains geographic data that is not supported with the built-in roles you can use custom geocoding to plot the data on a map.

For example, let’s say your data contains Country, States, and Cities. The built-in geocoding hierarchy recognizes country and state so your data will be aggregated to that level. However you could import the latitude and longitude information so that Tableau can plot data at the city level too. Another example is to use custom geocoding to extend the existing geographic roles to include more values such as postal codes for the United Kingdom, France, Germany, and so on. Finally, you can use custom geocoding to add new geographic roles for airport codes, custom sales geographies, and more.

To set up custom geocoding you must first create and import a geocoding file that defines the hierarchies and geographical roles. Learn more about custom geocoding in the following topics:

- Creating an Import File
- Importing Custom Geocoding
- Saving Custom Geocoding

### Creating an Import File

When adding custom geocoding in Tableau, you need to import a file that contains the new geographic roles and their latitude/longitude definitions. These files should be in Comma Separated Values (CSV) format. Setting up the file differs depending on the type of custom geocoding you are adding. There are three general types of custom geocoding:

- Extending an Existing Role (e.g., adding postal codes for countries outside of the U.S.)
- Adding New Roles (e.g., adding airport codes under the existing Country level)
- Adding New Hierarchies (e.g., adding custom sales territories that are defined by a hierarchy of district, region, territory, and theater)



### Extending an Existing Role

The built-in geocoding contains hierarchies that you can extend to include locations that are relevant to your data. For example, the existing hierarchy of Country > State only contains U.S. states. You can extend the state level to include states in other countries. The import file for this type of geocoding should contain all of the columns in every level of the hierarchy above the level you are extending. For example, the existing State level has both a Name and Abbreviation column. When adding new states make sure to include these two columns along with the parent column Country (Name). Below is an example.

	A	B	C	D	E	F
1	Country (Name)	State (Name)	State (Abbreviation)	Latitude	Longitude	
2	UNITED KINGDOM	ENGLAND	EN	52.5	1.5	
3	UNITED KINGDOM	SCOTLAND	SC	56.5	4	
4	UNITED KINGDOM	WALES	WL	52	3.75	
5	UNITED KINGDOM	NORTHERN IRELAND	NI	54.5	6.75	
6						
7						

In the import file, the names of the columns define the Geographic Roles. That means when you want to extend the existing geocoding, you must make sure your column names match the existing geographic roles in the hierarchy that you are extending. That way the new locations will be added to the Country (Name), State (Name), and State (Abbreviation) roles.

Save the file using a descriptive name such as UK States.

---

**Note** The highest level in the hierarchy is Country and cannot be extended to include higher levels such as Continent, etc.

---



## Adding New Roles

Sometimes you may want to add entirely new geographic roles that still fit into the existing hierarchies in the built-in geocoding. For example, you may have data that contains Airport codes or National Park names that you want to plot on a map. These roles fit into the existing hierarchy of Country > State. The import file for this type of geocoding must contain the columns for each level in the existing hierarchy you are adding the role to. For example, airport codes are added to the Country > State hierarchy so the import file must contain all of the columns for the Country and State levels. Below is an example of an import file containing airport codes.

	A	B	C	D	E	F	G
1	Airport (ICAO)	Airport (IATA)	Airport (City)	Country (Name)	Latitude	Longitude	Altitude
2	AYGA	GKA	GOROKA	PAPUA NEW GUINEA	-6.081667	145.391667	5282
3	AYLA	LAE	LAE	PAPUA NEW GUINEA	0	0	0
4	AYMD	MAG	MADANG	PAPUA NEW GUINEA	-5.206944	145.788611	20
5	AYMH	HGU	MOUNT HAGEN	PAPUA NEW GUINEA	-5.826111	144.296111	5388
6	AYNZ	LAE	NADZAB	PAPUA NEW GUINEA	-6.569722	146.726111	239
7	AYPY	POM	PORT MORESBY	PAPUA NEW GUINEA	-9.443333	147.22	146
8	AYRB	RAB	RABAUL	PAPUA NEW GUINEA	0	0	0
9	AYWK	WWK	WEWAK	PAPUA NEW GUINEA	-3.583611	143.669167	19
10	BGAM	N/A	ANGMAGSSALIK	GREENLAND	0	0	0
11	BGAS	N/A	ANGISSOQ	GREENLAND	0	0	0
12	BGAT	N/A	APUTITEQ	GREENLAND	0	0	0
13	BGBW	UAK	NARSSARSSUAQ	GREENLAND	61.161111	-45.4275	112
14	BGCH	JCH	CHRISTIANSHAAB	GREENLAND	0	0	0
15	BGCO	N/A	NERLERIT INAAT	GREENLAND	70.739444	-22.645833	45
16	BGDB	N/A	DANEBOG	GREENLAND	0	0	0
17	BGDH	N/A	DANMARKSHAVN	GREENLAND	0	0	0
18	BGDU	N/A	DUNDAS	GREENLAND	0	0	0

Importing the file above would add the geographic roles Airport (ICAO), Airport (IATA), and Airport (City). Notice that again, the column name for country matches the existing Country (Name) geographic role.

Name the file something descriptive such as Airports.



### Adding New Hierarchies

When you extend an existing role or add a new role you are working within the already existing hierarchies. However, you may have an entirely new hierarchy that you want to add. For example, you may have custom sales territories that are defined by Theater, Region, District, and Territory instead of Country, State, and Zip Code. This case is similar to creating new roles except there is no column mapping to an existing parent role. Instead you will need to create multiple import files, each representing a level in the new hierarchy.

For this example, the top level in the hierarchy is theater so the import file would like like the example below.

	A	B	C	D
1	Theater	Longitude	Latitude	
2	Domestic	-103.4729087	41.74052214	
3	Americas	-50.88549607	-11.6730767	
4	Europe	10.01001427	51.1682355	
5	Africa	15.46324327	-1.115573283	
6	Asia	66.9236836	48.41704292	
7	Oceania	160.2164916	-9.614065555	
8				

The next level is Region so its import file should contain columns for both Theater and Region. An example is shown below.

	A	B	C	D	E
1	Theater	Region	Longitude	Latitude	
2	Domestic	East	-77.01	38.89	
3	Domestic	Central	-84.81	37.39	
4	Domestic	West	-120.58	43.89	
5	Americas	North	-99.66889701	59.68488772	
6	Americas	Central	-85.41245982	12.64791516	
7	Americas	South	-50.88549607	-11.6730767	
8	Europe	West	1.717900132	46.93931477	
9	Europe	Central	8.22751159	46.98008672	
10	Europe	East	24.96676058	45.7526647	
11	Africa	Mediterranean	9.53749919	33.83789177	
12	Africa	Sahara	-3.996166451	19.06482432	
13	Africa	Sub-Sahara	17.90777611	-12.75895424	

Continue down the hierarchy making sure that each file contains columns for the current level and all of the levels above it. For example, the District file would contain columns for District, Region, and Theater.

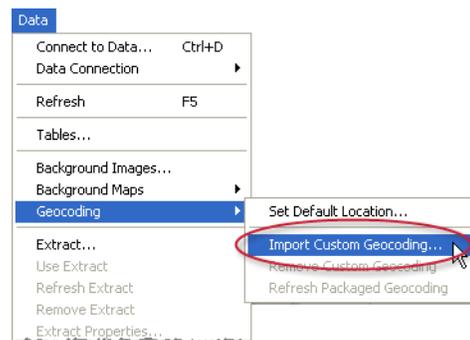


Save each import file into a single folder and name each file with the Level it represents. For this example there are four import files named Theater.csv, Region.csv, District.csv, and Territory.csv.

## Importing Custom Geocoding

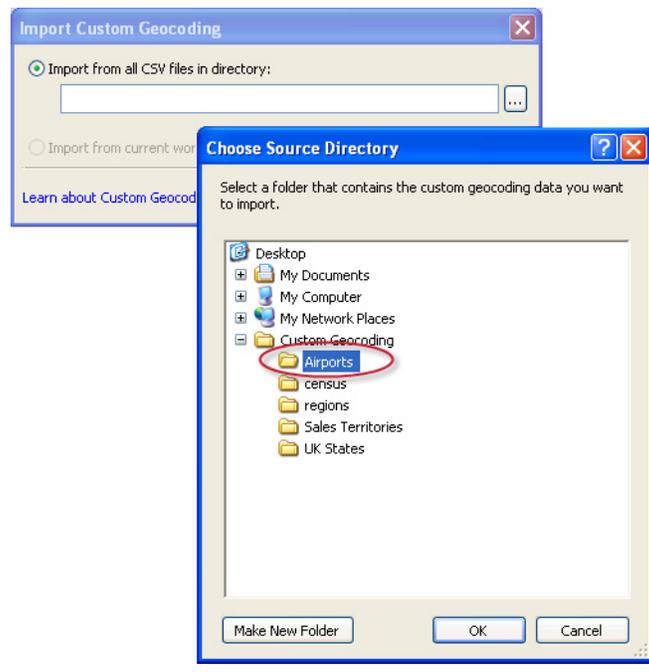
Once you have the import files set up you can import the custom geocoding into Tableau. Follow the steps below to import custom geocoding.

- 1 Select **Data > Geocoding > Import Custom Geocoding**.





- 2 In the Import Custom Geocoding dialog box, specify a directory that contains your import files. Refer to “Creating an Import File” on page 24-16 to learn how to create an import file. You can use the button to the right of the text field to browse to the directory.



- 3 When finished, click **Import**.

The custom geocoding data is imported into the workbook and the new geographic roles become available. Refer to “Geographic Roles” on page 24-3 to learn more about assigning roles and building map views.



## Saving Custom Geocoding

When you import custom geocoding, the data is stored in a Local Data folder inside your repository. Your repository must be on a local drive in order for custom geocoding to work. Every time you import custom geocoding, all of the CSV files within the chosen directory are scanned and replace the geocoding that already imported. That is why you should store your import files in a single Custom Geocoding directory and always import from there. The custom geocoding becomes available for all workbooks you open from your machine.

When you save your workbook as a packaged workbook, the custom geocoding data is packaged with the workbook.

---

**Note** When you open a packaged workbook you can import the custom into your own repository, but be careful because it will replace any other custom geocoding you may have imported.

---

You can remove the custom geocoding stored in your repository by selecting **Data > Geocoding > Remove Custom Geocoding**. Keep in mind that this will not remove the geocoding from a packaged workbook, it simply removes it from your repository.



## Background Map Sources

Tableau comes with a set of online and offline maps that you can access to create map views. In addition, Tableau supports connecting to a Web Map Server (WMS) to support custom maps that are specific to your industry. You can specify the map source, import new maps sources, or export a map source to share with others.

**To select a new map source:**

- Select **Data > Background Maps** and then select the map source you want to use.

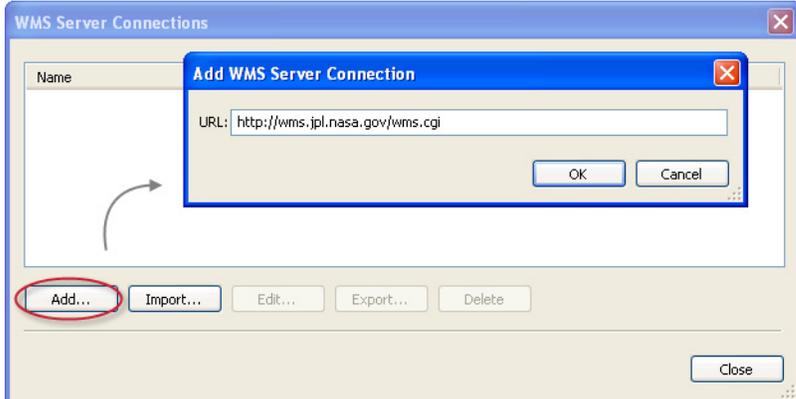
By default Tableau connects to an online map provider, offline map tiles that ship with the product, or you can add a WMS server.

### Working with WMS Servers

If you have a WMS server that provides custom maps that are specific to your industry you can add the server as a map source. After you added a WMS map server, you can export the map source to share with others or import a map source someone has shared with you.

**To add a WMS server:**

- 1 Select **Data > Background Maps > WMS Servers**
- 2 Click **Add**.
- 3 In the subsequent dialog box, type the URL for the server and click **OK**.

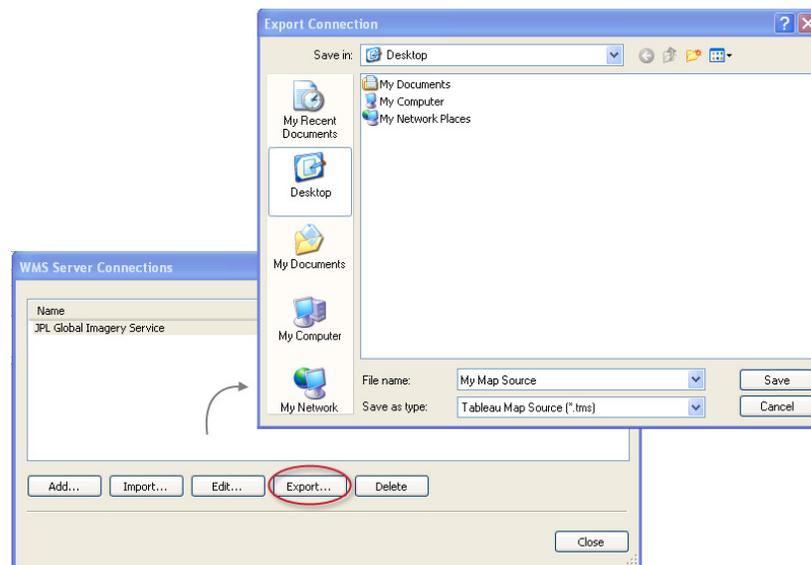




You can add as many map servers as you need. Each map source is displayed in the list of map sources on the Data menu. The map sources will be saved with the workbook and available to anyone you share the workbook with. You can also export the map source to a file that you can share with others so they can quickly connect to it.

**To export a WMS server:**

- 1 Select **Data > Background Maps > WMS Servers**.
- 2 Select the server you want to export as a map source and click **Export**.
- 3 Type a name for the file and choose a location. The file will be saved as a Tableau Map Source (.tms).
- 4 Click **Save**.

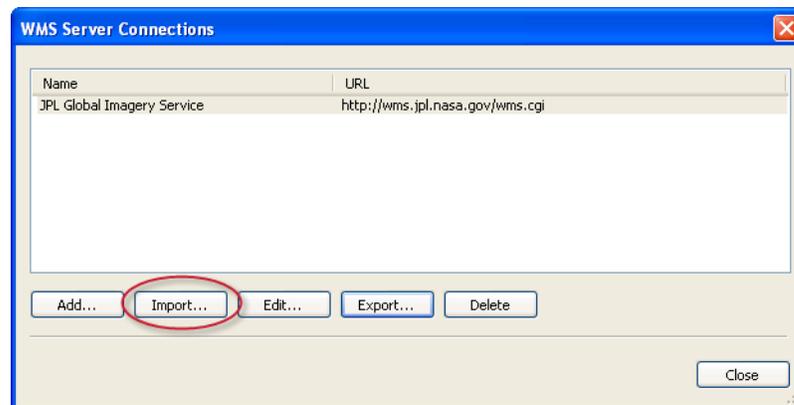




The Tableau Map source includes any default map options settings you have specified. For example, the map source will include any custom set of map layers you have specified to show by default. If you change the map options defaults you should export the map source again to include the new settings in the map source file.

**To import a map source:**

- 1 Select **Data > Background Maps > WMS Servers**.
- 2 Click **Import**.
- 3 Navigate to the saved map source file (.tms), select it, and click **Open**.



## Setting a Default Map Source

By default all new worksheets connect to Tableau’s online map source. You can specify a different map source as the default.

**To specify a default map source:**

- 1 Select the map source on the **Data > Background Maps** menu.
- 2 Select **Data > Background Maps > Set as Default**.

---

## Map Storing and Working Offline

When you create map views using the online map provider, Tableau stores the images that make up the map in a cache. That way, as you continue your analysis you don't have to wait for the maps to be retrieved. In addition, by storing the maps you can do a certain amount of work when you are offline.

The cache for the maps are stored with your Internet Explorer Temporary Internet Files and can be cleared at any time simply by deleting the temporary files in Internet Explorer.

When you are working offline and open a view that uses a map that is stored in the cache, the stored image will be used automatically. However, there are several actions that require Tableau to retrieve a new map. If the new map is not also stored in the cache you won't be able to load the map until you go online. The main actions to be aware of are listed below.

- **Turning layers on or off** - if you decide to turn on a layer that isn't stored in the cache, Tableau will need to connect to retrieve the necessary information.
- **Zooming** - zooming in or out on a map requires different map images. If the images at the specified zoom level don't exist in the cache, Tableau will need to retrieve the updated maps.
- **Panning** - panning sometimes requires new map images. If you are working offline and don't have the necessary map images and legends stored in the cache, the new images and legends will not load.

Stored map images and legends remain valid for about thirty days. After that time, Tableau will not use the stored image, instead it will require you to reconnect and fetch an updated map. This is to prevent the map images from becoming outdated.

# Trend Lines & Statistics

---

<b>Overview</b> . . . . .	<b>25-3</b>
<b>Adding Trend Lines</b> . . . . .	<b>25-4</b>
Add Trend Lines to the View . . . . .	25-4
Why can't I add Trend Lines? . . . . .	25-6
Remove Trend Lines . . . . .	25-7
<b>The Trend Line Model</b> . . . . .	<b>25-8</b>
<b>Removing Factors from the Model</b> . . . . .	<b>25-9</b>
<b>Testing Significance</b> . . . . .	<b>25-12</b>
Entire Model Significance . . . . .	25-12
Significance of Specific Fields . . . . .	25-13
<b>Trend Lines Example</b> . . . . .	<b>25-14</b>
<b>Assumptions</b> . . . . .	<b>25-18</b>
<b>Trend Line Model Terms</b> . . . . .	<b>25-19</b>
<b>Commonly Asked Questions</b> . . . . .	<b>25-21</b>





## 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:

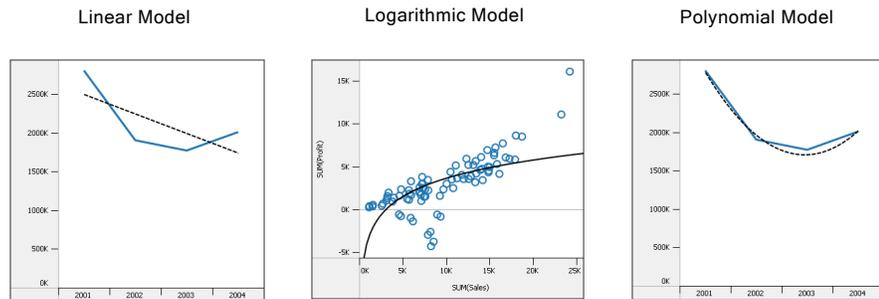
- Add Trend Lines to the View
- Why can't I add Trend Lines?
- Remove Trend Lines

### 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.

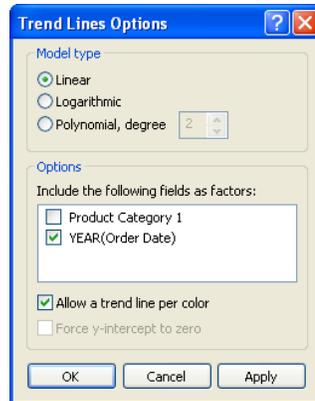
- 2 Right-click on the pane and select **Edit Trend Lines** to open a dialog box where you can specify the following options:
  - Select either a **Linear**, **Logarithmic**, or **Polynomial** model type.



- 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



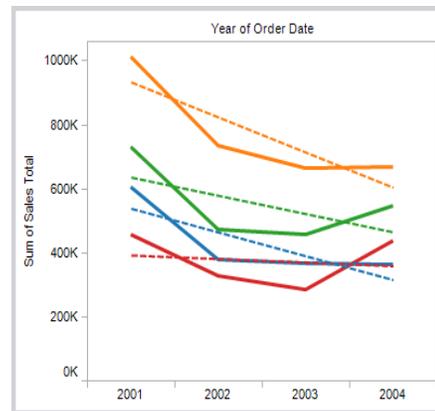
trend across all products rather than a different line for each product. In this case, you would exclude the product field as a factor. As shown here:



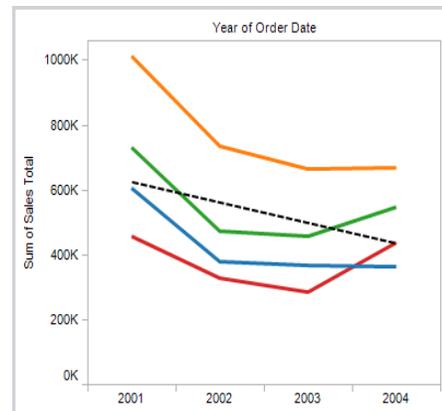
For more information on excluding factors refer to “Removing Factors from the Model” on page 25-9.

- 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





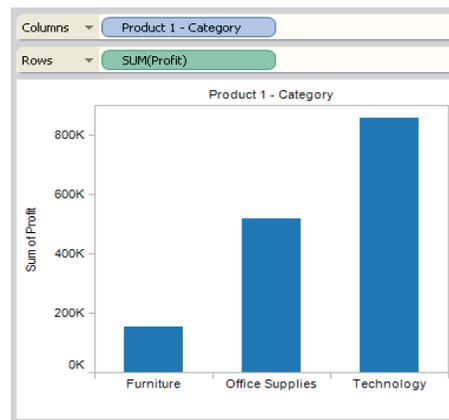
- Select whether to force the intercept to zero. This option is useful when you know that you want your trend line to begin at zero. For example, you may have an analysis of the number of products sold over time. You know that you started out with zero products sold, so you can force the trend line to begin there.

3 When finished click **OK**.

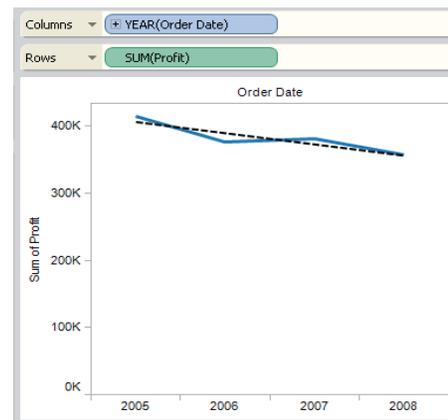
### Why can't I add Trend Lines?

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 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 profit of a company over four years.

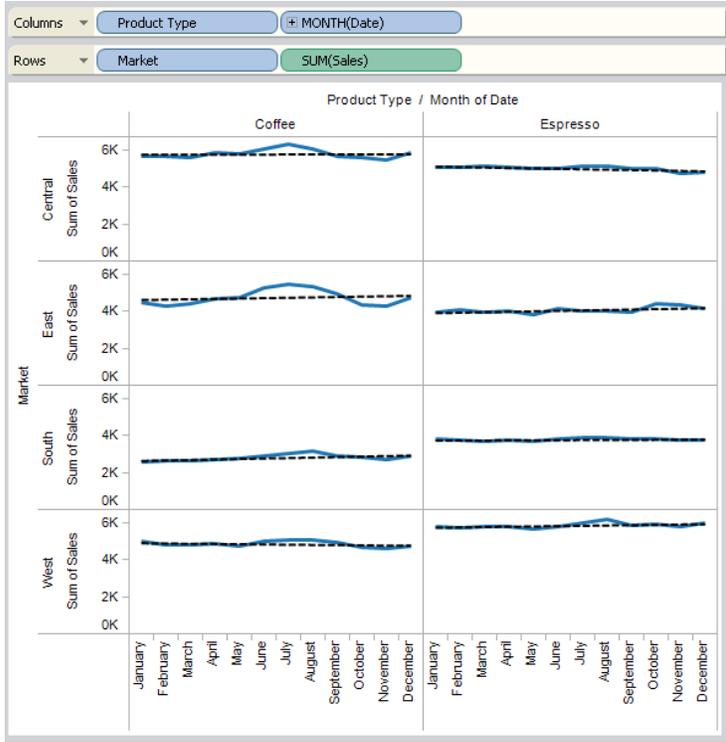
When you add a trend line in Tableau, the line you see is defined by a mathematical formula of the form:

$$Y = \text{factor 1} * \text{factor 2} * \dots * \text{factor N} * f(x) + \text{random error}$$

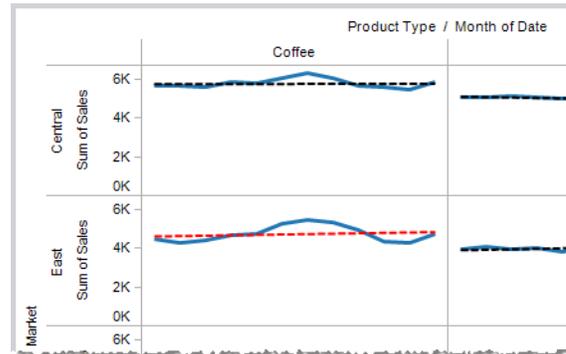
The term Y in the above expression is called the response variable and corresponds to the value you are trying to predict while the term X is the independent variable. The factors in the expression correspond to the categorical fields in the view (time in the profit example above). In addition, each 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 factor 1 \* 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

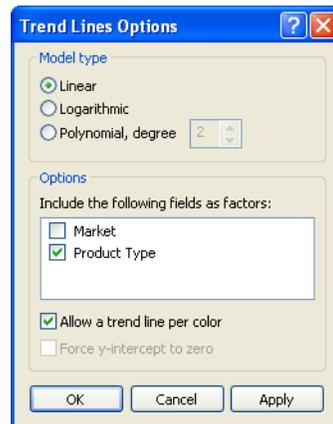
You can remove factors from the model using the Trend Lines Options dialog box. Often you will want to remove factors because you want the trend line model to be based on the entire row in the table rather than broken up by the members or values of a field. Consider the following example. The view below shows the monthly sales for two different products, broken down by region.



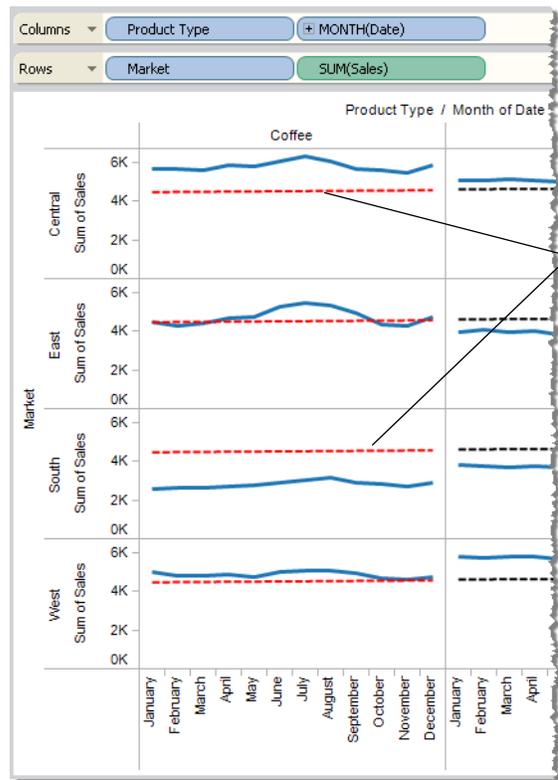
When you select a trend line in the view, you can see that a separate model is created for each customer segment.



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.



The same trend line model is used for each row because market is no longer a factor.

## Testing Significance

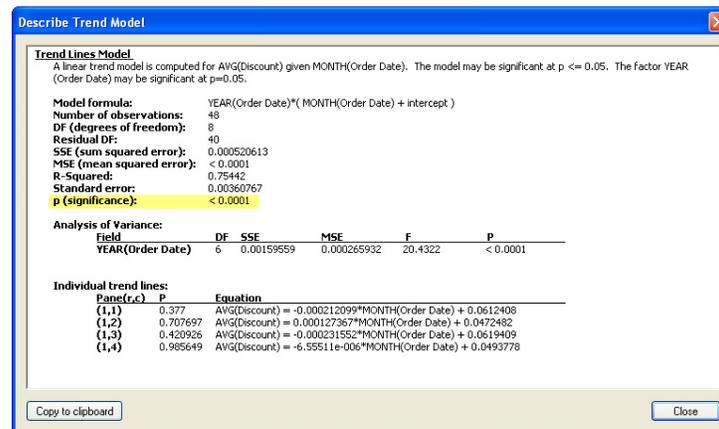
Once you've added a trend line to the view, you typically want to know the goodness of fit of the model, which is a measure of the certainty of the model's predictions. In addition, you may be interested in the significance of each factor. You can view these numbers using the Describe command.

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.



The screenshot shows the 'Describe Trend Model' dialog box with the following content:

**Trend Lines Model**  
A linear trend model is computed for AVG(Discount) given MONTH(Order Date). The model may be significant at  $p \leq 0.05$ . The factor YEAR (Order Date) may be significant at  $p=0.05$ .

**Model formula:** YEAR(Order Date)\*( MONTH(Order Date) + intercept )

**Number of observations:** 48  
**DF (degrees of freedom):** 8  
**Residual DF:** 40  
**SSE (sum squared error):** 0.000520613  
**MSE (mean squared error):** < 0.0001  
**R-Squared:** 0.75442  
**Standard error:** 0.00360767  
**p (significance):** < 0.0001

**Analysis of Variance:**

Field	DF	SSE	MSE	F	P
YEAR(Order Date)	6	0.00159559	0.000265932	20.4322	< 0.0001

**Individual trend lines:**

Parms(r,c)	P	Equation
(1,1)	0.377	AVG(Discount) = -0.000212099*MONTH(Order Date) + 0.0612408
(1,2)	0.707697	AVG(Discount) = 0.000127367*MONTH(Order Date) + 0.0472482
(1,3)	0.420926	AVG(Discount) = -0.000231552*MONTH(Order Date) + 0.0619409
(1,4)	0.985649	AVG(Discount) = -6.55511e-006*MONTH(Order Date) + 0.0493778

Buttons: Copy to clipboard, Close

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.

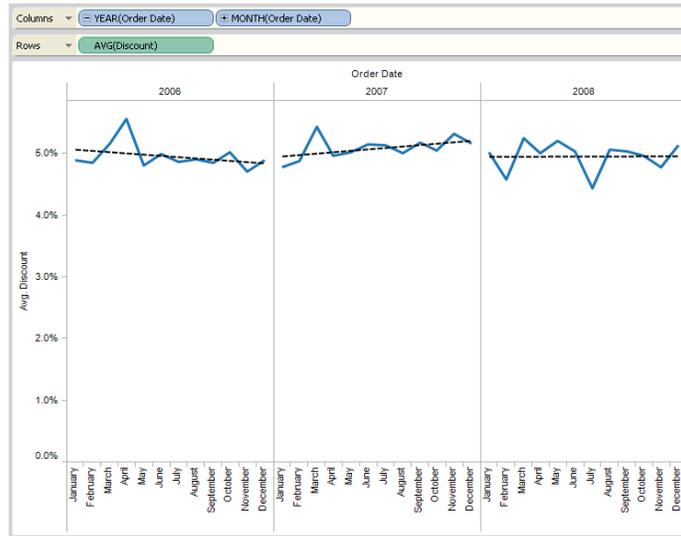
The Individual Trend Line table in the Describe Trend Line Model dialog box lists of all the trend lines in your view along with their p-value and the formula. This is an easy way to quickly determine which (if any) trend lines are statistically significant.

---

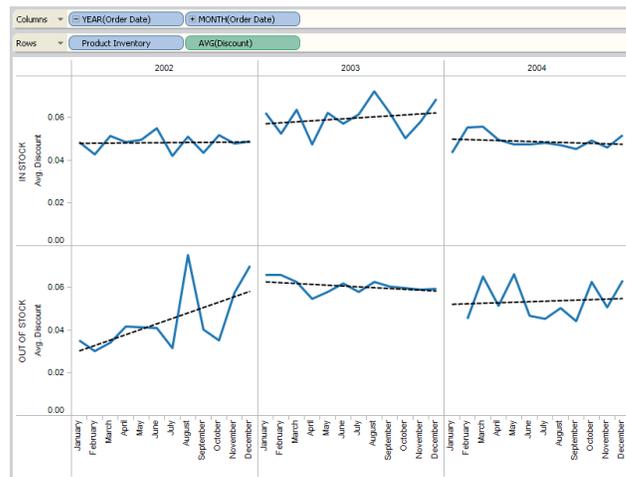
## Trend Lines Example

This example walks through the steps of incrementally testing the significance of various factors to solve a problem.

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 25-4 to learn how to add and remove trend lines.



Our first thought is that the discount rates may be high on out of stock items and employees are offering discounts for the inconvenience. When we put the Order Priority 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 Variance:**

Field	DF	SSE	MSE	F	P
Product Category 3	426	0.415266	0.000974803	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 25-12 for more information.

**Describe Trend Model**

**Trend Lines Model**  
 A linear trend model is computed for AVG(Discount) given MONTH(Order Date). The model may be significant at  $p \leq 0.05$ . The factor Product Category 3 may be significant at  $p=0.05$ . The factor YEAR(Order Date) may be significant at  $p=0.05$ .

**Model formula:** Product Category 3\*YEAR(Order Date)\*( MONTH(Order Date) + intercept )  
**Number of observations:** 1654  
**DF (degrees of freedom):** 434  
**Residual DF:** 1220  
**SSE (sum squared error):** 0.817705  
**MSE (mean squared error):** 0.00067025  
**R-Squared:** 0.37031  
**Standard error:** 0.0258892  
**p (significance):** < 0.0001

**Analysis of Variance:**

Field	DF	SSE	MSE	F	P
Product Category 3	426	0.415266	0.000974803	1.45439	< 0.0001
YEAR(Order Date)	316	0.329282	0.00104203	1.55469	< 0.0001

**Individual trend lines:**

Copy to clipboard Close



## 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 25-8) 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

---

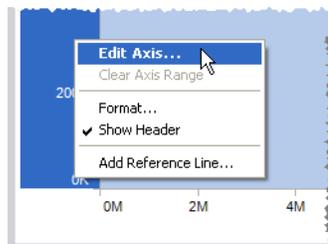
---

## 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:

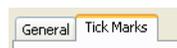
- 1 Right-click the axis in the view and select **Edit Axis**.



- 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**.

You can also reverse the axis by selecting **Reverse** in the Scale area on the General tab of the Edit Axis dialog box.

---

**Note** If your data contains negative values Tableau cannot plot them on a logarithmic scale. All values with a negative value will be displayed at 1 on the axis. You can then filter these records to exclude them from the view.

---



# Formatting

---

<b>Overview</b> . . . . .	<b>27-3</b>
<b>Worksheet Level Formatting</b> . . . . .	<b>27-4</b>
Opening the Format Window . . . . .	27-4
Fonts . . . . .	27-6
Alignment . . . . .	27-9
Shading . . . . .	27-10
Borders . . . . .	27-12
Lines . . . . .	27-12
<b>Formatting Specific Parts of the View</b> . . . . .	<b>27-14</b>
Fields . . . . .	27-14
Numbers . . . . .	27-16
Field Labels . . . . .	27-20
Legends . . . . .	27-21
Title and Caption . . . . .	27-23
Tooltips . . . . .	27-25
Null Values . . . . .	27-27
Reference Lines and Bands . . . . .	27-29
<b>Copying and Pasting Formatting</b> . . . . .	<b>27-30</b>
<b>Clearing Formatting</b> . . . . .	<b>27-31</b>
<b>Workbook Themes</b> . . . . .	<b>27-32</b>
<b>Resizing the Table</b> . . . . .	<b>27-33</b>
Resizing Rows and Columns . . . . .	27-33
Resizing the Entire Table . . . . .	27-35
Resizing Cells . . . . .	27-37
Reordering Rows and Columns . . . . .	27-39
Including and Excluding Rows and Columns . . . . .	27-39
Resize Keyboard Shortcuts and Commands . . . . .	27-40
Clearing Manual Sizing . . . . .	27-40

<b>Editing Axes . . . . .</b>	<b>27-43</b>
Changing the Axis Range . . . . .	27-43
Changing the Axis Appearance . . . . .	27-50
Formatting Tick Marks . . . . .	27-52

---

## 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



## 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 mark 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

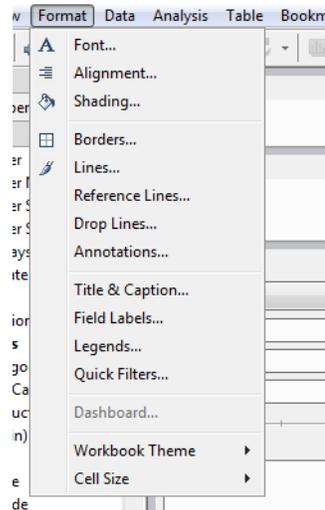
### Opening the Format Window

Use the Format menu to open the Format window.

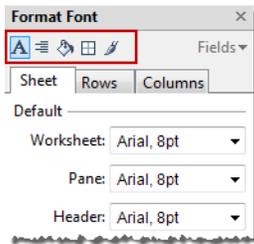
**To open the Format window:**



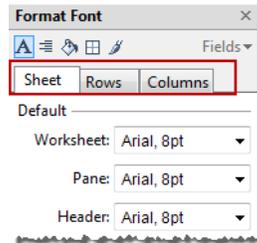
- Select Format and then select the part of the view you want to format.



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.

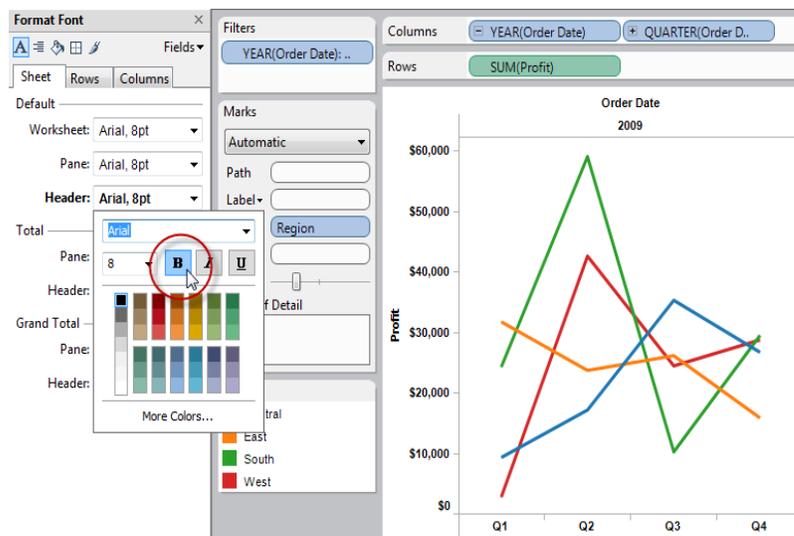


## Fonts

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.

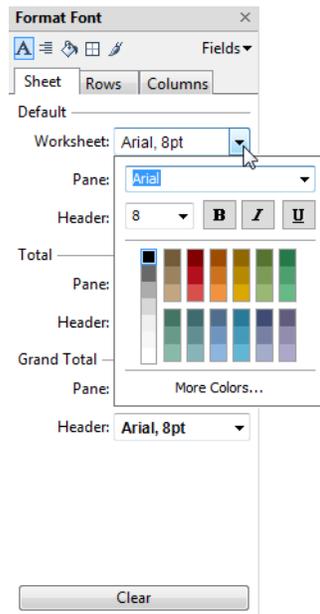




If you have totals or grand totals in the view, you can specify special font settings so that these values can stand out from the rest of the data. This is particularly useful when you are working with a text table. The view below shows a text table in which the grand totals are formatted to be bold and dark red.

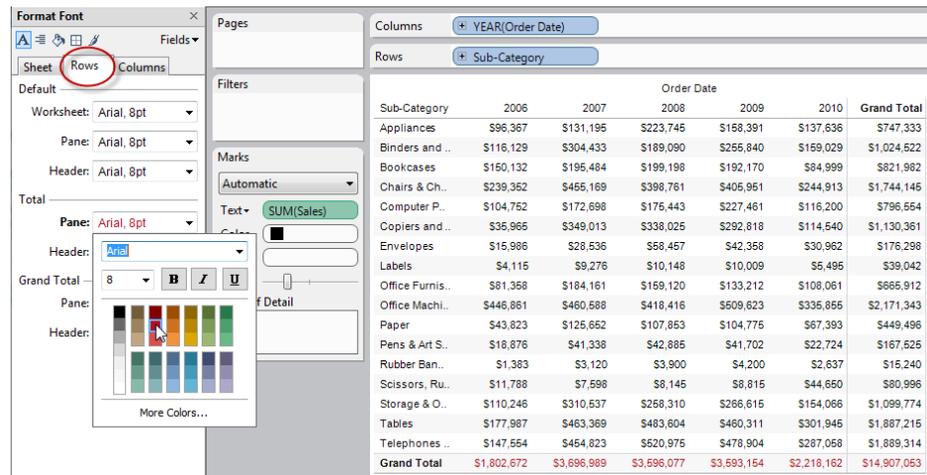
Sub-Category	Order Date					Grand Total
	2006	2007	2008	2009	2010	
Appliances	\$96,367	\$131,196	\$223,746	\$158,391	\$137,636	\$747,333
Binders and ...	\$116,129	\$304,433	\$189,090	\$255,840	\$159,029	\$1,024,522
Bookcases	\$150,132	\$196,484	\$199,198	\$192,170	\$84,999	\$821,982
Chairs & Ch...	\$239,352	\$466,169	\$398,761	\$406,951	\$244,913	\$1,744,145
Computer P...	\$104,752	\$172,698	\$175,443	\$227,461	\$116,200	\$796,554
Copiers and ...	\$36,966	\$349,013	\$338,026	\$292,818	\$114,540	\$1,130,361
Envelopes	\$15,986	\$28,536	\$58,457	\$42,358	\$30,962	\$176,298
Labels	\$4,115	\$9,276	\$10,148	\$10,009	\$5,495	\$39,042
Office Furnis...	\$81,358	\$184,161	\$159,120	\$133,212	\$108,061	\$665,912
Office Machi...	\$446,861	\$460,588	\$418,416	\$509,623	\$335,856	\$2,171,343
Paper	\$43,823	\$126,652	\$107,853	\$104,775	\$67,393	\$449,496
Pens & Art S...	\$18,876	\$41,338	\$42,885	\$41,702	\$22,724	\$167,525
Rubber Ban...	\$1,383	\$3,120	\$3,900	\$4,200	\$2,637	\$15,240
Scissors, Ru...	\$11,788	\$7,598	\$8,145	\$8,815	\$44,650	\$80,996
Storage & O...	\$110,246	\$310,537	\$258,310	\$266,615	\$154,066	\$1,099,774
Tables	\$177,987	\$463,369	\$483,604	\$460,311	\$301,945	\$1,887,215
Telephones ...	\$147,554	\$464,823	\$520,975	\$478,904	\$287,068	\$1,889,314

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

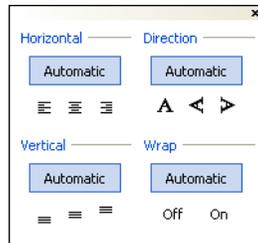
In the Format window, select the Alignment command  on the toolbar to see the Alignment format settings. You can specify alignment properties for the entire worksheet, just the rows, or just the columns using the tabs across the top of the Format window.

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	-0.0801	-0.0528
	CHAIRS & CHAIRMATS	0.1636	0.0586	0.1376
	OFFICE FURNISHINGS	0.0470	0.0396	-0.0166
	TABLES	0.0289	-0.1594	-0.1763
OFFICE SUPPLIES	APPLIANCES	0.1986	-0.0063	0.2959
	BINDERS AND BINDER AC..	0.1706	0.2670	0.1663
	ENVELOPES	0.3755	0.4326	0.4148
	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	0.1006	0.1152	0.1097
	SCISSORS, RULERS AND ..	-0.1473	-0.0740	-0.3197

		Year of Order Date		
		2001	2002	2003
FURNITURE	BOOKCASES	0.0001	-0.0801	-0.0528
	CHAIRS & CHAIRMATS	0.1636	0.0586	0.1376
	OFFICE FURNISHINGS	0.0470	0.0396	-0.0166
	TABLES	0.0289	-0.1594	-0.1763
OFFICE SUPPLIES	APPLIANCES	0.1986	-0.0063	0.2959
	BINDERS AND BINDER AC..	0.1706	0.2670	0.1663
	ENVELOPES	0.3755	0.4326	0.4148
	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	0.1006	0.1152	0.1097
	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							
	2003				2004			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
APPLIANCES	194	1,800	4,050	1,341	2,547	1,768	392	777
BINDERS AND B..	1,043	1,683	2,125	20	2,192	1,138	904	1,153
BOOKCASES	-201	-481	-603	128	-773	-102	-355	-1,433
CHAIRS & CHAI..	6,152	-224	6,771	2,916	2,578	4,856	3,241	5,896
COMPUTER PE..	12,433	7,912	9,703	3,047	5,627	7,770	4,272	7,839
COPIERS AND F..	6,562		4,439	8,796	957	8,855		23,401
ENVELOPES	3,389	1,698	-51	2,377	230	797	2,332	-5
LABELS	2,169	2,622	1,974	1,430	2,238	2,667	1,946	1,671
OFFICE FURNIS..	146	242	-638	-170	-1,168	-125	-1,056	-629
OFFICE MACHI..	4,157		-13	14,447				-26
PAPER	12,738	7,957	14,292	5,628	7,160	6,438	9,155	10,369
PENS & ART SU..	46	-124	271	-139	-107	570	448	493
RUBBER BANDS	142	115	47	0	38	-96	67	-23
SCISSORS, RUL..	-47	-471	-17	-245	117	153	-252	-276
STORAGE & OR..	3,499	-2,734	-170	38	436	-1,916	4,866	-1,362
TABLES	-2,729	-1,655	-5,896	-2,237	-3,017	-3,253	-4,255	-997
TELEPHONES A..	81,779	95,943	96,955	80,655	94,679	106,040	116,700	121,489

	Year of Order Date / Quarter of Order Date							
	2003				2004			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
APPLIANCES	194	1,800	4,050	1,341	2,547	1,768	392	777
BINDERS AND B..	1,043	1,683	2,125	20	2,192	1,138	904	1,153
BOOKCASES	-201	-481	-603	128	-773	-102	-355	-1,433
CHAIRS & CHAI..	6,152	-224	6,771	2,916	2,578	4,856	3,241	5,896
COMPUTER PE..	12,433	7,912	9,703	3,047	5,627	7,770	4,272	7,839
COPIERS AND F..	6,562		4,439	8,796	957	8,855		23,401
ENVELOPES	3,389	1,698	-51	2,377	230	797	2,332	-5
LABELS	2,169	2,622	1,974	1,430	2,238	2,667	1,946	1,671
OFFICE FURNIS..	146	242	-638	-170	-1,168	-125	-1,056	-629
OFFICE MACHI..	4,157		-13	14,447				-26
PAPER	12,738	7,957	14,292	5,628	7,160	6,438	9,155	10,369
PENS & ART SU..	46	-124	271	-139	-107	570	448	493
RUBBER BANDS	142	115	47	0	38	-96	67	-23
SCISSORS, RUL..	-47	-471	-17	-245	117	153	-252	-276
STORAGE & OR..	3,499	-2,734	-170	38	436	-1,916	4,866	-1,362
TABLES	-2,729	-1,655	-5,896	-2,237	-3,017	-3,253	-4,255	-997
TELEPHONES A..	81,779	95,943	96,955	80,655	94,679	106,040	116,700	121,489



## 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: Customer Segment, Region, and Product Category 1. When adding a row divider you can switch between dividing every product, every region, or every customer segment by sliding the level selector right and left.

