



Introducing Microsoft Power BI

Alberto Ferrari and Marco Russo

PUBLISHED BY
Microsoft Press
A division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 2016 by Microsoft Corporation

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

ISBN: 978-1-5093-0228-4

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Support at mssupport@microsoft.com. Please tell us what you think of this book at <http://aka.ms/tellpress>.

This book is provided “as-is” and expresses the author’s views and opinions. The views, opinions and information expressed in this book, including URL and other Internet website references, may change without notice.

Some examples depicted herein are provided for illustration only and are fictitious. No real association or connection is intended or should be inferred.

Microsoft and the trademarks listed at <http://www.microsoft.com> on the “Trademarks” webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

Acquisitions and Developmental Editor: Rosemary Caperton

Editorial Production: Dianne Russell, Octal Publishing, Inc.

Copyeditor: Bob Russell, Octal Publishing, Inc.

Technical Reviewer: Ed Price; Technical Review services provided by Content Master, a member of CM Group, Ltd.

Cover: Twist Creative • Seattle

Visit us today at

microsoftpressstore.com

- **Hundreds of titles available** – Books, eBooks, and online resources from industry experts
- **Free U.S. shipping**
- **eBooks in multiple formats** – Read on your computer, tablet, mobile device, or e-reader
- **Print & eBook Best Value Packs**
- **eBook Deal of the Week** – Save up to 60% on featured titles
- **Newsletter and special offers** – Be the first to hear about new releases, specials, and more
- **Register your book** – Get additional benefits




Contents

Introduction	v
Downloads.....	vi
Installing the companion content	vi
Acknowledgments.....	vi
Free ebooks from Microsoft Press	vii
Errata, updates, & book support	vii
We want to hear from you.....	vii
Stay in touch.....	vii
Introducing Power BI	1
Getting started with Power BI.....	2
Uploading data to Power BI	5
Introducing natural-language queries	6
Introducing Quick Insights.....	8
Introduction to reports.....	10
Introducing Visual Interactions.....	14
Decorating the report	17
Saving the report.....	19
Pinning a report.....	20
Refreshing the budget workbook	20
Filtering a report	23
Conclusions	25
Sharing the dashboard	27
Inviting a user to see a dashboard	27
Inviting users outside your organization	31
Creating a group workspace in Power BI.....	33
Turning on sharing with Microsoft OneDrive for Business	35
Viewing reports and dashboards on mobile devices	43
Conclusions	47

Understanding data refresh.....	49
Introducing data refresh	50
Introducing the Power BI refresh architecture	50
Introducing Power BI Desktop.....	52
Publishing to Power BI.....	55
Installing the Power BI Personal Gateway.....	57
Configuring automatic refresh	60
Conclusions	61
Using Power BI Desktop.....	62
Connecting to a database	63
Loading from multiple sources	65
Using Query Editor.....	68
Hiding or removing tables	73
Handling seasonality and sorting months.....	75
Conclusions	82
Getting data from services and content packs.....	84
Consuming a service content pack.....	85
Creating a custom dataset from a service	91
Creating a content pack for your organization.....	98
Consuming an organizational content pack.....	101
Updating an organizational content pack	103
Conclusions	106
Building a data model	107
Loading individual tables.....	108
Implementing measures	109
Creating calculated columns.....	111
Improving the report by using measures.....	112
Integrating budget information.....	113
Reallocating the budget.....	119
Conclusions	122
Improving Power BI reports	123
Choosing the right visualizations	124
Choosing between standard visuals.....	126
Using custom visualizations	130
First steps with custom visualizations	130
Improving reports by using custom visualizations	134
Identifying conditions when custom visualizations are required	139
Using DAX in data models	141

Creating high-density reports	144
Conclusions	149
Using Microsoft Power BI in your company	151
Getting data from existing systems.....	152
Understanding differences between data refresh and live connections.....	152
Using relational databases on-premises.....	153
Using relational databases in the cloud.....	155
Using live connections to Analysis Services	156
Integrating Power BI with Office.....	157
Publish Excel data models in Power BI	157
Consume Power BI content from Excel	158
Using Power BI Tiles from Office Store	162
Managing security to access data.....	166
Using row-level security	167
Extending and customizing Power BI	170
Creating custom visualizations for Power BI.....	170
Introducing the Power BI REST API.....	170
Pushing real-time data to Power BI dashboards.....	172
Power BI embedded in applications.....	174
Conclusions	175
About the authors	176



Exclusive offer for *Introducing Microsoft Power BI* readers!

Available this fall! Be the first to receive an **exclusive pre-order discount** for *Analyzing Data with Microsoft Power BI and Power Pivot for Excel* by expert trainers Alberto Ferrari and Marco Russo. It's easy:

- Sign up to receive **special offers from Microsoft Press**
- Watch your inbox for an exclusive email offer during the first week of September

Visit MicrosoftPressStore.com/PowerBI to get started. Even better? You'll receive a code* to save 35% on your next purchase at the Microsoft Press Store.

Learn more about Power BI at powerbi.microsoft.com.

Discount code valid on single book or eBook purchase from microsoftpressstore.com. Code cannot be combined with eBook Deal of the Day, Official Microsoft Practice Tests fulfilled by MeasureUp, or any other offer. Offer expires November 31, 2016.

Introduction

Microsoft introduced the idea of Self-Service Business Intelligence (BI) back in 2009, announcing Power Pivot for Microsoft Excel 2010. Strangely, at that time, it did not make big announcements, hold conferences, or undertake a big marketing campaign for it. Everything started slowly, with some enthusiastic users adopting the new technology, but the vast majority of people did not even know about its existence. As part of the community of BI professionals, we were very surprised by that approach. At the time, we could clearly see the advantages for users to begin adopting Power Pivot as a tool for gathering insights from data, so this complete lack of marketing was somewhat disappointing.

Thus, for several years we (as a community) kept asking Microsoft what they were waiting for; what was the delay in promoting Self-Service BI to the greater audience of data analysts, data scientists, decision makers, and BI enthusiasts all over the planet. We asked for the ability to share reports with a team, and the answer was to use SharePoint, either on-premises or the online version, with the first release of Power BI—an experience that was still not completely satisfactory. While we were waiting for Microsoft to fix the issues with the previous versions and to begin advertising the current products, it was doing something different that, with the benefit of hindsight, looks to have been the perfect choice. Microsoft collected the feedback of users, carefully considered what was missing in the world of end-user BI, and then crafted the version of Power BI that's available to you today.

Power BI is an evolution of the add-ins previously available in Excel: Power Pivot, Power Query, and Power View. You can use Power BI with or without Excel—you no longer are dependent on the version of Microsoft Office installed at your company. People did not like to share reports by using only SharePoint, and Microsoft moved away from it. Users wanted a mobile experience, and the development team created it. Data analysts wanted power, simplicity, new visualizations, and all of this is now available in Power BI. In addition, a lot of effort went into the creation of a seamless experience in loading data from many different cloud sources and building the infrastructure needed to provide all BI enthusiasts with a framework with which they can grow their reports, share them with their teams, and refresh the data in a simple yet effective way.

To make a long story short, Microsoft heard the feedback of users and built a great set of tools for the adoption of Self-Service BI. And, now—only now—it has begun marketing it.

Suddenly, in the last few months, Power BI has become the hottest topic at conferences, webinars, talks, and courses. As expected, people like you gathered interest in Power BI and began to search for resources to learn it. This book is one of these resources and its goal is to provide you with an effective introduction to the features available in the new Power BI.

We wanted to write an introduction to Power BI that covers the basics of the tool and, at the same time, shows you what the main capabilities of Power BI are. Thus, it is fair to say that the content of the book is somewhat unbalanced. At the beginning, we go for an easy introduction of the concepts along with an educational approach that lets you follow on your PC the same steps we show in the book. However, if we continued with that same mindset for the entire book, its size would quickly

become intimidating. Thus, after the first chapters, we begin to run a bit faster, knowing that we are no longer guiding you step by step. Instead, we show you available features; if you want to learn the details, you will need to read and study more.

This book is targeted to a variety of readers. There are information workers and people who are totally new to the BI world. For those readers, the book acts as a simple introduction to the concepts that are the foundation of BI. Yet, another category of we wanted to target is that of IT professionals and database administrators who might need to drive the decisions of the company in adopting Power BI, because their users are asking for it. If this is you, this book acts as both a simple introduction to the basic concepts, to help you understand why users are so interested in Power BI, and as an overview of the capabilities and tools available in Power BI, so that you can make educated choices in adopting it. Power BI is not just a tool: it is an ecosystem that can integrate existing corporate BI with Self-Service BI. The last chapter of the book gives you an overview of these capabilities.

We hope you enjoy reading the book as much as we enjoyed writing it. Keep in mind that this is probably your first step in the fascinating world of Self-Service BI, the first step of a long journey in gathering insights from your data.

Downloads

All of the chapters in this book include workbooks and databases that let you interactively try out new material learned in the main text. All sample content can be downloaded from the following page:

<http://aka.ms/IntroPowerBI/downloads>

Follow the instructions to download the IntroPowerBI_302284_CompanionContent.zip file.

Installing the companion content

Follow these steps to install the companion content on your computer so that you can use them with the exercises in this book.

1. Unzip the IntroPowerBI_302284_CompanionContent.zip file that you downloaded from the book's website (name a specific directory along with directions to create it, if necessary).
2. If prompted, review the displayed end user license agreement. If you accept the terms, select the accept option, and then click Next.

Acknowledgments

As usual with projects of this sort, there are too many people to thank, and a complete list of everyone who contributed to this book would be impossible to write.

Nevertheless, there are certain people we must mention personally, because of their particular contributions.

We want to thank Microsoft Press and all the people there who worked on the project. Rosemary Caperton has been a great editor who helped us immeasurably with the process of book writing. Many others behind the scenes guided us through the complexity of authoring a book—thanks to you all.

Finally, a special mention goes to our technical reviewer, Ed Price. He double-checked all the content of our original text, searching for errors and sentences that were not clear; giving us invaluable suggestions on how to improve the book. Without his meticulous work, the book would have been much harder to read and would contain more mistakes!

If the book contains fewer errors than our original manuscript, it is only because of them. If it still contains errors, it is our fault, of course.

Free ebooks from Microsoft Press

From technical overviews to in-depth information on special topics, the free ebooks from Microsoft Press cover a wide range of topics. These ebooks are available in PDF, EPUB, and Mobi for Kindle formats, ready for you to download at:

<http://aka.ms/mspressfree>

Check back often to see what is new!

Errata, updates, & book support

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

<http://aka.ms/IntroPowerBI/errata>

If you discover an error that is not already listed, please submit it to us at the same page.

If you need additional support, email Microsoft Press Book Support at mspinput@microsoft.com.

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to <http://support.microsoft.com>.

We want to hear from you

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

<http://aka.ms/tellpress>

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in touch

Let's keep the conversation going! We're on Twitter: <http://twitter.com/MicrosoftPress>.

Introducing Power BI

David is the manager of budgeting at Contoso, a company that sells electronic products worldwide through several retail shops and a website. Around the globe, country/region managers are responsible for producing figures for next year's budget for their respective countries/regions, which David then aggregates to produce the big picture to show to his boss.

Our scenario begins in October 2015, when David commences working on the budget for 2016. As always, David has a Microsoft Excel workbook containing the relevant information to produce the budget. Based on the results of the workbook, he would typically create a Microsoft PowerPoint presentation to share the results during internal meetings. This year, however, David wants to take advantage of the new Power BI service provided by Microsoft.

This entire book is a journey that we'll take along with David as he discovers how Power BI can help to build a rather sophisticated reporting solution; in this case, based on a budgeting system. But because this book is about Power BI, not budgeting, we will not focus on the complexity of building a budget. Instead, we will keep the budgeting considerations fairly basic, focusing on the complexity of teamwork, data modeling, and reporting.

We provided all the workbooks and databases that we used to build the demonstrations in the companion content for this book. If you are interested in learning the basics of Power BI, you can replicate David's activities on your computer so that you can augment your learning experience by following the examples. Be aware, though, that the results you obtain by running the demonstrations might be slightly different, and the appearance of webpages and the user interface might not be identical, either. Power BI is evolving very quickly, and we tried our best to show examples that will

last some time. Nevertheless, differences might occur; thus, you should concentrate on learning the features of Power BI, not the demonstrations. So, even if the numbers end up being different, what's important is to absorb how to do something, not just replicate what you read in the book.

Moreover, we strongly encourage you to test Power BI using your own data. You can perform the same operations on your personal information that we describe in the book, thus reaping the combined benefits of learning the Power BI tools while simultaneously gaining insights into your data.

Getting started with Power BI

Any journey begins with the first step, so let's take that step together.

David obtained from IT an Excel report that contains the sales for the past three years, divided by country/region, brand, and month. Sales in Contoso are strongly brand-oriented, and some brands are prone to seasonal effects that David wants to take into account. For this reason, he uses data grouped by month. Figure 1-1 shows a small portion of the resulting data, which he stores in an Excel file. If you would like to become more familiar with David's data, you can open 2015 Sales.xlsx from the book's companion content.

CountryRegion	Brand	Month	Sale 2013	Sale 2014	Sale 2015
China	Contoso	January	1,271.55	3,236.31	9,945.86
China	Contoso	February	7,372.34	5,995.55	1,303.01
China	Contoso	March	11,364.07	2,767.32	3,437.26
China	Contoso	April	13,590.94	17,530.47	25,007.74
China	Contoso	May	11,780.08	13,501.47	5,270.05
China	Contoso	June	8,562.00	8,338.87	7,357.26
China	Contoso	July	7,001.78	13,874.09	31,571.42
China	Contoso	August	605.89	3,340.02	1,386.20
China	Contoso	September	3,958.70	1,709.68	2,224.36
China	Contoso	October	8,785.89	5,787.57	10,792.05
China	Contoso	November	7,631.64	2,412.91	
China	Contoso	December	21,968.86	5,657.70	
China	Fabrikam	January	5,794.38	9,244.97	
China	Fabrikam	February		6,230.00	
China	Fabrikam	March	3,415.96	1,044.89	
China	Fabrikam	April			5,155.78
China	Fabrikam	May	3,657.95	2,954.00	20,177.94

Figure 1-1: An excerpt from the initial Excel workbook for David's budget plan.

Every year, David makes some considerations on these numbers and then he shares his findings with the country/region managers, who then send back to him workbooks with their numbers for the next year. Figure 1-1 shows some data from China, but there are several other countries/regions, as well. During the process of computing those numbers, there are many meetings and discussions in which the managers bring their experience and knowledge to bear on the process, adding their own versions of the original workbooks, each displaying various charts and calculations, which must all be explained to others. This is a daunting task, to be certain, and one that David would like to streamline.

Fortunately, David heard about an interesting tool called Power BI that Microsoft created in 2015 that might be helpful toward creating a collaborative environment in which any stakeholder of the budgeting process can share his findings with others, working together on the goal. But, at this point, the name and maybe a marketing video is all that David knows about Power BI.

Driven by curiosity, he navigates to www.powerbi.com and starts down his learning path. Figure 1-2 depicts the welcome page of the Power BI website.

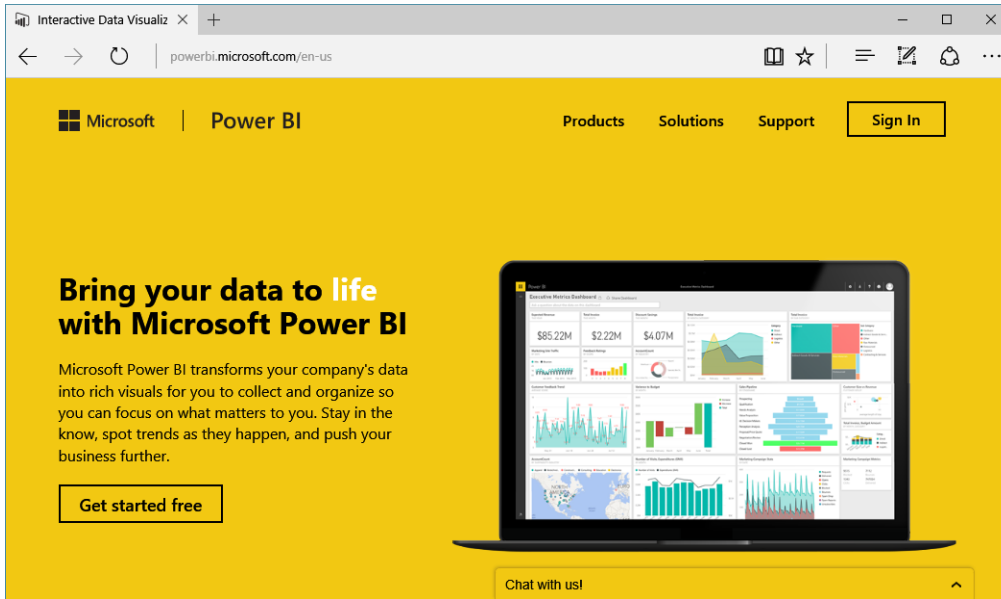


Figure 1-2: The welcome page of Power BI, the starting point of David's journey.

To begin, David clicks the Get Started Free button. He is then offered a choice as to which experience he would prefer to use: he can choose either Power BI Desktop For Windows or just Power BI, as shown in Figure 1-3.

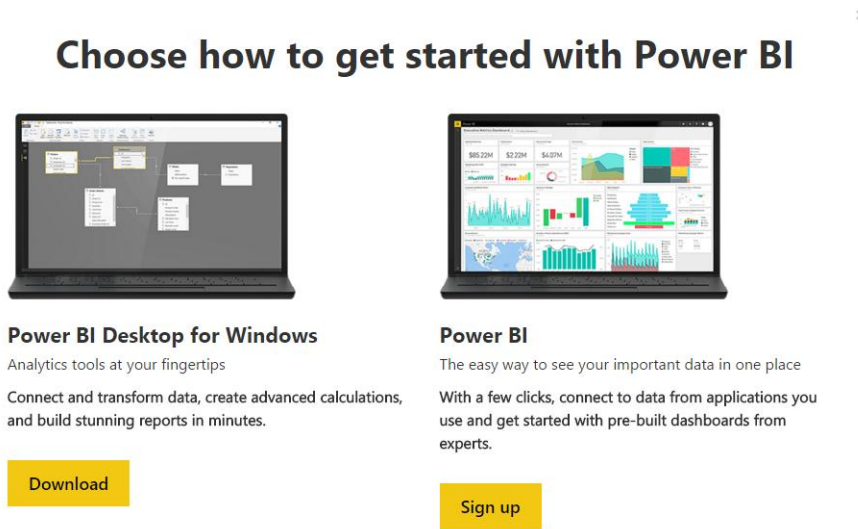


Figure 1-3: You can start with Power BI by using either of the two main experiences.

Actually, there is very little difference between the two. In fact, Power BI Desktop and Power BI are two sides of the same coin: Power BI Desktop is a Windows application running on your PC, whereas Power BI is a cloud service that you use through the web browser. In both cases, you will be able to perform the same operations, albeit with some subtle differences. Moreover, the two tools complement each other, and you are likely to use both to build your dashboards.

After reading the descriptions, David correctly concludes that Power BI Desktop is designed for more advanced tasks. Given that he's just beginning to learn about it, he opts for plain, vanilla Power BI.

When David clicks the Sign Up button, the screen shown in Figure 1-4 appears. Power BI is a web service to which you can upload data and build insightful dashboards and charts. As with any web service, you need to sign in, but Power BI does not require much in the way of credentials: to get started, all you need is a valid email address, which David provides.




Figure 1-4: You need only a valid email address to gain access to Power BI.

After clicking Sign Up, Power BI informs David that he already has a subscription to Microsoft Office 365; those credentials are sufficient to gain access to Power BI.

Note If you do not have an Office 365 account, Power BI will send you an email with a link to complete the registration process. (Be aware that you cannot use a personal email service such as Hotmail, Yahoo, or Gmail.) This is to ensure that you actually own the email address. Following the link directs you to the registration page, where you provide some basic details such as first name, last name, and so on. In both cases, no credit card or any other form of payment is required, because most of the features of Power BI are totally free.

On the same No Need To Sign Up page, David could click OK, Got It to sign in without any additional steps. Rather than do that, he goes back to the sign-in page, but instead of clicking Sign Up, he signs in to the portal by using the Sign In button (see Figure 1-2) and then provides his Office 365 credentials. The system takes a few seconds to prepare his account, after which David gets his first glimpse at the Power BI portal, as shown in Figure 1-5.

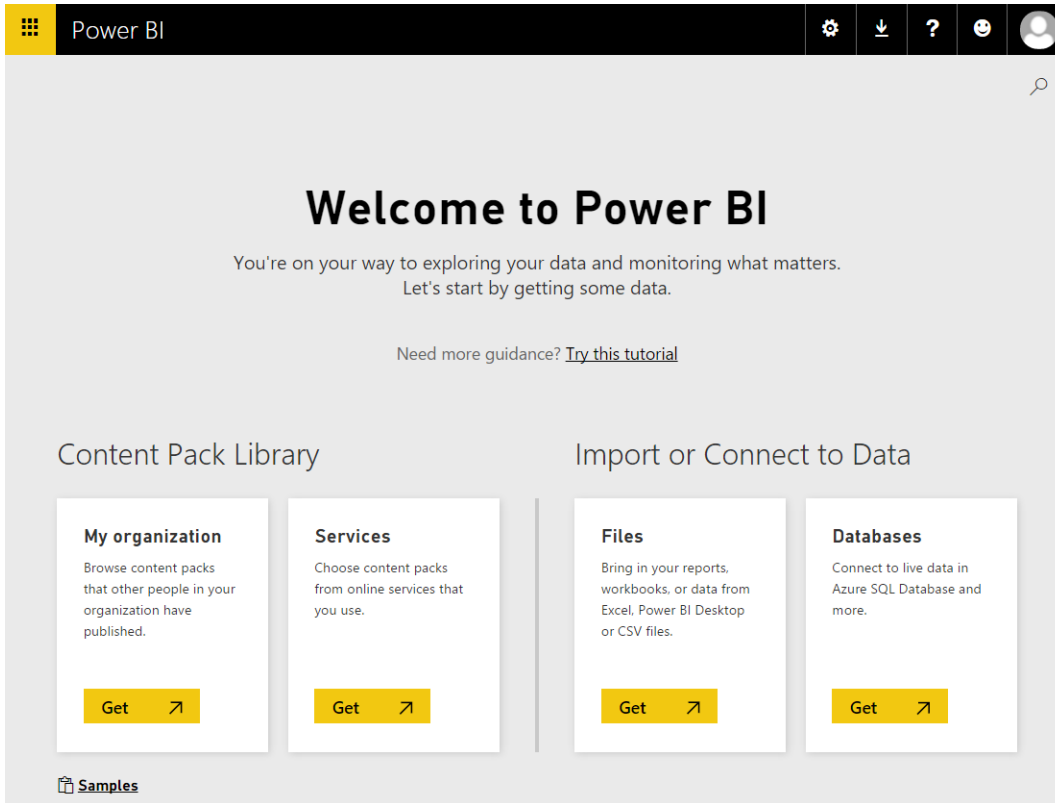


Figure 1-5: The introduction page of the Power BI portal.

Uploading data to Power BI

David has an Excel workbook that he wants to upload to Power BI to see what it has to offer. Because the data is stored in a local file on his laptop, he clicks the Get button on the Files tile (see Figure 1-5). This displays the screen in Figure 1-6, where he can then choose from among several upload options.



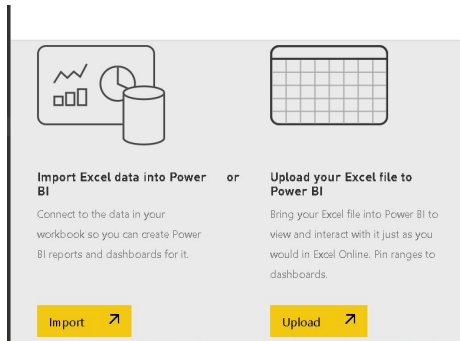


Figure 1-6: Some of the file uploading options in Power BI.

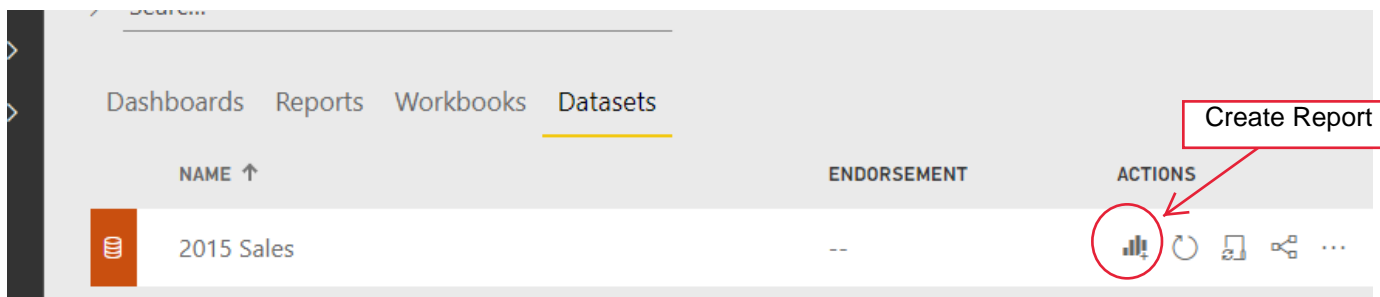
We will explore these options at greater length in the chapters that follow. For now, David chooses **Local File**, navigates to a file on his laptop named **2015 Sales.xlsx**, and then clicks **Open**. You will be given 2 choices, **Import Excel data into Power BI** or **Upload your Excel file to Power BI**. Click the **Import** button 



A message will display at the top right of the screen, confirming that your file has been updated, close this box.



You will see your Excel file listed as below under the heading **Datasets**. There are 4 headings as below:



- **Dashboards** This lists all of the dashboards you have created. After loading a single workbook, Power BI creates a dashboard for you, using the same name as that of the original workbook.
- **Reports** Here, you will see the reports based on your data. In Figure 1-7, there is no default report, but we'll follow along as David creates one very soon.
- **Workbooks** This lists all the workbooks uploaded to Power BI
- **Datasets** This lists all of the data sources that you connected to Power BI. In our narrative thus far, the only workbook David loaded is 2015 Sales.

The Power BI experience is all about gaining insights from data. You begin with a dataset (2015 Sales, in this example), you then build reports on the data, and, finally, you organize visualizations of the reports into dashboards. You will get an overview on how to perform all of these operations in detail in this book. For the moment, we want only for you to become acquainted with the basic operations.


Introducing natural-language queries

With Power BI, you have the ability to carry out analysis of your data by asking it questions, in plain English—no special code or syntax is required. This feature is called *natural-language queries*, in simple English, translates to 'Ask Power BI A Question about your data'. With it, you can ask Power BI to perform tasks in much the same way you would ask one of your colleagues. Let's take a look at an example of how David uses natural-language queries in Power BI.

To open the Excel file in Power BI, **click on the Create Report** button  under **Actions**.

The 'Ask a question tile' opens, offering a few report suggestions as below, to help him get started, but David types his own simple query, **"Show sales 2015 by brand."**

The Power BI screen will appear displaying the central pane with a diagram of a blank line chart or instructions on how to Build visuals with your data. At the top of the screen you will see the **Ask a Question** button.



As he is typing various suggestion appear. When he is finished a small **Clustered Bar Chart**, called a **Visualization** appears in a tile, as below:

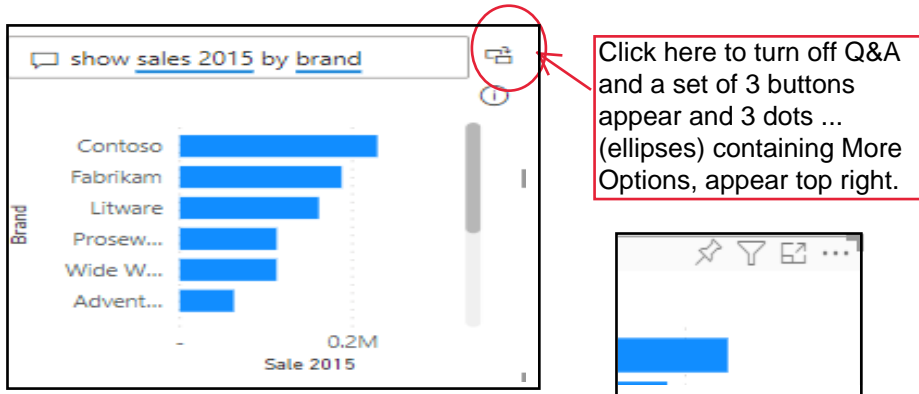
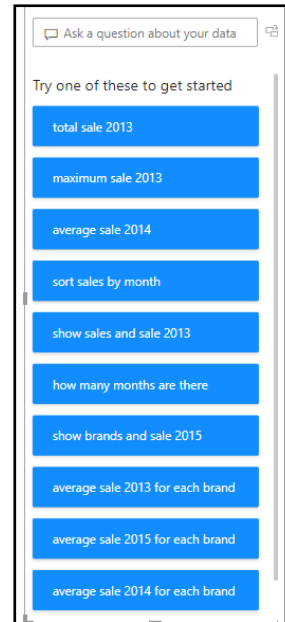


Figure 1-8: Power BI understands queries in natural language and displays the data you requested.

Click **File / Save as Report** and name it **Sales 2015**. Click **My Workspace** on left side of the screen, click **Reports**. Double-click on your report **Sales 2015** and you will see a small **pushpin icon - Pin to a Dashboard**, at the top of the screen. Click this to **"pin"** the currently displayed visualization to the **Dashboard**; it will be listed under Dashboards when you connect to Power BI. When the **pushpin is clicked**, Power BI opens the **Pin To Dashboard** dialog box shown in Figure 1-9. Type the Dashboard name **'Sales 2015'** and click **Pin live**.

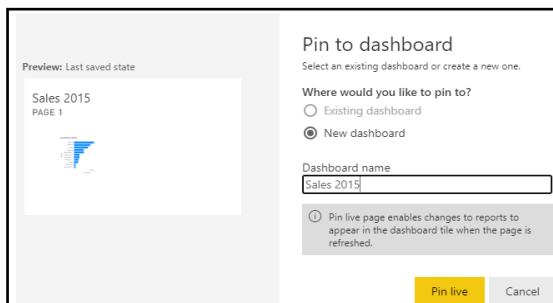




Figure 1-9: Using the **Pin To Dashboard** dialog box, you can choose to pin a visualization to an existing or a new dashboard.


The following message will appear at the top right of the screen, click **Go to Dashboard**.





Click **My workspace** at the top or left of the screen.

We want to do more work on the Visualization, click **Reports**



The report name will appear, double-click it to return to the Central pane, displaying the **Clustered Bar Chart**.

Click **Edit report** at the top of the screen, to continue working.

Figure 1-10: The dashboard is a container for visualizations which appear in a tile.

Using **natural-language queries** is quite impressive, but it is only one of the many ways in which Power BI can analyze your data.

Introducing Quick Insights

Another feature that is worth learning as soon as you begin using Power BI is **Quick Insights**. With this feature, Power BI can search a dataset for interesting patterns and provide you with a list of charts that help you to better understand your data.

To check out Quick Insights, click **My workspace** and click **Datasets**. To the right of the Dataset name click on the **... More Options** and click **Get Quick Insights**. A message box will appear, as below, click **View insights**.




Figure 1-11: The Quick Insights message box

The first time you run Quick Insights on a dataset, Power BI schedules an analysis of that dataset. This might last for some seconds or minutes, depending on the size of the data. When the search for insights is complete, Power BI notifies you. Of course, whenever you update your dataset, this search operation will need to be repeated. However, as long as the dataset remains unchanged, the insights will be immediately available.

But, what are these insights?

The basic idea is that Power BI can use artificial intelligence to analyze your data, searching for some useful or interesting patterns. It uses very sophisticated algorithms whose speed depends on the size and complexity of the dataset. Obviously, on a small dataset such as the one David uploaded, finding insights takes no more than a few seconds. Figure 1-12 presents the first two insights that Power BI found on David's file. Many others are within the list, so many, in fact, that they would not fit on the page in this book. David scrolls down the right side of the screen to view them all.

Power BI presents an impressive range of what it has extracted from your data into various charts, depicting many different scenarios, that could be used by you in your report.

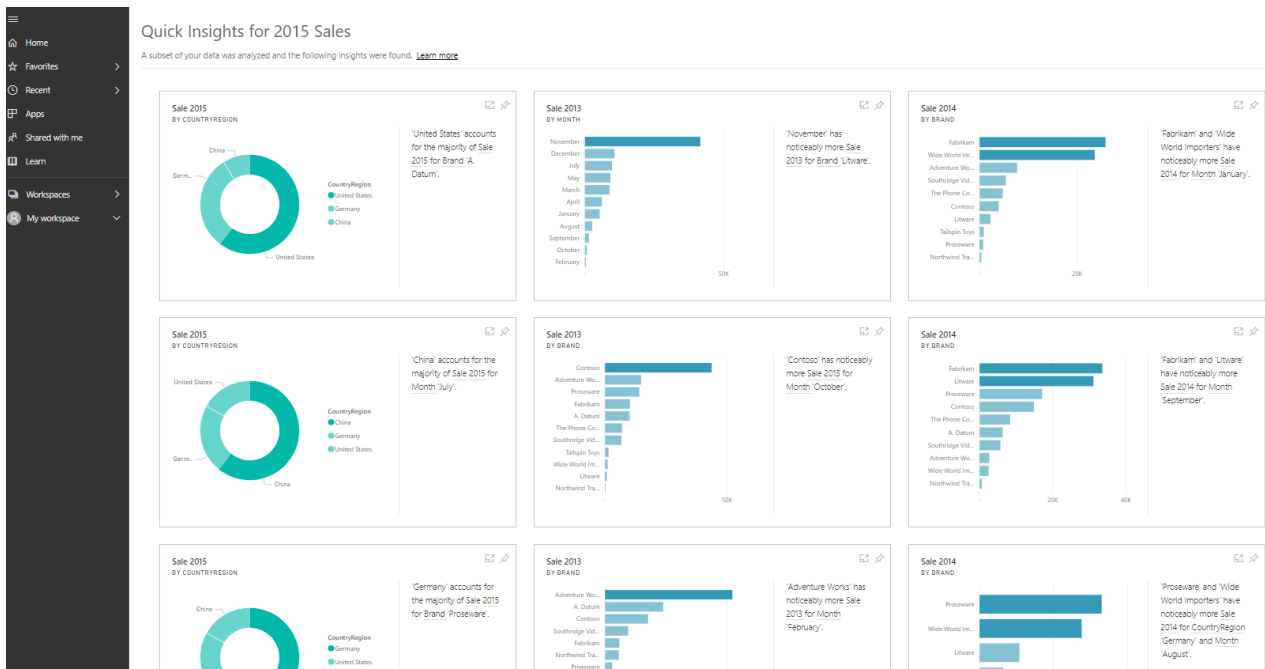


Figure 1-12: Quick Insights are a powerful analytical tool to glean information from your data.

The first insight shows Sales for 2015 by country and region and that the United States accounts for most of the sales of the A. Datum brand, compared with China and Germany.

As you scroll down the insights you will see that they are for the years 2-13, 2-14 and 2015 by Brand, Month or CountryRegion

Of course, **insights are gathered by Power BI without it having any knowledge of your business or the economic scenario as a whole**, so there might be many different reasons that explain the data and findings. Power BI cannot replace your brain when it comes to interpreting the numbers, but it can prove extremely useful because it can easily find some points of interests in your data by using the brute force of algorithms.


The best way you can use Quick Insights is to browse through them, looking for the confirmation of what you already know about your data and, at the same time, for fresh ideas.

It might be the case that some of the insights are not really meaningful, but, with the sheer number of insights that Power BI finds for you, it's likely that there are some real hidden gems that might improve your knowledge of your numbers.

More info You can find a more complete description of the algorithms used by Power BI and the types of insights that it can reveal by going to <https://powerbi.microsoft.com/en-US/documentation/powerbi-service-auto-insights-types/>. Of course, with newer versions of the analytics engine, the numbers and the quality of insights might change and improve.

Clicking any insight will enlarge it. To return to the previous screen, click **Exit Focus mode** at the top left of the screen. If you hover over one, the same pushpin button as that of the natural-language query appears so that you can pin the insight to the dashboard if you want.

David clicks the **Pin** on the top-right corner of the insight named, **Sale 2015 BY COUNTRYREGION**. When the **Pin To Dashboard** dialog box is displayed, he leaves **Existing Dashboard** selected and clicks **Pin** to pin it to his dashboard.



Pinning one of the insights to the dashboard makes it more interesting. Moreover, by doing that, you will learn that you can move and resize visualizations pinned to a dashboard by using a convenient grid, making them more aesthetically appealing. David returns to look at his expanded **2015 Sales dashboard**, by clicking **My Workspace**, and from under **Dashboards** he clicks **Sales 2015**. The new insight appears a. Figure 1-13 demonstrates David's dashboard, which now contains two visualizations, which can be resized and moved around the screen.

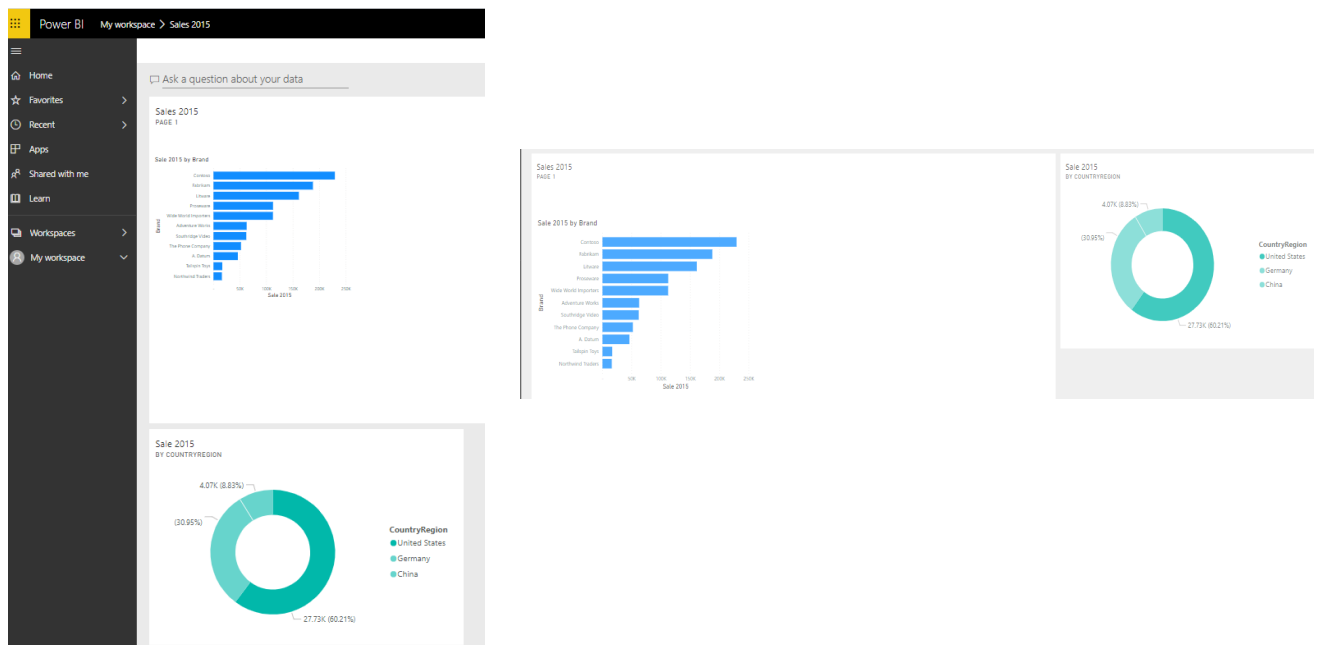



Figure 1-13: A dashboard can contain multiple visualizations organized in tiles in a grid, and can be individually moved by dragging it around the central pane. When a visualization is selected drag from the corners to resize.


Introduction to reports


So far, David has used only automated report building, using a natural-language query as well as the Quick Insights feature. As you might imagine, he only scratched the surface of Power BI's reporting capabilities. In fact, he can build reports manually, unleashing the full potential of Power BI visualizations.

To create a **new Report**, click **My workspace** at the top left of the screen.



When the 4 options appear, click the **Datasets** button. 2015 Sales appears, as below.



Under **Actions**, David clicks the **New Report Icon** 

Power BI opens an empty report based on that dataset, as illustrated in Figure 1-14.

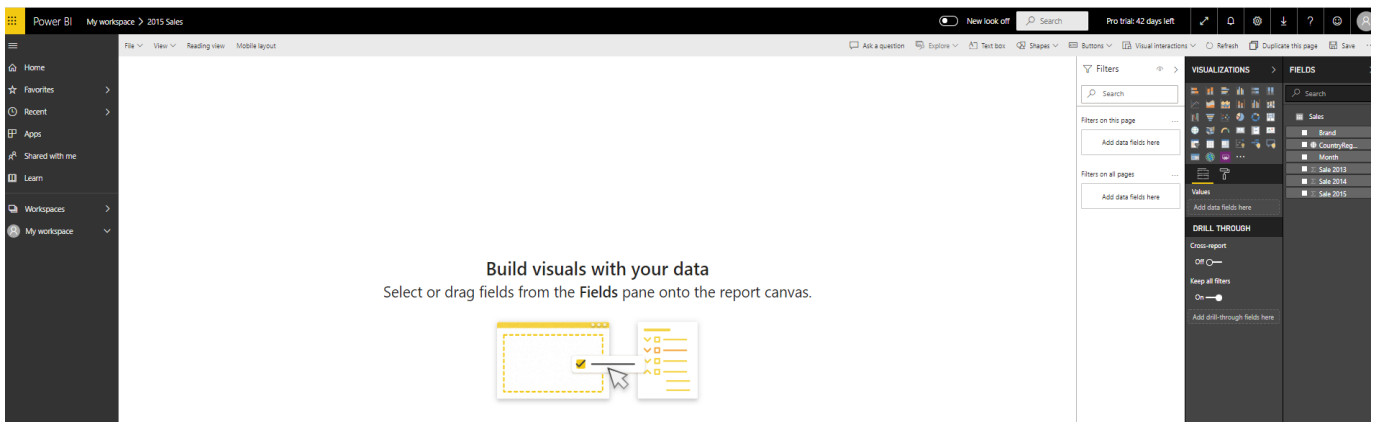


Figure 1-14: An empty report displays, based on the current dataset. The Excel file fields are listed on the right





Figure 1-15 shows an enlarged view of the **FIELDS** pane.

Figure 1-15: The names of the columns from the Excel file appear in the **Fields** list and a small **Summarization** icon  appears to the left of any field that has a **numerical** value. By placing a tick in the box to the left of the field name, Power BI will use that data and display in the central pane as a Visualization Table, Graph or Map.

To create a report, you must select the fields that you want to appear in the report. For example, referring back to Figure 1-14, in the Fields pane, David clicks **Brand** and then **Sale 2015**. Because **Brand has no summarization icon, the text from it is used to slice data**, whereas **Sale 2015**, which displays a summarization icon, will present each brand name's **value** and a total sum at the bottom, generating the report shown in Figure 1-16. You may find the table to be too small to read, at the top of the screen there is **View** option with a drop down list, select **Actual size**, making it easier to read. There are some other interesting options you can try, like high contrast colors etc.

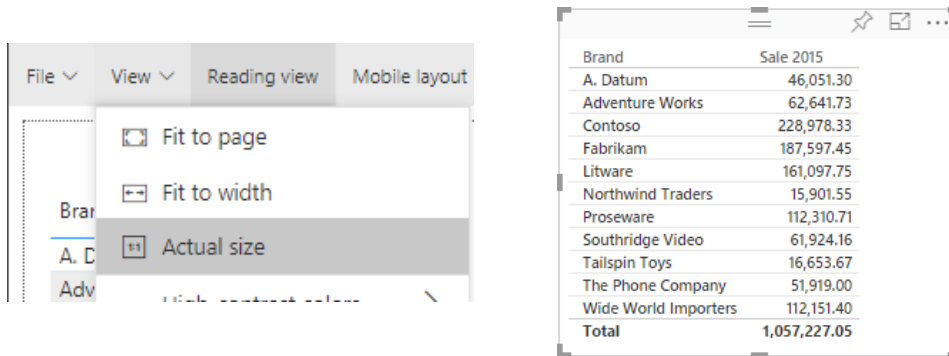


Figure 1-16: A visualization based on Brand and Sale 2015.

