Tableau Desktop Help

This PDF contains the complete documentation for Tableau Desktop. 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 regularly.

Version 8.3; Last Updated in 2014
Copyright © 2014 Tableau Software, Incorporated and its licensors. All rights reserved.
This product is Client Software as defined in Tableau Software's End User Software License Agreement.

Welcome to Tableau Desktop

This online Help system contains the complete product documentation for using Tableau Desktop on Windows. Features such as full text search make it easy for you to find what you are looking for. The online Help is divided into sections, which are displayed as books in the Contents tab on the left.

Getting Started on page 4 - learn the basic concepts in Tableau and complete a step-by-step tutorial.

Connect to Data on page 52 - connect to data, join tables, extract data, and manage queries.

Building Data Views on page 232 - drag and drop fields to create views of your data, then refine the view.

Advanced Analysis on page 543 - create custom fields, use the built in statistical tools, add maps, and more.

Publish and Share on page 732 - format and annotate the view, then export or publish it to share with others.

Reference on page 880 - understand Tableau functions, shortcuts, tips and tricks, the glossary, and more.

Build-It-Yourself Exercises on page 484 - browse examples to learn how to build specific types of views.

Copyright Notices

For more information about product Documentation,
see our End User Licensing Agreement.

Version: 8.3 | Updated on 11/29/2014

Getting Started

Welcome to Tableau Desktop. Learn more about the product and what it can do. Then explore the Tableau workspace to get familiar with the environment. Finally, follow a step-by-step tutorial that guides you through connecting to data and building your first view.

The Tableau Environment

This section introduces 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.

Opening and Closing the Application

The first thing to understand is how to open and close the application.

Open Tableau

Open Tableau Desktop by doing any 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.
- 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. Close the application by doing one of the following:

- Click the **Close** icon located in the right corner of the application title bar.

- Select **File > Exit**.

If your workbook has not been saved, you will be asked whether you want to save it.

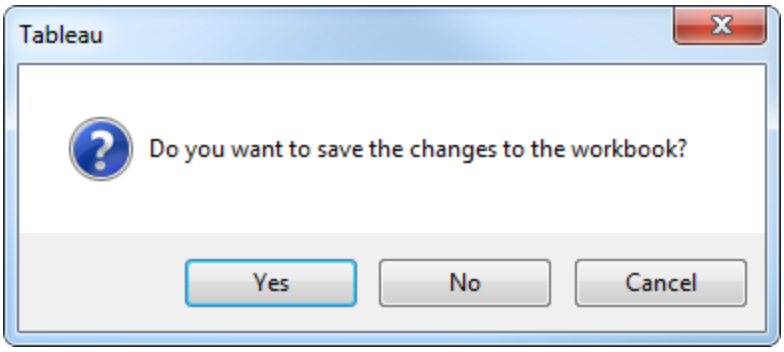

Start Page

When you first open Tableau, 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.

The Start Page displays when you open Tableau without specifying a workbook.

When Tableau is already open, you can navigate to the start page by clicking the Home icon in the top-right corner of the workbook:

The start page is split into the following sections:

- **Data** – The data section lets you either connect to a new data source or quickly open a saved connection. By default 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 a save new workbooks, the nine most recently opened are available on the start page.

  You can pin workbooks to the start page by clicking the pin icon that shows in the lower left corner when you hover the mouse over the workbook image. Pinned workbooks will always show on the start page even if they weren't opened recently. Remove the workbook by clicking the red "x" icon. The workbook is removed immediately but will be shown again with your most recently used workbooks the next time you open it.

Pin a Workbook **Remove a Workbook**

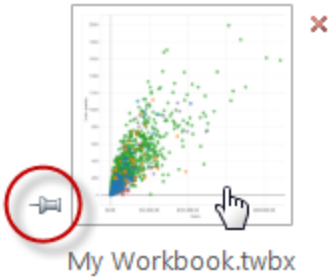

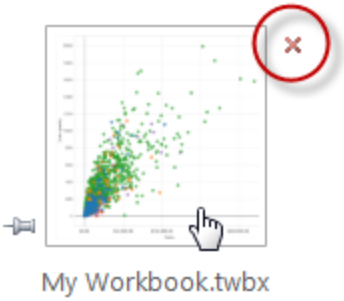

- **Getting Started** – The getting started section contains links to resources that can help you learn how to use Tableau. You can watch training videos and learn about other Tableau products.
- **Samples** – The samples section contains several sample workbooks that show off the types of views you can create with Tableau Desktop. 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.

The view below shows the Tableau workspace for a view. The workspace for creating dashboards or stories is quite a bit different. See **Adding Views to a Dashboard** on page 785 for information on the dashboard workspace, or **The Story Workspace** on page 812.

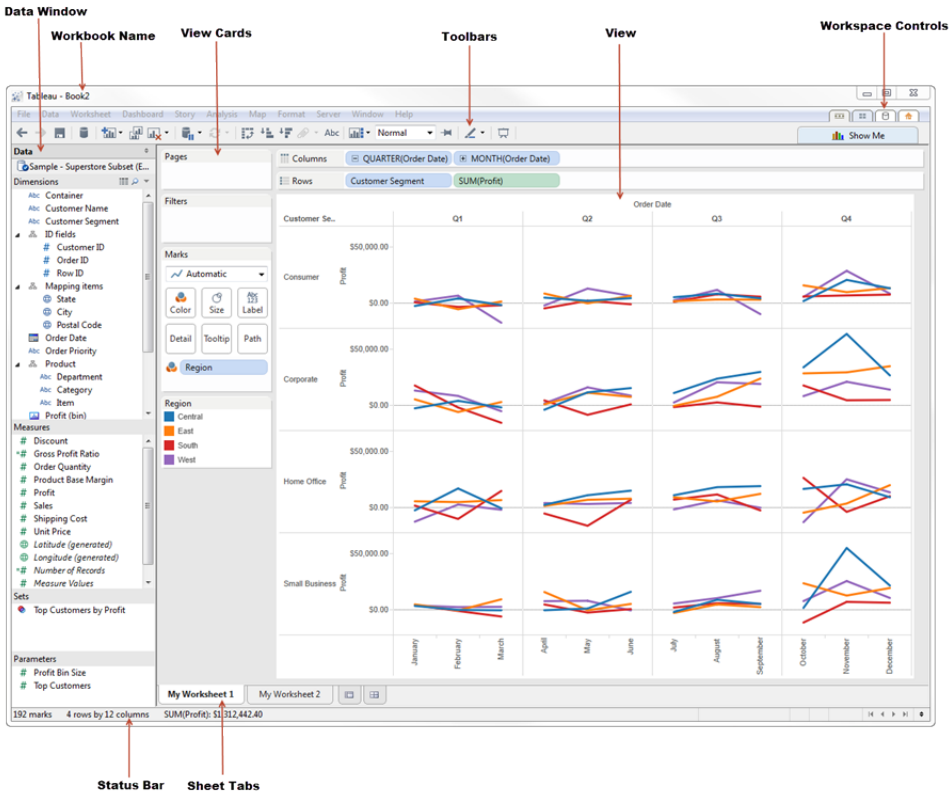

Data Window

Data fields appear on the left side of the workspace in the Data window. You can hide and show the Data window by clicking the minimize button in the upper right corner of the Data window. The Data window collapses to the bottom of the workbook. Click the minimize button again to show the Data window.

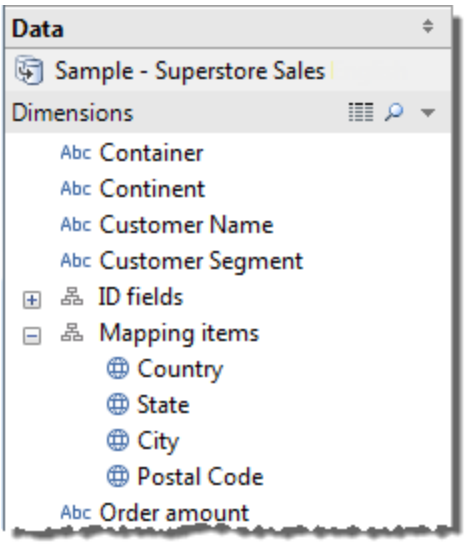

Current data sources are listed at the top of the Data window. Right-click a data source to view a context menu for the data source. See **Connecting to Data and Editing Data Sources** on page 52.

Below the data sources the various fields available in the currently selected data source are listed. You can search for fields in the Data window by clicking the magnifying class icon and then typing in the text box. Right-click the fields in the Data window to access important commands.

Click the **View Data** icon at the top of the Data window to see the underlying data. See **View Data** on page 479 to learn more about viewing the underlying data in a workbook.

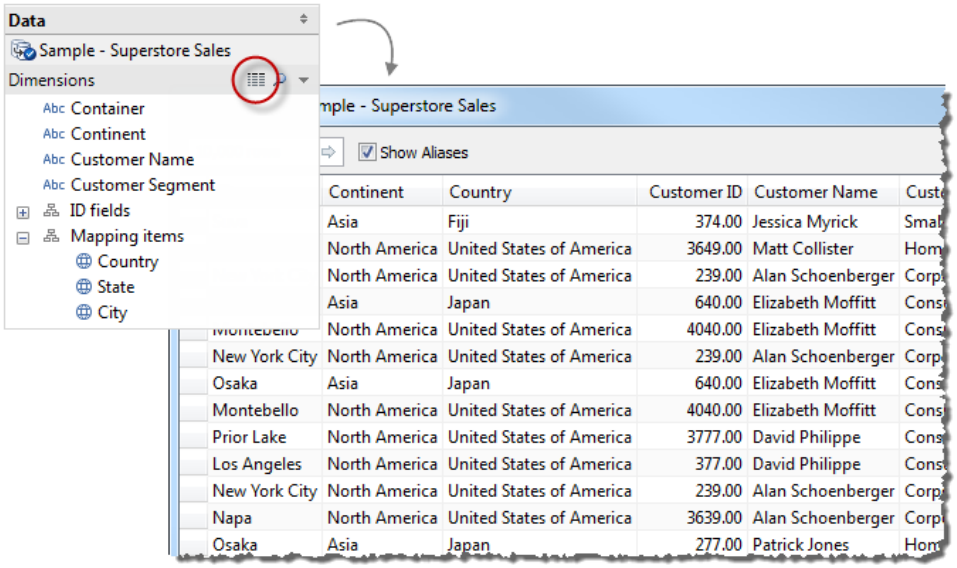

The Data window is separated into several areas:

- **Dimensions** - fields that contain category data such as text and dates. See **Data Roles: Dimension vs. Measure** on page 222.
- **Measures** - fields that contain numbers that can be aggregated. See **Data Roles: Dimension vs. Measure** on page 222.
- **Sets** - subsets of data that you define. See **Sets** on page 416.
- **Parameters** - dynamic placeholders that can replace constant values in calculated fields and filters. See **Parameters** on page 663.

Toolbar

Tableau's toolbar contains commands such as Connect to data, New Worksheet, and Save. In addition, the toolbar contains analysis and navigation tools such as Sort, Group, and Highlight. You can hide or display the toolbar by selecting **Window > Show Toolbar**.

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

Toolbar Button	Description
	Undo: reverses the most recent action in the workbook. You can undo an unlimited number of times, back to the last time you opened the workbook, even after you have saved. See **Undo and Redo** on page 476.
	Redo: repeats the last action you reversed with the Undo button.You can redo an unlimited number of times.
	Save: saves the changes made to the workbook. See **Saving Your Work** on page 860.
	Connect to Data: opens the Connect to Data page where you can create a new connection or open a saved connection from your repository. See **Connect to Data** on page 52.
	New Worksheet: creates a new blank worksheet. Use the drop-down menu to create a new worksheet, dashboard, or story. See **Creating New Sheets** on page 18.
	Duplicate Sheet: creates a new worksheet containing the exact same view as the current sheet. See **Duplicating Sheets** on page 20.
	Clear: clears the current worksheet. Use the drop-down menu to clear specific parts of the view such as filters, formatting, sizing, and axis

Toolbar Button	Description
	ranges.
	Automatic Updates: controls whether Tableau automatically updates the view when changes are made. Use the drop-down list to automatically update the entire sheet or just use quick filters. See **Automatic Updates** on page 191.
	Run Update: runs a manual query of the data to update the view with changes when automatic updates are turned off. Use the drop-down menu to update the entire sheet or just use quick filters.
	Swap: moves the fields on the Rows shelf to the Columns shelf and vice versa. The Hide Empty Rows and hide Empty Columns settings are always swapped with this button.
	Sort Ascending: applies a sort in ascending order of a selected field based on the measures in the view. See **Sorting** on page 389.
	Sort Descending: applies a sort in descending order of a selected field based on the measures in the view. See **Sorting** on page 389.
	Group Members: creates a group by combining selected values. When multiple dimensions are selected, use the drop-down menu to specify whether to group on a specific dimension or across all dimensions. See **Groups** on page 401.
Abc	**Show Mark Labels**: switches between showing and hiding mark labels for the current sheet. See **Mark Labels** on page 781.
	Presentation Mode: switches between showing and hiding everything except the view (i.e., shelves, toolbar, Data window). See **Reorganizing the Workspace** on page 15.
	View Cards: shows and hides specific cards in a worksheet. Select each card that you want to hide or show on the drop-down menu.
Normal	**Fit Selector**: specifies how the view should be sized within the application window. Select Normal fit, Fit Width, Fit Height, or Entire View.
	Fix Axes: switched between a locked axis that

Toolbar Button	Description
	only shows a specific range and a dynamic axis that adjusts the range based on the minimum and maximum values in the view. See **Editing Axes** on page 760.
	Highlight: turns on highlighting for the selected sheet. Use the options on the drop-down menu to define how values will be highlighted. See **High-light Toolbar Button** on page 550.
Show Me	**Show Me**: displays alternative ways to look at the data. The types of views available are dependent on the fields already in the view as well as any selections in the Data window. See **Show Me** on page 322

Tooltips

Tooltips are additional data details that display when you rest the pointer over one or more marks in the view. Tooltips also offer convenient tools to quickly filter marks, view underlying data, group data, and create sets. Tooltips consist of a body, action links, and commands.

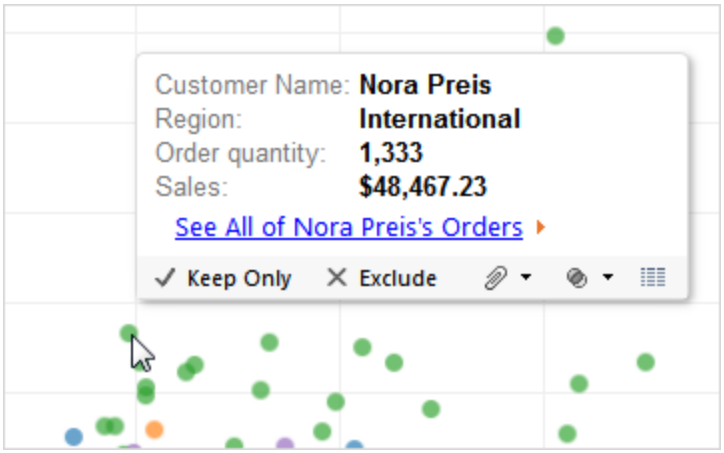

Body

The body of a tooltip contains details about a specific mark or a selection of multiple marks. For example, in a bar chart showing sales by region, the tooltip body may include the actual sales amount and the region name. The default tooltip is based on the fields used in the view. You can customize what is shown in the tooltip by dragging fields to Tooltip on the marks card. Click **Tooltip** on the marks card to further customize the tooltip including how the text is formatted. Alternatively, you can select **Worksheet > Tooltip**. See **Tooltips** on page 748 to learn more about formatting the body of the tooltips.

Action Links

If the sheet has any actions, the action links will be listed below the body of the tooltip. An action adds context and interactivity to your data through filters, highlighting, and links to external

resources. See **Actions** on page 543 to learn more about adding actions to your workbook.

Commands

The bottom of the tooltip lists commands for quickly filtering data and viewing the underlying data. For example, you can use the tooltip to quickly remove an outlier in a scatter plot. Each of the commands are described below.

- **Keep Only** - creates a filter that removes all other data. See **Selecting Data to Filter** on page 334 to learn more.
- **Exclude** - creates a filter that removes the selected data. See **Selecting Data to Filter** on page 334 to learn more.
- **Group Members** - creates a group based on the selection. If the selection contains multiple dimensions, you can group on one dimension or all dimensions. See **Groups** on page 401 to learn more.
- **Create Set** - creates a new set containing the selected members. You can create a new set or add members to an existing set. See **Sets** on page 416 to learn more.
- **View Data** - opens a window displaying the data. You can view the summarized data or the underlying data.See **View Data** on page 479 to learn more.

These commands are visible by default. You can disable the commands in the Edit Tooltip dialog box. Click Tooltip on the Marks card or select **Worksheet > Tooltip** and then clear the **Include command buttons** checkbox in the bottom left corner. Tooltip settings apply to the active worksheet and can be different for each worksheet in the workbook.

Status Bar

The status bar is located at the bottom of the Tableau workbook. It displays descriptions of menu items as well as information about the current view. For example, the status bar below shows that the view has 1864 marks shown in 1 row and 1 column. It also shows that the SUM (Sales) for all the marks is $30,434,456.08.

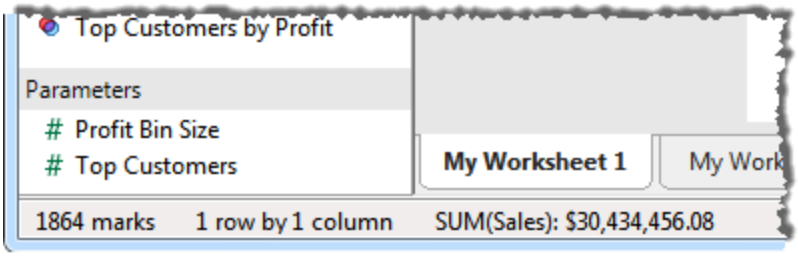

You can hide the status bar by selecting **Window** > **Show Status Bar**.

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

Warning Icon	Description
1	**Cancel Query Indicator**: When you cancel multiple queries, an indicator appears to show you how many queries are still

Warning Icon	Description
	running on the database and using resources. See **Abandoned Queries** on page 192 for more information about managing these queries.
	Precision Warning: 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. See **Precision Warnings** on page 193 for more information about this warning.
1 unknown	**Special Values Indicator**: If your data contains null values, unknown geographic locations, or negative or zero values on a logarithmic axis; the values are shown with an indicator in the lower right corner of the view. Click the indicator for options for handling these values. See **Special Values** on page 229 to learn more about this indicator and how to handle these values.

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 is where you control mark properties. It contains the mark type selector along with controls for Color, Size, Label, Detail, Tooltip, Shape, and Angle. The controls that are available depend on the mark type.

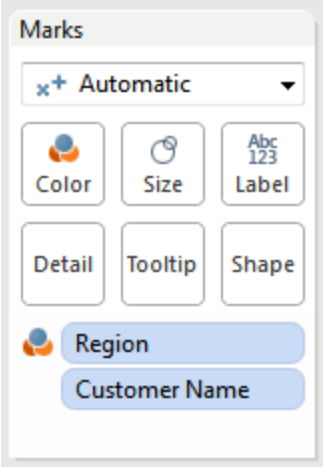

Cards can be shown and hidden as well as rearranged around the worksheet.

The following list describes each card and its contents.

- **Columns Shelf** – contains the Columns shelf where you can drag fields to add columns to the view. See **Columns and Rows Shelves** on page 298.
- **Rows Shelf** – contains the Rows shelf where you can drag fields to add rows to the

view. See **Columns and Rows Shelves** on page 298.

- **Pages Shelf** – contains the Pages shelf where you can split a view into several pages based on the members of a dimension or the values in a measure. See **Pages Shelf** on page 301.
- **Filters Shelf**– contains the Filters shelf; use this shelf to specify the values to include in the view. See **Filters Shelf** on page 306.
- **Measure Values Shelf** – contains the Measure Values shelf; use this shelf to blend multiple measures on a single axis. This shelf is only available when there is a blended axis in the view. See **Blended Axes** on page 328.
- **Color Legend** – contains the legend for the colors in the view and is only available when there is at least one field on Color. See **Color** on page 257.
- **Shape Legend** – contains the legend for the shapes in the view and is only available when there is at least one field on Shape. See **Shape** on page 274.
- **Size Legend** – contains the legend for the size of the marks in the view and is only available when there is at least one field on Size. See **Size** on page 265.
- **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 filter in the view. Use these cards to easily include and exclude values from the view. See **Using Quick Filters** on page 369.
- **Parameters** – a separate parameter card is available for every parameter in the workbook. The parameter card contains controls for changing the parameter values. See **Parameter Controls** on page 672.
- **Marks** – controls the mark properties in the view. There is a mark type selector where you can specify the mark type (e.g., bar, line, area, etc.). In addition, the Marks card contains controls for Color, Size, Label, Text, Detail, Tooltip, Shape, Path, and Angle. The availability of these controls is dependent on the fields in the view and the mark type.
- **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, Median, Sum, and Average.
- **Current Page** – contains the playback controls for the Pages shelf and indicates the current page that is displayed. This card is only available when there is at least one 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.

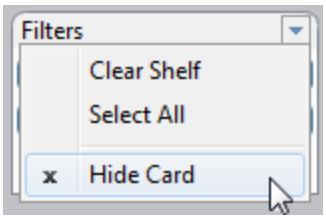

Reorganizing the Workspace

Every worksheet contains a variety of cards, shelves, legends, and so on. You can reorganize the workspace by rearranging cards, hiding and showing specific parts of the workspace, and hiding everything but the view using Presentation Mode.

Rearranging Cards

A worksheet contains several cards that contain shelves, legends, and other controls. 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.

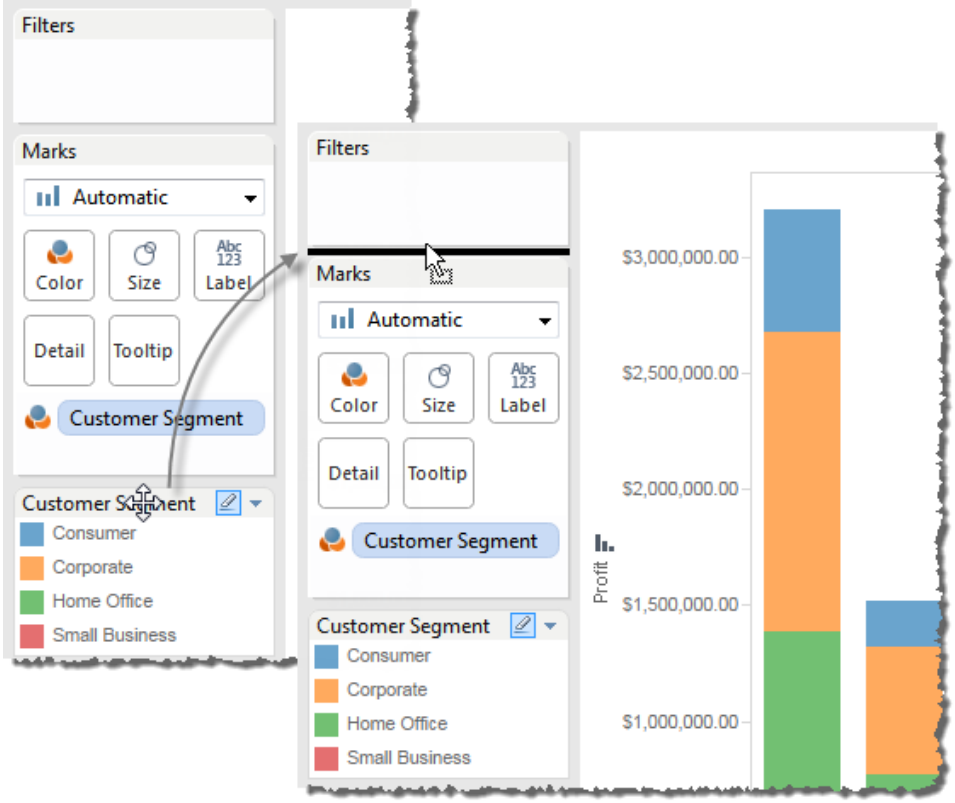

You can restore the worksheet windows to their default state by selecting **Reset Cards** on the Show/Hide Cards toolbar control.

Showing and Hiding Parts of the Workspace

Just about everything in the workspace can be turned on and off so you can avoid cluttering the worksheet with unnecessary cards, shelves, etc.

- To show and hide the toolbar or status bar, select **Window** and then select what you want to hide.
- To show and hide the window on the left side, which is either the Data window (for worksheets), the Dashboard window (for dashboards), or the Dashboards and Sheets window (for stories); click the Minimize button in the upper right corner of the window. The window is minimized to the bottom left corner of the workbook. Click the same button again to restore the window.
- To show or hide a card click **Show/Hide Cards** on the toolbar and then select the card you want to show or hide.

You can restore the worksheet windows to their default state by selecting **Reset Cards** on the **Show/Hide Cards** toolbar control.

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, quick filters, parameter controls, and worksheet tabs.

- To switch in and out of Presentation Mode, click the Presentation Mode button on the toolbar or select **Window > Presentation Mode**.

Use the presentation mode controls in the bottom right corner to move between sheets and more. Each presentation mode control is described below.

	Show/Hide Filmstrip - switches between showing the sheet tabs as thumbnails and as tabs.
	Previous/Next Sheet - advances forward or backward through the sheets in a workbook.
	Enter/Exit Full Screen - switches between expanding the workbook to fill the entire screen or showing it in a window.
	Exit Presentation Mode - returns the workbook to showing the entire workspace including the menus, toolbar, Data window, and so on.

Language and Locale

Tableau Desktop is localized into several languages.

When you first run Tableau, it recognizes your computer locale and uses the appropriate language if it is supported. If you are using an unsupported language, the application defaults to English.

You can configure Tableau to display the user interface (menus, messages, etc.) by choosing **Help** > **Choose Language**. After you change this setting, you'll need to restart the application for the changes to take effect. You do not need to change this setting for every workbook.

To configure date and number formatting, choose **File** > **Workbook Locale**. By default, the locale is set to **Automatic**, which means the locale will match the locale when the workbook is opened. This can be useful if you are authoring a workbook that will be viewed in many different languages and you want the dates and numbers to update accordingly. When you select a specific locale, the workbook will not change regardless of who opens it.

Tableau checks the following, in order, to determine the workbook locale:

- Workbook Locale (explicit setting)
- Windows Locale
- Tableau Language

If none of the above is set, then the workbook locale defaults to English.

Workbooks and Sheets

Tableau uses a workbook and sheet file structure, much like Microsoft Excel. A workbook contains sheets, which can either be a worksheet or a dashboard. A worksheet contains a single view along with shelves, legends, and the Data window. A dashboard is a collection of views from multiple worksheets. Within a workbook, you can create new sheets, clear an entire worksheet, duplicate sheets, hide or show a worksheet, and delete a sheet. Tableau has several ways to view and organize the sheets in your workbook.

Workbooks

Tableau workbook files are much like Microsoft Excel workbooks. They contain one or more sheets, which can be worksheets, dashboards, or stories. They allow you to organize, save, and share your results. The workbook is the container for all of your work in Tableau.

When you open Tableau, a blank workbook is automatically created. You can also create a new workbook by selecting **File > New** or by pressing Ctrl + N on your keyboard. You can open an existing workbook by doing one of the following:

- Click the thumbnail image of the workbook on the start page. The start page shows workbooks that you've recently opened.
- Select **File > Open** and navigate to the location of your workbook using the Open dialog box. Tableau workbooks have a .twb or .twbx file extension.
- Double-click on any workbook file in the Windows explorer.
- Drag any workbook file onto the Tableau desktop icon or onto the running application.

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

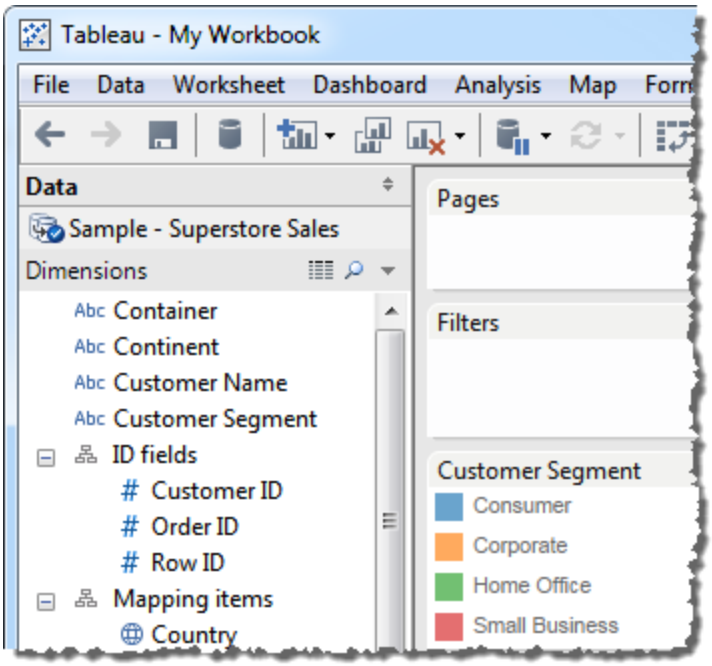

You can open multiple workbooks simultaneously. Each workbook is shown in its own application window.

Sheets

Each workbook can contain different types of sheets: views (also known as worksheets), dashboards, and stories. 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 that you can arrange for presentation or to monitor. A story is a sequence of views or dashboards that work together to convey information. The sheets display along the bottom of the workbook as tabs. In this section you'll learn how to create, open, duplicate, hide, and delete sheets. You'll also learn how to organize sheets in a workbook.

Creating New Sheets

There are several ways to create new sheets in a workbook. Follow the steps below to create a worksheet, a dashboard, or a story. While you can create as many sheets in a workbook as you like, it is best to create a new workbook for conceptually different analyses.

Creating Worksheets

Create a new worksheet by doing one of the following:

- Selecting **Worksheet > New Worksheet**
- Click the **New Worksheet** tab at the bottom of the workbook.

- Click the New Worksheet

icon on the toolbar and then choose **New Worksheet**.

- Press Ctrl + M on your keyboard.

Creating Dashboards

Create a new dashboard by doing one of the following:

- Select **Dashboard > New Dashboard**.
- Click the New Dashboard tab at the bottom of the workbook.

- Click the New Worksheet icon on the toolbar and then choose **New Dashboard**.

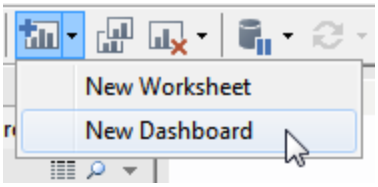

Creating Stories

Create a new story by doing one of the following:

- Select **Story** > **New Story**.
- Right-click the sheet tab and choose **New Story**.
- Click the New Worksheet icon on the toolbar and then choose **New Story**:

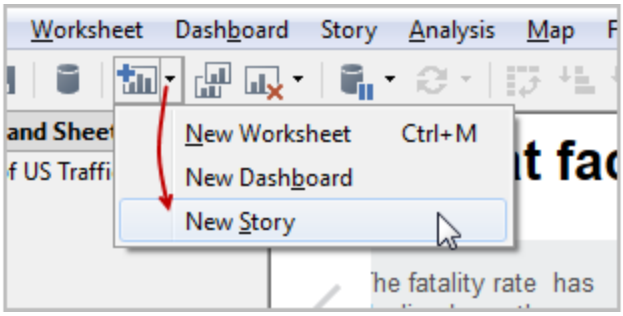

Tableau automatically generates sheet names. The first worksheet is named Sheet 1, the second worksheet is named Sheet 2, and so on. The same is true for dashboards (Dashboard 1, Dashboard 2, etc.) and stories (Story 1, Story 2, etc.) To rename a sheet, right-click the sheet tab and select **Rename Sheet**. 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 steps you have performed on the worksheets, dashboards, and stories in that workbook. To move backward through the history click **Undo** on the toolbar or press Ctrl + Z on your keyboard. Similarly, move forward through the history by clicking **Redo** on the toolbar or by pressing Ctrl + Y on your keyboard.

You can remove all fields, formatting, sizing, axis ranges, filters, sorts, and context filters 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, sorts, etc.

> 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, dashboard, or story. You can then modify the sheet without losing the original version. To duplicate the active sheet, right-click the sheet tab and select **Duplicate Sheet**.

A crosstab (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 the data view.

To quickly create a cross-tab from a view, do the following: right-click the sheet tab and select **Duplicate as Crosstab**. You can also select **Worksheet > Duplicate as Crosstab**. 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. (Dashboards and stories cannot be duplicated as crosstabs.)

There are other ways to see the numbers behind the data views. For example, you can mouse-over any mark to display the associated numbers in a tooltip. Click the **View Data** command at the bottom of the tooltip to view underlying data. You can copy and paste the data into Excel or another application.

Hiding and Showing Worksheets

A worksheet that is used in a dashboard or story cannot be deleted, but it can be hidden. You may want to hide a worksheet if you are sharing the dashboard or story with others and don't want to clutter the workbook with all the supporting worksheets. To hide a worksheet, right-click the worksheet tab and select **Hide Sheet**.

Dashboards

You can show a hidden sheet from a dashboard that uses it by choosing **Go to Sheet** on the dashboard view menu. The hidden sheet is shown until you switch to another sheet. While the hidden sheet is showing, you can right-click the sheet tab and select **Unhide** to unhide it permanently.

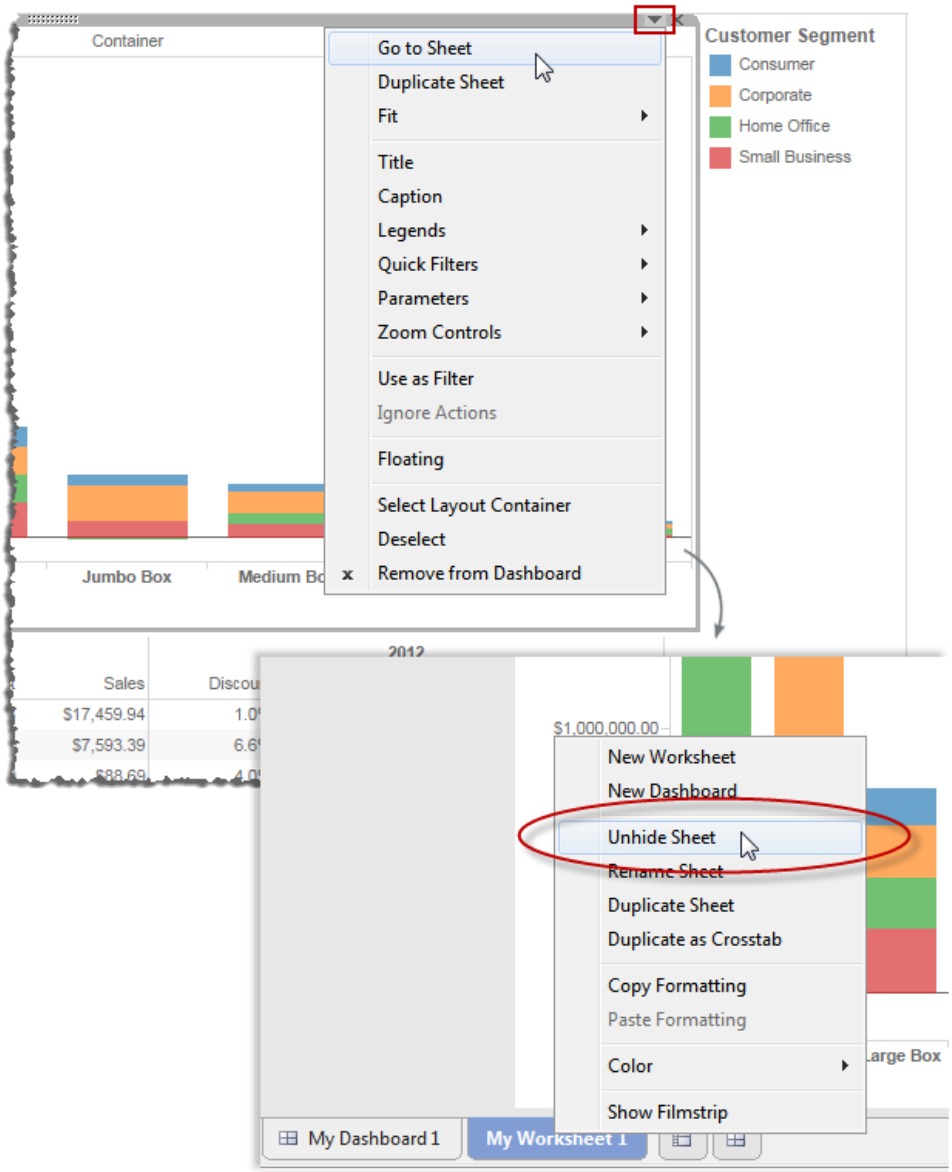

Stories

As you create stories, you can see the workbook's dashboards and views in the Dashboards and Worksheets area. To view a hidden dashboard and worksheet, hover over the icon to the right of the sheet name and click to go to the underlying sheet.

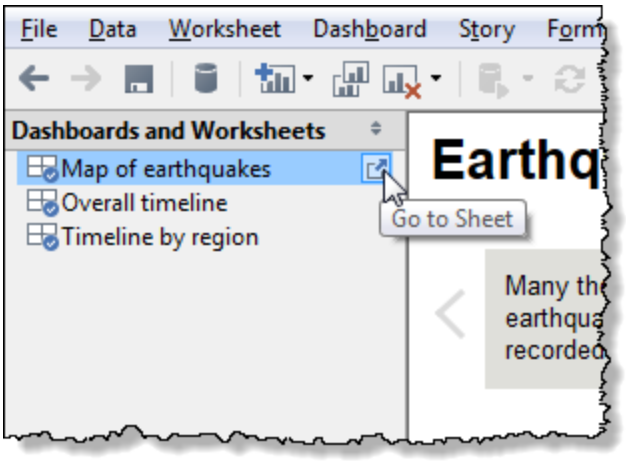

Though you can now see the sheet, it is still officially hidden, which is indicated by the blue background:

The sheet will not show if you publish the workbook, or after you close and reopen it. To unhide the sheet, right-click it and choose **Unhide Sheet**.

If a story uses a hidden dashboard which in turn uses a hidden worksheet, to see that hidden worksheet you must first go to the hidden dashboard from the story, and then go to the hidden worksheet from the dashboard.

Deleting Sheets

Deleting a sheet removes it from the workbook. To delete the active sheet, right-click the sheet in the sheet tab along the bottom of the workbook and select **Delete Sheet**. Worksheets used in a dashboard or story cannot be deleted, but they can be hidden. See **Hiding and Showing Worksheets** on page 20.

> There must always be at least one worksheet 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 (view, dashboard, or story) is represented as a tab along the bottom of the workbook. Select any tab to open the corresponding worksheet.

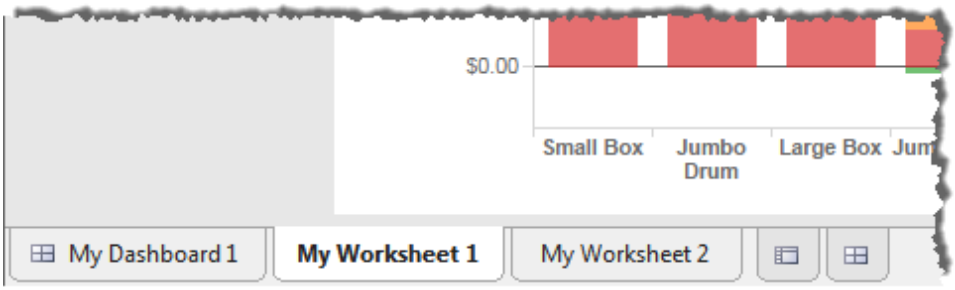

In the bottom right corner of the application window, there are several controls that you can use to advance through each sheet or quickly jump to the first or last sheet in the workbook. These controls are only available when there are too many sheet tabs to show across the bottom of the application window.

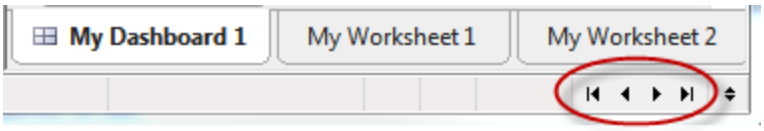

You can also navigate between sheets using the window menu or move through the multiple worksheets by pressing Ctrl + F6, or pressing the Right or Left Arrow keys on your keyboard.

To navigate through multiple worksheets, a worksheet tab at the bottom of the workbook must be selected.

Navigating within a story is a different matter—there, you use the Navigator to move between story points. See **Updating a Story** on page 814.

To make it easier to identify and group sheets, you can assign each sheet tab a color. Right-click the tab and select **Color**. You can select from seven different colors. Selecting **None** resets the color back to the default.

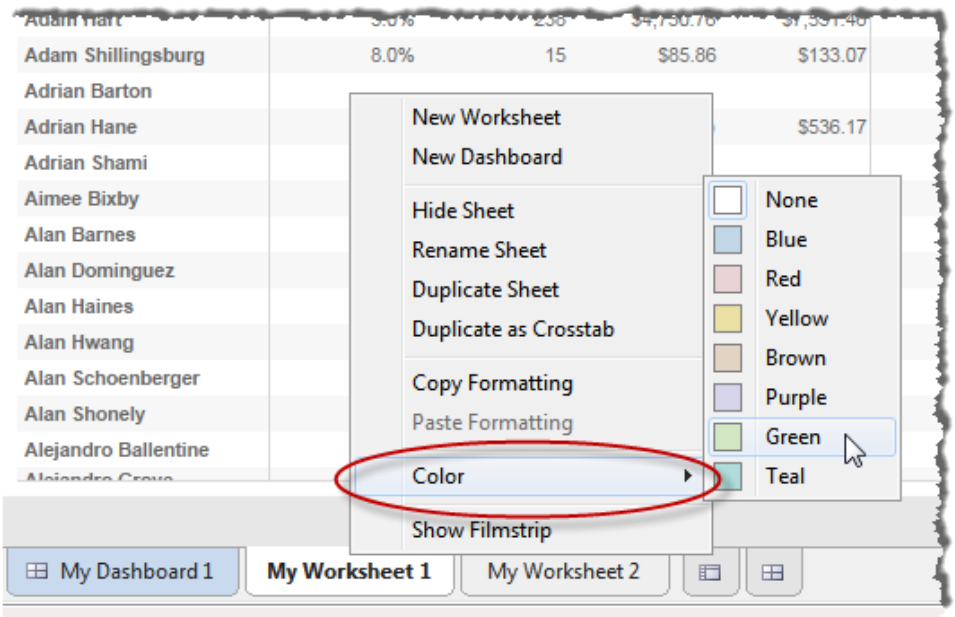

In addition to color, you can use the right-click menu to specify commands that apply to the selected sheet. For example you can create new sheets, duplicate sheets, copy formatting, and delete the sheet entirely. Hold the control key when selecting sheets to select and apply settings to multiple sheets all at once.

> Colors only display when viewing the sheets as tabs. The colors are not associated with the filmstrip and sheet sorter.

Filmstrip

Similar to the sheet tabs, the filmstrip 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 in Presentation mode.

Open the filmstrip by clicking the arrows on the bottom right corner of the workbook. Just like with the tabs, select the thumbnail image for the sheet you want to open. You can right-click the images to specify commands that apply to each sheet.

> **Note:** When viewed from a Retina display, workbooks that are created on a standard resolution device will show only the first thumbnail in the filmstrip in high resolution. To display all thumbnails in the filmstrip in high resolution, resave the workbook on a computer with Retina display.

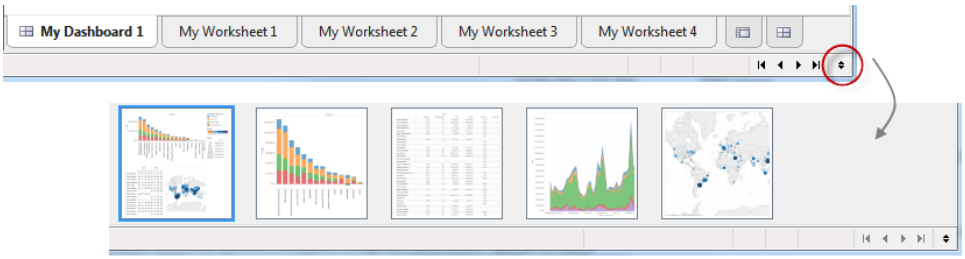

Sheet Sorter

The sheet sorter shows all 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 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.

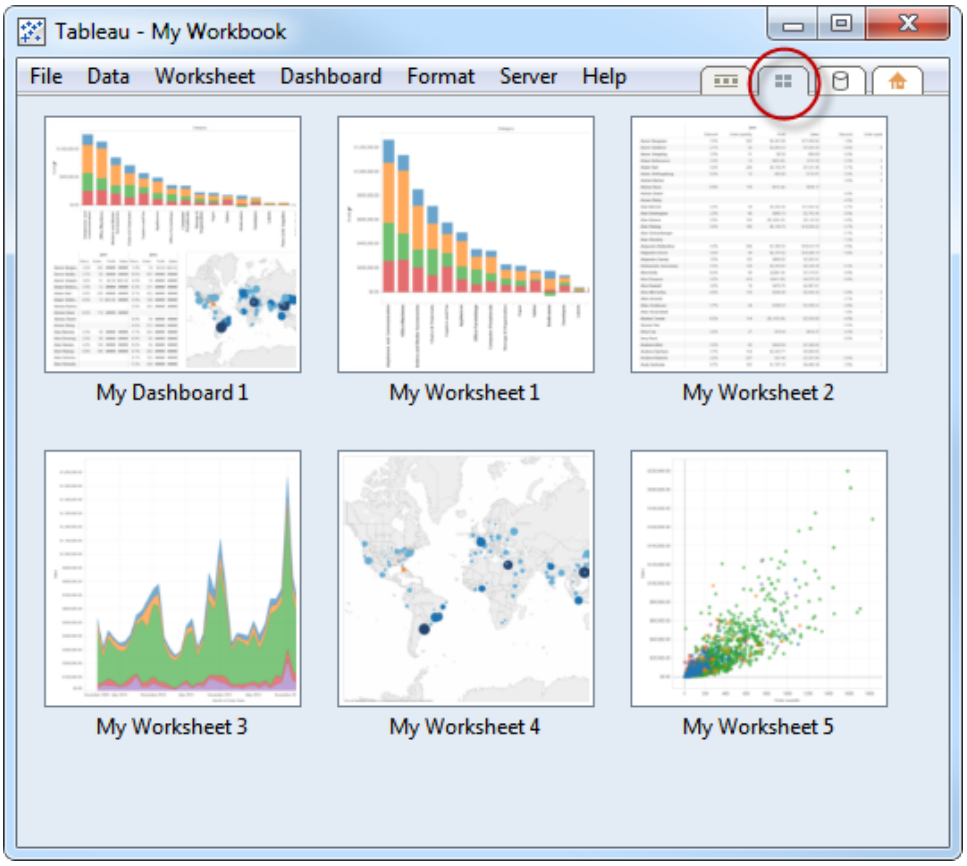

From the sheet sorter you can drag and drop to reorder the sheets, create new sheets, and duplicate or delete existing sheets. Right-click a sheet to see these commands. You can also right-click to refresh the thumbnail image of a particular sheet or **Refresh All Thumbnails** at once.

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 (.twb)** – Tableau workbook files have the .twb file extension. Workbooks hold one or more worksheets, plus zero or more dashboards and stories.
- **Bookmarks (.tbm)** – Tableau bookmark files have the .tbm file extension. Bookmarks contain a single worksheet and are an easy way to quickly share your work. See **Bookmarks** on page 863.
- **Packaged Workbooks (.twbx)** – Tableau packaged workbooks have the .twbx file extension. A packaged workbook is a single zip file that contains 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. See **Packaged Workbooks** on page 861.
- **Data Extract (.tde)** – Tableau data extract files have the .tde file extension. 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.
- **Data Source (.tds)** – Tableau data source files have the .tds file extension. Data source files are shortcuts for quickly connecting to data sources that you use often. Data source files do not contain the actual data but rather the information necessary to connect to the data source as well as modifications you've made in the Data window such as default properties, calculated fields, groups, and so on. See **Exporting Data Sources** on page 149.
- **Packaged Data Source (.tdsx)** – Tableau packaged data source files have the .tdsx file extension. A packaged data source is a zip file that contains the data source file (.tds) described above as well as any local file data sources such as Extract files (.tde), text files, Excel files, Access files, and local cube files. Use this format to create a single file that you can then share with others who may not have access to the original data stored locally on your computer. See **Exporting Data Sources** on page 149.

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.

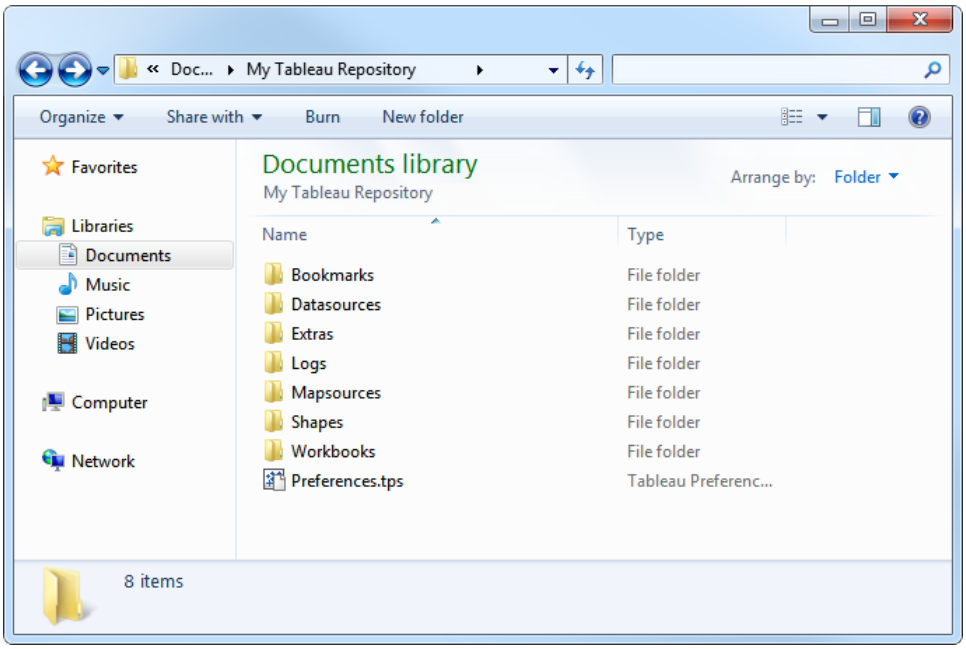

Changing the Repository Location

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.

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.

Changing the repository location does not move the files contained in the original repository. Instead, Tableau creates a new repository where you can store your files.

Learning to Use Tableau

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

Open Tableau

You can open Tableau by selecting **All Programs** > **Tableau 8.3** on the Windows Start menu or by double-clicking the desktop shortcut.

Tableau opens to 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

have started working in the worksheet by clicking the start page tab in the upper-right corner of the workbook.

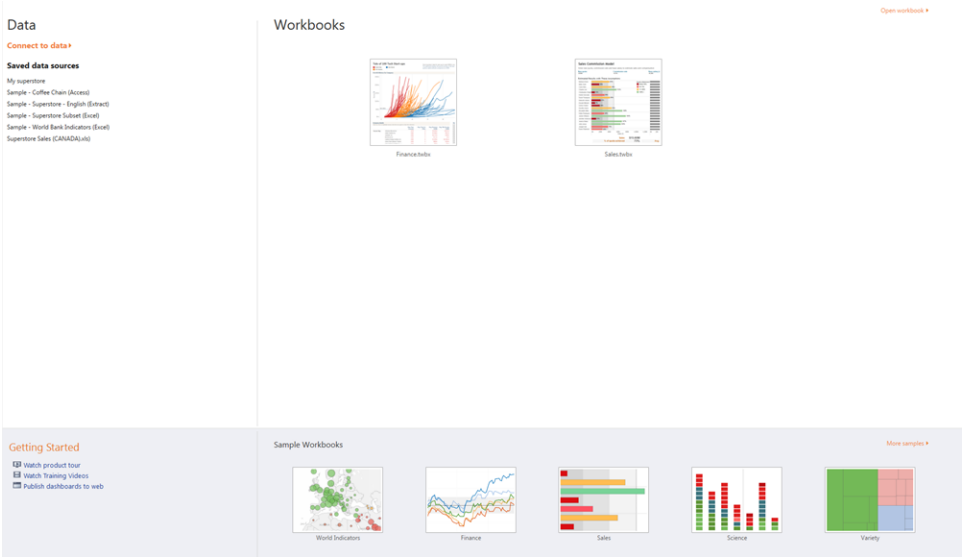

Open Data

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 Sample - Superstore data source that is installed with the application. 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.

1. Click **Connect to Data** on the left side of the start page. The Connect page opens showing the types of data sources you can connect to.
2. For this example, select **Microsoft Excel**.

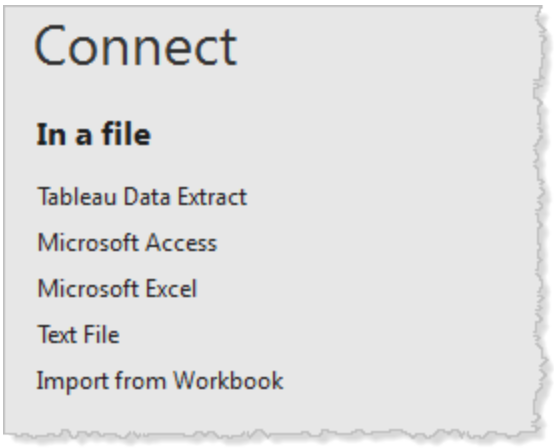

3. Navigate to an Excel file on your computer and open it.

The Excel data for this example is located in the Datasources directory of your Tableau Repository. By default, your Tableau Repository is created in your My Documents folder when Tableau is installed.

Once you have connected to the Excel data, the data source page displays the sheets in your data. You can then drag these sheets to the join area and combine them to form the data source from which you will build your view. See the **Combining Multiple Tables** on page 169 topic to learn how to join data. If there are more than 1 million rows in the data source, you will need to filter the data.

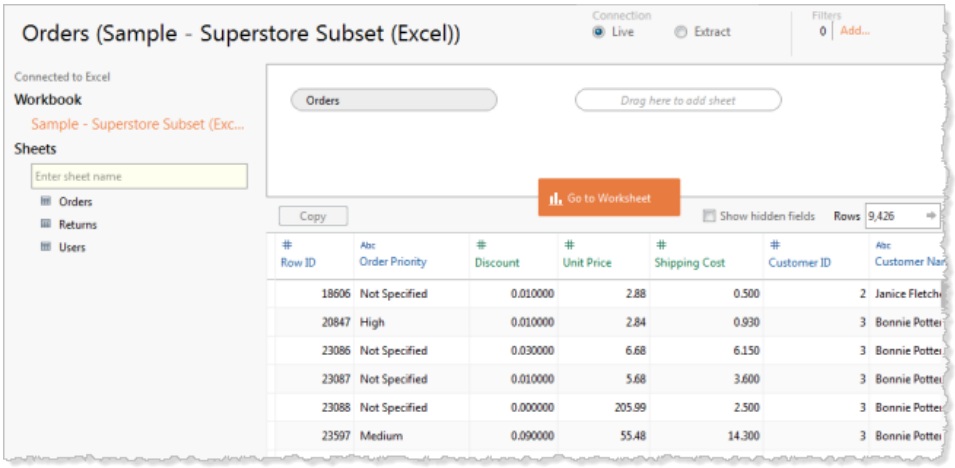

4. At the top of the data source page, select how you want to connect to the data and the data source filters you want to apply. You can select from the following options:
   - **Live** – Creates a direct connection to your data. The speed of your data source will determine performance.
   - **Extract** – By default, this option imports the entire data source into Tableau fast data engine as an extract. The extract is saved with the workbook. If you prefer to import a subset of the data, click the **Edit** link. This option requires you specify what data you want to extract using filters.
   - **Add** - Add data source filters to limit the visibility of fields contained in the workbook.
5. At the bottom of the data source page, preview the data in your data source. You can also make the following modifications:
   - Hide or rename the column. Click the header drop-down arrow, and select **Hide** or **Rename**.
   - Change the data type. Click the data type icon in the header, and select the default data type for the column.
6. Click the **Go to Worksheet** button to begin building your view.

Once connected, the columns from the data source (e.g., Customer Name, Order Date, and Sales) are shown as fields 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 in Tableau is that you can build views of your data by dragging fields from the Data window to the shelves in the view.

Building Views (Manually)

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

This section presents five data views using the Sample-Superstore Sales data source that is installed with the application.

Example 1 – Basic View

In this example, you will build a basic view that shows yearly profits. Follow the steps below to build this view.

1. Click and drag the **Order Date** field from the dimensions area of the Data window to the Columns area of the view.

   When you place a field on columns area of the view it is also added to the Columns shelf. You can also drag directly to these shelves. When you drag a field over a shelf, a blue arrow indicates that the shelf can accept the field.

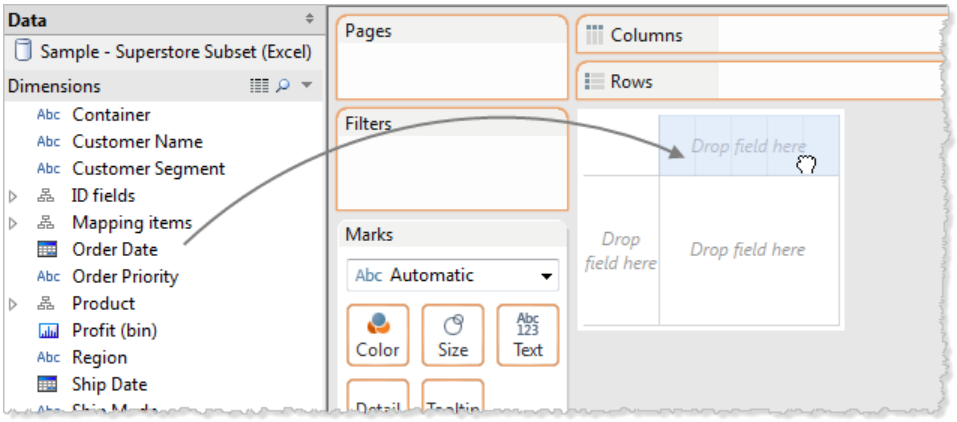

   The resulting table has four columns and one row. Each column header represents a member of the Order Date field (2010, 2011, 2012, and 2013). Each cell contains an “Abc” label, which indicates that the current mark type for this view is text.

Notice that the field is colored blue indicating it is discrete. 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.

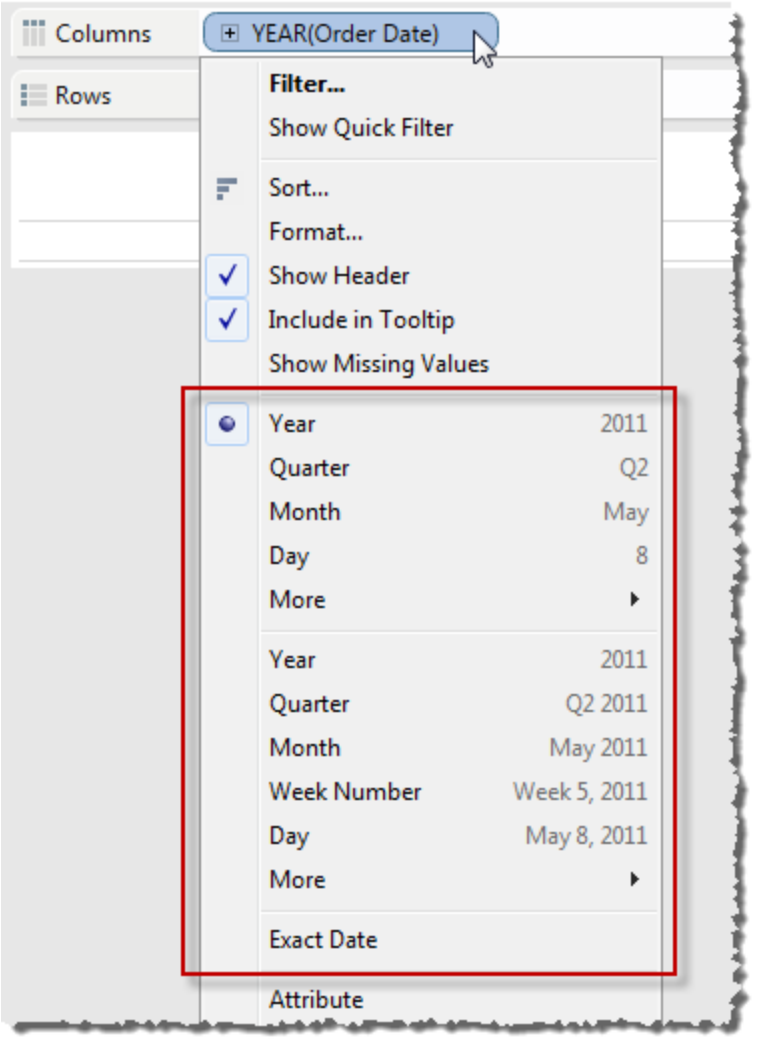

2. Drag the Profit field from the Measures area of the Data window and drop it on the Rows area of the view.

   The table is automatically transformed into a line chart and a vertical axis is created for the measure. A line is used as a great way to compare data over time, and allow you to visually compare data and identify trends effectively.

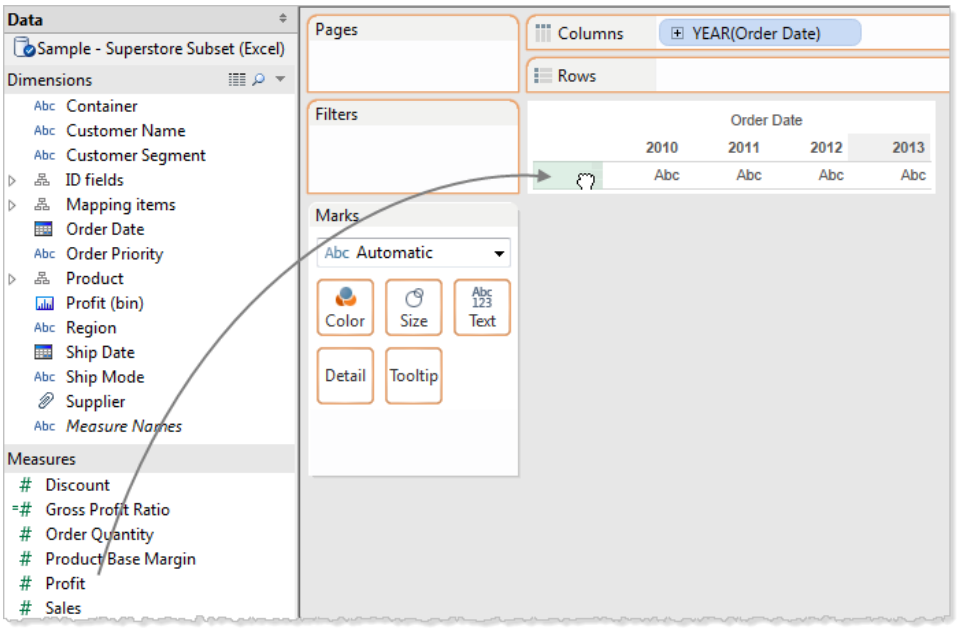

The line chart shows profit over time. Each point along the line shows the sum of profit for the corresponding year.

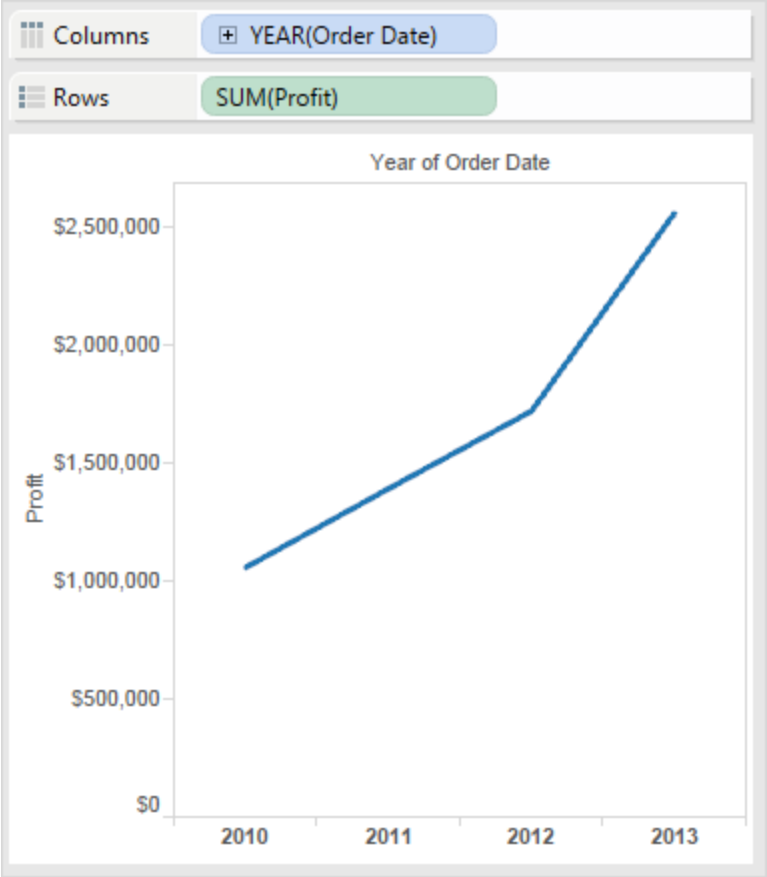

The Profit field is colored green, indicating that it is continuous. Also the field name changed to SUM(Profit). That's because the default aggregation for measures in a

relational database is SUM. You can change the aggregation using the field menu.

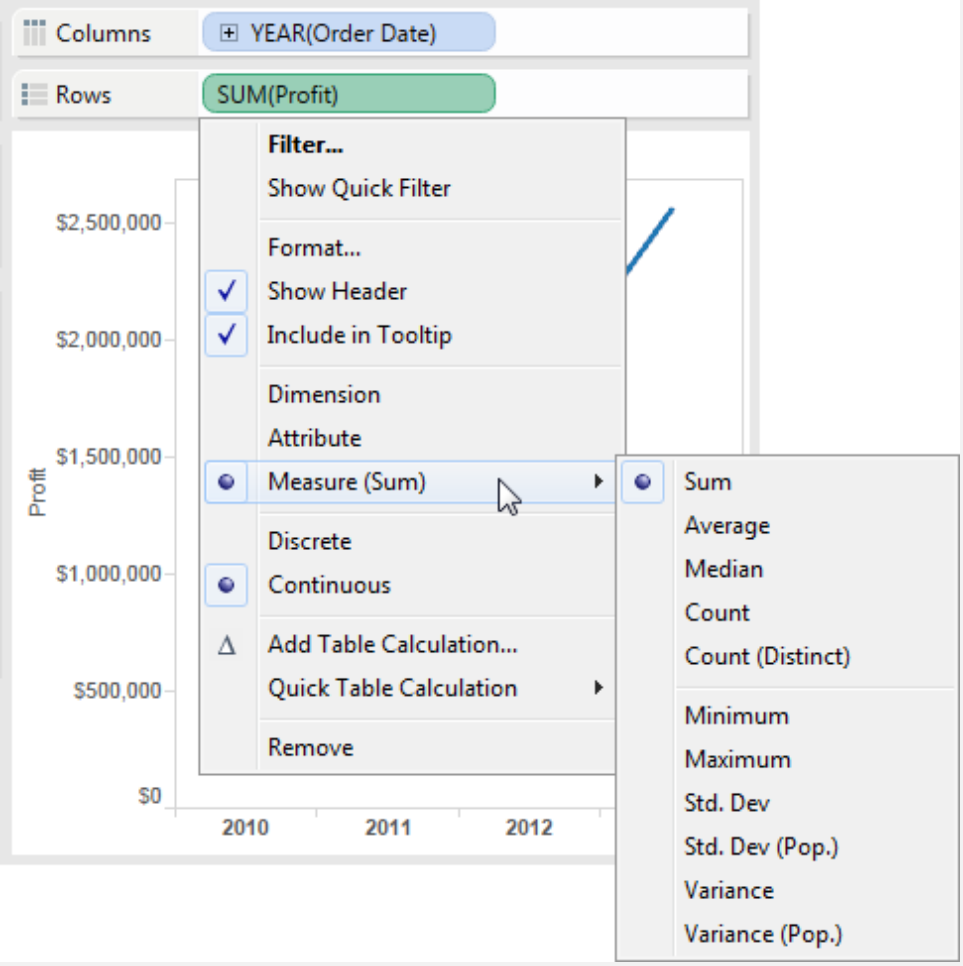

Example 2 – Nested Table

In this example you will modify the view from **Example 1 – Basic View** on page 30 to show quarters in addition to years. Follow the steps below to build this view.

Show quarters using one of the following methods:

- Drill down on the Year(Order Date) field by clicking the plus button ⊞ on the left side of the field.

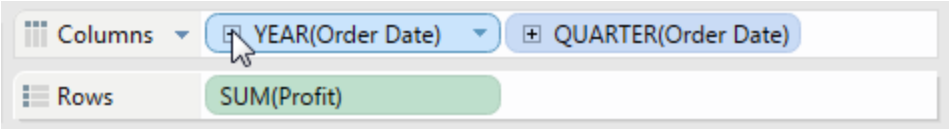

- Drag the Order Date field from the Data window and drop it on the Columns shelf to the

right of the Year(Order Date) field.

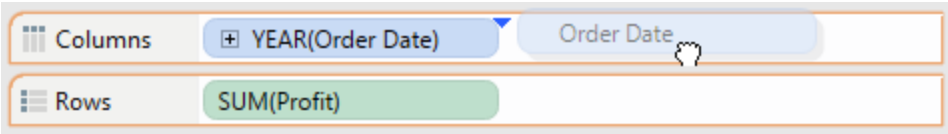

The new dimension divides the view into separate panes for each year. Each pane has columns for the quarters of the given year. This view is called a nested table because it displays multiple headers, with quarters nested within years.

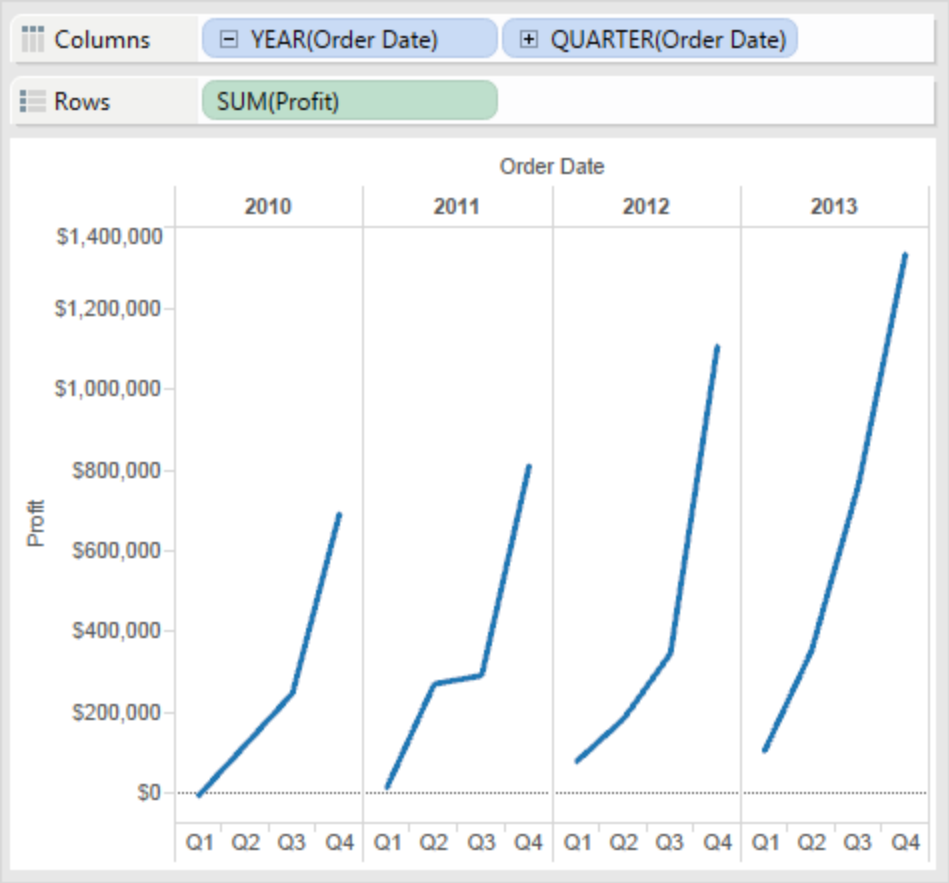

Example 3 – Small Multiples

In this example you will modify the view from **Example 2 – Nested Table** on the previous page to show quarterly profit by year and customer segment.

Drag the **Customer Segment** dimension from the Data window and drop it just to the left of the Profit axis in the view.

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

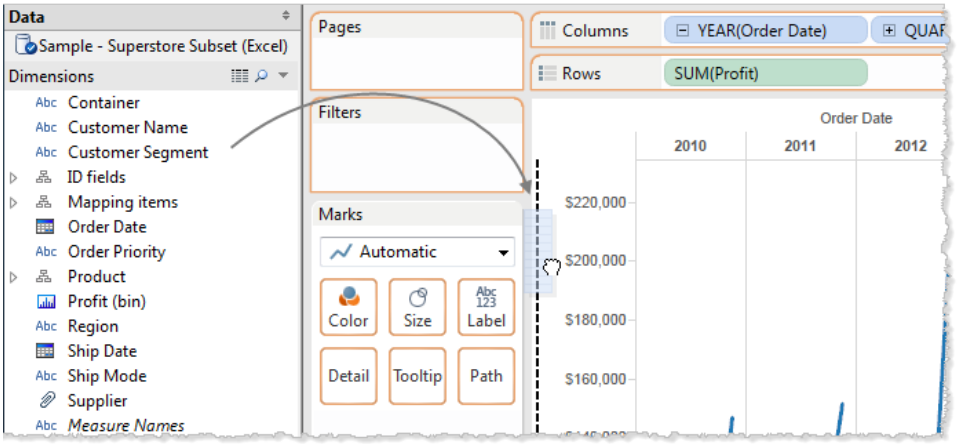

Tableau does not allow you to place a dimension to the right of a measure on either the Rows or Columns shelves.

The new dimension divides the view into 16 panes: one for each combination of year and customer segment. This view is a more complex example of a nested table and is often referred to as a small multiples view.

Example 4 – Filter Data

In this example you will modify the view from **Example 3 – Small Multiples** on page 36 to only show data for orders in 2011 and 2012 and made a profit.

1. Drag the **Profit** measure from the Data window and drop it on the Filters shelf.

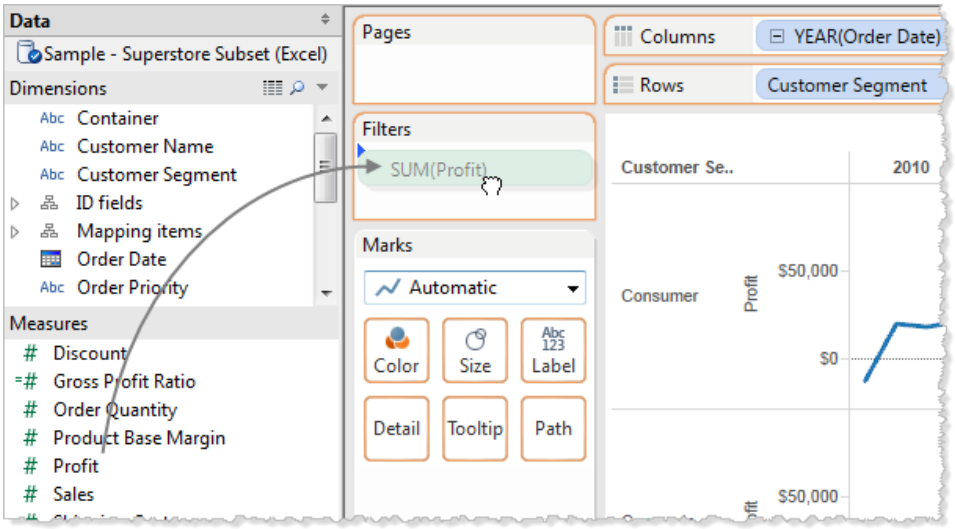

2. Select the aggregation you want to use for the filter. For this example, select **SUM** to create a filter on the SUM of the profit data.

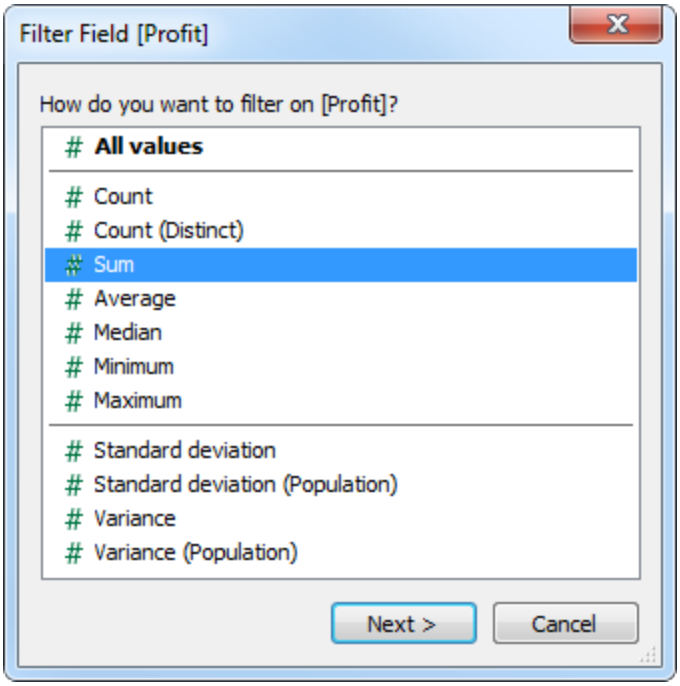

Filtering aggregated data means that the selected aggregation function (sum, average, etc.) is applied to the data and then it is filtered. Filtering disaggregated data (All values) means that the individual data rows are filtered before any aggregation function is applied. Aggregations do not apply to multidimensional data because in those databases, the data has already been aggregated in the database.

3. In the Filter dialog box, use the sliders or type into the text box to change the lower limit

on the range of values to 0.

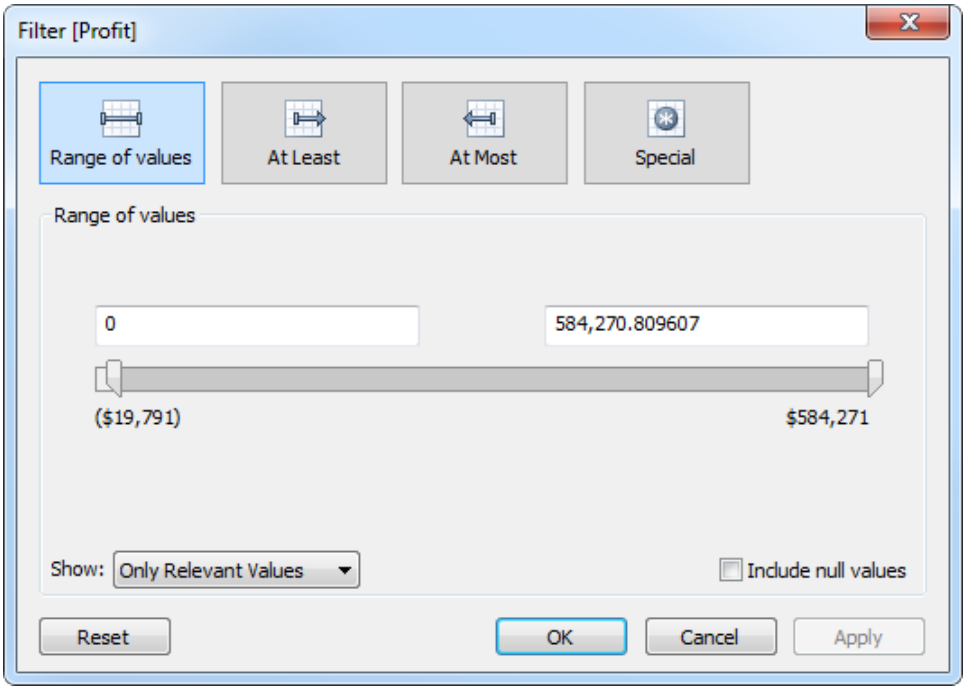

4. When finished, click **OK**.

   The view updates to only show orders that had a SUM(Profit) over $0. No new axes were created because the field was not added to the Rows or Columns shelves.

5. Select **Filter** on the Year(Order Date) field menu.

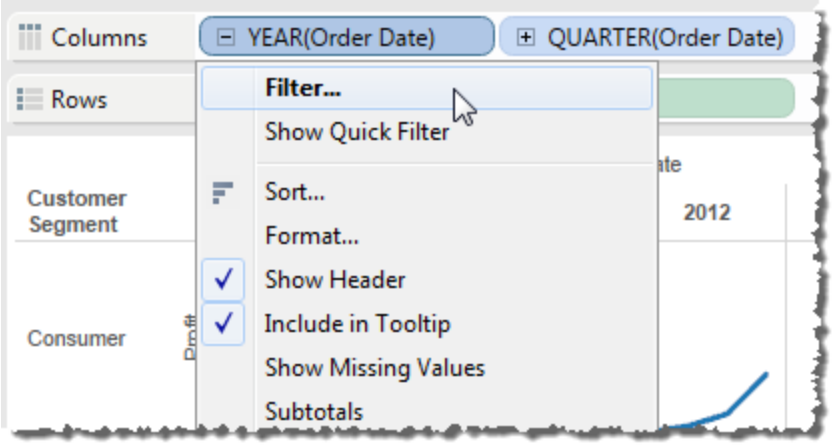

6. In the Filter dialog box, deselect all years except 2011 and 2012 and then click **OK**.

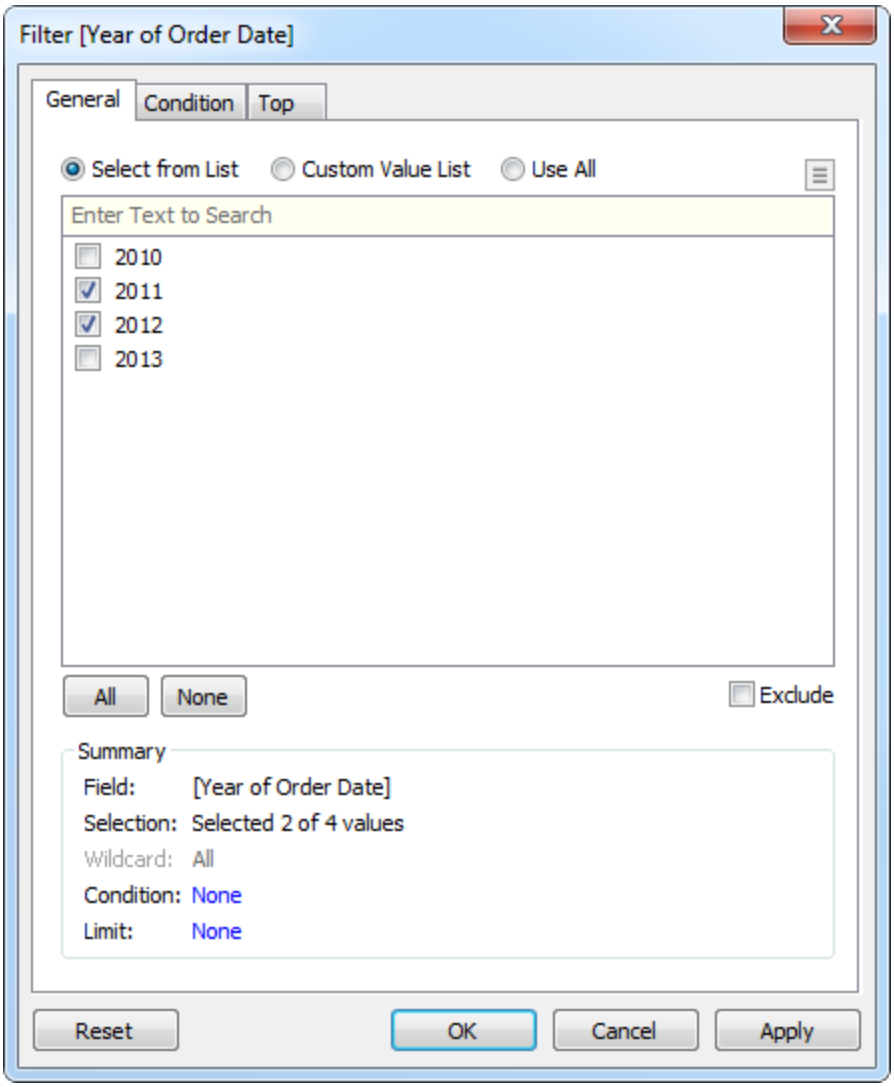

The SUM(Profit) is calculated only for data rows where profit is greater than or equal to $0 and Order Date is 2011 or 2012.

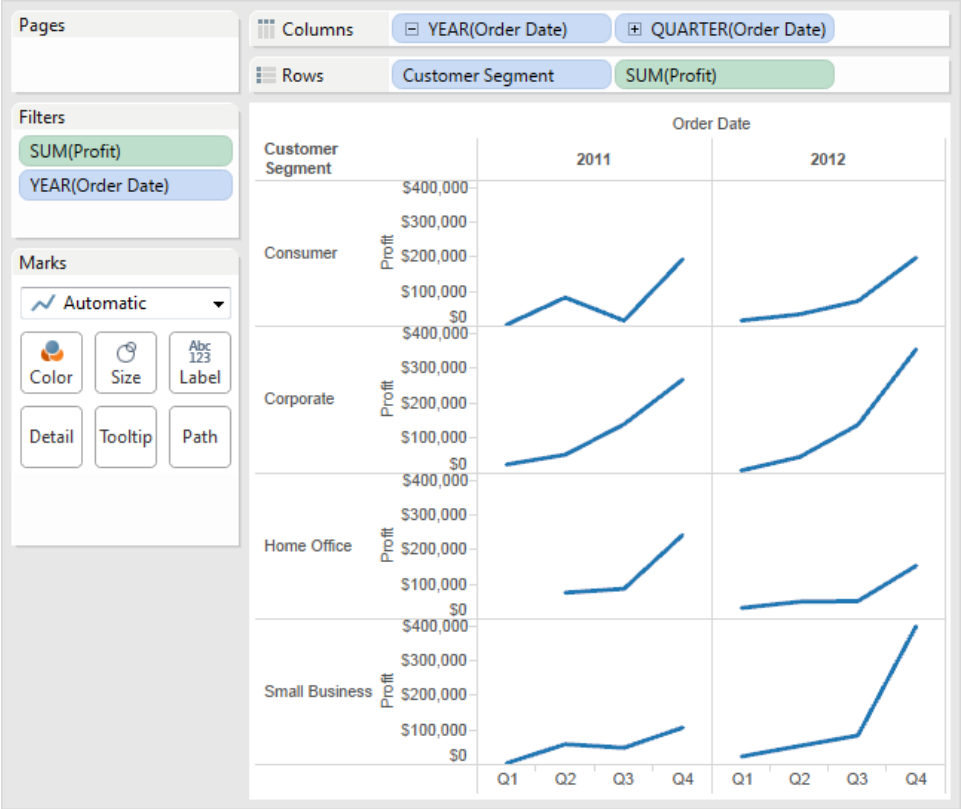

Example 5 – Color Encoding

In this example you will modify the view from **Example 4 – Filter Data** on page 38 to color the marks by region. Follow the steps below to build this view.

Drag the **Region** dimension from the Data window and drop it on Color.

Placing a dimension on Color 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.

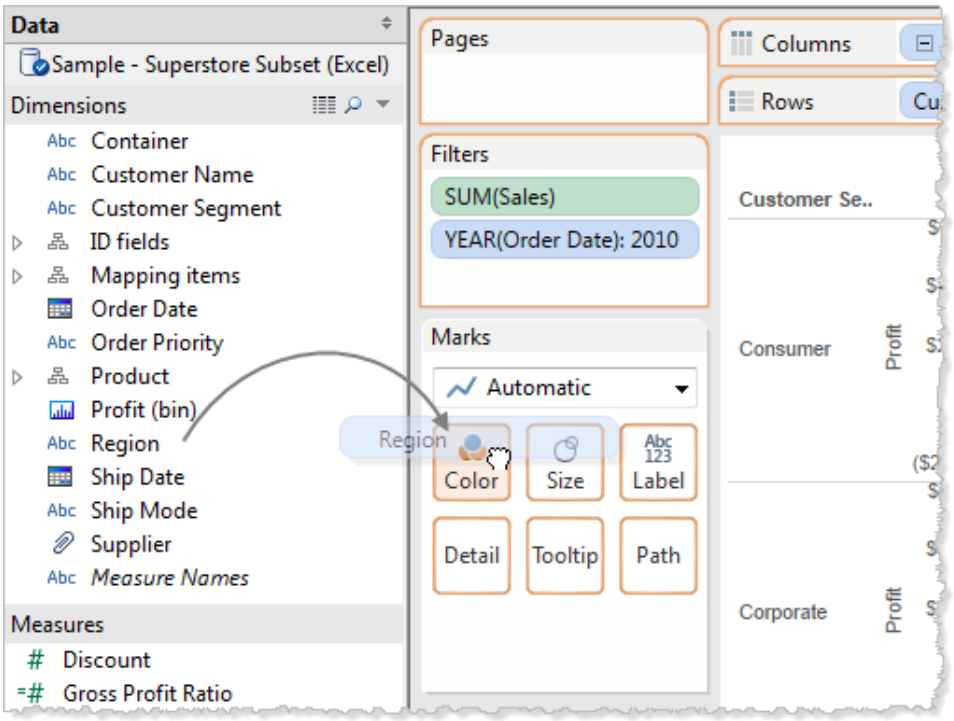

Each pane now has four lines, one for each region. This view now shows profit for each customer segment and region for 2011 and 2012 orders with a positive profit.

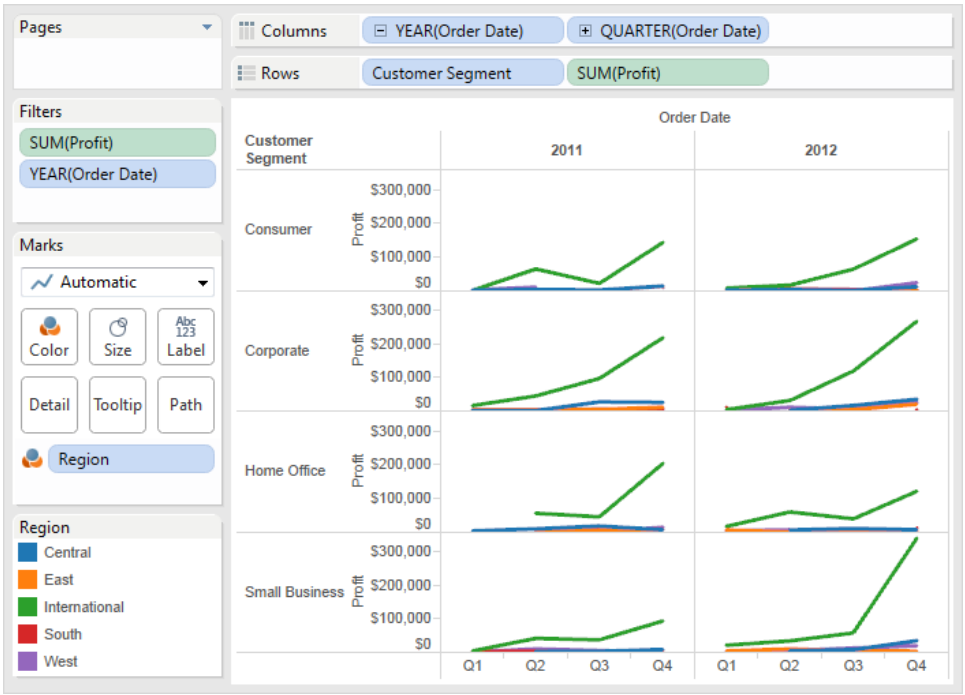

Building Views (Automatically)

Rather than building views by dragging and dropping fields, you can use Show Me™ to create views automatically.

This section presents two examples using the Sample-Superstore Sales (Excel) data source that comes with the application.

Example 1 – Show Me with Two Fields

In this example, you will create a line chart that displays profit as a function of time. Follow the steps below to create this view.

1. Select **Order Date** and **Profit** in the Data window. Hold the Control key (Ctrl) on your keyboard to select multiple fields.

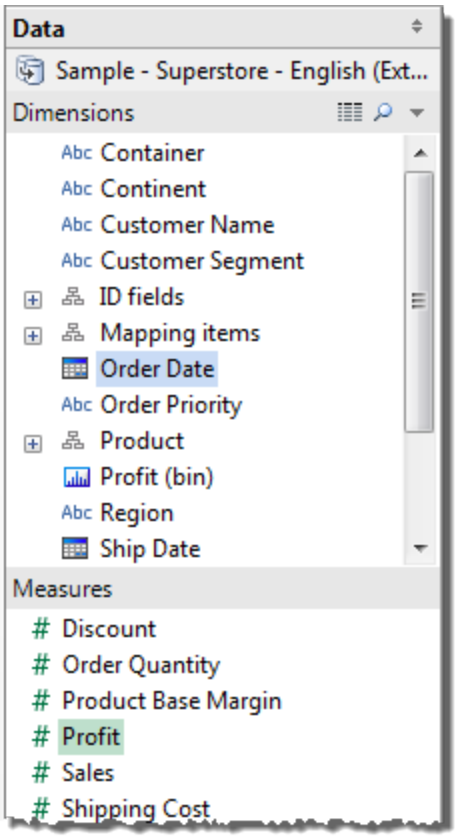

2. Click **Show Me** on the toolbar.

3. In the Show Me pane, select the type of view you want to create.

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

The view below shows SUM(Profit) over time. Each point on the line represents the sum of profit for the corresponding year.

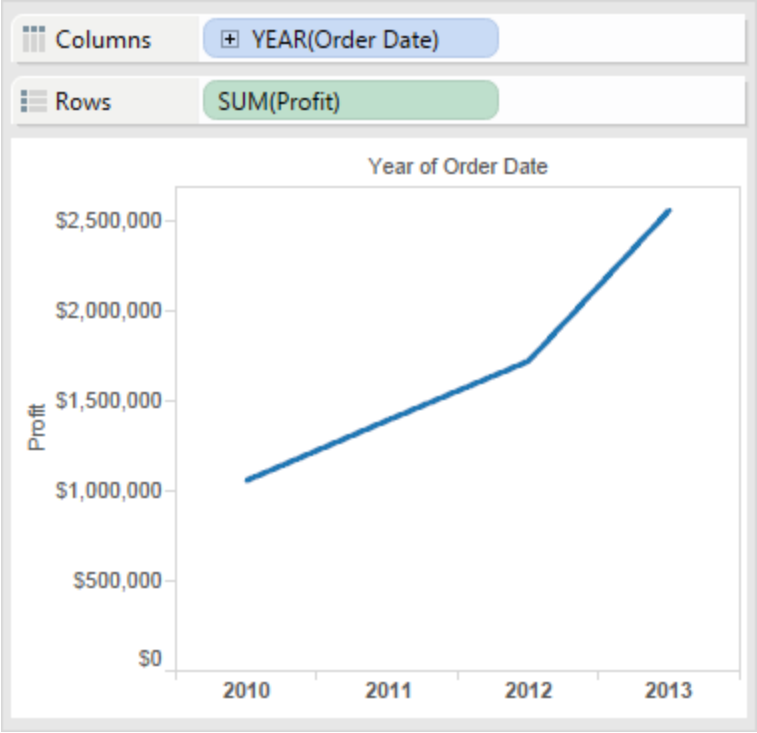

You can see the values for each year by turning on Mark Labels. Click the **Mark Labels** button on the toolbar.

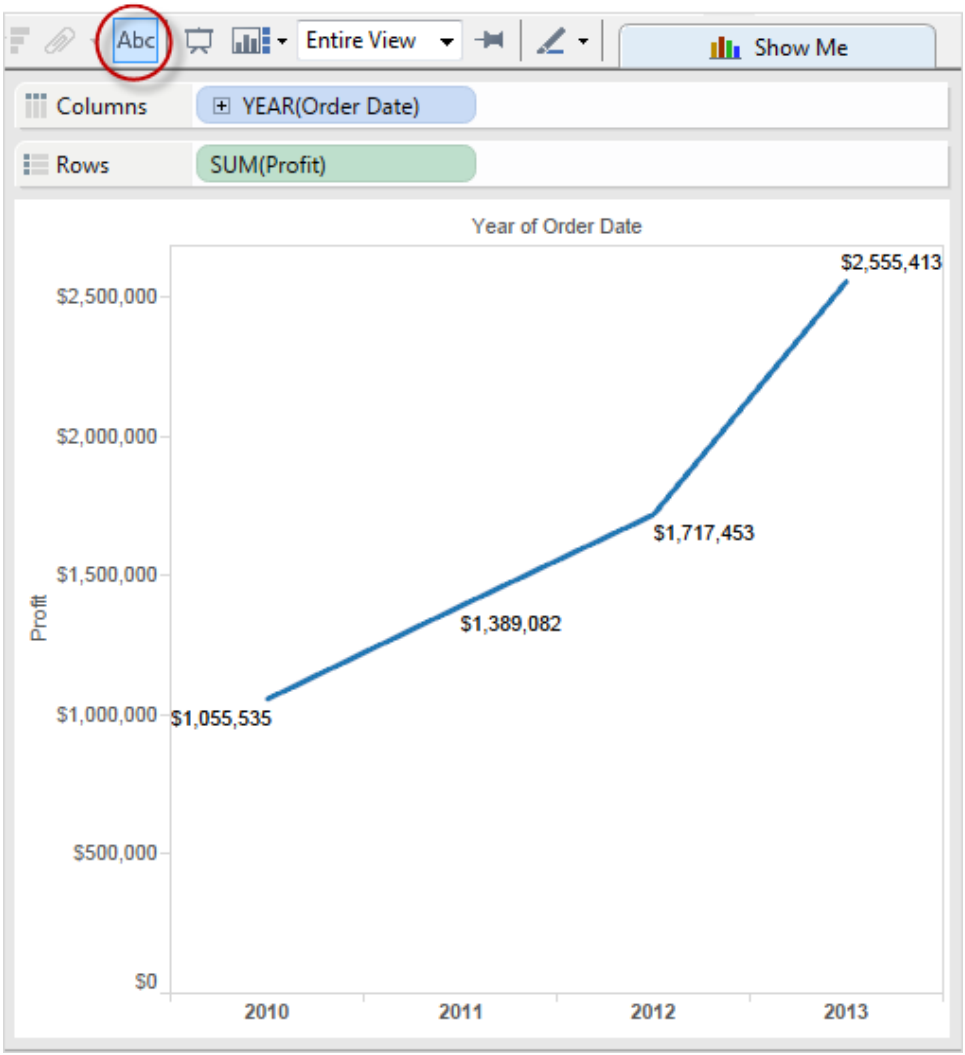

Example 2 – Show Me with Many Fields

In this example you will use Show Me to build a scatter plot that shows sales versus profit for each product and customer.

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

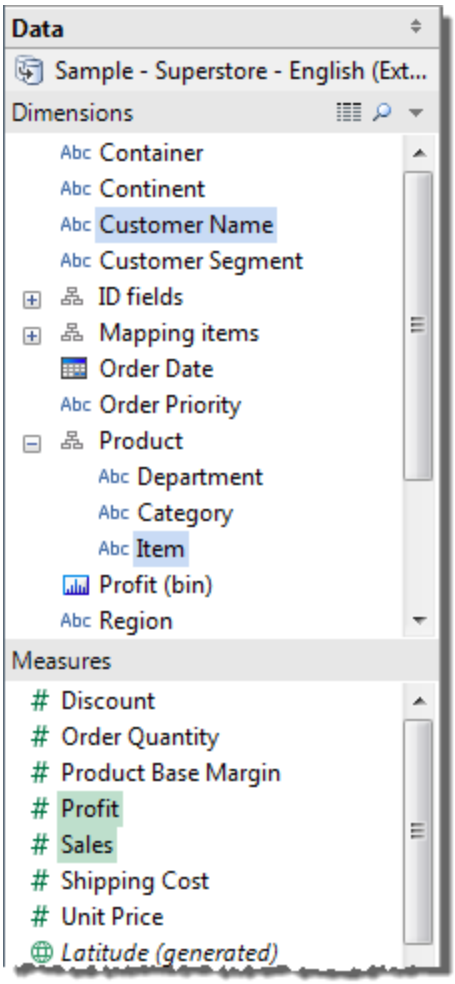

2. Click **Show Me** on the toolbar.

3. In the Show Me pane, select the scatter view.

Show Me automatically creates a scatter plot with the fields you selected. You can now manually start dragging fields to further refine the view.

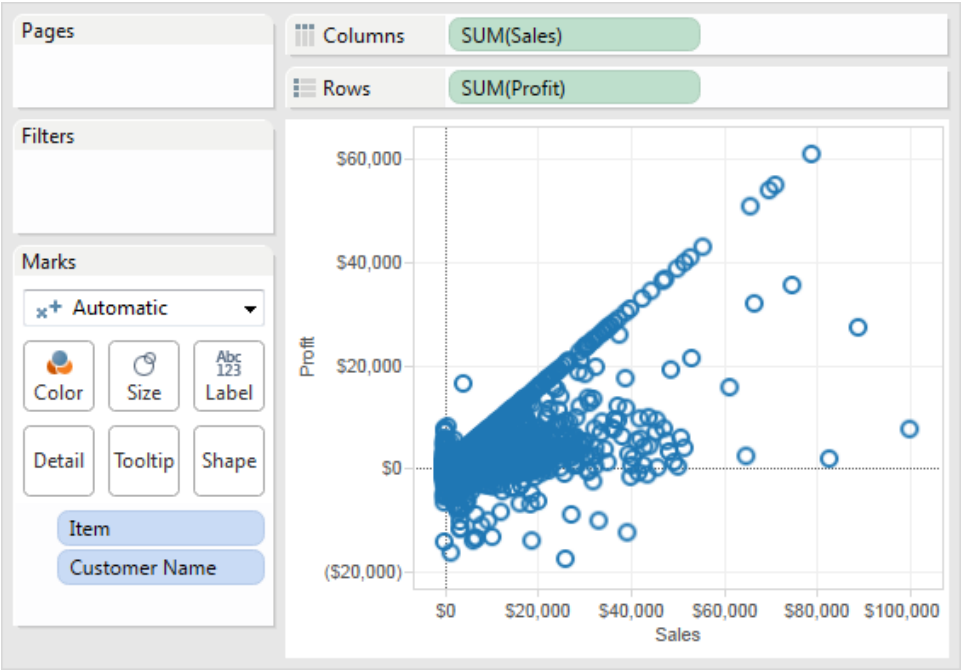

Save your Work

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 your worksheets for later use. It also allows you to share your results using a convenient file. Follow the steps below to save your workbook.

1. Select **File > Save** or press Ctrl + S on your keyboard.

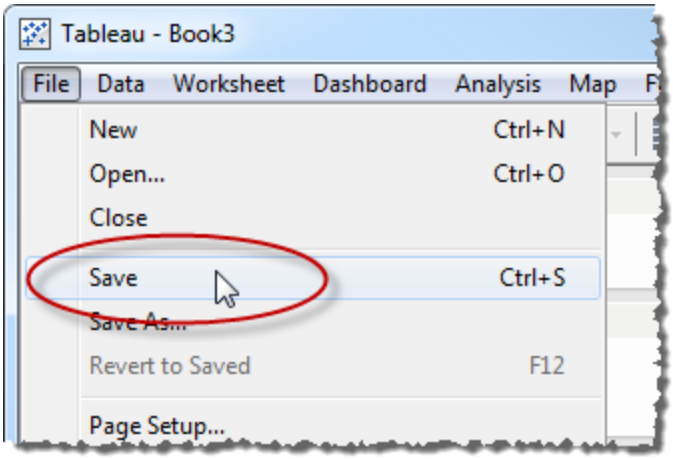

2. Browse to a file location to save the workbook.

   By default, Tableau saves workbooks in the Workbooks directory in your Tableau Repository.

3. Specify a file name for the workbook.

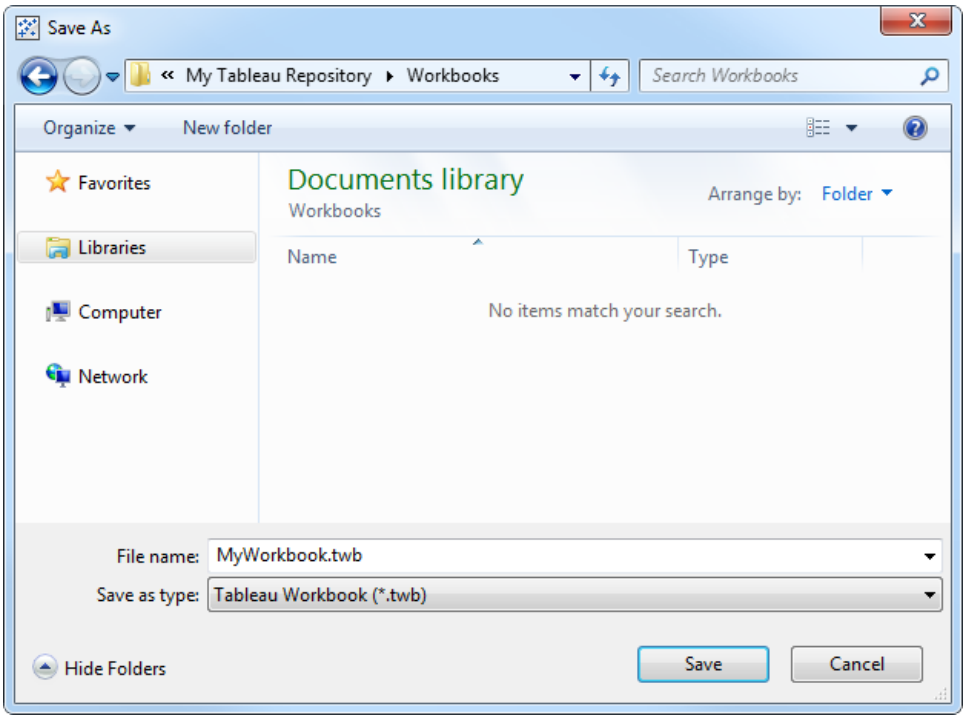

4. Specify a file type. You can select from the following options:

   Tableau Workbook (.twb) – Saves the all the sheets and their connection information in a workbook file. The data is not included.

   Tableau Packaged Workbook (.twbx) – Saves all the sheets, their connection information and any local resources (e.g., local file data sources, background images, custom geocoding, etc.).

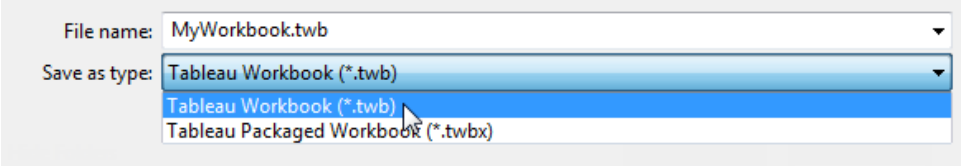

5. When finished, click **Save**.

Connect to Data

This section explains how to create and manage both basic and advanced connections to all of the supported data sources. Follow a step-by-step tutorial for connecting to each data source. Also, learn how to paste data into Tableau, join new tables, create and manage extracts, and monitor queries. In addition to the following, see **Date Properties for a Data Source** on page 438.

Basic Connection

To begin building your view and analyzing the 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 and setting up the data source in Tableau, to the data source and preparing the data source for analysis, 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.

Supported Data Source Types

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

The data source types supported by your copy of Tableau Desktop is determined by the version purchased. For more information on data source compatibility requirements, go to the Technical Specifications page on the Tableau website.

Connecting to Data and Editing Data Sources

Connect to Data

To build views of your data, you must first connect Tableau to data and set up the data source.

To connect to Data

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.

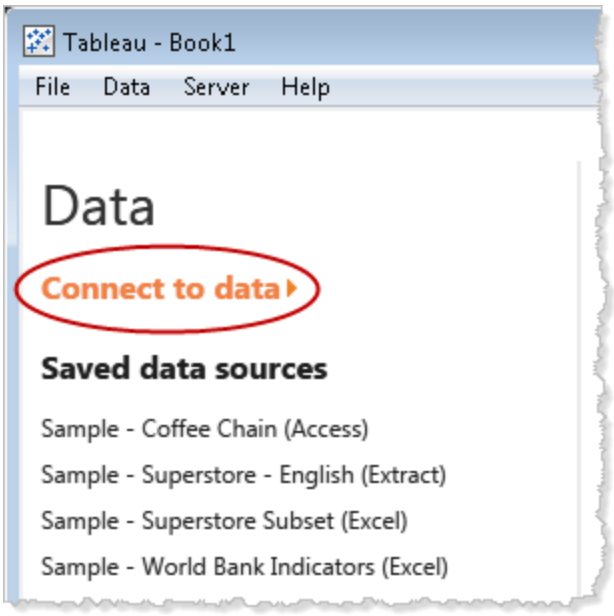

2. On the Connect page, select the type of data you want to connect to.

After the initial connection is established, the database and schemas or tables display on the left side of the page. Here you can combine different sheets or tables to create joins, preview the data in your data source, and then go to the worksheet. When you go to the worksheet, the data source fields display on the left side of the workbook in the Data window.

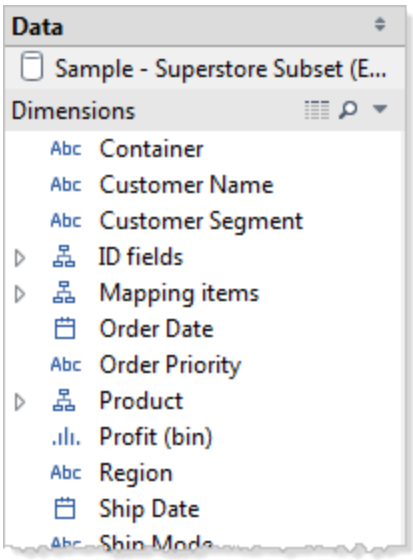

You can display information about the connection by selecting the data source on the **Data** menu and then selecting **Properties**. The properties of an example data source are shown below.

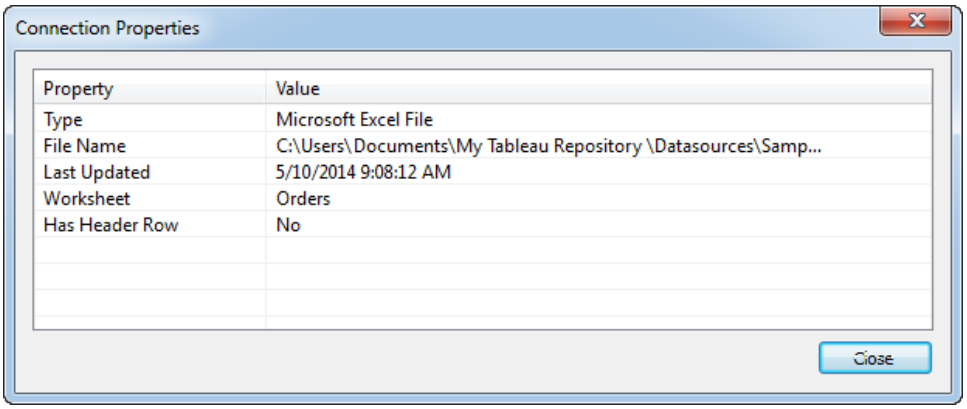

Edit the Data Source

At anytime during your analysis you can edit the data source used in the workbook. You might want to edit the data source to:

- Specify a new location for the data source. For example, 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, you can edit the original data connection and specify a new data source.

To edit the data source

1. On the **Data** menu, select a data source, and then select **Edit Data Source**.
2. On the data source page, make the changes to the data source.

   For example, if you are using a **Microsoft Excel** data source, you can specify a new file by clicking the Excel file name under **Workbook**. Alternatively, you can select a different table to analyze.

Replace Field References

When you successfully connect to a new data source, all worksheets in the workbook 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 as the original workbook, the fields become invalid and are marked with an exclamation point . 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 data source connection to point to a new data source that has all the same data but instead of Product Category, the field name has been changed to Product

Type. The Product Category field remains in the Data window but is marked as invalid. To make the field valid, you can replace the references, which means you can map the invalid field to a valid field in the new data source (for example, Product Category corresponds to Product Type).

To replace field references

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

Duplicate the Data Source

Sometimes you'll want to make changes to a data source 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 source. You can duplicate the data source so that you can make the changes without affecting the existing sheets. To duplicate a data source, on the **Data** menu, select a data source , and then select **Duplicate**. When you duplicate a data source, the duplicate source name has "(copy)" appended to the end.

Rename the Data Source

When you connect to a data source, you are given the option to give it a name for use in Tableau Desktop. You can change the name you specified by selecting the data source on the **Data** menu and then selecting **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 a data source on the **Data** menu, and then selecting **Properties**.

Using OAuth for Data Connections

For Salesforce.com, Google Analytics, and Google BigQuery data sources, you can connect to data through the OAuth 2.0 open authentication process. OAuth keeps your data source credentials secure with the data provider, and it can simplify the management of embedded credentials and extract refreshes.

When you perform tasks that require signing in to an OAuth-enabled data source, you automatically create an OAuth connection. For example:

- Connect to Salesforce.com, Google Analytics, or Google BigQuery data.
- Open a published workbook that connects to the OAuth-enabled data source.
- Create a new workbook from a data source published to Tableau Server or Tableau Online.

Connecting to Data Using OAuth

On the Tableau Desktop Connect page, when you select an OAuth-enabled data source, you are directed to the data provider's sign-in form, where you provide your credentials and approve Tableau access.

The data provider returns an access token that is used for authentication.

Workbooks with OAuth Connections

When you open a published workbook that has a connection to an OAuth data source, you are prompted to authenticate to the data source. If you have an existing OAuth access token for that data, you can select the existing token to open the workbook. Otherwise you can create a new access token (credential).

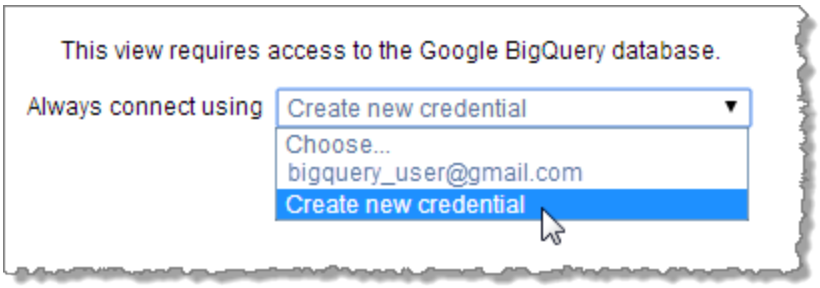

Selecting **Create new credential** takes you through the authentication process as described earlier in this topic in **Connecting to Data Using OAuth** on the previous page.

> **Note:** If you have not published your workbook or data source to a Tableau server, you are not able to embed OAuth credentials to bypass authentication and connect directly to the data. For information about using embedded OAuth credentials with published workbooks and data sources, see the next section, **OAuth on Tableau Server** below, and the related topics.

OAuth on Tableau Server

If you have an account on Tableau Server or Tableau Online, you can save OAuth tokens with your user preferences. You can then use the saved tokens when you open a published workbook and are prompted for credentials.

You can also edit data connections on the server to embed access tokens, to allow direct access to the data, or to facilitate scheduled extract refresh tasks.

Single Source for Embedded Credentials

If you embed the same OAuth credentials into multiple published data sources, you do not need to update each location manually when credentials expire. When it's time to replace or refresh an OAuth access token, if you use the same user credentials to obtain the update, the embedded locations are updated all at once.

For information about working with OAuth tokens on the server, see the following topics:

- Manage Credentials
- Edit Data Connections, specifically Authentication Options for Google and Salesforce.com

Connecting to Data Sources

Click one of the following for information on how to connect to a data source.

Tableau Data Extract

This topic describes how to connect Tableau to a Tableau Data Extract file.

1. Click **Connect to Data**.
2. Click **Tableau Data Extract**.
3. In the Open dialog box, navigate to and select a Tableau Data Extract file.
4. Click **Open**.
5. Give the connection a name for use in Tableau.

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

A completed Connection dialog box for Tableau Data Extract files is shown below.

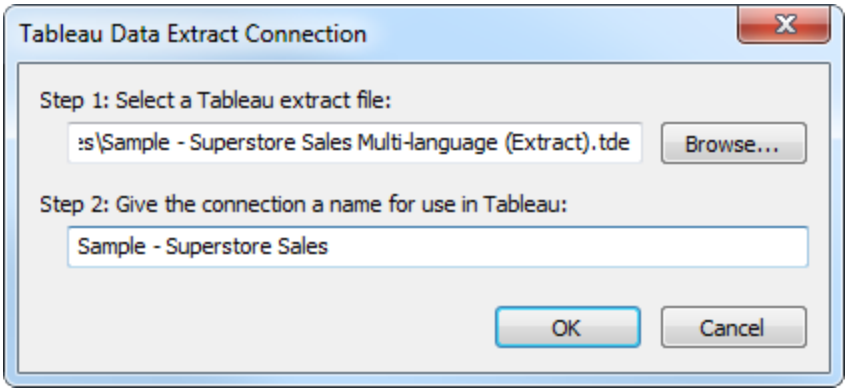

Microsoft Access

This topic describes how to connect Tableau to a Microsoft Access file and set up the data source. Tableau supports all Access data types except OLE Object and Hyperlink.

1. Click **Connect to Data**.
2. On the Connect page, click **Microsoft Access**, select the Access file that you want to connect to, and then click **Open**. The data source page displays.

   If the Access file is password protected or it is a file that is protected by workgroup security, a dialog box opens where you can enter your database password or select **Workgroup Security**. If the file is protected by workgroup security, type the Workgroup Information file name, User, and Password into the corresponding text fields.
3. On the data source page, do the following:
   1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.The default name is automatically generated based on the file name.
   2. Drag a table to the join area. You can drag a table or query.
   3. If your Access file contains one table, click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

Optional Settings

You can set the following options before building your view.

Preview the data and change data types

The first 10,000 rows of the data in the data source is automatically displayed below the join area. If you add tables, remove tables, or make changes to the join conditions, the preview updates. You can also do the following in the preview area:

- Change the data type of a column by clicking the data type icon.
- Rename or hide the column by clicking the column header drop-down menu and selecting the respective option.

Connect live or use an extract

At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link will display allowing you to set up filters that define the subset of the data to include in the extract.

Add data source filters

Add data source filters to restrict the visibility and use of the fields in the data source.

An example of an Access data source is shown below.

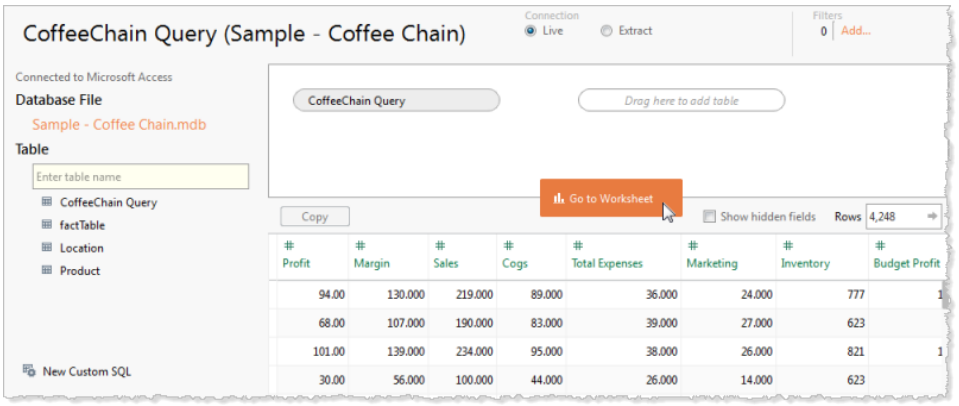

Note: If the Access file contains columns that are more than 254 characters wide, Tableau will not be able to use these fields. Either remove the columns from the table or modify them to fit within 254 characters prior to connecting with Tableau.

Microsoft Excel

This topic describes how to connect Tableau to a Microsoft Excel file and set up the data source. Tableau can connect to .xls and xlsx files.

1. Click **Connect to Data**.
2. On the Connect page, click **Microsoft Excel**, select the Excel workbook you want to connect to, and then click **Open**.

   To connect to the Excel file using the Microsoft Jet-based connection, in the Open dialog box, click the **Open** drop-down menu, and then select **Open with Legacy Connection**. The data source page displays.
3. On the data source page, do the following:
   1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
   2. If your Excel file has one table, click **Go to Worksheet** to start building the view.

      Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

      You can also connect to a named range the same way you connect to a worksheet. Named ranges are created in Excel by highlighting a range of cells and then selecting **Define Name** on the **Formulas** tab.

      You can connect to multiple Excel worksheets at the same time as long as each worksheet has a unique connection name.

Optional Settings

You can set the following options before building the view.

Set Excel table options

In the join area, hover over the table name until the gear icon displays, click the icon, and then specify whether the data includes field 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.

Preview the data and change data types

The first 10,000 rows of the data in the data source is automatically displayed below the join area. If you add tables, remove tables, or make changes to the join conditions, the preview updates. You can also do the following in the preview area:

- Change the data type of a column by clicking the data type icon.
- Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.

Connect live or use an extract

At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, the Edit link displays allowing you to set up filters that define a subset of the data that you want to include in the extract.

Add data source filters

Add data source filters to restrict the visibility and use of the fields in the data source.

Additional Information

An example of a Microsoft Excel data source is shown below.

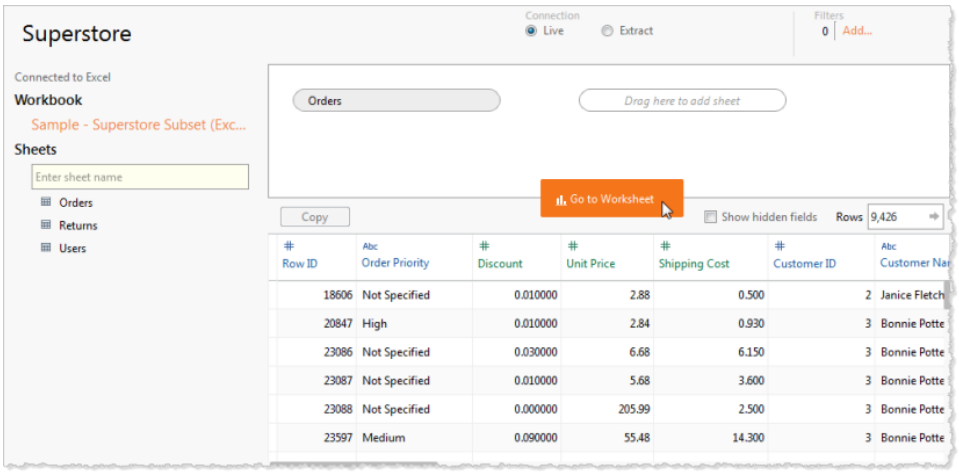

Note: If the Excel file contains columns that are more than 254 characters wide, Tableau is not able to use these fields for workbooks that were created before Tableau Desktop 8.2 or use the legacy connection. Either remove the columns, modify them to fit within 254 characters prior to connecting in Tableau, or upgrade the Excel data source.

Text File

This topic describes how to connect Tableau to a text file data and set up the data source. Tableau connects to delimited text files.

1. Click **Connect to Data**.
2. On the Connect page, click **Text File**, select the text file that you want to connect to, and then click **Open**.

   To connect to the text file using the Microsoft Jet-based connection, in the Open dialog box, click the **Open** drop-down menu, and then select **Open with Legacy Connection**. The data source page displays.
3. On the data source page, do the following:
   1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.The default name is automatically generated based on the file name.
   2. Click **Go to Worksheet** to start building the view.

      Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

Optional Settings

You can set the following options before building the view.

Set text file options

In the join area, hover over the table name until the gear icon displays, click the icon, and then specify the following:

- 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. Select from the list of characters or select **Other** to type in a custom character.
- Select the text qualifier that encloses values in the text file.
- Select a character set that describes the text file encoding. The available encodings are based on the operating system you are using. For example, on Windows, ANSI is listed as windows-1252 and OEM is listed as 437.

> **Note:** In workbooks created prior to Tableau Desktop 8.2 or use the legacy connection, 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 is verified when the connection is attempted.

- Select the locale by which the file should be parsed. This option tells Tableau which decimal and thousands separator to use.

Preview data and change data types

The first 10,000 rows of the data in the data source is automatically displayed below the join area. If you add tables, remove tables, or make changes to the join conditions, the preview updates. You can also do the following in the preview area:

- Change the data type of a column by clicking the data type icon.
- Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.

Connect live or use an extract

At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link will display allowing you to set up filters that define the subset of the data to include in the extract.

Add data source filters

Add data source filters to restrict the visibility and use of the fields in the data source.

Best Practices

If all text files for a multi-table connection are in a separate directory, there is no possibility that users will inadvertently select a file that is not appropriate for the connection.

An example of a text file data source is shown below.

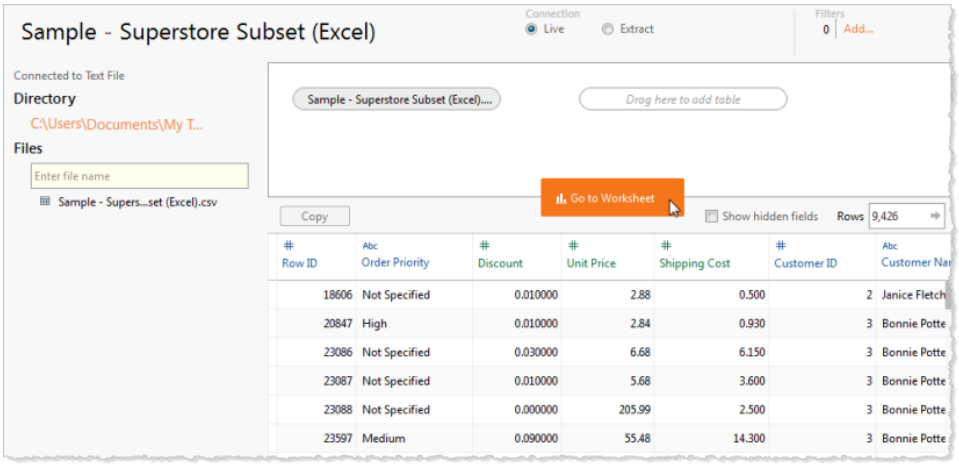

Note: If the text file contains columns that are more than 254 characters wide, Tableau is not able to use these fields for workbooks that were created before Tableau Desktop 8.2 or use the legacy connection. Either remove the columns, modify them to fit within 254 characters prior to connecting in Tableau, or upgrade the text file data source. Large text files often perform poorly as a data source, because the queries are slow.

Tableau Server

This topic describes how to connect Tableau to a Tableau Server data source.

1. Click **Connect to Data**.
2. Click **Tableau Server**.
3. Enter the name of the Tableau Server and click the **Connect** button.

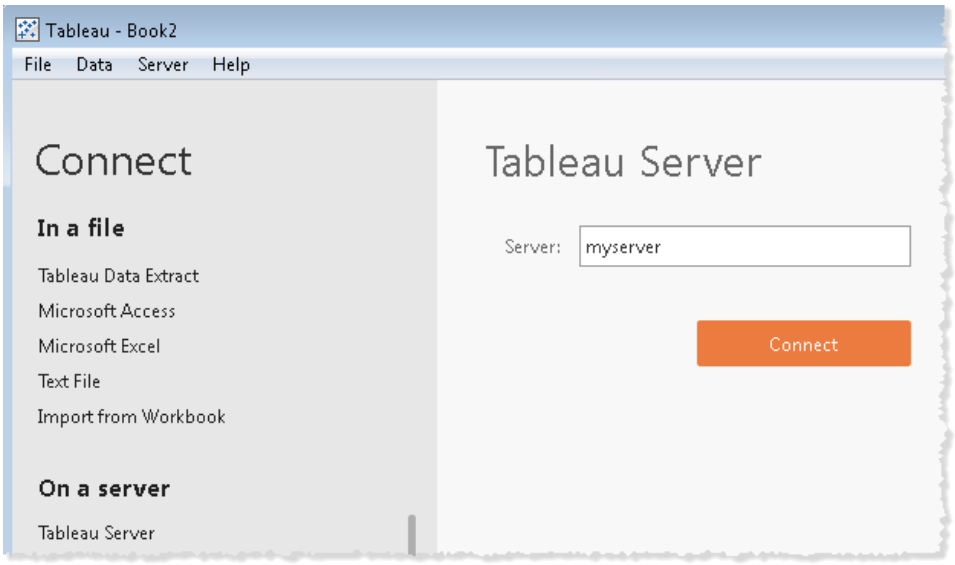

If Tableau Server is Kerberos-enabled and you are on a machine with valid Active Directory credentials, you will connect automatically to the server. In this case, you can skip step 4. If Kerberos authentication fails, you will be prompted to provide a user name and password for authentication. If you need to change the user you're signed in as, see Switch Kerberos User.

4. Enter your user name and password, and the click the **Sign In** button.

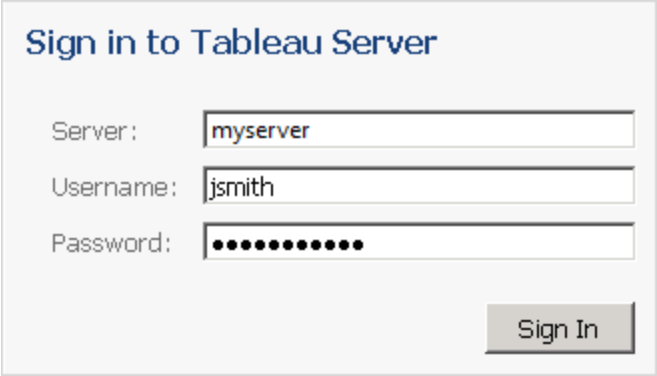

5. You can sort the list by clicking on the headers in the table. Alternatively, search for a data source using the search box in the upper right corner. Click the **Refresh** button to refresh the list and show any new data sources. Select a data source in the list of published data sources. Tableau Server data sources are shown in the Data window

with a Tableau icon.

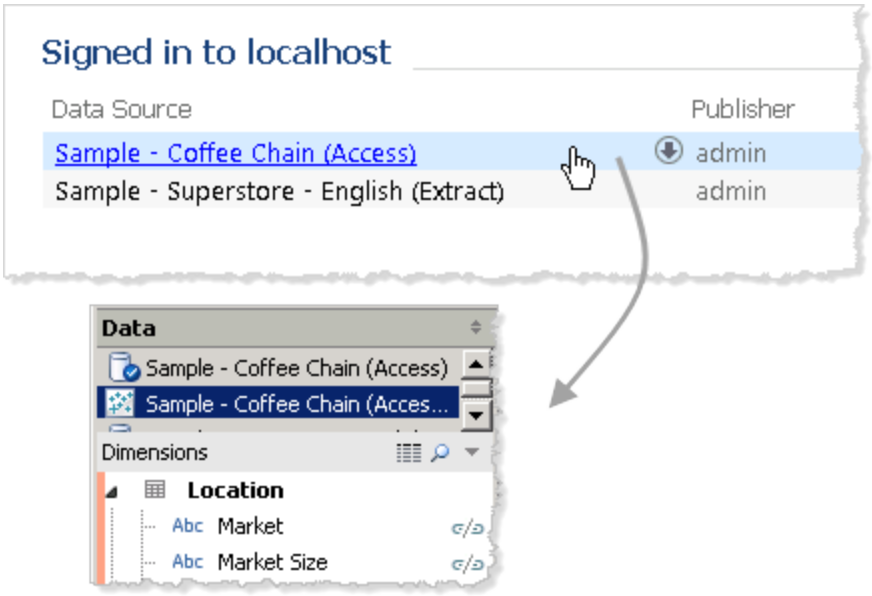

When you're connected to a Tableau Server data source, you can download a local copy by selecting the data source on the **Data** menu and then selecting **Create Local Copy**. A duplicate of the data source is added to the Data window.

Actian Vectorwise

This topic describes how to connect Tableau to an Actian Vectorwise database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **Actian Vectorwise**, and then do the following:
    1. Type the name of the virtual node for the database and the name of the database you want to connect to.
    2. Select how you want to sign in to the server. Specify whether to use Windows Authentication or a specific user name and password. If the server is password protected, you must type the user name and password.
    3. Click **Connect**.

    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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Schema drop-down list, select a schema.

3. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

   Optionally, you can make the following configurations before building the view:

   - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
   - **Change data types** – Change the data type of a column by clicking the data type icon.
   - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
   - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
   - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of an Actian Vectorwise data source is shown below.

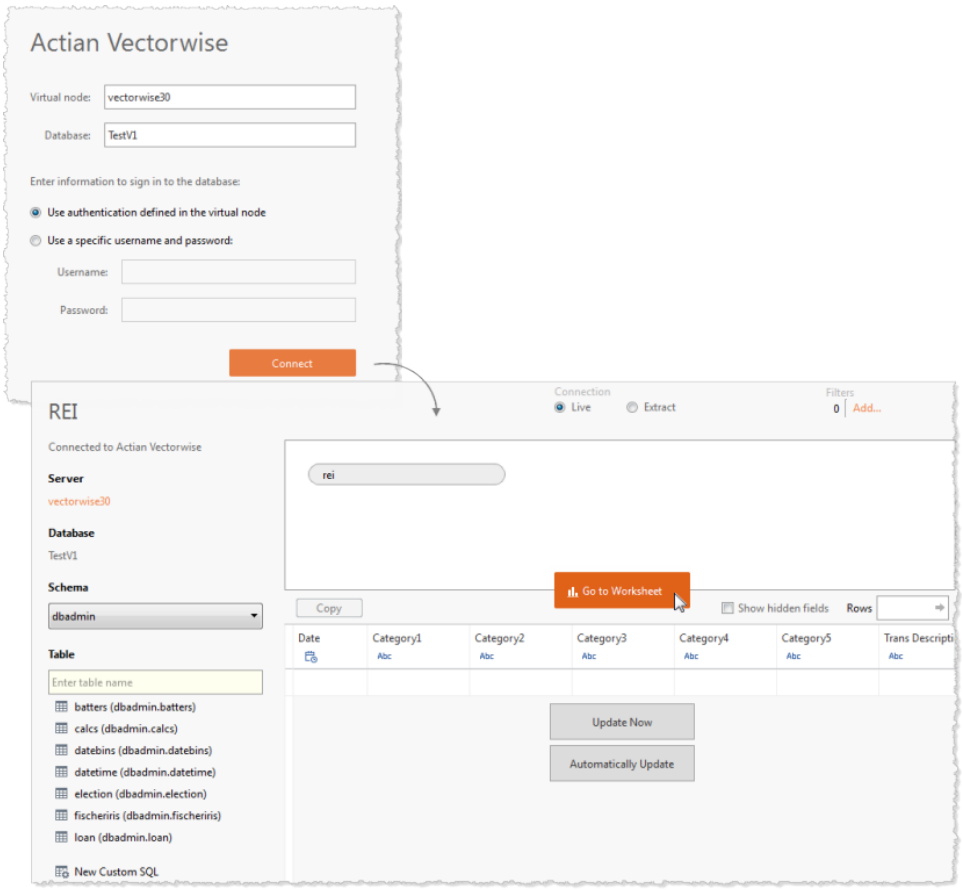

Amazon Redshift

This topic describes how to connect Tableau to an Amazon Redshift database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **Amazon Redshift**, and then do the following:
    1. Type the name of the server that hosts the database and the name of the database you want to connect to.
    2. Type the user name and password, and then click **Connect**.

       Select the **Require SSL** check box when connecting to an SSL server.

       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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.

2. From the Schema drop-down list, select a schema.
3. Drag a table to the join area, and then click **Go to Worksheet** to start building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

   Optionally, you can make the following configurations before building the view:

   - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
   - **Change data types** – Change the data type of a column by clicking the data type icon.
   - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
   - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
   - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of an Amazon Redshift data source is shown below.

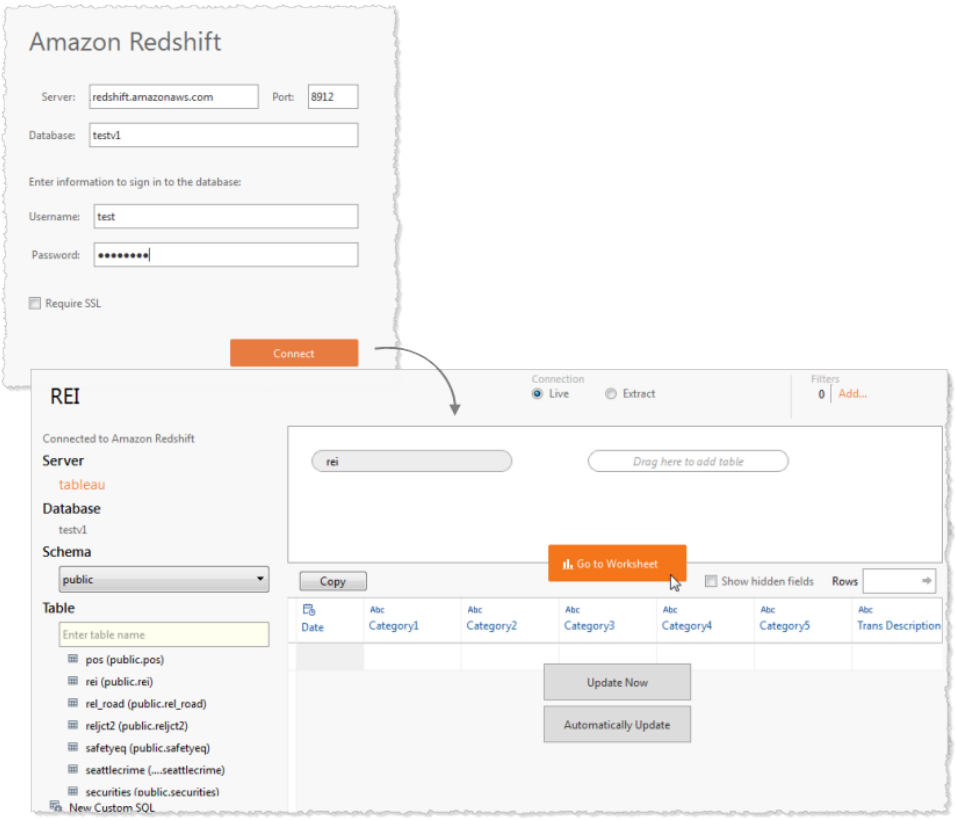

Aster Database

This topic describes how to connect Tableau to an Aster Database database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **Aster Database**, and then do the following:
    1. Type the name of the server that hosts the database and the name of the database that you want to connect to.
    2. Type the user name and password, and then click **Connect**.

       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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Schema drop-down list, select a schema.
    3. Drag a table to the join area, and then click **Go to Worksheet** to start building the view. Alternatively, drag additional tables to the join area and use conditions to join

them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

Optionally, you can make the following configurations before building the view:

- **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
- **Change data types** – Change the data type of a column by clicking the data type icon.
- **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
- **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
- **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

Initial SQL

When connecting to Aster databases, you can optionally specify a SQL command that will be run once upon connection. To do this, select **Data** > **Initial SQL**. See **Initial SQL** on page 142 to learn more about adding these commands to your connection.

In Aster, the Initial SQL can be used to generate an output table. This can improve performance of subsequent database access. Intial SQL in Aster also supports SQL-MapReduce, a framework created by Teradata Aster to allow developers to write powerful and highly expressive SQL-MapReduce functions. For details, see MapReduce, SQL-MapReduce Resources & Learning.

An example of an Aster Database data source is shown below.

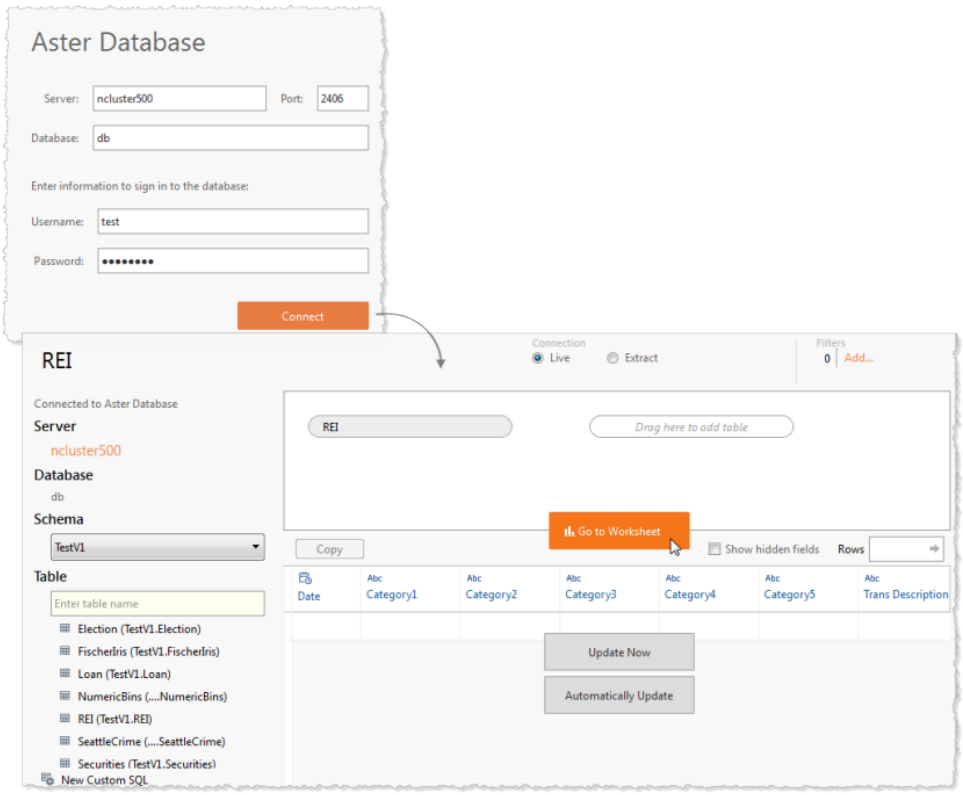

Cloudera Hadoop

This topic describes how to connect Tableau to a Cloudera Hadoop database and set up the data source. Refer to the Connecting to Cloudera Hadoop knowledge base article for more details about connecting to Hadoop data.

1. Click **Connect to Data**.
2. On the Connect page, click **Cloudera Hadoop**, and then do the following:
    1. Enter the name of the server that hosts the database and the port number to use. If you are connecting using Cloudera Impala, you must use port 21050; this is the default port if you are using the 2.5.x driver (recommended).
    2. Select how to connect to the database. Depending on the version of Hadoop and the drivers you have installed, you can connect using one of the following:
        - Hive Server
        - Hive Server 2
        - Impala
    3. Click **Connect**.

       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.
3. On the data source page, do the following:

1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
2. From the Schema drop-down list, type the schema name in the text box, click the search icon, and then select the schema.
3. Type the table name in the Table text box, click the search icon, drag the table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

   Optionally, you can make the following configurations before building the view:

   - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
   - **Change data types** – Change the data type of a column by clicking the data type icon.
   - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
   - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
   - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

Initial SQL

When connecting to Cloudera Hadoop, you can optionally specify a SQL command that will be run once upon connection. See **Initial SQL** on page 142 to learn more about adding these commands to your connection.

Hadoop and Impala

When connecting to Hadoop using Impala, the queries are limited to 100,000 rows of data when the query includes an ORDER BY clause. This limit can be configured either by modifying the `"order-by-limit"` value in the workbook xml or using a TDC customization such as `<customization name="order-by-limit" value='10' />`.

Cloudera Hadoop Data Source Example

An example of a Cloudera Hadoop data source is shown below.

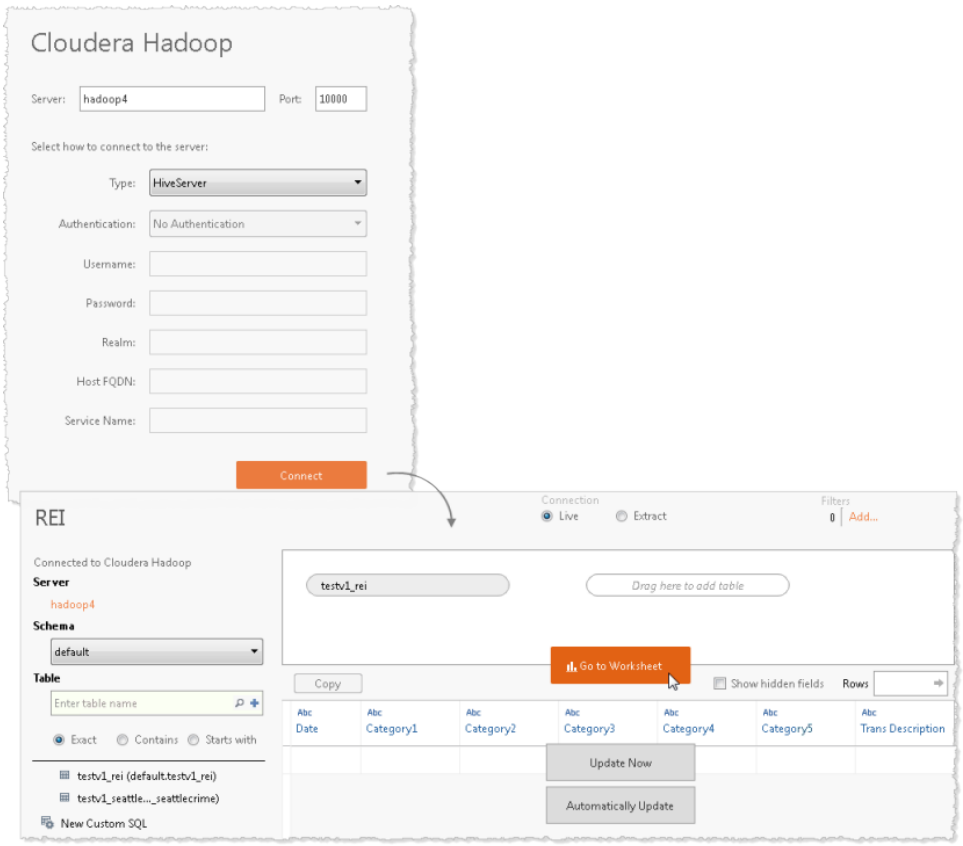

DataStax Enterprise

This topic describes discusses how to connect Tableau to DataStax Enterprise database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **DataStax Enterprise**.
3. Type the name of the server that hosts the database, and then click **Connect**.

   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.
4. On the data source page, do the following:
   1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
   2. From the Schema drop-down list, type the schema name in the text box, click the search icon, and then select the schema.
   3. Type the table name in the Table text box, click the search icon, drag the table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join

them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

Optionally, you can make the following configurations before building the view:

- **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
- **Change data types** – Change the data type of a column by clicking the data type icon.
- **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
- **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
- **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of a DataStax Enterprise data source is shown below.

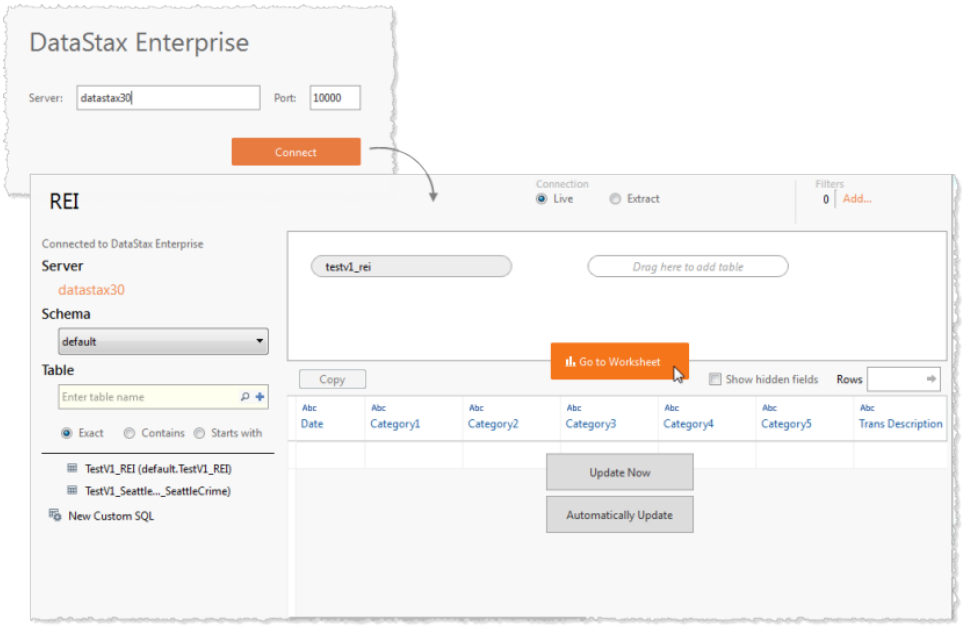

Pivotal Greenplum

This topic describes how to connect Tableau to a Pivotal Greenplum database and set up the data source.

1. Click **Connect to Data**..
2. On the Connect page, click **Pivotal Greenplum**, and then do the following:
    1. Type the name of the server that hosts the database and the name of the database you want to connect to.
    2. Type the user name and password, and then click **Connect**.

       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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

       Optionally, you can make the following configurations before building the view:

       - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
       - **Change data types** – Change the data type of a column by clicking the data type icon.
       - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
       - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
       - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of a Pivotal Greenplum data source is shown below.

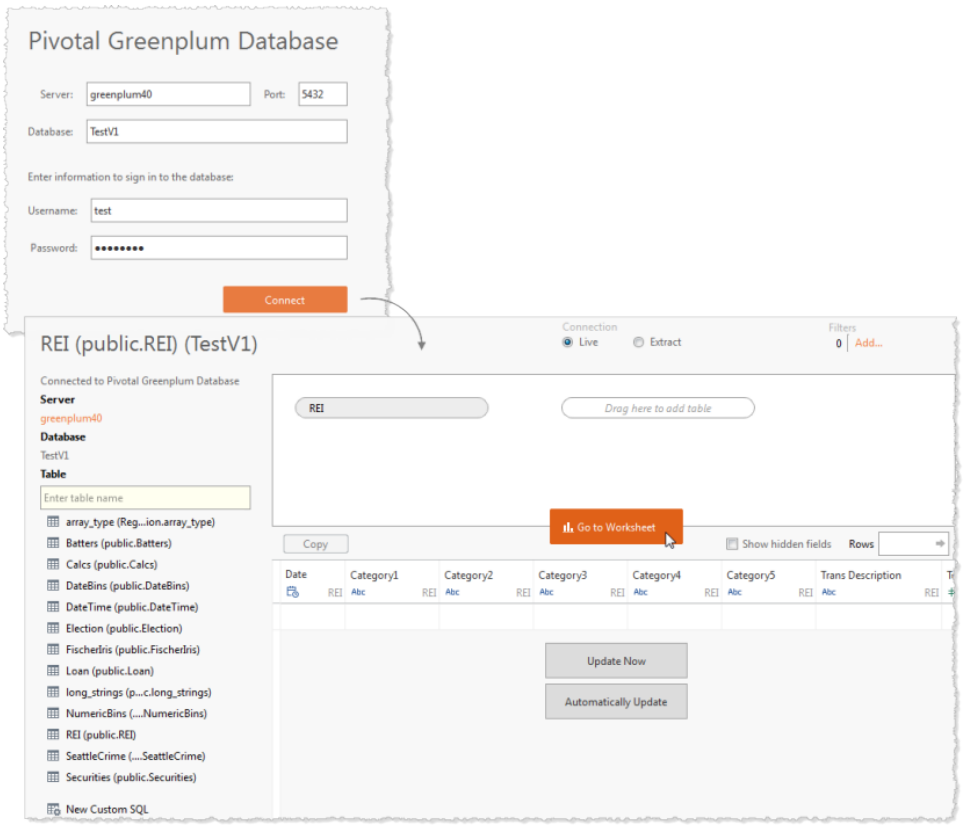

EXASolution

This topic describes how to connect Tableau to data stored in the EXASolution platform and set up the data source. Tableau can connect to EXASolution version 4.2 and later.

1. Click **Connect to Data**.
2. On the Connect page, click **EXASolution**, and then do the following:
    1. Type the name of the server that you want to connect to.
    2. Type the user name and password, and then click **Connect**.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Schema drop-down list, select a schema.
    3. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

    Optionally, you can make the following configurations before building the view:

- **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
- **Change data types** – Change the data type of a column by clicking the data type icon.
- **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
- **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
- **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of an EXASolution data source is shown below.

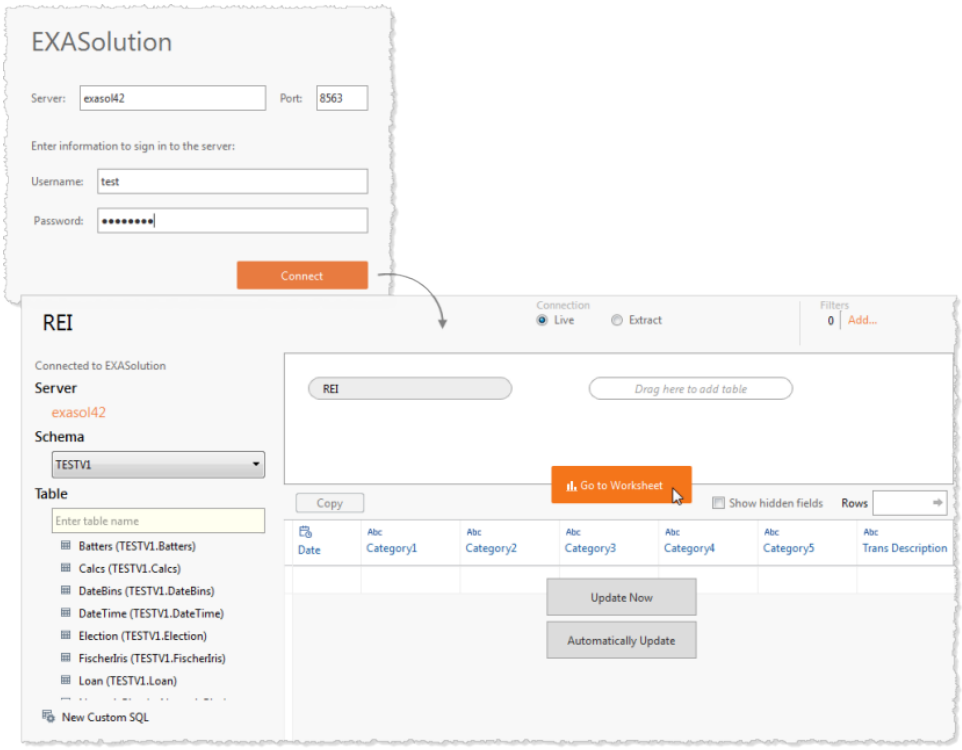

Firebird

This topic describes how to connect Tableau to a Firebird database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **Firebird**, and then do the following:
    1. Type the name of the server that hosts the database, and specify the location of the database.
    2. Type the user name and password, and then click **Connect**.

       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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

       Optionally, you can make the following configurations before building the view:

       - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
       - **Change data types** – Change the data type of a column by clicking the data type icon.
       - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
       - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
       - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of a Firebird data source is shown below.

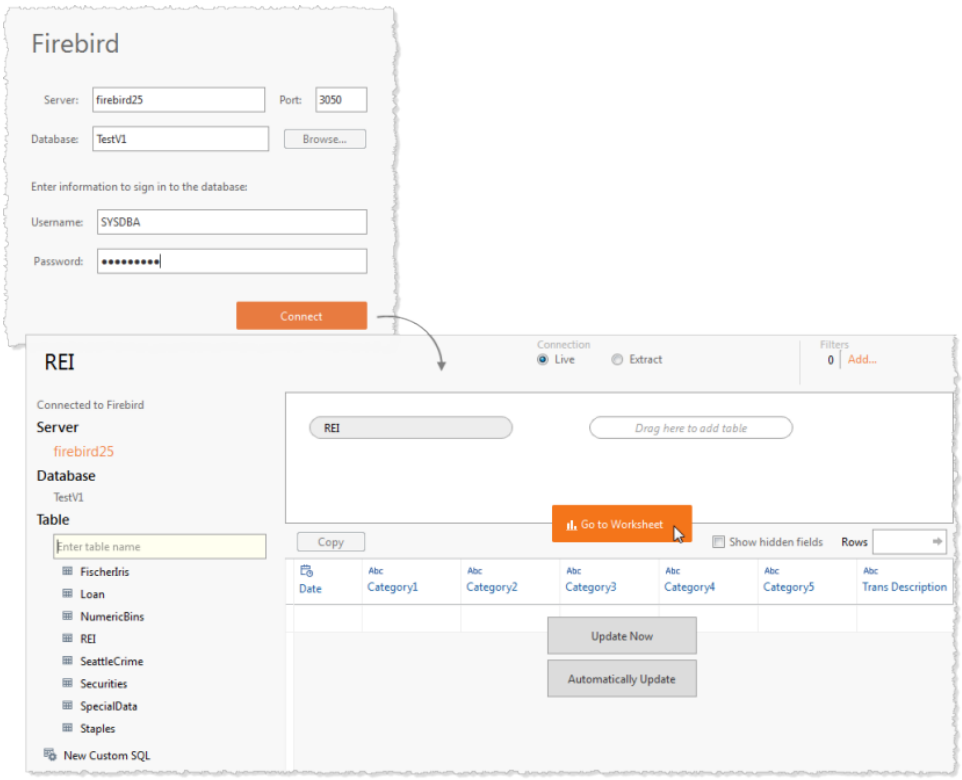

Note: In order to take full advantage of the calculation language in Tableau, you need to load some extra functions into your Firebird database. These functions are located in three .dll files in the \bin\udf folder of the install directory. Load the files using DEFINE FUNCTION commands. You can use 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.

Google Analytics

This topic describes how to connect Tableau to Google Analytics.

1. Click **Connect to Data**..
2. Select **Google Analytics**.
3. Sign in to Google Analytics using your email address and password.

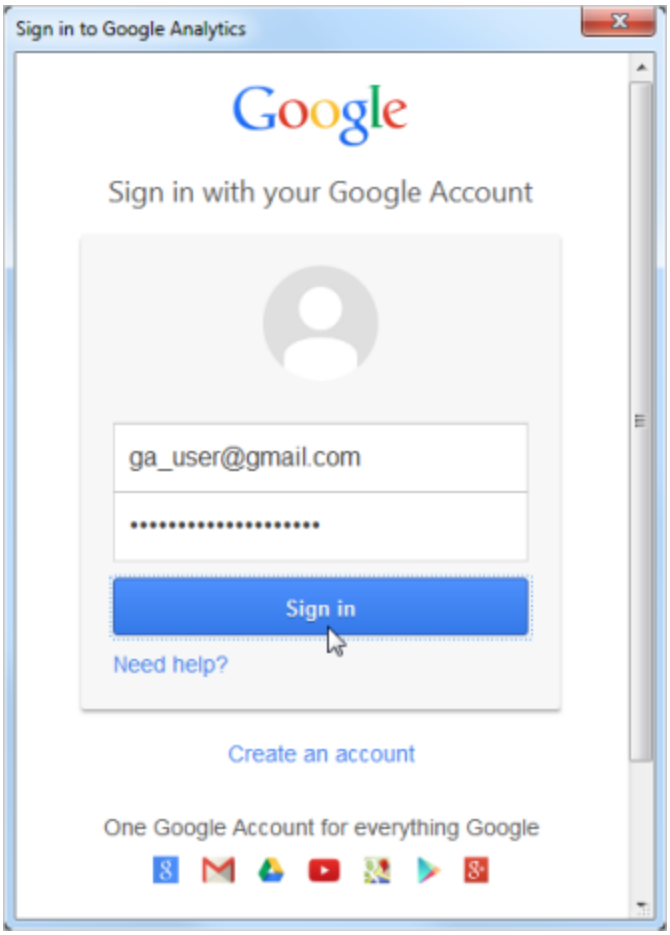

Note: If Google bypasses the sign-in screen and goes directly to the approval screen shown in the next step, it means you are already signed in to a Google account. If you are signed in, make sure you are using the account that provides access to your Google Analytics information.

4. Allow Tableau Desktop permission to access your Google Analytics account.

5. Follow the steps in the Google Analytics Connection dialog box to complete the connection.

   Step 1 – Select an Account, Property, and Profile using the drop-down menus.

   Step 2 – Select a date range. You can select one of the predefined date ranges or select specific dates.

   When selecting a date range, Google Analytics can only provide complete data up to the previous full day. For example, if you choose Last 30 days, data will be retrieved for the last 30-day period ending yesterday.

   Step 3 – Select up to 7 dimensions and 10 measures. To add a dimension or measure, select it from the drop-down menus. Select a predefined set of dimensions and measures using the Measure Group drop-down menu.

   Some dimensions and measures cannot be used together. See the Google Analytics Reference documentation.

   Step 4 – Select a segment to filter your data.

Segments are preset filters that you can set for a Google Analytics connection. Default Segments are defined by Google, and Custom Segments are defined by the user on Google Analytics website. Segments help prevent downsampling. When you attempt to retrieve data without applying a segment, results can often be more than what Google allows in a single call, resulting in random downsampling. So what you see in Tableau will be different from the actual data. With downsampling, you can miss interesting outliers, and aggregations can be inaccurate. With segments, you can get results for a specifc platform, such as tablets, or for a particular search engine, instead of having to download more data than you need and ending up with downsampled data.

Step 5 – Give the connection a name for use in Tableau.

Specify a unique name for the connection. The default name is automatically generated based on the database and table.

A completed Connection dialog box for Google Analytics is shown below.

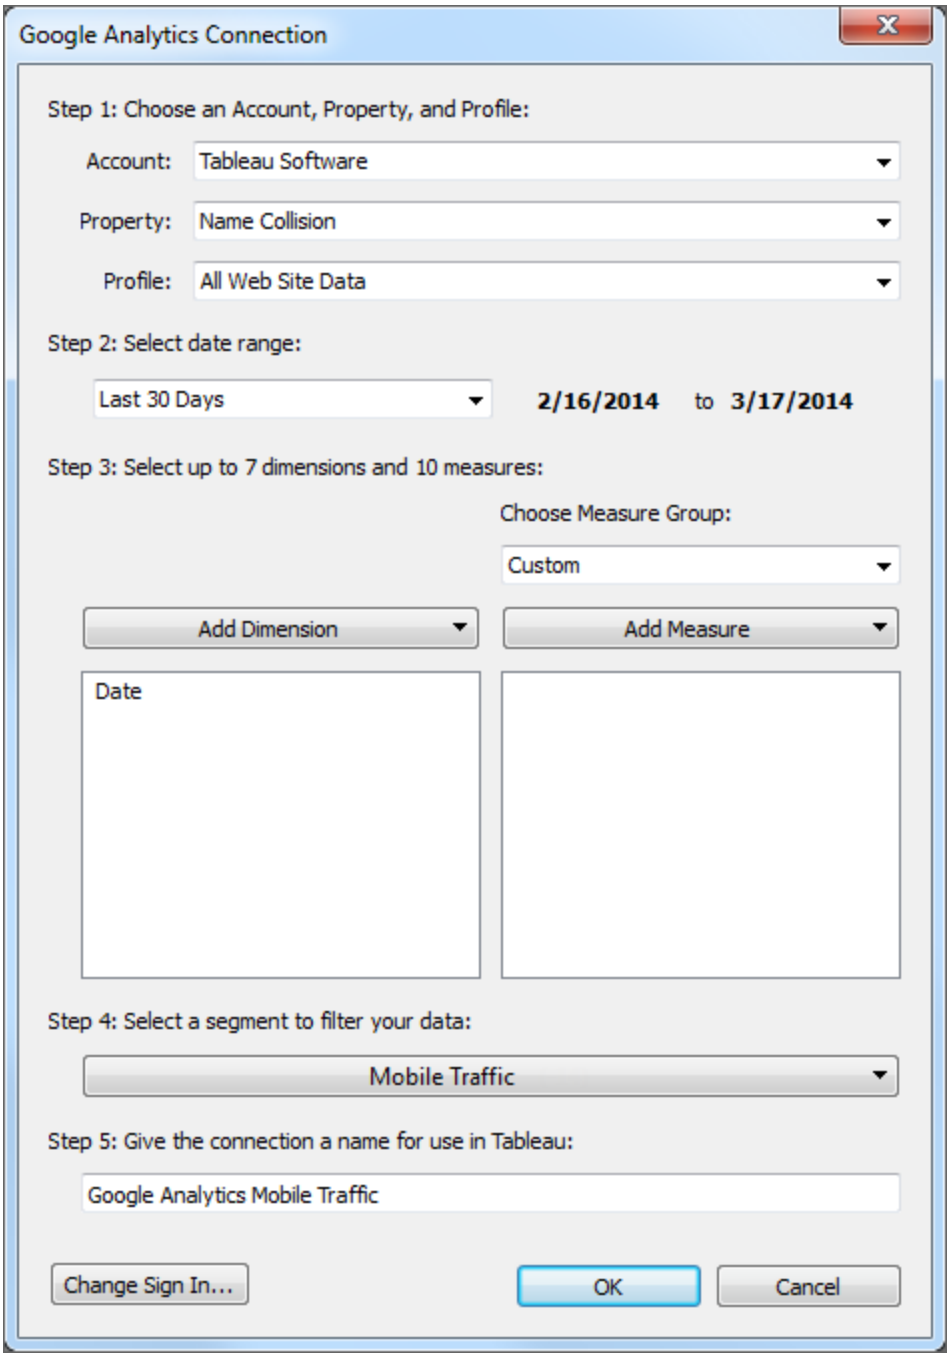

After you click **OK**, Tableau imports the data by creating an extract. You can update the data by refreshing the extract. See **Extracting Data** on page 173 to learn more.

Google BigQuery

This topic describes how to connect Tableau to Google BigQuery and set up the data source.

1. Click **Connect to Data**..
2. On the Connect page, click **Google BigQuery**.
3. Sign in to Google BigQuery using your email address and password.

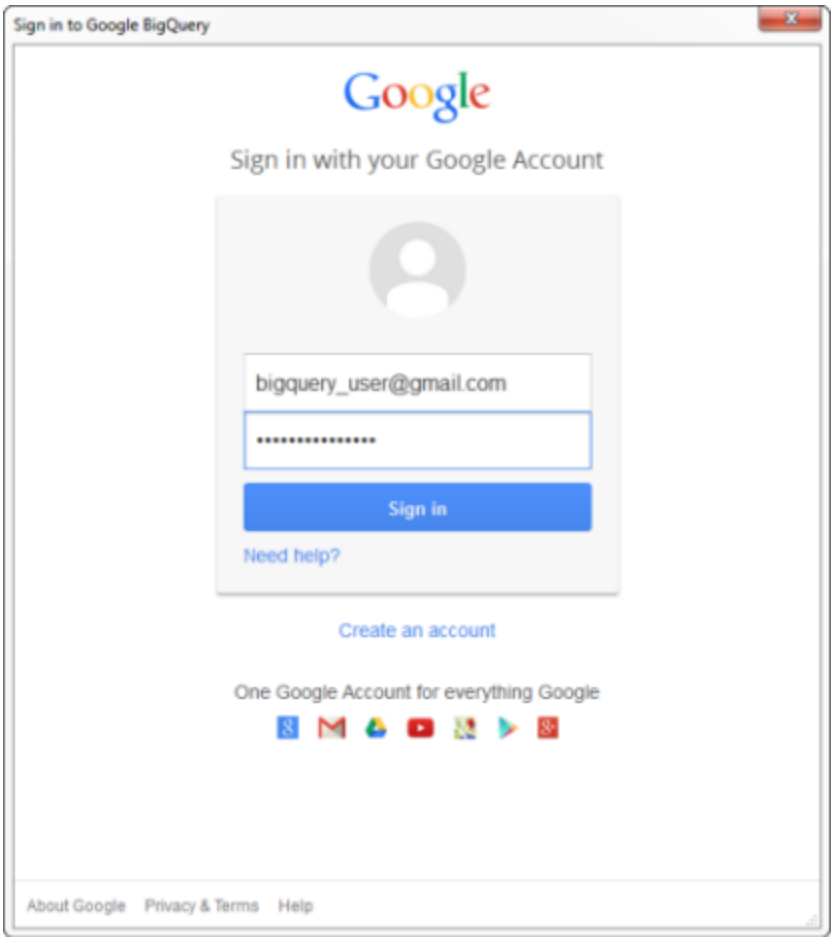

> **Note:** If Google bypasses the sign-in screen and goes directly to the approval screen shown in the next step, it means you are already signed in to a Google account. If you are signed in, make sure you are using the account that provides access to your Google BigQuery data.

4. Allow Tableau Desktop permission to access your Google BigQuery account.

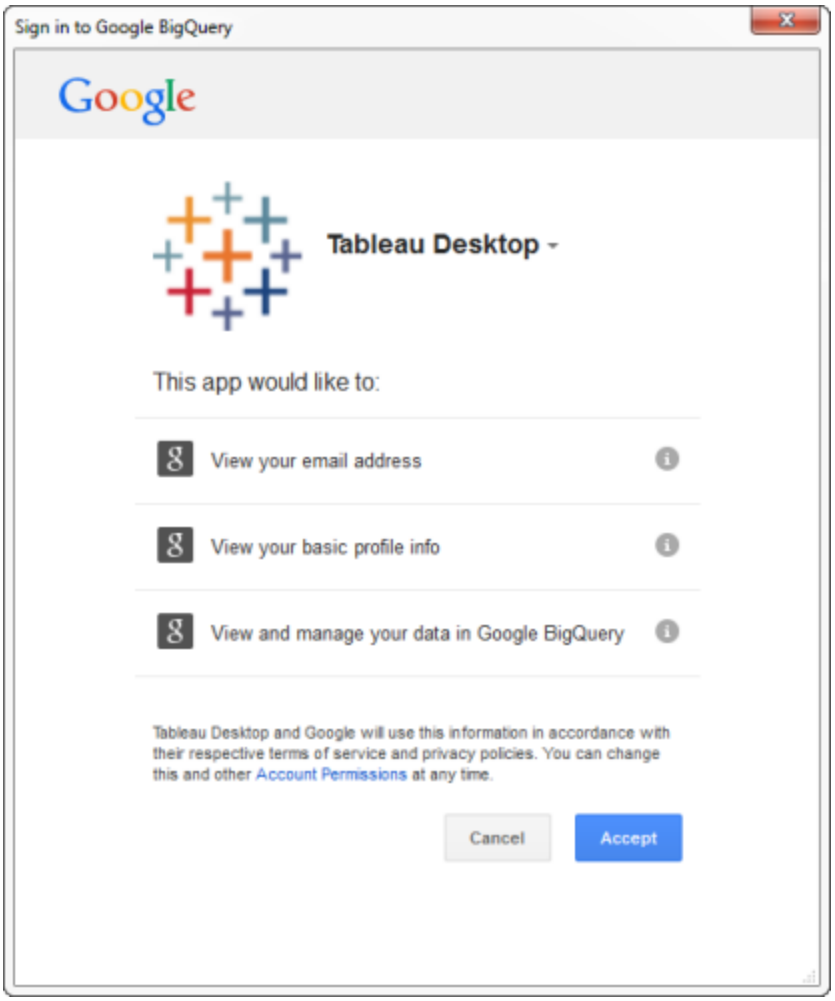

5. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Project drop-down list, select a project. Alternatively, select **publicdata** to connect to sample data in BigQuery.
    3. From the Dataset drop-down list, select a data set.
    4. Drag a table to the join area and click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

       Optionally, you can make the following configurations before building the view:

       - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want

changes to automatically reflect in the preview area, click **Automatically Update**.

- **Change data types** – Change the data type of a column by clicking the data type icon.
- **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
- **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
- **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

> **Note:** Because of the large volume of data in BigQuery, Tableau recommends you connect live.

An example of a Google BigQuery data source is shown below.

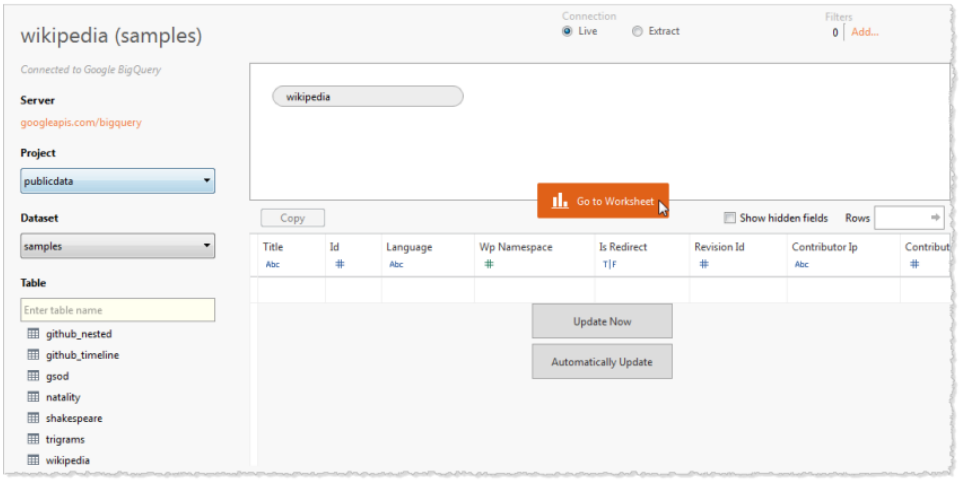

Hortonworks Hadoop Hive

This topic describes how to connect Tableau to a Hortonworks Hadoop Hive database and set up the data source. See the Connecting to Hadoop Hive knowledge base article for more detail about connecting to Hadoop data.

1. Click **Connect to Data**.
2. On the Connect page, click **Hortonworks Hadoop Hive**, and then do the following:
    1. Type the name of server that hosts the data base.
    2. Select the type of database to connect to. You can select one of the following:

- **HiveServer**
- **HiveServer2**
- **Hortonworks Hadoop Hive**

3. Click **Connect**.

   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.

3. On the data source page, do the following:
   1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
   2. From the Schema drop-down list, type the schema name in the text box, click the search icon, and then select the schema.
   3. Type the table name in the Table text box, click the search icon, drag the table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

      Optionally, you can make the following configurations before building the view:

      - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
      - **Change data types** – Change the data type of a column by clicking the data type icon.
      - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
      - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
      - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

Note: When connecting to Hortonworks Hadoop Hive, you can optionally specify a SQL command that will be run once upon connection. See **Initial SQL** on page 142 to learn more about adding these commands to your connection.

An example of a Hortonworks Hadoop Hive data source is shown below.

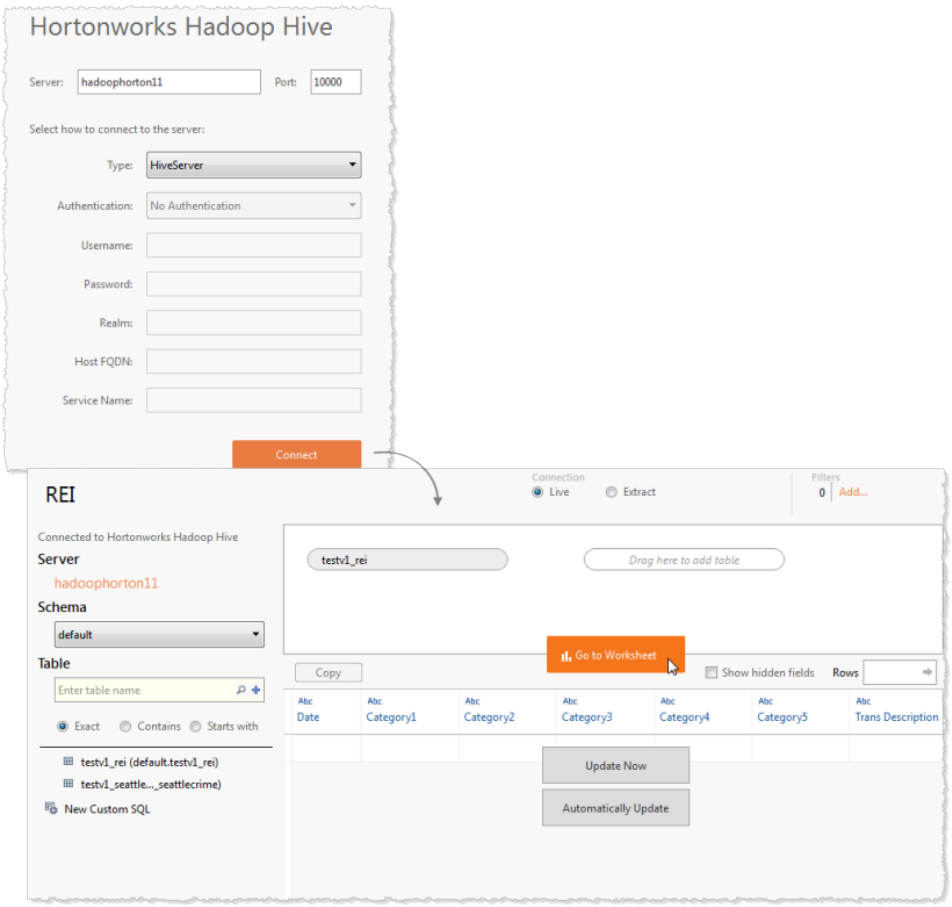

HP Vertica

This topic describes how to connect Tableau to an HP Vertica database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **HP Vertica**, and then do the following:
    1. Type the name of the server that hosts the database and the name of the database that you want to connect to.
    2. Type the user name and password, and then click **Connect**.
    3. Click the **Connect** button to establish the connection.

       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.
3. On the data source page, do the following:

1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
2. From the Schema drop-down list, select a schema.
3. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

   Optionally, you can make the following configurations before building the view:

   - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
   - **Change data types** – Change the data type of a column by clicking the data type icon.
   - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
   - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
   - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

Note: When connecting to HP Vertica databases, you can optionally specify a SQL command that will be run once upon connection. See **Initial SQL** on page 142 to learn more about adding these commands to your connection.

An example of an HP Vertica data source is shown below.

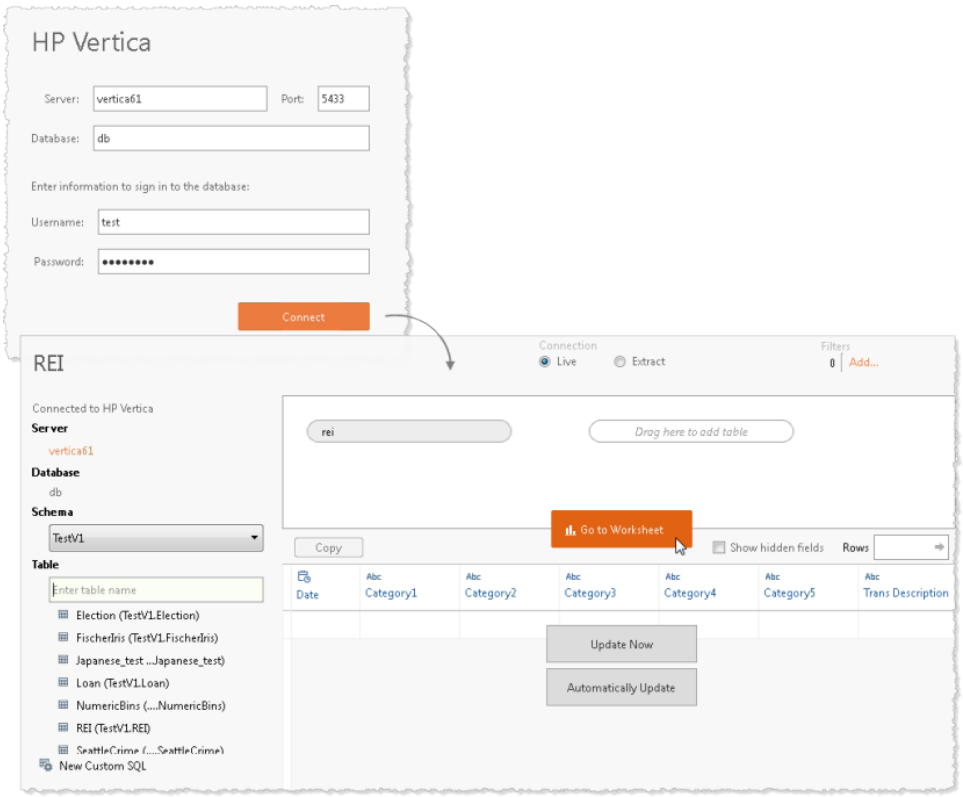

IBM BigInsights

This topic describes how to connect Tableau to an IBM BigInsights database.

1. Click **Connect to Data**.
2. On the Connect page, click **IBM BigInsights**, and then do the following:
    1. Enter the name of the server that hosts the database.
    2. Enter the port number.
    3. Enter the name of the database you want to connect to.
    4. Enter your user name and password, and then click **Connect**.

       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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the **Schema** drop-down list, select a schema or use the text box to search for a schema by name.
    3. Under **Table**, select a table or use the text box to search for a table by name.

4. Drag the table to the join area, and then click **Go to Worksheet** to begin building the view.

   Use custom SQL to filter data when connecting to multiple or large tables. For more information, see **Connecting to a Custom SQL Query** on page 137. For information on connecting to more than one table, see **Joining Tables** on page 168.

   Optionally, you can make the following configurations before building the view:

   - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
   - **Change data types** – Change the data type of a column by clicking the data type icon.
   - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
   - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
   - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

Note: When connecting to IBM BigInsights databases, you can optionally specify a SQL command that will be run once upon connection. See **Initial SQL** on page 142 to learn more about adding these commands to your connection.

An example of an IBM BigInsights data source is shown below.

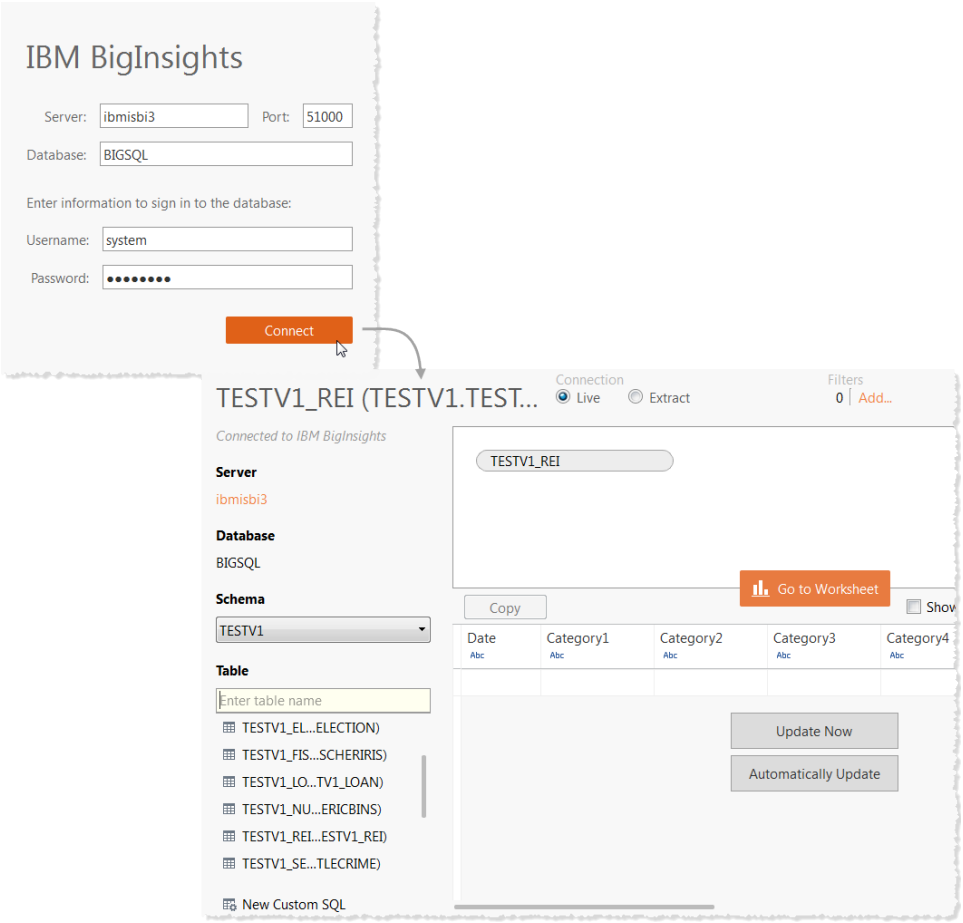

IBM DB2

This topic describes how to connect Tableau to an IBM DB2 database and set up the data source. Refer to the Technical Specifications to confirm which DB2 databases are supported.

1. Click **Connect to Data**.
2. On the Connect page, click **IBM DB2**, and then do the following:
    1. Type the name of the server that hosts the database and the name of the database that you want to connect to.

       The port is dependent on the type of server you are connecting to and whether you are connecting to an encrypted port. Generally, use 50000 for a non-encrypted port and 60000 for an encrypted port. It is possible that your server is configured to use a non-standard port. Contact your administrator if you don't know which port to connect to.
    2. Type your user name and password, and then click **Connect**.

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.

3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

       Optionally, you can make the following configurations before building the view:

       - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
       - **Change data types** – Change the data type of a column by clicking the data type icon.
       - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
       - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
       - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

Note: When connecting to IBM DB2 databases, you can optionally specify a SQL command that will be run once upon connection. See **Initial SQL** on page 142 to learn more about adding these commands to your connection.

An example of an IBM DB2 data source is shown below.

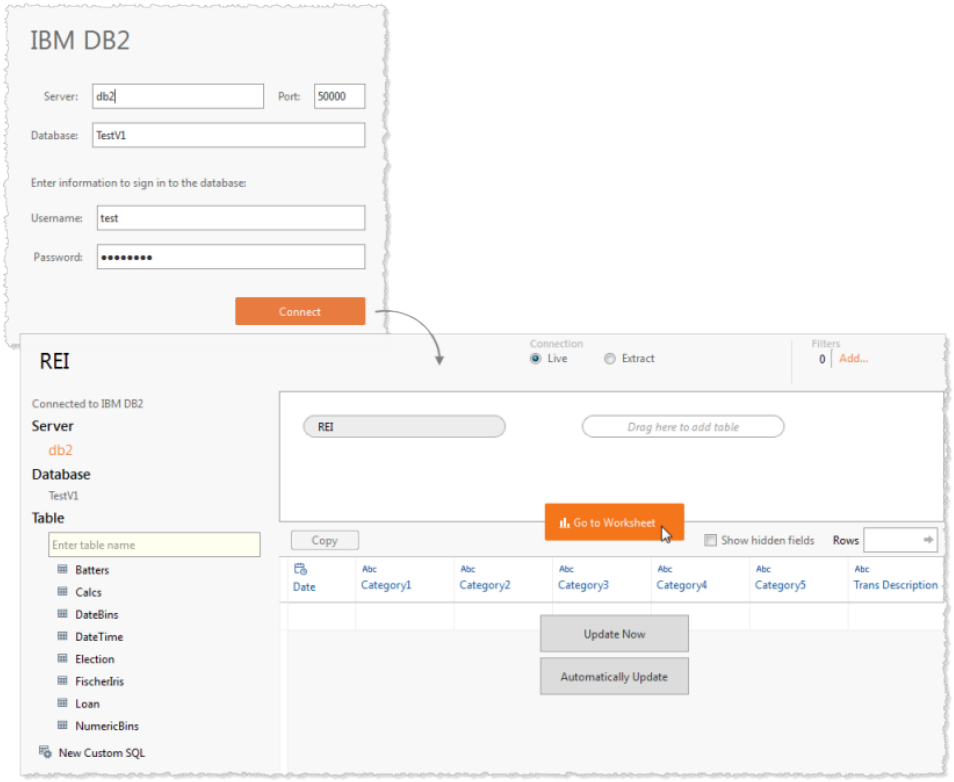

IBM Netezza

This topic describes how to connect Tableau to an IBM Netezza database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **IBM Netezza**, and then do the following:
    1. Type the name of the server that hosts the database and the name of the database that you want to connect to.
    2. Type the user name and password, and then click **Connect**.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Schema drop-down list, select a schema.
    3. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

    Optionally, you can make the following configurations before building the view:

- **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
- **Change data types** – Change the data type of a column by clicking the data type icon.
- **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
- **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
- **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

And example of an IBM Netezza data source is shown below.

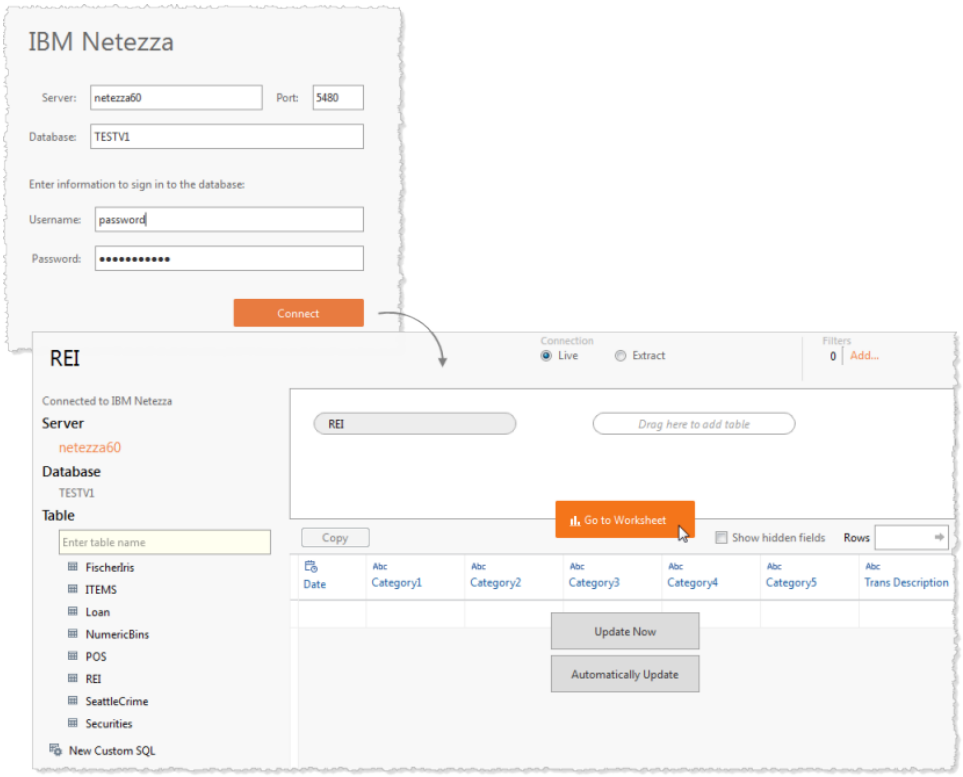

MapR Hadoop Hive

This topic describes how to connect Tableau to a MapR Hadoop Hive database and set up the data source. Refer to the Connecting to Hadoop Hive knowledge base article for more detail about connecting to Hadoop data.

1. Click **Connect to Data**.
2. On the Connect page, click **MapR Hadoop Hive**, and then do the following:
    1. Select the type of database to connect to. You can select one of the following:
        - **HiveServer**
        - **HiveServer2**
    2. Type the name of the server that hosts the database, and then click **Connect**.

       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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Schema drop-down list, type the schema name in the text box, click the search icon, and then select the schema.
    3. Type the table name in the Table text box, click the search icon, drag the table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

       Optionally, you can make the following configurations before building the view:
        - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
        - **Change data types** – Change the data type of a column by clicking the data type icon.
        - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
        - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.

- **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

> **Note** When connecting to MapR Hadoop Hive, you can optionally specify a SQL command that will be run once upon connection. See **Initial SQL** on page 142 to learn more about adding these commands to your connection.

An example of a MapR Hadoop Hive data source is shown below.

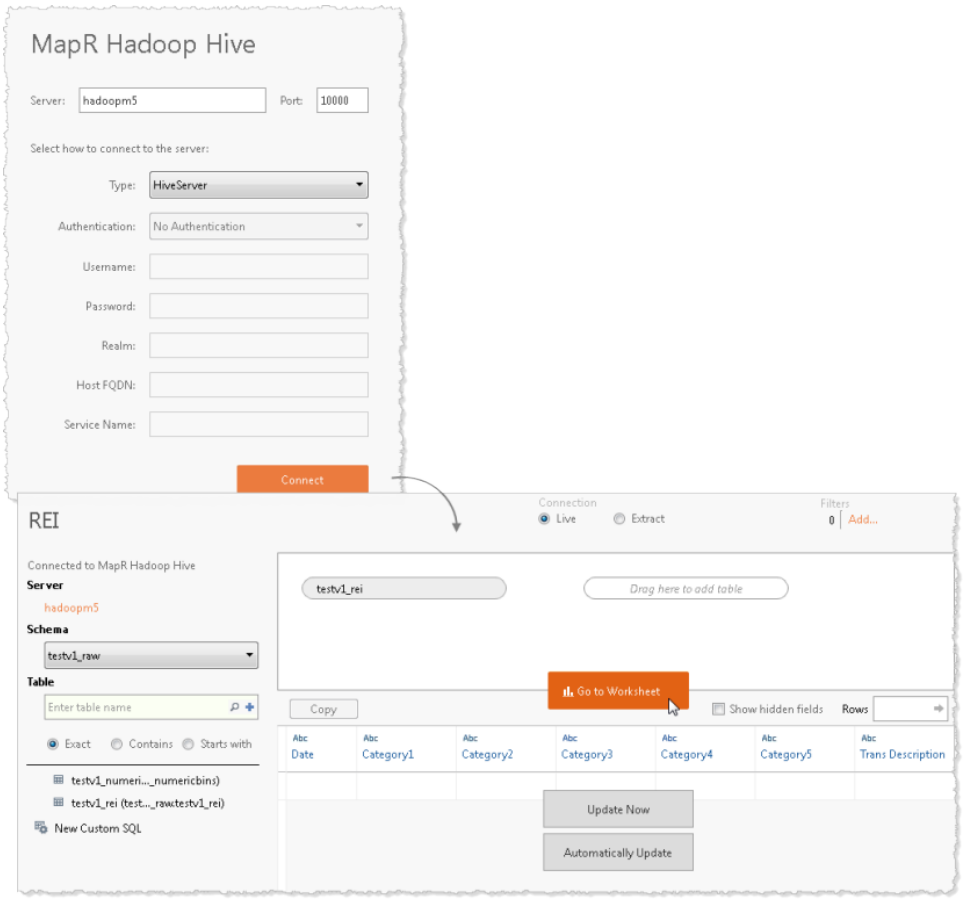

MarkLogic

This topic describes how to connect Tableau to a MarkLogic database.

1. Click **Connect to Data**.
2. On the Connect page, click **MarkLogic**, and then do the following:
    1. Enter the name of the server that hosts the database you want to connect to.
    2. Enter the port number for the ODBC server process of the database you want to

connect to.

3. Enter your user name and password, and then click **Connect**.

   If the connection is unsuccessful, verify that your user name and password are correct and that the port number correctly identifies the MarkLogic database configured by your database administrator. If the connection continues to fail, your computer is having trouble locating the server. Contact your network administrator or database administrator.

3. On the data source page, do the following:

   1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
   2. From the **Schema** drop-down list, select a schema or use the text box to search for a schema by name.
   3. Under **Table**, select a table or use the text box to search for a table by name.
   4. Drag the table to the join area, and then click **Go to Worksheet** to begin building the view.

      Use custom SQL to filter data when connecting to large tables. For more information, see **Connecting to a Custom SQL Query** on page 137. For information on connecting to more than one table, see **Joining Tables** on page 168.

      Optionally, you can make the following configurations before building the view:

      - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
      - **Change data types** – Change the data type of a column by clicking the data type icon.
      - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
      - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
      - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of a MarkLogic data source is shown below.

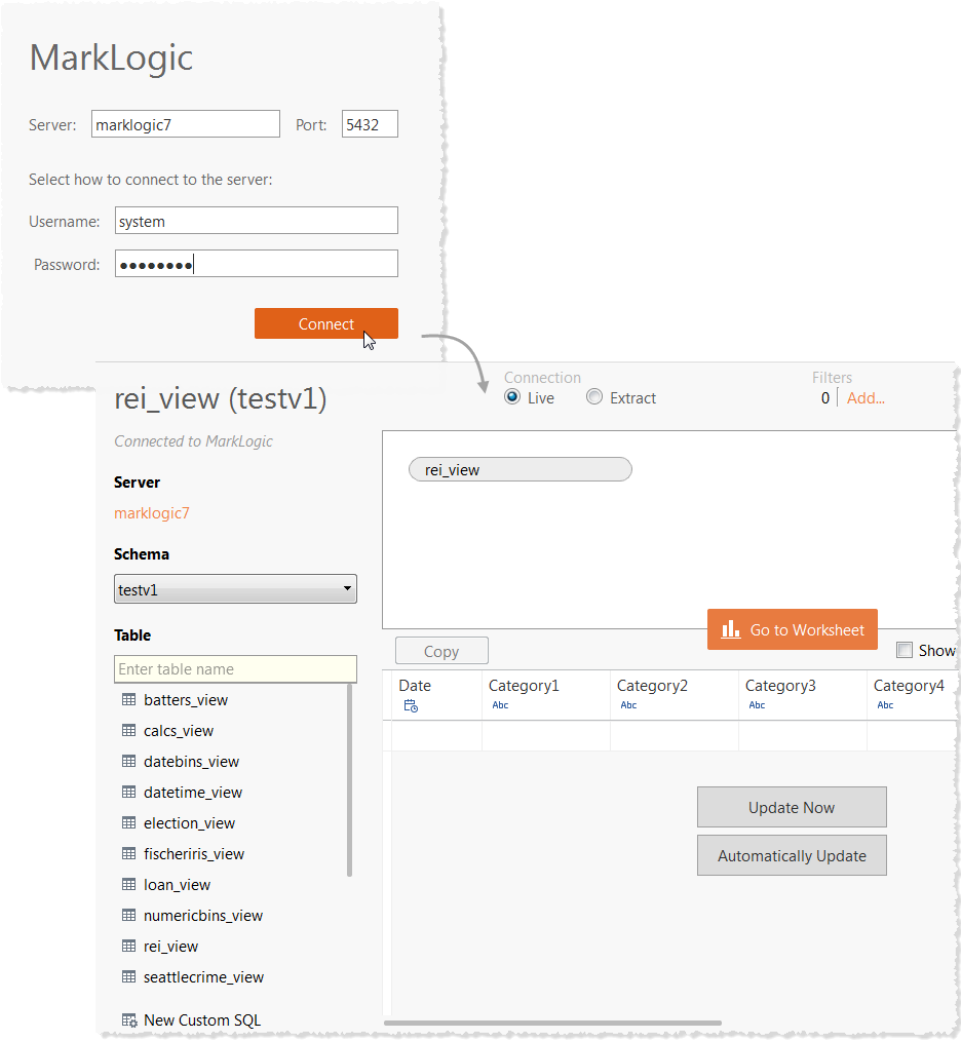

Microsoft Analysis Services

This topic describes how to connect Tableau to a Microsoft Analysis Services database.

1. Click **Connect to Data**.
2. Click **Microsoft Analysis Services**.
3. Follow the steps in the Microsoft Analysis Services Connection dialog box to complete the connection.

   Step 1 – Select whether to connect to a cube file on the server or locally and then locate the cube you want to connect to.

   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 sign-in information for the server.

Specify whether to use Windows Authentication or a specific user name and password. If the cube is password protected, and you are not in a Kerberos environment, you must enter your user name and password.

Step 3 – Establish the connection.

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 4 – Use the drop-down list to select the database that you want to connect to.

Step 5 – Select a cube from the database.

Microsoft Analysis Services databases can contain multiple tables and views. Specify which table or view within the database you want to connect to.

Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

Step 6 – Give the connection a name for use in Tableau.

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

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

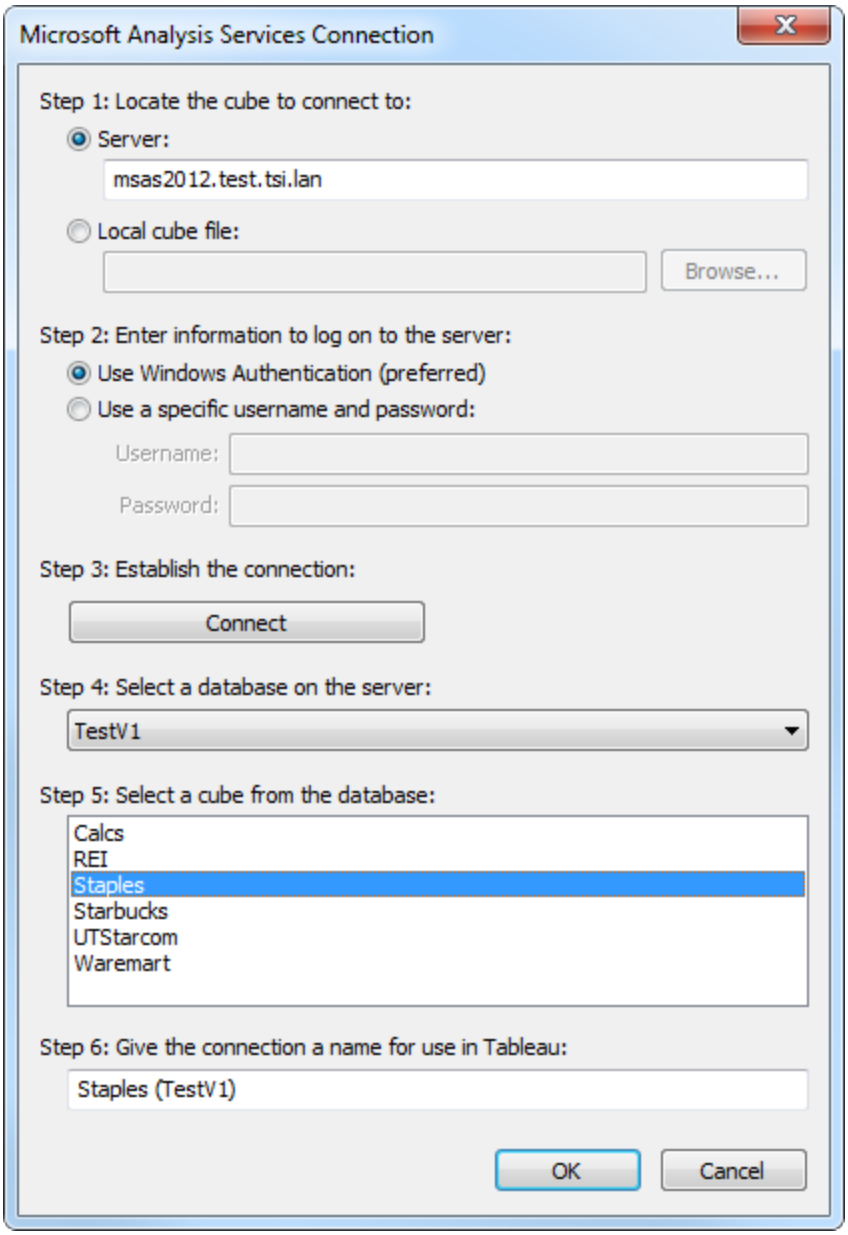

If you are connecting to a Microsoft Analysis Services 2000 data source, you are given the option to enable NonEmptyCrossjoin Optimization. When selected this option ensures that virtual dimensions are shown.

Microsoft PowerPivot

This topic describes how to connect Tableau to a Microsoft PowerPivot database.

1. Click **Connect to Data**..
2. Click **Microsoft PowerPivot**.

3. Follow the steps in the Microsoft PowerPivot Connection dialog box to complete the connection.

   Step 1 – Select whether to connect to a PowerPivot file on SharePoint or a local Excel file.

   If you are connecting to SharePoint type the URL or specify a UNC File Path. You can use the Browse button to select the file. If you are connecting to a local Excel file, click the Browse button to navigate to the file you want to connect to.

   If you plan to publish the workbook to Tableau Server, make sure to connect to a PowerPivot file on SharePoint. Tableau Server does not support connections to local PowerPivot files.

   Step 2 – Establish the connection.

   Step 3 – Select a perspective.

   PowerPivot data sources may contain one or more perspectives. Perspectives are subsets of objects form the model that define sets of data. Typically perspectives are defined for a particular group of users or business scenario. Using perspectives can make it easier to navigate large data sources.

   Step 4 – Give the connection a name for use in Tableau.

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

A completed Connection dialog box for Microsoft PowerPivot is shown below.

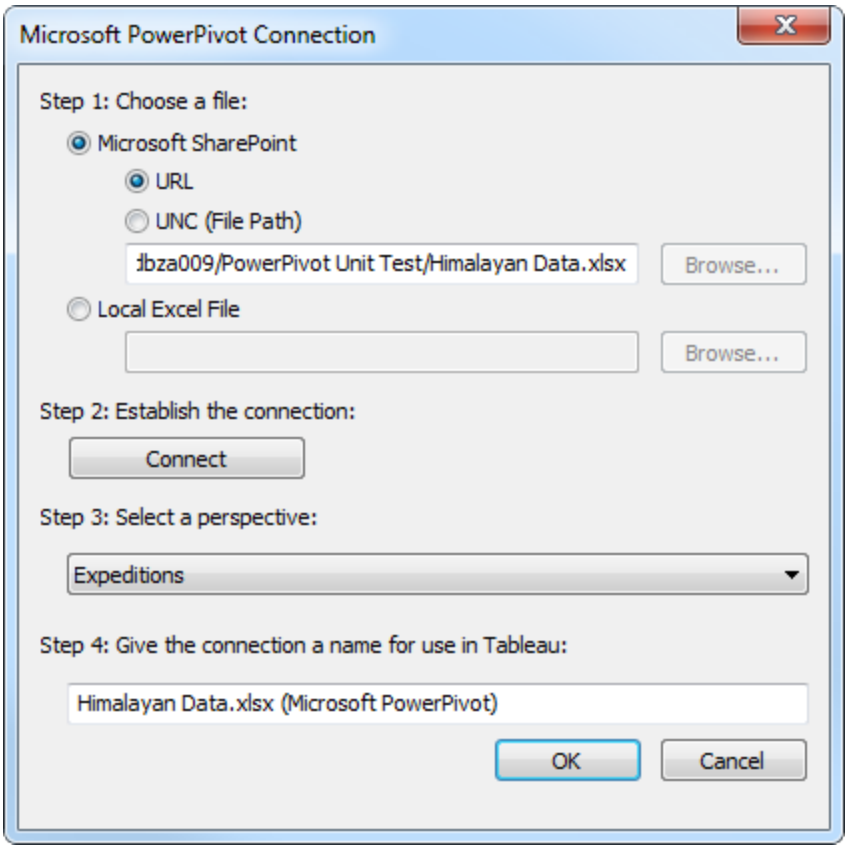

Microsoft SQL Server

This topic describes how to connect Tableau to a Microsoft SQL Server database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **Microsoft SQL Server**, and then do the following:
    1. Enter the name of the server you want to connect to.
    2. Select how you want to sign in to the server. Specify whether to use Windows Authentication or a specific user name and password. If the server is password protected, and you are not in a Kerberos environment, you must enter the user name and password.
    3. Specify whether to **Read uncommitted data**. This option sets the database isolation level to Read Uncommitted. Long queries from Tableau, including extract refreshes, can lock the database and delay transactions. Select this option to allow queries to read rows that have been modified by other transactions even when they have not been committed yet. When this option is cleared, Tableau uses the default isolation level specified by the database.
    4. Click **Connect**.

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.

3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Database drop-down list, select a database.
    3. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

       You can also specify a stored procedure in the database. See **Stored Procedures** on page 143 for details and for a list of constraints specific to SQL Server databases.

       Optionally, you can make the following configurations before building the view:

       - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
       - **Change data types** – Change the data type of a column by clicking the data type icon.
       - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
       - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
       - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of a Microsoft SQL Server data source is shown below.

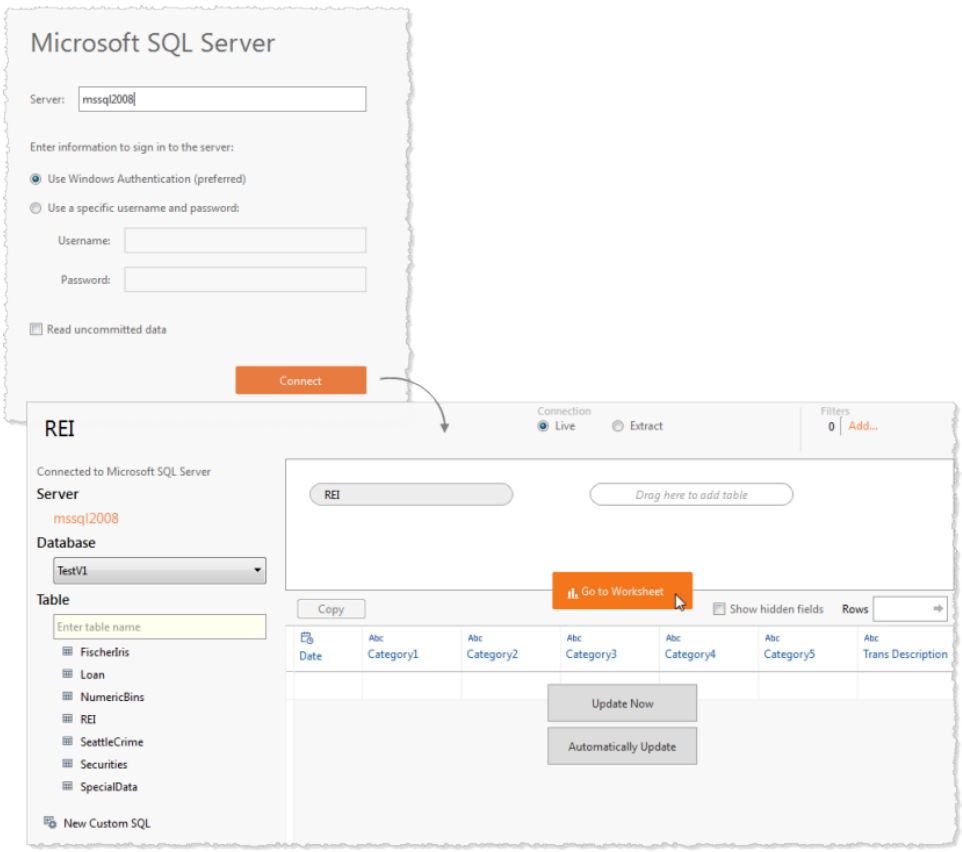

Note: Tableau Desktop does not support the Microsoft SQL Server TIME data type. Fields of this data type are not imported and do not appear in Tableau Desktop. If included in stored procedures, TIME data type fields will not appear in Tableau Desktop. For more information on stored procedures, see **Stored Procedures** on page 143.

MySQL

This topic describes how to connect Tableau to a MySQL database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **MySQL**, and then do the following:
    1. Type the name of the server that hosts the database.
    2. Type the user name and password, and then click **Connect**.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the **Database** drop-down list, select a database.
    3. Drag a table to the join area, and then click **Go to Worksheet**to begin building the

view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

Optionally, you can make the following configurations before building the view:

- **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
- **Change data types** – Change the data type of a column by clicking the data type icon.
- **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
- **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
- **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of a MySQL data source is shown below.

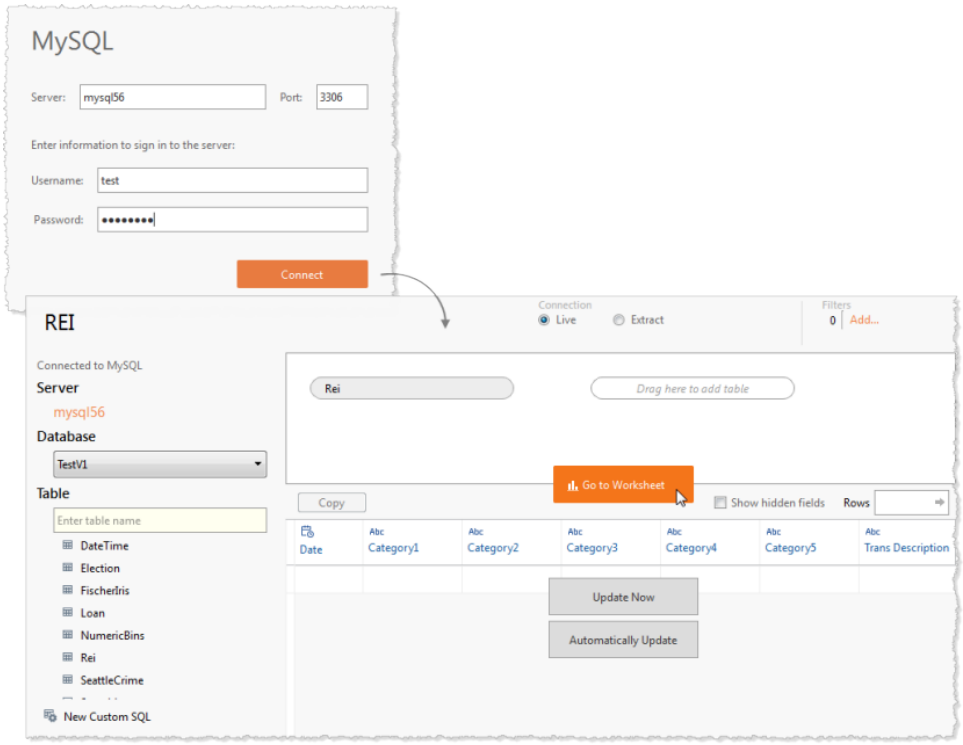

OData

This topic describes how to connect Tableau to an OData data source.

1. Click **Connect to Data**.
2. Click **OData**.
3. Follow the steps in the OData Connection dialog box to complete the connection.

   Step 1 – Select or enter a URL to the data you want to connect to. 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 – Establish the connection.

   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. A default name is automatically generated.

A completed Connection dialog box for OData is shown below.

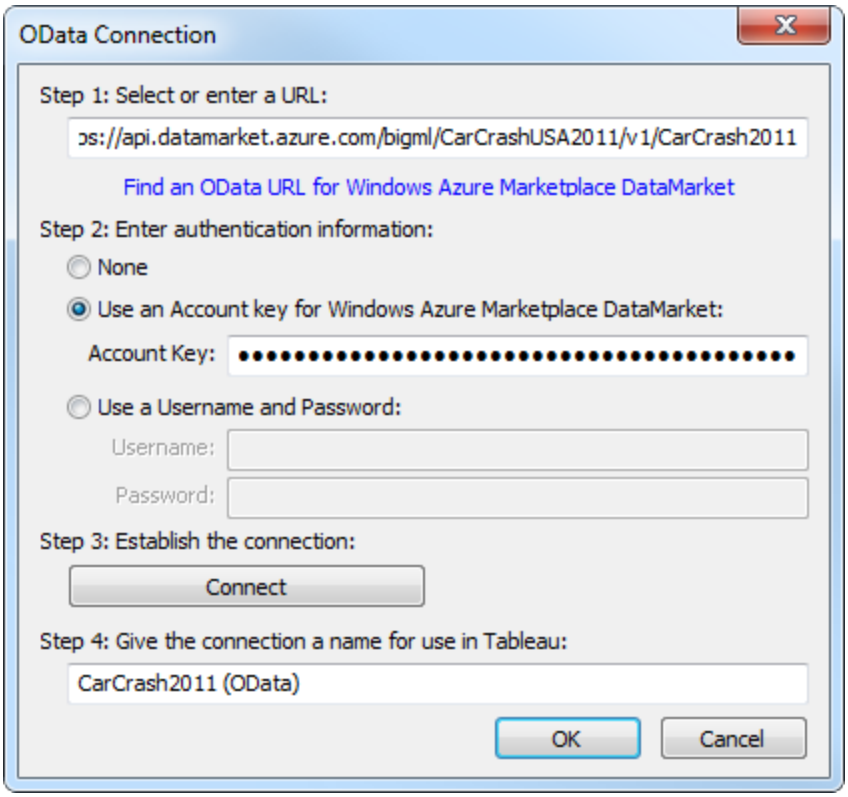

After you connect to an OData data source, you will be prompted to Import the data into Tableau's data engine (see **Extracting Data** on page 173). Unlike many extract connections, extracts from OData data sources cannot be refreshed. To get the updated data, you must reconnect to the data source.

Oracle

This topic describes how to connect Tableau to an Oracle database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **Oracle**, and then do the following:
    1. Type the server name. Optionally, you can specify the Oracle service name and port.
    2. Select how you want to sign in to the server. Specify whether to use Windows Authentication or a specific user name and password. If the server is password protected, you must type the user name and password.
    3. Click **Connect**.

    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.

3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Schema drop-down list, select a schema.
    3. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

       Optionally, you can make the following configurations before building the view:

       - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
       - **Change data types** – Change the data type of a column by clicking the data type icon.
       - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
       - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
       - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of an Oracle data source is shown below.

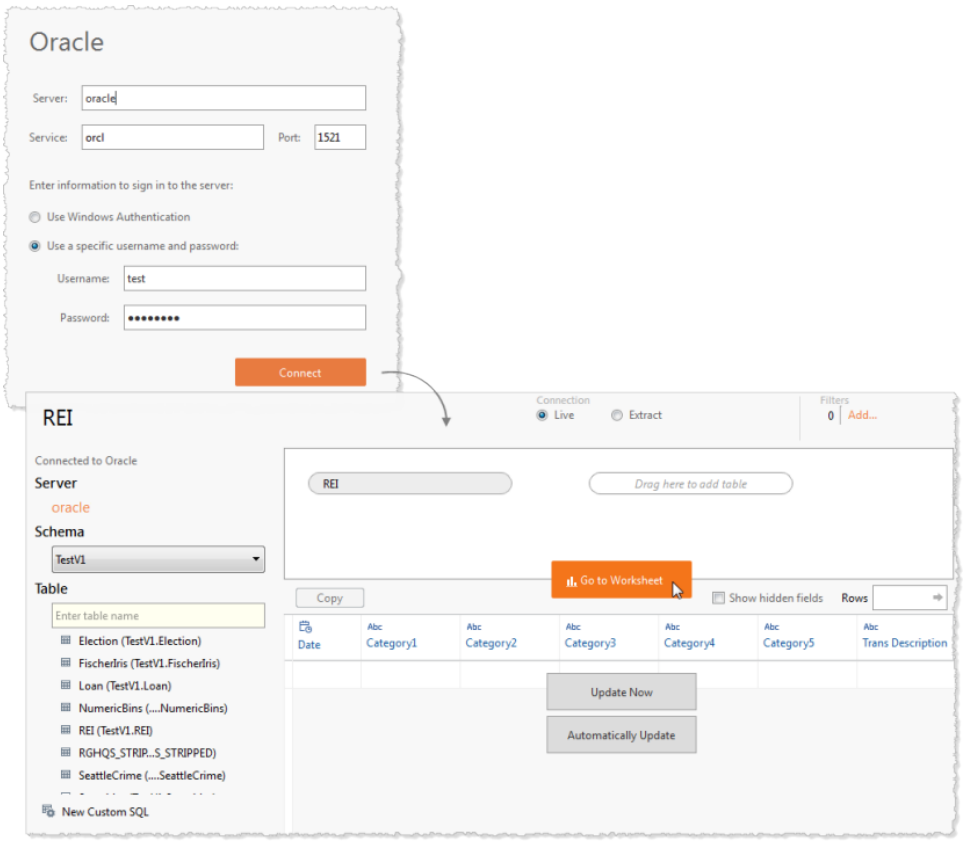

In order to use your net services definitions in Tableau, you must set either TNS_ADMIN or ORACLE_HOME as an environment variable. To set TNS_ADMIN as the environment variable use the full path of the directory that contains the tnsnames.ora file. To set ORACLE_Home as an environment variable use the path to the main Oracle directory.

Oracle Essbase

This topic describes how to connect Tableau to an Oracle Essbase database.

1. Click **Connect to Data**.
2. Click **Oracle Essbase**.
3. Follow the steps in the Oracle Essbase Connection dialog box to complete the connection.

   Step 1 – Type the name of the server that hosts the database.

   Step 2 – Enter your user name and password to sign in to the server.

   Step 3 – Establish the connection.

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 4 – Use the drop-down list to select an application on the server that you want to connect to.

Step 5 – Select a database from the application.

Oracle Essbase databases can contain multiple tables and views. Specify which table or view within the database you want to connect to.

Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

Step 6 – Give the connection a name for use in Tableau.

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

A completed Connection dialog box for Oracle Essbase is shown below.

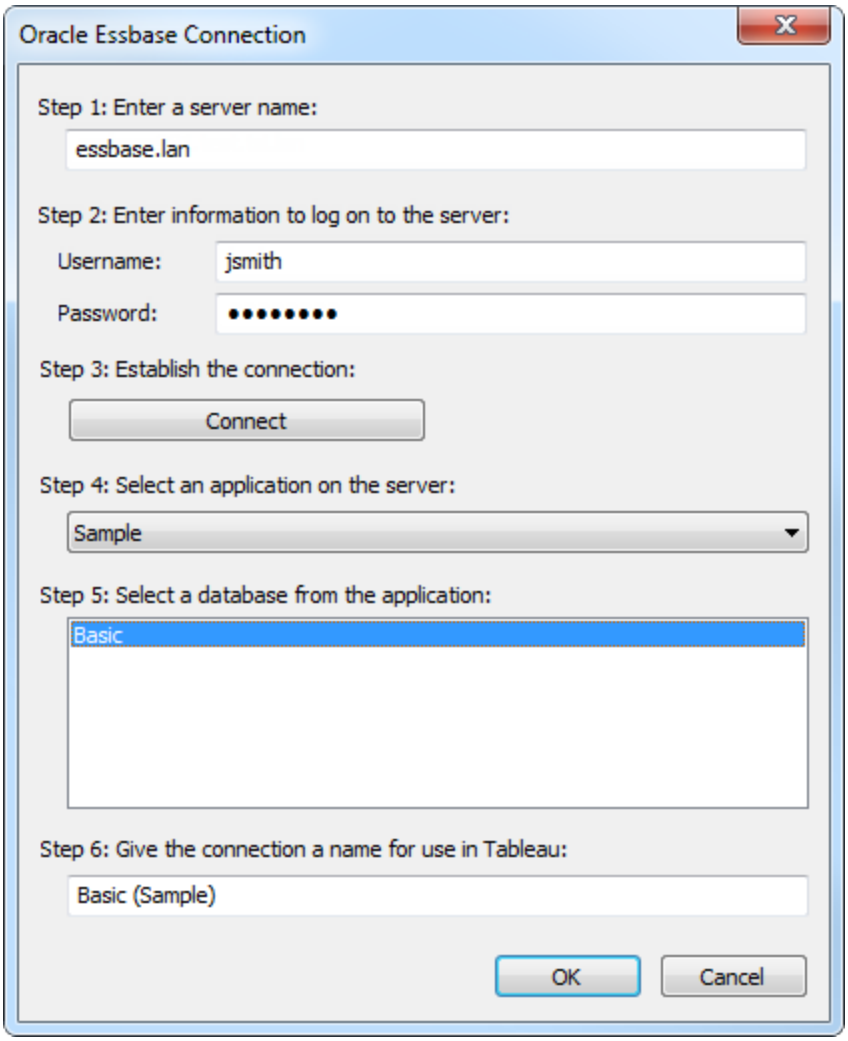

Set an Accounts Dimension

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.

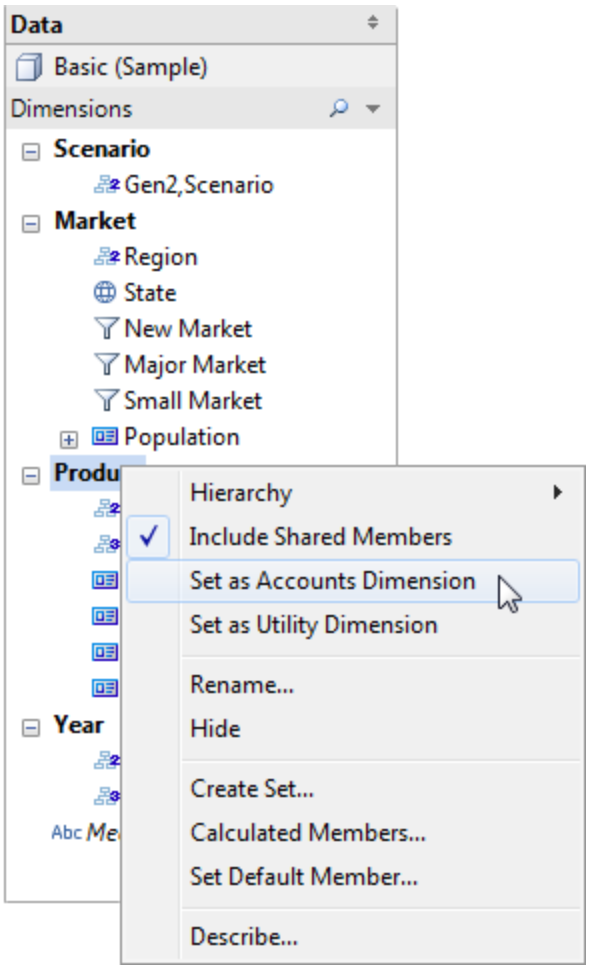

ParAccel

This topic describes how to connect Tableau to a ParAccel database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **ParAccel**, and then do the following:
    1. Type the name of the server that hosts the database and the name of the database that you want to connect to.
    2. Type the user name and password, and then click **Connect**.

       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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.

2. From the Schema drop-down list, select a schema.
3. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

   Optionally, you can make the following configurations before building the view:

   - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
   - **Change data types** – Change the data type of a column by clicking the data type icon.
   - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
   - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
   - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of a ParAccel data source is shown below.

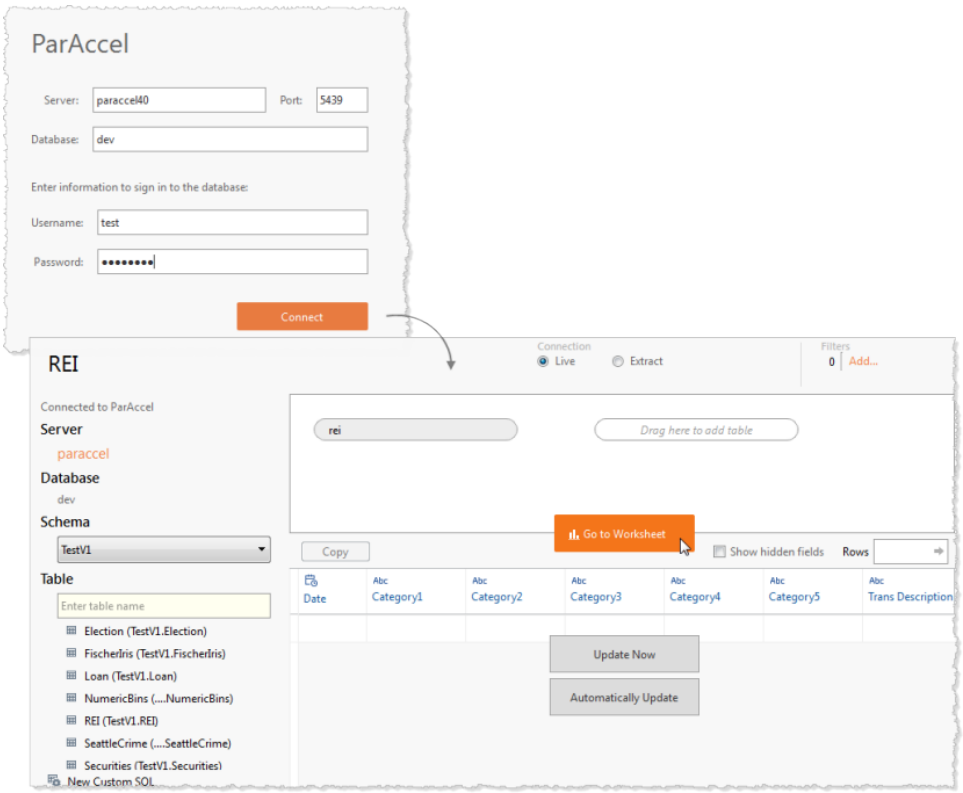

PostgreSQL

This topic describes how to connect Tableau to a PostgreSQL database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **PostgreSQL**, and then do the following:
    1. Type the name of the server that hosts the database and the name of the database that you want to connect to.
    2. Type the user name and password, and then click **Connect**.

       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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

Optionally, you can make the following configurations before building the view:

- **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
- **Change data types** – Change the data type of a column by clicking the data type icon.
- **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
- **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
- **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of a PostgreSQL data source is shown below.

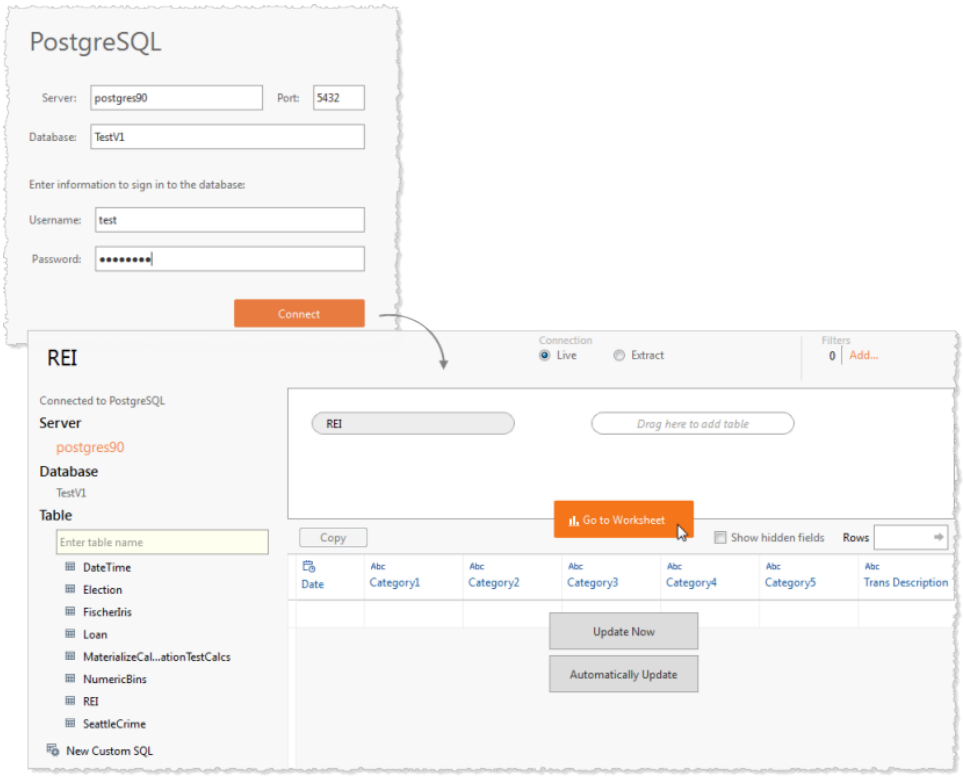

Progress OpenEdge

This topic describes how to connect Tableau to a Progress OpenEdge database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **Progress OpenEdge**, and then do the following:
    1. Type the name of the server that hosts the database and the name of the database you want to connect to.
    2. Type the user name and password, and then click **Connect**.

       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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

       Optionally, you can make the following configurations before building the view:

       - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
       - **Change data types** – Change the data type of a column by clicking the data type icon.
       - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
       - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
       - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of a Progress OpenEdge data source is shown below.

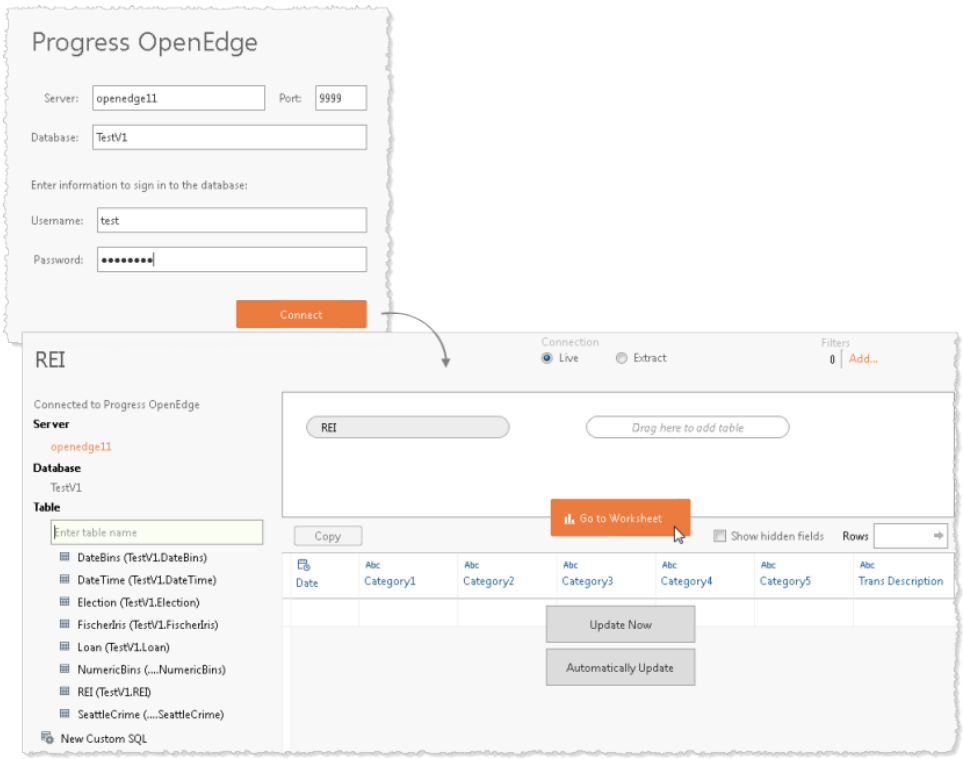

Salesforce

This topic describes how to connect Tableau to Salesforce.com data.

1. Click **Connect to Data**.
2. Click **Salesforce**.
3. Follow the steps in the Salesforce Connection dialog box to complete the connection.

   Step 1 – Sign in to Salesforce.

   Type your username and password for Salesforce.com. Depending on how your Salesforce server is set up, you may need to append a special security token to the end of your password. See **Security Token Requirements** on page 119 below. Also type the URL for the Salesforce server you are connecting to. By default, the server is `login.salesforce.com`.

   Step 2 – Establish the connection.

   If the connection is unsuccessful, verify that the server URL and authentication information are correct. If the connection continues to fail, your computer is having trouble locating the server or you may not have permission to access the data. Contact your network administrator or Salesforce administrator.

   Step 3 – Choose a query.

   You can select from a list of predefined **Standard** queries, which represent commonly used objects. Alternatively, you can connect to a **Single Table**, which

includes any custom objects your company has created. Finally, you can connect to **Multiple Tables** using relational joins. See **Joining Tables** on page 168. Note, though, that only Left and Inner joins are supported for Salesforce connections. This constraint also applies to extracts created from Salesforce connections.

Step 4– Give the connection a name for use in Tableau.

Specify a unique name for the connection. The default name is automatically generated based on the database and table.

A completed Connection dialog box for Salesforce is shown below.

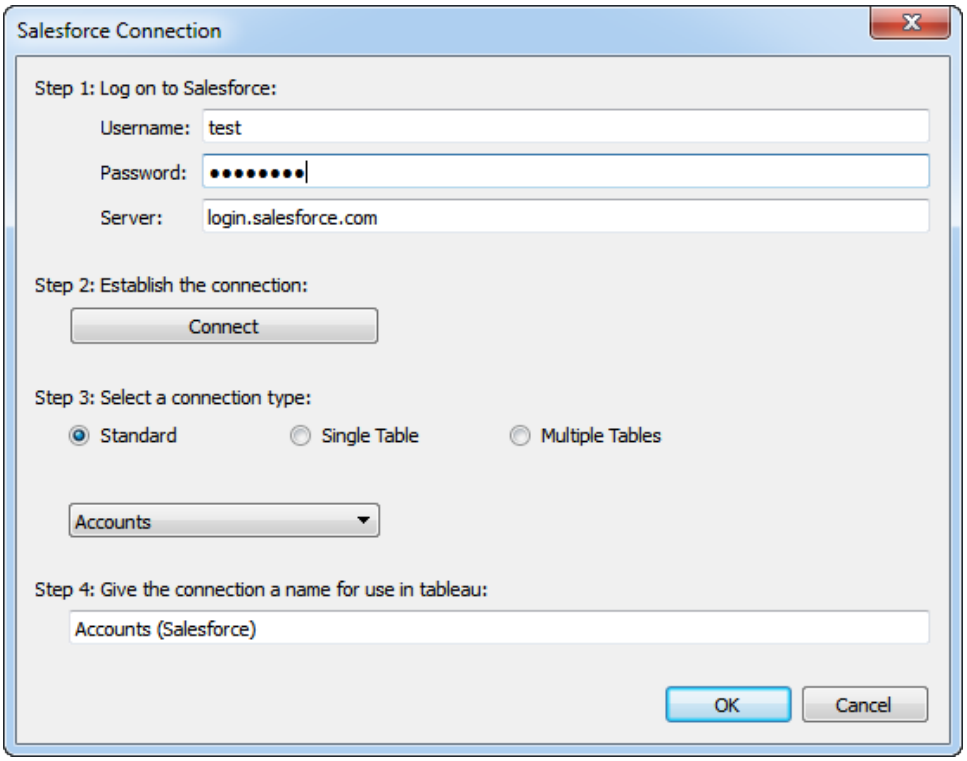

After you click **OK**, Tableau imports the data by creating an extract. You can update the data by refreshing the extract. See **Extracting Data** on page 173 to learn more.

> The initial extract may take some time depending on the amount of data that is included. After the initial extract, you can incrementally update objects by refreshing the extract.

Troubleshooting Salesforce Connections

Tableau leverages the various force.com APIs to connect to Salesforce.com, Force.com, and Database.com data. These services have some restrictions on the type of data and the amount of data that you can access at a given time. If you are having trouble connecting, the following list of common restrictions may help you find a solution.

Company Account and User Profile Configuration

In order for Tableau to connect to Salesforce data, all of the following must be enabled on both your company's account and your user profile:

- SOAP API for signing in.
- REST API for getting meta data.
- BULK API for downloading objects.
- REST API for downloading objects that the BULK API does not support.
- Replication SOAP APIs for retrieving changes in the data.

In order to optimize for performance and ensure the APIs are available for all of their customers, Salesforce.com balances the load by limiting the number of concurrent API requests as well as limiting the total number of API requests overall. An error occurs if these limits are reached while connected using Tableau. See the Salesforce.com developer documentation to learn more about these limits.

> Some editions of Salesforce may not allow API access at all. Contact your Salesforce administrator to verify that your account has access to the above list of APIs.

Security Token Requirements

Some security models may require you to append a special security token to the end of your password when signing in. This token is usually required if your IP Address of your computer is not listed in the trusted range specified by a Salesforce administrator. If you receive an error stating that you must include a security token, make sure to type your password followed by the security token (without spaces) in the Salesforce Connection dialog box.

Errors During Extract

When you connect to Salesforce using Tableau, the data is automatically extracted into a Tableau Data Extract File. In some cases, certain fields cannot be extracted because of character limits. Specifically, text fields that are greater than 4096 characters and calculated fields will not be included in the extract. If you have calculated fields in your data, you will need to recreate them in Tableau after creating the extract.

In addition, the Force.com API restricts queries to 10,000 total characters.

Splunk

This topic describes how to connect Tableau to Splunk data.

1. Click **Connect to Data**.
2. Click **Splunk**.
3. Follow the steps in the Splunk Connection dialog box to complete the connection.

   Step 1 – Enter the server URL.

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

Type the name of a specific database on the server. Database names are case sensitive.

Step 2 – Enter sign-in information for the server.

Type your user name and password.

Step 3 – Establish the connection.

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 4 - Select a table the database.

Splunk databases can contain multiple tables and views. Specify which table or view within the database you want to connect to.

Step 5– Give the connection a name for use in Tableau.

Specify a unique name for the connection. The default name is automatically generated based on the database and table.

A completed Connection dialog box for Splunk is shown below.

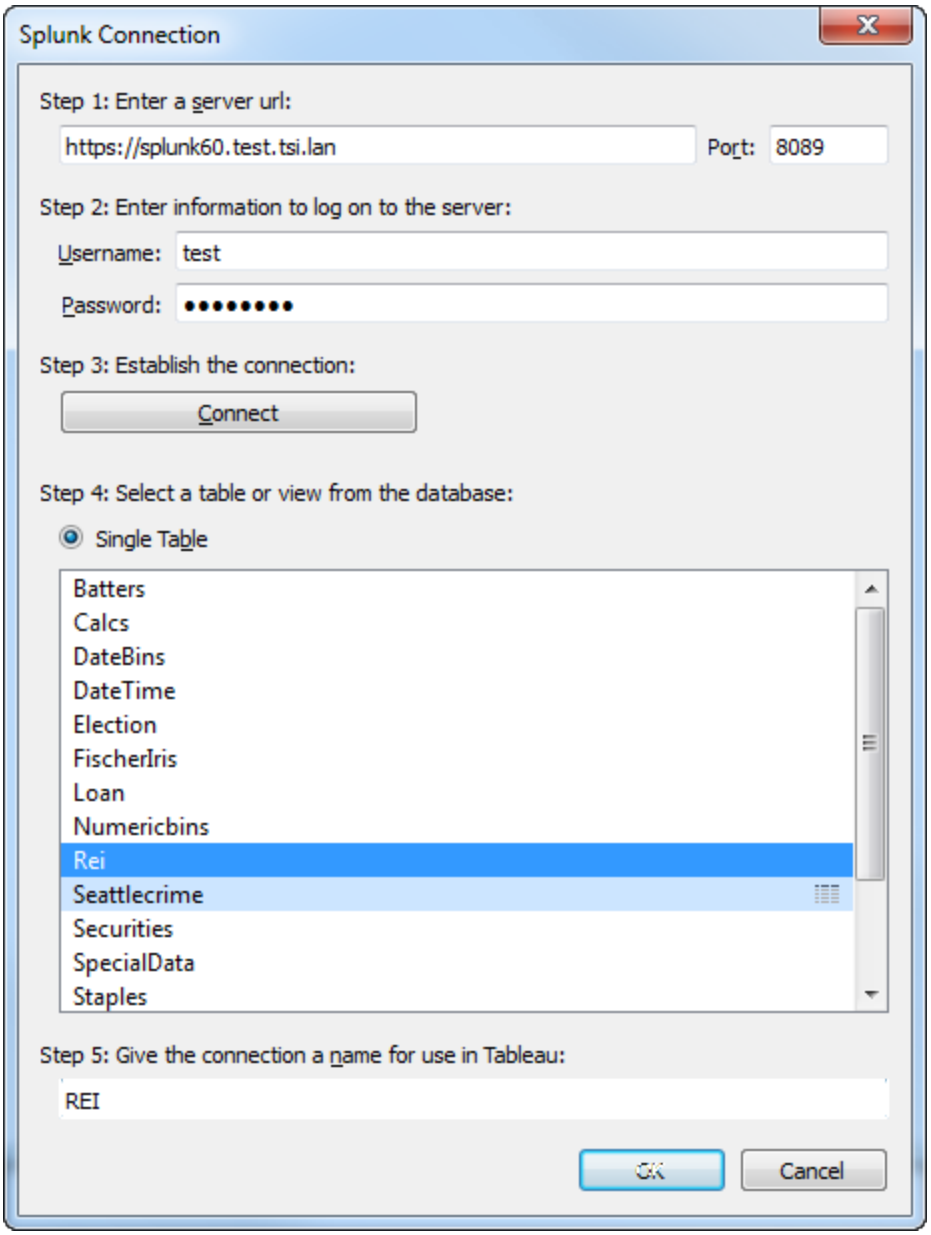

SAP HANA

This topic describes how to connect Tableau to an SAP HANA database.

1. Click **Connect to Data**.
2. On the Connect page, click **SAP HANA**, and then do the following:
    1. Type the name of the server that hosts the database you want to connect to.
    2. Type the user name and password, and then click **Connect**.

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.

3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Schema drop-down list, type the schema name in the text box, click the search icon, and then select the schema.
    3. Type the table name in the Table text box,click the search icon, drag the table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

    > **Note:** If the table you use includes required or optional variables or parameters, the Variables and Input Parameters window opens. Required variables and parameters display their current value or ***Required**. Optional variables and parameters display their current value or are blank. Select a variable or parameter and type or select a value for it. Repeat for all required values and any options values that you want to include, and then click **OK**.

    Optionally, you can make the following configurations before building the view:

    - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
    - **Change data types** – Change the data type of a column by clicking the data type icon.
    - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
    - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
    - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

> **Note** Learn more about connecting to SAP databases in the SAP FAQ knowledge base article.

An example of an SAP HANA data source is shown below.

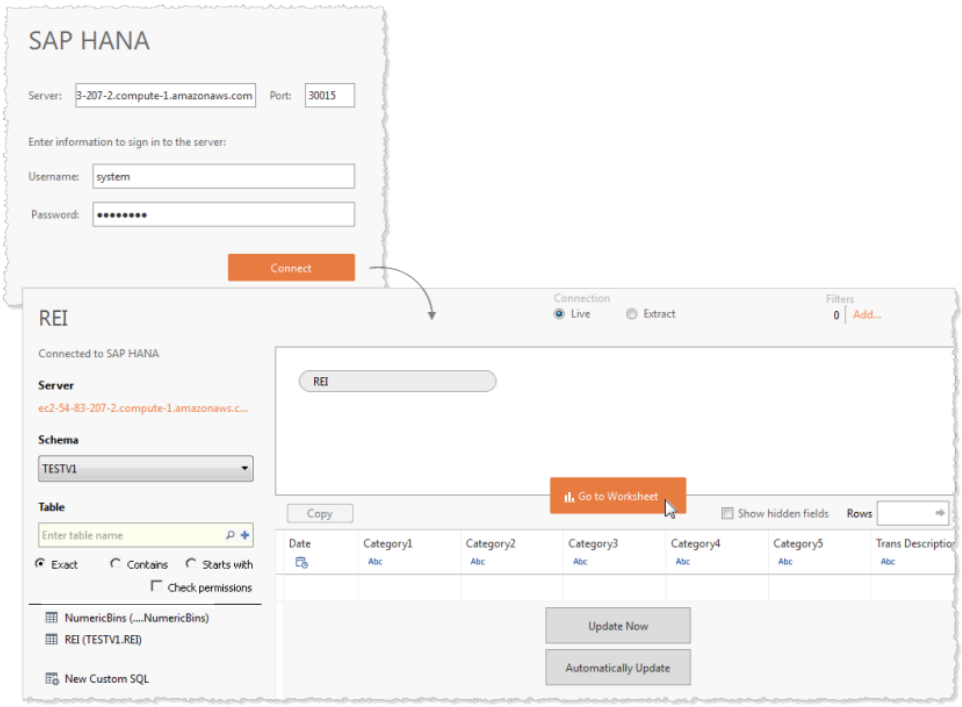

SAP NetWeaver Business Warehouse

This topic describes how to connect Tableau to a SAP NetWeaver Business Warehouse data source.

1. Click **Connect to Data**.
2. Click **SAP NetWeaver Business Warehouse**.
3. Follow the steps in the SAP NetWeaver Business Warehouse Connection dialog box to complete the connection.

   Step 1 – Select an SAP Logon connection.

   Step 2 – Enter information to sign in to the server. Type the Client ID of the BW system, username, password, and optional language setting.

   Step 3 – Establish the connection.

   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 4 – Select an InfoProvider.

   Step 5– Select a Query or InfoCube.

   Step 6– Give the connection a name for use in Tableau.

Specify a unique name for the connection. The default name is automatically generated based on the database and table.

Note: Learn more about connecting to SAP NetWeaver Business Warehouse databases in the Connecting to SAP BW Knowledge Base article.

When you connect to SAP NetWeaver Business Warehouse (SAP BW) databases, you have the option to select which alias to use for dimensions. SAP BW offers a number of different options including long name, short name, medium name, and key. Sometimes the alias used by default is not very useful. For example, in the view below, Key is being shown for the Customer dimension. You can right-click the dimension and change it to another option such as Short Name, which changes the labels to something more meaningful.

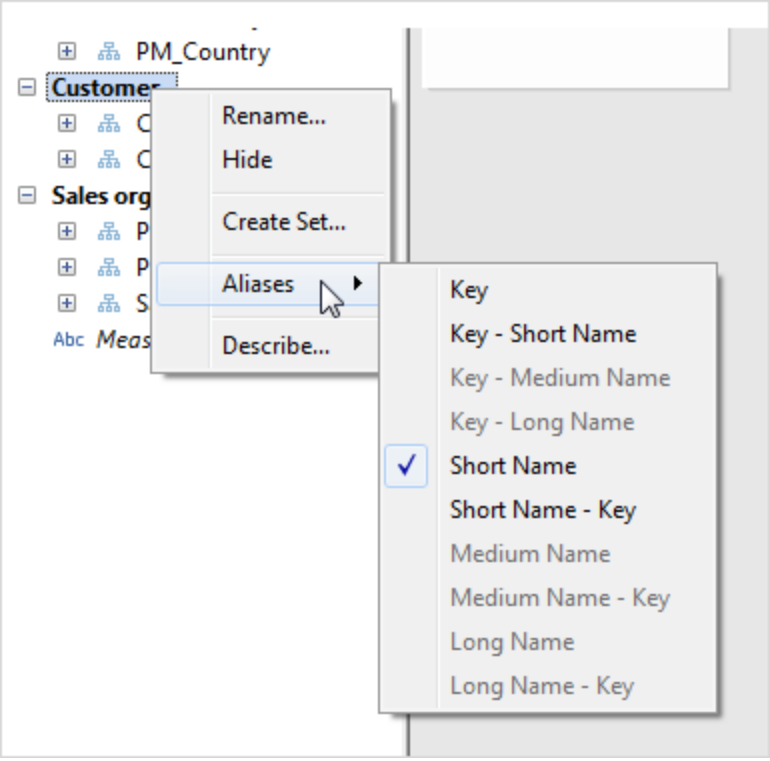

A completed Connection dialog box for SAP BW is shown below.

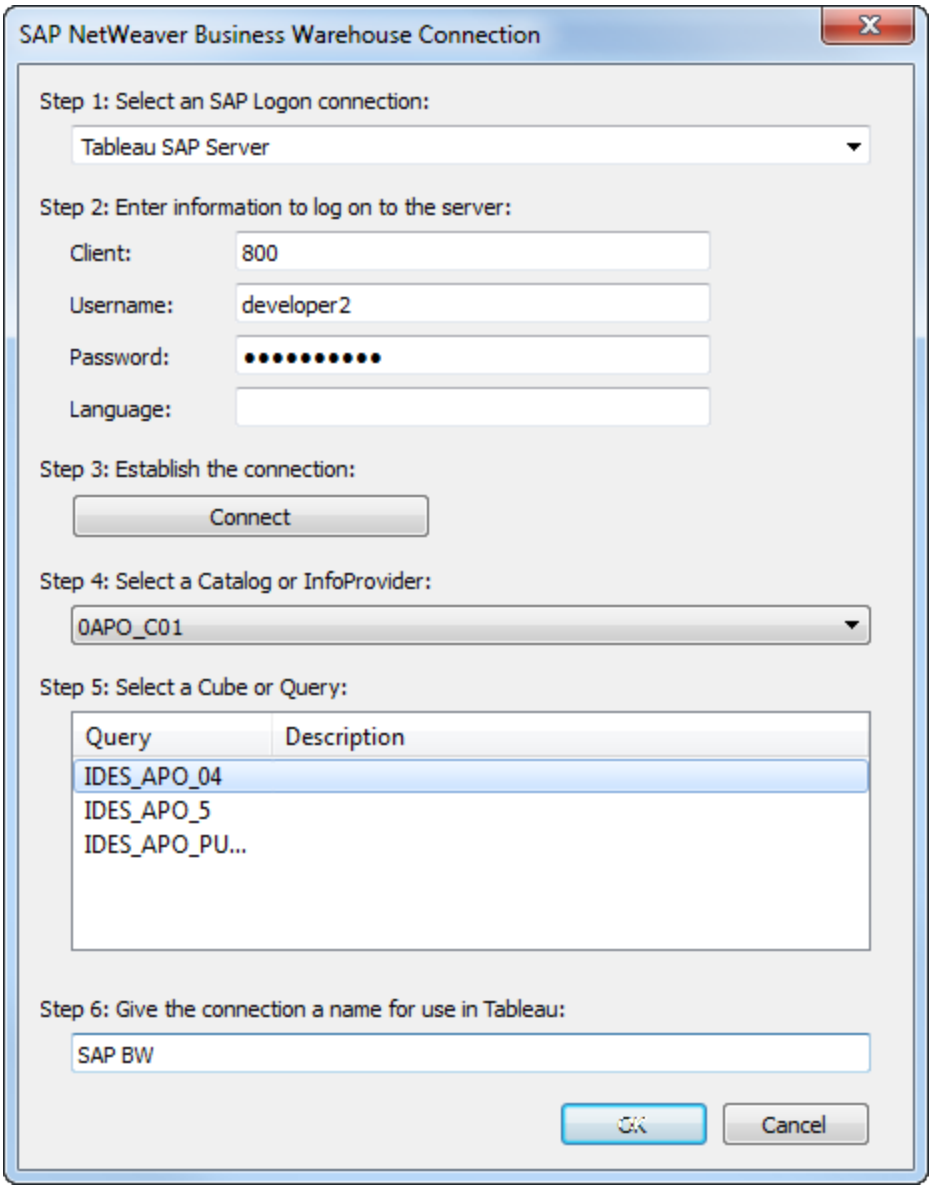

SAP Sybase ASE

This topic describes how to connect Tableau to an SAP Sybase ASE database and set up the data source. Tableau connects to Sybase ASE version 15.2 and later. It is recommended that you connect using the Sybase 15.5 driver at the ESD10 maintenance level.

1. Click **Connect to Data**.
2. On the Connect page, click **SAP Sybase ASE**, and then do the following:
    1. Type the name of the server you want to connect to, and specify the port to use.
    2. Type the user name and password, and then click **Connect**.

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.

3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Database drop-down list, select a database.
    3. From the Schema drop-down list, select a schema .
    4. Drag a table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

       Optionally, you can make the following configurations before building the view:

       - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
       - **Change data types** – Change the data type of a column by clicking the data type icon.
       - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
       - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
       - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

Note: When connecting to SAP Sybase ASE databases, you can optionally specify a SQL command that will be run once upon connection. See **Initial SQL** on page 142 to learn more about adding these commands to your connection. To access stored procedures, the database must be appropriately configured with a linked server, which could be a remote server or a pointer to the database itself (known as a *loopback*).

An example of an SAP Sybase ASE data source is shown below.

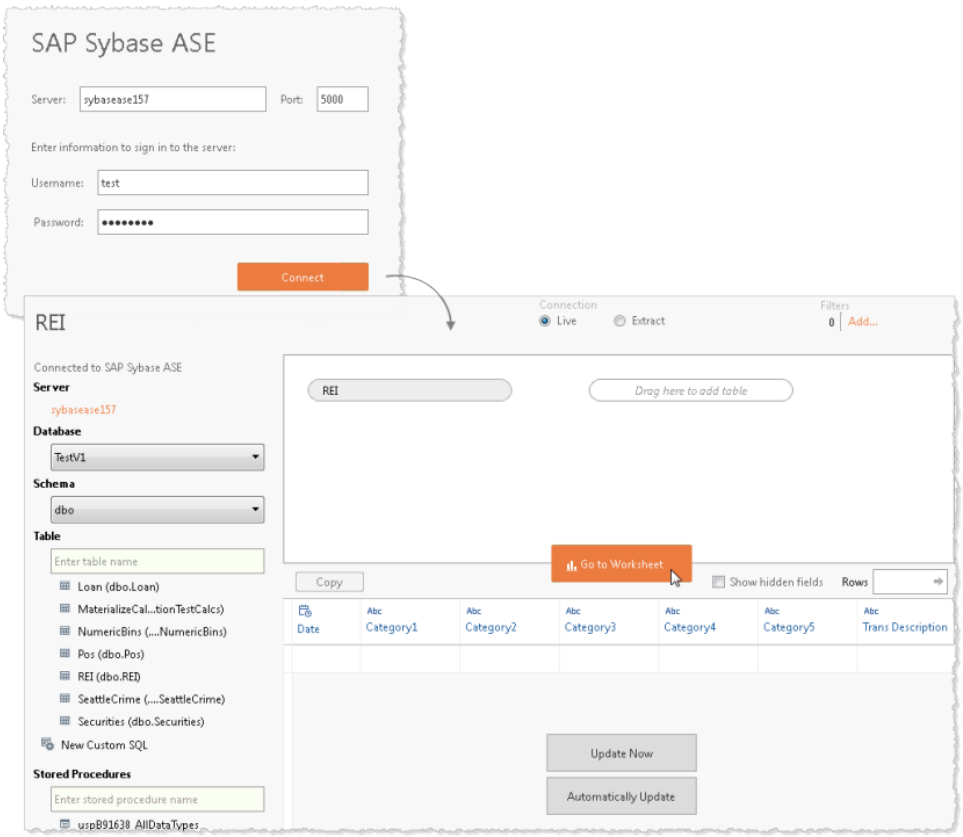

SAP Sybase IQ

This topic describes how to connect Tableau to a SAP Sybase IQ database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **SAP Sybase IQ**, and then do the following:
    1. Type the server name that hosts the database and the name of the database that you want to connect to.
    2. Select how you want to sign in to the server. Specify whether to use Windows Authentication or a specific user name and password. If the server is password protected, you must type the user name and password.
    3. Click **Connect**.

    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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.

2. From the Schema drop-down list, select a schema.
3. Drag a table to the join area, and then click **Go to Worksheet** to begin building your view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

   Optionally, you can make the following configurations before building the view:

   - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
   - **Change data types** – Change the data type of a column by clicking the data type icon.
   - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
   - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
   - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

An example of an SAP Sybase IQ data source is shown below.

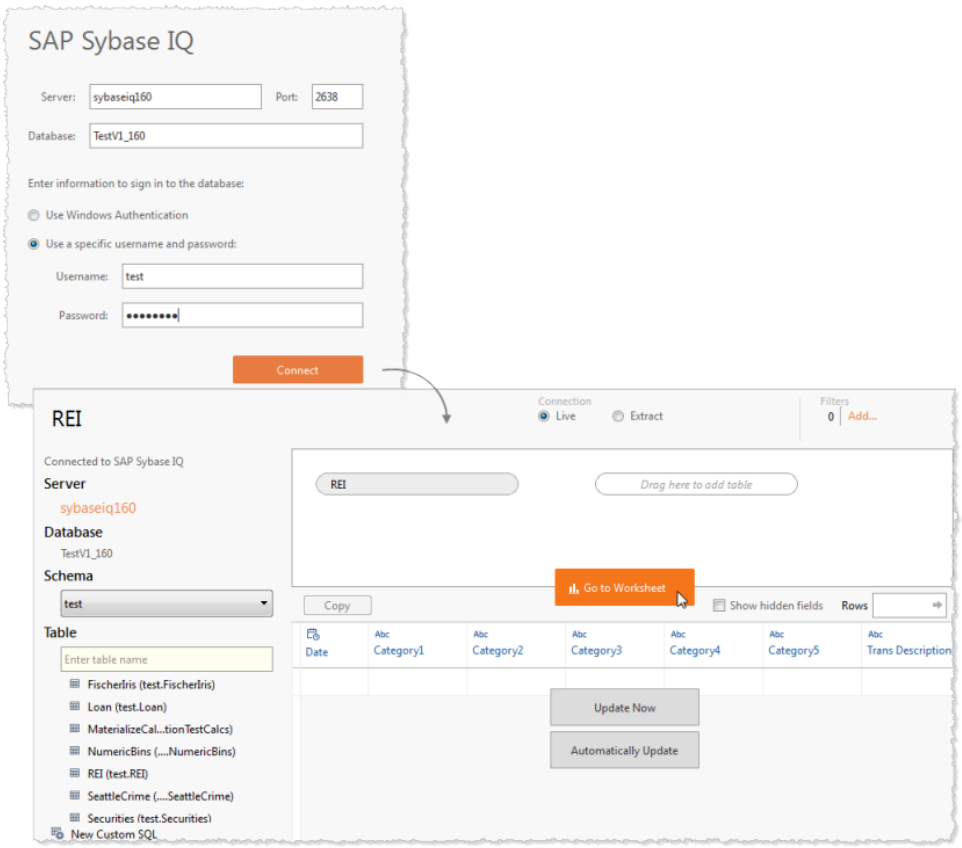

Teradata

This topic describes how to connect Tableau to a Teradata database and set up the data source.

1. Click **Connect to Data**.
2. On the Connect page, click **Teradata**, and then do the following:
    1. Type the name of the server that hosts the database you want to connect to.
    2. Select the authentication method. Tableau supports connecting with the built-in Teradata authentication or LDAP.
    3. Type the user name and password, and then click **Connect**.

       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.
3. On the data source page, do the following:
    1. (Optional) Click the default data source name at the top of the page, and then type a unique data source name for use in Tableau.
    2. From the Database drop-down list, type the database name in the text box, click

the search icon, and then select the database.

3. Type the table name in the Table text box, click the search icon, drag the table to the join area, and then click **Go to Worksheet** to begin building the view. Alternatively, drag additional tables to the join area and use conditions to join them. You can choose to combine tables or use Custom SQL when you are working with multiple tables. You can also add joins later.

   You can also specify a stored procedure in the database. See **Stored Procedures** on page 143 for details and for a list of constraints specific to Teradata databases.

   Optionally, you can make the following configurations before building the view:

   - **Preview the data** – In the preview area, select **Update Now** to preview the first 10,000 rows of the data source. If you add tables, remove tables, or make changes to the join conditions, click **Update Now**. If you want changes to automatically reflect in the preview area, click **Automatically Update**.
   - **Change data types** – Change the data type of a column by clicking the data type icon.
   - **Rename or hide columns** – Rename or hide a column by clicking the column header drop-down menu and selecting the respective option.
   - **Connect live or use an extract** – At the top of the data source page, select a live or extract connection to the data source. If you choose to take an extract, an Edit link displays allowing you to set up filters that define a subset of the data to include in the extract.
   - **Add data source filters** – Add data source filters to restrict the visibility and use of the fields in the data source.

Note: When connecting to Teradata databases, you can optionally set up Query Banding. See **Advanced Teradata Options** on the next page. Additionally, you can specify a SQL command that will be run once upon connection. See **Initial SQL** on page 142 to learn more about adding these commands to your connection.

An example of a Teradata data source is shown below.

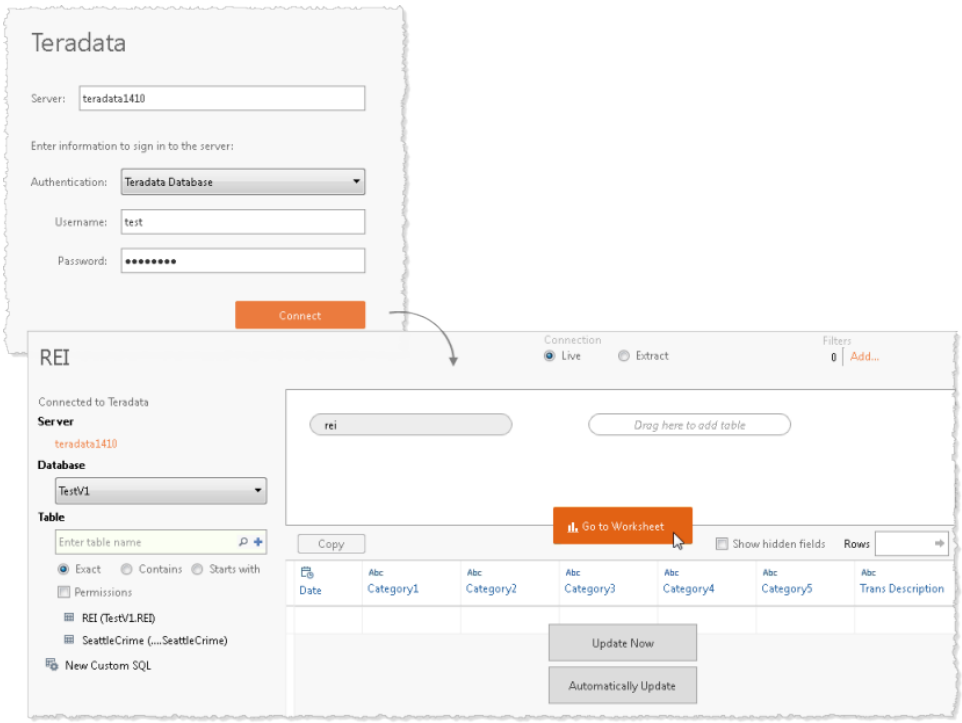

Advanced Teradata Options

When connecting to Teradata databases you can optionally set up query bands and initial sql. These advanced options are used to increase performance and take advantage of the built in security rules of the database.

Query Bands

When connecting to a Teradata database, you can optionally define query band statements that run during connection. Query banding allows you to pass parameters into the Teradata environment. Use these to set up a workbook to filter the data based on security rules that exist in the database. For example, you can pass in the Tableau Server user name for the current user so when the view is loaded it only shows the data specific to that user. Query bands can also be used to improve performance. When connecting to Teradata, you can define a map between the name of the attributes passed into the query band and the corresponding values from Tableau.

To set up query banding:

1. On the data source page, select **Data** > **Query Banding and Initial SQL**.
2. In the subsequent dialog box, specify name/value pairs in the top text box labeled Query Banding. You can use the **Insert** drop-down menu to add Tableau values. The Tableau values are described in the table below.

Value	Description	Example
`<TableauMode>`	The mode Tableau is operating in when generating queries. This value will either be “Connect” when	Connect or Analytical

Value	Description	Example
	retrieving metadata or “Analytical” when retrieving actual data.	
`<LoginUser>`	The username of the person signed in to the database.	jsmith
`<ServerUser>`	The signed-in Server user. Does not include domain name.	jsmith
`<ServerUserFull>`	The signed-in server user including the domain name (if the server is using Active Directory).	domain.lan\jsmith
`<ProxyUser>`	Used when setting up impersonation on the server. Provides the username of the current server user.	jsmith
`<ProxyUserFull>`	Used when setting up impersonation on the server. Provides the username and domain name of the current server user.	domain.lan\jsmith
`<TableauApp>`	The name of the Tableau application.	Tableau Desktop Professional or Tableau Server
`<TableauVersion>`	The version of the Tableau application	6100.11.0428.0300
`<WorkbookName>`	The name of the workbook.	Financial-Analysis

An example query band statement is shown below. This example passes the user name for the current server user.

```
Applic-
ationName-
=<TableauApp>Version=<TableauVersion>ProxyUser=<ProxyUser>TableauMode=<Tableau
```

Tableau checks the statement for errors as you type. When it is valid a green check mark displays at the bottom of the text box.

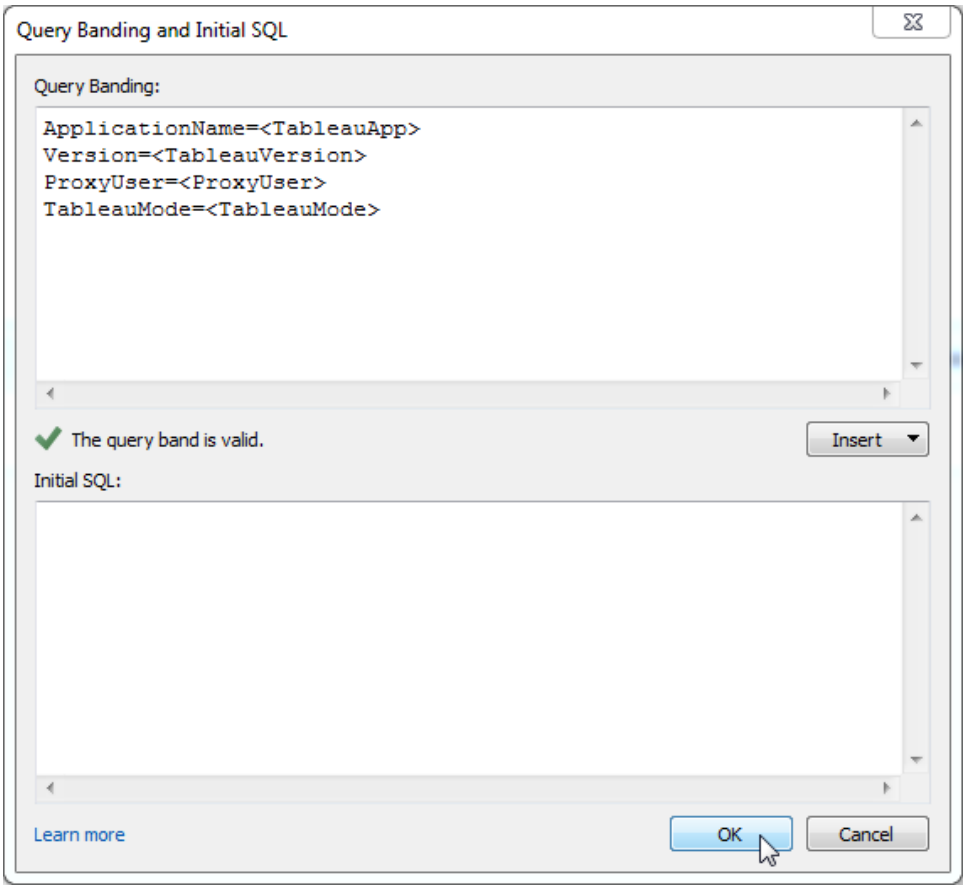

Initial SQL

When connecting to Teradata databases, you can optionally specify a SQL command that will be run once upon connection. See **Initial SQL** on page 142 to learn more about adding these commands to your connection.

> If Tableau detects that the Initial SQL payload cannot be delivered in one query, it attempts to split the query into separate statements, where each statement must end with a semicolon followed by a newline character, with no additional characters (such as an extra space) in between. This requirement is important because if a semicolon exists within the interior of a statement, for example within a string literal, it may not be safe to split the statement at that point.If the Initial SQL code that you specify in Tableau returns an error, and the same Initial SQL is deemed valid by the Teradata databases, it may be that Tableau has split the query incorrectly. If this happens, try reformatting the SQL code in Tableau manually.

Teradata OLAP Connector

This topic describes how to establish a Teradata OLAP connection.

1. Click **Connect to Data**.
2. Click **Teradata OLAP Connector**.
3. Follow the steps in the Teradata OLAP Connection dialog box to complete the connection.

   Step 1 – Type the name of the server that hosts the database.

   Step 2 – Enter your username and password for the server.

   Step 3 – Establish the connection.

   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 4 – Select a catalog.

   Step 5 – Select a cube from the database.

   Step 6 – Give the connection a name for use in Tableau.

   Specify a unique name for the connection.

Named sets from a Teradata OLAP data source are displayed in the Sets area of the Data window in Tableau. You can interact with these named sets in the same way you interact with other custom sets in Tableau. See **Understanding the Data Window** on page 193. You can view underlying data for Teradata OLAP data sources, provided the database administrator has enabled this functionality. See **View Data** on page 479.

A completed Connection dialog box for Teradata OLAP Connector is shown below.

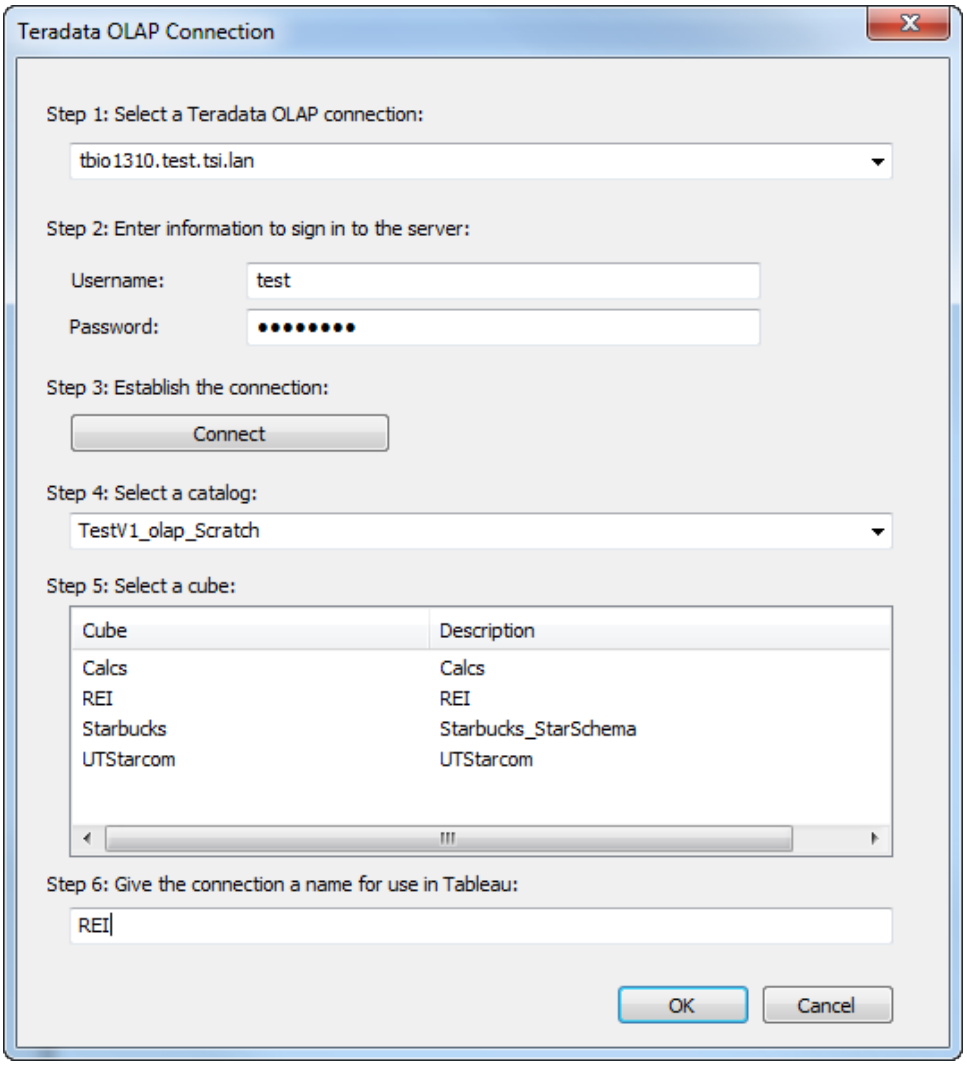

Window Azure Marketplace DataMarket

This topic describes how to connect Tableau to an Window Azure Marketplace DataMarket data source.

1. Click **Connect to Data**.
2. Click **Window Azure Marketplace DataMarket**.
3. Follow the steps in the Window Azure Marketplace DataMarket Connection dialog box to complete the connection.