Level: 

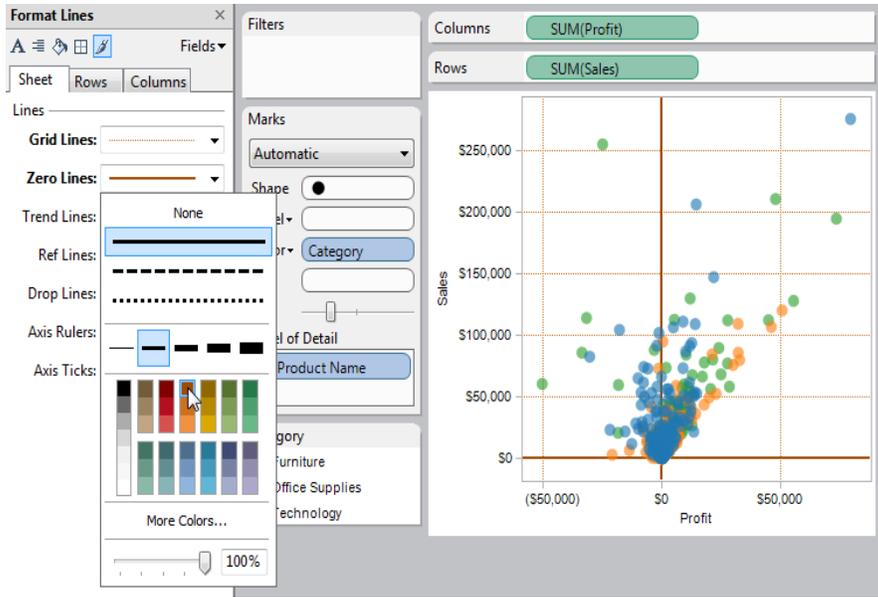
		Year of Order Date				
		2001	2002	2003	2004	
CONSUMER	CENTRAL	FURNITURE	2,816	3,405	-2,055	-120
		OFFICE SUPPLIES	5,537	7,789	4,244	2,423
		TECHNOLOGY	43,292	22,404	23,635	17,338
	EAST	FURNITURE	4,906	-2,677	5,776	-2,653
		OFFICE SUPPLIES	8,411	-1,899	5,705	1,107
		TECHNOLOGY	53,459	34,017	32,734	47,063
	WEST	FURNITURE	4,973	2,434	655	920
		OFFICE SUPPLIES	6,491	5,857	6,138	185
		TECHNOLOGY	50,902	29,598	29,861	31,057
CORPORATE	CENTRAL	FURNITURE	5,243	-1,050	499	2,702
		OFFICE SUPPLIES	6,615	8,989	3,316	4,678
		TECHNOLOGY	44,820	23,049	30,475	30,768
	EAST	FURNITURE	6,515	-2,293	-38	1,215
		OFFICE SUPPLIES	22,760	8,781	20,454	11,939
		TECHNOLOGY	130,141	101,000	93,927	85,138
	WEST	FURNITURE	-276	-3,260	-2,941	-89
		OFFICE SUPPLIES	1,636	9,640	1,513	8,402
		TECHNOLOGY	70,768	45,469	45,321	40,124
CENTRAL	FURNITURE	5,448	-714	2,338	-373	
	OFFICE SUPPLIES	6,615	8,989	3,316	4,678	
	TECHNOLOGY	44,820	23,049	30,475	30,768	

Level: 

## 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.



---

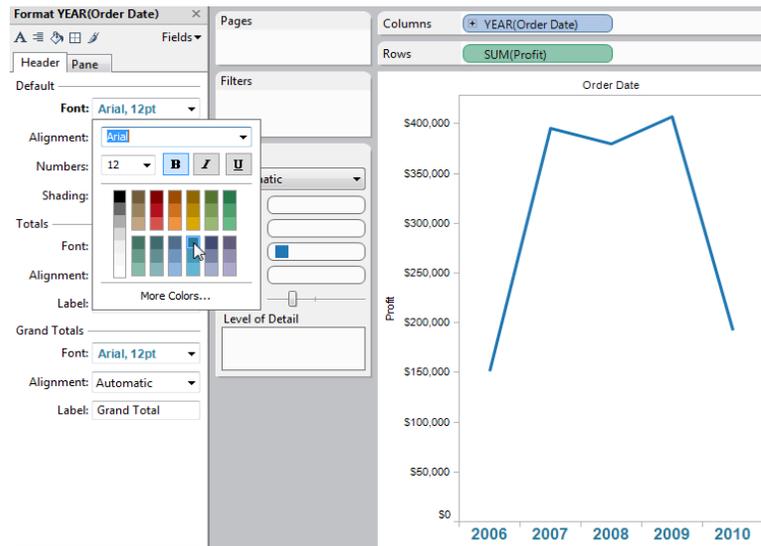
## Formatting Specific Parts of the View

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
- Legends
- Title and Caption
- Tooltips
- Null Values
- Reference Lines and Bands

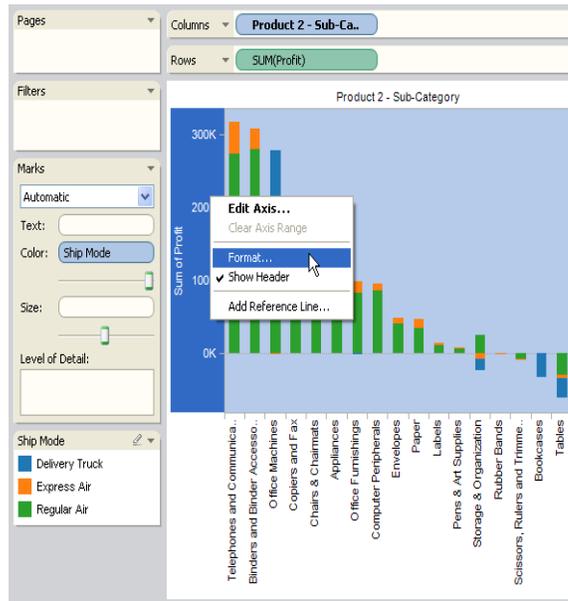
### Fields

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 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 27-43 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

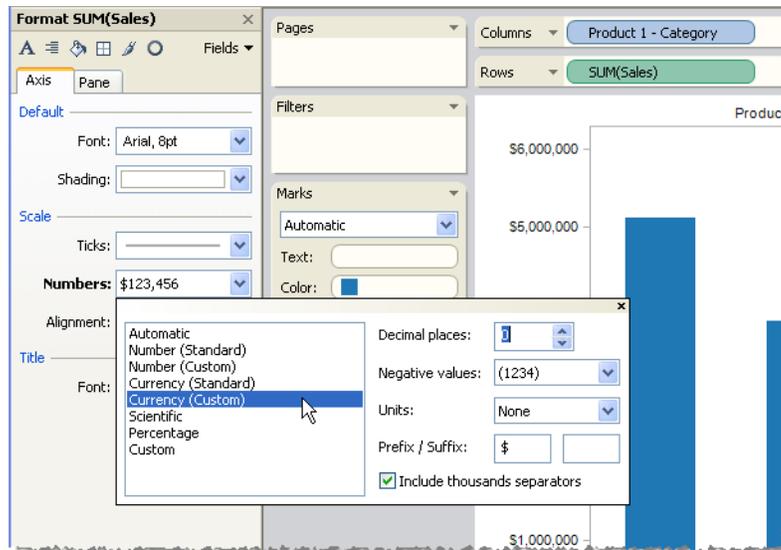
When you format a measure, you can specify the number format for both the axis and the pane text. You can select from a set of standard formats, such as number, currency,

---

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.	<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>Currency (Standard):</b> format and currency symbol is based on locale selected.	<b>Locale:</b> currency format based on the geographical location selected.



<b>Number Format</b>	<b>Format Options</b>
<b>Currency (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.

To set the default number format for a specific field, right-click the field in the Data window and select **Field Properties > Number Format**. In the subsequent dialog box you can specify a number format that will always be used when the field is added to the view. The default number format is saved with the workbook. It is also exported when you export the connection information. Refer to “Replacing Field References” on page 5-71 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

		Region / Order Date								
		EAST				WEST				
Product Family	Product Categories	2005	2006	2007	2008	2005	2006	2007	2008	
Consumer	Office Supplies	Appliances	9,818	500	9,591	4,383	9,852	4,331	5,324	2,352
		Binders and ...	7,040	5,356	642	5,283	28,665	4,997	10,507	3,312
		Envelopes	634		693	570	2,499	393	705	3,924
		Labels	88	110	285	643	323	1,380	59	106
	Technology	Computer P...	4,297	4,063	5,440	4,428	20,939	10,126	6,389	10,771
Corporate		Copiers and ...	32,238	6,950	57,919		15,782	38,231	23,278	
		Office Machi...	39,420	16,899	4,827	22,034	73,273	38,625	36,037	4,340
	<b>Total</b>		93,535	33,877	79,398	37,341	135,551	75,633	97,252	48,083
	Office Supplies	Appliances	18,544	6,326	10,726	18,618	25,340	29,183	2,934	48,854
		Binders and ...	34,390	11,756	1,577	35,347	47,241	22,954	23,642	35,905
Corporate		Envelopes	6,594	2,795	145	2,721	1,501	1,303	7,028	5,535
		Labels	1,506	873	736	542	593	400	3,380	601
	Technology	Computer P...	7,348	10,966	12,524	21,136	22,497	23,515	29,932	14,245
		Copiers and ...	12,616	37,593	17,388				3,458	3,378
		Office Machi...	8,511	44,831	60,100	61,521	75,819	55,797	6,617	98,233
<b>Total</b>		89,510	115,140	103,196	139,885	172,992	133,153	76,992	206,752	

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. You can also edit the titles that display on each legend.

### 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.

---

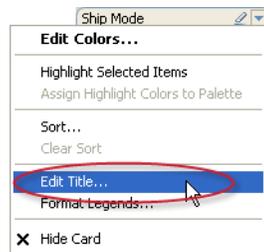
**Note** The legend format settings apply to all legends, you cannot format individual legends separately.

---

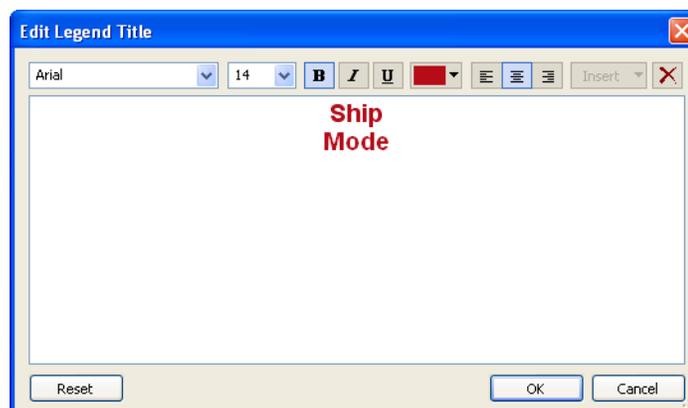
### To edit legend titles:



- 1 On the legend's card menu select **Edit Title**.



- 2 In the subsequent dialog box, type a new name for the legend and format it using the formatting options along the top of the dialog box. When finished, click **OK**.



You can click the **Reset** button in the Edit Title dialog box to return to the default title.

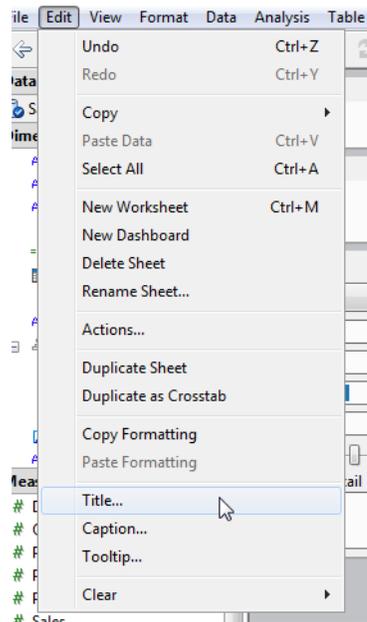


## Title and Caption

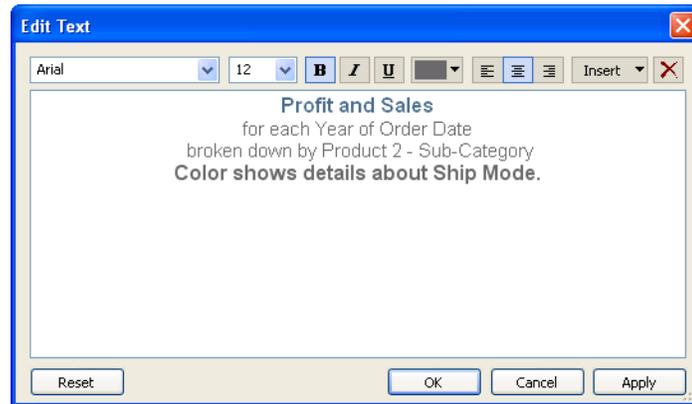
You can add titles and captions to any worksheet using the card menu on the toolbar. Refer to “Showing and hiding parts of the workspace” on page 3-15 to learn more about turning these parts of the view on and off. After you add a title or caption you can edit and format the text as well as the shading and border.

### To edit titles and captions:

- 1 Select **Edit > Title** or select **Edit > Caption**.

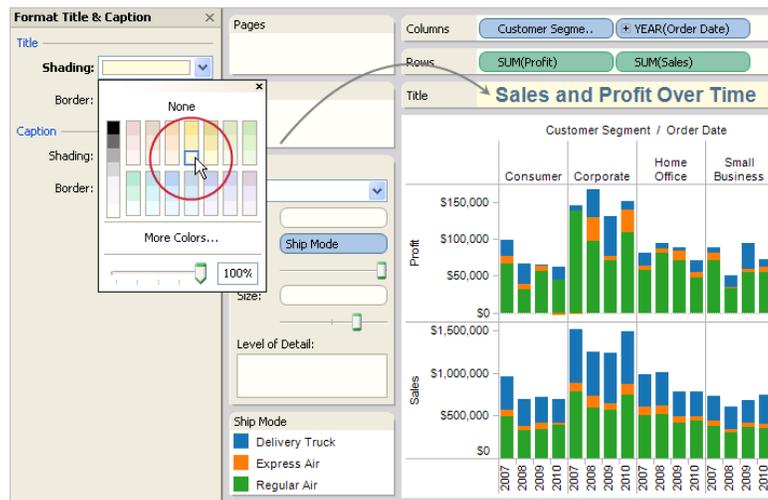


- 2 In the subsequent dialog box, modify the text and format the font, size, style, color, and alignment. Use the Insert menu to add dynamic text such as sheet properties and field values. When finished, click **OK**.



**To format title and caption borders and shading:**

- 1 Select **Format > Title & Caption** or right-click the title or caption in the view and select **Format**.
- 2 In the format window, use the drop-down controls to add shading and a border.

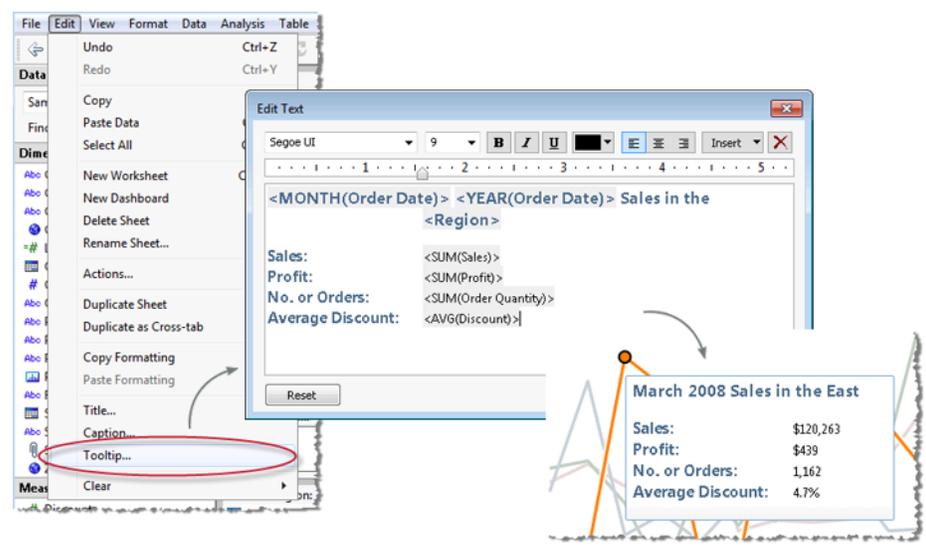




## Tooltips

Tooltips are additional data details that display when you rest the pointer over one or more marks in the view. You can edit the tooltip to include both static and dynamic text. You can also modify which fields are included in the automatic tooltip.

You can edit tooltips by selecting **Edit > Tooltip**. Tooltips are specified for each sheet and can be formatted using the formatting tools along the top of the Edit Tooltip dialog box. Use the Insert menu at the top of the dialog box to add dynamic text such as field values, sheet properties, and more.

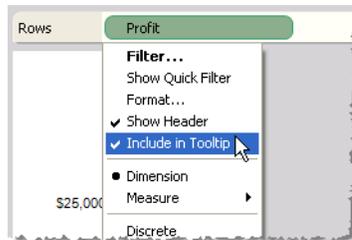


The All Fields option on the Insert menu adds all field names and values that are used in the view. Inserting the All Fields parameter will automatically update the tooltip as you change the view. You can exclude unnecessary fields from the All Fields option.

**To remove a field from the automatic tooltip:**



- Right-click the field on one of the shelves in the view and select **Include in Tooltip**.



You can show a field in the tooltip using the same menu option.

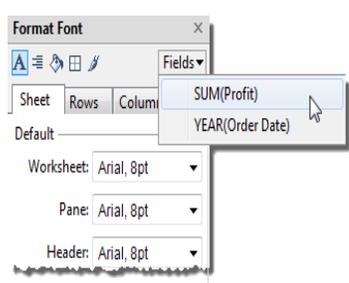


## Null Values

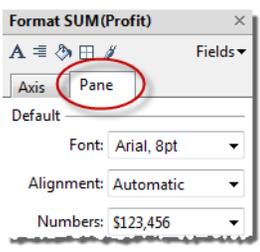
When a measure contains null values, they are usually plotted in a view as zero. However, sometimes that changes the view and you'd rather just suppress null values altogether. You can format each measure to handle null values in a unique way.

### To format a null values for a specific field:

- 1 Open the Format window and select the field using the drop-down menu in the upper right corner.



- 2 Make sure that the **Pane** tab is selected.



- 3 Optionally specify the text to label null values with when mark labels are turned on.



- 4 In the Special Values area, optionally specify whether to show or hide null values. You can alternatively hide null values but break line marks to indicate that null values exist.



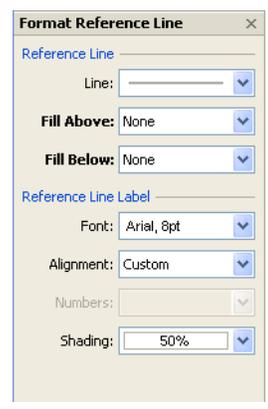


## Reference Lines and Bands

Tableau lets you format the reference lines and bands 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.

### To format a specific reference line:

- 1 Right-click on the line or bands and select **Format**.
- 2 In the Format window you can specify formatting properties associated with reference lines, bands, or distributions. For example, you can change the line style and color as well as the label font properties.



- 3 When finished click **OK**.

Many of the formatting options in the Format window are also available when adding and editing reference lines. Refer to “Reference Lines and Bands” on page 22-1 to learn more about reference lines and bands.

---

## 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**.

---

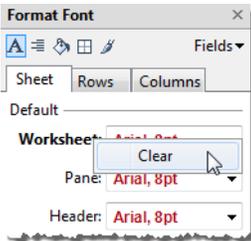


# Clearing Formatting

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.

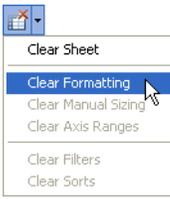
### To clear individual settings in the Format window:

- In the Format window, right-click the label of the setting you want to clear and select Clear.

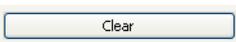


### To clear all custom formatting in the worksheet:

- 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
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
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: YEAR(Order Date) + QUARTER(Order Date)  
Rows: Customer Segment Region

Customer S.	Region	Order Date			
		2007	2007	2007	2007
		Q1	Q2	Q3	Q4
Consumer	Central	52,488	53,692	43,729	36,678
	EAST	24,175	27,161	60,526	48,296
	South	68,402	27,350	17,303	51,275
	WEST	46,099	27,849	65,259	65,036
Corporate	Central	106,046	42,721	71,616	68,736
	EAST	133,304	50,765	95,779	74,477
	South	55,454	64,284	83,726	150,323
	WEST	42,081	75,764	54,626	72,826
Home Office	Central	47,746	38,696	33,901	66,338
	EAST	19,410	74,152	49,270	34,465
	South	30,093	39,625	65,511	45,000
	WEST	80,462	72,399	30,410	46,966
Small Business	Central	16,818	27,456	45,221	55,074
	EAST	20,462	24,629	61,091	46,998
	South	59,309	21,221	54,502	87,684
	WEST	44,699	73,829	10,025	42,924

Columns: YEAR(Order Date) + QUARTER(Order Date)  
Rows: Customer Segment Region

Customer S.	Region	Order Date			
		2007	2007	2007	2007
		Q1	Q2	Q3	Q4
Consumer	Central	52,488	53,692	43,729	36,678
	EAST	24,175	27,161	60,526	48,296
	South	68,402	27,350	17,303	51,275
	WEST	46,099	27,849	65,259	65,036
Corporate	Central	106,046	42,721	71,616	68,736
	EAST	133,304	50,765	95,779	74,477
	South	55,454	64,284	83,726	150,323
	WEST	42,081	75,764	54,626	72,826
Home Office	Central	47,746	38,696	33,901	66,338

**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  $\leftrightarrow$ , 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.



The image shows a PivotTable in Excel with the following structure:

- Columns:** YEAR(Order Date), QUARTER(Order Date)
- Rows:** Customer Segment, Region

The PivotTable displays quarterly sales data for 2007 and 2008, categorized by Customer Segment and Region. The data is as follows:

Customer S..	Region	2007				2008			
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Consumer	Central	52,488	53,692	43,729	36,678	61,053	75,505	53,738	52,050
	EAST	24,175	27,161	60,526	48,296	60,526	48,296	60,526	48,296
	South	68,402	27,350	17,303	51,275	17,303	51,275	17,303	51,275
	WEST	46,099	27,849	66,269	65,036	66,269	65,036	66,269	65,036
Corporate	Central	106,046	42,721	71,616	68,736	106,046	42,721	71,616	68,736
	EAST	133,304	50,765	95,779	74,477	133,304	50,765	95,779	74,477
	South	55,454	64,284	83,726	150,323	55,454	64,284	83,726	150,323
	WEST	42,081	75,764	54,626	72,826	42,081	75,764	54,626	72,826
Home Office	Central	47,746	38,696	33,901	66,339	47,746	38,696	33,901	66,339
	EAST	19,410	74,152	49,270	34,465	19,410	74,152	49,270	34,465
	South	30,093	39,625	65,511	45,009	30,093	39,625	65,511	45,009
	WEST	80,482	72,399	30,410	46,966	80,482	72,399	30,410	46,966
Small Business	Central	16,818	27,456	46,221	55,074	16,818	27,456	46,221	55,074
	EAST	20,462	24,629	6	6	20,462	24,629	6	6
	South	59,309	21,221	5	5	59,309	21,221	5	5
	WEST	44,599	73,829	1	1	44,599	73,829	1	1

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.

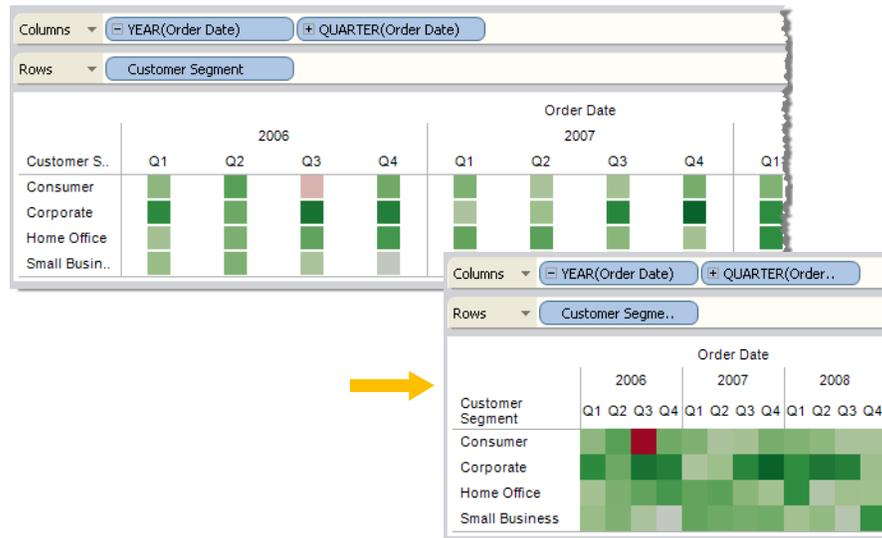


- **Text Cell** – Adjusts the view so the cell has a 3:1 aspect ratio. This is particularly useful for text tables.

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.

Customer S...	Order Date			
	2005	2006	2007	2008
Consumer	954,274	696,627	715,319	697,390
Corporate	1,518,516	1,250,647	1,242,528	1,489,335
Home Office	980,315	1,012,356	774,464	797,630
Small Business	736,818	601,001	690,842	748,991

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.



**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 17-14.

## 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 “Selecting Data to Filter” on page 16-3.



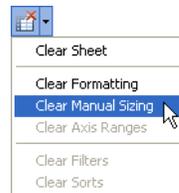
## Resize Keyboard Shortcuts and Commands

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 31-9.

## 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 27-33). 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 Profit over four years. The automatic axis may range from 0 to \$400,000 but your profits never went below \$300,000. You could adjust the Axis Range so that it starts at \$300,000 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:

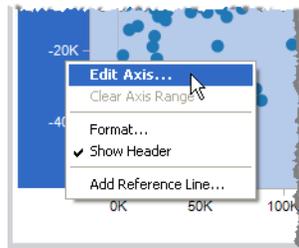
- Changing the Axis Range
- Changing the Axis Appearance
- Formatting Tick Marks

### Changing the Axis Range

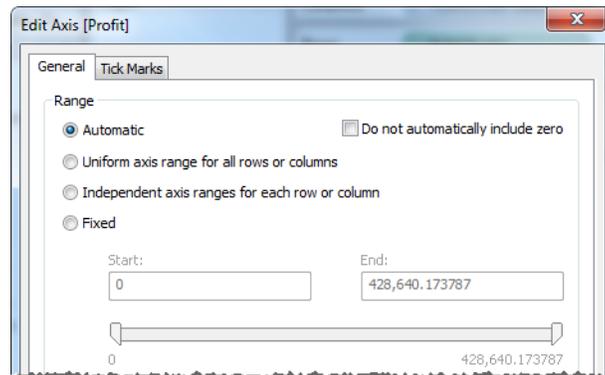
An axis shows data points that lie within a range of values. You can limit the axis range in order to focus the view to where the data points lie.

**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 You can also specify whether to automatically include zero. When the option to not include zero is selected, the axis range will adjust to just show the range of values in the data.



- 4 When finished, click **OK**.

Learn more with “Example – Changing the Axis Range” on page 27-45.

### **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.

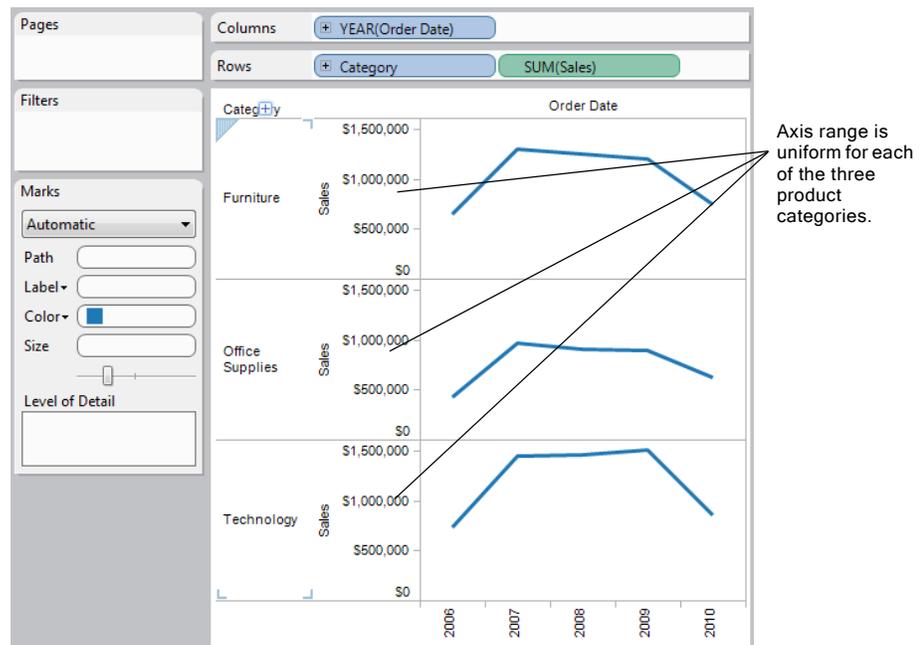
#### **View 1- Building a view with a uniform axis range**

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 1 - Category** dimension on the **Rows** shelf.

2 Place the **Sales** 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.

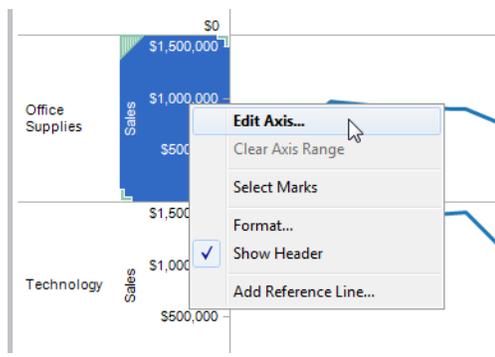


### View 2-Building a View using independent axis ranges

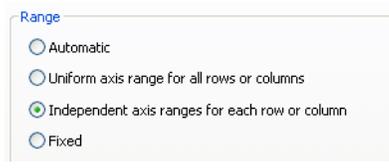
When you use an independent axis range, each row or column will have its own axis range based on the underlying data values.



- 1 Right-click the SUM(Sales) axis in the view and select **Edit Axis**.

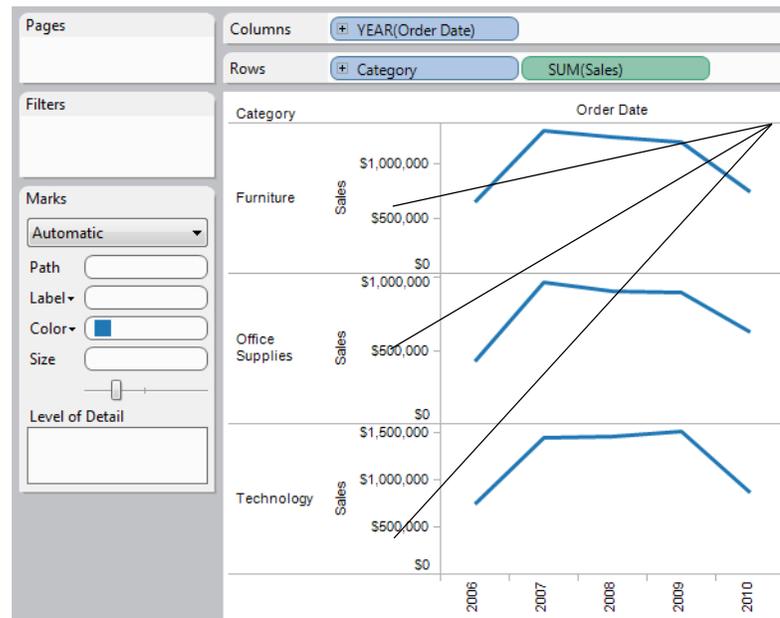


- 2 In the Edit Axis dialog box, select **Independent axis range for each row or column**.



- 3 When finished, click **OK**.

The axis range for each product category are now independent from each other. The Technology and Furniture categories still ranges from 0K to 1500K but the Office Supplies category ranges from 0K to a little over 1000K.



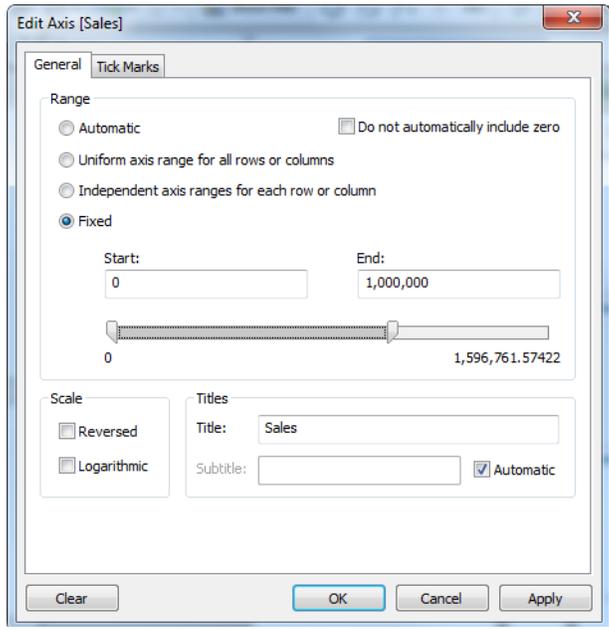
The axes ranges for each product category are independent of each other.

### 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.

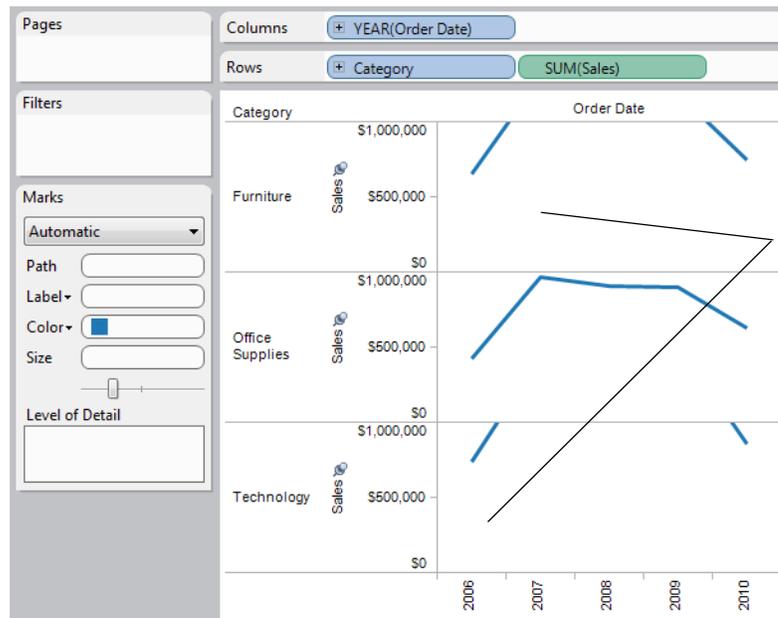


- 1 Right-click on the SUM(Sales) axis in the view and select **Edit Axis**.
- 2 In the Edit Axis dialog box, select **Fixed**. Then define **Start** and **End** values either by typing into the text boxes or by dragging the sliders toward each other. For this example, type 0 as the **Start** and 1000000 as the **End**.



- 3 When finished, click **OK**.

Notice that the Technology and Furniture categories don't have any transactions with Sales below 1000K so nothing displays in the panes. The axes are marked with a pin symbol  indicating that you have limited the axis range and that some data may not be showing.



There is no record of any Technology transaction with the total sales less than the 1,000,000 limit placed

## Changing the Axis Appearance

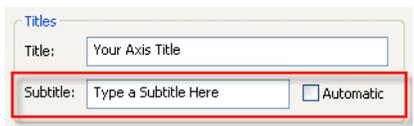
Every axis has a title that is automatically generated based on the fields in the view. You can specify a custom axis title and add a subtitle using the Edit Axis dialog box. In addition, you can specify the scale of the axis such as whether to use a logarithmic scale and whether to reverse the axis.

### 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.



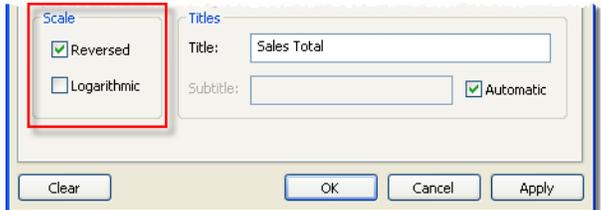
3 Deselect the **Automatic** check box to add a custom subtitle.



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 2626-1 to learn more.



3 When finished, click **OK**.

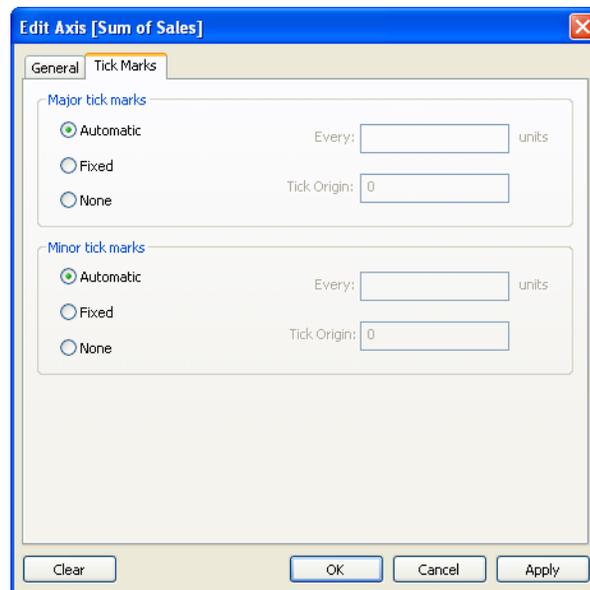


## Formatting Tick Marks

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**.



# Annotations and Mark Labels

---

<b>Overview</b> . . . . .	<b>28-2</b>
<b>Annotations</b> . . . . .	<b>28-3</b>
Adding Annotations . . . . .	28-3
Rearranging Annotations . . . . .	28-5
Formatting Annotations . . . . .	28-11
Removing Annotations . . . . .	28-15
<b>Mark Labels</b> . . . . .	<b>28-16</b>
Showing and Hiding Mark Labels for the Worksheet . . . . .	28-16
Showing and Hiding Individual Mark Labels . . . . .	28-17
Moving Mark Labels . . . . .	28-18
Editing Mark Labels with Aliases . . . . .	28-19

---

## 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 mark labels to call out marks of interest or more commonly to label the view to make it more understandable. You can show mark labels for all the marks in the view, or selectively show and hide individual labels. This section discusses the following topics:

- Annotations
- Mark Labels

---

## Annotations

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:

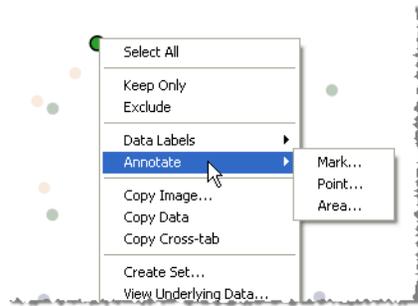
- Adding Annotations
- Rearranging Annotations
- Formatting Annotations
- Removing Annotations

### Adding 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:**

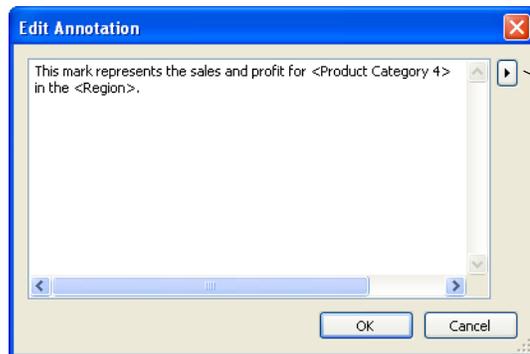
- 1 Right-click the view where you want to add an annotation and select **Annotate**.



- 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.

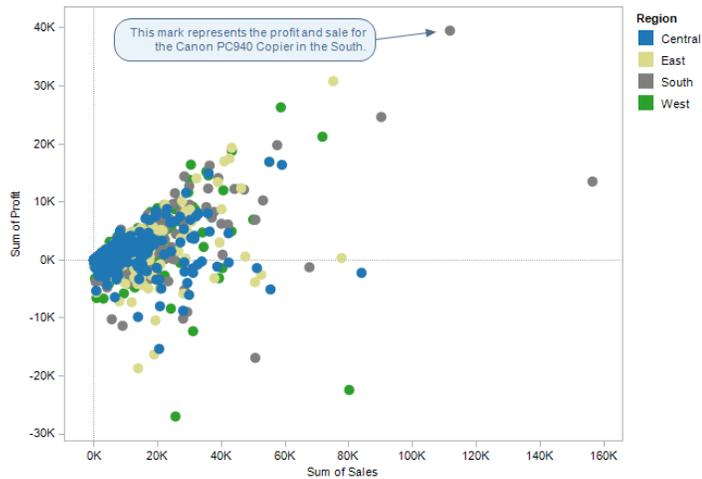


- **Area** - select this option to annotate an area in the view such as a cluster of outliers or a targeted region of the view.
- 3 In the Edit Annotation dialog box, type the text you want to show in the annotation. Use the arrow button to insert dynamic parameters into the annotation text. For example, 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**.

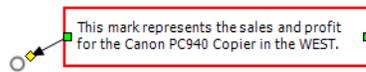


## Rearranging Annotations

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

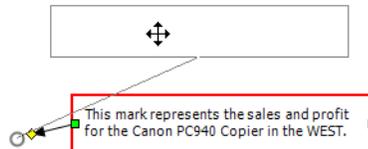
When you select a mark annotation the body and line are selected and several resize handles display. Use these handles to resize the body and line.



**To reposition the body:**

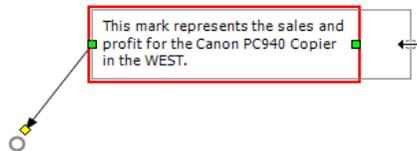


- Click and drag the body of the selected annotation to a new position.



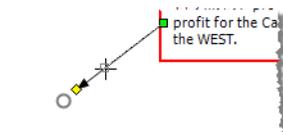
#### To resize the body:

- Click and drag the body resize handle ■ left and right. The text and height are automatically adjusted to fit the width of the body.



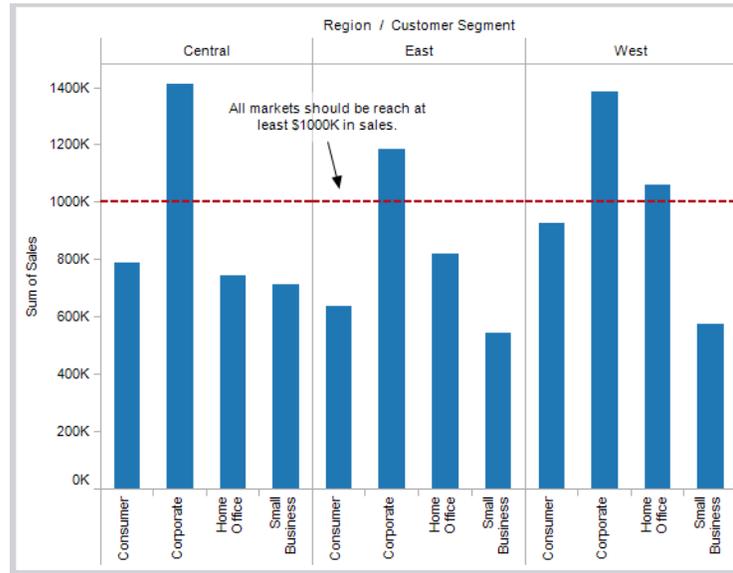
#### To resize the line:

- Click and drag the line resize handle ◆.



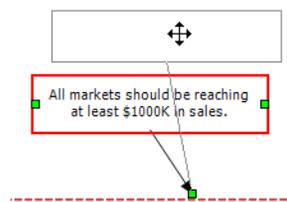
### Point Annotations

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:**

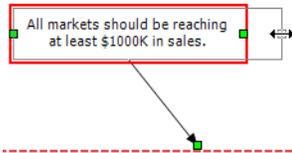
- Click and drag the body of the selected annotation to a new position. As you move the body, the line is automatically resized so that it continues to point at the specific point you selected.



**To resize the body:**

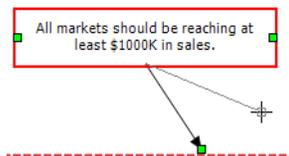


- Click and drag the side resize handles  left and right. The text and height are automatically adjusted to fit the width of the body.