What David just created is the **default visualization**; that is, **it's a grid with the brands and the sum of Sale 2015 on the rows, showing raw numbers**. Numbers are very interesting, but they do not give a clear idea of the relationship among them. In fact, at first glance it's not evident which brand is the most important one, which ones are the smallest, and what the relative importance of the numbers is. Charts, on the other hand, can give viewers a much quicker understanding of the data.

You can modify the visualization of a tile by choosing one of the many available types of charts in the Visualizations pane. For example, with the table visualization selected, you can change it to a **Clustered Bar Chart** visualization, by clicking the icon under Visualizations, which is highlighted below.

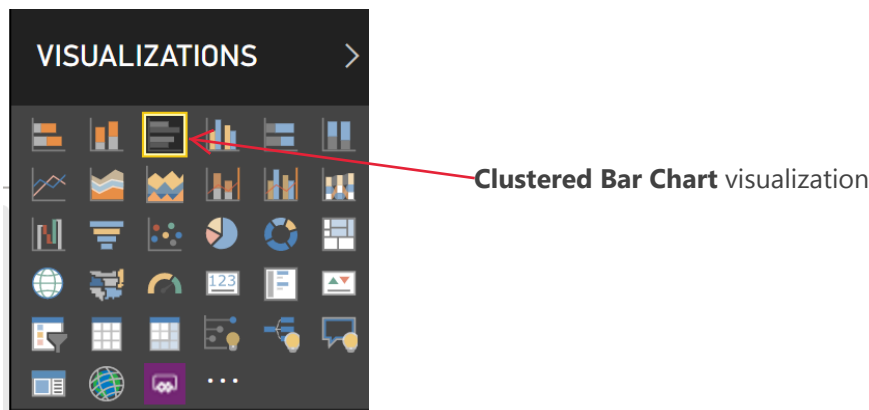



Figure 1-17: The Visualizations pane offers many different visualizations to use in your reports.

Note If you click a visualization type but do not have a specific tile selected, Power BI inserts a new, empty visualization. If this happens to you, do not worry: just select the empty chart and delete it by pressing the Delete key. Then, select the tile that you want to change and try again.

As **Figure 1-18** on the following page, so clearly demonstrates, the table visualization, **Sale 2015 by Brand** can instantly be shown as a **Clustered Bar Chart** making the data much easier to understand.



Note By clicking on the **More Options** (ellipse dots ...) at the top right of the tile, you can, select to **sort** the data by Brand or Sale 2015, try Brand in **ascending** order

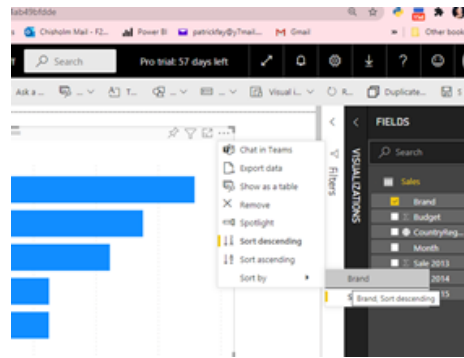
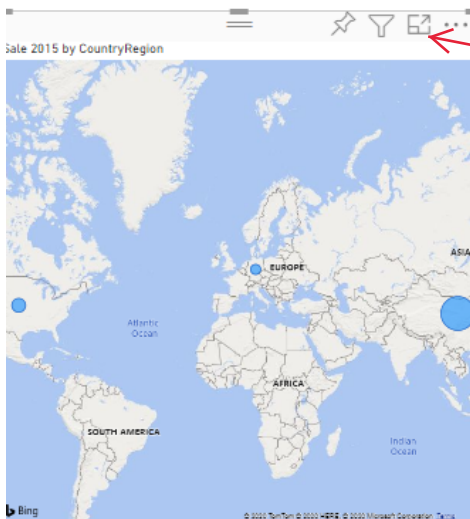


Figure 1-18: With the correct visualization, numbers are much more meaningful, e.g. Contoso has the most sales in 2015.

Note Before proceeding further, feel free to experiment by using different visualizations for the same data. As you will discover, each visualization offers a different insight from the same numbers. With Power BI you can use different visualizations to find the best way to tell a story about your data, using the same numbers.

So far, you've learned how to create an individual chart. You will add another visualization. **Click on an empty area in the central canvas** and place a tick next to the **CountryRegion**, a world map appears. Tick **Sale 2015**, you will generate a **new tile**, this time displaying a map with sales in the three countries/regions contained in the demonstration dataset, as below. Before doing this, return to **View/Fit to Page**



The **Focus Mode** button allows you to view the tile in Full Screen mode.

Click **Back to Report** to return to the Report view

Figure 1-19: A map of the world showing sales in different countries/regions, highlighted as bubbles.

Now, maps are powerful charting tools, but, as Figure 1-19 demonstrates, by having only three values for China, United States and Germany (Europe), they look dispersive. There are too many details in the map, whereas the goal is to show only the relative size of three areas. In this case, a column chart does this job well. You can transform the map by keeping it selected then select the **Clustered Column chart** from the **Visualizations** pane, then move the two visualizations so that they look like those shown in **Figure 1-20**.

Note you can collapse the panes on the right by clicking the small arrow on the top right of each pane.

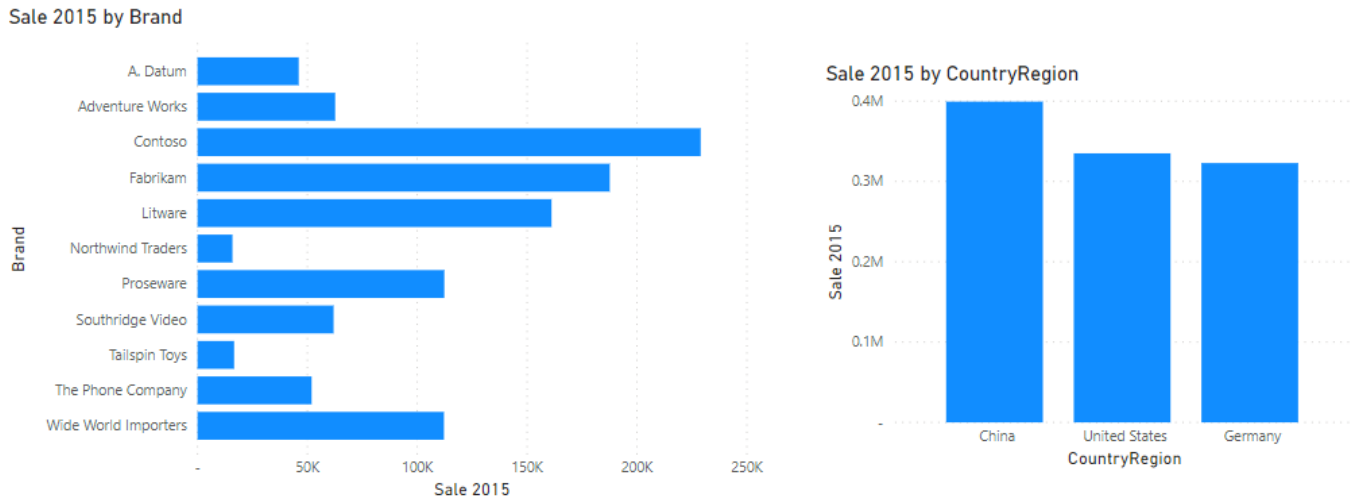


Figure 1-20: A report can contain multiple visualizations. Rearrange the tiles as above, or to best suit your screen

As you have seen, a report is a collection of visualizations organized in such a way as to communicate insights about the data. In Figure 1-20, a reader has an immediate feeling that sales in China, Germany, and the United States are nearly the same. Also it is clearly evident that there are only a few brands that make up most of the sales (Contoso, Fabrikam, and Litware), whereas others (Northwind Traders and Tailspin Toys) produce only a relatively tiny amount of sales.

Introducing Visual Interactions

This feature is very similar to what David could have achieved by using Excel and a couple of pivot tables on top of the table containing sales, yet there are some important differences between a report created in Excel and the same report done by using Power BI. We will look at those as you proceed through the book, but, for the moment, let's look at the interactive nature of Power BI reports.

In the top chart from Figure 1-20, **click the column for Germany**. As soon as you click an element within the chart, the entire report is **filtered** showing the contribution of Germany to sales of different brands, by means of coloring with two shades the Sale 2015 By Brand visualization, as depicted in Figure 1-21.





Figure 1-21: Clicking on a column in the column chart filters the bar chart, highlighting the contribution against the total.

By doing this simple operation, **David notes that sales of Northwind Traders in Germany are tiny when compared with China and the United States.** Clearly, **that brand is not popular in Germany,** and David is curious to see whether it is sold in equal volumes in China and United States or whether one of those countries/regions has much more sales than the other one. To perform this analysis, he clicks the bar for Northwind Traders. By doing so, the filter will move from the country/region to the brand and, as it happened before, the country/region chart will highlight the contribution of Northwind Traders to the total sales, as shown in Figure 1-22.






Figure 1-22: Filtering one brand shows the contribution of the brand against the total of sales by country/region.

The chart with sales by country/region already shows that a majority of sales are in United States, but because David is analyzing a very small brand, the chart is not clear in terms of relative importance of sales in different countries/regions.

Before we move on, we now need to be a bit more accurate in describing what we are seeing. Any chart produces graphical visualizations of the underlying numbers. **Any of those visualizations can behave as a filter**, and such a filter is activated by simply clicking the chart. So far, you have seen that a filter—when applied to other charts—highlights the relative contribution of the filtered item against the grand total by using two colors. This behavior is known as **visual interaction**, and it is extremely interesting. Yet, there are scenarios, like the one David is experimenting with, for which it would be better to compare the differences between countries/regions more than the overall contribution of a brand against the other brands.

You can configure visual interactions in a highly precise way. Namely, you can configure how the filtering on a chart behaves with respect to all of the other ones. The scenario we are looking at—with only two charts—is perfect for experimenting because it is very simple. To configure visual interactions, on the top menu bar of the report, click the Visual Interactions button, which you can see highlighted on the right in Figure 1-23 *Note: if you cannot see this button Click **Edit Report**.*

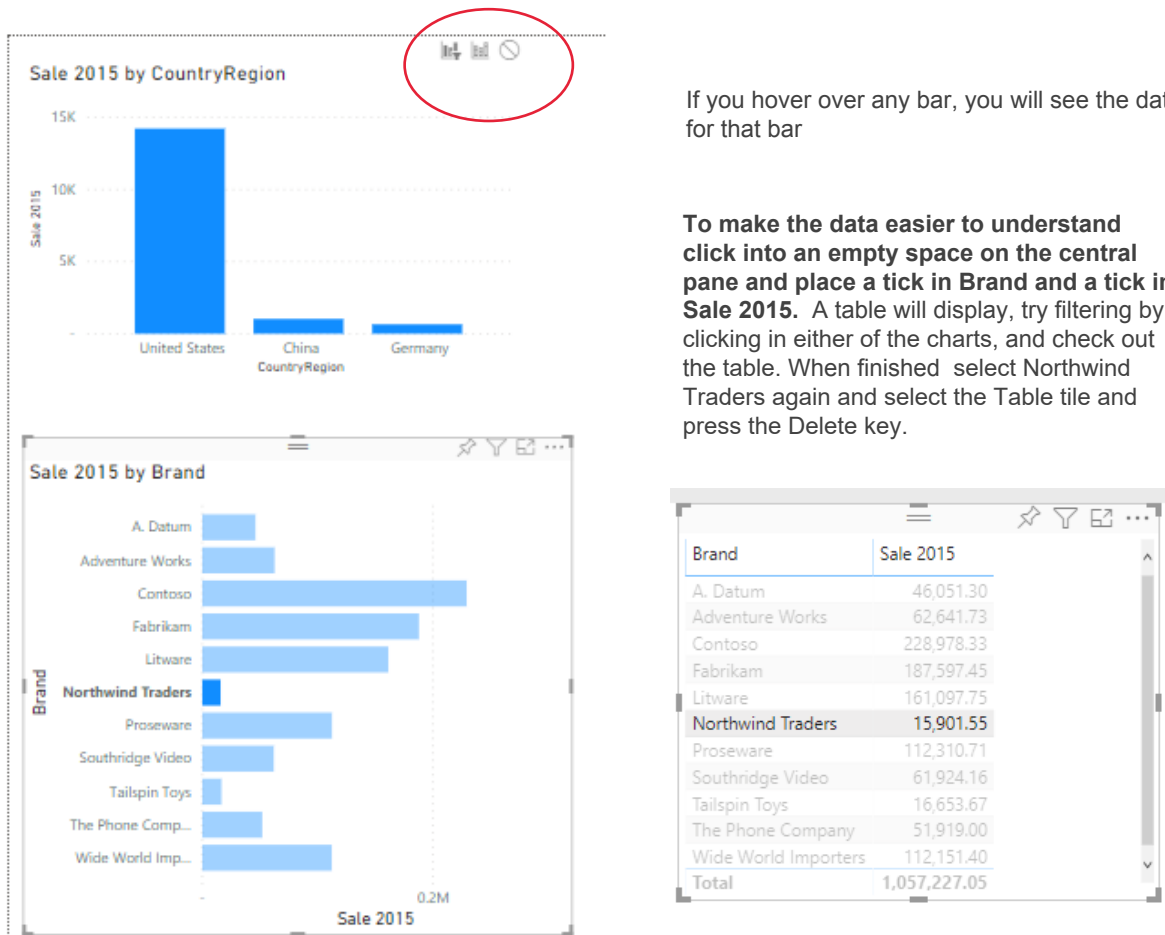


- The first button is a **Filter** button (a small funnel icon on a column chart). When you click this, filtering the selected chart will place the very same filter on the destination chart. In such a case, you will not see the contribution of the selection to the total. Instead, you will see only the selection in the chart, excluding the values (and corresponding areas) related to unselected items.
- The second interaction is (small column chart icon) is **Highlight** that is, the relative contribution from the selection on Sale 2015 will be highlighted in Sale 2015 by CountryRegion. This is the default filtering behavior, where the filtering on one chart shows, on the destination chart, the relative contribution of the selection against the total.
- The third is the **No filtering interaction** - When this is the selected behavior, filtering the selected chart has no effect on the target chart.

When you **turn on visual interactions**, each chart shows a different set of icons. The one you select (in this example, Sale 2015 By Brand) shows the standard selection icon, whereas the Sale 2015 By CountryRegion shows the three different kinds of interactions you can choose:

(Note: To remove a filter, click on the bar a second time). Move the Sale 2015 by Region to the top and Sale 2015 by Brand to the bottom, as in Figure 1-24

For example, you can select the filtering interaction from the sales by brand to the sales by country/region. By doing so, when you select Northwind Traders, the resulting report will show a different result, as shown in Figure 1-24 where the United States has the highest sales for the Northwind Traders brand.




If you hover over any bar, you will see the data for that bar

To make the data easier to understand click into an empty space on the central pane and place a tick in Brand and a tick in Sale 2015. A table will display, try filtering by clicking in either of the charts, and check out the table. When finished select Northwind Traders again and select the Table tile and press the Delete key.

Figure 1-24: By using the filtering behavior, the relative size of the bars in the Sale 2015 By CountryRegion chart is more meaningful.

Now, when you browse the report, you can quickly click a brand and see which country/region sold more than the others. Because of automatic determination of the scale, the insights are much clearer.

To switch off **Visual Interactions**, click the tab at the top of the screen and switch off **Edit interactions**.



Decorating the report

In the previous sections, David performed some analysis on the data, and now he thinks that his first report, although simple, contains some findings that are worth sharing. He can obviously make a screenshot and attach it to an email with some description, but Power BI offers some tools that make it possible for him to annotate a report with remarks.

David can add text to the report and decorate it with shapes. For example, he can add a colored arrow to Northwind Traders and a text box with some remarks about what he found. When David does this, the report is easier to read, as demonstrated in Figure 1-25.



Figure 1-25: Decorating the report makes it easier for the reader to immediately understand the insights.

To add the text, above the central pane, David clicks **Text Box**. When the text box appears, he types and can use the pop-up format text options to change the appearance of the text. To **add the arrow**, again above the central pane, David clicks **Shapes** and then select the **Arrow shape**. He then moves the arrow into position and resizes it.

When you add the text box or the arrow, you need to set some properties for these objects. These are under the **Format Shape** pane. In fact, the appearance of each object (either a decoration or a full chart) in a Power BI report is controlled by a set of properties that you can access by clicking the object in the central pane. The properties you set appear in the Visualizations pane, as shown below.

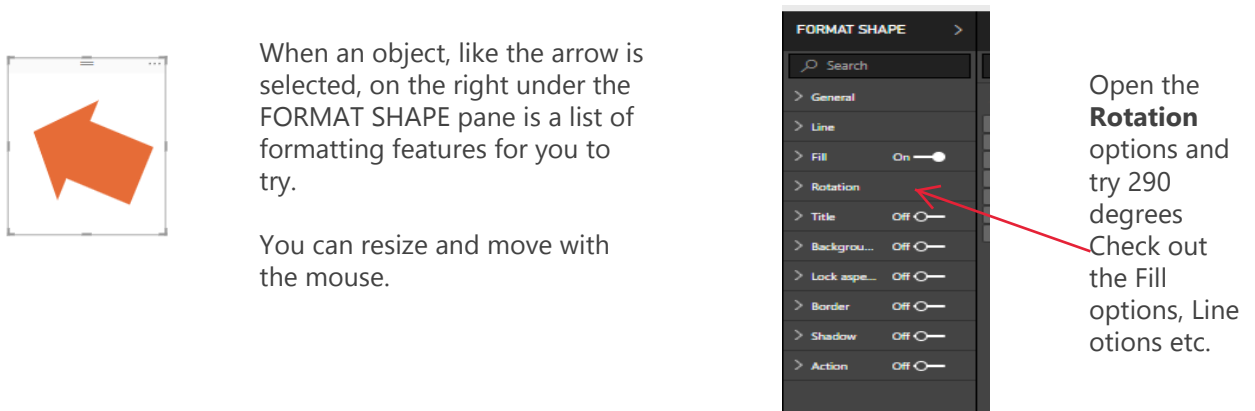


Figure 1-26: Each visualization has a set of properties that you can adjust to customize it.

For example, to rotate the arrow, David selects it and then, in the **Format Shape** pane on the right, he clicks **Rotation** and **drags the slider**. Similarly, he also **changes the fill color**.

Finally, keep in mind that visual filters on reports (that is, filters that you set by clicking a chart item) are not saved as part of the report. Thus, when you look at the report again, the arrow is useful to tell you where to apply the filter to see the data. Later in this chapter, you will learn how to place a permanent filter on a report.

Saving the report

At this point, David can save his report so that he can continue working on it later. To save a report, go to File, click Save (see Figure 1-27), and then provide a name for the report. Here, David saves his report with the name **Northwind Traders**.




Figure 1-27: When you finish editing a report, saving it is always a good idea.

After you save the report, it appears in the **My Workspace** pane, in the **Reports** section. You can now access it any time you sign in to Power BI.

Click on **My workspace**, under Reports, click **Northwind Traders**.

When you select a saved report, it opens and remains in **read-only mode** until you explicitly activate it for editing by clicking the **Edit Report** button highlighted in Figure 1-28.




Figure 1-28: You need to click the **Edit Report** button to bring a saved report in edit mode.

This behavior is useful to avoid unintentional editing of the report. A saved report can be easily viewed, and it always reflects the latest data. If you need further filtering or you want to perform a different analysis on the same report, you need to turn on edit mode.

Pinning a report

When you open a report in read-only mode, the menu bar at the top of the screen offers you several actions: for example under File you can choose to save a copy of the report, print it, export it etc. One of the menu bar items is worth a few moments of our attention: **Pin to dashboard** (see Figure 1-29). You can choose to pin a single visualization or every visualization on the page

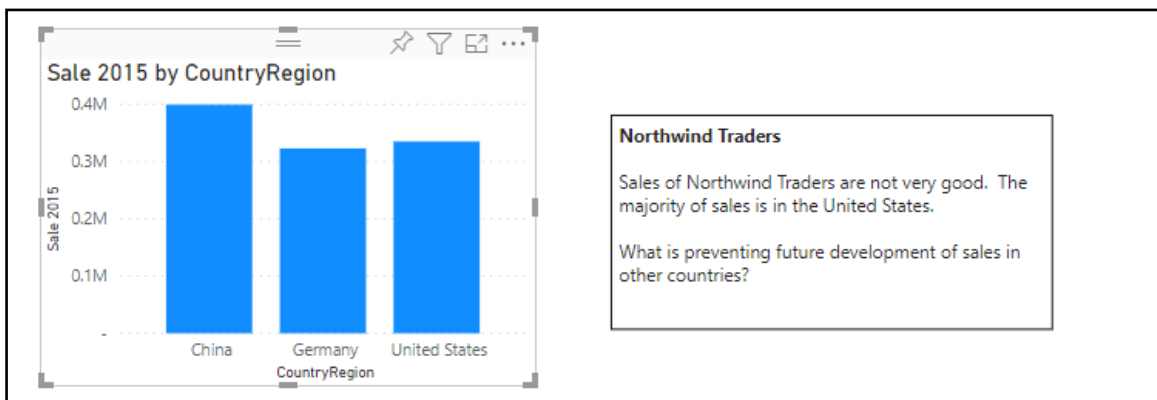



Figure 1-29: Visualizations aren't the only things that you can pin to the dashboard: You can pin reports, too.

What is the **difference between pinning a visualization and pinning a full report**? When you **pin a visualization**, Power BI **saves it as it is with all the visualizations, text and objects either into an existing visualization or allows to create a new one**. To try this, go to My Workspace, Northwind Traders and click Pin to a dashboard from the top of the screen,

When you **Edit the report** and pin the selected visualization, it is disconnected from any others in the same dashboard. Thus, any visualization in a dashboard does not include the visual interactions of other visualizations. This is usually good, because a dashboard is not intended for interaction. If you need interaction or further analysis, you can always click a visualization from the dashboard to open the source report.

To try this one, click **Edit report from the top of the screen** and select only the **Sales 2015 by CountryRegion** chart, the columns return to full color, click the **pin** and create a **New Dashboard** named **Sales 2015 by Country and Region**. Click **Pin**. Click **Go to Dashboard** on top right, and it will be displayed in its **own dashboard**. Click **My Workspace** and select **Dashboards** and click **Sales 2015 by Country and Region**. (note a single click will display the Chart but a double-click returns to the original charts.)



To return to your report **Northwind Traders**, go to **My Workspace** and click on **Reports** and double click

Northwind Traders with the 2 visualizations (charts), a text box and an arrow.

Refreshing the budget workbook

David receives figures from the country/region managers in a very simple form: They each send him a workbook with forecasts based on the brand, with no monthly details. Figure 1-30 shows what the forecasts look like for China.

CountryRegion	Brand	Budget
China	A. Datum	5,000.00
China	Adventure Works	36,500.00
China	Contoso	94,500.00
China	Fabrikam	90,800.00
China	Litware	65,000.00
China	Northwind Traders	12,000.00
China	Proseware	20,000.00
China	Southridge Video	22,500.00
China	Tailspin Toys	8,000.00
China	The Phone Company	30,000.00
China	Wide World Importers	45,000.00

← Figures in Budget are yearly

Figure 1-30: An example of forecasts received from the China manager.

The Budget forecasts are set at the year level, but David sets them up at the month level, because of he decides to divide the yearly sales by 12 and copies the result to his own Excel file named Sales and Budget manually combined, into a new column called Budget. Open the Excel files located in Chapter 1, as below and check them out.

The workbook with the new Budget column now looks like Figure 1-31. Looking at the numbers, it is clear that David could have done it differently, because the numbers do not reflect the seasonal nature of sales and, more important, they are not correct. We will fix this and use a better technique later in the book. Right now let's focus on Power BI. Remember, this is not a book about budgeting techniques.

Budget.xlsx

Brand	Budget	Column1
A. Datum	5,000.00	\$ 416.67
Adventure Works	36,500.00	\$ 3,041.67
Contoso	94,500.00	\$ 7,875.00
Fabrikam	90,800.00	\$ 7,566.67
Litware	65,000.00	\$ 5,416.67
Northwind Traders	12,000.00	\$ 1,000.00
Proseware	20,000.00	\$ 1,666.67

Sales and Budget manually combined.xlsx

CountryRegion	Brand	Month	Sale 2013	Sale 2014	Sale 2015	Budget
China	A. Datum	January	3,234.00		1,935.00	416.67
China	A. Datum	February	6,270.00	7,059.00		416.67
China	A. Datum	March	4,352.00			416.67
China	A. Datum	April	3,814.00			416.67
China	A. Datum	May	6,234.00			416.67
China	A. Datum	June	5,571.00	3,216.00		416.67
China	A. Datum	July	7,424.00			416.67
China	A. Datum	August			800.00	416.67
China	A. Datum	September	1,254.00	1,617.00	396.00	416.67
China	A. Datum	October	1,881.00	3,042.00	936.00	416.67

Check out the formula in Column1

The pasted values in the Budget column are not formulas

Figure 1-31: The 2015 Sales workbook now contains monthly sales in the last column, named Budget.

David exits Excel and returns to Power BI. He uses Get Data and locates the Excel file Sales and Budget Manually Combined and imports it into Power BI, Figure 1-32. A message will appear informing you that the dataset is ready, click View dataset.





Figure 1-32: The message box that offers you to View Dataset or Get Quick Insights.

A blank canvas appears, with the Column Names listed on the right.



Note Be aware that there might be a delay of several minutes from when you upload a new version of a dataset until the new columns and tables appear in Power BI. The exact timing depends on whether there is a recent release of Power BI and is subject to change in the future. If, for any reason, you do not see updated information after a new upload, just wait a few minutes and try it again; Power BI is being refreshed and nothing is going wrong.

Using the procedures you used earlier in this lesson, use the **Visualizations** and **Fields** construct the following report depicted in Figure 1-33. A lot of formatting has been used, you can chose your own colors, fonts, sizes, etc which are available to your for each visualization tile.

It is a good idea to select the chart visualization first, then populate with the required fields.

Hint - for the first chart in the report, the **Donut Chart**, using the **Budget** and **CountryRegion** fields. You will find all the options you need In the **VISUALIZATIONS** pane, under the **Fields** and **Format** buttons. **Details** are **CountryRegion**, **Values** are **Budget**, **Data labels** are on, the **Label Style** should be **Category**. **Legend** is off. The **Title** can be modified also. Experiment and try to match the sample below.





Figure 1-33: With the Budget in the model, reports are richer and provide better information.

Observe that the two visualizations at the bottom use a different visual interaction method. **The Clustered bar chart on the bottom left shows the contribution of a brand to the overall sales and budget, whereas the one on the right is more useful to perform a comparison of budget and sales in different countries/regions.** Both are useful and provide different insights. The technique of adding multiple copies of the same visualization with different visual interactions is common, and we encourage you to learn to use it.

To enhance the report, David tries out different titles for the visualizations (using color, font style and sizes). You can manage these visualization details by using the Format brush icon highlighted in Figure 1-34, which shows the many options to configure a visual. In the example, we used a custom title and changed the font size.




Figure 1-34: You can configure many aspects of a visualization by using the formatting options.

You will learn many more details about visualization formatting as the book progresses. However, for an introductory chapter, this is enough. It's time now to draw some conclusions.

Filtering a report

You already learned about the visual interactions feature, which makes filtering a report easy, simply select a bar and all visualizations change to show only that data highlighted, if you hover over the highlighted bar you can view the related data.

You can set all of the filters in the Filters section in the **Visualizations** pane. Figure 1-35 illustrates that for David's report, there are three kinds of filters, **the Brand, the Budget data and Sale 2015 data.**

You can set all of the filters in the Filters section in the **Visualizations** pane. Figure 1-35 illustrates that for David's report, on Sale 2015 and Budget by Country Region, there are three kinds of filters, listed on the right side of the screen under the **Filters** pane, **Brand**, **Budget data** and **Sale 2015** data.

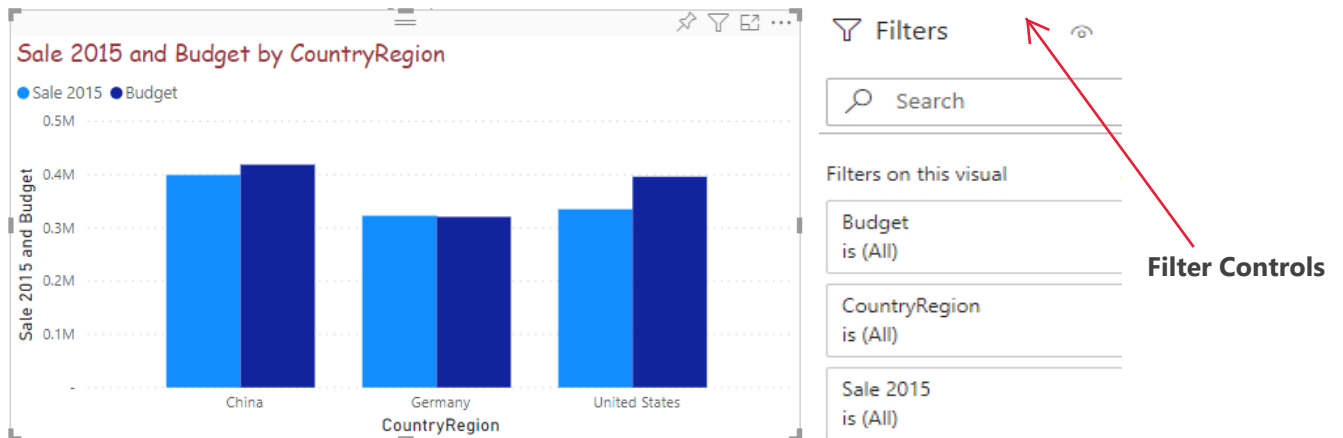


Figure 1-35: You can configure filters in the Filters pane.

For example, Figure 1-36 below, this shows the result if you add a page-level filter to the report, selecting only China and Germany. Click in an empty part of the screen, to de-select all tiles. Under the **Filters** pane, drag **CountryRegion** into the **Filters on this page** box. Remove **USA** from **CountryRegion** list and all the Visualizations on the page will display data for the 2 countries.

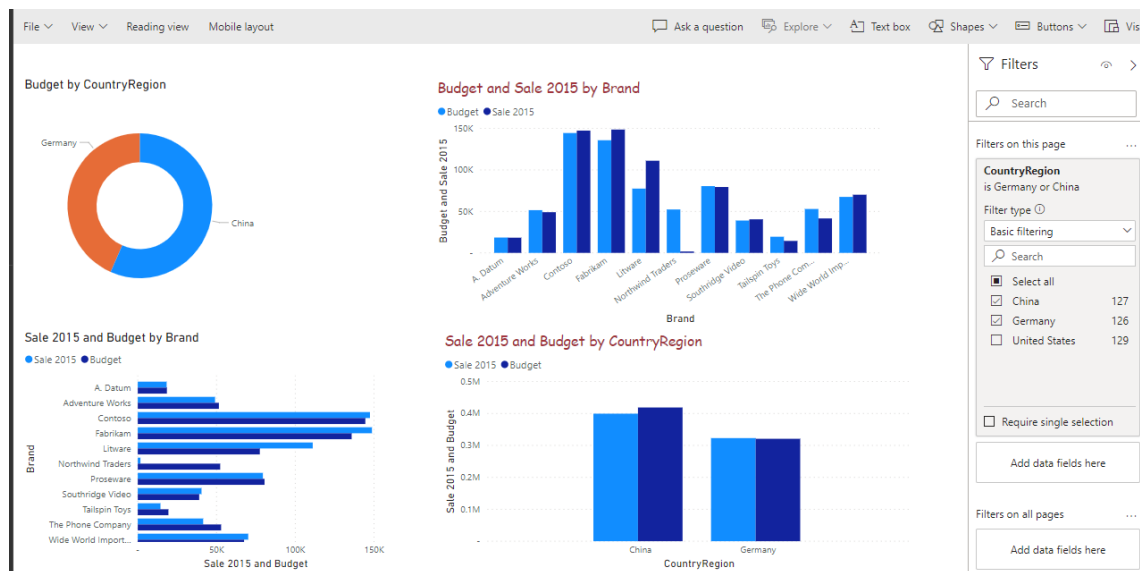


Figure 1-36: The same report as the report presented in Figure 1-33, this time filtered using only **China** and **Germany** applied to all visualizations on the page.

Filters at the report and page level behave the same way. Filters on a visualization, on the other hand, have an additional feature: they can filter both data (as was the case for the country/region) or the metric associated with the chart. A good time to save the file. Click File and save as **Sales and Budget**.

For example, select the chart, **Budget and Sales 2015 by Brand**. You can filter this chart to include only values for which the **Budget is greater than 50,000**. Figure 1-37 presents the result.

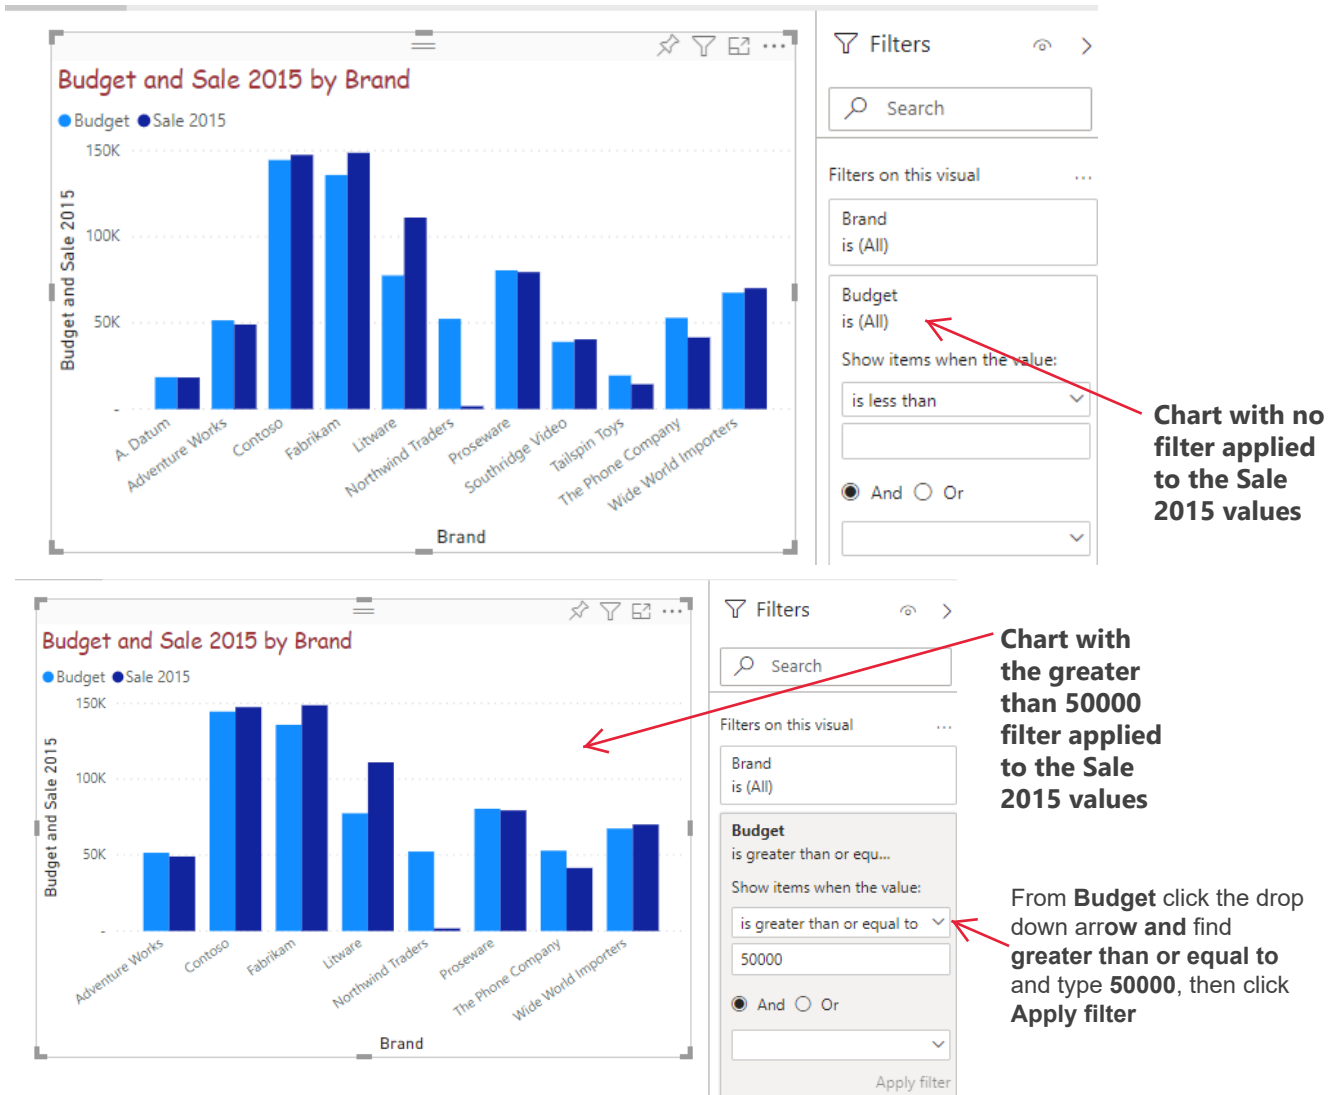


Figure 1-37: Chart after the greater than 50000 for Sale 2015 has been applied

Notice in Figure 1-37 that the number of brands is much less than that of Figure 1-36. This is because the latter report shows only brands that have a Sale 2015 measure greater than 50,000.


All of these filters are saved as part of the report, and they are not shown in any visual way. For this reason, if you prepare a report that, for example, filters only 2015, it is always useful to add a description of the filter as part of the report title—"Sales in 2015" instead of "Sales."

Conclusions

After this first tour in Power BI, it's now time to take a breath and describe what we've learned so far.

- Power BI is a cloud service that provides tools to perform analysis of data and gain insights from your numbers.
- To build a dashboard, you need a dataset, a report, and, finally, the dashboard. The dataset is the source of data, reports are useful to create visualizations that might be connected through visual interactions, and a dashboard is a collection of visualizations and/or reports.
- You can create visualizations by using natural-language queries, Quick Insights, or full reports.
- You can decorate a report by using text boxes, shapes, and pictures.
- Visualizations in a dashboard are not connected through visual interactions, which work only among visualizations in a report. If needed, you can pin a report as a live page in a dashboard to maintain the interaction capability.
- You can load data in Power BI from many different sources. So far, David has used only an Excel workbook, but there are many other sources that he will learn to use before becoming a Power BI expert (and that you will learn about later in this book).
- You can refresh the content of your workbook by uploading a new version of it. But, as Chapter 2 describes, there are better ways to refresh data.
- You can apply filters by using visual filters, which produce highly interactive reports, or you can use static filters, which you can apply at the visual level, page level, and report level. Static filters are saved as part of the report, whereas visual filters are not.

You can refresh a dataset immediately by clicking **Refresh** button.



When you ask for a refresh, Power BI prepares for the data refresh and then starts it. Depending on the size of the dataset and the speed of the Internet connection, the time to carry out the refresh can range from a few seconds or much longer.


Sharing the dashboard

In Chapter 1, David, our manager of budgeting at Contoso, created his first dashboard with analysis of sales. But that dashboard is only the starting point for creating the budget for 2016. David will involve his colleagues in this process, so the first thing he needs to do is share the work he has done thus far with these colleagues. After that, he must collect feedback and numbers from other managers in order to complete the entire budget. Also, he needs to choose how to share the resulting data and reports between members of the team.

In this chapter, you will see how David can use the features of Microsoft Power BI, along with other services, to achieve his goals.

Inviting a user to see a dashboard

David wants to **share** the dashboard he created with a colleague, the country/region manager for Germany. To do that, he clicks **Recent** or **Favorites** on the left of the screen, and double clicks **Northwind Traders** to open the dashboard and then clicks the **Share** button located at the top of the screen as below. To see the **Share** option, you must turn on the new look by clicking the button at the top of the monitor.



When the **New Look On** is switched on you will see the Share button on the left along with other buttons.




Figure 2-1: To share a dashboard, click the Share button in its upper-right corner.

This opens the **Share Dashboard dialog box** in which David can send Wendy an invitation. Before we continue along with David, let's take a closer look at this dialog box. The Share Dashboard dialog box has two tabs: Invite and Shared With. The Invite tab (see Figure 2-2) is where you provide the email address (or addresses) of the people with whom you want to share a dashboard. You also can include an optional message that you would like your invitees to receive from the Power BI service. At the bottom of the tab are the Allow Recipients To Share Your Dashboard check box and the Send Email Notification To Recipients check box, which are self-explanatory. If you decide to send an email as a notification to recipients, Power BI will automatically include a link to open the dashboard in the email that is sent.

Notice in Figure 2-2 that as David types Wendy's name in the name box, Power BI suggests email addresses. Because Wendy is a coworker in the same organization as David, Power BI is able to offer these suggestions (more details on this later). However, you also can type an email address in this box if the email address of the person with whom you are sharing does not have the same domain name as your address (for more details about this, read the section "Inviting users outside your organization," later in this chapter). You can add more than one person if required.

Share report
NORTHWIND TRADERS

Share Access

Only users with Power BI Pro will have access to this report. Recipients will have the same access as you unless row-level security on the dataset further restricts them. [Learn more](#)

Grant access to

Enter email addresses

Include an optional message...

- Allow recipients to share your report
- Allow recipients to build new content using the underlying datasets
- Send an email notification to recipients

Report link ⓘ

https://app.powerbi.com/groups/me/reports/c1149139-deb5-4a0e-bf65-3f3be7c

Share Cancel

Figure 2-2: In the Share Dashboard dialog box, you specify the list of people invited and, optionally, include a message for them.

If you clear the Send Email Notification To Recipients check box (see Figure 2-3), you will need to go on the Shared With tab to copy the link to your dashboard that you will then send to recipients after you click the Share button. The email addresses you provide will give those accounts access to your dashboard, but those individuals will not receive an email notification.

Allow recipients to share your dashboard
 Send email notification to recipients

Share Cancel

Figure 2-3: You can share a dashboard by copying a link instead of sending an email from the Power BI service.

On the Access tab (see Figure 2-4), the Dashboard Link box contains the URL for the dashboard. If you choose to not send a message via the Power BI service to your recipients or you simply prefer to use your own email account to do so, you will need to copy this URL and send it to your invitees. Using your own account can be helpful when you want to ensure that your recipients recognize that the incoming email is from you, as opposed to seeing “no-reply@powerbi.com,” which they, or their email client, might filter out as spam. Also on the Access tab is a list of users with whom you have shared the dashboard, along with their assigned privileges.

Share dashboard

Share Access

The following have access to this dashboard

David Bradley	Owner
Wendy Kahn	Can view ...

Dashboard Link ⓘ

<https://app.powerbi.com/groups/me/dashboards/415b3db6-906b-457c-9d58>

Figure 2-4: The Access tab shows the users who have access to the dashboard and provides a link to share it.

The email message that Wendy receives looks similar to that shown in Figure 2-5. If she has not previously used Power BI, when she clicks the link to open 2015 Sales.xlsx, she will be directed to the Power BI website where she will need to register with the service, just as David had to do when he first used Power BI (see Chapter 1 for a refresher on getting started with Power BI). If she is already enrolled as a Power BI user, she will go straight to the dashboard that David is sharing with her.