   Step 1 – Select or enter a URL to the data you want to connect to. Click the link to find an Window Azure Marketplace DataMarket data source on Windows Window

Azure Marketplace DataMarket DataMarket.

Step 2 – Enter authentication information.

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

Step 3 – Establish the connection.

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. The default name is automatically generated.

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

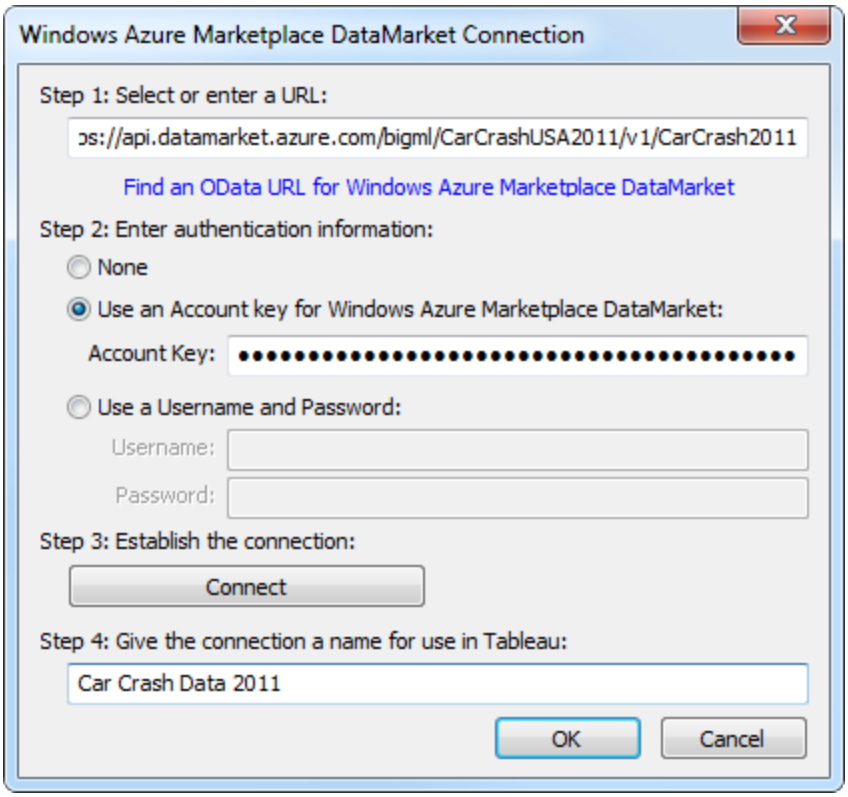

Other Databases (ODBC)

Tableau Desktop allows you to connect to several different files and databases. You can see the file and database types that are supported by Tableau Desktop listed in the left-pane of the Connect page. For the listed files and database types, Tableau has implemented customized techniques and capabilities so that you can get the most from your data when connected to Tableau Desktop.

If your database is not listed, you have the option to use the general ODBC standard for connecting Tableau Desktop to your data. As an industry-defined standard, many database vendors make ODBC drivers available for connecting to their database. To connect to your data using an ODBC driver, click **Other Databases (ODBC)** on the Connect page. For more information, see the Tableau and ODBC knowledge base article.

> **Important:** When you use an ODBC driver to connect to an unsupported database, the outcome may vary and compatibility with Tableau Desktop features is not guaranteed.

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.

> **Note:** For Microsoft Excel and text file data sources, this option is available only when using the legacy connection or in workbooks that were created before Tableau Desktop 8.2.

To connect to a custom SQL query

1. After connecting to data, double-click the **New Custom SQL** option in the data source page.

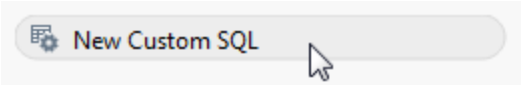

2. Type or paste the query into the text box.

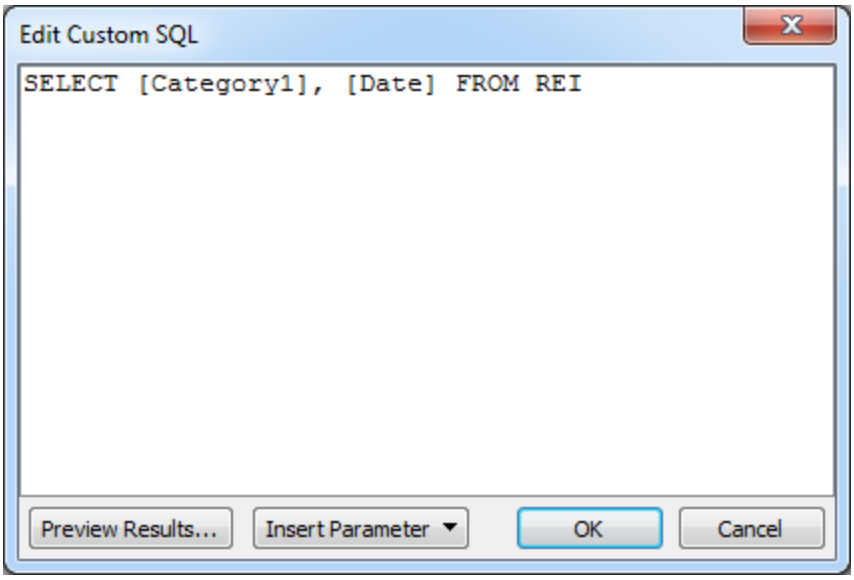

3. When finished, click **OK**.

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 = titleau-
thor.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.

Using Parameters in a Custom SQL Query

You can use parameters in a custom SQL query statement to replace a constant value with a dynamic value. You can then update the parameter in the workbook to modify the connection. For example, you may connect to a custom SQL query that provides web traffic data for a particular page that is specified by a pageID. Instead of using a constant value for the pageID value in the SQL query, you can insert a parameter. Then after finishing the connection, you can show a parameter control in the workbook. Use the parameter control to switch out the pageID and pull in data for each page of interest without having to edit or duplicate the connection.

You can create a parameter directly from the Custom SQL dialog box or use any parameters that are part of the workbook. If you create a new parameter, it becomes available for use in the workbook just like any other parameter. See **Parameters** on page 663 to learn more.

To add a parameter to a custom SQL query

1. In the join area of the data source page, hover over the table until the edit icon displays, and then click the edit icon.
2. At the bottom of the dialog box, click **Insert Parameter**.
3. Select a constant value in the SQL statement and then, from the **Insert Parameter** drop-down menu select the parameter you want to use instead. If you have not created a parameter yet, select **Create a new parameter**. Follow the instructions in **Creating Parameters** on page 663 to create a parameter.

> **Note:** Parameters can only replace literal values. They cannot replace expressions or identifiers such as table names.

The workbook in the example below connects to a Custom SQL query that returns all orders with that are marked as Urgent priority. In the SQL statement, the order priority is the constant value. If you want to change the connection to see the High priority orders, you would have to edit the data source.

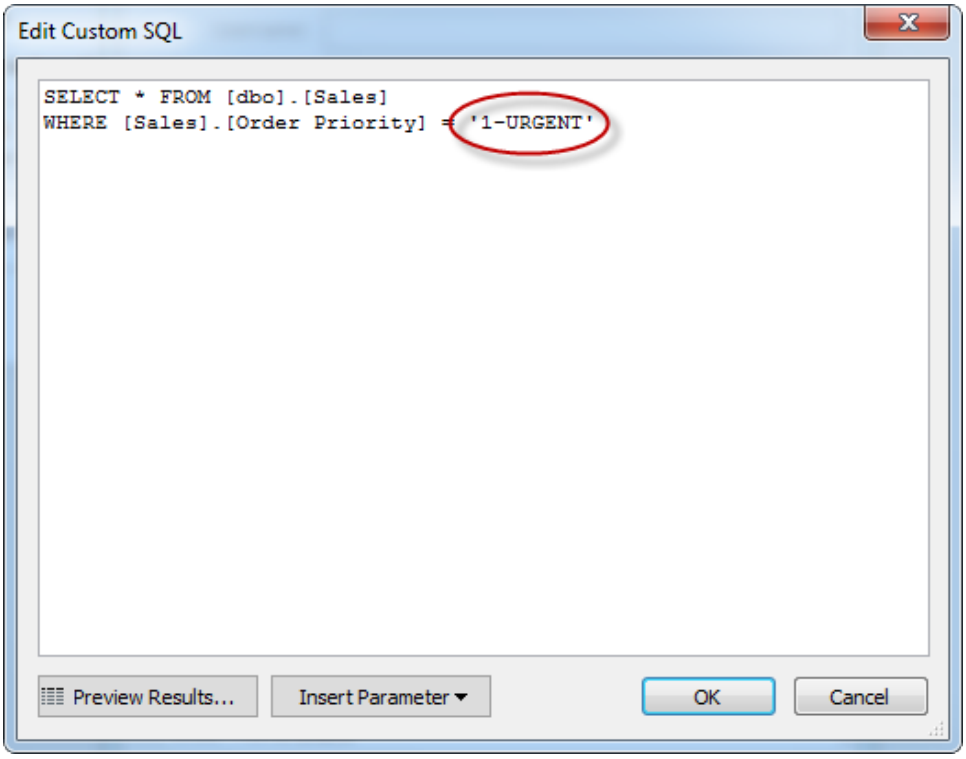

Instead of creating and maintaining many variations of the same query, you can replace the constant order priority value with a parameter. The parameter should contain all of the possible values for Order Priority.

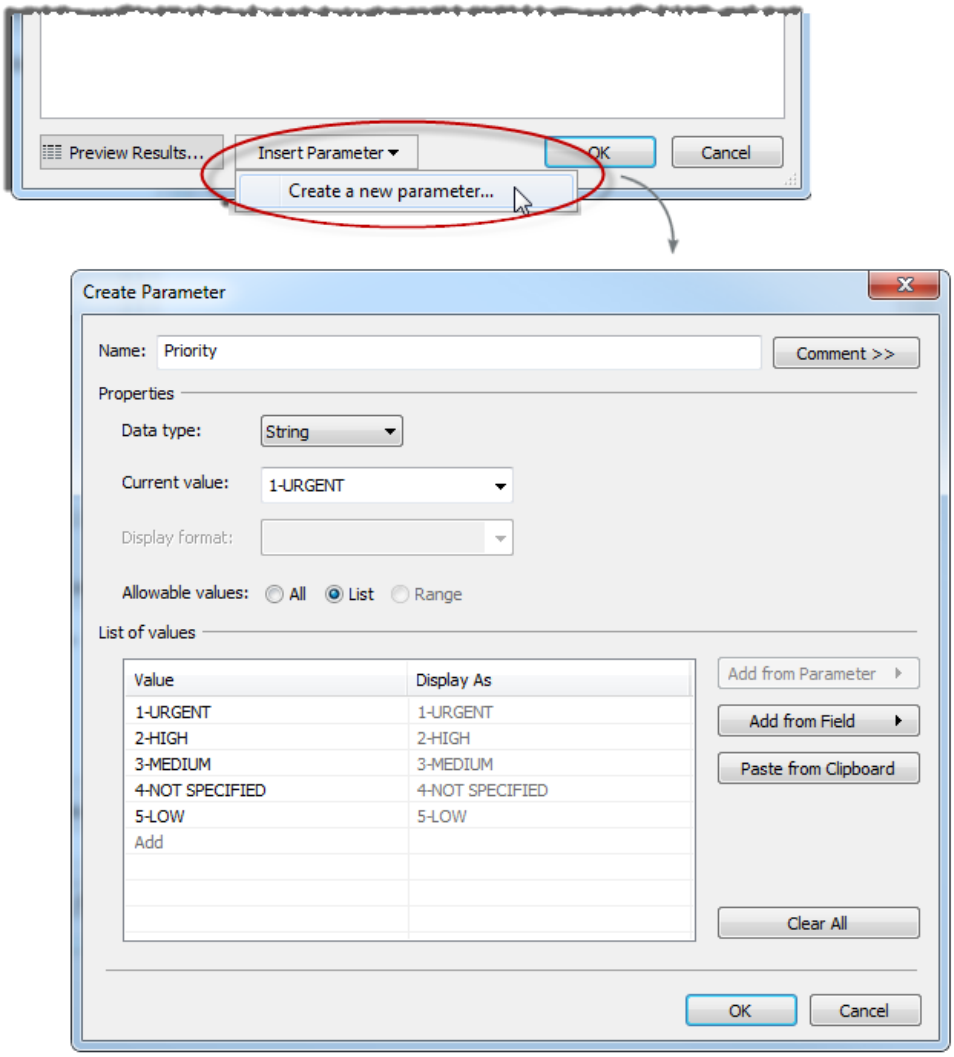

After you create a parameter, you can insert it into the SQL statement to replace the constant value.

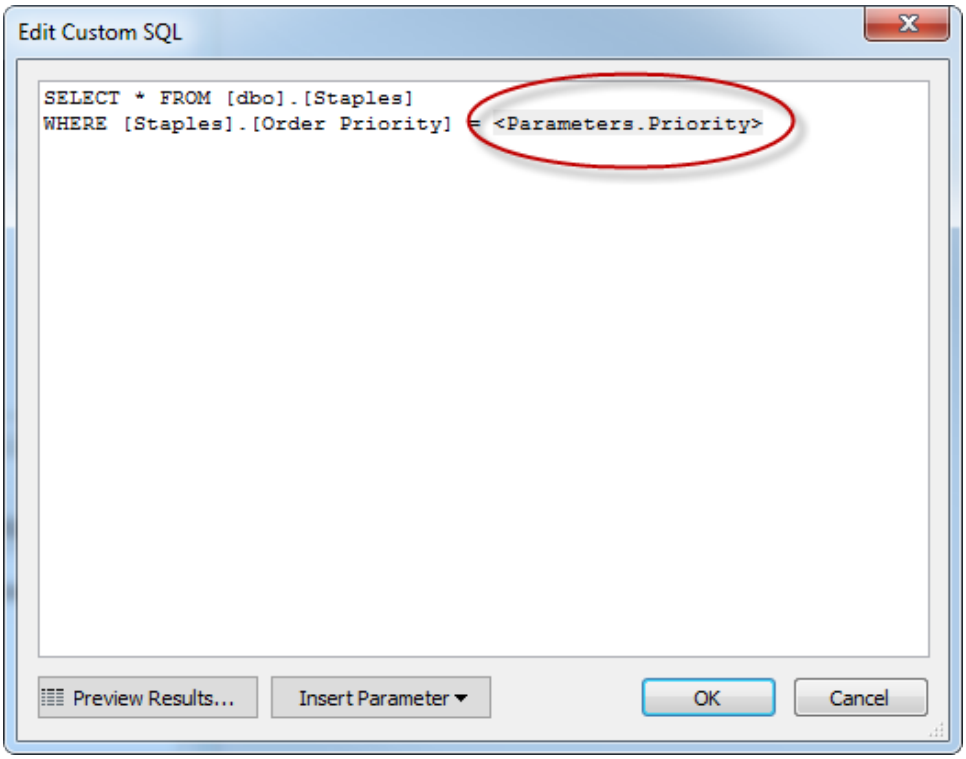

After you finish editing the connection, the new parameter is listed in the Parameters area at the bottom of the Data window and the parameter control displays on the right side of the view. As you select different values, the connection updates.

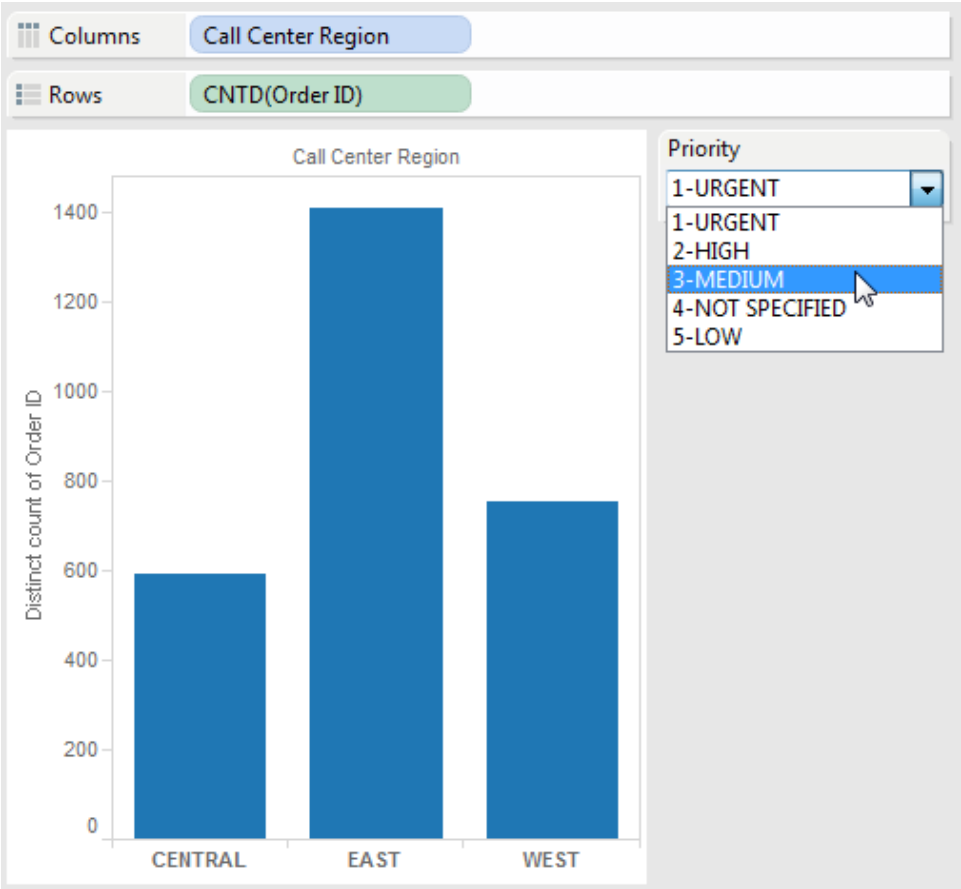

Note: If you are using an extract, you must refresh the extract in order to reflect changes to the parameter. Publishing a data source that uses Custom SQL parameters includes the parameters. The parameters are transferred to any workbooks that connect to the data source.

Initial SQL

When connecting to some databases, you can optionally specify a SQL command that will be run once, when you open the workbook. This initial SQL is different than a custom SQL connection, which affects every query to retrieve rows from the database.

You can use this command to:

- Set up temporary tables that will be used during the session.
- Set up a custom data environment.

After the initial connection, on the data source page, you have the option to use initial SQL.

To use initial SQL

1. On the Data Source page, select **Data** > **Initial SQL** or **Data** > **Query Banding and Initial SQL** depending on the database you are connecting to.
2. Type the SQL command into the **Initial SQL** text box.

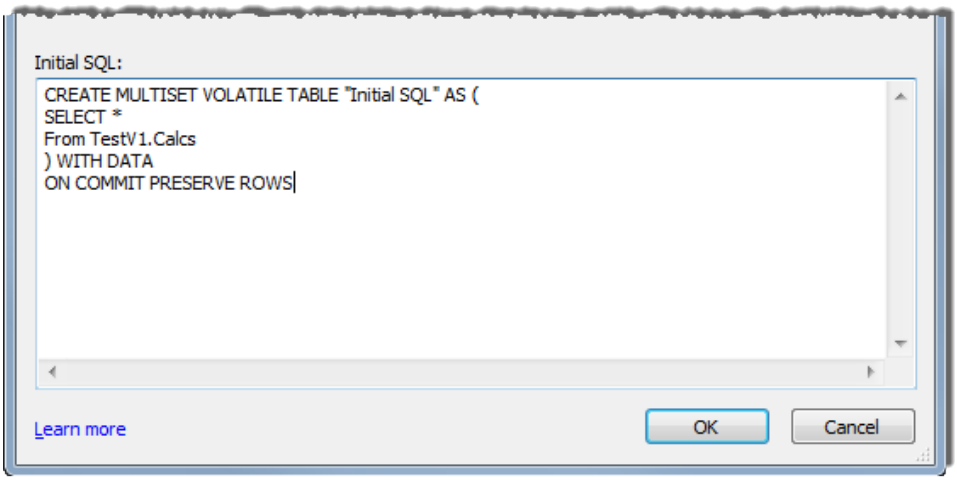

Note: Tableau does not examine the statement for errors. This SQL statement is simply sent to the database upon connection.

Your software license may restrict you from using initial SQL with your connection. If you are publishing to Tableau Server, the server must be configured to allow Initial SQL statements. By default, the server software is configured to allow these statements to run when the workbook is loaded in a web browser. Administrators can disable the functionality on the **Data Connections** tab of the Tableau Server Configuration utility. If the server does not allow Initial SQL statements, the workbook opens, but the Initial SQL commands are not sent.

Stored Procedures

A stored procedure is a subroutine available to applications that access a relational database system. When you connect to a SAP Sybase ASE, Microsoft SQL Server, or Teradata database with Tableau, you can use a stored procedure to define the connection.

When you are creating a data source using one of these data types, available procedures are listed under **Stored Procedures** .

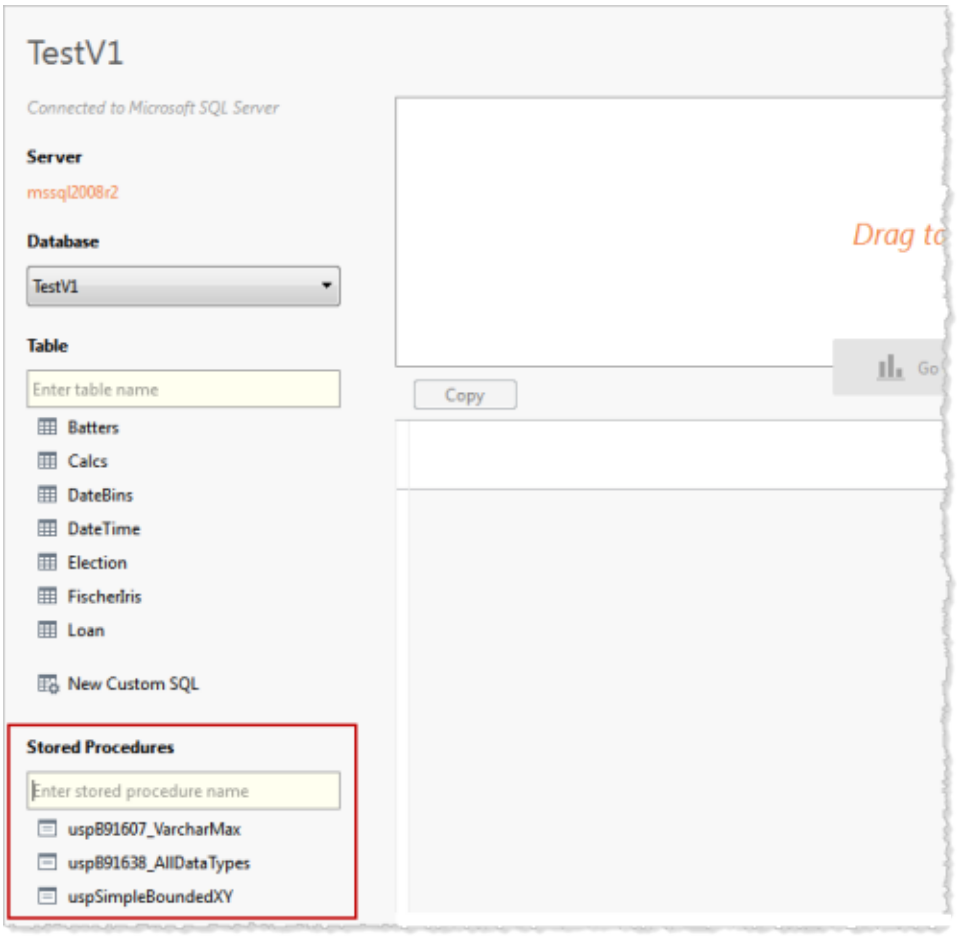

From the left pane, drag the procedure to the join area or double-click one of the listed procedures. If parameters are available in the procedure, the Parameters dialog box automatically displays:

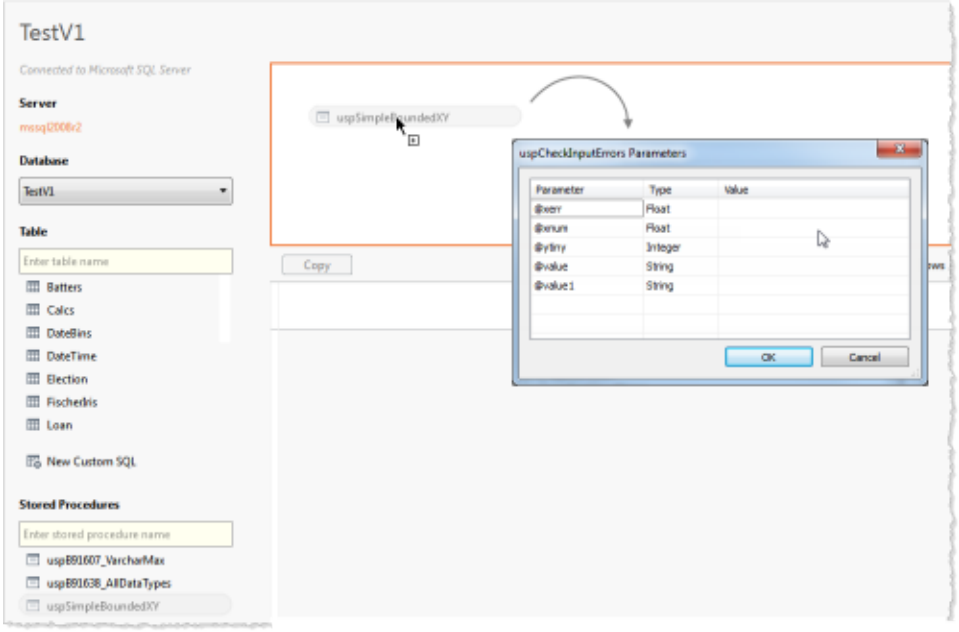

Instead of entering a value, you can use an existing Tableau parameter, or create a new Tableau parameter for the value:

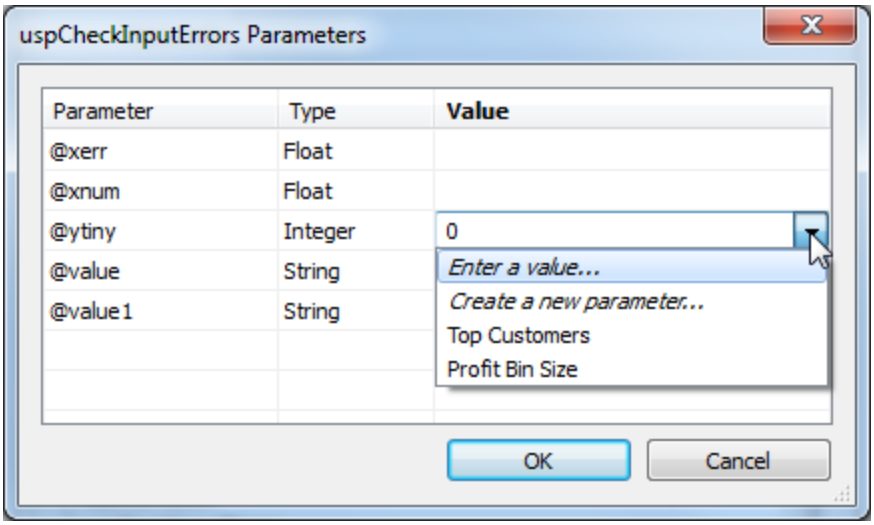

If you then expose the Tableau parameter in the view, users are able to change the value of the parameter in the procedure interactively.

Notes on Stored Procedures

If you use stored procedures to define a data source for Tableau, keep the following in mind:

- If a stored procedure returns more than one result set, Tableau reads the first one and ignores the rest.
- If a stored procedure has output parameters, Tableau filters out the stored procedure.

- Stored procedures that have parameters of a non-scalar type are excluded.
- Result set columns that don't have matching types in Tableau (such as varbinary, geometry, and hierarchyid) are logged. If all result set columns map to unknown data types, Tableau displays a message:

  "The result set... has no usable columns."
- Stored procedures that return no result sets are listed on the data source page but fail if selected.
- If no value is provided for a parameter that the stored procedure requires, an error occurs. Tableau cannot determine in advance whether parameters are required.
- Tableau does not perform any transaction management for stored procedures. That is, stored procedure writers must not depend on Tableau to start transactions before invoking stored procedures, or to commit them afterward.
- Column names must be unique for stored procedures to work. If two columns have the same name, or if no name is provided, the procedure can result in an error message:

  "InsertData: unbound column error"
- If there are multiple queries in a stored procedure (for example, to read values from another table or to hold temporary combinations) each of the queries must return the same sets of columns in the same order (same names and data types). To ensure that column order and names match in the query results, you may need to explicitly `CAST` to ensure the data type is correct, for example `CAST(Username as VARCHAR(20))`, and explicitly name the columns. If a stored procedure does not follow these guidelines, an error message can result:

  "InsertData: unbound column error"
- If there are multiple queries in a stored procedure (for example, to read values from another table or to hold temporary combinations) and the procedure is generating an error, try adding `SET NOCOUNT ON` to the top of the procedure.This prevents the message which shows the count of number of rows affected by a Transact-SQL statement from being returned as part of the result set for a query.

In addition, the following constraints apply for specific databases.

Stored Procedure Constraints for Teradata Databases

The following constraints apply for stored procedures on Teradata databases.

- Values must be provided for every parameter. If the user does not provide a value for one or more parameters, Tableau displays a Teradata database error stating there are too few values provided for the stored procedure.

Stored Procedure Constraints for SQL Server Databases

The following constraints apply for stored procedures on SQL Server databases.

- If the result set for a stored procedure contains columns of type IMAGE or TEXT, the stored procedure will fail with an "Incorrect syntax" error message.

- If the total width of the result set (number of bytes in each row) exceeds 8060, the stored procedure fails. This can occur with very wide tables (hundreds of columns) or with tables having large text columns, intended to hold thousands of characters of text.
- Tableau does not display stored procedures from schema "sys".
- If the user does not provide a value for one or more parameters that the procedure requires, Tableau displays a SQL Server database error in the form "The procedure requires a value for parameter @x but one was not provided."
- Stored procedures that contain multiple queries should follow the guidelines listed in Notes on Stored Procedures (above).
- Tableau Desktop does not support the Microsoft SQL Server TIME data type. When fields of this type are included in a stored procedure on a Microsoft SQL Server database Tableau Desktop will not import them.

Stored Procedure Constraints for SAP Sybase ASE Databases

The following constraints apply for stored procedures on SAP Sybase ASE databases.

- The database must have a properly configured remote server.
- If the user does not provide a value for one or more parameters that the procedure requires, Tableau displays a Sybase ASE database error in the form "The procedure requires a value for parameter @x but one was not provided."

Replacing Data Sources

There are times when you may want to update a workbook or sheet to use a different data source. For example, you may have a workbook that was built against a local text file. However, the data source has since been published to Tableau Server and you want to connect to the published version. Rather than rebuild your workbook using the new connection, you can replace the data source.

Replacing the data source is available between:

- Two relational data sources (on a server, local files, or extracts).
- Two multidimensional data sources of the same type. In Tableau, multidimensional data sources are supported only in Windows.

The two data sources do not have to be identical, however, any differences will affect the sheets in the workbook. Any fields, groups, sets, and calculated fields that don’t exist in the new data source (or have a different name) are removed from the Data window. If these fields were used by any sheets in the workbook, they will remain in the Data window but will be marked invalid. Additionally, you may see changes to custom sets, groups, and calculated fields that

depend on the missing fields. See **Connecting to Data and Editing Data Sources** on page 52 to learn how to replace field references to correct invalid fields.

> **Note:** Replacing a data source does not merge or edit the data sources. Rather, replacing a data source simply redirects fields in the worksheet to map to corresponding fields in the new data source. To successfully replace a data source, any fields in the original data source (e.g., calculated, groups, sets parameters, and so on), should also exist in the new data source. If not, then you may want to manually copy and paste these fields to the new data source before replacing the old one.

Follow the steps below to replace a data source:

1. Open a workbook that connects to the *old* data source.
2. Select **Data** > **Connect to Data** and then connect to the *new* data source.
3. Select **Data** > **Replace Data Source**.

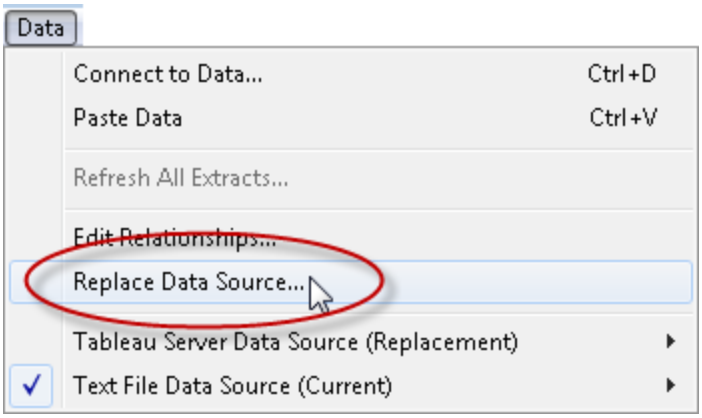

4. In the Replace Data Source dialog box, select the **Current** data source and the **Replacement** data source.

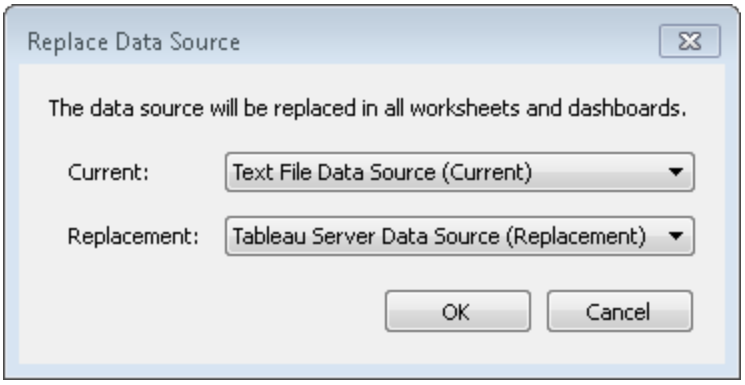

5. When finished, click **OK**.

All worksheets, dashboards, and stories that used the original connection are updated to use the new data source. Click **Undo** on the toolbar to revert the change and return to your original connection.

Exporting Data Sources

At any time while connected to a data source, you can export data source information as a shortcut that allows you to quickly connect to the data source in the future. You might want to do this if you often connect to the same data source multiple times or if you've added joined tables, default properties, or custom fields, such as groups, sets, calculated fields, and binned fields, to the Data window.

> **Note:** You can also save custom fields by saving the workbook or by creating a bookmark file.

You can export the data source in one of the following two formats:

Data Source (.tds) - contains just the information you need to connect to the data sources such as data source type, location, and custom fields. If you connect to local file data sources (Excel, Access, text, extracts), the file path is stored in the data source file.

Data source files contain the following types of information:

- data source type
- data source connection information specified in the data source page (e.g., server, port, location of local files, tables, and so on)
- groups
- sets
- calculated fields
- bins
- default field properties (e.g., number formats, aggregation, sort order, etc.)

Packaged Data Source (.tdsx) - contains all the information in the Data Source (.tds) file as well as any local file data sources (Excel, Access, text, and extracts). This file type is a single zipped file and is good for sharing a data source with people who may not have access to the original data that is stored locally on your computer.

After you export, the data sources are available on the Data page.

By default, the data source files are stored in the Datasources folder of the Tableau Repository. Data source files stored in another location do not display on the Data page. You can connect to data source files by selecting **File** > **Open** and navigating to the file. You can also connect by dragging the data source file onto Tableau Desktop icon or onto the running application.

Note: If you move a local file data source that is referenced by a .tds file, you will be prompted to locate or replace the original data source when you try to open the data source file. To avoid saving a specific file path, save the data source as a TDSX file, which packages a copy of the original local file data source with the .tds file. If you choose to replace the original data source, the replacement data source must be of the same type (Excel, MySQL, and so on) as the original.

To export a data source

1. On the **Data** menu, select a data source, and then select **Add to Saved Data Sources**.
2. Complete the **Add to Saved Data Sources** dialog box by specifying a file name and selecting the type of data source file.

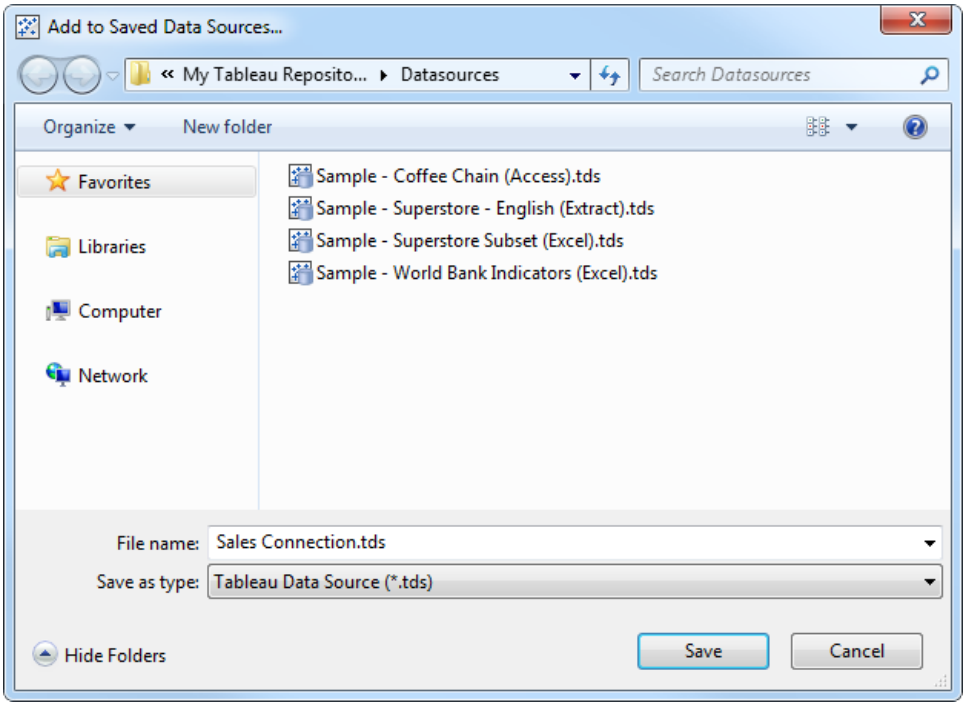

The new .tds or .tdsx file is listed in the Saved data sources section of the Data page.

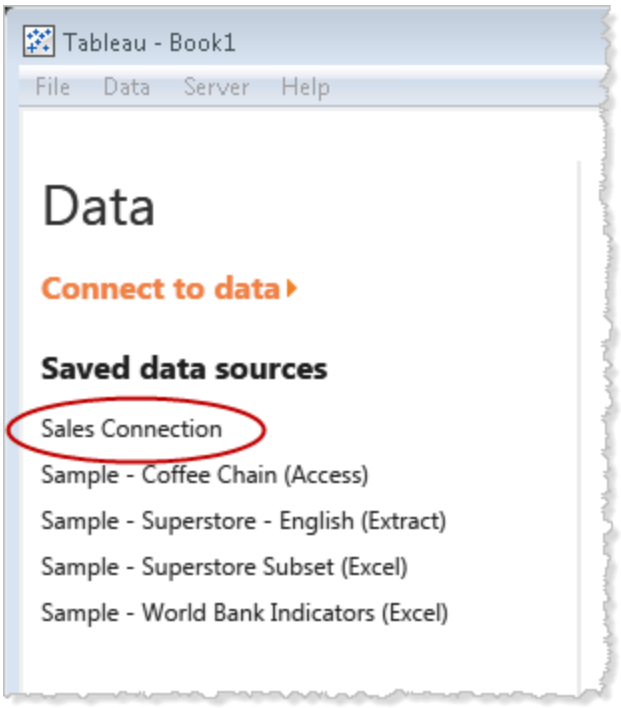

Upgrading Data Sources

If you have workbooks that were created before Tableau Desktop 8.2 and that use Microsoft Excel or text file data sources, or you are using the Excel or text file legacy connection, you might have the option to upgrade the data sources in your workbook. By upgrading your Excel and text file data sources you can take advantage of better data interpretation and compatibility on the Mac.

After upgrading these data sources, there may be changes to your workbook. For example, you may notice some fields have turned red, field names have changed, field types have changed, or the data values have changed. Resolve these changes before saving the workbook. Once the workbook is saved, the changes cannot be undone. If you encounter any of these types of changes, see the Resolving Data Source Upgrade Issues knowledge base article for specific examples of these changes and how to resolve them.

In some cases, the data source cannot be upgraded and you see a message indicating which data sources could not be upgraded. In these cases, do not save the workbook. See the Resolving Data Source Upgrade Issues knowledge base article to review alternative options that will allow you to continue using the workbook even when the Excel and text file data sources cannot be upgraded. Alternatively, close and reopen the workbook without upgrading the data sources.

Refreshing the Data

If a data source changes—for example, if new fields or rows are added , data values or field names are changed, or data is deleted, 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, provided the changes have been saved in the data source first.

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

If a field that is used in a Tableau worksheet is removed from the underlying data of the data source and then the data source is refreshed, 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 Data Sources

You can close a data source 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 data source is cleared from the Data window and all open worksheets associated with the data source are cleared. If you accidentally close a data source, use the Undo button to reopen. Close a data source by doing one of the following:

- Right-click the data source at the top of the Data window and select **Close**.

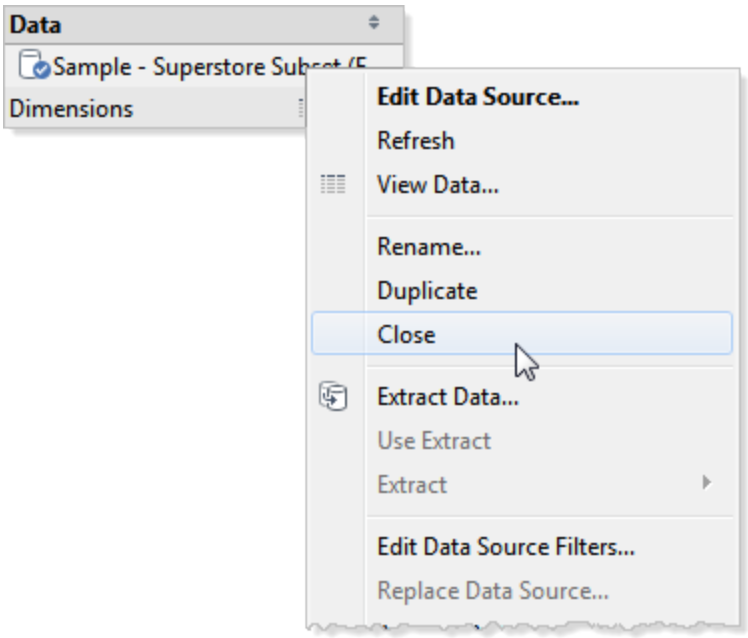

- Select a data source on the **Data** menu and then select **Close**.

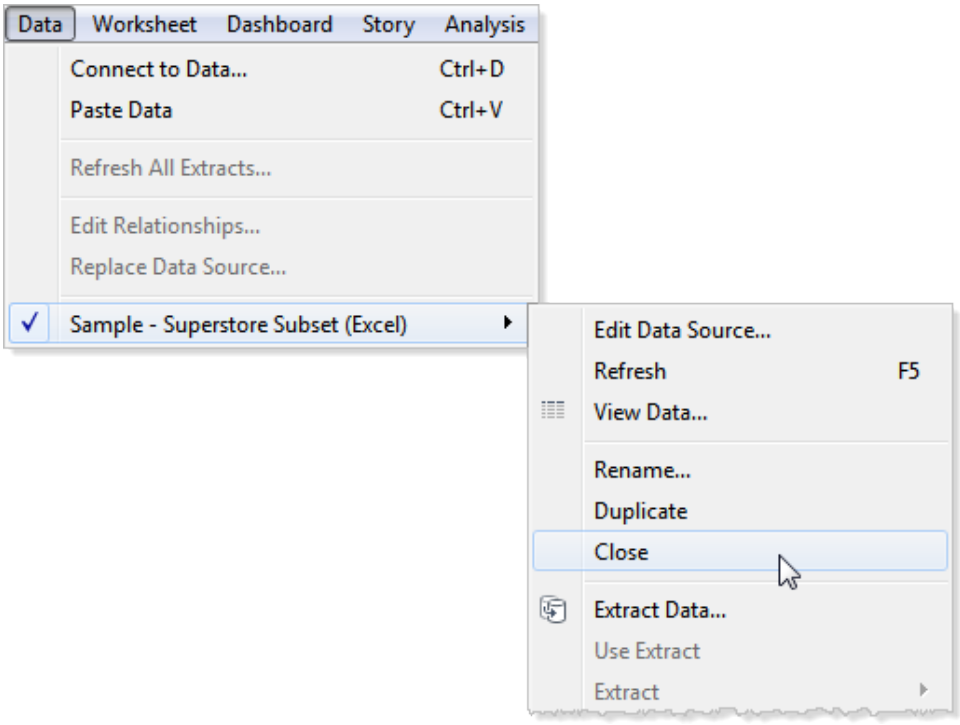

Clipboard Data Sources

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.

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

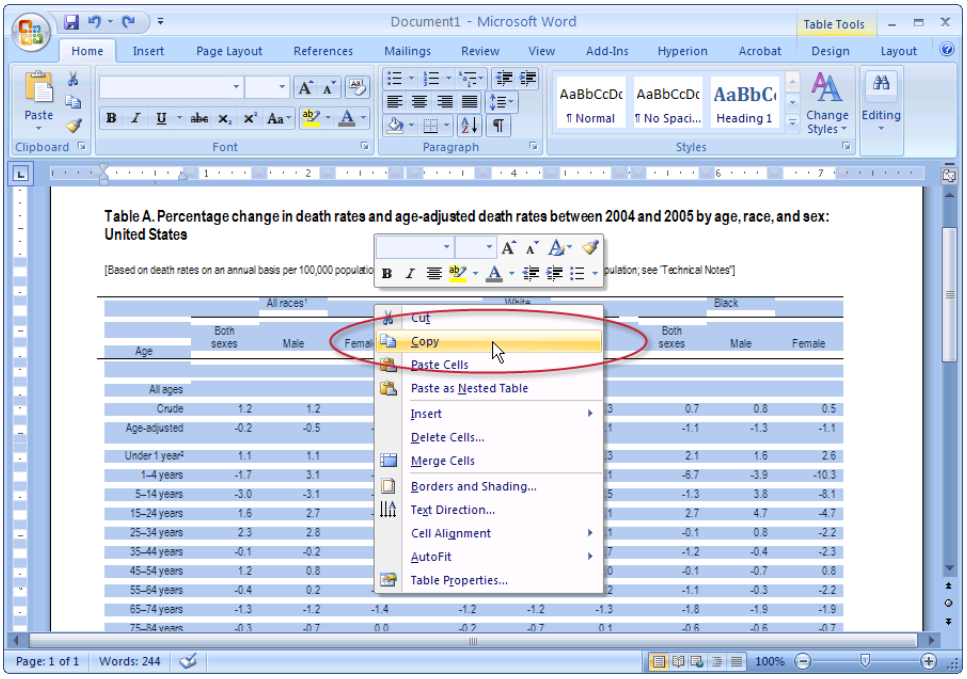

2. Open Tableau Desktop and select **Data** > **Paste Data**.

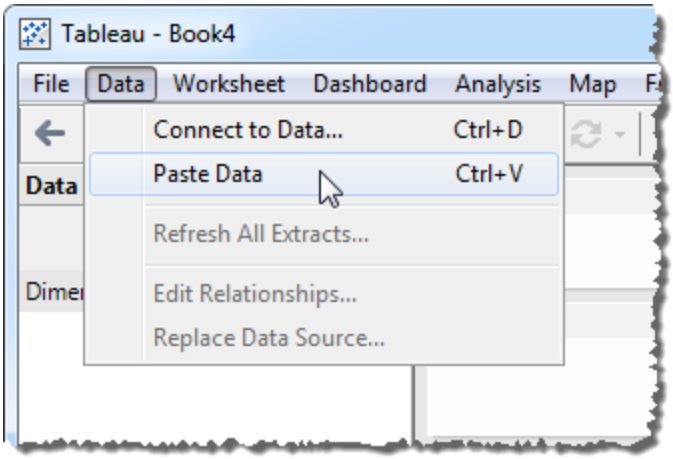

3. Select **File** > **Save** to save the data source. When you save the workbook the data source is automatically put into your repository. If you save as a packaged workbook the data sources is saved with the workbook instead.

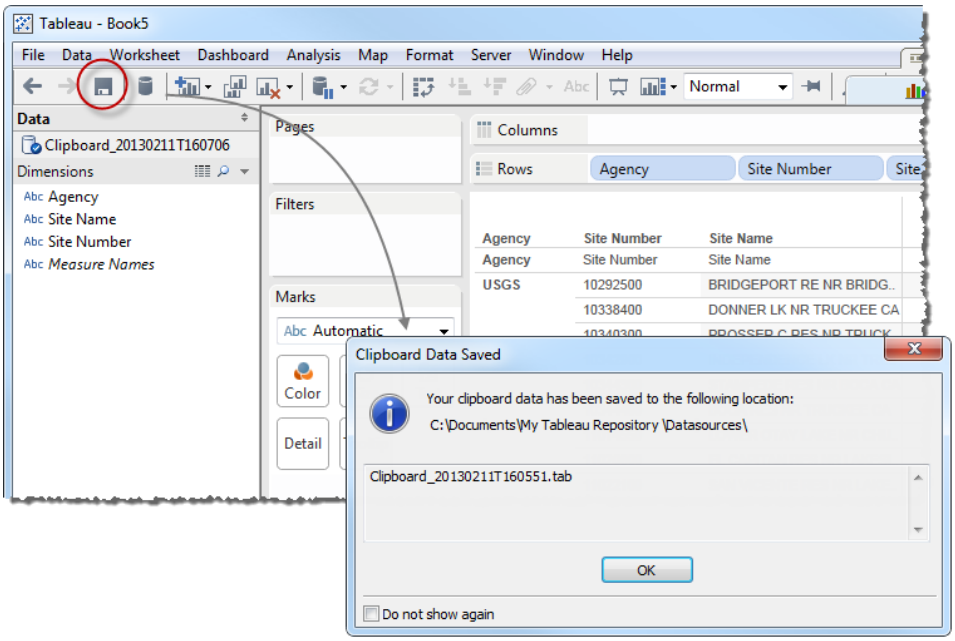

Working with Multiple Connections

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 blending. 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.

Understanding Data Blending

Data blending is when you blend data from multiple data sources on a single worksheet. The data is joined on common dimensions. Data Blending 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 168 to learn how to create those types of joins. Instead, data blending 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 define common dimensions between the primary and secondary data sources. 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 specified as a linking field. If the two dimensions 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.

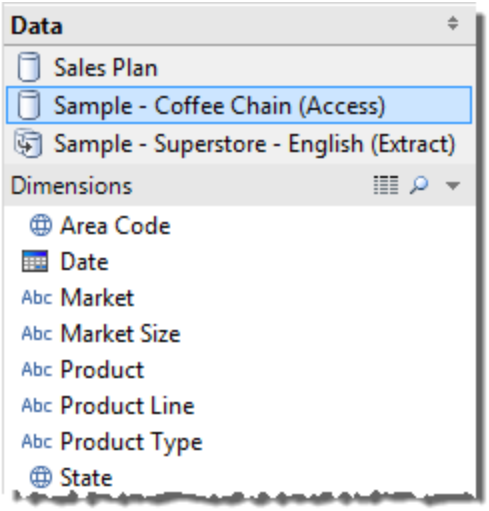

You can resize the connection list area in the Data window to save space. When you resize to a limited vertical height, the list is converted to a drop-down list.

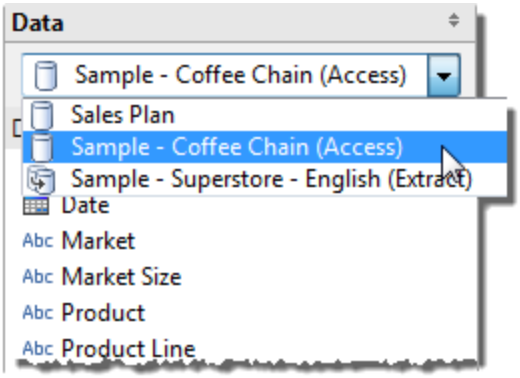

Each connection has an icon to indicate the type of connection. For example, the icon can indicate whether the data source is relational, multidimensional (cube), or a data extract. In Tableau, multidimensional data sources are supported only in Windows.

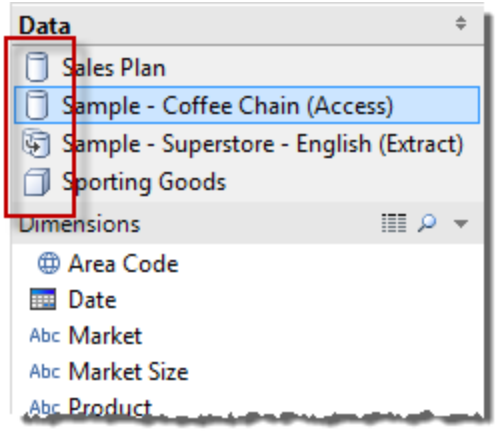

You can right-click the connections to access the commands that are on the Data menu. For example, you can right-click a connection and rename, export, or close it.

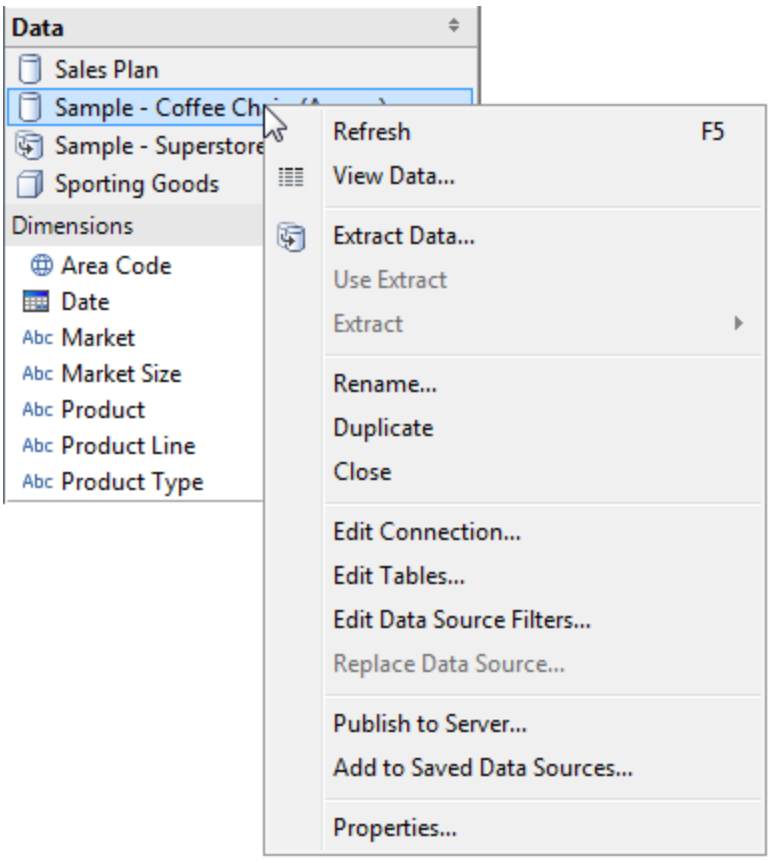

Adding a Secondary Connection

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.

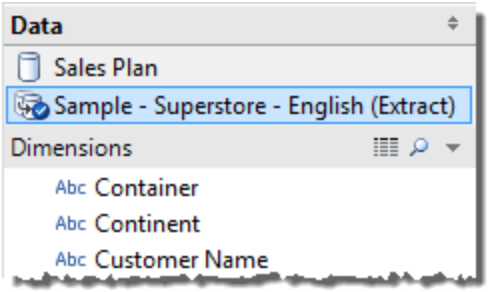

If you switch to another connection, you'll notice that the Data window is marked with an orange color to indicate that if you use fields from this connection, it will become the secondary connection.

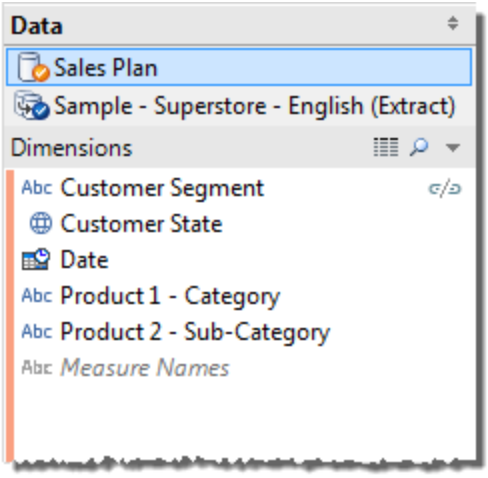

> Multidimensional data sources (cubes) cannot be used as the secondary data source. They can only be used as the primary data source.

To add a secondary connection:

1. Connect to the primary data source and build a view.
2. When you need additional data from a separate data source, select **Data** > **Connect to Data**.
3. Follow the steps in **Connecting to Data and Editing Data Sources** on page 52 to connect to the secondary data source.
4. 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.

Tableau automatically identifies fields in the secondary data source that also exist in the primary data source. These are potential linking fields are marked with a broken link icon . Click these link icons next to the fields that you want to blend. The icon is now orange and connected . You can also define your own custom relationships to handle columns that don't have matching names. See **Defining Relationships** below.

5. 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.

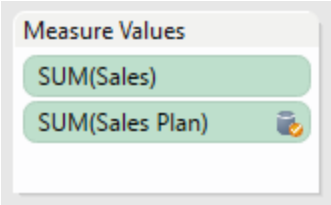

To remove a secondary connection:

- Right-click the secondary connection at the top of the Data window and select **Close**.

Defining Relationships

Tableau automatically recognizes when a field from the primary data source also exists in a secondary data source. These fields are marked with a gray broken link icon in the Data window to indicate that they are potential linking fields. When you click the broken link icon, the link becomes orange and is connection. An automatic relationship is created between the two data sources. The relationship allows you to blend data from both data sources on a single sheet. You must have a linked field in order to use data from the secondary data source. For example, a workbook may have 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 (blended) until one of those common fields has been linked.

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

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

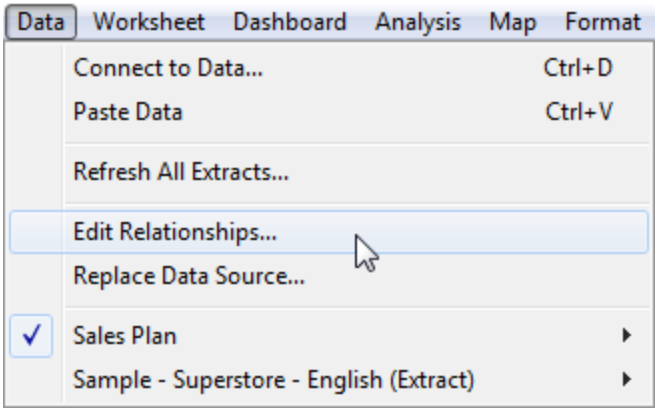

2. Make sure that the primary data source is selected from the down list.

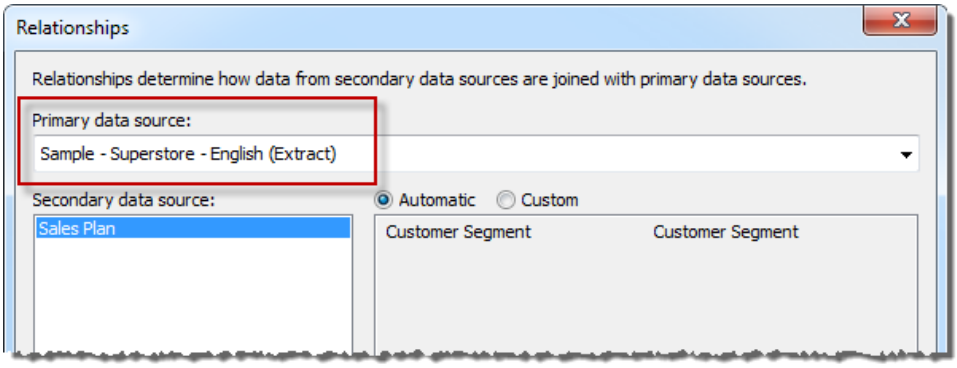

3. Select a secondary data source from the list on the left.

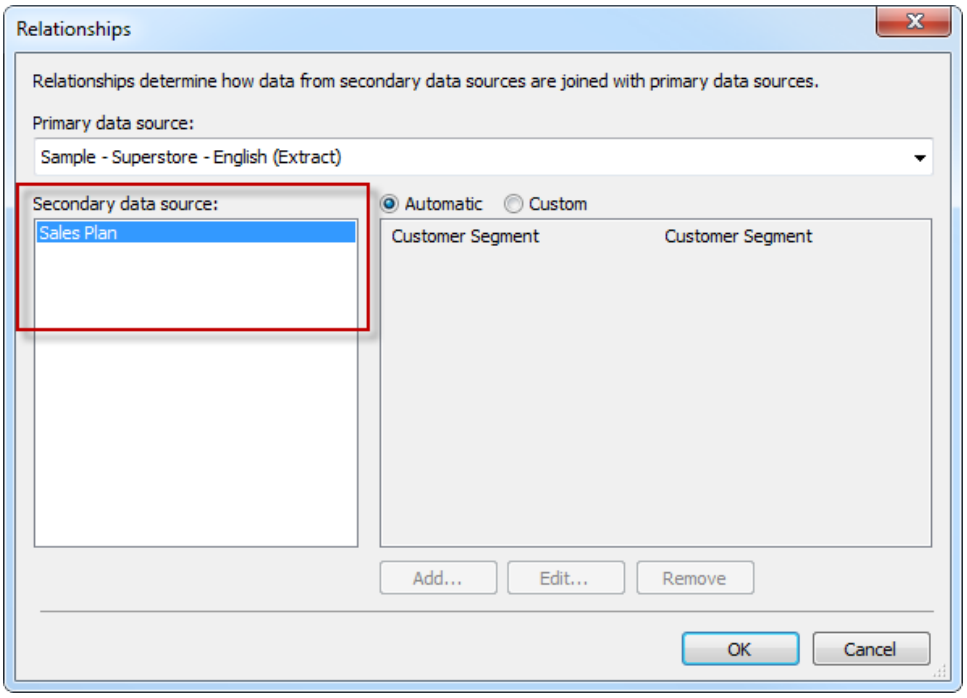

4. Select **Custom** at the top of the relationships list.
5. Click the **Add** button or select an existing relationship and click the **Edit** button.

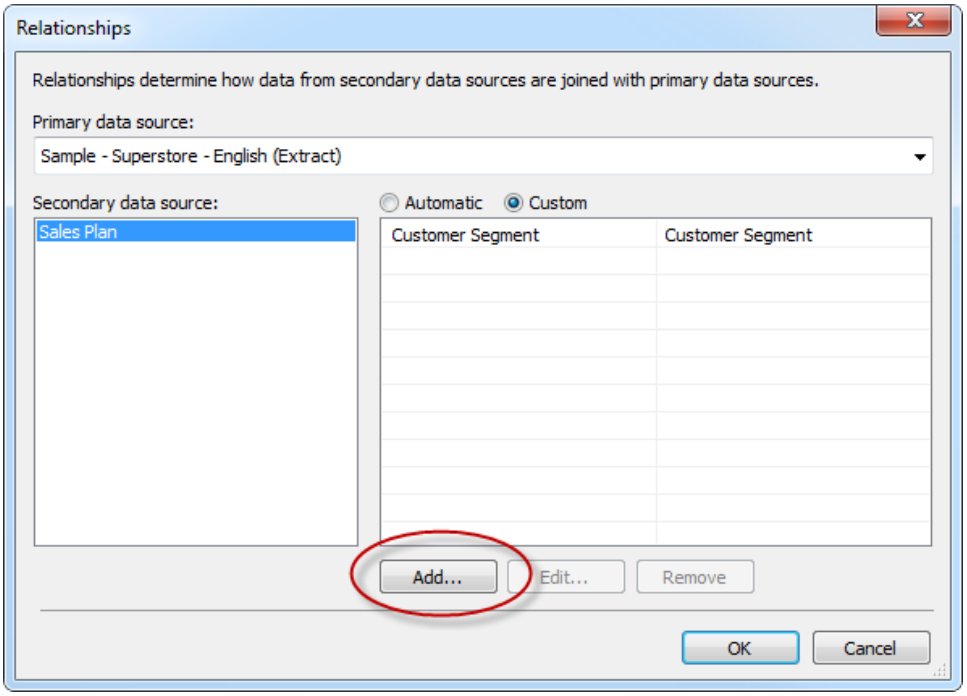

6. Select a field in the primary data source and map it to matching field in the secondary data source. In this example, the primary data source has a field called "State" while the

secondary data source has a field called "Customer State." These two fields can be mapped together to create a relationship even though they don't have the same name.

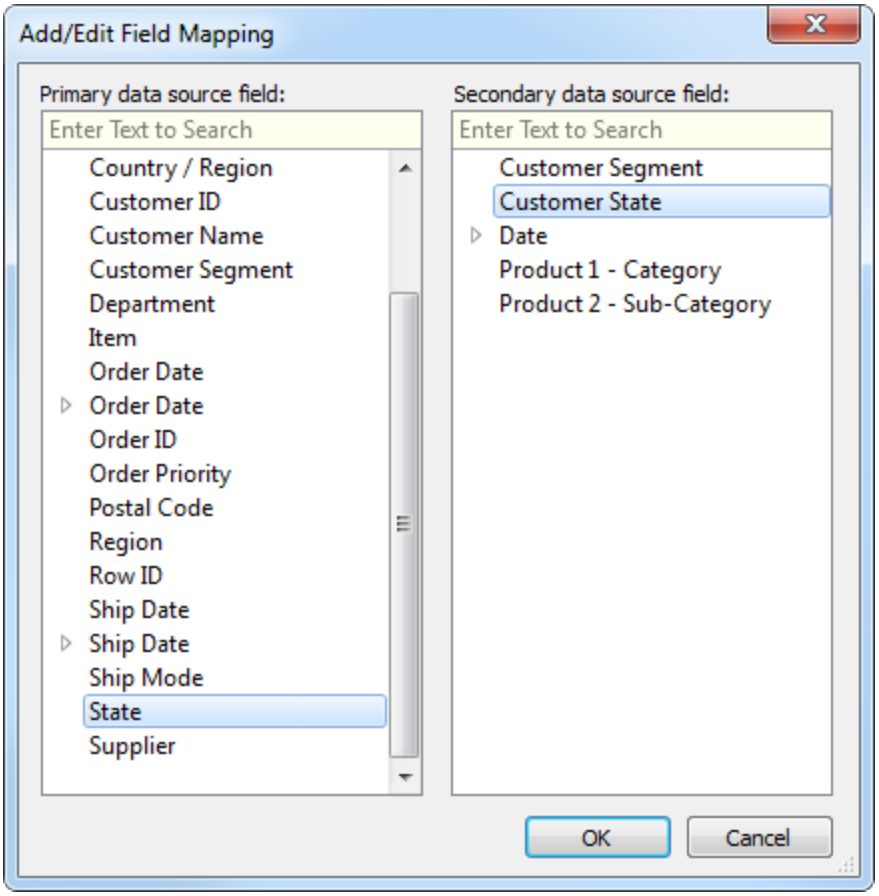

7. When finished, click **OK**.
8. Add and Remove as many relationships as necessary, and when finished, click **OK** again.

   The related fields are shown in the secondary data source as potential linking fields. Click the link icon next to these fields in the data window to make the links active. Linked fields are indicated with an orange and connected link icon in the Data window. When the related field from the primary data source is used in the view, the link becomes active automatically.

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 Blending on a Worksheet

This example uses two data sources: Sample - Superstore that comes with the application and an auxiliary Excel file that contains forecasted sales information. This example shows how to

blend data from both of these data sources to create a single view showing actual vs. planned sales.

1. Connect to the Sample - Superstore extract and build a view that shows Sales by Customer Segment and Department.

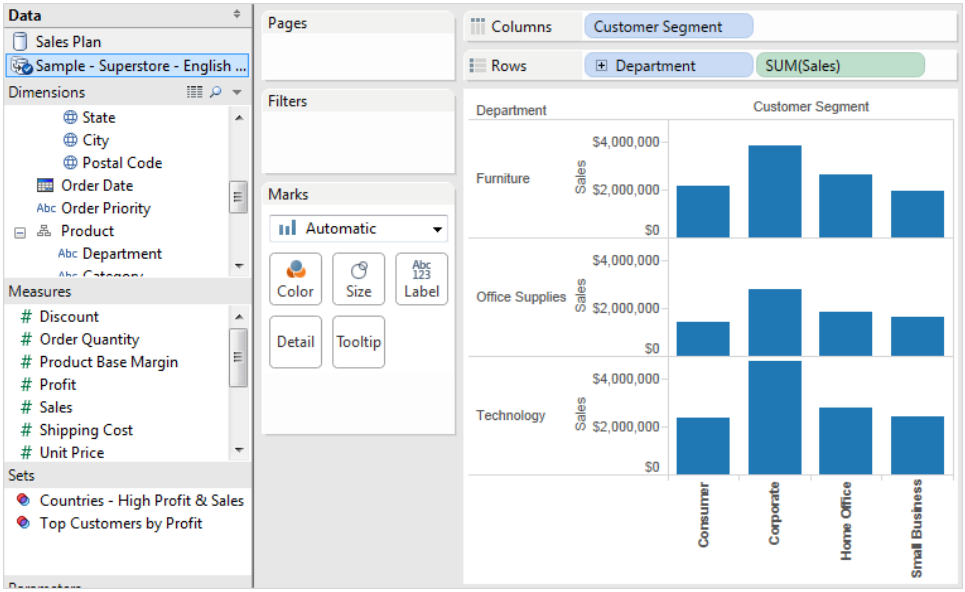

2. Select **Data** > **Connect to Data** and connect to the Sales Plan spreadsheet.
3. Switch to the Sales Plan connection and click the link icon next to Customer Segment.

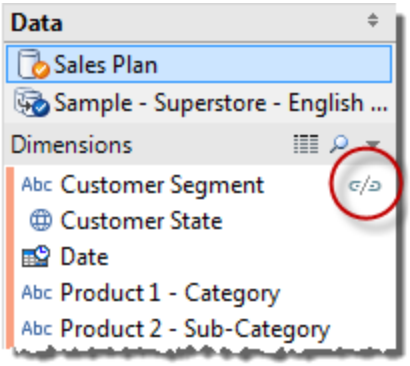

4. Drag the **Sales Plan** measure to Detail on the Marks card.

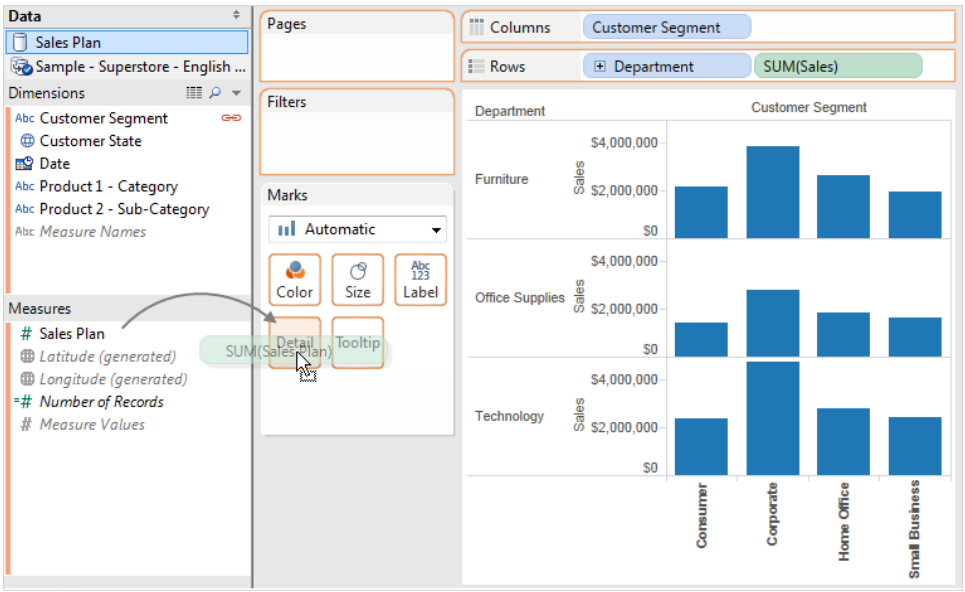

5. Right-click the Sales axis and select **Add Reference Line**.

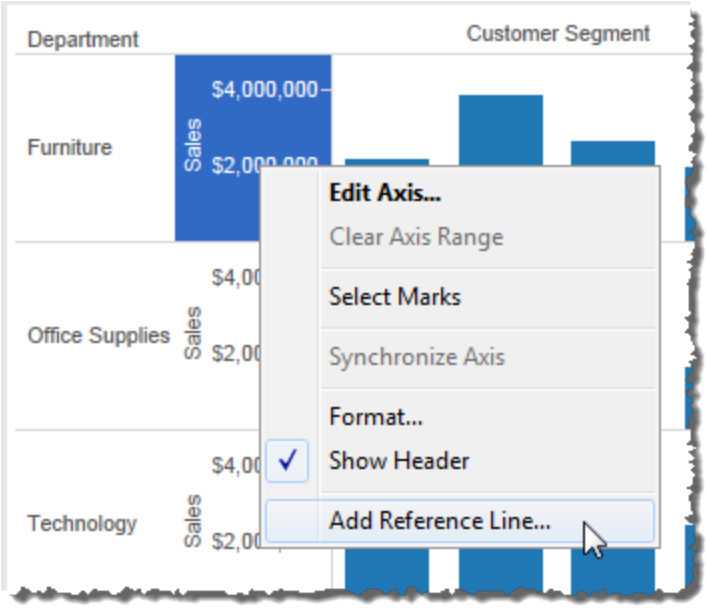

6. In the Reference Line dialog box, add a reference line that shows SUM(Sales Plan) per cell. When finished, click **OK**.

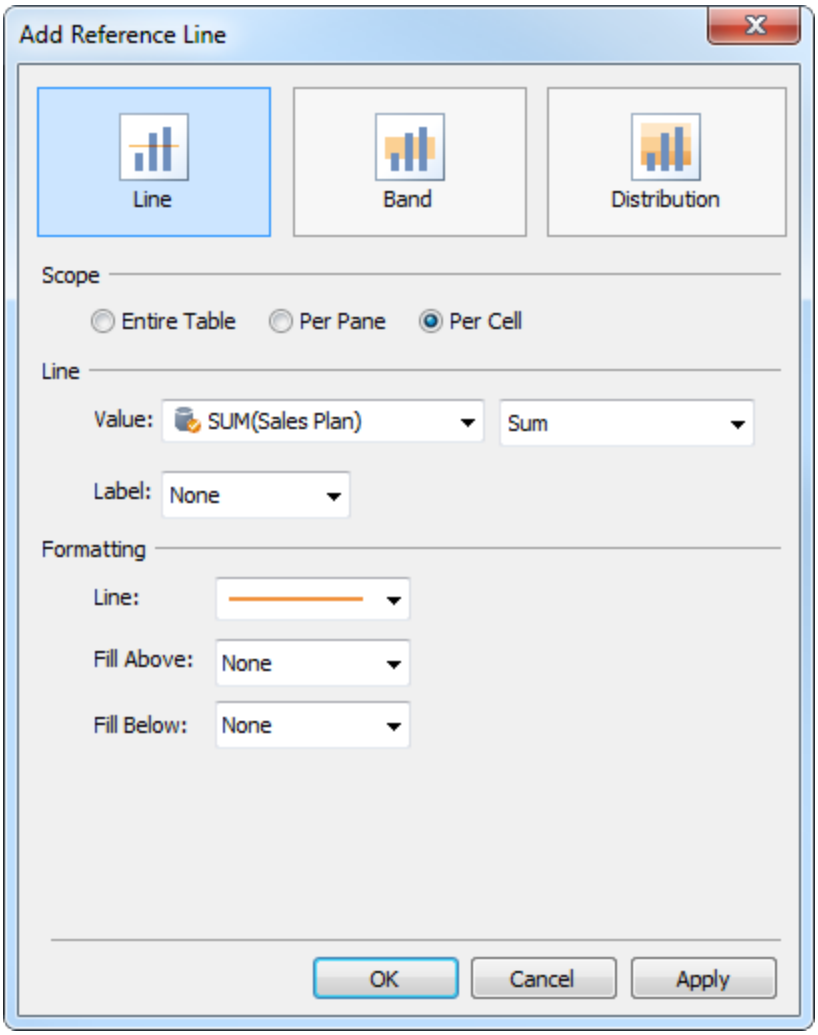

7. The Worksheet is now pulling data from the secondary data source (Sales Plan) to show how actual sales compared to the planned sales.

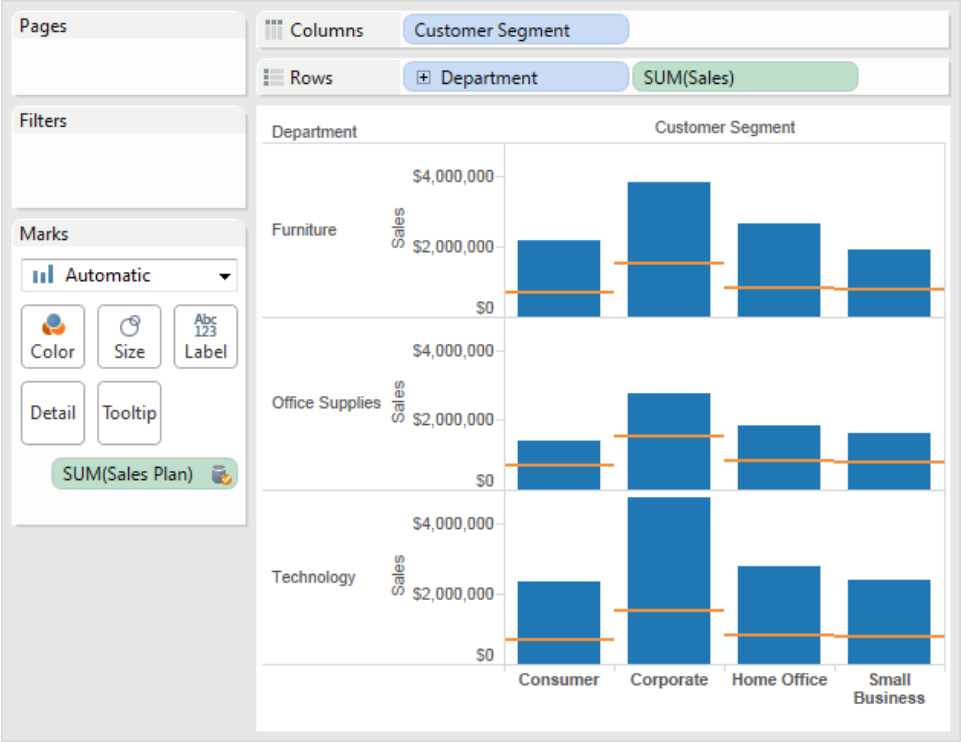

Troubleshooting Data Blending

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

Warning: No Relationship to the Primary Data Source

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

Fields cannot be used from the *[name of secondary data source]* data source, because there is no relationship to the primary data source. in the Data window, switch to the *[name of secondary data source]* data source, and click at least one link icon to blend these data sources.

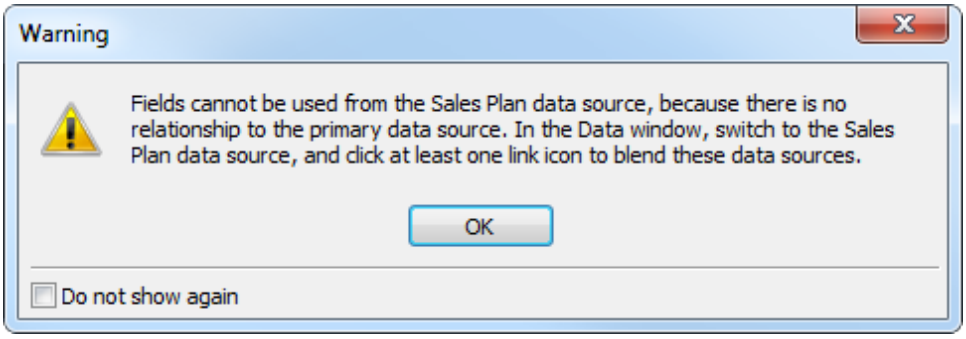

This warning occurs when you have no active links in the secondary data source. For example, if you have two data sources with the related dimensions State and Date, at least one of these fields needs to have the active orange link icon next to them in the secondary data source. You can activate a link by clicking the link icon in the data window or by using the related field from the primary data source in the view.

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. See **Defining Relationships** on page 159.

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.

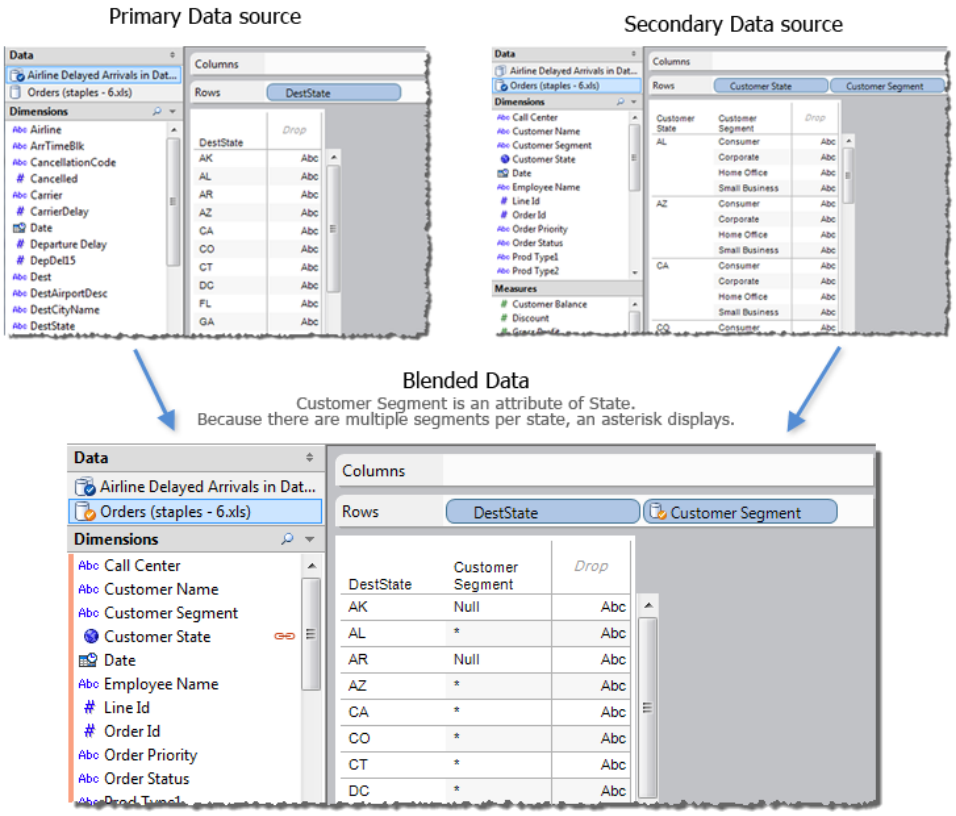

All secondary fields are aggregated. Dimensions are aggregated as Attributes (ATTR), which means that if there's only one member it will show the member value but if there are multiple members it will show an asterisk.

Nulls Appear After Blending Data

Nulls can sometimes appear for some records when you use data blending. This happens when no corresponding record exists in the secondary data source for some records in the primary data source or if the data types for some records in the two data sources are different. You can resolve this problem by inserting data in the secondary data source such that all records in the primary data source have a match, and then verifying that the data types in the data sources match. For example, when joining on State, if Database 1 has values all 50 U.S. states and Database 2 has values for only 30 U.S. states, add values for the remaining 20 U.S. states to Database 2 before blending the data.

This problem can happen because data blending works by performing a left join on the linking field, which means that Tableau will take all the data in the left table (the primary data source) and match the data in the right table (the secondary data source). The join result will be null for the secondary data source when no match exists for a particular member of the primary data source.

Using a Multidimensional (Cube) Data Source

When blending data in Tableau, multidimensional data sources cannot be used as secondary data sources. These types of databases can only be used as the primary data source.

Joining Tables

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.

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.

Adding or Editing Tables

You may want to add or 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 underlying data.

1. On the **Data** menu, select a data source, and then select **Edit Data Source**.
2. Drag a table to the join area.

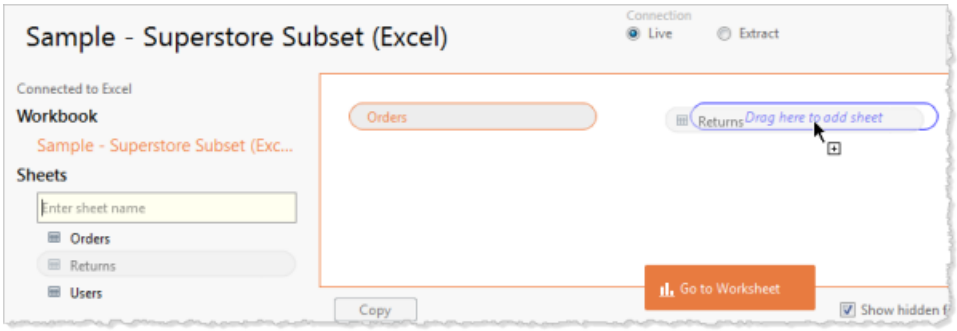

Alternatively, click the join icon to modify an existing join.

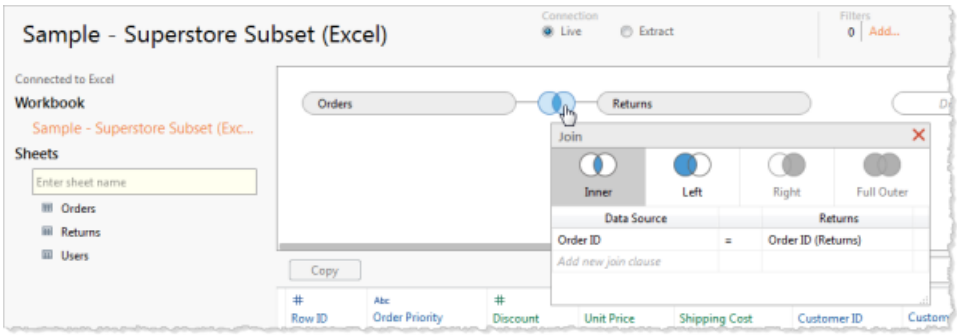

When you connect to multiple tables you are essentially connecting to a denormalized view of the data source. This means that in some cases, queries are run against all tables and it is possible for some measures to be over-counted. For example, suppose you have an employees table and an orders table. 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.

Combining Multiple Tables

When working with relational databases (for example, PostgreSQL) or files (for example, a text file), you can combine data that exists across multiple tables or files by creating joins. Using joins to combine tables allows you to analyze data that have a relationship with each other.

> **Note:** The tables you can join and the different join types you can use depend on database or file your workbook is connected to.

1. Connect to data and create your data source according to the examples in Examples - Connecting to Data Sources.
2. After you select the file, database, or schema, double-click or drag a table to the join area

of the data source page.

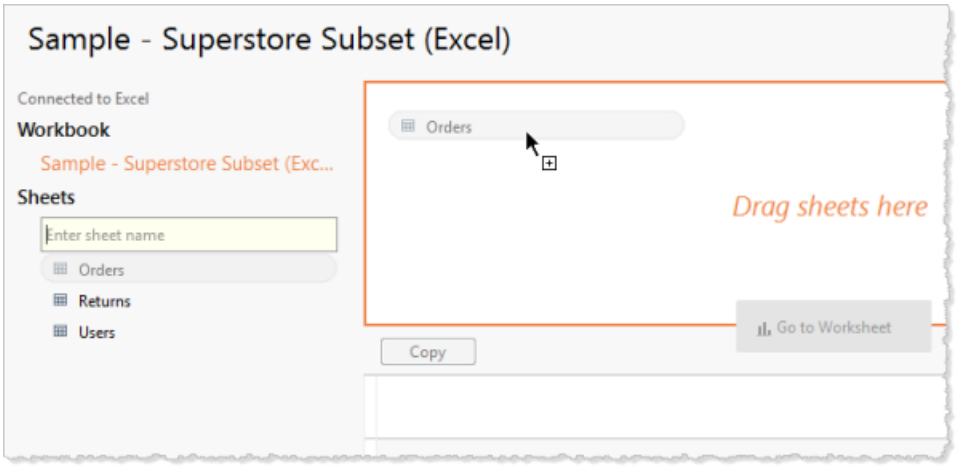

3. Double-click or drag another table to the join area. The join dialog box opens.

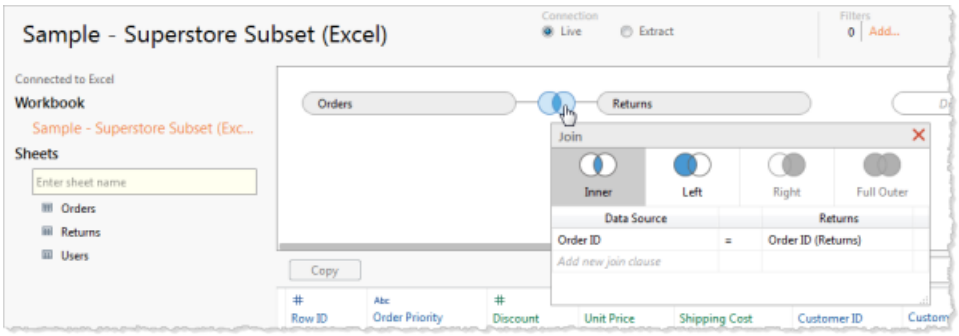

4. Add one or more join conditions by selecting a field from one of the available tables used in the data source, a join operator, and a field from the added table. Inspect the join condition to make sure it reflects how you want to connect the tables.

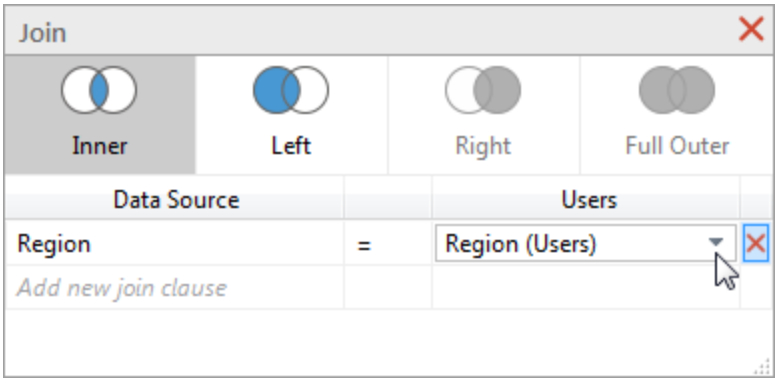

For example, in a data source that has a table of order information and another for users information, you could join the two tables based on the Region field that exists in both tables. Select the type of join.

Note: You cannot nest inner joins within left or right joins. These joins will cause a "join expression not supported" error.

5. Select Region in both lists of fields, and then select the equal sign as the operator.

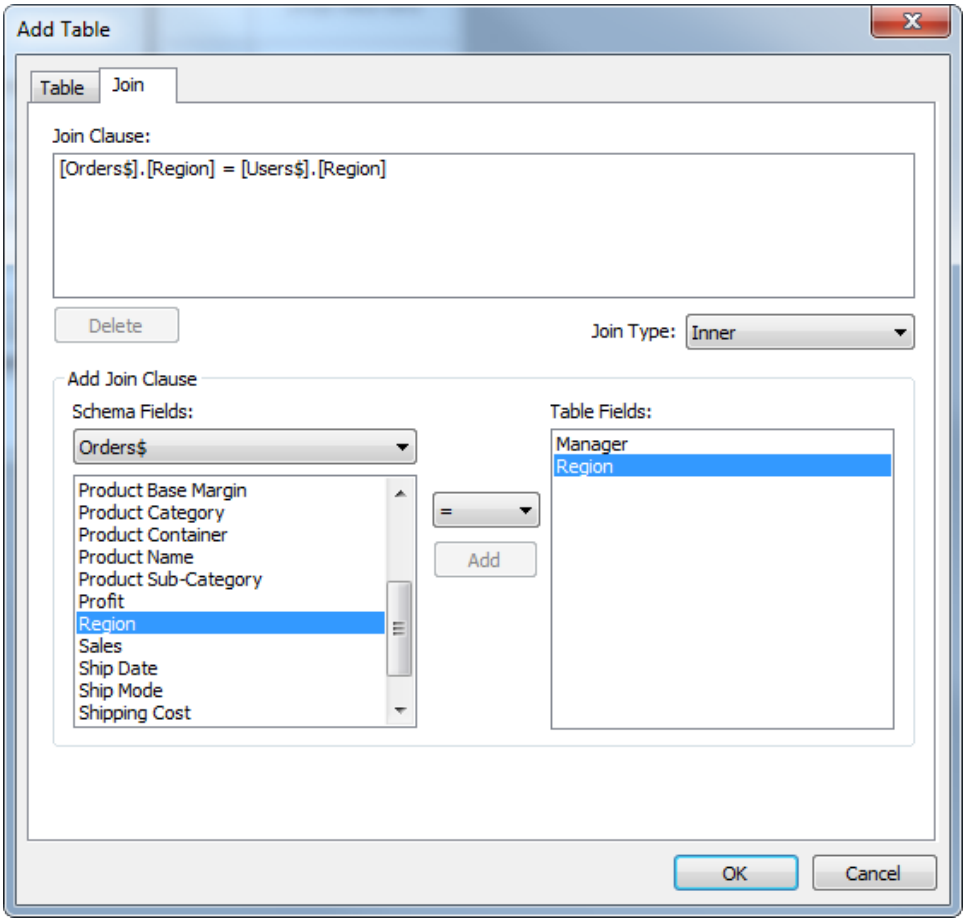

6. When finished, click the "x" icon to close the join dialog box.

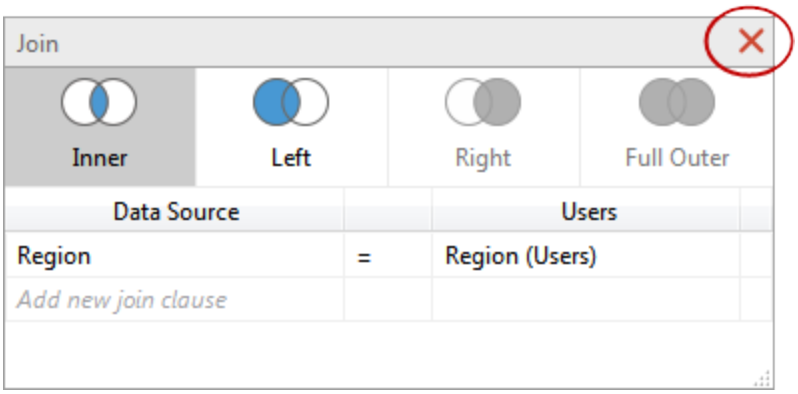

> **Note:** You can delete an unwanted join condition by click the red "x" that displays when you hover over the right-side of the condition.

When you have joined tables in your data source, 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 by selecting one of the **Sort by** options on the Data window menu. This option is only available if you have connected to the database live. If you have imported the data, the fields are no longer grouped by table.

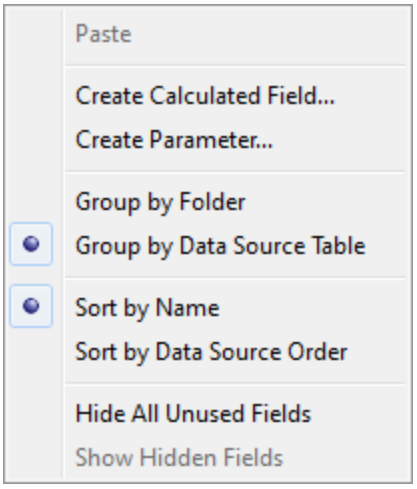

Assuming Referential Integrity

In some cases, you can improve query performance by selecting the option to **Assume Referential Integrity**. When selected, Tableau will only include the joined table in the query if it is specifically referenced by fields in the view. For example, imagine connecting to a sales database that has two tables: Product Catalog and Sales. These two tables are shown below:

Product Catalog

Product ID (Primary Key)	Product Name
1	10 Inch Tablet
2	Smart Phone
3	Desk Lamp
4	Memory Stick

Sales

Product ID (Foreign Key)	Sale Amount	Transaction Date
1	100	10/1/2012
1	2000	10/2/2012
2	50	9/30/2012
3	10	8/21/2012

Because all products that are sold must have a listing in the Product Catalog, every row in the Sales table has a matching row in the Product Catalog table. When these two tables are joined on Product ID, you end up with a table that looks like this:

Product ID	Product Name	Sale Amount	Transaction Date
1	10 Inch Tablet	100	10/1/2012
1	10 Inch Tablet	2000	10/2/2012
2	Smart Phone	50	9/30/2012
3	Desk Lamp	10	8/21/2012

Now let's say you build a view to look at Sale Amount by Region. By default, dragging the Sale Amount field to the view may create a query like this:

```
SELECT SUM([Sales Amount] FROM [Product Catalog] P INNER JOIN S
ON S.ProductID = P.ProductID
```

However, when you select the option to **Assume Referential Integrity**, you are telling Tableau that the joined tables have referential integrity. In other words, you are confirming that the Sales table will always have a matching row in the Product Catalog table. Because that is true, Tableau doesn't need any information from the Product Catalog table in order to return these results. When you drag the Sales Amount field into the view, Tableau can simplify the query to:

```
SELECT SUM([Sales Amount]) FROM [Sales]
```

This simplified query can often return quicker results because it removes the join operation. This option only impacts inner joins and does not affect single table connections.

> This option will not help improve query performance if used in combination with context filters or data source filters.

Extracting Data

Extracts are saved subsets of a data source that you can use to improve performance, upgrade your data to allow for more advanced capabilities, and analyze offline. You can create an extract by defining filters and limits that include the data you want in the extract. After you create an extract you can refresh it with data from the original data source. You can either fully refresh the data, replacing all of the extract contents; or you can increment the extract; which only adds rows that are new since the last refresh.

Extracts can:

- Improve performance. For file based data sources such as Excel or Access, a full extract takes advantage of the Tableau data engine. For large data sources, a filtered extract

can limit the load on the server when you only need a subset of data.

- Add functionality to file based data sources, such as the ability to compute Count Distinct.
- Provide offline access to your data. If you are traveling and need to access your data offline, you can extract the relevant data to a local data source.

Creating an Extract

1. Select a data source on the **Data** menu and then select **Extract Data** to open the Extract Data dialog box.

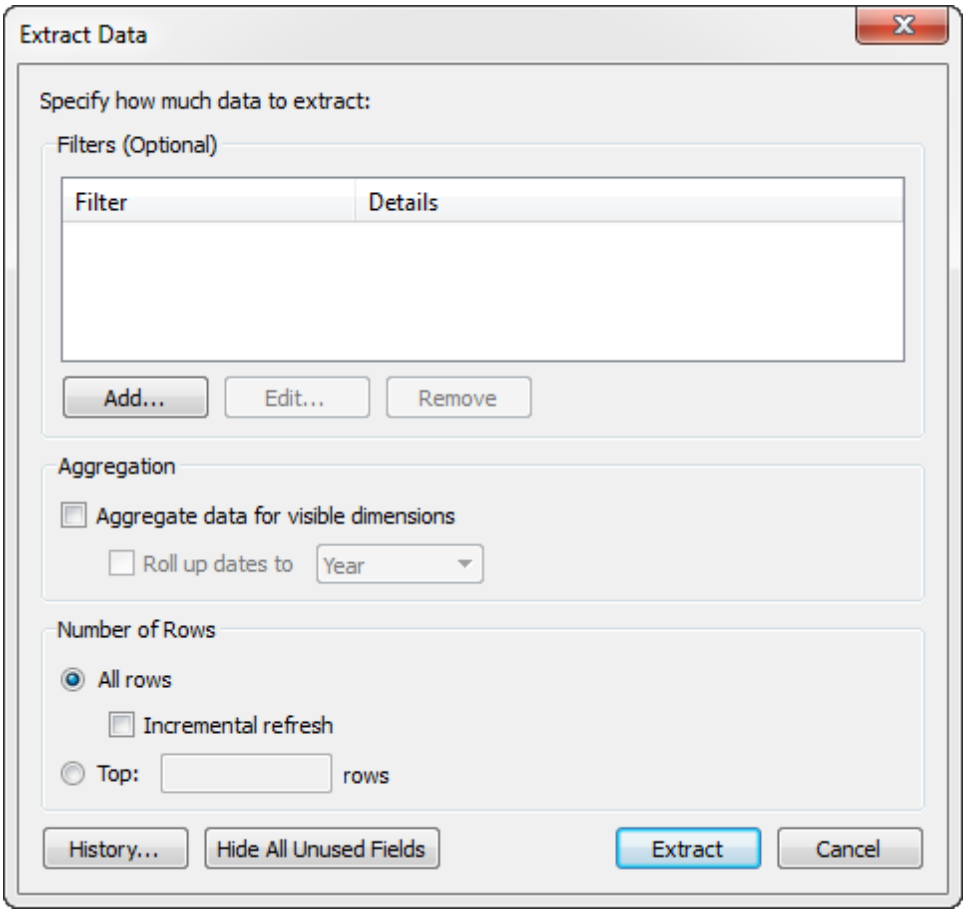

2. Optionally define filters to limit the data that will be extracted. Any fields that are hidden in the Data window will be automatically excluded from the extract. Click the **Hide All Unused Fields** button to quickly remove them from the extract.

   To add filters, click the **Add** button under the Filters list.

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.

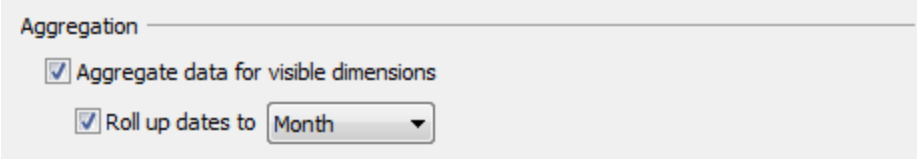

The examples below show how the data will be extracted for each aggregation option.

<table>
<tr><td>Original Data</td><td>
<table>
<tr><th></th><th>A</th><th>B</th><th>C</th></tr>
<tr><td>1</td><td>Date</td><td>Region</td><td>Sales</td></tr>
<tr><td>2</td><td>1/1/2009</td><td>South</td><td>$500</td></tr>
<tr><td>3</td><td>1/1/2009</td><td>West</td><td>$200</td></tr>
<tr><td>4</td><td>1/1/2009</td><td>West</td><td>$100</td></tr>
<tr><td>5</td><td>1/1/2009</td><td>East</td><td>$300</td></tr>
<tr><td>6</td><td>1/2/2009</td><td>South</td><td>$600</td></tr>
<tr><td>7</td><td>1/2/2009</td><td>South</td><td>$400</td></tr>
<tr><td>8</td><td>1/2/2009</td><td>East</td><td>$100</td></tr>
<tr><td>9</td><td></td><td></td><td></td></tr>
</table>
</td><td>Each record is shown as a separate row. There are 7 rows in the data source.</td></tr>
<tr><td>Aggregate Data
(no roll up)</td><td>
<table>
<tr><th></th><th>A</th><th>B</th><th>C</th></tr>
<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>
<tr><td>7</td><td></td><td></td><td></td></tr>
</table>
</td><td>Records with the same date and region have been aggregated into a single row.There are 5 rows in the data source.</td></tr>
<tr><td>Aggregate Data
(roll up dates to Month)</td><td>
<table>
<tr><th></th><th>A</th><th>B</th><th>C</th><th>D</th></tr>
<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>
</table>
</td><td>Dates have been rolled up to the Month level and records with the same region have been aggregated</td></tr>
</table>

		into a single row.There are 3 rows in the data source.

4. Select the number of rows you want to extract. You can extract All, and the top N rows. Tableau first applies any filters and aggregation and then extracts the number of rows from the filtered and aggregated results.

   The number of rows options depend on the type of data source you are extracting from. For example, not all data sources support sampling so that option is not always available.

5. When finished, click **Extract**.
6. In the subsequent dialog box, select a location to save the extract into and give the file a name. Then click **Save**.

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 may 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, thus avoiding long queries every time you place a field on the shelf. You can toggle between using the extract and using the entire data source by selecting a data source on the **Data** menu and then selecting **Use Extract**.

You can remove an extract at anytime by selecting a data source on the **Data** menu and then selecting **Extract** > **Remove**. When you remove an extract you can choose to Remove the extract from the workbook only or Remove and delete the extract file, which will delete the extract from your hard drive.

You can see when the extract was last updated and other details by selecting a data source on the **Data** menu and then selecting **Extract** > **History**.

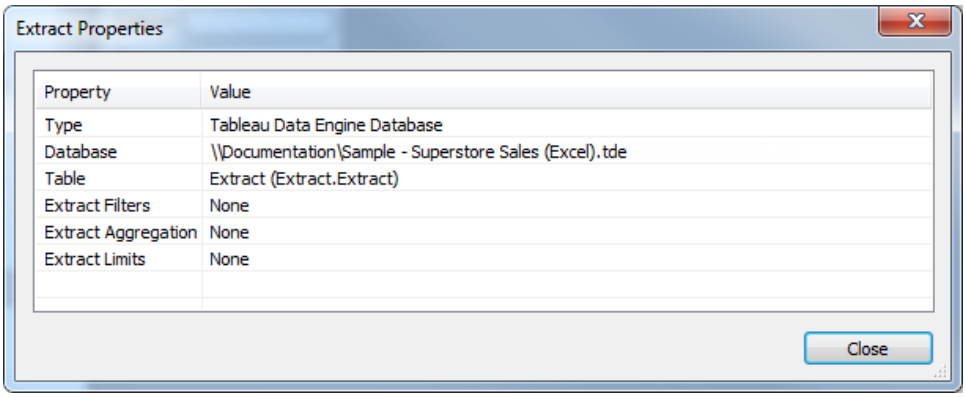

Refreshing Extracts

When the underlying data changes, you can refresh the extract by selecting a data source on the **Data** menu and then selecting **Extract** > **Refresh**. Extracts can be configured to be fully refreshed, replacing all of the data with what's in the underlying data source, or incrementally refreshed, adding just the new rows since the last refresh.

Full Extracts

By default, extracts are fully refreshed. That means that every time you refresh the extract, all of the rows are replaced with the data in the underlying data source. While this kind of refresh ensures you have an exact copy of what is in the underlying data source, it can sometimes take a long time and be expensive on the database depending on how big the extract is.

If the extract is not set up for incremental extract, selecting refreshing the extract will fully refresh the extract. If you're publishing the data source to Tableau Server, you can specify the type of refresh in the Scheduling & Passwords dialog box.

Incremental Extracts

Rather than refreshing the entire extract, you can set it up to only add the rows that are new since the last time you extracted data. For example, you may have a data source that is updated daily with new sales transactions. Rather than rebuild the entire extract each day, you can just add the new transactions that occurred that day. Then once a week you may want to do a full refresh just to be sure you have the most up to date data.

Follow the steps below to set up an extract to be incrementally refreshed.

1. Select a data source on the **Data** menu and then select **Extract**.
2. In the Extract Data dialog box, select **All rows** as the number of Rows to extract. Incremental refresh can only be defined when you are extracting all rows in the database. You cannot increment a sample extract.
3. Select **Incremental refresh** and then specify a column in the database that will be used to identify new rows. For example, if you select a Date field, refreshing will add all rows whose date is after that last time you refreshed. Alternatively, you can use an ID column

that increases as rows are added to the database.

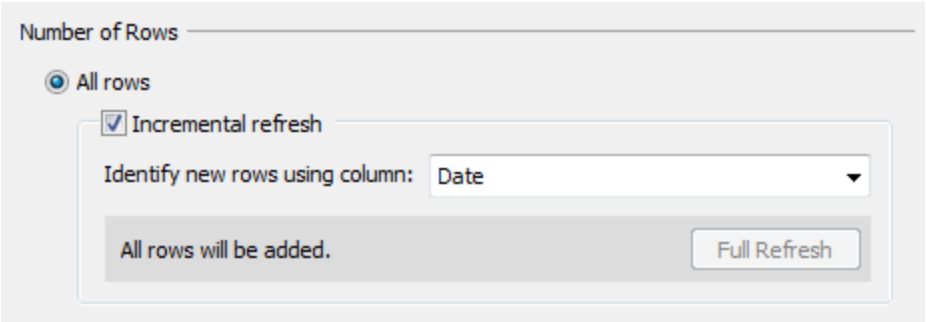

4. When finished, click **Extract**.

The steps above can be used to define a new extract or configure an existing extract for incremental refresh. If you are editing an existing extract, the last refresh is shown so you can be sure you are updating the extract with the correct data.

If you publish the data source to Tableau Server you can specify a schedule for incremental refresh as well as full refresh in the Schedules & Passwords dialog box.

Extract History

You can see a history of when the extract was refreshed by selecting a data source on the **Data** menu and then select **Extract** > **History**.

The Extract History dialog box shows the date and time for each refresh, whether it was full or incremental, and the number of rows that were added. If the refresh was from a file, it also shows the source file name.

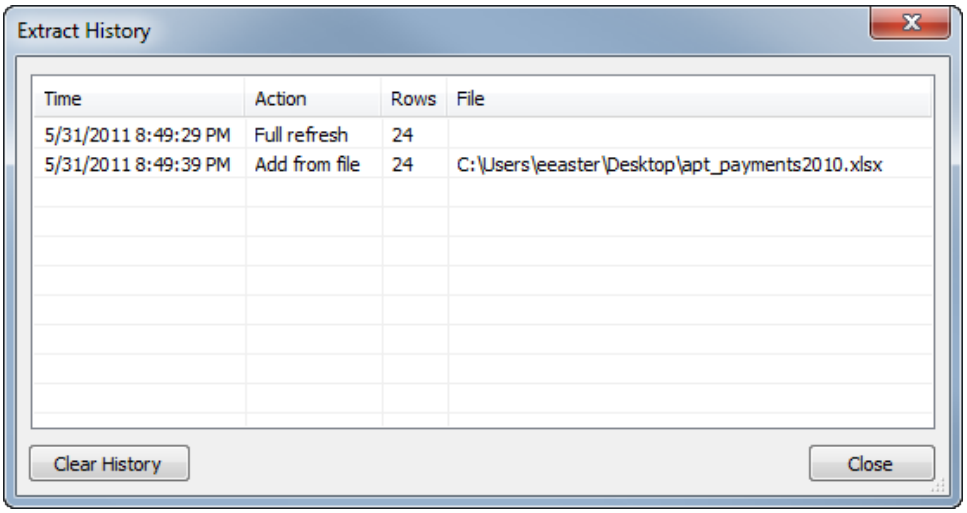

Adding Data to Extracts

There are two ways you can add new data to an extract: from a file or from a data source. However, to add new data, the columns in the file or data source must match the columns in

the extract.

Add data from a file

You can add new data to an extract from a file-based data source. Use this option when the file type of the extract is the same asthe file type of the data that you want to add. Alternatively, you can add data from a Tableau data extract (.tde) file. For example, you may have text files that are generated for a task that is performed every day. To add each day's worth of information to your extract whose original data source is also a text file, use the **Append Data from File** command.

To add data from a file

1. On the **Data** menu, select a data source, and then select **Extract** > **Append Data from File**.
2. Browse to and select the file that has the new data.

   > **Note:** By default, the file format of the extract's original data source is used. To add data from a Tableau data extract, click the file format drop-down list, and then select **Tableau Data Extract (.*tde)**.

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

Add data from a data source

You can also add new data to an extract from another data source in the workbook. Use this option when the file type of the extract is different from the file type of the data you want to add. For example, you created an extract from a data warehouse that has the past ten years worth of data. However, new data has been kept in an Excel workbook. You can add new data to the extract by using the **Append Data from Data Source** command.

> **Note:** Joins or custom SQL should be specified in the data source before adding data to the extract.

To add data from a data source

1. On the **Data** menu, select a data source, and then select **Extract** > **Append Data from Data Source**.
2. In the dialog box, select the data source that you want to append.
3. When finished, click **OK**.

Using either option will add new rows to the extract. To see a summary of the number of rows that were added, select a data source on the **Data** menu and then select **Extract** > **History**.

> **Note:** When you refresh this extract, the data will be replaced with the data from the original data source.

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.

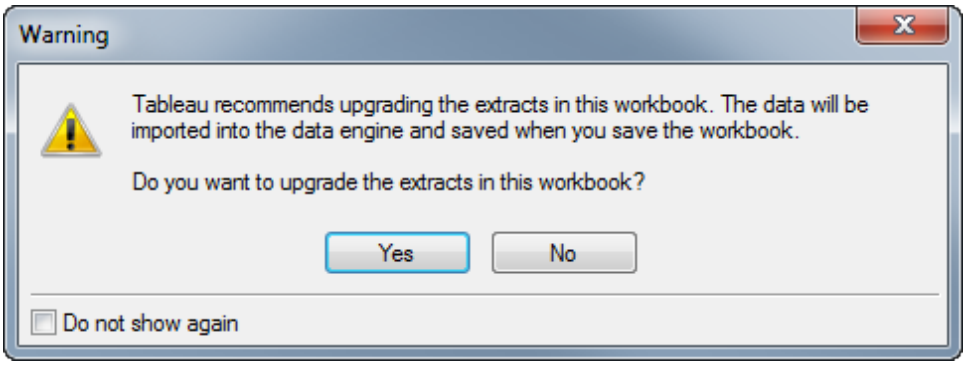

You can also upgrade the extracts by selecting a data source on the **Data** menu and then selecting **Upgrade Extract**.

Optimizing Extracts

To improve performance when working with extracts you can optimize the extract. Optimizing an extract creates secondary structure in the extract that speed up future queries.

Optimize the extract by selecting a data source on the **Data** menu and then selecting **Extract** > **Optimize**.

The following types of optimizations are made:

Materialized Calculated Fields

Calculated fields are computed in advance and stored in the extract. In future queries, Tableau can look up the already computed value rather than running the computation again. The following types of calculated fields ARE NOT materialized:

- Calculations that use unstable functions such as NOW() and TODAY()
- Calculations that use external functions such as RAWSQL and R
- Table calculations

In addition, if the formula for a materialized calculation changes or the calculation is deleted from the data source, the materialized calculation is dropped from the extract until the extract is optimized again.

Acceleration Views

When a workbook contains filters that are set to show only relevant values, computing the available values for that filter can be an expensive query. For these filters, Tableau must evaluate the other filters in the workbook first and then compute the relevant values based on their filter selections. To speed up these queries, a view can be created that computes the possible filter values and caches them for faster lookup later.

Updating Extracts on Tableau Server

You have the following options for updating extracts published to Tableau Server or Tableau Online:

- You can add the extract or a workbook that connects to it to a refresh schedule in Tableau Server or Tableau Online (cloud-based data sources only).
- You can update the extract in Tableau Desktop and then republish it.
- You can add to or refresh the extract in Tableau Server or Tableau Online without first adding to or refreshing the extract in Tableau Desktop.

The remainder of this topic describes the third option.

Refreshing Extracts Using Tableau Desktop

Before you attempt to update an extract, verify the following:

- The data source was originally published as an extract.
- Tableau Desktop is connected to the published data source, as indicated by the Tableau Server icon next to the data source name in the Data window:

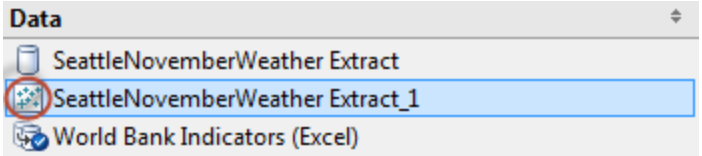

To refresh an extract on Tableau Server or Tableau Online from Tableau Desktop, right-click the data source in the Tableau Desktop Data window, select **Tableau Data Server**, and choose one of the following options:

- **Refresh from Source**

  Refreshes the extract (full or incremental) using the data in the original data source.

  This command is available only for extracts that include a connection to the original data source. If you connected directly to a Tableau Data Extract file (.tde) and then published it, the connection to the original data source is not included.

- **Append Data from Data Source**

  Updates the extract from another data source in the workbook.

- **Append Data from File**

  Updates the extract from the contents of a file if the original data source type of the extract is the same file-based data source or a Tableau data extract (.tde) file.

Note: If you do not see the Tableau Data Server option, your data source may not be on Tableau Server or Tableau Online (in which case it will not show the icon above). If you

see the Tableau Data Server option, but both commands are unavailable, the data source exists on the server, but it is not an extract.

It is also possible to update an extract on Tableau Server using a command-line utility. See **Tableau Data Extract Command-Line Utility** below.

Tableau Data Extract Command-Line Utility

The Tableau Data Extract Utility is installed with Tableau Desktop. You can use this utility at a Windows command prompt to refresh or add files to extracts published to Tableau Server or Tableau Online.

> **Note**: The Tableau Data Extract Utility is available only with Tableau Desktop on Windows.

To run the utility:

1. Open the Command Prompt as an administrator and change to the Tableau Desktop bin directory. For example:

   ```
   cd C:\Program Files\Tableau\Tableau 8.3\bin
   ```

2. Use either of the following commands, adding parameters from the command-options tables below.
   - `tableau refreshextract`
   - `tableau addfiletoextract`

> **Note:** When using the utility, always specify `tableau` on the command line or in scripts, never `tableau.exe`.

Syntax and parameters for the tableau refreshextract command

Use `tableau refreshextract` to refresh an extract on Tableau Server or Tableau Online. Refreshing an extract updates an existing extract with any modifications that have been made to the data source since the last refresh.

To see help for this command, at the Windows command prompt, type the following command:

```
tableau refreshextract --help
```

All options have a full form that you use with a double hyphen (for example, `--server`). Some options also have a short form that you use with a single hyphen (for example, `-s`). If the value for an option contains spaces, enclose it in quotation marks.

tableau refreshextract command options

Short Form	Full Form	Description
	`--source-username <username>`	A valid user name for the data source connection. Use this option with `--source-password`, or use `--original-file` instead of the username and password options. **Note**: You must provide the user name and password when refreshing a published extract, even if the data source was originally published with embedded credentials.
	`--source-password <password>`	The password for the data source user.
	`--original-file <path and filename>`	Path and filename information for the data source that is to be refreshed on the server.
	`--force-full-refresh`	If the data source is set up for incremental refreshes, use this option to force a full extract refresh. If this option is not included, an incremental refresh is performed. Not all data sources support incremental refresh.
`-s <server http address>`	`--server <URL>`	The URL for the Tableau server on which the data is published. To specify Tableau Online, use `http://online.tableausoftware.com`.
`-t <siteid>`	`--site <siteid>`	In a multiple-site environment, specifies the site to which the command applies. For Tableau Online, use this argument if your username is associated with more than one site. For Tableau Server, if you do not specify a site, the default site is assumed. The site id is independent of the site name, and it is indicated in the URL when you view the site in a browser. For example, if the URL for the page you see after signing in to Tableau Online is `https://online.tableausoftware.com/t/vernazza/views` the site id is *vernazza*.
	`--data-`	The name of the data source, as published to Tableau Server

	source <data-source>	or Tableau Online.
	--project <pro-jectname>	The project to which the data source belongs. If this option is not included, the default project is assumed.
-u <user-name>	--username <username>	Valid Tableau Server or Tableau Online user.
-p <pass-word>	--password <password>	The password for the specified Tableau Server or Tableau Online user.
	--proxy-username <username>	The username for a proxy server.
	--proxy-password <password>	The password for a proxy server.
-c <path and file-name>	--config-file <path and file-name>	Path and filename information for a file containing con-figuration options for the command. See **Using a Config File** on page 186 below for details.

Sample tableau refreshextract command

The following command refreshes an extract named CurrentYrOverYrStats that has been published to an on-premiseTableau Server. This command specifies the following:

- The name of your Tableau Server.
- Server user name and password.
- Project name.
- The name of the data source to refresh, along with the data source username and pass-word.

```
C:\Program Files\Tableau\Tableau 8.3\bin>tableau
refreshextract --server https://our_server_name --username
OurServerSignIn --password OurServerPwd --project "New
Animations" --datasource "CurrentYrOverYrStats" --source-
username OurDatabaseSignIn --source-password
OurDatabasePassword
```

The following command refreshes an extract named CurrentYrOverYrStats that has been published to Tableau Online. This command specifies the following:

- Tableau Online user and password.
- Tableau Online site and project names.

- The data source, which in this case is hosted by a cloud-based data source provider (for example, Salesforce.com), and the username and password to sign in to the hosted data source.

```
C:\Program Files\Tableau\Tableau 8.3\bin>tableau
refreshextract --server https://online.tableausoftware.com -
-username email@domain.com --password OurServerPwd --site
vernazza --project "New Animations" --datasource
"CurrentYrOverYrStats" --source-username database_
user@hosted_datasource_provider.com --source-password db_
password
```

Syntax for tableau addfiletoextract

Use `tableau addfiletoextract` to append file content to an extract that has been published to Tableau Server or Tableau Online. This command combines the two files.

If you want simply to update an existing extract with the latest changes, use the `refreshextract` command instead. Using `addfiletoextract` to update an existing extract will duplicate data instead.

To see help for this command, at the Windows command prompt, type the following command:

```
tableau addfiletoextract --help
```

All options have a full form that you use with a double hyphen (for example, `--server`). Some options also have a short form that you use with a single hyphen (for example, `-s`). If the value for an option contains spaces, enclose it in quotation marks.

tableau addfiletoextract command options

Short Form	Full Form	Description
	`--file <path and filename>`	Path and filename information for the data file containing data to append. The file can be from Excel, Access, a Tableau data extract, or a delimited text file. It cannot be password protected. Use UNC format if the file is on a network share. For example, `\\server\path\filename.csv`
`-s <server http address>`	`--server <URL>`	The URL for the Tableau server on which the data is published. To specify Tableau Online, use `http://online.tableausoftware.com`.
`-t <siteid>`	`--site <siteid>`	In a multiple-site environment, specifies the site to which the command applies. For Tableau Online, you must include this argument if your

		username is associated with more than one site. For Tableau Server, if you do not specify a site, the default site is assumed.
	`--datasource <datasource>`	The name of the data source, as published to Tableau Server or Tableau Online.
	`--project <projectname>`	The project to which the data source belongs. If this option is not included, the default project is assumed.
`-u <username>`	`--username <username>`	Valid Tableau Server or Tableau Online user.
`-p <password>`	`--password <password>`	The password for the specified Tableau Server or Tableau Online user.
	`--proxy-username <username>`	The username for a proxy server.
	`--proxy-password <password>`	The password for a proxy server.
`-c <path and filename>`	`--config-file <path and filename>`	Path and filename information for a file containing configuration options for the command. See **Using a Config File** below below for details.

Sample tableau addfiletoextract command

```
C:\Program Files\Tableau\Tableau 8.3\bin>tableau
addfiletoextract --server https://our_server_name --username
OurServerSignIn --password OurServerPwd --project “New
Animations” --datasource "CurrentYrOverYrStats" --file
"C:\Users\user1\Documents\DataUploadFiles\AprMay.csv"
```

```
C:\Program Files\Tableau\Tableau 8.3\bin>tableau
addfiletoextract --server https://online.tableausoftware.com
--username email@domain.com --password OurServerPwd --site
vernazza --project "New Animations" --datasource
"CurrentYrOverYrStats" --file
"C:\Users\user2\Documents\DataUploadFiles\AprMay.csv"
```

Using a Config File

You can use a plain text editor, such as Notepad or Text Edit, to create a config (configuration) file that you can use with either `tableau refreshextract` or `tableau addfiletoextract`. A config file can be useful if you expect to update the same data source regularly over time. Instead of having to type the same options each time you run a command,

you specify the config file. A config file also has the advantage of not exposing user names and passwords on the command line.

Create the Config File

For example, say you created a file called config.txt and saved it to your Documents folder. And in the file, you included the parameter information shown below.

For an extract published to an on-premise Tableau Server:

```
server=https://our_server_name
username=OurServerSignIn
password=OurServerPwd
project=New Animations
datasource=CurrentYrOverYrStats
```

For an extract from a hosted data source, published to Tableau Online:

```
server=https://online.tableausoftware.com
username=email@domain.com
password=OurPassword
project=New Animations
datasource=CurrentYrOverYrStats
source-username=database_user@hosted_datasource_provider.com

source-password=db_password
```

Reference the Config File from the Command Line

After you create the config file, you run the tableau refreshextract or tableau addfiletoextract command, pointing to the config file as the only option you use on the command line.

For example, to refresh the extract specified in the sample in the **Create the Config File** above section, you would run the following command (making sure that you are working in the bin directory for your version of Tableau Desktop):

```
C:\Program Files\Tableau\Tableau 8.3\bin>tableau
refreshextract --config-file
C:\Users\user1\Documents\config.txt
```

Syntax Differences for Config Files

The syntax for specifying options inside a config file differs from the syntax you use on the command line in the following ways:

- Option names do not begin with dashes or hyphens.
- You use an equals sign (with no spaces) to separate option names from option values.
- Quotation marks are not necessary (or allowed) around values, even when they include spaces (as for the `project` option in the example shown earlier).

Use Windows Task Scheduler to Refresh Extracts

You can use Windows Task Scheduler, in combination with the Tableau Data Extract Command-Line Utility, to automate regular updates to Tableau Desktop data sources from within your corporate firewall. You can configure a task to occur once per day, week, or month, or after a specific system event. For example, run the task when the computer starts.

To learn more, see the Task Scheduler How To... page in the Microsoft TechNet library.

Tableau Data Extract API

Use the Tableau Data Extract API to connect to data that is not currently a supported Tableau data source. With the Tableau Data Extract API, you create a program that accesses and processes your data. You then use that program to create a Tableau Data Extract (TDE) file.

The Data Extract API is available for developers on Windows and Linux platforms. Go to the Get the Data Extract API page on the Tableau website, and choose the appropriate version for your platform and programming language:

- Data Extract API – Python – 32-bit
- Data Extract API – Python – 64-bit
- Data Extract API - C/C++/Java – 32-bit
- Data Extract API - C/C++/Java - 64bit

> **Note**: A single-threaded application is recommended because the Java thread is not safe.

Notes for Developers

The Data Extract API includes a sample program, `makeorder`, coded in each supported language to demonstrate a typical usage scenario: creating an extract containing product orders. The application creates the extract order.tde with several columns of different types. The general flow of the sample programs is:

1. Open an Extract object to create a new file.
2. Define the extract's schema using a TableDefinition.
3. Add the Extract table.
4. Insert rows.
5. Close all objects.

It is important to free memory by closing all objects, and it is particularly critical to ensure Extract objects are cleaned up properly, particularly in non-native execution environments. See the note sections below for language-specific details.

String columns in a Data Extract can be 8- or 16-bit and can be sorted according to many available collations. By default, strings are sorted according to their binary representation, though this can be changed on a per-table or per-column basis.

Python Notes

Objects in the Data Extract API are automatically closed by `_del_` when necessary. While garbage collection handles the vast majority of concerns related to releasing resources, it is important to note that the virtual machine provides no guarantee that any particular object will ever be freed. While most objects are merely memory, Extract objects represent physical files created when `close` is invoked. Therefore, it is not safe to rely on garbage collection to close Extract objects. We recommend using `with` statements to ensure Extract instances are cleaned up. Alternatively, you can explicitly call `close`.

Java Notes

Data Extract API objects are automatically closed by `finalize()` as necessary. The Java Virtual Machine does not guarantee that any particular object is ever garbage collected. While most objects are merely memory that can be safely reclaimed by the operating system at JVM shutdown, Extract objects represent physical files that are created when `close()` is invoked. Therefore, it is important to invoke `Extract.close()` for all Extract instances. We recommend using the try-with-resources construct introduced in Java 7. For earlier versions of Java, you must call *close()* explicitly.

C++ Notes

Data Extract API objects should be managed according to standard memory management best practices, such as using stack variables or smart pointers. As in other languages, all objects have a `Close()` method to free internal resources. `Close()` is invoked by the destructor when necessary. However, it is important to note that `Extract::Close()` may throw an exception, so it is safer to call it explicitly, rather than allowing an exception to potentially escape the destructor.

C Notes

Objects in the Data Extract C API are managed through opaque `TAB_HANDLE`s. Every created object must be closed. It is advisable to free objects in the reverse order of creation.

Using the Data Extract API in Microsoft Visual Studio 2010

Follow these steps to build a C or C++ project in Visual Studio 2010:

1. Extract the non-Python package to an installation directory—for example, `C:\dataextract-8.0`. This directory is the `$(InstallRoot)`.
2. Add `$(InstallRoot)\bin` to the %PATH% environment variable.
3. Open Visual Studio and create a new console project.
4. Open the Property Manager and create a new Property Sheet for use with all targets in the project. Add these paths:

Path	Description
`$(InstallRoot)\include`	C/C++|General|Additional Include Dir-

	ectories
`$(InstallRoot)\lib`	VC++ Directories|Library Directories
`dataextract.lib`	Linker|Input|Additional Dependencies

5. Add `MakeOrder.c` or `MakeOrder.cpp` from `$(InstallRoot)\docs\samples`.
6. Compile and run the application.

Using the Data Extract API in Eclipse

Follow these steps to build a Java project in Eclipse:

1. Extract the non-Python package to an installation directory—for example, `C:\dataextract-8.0`. This directory is the `$(InstallRoot)`. Verify that the package is compatible with your Java platform; 64-bit libraries do not work with a 32-bit JVM and vice-versa.
2. Open Eclipse and create a new project. Add `dataextract.jar` and `jna.jar` to the build path as External JARs.
3. Right-click on the project's src package and import `$(InstallRoot)\docs\samples` as a local file system resource. Verify that `samples/com/tableausoftware/demos/MakeOrder.java` is selected before dismissing the dialog.
4. In the Run Configuration for MakeOrder, add an Environment Variable `PATH=$(InstallRoot)\bin`.

Using the Data Extract API with Python

Follow these steps:

1. Verify that your installed Python is version 2.x, where x is 7 or higher. Version 3 is not supported. Also verify that it matches the package you have downloaded: the 64-bit module is incompatible with 32-bit Python and vice-versa.
2. Extract the Python package, open a command prompt, and navigate to the directory that contains setup.py.
3. Run setup.py to install the module into site-packages.

Managing Queries

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.

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 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 Profit 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.

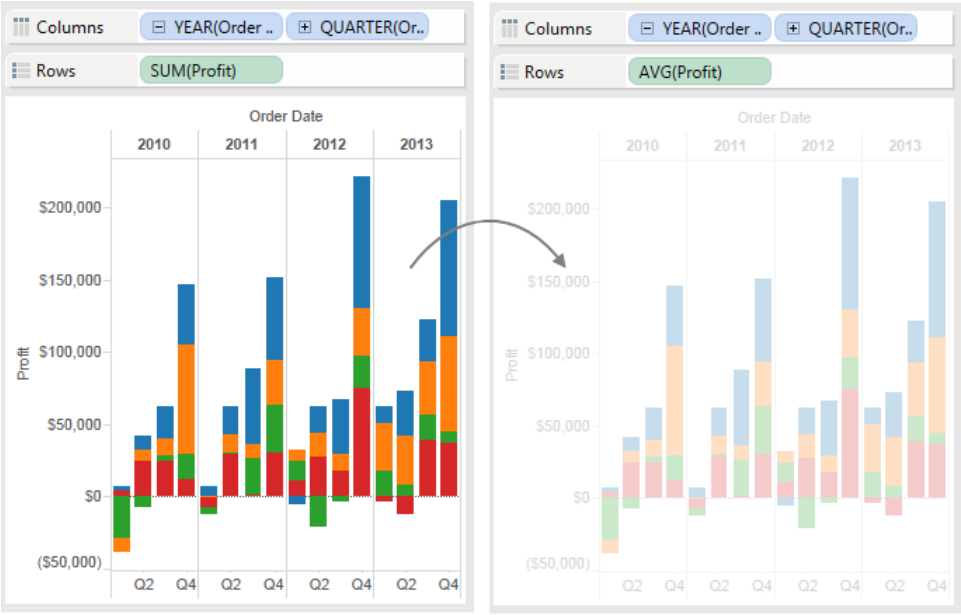

How Automatic Updates Affect Dashboards and Stories

You configure automatic updates on a per-view basis. This means that you can have a dashboard where some views are updating automatically, and others are not. Similarly, you can have a story where some story points are updating automatically, and others are not. But

when dashboards or stories are published to Tableau Server, automatic updates affect either none of the contributing views in a story or dashboard, or all of the views.

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.

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

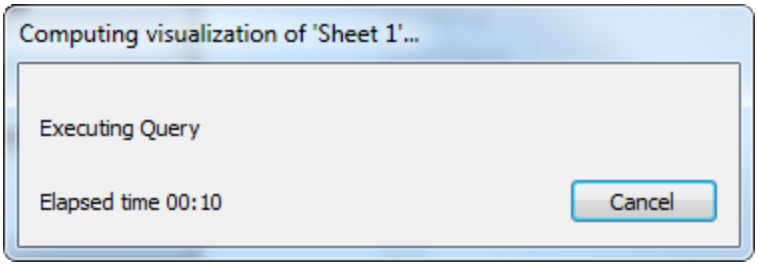

After canceling a query the view becomes invalid because it is in an in-between state. The result is a blank view although all your fields are still on the shelves. To resume working with Tableau, alter the view in anyway and allow the query to complete.

> 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.

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, Microsoft Analysis Services 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.

> Text, Microsoft Excel, and Microsoft Access data sources may be temporarily unavailable after canceling 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

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.

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. In Tableau, multidimensional data sources are supported only in Windows. 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	K	L	
1	Row ID	Product Sub-Category	Product Container	Product Name
2	18606	Labels	Small Box	Avery 49
3	20847	Pens & Art Supplies	Wrap Bag	SANFORD Liquid Accent
4	23086	Paper	Small Box	Xerox 1968
5	23087	Scissors, Rulers and Trimmers	Small Pack	Acme® Preferred Stainle
6	23088	Telephones and Communication	Small Box	V70
7	23597	Paper	Small Box	Xerox 194
8	25549	Office Machines	Jumbo Drum	Canon S750 Color Inkjet
9	20228	Chairs & Chairmats	Jumbo Drum	Global Troy™ Executive
10	19483	Paper	Small Box	Xerox 1930
11	24782	Appliances	Small Box	Kensington 6 Outlet Mas
12	24563	Paper	Small Box	Xerox 21
13	24564	Pens & Art Supplies	Wrap Bag	*Staples* Highlighting M
14	24565	Telephones and Communication	Wrap Bag	Accessory34
15	21866	Paper	Small Box	Xerox 1933
16	20876	Bookcases	Jumbo Box	Sauder Forest Hills Libra
17	20877	Tables	Jumbo Box	Riverside Furniture Stan
18	22241	Envelopes	Small Box	Park Ridge™ Embossed

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.

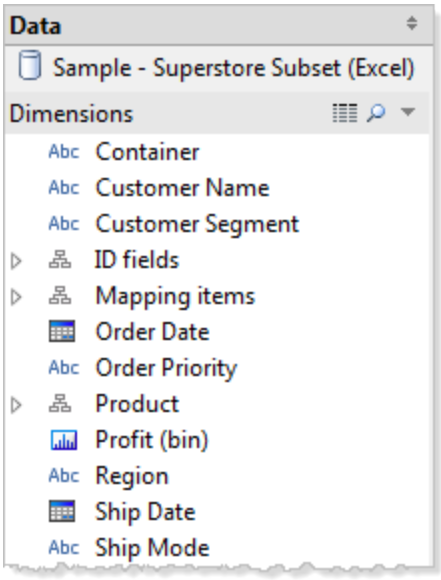

The Data window organizes fields into three areas:

- **Dimensions** – Fields that typically hold discrete qualitative data. Examples of dimensions include dates, customer names, and customer segments.
- **Measures** – Fields that typically hold numerical data that can be aggregated. Examples of measures include sales, profit, number of employees, temperature, frequency, and

pressure.

- **Sets** – An additional area that stores custom fields based on existing dimensions and criteria that you specify. Named sets from an MS Analysis Services server or from a Teradata OLAP connector 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 other custom sets in Tableau.
- **Parameters** – An additional area that stores parameters that you have created. Parameters are dynamic variables that can be used as placeholders in formulas.

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.

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.

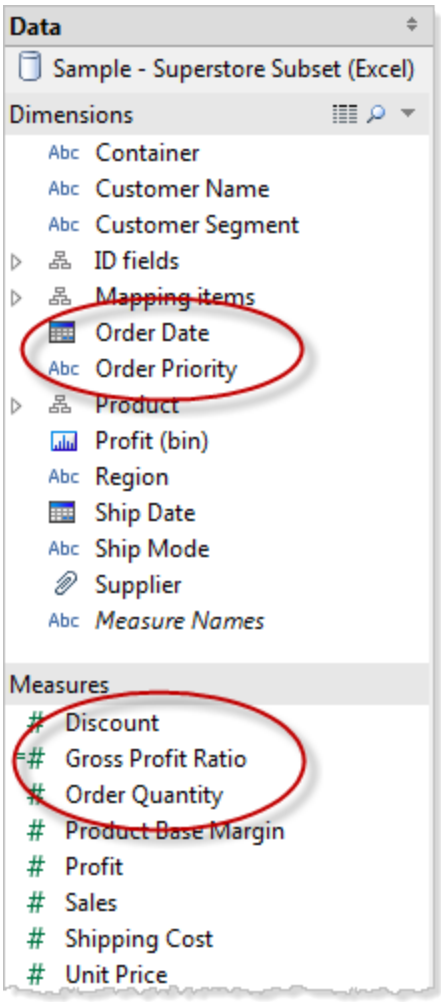

By default the field names defined in the data source are displayed in the Data window. You can rename fields as well as member names.

Relational and Multidimensional Data

In Tableau, multidimensional data sources are supported only in Windows.

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

Relational

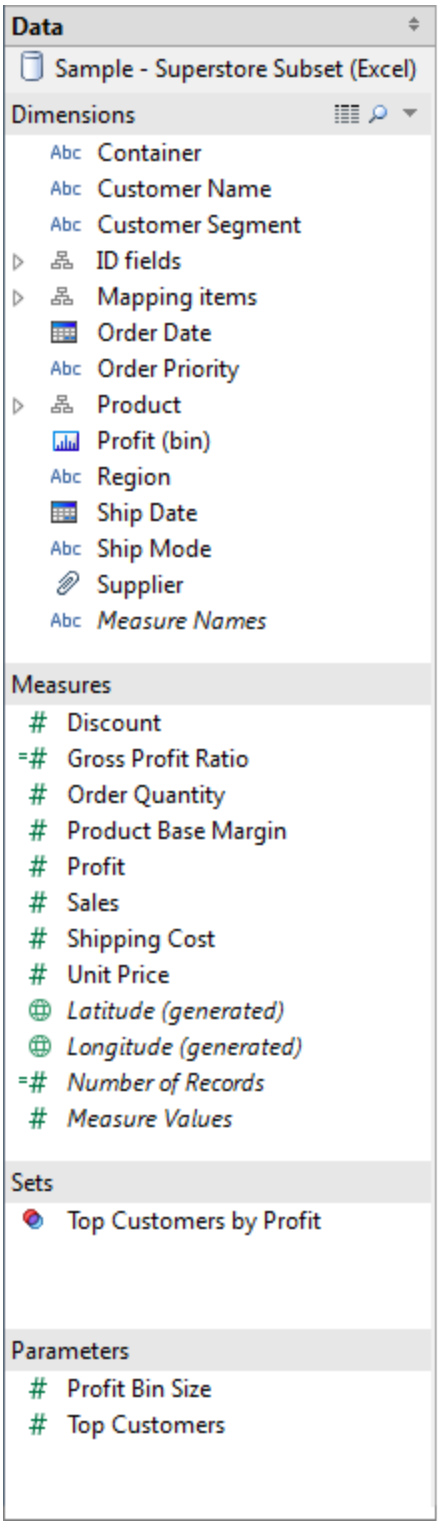

Multidimensional

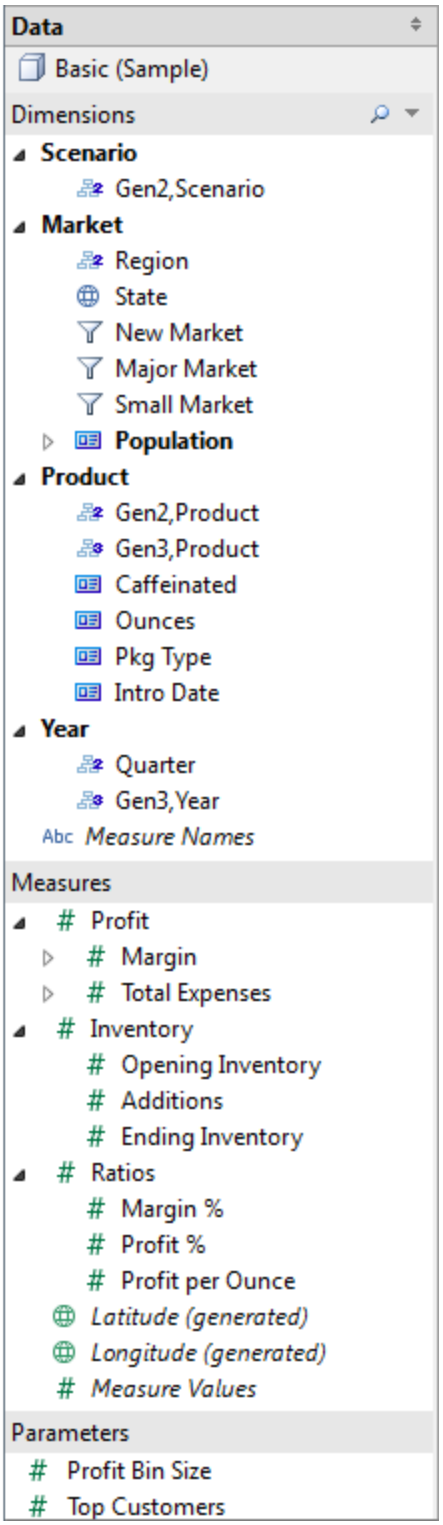

You can expand or collapse the various hierarchies in both relational and multidimensional Data windows by clicking the plus button. You can hide the Data window all together by clicking the minimize button in the upper right corner of the Data window.

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 relational hierarchies by dragging and dropping in the Data window.

1. Drag a field in the Data window and drop it directly on top of another field.

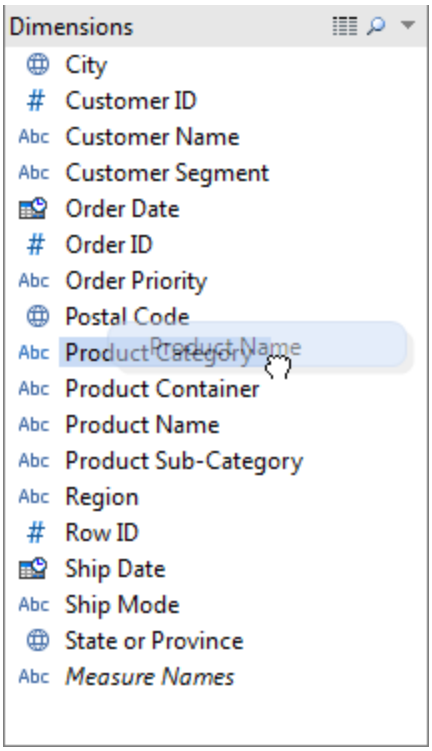

> When you want to create a hierarchy from a field inside a folder, right-click the field and then select **Create Hierarchy** instead.

2. When prompted, type a name for the relational hierarchy and click **OK**.

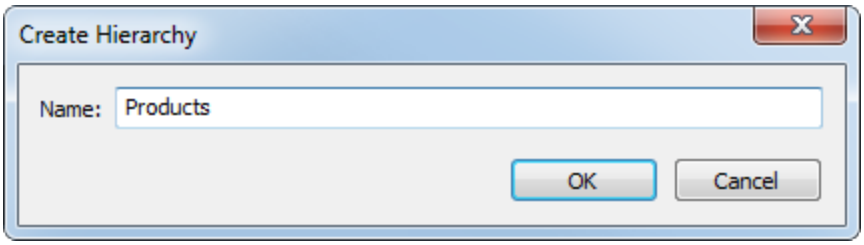

3. Drag any additional fields into the hierarchy. You can also re-order fields in the hierarchy by dragging them to a new position.

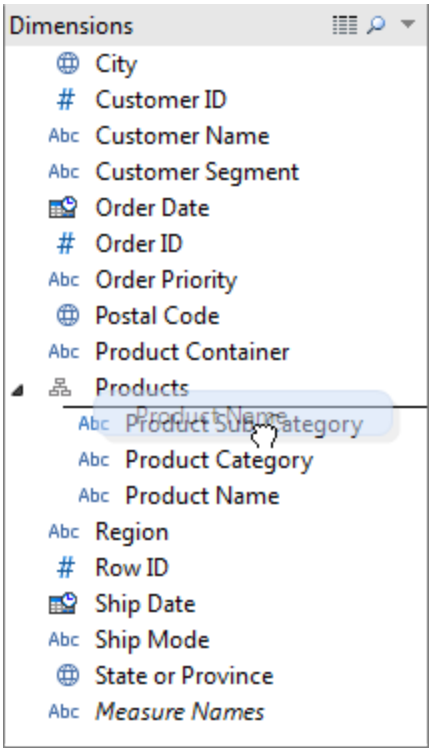

Hierarchies support single click navigation up and down the levels. When you use the fields in the view, a plus button displays on the field so you can drill down and up in the hierarchy.

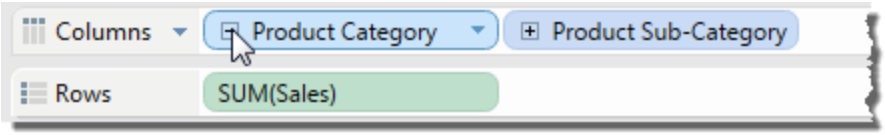

Measure Values and Measure Names

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 involved 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.

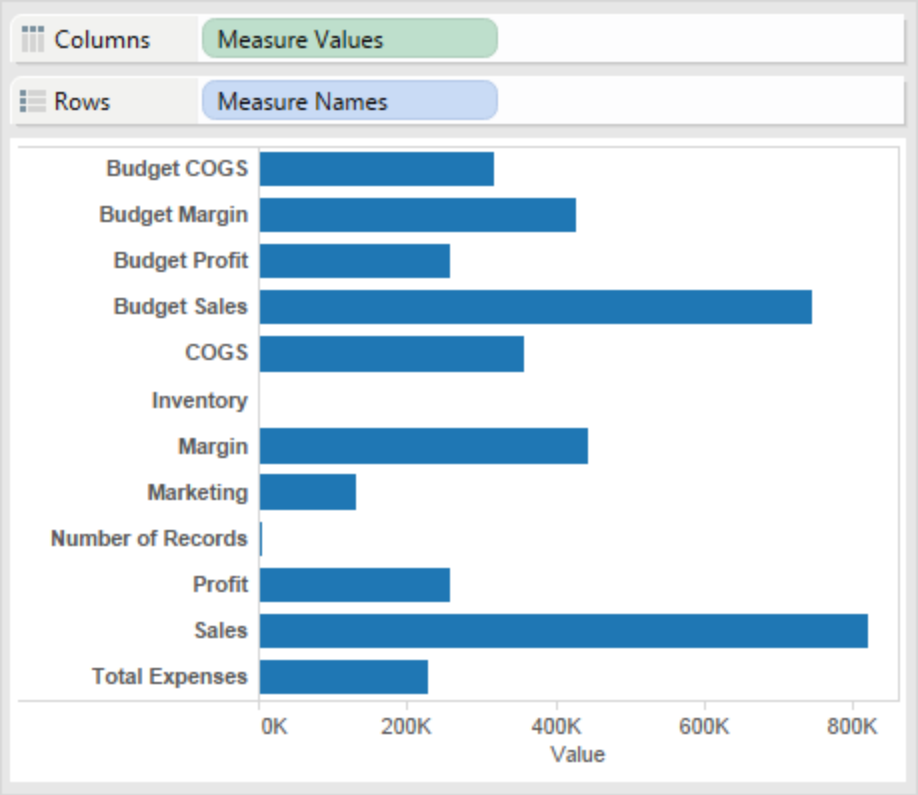

Number of Records

In addition to 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, 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.

Data Window Features and Functions

The Data window has many features and functions to help you organize your data fields, find specific fields, and hide others.

Organize the Data Window

You can reorganize the items in the Data window from its default layout using folders or through sorting.

Folders

To make data sources with many fields easier to work with, you can organize the Data window items into folders. Items like fields, parameters, and sets can be grouped into folders. The **Group by Folder** option can be accessed from the Data window menu. Alternatively, you can also access this option from the field context menu.

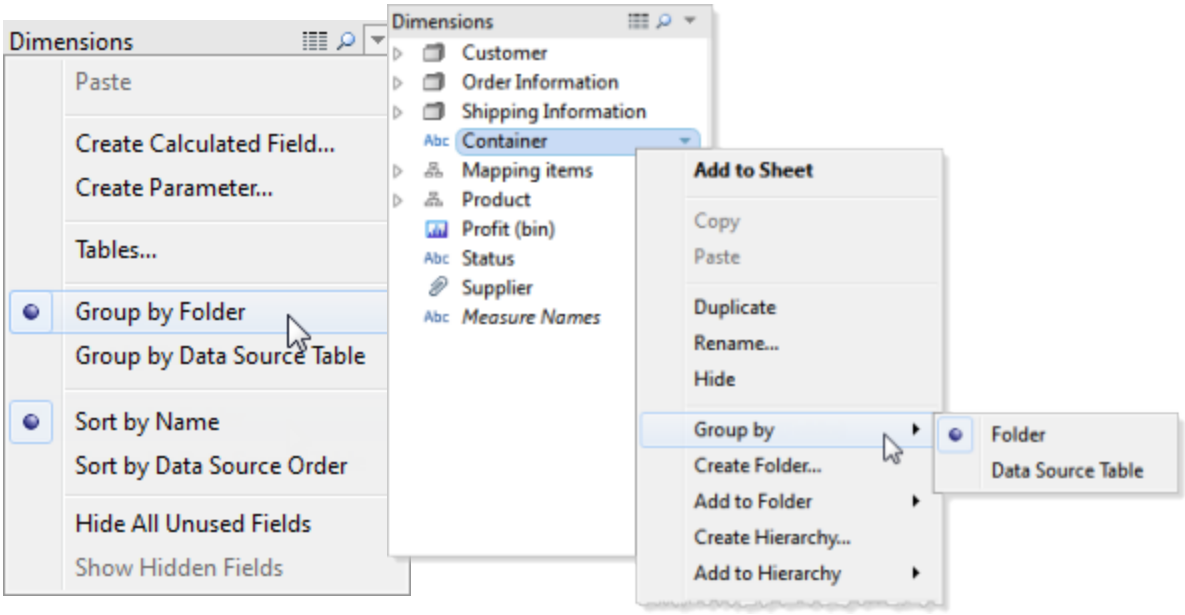

> When you connect to a single table in your data source, grouping by folder is enabled by default. When you connect to a data source with multiple tables, grouping by table is enabled. When the **Group by Data Source Table** option is selected, the dimensions and measures are grouped according to the database table they belong to. This is especially useful when you have several joined tables.

To group fields into folders:

1. In the Data window, select the fields you want to group together or right-click an empty area in the Data window.
2. Select **Create Folder**.
3. When prompted, name the new folder.

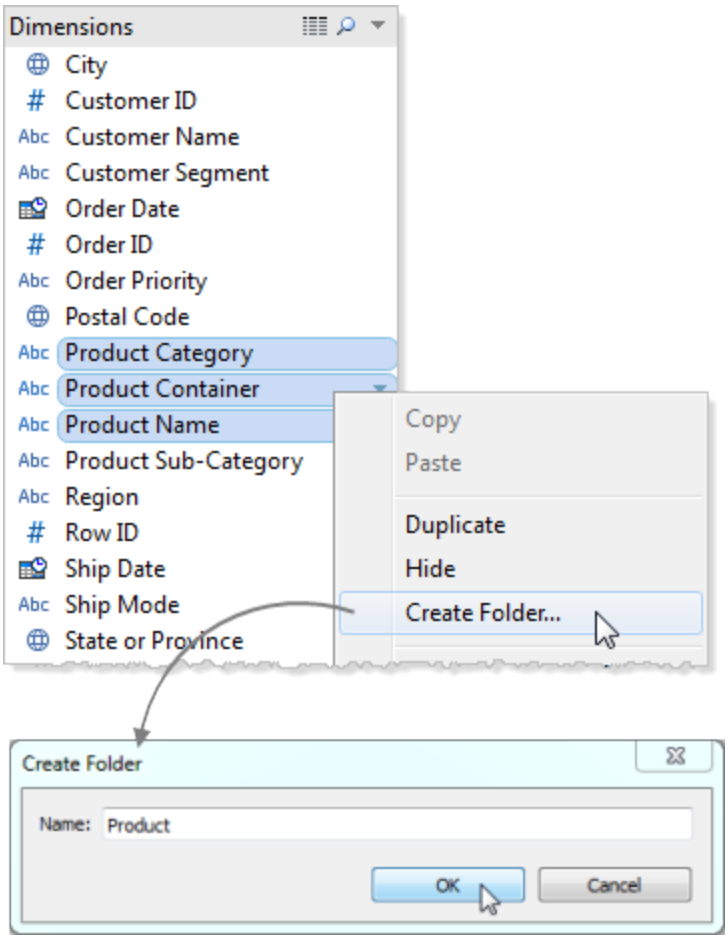

After you create a folder structure, you can drag fields from one folder to another or duplicate a field you want to have available in more than one folder.

To add a field to a folder:

- Drag a field on top of the folder name to add the field to the folder. If the folder is expanded, you can drag the field into the general area of the folder.

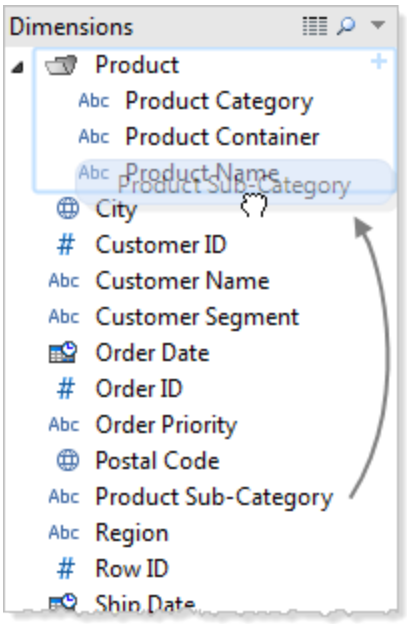

Sort

When organizing the Data window with or without folders, you can have Tableau sort the items. These **Sort by** options are also located in the Data window menu.

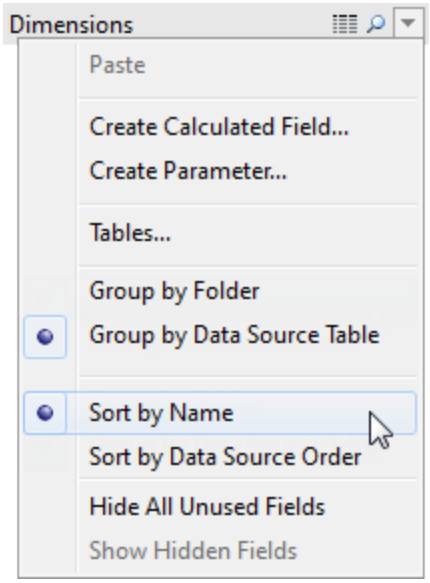

You can sort by one of the following options:

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

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.

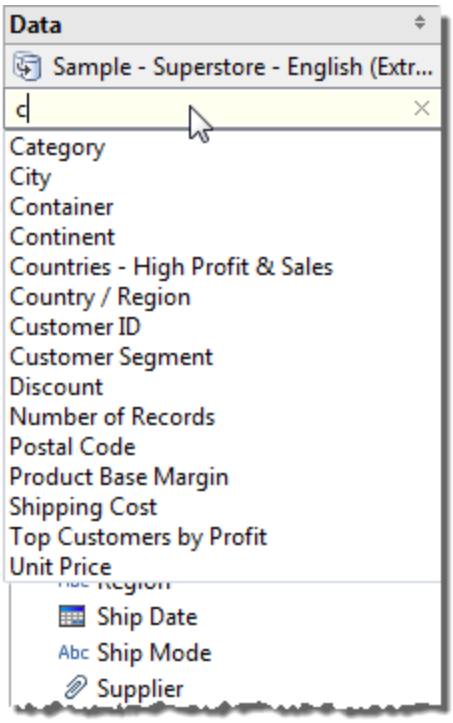

Rename Fields

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.

Renaming 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.

Reverting to the Default Field Name

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

Combine Fields

Combine fields to create a cross product of members from different dimensions. You would combine dimensions if you want to encode a data view using multiple dimensions.

To combine the fields, select multiple dimensions in the Data window and then right-click the fields and select **Combine Fields**.

> For multidimensional data sources, you must select levels from different hierarchies. In Tableau, multidimensional data sources are supported only in Windows.

For example, the selections shown below will produce a new field that consists of the Container and Customer Segment dimensions.

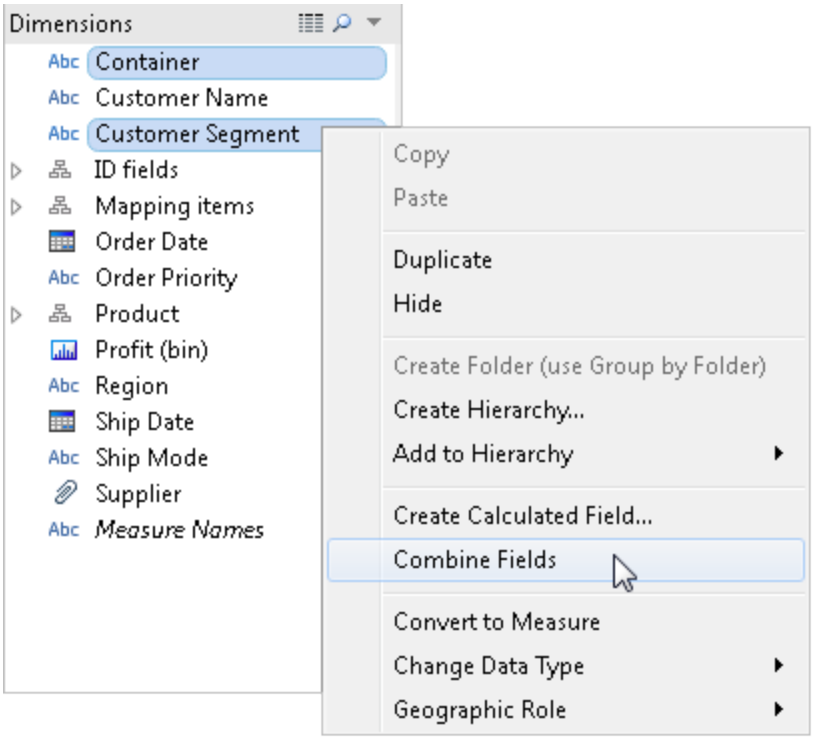

The two dimensions are combined into a new dimension. The name of the field is automatically created from the names of the original fields. Right-click the new field and select **Rename** to change the name.

When you use the new field in a view, a header is created for each combination of the two original dimensions. For example, the view below shows the members of the combined Customer Segment and Container fields.

Columns Region

Rows Segment and Container

Segment and Container	Region Central	 East
Consumer, Jumbo Box	$70,604	$90,512
Consumer, Jumbo Drum	$84,796	$138,972
Consumer, Large Box	$70,790	$70,746
Consumer, Medium Box	$12,782	$10,720
Consumer, Small Box	$137,953	$144,546
Consumer, Small Pack	$20,640	$14,071
Consumer, Wrap Bag	$6,727	$11,239
Corporate, Jumbo Box	$173,900	$143,674
Corporate, Jumbo Drum	$264,531	$173,960
Corporate, Large Box	$130,634	$59,860
Corporate, Medium Box	$82,805	$19,766
Corporate, Small Box	$331,450	$375,742
Corporate, Small Pack	$39,167	$30,447
Corporate, Wrap Bag	$19,996	$14,652
Home Office, Jumbo Box	$91,680	$88,040
Home Office, Jumbo Drum	$144,728	$113,118
Home Office, Large Box	$96,027	$88,739

For multidimensional data sources, to choose to display the fully qualified name, right-click the combined field in the Data window and select **Qualify Member Names**.

Hide or Unhide Fields

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

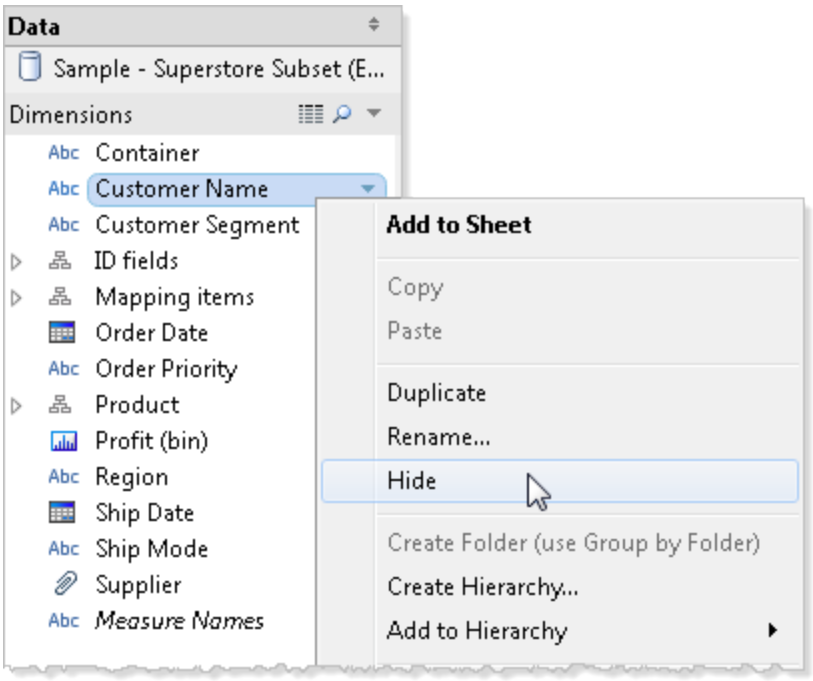

When you want to change your fields from hidden to visible, select **Show Hidden Fields** on the Data window menu.

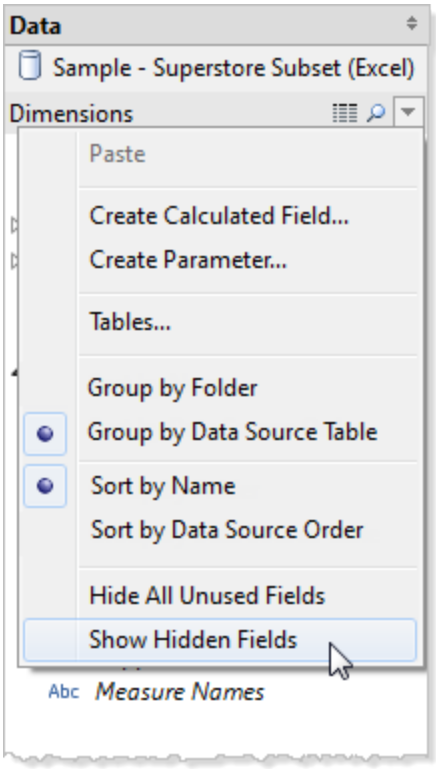

The hidden fields are shown in gray in the Data window. You can then select one or more hidden fields, right-click and select **Unhide**.

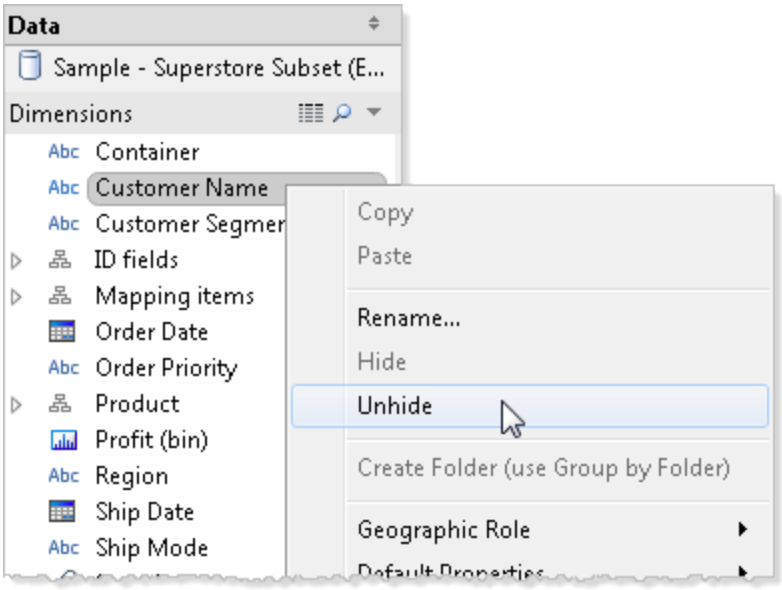

Select **Hide All Unused Fields** on the Data window menu to quickly hide all of the fields that are not being used in the workbook.

> 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**.

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, shapes, default aggregation, and so on.

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.

Adding a Comment to a Field

1. Right-click a field in the Data window and select **Default Properties** > **Comment**.
2. Write a comment in the subsequent dialog box. Comments support rich text formatting that will be represented in the tooltip.

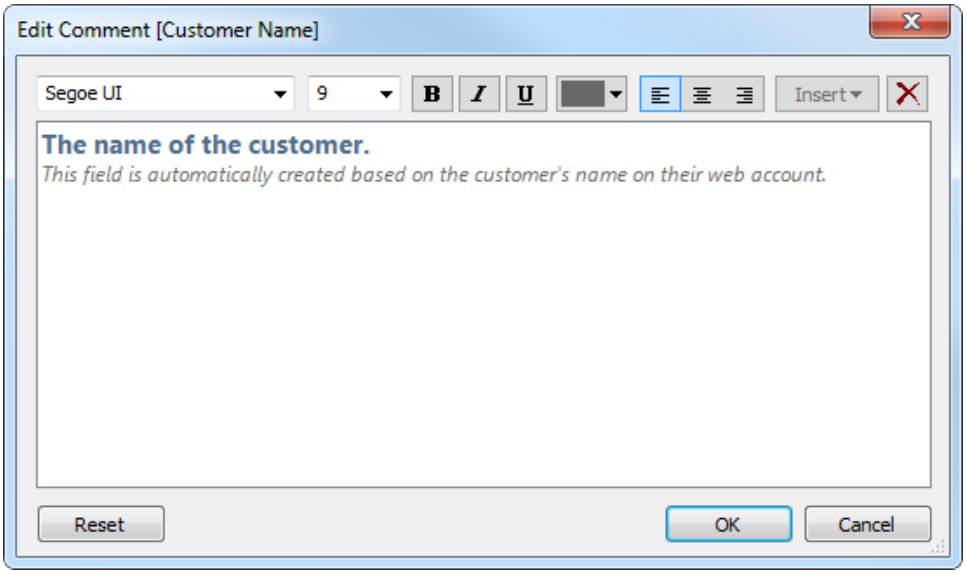

3. When finished, click **OK**.

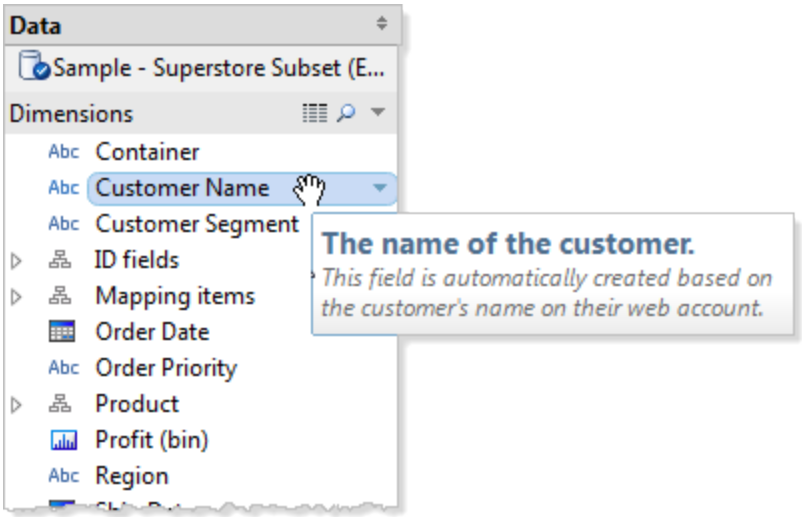

Aliases

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

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

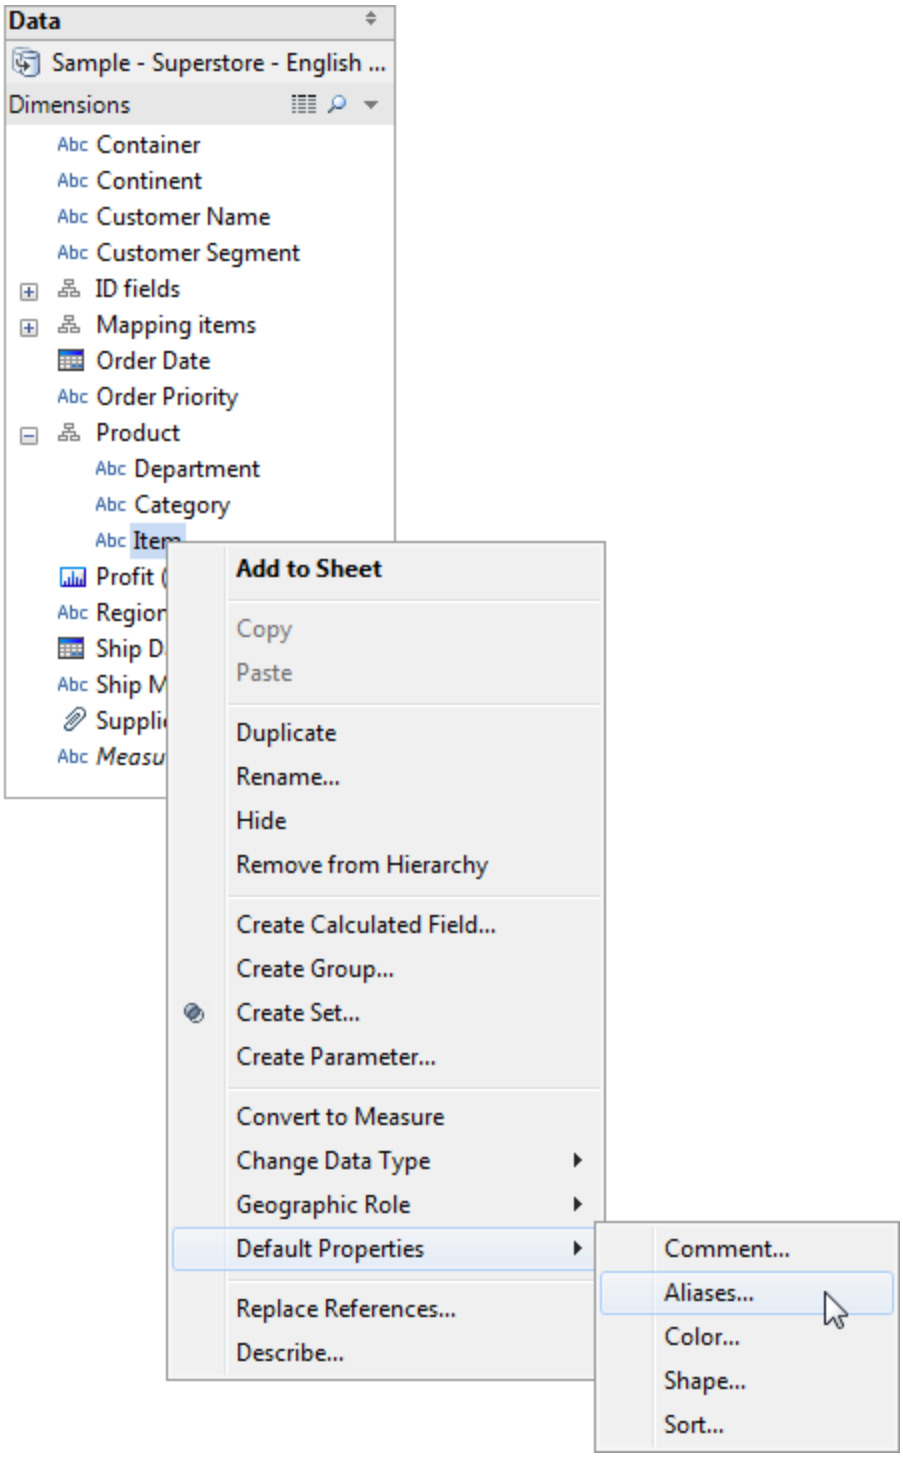

A dialog box opens allowing you to define aliases for each value within the selected dimension. You can reset the member names back to their original names by clicking the **Clear Aliases** button in the bottom right corner of the Edit Aliases dialog box.

Aliases with a Multidimensional Data Source

> In Tableau, multidimensional data sources are supported only in Windows.

Aliases for multidimensional 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 database administrator to find out whether your database has aliases available. Aliases are not support 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 member names are displayed (example on the left). As you can see, these names are not very intuitive. By selecting **Data** > **Aliases File** and selecting an appropriate alias file set up by the database administrator, meaningful names are displayed in the headers.

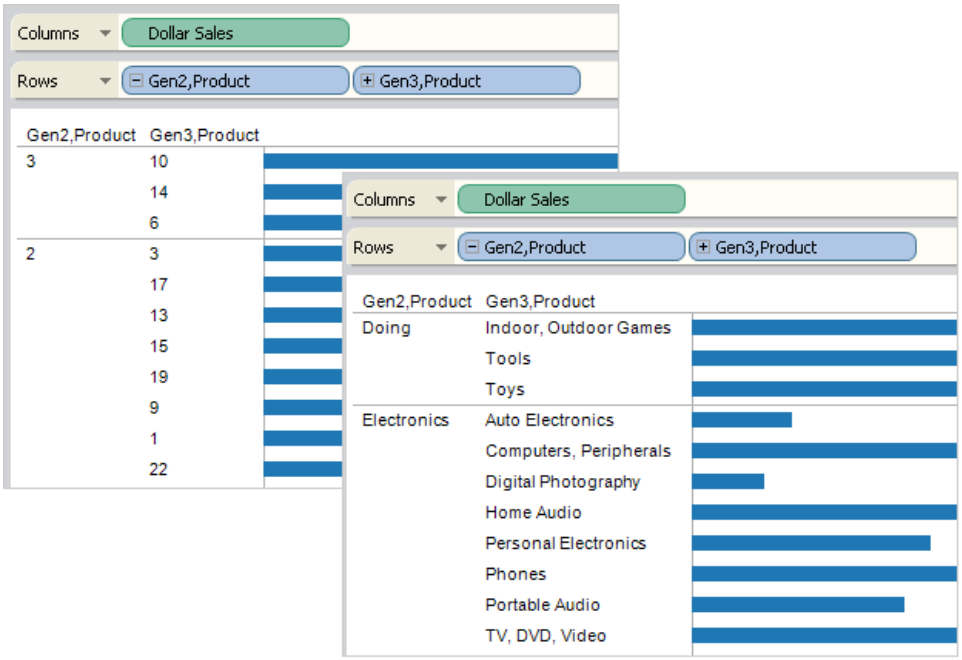

Example – Editing Aliases

The Sample - Superstore 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.

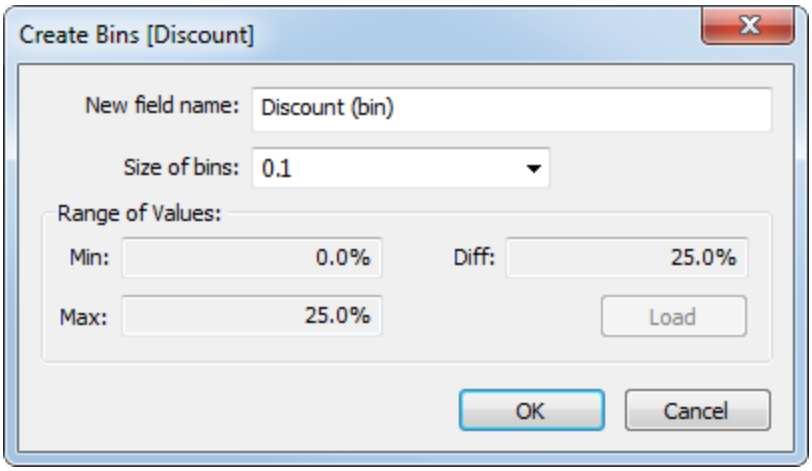

The new binned field is named **Discount (bin)** and appears in the Dimensions area of the Data window.

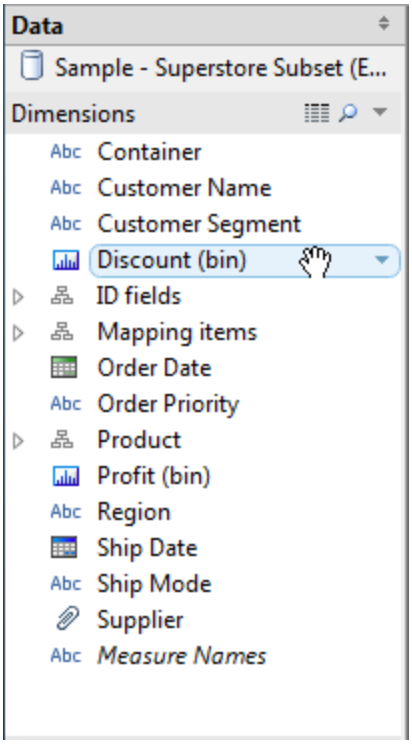

When you place **Discount (bin)** on the Rows or Columns shelf, the default aliases for the bins are given by the lower limit of the bin's numerical range

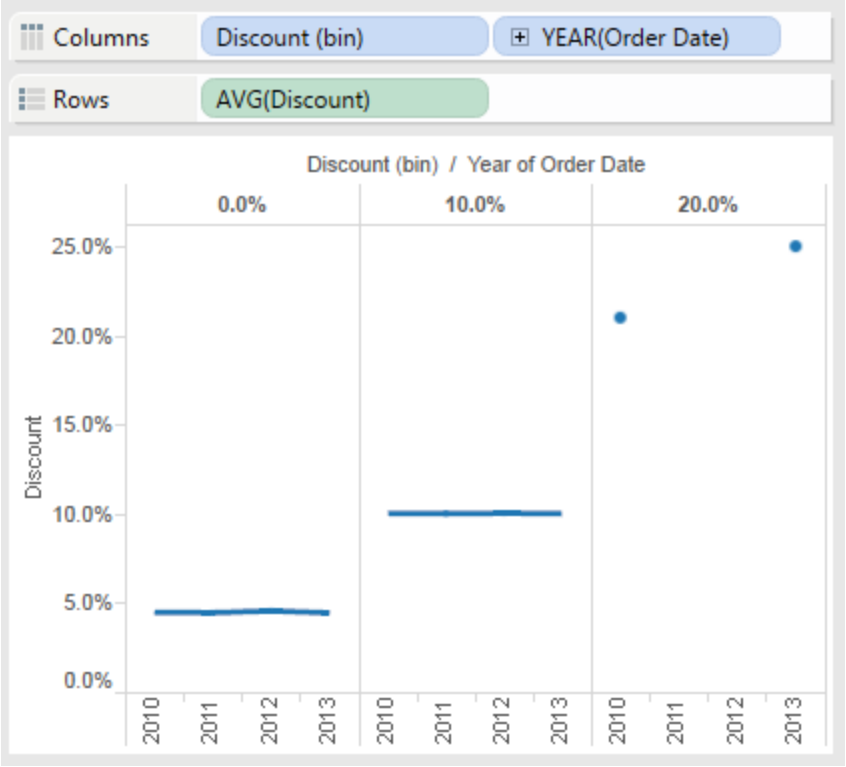

To improve the readability of the bins when they are displayed in Tableau, you can define aliases such as “Low discount”, “Medium discount”, and “High discount.”

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

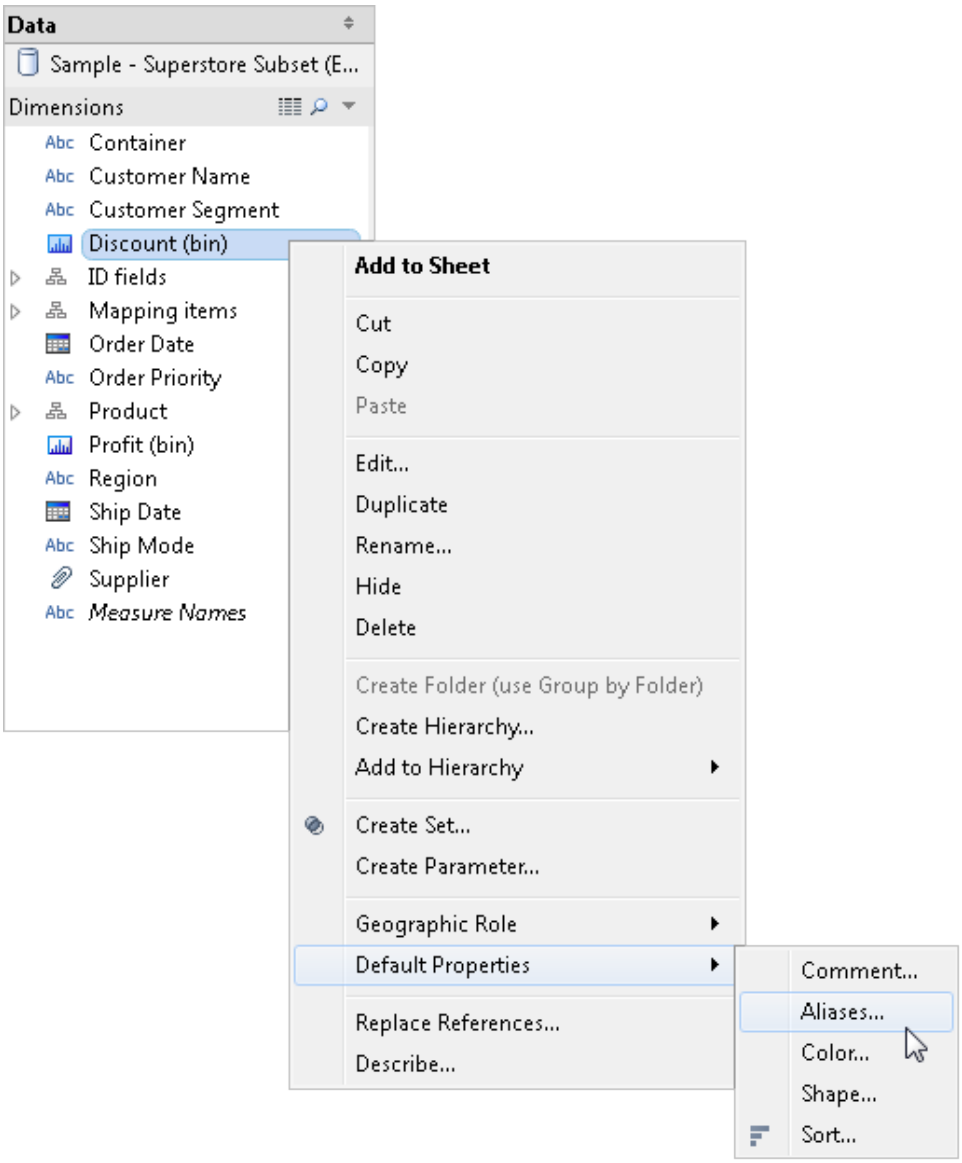

2. Assigning an alias to every member of the **Discount (bin)** fields. For example, the member originally labeled as 0.0% is now labeled as “Low Discount.”

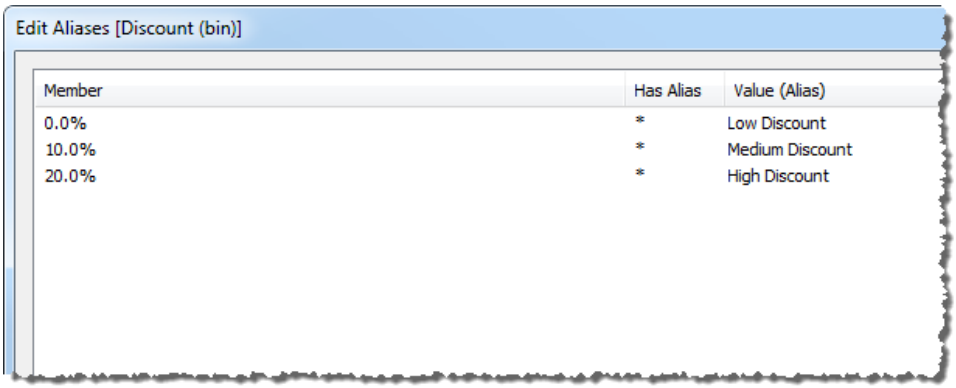

You can change aliases at any time using the **Edit Aliases** dialog box. To do so, click on the alias you want to change and specify the new name. Use the Tab key to advance from one value to the next. To restore the original aliases, click the **Clear Aliases** button in the bottom right corner of 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. To set the default color encodings for a field, right-click the field in the Data window and select **Default Properties** > **Color**.

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 that Furniture products are represented with a square mark, it will automatically be changed to a square mark in all other views in the workbook. To set the default shape encodings for a field, right-click the field in the Data window and select **Default Properties** > **Shape**.

Formats

You can set the default number 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 places. On the other hand, you may want to always show Discount as a percentage. To set the default formats, right-click a date or number field and select either **Date Format** or **Number Format** on the Default Properties menu. A dialog box opens where you can specify a default format.

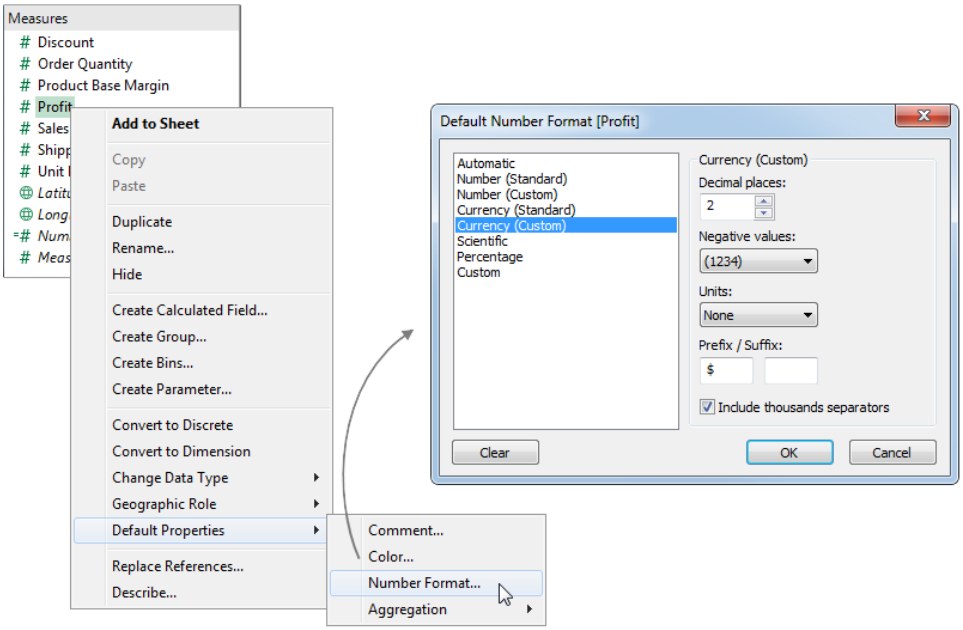

Sort

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, the 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 **Default Properties** > **Sort**. Then use the sort dialog box to specify a sort order.

> The default sort order also controls how the 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.

1. Right-click any measure in the Data window and select **Default Properties > Aggregation**.
2. On the Aggregation list, select an aggregation.

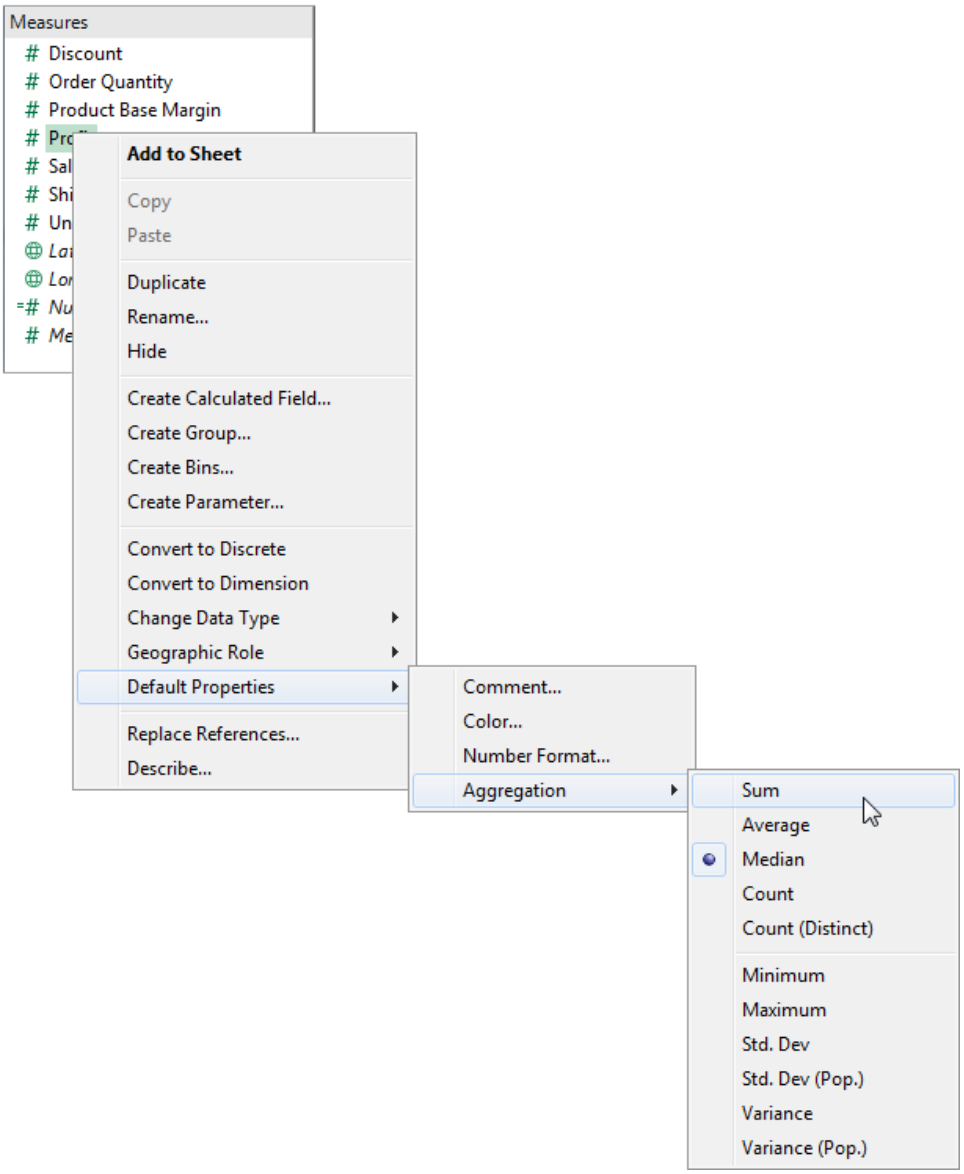

Whether you are specifying the aggregation for a field on a shelf or the default aggregation in the Data window, you can select from several aggregations. See **Aggregations** on page 563 to learn about each type of aggregation.

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 department by region.

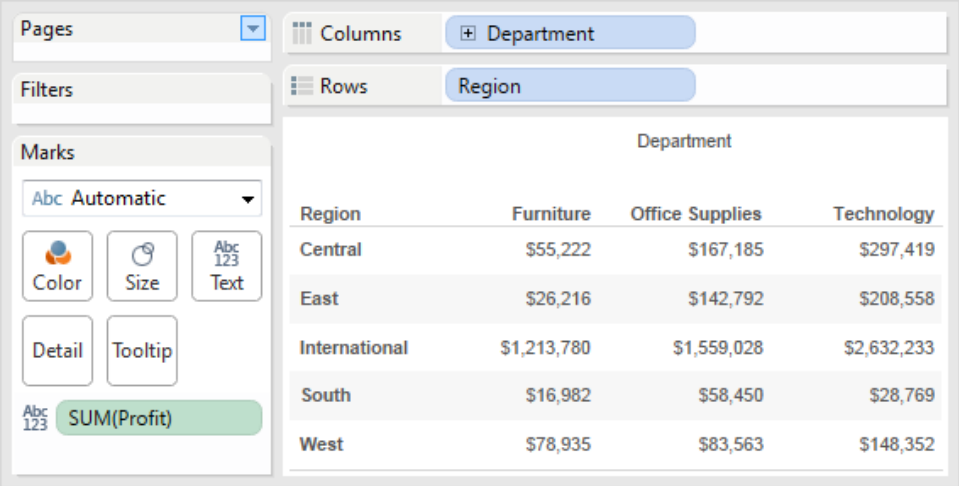

Now suppose you want to show both the Profit and the Sales for each department and region. When you add the Sales measure to the text table, the measures are combined and the Measure Values field is placed on Text. Additionally, the Measure Names field is added to the Rows shelf.

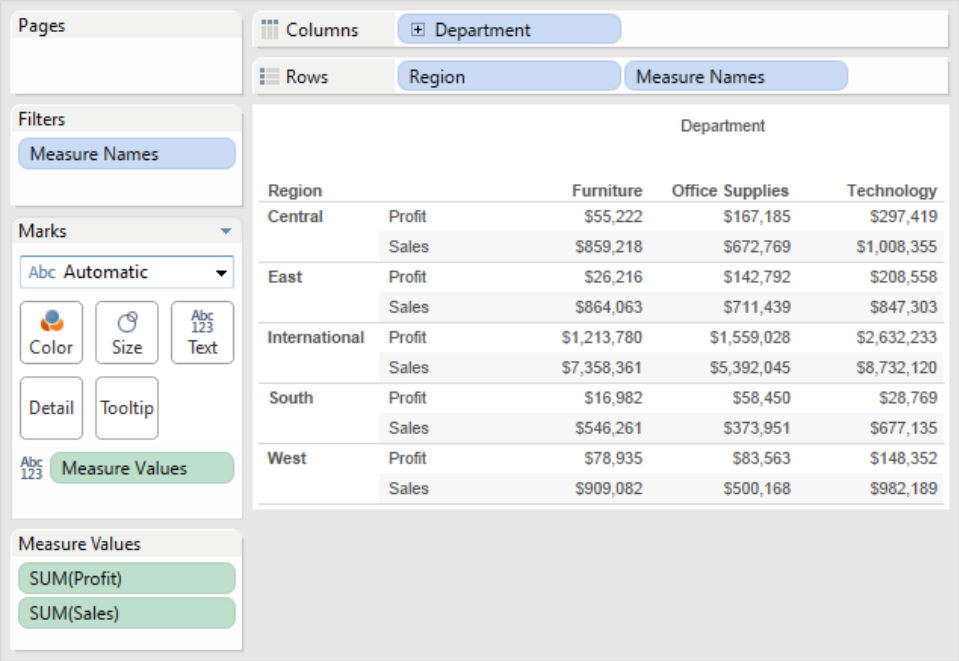

Notice how the header names removed the aggregation label by default. You may want to include the aggregations or call it "Total Profit" and "Total Sales." To change the measure names, right-click the Measure Names field on the Rows shelf and select **Edit Aliases**. Make the changes and click **OK**.

		Department		
Region		Furniture	Office Supplies	Technology
Central	Total Profit	$55,222	$167,185	$297,419
	Total Sales	$859,218	$672,769	$1,008,355
East	Total Profit	$26,216	$142,792	$208,558
	Total Sales	$864,063	$711,439	$847,303
International	Total Profit	$1,213,780	$1,559,028	$2,632,233
	Total Sales	$7,358,361	$5,392,045	$8,732,120
South	Total Profit	$16,982	$58,450	$28,769
	Total Sales	$546,261	$373,951	$677,135
West	Total Profit	$78,935	$83,563	$148,352
	Total Sales	$909,082	$500,168	$982,189

Data Types and 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.

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
Abc	Text values
	Date values
	Date & Time values
#	Numerical values
T|F	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. To change the data type in Tableau, right-click the field in the Data window, select **Change Data Type**, and then select the appropriate data type.

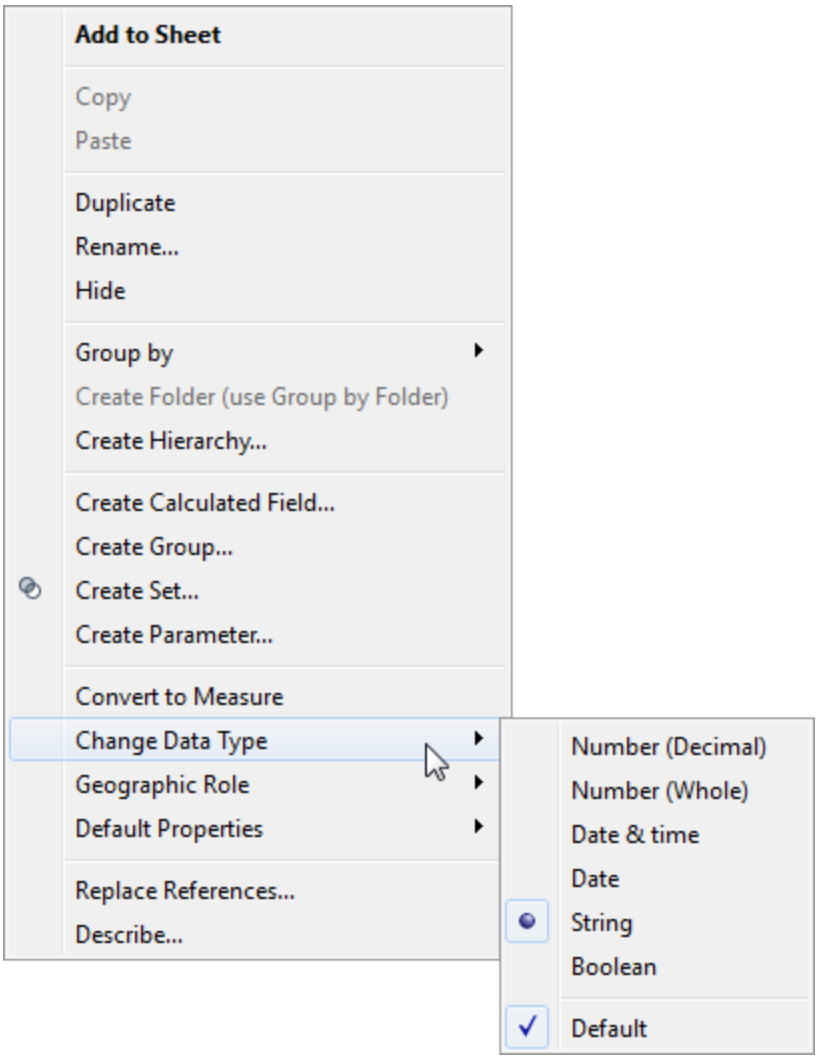

Sometimes the data in your database is more precise then Tableau can model. When you add these values to the view a precision warning will appear in the right corner of the status bar. See **Precision Warnings** on page 193.

Mixed Data Types for Excel, Access, and CSV Files

Most columns in an Excel, Access, or CSV (comma separated value) file contain values of the same data type (booleans, dates, numbers, and 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 10,000 rows of an Excel data source, and the first 1,024 rows of a

text file data source (first 16 rows if you're using a workbook that was created before Tableau Desktop 8.2 or the legacy connection). For example, if most of the first 10,000 rows are text values, 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 that Tableau determines for each field, the field might contain Null values for the other (non matching) records as described in the table below.

Mapped data type	Treatment of other data types in the field.
Text	Dates and numbers are treated as text. Nulls are not created.
Dates	Text is treated as Null. A number is treated as the day in numeric order from 1/1/1900.
Numbers	Text is treated as Null. A date is treated as the number of days since 1/1/1900.
Boolean	Text, dates, and numbers are treated as Null.

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 that does not contain the 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 a measure, and whether it is continuous or discrete.

> 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.

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

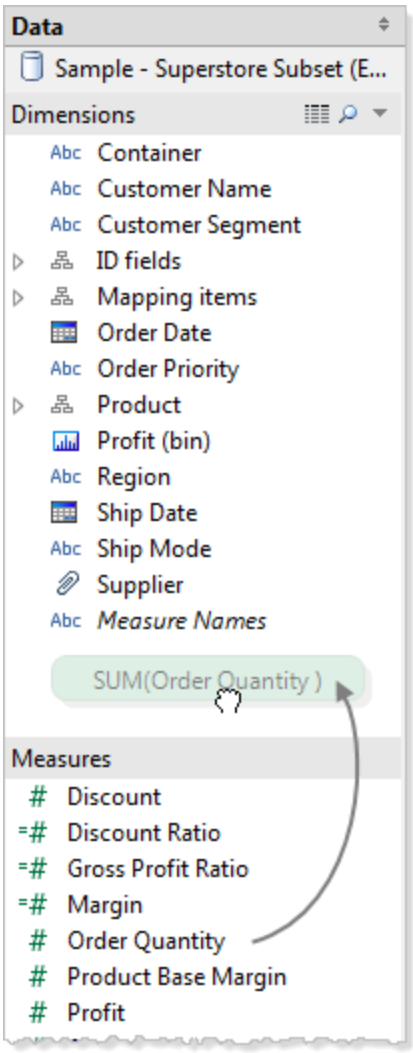

- Right-click the measure in the Data window and select **Convert to Dimension**.

The Order Quantity field is now displayed in the Dimensions area of the Data window and is a discrete quantity as indicated by the blue number icon.

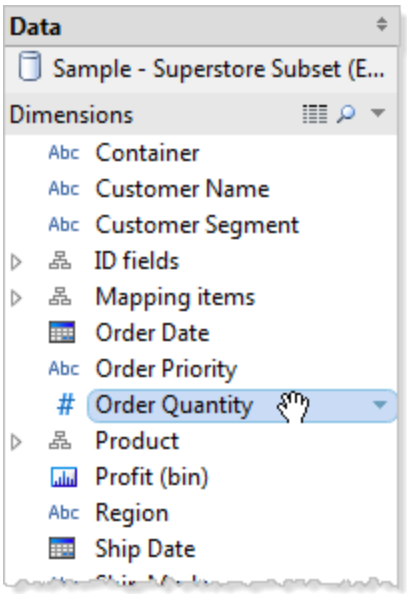

If you place the converted field on a shelve, it produces headers instead of an axis.

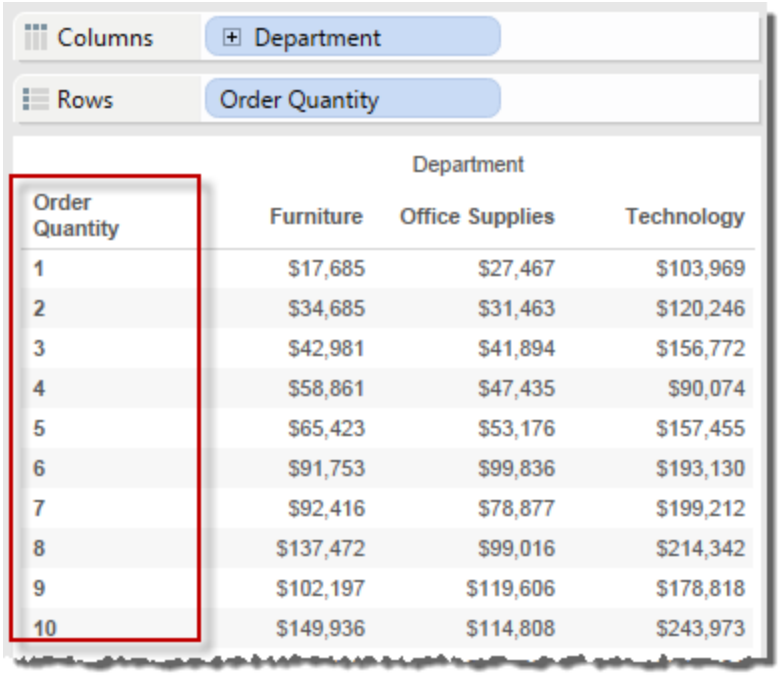

	Department		
Order Quantity	Furniture	Office Supplies	Technology
1	$17,685	$27,467	$103,969
2	$34,685	$31,463	$120,246
3	$42,981	$41,894	$156,772
4	$58,861	$47,435	$90,074
5	$65,423	$53,176	$157,455
6	$91,753	$99,836	$193,130
7	$92,416	$78,877	$199,212
8	$137,472	$99,016	$214,342
9	$102,197	$119,606	$178,818
10	$149,936	$114,808	$243,973

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 Sales as a function of Order Quantity. It is the same information presented in two different ways.

Discrete

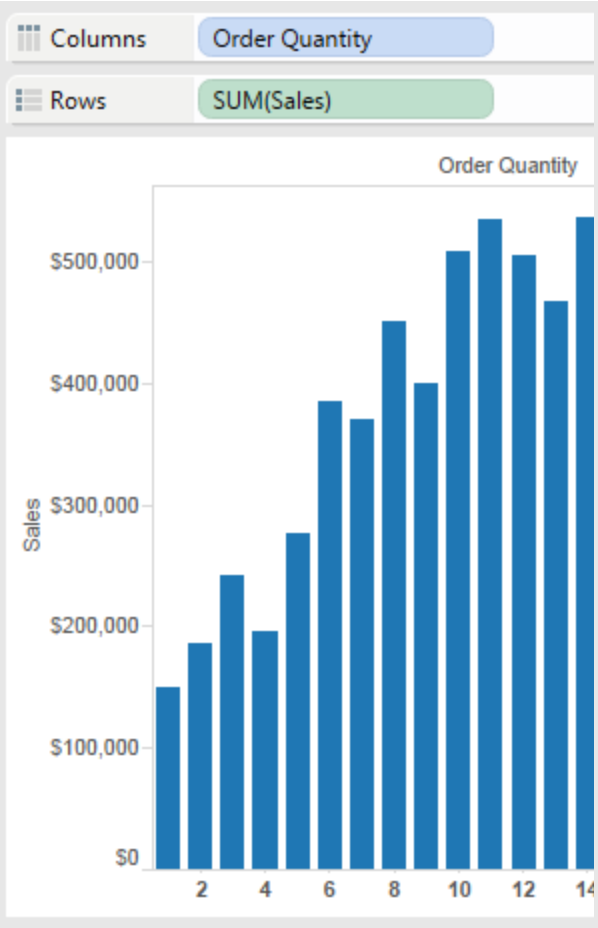

Each distinct order quantity is displayed as a header along the bottom of the bar chart. The Order Quantity field is colored **blue** on the Columns shelf.

Continuous

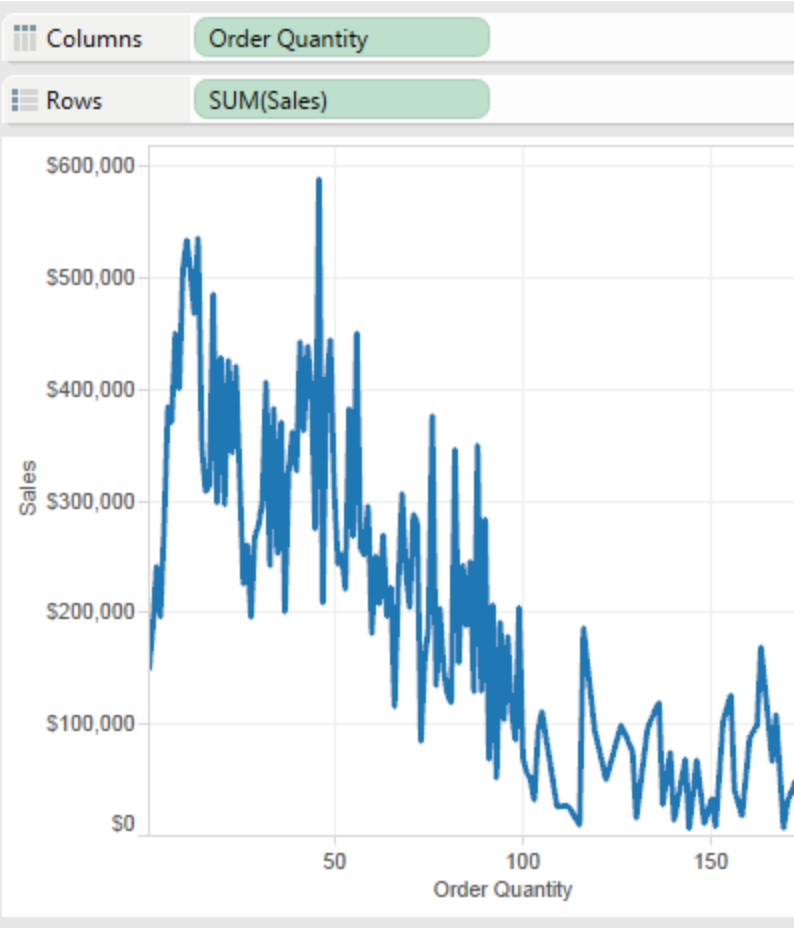

The order quantity values are drawn along a continuous axis along the bottom of the line chart. The Order Quantity field is colored **green** on the Columns shelf.

Whether a field is continuous or discrete is reflected in the color of the on the shelf. Blue fields are discrete and green fields are continuous fields.

Discrete fields always result in headers being drawn whenever they are placed on the Rows 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 see and the data itself. You can switch between continuous and discrete data roles using the field menu.

Converting Discrete Fields to Continuous

When you are using a relational data source you can convert any numeric or date field into a continuous field. Convert fields by doing one of the following

- Select **Continuous** on the field menu after dragging it to a shelf.

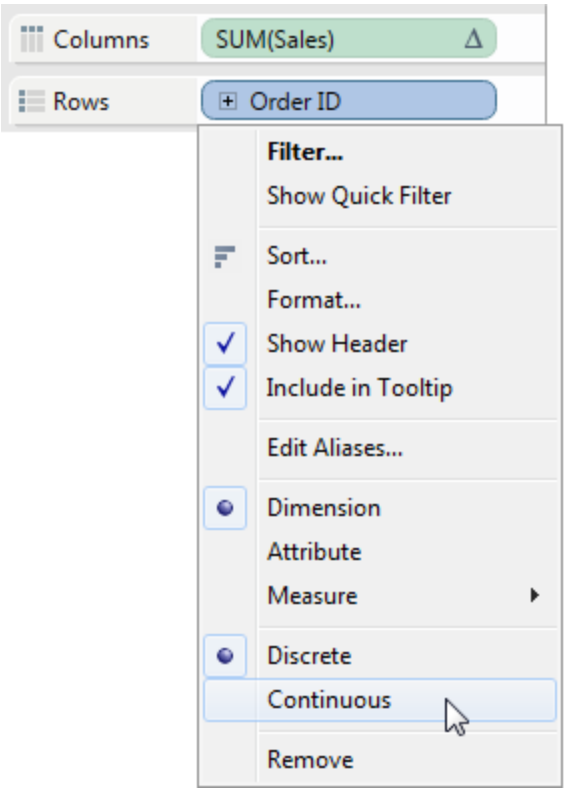

Converting a field that has already been placed on a shelf does not change its role in the Data window. Use this method when you want to convert the field for a specific view but don't want it to change for the whole workbook.

- Right-click a field in the Data window and select **Convert to Continuous**.

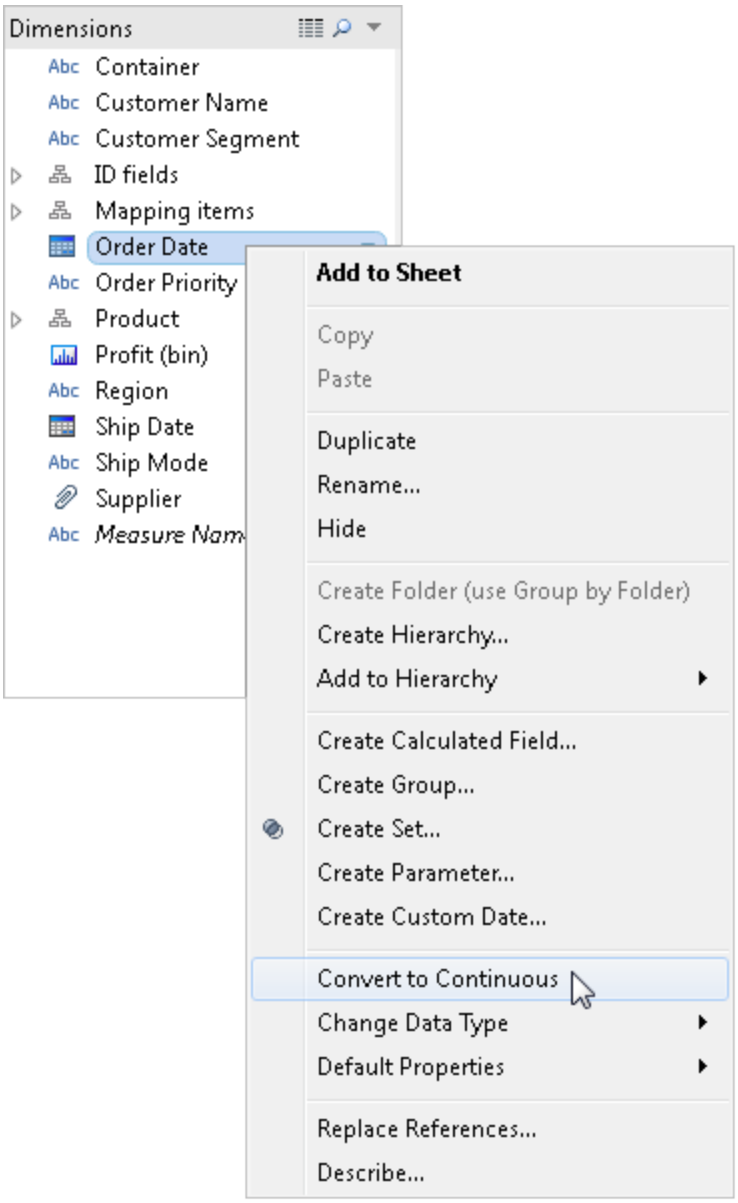

The Order Date field is still displayed in the Dimensions area of the Data window, but now uses a green icon, which indicates it is a continuous quantity.

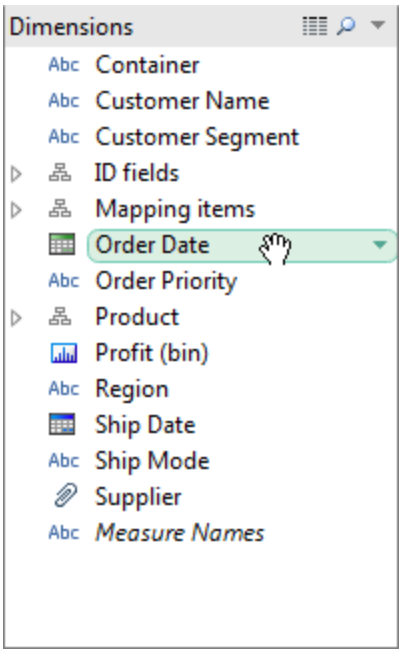

Placing the continuous dimension on a shelf produces an axis. However, the field is not a measure because you cannot aggregate it using the usual set of aggregation functions such as SUM and AVG.

Special Values

Some data requires special handling in Tableau. Specifically null values, unrecognized or ambiguous geographic locations, negative or zero values when working with a logarithmic scale, and negative or zero values when working with treemaps. When your data contains these special values, Tableau cannot plot them in the view. Instead, an indicator displays in the lower right corner of the view. Click the indicator to see more options for how to handle these values.

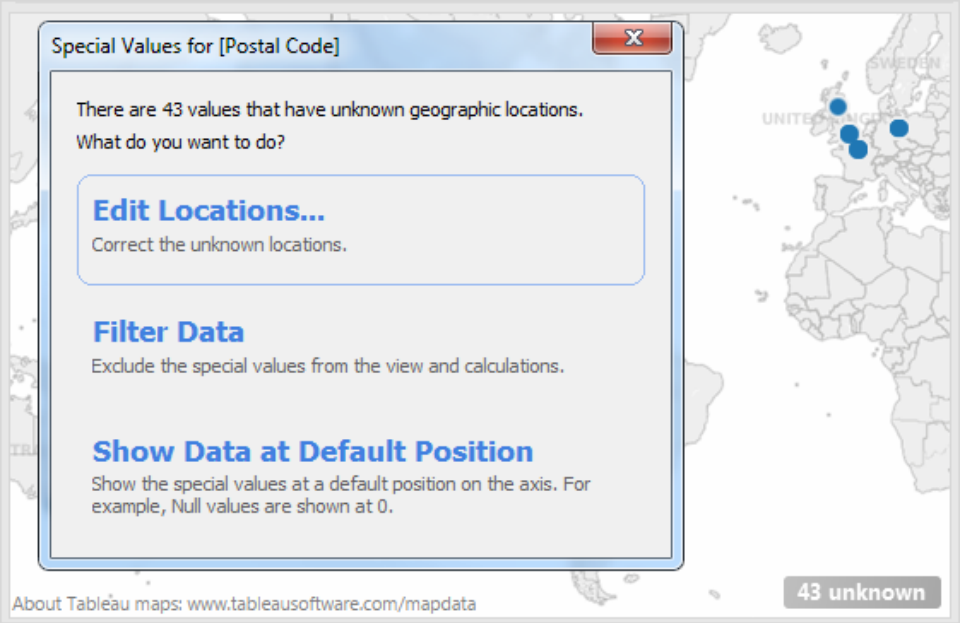

Null Numbers and Dates and Negative Values on Log Axes

When you drag a measure or continuous date to the view, the values are shown along a continuous axis. If the field contains null values or there are zero or negative values on a logarithmic axis, Tableau cannot plot them. Instead, these values are shown using an indicator in the lower right corner of the view. Click the indicator and choose from the following options:

- **Filter Data** - exclude the null values from the view using a filter. When you filter data, the null values are also excluded from any calculations used in the view.
- **Show Data at Default Position** - show the data at a default location on the axis. The null values will still be included in calculations. The default position depends on the data type. The table below defines the defaults.

Type of Data	Default Position
Numbers	0
Dates	12/31/1899
Negative Values on a Log Axis	1
Unknown Geographic Location	(0,0)

Unknown Geographic Locations

When working with maps and geographic fields, any unknown or ambiguous locations display in the indicator in the lower right corner of the view. Click the indicator and choose from the following options:

- **Edit Locations** - correct the locations by mapping your data to known locations.
- **Filter Data** - exclude the unknown locations from the view using a filter. The locations will not be included in calculations.
- **Show Data at Default Position** - show the values at the default position of (0, 0) on the map.

Zero and Negative Values in Treemaps

When working with treemaps, any null or zero values display in the indicator in the lower right corner of the view. Click the indicator and choose from the following options:

- **Filter Data** - exclude the unknown locations from the view using a filter. The locations will not be included in calculations.
- **Use Absolute Values** - use the absolute value to determine the size of the corresponding area in the view. For example, both values of 5 and -5 are shown as the same size.

If you don't know how to handle the values, you can choose to leave the special values indicator. Generally, you should continue to show the indicator so that you know there is data that is not being shown in the view. However, to optionally hide the indicator, right-click it and select **Hide Indicator**.

Building Data Views

This section discusses the basics of using Tableau to build views of your data. You will learn how to build views both manually and automatically; how to sort, filter, group, create sets; and how to work with dates and times, add reference lines, bands, and boxes, and view your data. The section concludes with a set of build-it-yourself exercises that show how you can create basic views and then enhance them to add more information and generate insight.

Parts of the View

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

Table Components

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

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

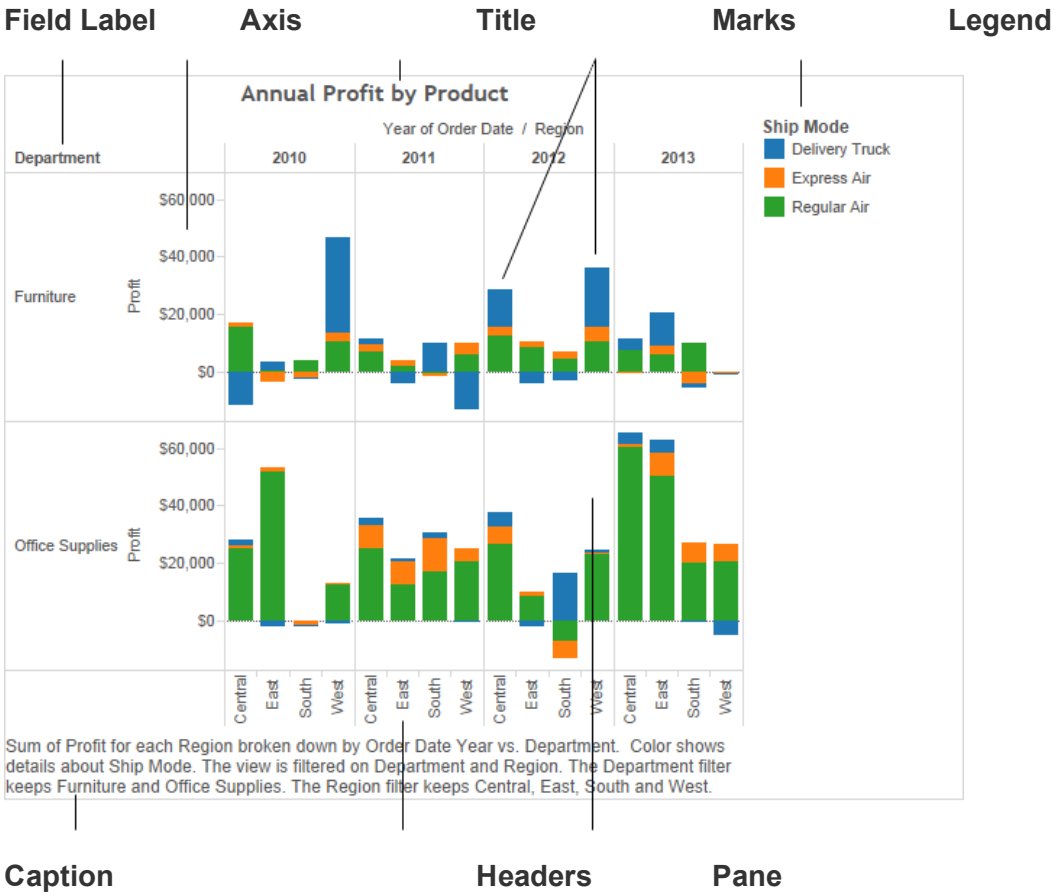

Headers

Headers are created when you place a dimension or discrete field on the **Rows** shelf or the **Columns** shelves. 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 the **Order Date** field and the row headers show the members of the **Category** field.

Columns: YEAR(Order Date)

Rows: Category

	Year of Order Date			
Category	2010	2011	2012	2013
Appliances	$231,824	$403,570	$329,008	$555,645
Binders and Binder Accesso..	$509,132	$317,544	$448,328	$822,498
Bookcases	$302,055	$455,991	$392,549	$524,182
Chairs & Chairmats	$751,254	$680,509	$932,647	$1,166,779
Computer Peripherals	$317,526	$330,973	$414,840	$567,921
Copiers and Fax	$405,263	$668,684	$585,012	$690,846
Envelopes	$40,519	$103,506	$82,191	$158,639
Labels	$13,571	$12,360	$24,464	$23,893
Office Furnishings	$283,638	$293,482	$357,810	$468,135
Office Machines	$1,040,107	$851,440	$958,732	$1,529,886
Paper	$183,361	$215,934	$206,196	$288,546
Pens & Art Supplies	$65,125	$81,250	$89,005	$102,846
Rubber Bands	$5,578	$7,252	$9,023	$10,396
Scissors, Rulers and Trimme..	$13,293	$13,048	$20,895	$135,388
Storage & Organization	$458,290	$458,368	$504,833	$705,052
Tables	$837,612	$765,507	$1,150,002	$1,174,831
Telephones and Communicat..	$687,072	$814,962	$990,268	$1,393,570

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

To hide headers:

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

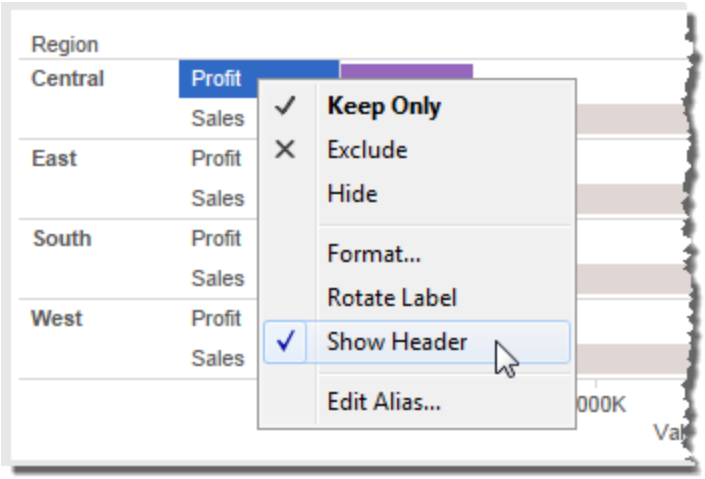

To show headers:

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

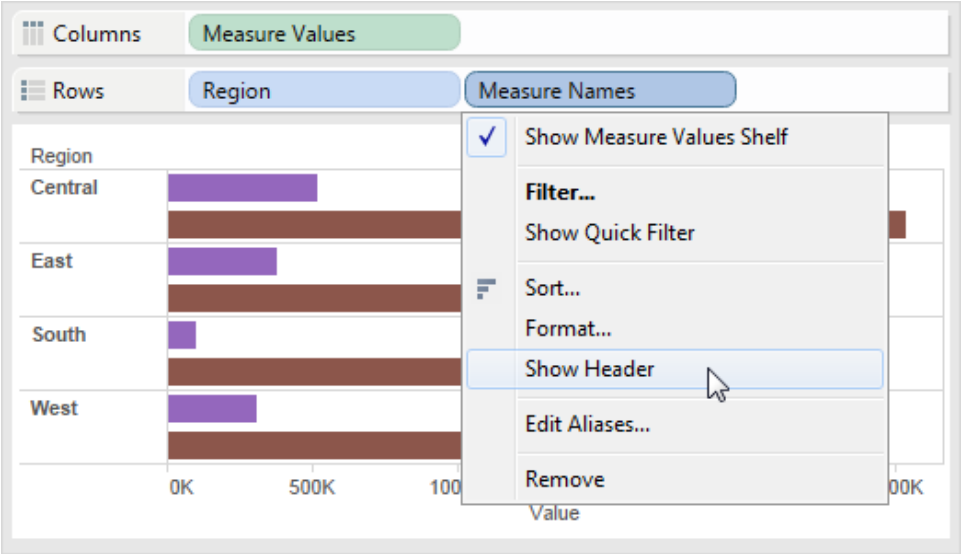

Hiding headers can be useful when you are working with multiple measures. For example, the view below shows both the sales and profit for each region along a single axis. You can see the view looks cluttered with the Measure Names headers showing. Because Measure Names is also indicated by the mark color,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

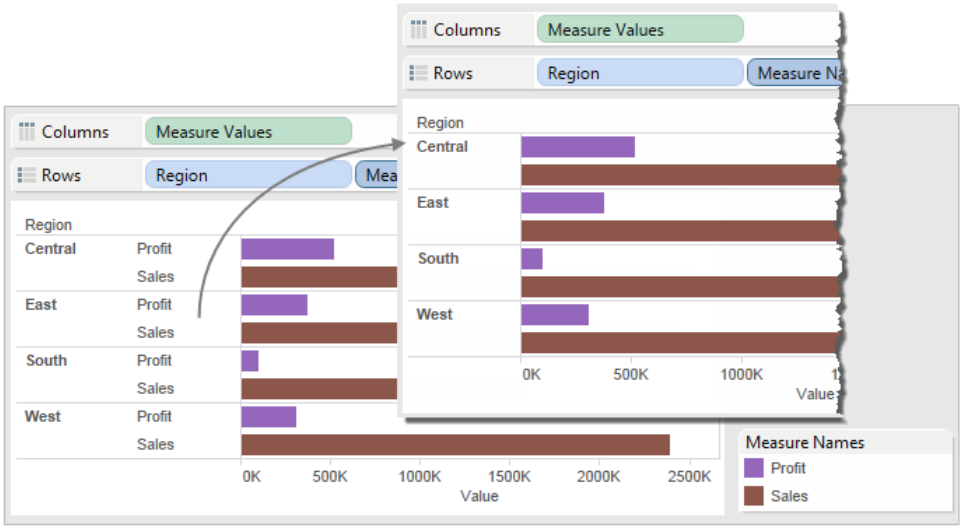

Axes

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

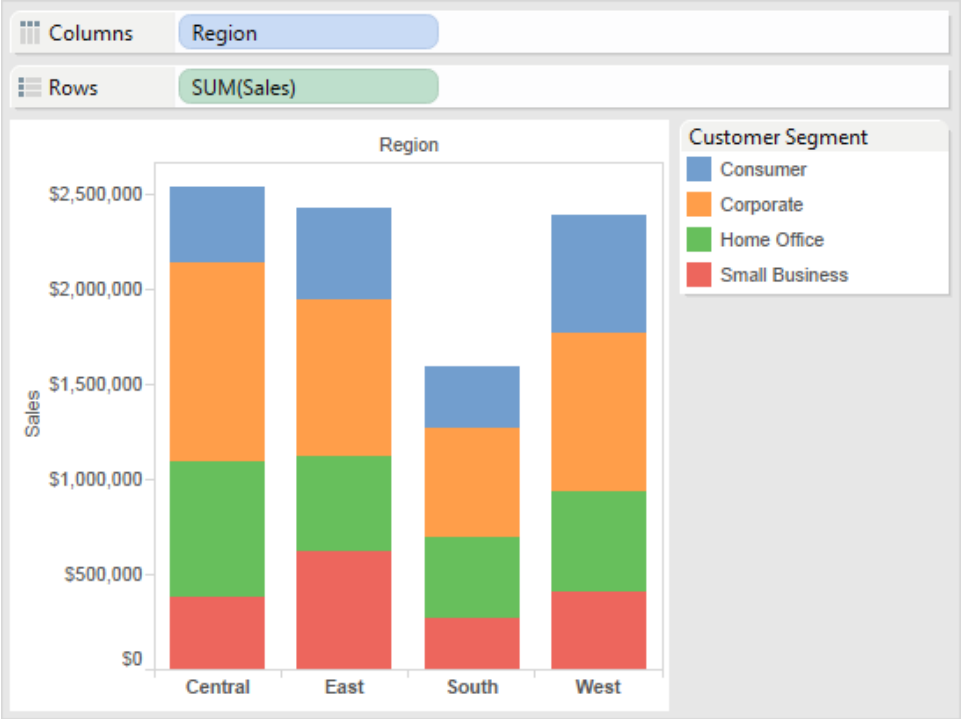

You can show and hide axes at anytime.

To hide axes:

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

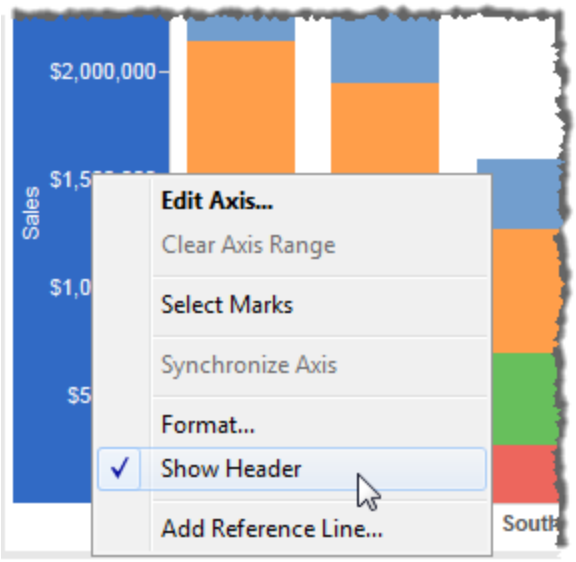

To show axes:

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

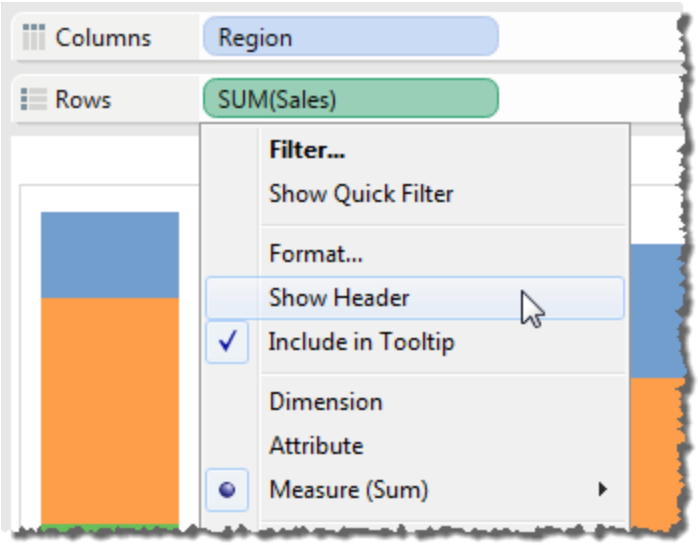

Panes

Panes are created by the intersection of rows and columns in a table. An example of a pane is highlighted in the view below.

Note: Tableau can display up to 1 million panes. If your view contains more than 1 million panes, you will need to filter your data so there are fewer panes in the view. For more information on filtering, see **Filtering** on page 334.

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 example, four of the 68 cells are highlighted in the view below.

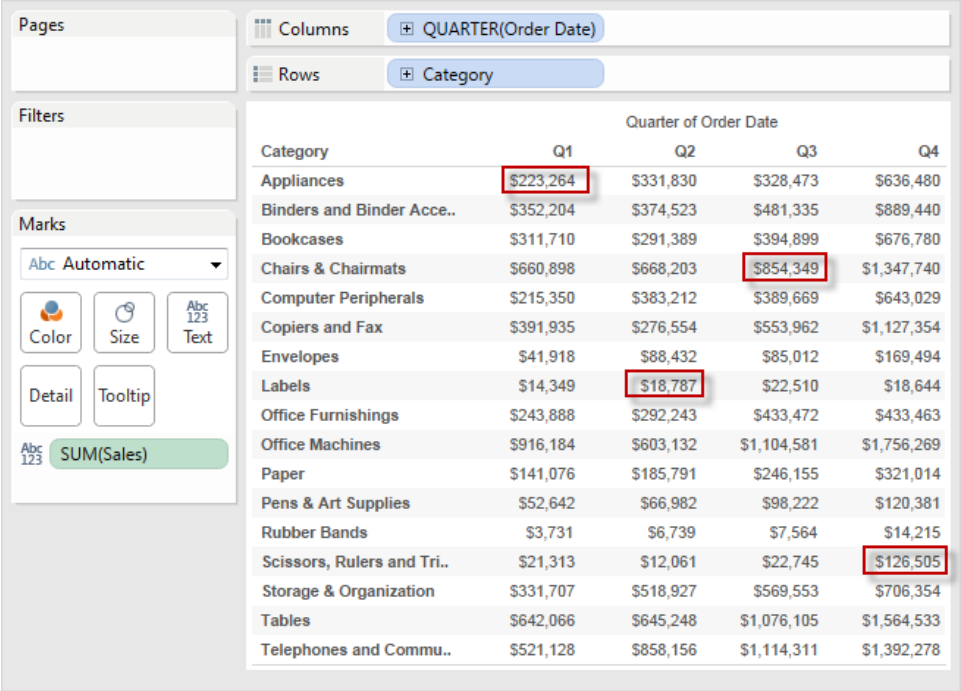

For other view types such as bar charts and scatter plots, identifying the cell is not always possible or useful.

Marks Card

When you drag fields to the view, the data are displayed using marks. Each mark represents the intersection of all of the dimensions in the view. For example, in a view with Region and Year dimensions, there is a mark for every combination of those two field (East 2011, East 2012, West 2011, West 2012, etc.).

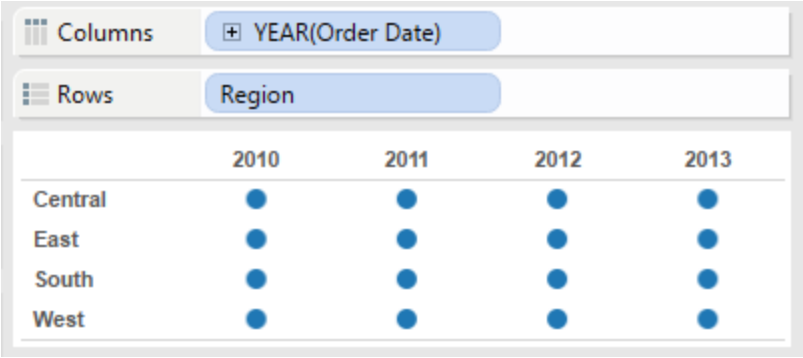

Marks can be displayed in many different ways including lines, shapes, bars, maps, and so on. You can show additional information about the data using mark properties such as color, size, shape, labels, etc. The type of mark you use and the mark properties are controlled by the Marks card. Drag fields to the Marks card to show more data. For example, the same view above is shown again below but this time with **Profit** on Color. With this additional information, it is clear that the Southern region was not profitable in 2010.

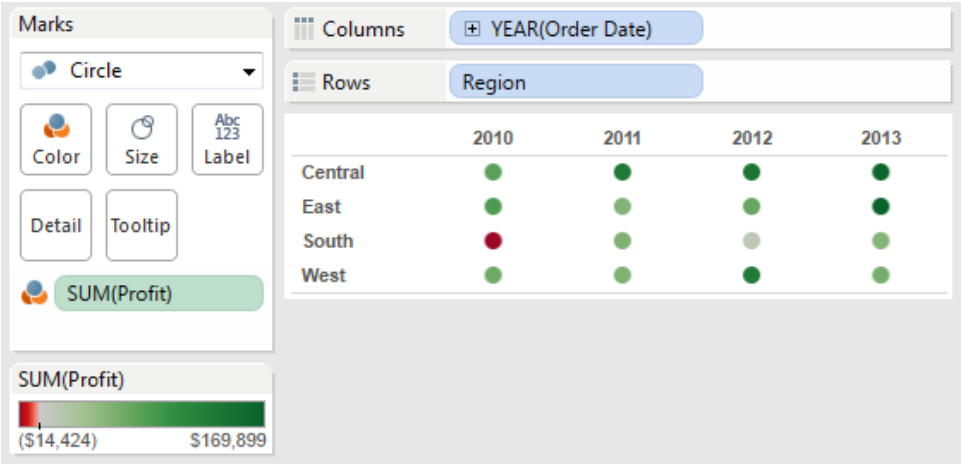

Control the marks in the view using the Marks card. Use the drop-down menu to specify the type of mark to show. Drag fields to the Marks card and use the drop-down controls to add more information to the view and control the color, shape, size, labels, and number of marks in the view.

Mark Types

Mark types are available from the **Marks** menu at the top of the Marks card. After selecting a mark type, you can further modify the marks by adding fields to Color, Size, Shape, and so on. The available mark properties are dependent on the type of mark you've selected.

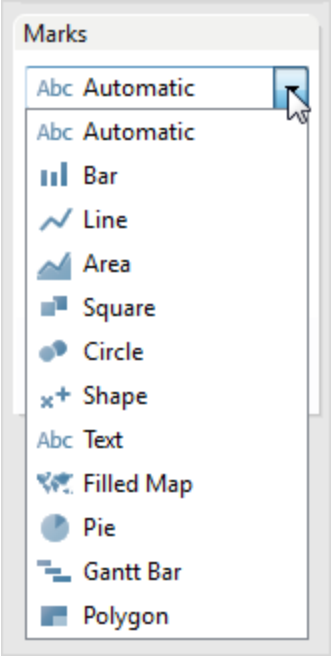

Automatic Mark

When the **Mark** menu is set to Automatic, Tableau automatically selects the best mark type for your data view. The mark type that is automatically selected is determined by the inner fields on

the **Rows** and **Columns** shelves—that is, the fields farthest to the right on the shelves. The icon in the Mark menu indicates which type of mark was automatically selected.The table below describes the type of mark that is selected in each scenario.

Type of fields in the view	**Auto-matic Mark Type**	**Example**
Dimen-sions as inner fields on both the **Rows** and **Columns** shelves.	Text	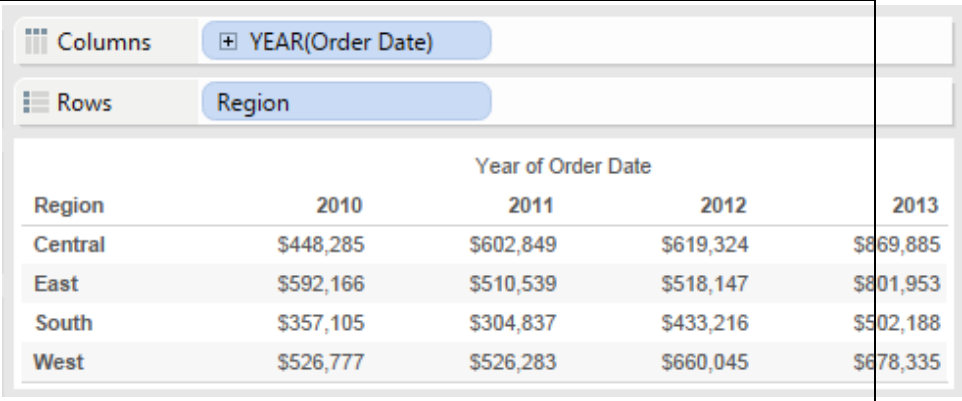
Measures as inner fields on both the **Rows** and **Columns** shelves.	Shape	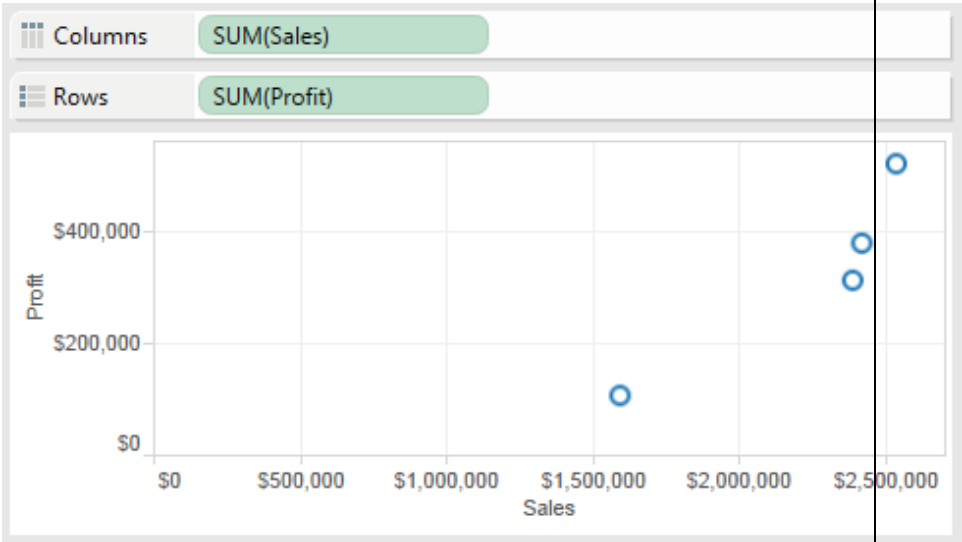

Type of fields in the view	Auto-matic Mark Type	Example
Dimen-sion and measure as inner fields on **Rows** and **Columns** shelves.	Bar	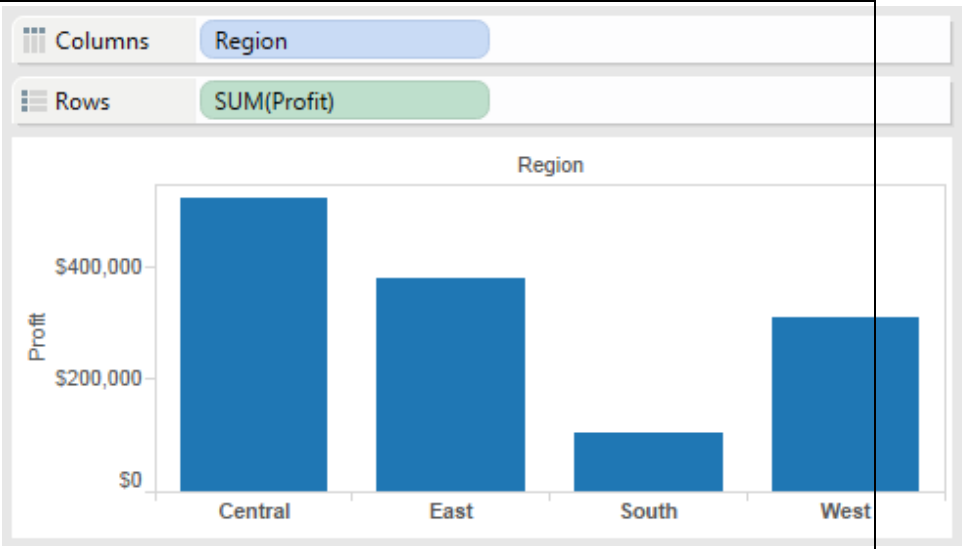
A date field and a measure as the inner fields on the **Rows** and **Columns** shelves.	Line	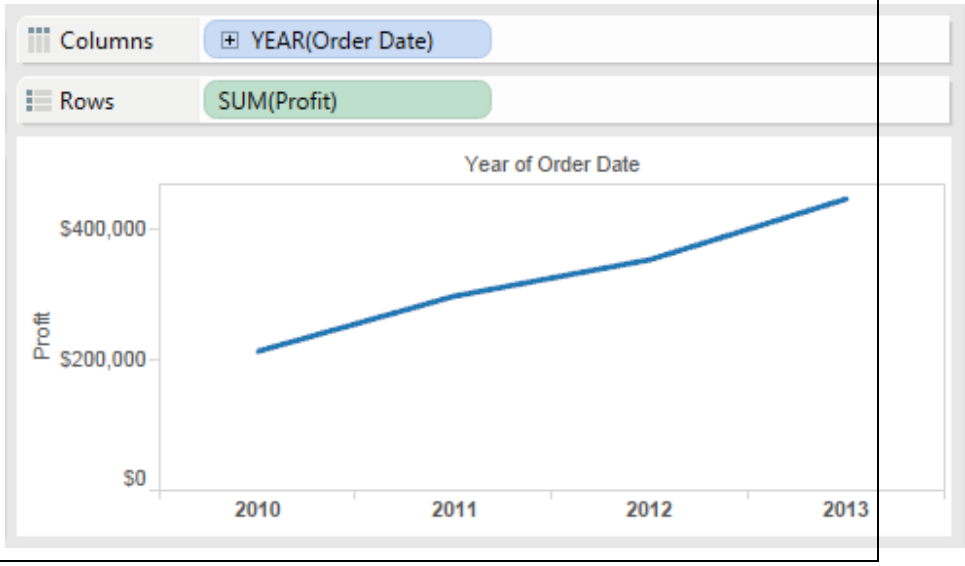

You can override the default selection and use any mark type that provides insight into your data. However, you should exercise some caution when manually selecting a mark type because the resulting view might hide important information about your data.

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). The inner fields are the

fields furthest to the right on the shelves. If the dimension is a date dimension, the Line mark is chosen instead.

- You select **Bar** from the **Mark** menu.

Marks are automatically stacked.

The data view shown below displays a dimension as the columns of the table, and a measure as the rows of the table. It is also color-encoded by a dimension (**Department**). Because the **Mark** menu is set to Automatic, data is displayed using bars.

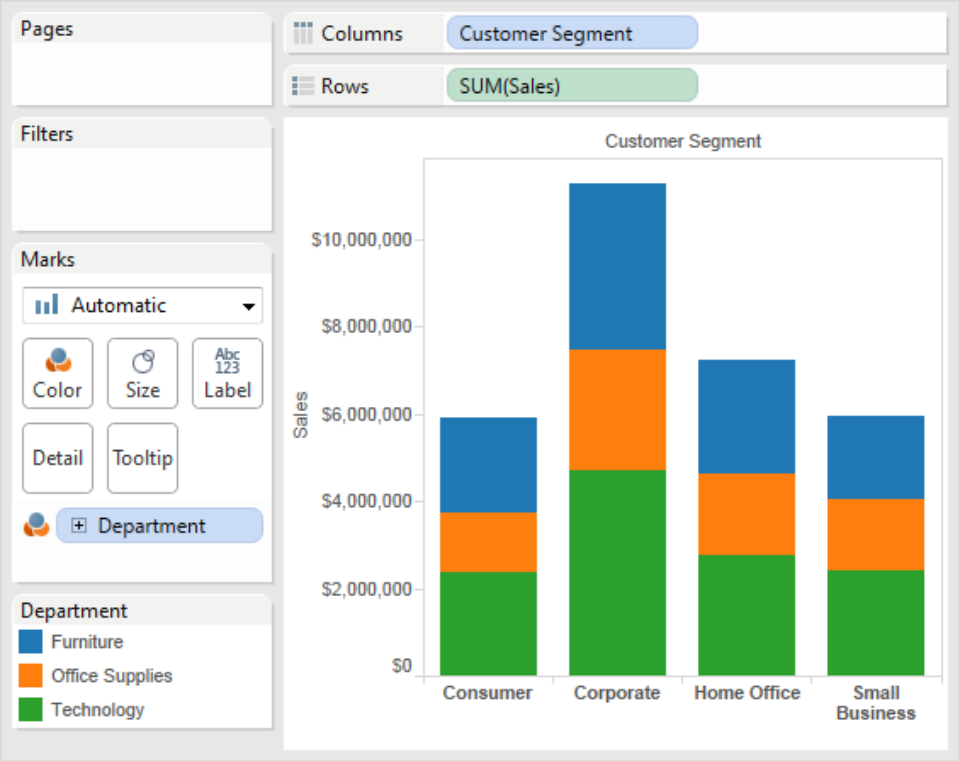

To try some hands-on exercises for building bar charts, see **Building Bar Charts** on page 484.

.

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 date dimension as the columns of the table and a measure as the rows of the table.

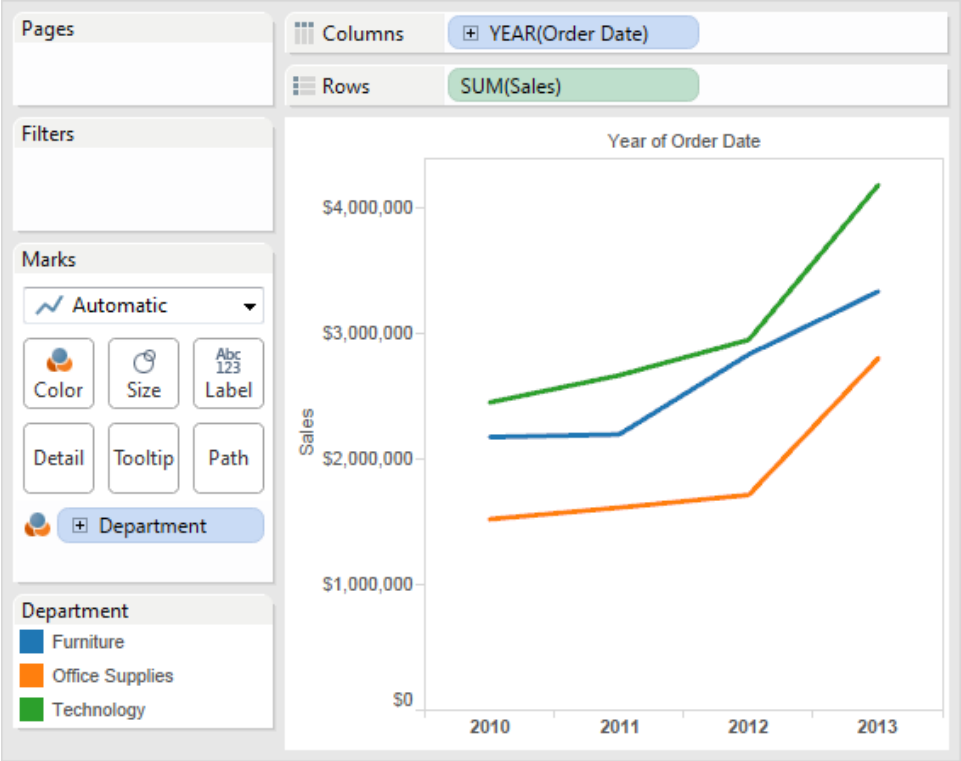

With the line mark type, you can specify the drawing order of the line by placing a field on **Path**.

As the density of data increases, trends are often easier to see when using lines. This view shows 93 data points.

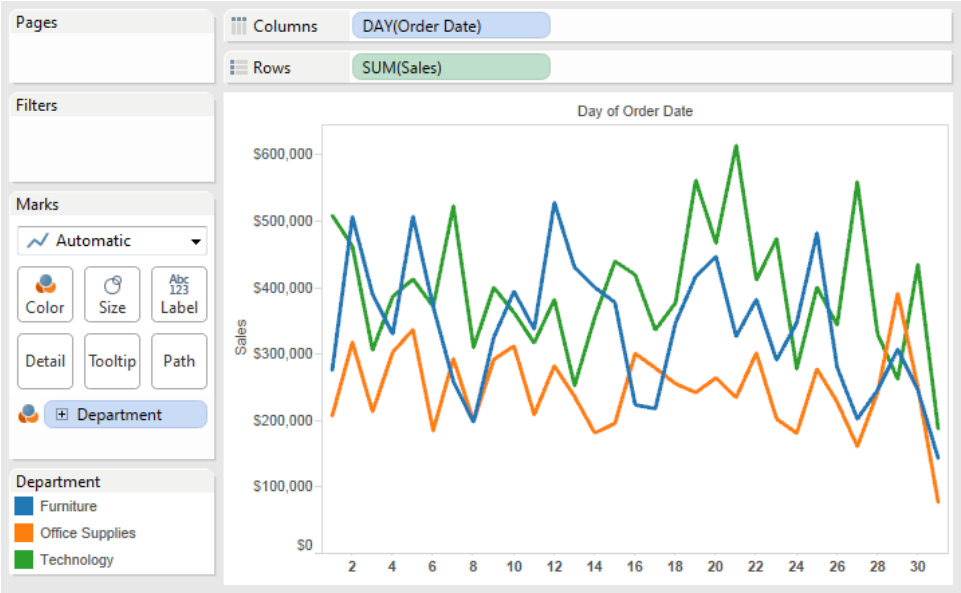

To try some hands-on exercises for building line charts, see **Building Line Charts** on page 499.

.

Area Mark

The Area mark type is useful when you have a view where the marks are Stacked and do not overlap. In an area chart, the space between each line and the next line is filled with a color. This type of chart is not the best way to show specific values along the line but it can clearly show the total values so you can get an idea of how a dimension is contributing to an overall trend. Tableau displays your data using areas when:

- The Mark menu is set to Automatic with one or more measures plotted against a date dimension or continuous dimension. The **Analysis > Stack Marks** option must also be selected.

The view below shows a date dimension on the **Columns** shelf and a measure on the **Rows** shelf. Each line represents the sales for a product department. The lines have been stacked so that they do not overlap.

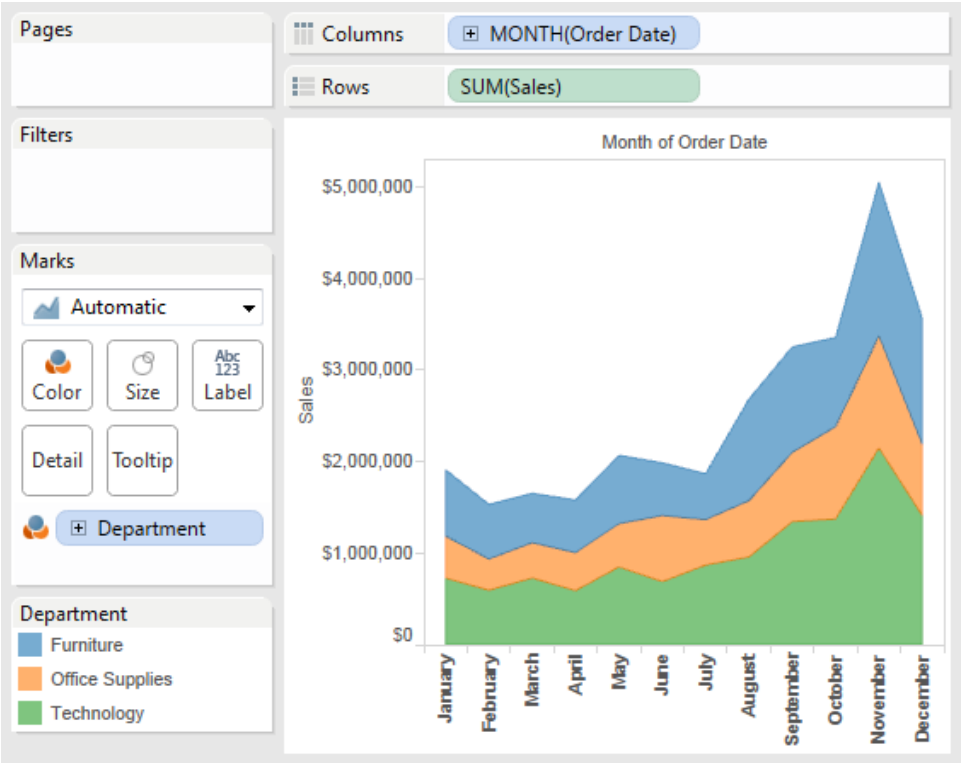

> You can add labels, edit the color legend, and highlight areas. In addition, you can turn on **Color** on page 257 effects such as mark borders.

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, you create a completely different view. In particular, by placing a measure on Color, you can use square marks to create a heat map. To try some hands-on exercises for building heat maps, see **Building Heat Maps** on page 508.

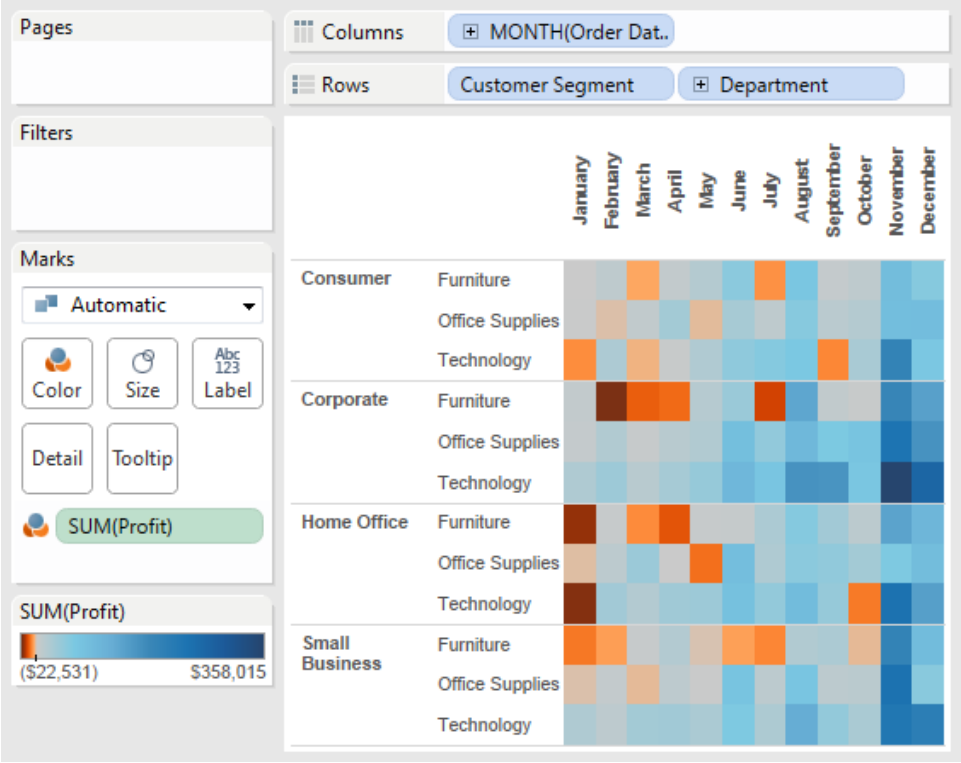

> To reproduce this view, you may need to select **Cell Size > Square Cell** from the **Format** menu on the toolbar, and then adjust the size of the squares using the Size slider.

When you add additional levels of detail by placing a dimension on the Detail, Color, Shape, Size, or Text target on the Marks card, the squares are shown side-by-side and wrap to fill the cell. If the application window is too small to show all of the squares, an ellipses is shown to indicate that there are more values than can be shown.

You can also use the square mark to create treemaps. A treemap displays hierarchical data as a set of nested rectangles. To try a hands-on exercise for building a treemap, see **Building Treemaps** on page 528.

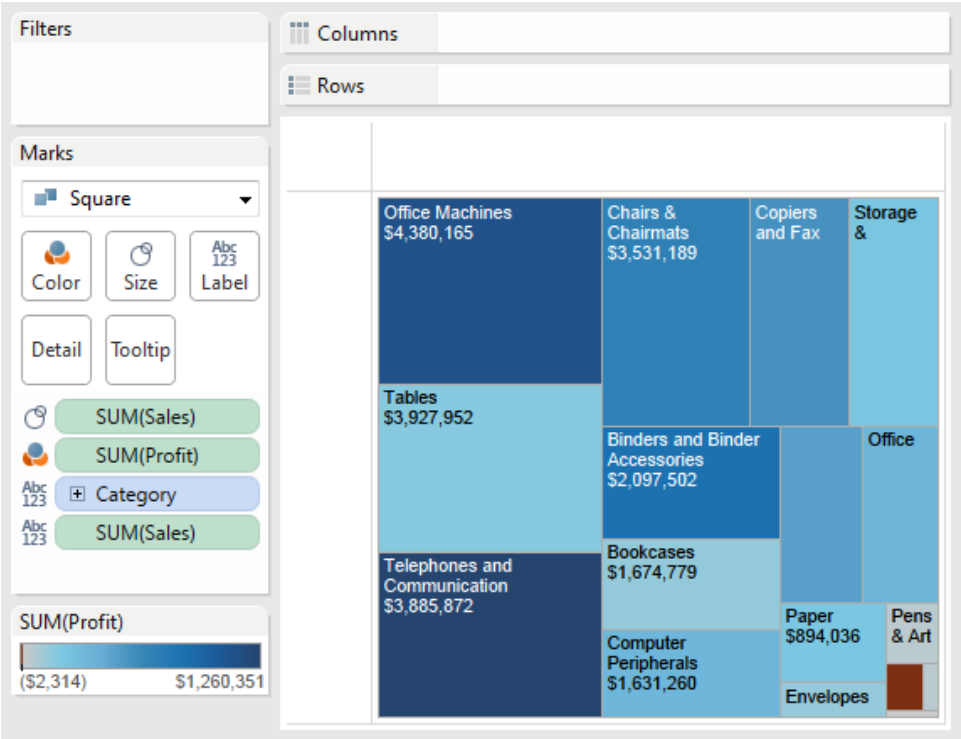

Circle Mark

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

If the mark type was set to Automatic, Tableau would display the data using a shape (that is, an open circle).

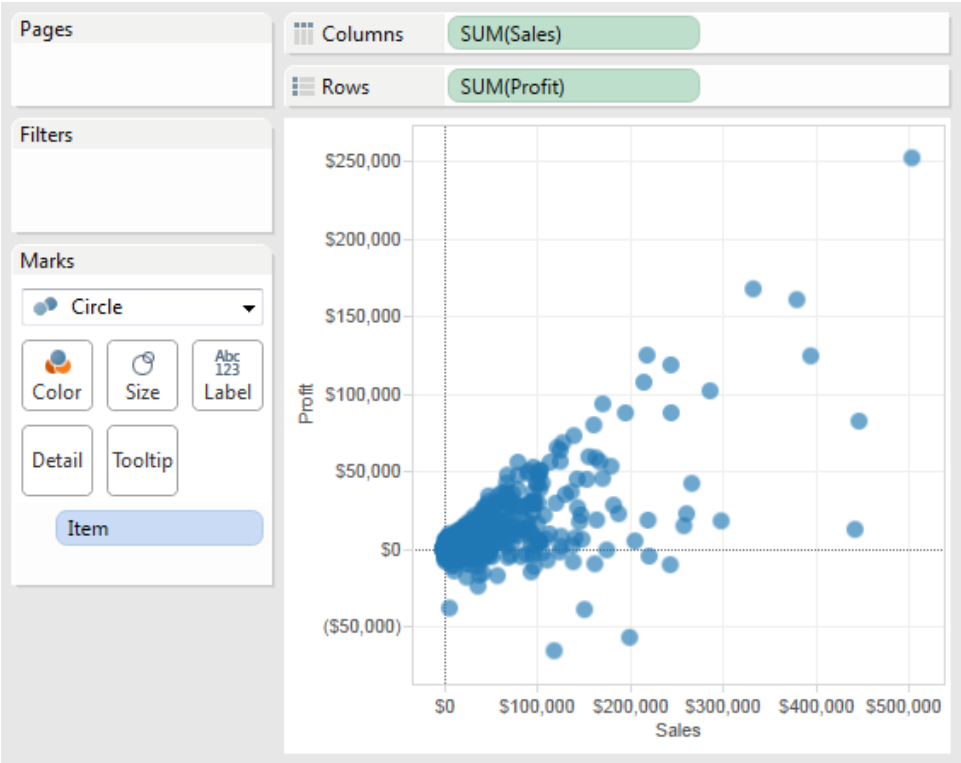

Shape Mark

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 **Marks** menu.

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

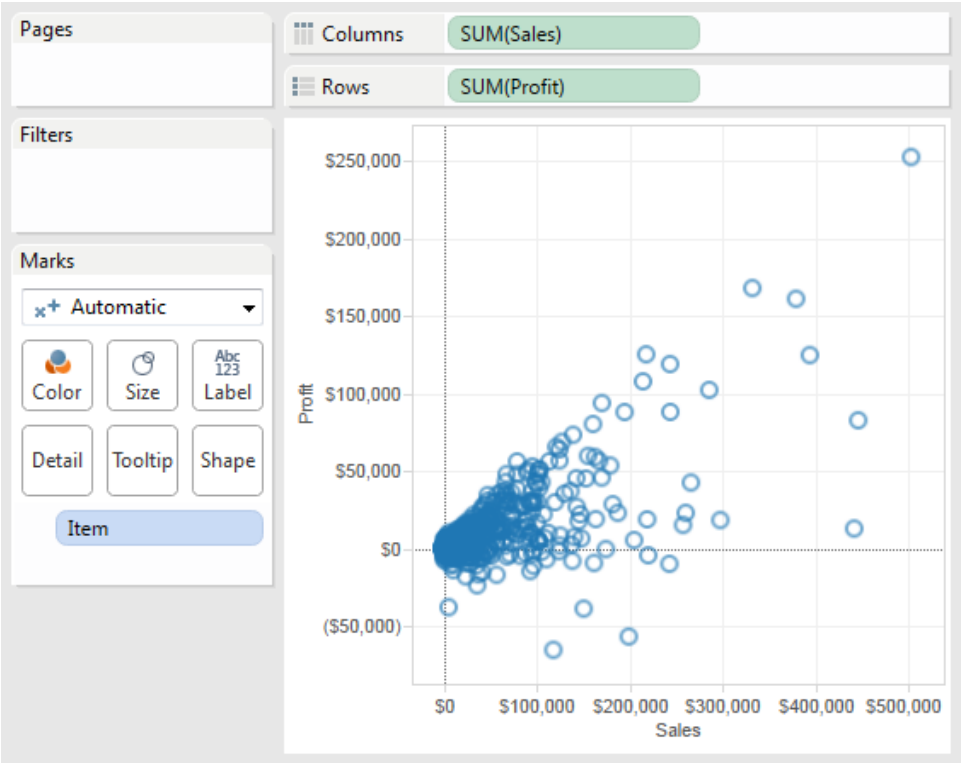

By default, the shape used is an open circle. You can select a different shape by clicking Shape on the Marks card. Twenty unique shapes are available:

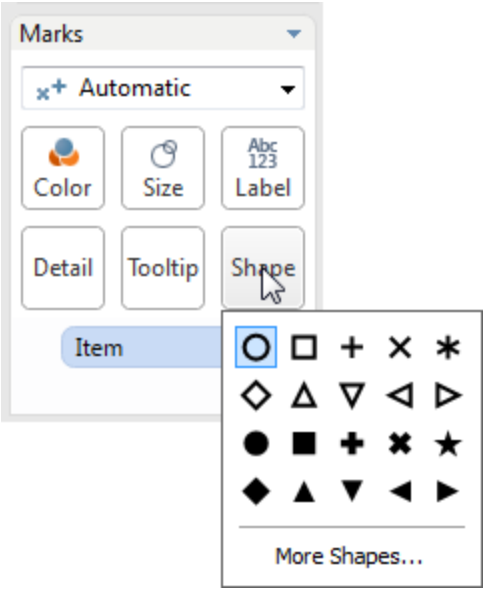

To add more data to the view, you can place a field on **Shape**. Tableau separates the marks according to the values in the field. If the field is a dimension, each member is assigned a unique shape. If the field is a measure, the measure is automatically binned into distinct

buckets and each bucket is assigned a unique shape. The shape legend displays how shapes are distributed.

As shown below, the **Ship Mode** dimension is used on Shape to encode each mark with information about how the order was shipped.

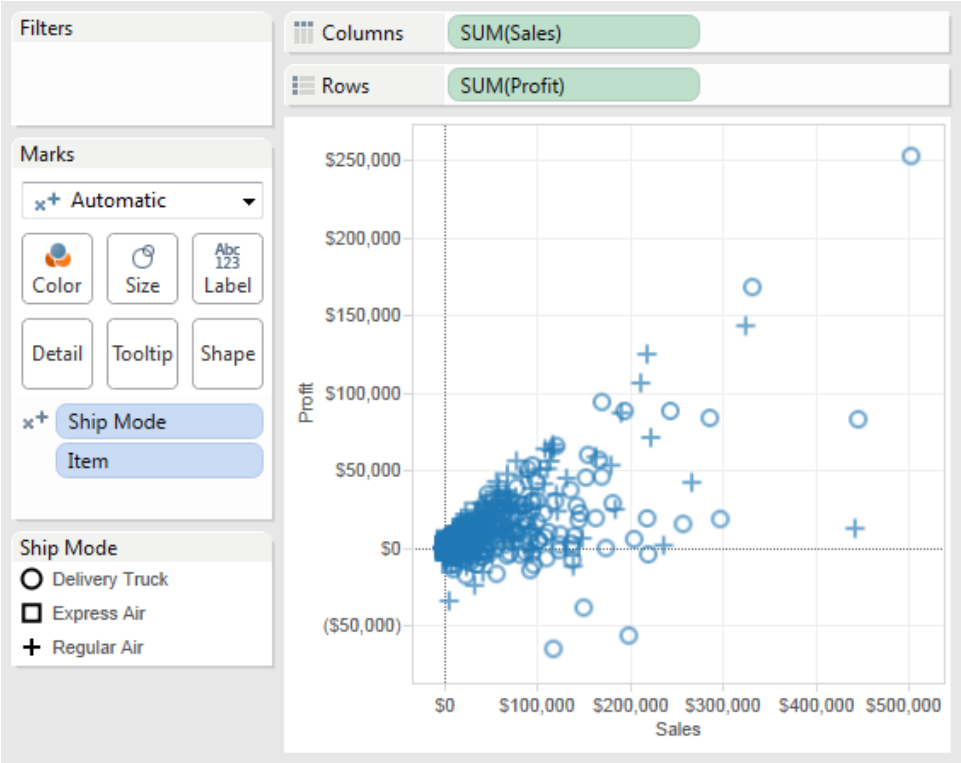

Text Mark

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. The inner fields are the fields farthest to the right on the shelves.
- You select Text from the **Mark** menu.

Initially, the data are displayed using the Abc icon.

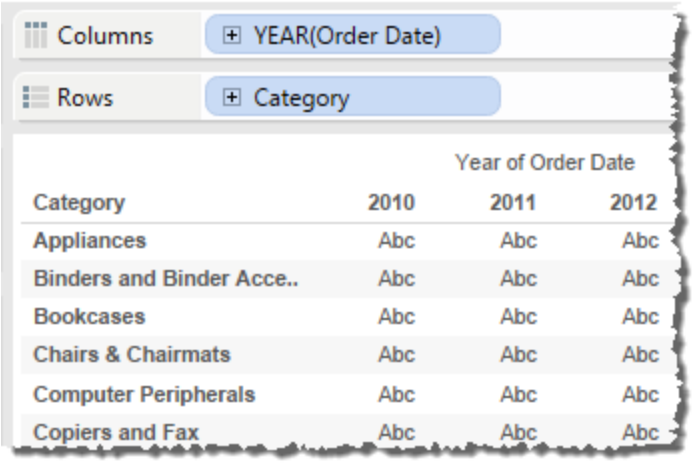

To complete the view, you must place a field (typically a measure) on Text. In the view below, the **Sales** measure, which is aggregated as a summation, is used to complete the table.