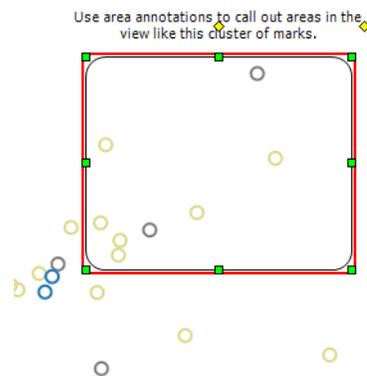
**To move the line end point:**

- Click and drag the end point of the line  so that it points at a new location.



**Area Annotations**

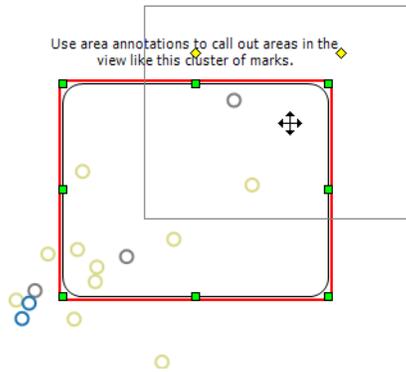
An area annotation is a way to highlight or call out an area in the view. Area annotations are not associated with any particular mark, in fact, these annotations are commonly used to call out several marks. When you select an area annotation, several resize handles and two text handles display. Use these handles to reposition and resize the box and text.





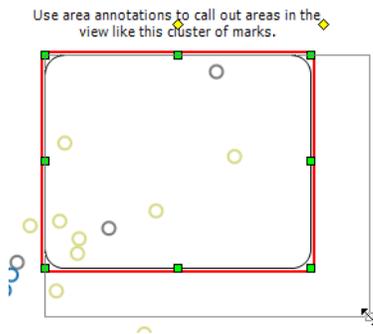
**To reposition the box:**

- Click and drag the box of the selected annotation to a new position.



**To resize the box:**

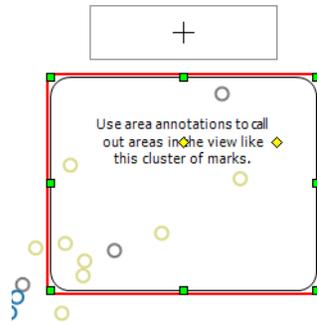
- Click and drag one of the box resize handles ■.



**To reposition the text:**

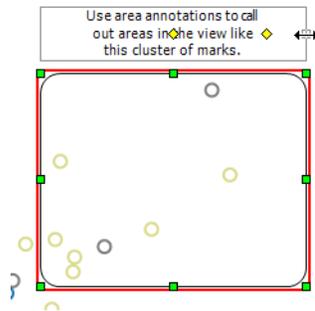


- Click and drag the center text handle  to a new position.



**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.

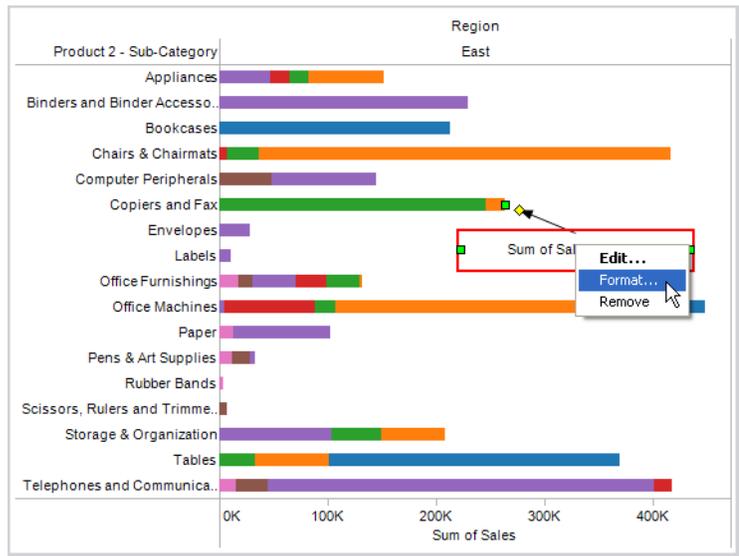


## Formatting Annotations

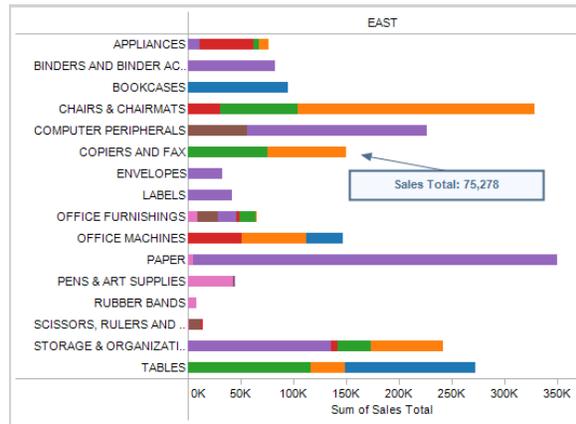
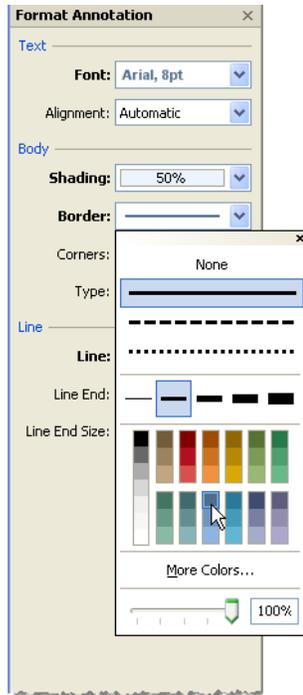
For each annotation you can modify the text, body, and line. For example, you can specify whether the body should be a box, a single edge, or not shown at all. Additionally, you can specify whether the lines on mark and point annotations end with an arrow, dot, or a simple line.

### 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.

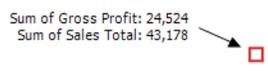
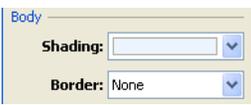
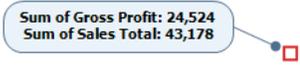
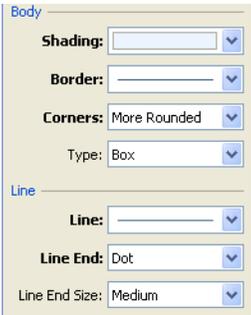
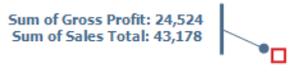
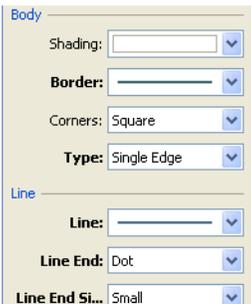


- 2 In the Format window, use the drop-downs to specify font properties, text alignment, line style, and shading.

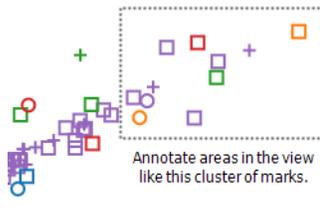




### Example Formatting Options

Annotation Style	Format Window Settings
 <p>Sum of Gross Profit: 24,524 Sum of Sales Total: 43,178</p>	Default format settings for point and mark annotations.
 <p>Sum of Gross Profit: 24,524 Sum of Sales Total: 43,178</p>	
 <p>Sum of Gross Profit: 24,524 Sum of Sales Total: 43,178</p>	
 <p>Sum of Gross Profit: 24,524 Sum of Sales Total: 43,178</p>	



Annotation Style	Format Window Settings
	<p>Body</p> <p>Shading: <input type="text"/></p> <p><b>Border:</b> <input type="text" value="Dashed"/></p> <p>Corners: <input type="text" value="Square"/></p> <p><b>Type:</b> <input type="text" value="Box"/></p> <p>Line</p> <p><b>Line:</b> <input type="text" value="Dashed"/></p> <p><b>Line End:</b> <input type="text" value="Open Arrow"/></p> <p><b>Line End Si...</b> <input type="text" value="Small"/></p>
	<p>Body</p> <p><b>Shading:</b> <input type="text" value="None"/></p> <p><b>Border:</b> <input type="text" value="None"/></p> <p><b>Corners:</b> <input type="text" value="Rounded"/></p> <p>Type: <input type="text" value="Box"/></p>
	<p>Body</p> <p>Shading: <input type="text"/></p> <p><b>Border:</b> <input type="text" value="Dotted"/></p> <p><b>Corners:</b> <input type="text" value="Square"/></p> <p>Type: <input type="text" value="Box"/></p>

To learn more about the Format window refer to “Formatting” on page 27-1.



## Removing Annotations

At any time you can remove one or more annotations.

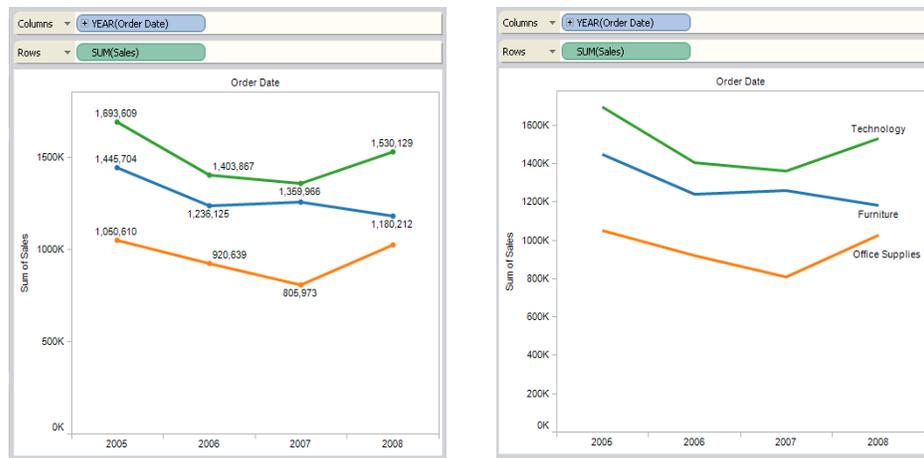
### To remove an annotation:

- 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.

---

## Mark Labels

Mark 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 mark labels so the sales values display next to each point along the lines. Alternatively, mark labels don't have to be measure values. In the same example, you could turn on mark labels that display the name of each product category next to each line.

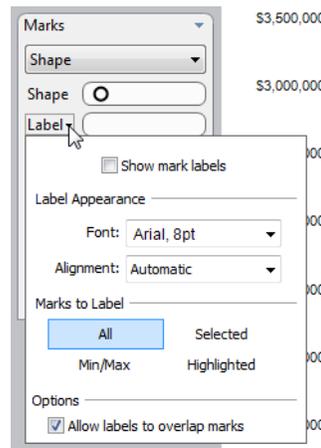


Mark Labels can be turned on for all marks, selected marks, highlighted marks, the minimum and maximum values, or just the line ends in a line chart. Additionally, you can turn on mark labels for an individual marks. This section discusses the following topics:

- Showing and Hiding Mark Labels for the Worksheet
- Showing and Hiding Individual Mark Labels
- Moving Mark Labels
- Editing Mark Labels with Aliases

### Showing and Hiding Mark Labels for the Worksheet

You can show and hide mark labels for the whole worksheet by clicking the Show Mark Labels  button on the toolbar. Use the Label shelf drop-down control to specify font properties, alignment, when to show the label, and other options.



## Showing and Hiding Individual Mark Labels

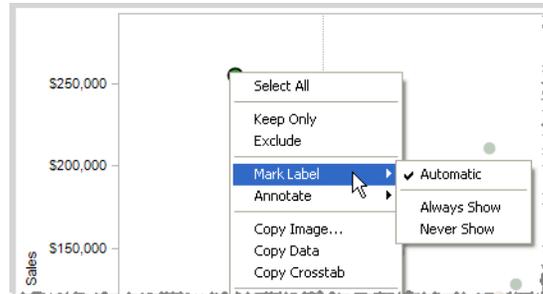
Rather than showing all mark labels or dynamically showing labels based on the view, you may want to show labels for a selection of individual marks. You can use mark labels to call out the values of specific marks of interest as well as hide overlapping mark labels. You can show and hide individual mark labels using the right-click context menus in the view.

### To show or hide individual mark labels:

- 1 Right-click the mark you want to show or hide a mark label for.
- 2 On the context menu select **Mark 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 settings in the Format window.
  - **Always Show** - select this option to show the mark label even when it would otherwise be hidden (based on the settings in the Format window).



- **Never Show** - select this option to hide the mark label even when it would otherwise be shown (based on the settings in the Format window).

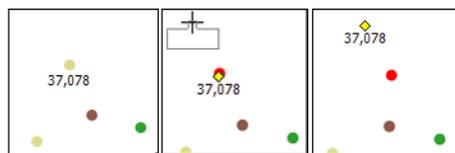


## Moving Mark Labels

After you show a mark label in a view, you can reposition it to best fit your view and presentation. For example, in a stacked bar chart, the mark 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 mark label:

- 1 Select the mark whose mark label you want to move.
- 2 Click and drag the yellow move handle to a new location.





## **Editing Mark Labels with Aliases**

Another way to modify mark 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 mark labels displayed in a the view. For more information about Aliases and how to edit them refer to “Aliases” on page 11-20.



# Dashboards

---

<b>Overview</b> . . . . .	<b>29-2</b>
<b>Creating Dashboards</b> . . . . .	<b>29-3</b>
How to Create a Dashboard . . . . .	29-3
Adding Views to a Dashboard . . . . .	29-3
Adding Dashboard Objects . . . . .	29-6
Removing Views and Objects from a Dashboard . . . . .	29-8
<b>Organizing Dashboards</b> . . . . .	<b>29-10</b>
Showing and Hiding Parts of a Worksheet . . . . .	29-15
Rearranging Dashboard Views and Objects . . . . .	29-15
Setting the Dashboard Size . . . . .	29-18
<b>Understanding Dashboards and Worksheets</b> . . . . .	<b>29-20</b>

---

## Overview

A dashboard is a collection of several worksheets and supporting information 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. You can also add a variety of supporting objects such as text areas, web pages, and images. 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

## Creating Dashboards

Create a dashboard just like any other worksheet, then add and remove views. This section discusses the following topics:

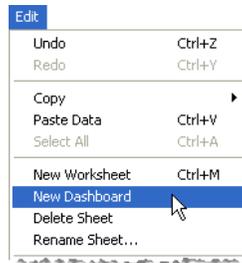
- How to Create a Dashboard
- Adding Views to a Dashboard
- Adding Dashboard Objects
- Removing Views and Objects from a Dashboard

### How to Create 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 and objects.

**To create a dashboard:**

- Select **Edit > New Dashboard**.



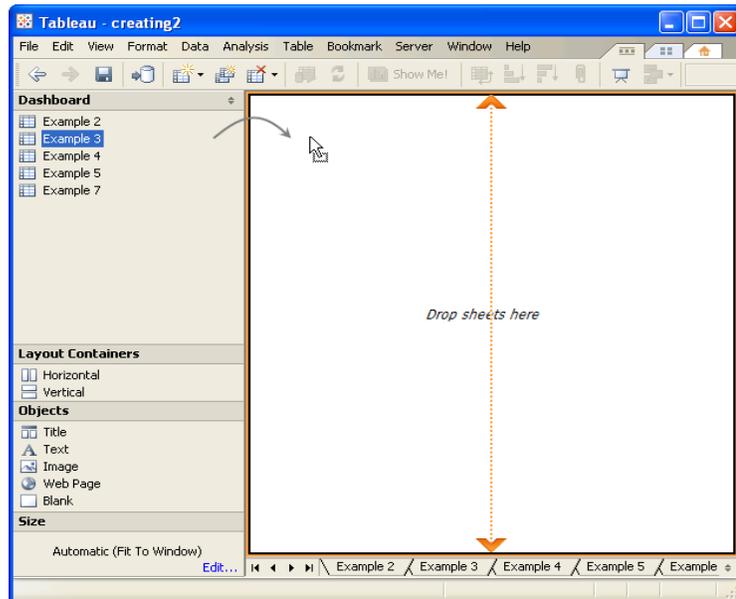
A new tab for the dashboard is added along the bottom of the workbook. Switch to the new dashboard to add views and objects.

### Adding Views to a 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.

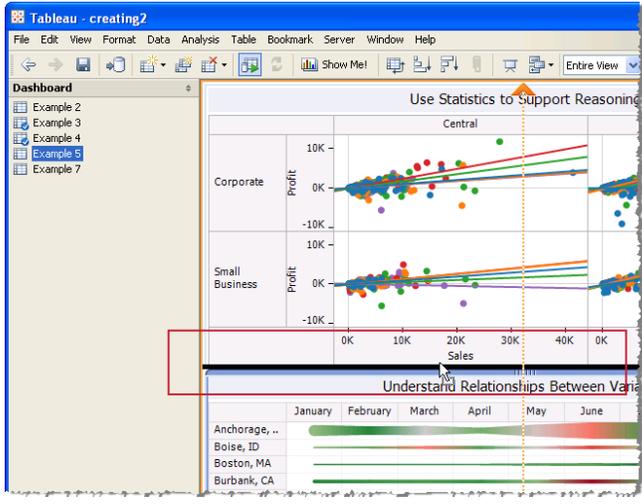
**To add a view to a dashboard:**

- Click and drag a worksheet from the Dashboard window to the dashboard on the right.

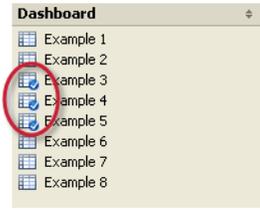




Continue to drag as many of the worksheets to the dashboard as you like. Notice as you drag worksheets around over the dashboard a black bar displays, indicating the various places you can drop it.



After a view is added to the dashboard, the worksheet is marked with a check mark in the Dashboard window. Also, any legends or quick filters that are turned on for the sheet are automatically added to the dashboard.





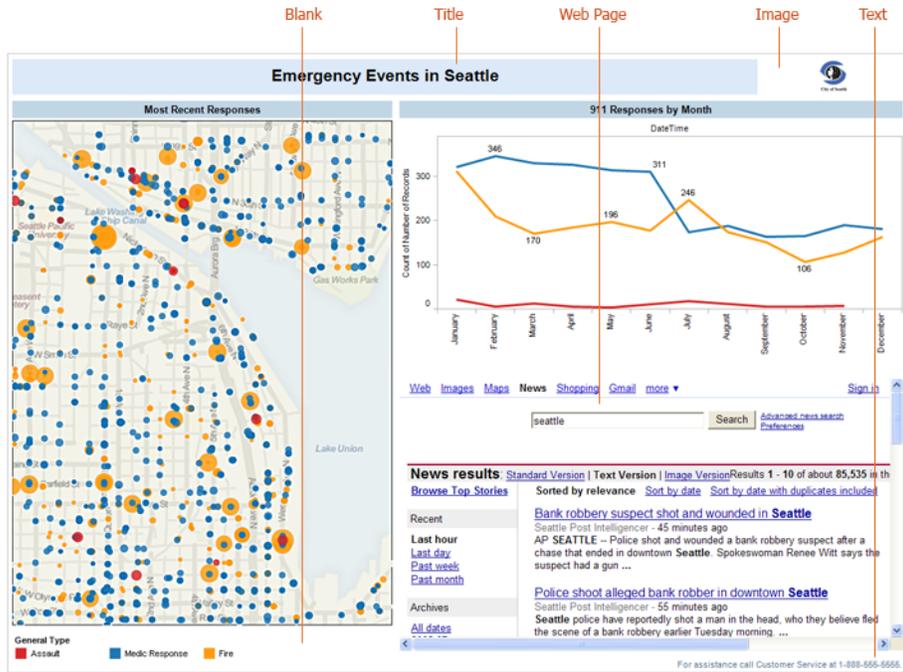
## Adding Dashboard Objects

Dashboards are used to monitor and analyze a collection of related views and information. A dashboard object is an area in the dashboard that can contain supporting information that is not a Tableau view. For example, you can add a text area to include a detailed description. Or you may want to add a web page that is the destination of your hyperlinks. Dashboard objects are listed at the bottom of the dashboard window. You can add titles, text, images, web pages, and blank areas.

### To add a dashboard object:

- Click and drag a dashboard object and drop it on the dashboard. Each of the types of objects are described below:
  - **Title** - The title object allows you to add a title that describes the dashboard. You can place the title object anywhere in the dashboard. The default tile text is the name of the dashboard. Edit the title by double-clicking it in the dashboard. You can also format the title by selecting Format > Dashboard.
  - **Text** - The text object allows you add a block of text to the dashboard. This is useful for adding captions, descriptions, and even copyright information. The text object will automatically resize to best fit where you place it in the dashboard. However you can also resize the text area manually by dragging the edges of the text object.
  - **Image** - You can add static image files to the dashboard. For example you may want to add a logo or descriptive diagram. When you add an image object you are prompted to select an image from your computer. You can also type a URL for an image that is hosted online.
  - **Web Page** - The web page object allows you to embed a web page into your dashboard so you can combine your Tableau content with information from other applications. The web page object is especially useful when you have hyperlinks set up using the Data > Hyperlinks command. If your views include hyperlinks to web pages, you can display those pages in a dashboard by adding a web page object -- the links will automatically open in the dashboard instead of opening a browser window. When you add a web page object you are asked to specify a URL. Refer to “Actions” on page 20-1 to learn more about links.
  - **Blank** - The blank object lets you add blank areas to your dashboard so you can get the layout just right. Resize the blank object by clicking and dragging the edges of the area.

Below is a dashboard that uses the different types of dashboard objects.



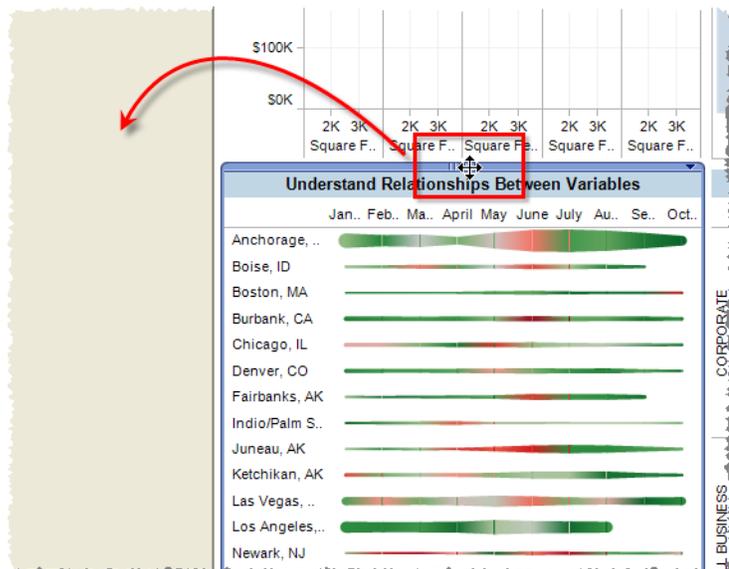


## Removing Views and Objects from a Dashboard

After you add a worksheet or object 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 or object by dragging:

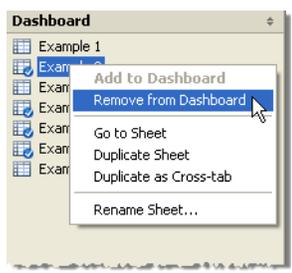
- 1 Select the view you want to remove from the view.
- 2 Click the move handle at the top of the view and drag it off the dashboard.





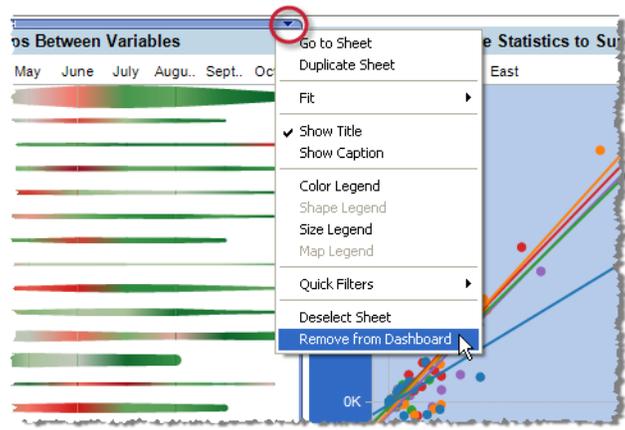
**To remove a worksheet using the Dashboard window:**

- Right-click the worksheet in the Dashboard window and select **Remove from Dashboard**.



**To remove a worksheet or object using the dashboard view menu:**

- 1 Select the view or object in the dashboard that you want to remove.
- 2 Select **Remove from Dashboard** on the dashboard view menu.



---

## Organizing Dashboards

A dashboard can be made up of several views, objects, 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:

- Layout Containers
- Showing and Hiding Parts of a Worksheet
- Rearranging Dashboard Views and Objects

### Layout Containers

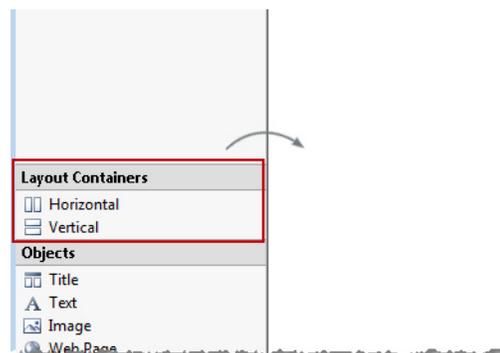
After you create a dashboard you can add sheets and other objects to the dashboard. One type of dashboard object is a layout container. Layout containers help you organize sheets and other objects on a dashboard. These containers create an area in the dashboard where objects automatically adjust their size and position based on the other objects in the container. For example, a dashboard that has a master-detail filter that changes the size of the target view can use a layout container to automatically adjust the other views when the filter is applied.

### Adding Layout Containers

Add a Horizontal Layout Container automatically adjusts the width of dashboard objects. Add a Vertical Layout Container to automatically adjust the height of dashboard objects.

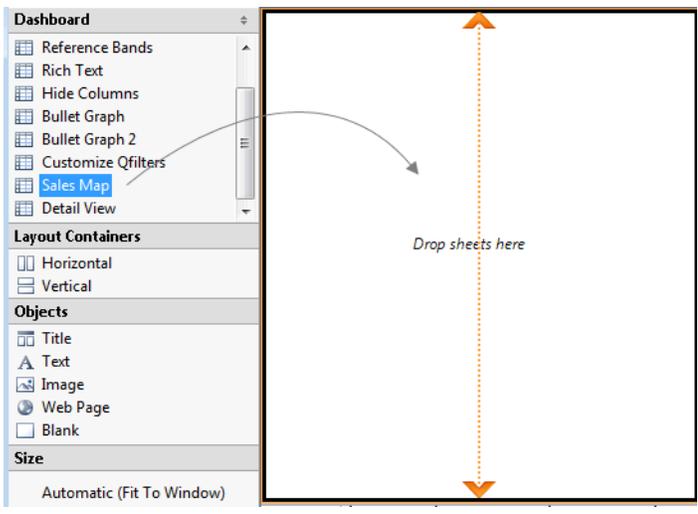
#### To add a layout container:

- 1 Drag a horizontal or vertical layout container to the dashboard.

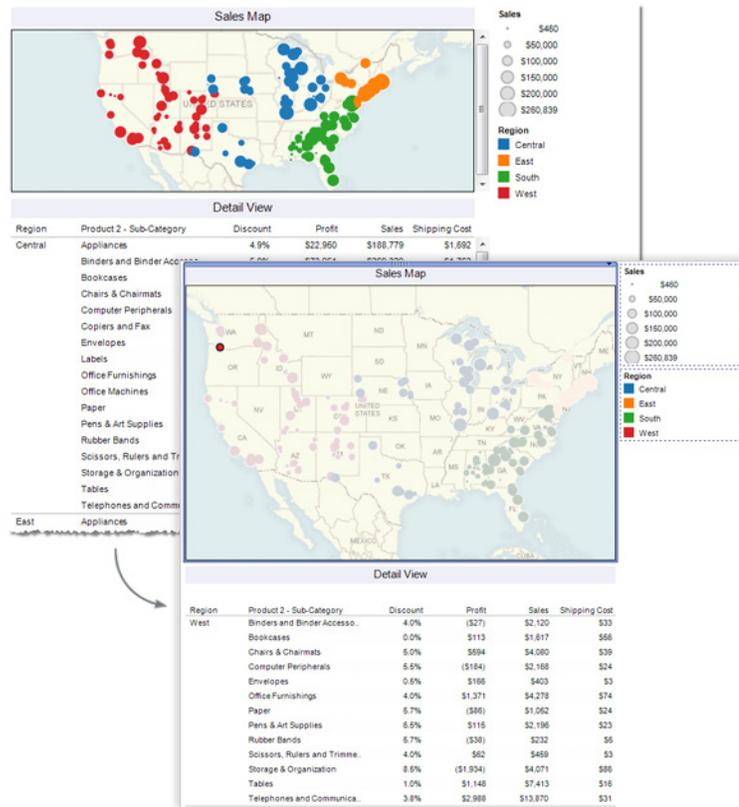




- 2 Add sheets and objects to the layout container. As you hover over the layout container, an arrow indicates how the objects will flow within the container.



- Watch as the objects move and resize. For example, in the dashboard below, a filter is applied that causes the text table to get shorter. Because both views are in a vertical layout container, the map automatically resizes into the new space.



You can add as many layout containers are you want and you can even add layout containers within other containers.

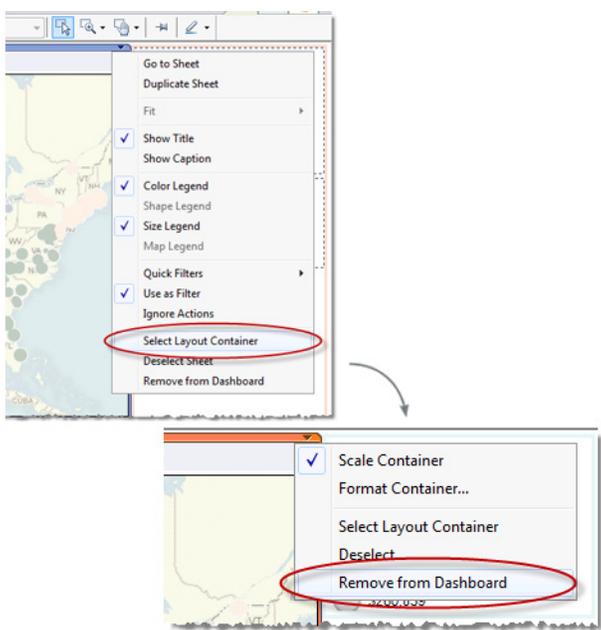


## Removing Layout Containers

When you remove a layout container, it and all of its contents are deleted from the dashboard.

### To remove a layout container:

- 1 Select an object in that layout container you want to delete.
- 2 Open the drop-down menu in the upper right corner of the selected object and select **Select Layout Container**.
- 3 Open the drop-down menu for the selected layout container and select **Remove from Dashboard**.



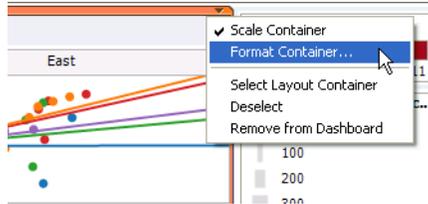
## Formatting Layout Containers

You can specify a border style for layout containers so that you can visually group objects in the dashboard.

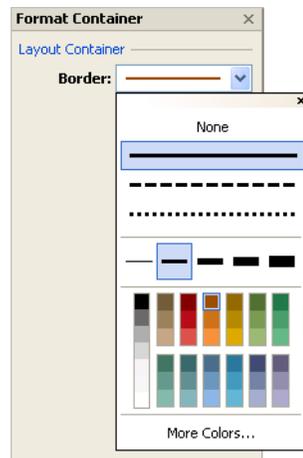
### To format layout container borders:



- 1 Open the drop-down menu for the layout container you want to format and select **Format Container**.



- 2 In the Format window, specify a line style, thickness, and color for the border.



Layout containers are useful for all kinds of dashboards and give you increased control for how the objects automatically flow in the dashboard as you apply filters. Also, use layout containers when comparing multiple bar charts or bullet graphs. In this case, the bar heights are automatically adjusted so that the bars in both sheets stay aligned.

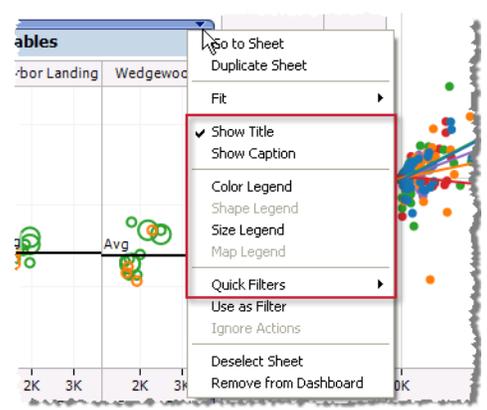


## Showing and Hiding Parts of a Worksheet

As you drag worksheets to the dashboard, the view from the worksheet and its legends and quick filters are automatically displayed. However, you may want to hide some parts of the worksheet such as legends, titles, captions, and quick filters. You can show and hide these parts of the worksheet using the drop-down menu in the upper right corner of the views in the dashboard.

### To show and hide parts of the worksheet:

- 1 Select a view in the dashboard.
- 2 Open the drop-down menu in the upper right corner of the selected view and select the items you want to show. 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.

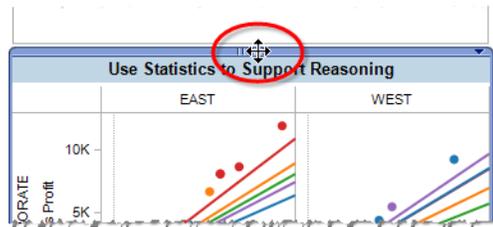
## Rearranging Dashboard Views and Objects

Rearrange the views, objects, 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, object, legend, or quick filter:



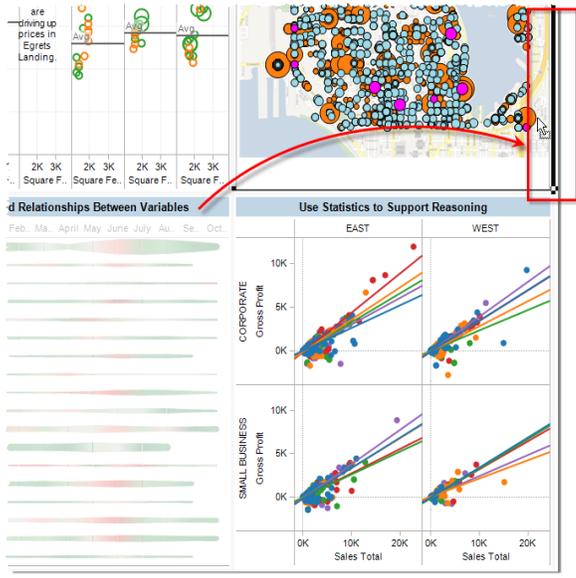
- 1 Select the view, legend, or quick filter that you want to move.
- 2 Click and drag the move handle at the top of the selected item to a new location.





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.





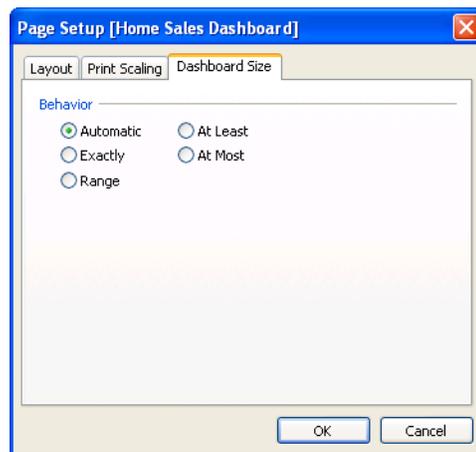
## Setting the Dashboard Size

By default, the views and objects 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 specify a specific 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.
  - **Range** - the dashboard scales between the specified minimum and maximum sizes, after which it scroll bars or white space will display.
  - **At Least** - the dashboard automatically resizes to fill the application window, but will not scale smaller than a specified minimum size. When viewed smaller, scroll bars will display.
  - **At Most** - the dashboard automatically resizes to fill the application window, but will not scale larger than the specified maximum size. When viewed larger, white space will be added.
- 4 If you select any option other than Automatic, you must define a size. 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**.

---

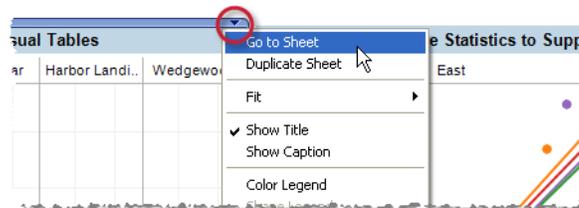
## Understanding Dashboards and Worksheets

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 in the filmstrip, sheet sorter, or in the tabs along the bottom of the workbook.

### To jump to a worksheet from a dashboard view:

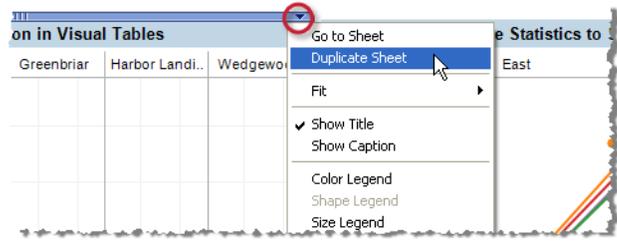
- 1 Select the view you want to see full size.
- 2 Select **Go to Sheet** on the dashboard view menu.





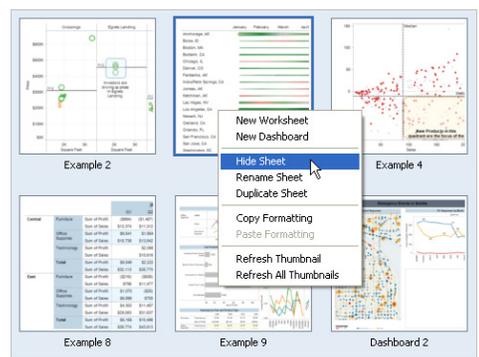
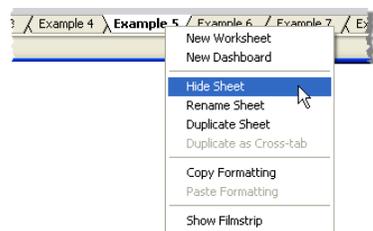
**To duplicate a worksheet from a dashboard view:**

- 1 Select the view you want to duplicate.
- 2 Select **Duplicate Sheet** on the dashboard view menu



**To hide a worksheet:**

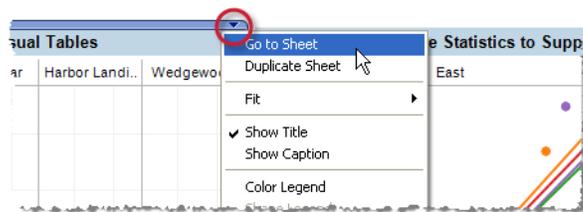
- Right-click the worksheet tab or thumbnail and select **Hide**.





### To Show a Hidden Sheet:

- 1 Open the dashboard that uses the hidden sheet.
- 2 Select the hidden sheet in the dashboard and then select **Go to Sheet** on the dashboard view menu.



Alternatively you can right-click the hidden sheet in the Dashboard window and select Go to Sheet. The sheet opens and its tab appears along the bottom of the workbook once again.

# Publishing to Tableau Server

<b>Overview</b> . . . . .	<b>30-2</b>
<b>Publishing Workbooks</b> . . . . .	<b>30-3</b>
How to Publish Workbooks to the Server . . . . .	30-3
Saving Workbooks to Tableau Public . . . . .	30-7
Specifying Permissions . . . . .	30-9
Showing and Hiding Worksheets . . . . .	30-15
Embedding Passwords . . . . .	30-16
Schedules . . . . .	30-17
<b>Publishing Data Sources</b> . . . . .	<b>30-19</b>
<b>Opening Workbooks from the Server</b> . . . . .	<b>30-23</b>
<b>Importing Data Sources from the Server</b> . . . . .	<b>30-26</b>
<b>User Filtering</b> . . . . .	<b>30-29</b>
How to Create a User Filter . . . . .	30-29
Previewing User Filters . . . . .	30-34
Editing User Filters . . . . .	30-35
Publishing with User Filters . . . . .	30-37
Example - Setting User Filters . . . . .	30-38
Filter Using a List of Users in your Data Source . . . . .	30-45

---

## Overview

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.

---



## Publishing Workbooks

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 embed database passwords for automatic authentication on the web. This section discusses the following topics:

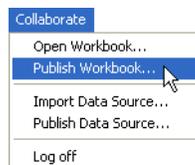
- How to Publish Workbooks to the Server
- Saving Workbooks to Tableau Public
- Specifying Permissions
- Showing and Hiding Worksheets
- Embedding Passwords
- Schedules

### How to Publish Workbooks to the Server

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 or URL (Examples: sales\_server, https://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.





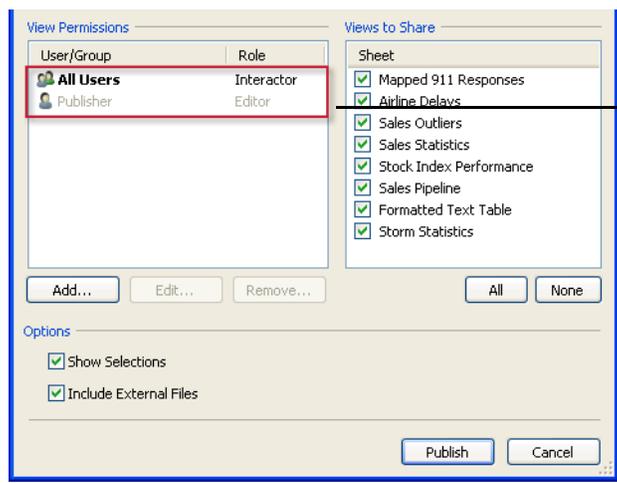
**Note** Use the drop-down list to select an existing workbook on the server. When you publish using an existing workbook name, the workbook on the server is overwritten with your workbook. You must be allowed the **Write** capability to overwrite workbooks on the server.

- 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”).



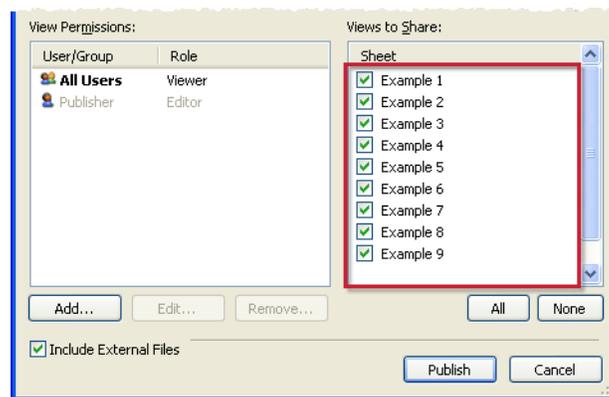
- 6 Optionally specify permissions to allow or deny access to the workbook on the server. By default all users can interact with the workbook and you, as the publisher, are allowed all capabilities. Refer to “Understanding Capabilities” on page 30-13 to learn more.



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 30-15 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 “Publishing with User Filters” on page 30-37 to learn more.
- 9 Select whether to **Show Selections**. When this option is selected, any selections you’ve made in the workbook will be published to the server. Refer to “Schedules” on page 30-17 to learn more about schedules.
- 10 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.
- 11 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 30-16 to learn more.
- 12 If the workbook uses Extract connections you can optionally add the workbook to a refresh schedule.
- 13 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.

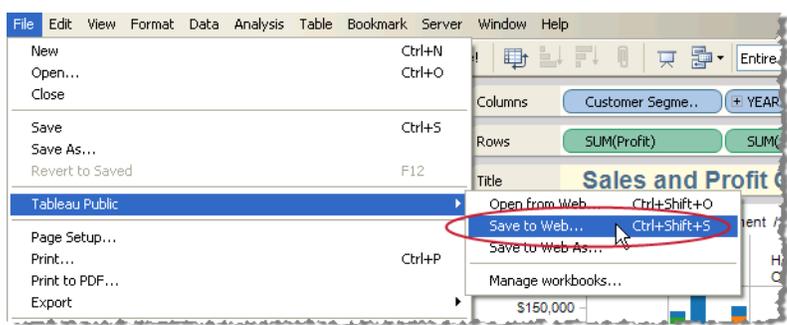
---

## Saving Workbooks to Tableau Public

In addition to saving to an internal server, you can also save workbooks to Tableau Public. Tableau Public is a free service that lets anyone publish interactive data to the web. Once on the web, anyone can interact with the data, download it, or create their own visualizations of it. Workbooks saved to Tableau Public must have fewer than 100,000 rows of data. Use Tableau Data Extracts to reduce the size of your data before saving. Refer to “Extracting Data to the Data Engine” on page 9-1 to learn more about creating extracts.

### To save workbooks to Tableau Public:

- 1 Select **File > Tableau Public > Save to Web**.





- 2 When prompted, login using your Tableau Public account. If you don't have one, you can create one for free.

Tableau Public Login

tableau+public

Login

If you already registered, enter your email address and password below to login.

Email: name@email.com

Password: ●●●●●●

Login Forgot your password?

Not registered? Create an account for FREE!