David Bradley has shared Power BI Dashboard '2015 Sales.xlsx' with you 

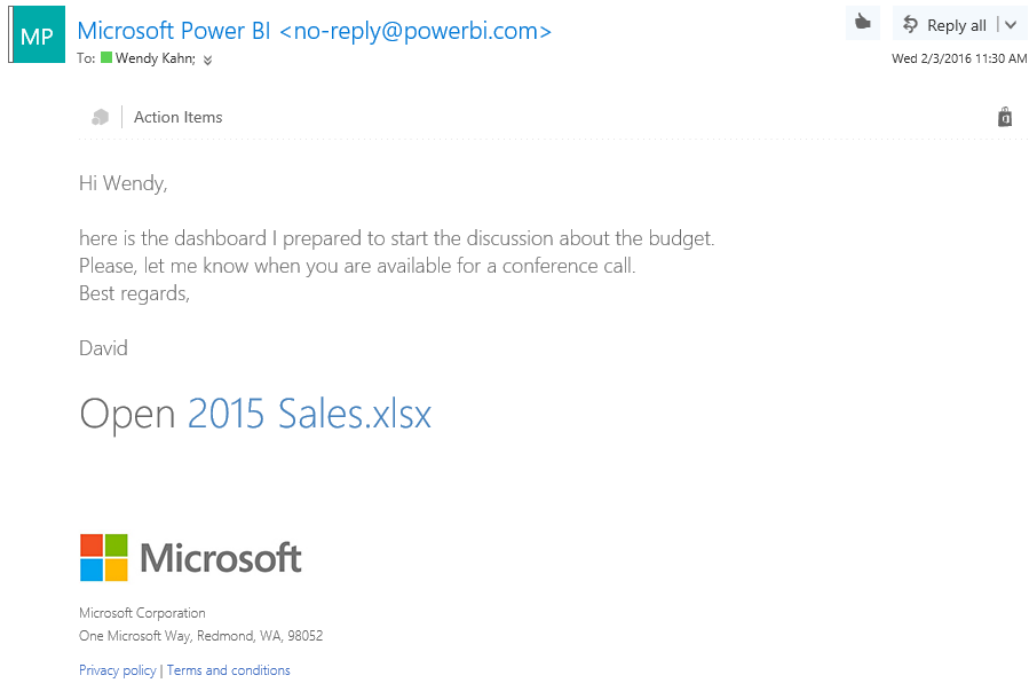


Figure 2-5: Power BI sends an email message to users invited to share a dashboard.

Now, Wendy is looking at the same dashboard as David, including the reports underlying the visualizations pinned to the dashboard. However, Wendy cannot modify either the dashboard layout itself or any of the single reports; at this point, she has read-only permission. For David to make it possible for other users to edit his dashboards and reports, he needs to create a group workspace, which we will cover a bit later in the chapter.

After David invites Wendy, she can open the dashboard in its own Power BI session. Dashboards shared by another user appear on a guest user's Workspace pane with a "shared" icon adjacent to the dashboard name, as shown in Figure 2-6. This indicates that the dashboard is read-only. So in this case, Wendy cannot change or modify the content of 2015 Sales.xlsx, but she can interact with reports pinned to the dashboard and can open the reports underlying each visualization by simply clicking a visualization. Even if such reports are not listed in the workspace, Wendy can open them through the shared dashboard, but she cannot modify their content (the Edit Report feature is turned off in these cases).

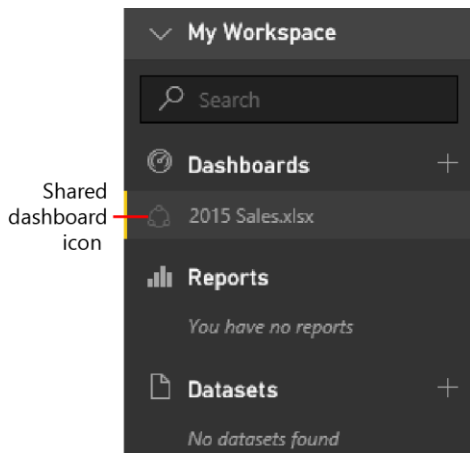


Figure 2-6: Shared dashboards display a “shared” icon before the name in the list of dashboards.

Sharing via content packs

An additional technique to share reports and dashboards within an organization is through a *content pack*, which is a set of datasets, reports, and dashboards that a user can copy within his personal workspace. David might consider using this feature to deploy a report to other users, but he does not use this system in our scenario because it is a technique that is better suited to distribute a set of predefined reports and dashboards that other users can customize in their personal copy. Content packs are not designed to share reports between users in an active way, as David needs at the moment. Chapter 5 shows you how to create and consume content packs from the public gallery and within an organization.

Inviting users outside your organization

Thus far, David has invited Wendy to view his dashboard; she works in the same company and has an email with the same domain (@contoso-bi.com). But what happens when David wants to invite someone who is not part of the same company? Answering this question requires some explanation.

Although Power BI is designed for you to share a dashboard with users who are within the same organization, you can also share dashboards with people from other organizations. The way Power BI identifies “an organization” can be described as follows:

- Every user requires an email address within the domain of the company.
- Power BI does not accept generic email domains such as hotmail.com, gmail.com, and so on. Your company needs a unique domain name, and all of the users must have an email address within that domain. All of the users having an email within the same domain are considered part of the same organization.
- If you use Microsoft Office 365 and/or Microsoft Azure Active Directory, you might have different domains belonging to the same organization. This is the only case for which users having email with a different domain name belong to the same organization for Power BI.

Note If you are not sure whether your organization already uses Office 365 and Azure Active Directory, ask your IT administrator, and if he would like to read more technical details about authentication in Power BI, refer him to the following documents:

<https://powerbi.microsoft.com/documentation/powerbi-admin-power-bi-security> and <http://go.microsoft.com/fwlink/?LinkId=619090> (which downloads the Power BI Security white paper)

On the surface, that seems rather restrictive, but in reality, you also can share a dashboard with users in other organizations, using the same method as that described in the previous section. However, when you specify an email address with a domain other than that of your organization, you will see a message similar to one shown in Figure 2-7, which David receives when he tries to share a dashboard with a vendor.

Share dashboard




Figure 2-7: The message that displays when you try to share a dashboard with someone outside of your organization.

It is important to understand the difference between a user within your organization (internal users) and outside of it (external users):

- **Internal users** You can invite internal users to share a dashboard by email or by sending them the URL of the dashboard. In the latter case, users must be authorized. If a user does not have authorization, she can ask for permission when she clicks the dashboard URL.
- **External users** You can share a dashboard with external users only by inviting them by email. When an external user receives the email, she must sign in to Power BI using the same email account used in the invitation. If she never previously used Power BI, she can create a free account the first time she signs in.

Finally, you can publish a report (but not a dashboard) on the web. To do so, select the report, click the File menu, and then click **Embed** as depicted in Figure 2-8. In the Embed In A Public Website (Preview) dialog box, click Create Embed Code. For this reason, you should use this technique only when you want to publish information intended for public consumption; for example, a report embedded in the public website of your company.

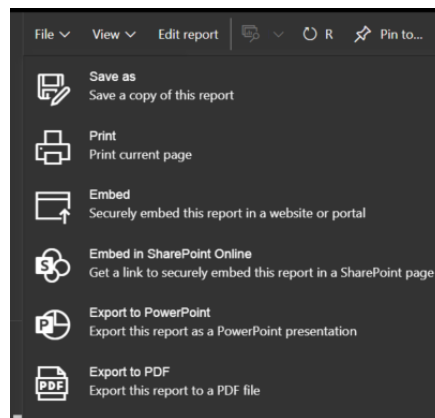


Figure 2-8: The File menu includes the Embed command, which create a code to link a report available on the Internet.

The Publish To Web feature guides you in creating a public webpage, getting a URL that you can send in an email, or the HTML code required to embed the report in a page of a website you own. For more technical information about publishing a report to the web and to get a detailed step-by-step guide, go to <https://powerbi.microsoft.com/documentation/powerbi-service-publish-to-web/>.

Creating a group workspace in Power BI

Let's return to David and Wendy. After David invited Wendy, he realizes that he will need to repeat the same share operation for every dashboard he creates. Moreover, as soon as other people become involved in the budgeting process, he will need to send them the invitation for all of the dashboards he shares with the group. Fortunately, David discovers that he can create a group of users with whom he can automatically share all of his dashboards, and also provide editing rights to certain users within that group. By creating groups of users in Power BI, you increase the level of collaboration among them.

The only caveat is that you must have a Power BI Pro license to have access to the group feature; you cannot create groups by using the free version of Power BI. However, you can try Power BI Pro for 60 days free of charge, giving you an opportunity to evaluate this feature and determine whether it is good for your company.

Assuming that you—and David's organization—opt to purchase a Power BI Pro license, to create a group, in the My Workspace pane, immediately below the list of workspaces, click the "+" button to the right of Create A Group, as demonstrated in Figure 2-9. (You might need to click My Workspace to open that pane.)

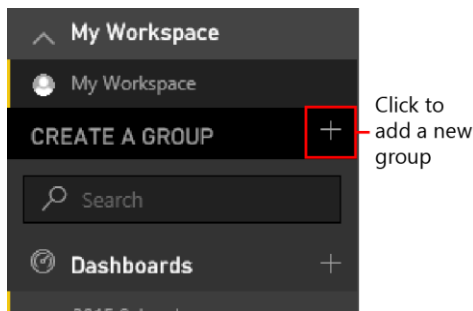


Figure 2-9: You can create a group by clicking the plus symbol (+) to the right of Create A Group.

David clicks the "+" button, which opens the Create A Group dialog box in which he creates a group named Budget 2016. This group will initially include himself as administrator and Wendy as a member. In the Privacy section, shown in Figure 2-10, you can define the privacy levels of the group.

Save
Cancel

Create a group

Name your group

Group ID

 ✎

Available

Privacy

Private - Only approved members can see what's inside ▾

Members can edit Power BI content ▾

Add group members

Add

david.bradley@contoso-bi.com	Admin ▾	✕
wendy.kahn@contoso-bi.com	Member ▾	✕

Figure 2-10: The group includes a list of members and privacy settings that group administrators can change later.

Every group has two privacy settings. The first determines whether the group is visible only to its members or also by other users within the organization who are not members of the group. Here are the possible choices:

- **Private** Only approved members can see the results of the group’s activities.
- **Public** Anyone can see what the group is doing.

The second setting specifies whether all the group members can modify the contents of reports and dashboards. There are two choices:

- Members can edit Power BI content.
- Members can only view Power BI content.

If you select view-only for this second setting, only group administrators can edit dashboards.

With David having configured the group as shown in Figure 2-10, Wendy will be able to edit the content of dashboards and reports included in the group. David and Wendy will see the public group in their list of group workspaces, as illustrated in Figure 2-11.

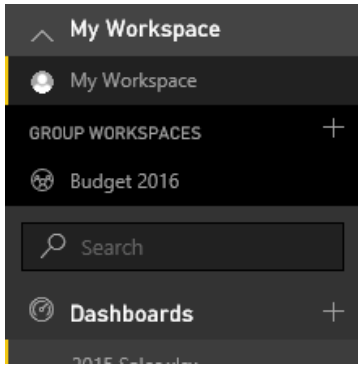


Figure 2-11: The list of group workspaces includes all the groups of which the user is a member.

Now that David has created a group, he can create reports and dashboards in the group that will be immediately visible by Wendy. However, he must import the data for these reports in the same group; he cannot move into the group what he already created in his personal workspace. Importing data and creating the reports will require some time at this point, repeating the same operation he has already done. Thus, when you know you will work with a team, it is a good idea to create the group at the outset, and save yourself a lot of redundant work later on.

Turning on sharing with Microsoft OneDrive for Business

Before moving forward, David wonders whether he will be able to share the data sources, not just the results (reports and dashboards). In particular, he wants to allow other colleagues to enter their data in Microsoft Excel files, so that he will be able to create a budget using data from workbooks modified by a number of other people.

In Chapter 1, David was forced to copy the budget data from Excel files received by country/region managers, and then he split those values by month in the table used to feed his Power BI report. Now, he wants to give other users the ability to modify the content of that Excel file directly so that he does not need to do all that work himself. For this reason, David creates a single Budget table in another worksheet of the same Excel file and copies to it the budget values by Country/Region and Brand, as depicted in Figure 2-12.

CountryRegion	Brand	Budget 2016
China	A. Datum	5000
China	Adventure Works	36500
China	Contoso	94500
China	Fabrikam	60000
China	Fabrikam	30000
China	Litware	60000
China	Litware	5500
China	Northwind Traders	9000
China	Proseware	20000
China	Southridge Video	22500
China	Tailspin Toys	8000
China	The Phone Company	30000
China	Wide World Importers	37500
Germany	A. Datum	13500
Germany	Adventure Works	15000
Germany	Contoso	50000

Figure 2-12: David's budget table now contains at least one row for each Country/Region and Brand.

David created a formula in the table used by Power BI that allocates the budget value over 12 months. At this point, David wants to share this workbook so that his colleagues will be able to directly modify the content of the budget table, and this should automatically apply the new values to reports and dashboards published in Power BI.

Using OneDrive for Business is the best way to share his source Excel workbook with other colleagues. You might already know OneDrive, which is the personal service with which you can store files in the cloud. But, even though you can share files on OneDrive, there are limitations when using it as a data source for Power BI, especially when it entails automatic data refresh. OneDrive for Business removes those limitations and provides more control, as well. Moreover, OneDrive for Business is integrated with Office 365 and directly supports groups, making it easier to share documents across your organization. To access OneDrive for Business, in the upper-left corner of the Power BI site, click the yellow button with the nine small squares in it (see Figure 2-13), and then click the OneDrive tile.

Note OneDrive for Business is a feature included in Office 365. If you do not have an Office 365 plan, you can subscribe to OneDrive for Business separately, without activating Office 365. If you are interested in using this feature, contact your IT administrator to determine which licensing option better fits your requirements.

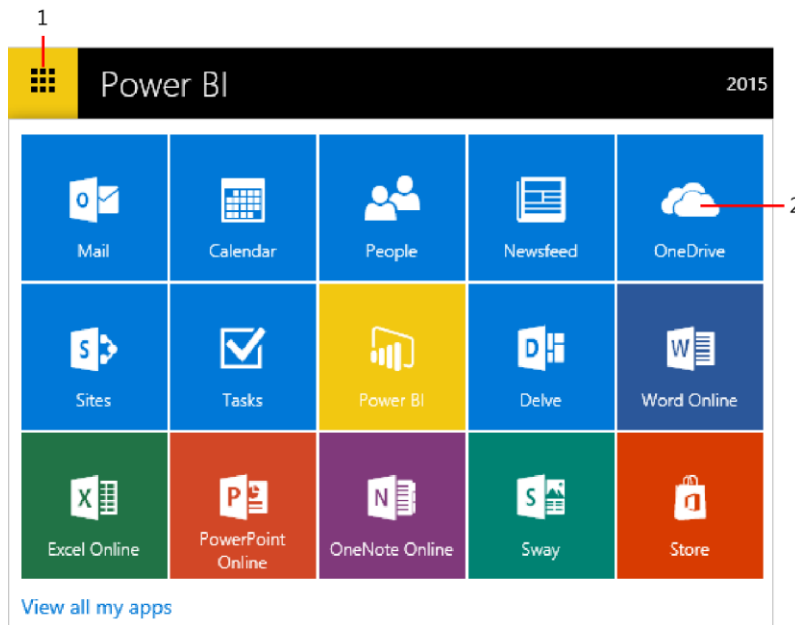


Figure 2-13: You can access OneDrive by clicking the button in the upper-left corner of the Power BI website and then clicking the OneDrive tile.

On the OneDrive webpage that opens, David uploads the workbook to the Budget 2016 group he previously created to share reports and dashboards. A group defined in Power BI corresponds to a group in Office 365, so you have an associated OneDrive for Business folder where you can place files to share. Figure 2-14 shows the sequence of steps that David needs to do to upload the document.




Figure 2-14: The actions required to upload a document in a group folder in OneDrive for Business.

When David finishes uploading the file, he can see it in the list of the files of the Budget 2016 folder, as depicted in Figure 2-15. Also in the figure, there is a Sync button which you can use to get information about how to synchronize the folder with a local computer, so that you can edit the file in a local folder of your PC and automatically upload any update to the shared folder on OneDrive. Thus, all of the files available in this OneDrive folder can now be shared among members of the Budget 2016 group, and the files will be available as a possible data source for reports created in the corresponding Power BI workspace.




Figure 2-15: After upload, the “Sales 2015 and Budget 2016” file is available and listed in the Budget 2016 group folder.

David has now shared a file to the Budget 2016 group. Later, he will ask other group members to update their budget data themselves. Before doing that, David wants to prepare a report that will display the aggregated total of the budget for every country/region, comparing it with the sales made in previous years.

Going back to Power BI, David opens the list of workspaces available (refer to Figure 2-11) and selects Budget 2016. Power BI displays dashboards, reports, and datasets for that group. Of course, initially all of these lists are empty, as shown in Figure 2-16.




Figure 2-16: The initial list of dashboards, reports, and datasets is empty for a new workspace group.

David wants to create a dashboard and a report based on data stored in the Excel file that now resides in the group folder on OneDrive for Business. The approach is similar to what he did when he made his first foray into Power BI (see Chapter 1), but instead of uploading a file from his own local computer (Local File), this time David selects the OneDrive tile to specify his data source, as shown in Figure 2-17. Notice, though, that the name on the tile is “OneDrive – Budget 2016.” Because David is using the Budget 2016 workspace, the associated OneDrive folder is automatically proposed as a possible data source. From a Power BI perspective, the biggest difference between a local file (one stored on your computer) and a file on OneDrive (stored in the cloud) is that the former cannot automatically update a report based on it, whereas the latter can propagate changes to data to Power BI reports without user intervention.

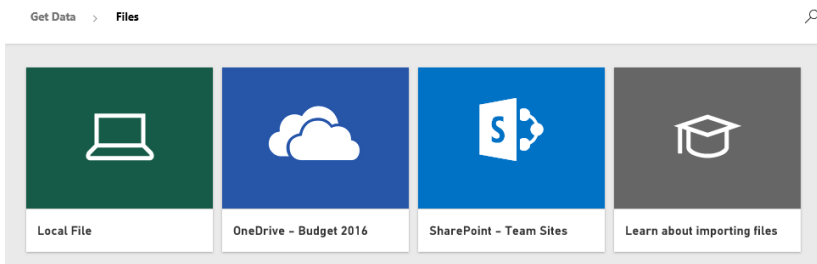


Figure 2-17: Possible sources of files for Power BI include local files and OneDrive.

After David selects OneDrive – Budget 2016, a message asks him to select the file to which he wants to connect to Power BI. There is only one file in this folder, so David clicks it and then clicks Connect, as shown in Figure 2-18. The file selected is highlighted in a different color so that it is recognizable when there are multiple files available.

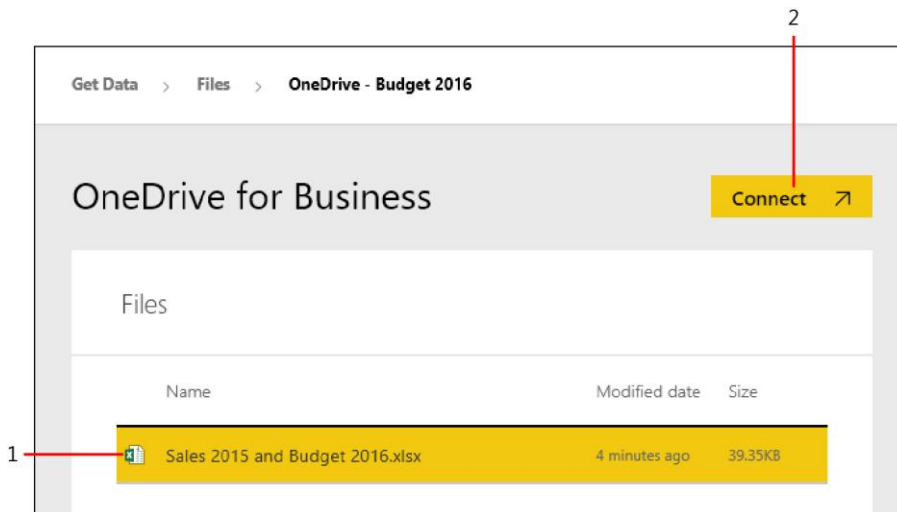


Figure 2-18: The list of files available in OneDrive for Business.

After choosing the Excel file on OneDrive for Business, David must then decide how he wants to use that file. There are two options (see Figure 2-19): Import Excel Data Into Power BI, and Connect, Manage And View Excel In Power BI. If you want to use the Excel file just as a “raw” data source for your reports, select Import Excel Data Into Power BI. Or, you might prefer to copy an existing Excel file as is, using both the data model (if you have one in Power Pivot) and all of the Excel features, such as PivotTables, PivotCharts, and other visualizations available in Excel. If this is the case, choose the Connect, Manage And View Excel In Power BI option. You will see practical examples of this choice later in this book.

David selects Import Excel Data Into Power BI by clicking the Import button in that section.

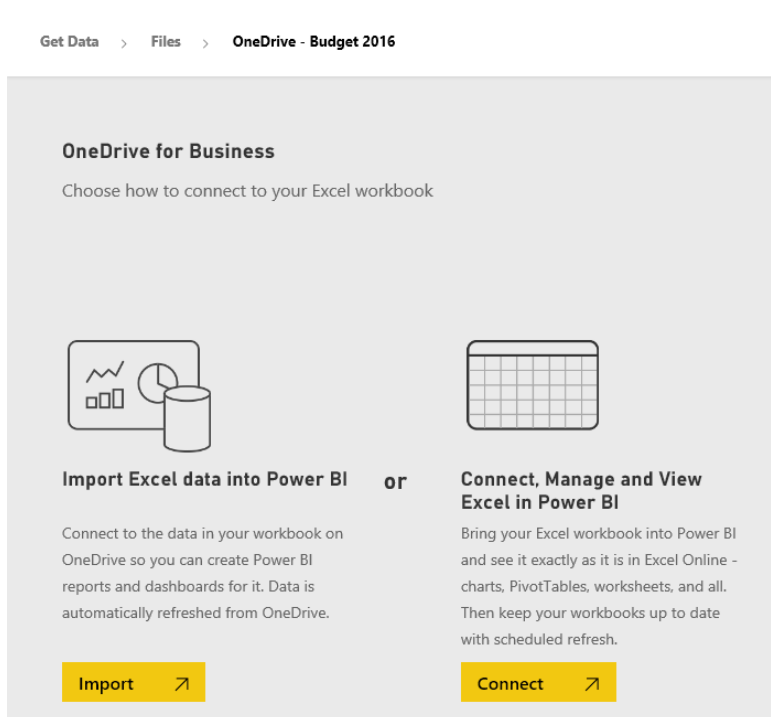


Figure 2-19: Click the Import button to bring in the content of an Excel file stored in OneDrive for Business.

David now has a dataset available, named Sales 2015 And Budget 2016. The dataset contains two tables, Sales and Budget2016, because the Excel file imported has two worksheets with one table each. The Budget2016 table (refer to Figure 2-12) is the one that needs to be modified by other managers, inserting updated numbers for their individual budgets. The Sales table is still the same as it was in Chapter 1 (refer to Figure 1-29), in which David has updated only the Budget column using a formula that searches the corresponding value in the Budget2016 table and allocates it by month. Thus, when a manager updates a row in the Budget2016 table, Excel automatically updates the Sales table.

Using the new dataset, David creates a report by dragging available fields to the report's central pane. The goal of this report is to show a quick recap of the overall budget divided by brand and country/region. For this reason, David chooses a matrix visualization in which he inserts fields from the Sales table (Brand on rows, CountryRegion on columns, and Budget on values, as shown in Figure 2-20). While doing this, he realizes that the two tables have identical fields, and this could be confusing. Because the budget is allocated in the Sales table, it would be nice to hide the Budget2016 table from the Fields list.

When you import an Excel file, all of the tables become part of the Power BI data model and are visible. Later on in the book, you will see how to control the Power BI data model in more detail. For the moment, David just wants to create a first report, and he knows that using the fields from the Sales table is the right choice because he can also create a clustered column chart, below the matrix, that compares the budget with the sales of previous years.

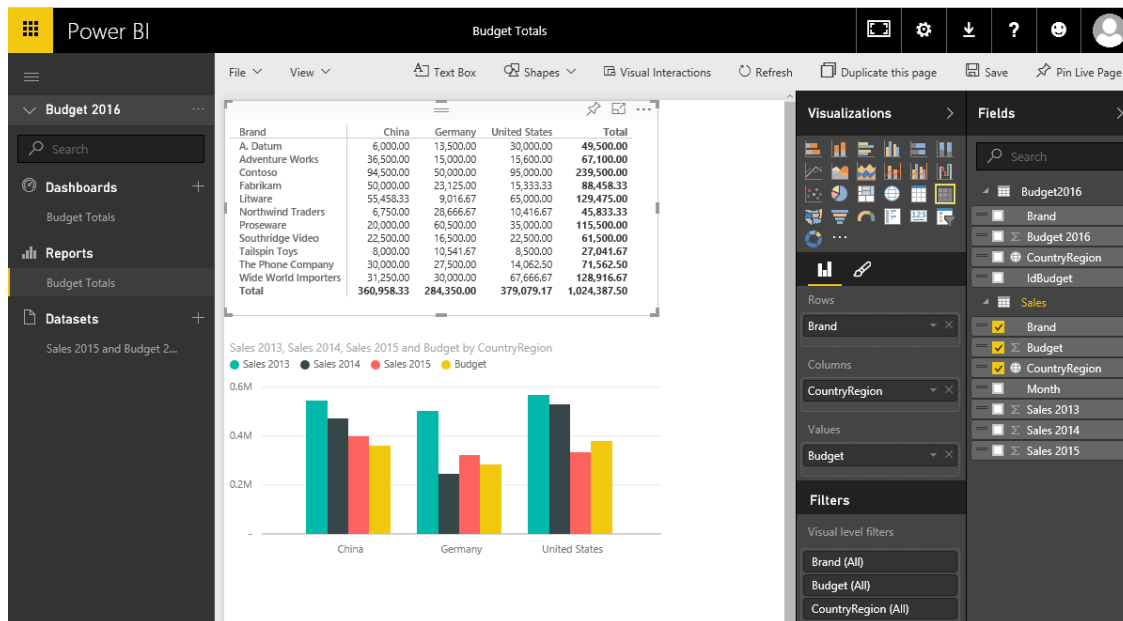


Figure 2-20: The Budget Totals report has both matrix and clustered column chart visualizations.

David saves the new report as Budget Totals and he also pins both visuals to a new dashboard with the same name. David already included Wendy in the group, so he sends her an email asking for a review and edit of budget numbers for Germany (Wendy is the country/region manager for Germany).

Wendy receives David's email, signs in to Power BI, and sees the dashboard. She chooses the Budget 2016 workspace, which she can see because David added her as a member of the group. She opens OneDrive and clicks the Sales 2015 And Budget 2016 file name. Now, she can see the content of the workbook in her browser, and she chooses the Budget worksheet. Next, on the menu bar, Wendy

clicks Edit Workbook and then clicks Edit In Excel Online to edit the file in her browser, as illustrated in Figure 2-21.




Figure 2-21: Wendy can edit the Excel workbook stored in OneDrive for Business in her browser.

Note It is worth remembering that Excel Online is available on many platforms. This means that Wendy can edit the workbook directly on her iPad.

Wendy corrects the budget for Germany for the A. Datum and Contoso brands (see Figure 2-22), because she knows that different marketing conditions and product lifecycles will affect the previous estimation.

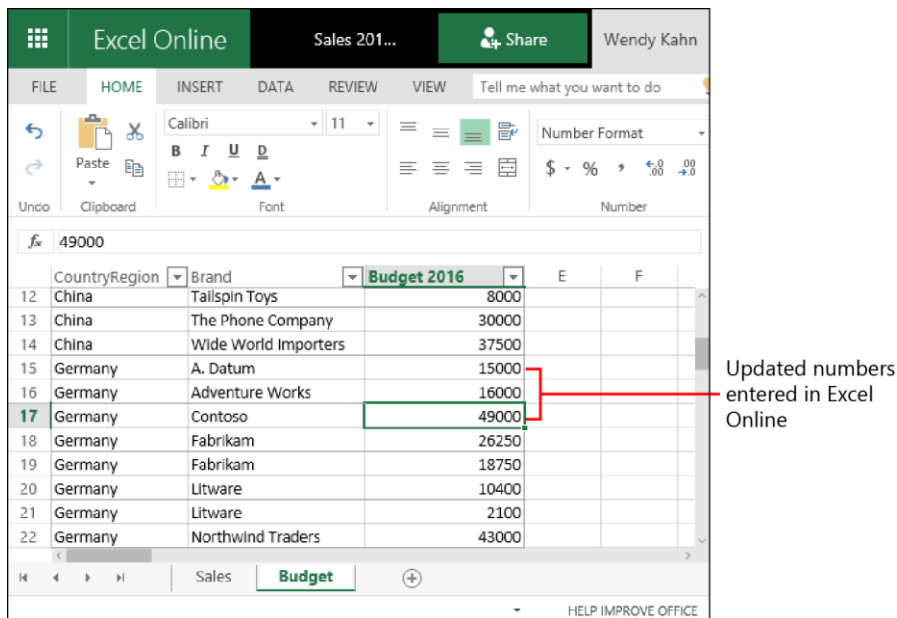


Figure 2-22: The numbers entered in Excel Online are immediately stored, updating the underlying workbook in OneDrive.

Wendy wants to see whether the new values have been correctly allocated, so she clicks the Sales worksheet and sees that the new value (15,000) for the A. Datum brand in Germany is divided into 12 identical parts (1,250 each) for each month, as shown in Figure 2-23. She makes a mental note that such an allocation does not correctly represent the seasonality of certain products, so she will discuss this issue with David during the next meeting.

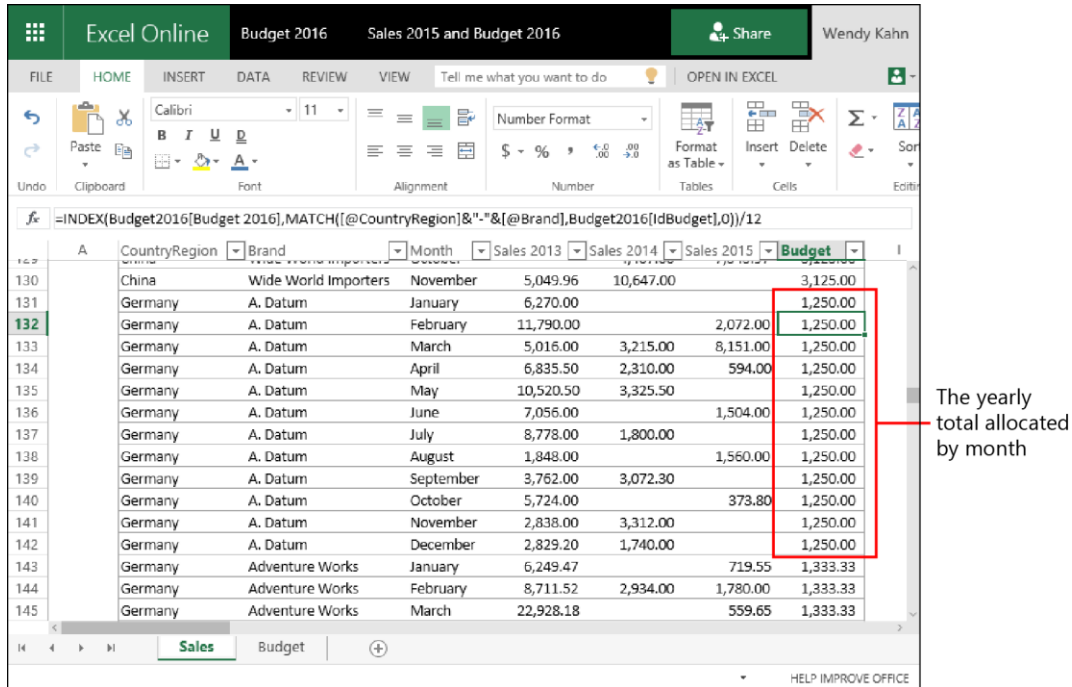



Figure 2-23: The Sales table allocates the budget by month in the Budget column.

Now, Wendy goes back to the Budget Totals dashboard in Power BI and can see the updated data (see Figure 2-24) correctly represented in the matrix. She also can see that all of the other totals are updated, too. Wendy is pleased and feels confident that the reports built by David will be accurate.

Note The Power BI report and dashboard might exhibit some latency when displaying data that was updated in files residing on OneDrive. For performance reasons, the reports are not refreshed every second; thus, if you try to do the same operation, you might not see updated numbers immediately when you go back to Power BI from Excel Online. However, if you wait a few minutes, you will likely see updated numbers. You also might need to refresh the page in the browser to see those updated numbers.



Updated totals in the source Excel file

Figure 2-24: The dashboard in Power BI automatically updates numbers when someone modifies the underlying Excel workbook stored in OneDrive.

You have seen how David created a collaborative environment in Power BI by using groups and OneDrive for Business. Such a collaboration requires licenses for Power BI Pro and OneDrive for Business (remember that an Office 365 subscription usually includes OneDrive for Business, too).

Using groups, you can share files in OneDrive for Business only with group members. If you want to share a file with someone outside the group, you must use the personal workspace in Power BI and the personal folder in OneDrive for Business. Remember: only group members can see everything you store in a group.

If you do not have OneDrive for Business available, you can still use the personal edition of OneDrive (which is available at no charge), but the scenario described in this chapter will be slightly different. You can share an Excel workbook in OneDrive, and other people can edit its content if you share the workbook with them. However, you can import such a workbook only to the personal workspace in Power BI, not to a group workspace. Thus, it is possible to do the following:

- Create an Excel workbook on OneDrive.
- Share the Excel file with other users, even those outside your organization (there are several limitations when doing so with OneDrive for Business).
- Import the Excel file stored on OneDrive in a Power BI dataset.
- Use the dataset in your reports and dashboards. However, you can share them only within your organization, not external users.

You can use the personal OneDrive to create a very similar scenario, but you will not gain the benefit of automatic visibility of files to group members as you do by using OneDrive for Business.

Viewing reports and dashboards on mobile devices

All users who can access Power BI can see all of the reports and dashboards, even on native applications for mobile devices. There are native Power BI apps for Windows (available at the Windows store), iOS (available at the Apple store), and for Android (available at Google Play). These

native apps also have additional features, such as annotations. They are updated with new features regularly.

Getting back to David, he might, for example, want to show the dashboard on his Windows 10 tablet device to a colleague during a meeting. Because the meeting room does not have good Internet connectivity, he takes advantage of the offline availability of dashboards by using the mobile app. This feature also makes it possible to view reports offline, even if you have no interactive capabilities. Thus, David can see the data available with the last refresh of the report, and this is enough for the meeting he planned. Figure 2-25 shows that the rendering of the Budget Totals report is pretty similar to the one available in the browser.

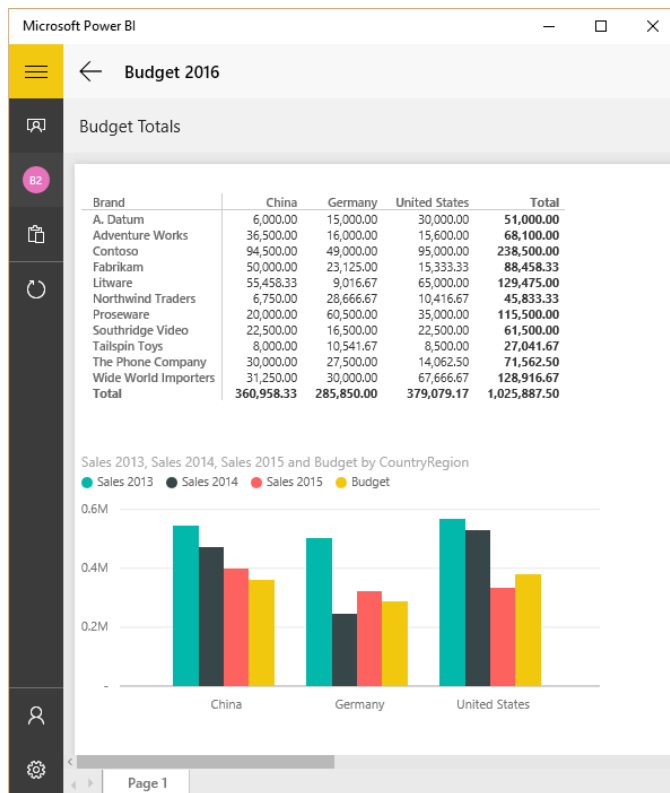


Figure 2-25: Rendering of a report in the Power BI app for Windows 10.

You might also experience different visual presentations in the Power BI app on tablets and smartphones. For example, Figure 2-26 demonstrates how an iPad displays the dashboard created by David in Chapter 1. You can zoom in to each visualization and navigate with the benefit of larger graphics, making it easier to read numbers, as illustrated in Figure 2-27.

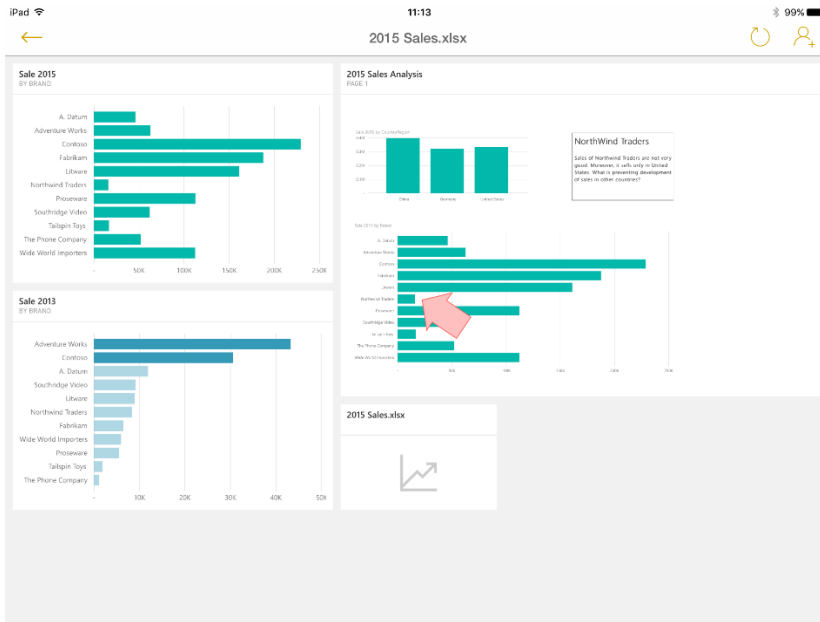


Figure 2-26: Rendering of a dashboard in the Power BI app for iPad.




Figure 2-27: An enlarged view of a dashboard's visualization in the Power BI app for iPad.

That same dashboard is rendered differently on the smaller screen of an Android smartphone, as depicted in Figure 2-28. Visualizations are organized in a vertical column, and you can zoom in to each one, as shown in Figure 2-29.



Figure 2-28: Rendering of a dashboard in the Power BI app for an Android smartphone.




Figure 2-29: An enlarged dashboard visualization in the Power BI app for an Android smartphone.

Users who need to enter budget data in Excel files can also choose to use the Excel app on mobile devices. This app provides a user interface that is optimized for the device, which offers a better user experience than the generic web user interface of Excel Online. Figure 2-30 and Figure 2-31 show how the two worksheets of the Sales 2015 And Budget 2016 workbook display on an Android smartphone.

	B	C	D	E	F	G	H
2							
3	CountryRegion	Brand	Month	Sales 2013	Sales 2014	Sales 2015	Budget
4	China	A. Datum	January	3,234.00		1,935.00	500.00
5	China	A. Datum	February	6,270.00	7,059.00		500.00
6	China	A. Datum	March	4,352.00			500.00
7	China	A. Datum	April	3,814.00			500.00
8	China	A. Datum	May	6,234.00			500.00
9	China	A. Datum	June	5,571.00	3,216.00		500.00
10	China	A. Datum	July	7,424.00			500.00
11	China	A. Datum	August			800.00	500.00

Figure 2-30: The Excel app on an Android smartphone, displaying the Sales worksheet.

	CountryRegion	Brand	Budget 2016	E	F	G	H
4	China	Contoso	94500				
5	China	Fabrikam	60000				
6	China	Fabrikam	30000				
7	China	Litware	60000				
8	China	Litware	55000				
9	China	Northwind Traders	9000				
10	China	Proseware	20000				
11	China	Southridge Video	22500				
12	China	Tailspin Toys	8000				
13	China	The Phone Company	30000				

Figure 2-31: The Excel app on an Android smartphone, displaying the Budget worksheet. Users can enter data here.

Different operating systems and devices might provide a slightly different user interface. The goal here is just to inform you of the options available for tools that complement the Power BI service and app. We suggest that you try these applications for yourself, with your own data, so that you can evaluate which of your data for Power BI can be displayed/edited using a mobile device.

Conclusions

In this chapter, you saw how to share dashboards, reports, and raw data with people within an organization.

- You can invite a user within your organization to see a dashboard. That person can also see the reports underlying the dashboard.
- To share a dashboard with users outside of your organization, you must provide them an account within the same tenant.

- You can publish a report on a webpage; however, anyone can view that report. You should consider this option only for data intended for public consumption; you should not use it for sensitive, company-private data.
- You can create groups of users within your organization and share datasets, reports, and dashboards.
- You can use the same groups with OneDrive for Business to make it possible for other users to directly enter data in shared reports, which are then automatically updated after each change.
- You can display dashboards and reports on mobile devices by using native apps. These apps typically provide a better user experience if you are using smaller displays and have limited offline access to dashboards.

Hear about it first.



Get the latest news from Microsoft Press sent to your inbox.

- New and upcoming books
- Special
- Free eBooks
- How-to articles

Sign up today at MicrosoftPressStore.com/Newsletters

Understanding data refresh

Let's continue following David, the manager of budgeting at Contoso. While he explores Microsoft Power BI, he now has begun sharing his dashboards and reports with Contoso's country/region managers around the globe. Everybody likes the idea of being able to view a report on any device and share considerations while looking at the same figures. Nevertheless, the managers are concerned that they are making decisions based on sales as of October 2015, and it is now mid-December. The figures are no longer the best data upon which to forecast sales, a fact that is even more pertinent for the products that show a clear seasonality.

Thus, David needs to retrieve the latest sales data set and refresh the Microsoft Excel model. Urged on by the country/region managers, he ends up doing this each and every morning until he begins to wonder whether this process can be automated in some way. In fact, Power BI offers an option with which he can schedule refreshes.

Introducing data refresh

In Chapter 1 and Chapter 2, you learned the basics of Power BI: how to upload a workbook containing some data, build reports and dashboards, and how to share the content with other users in your organization or outside of it. Those chapters also touched upon the basics of data refresh: uploading a new version of a workbook containing data, and using Microsoft OneDrive for Business to automate the uploading process.

For both of these scenarios, updating data meant that David needed to refresh the content of the Excel file manually and then upload it to Power BI, either via another web service, such as OneDrive for Business, or by using the Power BI user interface (UI). When it comes to refreshing your data, Power BI offers much more control. Learning it requires some attention to details. We suggest that you read this chapter in its entirety because there are some small but important details that you need to know before making any decision about your future data refresh strategy.

More important, in the first two chapters, we kept the presentation deliberately simple, trying to show only the basic concepts. In this chapter, however, we begin to delve into the details of working with Power BI. Hereinafter, the details become important.

Anyway, first, we need to focus on what “refresh” really means. In the context of this chapter, refreshing data does not mean to manually update the Excel workbook and save it as another version of the same file. Instead, we want the workbook to automatically update its content by using a connection to the source database from which we originally query the data.

Let’s review the steps in David’s simple data-processing system:

1. Data is retrieved from the database by IT. IT then gives David an Excel file containing the latest figures of sales.
2. David manually copies the information to his own version of the Sales 2015 workbook.
3. He saves it so that OneDrive uploads the content to the cloud.
4. Power BI loads the content of the file from OneDrive and updates its own internal data model.

As part of step 3, the Excel file automatically computes the forecasts generated by the country/region managers by using formulas. These results are saved in OneDrive by the managers, so the results are immediately available to Power BI, too.