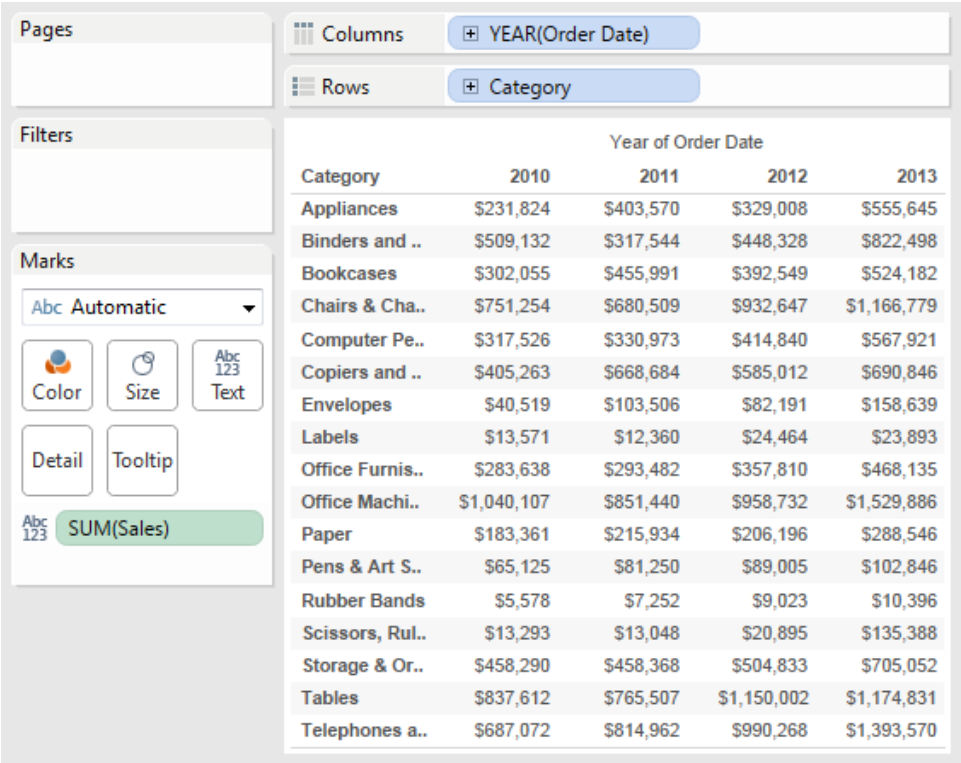

Category	2010	2011	2012	2013
Appliances	$231,824	$403,570	$329,008	$555,645
Binders and ..	$509,132	$317,544	$448,328	$822,498
Bookcases	$302,055	$455,991	$392,549	$524,182
Chairs & Cha..	$751,254	$680,509	$932,647	$1,166,779
Computer Pe..	$317,526	$330,973	$414,840	$567,921
Copiers and ..	$405,263	$668,684	$585,012	$690,846
Envelopes	$40,519	$103,506	$82,191	$158,639
Labels	$13,571	$12,360	$24,464	$23,893
Office Furnis..	$283,638	$293,482	$357,810	$468,135
Office Machi..	$1,040,107	$851,440	$958,732	$1,529,886
Paper	$183,361	$215,934	$206,196	$288,546
Pens & Art S..	$65,125	$81,250	$89,005	$102,846
Rubber Bands	$5,578	$7,252	$9,023	$10,396
Scissors, Rul..	$13,293	$13,048	$20,895	$135,388
Storage & Or..	$458,290	$458,368	$504,833	$705,052
Tables	$837,612	$765,507	$1,150,002	$1,174,831
Telephones a..	$687,072	$814,962	$990,268	$1,393,570

> You can create a cross-tab of any data view by selecting the **Worksheet** > **Duplicate as Cross-tab** menu item.

When you add additional levels of detail by placing a dimension on the Detail, Color, Shape, Size, or Text target on the Marks card, the values are shown side-by-side and wrap to fill the

cell. Allowing the text to stack up makes it possible to create word cloud visualizations like the one shown below.

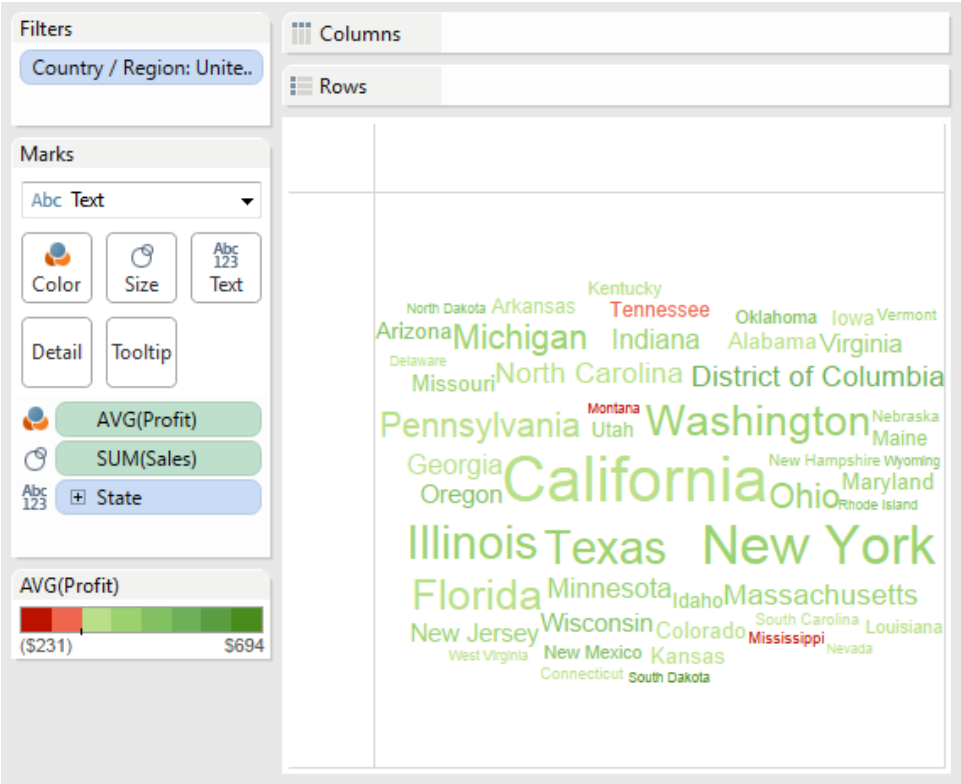

If the application window is too small to fill all of the text, the cell displays an ellipses to indicate that there are too many values than can fit.

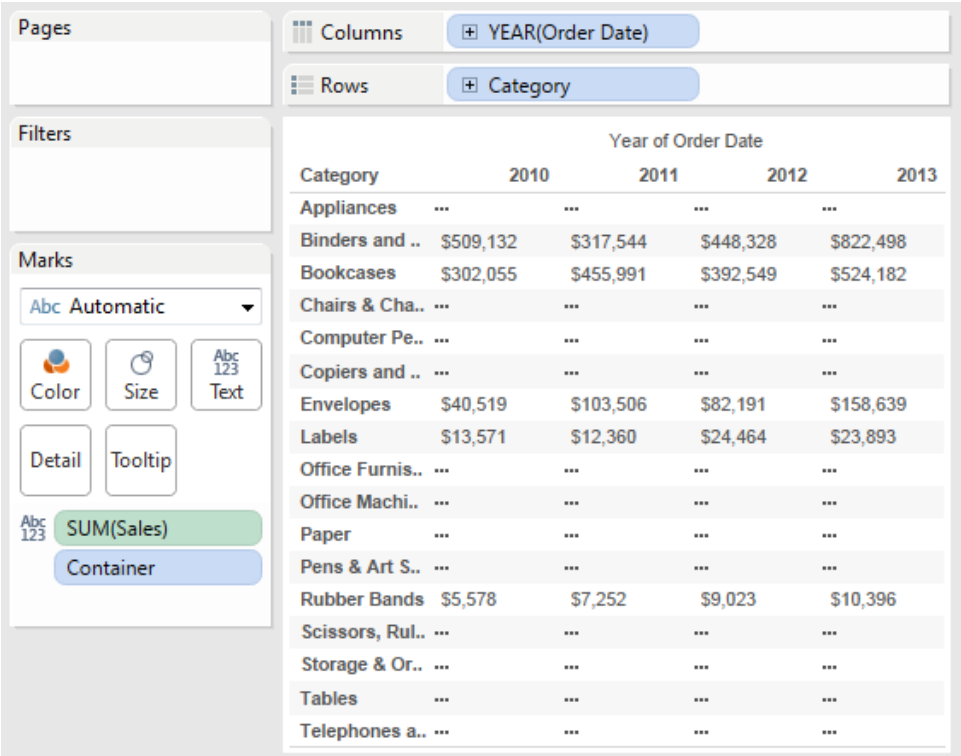

To try some hands-on exercises for building text tables, see **Building Text Tables** on page 492.

Filled Map Mark

The filled map mark type uses geocoding to fill a polygon with a color based on data. The primary use of the filled map mark type is for creating choropleth maps (also known as thematic maps or data maps). The area to be filled is defined by the geographic fields used in the view. Tableau displays data using a filled map when:

- A geographic dimension is on the Detail target on the Marks card, along with Latitude and Longitude on the **Rows** and **Columns** shelves. A measure or continuous dimension is added to the Color target.
- You select Filled Map from the Mark menu.

The data view below shows a map with the **State** geographic dimension on the Detail target of the Marks card. The **Profit** measure is on the Color target.

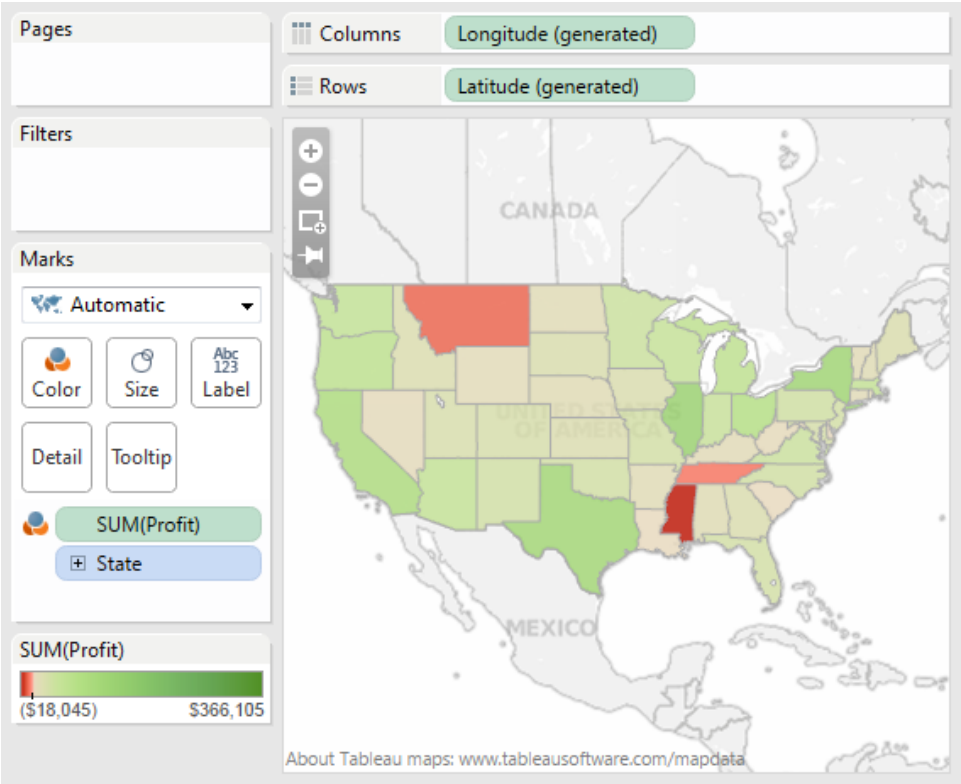

The polygon for each state is filled with a color based on the profit for that state. The profit in Montana is low (in fact, negative) while the profit in Texas is high.

Another way to use the Filled Map mark type is when you are plotting two measures against each other. By default, placing measures on both the **Rows** and **Columns** shelves creates a scatter plot. However, adding a geographic dimension allows you to change the mark type to Filled Map. Each mark becomes the area defined by the geographic field. For example, the view below shows **Sales** versus **Profit** for a selection of states. Rather than showing a round mark with a label for each state, the view uses the Filled Map mark type to draw the outline of each state. While this works well for recognizable areas, it is not the best choice for areas that are similar in shape or difficult to recognize.

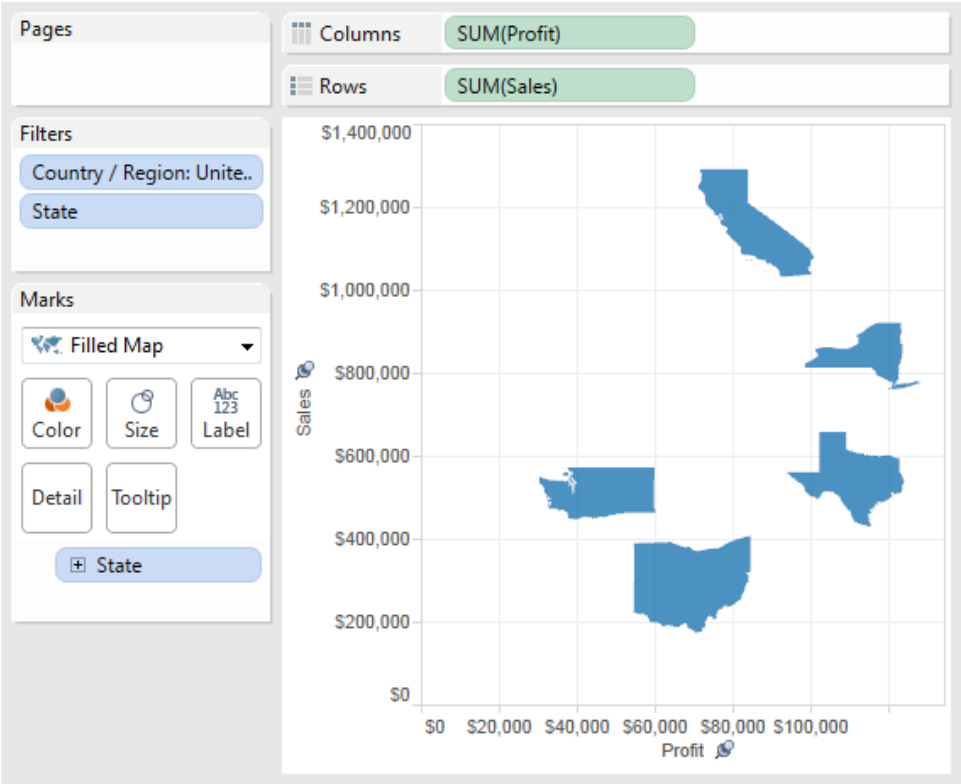

Pie Mark

You can use the Pie mark to show proportions. Although generally this type of information can be better shown using stacked bar charts, there are cases where pie marks can be very effective—for example, 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 from the **Mark** menu.

When you select the Pie mark type, an additional target named Angle is displayed on the Marks card. The Angle target determines the angular measure of the pie wedges. For example, if you place a measure such as **Sales** on the Angle target, the 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 target.

The view below shows the time it took to ship products by various ship modes. The data overlays a map and shows data by state. We can see that delivery by truck takes the longest to ship in this particular region. That is probably why orders are always shipped by air to the 48858 zip code.

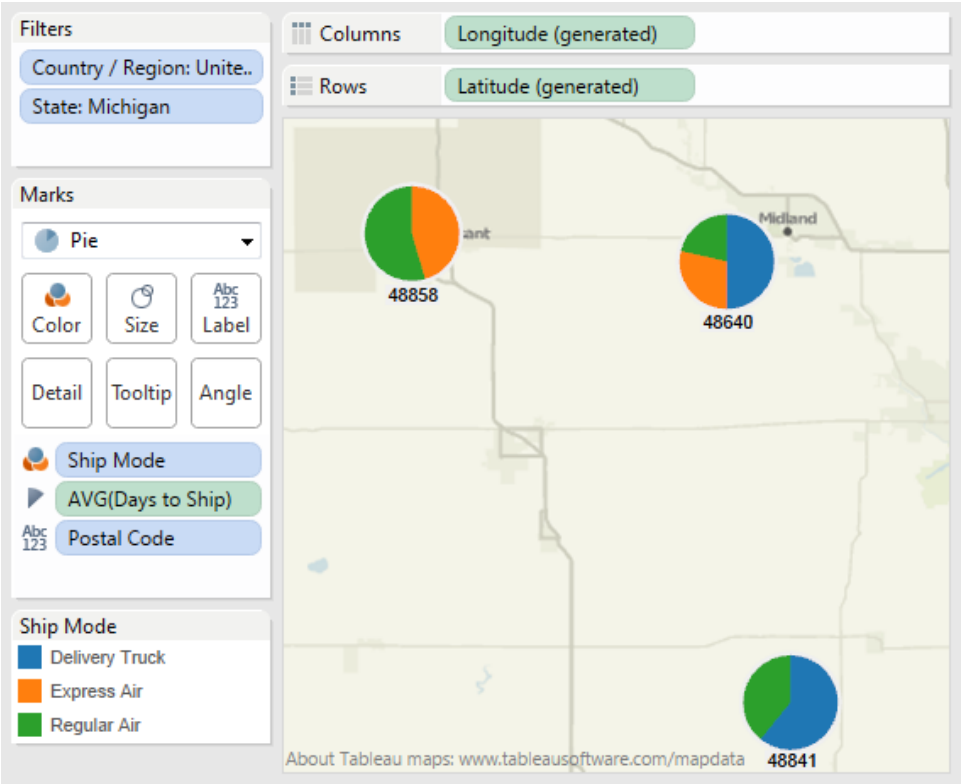

To try some hands-on exercises for pie charts, see **Building Pie Charts** on page 520.

Gantt Bar Mark

The Gantt Bar mark type is useful for viewing 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 target.

The view 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. Selecting **Gantt Bar** and adding additional fields will display a view like the following.

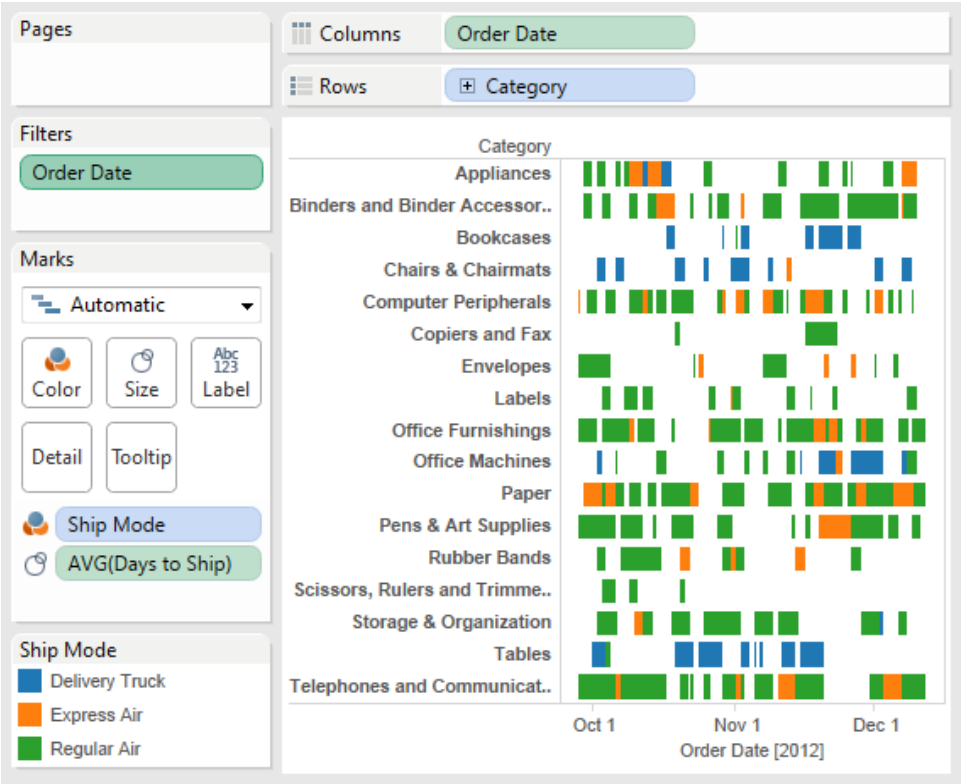

In particular, placing the **Days to Ship** measure on the Size target causes every bar in the view to be drawn with a length that indicates the delivery time of an order. Additionally, placing the **Ship Mode** dimension on the Color target causes each bar to be colored according to the ship mode.

To try a hands-on exercise for building a Gantt Bar chart, see **Building Gantt Bar Charts** on page 513.

Polygon Mark

Polygons are points connected by lines enclosing an area. 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.

> 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. Manually selecting Polygon and adding additional fields to the view causes a different view to be created.

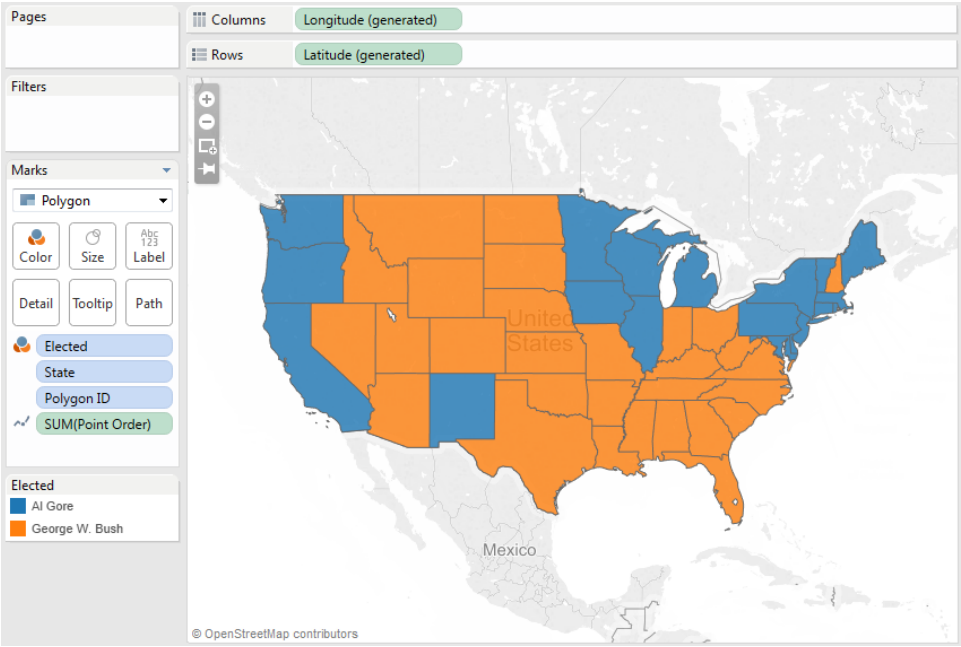

Every state is considered to be a polygon in the data source. The **PolygonID** field on the **Detail** target is distinct for each US state. You can remove states from the view by filtering this field.

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

Mark Properties

You can control the mark properties using the Marks card. For example, you can control the colors, size, shape, etc. of the marks in the view. Drag fields to each property to encode the marks using your data. Click each property on the marks card to open a drop-down control where you can further tune the mark properties.

Color

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

The Color property 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.

> 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. To set the default color encodings for a field, right-click the field in the Data window and select **Default Properties >Color**.

Categorical Colors

When you add a dimension to the Color target on the Marks card, a categorical legend is added based on the members of the field. To modify the colors used in the legend, right-click on the legend and select **Edit Colors** or double-click on the legend. The Edit Colors dialog box for a categorical legend is shown below.

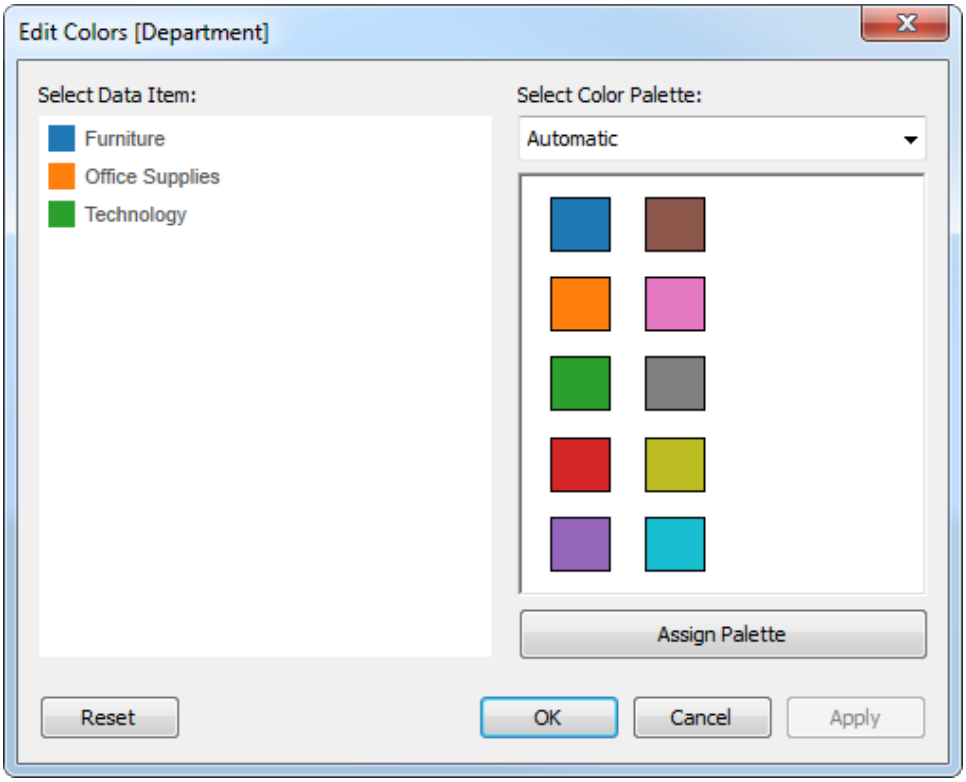

To change the color of a member:

1. Select the member on the left and then select the new color in the palette on the right.
2. When finished, click **OK** to close the Edit Colors dialog box.

You can select a different color palette from the drop-down menu in the upper right of the Edit Colors 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.

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 box.

To return to the automatic color settings, click **Reset** in the Edit Colors dialog box.

Quantitative Colors

When you add a measure to the Color target on the Marks card, 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 Colors 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.

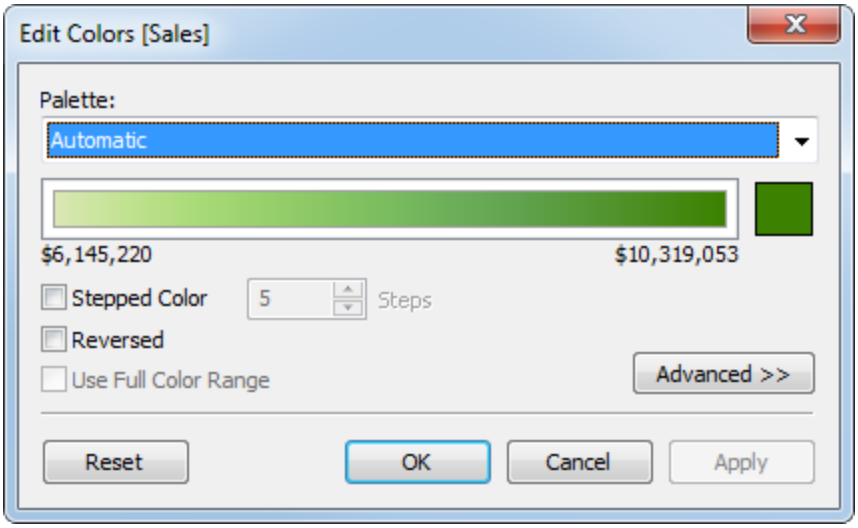

To change the color used in the range, click on the color indicator to the right of the range and select a new color in the spectrum. You can select a new palette from the Palette drop-down menu. You can choose between a sequential palette and a diverging palette. A sequential palette shows a simple 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.

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.

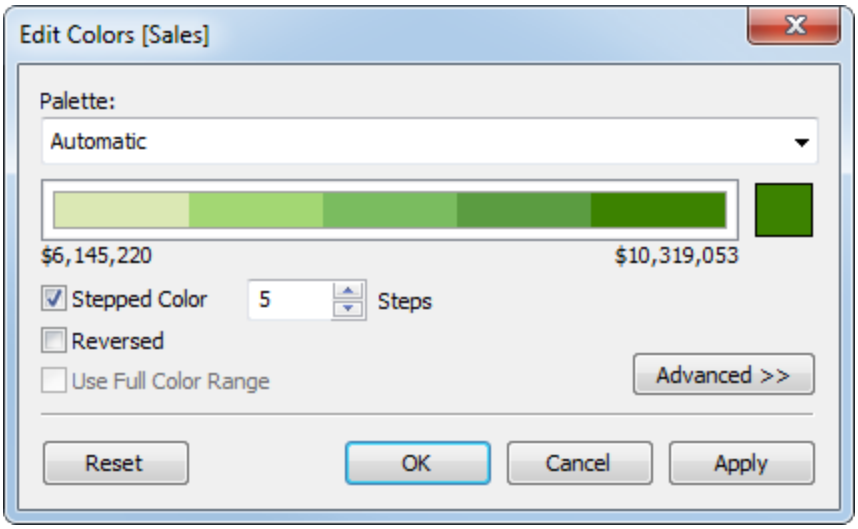

If a diverging color palette is selected, the center point is shown on the color ramp with a small black mark. When the number of steps is odd, the center mark is placed in the middle of the center step. When the number of steps is even, the center mark is placed at the boundary of the center-most two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

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 image below shows a Red-Blue diverging color palette for values from -231 to 55,000. Without using the full color range, -231 (the value for Green Tea) shows as gray. When the full color range is used, -231 is represented by a full red.

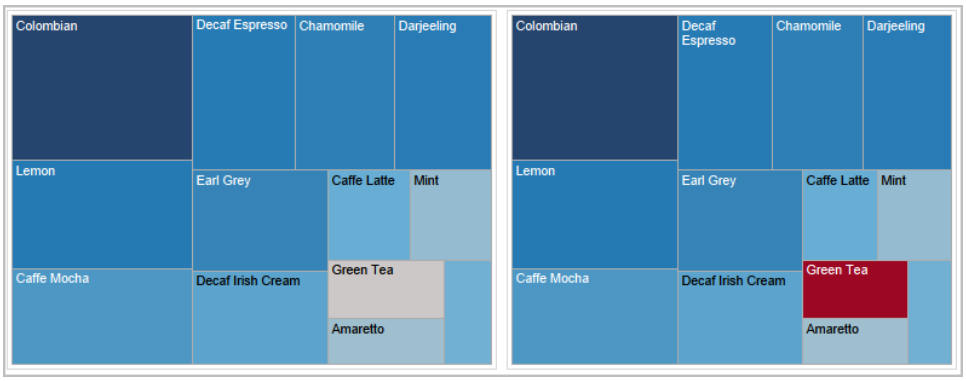

Limiting the Color Range

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 choose to specify the start, end, and center values on the range by selecting the check box and typing a new value into the text box. The Start value is the lower limit in the range, the End value is the upper limit, and the Center value is where the neutral color is located on a diverging color palette.

Resetting the Color Range

To return to automatic color settings, click **Reset** in the Edit Colors dialog box.

Transparency

You can modify the transparency of the marks using the drop-down control for Color on the Marks card. Changing the transparency 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

Click Color on the Marks card to open a drop-down control where you can modify other color properties. You can set up Mark Borders, Mark Halos, and Markers.

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. On the Color drop-down control, select a mark border color.

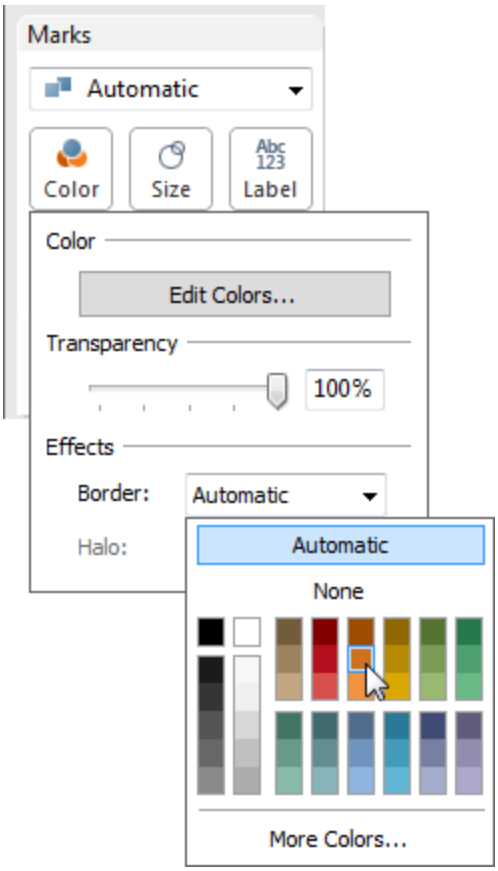

Borders are often useful for distinguishing between closely spaced marks. For example, the view shown below has mark borders turned on (left) and turned off (right). As you can see, when borders are turned on, marks become more easy to distinguish in the areas where they are tightly clustered.

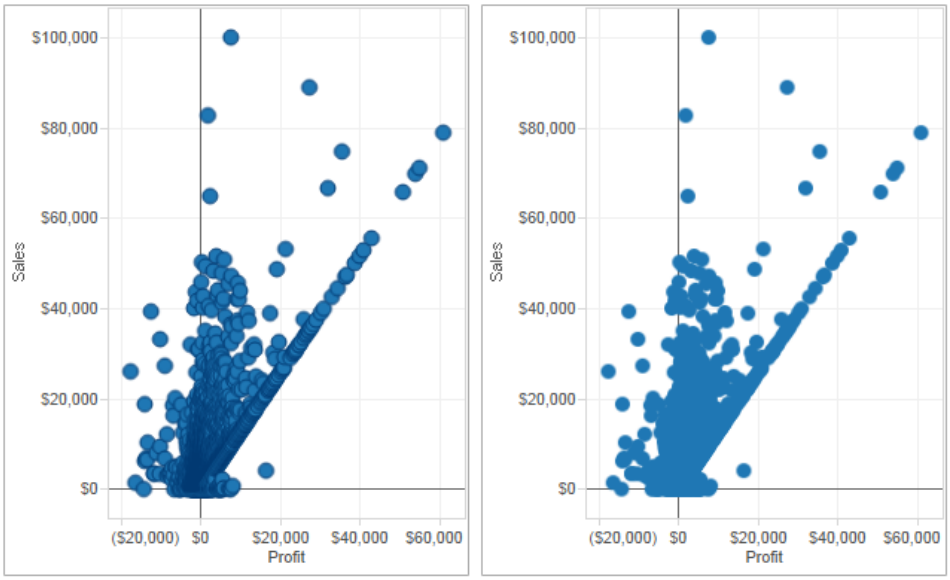

You can also use transparency to show the density of marks.

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. 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.

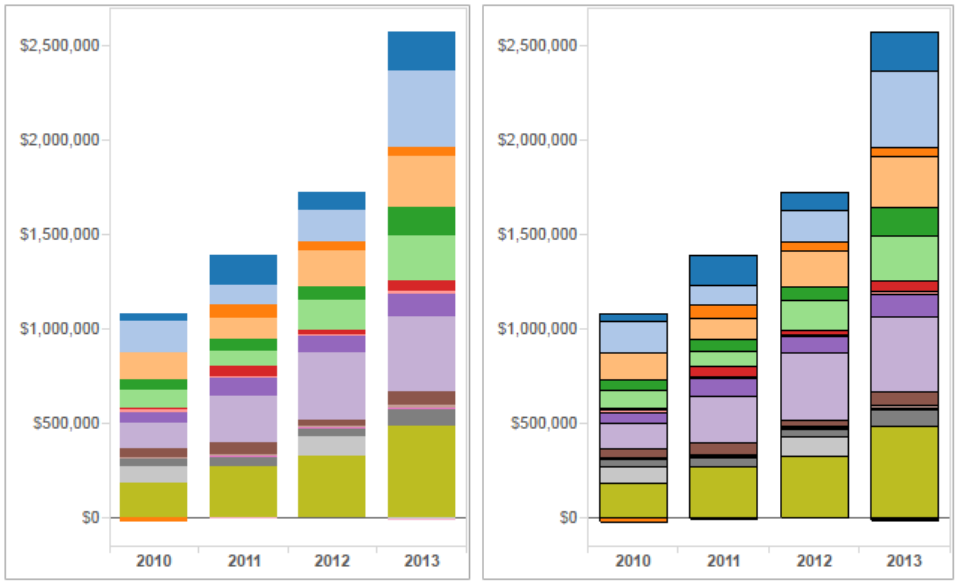

Mark Halos

In order to make the marks in a view more visible against a background image or map, you can surround each mark with a solid contrasting color called a halo. Mark halos are available when you have a background image or a background map. On the Color drop-down control, select a mark halo color.

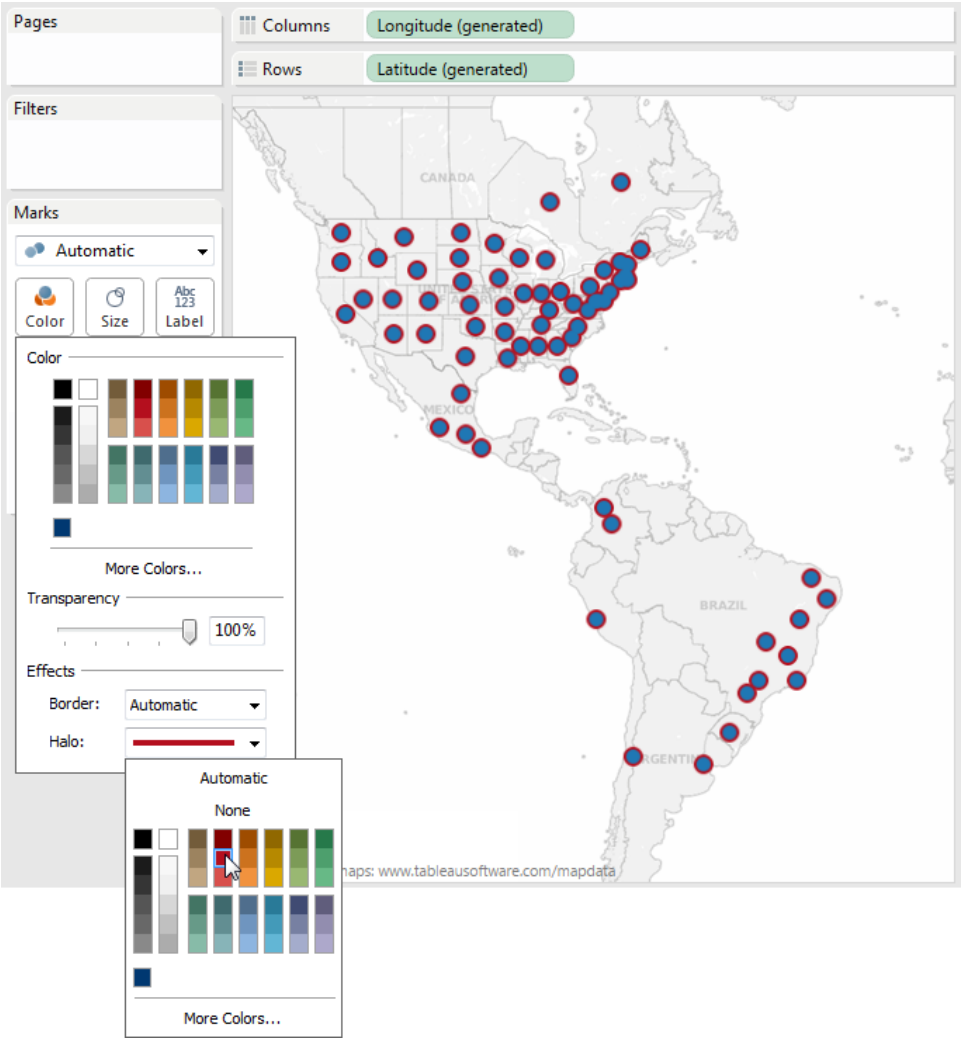

Markers

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. On the Color drop-down control, select one of the Marker effects.

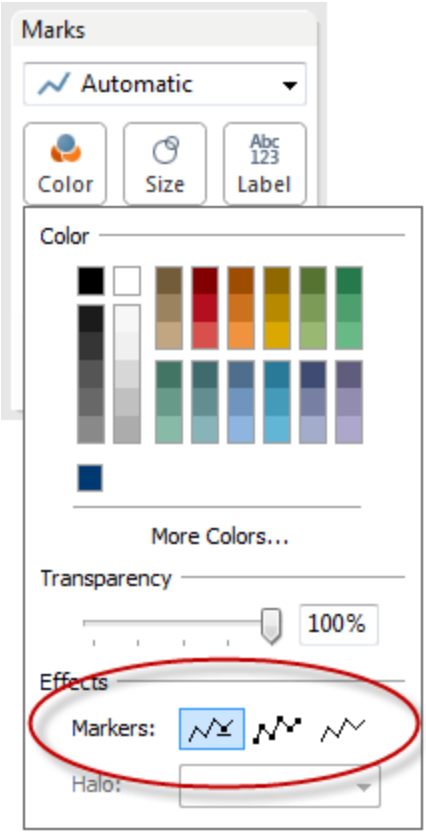

Size

The Size property 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.

Categorical Sizes

When you place a discrete field on the Size target on the Marks card, 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 (small to big), categorical sizes work best for ordered data like years or quarters.

Size-encoding data with a discrete field separates the marks in the same way as the **Detail** on page 271 property does, and then provides additional information (a size) for each mark. When you add categorical size encoding to a view, Tableau displays a legend showing the sizes assigned to each member in the field on the Size target. You can modify how these sizes are distributed using the Edit Sizes dialog box.

1. Double-click the legend or select **Edit Size** from the legend's menu.
2. In the Edit Sizes dialog box, 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.

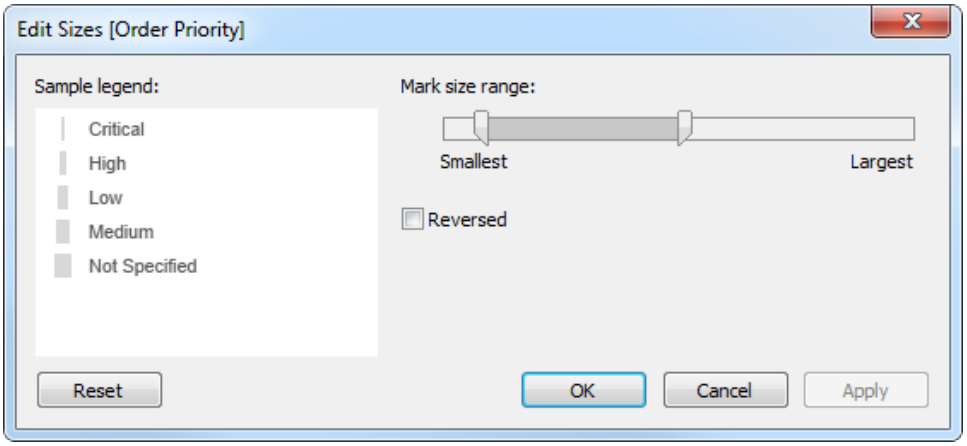

3. When finished, click **OK**.

The view below shows profit and order quantity broken out by region and order date. The order priority is indicated by the size of the mark.

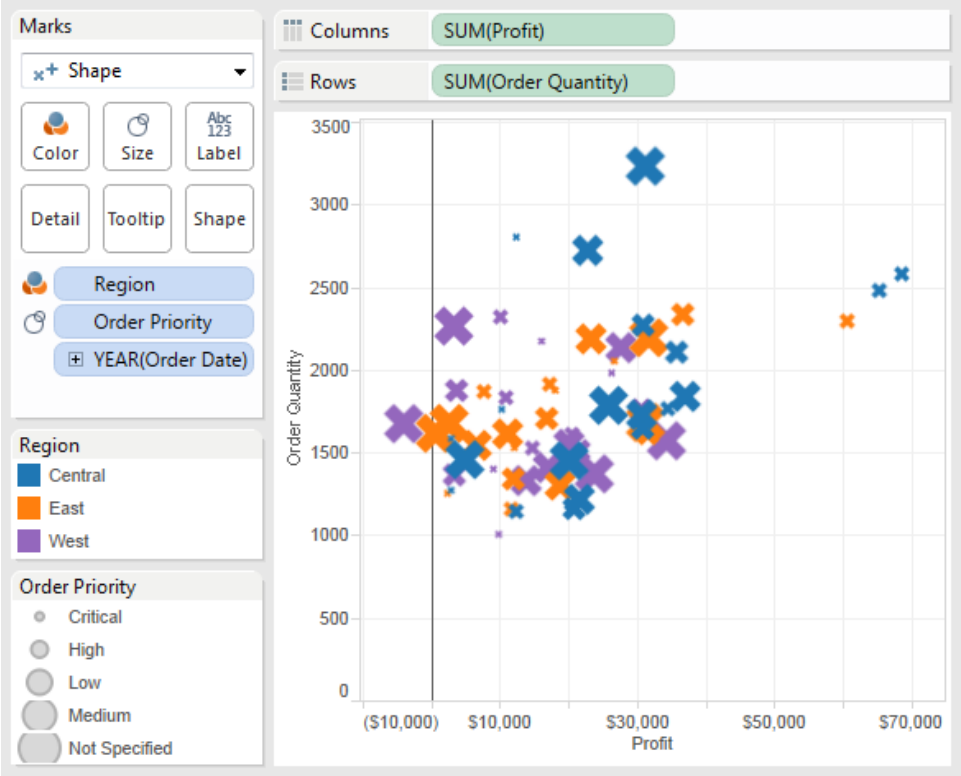

Quantitative Sizes

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

When you add quantitative size encoding to the view, Tableau displays a legend showing the range of values over which sizes are assigned. You can modify how these sizes are distributed using the Edit Sizes dialog box.

1. Double-click on the size legend or select **Edit Size** from the legend's menu to open the Edit Sizes dialog box.
2. 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 (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, assigning the maximum mark size to the absolute value of the data value that is farthest from zero.

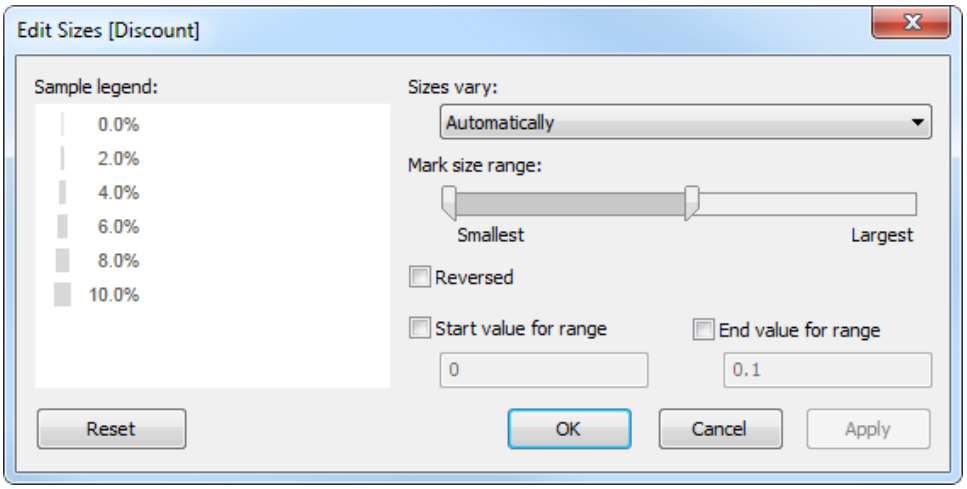

3. Use the range slider to adjust the distribution of sizes. When the **From zero** mapping is selected, the lower slider is disabled because it is always set to zero.
4. Select **Reversed** to assign the largest mark to the smallest value and the smallest mark to the largest value. This option is not available if you are mapping sizes from zero because the smallest mark is always assigned to zero.
5. To modify the distribution of sizes, select the **Start value for range** and **End value for range** checkboxes and type in beginning and end values for the range.
6. When finished, click **OK**.

The view below analyzes the time it takes to ship products based on their ship mode, order date, and the size of order. The size of each mark represents the order quantity while the color represents the ship mode. You can see that most products ship within one or two 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.

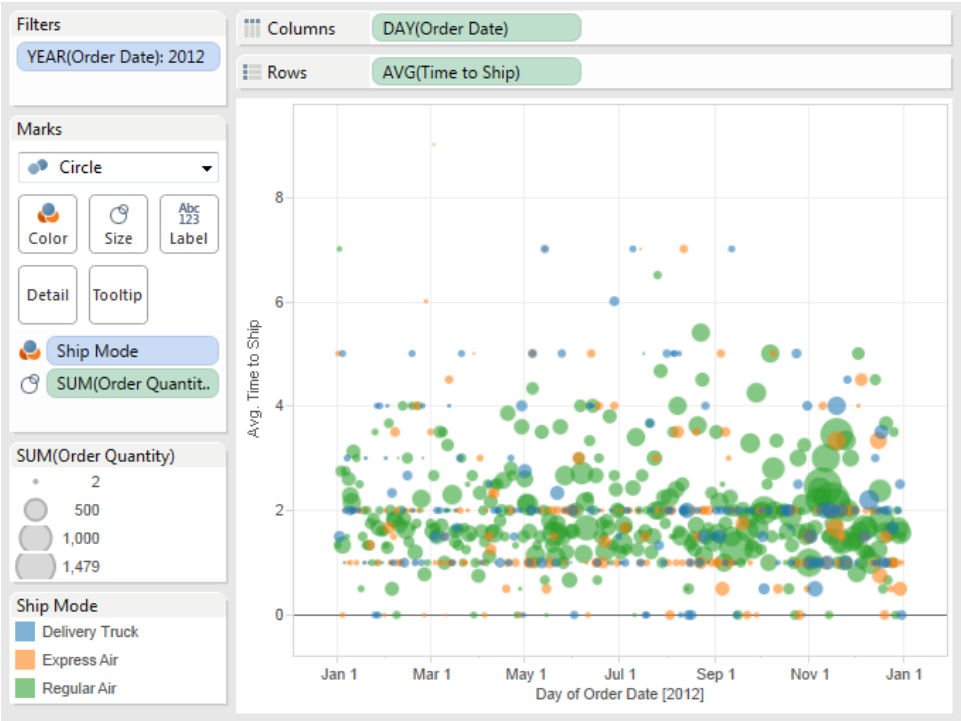

You can also change the size of the marks using the Size slider on the Size drop-down control.

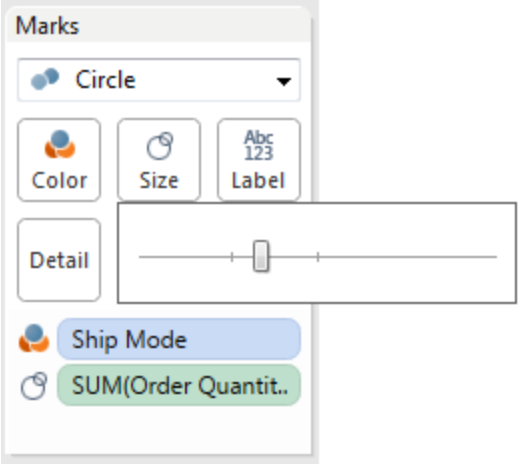

The Size slider affects different marks in different ways, as described in the following table.

Mark Type	Description
Circle, Square, Shape, Text	Makes the mark bigger or smaller.
Bar, Gantt Bar	Makes bars wider or narrower.
Line	Makes lines thicker or thinner.

Mark Type	Description
Polygon	You cannot change the size of a polygon.
Pie	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.

Label/Text

The Label property allows you to encode data by assigning text labels to the marks. When working with a text table, this property is called Text, which allows you to view the numbers associated with a data view. 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 Label or Text target on the Marks card, 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 Label or Text target on the Marks card, 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 (farthest to the right) for both the **Rows** shelf and the **Columns** shelf. This type of view is called a text table, which is also referred to as cross-tab or a PivotTable.

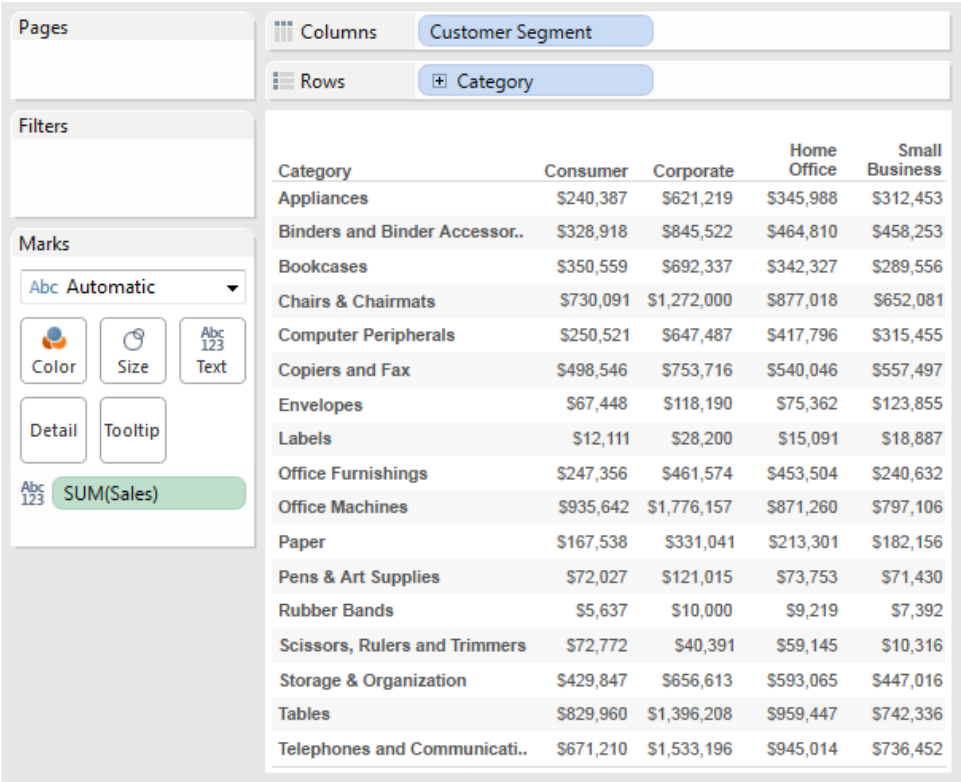

Category	Consumer	Corporate	Home Office	Small Business
Appliances	$240,387	$621,219	$345,988	$312,453
Binders and Binder Accessor..	$328,918	$845,522	$464,810	$458,253
Bookcases	$350,559	$692,337	$342,327	$289,556
Chairs & Chairmats	$730,091	$1,272,000	$877,018	$652,081
Computer Peripherals	$250,521	$647,487	$417,796	$315,455
Copiers and Fax	$498,546	$753,716	$540,046	$557,497
Envelopes	$67,448	$118,190	$75,362	$123,855
Labels	$12,111	$28,200	$15,091	$18,887
Office Furnishings	$247,356	$461,574	$453,504	$240,632
Office Machines	$935,642	$1,776,157	$871,260	$797,106
Paper	$167,538	$331,041	$213,301	$182,156
Pens & Art Supplies	$72,027	$121,015	$73,753	$71,430
Rubber Bands	$5,637	$10,000	$9,219	$7,392
Scissors, Rulers and Trimmers	$72,772	$40,391	$59,145	$10,316
Storage & Organization	$429,847	$656,613	$593,065	$447,016
Tables	$829,960	$1,396,208	$959,447	$742,336
Telephones and Communicati..	$671,210	$1,533,196	$945,014	$736,452

You can display text labels with other mark types by dragging a field to the Label target on the Marks card. See Mark Labels to learn more about showing and hiding mark labels.

In the view below, the heights of the bars are given by the **Sales** measure and the labels show the aggregation (SUM) of the **Profit** measure.

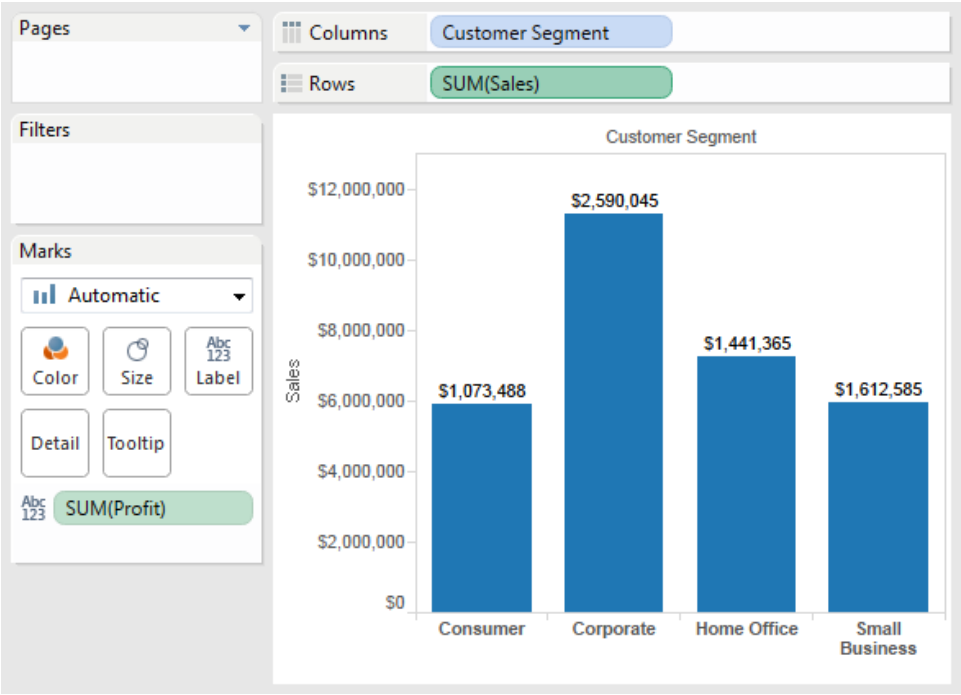

Detail

Whenever you place a dimension on the **Rows** or **Columns** shelf, Tableau uses the categorical members of the dimension to create table headers. The headers show how Tableau is sorting the underlying row data into specific categories. For example, the **Customer Segment** dimension separates the data source rows into four levels of detail: Consumer, Corporate, Home Office, and Small Business.

Drop a dimension on the Detail target on the Marks card to separate the marks in a data view according to the members of a dimension. Unlike dropping a dimension on the **Rows** or **Columns** shelf, dropping it on the Detail target is a way to show more data without changing the table structure.

In the view below, the bars are separated into segments according to the members of the **Category** dimension. The size of each segment reflects the contribution to the profit for a particular member. For example, the Appliances category in the Corporate market has a profit of $71,392.

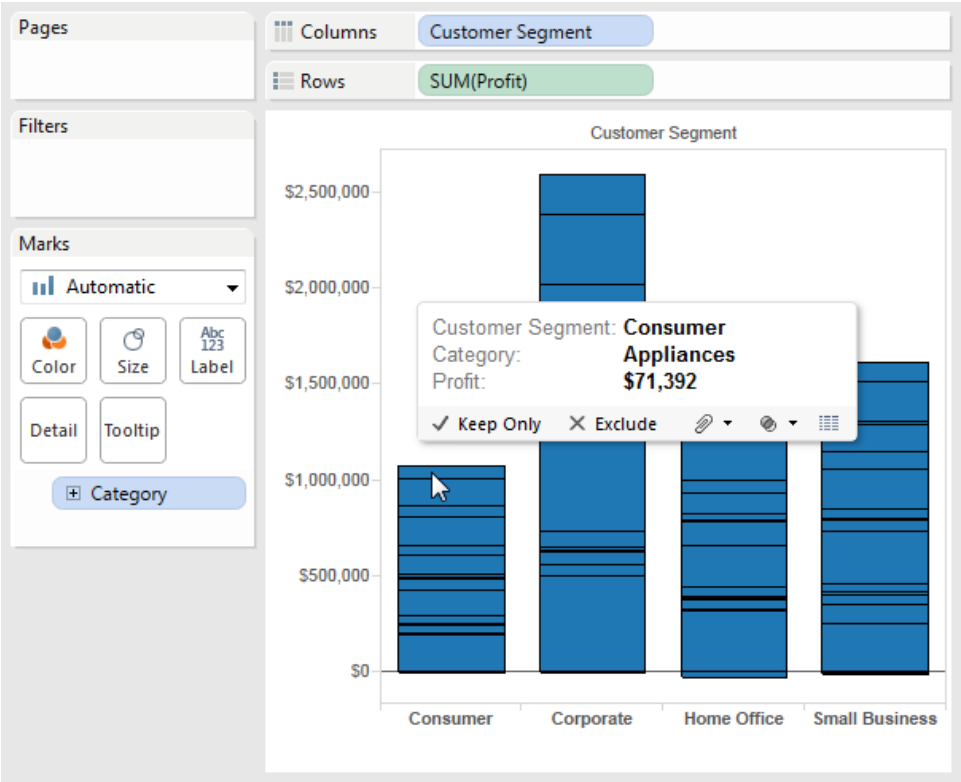

You can place any number of dimensions on the Detail target. In fact, placing all dimensions on this shelf is one way to display all the rows of your data source.

> The Detail property 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.

Placing a measure on Detail 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 Clipboard, or use the values in actions.

Tooltip

Tooltips are data details that are displayed when you hover over one or more marks in the view. Tooltips are also convenient for quickly filtering or removing a selection, or viewing underlying data. You can edit a tooltip to include both static and dynamic text. You can also modify which fields are included in the automatic tooltip.

Add data to the tooltip by dragging it to the Tooltip target on the Marks card.

> Dimensions are added to the tooltip using the ATTR aggregation, which means the tooltip may sometimes display as an asterisk. The asterisk indicates that there are multiple dimension members that apply to the mark you are pointing at. For example, a mark may represent the aggregated sales for all regions. Adding the Region field to the

tooltip results in an asterisk because the mark represents more than one region. To avoid showing an asterisk, add the dimension to the Detail target or use it elsewhere in the view to ensure the marks are at the same level of detail.

Click Tooltip on the Marks card to open a rich text editor where you can add additional text, rearrange the tooltip contents, and insert more fields.

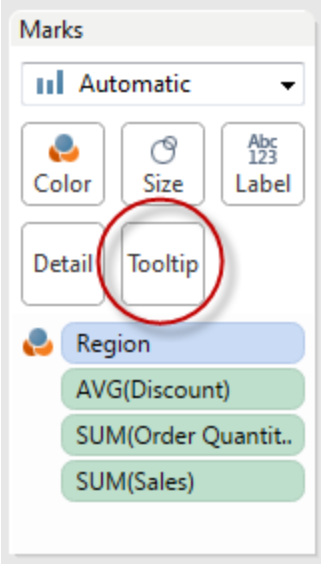

Tooltips are specified on a per-sheet basis and can be formatted using the tools along the top of the Edit Tooltip dialog box. Use the **Insert** menu at the top of the dialog box to add dynamic text such as field values, sheet properties, and more.

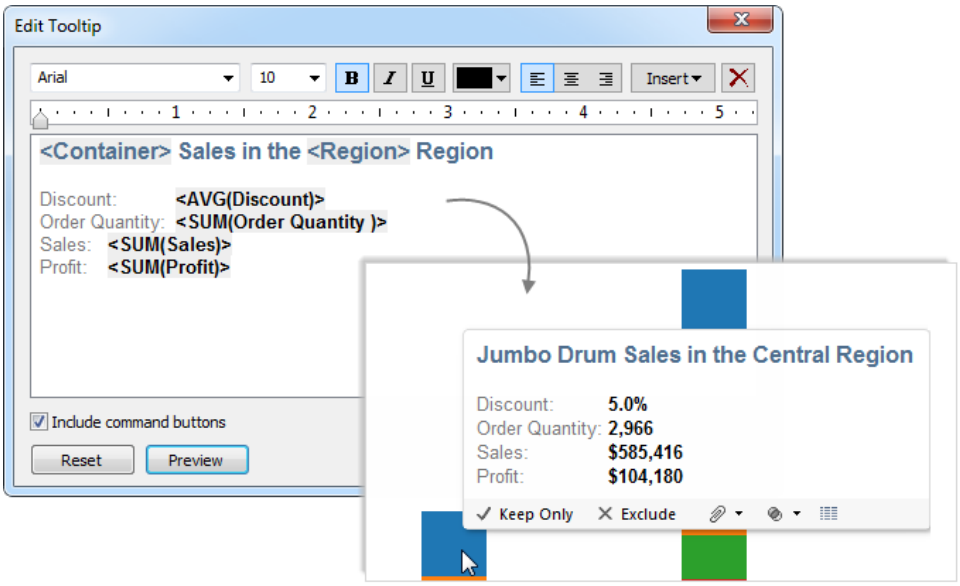

The **All Fields** command on the **Insert** menu adds all field names and values that are used in the view to the tooltip for any mark in the view.

Select the **Include command buttons** option to add **Keep Only**, **Exclude**, and **View Data** buttons to the bottom of the tooltip. These command buttons are available both in Tableau Desktop and when the view is published to the web or viewed on a mobile device.

To add and remove fields when using the automatic tooltip, right-click the field on one of the shelves in the view and select **Include in Tooltip**.

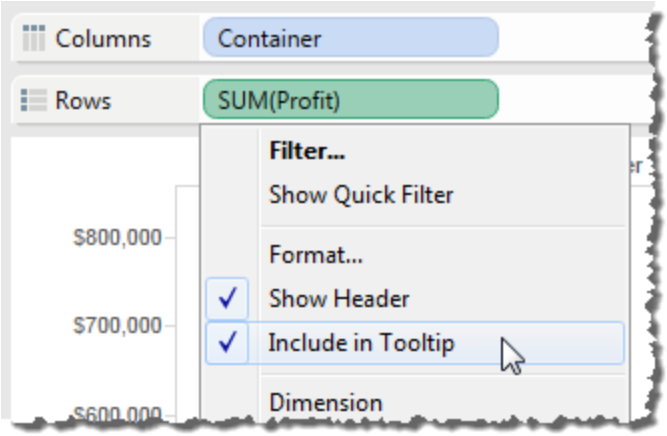

> The **Include in Tooltip** option is only available if you have not customized the tooltip. If you have customized the tooltip, you can return to the automatic tooltip by clicking **Tooltip** on the Marks card and then clicking **Reset** in the Edit Tooltip dialog box.

Shape

The Shape property allows you to encode data by assigning different shapes to the marks in a data view.

When you place a dimension on the Shape target on the Marks card, Tableau separates the marks according to the members in the dimension, and assigns a unique shape to each member. Tableau also displays a shape legend, which shows each member name and its associated shape. When you place a measure on the Shape target, the measure is converted to a discrete measure.

Shape-encoding data separates the marks in the same way as the Detail property does, and then provides additional information (a shape) for each mark. Shape is the default mark type when measures are the inner most (farthest right) fields for both the **Rows** shelf and the **Columns** shelf.

In the view 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 profit and sales.

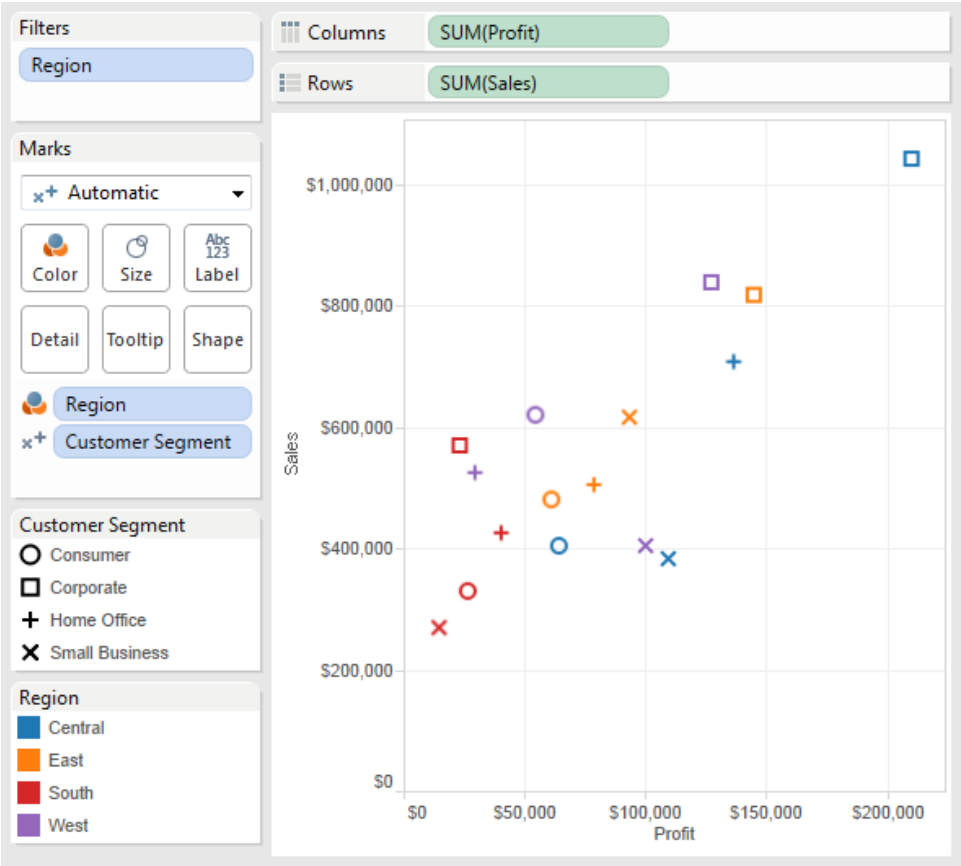

Editing Shapes

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, including filled shapes, arrows, and even weather symbols.

1. Click the Shape target on the Marks card, or select **Edit Shape** on the legend's card menu.
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 **Assign Palette** to quickly assign the

shapes to the members of the field.

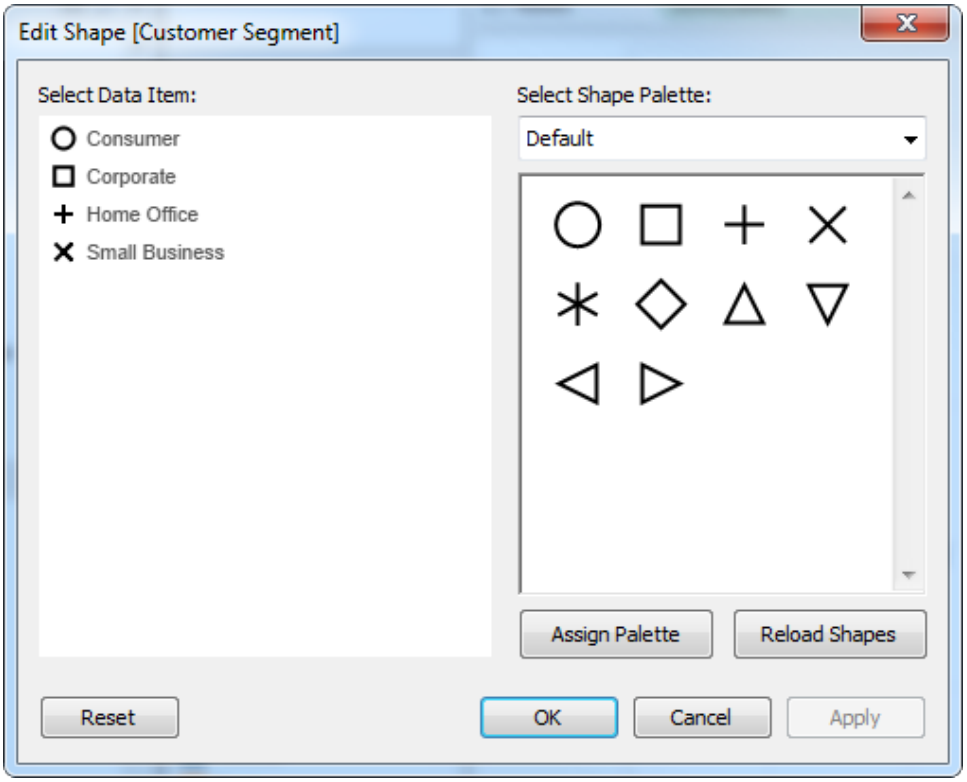

Select a different shape palette using the drop-down menu in the upper right.

> Shape encodings are shared across multiple worksheets that use the same data source. 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. To set the default shape encodings for a field, right-click the field in the Data window and select **Default Properties >Shape**.

Custom Shapes

You can add custom shapes to a workbook by copying shape image files to the Shapes folder in your Tableau Repository, which is 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.

1. Create your shape image files. Each shape should be saved as its own file and can be in any of several image formats including bitmap (.bmp), portable network graphic (.png), JPEG, graphics interchange format (.gif).
2. Copy the shape files to a new folder in the My Tableau Repository\Shapes folder in your Documents 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.

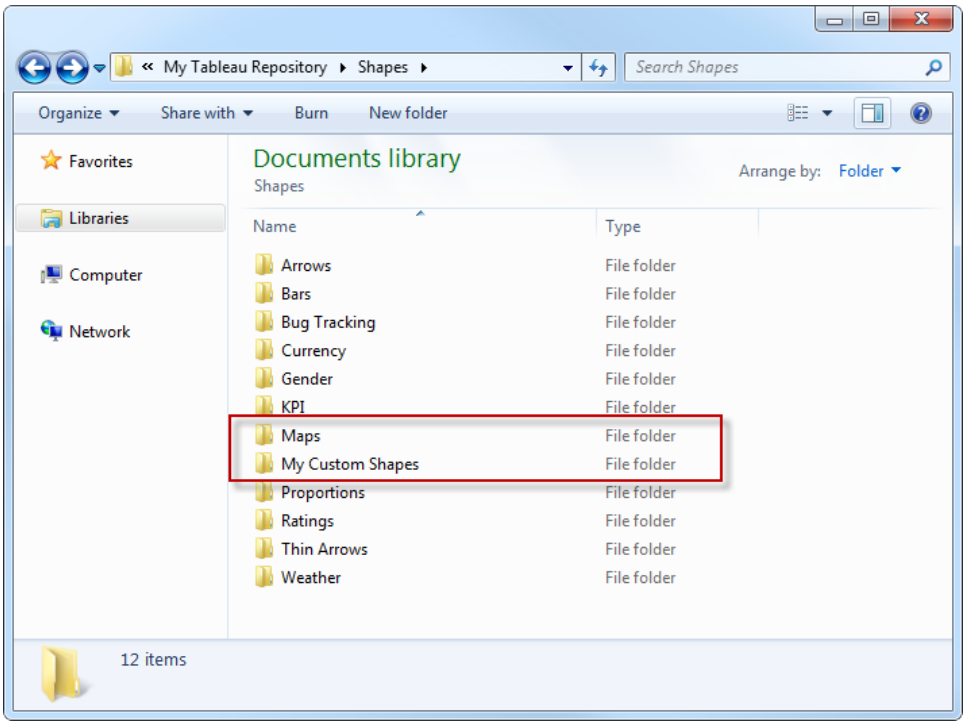

3. In Tableau, click to the right of the dimension name on the shape legend, and choose **Edit Shape**.

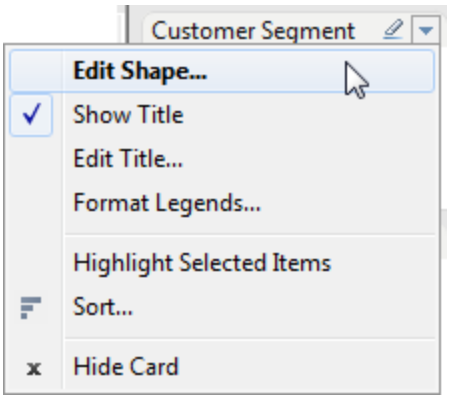

4. Choose the new custom palette in the drop-down list. If you modified the shapes while Tableau was running, you may need to click **Reload Shapes**.

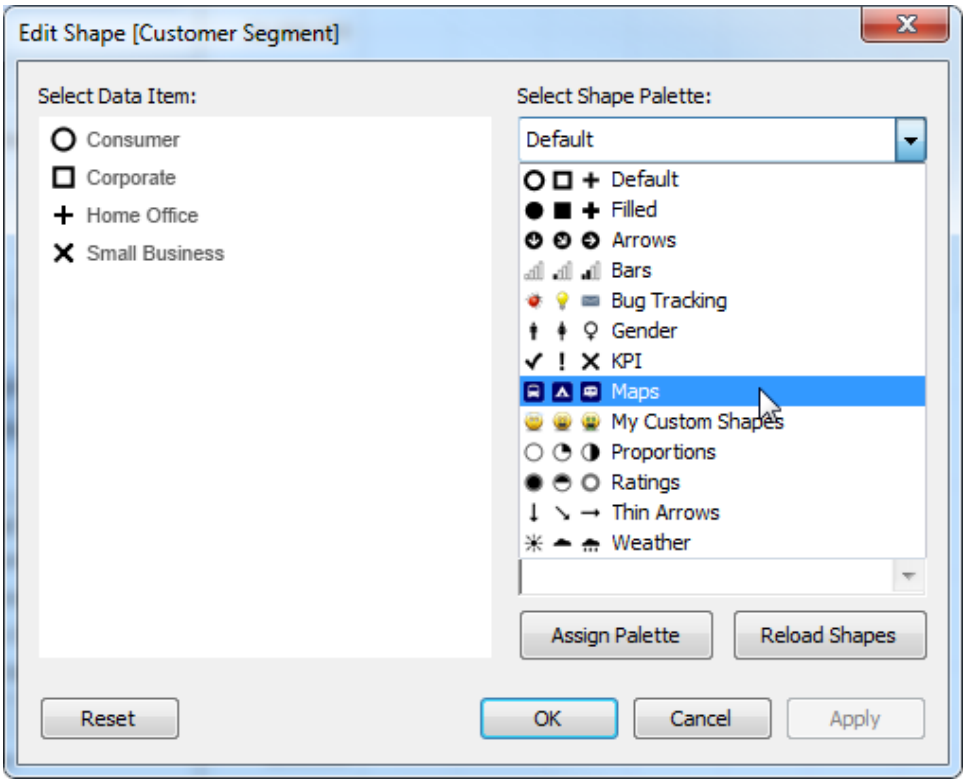

5. You can either assign members shapes one at a time, or click **Assign Palette** to automatically assign the shapes to the members.

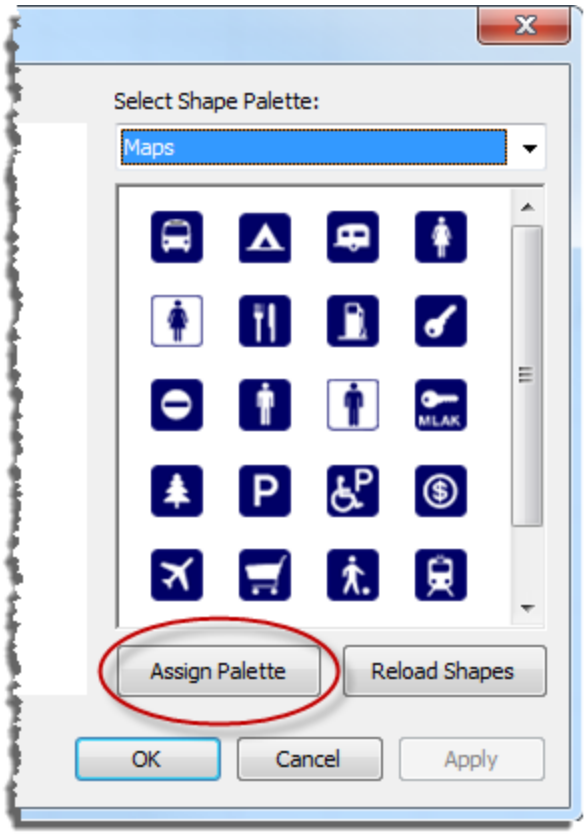

You can return to the default palette by clicking the **Reset** button. If you open a workbook that uses custom shapes that you don't have, the workbook will show the custom shapes because the shapes are saved as part of the workbook. However, you can click **Reload Shapes** in the Edit Shapes dialog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

Below are some examples of views that use both the default and custom shape palettes.

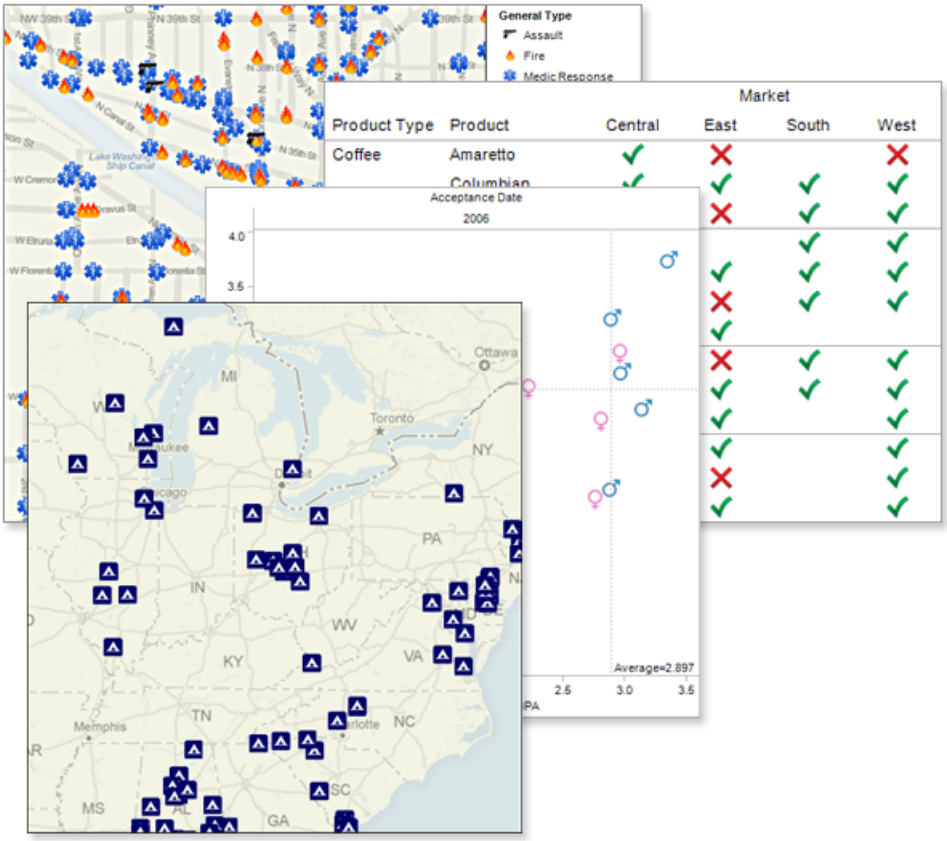

Tips for Creating Custom Shapes

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. 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 target 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 by clicking Size on the Marks card, or by using the cell size options on the **Format** menu.
- **Adding Color Encoding** - If you plan to also use Color to encode shapes, 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 files support alpha channels with a range of transparency levels available on every pixel in the image. When Tableau color encodes a 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.

Path

The Path property 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 target on the Marks card, 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.
- Measure – When you place a measure on the Path target, Tableau connects the marks according to the values of the measure. The measure can be aggregated or disaggregated.

The **Path** property is available only when you select the Line or Polygon mark type from the **Mark** menu. See **Mark Types** on page 239 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. This lets you see the path of the storm. By placing the continuous date on the Path target, you specify that the lines are to be drawn in chronological order.

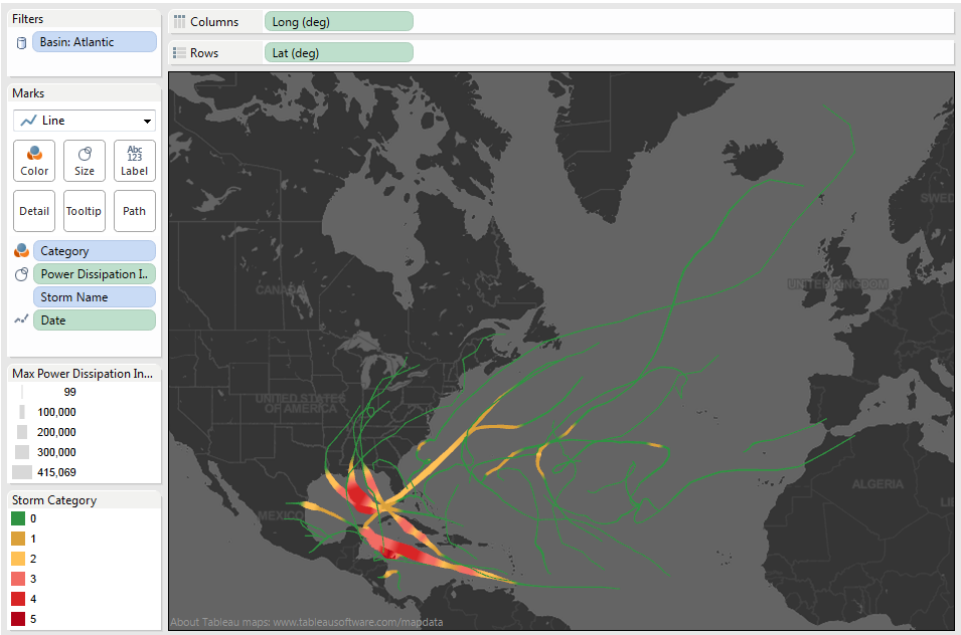

Stacking Marks

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** shelves. 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 setting is **Automatic**. When you are in automatic mode, the Stack Marks menu shows whether stacked marks is on or off.

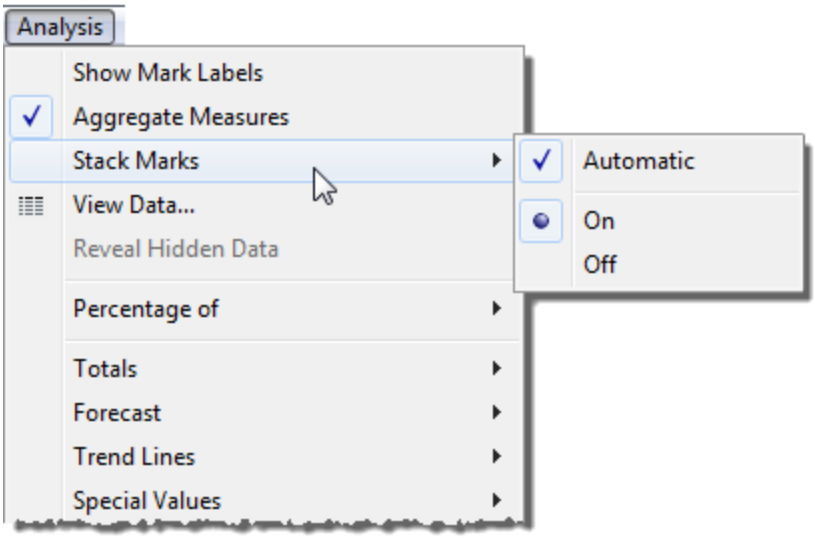

Selecting **On** or **Off** on the **Stack Marks** menu,switches into manual mode. Your selection remains throughout any changes you make to the view.

The following examples illustrate stacking marks.

Example – Stacking Bars

Consider the stacked bars view shown below. It was created by placing a dimension on the **Columns** shelf, placing a measure on the **Rows** shelf, and color-encoding the data by a dimension (that is, dropping a dimension on the Color target on the Marks card).

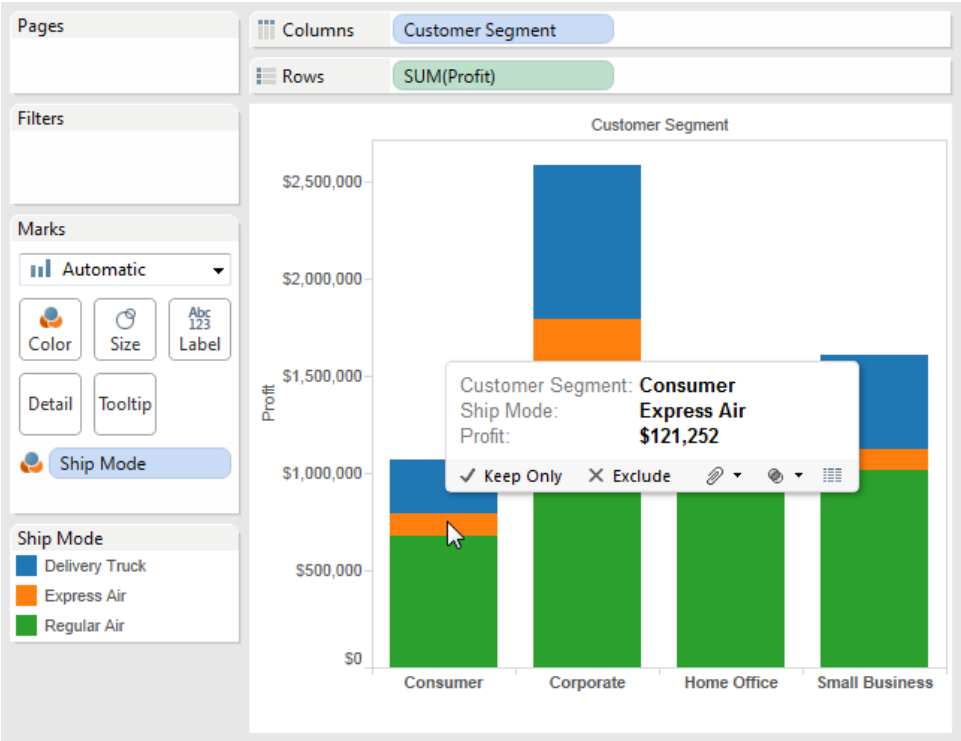

Because the mark type is a bar, Tableau automatically stacks the marks. This means that the marks are drawn cumulatively and the height of each stacked segment within each bar represents the value for that segment. For example, the sum of the profit for products shipped by Express Air (orange bar segment) in the Consumer market is $121,252.

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.

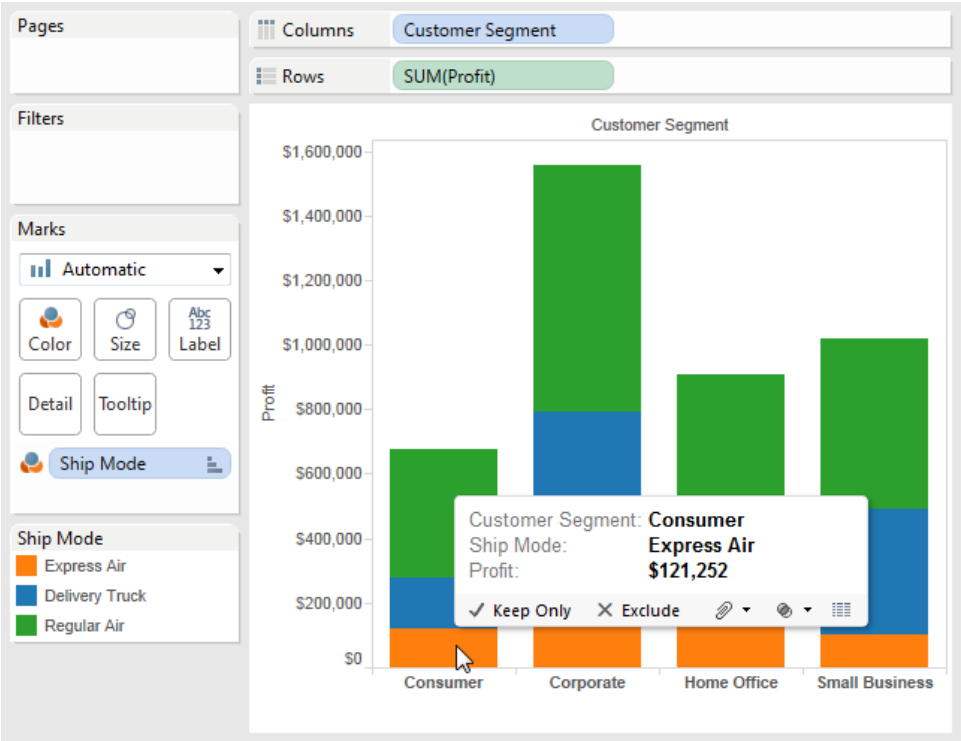

Example – Stacking Lines

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 (that is, dropping a dimension on the Color target on the Marks card). Because the mark type is Line, the marks are not automatically stacked. Instead, they are drawn independently from the horizontal axis.

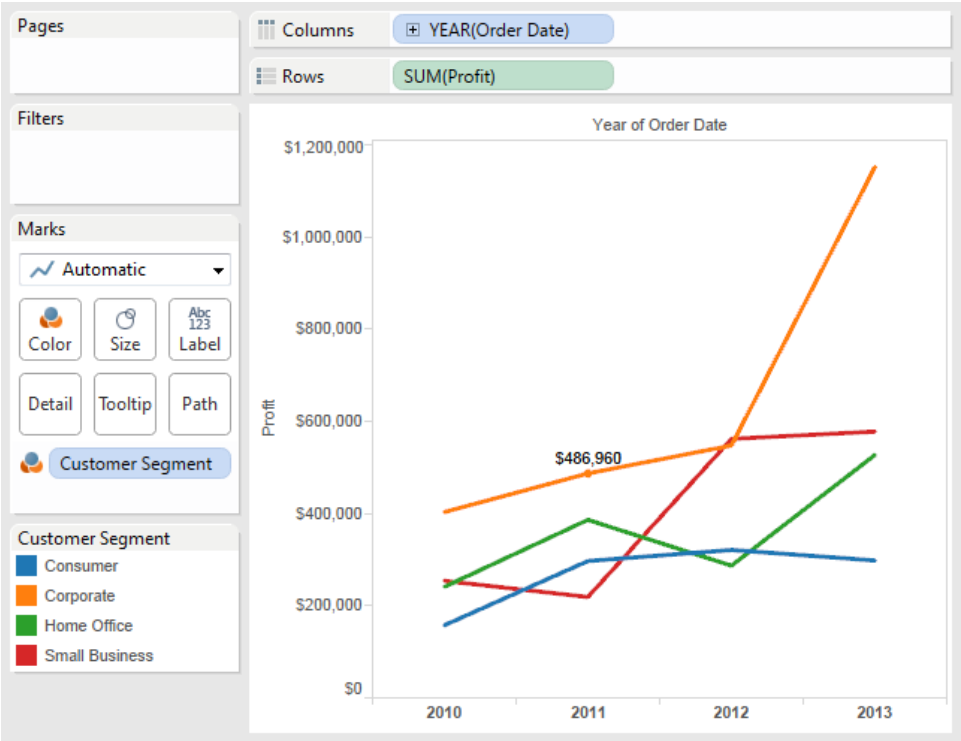

Interpret any data point by reading the associated values from the horizontal and vertical axes. For example, in the year 2010, the Corporate (orange) sales totaled $486,960. That is, the space between that data point and the horizontal axis is equal to the sum of the sales for the Corporate market.

Now, stack the marks by selecting the **Analysis** > **Stack Marks > On** menu item. Tableau automatically switches to the Area mark type.

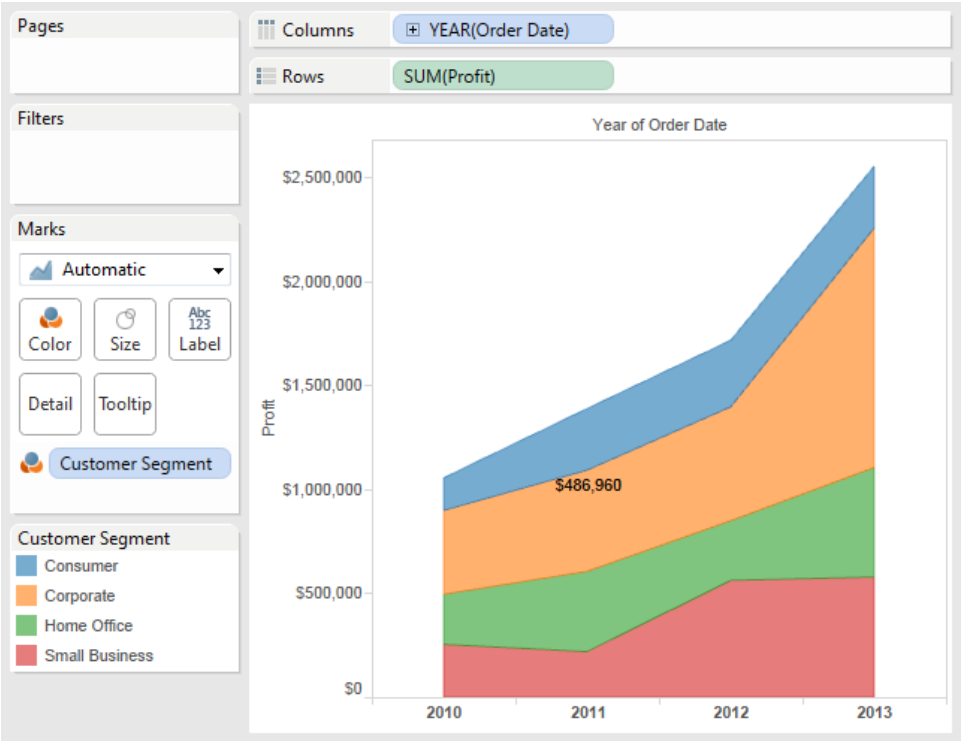

In this view, the lines are no longer independent of each other. Instead, they are drawn cumulatively. The stacking order is given by the order of the dimension members in the data source. This order is reflected in the color legend, from bottom to top.

Therefore, the stacked Small Business (red) area is the same as its un-stacked version because it's at the bottom of the stacking list. The stacked Home Office (green) area is derived by adding its un-stacked values to the un-stacked Small Business values. The stacked Corporate (orange) area is derived by adding its un-stacked values to the stacked Home Office data. The stacked Consumer (blue) area 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 filled area as the sum of the profit.

For example, notice that the label for the 2011 Corporate data still shows the profit as $486.960. 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.

Titles

You can add a title to any worksheet or dashboard. For stories, a title is displayed by default but can be removed.

The title for any kind of sheet is displayed on aTitle card.

Show and Hide Titles

For a worksheet, select **Worksheet** > **Show Title** or click View Cards

on the toolbar and then select the **Title** card.

On a dashboard, select **Dashboard** > **Show Title**.

By default, the title is the name of the sheet. Edit the title to change the text and include dynamic values such as page number and sheet name.

For a story, select **Story** > **Show Title** to add or remove a check mark. Story titles are displayed by default.

Edit Titles

1. Right-click on the title and select **Edit Title** or double-click the title itself.
2. In the Edit Title dialog box, type a new title. Use the Insert menu to add automatic text such as page number, sheet name, parameter values, and so on. Use the formatting

options along the top of the dialog box to change the font, color, style, and alignment.

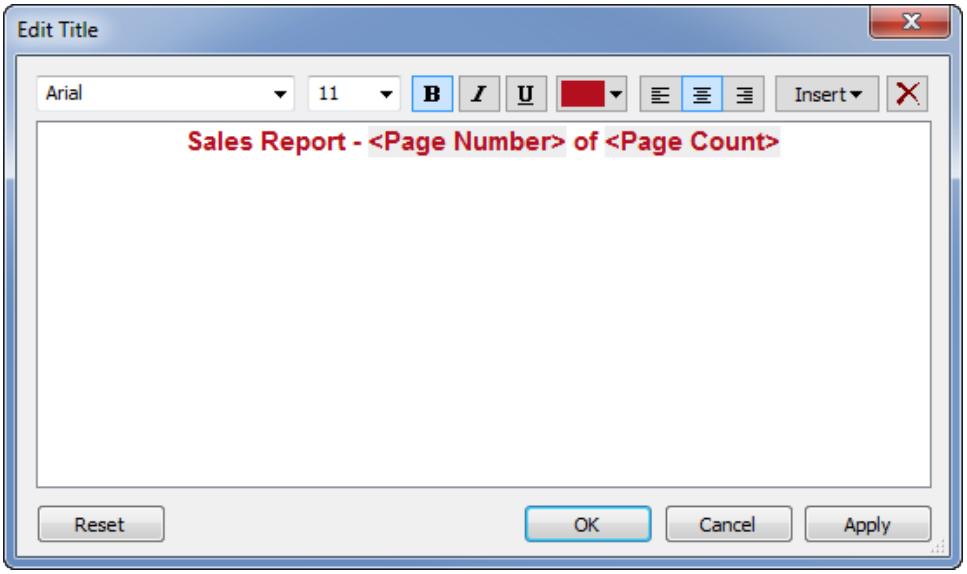

Reset the title back to the default by clicking **Reset** in the Edit Title dialog box.

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 **Worksheet** > **Show Caption**.

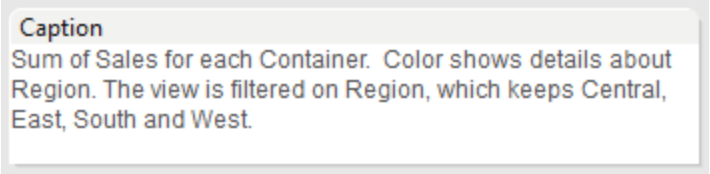

The caption is automatically generated by default, however, you can edit the caption by double clicking the Caption card. In Edit Caption dialog box, you can use change the font, size, color, and alignment and style.

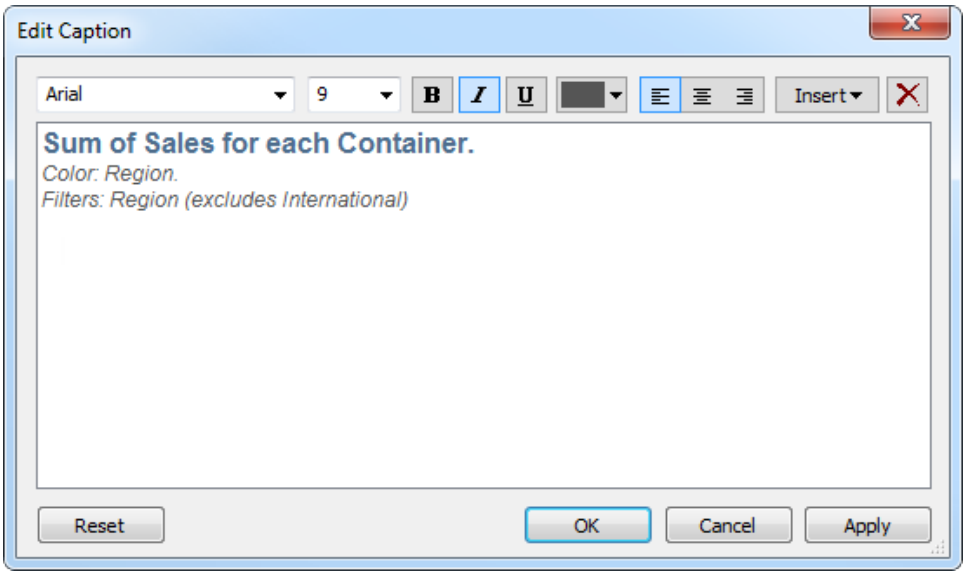

Use the Insert menu to add automatic text such as page number, sheet name, and field and parameter values.

The caption can optionally be included when printing, printing to PDF, and publishing to Tableau Server. When you export the view as an image to another application like Microsoft PowerPoint, you can optionally include the caption.

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 **Department**, thus indicating that the discrete department names are members of the **Department** field.

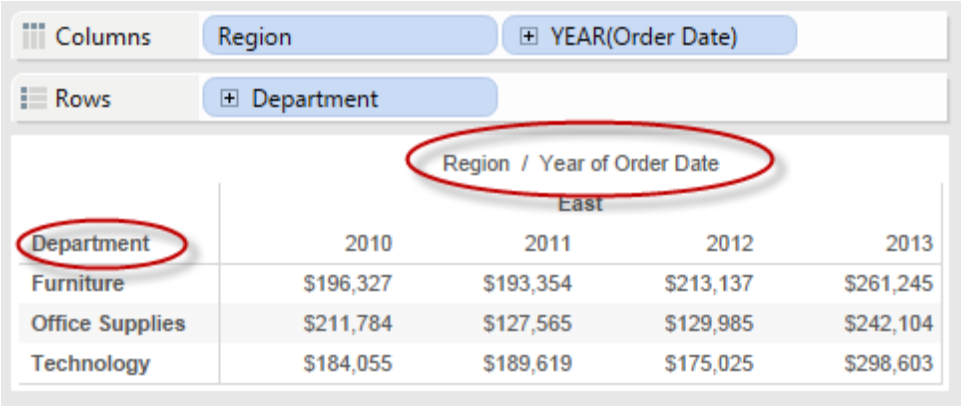

Field labels apply only to discrete fields. When you add continuous fields to the view, Tableau creates an axis. The axis is labeled with a header.

By default, field labels are shown. You can hide or show field labels at anytime by selecting **Analysis > Table Layout > Show Field Labels for Rows** or **Show Field Labels for Columns**.

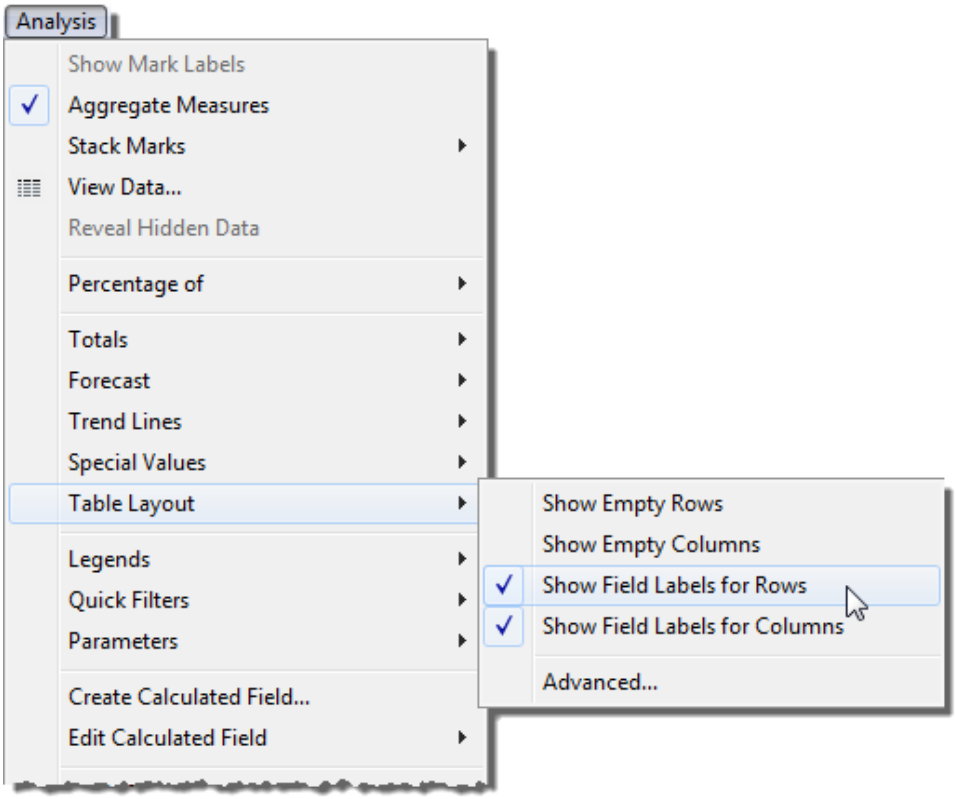

You can format the fonts, alignment, shading, and separators for field labels.

Legends

When you add fields to Color, Size, and Shape on the Marks card, a legend displays to indicate how the view is encoded with relation to your data.

Color Legend

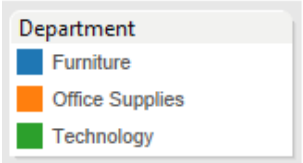

Size Legend

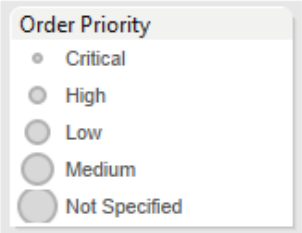

Shape Legend

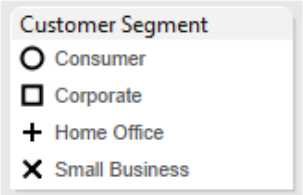

Not only do legends help you understand encodings, you can also use legends to sort, filter, and highlight specific sets of data.

Building Views Manually

Building views in Tableau can be really easy if you understand some basic concepts of how it all works.

To try some hands-on exercises for building views, see **Build-It-Yourself Exercises** on page 484.

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 and cards that are part of every worksheet.

The Basics

When you begin creating a new data view on a blank worksheet, you can drag a field from the Data window and drop it directly into the view, rather than to a shelf.

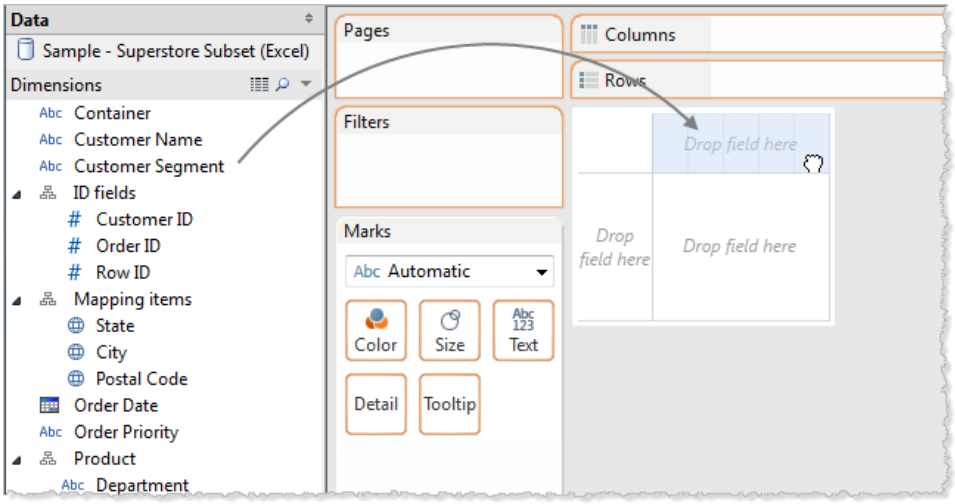

While dragging fields you can hover over the different areas in the view to see how the field will be incorporated into the structure of the view. For example,dimensions typically add a row and column headers to the view, while measures add continuous axes. Below are some examples of how fields can be added to the view.

Dimensions add headers

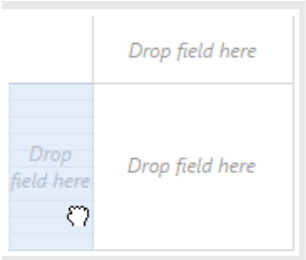

Measures add axes

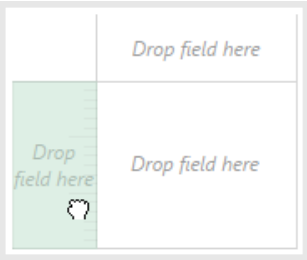

Show Me automatically adds the field according to best practices

For a more advanced discussion of dimensions and measures, see **Data Roles: Dimension vs. Measure** on page 222.

When you drop a field on one of the active areas in the view, the field is added to the view and also to one of the shelves or cards. For example, in the view below the **Customer Segment** dimension was added to the **Rows** shelf and the **Profit** measure was added to the **Columns** shelf—and automatically aggregated so that the values in the data are summed.

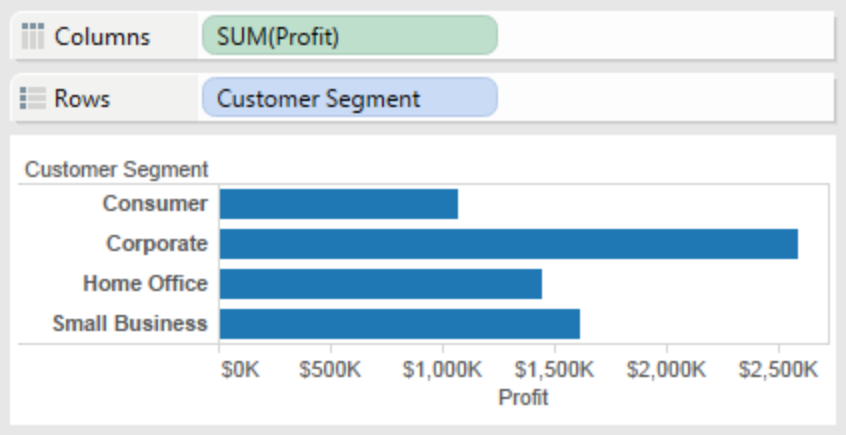

You can of course also drag fields directly to the shelves instead of to the active areas in the view. You can also drag fields from one shelf to another shelf. To remove a field from a shelf, drag it off the worksheet or choose **Remove** from the field's menu (available when you right-click a field in the view. To quickly remove multiple fields from a shelf, right-click the shelf and choose **Clear Shelf**.

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 options available. For example you can replace fields by dropping them on existing headers and axes in the view. Or, instead of replacing a field you can add measures to an existing axis. Finally, you can rearrange the rows and columns in the view.

Adding Headers Using Dimensions

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

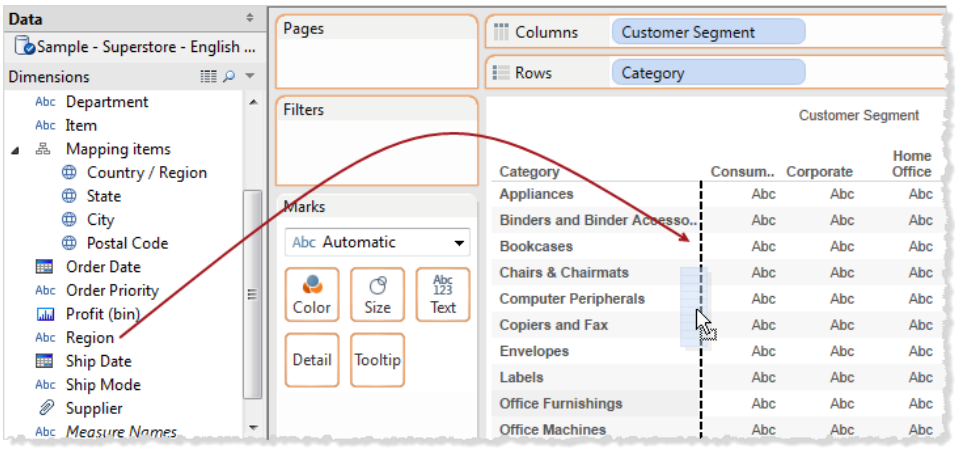

As you hover over the view, a dotted black line indicates active areas where you can add headers.

If you drop Region, it adds a column for Region along the top of the table, and slices the data so that you see results for each region within each category.

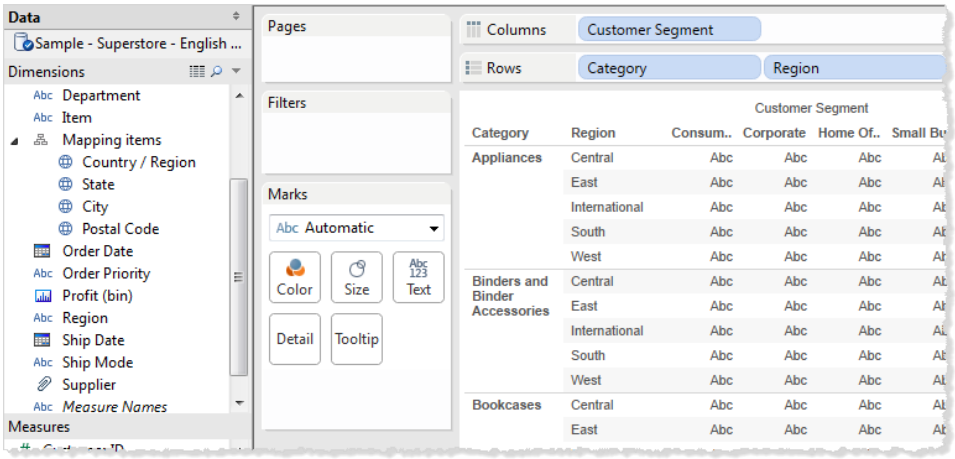

Adding Axes Using Measures

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. In this case you are replacing one measure (**Sales**) with another (**Profit**).

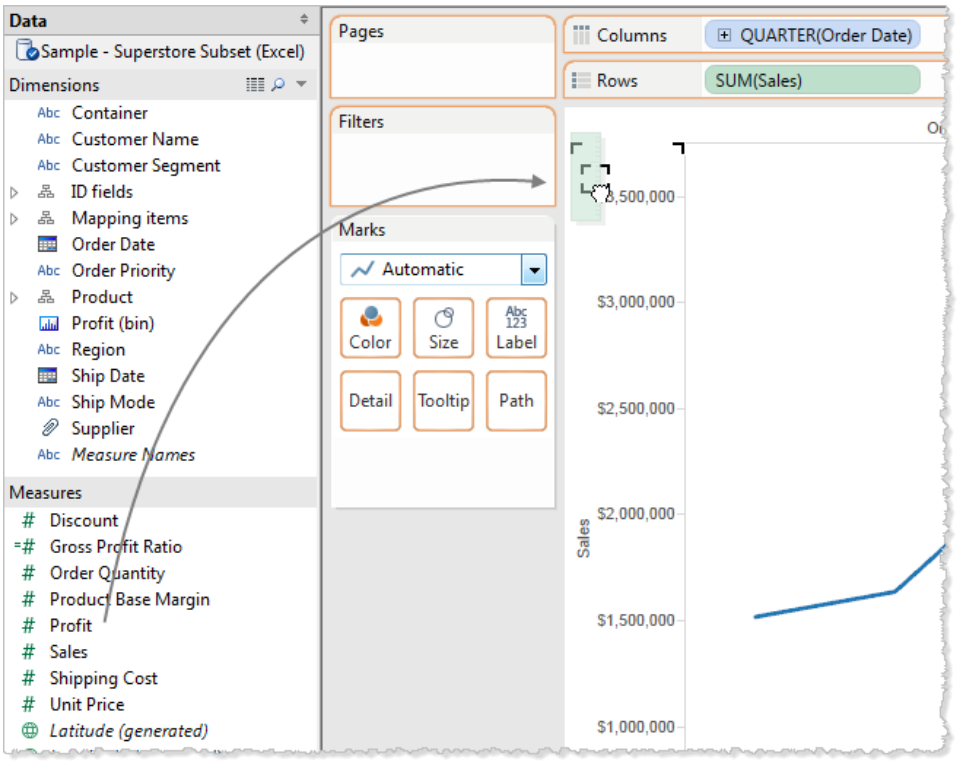

Blend the Measures on Single Axis

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. See **Measure Values and Measure Names** on page 199. In the image below you are adding a second measure (**Profit**) to the existing measure (**Sales**) on the axis.

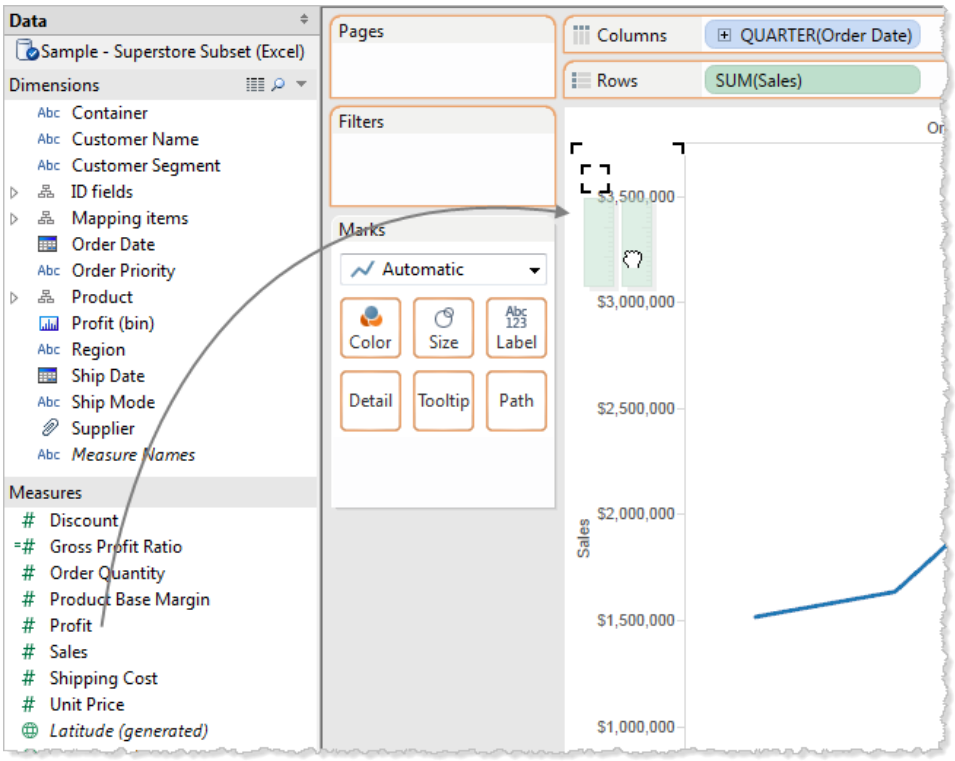

The result is a view with both measures plotted along a single axis. In this case, Sales is shown in gray and Profit is shown in purple.

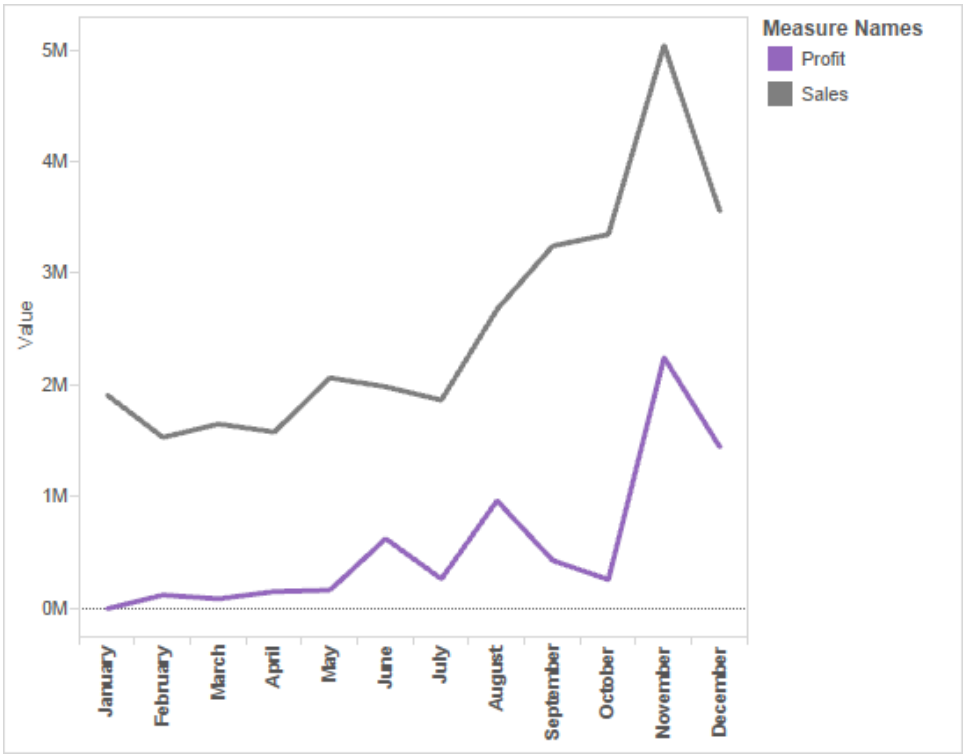

Add a Dual Axis

Drag the field to the right side of the view to add the measure as a dual axis. Dual Axes are useful when you want to compare two fields that have different scales. In this case, blending the two axes would distort the view. Instead you can add a dual axis.

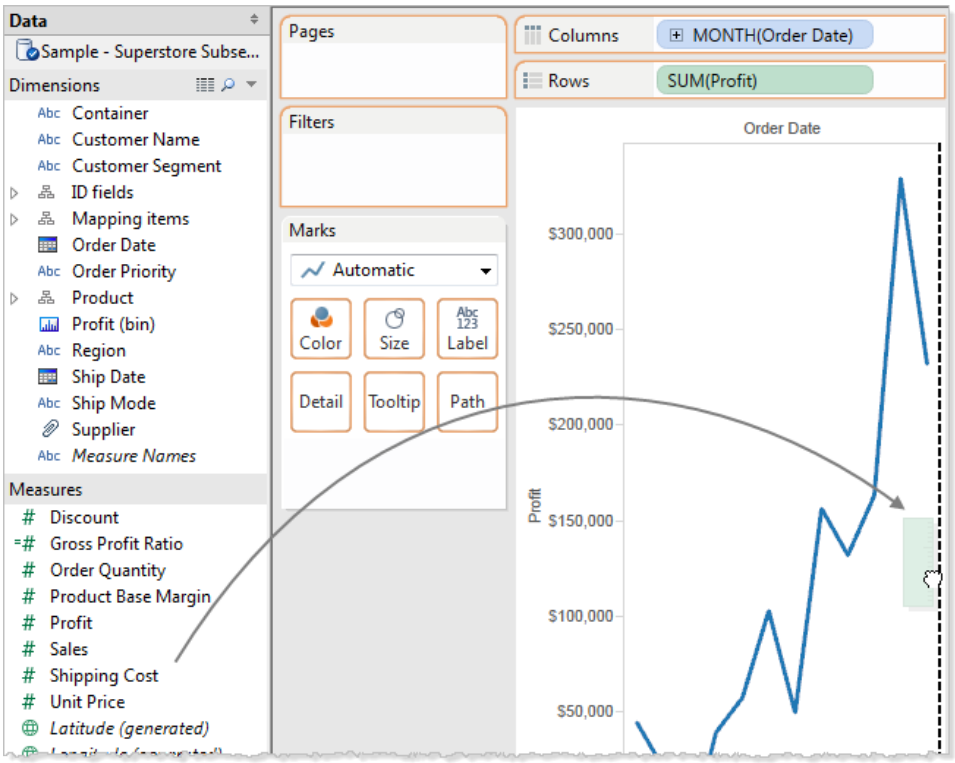

The result is a dual axis view where the Profit axis corresponds to the purple line and the Shipping Cost axis corresponds to the brown line.

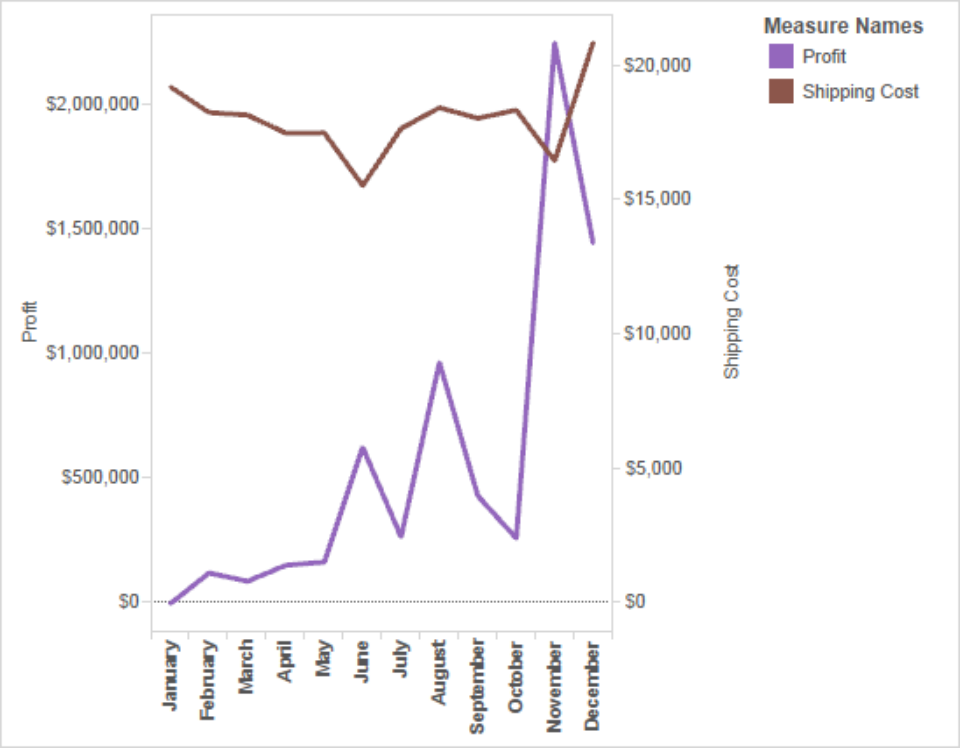

You can add up to four axes to the view: two on the Columns shelf and two on the Rows shelf. You can also turn a measure into a dual axis by selecting **Dual Axis** on the field menu after it has been added to a shelf.

Rearrange 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.

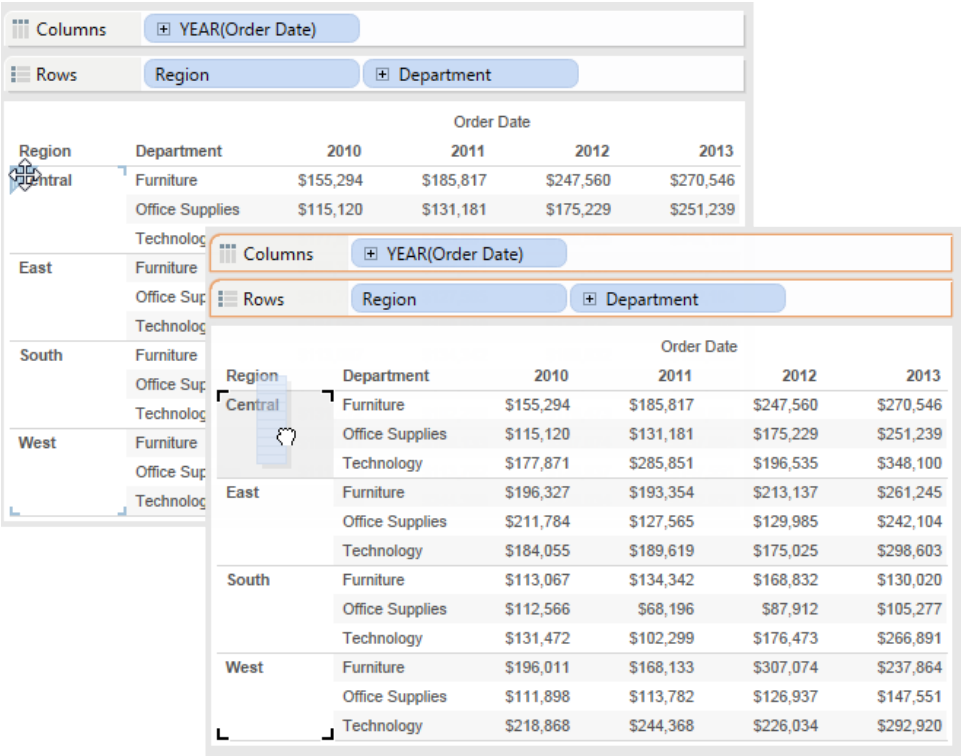

Shelves and Cards

Every worksheet in Tableau contains shelves and cards. By placing fields on shelves or cards, you can create the rows and columns of a data view, exclude data from the view, create pages, and control mark properties. You should experiment placing fields on different shelves and cards to find the optimal way to look at your data. Tableau can also help you determine the best way to display your data using Show Me! See **Building Views Automatically** on page 322.

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 any number of fields on these shelves.

When you place a dimension on the **Rows** or **Columns** shelves, 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 add more fields to the view, 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 axis.

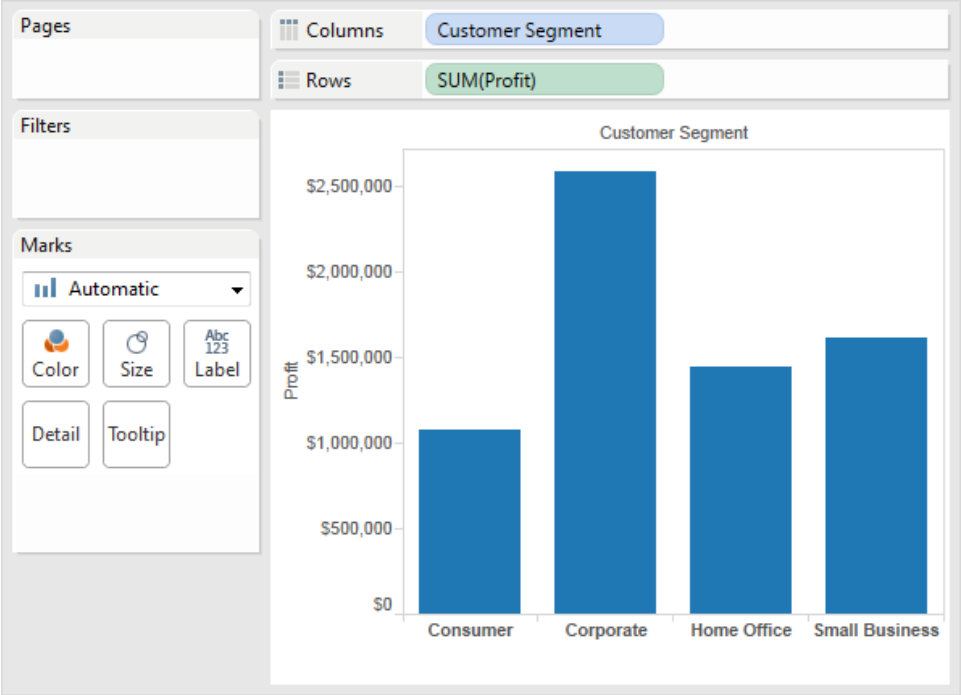

Tableau displays data using marks, where every mark corresponds to a row (or a group of rows) in your data source. The inner fields (that is, the fields farthest to the right) on the **Rows** and **Columns** shelves determine the default mark type. For example, if the inner fields are a measure and a dimension, the default mark type is a bar. You can manually select a different mark type using the Mark menu. See **Mark Types** on page 239 for more information.

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

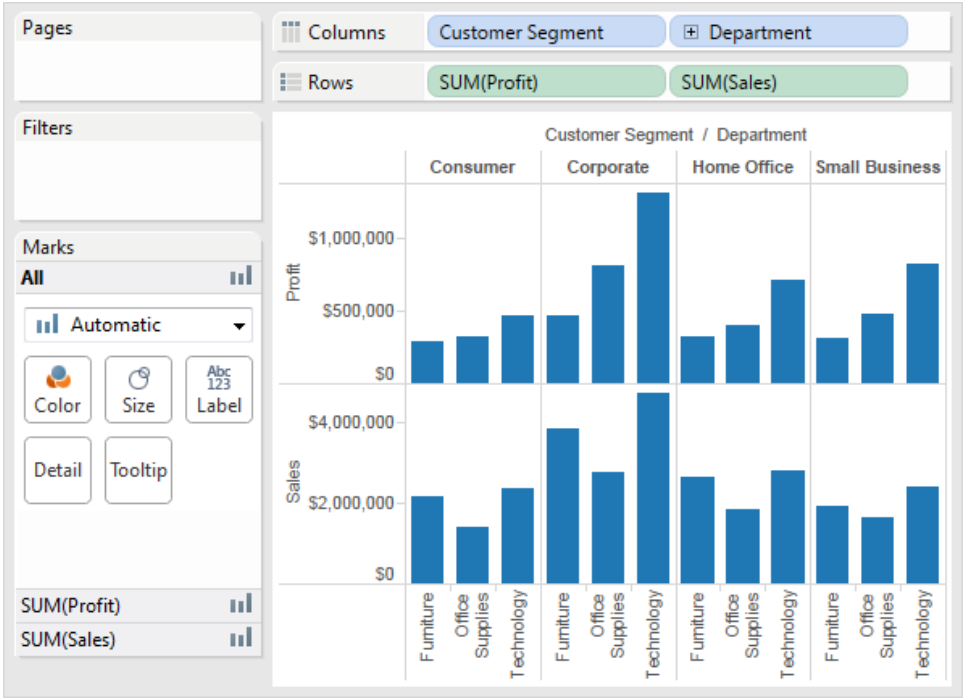

Hide Rows and Columns

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. Instead of filtering the data, you can hide the row or column so it doesn't display in the view but it is still included in calculations.

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. For example, when calculating year-over-year growth, the first year doesn't have a previous year to compare to, so the column is left blank. Filtering the first year will remove it from the view but it will also remove it from the calculation so now the second year doesn't have a previous year to compare to and is left blank. Instead of filtering, you can hide the column 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**.

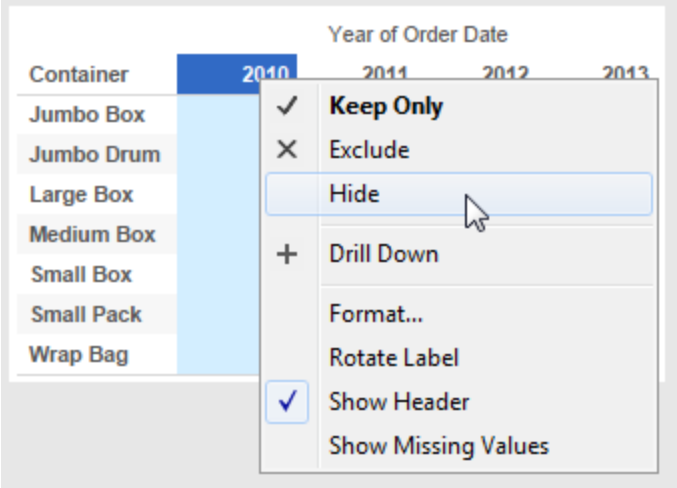

To show hidden data:

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

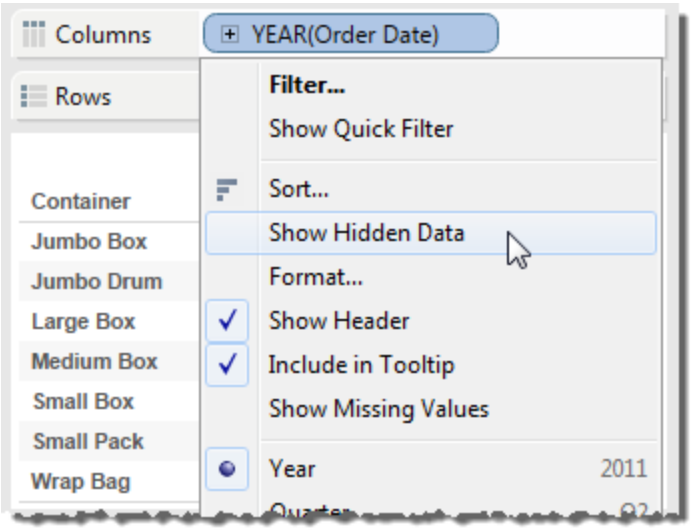

Pages Shelf

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 **Pages** shelf you are adding a new row for each member in the dimension. When you place a measure on the **Pages** shelf, Tableau automatically converts the measure into a discrete measure.

The **Pages** shelf creates a set of pages, with a different view on each page. Each view is based on a member of the field you placed on the **Pages** shelf. You can easily flip through the views and compare them on a common axis, using the controls that get added to the view when you move a field to the **Pages** shelf. For example, the view below shows the **Profit** vs. **Sales** by **Region** for each day throughout the month. The image below shows days 1, 2, and 3. You would have to scroll down to see other days in the month.

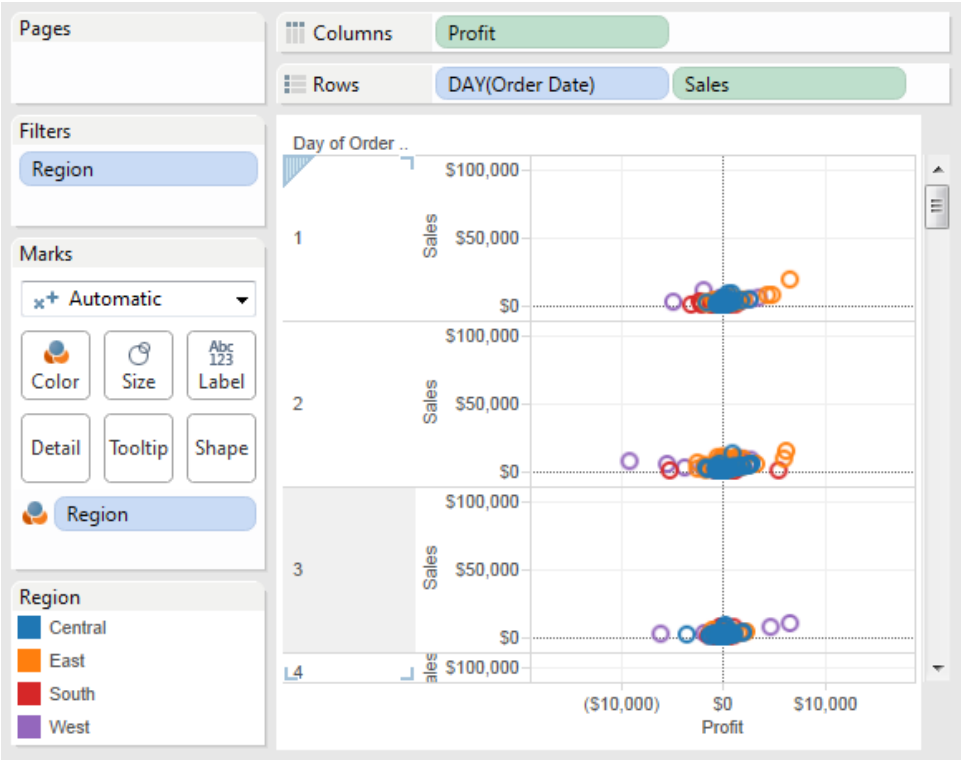

To make this view more user-friendly, move **DAY(Order Date)** to the **Pages** shelf and use the associated control to flip through the pages (one for each day). You can quickly discover hidden insights. In this example, it is interesting that the 19th is an especially big day in terms of sales and profit in the Western region.

Discover hidden insights and outliers.

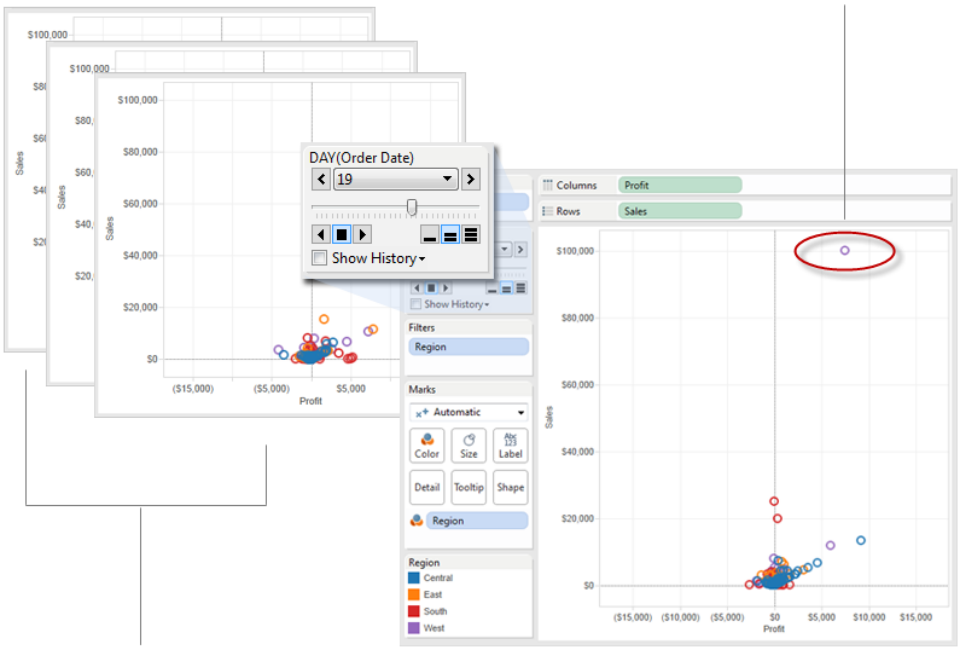

The data for each day is displayed as a separate page.

When you add a field to the **Pages** shelf, a page control is automatically added, just below the **Pages** shelf:

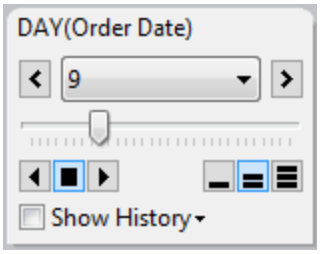

Use this control to navigate through the pages. There are three ways to navigate through the pages in a view:

- Jump to a specific page
- Manually advance through the pages
- Automatically advance through pages

Jump to a specific page

Select the member or value you want to view from the drop-down list to display a specific page.

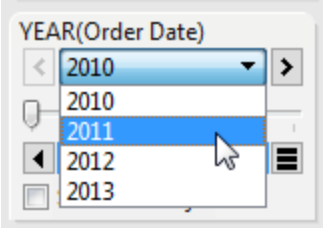

Manually Advance through the pages

You can manually advance through the sequence of pages by doing any 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 any of the following keyboard shortcuts 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 or backward, and stop the playback at any time. You can control the speed of playback with the speed controls in the bottom right corner of the control. The smallest bar indicates the slowest playback speed.

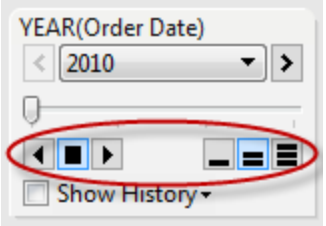

Page History

Show page history using the Show History checkbox. With page history, marks from previous pages are shown on the current page. Open the drop-down control for history to specify what marks to show and when to show them.

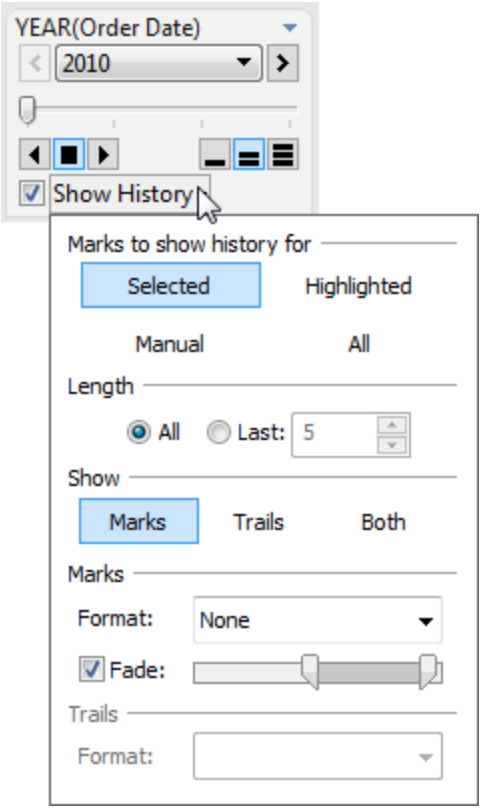

The history drop-down control has the following options:

- **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. To manually show history for marks, right-click the mark in the view and select 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.

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** shelf and on the Color target. Also, trails are only supported for discrete mark types such as squares, circles, or shapes. They are not supported when the mark type is Automatic.

Pages on Dashboards

When a dashboard contains multiple views that use the same field on the **Pages** shelf, you can control all of the views with a single page control by selecting the **Synchronized** option. This option is only available on the page control shown on a dashboard.

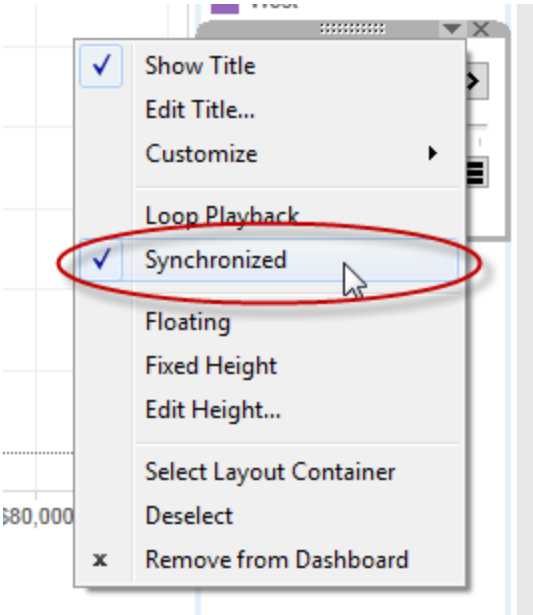

Filters Shelf

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.

> This section presents a brief overview of filtering. Refer to **Filtering** on page 334 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.

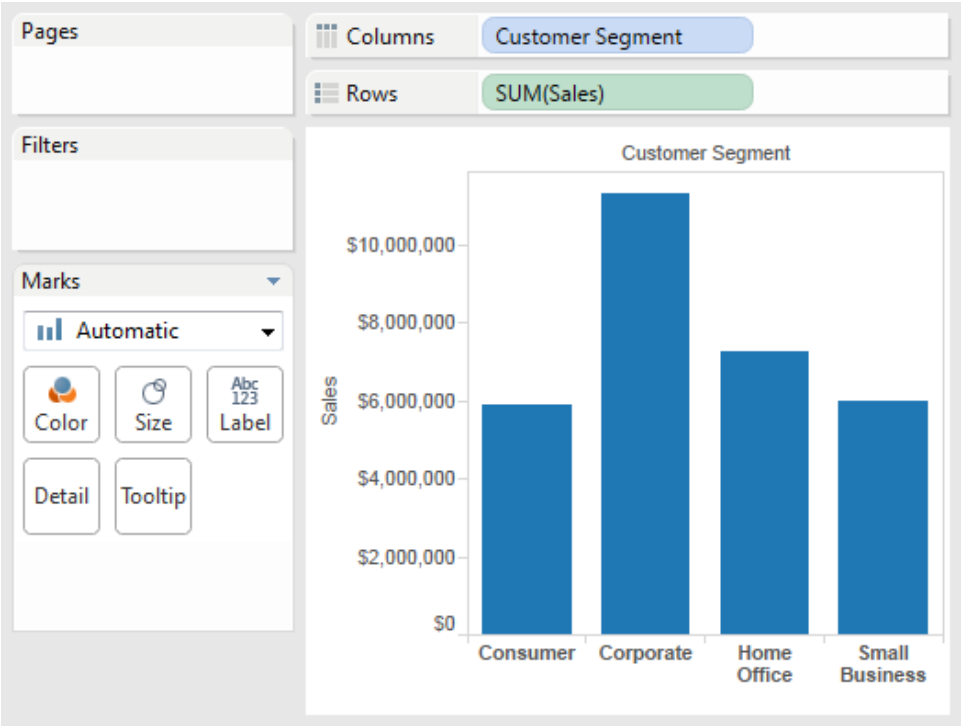

Suppose you are not interested in the Small Business data. You can remove this column from the view by filtering the Customer Segment dimension. To do so, select **Filter** on the field menu or drag the **Customer Segment** dimension to the **Filters** shelf. The Filter dialog box opens. By default all members are selected. Clear the checkbox for Small Business to exclude it from the view. All selected members will be included.

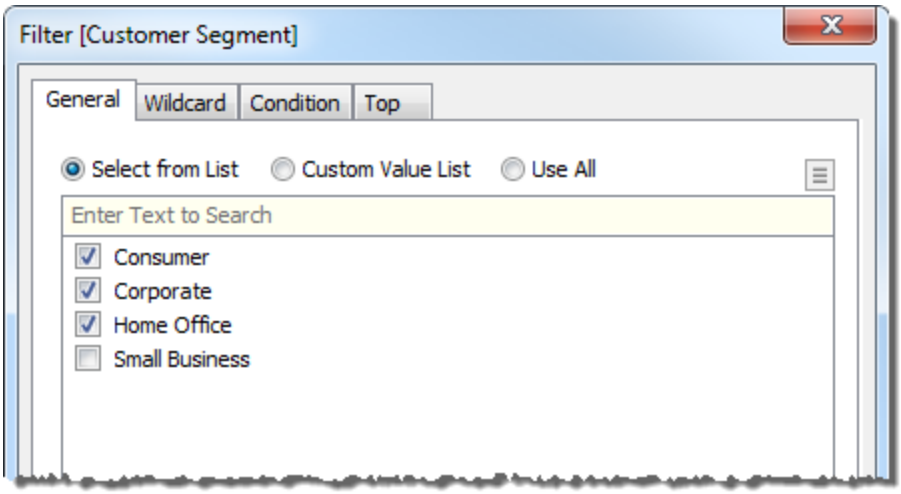

As shown below, the view updates and the Small Business column is removed. The filter is indicated by the Customer Segment field on the Filters shelf.

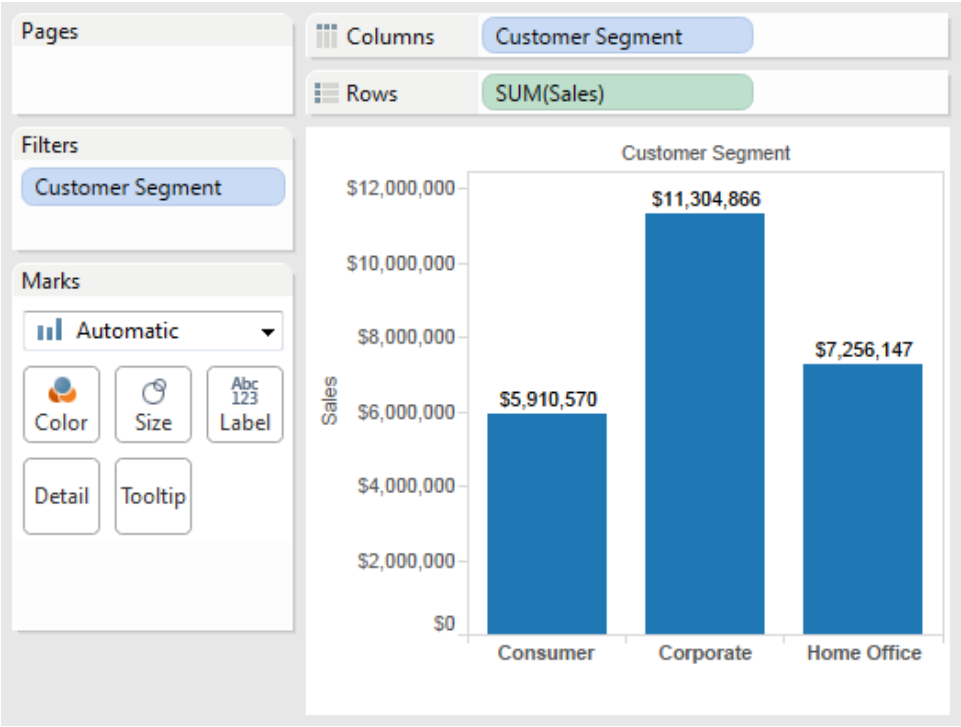

Suppose you want to only view sales for products that were shipped in boxes. Even though the Container field is not used on the Rows and Columns shelves or on the Marks card, you can still add a filter. Drag the **Container** dimension to the **Filters** shelf. This is an example of an external filter because Container is not part of the view.

The Filter dialog box automatically opens. By default, none of the members are selected. Select the members you want to keep as part of the view. All cleared members are excluded.

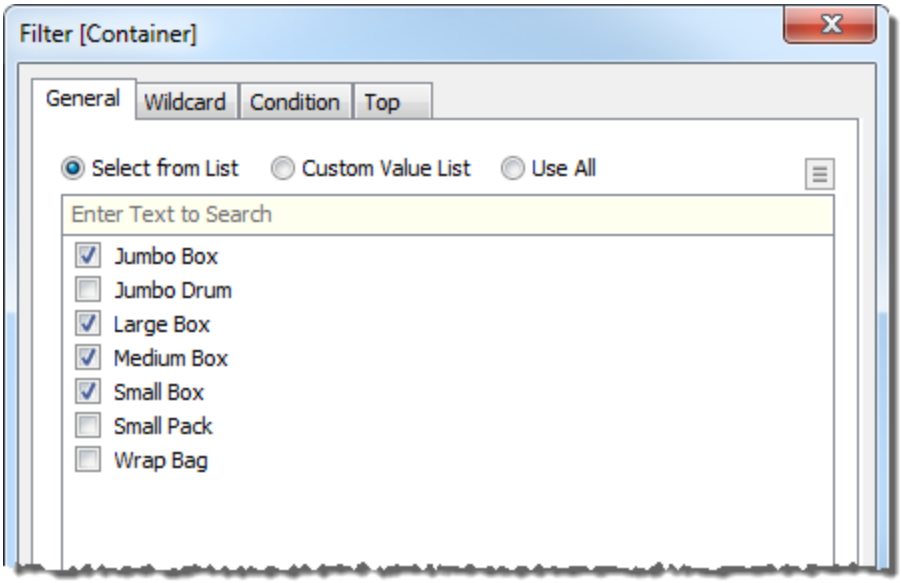

The modified data view is shown below. The tooltip shows that the sum of the sales for the Consumer segment has decreased to $4,331,503. This number is derived by summing all the rows in the data source that are associated with the Corporate market and that use a box as a shipping container.

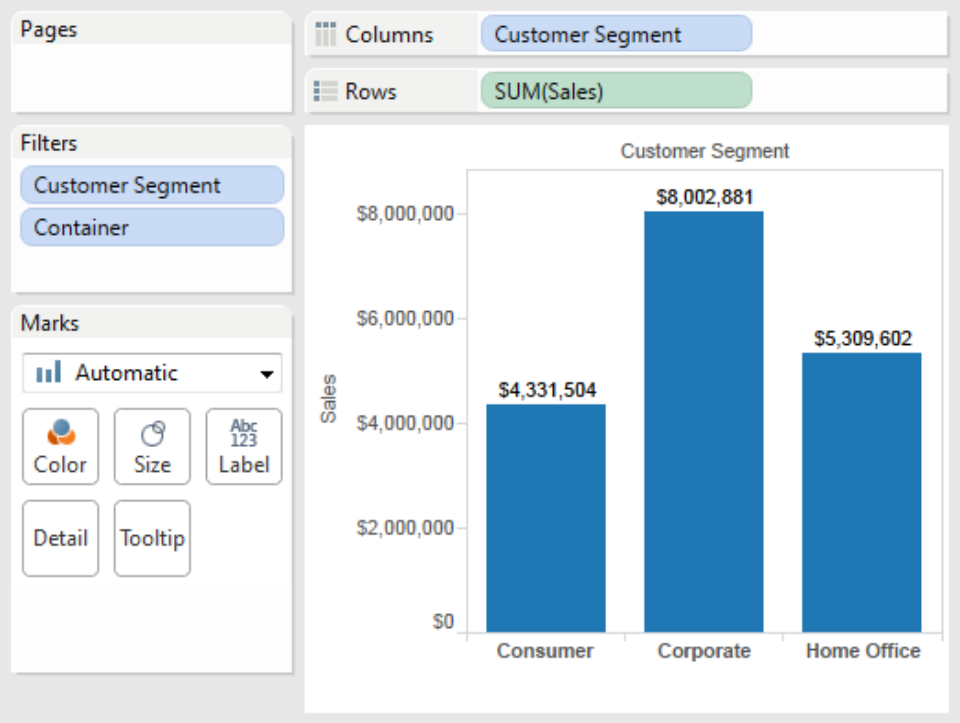

The order of fields placed on the Filters shelf does not affect the data view because the filters are independent. In other words, the result of filtering by customer segment, and then by container is the same as filtering by container and then by customer segment.

Marks Card

The Marks card is where you drag fields to control mark properties such as type, color, size, shape, and so on. The fields on the Marks card are listed at the bottom of the card. Each field has an icon next to it to identify the mark property it is setting. For example, the Marks card shown below has three fields: Customer Segment is on Color, Region is on Shape, and Container is on Size. Learn more about **Marks Card** on page 238, **Mark Types** on page 239, and **Mark Properties** on page 257.

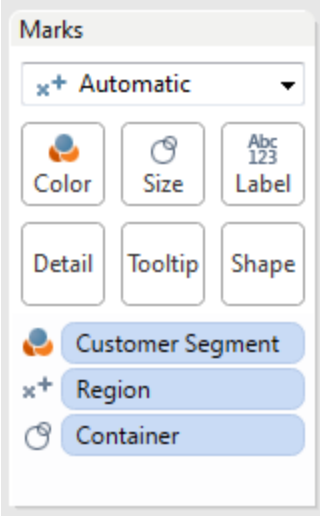

After a field has been added to the Marks card, you can click the icon next to the field to change the property it is controlling. The Detail and Tooltip properties don't display with an icon. However, you can click the white space where the icon would display to access the menu.

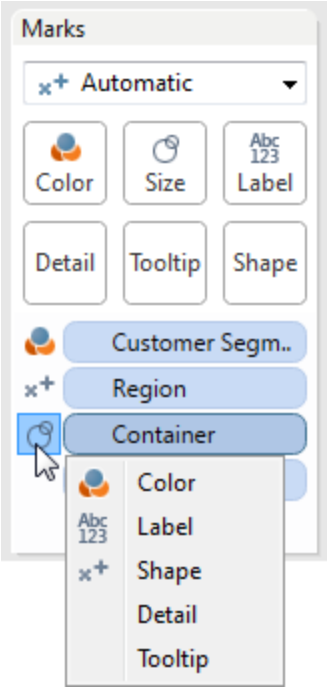

Many properties can have multiple fields. For example, you can add multiple fields to Label, Detail, Tooltip, and Color. Size and Shape can only have one field at a time.

> By default, dragging a new field to Color, replaces the existing fields. Hold the SHIFT key on your keyboard while dragging a new field to add it to Color without replacing the existing fields.

Multidimensional Hierarchies

In Tableau, multidimensional data sources are supported only in Windows.

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.

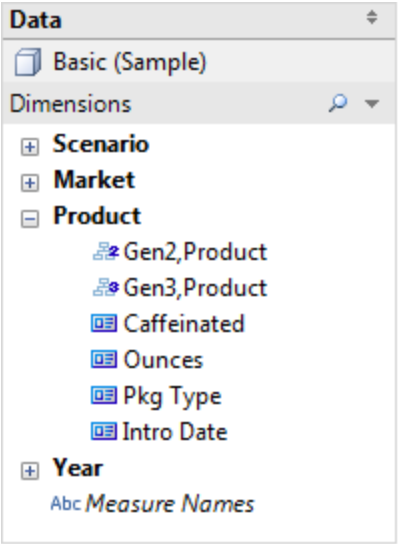

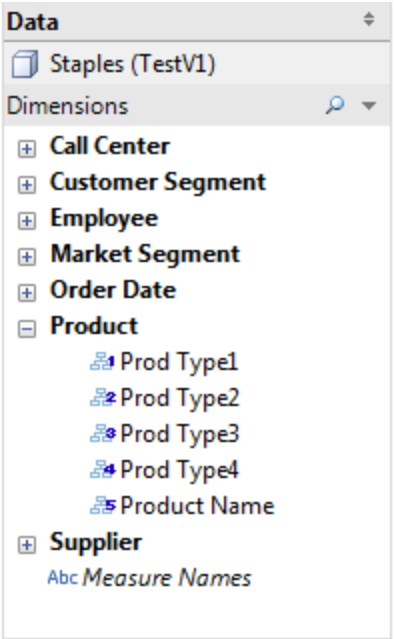

Drilling Down and Up in a Hierarchy

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 within each region.

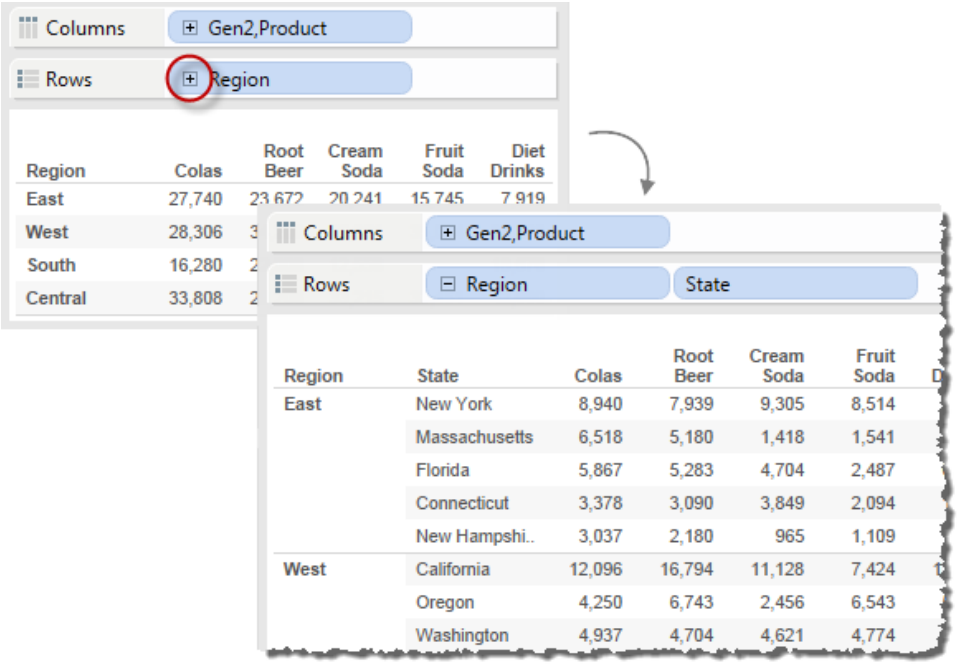

Using Headers

To drill down and drill up for individual dimension members in a hierarchy, right-click a table header and select **Drill Down** or **Drill Up** from the context menu. This is often referred to as non-uniform 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.

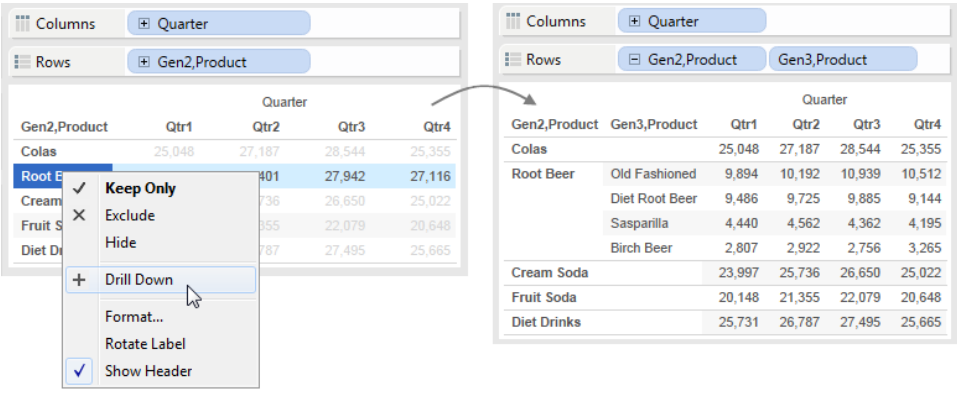

One reason to use non-uniform drill down is if your data source has a ragged hierarchy (asymmetric layout). You also might want to view the children for just the member of interest.

> Drilling down and drilling up results in filtering the data.

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

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**.

> 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.

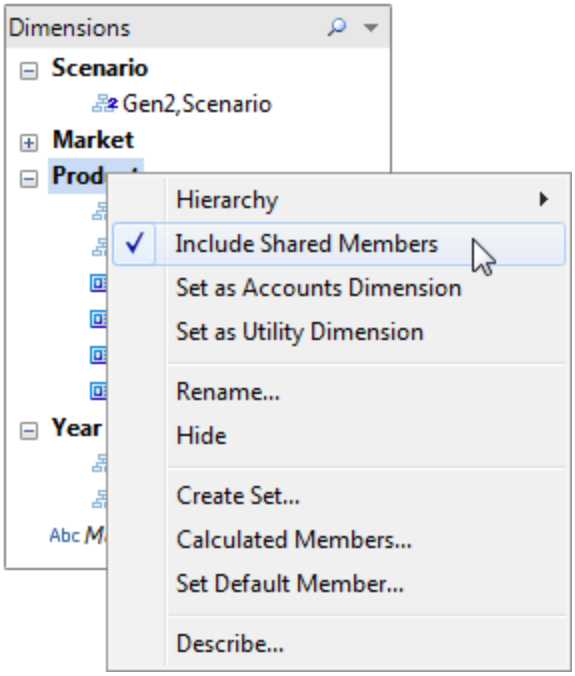

The figure below shows part of a data view where shared members are included (left) and excluded (right). Notice that diet drinks are shared members.

Includes Shared Members

The diet drinks are listed in both the Diet Drinks hierarchy and their respective drink hierarchies.

Does Not Include Shared Members

The diet drinks are only listed once, in their respective drink hierarchies.

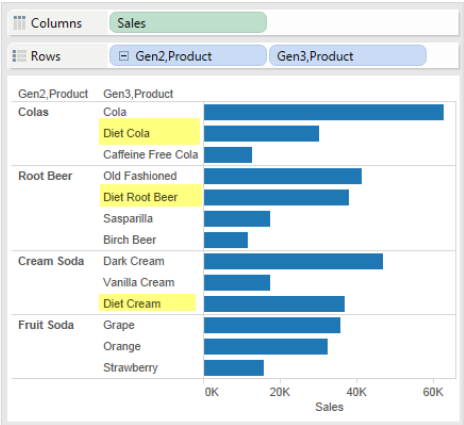

Setting the Default Member

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**.

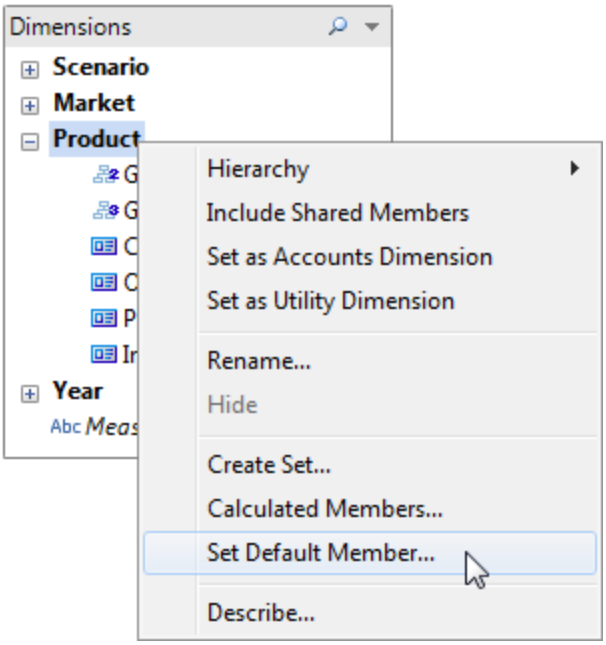

In the subsequent dialog box, select from the following options:

- 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

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.

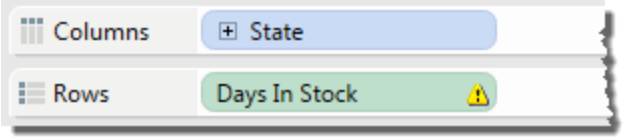

Perfect Pivoting

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 Color).
- 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 State level as Rows and the Region level as Color.

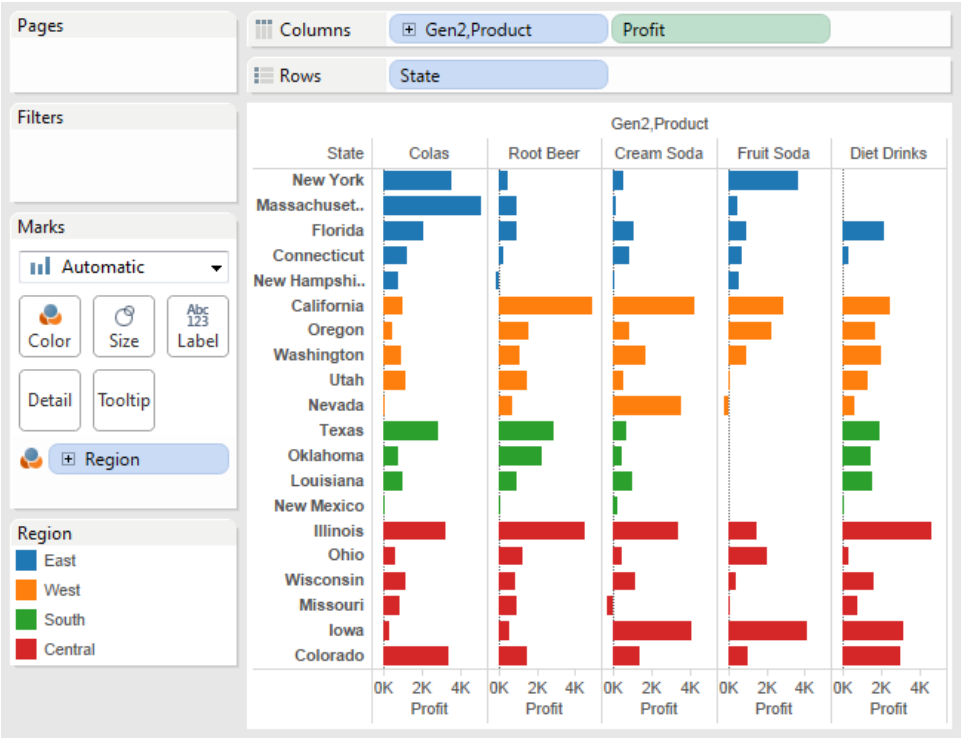

Defining Unique Values

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).

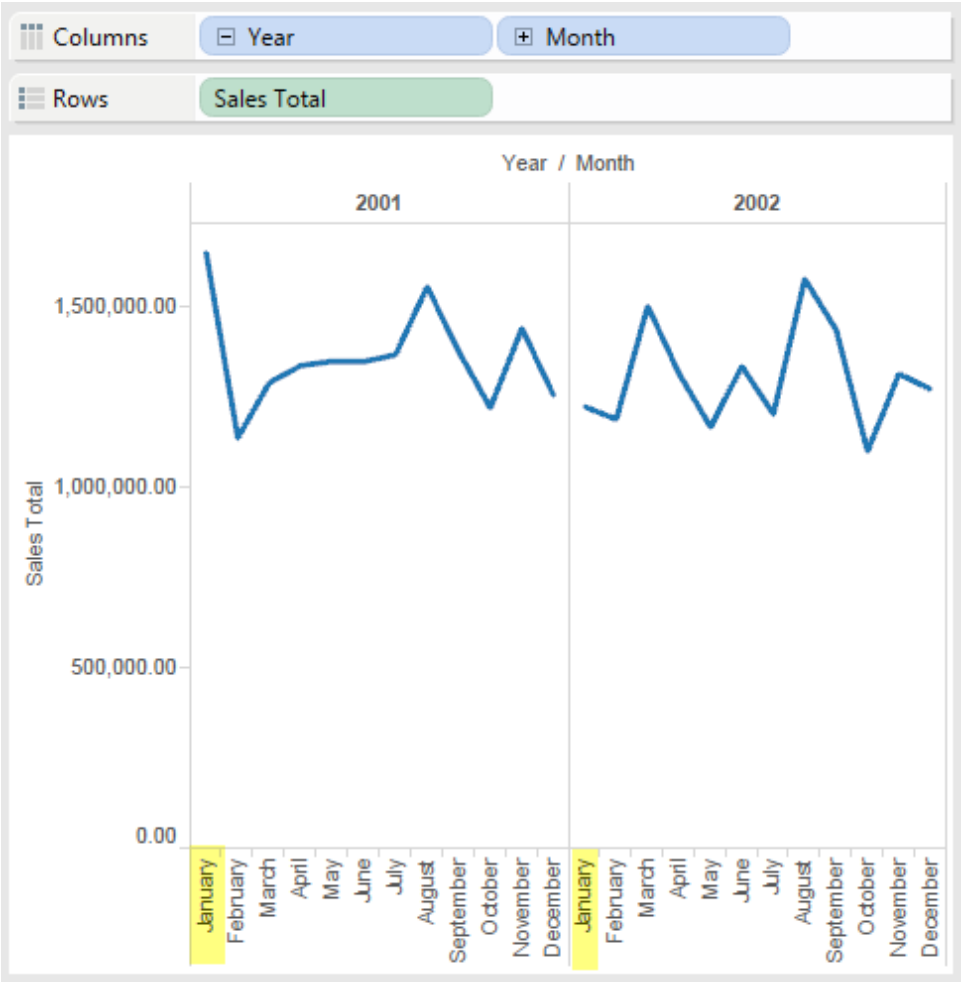

While the name, January, is repeated, each instance of January can either be considered similar or unique. If you consider them similar, they will appear in the same column if you decided to move the Year field to the Color. However, if you consider them unique, they will be treated as two different values.

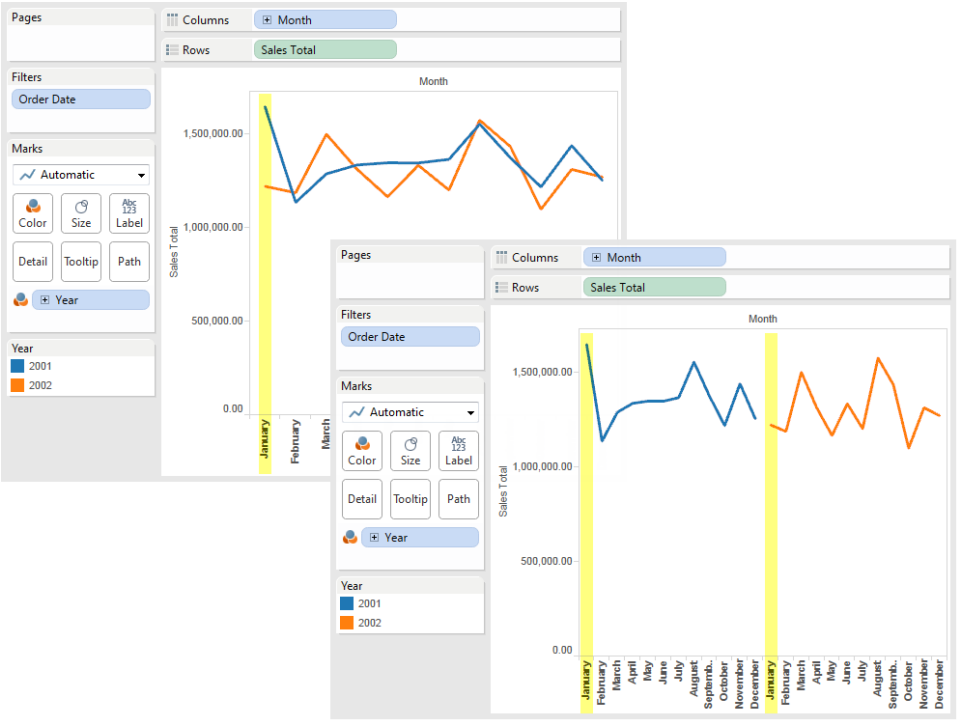

It is generally okay to consider repeated names within date and time fields (like in the previous example) similar but if there are repeated names in the Customer Name field, you won't want to consider the two customers as the same person.

To define how you want Tableau to determine whether repeated values are unique, right-click on the dimension and select 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.

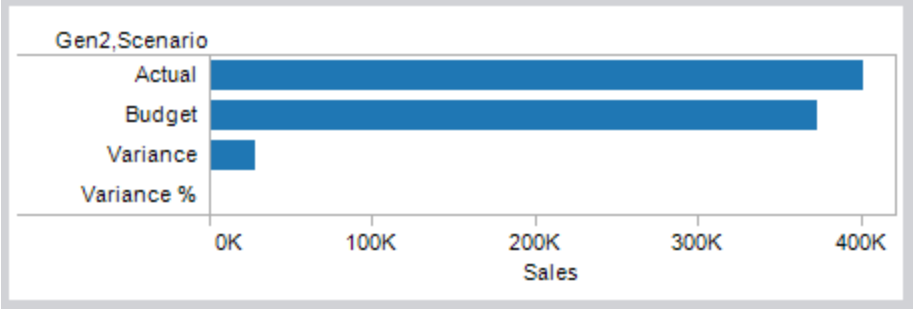

In the above view, you can see Actual Sales, Budgeted Sales, and so on. However, what if you wanted compare Actual Sales to Budgeted Sales in a bullet graph? In that case you need to set the Scenario dimension to be used as the utility dimension. When you set a dimension as the utility dimension you 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.

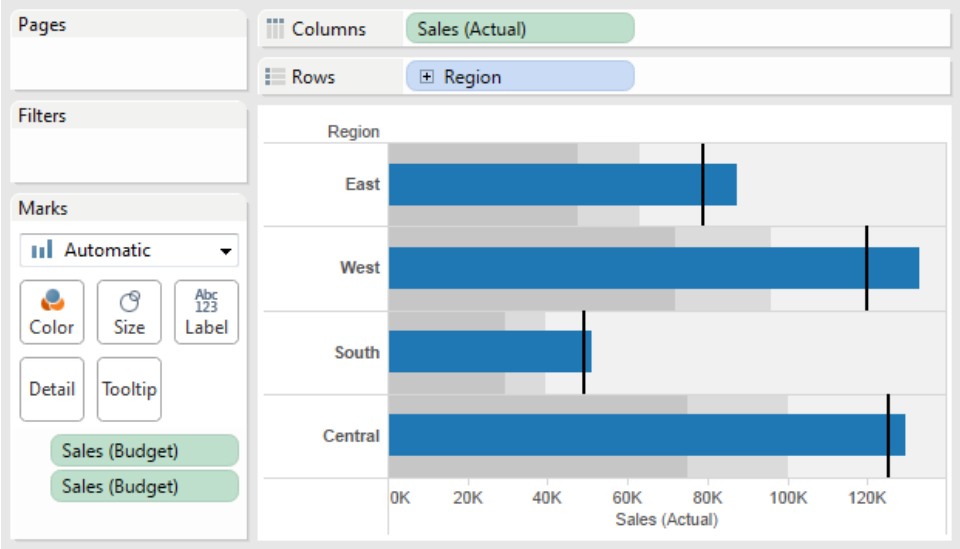

You can see that the Sales measure is used twice in the view: once to show actual and once to show budgeted.

To use a dimension as the utility dimension:

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

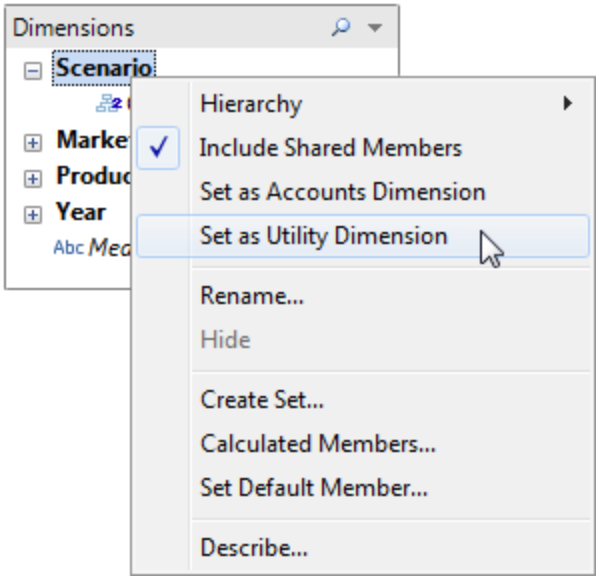

The dimension is hidden in the Data window and can no longer be used as a dimension field in the view. The Measures area of the Data window indicates that there is a utility dimension.

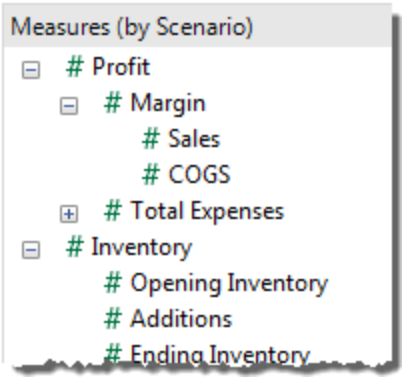

2. Drag a measure to the view.
3. Open the field menu for the measure in the view and select the member of the utility dimension you want to anchor the measure to.

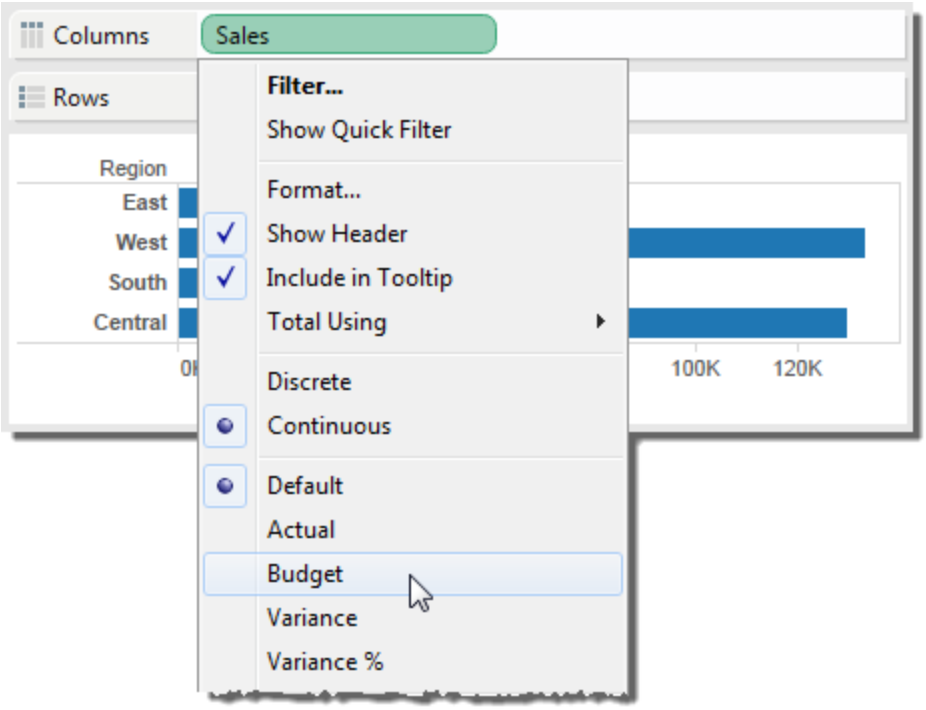

To remove a utility dimension:

Open the drop-down menu at the top of the Data window and select **Clear Utility Dimension**.

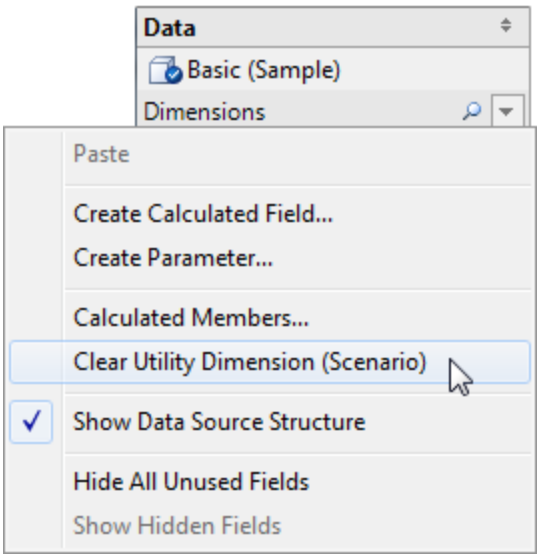

When you remove the utility dimension, measures that reference it in the view are no longer valid.

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. In some cases, you may be presented with a dialog box giving you the option to filter the members before adding it to the shelf.

If you are building a view that involves a large amount of data, you may find it easier to turn off automatic updates, build the view, and then resume updates. Turning off automatic updates queues the queries instead of sending a separate query to the database every time you drag a field. See **Automatic Updates** on page 191.

Follow the steps below to pause automatic updates while building the view.

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.
4. Turn on automatic updates by clicking the **Resume Automatic Updates** button on the toolbar.

Any time a query is taking too long, you can cancel the query by clicking **Cancel** in the progress dialog box.

Building Views Automatically

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.

Automatically generate views using Show Me or Double-Click.

Show Me

Show Me creates a view based on the fields already used in the view and any fields you've selected in the Data window. Open Show Me by clicking **Show Me** on the toolbar 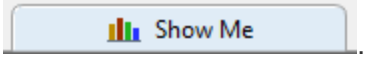

. When you use Show Me simply select fields you want to analyze in the Data window and then select the type of view you want to create. Tableau automatically evaluates the selected fields and gives you the option of several types of views that would be

appropriate for those fields. Further, Show Me highlights the visualization type that best matches the data.

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

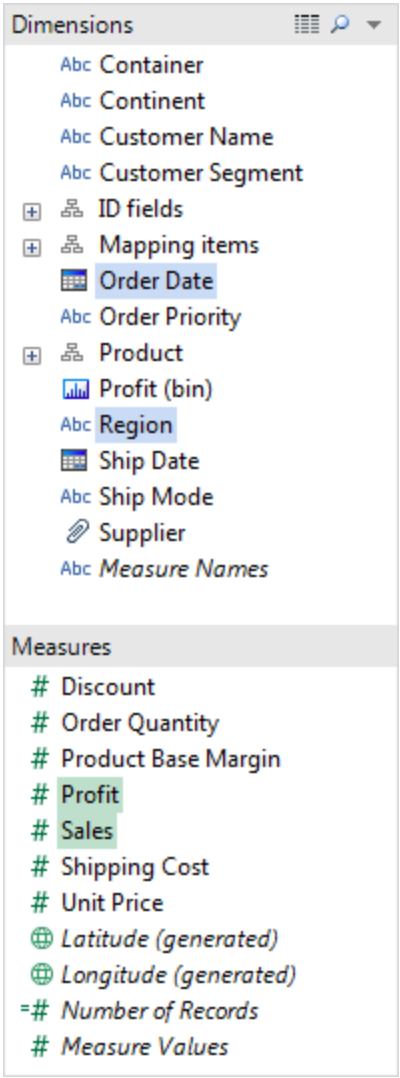

2. Click Show Me on the toolbar and then select the type of view you want to create.

Any view type that is not gray will generate a view of your data. As you hover over each view type, the description at the bottom shows the minimum requirements.

3. View the Result. Tableau automatically creates a view of the data.

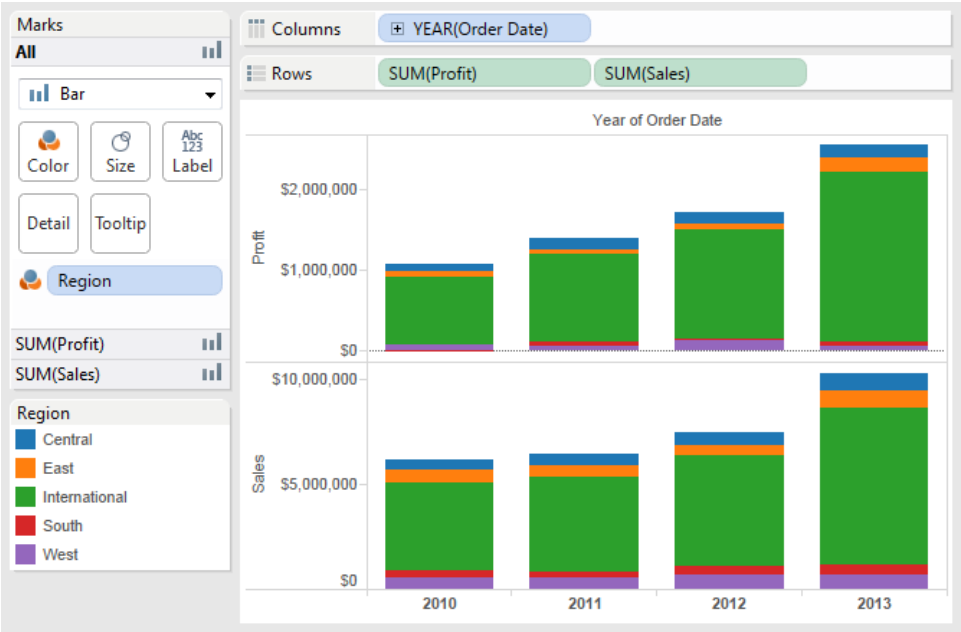

Add to Sheet: Double-Click

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. Double-clicking sometimes creates a view that you can use without further modification. More often, it can provide a starting point, which you can then modify to get the ideal result.

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

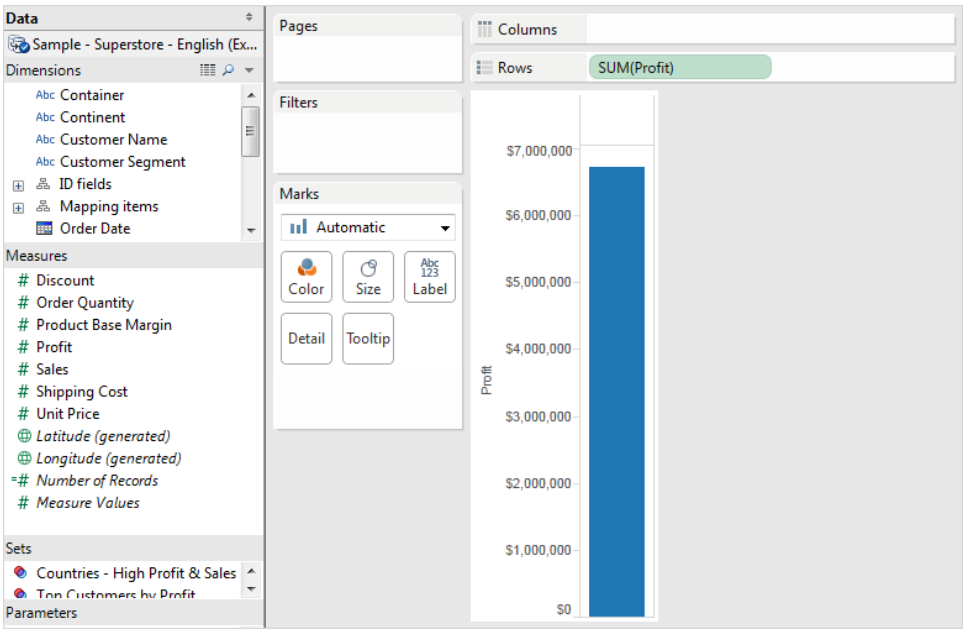

2. Double-clicking the Department dimension in the Data window automatically adds that field to the view based on the fact that Profit is already on the Rows shelf.

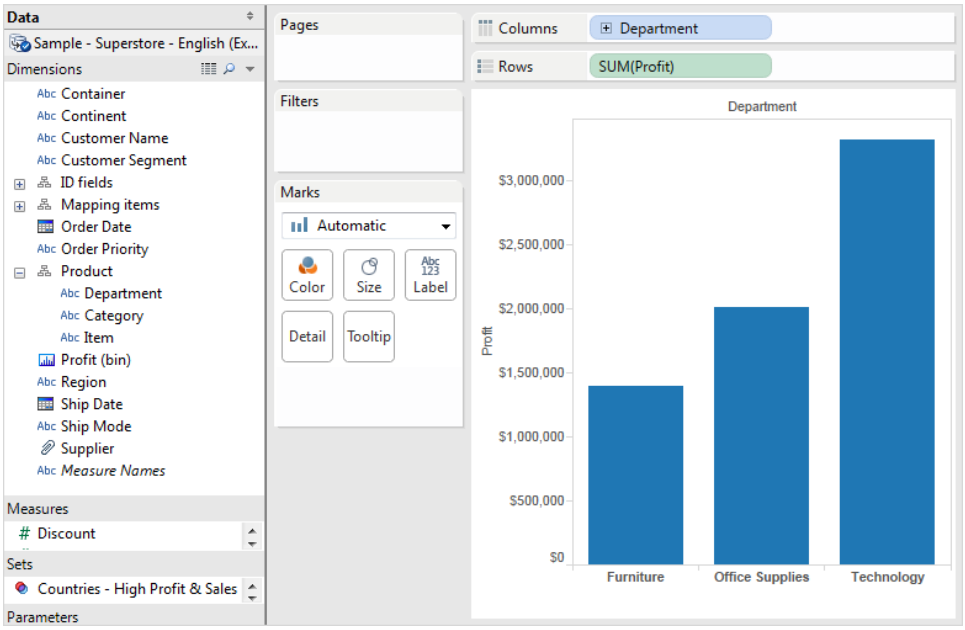

3. Double-clicking the Order Date dimension in the Data window automatically adds that field to the Columns shelf based on the placement of other fields in the view. As you double-click fields they are successively added to the view. The order in which you click fields determines the type of view created.

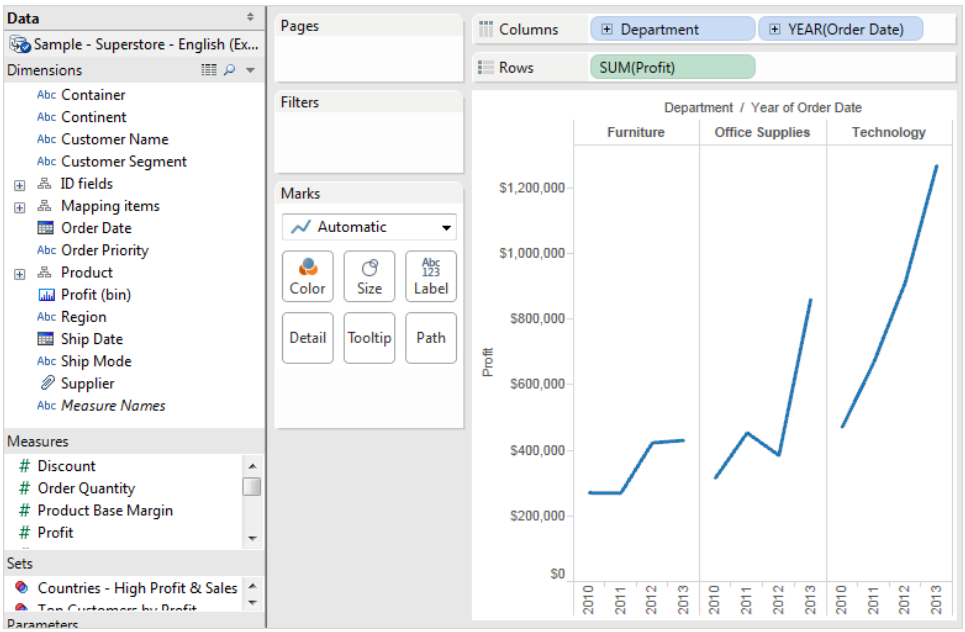

The following table describes some of the rules used in creating automatic views by double-clicking fields in the Data window.

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

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.

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.

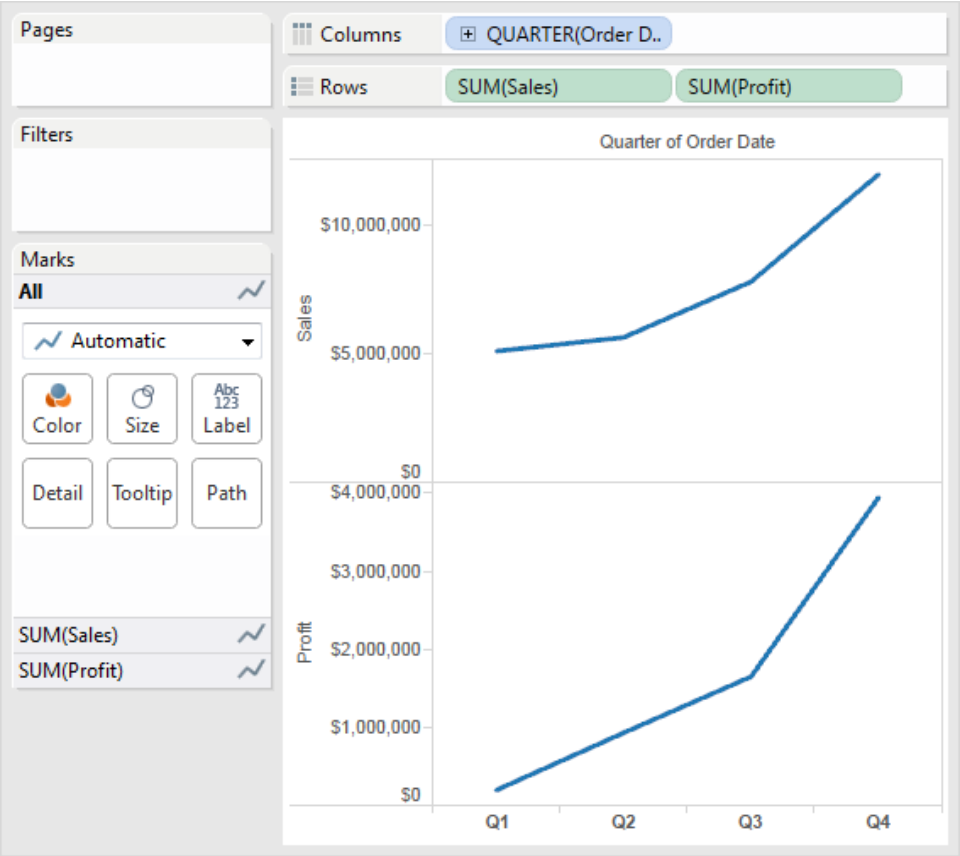

Blended Axes

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.

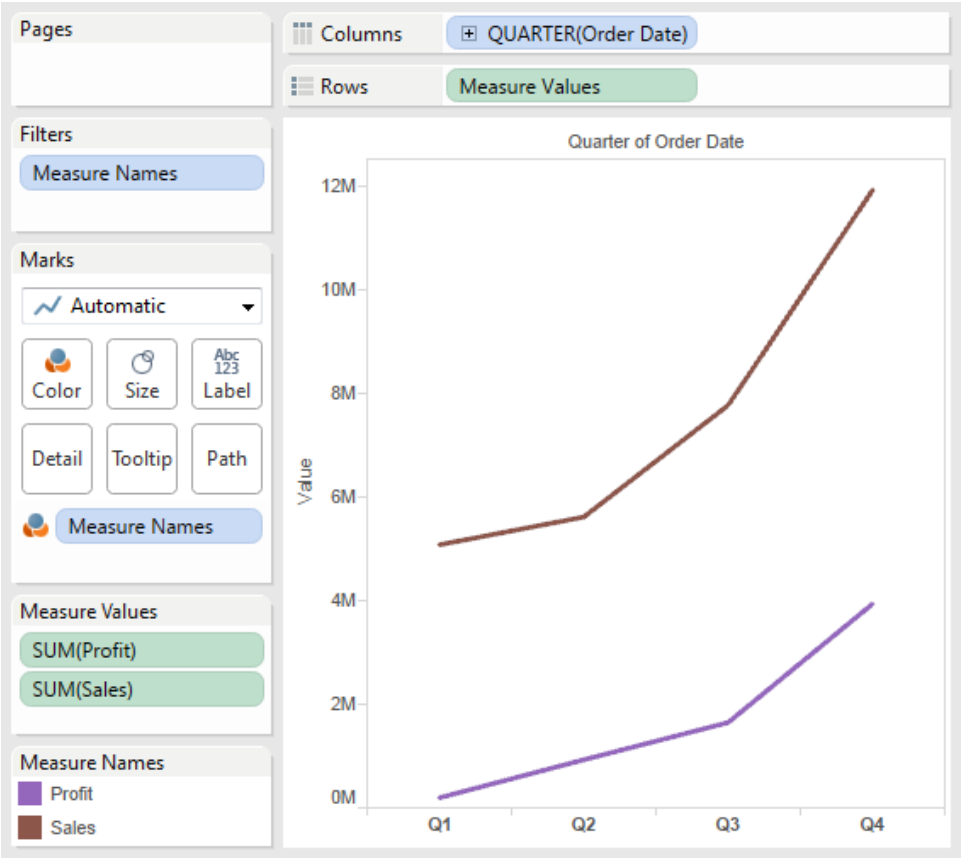

To blend multiple measures, simply drag one measure or axis and drop it onto an existing axis.

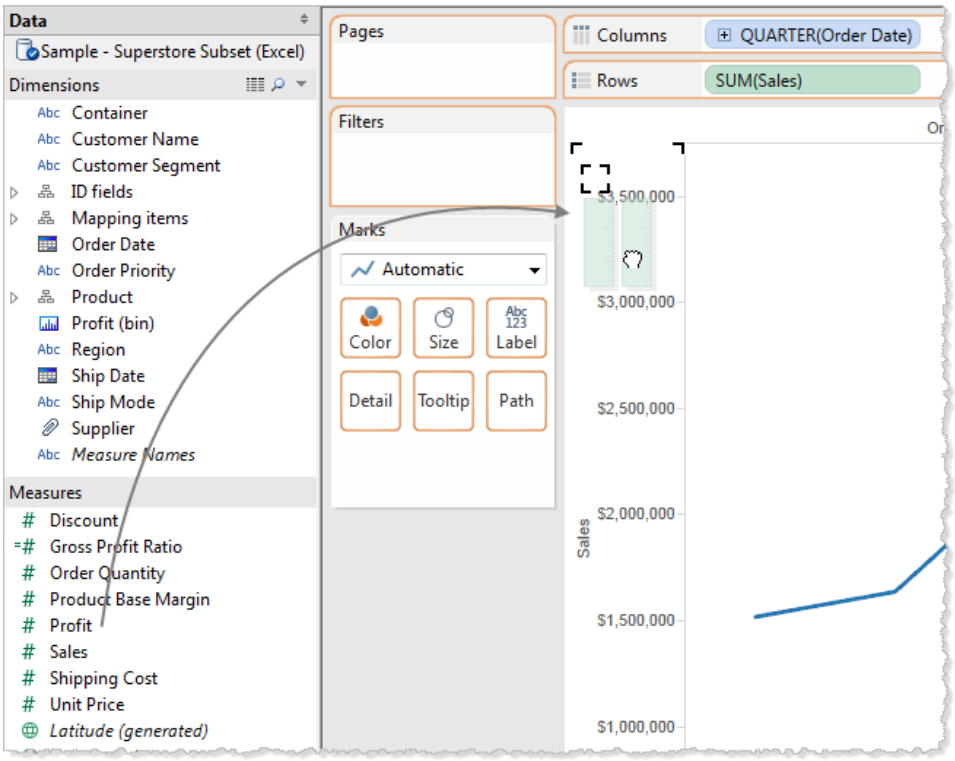

Blending measures uses the Measure Names and Measure Values fields, which are generated fields that contain all of the measure names in your data source and all of the measure values. The shared axis is created using the Measure Values field. The Measure Names field is added to the Color shelf so that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

> 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. 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.

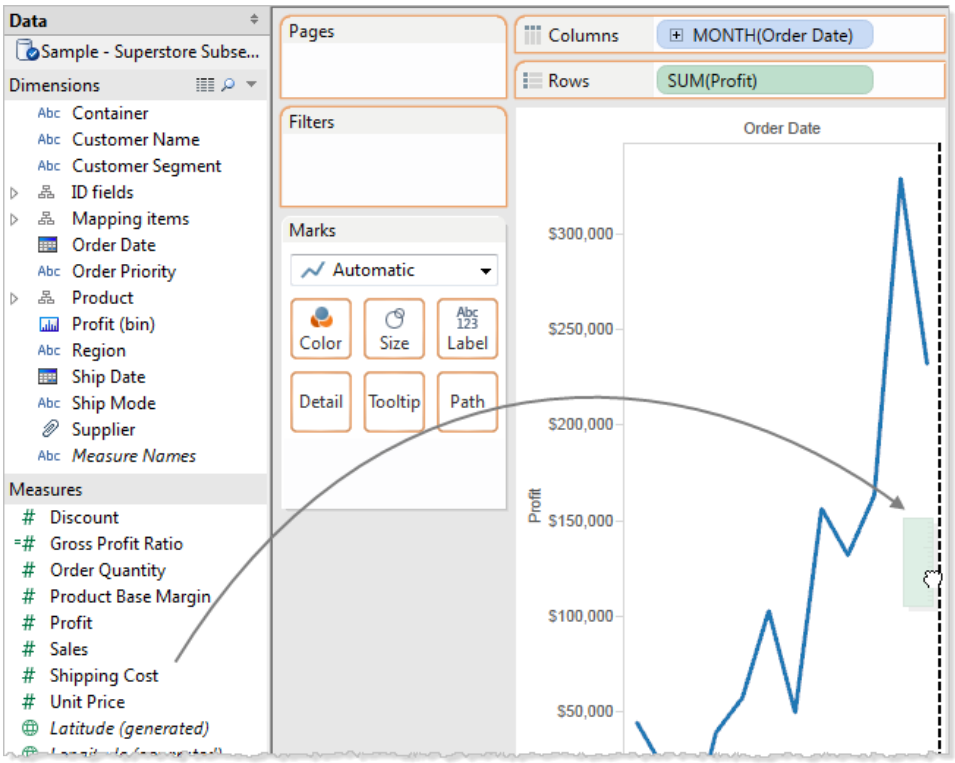

The result is a dual axis view where the Profit axis corresponds to the purple line and the Shipping Cost axis corresponds to the brown line.

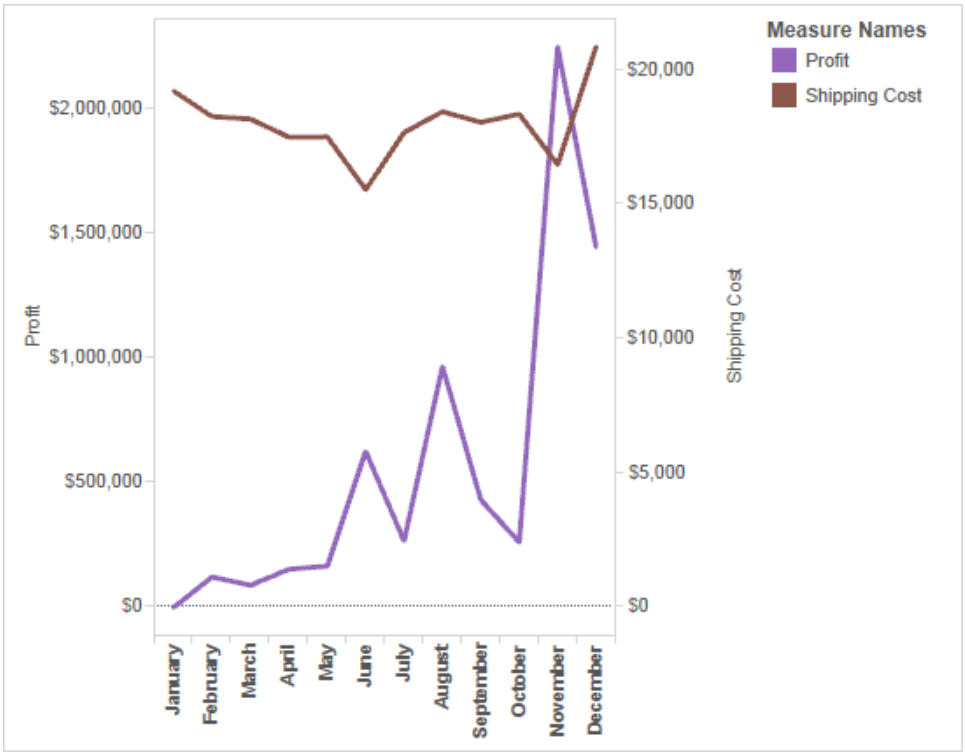

You can add up to four layered axes: two on the Columns shelf and two on the Rows shelf.

> To align the two axes in a dual axis up, right-click the dual axis and select **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.

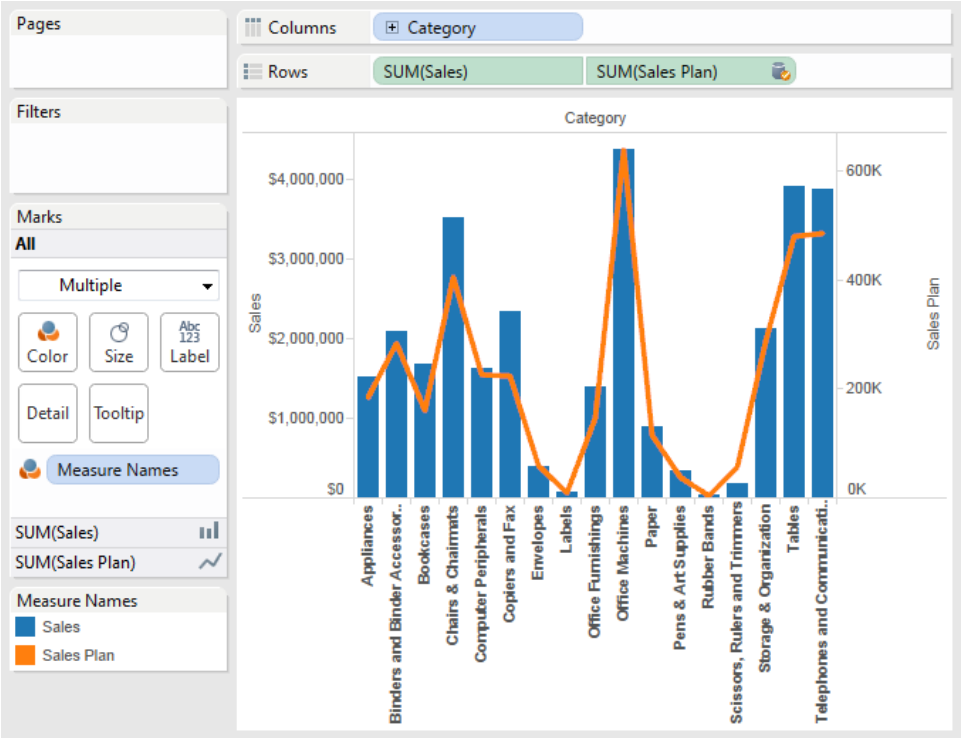

To customize the marks for a measure:

1. Select the Marks card for the measure that you want to customize. There is a Marks card for each measure on the Rows and Columns shelves.

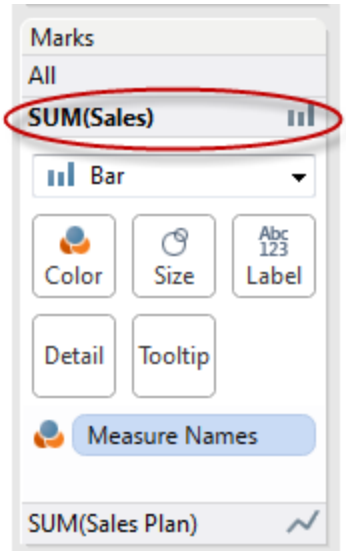

2. Select a new mark type for the measure. Any changes to the mark type, shape, size, color, detail and other mark properties will be applied to the selected measure. For example, in the view below the SUM(Sales) Marks card is active. The Mark type has been changed to Bar and when Department is placed on Color, the encoding and level of detail is only applied to the SUM(Sales) marks. The SUM(Sales Plan) is not broken down by Department.

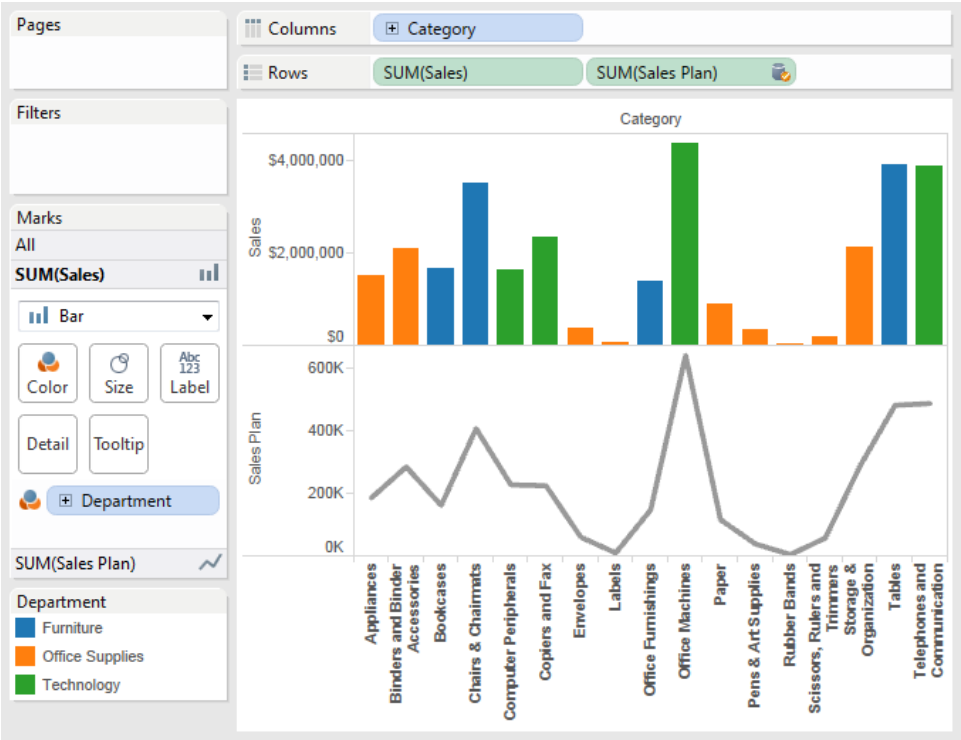

Select the **All** Marks card to modify properties for all measures at once.

Filtering

Narrow the data shown in a view using Filters. 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, making it easy to 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.

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.

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 or the tooltip. The dimension members are removed from the view and the filtered fields are added to the Filters shelf.

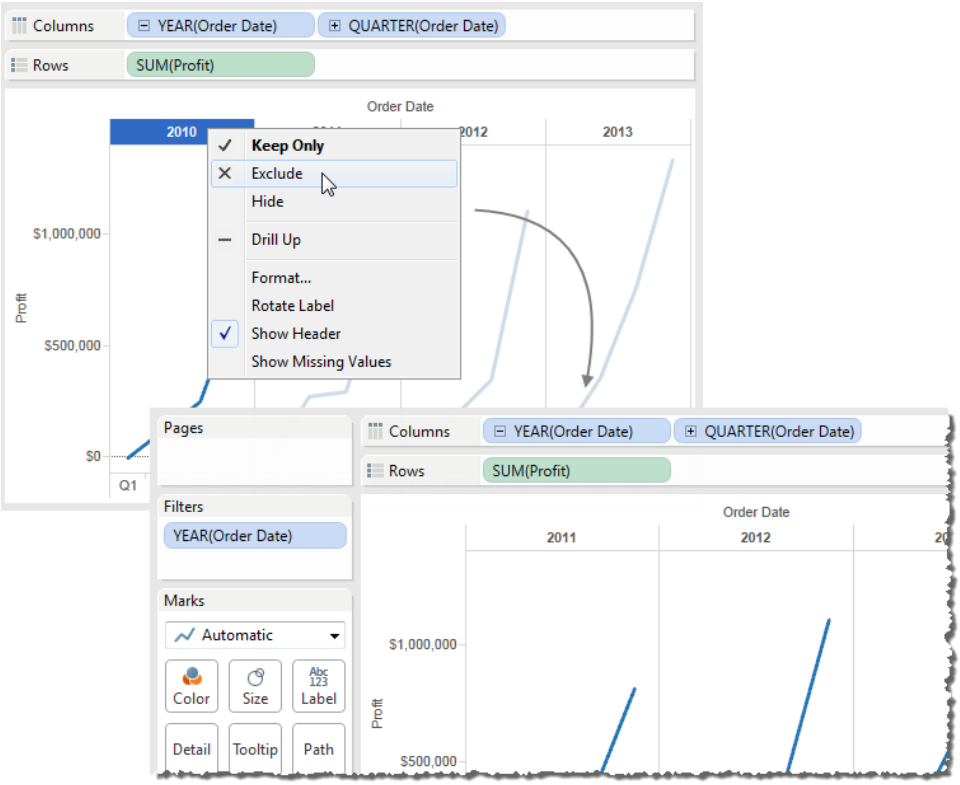

Selecting Headers

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.

Columns | Gen2,Product | Pkg Type

Rows | Region | State

Region	State	Colas Bottle	Colas Can	Root Beer Bottle	Cream Soda Bottle	Cream Soda Can	Fruit Soda Bottle	Diet Drinks Bottle	Diet Drinks Can
East	New York		8,940	7,939	9,305		8,514		
	Massachusetts		6,518	5,180	1,418		1,541		
	Florida		5,867	5,283	3,630	1,074	2,487	2,393	4,142
	Connecticut		3,378	3,090	3,849		2,094	1,384	
	New Hampshi..	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				1,291	1,515		2,276	3,378
	Oklahom			543	2,890	1,562		2,799	2,427
	Louisiana			774	2,085	1,797		1,371	3,214
	New Mex			267	825	593		1,087	2,124
Central	Illinois			026	4,918	3,183	4,241	5,875	6,834
	Ohio			228	865	676	5,867	1,418	2,109
	Wisconsi			142	1,417	1,697	3,077	1,074	4,090
	Missouri			440	1,531	1,368	2,864	1,125	3,132
	Iowa	593	1,331	1,129	3,651	7,151	12,943	516	7,657
	Colorado	3,183	4,918	4,241	2,948	3,810	4,459	1,965	6,865

✓ Keep Only
✕ Exclude
Hide
Drill Up
Format...
Rotate Label

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 | Gen2,Product | Pkg Type

Rows | Region | State

Region	State	Colas Bottle	Colas Can	Root Beer Bottle
East	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
South	Texas		8,073	7,424
	Louisiana		2,660	4,774

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: Gen2,Product | Pkg Type

Rows: Region | State

		Cream Soda		Fruit Soda	Diet Drinks	
Region	State	Bottle	Can	Bottle	Bottle	Can
West	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
South	Oklahoma	2,890	1,562		2,799	2,427
	New Mexico	825	593		1,087	2,124
Central	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

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** or click the filter options on the tooltip.

> The filtering options are not available if a Wildcard Match filter is already specified for the same field. Refer to **Filtering Dimensions** on the next page to learn more about Wildcard Match filters.

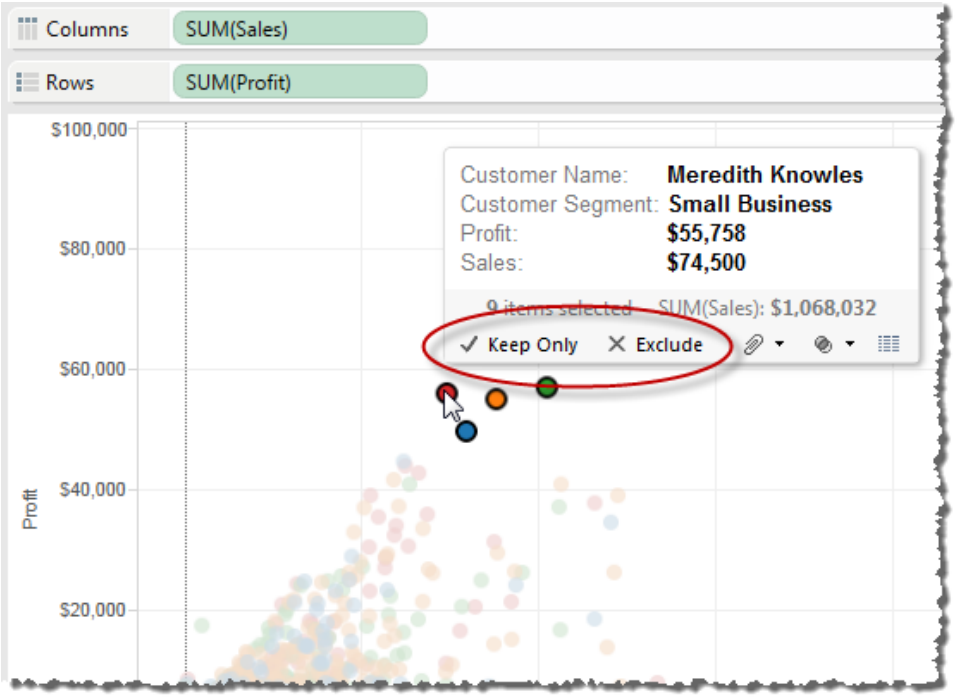

Dragging Fields to the Filters Shelf

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.

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.

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**.

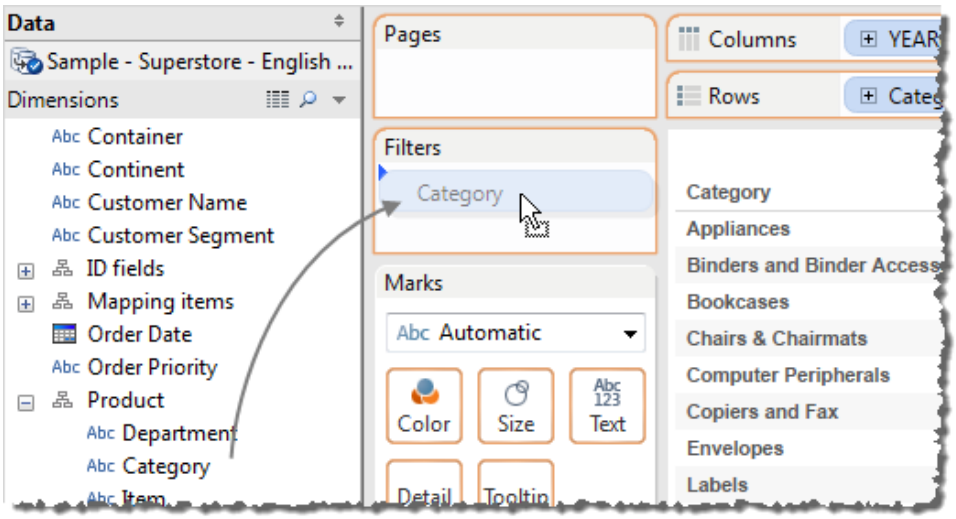

2. Use the General Tab of the Filter dialog box to select the values you want to include or exclude.

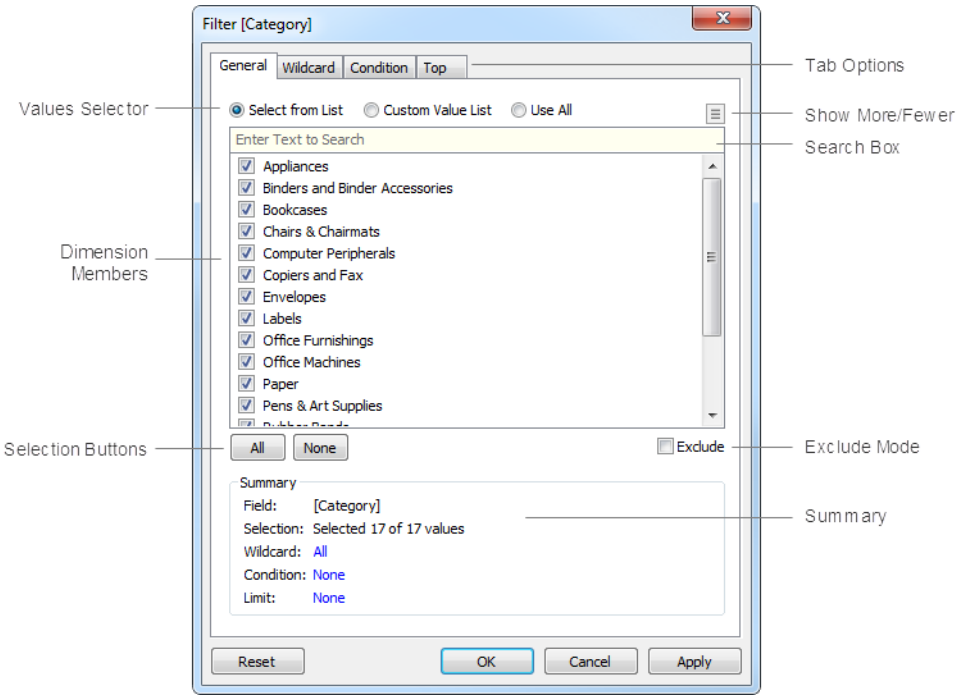

Each option on the General tab is described below:

Show More/Fewer - The contents of the Filter dialog box is affected by the filters that are already set in the view. For example, two filters are shown below: Writing Utensils and Colors. When the Colors filter is set to Show Fewer Values, selecting Pen in the writing utensils filter modifies the Colors filter to only show Black and Blue. That's because the underlying data only contains black and blue pens. When the Colors filter is set to Show More Values, all colors in the underlying data are always shown regardless of what Writing Utensil is selected. The

Show More/Fewer option includes and excludes data from displaying in the Filter so you can find what you are looking for quicker.

Colors is set to Show Fewer Values

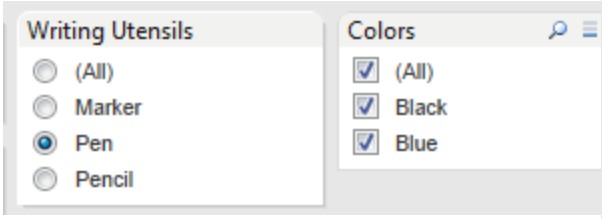

Because the data only contains black and blue pens, the Colors filter only shows values that pass the Writing Utensils filter.

Colors is set to Show More Values

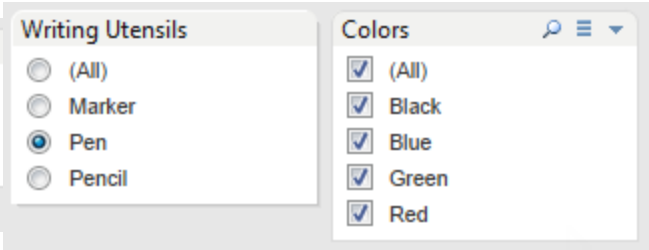

All colors in the data are shown regardless of what is selected in the Writing Utensils filter. In this case, changing the filter to exclude Red will have no affect on the view because there are no red pens in underlying data.

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)
- **Custom Value List** - 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.

Search Box - 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. As you type into the search box, matching values show below the search box. Select the values you want. Each search adds to the selection.

Dimension Members - Select members from the list when using the Select from List mode. When in Custom Value List mode this area displays the list of dimension members you have typed in manually.

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.

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 to 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.

Adding Wildcard Match to Filters

Use the Wildcard Tab in the Filter dialog box to define a pattern to filter on. For example, when filtering on email addresses you may want to only include emails from a specific domain. You can define a wildcard filter that ends with "@gmail.com" to only include Google email addresses.

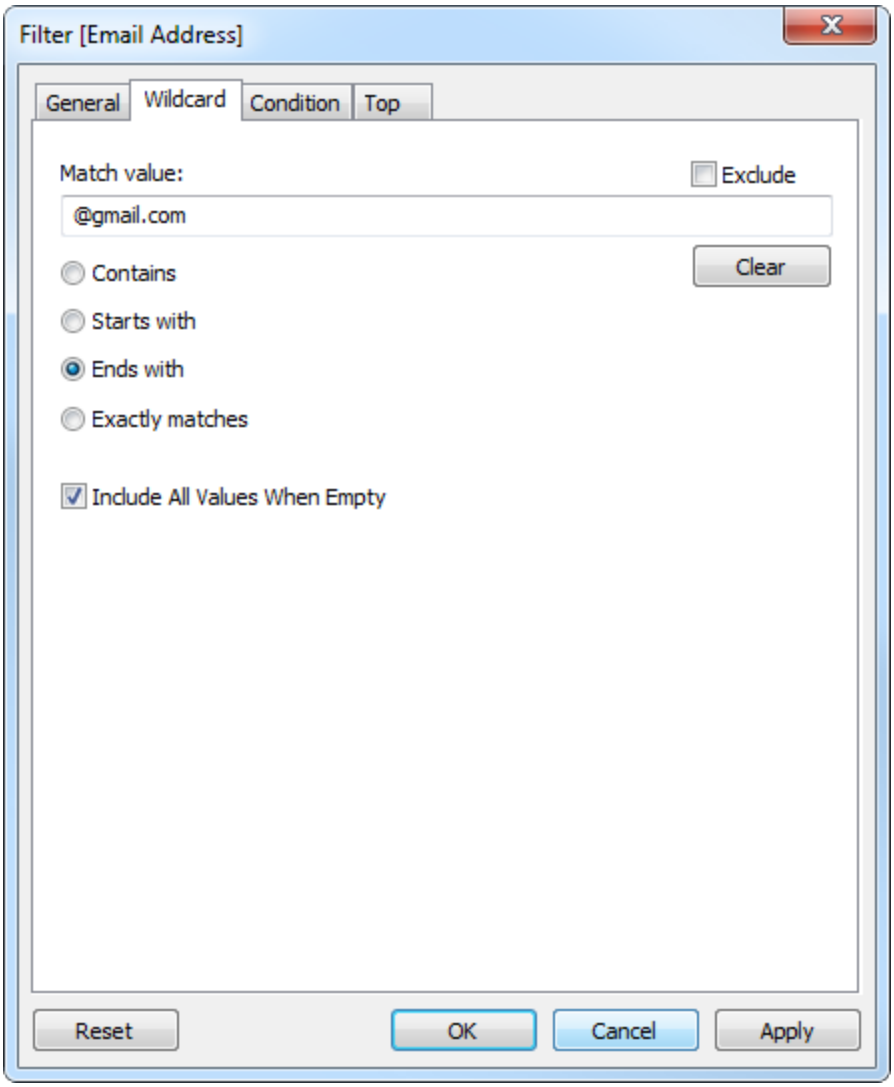

Each option on the Wildcard tab is described below:

Match Value - type a pattern that you want to match dimension members to.

Match Type - Select one of the options below for how to match the pattern.

- **Contains** - returns all members that contain the match value anywhere in the string.
- **Starts with** - returns members that have the match value at the beginning of the string.
- **Ends with** - returns member that have the match value at the end of the string.
- **Exactly matches** - returns the members that contain only the match value and nothing else.

Exclude - By default, members that are returned are included in the filter. The reverse this behavior, select the Exclude option. Matched values are excluded from the filter instead.

Include All Values When Empty - When this option is selected, the filter includes all dimension members when the match value couldn't be found at all. This option is selected by default. When cleared, the filter can potentially return zero results and result in a blank sheet.

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 Unit Price for a collection of products, you may want to only show the Products that have an average unit price that is greater than or equal to $25. You can use the built-in controls to write a condition or you can write a custom formula.

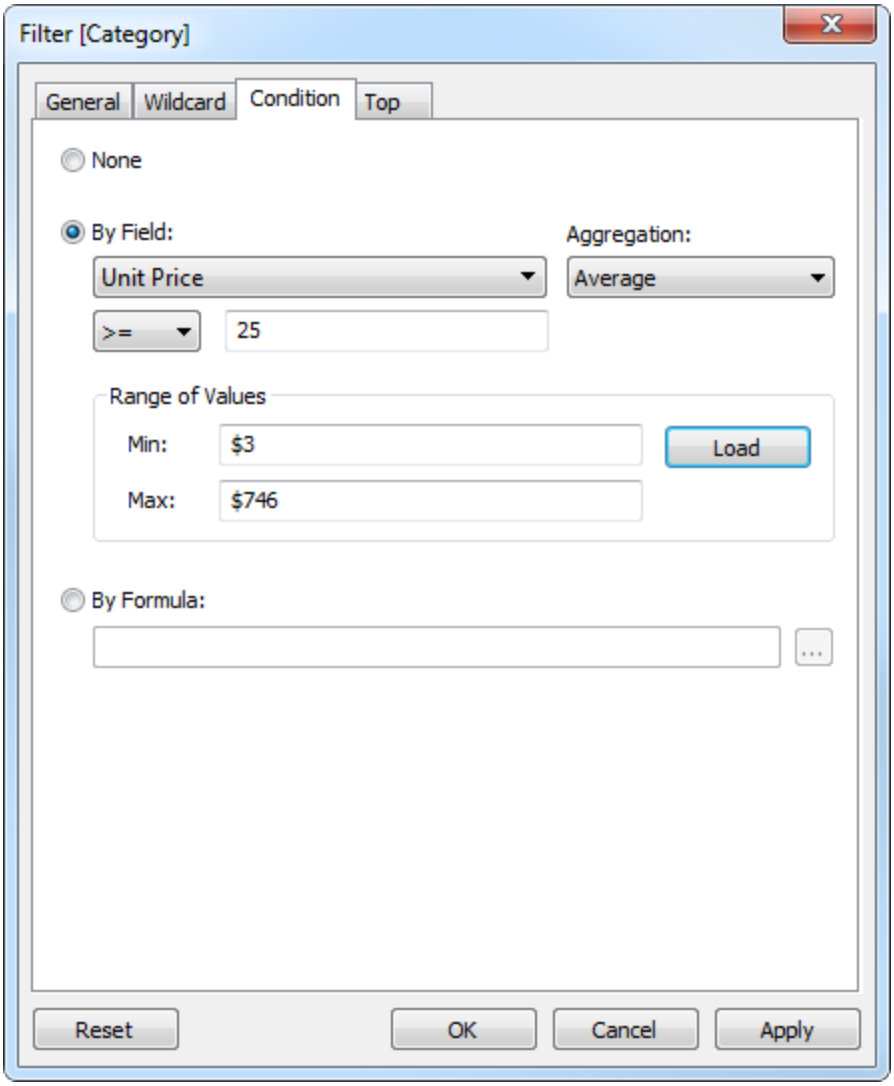

Each option on the Condition tab is described below:

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 menus 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 Unit Price and AVG from the first two drop-down menus. Then select the greater than or equal to symbol (>=) from the operator list and type 25 into the final text box.

> 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 had an average unit price under $25. Using the Range of Values box helps you decide a value that makes sense based on 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 more button ... on the right of the text box.

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 then 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.

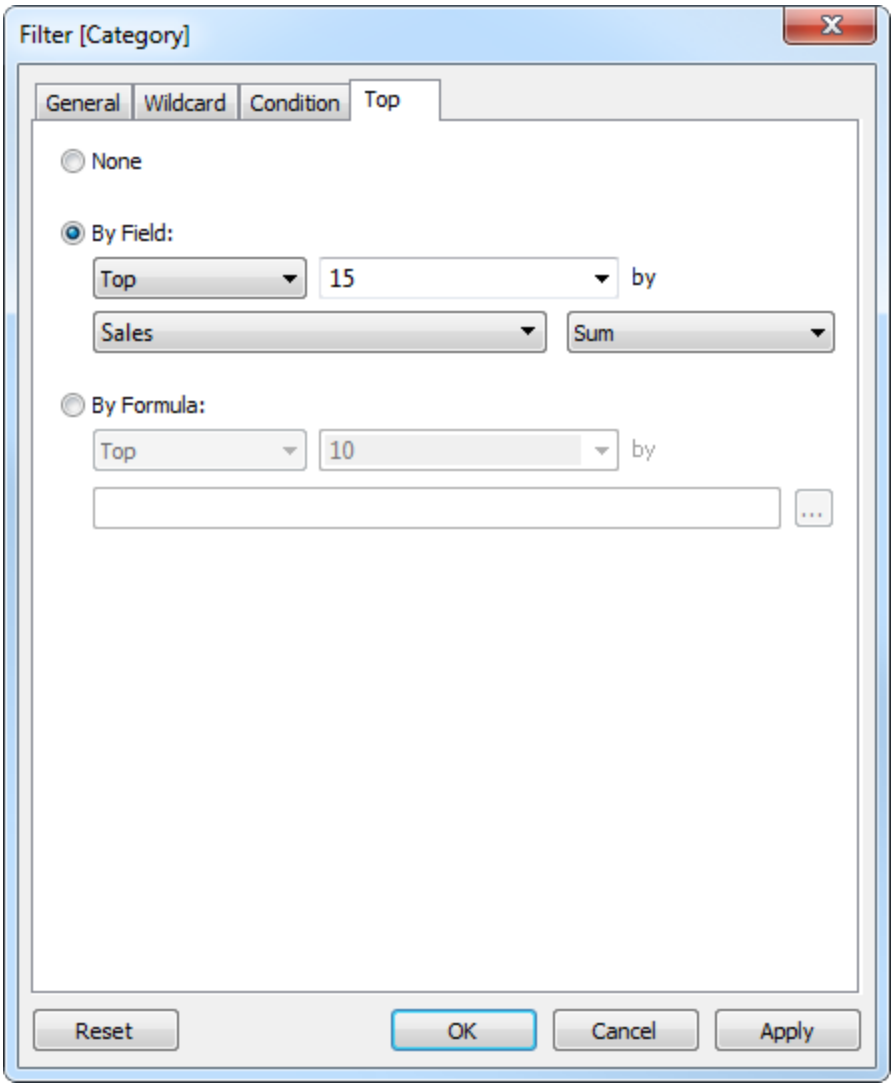

Each option on the Top tab is described below:

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.

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 view shown below. It was created using the Superstore sample data source that comes with the product. The view shows the profitability for three product departments broken down by their ship mode and container.

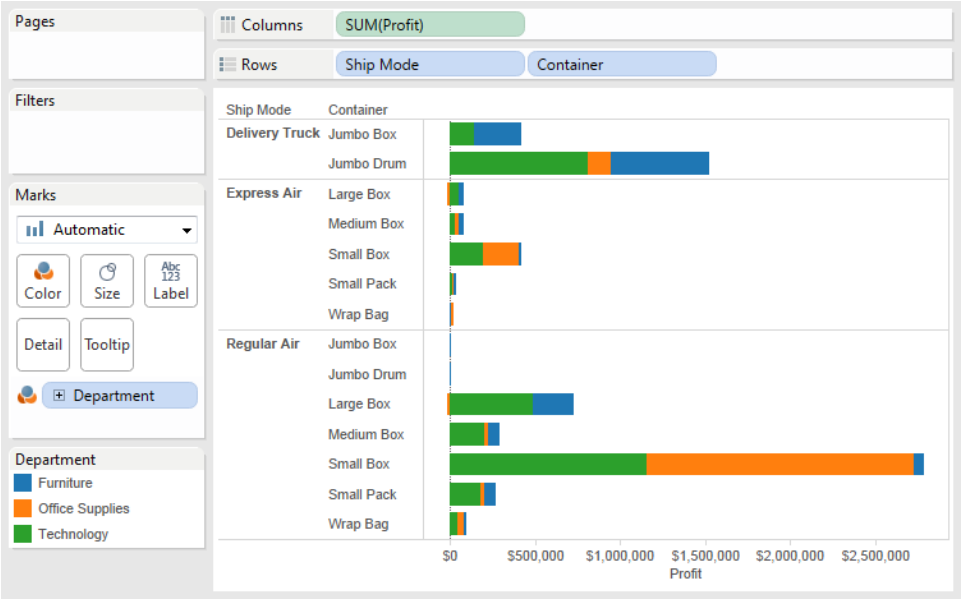

2. Create a basic filter on the Container dimension that excludes the Small Pack and Wrap Bag shipping containers.

   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 lower right corner of the dialog box. Finally, select **Small Pack** and **Wrap Bag**. When finished click **OK**.

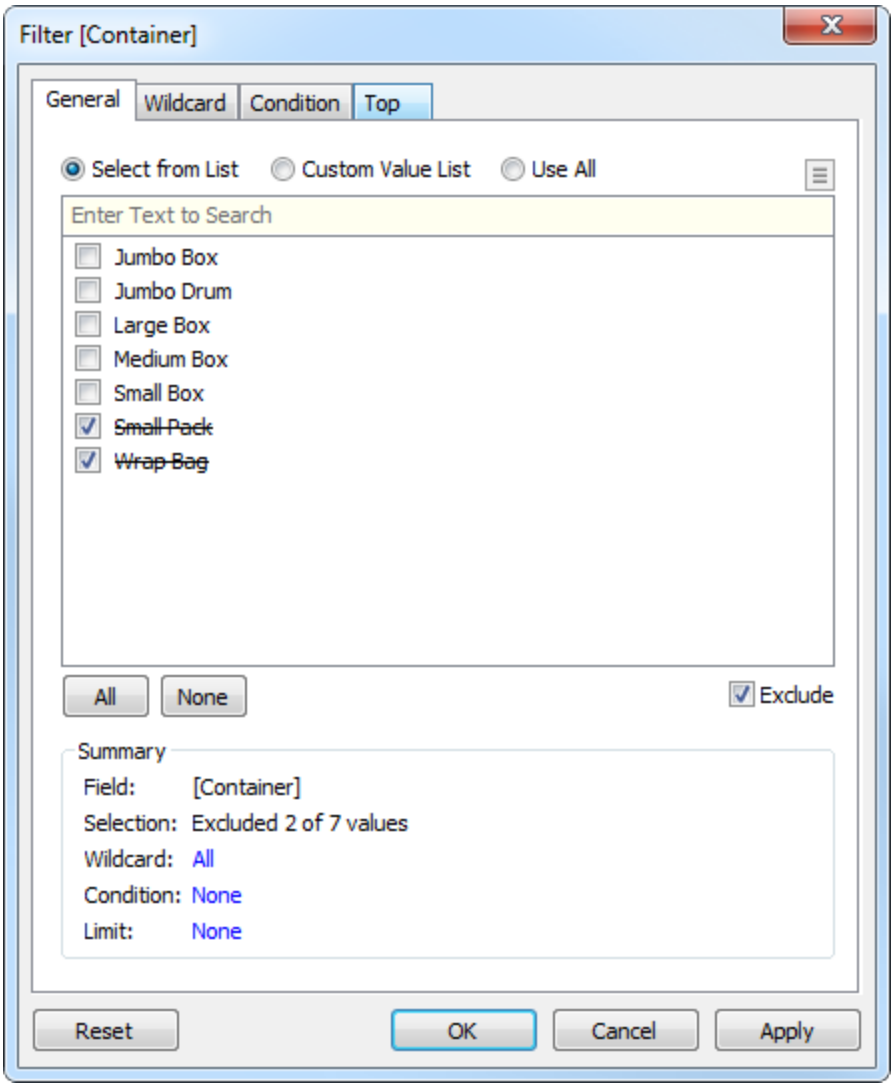

3. The view updates to only show orders that were not shipped in a Small Pack or Wrap Bag.

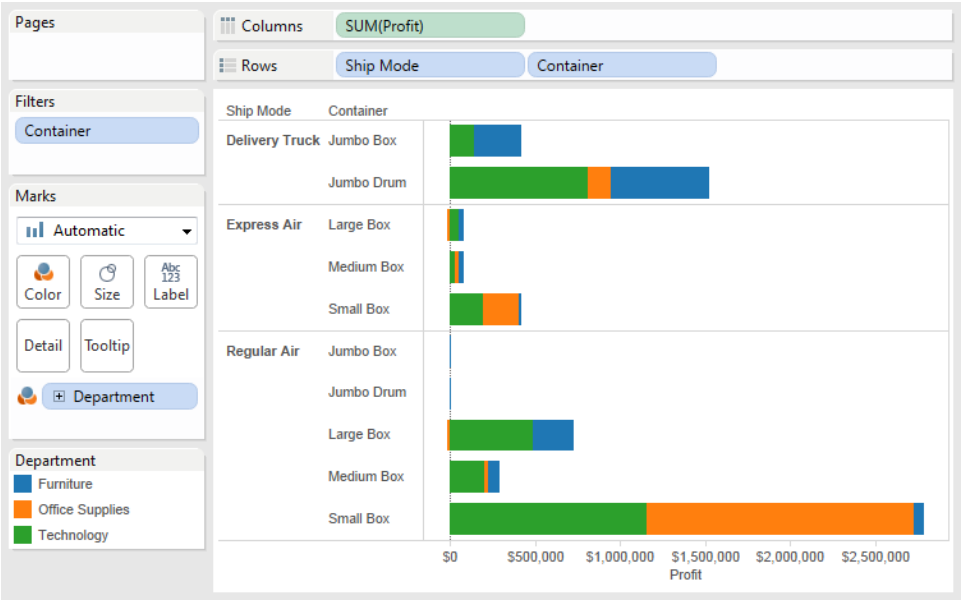

4. Now let's refine the filter on Container by adding a limit. Right-click the Container field on the Filters shelf and select **Filter**. The Filter dialog box opens again. Leave the selections as they are.
5. Switch to the **Top** tab and select By Field. Select Top 3 from first two drop-down lists. 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, which mean 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 Regular Air.

   Right-click the **Regular Air** row header and select **Exclude**. The Regular Air ship mode is removed from the view.

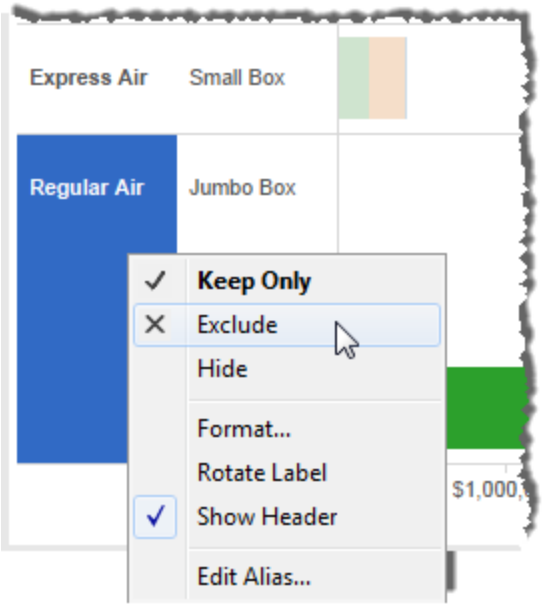

7. Finally, let's filter the Department dimension. Drag the Department dimension to the Filters shelf.
8. In the Filter dialog box, clear the checkbox for Office Supplies.

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.

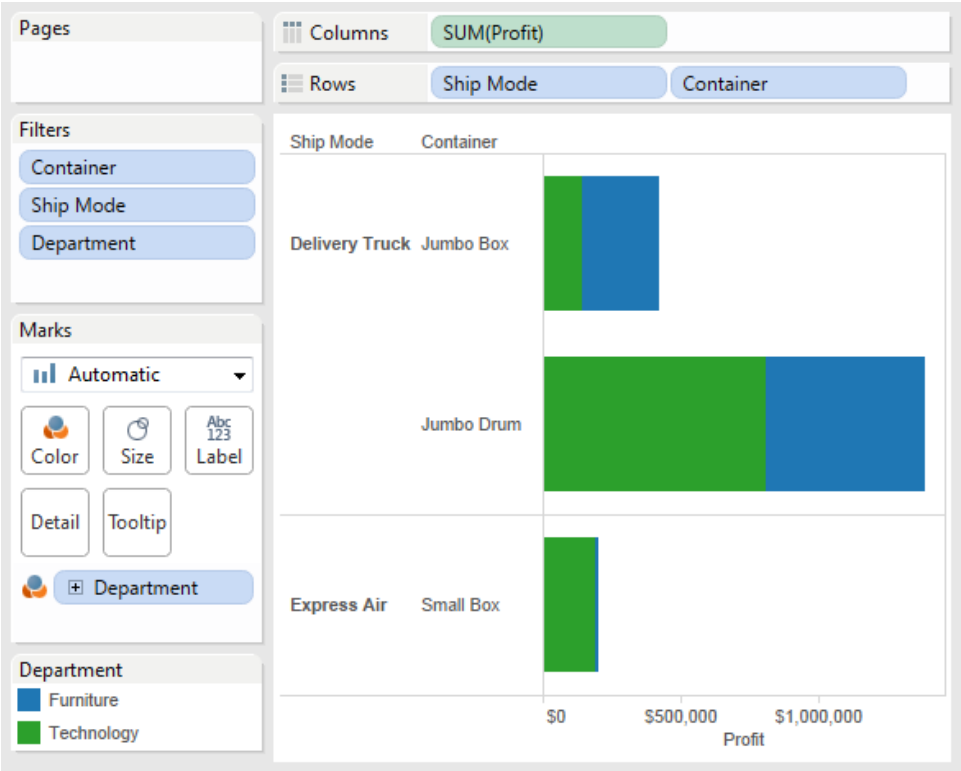

Filtering Measures

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.

> 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 apply a filter to the set.

Basic Quantitative Filters

1. Open the Filter dialog box by dragging a measure to any shelf.

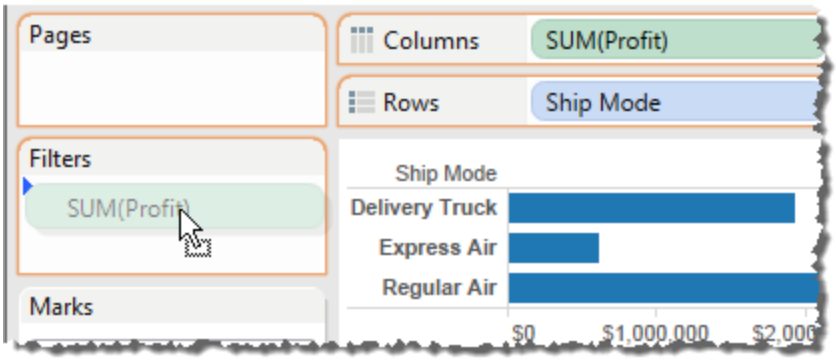

2. The Filter Field dialog box opens where you need to specify an aggregation. When finished, click **Next**.

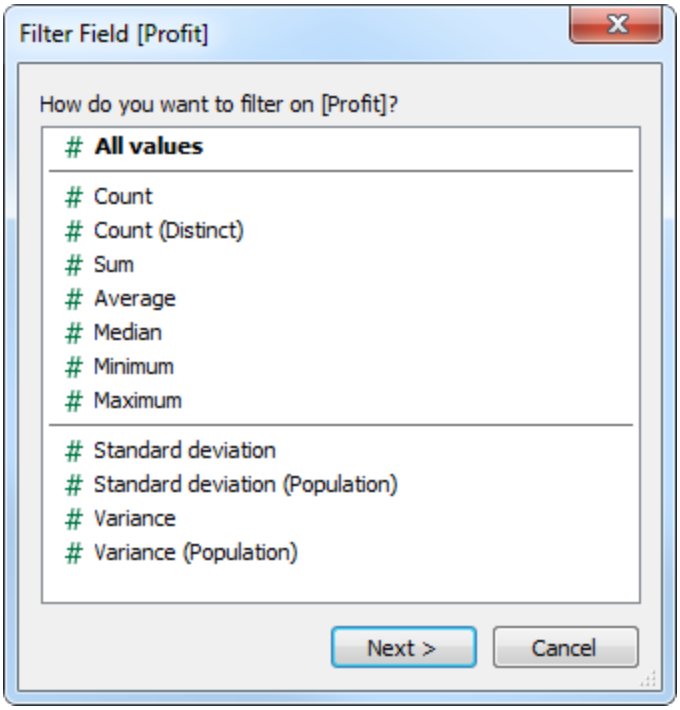

3. The Filter dialog box opens. There are four types of quantitative filters: Range of Values, At Least, At Most, and Special. Each of these types of filters are described below:

   - **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.

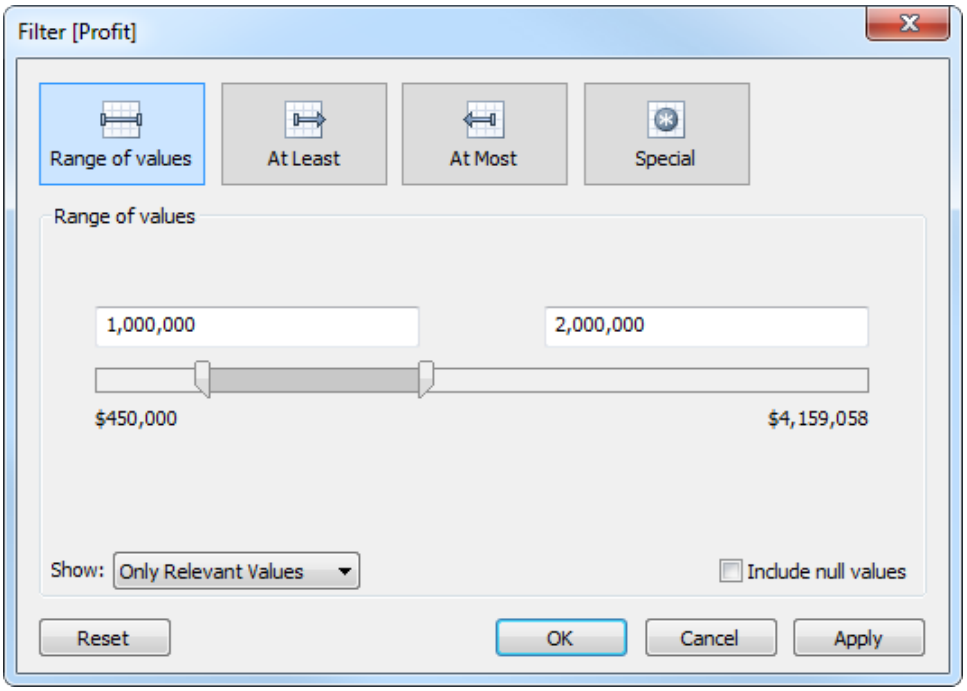

- **At Least** - Include all values that are greater than or equal to a specified minimum value. This type of filter is useful when the data changes often so specifying an upper limit may not be possible.

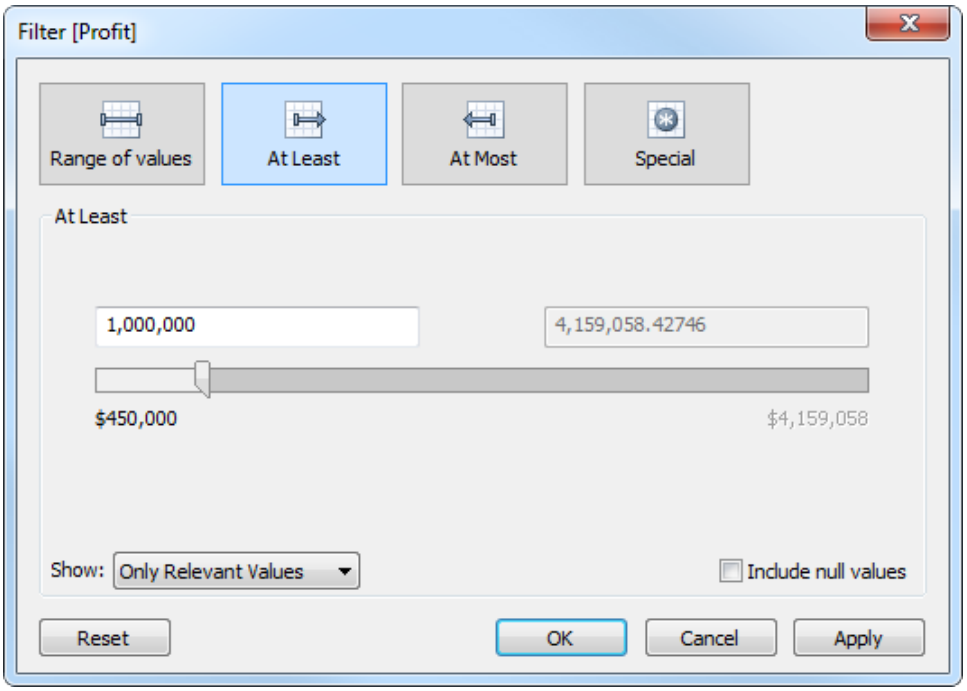

- **At Most** - Include all values that are less than or equal to a specified maximum value. This type of filter is useful when the data changes often so specifying a

lower limit may not be possible.

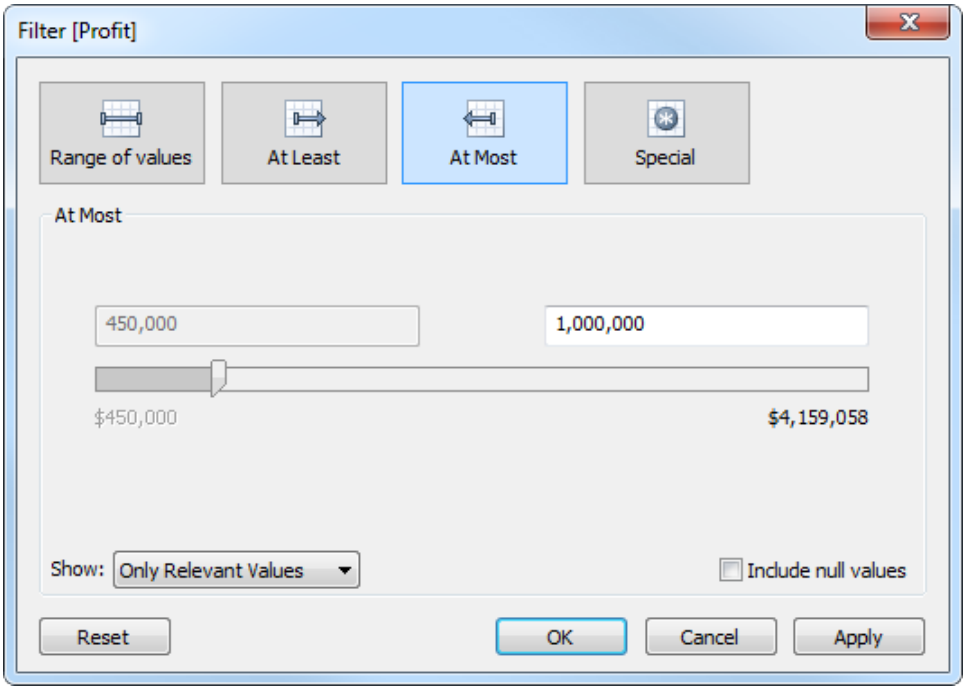

- **Special** - This special type of filter helps you filter on Null values. Include only **Null values**, **Non-null values**, or **All Values**.

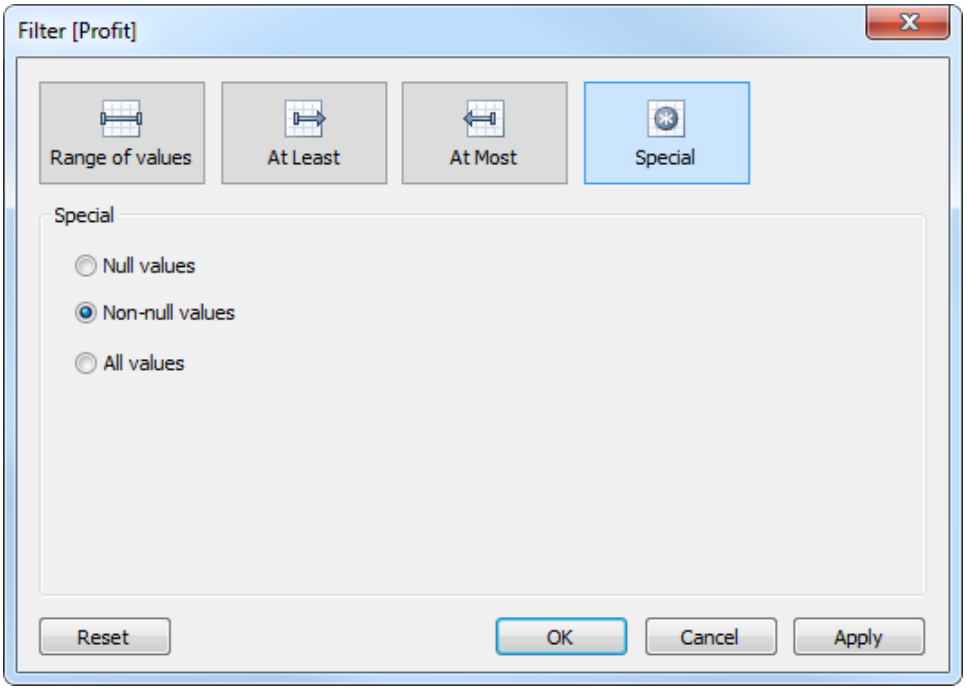

4. When finished defining the filter click **OK**.

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.

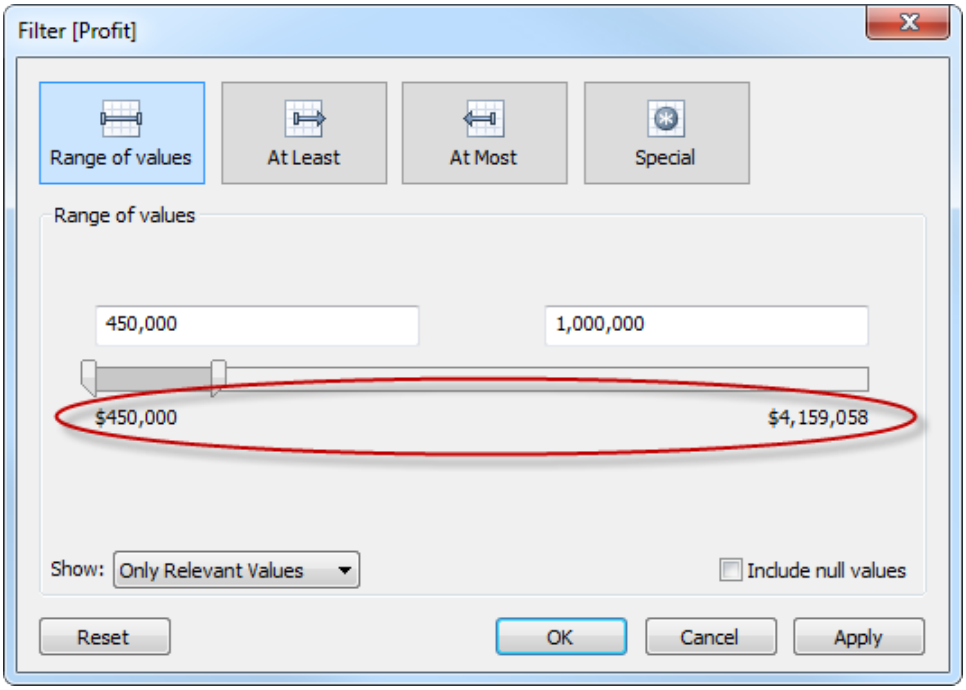

These minimum and maximum values are affected by the other filters set on the view. For example, a database may include records with sales ranging from $0 to $89K. If you created a filter on the Sales field the minimum and maximum values shown in the filter dialog box would indicate this range. However, 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 don't change how the filter will be applied to the view.

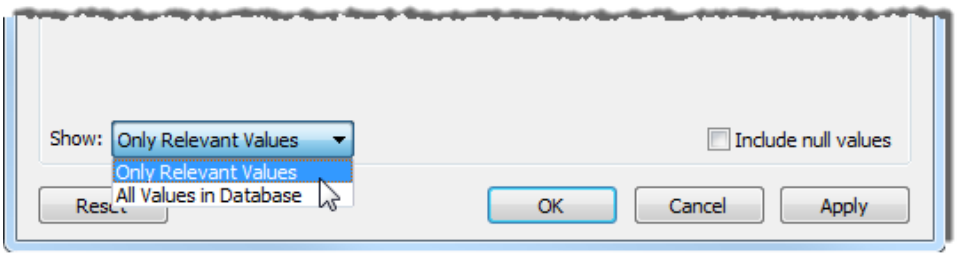

Example – Filtering Measures

This example filters a text table using an aggregated measure, and then filters the table using the same measure in a disaggregated state.

1. Create the initial view using the Superstore sample data source. The text table is shown below.

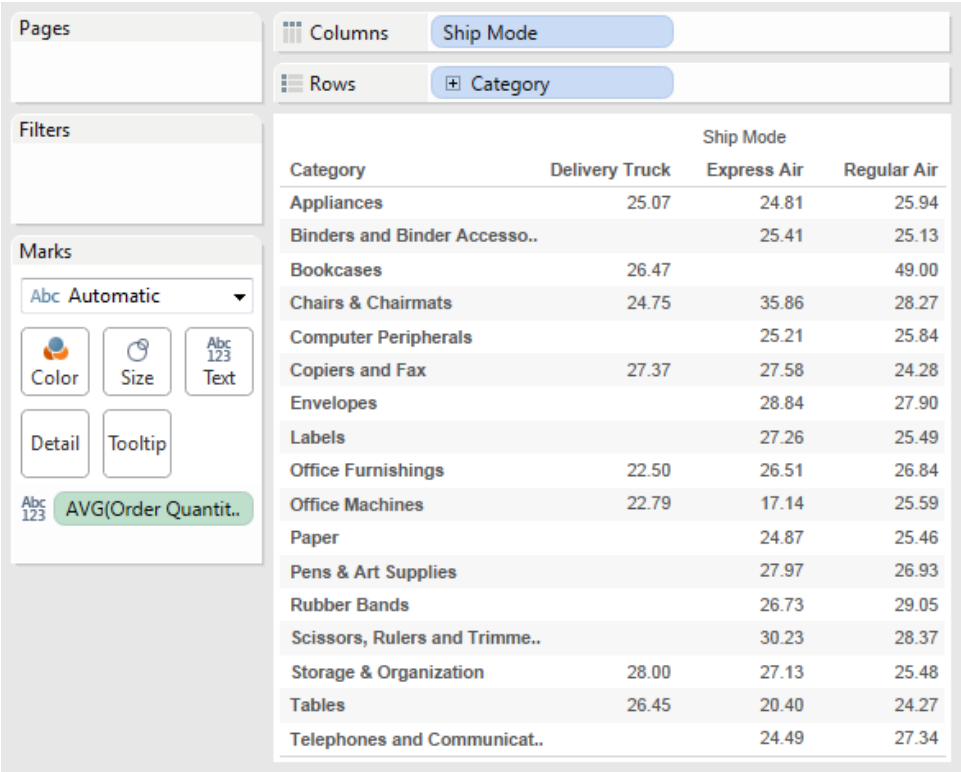

	Ship Mode		
Category	Delivery Truck	Express Air	Regular Air
Appliances	25.07	24.81	25.94
Binders and Binder Accesso..		25.41	25.13
Bookcases	26.47		49.00
Chairs & Chairmats	24.75	35.86	28.27
Computer Peripherals		25.21	25.84
Copiers and Fax	27.37	27.58	24.28
Envelopes		28.84	27.90
Labels		27.26	25.49
Office Furnishings	22.50	26.51	26.84
Office Machines	22.79	17.14	25.59
Paper		24.87	25.46
Pens & Art Supplies		27.97	26.93
Rubber Bands		26.73	29.05
Scissors, Rulers and Trimme..		30.23	28.37
Storage & Organization	28.00	27.13	25.48
Tables	26.45	20.40	24.27
Telephones and Communicat..		24.49	27.34

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.

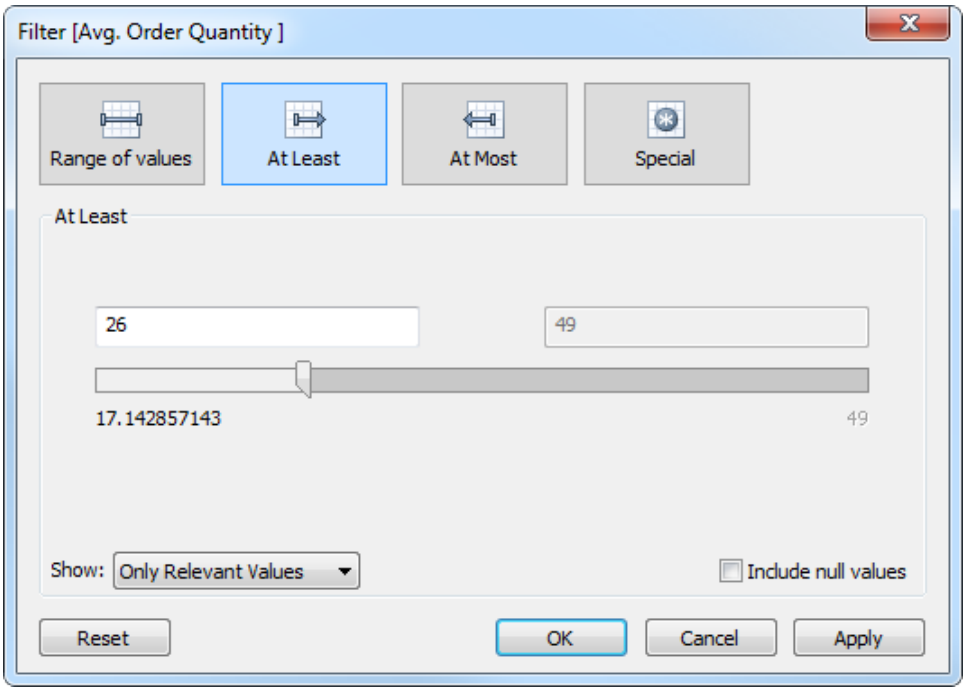

3. When finished, click **OK**.

   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, the Copiers & Faxes category shipped by Regular Air is removed from the view because the average order quantity is less than 26, while Copiers & Faxes shipped by Delivery Truck and Express Air remains in the view because the average order quantity is greater than 26.

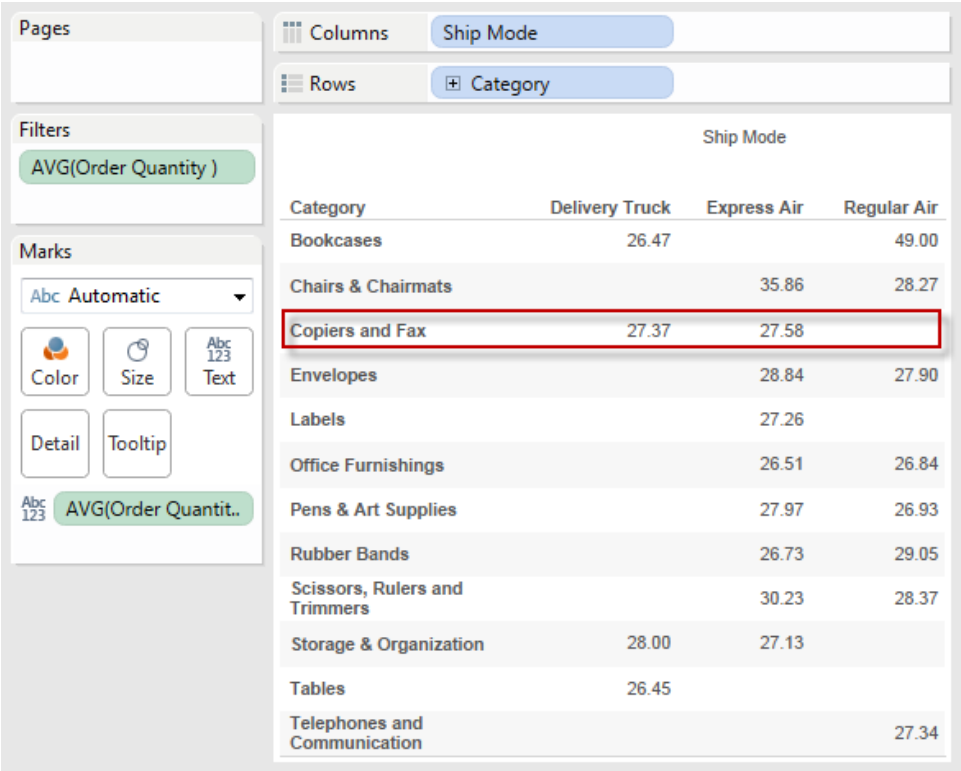

4. Now let's filter the same view using a disaggregated measure. Suppose you want to filter the view using the disaggregated **Order Quantity** measure. To do this, select **Dimension** on the context menu of the AVG(Order Quantity) field on the Filters shelf.

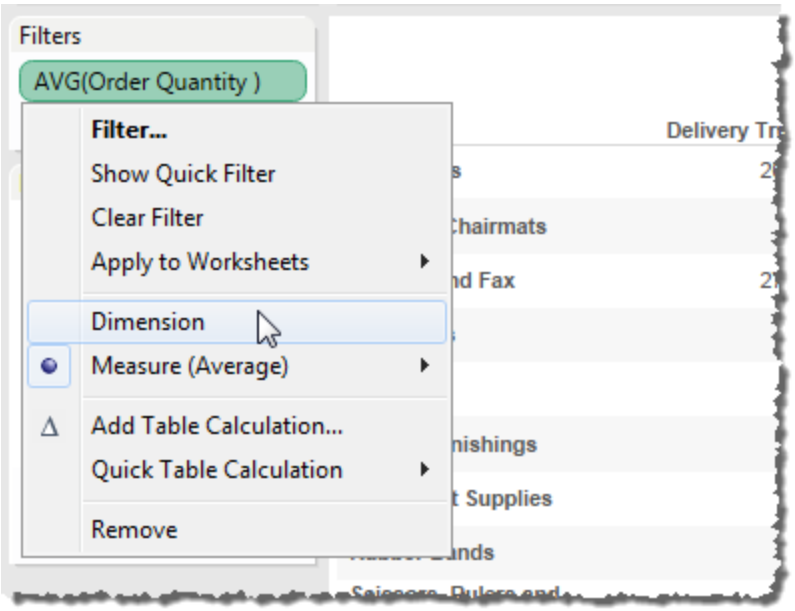

The **Filter** dialog box is shown below. It displays the limits of the individual rows for the Order Quantity measure. Specify a new lower limit of `26`.