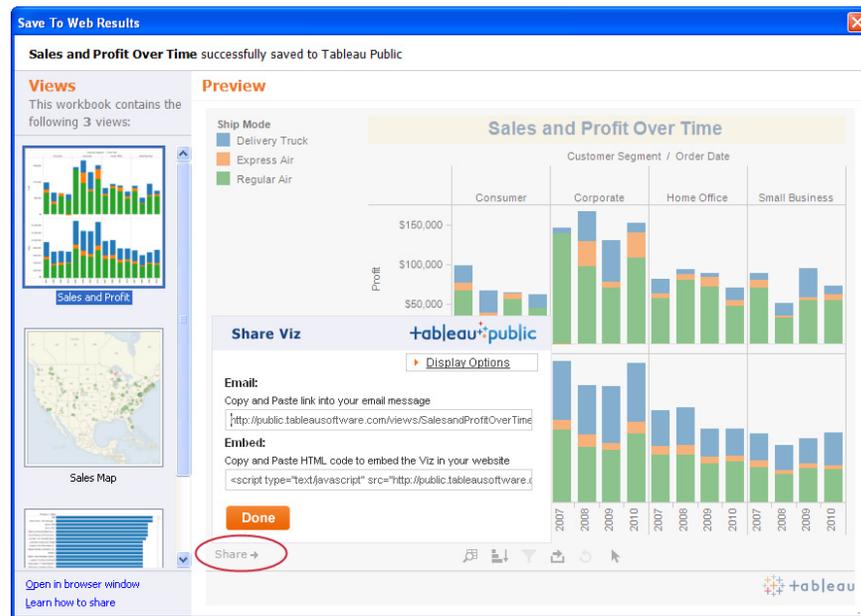
- 3 Type a name for the workbook and click **Save**.

Save Workbook to the Public Web

Name: Sales and Profit Over Time

Save Cancel

The published workbook displays so you can preview all of the saved sheets. Select a sheet and click the Share button in the lower left corner of the view to get a link you can email or embed into a webpage.



Workbooks and the underlying data saved to Tableau Public are publicly available. You can manage your content saved to the web online at <http://public.tableausoftware.com>.

Learn more about Tableau Public online at [www.tableausoftware.com/public](http://www.tableausoftware.com/public).

## Specifying Permissions

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

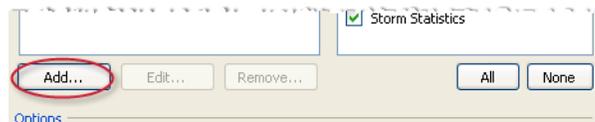


## Adding Permissions

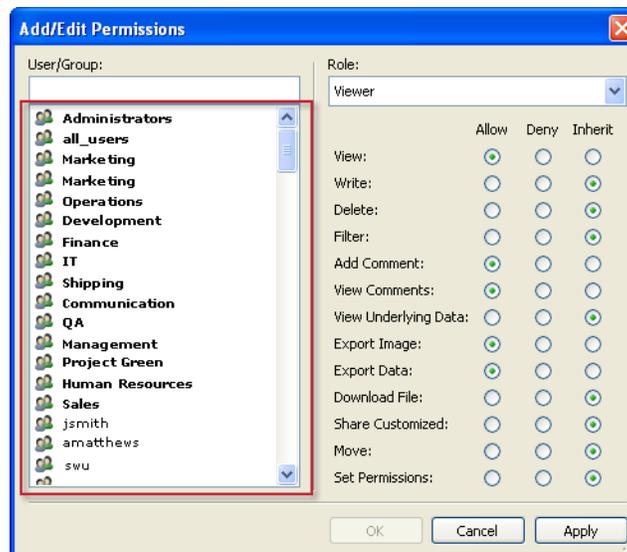
You can add permissions to a workbook in the Publish Workbook dialog box. After a workbook has been published, you can also modify and add permissions directly on Tableau Server. Refer to the Tableau Server online help to learn more.

### To add permissions:

- 1 In the Publish Workbook to Tableau Server dialog box, click the Add button in the bottom left corner.



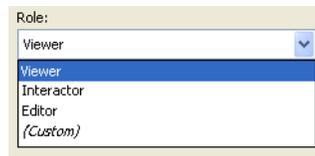
- 2 Select who you want authorize. You can select a group of users or a specific user in the User/Group list.



- 3 Select a role from the drop-down list to assign common sets of capabilities. Tableau Server has the following three pre-set roles:

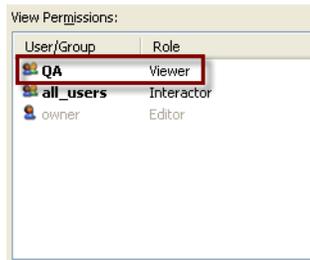


- **Viewer** - allows users to view the workbook on the server, as well as add and view comments.
- **Interactor** - allows users to view the workbook on the server, apply filters, view the underlying data, export the image, and export the data. All other capabilities are inherited from the user's group 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 30-13 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.

---

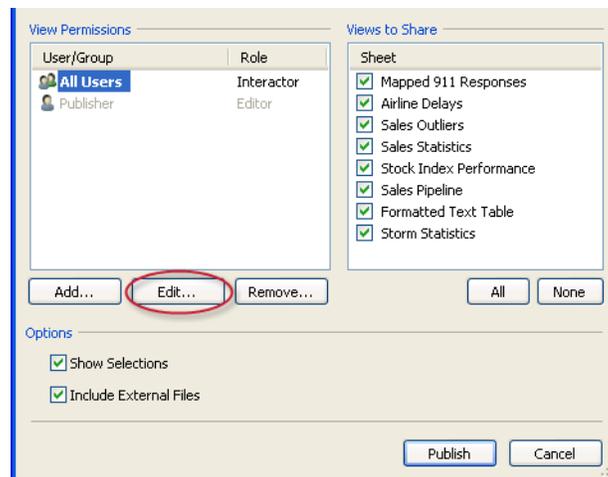


## Editing and Removing Permissions

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 30-10 to learn more about adding permissions.

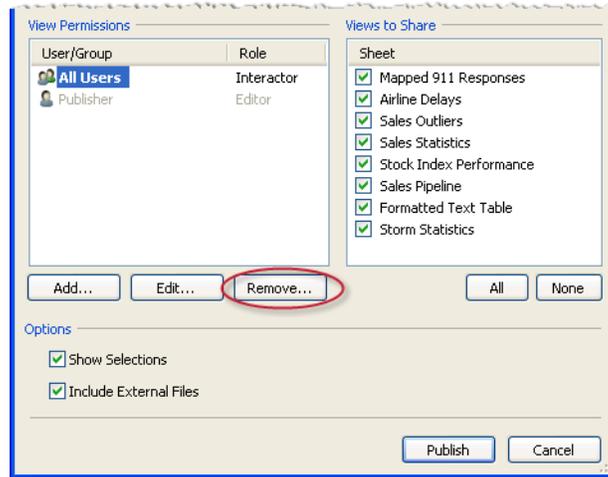
### To edit existing permissions:

- 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



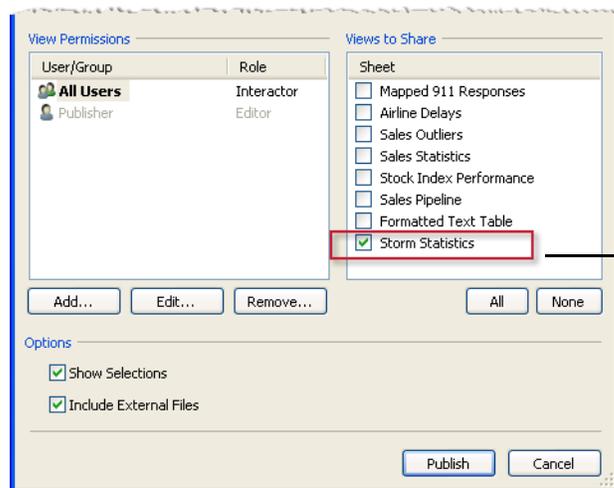
<b>Capability</b>	<b>When allowed, users can...</b>
View Underlying Data	View the underlying data for the worksheets in the workbook.
Export Image	Export each worksheet as an image.
Export Data	Export the data to a text file.
Download file	Open the workbook from the server using Tableau Professional
Share Customized	Share a customized view, making their changes public.
Move	Move workbooks between projects.
Set Permissions	Specify permissions for the workbook and all of the contained views.



## Showing and Hiding Worksheets

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 30-10 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.

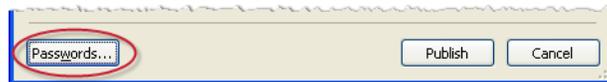


## Embedding Passwords

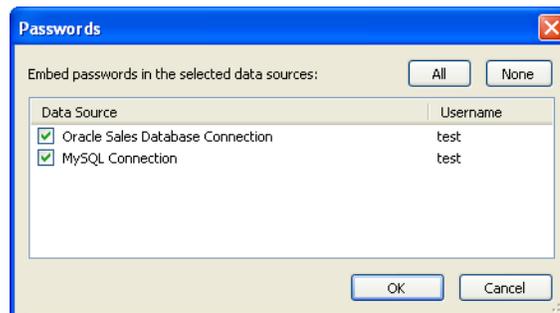
If your workbook connects to a data source that requires a user name and password, by default people viewing it on the server will also be prompted for a user name and password. However, as the author you can choose to specify a database user and password that will always be used when a view is opened 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.

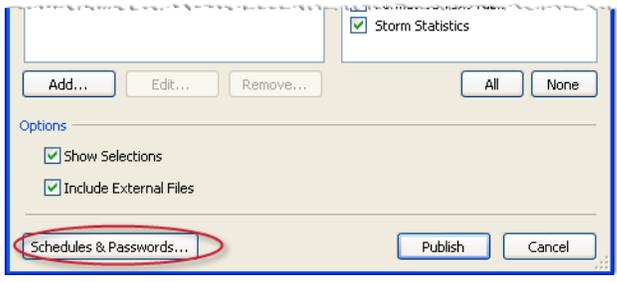
---



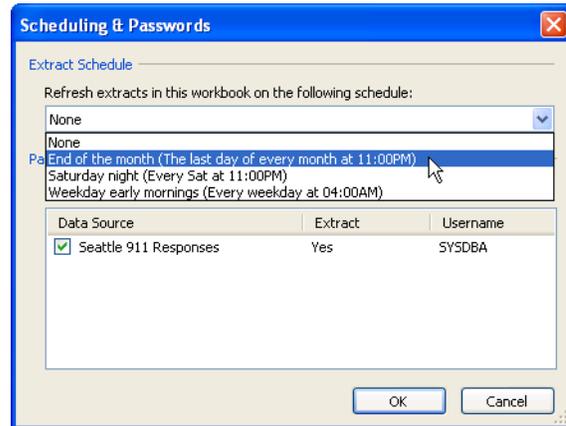
## Schedules

When you publish workbooks that connect to extracts you can schedule the extracts to be refreshed automatically. That way you don't have to republish the workbook every time the underlying data has updated and you can still get the performance of a data extract. For example, let's say you have a workbook that connects to a large data warehouse that is updated weekly. Instead of publishing a workbook that queries the live data, you can create an extract including just the data necessary. This increases performance and limits queries to the live database. Then you can add that workbook to a schedule so that the extract gets refreshed each week with the updated data from the data warehouse.

Schedules are created and managed on the server by an administrator. However, an administrator can allow you to add a workbook to a schedule when you are publishing from Tableau Desktop. If this option has been enabled, a Schedules & Passwords button shows in the publish dialog when you are publishing a workbook that connects to a data extract.



In the Scheduling & Passwords dialog box, select a schedule to add the workbook to. All data sources that require authentication must have an embedded password so that the extract can be refreshed. This includes data sources that are not extracts.



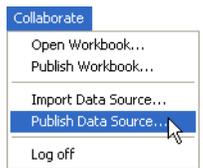


## Publishing Data Sources

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 or URL (Examples: sales\_server, https://sales\_server)
  - **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

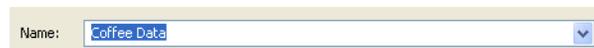


- 3 In the Publish Data Source dialog box, select a project to publish the data source into.

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.

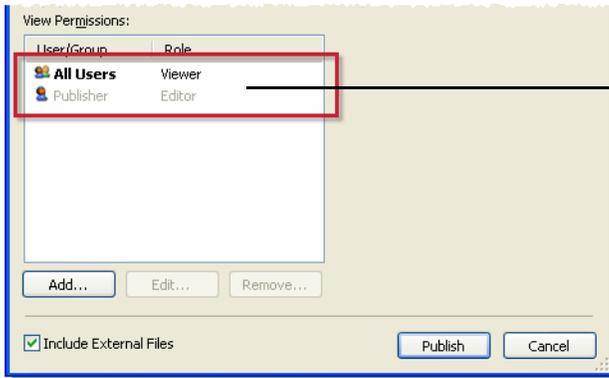
---

- 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”).



- 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 30-9 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 30-37 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**.

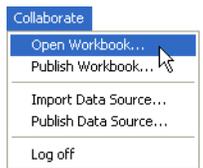


## Opening Workbooks from the Server

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**.



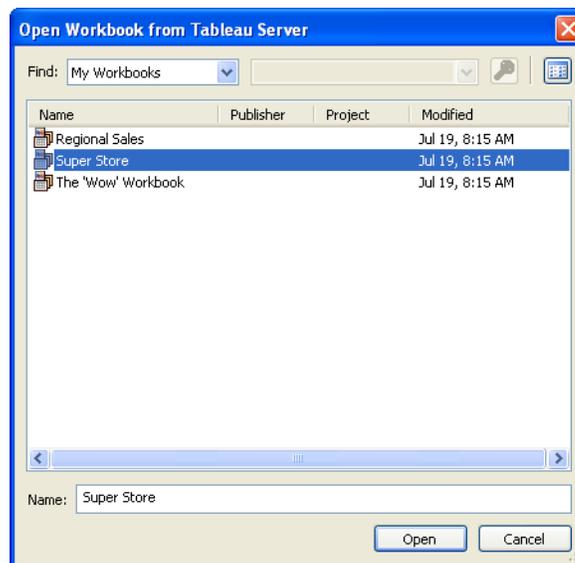
- 2 Type the following into the Tableau Server Login dialog box:
  - **Server:** the server name or URL (Examples: sales\_server, https://sales\_server)
  - **User Name:** your user name. If Tableau Server is configured to use Active Directory, type your Windows user name, otherwise, type your Tableau 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.



---

## Importing Data Sources from the Server

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:**

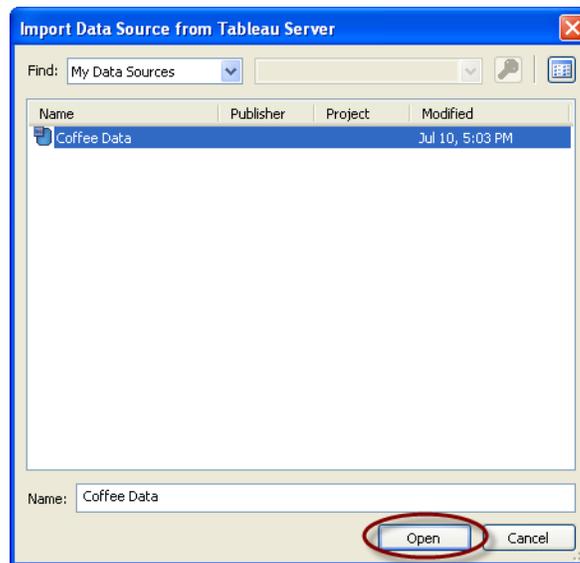
- 1 Select **Collaborate > Import Data Source**.
- 2 Type the following into the Tableau Server Login dialog box:
  - **Server:** the server name or URL (Examples: sales\_server, https://sales\_server)
  - **User Name:** your user name. If Tableau Server is configured to use Active Directory, type your 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

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 manager, 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 view is adjusted based on who is logged in and looking at it.

This section discusses the following topics

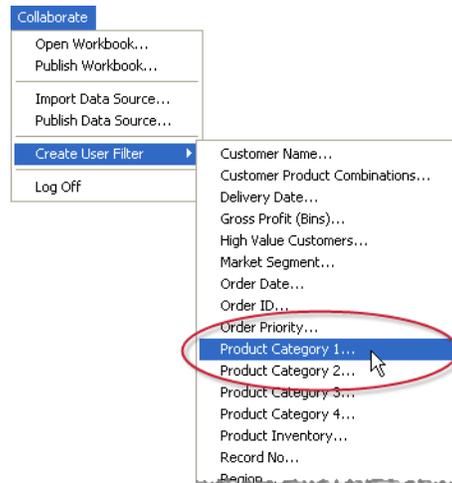
- How to Create a User Filter
- Previewing User Filters
- Editing User Filters
- Example - Setting User Filters

### How to Create a User Filter

You can create as many user filters as you like for a given view.



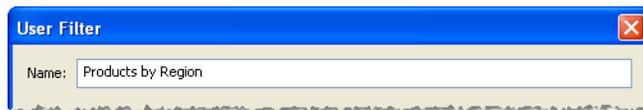
- 1 Select **Collaborate > Create User Filter** and then select a field that you want to filter on. For example, if you are limiting product data each person can access select the Product field.



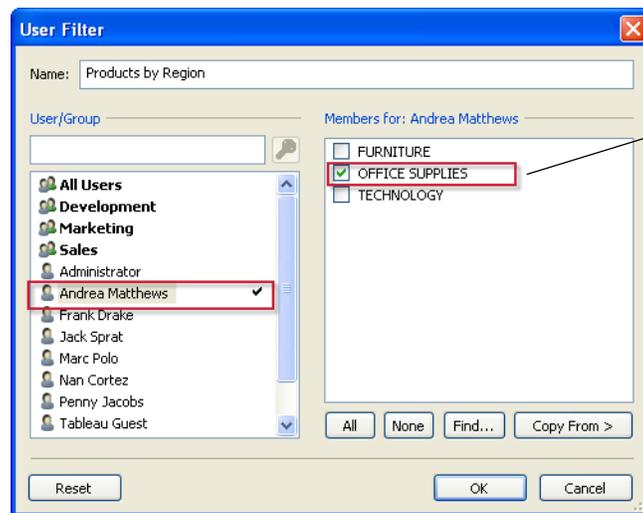
- 2 If necessary log in to Tableau Server. Refer to “How to Publish Workbooks to the Server” on page 30-3 to learn more about logging in.



- 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.

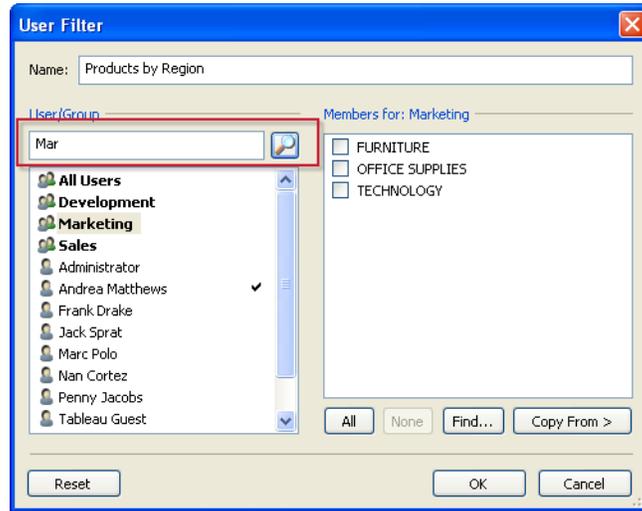


Here Andrea Matthews will only be able to see the Office Supplies data in the view.

- 5 When finished click **OK**.
- 6 The user filter now shows in the Sets area of the Data window. Drag the new set to the Filters shelf to begin using it in a view. Refer to “Example - Setting User Filters” on page 30-38 for a step-by-step example of defining user filters.

### 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.

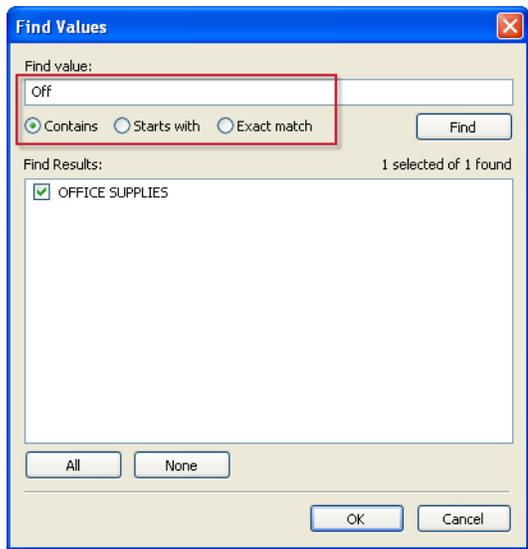


- 2 Type all or part of the member name into the **Find value** text box at the top of the Find Values dialog box and then click **Find**. You can change the search criteria by selecting



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.

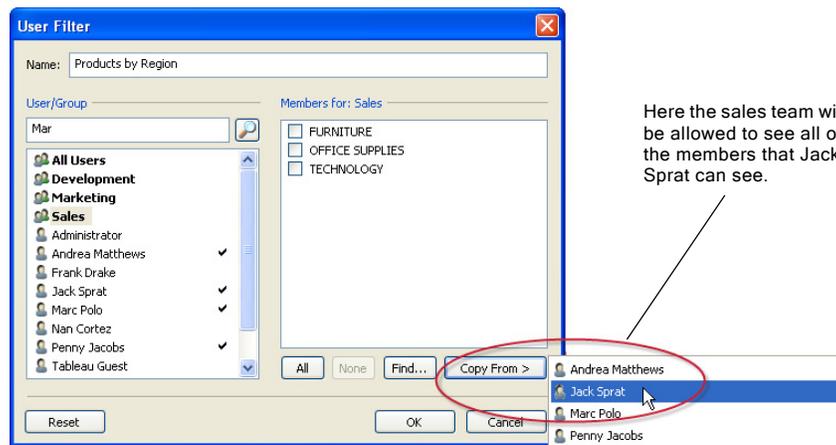
### Copying member selections from other users

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.



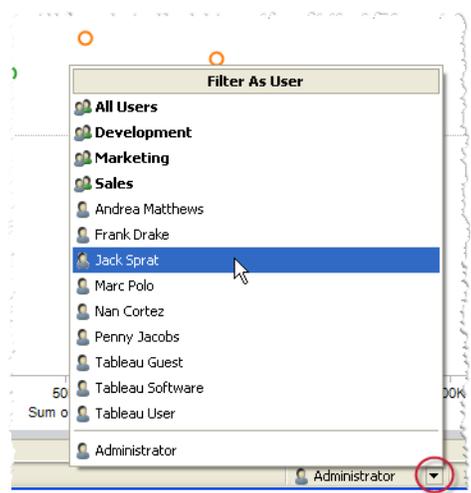
## Previewing User Filters

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.

---

## Editing User Filters

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 30-3 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.

---

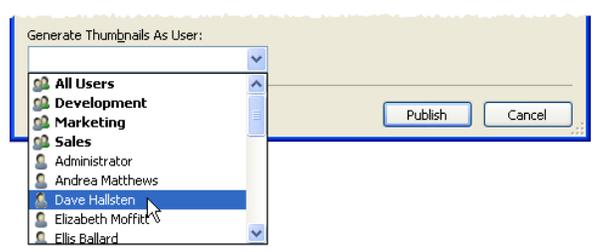


## Publishing with User Filters

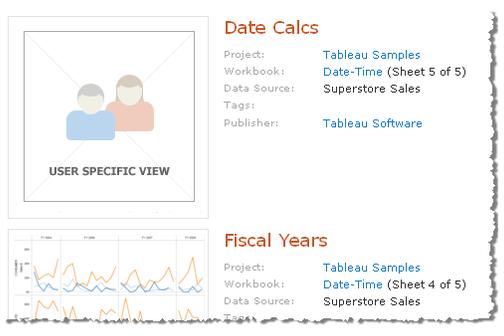
When you publish a workbook it and its contained sheets are represented with thumbnail images on the server. You can specify what these thumbnail images will look like by selecting to generate thumbnails as a specific user. For example, if you want the thumbnail image to show all three regions of a sales 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.

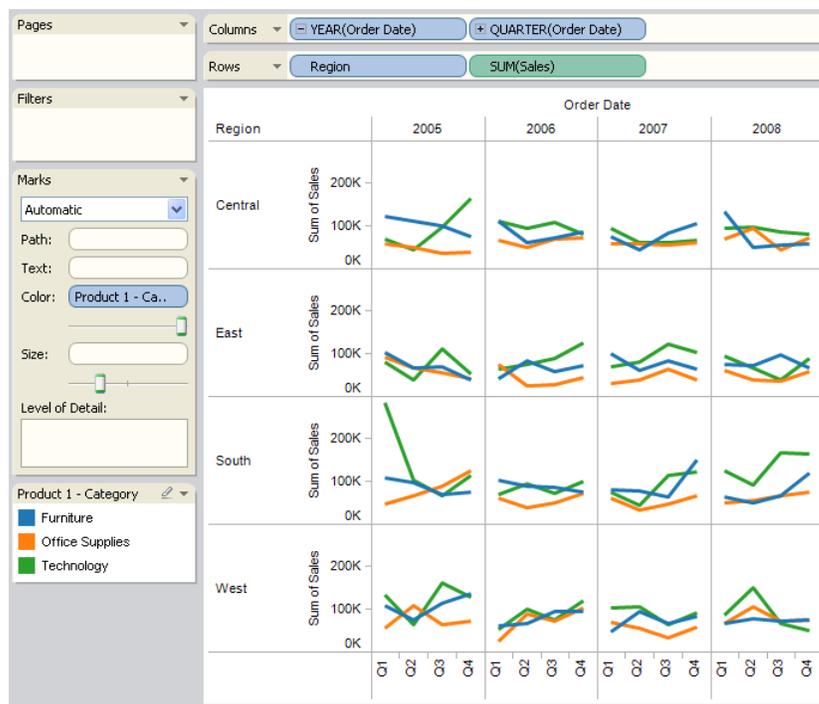
---



## Example - Setting 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.

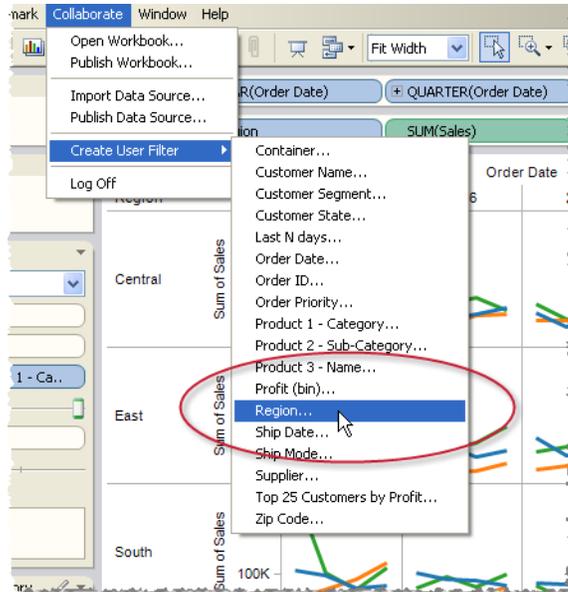
The sales report is shown below.



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.



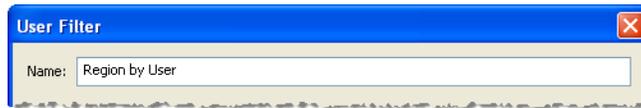
- 1 Select **Collaborate > Create User Filter**. Then select **Region** because that the field you want to use for filtering the view.



- 2 You may need to log into the server if are not already logged in. Log in by typing the Server name or URL, your user name, and password. You user name and password is either specific for the server or is your Windows user name and password.



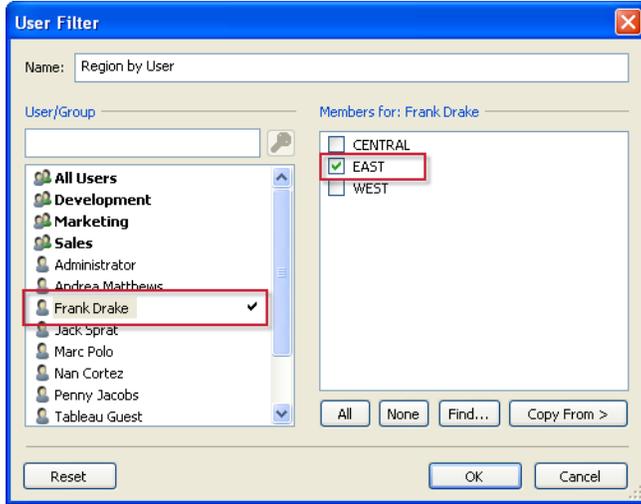
- 3 Type a name for the user filter. In this case we'll name it Region by User. After specifying the filter, it will display in the Sets area of the Data window by this name.



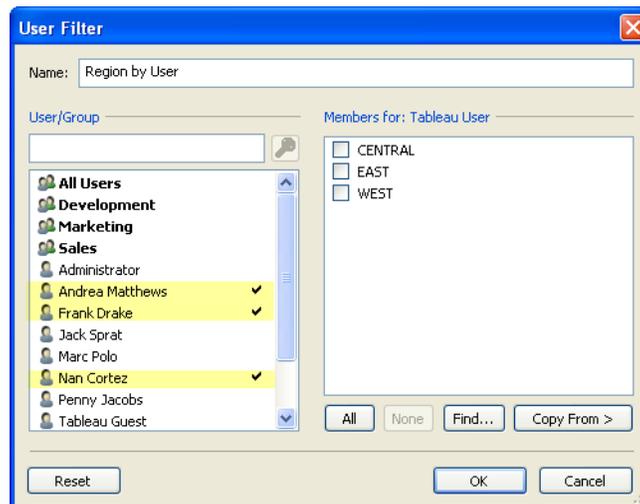
- 4 Then select a user or group on the left side for who you want to filter data. Then on the right side of the dialog box, select the members of the Region field that the selected user



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.

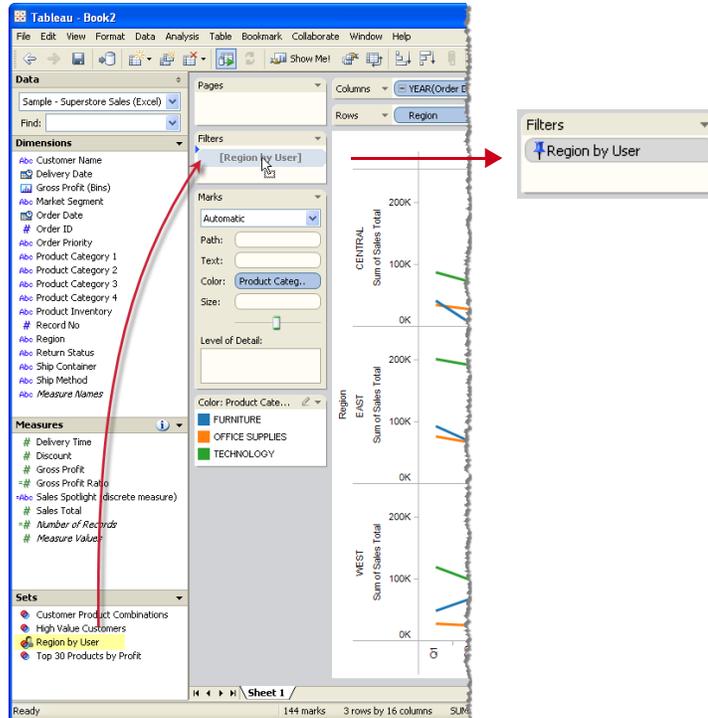


- 5 Repeat this process as many times as possible. Here we have three managers so Frank Drake gets to see the East, Andrea Matthews gets to see the Central region, and Nan Cortez gets to see the West region.



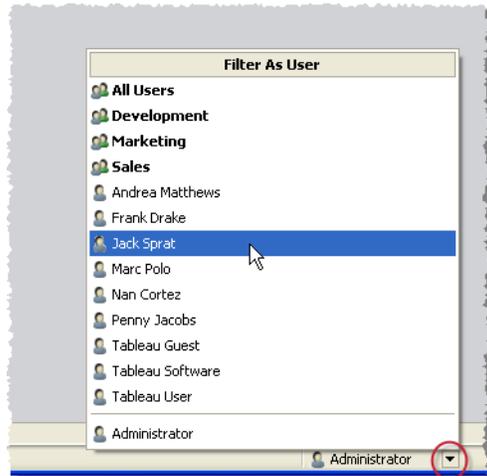
- 6 When finished click **OK**.
- 7 The user filter displays in the Sets area of the Data window. Now that it is defined you need to apply it to a view. Simply drag the Region by User user filter to the Filters shelf.

The filter becomes a context filter. Refer to “Context Filters” on page 16-50 to learn more.

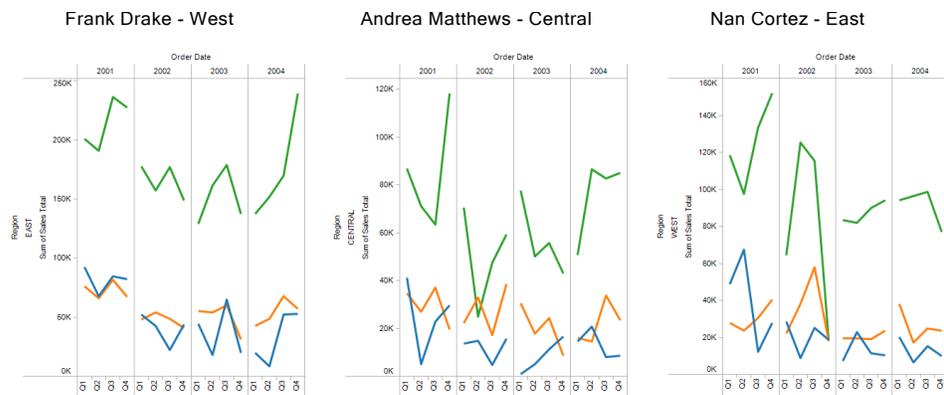


**Note** When you apply a user filter to a view, it’s possible that you don’t see anything. For example, in this case you have not allowed yourself to see either of the three regions. You can edit a user filter by right-clicking it in the Data window and selecting **Edit**.

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.



Now when you publish the view to the server each users will only see the data you have allowed them to see.

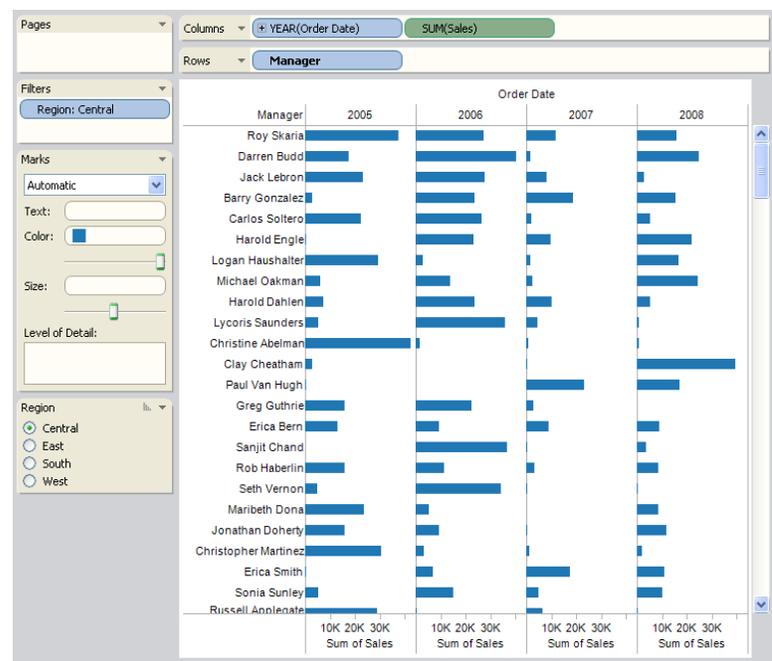


## Filter Using a List of Users in your Data Source

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 User functions.

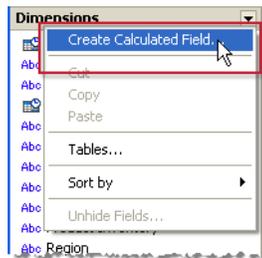
### Example of a user filter using users from the data source:

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.

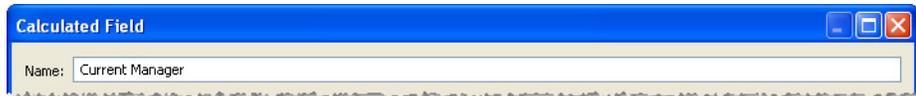




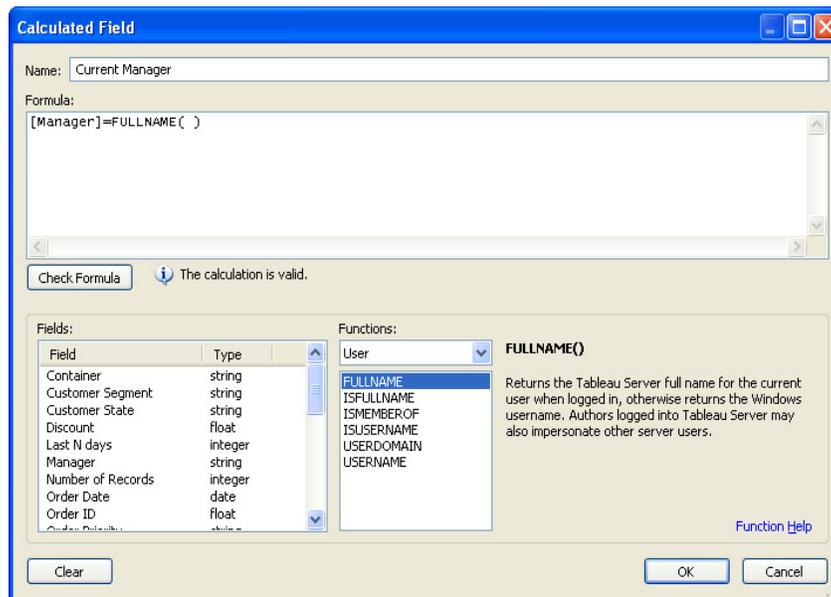
- 1 If you have not already logged in, select **Collaborate > Log In** to log into Tableau Server.
- 2 Create a calculated field by selecting **Create Calculated Field** on the Data window menu.



- 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 User 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

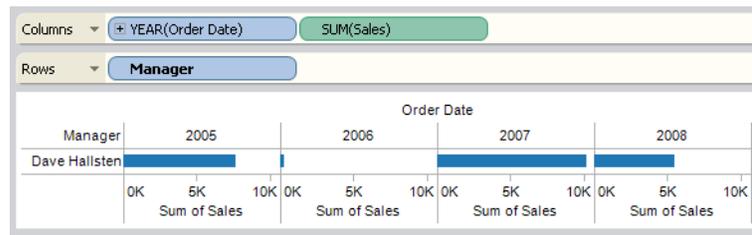


that case you should select nothing and simply click **OK**. Then you can preview the view as another user and edit the filter to include True.

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 30-34 to learn more.

The view below shows what Dave Hallsten would see if he was logged in.



Refer to “User Functions” on page 34-23 for more information about each special function available.

# Saving and Exporting

---

<b>Saving Your Work</b> . . . . .	<b>31-3</b>
Workbooks . . . . .	31-3
Bookmarks . . . . .	31-7
Packaged Workbooks . . . . .	31-5
<b>Reverting Workbooks</b> . . . . .	<b>31-9</b>
<b>Exporting Your Work</b> . . . . .	<b>31-10</b>
Export Data . . . . .	31-10
Export as an Image . . . . .	31-21
<b>Exporting the Data Source Connection</b> . . . . .	<b>31-24</b>

---

## Overview

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.



## Saving Your Work

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.

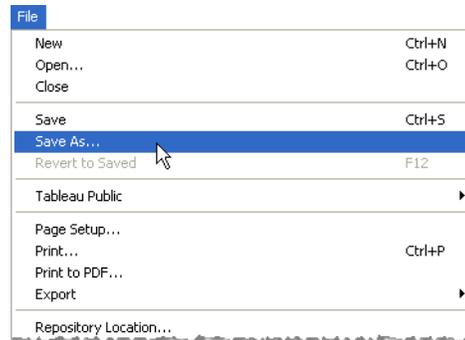
### Workbooks

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.

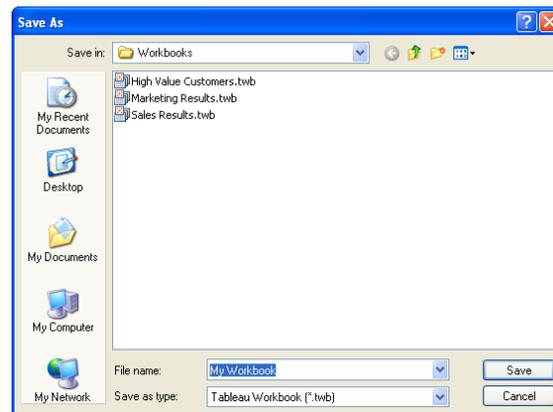


Tableau saves the file with the .twb extension. The default location is the **Workbooks** folder of the Tableau repository. However, you can save Tableau workbooks to any directory you choose.



---

**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.

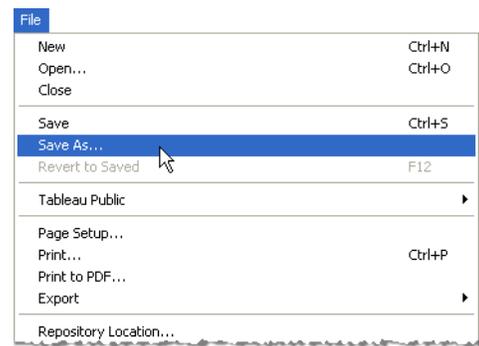
## Packaged Workbooks

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 and images, 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.

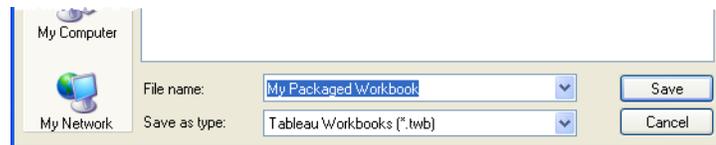
### To save a packaged workbook:

- 1 Select **File > Save As**.

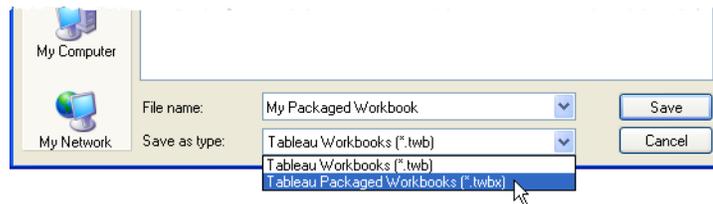




- 2 Specify a file name for the packaged workbook in the Save As dialog box.



- 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
- Custom Shapes
- Custom Geocoding



---

**Note** If you are sharing packaged workbooks that contain Microsoft Excel or Access 2007 data sources, the people opening the workbook must either have Microsoft Excel and Access 2007 or the Office 2007 Data Connectivity Components installed on their machines. The data connectivity components are available on the Tableau Download Drives page.

---

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.

## Bookmarks

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.

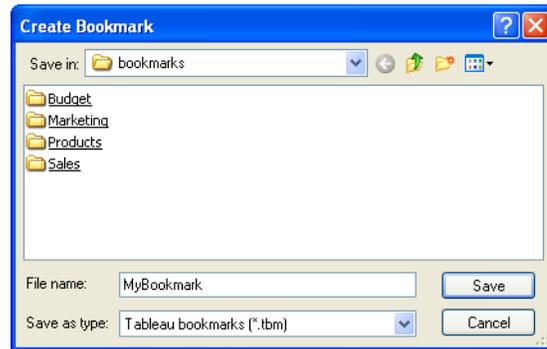


Tableau saves the file with the .tbn extension. The default location is the **Bookmarks** folder of the Tableau Repository. However, you can save bookmarks to any location you choose. Bookmarks that are not stored in the Tableau repository do not display in the **Bookmark** menu.

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.

---

**Note** While bookmarks are generally a snap shot of the worksheet and include the data connection, formatting, etc. A bookmark does not include parameter values and the current page setting on the Pages shelf.

---

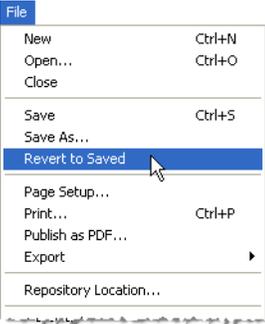


## Reverting Workbooks

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.



---

**Note** The Revert to Saved option is not available for workbooks that use data extracts.

---

---

## Exporting Your Work

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 32, “Printing”.

---

**Note** You can also use Tableau to present your data. Refer to “Presentation Mode” on page 3-16 to learn more about presenting your work.

---

### Export Data

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.

---

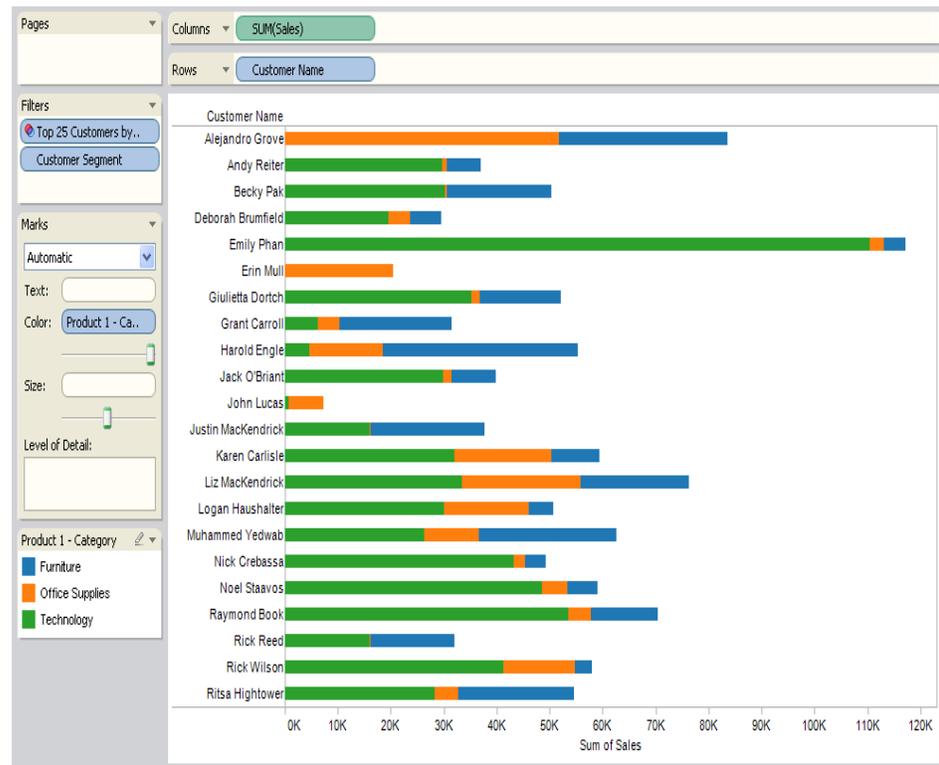
**Note** Excel 2003 worksheets hold 64,000 rows of data. If you copy more than this number of rows, Tableau pastes only the first 64,000 rows of the selected set into Excel. For datasets larger than 64,000 it is recommended that you export to Access instead.