The data refresh mechanisms David learned so far are useful to automate step 4. But now, he wants to automate steps 1 and 2. This requires some more understanding of how Power BI works internally.

Introducing the Power BI refresh architecture

What happens when you upload a workbook to Power BI? Let’s consider the workbook that David uploaded to look more closely at the process. Recall that his workbook contained a table. Figure 3-1 shows the flow of data.

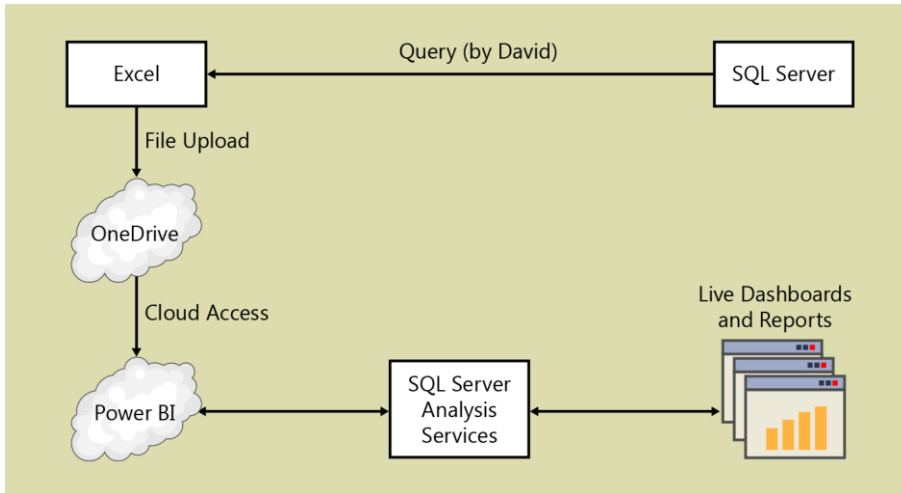


Figure 3-1: The data flow, from the original database up to Power BI.

David's data moves from SQL Server into Excel (David does that). Next, the Excel file is uploaded to OneDrive where Power BI reads it. After reading the file, Power BI generates a SQL Server Analysis Services (SSAS) database that ultimately computes, through the Power BI UI, the dashboards and reports. Sounds complicated, right? In fact, it is, but luckily, Power BI hides all of this complexity, making it easy for you to generate reports from Excel files. Nevertheless, to understand how data refresh works, you need to have a clear picture of the complete flow of information.

To use data refresh, you need a way to pull data from the data source (SQL Server, in this example) and push it directly into the SSAS model generated by Power BI. In other words, you want to create a data flow that circumvents Excel and OneDrive (both operations are done outside of Power BI) to make it flow as shown in Figure 3-2, in which the steps that we want to remove appear in a blue box.




Figure 3-2: We want to remove the steps enclosed in the blue box to make data refresh work more efficiently.

Here is what you need to do to make this happen:

- The data set cannot be a plain Excel table, because Power BI needs to know how to query the source database (SQL Server, in this example) in order to refresh the data. Obviously, it cannot rely on asking you or asking IT. The method must be formalized and use a language that Power BI can understand.

- The SSAS engine running in Power BI needs a way to access the source database. Such a database is usually located within your company (or on your laptop). Thus, you will need software that implements the connection with Power BI.

Be sure that you understand these architectural requirements well before moving on with the rest of the chapter. As you will see by reading the next sections, we stripped away most of the technical complexities that comprise data refresh, but you need to keep in mind that the aim is to create the scenario depicted in Figure 3-2.

Introducing Power BI Desktop

You might remember from Chapter 1 that Power BI offers you two ways of interacting with it: direct access to the web service, or by using Power BI Desktop. In this section, we are going to show you how to use Power BI Desktop. Power BI Desktop is an application that runs locally on your computer but offers you all of the features available on the web, plus many more options for building a data model.

By using Power BI Desktop, you do not rely on the web service to create the data model for you. Instead, you have the full modeling capabilities of Power BI available at your disposal and control.

Why should you worry about Power BI Desktop at all if the web service is capable of building a model for you? There are several reasons for this, but for now, let's concentrate on one of those reasons: you can describe the details of your dataset, which accomplishes the first requirement of our data-refresh scenarios from the preceding section. But, before we can get to that, first you need to download Power BI Desktop from the Power BI website and install it.

On the Power BI website, on the right side of the menu bar at the top, click the download button (see step 1 in Figure 3-3), and then click Power BI Desktop (see step 2).




Figure 3-3: Downloads in Power BI are available in the download list, on the right side of the menu bar.

After you have downloaded Power BI Desktop, install it by following the instructions provided in the Microsoft Power BI Desktop Setup Wizard, and then start the application. A welcome screen greets you, and then you see the main Power BI Desktop window, as shown in Figure 3-4.

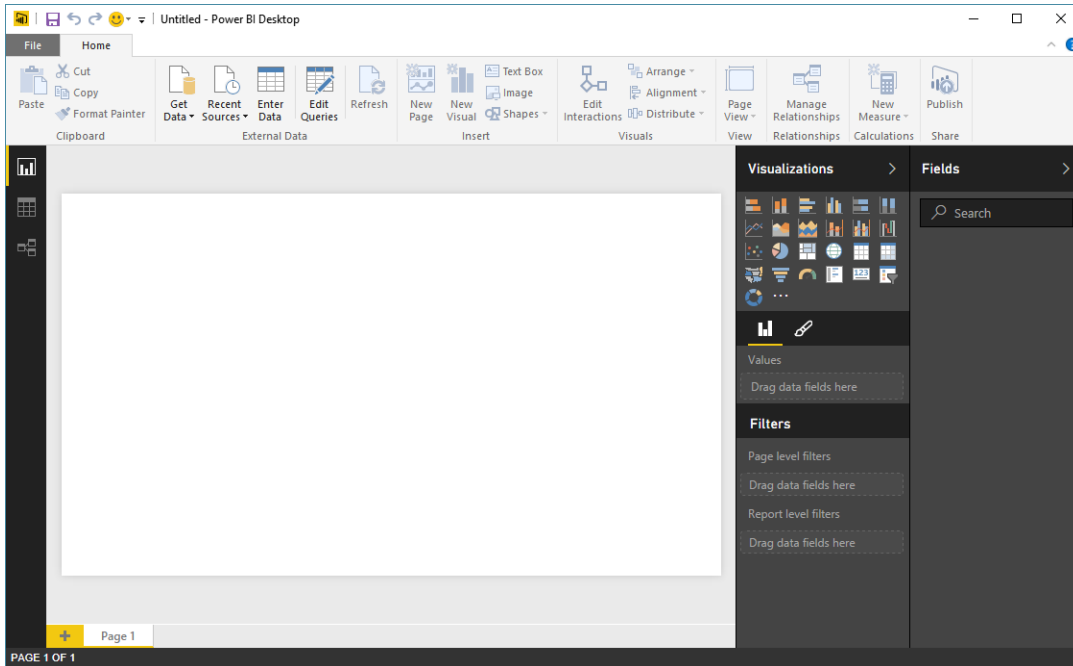


Figure 3-4: The Power BI Desktop UI resembles the Power BI website.

The first thing you will probably notice is that the UI of Power BI Desktop is very similar to that of Power BI. Nevertheless, as you will learn, there is more to Power BI Desktop, and it is a bit more complex to use than its web-based counterpart, but it exploits the full power of the Power BI engine.

As with Power BI, the first step is to provide some data to Power BI Desktop. Because the original data is in Excel, you can begin practicing with Power BI Desktop by using the same Excel file you used earlier for the website. On the ribbon, in the Data group, click Get Data, and then click Excel (see Figure 3-5).

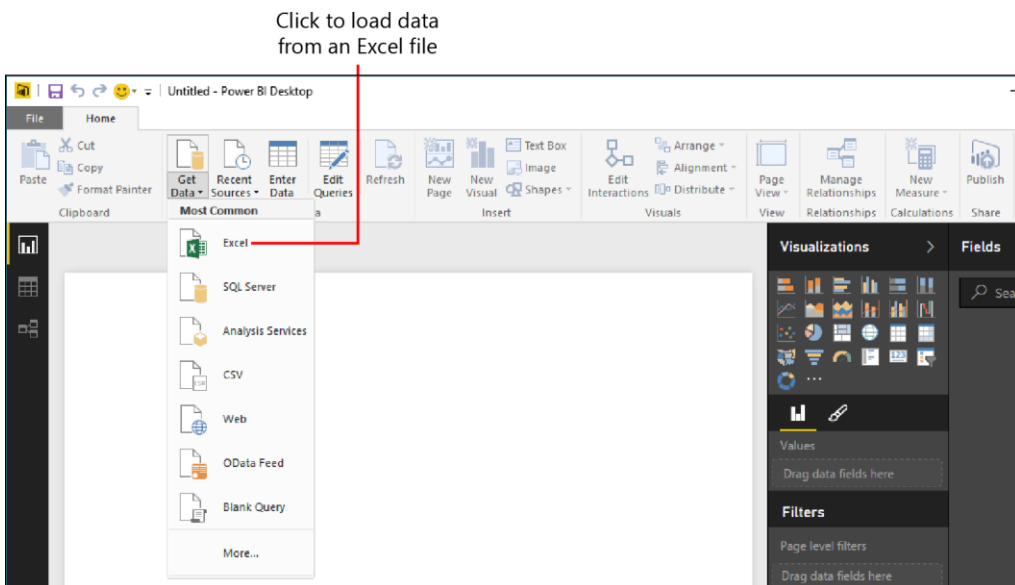


Figure 3-5: Your first step when using Power BI Desktop is to load some data; in this example it's from Excel.

In the Open dialog box, select an Excel file and then click Open. The Navigator dialog box opens, in which you select the source file. When you load data from Excel, you can choose to load from tables

or from worksheets. Figure 3-6 shows the two sample items available. Note that for this exercise we renamed the worksheet “Sales” to “Sales Worksheet” to make the figure clearer. In your models, it is likely they will have the same name, and only the icons adjacent to the names will differentiate tables from worksheets. For this example, let’s work with the table.




Figure 3-6: The Power BI Navigator helps you to choose the source for importing data in Power BI Desktop.

Select **Sales**, click **Load** to import the table into Power BI Desktop, and then close the Navigator dialog box. At this point, you will likely feel at home and in familiar surroundings. In fact, by using the Fields and Visualizations panes, you can build a report in Power BI Desktop in the very same way you built the report earlier, on the website.

For example, Figure 3-7 shows that you can build a report similar to the one you built on the website. The main difference is that, now, you are doing the work directly on your PC instead of interacting with a cloud service.

Note We suggest that you become familiar with the UI of Power BI Desktop by experimenting with some reports. In this chapter, we are not showing you a step-by-step guide to Power BI Desktop. Instead, we focus on the new features available in Power BI Desktop that are not in the cloud service.

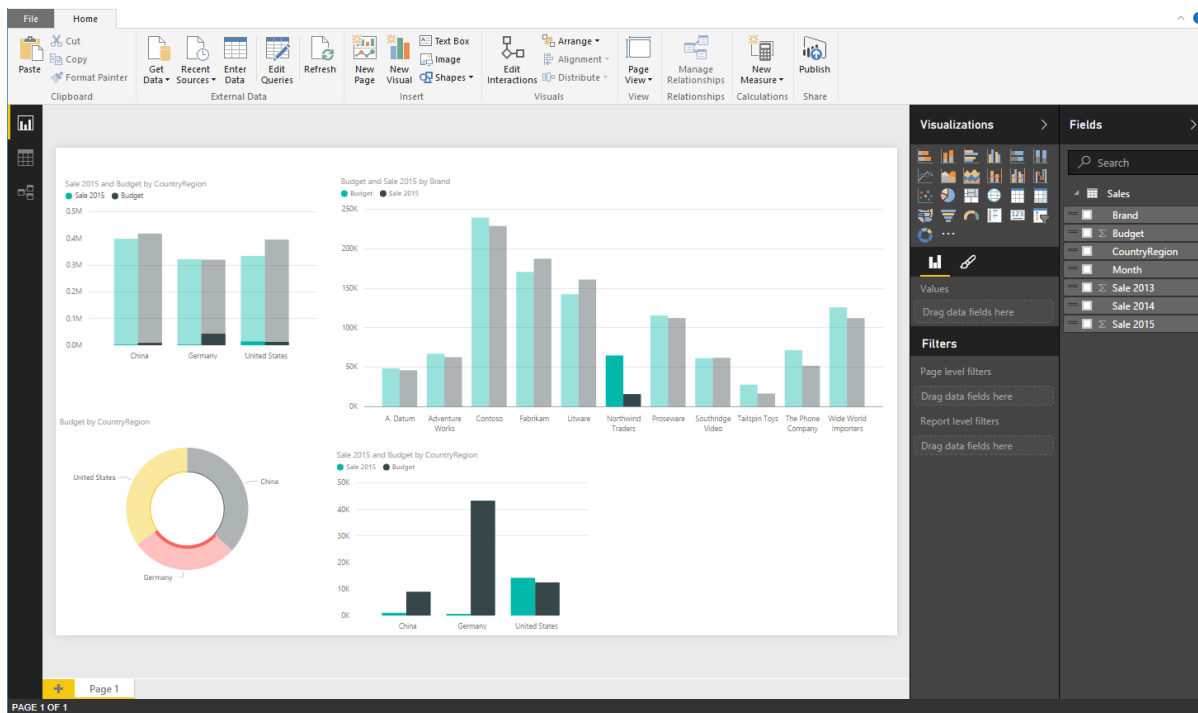


Figure 3-7: This is a sample report that you can build by using Power BI Desktop.

Building the report, you can appreciate how user-friendly the Power BI Desktop environment is, with features such as copy-and-paste available. So, for example, if you need a visualization similar to one you have already created, copy it, paste it, et voilà: the job is done. Of course, there is more to it than that, but Power BI Desktop offers little things like this that make life much better.

Publishing to Power BI

Let's get back to our exercise. When the model is ready, save it on your computer using the name **Sales PBD**. Of course, at this point, the report is local to your PC, and no one else can view it. However, you ultimately want to publish this model on the Power BI cloud service, to take advantage of all the features of Power BI, including sharing and viewing on a mobile device.

To do this, go to the Power BI Desktop ribbon, and then, on the Home tab, in the Share group, click Publish. Power BI Desktop then asks you to sign in to the Power BI service (and might ask if you want to modify your changes). As previously mentioned, Power BI Desktop is an application that runs locally on your PC, and it can work without a Power BI account, but as soon as you want to publish your data, you need to have an account (or create one) and sign in. After you sign in, Power BI Desktop shows the message depicted in Figure 3-8, confirming the operation and providing a link to the published report.

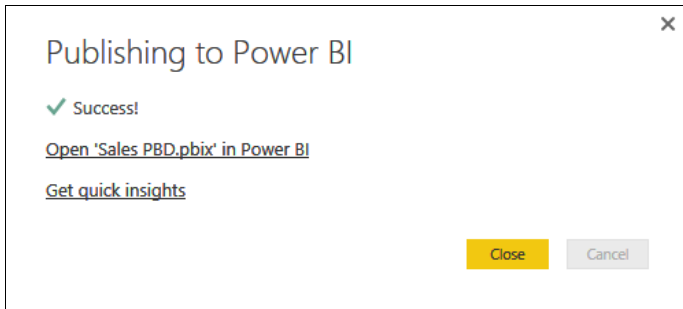


Figure 3-8: When the model is published, you can immediately see it using the appropriate link.

If you open the file in the Power BI website, you will find a dataset named Sales PBD and a report with the same name. By using Power BI Desktop, you created both a model and a report, and, when publishing it to Power BI, it created both objects.

Note The model is copied from your local file to Power BI. When the model is in Power BI, you can further enhance it, but the two versions are disconnected. Changes that you apply to the published model from within your browser are not be applied to your local version on your PC, and any subsequent publish operation that you initiate from Power BI Desktop will overwrite the changes that you made via the web browser. Thus, when you begin building models with Power BI Desktop, it is a good practice to continue updating them locally and then republish them. Do not modify the online version. You might want to use different version names for your local files in order to publish different, varied reports in the Power BI cloud service.

After you get used to this method of developing reports, you will find it extremely convenient. In fact, working with Power BI Desktop is more productive because you do not need an Internet connection and you have the full power of a Windows application. When the model is ready, you publish it, overwriting any previous version that you might already have done.

Let's recap what we have seen so far:

- Power BI Desktop is a Windows application that offers the same features of the cloud service, but it runs on your local PC.
- You can build a model with Power BI Desktop and save it to your PC.
- You can publish a Power BI Desktop model to Power BI, but you'll need to have or create an account and sign in first.

You might remember that this chapter is about data refresh. Why is Power BI Desktop relevant to data refresh? A Power BI Desktop file contains all the information needed to refresh the model. In fact, in the Power BI Desktop file, we created a link between the original Excel file containing figures for the budget and the Power BI Desktop model.

In Chapter 1, when we uploaded the file to Power BI, we simply copied it. However, using Power BI Desktop, we create a link between the Excel file and the Power BI Desktop file by writing a query. In reality, we do not actually author any query, but as you will learn, Power BI Desktop created the query for us, making the task transparent.

All that is missing at this point is to provide a way for Power BI to access the Excel file that is the ultimate source of our data. In fact, the original Excel file is stored on the local PC, and the Power BI cloud service cannot access it. Solving this problem is the topic of the next section.

Installing the Power BI Personal Gateway

The Power BI Personal Gateway is another piece of software that can connect with the Power BI cloud service and carry out the queries stored in the Power BI Desktop file. You can download it from the same webpage from which you downloaded the Power BI Desktop application, but this time select Power BI Gateways after you click the download button. Power BI then asks you to choose between the two Power BI gateways:

- **Personal Gateway** This version is intended for use with personal datasets. It is simple to use and install but offers limited features regarding monitoring and security for multiple users.
- **Enterprise Gateway** This version offers more functionality but, at the same time, involves more complexity in its setup and usage and usually requires involving your IT department.

Note In the last chapter of this book, we briefly outline the differences between the personal and the enterprise versions of the gateway. For the purposes of this discussion on how data refresh works, those differences are negligible.

Let's go back and visit David to see how he is progressing.

David is still experimenting with Power BI, so he has no intention of installing a complex system. For this reason, he chooses the Personal Gateway. Before jumping into the setup of the Personal Gateway, though, he needs to understand how it will be implemented.

If David runs the setup with administrative privileges (that is, he installs it as an administrator), the gateway will run as a service. On the other hand, if he installs it as a standard user, the gateway runs as a normal program. What is the difference? When the gateway runs as a service, it runs even when no user (or another user) is signed in to his PC. Conversely, if it runs as a standard program, the gateway will run only if David is signed in to his PC. This might be relevant if at some point he wants to refresh his data while at home, accessing the Power BI cloud service, but his laptop is still in the office, turned on but without anybody using it. In such a scenario, if the gateway runs as a service, the refresh operation will succeed, whereas if it runs as a normal program, it will fail.

The choice of which to run is up to you. Figure 3-9 shows David choosing to run the installer as an administrator.




Figure 3-9: Right-click the downloaded installer to install the gateway as an administrator.

After you install the gateway, you must start it to complete the configuration. In fact, when you start it, it requires the credentials to access Power BI. The gateway requires them because it must contact the Power BI cloud service and begin answering queries coming from the service.

After David provides the correct sign-in, the Power BI Gateway – Personal dialog box opens, in which he must provide another set of credentials, as illustrated in Figure 3-10.




Figure 3-10: As part of the configuration, the gateway asks for the credentials to use when running as a service.

Why does the gateway require these credentials? Because David wants to run it as a service. Being a service, it will run even if no one is connected to the PC, and it requires credentials to access files and connections from David's PC. In other words, the gateway will have the same permissions that David has on his PC, and those are required to let it access his files and data sets as if it were him.

When everything is complete, the gateway informs you that one operation still remains: you need to go to powerbi.com to complete the setup of the data sources. When you go there, you will reenter your credentials.

Note At first sight, it might seem cumbersome to be required to enter the credentials so many times. Indeed, it is important that you follow the procedure in the correct way. The gateway connects the Power BI website to your PC, and it will be able to access all of your files. Security is always important, and Microsoft takes it very seriously.

David is now near the end of the gateway configuration. The last step is to visit the Power BI cloud service. To complete the setup, in Power BI, on the menu bar, he clicks the configuration button (the small gear icon; step 1 in Figure 3-11) and then chooses Settings (step 2). This opens the Settings page on which he can configure the settings for each data set, as demonstrated in Figure 3-11.




Figure 3-11: The final step in the configuration of the gateway is to grant permission for a Power BI dataset to access it.

On the Settings page, David sees an alert with two important pieces of information:

- The gateway is online and running (in this case) on a PC named HARRY (which you can see in Figure 3-11, in the line below the Gateway Status section header).
- This data set is not yet ready. In fact, even if the gateway is set up, you need to specify the credentials specific for the single data source.

Why do you need to enter credentials again for each data source? The reason is that every data source might require different user credentials. In our example, the Sales PBD file is stored on David's local PC, so it is automatically accessible by using the Windows credentials that David stored in the Personal Gateway, too. In this case, when David clicks Edit Credentials, he is asked to choose an authentication method, as illustrated in Figure 3-12. The Windows authentication method is the only option, and when David clicks Sign In, he will not need to provide his user name and password again. However, for other data sources not using Windows authentication, clicking Sign In will require a user name and password to make it possible for the Personal Gateway to connect to that data source during a refresh operation.




Figure 3-12: You have a choice of authentication methods for each data source.

Configuring automatic refresh

After you set the credentials, the data source is ready to be refreshed. Power BI now has all of the information it needs to refresh the data, both on demand and on a scheduled basis. Expanding the Schedule Refresh section, you can define when Power BI attempts to refresh the dataset. Figure 3-13 shows an example of a data refresh scheduled twice daily, at 9:00 AM and 4:00 PM.




Figure 3-13: You can configure automatic data refresh to run daily or weekly and at different times.

Note You need a license for Power BI Pro to schedule more than one refresh per day. Using the free Power BI license, you can schedule only one daily refresh.

In case the refresh fails, you have the option to receive an email alert so that you can take the needed remedial actions.

At this point, you can either wait for the refresh to happen or, if you want to be sure that everything is set up correctly, you can force an immediate refresh. To perform this operation, in the navigation pane on the left, in the Datasets section, click the ellipsis to the right of the data source (Sales PBD, in our example; step 1 in Figure 3-14), and then click Refresh Now, shown as step 2 in Figure 3-14.

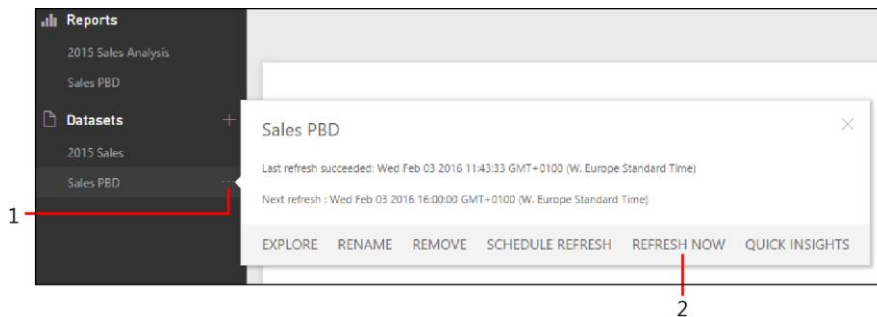


Figure 3-14: You can refresh a dataset immediately by clicking Refresh Now.

When you ask for a refresh, Power BI prepares for the data refresh and then starts it. Depending on the size of the dataset and the speed of the Internet connection, the time to carry out the refresh can range from a few seconds, as in the case of David, to a much longer duration. You can see when a dataset was last refreshed by looking in the same window where you asked for the immediate refresh (click the ellipsis to the right of the data set).

Note Although it might be obvious, it is useful to state an important fact: the data refresh happened on the model in the cloud, not in the model on David's PC. In fact, if David opens Sales PBD by using Power BI Desktop, he will see the data as he saved it. Power BI does not change any file on your PC, it only uses the gateway as a medium to access the data sets upon which the model is based.

Conclusions

In this chapter, you learned the basics of data refresh. Let's recap them briefly:

- You can upload simple data models, based on an Excel file, to Power BI and refresh them by using the Personal Gateway, which makes it possible for Power BI to access your local datasets.
- For data refresh to work on more complex models—for example those that load data from SQL Server—you need to use Power BI Desktop.
- With Power BI Desktop, you build models containing the needed information to let the cloud service connect to the Personal Gateway and retrieve the dataset.
- You can refresh your data daily with the free license, whereas you need a professional license if you need to refresh your model multiple times each day.

Power BI Desktop offers many more features and will let you move to the next level in the learning path of data modeling. This is the topic for Chapter 4.

Using Power BI Desktop

In this chapter, David will move a few steps further in building his Power BI model. But, the solution he has built so far still depends on IT providing him with the sales figures for the last three years. Fortunately, David discovers that Power BI Desktop can load the numbers directly from the Contoso data warehouse via Microsoft SQL Server. Now, when Power BI carries out data refresh, it will automatically get the latest sales data, updating the entire model.

Let's take a brief look at what David needs to do:

- Load sales figures from the data warehouse instead of using his Excel file. This requires accessing the corporate database, but the IT department can give him access to the data he needs.
- Load forecasts for the next year from the Excel file that the country/region managers update every day.

Power BI Desktop offers all of the functionalities that David requires. Now, let's go deeper and learn how to take advantage of this extremely useful application.

The following notes below are for information only as you need to be able to connect to an SQL server. You should be in the Power BI Desktop program and use Get Data to upload the Excel file Sales 2015. You can download the Power BI Desktop directly from Power Bi, lcik the Download button at the top right of the screen.

Connecting to a database

David already knows how to load data into Power BI Desktop from Excel; he always did it using the Excel file that he received from IT. Now he needs to work from a database. To access the database, he will ask the IT department to provide him the credentials to query the database and read the information he needs. Karin, the database administrator at Contoso, grants him read access to a view that returns the same dataset that she has been sending him every day.

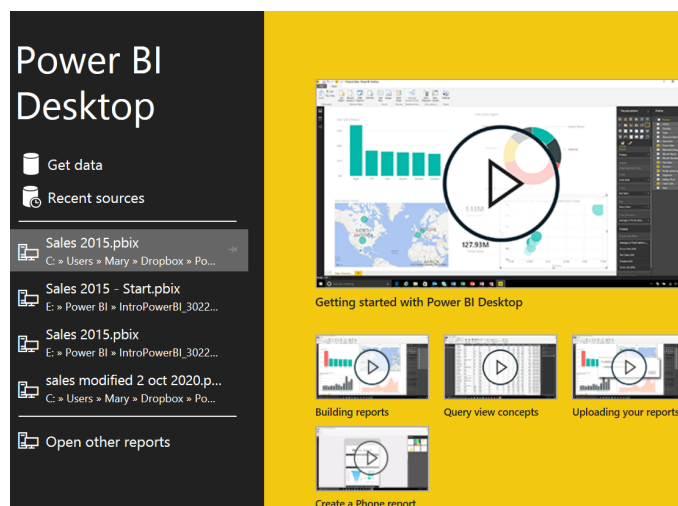
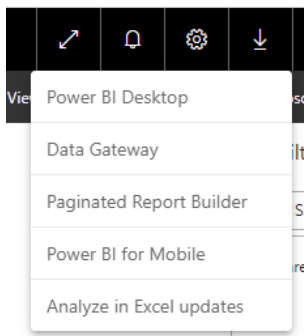
Karin is happy to do so because she will no longer need to fill the Excel workbook. If David can load the data and manage to grab insights by himself, her daily list of chores will become lighter. Thus, Karin tells David, "You can access the view named Sales2015 using your Windows credentials on the server ContosoDbServer." David's account can only read that view, so there is no potential danger. Database security will guarantee that nothing bad can happen to the database.

We will **not be connecting to an SQL Database**, as we do not have access, so we will build a report from an **Excel file, Sales 2015 that has been saved, in Chapter 3, 2015 Sales**. Open Power BI Desktop which may appear similar to below. You can also open Reports, view recent sources and files etc. The difference between Power Bi and Power BI Desktop is Power BI runs online, PowerBI Desktop runs locally. Files are interchangeable

Click **Get data**, select **Excel**, click **Connect**, go to Chapter 3 and open the **2015 Sales Excel file**.

Continue to Page 65

Note: You can download Power BI Desktop from the Power BI program.






Figure 4-2: Use the Navigator to choose from among the tables available on the database.

After David makes his selection and clicks Load, Power BI Desktop asks how he wants to interact with the data. He selects Import and then clicks OK, as illustrated in Figure 4-3.

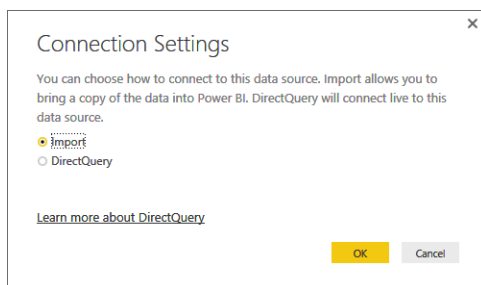


Figure 4-3: Before importing data from an SQL database, you need to choose the loading method.

Let's take a moment to learn about this connection option because it is an important one and will help shed more light on how Power BI connections work.

When you choose Import, Power BI Desktop connects to the database, loads the information, and stores it within its internal data model. You can then work on your data in Power BI Desktop without being connected to the database. You will only need a connection when you want to refresh the data.

With DirectQuery, Power BI Desktop does not load the data into its internal database. Instead, it runs a query to the original database every time it needs to draw a chart or, in general, run a query. Thus, the connection between Power BI Desktop and the database will be permanent.

The contrast in the query timings reflects a key difference: when you use Import, you are working with data that is only as current as the latest refresh, whereas with DirectQuery you always see the latest information available when you create the report.

At first glance, it looks like DirectQuery is the most convenient method for loading data, but this is not totally true. If the data is updated frequently, it is very likely that one minute you will see a report with

a set of figures, but when you open it again a few minutes later, the numbers might no longer be the same. This is frustrating if you are analyzing information over the span of an entire year (which is what David is doing. Numbers that change too frequently can become disturbing. Also, although real-time data might sometimes be useful, it comes at the cost of query speed; DirectQuery by its very nature is much slower than working with data that is resident on your device and directly accessible by Power BI Desktop.

As a final note, keep in mind that DirectQuery works fine when you use Power BI Desktop on your laptop, but when you publish the model to Power BI, the cloud service needs a way to communicate with the internal database server. This is accomplished by using the Enterprise Gateway, which is the advanced version of the Personal Gateway, to which you were introduced in Chapter 3.

Continue...

David wants to analyze data over the span of a year. The **Navigator** dialog box opens, he select **Sales 2015** and a preview of the file appears on the right, he selects **Load** and a blank report page appears. with Fields listing the column names from the Excel file on the right. We will create 4 Visualizations to look like the



Figure 4-4: The dashboard with 4 visualizations

David notices there seems to be little difference between this model and the previous ones that he created using Power BI. The main difference is that in **Power BI Desktop the data can be shaped by using queries using the Power Query Editor window.**

Loading from multiple sources

David experienced an unpleasant surprise: by using Excel, he was able to integrate into the same table both the sales, which came from SQL Server, and the **budget forecasts**, which came from an Excel file. However, the Excel file containing the forecast cannot be gathered from the SQL Server database, because the country/region managers update the Excel file whenever they want to share some new figures for 2016.

To solve this problem, we need to dive a bit more into the internal structure of the Power BI Desktop model. In Chapter 3, we said that a Power BI Desktop model contains an internal query, created by Power BI Desktop, for each dataset.

This internal query is not visible if you perform basic operations, such as loading data from an Excel file or from a SQL Server database. Yet, it is there, and if you need to modify it, you can.

The query language of Power BI Desktop is used by Query Editor, and discussion of that language alone would fill an entire book of several hundred pages. As you might imagine, we cannot cover it in a mere few pages here. Instead, we want to show you some basic features of Power Query Editor so that you better understand its capabilities.

The next step in Power BI is queries, it has its own screen called the Power Query Editor screen

To go to the **Power Query Editor** window, from the **Home** tab, click on **Transform data** button in the **Queries** group (in other versions it may have **Edit Queries** in the **External data** group). The Queries will be listed on the left of the screen.




Figure 4-5: Click **Transform data** or **Edit Queries** to access the Query Editor window.

The **Power Query Editor** window opens, you may get an error message about the location of the Sale 2015 file, browse to Chapter 3 and select the Excel file Sale 2015, the table will appear as below:




Figure 4-6: Power BI Desktop’s Query Editor is a complete development environment in and of itself.

Let’s take a quick tour of the Power Query Editor window. Along the top is the ribbon, which has tabs and we will check out some of these and see the effect on the table. This is a virtual copy of the Excel file and we can make any changes we wish, but they are not changing the original Excel file. On the right there is a list of **APPLIED STEPS** which is like a recording of every change you make and which can be deleted, very like the UNDO command.

Below the ribbon, on the left side, is the **Queries** pane, which displays the current query **Sales 2015**.

The **Query Settings** pane on the right displays the query properties with the **Name** of the Query **Sales 2015** and the history list, **Applied Steps**.

In David’s scenario, his objective now is to create a new query that also retrieves the budget forecast information from the Excel workbook.

To load data from a new dataset, on the ribbon, in the **New Query group** of the **Home** tab, click **New Source**, and then **locate and load the data** from the **Excel** file **Budget** in Folder **Chapter 4**. (You may see 2 files, **select Budget by placing a tick to the left and click ok**) This results in two tables displaying in the **Queries** pane on the left side of the **Query Editor** window and the data shows in the center pane. Notice the Budget 2016 figures, compare them to the budget figures in Sales 2015 which are monthly

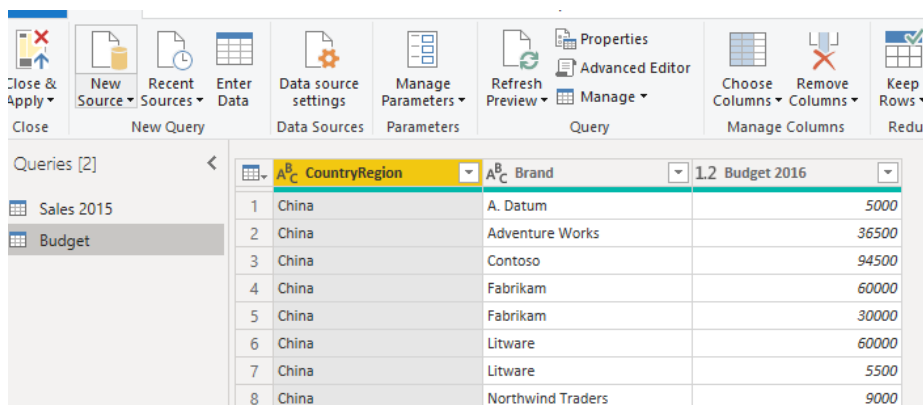


Figure 4-7: With the Query Editor can create multiple datasets, as you can see in this figure.

Now we are ready to do some editing. on the **Home** tab, click **Close & Apply** in the **Close** group to load the data into Power BI Desktop visualizations. When this is done, the **Power BI Desktop Fields** pane shows the two sources as depicted in Figure 4-8.

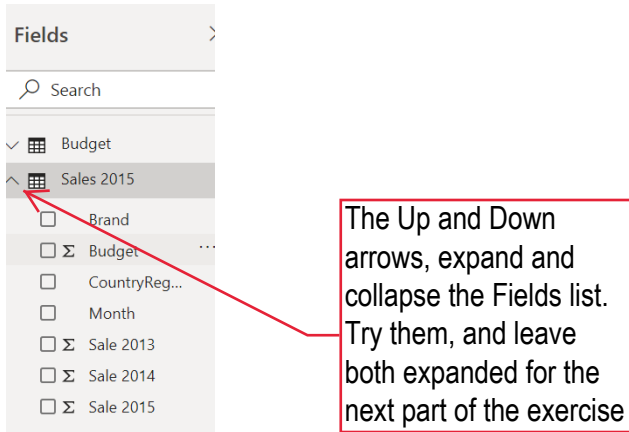
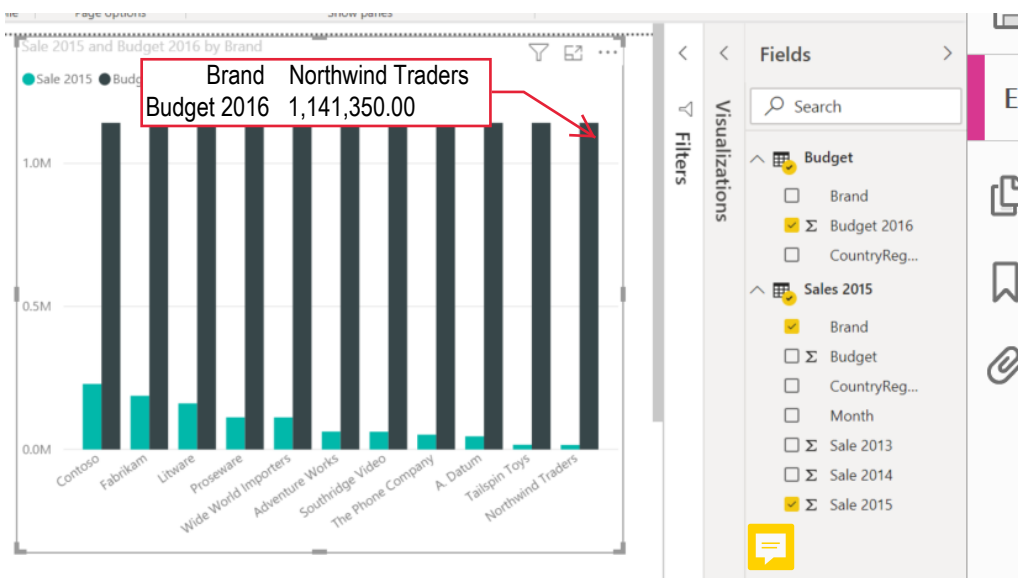


Figure 4-8: The Fields pane lists all of the tables (and columns) that the Power BI Desktop model is using.

Using Query Editor

Now David wants to build a report containing both the **Budget and Sales 2015 tables sliced by the Brand column**. At the bottom left of the page click the **Plus** symbol to add a 2nd page.

From the **Visualizations**, select the **Cluster Column Chart** icon. **Place ticks** next **Sale 2015** and **Brand** from the **Sales 2015 data** then select **Budget 2016** from the **Budget** data. As shown in Figure 4-9, David sees that the value of the budget is the same for all the columns. If you hover over any **Budget** column you will see that each value is **\$1,141,350.00**, which is the **Grand total** of **Budget 2016** in the Excel worksheet.



Note: The sorting is by Sales 2015, Descending order. You can enlarge the screen view if your tile is too small, select the **View** tab then select **Page view** and select **Actual size**.

Figure 4-9: Note In this chart, the value for Budget 2016 is the same for all the columns, and it is too high being the Grand Total of the Budget. **We have 2 tables but they are not yet joined.**

The chart is confusing, at the very least. David used the **Brand** and **Sales2015** from the **Sales2015** table, and the **Budget 2016** column from the Budget table.

The problem here is that if you use the **Brand** column from the **Sales2015** table, despite having the same name, it is not the same as using the **Brand** column from the **Budget** table. The two columns have the same values and the same name, but they are not the same column. In fact, if you were to try replacing Brand in the chart with the Brand column from the **Budget** table, the result would be similar, but with the opposite behavior, i.e. the Sales are now all the same 1,057,227.05 which is the Grand Total of all the Sales for 2015




Figure 4-10: Here the **Budget values are correct** but the **Sales are not**, hover over any sale bar and view the data, this time each Sale 2015 shows a value of **1,507,227.05**

Figure 4-10 shows that using the Brand column from Budget, the **budget values are now correct, but the sales are not.**

The main problem here is that **there is no relationship set up yet between the 2 tables** and you **cannot use columns coming from different tables, unless the tables share some kind of relationship.**

Note: In this specific case, **our two tables have no relationship**, and there are times when a relationship cannot be established between 2 tables as there is no common field to link them together. When this is the case, you **need to use a third table**—in the middle—that can then link them correctly together.

To solve the problem, you can use the power of Power BI and bring the **Budget 2016 column from the Excel Budget table into the Sales2015 table.** It turns out that **Power BI Desktop’s Query Editor** is the perfect tool to perform such an operation, **Merge Queries.** From the **Home** tab, click **Transform Data** (or **Edit Queries**)

With **Query Editor** you can manage tables and columns. To tidy up the data David **right clicks** the column name **Budget** and **removes it** as he wants to redo the calculation in Power BI Desktop. He selects the **Home** tab, then the **Transform data** button from the **Queries** group (or **Edit Queries** in the **External data group** and goes back to the **Power BI Desktop Query Editor** window.

He selects the **Sales 2015** query from the left side of the screen and then, on the **Home** tab, he clicks **Combine** button, then **Merge Queries (or Merge Queries** from the **Combine** group (see Figure 4-11).




Figure 4-11: Click **Merge Queries** to **join our 2 queries** Sale 2105 and Budget into a **single query using selected data.**

The **Merge dialog box**, in which you need to specify the destination table (the source is the one selected) and which columns to use to join them together. In David’s case, the columns to use are **CountryRegion and Brand**, in both tables, as illustrated in Figure 4-12 on the following page.

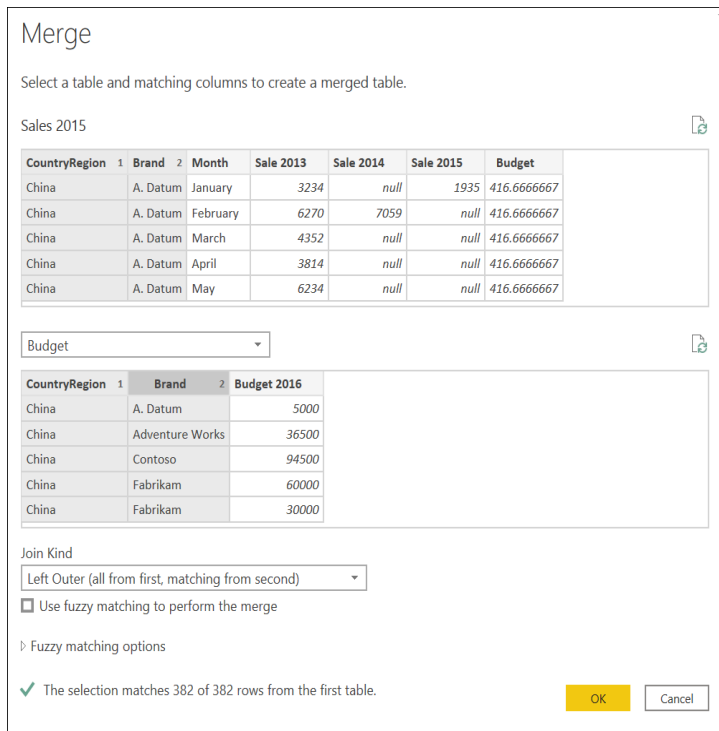


Figure 4-12: In the Merge dialog box, you can select which columns to use when merging two tables. Note the tick mark indicating that the **merge has been successful and 382 rows were matched**. Click **OK**

When you are selecting the columns to merge, Query Editor **may** display a dialog box similar to that shown in Figure 4-13, asking you for the privacy level of the data sources.

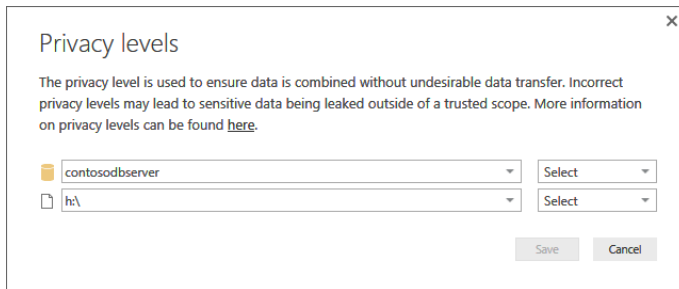


Figure 4-13: Query Editor may need to know the privacy levels of your data sources in order to merge them.