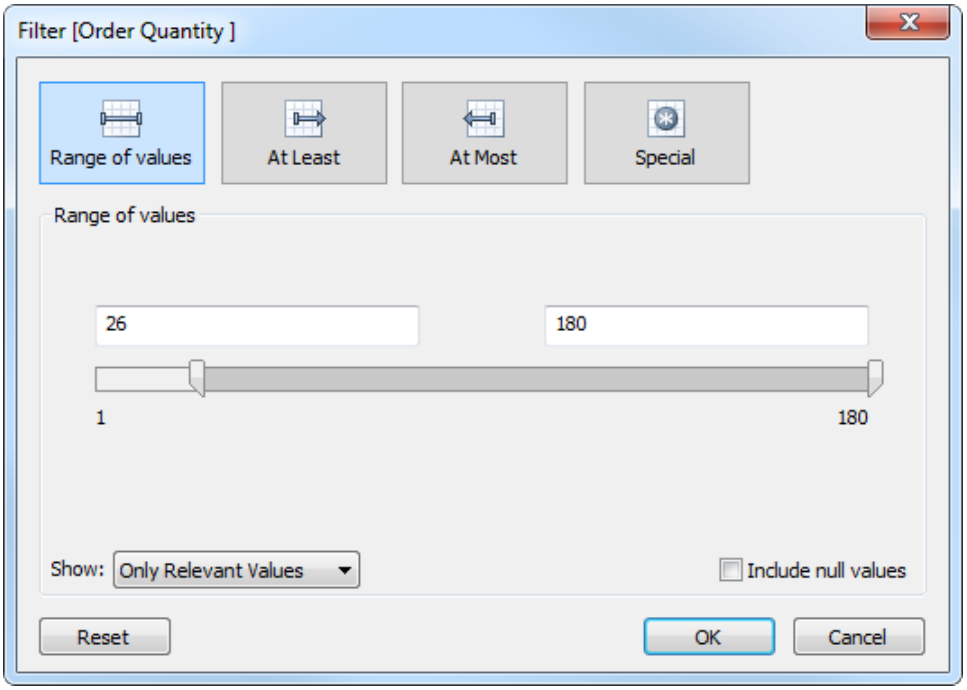

The filtered data view is shown below. Notice that the numbers are very different from the original, unfiltered view. This is because Tableau excludes each row in the data source that has an order quantity that is less than 26, and then aggregates the remaining rows as an average.

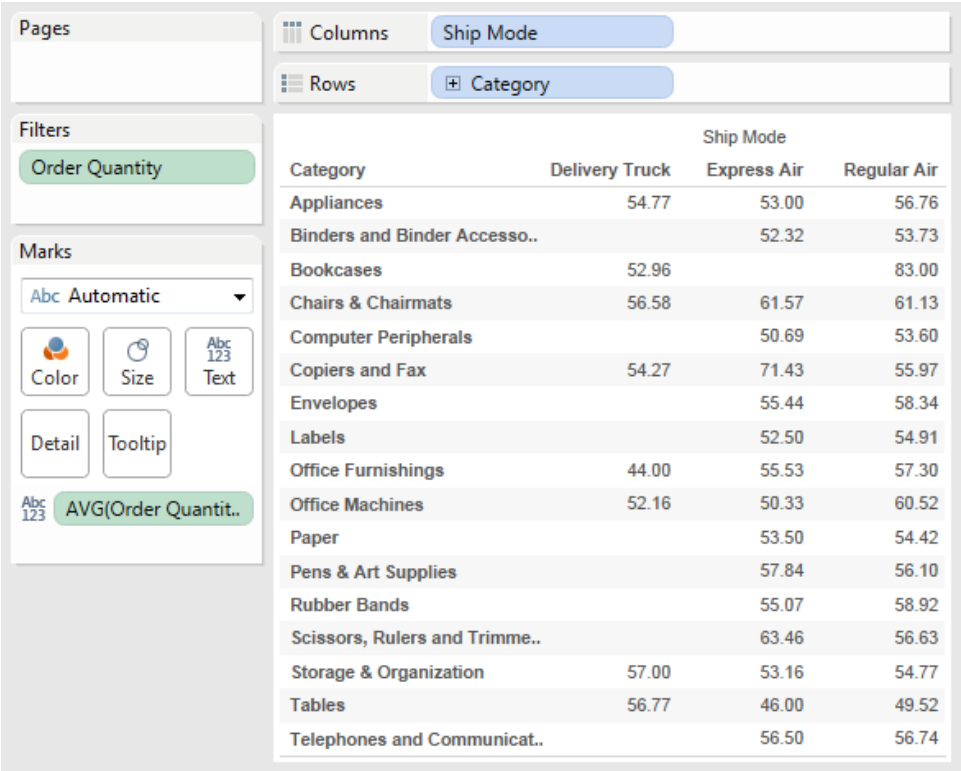

Filtering Dates

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.

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.

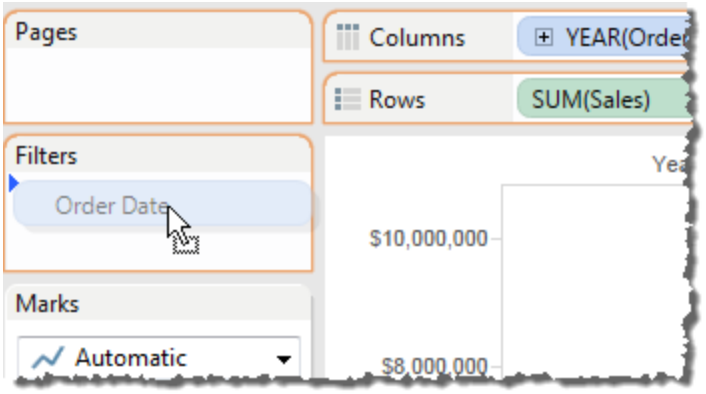

2. In the Filter Field dialog box, select **Relative date** and then click **Next**.

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 Quarters as the time unit.

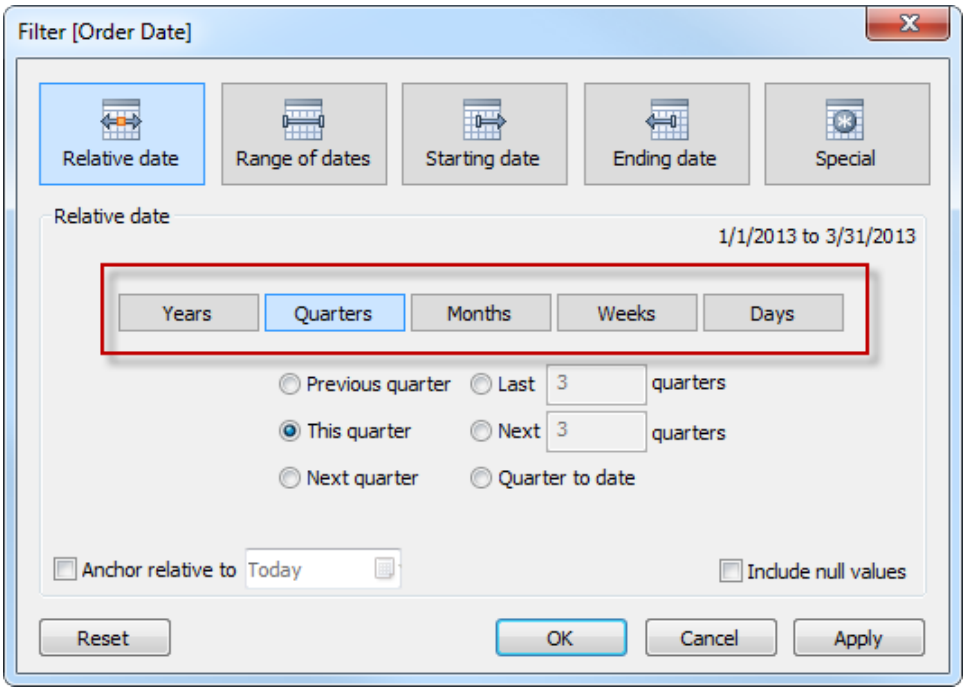

4. Use the rest of the controls to define the date filter. You can select from a variety of common options including current, previous, and next. By default, the filter is relative to today. To make the filter relative to an alternate date, select the **Anchor relative to** option in the bottom left corner 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.

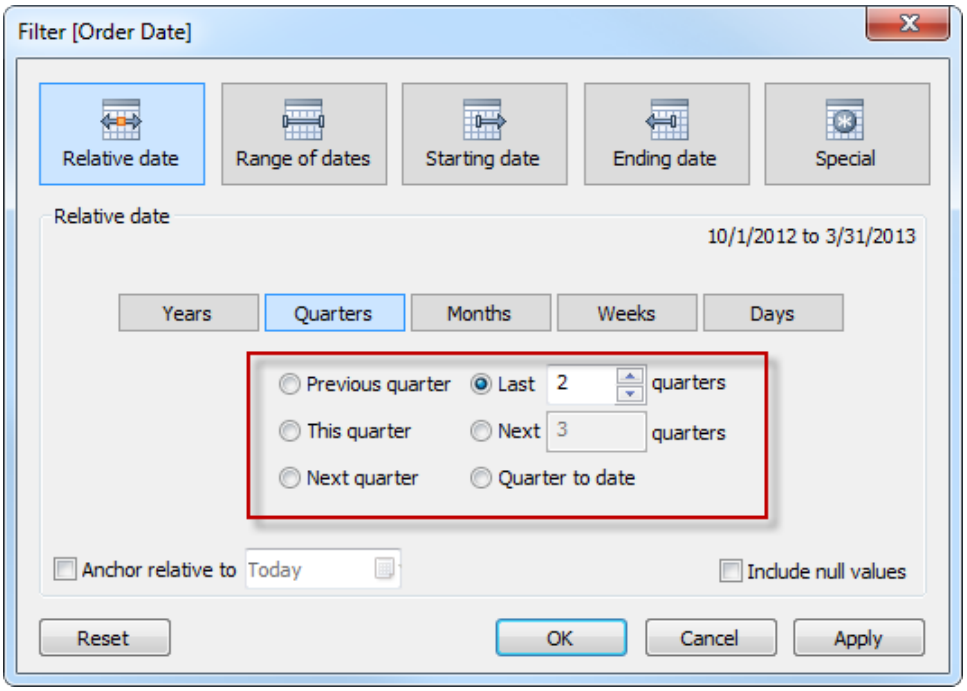

5. When finished, click **OK**.

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. 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.

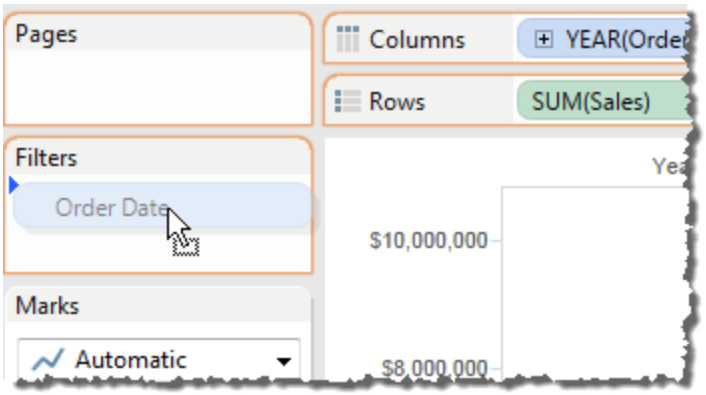

2. In the Filter Field dialog box, select **Range of Dates** and then click **Next**.

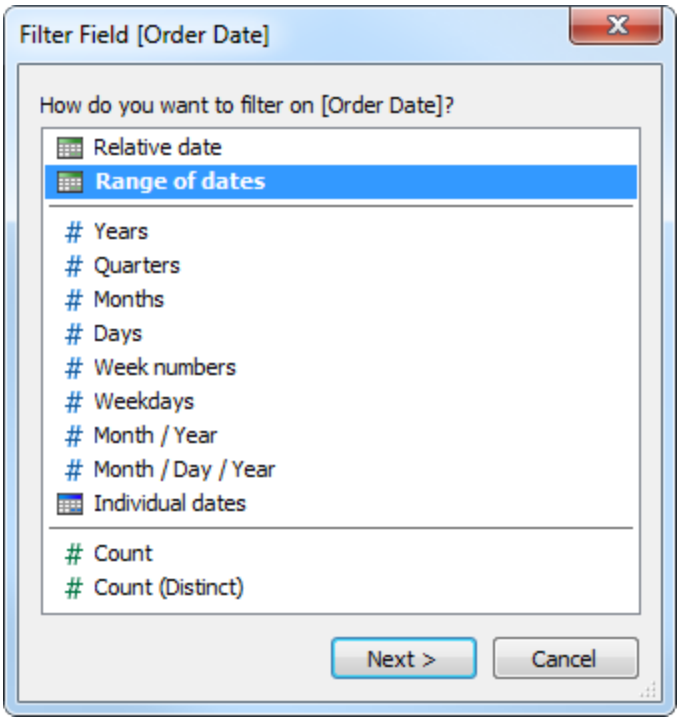

3. The Filter dialog box opens showing the Range of Dates options. Use the slider or the drop-down date controls to select minimum and maximum dates for the range you want to include. The range is inclusive, which means that the minimum and maximum dates are included in the filter.

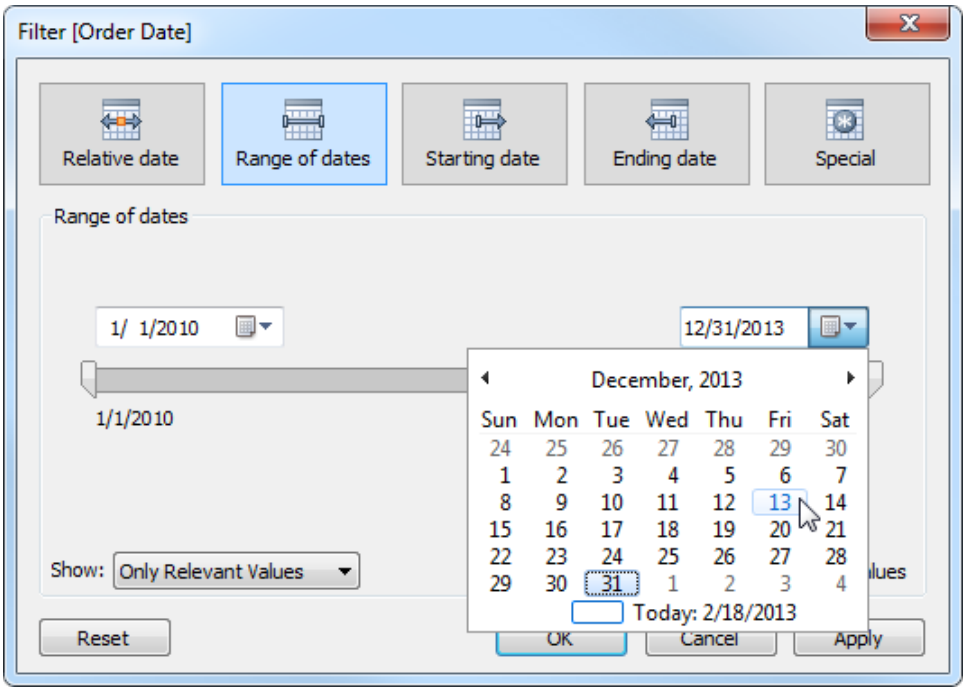

4. When finished, click **OK**.

> If the field also includes Time you can select the Show Times option to further refine your filter range.

Other Types of Date 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.

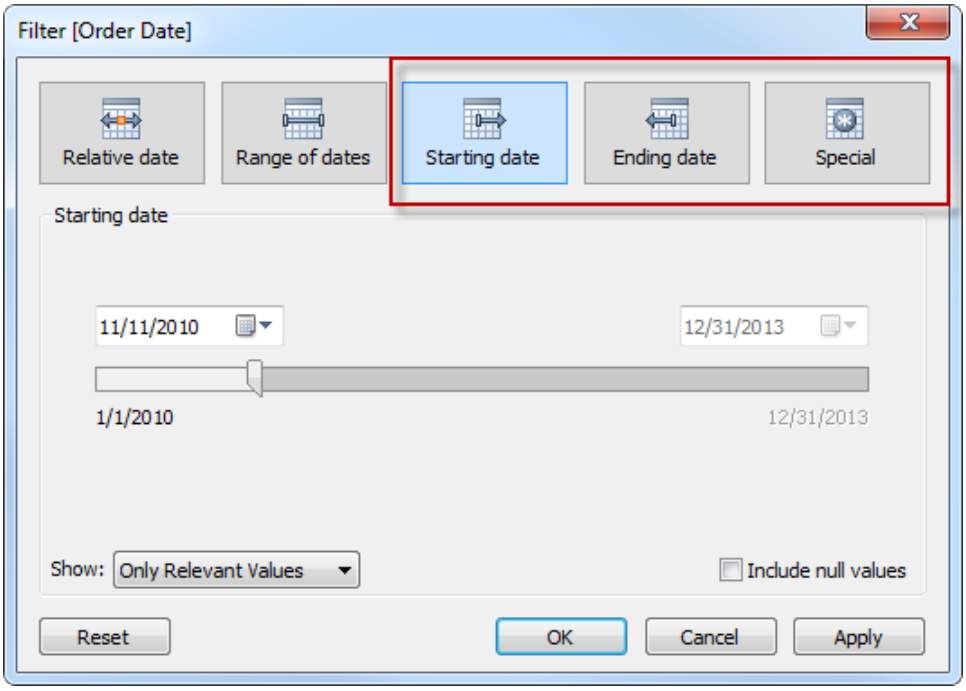

Discrete Date Filters

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.

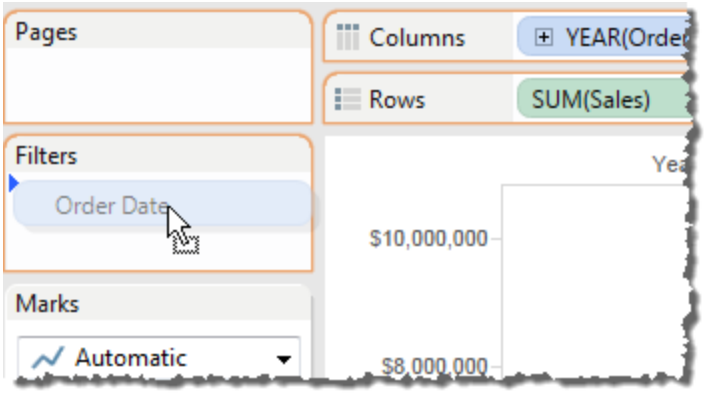

2. In the Filter Field dialog box, select a date level or select **Individual dates** and then click **Next**.

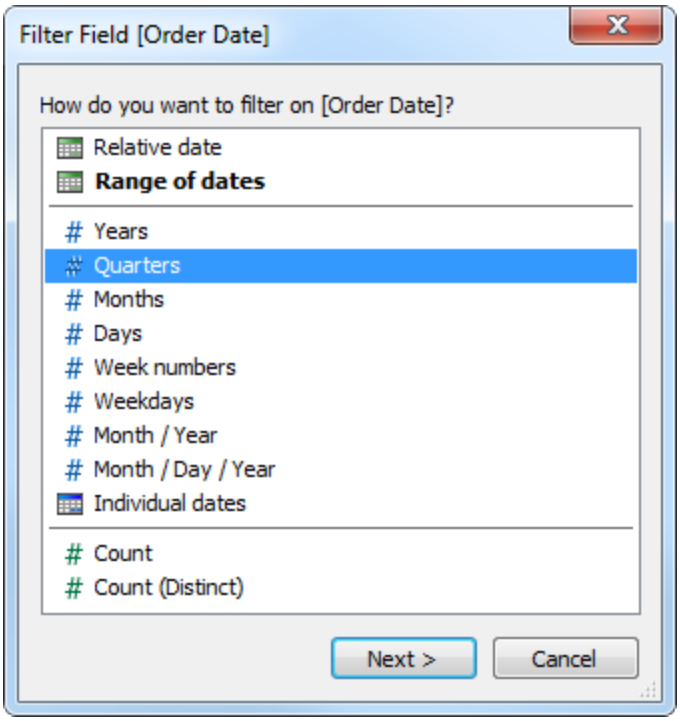

3. In the Filter dialog box, select the dates you want to include.

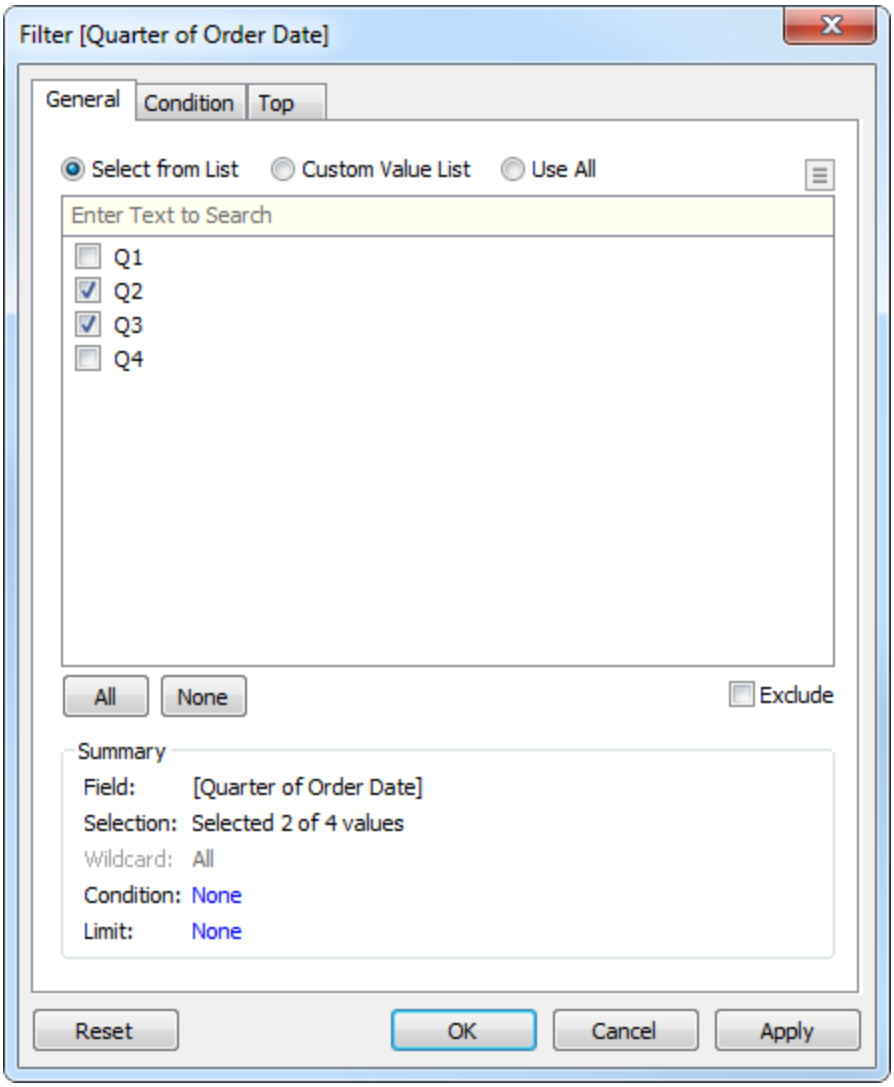

4. When finished, click **OK**.

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 sample data source. Place Order Date on to the Columns shelf and select Exact Date as the aggregation. Then place Profit onto the Rows shelf.

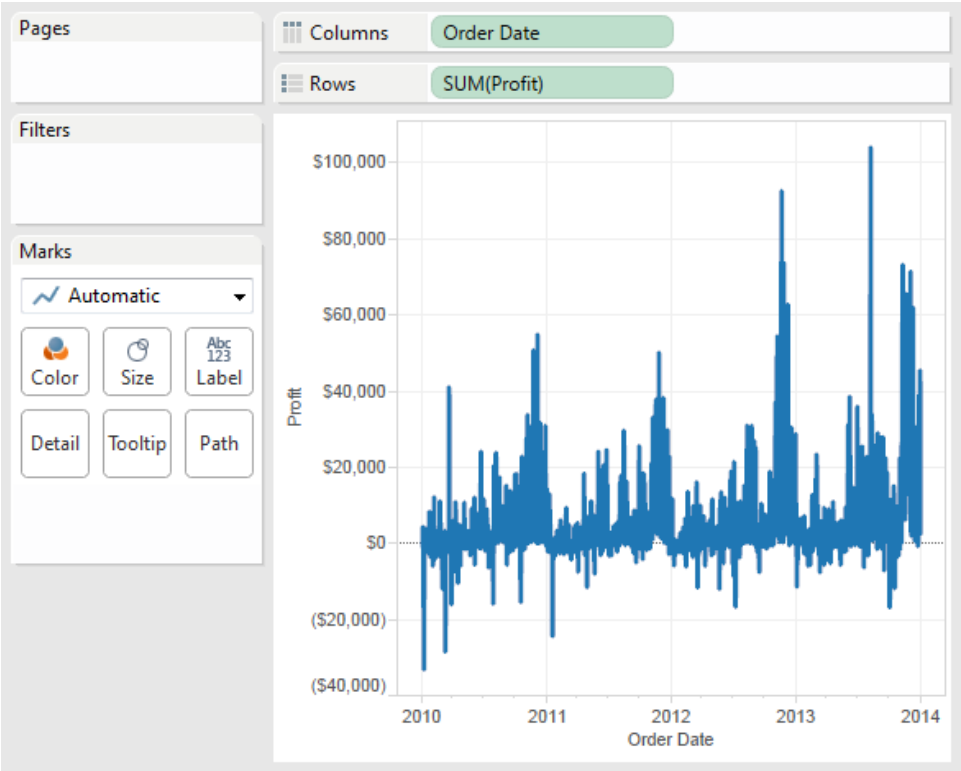

2. Now let's filter the view to include only orders that were place between August 2, 2011 and May 1, 2009. To create this filter drag the Order Date field to the Filters shelf and select Range of Dates in the Filter Field dialog box. Then click Next.

   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.

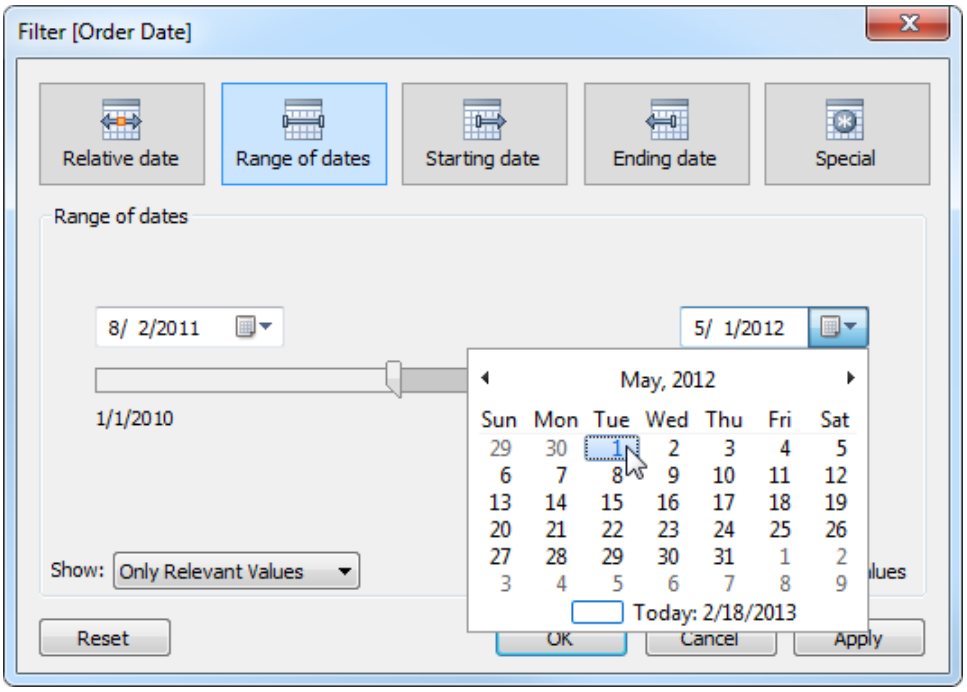

The filtered view is shown below.

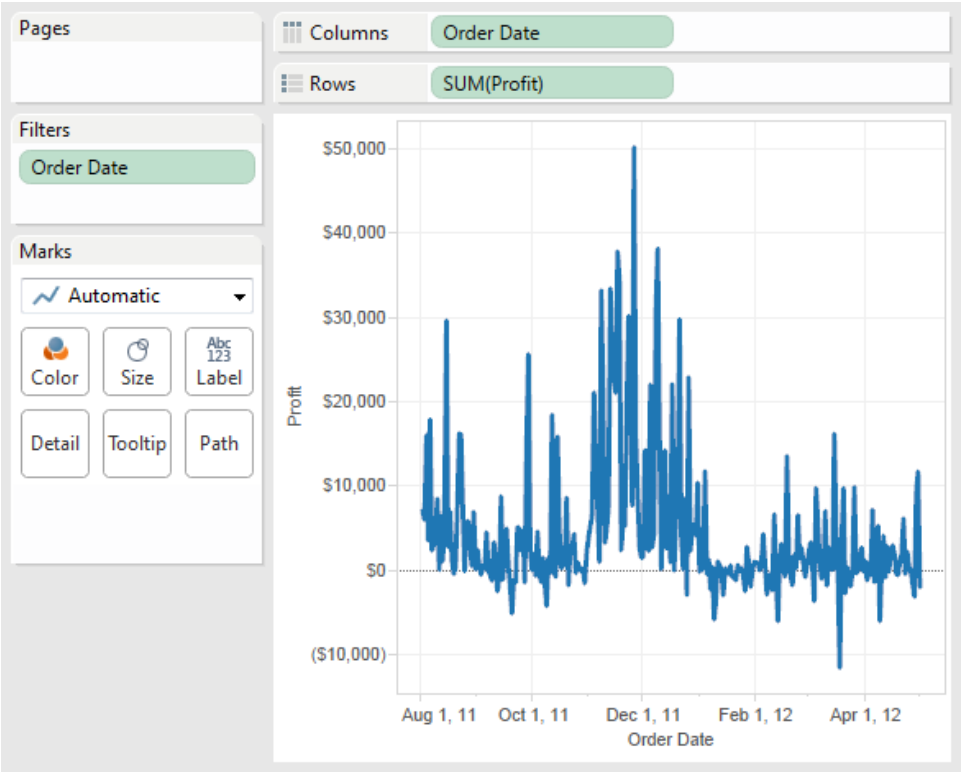

Using Quick Filters

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.

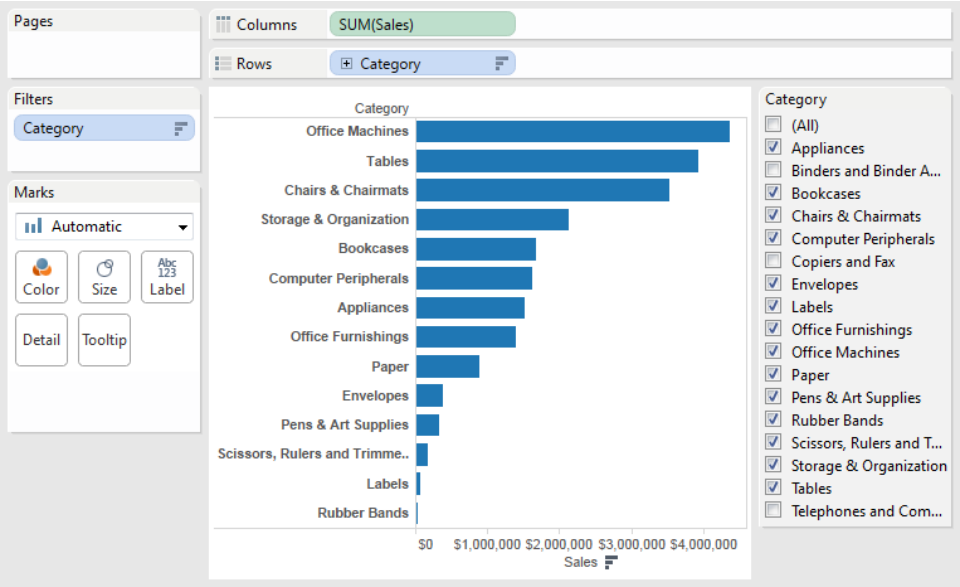

> You can add a quick filter for a field that is not current used in the view. To do so, right-click the field in the Data window and choose **Show Quick Filter**. The field is automatically added to the **Filters** shelf.

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.

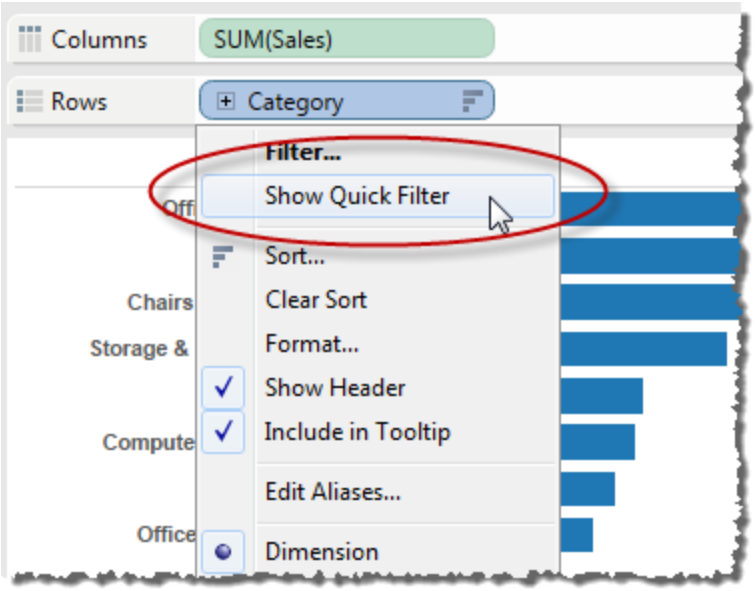

Quick Filter Options

After you've turned on a quick filter there are many different options that let you control how the filter works and appears. You can access these options using the card menu in the upper right corner of the quick filter card.

Note: To access the quick filter card, go to **Analysis** and click **Quick Filters**. Select which item you would like to filter from the menu.

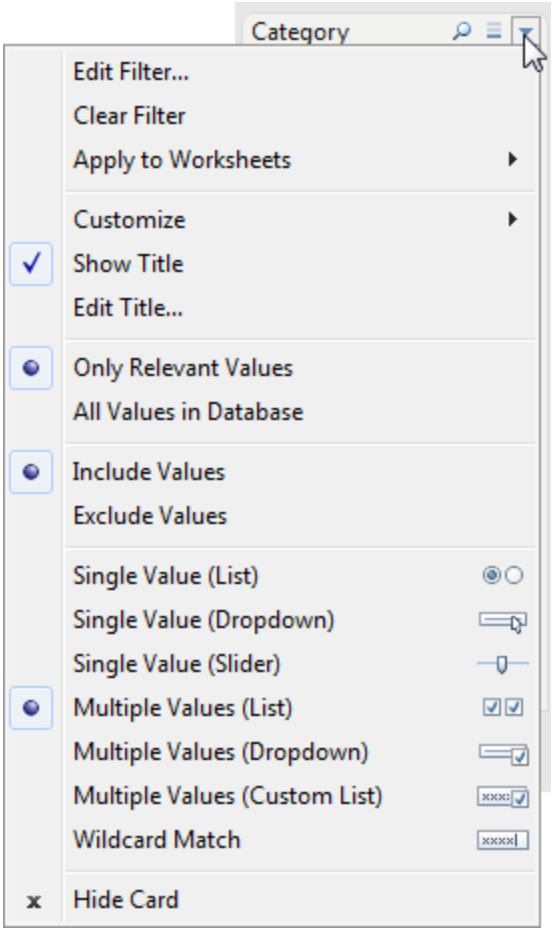

Some options are available for all types of filters and others depend on whether you're filtering a Categorical field (dimensions) or a Quantitative field (measures). Finally, you can customize how quick filters display on the sheet, in dashboards, or when published to Tableau Server.

General Quick Filter Options

- **Edit Filter** - 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.
- **Apply to Sheets** - Specifies whether the filter should apply to only the current sheet or be shared across multiple sheets. See **Sharing Filters** on page 377.
- **Customize** - Select from a list of customization options. See **Customizing Quick Filters** on page 375.
- **Show Title** - Specifies whether to show the title on the quick filter card.
- **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.
- **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.
- **Include** - When this option is selected, the selections in the quick filter are included in the view.
- **Exclude** - When this option is selected, the selections in the quick filter are excluded from the view.
- **Quick Filter Modes** - Select from several different quick filter modes. The modes differ depending on whether the filter is for a dimension or measure. See **Categorical Quick Filter Options** below and **Quantitative Quick Filter Options** on the next page.
- **Hide Card** - Hides the quick filter card but does not remove the filter from the Filters shelf.

Categorical Quick Filter Options

- **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.

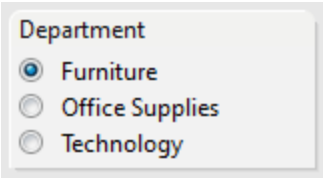

- **Single Value (Dropdown)** - Displays the values of the quick filter in a drop-down list where only a single value can be selected at a time.

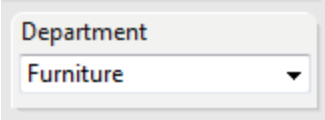

- **Single Value (Slider)** - Displays the values of the quick filter along the range of a slider. Only a single value can be selected at a time. This option is useful for dimensions that have an implicit order such as dates.

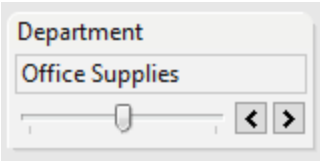

- **Multiple Values (List)** - Displays the values in the quick filter as a list of checkboxes where multiple values can be selected.

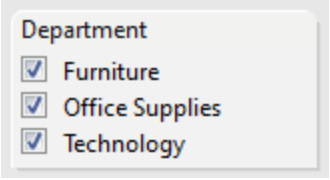

- **Multiple Values (Dropdown)** - Displays the values of the quick filter in a drop-down list where multiple values can be selected.

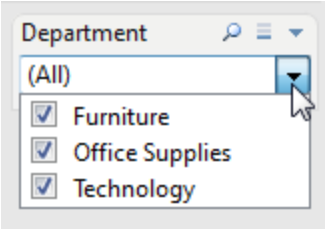

- **Multiple Values (Custom List)** - Displays a text box where you can type a few characters and search for the value. Alternatively, you can type or paste a list of values into the text box to create a custom list of values to include.

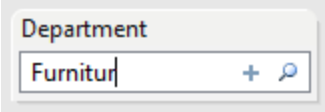

- **Wildcard Match** - Displays a text box where you can type a few characters. All values that match those characters are automatically selected. You can use the asterisk character as a wildcard character. For example, you can type “tab*” to select all values that begin with the letters “tab”. Pattern 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.

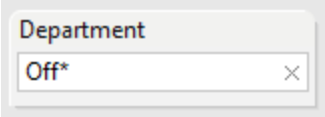

Quantitative Quick Filter Options

- **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.

Data bars only show in filters where the filtered field is also used in the view (e.g., on Columns, Rows, or on the Marks card, and son on) and are at the same aggregation level as the field on the Filters shelf. For example, a quick filter on SUM(Sales) will only display data bars if the SUM(Sales) field is used in the view. It won't show if AVG(Sales) is used in the view. Even though in both scenarios, the filtered field, Sales is used in the view; in the latter case the aggregation is different than the aggregation of the filter.

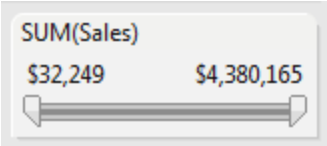

- **At Least/Starting Date**- shows a single slider with a fixed minimum value. Use this option to create a filter using an open ended range.

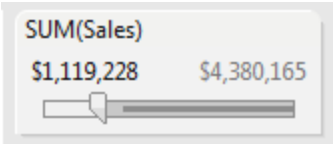

- **At Most/Ending Date** - shows a slider with a fixed maximum value. Use this option to create a filter using an open ended range.

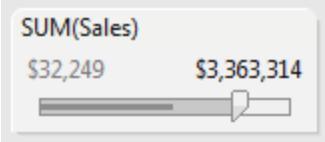

- **Relative to Now** - shows a control where you can define a dynamic date range that updates based on when you open the view. The option is only available for filters on continuous date fields.

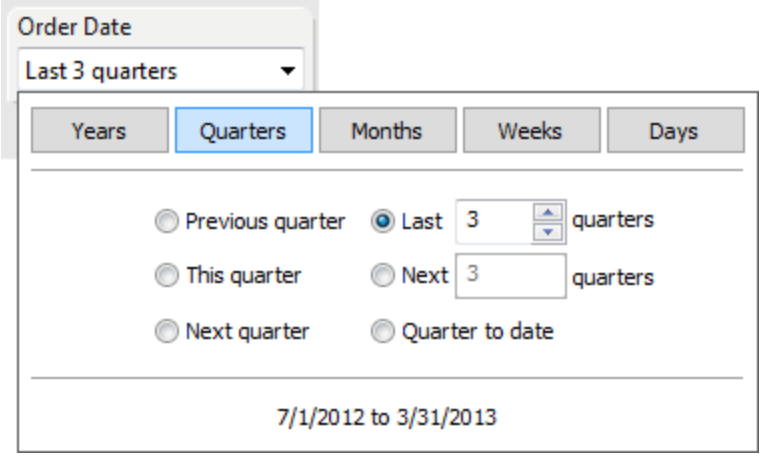

- **Browse Periods** - shows common date ranges such as past day, week, month, three months, one year, and five years. This option is only available for filters on continuous date fields.

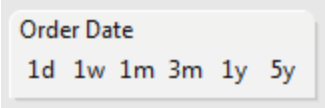

Customizing Quick Filters

You can control how a quick filter appears on the sheet, in dashboards, or when published to Tableau Server. Customize quick filters by selecting Customize on the quick filter card menu.

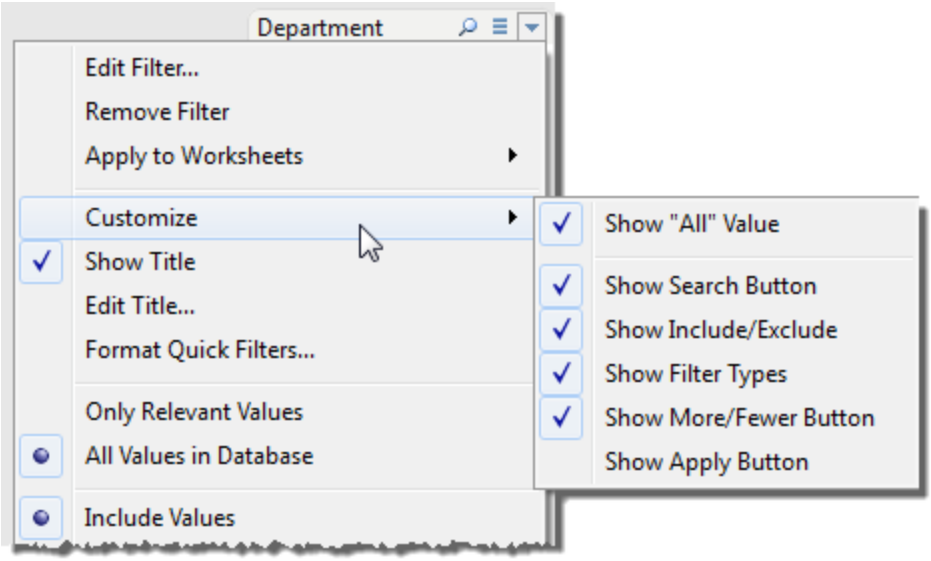

> Not all of the following options are available for views published to Tableau Server.

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 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 Apply Button** - toggles whether to show the Apply button at the bottom of the quick filter. When shown, changes to the filter are only applied after you click the button. Pending changes are indicated with a green color. This option is only available in multiple values lists and dropdowns.
- **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.

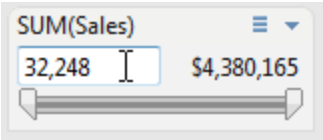

- **Show Slider** - controls whether the slider displays. When this option is cleared, the quick filter only displays the readouts.
- **Show Null Controls** - shows a drop-down list that lets you control how the filter handles null values. You can select from the following options:
    - Values in Range - 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.

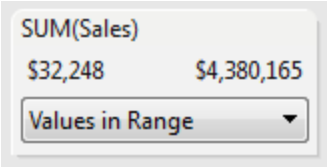

Searching Quick Filters

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 a keyword or search term. 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 “BI*” will find all values that start with the characters B and I. Search is not case sensitive.

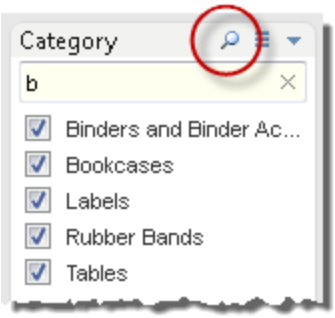

Search is also available when using a drop-down quick filter. Use the search box at the top of the drop-down list to type keywords or search terms. The results are shown directly below where you can select or deselect them as needed.

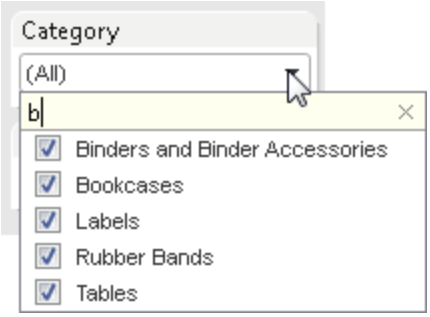

Sharing Filters

When you add a filter to a worksheet, by default it only applies to the current worksheet. Sometimes you may want to apply the filter to other worksheets in the workbook. You can select specific sheets to apply the filter to or apply it globally to all worksheets that use the same data source. Filters cannot be applied across multiple data sources. For example, you may have a filter that only includes a specific region or product of interest. Rather than adding this filter every time you create a new worksheet, you can simply create the filter once and then apply it to all worksheets using the same data source.

Select an option on the Filter menu for how to **Apply to Worksheets**.

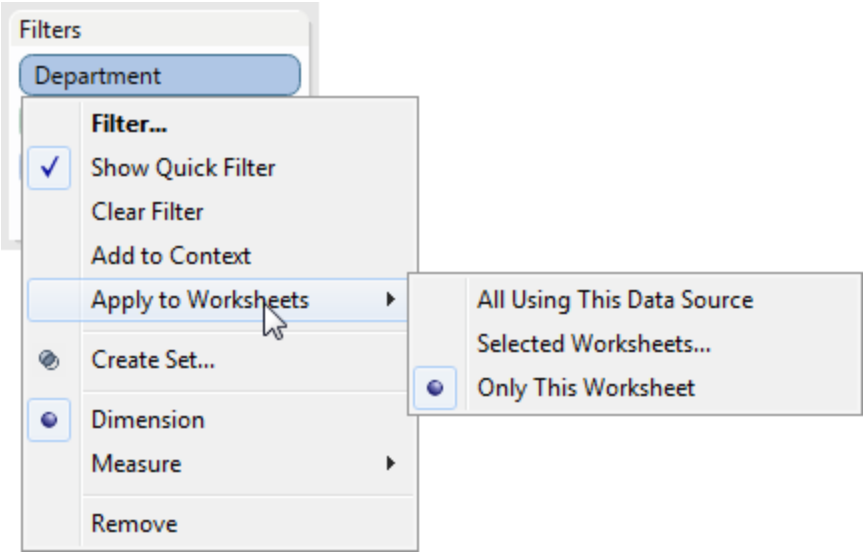

Select one of the following options:

- **All Using This Data Source** - applies the filter to all worksheets that use the current connection. Filters that use this option are global across the workbook. Filters that apply to all worksheets are marked with a data source icon . The filter is automatically created on any new worksheets you create after you drag a field to the view. Any changes you make to the filter affects all of the worksheets.
- **Selected Worksheets** - opens a dialog box where you can select from a list of worksheets that use the same connection. Select the worksheets to apply the filter to. If any of the sheets already contain a filter on the same field, the Apply Filter to Worksheets dialog box will provide details about the filter. If you select the sheet, the current filter will override any existing filter selections.

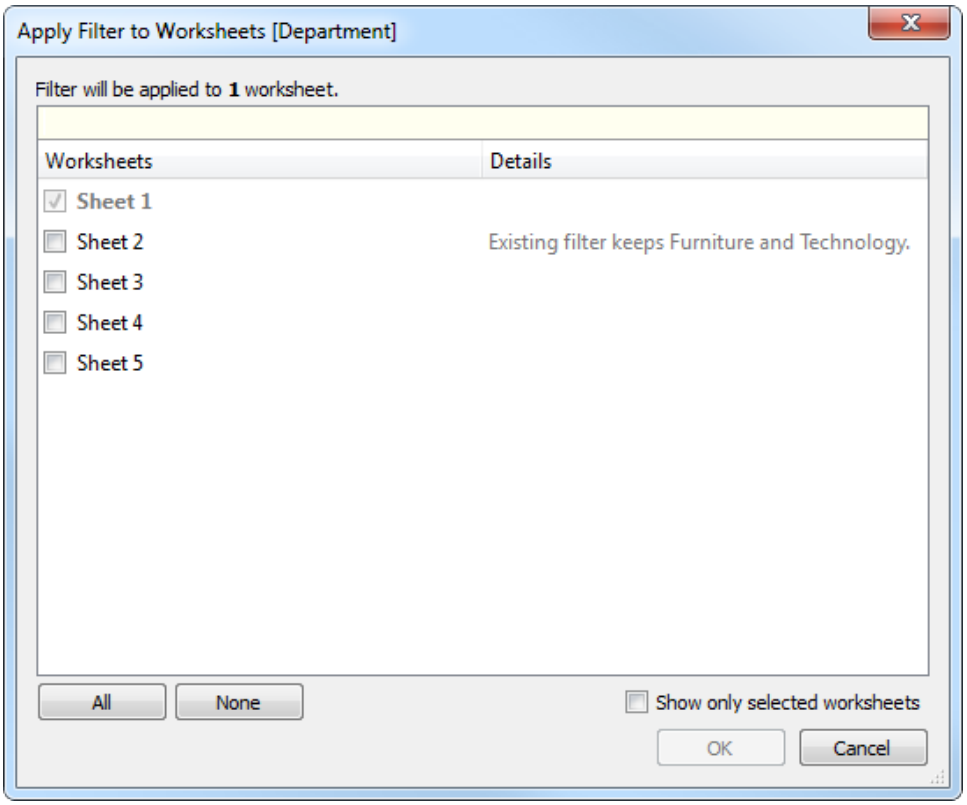

Filters that apply to a selection of worksheets are marked with the worksheet icon . Any changes you make to the filter affect all of the selected worksheets.

- **Only This Worksheet** - only applies to the current worksheet. This option is selected by default when you create new filters. Filters that are local to the current sheet are shown without any additional icons.

If you apply a filter to all sheets or selected sheets and then change the setting to apply the filter to just the current sheet, the filters are not removed from all other sheets. Rather, the filters are disconnected and they are all made local to their respective worksheets. You can go to each sheet and remove the filter or modify selections.

> In previous versions of Tableau Desktop, the **All Using This Data Source** option was called Make Global and the **Only This Worksheet** option was Make Local.

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.

> 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.

Creating Context Filters

To create a context filter, select **Add to Context** from the context menu of an existing categorical filter. 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 gray color on the Filters shelf.
- Cannot be rearranged on the shelf.

As shown below, the **Department** dimension is set to be the context for a view. The **Container** filter is computed using only the data that passes through **Department**.

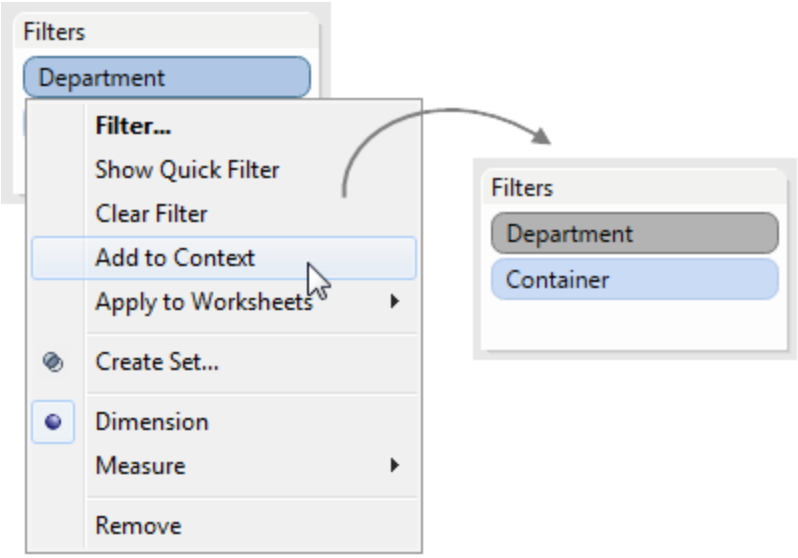

You can modify a context filter by:

- 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 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 - English (Extract) 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.

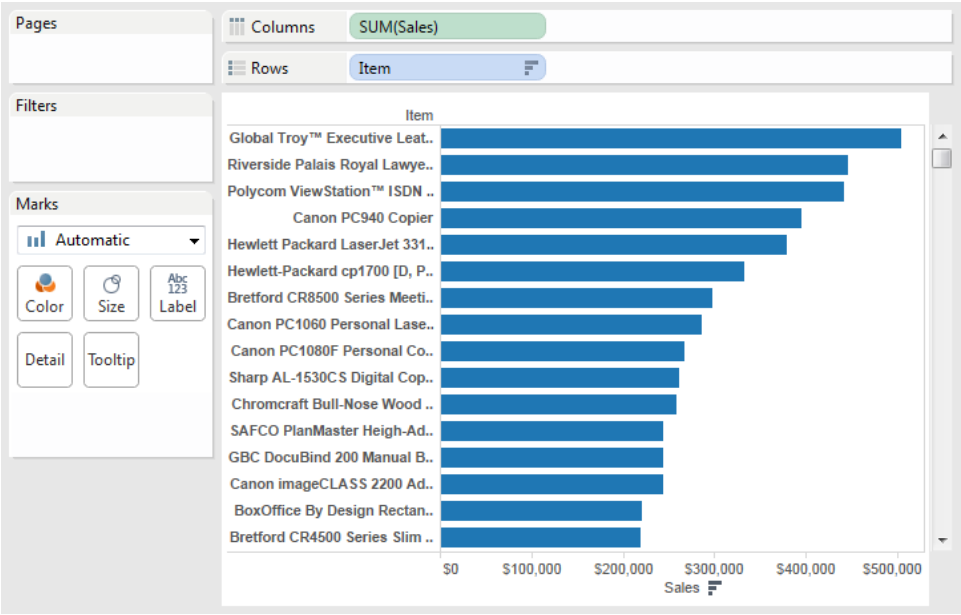

2. Now create a Top 10 filter to just show the top selling products. You can create this filter by dragging the **Item** field to the Filters shelf. In the Filter dialog box, switch to the **Top** tab and define a filter that is Top 10 by Sum of Sales. See **Adding Limits to Filters** on page 344 to learn more about defining a Top N filter.

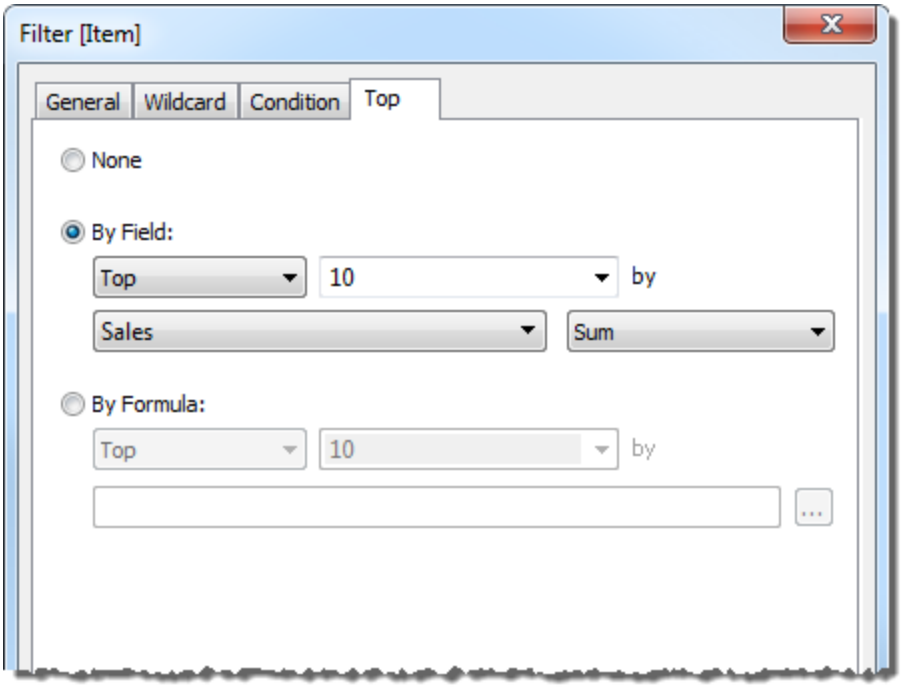

3. When you click **OK**, you'll see that the view is filtered to show the top 10 products in terms of sales.

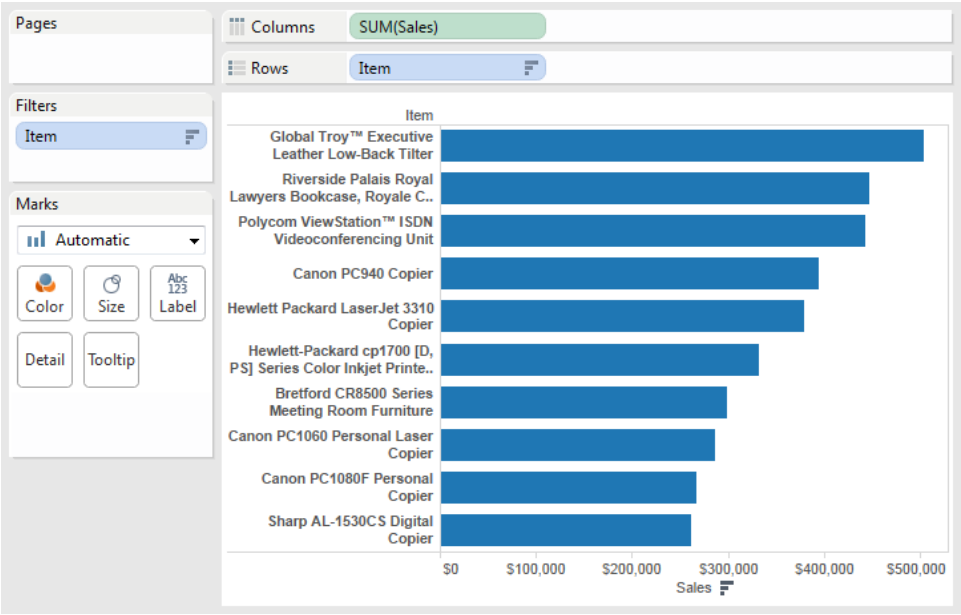

4. Now, let's add another filter to only show only furniture products. Drag the **Department** field to the Filters shelf and select **Furniture**. When finished, click **OK**.

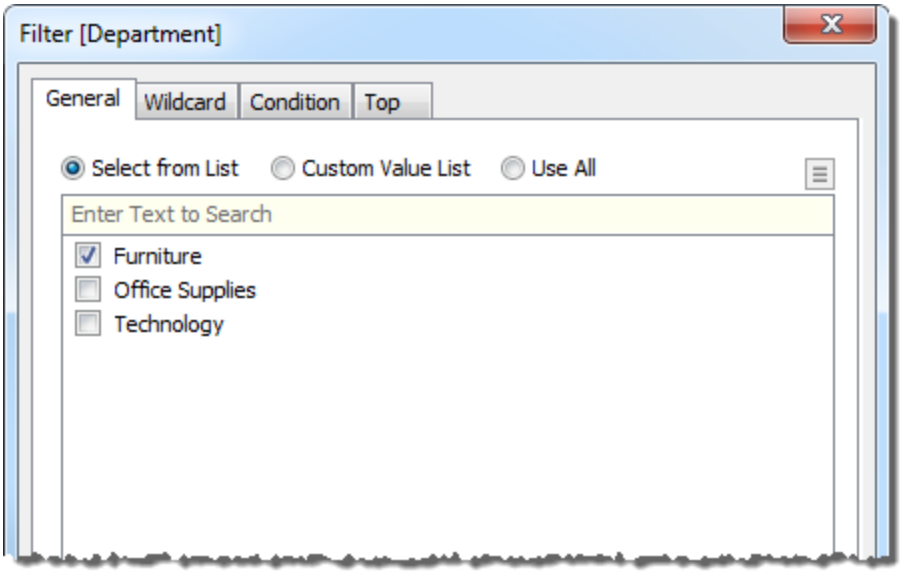

5. The view is filtered but instead of 10 products, it now only shows 3. The reason is because by default all filters are evaluated separately and the view shows the intersection of the results. So this view shows that three of the top 10 overall products are furniture products.

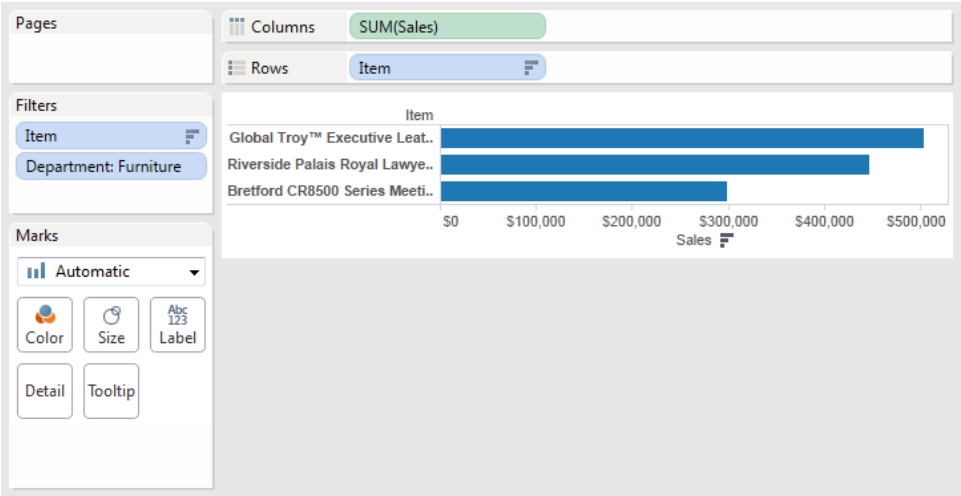

6. To find out what the top 10 furniture products are we need to make the Department filter a context filter. Right-click the field on the Filters shelf and select **Add to Context**.

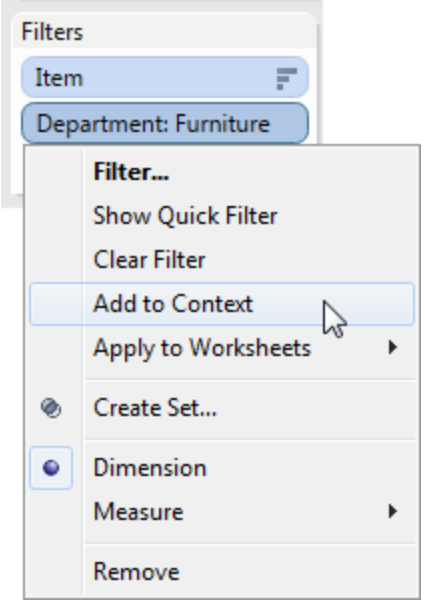

7. The filter is marked as a context filter and the view updates to show the top 10 furniture products. Tableau has first evaluated the data source and identified all of the furniture products. Then the Top 10 filter is evaluated on the results of that context.

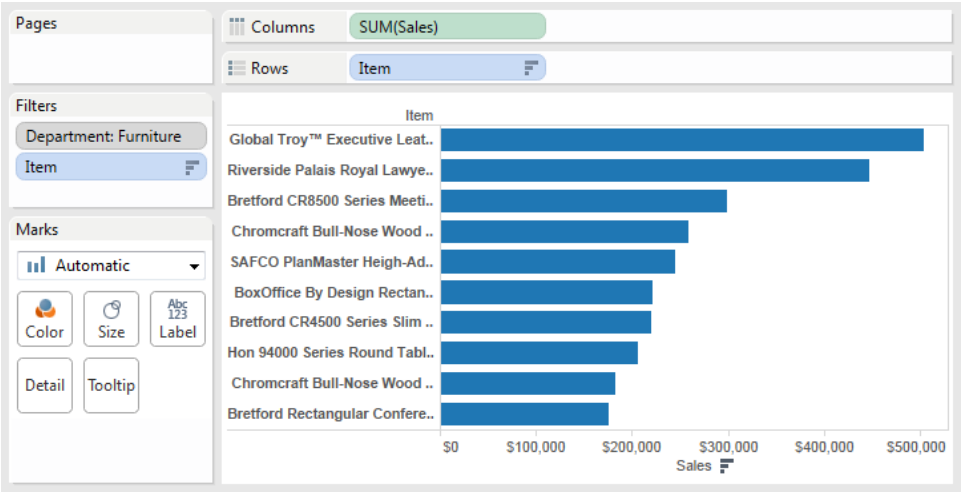

Calculation Filters

With a calculation filter, Tableau performs a calculation on the selected dimension members. By definition, a calculation filter exists when:

- The dimension is used only on the Filters shelf (that is, the dimension does not appear on other shelves).
- The filter is defined to include multiple values.

The calculation that is performed depends on the data source. For relational data sources, as shown below, the calculation matches the aggregation for each measure used in the view. For multidimensional data sources, as shown in **Calculation Filters: Multidimensional Data Source** on page 387, the calculation is always a summation. In Tableau, multidimensional data sources are supported only in Windows.

Consider the following view.

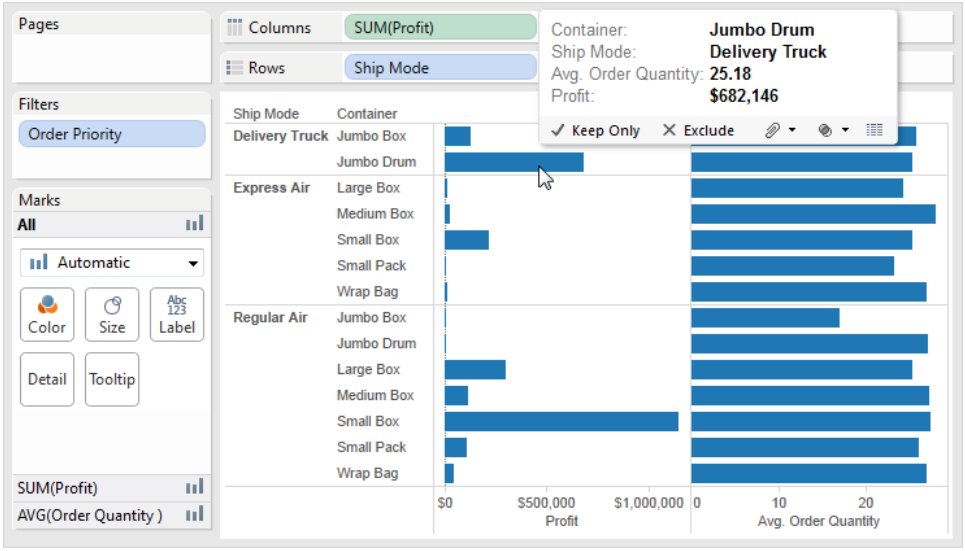

On the **Columns** shelf, the **Profit** measure is aggregated as a summation and the **Order Quantity** measure (hidden beneath the tooltip) is aggregated as an average. These measures are displayed with the **Ship Mode** and **Container** dimensions (on the **Rows** shelf). A filter that limits the view to two members of the **Order Priority** dimension is applied to the data:

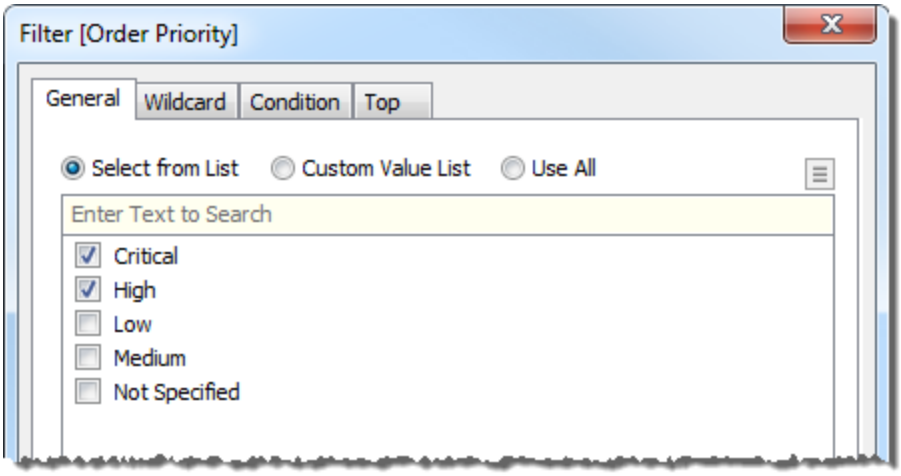

Tableau automatically applies the appropriate calculation to the members of the filter based on the aggregation of each measure. Therefore, a summation is performed for **Profit** and an average is performed for **Order Quantity**.

For example, the tooltip shows the data for Jumbo Drums delivered by truck. The average order quantity is 25.18. This number was calculated by averaging the order quantities for all the rows that have a Critical or High order priority. Similarly, the sum of profit is $682,196. This number was calculated by summing the profit for all the rows that have a Critical or High order priority.

Calculation Filters: Multidimensional Data Source

In Tableau, multidimensional data sources are supported only in Windows.

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, which shows profit by region and state.

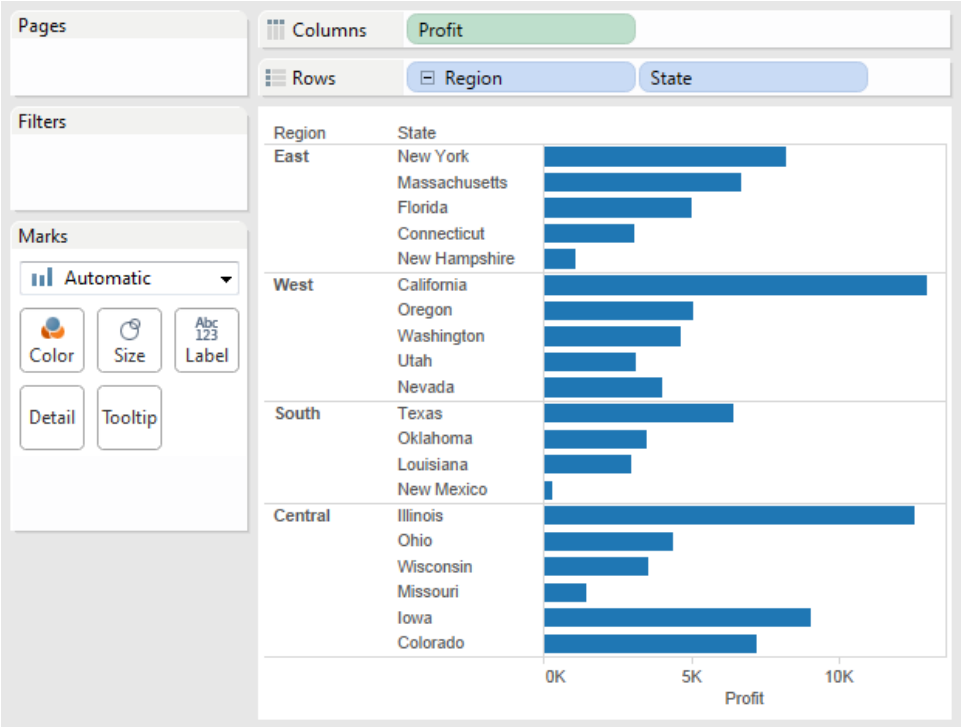

An external filter that consists of the Q3 and Q4 members of the **Quarter** dimension is applied to the data.

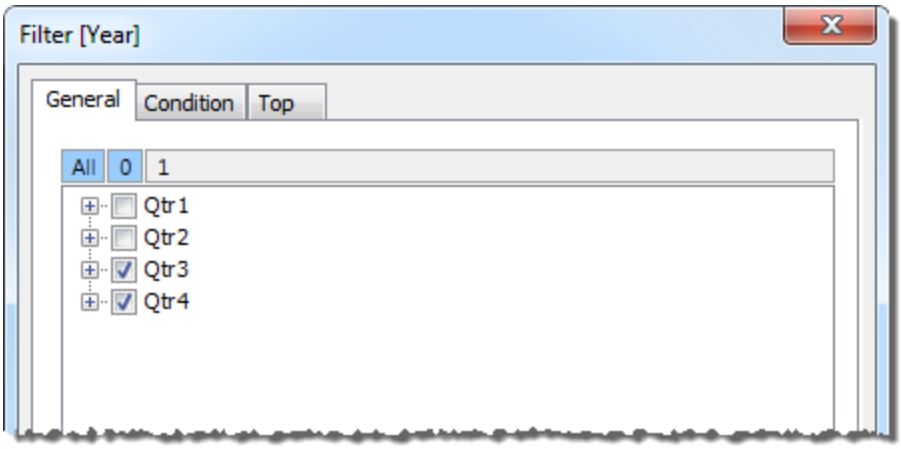

For example, tin the view below the profit for Oregon is now $2,567. This number was calculated by summing the data for Qtr3 and Qtr4.

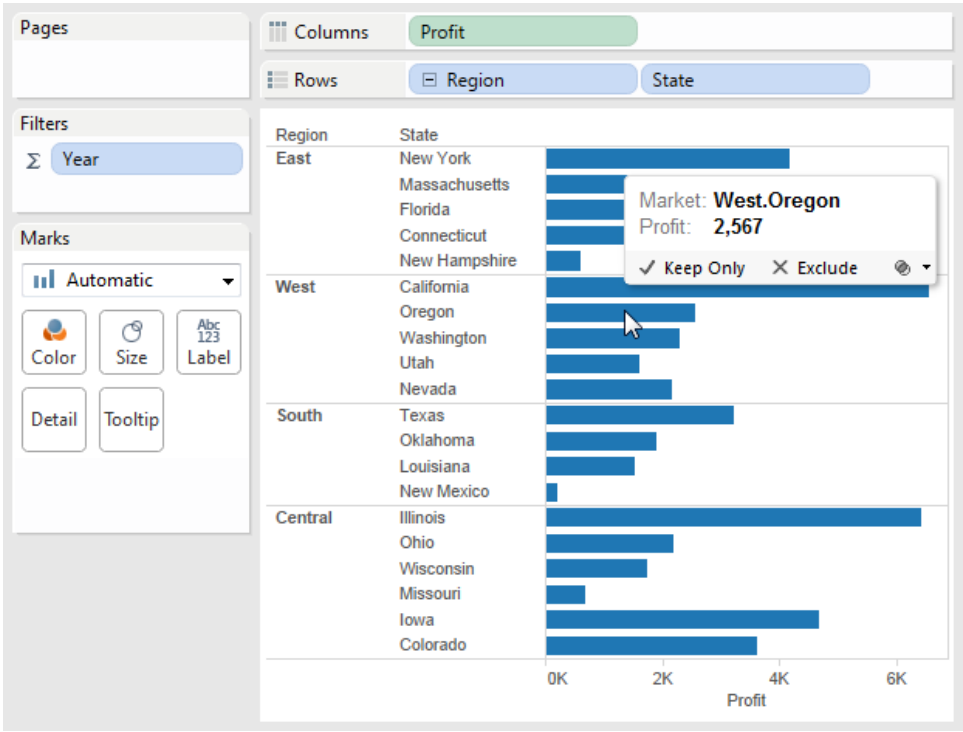

Sorting, Grouping, and Sets

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

In Tableau, sorting a data view means arranging dimension members in a specified order. Tableau supports computed sorting and 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.

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 401 for more information.
- Sorted fields are identified by a sort icon

on the right side of the field.

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.

How to Sort Data (Computed Sorts)

Computed sorts can be applied either directly on an axis in the view or using the sort dialog box to apply computed sorts to specific fields in the view.

Sorting on an Axis

A quick way to create a computed sort is to use the sort buttons on an axis. The sort buttons on an axis will automatically create a sort that makes sense for the view you’ve created. The sort will be computed, which means it will update correctly if the underlying data changes.

When you hover over an axis with your cursor, a sort icon displays. Click it once to sort Ascending, again to sort Descending, and a third time to clear the sort.

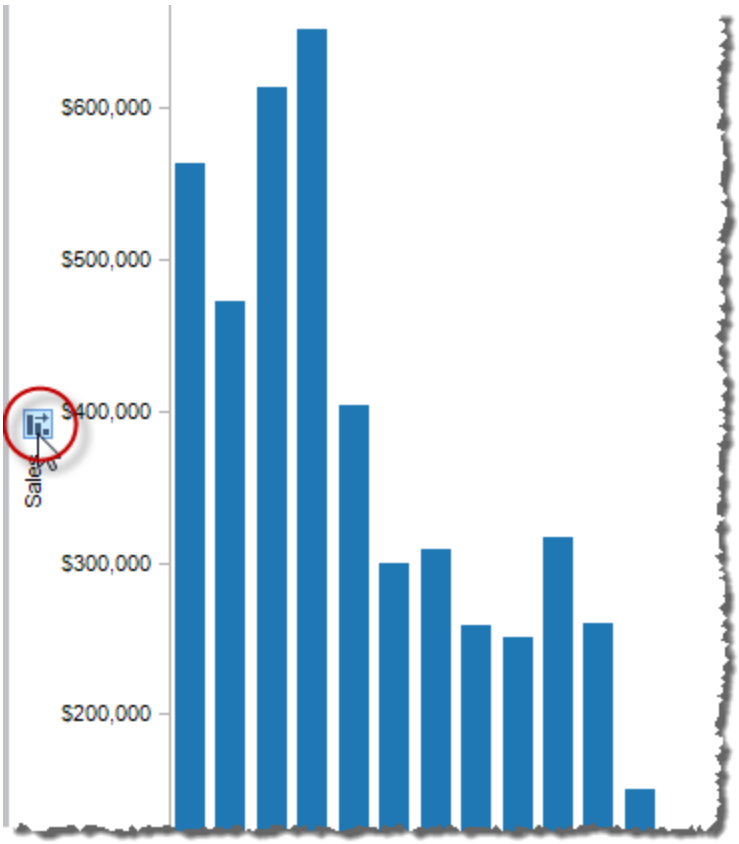

Sorting Specific Fields

You can sort a specific field using the Sort dialog box.

1. Open the Sort dialog box.

   Right-click on the field that you want to sort and select **Sort** from the context menu.

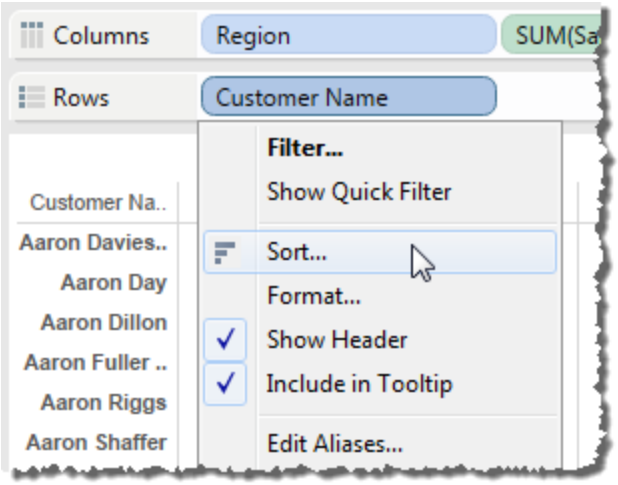

2. Specify the sorting options.

Complete the **Sort** dialog box by specifying the following criteria:

- Sort order - Displays 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. In Tableau, multidimensional data sources are supported only in Windows.

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.

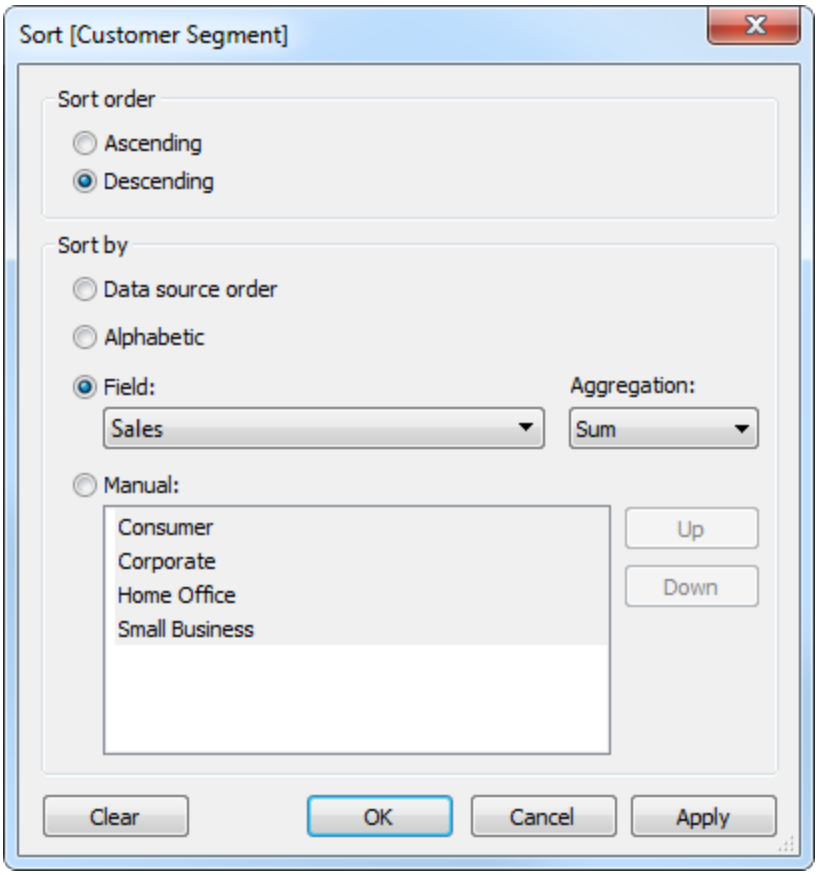

You should keep the following rules in mind when interpreting the sort results:

- Tableau computes the sort across the entire table using the specified criteria. Refer to **Example – Sorting a Text Table** below 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.

  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 Superstore sample 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 **Category** dimension on the Rows shelf.

Complete the text table by placing **Discount** on **Text**. By default, the table is sorted in alphabetical order.

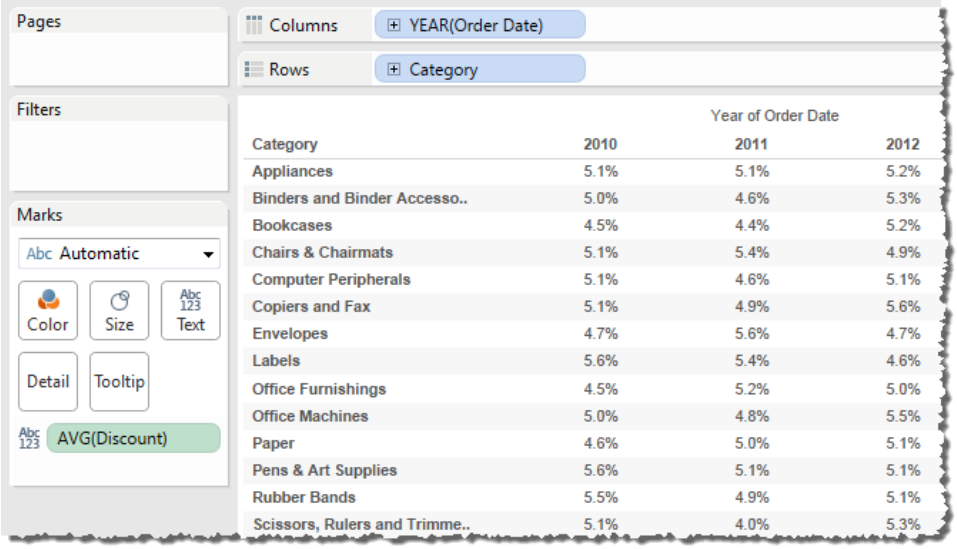

Category	2010	2011	2012
Appliances	5.1%	5.1%	5.2%
Binders and Binder Accesso..	5.0%	4.6%	5.3%
Bookcases	4.5%	4.4%	5.2%
Chairs & Chairmats	5.1%	5.4%	4.9%
Computer Peripherals	5.1%	4.6%	5.1%
Copiers and Fax	5.1%	4.9%	5.6%
Envelopes	4.7%	5.6%	4.7%
Labels	5.6%	5.4%	4.6%
Office Furnishings	4.5%	5.2%	5.0%
Office Machines	5.0%	4.8%	5.5%
Paper	4.6%	5.0%	5.1%
Pens & Art Supplies	5.6%	5.1%	5.1%
Rubber Bands	5.5%	4.9%	5.1%
Scissors, Rulers and Trimme..	5.1%	4.0%	5.3%

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 **Category**.

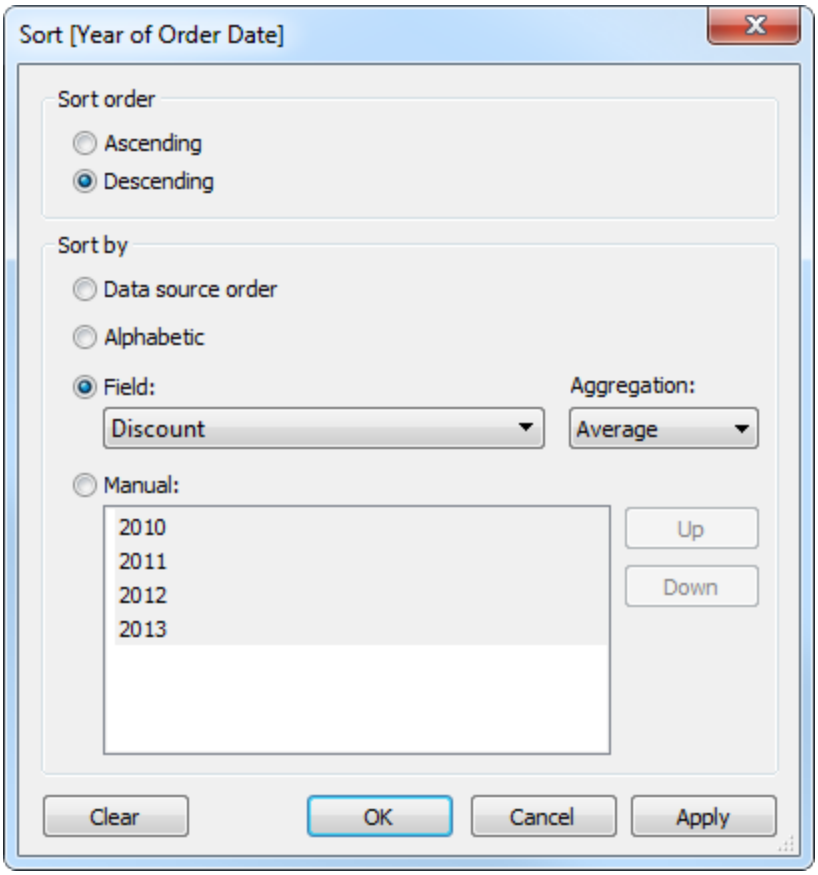

The view is shown below. Rubber Bands is the top row in the table because it has the largest average discount across all years, while Telephones and Communications are at the bottom in the table because that category has the smallest average discount across all years. Similarly, 2012 is the left most column because it has the largest average discount for all products, while **2011** is the right most column because it has the smallest average discount for all products.

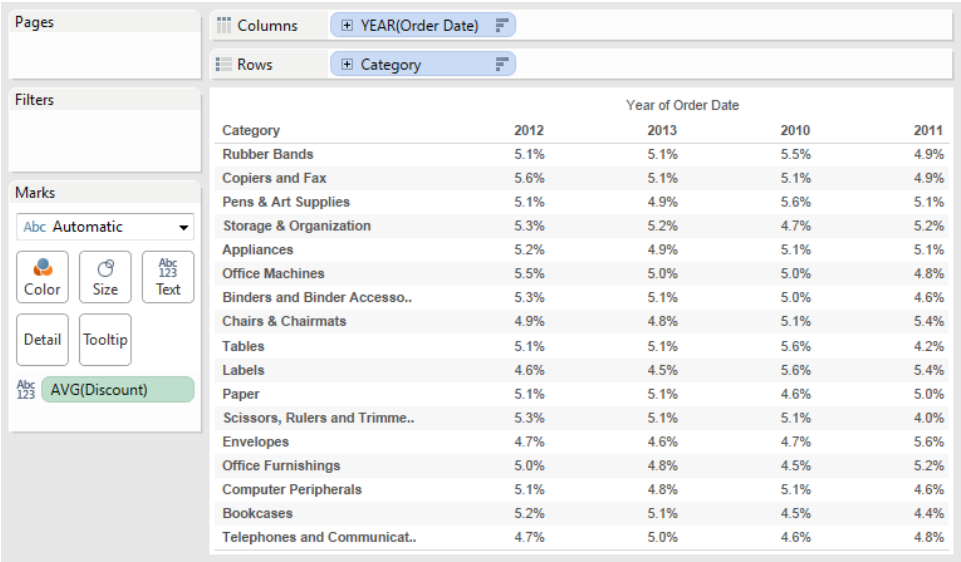

	Year of Order Date			
Category	2012	2013	2010	2011
Rubber Bands	5.1%	5.1%	5.5%	4.9%
Copiers and Fax	5.6%	5.1%	5.1%	4.9%
Pens & Art Supplies	5.1%	4.9%	5.6%	5.1%
Storage & Organization	5.3%	5.2%	4.7%	5.2%
Appliances	5.2%	4.9%	5.1%	5.1%
Office Machines	5.5%	5.0%	5.0%	4.8%
Binders and Binder Accesso..	5.3%	5.1%	5.0%	4.6%
Chairs & Chairmats	4.9%	4.8%	5.1%	5.4%
Tables	5.1%	5.1%	5.6%	4.2%
Labels	4.6%	4.5%	5.6%	5.4%
Paper	5.1%	5.1%	4.6%	5.0%
Scissors, Rulers and Trimme..	5.3%	5.1%	5.1%	4.0%
Envelopes	4.7%	4.6%	4.7%	5.6%
Office Furnishings	5.0%	4.8%	4.5%	5.2%
Computer Peripherals	5.1%	4.8%	5.1%	4.6%
Bookcases	5.2%	5.1%	4.5%	4.4%
Telephones and Communicat..	4.7%	5.0%	4.6%	4.8%

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.

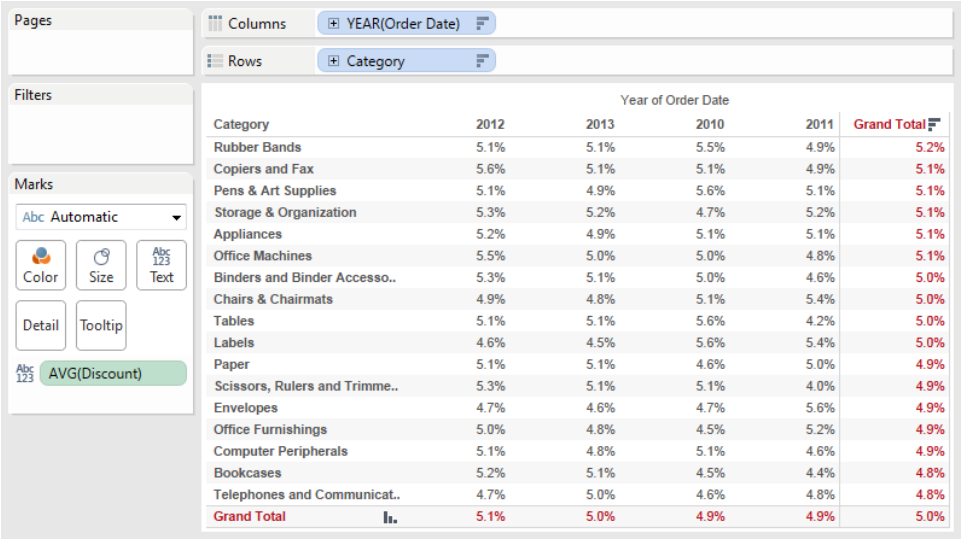

	Year of Order Date				
Category	2012	2013	2010	2011	Grand Total
Rubber Bands	5.1%	5.1%	5.5%	4.9%	5.2%
Copiers and Fax	5.6%	5.1%	5.1%	4.9%	5.1%
Pens & Art Supplies	5.1%	4.9%	5.6%	5.1%	5.1%
Storage & Organization	5.3%	5.2%	4.7%	5.2%	5.1%
Appliances	5.2%	4.9%	5.1%	5.1%	5.1%
Office Machines	5.5%	5.0%	5.0%	4.8%	5.1%
Binders and Binder Accesso..	5.3%	5.1%	5.0%	4.6%	5.0%
Chairs & Chairmats	4.9%	4.8%	5.1%	5.4%	5.0%
Tables	5.1%	5.1%	5.6%	4.2%	5.0%
Labels	4.6%	4.5%	5.6%	5.4%	5.0%
Paper	5.1%	5.1%	4.6%	5.0%	4.9%
Scissors, Rulers and Trimme..	5.3%	5.1%	5.1%	4.0%	4.9%
Envelopes	4.7%	4.6%	4.7%	5.6%	4.9%
Office Furnishings	5.0%	4.8%	4.5%	5.2%	4.9%
Computer Peripherals	5.1%	4.8%	5.1%	4.6%	4.9%
Bookcases	5.2%	5.1%	4.5%	4.4%	4.8%
Telephones and Communicat..	4.7%	5.0%	4.6%	4.8%	4.8%
Grand Total	5.1%	5.0%	4.9%	4.9%	5.0%

Example – Sorting a 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**.

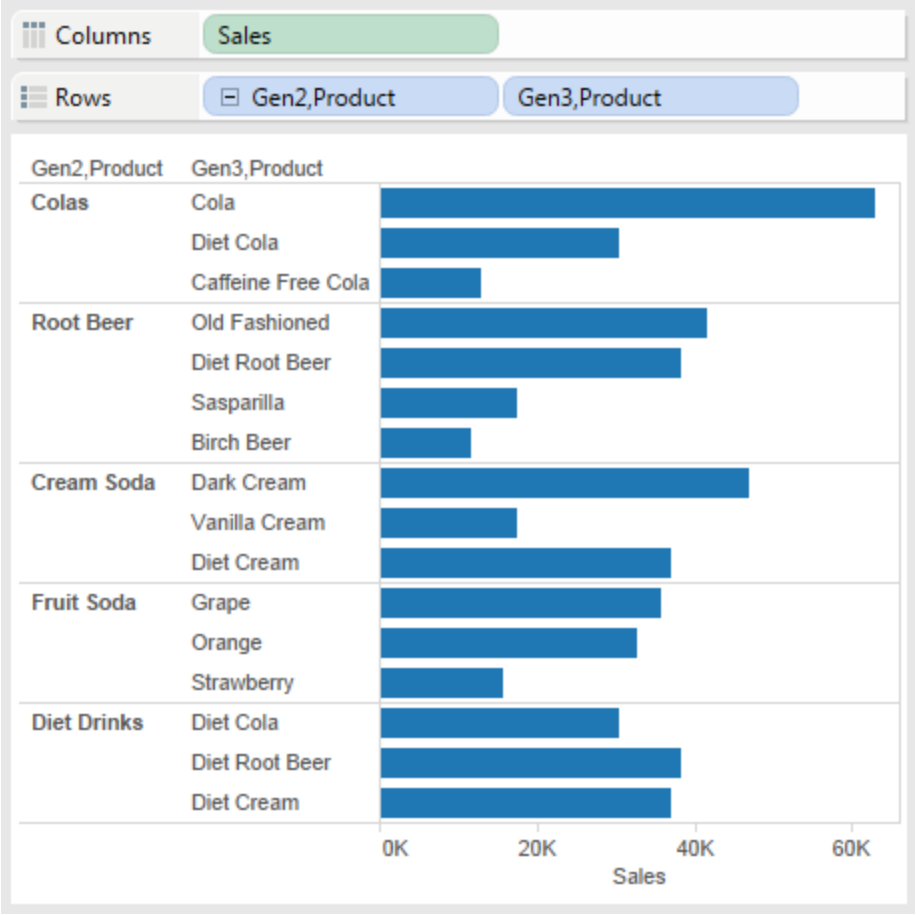

2. Sort **Gen3,Product** in descending order by the **Sales** measure.

   Right-click on Gen3,Product and select **Sort** from the field's context menu. In the Sort dialog box, select **Descending** as the Sort order and sort by the **Sales** field.

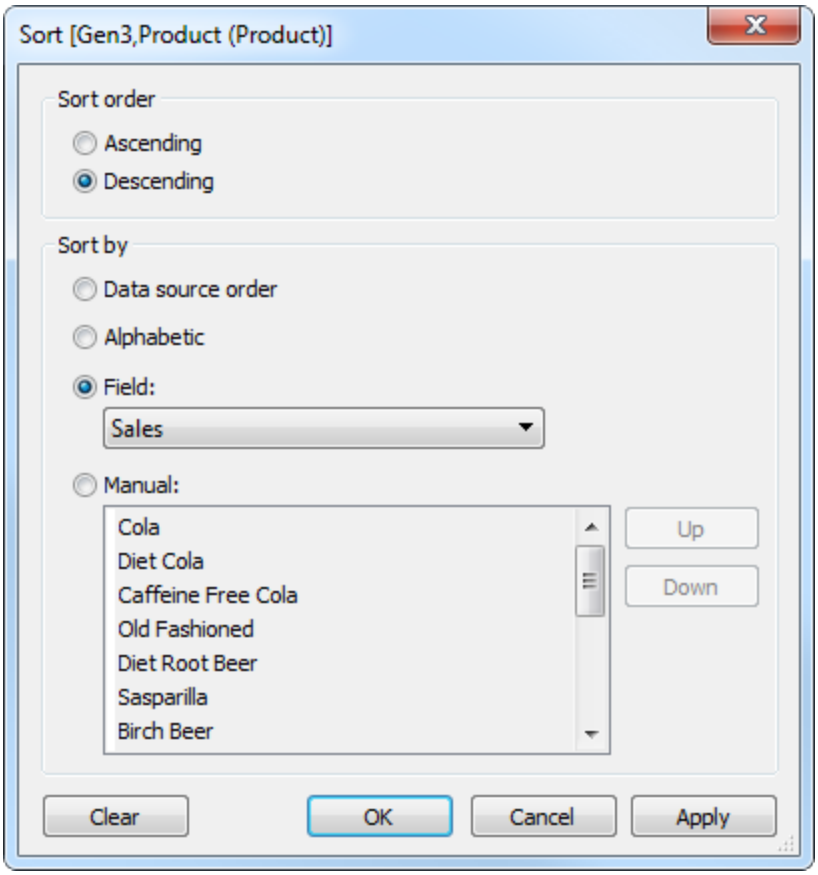

The view is shown below. Notice that the **Gen3,Product** members are sorted within each parent member. For example, Cola, Diet Cola, and Caffeine Free Cola are sorted only within the Colas level. Tableau does not rearrange headers that appear before the sorted field.

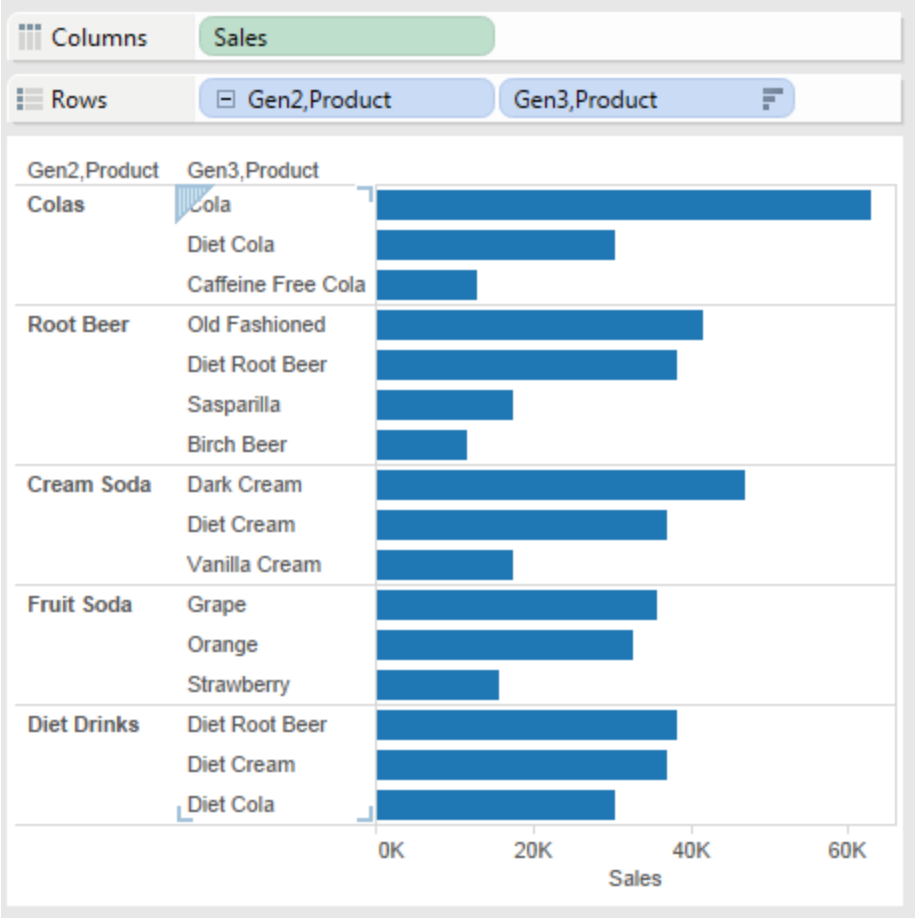

3. If you want to order dimension members without regard to its parent, you should remove Gen2,Product from the Rows shelf. The sorted data are shown below.

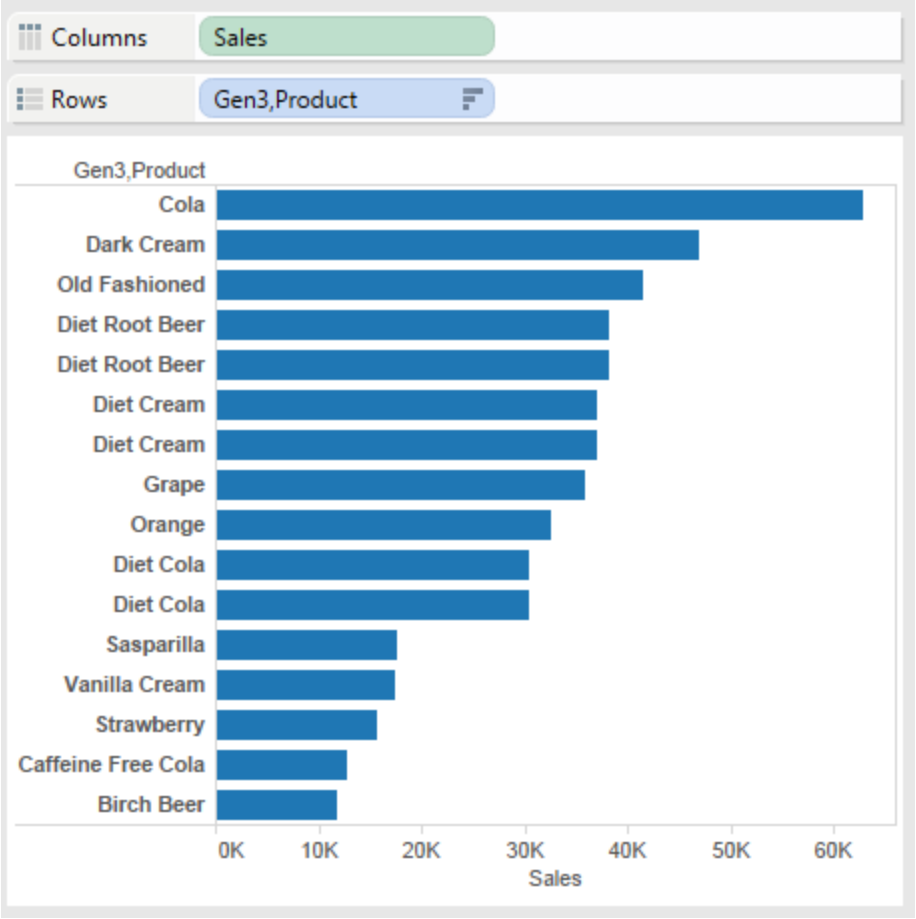

Manual Sorting

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.

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.

Sorting Using the Toolbar and Tooltips

The two sort buttons on the toolbar and in tooltips 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 category and customer segment. When you select the Corporate Market column, thus selecting all of the products in that segment, the quick sort buttons sorts the product field by SUM(Sales), which is the measure in the view.

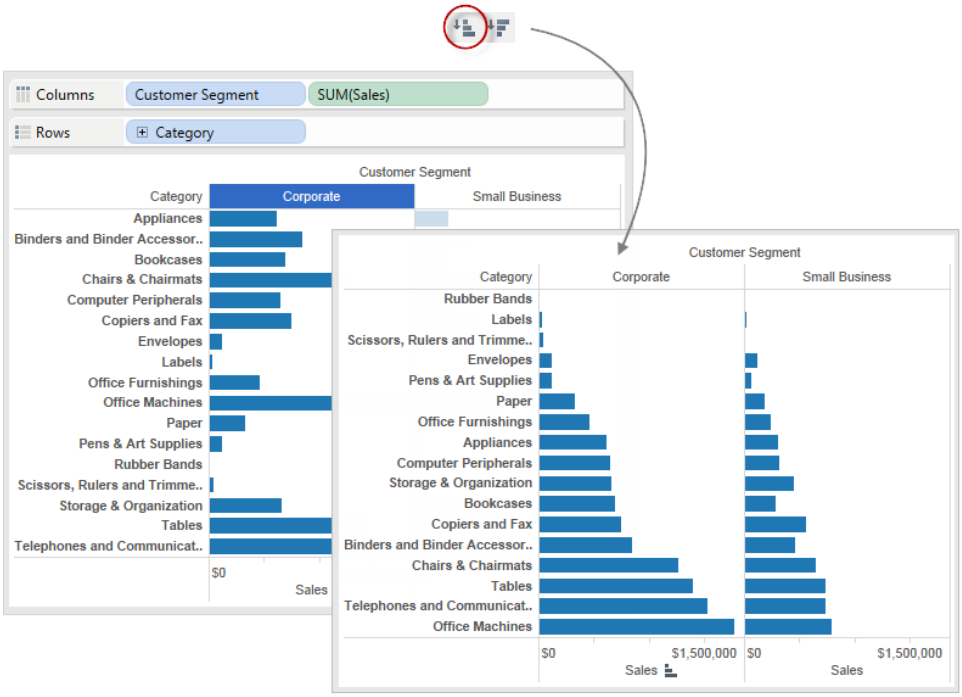

An easy way to anticipate how a selection will be sorted is to make a selection in the view and hover over the ascending or descending sort buttons on the toolbar. The tooltip for each button describes how the selection will be sorted.

Using the quick sort buttons on the toolbar or in the tooltip 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.

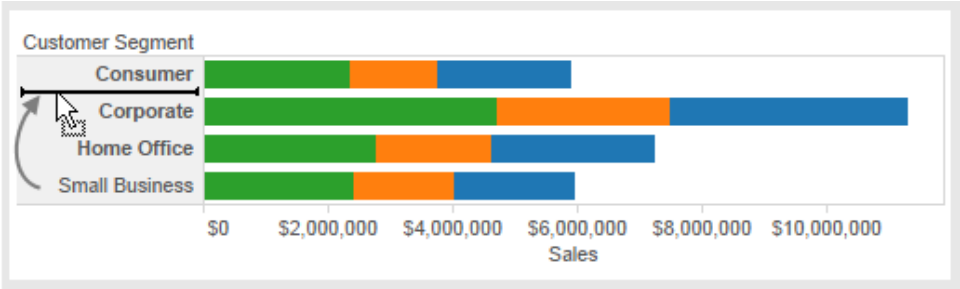

Example- Manually Sorting Drawing Order

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.

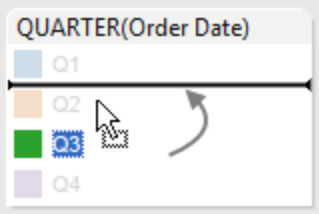

Sorting the drawing order is not restricted to color legends. You can reorder shape legends as well. If you have multiple valid legends, the drawing order is defined first by shape, then by color. For example, suppose you have both a shape legend and a color legend. If you have a red circle on top of 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 combination of dimension members that make 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 while Biology and Physics may be grouped as Science Majors. Groups are useful for both correcting data errors (e.g., combining CA, Calif., and California into one) as well as answering "what if" type questions (e.g., "What if we combined the East and West regions?).

Creating Groups

You can create a group by selecting headers in the view, selecting marks, or based on a dimension in the Data window. Regardless of how you create the group, a new group field is added to the Data window. You can use the group field like any other field in the view including on the Columns and Rows shelves, on the Marks card, and the Filter shelf. When you create a group by selecting headers or marks in the view, the view is updated to use the grouped field.

Correct Data Errors or Combine Dimension Members

If you are using groups to correct data errors (combine CA and California) or to combine dimensions members (combine East and West regions), the easiest way to create the group is by selecting headers in the view.

1. Hold the **CTRL** or **Shift** key on the keyboard to multi-select headers in the view.
2. Click the **Group** button on the toolbar.

   Alternatively, you can right-click and select **Group** or click the **Group** icon on the tooltip.

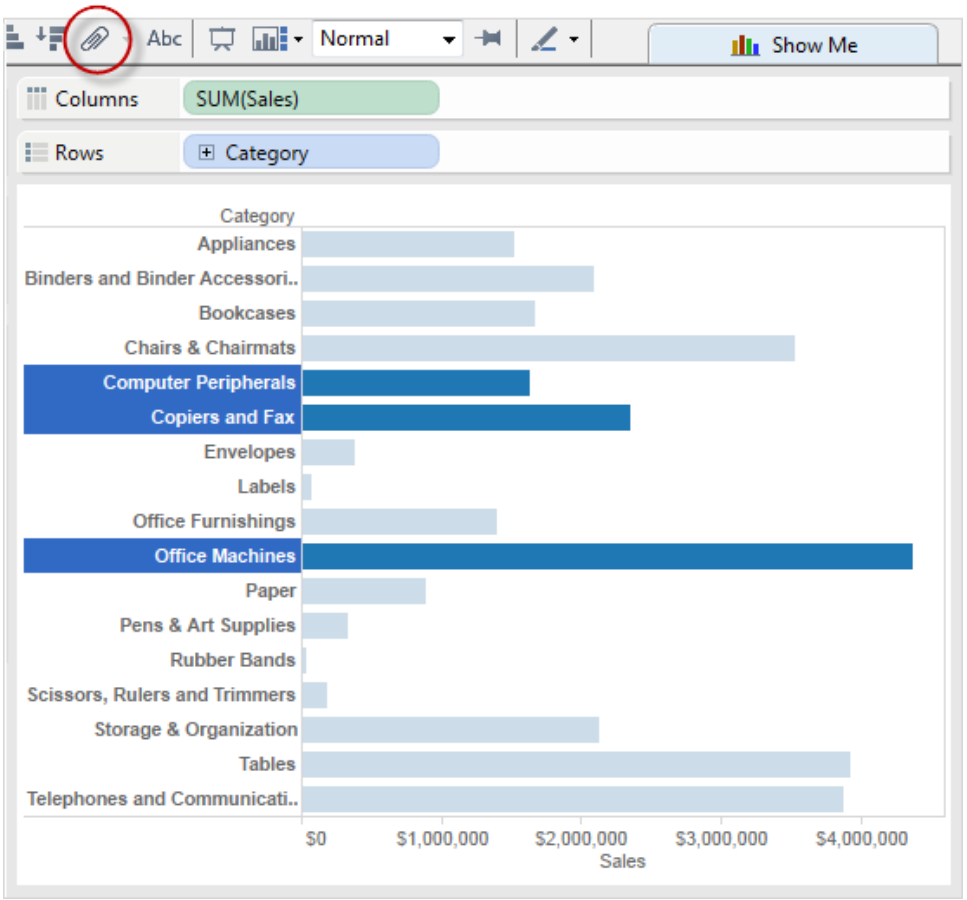

The selected members are combined into a single member. In this example, the view now shows the SUM(Sales) across all Computer Peripherals, Copiers and Fax, and Office Machines in a single bar. A default is automatically constructed using the combined member names.The dimension on the Rows or Columns shelf is replaced with the new grouped field.

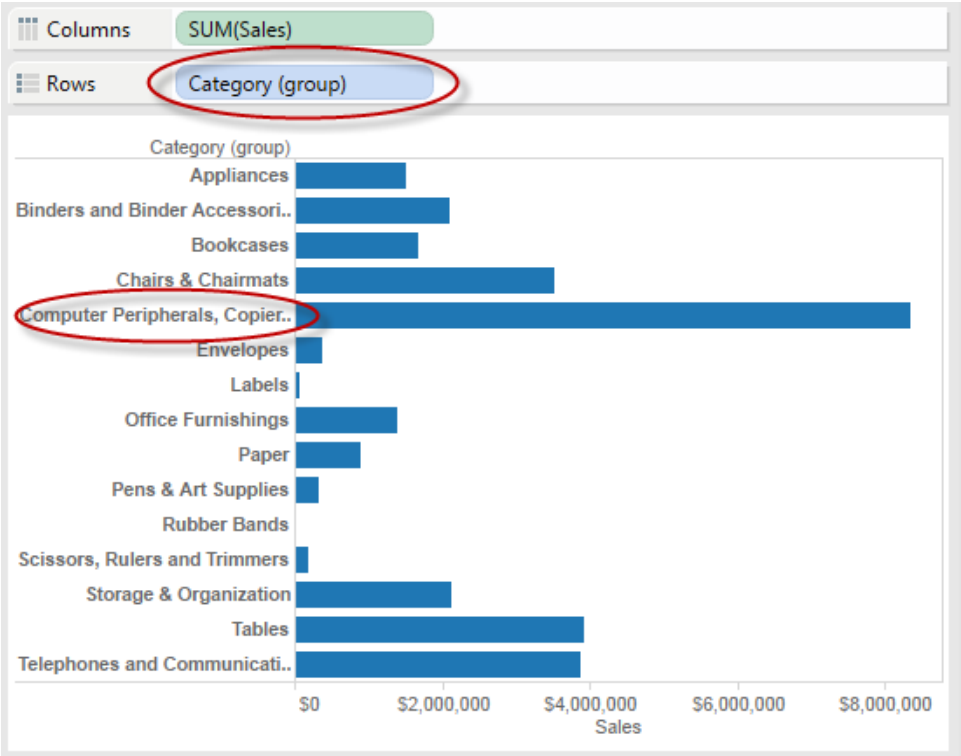

The grouped field is also added to the Data window. You can use it just like any other field in the view. For example, you can drag it to the shelves, the Marks card, and add it to relational hierarchies. Group fields cannot be used in calculated fields.

Color a View Using Groups

In addition to correcting data errors and combining dimensions, you can visually identify groups of related marks by selecting marks directly. This technique is especially useful when working with a scatter plot or a view that doesn't have headers that you can select when defining the group.

1. Hold the **CTRL** or **Shift** key on the keyboard to select one or more marks in the view.
2. Click the **Group** button on the toolbar.

Alternatively, you can right-click and select **Group** or click the **Group** icon on the tooltip.

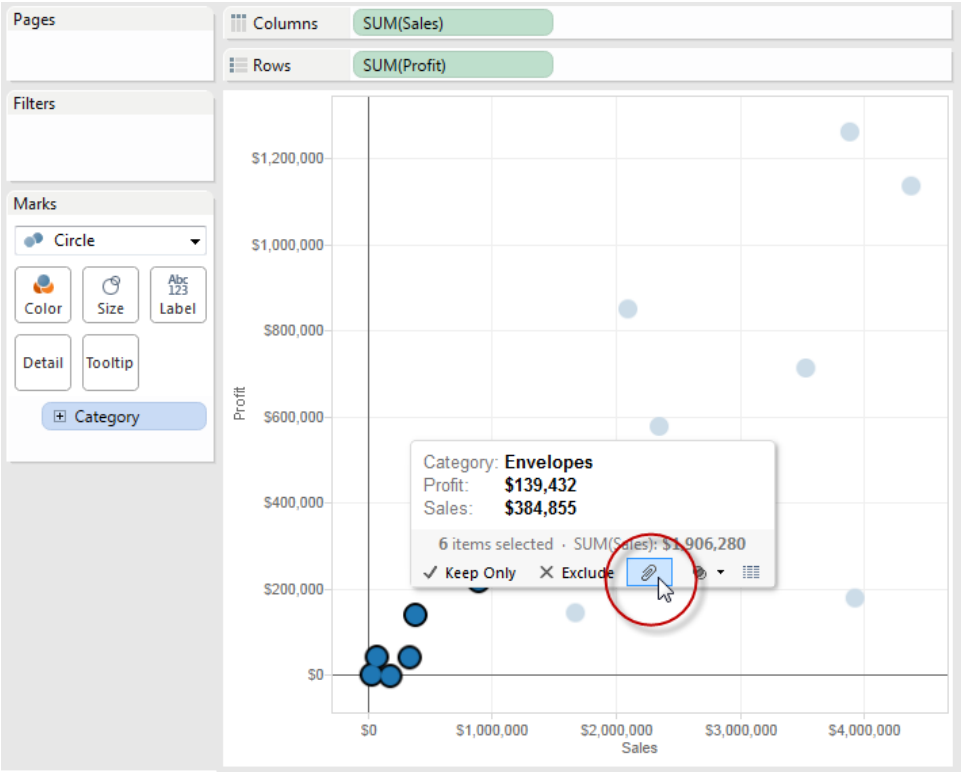

The selected marks are grouped and all other members are combined into an "Other" category. The new group field is added to Color. See **Including an Other Group** on page 413 to learn more about grouping fields using Other.

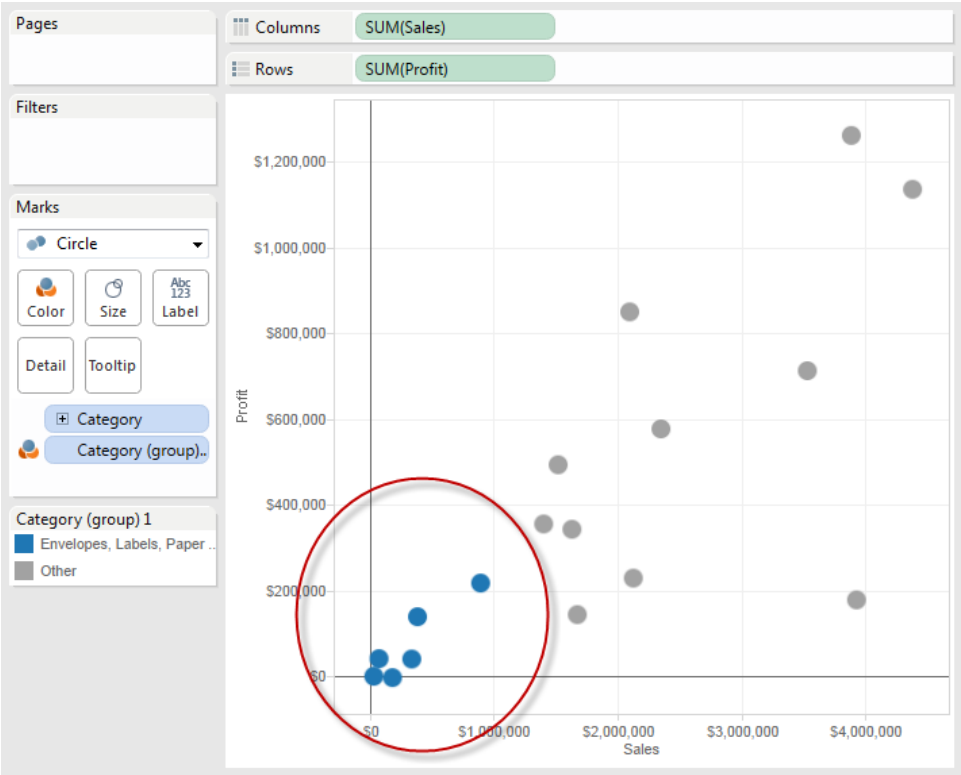

If there were already fields on Color, they are moved to Detail and replaced with the group field.

When you create groups by selecting marks, it's possible that the marks will represent multiple dimensions. For example, you may have a scatter plot that shows the Sales vs. Profit by Region and Category. In this case, a selection of marks will represent members of both the Region and Category dimensions. When the selection represents multiple dimensions, the Group menu lets you choose to group on All Dimensions or a particular dimension.

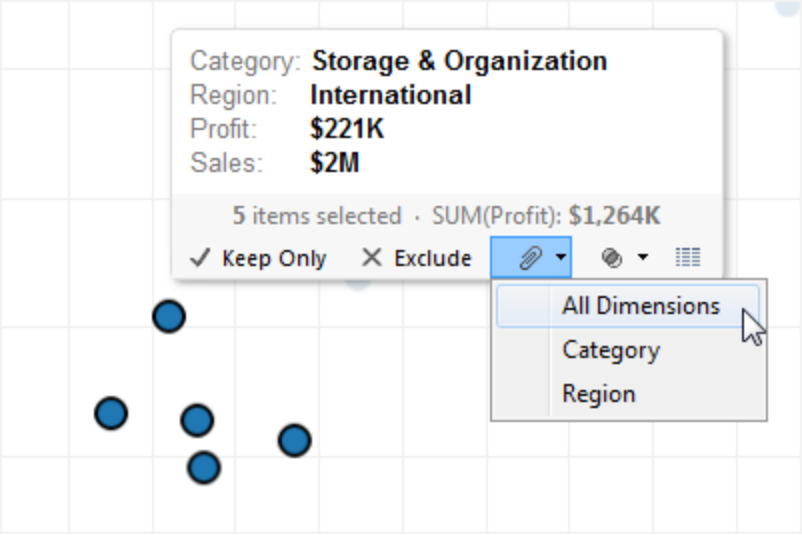

In the example above, the five selected marks represent a region and the following category combinations .

- Storage & Organization, International
- Bookcases, International
- Computer Peripherals, International
- Office Furnishings, International
- Appliances, International

The views below show the results of grouping these measures on All Dimensions, Category, and Region.

All Dimensions	Category	Region

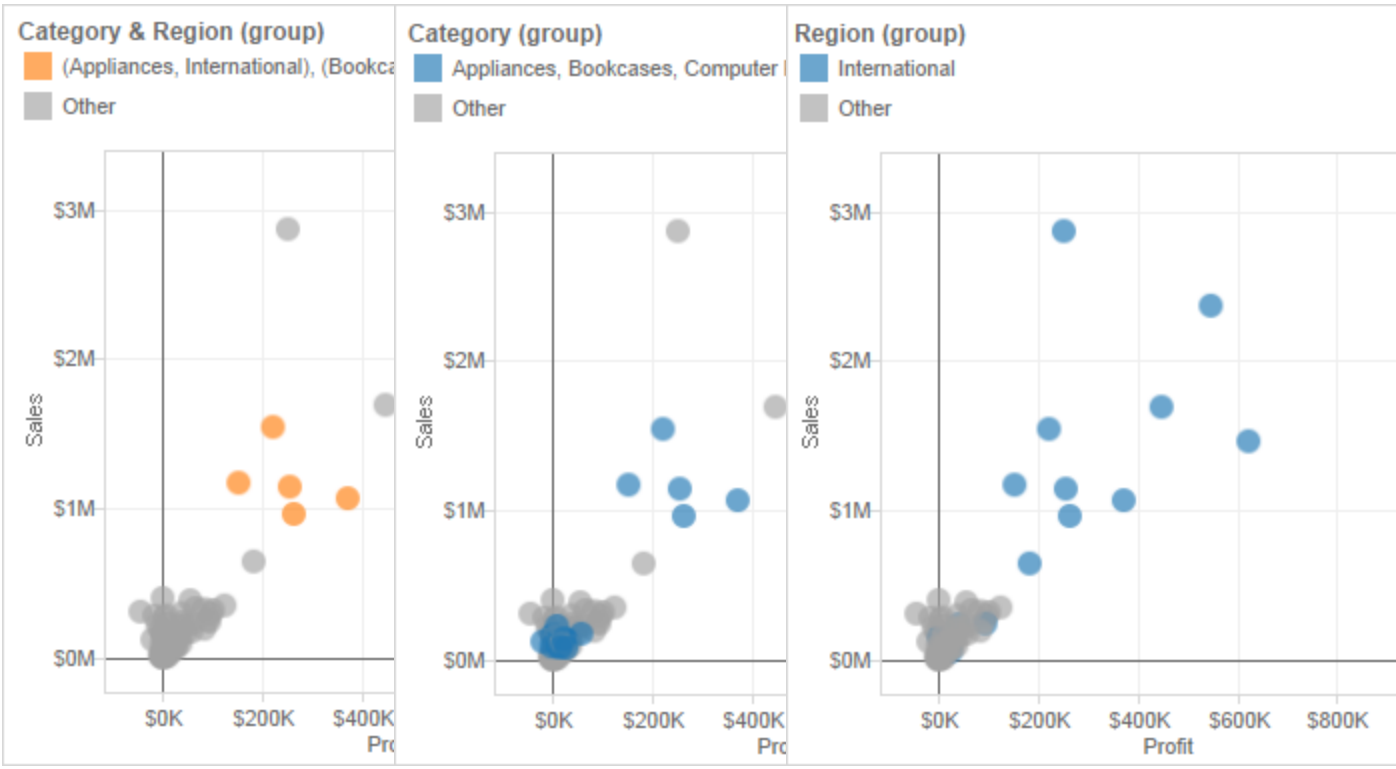

The five marks are combined and the rest of the marks are added to an "Other" category.

All marks associated with any of the five categories are combined and everything else is added to an "Other" category.

All marks associated with the five regions (in this case just International) are combined and everything else is added to an "Other" category.

Create Groups in the Data Window

Whenever you create a group, a group field is added to the Data window. When you create a group by selecting headers or marks in the view, the new group field is also added to the view either on the Rows or Columns shelf or to Color. Instead of creating a group by selections in the view, you can create a group directly from the Data window.

1. Right-click a dimension in the Data window and select **Create Group**.

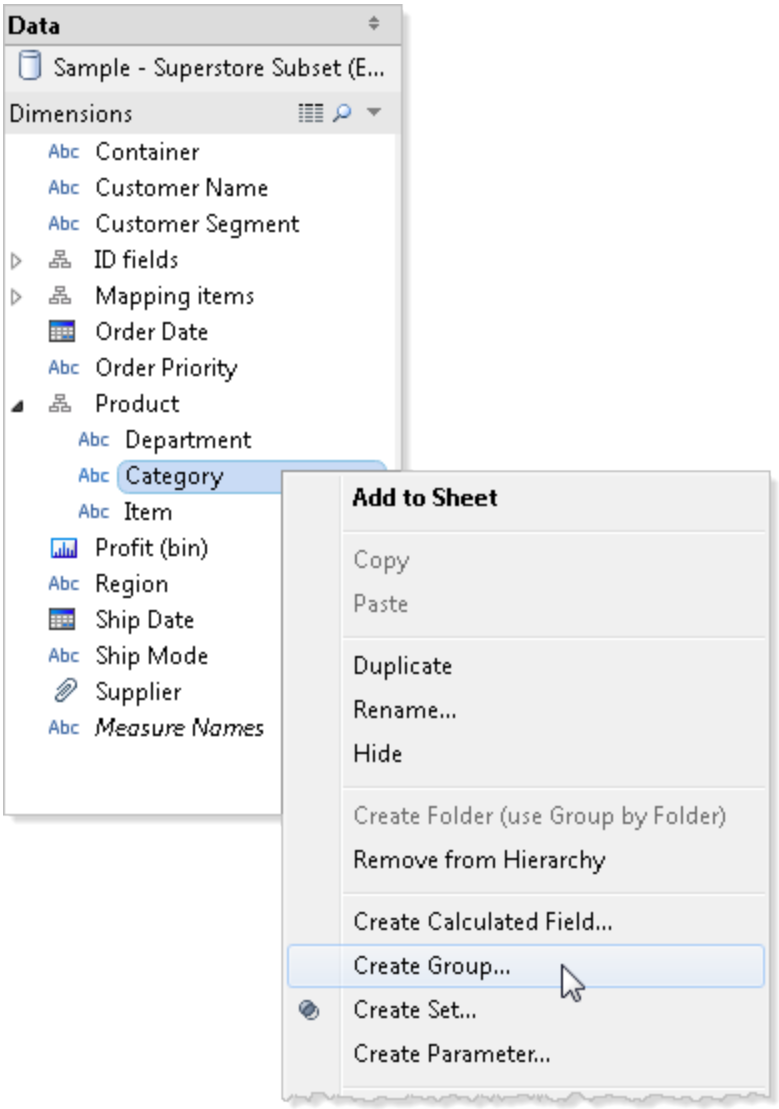

2. In the Create Group dialog box, select several members that you want to group. Hold the **CTRL** key on your keyboard to select multiple members.

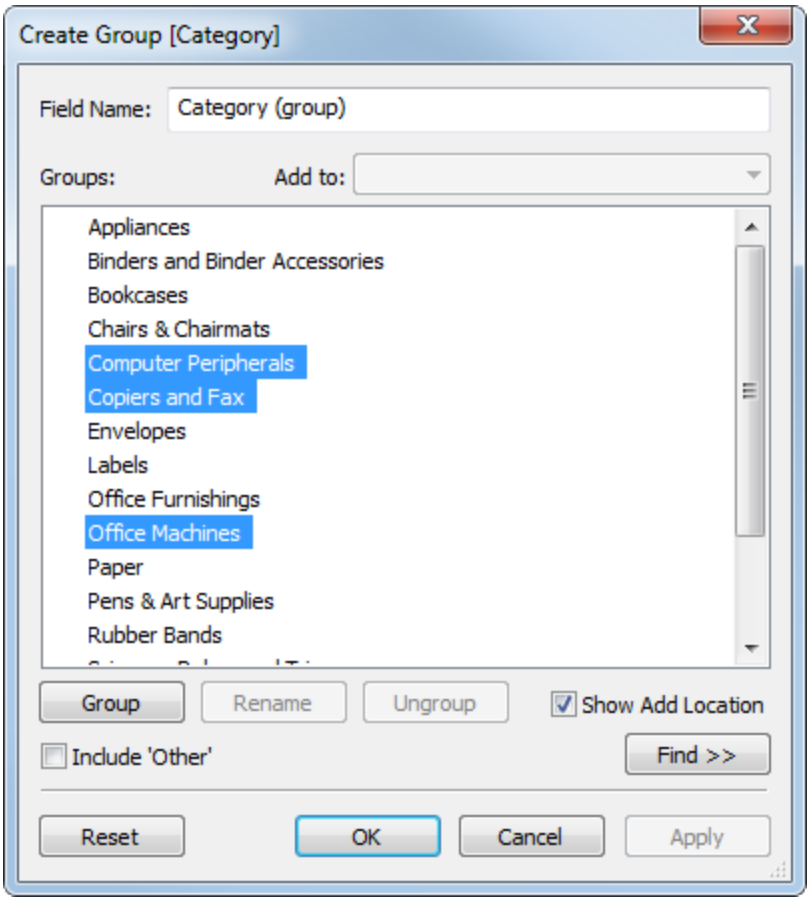

3. Click the **Group** button at the bottom of the dialog box.

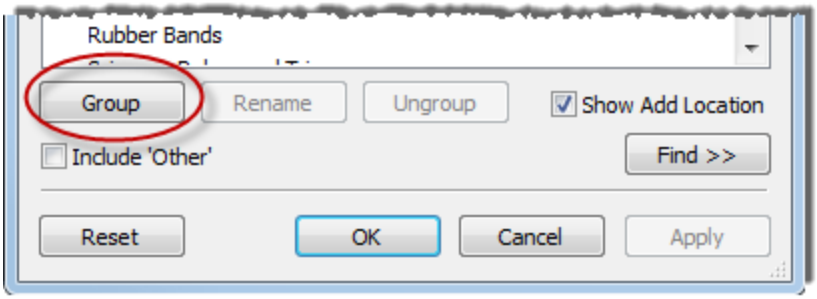

The selected members are combined into a single member. A default name is automatically constructed using the combined member names. Rename the group by selecting it in the list and clicking the **Rename** button at the bottom of the dialog box. You can search for members using the Find options. See **Finding Members in Group Dialog Box** on page 415.

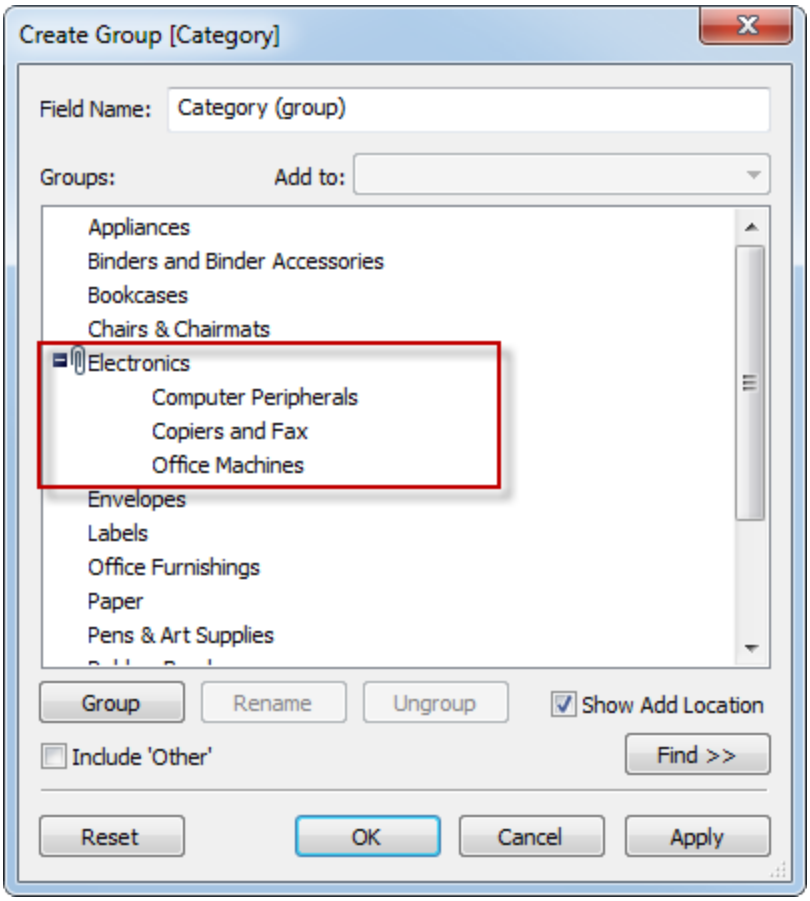

To add to or remove members from a group, right-click the grouped field in the Data window and select **Edit Group**. In the Edit Group dialog box you can also change the default name of the group and combine fields into new groups. Refer to **Editing an Existing Group** below to learn more. Alternatively, you can ungroup members by selecting them in the view and clicking the Ungroup button on the toolbar or tooltip.

Editing an Existing Group

After you have created a grouped field, you can add and remove members from the groups, create new groups, change the default group names, and change the name of the grouped field. Some of these changes can be made directly in the view while others require you to open the Edit Group dialog box.

Edit Groups Directly in the View

You can add and remove members from an existing group and create new groups by selecting marks directly in the view. Do one of the following:

- **Add to an Existing Group** - Select one or more marks you want to add to an existing group. Then hold **SHIFT** or **CTRL** on your keyboard to also select at least one member

of the group you want to add to. Click the **Group** button on the toolbar or tooltip.

- **Remove from an Existing Group** - Select one or more marks you want to ungroup. Then click the Ungroup button on the toolbar or tooltip. Make sure that the selection only contains members from one group at a time.
- **Create a new group** - Select one or more marks that you want to create a new group with. Then click the Group button on the toolbar or tooltip.

Open the Edit Group Dialog Box

While editing groups directly in the view is useful for rapidly defining ad hoc groups, you can more finely control the members of each group using the Edit Group dialog box. Right-click the grouped field in the Data window or in the view and select Edit Group.

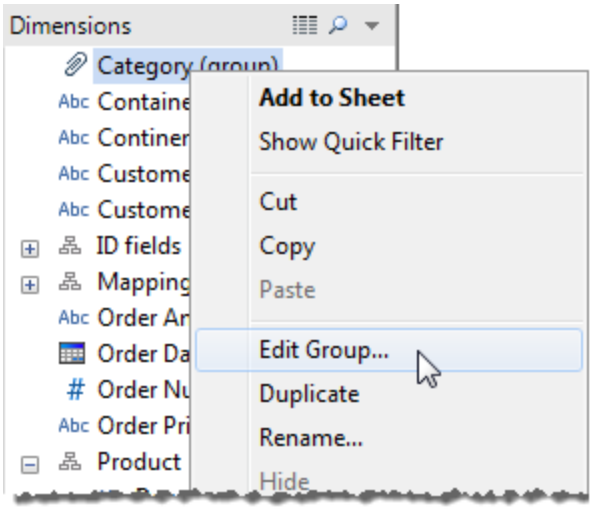

Then do one of the following:

- 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.

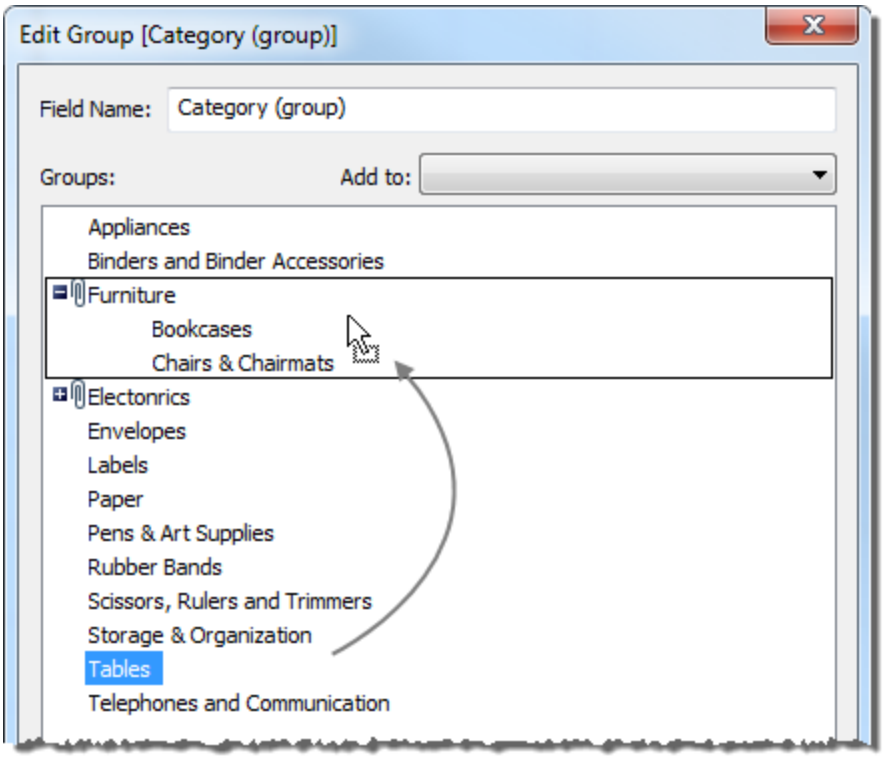

- Select one or more members, right-click and select **Add To**. In the subsequent dialog box, select the group you want to add the selected members to and click **OK**.

- Select one ore more members and select the group in the **Add to** drop down list at the top of the dialog box.

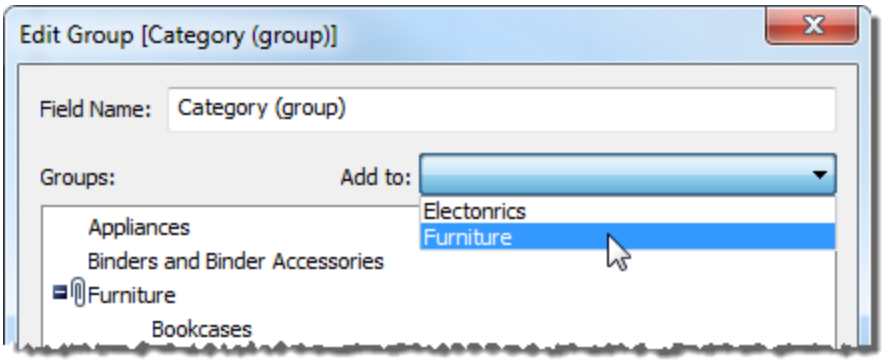

Rename a group by selecting the group in the Edit Group dialog box and then clicking the Rename button.

Including an Other Group

When you create groups in Tableau, you have the option to **Include Other**. When you include an other category, the groups you have defined are shown and all other dimension members are combined into an "Other" category. This option is useful for highlighting certain groups or comparing specific groups against everything else. For example, you may have a grouped field that combines product categories based on their sales performance. You may want to highlight High Performers and Low Performers and group all the other categories into an "Other" category.

Includes Other **Does Not Include Other**

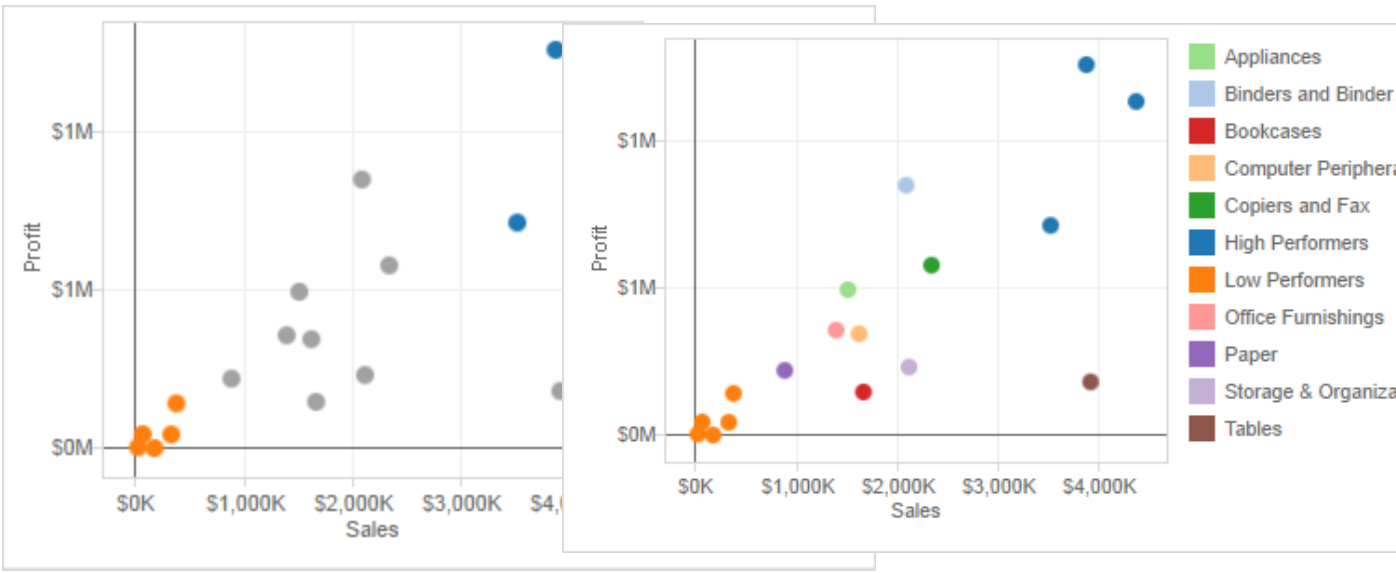

Define whether to include the other category by selecting the **Include Other** option on the field menu.

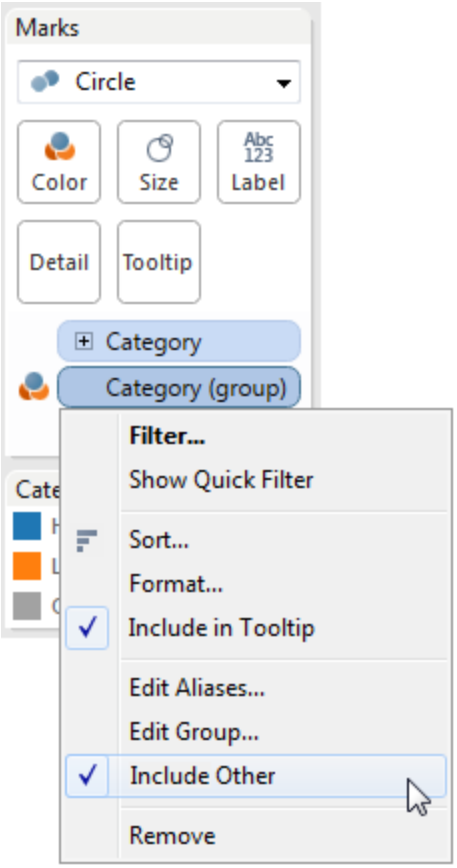

Alternatively, you can right-click the field in the Data window or on the sheet and select **Edit Group**. Then select the **Include Other** option near the bottom of the Edit Group dialog box.

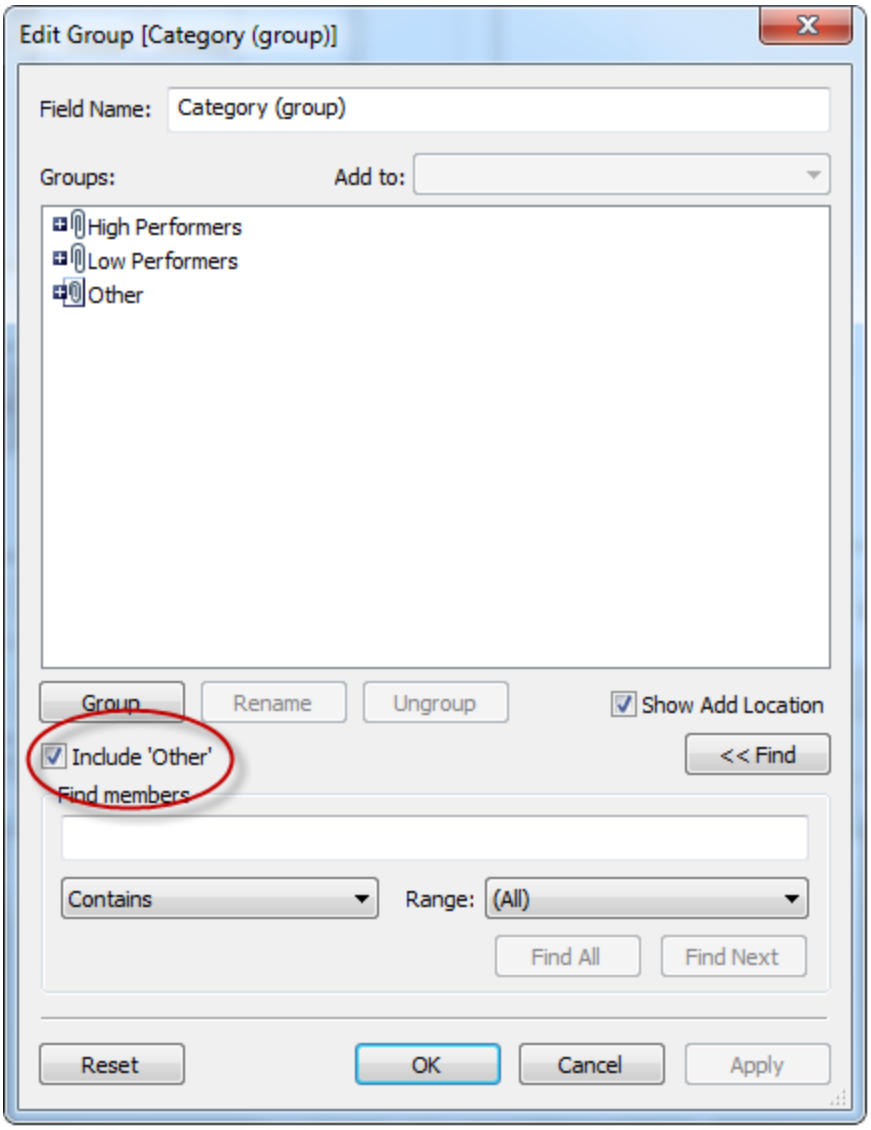

Finding Members in Group Dialog Box

When you create groups from a large dimension with many members, use the Find option to quickly select the members you are looking for and add them to an existing group.

1. Show the find options by clicking the **Find** button at the bottom of the dialog box.

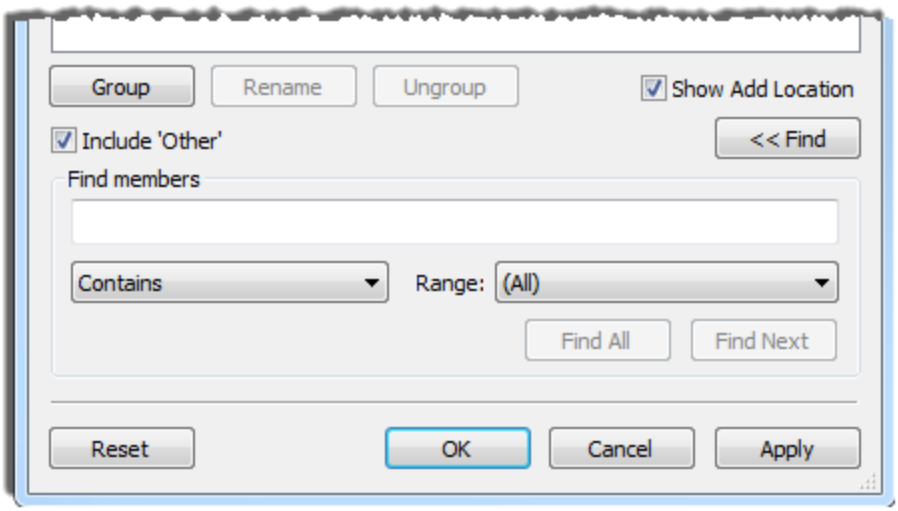

2. Type all or part of the member name into the text box and select an appropriate result criteria from the drop-down menu. You can select whether to find members that start with, contain, or are an exact match to the search term.
3. Select a Range to search in. You can select to search all members, 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 on the **Add to** drop-down menu at the top of the dialog box.

Sets

Sets are custom fields that define a subset of data based on some conditions. A set can be based on a computed condition, for example, a set may contain customers with sales over a certain threshold. Computed sets update as your data changes. Alternatively, a set can be based on specific data point in your view. You can use sets to compare and ask questions about a subset of data. For example, in a scatter plot showing product sales, you may select the products with the most sales and add them to a set. You may then create another set that contains the products with the highest return rates. With these two sets you can ask questions like "What percent of my total sales is from high returned products?" Or you may go further and compare the two sets to each other to create a new set that contains only high sales products that are also returned frequently.

Tableau displays sets at the bottom of the Data window and labels them with the set icon .

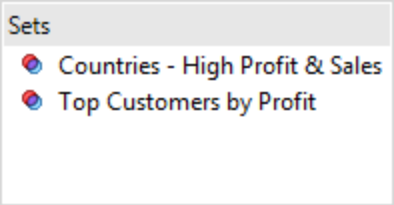

When you drag a set to the view, you can choose to aggregate the members using In/Out mode or to list all of the individual members of the set. See Using Sets to learn more about these modes.

If you are connected to a multidimensional data source, you may see additional sets that were created in your data source. These sets are also listed in the Sets area of the Data window but are indicated with the server named set icon. A server named set can be used in the same way as any other sets you create in Tableau.

Sets can be automatically generated as a result of an action. These sets are indicated by the action set icon . See **Actions** on page 543.

Finally, when working with workbooks connected to Tableau Server or Tableau Online, User Filters display in the Sets area of the Data window using the user filter set icon user. See **User Filtering** on page 843.

> If you are connected to a relational data source, creating a set that is based on a continuous date field will use the exact dates instead.

Creating a Set

There are many ways to create and edit sets in Tableau. To create a set, select marks or headers in the view or right-click a field in the data window. You can also create a set that is based on a filter you've already defined. Regardless of how you create the set, it is either constant or computed. A constant set contains a specified list of members based on one or more dimensions. For example, a constant set might contain the specific names of products that you are keeping an eye on. A computed set is dynamic and the members change as the underlying data changes. For example, a computed set may contain the top 5 products by total sales. As the sales data changes, the companies included in the set will also change.

Constant Sets

The members of a constant set are fixed and do not change. To create a constant set, you need to select the members you want to include. A constant set can be based on a single dimension or multiple dimensions.

1. Select one or more marks or headers in the view.
2. Right-click and select **Create Set** or click the **Create Set** option on the tooltip.

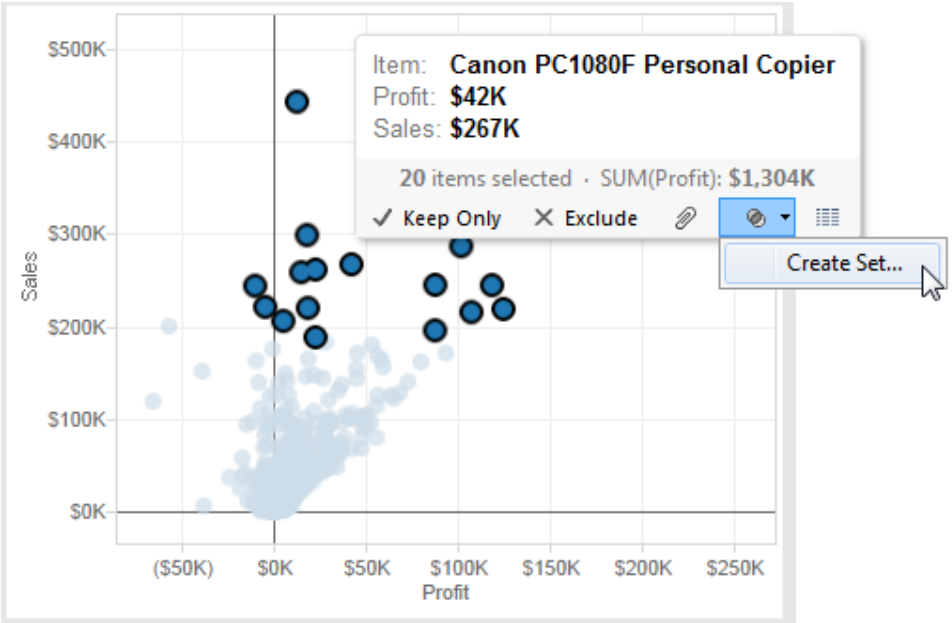

3. In the Create Set dialog box, type a name for the set.
4. Optionally complete any of the following:
   - By default, the set includes the members listed in the dialog box. You can select the option to **Exclude** these members instead. When you exclude, the set will include all of the members you didn't select.
   - Remove any dimensions that you don't want to be considered by clicking the red "x" icon that displays when you hover over a column heading 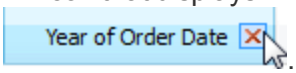

.
   - Remove any specific rows that you don't want to include in the set by clicking the red "x" icon that displays when you hover over the row ✕.
   - If the marks you selected represent multiple dimensions, each member of the set will be a combination of those dimensions. You can specify the character that separates the dimension values.
   - Select **Add to Filters shelf** to automatically move the set to the Filters shelf once it is created.

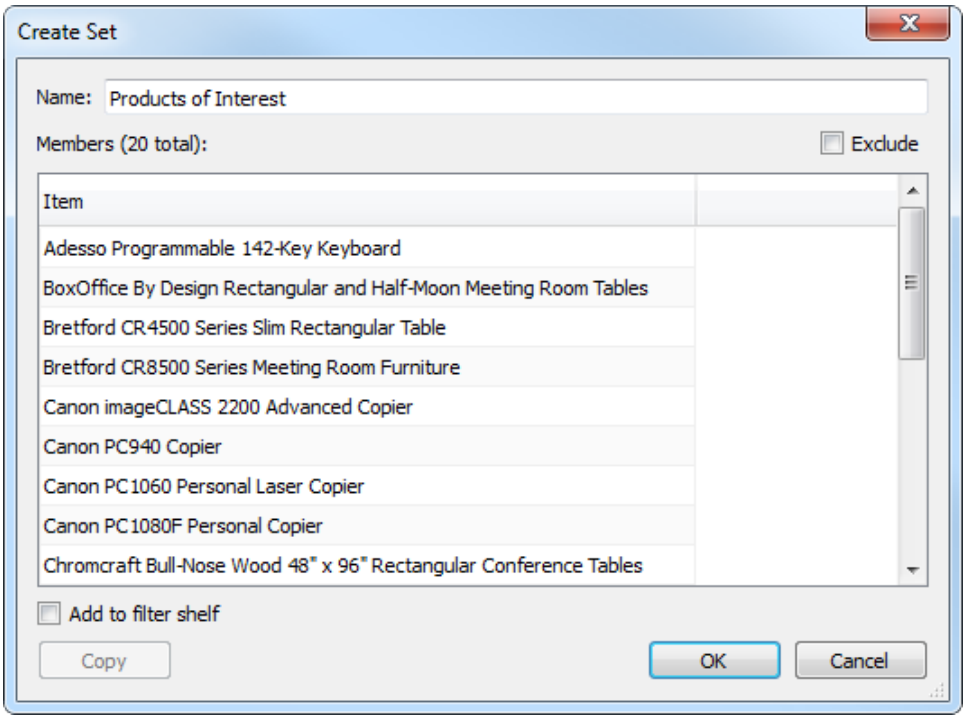

5. When finished, click **OK**.

Computed Sets

The members of a computed set are dynamic and change when the underlying data changes. Computed sets can only be based on a single dimension.

1. Right-click dimension in the Data window.
2. Select **Create Set**.
3. In the Create Set dialog box, type a name for the set.
4. On the General tab in the Create Set dialog box, select one or more values that will be considered when computing the set. You can alternatively select the **Use All** option to always consider all members even when new members are added or removed.

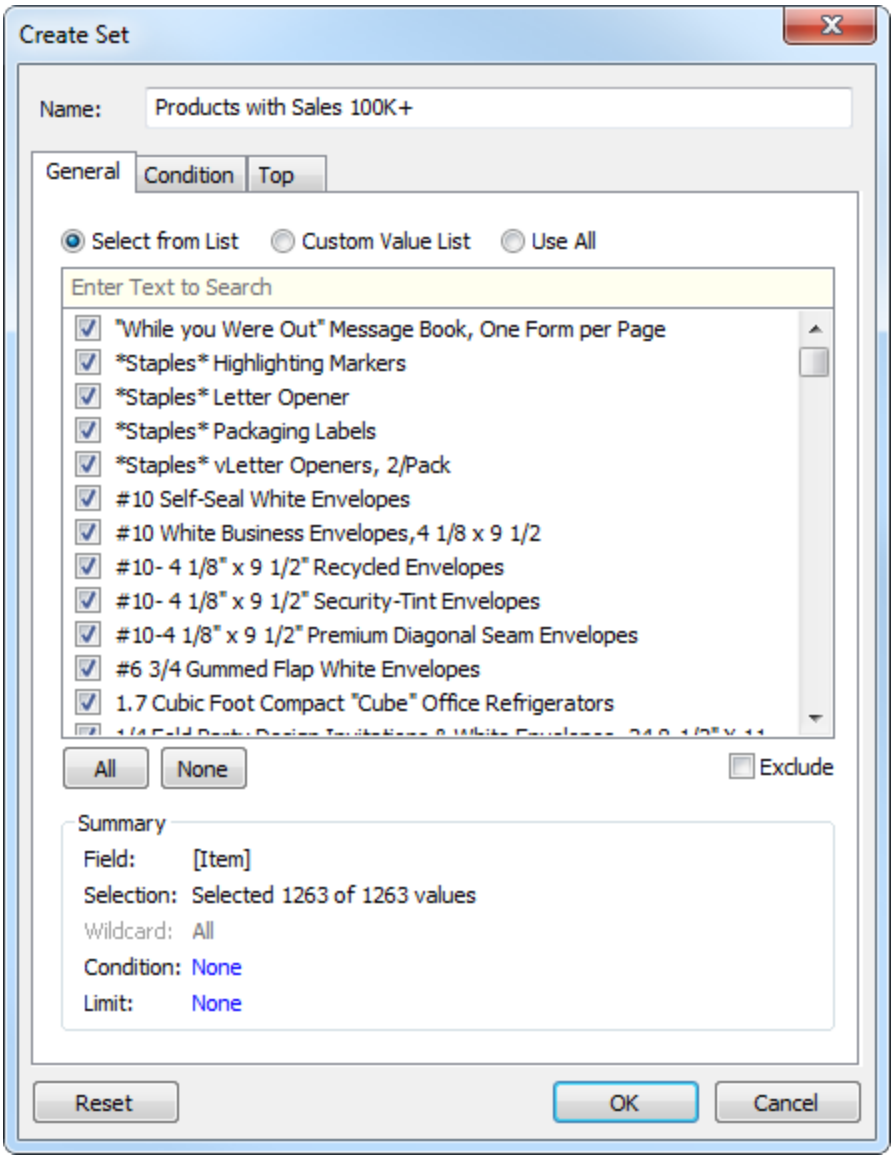

5. Define how to compute the members by adding conditions and limits:
    - **Condition**: Use the Condition tab to define rules that determine what members to include in the set. For example, you may specify a condition that is based on the total sales that only includes products with sales over $100,000.

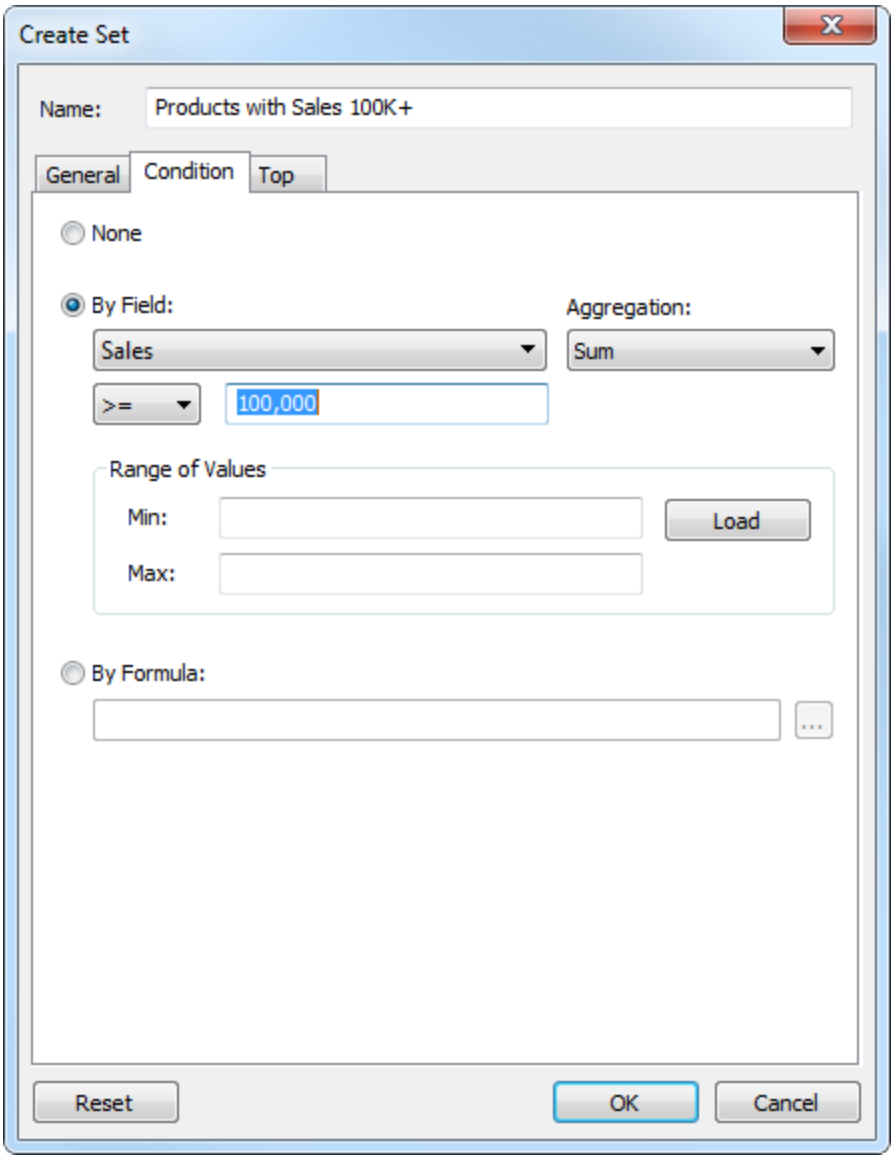

Set conditions work the same as filter conditions. See **Adding Conditions to Filters** on page 343 to learn more.

- **Top**: Use the Top tab to define limits on what members to include in the set. For example, you may specify a limit that is based on total sales that only includes the top 5 products based on their sales.

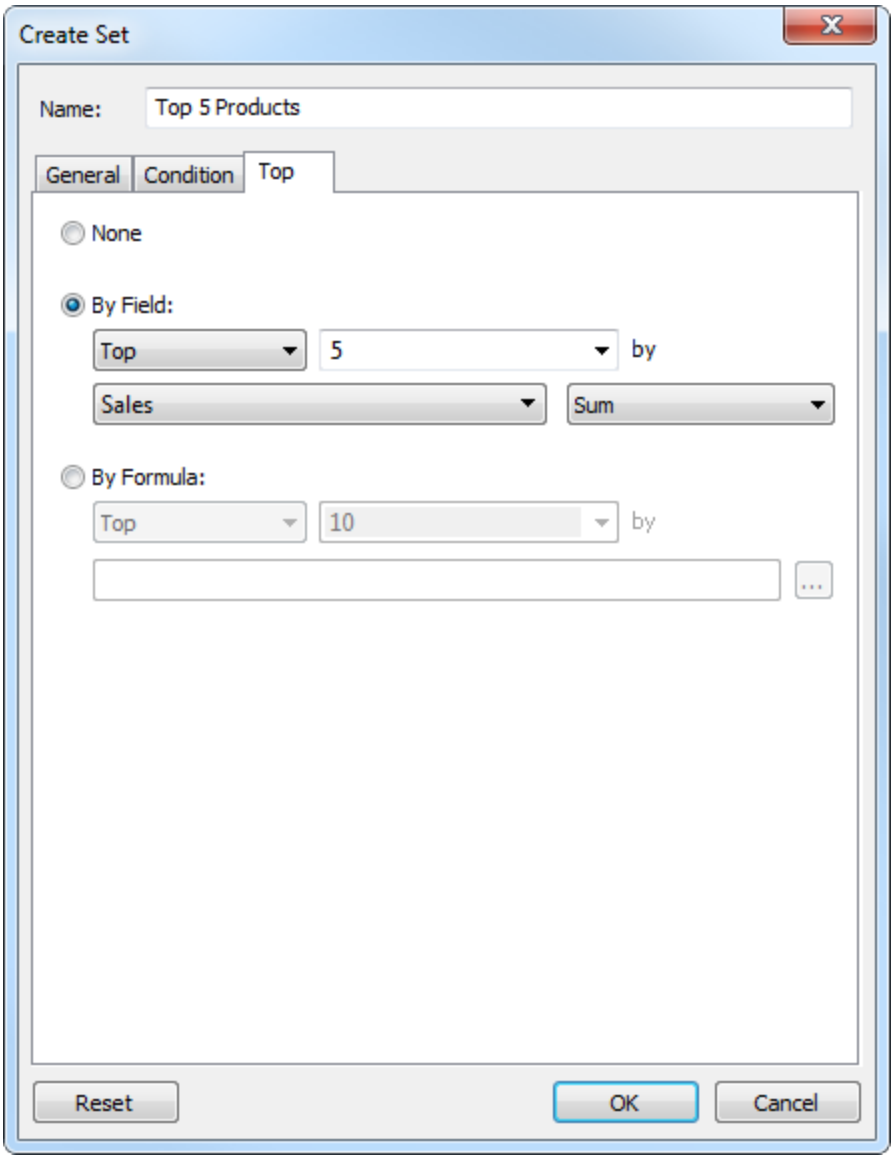

Set limits work the same as Filter limits. See **Adding Limits to Filters** on page 344 to learn more.

6. When finished, click **OK**.

Using Sets

After you create a set, it displays at the bottom of the Data window in the Sets area. You can then drag sets into the view like any other field. For example, you can drag the set to the Filter shelf to quickly filter the view to only show the members of the set. Alternatively, you can drag the set to the Marks card or the Rows and Columns shelves. When you drag a set to the view, you can choose to show the members of the set or aggregate the members into In/Out categories.

Show In/Out of Set

In most cases, when you drag a set to the view, Tableau displays the set using the In/Out mode. This mode separates the set into two categories: In, which contains the members in the set; and Out, which contains any members that are not part of the set. For example, in a set defined as the top 25 customers, the top customers would be part of the In category and all other customers would be part of the Out category. Using the In/Out mode makes it easy to compare the members in the set to everything else. You can use this mode to answer questions like "What percent of my total revenue comes from members of this set?"

Switch a set to use In/Out mode by selecting **Show In/Out of Set** on the field menu.

When a set is in In/Out mode, the field on the shelf is prefaced by the text, "IN/OUT" followed by the set name.

> **Note:**In/Out mode is not available in workbooks created before version 8.2 that use Microsoft Excel or text file data sources, workbooks that use the legacy connection, or workbooks that use Microsoft Access data sources.

Show Members in Set

As an alternative to showing the set using In/Out mode, you can list the members in the set. Showing the members in the set automatically adds a filter to the view that includes only the members of the set.

Switch a set to list the individual members by selecting **Show Members in Set** on the field menu.

> **Note:** For cubes, you can choose to display the fully qualified member names by right-clicking the set in the Data window and selecting **Qualify Member Names**.

Combining Sets

You can combine two sets to compare the members. When you combine sets you create a new set containing either the combination of all members, just the members that exist in both, or members that exist in one set but not the other. Combining sets allows you to answer complex questions to understand cohorts of data. For example, to determine the percentage of customers who purchased both last year and this year, you can combine two sets containing the customers from each year and return only the customers that exist in both sets. Another example would be to determine what products sell the most but also have the highest return rate.

To combine two sets, they must be based on the same dimensions. That is, you can combine a set containing the top customers with another set containing the customers that purchased last year. However, you cannot combine the top customers set with a top products set. In the latter case, the one set is based on the Customer dimension while the other is based on the Product dimension.

1. Select two sets in the Data window that you want to combine.
2. Right-click the sets and select **Create Combined Set**.
3. In the Create Set dialog box, type a name for the new combined set.
4. Verify that the two sets you want to combine are selected in the two drop-down menus.
5. Select one of the following options for how to combine the sets:
    - **All Members in Both Sets** - the combined set will contain all of the members from both sets.
    - **Shared Members in Both Sets** - the combined set will only contain members that exist in both sets.
    - **Except Shared Members** - the combined set will contain all members from the specified set that don't exist in the second set. These options are equivalent to subtracting one set from another. For example, if the first set contains Apples, Oranges, and Pears and the second set contains Pears and Nuts; combining the first set except the shared members would contain just Apples and Oranges. Pears is removed because it exists in the second set.
6. Optionally specify a character that will separate the members if the sets represent multiple dimensions.
7. When finished, click **OK**.

Note: This functionality is not available in workbooks created before version 8.2 that use Microsoft Excel or text file data sources, workbooks that use the legacy connection, or workbooks that use Microsoft Access data sources.

Examples of Sets

There are many ways you can use sets to answer complex questions and compare cohorts of data. Below are some examples of ways you can use sets to define and compare subsets of data.

How do members of a set contribute to the total?

You may have all kinds of questions surrounding how the members in a set contribute to the overall total. For example, what percent of total sales come from repeat customers? You can answer these types of questions using the IN/OUT mode for set. The example below uses sales data to calculate the percent of total sales from customers who have purchased at least 20 times.

1. Right-click the **Customer Name** dimension in the Data window and select **Create Set**.

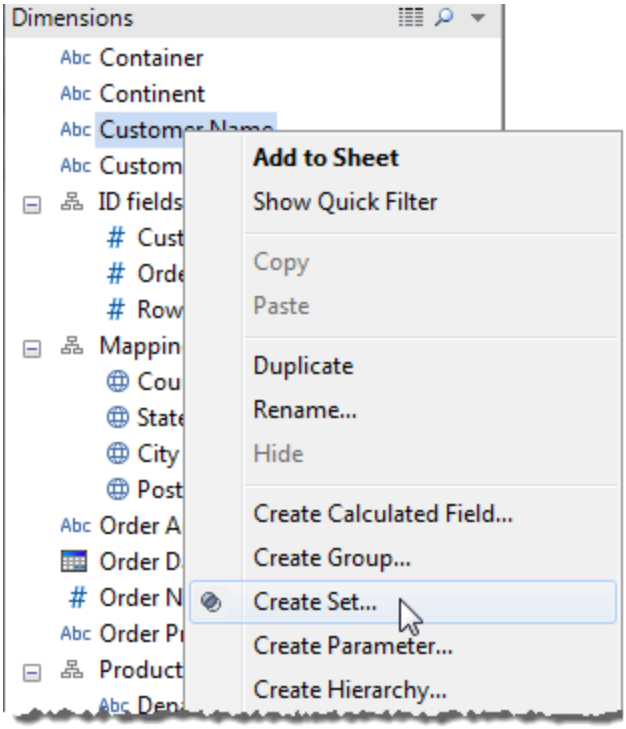

2. In the Create Set dialog box, type a name for the set. In this example, we'll call the set, "Frequent Customers"
3. Select the **Use All** option so the condition always applies to all values even when new customers are added.

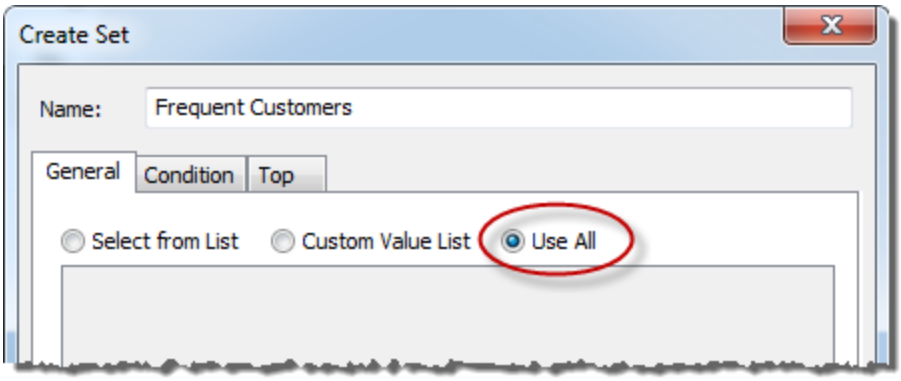

4. On the Condition tab, select **By Field** and define a condition that only includes customers when Count (Distinct) of Order ID is greater than or equal to 20.

5. When finished, click **OK**.
6. Drag the new set from the Sets area at the bottom of the Data window to the Rows shelf.
7. Drag the **Sales** measure to the Columns shelf. The view now shows the total sales for frequent customers and total sales for all other customers.

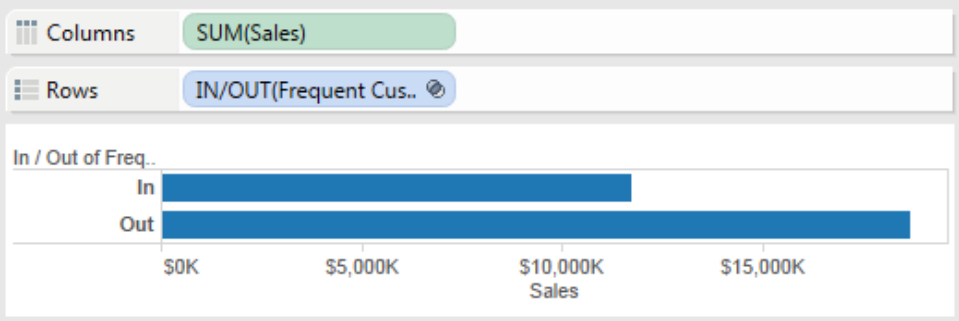

8. Finally, on the field menu for Sales, select **Quick Table Calculation** > **Percent of Total**.

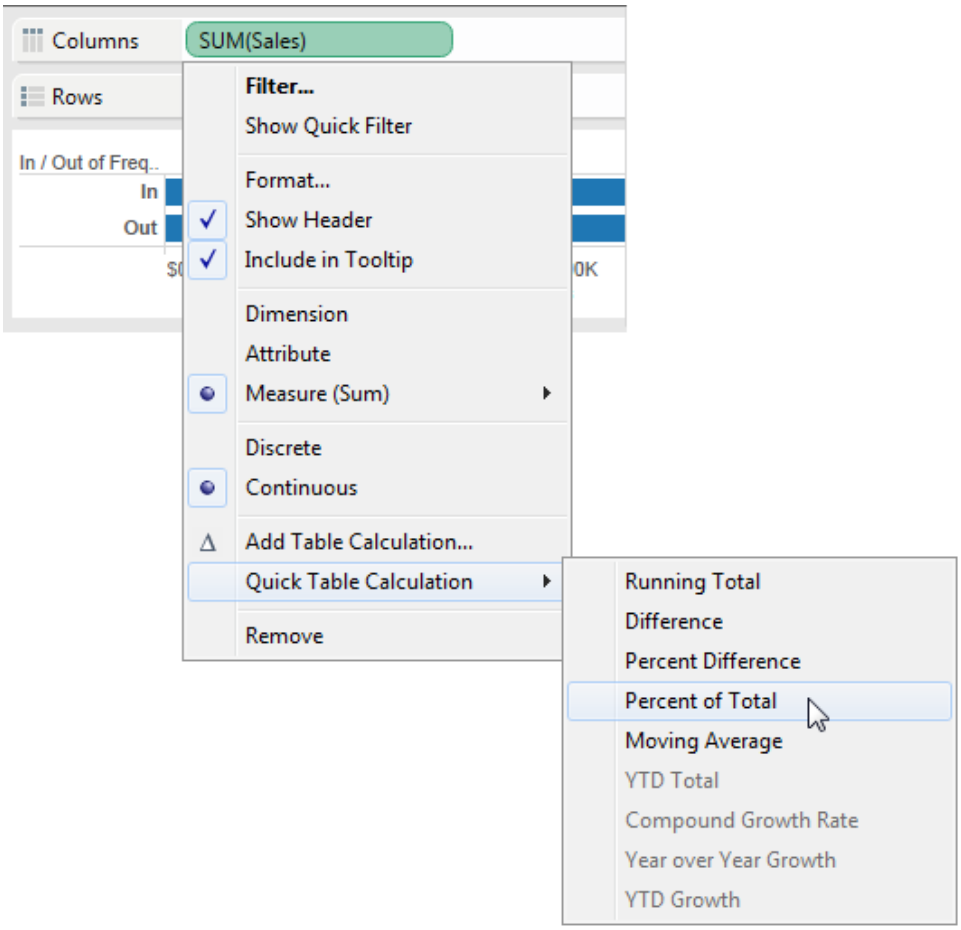

The view now shows that frequent customers make up about 39% of the overall sales.

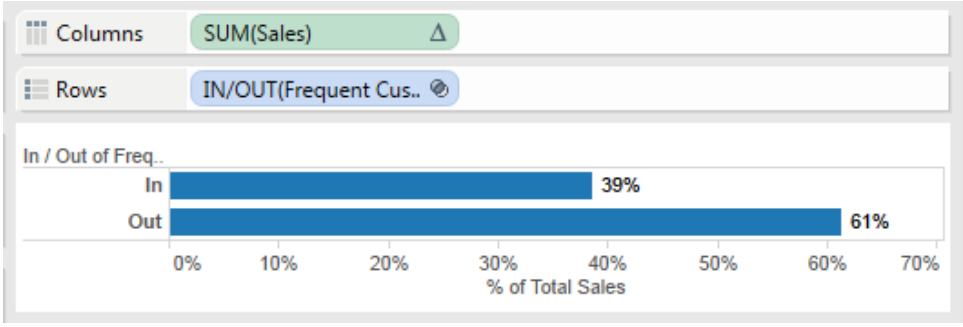

How many members of a set exist in another set?

Another common use of sets is to compare subsets of data or cohorts. For example, you may wonder how many customers that purchased last year also purchased this year. Or if a customer purchased a specific product, what other products did they buy? You can answer these types of questions by creating multiple sets and combining them. The example below uses sales data to determine how many customers that purchased in 2010 also purchased in 2011.

1. Drag the **Customer Name** field to the Rows shelf.
2. Drag the **Order Date** field to the Filters shelf.
3. In the Filter Field dialog box, select **Years** and click **Next**.

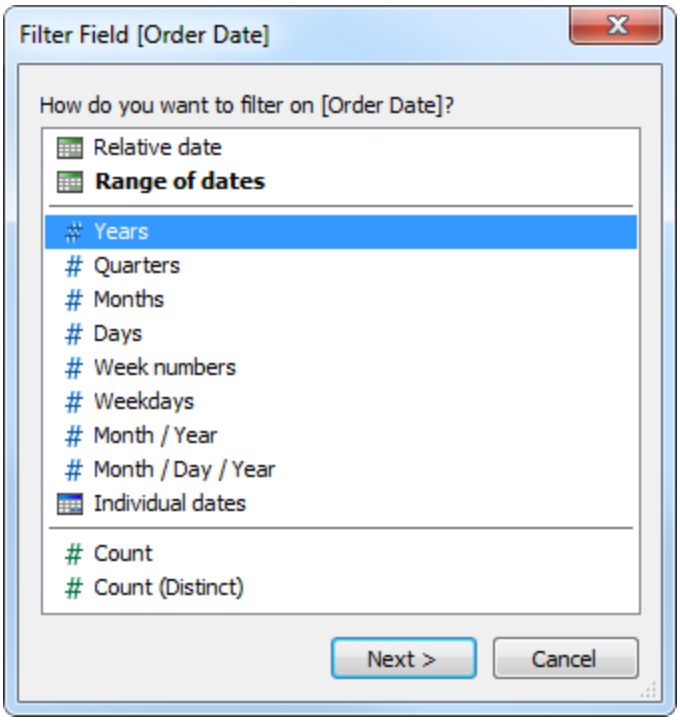

4. In the Filter dialog box, select **2010** and click **OK**.

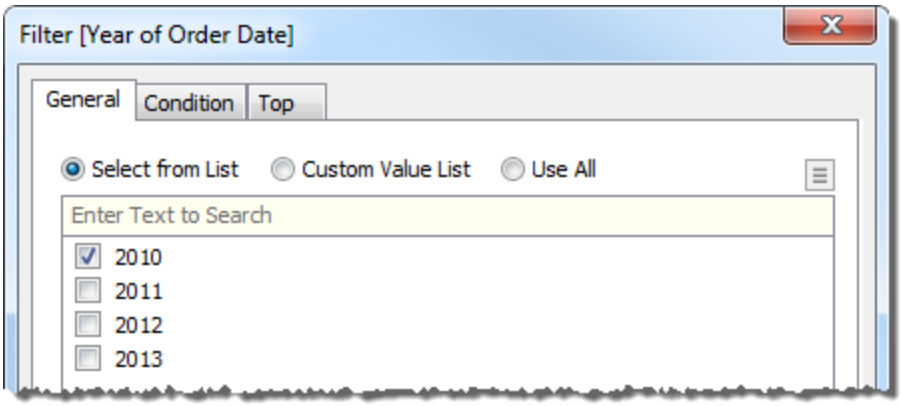

5. Back in the view, press **CTRL + A** on your keyboard to select all of the customers.
6. Right-click the selection and select **Create Set**.

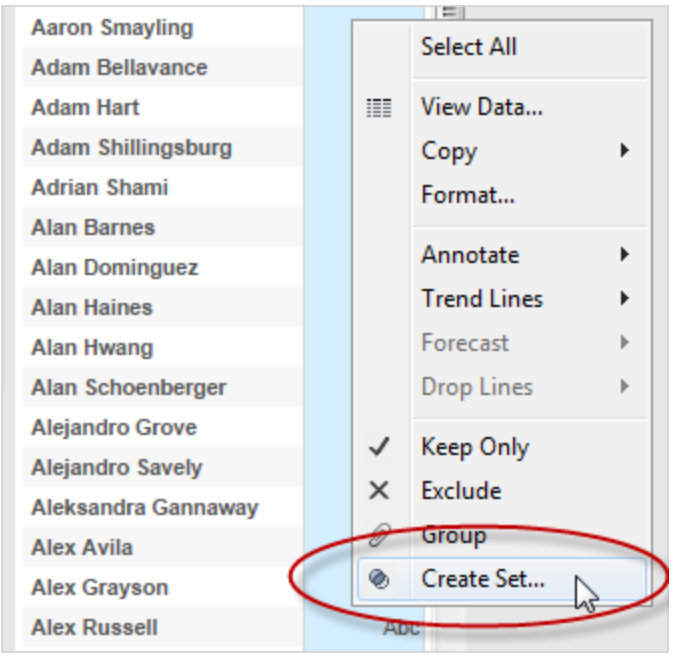

7. In the Create Set dialog box, type a name for the set. In this example, we'll call the set "Customers (2010)".

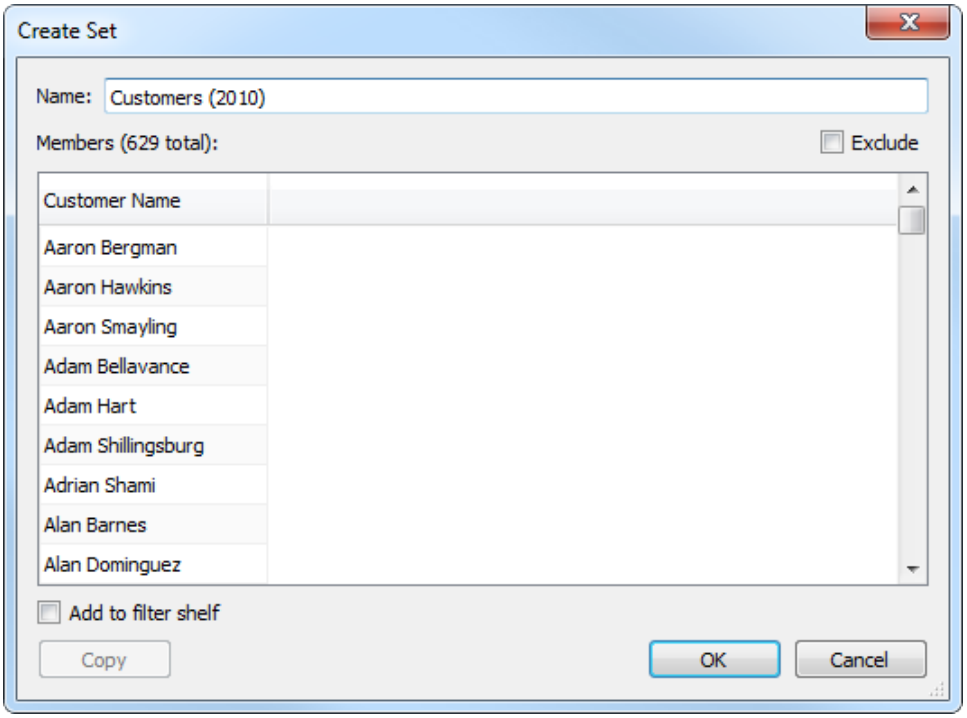

8. When finished, click **OK**.
9. On the field menu for the Order Date field on the Filters shelf, select **Filter**.

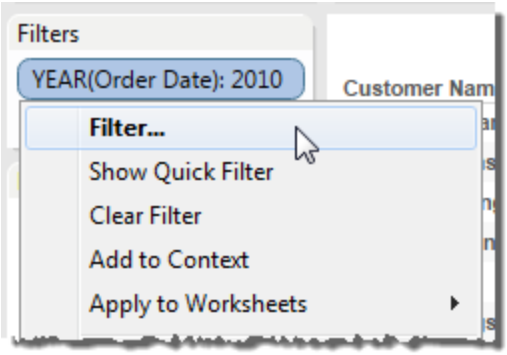

10. Change the filter to only include 2011 instead of 2010. Then click **OK**.

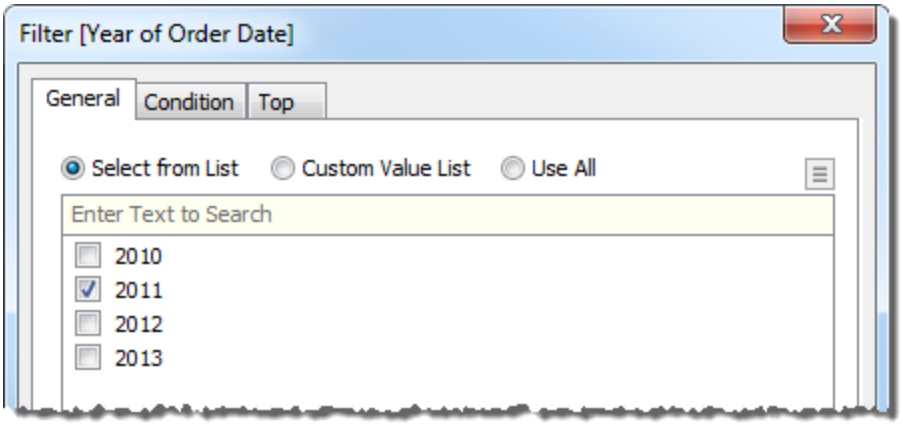

11. Again, press **CTRL + A** on your keyboard to select all of the customers.
12. Right-click the selection and select **Create Set**.
13. In the Create set dialog box, type a name for the set. This set will be called "Customers (2011)".

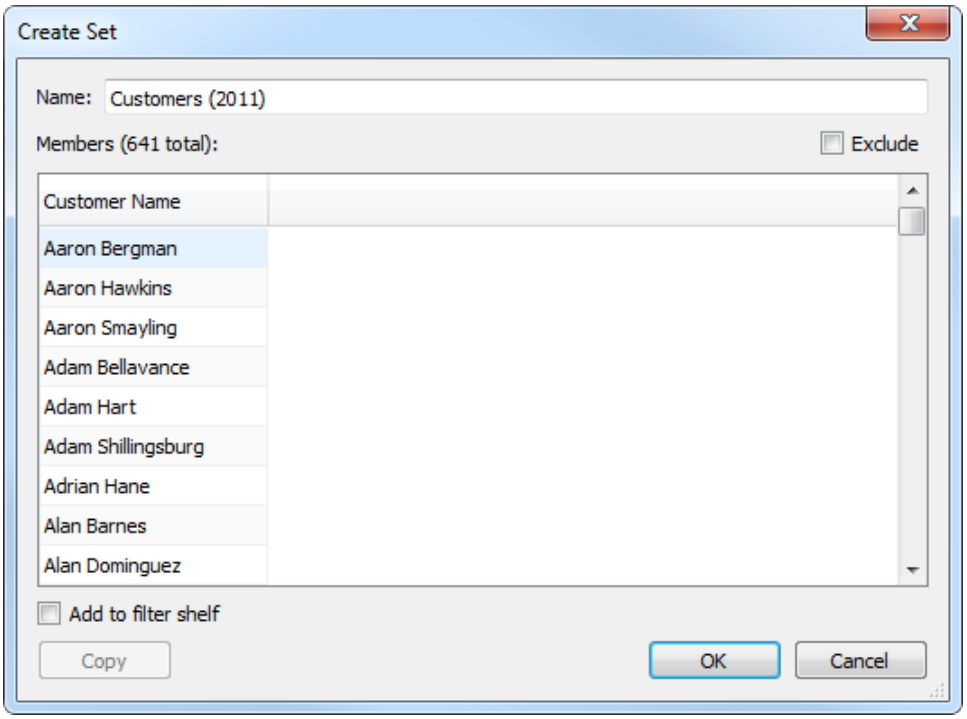

14. When finished, click **OK**.
15. Hold the CTRL key on your keyboard to select both the **Customers (2010)** set and the **Customers (2011)** set in the Data window.
16. Right-click the selection and select **Create Combined Set**.

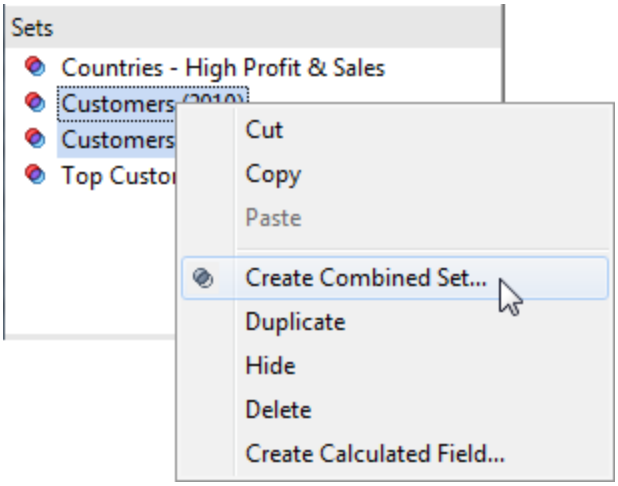

17. In the Create Set dialog box, type a name for the new set. In this example, we'll call the set "Customers (2010 & 2011)".
18. Make sure the correct two sets are selected in the drop-down menus.
19. Select the option to include **Shared Members in Both Sets**.

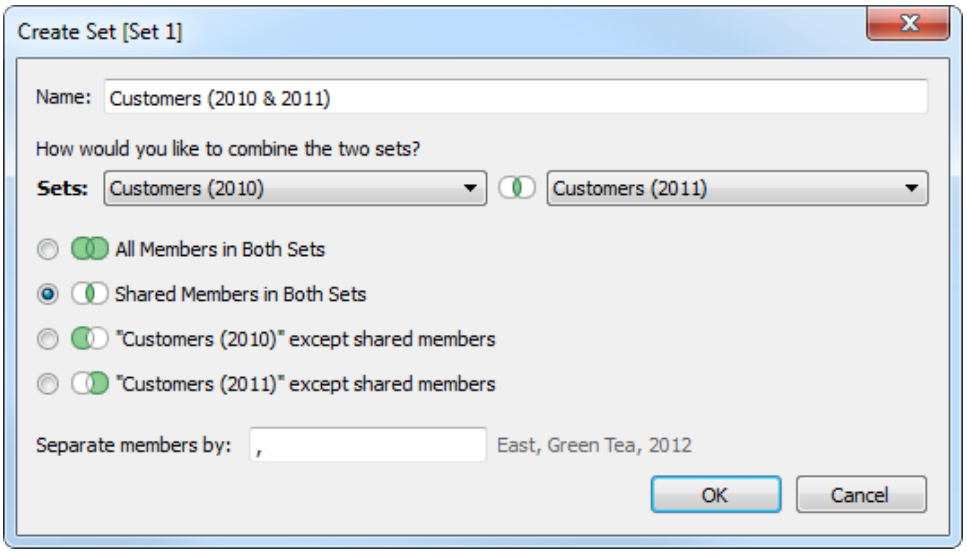

20. When finished, click **OK**.
21. Create a new sheet by clicking the **New Worksheet** tab at the bottom of the workbook.

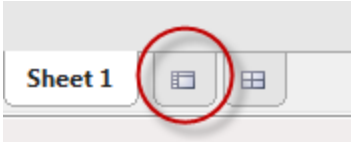

22. Drag **Customer Name** to the Rows shelf.
23. On the Customer Name field menu, select **Measure** > **Count (Distinct)**.
24. Finally, drag the **Customers (2010 & 2011)** field to the Filters shelf. You can see that 511 customers who purchased in both 2010 and 2011.

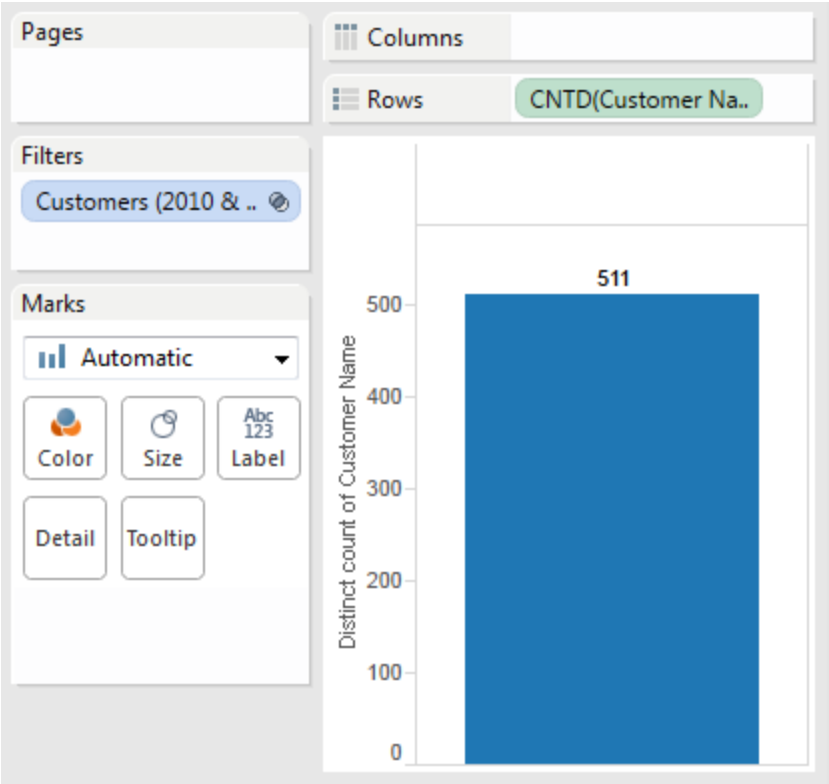

Hierarchical Sets and Descendants

A hierarchical set filters data to the selected members and all of their descendants. For example, a set named **Dairy** is created from the **Product** hierarchy. As shown below, it includes only the Dairy product department.

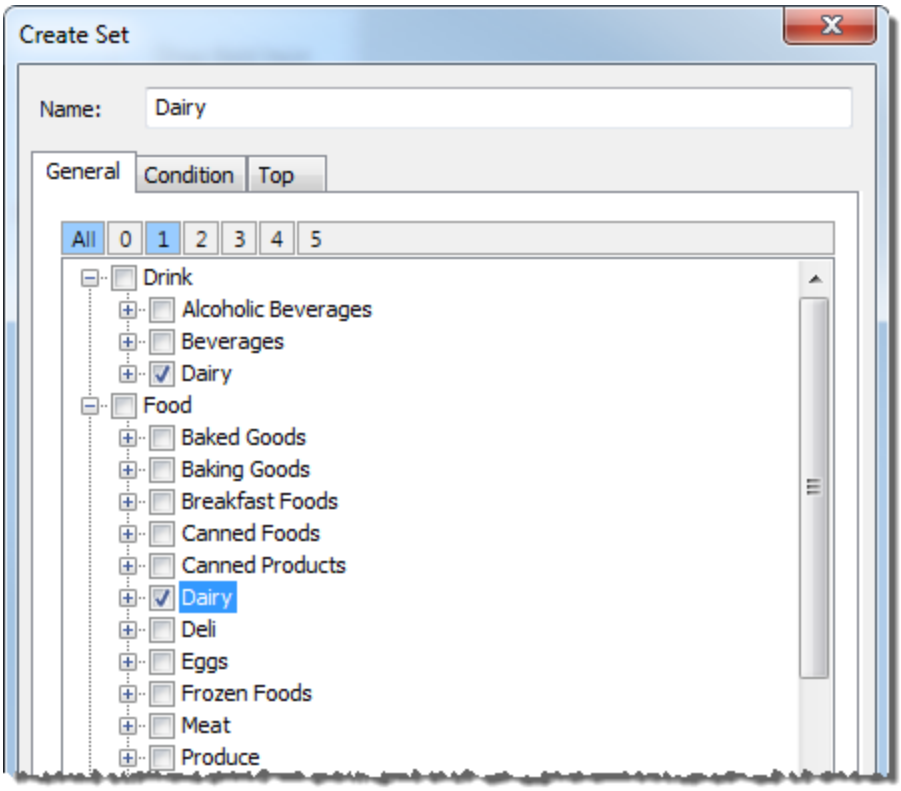

Consider the following view. The **Product Category** dimension is placed on the **Rows** shelf and the **Store Sales** measure is placed on the **Columns** shelf.

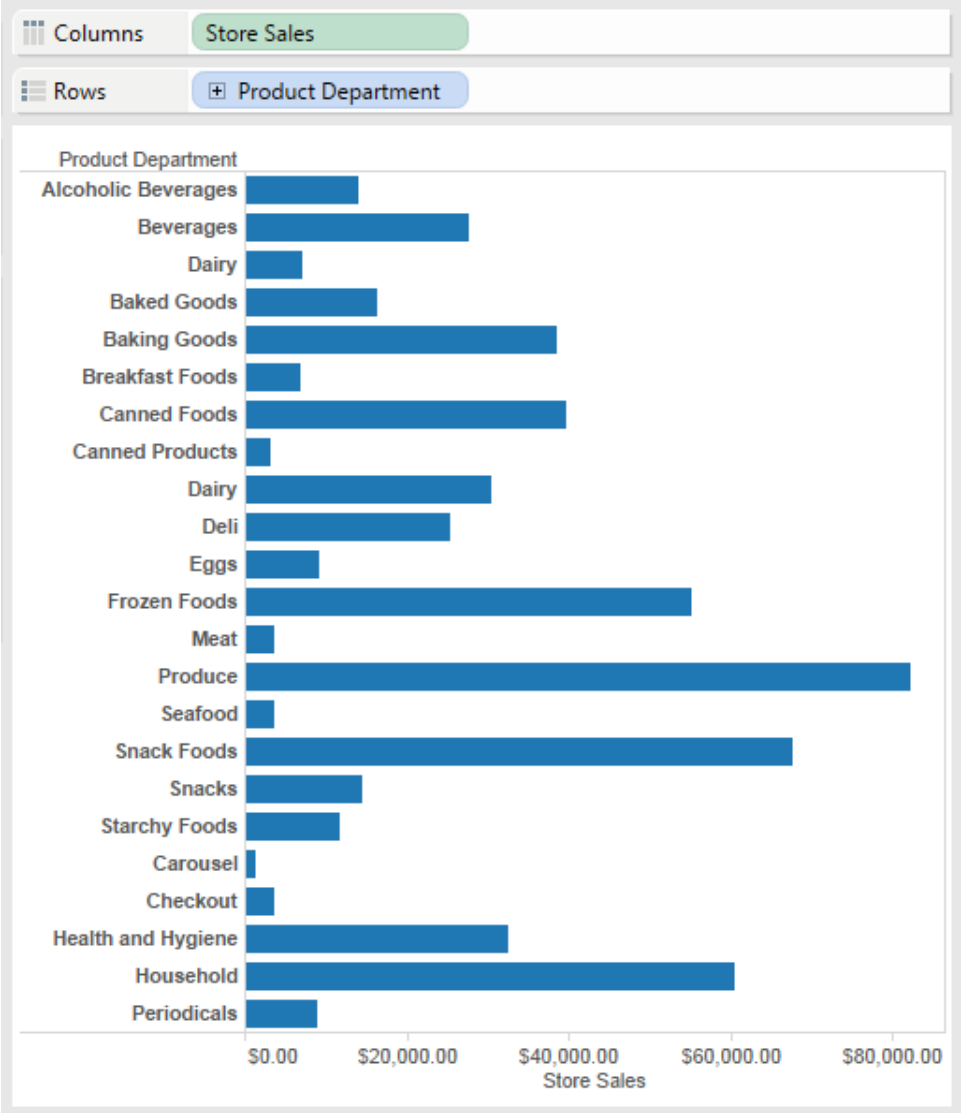

If you place the **Dairy** set on the **Filters** shelf, you can see that the view is filtered to include only the Dairy product categories.

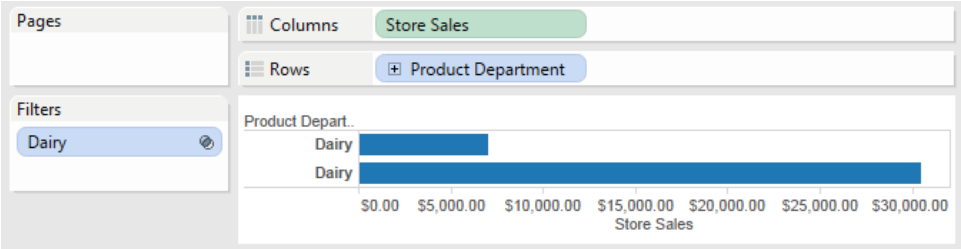

As shown below, you can drill down into **Product Department** to reveal the **Product Category**, **Product Subcategory**, and **Brand Name** levels. As these descendants are revealed, row headers are added to the view. This is because a set filter allows you to view the levels of detail contained within the filtered members.

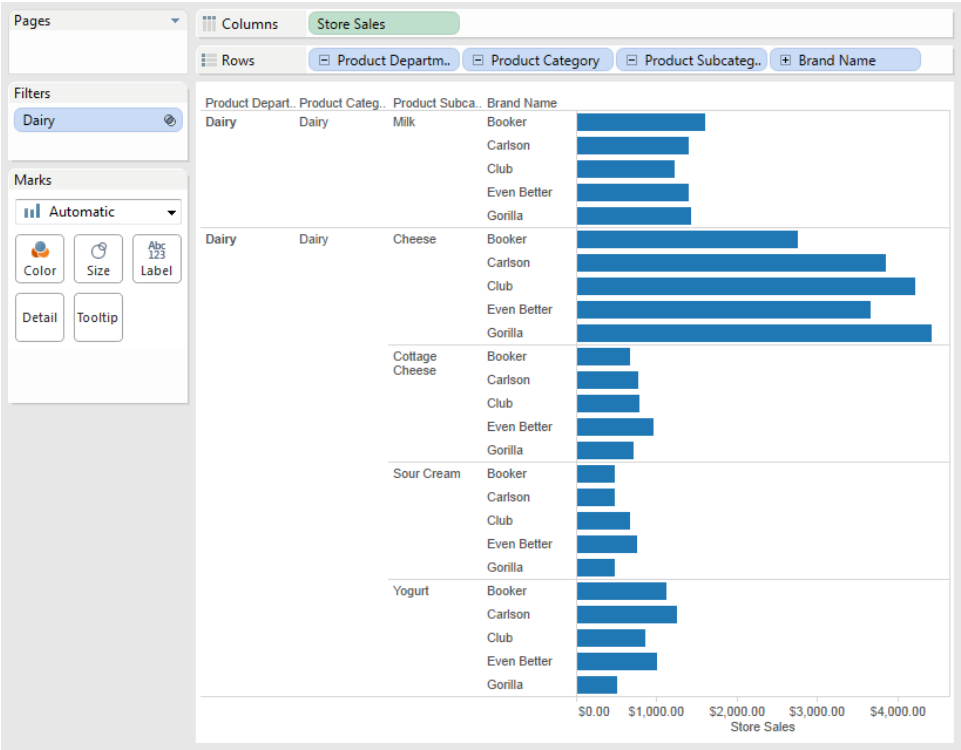

Dates and Times

Working with dates in Tableau differs depending on whether you are using a Relational or Multidimensional Data Source. This section below discusses these differences.

Dates in Multidimensional Data Sources

> In Tableau, multidimensional data sources are supported only in Windows.

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 makes it possible to look at data levels different ways, such as Months by Year, Months by Quarter, Weekends, etc. These levels are represented as attributes of the hierarchy. Hierarchies and attributes are defined when the cube is created and you cannot modify them in Tableau. For example, the **Year** dimension from an Oracle Essbase data source is shown below.

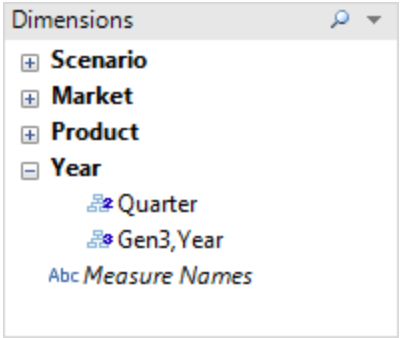

When you place a multidimensional date on a shelf, the field is treated like any other dimension. For example, you can drill down, drill up, and so on.

Dates in 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 date or date-time icon. For example, the Order Date and Ship Date dimensions from an Excel data source are shown below.

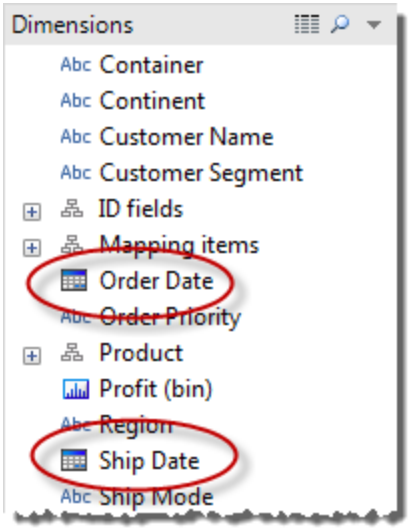

When you place a relational date on a shelf, the field name is automatically modified to reflect the default date level. Tableau defines the default date level to be the level at which there are multiple instances. For example, if the date field includes multiple years, the default level is year. However, 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 then turns green in the Data window; now when you use that dimension in a view, 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 field in the view to continuous while it is on a shelf by selecting **Continuous** on its context

menu (which you can see when you right-click the field). The field on the shelf turns green but the field in the Data window is still discrete.

Date Properties for a Data Source

You can set date properties for a data source. To do so, right-click a data source in the Data window and choose Date Properties:

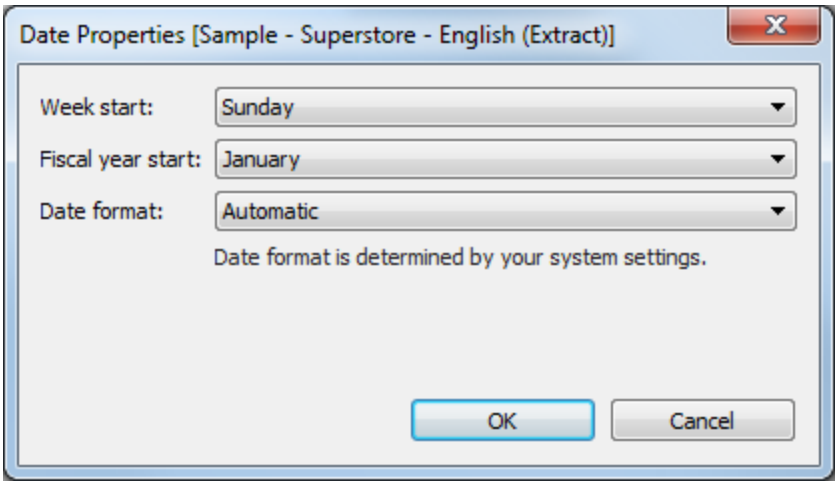

The options in the Date Properties dialog box are:

- **Week start** - Specifies which day is considered the first day of the week. You can override the **Week start** value you set at the data source level by including a `start_of_week` value with certain date functions (`DATEADD`, `DATEDIFF`, `DATENAME`, `DATEPART`, and `DATETRUNC`). See **Date Functions** on page 887.

  The initial value for this setting is in some cases determined by your data source. If the data source does not initialize the start-of-week day, then Tableau consults the system's locale settings.The day which is considered the first day of the week varies from region to region. For example, Sunday is the first day of the week in the US, while Monday is the first day in the EU.

  This setting is disabled for cubes because this information is defined when cube designer creates the date/time dimension.

  After a data source is created, the **Week start** value does not automatically change if for any reason the default start-of-week setting changes in the host operating system (for example, if the data source is uploaded to Tableau Server in a different country). Calendar controls reflect the workbook locale rather than the data source's **Week start** setting. The exception is that a week trunc start of week that shows the week number in it will use the data source's **Week start** setting in order to provide a consistent week number value in the calendar.

- **Fiscal year start** - Specifies which month is considered the first month of the fiscal year. To specify whether a date dimension uses the standard calendar (Jan. 1 - Dec. 31) or the fiscal calendar, right-click the field in the Data window and select **Default**

Properties > Calendar Type > and then select **Standard Calendar** or **Fiscal Calendar**. For any date field in the view, you can then specify whether to use the standard calendar or the fiscal calendar. For details, see **Fiscal Dates** on page 445.

- **Date format** - Specifies the default format for data dimensions, as displayed, for example, in tooltips.To override the default date format for a date dimension, right-click the field in the Data window and select **Default Properties > Date Format >** and then select one of the available formats.

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 multidimensional dates, the levels available in the context menu are given by the levels defined in the date hierarchy. For relational dates, you can select between the discrete date levels at the top of the menu or the continuous date at the bottom. A preview of each date level is shown. Below is the field menu for a relational date with the discrete Year level selected.

In Tableau, multidimensional data sources are supported only in Windows.

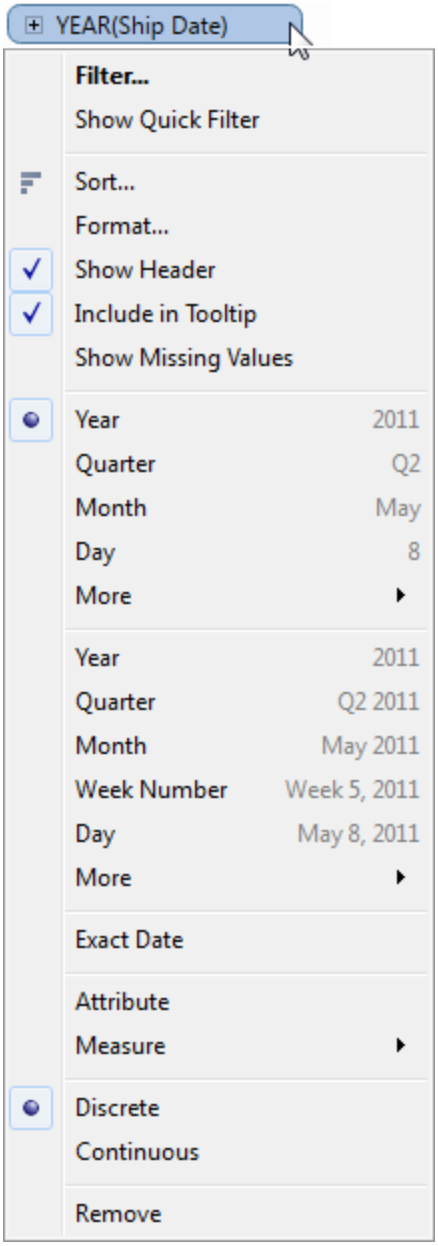

For relational dates, there are many different date levels available, as you can see above. When you select a particular level, Tableau asks the data source to perform a computation on the date field. For example, suppose a particular row in your data source has a date entry of 01/23/07. The year is 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 the configured Date Properties for the data source. See **Date Properties for a Data Source** on page 438.

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. You do this by drilling 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.

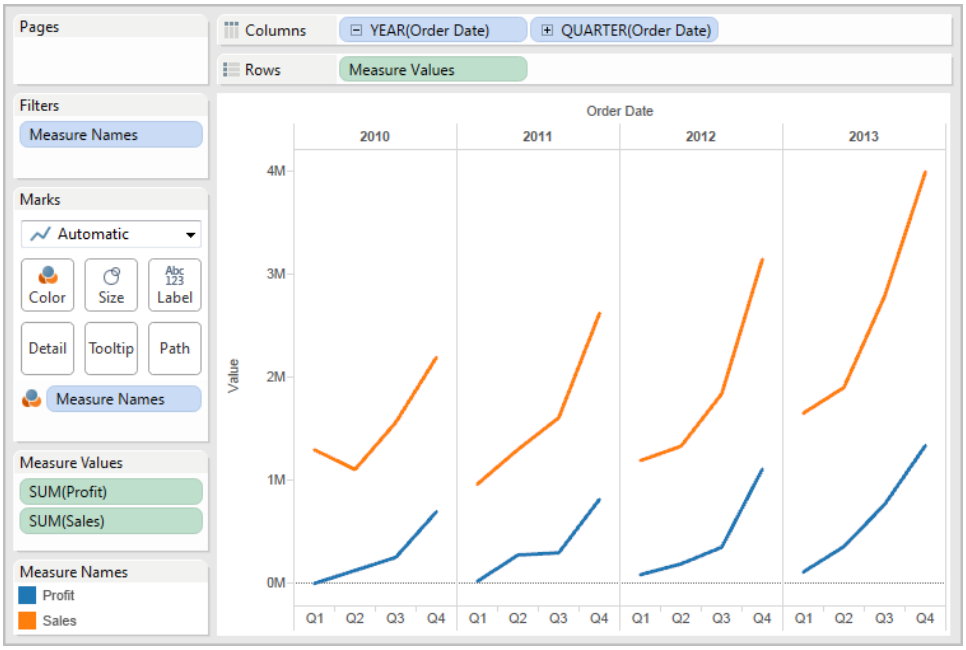

By removing one of the date fields from the Columns shelf, and then selecting **Month** from the remaining date field's context menu, you can display the data for each month across all years.

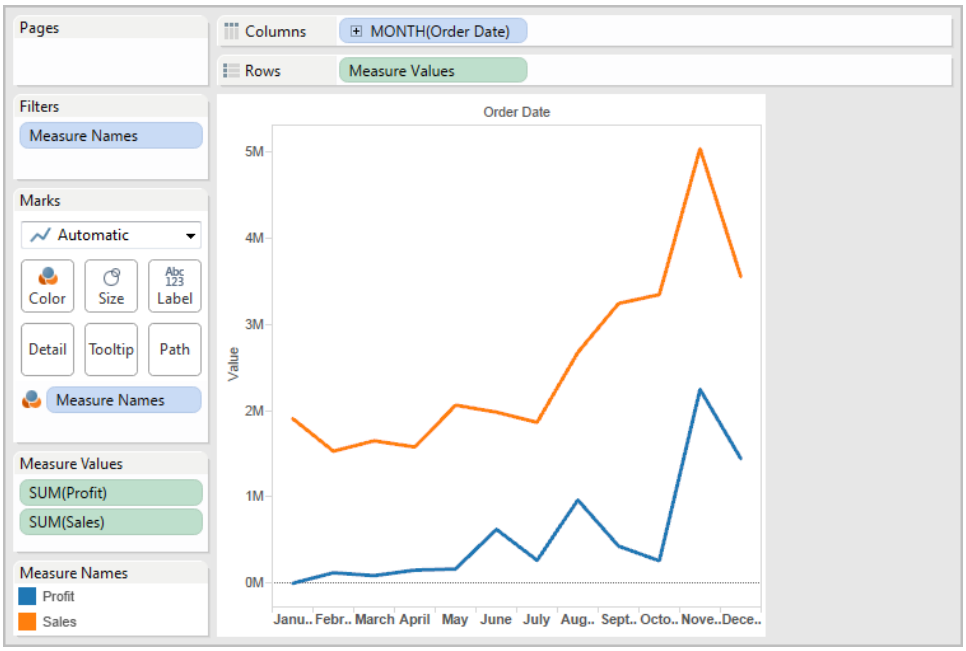

To display even finer granularity, select Month from the lower, continuous section of the field's context menu.

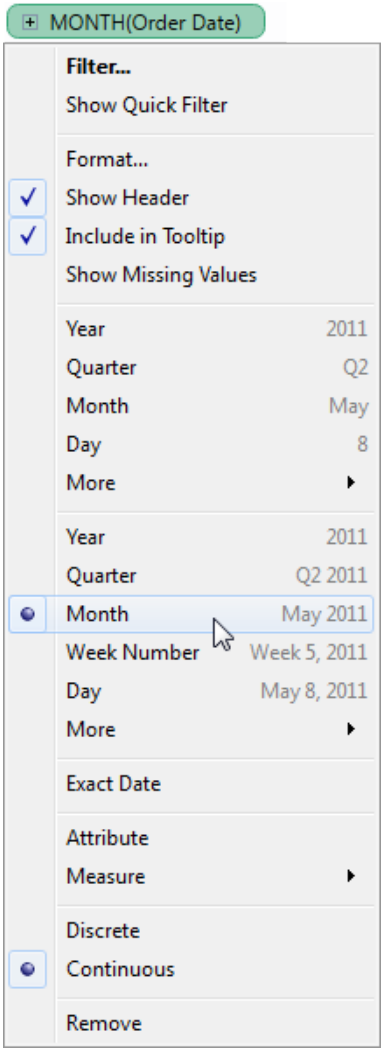

Tableau then displays the dates across the range of available years, at a month-by-month level of granularity.

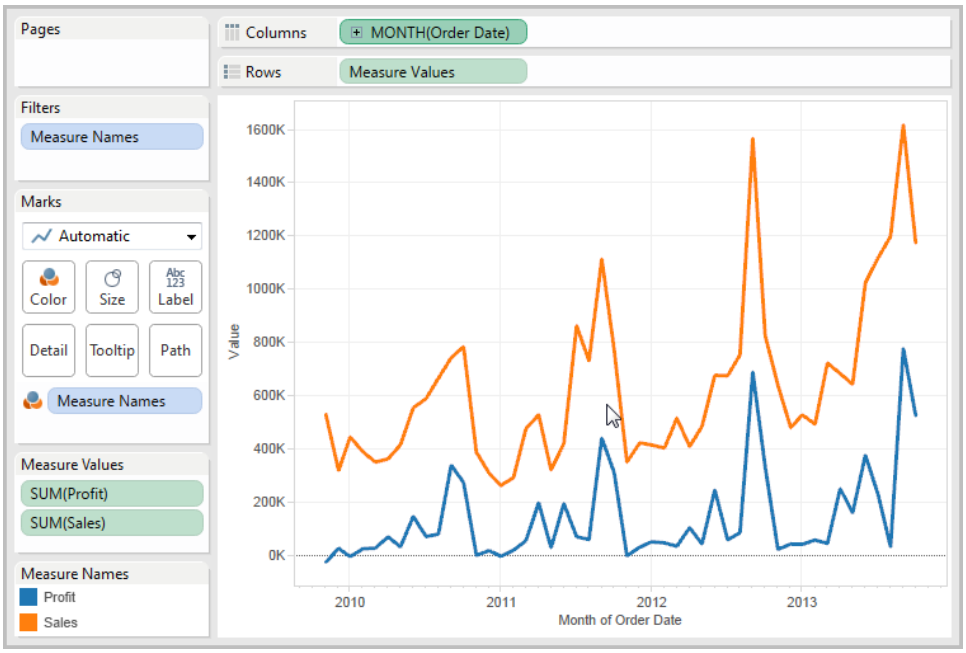

Custom Dates

You can create a custom date that is binned or truncated to a specific level. The custom date becomes a new field in the Data window. Custom dates are useful if you always use a date at a specific level or you want to create calculations that rely on a binned or truncated date. Custom dates can also be created using the `DATEPART` and `DATETRUNC` functions in a calculated field. Follow the steps below to create a custom date:

1. Right-click a date field in the Data window and select **Create Custom Date**.
2. In the Create Custom Date dialog box, type a name for the custom date. For example, Date (Quarter, Year).
3. Select the level you want to show the date at.
4. Specify whether you want the custom date to be discrete **Date Parts** or continuous **Date Values**.

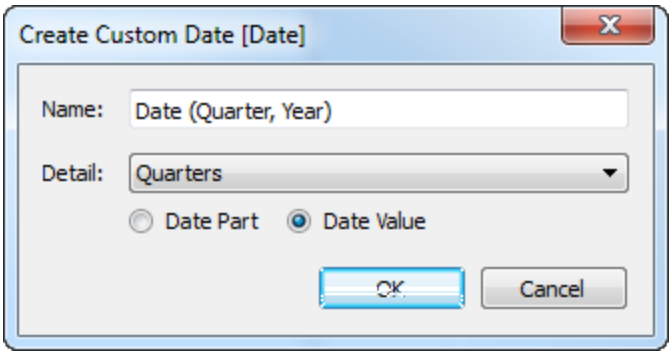

5. When finished, click **OK**.

Fiscal Dates

In some situations, a date field needs to be expressed in terms of an organization's fiscal year. Calendar years run from January 1st until December 31st, but an organization's fiscal year might start in a different month. For example, a company's fiscal year might run from June 1st through May 31st. In such cases, it's appropriate to display some date values in a view using fiscal equivalents (Fiscal Year, Fiscal Quarter, and Fiscal Week Number) rather than their calendar equivalents.

To express date fields in fiscal terms, follow these steps:

1. Right-click the data source in the Data window to open the Date Properties dialog box.

   This change affects the fiscal year for all date dimensions in the data source.
2. Set the **Fiscal year start** field to the appropriate month.
3. Right-click the date dimension in the Data window and select **Default Properties > Calendar Type > Fiscal Calendar**. This option is only available on fields that are classified as date dimensions.

Whether a given level of a date dimension is affected by the use of the Fiscal Calendar depends on the specific case. Consult the following table:

Date Level	When Converted to Fiscal
Year	Reflects the fiscal year. For example, if **Fiscal year start** is set to April, the year for the date June 1, 2004 would be shown as FY 2005.
Quarter	The Quarter reflects the fiscal quarter. For example, if **Fiscal year start** is set to April, the quarter for the date June 1, 2004 would be 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.
Hour	No change in behavior. The calendar hour is the same as the fiscal hour.
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.
Week Number	The Week Number reflects the fiscal week number. For example, if **Fiscal year start** is set to April, the week number for the date April 1, 2004 would be 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

Date Level	When Converted to Fiscal
	has been assigned.

The only date levels that explicitly indicate that the fiscal calendar is in use are the Year and Quarter level. Specifically, fiscal years and quarters are shown with the FY prefix. This is not true for 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.

Fiscal year formatting is applied to all date formats that include a year, or a year and a quarter. In particular, if you apply a custom date format, and only use the “y” and “q” placeholders, then `FY` will be prepended to each year.

Date Truncations

When a date dimension is using a fiscal calendar, only the following date parts and truncations will reflect the fiscal calendar:

Date part or truncation	When using a fiscal calendar
Year date part	Reflects the fiscal year. For example, if **Fiscal year start** is set to April, the year for the date June 1, 2004 would be shown as FY 2005.
Quarter date part	Reflects the fiscal quarter. For example, if **Fiscal year start** is set to April, the quarter for the date June 1, 2004 would be Q1.
Week Number date part	Reflects the fiscal week number. For example, if **Fiscal year start** is set to April, the week number for the date April 1, 2004 would be 1.
Year date truncation	Truncates to the start of the fiscal year. For example, if **Fiscal year start** is set to May, the date June 1, 2004 would become May 1, 2004.
Quarter date truncation	Truncates to the start of the fiscal quarter. For example, if **Fiscal year start** is set to July, the date June 1, 2004 would become April 1, 2004.

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.

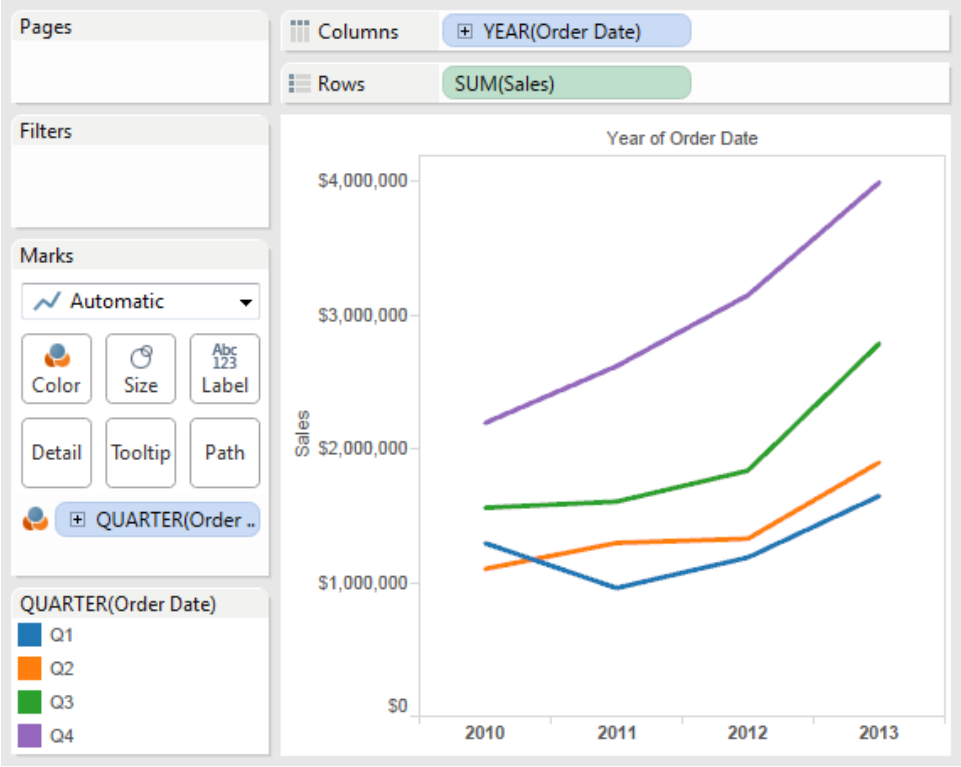

You can separate the marks by month and by quarter as shown below.

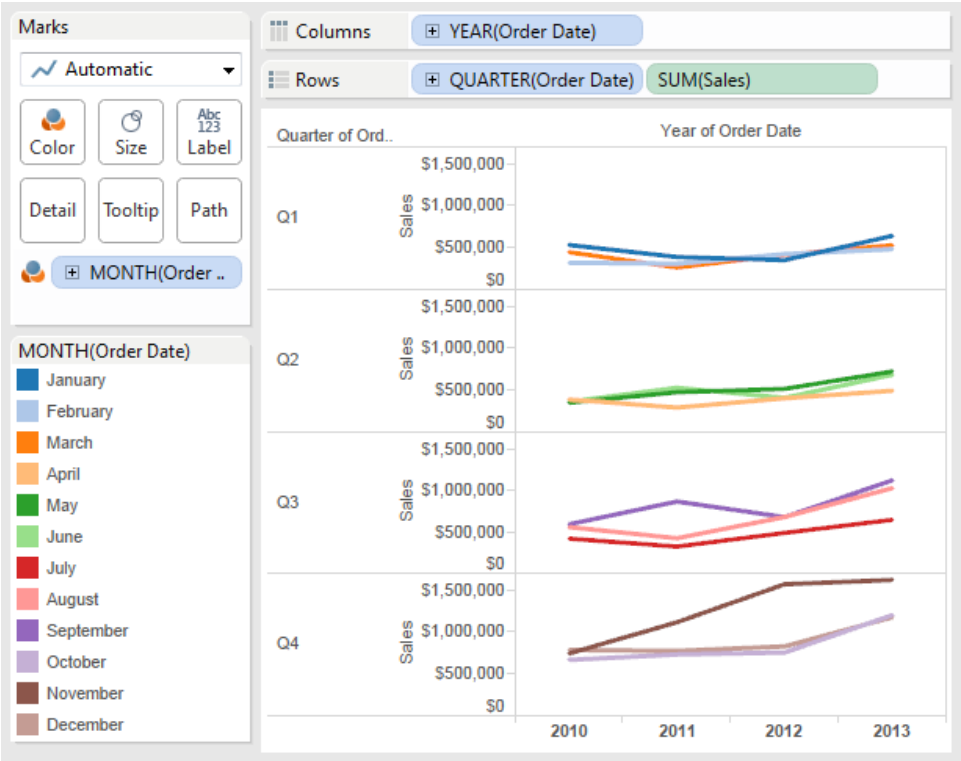

Continuous Dates

You can treat a date as a continuous quantity after placing the field on a shelf. You do this by selecting one of the Continuous date options on the field's context menu (lower list of date levels). Continuous dates draw a quantitative axis for the date values.

For example, the view below displays the sales as a function of a continuous order date and is color-encoded by department. As you can see, the color of the **Order Date** field changes from blue to green after it is converted to a continuous quantity.

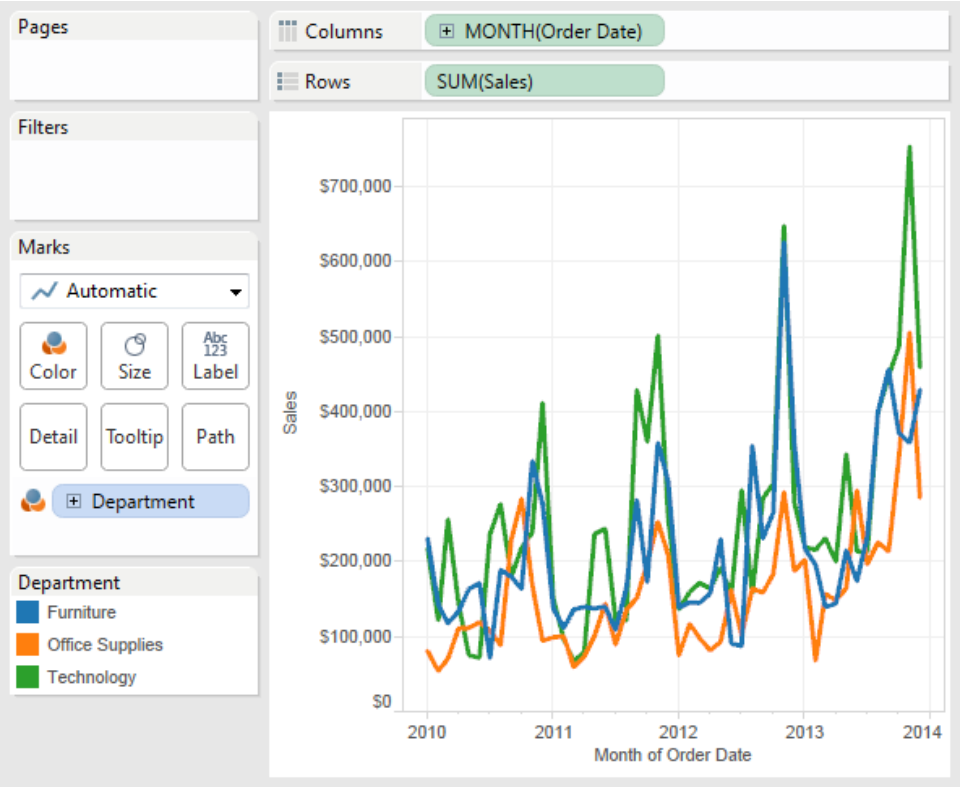

Treating dates as a continuous quantity is particularly useful when you use Gantt bars or want to see trends using line charts as shown above.

By default, date dimensions are discrete fields for which Tableau automatically selects a date level when it is placed on a shelf. To make a date dimension continuous by default, right-click the field in the Data window and select 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.

Missing Values

When you're working with dates or numeric bins, Tableau only shows the values that are represented in your data. If your data does not contain the complete range of values, the missing values will not be shown. For example, your data may contain data for January through May and September through December. However, there was no data recorded for June, July, and August. If you create a line chart in Tableau, the missing months will not be shown. You

can optionally show the missing months to make it clear that there was no data recorded during that time.

Missing Values Hidden (Default)

By default, missing values in a date range or numeric bins are not shown.

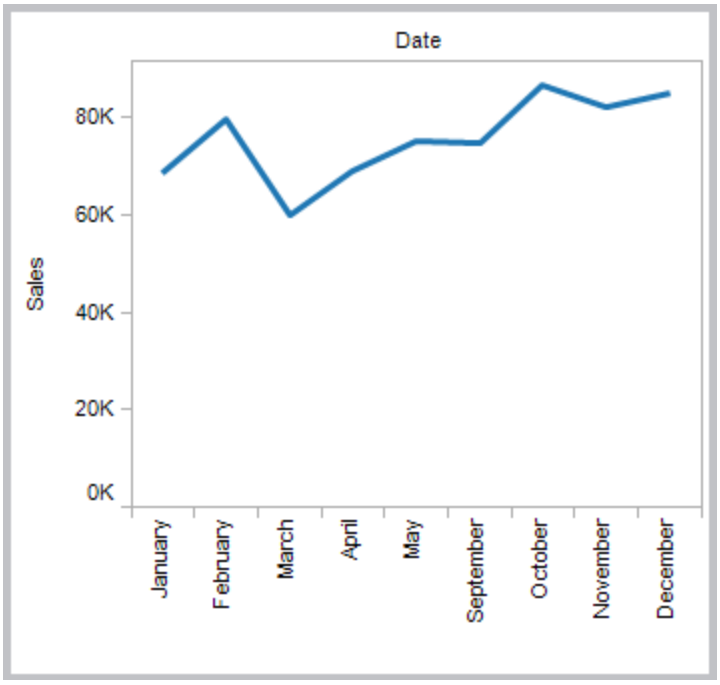

Missing Values Shown

You can show the missing values to indicate incomplete data.

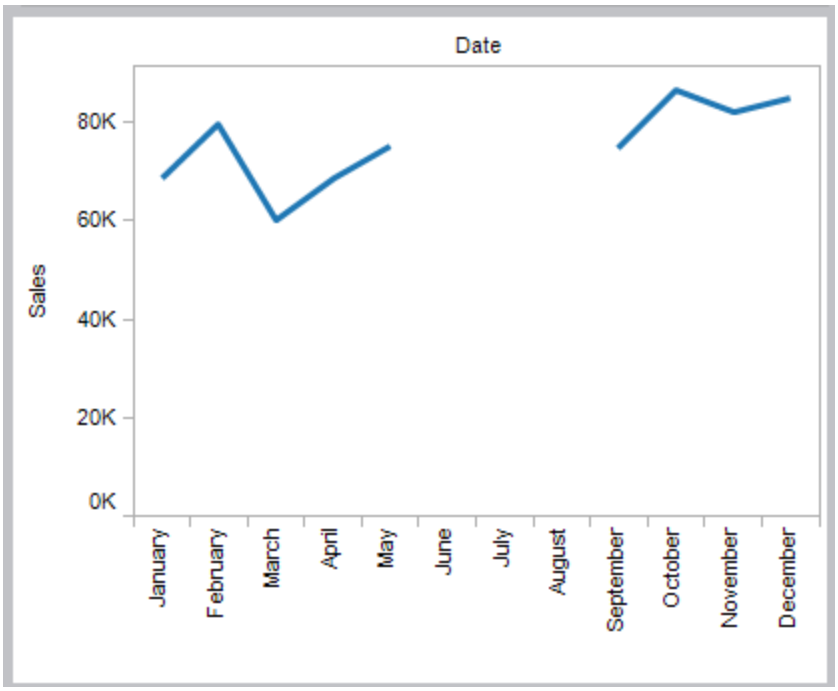

To show missing values in a range, right-click the date or bin headers and select **Show Missing Values**.

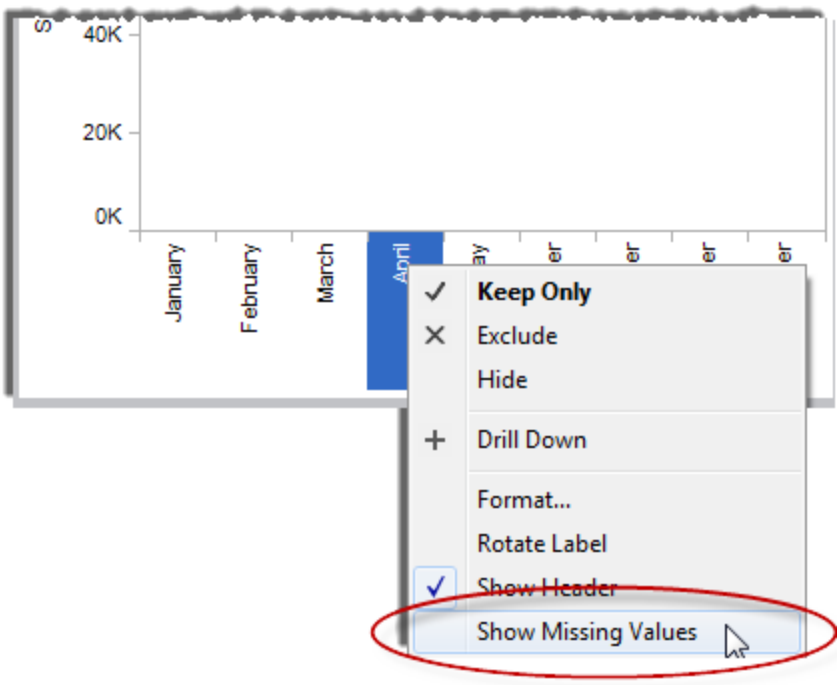

Show and Hide Empty Rows and Columns

When you are working with fields that are not dates or numeric bins, Tableau hides missing values by default. For example, you may be showing workers and the number of hours worked in a particular month. If a worker didn't work at all that month, there may not be a row in the database for that worker. When you drag the Worker field to the Rows shelf, the workers that didn't work are hidden by default. You can show the empty rows by selecting **Analysis** > **Table Layout** > **Show Empty Rows**. Similarly, show the empty columns by selecting **Analysis** > **Table Layout** > **Show Empty Columns**.

Empty Rows Hidden (default)

Bob did not work in January so there are no records in the database for him. By default he is not listed.

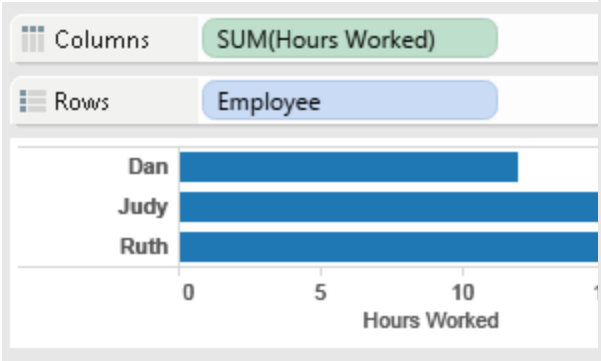

Empty Rows Shown

Even though Bob did not work in January, he is listed but no marks are drawn.

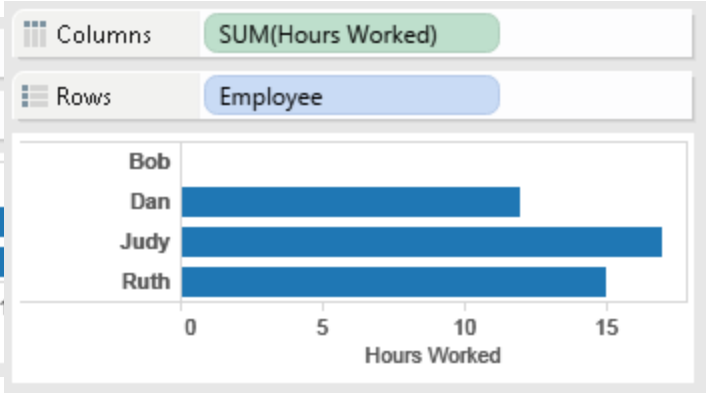

Reference Lines, Bands, and Boxes

Add a reference line, band, or box to mark a specific value or region on an axis. For example, if you are analyzing the monthly sales for several products, you may want to include a reference line at the average sales mark so you can see how each product performed against the average. Alternatively you may want to shade a particular area along the axis, show a distribution, or show the full range of values, using a box plot.

Tableau lets you add an unlimited number of reference lines, bands, and boxes. Right-click on an axis and select **Add Reference Line, Band or Box** to open the Add Reference Line, Band, or Box dialog box:

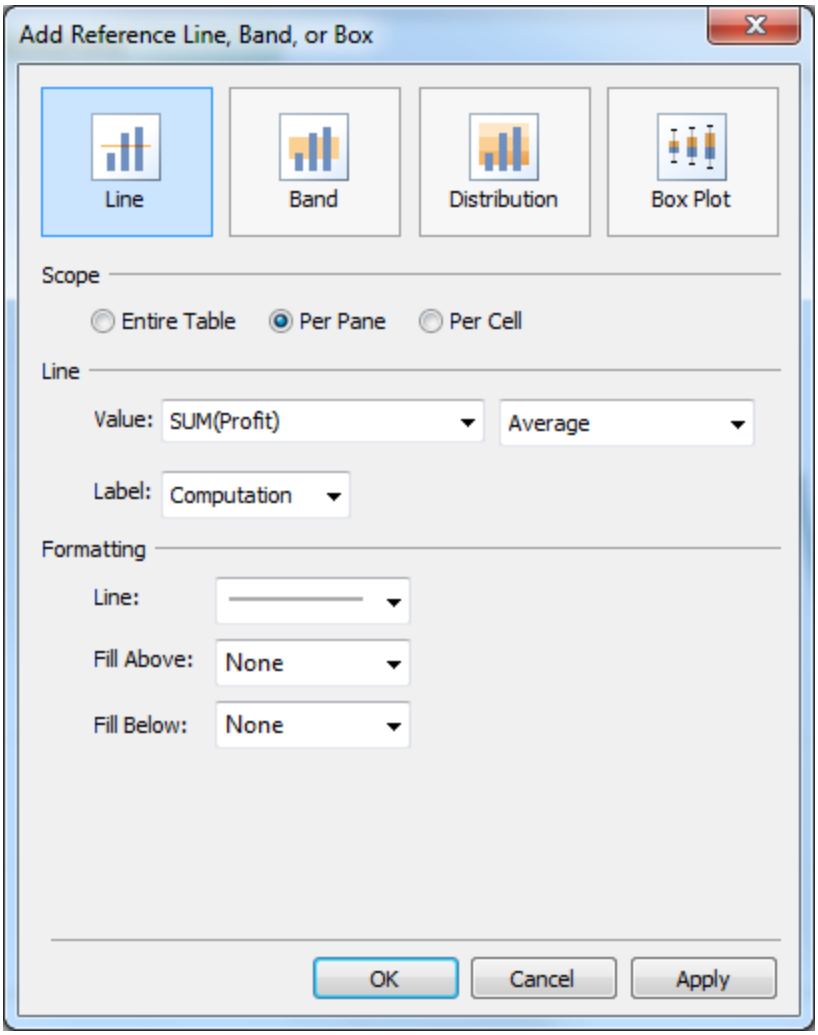

Reference lines, bands, and boxes are not available when the view is a map using online or offline maps.

Types of Reference Lines, Bands, and Boxes

Tableau offers four types of reference lines, bands, and boxes:

- **Line** - adds a line at a constant or computed value on the axis. Computed values can be based on a specified field.

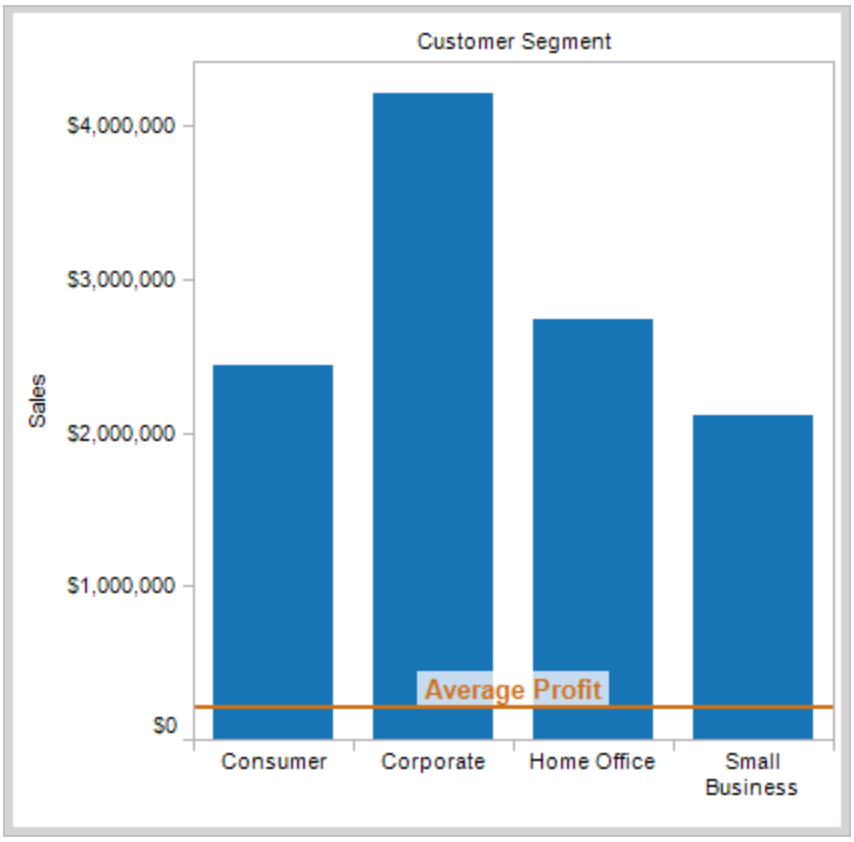

- **Band** - shades an area behind the marks in the view between two constants or computed values on the axis.

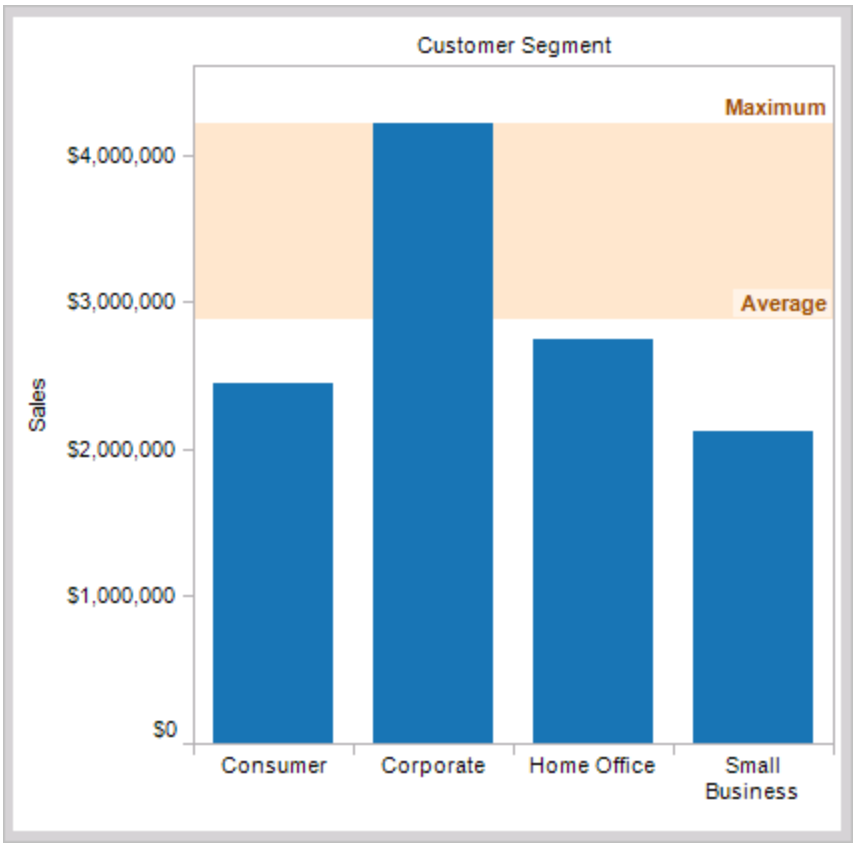

- **Distribution** - adds a gradient of shading to indicate the distribution of values along the axis. Distribution can be defined by confidence intervals, percentages, percentiles, quantiles, or standard deviation. In addition to the shading, you can add a line to mark a constant or computed value along the axis. This type of reference can be used to create bullet charts.

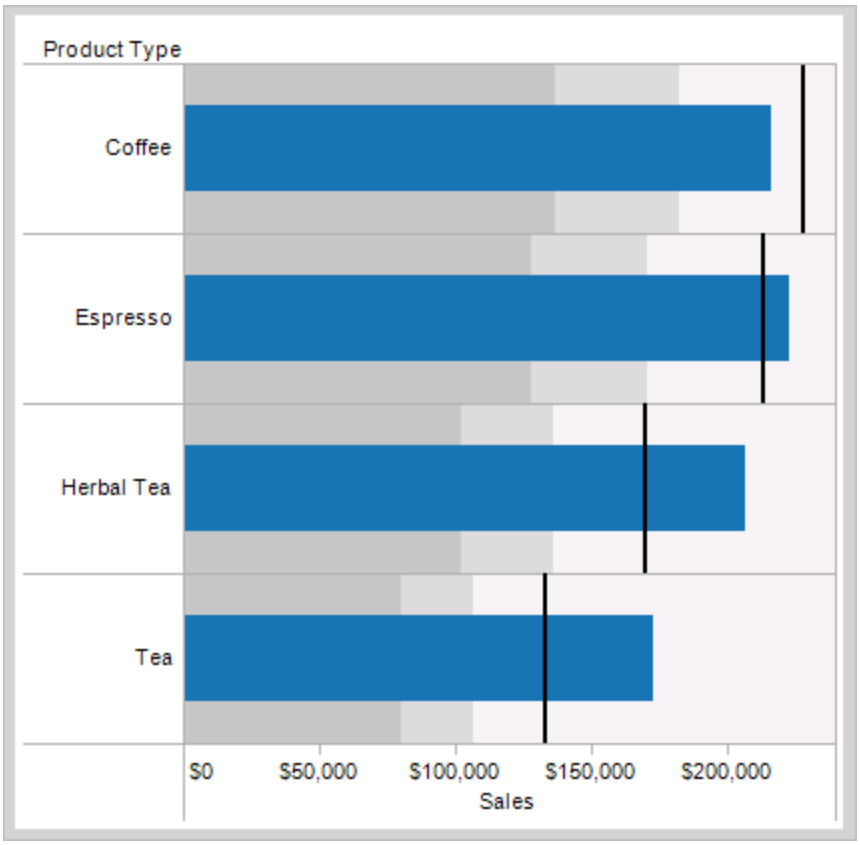

- **Box Plot** - adds a box plot that describes the distribution of values along the axis. Box plots show quartiles and whiskers. Tableau provides different box plot styles, and allows you to configure the location of the whiskers and other details. See **Adding Box Plots** on page 468.

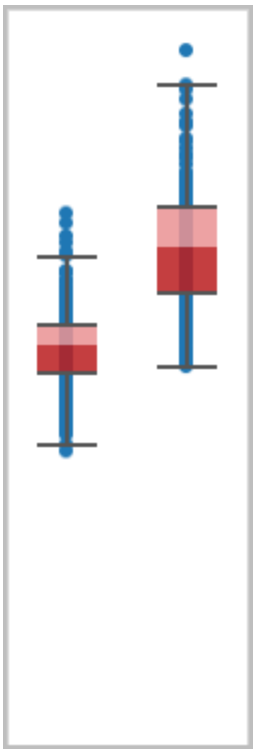

Adding Reference Lines

You can add a reference line to any continuous axis.

To add a reference line:

1. Right-click on a quantitative axis and select **Add Reference Line, Band, or Box**.

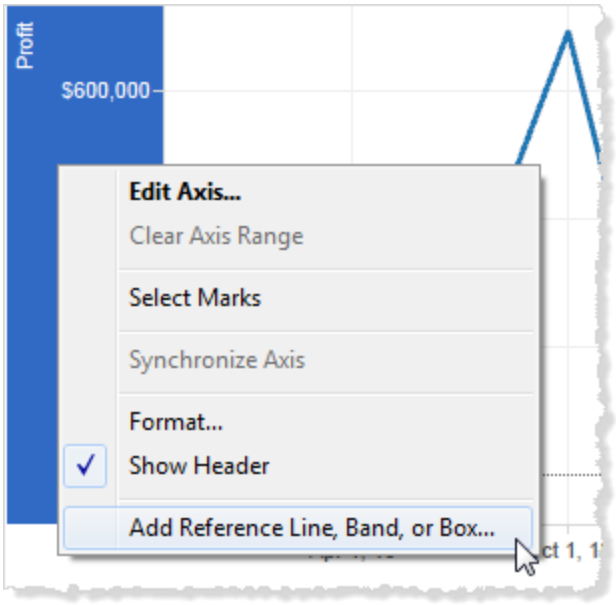

2. In the Add Reference Line, Band, or Box dialog box, select **Line**.

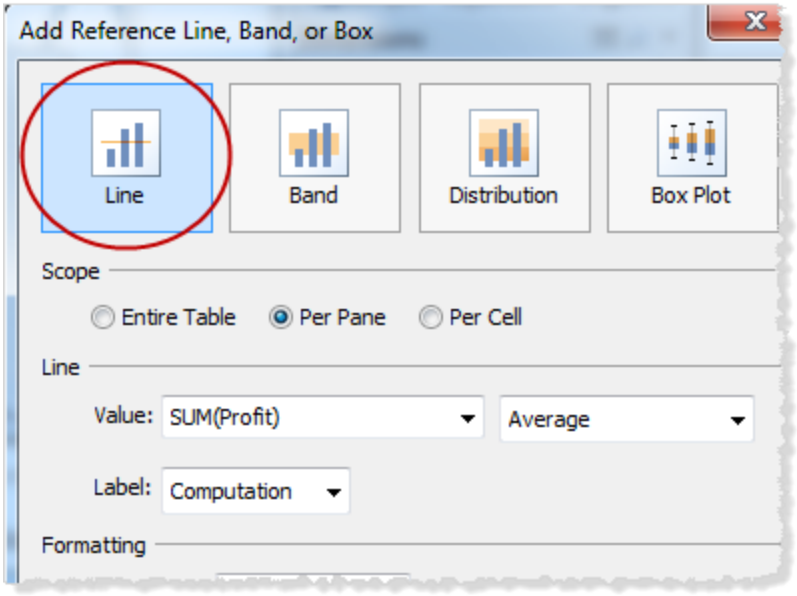

3. In the Add Reference Line, Band, or Box dialog box, select one of the following scopes:

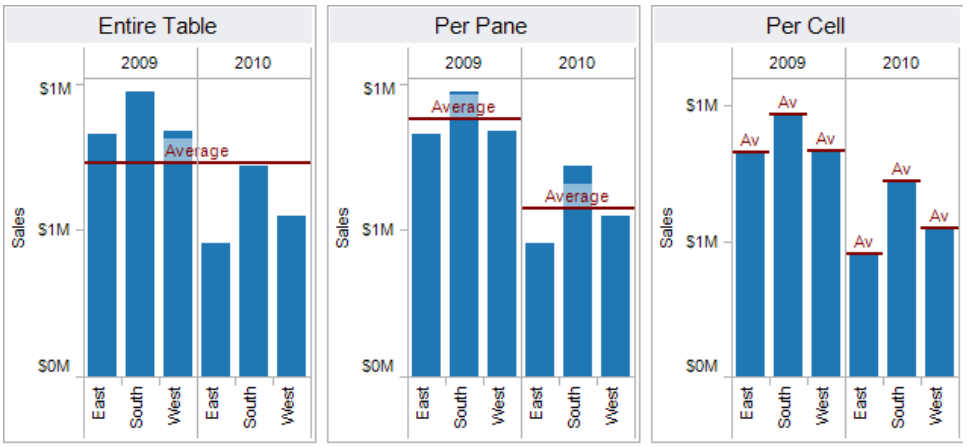

Adds a reference line to the entire table across all panes.

Adds a reference line on a per pane basis. Computed reference lines are recalculated for each pane in the view.

Adds a reference line within each cell. Computed reference lines are recalculated for each cell in the view.

4. Select a measure from the Value field to use as the basis for your reference line. You can also select an existing parameter, or create a new parameter on the fly.

   You cannot select a measure that isn't currently in the view as the basis for your reference line. If this is what you want to use, close the Add Reference Line, Band, or Box dialog box and then drag the measure you want to use from the Data window to the Details target on the Marks card and change the aggregation if necessary. This will not change the view, but it will allow you to use that measure as the basis for your reference line. Now re-open the Add Reference Line, Band, or Box dialog box and resume at step 1, above.

5. Select one of the following aggregations:
   - **Average** - places a line at the average value along the axis.
   - **Constant**- places a line at the specified value on the axis.
   - **Maximum** - places a line at the maximum value.
   - **Median**- places a line at the median value.
   - **Minimum** - places a line at the minimum value.
   - **Sum** - places a line at the SUM of all the values in either the cell, pane, or entire view.
   - **Total** - places a line at the aggregate of all the values in either the cell, pane, or the entire view. This option is particularly useful when computing a weighted average rather than an average of averages. It is also useful when working with a calculation with a custom aggregation. The total is computed using the underlying data and behaves the same as selecting one of the totals option the Analysis

menu.

6. Select how you want to label the line. You can select from the following options:
   - **None** –select this option to not include a label for the reference line.
   - **Value** – select this option to include a label that is the corresponding value on the axis.
   - **Computation** – select this option to display an automatic label. The label is based on the computation and the measure that is selected.
   - **Custom** – select this option to type a custom label into the text box. You can use the menu to the right of the text box to insert values such as the computation or the value.
7. Specify Formatting options for the line. You can change the style, thickness, and color.

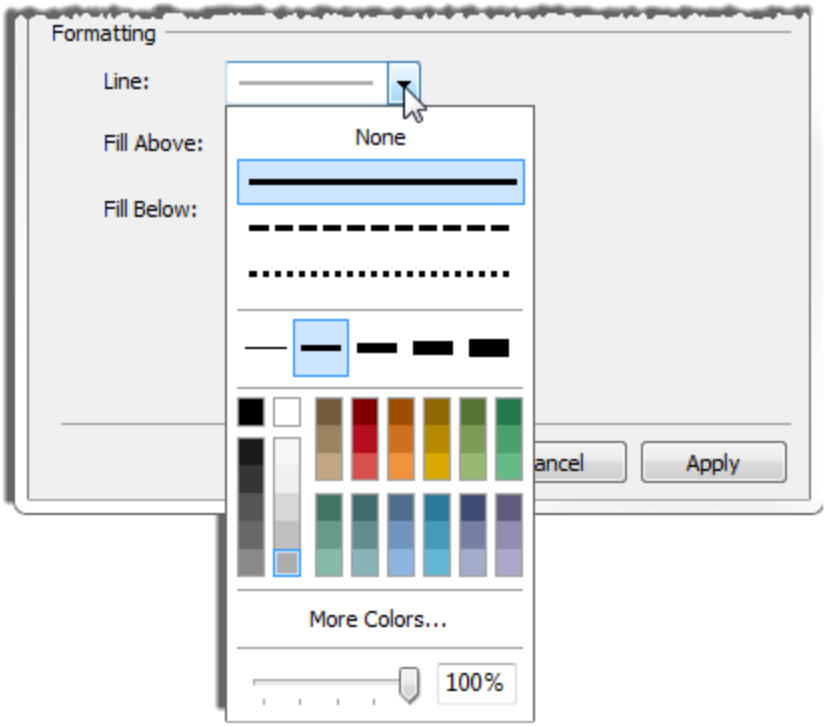

8. Optionally, add a Fill color **Above** and **Below** the line.

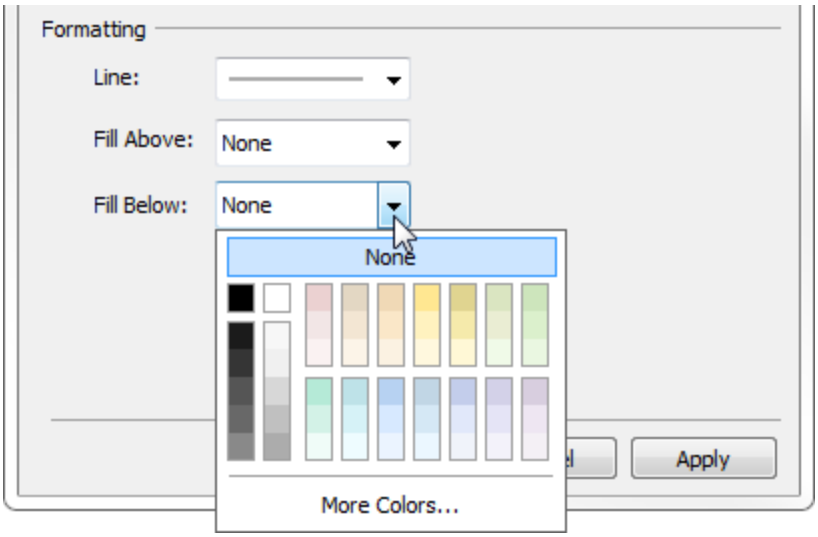

Adding Reference Bands

Reference bands are shaded areas behind the marks in the view between two constant or computed values on the axis. You can add reference bands to any continuous axis.

To add a reference band:

1. Right-click on a quantitative axis and select **Add Reference Line, Band, or Box**.

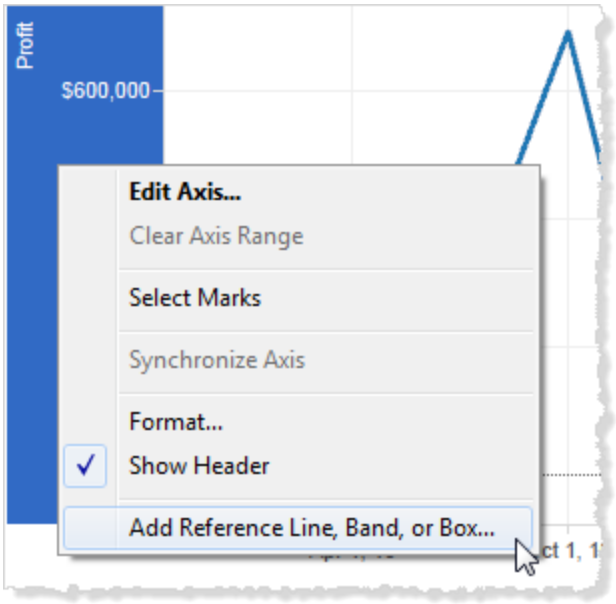

2. In the Add Reference Line, Band, or Box dialog box, select **Band**.

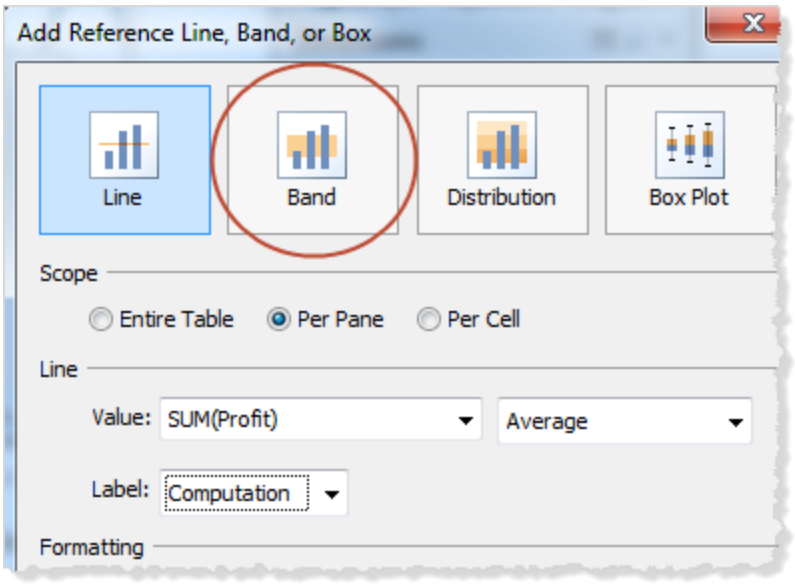

3. Select one of the following scopes:

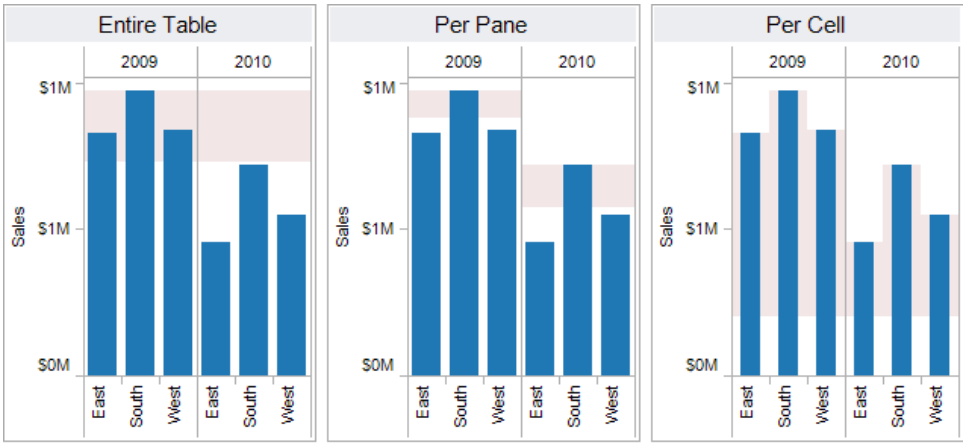

Adds a reference band to the entire table across all panes.

Adds a reference band on a per pane basis. Computed reference bands are recalculated for each pane in the view.

Adds a reference band within each cell. Computed reference bands are recalculated for each cell in the view.

4. Select two measures from the Value field—one in the **Band From** area and one in the **Band To** area—to use as the basis for your reference band. You can also select an existing parameter, or create a new parameter on the fly.

> You cannot select a measure that isn't currently in the view as the basis for your reference band. If this is what you want to use, close the Add Reference Line, Band, or Box dialog box and then drag the measure you want to use from the Data window to the Details target on the Marks card and change the aggregation

if necessary. This will not change the view, but it will allow you to use that measure as the basis for your reference band. Now re-open the Add Reference Line, Band, or Box dialog box and resume at step 1, above.

Select measures or parameters from the first drop-down menus and then select one of the following aggregations:

- **Average** - places a line at the average value along the axis.
- **Constant**- places a line at the specified value on the axis.
- **Maximum** - places a line at the maximum value.
- **Median**- places a line at the median value.
- **Minimum** - places a line at the minimum value.
- **Sum** - places a line at the SUM of all the values in either the cell, pane, or entire view.
- **Total** - places a line at the aggregate of all the values in either the cell, pane, or the entire view. This option is particularly useful when computing weighted average rater than an average of averages. It is also useful when working with a calculation with a custom aggregation. The total is computed using the underlying data and behaves the same as selecting one of the totals options in the Analysis menu.

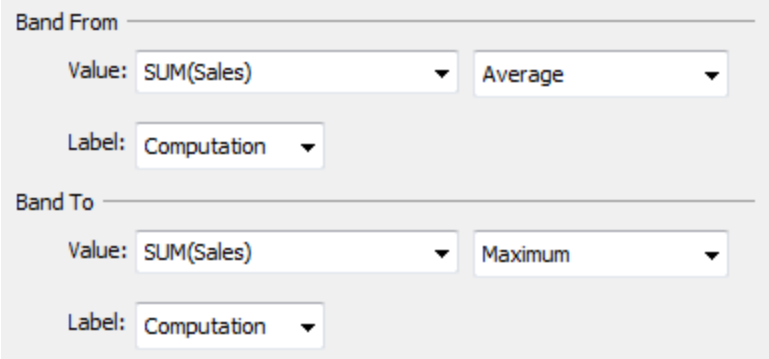

5. Format the reference band. You can mark the two values with a line and select the color to shade between them with.

6. When finished, click **OK**.

Adding Reference Distributions (Bullet Graphs)

Reference distributions are a variation of reference bands. A reference distribution adds a gradient of shading to indicate the distribution of values along the axis. Distributions can be defined by confidence interval, percentages, percentiles, quantiles, or standard deviation. In addition to the shading, you can add a line to mark a constant or computed value along the axis.

Basic Reference Distributions

To add a reference distribution:

1. Right-click on a quantitative axis and select **Add Reference Line, Band, or Box**.

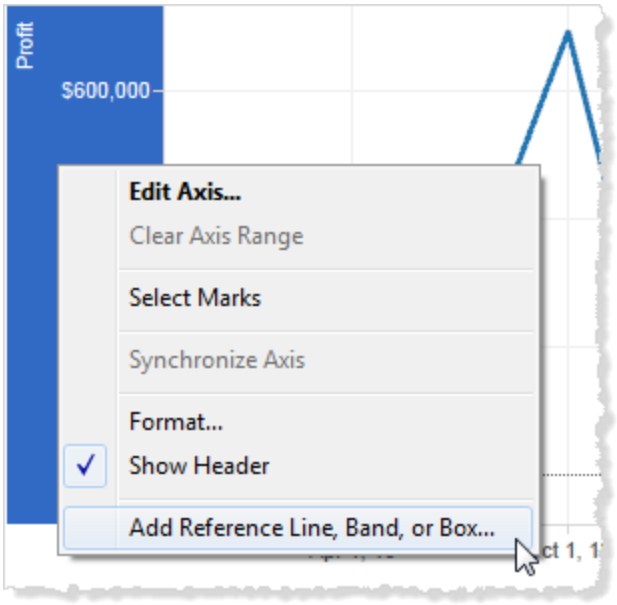

2. In the Add Reference Line, Band, or Box dialog box, select **Distribution**.

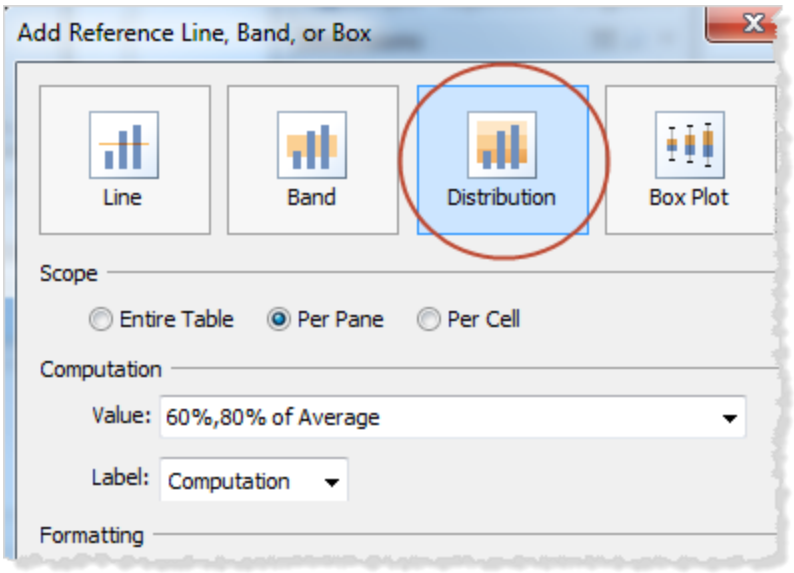

3. Select one of the following scopes:

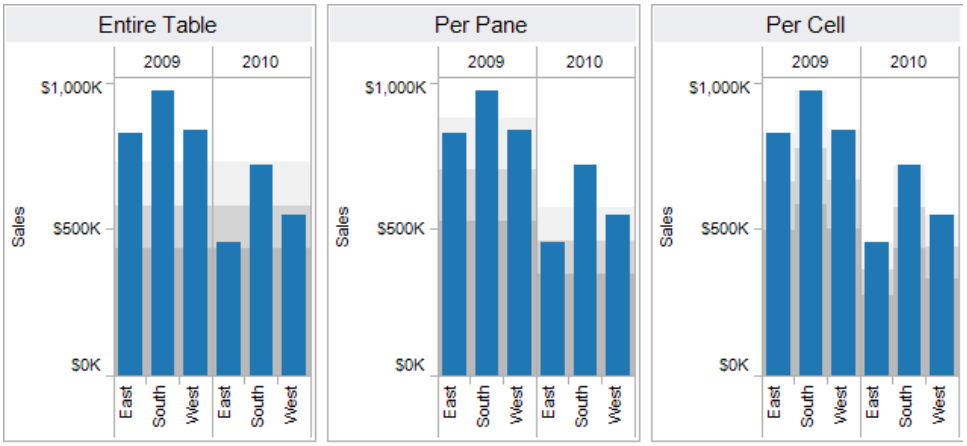

Adds a reference distribution to the entire table across all panes.