---

When exporting data, you should keep these rules in mind:

- You can select any portion of a data view to export. Refer to Chapter 19, “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.



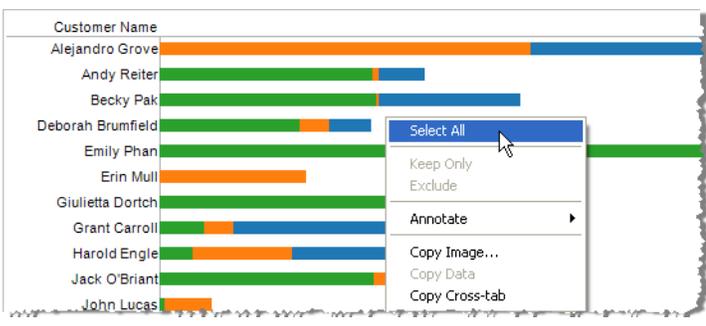
**Note** When exporting data to a Microsoft Office application, remember that Tableau only supports Office 2000 or higher (including Office 2007).

### Copy Records To Clipboard

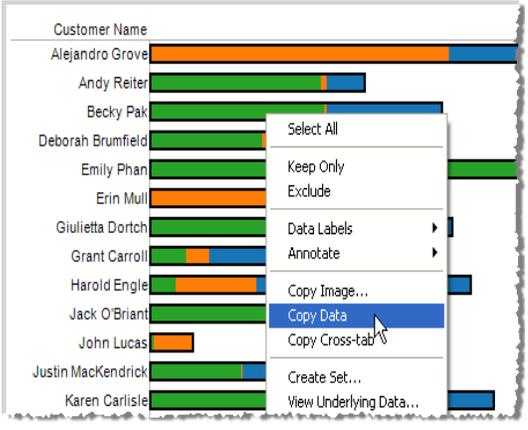
Typically this function is used to copy records from Tableau into Microsoft Excel. To create an Excel spreadsheet from Tableau data, follow these three steps:



- 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 **Customer 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	12959.78	
4	Deanra Eno	TECHNOLOGY	EAST	10455.7	
5	Speros Goranitis	TECHNOLOGY	EAST	14272.76	
6	Meg Tillman-Sach	FURNITURE	WEST	14584.2	
7	Skye Norling-Chri	TECHNOLOGY	WEST	11106.19	
8	Noel Staavos	TECHNOLOGY	WEST	29689.57	
9	Stephanie Ulgrigh	TECHNOLOGY	EAST	51280.73	

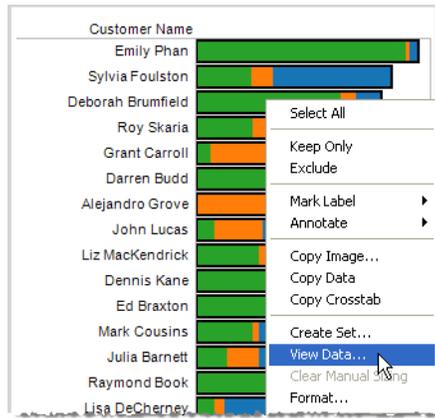
### Copy Underlying Records to Clipboard

Copying underlying data can be used to copy the disaggregated data behind a view. For a full discussion of drill-through, refer to “View Data (Drill-Through)” on page 19-12.

To copy underlying records:



- 1 Select the desired data in Tableau.
- 2 Right-click the selected records and select **View 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**.



- 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	996.66
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	896.31
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:



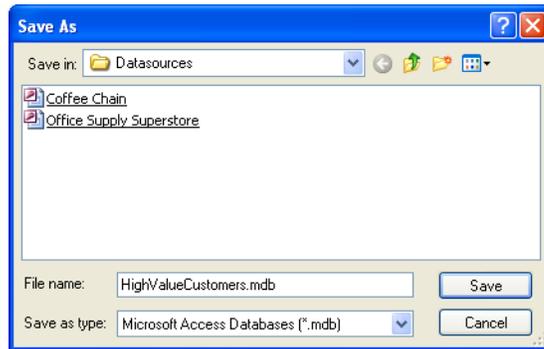
1 Select the desired data in Tableau. For this example, all the data are selected.



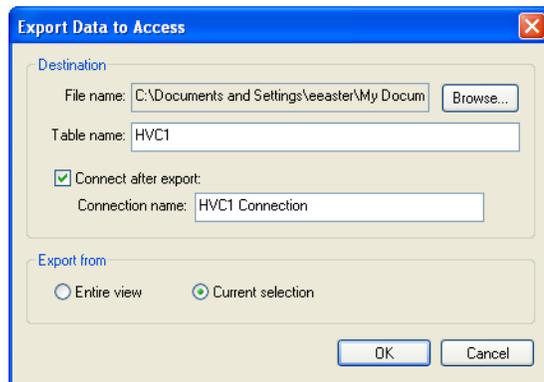
2 Select **File > Export > Data**.



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.

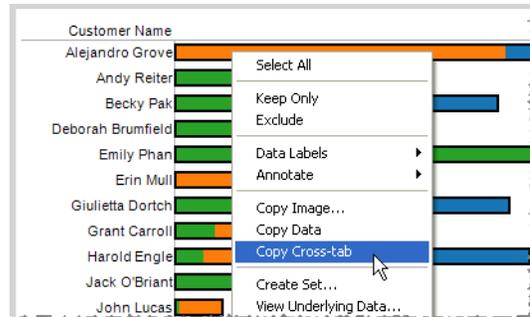


- You can use this feature for aggregated views only. It cannot be used on disaggregated views of data, because a cross-tab is by definition an aggregated view of data. In other words, the **Aggregate Measures** option on the **Analysis** menu must be on in order for this function to work properly. For a full discussion of Aggregated and Disaggregated data refer to “Aggregations” on page 21-4

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**.

A screenshot of Microsoft Excel showing a worksheet with pasted data. The data is in a cross-tab format with columns for categories and values. The data is as follows:

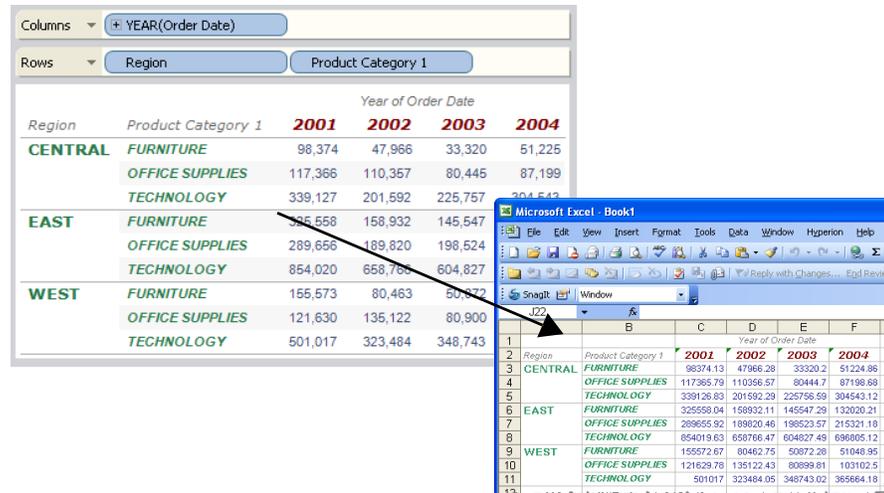
	A	B	C	D	E
1			CENTRAL	EAST	WEST
2	Barbara Fisher	FURNITURE		14,966	
3	Barbara Fisher	OFFICE SUPPLIES	4,211	11,473	1,645
4	Barbara Fisher	TECHNOLOGY	2,362	8,722	2,759
5	Dario Medina	FURNITURE	6,265	2,782	6,786
6	Dario Medina	OFFICE SUPPLIES	5,595	4,403	3,076
7	Dario Medina	TECHNOLOGY	8,782	40,055	7,560
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	David Flashing	FURNITURE	58	4,233	2,345
12	David Flashing	OFFICE SUPPLIES	9,860	49,800	15,322
13	David Flashing	TECHNOLOGY	8,074	17,687	26,059

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 to the Data Engine" on page 9-1.

## Export as an Image

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

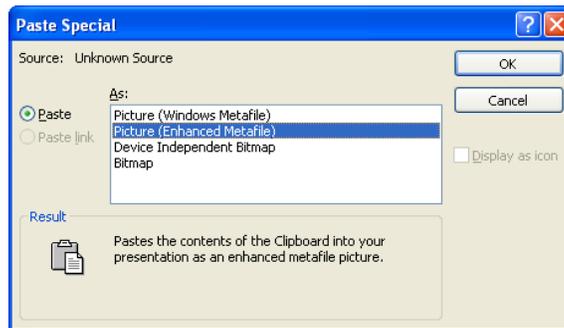
- Publish as PDF

### 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 32-8.



## Exporting the Data Source Connection

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 "Replacing Field References" on page 5-71.

# Printing

---

<b>Overview</b> . . . . .	<b>32-2</b>
<b>Printing Results</b> . . . . .	<b>32-3</b>
Page Setup . . . . .	32-3
Printing . . . . .	32-6
<b>Publish as PDF</b> . . . . .	<b>32-8</b>
<b>Printing the Help</b> . . . . .	<b>32-10</b>
Printing the Online Help . . . . .	32-10

---

## Overview

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.

## Printing Results

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.

### Page Setup

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:

- General – select which elements you want to show on the printed page, specify how to treat headers and breaks, and specify 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.

### General

Use the General 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, size legend, and map legend.

Specify how to handle headers and breaks. The headers refer to the headers in each of your views. When you select the option to repeat the headers and legends on each page the row and column headers will show when a view breaks across several pages.

Select the break pages on pane boundaries to avoid page breaks in the middle of a cell in table.

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 13-14.



## Layout

Use the Layout tab to specify the page margins, centering, and legend layout.

**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.

**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.

## 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 from Tableau Desktop. However, the orientation settings will be used as the default when you publish the workbook to Tableau Server or Tableau Public. For more information on publishing to a PDF refer to “Publish as PDF” on page 32-8.

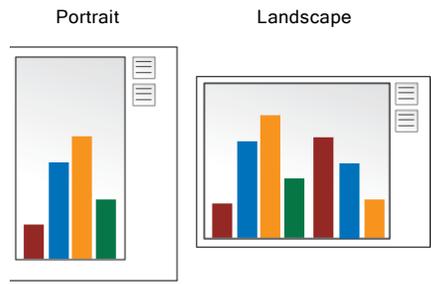
**Print Scaling.** You can scale your view to fit within a single page or scale it across multiple pages. Select from the following options:

- 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. Leaving the text box blank will allow the view to expand across as many pages as necessary in the given direction.

**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.



These page orientation settings are used as the default settings when you publish the workbook to Tableau Server or Tableau Public.

### 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:

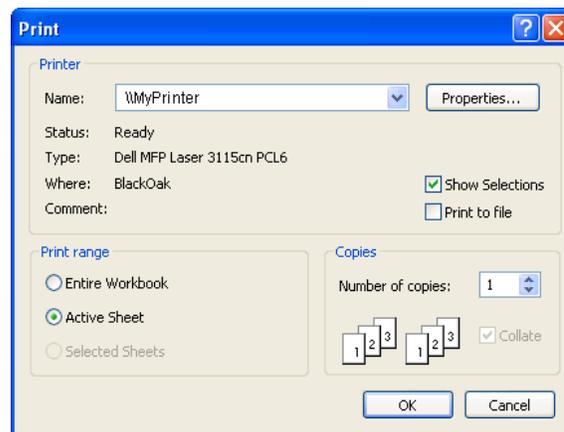
- Data Update Time – displays the last time the data was updated. This is especially useful when you are using data extracts.
- Sheet Name – displays the name of the worksheet.
- Workbook Name – displays the name of the workbook.
- PageCount – displays the total number of pages in a worksheet based on the fields on the Pages shelf.
- 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.
- Full Name – displays the user name and domain of the person logged in. If you are logged into Tableau Server, this is the Server username otherwise the Windows account is use.
- User Name – displays the user name of the person logged in. If you are logged into Tableau Server, this is the Server username, otherwise the Windows account information is used.



- **Field Values** – displays the relevant field value for the selected field. This is useful when you are working with filters to that limit the data shown in the view. If multiple values are relevant this option displays as “Multiple.”

## Printing

After you have specified the Page Setup settings (refer to “Page Setup” on page 32-3), you can print by selecting **File > Print**. In the Print dialog, select a printer, decide whether to show selections, specify a print range, and select the number of copies you want to print.



### Show Selections

When you print a workbook, you can choose to Show Selections. When this option is selected any selections you’ve made in the views will be maintained while printing.

### 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.

---

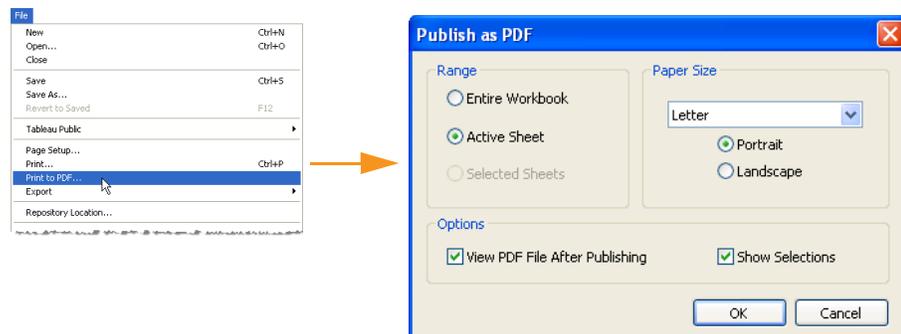
---

## Publish as PDF

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 32-3).
- 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.



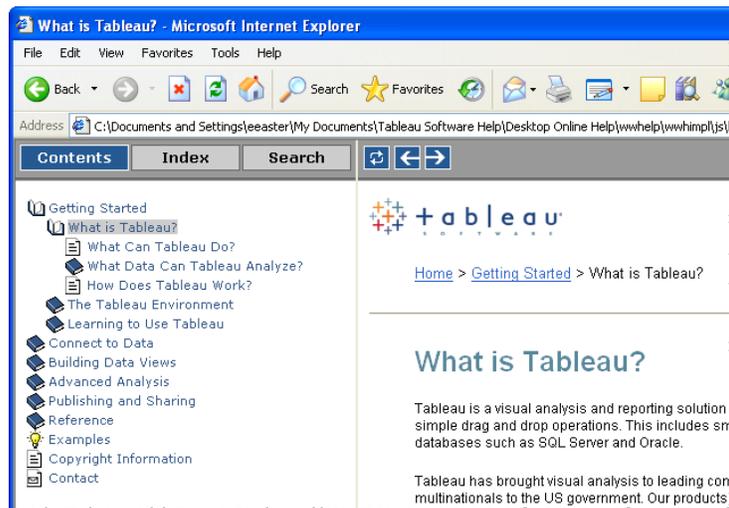
- Selected Sheets - Publishes the selected sheets.
- 4 Select a paper size. If you select Unspecified the paper size will expand to the necessary size to fit the entire view on a single page.
- 5 Optionally select **View PDF file after publishing** if you want to automatically open the PDF when you are done publishing. This option is only available if you have Adobe Acrobat Reader or Adobe Acrobat installed on your computer.
- 6 Optionally select whether to Show Selections. When this option is selected the selections in the views are maintained in the PDF.
- 7 Click **OK** and select where you want to save the PDF. Then click **Save**.

---

## Printing the Help

### 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 45-1 to learn more.

# Building Data Views

## Examples

<b>Bar Chart</b> . . . . .	<b>33-4</b>
Bar Chart–Basic . . . . .	33-5
Bar Chart–Color-Encoded . . . . .	33-7
Bar Chart–Color Encoded Nested Table . . . . .	33-9
Bar Chart–Filtered and Color-Encoded . . . . .	33-10
Bar Chart–Color Encoded and Sorted . . . . .	33-14
Bar Chart–Filtered, Color Encoded, and Level of Detail . . . . .	33-19
<b>Text Table</b> . . . . .	<b>33-22</b>
Text Table–Basic . . . . .	33-22
Text Table–Filter and Color Encode . . . . .	33-26
Duplicate as Crosstab . . . . .	33-28
<b>Line Chart</b> . . . . .	<b>33-30</b>
Line Chart–Basic . . . . .	33-31
Filter and Color Encode a Nested Table . . . . .	33-33
Line Chart–Filter and Color-Encode Multiple Measures . . . . .	33-35
Line Chart–Filter and Path Encode . . . . .	33-37
<b>Scatter Plot</b> . . . . .	<b>33-39</b>
Scatter Plot–Basic . . . . .	33-40
Scatter Plot–Filter and Color Encode . . . . .	33-42
Scatter Plot–Filter and Level of Detail . . . . .	33-45
Scatter Plot–Filter, Level of Detail, and Color Encode . . . . .	33-47
Scatter Plot–Filter, Level of Detail, Color Encode, and Shape Encode . . . . .	33-52
<b>Heat Map</b> . . . . .	<b>33-55</b>
Heat Map–Basic . . . . .	33-56
Heat Map–Size Encoded . . . . .	33-58
Heat Map–Shape and Size Encoded . . . . .	33-61
<b>Gantt Bar Chart</b> . . . . .	<b>33-63</b>
Gantt Bar Chart–Standard . . . . .	33-63
Gantt Bar Chart–Nonstandard . . . . .	33-66



**Pie Chart . . . . . 33-68**  
Pie Chart- Maps . . . . . 33-68



## Overview

This section includes detailed examples that show you how to build data views. The examples are categorized first by chart type, and then by other criteria such as encodings, sorting, and so on. All examples use the Superstore Sales Excel data source that's included with Tableau.

There are many types of data views you can create in Tableau, and this collection of examples cannot include every type. However, many basic types are included in one form or another. The following types of views are discussed:

- Bar Chart
- Text Table
- Line Chart
- Scatter Plot
- Heat Map
- Gantt Bar Chart
- Pie Chart

If you have specific views in mind, you should browse the examples to find those that most closely meet your needs.

---

**Note** For simplicity, the terms “Dimension” and “Measure” in this section are shorthand for “Discrete Dimensions” and “Continuous Measures.” The features described here, however, may also apply to Continuous Dimensions and Discrete Measures. For a full discussion of data types and data roles refer to “Data Types & Roles” on page 11-34.

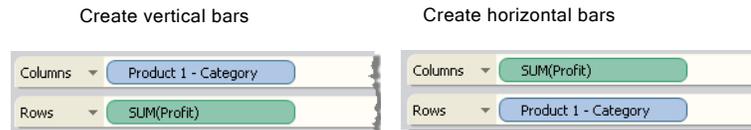
---

---

## Bar Chart

Bar charts are a great way to compare data across categories, or when you want to break data down into stacked bars. In Tableau, you typically create a bar chart by placing a dimension and a measure as the inner fields on the **Rows** and **Columns** shelves (or vice versa).

A bar chart uses the bar mark type. Tableau can automatically select this mark type if the data view matches one of the two types shown below and the **Mark** menu is set to Automatic. You can have as many fields as you like on these shelves provided a dimension and a measure are innermost.



Alternatively, you can manually select Bar from the **Mark** menu for any data view. Refer to “Bar Mark” on page 12-16 for information about this mark type.

The following examples show you how to use bar charts to explore your data:

- Bar Chart–Basic – Create a bar chart with vertical marks using one dimension and one measure.
- Bar Chart–Color-Encoded – Create a bar chart with vertical marks using one dimension and one measure, and color encode the data using a dimension.
- Bar Chart–Color Encoded Nested Table – Create a nested bar chart with vertical marks using multiple dimensions and one measure. Color encode the data using a dimension and remove mark borders to better see closely-spaced marks.
- Bar Chart–Filtered and Color-Encoded – Create a nested bar chart with vertical marks using multiple dimensions and multiple measures. Color encode the data using a dimension and filter the data by categories.
- Bar Chart–Color Encoded and Sorted – Create a bar chart with horizontal marks using one dimension and one measure. Color encode the data using a measure and sort the data by a categorical field.
- Bar Chart–Filtered, Color Encoded, and Level of Detail – Create a bar chart with horizontal marks using one dimension and one measure. Color encode the data using a

---

dimension, filter the data numerically, and provide an additional level of detail to each colored bar segment.

## **Bar Chart–Basic**

This view displays the sum of the sales for each year as a bar chart. You can create the view with the following procedure.

- 1 Place the **Order Date** dimension on the **Columns** shelf.

The date is automatically aggregated by year and headers are created with labels given by the dimension member names.

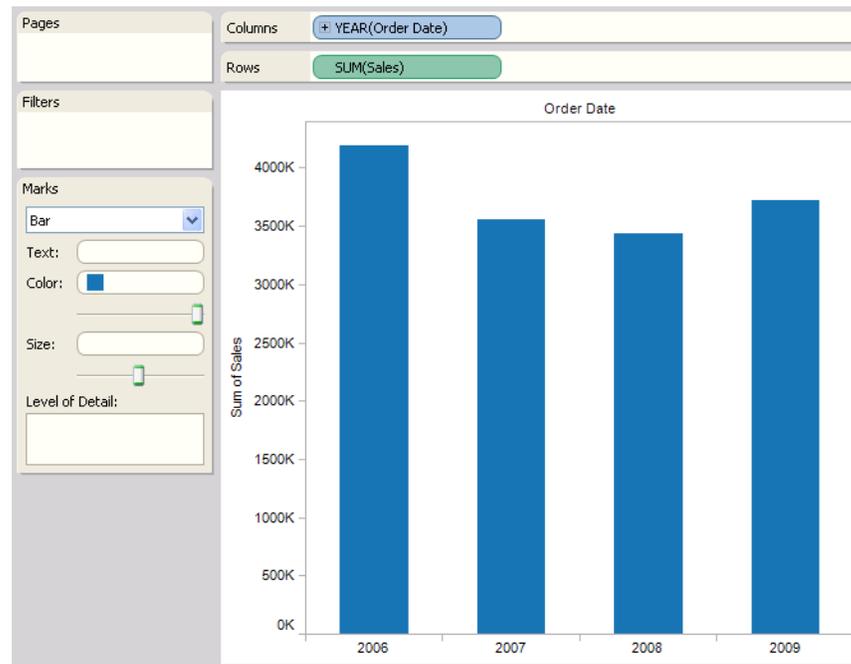
- 2 Place the **Sales** measure on the **Rows** shelf.

The measure is automatically aggregated as a summation and an axis is created with a label given by the name of the field.

- 3 Select **Bar** from the Mark menu.

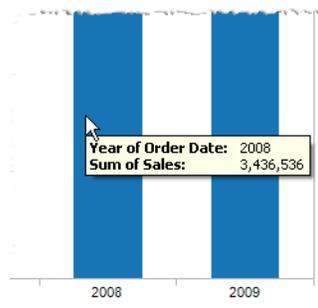
Tableau automatically selects **Line** as the mark type because a date dimension is displayed. If a different dimension was used that was not a date, Tableau would have automatically selected **Bar**. For this particular example, the **Bar** mark makes it easier to analyze the data.

The view is shown below.



The marks are vertical because the axis is vertical. The length of each mark represents the sum of the sales for each year. For example, the sum of the sales in 2008 is \$3,436,536.

You can display this number with tool tips by placing the mouse cursor over the mark.





## Bar Chart–Color-Encoded

This view displays the sum of the sales for each product category and is color-encoded by ship mode. You can create the view with the following procedure.

- 1 Place the **Product 1 - Category** dimension on the **Columns** shelf.

The date is automatically aggregated by year, and headers are created with labels given by the dimension member names.

- 2 Place the **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.

- 3 Place the **Ship Mode** dimension 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.

The view is shown below.

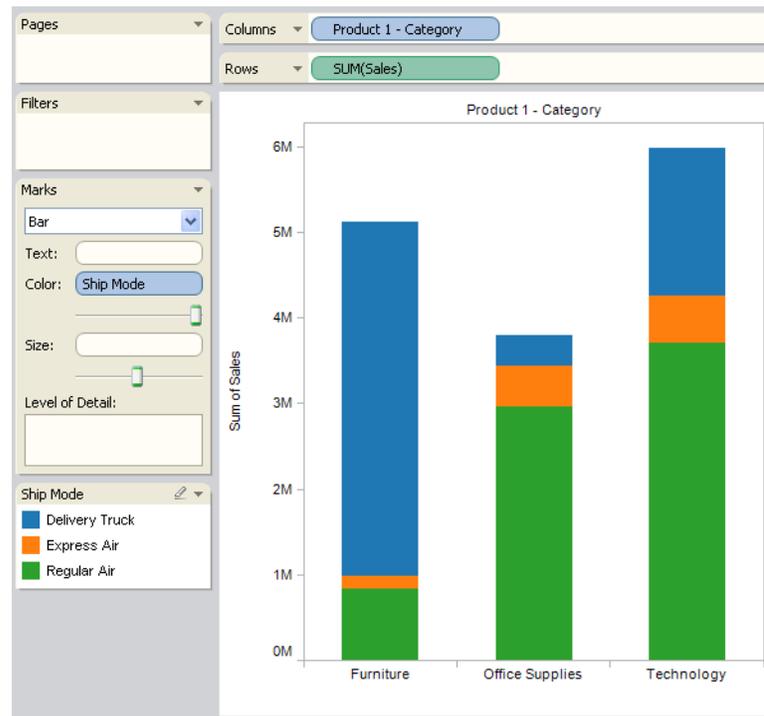
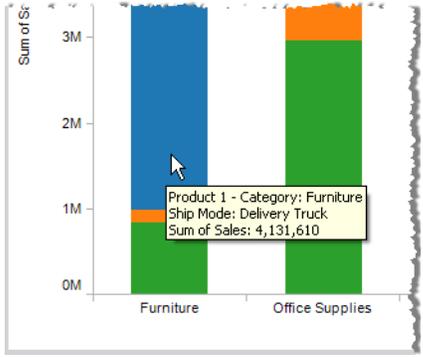


Tableau automatically stacks the marks, which means the bars are drawn cumulatively and “stacked” on each other. The stacking order is reflected in the color legend. Refer to “Stacking Marks” on page 12-27 for more information. The length of each colored bar segment represents the sales for the products that are categorized in one of the three



categories and were shipped using one of the three ship modes. For example, the sum of the sales for furniture shipped by delivery truck is \$4,131,610.



### Bar Chart–Color Encoded Nested Table

This view displays the sales in a nested table as a function of year, region, and customer segment. Additionally, the data are color-encoded by product category. You can create the view with the following procedure.

- 1 Place the **Order Date** and **Region** dimensions on the **Columns** shelf.

The date is automatically aggregated by year, and headers are created for each dimension with labels given by the dimension member names.

- 2 Place the **Sales** measure and the **Customer Segment** dimension 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. Headers are created for the dimension with labels given by the dimension member names.

- 3 Place the **Product 1 - Category** dimension 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.

The view is shown below.



## Bar Chart–Filtered and Color-Encoded

This view displays the sales and profit for all regions for the consumer market for each quarter in the year 2008. Additionally, the data are color-encoded by product category.

You can create the view with the following steps.



- 1 Place the **Order Date** and the **Region** dimensions on the **Columns** shelf.

The date is automatically aggregated by year, and headers are created for each dimension with labels given by the dimension member names. Change the date aggregation to quarter by selecting Quarter from the context menu.

- 2 Place the **Sales** and **Profit** measures on the **Rows** shelf.

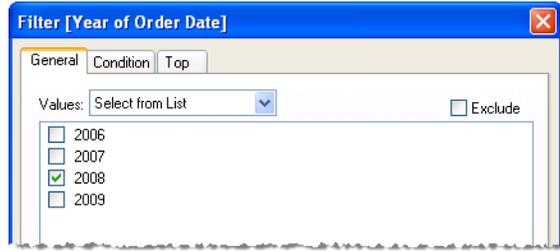
The measures are automatically aggregated as a summation, and axes are created with labels given by the field names.

- 3 Place the **Product 1 - Category** dimension 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.

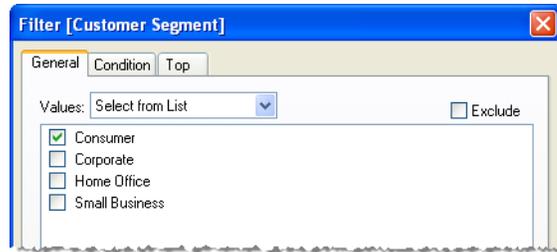
- 4 Place the **Order Date** and **Customer Segment** dimensions on the **Filters** shelf.

Filter **Order Date** to include only the year 2008. You open the filter dialog box by selecting Filter from the field's context menu.

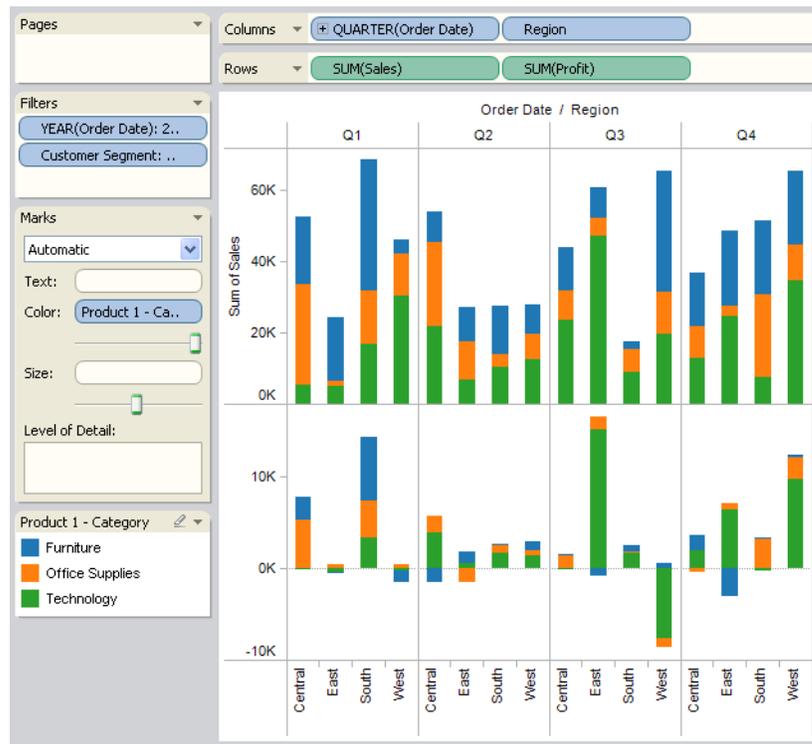




Filter **Customer Segment** to include only the consumer market. You open the filter dialog box by selecting Filter from the field's context menu.



The view is shown below. Note that the mark borders are not displayed so that color encodings are more easily identified.





## **Bar Chart–Color Encoded and Sorted**

This view displays the sum of the profit for each customer and is color-encoded by the average sales. The data are sorted from highest to lowest profit for the top 15 customers. You can create the view with the following procedure.

- 1 Place the **Profit** measure on the **Columns** shelf.

The measure is automatically aggregated as a summation and an axis is created with a label given by the name of the field.

- 2 Place the **Customer Name** dimension on the **Rows** shelf.

Row headers are created with labels given by the dimension member names.

- 3 Place the **Sales** measure on the **Color** shelf.

The measure is automatically aggregated as a summation. Change the aggregation to an average by selecting Average from the field's context menu. The color legend reflects the continuous data range.

- 4 Sort **Customer Names**.

Complete the **Sort** dialog box so that **Sort order** is **Descending** and **Sort by** is **Profit** aggregated as a summation.



**Sort [Customer Name]**

**Sort order**

Ascending

Descending

**Sort by**

Data source order

Alphabetic

Field:

Field: Profit Aggregation: Sum

Manual:

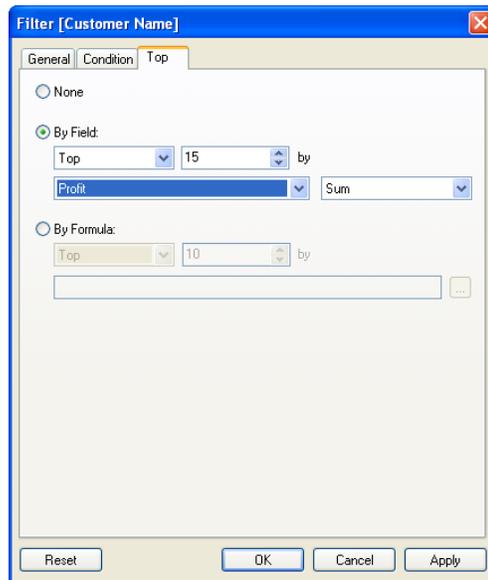
Emily Phan	Up
Deborah Brumfield	
Grant Carroll	Down
Alejandro Grove	
Karen Carlisle	
Liz MacKendrick	
John Lucas	
John Stevenson	
Clivia Kelly	

Clear OK Cancel Apply



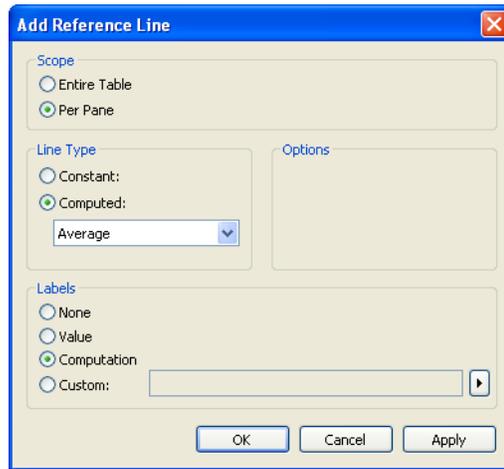
- 5 Filter the data to include only the top 15 customers based on profit.

You can create this filter using the Limit tab in the filter dialog box. Select By Field and complete the definition by making selections on the drop-down lists. To learn more about defining filters refer to “Dragging Fields to the Filters Shelf” on page 16-7.

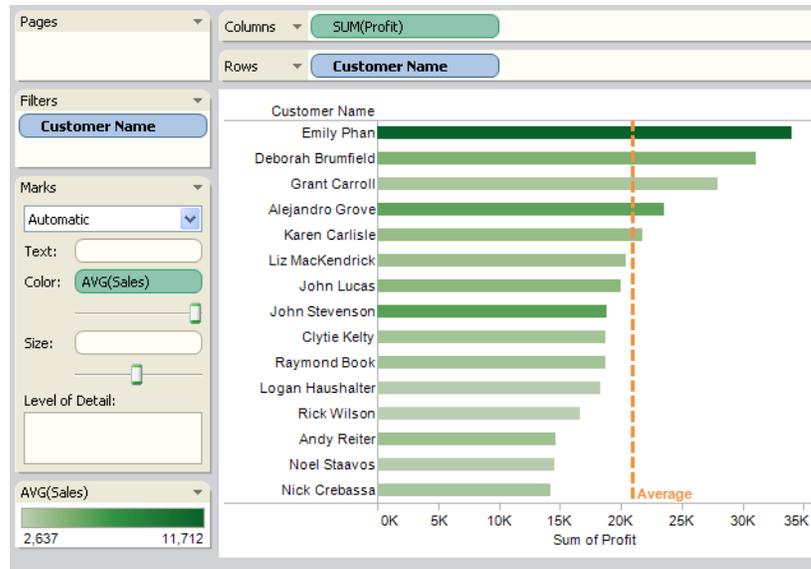


- 6 Add a reference line that indicates the average profit.

Right-click the SUM(Profit) axis and select Add Reference Line. In the subsequent dialog box, select Average for the computed line type and computation for the label.



The view is shown below.



The marks are horizontal because the axis is horizontal. The length of each mark represents the sum of the profit for customer and the color of each mark gives the average sales. For

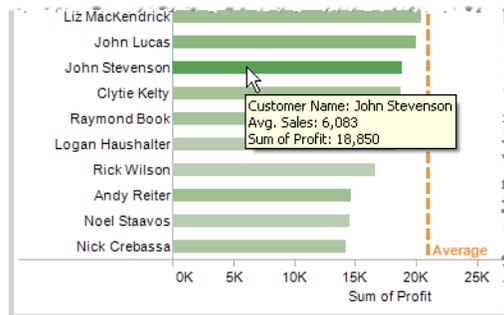


example, the average sales for John Stevenson is \$6,083 and his profit is \$18,850. You can display these numbers with tool tips by placing the mouse cursor over the mark.

---

**Note** You can format reference lines by right-clicking on the line itself and selecting **Format**. Refer to “Reference Lines and Bands” on page 22-1 to learn more.

---



Notice that for similar profit sums, there is a wide range of average sales as indicated by the bottom six rows in the table. For example, John Stevenson and Clytie Kelty have essentially the same profit, but John has a much higher average sales as indicated by the darker color.



## Bar Chart–Filtered, Color Encoded, and Level of Detail

This view displays the sum of the sales for a product category and is color-encoded by year. Additionally, the data display a level of detail given by the customer segment and are filtered to display a sales sum that is greater than \$100,000.

You can create the view with the following steps.

- 1 Place the **Sales** measure on the **Columns** shelf.

The measure is automatically aggregated as a summation and an axis is created with a label given by the name of the field.

- 2 Place the **Product 3 - Name** dimension on the **Rows** shelf.

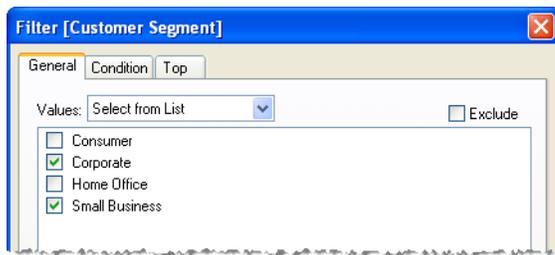
Row headers are created with labels given by the dimension member names.

- 3 Place the **Order Date** dimension 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.

- 4 Place the **Customer Segment** dimension on the **Level of Detail** shelf.

Filter the dimension to include only the Corporate and Small Business customer segments.



- 5 Filter **Product 3 - Name** by sales.

Filter the Product 3 - Name field to include only products whose sales are greater than \$100,000. To filter a dimension by a measure, use the Condition tab in the Filter dialog



box. Select By Field and define the filter using the drop-down lists. Refer to “Adding Conditions to Filters” on page 16-10 to learn more about defining filters.

Filter [Product 3 - Name]

General Condition Top

None

By Field:      Aggregation:

Sales      Sum

>      100,000

Range of Values:

Min:      Load

Max:     

By Formula:

Reset      OK      Cancel      Apply



The view is shown below.



The total length of a given mark represents the sum of the sales for the associated product. Additionally, each mark is color-encoded by year. The length of a colored segment represents the contribution to the sales for the associated year. Finally, each colored segment can be separated into two levels of detail: corporate and small business.

**Note** By default, mark borders are turned off. However, in this case it makes more sense to have them turned on so that the customer segment detail can be identified. You can turn off mark borders by selecting the **Format > Marks** menu item.

---

## Text Table

Text tables (also called cross-tabs or Pivot Tables) provide an easy way to display the numbers associated with categorical data. In Tableau, you typically create a text table by placing a dimension as the innermost field on both the **Rows** and the **Columns** shelf. Given this description, a typical text table will use only dimensions for the rows and columns of a table. You then complete the view by placing a field (usually a measure) on the **Text** shelf.

A text table uses the text mark type. Tableau can automatically select this mark type if the view is constructed using only dimensions as shown below, and the **Mark** menu is set to Automatic.



Alternatively, you can manually select Text from the **Mark** menu for any data view. Refer to “Text Mark” on page 12-12 for information about this mark type.

The following examples show you how to use text tables to explore your data:

- Text Table–Basic – Create a text table using several dimensions, and complete the table by displaying an aggregated measure.  
The example is extended to display percentages and totals for the table.
- Text Table–Filter and Color Encode – Create a text table using several dimensions, complete the table by displaying an aggregated measure, filter the data using a measure, and color-encode the data using a dimension.
- Duplicate as Crosstab – Duplicate an existing data view as a cross-tab.

### Text Table–Basic

This view displays the profit as a function of product category, year, and region as a text table.

You can create the view with the following procedure.



1 Place the **Order Date** dimension on the **Columns** shelf.

The date is automatically aggregated by year, and column headers are created with labels given by the dimension member names.

2 Place the **Product 2 - Sub-Category** dimension on the **Rows** shelf.

Row headers are created with labels given by the dimension member names.

3 Place the **Sales** dimension on the **Text** shelf.

The measure is automatically aggregated as a summation.

The view is shown below.

Product 2 - Sub-Category	Order Date			
	2006	2007	2008	2009
Appliances	170,657	216,660	168,657	191,359
Binders and Binder Accesso..	360,210	178,639	185,123	300,550
Bookcases	253,942	263,566	140,664	163,810
Chairs & Chairmats	518,011	380,704	451,249	394,181
Computer Peripherals	215,565	175,327	191,042	214,620
Copiers and Fax	273,287	339,713	280,821	236,541
Envelopes	31,376	51,093	39,308	54,521
Labels	10,364	7,393	12,937	8,350
Office Furnishings	194,770	162,197	164,979	143,966
Office Machines	732,505	446,780	428,749	563,308
Paper	118,093	119,516	100,490	111,397
Pens & Art Supplies	42,340	46,728	42,671	35,786
Rubber Bands	3,142	3,754	4,375	3,969
Scissors, Rulers and Trimme..	16,653	6,877	10,645	46,821
Storage & Organization	297,775	289,979	243,495	268,525
Tables	478,981	429,657	500,322	478,255
Telephones and Communica..	472,252	442,048	471,010	504,005



Tableau automatically selects the mark type as text and the mark color as black. Each cell in the table displays the sum of the sales.

The data show that office machines account for the largest sales almost all of the four years. Suppose that instead of displaying the raw numbers, you want to display:

- The percentage of each cell with respect to the entire table.
- The grand totals for the rows and columns.

When your data are displayed in a text table, it is easy to do. To display percentages, select one of the percentage options on the **Analysis > Percentage of...** menu, and to show grand totals select either **Table > Grand Totals for Columns** or **Table > Grand Totals for Rows**.

The view is shown below.

The screenshot shows the Tableau interface with a text table. The Columns shelf contains 'YEAR(Order Date)' and the Rows shelf contains 'Product 2 - Sub-Cate...'. The Marks card is set to 'Automatic' with 'SUM(Sales)' as the text and a black color. The Level of Detail is set to 'None'. The table displays sales data for various product categories from 2006 to 2009, with percentages and grand totals.

Product 2 - Sub-Category	Order Date				Grand Total
	2006	2007	2008	2009	
Appliances	1.14%	1.45%	1.13%	1.28%	5.01%
Binders and Binder Accesso..	2.42%	1.20%	1.24%	2.02%	6.87%
Bookcases	1.70%	1.77%	0.94%	1.10%	5.51%
Chairs & Chairmats	3.47%	2.55%	3.03%	2.64%	11.70%
Computer Peripherals	1.45%	1.18%	1.28%	1.44%	5.34%
Copiers and Fax	1.83%	2.28%	1.88%	1.59%	7.58%
Envelopes	0.21%	0.34%	0.26%	0.37%	1.18%
Labels	0.07%	0.05%	0.09%	0.06%	0.26%
Office Furnishings	1.31%	1.09%	1.11%	0.97%	4.47%
Office Machines	4.91%	3.00%	2.88%	3.78%	14.57%
Paper	0.79%	0.80%	0.67%	0.75%	3.02%
Pens & Art Supplies	0.28%	0.31%	0.29%	0.24%	1.12%
Rubber Bands	0.02%	0.03%	0.03%	0.03%	0.10%
Scissors, Rulers and Trimme..	0.11%	0.05%	0.07%	0.31%	0.54%
Storage & Organization	2.00%	1.95%	1.63%	1.80%	7.38%
Tables	3.21%	2.88%	3.36%	3.21%	12.66%
Telephones and Communica..	3.17%	2.97%	3.16%	3.38%	12.67%
<b>Grand Total</b>	<b>28.11%</b>	<b>23.89%</b>	<b>23.05%</b>	<b>24.95%</b>	<b>100.00%</b>



The interpretation of the data is as follows:

- Each data cell reflects the percentage of the entire table. For example, computer peripherals in 2006 accounted for 1.45% of the total sales across all years and all products.
- The grand totals for rows is the bottom-most row in the table. Each cell reflects the percent contribution for a given year. For example, 2007 contributed 23.89% of the total sales across all years and all products.
- The grand totals for columns is the right-most column in the table. Each cell reflects the percent contribution for a given product. For example, envelopes accounted for 1.18% of the total sales across all years and all products.

Refer to “Percentages” on page 21-87 for more information about displaying percentages.



## Text Table–Filter and Color Encode

This view is a text table that displays the profit as a function of product category, year, and region. Additionally, the data are color encoded by product category and filtered to exclude negative profits.

You can create the view with the following procedure.

- 1 Place the **Order Date** dimension on the **Columns** shelf.

The date is automatically aggregated by year, and headers are created with labels given by the dimension member names.

- 2 Place the **Product 2 - Sub-Category** dimension on the **Rows** shelf.

Row headers are created with labels given by the dimension member names.

- 3 Place the **Sales** measure on the **Text** shelf.

The measure is automatically aggregated as a summation.

- 4 Place the **Product 1 - Category** dimension 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.

- 5 Filter by **Profit**.

Place the **Profit** measure on the **Filters** shelf and filter the data so that every row with negative profits is excluded from the view.



**Filter [Sum of Profit]**

Range of values | **At Least** | At Most | Special

At Least

0 | 123,814.08

-23,947.08 | 123,814.08

Show: Only Relevant Values |  Include null values

Reset | OK | Cancel | Apply



The view is shown below.