The last option available in the **Merge dialog box** is **Join Kind**. You use this to choose what happens with rows in one table that have no corresponding rows in the other table. For example, if there are sales for a country/region but there is no budget for it, should that country/region be included in the resulting dataset?


The most typical kind of join is the default: a Left Outer join, which includes all the rows from the source table and only matching rows from the merged one. In other words, in David's case, **it retrieves all the sales, plus the budget for the countries/regions and brands with sales**.

Sales 2015 table

To select multiple columns select the first column, i.e. **CountryRegion**, then **hold down either the Shift or Control key** and select the second column, **Brand**.

Budget table

From the Drop down arrow in the box below Sales 2015, select **Budget** and repeat the above steps to select the **CountryRegion** and **Brand** columns



Here are the different types of the JOINS in SQL:

- (INNER) JOIN: Returns records that have matching values in both tables
- LEFT (OUTER) JOIN: **Returns all records from the left table, and the matched records from the right table**
- RIGHT (OUTER) JOIN: Returns all records from the right table, and the matched records from the left table
- FULL (OUTER) JOIN: Returns all records when there is a match in either left or right table

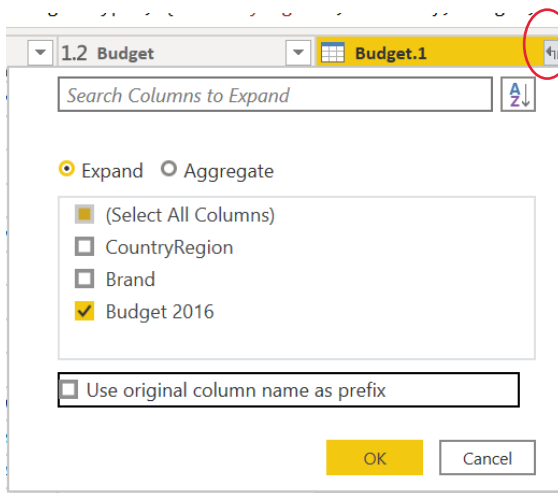
After David clicks **OK**, the table shows a new column with 'Table' showing in each cell as below:

	CountryRegion	Brand	Month	1.2 Sale 2013	1.2 Sale 2014	1.2 Sale 2015	Budget	Budget.1
1	China	A. Datum	January	3234	null	1935	Table	Table
2	China	A. Datum	February	6270	7059	null	Table	Table
3	China	A. Datum	March	4352	null	null	Table	Table
4	China	A. Datum	April	3814	null	null	Table	Table
5	China	A. Datum	May	6234	null	null	Table	Table
6	China	A. Datum	June	5571	3216	null	Table	Table

Figure 4-14: When you merge two tables, the result is the original column, plus a new column of type Table.

In fact, when you merge two tables, the result is the original column, plus a new column of type **Table**. This new column contains all of the rows that are related with the current row in the original table. In David's case, the table contains a single row, but, for more complex joins, it might contain many rows.

David is not interested in retrieving the full table. Instead, **he wants only the Budget column**. To do this, he needs to expand the new column so that it includes only the columns he wants. He can do this easily by clicking the two arrows adjacent to the column name, in the yellow box around the column. This opens the expand column drop down box, as illustrated in Figure 4-15.



This small button on the right side of the column when clicked will **expand to show all the column names** that came from the **Budget** table. By default all the columns are selected.

David only wants the last one. Remove the tick from (Select all and then select only **Budget 2016**. If need be remove the tick from **Use original column name as prefix**. Click **OK**

Figure 4-15: You can choose which columns to include in the result dataset in the expand column dialog box.

In the example, **David selected only the Budget 2016** column, as he does not need the other columns. The result of this is that you now have the **Budget 2016** column populated with the values from the Excel file, as depicted in Figure 4-16 on the following page.

	CountryRegion	Brand	Month	1.2 Sale 2013	1.2 Sale 2014	1.2 Sale 2015	Budget
1	China	A. Datum	January	3234			5000
2	China	A. Datum	February	6270	7059		5000
3	China	A. Datum	March	4352			5000
4	China	A. Datum	April	3814			5000
5	China	A. Datum	May	6234			5000
6	China	A. Datum	June	5571	3216		5000

Figure 4-16: When you expand a column displaying **Table**, it is replaced with the column or columns you chose.

David is nearly done. He wants to remove the **Budget** column and recalculate it in Power BI Desktop. He **right clicks on the name Budget** and selects **Remove**. The last step is that the column, as it is, shows the full yearly budget, whereas the **Sales table should contain only the monthly budget**. Back in Excel, David used a formula and **divided the budget value by 12**; hence, he is doing the same here. David decides to create a formula in Power BI Desktop. On the ribbon, David selects the **Add Column** Tab, he clicks **Custom Column** in the **General** group to create a new column **containing the value of Budget 2016 divided by 12**, as demonstrated below.




Figure 4-17: You can add new calculated columns to your query by using custom columns in **Power Query Editor**.

To create a new column, David must provide the correct **expression** that calculates it. He can do that in the **Add Custom Column** dialog box that opens. For the **New column name**, type **Budget**. Click to the right of the = symbol and from the list of **Available columns**, click **Budget 2016**, then click **Insert** (you can also double-click **Budget 2016**). Click to the right of the = **symbol** that has been inserted into the **Custom column formula box** and type / **12**

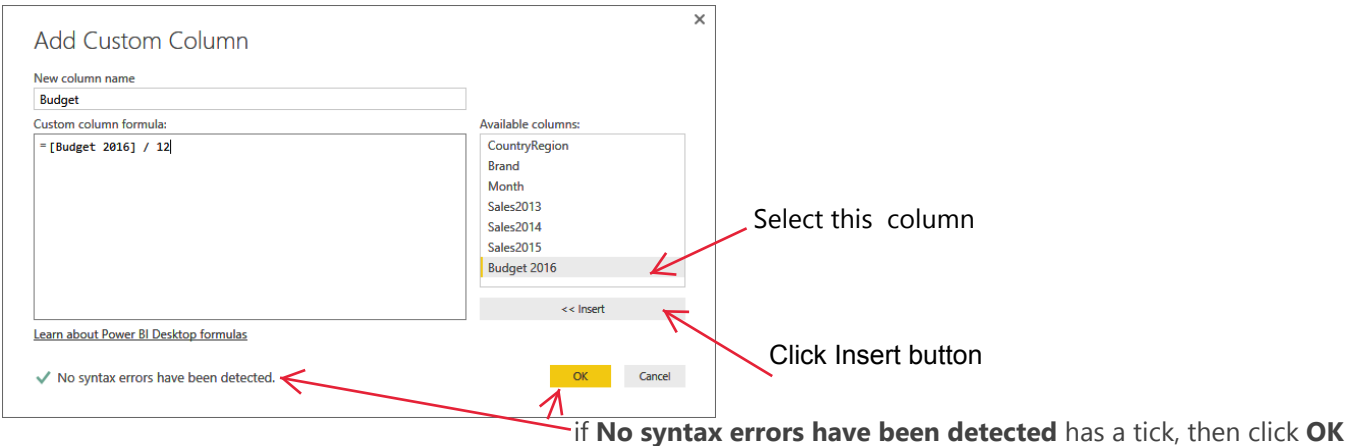


Figure 4-18: When you create a new column, you enter the formula in the Custom Column Formula box.

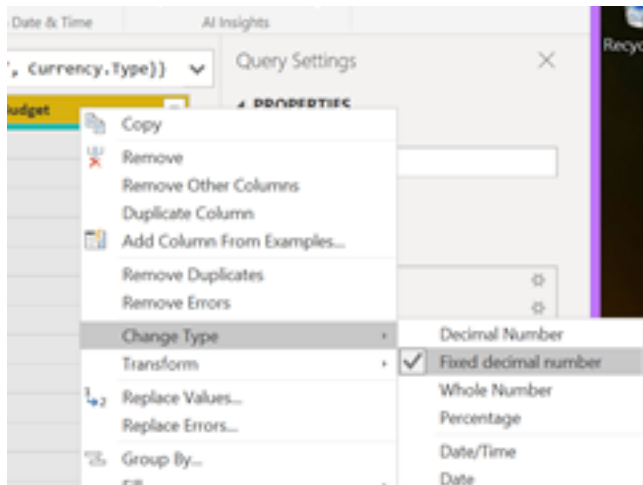
Figure 4-19 presents the result of all these steps, in which you can see the newly computed Budget , which has the same result as in the Excel worksheet.

	1.2 Budget 2016	ABC 123 Budget
35	5000	416.6666667
ull	5000	416.6666667
ull	5000	416.6666667
ull	5000	416.6666667
ull	5000	416.6666667
ull	5000	416.6666667

Figure 4-19: The new Budget custom column appears on the right of the table.

As a final step, David **right-clicks on the Column name Budget 2016** which has the yearly values and **removes** it as he no longer needs it for this report.

Before saving the query, David needs to perform a final step: he defines the data type of the column. **By default, custom columns are of the 'Any data type'**, meaning that the data type is not yet defined. But, David wants to use it to **aggregate values** (which are, numbers) later, so, he must change the data to **Fixed Decimal Number**, as shown in Figure 4-20. With the Budget column selected, right click on the column name **Budget** and select **Change Type**, select **Fixed Decimal Number**, the numbers will display with 2 decimals places.




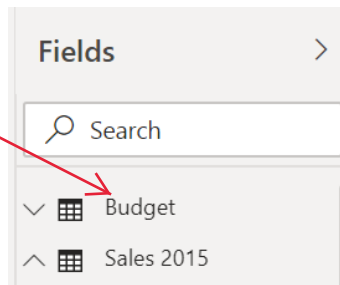
Hiding or Removing tables

David moved the Budget figures to the Sales 2015 table, using the Power Query Editor to create a single table with all the columns required for his report.

Next he clicks the **Close and Apply** button in the **Home tab**, and is returned to the **Dashboard** with the **Visualizations**. In the **Fields** pane on the right he can see that the **Budget** table name is still displayed.

To avoid confusion he can either **hide** the Budget table name, or **stop it from loading**. He initially decides to **hide it**.

Right click on the table name **Budget** in the **Fields** pane



From the pop-out menu, click **Hide**

Figure 4-21: The Fields list now longer show **Budget**

A hidden table, as its name implies, **is no longer visible in the Fields pane**, but you can always make it visible again by choosing View Hidden from the context menu of any table of the field pane and then clearing the Hide check mark. Keep in mind that hiding a table does not mean it is at all secure. A hidden table is only marked as not visible, but any user can see it by simply using the user interface; **hiding a table simply makes the model easier to browse and less error-prone.**

In David's case, **he has now decided to avoid loading the table altogether.** In fact, **all of the information needed to build the reports is now stored in the Sales 2015 table.** The Budget table is used by Query Editor to merge the **budget (divided by 12)** into Sales 2015. After the budget information is stored in Sales 2015, **the Budget table is redundant and does not need to be loaded.**

Select the **Transform data** button (or **Edit Query** button) on the Home tab of the Ribbon, to return to the **Power Query Editor** Window. To avoid loading the Budget table in the future, in the Queries pane, on the left side of the screen, right-click **Budget** query then, on the shortcut menu, **clear the check mark beside Enable Load**, this now means the Budget table is not automatically loaded when you are working with the data, the name will appear in Italics. If required you can right-click and place the tick back again.

Clear the check box next to **Enable Load** to turn of Table Loading





Figure 4-23 - Warning message

Figure 4-22: Like Hiding a table, you can also turn on or turn off loading for a table by using Query Editor.

When you turn off loading for a table, **Query Editor warns you with the message shown above.**

If you continue, the table will be removed from the model, which becomes a single-table model again, with Sales 2015 containing all the relevant information. Click **Continue**. The chart may reappear with a **Fix it** message, click this and you will see **Sales 2015 by Brand** chart

Note: Remember you added the data you needed from **Budget** into the **Sales 2015**, so it is now part of it, if it is not showing on the chart, click it to show the budget data on the chart.

Handling seasonality and sorting months

Remember in the formula created earlier, David **split the budget figures by 12**, but, in reality, many brands show some seasonal effect that is not taken into account while computing the budget. Moreover, because some brands do not have sales at all in some months, the final report does not contain all the months.

Click **Close & Apply** to return to the **Power BI Desktop** and **add a new Page**. From the **Visualizations**, select the **Clustered Bar Chart**. From the **Fields** list select **Budget and Brand**. Click on the **More options**, (the 3 ellipse dots on the right). **Select Sort by Brand**, then sort in **Ascending Order**. Under **Visualizations**, under the **Fields icon**, find the title **Budget by Brand** and format it. From the **Visualizations** find the **Slicer** icon and tick **CountryRegion**, a small box tile will appear with the three countries, listing the 3 countries. From the **Budget by Brand**, click **World Wide Traders** and you will see **China** is selected in the **Slicer for CountryRegion**

Click away from the chart into an empty area on the Central Pane. Select the **Table visualization** and select **Month, Sale 2015 and Budget**. When the 2 visualization tiles are on the dashboard, click **Wide World Importers in the chart for China**. You can spot the problem easily by looking at the budget data report for Wide World Importers in China, as depicted in Figure 4-24.

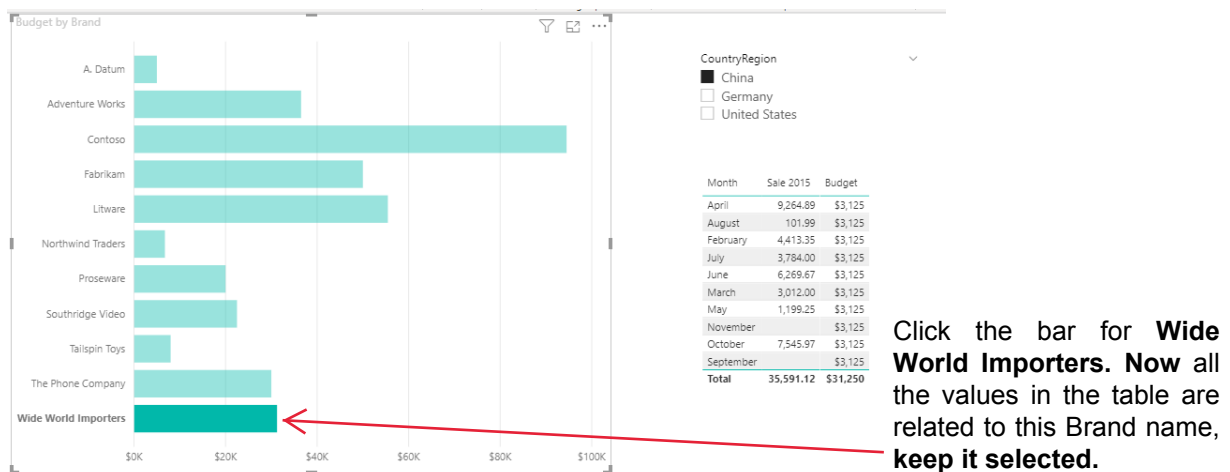


Figure 4-24: Note that some brands in China, for example, have no sales in September and November. There are two issues with the report in Figure 4-24 as below: **Save the file as Sales**

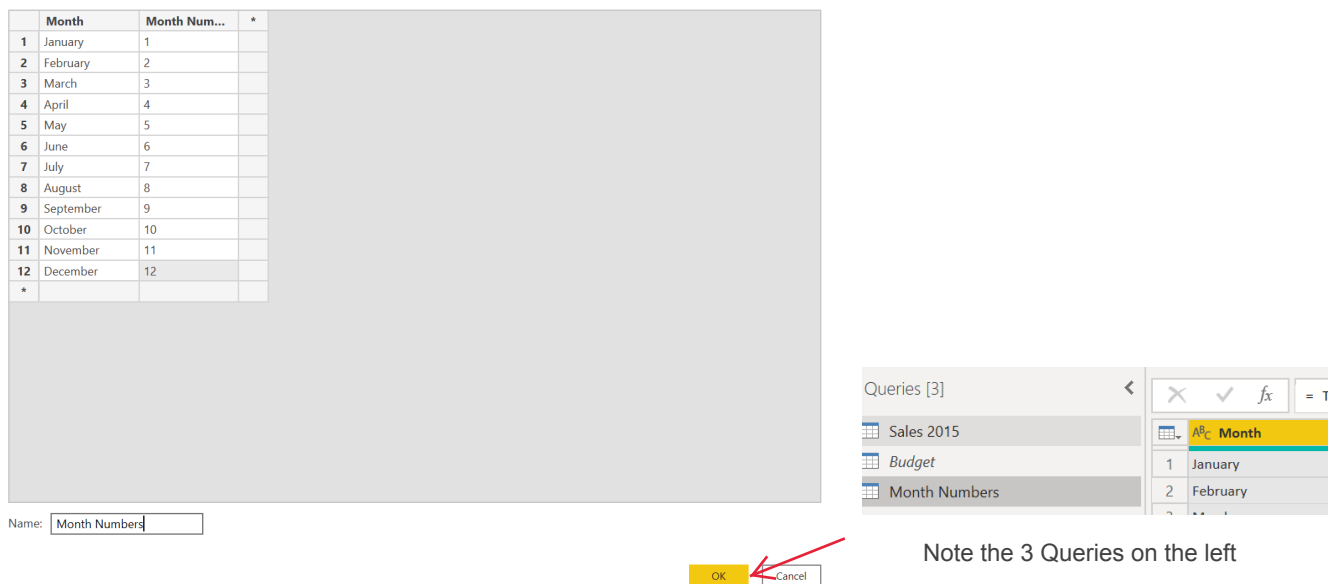
- The Sales figures for **January** and **December** are missing from the table data. **This is because there were no sales in January and December, so the corresponding rows are missing.** This is a big issue, because the budget is computed as the **total budget divided by 12**, but **only 10 rows are accounted for in the final figures. As a result, the budget values of the reports are wrong.** In fact, the **budget for World Wide Importers in 2016 was 37,500**, whereas the report shows a total of only **31,250**.
- While searching for the missing months, you might have noticed **that the months are not sorted in sequential order.** In fact, **by default, Power BI sorts each column alphabetically**, which, of course, is **not the correct way to sort months**.

Note When you have a date column in the data and you use it in Power BI Desktop, the month name proposed is already sorted alphabetically. However, in this case, the data source contains the month name and not a date, so you need to correct the sorting order.

David is determined to solve these two issues, beginning with the last one, which is somewhat easier. **To sort the months by sequential order**, he needs a **new column in the Sales table** containing a number, **ranging from 1 to 12**, which contains the sort order of the month. The problem is there is no such column available, and there is no predefined functionality to achieve this goal.

To go back to the **Power Query Editor** window, click **Transform Data** (or **Edit Queries**) Power BI Desktop offers you a great feature when you have some data to add to an existing model: you simply enter it. To do this, in the **Home** tab, in the **New Query** group, click **Enter Data**, and **Query Editor** shows you a grid in the **Create Table** dialog box, in which you can type (or paste) the data you want to add to the model. Figure 4-25 shows how David used this feature to **create a table containing the month names and numbers**. At the bottom of the table, type the name for the Table, **Month Numbers**, and click **OK**.

Create Table



Note the 3 Queries on the left

Figure 4-25: Using the Enter Data functionality, you can type or paste new datasets.

When you click **OK**, the table appears in the Power Query Editor window.

The next step, after saving and naming the table as **Month Numbers**, is to bring the **Month Number column** from this table into the **Sales 2015** table. The technique you use is very similar to what David already did with the budget: join the Sales 2015 table with Month Numbers. This time, the relationship is based on the month name. Figure 4-26 presents the Merge dialog box already prepared.

Keeping the **Power Query Editor** window open, select **Sales 2015** from the left side of the screen, in the Queries list. From the **Home** tab, click **Combine** button, then, click **Merge Queries** (or **Merge Queries** in the **Combine** group). The **Merge** dialog box opens as on the following page.

Merge

Select a table and matching columns to create a merged table.

Sales 2015

CountryRegion	Brand	Month	Sale 2013	Sale 2014	Sale 2015	Budget 2016	Budget
China	A. Datum	January	3234	null	1935	5000	417
China	A. Datum	February	6270	7059	null	5000	417
China	A. Datum	March	4352	null	null	5000	417
China	A. Datum	April	3814	null	null	5000	417
China	A. Datum	May	6234	null	null	5000	417

Month Numbers

Month	Month Numbers
January	1
February	2
March	3
April	4
May	5

Join Kind
Left Outer (all from first, matching from second)

Use fuzzy matching to perform the merge

Fuzzy matching options

✓ The selection matches 382 of 382 rows from the first table.

OK Cancel

Select **Month** from the **Sales 2015** table then select **Month Numbers** from the drop down list and select **Month**

Figure 4-26: Merging the Sales 2015 table with the **Month Numbers** table.

Click **OK** only if 'The selection matches 382 of 382 rows from the first table' After the merge, you still need to expand the **Month Numbers** column by clicking the **Expand** button to the right of the column and select **Month**. **Remove the tick from Use original column name as prefix** and Click **OK**

Expand Aggregate

(Select All Columns)

Month

Month Number

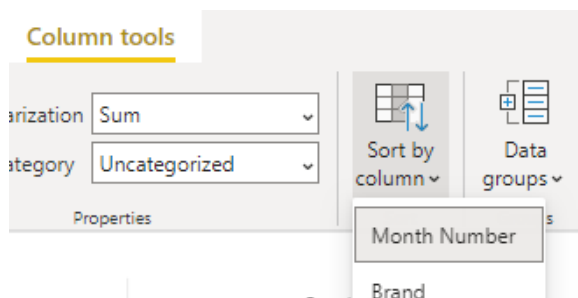
Use original column name as prefix

OK Cancel

We are now ready to go back to the desktop and modify the table on **Page 3**. Click the **Close and Apply** button.

Now that the **Month Number belongs to the Sales 2015 Table**, you need to instruct Power BI to **sort the month names by month number**. Select the **Table** that has the **months currently listed out of order**.

In the **Fields** pane, click the field name **Month Number**. Some extra tabs appear on the Ribbon, from **Column Tools** select **Sort by column** (or select the **Modeling** tab, click **Sort By Column** (number 1 in Figure 4-27), and then select Month Number (number 2 in Figure 4-27).



Using the Column Tools table to sort Month Number

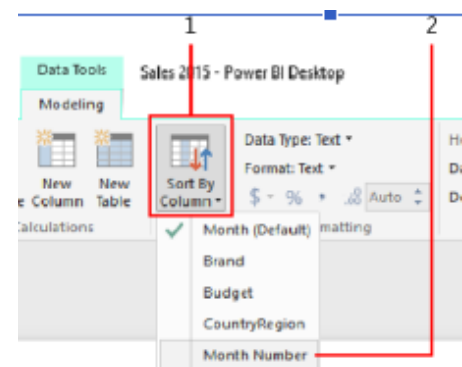



Figure 4-27: Using the Modeling tab to sort Month Number



Add the Month Number to the table. Sort by Month Order in Ascending Order. Check that you have select World Wide Importers so your data matches the above.

Figure 4-28: In this report, the month names are now sorted properly.

Now that David has the months displaying correctly, it is much clearer that **January** and **December** are missing, so he turns his attention to fix this. Because this requires a bit more work, there needs to be a bit of fore planning.


First, David needs to decide which year to use as a basis for demonstrating seasonality. In fact, **Wide World Importers** might have sales in January 2014 and no sales in January 2015. Should he consider 2014 or 2015 to decide what to allocate in January? **David goes for 2015** because it shows the best figures. You might make different decisions here, but as we have cautioned several times already, keep in mind that this is a book about Power BI, not a budgeting tutorial. So, please be patient; we are very naïve regarding choices like this one. Now that David has made his decision, he needs to build a table containing, for each country/region and brand, the number of months for which there are sales in 2015.

This requires several steps in Query Editor:

1. Start from Sales 2015, and then remove all the unwanted columns, to keep only CountryRegion, Brand, Month, and Sales2015.
2. Remove all the rows in Sales 2015 that are empty.
3. Remove the Sales2015 column.
4. For each brand, count the number of months.

The first part of **step 1** is easy: You need to go back to the **Power Query Editor** window. Click **Transform Data** in the **Queries** group of the **Home** tab (or **Edit Queries** in the **External data** group of the **Home** tab). On the left side of the window, under **Queries**, right-click **Sales 2015**, and then, on the shortcut menu, click **Duplicate to make a copy of the table**. Right click this copy and rename to **Months Count**, keep selected.

The second part of step 1 is also easy: on the right side of the screen, under **APPLIED STEPS**, using the small delete icon that appears on the left of each step (see Figure 4-29), remove all of the steps, keeping only the **first two (Source and Navigation)**, to return to the **original query** we used from the **Excel file**.




Click on the small X to the left of each step under Applied Steps, leaving only Source and Navigation.

Figure 4-29: Use the small delete icon to remove unwanted steps until only **Source** and **Navigation** remain.

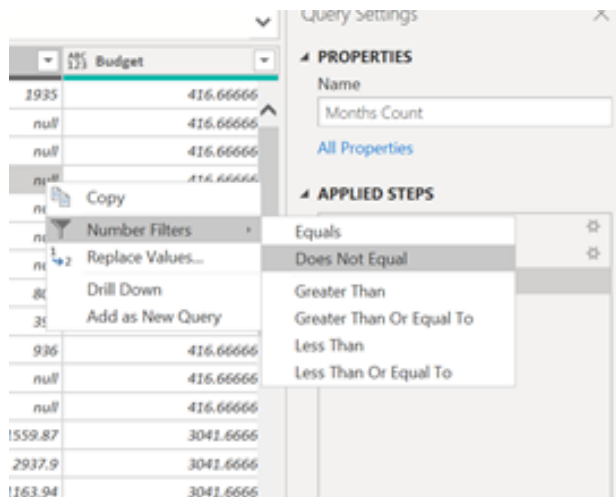
Remember, we are working on a copy of Sales 2015, so we are free to update it as needed. The original table remains untouched.

*Note: If the column header names have become Column numbers, and the first row are the original Column Heading names, in the Home tab in the **Transform** group, click **Use First Row as Headers***



At this point, **you can remove the Sales2013 and Sales2014 columns**, which are not needed for the purpose of this query. To do this, right-click the column header for each column and then, on the shortcut menu, click **Remove Columns**. **Step 1 is done!**

Moving on to **step 2**: David scrolls down the table and notices that there are **null values** in the **Sales2015**. He wants to remove these values, so he **right-clicks a null value to open its shortcut menu**, where he chooses **Number Filters**, then **Does Not Equal**, as shown in Figure 4-30. All the null values are now removed. *Note: The number of records at bottom left of screen is now 230*



Note: If you click on a row number on the left of the table, the bottom of the screen shows all its properties as below. Click on a column name heading to clear.

CountryRegion	China
Brand	A. Datum
Month	January
Sale 2015	1935
Budget	416.6666667

Figure 4-30: Right-click a cell value to filter rows using several criteria.

Step 3: At this point, the only columns required are **CountryRegion, Brand** and **Month**. David removes any other columns. **Steps 2 and 3 are now also done**, in just a few clicks.

Finally, **step 4**: David would like to **group** the current dataset by **CountryRegion** and **Brand**, then count, for each combination, the number of months. Because this is a very common operation on datasets, Query Editor offers a specific functionality: first, **select the columns to group by**, i.e. **CountryRegion** and **Brand**. Select the column head name, **CountryRegion**, then hold down the **Ctrl** or **Shift** key and select the column name **Brand**. In the **Transform** group in the **Home** tab, select **Group by**.

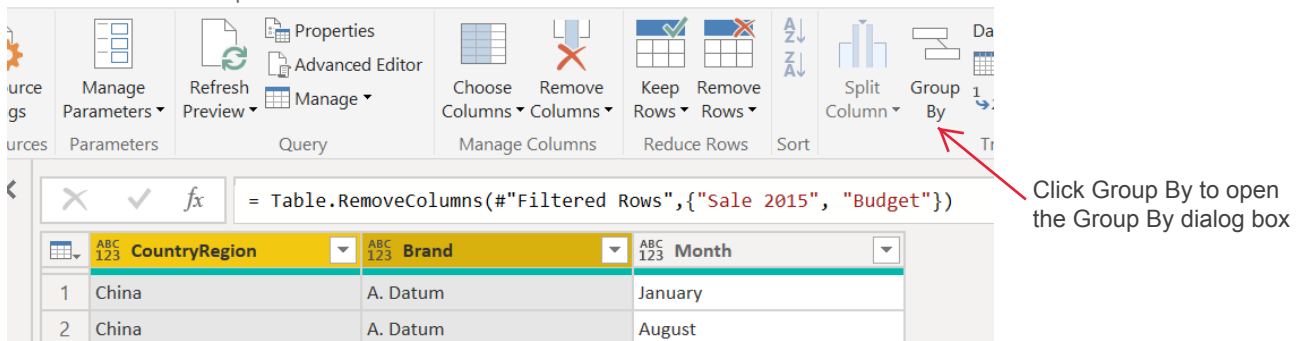
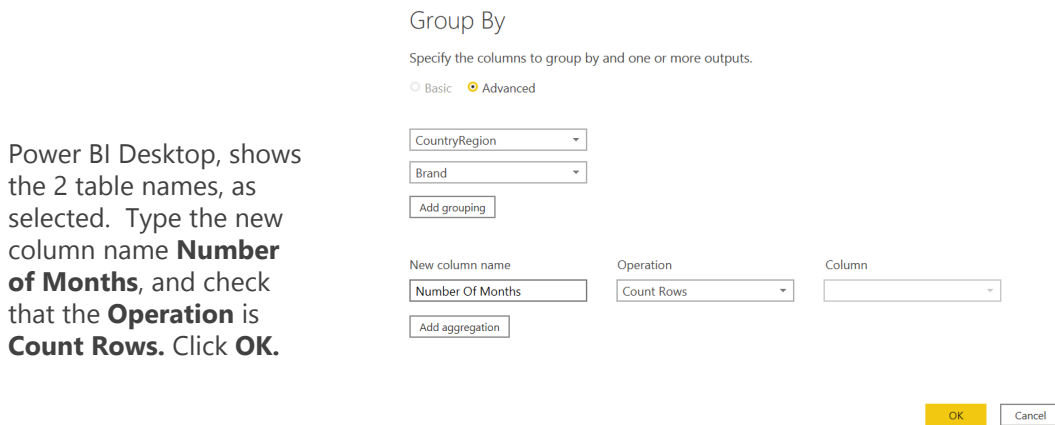


Figure 4-31: First, select the columns to group by, and then click **Group By** to open the **Group By** dialog box.



Power BI Desktop, shows the 2 table names, as selected. Type the new column name **Number of Months**, and check that the **Operation** is **Count Rows**. Click **OK**.

Figure 4-32: After he clicks **OK** to confirm the grouping choices are correct, David sees the dataset shown below

	CountryRegion	Brand	Count
1	China	A. Datum	4
2	China	Adventure Works	7
3	China	Contoso	10
4	China	Fabrikam	5
5	China	Fabrikam	2
6	China	Litware	6
7	China	Litware	1
8	China	Northwind Traders	3
9	China	Proseware	7
10	China	Southridge Video	7
11	China	Tailspin Toys	8
12	China	The Phone Company	5
13	China	Wide World Importers	8
14	Germany	A. Datum	6
15	Germany	Adventure Works	6
16	Germany	Contoso	10
17	Germany	Fabrikam	4

In this grouping result, each country is grouped together, China, Germany and United States, and the **Brand** is in Ascending order by country.

To view the **Brands** grouped for each country, change the sorting of **Brand** to **Ascending order**, then change the sorting of **CountryRegion** to **Ascending** also, which gives the above result.

Figure 4-33: Both tables show a **maximum of 10** for the **Number Of Months**. David wants the 2nd result, Brand grouped by CountryRegion

The new dataset indicates how many months, i.e. 10, are present for each brand and country/region. You need to use this number, instead of 12, to obtain the correct value to use for each month. The table is no longer needed. Turn off the loading for it, by right clicking on the **Months Count** query on the left of the screen and removing the tick from **Enable Load**. You may receive a **Data Loss Warning** message, click **Continue**

The last step is to **modify Sales 2015** to take this number into account. This time, you do not need to add further steps to an existing query; however, you need to replace some of them, and this requires a bit more attention. If you reopen the **Sales 2015** query in **Power Query Editor** and begin navigating through the Applied Steps panel (the history of all your steps), you notice that what is displayed in the results pane reflects what the query looks like after having applied the selected step. For example, in Figure 4-34 you can see that **when you select the Added Custom** step, the query shows the **Budget 2016** column, which will be removed by the next step.

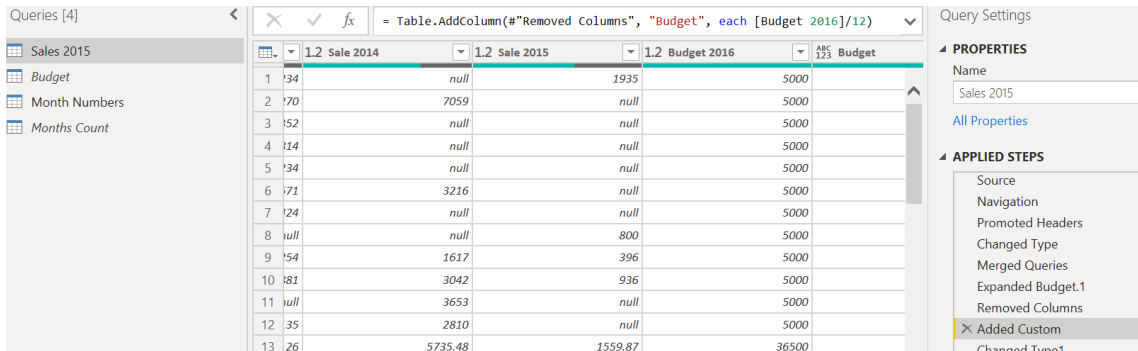
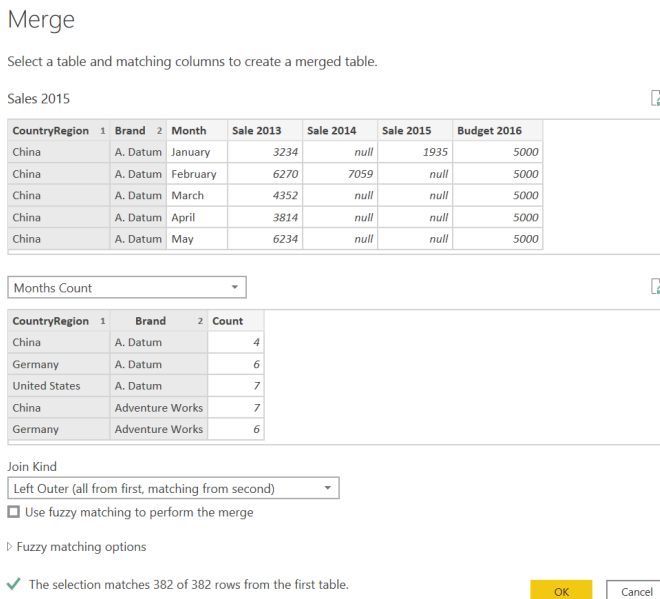


Figure 4-34: By navigating through the Applied Steps, you can view partial results of the final query.

You need to **add a few steps before the Added Custom step** (which is where you created a formula to divide the Budget by 12) and then **modify the calculation of the budget itself**. When you choose an operation from the toolbar, the step is added right after the currently selected one. **Select the fourth step (Removed Columns)**. In the Home tab, select **Merge Queries** from the **Combine** group

*Note: A warning message appears, to check that you want to do this Insert Step, click **Insert**.*

The **Merge** dialog box appears as below, use this as a guide to complete the merge. Click **OK** when your screen matches.



Note: The join is Left Outer (default) and there is a tick to the left of 'The selection matches 382 of 382 rows from the first table' an indication that the merge worked.

Figure 4-35: The Merge brings the Number Of Months column into Sales 2015.

After you add the column, it appears with 'Table' in each cell. Click on the expand button on the right and select **Number of Months** from the list, the column now populates with numbers. Next, the old formula must be updated. In fact, you want to **calculate the Budget 2016 column divided by the number of months**, for only the months for which there are sales in 2015. i.e. 10

To perform this operation, in the **Applied Steps** area of the **Query Settings** pane on the right, click the settings button (the small "gear" icon) to the right of the **Added Custom** step. The formula you created early appears. Change the expression accordingly, so that it appears like that shown in Figure 4-36. The formula is typed into the Custom Column formula box.

*Note: You do not need to type a column name, simply select it and click **Insert**.*



Note: As you type the formula, press the space-bar at the end of each line to stop the smart formula feature working

```
= if
[Sale 2015]<> null
then
[Budget 2016] / [Number of
Months]
else
0
```

Figure 4-36: To check that the report has updated click **Close & Apply**. The updated report appears as below




Figure 4-37: The report now shows the correct value of 37,500.00 for the budget for the Brand, Wide World Importers. Save the file.

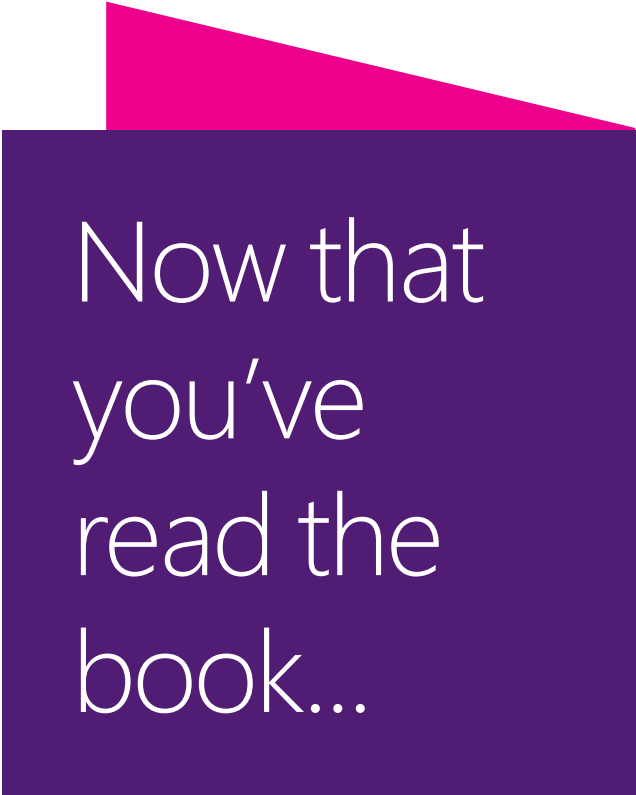
Conclusions

In this chapter, you learned the basics of Power BI Desktop, which is a desktop application that brings the full strength of Power BI to your PC. This was a basic introduction, and we will explore more features later in this book. In fact, there is a lot more to learn, but that would be beyond the scope of this book.

Here are the most relevant features:

- Power BI Desktop can load data from any database. In the example in this chapter, we used Microsoft SQL Server.
- Using Power BI Desktop, you can load data from multiple sources. In the example, we mixed data from Excel with data residing in a SQL Server database.
- Power BI Desktop uses Query Editor to load data. Query Editor offers many powerful features. We highlighted in particular the capability of merging different queries and adding calculations to the query.
- Some queries are loaded into the model; others are useful only to compute values in the main query. You can mark queries that should not go into the model as “do not load” so that they are used only in Query Editor.
- You can upload a Power BI Desktop model to the Power BI online service, and it retains the same refresh features: by using the Personal Gateway, you can refresh a Power BI model in the cloud, letting it access data on your PC.

At the beginning, it might look complex, but after you get used to it, Query Editor offers you a lot of power to build your models. Having successfully finished this chapter, you can call yourself a data modeler. Later in the book, we will introduce some more features of Power BI Desktop to further enhance your model with more advanced calculations.



Now that
you've
read the
book...

Tell us what you think!

Was it useful?

Did it teach you what you wanted to learn?

Was there room for improvement?

Let us know at <http://aka.ms/tellpress>

Your feedback goes directly to the staff at Microsoft Press,
and we read every one of your responses. Thanks in advance!



Getting data from services and content packs

Working on the forecast for the next year, David realizes that it would be useful to consider the statistics for the pages visited on the company website. Using these insights, he might anticipate which products will gain more traction in the upcoming months.

Microsoft Power BI provides many connectors and content packs that make it possible for David to easily access the data generated from different cloud services. Content packs are also a useful tool to deploy and share predefined models and reports within a company.

In this chapter, you will see how David imports data from Google Analytics into Power BI, using different techniques. You can use these same techniques to access many web services, the list of which is growing rapidly, week by week.

Note In the following examples in this chapter, we will show data gathered by using Google Analytics from the website www.daxformatter.com, which we created to help users format DAX expressions and queries. We will use this data as part of the scenario we create for David. Our assumption is that the data would make sense in the sales generated by Contoso, too. But, also, we would not have enough data to show if we created a fictitious website for these examples. For this reason, when you continue reading, assume that visitors from this website are meaningful for sales of Contoso products, even if we know this is not true. If you want to replicate the same example, you can use Google Analytics data for a website to which you have access.

Consuming a service content pack

David wants to analyze data regarding customer visits to the website in an effort to glean some early indicators of a potential growth in certain countries/regions. If the number of visitors to the website increases in certain countries/regions, there also could be a growth in sales for the same country/region. Looking at this information for the previous two years might help in the budgeting process by providing data that will help David define the sales target for each country/region. Comparing website visitors and sales by origin in the same report is an important step in the analysis that he wants to accomplish.

David knows that his company's website is monitored by Google Analytics, and he wonders whether Power BI supports it. Reading the Power BI documentation, he finds that Google Analytics is indeed supported as a service content pack in the Power BI service, and as a connector in Power BI Desktop. With that, he begins using the service content pack in the Power BI service.

David starts Power BI and then, in the lower left corner of the window, he clicks the Get button. He is then greeted by the Get Data page, as depicted in Figure 5-1.




Figure 5-1: The choices available to you to get data in Power BI.

When you load data from a service, Power BI automatically creates for you a data source, a report, and a dashboard connected to that service, using predefined templates for each of these elements. You can modify these objects later if you want. Obviously, Power BI will use your credentials to access the service. Thus, even if the report is the same, the numbers will represent your own data.

What is a content pack?

A content pack can include the following parts, which have dependencies among them:

Dataset One or more datasets provide data to reports and dashboards included in the content pack library.

Report Each report in a content pack connects to a dataset included in the same content pack. The datasets required by reports are always part of the pack.

Dashboard Each dashboard of a content pack includes visualizations of one or more reports included in the same content pack. The reports used by a dashboard are always part of the pack.

David wants to use a content pack for a service; so, in the Content Pack Library section of the Get Data page, on the Services tile, he clicks the Get button. This displays a list of the services that are available, as shown in Figure 5-2.

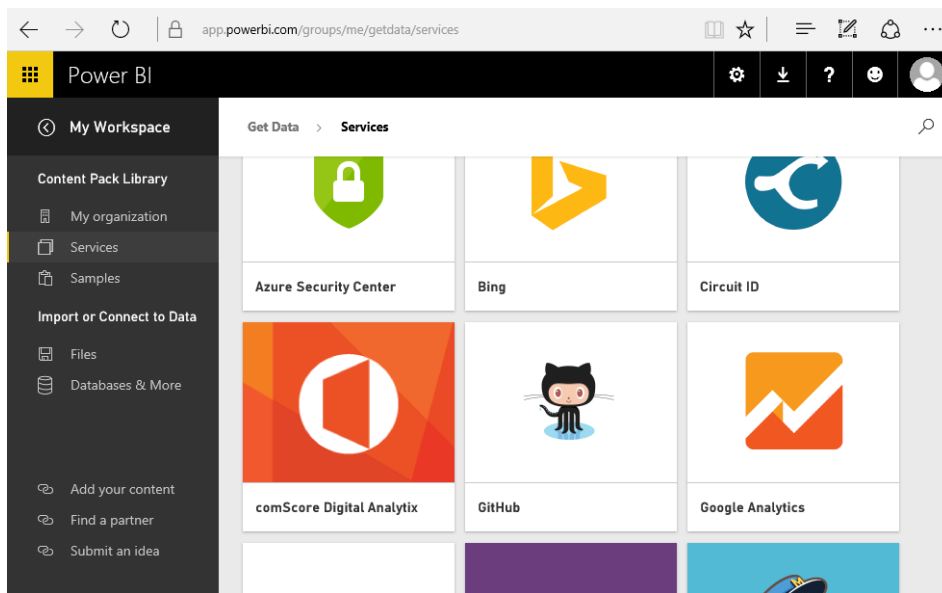


Figure 5-2: A partial list of the services that are available in Content Pack Library.