Adds a reference distribution on a per pane basis. Computed distributions are recalculated for each pane in the view.

Adds a reference distributions within each cell. Computed distributions are recalculated for each cell in the view.

4. Select the distribution values. You can select from the following options:
   - **Confidence Interval** - shades the interval between which lie the specified percentage of values.
   - **Percentages** - shades the interval between which lie specified percentages of values. Separate multiple percentage values with a comma (for example, `60, 80`).

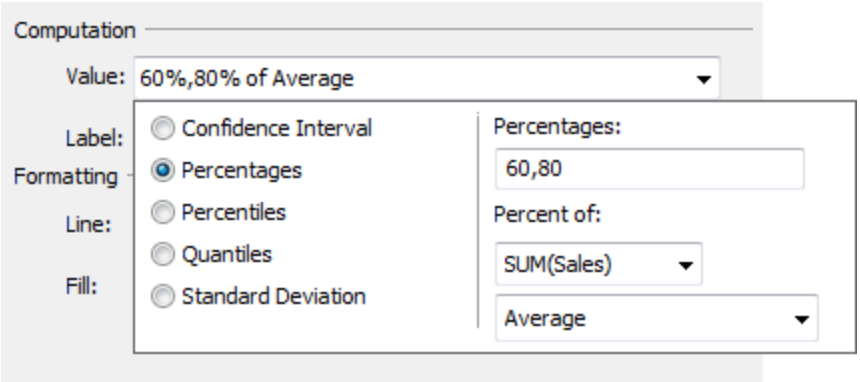

- **Percentiles** - shades intervals at the specified percentiles. When you select this option, you must also specify one or more numerical values. You will not actually see any shading unless you specify at least two values (for example, `60,80`).
- **Quantiles** - breaks the view into the specified number of tiles using shading and lines. When you select this computation, you must also specify the number of tiles (from 3 to 10, inclusive). For example, if you select 3, Tableau calculates the boundaries between the first, second and third terciles by calling the general quantile function and asking for the 33.33 and the 66.66 quantiles. It then shades the three terciles differently.

  Tableau uses estimation type 7 in the R standard to compute quantiles and percentiles. For further details, see Quantile (external link to wikipedia page, in English).
- **Standard Deviation** - places lines and shading to indicated the specified number of standard deviations above and below the mean. When you select this option you must specify the factor, which is the number of standard deviations and whether the computation is on a sample or the population.

5. Specify formatting options. You can format the lines (for example, style, thickness, and color) as well as the fill gradient. Select from a list of predefined gradients. Select Reverse to change the order of shading in the gradient and Symmetric to use a single color instead of a gradient. You can also specify whether to add additional shading above and below the defined distribution.

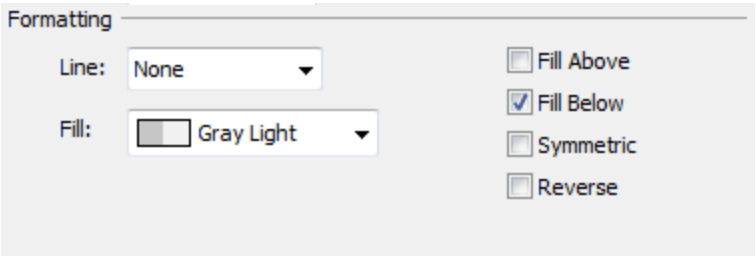

Bullet Graphs

Reference distributions can also be used to create bullet graphs. A bullet graph is a variation of a bar graph developed to replace dashboard gauges and meters. The bullet graph is generally used to compare a primary measure to one or more other measures in the context of qualitative ranges of performance such as poor, satisfactory, and good. You can create a bullet graph by adding two reference lines: a distribution to indicate the qualitative ranges of performance, and a line to indicate the target.

1. Select two measures in the Data window. These measures will be compared in the bullet graph. For example, budget vs. actual; actual vs. target; etc.

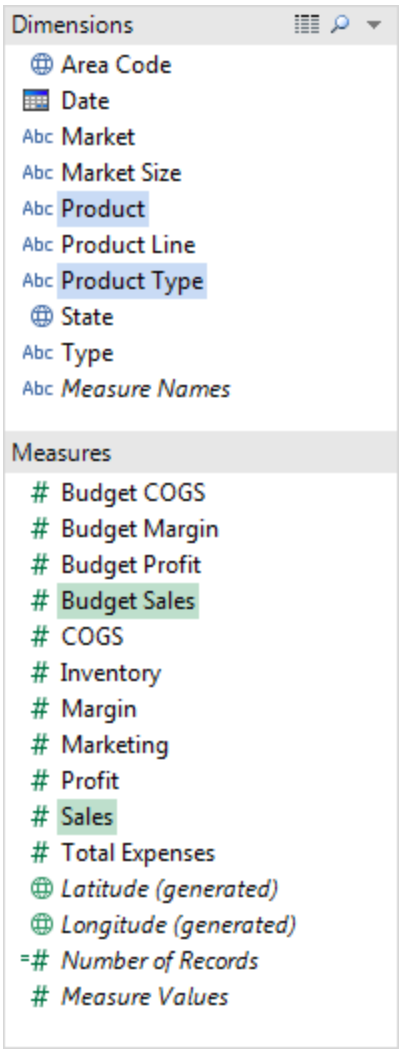

2. Click the **Show Me** button in the toolbar.

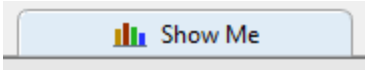

3. Select **Bullet Graph** in the Show Me pane.

Two reference lines are added. By default, Tableau adds a reference distribution that is defined at 60% and 80% of the Average of the measure on Detail. It also adds a reference line that marks the Average of that same measure. The other measure is placed on the Rows shelf.

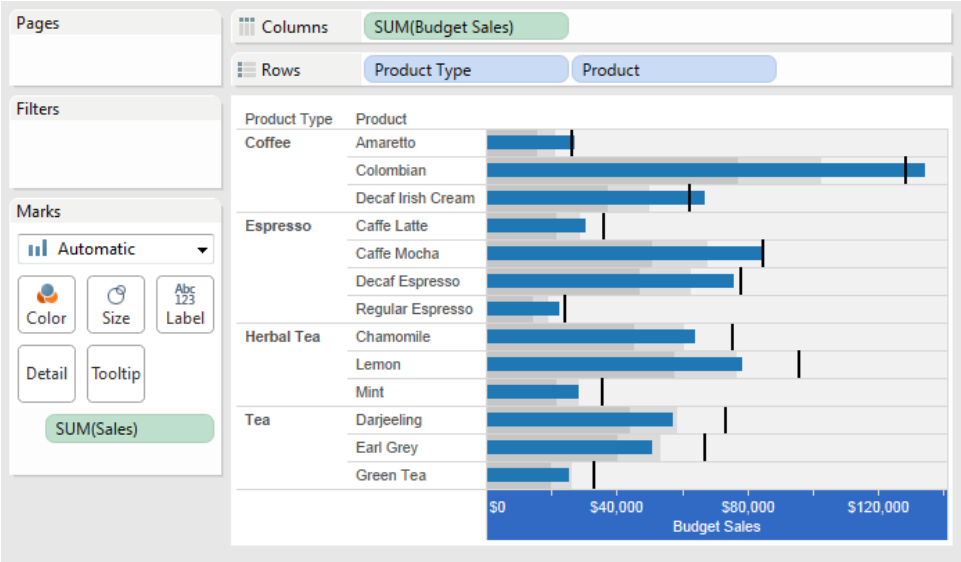

To quickly swap the two measures, right-click on the continuous axis and select Swap Reference Line Fields.

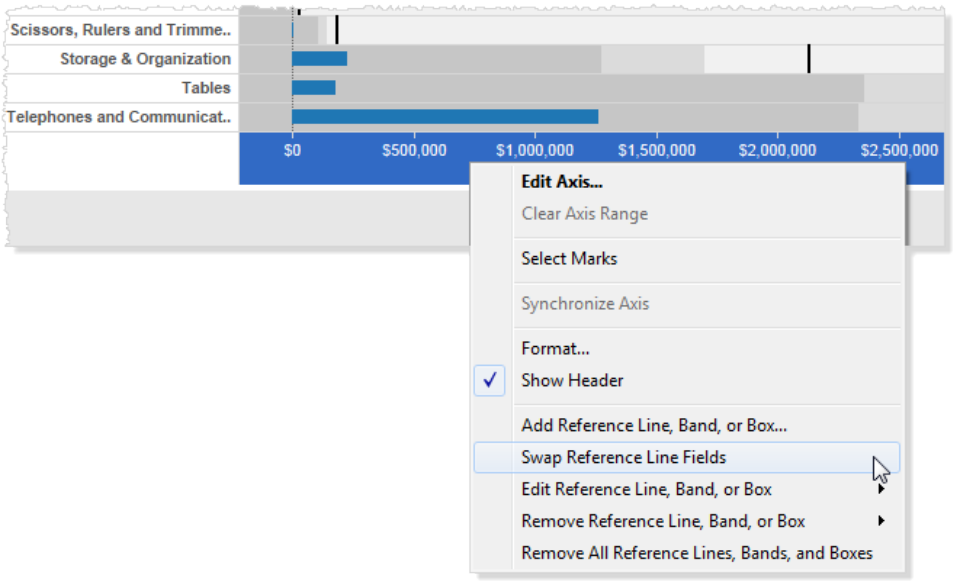

Edit each of the reference lines to change their definitions. For example, you may want to add 100%, or draw a line at a constant value.

Adding Box Plots

Box plots, also known as box-and-whisker plots, are a type of graph that shows the distribution of values along an axis. Boxes enclose the middle 50% of the data (that is, the middle two quartiles of the distribution). Lines, called whiskers, can be configured to display so as to include all points within 1.5 times the interquartile range (in other words, all points within 1.5 times the width of the adjoining box), or at the maximum extent of the data, as in the following image:

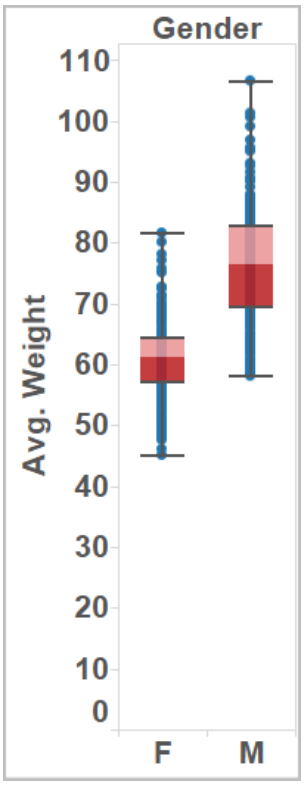

Boxplots are also available from the Show Me pane when you have at least one measure in the view:

For information on Show Me, see **Show Me** on page 322

To add a box plot:

1. Right-click on a quantitative axis and select **Add Reference Line, Band, or Box**.

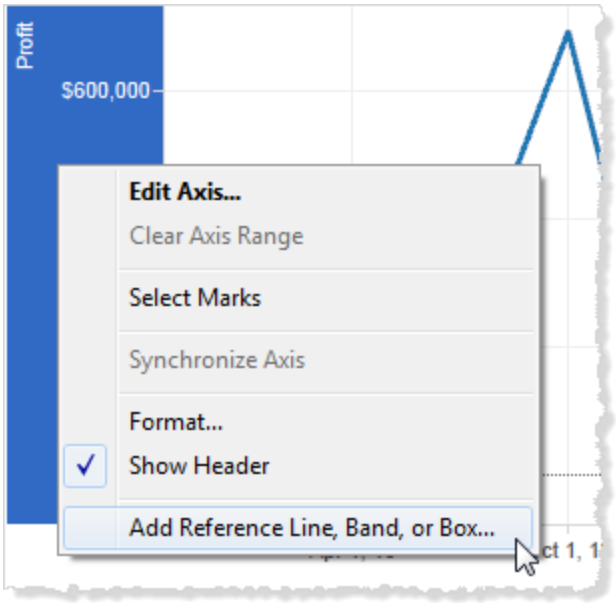

2. In the Add Reference Line, Band, or Box dialog box, select Box Plot.

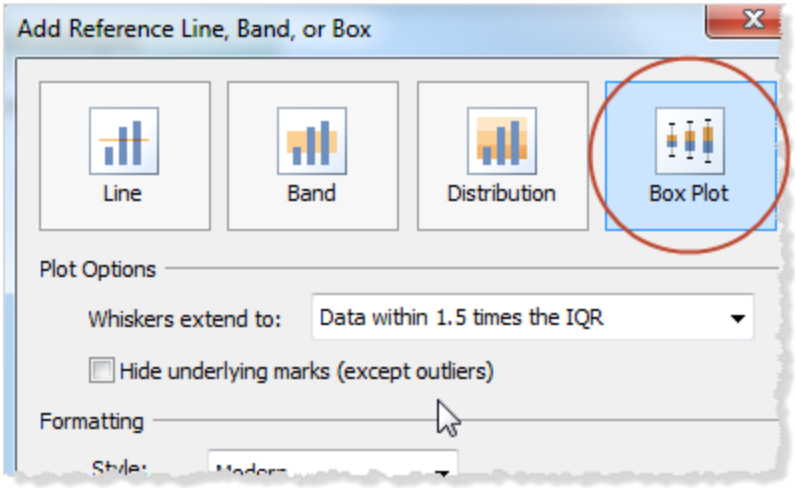

3. Under Plot Options, specify placement for the whiskers:
   - **Data within 1.5 times the IQR** - places whiskers at a location that is 1.5 times the interquartile range—that is, 1.5 times further out than the width of the adjoining box. This is also known as a schematic box plot.
   - **Maximum extent of the data** - places whiskers at the farthest data point (mark) in the distribution. This is also known as a skeletal box plot.
4. Specify whether to **Hide underlying marks (except outliers)**—that is, whether to hide all marks except those beyond the whiskers.

5. Configure the appearance of the plot by selecting a **Style** , **Fill**, **Border**, and **Whiskers**.

Show Me Vs. Add Reference Line, Band, or Box

The difference between adding a box plot using Show Me and adding a box plot using Add Reference Line, Band, or Box is that with Show Me, the box plot is your visualization, whereas with Add Reference Line, Band, or Box, you are adding a box plot to an existing visualization. For example, you could create the following view by first selecting a circle view in Show Me, and then adding a box plot from Add Reference Line, Band, or Box:

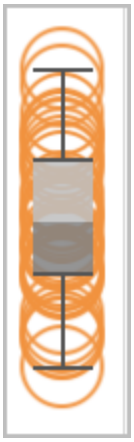

Editing Reference Lines, Bands, or Boxes

After you've added a reference line, band, or box, you can edit the definition. To do this, right-click the continuous axis and select **Edit Reference Line, Band, or Box**. If there are multiple reference lines, bands, or boxes in the view, use the additional menu to select the one you want to edit.

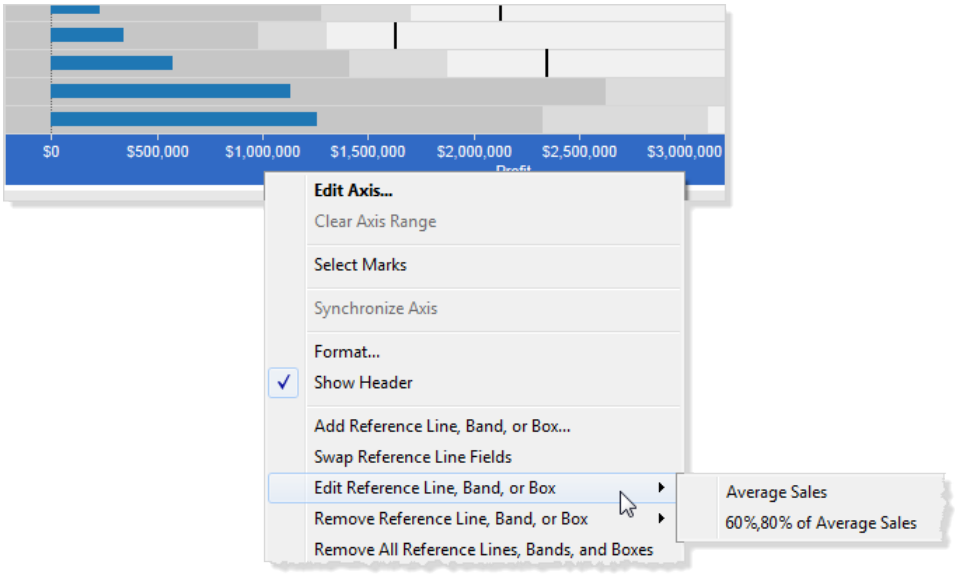

When you have multiple reference lines, bands, or boxes, you may want to change the order they are drawn in the view. To reorder a reference line, right-click the line and select **Move to Front** or **Move to Back**.

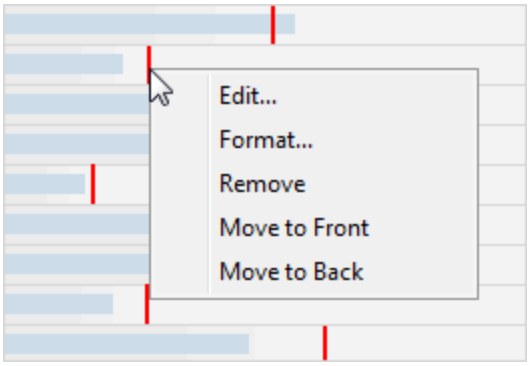

Removing Reference Lines, Bands, or Boxes

You can remove an individual reference line, band, or box, or remove them all at once.

To remove an individual reference line, band, or box, right-click the reference line, band, or box in the view and select **Remove**. If you are removing a reference band or distribution that doesn't include a line, right-click at the beginning or end of the shaded area. In distributions, you can also right-click between the different shades in the gradient.

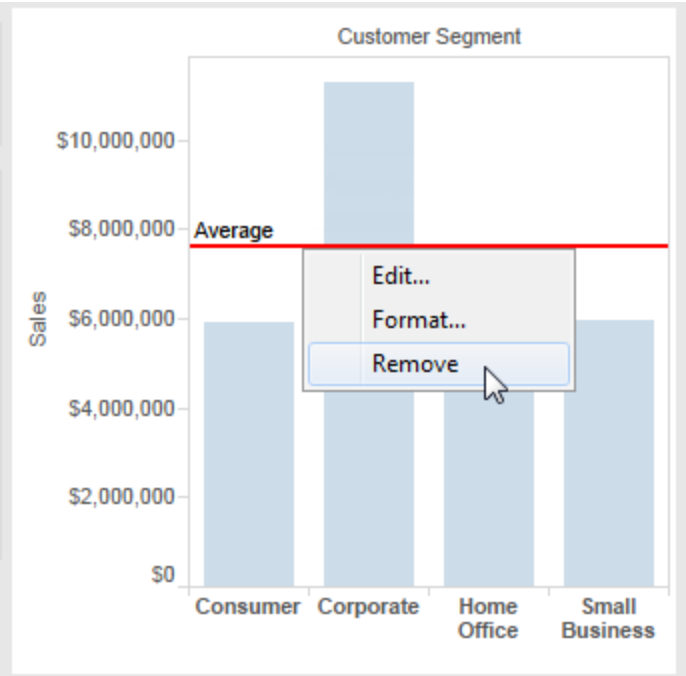

To remove all reference lines, bands, or boxes at once, right-click the continuous axis and select **Remove All Reference Lines, bands, and Boxes**.

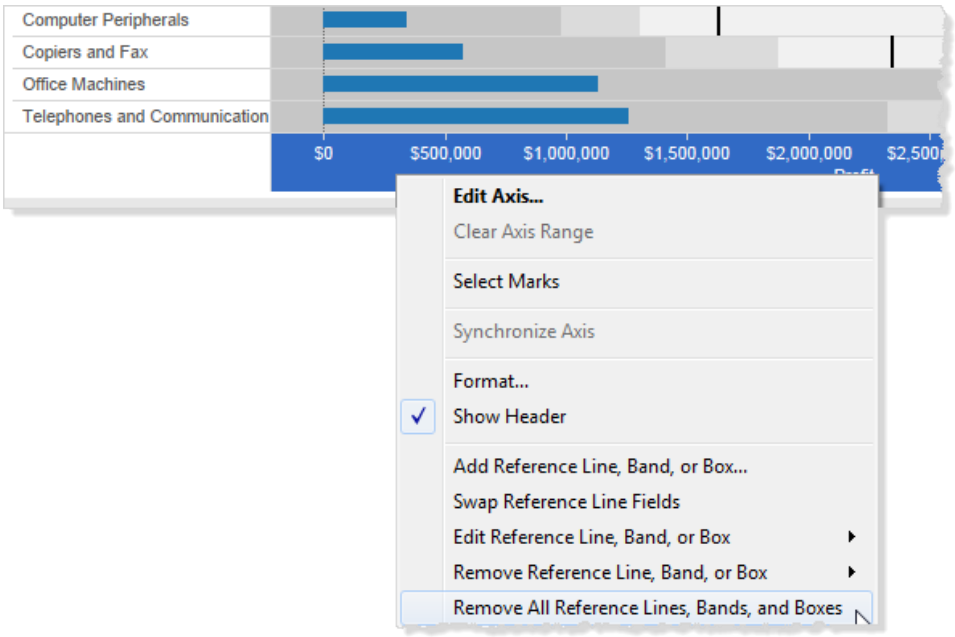

Inspecting Data

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 including the average, minimum, and maximum values.

Select

Selecting marks is useful when you want to visually identify a subset of the data view or you want to run an action.

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.

Zoom Controls

Tableau has a set of zoom controls that display in the upper left corner of the view. By default, these controls only display when you hover over a map view. You can control when the zoom controls display by selecting **Worksheet > Show Zoom Controls** and then select one of the following options:

- **Automatic** – displays when you hover the mouse over map views.
- **Show on hover** – displays when you hover the pointer over any kind of view.
- **Hide** – never displays.

These settings also apply to the view when it is opened in Tableau Reader or Tableau Server. You must specify a setting for each worksheet.

The zoom controls allow you to zoom in and out, zoom to a specific area, and fix or reset the axes. Each control is described below.

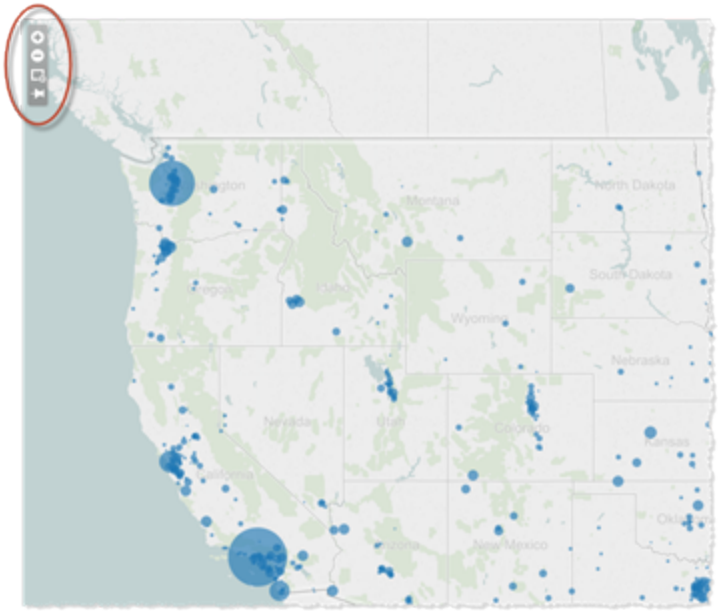

Zoom In and Out

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. Click the plus button to zoom in on the view and the minus button to zoom out. If the zoom controls are hidden, double click the view to zoom in and hold down SHIFT and double-click to zoom out.

Area Zoom

Rather then zooming in and out on the entire view, you can select a specific area to zoom to. When you zoom in on an area, the view is enlarged so that the selected area fills the window. Select the Area Zoom button and then click and drag in the view to select the area to zoom. If the zoom controls are hidden, hold down CTRL + SHIFT and then drag the pointer to select the area you want to zoom to.

Reset Axes

When you zoom in or out the axes in the view are locked to a specific range. You can quickly reset the view back to the automatic axis range by clicking the Reset Axes button in the zoom controls. This button is also available 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 holding SHIFT and then dragging the cursor across the view.

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.

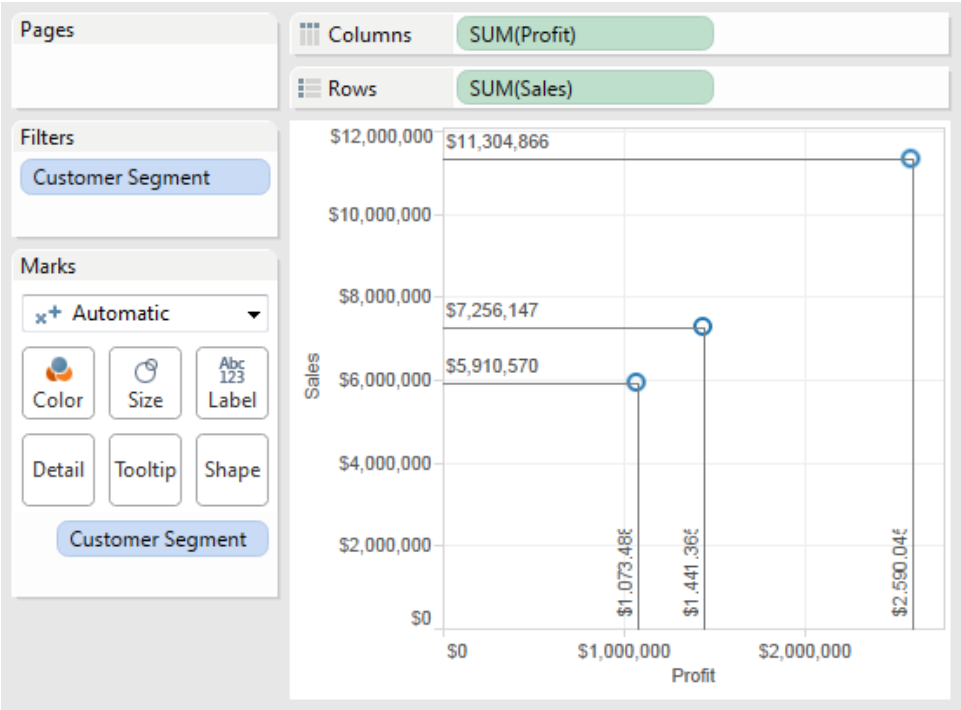

Drop lines do not display when you publish a view to Tableau Server or Tableau Online.

To add drop lines to the view, right-click on the pane and select **Drop Lines > Show 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.

Right-click on the pane and select **Drop Lines > Edit Drop Lines** to open the Drop Lines dialog box.

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.

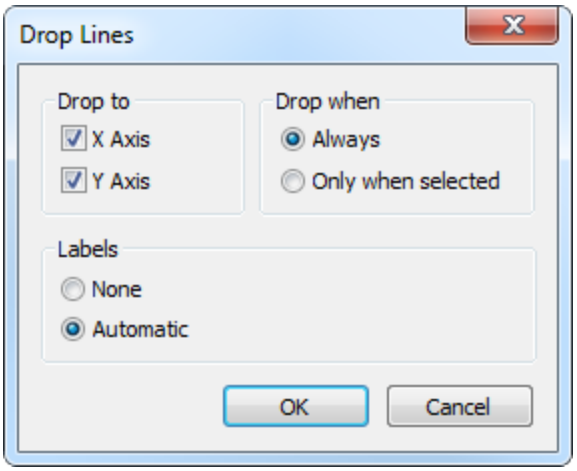

Summary Card

The summary card is a really quick way to view information about a selection or the entire data source. You can hide or show the Summary Card by selecting it on the View Cards toolbar menu . You can also select **Worksheet** > **Show Summary**.

When you select data in the view, the Summary Card updates to show you information only for the data within the selection:

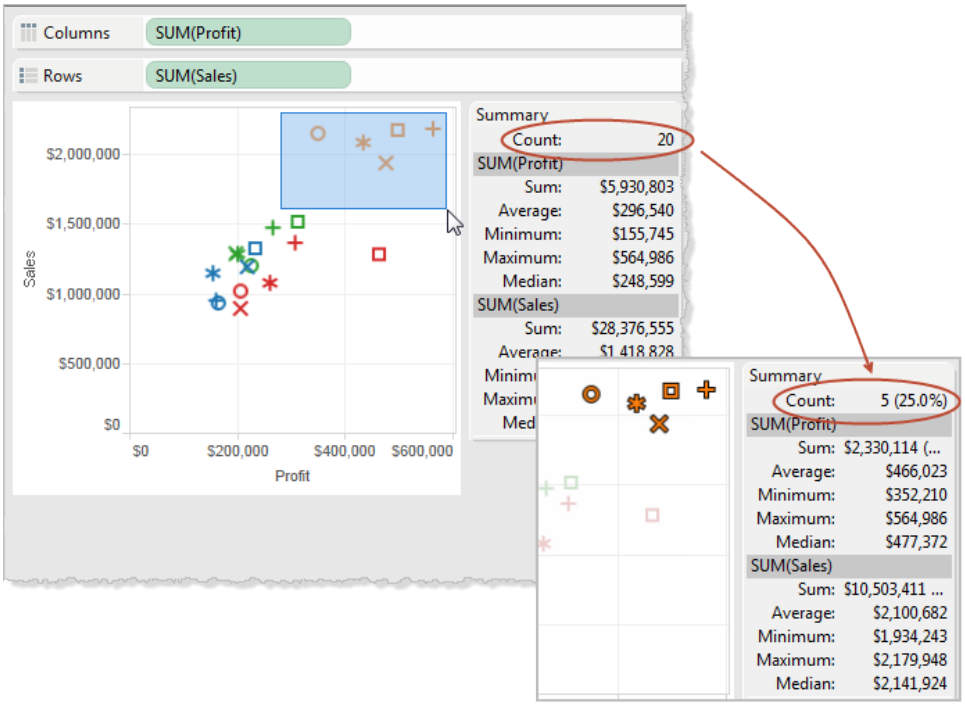

By default, the Summary Card shows Sum, Average, Minimum, Maximum, and Median values for the data in the view. The summary card values update to show these values for the current selection of data points. (Average is computed by summing all relevant values and then

dividing by the total number of values. Median is computed by sorting values from lowest to highest and then selecting the middle value.) The Count value at the top of the card indicates the number of marks in the view or selection.

You can use the drop-down menu for the Summary Card to show additional statistics:

- **Standard Deviation**

  A measure of data spread around its average, measured in the same units as the data itself. The sample standard deviation is an unbiased estimate of the population standard deviation given a slight correction. This standard deviation includes the correction.

- **First Quartile**

  A measure of location which is commonly used with other quartiles to provide a robust measure of spread. Robust in this case means not as sensitive to outliers as the standard deviation. The first quartile is the 25th percentile, typically the lower line in a boxplot

- **Third Quartile**

  A measure of location which is commonly used with other quartiles to provide a robust measure of spread. Robust in this case means not as sensitive to outliers as the standard deviation. The third quartile is the 75th percentile, typically the upper line in a boxplot.

- **Skewness**

  A measure of the tendency of your data to have extreme values to one side. Positive skewness means the extreme values are to the right, while negative skewness means the extreme values are to the left.

- **Excess Kurtosis**

  A measure of the tendency of your data to have more extreme or outlying values than a normal distribution. A normal distribution has a kurtosis of 3 so this value is kurtosis minus three.

View Data

The View Data command lets you display the values for all rows in the data source that underlie a set of marks in the view. It also shows you the summary data based on the aggregations in the view. You can 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.

You can view data for a selection of marks, for the fields in the Data window, and when you’re connecting to data.

The View Data command works with all relational and multi-dimensional databases except Oracle Essbase and SAP® Business Information Warehouse databases. While you can view data with the Microsoft Analysis Services and the Teradata OLAP connector multi-dimensional

databases, the database must be drill-through enabled; in addition there are some restrictions to the data you can see.

In the view below, sales for two product dimensions (**Department** and **Container**) are displayed as a bar chart. Suppose you wanted to view data for the largest marks in each pane. To do this, you would select those marks, right-click in the table, and select **View Data** on the context menu. Alternatively, you could select the **Analysis** > **View Data** menu item.

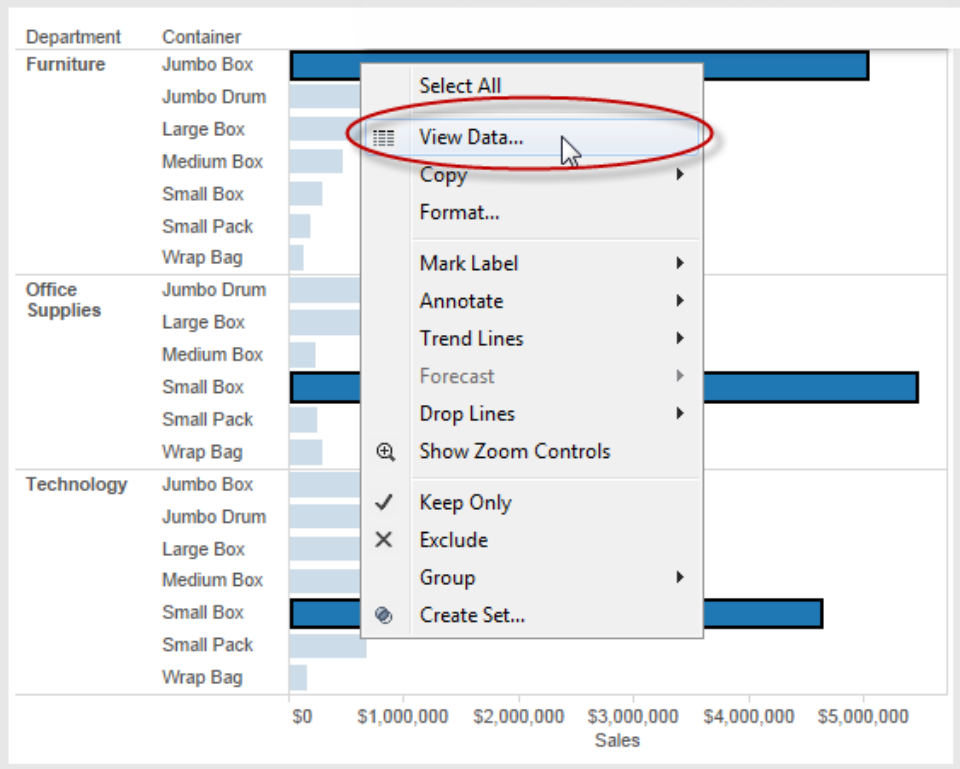

Viewing data may not return any records if you are using a field that contains floating point values as a dimension. This is due to the precision of the data source and mainly occurs when you are connected to Microsoft Excel, Microsoft Access, or text files.

Summary Data

Summarized data is shown on the **Summary** tab. Summarized data is a text table of the aggregated data for the fields shown in the view.

View Data: Sheet 1

Show Aliases | Copy

Category	Container	Sales
Office Supplies	Small Box	$2,695,506
Technology	Small Box	$2,254,595
Furniture	Jumbo Box	$2,448,408

Summary | Underlying | 3 rows

Underlying Data

Underlying data for the selected marks are displayed on the **Underlying** tab. In the lower right of the dialog box you can see the number of rows in the underlying data.

View Data: Sheet 1

Show Aliases | Show all fields | Copy

Last N days	City	Customer	Customer Segm...	Order Date
206	Elk Plain	Liz Pelletier	Corporate	7/15/2011 12:00...
107	Ames	Julie Creighton	Corporate	10/22/2011 12:0...
96	Albany	Sample Company A	Home Office	11/2/2011 12:00...
416	East Meadow	Helen Wasserman	Home Office	12/17/2010 12:0...
1026	Little Rock	Keith Dawkins	Home Office	4/16/2009 12:00...
739	Prescott Valley	Craig Yedwab	Consumer	1/28/2010 12:00...
-286	Moore	Pauline Chand	Corporate	11/18/2012 12:0...
-91	Cleveland	Roy Collins	Corporate	5/7/2012 12:00:...
-91	Cleveland	Roy Collins	Corporate	5/7/2012 12:00:...
606	Lewiston	Emily Phan	Consumer	6/10/2010 12:00...
148	Tamarac	Valerie Takahito	Consumer	9/11/2011 12:00...
964	Lemon Grove	Bart Folk	Corporate	6/17/2009 12:00...
964	Lemon Grove	Bart Folk	Corporate	6/17/2009 12:00...
405	Atlanta	Peter Fuller	Consumer	12/28/2010 12:0...
468	Burleson	Erin Creighton	Corporate	10/26/2010 12:0...
861	Hollywood	Sean Braxton	Small Business	9/28/2009 12:00...
405	Auburn	Liz Willingham	Home Office	12/28/2010 12:0...

Summary | Underlying | 4,586 rows

Sort the data by clicking one or more column headers. To restore the original sort order, click the header repeatedly until it is no longer highlighted with a sort arrow.

By default, **Show all fields** is selected. Clear this option to only show the columns used on shelves (or fields referenced by a calculation used on a shelf) in the current worksheet.

If you want to export one or more data source rows, select the data points of interest by selecting the row and then clicking **Copy** to copy the selected data.

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.

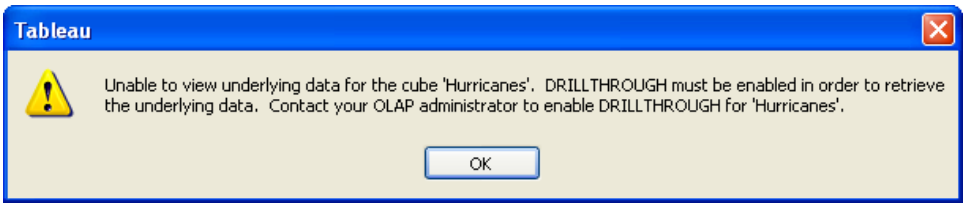

In addition, Microsoft Analysis Services databases limit viewing data to a single mark at a time. More precisely, viewing the data (which uses MDX drill-through) is not an option when the selection of mark(s) is defined by more than one value of a dimension.

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 **Worksheet** > **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 provides other important summary information. This information can be copied and exported to other applications using the Clipboard.

> If you have Trend Lines turned on, the Describe Sheet dialog box includes information about the trend line model. See **Assessing Trend Line Significance** on page 723. If you have Forecasting turned on, the Describe Sheet dialog box includes information about estimated data. See **Forecast Descriptions** on page 656.

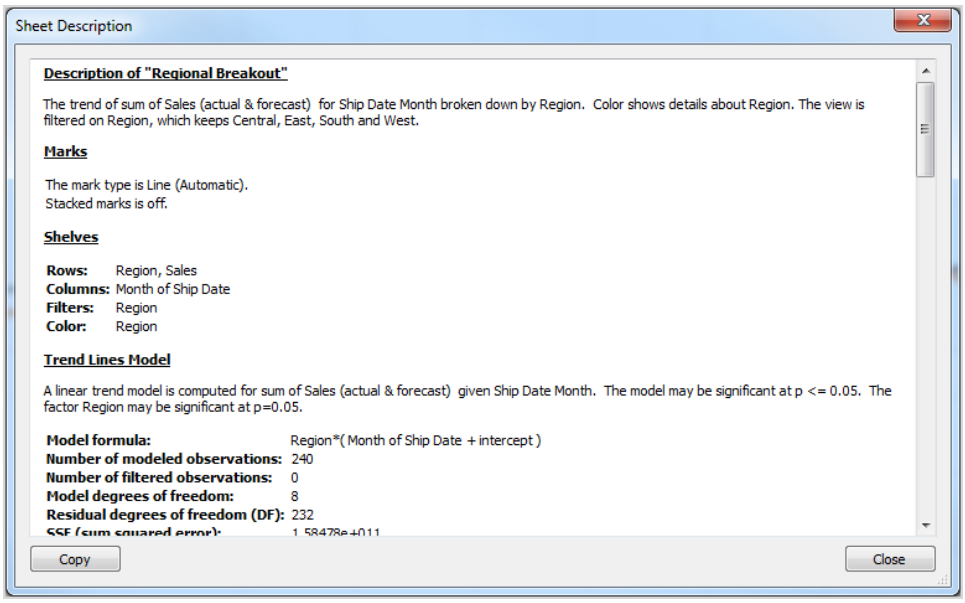

Sheet Description

Description of "Regional Breakout"

The trend of sum of Sales (actual & forecast) for Ship Date Month broken down by Region. Color shows details about Region. The view is filtered on Region, which keeps Central, East, South and West.

Marks

The mark type is Line (Automatic).
Stacked marks is off.

Shelves

Rows: Region, Sales
Columns: Month of Ship Date
Filters: Region
Color: Region

Trend Lines Model

A linear trend model is computed for sum of Sales (actual & forecast) given Ship Date Month. The model may be significant at p <= 0.05. The factor Region may be significant at p=0.05.

Model formula: Region*(Month of Ship Date + intercept)
Number of modeled observations: 240
Number of filtered observations: 0
Model degrees of freedom: 8
Residual degrees of freedom (DF): 232
SSE (sum squared error): 1.58478e+011

Copy Close

Build-It-Yourself Exercises

This section includes detailed exercises that guide you through the steps involved in building some common types of data views. All exercises use the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.You can create many types of data views in Tableau, and this collection of exercises does not attempt to show all types.

Building Bar Charts

Bar charts are a great way to compare data across categories. In Tableau, you typically create a bar chart by placing a dimension as the innermost field on the **Rows** shelf, and a measure as the innermost field on the **Columns** shelf, or vice-versa. The innermost field is the one farthest to the right.

A bar chart uses the bar mark type. Tableau automatically selects this mark type when the data view matches one of the two field arrangements shown below, provided the **Mark** type is set to **Automatic**. You can add additional fields on these shelves, and Tableau will not change the chart type, provided a dimension and a measure remain innermost.

Creates Vertical Bars

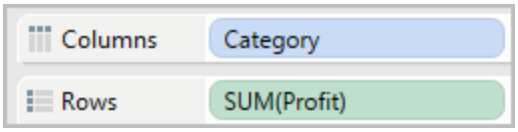

Creates Horizontal Bars

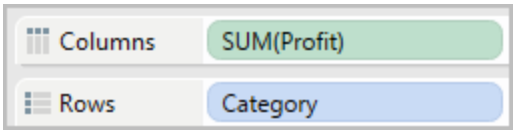

Alternatively, you can manually select **Bar** from the **Mark** menu for any data view. See **Bar Mark** on page 241 for information about this mark type.

To create a bar chart that displays total sales over a four-year period, follow these steps:

1. Connect to the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.
2. Drag the **Order Date** dimension to **Columns**.

   The data is automatically aggregated by year and column headers are displayed.
3. Drag the **Sales** measure to **Rows**.

   The measure is automatically aggregated as a sum and an axis is created. The column headers are relocated to the bottom of the view.

   Tableau automatically selects Line as the mark type because of the date dimension.

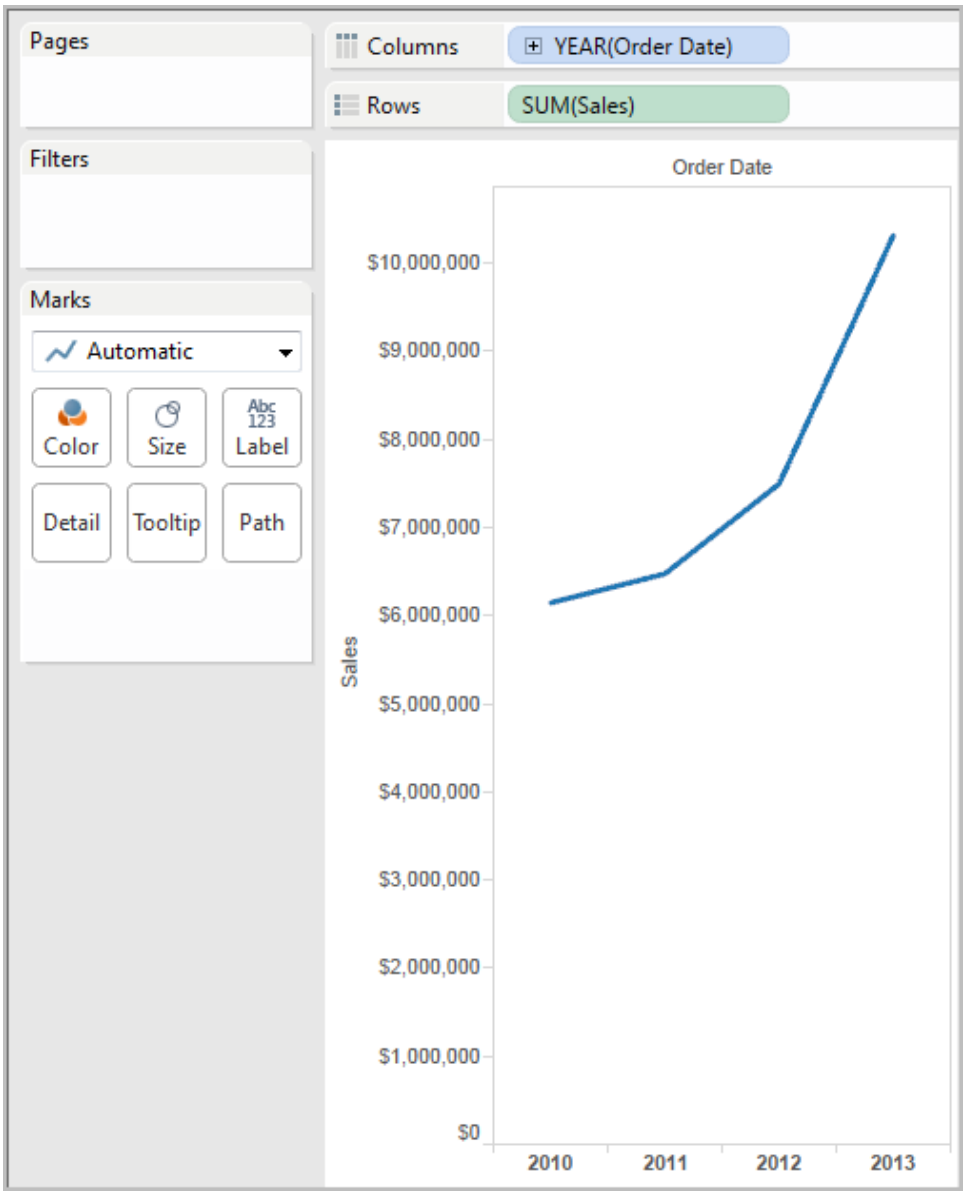

4. On the Marks card, from drop-down list of view types, select **Bar**.

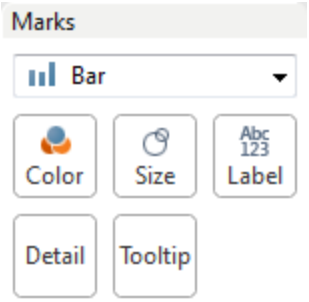

This changes the view to a bar chart (also known as a histogram).

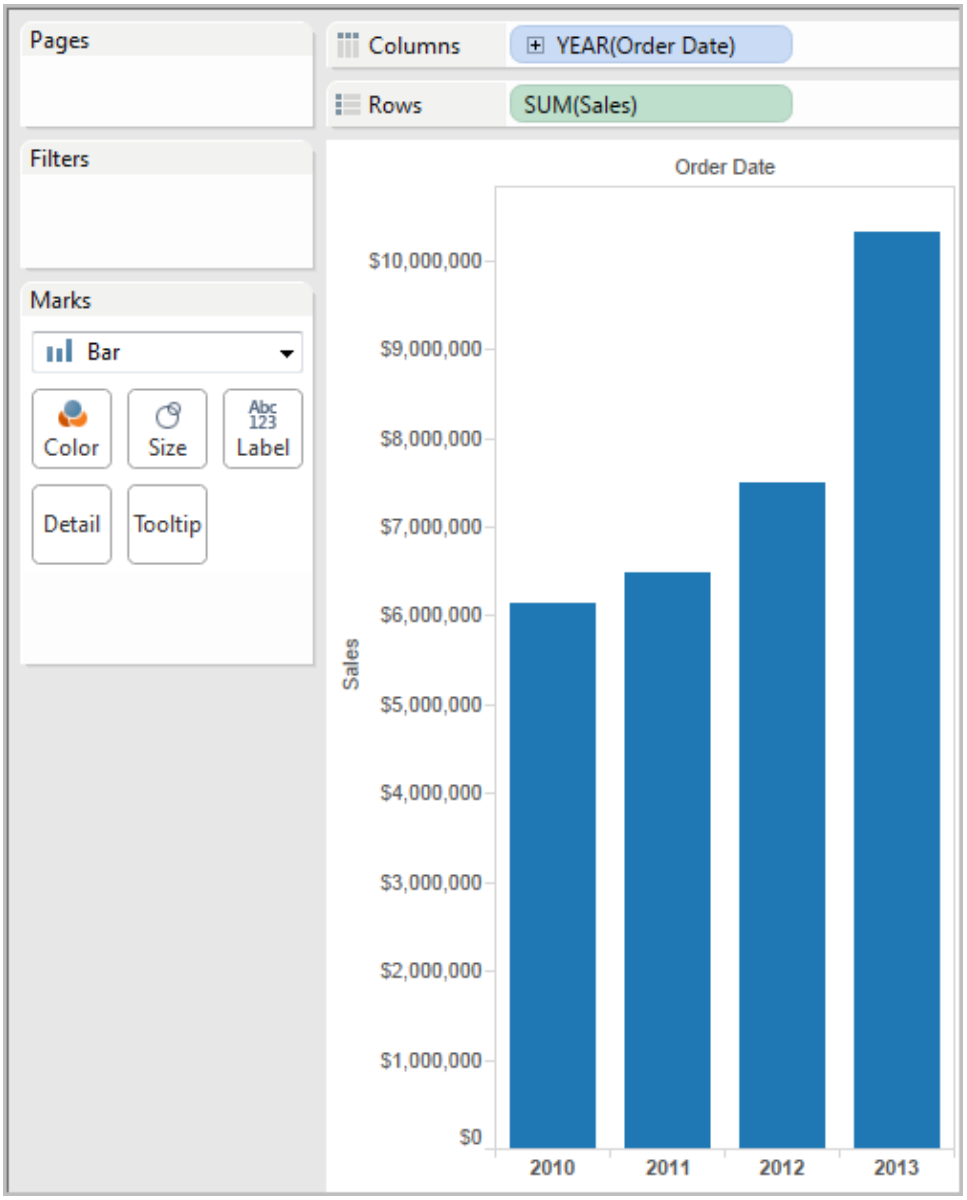

The marks (bars) are vertical because the axis is vertical. The length of each mark represents the sum of the sales for that year. For example, the sum of the sales in 2013 is $10,319,053, which you can verify by hovering the mouse cursor over that column:

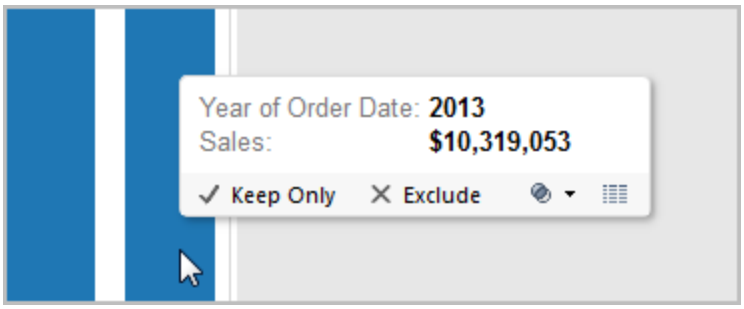

7. Drag the **Ship Mode** dimension to **Color** on the Marks card.

   The view shows how different shipping modes have contributed to total sales over time. The rations seem consistent from year to year:

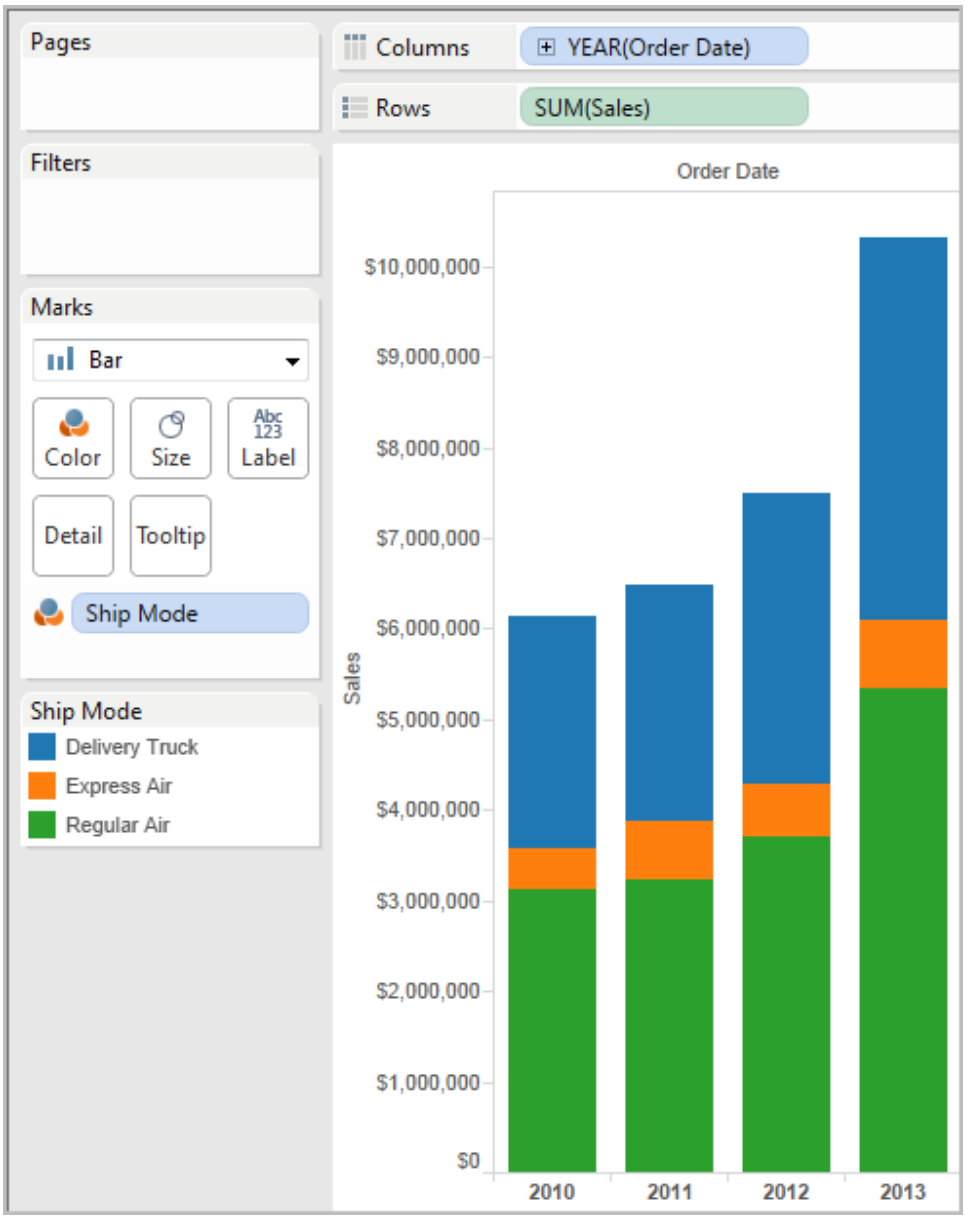

8. Drag the **Region** dimension to **Rows**, and drop it to the left of **Sales** to produce multiple axes for sales by region.
   Notice that sales for the International region exceed sales for other regions so much that sales for other regions appear too compressed to be readable:

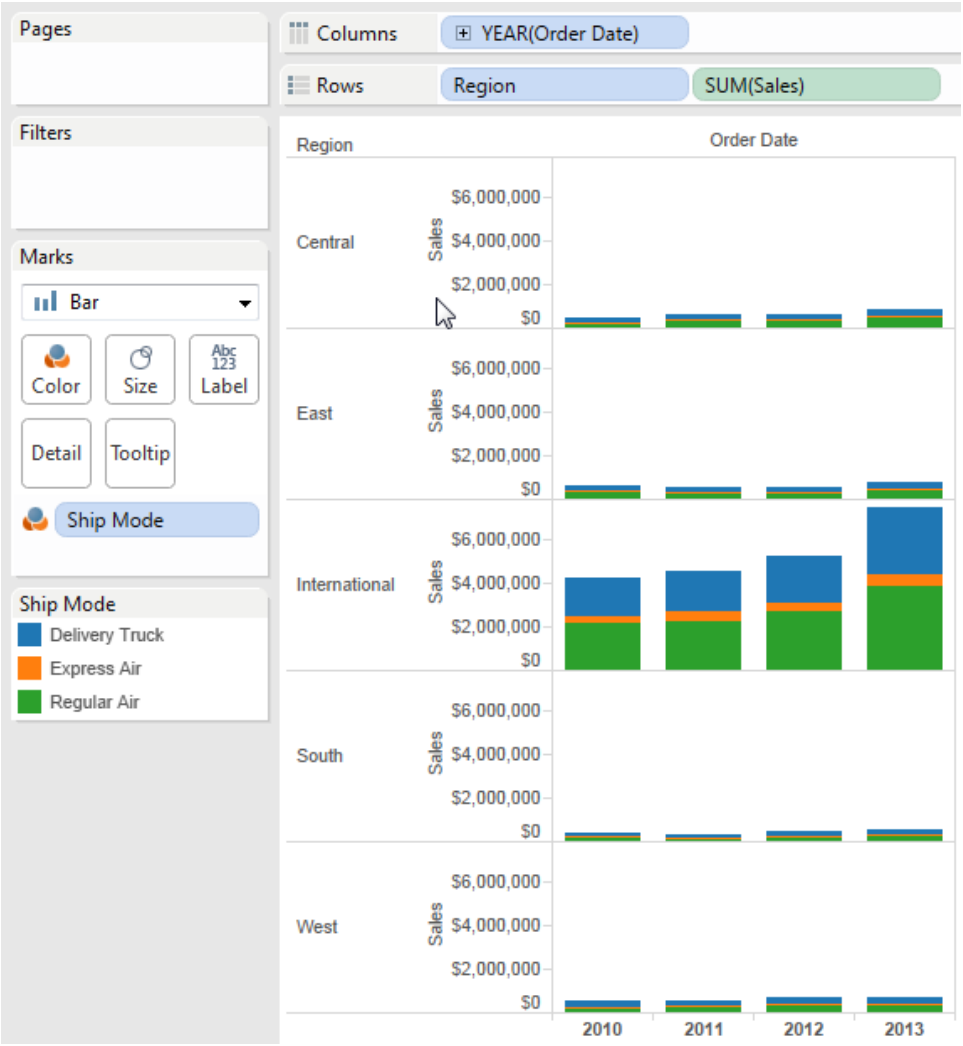

9. To better assess how sales vary by ship mode in the US, you can filter out the International region, either temporarily or permanently. To do this, drag the **Region** dimension again, this time from the Data window, to the **Filters** shelf.

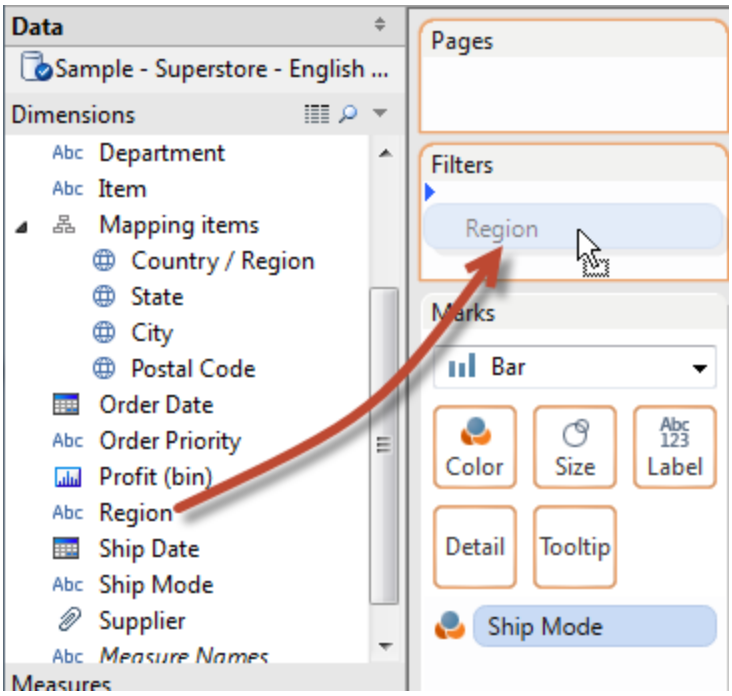

10. This opens the Filter [Region] dialog box. Clear the **International** check box and then click **OK**:

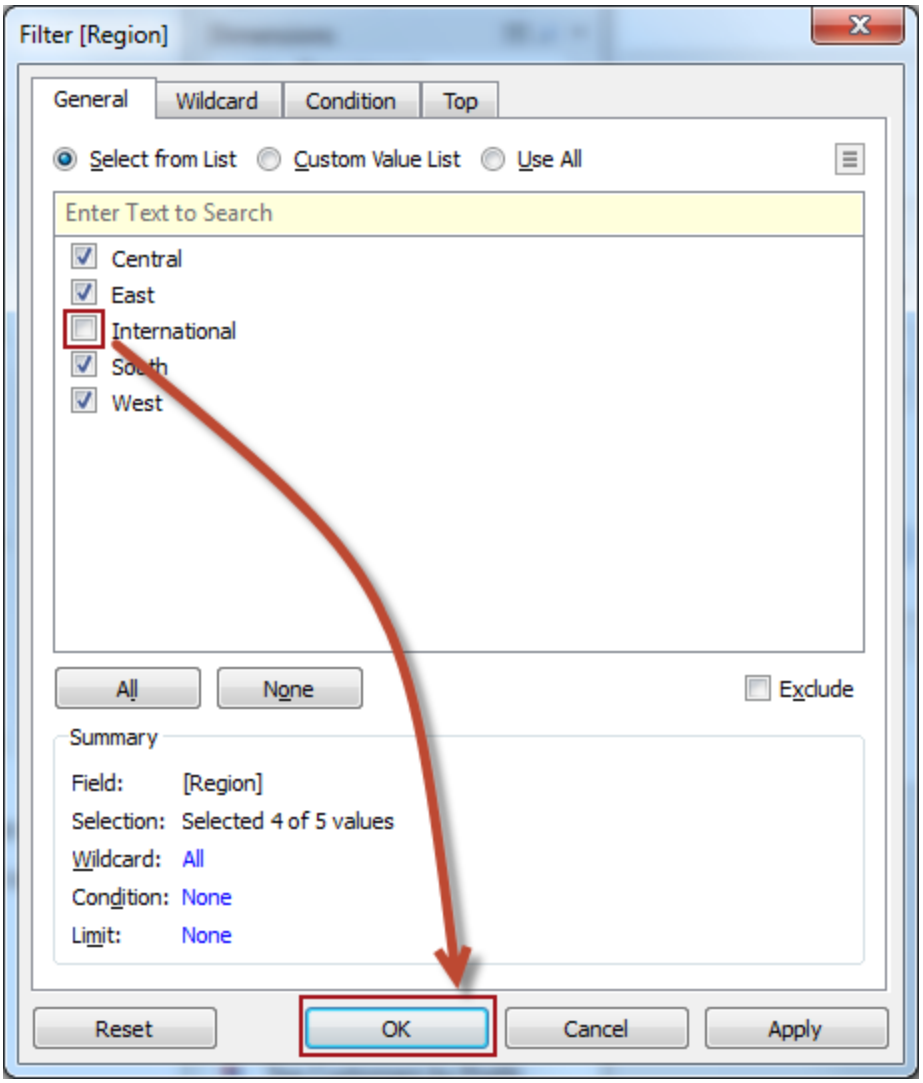

Your final view should look like this:

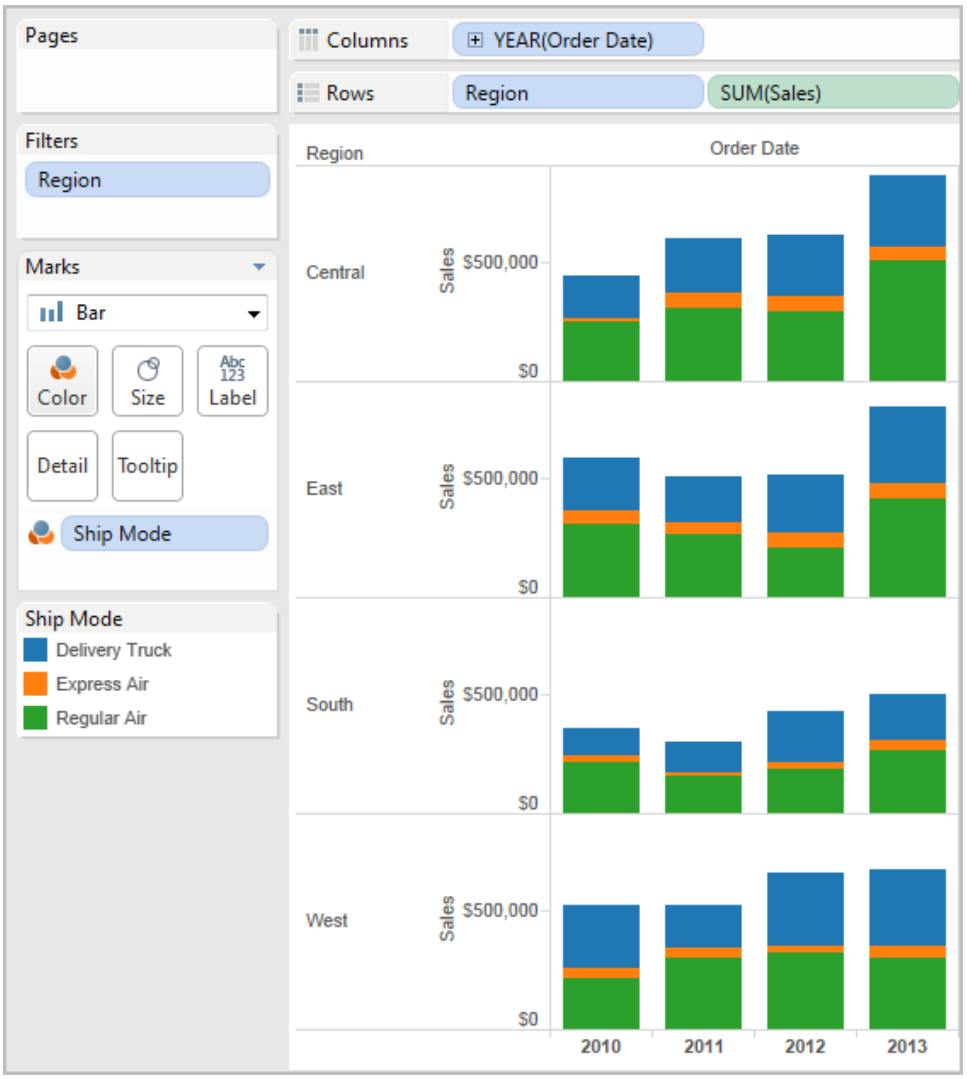

This view gives you enough information to answer some questions about your data—for example, which regions have shown consistent growth in sales over the four-year period. But answers only lead to more questions, and you can go on developing the view for as long as you like.

Building Text Tables

In Tableau, you typically create text tables (also called cross-tabs or pivot tables) by placing one dimension on the **Rows** shelf and another dimension on the **Columns** shelf. You then complete the view by dragging one or more measures to **Text** on the Marks card.

A text table uses the text mark type. Tableau chooses this mark type automatically if the view is constructed using only dimensions (assuming the Mark type is set to **Automatic**).

The following exercise walks you through creating a text table that shows sales totals by year and category.

1. Connect to the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.
2. Drag the **Order Date** dimension to **Columns**.

   The date is automatically aggregated by year, and column headers are created.
3. Drag the **Category** dimension to **Rows**.

   Row headers are created. Columns with headers plus rows with headers means that a valid table structure now exists. You must now add a measure to the view to see actual data (that is, numbers).
4. Drag the **Sales** measure to **Text** on the Marks card.

   The measure is automatically aggregated as a sum. The resulting view looks like this:

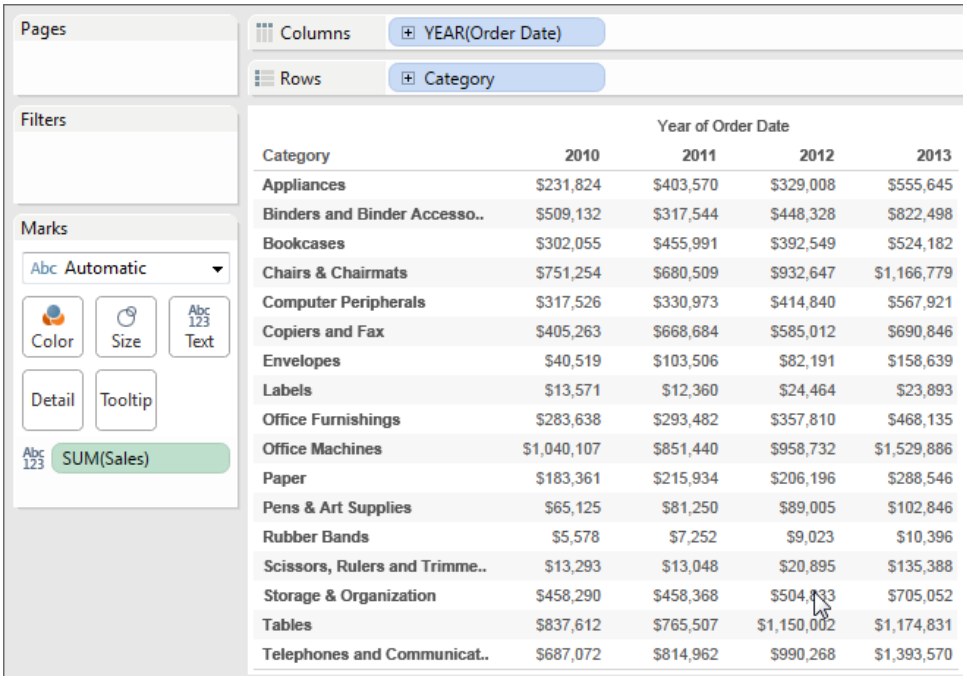

	Year of Order Date			
Category	2010	2011	2012	2013
Appliances	$231,824	$403,570	$329,008	$555,645
Binders and Binder Accesso..	$509,132	$317,544	$448,328	$822,498
Bookcases	$302,055	$455,991	$392,549	$524,182
Chairs & Chairmats	$751,254	$680,509	$932,647	$1,166,779
Computer Peripherals	$317,526	$330,973	$414,840	$567,921
Copiers and Fax	$405,263	$668,684	$585,012	$690,846
Envelopes	$40,519	$103,506	$82,191	$158,639
Labels	$13,571	$12,360	$24,464	$23,893
Office Furnishings	$283,638	$293,482	$357,810	$468,135
Office Machines	$1,040,107	$851,440	$958,732	$1,529,886
Paper	$183,361	$215,934	$206,196	$288,546
Pens & Art Supplies	$65,125	$81,250	$89,005	$102,846
Rubber Bands	$5,578	$7,252	$9,023	$10,396
Scissors, Rulers and Trimme..	$13,293	$13,048	$20,895	$135,388
Storage & Organization	$458,290	$458,368	$504,833	$705,052
Tables	$837,612	$765,507	$1,150,002	$1,174,831
Telephones and Communicat..	$687,072	$814,962	$990,268	$1,393,570

   Tableau automatically uses text as the mark type. Each cell in the table displays the sum of the sales for a particular year and category.

   We can see that Office Machines had the highest sales in every year but one.
5. Drag the **Continent** dimension to **Rows** and drop it to the left of **Category**. Do this by positioning **Continent** to the left of **Category** until you see a small triangle indicating that the new field will be inserted to the left of the existing field:

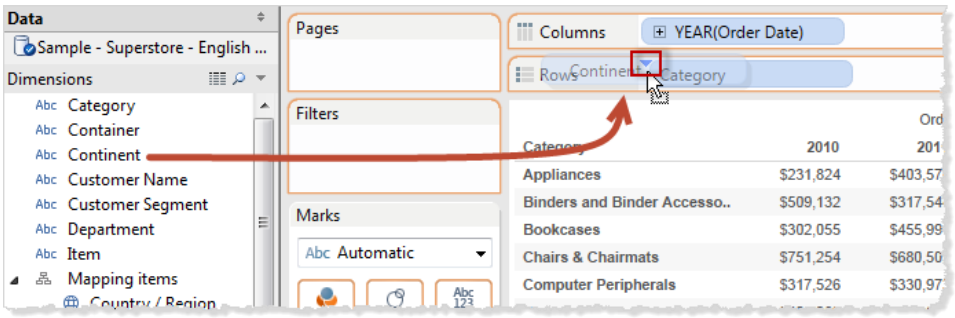

The view updates to break out sales by continent, in addition to year and category:

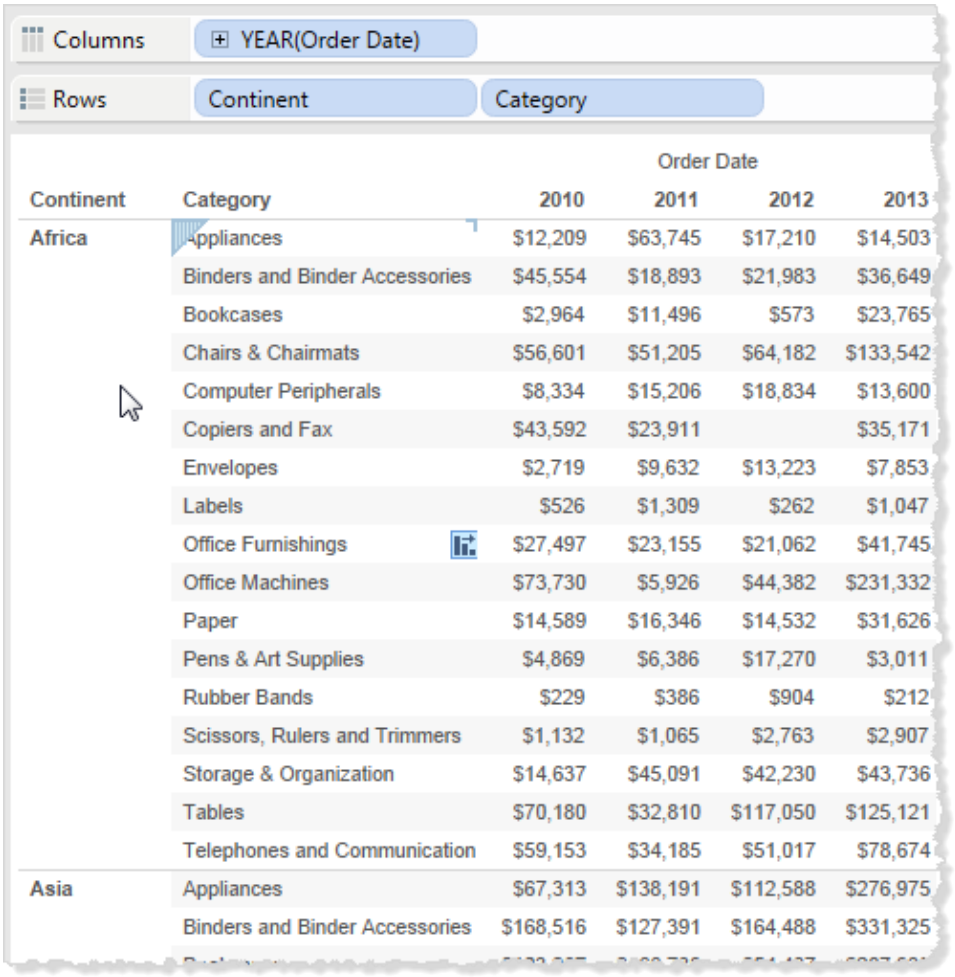

Continents are listed alphabetically, with Africa first. If you'd dropped **Continent** to the right of **Category**, the view would have been organized first by category, and then by continent:

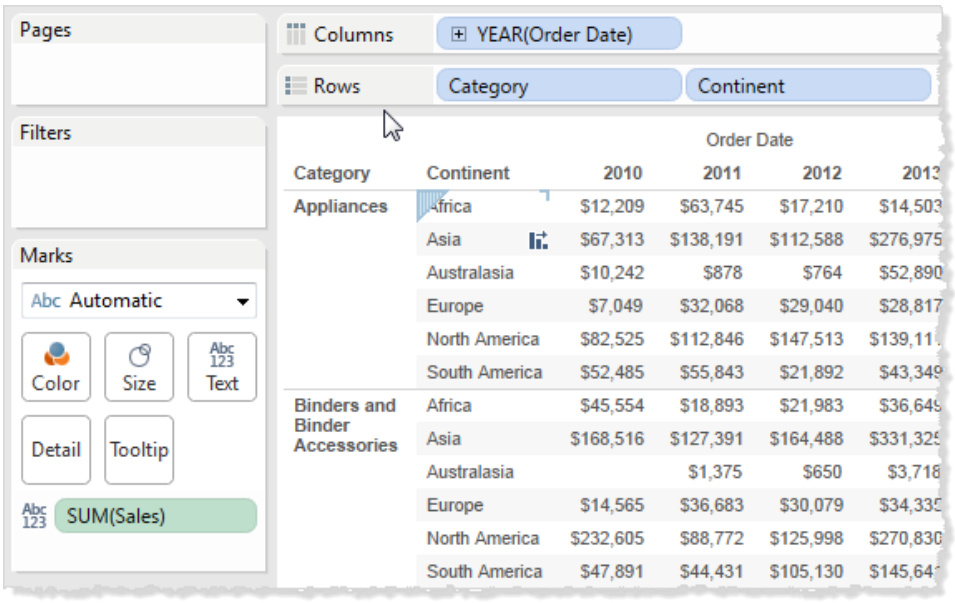

You can use a table calculation to show percentages of total instead of raw dollar values. The key question with a table calculation is how to frame the calculation--the more dimensions in your view, the more complicated this question can become. In the current case there are three dimensions in the view—**Order Date**, **Category**, and **Continent**. You could show percentages of total for a single dimension, but that can be unwieldy. For example, if you show percentages just by continent, the percentages would be calculated across the two remaining dimensions: **Category** (there are 17 categories) and **Year(Order Date)** (there are 4 years). So you would be dividing the total 17 x 4 = 68 ways. That would make for some tiny percentages. What might be useful is to show percentages using two dimensions: **Year(Order Date)** and **Continent**. Then the percentages would be calculated on the remaining dimension: **Category**. To put that in visual terms, you would be calculating percent of total within each red box in the following image:

Columns: YEAR(Order Date)

Rows: Continent, Category

Continent	Category	Order Date 2010	2011	2012	2013
Africa	Appliances	$12,209	$63,745	$17,210	$14,503
	Binders and Binder Accessories	$45,554	$18,893	$21,983	$36,649
	Bookcases	$2,964	$11,496	$573	$23,765
	Chairs & Chairmats	$56,601	$51,205	$64,182	$133,542
	Computer Peripherals	$8,334	$15,206	$18,834	$13,600
	Copiers and Fax	$43,592	$23,911		$35,171
	Envelopes	$2,719	$9,632	$13,223	$7,853
	Labels	$526	$1,309	$262	$1,047
	Office Furnishings	$27,497	$23,155	$21,062	$41,745
	Office Machines	$73,730	$5,926	$44,382	$231,332
	Paper	$14,589	$16,346	$14,532	$31,626
	Pens & Art Supplies	$4,869	$6,386	$17,270	$3,011
	Rubber Bands	$229	$386	$904	$212
	Scissors, Rulers and Trimmers	$1,132	$1,065	$2,763	$2,907
	Storage & Organization	$14,637	$45,091	$42,230	$43,736
	Tables	$70,180	$32,810	$117,050	$125,121
	Telephones and Communication	$59,153	$34,185	$51,017	$78,674
Asia	Appliances	$67,313	$138,191	$112,588	$276,975
	Binders and Binder Accessories	$168,516	$127,391	$164,488	$331,325

The dimensions that you use to frame your calculation are called the *addressing fields*, and the fields (in this case a single field) in which you run your calculation are the *partition fields*. For more on these concepts, see **Addressing and Partitioning** on page 594.

6. To create a table calculation to show percentages, right-click the **SUM(Sales)** field on the Marks card, and then choose **Add Table Calculation**:

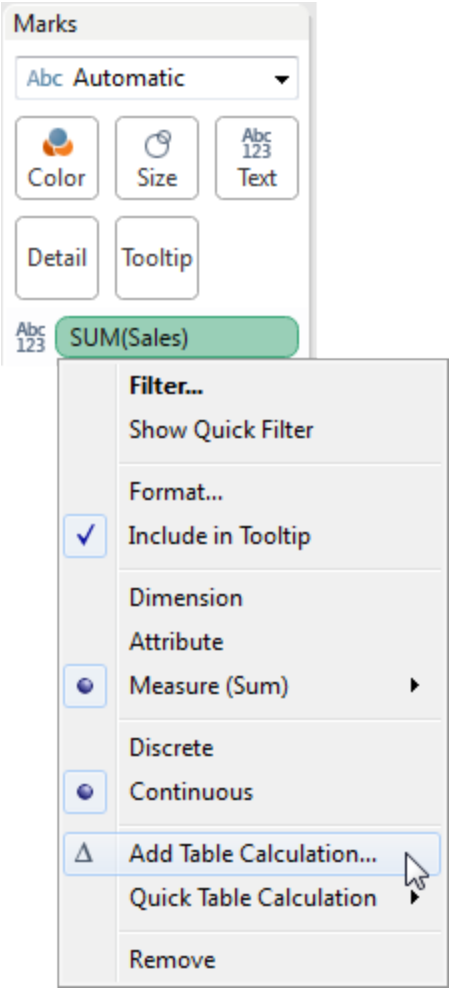

7. In the Table Calculation dialog box, set **Calculation Type** to **Percent of Total**.

   The options in the dialog box change depending on the type of calculation you choose:

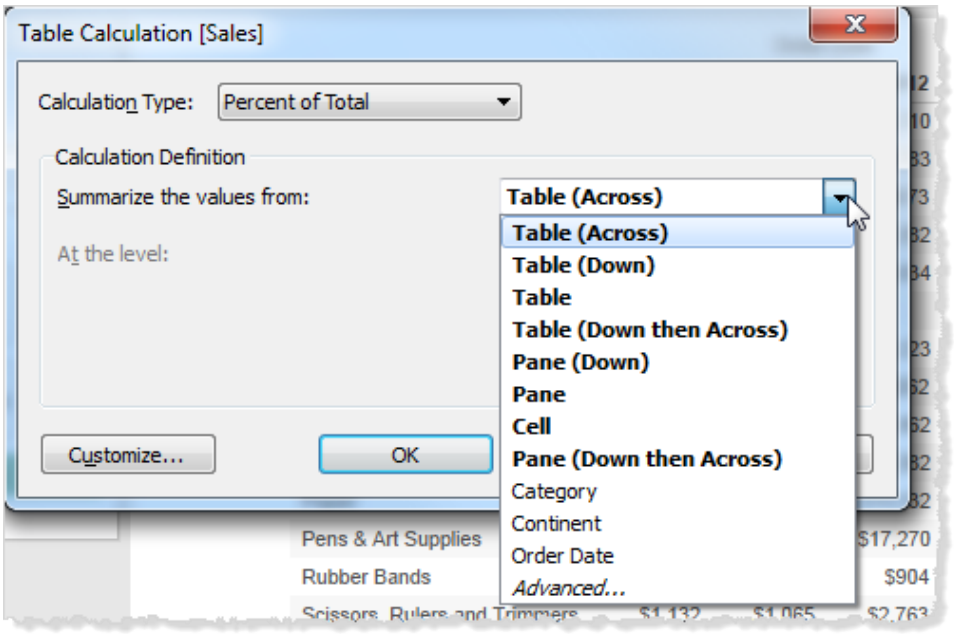

For more information on the options in the Table Calculation dialog box, see Using Table Calculations in the Tableau knowledge base.

8. Choose **Pane (Down)**, and then click **OK** to close the Table Calculation dialog box.

   Now we see percentages calculated within each category, duplicated for each year within each continent. The numbers within each red box add up to 100%:

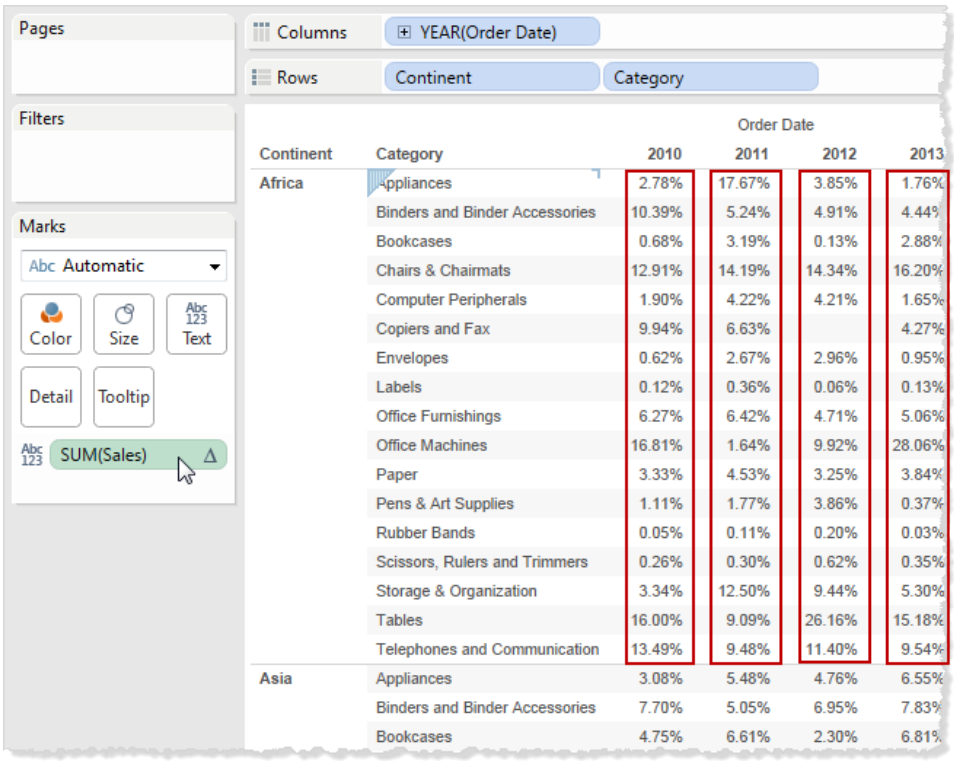

		Order Date			
Continent	Category	2010	2011	2012	2013
Africa	Appliances	2.78%	17.67%	3.85%	1.76%
	Binders and Binder Accessories	10.39%	5.24%	4.91%	4.44%
	Bookcases	0.68%	3.19%	0.13%	2.88%
	Chairs & Chairmats	12.91%	14.19%	14.34%	16.20%
	Computer Peripherals	1.90%	4.22%	4.21%	1.65%
	Copiers and Fax	9.94%	6.63%		4.27%
	Envelopes	0.62%	2.67%	2.96%	0.95%
	Labels	0.12%	0.36%	0.06%	0.13%
	Office Furnishings	6.27%	6.42%	4.71%	5.06%
	Office Machines	16.81%	1.64%	9.92%	28.06%
	Paper	3.33%	4.53%	3.25%	3.84%
	Pens & Art Supplies	1.11%	1.77%	3.86%	0.37%
	Rubber Bands	0.05%	0.11%	0.20%	0.03%
	Scissors, Rulers and Trimmers	0.26%	0.30%	0.62%	0.35%
	Storage & Organization	3.34%	12.50%	9.44%	5.30%
	Tables	16.00%	9.09%	26.16%	15.18%
	Telephones and Communication	13.49%	9.48%	11.40%	9.54%
Asia	Appliances	3.08%	5.48%	4.76%	6.55%
	Binders and Binder Accessories	7.70%	5.05%	6.95%	7.83%
	Bookcases	4.75%	6.61%	2.30%	6.81%

Pane (Down) is the correct choice because it specifies that the calculation should be performed from top to bottom within each pane of the table. The table has two vertical dimensions, so **Table (Down)** would have calculated the percent of total from top to bottom for the entire table, ignoring the **Continent** dimension. The pane is always the finest level of detail for the relevant direction (across or down). If you had three dimensions on the vertical axis, you might have had to use field names to define the calculation, since only the dimension furthest to the left on the **Rows** shelf (defined as Table) and the dimension furthest to the right (defined as Pane) could be captured with the structural options.

If you find these concepts daunting, you're not alone. It's not inappropriate to use trial-and-error to see the results of different definitions for table calculations. But you should have a clear idea of what result you want, so that you can recognize it when you see it.

Building Line Charts

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, or to forecast future values. For more information about the line mark type, see **Line Mark** on page 242.

The following exercise walks you through creating a view that displays the sum of the sales and the sum of the profit for all years and then uses forecasting to determine a trend.

1. Connect to the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.
2. Drag the **Order Date** dimension to **Columns**.

   The date is automatically aggregated by year, and column headers are created.
3. Drag the **Sales** measure to **Rows**.

   Tableau automatically aggregates **Sales** as SUM, and displays a simple line chart.
4. Drag the **Profit** measure to **Rows**, and drop it to the right of the **Sales** measure.

   Tableau creates separate axes along the left margin for **Sales** and **Profit**:

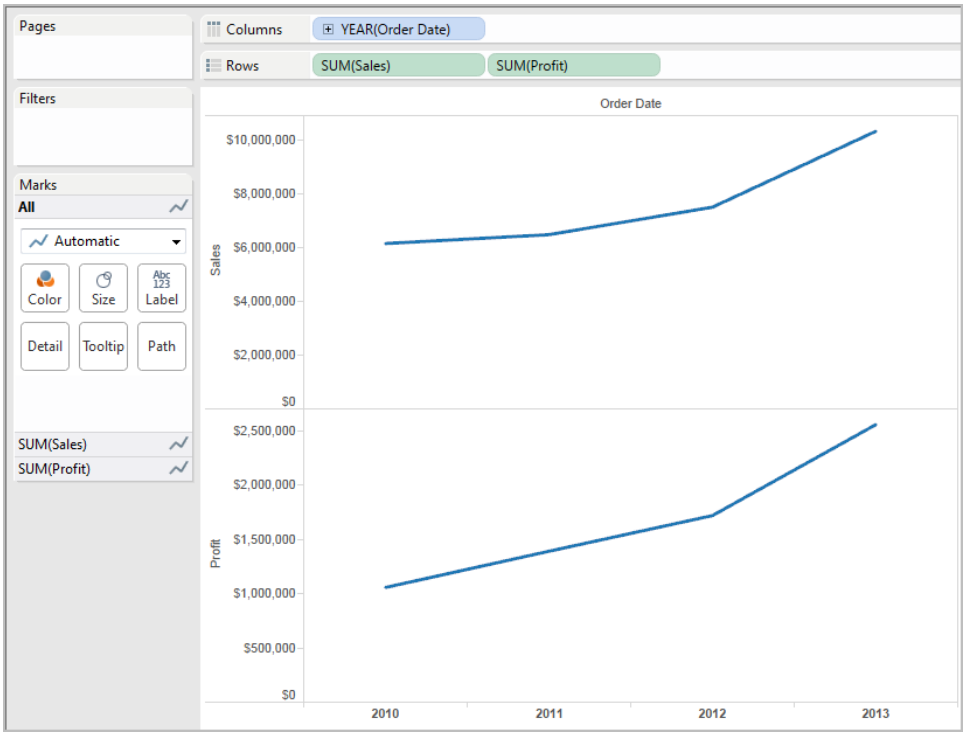

   Notice that the scale of the two axes is different—the **Sales** axis scales from $0 to $10,000,000, whereas the **Profit** axis scales from $0 to $2,500,000. This can make it hard to see that sales values are much greater than profit values. When you are displaying multiple measures in a line chart, you can align or merge axes to make it easier for users to compare values. Use **Dual Axes** on page 330 to align the axes or use **Blended Axes** on page 328 to enforce a single axis across multiple measures.

   With either of these options, you can create a combination chart to change the mark type for one of your measures. See **Combination Charts** on page 332.
5. Drag the **SUM(Profit)** field from **Rows** to the **Sales** axis to create a blended axis. The two pale green parallel bars (shown in the red box) indicate that **Profit** and **Sales** will use a blended axis when you release the mouse button:

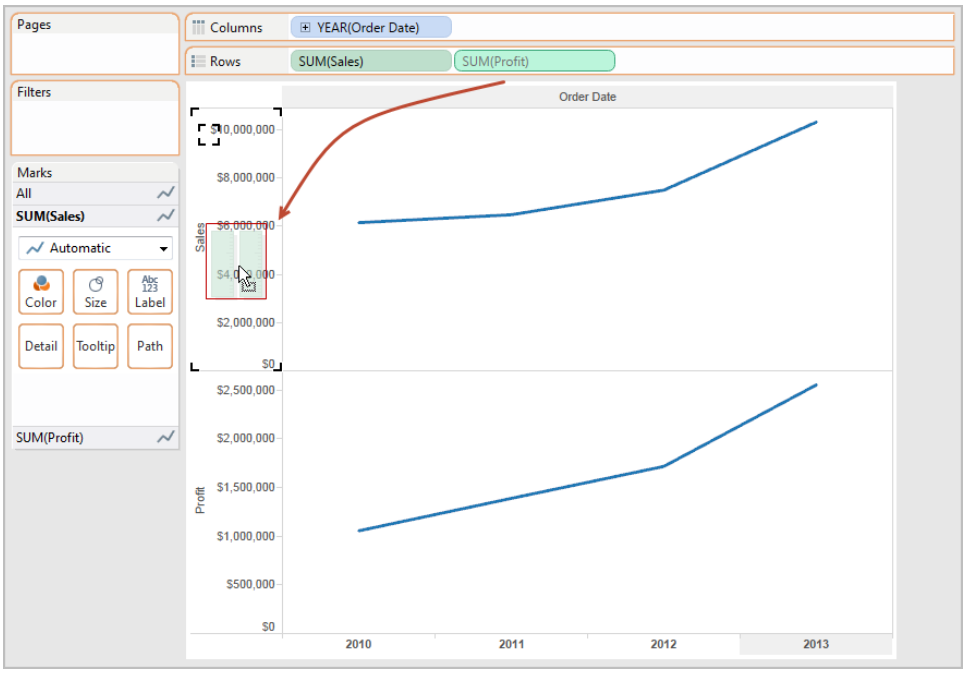

The view updates to look like this:

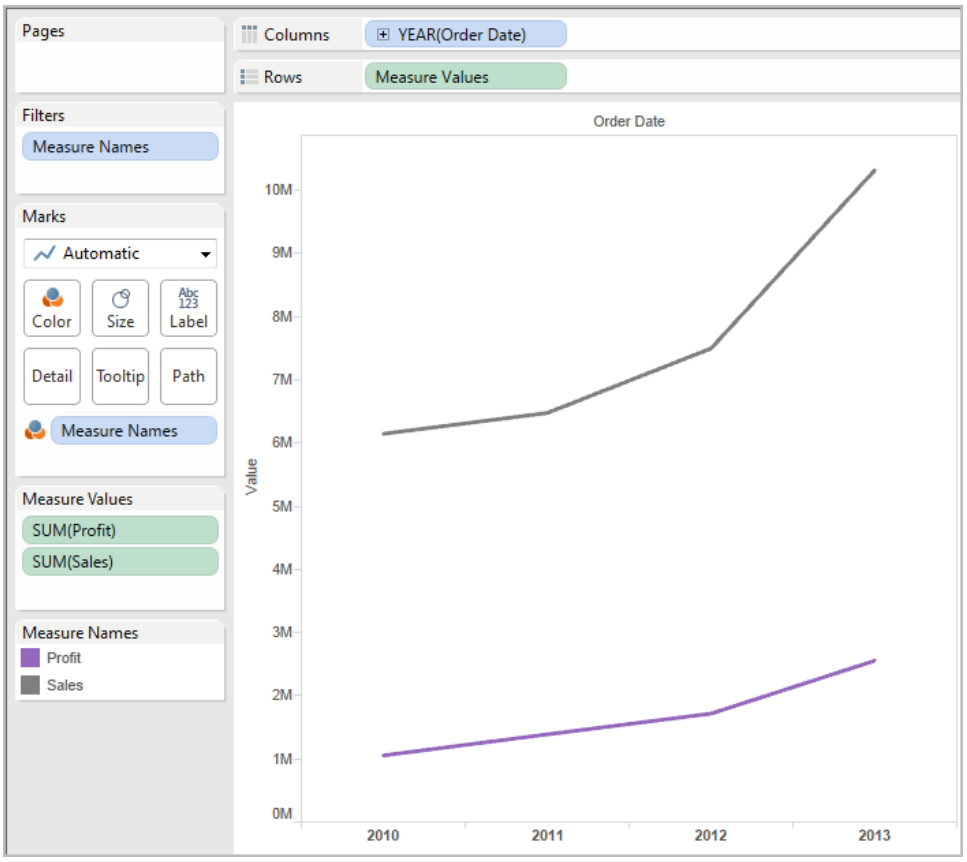

The view is rather sparse, because we are looking at a summation of values on a per-year basis.

6. Click the **Year(Order Date)** field in the view and select **Month** in the lower part of the content menu (see image below). This tells Tableau that you want to look at a continuous range of values over the four-year period:

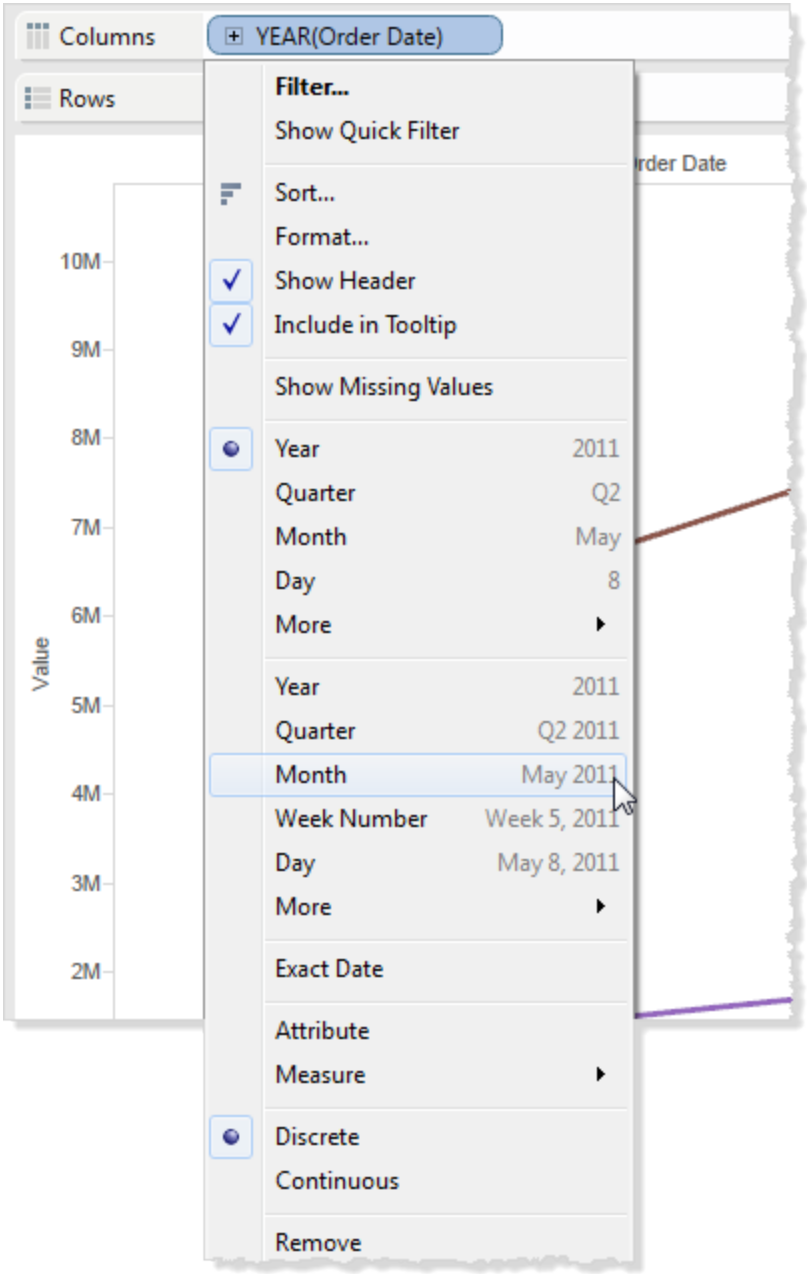

The resulting view is a lot more detailed than the original view:

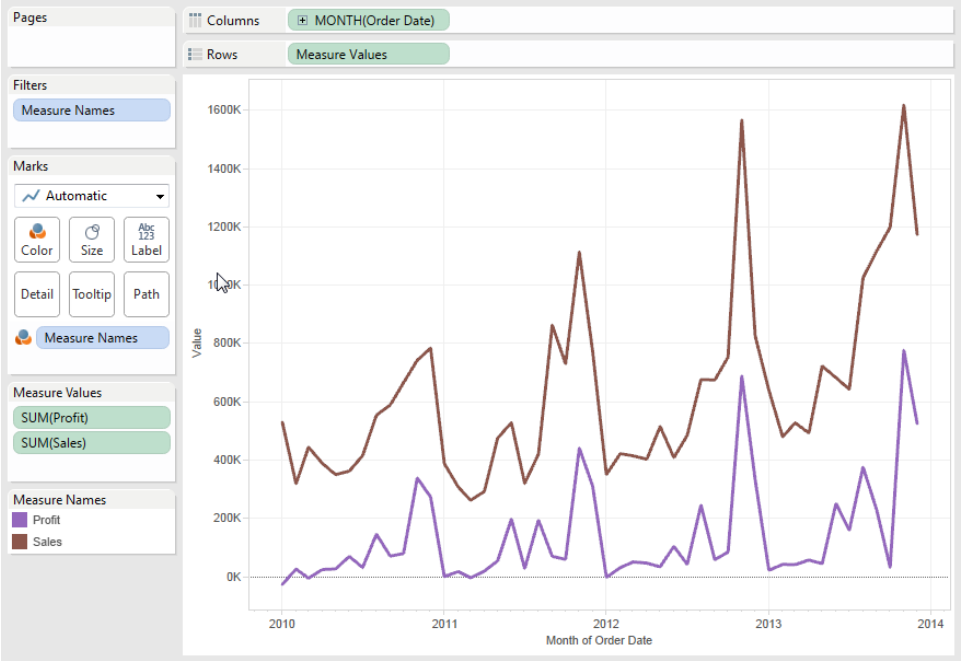

Notice that the values seem to go much higher just before the end of each year. A pattern like that is known as *seasonality*. If we turn on the forecasting feature in the view, we can see whether we should expect that apparent seasonal trend continuing into the future.

7. Right-click in the view and choose **Forecast >Show Forecast**. We then see that, according to Tableau forecasting, the seasonal trend does continue into the future:

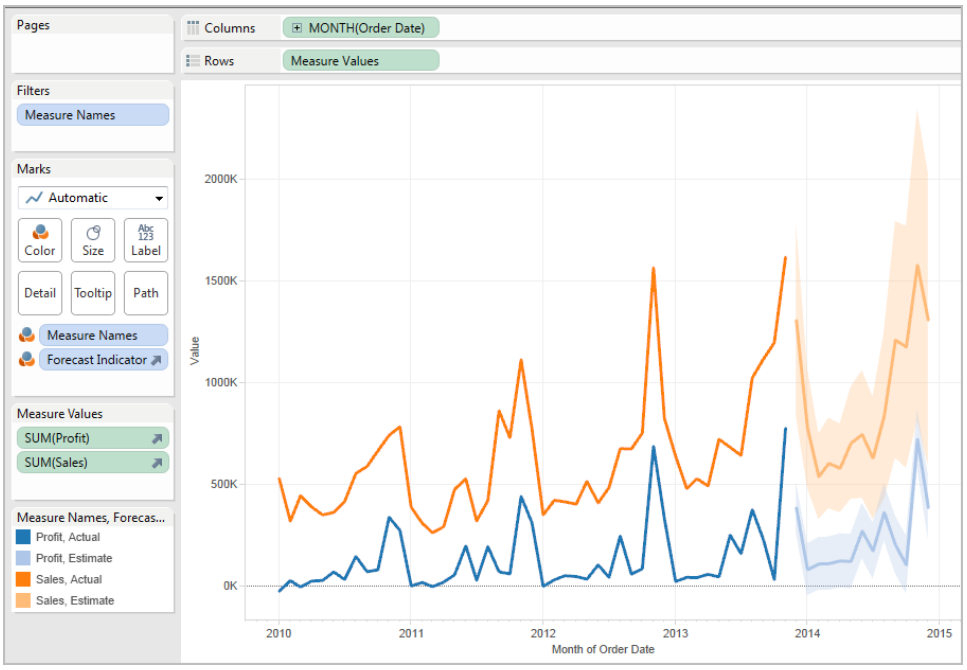

Building Scatter Plots

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 the **Columns** shelf and at least one measure on the **Rows** shelf. If these shelves contain both dimensions and measures, Tableau automatically places the measures as the innermost fields, which means that measures are always to the right of any dimensions you have also placed on these shelves. The word "innermost" in this case refers to the table structure.

Creates Simple Scatter Plot		Creates Matrix of Scatter Plots		
Columns	SUM(Sales)	Columns	Region	SUM(Sales)
Rows	SUM(Profit)	Rows	Category	SUM(Profit)

A scatter plot can use several mark types. By default, Tableau uses the shape mark type. Depending on your data, you might want to use another mark type, such as a circle or a square. For more information, see **Mark Types** on page 239.

The following exercise walks you through using scatter plots and trend lines to compare sales to profit:

1. Open the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.
2. Drag the **Profit** measure to **Columns**.

   The measure is automatically aggregated as a sum and a horizontal axis is created.
3. Drag the **Sales** measure to **Rows**.

   The measure is automatically aggregated as a sum and a vertical axis is created.

   Measures contain continuous numerical data. When you plot one number against another, you are comparing two numbers; the resulting chart is analogous to a Cartesian chart, with x and y coordinates.

   What you have at this point is the most basic kind of scatter plot: a one-mark scatter plot:

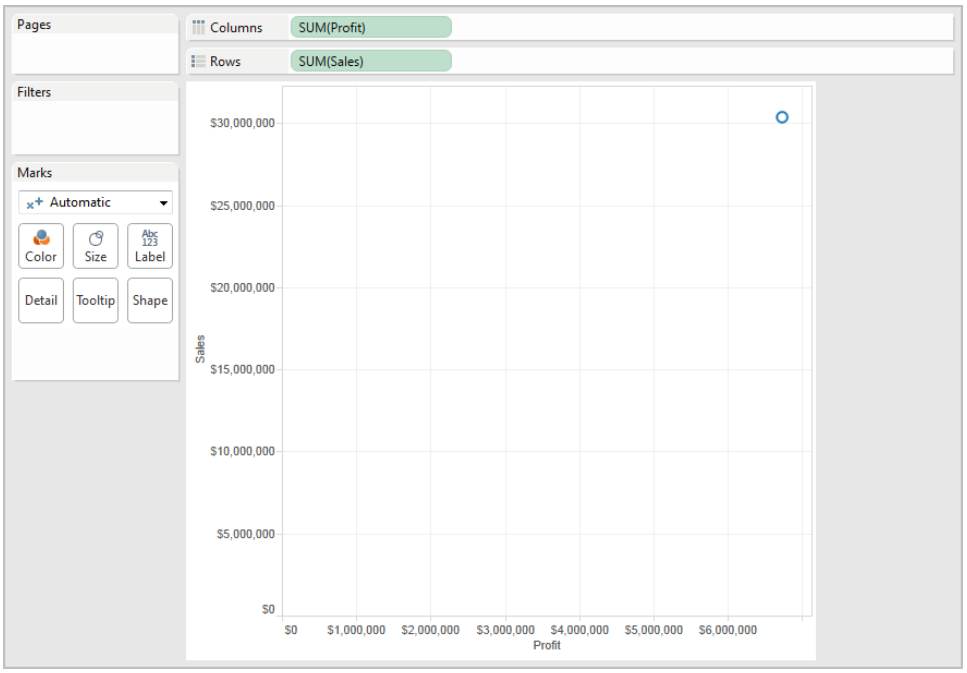

4. Drag the **Department** dimension to **Color** on the Marks card.

   This separates the data into three marks—one for each dimension member—and encodes the marks using color.

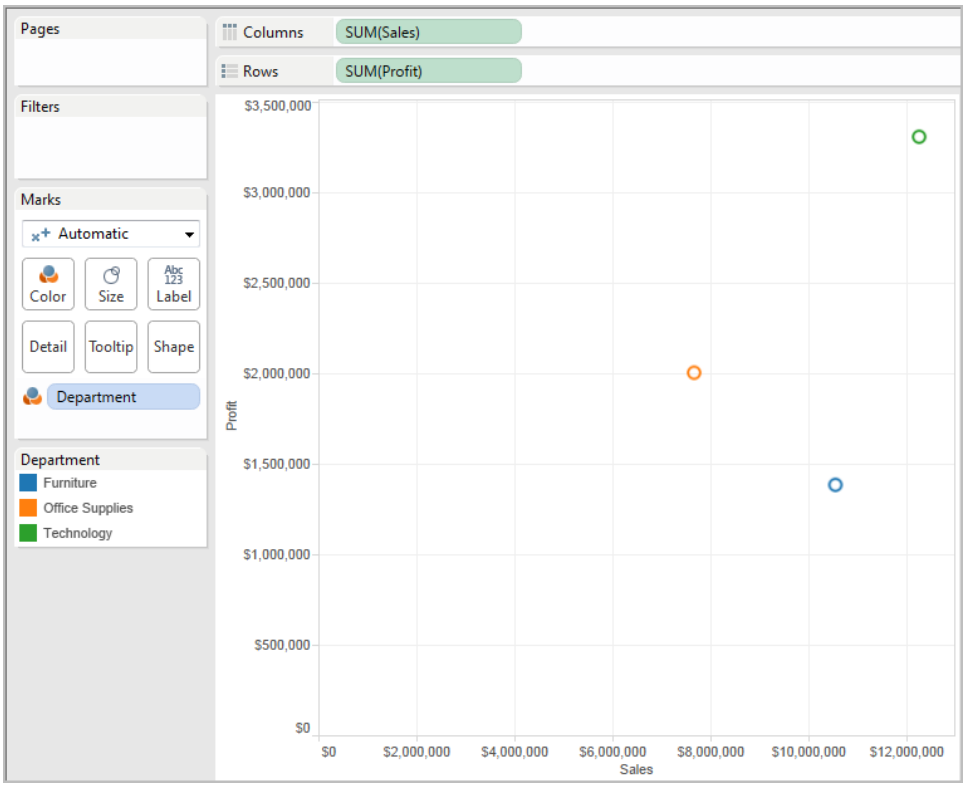

5. Drag the **Region** dimension to **Detail** on the Marks card.

   Now there are many more marks in the view. The number of marks is equal to the number of distinct country/regions in the data source multiplied by the number of departments. (If you're curious, see what would have happened if you'd dropped the **Region** dimension on **Shape** instead of Detail.)

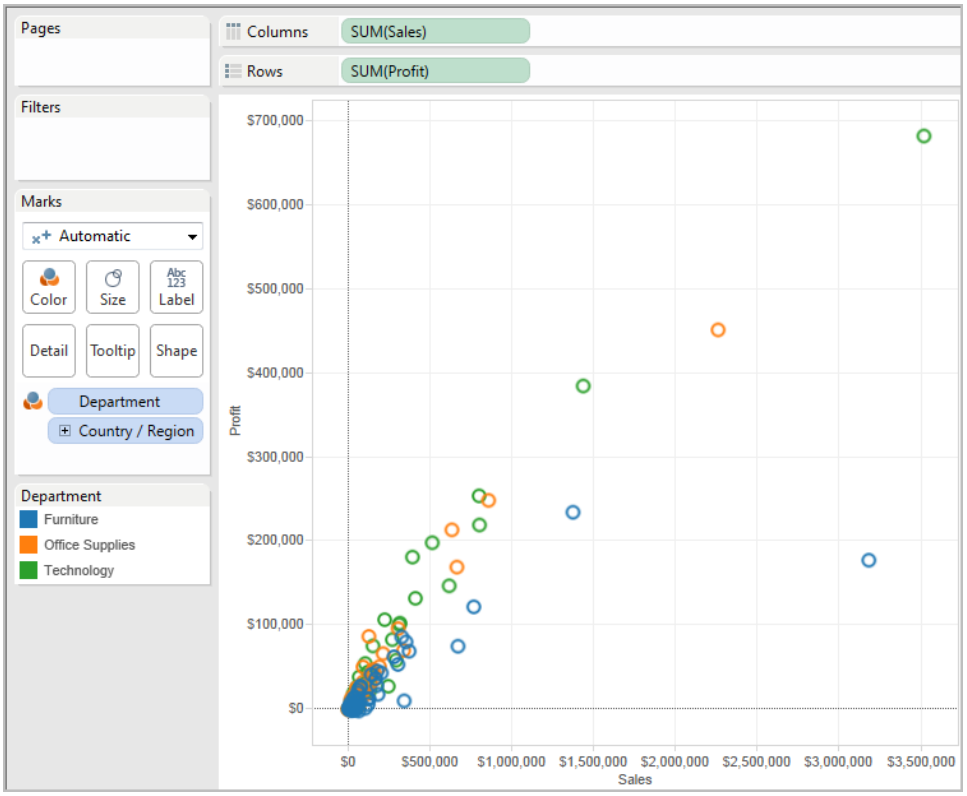

6. Right-click in the view and choose **Trend Lines >Show Trend Lines** to add trend lines.

   A trend line can provide a statistical definition of the relationship between two numerical values. To add trend lines to a view, both axes must contain a field that can be interpreted as a number—by definition, that is always the case with a scatter plot.

   Tableau adds three linear trend lines—one for each color that you are using to distinguish the three departments:

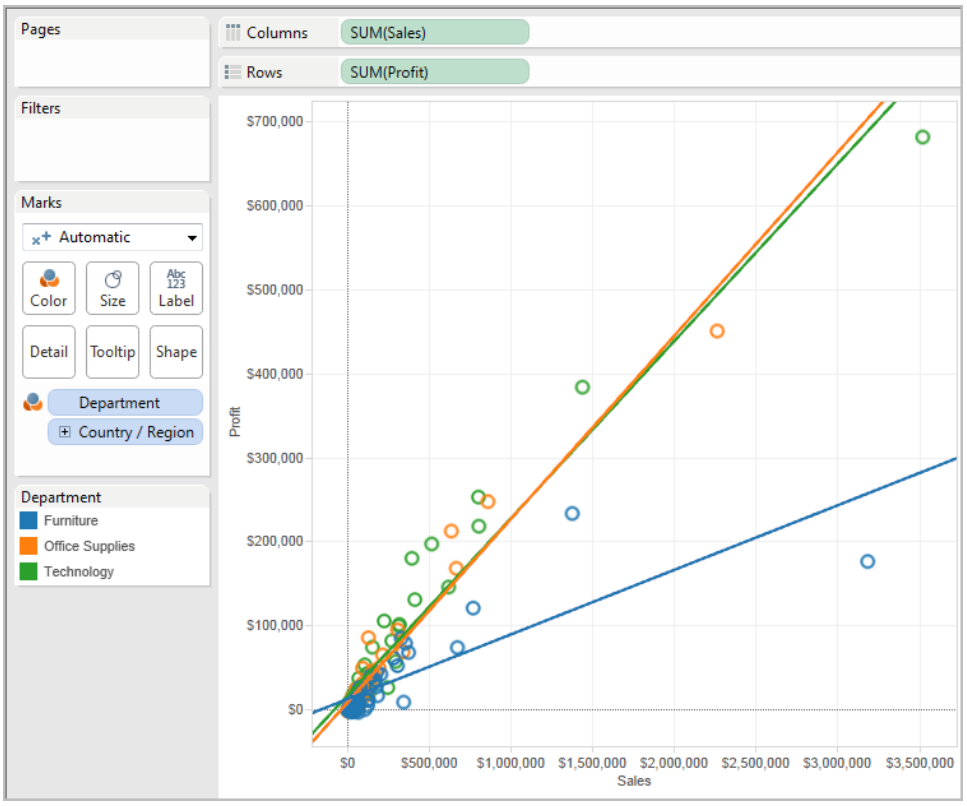

7. Hover the cursor over the trend lines to see statistical information about the model that was used to create the line:

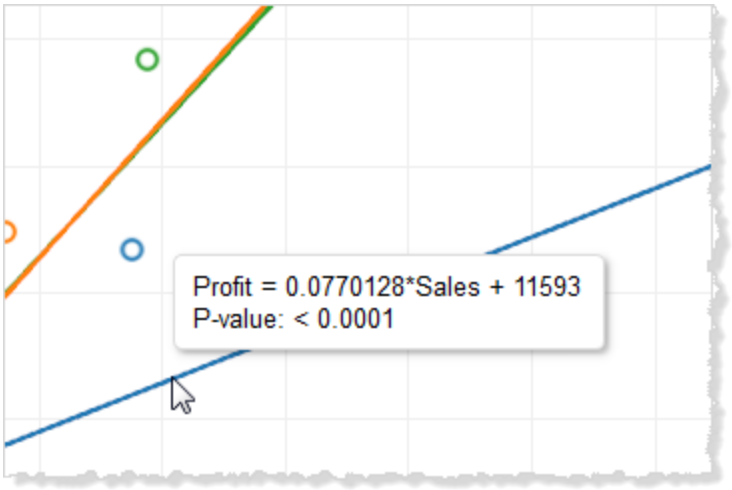

For more information, see **Assessing Trend Line Significance** on page 723. You can also customize the trend line to use a different model type or to include confidence bands. See **Adding Trend Lines** on page 717.

See Also

Example – Scatter Plots and Aggregation on page 572

Building Heat Maps

Heat maps are a great way to compare categorical data using 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 **Square** as the mark type and place a measure of interest on the **Color** shelf. You can enhance this basic heat map by size-encoding and/or shape-encoding the cells in the table.

The following exercise walks you through using a heat map to explore how profit varies across regions, product categories, and customer segments.

1. Connect to the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.
2. Drag the **Customer Segment** dimension to **Columns**.

   Headers are created with labels derived from the dimension member names.
3. Drag the **Region** and **Category** dimensions to **Rows**, dropping **Category** to the right of **Region**.

   Headers are created with labels derived from the dimension member names. You have now created a nested table of categorical data (the **Category** dimension is nested within the **Region** dimension).
4. Drag the **Profit** measure to **Color** on the Marks card.

   The measure is automatically aggregated as a sum. The color legend reflects the continuous data range.
5. Optimize the view format:
   - From the **Format** menu, select **Cell Size > Square Cell**.
   - Increase the mark size by pressing Ctrl+Shift+B. Hold down Ctrl+Shift and continue to press B until the squares are large enough.
   - Make the columns wider by pressing Ctrl+Right arrow. Hold down Ctrl and continue pressing the Right arrow key until the headings for **Category** are displayed in full:

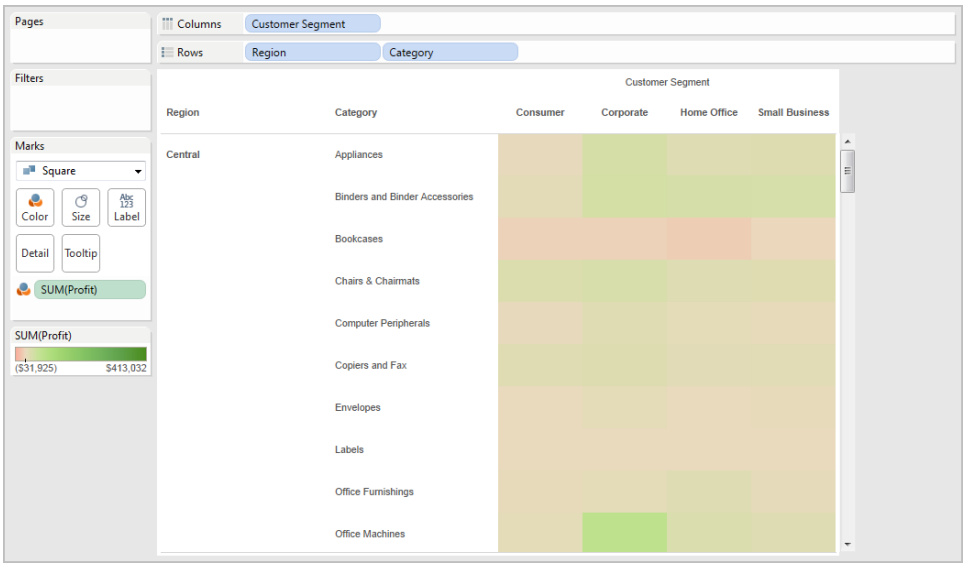

These cell formatting options are also available when you select **Cell Size** from the **Format** menu.

In this view, you can see data for only for the Central region. You must scroll down to see data for other regions. Or you can maximize the Tableau Desktop window.

In the Central region, office machines seem to be the most profitable category, and bookcases the least profitable.

6. Click **Color** on the Marks card to display configuration options. In the **Border** drop-down list and chose a medium gray color for cell borders, as in the following image:

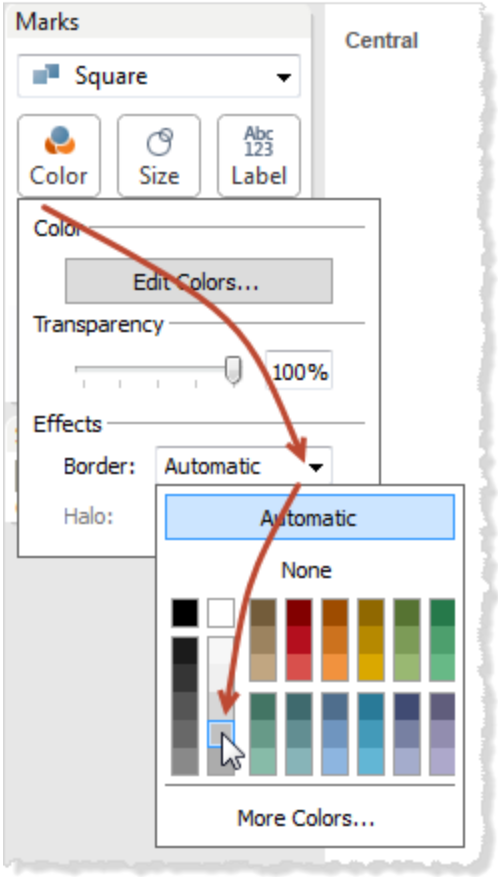

Now it's easier to see the individual cells in the view:

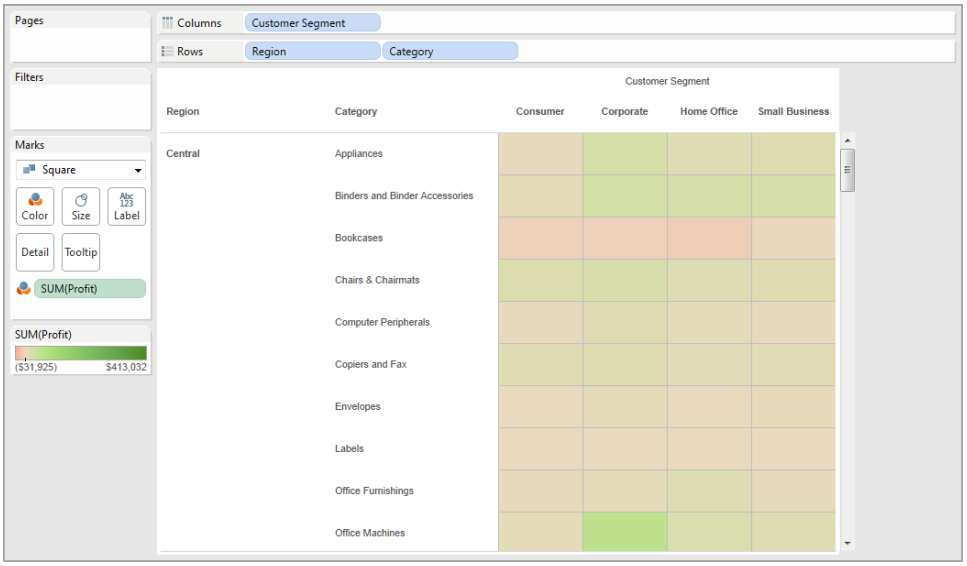

7. To make the colors more distinct, hover over the top right corner of the **SUM(Profit)**

color legend, and then click the downward triangle that appears. The following menu of options is displayed:

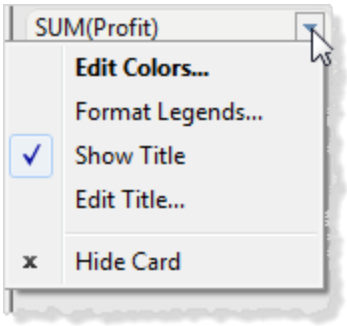

8. Click **Edit Colors**. In the Edit Colors dialog box, select **Use Full Color Range**:

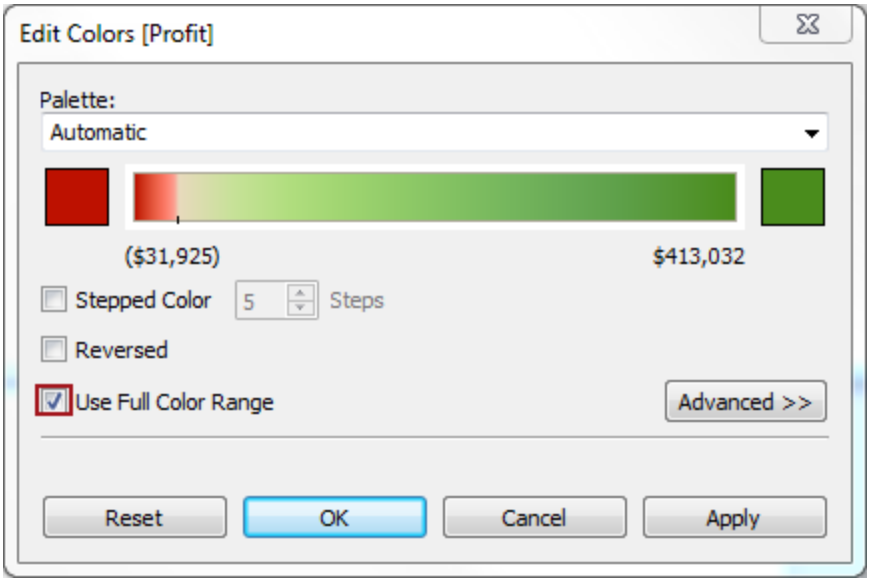

When you select this option, Tableau assigns the starting number a full intensity and the ending number a full intensity. If the range is from -10 to 100, the color representing negative numbers changes in shade much more quickly than the color representing positive numbers. If you do not select **Use 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. (See **Color** on page 257 for more on color options.) The effect is to make the color contrasts in your view much more distinct:

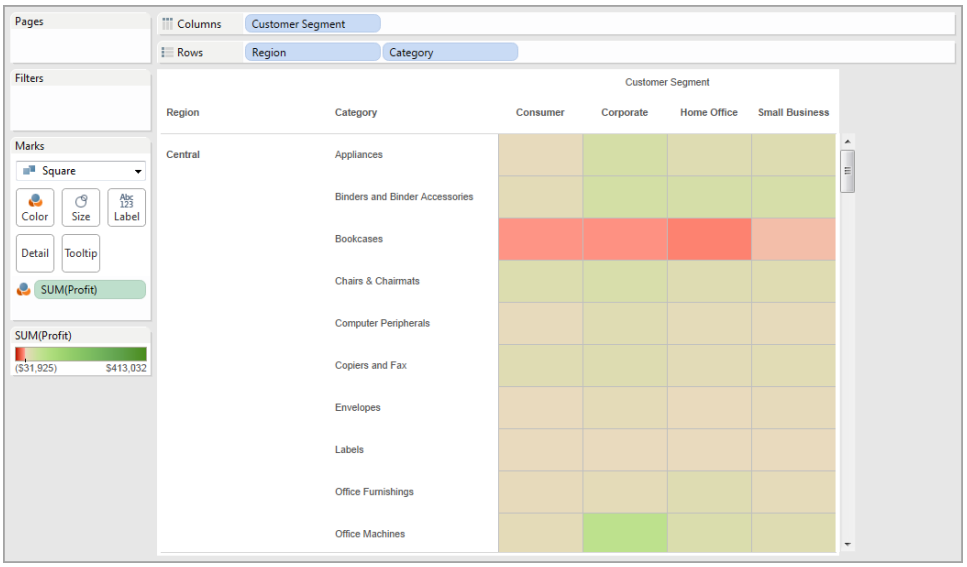

9. Drag the **Sales** measure to **Size** on the Marks card to control the size of the boxes by the Sales measure. This enables you to compare absolute sales numbers (by size of the boxes) and profit (by color).

   Initially, the marks are too small:

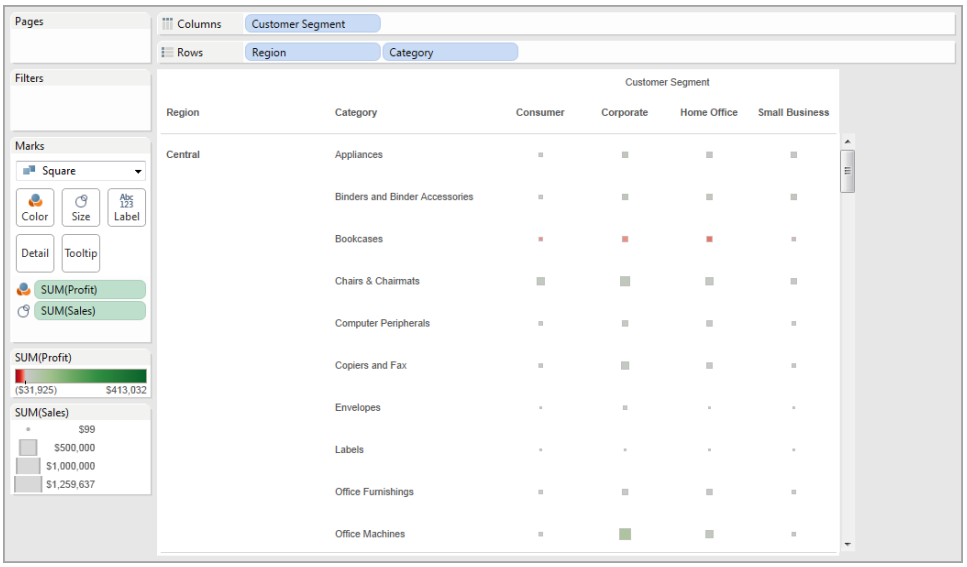

10. To adjust the size of the marks, click **Size** on the Marks card to display a size slider:

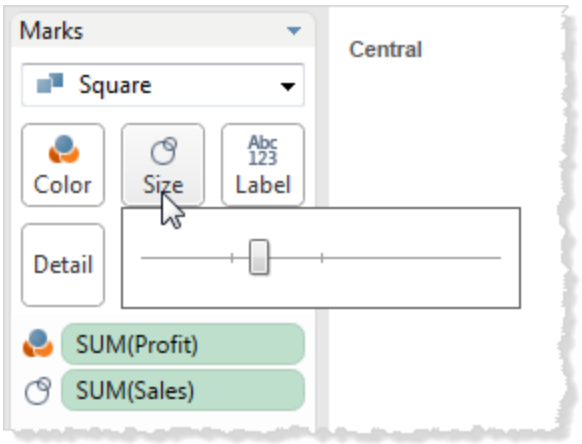

11. Drag the slider to the right until the boxes in the view are the optimal size. Now your view is complete:

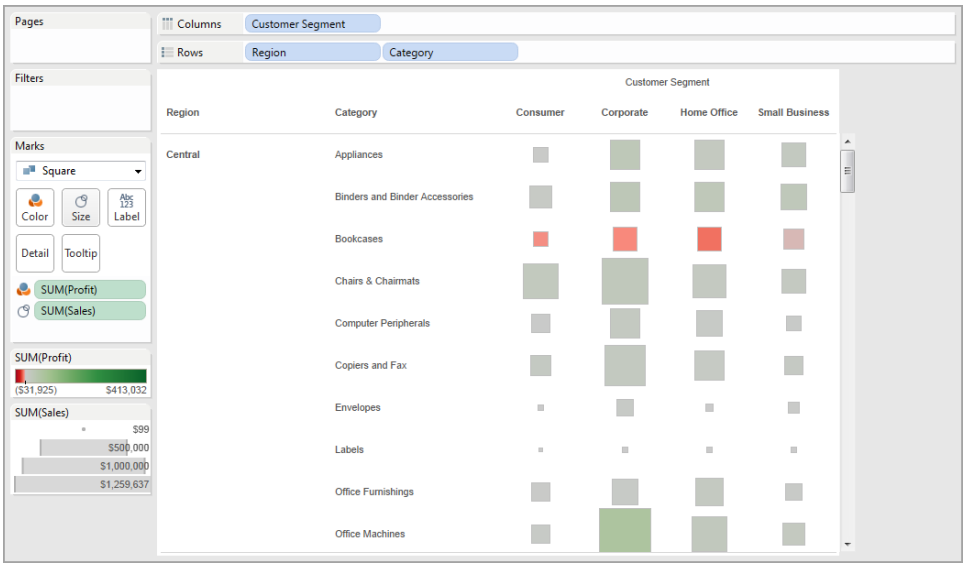

12. Use the scroll bar along the right side of the view to examine the data for different regions. Notice what's going on in the International region.

Building Gantt Bar Charts

Gantt charts are useful for displaying the duration of events or activities over time. In a Gantt Bar Chart, each separate mark (usually a bar) shows a duration. For example, you could use a Gantt chart to display average delivery time for a range of products. For more information about Gantt bar mark type, see **Gantt Bar Mark** on page 255.

The following exercise walks you through using Gantt charts to show how many days, on average, elapse between order date and ship date. The data is broken out by product category and ship mode. The view displays results on a weekly basis for a three-month period.

1. Connect to the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.
2. Drag the **Order Date** dimension to **Columns**.

   The date is automatically aggregated by year, and column headers are created with labels showing the years in your data.
3. Right-click the **Order Date** field on the **Columns** shelf and choose **Week Number**:

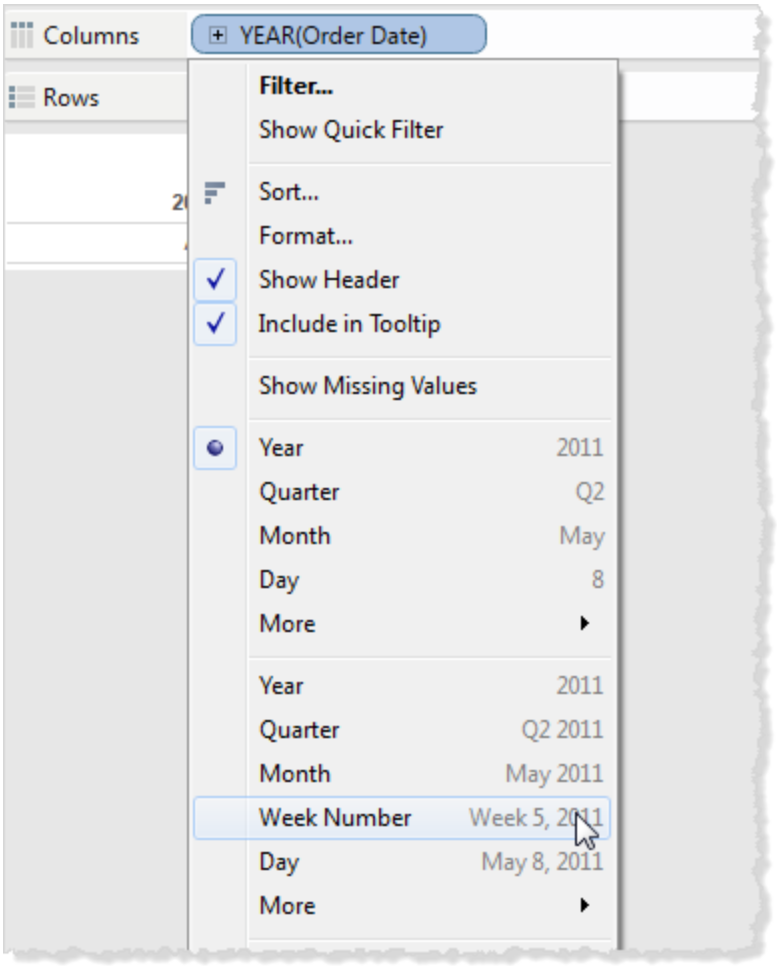

The column headers change—the individual weeks are indicated by tick marks because there are too many weeks in a four-year span (208) to show as labels in the view:

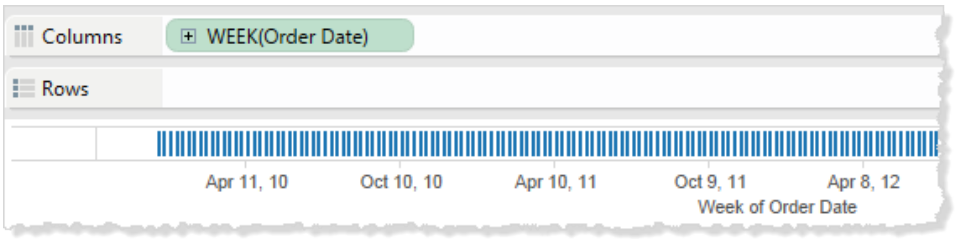

4. Drag **Category** and **Ship Mode** dimensions to the **Rows** shelf, dropping **Ship Mode** to the right of **Category**.

   This builds a two-level nested hierarchy of dimensions along the left axis:

   Now we want to size the marks according to the length of the interval between the order date and the ship date. But there are no measures that we can use to capture that interval. However, we can create such a measure by creating a calculated field.

5. Right-click anywhere in the Data window (on the far left) and choose **Create Calculated Field**:

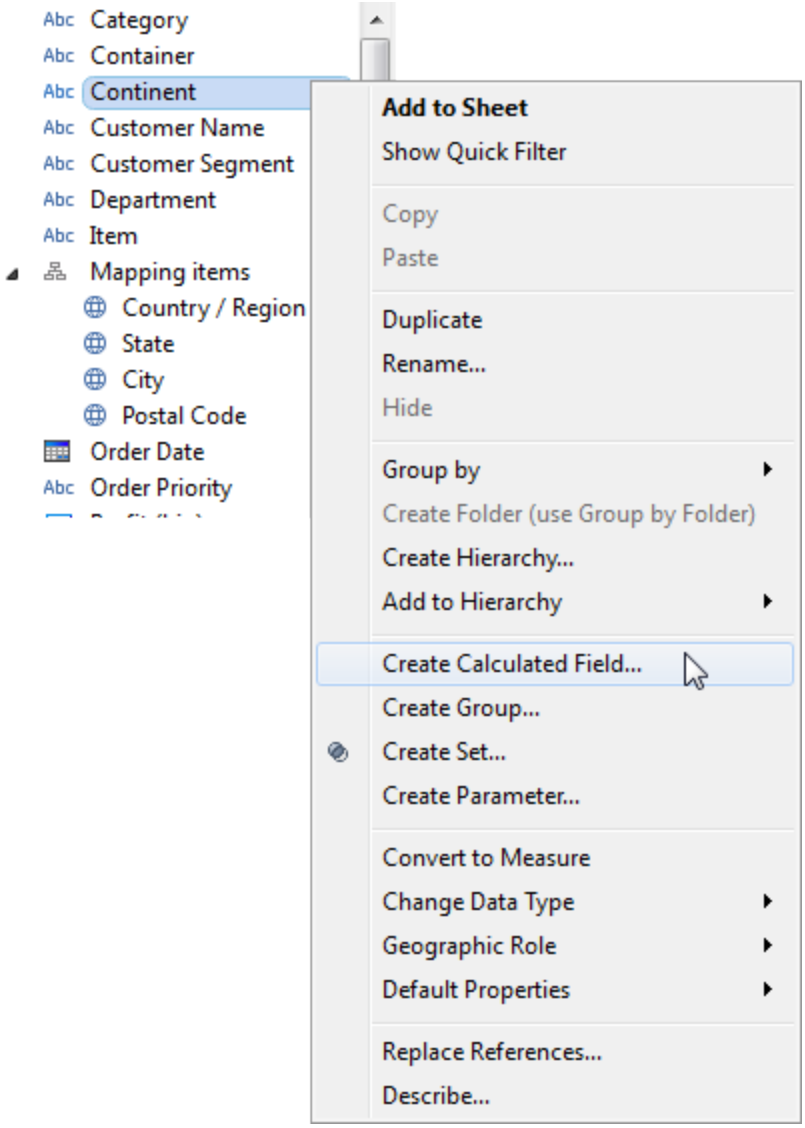

This opens the Calculated Field dialog box. This option is also available from the **Analysis** menu.

6. Name your calculated field **OrderUntilShip**.
7. Clear any content that may be in the **Formula** box by default.
8. Type (or copy) the following formula and click **OK**:

   ```
   DATEDIFF('day',[Order Date],[Ship Date] )
   ```

   This creates a custom measure that captures the difference between the **Order Date** and **Ship Date** values, in days.
9. Drag the **OrderUntilShip** measure to Size on the Marks card.

The default aggregation for **OrderUntilShip** is Sum. But we're not interested in adding all the intervals—it makes more sense to average them.

10. Right-click the **SUM(OrderUntilShip)** field on the Marks card, choose **Measure (Sum)**, and then choose **Average**:

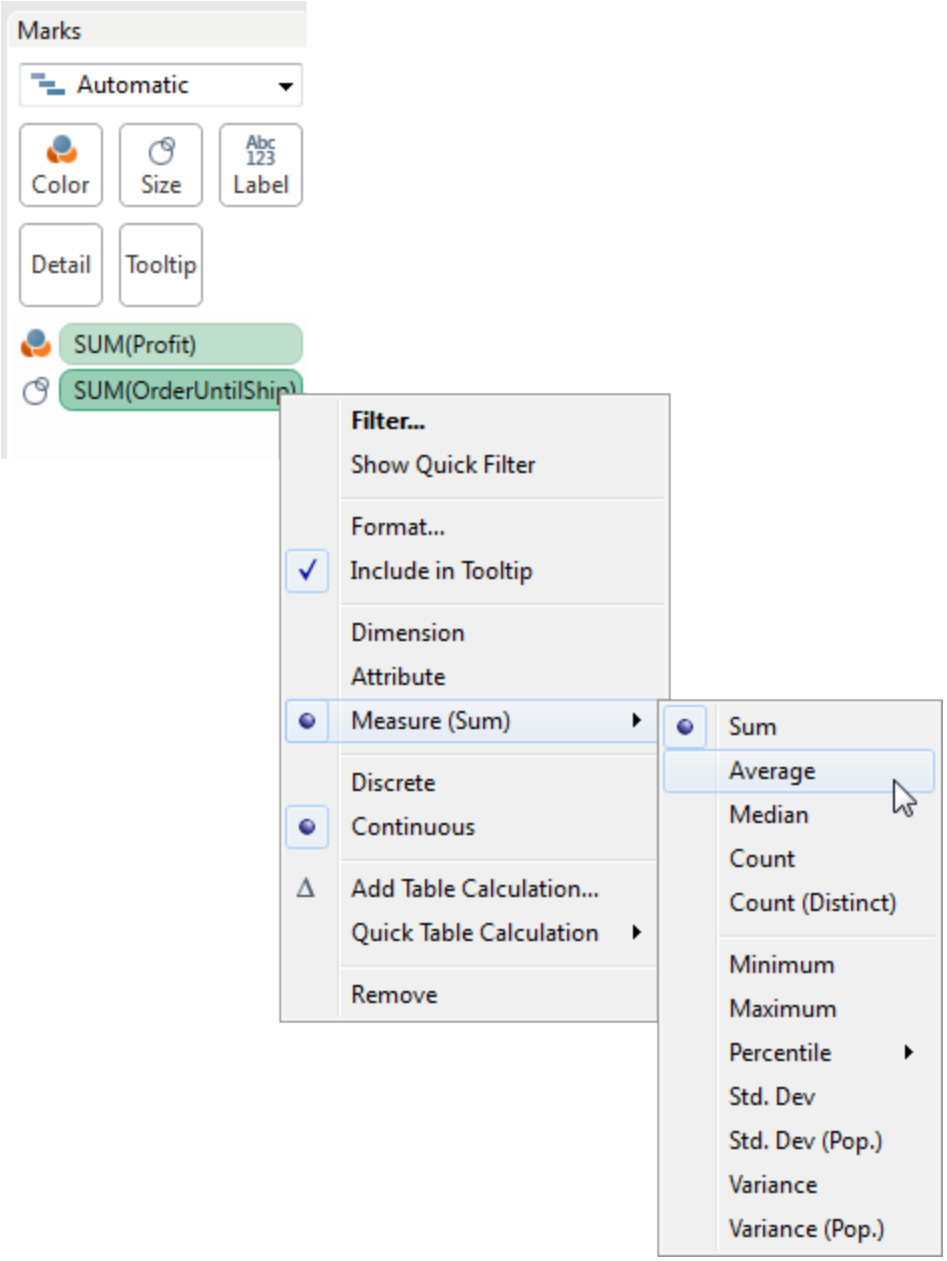

The view is coming along. But there are too many marks squeezed into the view:

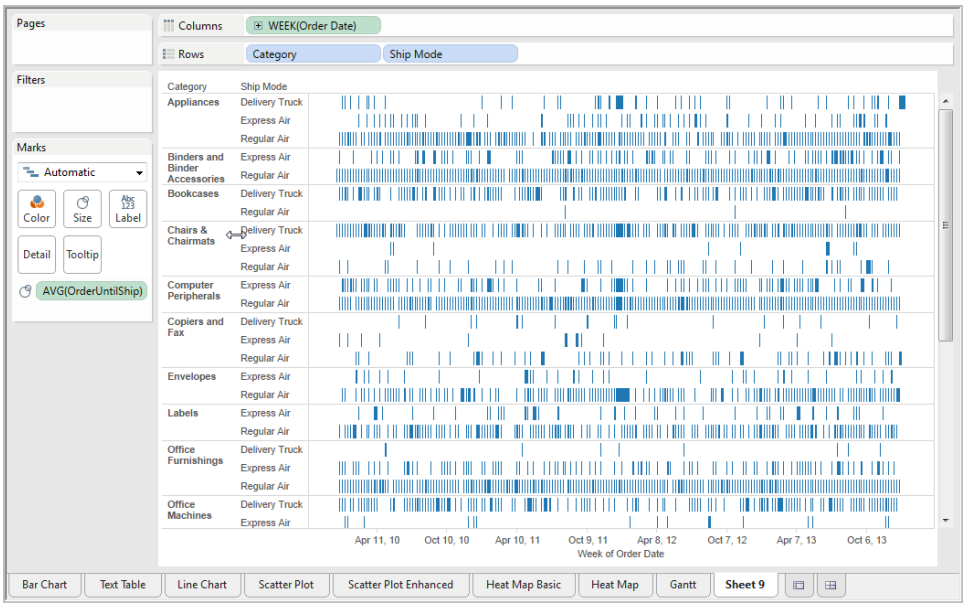

We can make our data more readable by filtering down to a smaller time window.

11. Hold down the Ctrl key and drag the **Week(Order Date)** field from the **Columns** shelf to the **Filter** shelf.

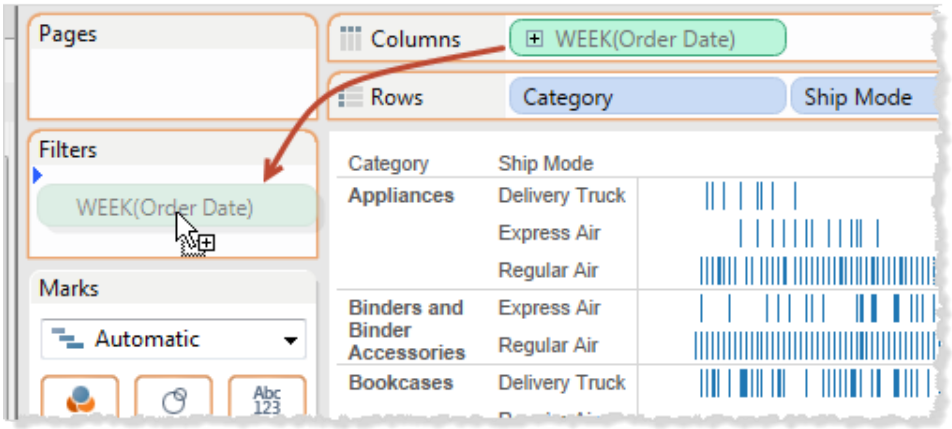

By holding down the Ctrl key, you tell Tableau that you want to copy the field to the new location, with whatever customizations you have added, without removing it from the old location.

A Filter Field dialog box opens. Click **Range of Dates** and then click **Next**, to display this dialog box:

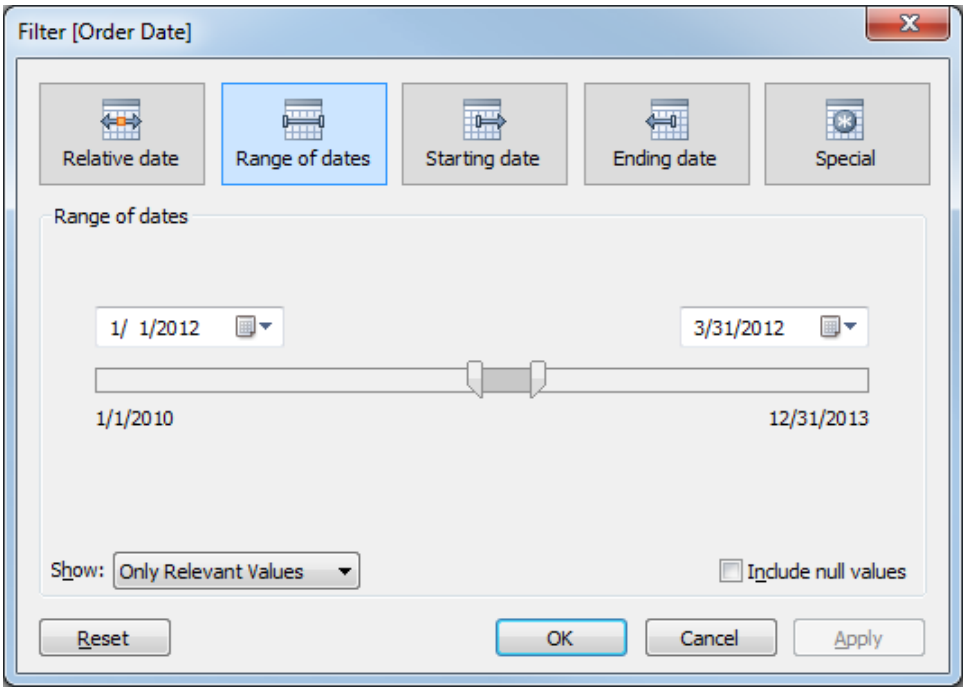

12. Set the range to a three-month time interval, such as 1/1/2012 to 3/31/2012, and the click **OK**.

    It can be difficult to get the exact date using the sliders—it's easier just to type the numbers you want directly into the date boxes.

13. Drag the **Ship Mode** dimension to Color on the Marks card.

    You now have a view that lets you see all sorts of information about the lag between order times and ship times:

Which ship modes are more prone to longer lag times? Do lag times vary by category? Are lag times consistent over time?

If you publish this view to Tableau Server, you can include Quick Filters that let users interact with the view by varying the time window, or filtering out various categories or ship modes. For more information, see **Publishing Workbooks** on page 822.

Building Pie Charts

In Tableau there are **pie charts**, which are visualizations in which data is divided into wedge sectors in a round graph, illustrating numerical proportion, and there are **pie marks**, which are pie charts that you use as a mark type in a visualization, such as a map. For more information about pie marks, see **Pie Mark** on page 254.

Pie charts can be as effective as bar charts for showing proportions. Pie marks are very useful, for example, to demonstrate a percentage of marketing expenses by state.

The first exercise in this topic shows you how to create a view that used a pie chart.

Building a Pie Chart

This exercise walks you through creating a pie chart view that shows how different product categories contribute to total sales.

1. Connect to the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.
2. Drag the **Sales** measure to **Columns**.

   The measure is automatically aggregated as a Sum.

3. Drag the **Category** dimension to **Rows**.

   The default chart type is a bar chart:

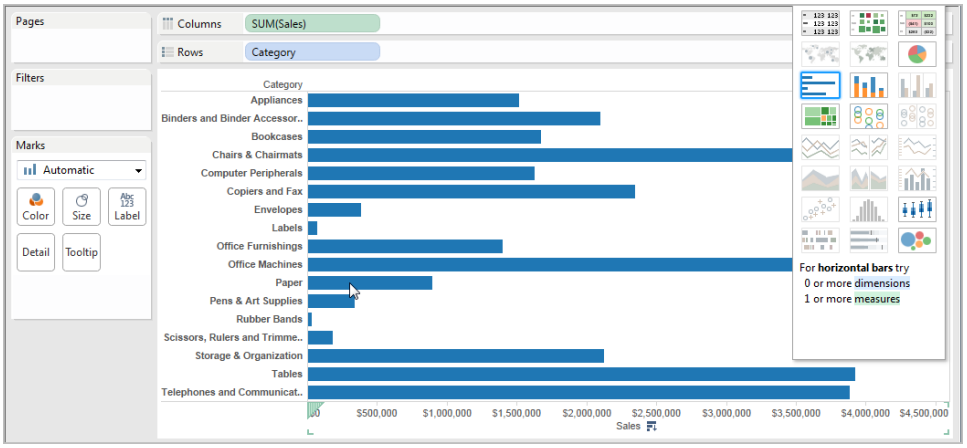

4. Choose the pie chart type from Show Me:

   The result is a rather small pie:

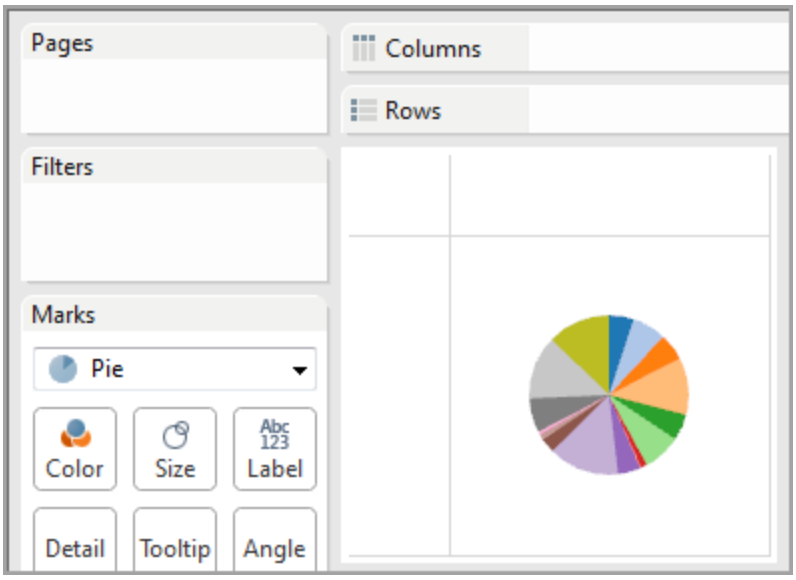

5. To make the chart bigger, hold down Ctrl + Shift and press B several times. This is equivalent to selecting **Format >Cell Size >Bigger**.
6. Press and hold Ctrl as you drag the **Category** dimension from the Data window to **Label** on the Marks card.

   Labels are necessary because your pie chart only makes sense if your users can know what the individual slices represent.

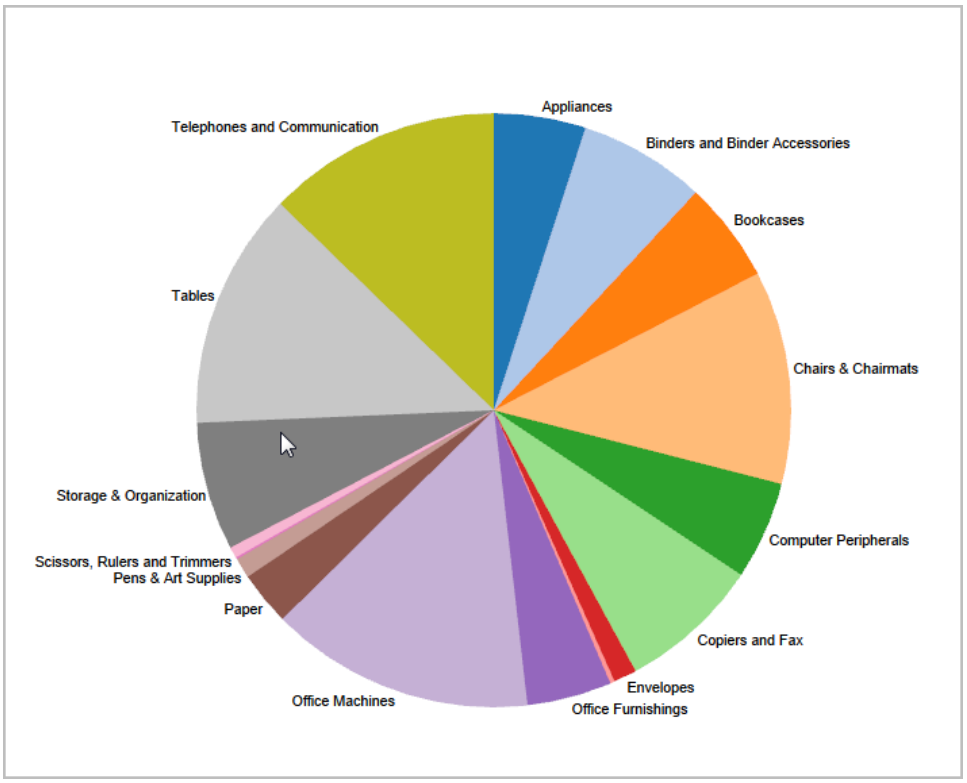

7. Press Ctrl + Shift + B to make sure most of the individual labels are visible.

A basic pie chart isn't all that interesting on its own, but it can be a useful ingredient in a dashboard, especially if you use actions to show data proportions (in the pie chart) for different values in a dimension. See **Actions and Dashboards** on page 558 for details.

Building a Map with Pie Marks

This exercise walks you through creating a map view that shows sales totals for the US states (just the 48 contiguous states). The mark type is Pie, with the mark for each state showing the proportion of total sales for each ship mode (delivery truck, express air, regular air).

1. Open the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.
2. Double-click the **State** dimension.

   Tableau adds **State** to Detail on the Marks card. Because Tableau has automatically assigned the **State** dimension a geographic role, it shows a map when you double-click it.

   For the **Sample - Superstore - English (Extract)** data source, Tableau has also automatically created a hierarchy of geographic roles (mapping items), which includes **Country/Region**, **State**, **City**, and **Postal Code**:

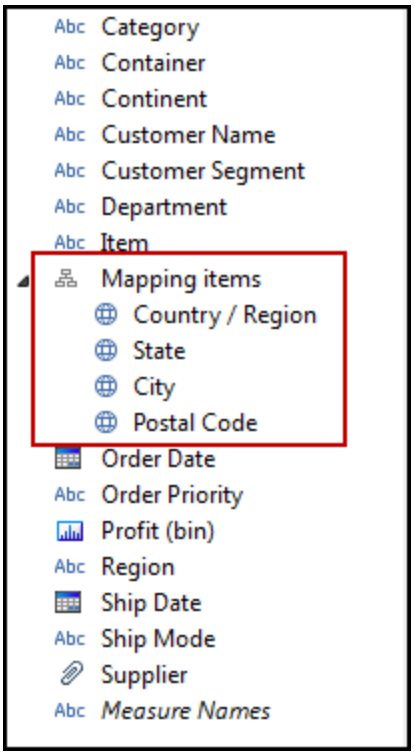

Tableau will look for these kinds of geographical fields in any data source.

When you double-click **State**, Tableau automatically adds **Country/Region** and **State** to Detail. See **Geographic Roles** on page 689 for information on how Tableau interprets geographical data in data sources.

3. To zoom the map in on the continental United States, drag **Country/Region** from the Data window to **Filters**. This opens a Filter dialog box:

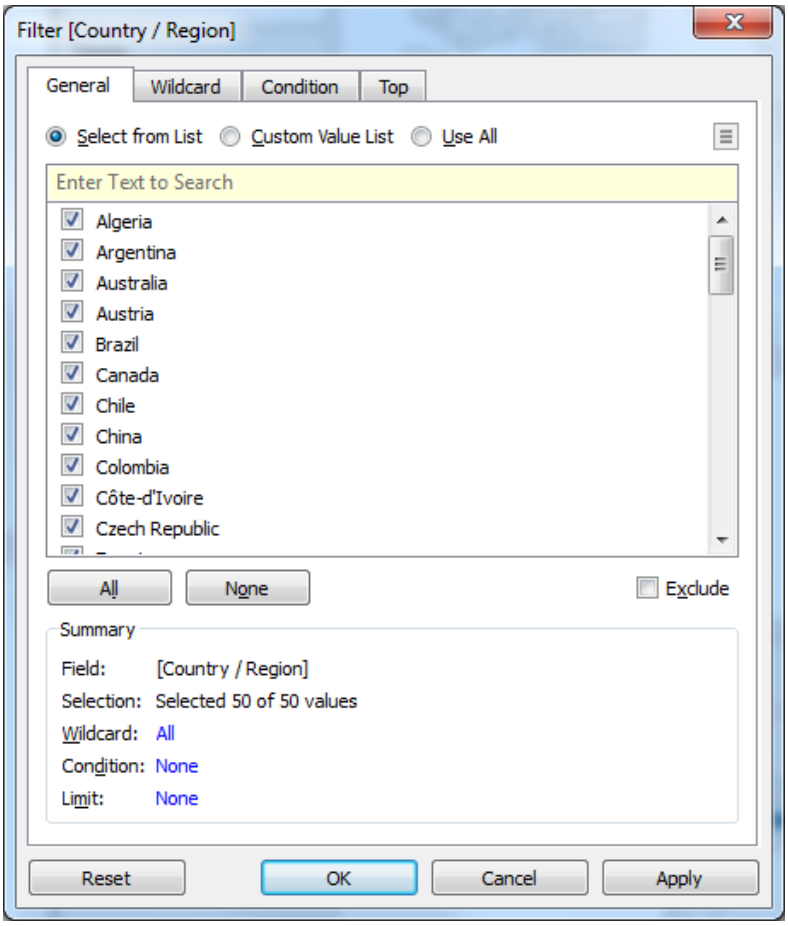

4. Click **None** to clear all selections, scroll down and select **United States of America**, and then click **OK**.

   The map rescales to show the continental United States:

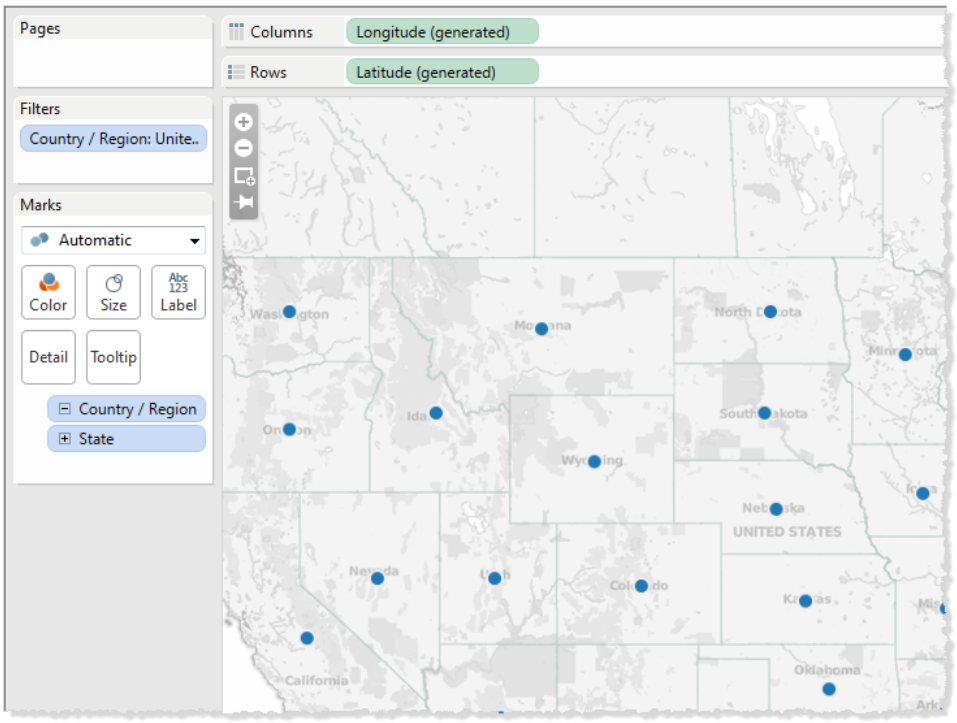

(This is just a partial image—the full image is too large to show here.)

5. Drag the **Sales** measure to **Size** on the Marks card.

   The marks in the view resize according to sales in each state.

6. From the Mark type drop-down list, selecting **Pie**:

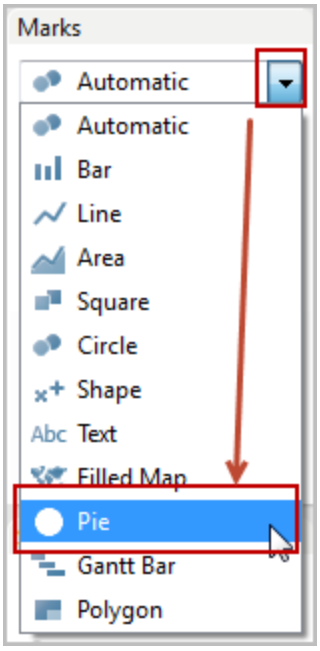

There is no obvious change in the view, but the next step will not work as intended unless you have done this.

7. Drag the **Ship Mode** dimension to **Color** on the Marks card.

   The view is now nearly finished—the only problem is that the marks (that is, the pie marks) are too small.

8. Click **Size** on the Marks card and drag the slider to the right about two-thirds of the way across:

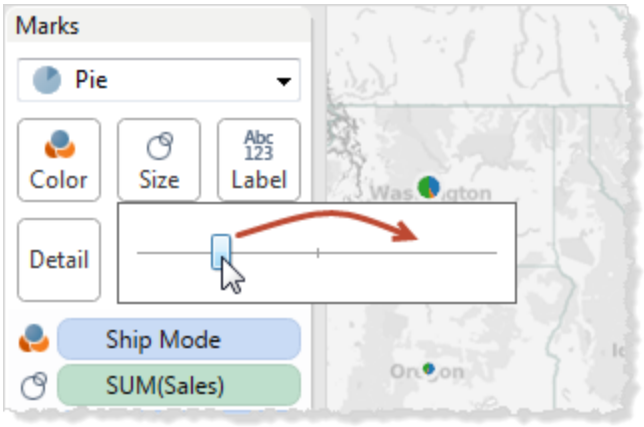

Your view is now complete--your map shows pie chart marks for each state. The size of each pie corresponds with the sales total for that state; the slices of the pie show what proportions of the total were provided by each of the three available ship modes:

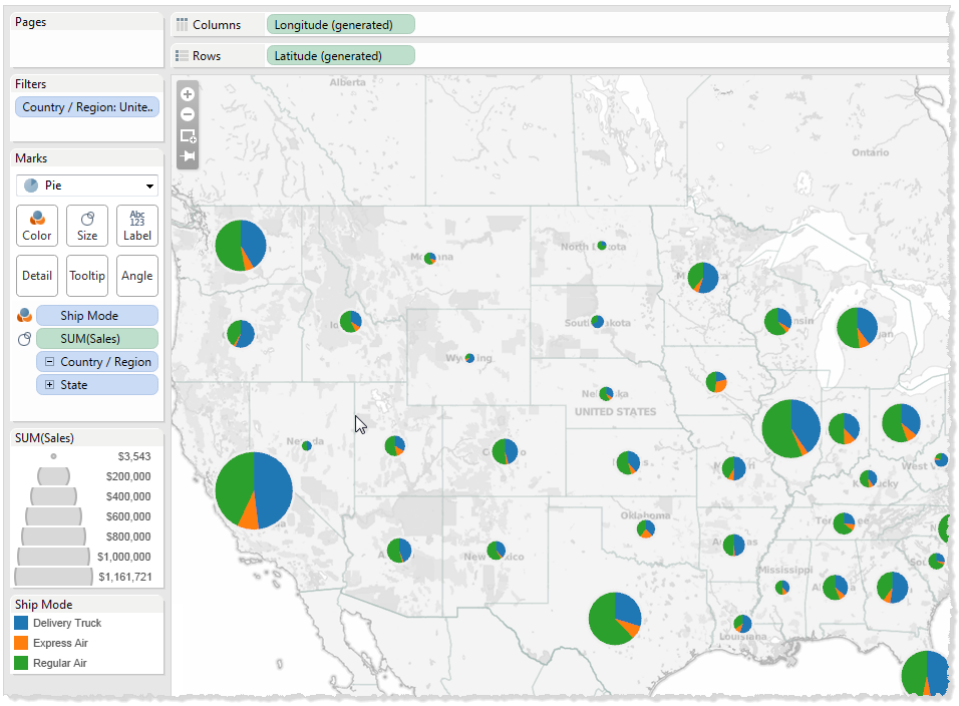

Building Treemaps

Treemaps let you display data in nested rectangles. You use dimensions to define the structure of the treemap, and measures to define the size and/or color of the individual rectangles.Treemaps are a relatively simple data visualization that can provide insight in a visually attractive format.

In Tableau, you create a treemap by first placing one or more dimensions on the **Columns** shelf and one or two measures on the **Rows** shelf. Then choose treemap from Show Me.

The treemap you will create in this exercise shows aggregated sales totals across a range of product categories. You then have two options for enhancing this basic treemap.

1. Connect to the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.
2. Drag the **Category** dimension to **Columns**.

   A horizontal axis is created showing product categories.
3. Drag the **Sales** measure to **Rows**.

   The measure is automatically aggregated as a sum and a vertical axis is created.

   Tableau displays a bar chart—the default chart type when there is a dimension on the **Columns** shelf and a measure on the **Rows** shelf.
4. Choose the treemap chart type from Show Me:

Tableau displays the following treemap:

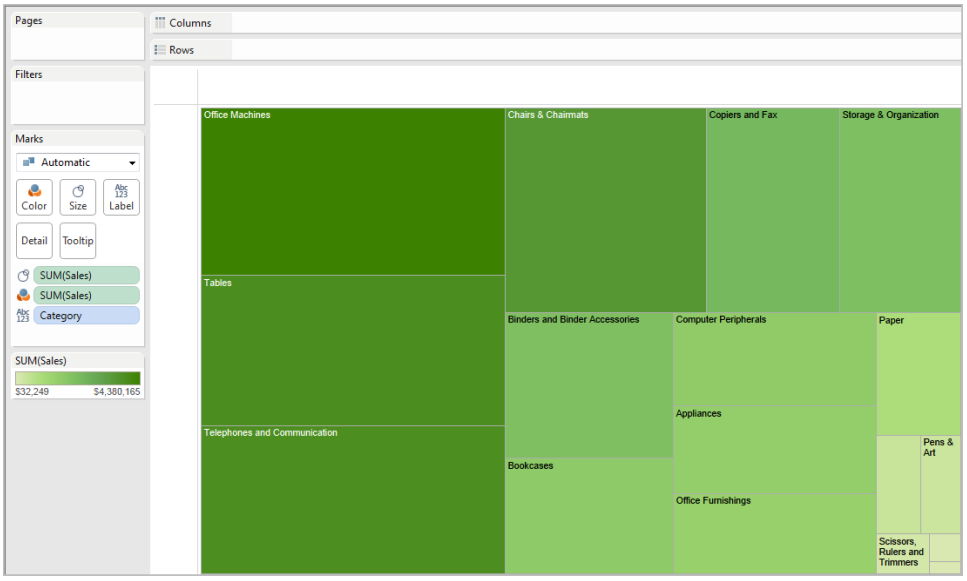

In this treemap, both the size of the rectangles and their color are determined by the value of **Sales**—the greater the sum of sales for each category, the darker and larger its box.

5. Drag the **Ship Mode** dimension to **Color** on the Marks. In the resulting view, **Ship Mode** determines the color of the rectangles—and sorts them into three separate areas accordingly—and **Sales** determines the size of the rectangles:

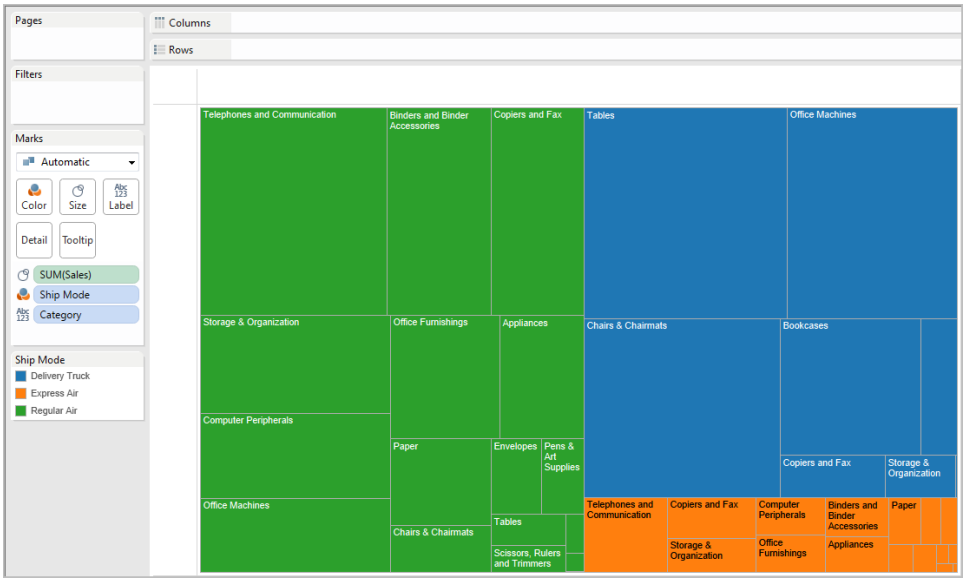

6. To examine another option to modify the treemap, click the **Undo** button to remove **Ship Mode** from view. For more information, see **Undo and Redo** on page 476.

7. Drag the **Profit** measure to **Color** on the Marks card. Now **Profit** determines the color of the rectangles, and **Sales** is determining their size:

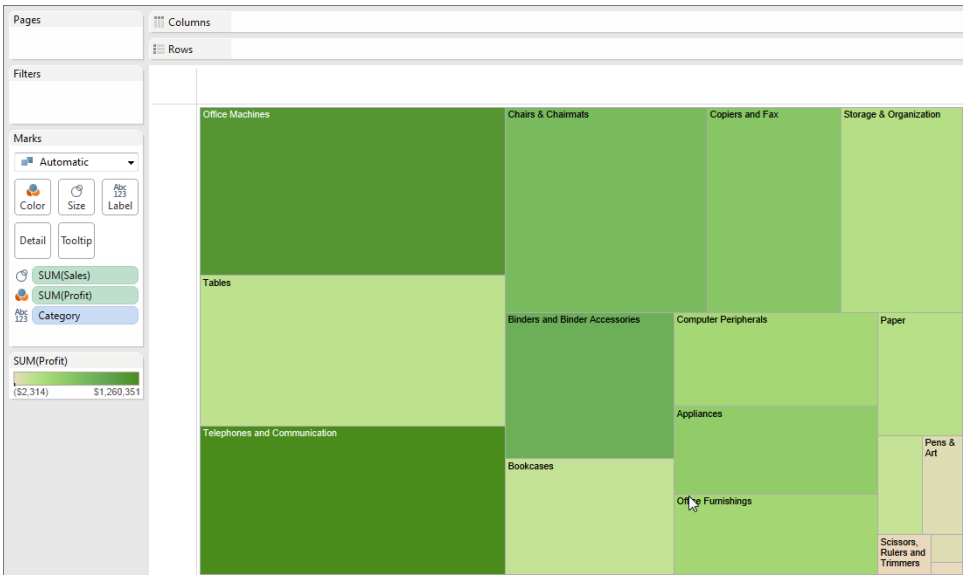

Treemaps cannot accommodate more than two measures—one to control size, the other to control color. Treemaps can accommodate any number of dimensions, but you can only use one dimension to diversify the view, by dragging it to color. Other dimensions can be used only

to multiply the number of rectangles in the view. Dimensions cannot be displayed in a hierarchy.

Building Box Plots

Box plots, also known as box-and-whisker plots, are a type of graph that shows the distribution of values along an axis. Boxes enclose the middle 50% of the data (that is, the middle two quartiles of the distribution). Lines, called whiskers, can be configured to display so as to include all points within 1.5 times the interquartile range (in other words, all points within 1.5 times the width of the adjoining box), or at the maximum extent of the data, as in the following image:

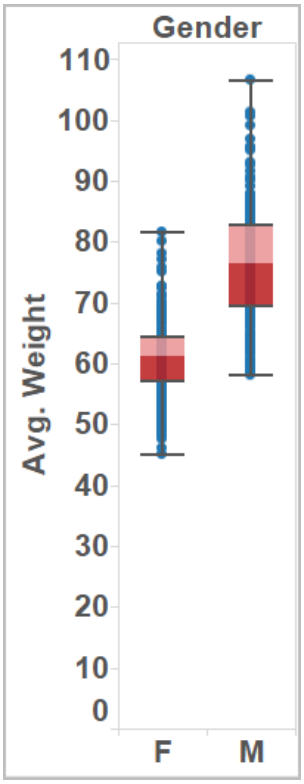

In Tableau, box plots are a chart type that you can select from Show Me, and also a type of reference line that you can add to an axis in a view. For more information about box plots, see **Reference Lines, Bands, and Boxes** on page 451. To add a box plot using Show Me and to configure that box plot, right-click the axis and then choose **Edit Reference Line, Band, or Box**.

The following exercise walks you through creating a set of box plots that show shipping costs on a per-customer basis, by continent and customer segment.

1. Connect to the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.

2. Drag the **Continent** dimension to **Columns**.

   The measure is automatically aggregated as a sum and row headers are displayed, identifying six continents.

3. Drag the **Shipping Cost** measure to **Rows**.

   Tableau creates a vertical axis.

   Tableau displays a bar chart—the default chart type when there is a dimension on the **Columns** shelf and a measure on the **Rows** shelf.

4. Drag the **Customer Segment** dimension to **Columns**, and drop it to the right of **Continent**.

   Now you have a two-level hierarchy of dimensions from left to right in the view, with **Customer Segment** nested within **Continent**.

5. Choose the box-and-whisker plot chart type from Show Me:

   Tableau displays the following box plot:

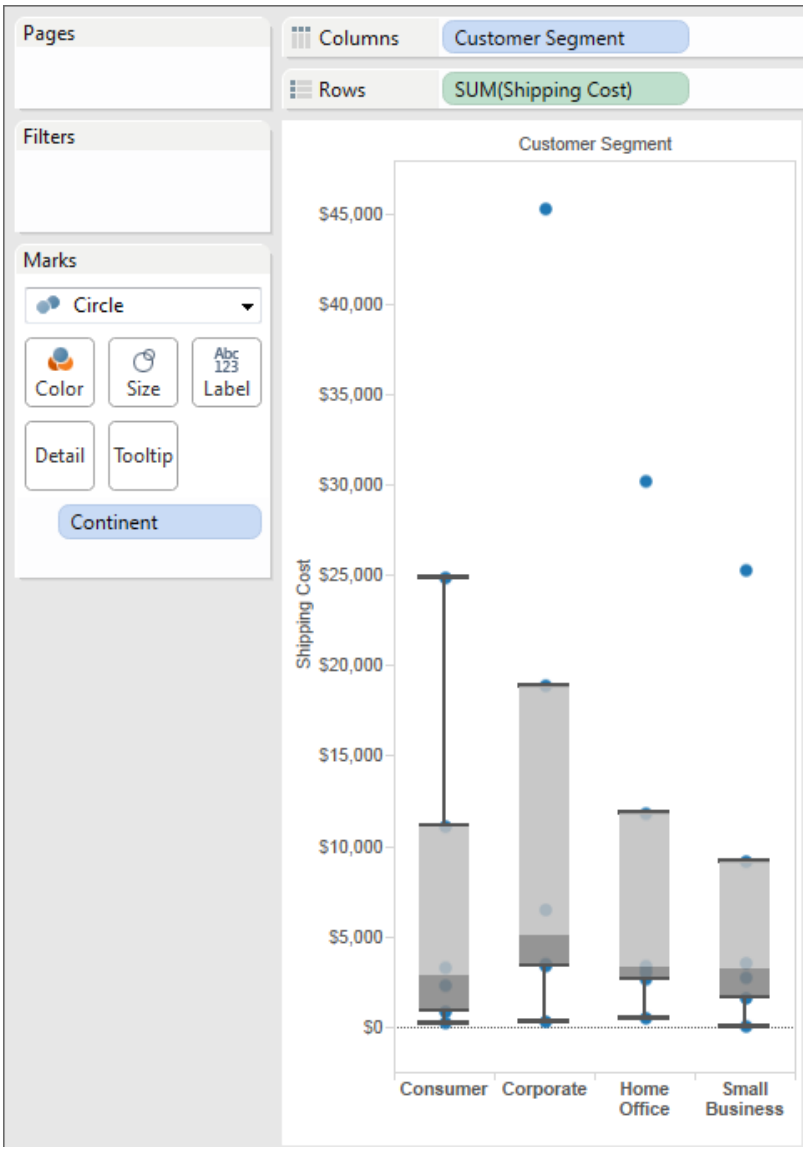

Notice that there are only a few marks in each box plot. Also notice that Tableau has reassigned **Continent** from the **Columns** shelf to the Marks card. When you changed the chart type to a box plot, Tableau needed to determine what the individual marks in the plot should represent. It decided that the marks should represent continents. This was a reasonable conclusion, but it is not what we wanted.

6. Drag **Continent** from the Marks card back to **Columns**.

   This is what the view now looks like:

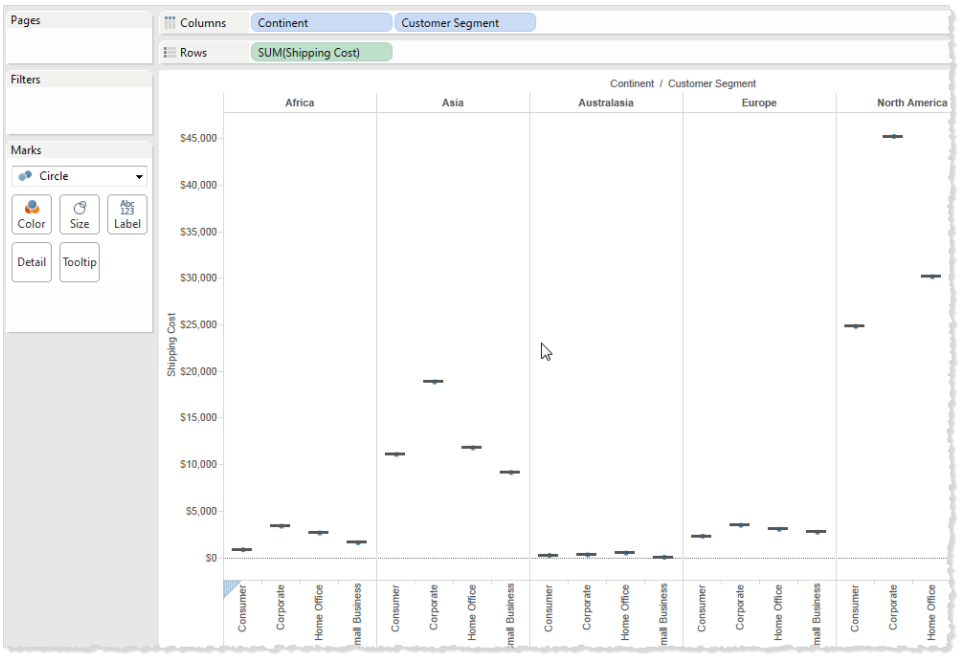

Those horizontal lines are flattened box plots, which is what happens when box plots are based on a single mark.

Box plots are intended to show a distribution of data, and that can be difficult when data is aggregated, as in the current view.

7. To disaggregate data, select **Analysis** > **Aggregate Measures**. This command is a toggle, and because data is aggregated by default in Tableau, the first time you choose this command it has the effect of disaggregating the data (that is, it removes the check mark from this menu item). For information on disaggregating data, see **Disaggregating Data** on page 571.

   Now, instead of having a single mark for each column in the view, you have a range of marks, one for each row (that is, each customer transaction) in your data source:

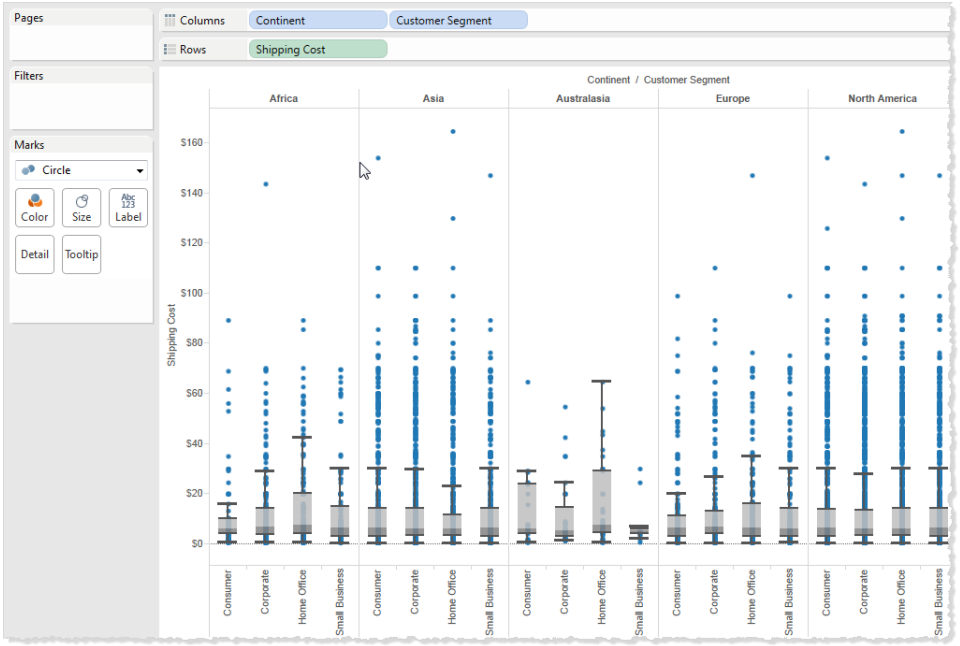

The view is now showing us the information we want to see. The remaining steps have to do with making the view more readable and more attractive.

8. Click the toolbar button for swapping axes:

The box plots now lay left-to-right, and we are able to see a lot more information in a more compressed space:

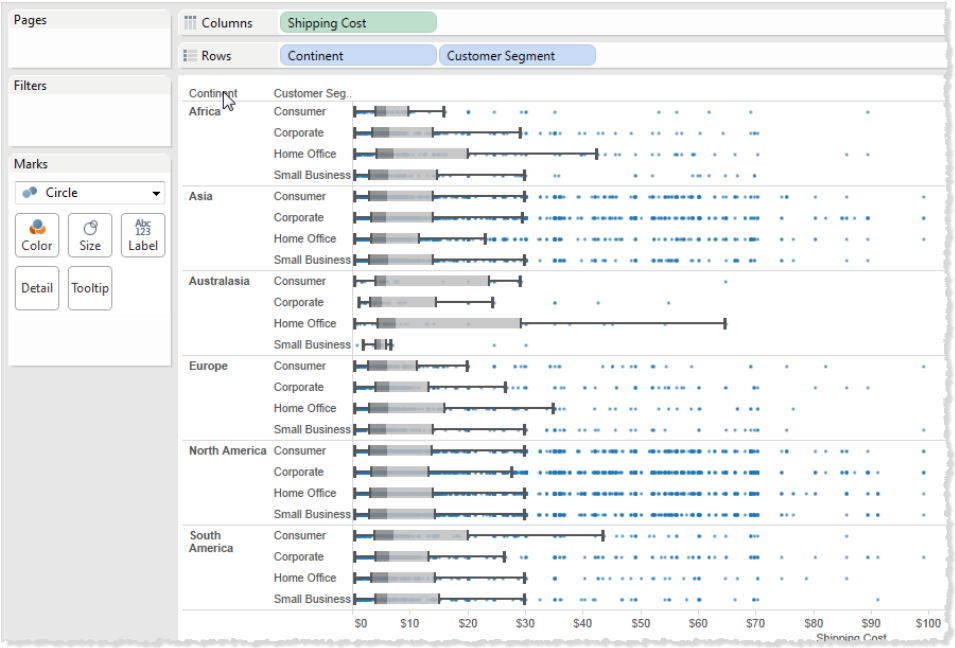

9. Right-click the bottom axis and choose **Edit Reference Line, Band, or Box**. The following dialog box opens:

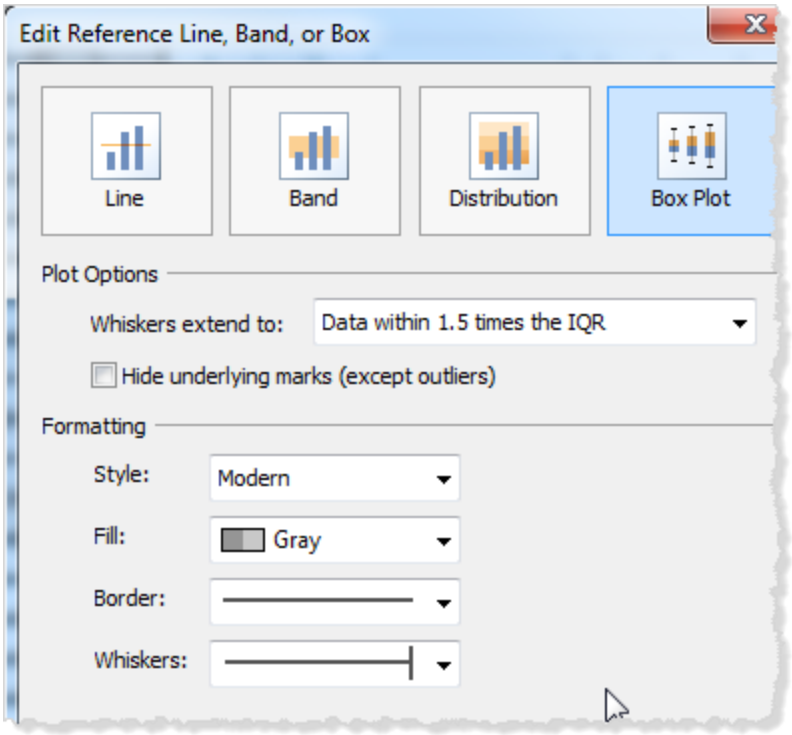

10. In the **Fill** drop-down list, select an interesting color scheme. For more on these options, see **Adding Box Plots** on page 468.

Now your view is complete:

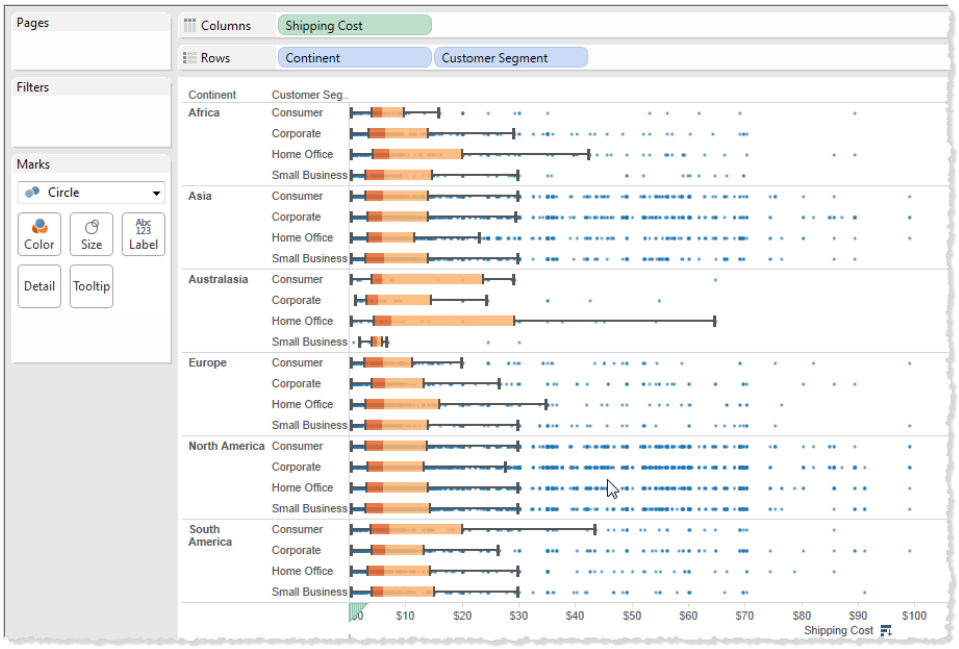

You can see from the density of the marks that sales were greatest in Asia and North America—considering sorting to list the continents in order of total sales revenue. You can also see that the inter-quartile range (from the 25th percentile to the 75th percentile) for shipping costs typically tops out around $30, with a few interesting outliers.

Building Packed Bubble Charts

Packed bubble charts let you display data in a cluster of circles. You use dimensions to define the individual bubbles, and measures to define the size and/or color of the individual circles. Packed bubble charts are a relatively simple data visualization that can provide insight in a visually attractive format.

In Tableau, you create a packed bubble chart by first placing one or more dimensions on the **Columns** shelf and one or two measures on the **Rows** shelf. Then choose packed bubbles from Show Me.

The following exercise walks your through creating a basic packed bubble chart that shows sales and profit information for different product categories.

1. Connect to the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.
2. Drag the **Department** dimension to **Columns**.

   A horizontal axis is created showing product departments.
3. Drag the **Sales** measure to **Rows**.

   The measure is automatically aggregated as a sum and a vertical axis is created.

Tableau displays a bar chart—the default chart type when there is a dimension on the **Columns** shelf and a measure on the **Rows** shelf.

4. Choose the packed bubbles chart type from Show Me:

Tableau displays the following packed bubble chart:

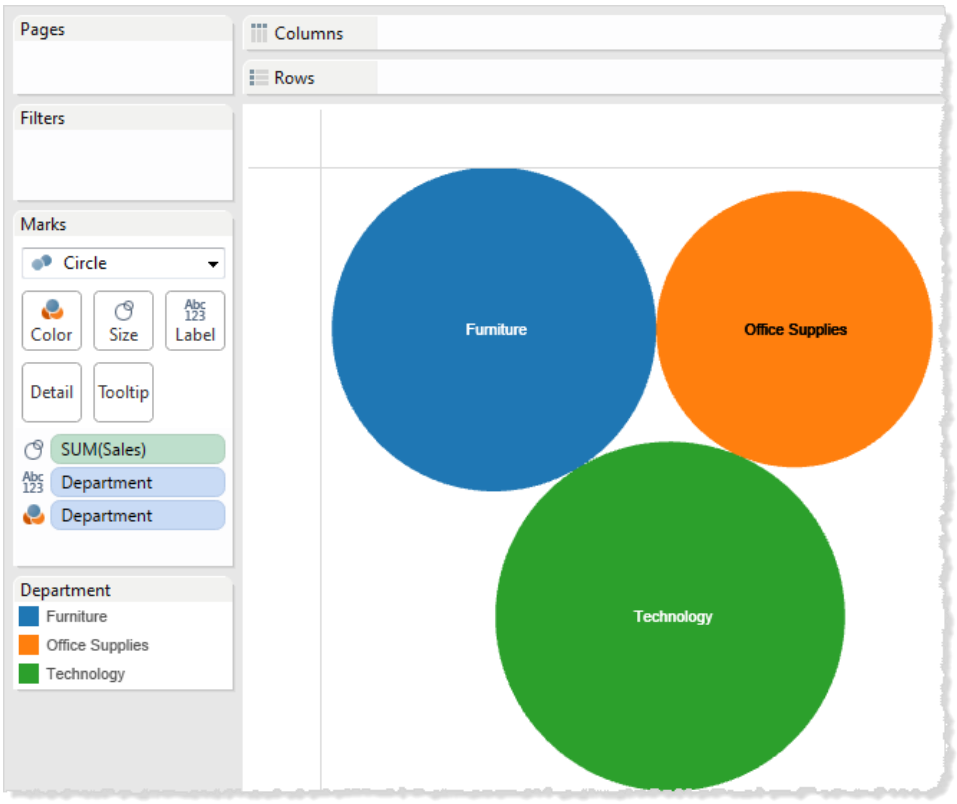

5. Drag **Region** to **Detail** on the Marks card to get more bubbles in the view. There isn't enough space on this page to make the image big enough to show labels for all the circles. With a bigger image, more labels would display, but it's often the case that some bubbles will be too small even when the view is fairly large. When this happens, hover over a bubble to see the tooltip.

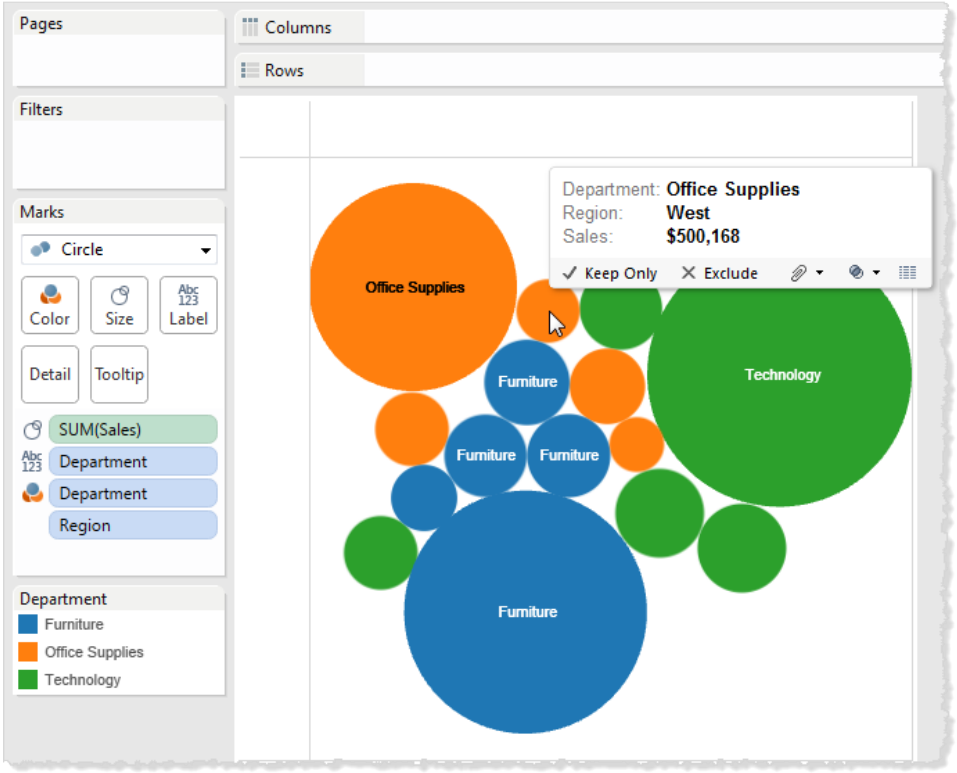

At this point you can continue adding dimensions to the view to multiply the number of circles—packed bubble charts that show hundreds of bubbles can be interesting for some purposes. Or you can drag a second measure to color, to add another layer of information to the view.

6. Drag **Profit** to **Color** on the Marks card:

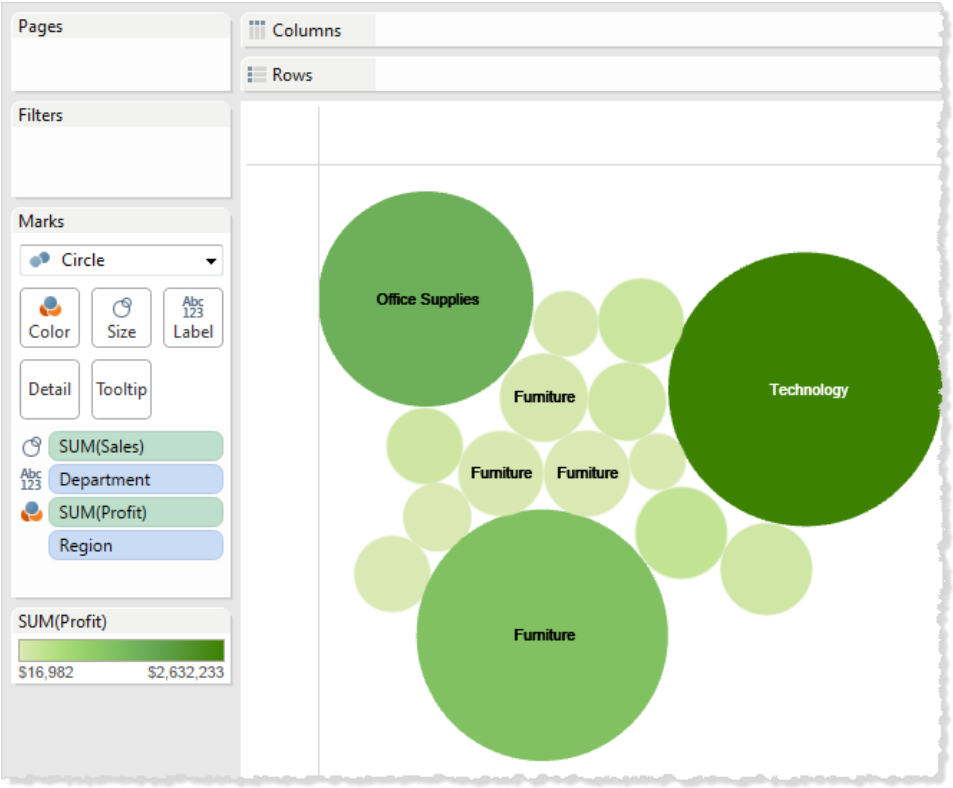

7. Drag **Region** to **Label** on the Marks card so that users won't be confused by seeing identical labels on different circles:

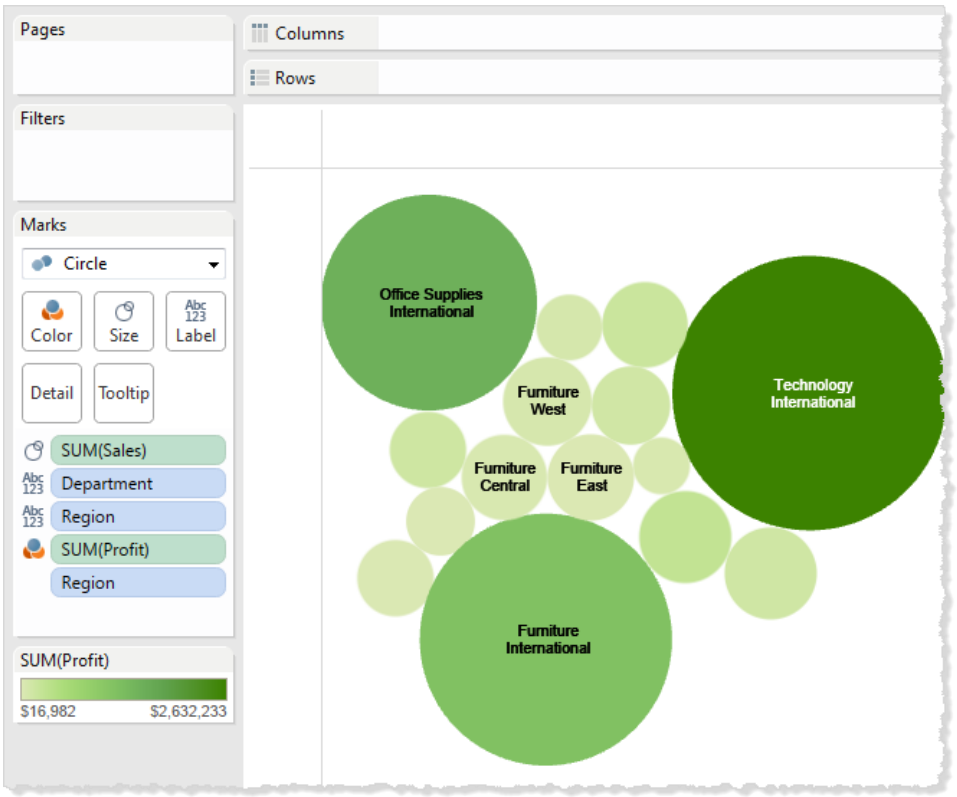

The size of the bubbles shows the sales for different combinations of region and department; color shows profit (the darker the green, the greater the profit).

For information about formatting mark labels, see **Showing and Hiding Mark Labels for the Worksheet** on page 781.

There is plenty more you could do to develop this view. You could edit the colors for **Profit** to use a diverging palette (so that negative profit shows in a different color), or you could create a calculated field that showed profit divided by sales (that is, profit margin) and drop that on **Color** instead of absolute profit. For more information, see **Color** on page 257.

Advanced Analysis

Now that you understand the basics of building views in Tableau, become an advanced user by learning how to create custom calculations, use the built in statistics tools, leverage dynamic parameters, map data, and more.

Actions

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 web page showing census data for the neighborhood.

> **Note:** Actions behave differently in workbooks that use cube data sources. Cube data sources do not accept actions from relational or other cube data sources. For example, suppose you have a workbook that contains a view that uses a MySQL data source, a second view that uses a cube data source A, and a third view that uses a cube data source B. Actions in a view that use the MySQL data source will not affect the views that use the cube data sources. However, actions in the view that use the cube A data source can affect the view that uses the MySQL data source. The view that uses the cube B data source will not be affected in this case.

There are three kinds of actions in Tableau: Filter, Highlight, and URL 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.

1. On a worksheet, select **Worksheet** > **Actions**. From a dashboard, select **Dashboard** > **Actions**.
2. In the Actions dialog box, click **Add Action** and then select **Filter**.

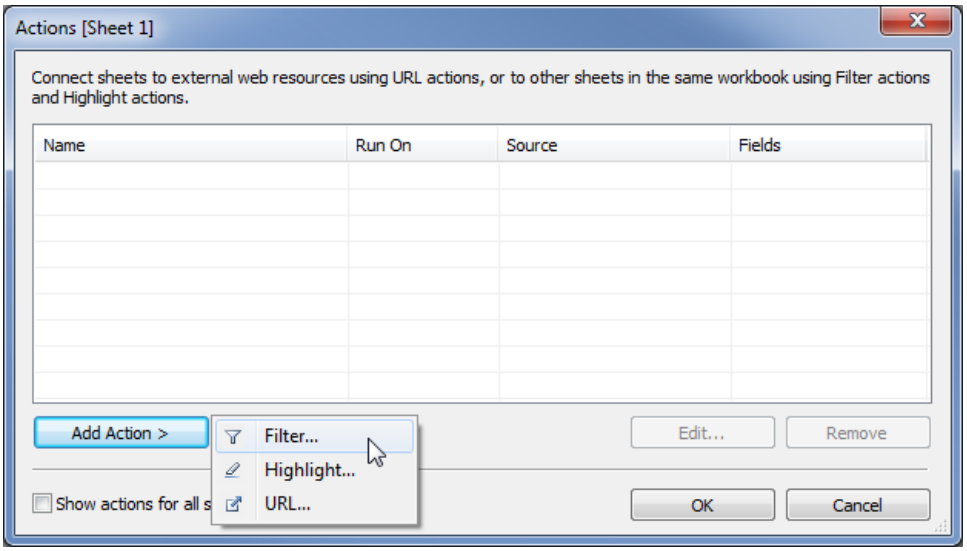

3. In the subsequent dialog box specify a name for the Action.

   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.

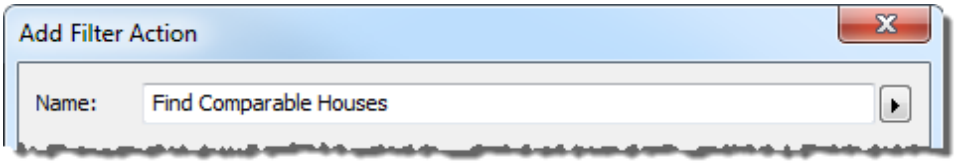

4. Use the drop-down list to select a source sheet or data source. When you select a data source or dashboard sheet you can further refine by selecting the individual sheets you want to launch the action from.

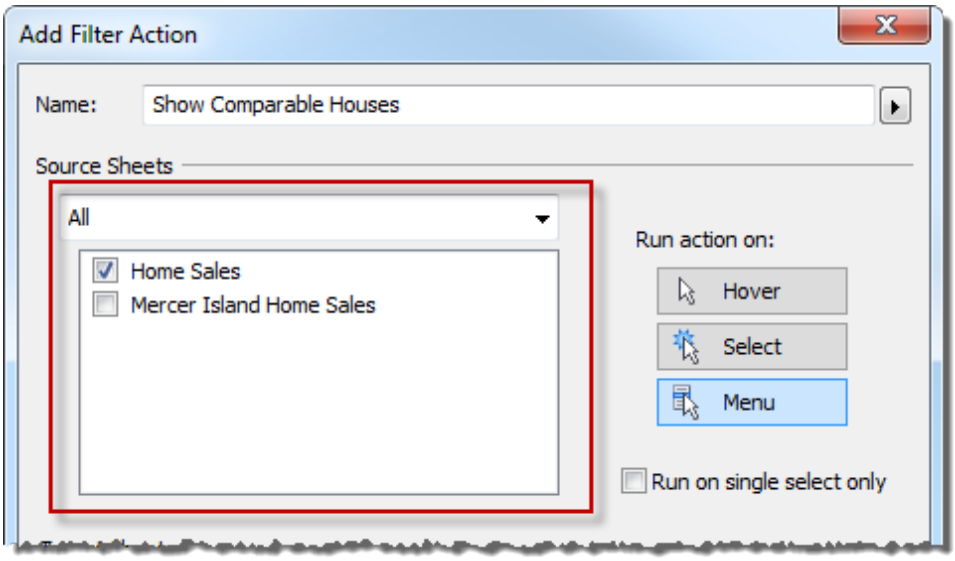

5. Then select how you want to launch the action. Select one of the following options:
   - **Hover** - rest the pointer over a mark in the view to run the action. This option works well for highlight and filter actions within a dashboard.
   - **Select** - click on a mark in the view to run the action. 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.

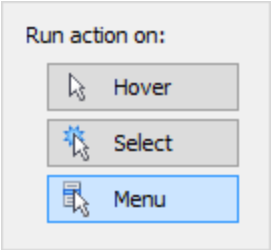

6. Use the second drop-down list to select a target sheet. When you select a dashboard sheet you can further refine the target by selecting one or more sheets within the dashboard.

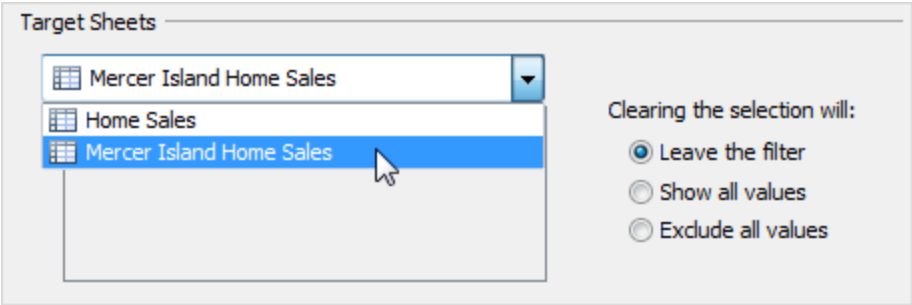

7. Specify what to do when the select is cleared in the view. You can select from the following options:
    - **Leave the filter** - leaves the filter on the target sheets. The target views in the dashboard will show the filtered results.
    - **Show all values** - changes the filter to include all 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**.

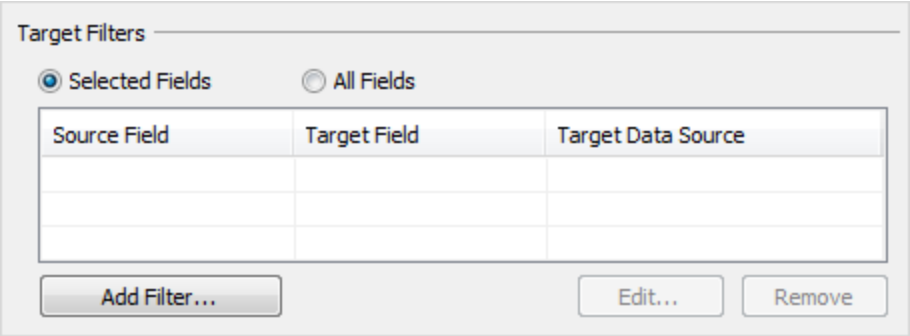

10. In the Add Filter dialog box, select a source and target data sources and fields. When you run the action from a specific mark on the source sheet, a filter is added to the target sheet that only includes values for the target field that match the values of the source field. In the comparable houses 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 three bedrooms, the destination worksheet will only show houses that also have three bedrooms.

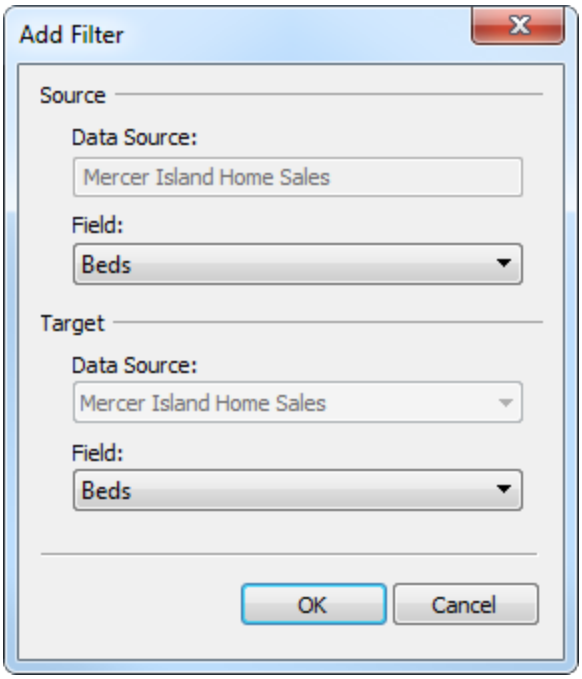

11. When finished, click **OK** three times to close all open dialog boxes and return to the view.

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. In Tableau, multidimensional data sources are supported only in Windows.

> 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.

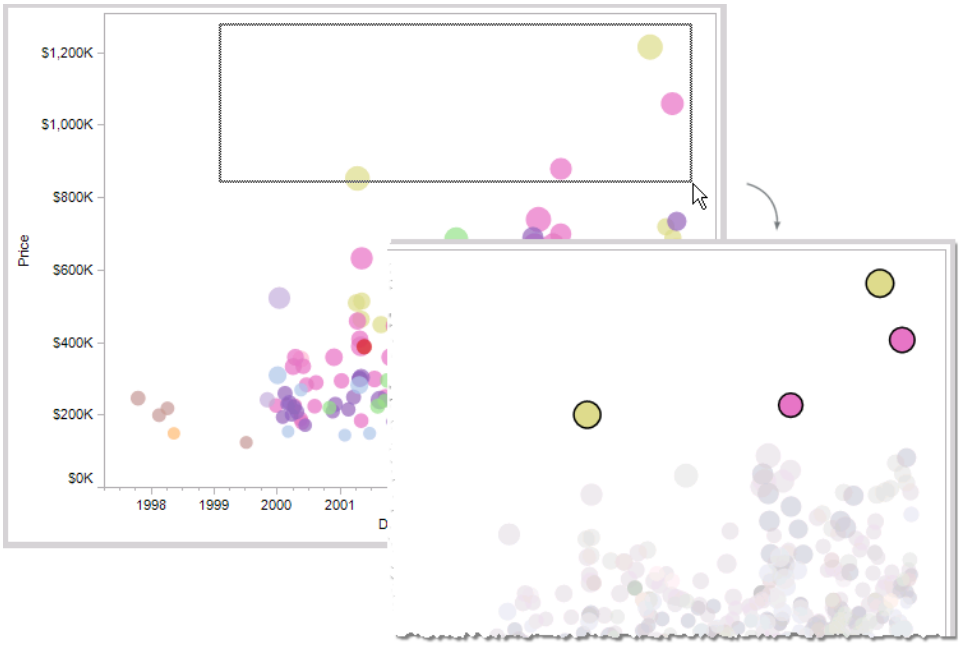

Color Legend Highlighting

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.

Normal Color Legend

Color Legend Highlighting Enabled

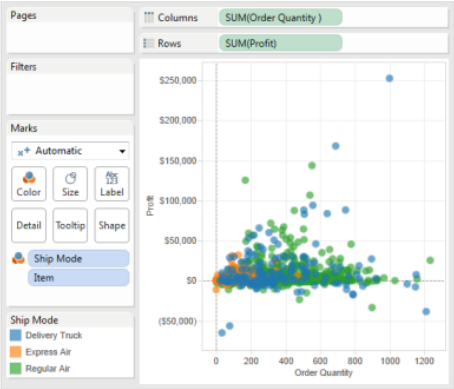

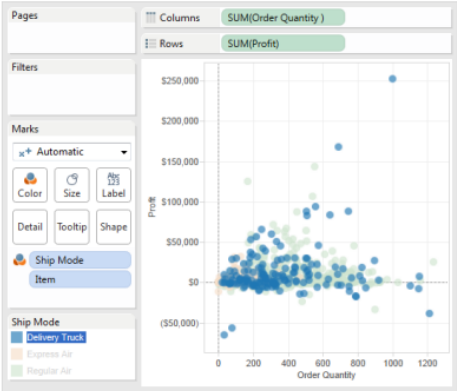

You can easily switch between legend highlighting and normal modes using the color legend card menu. Then, if you like how a view is highlighted, you can assign the highlight colors to the color palette. The old colors are replaced with the highlight colors.

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.

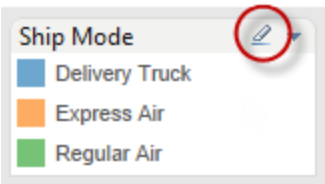

2. Select an item in the color legend.

   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.

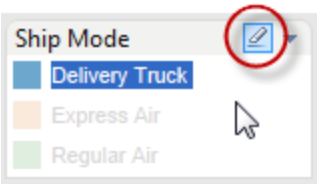

When you turn color legend highlighting off the action is removed from the Actions dialog box.

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.

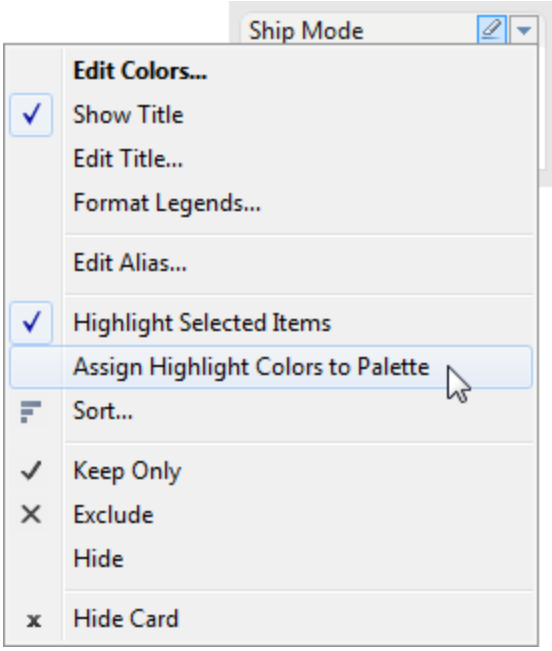

Highlight Toolbar Button

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.

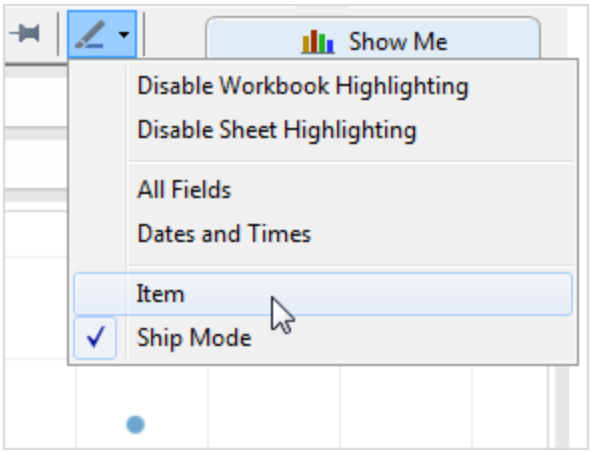

For example, the view below shows sales vs. profit by region. When a mark is selected, all other marks from that region that were shipped using the selected ship mode are highlighted.

In this case you can quickly see all products from the Western region that were shipped via Delivery Truck.

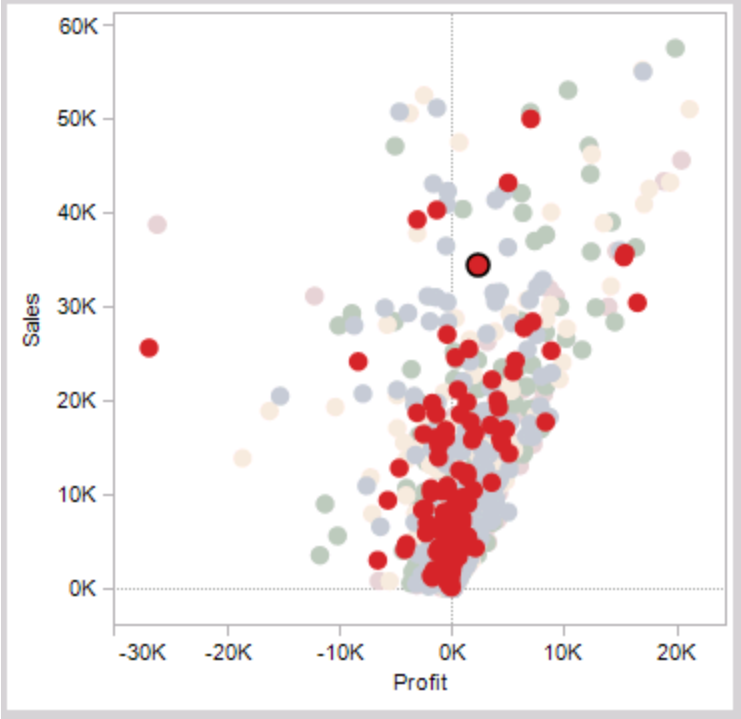

The toolbar menu also lets you highlight on **All Fields** or **Dates and Times**. **All Fields** will consider all fields when determining matching records; **Dates and Times** considers all date and time fields.

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.

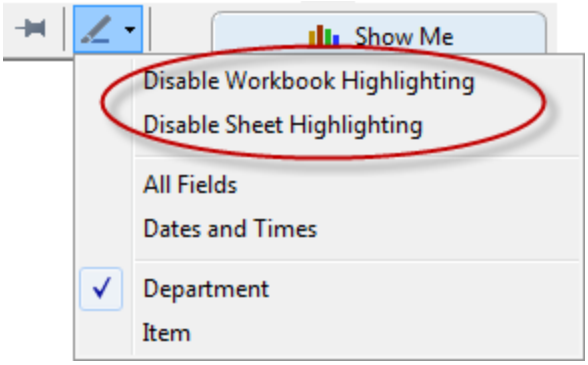

Creating Advanced Highlight Actions

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. On a worksheet select **Worksheet > Actions**. From a dashboard, select **Dashboard > Actions**.
2. In the Actions dialog box click the **Add Action** button and then select **Highlight**.

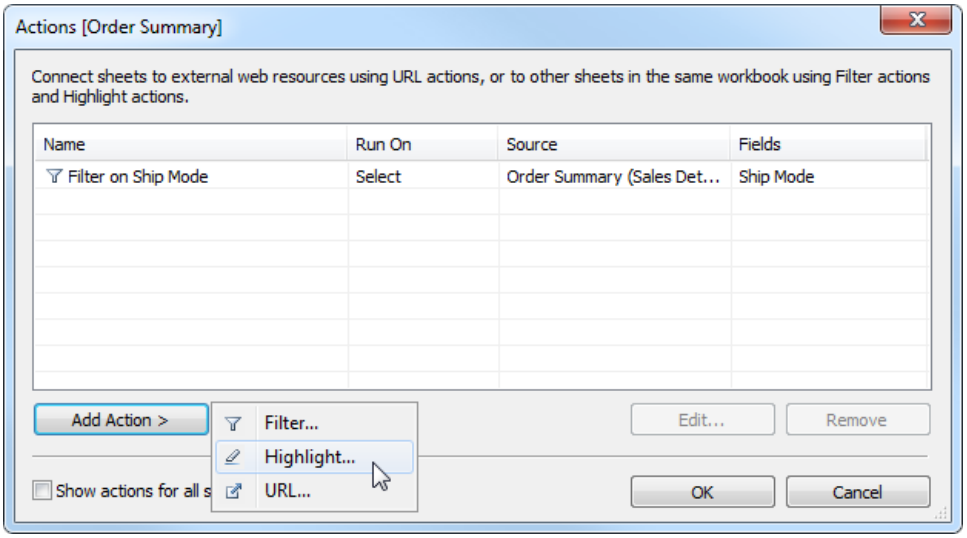

3. Give the action a name that will identify it in the Actions dialog. Try to make it descriptive. For example, *Highlight Products Shipped by Delivery Truck*. You can use variables in the name that will be filled in based on the values of the selected field.

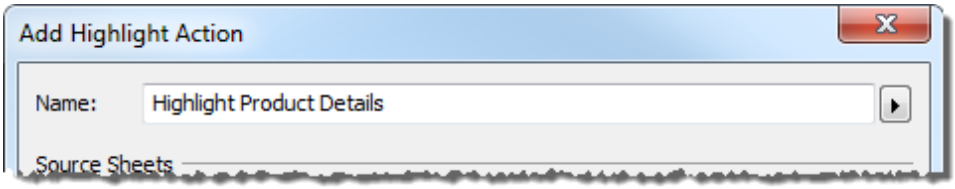

4. Use the drop-down list to select the Source sheet or data source. If you select a data source or a dashboard sheet you can further select individual sheets within them.

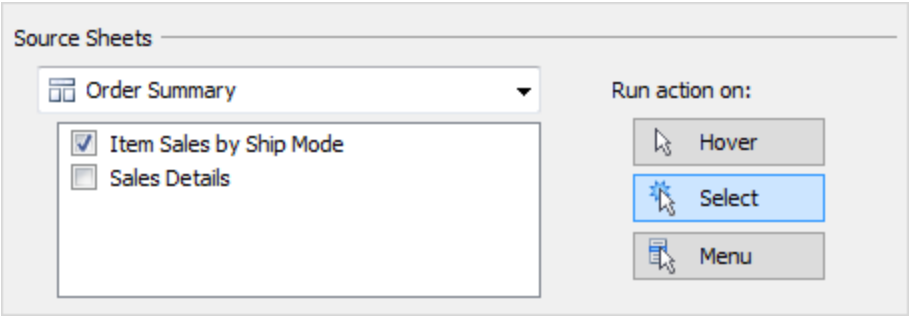

5. Select how you want to launch the action. You can select from the following options:

- **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.

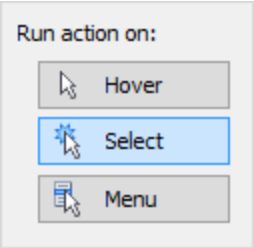

6. Select a Target sheet. If you select a dashboard you can further select individual sheets within the dashboard.

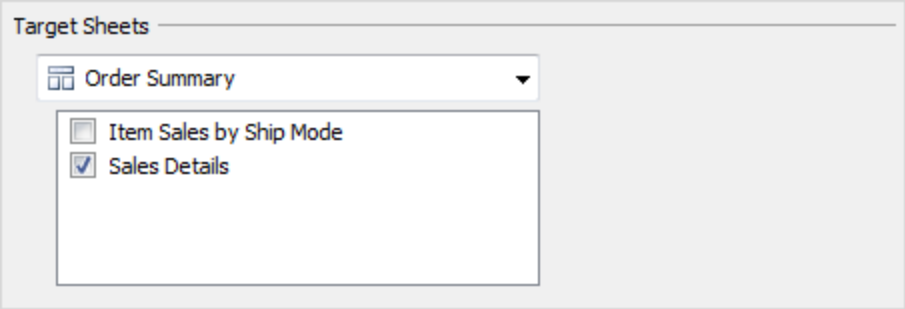

7. Select the fields you want to use for highlighting. Select from the following options:
   - **Selected Fields** - marks in the target sheet are highlighted based on select fields. For example, highlighting using the Ship Mode field will result in an action that highlights all marks in the target sheet 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.

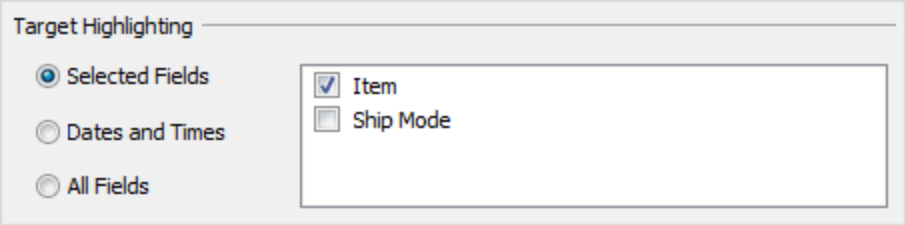

8. When finished, click **OK** twice to close the dialog boxes and return to the view.

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. On a worksheet, select **Worksheet > Actions**. From a dashboard, select **Dashboard >Actions**.
2. In the Actions dialog box, click **Add Action** and then select **URL**.

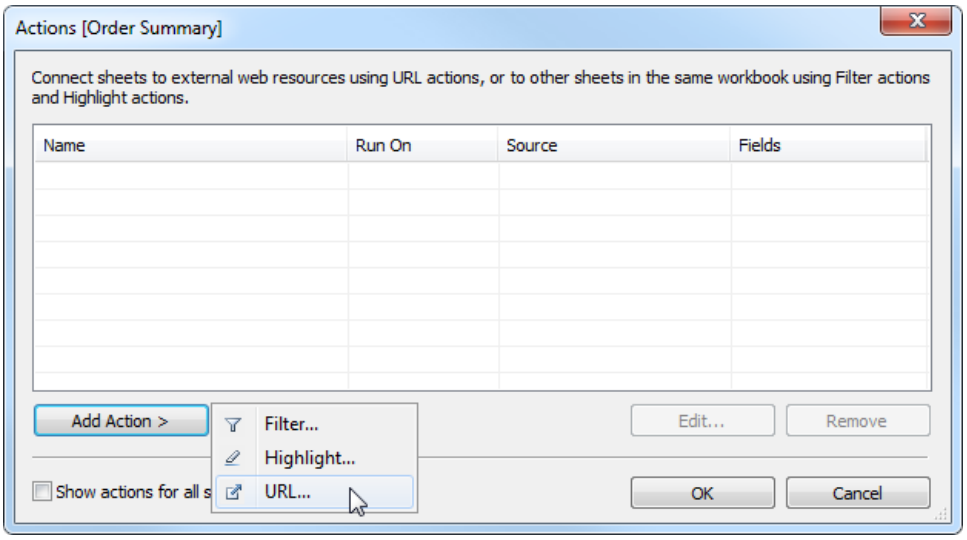

3. In the subsequent dialog box, specify a name for the link.

   Make the name descriptive of the action. If you choose to run the action using the menu, the name you specify here becomes the option that shows on the menu. For example, when linking to more product details, the name could be “Show More Details.” You can use variables in the name, which will be filled in based on the values of the selected field.

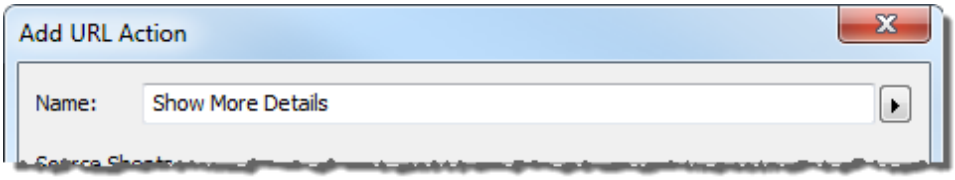

4. Use the drop-down list to select a source sheet or data source. If you select a data source or dashboard you can select individual sheets within it.

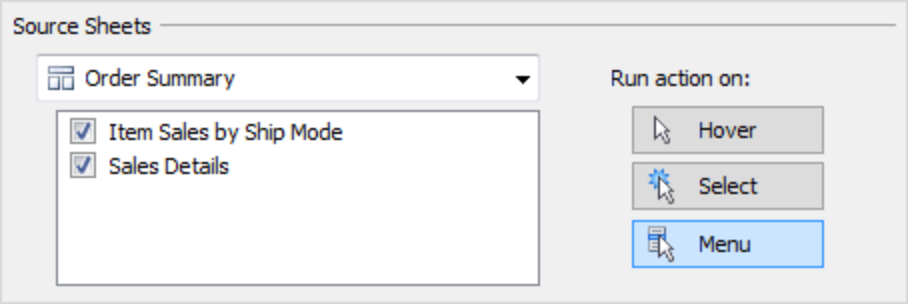

5. Select the fields you want to use for highlighting. Select from the following options:
   - **Hover** - rest the pointer over a mark in the view to run the action. This option works well for highlight and filter actions within a dashboard.
   - **Select** - click on a mark in the view to run the action. 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.

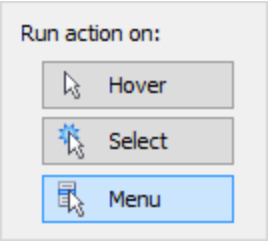

6. Specify the URL. You can use any URL that your browser can recognize including web pages, FTP resources, and files.

   Just as you can use variables in the name of the URL, you can also use field values and filter values as parameters in the URL. That means that you can send information about each selected mark or filter setting to a given website.

   In the URL you specify, include the appropriate prefix to ensure that the resulting hyperlink directs to the correct location. For example, if your URL links to a web page, include the `http://` prefix. If you publish workbooks to a Tableau Server, and your URL includes a protocol other than `http`, `https`, `gopher`, `news`, `ftp`, or `mailto`, you will need a server administrator to add the protocol to the server's whitelist. For more information, see tabadmin set options in the Tableau Server help.

   When using a URL action to point to an external file (rather than a web page), use the full universal naming convention (UNC) path for the URL action. A UNC path is the full path of a resource or file that is stored on the machine, which includes

the full computer name, drive, path to the file, and file name. For example, for a text file D:\myfile.txt, specify the following UNC path: \\workstation1\d$\myfile.txt, where "workstation1" is the computer name in your company domain.

7. (Optional) Select any 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.
   - **Allow Multiple Values** - 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 allow multiple values, you must also define the item delimiter, which is the character that separates each item in the list (for example, a comma). You must also define the Delimiter Escape, which is used if the delimiter character is

used in a data value.

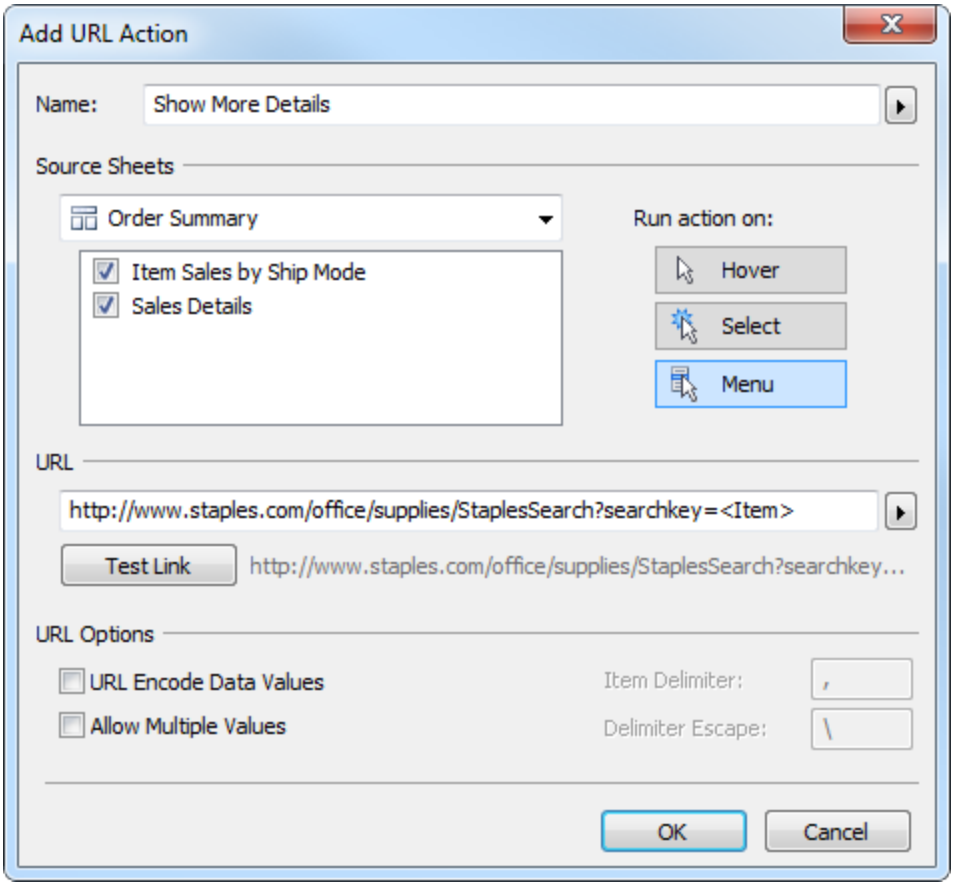

8. When finished, click **OK** twice to close the dialog boxes and return to the view.

URL actions can also point to a web page object in a dashboard. Refer to **Actions and Dashboards** on the next page 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.

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.

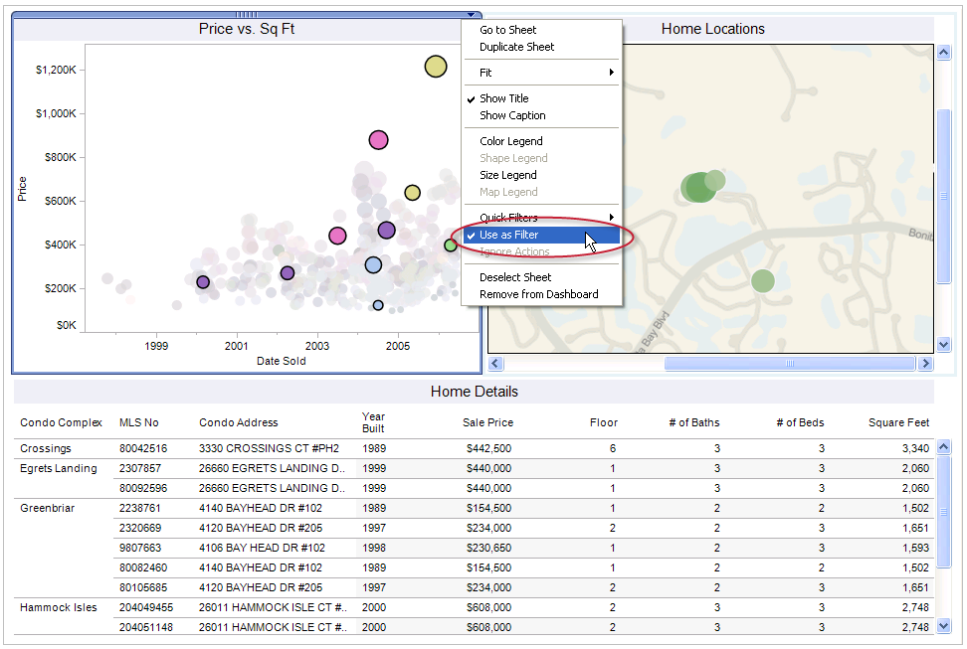

Condo Complex	MLS No	Condo Address	Year Built	Sale Price	Floor	# of Baths	# of Beds	Square Feet
Crossings	80042516	3330 CROSSINGS CT #PH2	1989	$442,500	6	3	3	3,340
Egrets Landing	2307857	26660 EGRETS LANDING D..	1999	$440,000	1	3	3	2,060
	80092596	26660 EGRETS LANDING D..	1999	$440,000	1	3	3	2,060
Greenbriar	2238761	4140 BAYHEAD DR #102	1989	$154,500	1	2	2	1,502
	2320669	4120 BAYHEAD DR #205	1997	$234,000	2	2	3	1,651
	9807663	4106 BAY HEAD DR #102	1998	$230,650	1	2	3	1,593
	80082460	4140 BAYHEAD DR #102	1989	$154,500	1	2	2	1,502
	80105685	4120 BAYHEAD DR #205	1997	$234,000	2	2	3	1,651
Hammock Isles	204049455	26011 HAMMOCK ISLE CT #..	2000	$608,000	2	3	3	2,748
	204051148	26011 HAMMOCK ISLE CT #..	2000	$608,000	2	3	3	2,748

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.

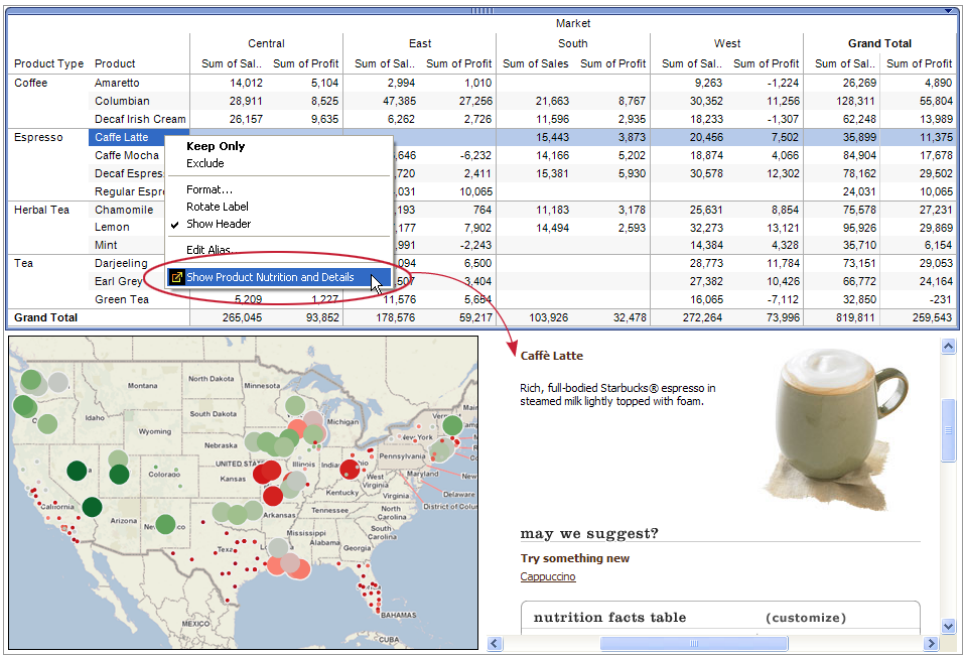

Using Field and Filter Values in Actions

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 and filter values as parameters in the URL of a URL action allows you to send information about a specific data point or filter setting to the destination webpage.

Using Field and Filter Values in URLs

Tableau lets you add field, filter information, and parameter values as variables into URL actions so when you follow the link the values of those fields and filters are included. For example, when linking to an online mapping service, you can insert the address field into the URL so launching the link from a specific data point shows the address associated with that record on a map.

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, parameter, or filter value.
3. Click the arrow to the right of the text box and select the field, parameter, or filter you want to add to the URL. The field, parameter, or filter name is added to the URL between angle brackets.You can continue adding field and filter values parameters as

many times as you need to create the URL.

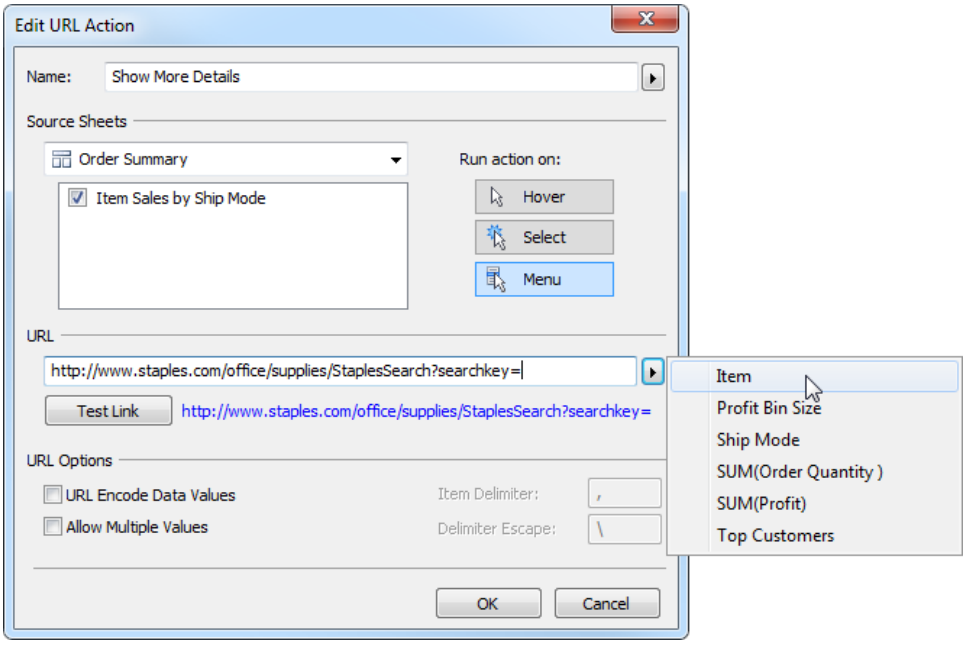

The list of available fields only includes non-aggregated fields. To use aggregated field values as a parameter in a link, you must first create a calculated field and then use the name of that field in the link. The calculated field must also be used in the view in order for the link to be available. A good way to use these fields is by placing them on the Level of Detail shelf.

When inserting parameter values, the Display As value is what is sent with the URL action. To send the actual value instead, you should add the character `~na` after the parameter name. For example, you may have a parameter that includes IP addresses. The Actual Value strings for the parameter include values such as 10.1.1.195 while the Display As strings include more friendly values such as `Computer A (10.1.1.195)`. Inserting the parameter into the URL like this: `http://<IPAddress>/page.htm` will send the display values so the final destination URL would be `http://Computer A (10.1.1.195)/page.htm`. Obviously, that is not likely to create a valid URL. To send the actual value, the parameter in the URL should look like this: `http://<IPAddress~na>/page.htm`.

Using Field and Filter Values in Action Names

In addition to using field, parameter, and filter values in URLs, you can use this information 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 and filter 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.

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 or filter value.
3. Click the arrow to the right of the text box and select the field or filter you want to add as a variable. The field or filter name is added between angle brackets.

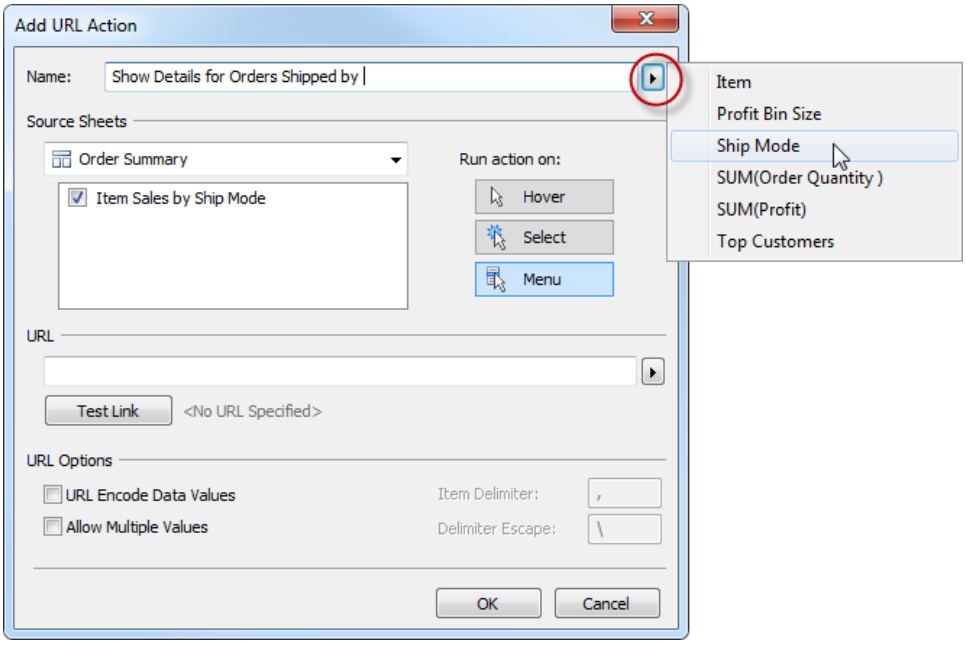

Calculations

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.

You can use all of the 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.

> You can use all of the following calculations if you are using a relational data source, however, multidimensional data sources do not support aggregations and binned data.

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 return a single value. For example, a measure that contains the values 1, 2, 3, 3, 4 aggregated as a sum returns a single value: 13. Or 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 have the highest revenue.

You can use Tableau to aggregate measures only with relational data sources. Multidimensional data sources contain aggregated data only.

> **Aggregations and floating-point arithmetic**: The results of some aggregations may not always be exactly as expected. For example, you may find that the Sum function returns a value such as -1.42e-14 for a column of numbers that you know should sum to exactly 0. This happens because the Institute of Electrical and Electronics Engineers (IEEE) 754 floating-point standard requires that numbers be stored in binary format, which means that numbers are sometimes rounded at extremely fine levels of precision. You can eliminate this potential distraction by using the ROUND function (see **Number Functions** on page 882) or by formatting the number to show fewer decimal places.

Tableau provides a set of predefined aggregations that are shown in the table below. You can set the default aggregation for any measure that is not a user-defined aggregation. See **How to Set the Default Aggregation for a Measure** on page 566. You can also set the aggregation for a field already in the view. See **Aggregating Data** on page 567.

Aggregation	Description	Result for measure that contains 1, 2, 2, 3
Attribute	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. This aggregation is particularly useful when aggregating a dimension. To set a measure in the view to this aggregation, right-click the measure and choose **Attribute**. The field then changes to show the text ATTR: ATTR(Sales)	N/A
Dimension	Returns all unique values in a measure or dimension.	3 values (1, 2, 3)
Sum	Returns the sum of the numbers in a measure. Null values are ignored.	1 value (8)

Aggregation	Description	Result for measure that contains 1, 2, 2, 3
Average	Returns the arithmetic mean of the numbers in a measure. Null values are ignored.	1 value (2)
Median	Returns the median of the numbers in a measure. Null values are ignored. This aggregation is not available for workbooks: • Created before Tableau Desktop 8.2 and that use Microsoft Excel or text file data sources. • That use legacy connections. • That use Microsoft Access data sources. If you are connected to one of these types of data sources, Median is unavailable and shows the message "Requires extract." To use this aggregation, extract your data. See **Extracting Data** on page 173.	1 value (2)
Count	Returns 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)	Returns 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 aggregation is not available for workbooks: • Created before Tableau Desktop 8.2 and that use Microsoft Excel or text file data sources. • That use legacy connections. • That use Microsoft Access data sources. If you are connected to one of these types of data sources, Count (Distinct) is unavailable and shows the message "Requires extract." To use this aggregation, extract your data. See **Extracting Data** on page 173.	1 value (3)
Minimum	Returns the smallest number in a measure or continuous dimension. Null values are ignored.	1 value (1)

Aggregation	Description	Result for measure that contains 1, 2, 2, 3
Maximum	Returns the largest number in a measure or a continuous dimension. Null values are ignored.	1 value (3)
Percentile	Returns the value at the specified percentile for the measure. When you select this aggregation, you must choose from a submenu offering a range of percentile values: 5, 10, 25, 50, 75, 90, 95. When you set this aggregation on a field in the view, the field shows PCT and the percent value assigned. For example: PCT95(BMI) If you want to use a percentage value other than the ones listed, use the PERCENTILE function in a calculation and specify the percentile you want. See **Aggregate Functions** on page 894. This aggregation is available for: • Non-legacy Microsoft Excel and text file connections. • Extracts and extract-only data source types (for example, Google Analytics, OData, or Salesforce). • Sybase IQ 15.1 and later data sources. • Oracle 10 and later data sources. If you are not connected to one of these types of data sources, Percentile is unavailable and shows the message "Requires extract." To use this aggregation, extract your data. See **Extracting Data** on page 173.	1 value. The value for PCT50 would be 2 for the given data.
Std. Dev	Returns 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)
Std. Dev (Pop.)	Returns 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)

Aggregation	Description	Result for measure that contains 1, 2, 2, 3
Variance	Returns 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 (Pop.)	Returns 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)
Disaggregate	Returns all records in the underlying data source. To disaggregate all measures in the view, select **Aggregate Measures** from the Analysis menu (to clear the check mark). Tableau allows you to view data in disaggregated form (relational databases only). 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 (see **Disaggregating Data** on page 571). Another way to look at disaggregated data is to view the underlying data for all or part of a view. See **View Data** on page 479.	4 values (1, 2, 2, 3)

You can also define custom aggregations as described in **Aggregate Calculations** on page 586. 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).

How to Set the Default Aggregation for a Measure

You can 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 you drag a measure to a view. Change the default aggregation by right-clicking a measure in the Data window and selecting **Default Properties > Aggregation** and then selecting one of the options.

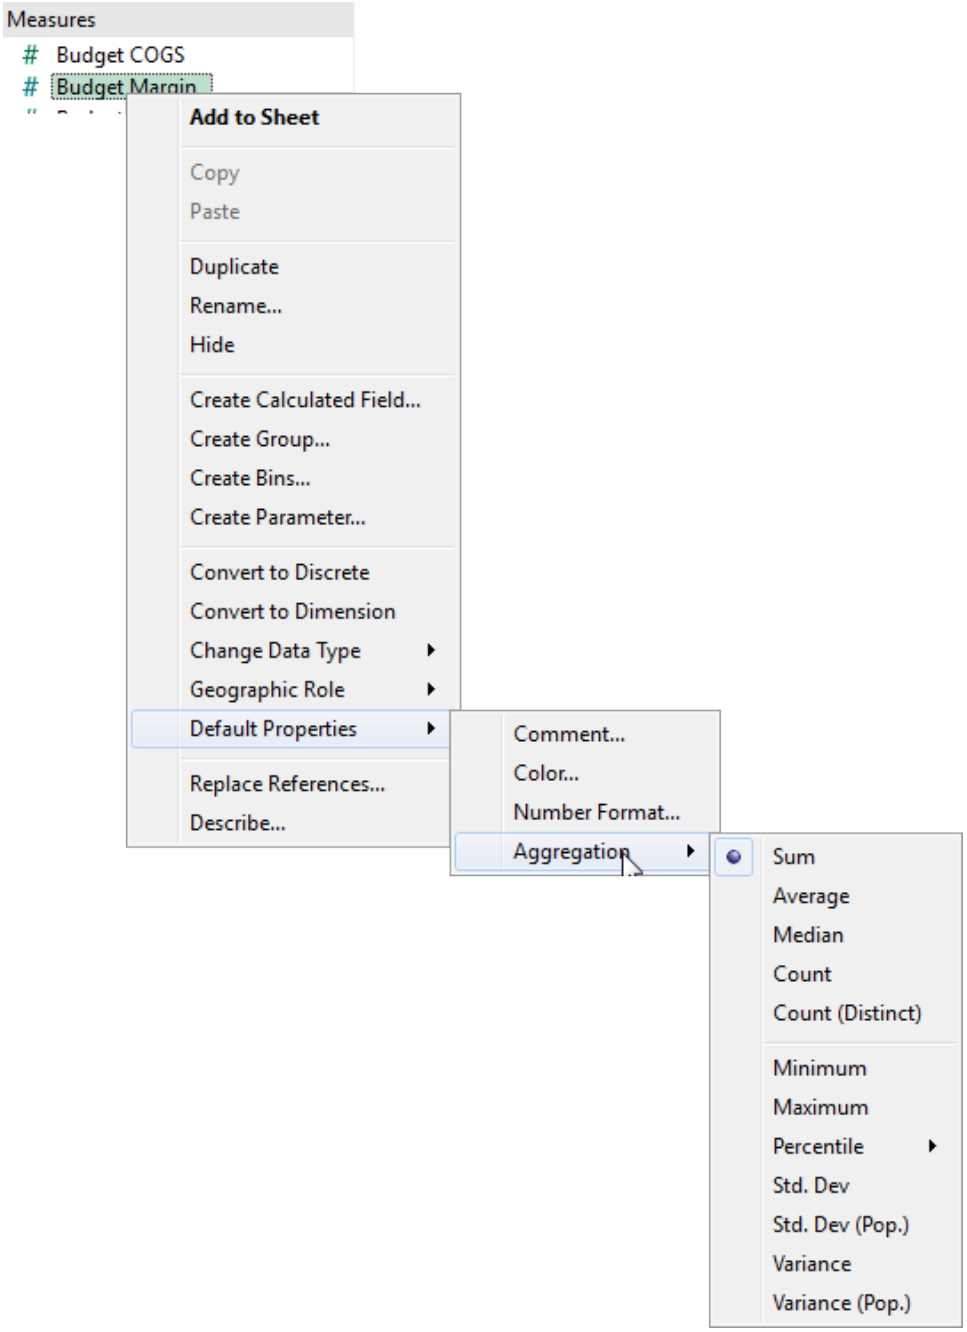

Aggregating Data

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, **Sales** becomes **SUM(Sales)**.

You can aggregate measures using Tableau only for relational data sources. Multidimensional data sources contain aggregated data only. In Tableau, multidimensional data sources are supported only in Windows.

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.

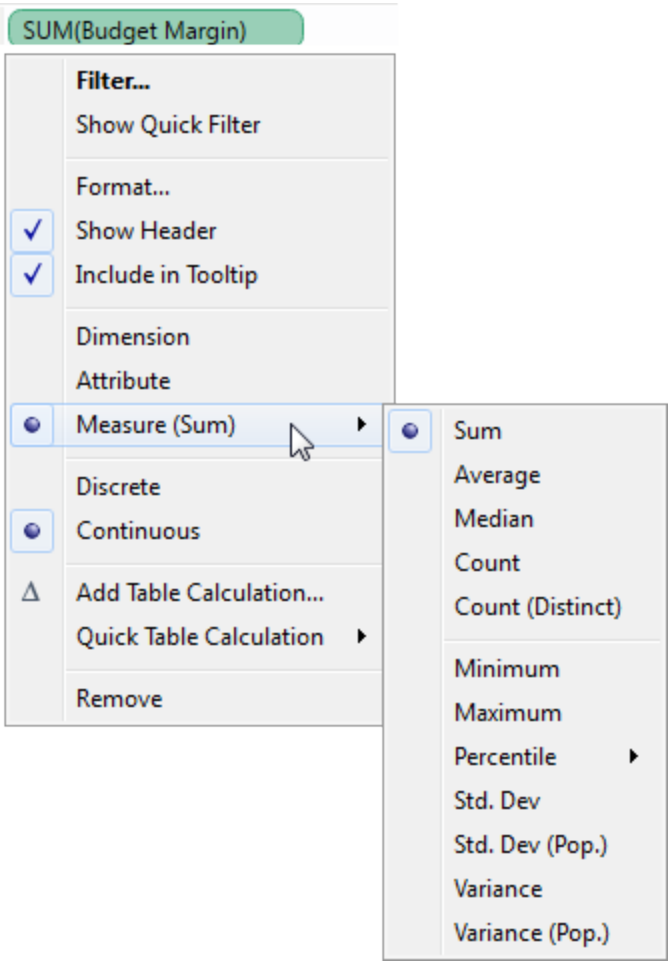

Aggregating Measures

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.

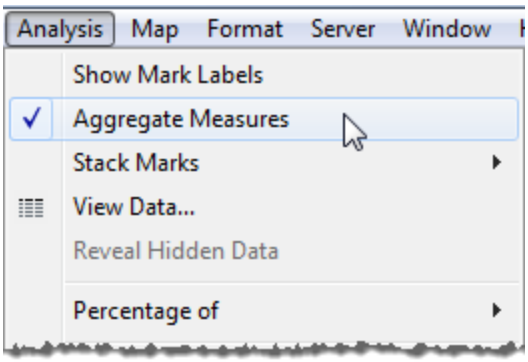

When all measures are disaggregated you see a mark for each row in the view. You cannot select specific marks to Keep Only, Exclude, or create a Set when all measures are disaggregated.

Aggregating Dimensions

A dimension can be aggregated as a measure using **Minimum**, **Maximum**, **Count**, or **Count (Distinct)**. When you aggregate dimensions, you create a new temporary measure column, so the dimension is now viewed as a measure.

> **Count (Distinct)** is not supported for Microsoft Access, Microsoft Excel, and text file data sources. If you are connected one of these types of data sources, the Count (Distinct) aggregation is unavailable and shows the remark "Requires extract.")

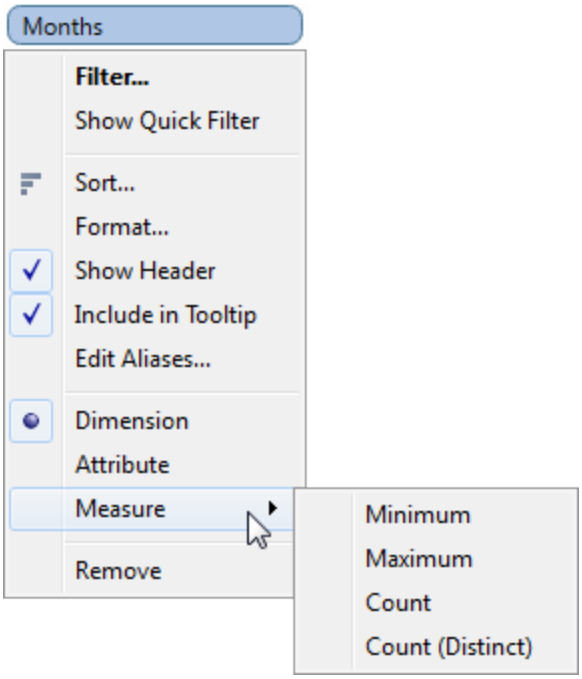

Another way to view a dimension is to treat it as an Attribute. Do this by choosing **Attribute** from the context menu for the dimension. When you aggregate a dimension as an attribute, it is

treated like a label instead of partitioning the data. The Attribute aggregation has several uses such as ensuring a consistent level of detail when blending multiple data sources, providing a way to aggregate dimensions when computing table calculations, which require an aggregate expression, and finally it can increase query performance because it is computed locally.

Tableau computes Attribute using the following formula:

```
IF MIN([dimension]) = MAX([dimension]) THEN MIN([dimension]) ELSE
"*" END
```

The formula is computed in Tableau after the data is retrieved from the initial query. The asterisk (*) is actually a visual indicator of a special type of Null value that occurs when there are multiple values. See **Troubleshooting Data Blending** on page 166 to learn more about the asterisk.

Below is an example of using Attribute in a table calculation. The table shows sales by market, market size, and state. Suppose you wanted to compute the percent of total sales each state contributed to the market. When you add a Percent of Total **Quick Table Calculations** on page 602 that computes along State, the calculation computes within the red area shown below. This is because the Market Size dimension is partitioning the data.

Columns: Measure Names

Rows: Market | Market Size | State

Market	Market Size	State	Sales	% of Total Sales along State
Central	Major Market	Colorado	$48,179	31.58%
		Illinois	$69,883	45.80%
		Ohio	$34,517	22.62%
	Small Market	Iowa	$54,750	48.68%
		Missouri	$24,647	21.92%
		Wisconsin	$33,069	29.40%
East	Major Market	Florida	$37,443	27.08%
		Massachusetts	$29,965	21.67%
		New York	$70,852	51.25%
	Small Market	Connecticut	$25,429	63.07%
		New Hampshire	$14,887	36.93%
South	Major Market	Texas	$37,410	100.00%
	Small Market	Louisiana	$23,161	34.82%
		New Mexico	$15,892	23.89%

When you aggregate Market Size as an Attribute, the calculation is computed within the Market (East, in this case), and the Market Size information is used purely as a label in the display.

Columns: Measure Names

Rows: Market | ATTR(Market Size) | State

Market	Market Size	State	Sales	% of Total Sales along State
Central	Major Market	Colorado	$48,179	18.18%
		Illinois	$69,883	26.37%
		Ohio	$34,517	13.02%
	Small Market	Iowa	$54,750	20.66%
		Missouri	$24,647	9.30%
		Wisconsin	$33,069	12.48%
East	Major Market	Florida	$37,443	20.97%
		Massachusetts	$29,965	16.78%
		New York	$70,852	39.68%
	Small Market	Connecticut	$25,429	14.24%
		New Hampshire	$14,887	8.34%
South	Major Market	Texas	$37,410	36.00%
	Small Market	Louisiana	$23,161	22.29%
		New Mexico	$15,892	15.29%

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.

> If your data source is very large, disaggregating the data can result in a significant performance degradation.

You can disaggregate all measures in the view by selecting **Analysis** > **Aggregate Measures**.

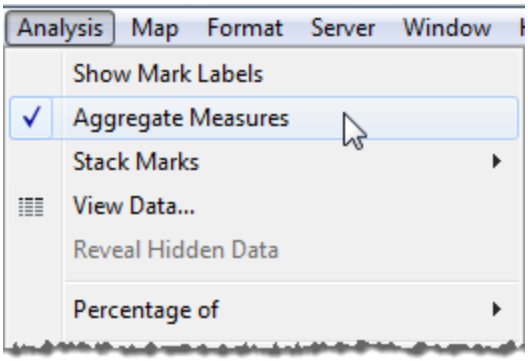

Example – Scatter Plots and Aggregation

If you place one measure on the **Rows** shelf and another measure on the **Columns** shelf, you are asking Tableau to compare two numerical values. Typically, Tableau chooses a scatterplot as the default visualization in such cases. The initial view can be disappointing—a single mark, showing the sum for all values for the two measures. There are various ways to develop such a scatterplot: you can **use dimensions to add detail**, you can add additional measures and/or dimensions to the Rows and Columns shelves, to create multiple one-mark scatter plots in the view, or you can **disaggregate the data**. (Or you can use any combination of these options.) This topic looks at these alternatives, and uses the **Sample – Superstore English (Extract)** data source. To create the initial view, follow these steps:

1. Place the **Sales** measure on the **Columns** shelf.
2. Place the **Profit** measure on the **Rows** shelf.

The measures are automatically aggregated as sums. The default aggregation (SUM) is indicated in the the field names. The values shown in the tooltip show the sum of sales and profit values across every row in the data source.

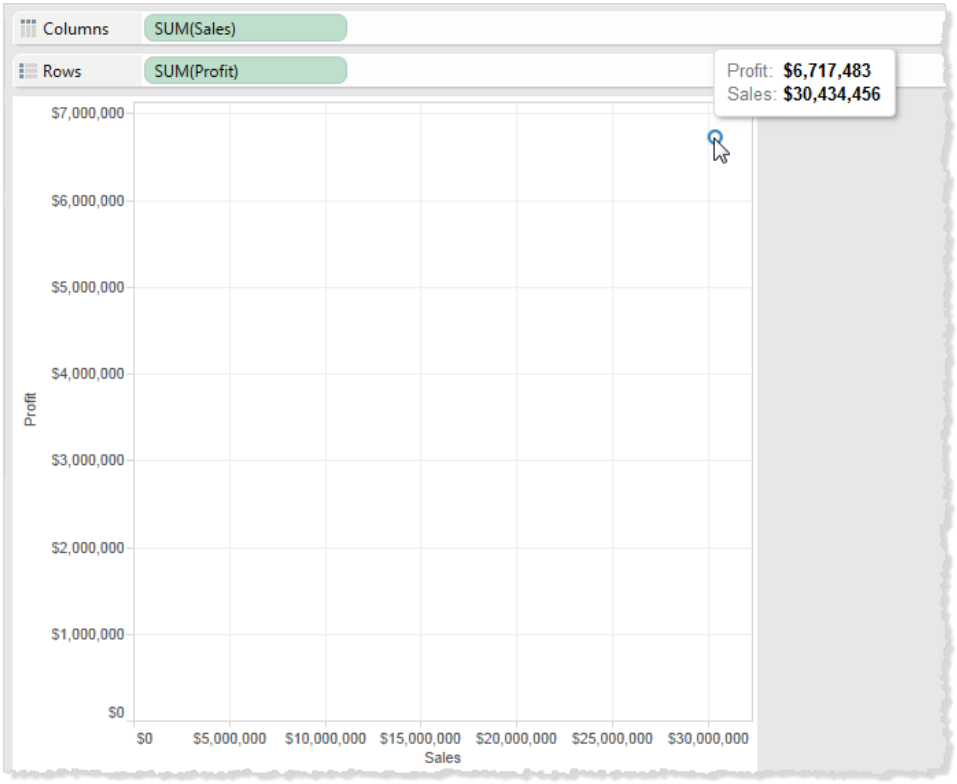

Using Dimensions to Add Detail

Follow these steps to develop the scatter plot view you created above by adding dimensions to show additional levels of detail.

1. Drag the **Department** dimension to the **Color** target on the marks card.

   This separates the data into three marks—one for each dimension member—and encodes the marks using color.

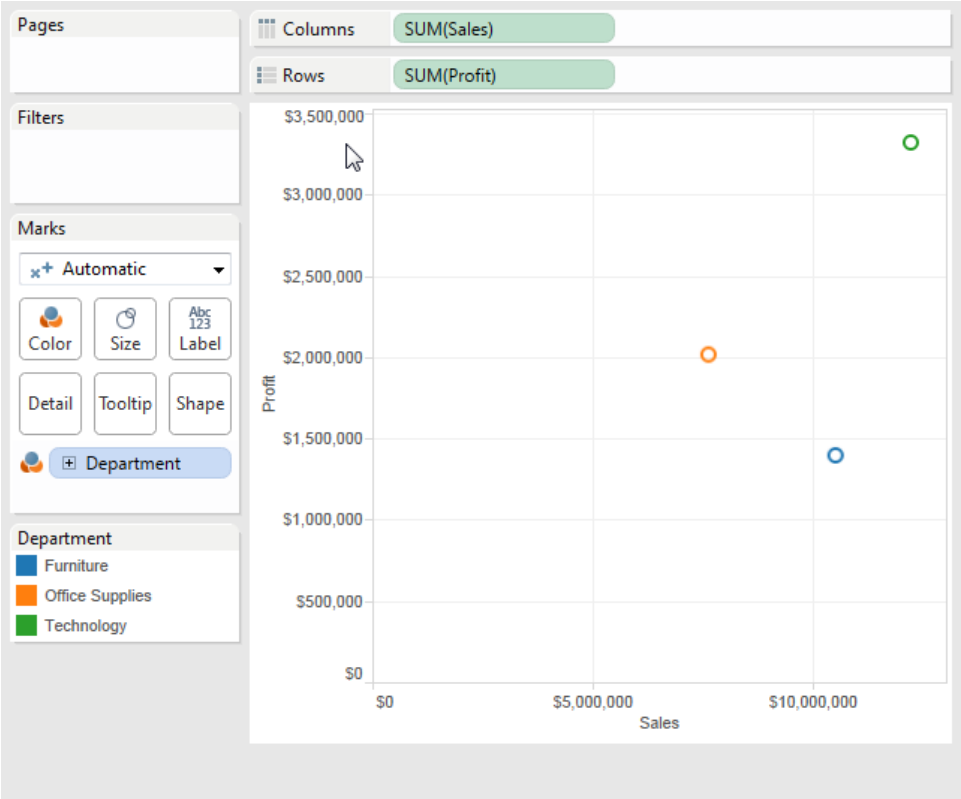

2. Drag the **Country/Region** dimension to the **Detail** shelf.

   Now there are many more marks in the view. The number of marks is equal to the number of distinct country/regions in the data source multipled by the number of departments.

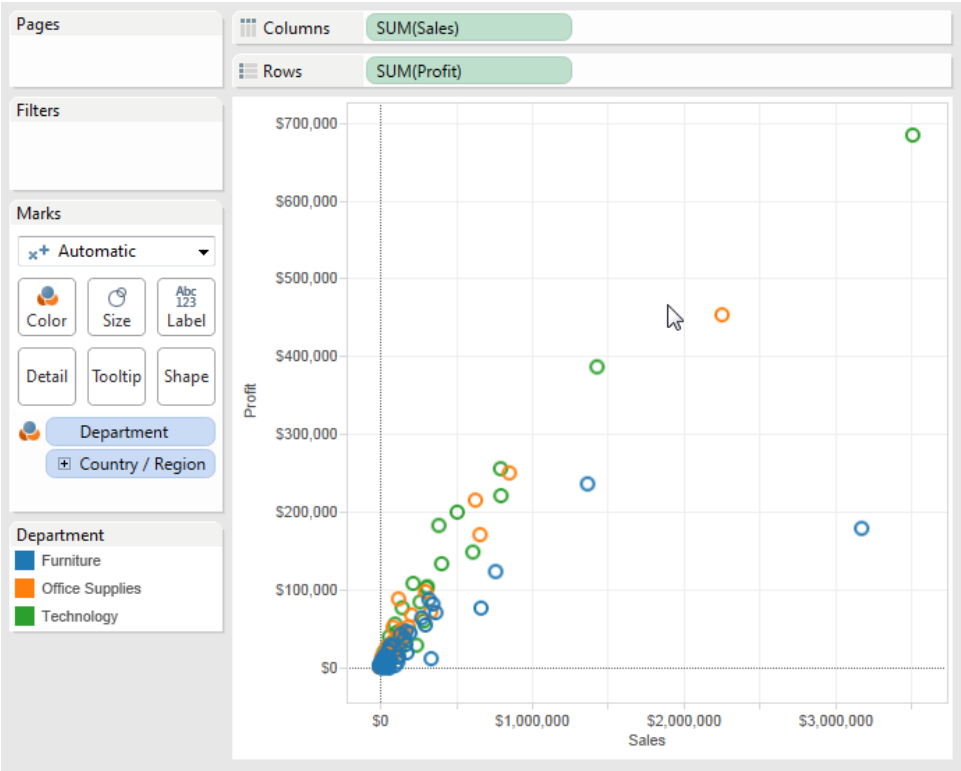

Although more marks are now displayed, the measures are still aggregated. So regardless of whether there is one row in the data source where Country/Region = China and Department = Furniture, or 100 such rows, the result is always a single mark.