The screenshot shows a Tableau view with the following configuration:

- Columns:** + YEAR(Order Date)
- Rows:** Product 2 - Sub-Cate...
- Filters:** SUM(Profit)
- Marks:** Automatic
- Text:** SUM(Sales)
- Color:** Product 1 - Ca...
- Size:** (slider)
- Level of Detail:** (empty)
- Product 1 - Category:** Furniture (blue), Office Supplies (orange), Technology (green)

Product 2 - Sub-Category	Order Date			
	2006	2007	2008	2009
Appliances	170,657	216,660	168,657	191,359
Binders and Binder Accessories	360,210	178,639	185,123	300,550
Chairs & Chairmats	518,011	380,704	451,249	394,181
Computer Peripherals	215,665	175,327	191,042	214,620
Copiers and Fax	273,287	339,713	280,821	236,541
Envelopes	31,376	51,093	39,308	54,521
Labels	10,384	7,393	12,937	8,350
Office Furnishings	194,770	162,197	164,979	143,966
Office Machines	732,505	446,780	428,749	563,308
Paper	118,093	119,516	100,490	111,397
Pens & Art Supplies	42,340	46,728	42,671	35,786
Rubber Bands				3,969
Storage & Organization	297,775	289,979		
Telephones and Communication	472,252	442,048	471,010	504,005

The data show the sales for all products excluding rows with negative profits. Note that Rubber Bands only had rows with positive profits in 2009, the other negative values are excluded from the view.

## Duplicate as Crosstab

The view shown below was generated as a crosstab from the bar chart example “Bar Chart–Filtered and Color-Encoded” on page 33-10. Additionally, the rows and columns were swapped.

You can duplicate any view as a crosstab by selecting **Edit > Duplicate as Crosstab**. The result will contain only dimensions in the **Rows** and **Columns** shelves. If measures are part of the original view, the measure names will be displayed in the crosstab using the **Measure Names** field. Views that are diagggregated cannot be duplicated as a crosstab.

Columns		Product 1 - Category					
Rows		Furniture		Office Supplies		Technology	
Quarter of Order Date	Region	Sum of Profit	Sum of Sales	Sum of Profit	Sum of Sales	Sum of Profit	Sum of Sales
Q1	Central	-1,228	10,322	978	10,574	8,780	33,540
	East	-246	22,020	762	14,203	264	5,173
	South	-122	268	-174	789	310	4,585
	West	7,997	40,836	-603	7,005	658	6,453
Q2	Central	3,065	22,927	3,677	18,452	1,703	11,339
	East	-845	5,009	1,455	9,798	3,181	23,372
	South	548	14,073	1,019	3,963	-534	1,937
	West	3,963	34,875	5,810	25,174	6,531	27,678
Q3	Central	-725	16,888	2,001	13,628	843	5,299
	East	67	6,898	-241	704	2,479	14,780
	South	-948	12,518	1,588	14,172	339	2,326
	West	-653	9,749	457	6,696	-9,269	30,462
Q4	Central	-1,626	10,729	4,872	14,262	-385	5,487
	East	1,575	23,649	808	8,740	-5	7,030
	South	-1,614	20,038	4,837	20,851	14,934	45,137
	West	-1,712	17,097	706	12,427	1,638	22,697

Notice that the fields **Measure Names** and **Measure Values** appear on the **Columns** shelf and the **Text** shelf, respectively. These fields were automatically included by Tableau so that the headers and text values are included for the **Profit** and **Sales** measures.



## Line Chart

Line charts connect individual data points in a data view. They provide a simple way to visualize a sequence of values, and are especially useful when you want to see trends over time.

In Tableau, you create a line chart by selecting the line mark type. This mark type is automatically selected if 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 continuous quantity. You can also plot the measures against a discrete or continuous date or time field. Refer to “Line Mark” on page 12-22 for more information about the line mark type.

The following examples show you how to use line charts to explore your data:

- **Line Chart–Basic** – Create a line chart by plotting a dimension and two aggregated measures.  
The example is extended by plotting multiple measures in the same pane and identifying the lines using mark labels.
- **Filter and Color Encode a Nested Table** – Create a nested line chart by plotting multiple dimension and an aggregated measure, filter the data by year, and color-encode the data by a dimension.
- **Line Chart–Filter and Color-Encode Multiple Measures** – Plot multiple measures in a single pane as a function of month, filter the data using several dimensions, and color-encode the data by the measure names.  
The example is extended by converting a date field to a continuous quantity and then filtering the date values.
- **Line Chart–Filter and Path Encode** – Plot disaggregated measures, filter the data using several dimensions, and path-encode the data by date.



## Line Chart–Basic

This view displays the sum of the sales and the sum of the profit for all years. You can create the view with the following procedure.

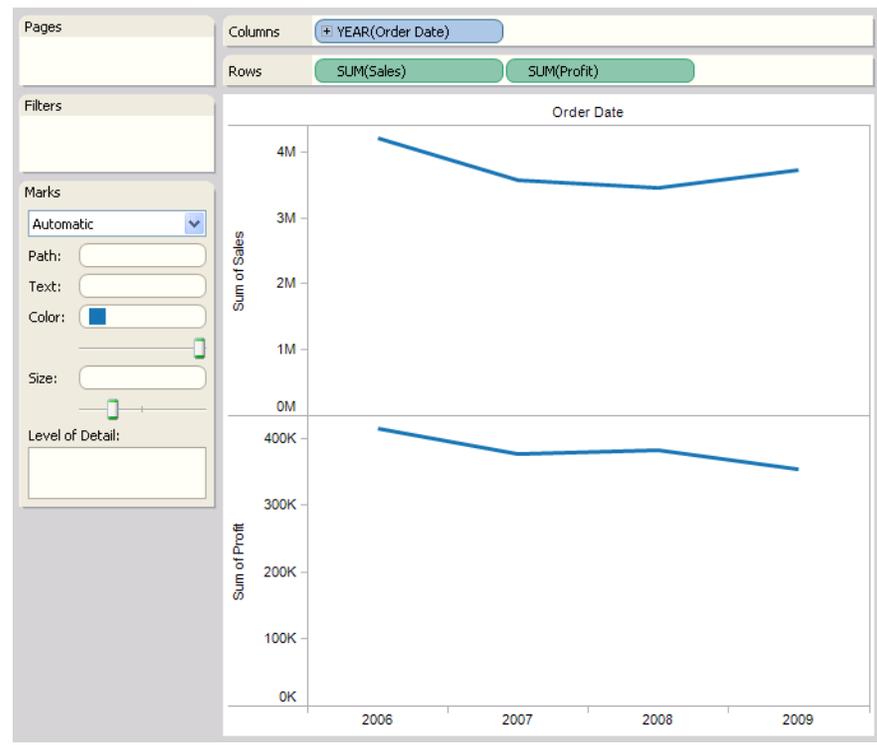
- 1 Place the **Order Date** dimension on the **Columns** shelf.

The date is automatically aggregated by year, and headers are created with labels given by the dimension member names.

- 2 Place the **Sales** and **Profit** measures on the **Rows** shelf.

The measures are automatically aggregated as summations, and axes are created with labels given by the field names.

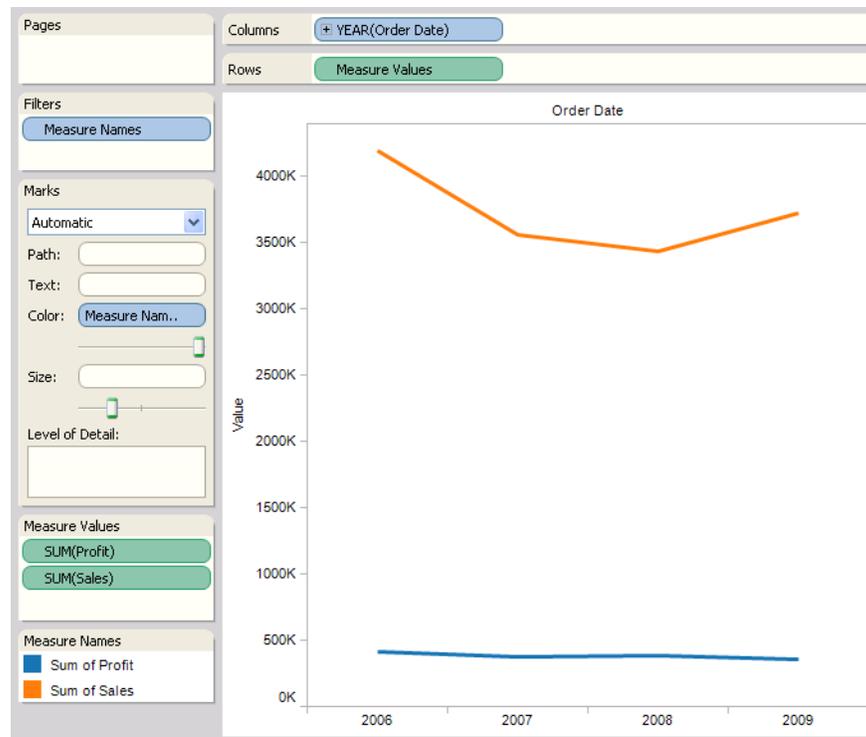
The view is shown below.





You might find that the view is more useful if you display both measures in the same pane. This way, the measures use the same scale and are easier to compare. To do this, drag one of the measures off the **Rows** shelf and then drag that measure from the Data window to the axis.

The view is shown below. Tableau automatically placed the **Measure Names** dimension on the **Color** and **Filters** shelf, and placed the **Measure Values** measure on the **Rows** shelf.





## Filter and Color Encode a Nested Table

This view displays the sum of profit in 2008 for all regions and customer segments as a nested table using the line mark type. Additionally, the data are color-encoded by product category.

You can create the view with the following procedure.

- 1 Place the **Customer Segment** and **Order Date** dimensions on the **Columns** shelf.

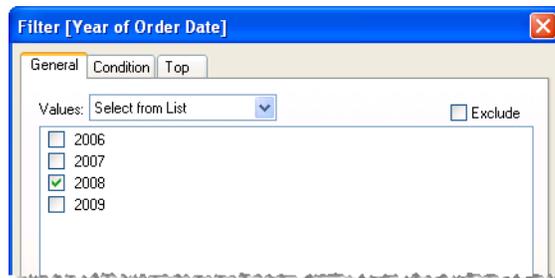
The date is automatically aggregated by year, and headers are created for each dimension with labels given by the dimension member names. Display the data by quarter by selecting Quarter from the field's context menu.

- 2 Place the **Profit** measure and the **Region** dimension 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. Headers are created for the dimension with labels given by the dimension member names.

- 3 Filter **Order Date**.

You open the **Filter** dialog box by selecting **Filter** from the field's context menu. As shown below, include only 2008.

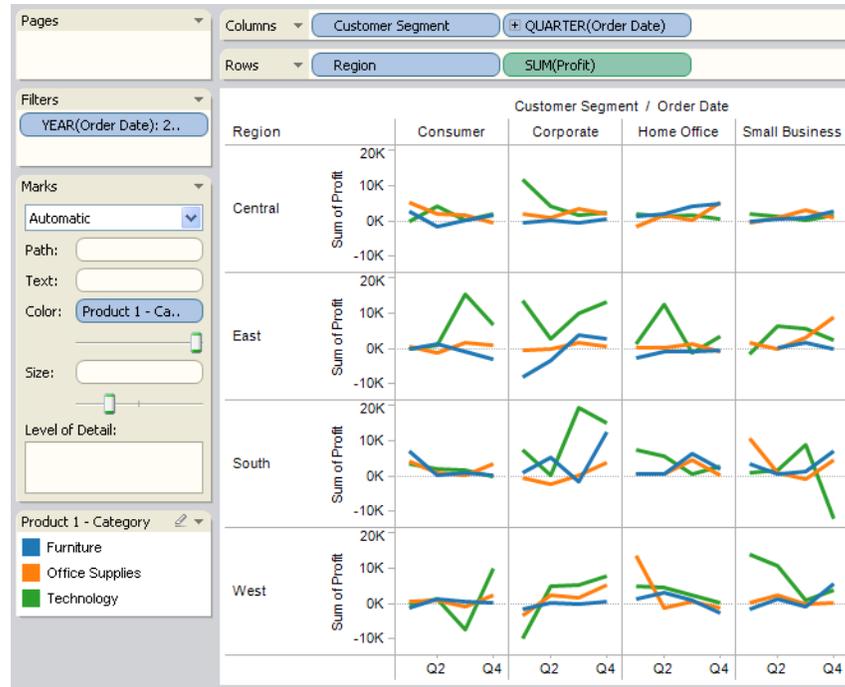


- 4 Place the **Product 1 - Category** dimension on the **Color** shelf.

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.



The view is shown below.





## Line Chart–Filter and Color-Encode Multiple Measures

It displays multiple measures in a single pane as a function of month. The data are filtered for a specific year, region, customer segment, and product category. Additionally, the data are color-encoded by measure names.

You can create the view with the following procedure.

- 1 Place the **Order Date** dimension on the **Columns** shelf.

The date is automatically aggregated by year, and headers are created with labels given by the dimension member names. Change the aggregation to month by selecting **Month** from the field's context menu.

- 2 Place the **Measure Values** measure on the **Rows** shelf.

By default, all measure are selected and an axis is created with the label **Value**, indicating multiple values are displayed.

- 3 Remove all measure except **Sales** and **Profit**.

Drag all measures except the **Sales** and **Profit** measures off the **Measure Names/Values** shelf.



Note that the **Measure Names** dimension was automatically placed on the **Filters** shelf after removing the measures.

- 4 Place the **Measure Names** dimension 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.

- 5 Place the **Order Date**, **Customer Segment**, **Region**, and **Product 1 - Category** dimensions on the **Filters** shelf.

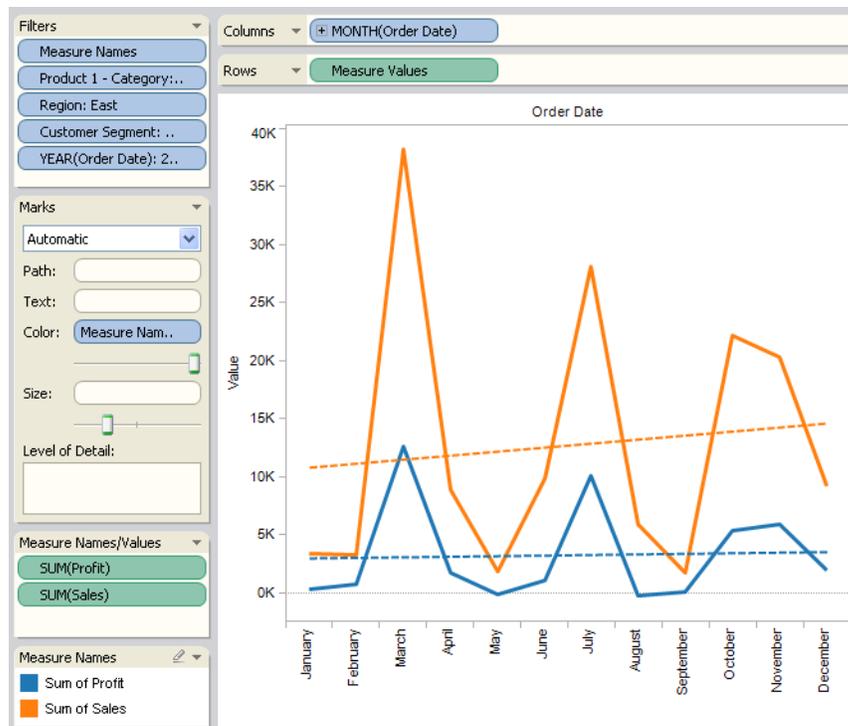
- Filter **Order Date** to include only the year 2008.

- Filter **Customer Segment** to include only the Corporate market.
- Filter **Region** to include only the East region.
- Filter **Product 1 - Category** to include only Technology

6 Add trend lines to the view.

Right-click on the view and select **Trend Lines** from the context menu. Two trend lines are shown, one for profit and the other for sales. You can edit the trend lines right-clicking and selecting **Edit Trend Lines**. Refer to Chapter 25, “Trend Lines & Statistics” to learn more.

The view is shown below.





Notice that the Month labels along the bottom of the view are oriented vertically. You can easily toggle between vertical and horizontal alignment for headers along the bottom of the view by pressing **Ctrl + L**.

## **Line Chart–Filter and Path Encode**

This view displays the profit and sales in a single table pane. The data are filtered for a specific year, region, customer segment, and product category. Additionally, the data are path-encoded by order date.

You can create the view with the following procedure.

- 1 Place the **Profit** measure on the **Columns** shelf.

The measure is automatically aggregated as a summation and an axis is created with a label given by the name of the field.

- 2 Place the **Sales** measure on the **Rows** shelf.

The measure is automatically aggregated as a summation and an axis is created with a label given by the name of the field.

- 3 Disaggregate the measures by selecting **Analysis > Aggregate Measures**.

- 4 Place **Order Date** (twice), **Customer Segment**, **Region**, and **Product Category 1** on the **Filters** shelf.

- Filter the first **Order Date** field to include only the year 2008.
- For the second **Order Date** field, select Quarter from the field's context menu and filter the field to include only Q4.
- Filter **Customer Segment** to include only the Corporate market.
- Filter **Region** to include only the East region.



- Filter **Product 1 - Category** to include only Furniture

Note that there are five data points that pass through all the filters.

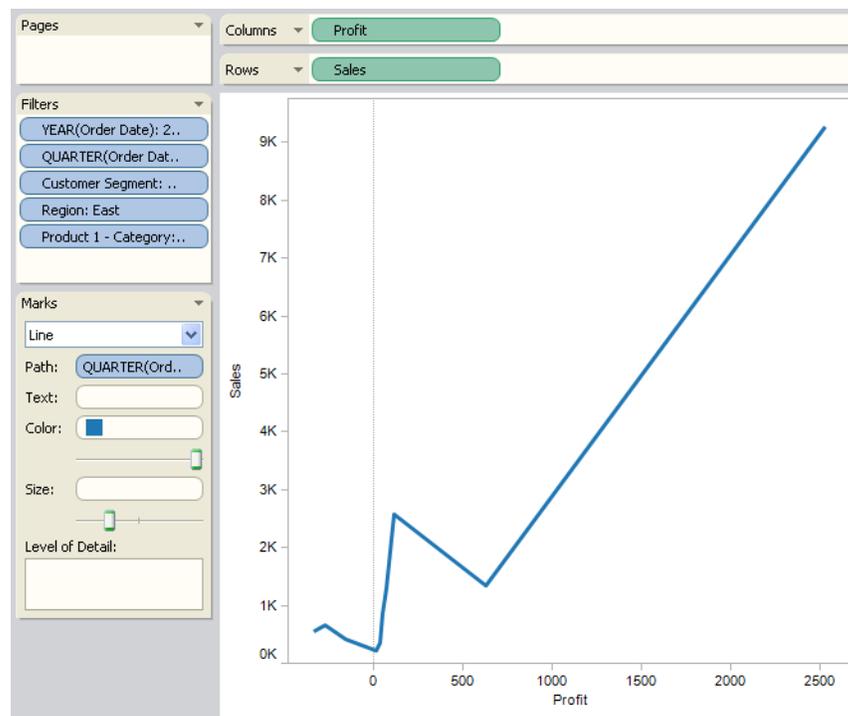
5 Select Line from the **Mark** menu.

Note that the data points are connected from lowest profit value to highest value, in order.

6 Place the **Order Date** dimension on the **Path** shelf and select **Quarter** as the aggregation.

The data points are now connected in the chronological order.

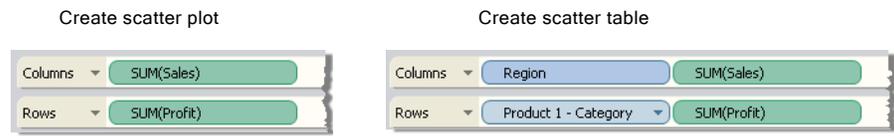
The view is shown below.





## Scatter Plot

Scatter plots provide an easy way to visualize relationships between numerical variables. In Tableau, you create a scatter plot by placing at least one measure on both the **Columns** and the **Rows** shelf. If these shelves contain both dimensions and measures, Tableau automatically places the measures as the innermost fields.



A scatter plot can use several mark types. By default, Tableau uses the shape mark type. However, depending on your data, you might want to use another mark such as a circle or a square. Refer to “Mark Types” on page 12-11 for more information.

By default, measures are aggregated as a summation when placed on a shelf. Therefore, Tableau displays only one data point in each pane of a data table. You can view more data by disaggregating the measures or by placing one or more dimensions on another shelf such as **Level of Detail**, **Shape**, or **Color**.

The following examples show you how to use scatter plots to explore your data:

- Scatter Plot–Basic – Plot two disaggregated measures using the default mark type.
- Scatter Plot–Filter and Color Encode – Plot four disaggregated measures using the default mark type, filter and color-encode the data, and look for relationships between the four measures by brushing the data.
- Scatter Plot–Filter and Level of Detail – Plot two aggregated measures using the circle mark type, filter the data, and display additional levels of detail.
- Scatter Plot–Filter, Level of Detail, and Color Encode – Plot two aggregated measures using the circle mark type, filter the data, display additional levels of detail, and color-encode the data.
- Scatter Plot–Filter, Level of Detail, Color Encode, and Shape Encode – Plot two aggregated measures using the shape mark type, filter the data, display additional levels of detail, color-encode the data, and shape-encode the data.



---

## Scatter Plot–Basic

This view displays the disaggregated profit and sales in a single pane as a scatter plot. You can create the view with the following procedure.

- 1 Place the **Profit** measure on the **Columns** shelf.

The measure is automatically aggregated as a summation and an axis is created with a label given by the name of the field.

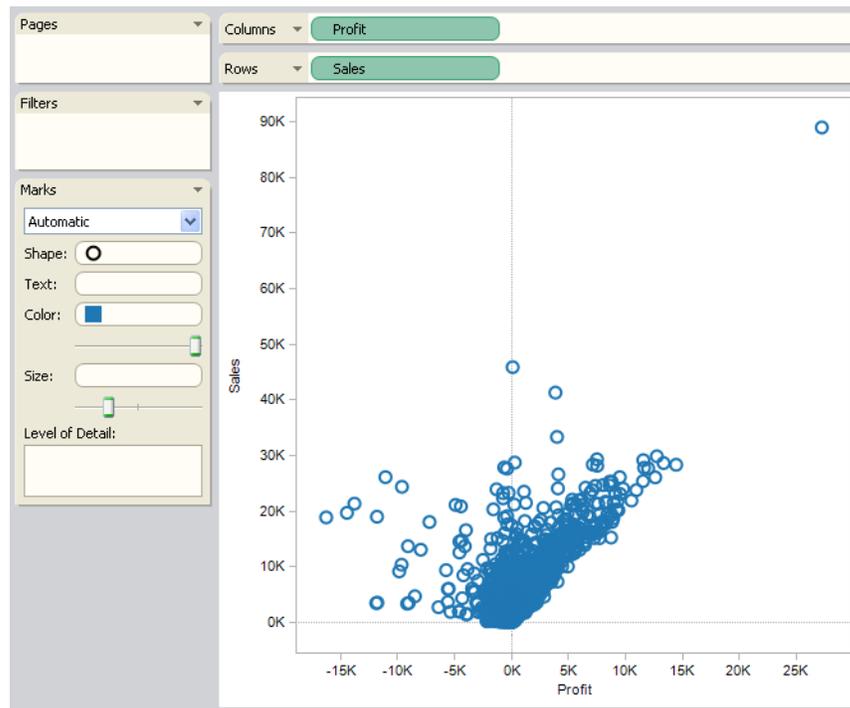
- 2 Place the **Sales** measure on the **Rows** shelf.

The measure is automatically aggregated as a summation and an axis is created with a label given by the name of the field.

- 3 Disaggregate the measures.

Select **Aggregate Measures** from the **Analysis** menu. This instructs Tableau to display all the data from your source as a scatter plot.

The view is shown below. Note that Tableau automatically selects the mark as **Shape**, the mark shape as an open circle, and the mark color as blue.





## Scatter Plot–Filter and Color Encode

This view displays the disaggregated profit, sales, ship time, and discount as a scatter table. Additionally, brushing is used to visualize the relationships among the data. You can create the view with the following procedure.

- 1 Place the **Profit** and **Time to Ship** measures on the **Columns** shelf.

The Profit measure is automatically aggregated as a summation while Time to Ship is aggregated as a maximum.

- 2 Place the **Sales** and **Discount** measures on the **Rows** shelf.

The measures are automatically aggregated as a summation and axes are created with labels given by the field names.

- 3 Disaggregate the measures.

Select the **Analysis > Aggregate Measures** menu item. This instructs Tableau to display all the data from your source as a scatter plot.

- 4 Place the **Product 1 - Category**, **Ship Mode**, and **Container** dimensions on the **Filters** shelf.

- Filter **Product 1 - Category** to include only furniture.
- Filter **Ship Mode** to include delivery truck and regular air.
- Filter **Container** to include large box.

- 5 Place the **Region** dimension on the **Color** shelf.

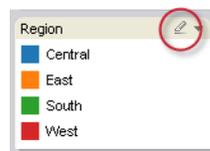
Placing a dimension on the **Color** shelf separates the marks according to the members in the dimension, and assigns a unique color to each member. The color legend displays each member name and its associated color.

The view is shown below.



You can easily highlight a specific set of data using color legend highlighting. Legend highlighting is a powerful analytical mode that allows you to focus on selected items in the color legend.

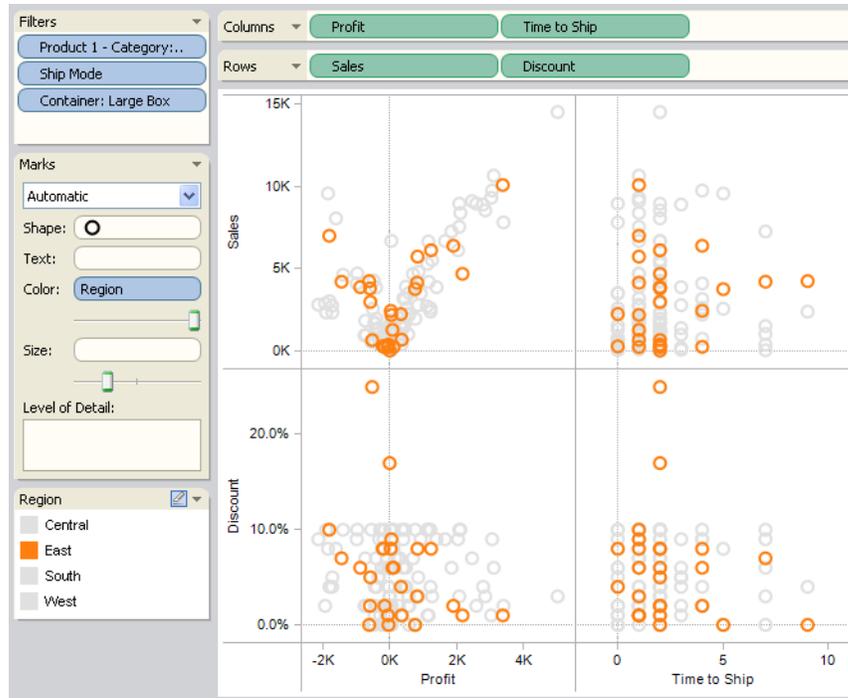
Turn on color legend highlighting by clicking the Highlight button in the title of the color legend.



Then select an item in the color legend. Marks that are associated with the selected member are colored while all other marks are gray.



The view below highlights items that are out of stock. You can easily see that all of orders with long delivery times are for products that are out of stock.





## Scatter Plot–Filter and Level of Detail

This view displays the aggregated profit and sales as a scatter plot using a circle as the mark type. The data are displayed with a level of detail given by the year, region, customer segment, and product category.

You can create the view with the following procedure.

- 1 Place the **Profit** measure on the **Columns** shelf.

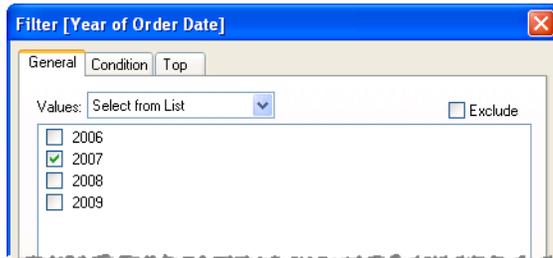
The measure is automatically aggregated as a summation and an axis is created with labels given by the field name.

- 2 Place the **Sales** measure on the **Rows** shelf.

The measure is automatically aggregated as a summation and an axis is created with labels given by the field name.

- 3 Place the **Order Date** dimension on the **Filters** shelf.

Filter **Order Date** to include only the year 2007. You open the **Filter** dialog box by selecting Filter from the field's context menu.



- 4 Place the **Region**, **Customer Segment**, and **Product 1 - Category** dimensions on the **Level of Detail** shelf.

Placing a dimension on the **Level of Detail** shelf instructs Tableau to display the data according to the level of detail given by the dimension members. For example, the

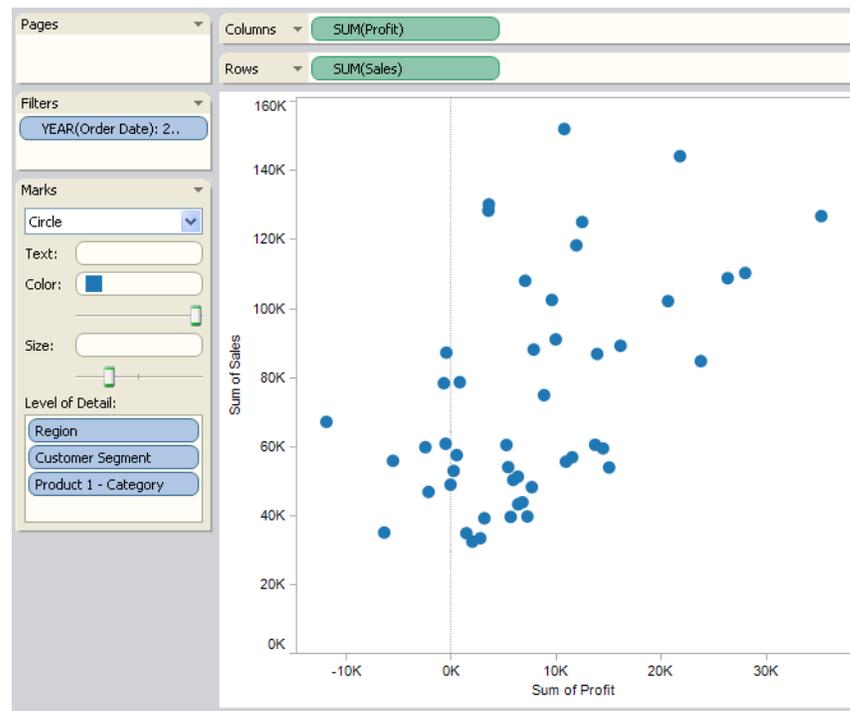


**Region** dimension separates the data points into four levels of detail: East, West, South and Central.

5 Select Circle from the **Mark** menu.

Using a circle as the mark type allows you to better distinguish between closely spaced data points.

The view is shown below.



**Note** By default, mark borders are turned off. However, in this case it makes more sense to have them turned on to distinguish between closely spaced data points. You can turn on mark borders by selecting the **Format > Show Mark Borders** menu item.



## Scatter Plot–Filter, Level of Detail, and Color Encode

This view displays the aggregated profit and sales as a scatter plot using a circle as the mark type. The data are displayed with a level of detail given by the year, region, customer segment, and product category. Additionally, the data are color-encoded by product category.

You can create the view with the following procedure.

- 1 Place the **Profit** measure on the **Columns** shelf.

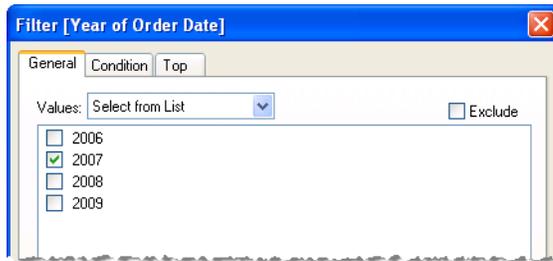
The measure is automatically aggregated as a summation and an axis is created with the label given by the field name.

- 2 Place the **Sales** measure on the **Rows** shelf.

The measure is automatically aggregated as a summation and an axis is created with the label given by the field name.

- 3 Place the **Order Date** dimension on the **Filters** shelf.

Filter **Order Date** to include only the year 2007. You open the **Filter** dialog box by selecting Filter from the field's context menu.



- 4 Place the **Region**, **Customer Segment**, and **Product 1 - Category** dimensions on the **Level of Detail** shelf.

Placing a dimension on the **Level of Detail** shelf instructs Tableau to display the data according to the level of detail given by the dimension members. For example, the



**Region** dimension separates the data points into three levels of detail: East, West, and Central.

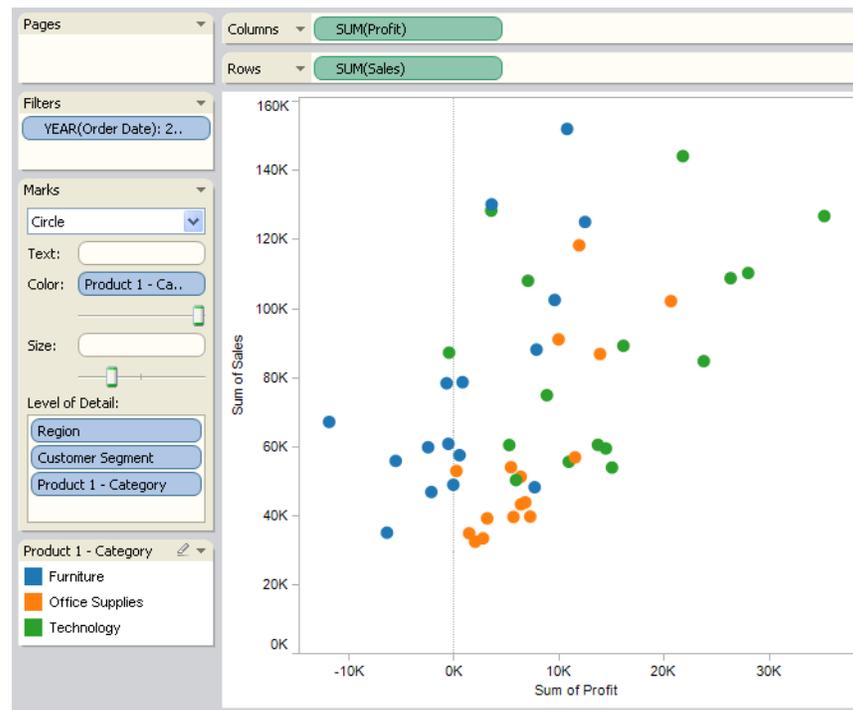
- 5 Select Circle from the **Mark** menu.

Using a circle as the mark type allows you to better distinguish between closely spaced data points.

- 6 Place the **Product 1 - Category** dimension 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.

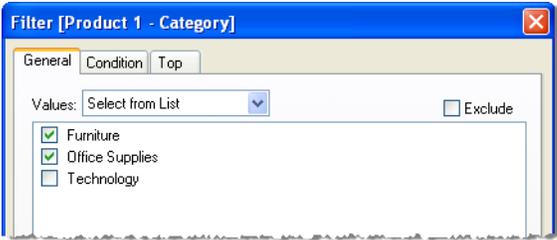
The view is shown below.



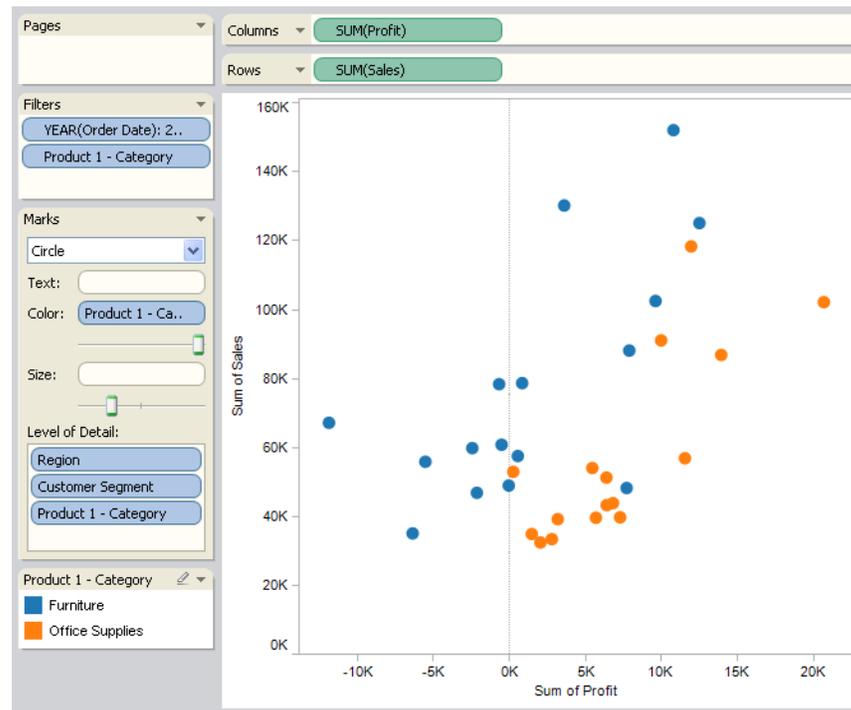


As you can see, the data points are clustered according to product type. For example, the profit for technology (green) is consistently positive, has the largest values, and appears to have a strong linear relationship to the sales. Office supplies appear to have mostly positive profits, while furniture appears to have mostly negative profits.

To gain additional insight into the data, filter **Product 1 - Category** to include only furniture and office supplies. Open the **Filter** dialog box by selecting Filter from the field's context menu.



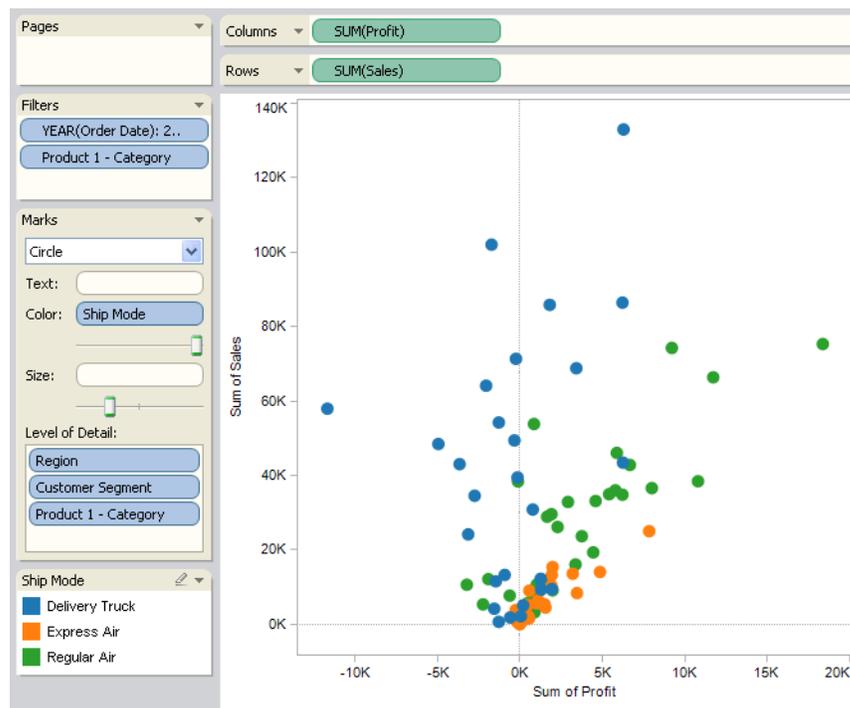
The view is shown below.



Now let's look at an additional level of detail by placing the **Ship Mode** dimension on the **Color** shelf.

Placing this 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.

The view is shown below.



It appears that the items with large positive profits are all shipped by regular air, while the items with large negative profits are all shipped by delivery truck.



## Scatter Plot–Filter, Level of Detail, Color Encode, and Shape Encode

This view displays the aggregated profit and time to ship as a scatter plot using a shape as the mark type. The data are displayed with a level of detail given by the year, region, customer segment, and product category. Additionally, the data are color-encoded by ship mode.

You can create the view with the following procedure.

- 1 Place the **Profit** measure on the **Columns** shelf.

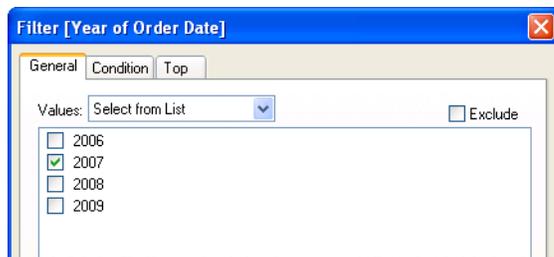
The measure is automatically aggregated as a summation and an axis is created with the label given by the field name.

- 2 Place the **Time to Ship** measure on the **Rows** shelf.

The measure is automatically aggregated as a maximum and an axis is created with the label given by the field name. Change the aggregation to an average by selecting Average from the field's context menu.

- 3 Place the **Order Date** dimension on the **Filters** shelf.

Filter **Order Date** to include only the year 2007. Open the **Filter** dialog box by selecting Filter from the field's context menu.



- 4 Place the **Region**, **Customer Segment**, and **Product 1 - Category** dimensions on the **Level of Detail** shelf.

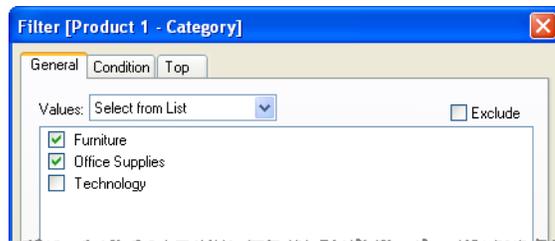
Placing a dimension on the **Level of Detail** shelf instructs Tableau to display the data according to the level of detail given by the dimension members. For example, the



**Region** dimension separates the data points into three levels of detail: East, West, and Central.

**5 Filter Product Category 1.**

You open the **Filter** dialog box by selecting Filter from the field's context menu. Include only furniture and office supplies as shown below.



**6** Select Shape from the **Mark** menu.

**7** Place the **Ship Mode** dimension on the **Shape** shelf.

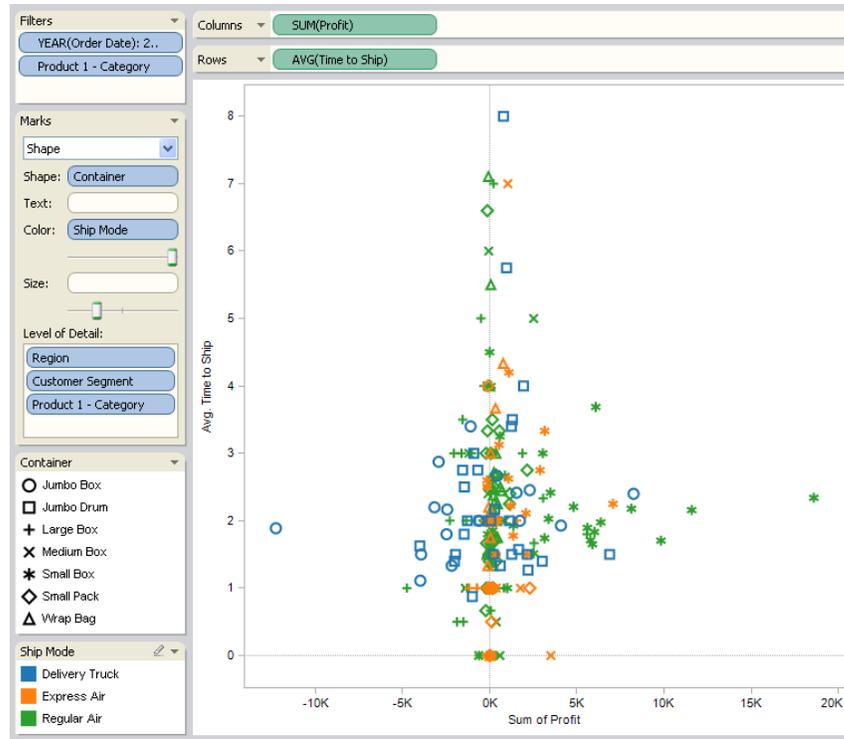
The **Shape** shelf separates the marks according to the members in the dimension, and assigns a unique shape to each member. The shape legend displays each member name and its associated shape.

**8** Place the **Container** dimension 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.



The view is shown below.



Here are a few things you can learn from the data:

- The large profit values are associated with items shipped in a small box and delivered via regular air. Additionally, these items have a consistent and relatively short average time to ship.
- There are several outliers with very long average time to ship.

## Heat Map

Heat maps are a great way to compare categorical data using color. They are typically constructed as a table using colored squares to represent the data, and a continuous range of colors. Heat maps allow you to see variations in the data via variations in color.

In Tableau, you create a heat map by placing one or more dimensions on the **Columns** shelf and one or more dimensions on the **Rows** shelf. You then select the mark type as Square and place a measure of interest on the **Color** shelf.

Of course, you are not limited to this type of heat map. You can build any view that is meaningful to you. For example, you can color-encode the data using a dimension, as well as size-encode and shape-encode the data. Some of these views certainly stretch the definition of a heat map. The common thread between them is the display of color-encoded categorical data.

The following examples show you how to use heat maps to explore your data:

- Heat Map–Basic – A basic heat map consisting of a table of categorical data, color-encoded by a measure, and displayed using a square mark type.
- Heat Map–Size Encoded – A basic heat map that is modified by size-encoding the data. Two examples are presented: size-encoding using a measure and size-encoding using a dimension.
- Heat Map–Shape and Size Encoded – A basic heat map that is modified by shape-encoding and color-encoding the data using a dimension, and size-encoding the data using a measure. You would probably not consider this view to be a heat map because it uses the shape mark type. However, it does illustrate the ease with which you can transform a basic heat map into a different view of your data.



## Heat Map–Basic

This view shows a basic heat map. It is constructed as a table of categorical data and is color-encoded by a measure. The view shows the profit across all markets, regions, and products.