He clicks the Google Analytics tile and sees a message containing a description of the service, as illustrated in Figure 5-3.




Figure 5-3: The description of the selected Google Analytics content pack.

Note For all available services, Power BI presents a similar description along with a Connect button. However, the steps that follow after you click Connect might differ from those for Google Analytics, depending on the security model and the implementation of the authentication of the user for the selected process. So, keep in mind that if you select another service, the experience might be altogether different from what we describe in this chapter.

David clicks Connect and is then prompted to choose the authentication method to use to connect to Google Analytics. The only choice available to him is OAuth2, as depicted in Figure 5-4.

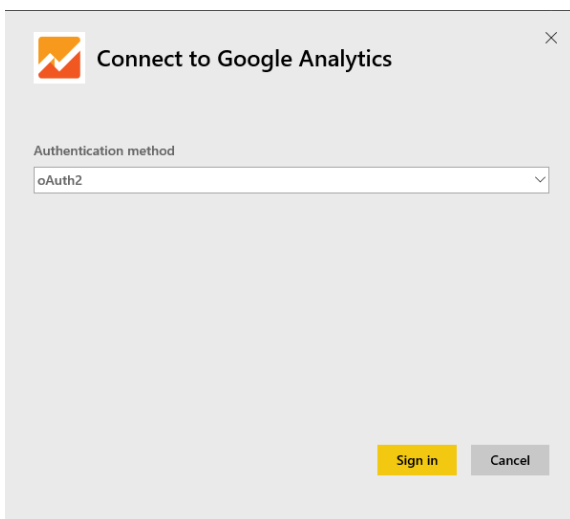


Figure 5-4: You must choose an authentication method when connecting to Google Analytics.

Be aware that after you click Sign In, you might be requested to sign in to Google Analytics (see Figure 5-5), unless your credentials are already stored because of previous access.




Figure 5-5: The Google sign-in page.

When David completes the sign-in, he is asked to provide offline access to the Power BI application, as shown in Figure 5-6. To authorize Power BI to retrieve data from the Google Analytics service on his behalf, he clicks Allow.

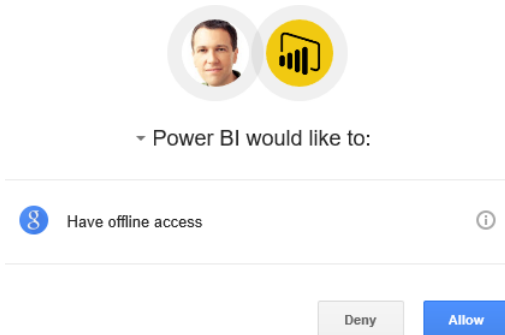


Figure 5-6: Confirming offline access to Google Analytics for Power BI.

Note By allowing “offline access” you are indicating that you want to allow the Power BI Desktop application to interact with Google Analytics even when you are not interacting with the Google Analytics service directly. For example, when you navigate in the data offered by Google Analytics in a web browser, you are engaging in an “online access,” because you directly interact with the service. However, when Power BI Desktop requests data from Google Analytics, it will act on your behalf, and you will not see the details of every request. This is true not only when you use Power BI Desktop interactively, but also whenever you schedule an automatic refresh of the dataset in Power BI.

On the next page of the Connect To Google Analytics dialog box, David can choose which part of the data available in Google Analytics that he wants to bring into the content pack. There are three options (see Figure 5-7):

- **Account** The account name for Google Analytics (a single user might have access rights to multiple accounts). Select from the list of available account names for the current user.
- **Property** The name of the property (which is a Google Analytics concept) within the data owned by the selected account.
- **View** The view name within the property. Oftentimes, the view corresponds to the property, unless you handle multiple websites within the same account.




Figure 5-7: Choosing what part of the data to import from Google Analytics.

Note In our example, for the DAX Formatter property, the View option is unavailable because there is only one view from which to choose.

When David clicks Import, the Power BI service copies the content library into the currently selected workspace, and then it updates the connection with the information provided in the previous steps. Finally, it populates the workspace with data read from the Google Analytics service. As a result, he obtains one dashboard, one report, and one dataset named Google Analytics, which he can rename if necessary (you should do that in case you import multiple copies of the same content pack in the same workspace—in this case, they would look the same, but they would contain different data).

The Google Analytics dashboard includes information about the traffic received in the past 30 days, as demonstrated in Figure 5-8.




Figure 5-8: A dashboard created by the Google Analytics content pack.

Let's take a moment to discuss this dashboard. If you click one of the dashboard's visualizations, you are moved to the underlying data in a corresponding report. All of the visualizations come from the same report, which has pages that filter the data from the past 30 days (such as Total Users), the past 90 days (such as Site Traffic), or the past 180 days (such as System Usage, Page Performance, and Top Pages). Figure 5-9 presents the System Usage page for David's report.




Figure 5-9: A report created by the Google Analytics content pack.

Back to David; he can now edit the report created by the content pack, or he can create a new report based on the same dataset used by the existing report. In both cases, one of the issues is that the dataset provided in this content pack does not include all the possible measures and slicers available

in the Google Analytics server. Also the available historical depth is limited to a maximum of 180 days. Figure 5-10 demonstrates that the number of tables, attributes, and measures available is just a fraction of the measures available in Google Analytics.




Figure 5-10: The tables, attributes, and measures available in the dataset of the Google Analytics content pack.

The reason for this is that the engineers who created this content pack tried to include the minimum amount of information required to create the desired reports. In this way, the size of the resulting Power BI file (.pbix) is kept to a minimum, improving the performance of many related operations. However, the level of details available might not satisfy your requirements, depending on what you are trying to accomplish.

The Google Analytics content pack has been useful for an initial overview of the data available, but it does not meet David's requirements. First, he needs greater historical depth to analyze the trends, and the limit of 180 days provided in the content pack is not sufficient for his requirements. Second, he would like to create a single report showing the relationship between data from Google Analytics and other data, such as past sales and forecast. This requires a single dataset with multiple connections, so getting data from the Google Analytics content pack is not very helpful, because you cannot change or customize the data model of a dataset copied from a content pack, whereas you can customize the reports and dashboard imported from a content pack.

For these reasons, David will create a new dataset using only the connection to Google Analytics in Power BI Desktop, without relying on the content pack he has used thus far.

Creating a custom dataset from a service

David wants to show in the same report some measures from Google Analytics related only to China, Germany, and the United States, which are the countries/regions interested in the budget process. He wants a better result than the one he can obtain by using the Google Analytics content pack on Power BI. Also, he needs greater historical depth than what is available through the service he tried earlier.

Instead of using the predefined (and read-only) content pack, David will use Power BI Desktop to connect to Google Analytics, and then include the Google dataset in the budget model he is developing. Of course, this requires some more effort, but it provides him with the advantage of flexibility in the definition of measures, calculations, new tables, and relationships in the data model.

Thus, David begins with the model he created by using Power BI Desktop in Chapter 4, which, to refresh your memory, has one table that contains the sales and budget in different columns. Such a table will be useful to analyze historical trends in sales. However, to import other tables from Google Analytics, he needs to use a special connector that imports data from Google Analytics directly into the model of Power BI Desktop.

David starts Power BI Desktop. On the ribbon, on the Home tab, he clicks Get Data and then selects the More option. In the Get Data dialog box that opens, in the pane on the left, he clicks the Other category, and then, in the pane on the right, he clicks Google Analytics, as shown in Figure 5-11.

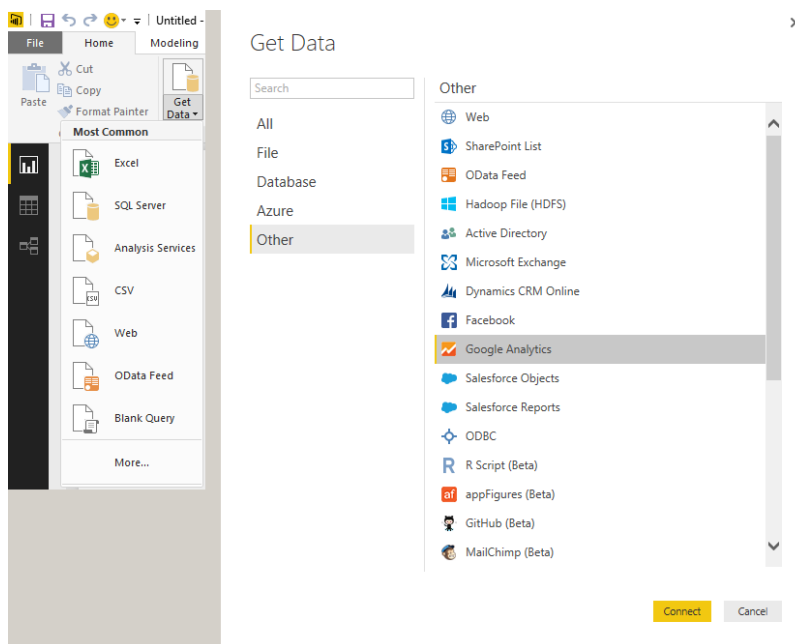


Figure 5-11: The Google Analytics connector is available in the Get Data dialog box.

When he clicks the Connect button, he is asked to sign in to his Google Account (Figure 5-12) so that Power BI Desktop will be able to access information in Google Analytics to which his Google Account has permission.

Note You might see the Connecting To A Third-Party Service message box, warning you that the features, updates, and availability change often. Click Continue to close the message and move on to the next step.

When you sign in for the first time, you must provide your user name and password. (The sign-in process might use 2-factor authentication if required.) After you complete the sign-in process, you can click the Connect button to move forward.

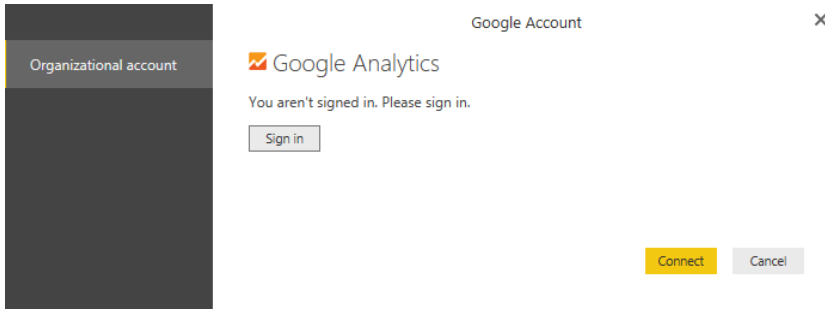


Figure 5-12: Connecting to Google Analytics from Power BI Desktop requires that you sign in to a related Google Account.

When David completes the sign-in process, he then needs to select the attributes and the measures to import in the data model. The Navigator dialog box provides a list of all the services monitored by the Google Account that he is using. After he selects the service, he is presented with a list of folders containing attributes and measures that he can select and import. The attributes are information collected by Google Analytics, such as date-related columns (year, month, day, etc.), demographic and geographical information about the visitors (such as age, gender, country/region, and city), and many other pieces of information that can be used to group and filter data. The measures are numeric information describing the frequency or the size of an event; for example, the number of users, the number of visits, the average time to load a page, and so on.

Querying a service such as Google Analytics is similar to querying a data model by using a pivot table. You select certain attributes, and measures are automatically aggregated at the **granularity** defined by the attributes included within the same report. For example, Figure 5-13 depicts the preview of the result obtained by selecting the Country/Region attribute from the Geo Network folder, the Year attribute from the Time folder, the Sessions measure from the Session folder, and the New Users and Users measures from the User folder.

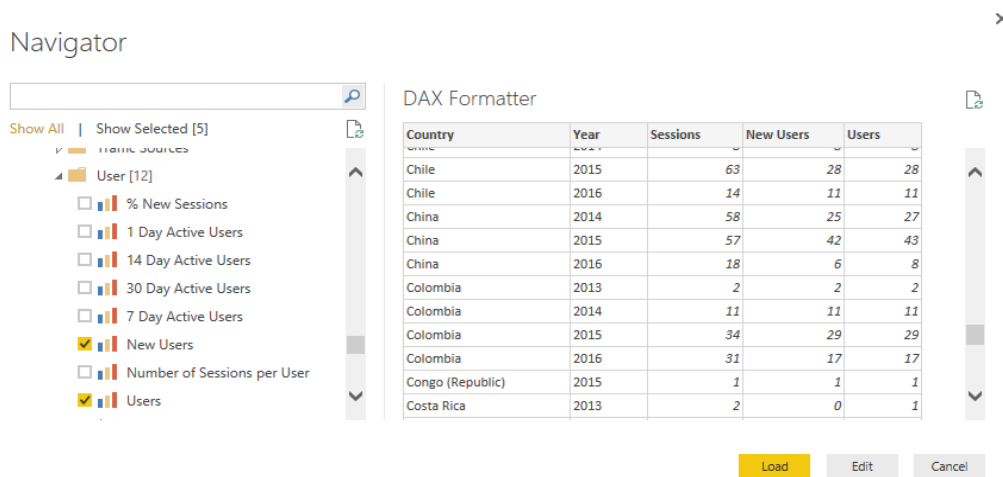


Figure 5-13: Connecting to Google Analytics, you select the measures and the attributes to import in the data model.

Choosing the right granularity

The approach for getting data from a service such as Google Analytics is different from the one you use when you collect data from a relational database (such as SQL Server) or from an external Excel file. In these cases, you always see the data at the maximum level of detail (also known as “fine granularity”). The selection of the attributes defines the granularity, and you can still change the query later in Query Editor, but you must add other attributes to the query in order to increase the granularity. You have a similar user interface in Query Editor when you import data from an external rich semantic model, such as the one provided by Analysis Services and SAP HANA.

In Power BI Desktop, on the Home tab, David clicks Edit Queries to open the Query Editor window. Because David wants to analyze a limited number of countries/regions, he needs to apply a filter in Query Editor to the Country/Region column, restricting the selection to China, Germany, and the United States. This obtains the result illustrated in Figure 5-14.

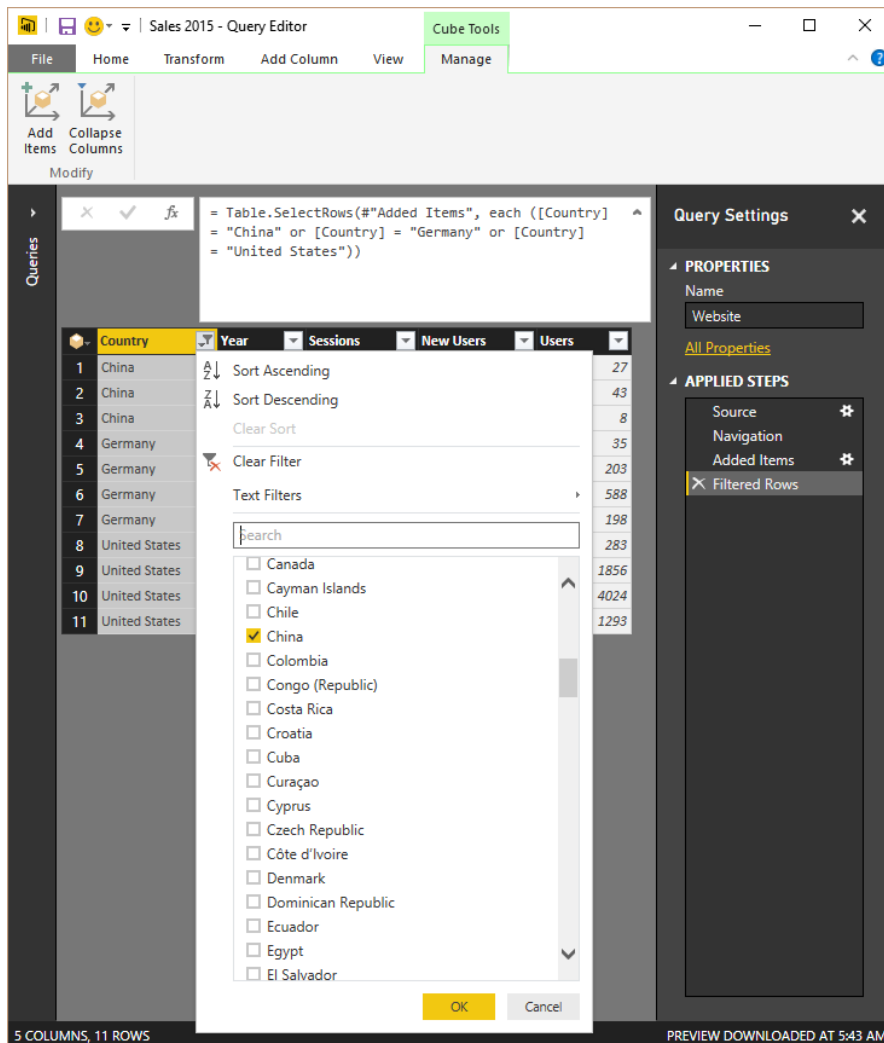


Figure 5-14: Using Query Editor to refine the query for getting data from Google Analytics.

Also note in Figure 5-14 that the ribbon includes a special tab: Cube Tools | Manage. The name “cube” references an external rich semantic model, and as we said earlier, it is the same approach available for Analysis Services and SAP Hana. When you click Add Items, you can select additional attributes and measures that will modify the query made to Google Analytics, enhancing the granularity (with

attributes) or more information (with measures). Figure 5-15 depicts the different graphical representation of attributes (in the Audience folder) and measures (in the DoubleClick Campaign Manager folder).

Note You should use the Collapse Columns button whenever you want to remove an attribute and obtain a corresponding granularity without the attribute you removed. If you just remove an attribute by removing the corresponding column in Query Editor, you do not modify the original query, and the cardinality will still include the attribute you removed. A detailed tutorial about how to use Query Editor with each data source is beyond the scope of this book, but you should be aware of this important difference compared to other types of data sources.




Figure 5-15: The Add Items dialog box in which you can add measures and attributes from a Google Analytics connection.

After David confirms the items to import, Power BI Desktop sends the query to Google Analytics and imports the result in a new table, as shown in Figure 5-16, in which you can see the content of the table named Website, containing the results of the query to Google Analytics defined so far.

Country	Year	Sessions	New Users	Users
China	2014	58	25	27
China	2015	57	42	43
China	2016	19	6	8
Germany	2013	55	31	35
Germany	2014	420	186	203
Germany	2015	1474	573	588
Germany	2016	525	172	198
United States	2013	365	245	283
United States	2014	4542	1827	1856
United States	2015	9677	3887	4024
United States	2016	2644	1127	1293

TABLE: Website (11 rows)

Figure 5-16: The table Website is the result of the query to Google Analytics.

Note For the purposes of this scenario, our intrepid budgeting manager, David, does not create relationships between the table with data from Google Analytics and the other table in the same data model. However, it is common to create relationships between tables in order to simplify data navigation. You will see more complex data models in Chapter 6.

The key metric that David wants to obtain is the growth in new users that each country/region experienced in 2015. The data available in the Google Analytics service content pack did not have the historical depth required for this analysis; David requires data for both 2014 and 2015. The data imported in Power BI Desktop does not have such a limitation, so it is possible to import such historical depth. However, to create the calculation of the growth percentage, David has to learn the language used by Power BI, which is called DAX.

DAX, which stands for Data Analysis Expressions, was introduced in Power Pivot for Excel in 2010 and is based on the Excel formula language. If you have experience with Excel, you will find many functions that have the same name and syntax as those in your favorite spreadsheet. But, there are also several new concepts and functions that would require a separate book to cover fully. Fortunately, these books exist, such as *The Definitive Guide to DAX*, published by Microsoft Press. You will find a very basic discussion about DAX “measures” (the DAX term used to refer to scripts) in Chapter 6 of this book.

Because David used Power Pivot for Excel in the past, he already knows how to write the measure he needs. So, on the ribbon, on the Home tab, he clicks New Measure and inserts the following DAX measure in the formula bar:

```

New Users Growth =
IF (
    HASONEVALUE ( Website[Year] ),
    DIVIDE (
        SUM ( Website[New Users] ),
        CALCULATE (
            SUM ( Website[New Users] ),
            Website[Year] = VALUES ( Website[Year] ) - 1
        )
    )
)

```

Then, he displays this measure in a separate visualization, under the New Users metric, grouped by country/region and year, as illustrated in Figure 5-17.

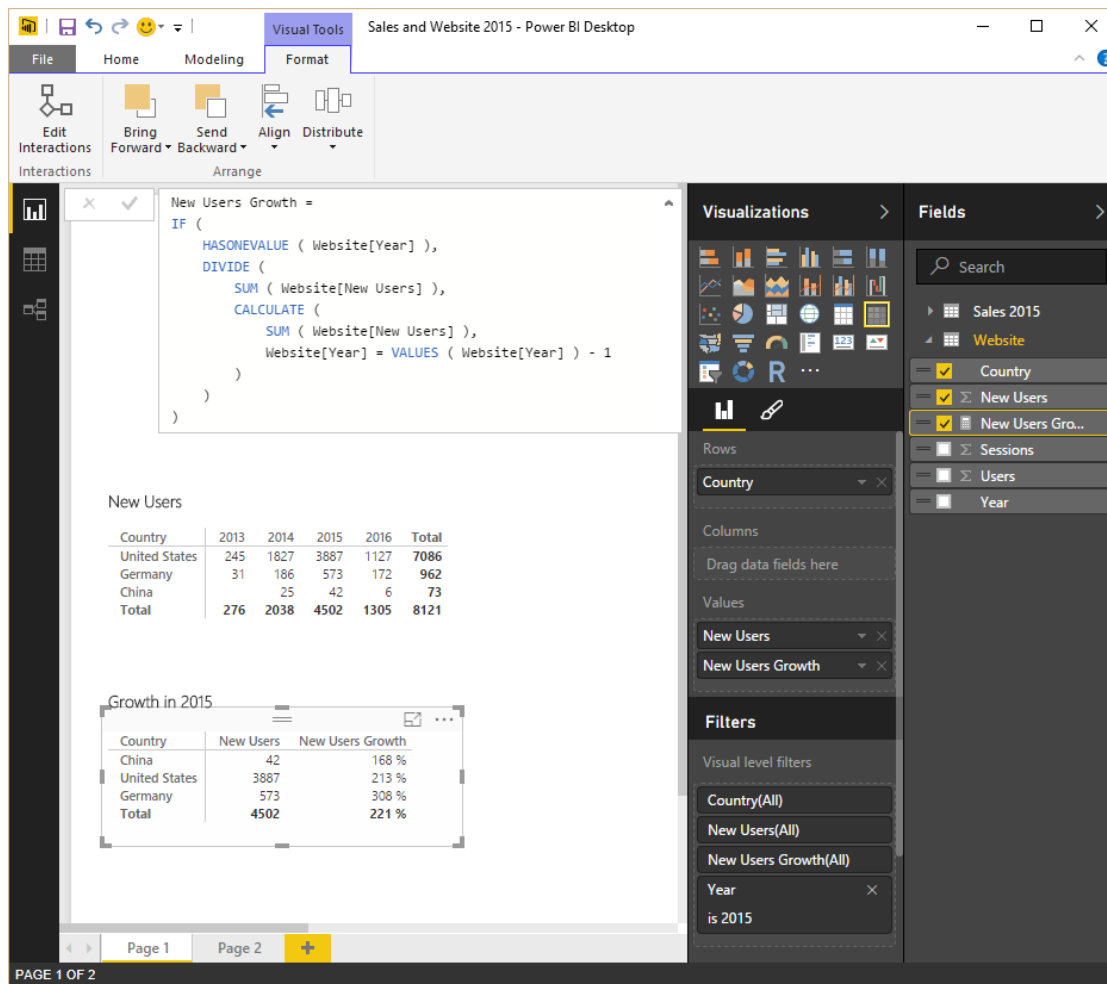


Figure 5-17: Two visualizations displaying the New Users measure and its growth, using data from Google Analytics.

At this point, David has the data he needed from Google Analytics, and he can consolidate that in a single report, together with data coming from other data sources. Chapter 4 shows you how to load data from different data sources, and Chapter 6 shows you how you can combine this data in a single model by using the DAX language. This helps to improve the browsing experience by providing a unique filter for each entity (such as the Country/Region in this report) instead of having similar columns in different tables that filter only the specific table to which they belong.

David can now publish the report he created in Power BI Desktop by using Power BI. By doing that, he is able to pin report content to a dashboard. Figure 5-18 depicts a dashboard built by pinning the two visualizations of the report he created. At this point, David has a dashboard and a report published on Power BI that are based on a custom dataset he created in Power BI Desktop, which gets a particular selection of data from Google Analytics.




Figure 5-18: A dashboard that includes two visualizations displaying data from Google Analytics.

Creating a content pack for your organization

After David created a report in Power BI Desktop using the Google Analytics connector, he wants to share the result of his work with other colleagues in the company, and he wonders what the best tool is in Power BI to do that.

David realizes that the report he created would be a good starting point for deeper insights created by his colleagues. By using the share feature, he is able to share only a dashboard and its underlying reports, but he would like to publish the report based on Google Analytics as a template for reports created by other colleagues. The sharing options you have seen so far do not satisfy David's requirements. He wants other users to be able to customize the reports he made and create new reports based on his work. Sharing a dashboard does not provide such flexibility. Using the group workspace, there would be a single copy of reports and dashboards shared among the group's users, which would not be visible to users who are external to the group. However, David wants to share the results of analysis based on Google Analytics with other users outside the budgeting group. Thus, sharing a dashboard and creating a group workspace are not viable options.

The content pack for an organization is a good solution for David's requirements. This content pack can contain datasets, reports, and dashboards. Users receive a copy of these objects that are automatically synchronized in case a new version of the same content pack is published. If users customize one of these objects, they will work on their own copy of the reports, which will no longer be synchronized with the original one.

To create a content pack, in the upper-right corner, click the Settings button (the small gear icon), and then, on the menu that opens, select Create Content Pack, as shown in Figure 5-19.




Figure 5-19: The Create Content Pack command on the Settings menu.

The Create Content Pack dialog box opens (see Figure 5-20). Here, David provides data about the new content pack. He can choose whether the content pack should be visible to any user in the organization or just to users belonging to specific groups. In this case, David selects the My Entire Organization option. He defines a title for the content pack (Sales And Website 2015, in this example) and a description of its content.

It is important to provide a clear description because this is what other users will read before importing a content pack in their own workspace. It is also possible to customize the content pack by using a specific image, which could be just the company logo or a more customized graphic. Selecting an image is optional; if you do not do that, a default one will be used, instead.

The more important part of the Create Content Pack dialog box is the area where you select the item to publish. There are three lists for all the dashboards, reports, and the datasets you have in your personal workspace. You can select any number of the available entities, even if there are a few constraints. If you select a dashboard, all of the reports and datasets used in the visualizations of that dashboard will be automatically included, too. For example, in Figure 5-20 you can see that the dashboard Sales And Website is selected, and because of that, the report and the dataset named Sales And Website 2015 are automatically selected, as well. You cannot remove the selection of a dataset or a report if it is used in a dashboard. The same is true when you select a report: the underlying dataset is automatically selected, too.

Create content pack

Choose who will have access to this content pack:

- Specific groups My entire organization

Title

Sales and Website 2015

Description

Statistics about new users growth and visits on our website, plus summary of sales until 2015.



Upload an image or company logo
Image size: 45 KB or less, 4:3 aspect ratio, JPG or PNG format

Select items to publish

Dashboards

- Google Analytics
 Sales and Website

Reports

- Google Analytics
 Sales 2015 - Analytics
 Sales and Website 2015

Datasets

- Sales 2015 - Analytics
 Google Analytics
 Sales and Website 2015

The content pack will be available in your organization's content gallery. [Learn more](#)

Publish

Cancel

Figure 5-20: The Create Content Pack dialog box requires you to select objects to publish in a new content pack.

When you click the Publish button, the content pack is published and displayed in the list of the content packs that you can obtain by selecting the View Content Pack item in the Settings menu (refer to Figure 5-19).

Figure 5-21 presents David's list of content packs.

Name	Published To	Date published	Actions
Sales and Website 2015	My Organization	Mar 08, 2016	Edit Delete

Figure 5-21: List of content packs published by the current user.

From this list, David can edit or delete each content pack. If he were to select Edit, he would return to the same configuration window shown in Figure 5-20; however, this time the window would be titled Update Content Pack instead of Create Content Pack.

At this point, other users in his organization are able to consume the content pack that David created. In the next section, you will see how this works and what the difference is between consuming an existing content pack as is and creating a personal copy that can be modified.

Consuming an organizational content pack

In this section, you will see how Wendy can use the content pack that David created.

When she clicks Get Data, Power BI opens the familiar Get Data page, shown in Figure 5-22.




Figure 5-22: The Get Data page in Power BI.

In the Content Pack Library section in Figure 5-22, Wendy clicks Get on the My Organization tile. This opens the list of content packs she can use, along with a tile to create a new content pack, as demonstrated in Figure 5-23. In this scenario, the only content pack available to Wendy is the one David published, Sales And Website 2015.

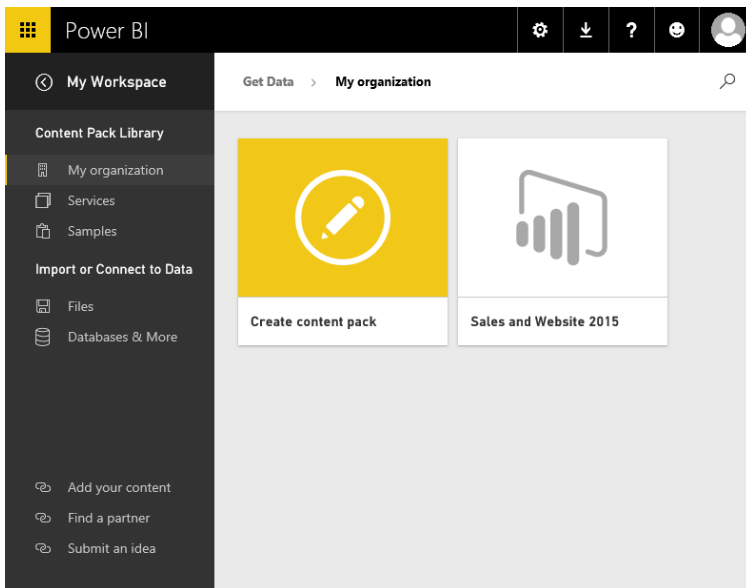


Figure 5-23: The list of content packs available in My Organization.

Wendy clicks the Sales And Website 2015 tile and then sees the information that David previously included when he published the content pack (refer back to Figure 5-20). This information includes a description of the content pack, the name of its publisher, and the time since it was last published, as shown in Figure 5-24.




Figure 5-24: Details about the selected content pack.

When Wendy clicks the Connect button, Power BI imports into her personal workspace the entities included in the content pack. She will find a new dashboard (Sales And Website) has been added to the left pane, as well as a new report and a new dataset (both named Sales And Website 2015), as illustrated in Figure 5-25.




Figure 5-25: Entities included in the content pack that are imported into a personal workspace.

Wendy can navigate in the dashboards and in the reports in view mode without any issue. If the content pack were to be updated by David, the new version of datasets, reports, and dashboards included in the content pack would automatically replace those available to Wendy. However, if She tries to pin something more in the Sales And Website dashboard, or edit the report Sales And Website 2015, or if she clicks on the Sales and Website 2015 dataset to create a new report based on such a dataset, a message will appear (Figure 5-26), asking if she wants to personalize the content pack.




Figure 5-26: The request to personalize a content pack.

The same option is also available as the Personalize button when she clicks the option of a dashboard, report, or dataset that she obtained from a content pack, as demonstrated in Figure 5-27. These options also include the ability to remove or to open the object.




Figure 5-27: The options available in a dashboard obtained from a content pack.

If Wendy removes any object belonging to a content pack, all of the other entities from the same content pack are removed, as well. If she proceeds with the personalization and clicks Save in response to the message shown in Figure 5-26, she creates a copy of the objects in the content pack, and she is able to modify or delete them as she pleases.

In other words, if Wendy does not create a personal copy of the content pack, she will automatically receive any update to that same content pack that David publishes. If she creates a personal copy, any future updates by David will not propagate to her workspace. She will be able to modify the dashboards and the reports in her personal copy of the content pack; however, she will not be able to modify the dataset, because a dataset obtained from a content pack is always a shared dataset, and only the package owner can modify and refresh its content.

To recap, we have this following possible behavior for each object type:

- **Dashboards and Reports** These can be shared from the content pack, or they can be copied from the content pack in a personal copy. All of the dashboards and reports of a content pack are copied to the personal workspace if the user wants a personal copy of the content pack, but those dashboards and reports no longer receive updates from the content pack publisher.
- **Datasets** These are always owned by the content pack publisher, who is the only person who can schedule refresh operations and change other definitions in the dataset.

Updating an organizational content pack

David is the owner of the content pack named Sales And Website 2015. The moment he publishes the content pack on the Power BI service, any ensuing changes to any object included in the content pack on the Power BI service will generate a special notification. For example, if David changes the dashboard by moving the visualizations to different positions, he receives the warning message depicted in Figure 5-28.




Figure 5-28: A warning message after changes are made to an object included in a published content pack.

Note If you edit and then publish a content pack for a data model created with Power BI Desktop, subsequent changes made to the local data model (the .pbix file) do not automatically propagate to the Power BI service and to the published content pack. You do not receive any update in Power BI Desktop. It is up to you to remember that you must publish the .pbix file on the Power BI service to update the data model in the cloud.

If David clicks View Content Packs, in the list of the content packs he published, he sees which ones are affected by one or more changes he made to the reports and dashboards. Figure 5-29 shows that the only content pack he published, named Sales And Website 2015, displays a warning icon next to the name. When David points to the icon, he can see the content of the warning message, specifying that the content pack must be updated in order to show the changes to other users.


Changes have been made to this published content pack. Click Edit to update it so others can see the changes.		Date published	Actions
Sales and Website 2015 	My Organization	Mar 08, 2016	Edit Delete

Figure 5-29: A content pack is marked with a warning icon when it includes unpublished changes.

In practice, any changes made to objects that are a part of a content pack are not automatically published in a new version of the content pack until the owner specifically performs such an update. When David clicks the Edit action, he can publish a new version of the content pack, which will include the current version of the objects (dashboards, reports, and datasets), replacing the version previously published with the same name. Figure 5-30 shows that the Update Content Pack dialog box is identical to that of Create Content Pack, the only difference being that the Publish button is now labeled Update.

Update content pack

Choose who will have access to this content pack:

- Specific groups My entire organization

Title

Sales and Website 2015

Description

Statistics about new users growth and visits on our website, plus summary of sales until 2015.



Upload an image or company logo
Image size: 45 KB or less, 4:3 aspect ratio, JPG or PNG format

Select items to publish

Dashboards

- Google Analytics
 Sales and Website

Reports

- Google Analytics
 Sales 2015 - Analytics
 Sales and Website 2015

Datasets

- Sales 2015 - Analytics
 Google Analytics
 Sales and Website 2015

The content pack will be available in your organization's content gallery. [Learn more](#)

Update

Cancel

Figure 5-30: The Update Content Pack dialog box is the same as the Create Content Pack dialog box, except the Publish button is now labeled Update.

When David clicks Update, he generates a new version of the content pack that overrides the previous one. This action automatically updates all of the objects within the content pack in all the workspaces of any users who have consumed the same content pack without any customizations. This type of consumption of a content pack corresponds to the share feature of a dashboard, which provides a read-only copy of the dashboard that other users can see but not modify. However, using the content pack, this capability is extended to reports and datasets, which can be a part of a content pack regardless of whether they're used in a dashboard.

If a user who consumed the content pack created her own copy of it, a warning message will display in her workspace, notifying her that a new version of the content pack is available, but that is all that will happen, because changes in the consumed content pack are no longer automatically published to the corresponding entities in her content pack. For example, earlier we saw that Wendy created her own copy of the content pack, so when David publishes the new version of the content pack, she does not see any changes applied to the Sales And Website dashboard. However, because she still uses some objects that were originally created from a content pack that now has been updated, she receives the message shown in Figure 5-31 in her Power BI window.




Figure 5-31: The warning message when there is an update to a content pack that has been copied to a personal copy.

Knowing that a new version of the content pack is available, Wendy can decide whether to get data from the same content pack again. If she decides to do so, a new copy of all the objects will be imported into her workspace. Because these objects will have the same name as the ones she previously copied, it would be a good idea to rename objects imported from a content pack when you decide to create your own personal copy; this way, you would not confuse them with the original copy in case you import the same content pack again in the future.

Conclusions

In this chapter, you learned how to consume and create content packs in Power BI. There are different types of content packs, each one with different behaviors available to the user. Moreover, you have also seen that you can create custom datasets from a service when the corresponding service content pack does not provide the data model that you need.

Here are the most important features you learned:

- A content pack contains a set of dashboards, reports, and datasets that a Power BI user can quickly import into his own personal workspace. He can also customize dashboards and reports imported from a content pack, but not a dataset.
- Content packs are available in the Power BI service only.
- A service content pack is published by Microsoft that you can use to connect to a service importing data that populates a set of predefined dashboards and reports. You must provide the credentials required to connect to the service from which you want to extract data.
- You can create a custom dataset in Power BI Desktop only, using a connector corresponding to the service content pack that you cannot customize. Be aware that not all service packs have a corresponding connector in Power BI Desktop, but there are several (such as Google Analytics) that have both.
- Any user can publish an organizational content pack. These contain predefined connections to data sources that cannot be changed by the user who consumes a content pack; only the content pack publisher can modify those connections.
- A user can consume an organizational content pack by just reading its content or by creating a personal copy that she then can modify.
- An organizational content pack publisher can update data. Such changes are automatically propagated to users who consume the content pack in a read-only mode.

Content packs are an important tool to quickly create a set of predefined reports and dashboards based on data coming from an existing external service, or from a dataset created within the organization.

Building a data model

In this chapter, David moves to the next level in Microsoft Power BI usage. We are probably cheating a bit now, but we wanted this book to show you what Power BI can do for you when you master it, not just demonstrate its basic features. To do that, we presume that David—encouraged by the good results so far—spent some time learning the basics of data modeling and the DAX language. Having learned more details about Power BI, he begins again building the budgeting solution, but this time he can trust his better knowledge of the tools.

David loads the sales in the previous years, but unlike the last time, he does not use the view that Karin, the database administrator at Contoso, created for him. Instead, he uses more basic views, on top of the Contoso data warehouse, that provide data in a more fragmented way. There is a table containing information on the stores, one with the sales data, one for the date, and, lastly, a table of the products themselves. Using this information, he builds a first sales analysis project. Finally, he adds the budget information from the Microsoft Excel workbook and writes some DAX code to prepare the dashboards.

We will not discuss in detail all of the formulas and intricacies of the code and the data model; it's not realistic to expect you to learn how to perform these operations by reading one chapter. Our goal is to build the full project together (remember, you can replicate it by using the companion content). If you like the final project, you will probably be better motivated to proceed further with your study and follow David's path in learning data modeling and DAX.

Loading individual tables

In Chapter 6 there is a report named Budget Start that contains the following 4 tables:

- **Products** This table contains information about the products sold by Contoso.
- **Sales** This one contains detailed sales, one row for each individual sale.
- **Stores** This table has information about the stores where the sales were transacted.
- **Date** This is a helper table that contains the calendar. David learned in a Business Intelligence class that such a table is of paramount importance when building a good data model.

Launch **Power BI Desktop** , the following screen will appear:


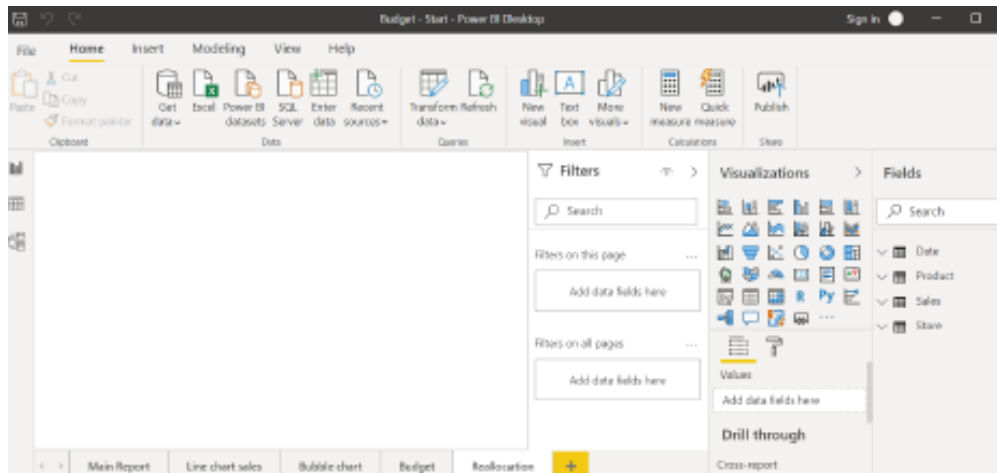


Figure 6-1: Click **Open other reports**

Instead using **Get Data** and loading the tables one by one from the Navigator dialog box, you can simply open the report.

To follow along and catch up to this point, locate and open the companion content file for **Chapter 6: The Report Budget – Start.pbix**. Notice the **Table names on the right** and the **Report page names** at the bottom. During this session we will populate each one of these reports.



On the left side of the screen David clicks the **Model View** icon. He then sees the model illustrated in Figure 6-2 and notices that a **relationship was not created between the Date and Sales tables**. Before working on reports it is important that you check the relationships between the tables that will be used to report on.

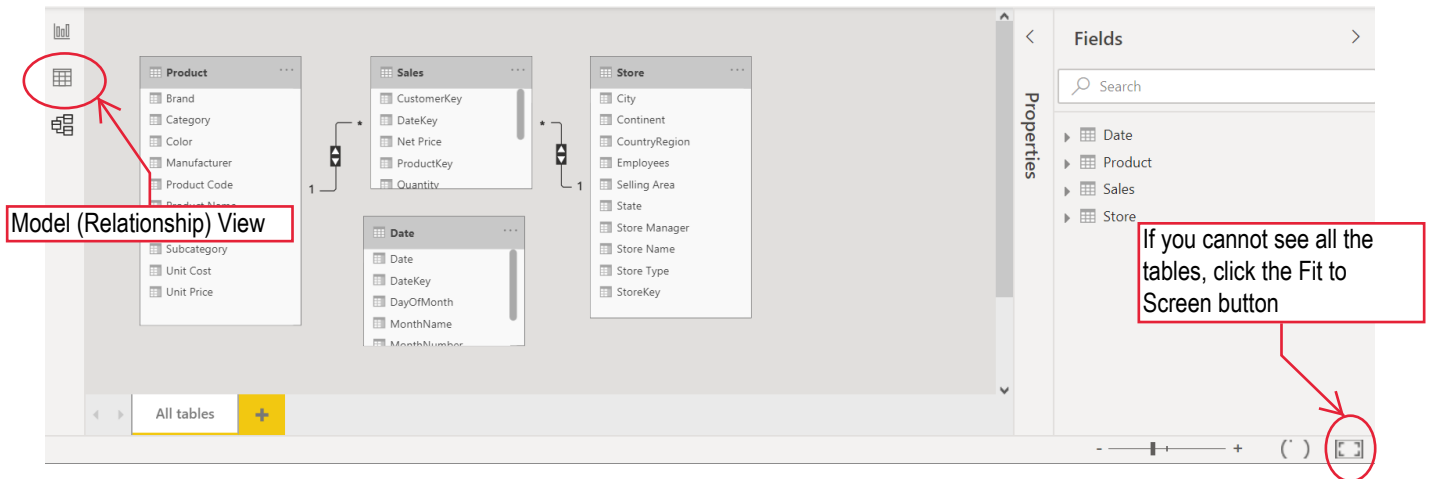


Figure 6-2: The Power BI Desktop relationship detector did not find the relationship between Sales and Date.