Maybe this process is developing the view in a direction you find useful, or maybe you prefer to go in a different direction—for example, by adding time dimension to the view, or by introducing trend lines or forecasting.You decide what questions to ask.

Adding More Fields to the Rows and Columns Shelves

Revert to the original one-mark view and follow these steps to develop the scatter plot view by adding fields to the **Rows** and **Columns** shelves.

1. Drag the **Continent** dimension to the **Columns** shelf.

   Even if you drop **Continent** to the right of **SUM(Sales)**, Tableau moves it to the left of **SUM(Sales)**. This is because you cannot insert a dimension within a continuous axis. Instead, your view shows a separate axis for each member of the dimension.

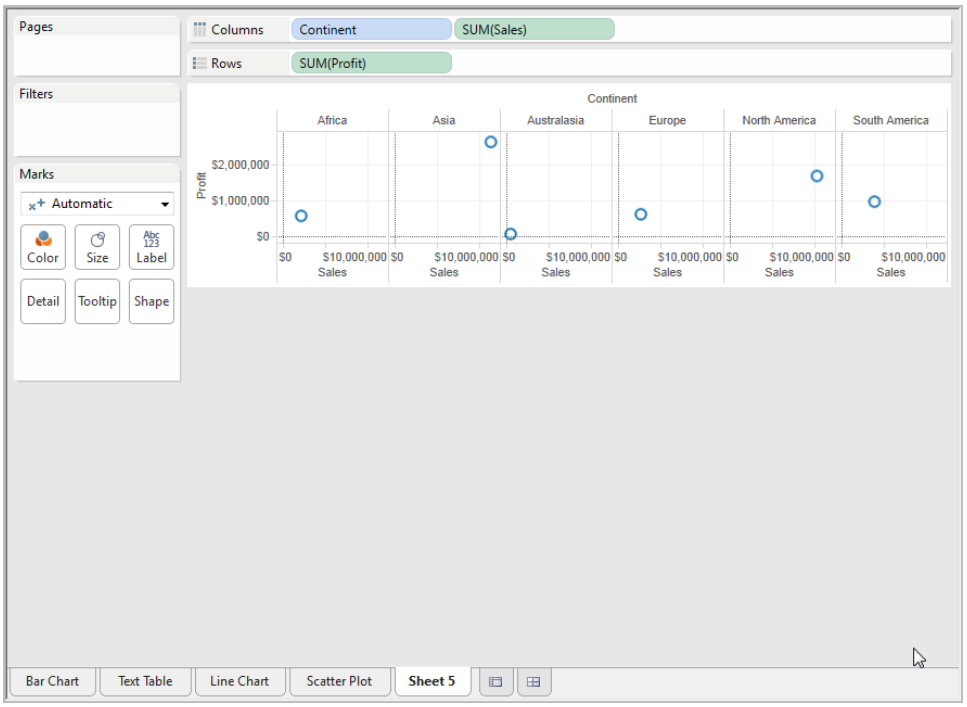

2. Drag the **Customer Segment** dimension to the **Rows** shelf.

   You now have a view that provides an overview of Sales and Profit across continents and customer segments. It can be interesting to hover over the marks in the view to see tooltip data for various segments:

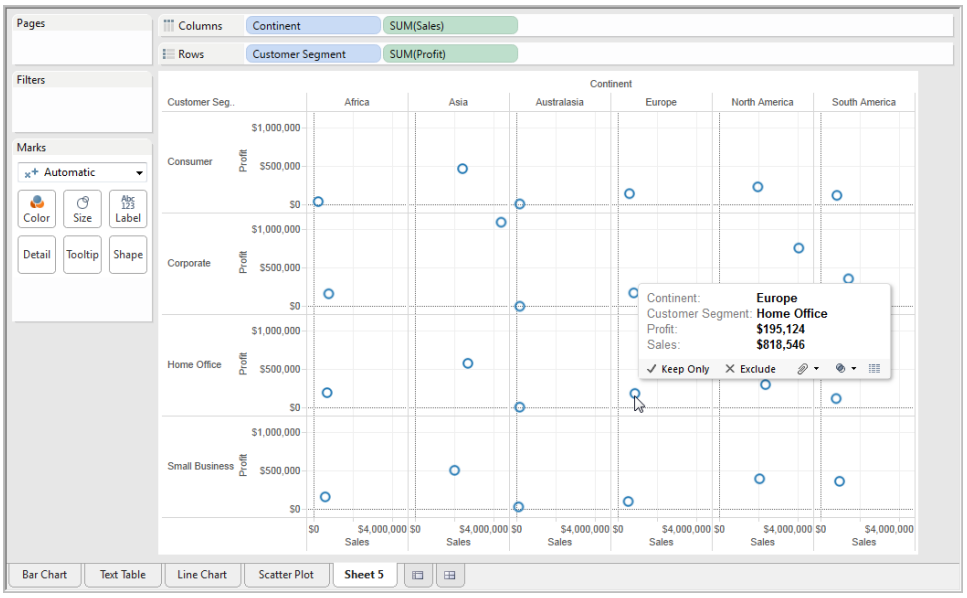

Disaggregate Data

Another way to modify your original one-mark scatter plot to display more marks is by disaggregating the data.

Select **Analysis > Aggregate Measures**.

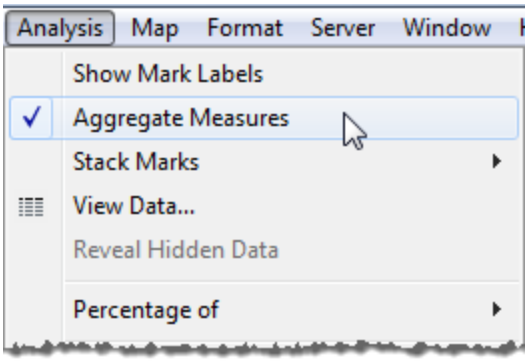

What you have actual done is to dis-aggregate the data, because this command is a toggle that was original selected (check mark present). Tableau aggregates data in your view by default.

Now you see a lot of marks--one for each row in your original data source:

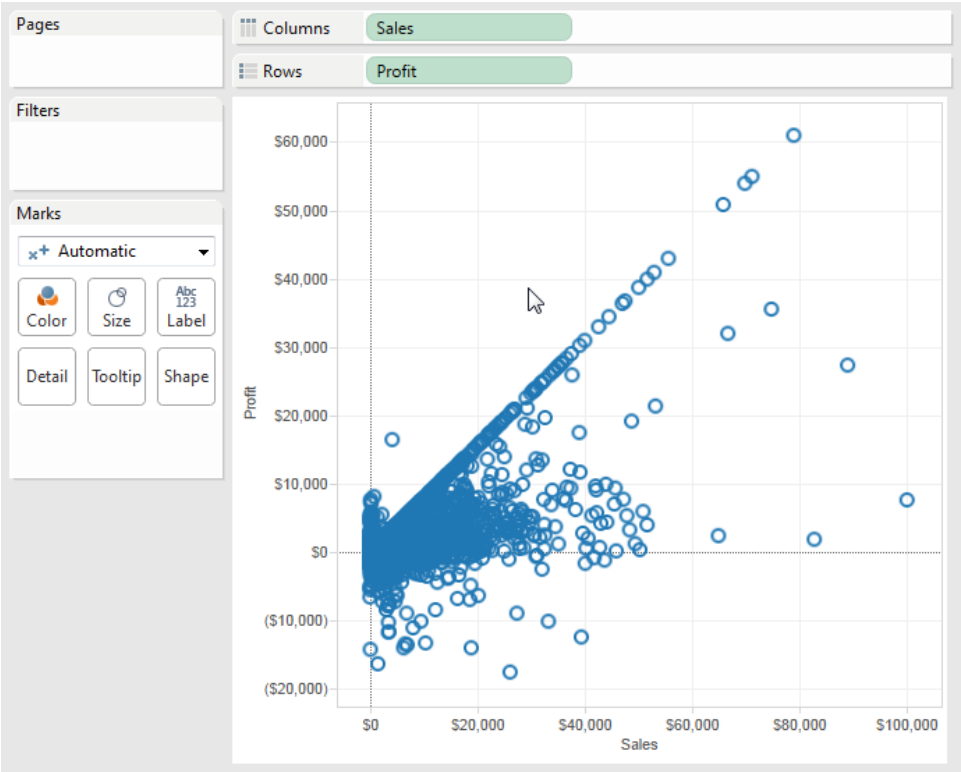

When you disaggregate measures, you no longer are looking at the average or sum for the values in the rows in the data source. Instead, the view shows a mark for every row in the data source. Disaggregating data is a way to look at the entire surface area of the data. It's a quick way to understand the shape of your data and to identify outliers. In this case, the disaggregated data shows that for many rows in the data, there is a consistent relationship between sales income and profit—this is indicated by the line of marks aligned at a forty-five degree angle.

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.

How to Create a Calculated Field

To create a new calculated field, select **Analysis** > **Create Calculated Field**, or right-click in the Data window and choose **Create Calculated Field** .

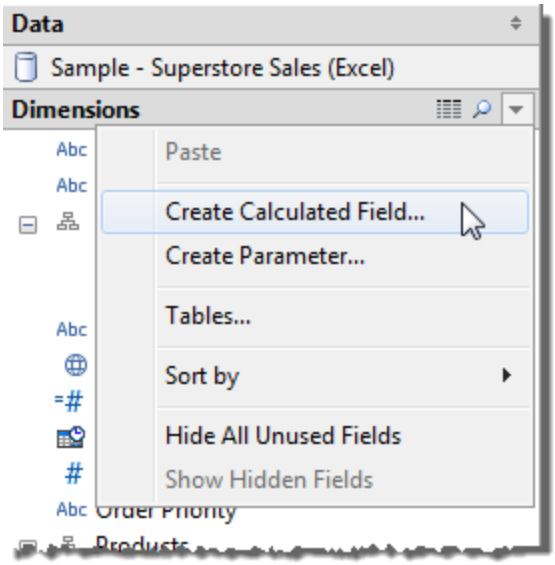

The **Calculated Field** dialog box opens.

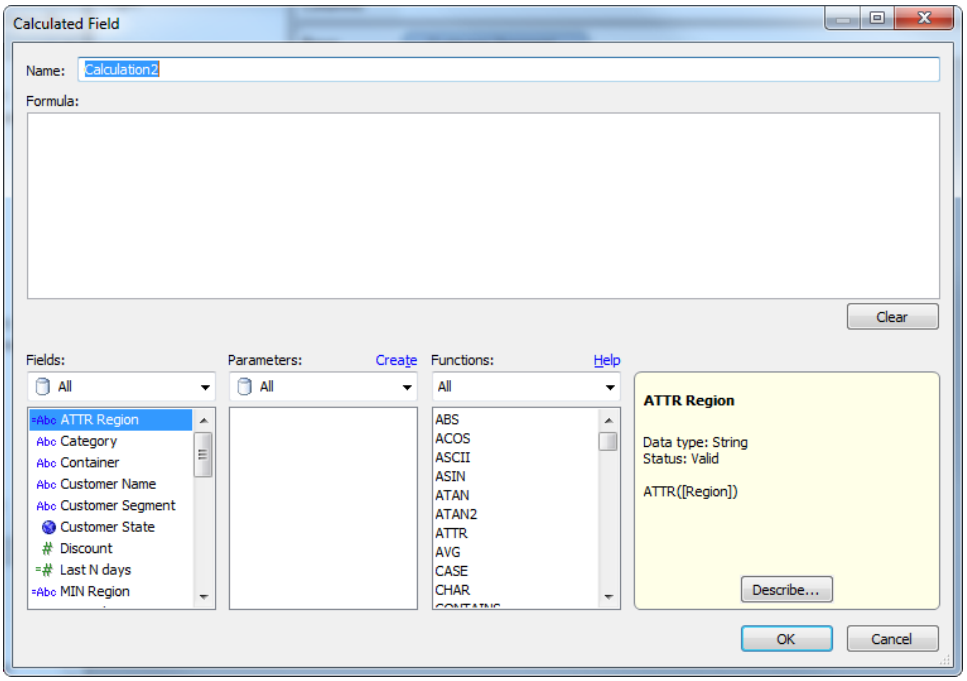

To define the calculation do the following:

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 580 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.

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: Soda, Coffee, and Crackers. You can define a new calculated member called "Beverages" that sums the Soda and Coffee members. Now when you place the Products dimension on the Rows shelf it displays four rows: Soda, Coffee, Crackers, and Beverages.

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.

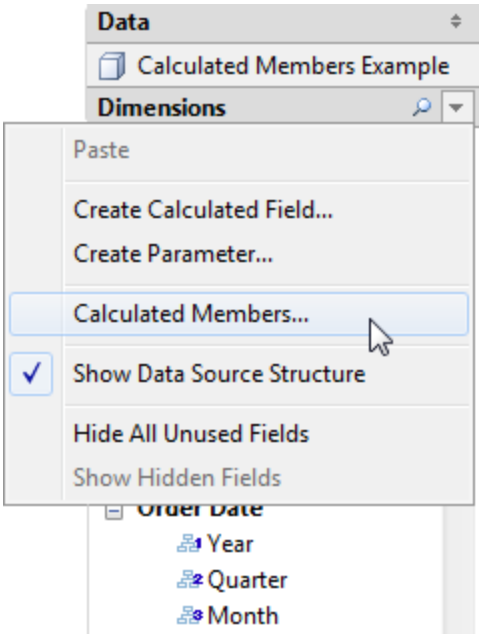

To create a new calculated member do the following:

1. Click New to add a new row to the list of calculated members at the top of the dialog box.

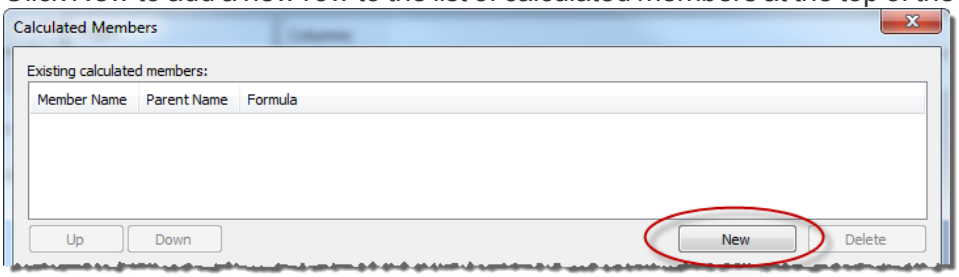

2. Type a Name for the new calculated member in the Member Definition area of the dialog box.

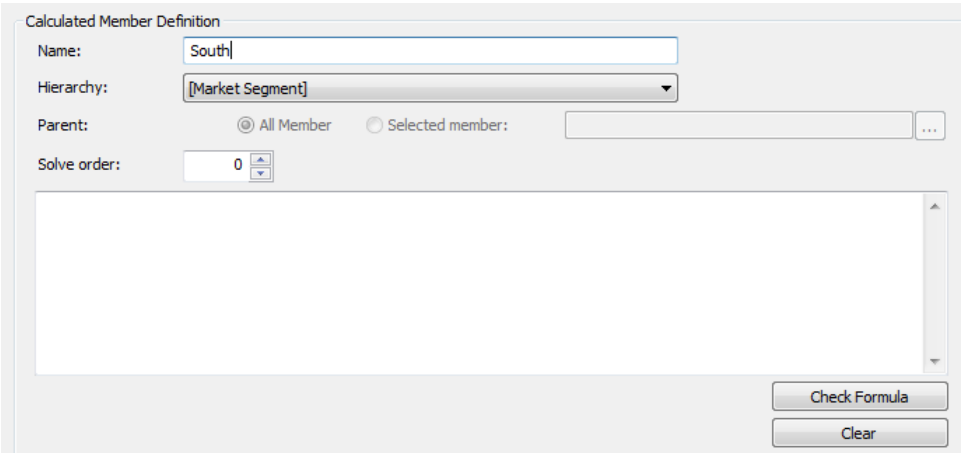

3. Select a hierarchy from the Hierarchy drop-down list. If you are creating a calculated measure, select [Measures] and define a result type using the **Result Type** drop-down list. Then skip to step five.
4. Specify the Parent member for the new calculated member. The All Member is selected by default, 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:

- **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.

- **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. 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.

- **Operators**

  Type any necessary operators into your formula. All standard operators such as addition (+), subtraction (–), multiplication (*), and division (/) are supported. Operators are colored black in the formula.

- **Parameters (optional)**

  Parameters are placeholders 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. Parameters are colored purple in the formula.

- **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 create a calculated field using Tableau formulas and use the new field in a data view. Then we edit the field's formula to create a new view, and delete the field from the Data window. the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop.

1. Create the view.

   Select **Create Calculated Field** on the Data window menu.

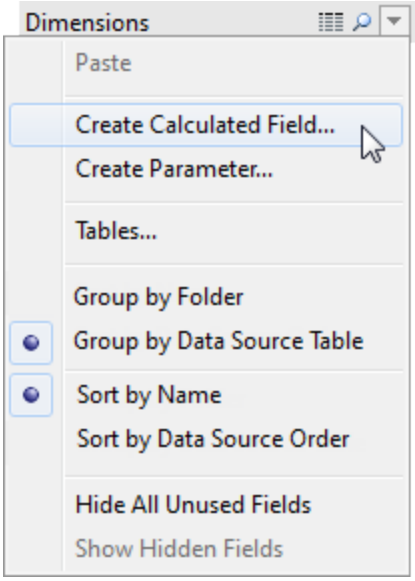

2. Name the new field **Discount Ratio** and enter the formula shown below.

   `IIF([Sales] !=0, [Discount]/[Sales],0)`

   Create the formula by double clicking field names in the **Fields** list, and functions in the **Functions** list. Type the operators (`!=` and `/`) manually.The `IIF` statement is used to avoid dividing by zero.

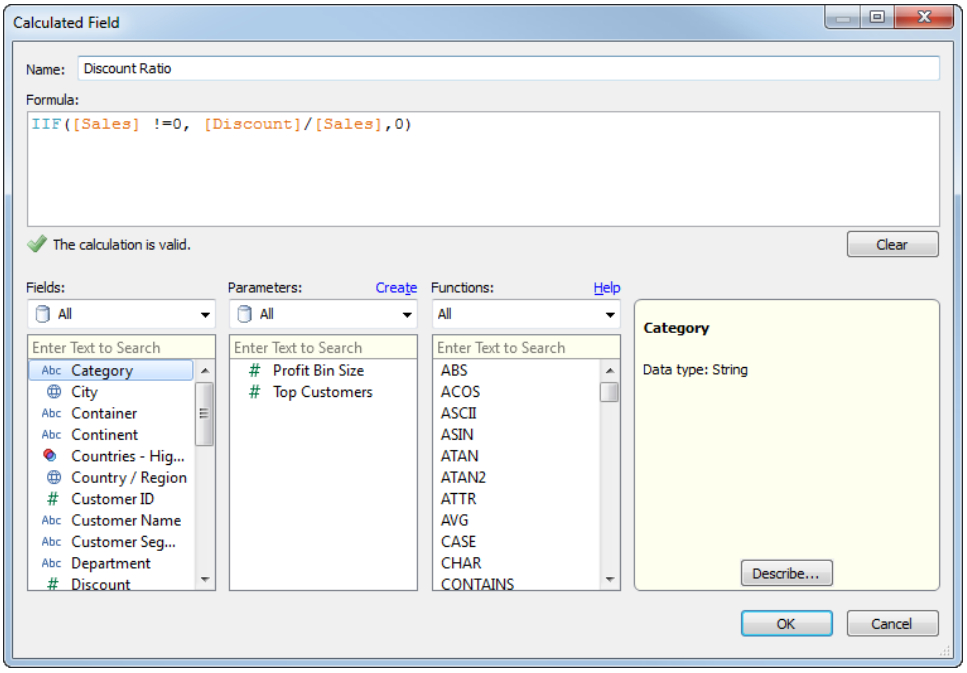

The new field displays in the Measures area of the Data window because the calculation returns a number. You can use this new field just like any other field.

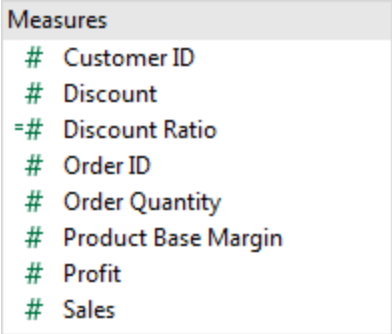

3. Add the calculation to the view.

   Place **Region** on the **Columns** shelf, **Ship Mode** on the **Rows** shelf, **Department** on Color in the Marks card. Then place the new calculated field, **Discount Ratio**, on the **Rows** shelf. You can treat the new calculated field just like any other measure. For example, you can apply an aggregation to it. Below, Discount Ratio is aggregated as an average.

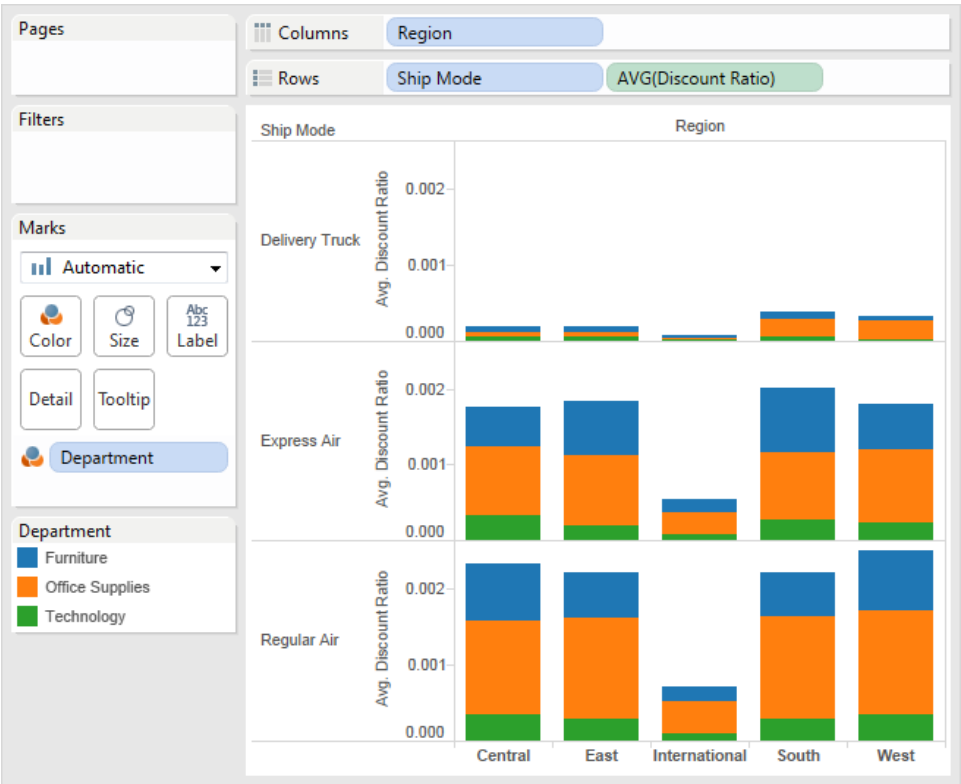

4. Edit the calculation.

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**.

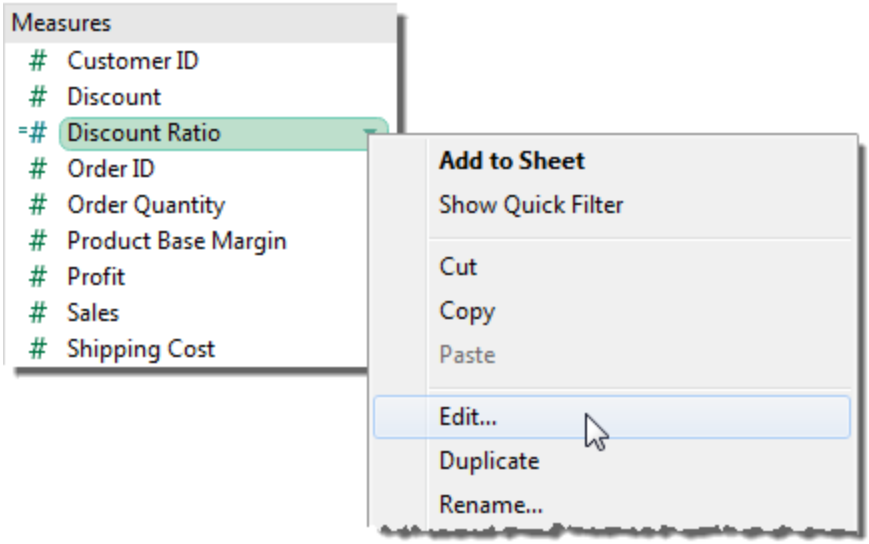

In the Calculated Field dialog box, change the formula to compute discount ratio only for sales over $2000.

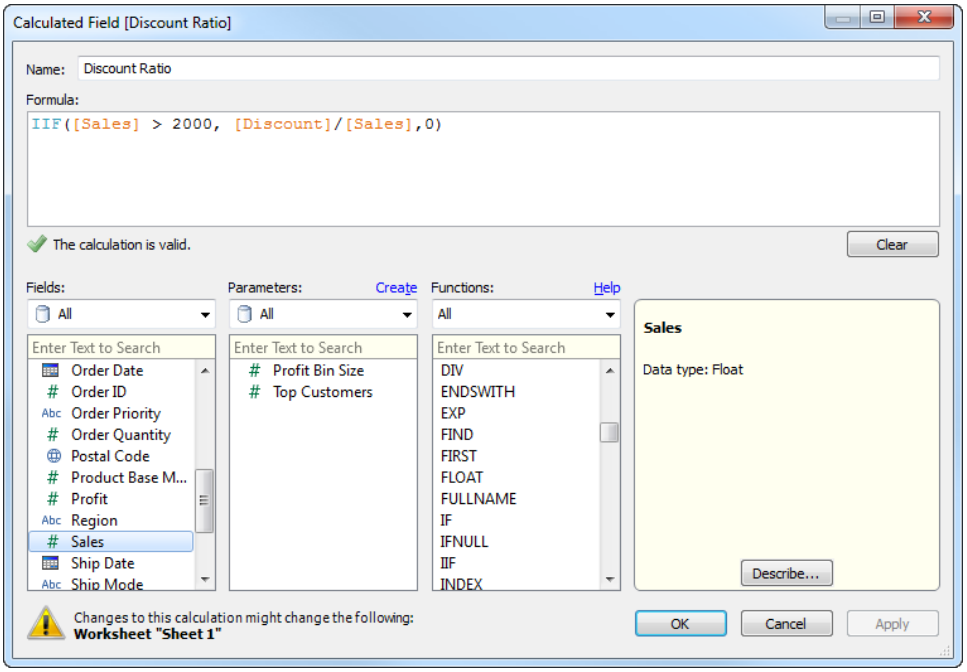

The view automatically updates after you click **OK** in the Calculated Field dialog box.

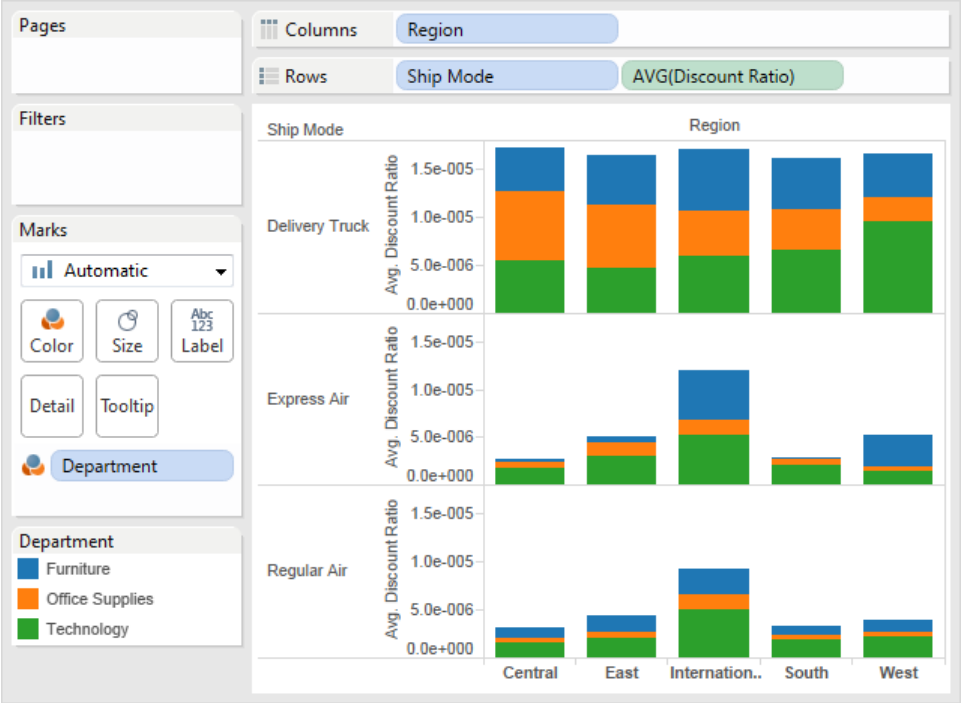

5. Delete the calculated field.

   You can delete a calculated field by right-clickng the field name in the Data window and selecting **Delete**. Before deleting the field, you might want to save your workbook. If you do not save your work, the calculated field will be lost.

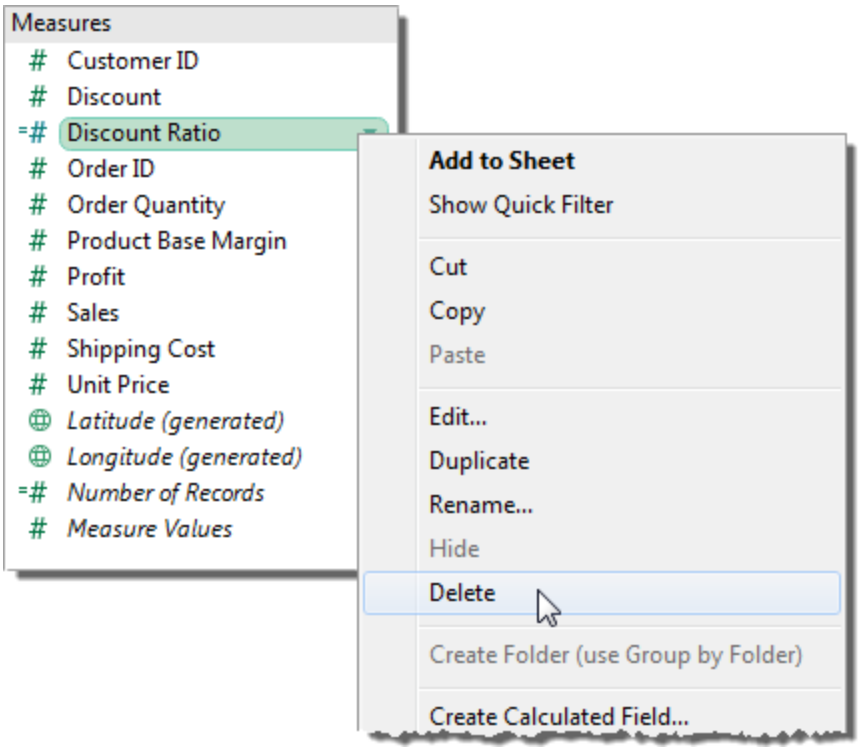

Aggregate Calculations

Aggregate functions allow you to summarize data. As described in **Aggregations** on page 563, 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.

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:

```
Margin = SUM([Profit]/ [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.

```
Margin = SUM( [Profit]) / SUM([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 577. 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 563.

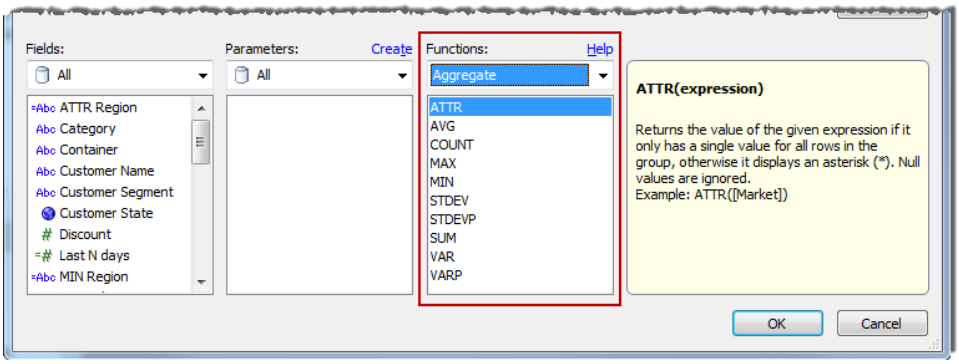

The aggregate calculation appears with the letters **AGG** in front of it when it is placed on a shelf.

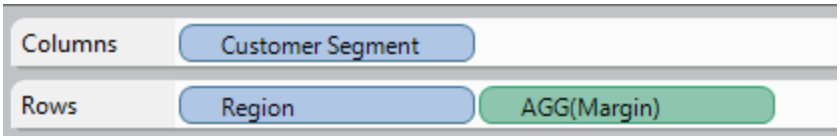

When you create an aggregate calculation, no further aggregation of the calculation is possible. Therefore, the field's context menu does not offer any aggregation choices. However, you can disaggregate the field.

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** on page 629 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 563 for descriptions of the aggregation functions.

Therefore, if you define an aggregate calculation called `Margin` that is equal to `SUM (Profit)/SUM(Sales)` and then disaggregate the data, it is interpreted as `Profit/Sales`.

Example – Aggregate Calculation

In this example you will use the **Sample - Superstore - English (Extract)** data source, which is included with Tableau Desktop. You will create an aggregate calculation called Margin, and use the new field in a data view.

1. Select **Create Calculated Field** on the Data window menu.

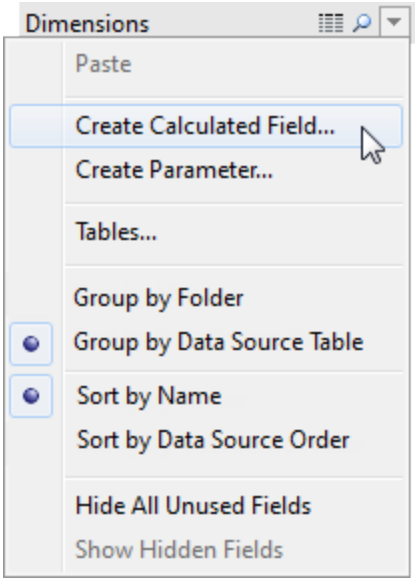

2. Name the new field **Margin** and enter the formula shown below.

```
IIF(SUM([Sales]) !=0, SUM([Profit])/SUM([Sales]), 0)
```

3. Create the formula by double-clicking functions from the **Functions** list, and field names from the **Fields** list. Type the operators (`!=` and `/`) manually. The `IIF` statement is used to avoid dividing by zero.

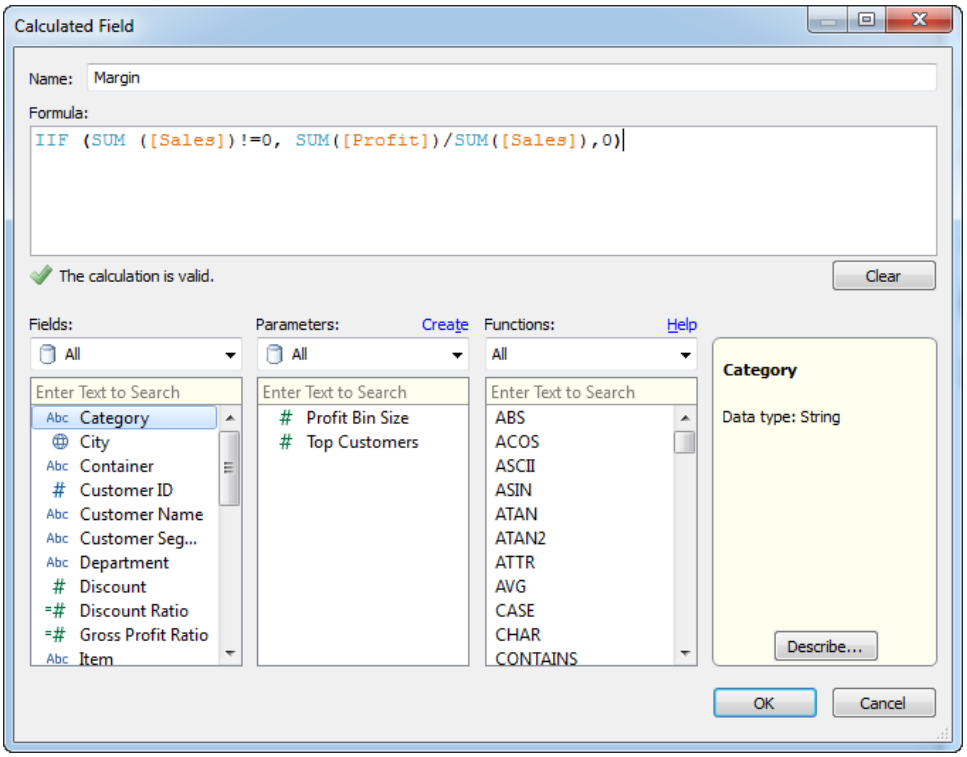

The new calculated field displays in the **Measures** area of the Data window where you can use it like any other measure.

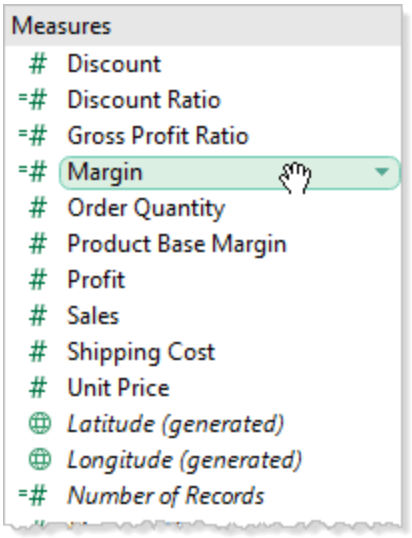

A view using the new aggregate measure is shown below.

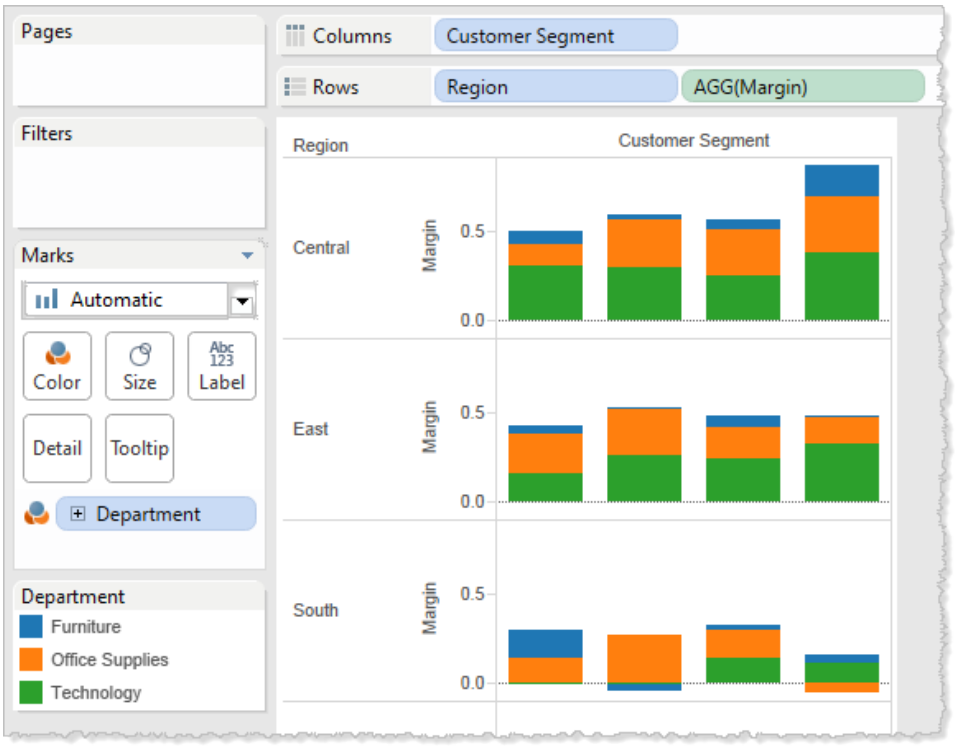

When **Margin** is placed on a shelf, its name is automatically changed to **AGG(Margin)**, which indicates it's an aggregate calculation. Additionally, the field's context menu does not include any aggregation choices because aggregating a field that's already aggregated is not possible.

Example - Spotlighting Using Calculations

Spotlighting is a term that applies a calculation that shows discrete thresholds based on the values of a measure. For instance, you might want to color-code sales so that those over 10,000 appear green and those below 10,000 appear red. A spotlighting calculation is just a special case of a calculation that results in a discrete measure. A discrete measure is a calculation that is a dependent variable (and therefore a measure), but which results in a discrete result (as opposed to a continuous result). Thus the name discrete measure. Here is an example:

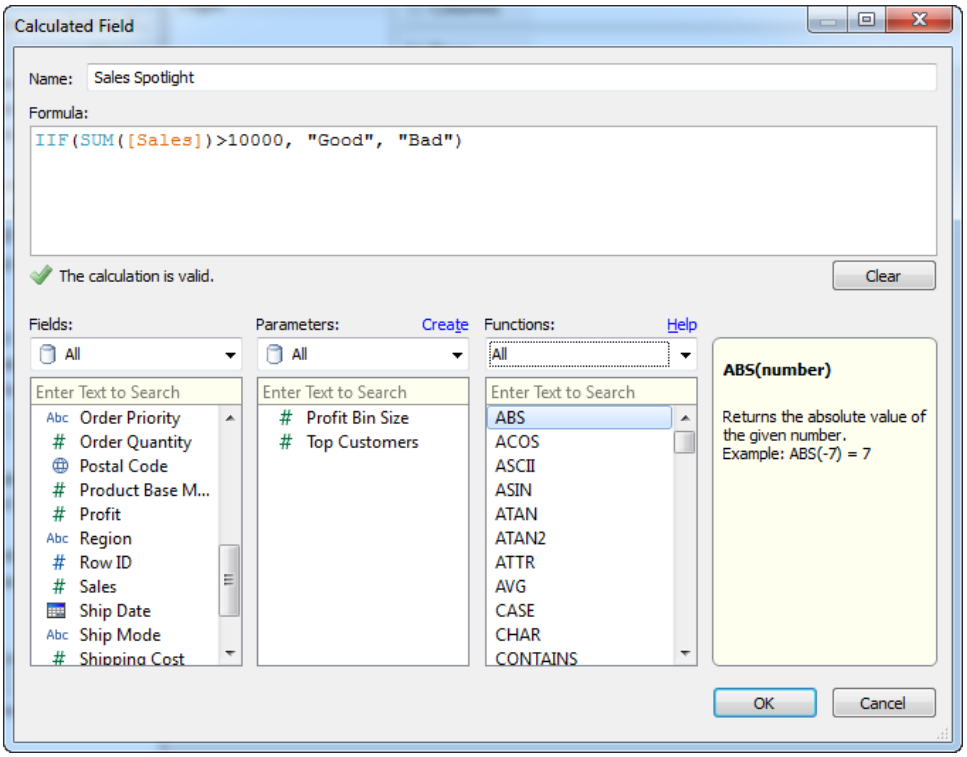

The formula in this example defines a discrete measure called “Sales Spotlight.” Discrete measures always appear with a blue “abc” icon in the Data window. The example above is a measure because it is a function of another measure. It’s discrete because it produces discrete values (“Good” and “Bad”) as a result rather than continuous values like numbers. Here is an example of this categorical measure in use:

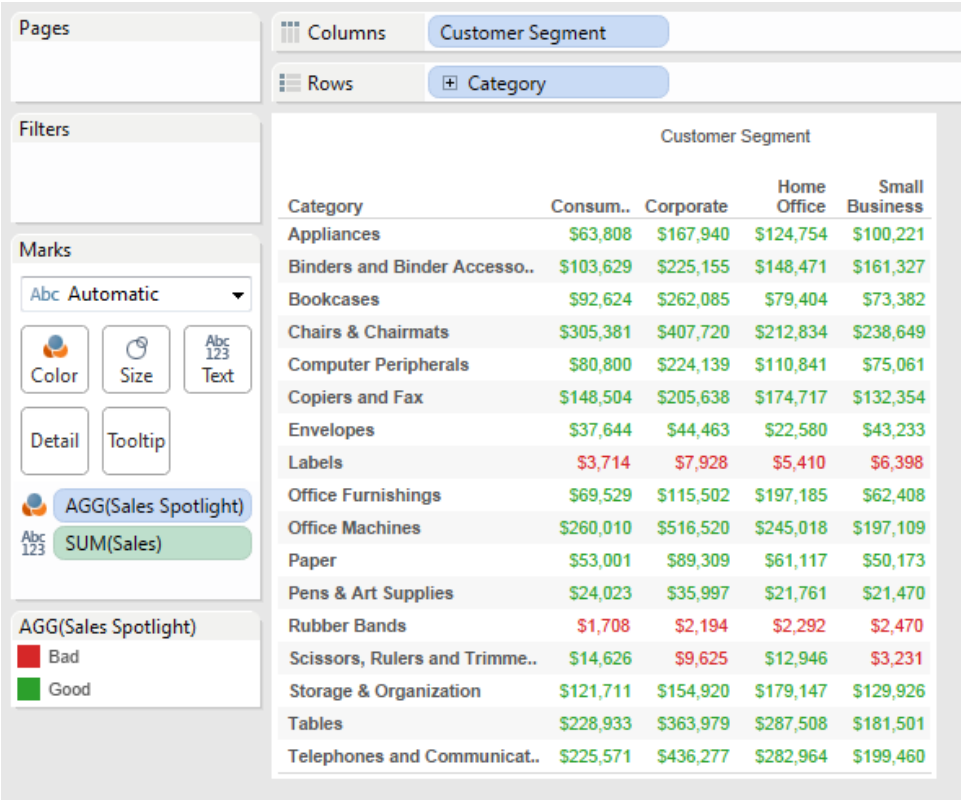

Here the “sales spotlight” field is on Color in the Marks card. It appears with the “AGG” prefix because it is an aggregate calculation. Values above 10,000 and below 10,000 are assigned different colors. This type of discrete highlighting is often called spotlighting.

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 values you want to compute and what values to compute along. These are defined in the Table Calculation dialog box using the Calculation Type and Calculate Along drop-down menus.

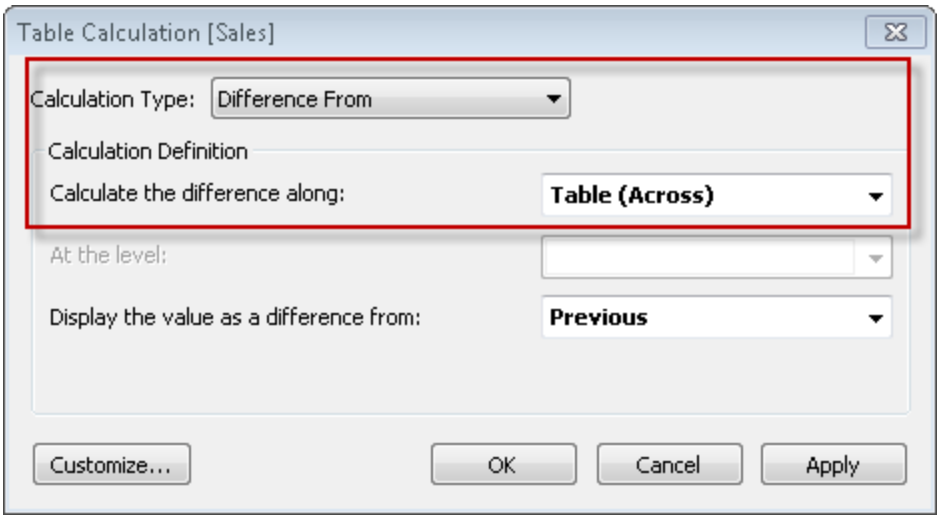

The definition of what to compute along has two parts: addressing fields and partitioning fields.

Addressing and Partitioning

Table calculations are performed on a single measure in the view. The dimensions that define the part of the table you are applying the calculation to (computing along) are called *addressing* fields, and the dimensions that define how to group the calculation are called *partitioning* fields. In the example of a running sum of product sales across several years, the addressing field is the Date field, and the partitioning field is the product field. When you define the addressing for a table calculation, all the other dimensions are used for partitioning.

You specify addressing when you create or edit a table calculation, in the Table Calculation dialog box. You can also update addressing for a field in the view that already has a table calculation by right-clicking the field and choosing one of the options under **Compute using**. For example:

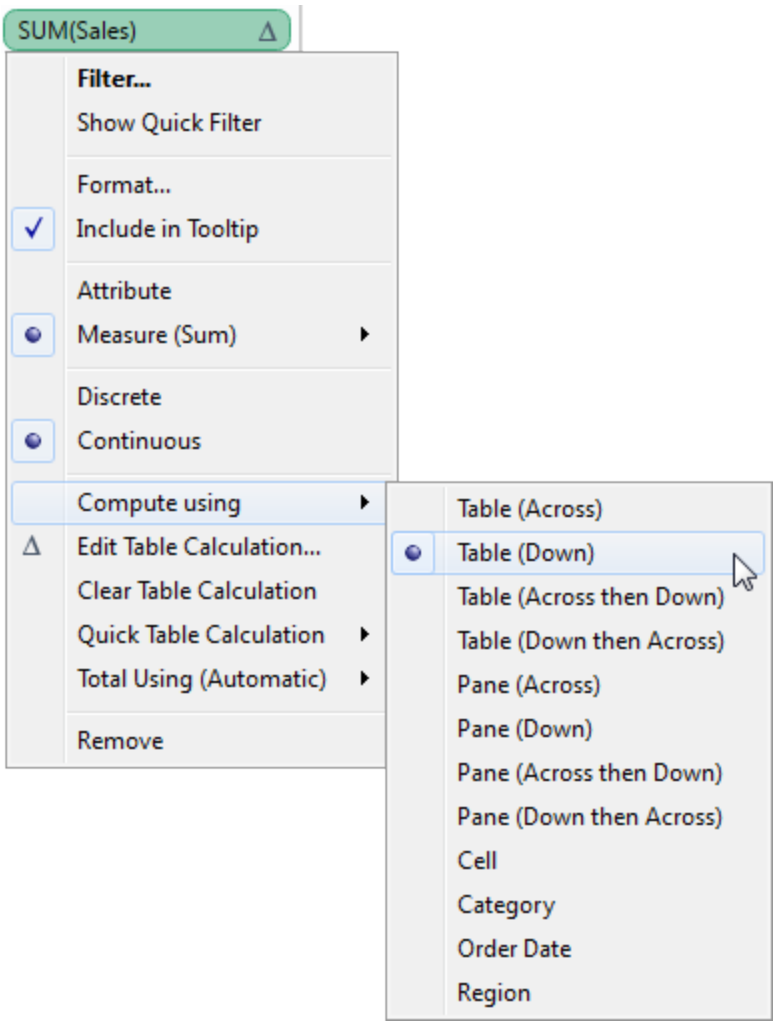

Addressing can be relative to the table structure (options beginning with **Table**, **Pane**, or **Cell**) or to a specific field (such as **Category**, **Order Date**, or **Region**). Addressing options based on table structure are described below.

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 dimensions that span horizontally across the table are the addressing fields (in the view below, **Category** and **Region**). All the other dimensions (**Year**, **Quarter**) are partitioning. The addressing dimensions are shown in orange, while partitioning dimensions are shown in blue.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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**. Any calculation that is performed is scoped to the partition. For example, if the calculation is percent of total, the calculation will be performed on the numbers within each of the orange boxes.

		Category / Region								
		Furniture				Office Supplies				
Year of Order Date	Quarter of Order Date	Central	East	South	West	Central	East	South	West	Partition
2008	Q1	$54,839	$41,843	$54,162	$98,604	$69,738	$20,531	$38,466	$85,317	1
	Q2	$71,341	$56,705	$103,741	$103,660	$43,070	$38,547	$65,251	$108,219	2
	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	

Table (Down)

This option sets the addressing to compute along the entire table moving vertically through each partition. For example, the 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**, **Region**). The addressing fields are shown in orange while partitioning fields are shown in blue.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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 entire table is the partition. The computation will compute across, move to the next row and continue to compute across, and so on.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	Central	East	South	West	Central	East	South	West
2008	Q1	$54,839	$41,843	$54,162	$98,604	$69,738	$20,531	[illegible]	[illegible]
	Q2	$71,341	$56,705	$103,741	$103,660	$43,070	$38,547	[illegible]	$108,219
	Q3	$104,219	$78,408	$91,407	$56,204	$70,649	$29,540	[illegible]	[illegible]
	Q4	$43,478	$107,494	$89,127	$85,453	$64,378	[illegible]	[illegible]	$52,238
2009	Q1	$41,344	$77,462	$29,104	$66,449	$62,809	$58,498	[illegible]	$47,086
	Q2	$78,894	$64,045	$163,505	$69,963	$46,951	$21,775	[illegible]	$50,694
	Q3	$157,720	$39,988	$75,475	$82,548	$65,696	$53,205	[illegible]	$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.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	Central	East	South	West	Central	East	South	West
2008	Q1	1 $54,839	$41,843	$54,162	$98,604	$69,738	$20,531	$38,466	$85,317
	Q2	2 $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**, **Year**) are partitioning fields. In addition, **Region** becomes a partitioning field and **Quarter** is the addressing field.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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**, **Quarter**). The partitioning fields are the fields that define the pane (**Category**, **Year**).

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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

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.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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 partition is the combination of **Category**, **Region**, **Year**, and **Quarter**.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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

Individual Fields

The individual dimensions in the view are listed below the option above in the Table Calculation dialog box. Use them to set the addressing to compute using the field you specify. The benefit of this option is that you get absolute control over how the calculation will be computed—if you change the orientation of your view, the table calculation will continue using the same fields for addressing and partitioning. 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, a 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.

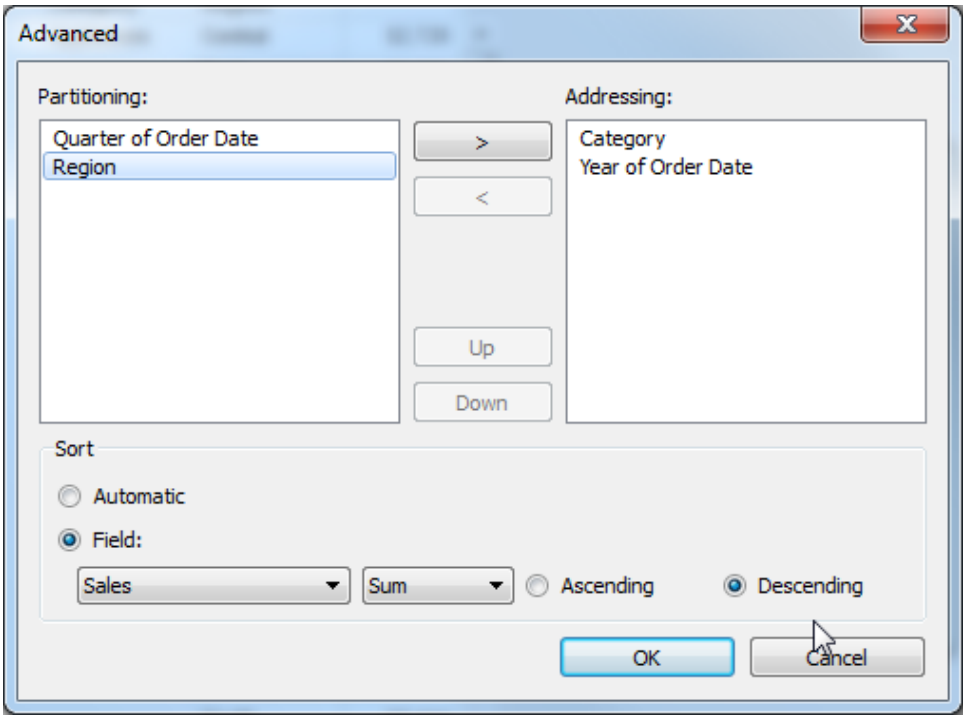

For example, in the view below the addressing fields are set to **Category** and **Year**. These are ordered by **SUM(Sales)**, in descending order (from greatest to least). That means that the combination of **Quarter** and **Region** create the partition. Q1 Central exists four times in the table, and that is the partition.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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 [illegible]	[illegible]	[illegible]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

Because the order is set to **SUM(Sales)**, the calculation is computed based on their **SUM (Sales)** values from highest to lowest.

		Category / Region							
		Furniture				Office Supplies			
Year of Order Date	Quarter of Order Date	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	$47,478	$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

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
    - Rank
    - Percentile
    - Moving Average
    - Year to Date (YTD) Total
    - Compound Growth Rate
    - Year over Year Growth
    - Year to Date (YTD) Growth

After adding a quick table calculation to the view, you can edit its definition by selecting **Edit Table Calculation** from 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:

Right-click the measure you want to use in the computation and choose **Add Table Calculation**.

1. 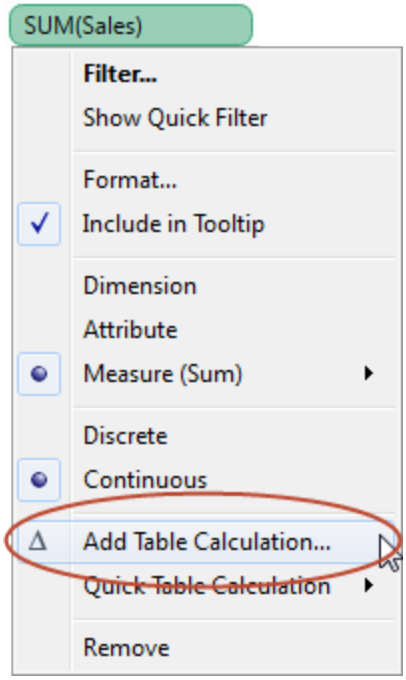

2. In the Table Calculation dialog box, select one of the types of calculations from the drop-down menu at the top.

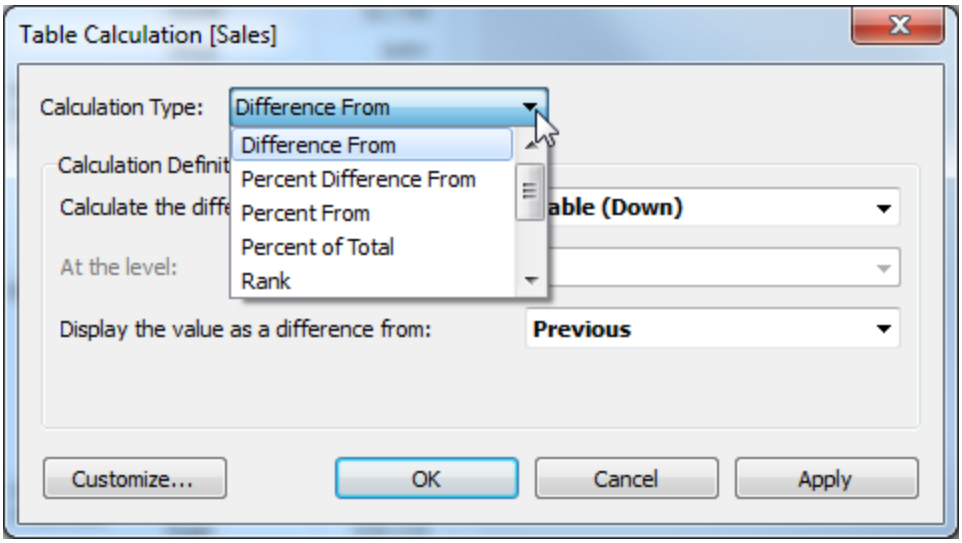

Define the formula using the drop-down lists further down in the dialog box. Learn more about how to define each type of calculation by selecting it from the **Calculation Type** field.

When finished, click **OK**. The measure is now marked as a table calculation (the triangle symbol) and all the relevant values in the view are computed using the table calculation.

3. SUM(Sales) Δ

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.

To define a Difference From calculation, 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.

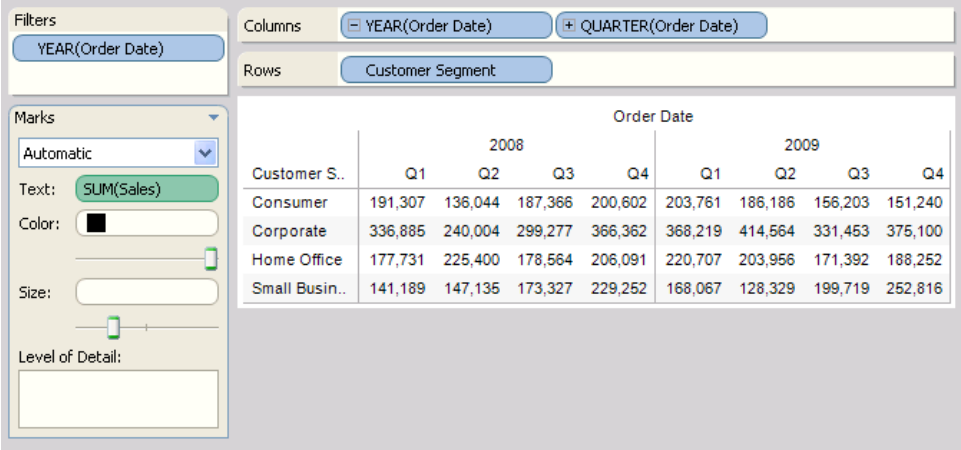

To compute the difference between 2008 and 2009 sales, use the following values in the Table Calculation dialog box.

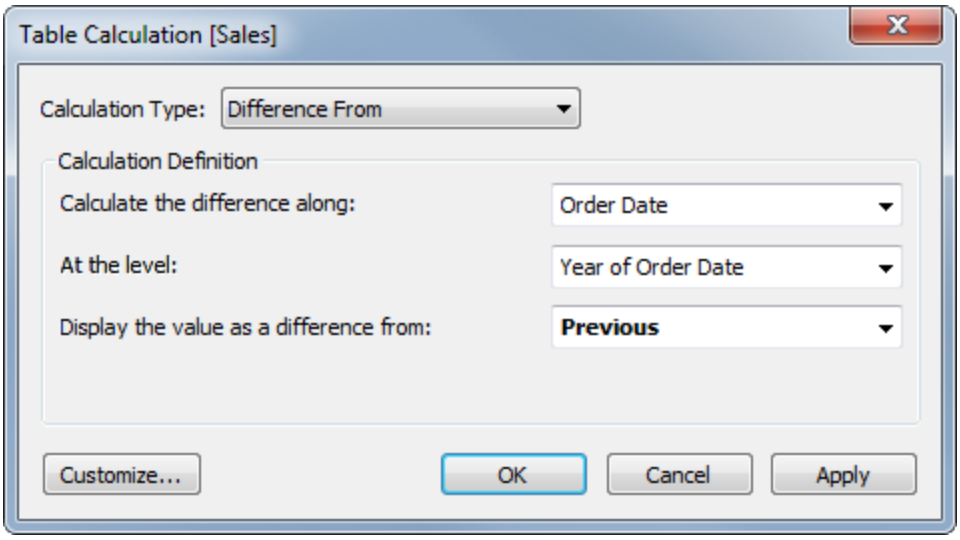

The difference is calculated along the **Order Date** dimension at the level of year because we are comparing 2008 sales to 2009 sales. The table now displays the difference between each quarter in 2009 and the corresponding quarter in the previous year. Notice that there are no values for 2008. That's because there are no previous years to compute the difference from. You can hide that column without affecting the calculation.

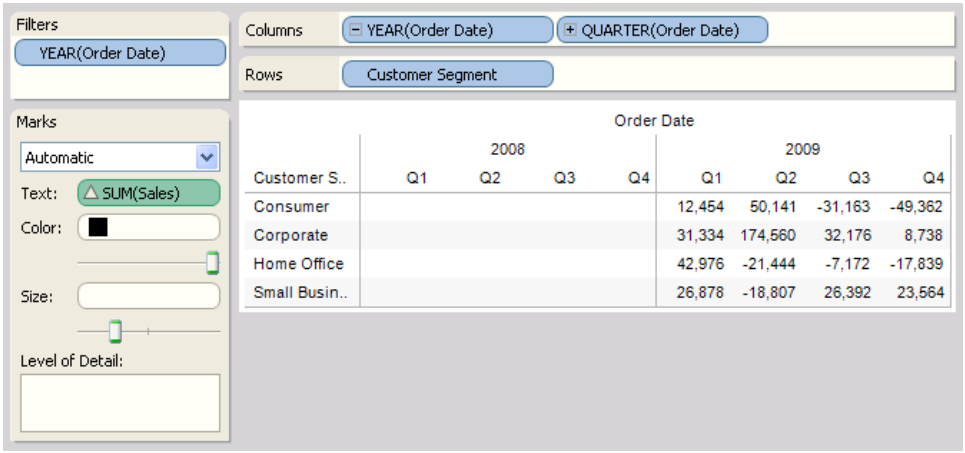

The view below may be more clear. It shows both the Difference From calculation and the **Total Sales** (before the computation). You can see that in the first quarter of 2009 the total sales was $203,761 while in the same quarter in 2008 the total sales was $191,307. The difference between these two values is $12,454.

Columns: + QUARTER(Order Date) | Measure Names

Rows: Customer Segment | + YEAR(Order Date)

		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	203,956	-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

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. See **Quick Table Calculations** on page 602.

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 over year (CAGR).

To define a Percent Difference From calculation 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.

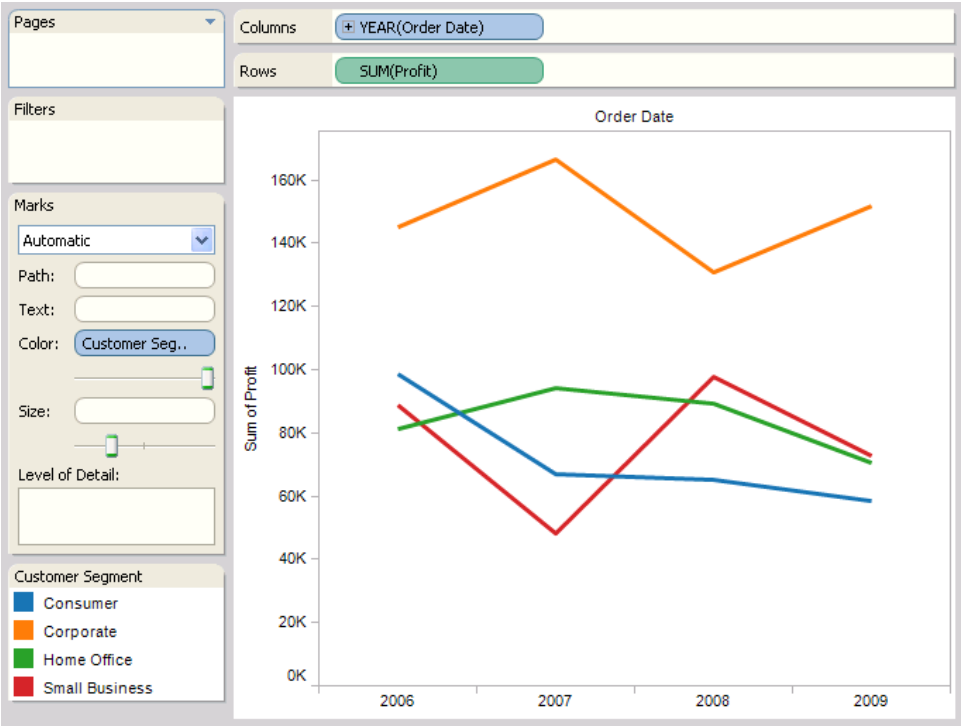

When we view this same view using a Percent Difference From calculation, it becomes clear that Small Business segment rebounded quite dramatically in 2008.

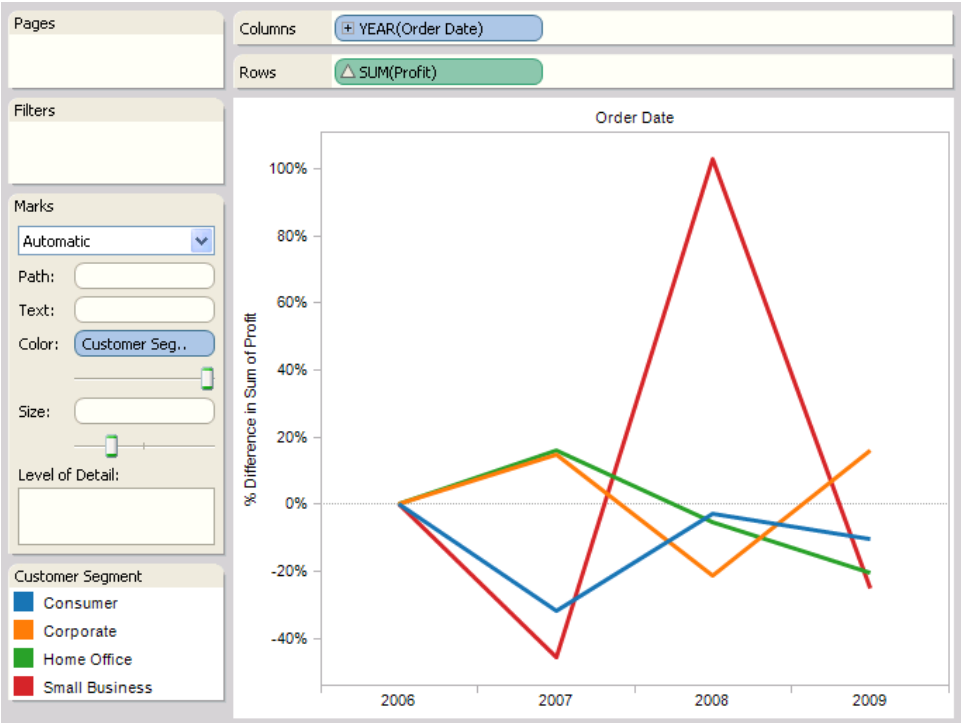

To define a Percent Difference From calculation, select **Percent Difference From** in the **Calculation Type** field in the Table Calculation dialog box.

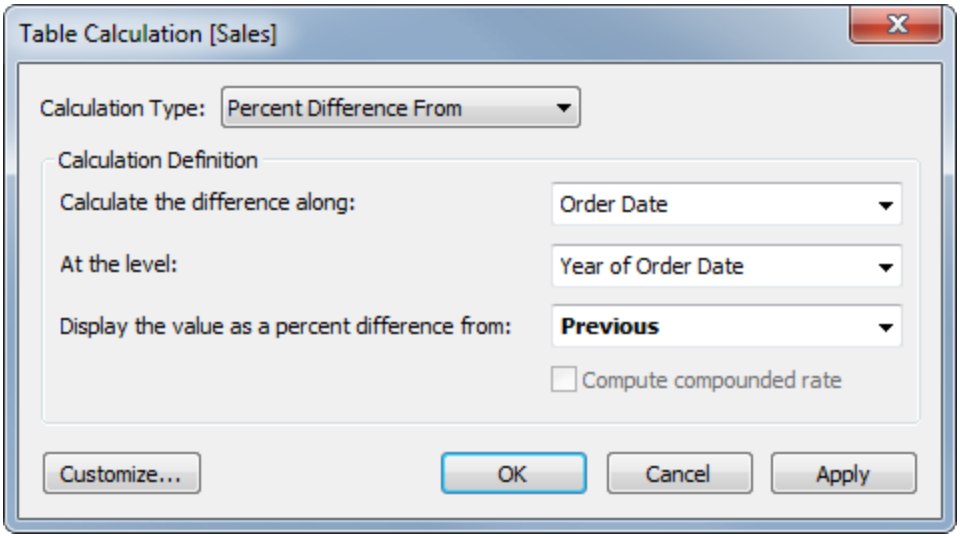

The difference is calculated along the **Order Date** dimension at the Year level because we are comparing year over year profit. Each value in the view is a percent difference from the previous year. The view below shows each year as a difference of 2006 Profit.

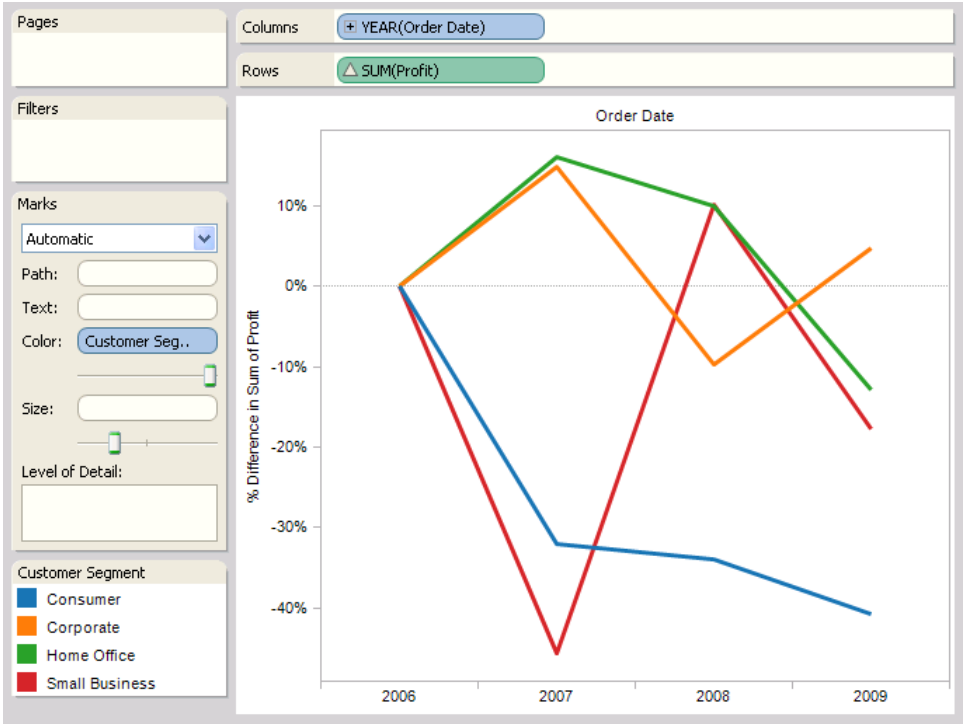

Percent Difference From calculations are commonly used to calculate compound growth rates and year over year growth. You can quickly add these calculations by right-clicking

the measure you want to use in the view and selecting **Percent Difference** from the Quick Table Calculation sub-menu.See **Quick Table Calculations** on page 602.

Percent From Calculation

A Percent From calculation is similar to a 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.

To define a Percent From calculation, 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.

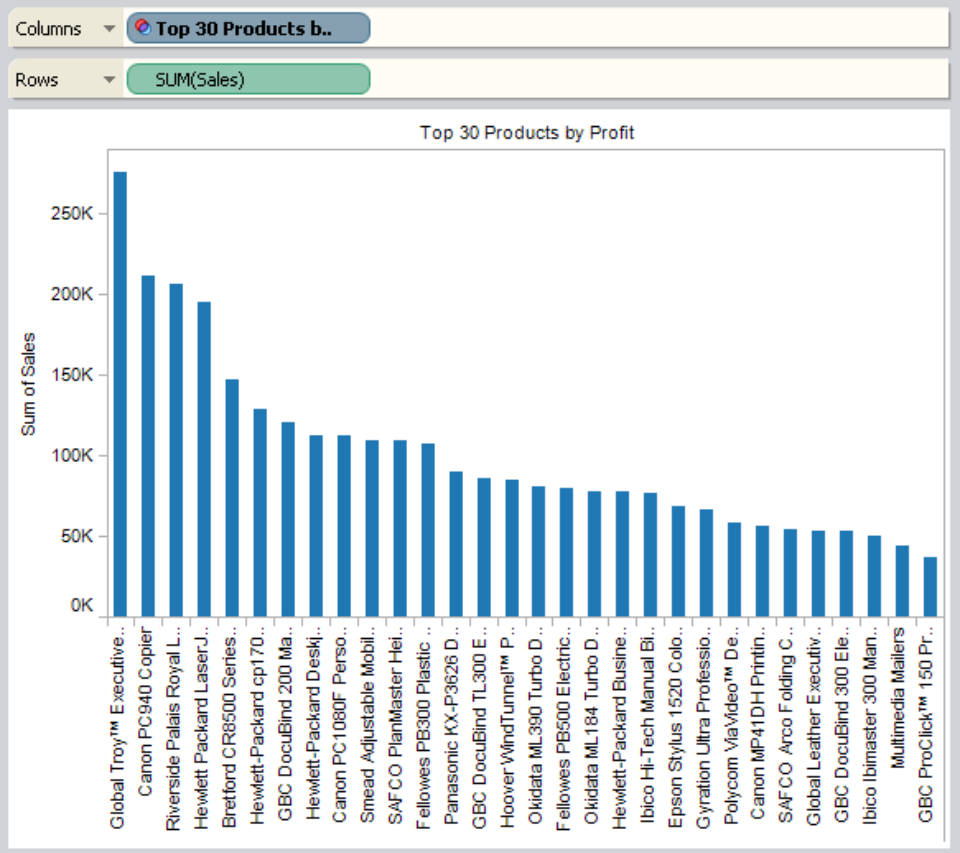

While it is generally clear how each product is performing when compared to the next in terms of sales, you can use a Percent From calculation to compute just how much better each product is from the next. The view below has this type of calculation placed on Detail on the Marks card. The tooltip for each product shows the value of the mark expressed as a

percentage of the next product in the view. For example, you can examine the tooltip for the Global Troy Executive Leather chair and see that it is 130% of the sales for the next item—the Canon Copier.

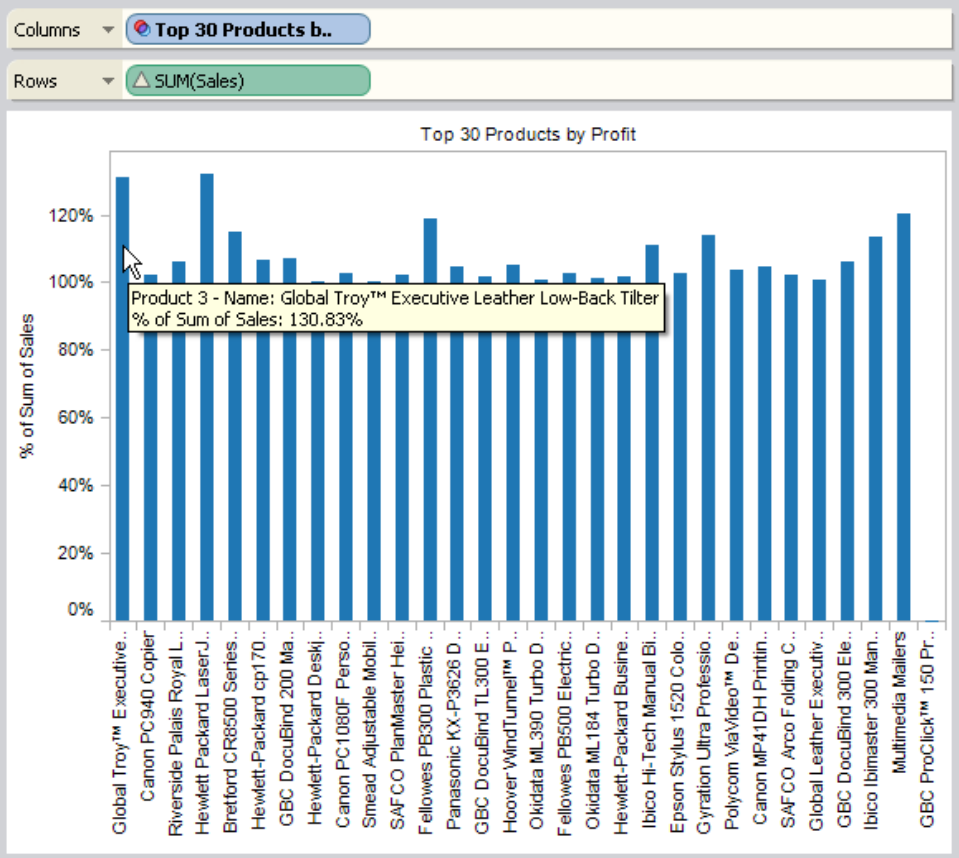

This table calculation was computed using the definition shown in the dialog box below. The percentage is calculated across the rows, and each product is displayed as a percentage of the next product.

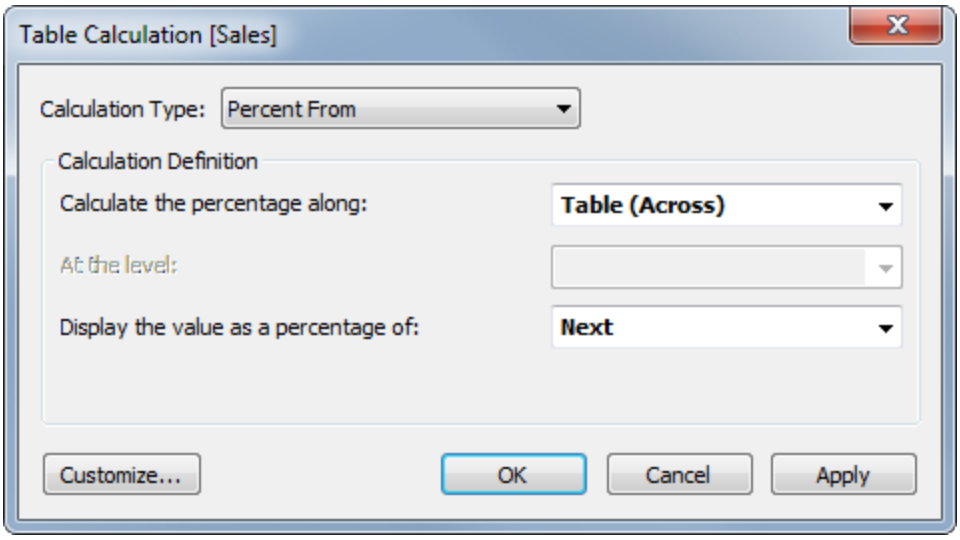

Rank Calculation

Use this type of calculation to compute rankings for a set of values.

To define a Rank calculation, 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.

Tableau ignores nulls in rank calculations. They appear as blank rows.

The following is an example of a Rank calculation.

Example: Rank Calculation

The view below shows sales over a four-year period by region and category.

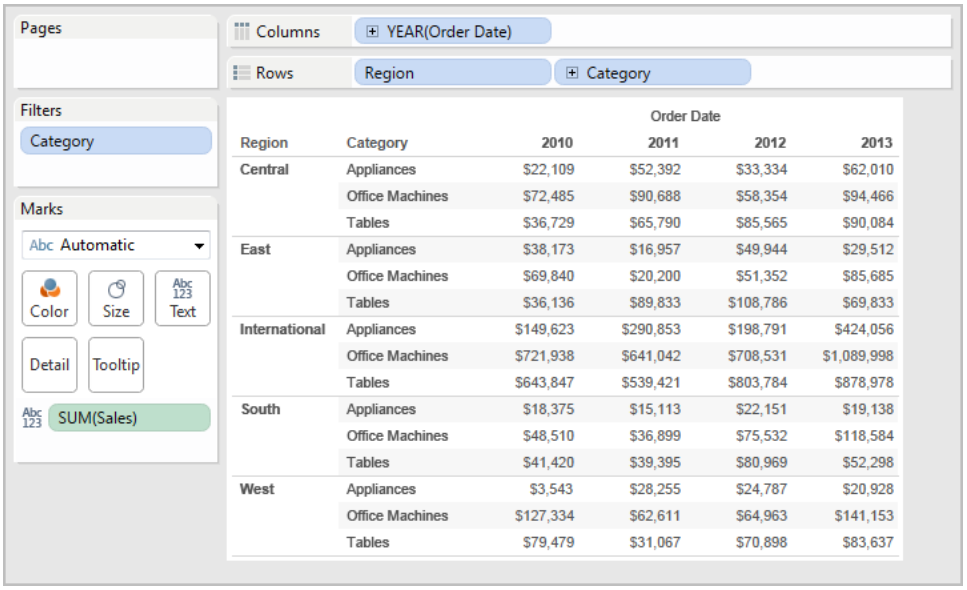

		Order Date			
Region	Category	2010	2011	2012	2013
Central	Appliances	$22,109	$52,392	$33,334	$62,010
	Office Machines	$72,485	$90,688	$58,354	$94,466
	Tables	$36,729	$65,790	$85,565	$90,084
East	Appliances	$38,173	$16,957	$49,944	$29,512
	Office Machines	$69,840	$20,200	$51,352	$85,685
	Tables	$36,136	$89,833	$108,786	$69,833
International	Appliances	$149,623	$290,853	$198,791	$424,056
	Office Machines	$721,938	$641,042	$708,531	$1,089,998
	Tables	$643,847	$539,421	$803,784	$878,978
South	Appliances	$18,375	$15,113	$22,151	$19,138
	Office Machines	$48,510	$36,899	$75,532	$118,584
	Tables	$41,420	$39,395	$80,969	$52,298
West	Appliances	$3,543	$28,255	$24,787	$20,928
	Office Machines	$127,334	$62,611	$64,963	$141,153
	Tables	$79,479	$31,067	$70,898	$83,637

We would like to know how different categories rank from year to year, without regard to region. That is, we would like to rank the 15 values in each row to generate the following result.

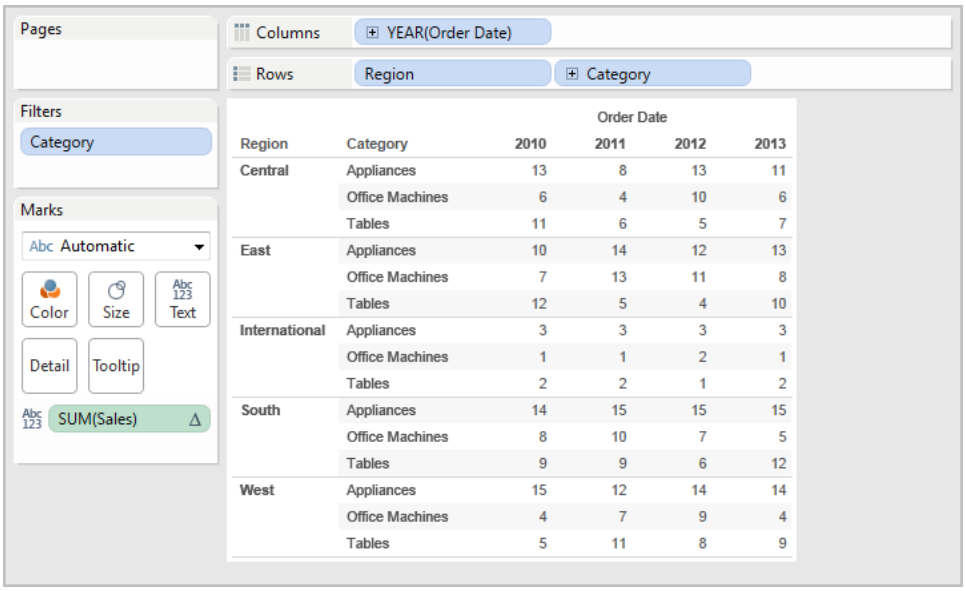

To obtain this view, follow these steps:

1. Click the SUM(Sales) field in the view and select **Add Table Calculation**.
2. In the Table Calc [Sales] dialog box, select **Rank** from the **Calculation Type** drop-down.

   The Table Calc [Sales] dialog box now shows fields relevant for ranking:

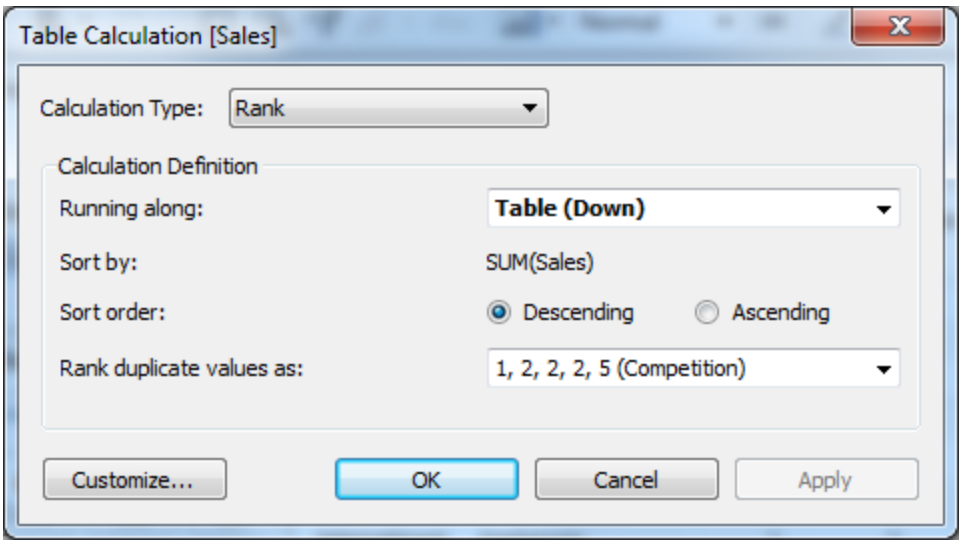

3. Select **Table (Down)** from the **Running along** field.
4. Choose **OK**.

Ranking Duplicate Values

One issue with Rank calculations is that there may be more than one cell in the relevant row or column with the same value. What would happen, for example, if Tables in the Central region

and Appliances in the South region both had sales of exactly $36,729? Tableau lets you specify how to handle such cases by including a **Rank duplicate values as** field in the Table Calculation dialog box when you set **Calculation Type** to Rank.

The choices are listed below. The number sequence at the beginning of each option (for example, 1, 2, 2, 2, 5) shows how the option would rank a hypothetical set of values 10, 15, 15, 15, 20:

Option	Result
1, 2, 2, 2, 5 (Competition)	Duplicate values are all given the same rank, which is the next number in the ranking sequence.The next value after the duplicate values is computed by adding 1 to the number of values already computed.
1, 4, 4, 4, 5 (Modified Competition)	Duplicate values are all given the same rank, which is computed by adding the number of values before the duplicate values to the number of duplicate values. The next value after the duplicate values is computed by adding 1 to the number of values already computed.
1, 2, 2, 2, 3 (Dense)	Duplicate values are all given the same rank, which is the next number in the ranking sequence. The next value after the duplicate values is computed as though the duplicate values were a single value.
1, 2, 3, 4, 5 (Unique)	Duplicate values are given unique rankings, according to the direction in which the ranking is being computed.

Tableau also provides a set of table calculation functions that let you create calculated fields that rank values. See **Table Calculation Functions** on page 899.

> You can add a Rank calculation 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** > **Rank**. By default, this quick calculation will add a competition rank calculation across the rows in the view. See **Quick Table Calculations** on page 602.

Percentile Calculation

Use this type of calculation to compute percentiles for a set of values.

To define a Percentile calculation, 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.

Tableau ignores nulls in percentile calculations. They appear as blank rows in a cross-tab and do not count towards the total number of items used in the calculation (%).

The following is an example of a Percentile calculation.

Example: Percentile Calculation

A school administrator is looking at four-year GPAs for a group of students:

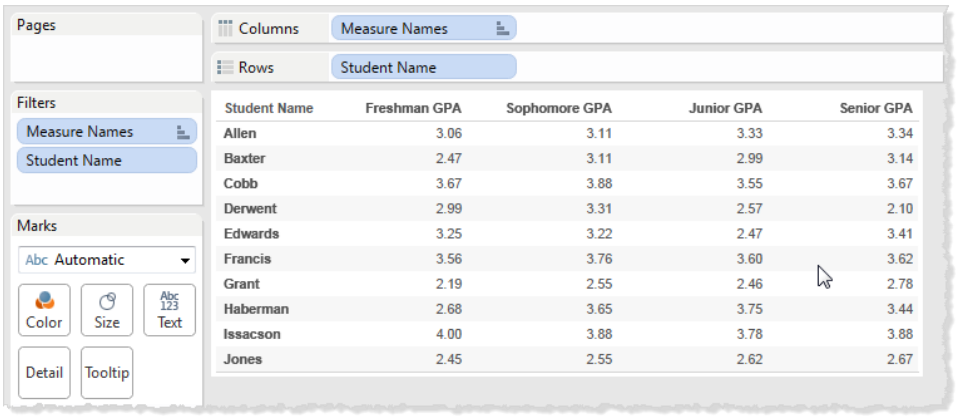

Student Name	Freshman GPA	Sophomore GPA	Junior GPA	Senior GPA
Allen	3.06	3.11	3.33	3.34
Baxter	2.47	3.11	2.99	3.14
Cobb	3.67	3.88	3.55	3.67
Derwent	2.99	3.31	2.57	2.10
Edwards	3.25	3.22	2.47	3.41
Francis	3.56	3.76	3.60	3.62
Grant	2.19	2.55	2.46	2.78
Haberman	2.68	3.65	3.75	3.44
Issacson	4.00	3.88	3.78	3.88
Jones	2.45	2.55	2.62	2.67

The administrator wonders how the students' performance has changed over time, relative to each other. To find out, the administrator applies a Percentile table calculation to the grades for each year, to obtain the following result:

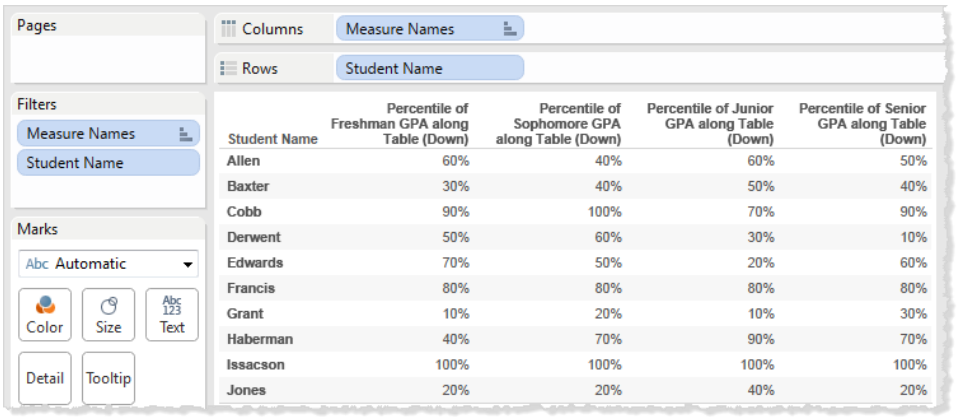

Student Name	Percentile of Freshman GPA along Table (Down)	Percentile of Sophomore GPA along Table (Down)	Percentile of Junior GPA along Table (Down)	Percentile of Senior GPA along Table (Down)
Allen	60%	40%	60%	50%
Baxter	30%	40%	50%	40%
Cobb	90%	100%	70%	90%
Derwent	50%	60%	30%	10%
Edwards	70%	50%	20%	60%
Francis	80%	80%	80%	80%
Grant	10%	20%	10%	30%
Haberman	40%	70%	90%	70%
Issacson	100%	100%	100%	100%
Jones	20%	20%	40%	20%

The administrator can now observe that despite some fluctuations, the students' performance was relatively stable relative to each other during the four years.

To obtain this view, the administrator had to apply a Percentile table calculation to each of four measures: Freshman GPA, Sophomore GPA, Junior GPA, and Senior GPA. For each of these measures the administrator followed these steps:

1. Click the measure in the view and select **Add Table Calculation**.
2. In the Table Calc dialog box, select **Percentile** from the **Calculation Type** drop-down.
3. The dialog box now shows fields relevant for Percentile:

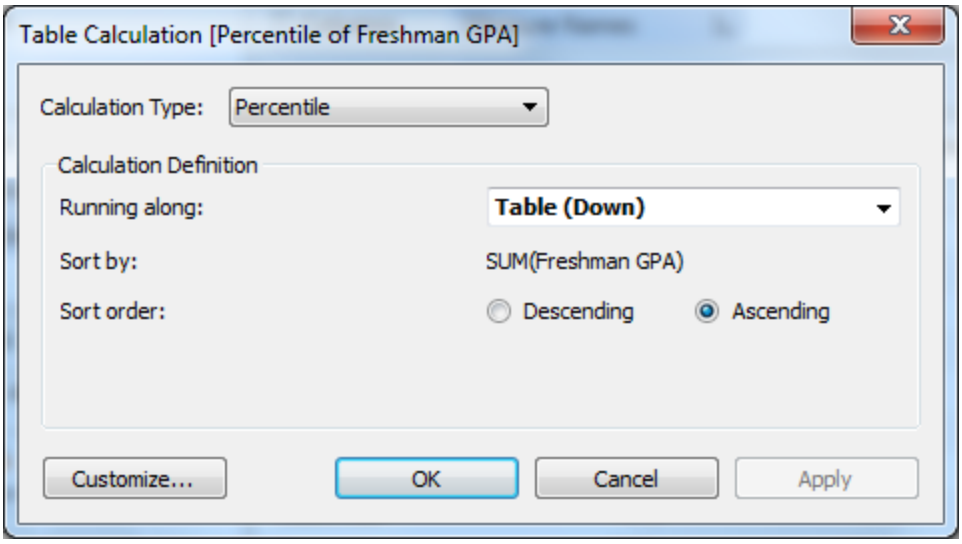

4. Select **Table (Down)** from the **Running along** field.
5. Select **Ascending** from the **Sort Order** field.
6. Click **OK**.

You can add a Percentile calculation 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** > **Percentile**. By default this quick calculation will add a percentile calculation across the rows in the view. See **Quick Table Calculations** on page 602.

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.

Columns: QUARTER(Order Date)

Rows: YEAR(Order Date)

	Order Date			
Year of Orde..	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.

Columns: QUARTER(Order Date)

Rows: YEAR(Order Date), **Measure Names**

		Order Date			
Year of Orde..		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

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.

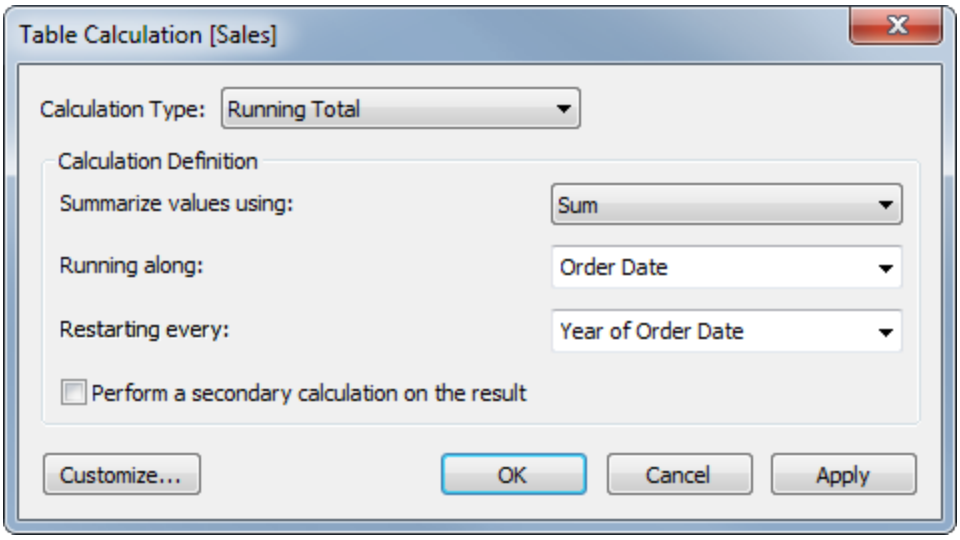

You can add a Running Total calculation to your view easily using the Quick Table Calculations menu. Right-click the measure you want to use in the calculation and select **Quick Table Calculations** > **Running Total**. See **Quick Table Calculations** on page 602.

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 all 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 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.

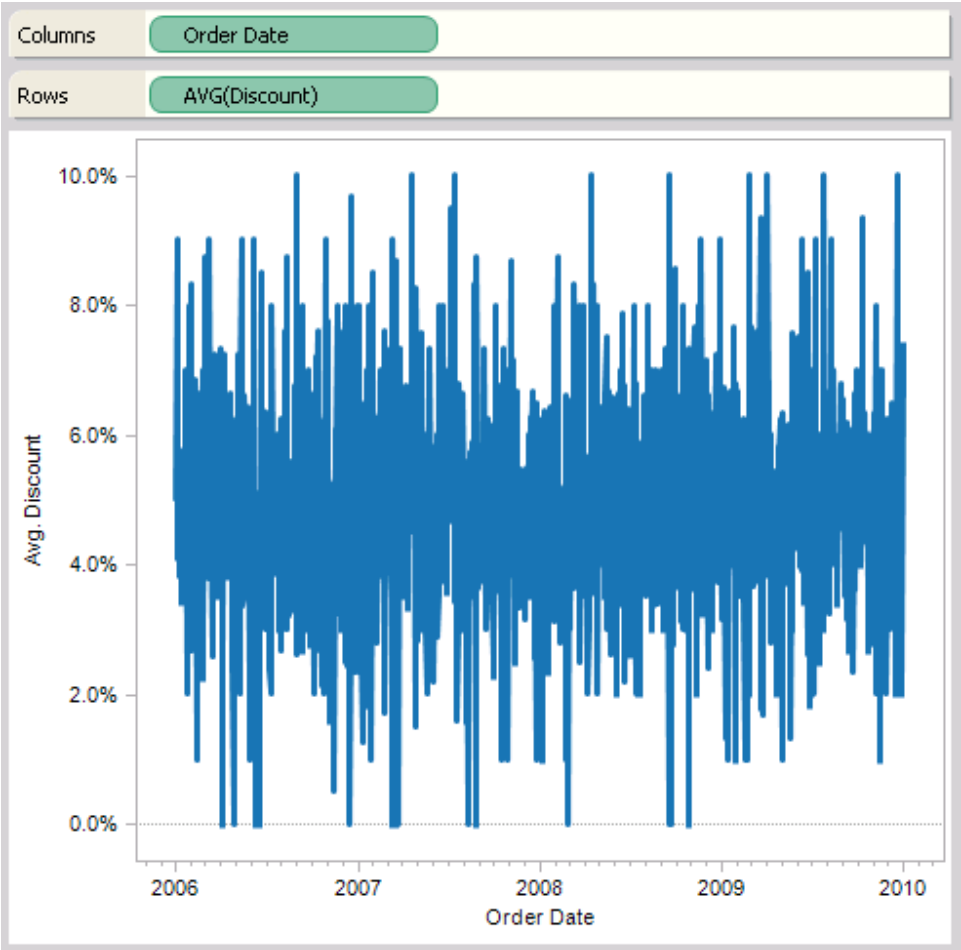

However, if you add a Moving Average, the view becomes much more manageable.

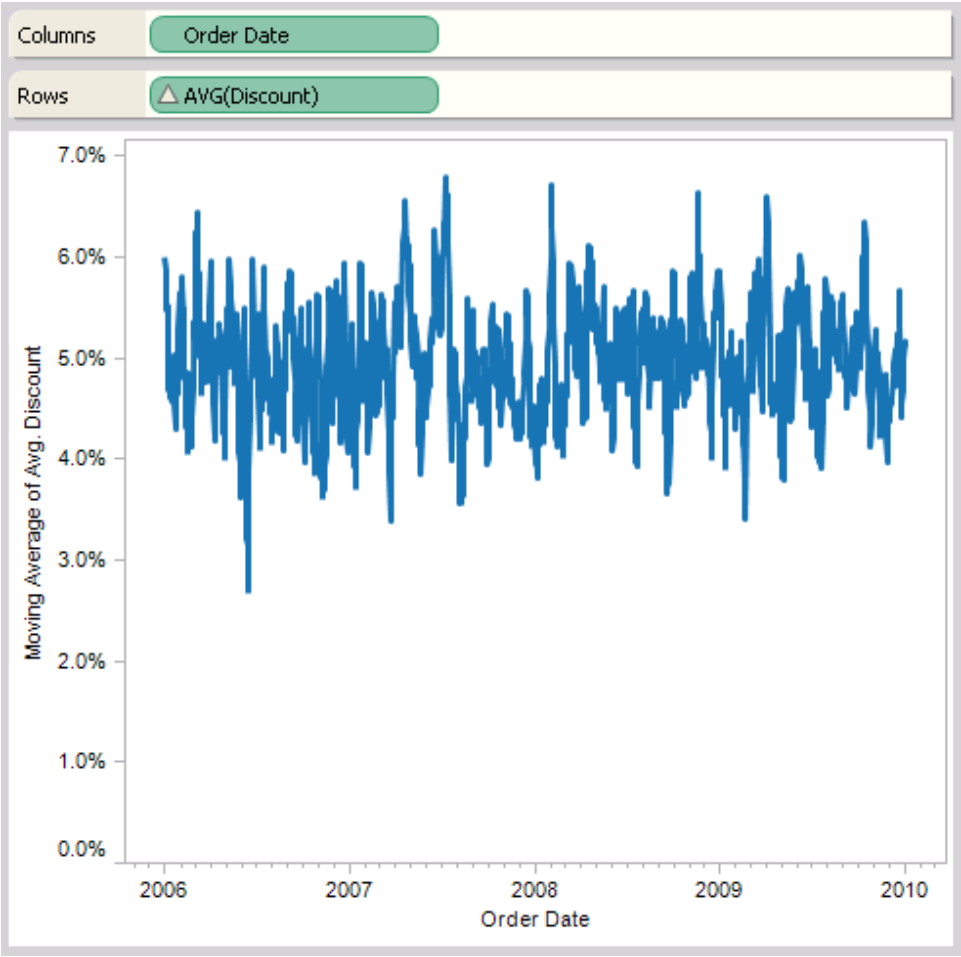

This calculation was defined by the formula shown below. The values are summarized as an average along the rows in the view. Each value is an average of the seven days surrounding the current value (four days before and three days after). Note that we have opted to include the current value.

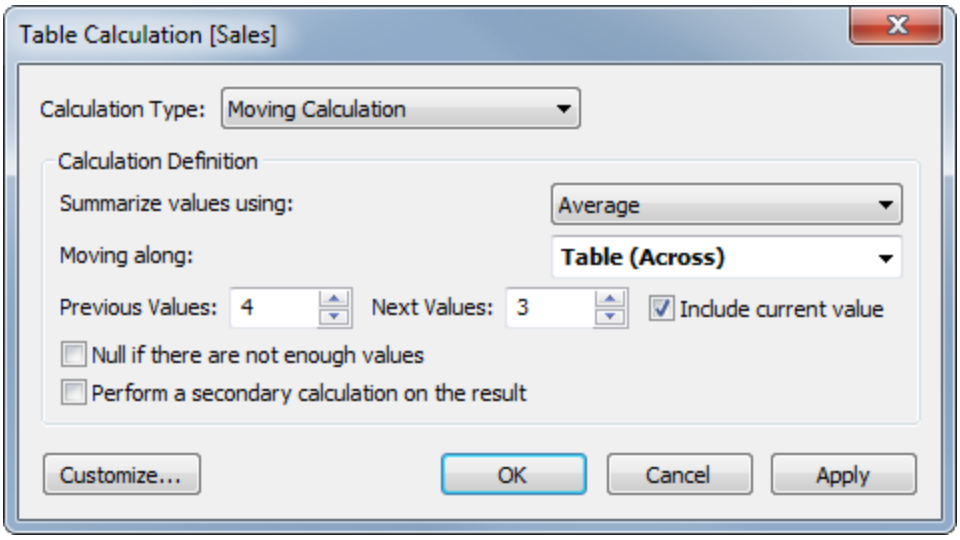

You can add a Moving Average to your view quickly using the Quick Table Calculations menu. Right-click the measure you want to use in the calculation and select **Quick Table Calculations** > **Moving Average**. By default this quick calculation will add a moving average across the rows in the view, summarizing the previous two values including the current value. Refer to **Quick Table Calculations** on page 602.

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. 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 when the **Calculation Type** is Running Total or Moving Calculation. Select **Perform secondary calculation** on the result in the Table Calculation dialog box.

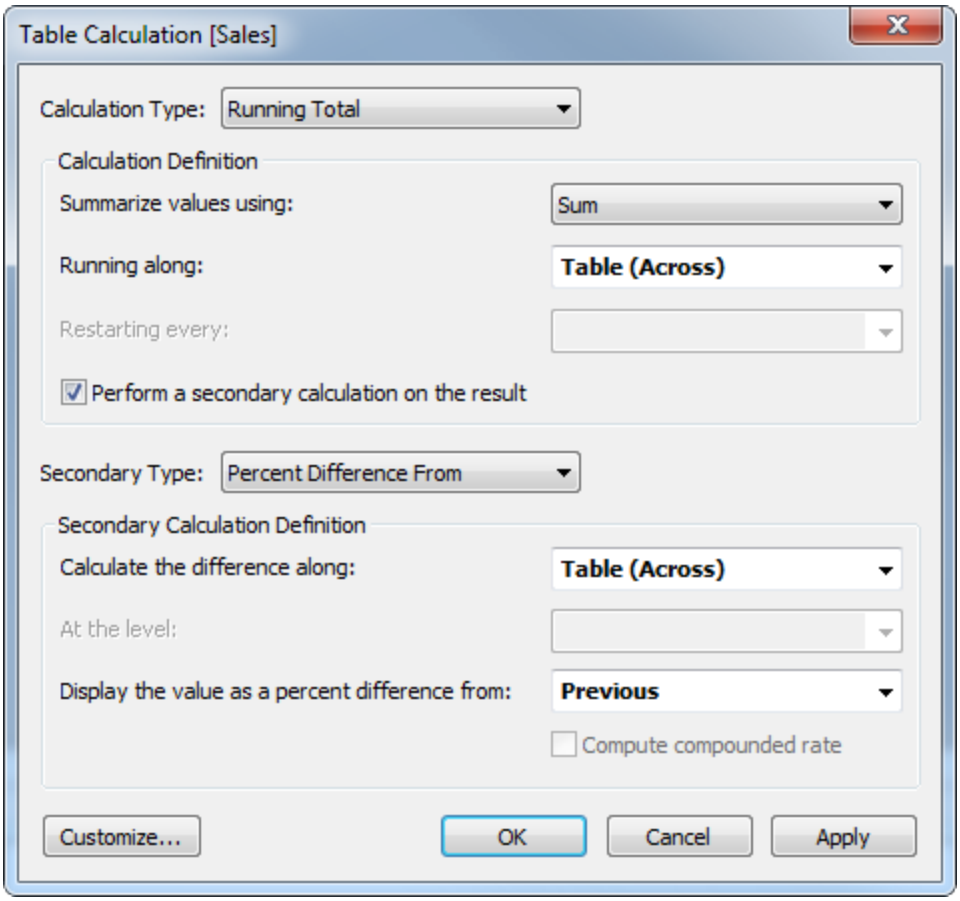

Customizing Table Calculations

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.

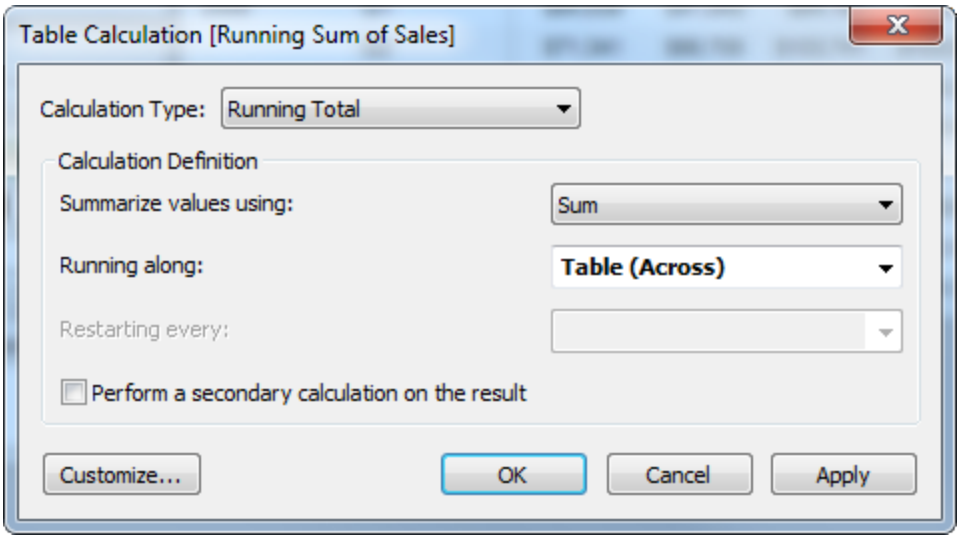

When you click **Customize**, the Calculated Field dialog box opens showing the formula for the calculation. You can see that it uses special functions.

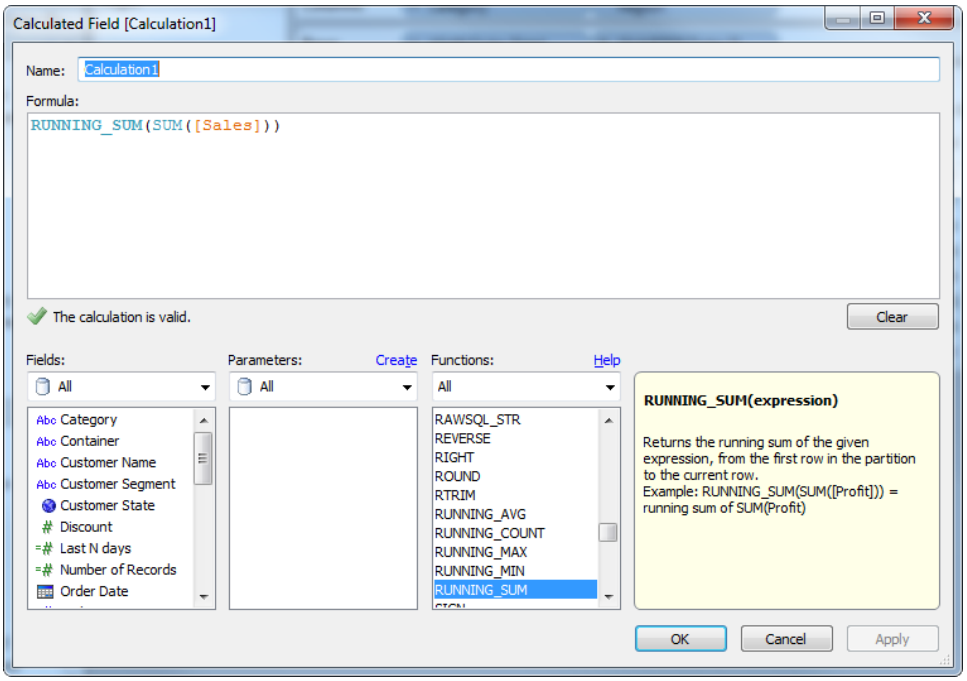

After you customize the calculated field, the changes are not saved until you click **OK** in the Calculated Field dialog box and in the Table Calculation dialog box. The new table calculation field is added to the Data window.

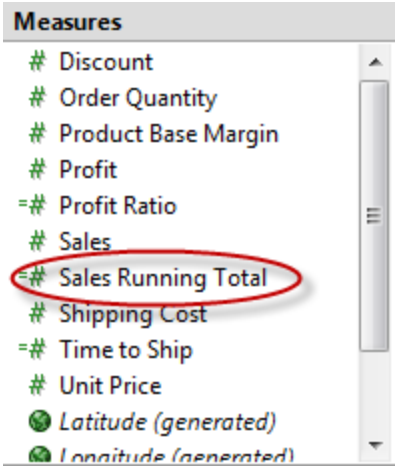

When you use that field in other views, it uses the default addressing and partitioning specified when the table calculation was created. You can change the addressing by right-clicking the field and selecting an option from the Compare To context menu.

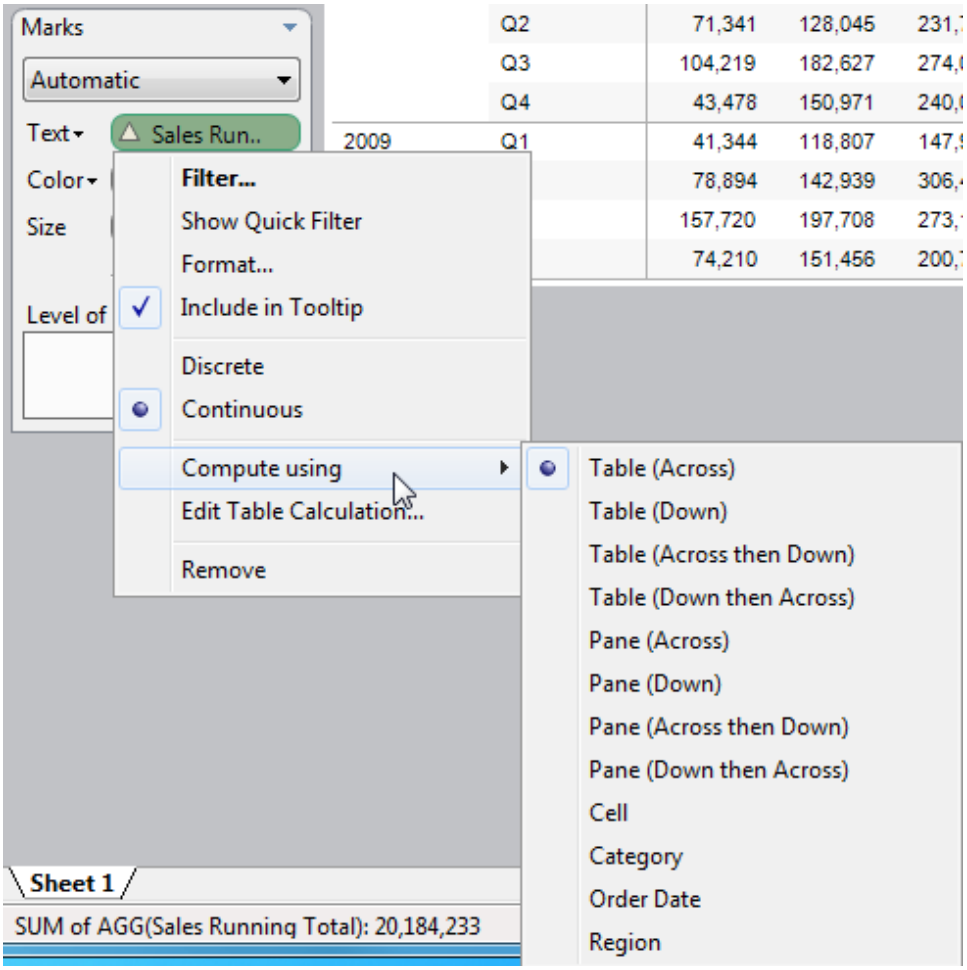

Table calculation formulas must use aggregated data throughout the entire formula. When you are referring to a dimension in the formula you can use the MIN([Dimension]) or MAX ([Dimension]) aggregations. However, this trick only works if the view is grouped by that dimension. That is, when there is a 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	Drop
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.

> 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.

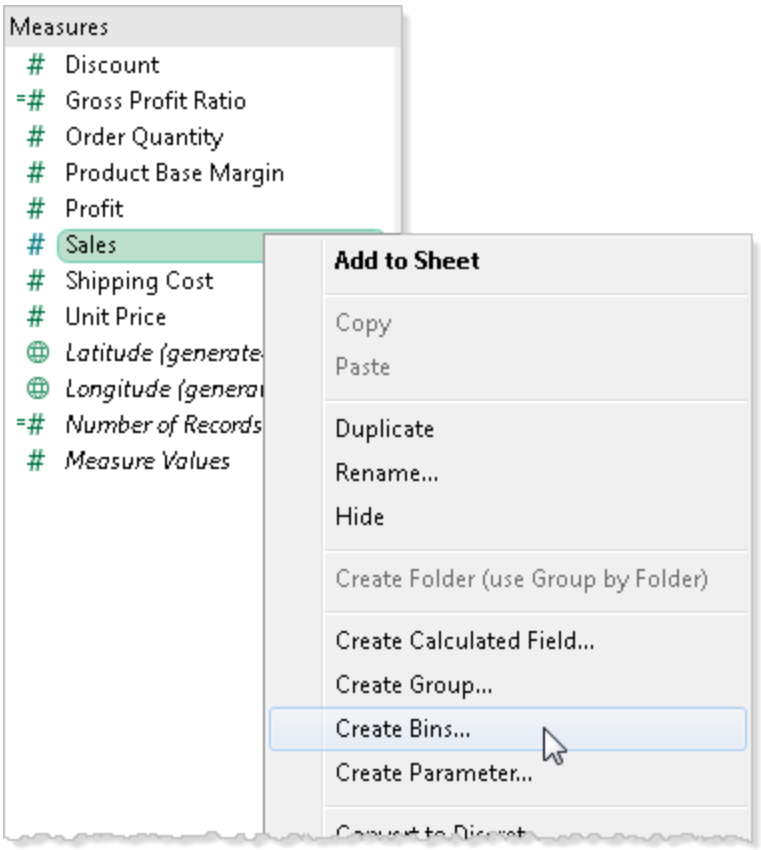

2. Complete the Create Bins dialog box.

   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.

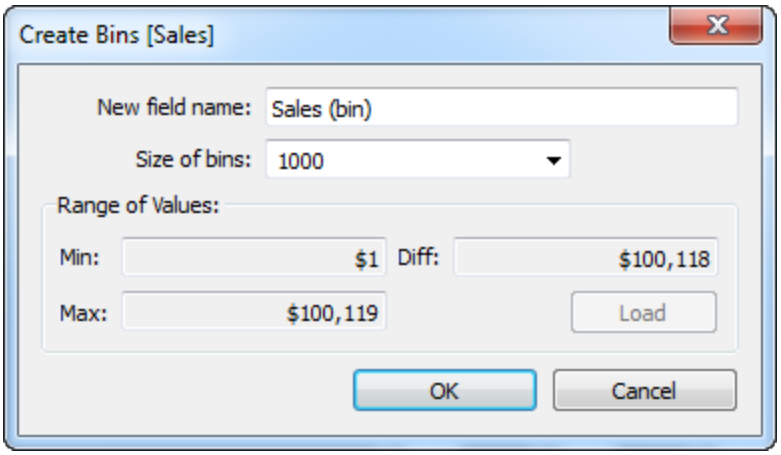

The binned field appears in the **Dimensions** area of the Data window because the bins are treated as discrete categories.

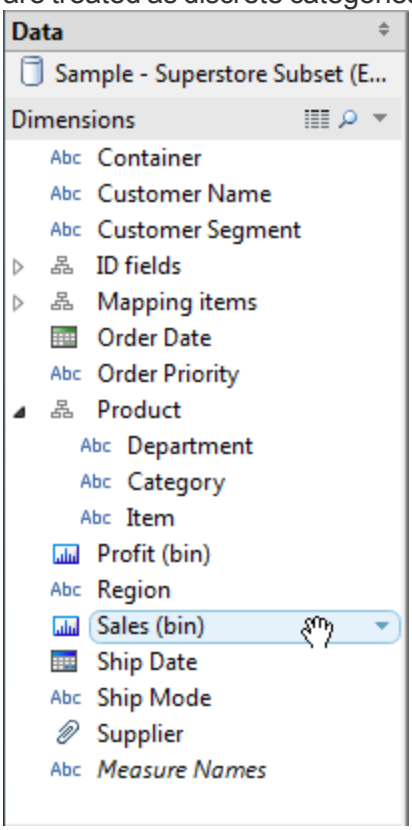

3. Place the **Sales** measure on the **Rows** shelf.

   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.

Notice that all bins are of equal size. If you want to create variable sized bins, you can create a calculation using the CASE function.

The view is shown below.

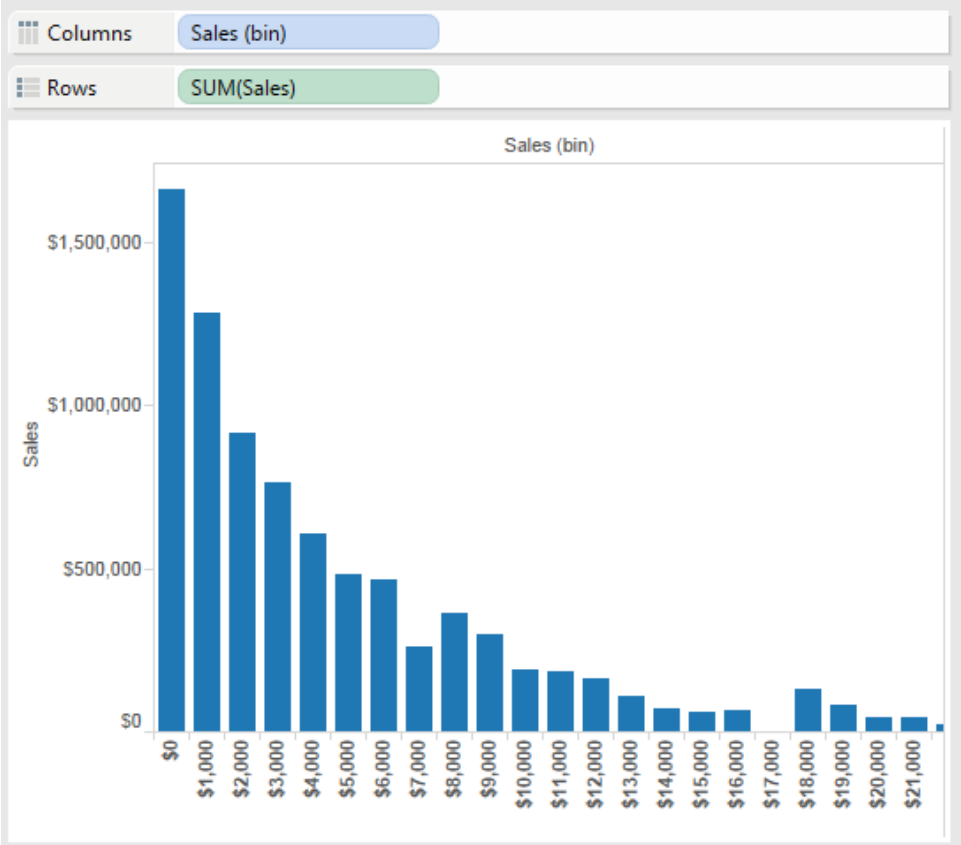

Each bin acts as an equal-sized container that summarizes data for a specific range of values. Each bin label designates the lower limit of the range of numbers that is assigned to the bin. Note that the lower limit is inclusive. For example, the bin labeled **1K** contains numbers greater than or equal to 1,000, but less than 2,000.

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. In Tableau, multidimensional data sources are supported only in Windows.

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 a Microsoft Analysis Services 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 for the measure you are using to calculate the subtotals.

   If you are not displaying totals, this menu item does not appear. See **How to Turn on Grand Totals** on page 630.

2. Then select an aggregation—any value other than **Server**.

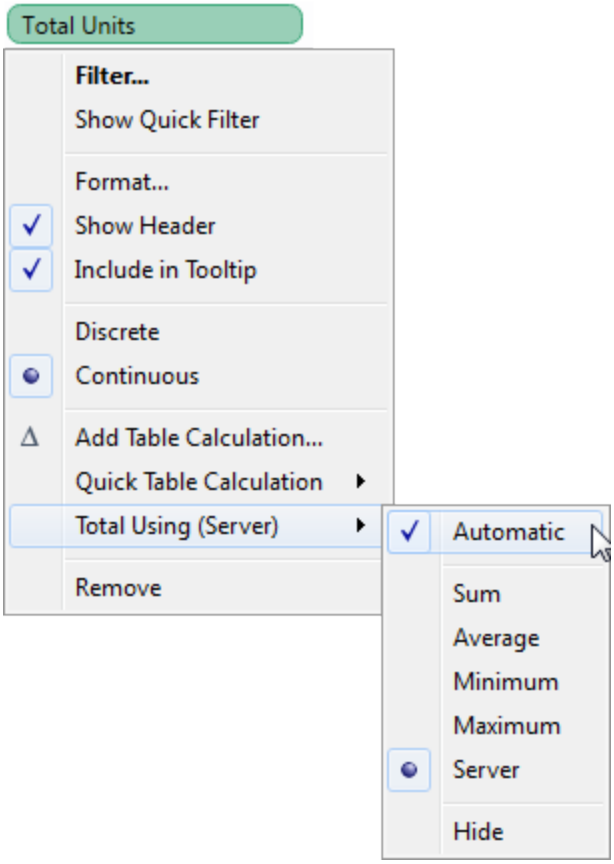

Server tells Tableau to use the underlying data in the data source—a server computation.

Hide tells Tableau to not show a total for the field.

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 total sales for each department and category for four years, you can turn on grand totals to also see the totals across all products and all years.

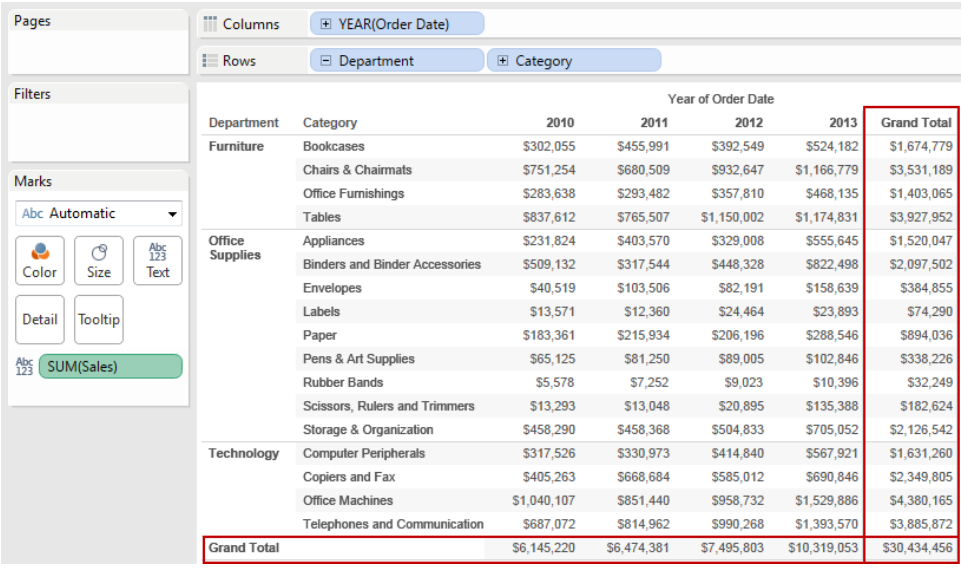

		Year of Order Date				
Department	Category	2010	2011	2012	2013	Grand Total
Furniture	Bookcases	$302,055	$455,991	$392,549	$524,182	$1,674,779
	Chairs & Chairmats	$751,254	$680,509	$932,647	$1,166,779	$3,531,189
	Office Furnishings	$283,638	$293,482	$357,810	$468,135	$1,403,065
	Tables	$837,612	$765,507	$1,150,002	$1,174,831	$3,927,952
Office Supplies	Appliances	$231,824	$403,570	$329,008	$555,645	$1,520,047
	Binders and Binder Accessories	$509,132	$317,544	$448,328	$822,498	$2,097,502
	Envelopes	$40,519	$103,506	$82,191	$158,639	$384,855
	Labels	$13,571	$12,360	$24,464	$23,893	$74,290
	Paper	$183,361	$215,934	$206,196	$288,546	$894,036
	Pens & Art Supplies	$65,125	$81,250	$89,005	$102,846	$338,226
	Rubber Bands	$5,578	$7,252	$9,023	$10,396	$32,249
	Scissors, Rulers and Trimmers	$13,293	$13,048	$20,895	$135,388	$182,624
	Storage & Organization	$458,290	$458,368	$504,833	$705,052	$2,126,542
Technology	Computer Peripherals	$317,526	$330,973	$414,840	$567,921	$1,631,260
	Copiers and Fax	$405,263	$668,684	$585,012	$690,846	$2,349,805
	Office Machines	$1,040,107	$851,440	$958,732	$1,529,886	$4,380,165
	Telephones and Communication	$687,072	$814,962	$990,268	$1,393,570	$3,885,872
Grand Total		$6,145,220	$6,474,381	$7,495,803	$10,319,053	$30,434,456

How to Turn on Grand Totals

You can calculate grand totals by selecting one of the **Grand Totals** options on the **Analysis > Totals** menu. The grand totals are added as an additional row or column to your table.

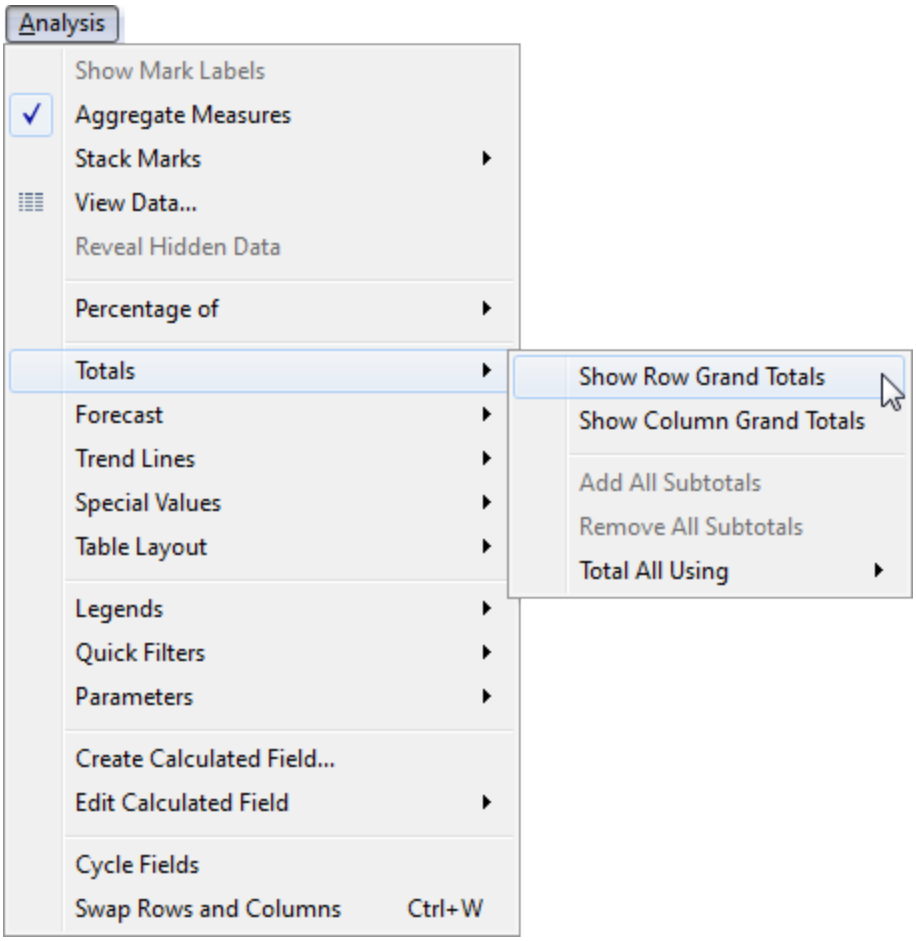

The following rules dictate whether you can turn on grand totals:

- 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. See **Grand Totals and Aggregations** on page 634 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.

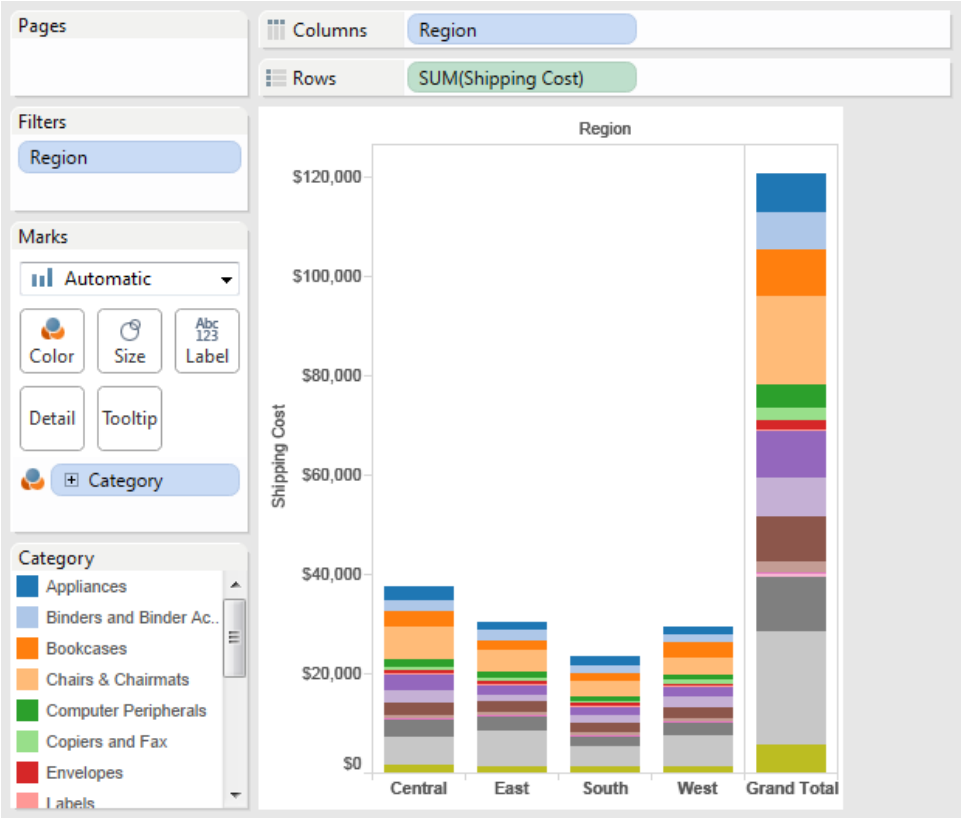

> By default totals are computed on the server if you are connected to a Microsoft Analysis Services data source and locally if you are connected to an Essbase data source using the aggregation specified in the cube. Refer to **Local vs. Server Computation** on page 628 to learn more.

Options for Calculating Grand Totals

When you first turn on grand totals (see **How to Turn on Grand Totals** on the previous page), the totals are computed using disaggregated data in the underlying data source. Consider the following view:

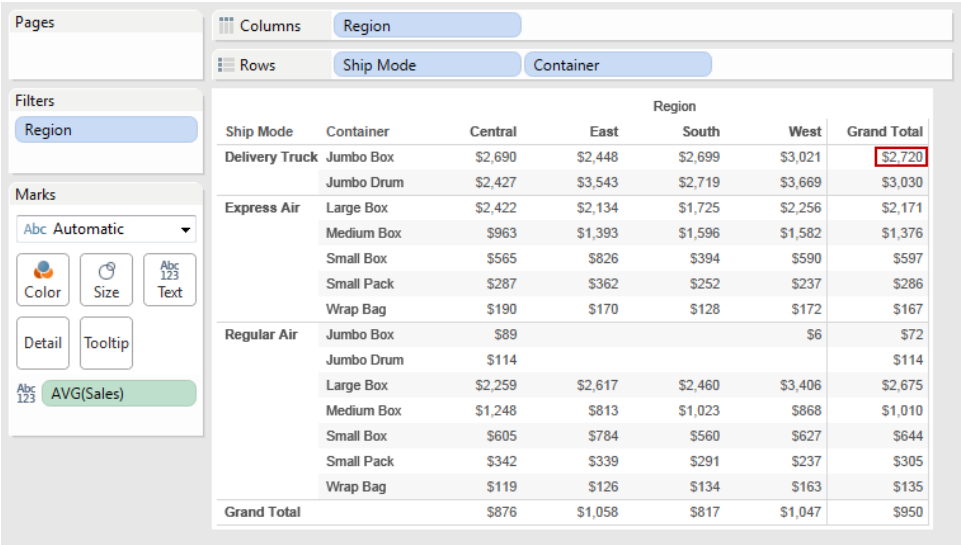

		Region				
Ship Mode	Container	Central	East	South	West	Grand Total
Delivery Truck	Jumbo Box	$2,690	$2,448	$2,699	$3,021	$2,720
	Jumbo Drum	$2,427	$3,543	$2,719	$3,669	$3,030
Express Air	Large Box	$2,422	$2,134	$1,725	$2,256	$2,171
	Medium Box	$963	$1,393	$1,596	$1,582	$1,376
	Small Box	$565	$826	$394	$590	$597
	Small Pack	$287	$362	$252	$237	$286
	Wrap Bag	$190	$170	$128	$172	$167
Regular Air	Jumbo Box	$89			$6	$72
	Jumbo Drum	$114				$114
	Large Box	$2,259	$2,617	$2,460	$3,406	$2,675
	Medium Box	$1,248	$813	$1,023	$868	$1,010
	Small Box	$605	$784	$560	$627	$644
	Small Pack	$342	$339	$291	$237	$305
	Wrap Bag	$119	$126	$134	$163	$135
Grand Total		$876	$1,058	$817	$1,047	$950

The average that you see at the right of the first row under Grand Total is $2720. But if you compute the average for the four values you see in the row ($2690, $2448, $2699, $3021), the result comes to $2714.50, not $2720. The discrepancy is due to the fact that Tableau is averaging the data in the underlying data source. In this case there are more than four numbers to average, perhaps many more. The result is derived by averaging all values that have the properties Ship Mode = Delivery Truck, Container = Jumbo Box, without regard to region.

To see totals that correspond to the numbers you see in your view requires two averaging operations in Tableau: first, values for the individual marks (or cells) in the view must be derived—for example, by averaging all values that have the properties Ship Mode = Delivery Truck, Container = Jumbo Box, and Region = Central. Then the results for each region must be derived, by averaging the results for the individual marks. Fortunately, you do not need to perform two operations. To display a result of this type, from the **Analysis** menu choose **Totals** > **Total All Using** > **Average** . Now the average is performed on the values you see, and not on the disaggregated date in the data source:

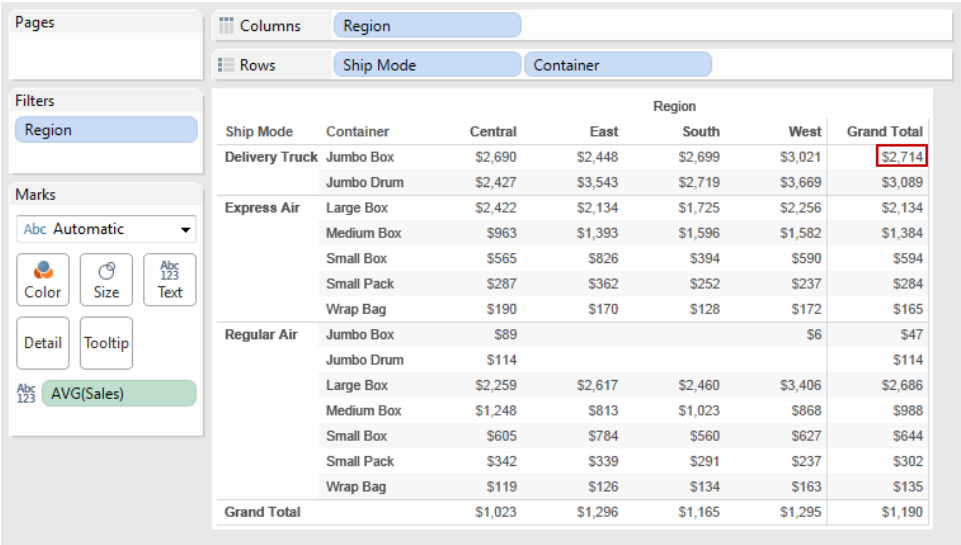

		Region				
Ship Mode	Container	Central	East	South	West	Grand Total
Delivery Truck	Jumbo Box	$2,690	$2,448	$2,699	$3,021	$2,714
	Jumbo Drum	$2,427	$3,543	$2,719	$3,669	$3,089
Express Air	Large Box	$2,422	$2,134	$1,725	$2,256	$2,134
	Medium Box	$963	$1,393	$1,596	$1,582	$1,384
	Small Box	$565	$826	$394	$590	$594
	Small Pack	$287	$362	$252	$237	$284
	Wrap Bag	$190	$170	$128	$172	$165
Regular Air	Jumbo Box	$89			$6	$47
	Jumbo Drum	$114				$114
	Large Box	$2,259	$2,617	$2,460	$3,406	$2,686
	Medium Box	$1,248	$813	$1,023	$868	$988
	Small Box	$605	$784	$560	$627	$644
	Small Pack	$342	$339	$291	$237	$302
	Wrap Bag	$119	$126	$134	$163	$135
Grand Total		$1,023	$1,296	$1,165	$1,295	$1,190

This kind of total is sometimes referred to as a two-pass total, because the average you see in the grand total column is aggregated twice--once to derive the column or row value, and then again across column or rows to derive the grand total.

Configuring Grand Totals

When grand totals are turned on in the view (either Row grand totals or Column grand totals), you can specify how totals should be computed.

To configure grand totals, from the **Analysis** menu choose **Totals** > **Total All Using** to display a submenu:

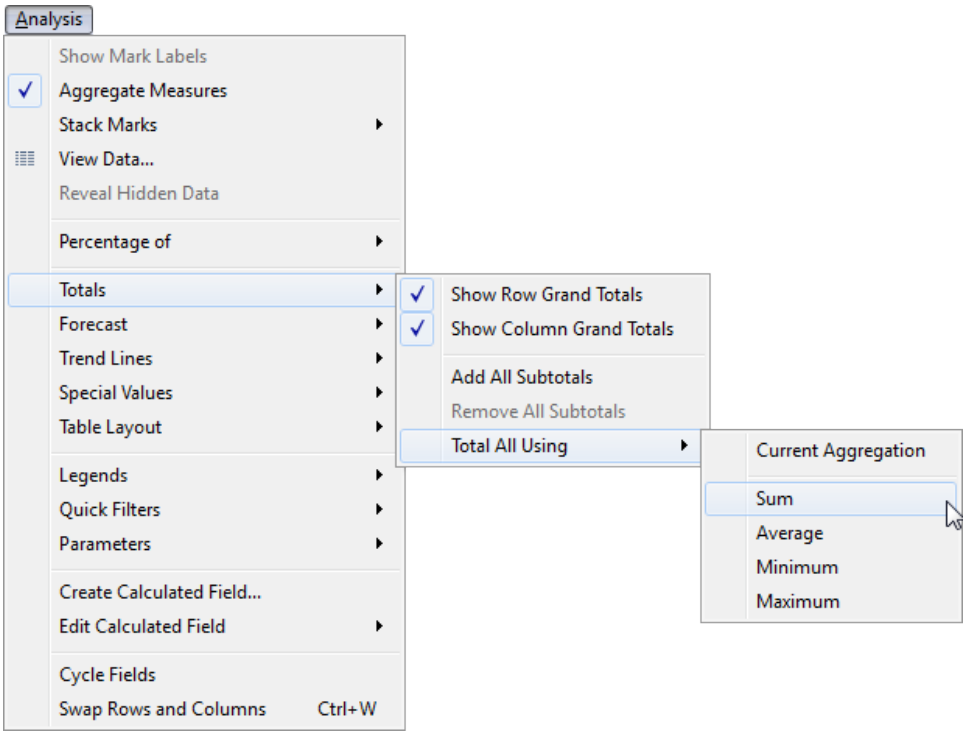

When you choose **Current Aggregation** from the submenu, totals are computed as they always were in earlier versions of Tableau: with column and row totals based on the configured aggregation for the field. In this case, computations are based on the underlying data, which is disaggregated, and not on the data in the view. See **Options for Calculating Grand Totals** on page 631. And for details on how Tableau computes totals using the current aggregation, see **Grand Totals and Aggregations** below.

When you choose any of the other values (**Sum**, **Average**, **Minimum**, or **Maximum**), all totals are computed using the selected aggregation. The computations are performed on the aggregated data you see in the view.

You can also configure an aggregation for a specific field in the view. To do this, right-click a field in the view and choose **Total Using**. With the first value, **Automatic**, Tableau uses the currently configured aggregation for the field. If you choose any of the other values (which are other types of aggregation), that aggregation is used for totals.

Grand Totals and Aggregations

When you turn on grand totals, the initial values are computing using the current aggregation for the fields in the view. In this case, totals are based on the underlying data rather than the data in the view.

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 details, See **Configuring Grand Totals** on the previous page. You can verify any calculation such as an aggregation or a grand total by viewing the underlying disaggregated data.

The following table summarizes the standard aggregations and the grand totals that are calculated by default when, from the Analysis menu, **Totals > Total All Using** is set to the default value **Automatic**.

Aggregation	Calculation Description
Sum	Shows the sum of the values shown in the row or column.
Average	Shows the average of the values shown in the row or column.
Median	Shows the median for the values shown in the row or column.
Count; Count Distinct	Shows how many values or distinct values are displayed in the rows and columns in the view.
Minimum	Shows the minimum value shown in the row or column.
Maximum	Shows the maximum value shown in the row or column.
Percentile	Shows the average percentile for all values shown in the row or column.
Standard	The grand total using standard deviation is the standard devi-

Aggregation	Calculation Description
Deviation	ation of the values shown in the row or column.
Variance	The grand total using variance is not the variance of the rows and columns in which they reside but rather of the underlying data behind the row or column.

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.

You can add subtotals for all fields or for selected fields. To add subtotals for all fields, choose **Totals > Add All Subtotals** from the **Analysis** menu. You then have the option of turning subtotals off for one or more individual fields. You do this by toggling subtotals off, as per the inststructions below for turning subtotals off for individual fields.

To calculate subtotals for an individual field, right-click the field in the view and choose **Subtotals** from the context menu.

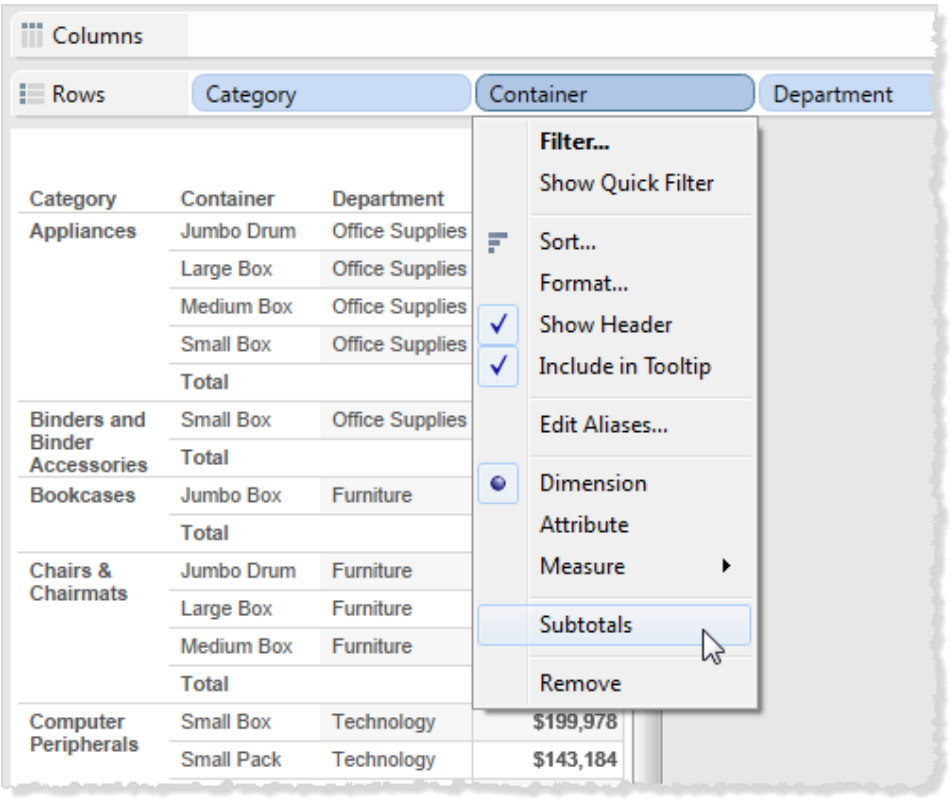

A check mark then appears next to Subtotals on the context menu.

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.

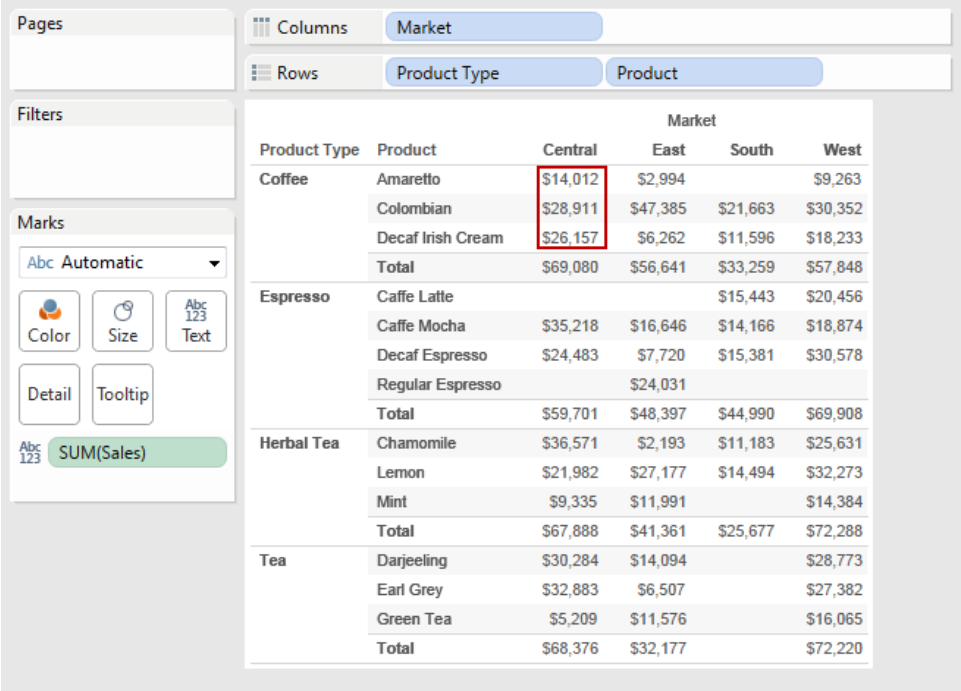

		Market			
Product Type	Product	Central	East	South	West
Coffee	Amaretto	$14,012	$2,994		$9,263
	Colombian	$28,911	$47,385	$21,663	$30,352
	Decaf Irish Cream	$26,157	$6,262	$11,596	$18,233
	Total	$69,080	$56,641	$33,259	$57,848
Espresso	Caffe Latte			$15,443	$20,456
	Caffe Mocha	$35,218	$16,646	$14,166	$18,874
	Decaf Espresso	$24,483	$7,720	$15,381	$30,578
	Regular Espresso		$24,031		
	Total	$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
	Total	$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
	Total	$68,376	$32,177		$72,220

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.

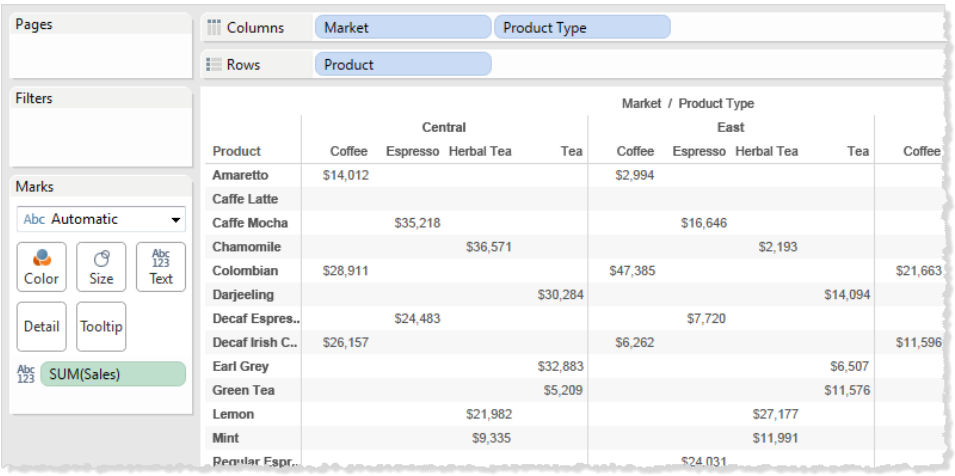

	Market / Product Type								
	Central				East				
Product	Coffee	Espresso	Herbal Tea	Tea	Coffee	Espresso	Herbal Tea	Tea	Coffee
Amaretto	$14,012				$2,994				
Caffe Latte									
Caffe Mocha		$35,218				$16,646			
Chamomile			$36,571				$2,193		
Colombian	$28,911				$47,385				$21,663
Darjeeling				$30,284				$14,094	
Decaf Espres..		$24,483				$7,720			
Decaf Irish C..	$26,157				$6,262				$11,596
Earl Grey				$32,883				$6,507	
Green Tea				$5,209				$11,576	
Lemon			$21,982				$27,177		
Mint			$9,335				$11,991		
Regular Espr..						$24,031			

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.

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.

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. See **Table Calculations** on page 593 for more information.

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 the next page 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 640 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

		Order Date			
Product 1 - Category	Region	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 563 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 586 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%.

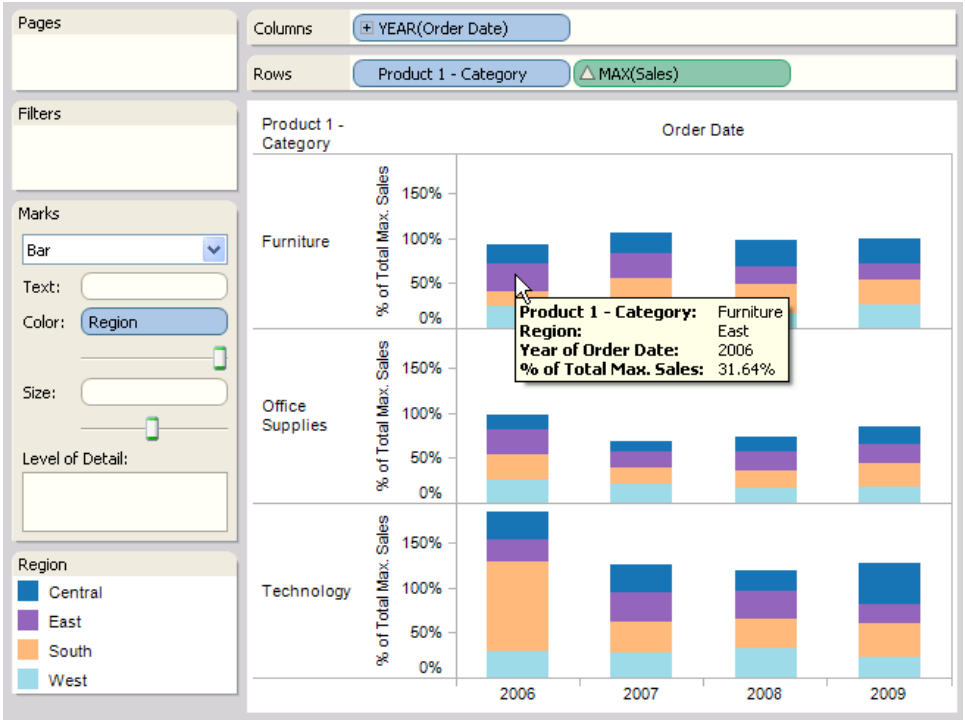

The next view displays two disaggregated measures as a scatter plot. Again, the default percentage calculation has been applied as reflected by the modified axis labels.

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.

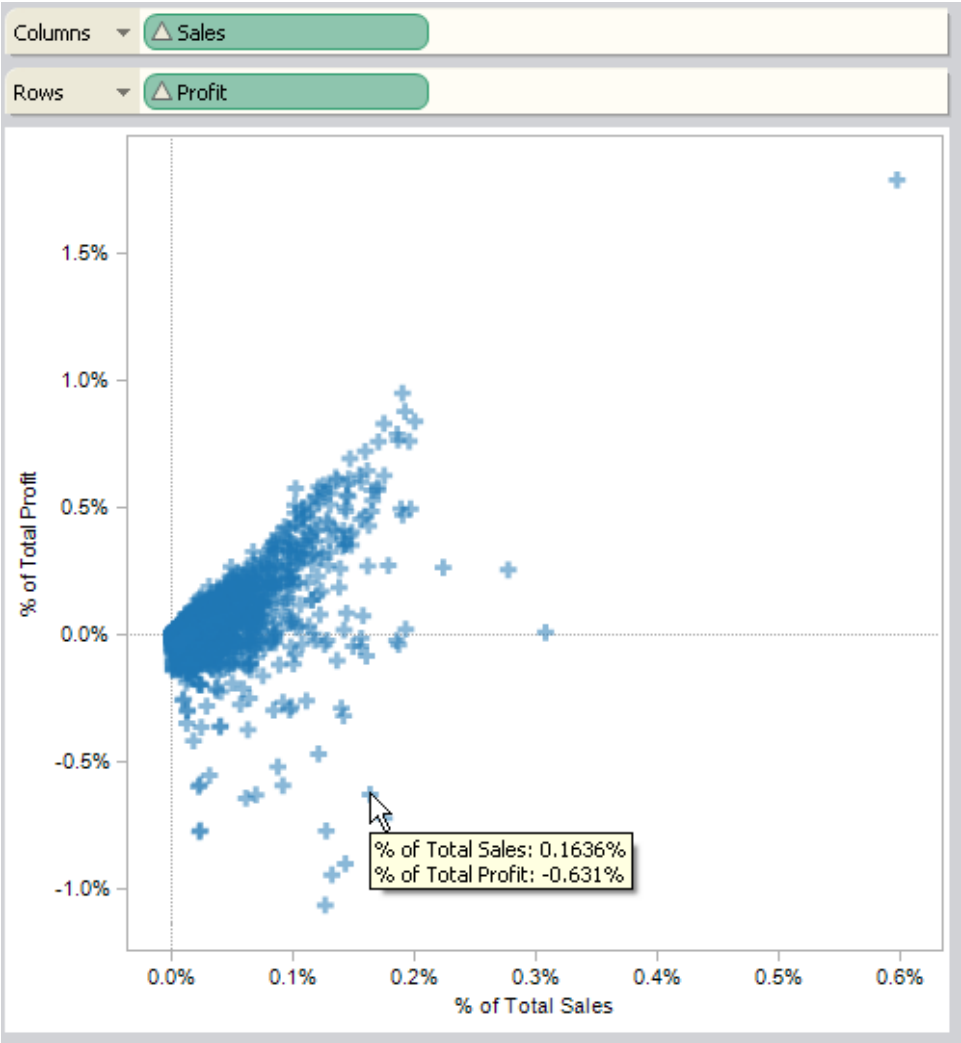

Percentage Options

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.

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 629 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.

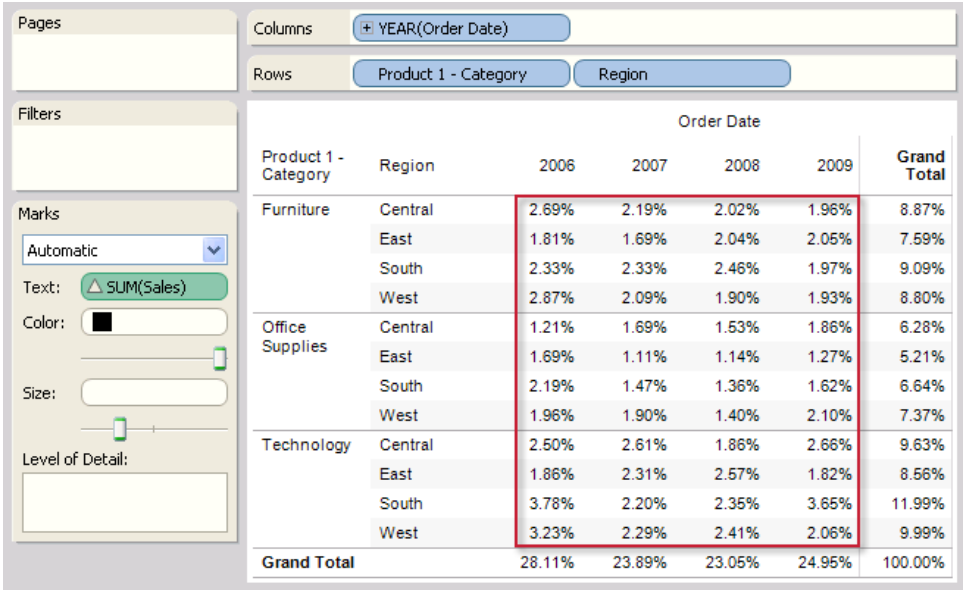

Product 1 - Category	Region	Order Date 2006	2007	2008	2009	Grand Total
Furniture	Central	2.69%	2.19%	2.02%	1.96%	8.87%
	East	1.81%	1.69%	2.04%	2.05%	7.59%
	South	2.33%	2.33%	2.46%	1.97%	9.09%
	West	2.87%	2.09%	1.90%	1.93%	8.80%
Office Supplies	Central	1.21%	1.69%	1.53%	1.86%	6.28%
	East	1.69%	1.11%	1.14%	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%
Grand Total		28.11%	23.89%	23.05%	24.95%	100.00%

An equivalent graphical table is shown below.

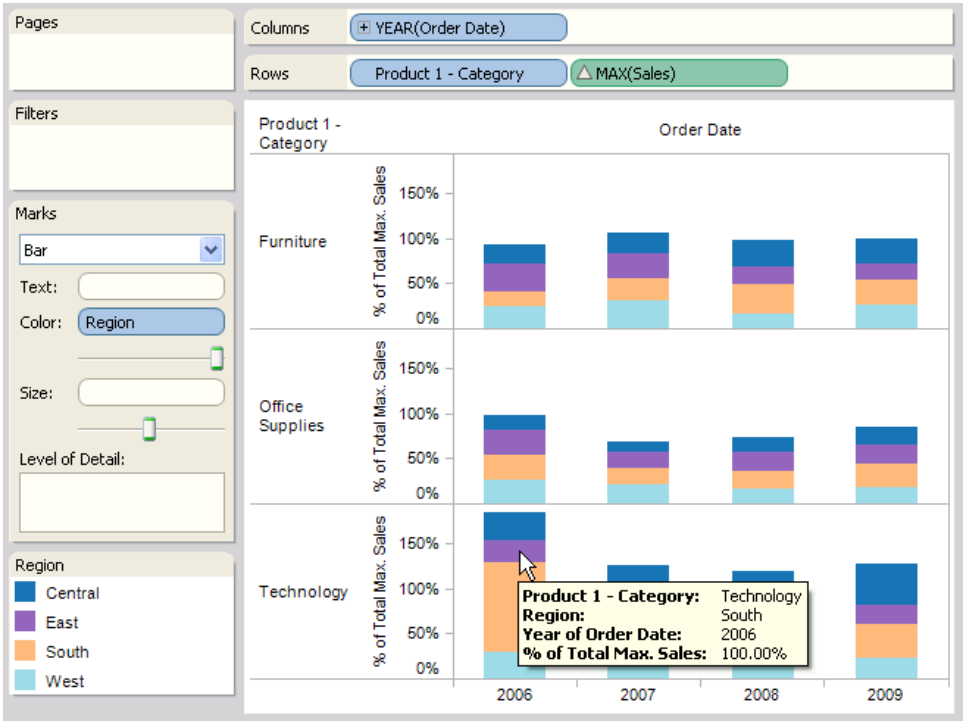

Percent of Column

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

		Order Date				
Product 1 - Category	Region	2006	2007	2008	2009	**Grand Total**
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%
Grand Total		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.

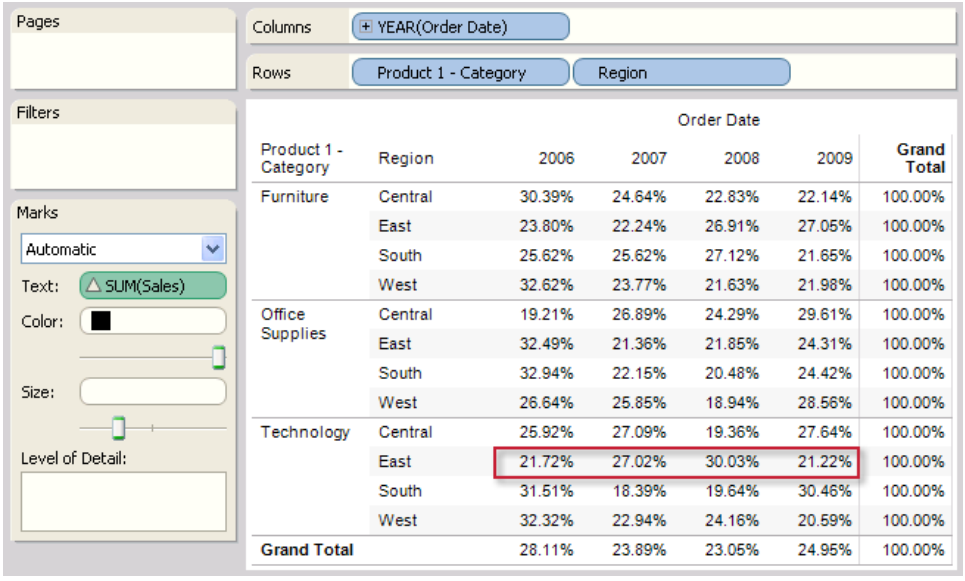

		Order Date				
Product 1 - Category	Region	2006	2007	2008	2009	Grand Total
Furniture	Central	30.39%	24.64%	22.83%	22.14%	100.00%
	East	23.80%	22.24%	26.91%	27.05%	100.00%
	South	25.62%	25.62%	27.12%	21.65%	100.00%
	West	32.62%	23.77%	21.63%	21.98%	100.00%
Office Supplies	Central	19.21%	26.89%	24.29%	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%
Grand Total		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.

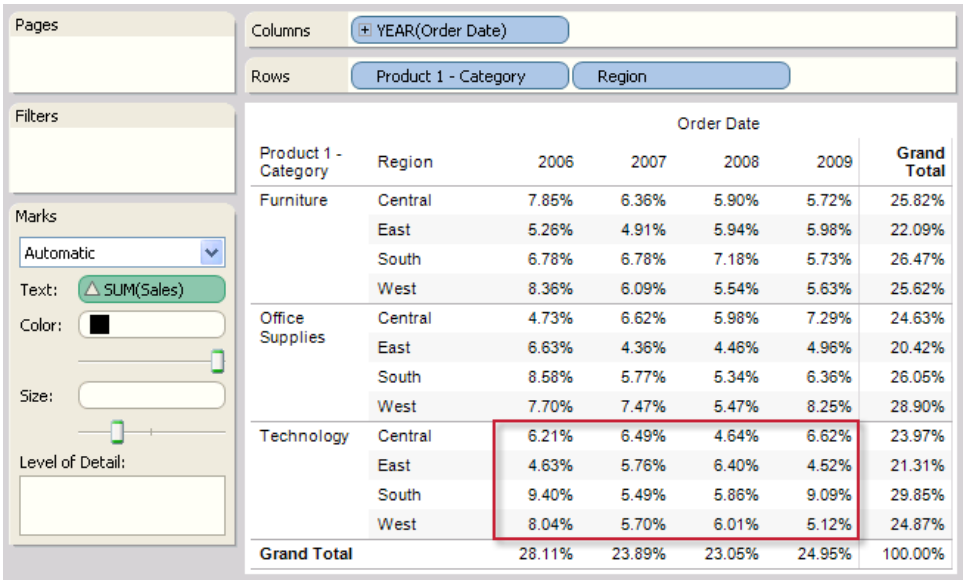

		Order Date				
Product 1 - Category	Region	2006	2007	2008	2009	Grand Total
Furniture	Central	7.85%	6.36%	5.90%	5.72%	25.82%
	East	5.26%	4.91%	5.94%	5.98%	22.09%
	South	6.78%	6.78%	7.18%	5.73%	26.47%
	West	8.36%	6.09%	5.54%	5.63%	25.62%
Office Supplies	Central	4.73%	6.62%	5.98%	7.29%	24.63%
	East	6.63%	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%
Grand Total		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 (that is, the dimension farthest to the right), 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

		Order Date				
Product 1 - Category	Region	2006	2007	2008	2009	**Grand Total**
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%
Grand Total		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.

> If you place **Measure Names** as the inner dimension on the **Rows** shelf (that is, the dimension farthest to the right on the 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.

Columns	⊞ YEAR(Order Date)
Rows	Product 1 - Category / Region

		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%
Grand Total		100.00%	100.00%	100.00%	100.00%	100.00%

Forecasting

You can forecast quantitative time-series data using exponential smoothing models in Tableau Desktop. With exponential smoothing, recent observations are given relatively more weight than older observations. These models capture the evolving trend or seasonality of your data and extrapolate them into the future. Forecasting is fully automatic, yet configurable. Many forecast results can become fields in your visualizations.

To turn forecasting on, there must be at least one date and one measure in the view. Choose **Analysis > Forecast > Show Forecast**.

The forecasted future values for the measure are now shown in the view next to the actual historical values. By default, forecast values are shown in a lighter version of the color used for actual historical values.

Forecasting Constraints

Forecasting is not supported for Multidimensional data sources. In Tableau, multidimensional data sources are supported only in Windows.

In addition, the view cannot contain any of the following:

- Table calculations
- Disaggregated measures

- Percent calculations
- Grand Totals or Subtotals
- Date values with aggregation set to Exact Date

A time series containing null values also imposes constraints—see **Resolving Forecasting Errors** on page 660.

How Forecasting Works in Tableau

Forecasting in Tableau uses a technique known as exponential smoothing. Forecast algorithms try to find a regular pattern in measures that can be continued into the future.

All forecast algorithms are simple models of a real-world data generating process (DGP). For a high quality forecast, a simple pattern in the DGP must match the pattern described by the model reasonably well. Quality metrics measure how well the model matches the DGP. If the quality is low, the precision measured by the confidence bands is not important because it measures the precision of an inaccurate estimate.

Tableau automatically selects the best of up to eight models, the best being the one that generates the highest quality forecast. The smoothing parameters of each model are optimized before Tableau assesses forecast quality. The optimization method is global. Therefore, choosing locally optimal smoothing parameters that are not also globally optimal is not impossible. However, initial value parameters are selected according to best practices but are not further optimized. So it is possible for initial value parameters to be less than optimal. The eight models available in Tableau are among those described at the following location on the OTexts web site: A taxonomy of exponential smoothing methods.

When there is not enough data in the visualization, Tableau automatically tries to forecast at a finer temporal granularity, and then aggregates the forecast back to the granularity of the visualization. Tableau provides prediction bands which may be simulated or calculated from a closed form equation. All models with a multiplicative component or with aggregated forecasts have simulated bands, while all other models use the closed form equations.

Exponential Smoothing, Trend, and Seasonality

Exponential smoothing models iteratively forecast future values of a regular time series of values from weighted averages of past values of the series. The simplest model, *Simple Exponential Smoothing*, computes the next level or smoothed value from a weighted average of the last actual value and the last level value. The method is exponential because the value of each level is influenced by every preceding actual value to an exponentially decreasing degree—more recent values are given greater weight.

Exponential smoothing models with trend or seasonal components are effective when the measure to be forecast exhibits trend or seasonality over the period of time on which the forecast is based. *Trend* is a tendency in the data to increase or decrease over time. *Seasonality* is a repeating, predictable variation in value, such as an annual fluctuation in temperature relative to the season. Tableau tests for a seasonal cycle with the length most

typical for the time aggregation of the time series for which the forecast is estimated. So if you aggregate by months, Tableau will look for a 12 month cycle; if you aggregate by quarters, Tableau will search for a four-quarter cycle; and if you aggregate by days, Tableau will search for weekly seasonality. Therefore, if there is a six-month cycle in your monthly time series, Tableau will probably find a 12-month pattern that contains two similar sub-patterns. However, if there is a seven-month cycle in your monthly time series, Tableau will probably find no cycle at all. Luckily, seven-month cycles are uncommon.

In general, the more data points you have in your time series, the better the resulting forecast will be. Having enough data is particularly important if you want to model seasonality, because the model is more complicated and requires more proof in the form of data to achieve a reasonable level of precision.On the other hand, if you forecast using data generated by two or more different DGPs, you will get a lower quality forecast because a model can only match one.

Model Types

In the Forecast Options dialog box, you can choose the model type Tableau users for forecasting. The **Automatic** setting is typically optimal for most views. If you choose **Custom** , then you can specify the trend and season characteristics independently, choosing either **None**, **Additive**, or **Multiplicative**:

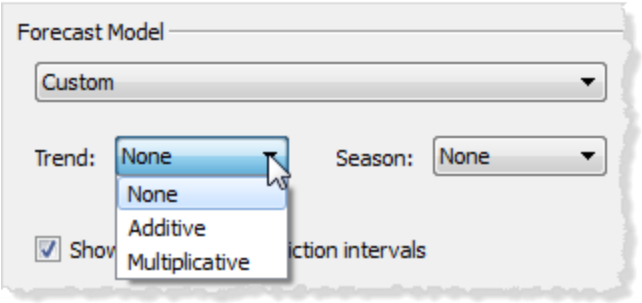

An additive model is one in which the contributions of the model components are summed, whereas a multiplicative model is one in which at least some component contributions are multiplied. Multiplicative models can significantly improve forecast quality for data where the trend or seasonality is affected by the level (magnitude) of the data:

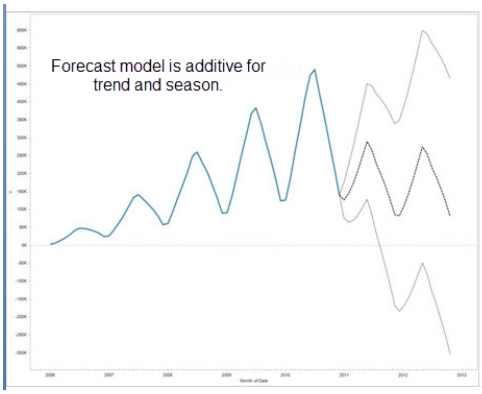

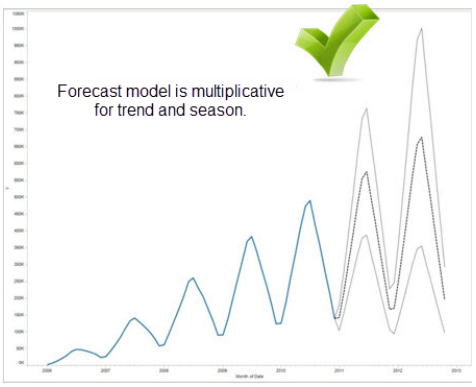

Keep in mind that you do not need to create a custom model to generate a forecast that is multiplicative: the **Automatic** setting can determine if a multiplicative forecast is appropriate for your data. However, a multiplicative model cannot be computed when the measure to be forecast has one or more values that are less than or equal to zero.

Granularity and Trimming

When you create a forecast, you select a date dimension that specifies a unit of time at which date values are to be measured. Tableau dates support a range of such time units, including Year, Quarter, Month, and Day. The unit you choose for the date value is known as the *granularity* of the date.

The data in your measure typically does not align precisely with your unit of granularity. You might set your date value to quarters, but your actual data may terminate in the middle of a quarter—for example, at the end of November. This can cause a problem because the value for this fractional quarter is treated by the forecasting model as a full quarter, which will typically have a lower value than a full quarter would. If the forecasting model is allowed to consider this data, the resulting forecast will be inaccurate. The solution is to trim the data, such that the trailing periods that could mislead the forecast are ignored. Use the Ignore Last option in the Forecast Options dialog box to remove—or *trim*—such partial periods. The default is to trim one period.

Getting More Data

Tableau requires at least five data points in the time series to estimate a trend, and enough data points for at least two seasons or one season plus five periods to estimate seasonality. For example, at least nine data points are required to estimate a model with a four quarter seasonal cycle (4 + 5), and at least 24 to estimate a model with a twelve month seasonal cycle (2 * 12).

If you turn on forecasting for a view that does not have enough data points to support a good forecast, Tableau can sometimes retrieve enough data points to produce a valid forecast by querying the datasource for a finer level of granularity:

- If your view contains fewer than nine years of data, by default, Tableau will query the data source for quarterly data, estimate a quarterly forecast, and aggregate to a yearly forecast to display in your view. If there are still not enough data points, Tableau will estimate a monthly forecast and return the aggregated yearly forecast to your view.
- If your view contains fewer than nine quarters of data, by default Tableau will estimate a monthly forecast and return the aggregated quarterly forecast results to your view.
- If your view contains fewer than nine weeks of data, by default, Tableau will estimate a daily forecast and return the aggregated weekly forecast results to your view.
- If your view contains fewer than nine days of data, by default, Tableau will estimate an hourly forecast and return the aggregated daily forecast results to your view.
- If your view contains fewer than nine hours of data, by default, Tableau will estimate an minutely forecast and return the aggregated hourly forecast results to your view.

- If your view contains fewer than nine minutes of data, by default, Tableau will estimate an secondly forecast and return the aggregated minutely forecast results to your view.

These adjustments happen behind the scene and require no configuration. Tableau does not change the appearance of your visualization, and does not actually change your date value. However, the summary of the forecast time period in the Forecast Describe and Forecast Options dialog will reflect the actual granularity used.

Tableau can only get more data when the aggregation for the measure you are forecasting is SUM or COUNT. See **Aggregations** on page 563 for information on available aggregation types and information on how to change the aggregation type.

To Create a Forecast

Forecasting requires a view that uses at least one date dimension and one measure. For example:

- The field you want to forecast is on the **Rows** shelf and a continuous date field is on the **Columns** shelf.
- The field you want to forecast is on the **Columns** shelf and a continuous date field is on the **Rows** shelf.
- The field you want to forecast on either the **Rows** or **Columns** shelf, and discrete dates are on either the **Rows** or **Columns** shelf. At least one of the included date levels must be Year.
- The field you want to forecast is on the Marks card, and a continuous date or discrete date set is on **Rows**, **Columns** or **Marks**.

To turn forecasting on, either right-click on the visualization and choose **Forecast > Show Forecast**, or choose **Analysis > Forecast > Show Forecast**.

With forecasting on, Tableau visualizes estimated future values of the measure, in additional to actual historical values. The estimated values are shown by default in a lighter shade of the color used for the historical data:

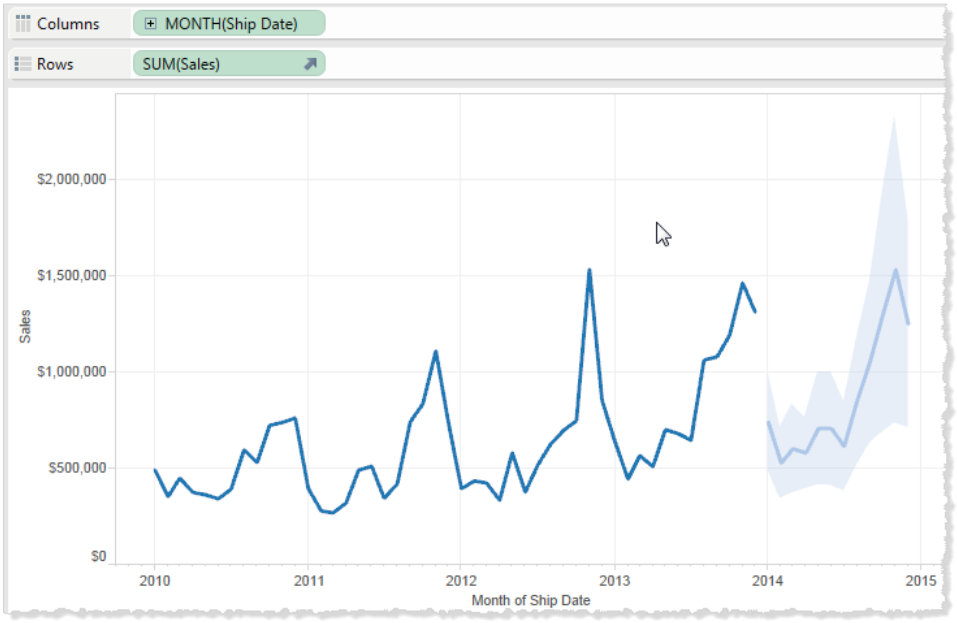

Prediction Intervals

The shaded area in the image above shows the 95% prediction interval for the forecast. That is, the model has determined that there is a 95% likelihood that the value of sales will be within the shaded area for the forecast period. You can configure the confidence level percentile for the prediction bands, and whether prediction bands are included in the forecast, using the Show prediction intervals setting in the Forecast Options dialog box:

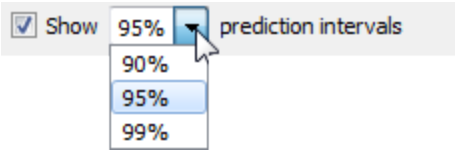

Clear the check box if you do not want to display prediction bands in forecasts. To set the prediction interval, select one of the values or enter a custom value. The lower the percentile you set for the confidence level, the narrower the prediction bands will be.

How your prediction intervals are displayed depends on the mark type of your forecasted marks:

Forecast mark type	Prediction intervals displayed using
Line	Bands
Shape, square, circle, bar, or pie	Whiskers

In the following example, forecast data is indicated by lighter shaded circles, and the prediction intervals are indicated by lines ending in whiskers:

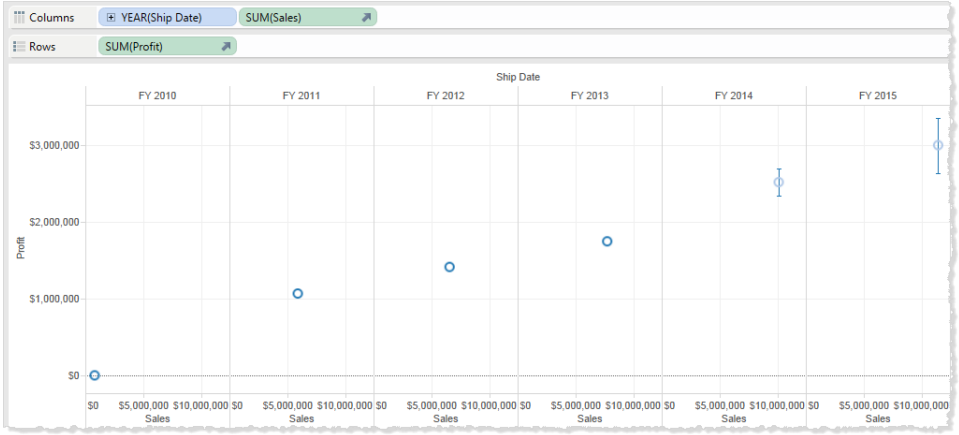

Enhancing Forecasts

For each forecast value, consider verifying the quality or precision of your forecast by dragging another instance of the forecast measure from the Data Window to the Detail shelf on the Marks card and then changing the result type to one of the following:

- **Quality of <*value*>**: Shows the quality of the forecast, on a scale of 0 (worst) to 100 (best).
- **Precision of <*value*>**: Shows the 95% prediction interval distance from the forecasted value.

You can repeat the process to add additional result types for each forecast value. See "Changing the Forecast Result" type in **Forecast Field Results** below for information on changing the result type.

By adding such result types to the Details shelf, you add information about the forecast to tooltips for all marks that are based on forecasted data.

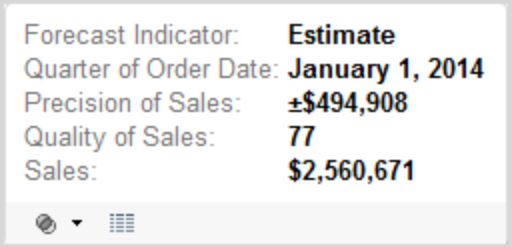

Forecast Field Results

Tableau provides several types of forecast results. To view these result types in the view, right-click on the measure field, choose **Show Forecast**, and then choose one of the options.

The options are:

- **Actual & Forecast**—Show the actual data extended by forecasted data.
- **Trend**—Show the forecast value with the seasonal component removed.

- **Precision**—Show the prediction interval distance from the forecast value for the configured confidence level.
- **Precision %**—Show precision as a percentage of the forecast value.
- **Quality**—Show the quality of the forecast, on a scale of 0 (worst) to 100 (best). This metric is scaled MASE, based on the MASE (Mean Absolute Scaled Error) of the forecast, which is the ratio of forecast error to the errors of a naïve forecast which assumes that the value of the current period will be the same as the value of the next period. The actual equation used for quality is:

  $$100 * \max(1 - MASE,\ 0)$$

  The Quality for a naïve forecast would be 0. The advantage of the MASE metric over the more common MAPE is that MASE is defined for time series which contain zero, wheras MAPE is not. In addition, MASE weights errors equally while MAPE weights positive and/or extreme errors more heavily.
- **Upper Prediction Interval**—Shows the value above which the true future value will lie confidence level percent of the time assuming a high quality model. The confidence level percentage is controlled by the Prediction Interval setting in the Forecast Options dialog box. See **Configure Forecast Options** on the next page.
- **Lower Prediction Interval**—Shows 90, 95, or 99 confidence level below the forecast value. The actual interval is controlled by the **Prediction Interval** setting in the Forecast Options dialog box.
- **Indicator**—Show the string **Actual** for rows that were already on the worksheet when forecasting was inactive and **Estimate** for rows that were added when forecasting was activated.
- **None**—Do not show forecast data for this measure.

Forecast description information is also included in the worksheet description. See **Describing the View** on page 482.

Forecasting a New Measure

When you add a new measure to a visualization that already has forecasting enabled, Tableau attempts to forecast future values.

Changing the Forecast Result Type

To change the forecast result type for a measure, right-click on the measure field, select **Show Forecast**, and then choose a result type.

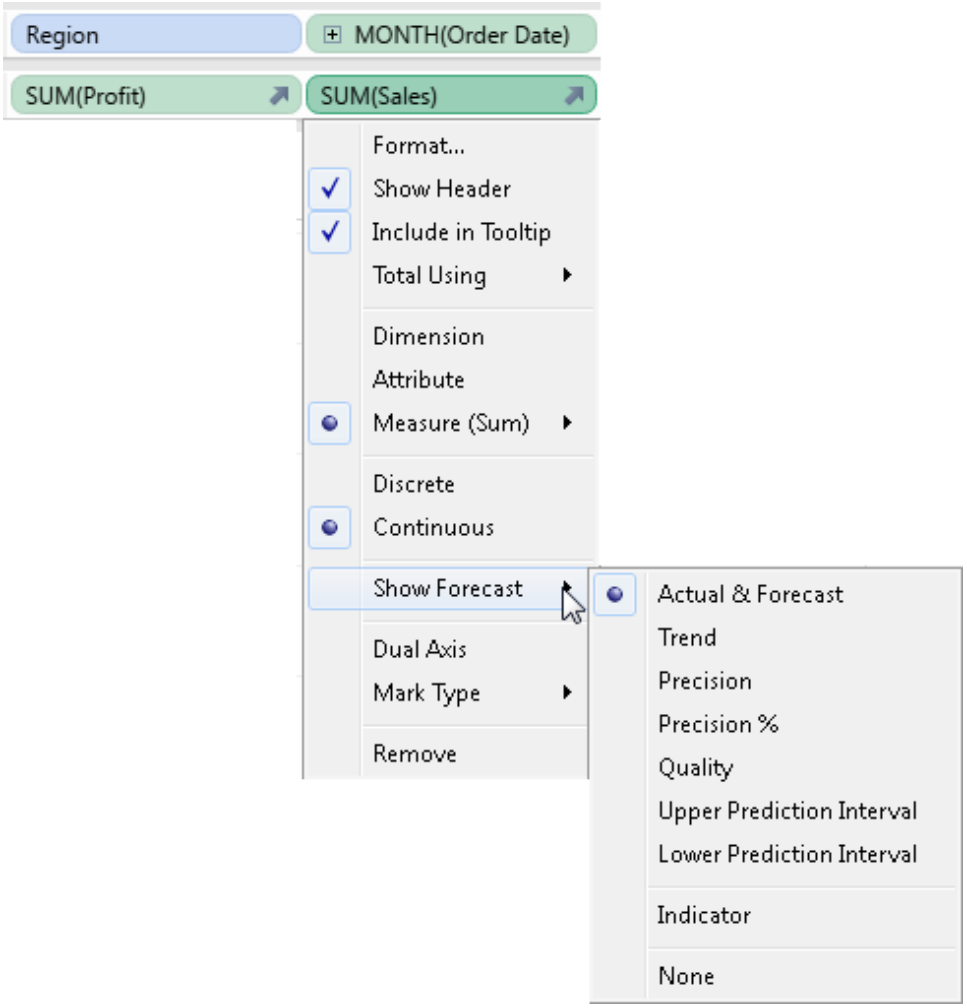

Configure Forecast Options

Use the Forecast Options dialog box to configure forecast options, including:

- The length of the forecast
- The range and temporal aggregation of source data from which to generate the forecast
- The forecast model
- Prediction intervals

When forecasting is enabled, you can open the Forecast Options dialog box by choosing **Analysis > Forecast > Forecast Options**.

Forecast Length

The **Forecast Length** section determines how far into the future the forecast extends. Select one of the following:

- **Auto**: Tableau automatically picks the forecast length.
- **Exactly**: Extends the forecast for the specified number of units.
- **Until**: Extends the forecast to the specified point in the future.

Source Data

Use the **Source Data** section to specify.

- **Aggregate by**: Specifies the temporal granularity of the time series. With the default value (**Automatic**), Tableau chooses the best granularity for estimation. This will typically match the temporal granularity of the visualization (that is, the date dimension that the forecast is based on). However, it is sometimes possible and desirable to estimate the forecast model at a finer granularity than the visualization when the time series in the visualization is too short to allow estimation.
- **Ignore last**: Specifies the number of periods at the end of the actual data that should be ignored in estimating the forecast model. Forecast data is used instead of actual data for these time periods. Use this feature to trim off unreliable or partial trailing periods which could mislead the forecast. When the estimation granularity specified in the **Source Data** section is finer than in the visualization, the trimmed periods are estimation periods. As a result, the trailing actual visualization period may become a forecast period, which is an aggregate of both actual and forecast periods of estimation granularity. In contrast, null values are not filled with zeros and must be filtered to allow forecast.
- **Fill in missing values with zeros**: If there are missing values in the measure you are attempting to forecast, Tableau cannot create a forecast. Select this option to prevent this error.

Forecast Model

The **Forecast Model** section specifies how the forecast model is to be produced.

Use the drop down to specify whether Tableau selects what it determines to be the best of all models (**Automatic**), the best of those with no seasonal component (**Automatic without seasonality**), or the model you specify (**Custom**).

When you choose the **Custom** option, two new fields appear in the Forecast Options dialog box, which you use to specify the trend and season characteristics for your model:

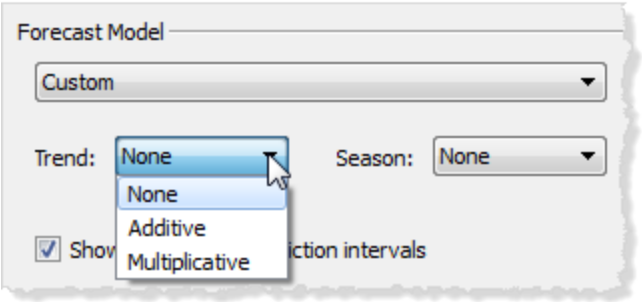

The choices are the same for both fields:

- **None**: When you select **None** for Trend, the model does not assess the data for trend. When you select **None** for Season, the model does not assess the data for seasonality.
- **Additive**: An additive model is one in which the combined effect of several independent factors is the sum of the isolated effects of each factor. You can assess the data in your view for additive trend, additive seasonality, or both.
- **Multiplicative** : A multiplicative model is one in which the combined effect of several independent factors is the product of the isolated effects of each factor. You can assess the data in your view for multiplicative trend, multiplicative seasonality, or both.

If there is more than one time series in your visualization, the Custom option forces them all to be forecast using the same custom model. Constraining the models in this way usually results in lower quality models than would be produced by automatic model selection.

Constraints on Multiplicative Models

- You cannot use a multiplicative model when the measure to be forecast has one or more values that are less than or equal to zero, or even when some of the data points are too close to zero, relative to other data points.
- You cannot specify a model with multiplicative trend and additive season because the result may be numerically unstable.

Prediction Interval

You can set the prediction interval to 90, 95, or 99 percent, or enter a custom value. This value is used in two locations:

- In the prediction bands displayed with a forecast.
- For the prediction interval options (**Upper Prediction Interva**l and **Lower Prediction**

Interval) that are available as forecast result types for a measure in the view:

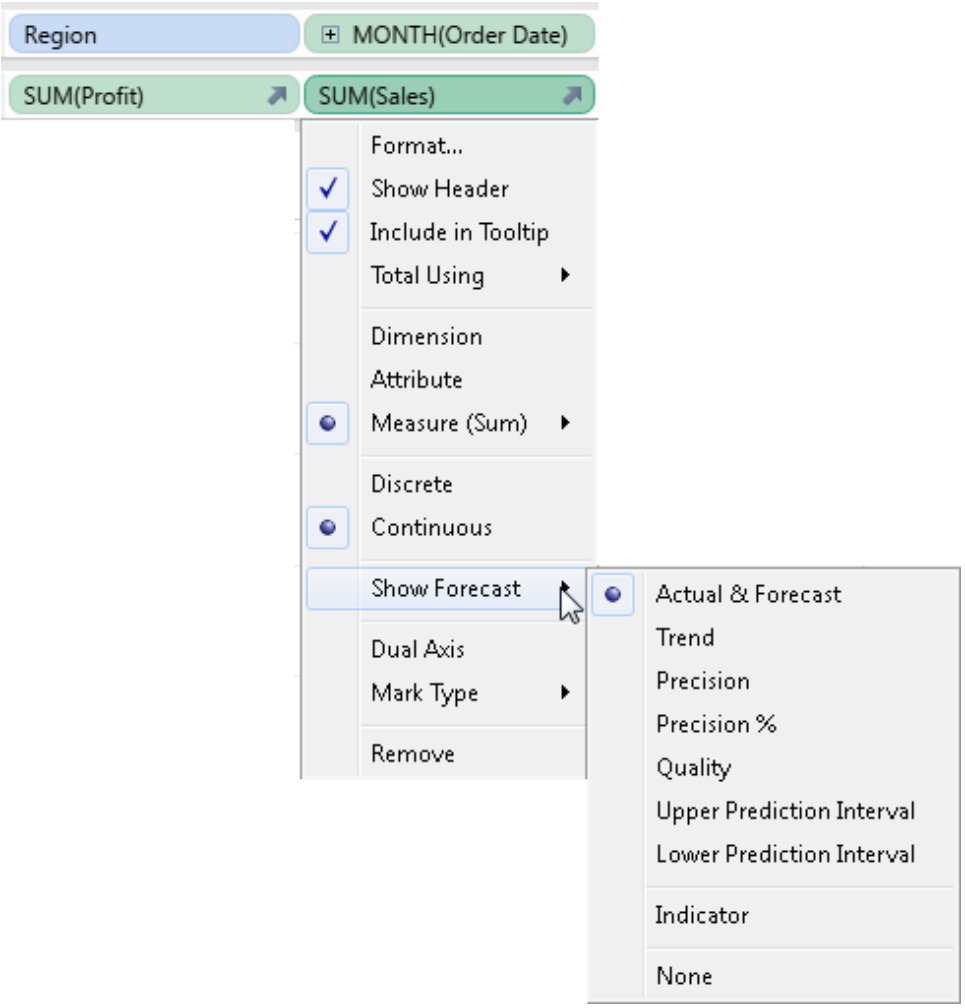

Forecast Summary

The text box at the bottom of the Forecast Options dialog box provides a description of the current forecast. The forecast summary updates whenever you change any of the forecast options above. If there is a problem with the forecast, the text box provides an error message that may help you resolve the issue. See **Resolving Forecasting Errors** on page 660.

Forecast Descriptions

The **Describe Forecast** dialog box describes the forecast models that Tableau computed for your visualization. When forecasting is enabled, you can open this dialog by selecting **Analysis > Forecast > Describe Forecast**. The information in the Describe Forecast dialog box is read-only, though you can click **Copy to Clipboard** and then paste the screen contents into a writeable document.

The **Describe Forecast** dialog box has two tabs: a **Summary** tab and a **Models** tab.

Describe Forecast – Summary Tab

The **Summary** tab describes the forecast models Tableau has created, as well as the general patterns Tableau discovered in the data.

Options Used To Create Forecasts

This section summarizes the options Tableau used to create forecasts. These options were either picked automatically by Tableau or specified in the Forecast Options dialog box.

- **Time series**—The continuous date field used to define the time series.
- **Measures**—The measures for which values are estimated.
- **Forecast forward**—The length and date range of the forecast.
- **Forecast based on**—The date range of the actual data used to create the forecast.
- **Ignore last**—The number of periods at the end of the actual data that are disregarded--forecast data is displayed for these periods.This value is determined by the **Ignore Last** option in the Forecast Options dialog box.
- **Seasonal pattern**—The length of the seasonal cycle that Tableau found in the data, or None if no seasonal cycle was found in any forecast.

Forecast Summary Tables

For each measure that is forecasted, a summary table is displayed describing the forecast. If the view is broken into multiple panes using dimensions, a column is inserted into each table that identifies the dimensions. The fields in summary forecast tables are:

- **Initial**—The value and prediction interval of the first forecast period.
- **Change From Initial**—The difference between the first and the last forecast estimate points. The interval for those two points is shown in the column header. When values are shown as percentages, this field shows the percentage change from the first forecast period.
- **Seasonal Effect**—These fields are displayed for models identified as having seasonality--that is, a repeating pattern of variation over time. They show the high and low value of the seasonal component of the last full seasonal cycle in the combined time series of actual and forecast values. The seasonal component expresses the deviation from the trend and so varies around zero and sums to zero over the course of a season.
- **Contribution**—The extent to which Trend and Seasonality contribute to the forecast. These values are always expressed as percentages and add up to 100%.
- **Quality**—Indicates how well the forecast fits the actual data. Possible values are GOOD, OK, and POOR. A naïve forecast is defined as a forecast that estimates that the value of the next period will be identical to the value of the current period. Quality is expressed relative to a naïve forecast, such that OK means the forecast is likely to have less error than a naïve forecast, GOOD means that the forecast has less than half as much error, and POOR means that the forecast has more error.

Describe Forecast – Models Tab

The **Models** tab provides more exhaustive statistics and smoothing coefficient values for the Holt-Winters exponential smoothing models underlying the forecasts. For each measure that is forecasted, a table is displayed describing the forecast models Tableau created for the measure. If the view is broken into multiple panes using dimensions, a column is inserted into each table that identifies the dimensions. The table is organized into the following sections:

Model

Specifies whether the components **Level**, **Trend**, or **Season** are part of the model used to generate the forecast. The value for each component is one of the following:

- **None**—The component is not present in the model.
- **Additive**—The component is present and is added to the other components to create the overall forecast value.
- **Multiplicative**—The component is present and is multiplied by the other components to create the overall forecast value.

Quality Metrics

This set of values provides statistical information about the quality of the model.

Value	Definition
RMSE: Root mean squared error	$\sqrt{\left(\frac{1}{n}\right)\sum e(t)^2}$
MAE: Mean absolute error	$\frac{1}{n}\sum\left\lvert e\left(t\right)\right\rvert$
MASE: Mean absolute scaled error. MASE measures the magnitude of the error compared to the magnitude of the error of a naive one-step ahead forecast as a ratio. A naive forecast assumes that whatever the value is today will be same value tomorrow. So, a MASE of 0.5 means that your forecast is likely to have half as much error as a naive forecast, which is better than a MASE of 1.0, which is no better than a naive forecast. Since this is a normalized statistic that is defined for all values and weighs errors evenly, it is an excellent metric for comparing the quality of different forecast methods. The advantage of MASE over the more common MAPE metric is that MASE is defined for time series that contain zero, whereas MAPE is not. Also, MASE weights errors equally, whereas MAPE weights positive and/or extreme errors more heavily.	$\frac{\frac{1}{n}\sum\left\lvert e\left(t\right)\right\rvert}{\frac{1}{(n-1)}\sum_{2}^{n}\left\lvert Y\left(t\right)-Y\left(t-1\right)\right\rvert}$
MAPE: Mean absolute percentage error. MAPE measures the magnitude of the error compared to the	$100\frac{1}{n}\sum\left\lvert\frac{e(t)}{A(t)}\right\rvert$

magnitude of your data, as a percentage. So, a MAPE of 20% is better than a MAPE of 60%. Errors are the differences between the response values, which the model estimates, and the actual response values for each explanatory value in your data. Since this is a normalized statistic, it can be used to compare the quality of different models computed in Tableau. However, it is unreliable for some comparisons because it weights some kinds of error more heavily than others. Also, it is undefined for data with values of zero.	
AIC: Akaike information criterion. AIC is a model quality measure, developed by Hirotugu Akaike, that penalizes complex models to prevent overfitting. In this definition, *k* is the number of estimated parameters, including initial states, and *SSE* is the sum of the squared errors.	$n * \log(SSE / n) + 2 * (k + 1)$

In the preceding definitions, the variables are as follow:

Variable	Meaning
t	Index of a period in a time series.
n	Time series length.
m	Number of periods in a season/cycle.
A(t)	Actual value of the time series at period t.
F(t)	Fitted or forecast value at period t.

Residuals are: e(t) = F(t)-A(t)

Smoothing Coefficients

Depending on the rate of evolution in the level, trend, or seasonal components of the data, smoothing coefficients are optimized to weight more recent data values over older ones, such that within-sample one-step-ahead forecast errors are minimized. Alpha is the level smoothing coefficient, beta the trend smoothing coefficient, and gamma the seasonal smoothing coefficient. The closer a smoothing coefficient is to 1.00, the less smoothing is performed, allowing for rapid component changes and heavy reliance on recent data. The closer a smoothing coefficient is to 0.00, the more smoothing is performed, allowing for gradual component changes and less reliance on recent data.

Troubleshooting Forecasting

This topic discusses some issues that may arise as you use forecasting in Tableau.

Null Forecasts

A null forecast results when you are using forecasting and modify the view in a manner that is incompatible with forecasting. The most obvious indication that you have a null forecast is that

no forecast data is displayed in your view and the text No Forecast is displayed. Hover the cursor over this message to see why Tableau is unable to create a forecast:

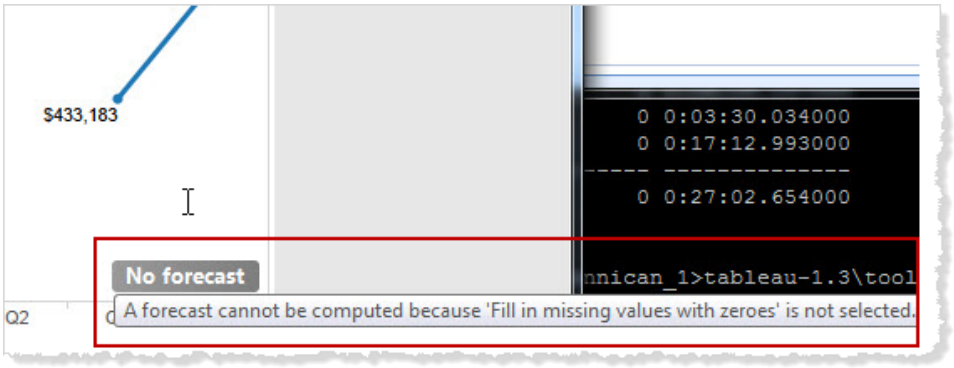

Another indication is that the **Forecast Indicator** field on the Marks card shows Actual values, but no Estimate values:

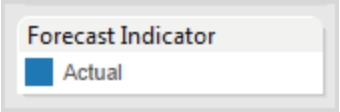

To diagnose a null forecast, open the Describe Forecast dialog box from the Analysis menu (**Analysis > Forecast > Describe Forecast**) to see the error message. Then see **Resolving Forecasting Errors** below for a suggested resolution.

Tableau also displays the error message in the Forecast Options dialog box (**Analysis > Forecast > Forecast Options**).

Because a forecast cannot be computed for a time series with null date values in the middle, actions that filter data, either explicitly or implicitly can trigger a null forecast. The **Keep Only** and **Exclude** commands on tooltips are examples of actions that can filter data implicitly—in some cases, these commands are removed when a forecast is shown. For example, if you have a time series of sales for each quarter from 2008 until 2012 and you exclude the Sales value for the second quarter of 2010, you will get a Null forecast because the time series is irregular. If, instead, you exclude the first quarter of 2008, you have shortened the time series but it remains regular. So, a valid forecast is still possible.

Resolving Forecasting Errors

If Tableau is unable to provide a forecast for your view, the problem can often be resolved by changing the Date value in the view (see **Changing Date Levels** on page 439).

Forecasting errors can result when the aggregation level of the time series (months, weeks, etc.) is either too fine or too coarse for the data to be forecast. This can lead to the "too much data" or "too little data" errors described below. Date aggregation can trigger a"too many Nulls" scenario when forecasting attempts to extract more data from the measure than is possible. For example, if the underlying granularity of the sales data is months but you aggregate by weeks, the result may be a significant number of Null values.

Other problems arise when the view's aggregation and the aggregation specified for the forecast (using the **Aggregate by** field in the Forecast Options dialog box) are not compatible. Tableau can create a forecast when the forecast aggregation is a finer level of detail than the view's aggregation, but not when it is at a coarser level of detail; even when it is finer, the two values are only compatible if there is a strict hierarchy that Tableau can use (for example, quarters can be evenly divided into three months, but months can't be evenly divided into weeks). Avoid these scenarios by setting **Aggregate by** to Automatic.

The following list shows errors that can be result from invalid forecasts in Tableau, and provides advice on how to resolve them.

Error message	Suggestion for Resolution
A continuous date cannot be derived from the date fields in the view.	Forecasting requires a date field that can be interpreted continuously. If the date field is not explicitly continuous, then one of the included date levels must be Year. This error is returned if there are no dates in the view, or if the dates in the view don't constitute a full hierarchy (for example, the date includes Year and Day, but not Month), or if they constitute a hierarchy that is not supported (for example, Year, Week, Day).
The time series is too short to forecast.	Expand the time series in your view to include more date values. This error is returned if there are fewer than four data points after trimming off unreliable or partial trailing periods which could mislead the forecast.
A forecast cannot be computed for a time series with Null date values.	Eliminate any Null values from the date field or fields in the view, either by filtering the date field or by using a less detailed date granularity (for example, by switching from months to quarters).
A forecast cannot be computed when the view contains multiple distinct date fields.	This error is returned if there are multiple date fields in the view. For example, if both Order Date and Ship Date are in the same view, forecasting is not supported.
The selected 'Aggregate by' value in Forecast Options is not compatible with the visualization.	The date in the view must be compatible with the value of **Aggregate by** in the Forecast Options dialog box. For example, if **Aggregate by** is set to Weeks and the date in the view is set to Months, this error occurs. Change one of the dates so that the two are compatible, or set **Aggregate by** to Automatic.
A forecast cannot be computed because 'Fill in missing values with zeroes' is not selected.	This error is returned if there are missing data points and the user has not selected **Fill in missing values with zeros** in the Forecast Options dialog box. Either select **Fill in missing values with zeros** or provide the missing values.

A forecast cannot be computed because there are too many missing values.	This error is returned if more than 40% of the data is missing. Selecting **Fill in missing values with zeros** in the Forecast Options dialog box will not resolve this error. You must modify the source data or use data from a different source.
A forecast cannot be computed because the measure to be forecast contains Null values.	This error is returned if the measure to be forecast contains one or more Null values.See **Null Values** on page 750 for options on how to suppress Null values.
There is no measure to forecast.	This error is returned if no measure that can be forecast is present in the view. Forecast measures must be on the Rows or Columns shelf, or on the Marks card.
The measure to forecast must be a number.	Some measures cannot be interpreted numerically and therefore cannot be forecast.
A forecast cannot be computed for a dimension.	The value to be forecast must be a measure, and not a dimension.
There is too much data to compute a forecast.	Forecasting is not possible when there are more than 1000 rows of data--that is, the measure to be forecast contains too many values. Update your data source to extract or sample fewer than 1000 rows.
A forecast cannot be computed because the data is divided into too many rows, columns, or colors.	Simplify the view to resolve the error by filtering or removing some of the dimensions.
A forecast cannot be computed because the view contains table calculations.	Create a version of the view that does not contain table calculations.
A forecast cannot be computed because there is a measure on the Filters shelf.	Remove the measure from the Filters shelf.
A forecast cannot be computed because **Aggregate Measures** is not selected.	**Aggregate Measures** is an option on the Analysis menu. See **Aggregating Data** on page 567.
A forecast cannot be computed because the view contains percent calculations.	**Percentage of** is an option on the Analysis menu. See **Percentages** on page 637.
A forecast cannot be computed because Grand Totals or Subtotals is enabled.	These options are controlled from the **Totals** command in the Analysis menu. See **Totals** on page 627.
A multiplicative model cannot be computed because the measure to be forecast has one or more values that	You have created a custom model with Trend or Seasonality set to **Multiplicative**. Change this value, or set the Forecast Model to **Automatic**.

are less than or equal to zero.	
A model with multiplicative trend and additive season is not allowed because it is numerically unstable.	You have created a custom model configured as described in the error message. Change the settings for the custom model, or set the Forecast Model to **Automatic**.
A seasonal model cannot be computed because the time series is too short.	Expand the time series in your view to include more date values.
The selected multiplicative model cannot be computed because some of the data is too close to zero relative to the rest of the data.	You have created a custom model configured as described in the error message. Change the settings for the custom model, or set the Forecast Model to **Automatic**.

Parameters

Parameters are dynamic values that can replace constant values in calculations, filters, and reference lines. 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. Alternatively, you may have a filter to show the top 10 products by profit. You can replace the fixed value “10” in the filter to by a dynamic parameter so you can quickly look at the top 15, 20, and 30 products.

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 any one of the following:
    - In the Calculated field dialog box, click the **Create** link at the top of the list of parameters.

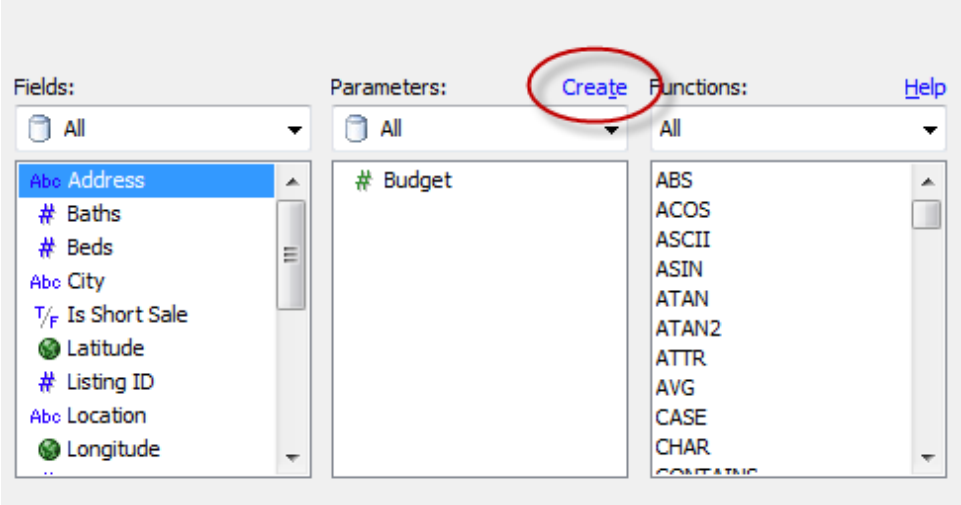

In the Data window, right-click a field to base the parameter on and select **Create Parameter**.

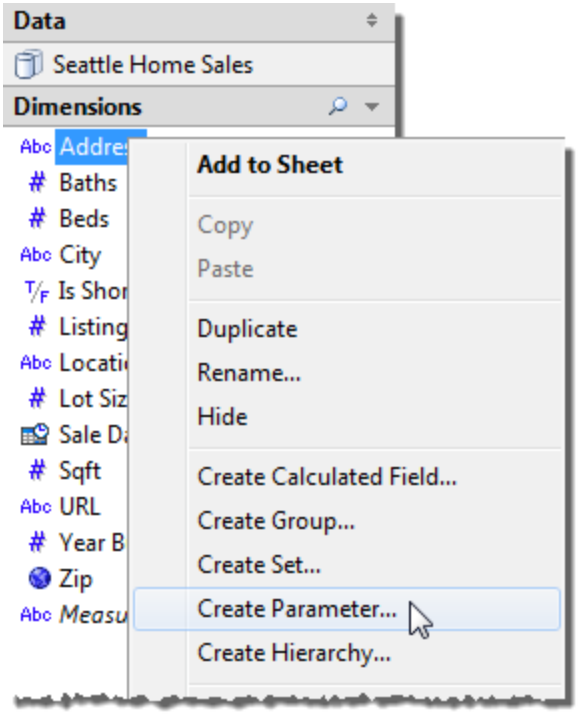

-
- Open the Data window menu using the drop-down arrow in the upper right corner and select **Create Parameter**.

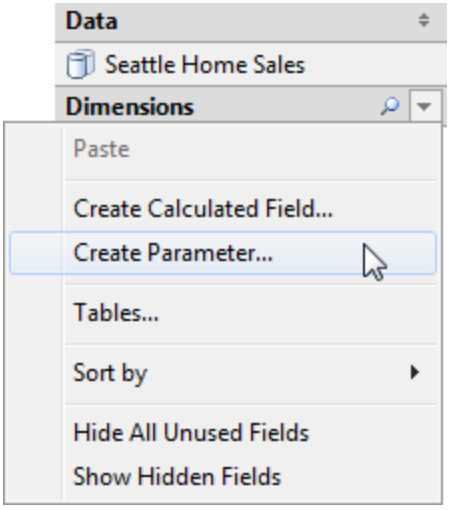

On the Top tab of the Filter dialog box, select **Create New Parameter** in the Top drop-down list (by field or by formula).

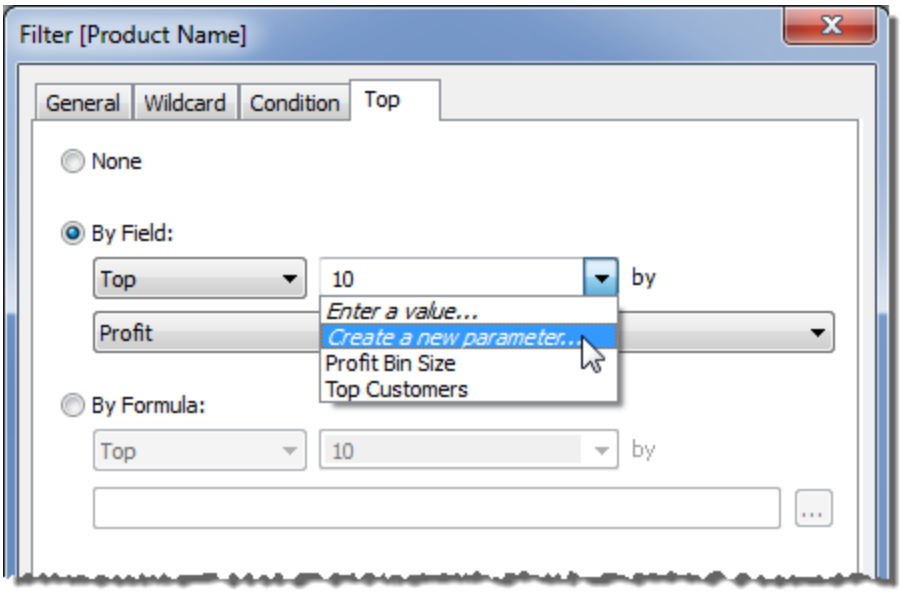

2. In the Create Parameter dialog box, give the field a **Name** and optionally write a **Comment** to describe the parameter.

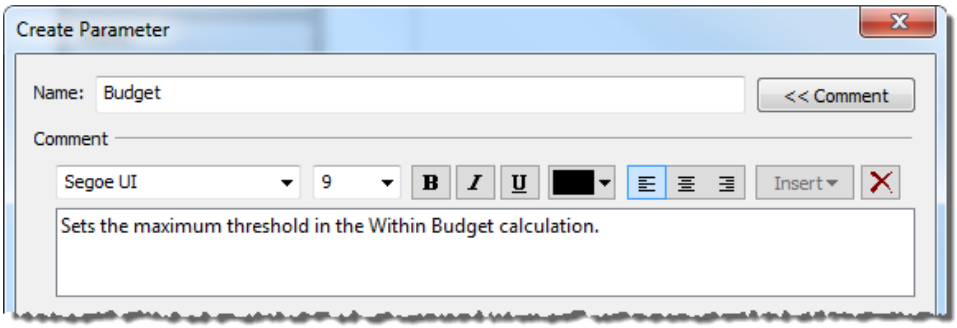

Specify the data type for the values it will accept:

3. 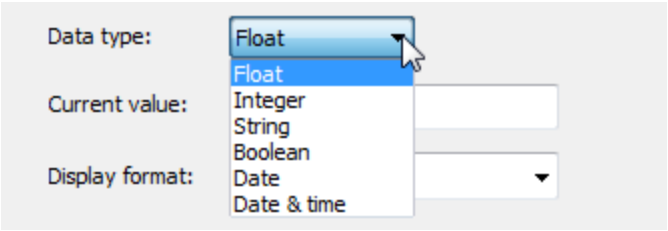

Specify a current value. This is the default value for the parameter.

4.

Specify the display format to use in the parameter control.

5. 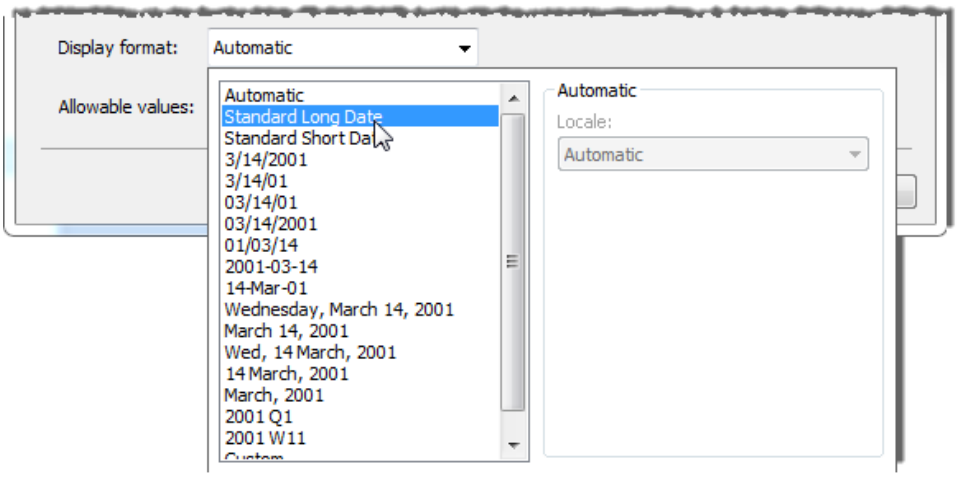

6. Specify how the parameter will accept values. You can select from the following options:

- **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**.

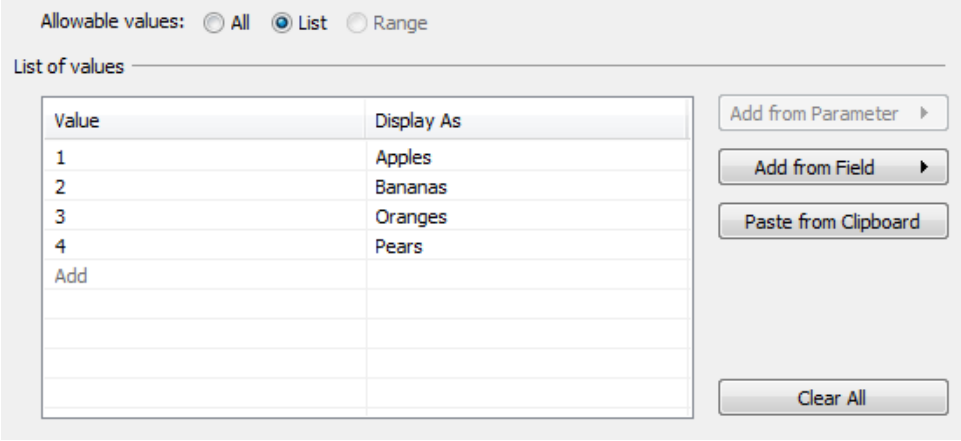

If you select **Range** you must specify a minimum, maximum and step size. For example, you can define a date range between January 1, 2010 and December 31, 2010, with the step size set to 1 month to create a parameter control that lets you select each month in 2010.

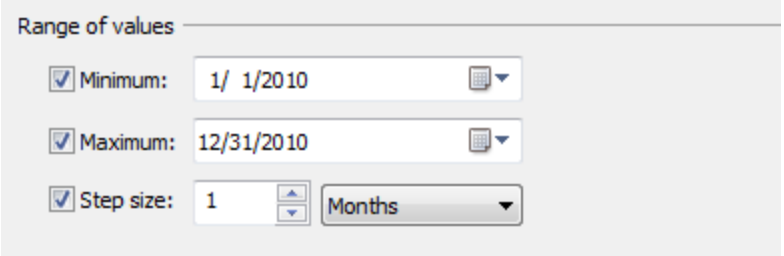

7. When finished, click **OK**.

   The parameter is listed in the Parameters section at the bottom of the Data window.

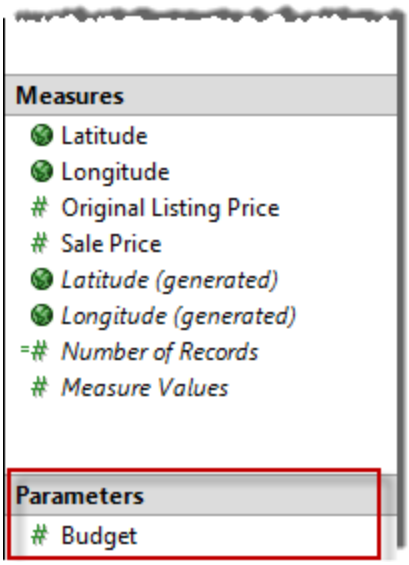

It is also available in the Parameters list in the Calculated Field dialog box, the Top tab of the Filter dialog box, and in the Reference Line dialog box. Parameters are global across the workbook and can be used in any worksheet.

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.

To delete a parameter, right-click it in the Data window and select **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. Double-click the parameter in the Parameters list to add it to your formula.

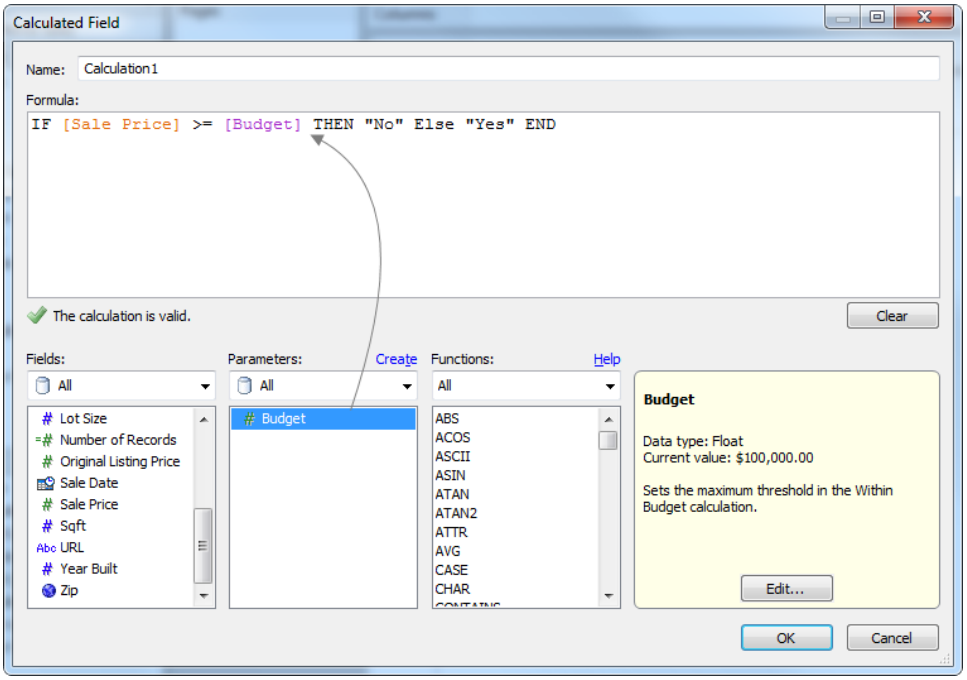

Because you often don't realize you need a parameter until you are authoring the calculation, you can create parameters directly from the Calculated Field dialog box. Click the **Create** link at the top of the Parameters list.

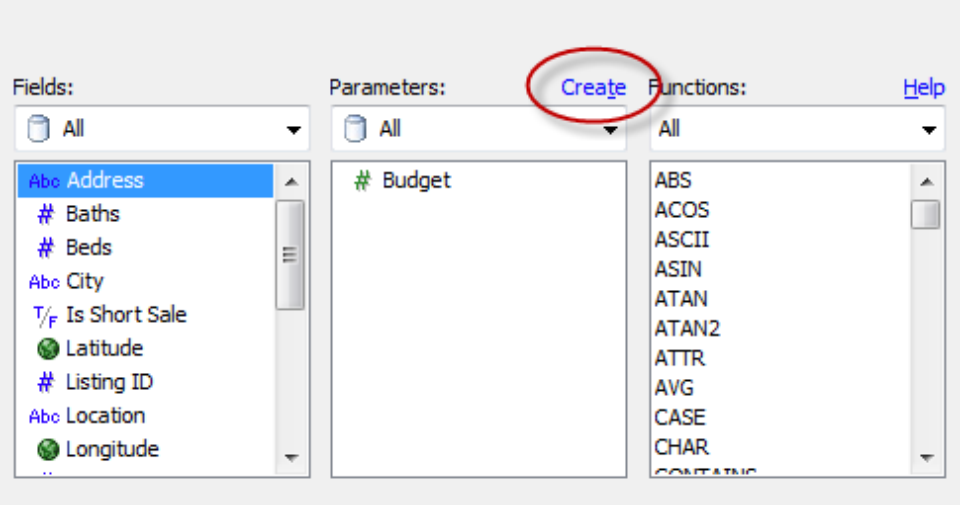

Similar to fields, parameters are written within square brackets in the formula. Parameters are shown as purple in the formula editor.

Using Parameters in Filters

Parameters give you a way to dynamically modify values in a TopN filters. Rather than manually setting the number of values you want to show in the filter, you can use a parameter. Then when you want to change the value, you open the parameter control and the filter updates. For example, when creating a filter to show the Top 10 products based on total profit, you may want to use a parameter instead of the fixed “10” value. That way, you can the quickly update the filter to show the top 10, 20, or 30 products.

A list of parameters is available in the drop-down lists on the Top tab of the Filter dialog box. Select the parameter you want to use in the filter.

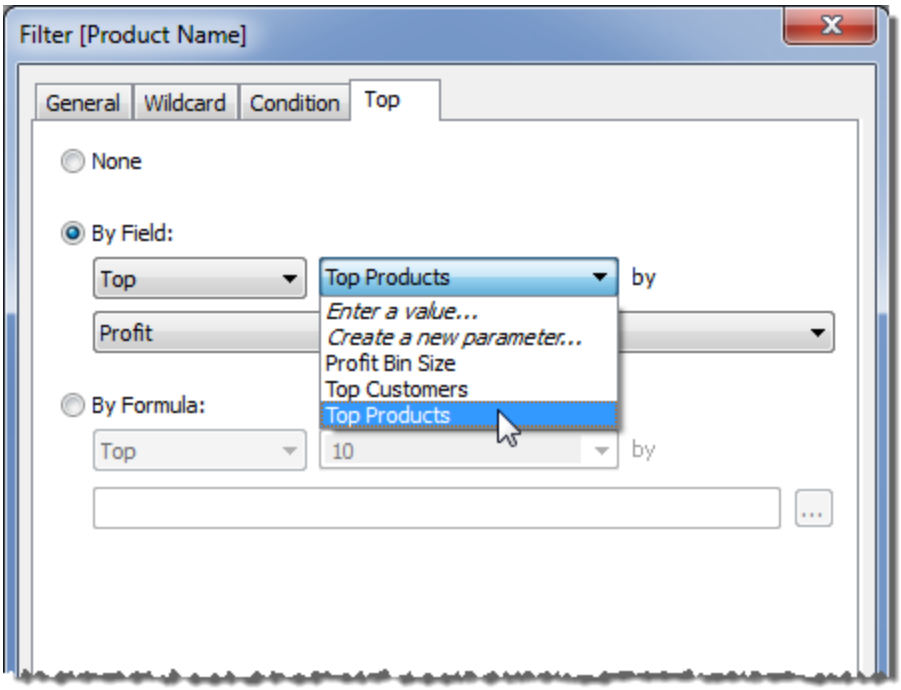

Then show the parameter control by right-clicking the parameter in the Data window and selecting **Show Parameter Control**. Use the parameter control to modify the filter to show the top 10 products, 15 products, 20 products, and so on.

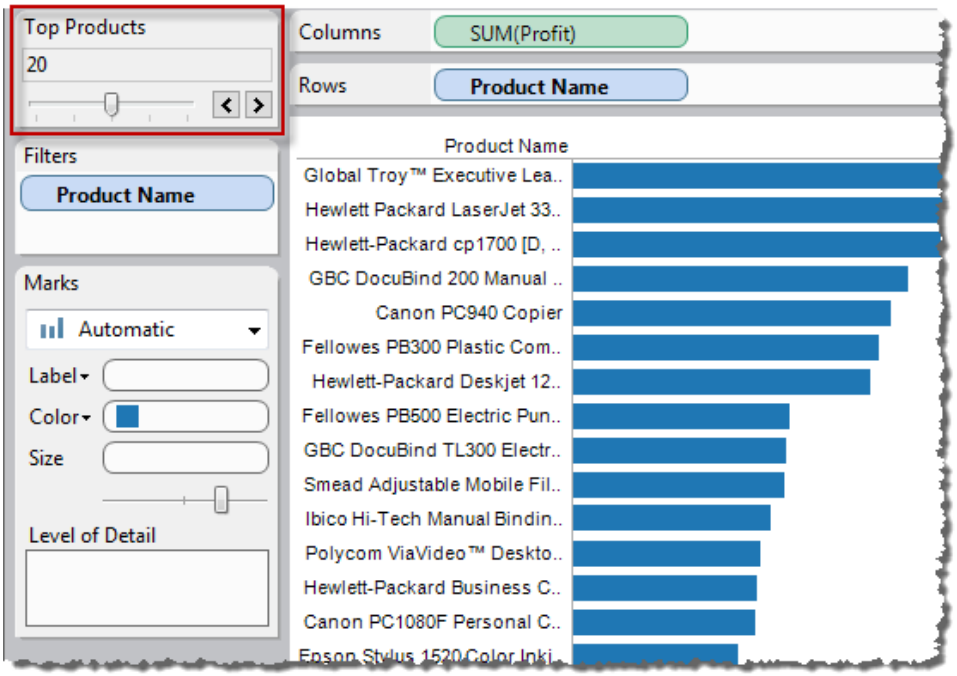

Using Parameters in Reference Lines

Parameters give you a way to dynamically modify a reference line, band, or box. For example instead of showing a reference line at a fixed location on the axis, you can reference a parameter. Then you can use the parameter control to move the reference line.

A list of parameters is available in the Value drop-down list in the Add Reference Line, Band, or Box dialog box. Select the parameter you want to use.

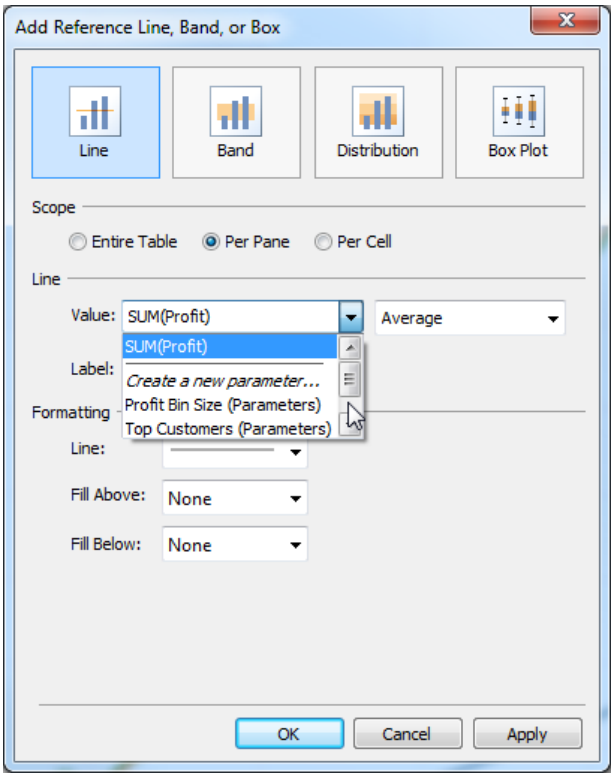

The reference line is drawn at the Current Value specified by the parameter. Open the parameter control by right-clicking the parameter in the Data window and then selecting **Show Parameter Control**. Use the parameter control to change where the reference line is drawn.

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**.

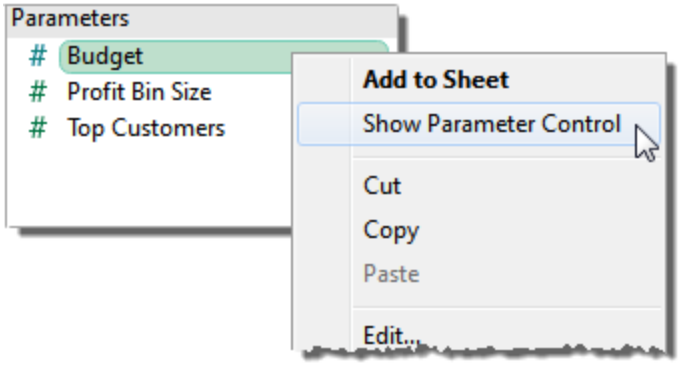

Like other cards, Parameter Controls have a menu that you can open using the drop-down arrow in the upper right corner of the card. Use this menu to customize the display of the control. For example, you can show a list of values as radio buttons, a compact list, a slider, or a type in field. The options 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

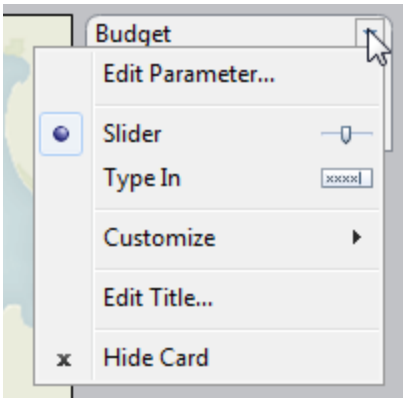

Example - Parameters

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.

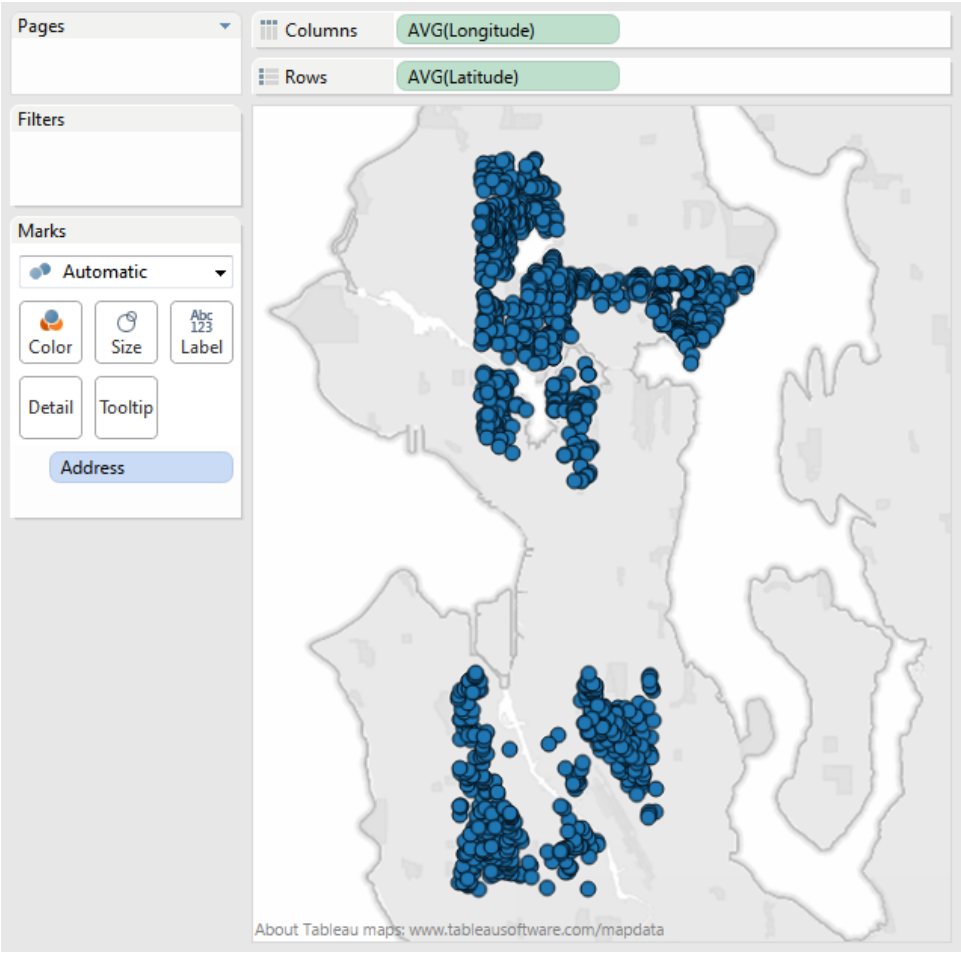

Follow the steps below to create a parameter that is used to color the houses based whether they fit into a specific budget.

1. Open the calculated field dialog box by selecting **Analysis** > **Create Calculated Field**.

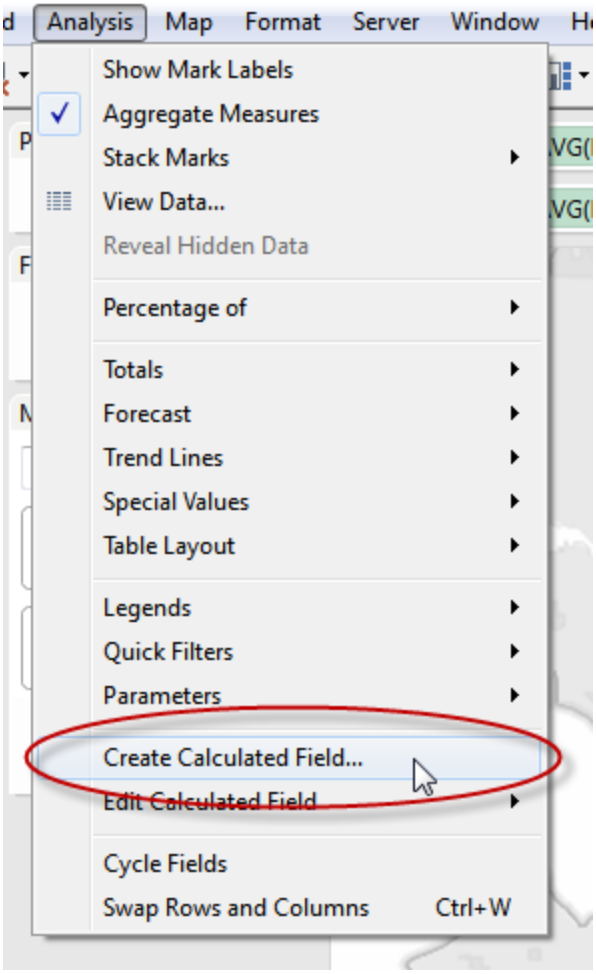

2. Name the field “Within Budget.”
3. In the formula type a boolean expression using the IF function. The formula is shown below:

```
IF [Sale Price] >= 500000 THEN "No" ELSE "Yes" END
```

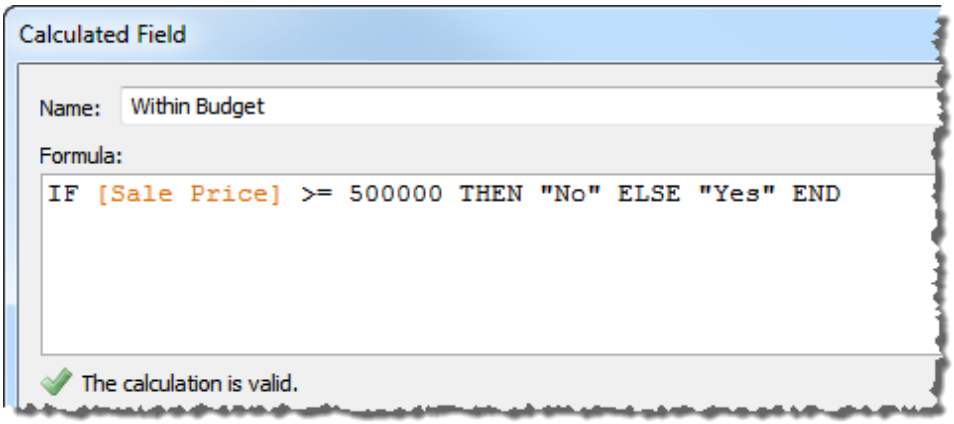

This formula will create a field that returns No if the sale price is over a budget of 500000. However, this threshold should be dynamic so that you can quickly change the budget without editing the calculation. We need to replace the constant value with a parameter.

4. Click the **Create** link at the top of the Parameters list in the Calculated Field dialog box.

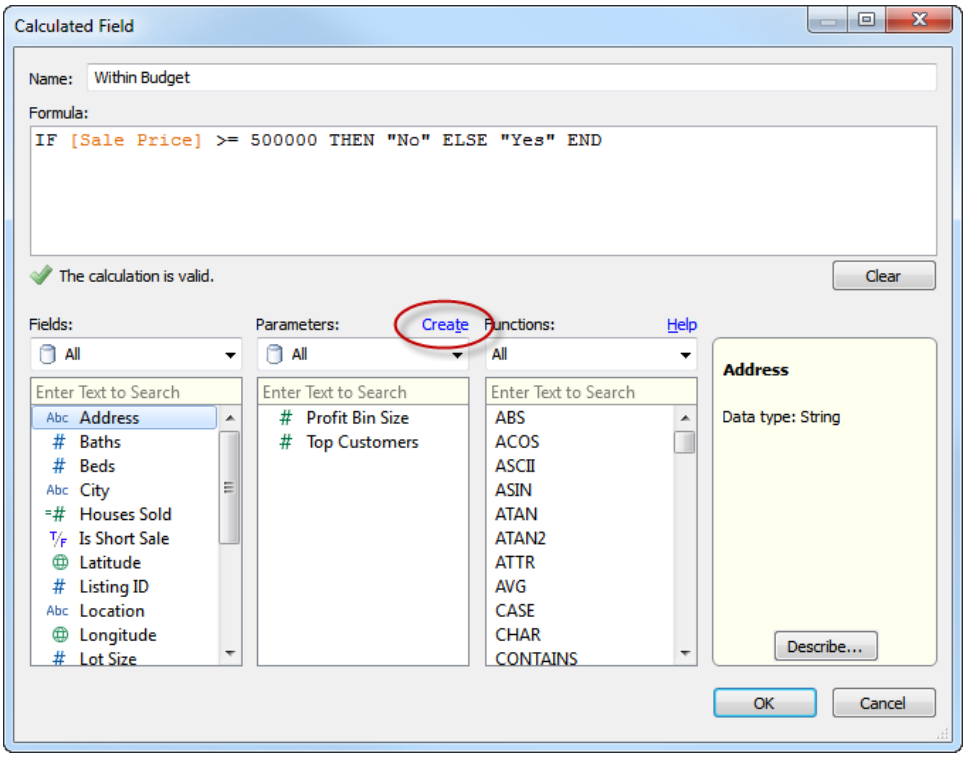

5. Name the parameter “Budget.”

6. Select the Float data type because the values will be currency.

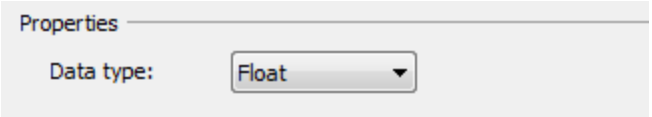

7. Set the **Current Value** to 500000. This is the default value for the parameter and just gives it a starting point.

8. Change the display format to **Currency (Standard)**.

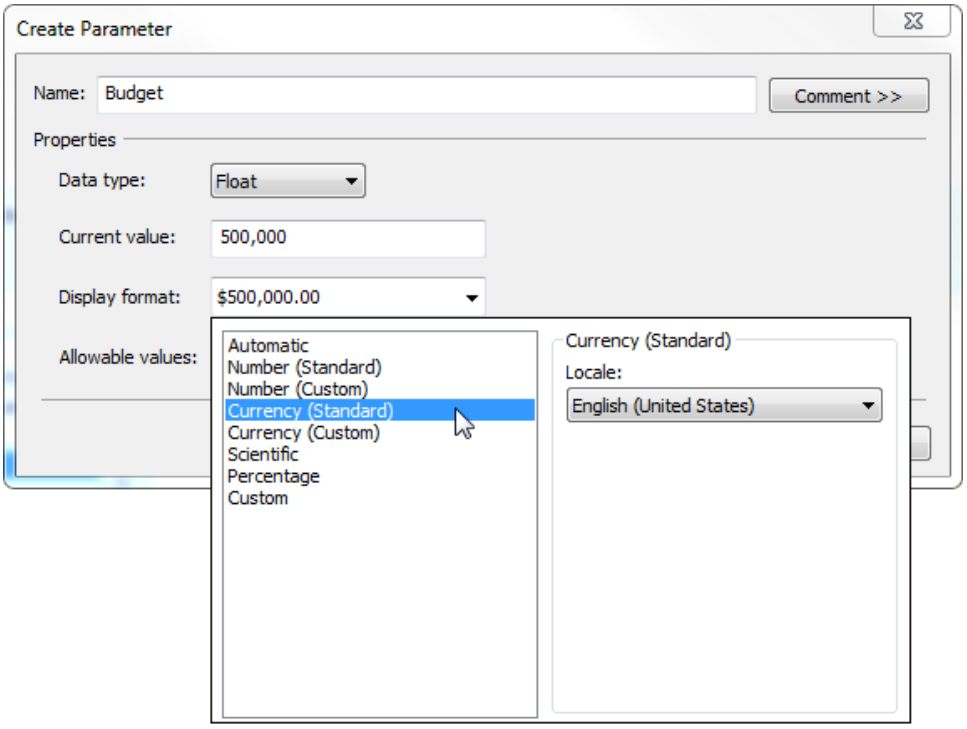

9. Allow a Range of values. Set the **Minimum** to 100,000 and the **Maximum** to 3,000,000. Set the **Step size** to 25,000.

   This definition means that the budget can be between $100K and $3M in increments of $25K.

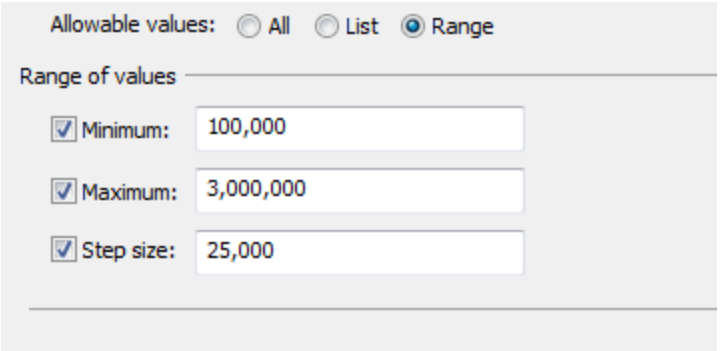

10. When finished, click **OK**.

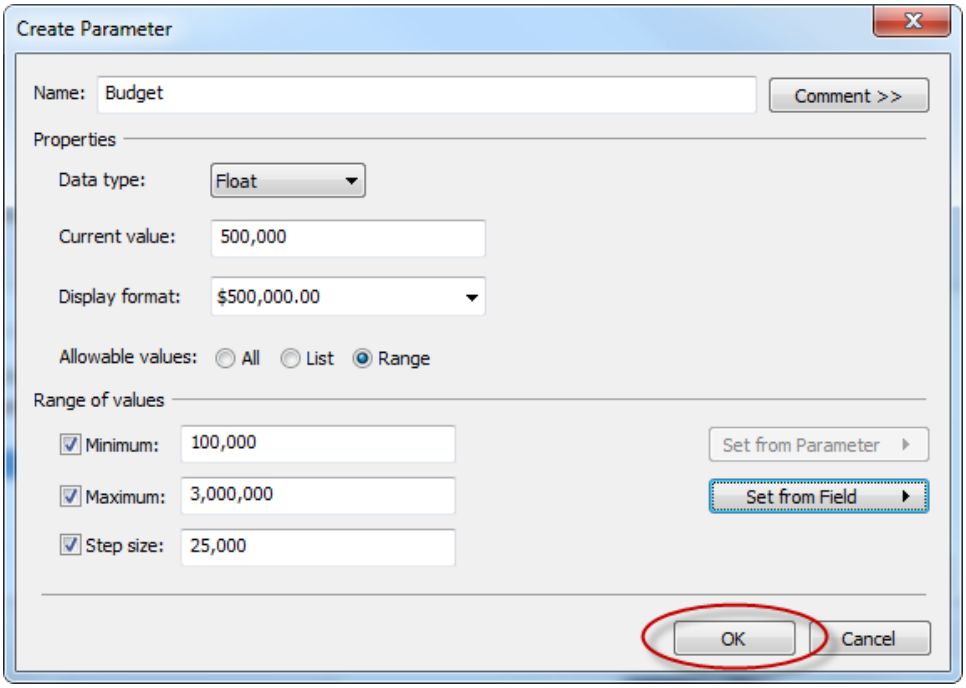

11. Back in Calculated Field dialog box, delete the constant value of “5000000” in the formula. Then double-click the Budget parameter in the Parameters list to add it to the formula. The formula now looks like this:

```
IF [Sale Price] >= [Budget] THEN "No" ELSE "Yes" END
```

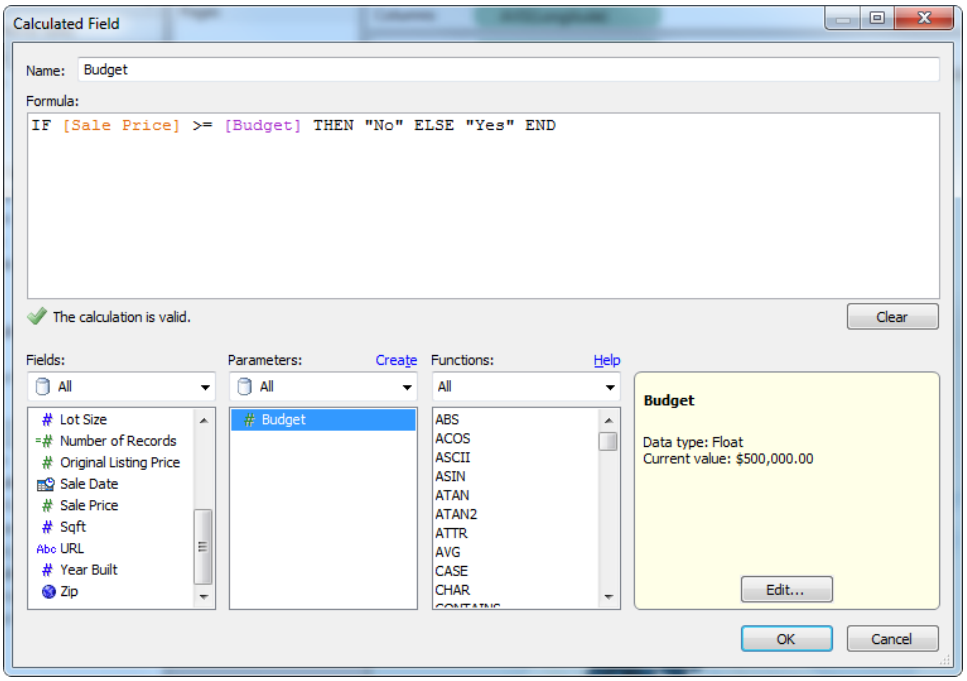

12. Make sure the calculation is valid and then click **OK**.
13. The new field is added to the Dimensions area of the Data window. Drag the **Within Budget** field to the **Color** shelf.

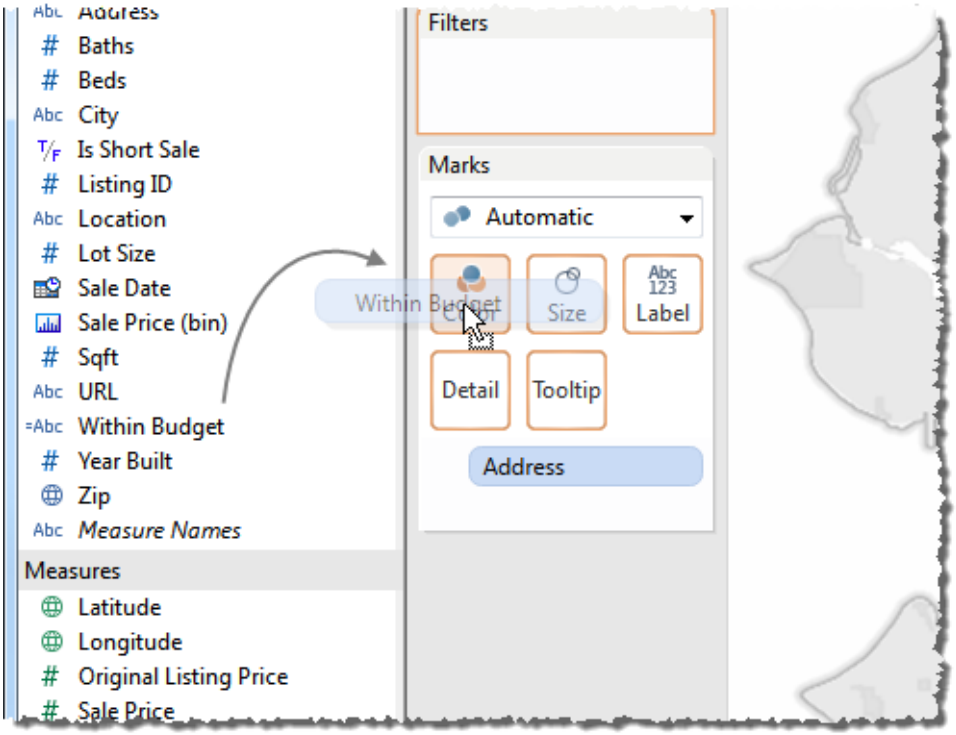

14. Right-click the Budget parameter at the bottom of the Data window and select **Show Parameter Control**.

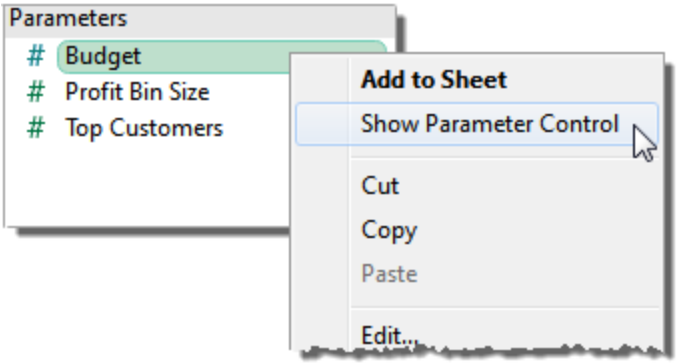

By default the control is shown as a slider. Clicking the right arrow in the control will increase the budget by $25K.

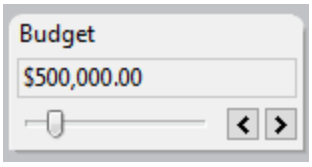

The view below shows the budget set to $400,000. Orange houses are within budget and blue are not.

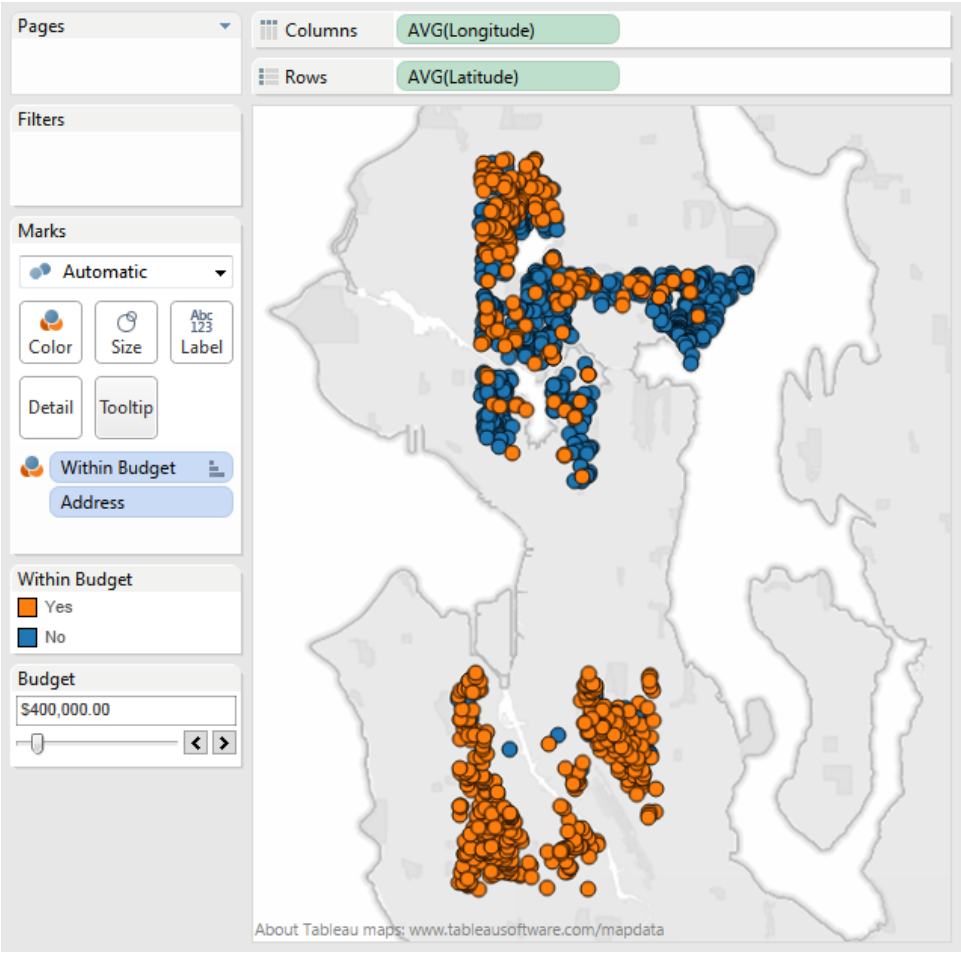

Background Images

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.

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 **Map** > **Background Images** and then select a data source.

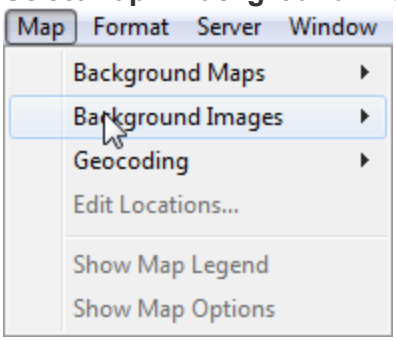

2. In the Background Images dialog box, click **Add Image**.

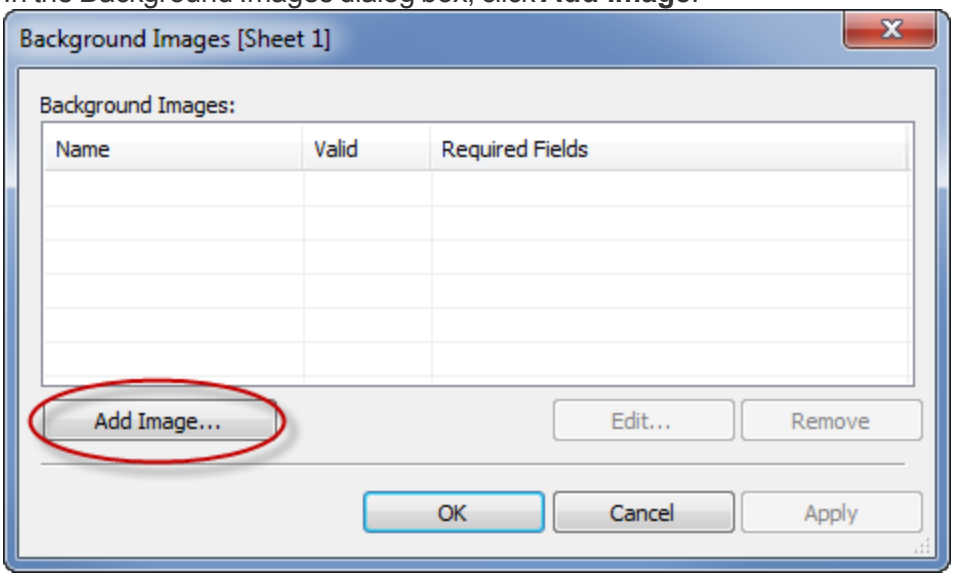

3. In the Add Background Image dialog box do the following:
   - Type a name for the image into the **Name** text box.
   - Click **Browse** to navigate to and select the image you want to add to the background. You can also type a URL to link to an image hosted online.
   - Select the field to map to 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.

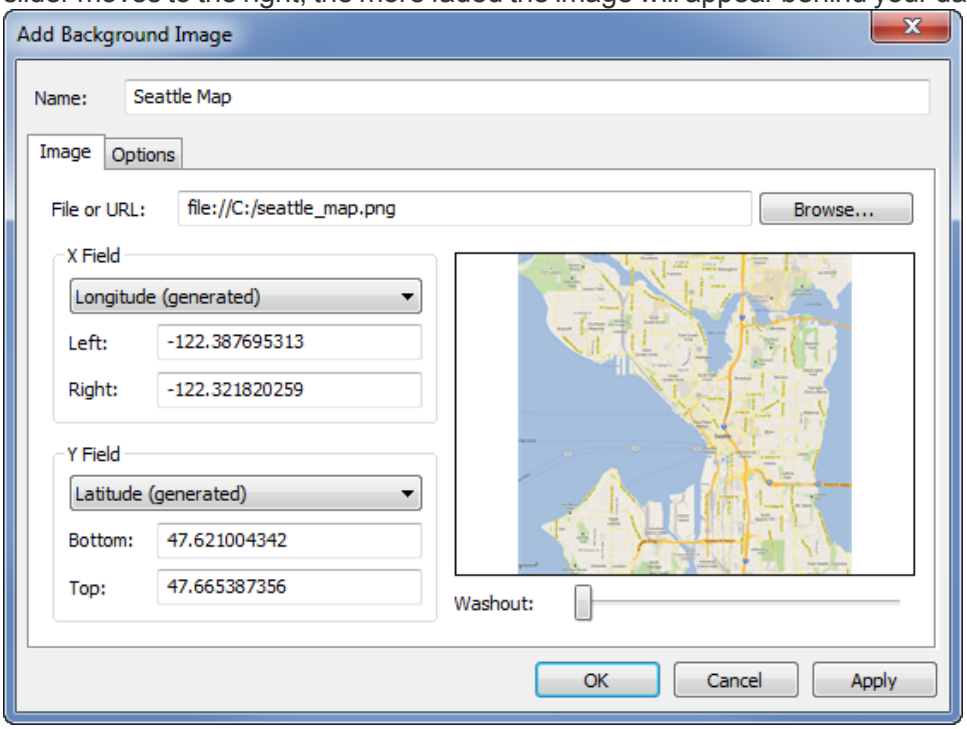

4. You can specify the following options using the **Options** tab:
   - **Lock Aspect Ratio** - select this option to maintain the original dimensions of the image for any manipulations of the axes. Deselecting this option allows the image's shape to be distorted.
   - **Always Show Entire Image** - 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 Conditions to Filters** on page 343 to learn more about defining conditions.

- 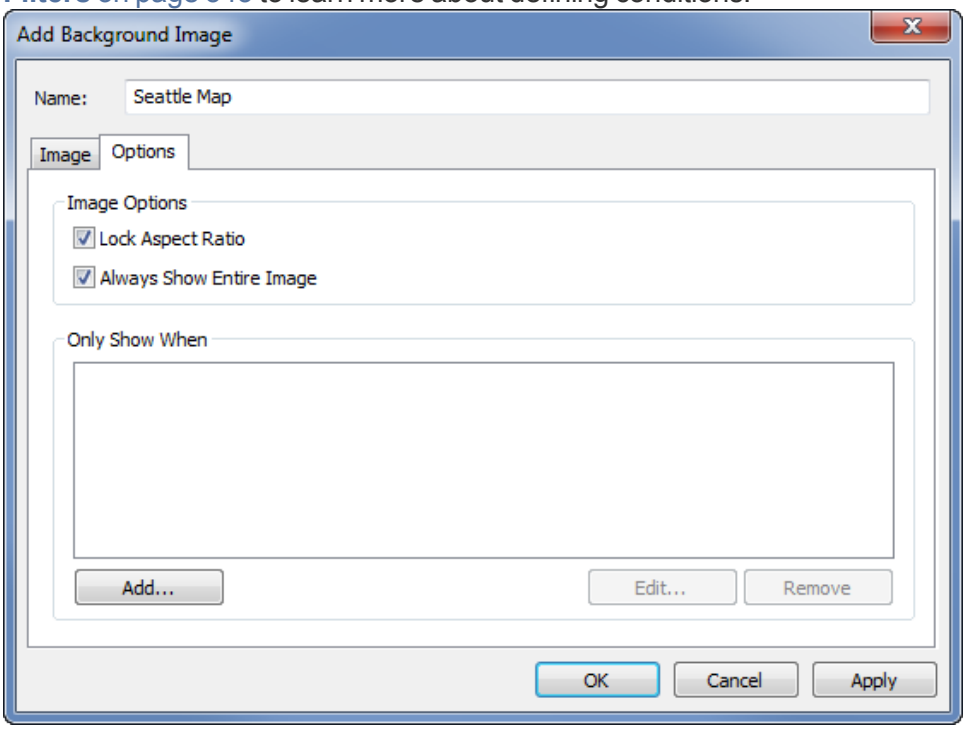

5. Click **OK**.

When you add the x and y fields to the Rows and Columns shelf in the view, the background image displays behind the data. If the background image does not display, make sure that you are using the disaggregated measures for the x and y fields. You can disaggregate all measures by selecting **Analysis > Aggregate Measures** or change each measure individually by right-clicking the field on the shelf and selecting **Dimension**. Finally, if you've used the generated Latitude and Longitude fields for the x and y fields, you'll need to disable the built in maps before your background image will display. Select **Map > Background Maps > None** to disable the built-in maps.

> 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**.

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.

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 **Map** > **Background Images**.

   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).

2. 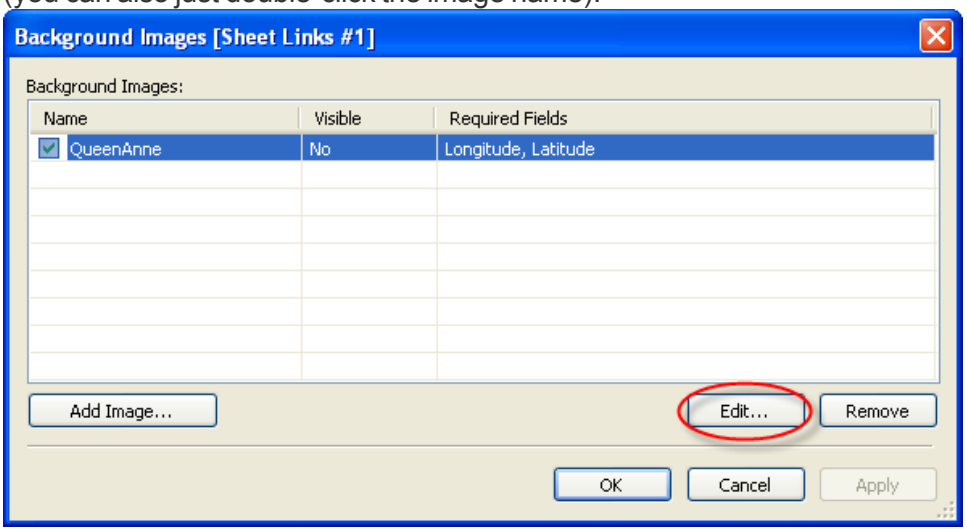

3. In the Edit Background Image dialog box, make the changes to the image and click **OK**.

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 **Map** > **Background Images**.

   In the Background Images dialog box, select the checkboxes next to the images you want enabled.

2. 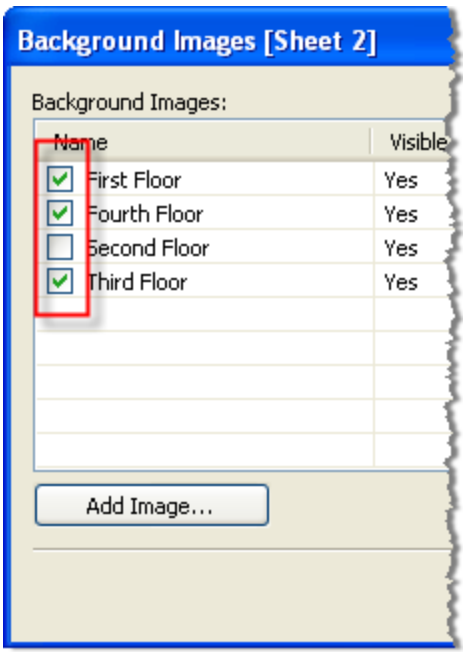

3. Click **OK**.

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 **Map > Background Images** and then select a data source.
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.

   Click the **Add** button at the bottom of the dialog box.

4. 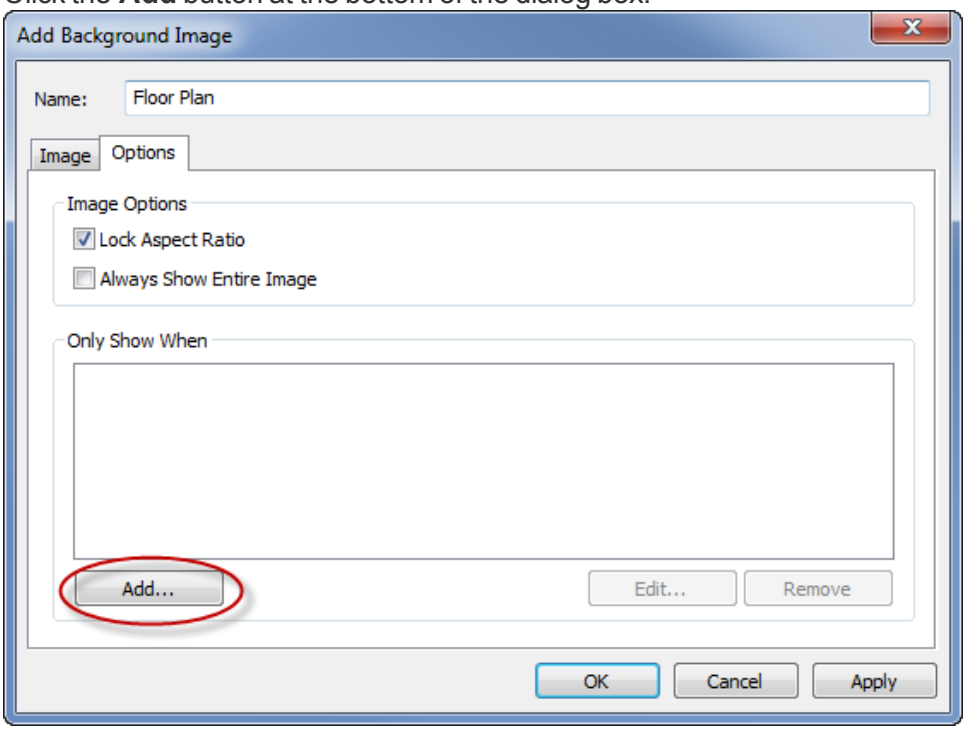

   Select a field to base the condition on. In the example described above, the field is Floor.