You can create the view with the following procedure.

- 1 Place the **Customer Segment** dimension on the **Columns** shelf.

Headers are created with labels given by the dimension member names.

- 2 Place the **Region** and **Product 1 - Category** dimensions on the **Rows** shelf.

Headers are created with labels given by the dimension member names. You have now created a nested table of categorical data.

- 3 Place the **Profit** measure on the **Color** shelf.

The measure is automatically aggregated as a summation. The color legend reflects the continuous data range.

- 4 Select the mark type to be Square.

- 5 Optimize the view format.

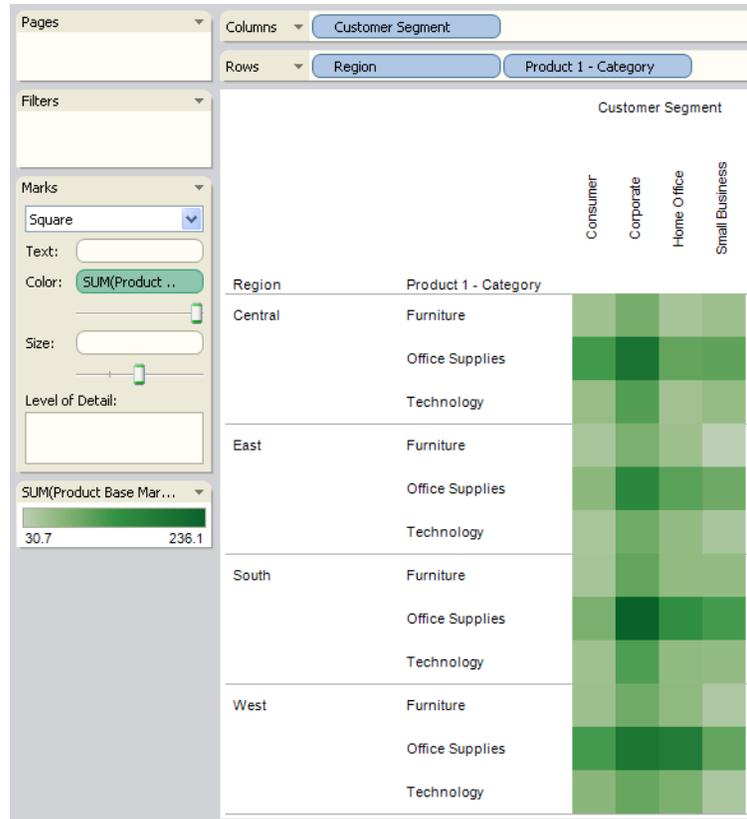
- Select the **Format > Cell Size > Square Cell** menu item.
- Increase the mark size with the **Size** slider so that the marks fill almost the entire table cell.



- Increase the size of the table using **Ctrl + Shift + B**.
- Rotate the column headers by right-clicking on one of the headers and selecting **Rotate Labels**.



The view is shown below.



You can see that technology has the highest profit across all regions and segments, while furniture has the lowest profit across all regions and markets.

The heat map allows you to easily pick out values that are near either edge of the continuous color map. For example, it's easy to see that technology in the eastern corporate market has the highest profit.



## Heat Map–Size Encoded

This view shows a basic heat map that is modified by size-encoding the data. It is constructed as a table of categorical data and is color-encoded by a measure.

Two examples are presented. The first example size-encodes the data using a measure. The second example size-encodes the data using a dimension.

You can create the view with the following procedure.

- 1 Place the **Customer Segment** dimension on the **Columns** shelf.

Headers are created with labels given by the dimension member names.

- 2 Place the **Region** and **Product 1 - Category** dimensions on the **Rows** shelf.

Headers are created with labels given by the dimension member names. You have now created a nested table of categorical data.

- 3 Place the **Profit** measure on the **Color** shelf.

The measure is automatically aggregated as a summation. The color legend reflects the continuous data range.

- 4 Place the **Sales** measure on the **Size** shelf.

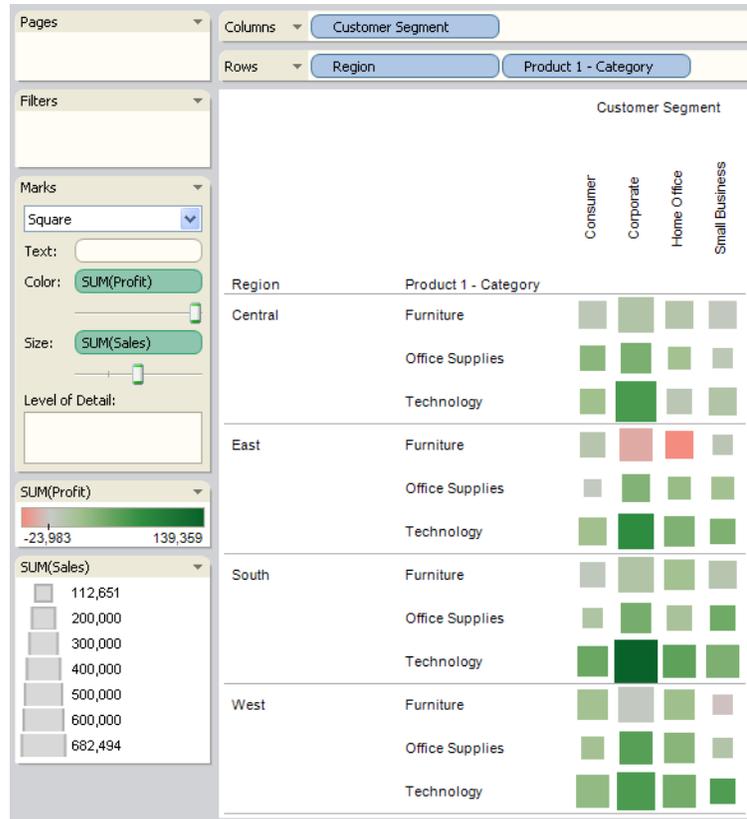
The measure is automatically aggregated as a summation. Change the aggregation to an average by selecting Average from the field's context menu.

- 5 Select the mark type to be Square.

- 6 Optimize the view format. Refer to the previous example “Heat Map–Basic” on page 33-56 for more information.



The view is shown below.



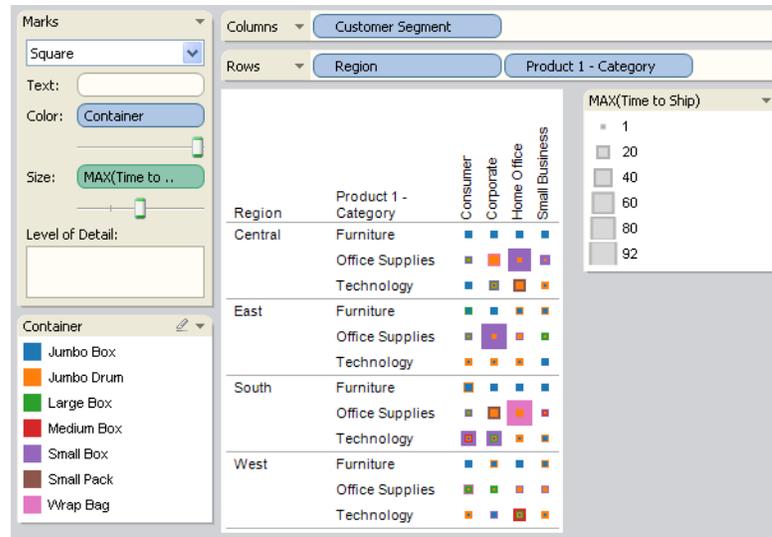
As in the previous example, you can see that technology has the highest profit across all regions and markets, while furniture has the lowest profit across all regions and markets.

Placing a measure on the **Size** shelf yields a continuous range of sizes. Because there is a wide range of average sales, there is a wide range of sizes. For example, the size encoding tells you that the average sales for office supplies in all regions and markets is much smaller than for technology and furniture. You can also conclude that the average sales for technology in the eastern region is about the same for all markets.



The size-encoded heat map allows you to easily pick out extremes in color and size. For example, it's easy to see that furniture in the central small business market has a negative profit and a small average sales. To create the second view, place the **Container** dimension on the **Color** shelf and the **Time to Ship** measure on the **Size** shelf. Aggregate the measure as an average by selecting Average from the field's context menu.

The modified view is shown below.



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. Because many of the marks are close together, their colors are obstructed by the mark borders. Therefore, mark borders are turned off.

You should use caution when building this type of data view because some of the marks might be obscured. For example, suppose the average times to ship for corporate furniture in the western region were the same for the jumbo box and the large box. In this case, the mark sizes would be the same and would overlap. Therefore, you would see only the last mark displayed in the view.



## Heat Map–Shape and Size Encoded

This view shows a basic heat map that is modified by shape-encoding and size-encoding the data. You can create the view with the following procedure.

- 1 Place the **Customer Segment** dimension on the **Columns** shelf.

Headers are created with labels given by the dimension member names.

- 2 Place the **Region** and **Product 1 - Category** dimensions on the **Rows** shelf.

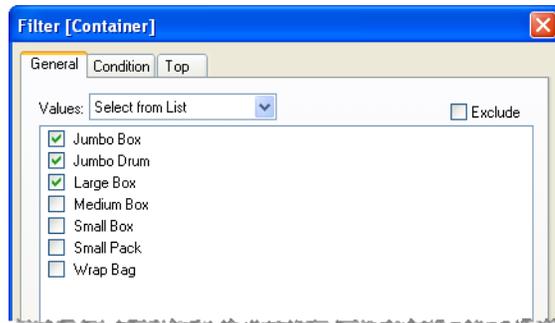
Headers are created with labels given by the dimension member names. You have now created a nested table of categorical data.

- 3 Place the **Container** dimension 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.

- 4 Filter **Container** so that only the first three members are displayed.

Filtering the data creates a less dense view so that you can more easily see individual marks. You open the **Filter** dialog box by selecting Filter from the field's context menu.



- 5 Place the **Time to Ship** measure on the **Size** shelf.

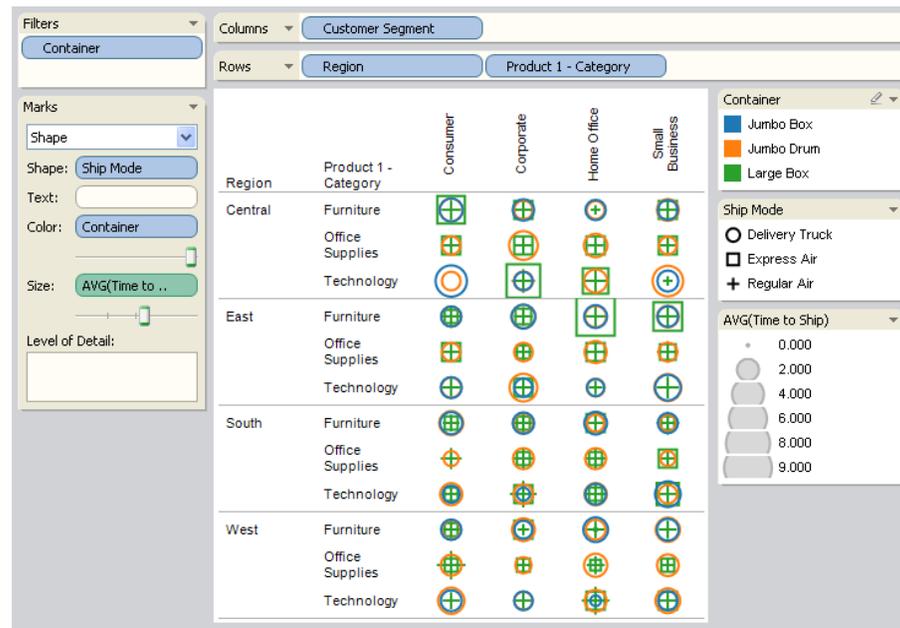
The measure is automatically aggregated as a maximum. Change the aggregation to an average by selecting Average from the field's context menu.

- 6 Select Shape as the mark type and place the **Ship Mode** dimension on the **Shape** shelf.

Placing a dimension on the **Shape** shelf 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.

- 7 Optimize the view format. Refer to the previous example “Heat Map–Basic” on page 33-56 for more information.

The view is shown below.



## Gantt Bar Chart

Gantt charts are typically used when you want to display the duration of one or more categories of interest against the progression of time. Each category uses a separate mark (usually a bar), where the length of the mark gives the duration of the category. For example, you could use a Gantt chart to display a product line versus order date with a mark length that gives the average delivery time.

You create a Gantt chart by selecting the Gantt bar mark type. This mark type is automatically selected if 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. Refer to “Gantt Bar Mark” on page 12-25 for more information about the this mark type.

In Tableau, a standard Gantt chart would have a continuous date on the horizontal axis and a dimension on the vertical axis. You would then complete the view by placing another time-related field on the **Size** shelf. Of course, you are not limited to this type of Gantt chart. You can build any view that is meaningful to you. For example, the length of the mark does not have to represent time. It can represent profit, sales, or any other quantity that provides insight into your data. Indeed, you do not need to include time anywhere in the view.

The following examples show you how to use Gantt charts to explore your data:

- **Gantt Bar Chart–Standard** – A standard Gantt chart that displays a dimension as a function of continuous time. The length of the bars is given by an aggregated time-based measure and the data are color-encoded by a dimension.
- **Gantt Bar Chart–Nonstandard** – A nonstandard Gantt chart that does not include time anywhere in the view. Instead, a measure is displayed against a dimension. The length of the Gantt bars is given by a measure and the data are color-encoded by a dimension.

### Gantt Bar Chart–Standard

This view shows a particular product category for the first half of 2007 as a Gantt chart. The view is size-encoded by the time to ship and is color-encoded by the ship mode. The view shows the time it takes to deliver certain products during a specified (continuous) time period. It is a standard Gantt chart application.

You can create the view with the following procedure.



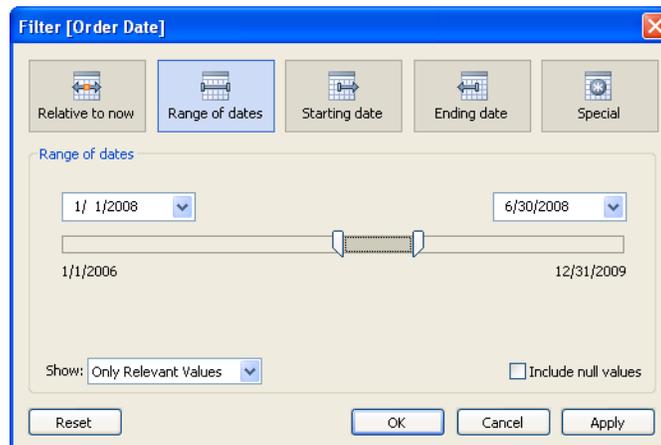
- 1 Place **Order Date** on the **Columns** shelf.

The date is automatically aggregated by year, and headers are created with labels given by the dimension member names. Display every row in the data source by selecting **All Values** from the field's context menu.

- 2 Place **Product 2 - Sub-Category** on the **Rows** shelf.

Headers are created with labels given by the dimension member names.

- 3 Filter **Order Date** to include only the first half of 2008.



- 4 Place **Time to Ship** on the **Size** shelf.

The measure is automatically aggregated as a summation. Disaggregate the data by selecting the **Analysis > Aggregate Measures** menu item.



---

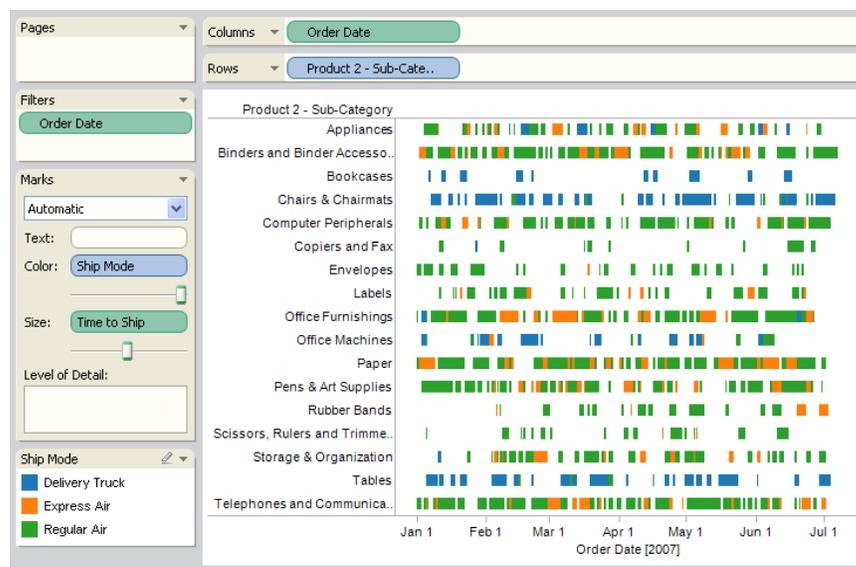
**Note** When a date is placed on the **Size** shelf, the units are automatically assumed to be days. For example, a value of 20 will map to a length corresponding to twenty days.

---

### 5 Place **Ship Mode** on the **Color** shelf.

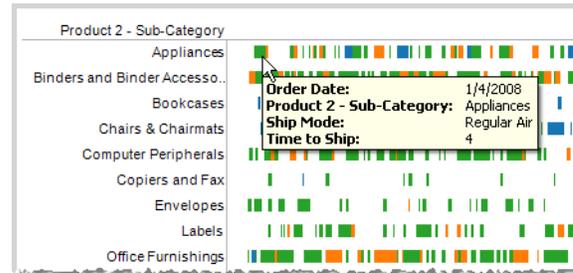
Placing this dimension on the **Color** shelf separates the marks according to the members in the dimension, and assigns a unique color to each member. The color legend displays each member name and its associated color.

The view is shown below.



Each row of the table gives you all the orders placed for a given product during the first half of 2007. The beginning (left edge) of each mark corresponds to the time the order was placed, the length of the mark corresponds to the time it took to ship, and the color of the mark corresponds to the delivery mode.

You can display the data associated with each mark using tooltips. For example, the data for the first appliance order is shown below. As you can see, the order was placed January 4, 2008 and the delivery time was 4days by regular air.



## Gantt Bar Chart–Nonstandard

This view shows the maximum sales for a particular product category as a Gantt chart. The data are color-encoded by the shipping mode and size-encoded by maximum profit. This is a nonstandard Gantt chart in that time is not included anywhere in the view.

You can create the view with the following procedure.

- 1 Place **Sales** on the **Columns** shelf.

The measure is automatically aggregated as a summation. Change the aggregation to a maximum by selecting **Maximum** from the field's right-click context menu.

- 2 Place **Product 2 - Sub-Category** on the **Rows** shelf.

Headers are created with labels given by the dimension member names.

- 3 Select the **Mark** type as **Gantt bar**.

- 4 Place **Profit** on the **Size** shelf.

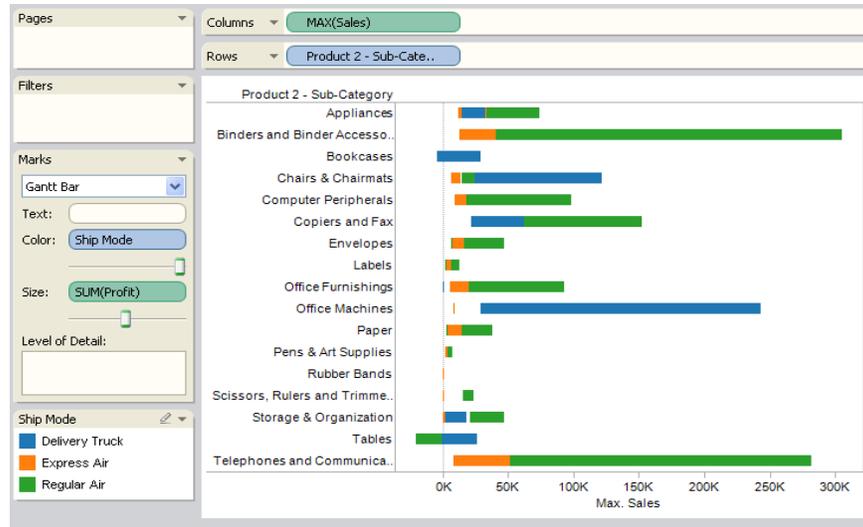
The measure is automatically aggregated as a summation. Change the aggregation to a maximum by selecting **Maximum** from the field's right-click context menu.

- 5 Place **Ship Mode** on the **Color** shelf.

Placing this 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.



The view is shown below.



Each row of the table gives you the maximum sales and profit totals for a given product. The left edge of each mark corresponds to the maximum sales, the length of the mark corresponds to the sum profit, and the color of the mark corresponds to the delivery method.

---

## Pie Chart

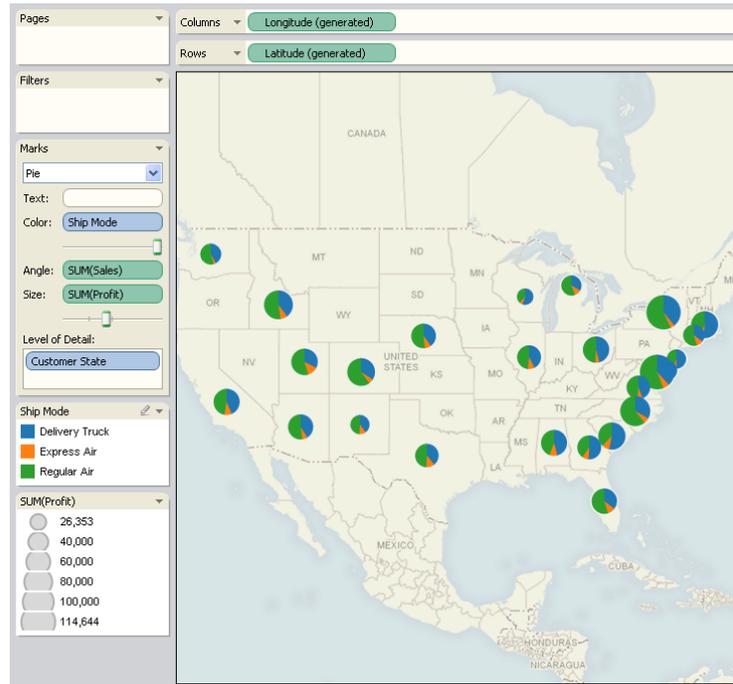
The pie mark can be used to show proportions. Although generally this type of information can be better shown using stacked bar charts (refer to “Stacking Marks” on page 12-27), there are cases where using pie marks can be very effective. For example, pie marks are very useful when you are trying to convey the percentage allocation of marketing expenses by state where the spending of geographically close states are very relevant. The following example shows you how to use a pie chart to explore your data.

### Pie Chart– Maps

This view shows the profit and sales by ship mode for a variety of stores across the United States. It uses the pie mark to show the proportion of sales for each ship mode and location. Then the overall size of the pie indicates the profit from those sales. Follow the steps below to build this view using the Sample - Superstore Sales data source.

- 1 Double-click the **State** field in the Data window. Note that it automatically assigned a geographic role. You can assign geographic roles by right-clicking the field and selecting **Geographic Role**.
- 2 Drag the **Ship Mode** field and drop it on the **Color** shelf.
- 3 Select the **Pie** mark type on the Mark menu.
- 4 Drag the **Sales** measure and drop it on the **Angle** shelf.
- 5 Drag the **Profit** measure and drop it on the **Size** shelf.

The final view is shown below. You can use the size slider on the Marks card to increase the overall scale of the marks.





# 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-20
Pass Through Functions (RAWSQL) . . . . .	A-21
User Functions . . . . .	A-23
Table Calculation Functions . . . . .	34-25
<b>Operators</b> . . . . .	<b>A-37</b>
Precedence . . . . .	A-39



## 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 34-15 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 10-6.



## **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 34-13). 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)
- User Functions
- Table Calculation 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 21-34.

### 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,



$\text{TAN}(\text{PI}() / 4) = 1.0$

### **ZN(expression)**

Returns the expression if it is not null, otherwise returns zero. Use this function to use zero values instead of null values.

$\text{ZN}([\text{Profit}]) = [\text{Profit}]$

## **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,

$\text{CHAR}(65) = 'A'$

### **Contains(string, substring)**

Returns true if the given string contains the specified substring.

$\text{CONTAINS}(\text{"Calculation"}, \text{"alcu"}) = \text{true}$

### **ENDSWITH(string, substring)**

Returns true if the given string ends with the specified substring. Trailing white spaces are ignored.

$\text{ENDSWITH}(\text{"Tableau"}, \text{"leau"}) = \text{true}$

### **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,

$\text{FIND}(\text{"Calculation"}, \text{"alcu"}) = 2$

$\text{FIND}(\text{"Calculation"}, \text{"Computer"}) = 0$

$\text{FIND}(\text{"Calculation"}, \text{"a"}, 3) = 7$



FIND("Calculation", "a", 2) = 2

FIND("Calculation", "a", 8) = 0

### **LEFT(string, number)**

Returns the left-most number of characters in the string. For example,

LEFT("Matador", 4) = "Mata"

### **LEN(string)**

Returns the length of the string. For example,

LEN("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) = " "

### **STARTSWITH(string, substring)**

Returns true if the given string starts with the specified substring. Leading white spaces are ignored. For example,

STARTSWITH("Joker", "Jo") = true

### **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 34-5 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. Make sure to use the `date_part` exactly as it is shown.

<code>date_part</code>	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(), DATETIME(), 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 second and third 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) = 2

INT(0.50/1.0) = 0

INT(0.501/1.0) = 0

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 “Groups” on page 17-18 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
```

**IFNULL(*expression1*, *expression2*)**

Returns the first expression if the result is not null, otherwise returns the second expression.