This is not an issue; **you can easily create the relationship between Date and Sales**. David uses the mouse to **drag the DateKey field name from the Sales table to DateKey field name in the Date table** to link the two tables. The final data model is presented in Figure 6-3. (Note that your view of the data models might be different; for example, the Date table might be to the left of Sales, not below it.)

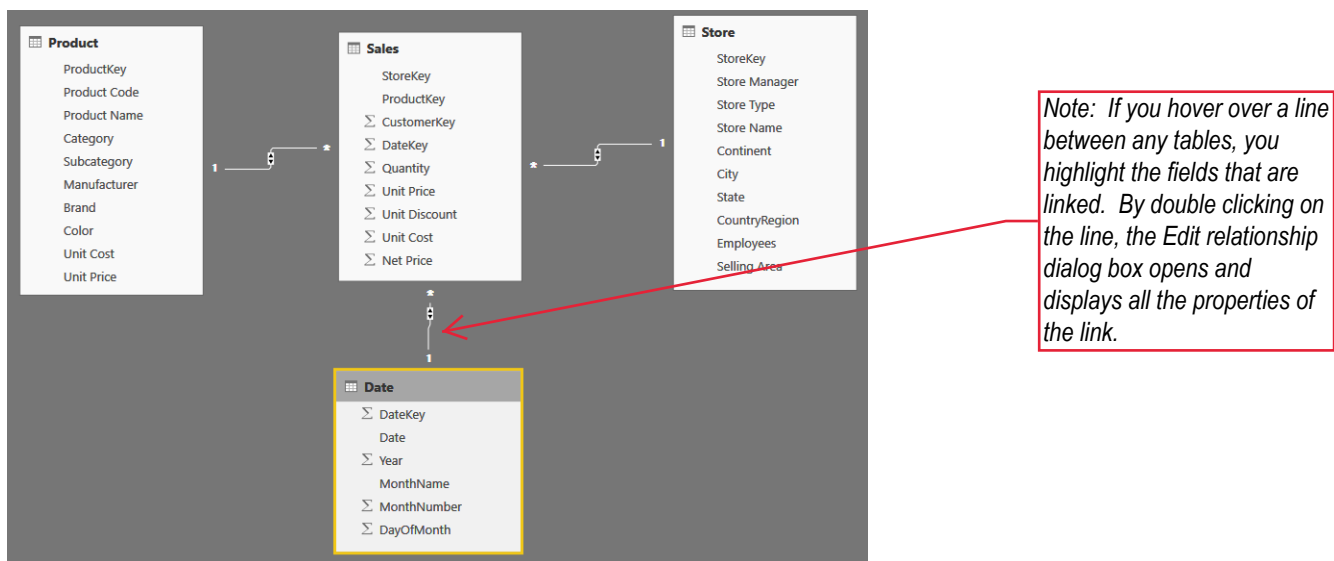


Figure 6-3: The data model structure is a very simple **schema**, with Sales in the middle and the other tables around it.

Implementing measures

The model, as it is, still requires some adjustments. First, when the Report was created, many columns were hidden in various table, so they would not be visible when creating reports. For example, **click the Data view, icon on the left, click the Sales table on the right**. Note the table contains **Quantity and Net Price columns**. (By default, Power BI offers to summarize Net Price by summing the values. In reality, this would be wrong because **summing the price would not take into consideration the quantity sold**.)

Right-click the Sales table, and then, on the shortcut menu that opens, click **Unhide All**. All the columns are now visible. We will make all columns in all the tables visible, right-click on the other three tables and select **Unhide All**.

The default summarization used by Power BI works perfectly well when you have a simple data model. But, as soon as you begin loading data from relational databases for which numbers are not stored in such a way as to be used in Excel workbooks, you need to stop using default summarization and begin writing **DAX measures**, instead. **Measures, in DAX parlance, are scripts that you write using DAX-specific syntax.** By using measures, you can author your own code and produce much more powerful data models. Go back to **Report** view.

David creates a simple measure to compute the **Sales Amount**. In the **Report View**, in the **Fields pane**, David **right-clicks the Sales table** and then clicks **New Measure**. In the **formula bar** above the canvas in the middle pane, he replaces **"Measure ="** with the following code:

Sales Amount = SUMX (Sales, Sales[Quantity] * Sales[Unit Price]) press the Enter key

On the right of the screen under the **Sales** table, you can see the name of the **Measure**, i.e. **Sales Amount** with a small **calculator icon** to its left. This measure alone is already extremely powerful. David decides to use this new measure and create the report below.

David starts by clicking **'Main Report'** page at the bottom of the screen and from the **Visualizations** select **Line chart**. He selects **Sales Amount, MonthName** and **Year**. The line chart is now populated. Clicking in an empty area of the screen, he next selects the **Clustered Column chart** from **Visualizations** and selects **Sales Amount** and **Brand** from the **Fields list**. **Can you complete the other two visualizations?**



Hint: At the top right of each Visualization there are 3 small dots (ellipse), called **More Options**. From here you can **Sort** by a field name, or **Ascending/Descending Order**. Can you **Sort the Sales Amount by Category** and **Sales Amount by Brand** as in the Report.

To the left of the ellipse dots is a button named **Focus Mode**, try it out. Also try **View** from the Ribbon.


Figure 6-4: When the reports is complete, David wants to see the related values for Computers. He clicks the **Computers** bar in the **Sales Amount by Category** chart. He can now see the related data in all the other visualizations. Using new columns available in the data model, the report become more powerful. Try selecting other categorizes and watch the visualization change, return to **Computers** category.

Analysis of sales in previous years becomes more interesting as soon as you have more columns available. **The report shows the relative contribution of different brands to the Computers category**, whereas the **line chart shows the behavior of sales over time**. Different years are highlighted by different line colors.

Creating calculated columns

Having more power typically raises the requirements of the data model. As an example, consider the line chart: showing the sales of the three years with different lines, this might be useful for a comparison of different years; however, if you want to analyze the behavior of sales over the three years, it would be much better to show a single line that spans all of the years.

David selects the 'Line chart sales' page. He then selects a Line chart from the Visualizations. Again he uses the Sales Amount measure he created earlier, and the MonthName field. A line chart displays as shown in Figure 6-5.



Hint: Try some of the Line formatting features available Under Visualizations Format Brush

Figure 6-5: Slicing sales by month shows the total sold over all years for each month.

David needs a column containing both year and month at the same time. Such a column is not available in the original database, so he decides to create a calculated column.

David clicks the Data icon. From the tables on the right he clicks Date. In the table in the Year column he selects the down arrow (filter button) and sees a list of years from 2011 to 2017. He de-selects (Select all) and selects only the years 2013, 2014 and 2015 and Sorts in Ascending order

Let's use this as an opportunity to learn how to use calculated columns. To create a new column in a table, David right clicks the Date table in the Fields list and clicks New Column.

In the Formula bar, David deletes the current content, 'Column =' (which is basically a default column heading name) and types the first formula/measure. He presses the enter key and the Column Name and values populate the column. He clicks on the Date table name again and creates the 2nd new column and types the formula/measure into the formula bar:

Month Year = FORMAT ('Date'[Date], "mmm YY")

Month Year Number = 'Date'[Year] * 100 + MONTH ('Date'[Date])


Breakdown for Month Year Number formula
2015*100 + 9 = 201509

Figure 6-6: The two calculated columns that have been added to the Date table

In the Date table select Month Year. In the Column tools tab, select Sort by Column and select Month Year Number

The **first column contains a shortened version of month and year** (we keep it short, to make it suitable for the line chart), **whereas the second column is used to sort the first one, using the Sort By Column.**

Go to **Report** view to the **Line chart sales page**. If you now **replace the Month Name with the Month Year** as the axis of the line chart, the Visualization is exactly what you want, showing the behavior of sales over three years, 2013, 2014 and 2015 as you can see in Figure 6-7.



Hint: With the line chart visualization selected, you can check out the various formatting features available, or check out the **View** tab's many themes.

Figure 6-7: Using a calculated column for the axis of the line chart leads to the desired visualization.

When building reports, you will typically need a calculated column to make the visualization look perfect. Sometimes the descriptions are too large. In other cases, such as this one, you need a column representing a specific behavior. Power BI is an environment in which you model the data while having the visualization in mind as the final goal.

Improving the report by using measures

When you use calculated columns and measures to perform analyses, you're limited only by your imagination. For example, with a few calculations you can easily build all types of reports.

Next David would like to create an interesting chart with the number of products versus the margin divided by category, and where the size of each bubble represents the amount sold. He clicks the **Bubble chart** page. He clicks the **Scatter chart** in the **Visualizations** and proceeds to create the measures required to complete the **visualization**. He clicks on the **Product** table name under **Fields**, then **right clicks** on **Product** and select **New Measure** and **create the following measures**, which will be needed to build the report.


NumOfProducts = COUNTROWS ('Product')
Gross Margin = SUMX (Sales, Sales[Quantity] * (Sales[Unit Price] - Sales[Unit Cost]))

The **NumOfProducts** measure counts the number of products and gives an idea of how many articles are in the portfolio, whereas **Gross Margin** computes the gross margin of sales by subtracting the cost from the unit price before multiplying that value by the quantity.



Now that the measures have been created David can proceed to create the interesting chart he had in mind. He checks that the Scatter chart is still selected then from the **Fields** list finds and ticks - **Gross Margin, NumOfProducts, Sales Amount and Subcategory**. As he ticks off each field name the Scatter chart start to populate with bubbles as below. It appears very overcrowded as every subcategory item is included as below.



He decides to create a second visualization so he can narrow down the data, in the scatter chart, so he clicks into an empty area in the Central pane and from the **Visualizations** selects a **Clustered bar chart** and ticks off **Gross Margin and Brand**, which he then **Sorts by Brand and Ascending order** and gets the following results.



In the **Clustered Bar Chart** David clicks the the bar for **Brand - Litware** and the Scatter chart updates to reflect the changes as below. and the scatter chart is less cluttered.



Hint: The final report on the left shows the **Litware** brand highlighted and the Scatter chart formatted. All the formatting features in the Scatter chart are available under the **Format brush icon** under **Visualizations**, and under **Data Colors**. Experiment with **color**, **transparencies** add labels etc.

Figure 6-8: The final bubble chart shows a large amount of information in a single chart, and they are gorgeous when used with visual interaction and automatic filtering.

Note As you have seen so far, we are not explaining how to write the DAX code. The goal of this chapter is not to teach DAX; a short chapter in a short book would not properly address the complexity of the language. Our goal here is to show you what you can do as soon as you begin learning the basics of data modeling and DAX coding. Incredible analytical power awaits you when you complete your journey learning DAX, so hurry and start learning it today. In the meantime, we will move on with more complex DAX code and present some more scenarios that you can solve by writing simple formulas.

Integrating budget information

So far, David is excited about the power of his analytical tool; so much so, in fact, that he's forgotten that the task is about budgeting, not sales analysis. This is one of the major drawbacks of using Power BI: it is so much fun to dive into data and analyze it that you might become lost in evocative reports. Now it's time to get back to business and integrate the budget information. Before moving on, click the **Budget Page tab** at the bottom of the screen to open it. It is currently blank.

Loading the budget information from Excel is straightforward, it is located in Chapter 4, and David clicks the **Model** tab, on the left of the screen, then clicks **Get data** and finds the Excel file, **2016 Budget**, he opens it. When the **Navigator** dialog box opens, he clicks **Budget2016**, and clicks **Load**, but a problem arises as soon as he looks at the data model. The new table containing the Budget does not have any relationship with the other tables, as you can see in Figure 6-9.



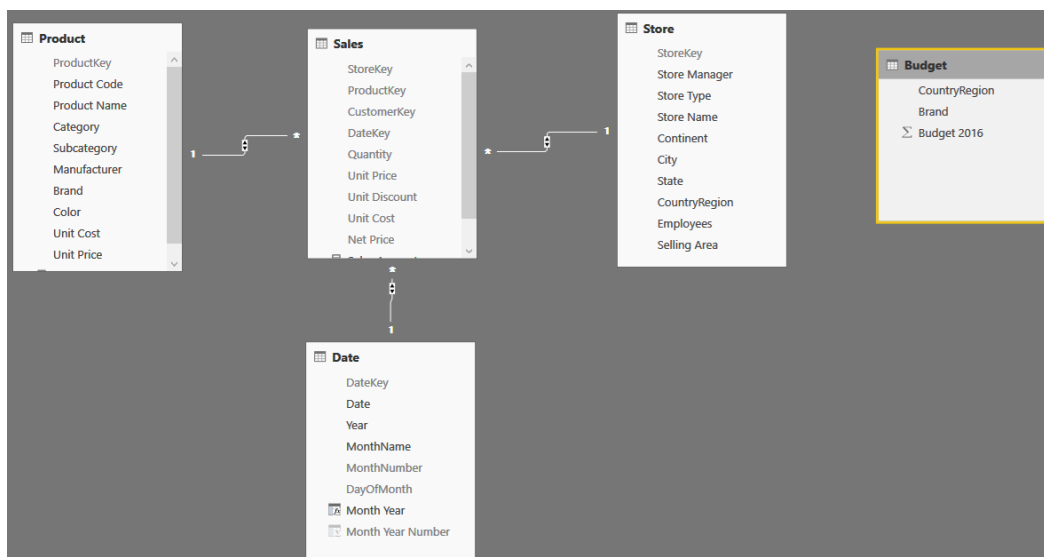


Figure 6-9: The Budget table does not have any relationships with the other tables in the model.

This time, it's not an issue of Power BI failing to detect the relationship, which earlier David was able to correct by using a simple drag-and-drop technique to establish a relationship. In this case, the relationship cannot be created this way. **When he tries to drag CountryRegion from the Budget2016 table to the Store table, the Create Relationship dialog box opens with a warning message at the bottom, as in Figure 6-10 below. Click Cancel**




Figure 6-10: Trying to create a relationship between the Budget and Store tables lead to this error.

But, before solving the issue, it's worth taking a few moments to understand it better.

You can **create a relationship between two tables if the column you use to create the relationship is a key in the destination table**. You can **create a relationship between the Sales and Date tables based on the DateKey column because DateKey has a different value for each row in Date**. Having a different value for each row is the requisite for a column to be a key. **In fact, when you have a given date, you can uniquely identify the entire row in Date. In the model with Budget, CountryRegion is neither a key in the Budget table, nor in Store. Thus, you cannot create such a relationship.**

There are multiple ways to solve this problem, based on the model or on an advanced usage of the DAX language. The solution based on the model is somewhat easier to learn and, by the way, it is the way suggested by the error message.

If you create a table containing all the possible values of CountryRegion, CountryRegion becomes a key for that table. At that point, you are able to create the relationships between Budget and Store.

You can build a table containing the possible values of CountryRegion by using Query Editor, as you did in Chapter 4, or by adopting a new technique: **a calculated table. Calculated tables are tables computed using the DAX language** that you can store in the model and use as any other table.

To resolve the problem, David needs to create a **calculated table**, he clicks the **Report** Icon then on the Power BI ribbon, on the **Modeling** tab, clicks **New Table**, in the formula bar, he deletes the 'Table =' and types the following **DAX** (Data Analysis Expressions) expression. He presses **Enter** when finished.

```
CountryRegions =
SUMMARIZE (
    UNION (
        DISTINCT ( Budget2016[CountryRegion] ), DISTINCT ( Store[CountryRegion] )
    ),
    [CountryRegion]
)
```

SUMMARIZE Function (DAX) Returns a summary table

UNION - Columns are combined

DISTINCT - Duplicate values are removed and only unique values are returned

Clicking on the **Data** icon and he selects **CountryRegions**, from the **Fields** list on the right, will display the **DAX expression** and the resulting table as below, where each CountryRegion is entered once.



Figure 6-11: You can populate a calculated table with the distinct values of CountryRegion.

Let's return to see how David is faring. To create the table, he takes the **distinct values of both the Budget[CountryRegion] and Store[CountryRegion] columns**, forms a union with the partial results, and finally the **Summarize function returns a summary table of the CountryRegion column**. In this way, all the possible values will be represented in the resulting table, either referenced by the Budget or Store tables.

To check the results he selects the **Model** icon on the left of the screen, and can see the new table name on the far right.

He needs to create 2 relationships, the first relationship between **Store, CountryRegion** and **CountryRegions, CountryRegion**, and the second one between **Budget2016 CountryRegion** and **CountryRegions, CountryRegion**, completing the model. The table is a technical table, which is useful only to propagate the filter from Store to the Budget. Figure 6-12 presents the data model with the new table in place.

Again David uses the mouse to drag the fields between the tables. This time it all goes smooth and no error message displays. To be sure, he decides to check the relationship properties between tables. He double clicks on the link lines he just created and the **Edit relationship** dialog box displays showing the relevant tables and the type of relationship. He clicks **Cancel**. The model is looking good now as on the next page.

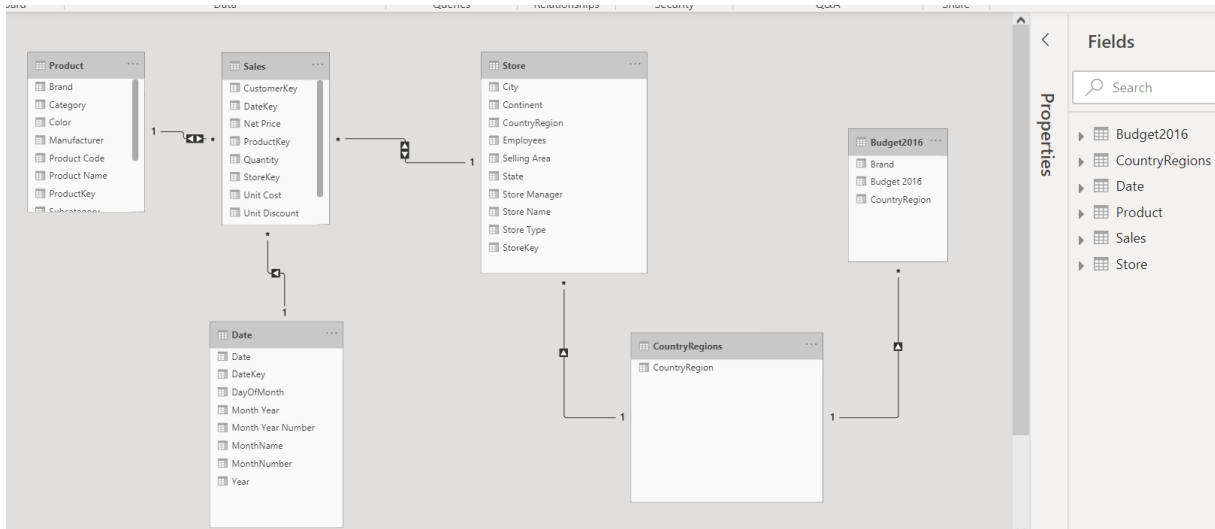


Figure 6-12: The CountryRegions table acts as an **intermediate** between **Store** and **Budget**.

From a technical point of view, David has created a relationship between the Budget and Store tables, using **CountryRegions as the bridge table between them**. To test that the model works well, he returns to the **Report** view and returns to the **Budget page**. He right-clicks on the **Budget2016 table name** under **Fields** and he creates the following simple **New Measure** that will **return the total of the Budget 2016 values**, contained in a single column:

$$\text{Budget Amount} = \text{SUM} (\text{Budget2016}[\text{Budget 2016}])$$

Select the **Table** visualization and add the **CountryRegion from the CountryRegions table** and **Budget Amount** from the **Budget2016 table** and from the **Sales table Sales Amount**. You can project this measure on a simple report containing **CountryRegion**, the **value of budget**, and the **value of sales**.

CountryRegion	Budget Amount	Sales Amount
China	418,500.00	1,416,712.00
United States	395,350.00	1,419,755.00
Germany	327,500.00	1,070,649.00
Total	1,141,350.00	3,907,116.00

The **Budget Amount**, the measure you created above shows the total for each of the 3 countries but shows the overall sales amount values for every year for each country, as in Figure 6-13. **Note: you can change the order of the columns under Visualizations, under Values**

Figure 6-13: The data model, created using the Budget Amount measure

The first problem, here, is that the report is not showing meaningful numbers. (You can check this out in the original Excel Sales table) This happened because no filter set for the year, therefore **it is accumulating sales over all available years and comparing them with the Budget2016 table, which contains forecasts for 2016 only**. David decides to solve the issue by **creating the following measure that computes the sales amount for only 2015**, using the same method above, he right clicks on the **Sales table** and selects **New Measure**, and types the measure as below:

$$\text{Sales 2015} = \text{CALCULATE} ([\text{Sales Amount}], 'Date'[Year] = 2015)$$

David replaces the **Sales Amount** with **Sales 2015** in the report, the numbers can be now be compared correctly as demonstrated in Figure 6-14.

CountryRegion	Budget Amount	Sales 2015
China	418,500.00	399,388.00
United States	395,350.00	324,186.00
Germany	327,500.00	322,863.00
Total	1,141,350.00	1,046,439.00

Figure 6-14: Using Sales 2015 in the report makes the numbers comparable.

The CALCULATE function, evaluates an expression, as an argument, in a context that is modified by the specified filters.

David will now create a new table for the other column available in Budget, which is **Brand**, to further analyze the data. He clicks an empty area in the Central pane, deselecting the current table visualization. He select the **Modeling** tab, selects **New Table**, and types the following **DAX code**, into the Formula Bar (first deleting Table =) this formula will obtain a list of all the possible brands:

```
Brands =
    SUMMARIZE (
        UNION ( DISTINCT ( Budget2016[Brand] ), DISTINCT ( Product[Brand] )
        ),
        [Brand]
    )
```

SUMMARIZE - Returns a summary table for the requested totals

UNION - Creates a union (join) table from a pair of tables.

DISTINCT - Returns a table by removing duplicate rows from another table or expression

The Brands table appears on the right under the Fields list. So far, so good. David clicks the **Model** icon. He create relationships between **Brand** in the **Product** table and **Brand** in the **Brands** Table and **Brand** in the **Budget2016** table to **Brand** in the **Brands** table. He ends up with a model like the one depicted in Figure 6-15.

The next step hides a problem that—again—requires a bit of theory to be explained.

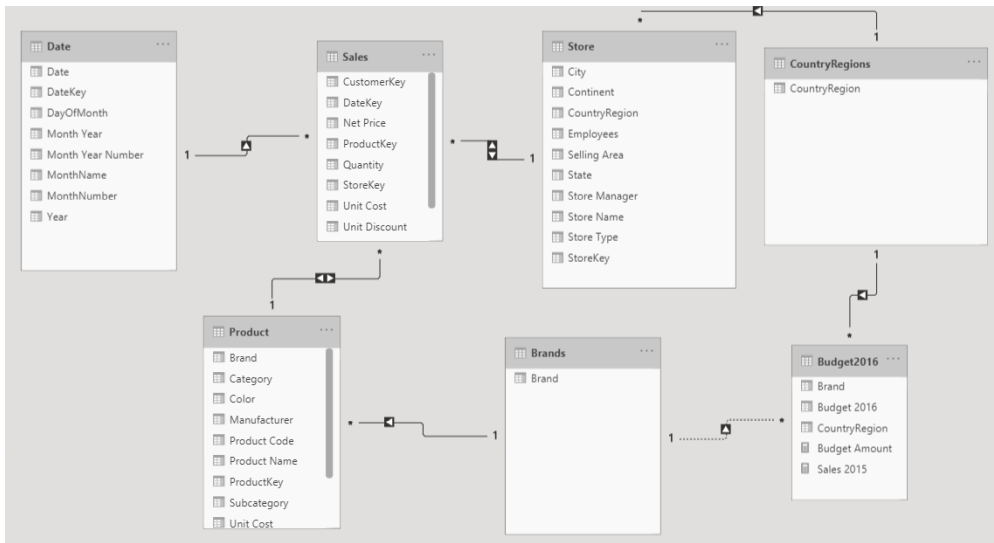


Figure 6-15: The tables are now all linked

To test the links, David clicks on an empty area in the central pane, then clicks the **Report** icon and creates a new **Matrix** (a cross-tab type table and groups and subtotals data) visualization using **Brand** from the **Brands** table and **CountryRegion** from the **Budget2016** table and **Sales 2015** from the **Sales** table

The **result is wrong**, as demonstrated in Figure 6-16 on the following page.

Brand	China	Germany	United States	Total
A. Datum	\$46,051.3	\$46,051.3	\$46,051.3	\$46,051.3
Adventure Works	\$62,641.73	\$62,641.73	\$62,641.73	\$62,641.73
Contoso	\$228,978.33	\$228,978.33	\$228,978.33	\$228,978.33
Fabrikam	\$187,597.45	\$187,597.45	\$187,597.45	\$187,597.45
Litware	\$161,097.75	\$161,097.75	\$161,097.75	\$161,097.75
Northwind Traders	\$15,901.55	\$15,901.55	\$15,901.55	\$15,901.55
Proseware	\$112,310.71	\$112,310.71	\$112,310.71	\$112,310.71
Southridge Video	\$61,924.16	\$61,924.16	\$61,924.16	\$61,924.16
Tailspin Toys	\$16,653.67	\$16,653.67	\$16,653.67	\$16,653.67
The Phone Company	\$51,919	\$51,919	\$51,919	\$51,919
Wide World Importers	\$101,363.4	\$101,363.4	\$101,363.4	\$101,363.4
Total	\$1,046,439.05	\$1,046,439.05	\$1,046,439.05	\$1,046,439.05

Figure 6-16: Slicing by Brand does not produce a meaningful result (all of the values are repeated) because the underlying relationship is of the wrong type, or may be inactive.

An **ambiguous model** is a data model within which there are **multiple paths linking two tables** because all of the relationships are set as bidirectional (that is, the filter applies in both directions).

(Note that you can tell when a relationship is bidirectional because there are two small arrows on the connector line, facing both directions.) So, where is the ambiguity? There are many here. For example, if you start from Product (see Figure 6-15), you can reach Budget following the bottom chain **Product/Brands/Budget** or the upper chain **Product/Sales/Store/CountryRegion/Budget**. If all of the relationships remained active, both would be legal paths, and you would end up with ambiguity.

You can solve ambiguity in most cases by just preventing the path from being traversed, but still maintaining the model features. For example, in the model examined, there is no need to make **Sales filter Store and Sales filter Product**. It is enough that the opposite direction works in both cases; that is, you can have both the **Store and Product tables filter Sales**.

David goes back to the **Model** icon and checks out the joins. When he **double-clicks the join between Sales and Store**, the **Edit Relationship** dialog box opens, as shown in Figure 6-17.

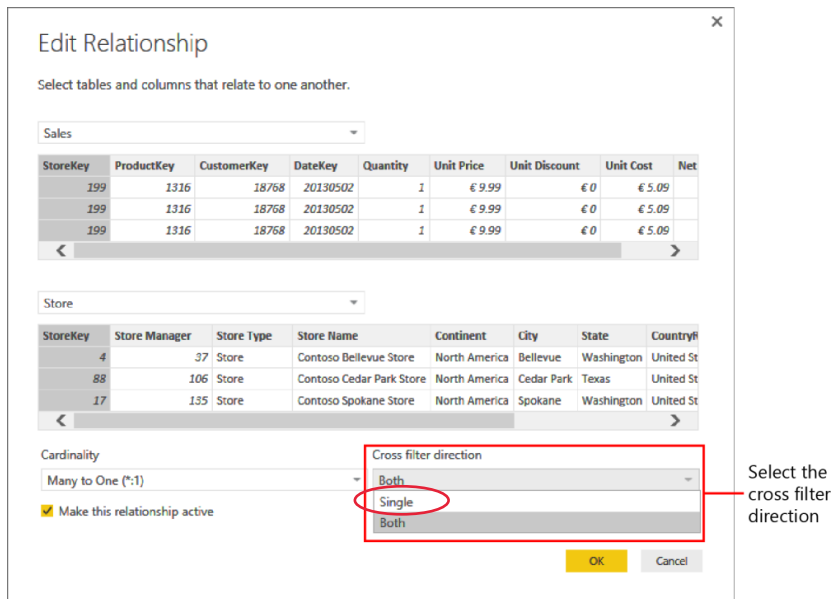


Figure 6-17: In the Edit Relationship dialog box, you can configure many properties of a relationship.

David sees that the **Cross filter direction** has defaulted to **Both**. He clicks the down arrow to the right of **Both** and sets it to **Single**, this will disable the bidirectional filtering, He clicks **Ok** then double-clicks the **join line** between **Sales** and **Product**, which also needs to be made **Single**.

The final model is represented in Figure 6-18.




Figure 6-18: Disabling bidirectional filtering on most relationships removes the ambiguity. After you remove ambiguity from the model, the model works fine.

David selects the **Report** icon and deletes the table with the repeating data. He tests the new links by creating a **Matrix** table, and selects both **Brand**, **CountryRegion** and **Budget 2016** from the **Budget2016** table. Now the budget is filtering correctly, as illustrated in Figure 6-19.

PLEASE NOTE WELL DESIGNED DATABASES WITH SOUND TABLE STRUCTURE DO NOT CAUSE THESE PROBLEMS, THEY WERE DELIBERATELY BUILT INTO THIS SCENARIO TO LET YOU SEE THE RESULTS OF BADLY DESIGNED DATABASE POORLY CONSTRUCT TABLES AND METHODS TO CHECK AND CORRECT THEM.



Figure 6-19: With the correct model, slicing is performed the right way.

The granularity of a DAX formula is described as being the level of detail that the formula uses to compute those result in a report, and the accuracy of those results.

Reallocating the budget

Budget numbers are **correct as long as you slice by brand or by country/region**, which is the granularity at which the budget is defined. However, if you add a column that is not a part of the Budget table (for example, if you sliced by **Country/Region and Color**), the result would be wrong. Figure 6-20 shows that values for **Black and Silver colors in United States have the same value as the grand total of all the colors in United States**.

David decides to add another **Matrix** table, this time using the **Reallocation** page.

He selects the **Matrix** icon from the **Visualizations** pane. From the **Product** table he ticks **Color** and **Category**. A list of colors will display as column labels and the categories will appear as row headings.

From the **Sales** table he ticks **Sales 2015**, and the values appear, in a table as below:

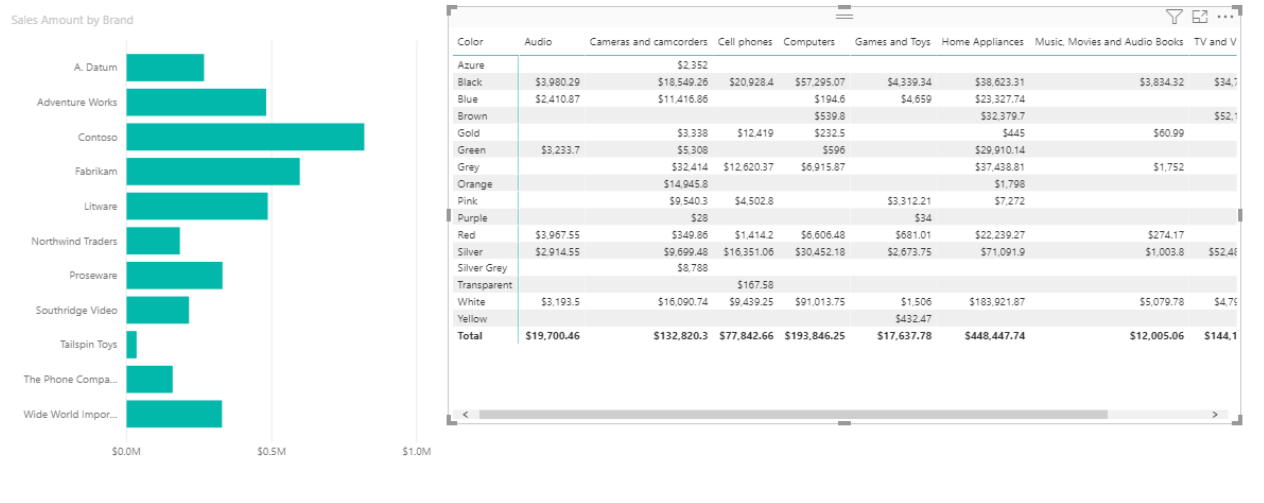


Figure 6-20: Try selecting a bar from the Brand chart and watch the changes in the Matrix table.

Conclusions

As we said in the introduction, in this chapter we cheated a bit. We did not want to show you another step-by-step guide to implement another dashboard. Instead, we wanted to give you a sneak preview of the capabilities of Power BI Desktop when you uncover the most advanced tools, namely:

- **Building a model** When you begin loading the raw tables from the SQL Server database, instead of predefined queries with aggregated values, you can perform much more powerful analyses. At the same time, you are responsible for handling the data model by yourself. Power BI Desktop offers you all the tools required to build a complex data model.
- **The DAX language** DAX is your best friend in the process of analyzing data. **In this chapter, we used it to create calculated columns, measures, and calculated tables. This book is not the proper venue for explaining how DAX works; that would fill an entire book by itself.** If you are interested in learning more about DAX, check out our book *The Definitive Guide to DAX* (Microsoft Press, 2015).
- **Building columns for specific charts** Sometimes, you need a column for an individual chart. There is nothing wrong with doing this; you can just build it.

By using some basic skills, you can take Power BI from a simple reporting tool to what it really is: an extremely powerful modeling tool with which you can build gorgeous analyses on top of your data.

DAX Glossary

Analytic query

Power BI visuals query a data model by using an analytic query. An analytic query strives to reduce potentially large data volumes and model complexities using three distinct phases: Filter, group and summarize. An analytic query is created automatically when fields are assigned to the wells of report visuals. Report authors can control the behavior of field assignments by renaming fields, modifying the summarization technique, or disabling summarization to achieve grouping. At report design time, filters can be added to the report, a report page, or a visual. In reading view, filters can be modified in the Filters pane, or by interactions with slicers and other visuals (cross-filtering).

BLANK

DAX defines the absence of a value as BLANK. It's the equivalent of SQL NULL, but it doesn't behave exactly the same. It's more closely aligned to Excel and how it defines an empty cell. BLANK is evaluated as zero or an empty string when combined with other operations. For example, $BLANK + 20 = 20$. Always use capital letters; the plural is BLANKs, with a lowercase "s".

Calculated column

A model calculation used to add a column to a tabular model by writing a DAX formula. The formula must return a scalar value, and it's evaluated for each row in the table. A calculated column can be added to an Import or DirectQuery storage mode table.

Calculated measure

In tabular modeling, there's no such concept as a calculated measure. Use measure instead. The word calculated is used to describe calculated tables and calculated columns. It distinguishes them from tables and columns that originate from Power Query. Power Query doesn't have the concept of a measure.

Calculated table

A model calculation used to add a table to a tabular model by writing a DAX formula. The formula must return a table object. It results in a table that uses Import storage mode.

Calculation

A deliberate process that transforms one or more inputs into one or more results. In a tabular data model, a calculation can be a model object; either a calculated table, calculated column, or measure.

Context

Describes the environment in which a DAX formula is evaluated. There are two types of context: Row context and filter context. Row context represents the "current row", and is used to evaluate calculated column formulas and expressions used by table iterators. Filter context is used to evaluate measures, and it represents filters applied directly to model columns and filters propagated by model relationships.

Cube

See Multidimensional model.

Data model

A data resource that's specifically prepared for reporting and analytics. It allows report users to browse and explore data in a simple and intuitive way. Importantly, it delivers high performance query results, even over large data volumes. It can integrate data from multiple sources and use calculations to transform data. It can enforce row-level permission to ensure different users have access to different data. Sometimes it's called a Semantic model or just a Model.

Data modeler

A skilled data professional who creates data models. In self-service BI (SSBI), they can be called business analysts. In corporate BI, they can be called BI developers.

DAX

Data Analysis Expressions (DAX) language is a formula language for Power Pivot in Excel, Power BI, Azure Analysis Services, and tabular modeling in SQL Server Analysis Services. You can also use DAX to add data model calculations and define row-level security (RLS) rules.

Dynamic security

When row-level security (RLS) rules are enforced by using the identity of the report user. Rules filter model tables by using the user's account name, which can be done with the USERNAME or USERPRINCIPALNAME functions. See Row-level security.

Expression

A unit of DAX logic that's evaluated and returns a result. Expressions can declare variables in which case they're assigned a sub-expression and must include a RETURN statement that outputs a final expression. Expressions are constructed by using model objects (tables, columns, or measures), functions, operators, or constants.

Field

Data model resource presented in the Fields pane. Fields are used to configure report filters and visuals. Fields consist of model columns, hierarchy levels, and measures.

Formula

One or more DAX expressions used to define a model calculation. Inner expressions are called sub-expressions. Plural is formulas.

Function

DAX functions have arguments that allow passing in parameters. Formulas can use many function calls, possibly nesting functions within other functions. In a formula, function names must be followed by parentheses. Within the parentheses, parameters are passed in.

Implicit measure

An automatically generated calculation achieved by configuring a Power BI visual to summarize column values. Numeric columns support the greatest range of summarization, including: Sum, Average, Minimum, Maximum, Count (Distinct), Count, Standard deviation, Variance, or Median. Columns of other data types can be summarized, too. Text columns can be summarized by using: First (alphabetically), Last (alphabetically), Count (Distinct), or Count. Date columns can be summarized by using: Earliest, Latest, Count (Distinct), or Count. Boolean columns can be summarized by using: Count (Distinct), or Count.

Iterator function

A DAX function that enumerates all rows of a given table and evaluate a given expression for each row. It provides flexibility and control over how model calculations summarize data.

MDX

Multidimensional Expressions (MDX) language is a formula language for SQL Server Analysis Services multidimensional models (also known as cubes). MDX can be used to query tabular models, however it can't define implicit measures. It can only query measures that are already defined in the model.

Measure

A calculation that achieves summarization. Measures are either implicit or explicit. An explicit measure is a calculation added to a tabular data model by writing a DAX formula. A measure formula must return a scalar value. In the Fields pane, explicit measures are adorned with a calculator icon. Explicit measures are required when the model is queried by using Multidimensional Expressions (MDX), as is the case when using Analyze in Excel. An explicit measure is commonly just called a measure.

Measure group

A model table that contains at least one measure, and has no hierarchies or visible columns. In the Fields pane, each measure group is adorned with a multi-calculator icon. Measure groups are listed together at the top of the Fields pane, and sorted alphabetically by name.

Model

See Data model.

Modeler

See Data modeler.

Model calculation

A named formula that's used to add a calculated table, calculated column, or measure to a tabular data model. Its structure is <NAME> = <FORMULA>. Most calculations are added by data modelers in Power BI Desktop, but measures can also be added to a live connection report. See Report measures.

Multidimensional model

A data model developed for SQL Server Analysis Services (multidimensional mode). It comprises dimensions and measures. It's often just called a cube.

Quick measures

A feature in Power BI Desktop that eliminates the need to write DAX formulas for commonly defined measures. Quick measures include average per category, rank, and difference from baseline.

Report measures

Also called report-level measures. They're added to a live connection report in Power BI Desktop by writing a DAX formula, but only for connections to Power BI models or Analysis Services tabular models.

Row-level security

Also called RLS. Design technique to restrict access to subsets of data for specific users. In a tabular model, it's achieved by creating model roles. Roles have rules, which are DAX expressions to filter table rows.

Semantic model

See Data model.

Summarization

An operation applied to the values of a column. See measure.

Tabular cube

There's no such concept as a tabular cube. Use Tabular model instead.

Tabular model

A data model developed in Power Pivot in Excel, Power BI, Azure Analysis Services, or SQL Server Analysis Services (tabular mode).

Time intelligence

Time intelligence relates to calculations over time, like year-to-date (YTD).

Time intelligence function

DAX includes many time intelligence functions. Each time intelligence function achieves its result by modifying the filter context for date filters. Example functions: TOTALYTD and SAMEPERIODLASTYEAR.

Value, values

Data to be visualized.

What-if parameter

A Power BI Desktop feature that provides the ability to accept user input through slicers. Each parameter creates a single-column calculated table and a measure that returns a single-selected value. The measure can be used in model calculations to respond to the user's input.

Improving Power BI reports

The previous chapters introduce many features of Power BI that are related to data modeling, publishing, and sharing. The focus in those chapters is on the data and the numbers, not on the presentation. Our friend David took advantage of these features to create a solution that uses Power BI to support a collaborative effort in the yearly budgeting process for his company, Contoso. Those chapters intentionally do not focus on the presentation layer or the visual options of Power BI, but now it is time to improve the reports and dashboards in your models.

This chapter is dedicated to the visual presentation capabilities of Power BI. We will use different dashboards with different datasets and requirements to show you the different scenarios and tools available. For this reason, we're going to look beyond David's requirements for his solution—we need a broader scope now. So, you will learn how to choose between the built-in visualizations available in Power BI, how to use custom visuals, how to use DAX to solve common reporting issues, and, finally, what the correct approach is for designing high-density reports.

This chapter is neither a step-by-step tutorial for using the visual editor in Power BI, nor a complete visual design patterns guide covering all the possible types of standard and custom visualizations. Its

goal is to show you the available options in Power BI and act as an initial guide to choosing visualization solutions, depending on your requirements. We will do that by providing different practical examples that are available in the companion content, so that you can open the files and analyze the data in more detail, reviewing all the properties we used. In these pages, we provide commentary, explaining the reason for certain design choices; you should be able to apply the same principles when designing your own reports.

Choosing the right visualizations

You have several standard visualizations available in Power BI. In addition, you will see that you can extend this set by using custom visualizations. But, before doing that, you need to know what you can and cannot do using the standard components, which you can see in Figure 7-1.

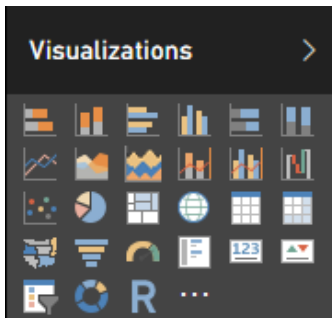


Figure 7-1: The standard visualizations available in Power BI.

The list of standard visuals includes 27 components in Power BI that are available as of this writing, but this might expand in future versions. Here are descriptions of the standard components:

- **Stacked bar chart** Use this when you want to compare different values of the same measure, side by side, or when you need to display different measures that are a part of the same whole. The bars are horizontally oriented rows.
- **Stacked column chart** The same as a stacked bar chart, but vertically oriented.
- **Clustered bar chart** Similar to the stacked bar chart, but instead of comparing different measures within the same bar, with a clustered bar chart you can compare different measures side by side.
- **Clustered column chart** The same as a clustered bar chart, but vertically oriented.
- **100% Stacked bar chart** Similar to the stacked bar chart, but with each measure using a slice of each bar, which always corresponds to the entire width available (100%).
- **100% Stacked column chart** The same as the 100% stacked column chart, but vertically oriented.
- **Line chart** Use this to display the trend of some measures over time. Usually the y-axis has a range that does not include zero.
- **Area chart** Similar to the line chart, use this when you want to display cumulative data rather than sequences of points. Usually, the y-axis range begins at zero, and there is only one measure. This looks like a line chart wherein the areas are filled with layers of colors.
- **Stacked area chart** Similar to the area chart, but with each measure cumulated to the others.

- **Line and stacked column chart** Use this when you need to display measures with different scales, such as currency and percentage, or different value ranges.
- **Line and clustered column chart** The same as the line and stacked column chart, but using clustered columns instead of stacked columns.
- **Waterfall chart** Use this to display cumulative data, highlighting for each value its positive or negative value. The initial and final value columns usually start on the horizontal axis, with color-coded floating columns between them, making it look like a waterfall or bridge.
- **Scatter chart** Use this when you want to show possible correlations between two measures.
- **Pie chart** Use this to display the distribution of values of one or more measures. The values appear as pieces of the pie, with the larger values taking up larger slices. However, using pie charts is not a best practice.
- **Treemap** Similar to a pie chart, but using a rather different graphical representation, wherein the values are represented by colored rectangles on a page. It could be an alternative to a pie chart, but it is equally unreadable when you have many elements in it.
- **Map** Use this to display geographical data with variable-sized circular shapes on Bing maps.
- **Table** Use this to display data in a textual form as a simple table, where every attribute and every measure is a single column in the result.
- **Matrix** This extends the table, making it possible to group the measures by rows and columns.
- **Filled map** Similar to the map, but the data is represented by colored overlay areas.
- **Funnel** Similar to the stacked bar chart, but with a single measure and a different graphical representation, wherein the rows are stacked in order, which makes the chart look like a funnel.
- **Gauge** Use this to display a single measure against a goal. This chart resembles a gauge style that is common in cars.
- **Multi-row card** Use this to display different measures and attributes for each instance of an entity, each placed on a different colored and graphical card.
- **Card** Use this to display a single numerical value of a measure textually, placed on a colored and graphical card.
- **KPI** Use this to display a single value with a trend line chart in the background, highlighting its performance with colors.
- **Slicer** Use this to filter one or more charts by selecting values of an attribute.
- **Donut chart** Similar to the pie chart, but with a donut or tire-like graphical representation. However, using donut charts is not a best practice.
- **R script visual** Use this to display charts generated by R-language code.