5. 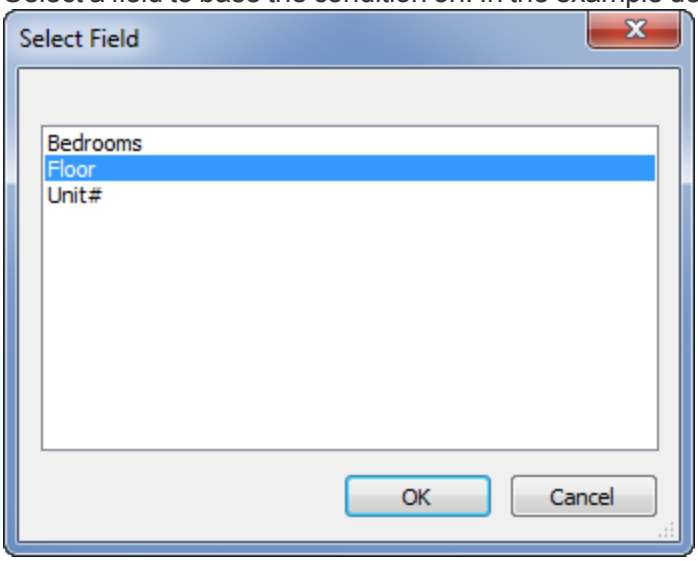

Specify when you want to show the image by selecting a value of the field. For this example, one is selected.

6. 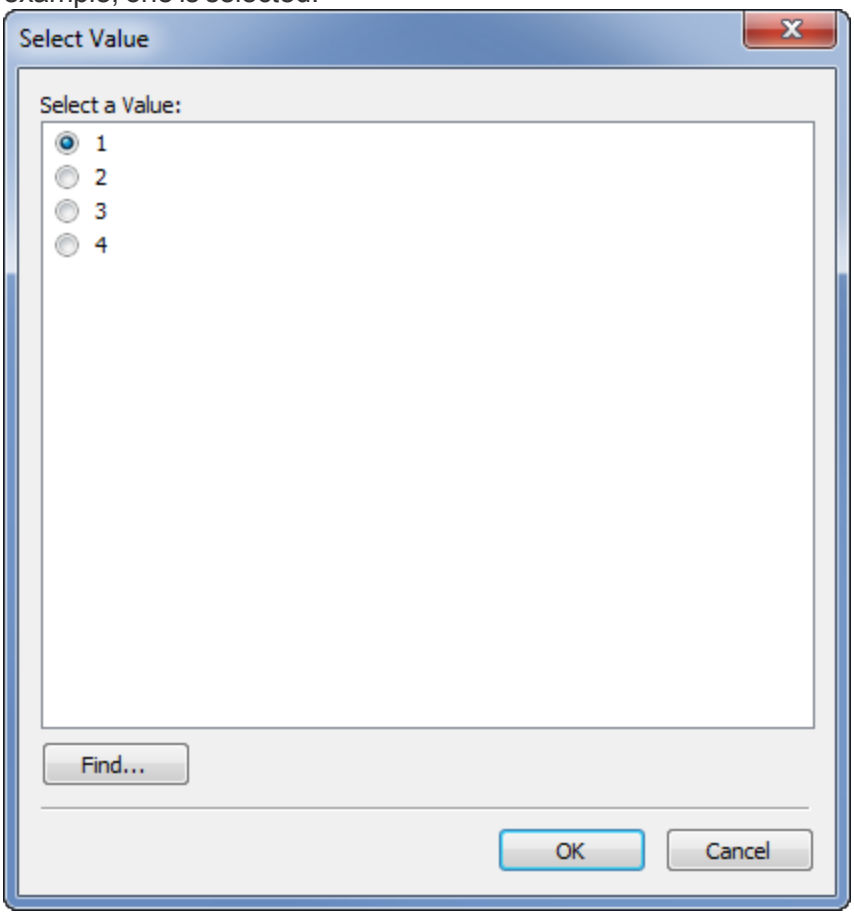

7. Click **OK**.

   A condition statement is added to the image. In the building floor plan example, the condition statement is Only show the image when Floor is equal to One.

8. Click **OK** twice to close the Background Image dialog boxes and apply the changes.

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 or remove it, making it unavailable to all worksheets.

To remove an image:

1. Select **Map** > **Background Images**.
2. In the Background Images dialog box, select the image you want to remove and click **Remove**.

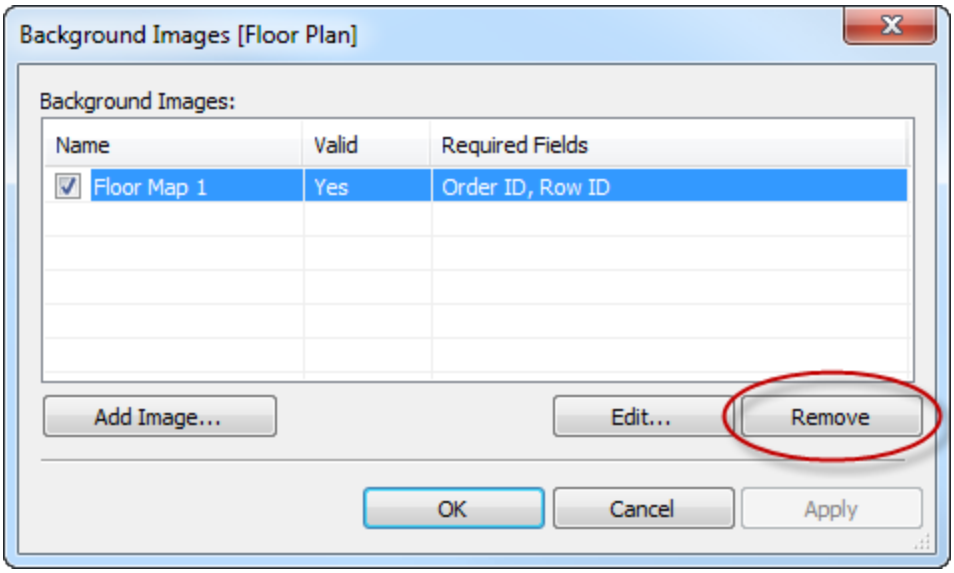

3. Click **OK**.

Background Maps

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. You can also learn more about maps by watching the following training videos from the Training and Tutorials page on Tableau's website:

Basic Maps

Background Images

Custom Geocoding

Polygon Maps

WMS Servers

Advanced Mapping Techniques

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 681 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.

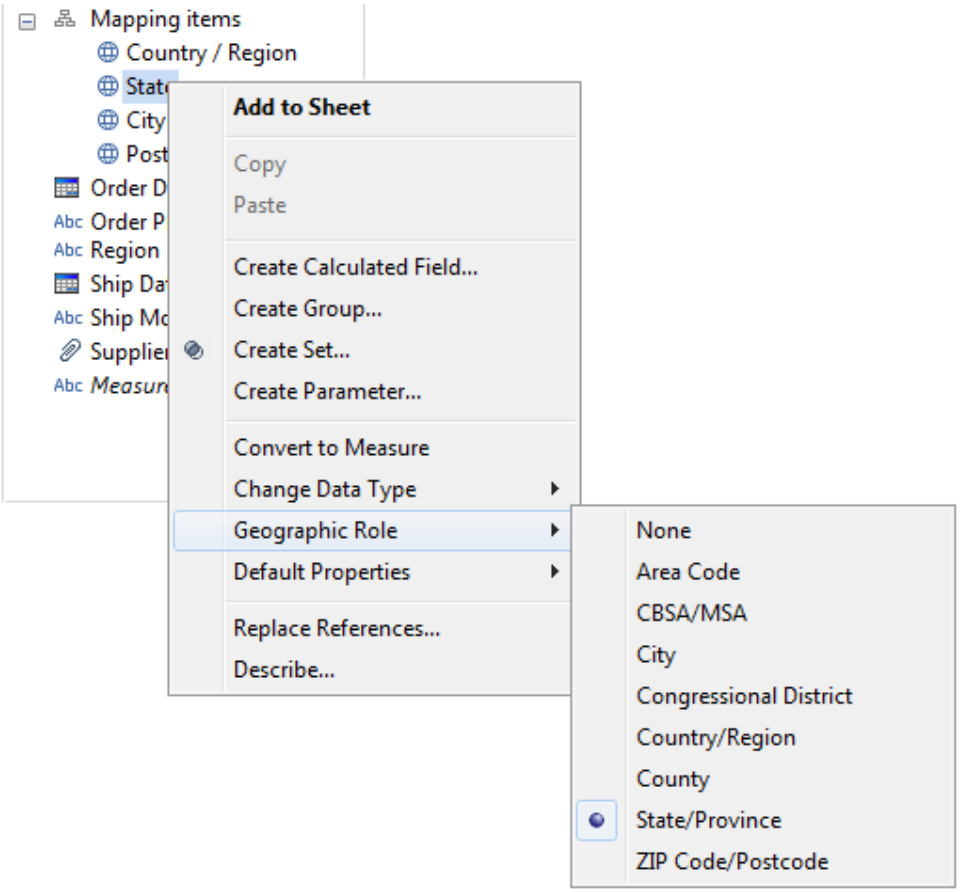

When you assign a field a geographic role it is marked in the Data window with a globe icon . This means that Tableau has automatically geocoded the information in that field and associated each value with a latitude and longitude value. Two automatic fields are added to the Measures area of the Data window: Latitude (generated) and Longitude (generated). These fields are available anytime you use Tableau to geocode your data.

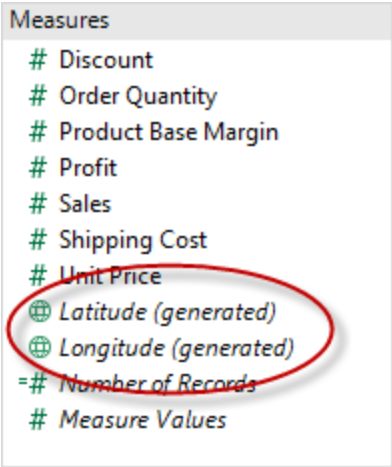

Below is a table describing some of the types of information Tableau can geocode.

Geographic Role	Description
Area Code	U.S. Area Codes; numbers only. Examples: 206, 650, 415
CBSA/MSA	U.S. Core Based Statistical Area or Metropolitan Statistical Area. Examples: Dallas-Fort Worth-Arlington, TX
City	Worldwide city names. Examples: Seattle, Bordeaux, Stuttgart
Congressional District	U.S. Congressional District names. Boundaries are determined by data provided by each state's redistricting commission. Examples: 1st District, 2, District 3, 4th
Country/Region	Worldwide countries and regions. Includes names, FIPS 10, 2 Character (ISO 3166-1), or 3 Character(ISO 3166-1). Examples: AF, CD, Japan, Australia, BH, AFG, UKR
County	Second-level administrative divisions for select countries. For example, U.S. counties, French départements, German kriese, etc. For more information about the second-level administrative divisions automatically geocoded in Tableau, see Mapping Basics in the Tableau Knowledge Base.
Latitude	Latitude in decimal degrees. Only available for numeric fields.
Longitude	Longitude in decimal degrees. Only available for numeric fields.
State/Province	Worldwide state, province, and other first-level administrative divisions. Names are in English, French, German, Spanish, Brazilian-Portuguese, Japanese, Korean, and Simplified Chinese.
ZIP Code/Post-code	ZIP Codes and Postcodes for the United States, France, Germany, United Kingdom, Canada, Australia, and New Zealand.

If your data contains geographic data that is not supported with the built-in roles, you can import custom geocoding to plot the data on a map. For more information on how to import custom geocoding into Tableau, see **Custom Geocoding** on page 702.

Building a Map View

When creating a map view, Tableau automatically adds the generated Latitude and Longitude fields to the Rows and Columns shelves and places a selected geographic field on Detail on the Marks card.

The marks in the view correspond to the field on Detail. For example, if the State field is used in the view, there is a single mark for each state in your data. As you add more geographic fields to Detail, the marks in the view will be broken down by the members of those fields. For more information, see **Building Views Manually** on page 291.

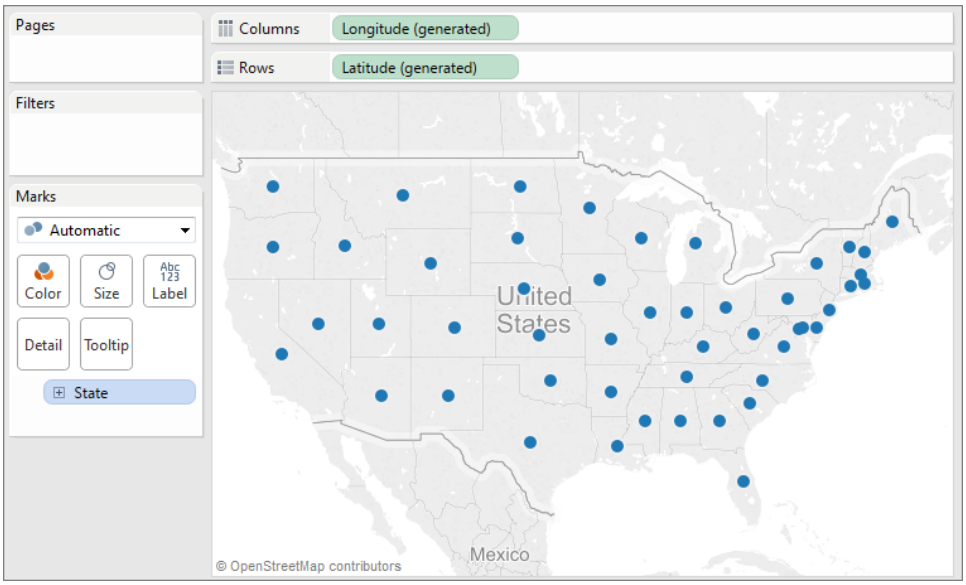

To build a map view:

- Double-click a geographic field in the Data window. A globe icon indicates fields that have been geocoded. For more information, see **Geographic Roles** on page 689.
- Select a geographic field in the Data window, and then select one of the map views in Show Me.

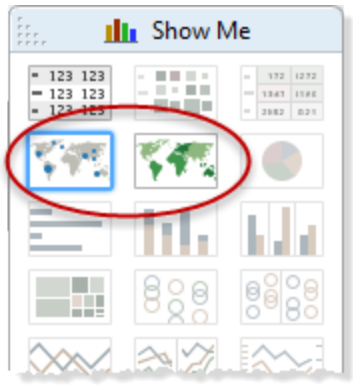

- In the Data window, double-click the measures **Latitude (generated)** and **Longitude (generated)**, and then drag a geographic field to **Detail** on the **Marks** card.

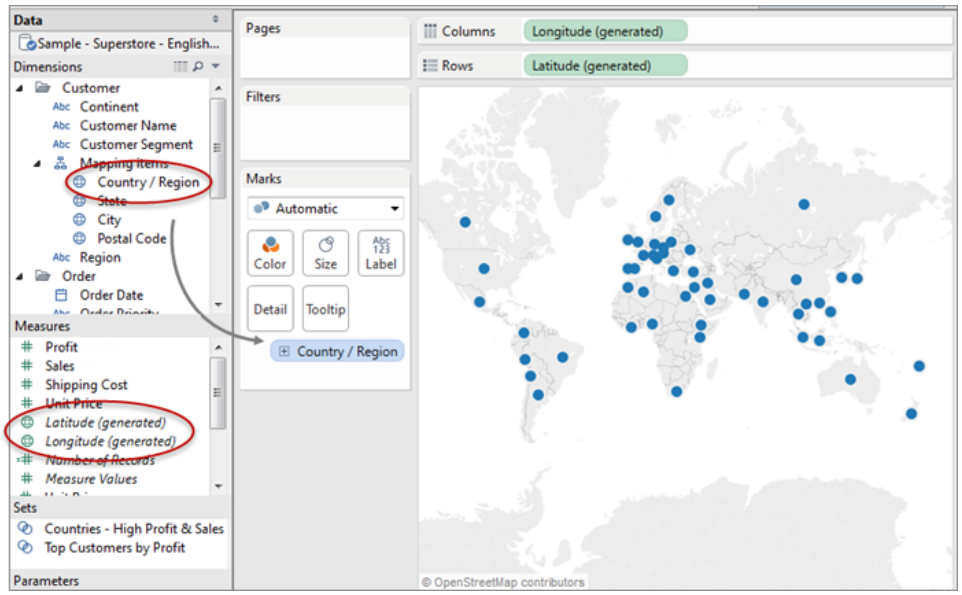

Map Options

When you load a map there are several options to help you control the display of the map. These options are available in the Map Options window. Open the Map Options window by selecting **Map** > **Map Options**.

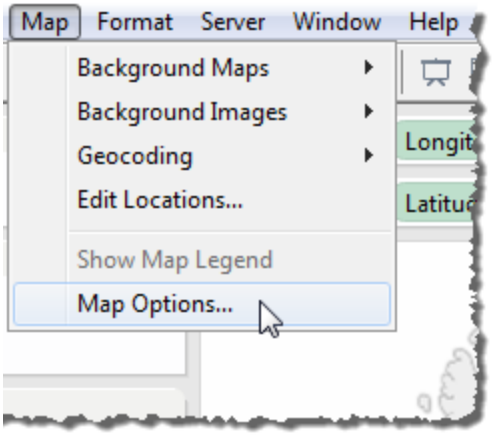

Use the Map Options window to modify how the map is displayed. You can change the map background, hide and show layers such as street names and places of interest, and add data layers.

Map Background

Select between at the top of the Map Options window.

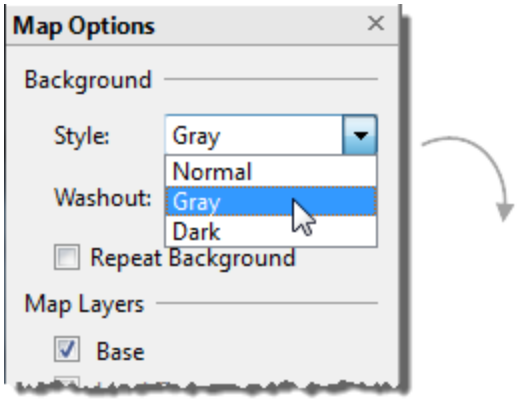

Normal **Gray** **Dark**

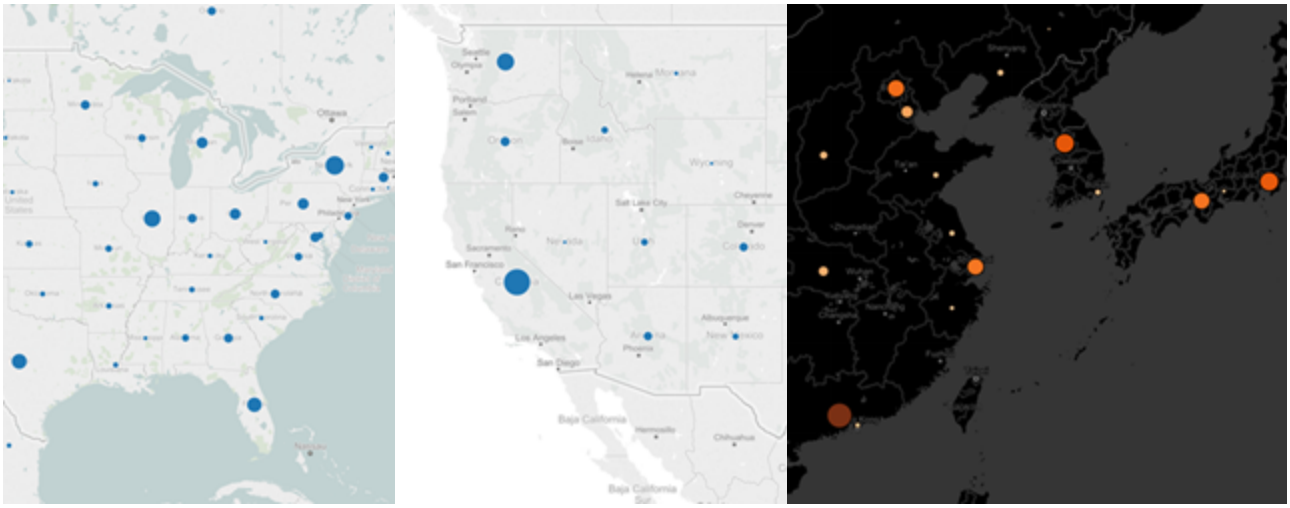

In addition to the map style, you can use the Washout slider to control the intensity, or luminance, of the background map. The farther the slider moves to the right, the more faded the map becomes.

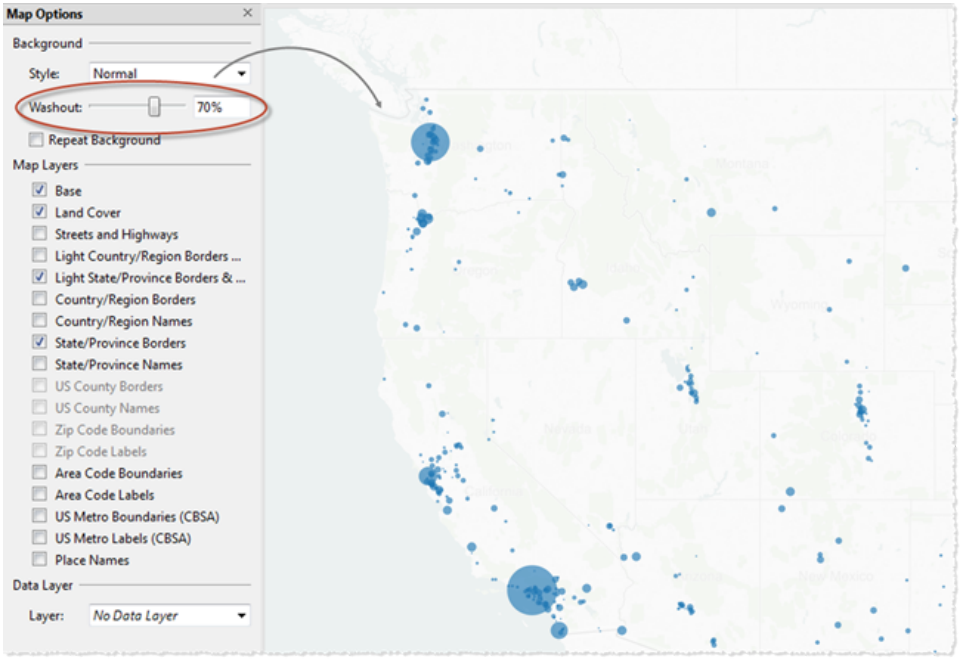

Finally, you can select whether to repeat the background. When the **Repeat Background** option it selected, the background map may show the same area multiple times depending on where the map is centered.

Map Layers

Tableau 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 window. Select **Map** > **Map Options** to show the Map Options window.

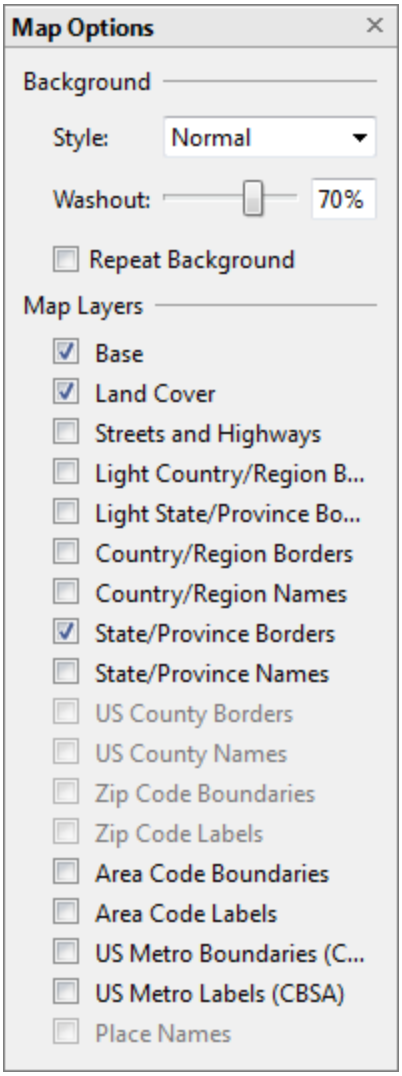

Many layers are only visible at specific zoom levels. Refer to **Zoom Controls** on page 474 to learn more about navigating the view. Each of the map layers are described in the table below.

Layer Name	Description
Base	Shows the base map including water and land areas.
Land Cover	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 Country/Region Borders	Shows a light gray outline of country/region borders and names.
Light Country/Region	Shows country and region names in a light gray.

Layer Name	Description
Names	
Light State/Province Borders	Shows a light gray outline of state borders and names.
Light State/Province Names	Shows state and province names in a light gray.
Country/Region Borders	Highlights country and region borders in a darker gray.
Country/Region Names	Highlights country and region names in a darker gray.
State/Province Borders	Highlights state and province borders in a darker gray.
State/Province Names	Highlights state and province names in a darker gray.
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 (CBSA)	Marks U.S. Metropolitan Statistical Areas and Micropolitan area boundaries.
US Metro Labels (CBSA)	Shows labels for the U.S. Metropolitan Statistical Areas and Micropolitan areas.
Place Names	Displays the names of places from country/region names and borders to city names, bodies of water, parks, universities, and more. This layer is dependent on the zoom level.
Coastlines	Displays a coastline only.

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.

1. Select the layers you want to show by default.
2. Click **Make Default** at the bottom of the Map Options window.

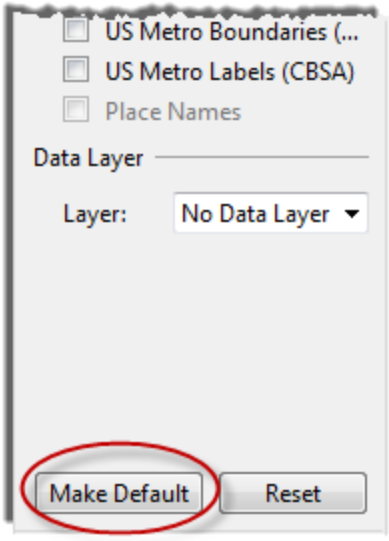

Map Data Layers

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 in the Map Options window. Hide and show the legend by selecting **Map** > **Show Map Legend**.

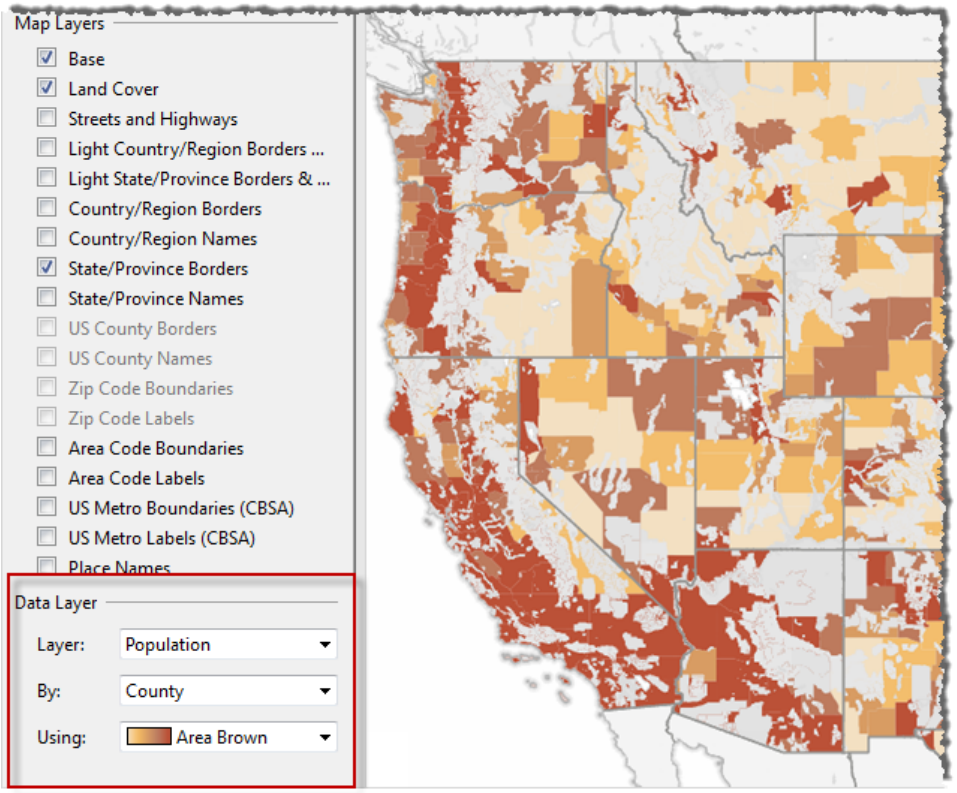

Setting a Default Location

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 edit the locations to tell Tableau which country and state your data refers to.

You can edit locations by selecting **Map > Edit Locations**. Then select the country and/or state that your data refers to.

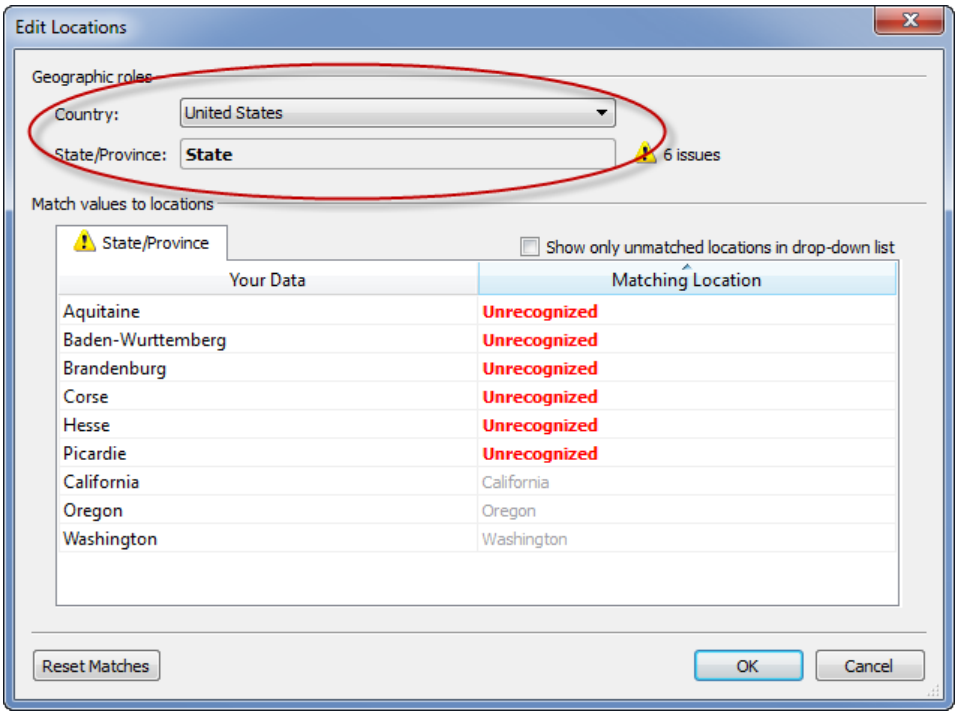

Editing Locations

Sometimes Tableau will not recognize one or more of the location names in your data. When this happens those values are marked as unrecognized in the lower right corner of the view. 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 clicking the special values indicator and the selecting **Edit Locations**. You can also select **Map > Edit Locations**.

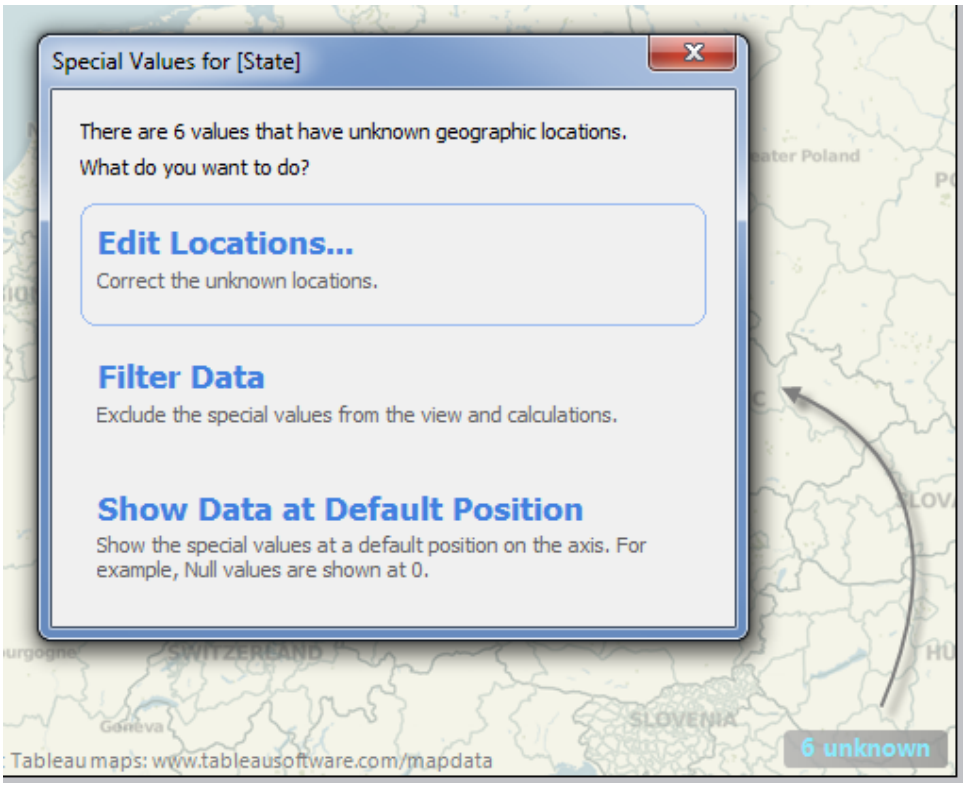

In the Edit Locations dialog box, both unknown and ambiguous locations are listed. You can usually fix ambiguous locations by specifying the country and/or state your data refers to. For each unrecognized value in your data, select a known location in the right column. For example, below Carson is an unrecognized city name in Nevada. Select Carson City to map it to the correct location. Alternatively, you can select **Enter Latitude and Longitude** to manually map the value to a location on the map. When you type a latitude and longitude, enter the values in decimal format.

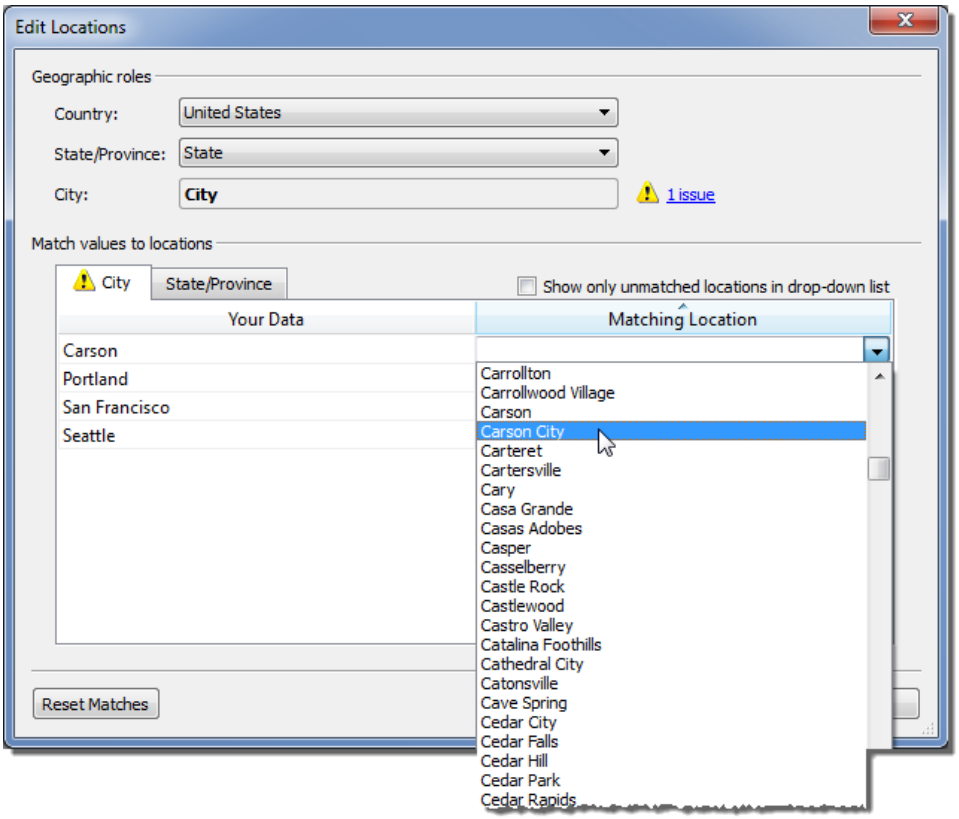

By default all possible locations are listed in the drop-down list. For example, above all cities in the United States are listed as options. Select the **Show only unmatched locations in drop-down list** option to restrict the list to locations that have not yet been matched to your data.

Custom Geocoding

Tableau Desktop recognizes a set of geographic roles that can be used to automatically geocode your data and create map views. For example, Tableau Desktop recognizes country names, state/province names, city names, and area codes. For more information, see **Geographic Roles** on page 689.

If your geographic data does not fit into the built-in geographic roles, you can create new geographic roles and assign them to the geographic fields in your data. For example, if your data contains country, state/province, and street address data, Tableau Desktop will aggregate your data to the country and state/province level, but will not recognize the street address data as a geographic role. In this case, you should create a custom geographic role for the street address data.

To create a new geographic role, you must create and import a custom geocoding file that defines the hierarchies and locations of your geographical data into Tableau.

Creating an Import File

A custom geocoding file is a .csv file. The contents of the custom geocoding file differ depending on whether you are extending an existing geographic hierarchy, adding a new role, or adding a new hierarchy to Tableau.

However, the following applies to all custom geocoding files:

- The .csv file must contain the Latitude and Longitude columns.
- Latitude and Longitude values must be real numbers. Make sure to include at least one decimal place when specifying these values.

> **Note:** A numeric field in your data should be treated as a text field on import (for example, some post codes). If you are importing numeric fields in to Tableau, you need to create a schema.ini file to tell Tableau that the numeric field you want to import should be treated as a text field. For more information on how to create a schema.ini file, see Importing Numeric Geographic Roles in the Tableau Knowledge Base.

Extending an Existing Role

The built-in geographic roles in Tableau contain hierarchies that can be extended to include locations relevant to your data. For example, the existing hierarchy of Country > State/Province may not contain all of the states or provinces in your data. You can extend this level to include missing states or provinces. The import file for this type of geocoding should contain every level of the hierarchy above the level you are extending. For example, the existing State/Province hierarchy has a column for Country and State/Province, along with Latitude and Longitude.

	A	B	C	D
1	**Country (Name)**	**State/Province**	**Latitude**	**Longitude**
2	England	United Kingdom	51.5000	0.1167
3	Scotland	United Kingdom	55.9500	3.1833
4	Wales	United Kingdom	51.4833	3.1833
5	Northern Ireland	United Kingdom	54.6000	5.9167

In the import file, the names of the columns define the geographic roles. When extending an existing role, the column names must match the existing geographic roles in the hierarchy that you are extending. This will ensure that the new locations are added to the proper roles and hierarchies. For more information on how to organize hierarchies in your import file, see **Built-In Hierarchies** on page 705.

Once you have created your import file, save the file as a Comma Delimited (.csv) file (Windows Comma Separated if on a Mac) in a folder on your computer.

> **Note:** A numeric field in your data should be treated as a text field on import (for example, some post codes). If you are importing numeric fields in to Tableau, you need to create a schema.ini file to tell Tableau that the numeric field you want to import should

be treated as a text field. For more information on how to create a schema.ini file, see Importing Numeric Geographic Roles.

Adding New Roles

To add new geographic roles to the existing geographic hierarchy in Tableau, format your import file to include the new roles along with their parent roles.

When adding new roles to an existing hierarchy, the import file for those roles needs to contain the columns for each level in the existing hierarchy. For more information on the columns to include in each geographic hierarchy, see **Built-In Hierarchies** on the next page.

Below is an example of an import file containing airport codes. Importing the file below would add the geographic roles Airport (ICAO), Airport (IATA), and Airport (City) to the existing Country (Name) hierarchy. Notice that the column name for country matches the existing Country (Name) geographic role.

	A	B	C	D	E	F
1	**Airport (ICAO)**	**Airport (IATA)**	**Airport (City)**	**Country (Name)**	**Latitude**	**Longitude**
2	AYGA	GKA	GOROKA	PAPUA NEW GUINEA	-6.08167	145.39167
3	AYLA	LAE	LAE	PAPUA NEW GUINEA	0	0
4	AYMD	MAG	MADANG	PAPUA NEW GUINEA	-5.20694	145.78861
5	AYMH	HGU	MOUNT HAGEN	PAPUA NEW GUINEA	-5.82611	144.29611
6	AYNZ	LAE	NADZAB	PAPUA NEW GUINEA	-6.56972	146.72611
7	AYPY	POM	PORT MORESBY	PAPUA NEW GUINEA	-9.44333	147.22
8	AYRB	RAB	RABAUL	PAPUA NEW GUINEA	0	0
9	AYWK	WWK	WEWAK	PAPUA NEW GUINEA	-3.58361	143.66917
10	BGAM	N/A	ANGMAGSSALIK	GREENLAND	0	0
11	BGAS	N/A	ANGISSOQ	GREENLAND	0	0
12	BGAT	N/A	APUTITEQ	GREENLAND	0	0

Once you have created your import file, save the file as a Comma Delimited (.csv) file (Windows Comma Separated if on a Mac) in a folder on your computer.

Note: A numeric field in your data should be treated as a text field on import (for example, some post codes). If you are importing numeric fields in to Tableau, you need to create a schema.ini file to tell Tableau that the numeric field you want to import should be treated as a text field. For more information on how to create a schema.ini file, see Importing Numeric Geographic Roles.

Adding New Hierarchies

To add a new hierarchy in Tableau, you must create multiple import files, each representing a level in the new hierarchy. For example, you may have custom sales that are defined by Theater, Country, Region, and City. To create a new hierarchy from these geographic roles, you 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 include Theater, Latitude, and Longitude.

	A	B	C
1	**Theater**	**Latitude**	**Longitude**
2	Shakespeare's Globe theater	51.508342	-0.096984
3	The Düsseldorfer Schauspielhaus	51.227420	6.783620
4	Balboa Theatre	32.714329	-117.161364

The next level is Country, so its import file should contain columns for Theater, Country, Latitude and Longitude.

	A	B	C	D
1	**Theater**	**Country**	**Latitude**	**Longitude**
2	Shakespeare's Globe theater	England	51.508342	-0.096984
3	The Düsseldorfer Schauspielhaus	Germany	51.227420	6.783620
4	Balboa Theatre	United States	32.714329	-117.161364

Make sure that each file in the hierarchy contains columns for the current level and all levels above it. For example, the City import file, the last geographic role in the hierarchy, would contain columns for Theater, Country, Region and City.

	A	B	C	D	E	F
1	**Theater**	**Country**	**Region**	**City**	**Latitude**	**Longitude**
2	Shakespeare's Globe theater	England	London	London	51.508342	-0.096984
3	The Düsseldorfer Schauspielhaus	Germany	North Rhine-Westphalia	Düsseldorf	51.227420	6.783620
4	Balboa Theatre	United States	Anza Borrego	San Diego	32.714329	-117.161364

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, country.csv, region.csv, and city.csv.

> **Note:** A numeric field in your data should be treated as a text field on import (for example, some post codes). If you are importing numeric fields in to Tableau, you need to create a schema.ini file to tell Tableau that the numeric field you want to import should be treated as a text field. For more information on how to create a schema.ini file, see Importing Numeric Geographic Roles.

Built-In Hierarchies

When extending an existing role or adding a new role, your custom geocoding import file should contain every level of the hierarchy above the level you are extending or adding. Below is a list of built-in hierarchies and how they should be organized in your import file.

Built-In Hierarchy	**Columns to Include in the .csv File (in order; left to right)**
Country (Name)	Country (Name) Latitude Longitude
Country (Name), State/Province	Country (Name) State/Province

	Latitude Longitude
Country (Name), State/Province, City	Country (Name) State/Province City Latitude Longitude
Country (Name), State/Province, County	Country (Name) State/Province County Latitude Longitude
Country (Name), ZIP Code/Post-code	Country (Name) ZIP Code/Postcode Latitude Longitude
Country (Name), Area Code	Country (Name) Area Code Latitude Longitude
Country (Name), CBSA	Country (Name) CBSA Latitude Longitude

In addition to the Country (Name) column, you can optionally include the following columns: Country 2 char (ISO 3166-1), Country 3 char (ISO 3166-1), and Country (FIPS 10). If you include these columns, they should be just to the right of the Country (Name) column in any order.

Latitude and Longitude values must be real numbers instead of integers. That is, make sure to include at least one decimal place when specifying these values.

Note: The highest level in the hierarchy is Country and cannot be extended to include higher levels such as Continent, etc.

Create a schema.ini File

A schema.ini file tells Tableau how to interpret the contents of text files. Creating a schema.ini file specifying that the numeric postcodes in your import file should be treated as text, for example, will allow you to plot the postcodes on a map. For more information on how to create and format a schema.ini file, see Importing Numeric Geographic Roles in the Tableau Knowledge Base.

Below is an example of a schema.ini file for a .csv file with the columns German Post Code, Latitude, and Longitude.

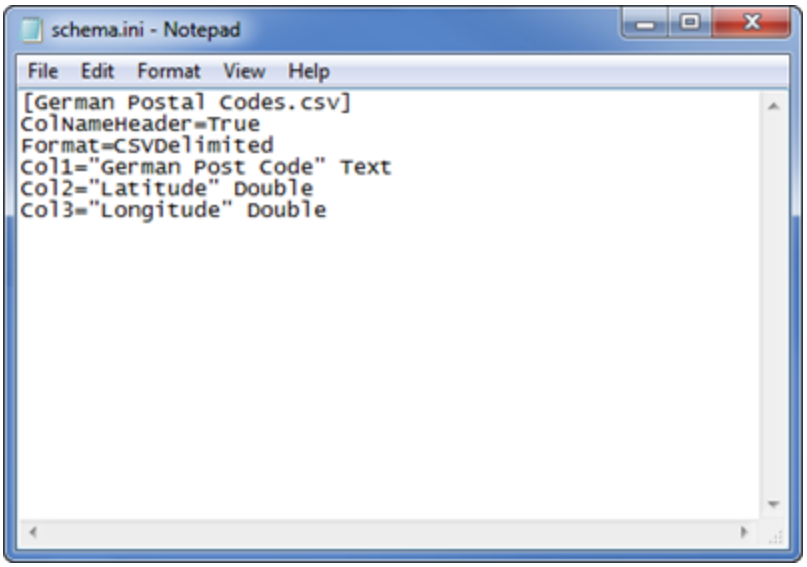

Save the file as schema.ini in the same folder as the .csv file.

Importing Custom Geocoding

Once you have created a .csv file with custom geocoding you can import that file into Tableau. Follow the steps below to import a new geographic role.

1. Click **Map** > **Geocoding** > **Import Custom Geocoding**.
2. In the Import Custom Geocoding dialog box, click the button to the right of the text field to browse to the folder your import file is saved in. All files in the folder will be imported into Tableau.

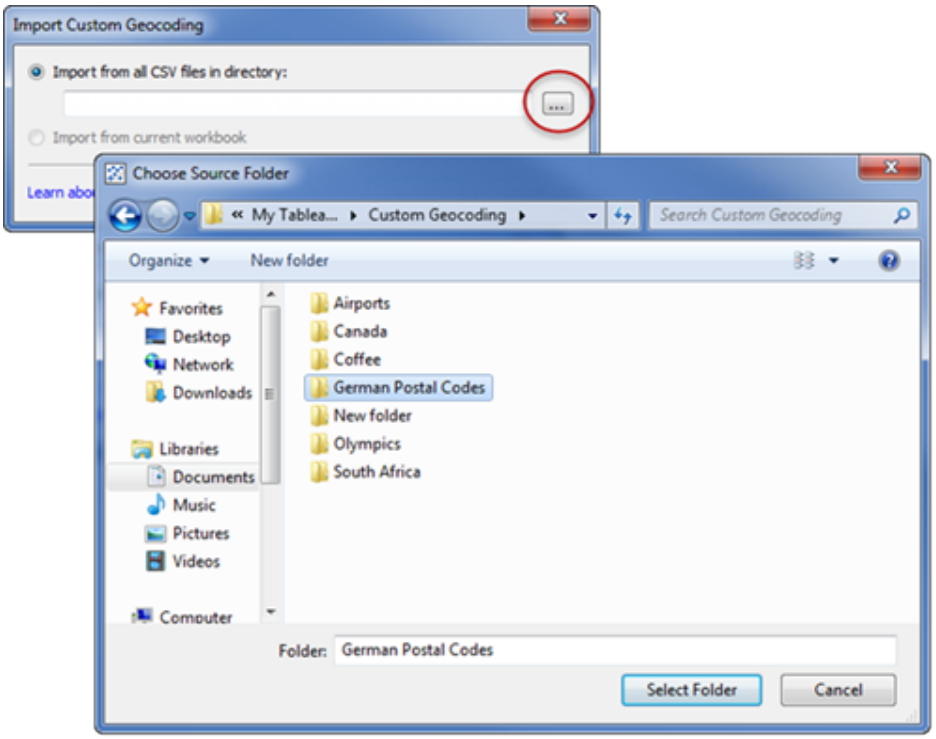

3. When finished, click **Import**.

The custom geocoding data is imported into the workbook and the new geographic roles become available. To assign a new geographic role to a field, right-click the field in the Data window,select **Geographic Role**, and then select the role you want to assign it. For more information on assigning geographic roles, see **Geographic Roles** on page 689.

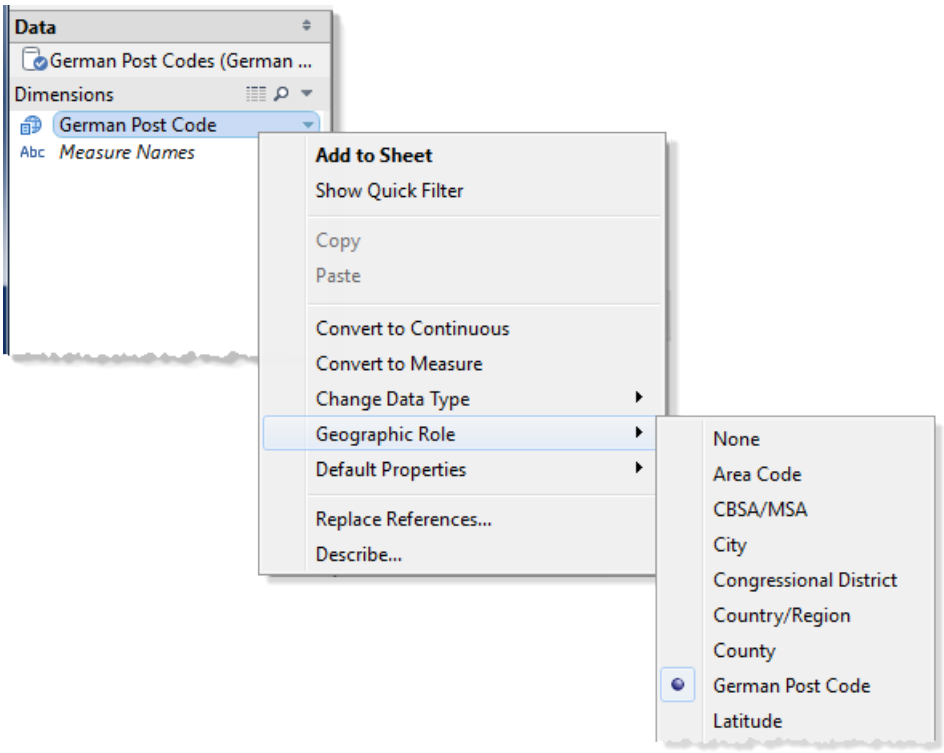

Custom Geocoding Data

When you import custom geocoding, the data is stored in the Local Data folder in your Tableau Repository. To import custom geocoding, your Tableau Repository must be on a local hard drive. The custom geocoding is then available for all workbooks.

Save custom geocoding files in a separate folder on your local computer. If you have several different sets of custom geocoding files, save each of them under their own folder names.

You can remove the custom geocoding stored in your Tableau Repository by clicking **Map** > **Geocoding** > **Remove Custom Geocoding**. This will not remove the geocoding from a packaged workbook, but it will remove it from the Local Data folder in your Tableau Repository.

When you save your workbook as a packaged workbook, the custom geocoding data is packaged with the workbook. When you open a packaged workbook you can import the custom geocoding from that workbook into your own Tableau Repository.

> **Note:** Importing a new custom geocoding file will replace any custom geographic roles previously imported into Tableau.

Example - Importing Numeric Postcodes

Follow the steps below to import numeric postcodes into Tableau and create a new postcode geographic role. If a field in your data contains only numerical geographical data, such as ZIP

codes or postcodes, you will need to create a schema.ini file before importing the geocoding file into Tableau.

1. Create a .csv file.

	A	B	C	D
1	German Post Code	Latitude	Longitude	
2	99998	51.2125	10.5625	
3	99996	51.28335	10.575	
4	99994	51.24165	10.675	
5	99991	51.15	10.55557	
6	99988	51.175	10.2917	
7	99986	51.14443	10.4389	
8	99976	51.2477	10.32917	
9	99974	51.2428	10.4501	
10	99958	51.09585	10.7583	
11	99955	51.17273	10.80151	
12	99947	51.14538	10.62537	
13	99898	50.8333	10.5833	
14	99897	50.8	10.6167	
15	99894	50.86305	10.5944	
16	99891	50.8854	10.459	

Note: To add German Postcodes to the existing Country hierarchy, include a column for Country (Name) in the .csv file.

2. Save the .csv file in a folder on your local computer.
3. Create a schema.ini file.

   An example of a schema.ini file for an import file titled German Postal Codes.csv with the columns German Post Code, Latitude, and Longitude is below.

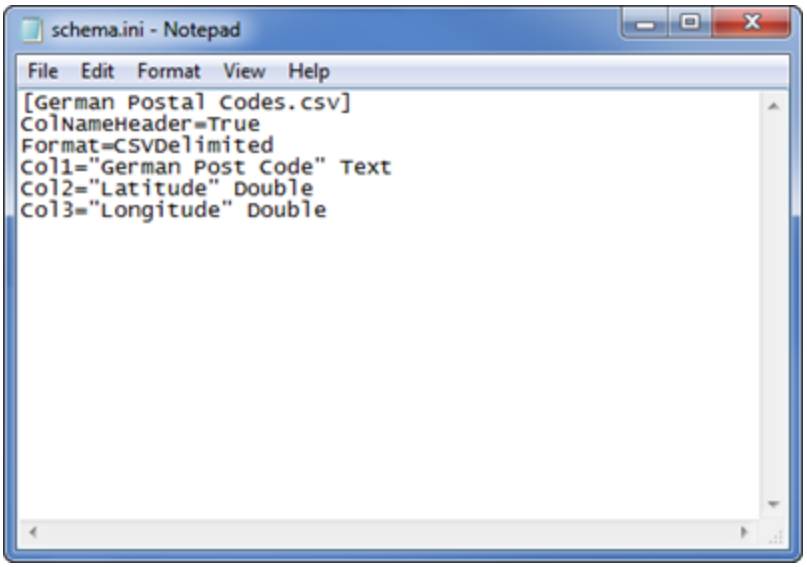

For more information on how to create a schema.ini file, see Importing Numeric Geographic Roles in the Tableau Knowledge Base.

4. Save the file as schema.ini in the same folder as the .csv file.
5. Import the custom geocoding files into Tableau. For more information, see Import Custom Geocoding.
6. Assign the new geographic role to a field. For more information, see Geographic Roles.
7. Plot the postcodes on a map.

   After assigning a new geographic role to a field, drag the field to a map view to plot the data. For more information on how to create map views, see Building a Map View.

   The German postcodes will be placed in your map view. Each mark in the view represents a postcode, whose location was specified in the custom geocoding data you imported.

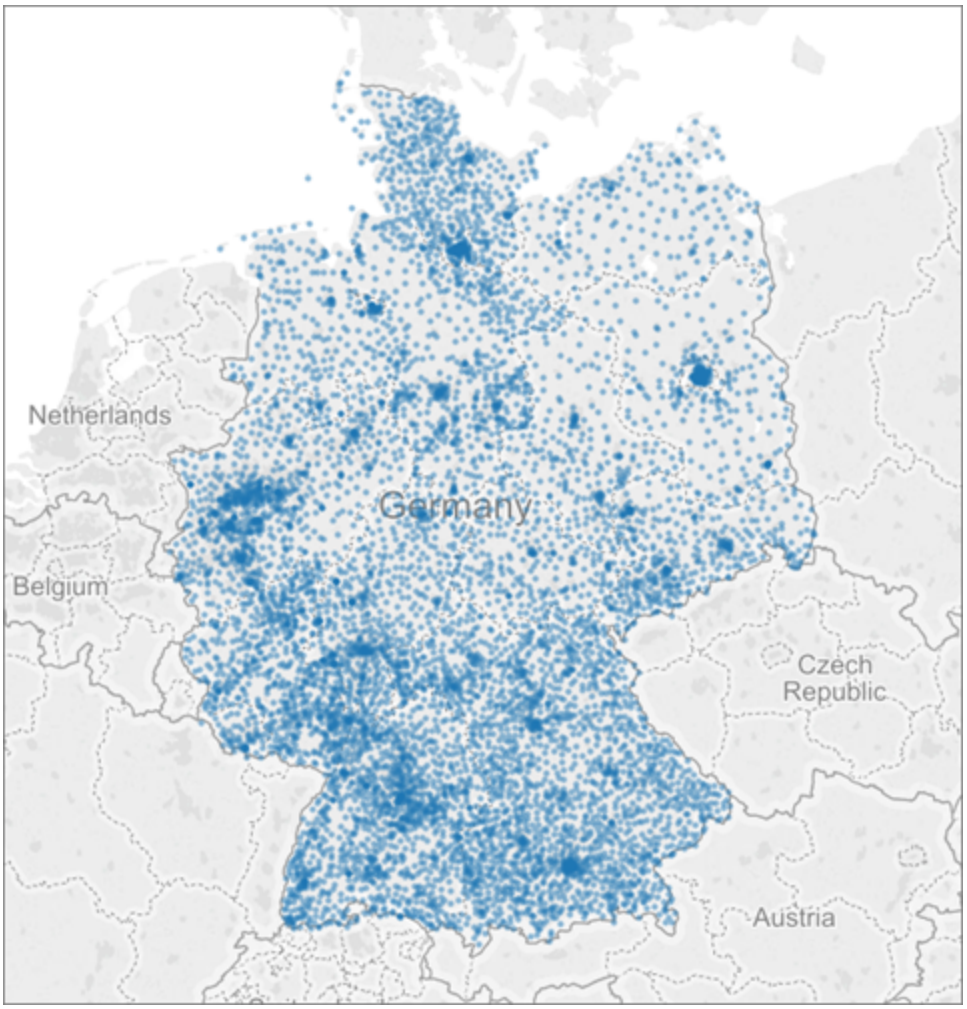

Background Map Sources

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 **Map > 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. For more background information, refer to the Mapping Data with WMS knowledge base article.

To add a WMS server:

1. Select **Map** > **Background Maps** > **WMS Servers**.
2. Click **Add**.
3. In the subsequent dialog box, type the URL for the server and click **OK**.

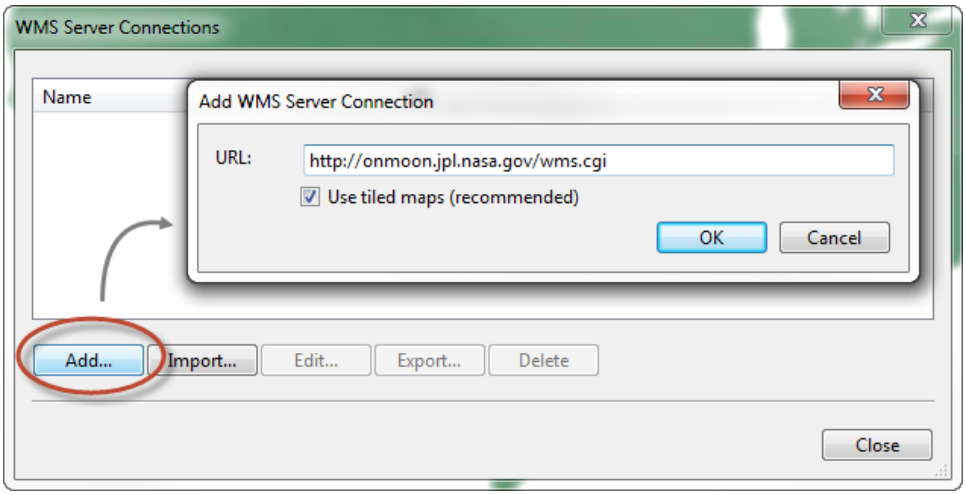

You can add as many map servers as you need. Each map source is displayed in the list of map sources on the Data menu. The map sources will be saved with the workbook and available to anyone you share the workbook with. You can also export the map source to a file that you can share with others so they can quickly connect to it.

To export a WMS server:

1. Select **Map** > **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**.

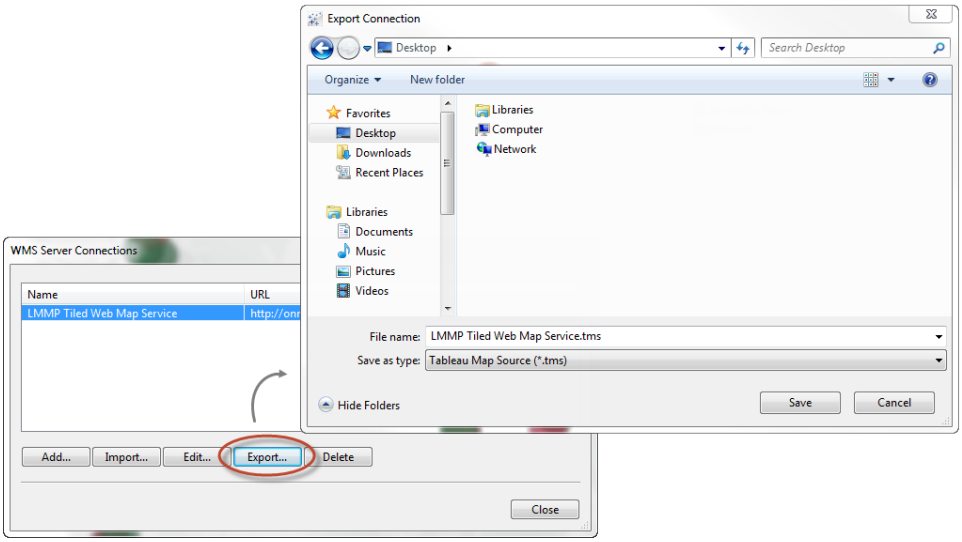

The Tableau Map source includes any default map options settings you have specified. For example, the map source will include any custom set of map layers you have specified to show by default. If you change the map options defaults you should export the map source again to include the new settings in the map source file.

To import a map source:

1. Select **Map** > **Background Maps** > **WMS Servers**.
2. Click **Import**.
3. Navigate to the saved map source file (.tms), select it, and click **Open**.

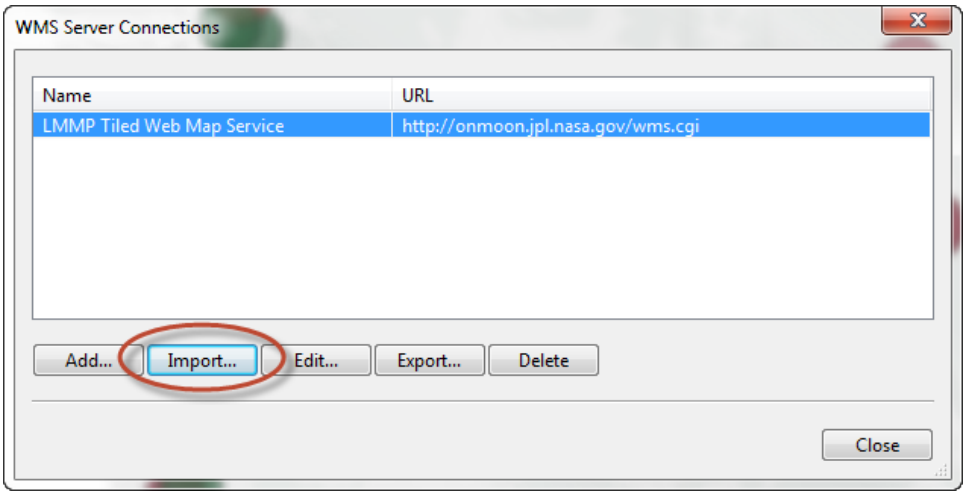

Setting a Default Map Source

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 a background map style on the **Map** > **Background Maps** menu.
2. Select **Map** > **Background Maps** > **Set as Default** to set the selected background map style to the default.

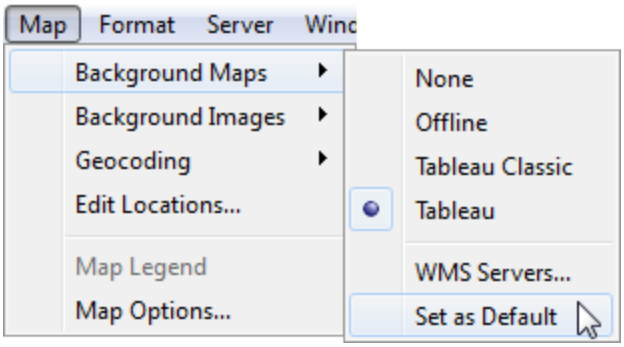

Tableau has four background map style options.

None	When selected, data is displayed between latitude and longitude axes.
Offline	Stores the images that make up the map in a cache with your Internet Explorer Temporary Internet Files for improved performance and offline access. See **Map Storing and Working Offline** on the next page for details.
Tableau Classic	When selected, this switches the background map to the Tableau classic style.
Tableau	The default background map unless selected otherwise, starting with Tableau 8.2. The fonts, details, and colors in this style differ from the Tableau classic background.

Tableau Classic **Tableau**

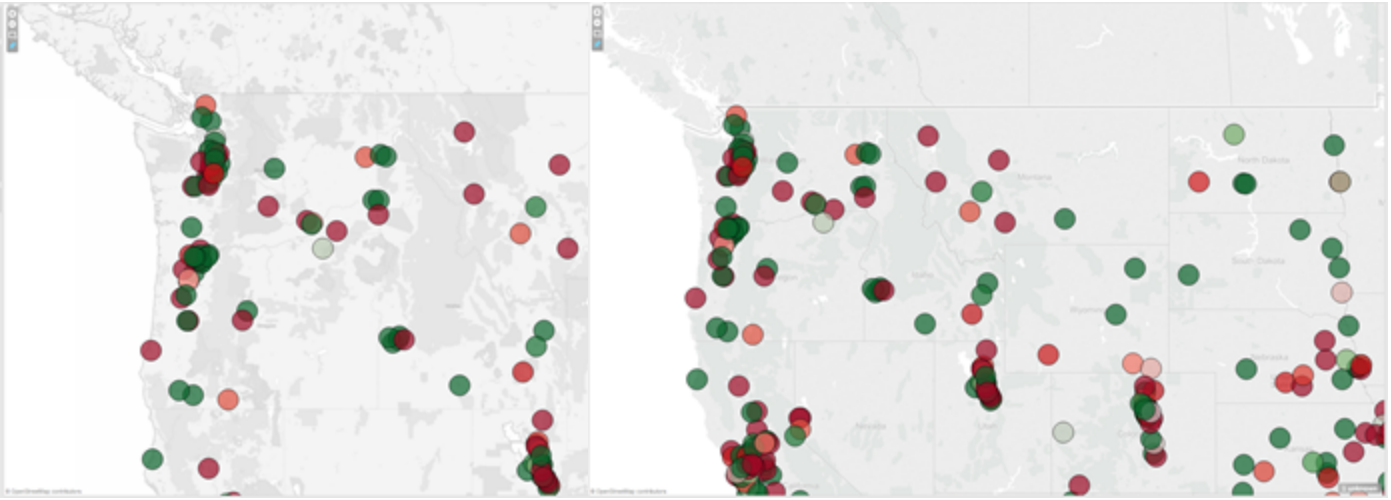

Map Storing and Working Offline

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 and Statistics

Use Tableau's trend line feature to incrementally construct interactive models of behavior. With trend lines you can answer such questions as whether profit is predicted by sales, or whether average delays at an airport are significantly correlated with the month of the year.

Adding Trend Lines

When you add trend lines to the view, you can specify how you want them to look and behave.

Add Trend Lines to the View

1. Select **Analysis > Trend Lines > Show Trend Lines**, or right-click on the pane and choose **Trend Lines > Show Trend Lines**.

   This command adds a linear trend line for each page, pane, and color on the worksheet. You can continue with the steps below to configure trend lines.

2. Select **Analysis > Trend Lines > Edit Trend Lines**, or right-click on the pane and choose **Trend Lines > Edit Trend Lines** to open the Trend Lines Options dialog box.

   a. Select either a Linear, Logarithmic, Exponential, or Polynomial model type.

   b. Select the fields you want to include as factors in the trend line model. For example, on a view of sales across categories and regions, you may want to see the overall sales trend across all categories, rather than a different line for each region. In this case, you would exclude the **Region** field as a factor:

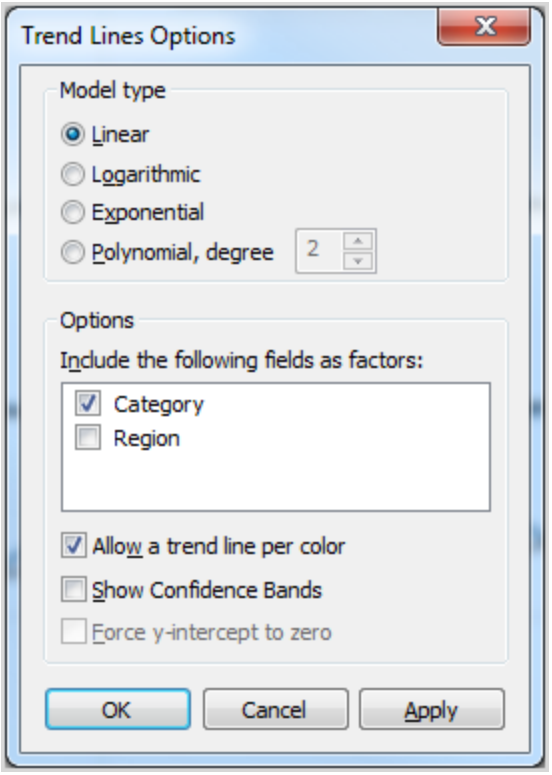

For more information on excluding factors, see **Removing Factors from the Model** on page 721.

c. Decide whether to exclude color, using the **Allow a trend line per color** option. 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 color encoding.

d. Decide whether to **Show Confidence Bands**. Tableau confidence bands show upper and lower 95% confidence lines. Confidence lines are not supported for Exponential models.

e. Select whether to **Force the y-intercept to zero**. This option is useful when you know that you want your trend line to begin at zero. This option is available only when both the **Rows** shelf and the **Columns** shelf contain a continuous field, as with a scatterplot.

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

When Can't You Add Trend Lines?

To add trend lines to a view, both axes must contain a field that can be interpreted as a number. For example, you cannot add a trend line to a view that has the **Product Category** dimension, which contains strings, on the Columns shelf and the **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.

For multidimensional data sources, the date hierarchies actually contain strings rather than numbers. Therefore, trend lines are not 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 do not show. When you change the view back to a state that allows trend lines, they reappear.

> 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 clearing the **Analysis** > **Stacked Marks** option.

Removing Trend Lines

To remove trend lines from the view, choose **Analysis > Trend Lines > Show Trend Lines** to remove the check mark, or right-click the pane and choose **Trend Lines > Show Trend Lines** to remove the check mark. Trend line options are retained for the next time you turn on trend lines. However, if you close the workbook with trend lines turned off, trend line options revert to defaults.

Trend Line Model Types

When you add a trend line to a view, you are building a statistical model. You are answering a question: Do the factors in your view predict a specific value (measure)? More precisely, do the categorical factors in your view split your data into meaningful samples and does the data in those samples allow you to predict values of the response measure given values of the explanatory measure or date? For example, you might be asking whether profit is predicted by time in a view that shows the profit of a company over four years.

Each trend line in the view visually represents a linear regression statistical model. Each model is estimated using data in the same pane and of the same color as the corresponding trend line. Although trend lines may be of type linear, logarithmic, exponential, or polynomial, this does *not* indicate that any of these models is not a linear regression. These types just identify transformations of either the explanatory or response variable in the linear regression formula. The word linear means that the formula is always of the form

```
y = b0 + b1 * x1 + ... + bn * xn
```

In other words, the formula is always linear in its coefficients: `b0 ... bn`.

Model Types

Choose from the following model types. In the following formulas, X represents the explanatory variable, Y to response variable, and (ε) (epsilon) represents random error. Random errors are uncorrelated with each other and with the explanatory variable, and have equal variance.

Linear

With the linear model type, no transformations are performed on either the explanatory or response variable. So the formula is

```
Y = b0 + b1 * X + ε
```

Logarithmic

With the Logarithmic model type, the explanatory variable is transformed by the natural log before estimation of the model. So the formula is

```
Y = b0 + b1 * ln(X) + ε
```

Because a logarithm is not defined for number less than zero, any marks for which the explanatory variable is negative are filtered before estimation of the model. Avoid using a model that discards some data unless you know that the data being filtered out is invalid. The trend line description reports how many marks were filtered before model estimation.

Exponential

With the exponential model type, the response variable is transformed by the natural log before estimation of the model. So the formula is

```
ln(Y) = b0 + b1 * X + ε
```

With an exponential model, your response axis does *not* become logarithmic. Instead, the marks plotted in your view are found by plugging in various explanatory values to find values of `ln(Y)`. These values are then exponentiated to plot the trend line. What you see is the exponential model:

```
Y = e ^ (b0 + b1 * X + ε)
```

Because a logarithm is not defined for numbers less than zero, any marks for which the response variable is negative are filtered before model estimation.

Polynomial

With the polynomial model type, the response variable is transformed into a polynomial series of the specified degree. So the formula is

```
Y = b0 + b1 * X + b2 * X^2 + … + ε
```

With a polynomial model type, you also select a **Degree** between 2 and 8. The higher polynomial degrees exaggerate the differences between the values of your data. So if your data increases very rapidly, the lower order terms may have almost no variation compared to the higher order terms, rendering the model impossible to estimate accurately. Also, more complicated higher order polynomial models require more data to estimate. So check the model description of the individual trends line for a red warning message indicating that an accurate model of this type is not possible.

Removing Factors from the Model

You can remove factors from a trend line model in 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 various products categories, broken out by region.

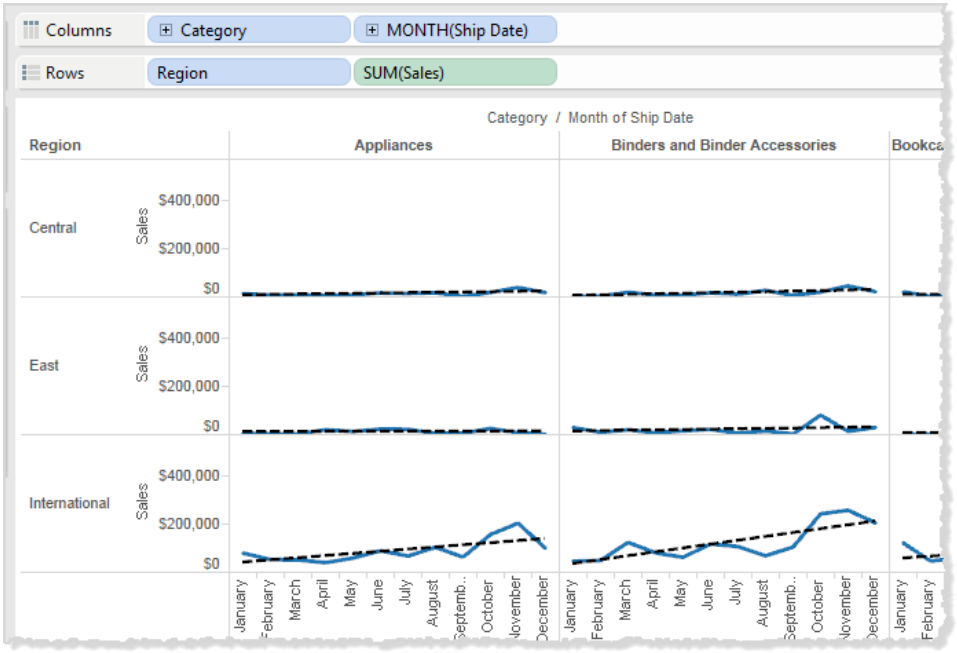

You can see that a separate model is created for each region.

Now remove **Region** as a factor in the model by deselecting it in the Trend Lines Options dialog box.

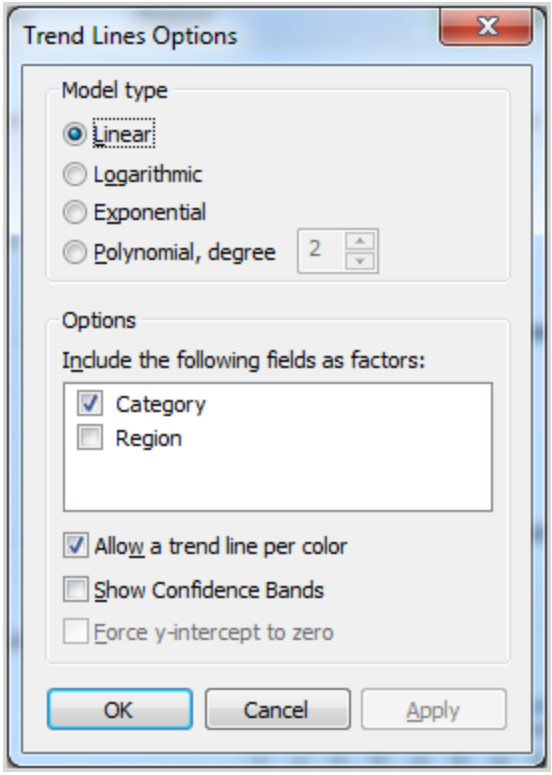

You can see that the trend line model within a category is now the same across all regions. This allows you to compare actual sales against a trend line that is the same for all regions.

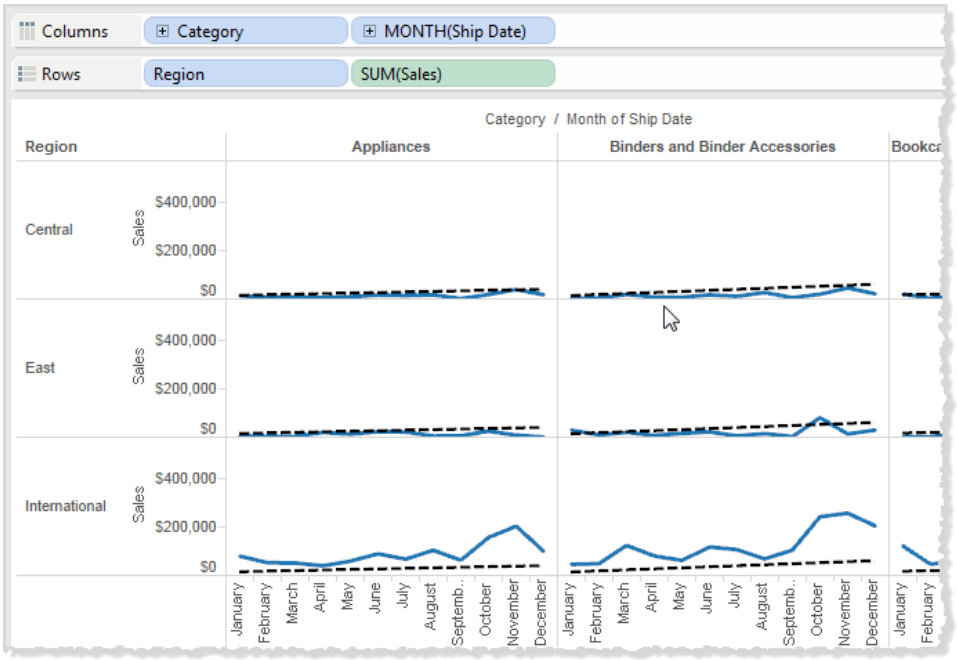

Assessing Trend Line Significance

To see relevant information for any trend line in the view, hover the cursor over it to see information about the trend line:

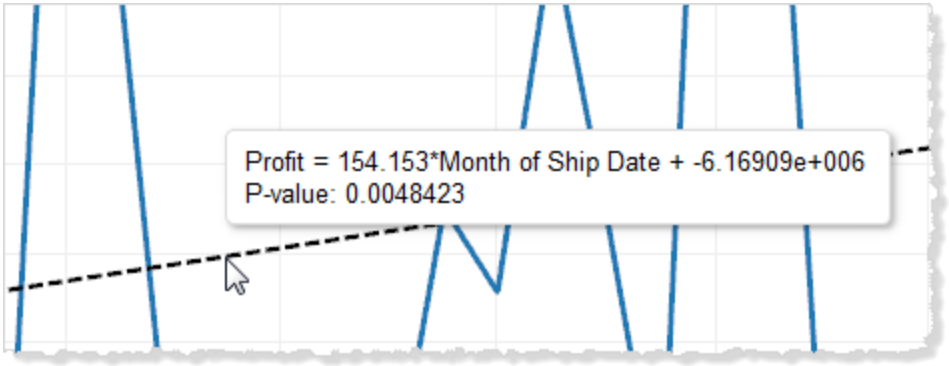

The first line in the tooltip shows the equation used to compute a value of Profit from a value of Month of Ship Date.The second line reports the probability that the equation in the first line was a result of random chance. The smaller the p-value, the more significant the model is. A p-value of 0.05 or less is often considered sufficient.

Entire Model Significance

Once you've added a trend line to a view, you typically want to know the goodness of fit of the model, which is a measure of the quality of the model's predictions. In addition, you may be interested in the significance of each factor contributing to the model. To view these numbers, open the Describe Trend Model dialog box by right-clicking in the view and selecting **Trend Lines > Describe Trend Model**.

When you are testing significance, you are 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 which contains an individual trend line or a term of an individual trend line that does not contribute to the overall significance.

Under Trend Lines Model, look for the line that shows the p-value (significance) of the model: The smaller the p-value, the less likely it is that the difference in the unexplained variance between models with and without the relevant measure or measures was a result of random chance.

Describe Trend Model

Trend Lines Model

A linear trend model is computed for sum of Sales given Ship Date Month. The model may be significant at p <= 0.05. The factor Category may be significant at p=0.05.

Model formula:	Category*(Month of Ship Date + intercept)
Number of modeled observations:	204
Number of filtered observations:	0
Model degrees of freedom:	34
Residual degrees of freedom (DF):	170
SSE (sum squared error):	6.32246e+011
MSE (mean squared error):	3.71909e+009
R-Squared:	0.850309
Standard error:	60984.4
p-value (significance):	< 0.0001

Analysis of Variance:

Field	DF	SSE	MSE	F	p-value
Category	32	3.08378e+012	9.6368e+010	25.9117	< 0.0001

Individual trend lines:

Panes		Line		Coefficients				
Row	Column	p-value	DF	Term	Value	StdErr	t-value	p-value
Sales	Appliances	0.0221712	10	Month of Ship Date	11355.9	4200.11	2.70372	0.0221712
				intercept	52856.9	30912	1.70992	0.118073
Sales	Binders and Binder Accessories	0.0048827	10	Month of Ship Date	21028.2	5848.16	3.5957	0.0048827
				intercept	38108.4	43041.3	0.885391	0.396736
Sales	Bookcases	0.0582459	10	Month of Ship Date	12823	5998.1	2.13785	0.0582459

This p-value for a model compares 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, it assesses the explanatory power of the quantitative term f(x) in the model formula, which can be linear, polynomial, exponential, or logarithmic with the factors fixed. It is common to assess significance using the "95% confidence" rule. Thus, as noted above, a p-value of 0.05 or less is considered good.

Significance of Categorical Factors

In the Analysis of Variance table, sometimes referred to as the ANOVA table, each field that is used as a factor in the model is listed. For each field, among other values, you can see the p-value. In this case, the p-value indicates how much that field adds to the significance of the entire model. The smaller the p-value the less likely it is that the difference in the unexplained variance between models with and without the field was a result of random chance. 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 Analysis of Variance 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-value
Category	160	1.65811e+012	1.03632e+010	20.9732	< 0.0001
Region	136	2.64345e+012	1.94371e+010	39.3372	< 0.0001

As you can see, the p-values for **Category** and **Region** are both quite small. Both of these factors are statistically significant in this model.

For information on specific trend line terms, see **Trend Line Model Terms** on page 728.

For ANOVA models, trend lines are defined by the mathematical formula:

```
Y = factor 1 * factor 2 * ...factorN * f(x) + ε
```

The term `Y` is called the *response variable* and corresponds to the value you are trying to predict. The term `X` is the *explanatory variable*, and `ε` (epsilon) is random error. The factors in the expression correspond to the categorical fields in the view. 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 are introduced. For example, if factor 1 and factor 2 both have three members, then a total of nine variables are introduced into the model formula by this operator.

Answering Questions with Trend Lines

We start with a question: what is causing high discount rates at a superstore? Using trend lines, we can find out which variables correlate with high discounts. The view below shows the average monthly discount rates for all stores over a four-year range.

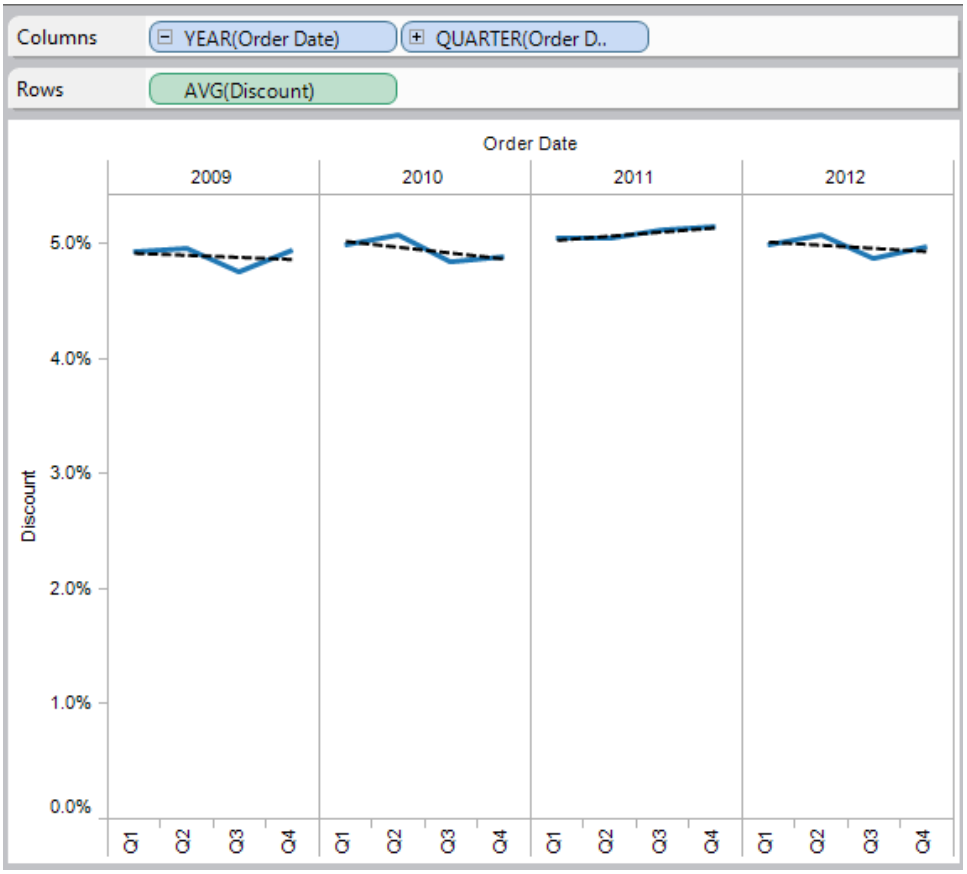

Our first thought is that discount rates may be high for specific ship modes. When we put the **Ship Mode** field on the **Rows** shelf, we see that there was a spike in discount rates for items shipped by Delivery Truck in 2009.

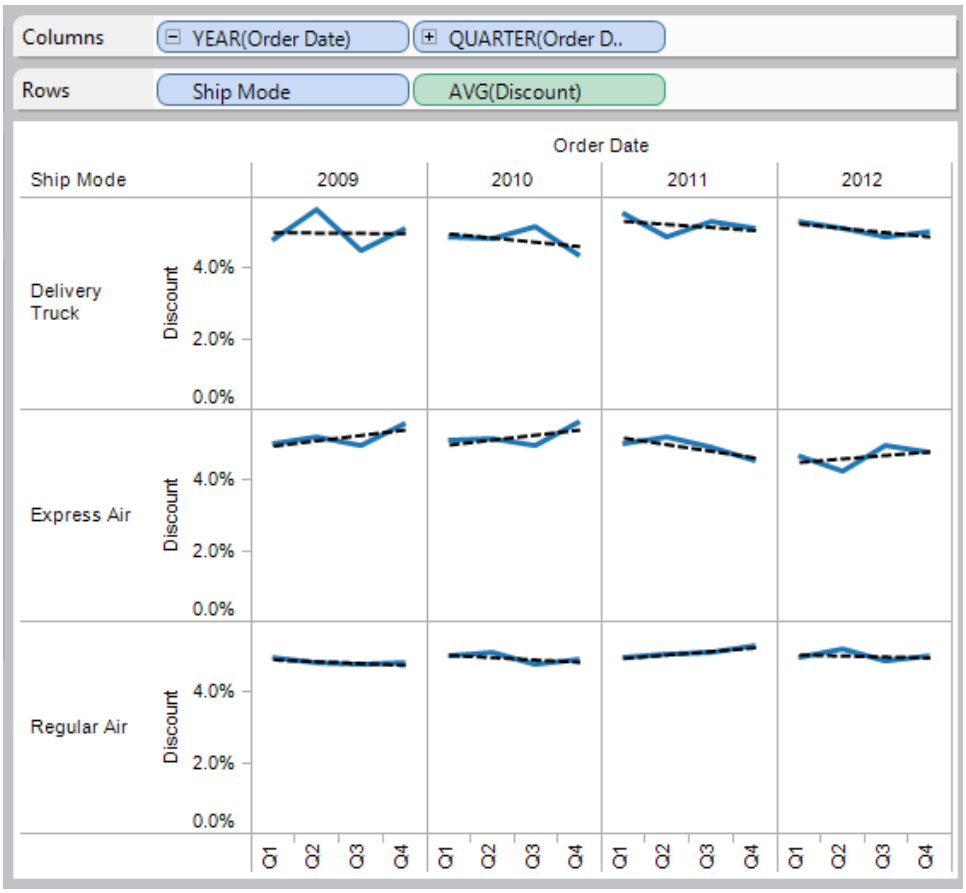

However, in the description of the trend line mode, the Analysis of Variance list shows that the p-value for Ship Mode is 0.22. (To see the description, right-click the trend line and choose **Describe Trend Model**). This number is too high to be significant, indicating that we cannot predict discount rates based on ship mode. (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-value
Ship Mode	16	0.000181235	1.133e-005	1.39986	0.222194
Year of Order Date	18	0.00018312	1.017e-005	1.25726	0.295853

The next variable we test is supplier. Suppliers sometimes offer substantial discounts. So we remove **Ship Mode** and drag **Supplier** to the **Rows** shelf.

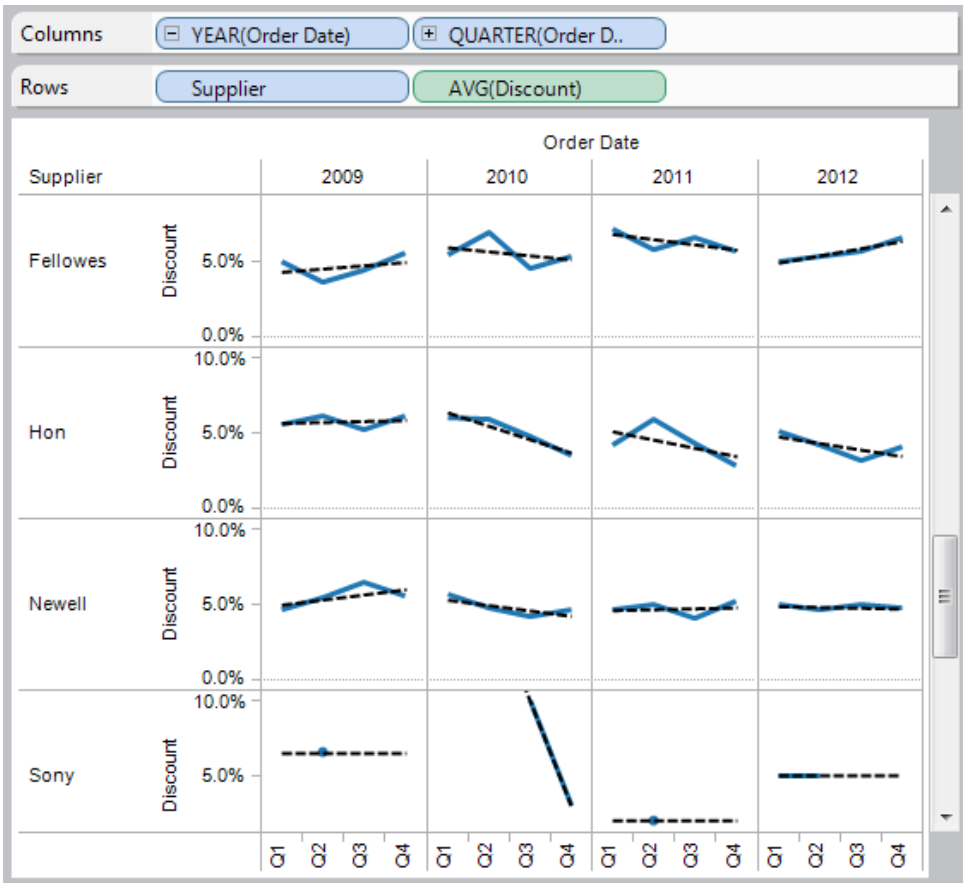

Now when we open the Describe Trend Model dialog box, we see that the p-value for Supplier is 0.0004, while the p-value for the entire model is less than 0.003. With that, we can be confident in concluding that changes in average discount rates over time are in fact quite different for different suppliers. In other words we can be statistically confident that the different trend lines slopes for each supplier isn't just due to randomness, but rather to a real correlation between a supplier and the slope of this line.

Analysis of Variance:

Field	DF	SSE	MSE	F	p-value
Supplier	126	0.0349464	0.0002774	1.60454	0.0045179
Year of Order Date	100	0.0318394	0.0003184	1.84197	0.000667

In addition to **Supplier**, we notice that the **Year of Order Date** field offers a statistically significant improvement to the model. See Assessing Trend Line Significance for more information.

Assumptions

The p-values reported in Tableau trend lines depend on some assumptions about the data.

The first assumption is that, whenever a test is performed, the model for the mean is (at least approximately) correct.

The second assumption is that the "random errors" referred to in the model formula (see **Trend Line Model Types** on page 719) are independent across different observations, and that they all have the same distribution. This constraint would be violated if the response variable had much more variability around the true trend line in one category than in another.

Assumptions Required to Compute Trend Lines

The Assumptions required to compute (using Ordinary Least Squares) each individual trend line are:

- Your model is an accurate functional simplification of the true data generating process (for example, no linear model for a log linear relationship).
- Your errors average to zero and are uncorrelated with your independent variable (for example, no error measuring the independent variable).
- Your errors have constant variance and are not correlated with each other (for example, no increase in error spread as your independent variable increases).
- Explanatory variables are not exact linear functions of each other (perfect multicollinearity).

Trend Line Model Terms

When you view the description for a 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 modeled observations

The number of rows used in the view.

Number of filtered observations

The number of observations excluded from the model.

Model degrees of freedom

The number of parameters needed to completely specify the model. Linear, logarithmic, and exponential trends have model degrees of freedom of 2. Polynomial trends have model degrees of freedom of 1 plus the degree of the polynomial. For example a cubic trend has model degrees of freedom of 4, since we need parameters for the cubed, squared, linear and constant terms.

Residual degrees of freedom (DF)

For a fixed model, this value is defined as the number of observations minus the number of parameters estimated in the model.

SSE (sum squared error)

The errors are the difference between the observed value and the value predicted by the model. In the Analysis of Variance 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

The ratio of variance in the data explained by the model to total variance in the data.

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

$$R^2 = \frac{(SSE(S) - SSE(F))}{SSE(S)}$$

measures the relative drop in variability of the response when using the full model as opposed to the smaller model. In the Describe Trend 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. An estimate of the standard deviation (variability) of the "random errors" in the model formula.

p-value (significance)

The probability that an F random variable with the above degrees of freedom exceeds the observed F in this row of the Analysis of Variance 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. Looking at the list you can see which, if any, are the most statistically significant. This table also lists coefficient statistics for each trend line. A row describes each coefficient in each trend line model. For example, a linear model with an intercept requires two rows for each trend line. In the Line column, the p-value and the DF for each line span all the coefficient rows. The DF column under the shows the residual degrees of freedom available during the estimation of each line.

Terms

The name of the independent term.

Value

The estimated value of the coefficient for the independent term.

StdErr

A measure of the spread of the sampling distribution of the coefficient estimate. This error shrinks as the quality and quantity of the information used in the estimate grows.

t-value

The statistic used to test the null hypothesis that the true value of the coefficient is zero.

p-value

The probability of observing a t-value that large or larger in magnitude if the true value of the coefficient is zero. So, a p-value of .05 gives us 95% confidence that the true value is not zero.

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. The p-value will measure the probability of obtaining the same trend result without taking the dimensions into account. For example, a trend of sales per time p-value of 0.05 means that there is 5% chance that the same value could be obtained without taking the time into consideration.

What does it mean if the p-value for the model is significant but the p-value for the specific field in the Analysis of Variance table is not significant?

The p-value in the Analysis of Variance table indicates whether the field adds or detracts from the significance of the entire model. The smaller the p-value the less likely it is that the difference in the unexplained variance between models with and without the field was a result of random chance. 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, for the situation where the p-value for the model is significant but the p-value for the specific field is not, you know that the model is statistically significant, but you cannot be confident that the specific field in question adds anything to it. Consider whether you might not be better off removing the factor from the model.

What does it mean if the p-value for the specific field in the Analysis of Variance table is significant but the p-value for the model is not 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

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**.

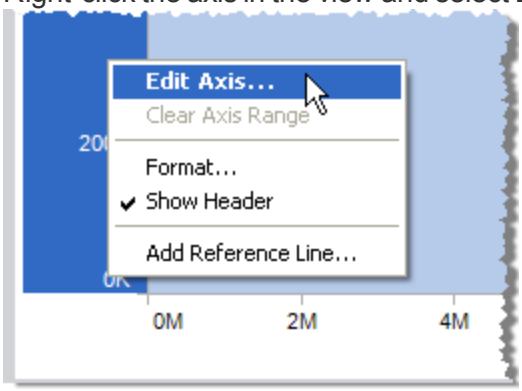

2. In the Edit Axis dialog box, select **Logarithmic** for the axis scale.

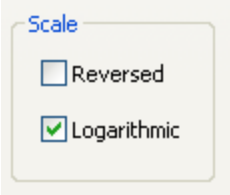

3. Select the **Tick Marks** tab.

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.

> 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.

Publish and Share

Close the loop in your analysis and share your findings with others. Start with rich formatting and finish with publishing the view, dashboard, or story.

Formatting

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.

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.

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.

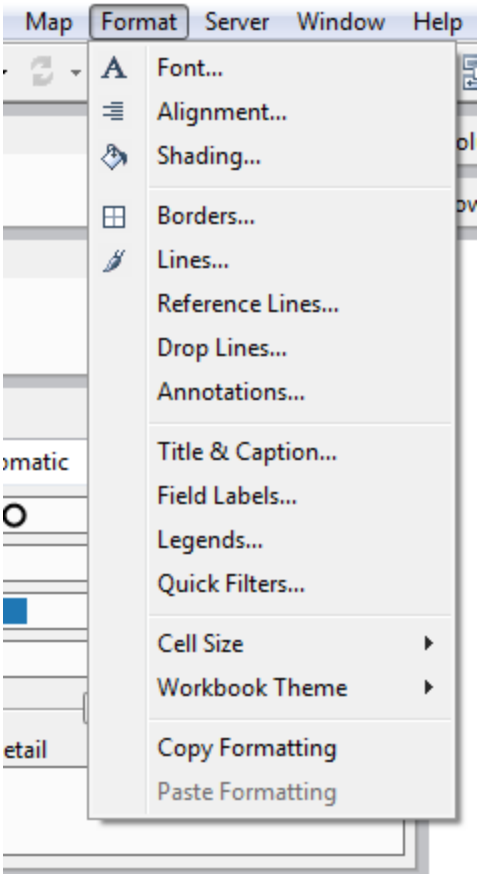

The Format window opens on the left side of the workbook, replacing the Data window. At the top of the Format window, there is a toolbar where you can quickly switch between each of the types of format settings available.

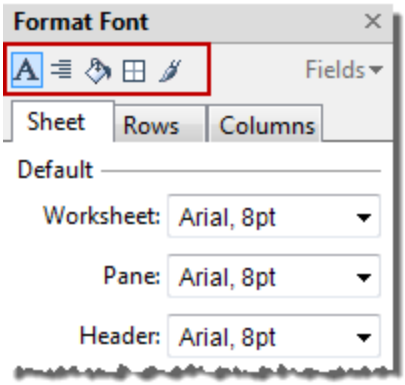

The Format window also contains three tabs: Sheet, Rows, and Columns. Switch between these tabs to apply formatting to the entire sheet, just the rows, or just the columns.

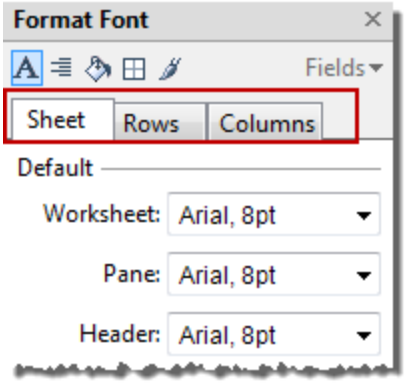

Fonts

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.

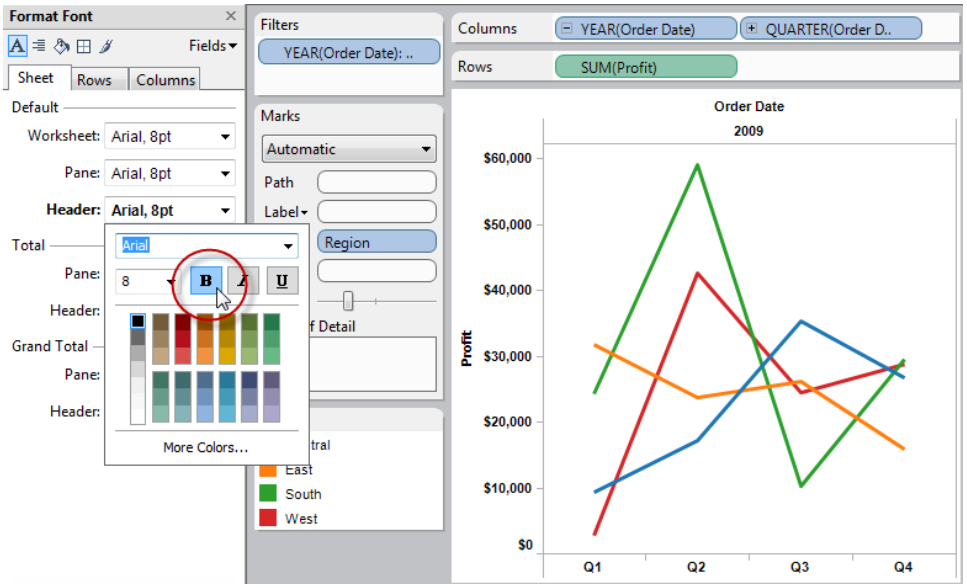

If you have totals or grand totals in the view, you can specify special font settings so that these values can stand out from the rest of the data. This is particularly useful when you are working with a text table. The view below shows a text table in which the grand totals are formatted to be bold and dark red.

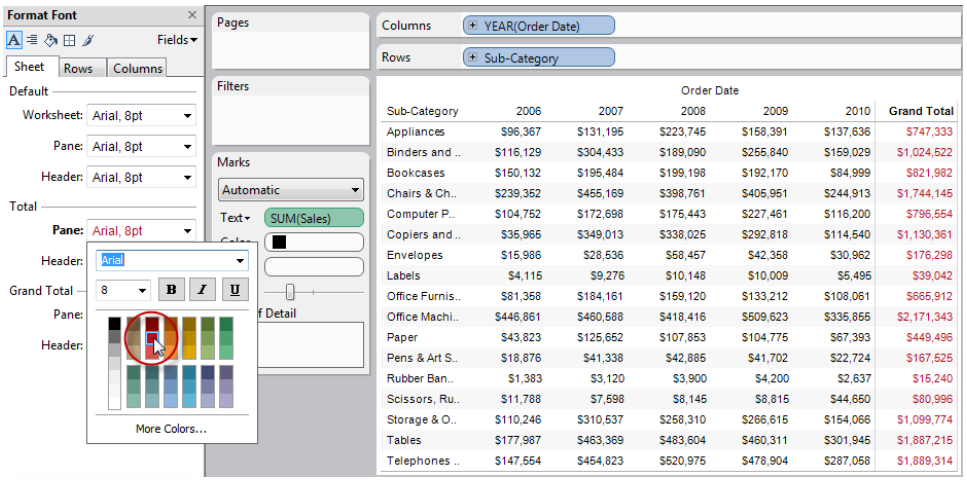

Finally, you can use the Worksheet drop-down to specify the properties of all text in the worksheet.

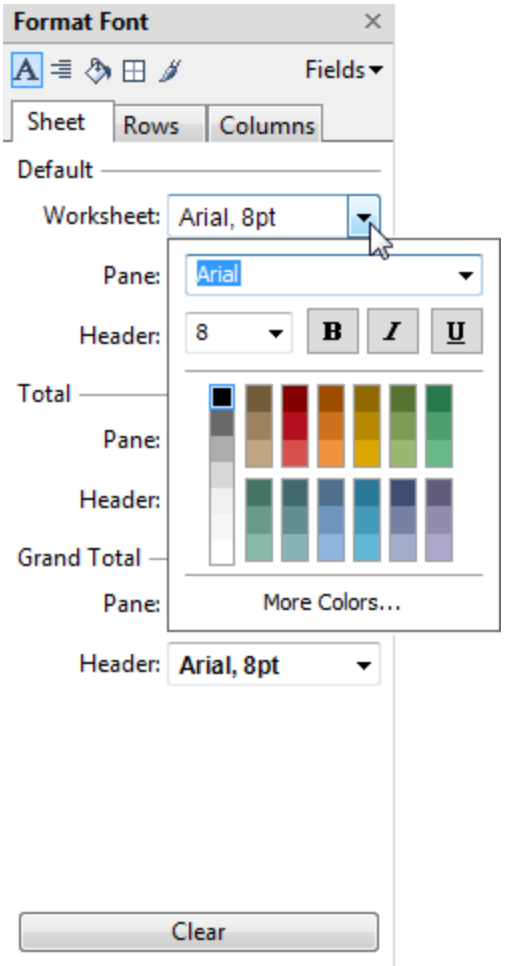

Row and Column Font Settings

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.

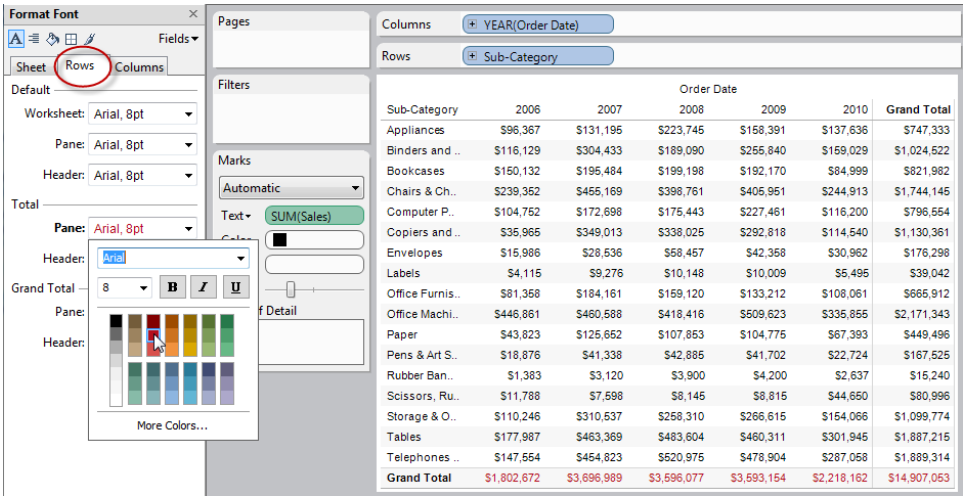

However, sometimes the settings don't make sense for particular views and are unavailable. For example, in the view above, the default pane text is in both rows and columns. Specifying font properties on a row or column level does not make sense, to change the pane text, switch back to the **Sheet** tab.

Alignment

From the **Format** menu, choose **Alignment** to see Format Alignment options. You can specify alignment properties for rows, columns, or the entire worksheet.

Use the drop-down menus on the **Format Alignment** pane to specify text alignments for either panes or headers. For each text areas you can specify the following alignment options:

- Horizontal - controls whether text aligns to the left side, right side, or in the center.
- Vertical Alignment - controls whether text aligns at the top, middle, or bottom.
- Direction - rotates the text so that it runs horizontally, top-to-bottom, or bottom-to-top.
- Wrap - controls whether long headers and pane text wrap to the next line or are abbreviated.

> **Note**: If cells are not large enough to show more than one row of text, turning on wrapping will have no visible effect. If this happens, you can hover the cursor over a cell until a double-sided arrow appears, and then click and drag down to expand the cell size.

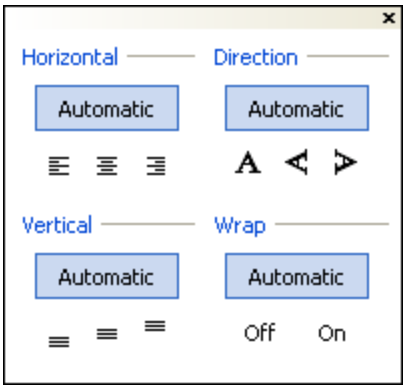

Shading

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 your eye to 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 is 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.
- 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.

  Slide the level selector right and left to specify the banding level.

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 serve 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.

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 are turned off.

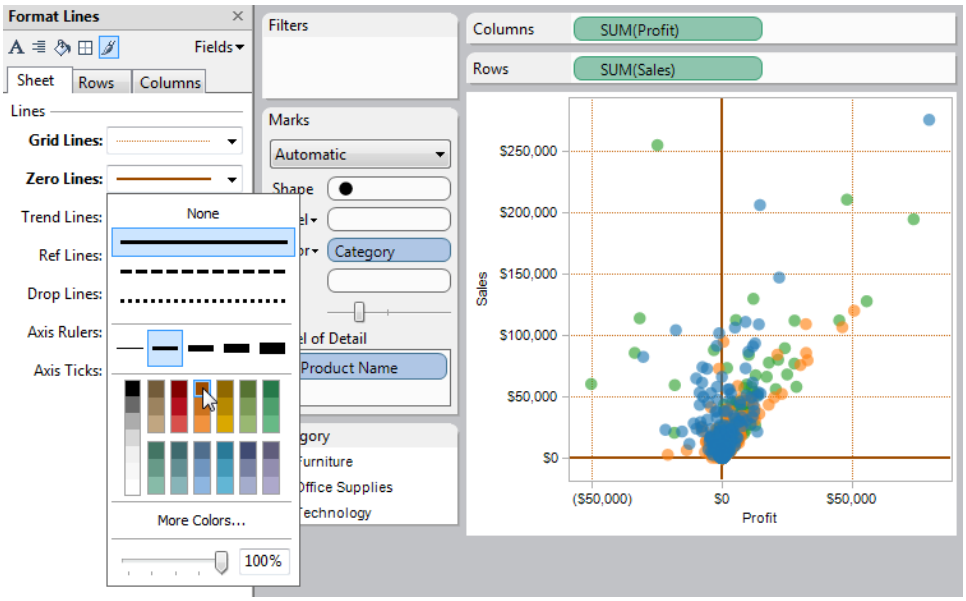

Formatting Specific Parts of the View

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. To specify individual format settings, right-click a specific part of the view and select **Format**.

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.

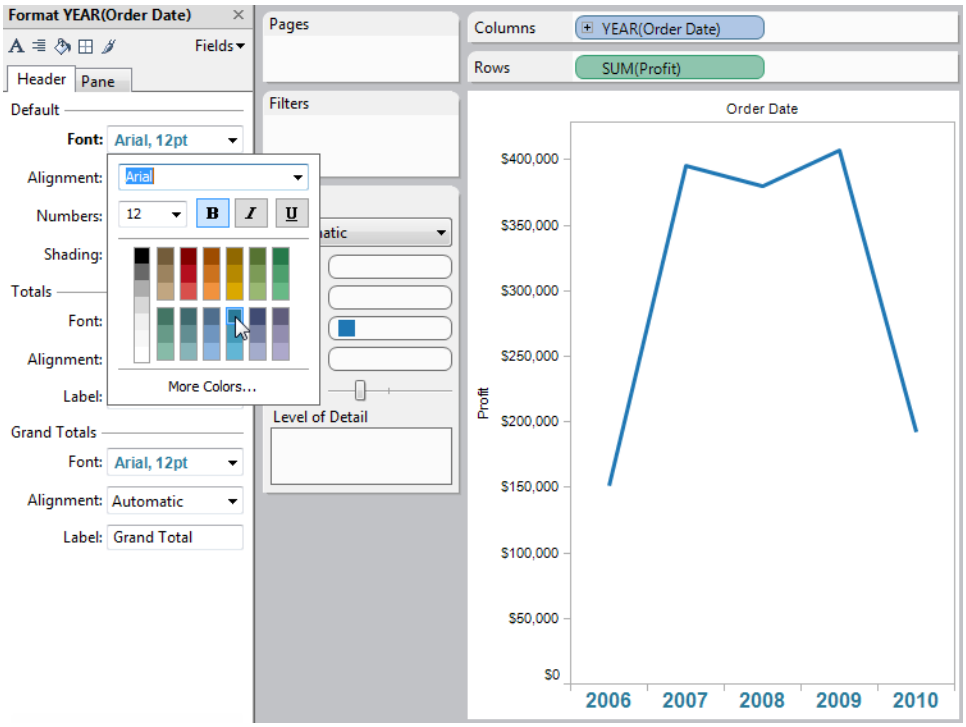

To format a specific field:

1. Right-click the part of the view you want to format and select **Format**.

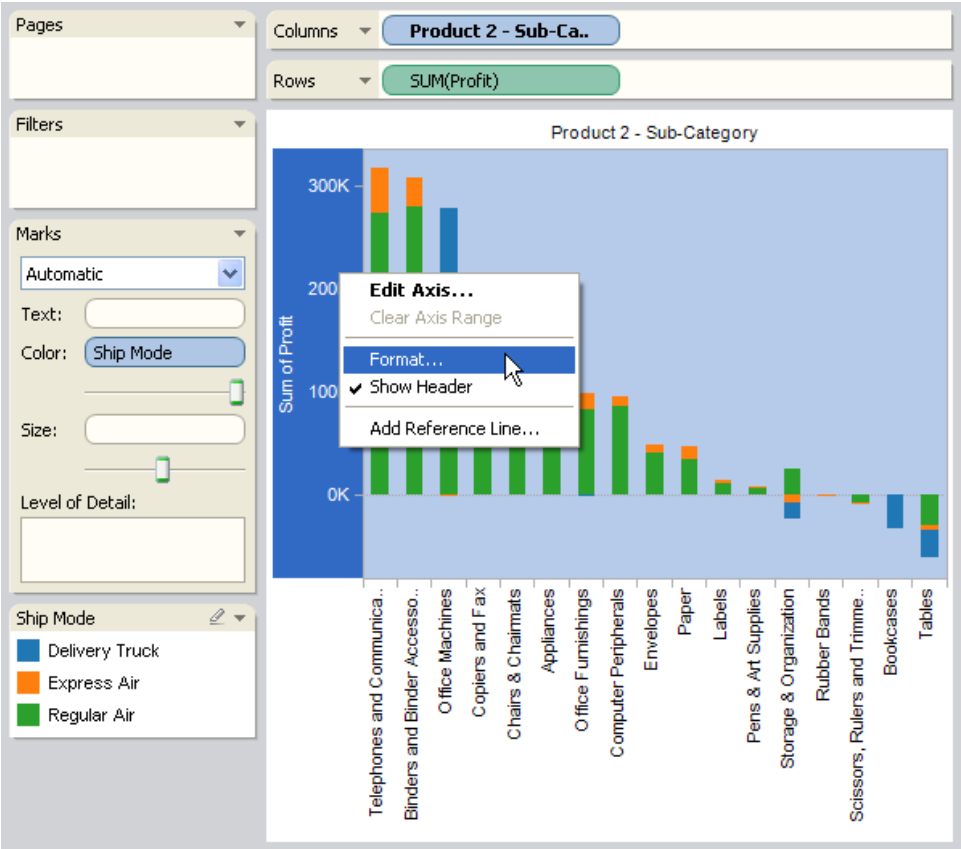

The Format window opens containing settings relevant to the selected field.

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 760 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.

   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.

3. 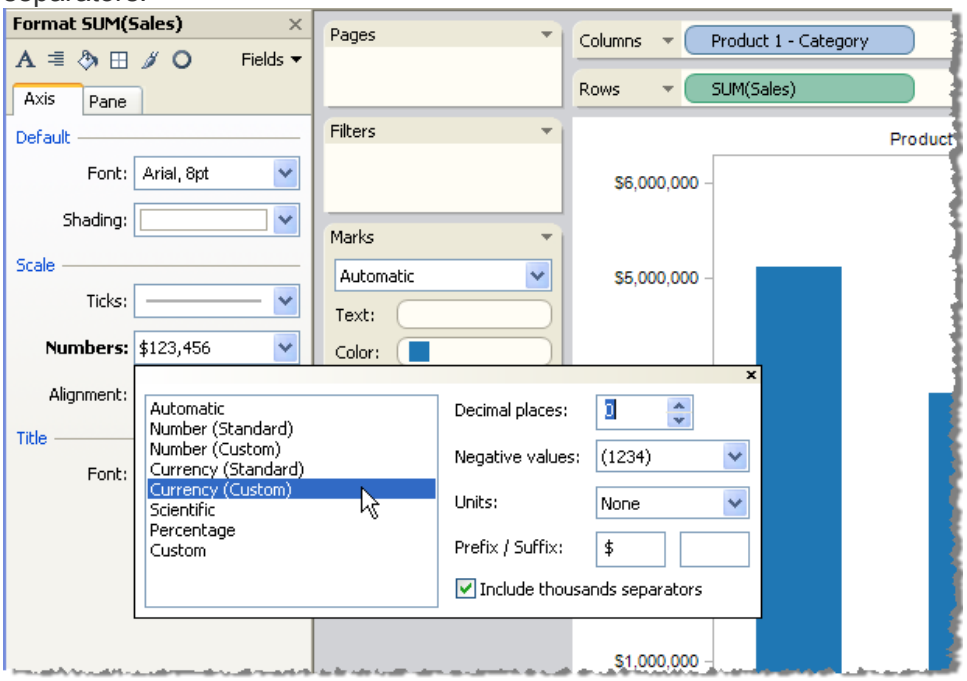

Below is a list of the number formats and associated options available in Tableau.

Number Format	Format Options
Automatic: format is automatically selected based on either the format specified by the data source or the data contained in the field.	None
Number (Standard): format is based on locale selected.	Locale: number format changes based on the geographical location selected.
Number (Custom): format is customized to your liking.	Decimal Places: the number of decimal places to display. Negative Values: how negative values will be displayed. Units: the number will be displayed

Number Format	Format Options
	using the specified units. For example, if the number is 20,000 and the units are thousands, the number will be displayed as 20K. Prefix/Suffix: characters that will precede and follow each displayed number. Include thousands separators: whether the number will show separators every thousand (example: 100,000 vs. 100000).
Currency (Standard): format and currency symbol is based on locale selected.	Locale: currency format based on the geographical location selected.
Currency (Custom): format and currency symbol is customized to your liking.	Decimal Places: the number of decimal places to display. Negative Values: how negative values will be displayed. Units: 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. Prefix/Suffix: characters that will precede and follow each displayed number. Include thousands separators: whether the number will show separators every thousand (example: 100,000 vs. 100000).
Scientific: numbers are displayed in scientific notation.	Decimal: the number of decimal places to display.
Percentage: numbers are displayed as a percentage with the percent symbol. The value of 1 is interpreted as 100% and 0 as 0%	Decimal: the number of decimal places to display.
Custom: format is based entirely on what is specified in the format options.	Custom: 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.

Note: Formatting numbers using the Format window overrides any number formatting that you might have applied elsewhere.

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.

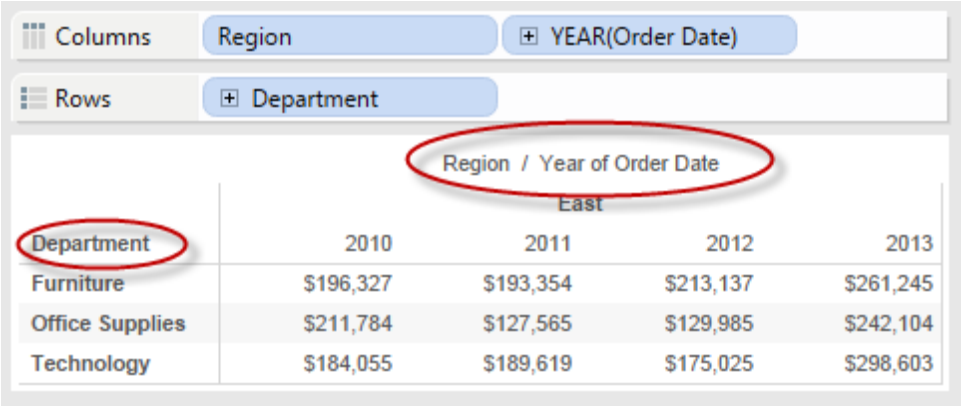

You can format the font, shading, alignment, and separators for each of these types of field labels.

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 the settings of the font, shading, and alignment field labels.

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 title and body of the legends.

By default, legend objects have transparent backgrounds. To change the opacity or legend color, follow the procedure listed above.

The legend format settings apply to all legends associated with a sheet. You cannot format individual legends separately.

To edit legend titles:

1. On the legend's card menu, select **Edit Title**.

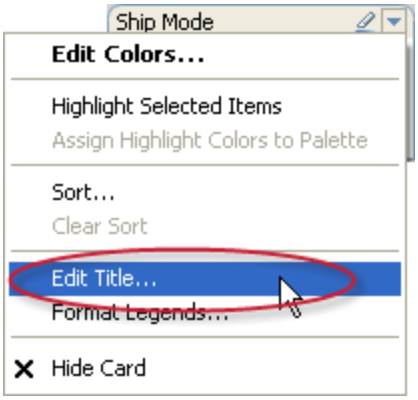

2. In the subsequent dialog box, type a new name for the legend and format it using the formatting options along the top of the dialog box. When finished, click **OK**.

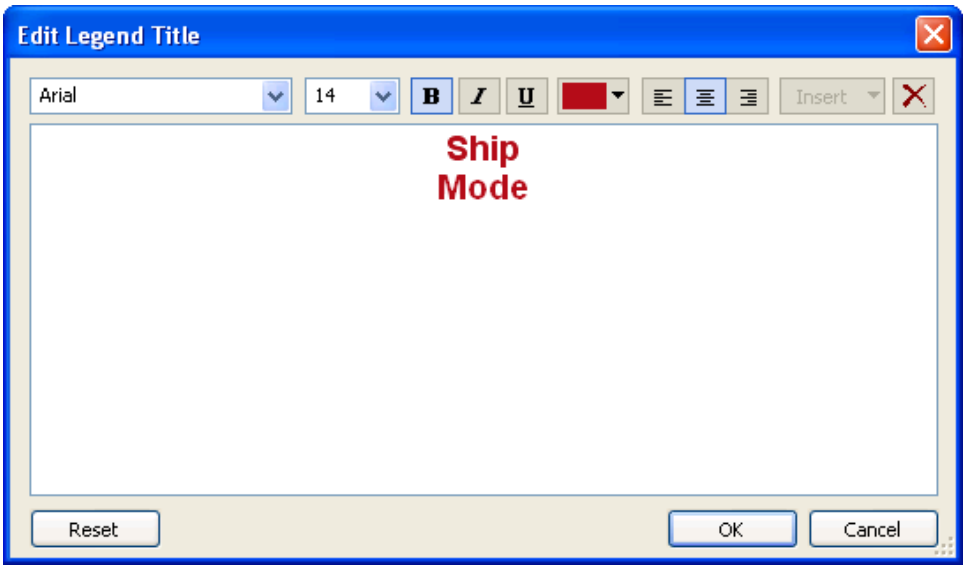

You can click the **Reset** button in the Edit Title dialog box to return to the default title.

Quick Filters

The quick filters on your worksheet can be formatted to use a different font, style, color, background color, font size, and border. Formatting quick filters allows you to better integrate the quick filters into your dashboard or worksheet style. The format settings you specify apply to all quick filters.

To format quick filters:

1. Select **Format** > **Quick Filters**.
2. In the Format window on the left side of the workbook, specify the following options:
    - **Title**: You can modify the **Font** and **Alignment** of the quick filter titles. The title formatting only displays when the quick filters are used on dashboards or when published to the web.
    - **Body**: You can modify the body **Font**, **Shading**, and **Border** of the quick filters. These settings display on both worksheets, dashboards, and when published to the web.

The quick filters in the dashboard below have custom formatting with orange titles and borders and a bold body font.

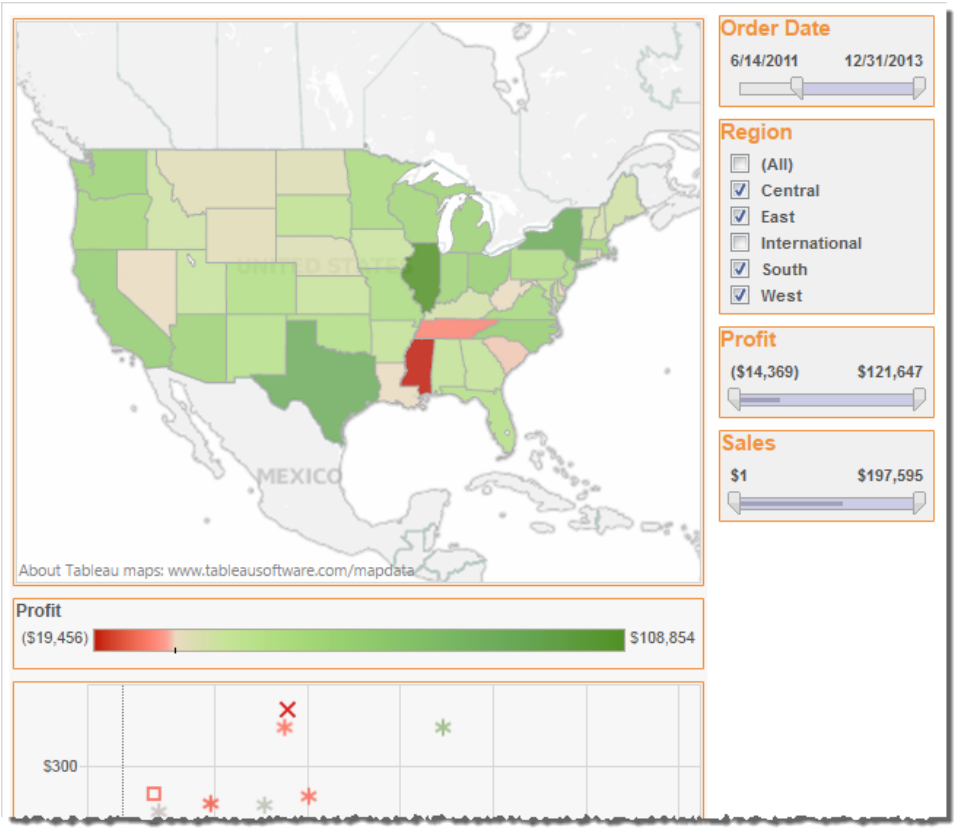

Parameters

The parameter controls on your worksheet can be formatted to use a different font, style, color, background color, font size, and border. Formatting parameters allows you to better integrate the parameter controls into your dashboard or worksheet style. The format settings you specify apply to all quick filters.

To format parameters:

1. Select **Format** > **Parameters**.
2. In the Format window on the left side of the workbook, specify the following options:
    - **Title**: You can modify the **Font** and **Alignment** of the parameters titles. The title formatting only displays when the parameter controls are shown on dashboards or when published to the web.
    - **Body**: You can modify the body **Font**, **Shading**, and **Border** around of the parameters. These settings display on both worksheets, dashboards, and when published to the web.

The Profit Bin Size parameter control in the dashboard below has custom formatting with a centered orange title, orange border, and a bold body font.

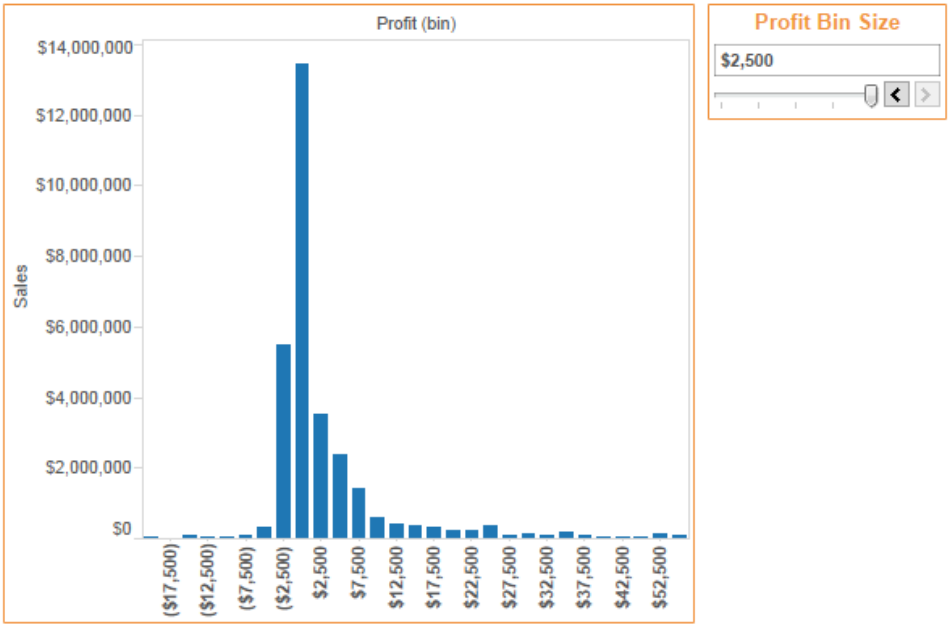

Title and Caption

You can add titles and captions to any sheet. After you add a title or caption you can edit and format the text as well as the shading and border.

If the title is not showing, to go the menu for the current sheet type (Worksheet, Dashboard, or Story) and choose **Show Title**.

To edit titles and captions:

1. Double-click the title.
2. In the Edit Title 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**.

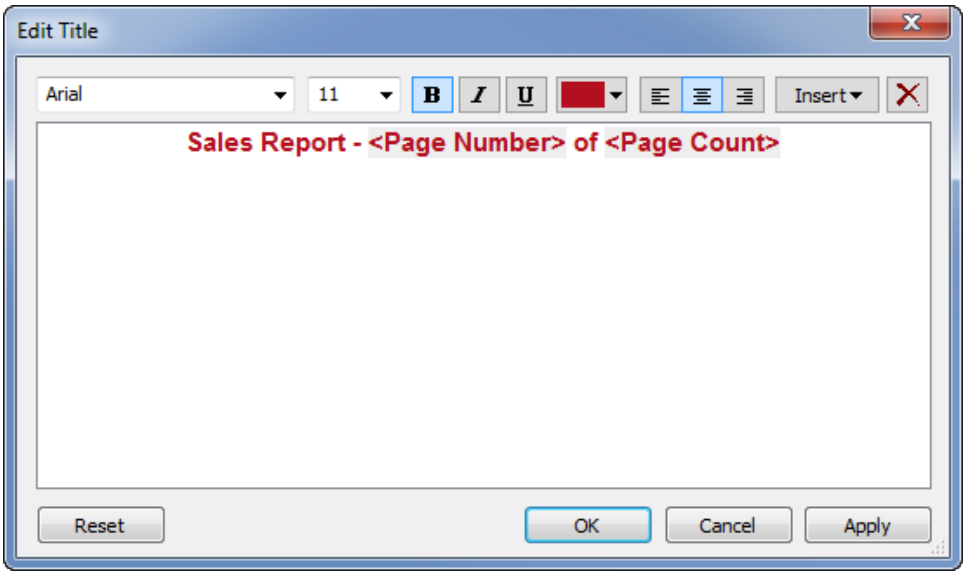

To format title and caption borders and shading:

1. Right-click the title or captionand select **Format**.

   In the format window, use the drop-down controls to add shading and a border.

2. 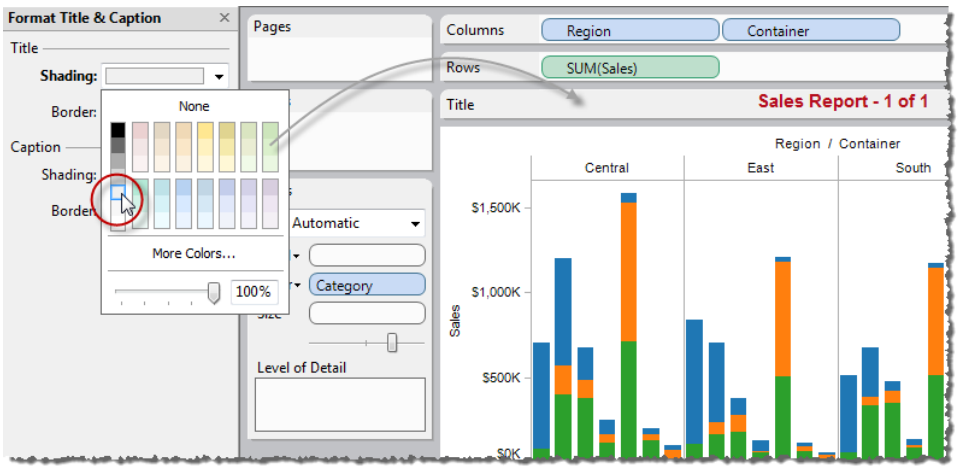

Tooltips

Tooltips are additional data details that display when you rest the pointer over one or more marks in the view. Tooltips also offer convenient tools to quickly filter or remove a selection or view underlying data. 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 **Worksheet** > **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.

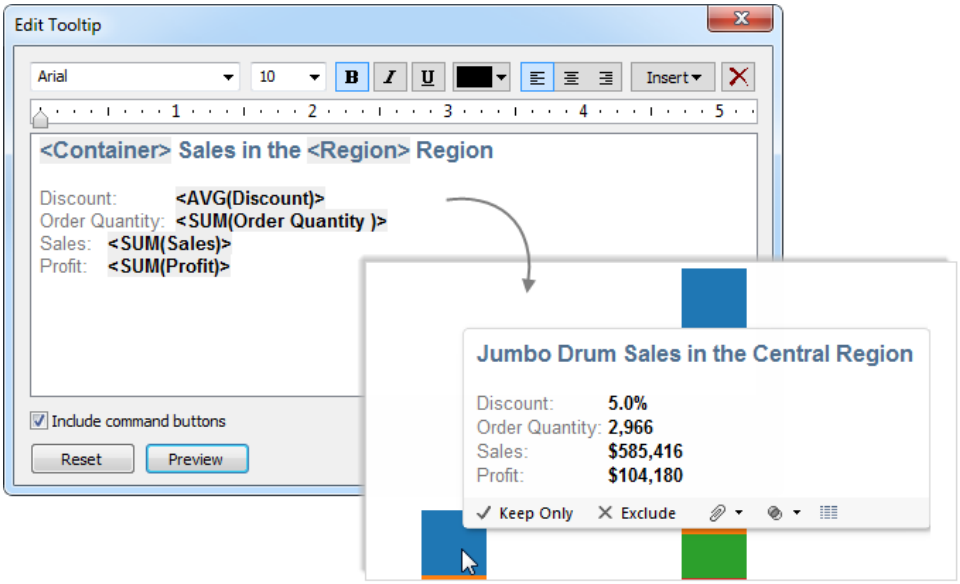

The All Fields option on the Insert menu adds all field names and values that are used in the view. Inserting the All Fields parameter will automatically update the tooltip as you change the view. You can exclude unnecessary fields from the All Fields option.

Select the **Include command buttons** option to show filtering and view data options in the tooltip. For example, including command buttons will add Keep Only, Exclude, and View Data buttons to the bottom of the tooltip. These command buttons are available both in Tableau Desktop and when the view is published to the web or viewed on a mobile device.

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**.

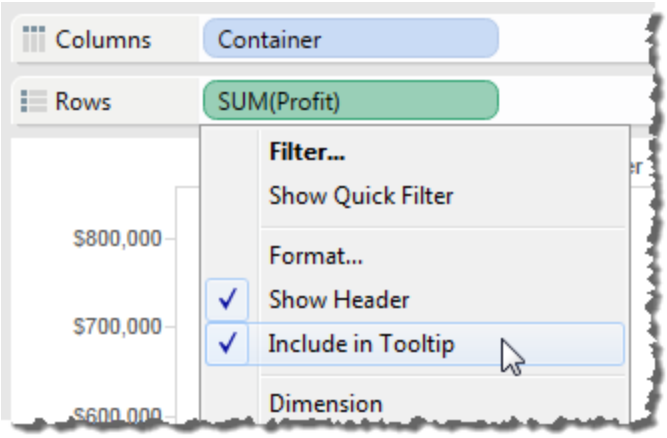

You can show a field in the tooltip using the same menu option.

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 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.

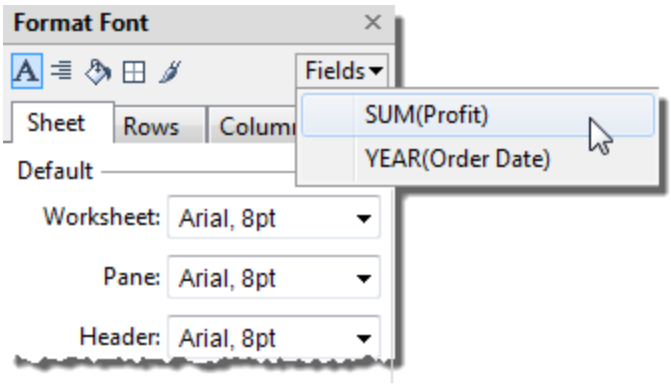

2. Make sure that the **Pane** tab is selected.

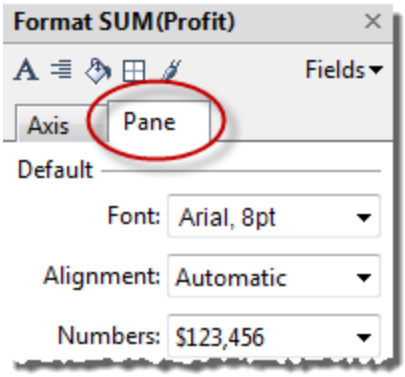

3. Optionally specify the text to label null values when mark labels are turned on.
4. In the Special Values area, optionally specify whether to show the null values using an indicator in the lower right corner of the view, plot them at a default value (e.g., zero for number fields), hide the values but connect lines, or hide and break lines to indicate that null values exist.

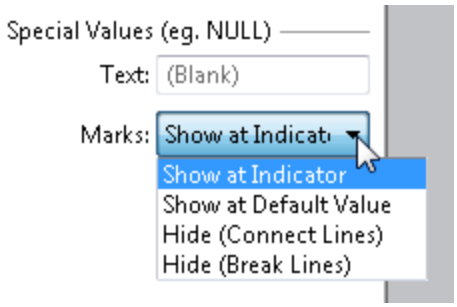

Reference Lines, Bands, and Boxes

Tableau lets you format the reference lines, bands, and boxes you add to an axis in 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, band, or box and select Format. For example:

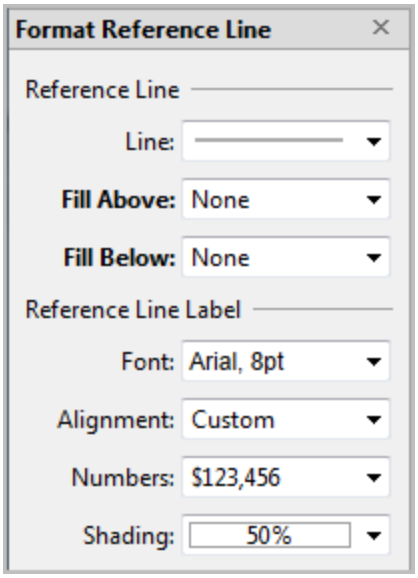

2. You can specify formatting properties associated with reference lines, bands, or distributions. For example, you can change the line style and color as well as the label font properties.
3. When finished, click **OK**.

Many of the formatting options in the Format window are also available when adding and editing reference lines, bands, and boxes.

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. Right-click the worksheet tab and select **Copy Formatting**.
3. Select the worksheet you want to paste the formatting into.
4. Right-click the worksheet tab and select **Paste Formatting**.

> You can also copy and paste formatting by selecting **Copy Formatting** on the Format menu.

Clearing Formatting

Any time you make changes to a setting in the Format window, the label of the setting is bold 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**.

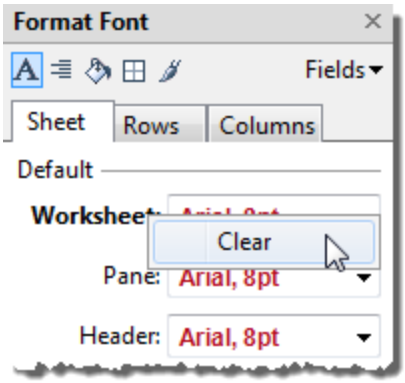

To clear all custom formatting in the worksheet:

- On the toolbar, click the clear menu and select Clear Formatting on the drop-down.

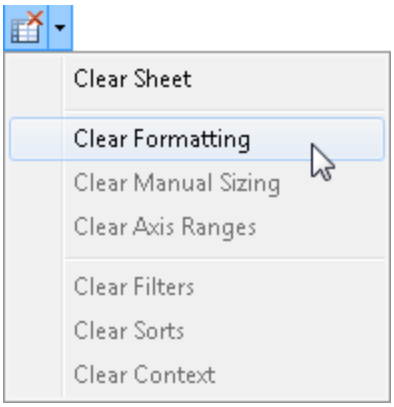

You can also click the Clear button at the bottom of the Format window to clear all format settings currently showing in the Format window.

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 use the Default workbook theme. However, older Tableau 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.
2. Select **Format** > **Workbook Theme** > **Default**.
3. Save the workbook.

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.

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.

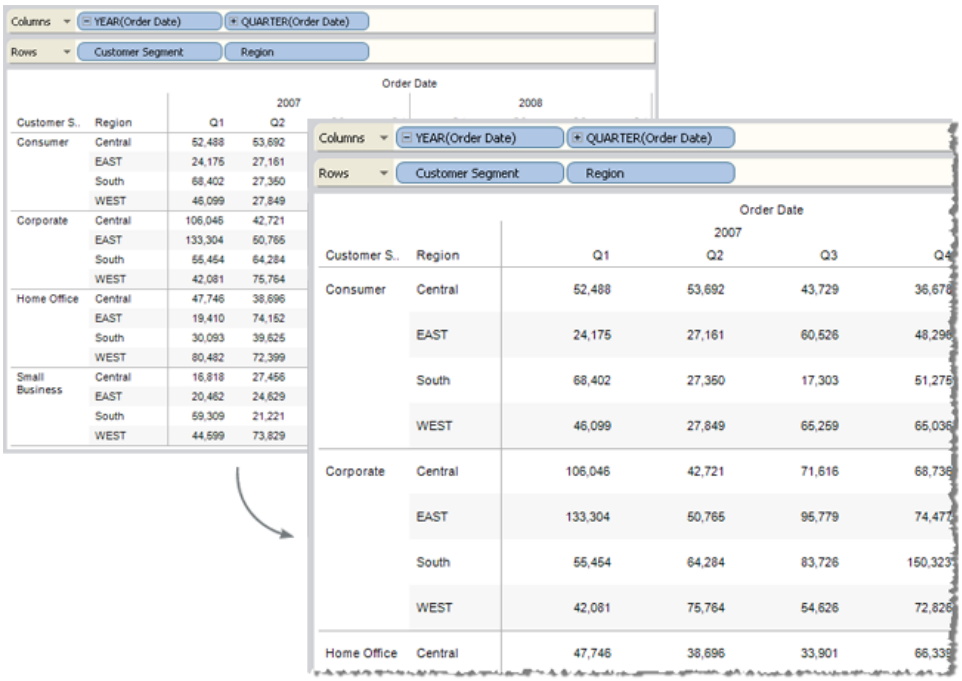

Note:

For a given field, all members will have the same width and the same height. That is, you cannot resize individual field members.

Manually Resizing Rows and Columns

To manually resize the widths or heights of row and column headers or axes:

1. Place your cursor over the vertical or horizontal border of a header or axis.
2. When you see the resize cursor ↔ , click and drag the border left and right or up and down.

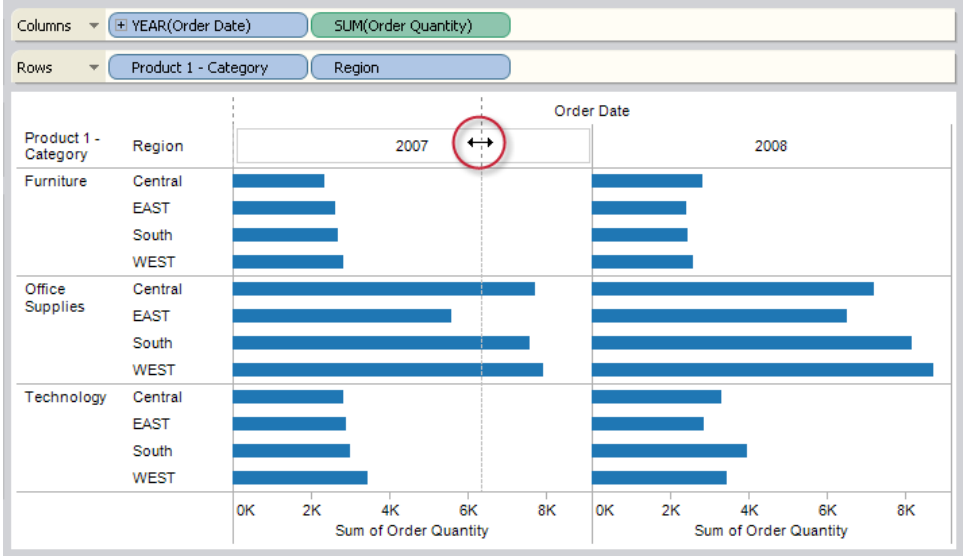

Resizing the Entire Table

You can increase or decrease the size of the entire table by selecting Bigger or Smaller on the Format > Cell Size menu. For example, to increase the width of the columns and the height of the rows for the view shown below, you can select Format > Cell Size > Bigger until the view is of the desired size. This option increases both the width and height of the panes in an intelligent way. Notice that the size of the row headers increase horizontally when you resize the table.

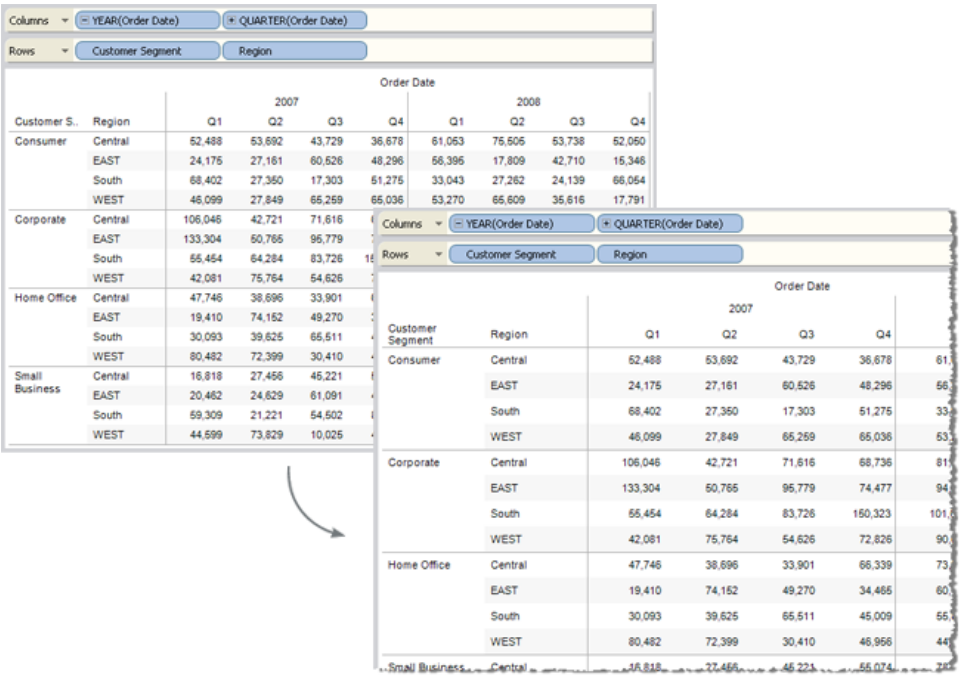

For the view shown below, you can select Format > Cell Size > Smaller to decrease the size of the table.

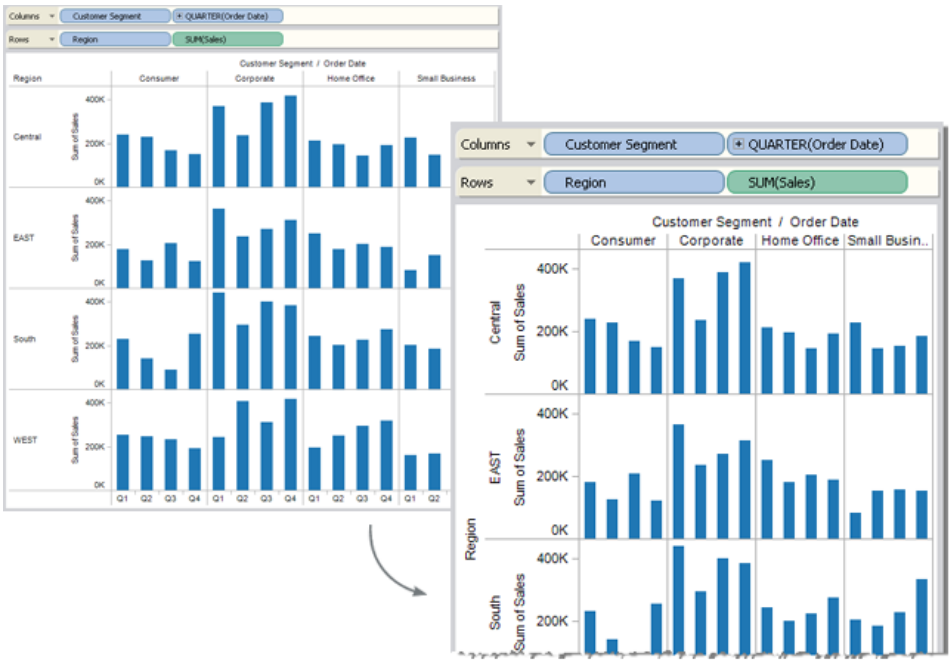

Resizing Cells

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 (that is, farthest to the right) on both the Rows and Columns shelves. In this case, there are two shortcuts you can select on the **Format** > **Cell Size** menu:

- **Square Cell** – Adjusts the view so the cell has a 1:1 aspect ratio. This results in a square cell, which is particularly useful for heat maps.
- **Text Cell** – Adjusts the view so the cell has a 3:1 aspect ratio. This is particularly useful for text tables.

For example, the text table shown below is modified by selecting **Text Cell** from the **Format** menu. This enforces a cell aspect ratio of 3:1 and results in a compact table that is easy to read.

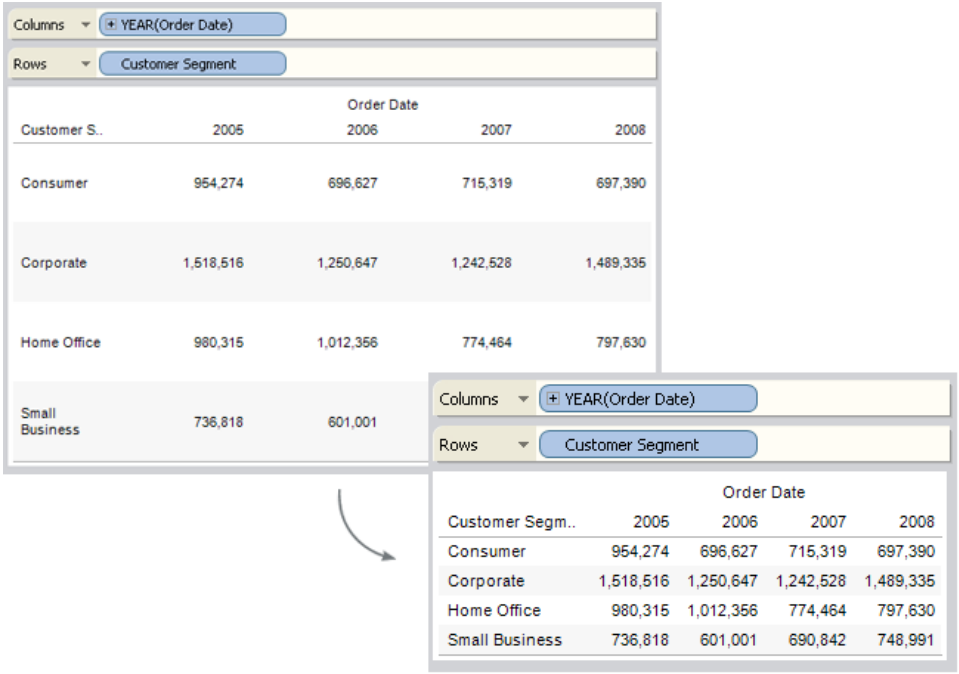

The heat map shown below is modified by selecting **Square Cell** on the **Format > Cell Size** menu. This enforces a cell aspect ratio of 1:1 and results in a compact table that is easy to analyze. You can also use the **Size** slider to adjust the size of each mark.

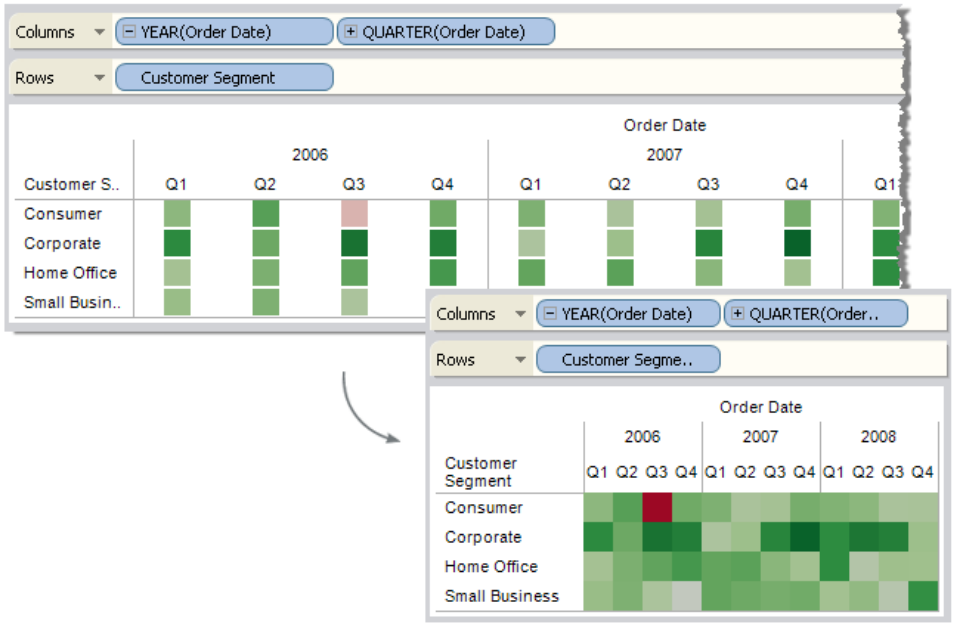

After changing the cell size, you can use **Ctrl**+**B** and **Ctrl**+**Shift**+**B** to decrease or increase the table size while maintaining the cell aspect ratio.

Reordering Rows and Columns

The members in the view may not always be ordered exactly how you want them. 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.

Including and Excluding Rows and Columns

Sometimes you will want to restrict certain members of a field from displaying. To easily exclude a row or column, right-click the header and select **Exclude** on the context menu. Excluding a row or column creates a filter.

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.

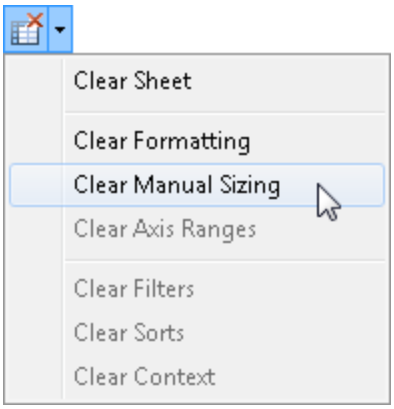

You can also revert to the last saved state by selecting **File** > **Revert to Saved**. This option discards all unsaved changes including manual sizing.

Miscellaneous Table Options

In addition to the standard formatting, there are some other settings that define the table structure. You can modify these settings by selecting **Analysis** > **Table Layout** > **Advanced** to open 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. 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.

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**.

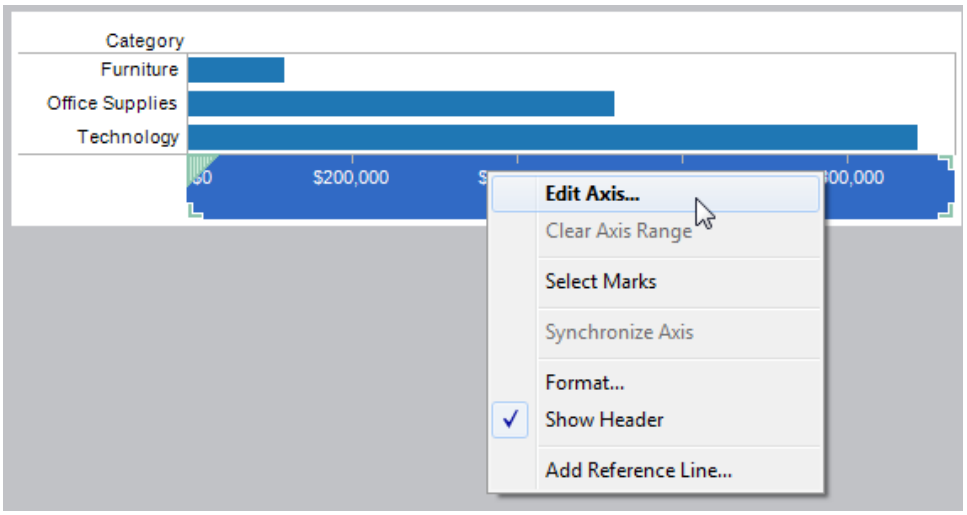

2. In the Edit Axis dialog box, select one of the following:
   - Automatic - The axis range is automatically decided based on the data used in the view.
   - Uniform axis range for all rows or columns - the axis range is the same across all panes in the view.
   - Independent axis ranges for each row 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.

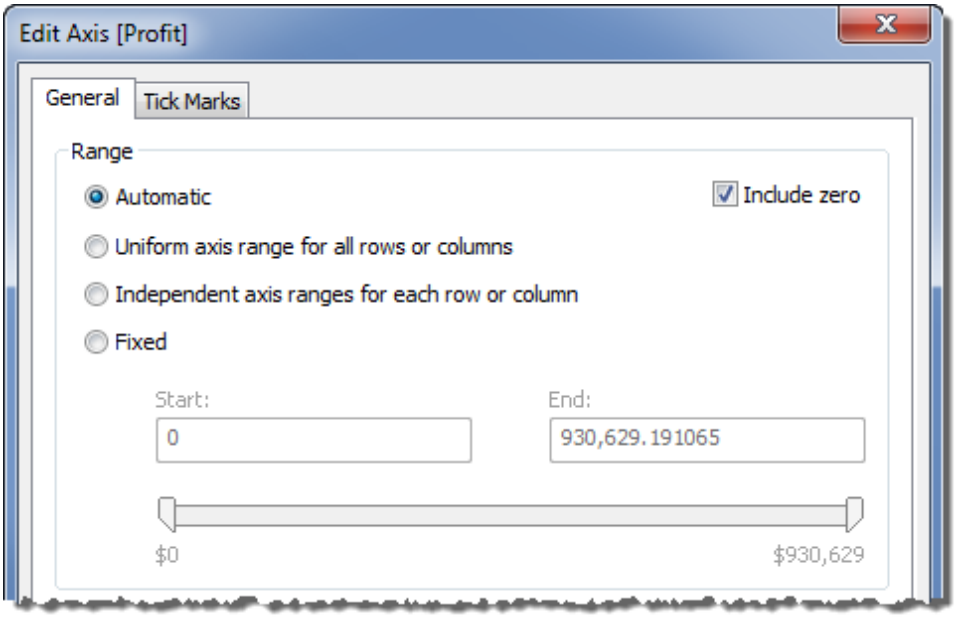

3. You can also specify whether to include zero. When the option include zero is cleared,

the axis range will adjust to just show the range of values in the data.

4. When finished, click **OK**.

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.

   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.

2. 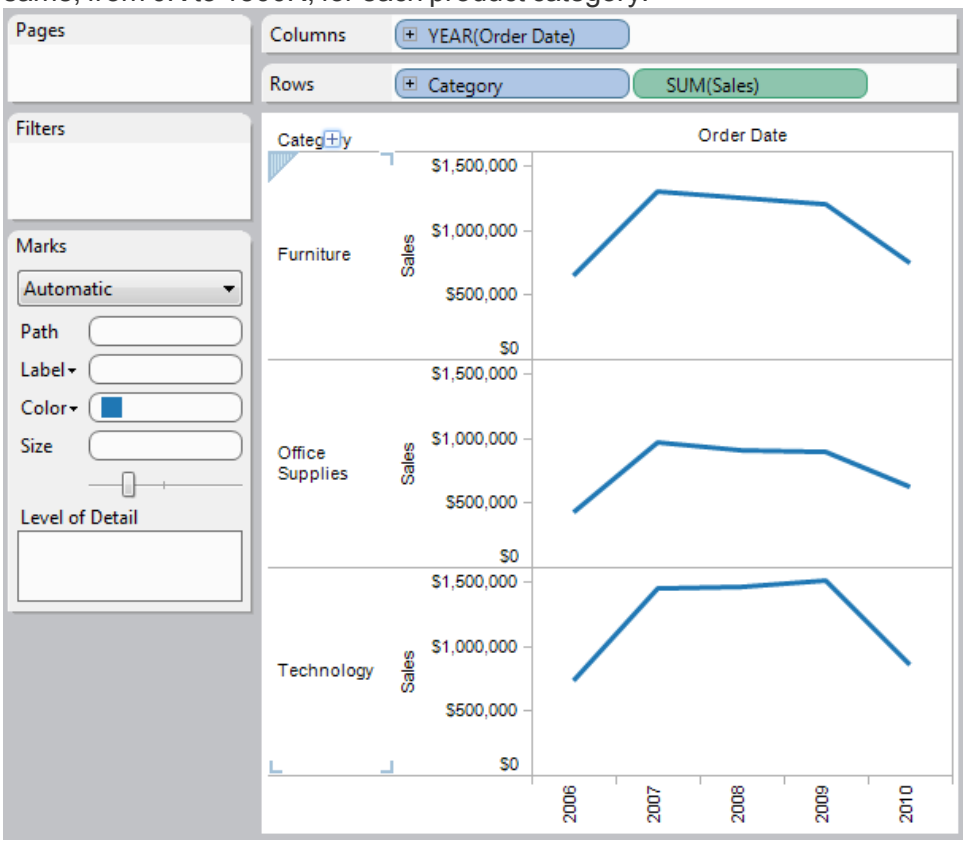

View 2-Building a View using independent axis ranges

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.

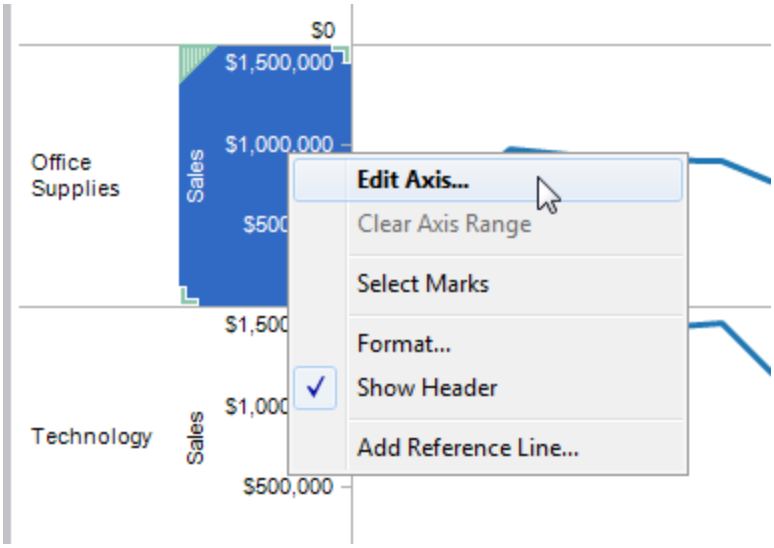

2. In the Edit Axis dialog box, select Independent axis range for each row or column.

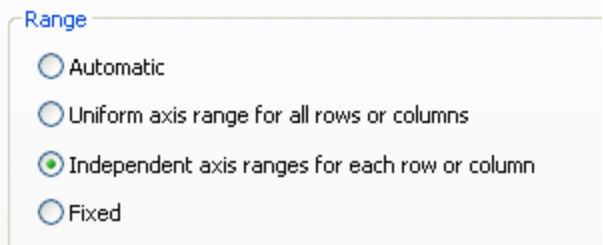

3. When finished, click OK.

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.

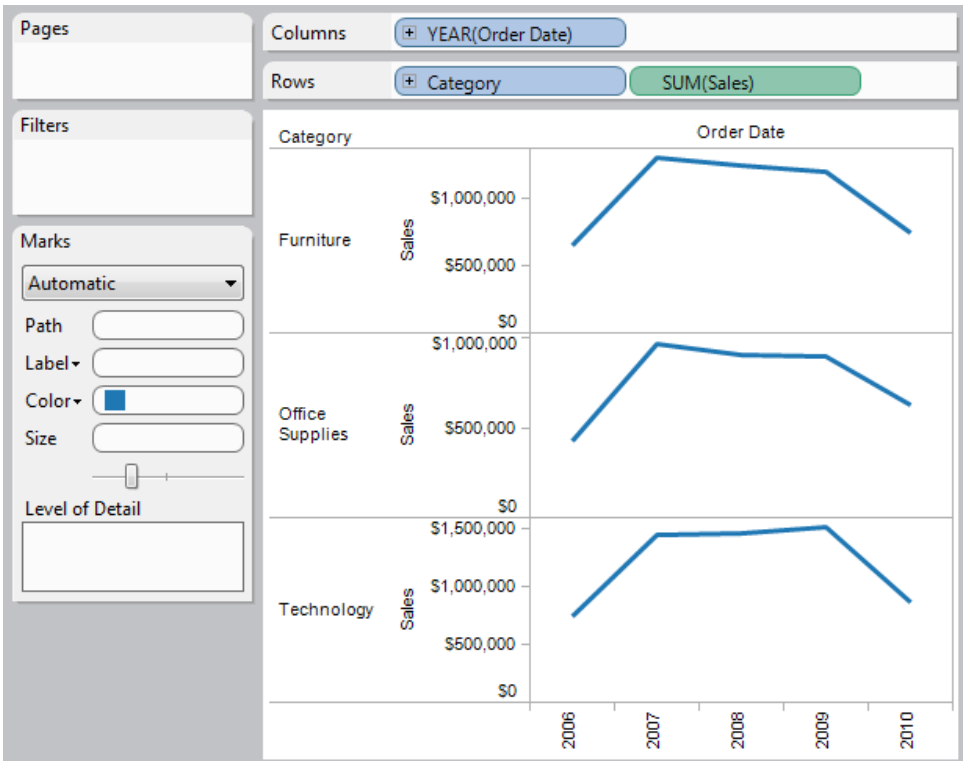

View 3- Building a view using a fixed axis range

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.

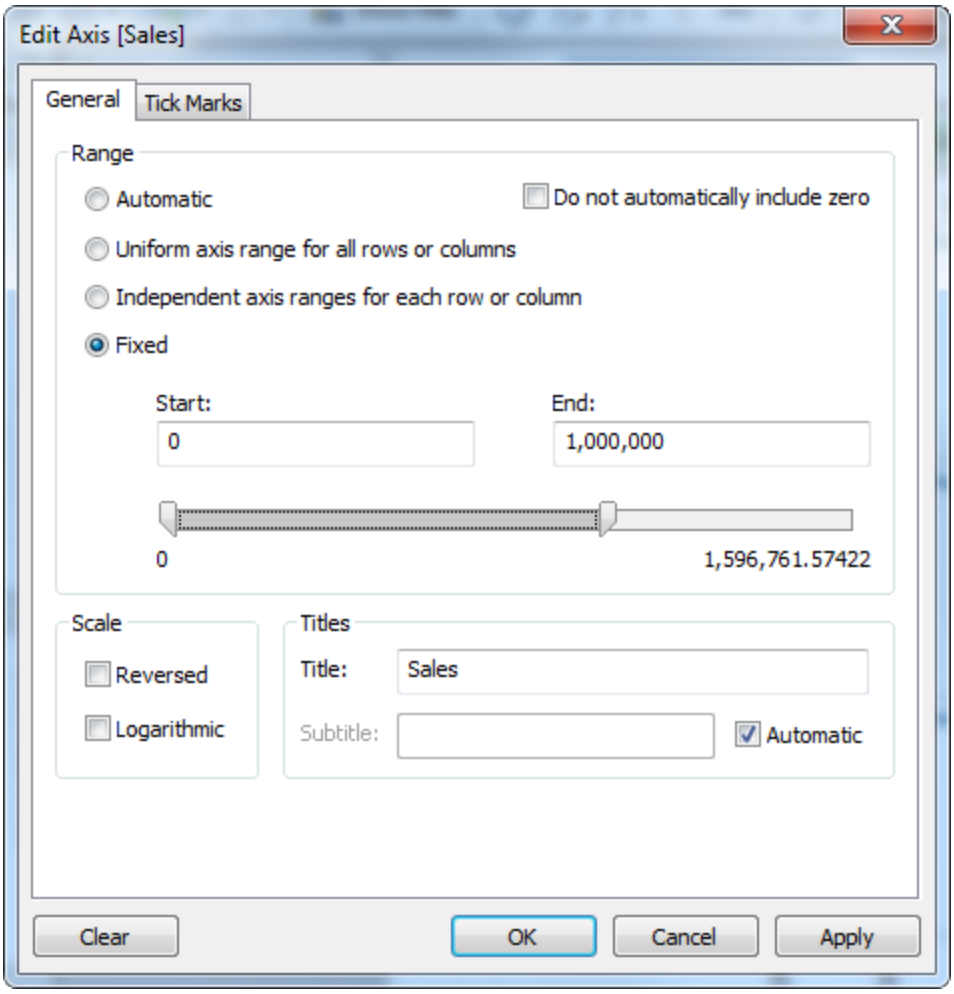

3. When finished, click OK.

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.

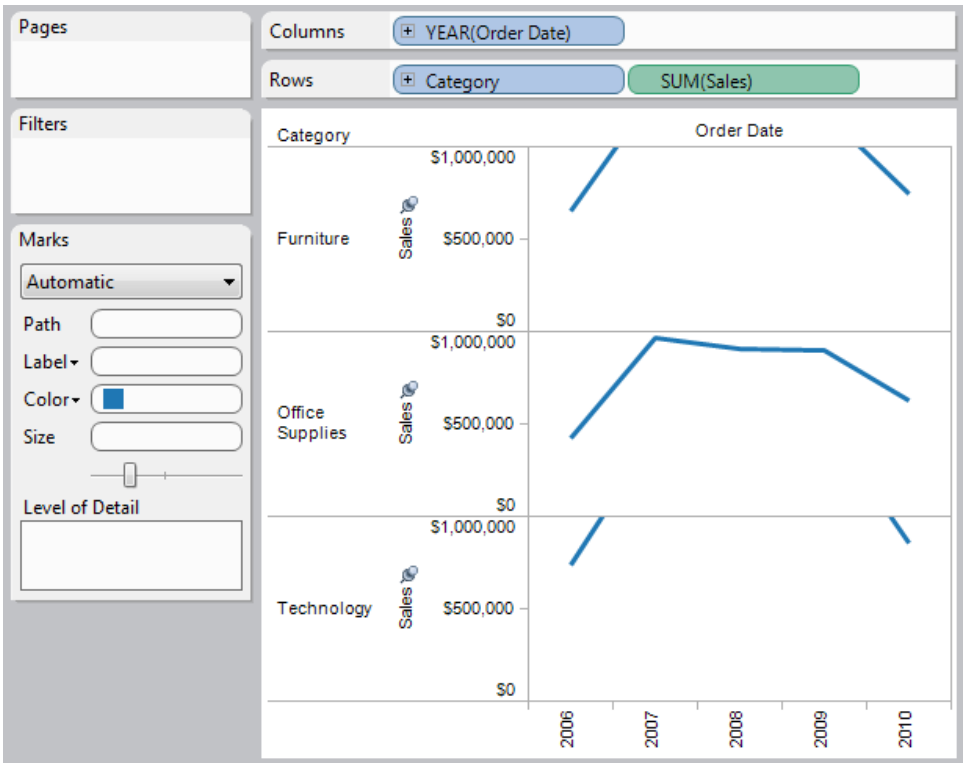

Changing the Axis Appearance

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**.

   In the Edit Axis dialog box, type a new title in the text box.

2. 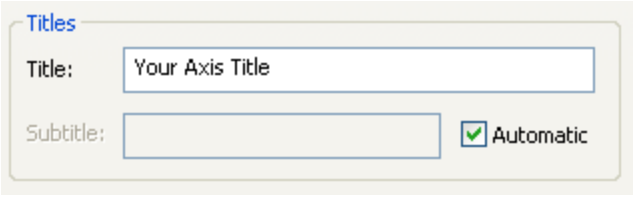

Deselect the **Automatic** check box to add a custom subtitle.

3. 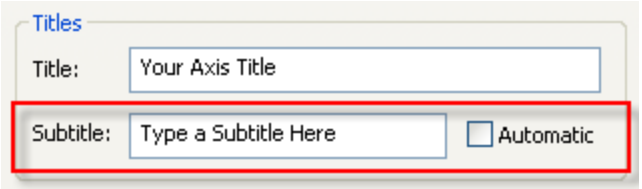

4. When finished, click **OK**.

To change the axis scale:

1. Right-click the axis 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.

     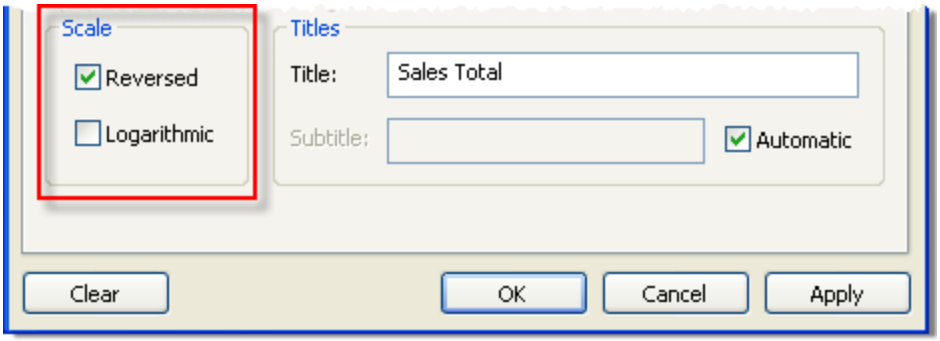

3. When finished, click **OK**.

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.

- 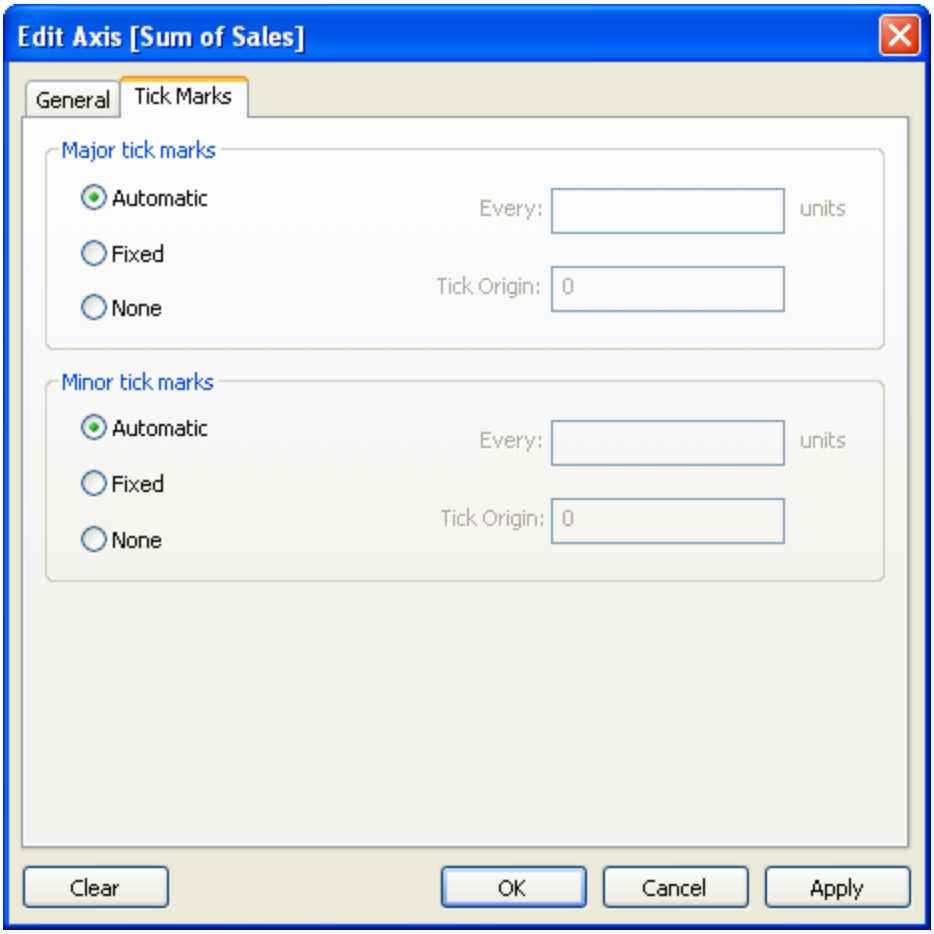

4. When finished, click **OK**.

Annotations and Mark Labels

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.

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.

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**.

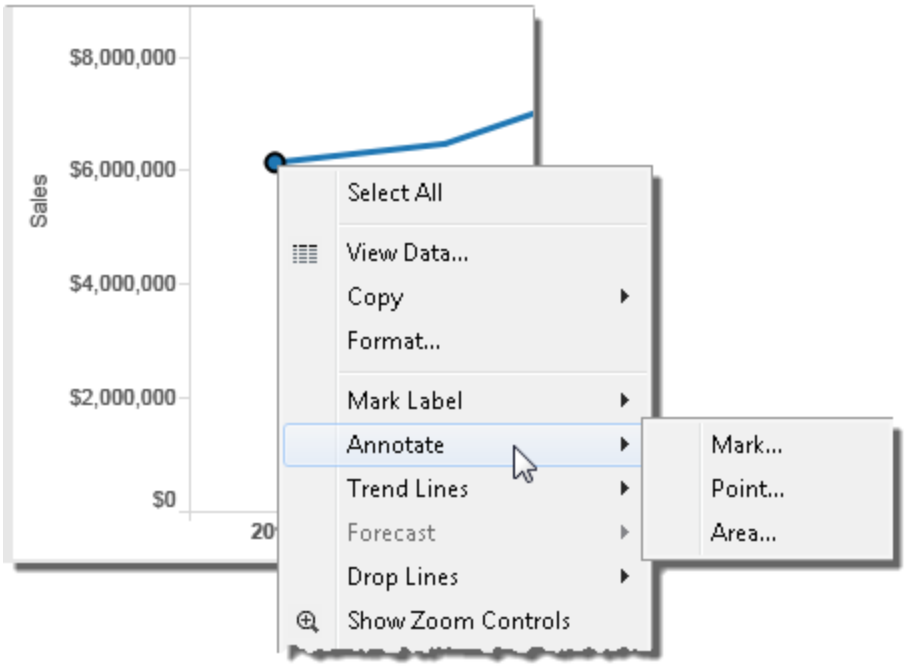

2. On the sub-menu select one of the following types of annotations:
   - **Mark** - select this option to add an annotation that is associated with the selected mark. This option is only available if a mark is selected.
   - **Point** - select this option to annotate a specific point in the view.
   - **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 **Insert** menu to insert dynamic variables into the annotation text. For example, the annotation can display data values that update as the underlying data changes. The dynamic variables that are available are dependent on whether you are annotating a mark, point, or area. Only relevant variables are available for each of these types of annotations. For example, mark annotations can include dynamic data about that specific data point such as dimension and measure values. Point annotations don't refer

to a specific data point but rather a place in the view. For that reason, point annotations can only include dynamic information about that point, such as measure values along the axes. Finally, when adding area annotations, only data relevant to the entire sheet is available, such as title and sheet name.

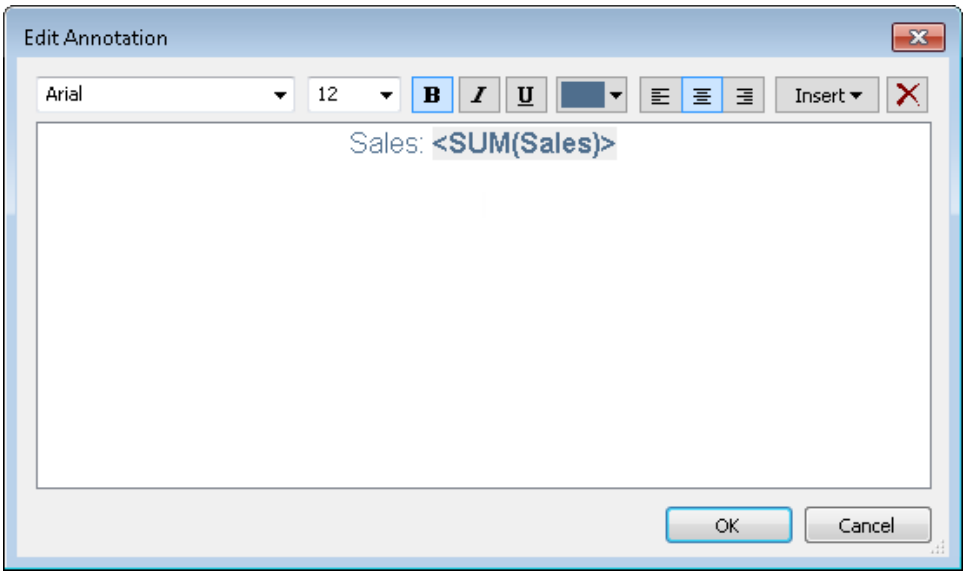

Insert the **All Values** option using the Insert menu to show the data from the fields in the view. As you add more detail to the view, this text is updated to show the live data.

When finished, click **OK**.

4. 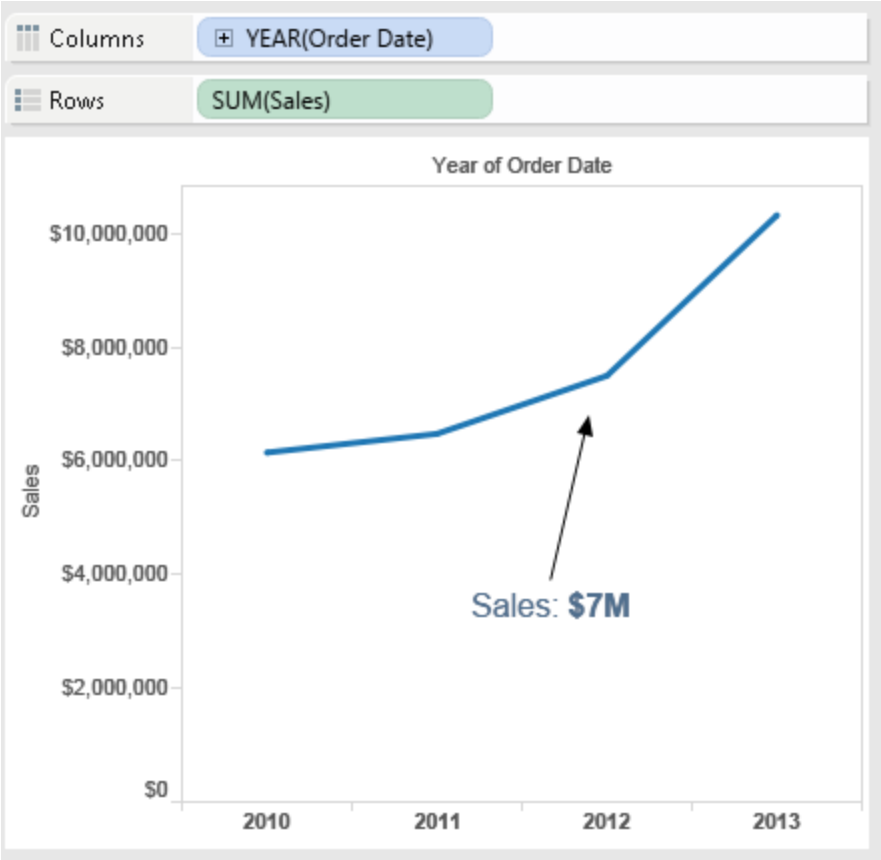

Rearranging Annotations

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.

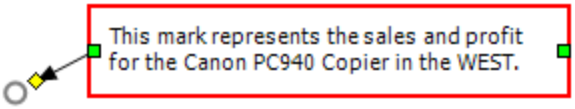

To reposition the body:

Click and drag the body of the selected annotation to a new position.

- 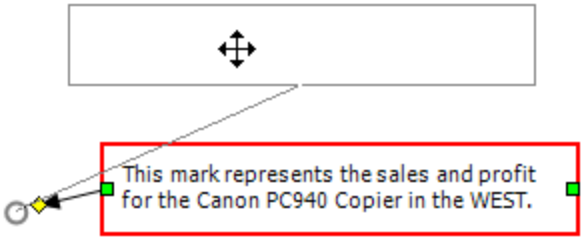

To resize the body:

- Click and drag the body resize handle ■ left and right. The text and height are automatically adjusted to fit the width of the body.

- 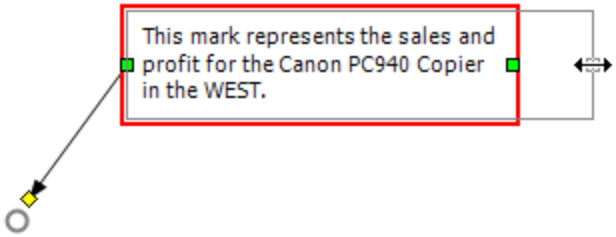

To resize the line:

- Click and drag the line resize handle ◆.

- 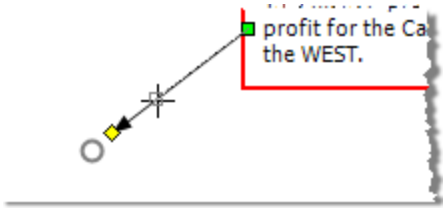

Point Annotations

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.

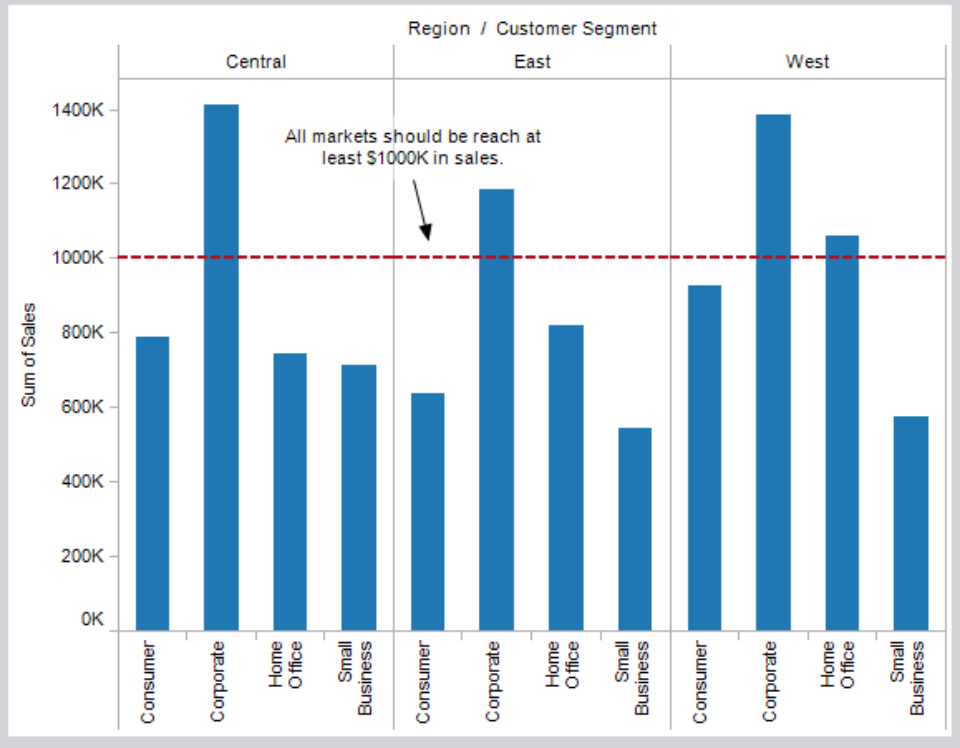

To reposition the body:

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.

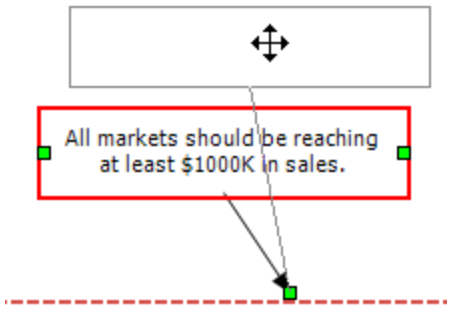

-

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.

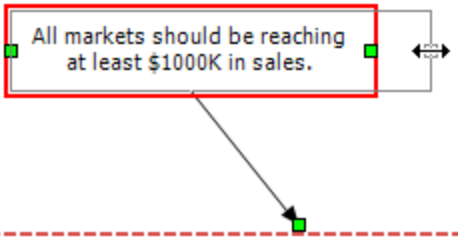

To move the line end point:

- Click and drag the end point of the line

  so that it points at a new location.

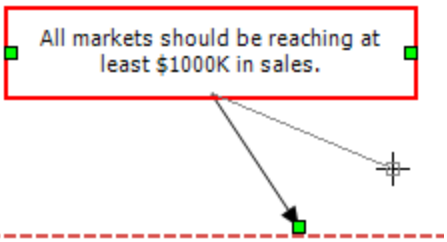

Area Annotations

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.

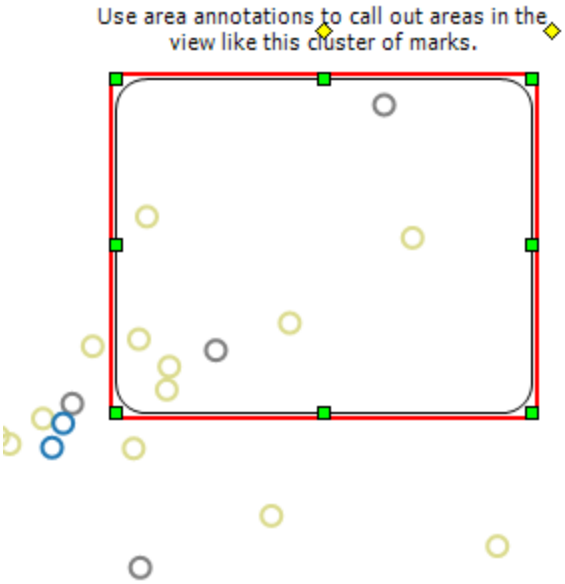

To reposition the box:

- Click and drag the box of the selected annotation to a new position.

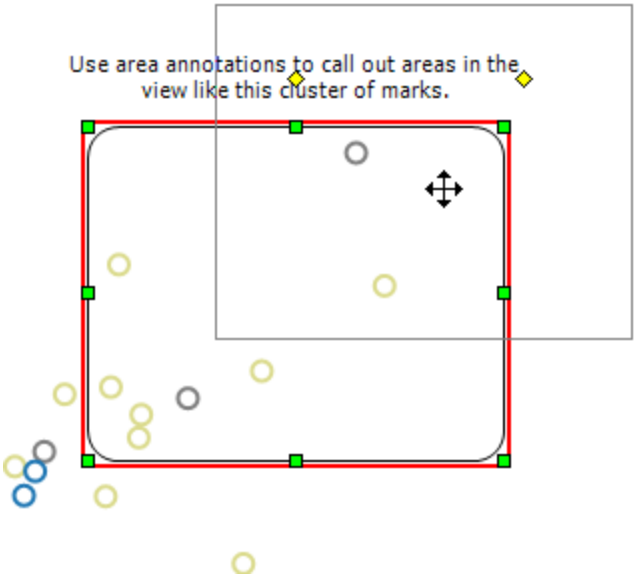

To resize the box:

- Click and drag one of the box resize handles ■

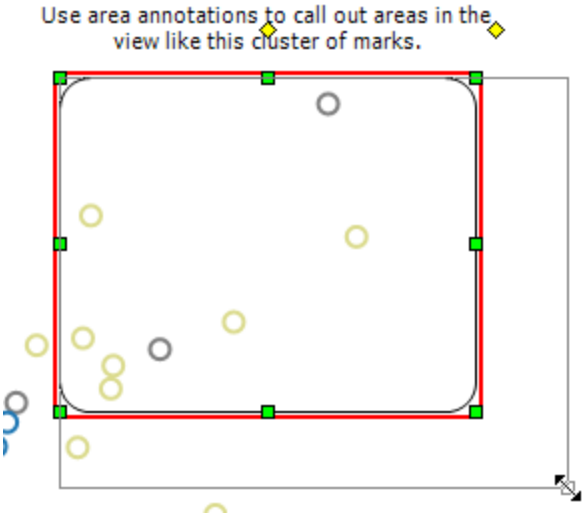

To reposition the text:

- Click and drag the center text handle ◆ to a new position.

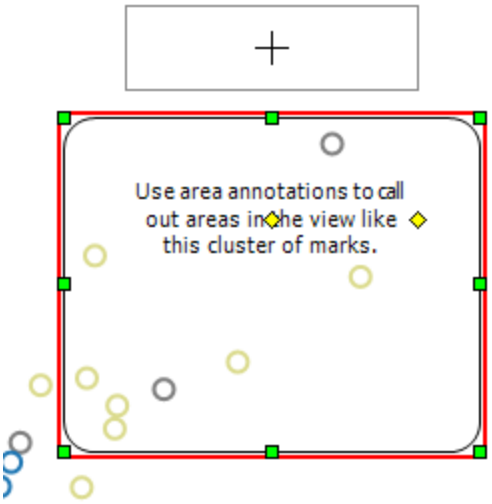

To resize the text width:

- Click and drag the right text handle ◇ left and right. The text height is automatically adjusted to fit the width.

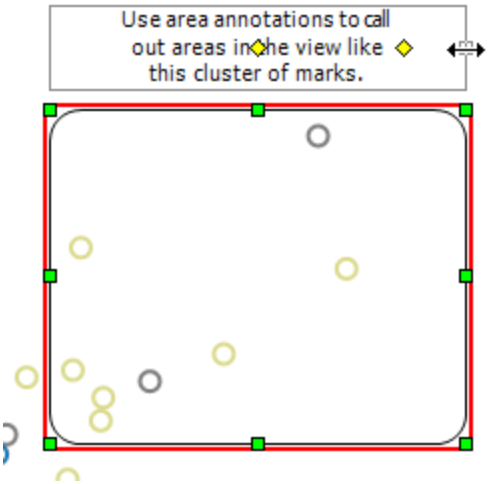

Formatting Annotations

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.

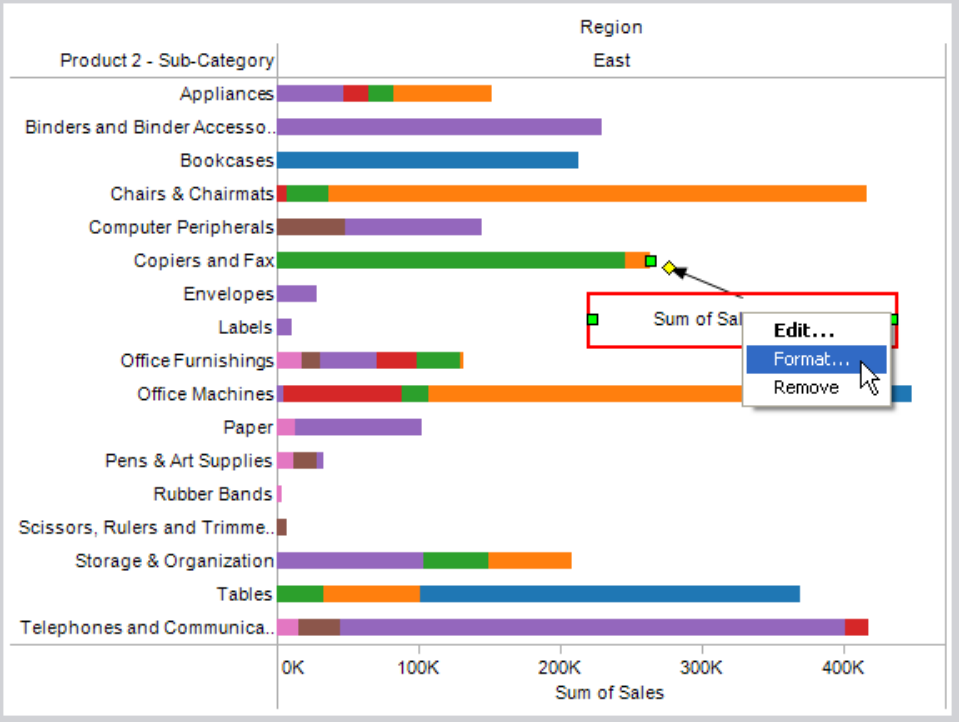

2. In the Format window, use the drop-downs to specify font properties, text alignment, line style, and shading.

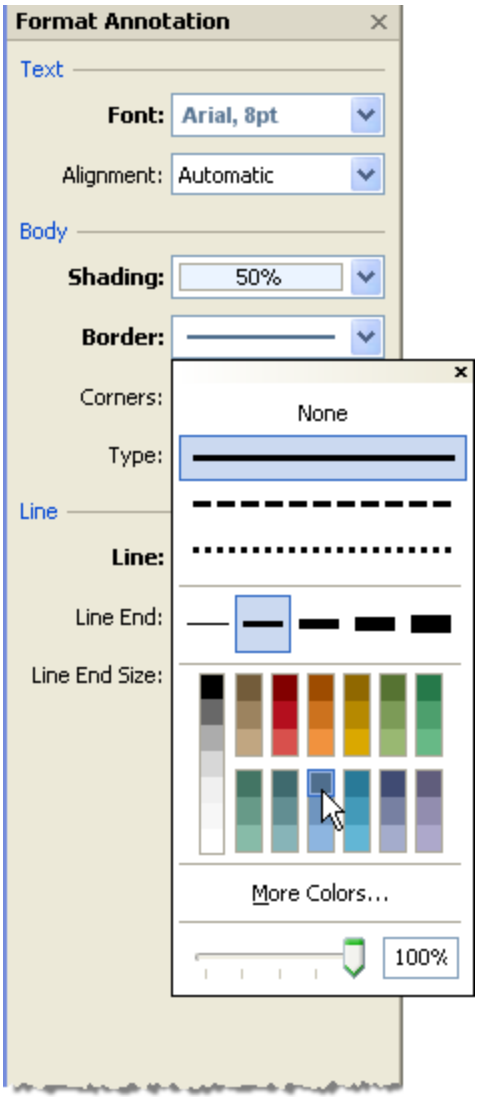

Example Formatting Options

Annotation Style	Format Window Settings
Sum of Gross Profit: 24,524 Sum of Sales Total: 43,178	Default format settings for point and mark annotations.

Annotation Style	Format Window Settings
Sum of Gross Profit: 24,524 Sum of Sales Total: 43,178	Body **Shading:** **Border:** None
Sum of Gross Profit: 24,524 **Sum of Sales Total: 43,178**	Body **Shading:** **Border:** **Corners:** More Rounded Type: Box Line **Line:** **Line End:** Dot Line End Size: Medium
Sum of Gross Profit: 24,524 **Sum of Sales Total: 43,178**	Body Shading: **Border:** Corners: Square **Type:** Single Edge Line **Line:** **Line End:** Dot **Line End Si...** Small

Annotation Style	Format Window Settings
	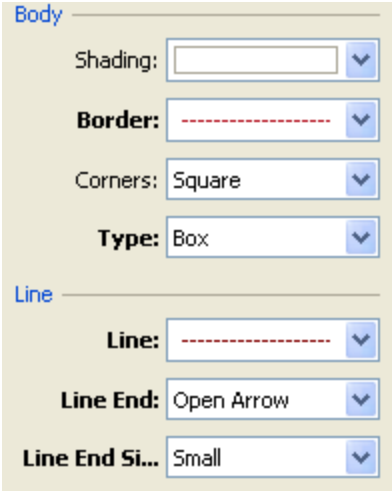
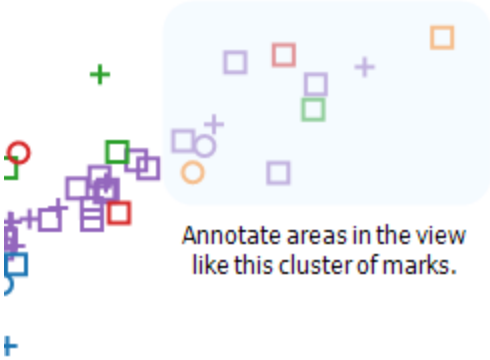	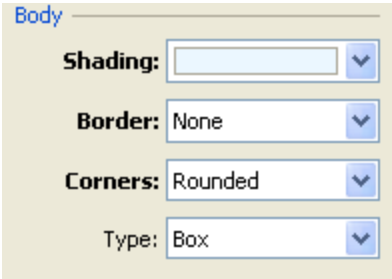
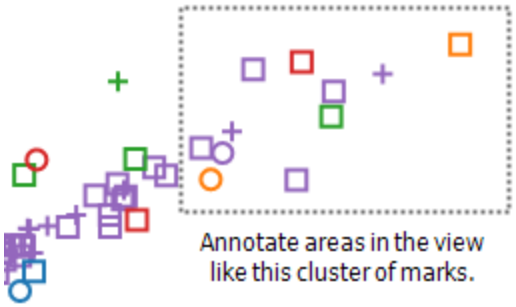	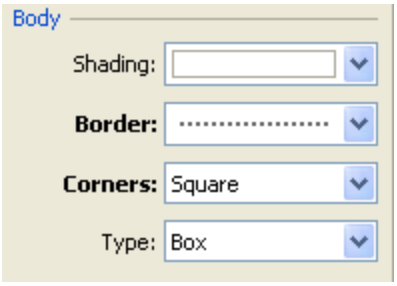

Removing Annotations

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 press 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.

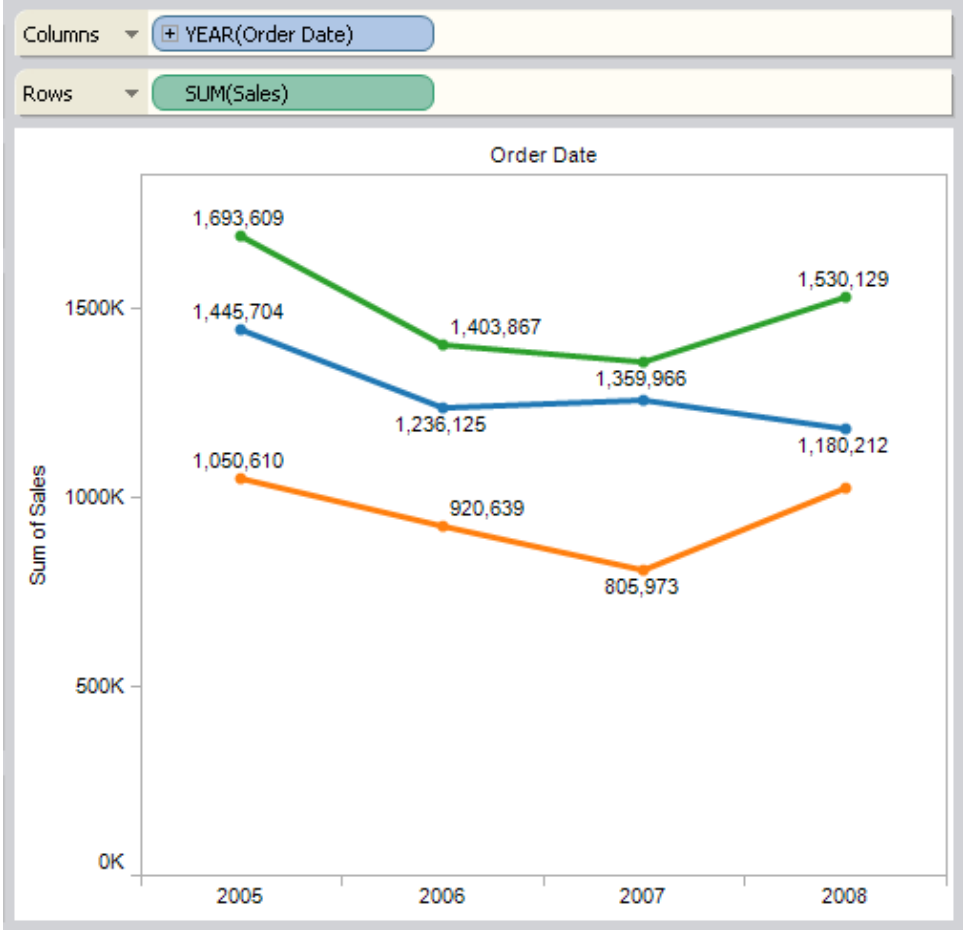

Mark Labels can be turned on for all marks, selected marks, highlighted marks, the minimum and maximum values, or just the line ends in a line chart. Additionally, you can turn on mark labels for an individual marks. This section discusses the following topics:

Showing and Hiding Mark Labels for the Worksheet

You can show and hide mark labels for the whole worksheet by clicking the Show Mark Labels button on the toolbar. If there are no fields on Label, the mark labels indicate the

corresponding measure values based on the Measures on the Rows and Columns shelves. Click **Label** on the Marks card to specify font properties, alignment, when to show the label, and even customize the text itself.

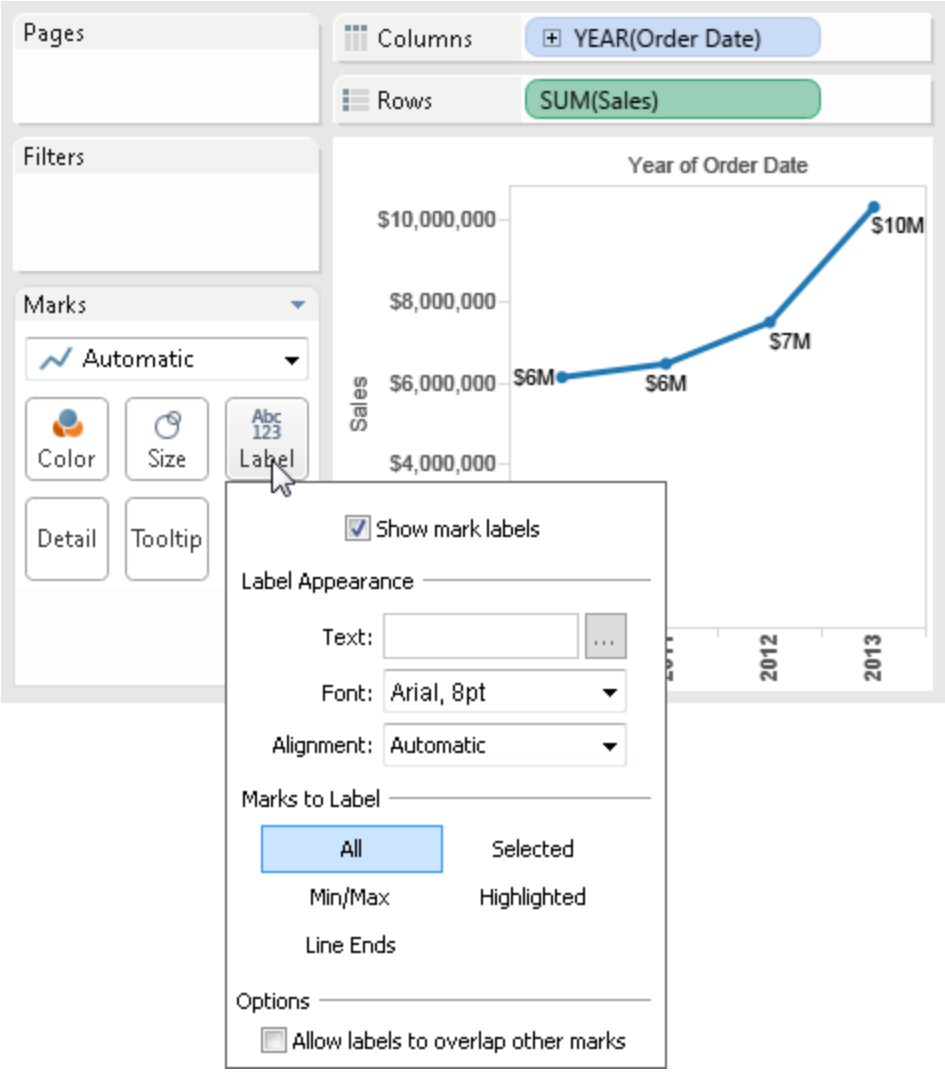

Showing and Hiding Individual Mark Labels

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 Label** 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).

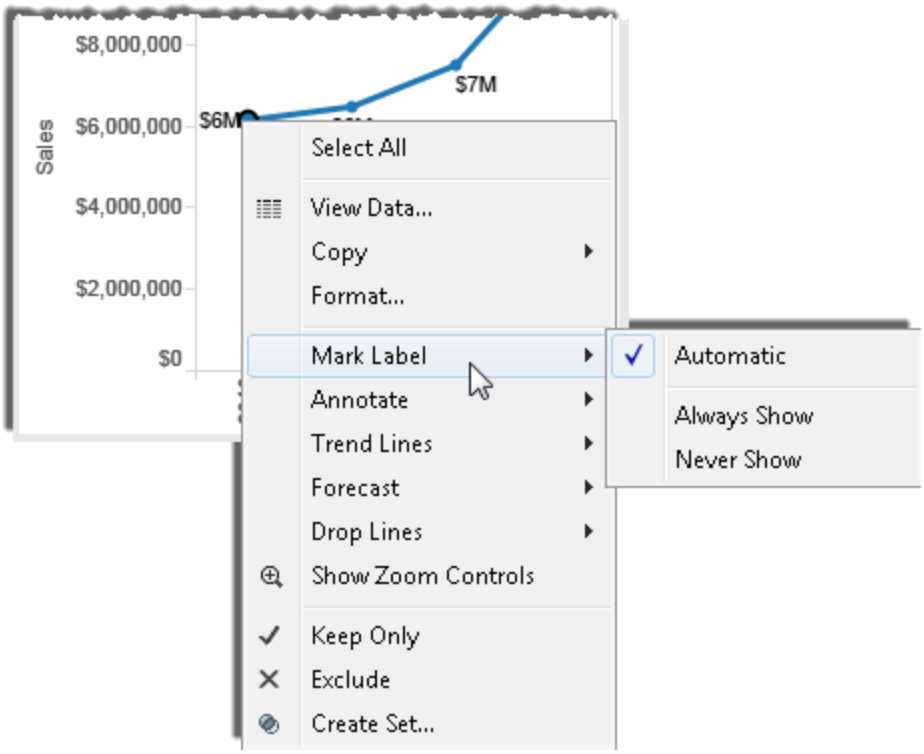

Moving Mark Labels

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.

1. Select the mark whose mark label you want to move.
2. Click and drag the move handle to a new location.

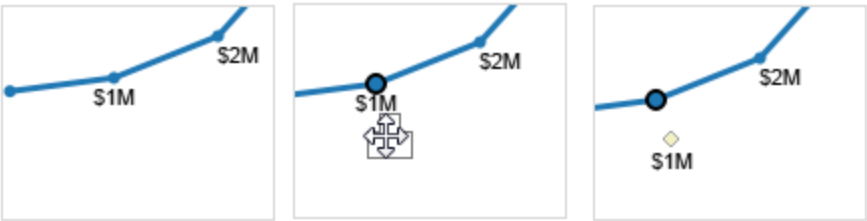

Editing Mark Labels with Aliases

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.

Dashboards

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.

Creating Dashboards

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.

Select **Dashboard** > **New Dashboard**.

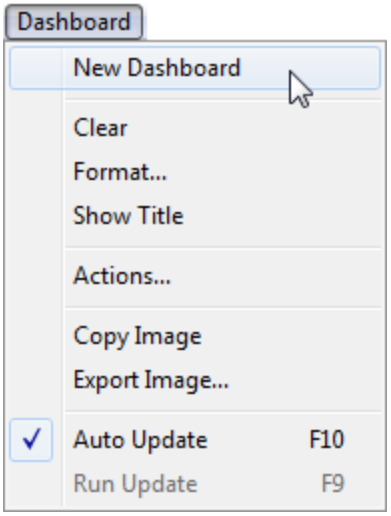

Alternatively, click the New Dashboard tab along the bottom of the workbook.

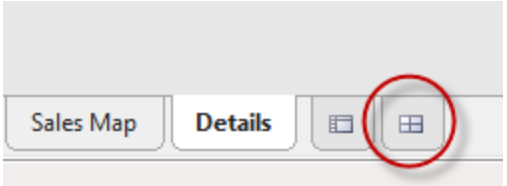

A new tab for the dashboard is added along the bottom of the workbook. Switch to the new dashboard to add views and objects.

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.

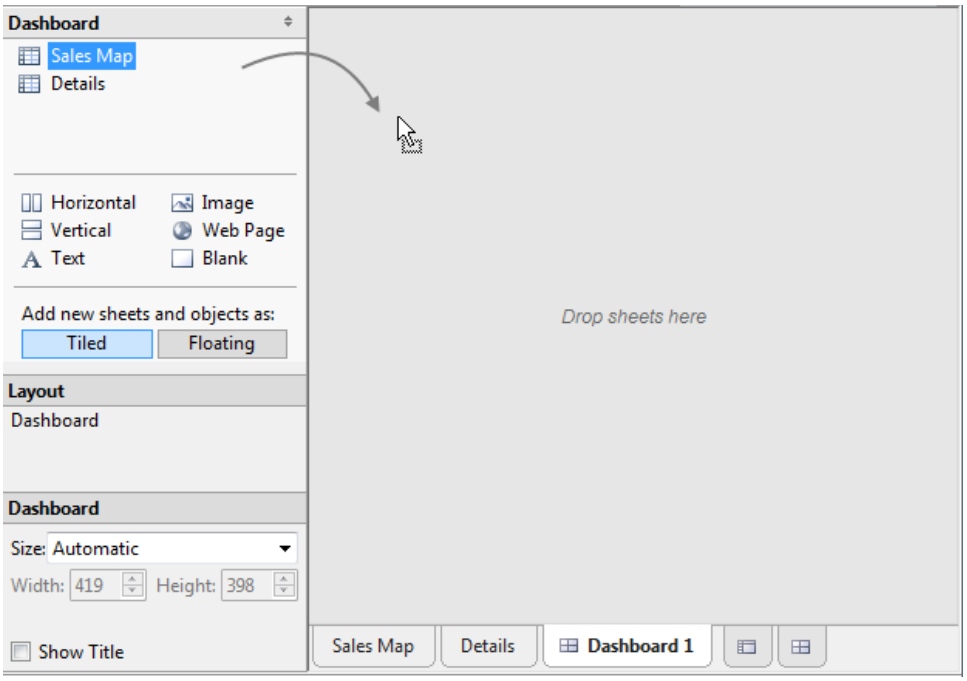

Continue to drag as many of the worksheets to the dashboard as you like. Notice as you drag worksheets around the dashboard, a gray shaded area indicates the various places you can drop it.

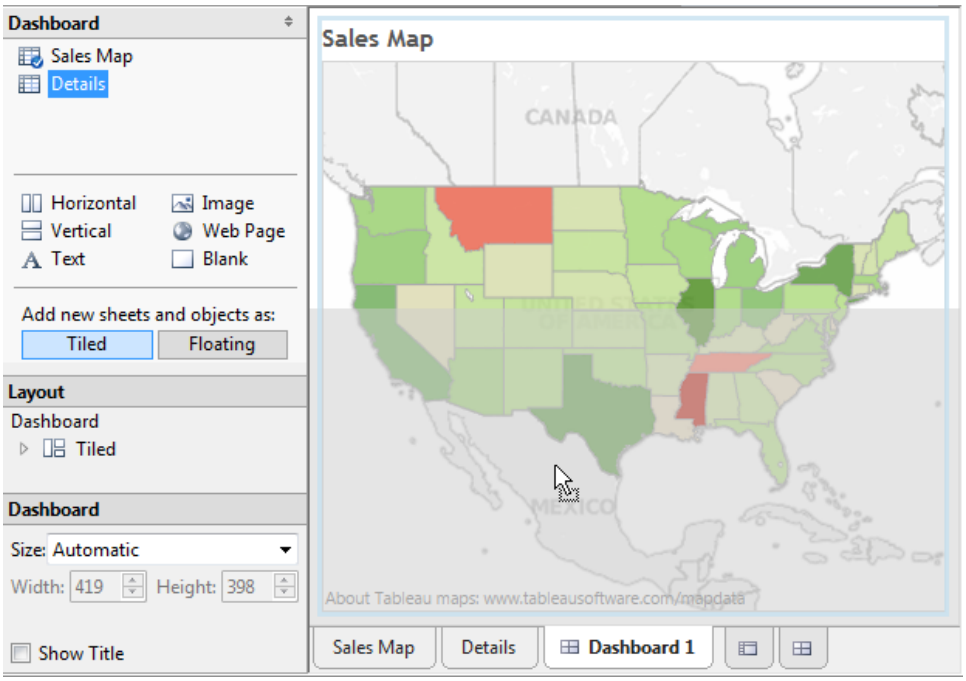

After a view is added to the dashboard, the worksheet is marked with a check mark in the Dashboard window. Also, any legends or quick filters that are turned on for the sheet are automatically added to the dashboard.

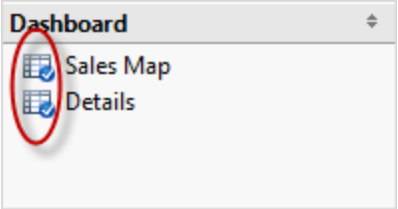

By default, dashboards use a Tiled layout, which means that each view and object is arranged into a single layered grid. You can change the layout to Floating to allow views and objects to overlap. See **Organizing Dashboards** on page 792 to learn more about these layouts.

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 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:
  - **Horizontal** - The Horizontal object adds a horizontal layout container. The widths of views and objects placed inside a horizontal layout container are automatically adjusted to equally fill the width of the container. See **Layout Containers** on page 795 to learn more about these objects.
  - **Vertical** - The Vertical object adds a vertical layout container. The heights of views and objects placed inside a vertical layout container are automatically adjusted to equally fill the height of the container. See **Layout Containers** on page 795 to learn more about these objects.
  - **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. By default, text objects are transparent. To change this, right-click the text object in the dashboard, and select **Format**.
  - **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. When you add an image to a dashboard, you can customize how the image displays by selecting an option on the image's menu.

For example, you can select whether to **Fit Image**, which scales the image to the size of image object on the dashboard. You can also select whether to **Center Image**, which aligns the image to the center of the image object on the dashboard. Finally, you can **Set URL**, which turns the image into an active hyperlink on the dashboard. If you are designing a dashboard that will be viewed on a Retina display, see **Displaying High Resolution Images** on page 791 for more information.

- **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 then open in the dashboard instead of in a browser window.

  When you add a web page object you are asked to specify a URL. If you will publish the dashboard to a server, a best practice is to to use the **https://** protocol with URL actions.

  If you use http:// in the URL action, and the server on which you publish the dashboard runs through https://, browsers will not be able to display the http:// page that the URL action points to.

  When printing the Dashboard to PDF, the contents of the web page are not included.

- **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 many different types of dashboard objects.

Image	Web Page	Blank	Text

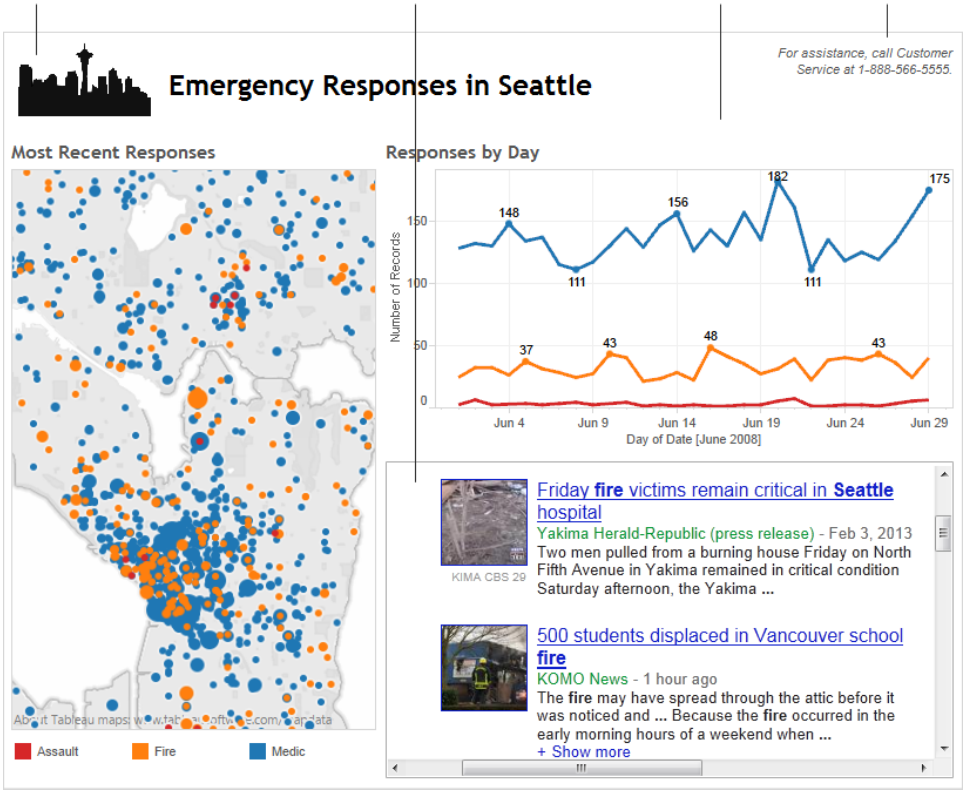

By default, dashboards use a Tiled layout, which means that each object is arranged into a single layered grid. You can change the layout to Floating to allow views and objects to overlap. See **Organizing Dashboards** on page 792 to learn more about these layouts.

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.

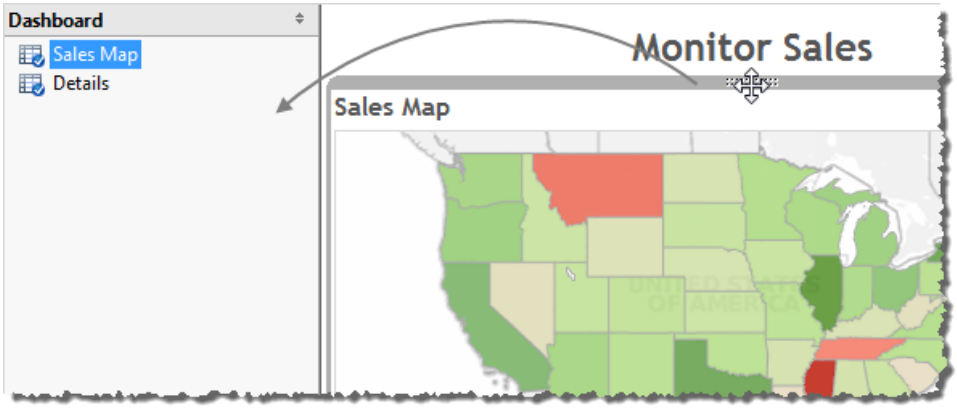

To remove a worksheet using the Dashboard window:

- Right-click the worksheet in the Dashboard window and select **Remove from Dashboard**.

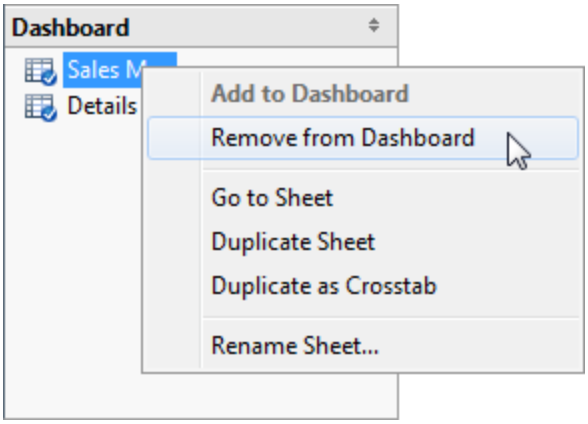

To remove a worksheet or object using the dashboard view menu:

1. Select the view or object in the dashboard you want to remove.
2. Either select **Remove from Dashboard** on the dashboard view menu or click the "x"

icon in the upper right corner of the selected object.

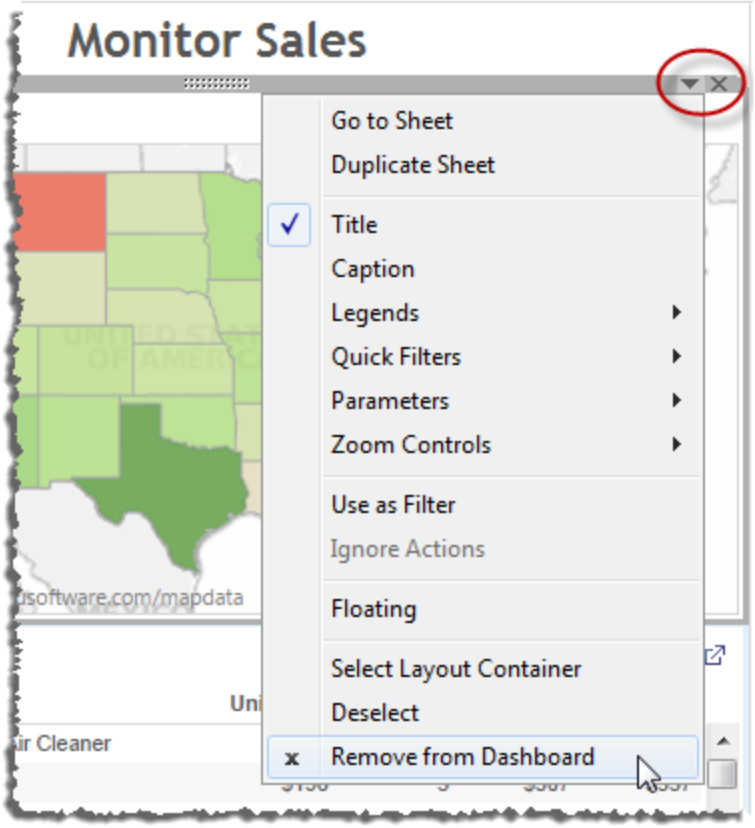

Displaying High Resolution Images

By default, high-resolution images that you use in a dashboard display in standard resolution. The default configuration ensures that images maintain the same size when viewed on either standard resolution or Retina displays. To take advantage of high-resolution images in a dashboard that are intended to be viewed on a Retina display, use the @2x naming convention.

To display a high resolution image on a dashboard:

1. Navigate to the high-resolution image on your computer with Retina display.
2. Rename your high-resolution image to include the @2x modifier by using <*image-name*>@2x<*file-extension*> format. For example, logo@2x.png.
3. In your workbook, click and drag an image object and drop it on the dashboard. You are prompted to select an image from your computer.
4. Select the image you just renamed using the @2x naming convention and then click **Open**.

Note: EMF image file formats are not compatible with Tableau Desktop on the Mac.

Organizing Dashboards

A dashboard is made up of several views, objects, legends, quick filters, and so on. Each view and object can be rearranged or hidden to help you highlight the most important information and tell a story with your data. For each view or object in the dashboard, you can specify its position, its size and whether it uses a tiled or floating layout. You can also use layout containers to build dashboards that dynamically adjust and resize based on the data showing. Layout containers are particularly useful when the dashboard has filtered data that results in a variable view size.

Tiled and Floating Layouts

Each object in a dashboard can use one of two types of layouts: tiled or floating. Tiled objects are arranged in a single layer grid that adjust in size based on the total dashboard size and the objects around it. Floating objects can be layered on top of other objects and can have a fixed size and position.

Tiled Layout

All objects are tiled on a single layer. The top three views are in a horizontal layout container.

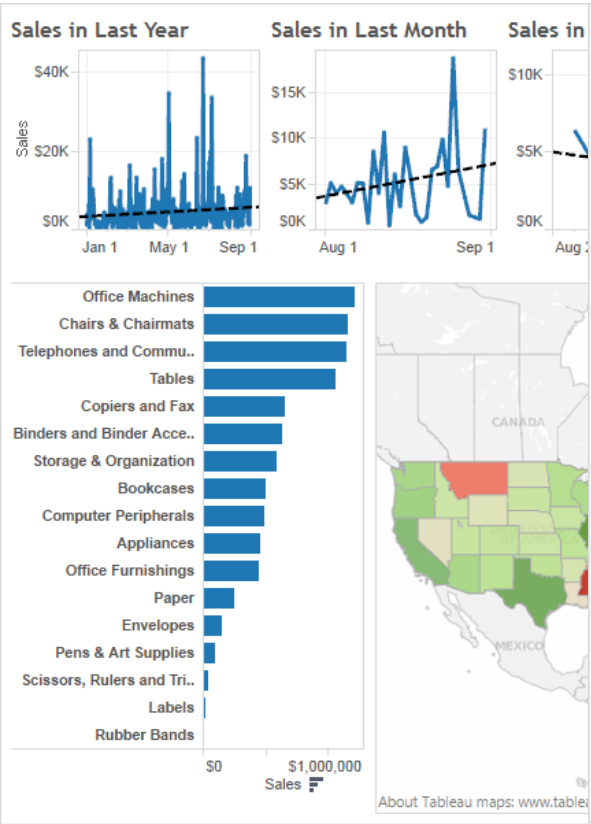

Floating Layout

While most objects are tiled on this dashboard, the map view and its corresponding color legend are floating. They are layered on top of the bar chart, which uses a tiled layout.

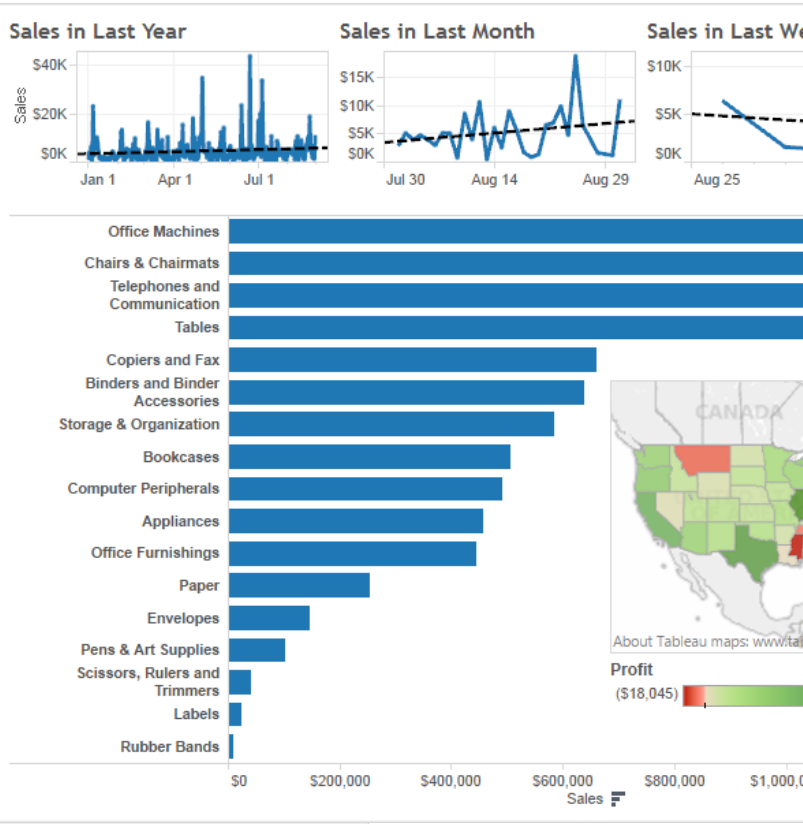

Switch Between Layouts

By default, dashboards are set to use the Tiled layout. All views and objects are added as tiled. To switch an object to floating, do one of the following:

- Select a view or object in the dashboard and then select the **Floating** option at the bottom of the Dashboard window.

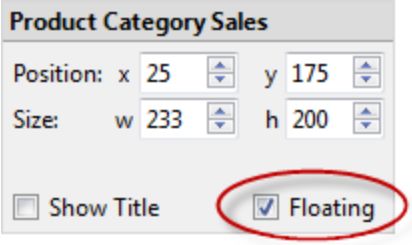

- Hold down **Shift** on your keyboard while dragging a new sheet or object to the view. Alternatively, select an existing view or object in the dashboard and hold **Shift** while dragging the object to a new location in the dashboard.

Similarly, you can convert a floating object back to tiled in the same ways listed above.

To change the default for the whole dashboard, click the **Floating** button in the middle of the Dashboard window. When the dashboard is set to Floating layout, any new sheets and objects are added as floating.

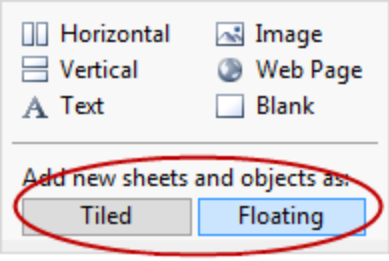

Reorder and Resize Floating Objects

All of the items in a dashboard are listed in the Layout section of the Dashboard window. The Layout section displays tiled objects and any floating objects in a hierarchy. You can drag and drop the items in the hierarchy to change the order that they are layered in the dashboard. Items shown at the top of the list are in the front while items at the bottom of the list are in the back.

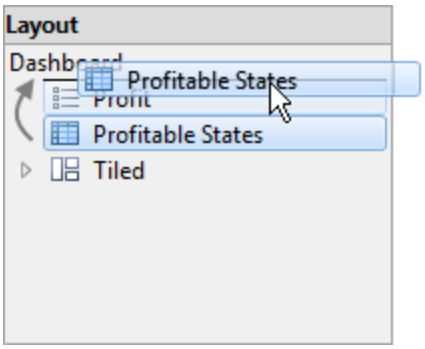

Right-click items in the Layout section to customize the object and hide and show parts of the worksheet. See **Showing and Hiding Parts of a Worksheet** on page 800.

Use the **Position** field at the bottom of the Dashboard window to specify exact location for floating objects. Define the position in pixels as an offset from the top left corner of the dashboard. The **x** and **y** values specify the location of the top left corner of the object. For example, to place an object in the top left corner of the dashboard, you would specify x = 0 and y = 0. If you wanted to move the object to the right 10 pixels, you would change the x value to 10. Similarly, to move it down 10 pixels you would change the y value to 10. The values you enter can be positive or negative but must be whole numbers.

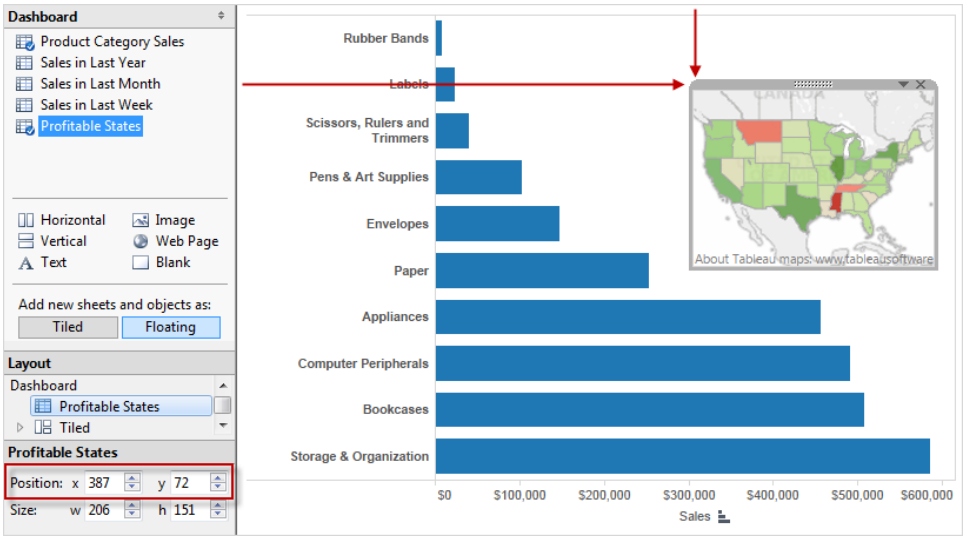

Use the **Size** field at the bottom of the Dashboard window to specify the exact dimensions for floating objects. Define the size in pixels where w = the width of the object and h = the height of the object. You can also resize floating objects by clicking and dragging an edge or corner of the selected object in the dashboard.

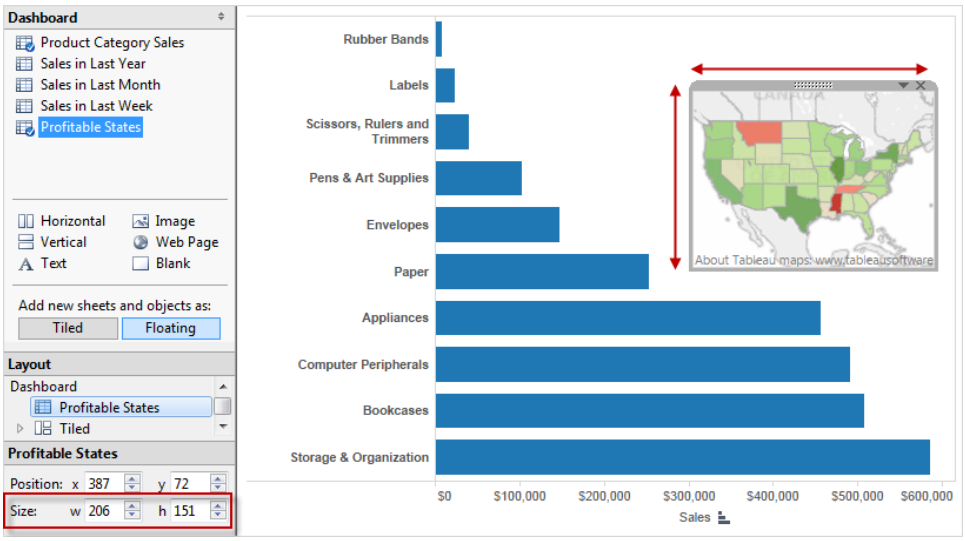

Layout Containers

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 to automatically adjust the width of dashboard objects. Add a Vertical Layout Container to automatically adjust the height of dashboard objects.

1. Drag a horizontal or vertical layout container to the dashboard.

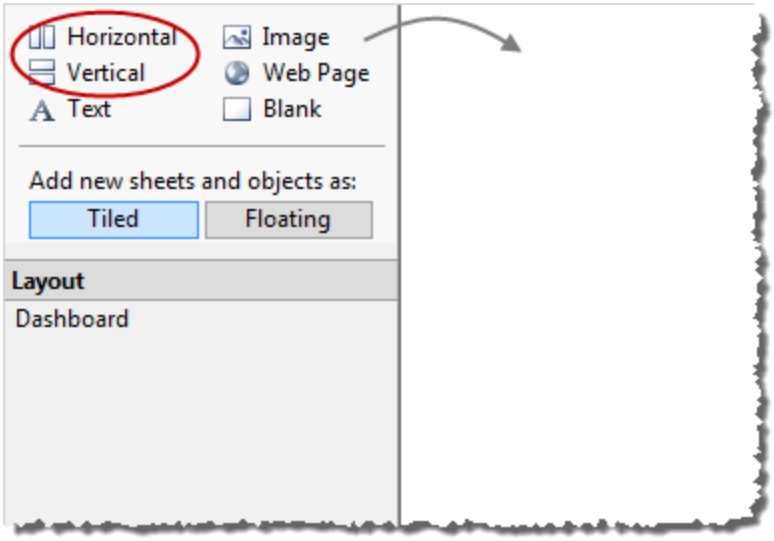

2. Add sheets and objects to the layout container. As you hover over the layout container, a

blue box indicates that the object is being added to the flow of the layout container.

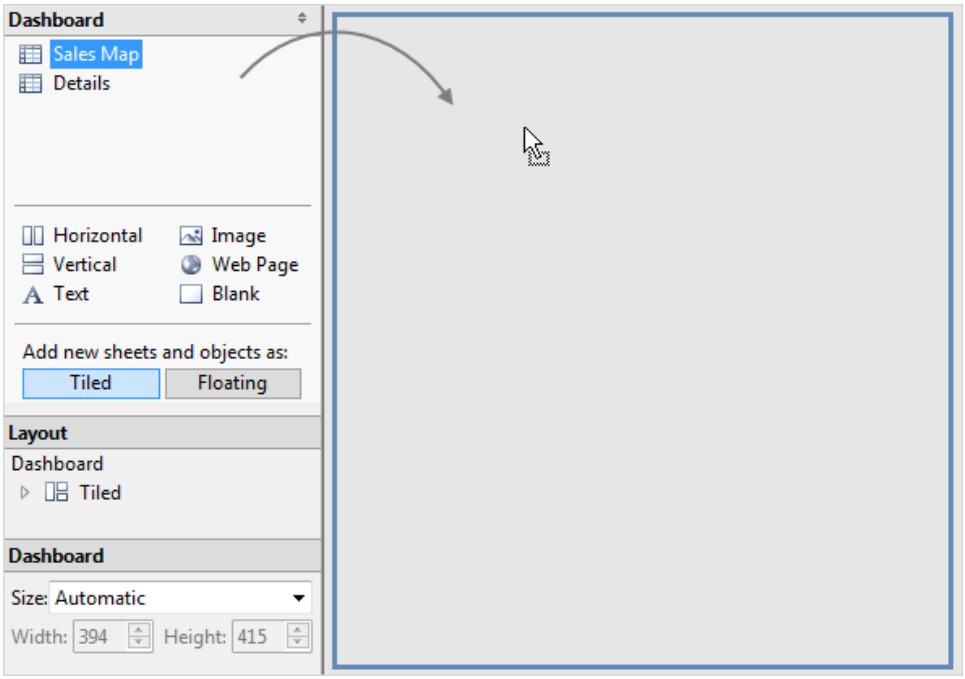

3. Watch as the objects move and resize. For example, in the dashboard below, a filter is applied that causes the text table to get shorter. Because both views are in a vertical layout container, the map automatically resizes into the new space.

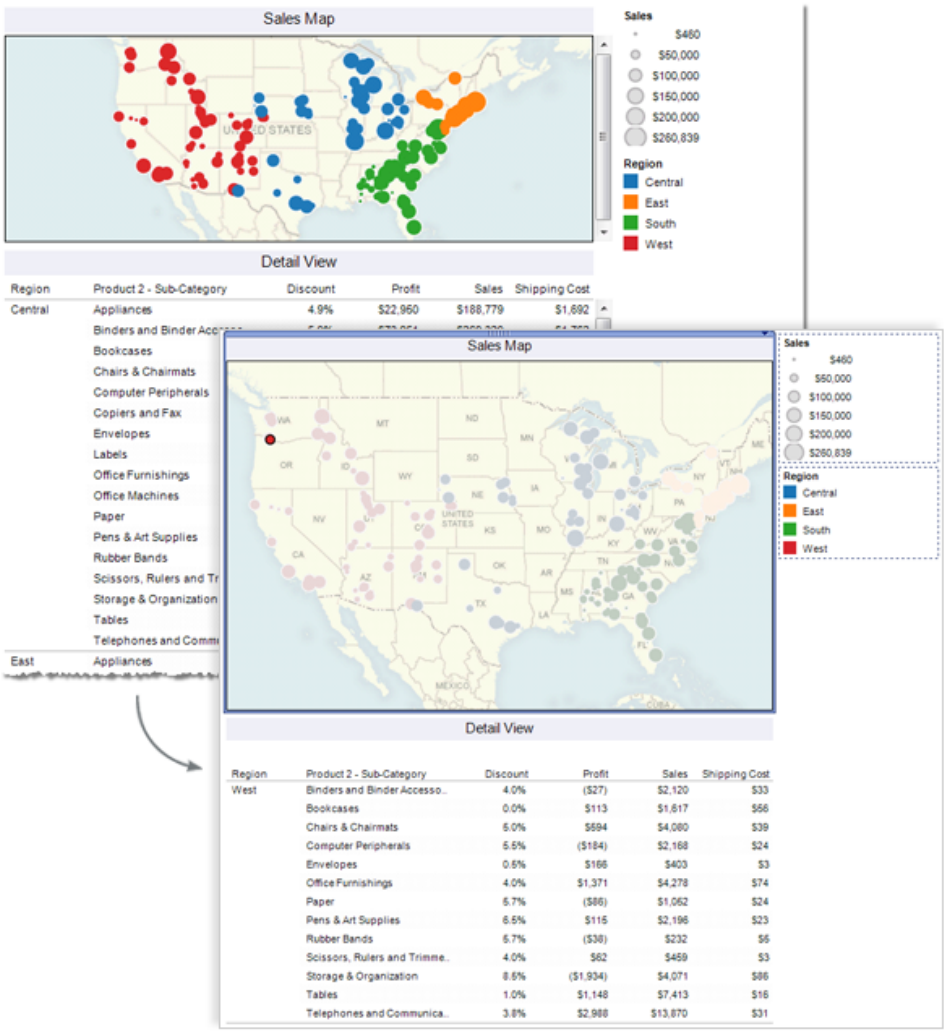

You can add as many layout containers are you want and you can even add layout containers within other containers.

Removing Layout Containers

When you remove a layout container, it and all of its contents are deleted from the dashboard.

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.

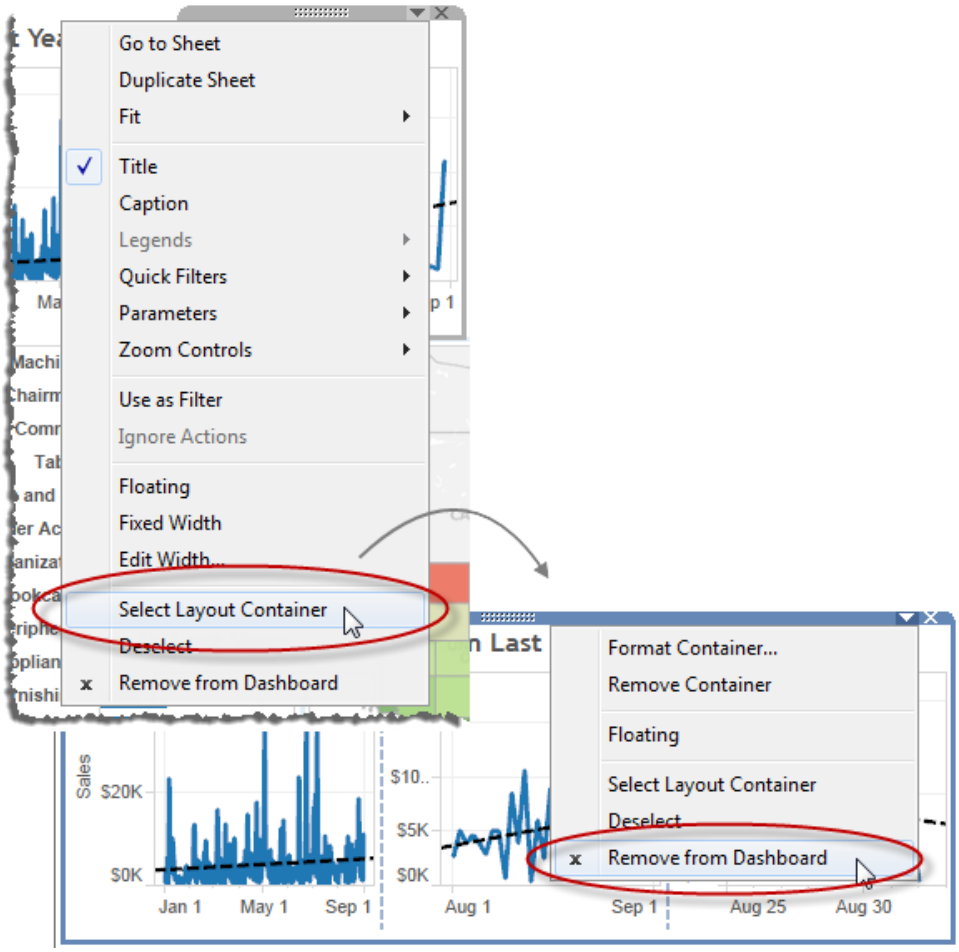

Formatting Layout Containers

You can specify shading and border style for layout containers so that you can visually group objects in the dashboard. By default, layout containers are transparent and have no border style.

1. Open the drop-down menu for the layout container you want to format and select **Format Container**.

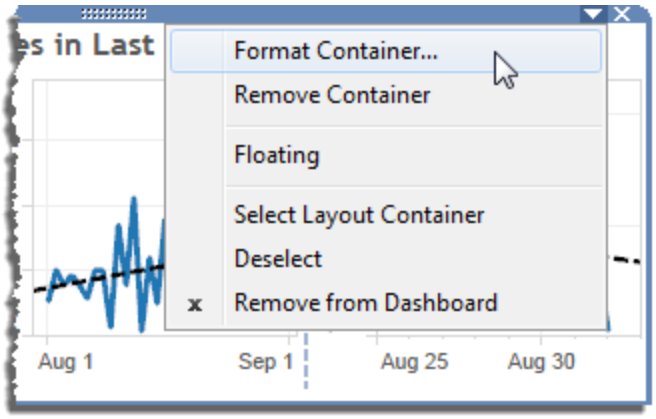

2. In the Format Container window, specify the color and opacity from the Shading control.
3. Specify a line style, thickness, and color for the border from the Border control.

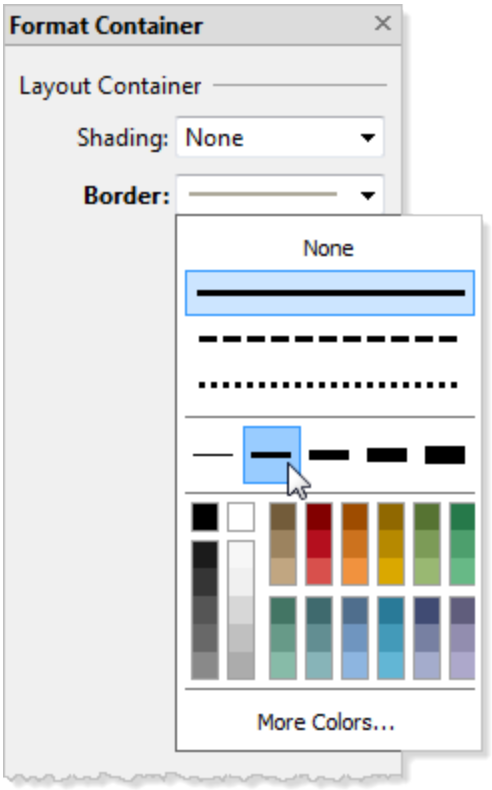

Scaling Layout Containers

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.

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.

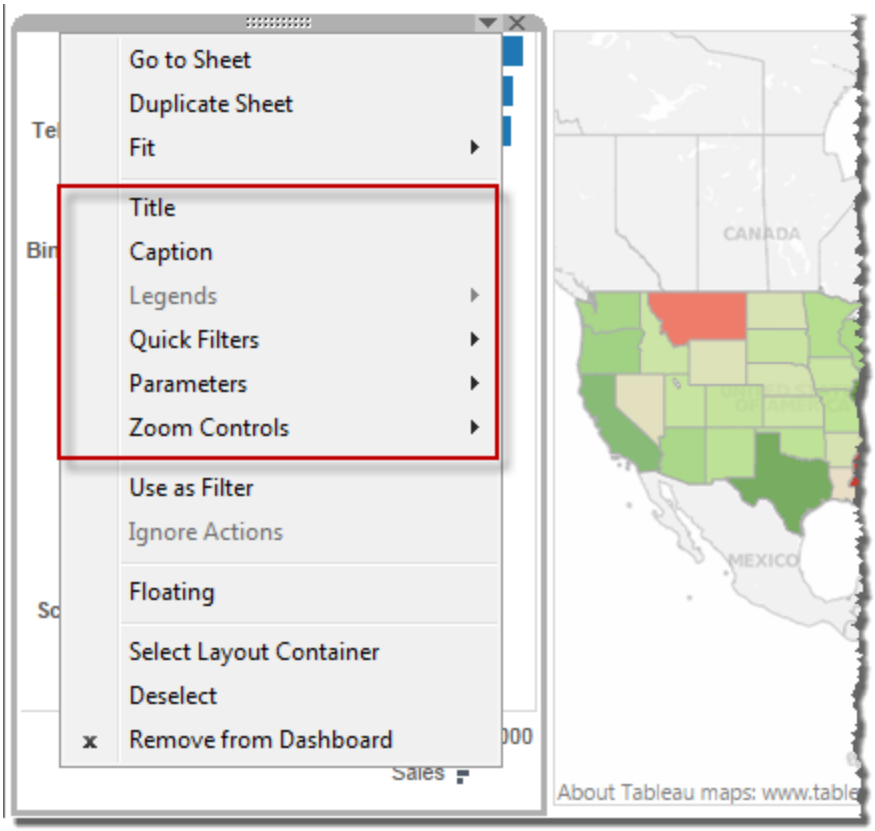

Alternatively, you can right-click an item in the Layout area of the Dashboard window to access all of these same commands.

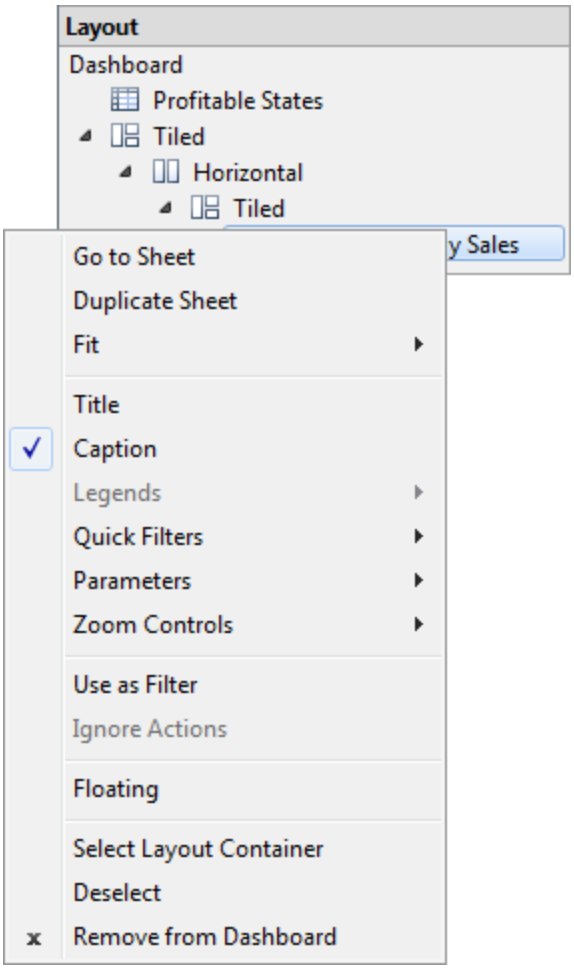

Quick filters are only available for the fields used in the original view.

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.

1. Select the view or object that you want to move.
2. Click and drag the move handle at the top of the selected item to a new location.

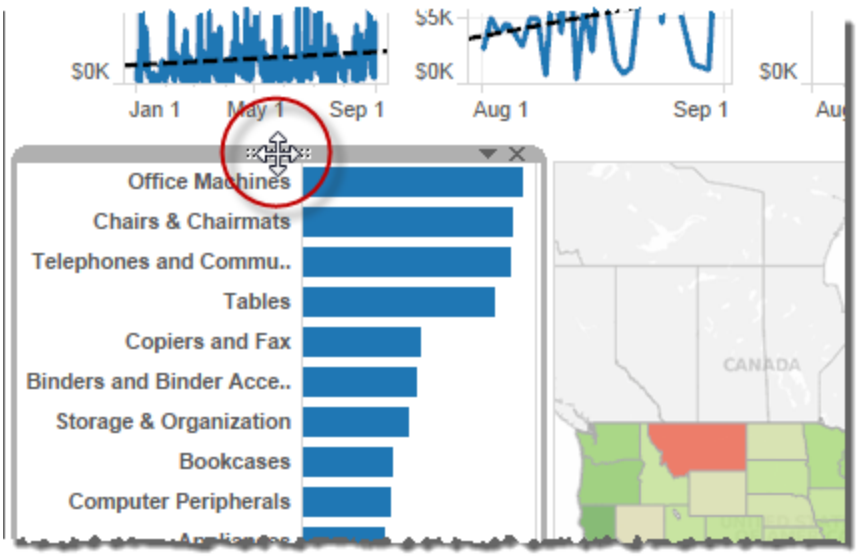

3. Drop the view or object in a new location.

   As you drag objects around the dashboard, the areas where you can drop them are shaded in gray.

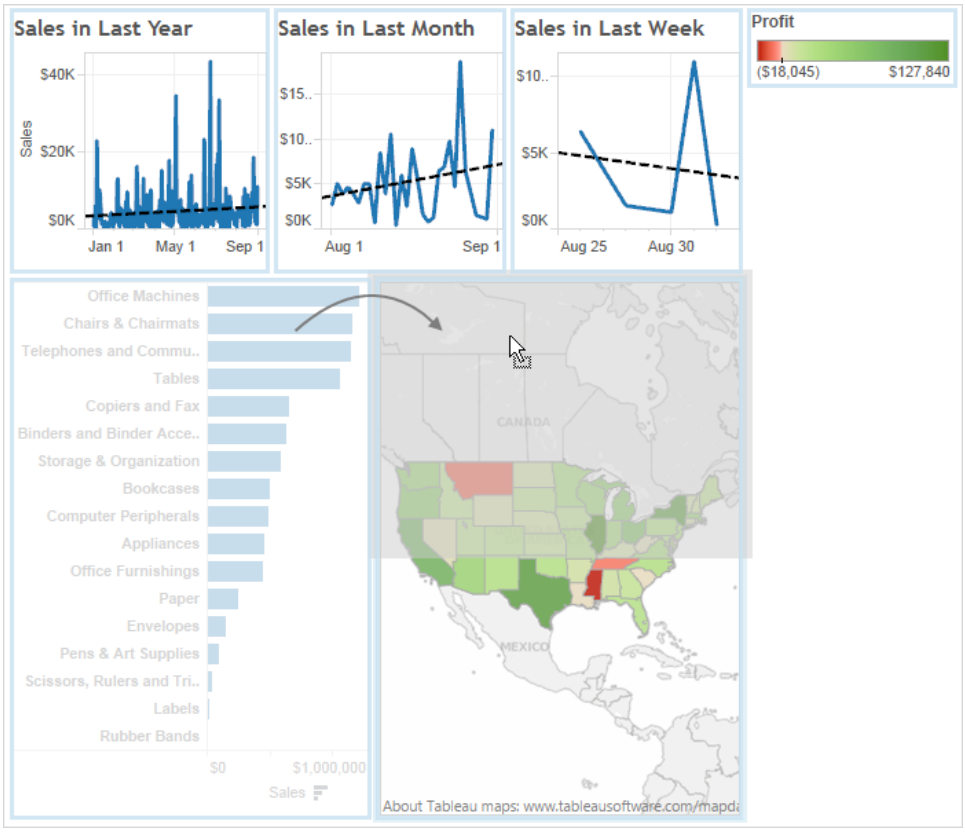

Setting the Dashboard Size

You can specify the overall size of the dashboard using the Dashboard area at the bottom of the Dashboard window. The Dashboard area displays when you've deselected all items in the dashboard. By default, the dashboard is set to the Desktop preset, which is 1000 pixels by 800 pixels. Select a new size using the drop-down menu.

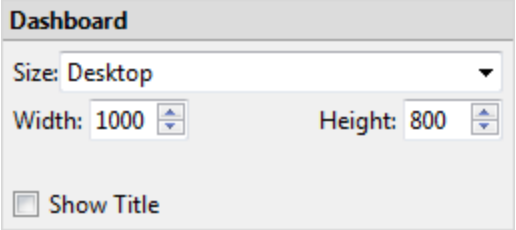

You can select from the following options:

- **Automatic** - the dashboard automatically resizes to fill the application window.
- **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 scroll bars or white space will display.
- **Presets** - select from a variety of fixed size presets such as Letter, Small Blog, and iPad. If you select a preset that is larger than the window, the dashboard becomes scrollable.

Understanding Dashboards and Worksheets

The views in a dashboard are connected to the worksheets they represent. That means when you make changes to the worksheet, the dashboard is updated and subsequently, many 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 and 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 they are not shown in the filmstrip, sheet sorter, or in the tabs along the bottom of the workbook.

Go to Sheet

1. Select the view you want to see full size.
2. Select **Go to Sheet** on the dashboard view menu.

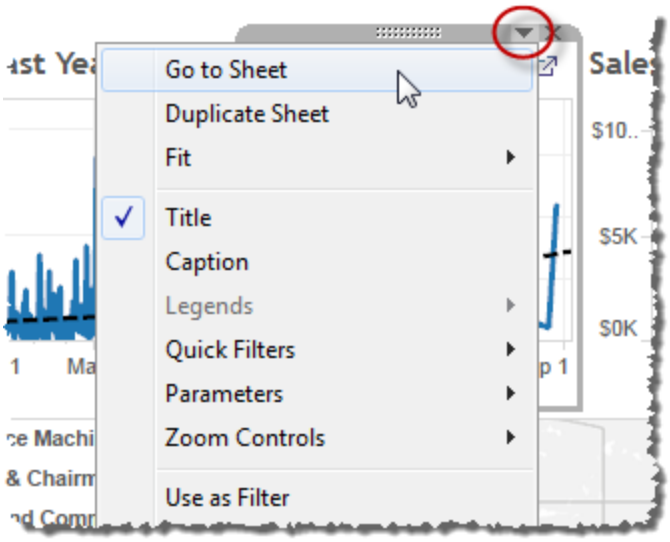

Duplicate Sheet

1. Select the view you want to duplicate.
2. Select **Duplicate Sheet** on the dashboard view menu

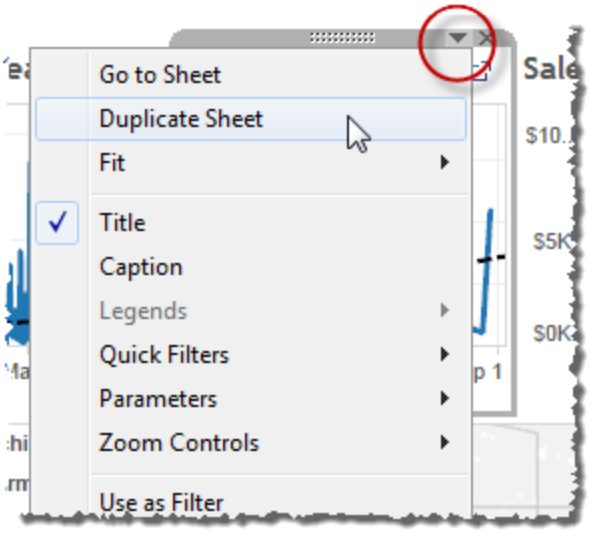

Hide Sheet

Right-click the worksheet tab or thumbnail and select **Hide Sheet**.

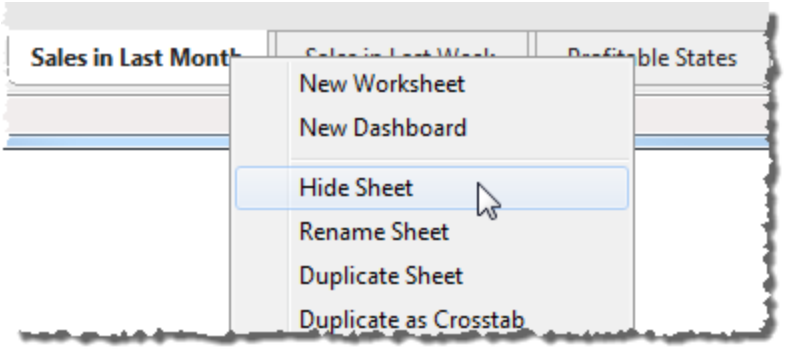

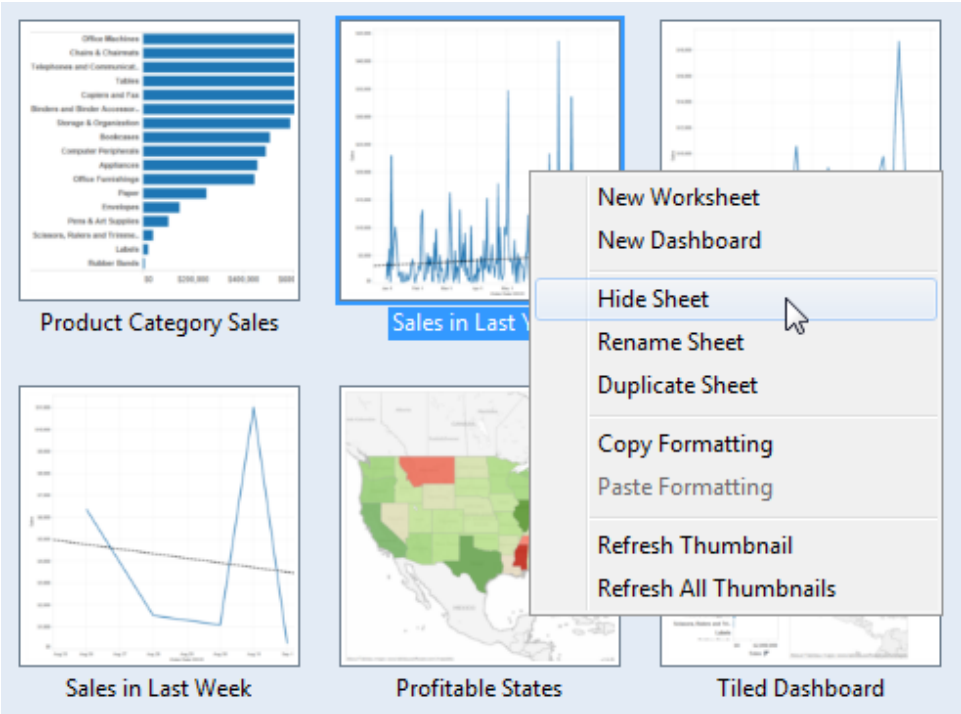

Show Hidden Sheets

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.

Stories

A **story** is a sheet that contains a sequence of views or dashboards that work together to convey information. You can create stories to show how facts are connected, provide context, demonstrate how decisions relate to outcomes, or simply make a compelling case.

A story is a sheet, so the methods you use to create, name, and otherwise manage worksheets and dashboards apply to stories (see **Sheets** on page 18). At the same time, a story is also a collection of sheets, arranged in a sequence. Each individual sheet in a story is called a **story point**.

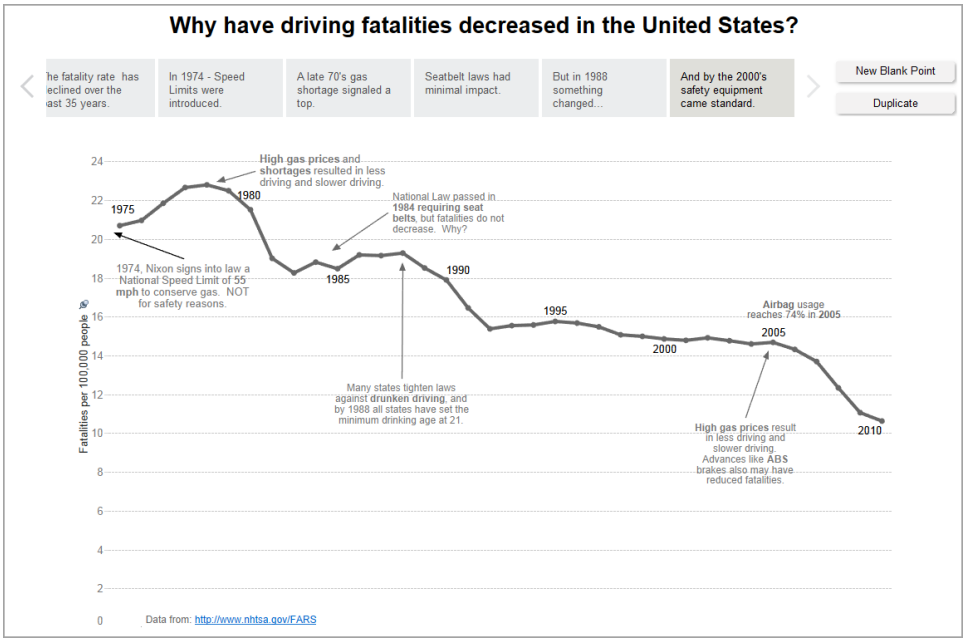

Tableau stories are not collections of static screen captures—your story points remain connected to the underlying data and change as the data changes—or as the views and dashboards that you use in your story change. When you share a story —for example, by emailing a workbook or by publishing a workbook to Tableau Server or Tableau Online—users can interact with the story, to reveal new findings or ask new questions of the data.

There are many different ways to use stories. For example:

- **Use stories for collaborative analysis**: You can use stories to assemble a sequenced analysis, for yourself or for collaboration with your colleagues. Visualize the effects of data changes over time, or perform what-if analysis.
- **Use stories as presentation tools**: You can use stories to present a narrative to an audience. Just as dashboards provide spatial arrangements of views that work together, stories present sequential arrangements of views or dashboards that create a narrative flow for your audience.

There are different ways to build a story. For example, each story point in a story can be based on a different worksheet or dashboard. Conversely, each story point can be based on a single worksheet or dashboard, which you customize for each story point, perhaps adding more information in each new story point. Often you will want to combine these approaches, using new sheets for some story points, and customizing the same sheet for other story points.

Creating a Story

You create a story from existing worksheets and dashboards.

To create a story:

1. Open the workbook that contains the views and dashboards you want to use in your story.

   Alternatively, create a new workbook and create the necessary worksheets and dashboards as you create your story.