```
IFNULL([Profit], 0) = [Profit] if it is not null, otherwise returns 0
```

**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 to the Data Engine” on page 9-1.

### **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(“MIN(%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(“MIN(%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])

## **User Functions**

The user 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 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.

### **FULLNAME( )**

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]==FULLNAME( )

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.

### **ISFULLNAME(string)**

Returns True if the full name of the person currently using Tableau matches the given string. The full name for the person currently using Tableau is the Tableau Server full name if they are logged in, otherwise it's the Windows username.

Example: ISFULLNAME("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

### **ISUSERNAME(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: ISUSERNAME("dhallsten")

This returns true if dhallsten is the current user, otherwise it returns false.



### **USERDOMAIN()**

Returns the domain for the person currently using Tableau. This is the Tableau Server domain if the user is logged in otherwise this functions returns the Windows domain. Use this function in combination with other user functions to create calculations that are dependent on the current user and domain.

Example: [Manager]=USERNAME() AND [Domain]=USERDOMAIN()

### **USERNAME()**

Returns the username 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]=USERNAME()

If the manager dhallsten was logged in, this function would only return true if the Manager field in the view is also equal to dhallsten. 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.

## **Table Calculation Functions**

Table calculation functions are used to customize table calculations. Table Calculations are computations that are applied to the values in the entire table and are often dependent on the table structure itself. Refer to “Table Calculations” on page 21-40 to learn more about creating and customizing table calculations.

### **First()**

Returns the number of rows from the current row to the first row in the partition. For example, the view below shows quarterly sales. When FIRST() is computed within the Date partition, the offset of the first row from the second row is -1.



Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	\$160,877	\$231,411	\$133,934	\$185,961
	Q2	\$197,213	\$204,914	\$337,813	\$213,507
	Q3	\$302,678	\$165,201	\$283,806	\$206,512
	Q4	\$297,208	\$226,983	\$214,845	\$230,291
2010	Q1	\$180,609	\$180,123	\$273,943	\$251,145
	Q2	\$195,785	\$224,882	\$251,391	\$195,976
	Q3	\$116,613	\$50,363	\$194,601	\$102,731

\$160,877	0
\$197,213	-1
\$302,678	-2
\$297,208	-3
\$180,609	-4
\$195,785	-5
\$116,613	-6

Example (when current row index is 3): FIRST() = -2

### INDEX()

Returns the index of the current row in the partition. The first row index starts at 1. For example, the table below shows quarterly sales. When INDEX() is computed within the Date partition, the index of each row is 1, 2, 3, 4...etc.

Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	\$160,877	\$231,411	\$133,934	\$185,961
	Q2	\$197,213	\$204,914	\$337,813	\$213,507
	Q3	\$302,678	\$165,201	\$283,806	\$206,512
	Q4	\$297,208	\$226,983	\$214,845	\$230,291
2010	Q1	\$180,609	\$180,123	\$273,943	\$251,145
	Q2	\$195,785	\$224,882	\$251,391	\$195,976
	Q3	\$116,613	\$50,363	\$194,601	\$102,731

\$160,877	1
\$197,213	2
\$302,678	3
\$297,208	4
\$180,609	5
\$195,785	6
\$116,613	7

Example (for the third row in the partition): INDEX() = 3

### LAST()

Returns the number of rows from the current row to the last row in the partition. For example, the table below shows quarterly sales. When LAST() is computed within the Date partition, the offset of the last row from the second row is 5.



Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	\$160,877	\$231,411	\$133,934	\$185,961
	Q2	\$197,213	\$204,914	\$337,813	\$213,507
	Q3	\$302,678	\$165,201	\$283,806	\$206,512
	Q4	\$297,208	\$226,983	\$214,845	\$230,291
2010	Q1	\$180,609	\$180,123	\$273,943	\$251,145
	Q2	\$195,785	\$224,882	\$251,391	\$195,976
	Q3	\$116,613	\$50,363	\$194,601	\$102,731

LAST()	
\$160,877	6
\$197,213	5
\$302,678	4
\$297,208	3
\$180,609	2
\$195,785	1
\$116,613	0

Example (when the current row index is 3 of 7): LAST() = 4

### LOOKUP(expression, [offset])

Returns the value of the given expression in a target row, specified as a relative offset from the current row. Use FIRST() + n and LAST() - n for a target relative to the first/last rows in the partition. If offset is omitted, the Compare To row may be set on the field menu. Returns NULL if the target row cannot be determined.

For example, the view below shows quarterly sales. When LOOKUP (SUM(Sales), 2) is computed within the Date partition, each row will show the sales value from 2 quarters in the future.

Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	\$160,877	\$231,411	\$133,934	\$185,961
	Q2	\$197,213	\$204,914	\$337,813	\$213,507
	Q3	\$302,678	\$165,201	\$283,806	\$206,512
	Q4	\$297,208	\$226,983	\$214,845	\$230,291
2010	Q1	\$180,609	\$180,123	\$273,943	\$251,145
	Q2	\$195,785	\$224,882	\$251,391	\$195,976
	Q3	\$116,613	\$50,363	\$194,601	\$102,731

Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	\$302,678	\$165,201	\$283,806	\$206,512
	Q2	\$297,208	\$226,983	\$214,845	\$230,291
	Q3	\$180,609	\$180,123	\$273,943	\$251,145
	Q4	\$195,785	\$224,882	\$251,391	\$195,976
2010	Q1	\$116,613	\$50,363	\$194,601	\$102,731
	Q2				
	Q3				

Example: LOOKUP(SUM([Profit]), FIRST()+2) = SUM([Profit]) in the third row of the partition.

### PREVIOUS\_VALUE(expression)

Returns the value of this calculation in the previous row. Returns the given expression if the current row is the first row of the partition.

Example: SUM([Profit]) \* PREVIOUS\_VALUE(1) = running product of SUM(Profit)

### RUNNING\_AVG(expression)

Returns the running average of the given expression, from the first row in the partition to the current row.

For example, the view below shows quarterly sales. When RUNNING\_AVG(SUM([Sales])) is computed within the Date partition, the result is a running average of the sales values for each quarter.

Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	\$160,877	\$231,411	\$133,934	\$185,961
	Q2	\$197,213	\$204,914	\$337,813	\$213,507
	Q3	\$302,678	\$165,201	\$283,806	\$206,512
	Q4	\$297,208	\$226,983	\$214,846	\$230,291
2010	Q1	\$180,609	\$180,123	\$273,943	\$251,145
	Q2	\$195,785	\$224,882	\$251,391	\$195,976
	Q3	\$116,613	\$50,363	\$194,601	\$102,731

Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	160,877	231,411	133,934	185,961
	Q2	179,045	18,162	235,873	199,734
	Q3	220,256	200,509	251,851	201,993
	Q4	239,494	207,127	242,599	209,068
2010	Q1	227,717	201,726	248,868	217,483
	Q2	222,395	205,586	249,289	213,899
	Q3	207,283	183,411	241,476	198,018

Example: RUNNING\_AVG(SUM([Profit])) = running average of SUM(Profit)

### RUNNING\_COUNT(expression)

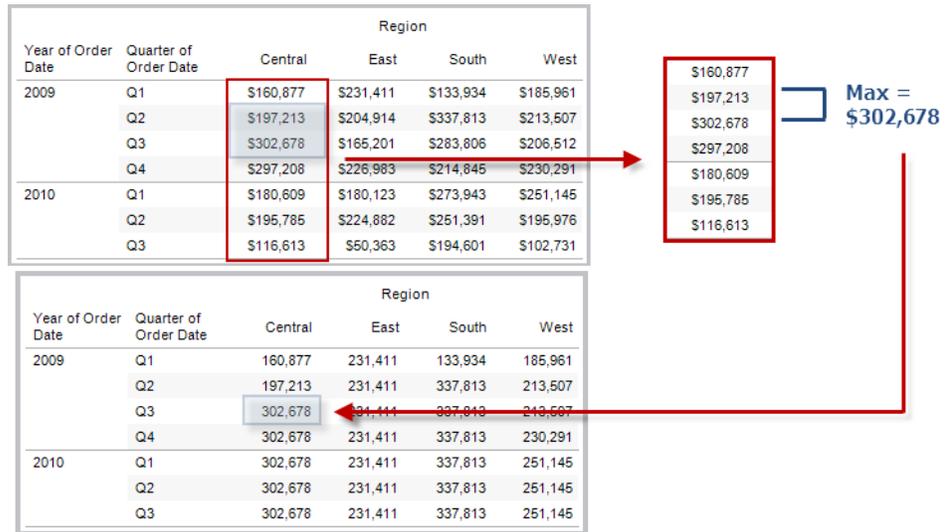
Returns the running count of the given expression, from the first row in the partition to the current row.



Example: `RUNNING_COUNT(SUM([Profit]))` = running count of `SUM(Profit)`

### **RUNNING\_MAX(expression)**

Returns the running maximum of the given expression, from the first row in the partition to the current row.



Example: `RUNNING_MAX(SUM([Profit]))` = running maximum of `SUM(Profit)`

### **RUNNING\_MIN(expression)**

Returns the running minimum of the given expression, from the first row in the partition to the current row.



Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	\$160,877	\$231,411	\$133,934	\$185,961
	Q2	\$197,213	\$204,914	\$337,813	\$213,507
	Q3	\$302,678	\$165,201	\$283,806	\$206,512
	Q4	\$297,208	\$226,983	\$214,845	\$230,291
2010	Q1	\$180,609	\$180,123	\$273,943	\$251,145
	Q2	\$195,785	\$224,882	\$251,391	\$195,976
	Q3	\$116,613	\$50,363	\$194,601	\$102,731

Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	160,877	231,411	133,934	185,961
	Q2	160,877	204,914	133,934	185,961
	Q3	160,877	165,201	133,934	185,961
	Q4	160,877	165,201	133,934	185,961
2010	Q1	160,877	165,201	133,934	185,961
	Q2	160,877	165,201	133,934	185,961
	Q3	116,613	50,363	133,934	102,731

Min = \$106,877

Example:  $\text{RUNNING\_MIN}(\text{SUM}([\text{Profit}]))$  = running minimum of  $\text{SUM}(\text{Profit})$

### **$\text{RUNNING\_SUM}(\text{expression})$**

Returns the running summation of the given expression, from the first row in the partition to the current row.



Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	\$160,877	\$231,411	\$133,934	\$185,961
	Q2	\$197,213	\$204,914	\$337,813	\$213,507
	Q3	\$302,678	\$165,201	\$283,806	\$206,512
	Q4	\$297,208	\$226,983	\$214,845	\$230,291
2010	Q1	\$180,609	\$180,123	\$273,943	\$251,145
	Q2	\$195,785	\$224,882	\$251,391	\$195,976
	Q3	\$116,613	\$50,363	\$194,601	\$102,731

Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	160,877	231,411	133,934	185,961
	Q2	358,090	436,325	471,747	399,469
	Q3	660,768	41,526	755,553	685,988
	Q4	957,976	828,508	970,398	836,272
2010	Q1	1,138,585	1,008,631	1,244,341	1,087,417
	Q2	1,334,369	1,233,513	1,495,732	1,283,392
	Q3	1,450,982	1,283,877	1,690,333	1,386,123

Example: `RUNNING_SUM(SUM([Profit]))` = running sum of `SUM(Profit)`

### Size()

Returns the number of rows in the partition. For example, the view below shows quarterly sales. Within the Date partition, there are seven rows so the `Size()` of the Date partition is 7.

Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	\$160,877	\$231,411	\$133,934	\$185,961
	Q2	\$197,213	\$204,914	\$337,813	\$213,507
	Q3	\$302,678	\$165,201	\$283,806	\$206,512
	Q4	\$297,208	\$226,983	\$214,845	\$230,291
2010	Q1	\$180,609	\$180,123	\$273,943	\$251,145
	Q2	\$195,785	\$224,882	\$251,391	\$195,976
	Q3	\$116,613	\$50,363	\$194,601	\$102,731

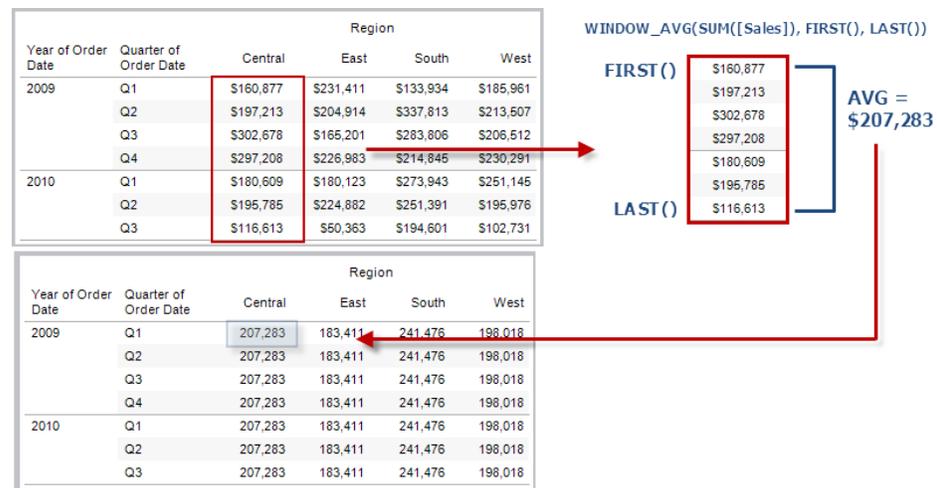
Example (partition has 5 rows): `SIZE() = 5`



### WINDOW\_AVG(expression, [start, end])

Returns the average of the expression within the window. The window is defined as offsets from the current row. Use FIRST()+n and LAST()-n for offsets from the first or last row in the partition. If the start and end are omitted, the entire partition is used.

For example, the view below shows quarterly sales. A window average within the Date partition returns the average sales across all dates.



Example: WINDOW\_AVG(SUM[Profit]), FIRST()+1, 0) = Average of SUM(Profit) from the second row to the current row

### WINDOW\_COUNT(expression, [start, end])

Returns the count of the expression within the window. The window is defined as offsets from the current row. Use FIRST()+n and LAST()-n for offsets from the first or last row in the partition. If the start and end are omitted, the entire partition is used.

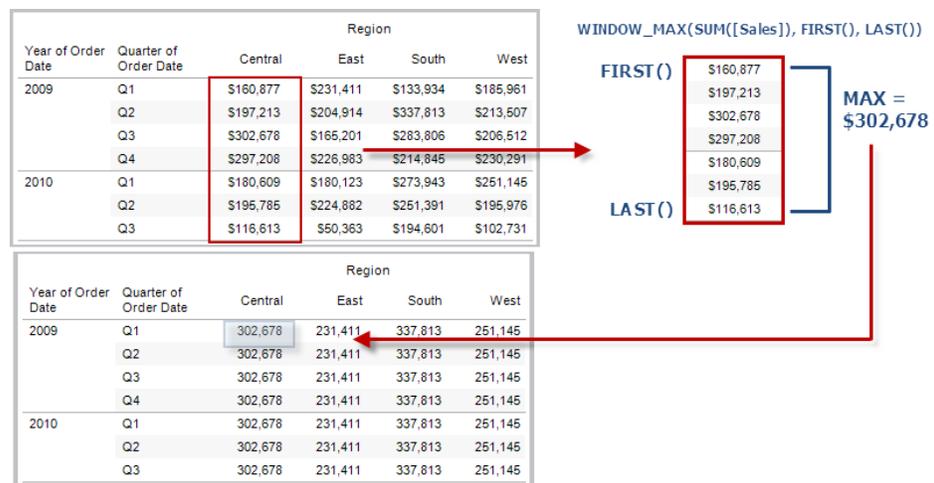
Example: WINDOW\_COUNT(SUM[Profit]), FIRST()+1, 0) = Count of SUM(Profit) from the second row to the current row

### WINDOW\_MAX(expression, [start, end])

Returns the maximum of the expression within the window. The window is defined as offsets from the current row. Use FIRST()+n and LAST()-n for offsets from the first or last row in the partition. If the start and end are omitted, the entire partition is used.



For example, the view below shows quarterly sales. A window maximum within the Date partition returns the maximum sales across all dates.



Example: WINDOW\_MAX(SUM[Profit]), FIRST()+1, 0) = Maximum of SUM(Profit) from the second row to the current row

### WINDOW\_MIN(expression, [start, end])

Returns the minimum of the expression within the window. The window is defined as offsets from the current row. Use FIRST()+n and LAST()-n for offsets from the first or last row in the partition. If the start and end are omitted, the entire partition is used.

For example, the view below shows quarterly sales. A window minimum within the Date partition returns the minimum sales across all dates.



Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	\$160,877	\$231,411	\$133,934	\$185,961
	Q2	\$197,213	\$204,914	\$337,813	\$213,507
	Q3	\$302,678	\$165,201	\$283,806	\$206,512
	Q4	\$297,208	\$226,983	\$214,845	\$230,291
2010	Q1	\$180,609	\$180,123	\$273,943	\$251,145
	Q2	\$195,785	\$224,882	\$251,391	\$195,976
	Q3	\$116,613	\$50,363	\$194,601	\$102,731

**WINDOW\_MIN(SUM([Sales]), FIRST(), LAST())**

<b>FIRST()</b>	\$160,877
	\$197,213
	\$302,678
	\$297,208
<b>LAST()</b>	\$180,609
	\$195,785
	\$116,613

**MIN = \$116,613**

Year of Order Date	Quarter of Order Date	Region			
		Central	East	South	West
2009	Q1	116,613	50,363	133,934	102,731
	Q2	116,613	50,363	133,934	102,731
	Q3	116,613	50,363	133,934	102,731
	Q4	116,613	50,363	133,934	102,731
2010	Q1	116,613	50,363	133,934	102,731
	Q2	116,613	50,363	133,934	102,731
	Q3	116,613	50,363	133,934	102,731

Example:  $\text{WINDOW\_MIN}(\text{SUM}[\text{Profit}], \text{FIRST}()+1, 0) = \text{Minimum of SUM(Profit) from the second row to the current row}$

### **WINDOW\_STDEV(expression, [start, end])**

Returns the sample standard deviation of the expression within the window. The window is defined as offsets from the current row. Use  $\text{FIRST}()+n$  and  $\text{LAST}()-n$  for offsets from the first or last row in the partition. If the start and end are omitted, the entire partition is used.

Example:  $\text{WINDOW\_STDEV}(\text{SUM}[\text{Profit}], \text{FIRST}()+1, 0) = \text{Standard deviation of SUM(Profit) from the second row to the current row}$

### **WINDOW\_STDEVP(expression, [start, end])**

Returns the biased standard deviation of the expression within the window. The window is defined as offsets from the current row. Use  $\text{FIRST}()+n$  and  $\text{LAST}()-n$  for offsets from the first or last row in the partition. If the start and end are omitted, the entire partition is used.

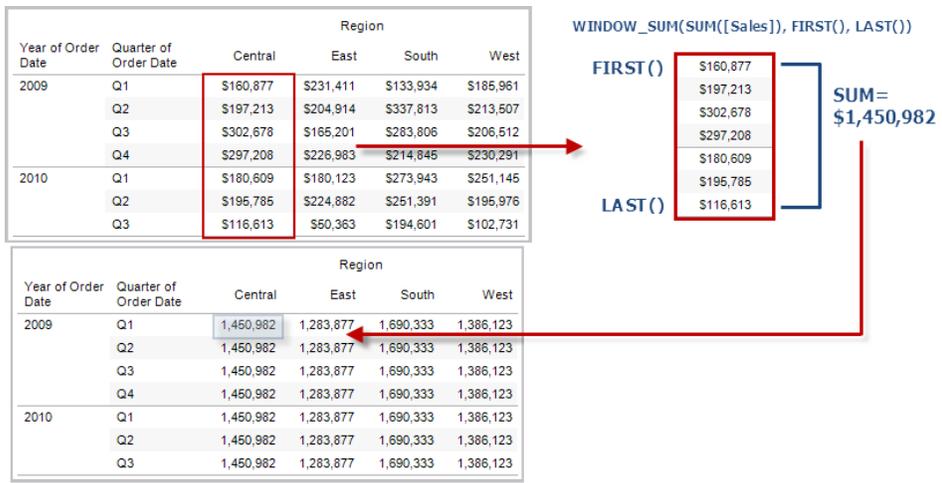
Example:  $\text{WINDOW\_STDEVP}(\text{SUM}[\text{Profit}], \text{FIRST}()+1, 0) = \text{Standard deviation of SUM(Profit) from the second row to the current row}$



### WINDOW\_SUM(expression, [start, end])

Returns the summation of the expression within the window. The window is defined as offsets from the current row. Use FIRST()+n and LAST()-n for offsets from the first or last row in the partition. If the start and end are omitted, the entire partition is used.

For example, the view below shows quarterly sales. A window sum computed within the Date partition returns the summation of sales across all quarters.



Example: WINDOW\_SUM(SUM[Profit], FIRST()+1, 0) = Summation of SUM(Profit) from the second row to the current row

### WINDOW\_VAR(expression, [start, end])

Returns the sample variance of the expression within the window. The window is defined as offsets from the current row. Use FIRST()+n and LAST()-n for offsets from the first or last row in the partition. If the start and end are omitted, the entire partition is used.

Example: WINDOW\_VAR(SUM[Profit], FIRST()+1, 0) = Variance of SUM(Profit) from the second row to the current row

### WINDOW\_VARP(expression, [start, end])

Returns the biased variance of the expression within the window. The window is defined as offsets from the current row. Use FIRST()+n and LAST()-n for offsets from the first or last row in the partition. If the start and end are omitted, the entire partition is used.



Example: `WINDOW_VARP(SUM[Profit], FIRST()+1, 0)` = Variance of SUM(Profit)  
from the second row to the current row



## 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: = (equal to), > (greater than), < (less than), >= (greater than or equal to), <= (less than or equal to), != (not 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)
3	*, /, %
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:

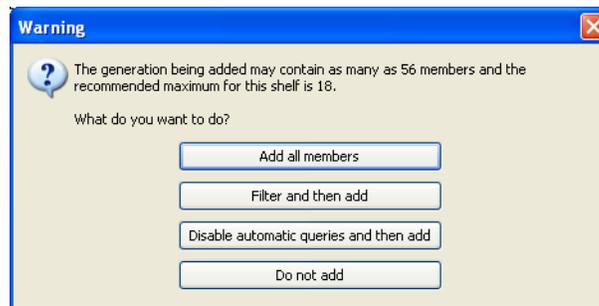
- **All Data Sources**- general performance tips that can be applied to all data sources.
- **Relational Data Sources**- performance tips specific to relational data sources.
- **Multidimensional Data Sources**- performance tips specific to multidimensional data sources.
- **Speeding up Context Filters**- tips on increasing the performance of context filters.
- **Extracting Large Text Files**- tips for working with large text files.

---

## Performance Tips: All Data Sources

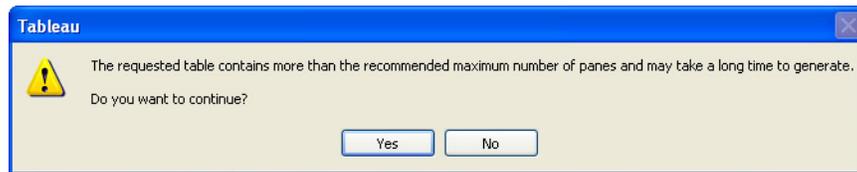
The following performance tips apply to all data sources supported by Tableau.

- **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 10-3 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 16-50 for how to create context filters. For more information about performance improvement with context filters refer to “Speeding up Context Filters” on page 35-8.
- **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 17-29, 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 35-9.
- **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 17-29, 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 5-8.
- 2 When you click **OK** you are prompted to create an extract file. You can also select **Data > Extract**. You can specify filters and options to define a subset of data or click **Extract** to include all the data. Refer to “Creating an Extract” on page 9-3 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.

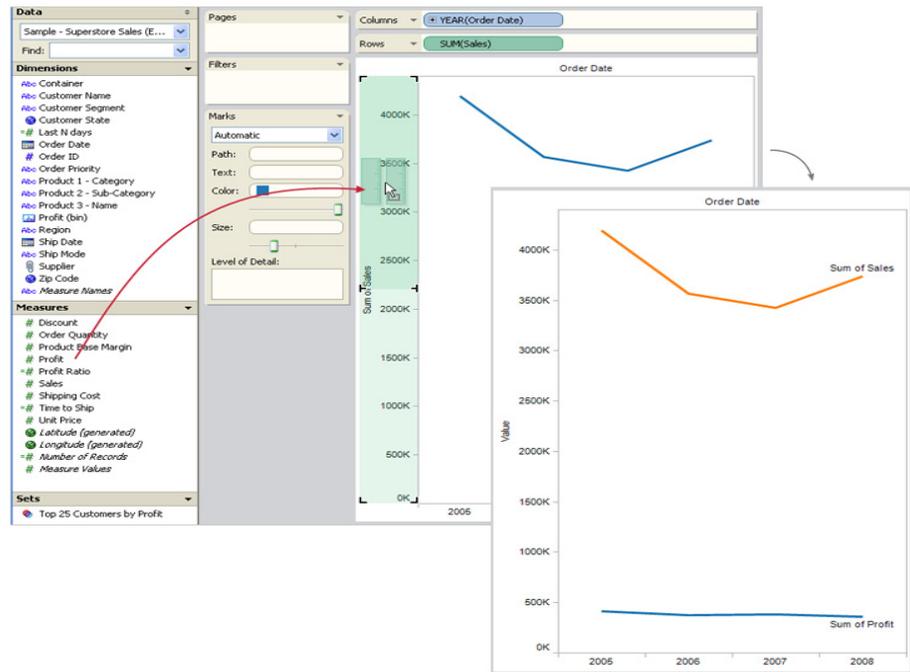
- Show Multiple Measures on a Single Axis
- Convert Measures to Dimensions
- Convert Dimensions to Measures
- Duplicate Any View as a Cross-tab
- Manually Sort Items in the Color Legend
- Duplicate Fields in the View
- Resize Individual Views in a Dashboard
- Use the Pages Shelf to Split Large PDF Reports
- 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.

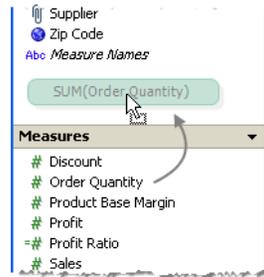
In the example below, both the Sales field and the Profit field share a common axis. This view was created by dragging the Sales field to the Row Shelf and then dragging the Profit field onto the axis created by the Sales field. The result is the creation of a “multiple measures” axis that is labeled ‘value.’ For a complete description refer to “Measure Values and Measure Names” on page 11-11.



---

## Convert Measures to Dimensions

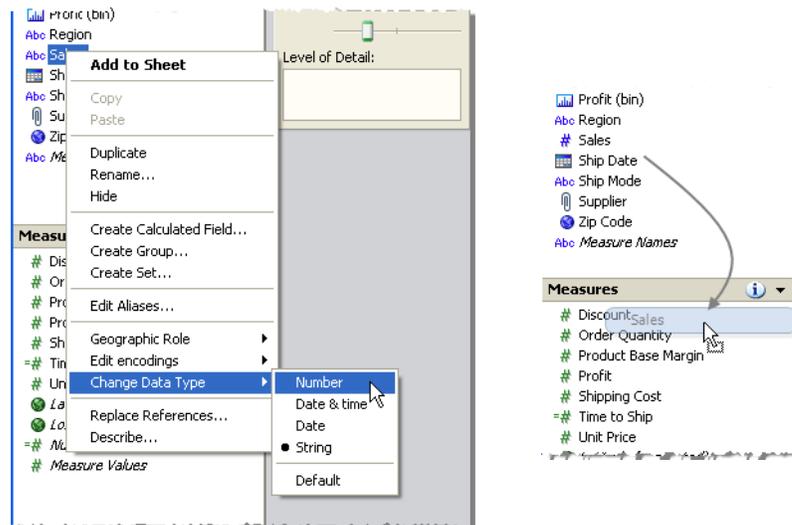
By default, Tableau treats all relational fields containing numbers as measures. However, you might decide that some of these fields should be treated as dimensions. For example, a field containing ages may be categorized as a measure by default in Tableau because it contains numeric data. However, if 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 11-38.



---

## 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.

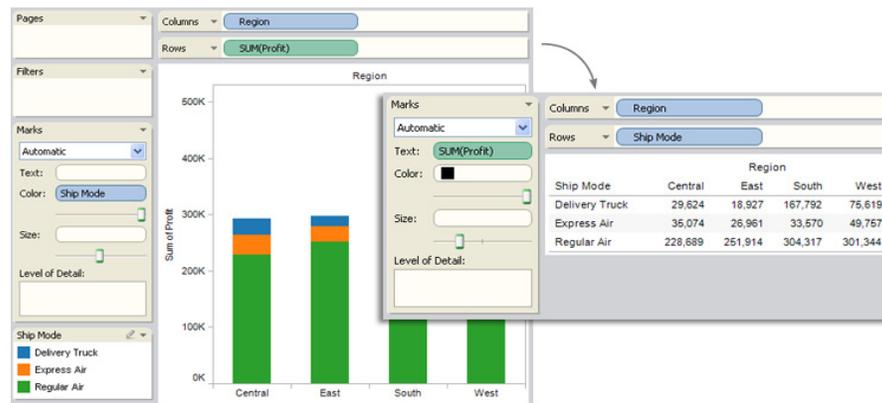


---

## 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.

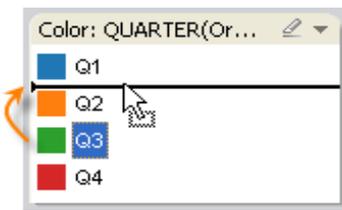
You can also export a cross-tab directly to Microsoft Excel by selecting **File > Export > Cross-tab to Excel**. For more information about exporting as a cross-tab refer to “Export Cross-tab to Excel” on page 31-20.





## Manually Sort Items in the Color Legend

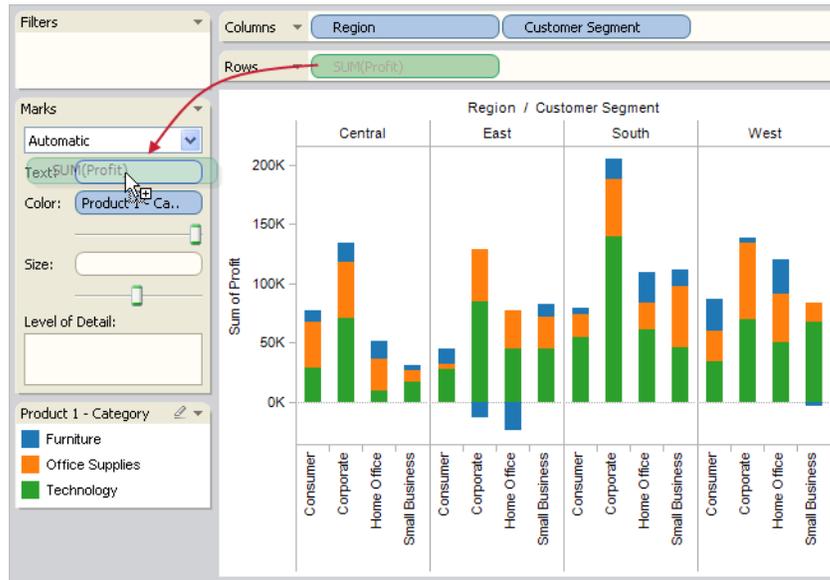
Manually sorting the items in the color legend change the drawing order of a field so you can see obscured data in your views. For instance, if you can't see green marks in a scatter plot because they are obscured by orange marks, you can manually sort the items in the color legend so that the green 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 17-16 to learn more.



---

## Duplicate Fields in the View

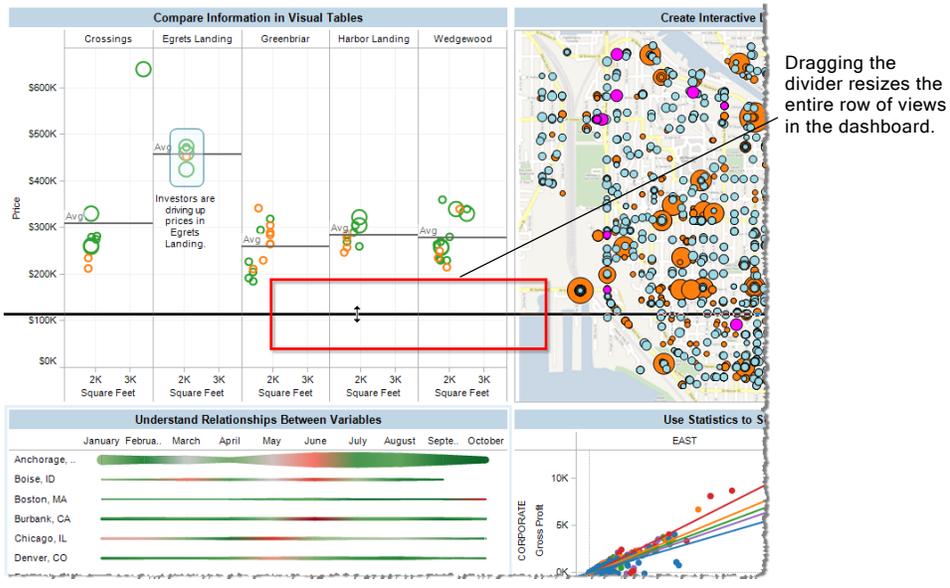
Once you add fields to the view, you can duplicate them simply by hold down the Ctrl key on your keyboard while you drag it to another shelf. For example, in the view below, the SUM(Profit) field is duplicated and placed on the Text shelf.



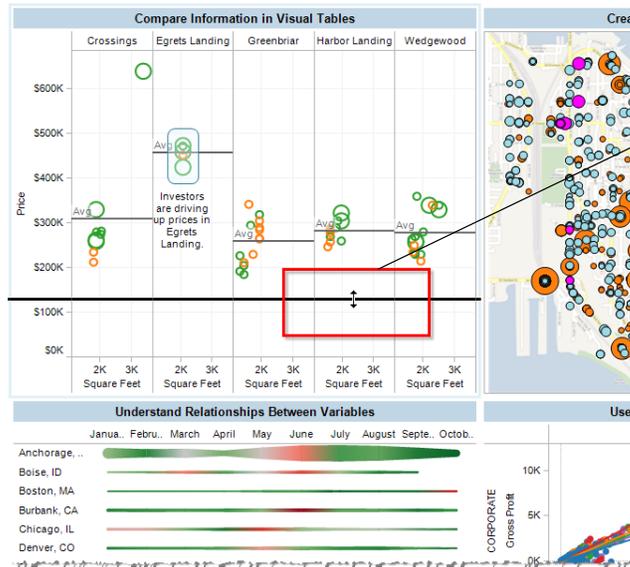


# Resize Individual Views in a Dashboard

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.

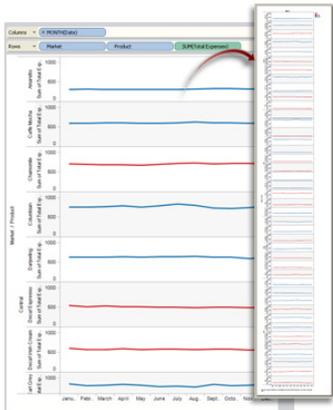


Hold down the Shift key while dragging the divider resizes an individual view. The views on the right are not resized.

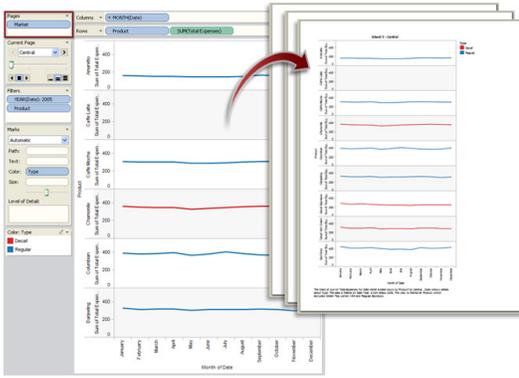


# Use the Pages Shelf to Split Large PDF Reports

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 32-3 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.

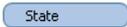
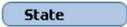
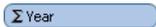
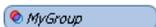
Visual Cue	Description
	The field contains text values.
	The field contains numeric values.
	The field contains only date values.
	The field contains both date and time values.
	The field contains geographical data and has been assigned a geographic role. Use these field when building map views.
	The field contains boolean (true or false) values.
	The field is a user-defined set.
	The field is a server named set.



Visual Cue	Description
	The field is a user filter, used when publishing to Tableau Server.
	The field is a numeric bin.
	The field is an ad-hoc group.
	The field is an attribute of a cube.
	The field is a varying 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

---



## General 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



<b>Keyboard Shortcut</b>	<b>Description</b>
Ctrl+Alt+P	Place selected field on pages shelf
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



<b>Keyboard Shortcut</b>	<b>Description</b>
Alt+F4	Closes the current workbook
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 forward through open worksheets
Ctrl+Shift+Tab, Ctrl+Shift+F6	Cycle backward through open worksheets
Shift+F6	Select mode
Shift+F7	Pan 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.

## **Navigation Tool Shortcuts**

In addition to the standard keyboard shortcuts there are several key combinations that can make switching between navigation tools fast and easy. The table below explains how you can extend each of the navigation tools located in the standard toolbar.



Navigation Tool	Keyboard/Mouse Action	Description
	Click	Selects the mark.
	Drag	Selects a group of marks.
	Ctrl + Click	Adds individual marks to the selection.
	Ctrl + Drag	Adds a group of marks to the selection.
	Shift + Click	Adds individual marks to the selection.
	Shift + Drag	Adds a group of marks to the selection.
	Alt + Click	Centers and focuses on the point you clicked. Focusing adds a filter and locks the axis to only show the relevant data.
	Alt + Drag	Focuses on the selected area.
	Alt + Ctrl + Click	Centers and focuses on the point you clicked.
	Alt + Ctrl + Drag	Focuses on the selected area.
Alt + Shift + Drag	Pans around the view.	



Navigation Tool	Keyboard/Mouse Action	Description
<b>Zoom Tool</b> 	Click	Centers and focuses on the point you clicked.
	Drag	Focuses on the selected area.
	Ctrl + Click	Centers and focuses on the point you click
	Ctrl + Drag	Focuses on the selected area.
	Shift + Drag	Pans around the view.
	Alt + Click	Centers and zooms in on the point you clicked.
	Alt + Drag	Zooms in on the selected area.
	Alt + Ctrl + Click	Centers and zooms out from the point you clicked.
	Alt + Ctrl + Drag	Zooms out from the selected area.
	Alt + Shift + Drag	Pans around the view.
<b>Pan Tool</b> 	Drag	Pans around the view.
	Ctrl + Drag	Pans around the view.
	Shift + Drag	Pans around the view.
	Alt + Drag	Scrolls across multiple panes in a large view.
	Alt + Ctrl + Drag	Scrolls across multiple panes in a large view.
	Alt + Shift + Drag	Scrolls across multiple panes in a large view.

## Glossary

---



### **Action**

An action is interactivity that you can add to your views. There are three kinds of actions: Filter, Highlight, and URL. Refer to “Actions” on page 20-1 to learn more.

### **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.

### **Custom Geocoding**

Tableau comes with built in geocoding data so that you can plot your data on a map. Custom Geocoding is when you add your own location data to extend the built in geocoding. Refer to “Custom Geocoding” on page 24-16 to learn more.

### **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 customer 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, 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.

### **Group**

A group is a field consisting of dimension members that have been combined into higher level categories. For example, a dimension that contains states could be grouped into regions using Groups. Groups are marked with a paper clip icon in the Data window.

### **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 an Oracle 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.

# Copyright Information

---

©2010 Tableau Software, Incorporated and its licensors. All rights reserved.

Protected by U.S. Patent 7,089,266; Patents Pending.

Portions of the code copyright ©2002 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.



# Contact

---

For sales inquiries, please contact:  
[sales@tableausoftware.com](mailto:sales@tableausoftware.com)

For customer support, please contact:  
[support@tableausoftware.com](mailto:support@tableausoftware.com)

For other inquiries, please contact:  
[info@tableausoftware.com](mailto:info@tableausoftware.com)

Telephone:  
(206) 633-3400



## A

- About Shelves 13-10
- Access database 5-9
- Accessing Data 5-4
- Accounts Dimension 5-49
- Actions 20-1
  - filter 20-3
  - highlight 20-8, 20-13
  - in dashboards 20-22
  - running 20-20
  - using fields as parameters 20-24
- Add to sheet
  - double-click 14-6
- Ad-hoc Groups 3-9, 17-18
  - creating 17-18
  - editing 17-22
  - finding members 17-25
- Aggregate Calculations 21-34
  - about 21-34
  - creating 21-35
  - disaggregated 21-36
- Aggregate Functions 34-20
  - AVG 34-20
  - COUNT 34-20
  - COUNTTD 34-20
  - MAX 34-20
  - MIN 34-20
  - STDEV 34-20
  - STDEVP 34-20
  - SUM 34-20
  - VAR 34-21
  - VARP 34-21
- Aggregate Measures 21-9
  - in extracts 9-4
- Aggregation Labels
  - removing from measure names 11-32
- Aggregations 21-4, 21-8
  - ATTR 21-4, 21-70
  - average 21-4
  - count 21-6
  - count distinct 21-6
  - dimension 21-4
  - grand totals 21-80
  - maximum 21-5
  - measures 21-8
  - minimum 21-4
  - percentages 21-88
  - standard deviation 21-5
  - standard deviation population 21-5
  - sum 21-4
  - variance 21-5
  - variance population 21-5
- Alerts 3-10
  - show/hide 12-13
- Aliases
  - changing measure names 11-32
  - editing data labels 28-19
  - multidimensional 11-22
  - Relational 11-21
- Alignment
  - formatting 27-9
- Analysis
  - aggregate measures 21-8
- Analysis Services 5-33
- Annotations 28-3
  - adding 28-3
  - formatting 28-11
  - rearranging 28-5
  - removing 28-15
- Aspect Ratio 27-41
- asymmetric layout 13-53
- ATTR Aggregation 21-70
- Automatic mark type 12-12

Automatic Updates 3-8

automatic updates 10-3

Automatic Views

Double-Click 14-6

Average 21-4

Axes

appearance 27-50

creating in a table 13-10

formatting 27-43

range 27-43

reverse 27-50

scale 27-43

tick marks 27-52

title 27-50

## B

Background Images 23-1

adding 23-3

aspect ratio 23-5

can't see marks 13-30

disabling 23-9

editing 23-8

enabling 23-9

removing 23-13

show/hide conditions 23-10

Background Maps 24-1

defaults 24-25

map sources 24-23

WMS 24-23

Banding 27-10

Bar

mark type 12-16

Bar Chart 33-4

basic 33-5

color encoded and sorted 33-14

filtered and color encoded 33-10

Binned Data 21-72

Blending Data 8-1

Bookmarks 31-7

deleting 31-8

removing 31-8

Boolean data type 11-34

Borders

formatting 27-12

brushing 33-42

Build Views Automatically

double-click 14-6

Building Views

Hyperion Essbase 13-54

Microsoft Analysis Services Cube 13-58

## C

Calculated Fields

creating 21-18

Calculated Members

creating 21-21

Calculations

aggregate calculations 21-34

aggregations 21-4, 21-8

binned data 21-72

grand totals 21-78

local vs. server 21-77

new fields 21-18

percentages 21-87

subtotals 21-84

user filters 30-45

writing formulas 21-24

Cancel Query

abandoned queries 10-5

reconnecting 10-5

Cards 3-11

caption 3-13

- card menu 3-14
- color legend 3-12
- columns 3-12
- current page 3-13
- filters 3-12
- map legend 3-12
- map options 3-13, 24-7
- marks 3-13
- measure names/measure values 3-12
- pages 3-12
- quick filters 3-13
- rows 3-12
- shape legend 3-12
- size legend 3-12
- summary 3-13, 19-11
- title 3-13
- Cell Size
  - clear 27-40
  - Narrower 27-33
  - shorter 27-33
  - square cell 27-37
  - taller 27-33
  - text cell 27-37
  - wider 27-33
- Chart Type
  - histogram 21-72
- chart type
  - bar 33-4
  - Gantt Bar Chart 33-63
  - heat map 33-55
  - line 33-30
  - scatter 33-39
  - text table 33-22
- Circle mark type 12-19
- Clear 3-8
  - manual sizing 27-40
- clear worksheet 3-19
- Clipboard Data Source 6-1
- Close
  - application 3-3
- Close Connection 5-77
- Collaborate
  - data sources 30-19
  - import data source 30-26
  - open workbooks 30-23
  - workbooks 30-3
- Color
  - full color range 13-27
  - limiting the range 13-28
  - resetting the range 13-28
  - reverse color palette 13-27
- Color Legend Highlighting 20-9
- Color shelf 13-23
- Column Attributes 27-41
- Column Labels
  - formatting 27-41
  - maximum number 27-41
- Columns shelf 13-10
- Comments. See Annotations
- Compute Using 21-69
- Connecting 5-4
  - Extract Data File 5-16
  - Firebird 5-22
  - Hyperion Essbase 5-49
  - MS Access 5-9
  - MS Analysis Services Cube 5-33
  - MS Excel 5-11
  - MySQL 5-41
  - OData 5-14
  - Oracle 5-46
  - PostgreSQL 5-54
  - SQL Server 5-38
  - SSAS 5-33
  - Tableau Data Extract 5-16

Text of CSV file 5-17  
Windows Azure Marketplace DataMarket 5-20

Connection  
adding tables 7-3  
close 5-77  
data source 5-4  
duplicate 5-73  
editing tables 7-8  
exporting 5-74  
properties 5-72  
rename 5-72  
saving 5-74

Continuous  
dates 18-12

Continuous and Discrete 11-37

Converting  
discrete to continuous 11-40  
measures to dimensions 11-38

Converting data types 34-15

DATE 34-16  
DATETIME 34-16  
FLOAT 34-16  
INT 34-16  
STR 34-16

Copy & Paste  
ad-hoc groups 21-24  
calculated fields 21-24  
data source 6-1  
sets 21-24  
user filters 21-24

copy image 31-21

Count 21-6

Count Distinct 21-6

Cross-Tab  
duplicating worksheets 3-19

cross-tab 33-22

CSV file 5-17

mixed data types 11-35

Custom SQL 5-68

Custom Table Calculations 21-68

## D

Dashboards 29-1

actions 20-22  
adding viewing 29-3  
creating 29-3  
organizing 29-10  
rearranging views 29-15  
removing views 29-8  
show/hide 29-10

Data Blending 8-1

Data Integration 8-1

Data Labels 28-16

aliases 28-19  
rearranging 28-18  
show and hide 28-16

Data Roles 11-34, 11-37

continuous Dimensions 11-46  
continuous dimensions 11-42  
continuous measures 11-47  
continuous vs. discrete 11-39  
converting discrete to continuous 11-40  
converting measures to dimensions 11-38  
dimension vs. measure 11-37  
Discrete Dimension 11-44  
Discrete Measures 11-45  
discrete measures 11-42

Data Sources

clipboard 6-1  
close 5-77  
connection 5-4  
connection examples 5-8  
duplicate 5-73

- export connection 5-74
- properties 5-72
- refresh 5-76
- rename 5-72
- supported 5-3
- Data Type 11-34, 34-3
  - BOOLEAN 34-4
  - changing 11-34
  - DATE/TIME 34-3
  - mixed 11-35
  - mixed (Excel and CSV) 11-35
  - NUMBER 34-3
  - STRING 34-3
- Data Window 3-7
  - features and functions 11-13
- Database
  - connecting 5-4
- Database Columns 3-7
- Date data type 11-34
- Date Functions 34-13
  - DATEADD 34-13
  - DATEDIFF 34-13
  - DATENAME 34-14
  - DATEPART 34-14
  - DATETRUNC 34-14
  - DAY 34-14
  - ISDATE 34-14
  - MAX 34-14
  - MIN 34-15
  - MONTH 34-15
  - NOW 34-15
  - TODAY 34-15
  - YEAR 34-15
- Dates 18-1
  - computation standards 18-4
  - continuous 18-12
  - levels 18-4
  - perfect pivoting 18-10
- Default Fields
  - Measure Names 11-11
  - Measure Values 11-11
  - number of records 11-11
- Default Map Source 24-25
- Default Member 13-54
- Default Number Format 27-19
- Defaults
  - colors 11-27
  - number format 11-28
  - shapes 11-27
  - sort order 11-29
- Deleting
  - annotations 28-15
  - bookmarks 31-8
- Dialog Box
  - Connect to Data Source 5-5
  - Open 5-4
- Dimensions 11-4
  - add to sheet 14-6
  - Columns shelf 13-10
  - Filters shelf 13-18
  - gray highlight 13-58
  - incompatible measures 13-58
  - Level of Detail shelf 13-21
  - Path shelf 13-49
  - Rows shelf 13-10
  - Shape shelf 13-39
  - Text shelf 13-47
- Dimensions and Measures 11-37
- Disaggregation 21-7
- Dividers
  - row and columns 27-12
- Double-Click
  - add to sheet 14-6
- Drill down

nonuniform 13-53  
Dual Axes 13-8  
Duplicate Connection 5-73  
Duplicate Sheet 3-8

## **E**

End  
    application 3-3  
Environment 3-1  
    workspace 3-6  
environment\_delete 3-20  
environment\_open 3-3  
environment\_organize 3-20  
environment\_repository 3-25  
Essbase 5-49  
Excel Data Source 5-11  
Excel workbook  
    mixed data types 11-35  
Export  
    data source connection 5-74  
    to print 32-3  
export 31-10  
    Access 31-16  
    image 31-21  
    measures 13-23  
Extract Data File  
    connecting 5-16  
Extracts  
    aggregate data 9-4  
    creating 9-3  
    defining 9-3  
    properties 9-7  
    refreshing 9-7  
    removing 9-7  
    roll up dates 9-4  
    using 9-7

## **F**

Field Defaults  
    formats 11-28  
    sort order 11-29  
Field Names  
    unknown 5-71  
Fields 3-7  
    formatting 27-14  
Filter  
    rows and columns 27-39  
    user filters 30-29  
filter  
    summation icon 16-57  
Filter Actions 20-3  
Filters  
    sets 17-26  
Filters shelf 13-18  
Firebird  
    connect to 5-22  
Fit Selector 3-9  
Fix Axes 3-10  
Fonts  
    formatting 27-6  
Formatting  
    aspect ratio 27-41  
    axis appearance 27-50  
    axis range 27-43  
    borders 27-12  
    cell size 27-37  
    centering text 27-9  
    clear from worksheet 3-19  
    column labels 27-41  
    default number format 27-41  
    defaults 27-4  
    dividers 27-12  
    fields 27-14  
    fonts 27-6

format window 27-4  
grid lines 27-12  
headers 27-41  
keyboard shortcuts 27-40  
lines 27-12  
mark halos 13-30  
maximum number 27-41  
panes 27-41  
reference lines 27-29  
row and column banding 27-10  
row and column order 27-39  
row and column size 27-33  
row labels 27-41  
shading 27-10  
table layout 27-33  
table size 27-35  
text 27-6  
text alignment 27-9  
tick marks 27-52  
worksheet 27-4  
zero lines 27-12

Functions 34-6  
  aggregate 34-20  
  converting data types 34-15  
  dates 34-13  
  logical 34-17  
  pass through 34-21  
  special 34-23  
  strings 34-10  
  table calculation 34-25  
  type conversion 34-15

## G

Gantt bar  
  mark type 12-25  
Gantt Bar Chart 33-63

  nonstandard 33-66  
  standard 33-63  
generations 13-54  
Geographic Roles 24-3  
Grand Totals 21-78  
  aggregations 21-80  
grand totals  
  sorting 17-8  
Group Members 3-9  
Grouping Measures 21-72

## H

Headers  
  creating in a table 13-10  
Heat Map 33-55  
  basic 33-56  
  overlapping data 12-18  
  size encoded 33-58  
Heat Maps  
  cell size  
    Cross Tabs  
      cell size 27-37  
hierarchy 11-6  
Highlight Actions 20-8, 20-13  
Highlighting  
  advanced 20-13  
  legends 20-9  
  selection 20-8  
Histogram 21-72  
Histograms 21-72  
history 19-8  
HTML Help  
  printing 32-10  
Hyperion Essbase 5-49  
  building views 13-54  
  set accounts dimension 5-49

- Hyperlinks
  - filter actions 20-3
- Hyperlinks. See Actions
  
- I**
- Images. See Background Images 23-1
- Import Data Source
  - from server 30-26
- import from workbook 5-6
- Invalid Fields 5-71
- ISO 8601 18-4
  
- J**
- Joins
  - adding tables 7-3
  - editing 7-8
  
- K**
- keyboard shortcut 37-1
- Keyboard shortcuts
  - table layout 27-40
  
- L**
- Layers 24-7
- Legends
  - sorting 17-16
- Level of Detail shelf 13-21
  - exporting measures 13-23
- levels 13-54
- Line
  - mark type 12-22
- Line Chart 33-30
  - filter and color encoded with multiple measures 33-35
  - filter and color encoded, nested table 33-33
  - filter and path encoded 33-37
- Lines
  - formatting 27-12
- Links. See Actions
- Literal expressions
  - booleans 34-5
  - dates 34-5
  - Null 34-5
  - numeric 34-5
  - strings 34-5
- Logical Functions 34-17
  - CASE 34-17
  - IF THEN ELSE END 34-18
  - IF THEN ELSEIF 34-19
  - IFNULL 34-19
  - IIF 34-18
  - ISDATE 34-19
  - ISNULL 34-19
  - MIN 34-19
  
- M**
- Map Options 24-7
- Map Source
  - set as default 24-25
  - WMS 24-23
- Map Sources 24-23
- Maps 24-1
  - cache 24-26
  - geographic roles 24-3
  - options 24-7
  - working offline 24-26
- Maps. See Background Images
- Mark Labels 28-16

- Mark Properties
    - colors 11-27
    - shapes 11-27
  - Marks 12-11
    - automatic 12-12
    - bar 12-16
    - changing size 12-32
    - circle 12-19
    - gantt bar 12-25
    - halos 13-30
    - line 12-22
    - polygon 12-24
    - shape 12-20
    - square 12-17
    - text 12-12
  - marks
    - stacking 12-27
  - Maximum 21-5
  - MDX 21-21
  - Measure Names 11-11, 11-32
    - editing 11-32
    - editing aliases 11-32
  - Measure names
    - removing 11-32
  - Measure Names and Measure Values 13-7
  - Measure Values 11-11, 11-32
  - Measures
    - add to sheet 14-6
    - caution symbol 13-58
    - Filters shelf 13-18
    - gray highlight 13-58
    - incompatible dimensions 13-58
    - setting as accounts dimension 5-49
  - measures 11-4
    - Columns shelf 13-10
    - Level of Detail shelf 13-23
    - Path shelf 13-49
    - Rows shelf 13-10
    - Shape shelf 13-39
    - Text shelf 13-47
  - Menu Commands
    - table layout 27-40
  - Message
    - show/hide 12-13
  - Microsoft Access database 5-9
  - Microsoft Analysis Services Cube
    - building views 13-58
  - Microsoft Analysis Services cube 5-33
  - Microsoft Excel workbook 5-11
  - Microsoft SQL Server database 5-38
  - Minimum 21-4
  - mixed data types 11-35
  - MSAS 5-33
  - Multidimensional
    - data sources 5-3
    - schema 11-6
  - multidimensional
    - hierarchy 11-6
    - slicer 16-58
  - Multiple Measures 11-11, 11-32
  - MySQL database 5-41
    - max\_allowed\_packet error 19-12
- N**
- New Sheet 3-8
  - new workbook 3-17
  - nonuniform drill down 13-53
  - Number Format 11-28, 27-41
    - default 27-19
  - Number functions
    - ABS 34-6
    - ACOS 34-6
    - ASIN 34-6

ATAN 34-7  
 ATAN2 34-7  
 COS 34-7  
 COT 34-7  
 DEGREES 34-7  
 EXP 34-7  
 LN 34-7  
 LOG 34-7  
 MAX 34-8  
 MIN 34-8  
 PI 34-8  
 POWER 34-8  
 RADIANS 34-8  
 ROUND 34-9  
 SIGN 34-9  
 SIN 34-9  
 SQRT 34-9  
 SQUARE 34-9  
 TAN 34-9  
 ZN 34-10  
 Number of Records 11-11  
 Numbers  
     default format 27-41  
 Numerical data type 11-34

## O

OData Data Source 5-14  
 Open  
     application 3-3  
 Open dialog box 5-4  
 Open Screen 3-4  
 Open Workbook  
     from server 30-23  
 Opening HVE 4-3  
 Opening Tableau 4-3  
 Operators 34-37

- (subtraction) 34-37  
 % (modulo) 34-37  
 \* (multiplication) 34-37  
 + (addition) 34-37  
 / (division) 34-37  
 AND 34-38  
     comparisons 34-38  
 NOT 34-39  
 OR 34-38  
 Oracle Database  
     configuring 5-46  
 Order of Operations  
     See Precedence  
 overlapping text 12-13  
 Overlay 24-7

## P

Pages Shelf 13-14  
 Pan 3-10  
 panning 19-7  
 Parameters 22-1  
 Pass Through Functions 34-21  
     RAWSQL\_BOOL 34-21  
     RAWSQL\_DATE 34-21  
     RAWSQL\_DATETIME 34-22  
     RAWSQL\_INT 34-22  
     RAWSQL\_REAL 34-22  
     RAWSQL\_STR 34-22  
     RAWSQLAGG\_BOOL 34-22  
     RAWSQLAGG\_DATE 34-23  
     RAWSQLAGG\_DATETIME 34-23  
     RAWSQLAGG\_INT 34-23  
     RAWSQLAGG\_REAL 34-23  
     RAWSQLAGG\_STR 34-23  
 Path shelf 13-49  
 PDF

- print to 32-8
  - publish 32-8
  - Percentages 21-87
    - about 21-87
    - aggregations 21-88
    - column 21-93
    - column in pane 21-97
    - options 21-91
    - pane 21-94
    - row 21-93
    - row in pane 21-95
    - table 21-91
  - perfect pivoting
    - date 18-10
  - Performance
    - all data sources 35-3
    - context filters 35-8
    - large text files 35-9
    - multidimensional data sources 35-7
    - relational data sources 35-5
    - tips 35-1
  - PivotTable 33-22
  - Polygon mark type 12-24
  - PostgreSQL
    - connect to 5-54
  - Precedence 34-39
  - Precision
    - warning 10-6
  - Presentation Mode 3-9
  - Printing 32-3, 32-6
    - HTML Help 32-10
    - page orientation 32-4
    - print range 32-6
    - show selections 32-6
    - to PDF 32-8
  - Publish as PDF 32-8
  - Publish to Server 30-1
  - Publishing Data Sources 30-19
  - Publishing Workbooks
    - hide sheets 30-15
    - passwords 30-16
    - permissions 30-9
    - server 30-3
    - user filters 30-29
- Q**
- qualified member names 17-32
  - Queries
    - abandoned 10-5
    - data source unavailable 10-5
    - precision warnings 10-6
  - Quit 3-3
- R**
- ragged hierarchy 13-53
  - RAWSQL 34-21
  - Redo 3-8
  - redo 19-8
  - Reference Lines
    - formatting 27-29
  - refresh data source 5-76
  - Relational
    - data sources 5-3
    - schema 11-6
  - Removing
    - bookmarks 31-8
  - Rename Connection 5-72
  - rename sheet 3-19
  - Re-Ordering
    - rows and columns 27-39
  - Replace References 5-71
  - Resizing

- cells 27-37
    - entire table 27-35
    - rows and columns 27-33
  - Roll up dates 9-4
  - Row attributes 27-41
  - Row Labels
    - formatting 27-41
    - maximum number 27-41
  - Rows and Columns
    - changing order 27-39
    - include/exclude 27-39
    - re-ordering 27-39
    - show/hide 27-39
  - Rows shelf 13-10
  - Run Update 3-8
  - Running Actions 20-20
  - Running Totals 21-49
- S**
- Save 3-8
  - save workbook 31-3
  - Scatter Plot 33-39
    - basic 33-40
    - filter and color encoded 33-42
    - filter and level of detail 33-45
    - filter, level of detail, and color encoded 33-47
    - filter, level of detail, color and shape encoded 33-52
  - Schema
    - add fields 11-17
    - find 11-14
    - hide/unhide fields 11-16
    - multidimensional 11-6
    - organize 11-13
    - relational 11-6
    - rename fields 11-15
  - schema
    - hiding 11-7
  - Secondary Axis 13-8
  - Select 3-9
  - Selecting Marks
    - highlighting 20-8
  - selecting marks 19-3
  - Server 30-1
    - publishing workbooks 30-3
  - Sets 17-26
    - about 17-26
    - creating 17-28
  - Shading
    - formatting 27-10
  - shape mark type 12-20
  - Shape shelf 13-38
  - shared members 13-55
  - Sheets
    - renaming 3-19
  - shelf
    - Color 13-23
    - Columns 13-10
    - Level of Detail 13-21
    - Path 13-49
    - Rows 13-10
    - Shape 13-38
    - Size 13-31
    - Text 13-47
  - Shelves 3-11
    - pages 13-14
    - types 13-10
  - Shortcut
    - data source connection 5-74
  - shortcut
    - keyboard 37-1
  - Show Mark Labels 3-9
  - Show/Hide
    - data labels 28-16

- parts of dashboard 29-10
  - rows and columns 27-39
  - Size
    - marks 12-32
  - Size shelf 13-31
  - slicer 16-58
  - Sort Ascending 3-9
  - Sort Descending 3-9
  - Sort Order 11-29
  - Sorting
    - manually 27-39
  - sorting 17-3
    - grand totals 17-8
  - Special Functions 34-23
    - CURRENTUSER 34-24, 34-25
    - ISCURRENTUSER 34-24
    - ISMEMBEROF 34-24
  - SQL
    - custom query 5-68
    - Server database 5-38
    - typing in 5-68
  - Square Cell 27-37
  - Square mark type 12-17
  - stacking marks 12-27
  - Standard Deviation 21-5
    - population 21-5
  - Start
    - application 3-3
    - page 3-4
  - Status Bar 3-10
    - precision warning 10-6
  - String Functions 34-10
    - ASCII 34-10
    - CHAR 34-10
    - CONTAINS 34-10
    - ENDSWITH 34-10
    - FIND 34-10
    - LEFT 34-11
    - LEN 34-11
    - LOWER 34-11
    - LTRIM 34-11
    - MAX 34-11
    - MID 34-11
    - MIN 34-12
    - RIGHT 34-12
    - RTRIM 34-12
    - SPACE 34-12
    - STARTSWITH 34-12
    - TRIM 34-12
    - UPPER 34-12
  - Subtotals 21-84
    - local vs. server 21-77
  - Sum 21-4
  - Summary Card 19-11
  - summation icon 16-57
  - Supported Data Sources 5-3
  - Symbols
    - caution 10-6
- T**
- Table
    - components 12-2
  - table
    - axes creation 13-10
    - header creation 13-10
  - Table Calculation Functions 34-25
    - FIRST 34-25
    - INDEX 34-26
    - LAST 34-26
    - LOOKUP 34-27
    - PREVIOUS\_VALUE 34-28
    - RUNNING\_AVG 34-28
    - RUNNING\_COUNT 34-28

- RUNNING\_MAX 34-29
- RUNNING\_MIN 34-29
- RUNNING\_SUM 34-30
- SIZE 34-31
- WINDOW\_AVG 34-32
- WINDOW\_COUNT 34-32
- WINDOW\_MAX 34-32
- WINDOW\_MIN 34-33
- WINDOW\_STDEV 34-34
- WINDOW\_STDEVP 34-34
- WINDOW\_SUM 34-35
- WINDOW\_VAR 34-35
- WINDOW\_VARP 34-35
- Table Calculations 21-40
  - compute using 21-69
  - cumulative total 21-63
  - customize 21-68
  - defining 21-50
  - difference from 21-52
  - moving average 21-49, 21-64
  - moving calculation 21-64
  - multiple calculations 21-67
  - percent difference from 21-54
  - percent from 21-57
  - percent of total 21-60
  - running total 21-49, 21-63
  - secondary calculation 21-67
- Table Layout
  - formatting 27-33
  - keyboard shortcuts 27-40
  - menu commands 27-40
- Table Options 27-41
- Tableau Server 30-1
- Tables
  - editing 7-8
- Tactician Maps 24-1
- Text
  - formatting 27-6
- text
  - mark type 12-12
  - overlapping 12-13
- Text Cell 27-38
- Text data type 11-34
- Text file 5-17
- Text shelf 13-47
- Text Table 33-22
  - basic 33-22
  - filter and color encoded 33-26
- Tick Marks 27-52
- Time 18-1
- Tips and Tricks 36-1
- Toolbar 3-8
  - automatic updates 3-8
  - clear sheet 3-8
  - connect to data 3-8
  - duplicate sheet 3-8
  - fit selector 3-9
  - fix axes 3-10
  - group 3-9
  - highlight 3-10
  - new sheet 3-8
  - pan 3-10
  - presentation mode 3-9
  - redo 3-8
  - run update 3-8
  - save 3-8
  - select 3-9
  - show mark labels 3-9
  - Show Me! 3-9
  - sort ascending 3-9
  - sort descending 3-9
  - swap 3-9
  - undo 3-8
  - view cards 3-9

zoom 3-9  
ToolTips 4-14  
Totals  
  grand totals 21-78  
  local vs. server 21-77  
  subtotals 21-84  
Type Conversion 34-15  
  DATE 34-16  
  DATETIME 34-16  
  FLOAT 34-16  
  INT 34-16  
  STR 34-16

## U

Undo 3-8  
undo 19-8  
Updates  
  automatic 10-3  
User Filters 30-29  
  creating 30-29  
  editing 30-35  
  example 30-38  
  previewing as user 30-34  
  publishing 30-37  
  special functions 30-45  
  using data source 30-45

## V

Variance 21-5  
  population 21-5  
View Cards 3-9  
View Statistics 19-11  
viewing underlying data 19-12

## W

Warnings  
  precision 10-6  
  show/hide 12-13  
Welcome Screen 3-4  
Windows Azure Marketplace DataMarket 5-20  
WMS Servers 24-23  
workbooks  
  importing from 5-6  
  saving 31-3  
Working Offline  
  maps 24-26  
Worksheets  
  duplicating as cross-tab 3-19  
worksheets  
  clearing 3-19  
Workspace 3-6  
  cards 3-11  
  data window 3-7  
  shelves 3-11  
  status bar 3-10  
  toolbar 3-8

## Z

Zoom 3-9  
Zoom mode 19-5  
Zoom tool 19-5