2. Choose **Story** > **New Story**.

   Alternatively, click the new sheet button on the toolbar, and then choose **New Story**, or right-click at the bottom of the window, in the sheet tab area, and choose **New Story**:

New story from toolbar	New story from sheet tab area
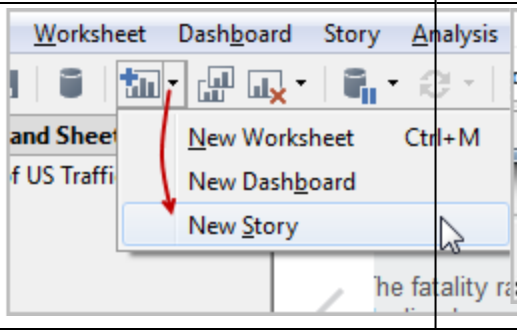	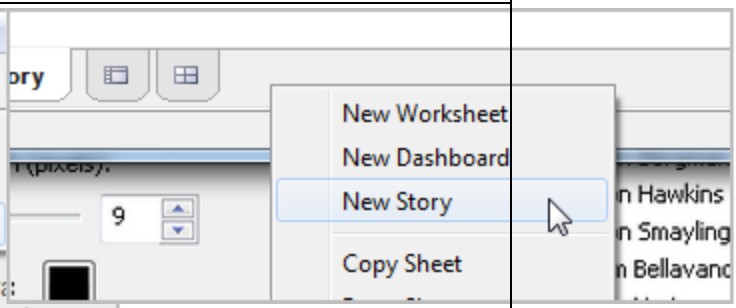

   Tableau opens a new story as your starting point:

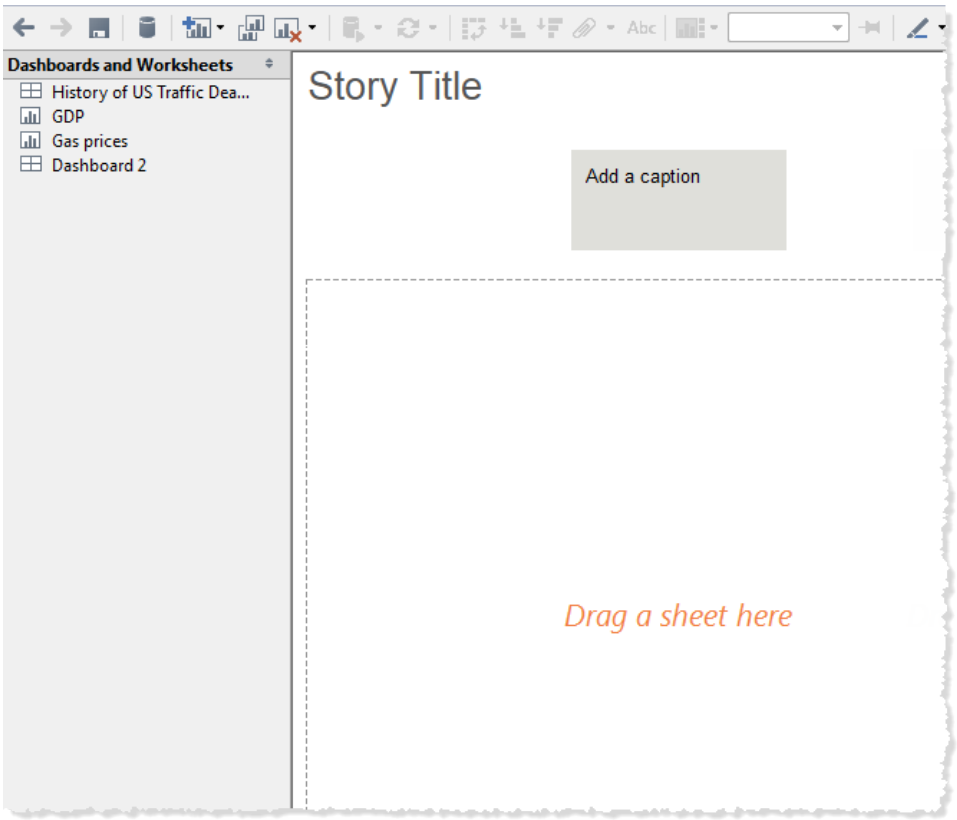

3. In the lower-left corner of the screen, choose a size for your story. Choose from one of the predefined sizes, or set a custom size, in pixels:

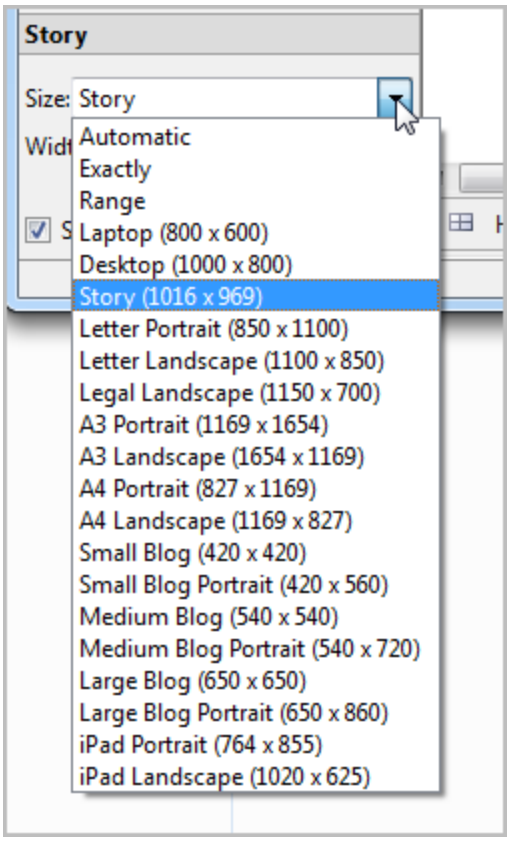

Choose a size with your target platform in mind, rather than on the platform on which you are creating the story.

4. To add a title to your story, double-click the default story title to open the Edit Title dialog box. You can type your title in the dialog box, as well as choose a font, color, and alignment for your title. Click **Apply** to view your changes.
5. Drag a sheet from the **Dashboards and Worksheets** area to the story, and drop it into the center of the view:

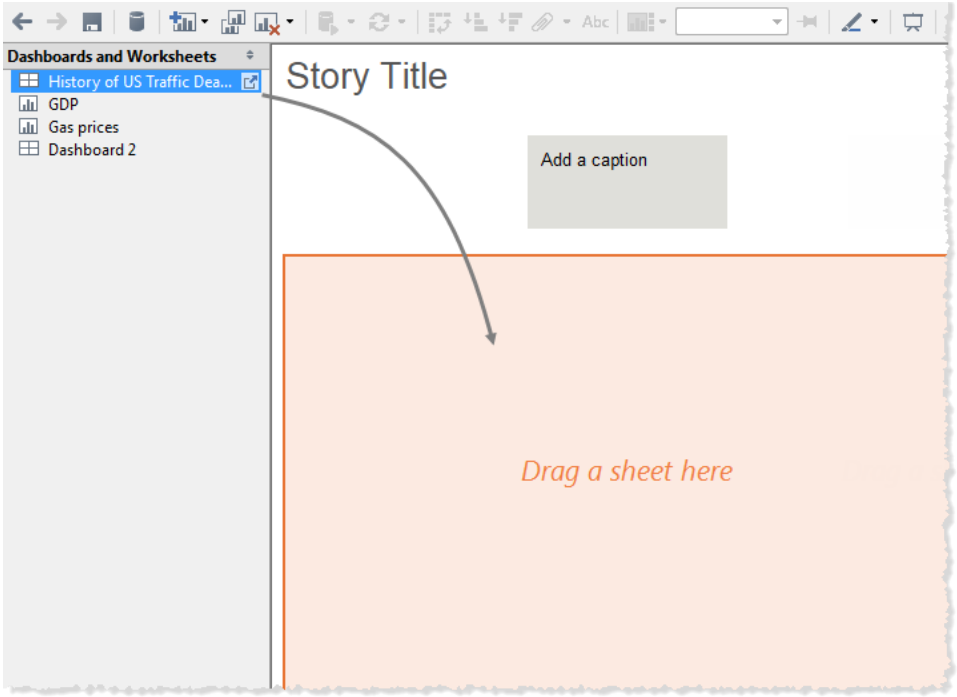

6. Click in the navigator box for the story point (the gray box just below the title that reads *Fill in the caption text*) and type some text to summarize the story point.

   Don't type more text than can fit in the box—you'll be able to add descriptions and annotations within each story point if you want to provide more information.

7. Customize the story point. For more information, see **Customizing a Story Point** on page 812.

8. If you have customized the view, click **Update** above the navigator box to save your changes to the story point:

9. When you're ready to add your next story point, there are different ways to proceed:
   - Click **New Blank Point**:

     Then drag a sheet from the Dashboards and Sheets area and customize the story point as needed (as per steps 6 - 8, above).
   - Click **Duplicate**.

     Do this if you want to use the current story point as the starting point for a new story point. Then customize the view or worksheet in the second story point—make it different in some way from the original story point.
   - Continue performing customizations on the original story point. After you do so, the **Duplicate** button becomes the **Save as a New Point** button. Click **Save as a New Point** to save your customizations as a new story point. The original story point remains unchanged.

     > **Note:**Adding a description box to a story point will not automatically change the **Duplicate** option to **Save as a New Point**. If you want to create a new

story point that differs from the previous story point only in having a new description box, start by duplicating the original story point and then add the description box.

10. Continue adding story points, until your story is complete.

The Story Workspace

As you work on a story, you can use the following controls, elements, and features.

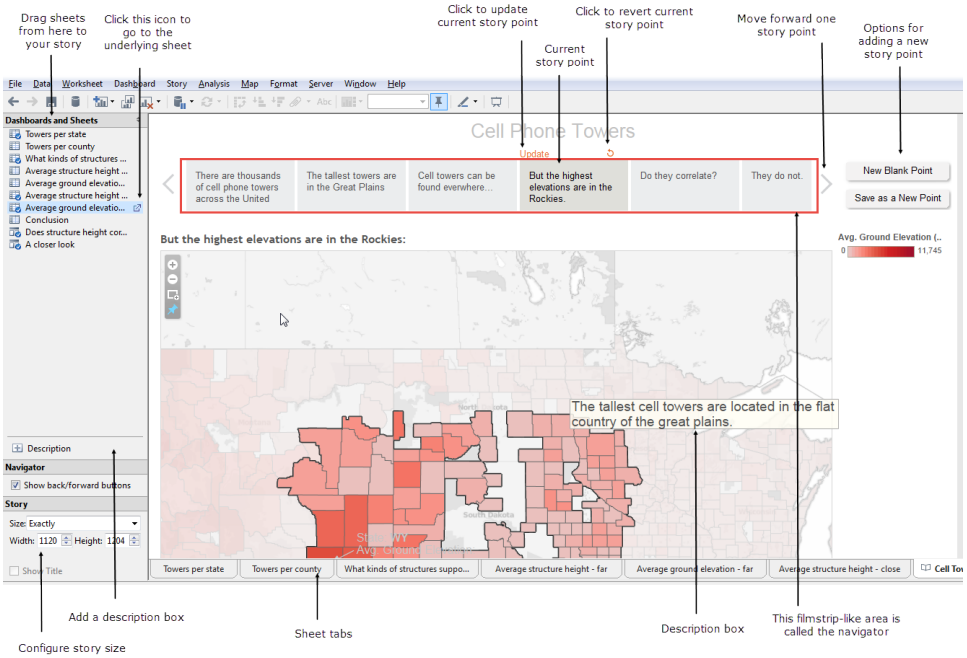

Customizing a Story Point

You can customize a story point in any of the following ways:

- By selecting a range of marks
- By filtering a field in the view
- By sorting on a field in the view
- By zooming in or panning on a map
- By adding a description box
- By adding annotations
- By changing the value of a parameter in the view

- By editing a dashboard text object
- By drilling down or up in a hierarchy in the view

After you modify a story point, you can either save your changes (by clicking **Update** above the story point) or revert the story point to its original state (by clicking the circular arrow):

> **Note:** When you drag a sheet from the **Dashboards and Sheets** section to a story point, that sheet remains connected to the original sheet. So if you modify the original sheet, your changes will be automatically reflected on the story points that use this sheet. But changes you make in a story point do not automatically update the original sheet. So be careful about changing an underlying sheet after customizing a story point—the sheet changes may not be compatible with the customizations you have made to the story point.

Adding a Description

Descriptions are a special type of annotation that you can add to a story point. To add a description, double-click **Description** under Story in the left pane. Creating, editing, or formatting a description is the same as creating, editing, or formatting an annotation (see **Adding Annotations** on page 769). But note the following:

- You can add as many descriptions to a story point as you want.
- Descriptions do not attach to marks, points, or areas in the story point. You can position them wherever you like.
- Descriptions exist only on the story point where you add them. They do not affect the underlying sheet nor any other story point in the story.

After you add a description box, click it to select and position it. When you select a description box, it acquires a frame and a menu. You can open the menu by clicking the inverted triangle in the upper-right of the frame:

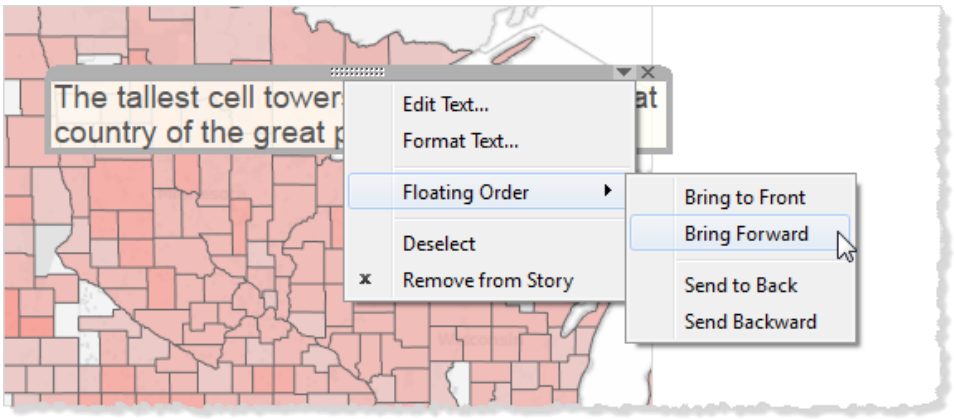

Use the commands on this drop-down menu to update the description, set its floating order relative to any other description boxes that it may overlap, deselect it, or remove it from the story point.

Updating a Story

You can update a story in any of the following ways.

Revising an Existing Story Point

To revise an existing story point, click it in the navigator, and then make changes. See **Customizing a Story Point** on page 812. You can even replace the underlying sheet by dragging a different sheet from the **Dashboards and Worksheets** area onto the story pane.

Deleting a Story Point

To delete a story point, click it in the navigator, and then click the orange **x** just above the box:

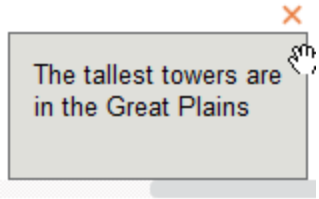

If you delete a story point by accident, click the Undo button to restore it.

Inserting a Story Point

To insert a new story point somewhere other than at the end of a story, add a story point, then drag it to the desired location in the navigator and drop it:

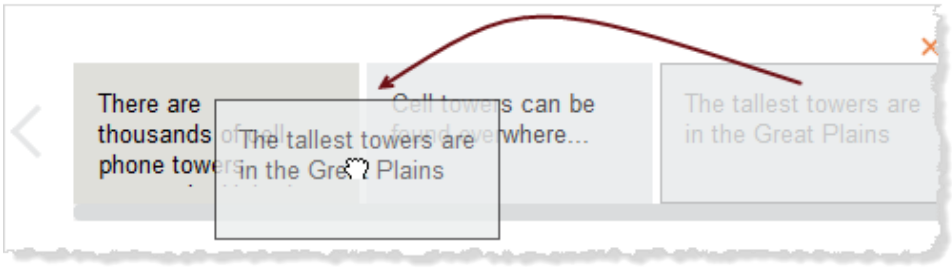

The story point is inserted at the specified location:

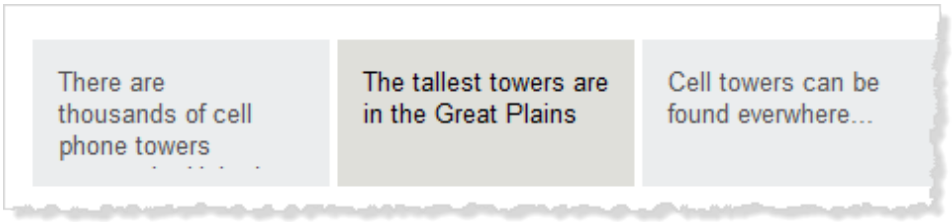

Alternatively, if you are dragging a sheet into the story, you can just drop it between two existing story points in the navigator.

Rearranging Story Points

You can use the navigator to drag and drop story points within a story as needed.

Presenting a Story

Typically you create a story for others to review or experience, either individually, at their desks, or as part of an audience for a presentation.

To present a story, use presentation mode. To switch in and out of presentation mode, click the Presentation Mode button on the toolbar, choose **Window** > **Presentation Mode**, or press F7. To exit Presentation Mode, press Esc or F7. See **Reorganizing the Workspace** on page 15 for more information on presentation mode.

You can also publish a workbook that contains a story to Tableau Server or Tableau Online. User can then open the story and navigate from story point to story point, or interact with stories just as they would interact with views and dashboards. Web users cannot, however, author stories, or permanently modify published stories.

Example – Earthquake Trend Story

This example shows how you can use Tableau to tell a story with data. The story described below considers the question: Are serious earthquakes becoming more common, or is it just that some strong earthquakes in recent years are creating that impression?

This story is available on Tableau Public: https://public.tableausoftware.com/views/EarthquakeTrendStory/Earthquakestory?:showVizHome=no#1.

The story was built from three simple dashboards:

- A map of earthquakes recorded around the world since 2004.
- A line chart showing the number of earthquakes recorded, by year, since 1973 (timeline).
- A different version of same timeline visualization, with earthquakes broken out by region, and with trend lines added.

Story Point 1

The first story point shows the map with all data included—all earthquakes, across the entire planet.

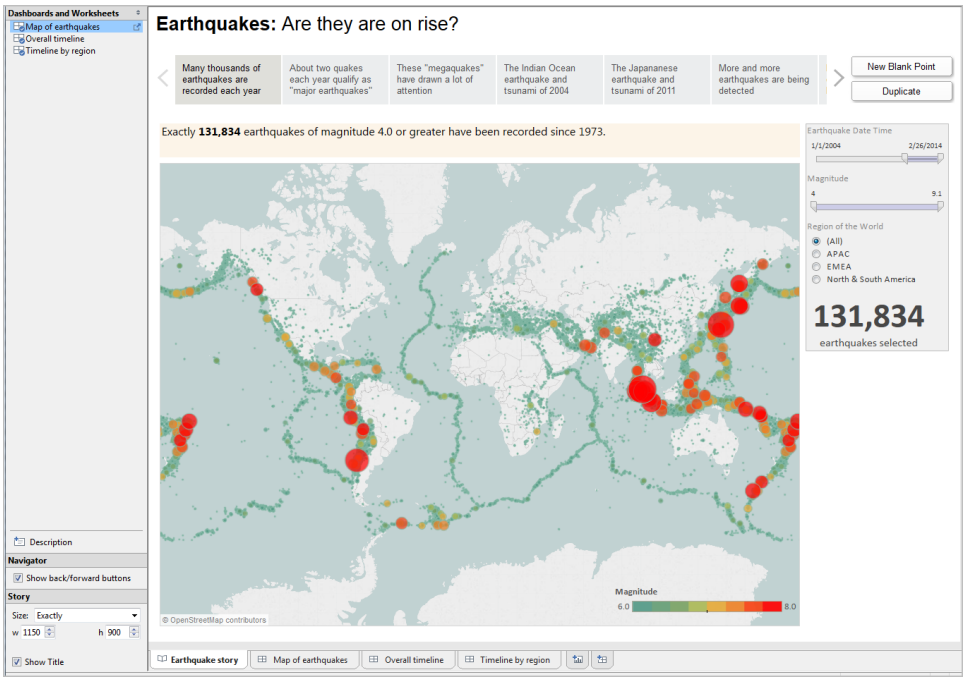

Story Point 2

The second story point compresses the time frame and filters out smaller earthquakes—everything below magnitude 7. The map pans to show the Pacific "Ring of Fire," where the majority of the large earthquakes occurred.

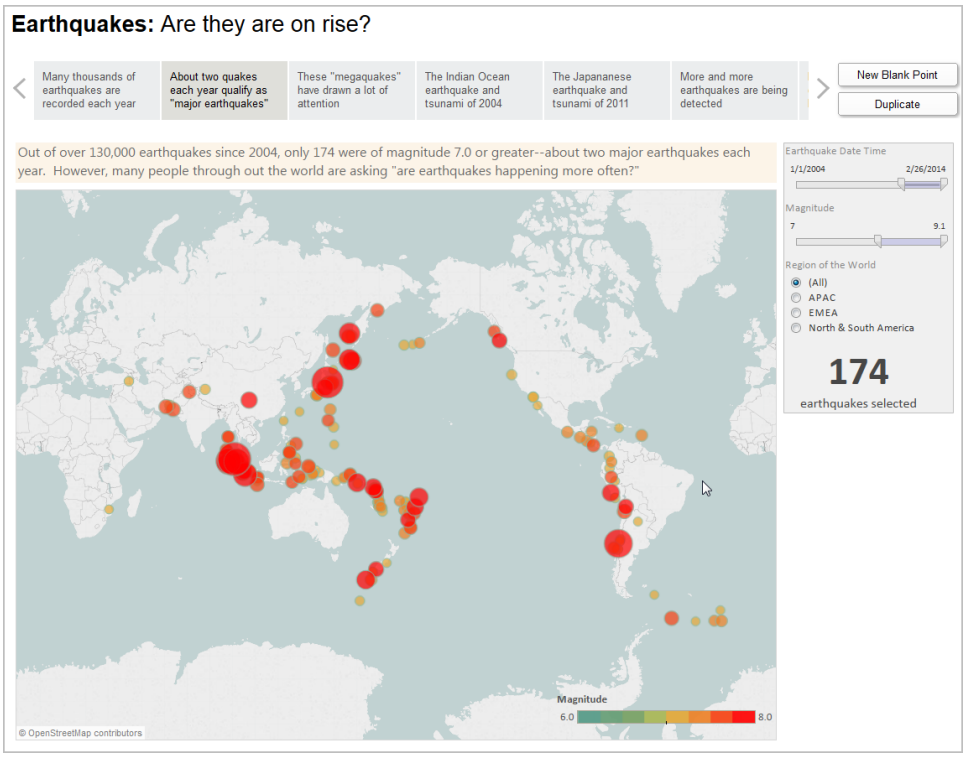

Story Point 3

In story point 3, the magnitude filter is adjusted to show only "megaquakes" (megathrust earthquakes of magnitude 8 or greater). The words work with the pictures to move the narrative forward.

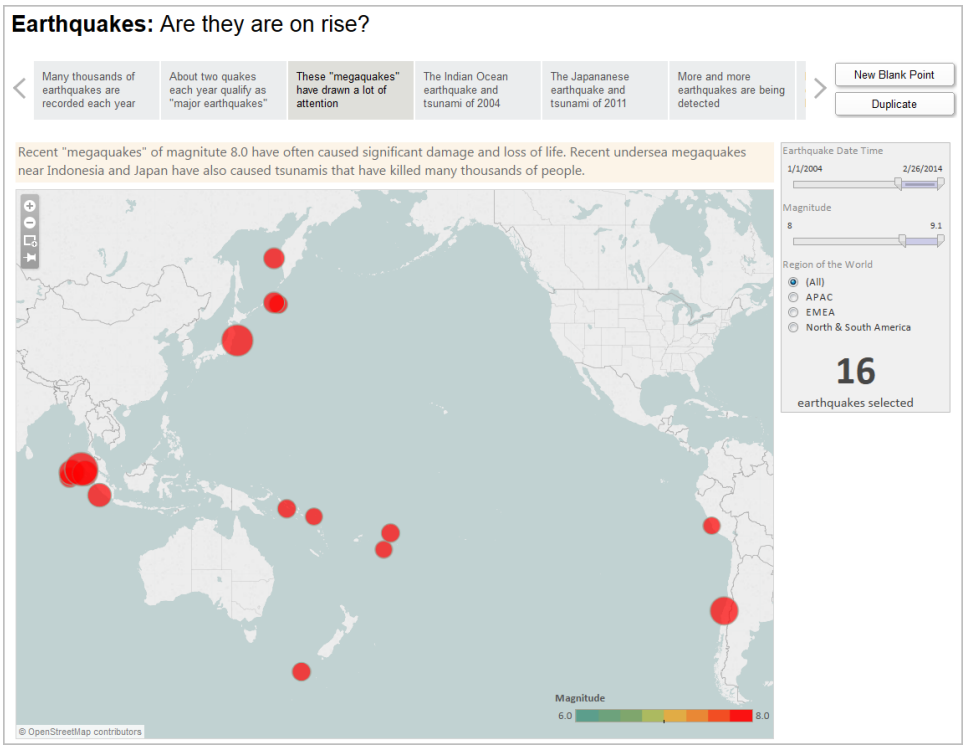

Story Points 4 and 5

Story points 4 and 5 revisit the two most deadly earthquakes in recent history: the 2004 Indian Ocean earthquake and tsunami, and the 2001 Japanese earthquake and tsunami.

The first five story points have all used the same underlying dashboard—the author is creating a compelling visual story just by filtering the data and zooming and panning the map.

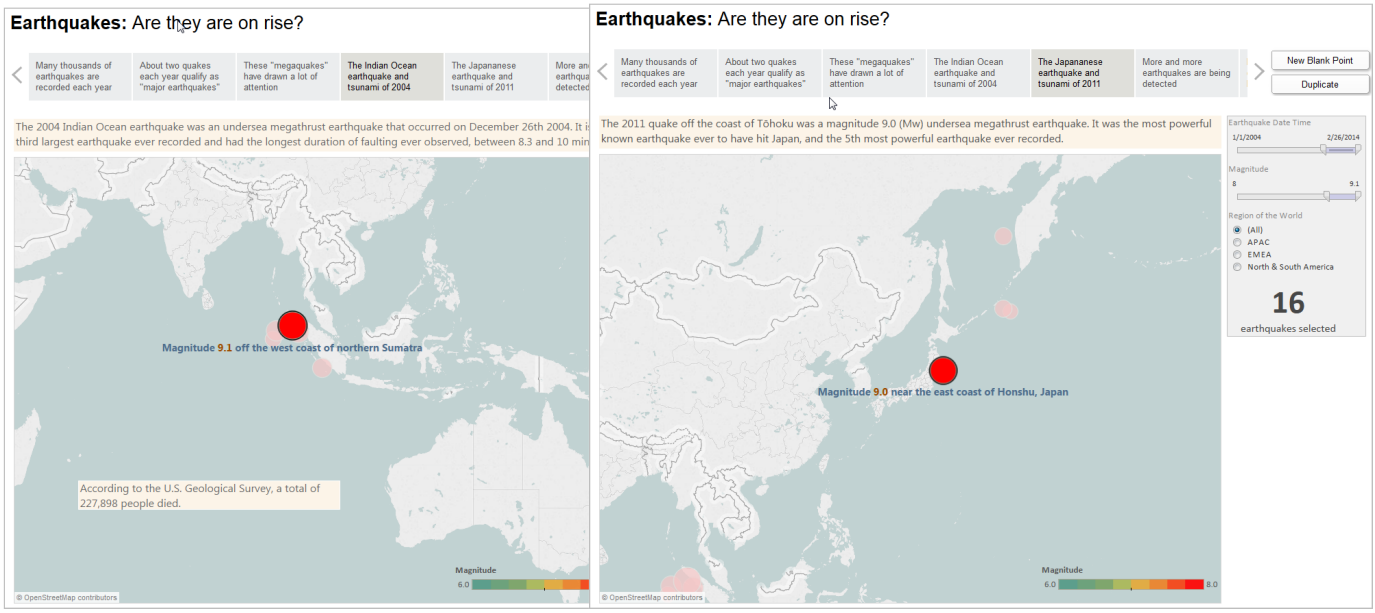

Story Point 6

Story point 6 uses a line chart to show that more quakes are being reported over time since 1973—a three-fold increase.

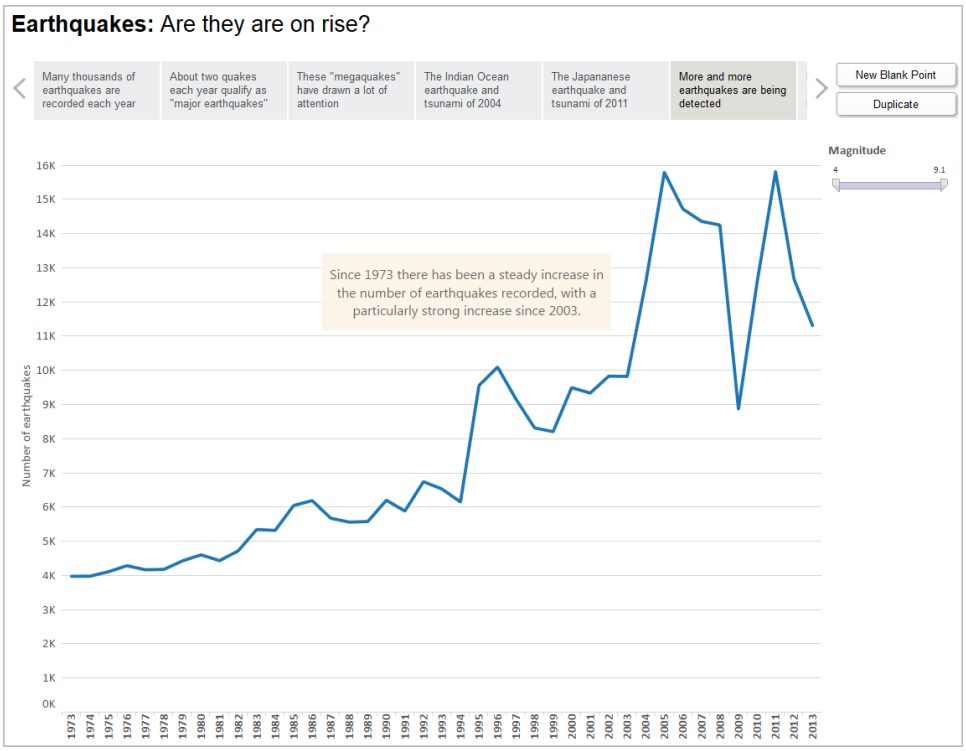

Story Point 7

Story point 7 adds two kinds of detail to the simple line chart from story point 6: it breaks out earthquakes by region, and it adds trend lines, which reduce the variability in the data.

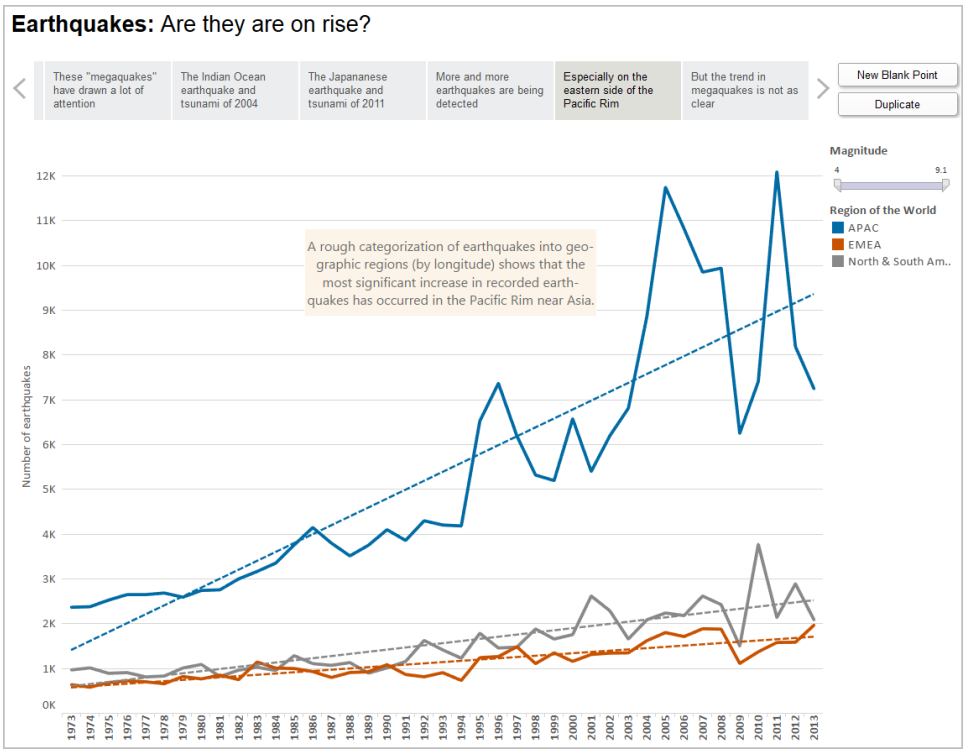

Specifically, the author has duplicated his original worksheet and then dragged region to color and added a trend line. Here is what the underlying worksheet looks like:

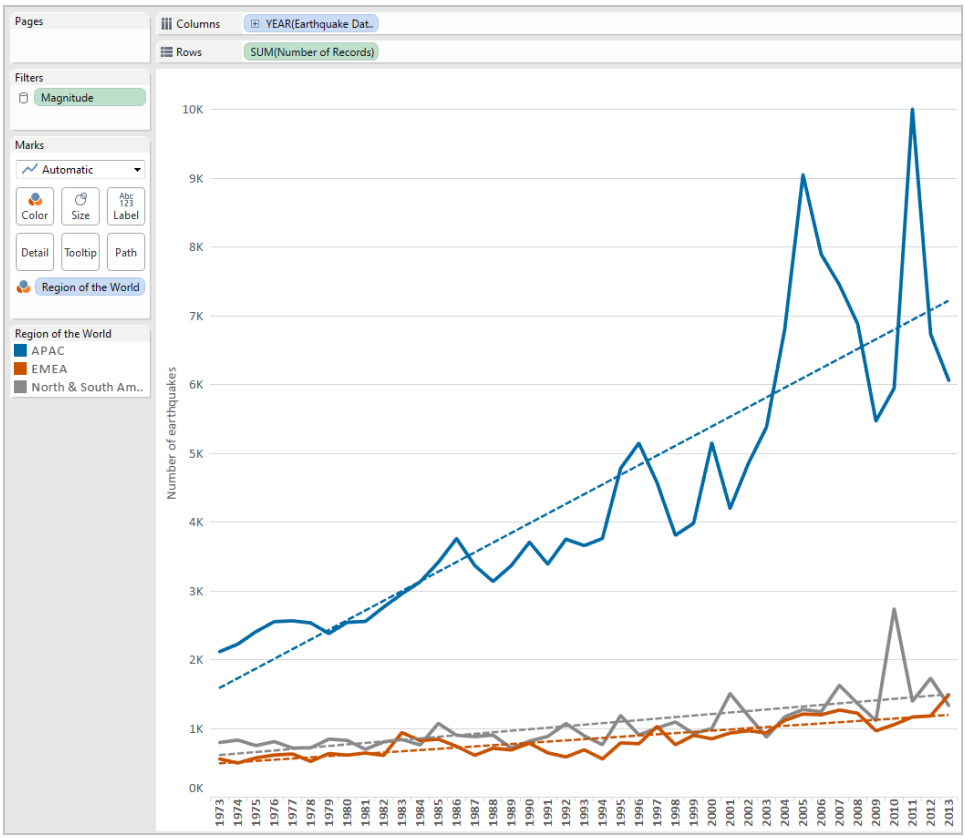

Story Point 8

The final story point filters out weaker earthquakes and shows that what appeared to be a very clear trend is maybe not so clear. Yes, more earthquakes have been reported in recent years, especially in the Asia-Pacific region, but could that just be natural variation? Can we be confident that the trend will continue into the future?

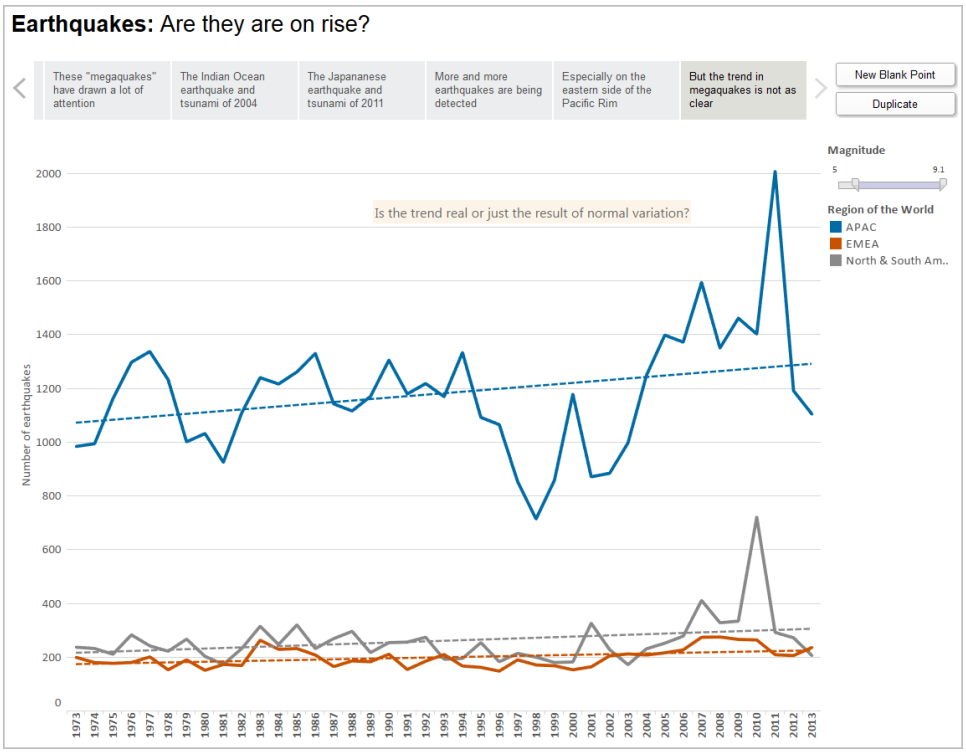

Publishing to Tableau Server

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 that others have published to the server.

- **Tableau Professional**: You must be added as a user on Tableau Server and be given publishing permission to publish views and data sources using Tableau Professional.
- **Tableau Personal**: Tableau Personal users can only publish to Tableau Public.

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.

> See the Tableau Server Online Help 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 a workbook. When you publish a workbook you can add it to a server project, hide some of the sheets, add tags to enhance searchability, specify permissions to regulate access to the workbook on the server, and choose to embed database passwords for automatic authentication on the web.

How to Publish Workbooks to a Tableau Server

After you create a workbook, you can publish to Tableau Server or Tableau Online by following the steps below.

1. Select **Server** > **Publish Workbook**.

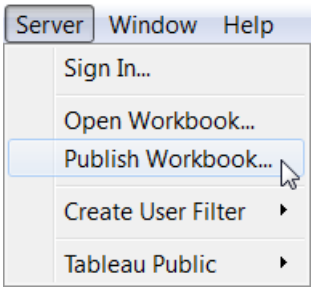

   If you are not already signed in to Tableau Server, you will see the Tableau Server Sign In dialog box.

2. Enter the server name or URL (for example, sales_server, or https://sales_server).

   To sign in to Tableau Online, enter `https://online.tableausoftware.com`.

   Click **Connect**.

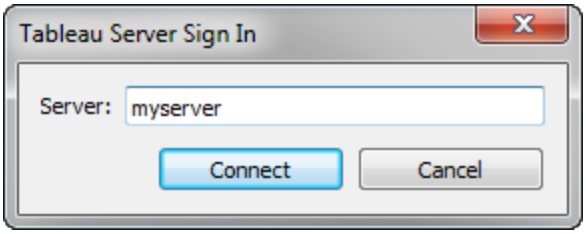

   If Tableau Server is Kerberos-enabled and you are on a machine with valid Active Directory credentials, you will connect automatically to the server. In this case, you can skip step 3. If Kerberos authentication fails, you will be prompted to provide a user name and password for authentication. If you need to change the user you're signed in as, see **Switch Kerberos User** on page 827.

3. Next, enter your user name for Tableau Server on-premise or email address for Tableau Online, enter your password, and then click **Sign In**.

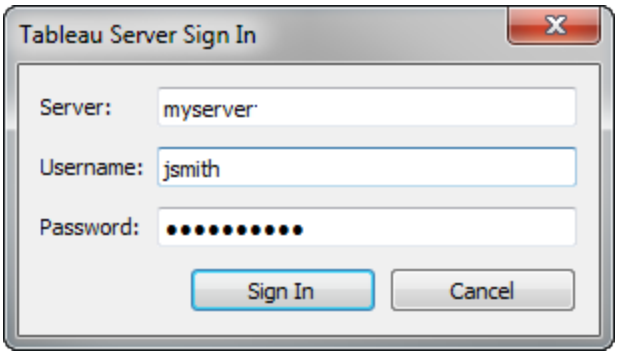

If Tableau Server on-premise is configured to use Active Directory, enter your Windows user name (the domain is not required—except in multi-domain environments where the user is not in the default domain); otherwise, enter your Tableau Server user name.

If Tableau Server is configured to use SAML for user authentication, you won't see the above dialog box. Instead, you'll see a sign in prompt from an external identity provider.

4. In the Publish Workbook to Tableau Server dialog box, specify the following:
   - **Project**: A project is like a folder that can contain workbooks and data sources. The default project on Tableau Server is named Default. All workbooks must be published into a project.
   - **Name**: Provide a name for the workbook in the **Name** text box.

     You can use the drop-down list to select an existing workbook on the server. In this case, the workbook on the server is overwritten. To do this, you must have the **Write** permission for the workbook on the server.
   - **Add Tags**: You can enter one or more keywords that describe the workbook into the Tags text box. Tags help users find related workbooks when browsing the server.

     Multiple tags should be separated using either a comma or a space. If any of the

tags contains a space, use quotations (e.g., “Sales Quotes”).

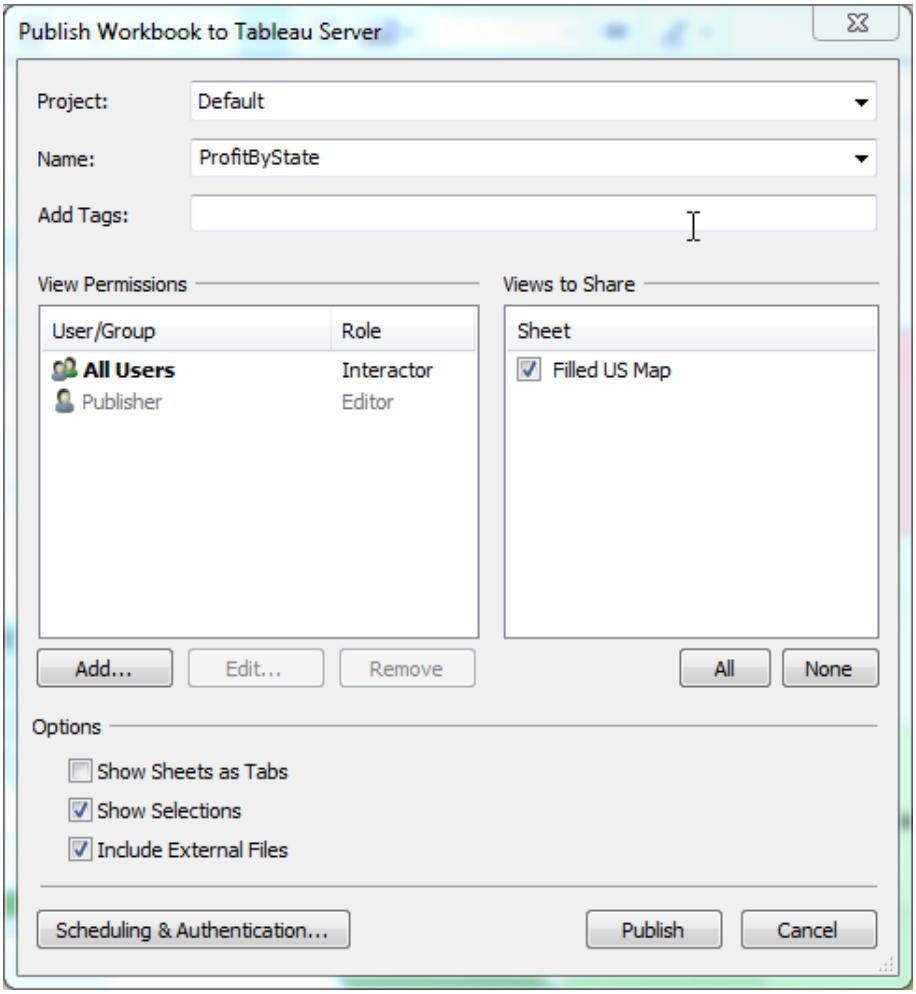

5. You can specify permissions to allow or deny access to the workbook on the server.

   By default all users can interact with the workbook and you, as the publisher, have all permissions. See **Setting Permissions** on page 836 for details, and see **Permissions** on page 839 for information on the different permissions you can assign to users and groups.

6. Under **Views to Share**, select the sheets you want to share on Tableau Server. Any sheets that are not selected are hidden on the server.

   Showing and hiding worksheets is useful when you want to publish a dashboard or story without publishing the worksheets that were used to create that dashboard or story. 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. Remember though, that anyone allowed the **Download/Web Save As** permission can download the workbook from the server and then access the hidden worksheets. So hiding worksheets is not a secure way of concealing information.

Hidden sheets can still be accessed when a workbook is opened from the server using Tableau Professional. Users must be given the **Download/Web Save As** permission to open the workbook from the server. Hidden sheets can also be used by users with appropriate credentials from edit mode . See Edit and Create Views in the Tableau Server help.

7. Select whether to **Show Sheets as Tabs**.

   When this option is selected, all workbook sheets are displayed on a row of tabs in each individual view. Users can then click to see the other views in the workbook.

   > When you select this option, all the sheets are shown regardless of any sheet level permissions you may have already set up. If you want to enforce sheet level security, you should make sure the **Show Sheets as Tabs** option is cleared.

8. Select whether to **Show Selections**. When this option is selected, any selections you've made in the workbook will be published to the server.
9. Select whether to **Include External Files**. When you include external files, a copy of any referenced external file data sources or background images are published along with the workbook. External files include Excel, Access, Text, Data Extract, and image files. If you don't include these files, others may not be able to see the worksheets online.

   If you are publishing to an on-premise Tableau Server, and the workbook references data sources or images on a mapped drive, be sure to check **Include External Files**. If you do not want to publish the external files to the server, for example because you want to preserve a live connection to a file-based data source after publishing to Tableau 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.

   If you are publishing to Tableau Online, and the workbook connects to a data source other than Google BigQuery or Amazon Redshift, you will need to include external files.

10. Optionally, click **Scheduling and Authentication**. Decide whether to embed user names and passwords, so that server users won't have to have an account on the database to see the views. For information about embedding passwords, see **Embedding Passwords and Authentication** on page 829.

    If the workbook uses Extract connections you can optionally add the workbook to a refresh schedule.

11. Click **Publish**.

> When publishing a single worksheet in a workbook containing multiple data sources, each data source, regardless of whether or not it is active in the worksheet, may require authentication.

Switch Kerberos User

If you're in a Kerberos environment and you need to change the user you're signed in as on Tableau Server (typically for testing purposes), on the Server menu, click **Switch User**.

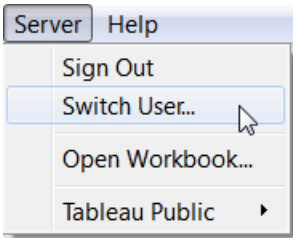

> If the **Switch User** option is not displayed on the server menu, the Tableau Server you are connected to may not be configured for Kerberos single sign-on (SSO).

Selecting **Switch User** opens the Tableau Server Sign In dialog box. You can sign in with a different user name and password. The Tableau Desktop status bar will reflect the user name change.

If you sign in as a different user and then need to sign back in as yourself, on the Server menu, click **Switch to Self**.

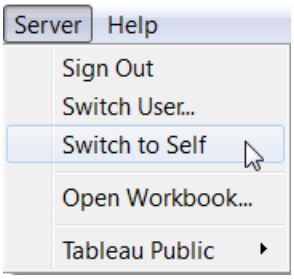

Saving Workbooks to Tableau Public

In addition to saving to an internal server, you can also save workbooks to Tableau Public, a free service that lets you publish interactive data to the web. Anyone can interact with data, download it, or create their own visualizations using Tableau Public. For details, see the Tableau Public web site. Workbooks you save to Tableau Public must not exceed 1 million rows of data.

You cannot publish a live connection with a workbook when you publish to Tableau Public. Tableau will create an extract of your data source automatically. See **Extracting Data** on page 173 to learn about extracts.

> Follow the same steps to publish workbooks to Tableau Public Premium. Make sure to use your Tableau Public Premium username and password when you sign in.

To save workbooks to Tableau Public:

1. Choose **Server** > **Tableau Public** > **Save to Web**.

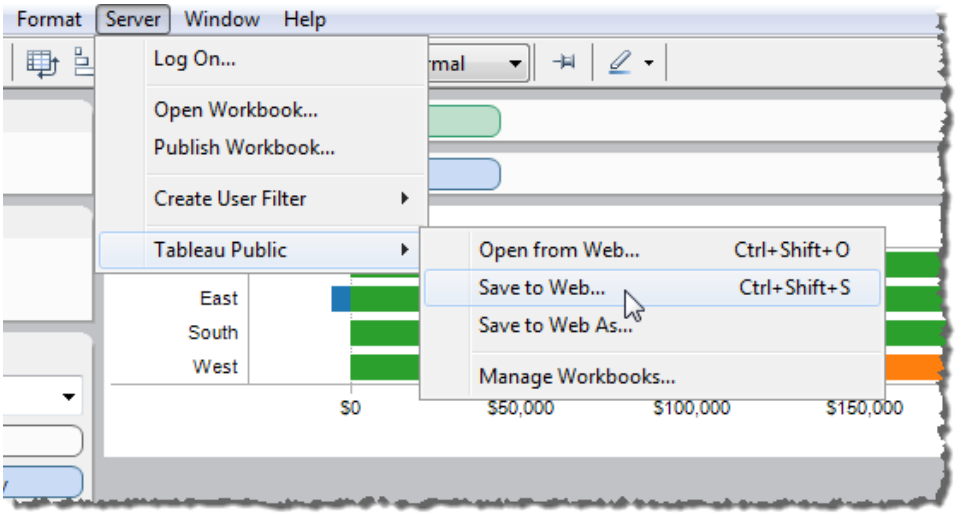

2. When prompted, sign in using your Tableau Public account. If you don't have one, you can create one for free.

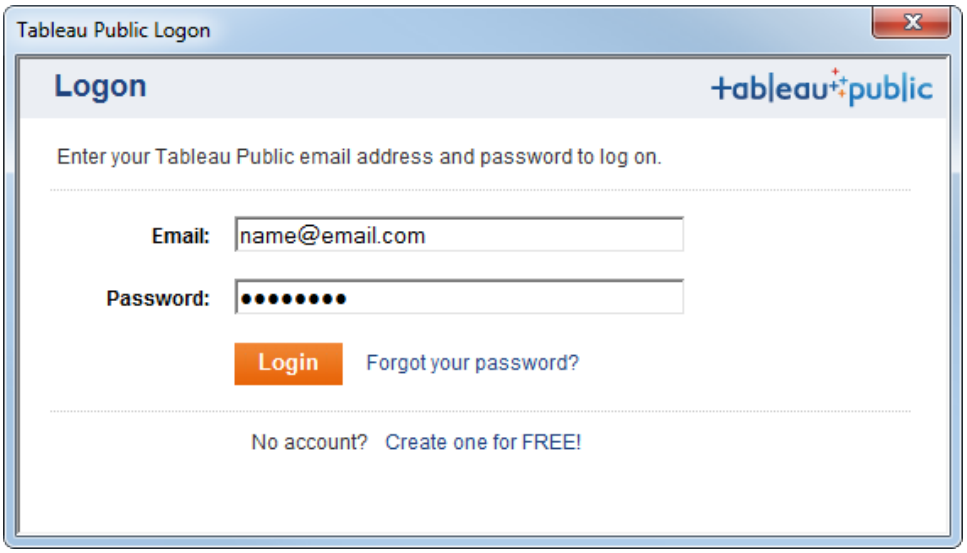

3. The published workbook displays so you can preview all of the saved sheets. Select a sheet and click the **Share** button in the lower left corner of the view to get a link you can email or embed into a web page.

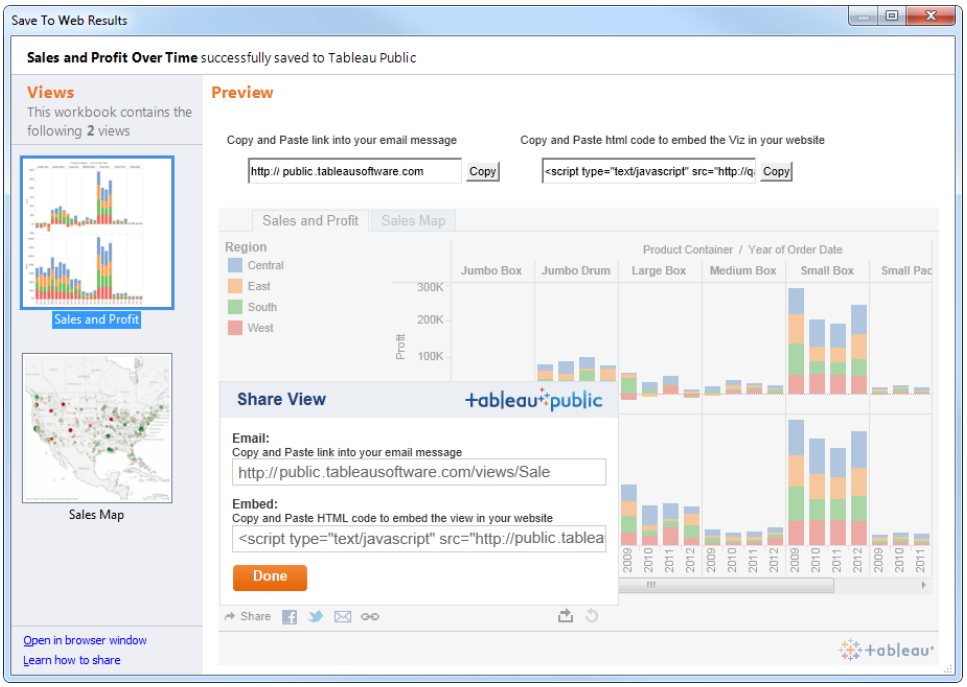

Workbooks and the underlying data saved to Tableau Public are publicly available. You can manage your content saved to the web online at http://public.tableausoftware.com.

Learn more about Tableau Public online at www.tableausoftware.com/public.

Embedding Passwords and Authentication

If your published workbook connects to data that requires a user name and password (credentials) for access, people viewing the workbook on the server are prompted for those credentials. As the workbook author, you can choose to provide a database user name and password that is always used when a view is opened on the server. Server users then bypass the sign-in step and are taken directly to the workbook.

> **Note** If your workbook connects to Salesforce.com, Google Analytics, or Google BigQuery data see **Using OAuth for Data Connections** on page 55 for more information about embedding credentials.

To embed passwords in a published workbook

1. In the Publish Workbook dialog box, click the **Scheduling & Authentication** button in the bottom left corner to open the Scheduling & Authentication dialog box.

In the Scheduling & Authentication dialog box, data source connections that require a user name and password are listed, along with your user name.

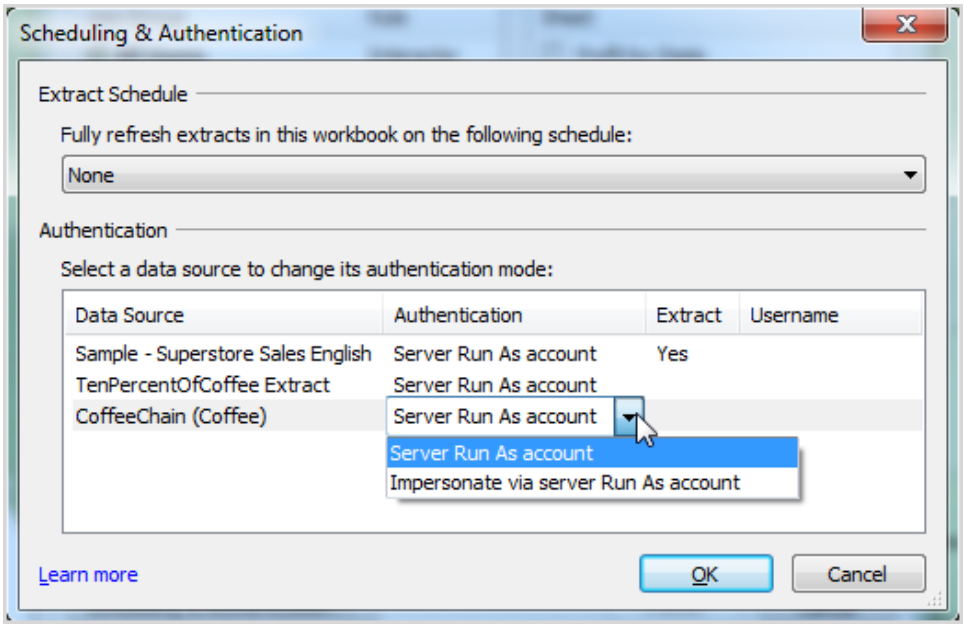

2. Select the connections for which you want to embed passwords, and then select an authentication type. Depending on the data source, one or more of the following options is available:

   - **Prompt user** - visitors on Tableau Server will be prompted to enter their own database username and password.
   - **Embedded Password** - your current user name and password will be used to automatically authenticate Tableau Server users when they load the view.
   - **Server Run As account** - the Server Run As account will be used to automatically authenticate Tableau Server users when they load the view. The Server Run As account is configured by your Tableau Server system administrator.
   - **Viewer Credentials** - in a Kerberos environment, the user's domain user name and password will be used to automatically authenticate Tableau Server users when they load the view.

     If you're using Impala with Cloudera Hadoop and Tableau Server is configured for single sign-on, select **Viewer Credentials** to automatically authenticate Tableau Server users when they load the view.
   - **Impersonate via server Run As account** - the Tableau Server Run As account will be used on behalf of a Tableau Server user visiting the view. This option is only available for Microsoft SQL Server data sources that have been configured for impersonation. See the Reference section of the Tableau Server Online Help.
   - **Impersonate via embedded password** - your username and password will be

used to authenticate Tableau Server users when they load the view. Tableau Server can be running under any type of account, but it will use these credentials, supplied by the publisher, to connect to the database. This option is only available for Microsoft SQL Server data sources that have been configured for impersonation. See the Reference section of the Tableau Server Online Help.

Whether you are allowed to embed passwords or impersonate is controlled by the Tableau Server administrator. Administrators can allow authors to embed passwords with the Settings on the Maintenance page of the server.

3. Click **OK** to close the Scheduling & Authentication dialog box.

Scheduling Extract Updates

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 avoids queries to the live database. Then you can add that workbook to a schedule so that the extract is refreshed at regular intervals with 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.

1. As you are publishing a workbook, in the Publish Workbook to Tableau Server dialog box, click **Scheduling & Authentication**.
2. In the Scheduling & Authentication dialog box, select a schedule for the workbook:

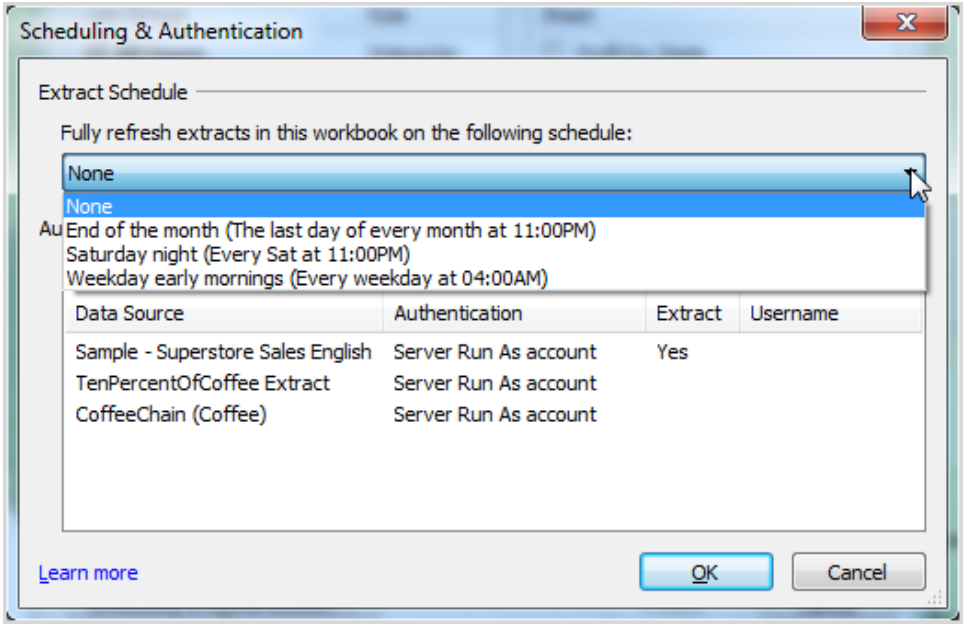

All data sources that require authentication must have an embedded password so that the extract can be refreshed. This includes data sources that are not extracts. For more information about embedding passwords, see **Embedding Passwords and Authentication** on page 829.

Publishing Data Sources

In addition to publishing workbooks, you can also publish data sources to Tableau Server (or Tableau Online). A data source is a reusable connection to data. Publishing data sources allows you to centrally manage and store data sources. The published data can be located in Tableau's data engine (extracts) or in a live, relational database. The published data source also contains field level customizations such as calculations, groups, sets and default properties. This topic describes how to publish a data source to Tableau Server.

To publish a data source to Tableau Server

1. Select the data source on the **Data** menu and then select **Publish to Server**:

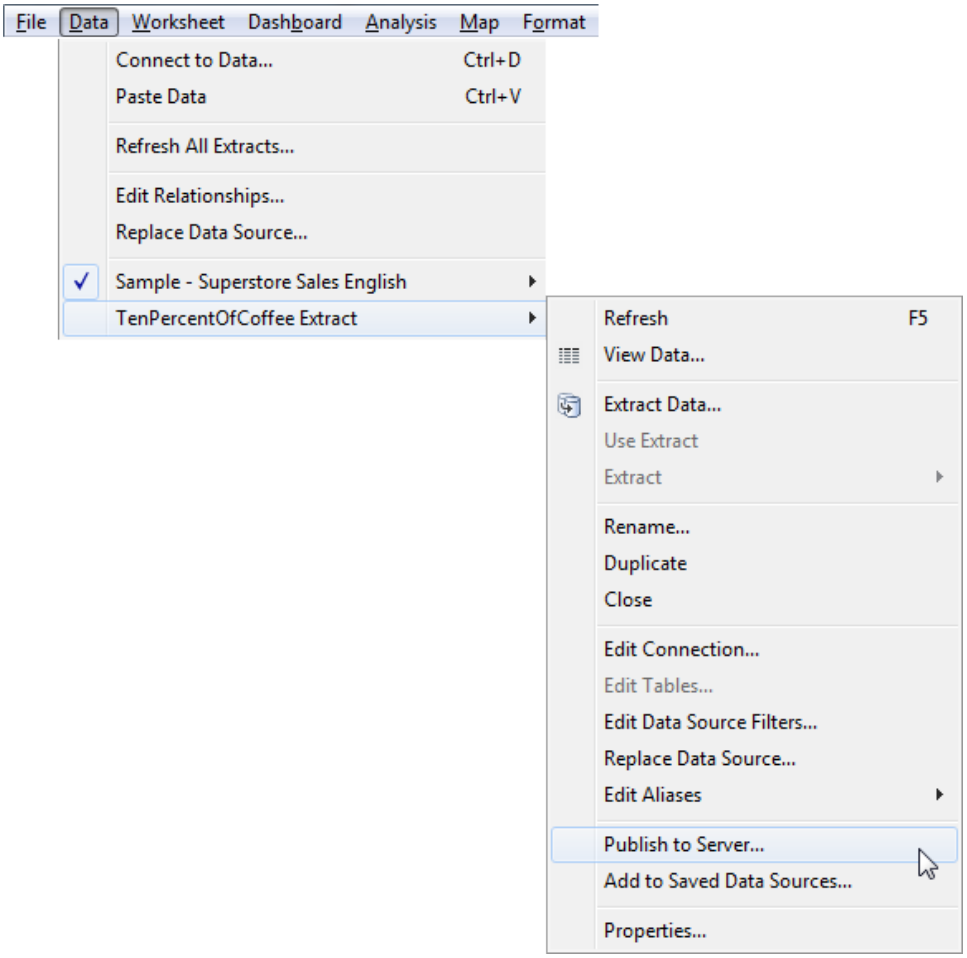

2. If you are not already signed in to Tableau Server, you see the Tableau Server Sign In dialog box.

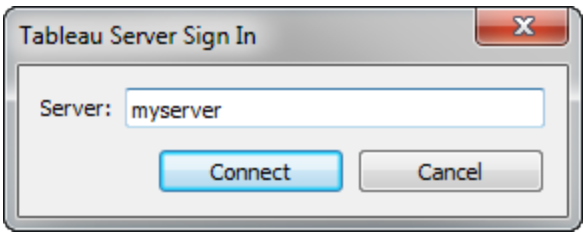

Type the server name or URL (for example, sales_server, or https://sales_server) and click **Connect**.

If Tableau Server is Kerberos-enabled and you are on a machine with valid Active Directory credentials, you will connect automatically to the server. In this case, you can skip step 3. If Kerberos authentication fails, you will be prompted to provide a user name

and password for authentication. If you need to change the user you're signed in as, see Switch Kerberos User.

3. Next, enter your user name and password and click **Sign In**.

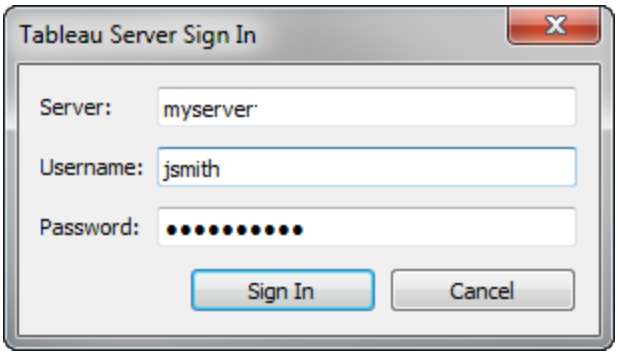

If Tableau Server is configured to use Active Directory, enter your Windows user name (the domain is not required—except in multi-domain environments where the user is not in the default domain); otherwise, enter your Tableau Server user name.

If Tableau Server is configured to use SAML for user authentication, you won't see the above dialog box. Instead, you'll see a sign in prompt from an external identity provider.

4. You now see the Publish Data Source to Tableau Server dialog box.

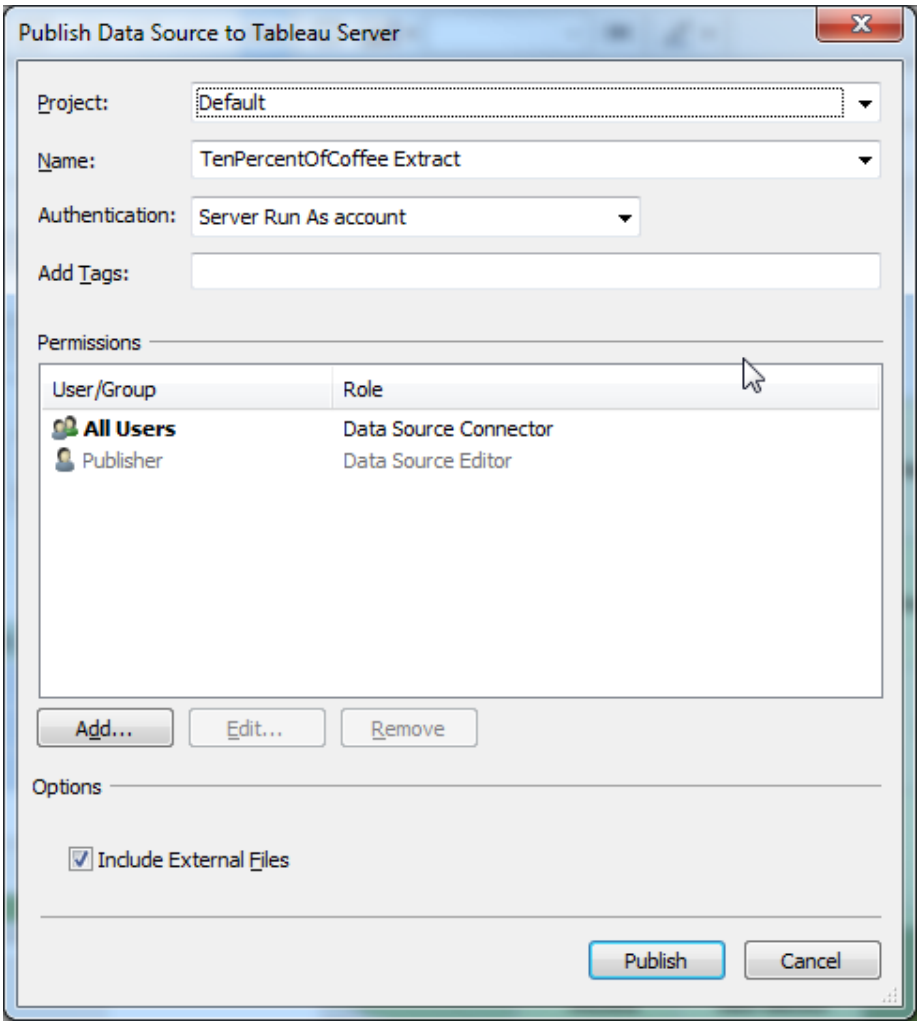

Specify the following:

- **Project**: A project is like a folder that can contain workbooks and data sources. Tableau Server comes with one project called Default. All data sources must be published into a project.
- **Name**: Provide a name for the data source in the **Name** text box.

  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/Web Save As** permission to overwrite data sources on the server.
- **Authentication**: If your data source requires a user name and password, you can specify how authentication should be handled when it is published to the server. The options available depend on the type of data source you are publishing.

Select **Viewer credentials** to use the user's domain user name and password to automatically authenticate Tableau Server users when they access a Kerberos-enabled Microsoft SQL Server or Microsoft Analysis Services data source.

If you're using Impala with Cloudera Hadoop and Tableau Server is configured for single sign-on (SSO), select **Viewer Credentials** to automatically authenticate Tableau Server users when they access the data source.

If you are publishing a Salesforce.com, Google Analytics, or Google BigQuery data source, see **Using OAuth for Data Connections** on page 55 for more information.

- **Add Tags**: You can enter 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 Data”).

5. You can specify permissions to allow or deny access to the data source on the server. See **Setting Permissions** below for details, and see **Permissions** on page 839 for information on the capabilities you can assign to users and groups.
6. If you are publishing an external file data source or a data source that is on a mapped drive select **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 data source online.

   If you do not want to publish the external files to the server, change the connection information so that the data source references a full UNC path. For example rather than connecting to D:\datasource.xls, you would connect to \\filesrv\datasource.xls.
7. Click **Publish**.

Setting Permissions

When you publish workbooks and data sources to Tableau Server, you have the option to specify permissions both for groups and specific users. Permissions allow or deny access for a specified object on Tableau Server to a specified user or group.

- For workbooks, the default is for all users to be able to view published workbooks and for you, as the publisher, to have all permissions. Server administrators can also edit workbooks on the web, as can users with the Interactor license level. Administrators assign a license level when they create users on Tableau Server.
- For data sources, the default is for all users to be able to see the data source in Tableau Server (and therefore to create new workbooks based on it) and for you, as the publisher, to have all permissions.

You can modify these permissions for specific users and groups in Tableau Server when you publish a workbook or data source. You can also modify them on Tableau Server if you have the appropriate authorizations. See **Permissions** on page 839.

To set permissions:

1. In the Publish Workbook to Tableau Server or Publish Data Source to Tableau Server dialog box, click **Add**.

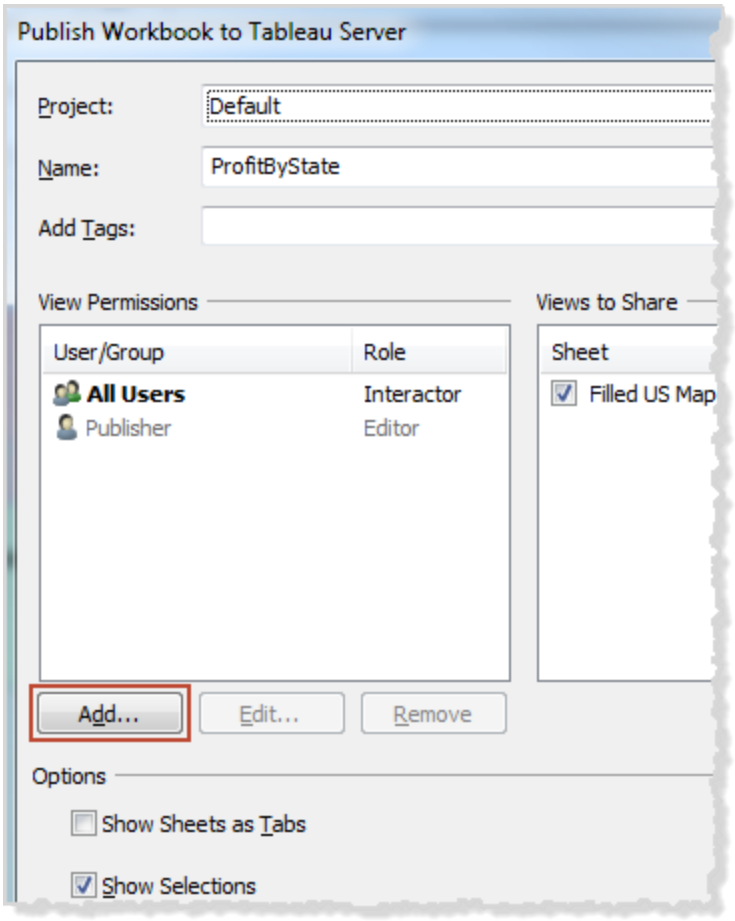

2. In the Add/Edit Permissions dialog box, select a user or group.
3. Select a predefined **Role** from the drop-down menu, or specify individual permissions. The list of permissions and the available roles vary depending on whether you are publishing a workbook or a data source. See **Permissions** on page 839 for a table that defines permissions and indicates whether they apply to workbooks, data sources, or both.

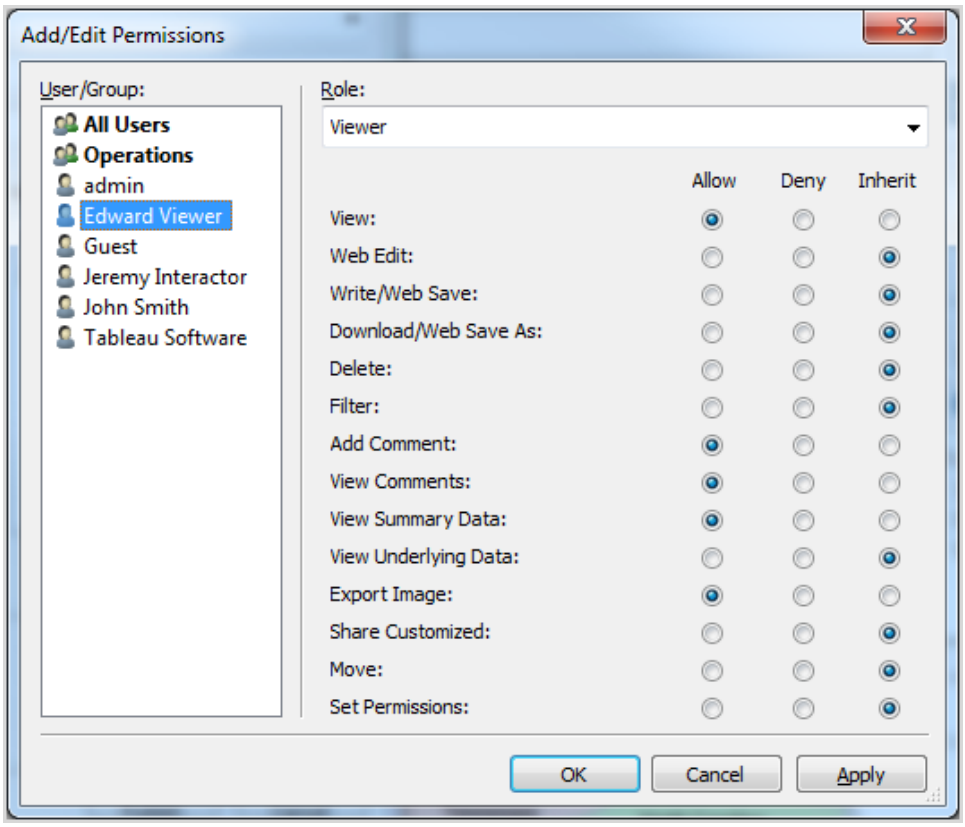

Tableau Server provides the following pre-set roles when you publish a workbook:

- **Viewer**: Allows the user or group to view the workbook on the server, as well as add and view comments.
- **Interactor**: Allows the user or group to view the workbook on the server, edit workbook views, apply filters, view underlying data, export images, and export data. All other capabilities are inherited from the user's group and project permissions.
- **Editor**: Allows all capabilities to the user or group.

For a data source, the predefined roles are:

- **Data Source Connector**: Allows the user or group to connect to the data source on the server.
- **Data Source Editor**: Allows the user or group to connect to the data source on the server and to publish, edit, download, delete, set permissions, and schedule refreshes for the data source.

When you specify individual permissions, the name assigned to the role is (Custom). In this case, the user's permissions do not align with a role. You cannot modify any of the defined roles.

If you set a permission to **Inherit**, that permission is determined by the user's group and project permissions.

4. Click **OK**.

The procedure for editing permissions for a user or group already listed in the Publish Workbook to Tableau Server dialog box would be the same, except you would select the user or group and then click **Edit**. Similarly, to delete permissions for a user, select the user and click **Remove**.

> If you are adding permissions for several users and groups, click **Apply** after specifying permissions for each user or group. The set of permissions is added and the Add/Edit Permissions dialog box remains open so you can continue.

Permissions

A permission is an action or set of actions that can either be allowed or denied to users and groups when publishing a workbook and/or data source from Tableau Desktop. Permissions can also be set in Tableau Server, on views, workbooks, data sources, and projects.

Permission	Affects...	When allowed, users can...
View	workbooks data sources	View the item on Tableau Server.
Web Edit	workbooks	Edit the workbook on Tableau Server. See Edit and Create Views in the Tableau Server help.
Write/Web Save	workbooks data sources	Overwrite the item on Tableau Server. When allowed, the user can re-publish a workbook or data source from Tableau Desktop, thereby becoming the publisher and gaining all permissions. Subsequently, the original publisher's access to the workbook is determined by that user's group permissions--and any further permissions the new publisher decides to set. This permission also determines whether the user can overwrite a workbook after editing it on Tableau Server. See Who Can Edit and Create Views in the Tableau Server help. Special Consideration for the **All Users** group: In the interest of protecting a publisher's content from being overwritten by another user (either via publishing from Tableau Desktop or saving a web-edited workbook on Tableau Server), whenever a user publishes into a project where the **All Users** group has permissions, the **Write/Web Save As** permission for the **All**

Permission	Affects...	When allowed, users can...
		Users group is changed from Allow to Inherit by default. You can manually modify this permission by following the steps in **Setting Permissions** on page 836 to change this from Inherit to Allow.
Download/Web Save As	workbooks data sources	Download the workbook or data source from Tableau Server. See Download Workbooks in the Tableau Server help. This permission also determines the user's or group's ability to save an edited workbook on Tableau Server. See Who Can Edit and Create Views in the Tableau Server help.
Delete	workbooks data sources	Remove the item from Tableau Server.
Filter	workbooks	Modify quick filters, keep only filters, and exclude data.
Add Comment	workbooks	Add comments to the views in the workbook.
View Comments	workbooks	View the comments associated with the views in the workbook. See Comment on Views in the Tableau Server help.
View Summary Data	workbooks	View the aggregated data in a view, or in the user's selection within the view, and download that data as a text file.
View Underlying Data	workbooks	View the raw data behind each row in the view, as restricted by any marks the user has selected, and download the data as a text file.
Export Image	workbooks	Export each worksheet as an image. See Export Views in the Tableau Server help.
Share Customized	workbooks	Make saved customizations to a view available for others to see. Users can create custom views using the **Remember my changes** option in Tableau Server. See Customized Views in the Tableau Server help.
Move	workbooks	Move workbooks between projects.
Set Permissions	workbooks data sources	Specify permissions for the data source or workbook. For workbooks, this permission extends to the views in a workbook.
Connect	data sources	Connect to the data source. A user who accesses a view that connects to a Tableau Server data source must have both View permission for the workbook and Connect permission for the data source.

Opening Workbooks from the Server

If you have been granted the **Download/Web Save As** permission 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/Web Save** permission, you can republish the workbook to the server.

To open a workbook from the server:

1. Select **Server** > **Open Workbook**.

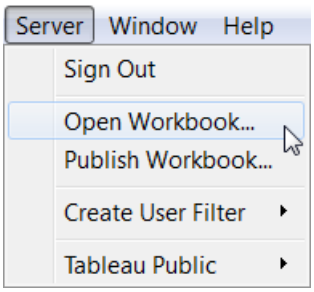

2. If you are not already signed in to Tableau Server, you see the Tableau Server Sign In dialog box.

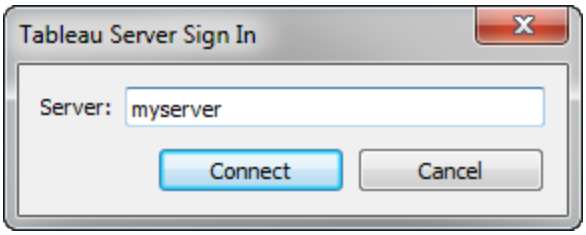

   Enter the server name or URL (for example, sales_server, or https://sales_server) and click **Connect**.

   If Tableau Server is Kerberos-enabled and you are on a machine with valid Active Directory credentials, you will connect automatically to the server. In this case, you can skip step 3. If Kerberos authentication fails, you will be prompted to provide a user name and password for authentication.

3. Next, enter your user name and password and click **Sign In**.

If Tableau Server is configured to use Active Directory, enter your Windows user name (the domain is not required—except in multi-domain environments where the user is not in the default domain); otherwise, enter your Tableau Server user name.

If Tableau Server is configured to use SAML for user authentication, you won't see the above dialog box. Instead, you'll see a sign in prompt from an external identity provider.

4. In the Open Workbook from Tableau Server dialog box, select the workbook you want to open.

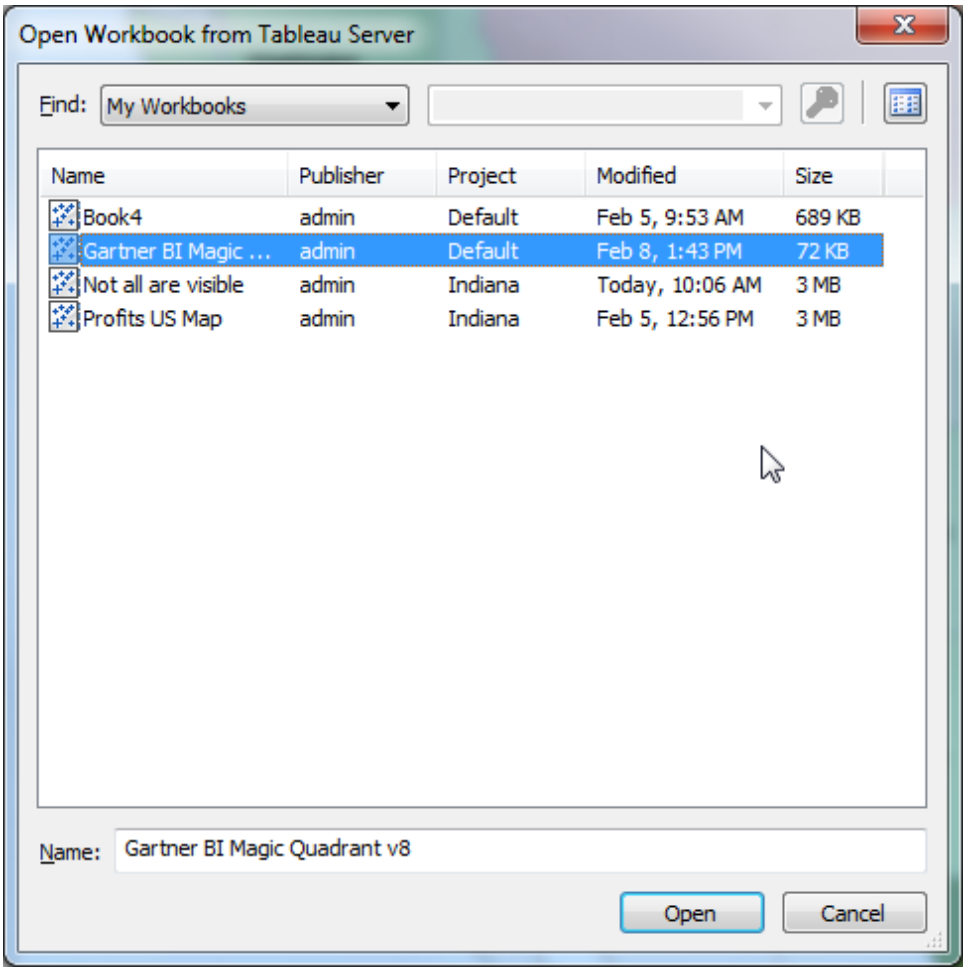

You can find workbooks using the **Find** drop-down lists. You can search all workbooks on the server or find by tags, publisher, project, or workbooks that you published.

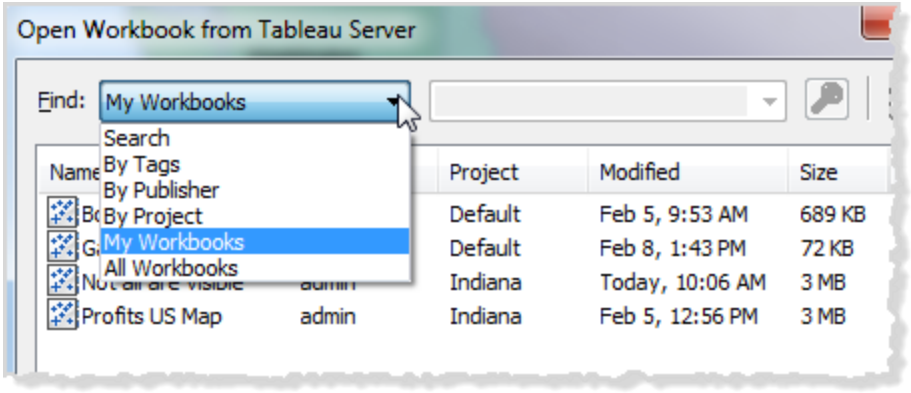

5. Click **Open**.

Importing Data Sources from the Server

You can connect to Tableau Server data sources on the Connect to Data page. See the connection example for **Tableau Server** on page 63 to learn how to connect to a published data source.

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. 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 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 signed in and looking at it.

How to Create a User Filter

You can create as many user filters as you like for a given view.

1. Select **Server** > **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 select the Product field.

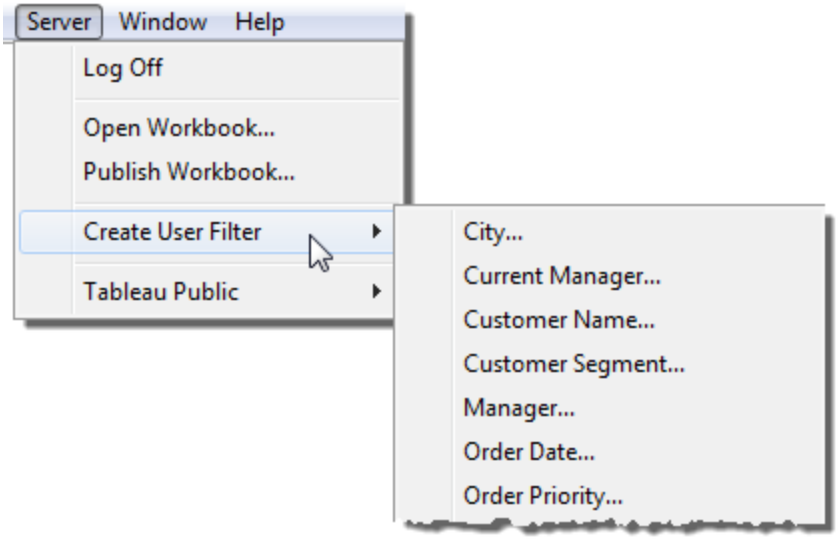

2. If prompted, sign in to Tableau Server. See **How to Publish Workbooks to a Tableau Server** on page 823 for information about signing in.
3. In the User Filter dialog box, type a **Name** for the set of rules you are creating. For example, if you are filtering on product information you could name 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. If there are a large number of users and groups, a search box is available on the left side of the dialog box.

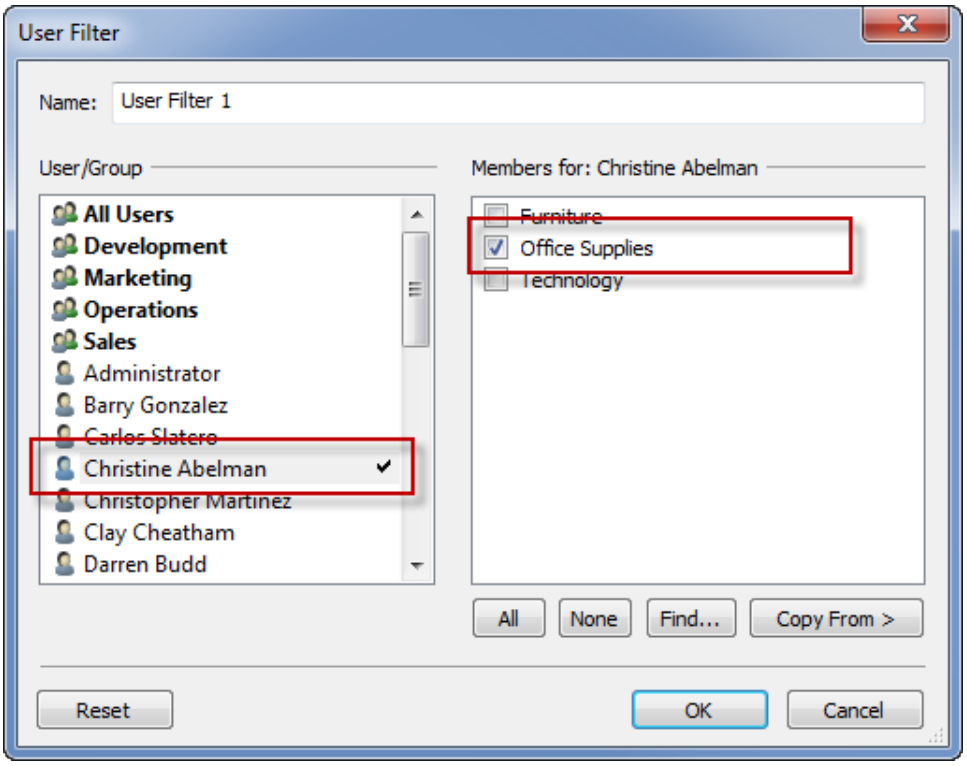

5. When finished, click **OK**.

   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. See **Example - Setting User Filters** on page 851 for a step-by-step example of defining user filters.

Finding Users in the User Filter Dialog Box

When there are many users and groups on the server, a search box is displayed on the left side of the User Filter dialog box. Begin typing the name of a user or group into the text box at the top of the list of users. When ready, press Enter on your keyboard or click the Search icon. Matching names and groups appear in a drop-down list. You can filter the results to only show Users or Groups, or show both User and Groups. Click on a user or group name to add it to the selection. Continue to search and add to the selection until you have the list you need.

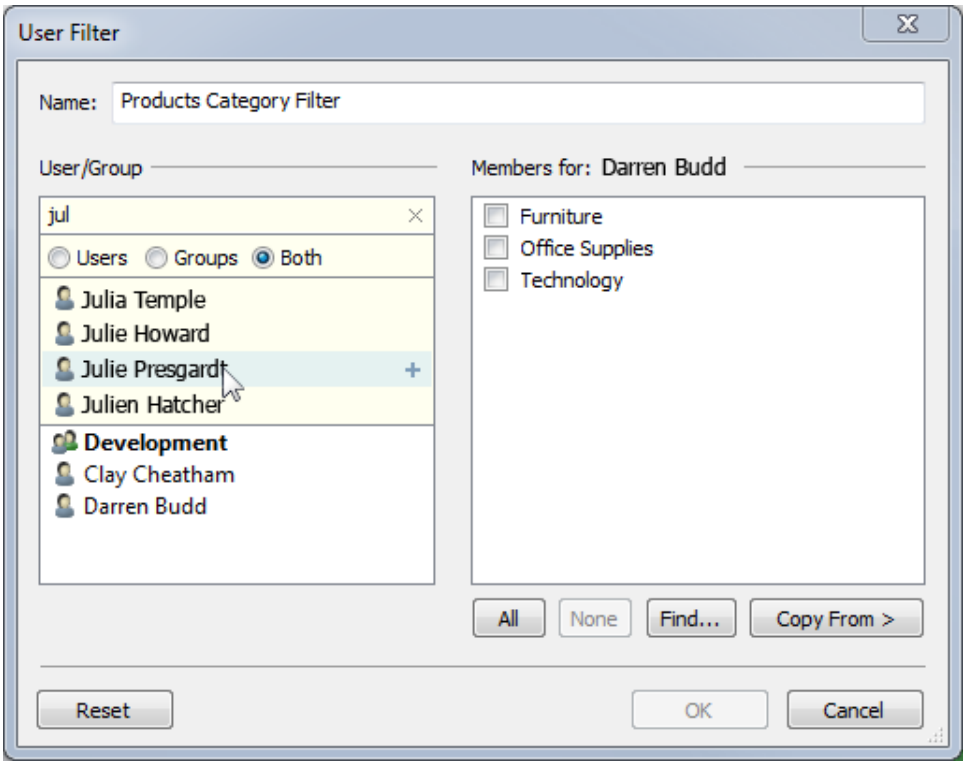

Finding Field Members in the User Filter Dialog Box

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.

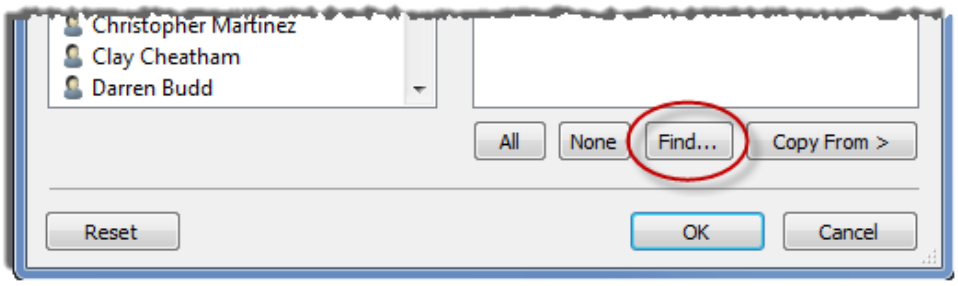

2. Type all or part of the member name into the **Find value** text box at the top of the Find Values dialog box. Then click **Find**. You can change the search criteria by selecting whether to return members that **Contain**, **Starts with**, or is an **Exact match** for the text you typed.
3. Select 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.

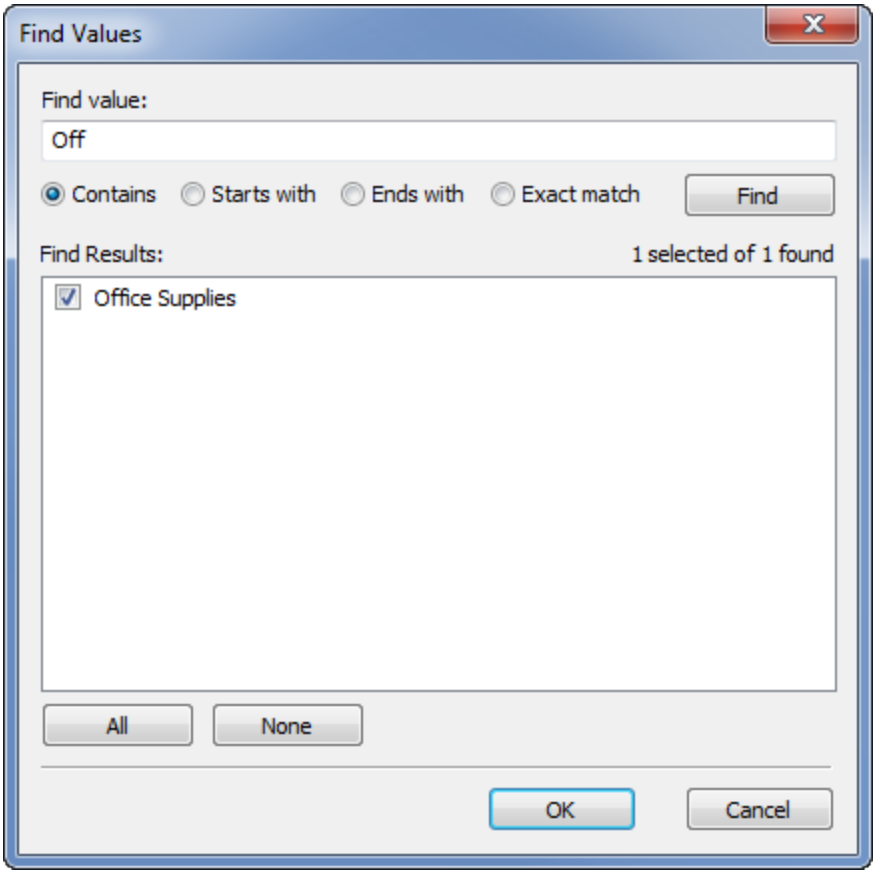

4. When finished, click **OK** to return to the User Filter dialog box.

Copying member selections from other users

As you specify which 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 set the same filter for everyone else on the team, you can simply duplicate the member selection instead of having to select the products again.

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.

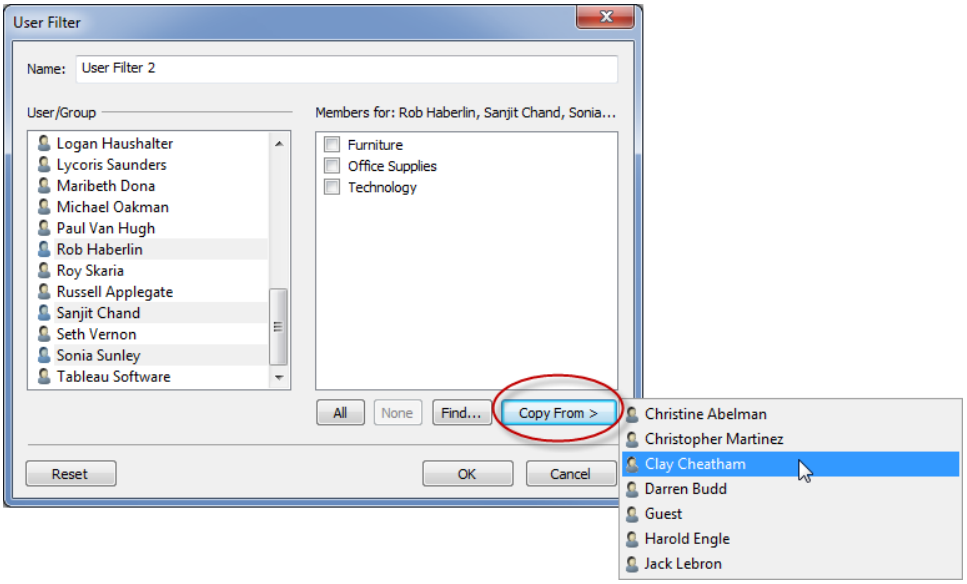

Previewing User Filters

When a workbook contains one or more user filters and you are connected to Tableau Server, the Filter as User menu appears in the bottom right corner of the workbook window. The Filter as User menu allows you to preview what each user or group will see when they look at the view on Tableau Server.

> Previewing user filters is not available for connections to Tableau Server Data Sources.

To preview a user filter:

1. Open the Filter as User 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.

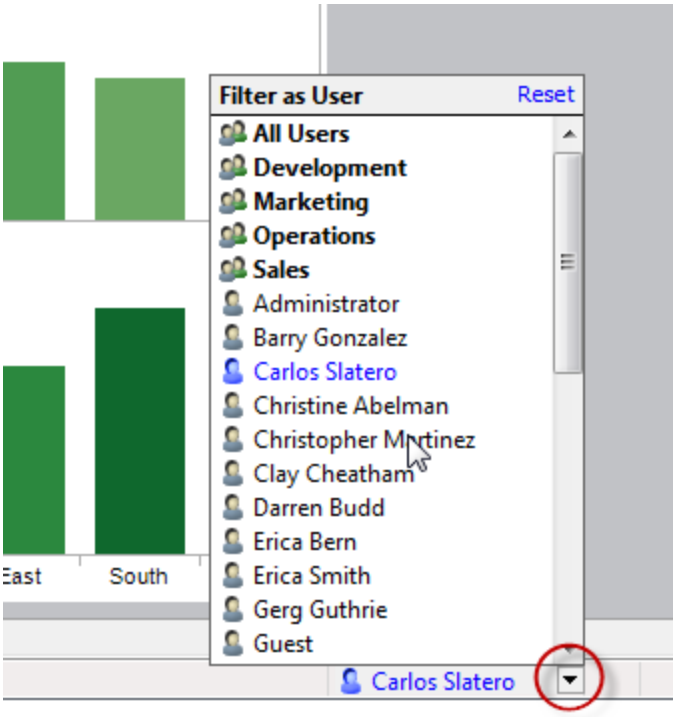

At any time you can return to viewing the workbook as yourself by clicking **Reset** in the top right corner of the Filter as user menu.

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 Set**.

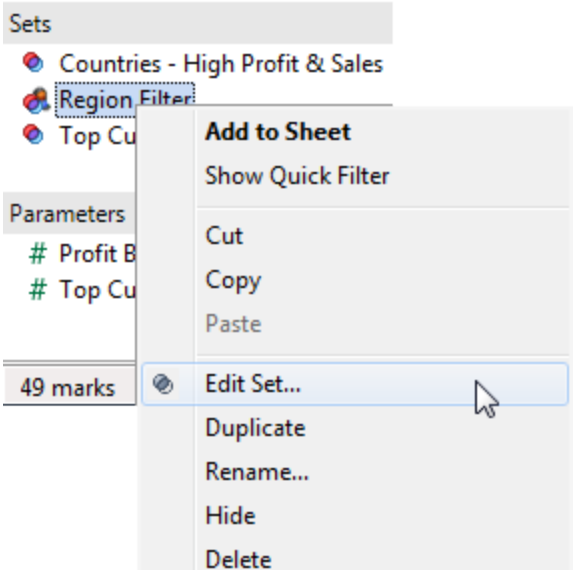

2. If necessary, sign in to Tableau Server. See **How to Publish Workbooks to a Tableau Server** on page 823 for details on signing in to the server.
3. In the User Filter dialog box, make the necessary changes and click **OK**.

If you modify a user filter while not signed in to the server, the set will only show users and groups that have a user filter specified. Sign 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.

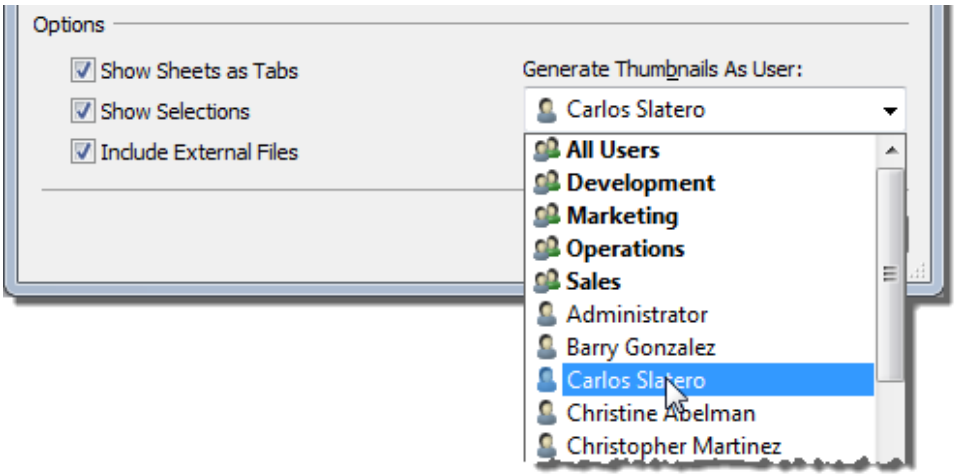

If the user you select cannot see any data, a “blank” thumbnail is displayed. For example:

> 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 data source that comes with Tableau.

The sales report is shown below.

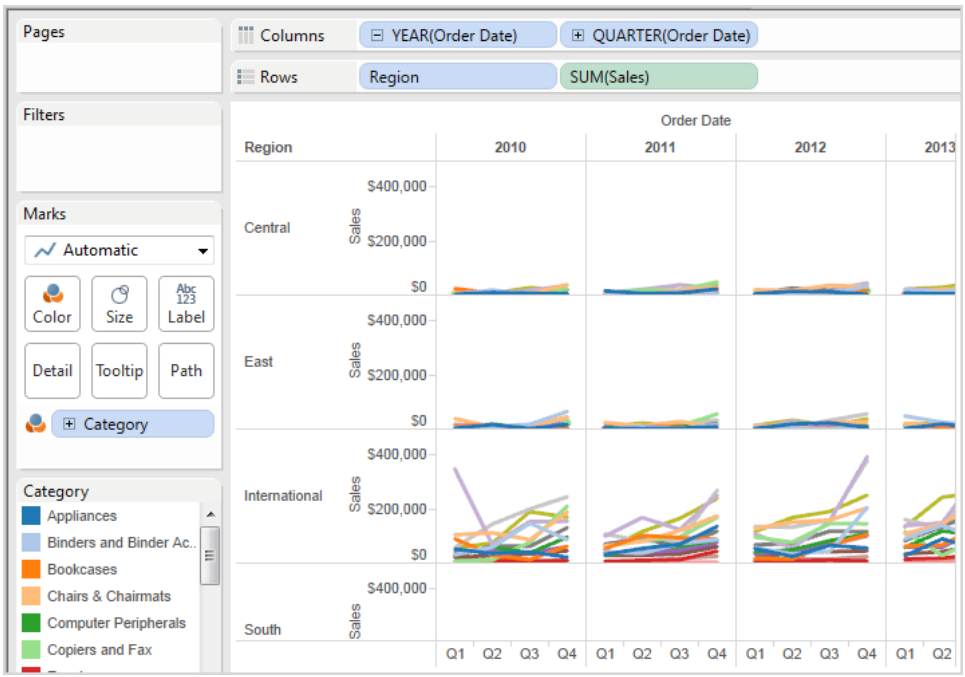

As you can see the view shows the quarterly sales over several years for each region and product category. Now suppose you are going to publish this view to the server so each of the three regional managers can track their sales. However, you want to limit each manager to only see data relevant to their region. To do that you need to set up a user filter.

1. Select **Server** > **Create User Filter**. Then select **Region** because that is the field you want to use for filtering the view.

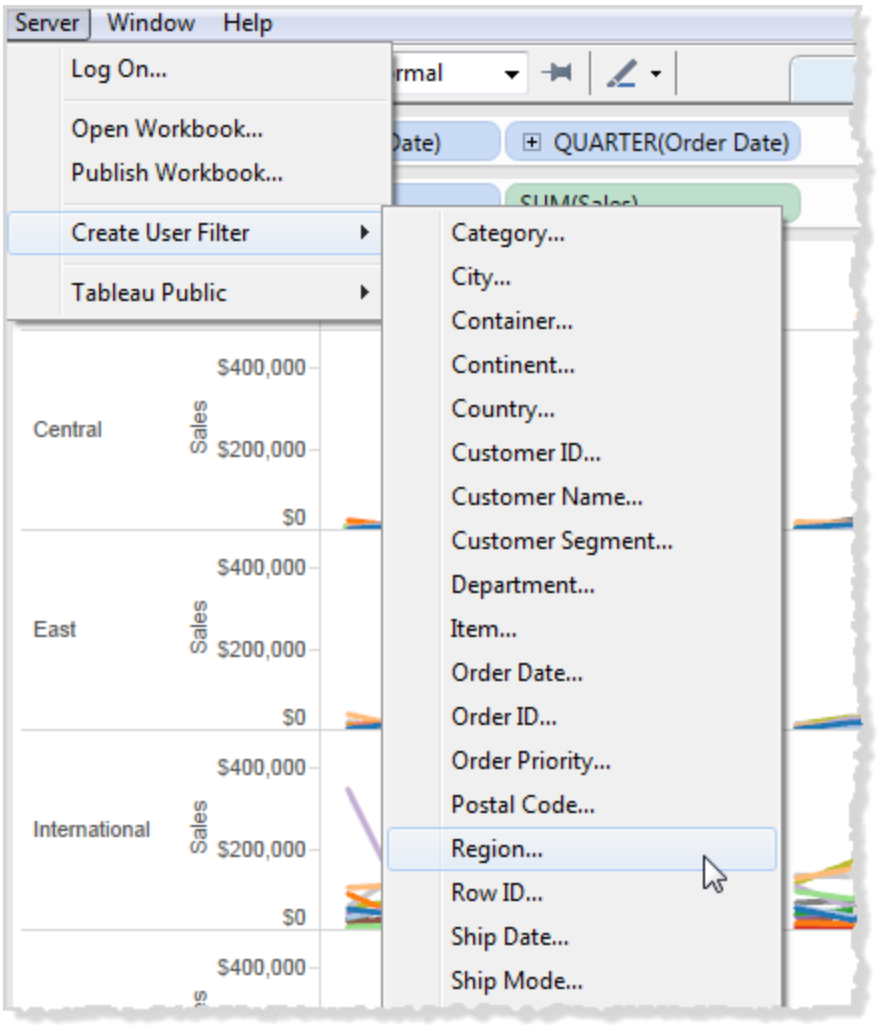

2. You may at this point be prompted to sign in to the server if are not already signed in. See **How to Publish Workbooks to a Tableau Server** on page 823 for details on what to do.
3. Type a name for the user filter. In this case we'll name it Regional Managers. After specifying the filter, it will display in the Sets area of the Data window by this name.

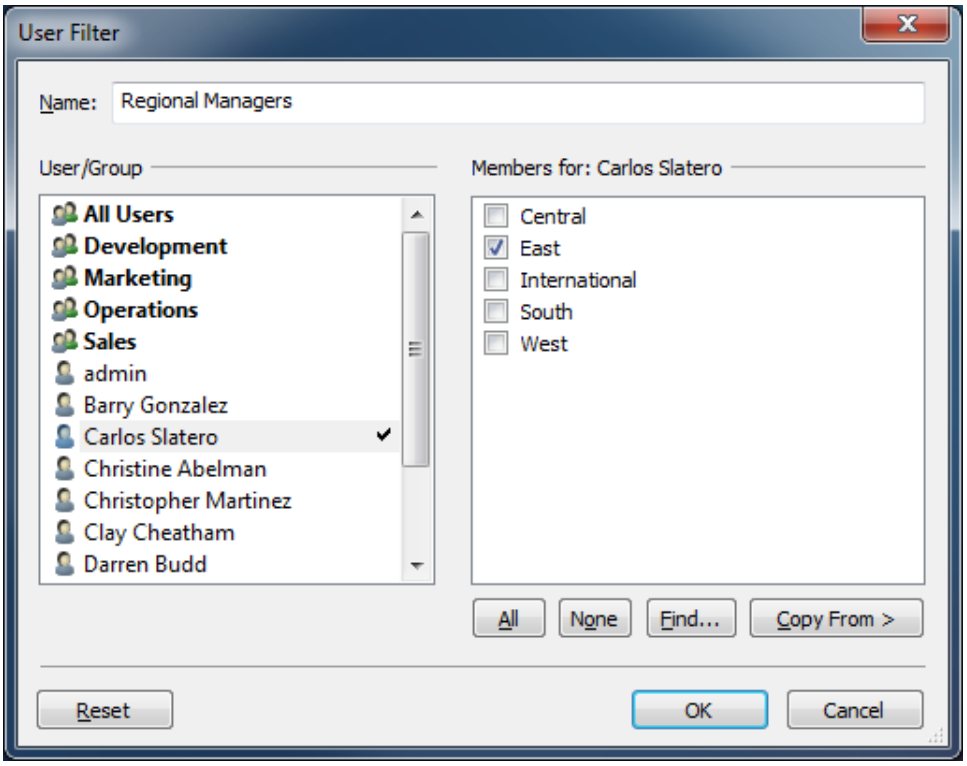

4. Then select one or more users and groups that are allowed to see the East region data. For example, above, Carlos Slatero is the Eastern Regional Manager, so his name is selected on the left side and East is selected on the right side.
5. Repeat this process as many times as necessary. In this example, there are five managers, each with responsibility for a single region. Carlo Slatero in the East, Clay Cheatham in the South, Barry Gonzalez in the West, Darren Budd in the Central, and Christopher Martinez in the International region. Nobody else will be able to see any of the regions.

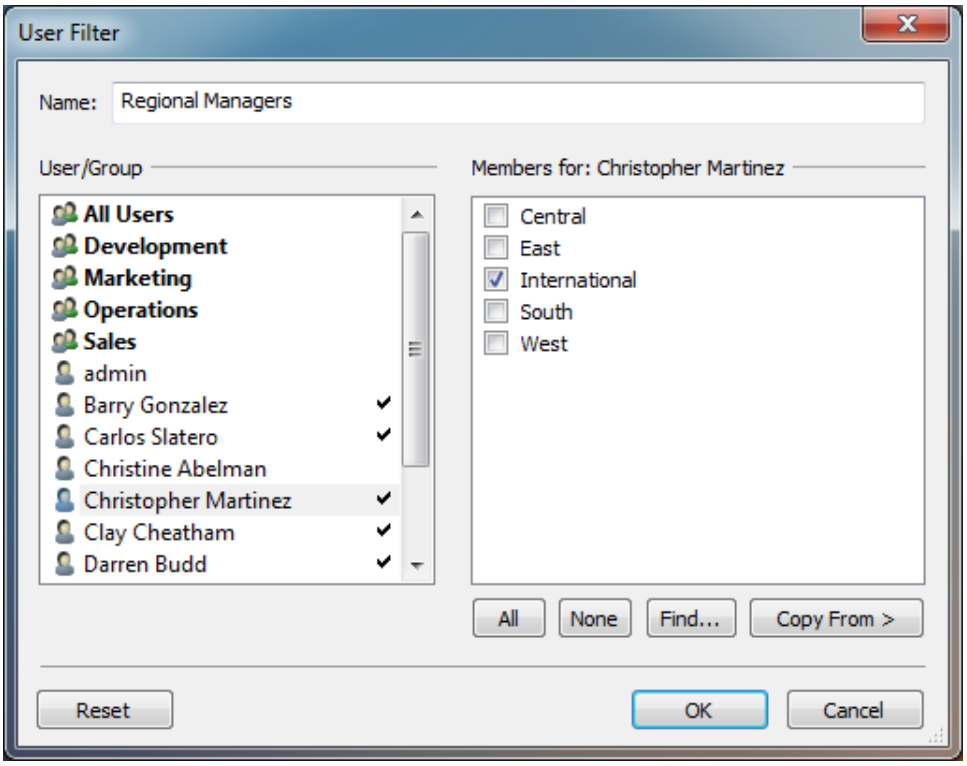

6. When finished click **OK**.
7. The user filter is now visible in the Sets area of the Data window. Now that it is defined you need to apply it to a view. Drag the Regional Managers user filter to the Filters shelf. The filter becomes a context filter. See **Filtering** on page 334 to learn more.

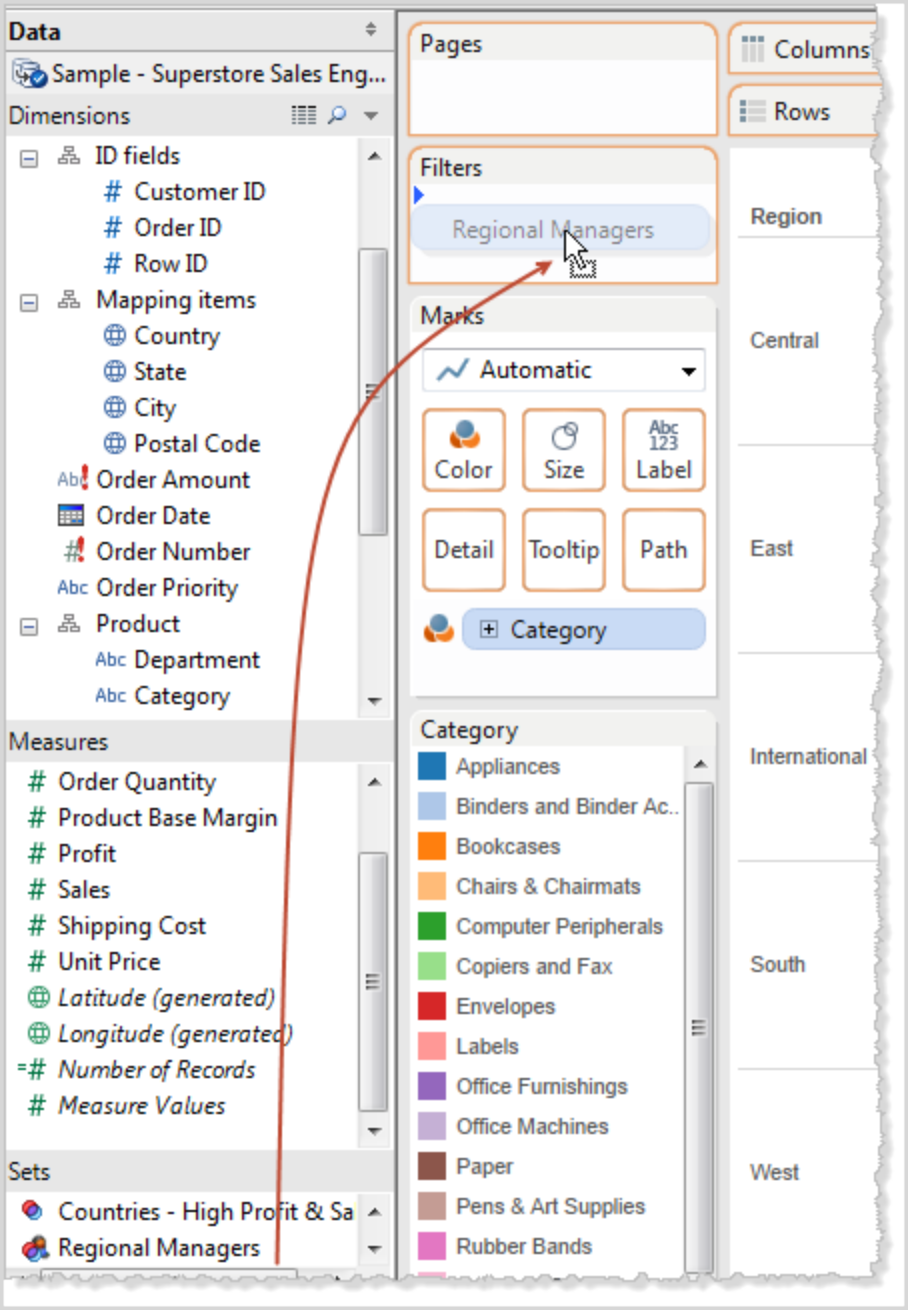

When you apply a user filter to a view, it's possible that you will not see anything. For example, in this case you have not allowed yourself to see any of the three regions. To edit a user filter, right-click it in the Data window and select **Edit Set**.

At anytime, you can preview what the view will look like for each user using the Filter as User 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.

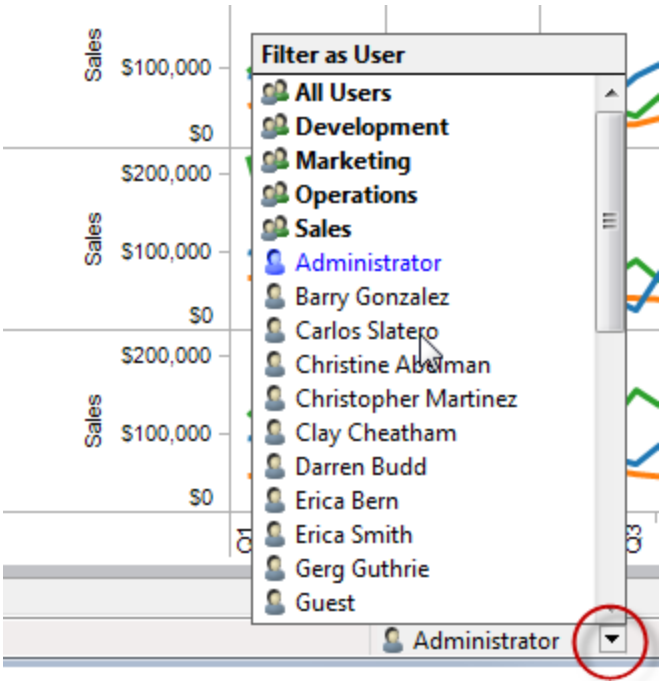

Filter Using a List of Users in your Data Source

Sometimes there is a field in the data source that contains your users. When this is the case, you can create user filters based on the names in the field rather than having to create a set based on server users. 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 signed in.

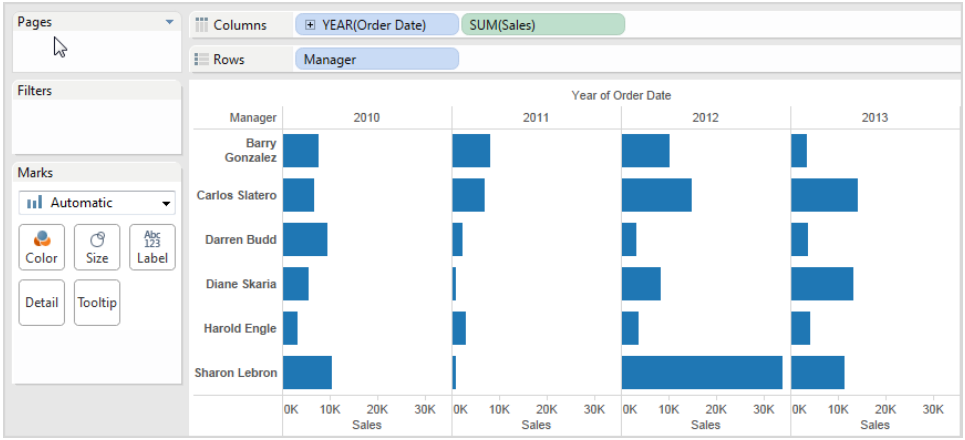

1. If you have not already signed in, select **Server** > **Sign In** to sign in to Tableau Server.
2. Create a calculated field by selecting **Create Calculated Field** on the Data window menu.

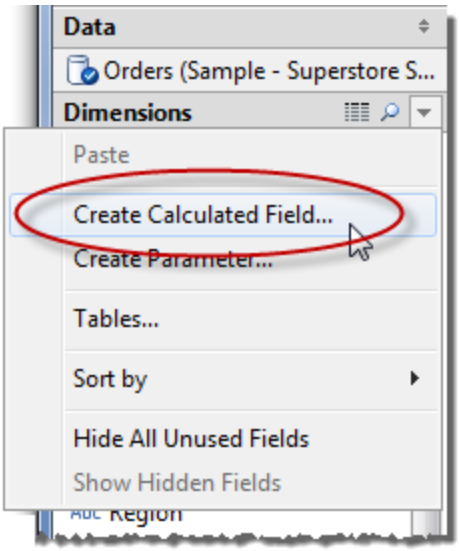

3. In the Calculated Field dialog box name the field. In this example we'll call it Current Manager.
4. Then use the Fields and User functions in the bottom half of the dialog box to create calculation that uses the following formula:

```
[Manager]=FULLNAME()
```

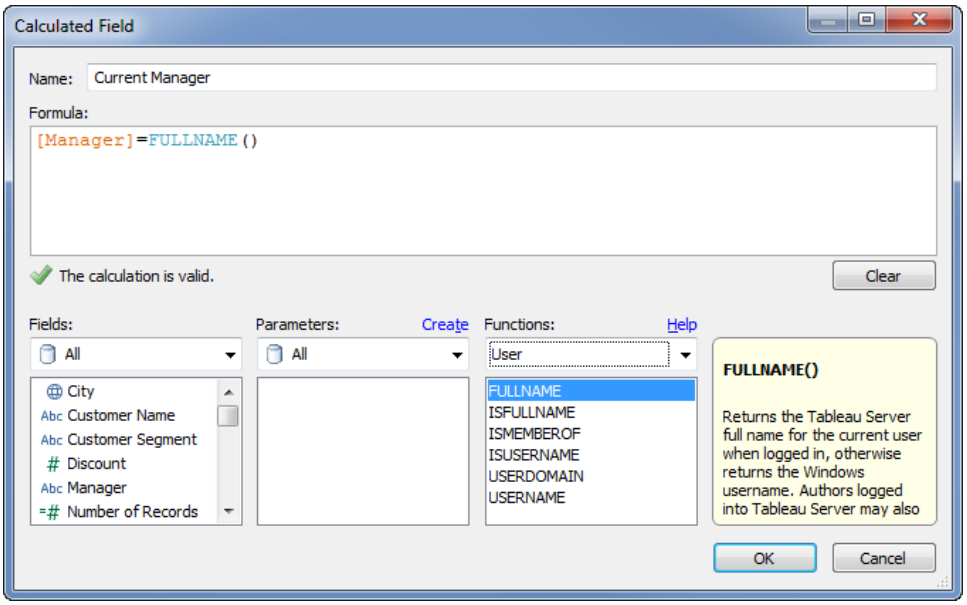

5. When finished click **OK**.

6. The field displays in the Dimensions area of the Data window.

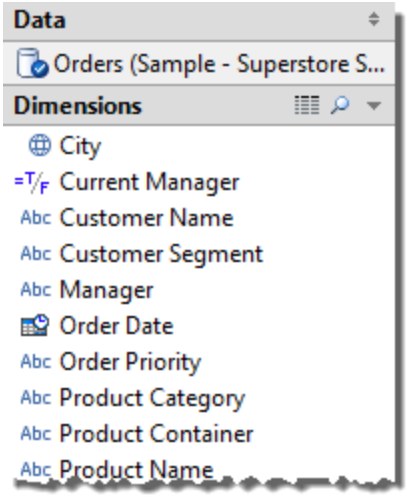

7. Place the new field on the Filters shelf. If you are not one of the users in the view (in this example one of the managers) the only option in the Filter dialog box will be False. In that case you should select nothing and click **OK**. Then you can preview the view as another user and edit the filter 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 Filter as User menu in the bottom right corner of the workbook. See **Previewing User Filters** on page 848 to learn more.

The view below shows what Carlos Slatero would see if he was signed in.

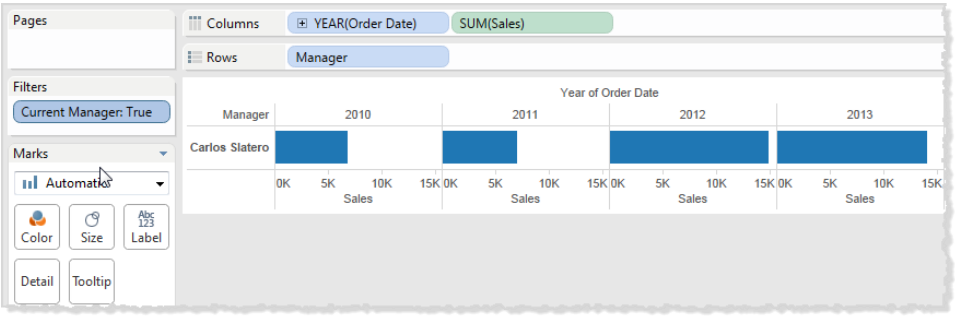

Saving and Exporting

After you create views in Tableau, you may 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.

Saving Your Work

After 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 workbook.

Custom fields such as binned measures, calculated fields, groups, and sets are saved with workbooks and bookmarks.

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**.

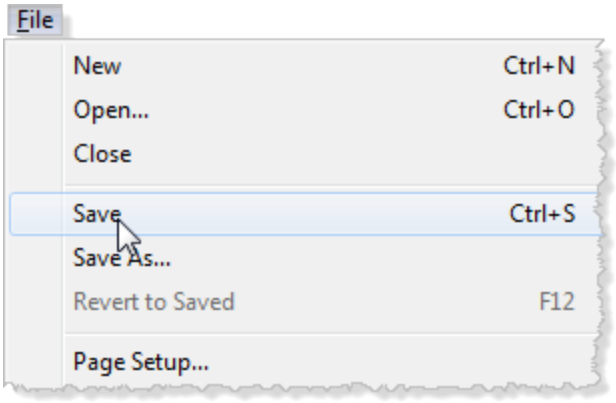

2. Specify the workbook file name in the **Save As** dialog box.

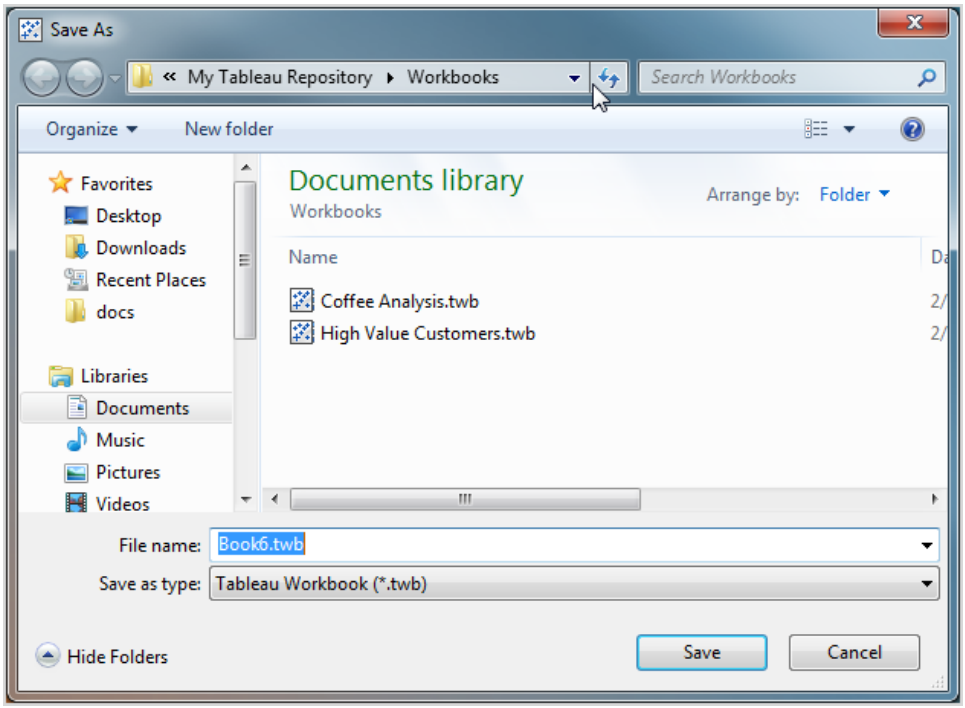

By default, Tableau saves the file with the .twb extension. The default location is the **Workbooks** folder in the Tableau repository. However, you can save Tableau workbooks to any directory you choose.

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 background images, or local file data sources such as Excel, Access, and Extract files.

When you save a workbook, links to these resources are also saved. The next time you open the workbook, the views are automatically updated with any changes that may have occurred to the data and images. In most cases, you will want to save the workbook in this way. But if you plan to share the workbook with someone who does not have access to the referenced resources, or to publish the workbook to Tableau Server, you might want to 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. These workbooks are saved with a .twbx file extension. Other users can open the packaged workbook using Tableau.

To save a packaged workbook:

1. Select **File > Save As**.

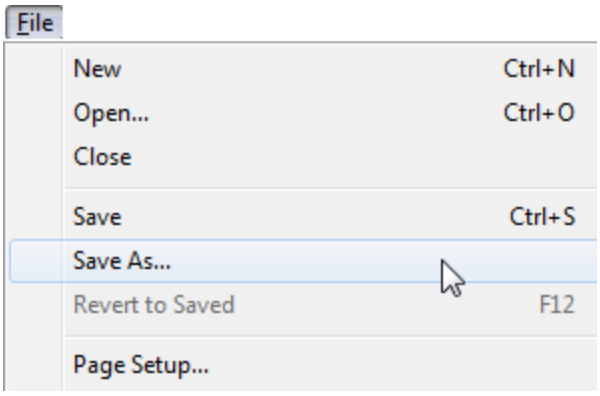

2. Specify a file name for the packaged workbook in the Save As dialog box.
3. Select **Tableau Packaged Workbooks** on the Save as type drop-down list.

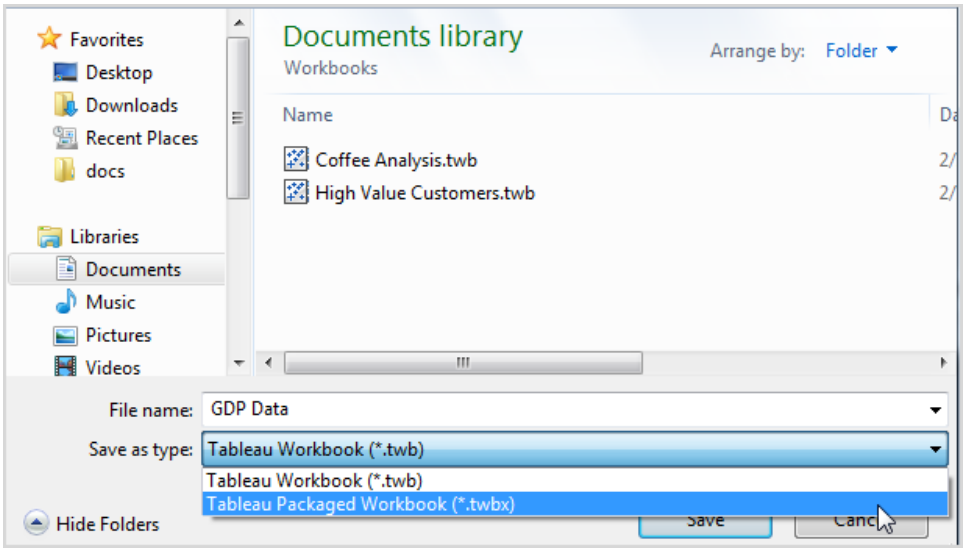

4. Click **Save**.

The default location is the **Workbooks** folder of the Tableau repository. However, you can save packaged workbooks to any directory you choose.

The following files are included in packaged workbooks:

- Excel Files
- Access Files
- Text Files
- Tableau Data Extract Files
- Local Cube Files

- Background Image Files
- Custom Geocoding

> 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 Tableau Drivers page.

Packaged Workbooks can be unpackaged.

To unpackage a workbook:

- Right-click the packaged workbook file (.twbx) in Windows Explorer and select **Unpackage**.

When you unpackage a workbook, you will get a regular workbook file (.twb), along with a folder, that contain the 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 worksheets that you use frequently.

To save a Tableau bookmark:

1. Select **Window** > **Bookmark** > **Create Bookmark**.
2. Specify the bookmark file name and location in the Create Bookmark dialog box.

Tableau saves the file with a .tbm extension. The default location is the Bookmarks folder in the Tableau Repository. However, you can save bookmarks to any location you choose. Bookmarks that are not stored in the Tableau repository do not appear on the **Bookmark** menu.

You can organize bookmarks into folders in the same way you organize files or documents. This can be useful when you have a large number of bookmarks to manage. For example, you might organize bookmarks based on employee name, product types, or sales results. You can organize bookmarks by creating a new folder, renaming an existing folder, renaming existing bookmark files, and so on.

Delete bookmarks 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 Tableau.

> While bookmarks are generally a snapshot 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

To revert a workbook is to undo all of the changes you've made to a workbook since you last saved it. You could click through your history using the **Undo** button, but there is an easier way.

To revert to the last saved state of a workbook:

1. Select **File** > **Revert to Saved**.

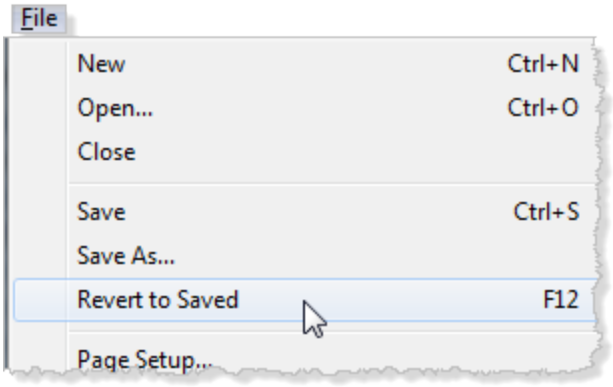

2. Click **Revert** in the warning dialog box.

   Or, just press F12.

> The **Revert to Saved** command is only available for workbooks (.twb) that do not have connections to Extract data sources.

Exporting Your Work

After you have created some 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. You can also use Tableau to present your data.

> **Note:** You can also export selected sheets to incorporate your results into other Tableau workbooks. For information, see **Export and Import Sheets** on page 872.

Export Data

If 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 options in Tableau.

When exporting data, you should keep these rules in mind:

- You can select any portion of a data view to export. If you want to export all data in a view, right-click in the view and choose **Select All**. Copying and exporting to a crosstab always exports all data in the view regardless of what you have selected.
- The fields that are exported to the new data source come from the fields on the worksheet shelves. The exception is fields that are external filters--that is, fields that 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, place those fields on the **Level of Detail** shelf.

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.

   Right-click the view and click **Select All**.

2. Select **Worksheet > Copy > Data**, or right-click the view and select **Copy > Data** from the context menu.
3. Open an Excel worksheet and paste the data into a new sheet. Notice that the fields placed on the **Rows**, **Columns**, and **Color** shelves are copied into the sheet. However, the **Customer Segment** field is not copied because it is an external filter (it appears only

on the **Filters** shelf).

Book1

	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 Ulprigh	TECHNOLOGY	EAST	51280.73	

Copy Underlying Records to Clipboard

Copying underlying data can be used to copy the disaggregated data behind a view.

To copy underlying records:

1. Select the desired data in Tableau.

   Right-click the selected records and select **View Data** on the context menu.

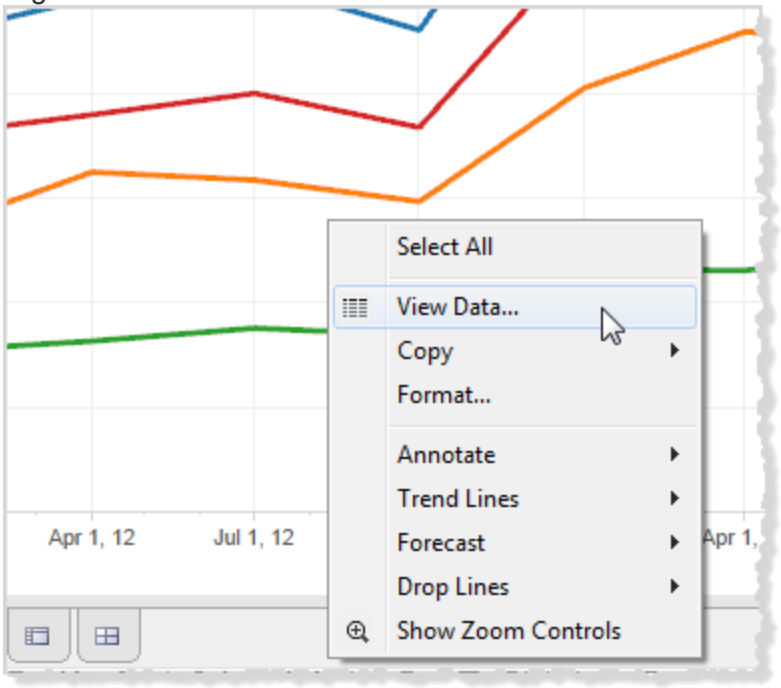

2.

3. The resulting dialog box shows the Summary data along with the Underlying Data. Select the data you want to copy and then click **Copy** in the upper right corner of the dialog box.

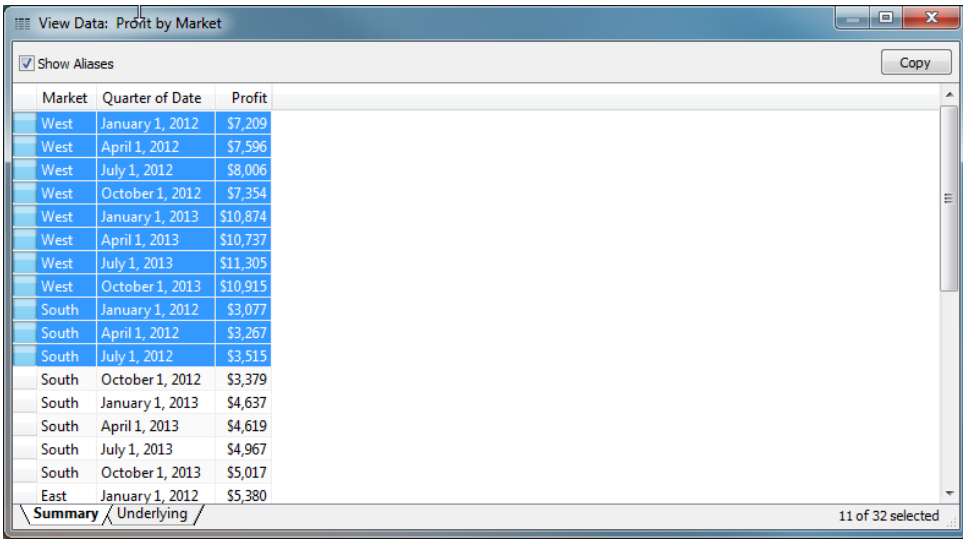

Market	Quarter of Date	Profit
West	January 1, 2012	$7,209
West	April 1, 2012	$7,596
West	July 1, 2012	$8,006
West	October 1, 2012	$7,354
West	January 1, 2013	$10,874
West	April 1, 2013	$10,737
West	July 1, 2013	$11,305
West	October 1, 2013	$10,915
South	January 1, 2012	$3,077
South	April 1, 2012	$3,267
South	July 1, 2012	$3,515
South	October 1, 2012	$3,379
South	January 1, 2013	$4,637
South	April 1, 2013	$4,619
South	July 1, 2013	$4,967
South	October 1, 2013	$5,017
East	January 1, 2012	$5,380

4. Open an Excel worksheet and paste the data into a sheet. Notice that the fields placed on the Rows, Columns and Color shelves are copied onto the worksheet.

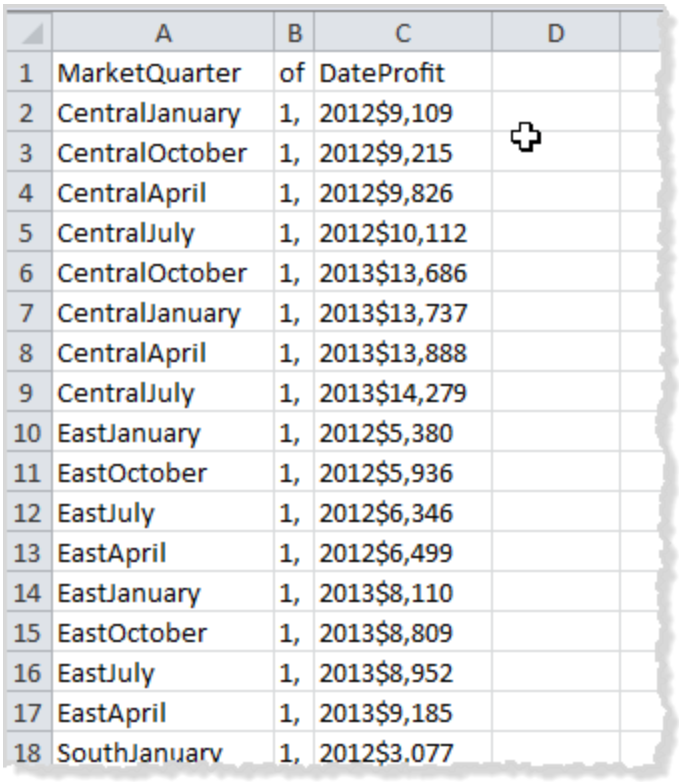

	A	B	C	D
1	MarketQuarter	of	DateProfit	
2	CentralJanuary	1,	2012$9,109	
3	CentralOctober	1,	2012$9,215	
4	CentralApril	1,	2012$9,826	
5	CentralJuly	1,	2012$10,112	
6	CentralOctober	1,	2013$13,686	
7	CentralJanuary	1,	2013$13,737	
8	CentralApril	1,	2013$13,888	
9	CentralJuly	1,	2013$14,279	
10	EastJanuary	1,	2012$5,380	
11	EastOctober	1,	2012$5,936	
12	EastJuly	1,	2012$6,346	
13	EastApril	1,	2012$6,499	
14	EastJanuary	1,	2013$8,110	
15	EastOctober	1,	2013$8,809	
16	EastJuly	1,	2013$8,952	
17	EastApril	1,	2013$9,185	
18	SouthJanuary	1,	2012$3,077	

Export Records To Microsoft Access

To create an Access database from Tableau data, follow these steps:

1. Select the desired data in Tableau. Right-click the view and then select **Select All**.

2. Select **Worksheet > Export > Data**.

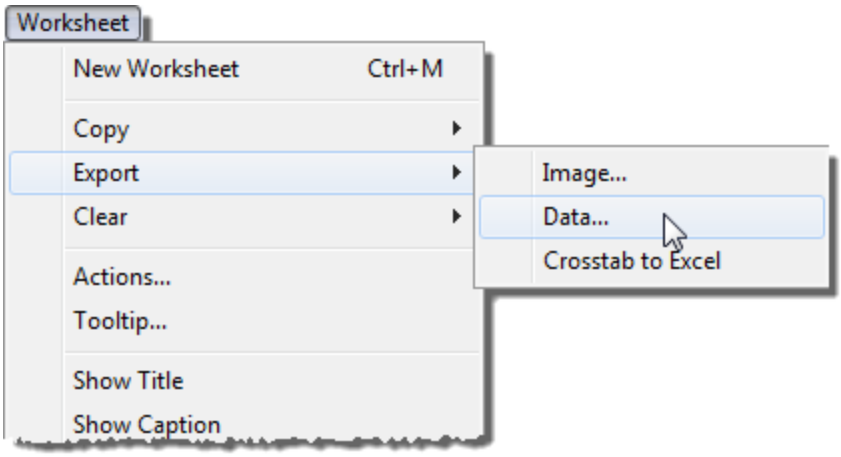

3. Select a location and name for your Access database. Access databases end with a .mdb file extension.

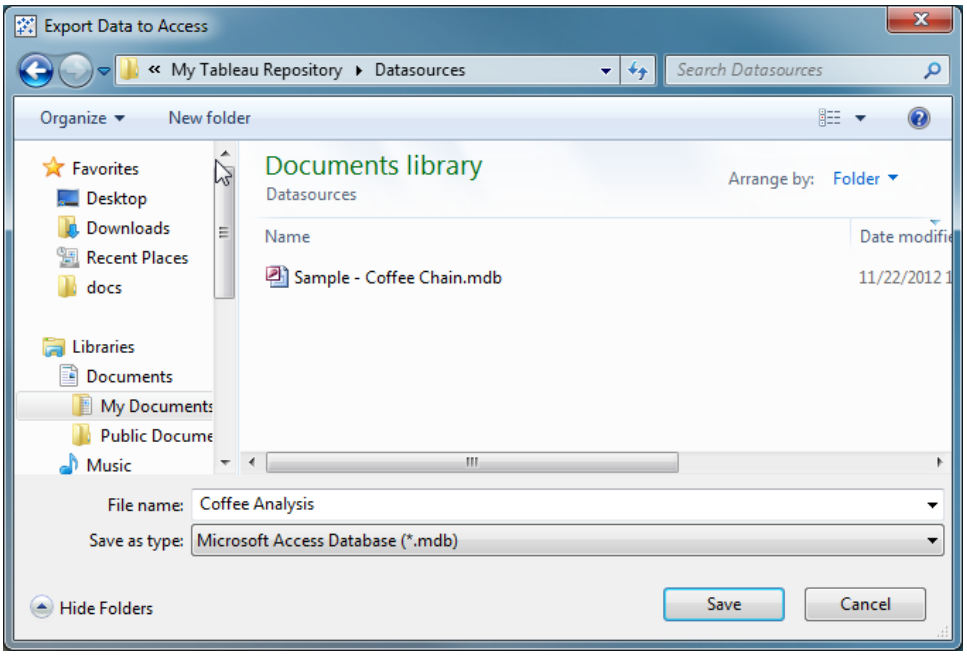

4. Type a name for the database and click **Save**. The Export Data to Access dialog box is now displayed. The Connect after export option allows you to immediately connect to the new data source and continue working in Access without interrupting your work flow.

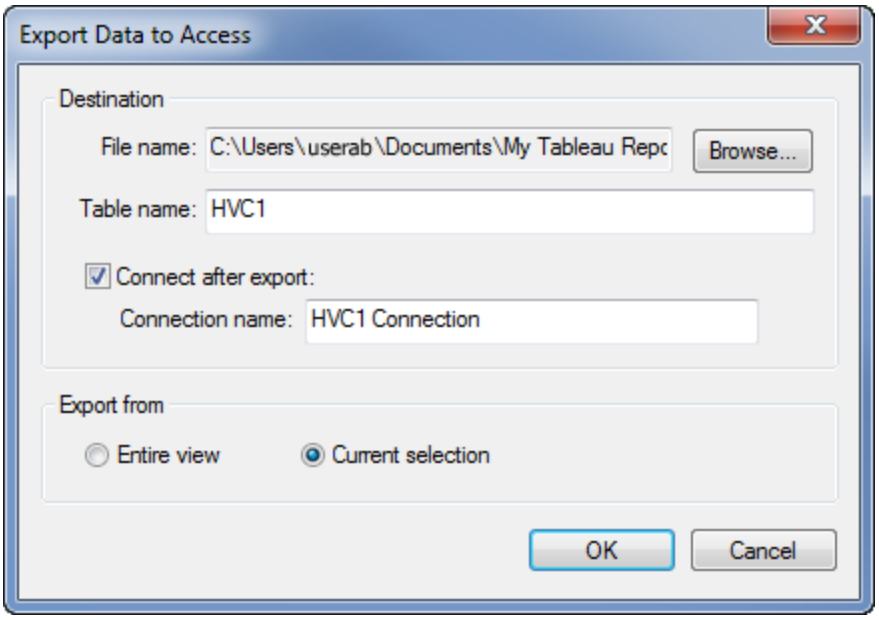

Copy Crosstab to Clipboard

You can copy a crosstab (text table) version of a view to your local Clipboard and transfer it from the Clipboard to another application. For instance, you might want to transfer a crosstab in Tableau to a crosstab in Microsoft Excel. Or you might want to transfer the data behind a

graphical view in Tableau to Excel in a crosstab format. Copying a crosstab to the Clipboard is restricted by the following general rules:

- You must copy all records in the view. You cannot copy a subset of records.
- This option is valid for aggregated views only. It cannot be used on disaggregated views of data, because a crosstab is by definition an aggregated view of data. In other words, the **Aggregate Measures** option on the Analysis menu must be selected in order for this option to work properly.

Other restrictions may apply depending on the data in your view. You cannot copy a crosstab if the view contains continuous dimensions such as continuous dates and times.

To copy a view as a crosstab to the clipboard:

1. Right-click any view in Tableau and select **Copy > Crosstab**.

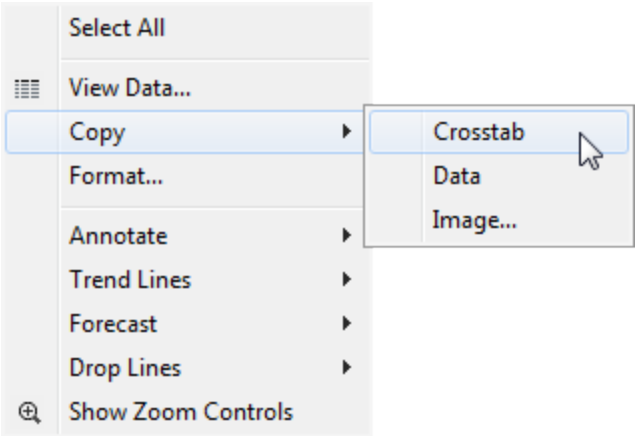

2. Open an Excel worksheet and select Paste from Excel's Edit menu or press Ctrl-V.

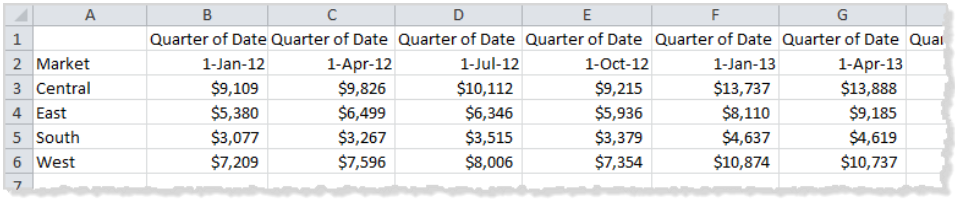

	A	B	C	D	E	F	G	
1		Quarter of Date	Quarter of Date	Quarter of Date	Quarter of Date	Quarter of Date	Quarter of Date	Quar
2	Market	1-Jan-12	1-Apr-12	1-Jul-12	1-Oct-12	1-Jan-13	1-Apr-13	
3	Central	$9,109	$9,826	$10,112	$9,215	$13,737	$13,888	
4	East	$5,380	$6,499	$6,346	$5,936	$8,110	$9,185	
5	South	$3,077	$3,267	$3,515	$3,379	$4,637	$4,619	
6	West	$7,209	$7,596	$8,006	$7,354	$10,874	$10,737	
7								

The pasted data always appears as a crosstab in Excel, even if the initial view of the data in Tableau was not in a crosstab format.

Export Crosstab to Excel

There is a more direct way to transfer a cross-tab view of data to Microsoft Excel. Select **Worksheet > Export > Crosstab to Excel**. Tableau automatically pastes a crosstab version of the current view into a new Excel workbook. This option automatically opens a new instance of the Excel application.

Although, copying a crosstab 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 can enhance performance.

Extracting Data

Another way to export all data or a subset of your data to a new data source is to use Tableau Extracts. For details, see **Extracting Data** on page 173.

Export as an Image

If you want to transfer your Tableau results into a presentation, report or web page, Tableau provides the following several options.

Copy to Another Application

To insert an image of your Tableau results into another application such as PowerPoint, Word, or Excel, follow these steps:

1. Select **Worksheet** > **Copy** > **Image**.
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. Click **Copy**. Tableau copies the current data view to the Clipboard.
4. Open the target application paste from the Clipboard.

   If the target application offers the option, 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 **Worksheet > Export > Image**.
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. Click **Save**.
4. 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.
5. Click **Save**.

Print to PDF

You can publish one or more views to PDF by selecting **File > Print to PDF**.

> When printing a sheet, quick filters are not included. To show quick filters, create a dashboard containing the sheet and print the dashboard to PDF. Note, when printing a dashboard to PDF, the contents of web page objects are not included.

Exporting the Data Source

After connecting to data, you can save the data source to your repository. Saving the data source creates a shortcut to the data and lets you avoid having to create a new connection every time you want to use that data source. Exporting the data source is useful if you have added custom fields such as groups or sets to the Data window. See **Exporting Data Sources** on page 149.

Copying Information Between Workbooks

You can combine resources from different workbooks into one, or pull information from a larger workbook into a standalone subset. To share or extract a subset of information, you can use any of the following actions on worksheets, dashboards, and stories:

- Copy and paste selected sheets to another workbook.
- Import a saved workbook into your current workbook.
- Export selected sheets to a new workbook.

What gets Copied or Saved with Selected Sheets

When you copy, save, or export selected sheets, other workbook items that the views in those sheets depend on are also included:

- Any associated visible or hidden sheets.
- Data sources for fields used in the selected sheets.
- Calculations, parameters, groups, sets, actions, custom shapes, and so on.

How Tableau Handles Duplicate Items

When you paste or import sheets from a different workbook, some items might already exist in the destination workbook, or some items might have the same name in both places. If Tableau encounters an exact duplicate item in the Data window, such as a calculation, it does not paste or import that item into the destination workbook. However, if an item in the Data window has the same name but is defined differently, Tableau imports and renames it.

Tableau also pastes or imports sheets and data sources with duplicate names—whether only the name is identical, or their names and contents are identical—and it renames the newer copy.

Export and Import Sheets

If you want to extract a subset of information from a larger workbook to maintain as a standalone file, you can export or save selected sheets to a new workbook. You can then import that workbook into an existing one to incorporate its sheets and other objects into the existing workbook.

Note: These steps describe how to share your work among Tableau workbooks. You can also export views to use outside of Tableau. For information, see **Exporting Your Work** on page 864.

1. Open the workbook that contains the sheets you want to export to a new file.
2. Using the tabbed worksheet or filmstrip view, right-click the sheet tabs or thumbnail views, and then select **Export Sheet** to export a single sheet. Select **File** > **Save As** to save all sheets.
3. Specify the file format you want to save (.twb or .twbx), select the location for the new workbook file, give it a name, and click **Save**.

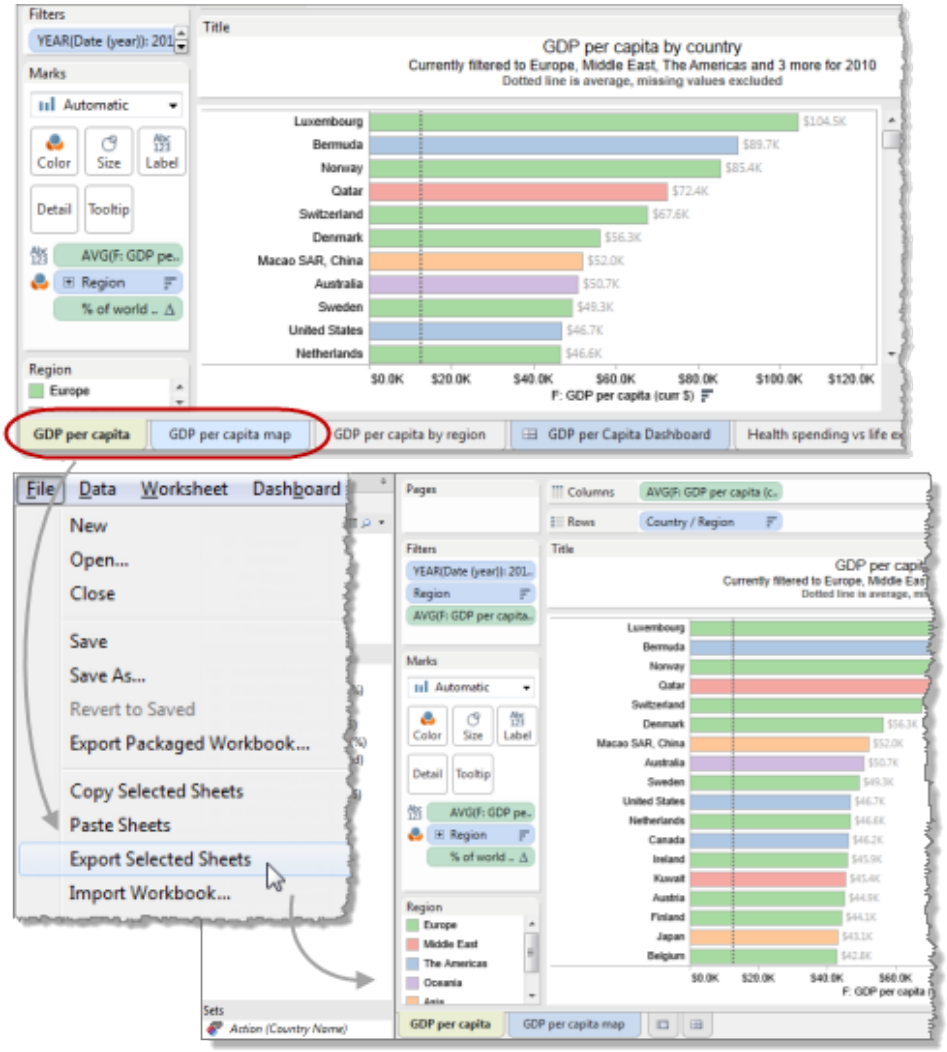

Import a Tableau Workbook

After you save or export selected sheets to a new workbook (.twb) file, you can import the information into another workbook.

1. With the existing workbook open, select **File** > **Import Workbook**.
2. Select the workbook that contains the sheets you saved from another workbook, and click **Open**.

The following image shows an import result. As shown, when the imported workbook contains a duplicate sheet name, Tableau adds a number after the name of the imported sheet.

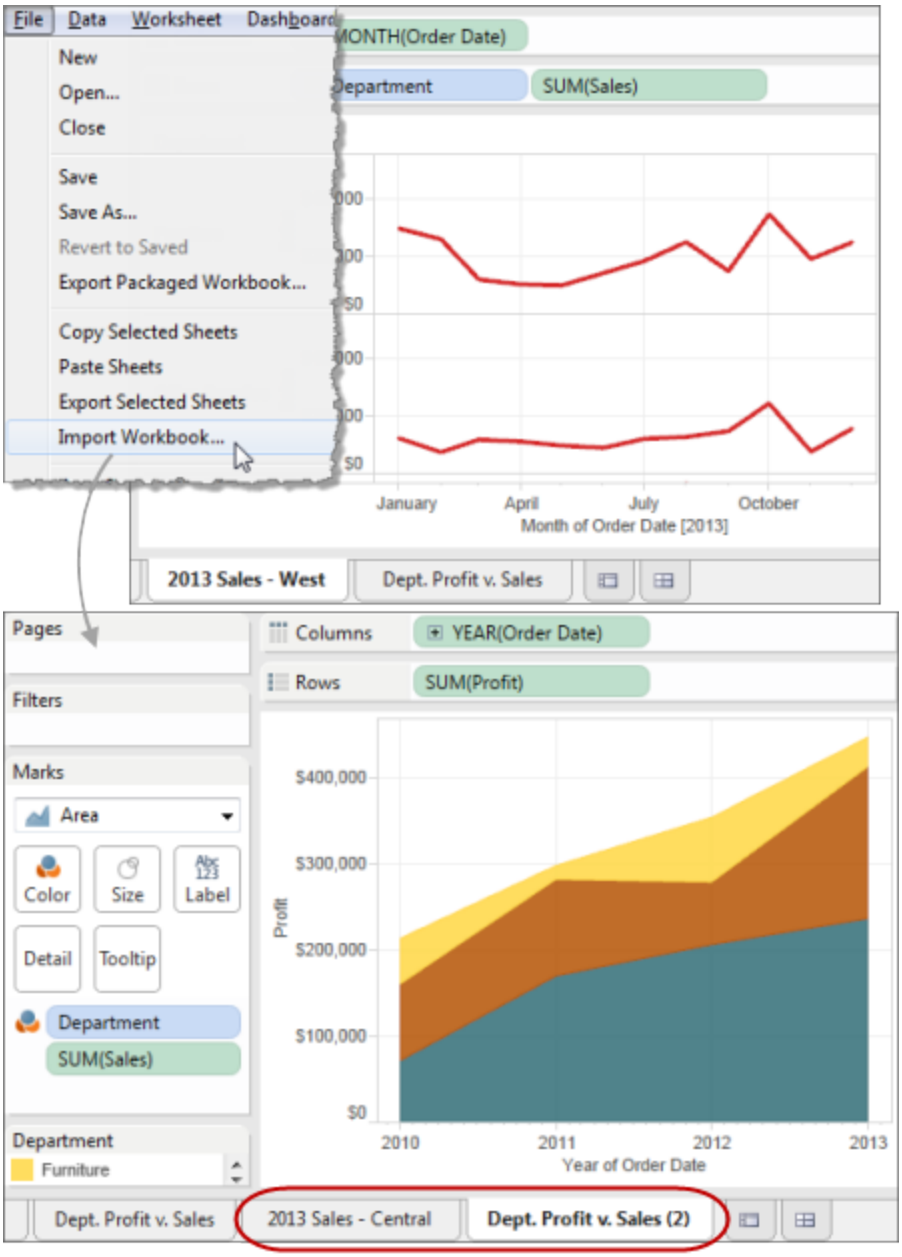

Copy and Paste Sheets

Copying and pasting sheets is a quick way to combine information from different workbooks or create a new workbook.

1. Open a workbook and display it in the filmstrip view.

2. Select the thumbnails of the sheets you want to copy, then right-click and select **Copy Sheet**.

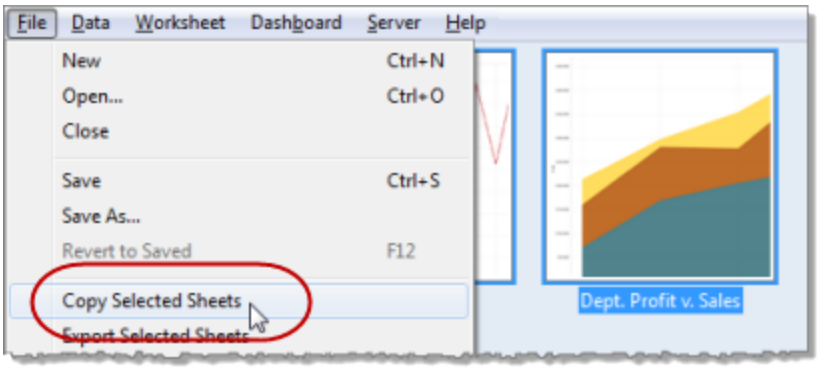

   Tableau copies the information in the file format (.twb or .twbx) of the workbook.

3. Open the destination workbook, or create a new workbook. Right-click a new sheet, and select **Paste Sheet**.

   Pasted sheets are placed after existing worksheets, dashboards, and stories.

4. Save the changes.

> **Note:** You can also display the workbook in the tabbed sheet view and select the tabs you want to copy.

For information about what gets copied when you copy and paste sheets, see **Copying Information Between Workbooks** on page 872
.

Printing

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 dialog box. 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 to a PDF.

Printing Results

Once you have a view or several views created in Tableau, you can print them or publish them as a PDF.

Page Setup

Before you print, 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, set the printed page orientation, specify where you want to put the legend, margins, and more. You define These settings 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**.

The Page Setup dialog box has the following tabs:

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 **Repeat headers and legends on each page**, the row and column headers will appear at the top of each printed page when a view breaks across several pages.

Select **Break pages on pane boundaries** to prevent page breaks in the middle of a cell in a 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.

Layout

Use the Layout tab in the Page Setup dialog box to specify the layout legend, page margins, and centering options for printing.

Legend Layout

If you include one or more legends, you can select one of the options for how you want the legends to appear on the printed page.

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--or both--on the page.

Print Scaling

Use the Print Scaling tab to fit the view to a certain size or to change the page orientation. These options only affect printed documents. The scaling options you specify here will not affect exported images or PDFs. However, the orientation settings will be used as the default when you publish the workbook to Tableau Server or Tableau Public.

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 original 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 three pages across by one page down.

Page Orientation

Use the page orientation settings to specify how you want the view oriented on the printed page. For example, if you have a view that is very wide but not very high, select **Landscape**. Select from the following page orientation options:

- Use Printer Setting – Use the page orientation that is already specified by the printer.
- Portrait – Presents the view so that it is oriented vertically on the printed page.
- Landscape – Presents the view so that it is oriented horizontally on the printed page.

The following diagram shows the difference between portrait and landscape page orientations.

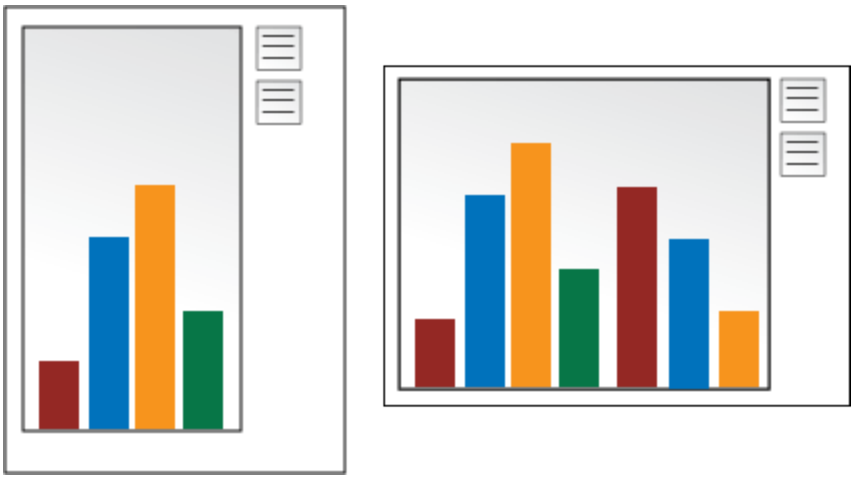

These page orientation settings are used as the default settings when you publish the workbook to Tableau Server or Tableau Public.

Printing

After you have configured the **Page Setup** on the previous page settings, 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.

The following options in the Print dialog box are unique to Tableau.

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 setup.

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.

> 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.

Print to PDF

In Tableau, you can publish views as PDF files rather than printing them as hard copies. You do not need to have Adobe Acrobat installed on your computer.

To publish as a PDF:

1. Specify **Page Setup** on page 876 options for each sheet in your workbook.
2. Select **File** > **Print to PDF**.

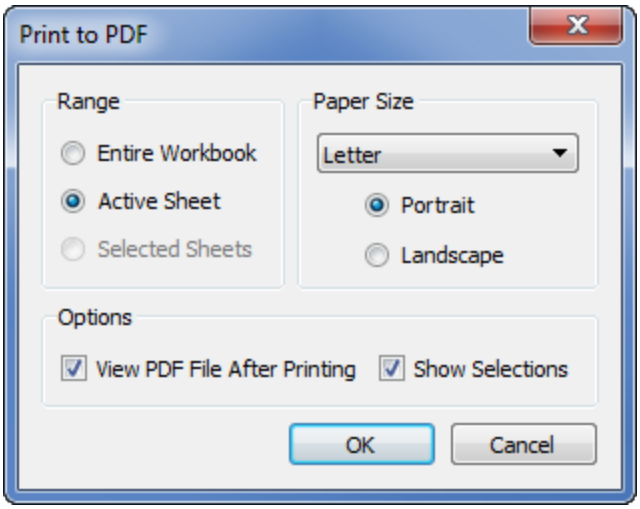

3. In the Print to PDF dialog box, select the print Range:
   - Entire Workbook - Publishes all the sheets in the workbook.
   - Active Sheet - Publishes only the sheet currently displayed in the workbook.
   - Selected Sheets - Publishes the selected sheets.
4. Select a **Paper Size**. If you select Unspecified, the paper size will expand to the necessary size to fit the entire view on a single page.
5. Select **View PDF File After Printing** if you want to automatically open the PDF after creating it. This option is only available if you have Adobe Acrobat Reader or Adobe Acrobat installed on your computer.
6. Select whether to **Show Selections**. When this option is selected the selections in the views are maintained in the PDF.

7. Click **OK** and specify where you want to save the PDF. Then click **Save**.

Printing the Help

You can print an individual Help topic using the print options in your Web browser.

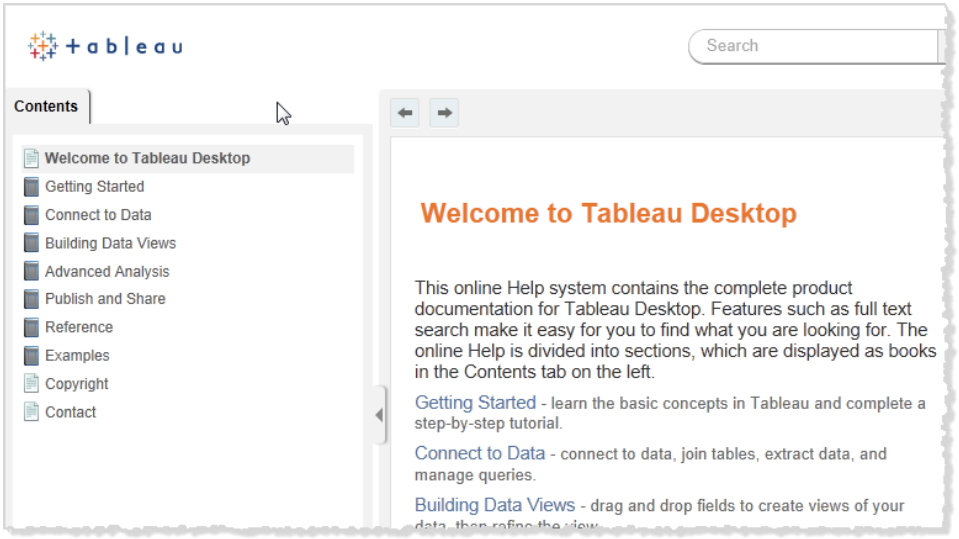

In addition to printing individual help topics, you can also download an offline help system and a printable PDF.

Reference

This section contains reference information for using Tableau. Learn how to use functions and operators when writing calculation formulas. Also, see tips and tricks that can help you become more efficient with the product.

Functions, Operators, & Data Types

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

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.

> 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 890 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 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.

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 then just type "May 1, 2005", which would be interpreted a 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 887). 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 categories. These are the same categories used in the **Calculation** dialog box. The aggregate functions such as sum, average, and so on are described in **Aggregations** on page 563.

Number Functions

ABS(number)

Returns the absolute value of the given number.

Examples

```
ABS(-7) = 7
```

```
ABS([Budget Variance])
```

The second example 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.

Example

```
ACOS(-1) = 3.14159265358979
```

ASIN(number)

Returns the arc sine of a given number. The result is in radians.

Example

```
ASIN(1) = 1.5707963267949
```

ATAN(number)

Returns the arc tangent of a given number. The result is in radians.

Example

```
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.

Example

```
ATAN2(2, 1) = 1.10714871779409
```

COS(number)

Returns the cosine of an angle. Specify the angle in radians.

Example

```
COS(PI( ) /4) = 0.707106781186548
```

COT(number)

Returns the cotangent of an angle. Specify the angle in radians.

Example

```
COT(PI( ) /4) = 1
```

DEGREES(number)

Converts a given number in radians to degrees.

Example

```
DEGREES(PI( )/4) = 45.0
```

EXP(number)

Returns e raised to the power of the given number.

Examples

```
EXP(2) = 7.389
```

```
EXP(-[Growth Rate]*[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.

Examples

```
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.

Examples

```
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.

Examples

```
POWER(5,2) = 52 = 25
```

```
POWER(Temperature, 2)
```

You can also use the ^ symbol:

```
5^2 = POWER(5,2) = 25
```

Radians (number)

Converts the given number from degrees to radians.

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 `decimals` is omitted, `number` is rounded to the nearest integer.

Example

Tjos example rounds every `Sales` value to an integer:

```
ROUND(7.3) = 7ROUND(-6.9) = -7ROUND(123.47,1) = 123.5ROUND(Sales)
```

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.

Example

If the average of the profit field is negative, then

```
SIGN(AVG(Profit)) = -1
```

SIN(number)

Returns the sine of an angle. Specify the angle in radians.

Example

```
SIN(0) = 1.0SIN(PI( )/4) = 0.707106781186548
```

SQRT(number)

Returns the square root of a number.

Example

```
SQRT(25) = 5
```

SQUARE(number)

Returns the square of a number.

Example

```
SQUARE(5) = 25
```

TAN(number)

Returns the tangent of an angle. Specify the angle in radians..

Example

```
TAN(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.

Example

```
ZN([Profit]) = [Profit]
```

String Functions

ASCII(string)

Return the ASCII code for the first character of `string`.

Example

```
ASCII('A') = 65
```

CHAR(number)

Returns the character encoded by the ASCII code `number`.

Example

```
CHAR(65) = 'A'
```

Contains(string, substring)

Returns true if the given string contains the specified substring.

Example

```
CONTAINS(“Calculation”, “alcu”) = true
```

ENDSWITH(string, substring)

Returns true if the given string ends with the specified substring. Trailing white spaces are ignored.

Example

```
ENDSWITH(“Tableau”, “leau”) = 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 ignores any instances of `substring` that appear before the index position `start`. The first character in the string is position 1.

Examples

```
FIND("Calculation", "alcu") = 2
FIND("Calculation", "Computer") = 0
FIND("Calculation", "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.

Example

```
LEFT("Matador", 4) = "Mata"
```

LEN(string)

Returns the length of the string.

Example

```
LEN("Matador") = 7
```

LOWER(string)

Returns `string`, with all characters lowercase.

Example

```
LOWER("ProductVersion") = "productversion"
```

LTRIM(string)

Returns the string with any leading spaces removed.

Example

```
LTRIM(" Matador ") = "Matador "
```

MAX(a, b)

Returns the maximum of `a` and `b` (which must be of the same type). This function is usually used to compare numbers, but also works on strings. 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`.

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.

Examples

```
MID("Calculation", 2) = "alculation"
```

```
MID("Calculation", 2, 5) ="alcul"
```

MIN(a, b)

Returns the minimum of `a` and `b` (which must be of the same type). This function is usually used to compare numbers, but also works on strings. With strings, `MIN` finds the value that is lowest in the sort sequence. It returns `Null` if either argument is `Null`.

Example

```
MIN ("Apple","Banana") = "Apple"
```

REPLACE(string, substring, replacement)

Searches `string` for `substring` and replaces it with `replacement`. If `substring` is not found, the string is not changed.

Example

```
REPLACE("Version8.5", "8.5", "9.0") = "Version9.0"
```

RIGHT(string, number)

Returns the right-most number of characters in `string`.

Example

```
RIGHT("Calculation", 4) = "tion"
```

RTRIM(string)

Returns `string` with any trailing spaces removed.

Example

```
RTRIM(" Calculation ") = " Calculation"
```

SPACE(number)

Returns a string that is composed of the specified `number` of repeated spaces.

Example

```
SPACE(1) = " "
```

STARTSWITH(string, substring)

Returns true if `string` stars with `substring`. Leading white spaces are ignored.

Example

```
STARTSWITH(“Joker”, “Jo”) = true
```

TRIM(string)

Returns the string with leading and trailing spaces removed. For example, `TRIM(" Calculation ") = "Calculation"`

UPPER(string)

Returns string, with all characters uppercase.

Example

```
UPPER("Calculation") = "CALCULATION"
```

Date Functions

Tableau provides a variety of date functions. Many of the examples use the `#` symbol with date expressions. See **Formatting Literals** on page 881 for an explanation of this symbol. Additionally, many date functions use `date_part`, which is a constant string argument. The valid `date_part` values that you can use are:

date_part	Values
`'year'`	Four-digit year
`'quarter'`	1-4
`'month'`	1-12 or "January", "February", and so on
`'dayofyear'`	Day of the year; Jan 1 is 1, Feb 1 is 32, and so on
`'day'`	1-31
`'weekday'`	1-7 or "Sunday", "Monday", and so on
`'week'`	1-52
`'hour'`	0-23

date_part	Values
'minute'	0-59
'second'	0-60

DATEADD(date_part, increment, date, start_of_week)

Returns an `increment` added to `date`. The type of increment is specified in `date_part`.

The `start_of_week` parameter is optional. If it is omitted, the start of week is determined by the data source. See **Date Properties for a Data Source** on page 438.

Example

```
DATEADD('month', 3, #2004-04-15#) = 2004-07-15 12:00:00 AM
```

This expression adds three months to the date `#2004-04-15#`.

DATEDIFF(date_part, date1, date2, start_of_week)

Returns the difference between `date1` and `date2` expressed in units of `date_part`.

The `start_of_week` parameter is optional. If it is omitted, the start of week is determined by the data source. See **Date Properties for a Data Source** on page 438.

Example

```
DATEDIFF('week', #2013-09-22#, #2013-09-24#, 'monday')= 1
```

```
DATEDIFF('week', #2013-09-22#, #2013-09-24#, 'sunday')= 0
```

The first expression returns 1 because when `start_of_week` is 'monday', then 22 September (a Sunday) and 24 September (a Tuesday) are in different weeks. The second expression returns 2 because when `start_of_week` is 'sunday' then 22 September (a Sunday) and 24 September (a Tuesday) are in the same week.

DATENAME(date_part, date, start_of_week)

Returns `date_part` of `date` as a string. The `start_of_week` parameter is optional.

If it is omitted, the start of week is determined by the data source. See **Date Properties for a Data Source** on page 438.

Examples

```
DATENAME('year', #2004-04-15#) = "2004"
```

```
DATENAME('month', #2004-04-15#) = "April"
```

DATEPARSE(format, string)

Converts a string to a datetime in the specified format. Support for some locale-specific formats is determined by the computer's system settings. Letters that appear in the data and do not need to be parsed should be surrounded by single quotes (' '). For formats that do not have delimiters between values (e.g., MMddyy), verify that they are parsed as expected. The format must be a constant string, not a field value. Returns `Null` if the data does not match the format.

This function is available for non-legacy Microsoft Excel and text file connections, MySQL, Oracle, PostgreSQL, and Tableau data extract data sources. Some formats may not be available for all data sources.

Examples

```
DATEPARSE ("dd.MMMM.yyyy", "15.April.2004") = #April 15, 2004#
```

```
DATEPARSE ("h'h' m'm' s's'", "10h 5m 3s") = #10:05:03#
```

Refer to the Understanding the DATEPARSE Function knowledge base article for more examples.

DATEPART(date_part, date, start_of_week)

Returns `date_part` of `date` as an integer.

The `start_of_week` parameter is optional. If it is omitted, the start of week is determined by the data source. See **Date Properties for a Data Source** on page 438.

> When the `date_part` is weekday, the `start_of_week` parameter is ignored. This is because Tableau relies on a fixed weekday ordering to apply offsets.

Examples

```
DATEPART('year', #2004-04-15#) = 2004
```

```
DATEPART('month', #2004-04-15#) = 4
```

DATETRUNC(date_part, date, start_of_week)

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. The `start_of_week` parameter is optional. If it is omitted, the start of week is determined by the data source. See **Date Properties for a Data Source** on page 438.

Examples

```
DATETRUNC('quarter', #2004-08-15#) = 2004-07-01 12:00:00 AM
```

```
DATETRUNC('month', #2004-04-15#) = 2004-04-01 12:00:00 AM
```

DAY(date)

Returns the day of the given date as an integer.

Example

```
DAY(#2004-04-12#) = 12
```

ISDATE(string)

Returns true if a given string is a valid date.

Example

```
ISDATE("April 15, 2004") = true
```

MAKEDATETIME(date, time)

Returns a datetime that combines a date and a time. The date can be a date, datetime, or a string type. The time must be a datetime. This function is available only for MySQL connections.

Examples

```
MAKEDATETIME("1899-12-30", #07:59:00#) = #12/30/1899 7:59:00 AM#
```

```
MAKEDATETIME([Date], [Time]) = #1/1/2001 6:00:00 AM#
```

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`.

Examples

```
MAX(#2004-01-01# ,#2004-03-01#) = 2004-03-01 12:00:00 AM
```

```
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`.

Examples

```
MIN(#2004-01-01# ,#2004-03-01#) = 2004-01-01 12:00:00 AM
```

```
MIN([ShipDate1], [ShipDate2])
```

MONTH(date)

Returns the month of the given date as an integer.

Example

```
MONTH(#2004-04-15#) = 4
```

NOW( )

Returns the current date and time.

Example

```
NOW( ) = 2004-04-15 1:08:21 PM
```

TODAY( )

Returns the current date.

Example

```
TODAY( ) = 2004-04-15
```

YEAR (date)

Returns the year of the given date as an integer.

Example

```
YEAR(#2004-15#) = 2004
```

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.

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.

Examples

```
DATE([Employee Start Date])
```

```
DATE("April 15, 2004") = #April 15, 2004#
```

```
DATE("4/15/2004")
```

```
DATE(#2006-06-15 14:52#) = #2006-06-15#
```

Quotation marks are required in the second and third examples.

DATETIME(expression)

Returns a datetime given a number, string, or date expression.

Example

```
DATETIME(“April 15, 2005 07:59:00”) = April 15, 2005 07:59:00
```

FLOAT(expression)

Casts its argument as a floating point number.

Examples

`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 results to the closest integer toward zero.

Examples

`INT(8.0/3.0) = 2`

`INT(4.0/1.5) = 2`

`INT(0.50/1.0) = 0`

`INT(-9.7) = -9`

When a string is converted to an integer it is first converted to a float and then rounded.

STR(expression)

Casts its argument as a string.

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

Use the CASE function to perform logical tests and return appropriate values. CASE is often easier to use than IIF or IF THEN ELSE. The CASE function evaluates `expression`, compares it to a sequence of values, `value1`, `value2`, etc., and returns a result. When a value that matches `expression` is encountered, CASE returns the corresponding return value. If no match is found, the default return expression is used. If there is no default return and no values match, then Null is returned.

Typically, you use an IF function to perform a sequence of arbitrary tests, and you use a CASE function to search for a match to an expression. But a CASE function can always be rewritten as an IF function , although the CASE function will generally be more concise.

Many times you can use a group to get the same results as a complicated case function.

Examples

```
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
```

If you need to include numeric comparisons in your conditions, use a nested IF clause. The CASE function compares strings only. For example, 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
```

IIF(test, then, else, [unknown])

Use the IIF function to perform logical tests and return appropriate values. The first argument, `test`, must be a boolean: either 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 IIF returns the `then` value. If `test` evaluates to FALSE, then IIF returns the `else` value.

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.

Examples

```
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(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

Use the IF THEN ELSE function to perform logical tests and return appropriate values. The IF THEN ELSE function 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: 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.

Examples

```
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

Use this version of the IF function to perform logical tests recursively. There is no built-in limit to the number of ELSEIFvalues you can use with an IF function, though individual databases may impose a limit on IF function complexity. While an IF function 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, whereas an IF statement only worries about TRUE and not true (which includes both FALSE and UNKNOWN).

Example

```
IF [Region]="West" THEN 1 ELSEIF [Region]="East" THEN 2 ELSE 3
END
```

IFNULL(expression1, expression2)

The IFNULL function returns the first expression if the result is not null, and returns the second expression if it is null.

Example

```
IFNULL([Proft], 0) = [Profit]
```

ISDATE(string)

The ISDATE function returns `TRUE` if the string argument can be converted to a date and `FALSE` if it cannot.

Examples

```
ISDATE("January 1, 2003") = TRUE
```

```
ISDATE("Jan 1 2003") = TRUE
```

```
ISDATE("1/1/03") = TRUE
```

```
ISDATE("Janxx 1 2003") = FALSE
```

ISNULL(expression)

The ISNULL function returns `TRUE` if the expression is `Null` and `FALSE` if it is not.

MIN(expression) or MIN(expression1,expression2)

The MIN function returns the minimum of an expression across all records or the minimum of two expressions for each record.

Aggregate Functions

> **Aggregations and floating-point arithmetic**: The results of some aggregations may not always be exactly as expected. For example, you may find that the Sum function returns a value such as -1.42e-14 for a column of numbers that you know should sum to exactly 0. This happens because the Institute of Electrical and Electronics Engineers (IEEE) 754 floating-point standard requires that numbers be stored in binary format, which means that numbers are sometimes rounded at extremely fine levels of precision. You can eliminate this potential distraction by using the ROUND function (see **Number Functions** on page 882) or by formatting the number to show fewer decimal places.

ATTR(expression)

Returns the value of the expression if it has a single value for all rows. Otherwise returns an asterisk. Null values are ignored.

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 in the following cases: workbooks created before Tableau Desktop 8.2 that use Microsoft Excel or text file data sources, workbooks that use the legacy connection, and workbooks that use Microsoft Access data sources. Extract your data into an extract file to use this function. See **Extracting Data** on page 173.

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.

MEDIAN(expression)

Returns the median of an expression across all records. Median can only be used with numeric fields. Null values are ignored. This function is not available in the following cases: workbooks that were created before Tableau Desktop 8.2 that use Microsoft Excel or text file data sources, workbooks that use the legacy connection, and workbooks that use Microsoft Access data sources.Extract your data into an extract file to use this function. See **Extracting Data** on page 173.

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.

PERCENTILE(expression, number)

Returns the percentile value from the given expression corresponding to the specified number.The number must be between 0 and 1 (inclusive)—for example, 0.66, and must be a

numeric constant. This function is available for non-legacy Microsoft Excel and text file connections, extracts, extract-only data sources (for example, Google Analytics, OData, and Salesforce), Sybase IQ 15.1 and later data sources, and Oracle 10 and later data sources. Extract your data into an extract file to use this function. See **Extracting Data** on page 173.

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)

These RAWSQL pass-through functions can be used to send SQL expressions directly to the database, without first being interpreted by Tableau. If you have custom database functions that Tableau doesn't know about, you can use these 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, if you had a function that computed the median of a set of values, you could call that function on the Tableau column [Sales] like this:

```
RAWSQLAGG_REAL("MEDIAN(%1)", [Sales])
```

Because Tableau does not interpret the expression, you must define the aggregation. You can use the RAWSQLAGG functions described below when you are using aggregated expressions.

RAWSQL pass-through functions will not work with published data sources.

These functions may return different results starting in Tableau Desktop 8.2 than they did in earlier versions of Tableau Desktop. This is because Tableau now uses ODBC for pass-through functions instead of OLE DB. ODBC truncates when returning real values as integer; OLE DB rounds when returning real values as integer.

RAWSQL Functions

The following RAWSQL functions are available in Tableau.

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 database. 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 database. 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 database. 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 database. 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

Use these user functions to create user filters based on user lists in your data source. For example, assume you have a view that shows the sales performance for each employee. When you publish that view, you may want to only allow employees to see their own sales numbers. You can use the function CURRENTUSER to create a field that returns True if the username of the person signed in to the server is the same as the employee name in the view. Then when you filter the view using this calculated field, only the data for the user who is currently signed in is shown.

FULLNAME( )

Returns the full name for the current user. This is the Tableau Server or Tableau Online full name when the user is signed in; otherwise the local or network full name for the Tableau Desktop user.

Example

```
[Manager]=FULLNAME( )
```

If manager Dave Hallsten is signed in, this example returns True only if the Manager field in the view contained 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 signed in to the server.

ISFULLNAME(string)

Returns true if the current user's full name matches the specified full name, or false if it does not match. This function uses the Tableau Server or Online full name when the user is signed in; otherwise it uses the local or network full name for the Tableau Desktop user.

Example

```
ISFULLNAME("Dave Hallsten")
```

This example 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 signed in, the group membership is determined by groups on Tableau Server or Tableau Online. If the person is not signed in, this function returns false.

Example

```
IF ISMEMBEROF("Sales") THEN "Sales" ELSE "Other" END
```

ISUSERNAME(string)

Returns true if the current user's username matches the specified username, or false if it does not match. This function uses the Tableau Server or Online username when the user is signed in; otherwise it uses the local or network username for the Tableau Desktop user.

Example

```
ISUSERNAME("dhallsten")
```

This example returns true if dhallsten is the current user; otherwise it returns false.

USERDOMAIN()

Returns the domain for the current user when the user is signed on to Tableau Server. Returns the Windows domain if the Tableau Desktop user is on a domain. Otherwise this function returns a null string.

Example

```
[Manager]=USERNAME() AND [Domain]=USERDOMAIN()
```

USERNAME( )

Returns the username for the current user. This is the Tableau Server or Tableau Online username when the user is signed in; otherwise it is the local or network username for the Tableau Desktop user.

Example

```
[Manager]=USERNAME( )
```

If the manager dhallsten was signed in, this function would only return True when the Manager field in the view is 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 signed in to the server.

Table Calculation Functions

Use table calculation functions 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. See **Customizing Table Calculations** on page 621 for information 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	First()
2009	Q1	$160,877	$231,411	$133,934	$185,961	0
	Q2	$197,213	$204,914	$337,813	$213,507	-1
	Q3	$302,678	$165,201	$283,806	$206,512	-2
	Q4	$297,208	$226,983	$214,845	$230,291	-3
2010	Q1	$180,609	$180,123	$273,943	$251,145	-4
	Q2	$195,785	$224,882	$251,391	$195,976	-5
	Q3	$116,613	$50,363	$194,601	$102,731	-6

Example

When the current row index is 3, `FIRST() = -2`.

INDEX( )

Returns the index of the current row in the partition, without any sorting with regard to value. 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

	INDEX()
$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 expression in a target row, specified as a relative offset from the current row. Use FIRST() + n and LAST() - n as part of your offset definition for a target relative to the first/last rows in the partition. If `offset` is omitted, the row to compare to can be set on the field menu. This function returns NULL if the target row cannot be determined.

The view below shows quarterly sales. When `LOOKUP (SUM(Sales), 2)` is computed within the Date partition, each row shows the sales value from 2 quarters into 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 +2	$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)` computes the 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)` computes the running product of SUM(Profit).

RANK(expression, ['asc' | 'desc'])

Returns the standard competition rank for the current row in the partition. Identical values are assigned an identical rank. Use the optional `'asc' | 'desc'` argument to specify ascending or descending order. The default is descending.

With this function, the set of values (6, 9, 9, 14) would be ranked (1, 2, 2, 4).

Nulls are ignored in ranking functions. They are not numbered and they do not count against the total number of records in percentile rank calculations.

For information on different ranking options, see **Rank Calculation** on page 611.

Example

The following image shows the effect of the various ranking functions (RANK, RANK_DENSE, RANK_MODIFIED, RANK_PERCENTILE, and RANK_UNIQUE) on a set of values. The data set contains information on 14 students (StudentA through StudentN); the **Age** column shows the current age of each student (all students are between 17 and 20 years of age). The remaining columns show the effect of each rank function on the set of age values.

Student	Age	RANKofAge	RANK_DENSEofAge	RANK_MODIFIEDofAge	RANK_PERCENTILEofAge	RANK_UNIQUEofAge
StudentA	19	4	2	7	79%	4
StudentB	18	8	3	12	50%	8
StudentC	19	4	2	7	79%	5
StudentD	18	8	3	12	50%	9
StudentE	17	13	4	14	14%	13
StudentF	18	8	3	12	50%	10
StudentG	19	4	2	7	79%	6
StudentH	20	1	1	3	100%	1
StudentI	19	4	2	7	79%	7
StudentJ	20	1	1	3	100%	2
StudentK	20	1	1	3	100%	3
StudentL	17	13	4	14	14%	14
StudentM	18	8	3	12	50%	11
StudentN	18	8	3	12	50%	12

RANK_DENSE(expression, ['asc' | 'desc'])

Returns the dense rank for the current row in the partition. Identical values are assigned an identical rank, but no gaps are inserted into the number sequence. Use the optional `'asc' | 'desc'` argument to specify ascending or descending order. The default is descending.

With this function, the set of values (6, 9, 9, 14) would be ranked (1, 2, 2, 3).

Nulls are ignored in ranking functions. They are not numbered and they do not count against the total number of records in percentile rank calculations.

For information on different ranking options, see **Rank Calculation** on page 611.

RANK_MODIFIED(expression, ['asc' | 'desc'])

Returns the modified competition rank for the current row in the partition. Identical values are assigned an identical rank. Use the optional `'asc' | 'desc'` argument to specify ascending or descending order. The default is descending.

With this function, the set of values (6, 9, 9, 14) would be ranked (1, 3, 3, 4).

Nulls are ignored in ranking functions. They are not numbered and they do not count against the total number of records in percentile rank calculations.

For information on different ranking options, see **Rank Calculation** on page 611.

RANK_PERCENTILE(expression, ['asc' | 'desc'])

Returns the percentile rank for the current row in the partition. Use the optional `'asc' | 'desc'` argument to specify ascending or descending order. The default is ascending.

With this function, the set of values (6, 9, 9, 14) would be ranked (25, 75, 75, 100).

Nulls are ignored in ranking functions. They are not numbered and they do not count against the total number of records in percentile rank calculations.

For information on different ranking options, see **Rank Calculation** on page 611.

RANK_UNIQUE(expression, ['asc' | 'desc'])

Returns the unique rank for the current row in the partition. Identical values are assigned different ranks. Use the optional `'asc' | 'desc'` argument to specify ascending or

descending order. The default is descending.

With this function, the set of values (6, 9, 9, 14) would be ranked (1, 2, 3, 4).

Nulls are ignored in ranking functions. They are not numbered and they do not count against the total number of records in percentile rank calculations.

For information on different ranking options, see **Rank Calculation** on page 611.

RUNNING_AVG(expression)

Returns the running average of the given expression, from the first row in the partition to the current row.

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.

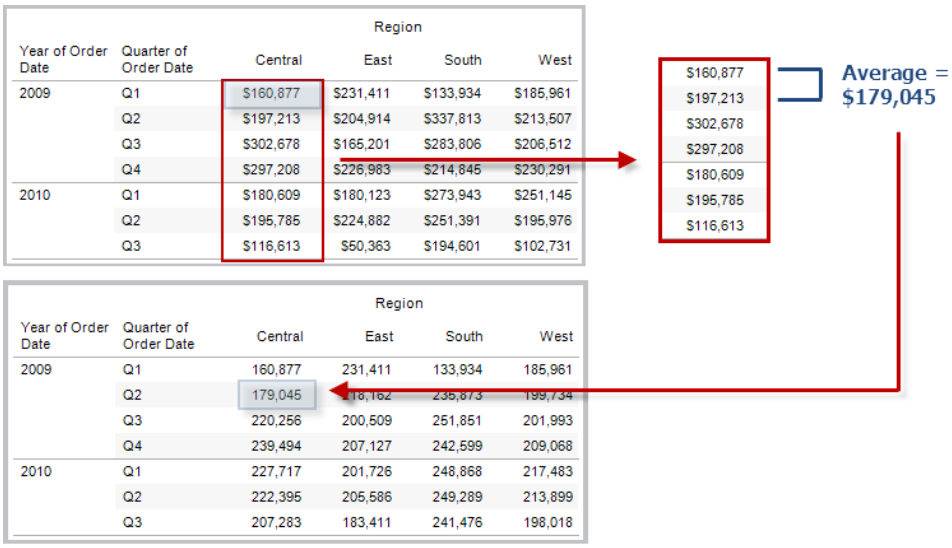

Example

`RUNNING_AVG(SUM([Profit]))` computes the 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]))` computes the 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.

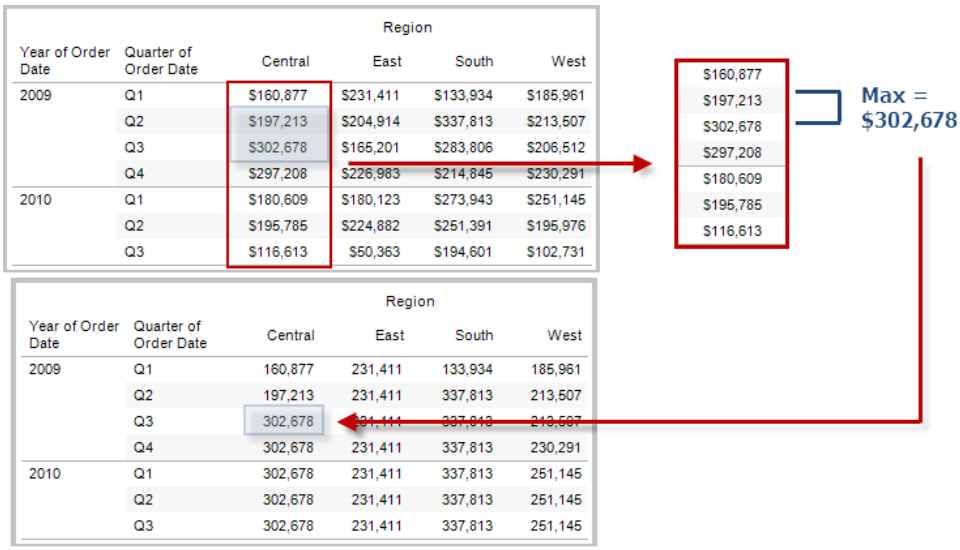

Example

`RUNNING_MAX(SUM([Profit]))` computes the 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.

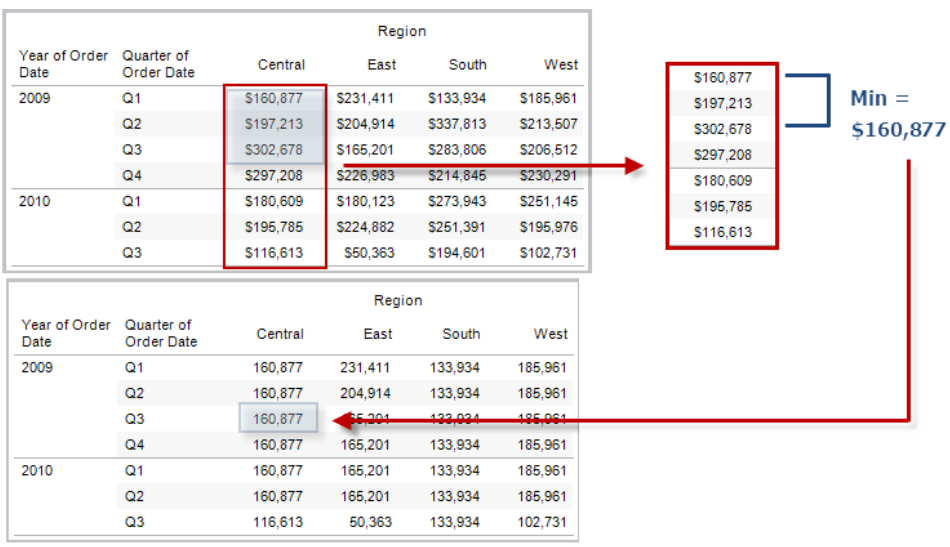

Example

`RUNNING_MIN(SUM([Profit]))` computes the running minimum of SUM(Profit).

RUNNING_SUM(expression)

Returns the running sum of the given expression, from the first row in the partition to the current row.

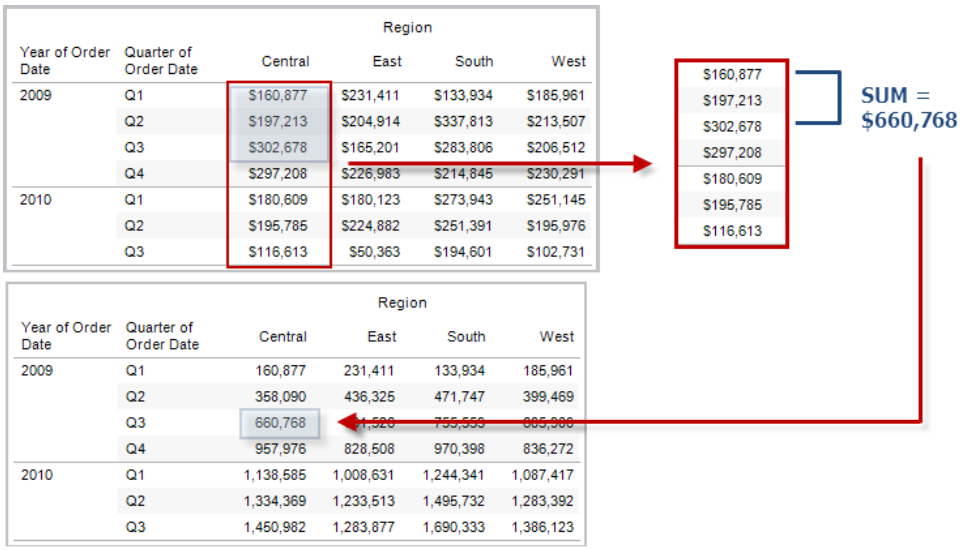

Example

`RUNNING_SUM(SUM([Profit]))` computes the 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.

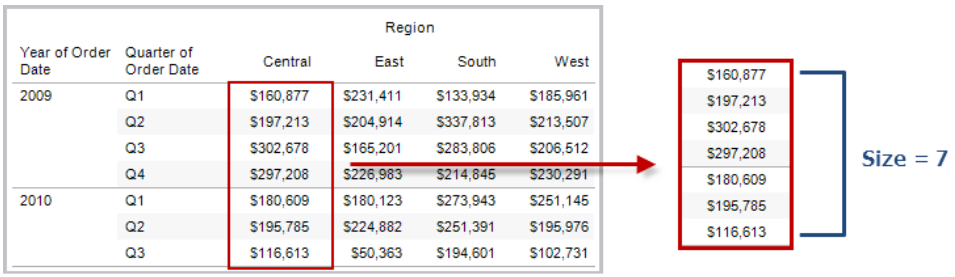

Example

`SIZE() = 5` when the current partition contains five rows.

SCRIPT_BOOL

Returns a Boolean result from the specified R expression. The R expression is passed directly to a running Rserve instance. Use .arg*n* in the R expression to reference parameters (.arg1, .arg2, etc.).

Examples

In the following example, .arg1 is equal to SUM([Profit]):

```
SCRIPT_BOOL("is.finite(.arg1)", SUM([Profit]))
```

The next example returns True for store IDs in Washington state, and False otherwise. This example could be the definition for a calculated field titled IsStoreInWA.

```
SCRIPT_BOOL('grepl(".*_WA", .arg1, perl=TRUE)',ATTR([Store ID]))
```

SCRIPT_INT

Returns an integer result from the specified R expression. The R expression is passed directly to a running Rserve instance. Use .arg*n* in the R expression to reference parameters (.arg1, .arg2, etc.).

Examples

In the following example, .arg1 is equal to SUM([Profit]):

```
SCRIPT_INT("is.finite(.arg1)", SUM([Profit]))
```

In the next example, k-means clustering is used to create three clusters:

```
SCRIPT_INT('result <- kmeans(data.frame(.arg1,.arg2,.arg3,.arg4),
3);result$cluster;', SUM([Petal length]), SUM([Petal width]),SUM
([Sepal length]),SUM([Sepal width]))
```

SCRIPT_REAL

Returns a real result from the specified R expression. The R expression is passed directly to a running Rserve instance. Use .arg*n* in the R expression to reference parameters (.arg1, .arg2, etc.).

Examples

In the following example, .arg1 is equal to SUM([Profit]):

```
SCRIPT_REAL("is.finite(.arg1)", SUM([Profit]))
```

The next example converts temperature values from Celsius to Fahrenheit.

```
SCRIPT_REAL('library(udunits2);ud.convert(.arg1, "celsius",
"degree_fahrenheit")',AVG([Temperature]))
```

SCRIPT_STR

Returns a string result from the specified R expression. The R expression is passed directly to a running Rserve instance. Use .arg*n* in the R expression to reference parameters (.arg1, .arg2, etc.).

Examples

In the following example, .arg1 is equal to SUM([Profit]):

```
SCRIPT_STR("is.finite(.arg1)", SUM([Profit]))
```

The next example extracts a state abbreviation from a more complicated string (in the original form 13XSL_CA, A13_WA):

```
SCRIPT_STR('gsub(".*_", "", .arg1)',ATTR([Store ID]))
```

TOTAL(expression)

Returns the total for the given expression. For example, the view below shows quarterly sales. The total within the Date partition is the SUM([Sales]) across all dates.

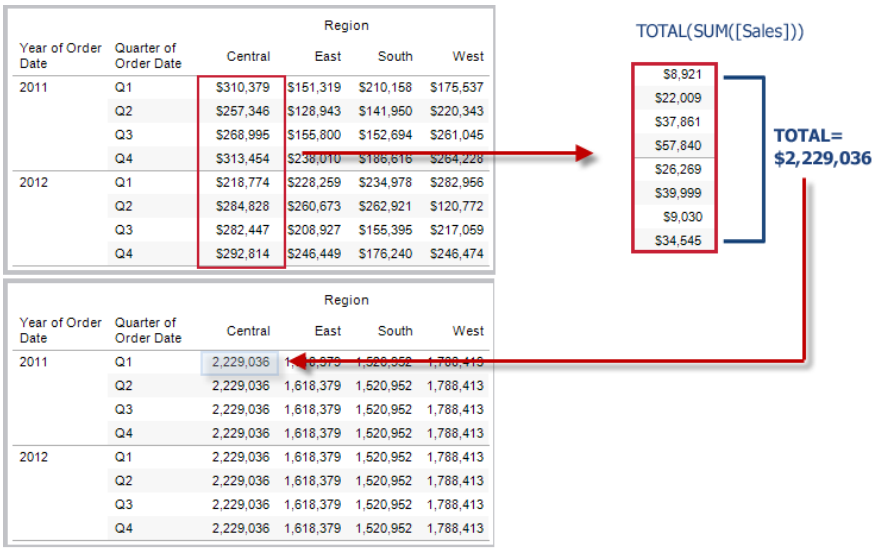

Example

`TOTAL(SUM([Profit]))` computes SUM(Profit) across all database rows in the partition.

WINDOW_AVG(expression, [start, end])

Returns the average of the expression within the window. The window is defined by means of 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.

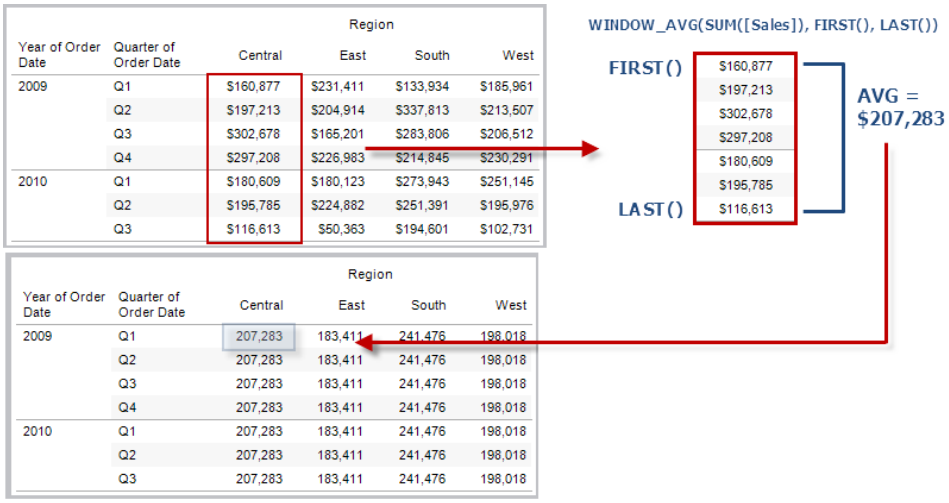

Example

`WINDOW_AVG(SUM[Profit]), FIRST()+1, 0)` computes the 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 by means of 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)` computes the count of SUM(Profit) from the second row to the current row

WINDOW_MEDIAN(expression, [start, end])

Returns the median of the expression within the window. The window is defined by means of 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 profit. A window median within the Date partition returns the median profit across all dates.

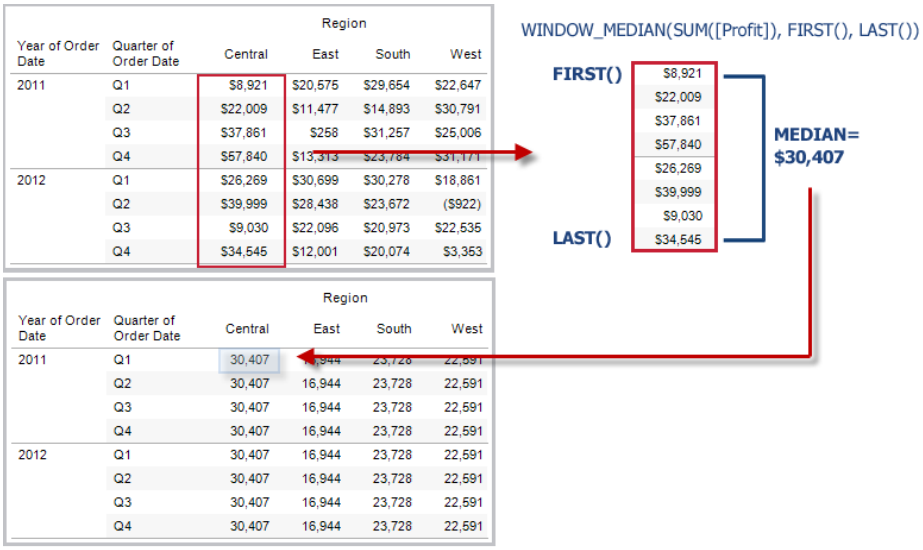

Year of Order Date	Quarter of Order Date	Region: Central	East	South	West
2011	Q1	$8,921	$20,575	$29,654	$22,647
	Q2	$22,009	$11,477	$14,893	$30,791
	Q3	$37,861	$258	$31,257	$25,006
	Q4	$57,840	$13,313	$23,784	$31,171
2012	Q1	$26,269	$30,699	$30,278	$18,861
	Q2	$39,999	$28,438	$23,672	($922)
	Q3	$9,030	$22,096	$20,973	$22,535
	Q4	$34,545	$12,001	$20,074	$3,353

Year of Order Date	Quarter of Order Date	Region: Central	East	South	West
2011	Q1	30,407	16,944	23,728	22,591
	Q2	30,407	16,944	23,728	22,591
	Q3	30,407	16,944	23,728	22,591
	Q4	30,407	16,944	23,728	22,591
2012	Q1	30,407	16,944	23,728	22,591
	Q2	30,407	16,944	23,728	22,591
	Q3	30,407	16,944	23,728	22,591
	Q4	30,407	16,944	23,728	22,591

Example

`WINDOW_MEDIAN(SUM[Profit]), FIRST()+1, 0)` computes the median 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 by means of 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.

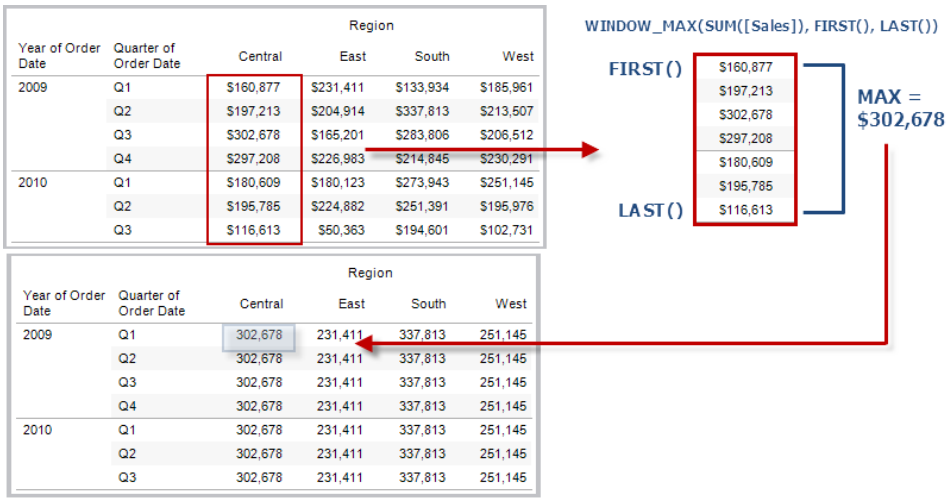

Example

`WINDOW_MAX(SUM[Profit]), FIRST()+1, 0)` computes the 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 by means of 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.

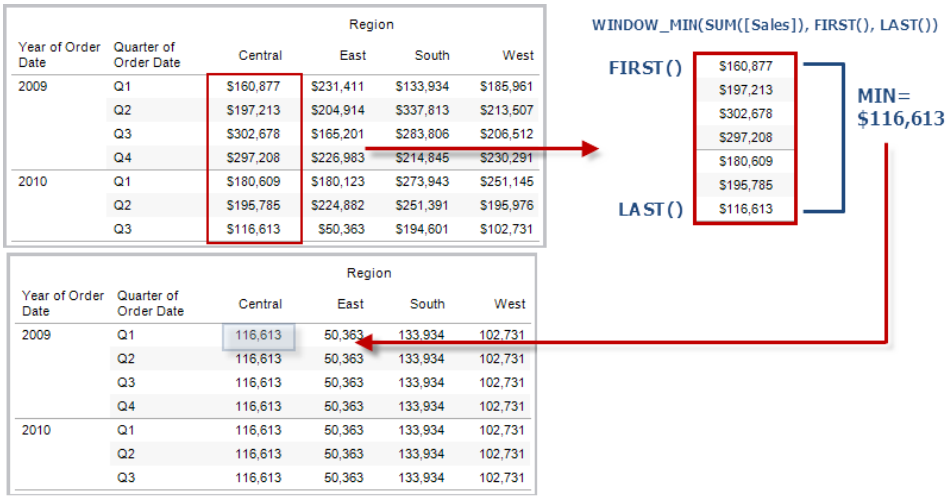

Example

`WINDOW_MIN(SUM[Profit]), FIRST()+1, 0)` computes the minimum of SUM(Profit) from the second row to the current row.

WINDOW_PERCENTILE(expression, number, [start, end])

Returns the value corresponding to the specified percentile within the window. The window is defined by means of 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_PERCENTILE(SUM([Profit]), 0.75, -2, 0))` returns the 75th percentile for SUM(Profit) from the two previous rows to the current row.

WINDOW_STDEV(expression, [start, end])

Returns the sample standard deviation of the expression within the window. The window is defined by means of 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_STDEV(SUM[Profit]), FIRST()+1, 0)` computes the 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 by means of 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_STDEVP(SUM[Profit]), FIRST()+1, 0)` computes the standard deviation of SUM(Profit) from the second row to the current row.

WINDOW_SUM(expression, [start, end])

Returns the sum of the expression within the window. The window is defined by means of 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.

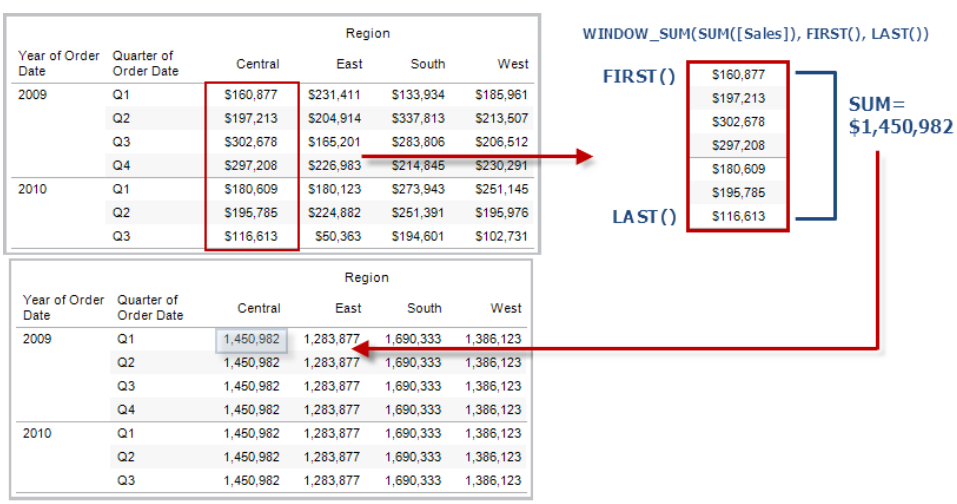

Year of Order Date	Quarter of Order Date	Region: Central	East	South	West
2009	Q1	$160,877	$231,411	$133,934	$185,961
	Q2	$197,213	$204,914	$337,813	$213,507
	Q3	$302,678	$165,201	$283,806	$206,512
	Q4	$297,208	$226,983	$214,845	$230,291
2010	Q1	$180,609	$180,123	$273,943	$251,145
	Q2	$195,785	$224,882	$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	1,450,982	1,283,877	1,690,333	1,386,123
	Q2	1,450,982	1,283,877	1,690,333	1,386,123
	Q3	1,450,982	1,283,877	1,690,333	1,386,123
	Q4	1,450,982	1,283,877	1,690,333	1,386,123
2010	Q1	1,450,982	1,283,877	1,690,333	1,386,123
	Q2	1,450,982	1,283,877	1,690,333	1,386,123
	Q3	1,450,982	1,283,877	1,690,333	1,386,123

Example

`WINDOW_SUM(SUM[Profit]), FIRST()+1, 0)` computes the sum 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 by means of 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)` computes the 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 by means of 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)` computes the variance of SUM (Profit) from the second row to the current row.

Additional Functions

CEILING(number)

Rounds a number to the nearest integer of equal or greater value. This function is available for Hadoop Hive and Google BigQuery data sources.

Example

CEILING(3.1415) = 4

FLOOR(number)

Rounds a number to the nearest integer of equal or lesser value. This function is available for Hadoop Hive and Google BigQuery data sources.

Example

FLOOR(3.1415) = 3

REGEXP_REPLACE(string, pattern, replacement)

Returns a copy of the given string where the matching pattern is substituted with the replacement string. This function is available for Hadoop Hive and Google BigQuery data sources. For more information on regular expression syntax, see your data source's documentation.

Example

REGEXP_REPLACE('abc 123', '\s', '-') = 'abc--123'

Hadoop Hive Specific Functions

GET_JSON_OBJECT(JSON string, JSON path)

Returns the JSON object within the JSON string based on the JSON path.

PARSE_URL(string, url_part)

Returns a component of the given URL string where the component is defined by url_part. Valid url_part values include: 'HOST', 'PATH', 'QUERY', 'REF', 'PROTOCOL', 'AUTHORITY', 'FILE' and 'USERINFO'.

Example

PARSE_URL('http://www.tableausoftware.com', 'HOST') = 'www.tableausoftware.com'

PARSE_URL_QUERY(string, key)

Returns the value of the specified query parameter in the given URL string. The query parameter is defined by the key.

Example

PARSE_URL_QUERY('http://www.tableausoftware.com?page=1&cat=4', 'page') = '1'

REGEXP_EXTRACT_NTH(string, pattern)

Returns the portion of a string of the given string using the regular expression pattern. The substring is matched to the nth capturing group, where n is the given index. For more information on regular expression syntax, see your data source's documentation.

Example

REGEXP_EXTRACT('abc 123', '[a-z]+\s+(\d+)') = '123'

XPATH_BOOLEAN(XML string, XPath expression string)

Returns true if the XPath expression matches a node or evaluates to true.

Example

XPATH_BOOLEAN('<values> <value id="0">1</value><value id="1">5</value>', 'values/value[@id="1"] = 5') = true

XPATH_DOUBLE(XML string, XPath expression string)

Returns the floating-point value of the XPath expression.

Example

XPATH_DOUBLE('<values><value>1.0</value><value>5.5</value> </values>', 'sum (value/*)') = 6.5

XPATH_FLOAT(XML string, XPath expression string)

Returns the floating-point value of the XPath expression.

Example

XPATH_FLOAT('<values><value>1.0</value><value>5.5</value> </values>','sum(value/*) ') = 6.5

XPATH_INT(XML string, XPath expression string)

Returns the numerical value of the XPath expression, or zero if the XPath expression cannot evaluate to a number.

Example

XPATH_INT('<values><value>1</value><value>5</value> </values>','sum(value/*)') = 6

XPATH_LONG(XML string, XPath expression string)

Returns the numerical value of the XPath expression, or zero if the XPath expression cannot evaluate to a number.

Example

XPATH_LONG('<values><value>1</value><value>5</value> </values>','sum(value/*)') = 6

XPATH_SHORT(XML string, XPath expression string)

Returns the numerical value of the XPath expression, or zero if the XPath expression cannot evaluate to a number.

Example

XPATH_SHORT('<values><value>1</value><value>5</value> </values>','sum(value/*)') = 6

XPATH_STRING(XML string, XPath expression string)

Returns the text of the first matching node.

Example

XPATH_STRING('<sites ><url domain="org">http://www.w3.org</url> <url domain="com">http://www.tableausoftware.com</url></sites>', 'sites/url[@domain="com"]') = 'http://www.tableausoftware.com'

Google BigQuery Specific Functions

DOMAIN(string_url)

Given a URL string, returns the domain as a string.

Example

DOMAIN('http://www.google.com:80/index.html') = 'google.com'

GROUP_CONCAT(expression)

Concatenates values from each record into a single comma-delimited string. This function acts like a SUM() for strings.

Example

GROUP_CONCAT(Region) = "Central,East,West"

HOST(string_url)

Given a URL string, returns the host name as a string.

Example

HOST('http://www.google.com:80/index.html') = 'www.google.com:80'

LOG2(number)

Returns the logarithm base 2 of a number.

Example

LOG2(16) = '4.00'

LTRIM_THIS(string, string)

Returns the first string with any leading occurrence of the second string removed.

Example

LTRIM_THIS('[-Sales-]','[-') = 'Sales-]'

REGEXP_MATCH(string, pattern)

Returns true if a substring of the provided string matches the regular expression pattern. For more information on regular expression syntax, see your data source's documentation.

Example

REGEXP_MATCH('-([1234].[The.Market])-','\[\s*(\w*\.)(\w*\s*\])')=true

RTRIM_THIS(string, string)

Returns the first string with any trailing occurrence of the second string removed.

Example

RTRIM_THIS('[-Market-]','-]') = '[-Market'

TIMESTAMP_TO_USEC(expression)

Converts a TIMESTAMP data type to a UNIX timestamp in microseconds.

Example

TIMESTAMP_TO_USEC(#2012-10-01 01:02:03#)=1349053323000000

USEC_TO_TIMESTAMP(expression)

Converts a UNIX timestamp in microsseconds to a TIMESTAMP data type.

Example

USEC_TO_TIMESTAMP(1349053323000000) = #2012-10-01 01:02:03#

TLD(string_url)

Given a URL string, returns the top level domain plus any country domain in the URL.

Example

TLD('http://www.google.com:80/index.html') = '.com'

TLD('http://www.google.co.uk:80/index.html') = '.co.uk'

Operators

To create calculated fields and formulas, you need to understand the operators supported by Tableau. This section discusses the basic operators that are available and the order (precedence) of operations.

+ (addition)

This means addition when applied to numbers and concatenation when applied to strings. When applied to dates, it can be used to add a number of days to a date. For example,

```
7 + 3
Profit + Sales
'abc' + 'def' = 'abcdef'
#April 15, 2004# + 15 = #April 30, 2004#
```

– (subtraction)

This means subtraction when applied to numbers and negation if applied to an expression. When applied to dates, it can be used to subtract a number of days from a date. Hence it can also be used to calculate the difference in days between two dates. For example,

```
7 - 3
Profit - Sales
-(7+3) = -10
#April 16, 2004# - 15 = #April 1, 2004#
#April 15, 2004# - #April 8, 2004# = 7
```

* (multiplication)

This means numeric multiplication. For example, `5 * 4 = 20`.

/ (division)

This means numeric division. For example, `20 / 4 = 5`.

% (modulo)

This calculates a numeric remainder. For example, `5% 4 = 1`.

= =, =, >, <, >=, <=, !=, <>(comparisons)

These are the basic comparison operators that can be used in expressions. Their meanings are as follows: = = or =(equal to), > (greater than), < (less than), >= (greater than or equal to), <= (less than or equal to), != and <> (not equal to).

Each operator compares two numbers, dates, or strings and returns a boolean (`TRUE` or `FALSE`). Booleans themselves, however, cannot be compared using these operators. For example, `TRUE=TRUE` is not a valid expression. To compare booleans in this way, use the logical operators `AND` and `OR`. For example, `TRUE AND TRUE` is a valid expression.

^ (power)

This symbol is equivalent to the POWER function. It raises a number to the specified power.

For example:

```
6^3 = 216
```

AND

This is a logical operator. An expression or a boolean must appear on either side of it. For example,

```
IIF(Profit =100 AND Sales =1000, "High", "Low")
```

If both expressions are `TRUE` (i.e., not `FALSE` and not `UNKNOWN`), then the result is `TRUE`. If either expression is `UNKNOWN`, then the result is `UNKNOWN`. In all other cases, the result is `FALSE`.

If you create a calculation in which the result of an AND comparison is displayed on a worksheet, Tableau displays TRUE and FALSE. If you would like to change this, use the Format area in the format dialog.

OR

This is a logical operator. An expression or a boolean must appear on either side of it. For example,

```
IIF(Profit =100 OR Sales =1000, "High", "Low")
```

If either expression is `TRUE`, then the result is `TRUE`. If both expressions are `FALSE`, then the result is `FALSE`. If both expressions are `UNKNOWN`, then the result is `UNKNOWN`.

If you create a calculation in which the result of an `OR` comparison is displayed on a worksheet, Tableau displays TRUE and FALSE. If you would like to change this, use the Format area in the format dialog. The `OR` operator employs "short circuit evaluation." This means that if the first expression is evaluated to be `TRUE`, then the second expression is not evaluated at all. This can be helpful if the second expression results in an error when the first expression is `TRUE`, because the second expression in this case is never evaluated.

NOT

This is a logical operator. It can be used to negate another boolean or an expression. For example,

```
IIF(NOT(Sales = Profit),"Not Equal","Equal")
```

Precedence

All operators are evaluated in a specific order. For example, 2*1+2 is equal to 4 and not equal to 6. The reason is that the * operator is always evaluated before the + operator.

The following table shows the order in which operators are evaluated. The first line has the highest precedence. Operators on the same line have the same precedence. If two operators have the same precedence they are evaluated from left to right in the formula.

Precedence	Operator
1	– (negate)
2	^ (power)
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

This section includes information on creating performance recordings, and suggestions for improving Tableau's performance. 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 best practice, use Tableau with databases that are suitable for real-time querying and analysis.

Create a Performance Recording

With the Performance Recording feature in Tableau, you can record performance information about key events as you interact with workbooks. You then view performance metrics in a performance workbook that Tableau creates automatically. The steps you follow to create and view performance recording vary somewhat between Tableau Desktop and Tableau Server. However, the resulting performance workbooks have the same format in both Tableau Desktop and Tableau Server.

Use performance workbooks to analyze and troubleshoot performance issues pertaining to different events that are known to affect performance, including:

- Query execution
- Geocoding
- Connections to data sources
- Layout computations
- Extract generation
- Blending data
- Server blending (Tableau Server only)

Tableau support may request that you create performance workbooks as they assist you with diagnosing performance issues.

To create a performance recording in Tableau Desktop

To start recording performance, follow this step in Tableau Desktop:

Help > Settings and Performance > Start Performance Recording

To stop recording, and then view a temporary workbook containing results from the recording session, follow this step:

Help > Settings and Performance > Stop Performance Recording

You can then save this workbook as a packaged workbook (.twbx) file, and send it to Tableau support.

Interpret a Performance Recording Workbook

A performance recording workbook is a Tableau dashboard that contains three views: **Timeline**, **Events**, and **Query**.

For information on how to create a performance recording in Tableau Desktop, see **Create a Performance Recording** on the previous page.

Timeline

The uppermost view in a performance recording dashboard shows the events that occurred during recording, arranged chronologically from left to right. The bottom axis shows elapsed time since Tableau started, in seconds.

In the Timeline view, the **Workbook**, **Dashboard**, and **Worksheet** columns identify the context for events. The **Event** column identifies the nature of the event, and the final column shows each event's duration and how it compares chronologically to other recorded events:

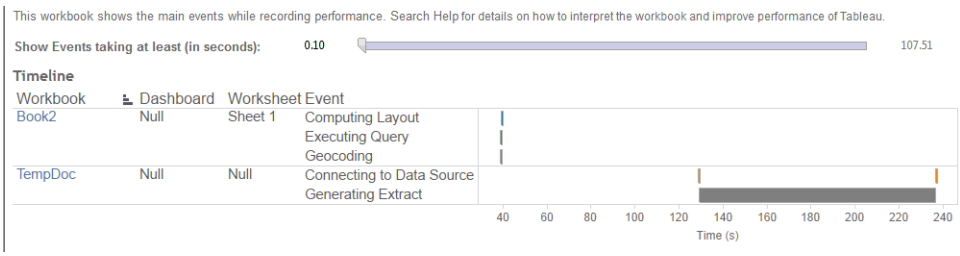

Events

The middle view in a performance recording workbook shows the events, sorted by duration (greatest to least). This can help you identify where to look first if you want to speed up your workbook.

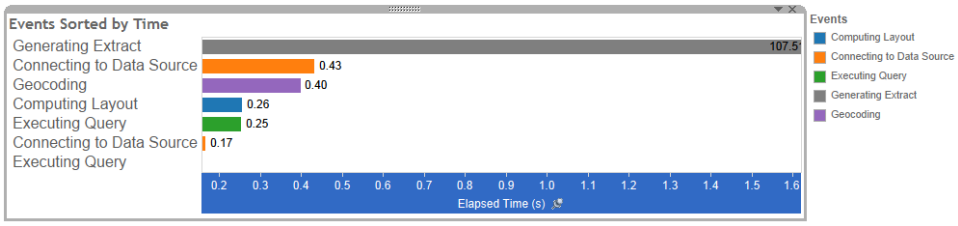

Different colors indicate different types of events. The range of events that can be recorded is:

- Computing layouts.

  If layouts are taking too long, consider simplifying your workbook.
- Connecting to data source.

  Slow connections could be due to network issues or issues with the database server.
- Executing query.

  If queries are taking too long, consult your database server's documentation.
- Generating extract.

  To speed up extract generation, consider only importing some data from the original data source. For example, you can filter on specific data fields, or create a sample based on a specified number of rows or percentage of the data.
- Geocoding.

  To speed up geocoding performance, try using less data or filtering out data.
- Blending data.

  To speed up data blending, try using less data or filtering out data.
- Server rendering.

  You can speed up server rendering by running additional VizQL Server processes on additional machines.

Query

If you click on an **Executing Query** event in either the **Timeline** or **Events** section of a performance recording dashboard , the text for that query is displayed in the Query section. For example:

Query

```
SELECT "State"."ID" AS "ID",
  "StateSynonyms"."Name" AS "State_Name",
  "State"."ParentID" AS "State_ParentID"
FROM "StateSynonyms"
  INNER JOIN "State" ON (("State"."ID" = "StateSynonyms"."ParentID") AND ("State"."MapCode" = "StateSynonyms"."MapCode
```

Sometimes the query is truncated and you'll need to look in the Tableau log to find the full query. Most database servers can give you advice about how to optimize a query by adding indexes or other techniques. See your database server documentation for details.

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 **Managing Queries** on page 190 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 you might experience a significant degradation in performance.

- 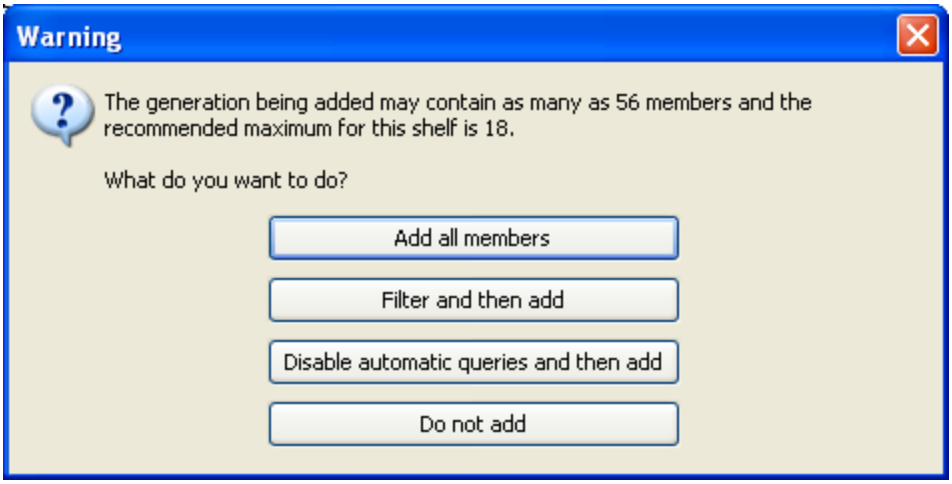

  Avoid Generating Too Many Panes – If you attempt to create many panes in a table, Tableau will display the warning dialog box shown below. In most cases, you should not display more than the recommended number of panes because such a view is not very useful.

- 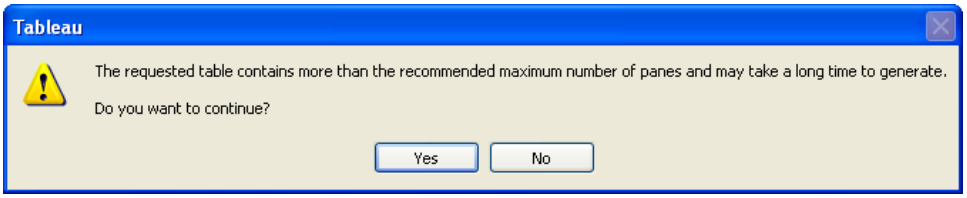

Performance Tips: Relational Data Sources

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 **Filtering** on page 334 for how to create context filters. For more information about performance improvement with context filters refer to **Speeding up Context Filters** on page 924.
- 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 **Sorting, Grouping, and Sets** on page 388, 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.
- 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

In Tableau, multidimensional data sources are supported only in Windows.

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 **Creating a Set** on page 417, 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.

Reducing Upload Times

It isn't always possible to reduce upload times, but here are some considerations to keep in mind.

Avoid Publishing Packaged Workbooks

It is a best practice to publish workbooks (.twb files) and data sources separately, rather than publishing them together as packaged workbook (.twbx files). This may not make your initial uploads of workbooks and data sources any faster than publishing them together, but it makes republishing workbooks faster (if you don't have to also republish data sources), and it makes data sources available online for new workbooks. For data sources that are extracts, it also allows you to schedule automated refreshes.

Make Extracts Smaller

For Tableau extracts, you can help speed uploads by making your extract smaller. As you create extracts, consider:

- Removing unused fields from extracts.
- Making extracts smaller with sampling and filtering.
- Aggregating data for visible dimensions.

See Also

Creating an Extract on page 174

Issues with Large Text and Excel Files below

Publishing Data Sources on page 832

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.
- 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.

> If your data set is heavily indexed, context filters may not provide performance improvement and may actually cause slower query performance.
>
> Context filters can adversely affect any query performance improvements when using the **Include joined tables only when referenced** option in the Tables dialog box. See Adding Tables to the Data window.

Issues with Large Text and Excel Files

You may encounter the following issues while working with large text and Microsoft Excel data sources:

- Performance issues
- Excel file errors

Performance Issues

If you have a workbook that was created before Tableau Desktop 8.2 and that uses large text or Excel data sources or the legacy connection, you may experience poor performance. This is due to queries running slowly. If your file has more than 20,000 rows, you should extract the data using the extract feature in Tableau Desktop.

> **Note:** This issue applies only to workbooks created on Windows.

To create an extract of a text or Excel file

1. Connect to a new text or Excel file. For information about connecting to a text or Excel file, see **Connecting to Data Sources** on page 57.
2. At the top of the data source page, select **Extract**. Alternatively, on the **Data** menu in the worksheet, select a data source, and then select **Extract Data**.
3. Specify filters and options to define a subset of data or just click **Extract** to include all the data. For more information, see **Extracting Data** on page 173.

   The data from the text or Excel file is converted and stored with the name you specify.

Excel File Errors

While simultaneously working with an Excel file and a Tableau workbook that uses the Excel file as its underlying data source, you may see an error in the Excel file indicating that changes to the Excel file could not be saved because of a sharing violation. You might see this error message when your Tableau workbook uses a large Excel file as its data source. In cases like this, the Excel file cannot be simultaneously edited or saved while the Tableau workbook is open.

Tips and Tricks

This section provides a collection of tips to help you use Tableau.

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.’

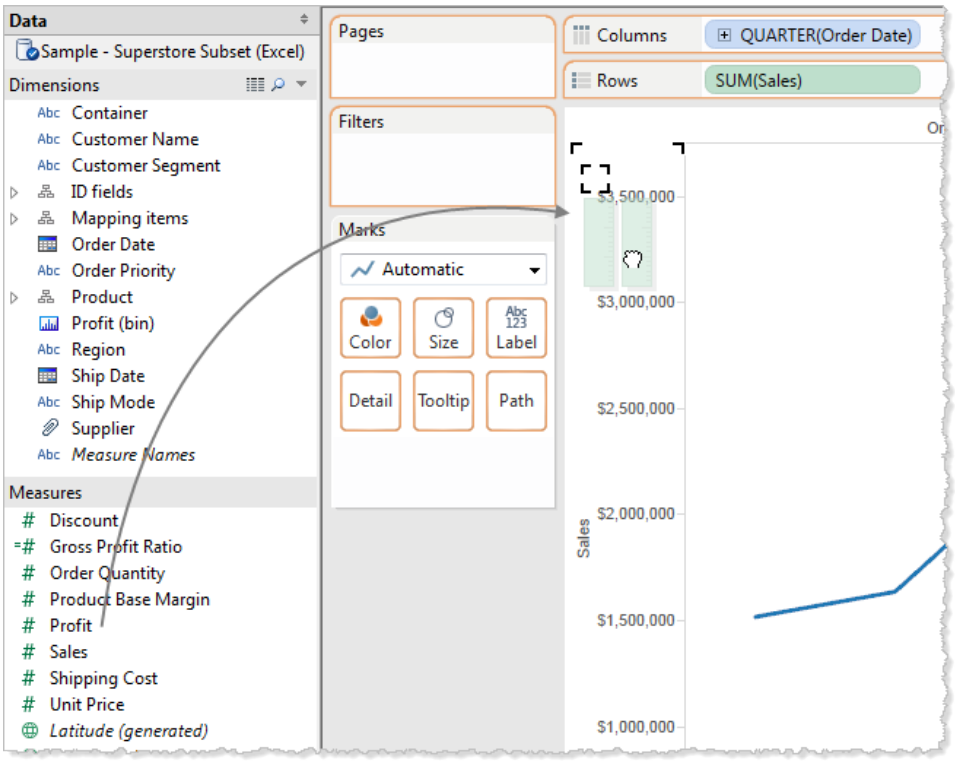

Convert Measures to Dimensions

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.

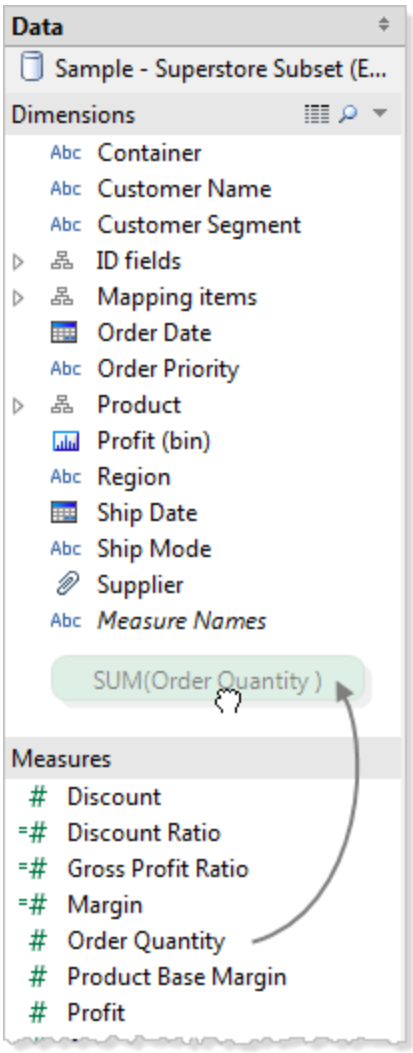

Convert Dimensions to Measures

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.

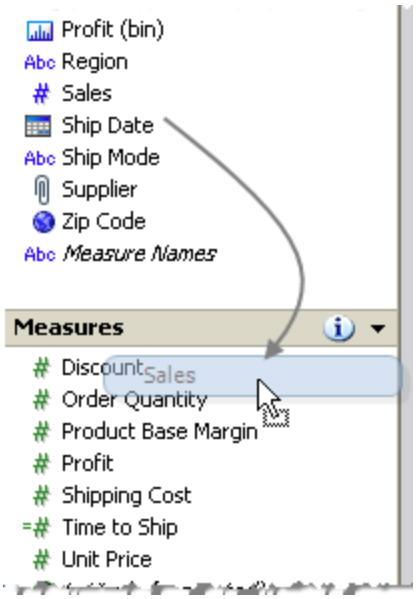

Duplicate Any View as a Cross-tab

To duplicate any worksheet in Tableau as a cross-tab, right-click the worksheet tab and select **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**

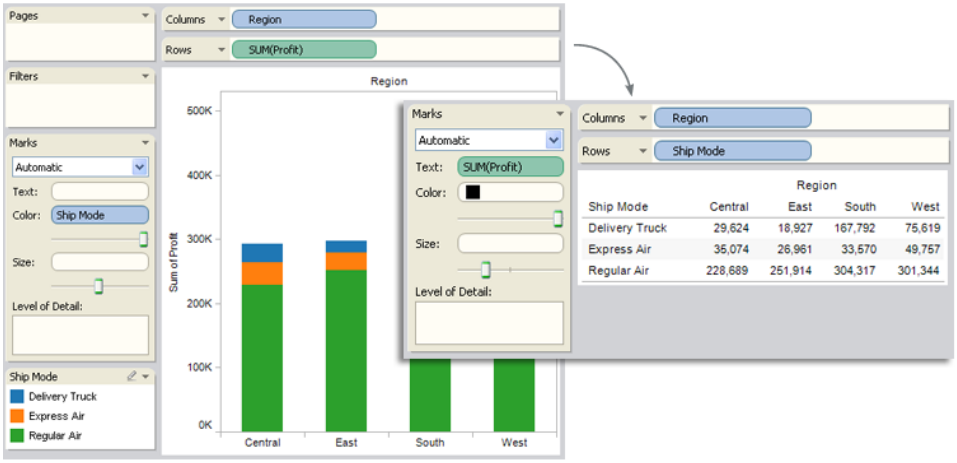

Manually Sort Items in the Color Legend

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.

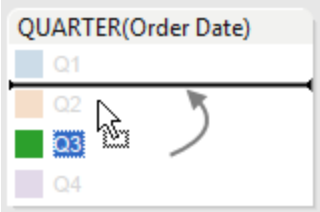

Duplicate Fields in the View

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.

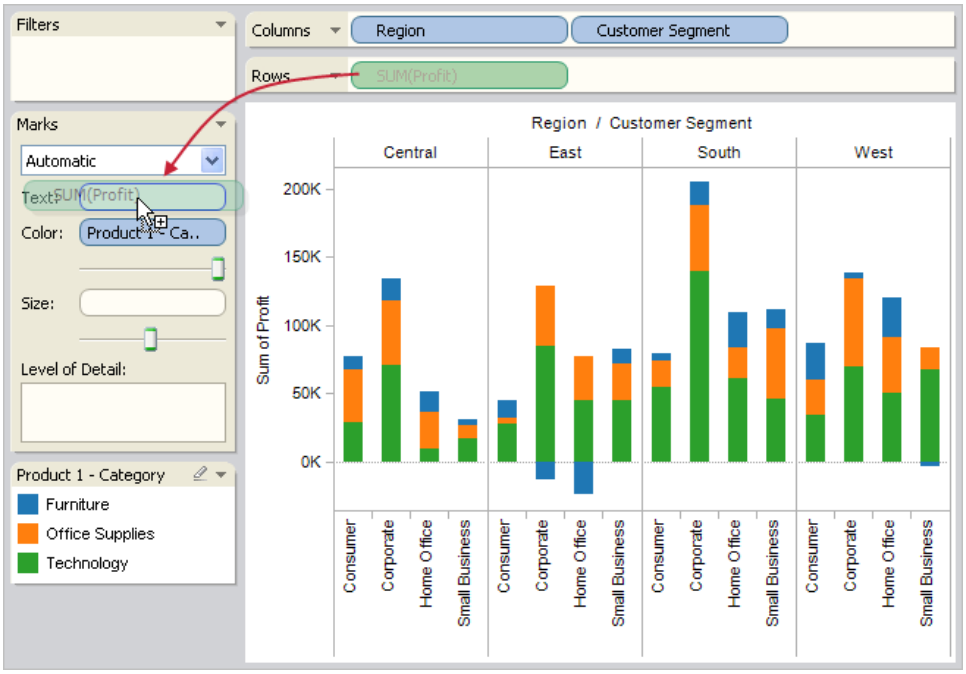

Resize Individual Views in a Dashboard

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:

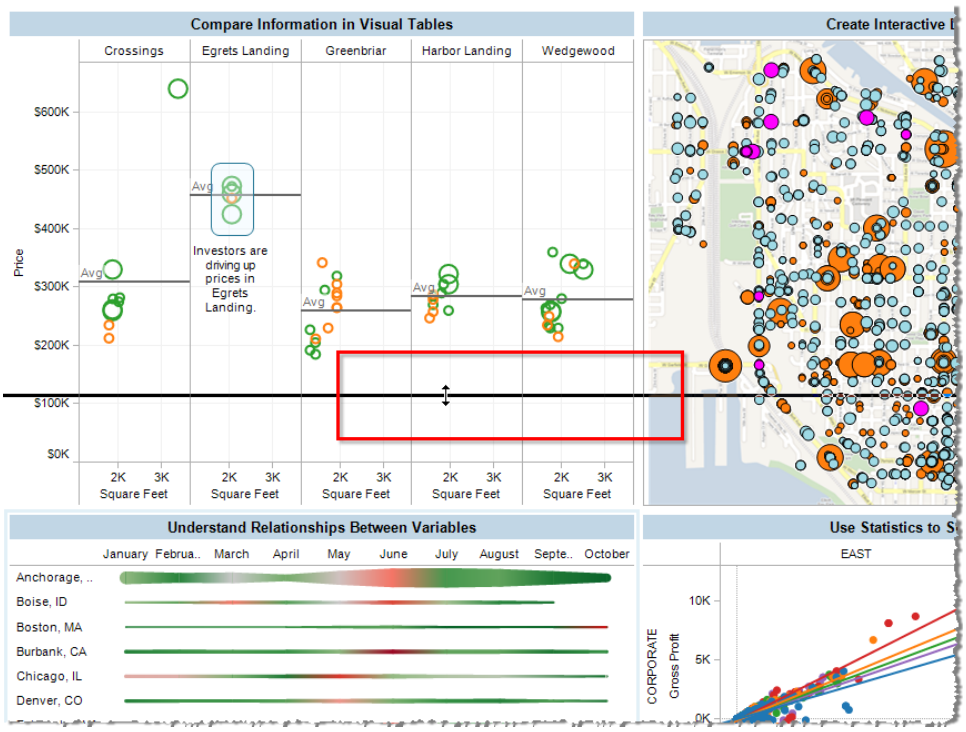

However, you can hold down the Shift key on your keyboard while dragging a divider line, to resize an individual view.

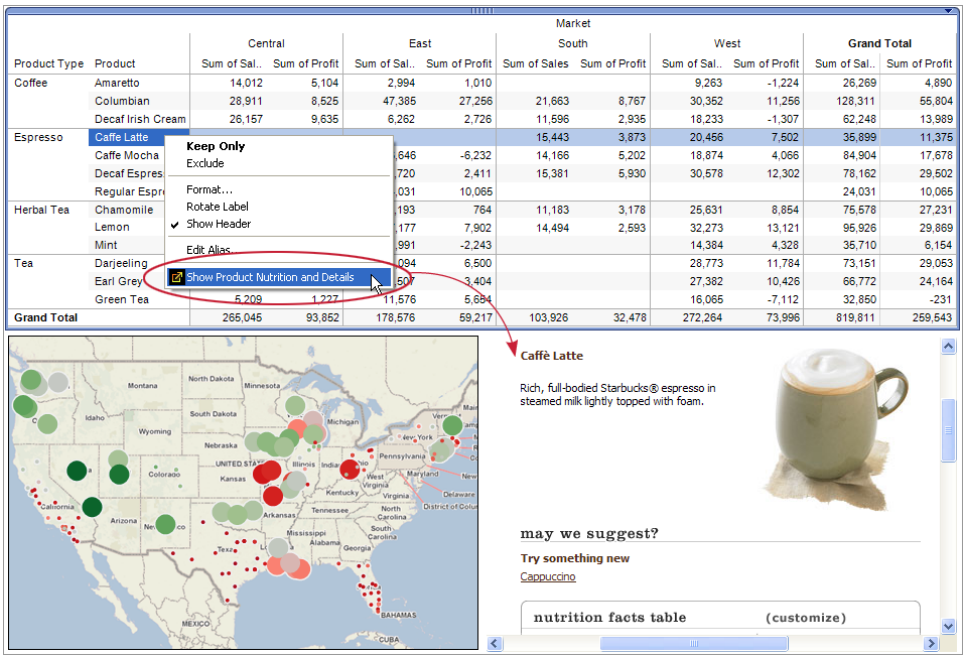

Use the Pages Shelf to Split Large PDF Reports

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 shows the total monthly expenses by Product and Market. When published to a PDF, the page is really long.

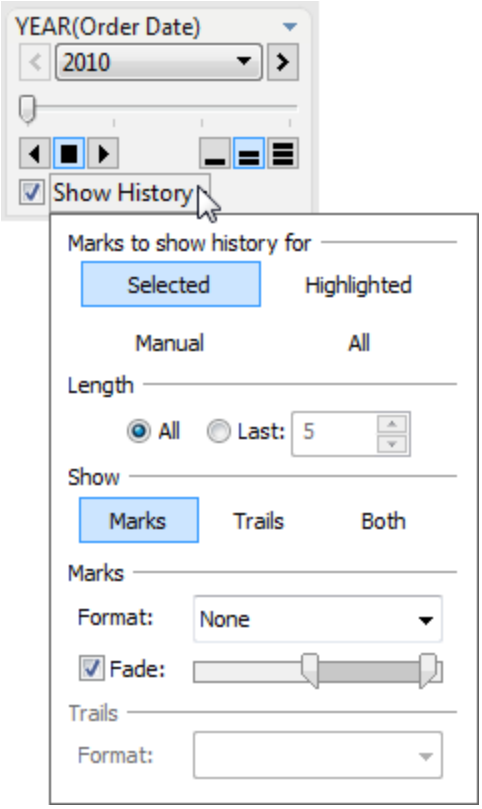

If you move the Market field to the Pages shelf, the published PDF fits on a standard letter sized paper and contains four pages — one for each market.

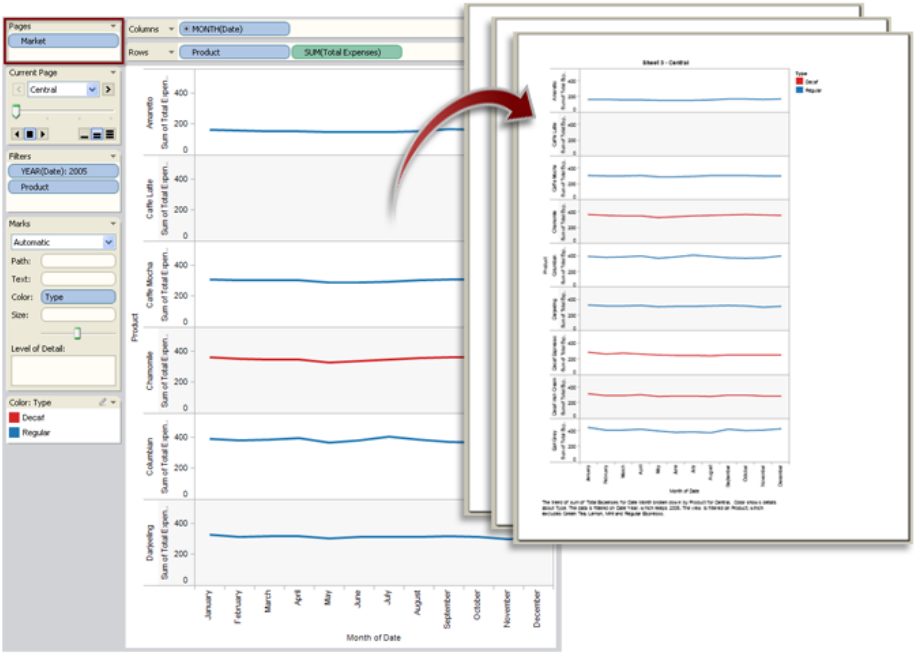

Specify Page Setup for Multiple Worksheets

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.

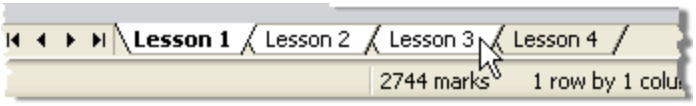

2. Select **File** > **Page Setup**.
3. Make the necessary changes in the Page setup dialog box and click **OK**.

Visual Cues and Icons in Tableau Desktop

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.

Data sources in the Data window

The following table explains each of the icons used to describe the type of data sources in the Data window. Each icon in the table can be modified by one of two indicators.

- A blue check mark indicates that the data source is the primary data source in the workbook.
- An orange check mark indicates that the data source is the secondary data source in the workbook.

Visual Cue	Description
	The workbook is connected to a relational data source.
	The workbook is connected to a multidimensional data source. In Tableau, multidimensional data sources are supported only in Windows.
	The workbook is connected to an extract.

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. Abc
- 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. =Abc
- Icons with an exclamation mark on them indicate that the field is invalid. !

Visual Cue	Description
Abc =Abc Abc!	The field contains text values.
# =# #!	The field contains numeric values.
=#	The field is a calculation defined on the server.
	The field contains only date values.

Visual Cue	Description
	The field contains both date and time values.
	The field contains geographical data and has been assigned a geographic role. Use these fields when building map views. See **Background Maps** on page 689.
	The field contains geographical data from an active custom geocoding file. See **Custom Geocoding** on page 702.
T|F =T|F	The field contains boolean (true or false) values.
=Abc	The field is a calculation that is defined in the database by an administrator. These fields are marked with a cylinder icon and are not available for all data sources.
	The field is a user-defined set. See **Sets** on page 416.
	The field is a server named set.
	The field is a set that was automatically created as a result of an action.
	The field is a user filter, used when publishing to the web. See **User Filtering** on page 843.
	The field is a numeric bin. See **Binned Data** on page 624.
	The field is an group. See **Groups** on page 401.
	The field is a relational hierarchy. See **Hierarchies (For Relational Databases)** on page 198.
	A folder that contains one or more fields. Folders are used to organize fields in the Data window. See **Organize the Data Window** on page 201.
	The field is an attribute of a multidimensional (cube) data source. In Tableau, multidimensional data sources are supported only in Windows.
	The field is a varying attribute of a multidimensional (cube) data source.
	The field is a level in a multidimensional hierarchy. Levels greater than five are shown without numbers.
	The field is a hierarchy in a dimension that contains multiple hierarchies.
	The field is blended with a field from another data source. See **Adding a Secondary Connection** on page 157
	The field is not blended with a field from another data source. See **Adding a Secondary Connection** on page 157

Fields on Shelves

Fields placed on shelves use a combination of icons, colors, and text styles as visual cues.

Visual Cue	Description
Container	A blue field on a shelf indicates a discrete field.
SUM(Sales)	A green field on a shelf indicates a continuous field.
Order Priority	The Sort icon indicates a field that has either a computed or manual sort order applied. See **Computed Sorting** on page 389.
Σ Call Center	The sigma icon indicates a calculation filter in a multidimensional (cube) data source. See **Calculation Filters: Multidimensional Data Source** on page 387.
MySet	The Venn diagram icon indicates a set. See **Sets** on page 416.
MySet	An field name shown in italics indicates a filtered set.
Region	A gray field on the Filters shelf indicates a context filter. See **Context Filters** on page 379.
SUM(Sales) Δ	The delta icon indicates that the field is a table calculation. See **Table Calculations** on page 593.
⊟ Country / Region ⊞ State	The plus and minus controls appear when the field is part of a hierarchy that you can traverse.
SUM(Sales)	The arrow icon indicates that a forecast is being displayed for the field. This icon is also used on the Forecast Indicator field, which is an automatic field used to distinguish between actual and forecast values. See **Forecasting** on page 645.
Status	The field is from a secondary data source. See **Working with Multiple Connections** on page 155.
Last Sale Price	The field is assigned to a specific worksheet.
Segment	The field is assigned to all worksheets with the same data source.
Internet Sales Amount	The field is incompatible with one or more other fields in the view.

Fields on the Marks card

Fields placed on the Marks card use specific icons to describe how they appear in the view. See **Mark Properties** on page 257.

Visual Cue	Description
Sales	The field is applied to Color on the Marks

Visual Cue	Description
	card.
Sales	The field is applied to Size on the Marks card.
Abc 123 Sales	The field is applied to Label on the Marks card.
Department	The field is applied to Shape on the Marks card.
YEAR(Order ..	The field is applied to Path on the Marks card. Path is only available when the Line or Polygon mark type is selected from the Marks drop-down menu. See **Path** on page 281.

Keyboard Shortcuts

General Keyboard Shortcuts

Tableau keyboard shortcuts are listed below.

Keyboard Shortcut		Description
Windows	**Mac**	
Ctrl+A	**⌘A**	Select all data
Ctrl+B	**⌘B**	Smaller cell size
Ctrl+Shift +B	**⇧⌘B**	Bigger cell size
Ctrl+C	**⌘C**	Copy selected data
Ctrl+Alt+C	**⌥⌘C**	Place selected field on Columns shelf
Ctrl+D	**⌘D**	Connect to data source
Ctrl+E	**⌘E**	Describe sheet
Ctrl+F	**⌘F**	Makes the find command in the Data window active
Ctrl+Alt+F	**⌥⌘F**	Place selected field on Filters shelf
	^⌘F	Enter/Exit Full Screen
F7 , Ctrl+H	**⌥Enter**	Switch in and out of Presentation Mode
Ctrl+Alt+I	**⌥⌘I**	Place selected field on Size
Ctrl+L		Flip orientation of column labels at bottom of view
Ctrl+Alt+L	**⌥⌘L**	Place selected field on Detail
Ctrl+M	**⌘T**	New worksheet

Ctrl+N	**⌘N**	New workbook
Ctrl+O	**⌘O**	Open file
Ctrl+Alt+O	**⌥⌘O**	Place selected field on Color
Ctrl+P	**⌘P**	Print
Ctrl+Alt+P	**⌥⌘P**	Place selected field on Pages shelf
Ctrl+Alt+R	**⌥⌘R**	Place selected field on Rows shelf
Ctrl+S	**⌘S**	Save file
Ctrl+Alt+S	**⌥⌘S**	Place selected field on Shape
Ctrl+Alt+T	**⌥⌘T**	Place selected field on Text/Label
Ctrl+V	**⌘V**	Paste clipboard
Ctrl+W	**^⌘W**	Swap rows and columns
Ctrl+X	**⌘X**	Cut text selection (e.g., in captions, titles, formulas, etc.)
Ctrl+Alt+X	**⌥⌘X**	Place selected field on Rows shelf
Ctrl+Y	**⌘Y**	Redo
Ctrl+Alt+Y	**⌥⌘Y**	Place selected field on Columns shelf
Ctrl+Z	**⌘Z**	Undo
Ctrl+Alt+Backspace	**⌥⌘delete**	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+!**	**⌘1**	Show Me!
Enter	**Return**	Add the selected field to the sheet. Only works with a single field
F1	**⇧⌘?**	Opens the Help
Ctrl+F4		Deletes the selected sheet (on a dashboard)
Alt+F4	**⌘W**	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	**⌘R**	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 ,	**⇧⌘{**	Cycle backward through open work-

Ctrl+Shift+F6		sheets
F9	**⇧⌘0**	Run update
F10	**⌥⌘0**	Toggles Automatic Updates on and off
F12	**⌥⌘E**	Reverts workbook to last saved state
Esc	**Esc**	Clears the selection (Desktop and Reader only)

Navigation and Selection Shortcuts

In addition to the standard keyboard shortcuts there are several key combinations that can make navigating and selecting marks fast and easy.

Keyboard/Mouse Action		Description
Windows	**Mac**	
Click	**Click**	Selects the mark
Drag	**Drag**	Selects a group of marks
Ctrl+Click	**⌘Click**	Adds individual marks to the selec-tion
Ctrl+Drag	**⌘Drag**	Adds a group of marks to the selec-tion
Shift+Drag , **Click and Hold+Drag**	**⇧Drag** , **Click and Hold+Drag**	Pans around the view
Double-click , **Ctrl+Shift+Click**	**Double-click** , **⇧⌘Click**	Zooms into a point in the view (requires zoom mode if not map)
Ctrl+Shift+Alt+Click	**⇧⌥⌘Click**	Zooms out from a point on a map (requires zoom mode if not map)
Shift+Double-click	**⇧Double-click**	Zooms out
Ctrl+Shift+Drag	**⇧⌘Drag**	Zooms into an area in the view (requires zoom mode if not map)
Ctrl+Scroll	**⌘Scroll**	Zooms in and out on a map (Desktop and Reader only)
Click+Drag to bottom of pane+Hold	**Click+Scroll**, **⌘Hold**	Drags a row and scrolls through a long list simultaneously

Field Selection Shortcuts

There are several key and mouse combinations that can make creating a view and selecting fields fast and easy.

Keyboard/Mouse Action	Description

Windows	Mac	
Right-click+Drag to shelf	**⌥Drag to shelf**	Opens the Drop Field menu
Ctrl+Drag	**⌘Drag**	Copies a field in the view to be placed on another shelf or card
Double-click	**Double-click**	Adds a field to the view

Updating the Software

When you upgrade to Tableau 8.3 from previous versions, your repository is upgraded. Any bookmarks, workbooks, and data sources that you had in your old repository will still be accessible by the application. In addition, the new sample data sources and workbooks will replace the old samples unless you have modified them and saved them as your own.

If you have participated in the Tableau Software Beta, you also have a beta repository. While this folder will still exist, the application will no longer access it. To make your beta workbooks accessible in Tableau 8.3, copy the workbooks from the beta repository to your new 8.3 repository.

Accessing the Help

Tableau offers multiple ways to access the help topics. First, you can always view the online help by selecting **Help** > **Get Help & Support** in the application. You must have an internet connection to view these pages. However, you can download a copy of the help files to your hard drive so that you can open them even when you are not connected to the internet. Finally, you can choose to download a printable PDF version of the help. This section explains each of the three ways to view the help.

Viewing the Online Help

The online help is the best way to access the most current documentation. You can open the help by selecting **Help** > **Get Help & Support** in the application or by pressing **F1** on your keyboard. You must have an internet connection to access this help.

Download the Offline Help

Although it is more efficient to view the help online, there may be times when you do not have access to the internet. In these cases, you may want a copy of the help stored on your computer. Download the offline help from the Product Help page on the Tableau website.

> **Note**: Offline help is only available for Tableau Desktop on Windows.

You must have an internet connection to initially download the help, but any subsequent times you want to access the help you are not required to be online.

If you choose to download the help and use the offline version, please remember to update your local files periodically in order to ensure you are viewing the most recent information. Anytime you want to update your offline help files, repeat the download and overwrite the existing files.

> If you download the offline help, Tableau will always open the local files when you select **Help** > **Get Help & Support** in the application. You can view the online help by visiting the Product Help page on the Tableau website

Download a PDF

You must have Adobe Reader to view this file. Adobe Reader is a free software that can be found on Adobe's Web site: www.adobe.com. You can download the PDF from the Product Help page on the Tableau website.

R Connection

R is an open source software programming language and a software environment for statistical computing and graphics. In Tableau Desktop, you can use a set of four functions to pass R expressions to an Rserve server and obtain a result. The functions are

`SCRIPT_BOOL`

`SCRIPT_INT`

`SCRIPT_REAL`

`SCRIPT_STR`

See **Table Calculation Functions** on page 899 for details and examples.

Before you can use any of the SCRIPT functions, you must establish a connection to an Rserve server, which is a server that allows applications to access R functionality. See Rserve for details. And for more information about using R with Tableau, see the blog post Tableau 8.1 and R.

> Tableau has been tested with R versions 3.1 and 3.2, and with Rserver version 0.6-8.

Configure an Rserve Connection

To configure an Rserve connection, follow these steps:

1. On the Help menu in Tableau Desktop choose **Settings and Performance > Manage R Connection** to open the Rserve connection dialog box:

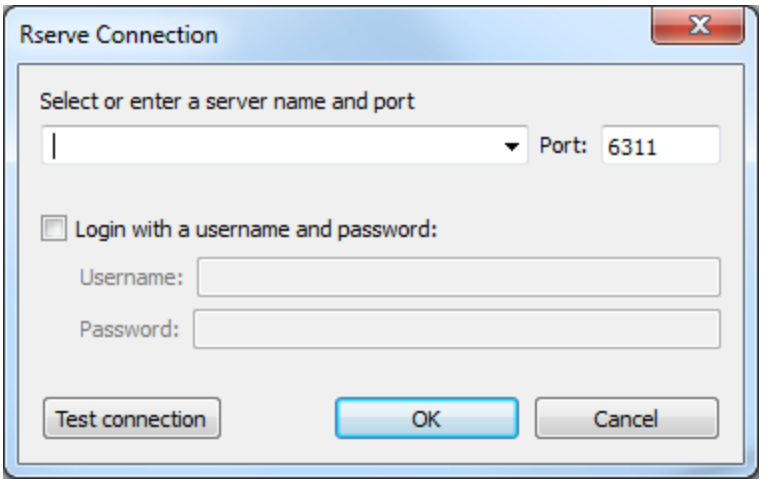

2. Enter or select a server name using a domain or an IP address. The drop-down list includes **localhost** and the server you most recently connected to.
3. Specify a port. Port 6311 is the default port for Rserve servers.
4. If the server requires credentials, specify a Username and Password.
5. Click **Test Connection**.
6. Click **OK**.

If no connection can be established, an error message is displayed. Click **Show Details** in the message to see any diagnostic information returned by the server.

Share a Workbook That Requires an Rserve Connection

You may need to send a workbook that contains R functionality to other users, who may be using different copies of Tableau Desktop on other computers. Or, users may download a workbook from Tableau Server that contains R functionality. Before users will be able to use the R functionality in workbooks they have received or downloaded to Tableau Desktop, they must configure Rserve connections on their computers.

Publish a Workbook That Requires an Rserve Connection

Before you publish a workbook that relies on an Rserve connection to Tableau Server, you should configure Tableau Server to have its own Rserve connection. You do this by configuring settings with `tabadmin set`. The settings are equivalent to the values you set in the Rserve Connection dialog box. They are:

- `vizqlserver.rserve.host`
- `vizqlserver.rserve.port`
- `vizqlserver.rserve.username`
- `vizqlserver.rserve.password`

> **Note**: The settings vizqlserver.rserve.username and vizqlserver.rserve.password should be omitted if Tableau Desktop connects to RServe without a username and password.

For information on how to use the above `tabadmin set` options to configure Tableau Server, see Change Tableau Server's Configuration from the Command Line.

Tableau cannot verify that workbooks that use R will render properly on Tableau Server. There might be scenarios where a required statistical library is available on a user's machine but not on the Rserve instance that Tableau Server is using.

For views that cannot be rendered in Tableau Server because of an R script error, you see a warning error when you publish the workbook:

> This worksheet contains R scripts, which cannot be viewed on the target platform until the administrator configures an Rserve connection.

You cannot publish a workbook that contains R scripting to Tableau Online.

> Because Tableau Server provides an authentication mechanism, it can be more secure to expose Rserve functionality to users through Tableau Server than in Tableau Desktop.

Troubleshooting R Connections

This topic describes errors you can receive when Tableau is connected to an Rserve server—or attempting to connect to an Rserve server. It also lists issues that you may encounter as you use R with Tableau.

Error Messages below

Other Issues on page 945

Error Messages

Error Message	Comments
Unsupported type passed as an argument to SCRIPT func-	Tableau can only export the following Tableau data types to R: • Number (Decimal)

tion.	• Number (Whole) • Boolean • String • Date • Date/Time
Unexpected number of results returned by SCRIPT function. Function expected %2 values; %1 values were returned.	The R script result must be either a scalar or vector of length one that is replicated for all rows, or a vector of length equal to the number of rows in the Tableau result table.
The result returned by the SCRIPT function is of an unexpected type.	Occurs when an invalid data type is received. Tableau can only import the following data types from R: • RDouble • RInteger • RLogical • RCharacter This error is also reported if the result is null or if there was a script execution error for which Tableau could not collect an explanation from R.
This worksheet contains R scripts, which are not supported on the target platform.	This error is reported when you attempt to publish a worksheet containing an R script to a server that does not allow R scripts because `vizqlserver.script.disabled` is set to true.
This worksheet contains R scripts, which cannot be viewed on the target platform until the administrator configures an Rserve connection.	This error is reported when you try to publish a worksheet containing an R script to a server is not currently configured to support R scripts. Specifically, the `vizqlserver.rserve.host` option, and associated `vizqlserver.rserve` parameters, for tabadmin set, must be configured properly. See tabadmin set options, in the Tableau Server online help, for details.
An error occurred while communicating with the Rserve service.	Tableau runs all R scripts inside of the "try" R function. This error is displayed along with an R generated error message when the “try” function traps an R evaluation error.
This Rserve connection does not support authentication. Try connecting	

without specifying a password.	
Authentication failed. Please provide a valid Rserve user-name and password.	
An unsupported authentication type is enabled in Rserve.	Either disable Rserve authentication or change to plaintext password authentication.
No Rserve con-nection configured. Specify a server name and try again.	See **R Connection** on page 940.
The calculation '%1' contains a disabled R script. Configure an Rserve connection to enable custom R scripts.	See **R Connection** on page 940.
The workbook you are attempting to pub-lish contains custom R scripts. Custom R scripts are not allowed in public workbooks.	You cannot publish workbooks containing R scripts to Tableau Public.
Tableau Public does not support running R scripts. To take advantage of R integ-ration, upgrade to Tableau Pro-fessional.	The "Tableau Public" in this error refers to Tableau Desktop Public.
Tableau Reader can-not load custom R scripts. To take advantage of R integ-ration, upgrade to Tableau Pro-fessional.	You cannot view workbooks containing R scripts in Tableau Reader.
Rserve is busy or not responding. Failed to	Tableau has timed out a read pending on connection to Rserve – the timeout is 250ms. An IPC connection Read of the Rserve protocol header has thrown an exception.

create socket to Rserve.	
Unrecognized Rserve signature.	The Rserve header signature must be "Rsrv".
Unrecognized Rserve version.	The Rserve header version must be "0103".
Unrecognized Rserve protocol.	The Rserve header protocol must be "QAP1".
Authentication failure when connecting to R.	Tableau attempted and failed to authenticate with Rserve. Verify that you entered a valid password.
Incorrect number of bytes in parameter/body.	
The length of data which Tableau read does not equal the length promised by the header.	
Unrecognized response type.	The transport protocol type when reading a result was not SEXP as expected.
Excessively long <type>vector.	Unreasonably large number measuring the length of data sent to or from Rserve possibly indicating a corrupt protocol header.
Invalid Rserve command.	Tableau may have improperly implemented the Rserve protocol.
Response from server was Error " << (uint32_t)status << ". See Rsrv.h for details.	Various error conditions are documented in the comments in Rsrv.h.
Excessively long header offset.	This may be due to a garbled header with an unreasonable offset to response data.
Rserve socket failed.	A login, script evaluation, read pending check, result read, or Tableau field to R script argument assignment threw a non-standard exception.

Other Issues

SCRIPT Functions Run Even in Logical Statements That Evaluate as False

A function that sends an expression to a running Rserve instance will be executed even when it is within a logical statement that would otherwise prevent it from being executed. This is true for logical functions such as IF, IIF, and CASE. For example:

```
IF 1==0 THEN

[[R script code]]

ELSE

"1 does not equal 0"

END
```

Glossary

A

Action

An action is interactivity that you can add to your views. There are three kinds of actions: Filter, Highlight, and URL.

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.

B

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.

C

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 Worksheet > Show 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.

Crosstab

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.

D

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 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.

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.

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.

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.

E

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.

F

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.

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.

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.

G

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.

H

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.

M

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.

P

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.

Q

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.

R

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.

S

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.

T

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.

U

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.

W

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

©2014 Tableau Software, Incorporated and its licensors. All rights reserved.

Patent www.tableausoftware.com/ip.

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

For a listing of third party copyright notices please refer to the following file that is installed with Tableau Server:

```
C:\Program Files\Tableau\Tableau Server\8.3\COPYRIGHTS.rtf
```

> Note: If you installed 32-bit Tableau Server on a 64-bit operating system, it will be in `C:\Program Files (x86)\Tableau\Tableau Server\8.3\COPYRIGHTS.rtf`

This product includes software developed by Andy Clark.

This product includes software developed by the Apache Software Foundation (http://www.apache.org/).

This product is Client Software as defined in Tableau Software's End User Software License Agreement.

Contact Us

Directory of worldwide offices

Sales

Contact

+1 888 365 0853

Support

1. Search our support resources.
2. Review the search results to see if your question is answered.
3. If you can't find what you need, scroll to the bottom of the search results, and click **Continue and Create Case**.

Other Inquiries

Email: info@tableausoftware.com

+1 206 633 3400