The first design principle is simple: just because you have many components, there is no reason for you to use all of them. In a single report, the presence of many different types of components can be confusing. Thus, do not put too many different visualizations in the same report without a good reason. Moreover, the default properties of a component are not necessarily the right choice for your report, so you should consider modifying their value to better display your data. You will see several examples of these principles applied to our first report example.

For instance, consider a dashboard displaying Sales for Contoso. By choosing the right visualization and setting the right colors, you can obtain a good presentation of your data, focusing your viewer’s attention on the data instead of the visualization itself. You can see a first version of this dashboard in Figure 7-2, in which sales amount, margin, and target values are sliced by date, brand, subcategory, and class. This dashboard is available on the page “Sales 2015” of the Sample-Sales.pbix file included in the companion content.

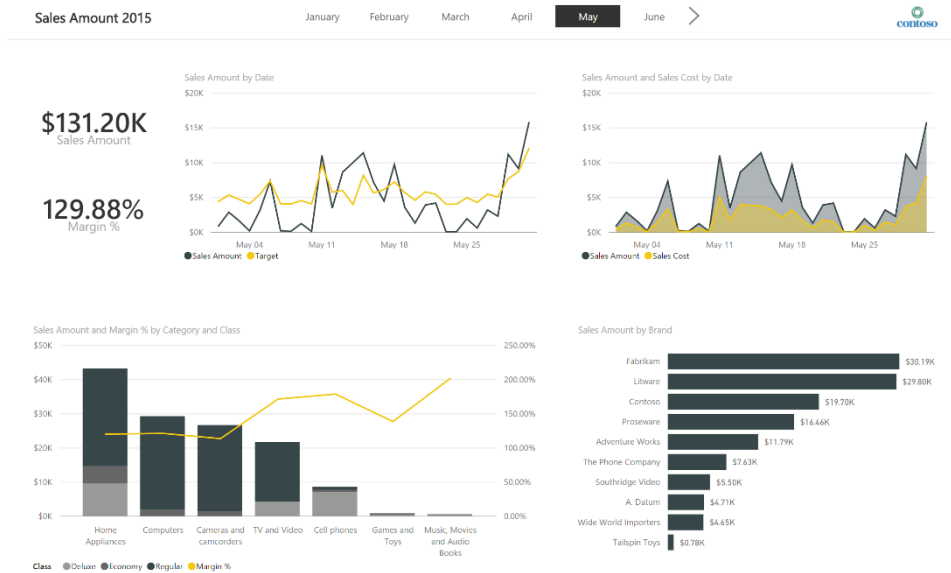


Figure 7-2: First version of a report displaying Contoso’s sales in 2015 using standard visualizations.

The first thing to notice in Figure 7-2 is that there are only two colors used in the report: black and yellow. Items in black identify the sales amount measure, whereas items in yellow are for other comparison measures (target, sales cost, and margin percent, depending on the visualization).

The four charts in the report use a limited number of common visualizations: line chart, bar chart, and columns chart. The rationale behind each visualization choice is described in the next section, but in general it is better to use simple, well-known visualizations when that is sufficient to produce the display of useful and understandable information.

Choosing between standard visuals

The first chart in Figure 7-2 compares Sales Amount and the Target values using a simple line chart. Figure 7-3 shows this in more detail.




Figure 7-3: A line chart of sales amount by date.

The line chart is the primary choice when you display a measure over a date range or time, using the x-axis for the temporal dimension. You can select the colors of the different measures by using the Data Colors properties, as demonstrated in Figure 7-4. Here you can select the color of each measure included in the line chart.

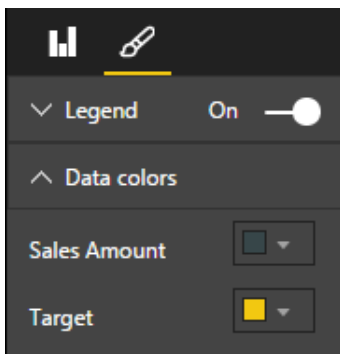


Figure 7-4: The Data Colors properties for a line chart visualization.

You can use a slight variation of the line chart when a measure that you want to display is part of another measure. For example, consider the Sales Amount and Sales Cost measures. Hopefully, Sales Cost is always lower than Sales Amount, and the graphical distance between them represents this margin. With the line chart, the gap between the two measures might not be clear, so you might want to “paint” the area below the line by using the values of the measures along time. The area chart does exactly this, as illustrated in Figure 7-5.




Figure 7-5: An area chart of sales amount versus sales cost by date.

The y-axis must begin at zero; otherwise, the area would not be fully representative of the two values. The visible gray area expresses the delta between the two measures in a graphical way and corresponds to the margin. You should not use an area chart when you have several intersections between different lines. You should consider it only when measures do not intersect often. The example in Figure 7-5 demonstrates one of the few cases for which you can consider using it.

Note For the sample reports in this chapter, for the most part we do not use pie charts and donut charts. This is because they are not considered a best practice, with an exception that you will see in the last section. The human brain can more easily make comparisons between lengths (as in a bar chart) than between angles (as in pie and donut charts).

Comparing measures with different scales requires particular visualizations. You need to display two y-axes, and you need a way to easily associate the axis corresponding to each measure. For example, Figure 7-6 shows a line and stacked column chart that yields more details by displaying the sales amount measure divided by category and class, compared with the margin percent divided by category.

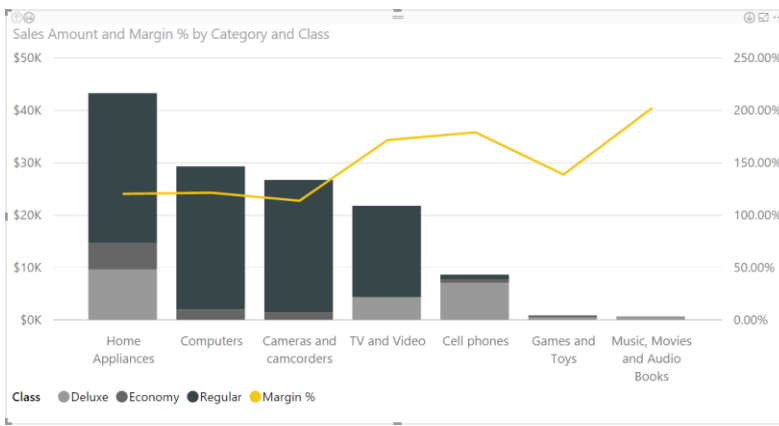


Figure 7-6: A line and stacked column chart of sales amount and margin percent by category and class.

The scale of the main measure (sales amount) is represented on the left y-axis, and the other measure (margin percent) is on the right y-axis. The x-axis shows the name of the category corresponding to each column, which is also divided by class using different shades of gray. Figure 7-7 shows the properties of this component used to bind data. The x-axis is called the *shared axis*, and it can include more than one attribute. In this case, we used two product attributes: Category and Subcategory. This makes it possible to perform an interactive drill-down of data in Power BI.




Figure 7-7: Data binding properties for a line and stacked column chart.

You can activate the drill-down feature for the selected column chart by clicking the drill down button (the down-arrow icon) located in the upper-right corner of the visualization. When drill-down mode is turned on, the drill-down button changes to display a black background, as depicted in Figure 7-8.



Figure 7-8: The drill-down button in a visualization. The black background signifies that drill-down mode is turned on.

With drill-down mode activated, when you click a column within the chart, you can drill down to that column's respective subcategories. Figure 7-9 shows the resulting chart when you click the Computers column in Figure 7-6. Note all of the subcategories that are related to the Computers category.




Figure 7-9: A line and stacked column chart of sales amount and margin percent by subcategory and class for the Computers category.

To drill up, in the upper-left corner of the visualization, click the drill-up button (the up-arrow icon), as highlighted in Figure 7-10.




Figure 7-10: The drill-up button in a visualization.

After you move up to the product category again, you can drill down to all of the subcategories for every category by selecting the double-arrow button in the upper-left corner of the visualization, as illustrated in Figure 7-11.

Click to drill down
to all subcategories




Figure 7-11: The drill-down button in a visualization.

When you drill down to the subcategory level for all the categories, you obtain a chart similar to that shown in Figure 7-12. Both the Sales Amount and Margin percentage measures have the subcategories granularity in the chart, so you can still compare them.




Figure 7-12: A line and stacked column chart of sales amount and margin percent by subcategory and class, for all the categories.

More info You can find an animated guide on how to use the drill-down feature in Power BI at <https://powerbi.microsoft.com/documentation/powerbi-service-drill-down-in-a-visualization/>.

Using custom visualizations

Power BI provides a number of embedded visualizations that are ready to use. By changing properties of the existing components, you can create solid presentations of your data. However, sometimes you might want to present data in a way that is not possible using the standard components. Power BI provides you with a gallery of custom visualizations that have been created by members of the Power BI community. You can download and install one or more of these custom visualizations in your report. To view a gallery of custom visualizations, go to <https://app.powerbi.com/visuals/>.

In this section, we will improve Power BI reports using features available in selected custom visualizations. The goal is to show you how to include custom visualizations in your report and what improvement you can achieve by using them. We suggest that you visit the gallery of available components because we cannot cover all the custom visualizations that are available, and new custom visualizations are published weekly.

First steps with custom visualizations

The first improvement we want to apply to the dashboard in the Sample-Sales.pbix file is coloring Sales Amount and Margin % according to their value. Using the standard card visualization, these values have a fixed color, as depicted in Figure 7-13.

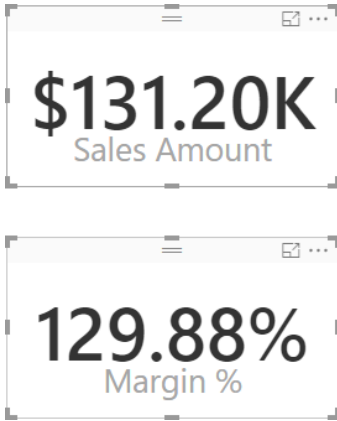


Figure 7-13: Sales Amount and Margin % displayed in a standard card visualization.

You can change the color in the Data Label properties, but the choice would be static. We would like to dynamically change the color so that the Sales Amount is green when it increases more than 20 percent compared to the value of the previous year, and the margin percentage is green when it is higher than 130 percent; yellow when it is between 100 percent and 130 percent; and red when it is lower than 100 percent (you might want to use a red color only for negative margins in other scenarios; we use ranges that adapt to this specific example).

To achieve this goal, you need a custom visualization that dynamically changes the color of the displayed value, based on a particular state. In the custom visual gallery, you can choose the Card With States By SQLBI visualization, as illustrated in Figure 7-14.

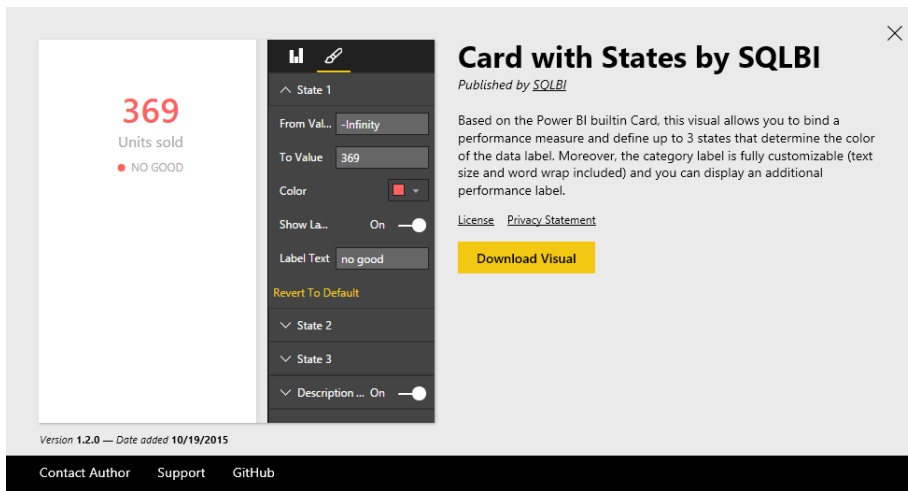


Figure 7-14: A description of the custom visualization Card With States By SQLBI.

After you download the visualization, you save the file using the .pbviz extension (Power BI Visual). Then, in the Visualizations pane, you can import that visualization in your report by clicking the Import From File button (the ellipsis), which is highlighted in Figure 7-15.

More info You can find a more detailed guide describing how to import a custom visualization at <https://powerbi.microsoft.com/documentation/powerbi-custom-visuals-use/>.

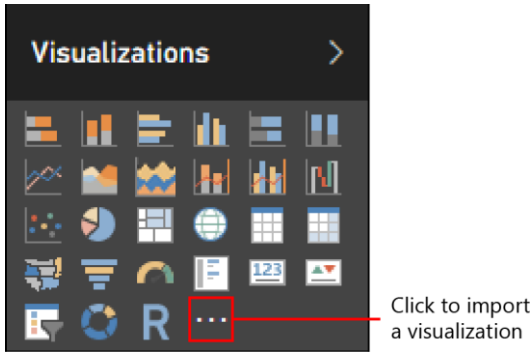


Figure 7-15: The Import From File button (ellipsis) in the Visualizations pane.

Note The file Sample-Sales.pbix already includes the custom visualizations used in this example. You should already see the corresponding icons without having to install them.

If you select the card visualization displaying Sales Amount, you can change it to Card With States By SQLBI by clicking its button in the Visualizations pane, which became available after you imported the custom visual. Figure 7-16 shows the new button.

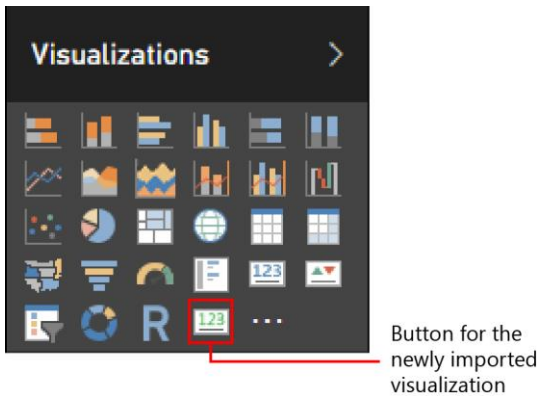


Figure 7-16: The button for the imported Card With States By SQLBI visualization.

In the Fields pane of the component (left side, Figure 7-17), you set State Value to YOY %, which is a measure defined in the data model that provides the year-over-year growth as a percentage. We want a red color if the growth is negative, yellow if the growth is positive and less than or equal to 20 percent, and green if the growth is greater than 20 percent. Set the To Value property (shown in the right pane in Figure 7-17) of the State 2 section to 0.2 (which corresponds to a 20 percent increase), and the From Value property of the State 3 section to 0.2, as well.

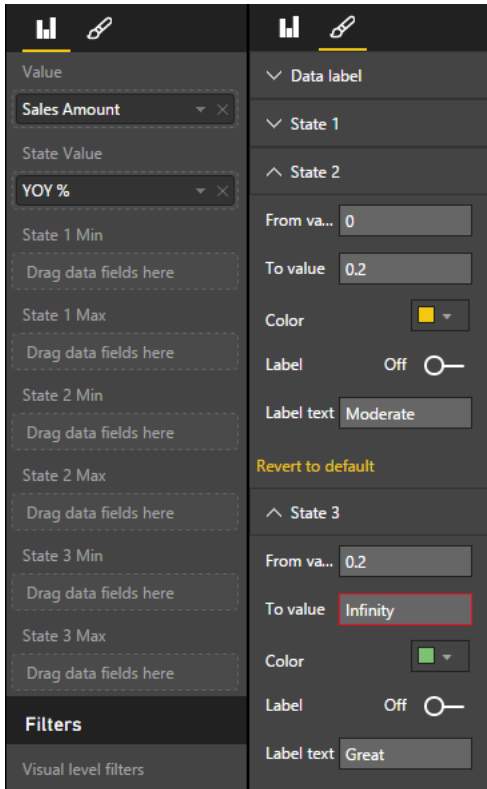


Figure 7-17: The Fields and Format panes for the Card With States By SQLBI visualization.

You repeat the same operation for the card visualization displaying Margin %: changing it to Card With States By SQLBI, setting the State Value field to the Margin % measure, and setting To Value in the State 2 properties to 1.3, and also changing From Value in State 3 to 1.3. In this way, you will display the Margin % using the same value displayed, whereas the Sales Amount will be colored, based on a year-over-year growth percentage. Figure 7-18 shows the final result. By changing the selection of months, categories, and brands, you might see different colors.




Figure 7-18: The Sales Amount and Margin percentage using Card With States By SQLBI.

We used this simple example to explain the process of importing and using custom visualizations in your report. In the following sections, we will focus more on considerations about when a custom visualization is useful or even required.

Improving reports by using custom visualizations

One of the charts available in the initial dashboard (the clustered bar chart in the lower-right corner of Figure 7-2) shows the sales amount divided by brand. This chart, as with all of the others in the dashboard, is dynamically updated whenever you select a different month or an element in other charts in the same dashboard. For this reason, the values displayed correspond to the sales filtered by the current selections in the charts and slicers in the dashboard. However, you might want to compare the sales amount of each brand with the corresponding sales made one year before as well as with the goal defined for the same selection. This initial chart displays only the sales amount, as shown in Figure 7-19.




Figure 7-19: A clustered bar chart of sales amount by brand.

If you want to display other measures in the same chart, you need to add them. Each measure will have a different bar, and you should set different colors to recognize each measure. For example, Figure 7-20 shows the goal measure and the sales amount measures for the years 2014 and 2015 (previously, we were displaying the sales amount only for the year 2015).

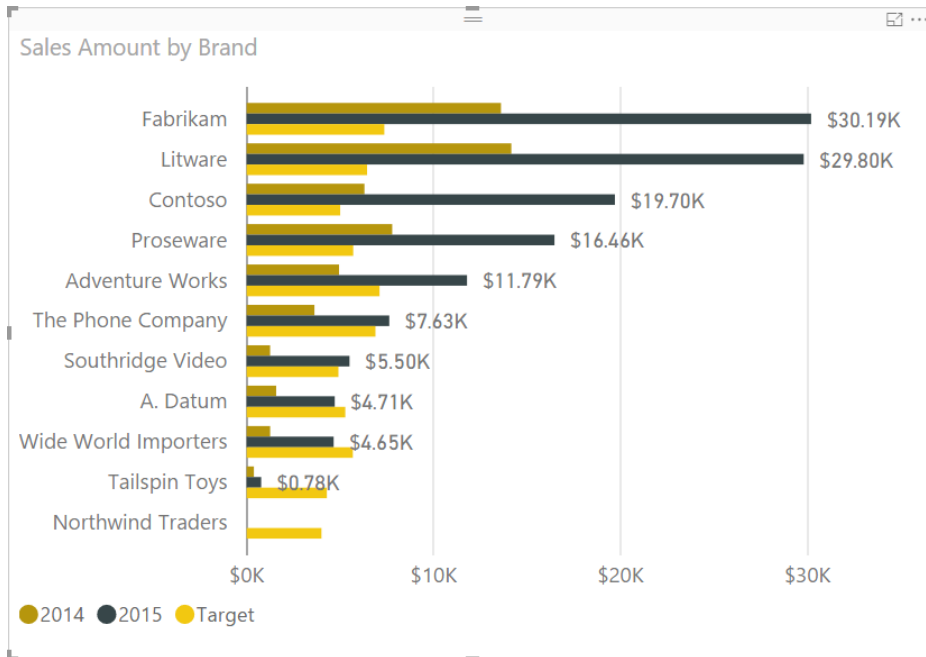


Figure 7-20: A clustered bar chart of sales amount and target by brand.

Note You might wonder how we displayed two values of the same measure (sales amount) in the same chart with two different names. In Power BI Desktop, you can create new measures in the data model, so you can simply assign an existing measure to a new measure with a different name and then assign the new measure to the visualization. You will find other considerations about using DAX to improve reports in the section “Using DAX in data models,” later in this chapter.

You might consider using a different visualization to improve the readability of this chart. For example, by using the custom Bullet Chart By SQLBI, you can obtain the results shown in Figure 7-21. The value of the sales amount for 2015 is a bar in the middle, surrounded by a shorter bar (overlapping on top of it) that has the sales amount for 2014, along with a short, black vertical line that acts as a marker for the goal value. The different graphics simplify the way the reader recognizes the more important value and the terms of comparison.




Figure 7-21: The Bullet Chart By SQLBI visualization, which displays actual sales amounts and goals by brand.

The only issue in this visualization is that we used black for the Target marker and yellow for the 2015 sales amount, which inverts the choice made in other charts of the same report. The reason for this is if we applied this color choice to other charts, it would have produced a barely readable marker for the goal value. In this case, the custom visualization can slightly improve the visualization, but it is not strictly necessary. Figure 7-22 demonstrates the final result of the Sample-Sales.pbix report, using the two custom visualizations that we've used thus far.

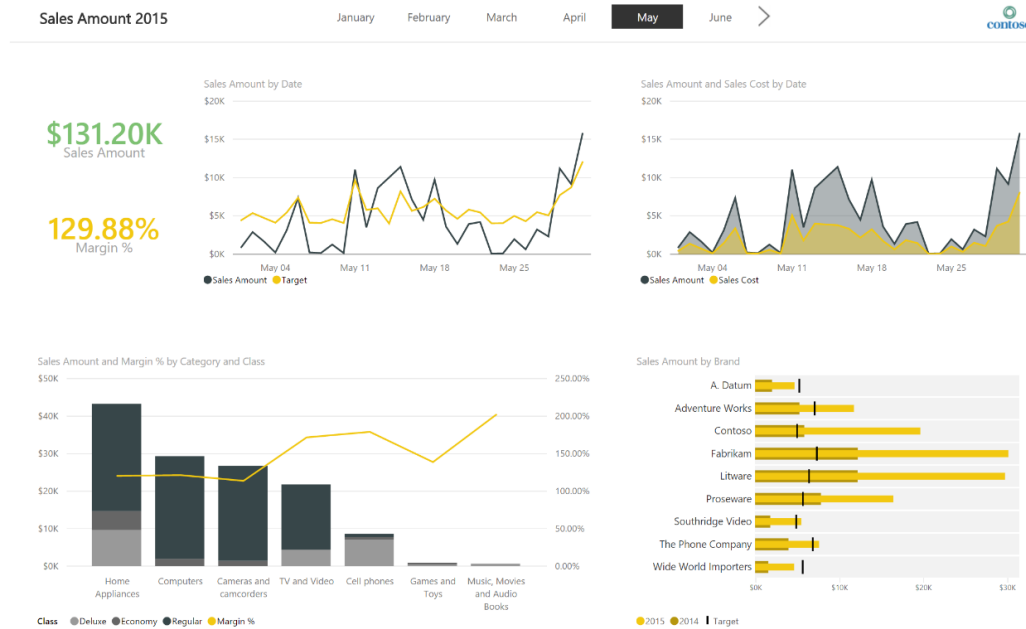


Figure 7-22: The final version of the report, displaying Contoso's sales in May, 2015. This version uses standard and custom visualizations.

Let's look at another example to see how using a custom visualization can improve a report. Figure 7-23 shows a report that represents the density of the population in each state within the United States, using the standard filled-map visualization. Darker shades correspond to higher-density values; lighter shades correspond to lower-density values.




Figure 7-23: A report displaying population density using a filled-map visualization.

This particular visualization (on the right of Figure 7-23) does not include the name of each state, because more than 95 percent of the states are condensed in less than 20 percent of the available real estate in the chart, so there simply is not enough room.

You can represent the same information by using a custom map and moving Alaska and Hawaii in a different position, changing also their size and making the map more readable. You can do this by using the Synoptic Panel By SQLBI visualization component that you can find in the visualization gallery. With this component, you can draw custom areas over any map image. (There is a gallery of maps ready to use at <http://synoptic.design>.) Figure 7-24 shows where you can find and download a map of the United States that includes the state names.




Figure 7-24: The gallery of country/region maps that are available on the Synoptic Designer website.

By using the Synoptic Design panel with the customized map of the United States, you can create the map shown in Figure 7-25, in which each state displays its name as well as its respective shade of gray. Moreover, the Synoptic Design panel also can work offline, whereas displaying the standard map component requires an active Internet connection to query the Bing service.




Figure 7-25: The same report displaying population density, but this time using a panel from Synoptic Design with a custom map of the United States.

Here again, you have seen how to improve a report by using a custom visualization, but until now we never had the need for a custom visualization to show the desired data. The standard components always offered an alternative way for you to display the same data that, even if not ideal, you could use if custom visualizations are not available. In the next section, you will see some examples for which a custom visualization is required to achieve a minimum goal.

Identifying conditions when custom visualizations are required

In this section, we use a different report to show conditions for which different visualization choices can produce more- or less-meaningful results. The report we want to consider shows the status of a stocks portfolio, using historical prices of stocks to display charts describing the behaviors of different shares and of the entire portfolio. Figure 7-26 depicts the initial version of this report.

Note You can find this report on the page Stocks in the Sample-Stocks.pbix file.




Figure 7-26: A report displaying stock portfolio performance over time using standard visualizations.

The report includes a line chart for each share, representing the closing price of the corresponding ticker. The details of the portfolio are included in a simple table in the upper-left corner, which displays for each share the quantity owned in the portfolio and its corresponding value at current prices. There is also a line chart for the entire portfolio, in which each share is represented by a single line of a different color. The line represents the total value of such a share in the portfolio over time. This chart can be useful to identify which stock has the largest value in the portfolio over time, but it does not provide in any way an indication of the total value of the portfolio over time. However, you can change this visualization to the stacked area chart depicted in Figure 7-27 to obtain a better representation of the portfolio's value.




Figure 7-27: The same portfolio value shown in Figure 7-26, but now shown over time, divided by share name.

Every color represents a different share, so you also have a rough estimation of the weight of each stock in the portfolio. The main difference between a stacked area chart and a line chart is that the former cumulates the value of each share name, whereas the latter displays each share name individually. In this case, an accurate choice of the visualization between the existing ones can satisfy the presentation requirements. However, in the other charts, a standard Power BI visualization is not very useful.

The four charts displaying the stock price for each day do not provide complete information. The data model actually contains different price values for each day: open price, minimum price, maximum price, and close price. It is very common to display these four values for each period considered (a day, in this example) by using a candlestick chart. This chart type is not available in the standard Power BI visualizations, so you need to install a custom visualization for that. For example, the Candlestick by SQLBI visualization provides a basic visualization that includes four measures for each period: Open, Close, High, and Low. Figure 7-28 presents the final result of the report after applying the two changes described in this section.

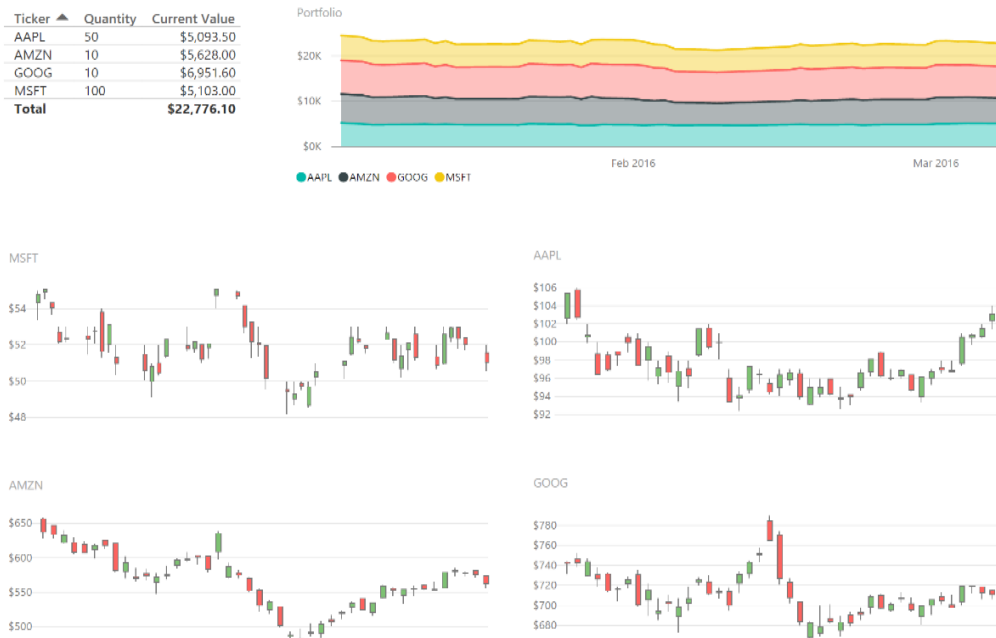


Figure 7-28: The same stock portfolio report displaying performance over time using custom visualizations.

You have seen how custom visualizations can improve the report, and sometimes they are necessary to achieve the desired graphical result. From time to time, though, you might still need to massage the data model to present measures and attributes to visualizations in the expected way, using the right granularity, and the expected formatting. The DAX language is your best friend here, as you will see in the next section.

Using DAX in data models

So far in this chapter, we have presented several examples of visualizations. We've demonstrated how you can improve your reports by choosing the right visualization, setting the properties in an appropriate way, and even installing custom visualizations when required. However, there are a number of improvements that you can achieve that do not require a direct action on the visualizations. You can instead create *measures* or calculated columns in DAX. Usually, you use DAX to obtain a certain numeric result, but sometimes you can take advantage of a DAX expression to control the report layout.

Our first example concerns the measures used in the candlestick charts of the report shown in Figure 7-28. Every time period displayed involves four measures: Open, Close, High, and Low. The data recorded in the data model has four corresponding columns for each day and stock. However, you might use the candlestick chart to display data by week or by month instead of by day. The aggregation required for each measure depends on the measure itself. The Open measure's price must be the DayOpen value of the first day in the period, the Close price must be the DayClose value of the last day in the period, the High price must be the maximum value of DayHigh in the period, and

the Low price must be the minimum value of DayLow in the period. You can write these four measures by using the following DAX expressions:

```
Open =
IF (
    HASONONEVALUE( StocksPrices[Date] ),
    VALUES ( StocksPrices[DayOpen] ),
    CALCULATE ( VALUES ( StocksPrices[DayOpen] ), FIRSTDATE ( StocksPrices[Date] ) )
)

Close =
IF (
    HASONONEVALUE( StocksPrices[Date] ),
    VALUES ( StocksPrices[DayClose] ),
    CALCULATE ( VALUES ( StocksPrices[DayClose] ), LASTDATE ( StocksPrices[Date] ) )
)

High = MAX ( StocksPrices[DayHigh] )

Low = MIN ( StocksPrices[DayLow] )
```

As mentioned earlier in this chapter, another useful tip is to create a measure just to display a measure with a different name. The reason is that many visualization components use the measure name in a legend or other description, and you do not have a way to rename that by using visualization properties. For instance, if you have two measures, Current and Previous, but you want to display the exact year in a particular visualization, you might create two measures with the exact year that you want to display in the legend, as in the following example:

```
[2014] = [Previous]
[2015] = [Current]
```

In a report that you will see in the next section, we will use data extracted from Google Analytics. The Website table contains the columns Users and Country ISO Code. In the dashboard, it will be interesting to show on a map the ratio between the number of users visiting a website and the population of the country/region where users come from, but such information is not available directly from Google Analytics. You can import the information about countries'/regions' populations in a separate Countries/Regions table and link it to the Website table using the Country ISO Code column, as illustrated in Figure 7-29.

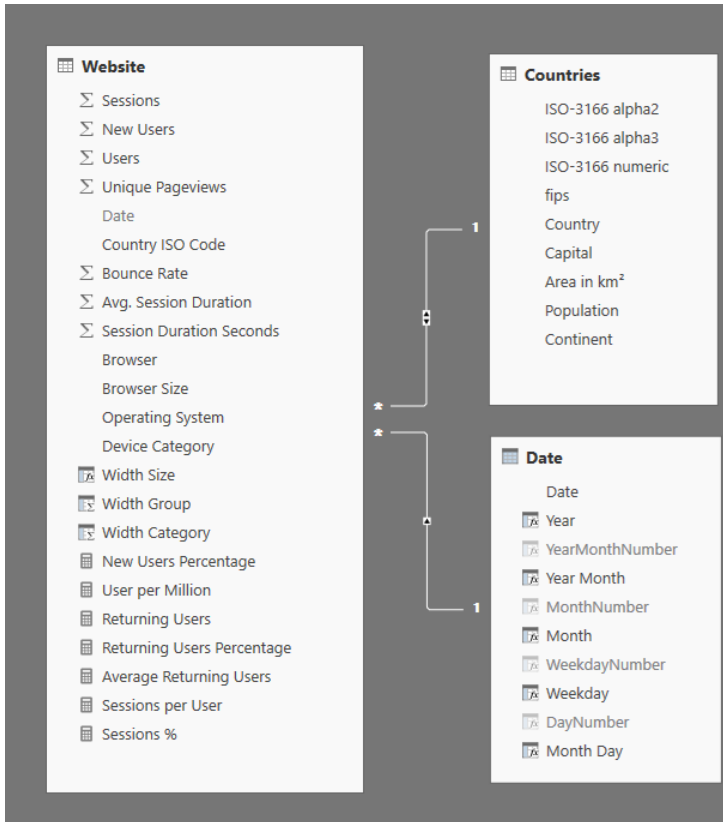


Figure 7-29: A schema of tables and relationships of a data model that extends Google Analytics website data.

Because the ratio would be a decimal number, you can create a metric called Users Per Million by using the following measure definition in DAX:

```
Users per Million =
DIVIDE (
    SUM ( 'Website'[Users] ),
    SUM ( 'Countries'[Population] )
) * 1000000
```

You should not expect that a visualization component would directly do a calculation, such as a ratio or a difference. It is always better to define the corresponding calculation in a DAX measure, and then you bind that measure to the visualization.

Finally, you should consider using a calculated column when you need a classification that groups existing data with a high granularity to a lower number of unique values that are easier to display in a chart. For example, Figure 7-30 shows the existing values in the Browser Size column that is provided by Google Analytics. There are more than 5,000 unique combinations of width and height, and this fragmentation of values makes the analysis harder. Moreover, each resolution is a single string, and the reports should group different resolutions by width, ignoring the height.

Browser Size
1260x680
1350x670
1520x770
1660x920
1010x670
1590x800
1350x650
1420x750
1350x650

Figure 7-30: A partial list of values included in the Browser Size column.

You can split the problem into two steps. First, you extract the width size from the string, converting the digits before the "x" character in an integer number. Then, you compare this number with a list of predefined values representing the interesting range of resolutions you want to analyze, such as 1024, 1280, 1440, 1920, and 2560. You can see the two calculations implemented in the following two measures, Width Size and Width Category, respectively:

```
Width Size =
VAR xPos = FIND ( "x", Website[Browser Size], , 0 )
VAR widthString = IF ( xPos > 1, LEFT ( Website[Browser Size] , xPos - 1 ), "" )
RETURN IF ( widthString <> "", INT ( widthString ) )

Width Category =
SWITCH (
    TRUE(),
    Website[Width Size] <= 1024, 1024,
    Website[Width Size] <= 1280, 1280,
    Website[Width Size] <= 1440, 1440,
    Website[Width Size] <= 1920, 1920,
    Website[Width Size] <= 2560, 2560,
    CALCULATE ( MAX ( Website[Width Size] ), ALL ( Website ) )
)
```

Figure 7-31 presents the results of the two calculated columns. We will use the Width Category in one of the visualizations described in the next section.

Browser Size	Width Size	Width Category
1260x680	1260	1280
1350x670	1350	1440
1520x770	1520	1920
1660x920	1660	1920
1010x670	1010	1024
1590x800	1590	1920
1350x650	1350	1440
1420x750	1420	1440
1350x650	1350	1440

Figure 7-31: A partial list of values included in the Browser Size, Width Size, and Width Category columns.

Creating high-density reports

In the last section of this chapter, we want to discuss the challenges that you face when creating high-density reports. When you include many visualizations in a single report, you need to carefully balance the amount of information provided in each of those visualizations, removing all the unnecessary elements that would reduce the attention of the user. You want your user to focus only on the data.

As you will see, having an idea of the overall structure and tuning properties of each visualization is much more important than using particular custom visualizations, which are usually the icing on the cake.

Figure 7-32 depicts a first version of a report showing website data from Google Analytics for the DAX Formatter website, which is available in the companion content, in the file Sample-DAXFormatter-Analytics.pbix file. The report contains 28 visualizations with data, plus one slicer and eight components without data that are there only for aesthetic reasons (title, logo, pictures, and separators). The 28 visualizations only use 7 different visualization types, and some of them are simple variations of the same concept, such as stacked/clustered bar/column charts. You can obtain a complete and complex report using only a few component types.



Figure 7-32: A high-density report based on Google Analytics data.

The entire report is organized in three zones: left, center, and right. The left zone contains metrics regarding the number of users, the center zone shows data about the sessions, and the right zone includes technical information about average page load time, device type, operating system, browser, and resolution used by website visitors.

If you were to try to create a similar report starting from scratch, you would need to apply the following guidelines:

- **Reduce text** Include only the minimum necessary of textual elements, avoiding repetitive or verbose descriptions.
- **Remove legends** Whenever possible, avoid including a legend, especially whenever there is only one measure displayed in the chart.

- **Remove axes** In a compact visualization for which you already included the data labels (setting the Data Labels property to On), you can remove corresponding axes. All the clustered bar charts in the report are formatted in this way.
- **Use images to explain concepts** Use an icon or a meaningful image related to the data you show. Remember, a picture is worth a thousand words. In the example in Figure 7-32, at the top of the report is one different image for each of the three zones (Users, Sessions, and Average Page Load Time).

In the report, we used a donut chart. Previously in this chapter, we mentioned that using a pie chart or a donut chart is not a good idea, so you should avoid doing that. The exception we wanted to include in this report is when you compare only two values. In this example, we used a donut chart for the search engines distribution, showing the percentage of sessions coming from Bing searches versus Google searches. It is clear that Google has a clear leadership for directing visitors to this website, with Bing producing only a marginal contribution. For this difference, looking at the exact number or percentage is not relevant.

There are three visualizations that we wanted to improve, starting from this initial example. The sparklines used at the top of the report are simple line charts without any axes, legend, data labels, or border. Figure 7-33 illustrates that we set to Off most of the format properties of those line charts.

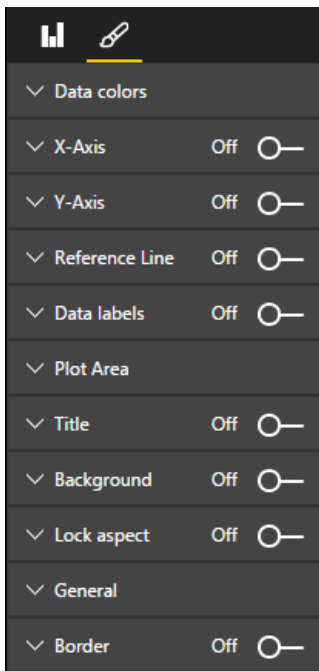


Figure 7-33: The Properties pane of a line chart used as a sparkline.

However, you cannot change the line width of the line chart. When you use the line chart in a small area, this produces a result that is difficult to read. You can replace the line chart with the Sparkline custom visualization that you can download from the Power BI visualization gallery. Figure 7-34 shows a side-by-side comparison of the same chart displayed by using a line chart (on the left) and a sparkline custom visualization (on the right). The custom visualization draws a line with a smaller width, generating a final result that is more readable.




Figure 7-34: A couple of examples of the same chart, using a standard line chart and a sparkline custom visualization.

Another visualization to improve is the countries/regions penetration displayed in the lower-left corner of the report. Instead of using a standard map, which shows a pie for each country/region with a size depending on the population density, you can use the Synoptic Designs panel, loading the world map from the gallery shown back in Figure 7-24. The result of displaying the measure Users Per Million in a Synoptic Designs pane is visible in Figure 7-35.



Figure 7-35: Synoptic Designs panel with a world map that displays the ratio between users and population.

The last visualization to improve is the one in the lower-right corner, showing the number of sessions by browser resolution. In the original data, we have a very fragmented number of different resolutions, so the bar chart only displays the first values, but the number of sessions of the most common resolution is just five percent of the total number of sessions, so there are a wide number of different resolutions considered, most of them not visible in the report. We solved this problem in two steps. First, we created a column containing the width category, which classifies the width extracted from the resolution string, as you have seen in the previous section about DAX. Then, we changed the clustered bar chart to a waterfall chart, using the Sessions % measure instead of the Sessions measure, as shown in Figure 7-36.




Figure 7-36: A waterfall chart with the distribution of sessions by browser resolution.

We do not use any decreasing step in the waterfall chart, but the final result clearly shows the distribution of the resolution in a meaningful way. The Sessions % measure is a DAX expression created just for this report, using the following definition:

```
Sessions % =
DIVIDE (
    SUM ( Website[Sessions] ),
    CALCULATE (
        SUM ( Website[Sessions] ),
        ALL ( Website[Width Category] )
    )
)
```

Figure 7-37 presents the final result, after these improvements have been incorporated. As you can see, the overall difference is an incremental improvement and not a substantial change from the result we got initially by using standard visualization components.

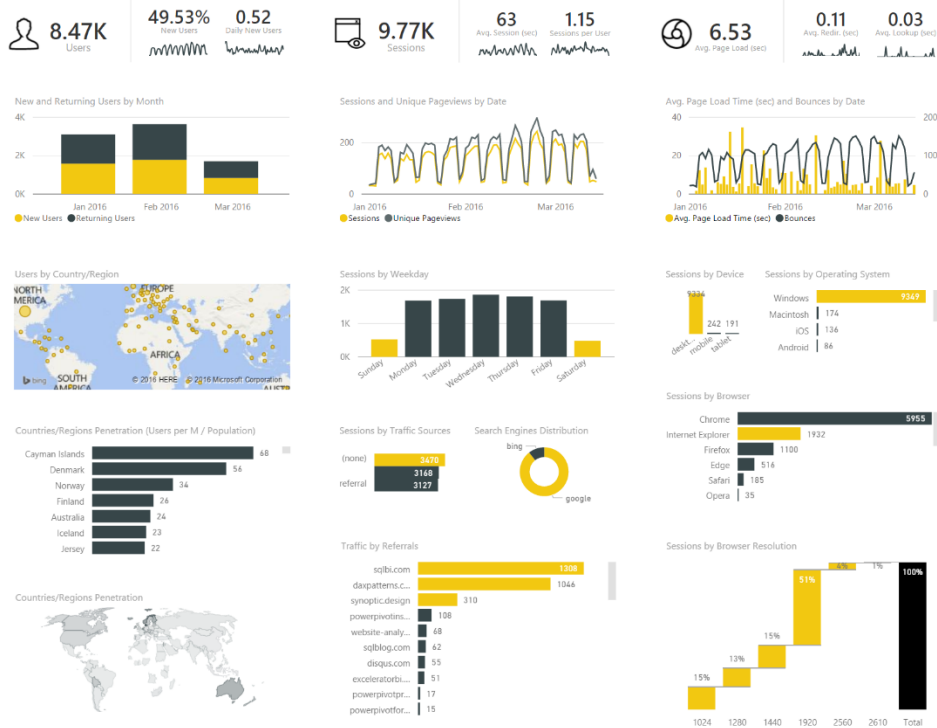


Figure 7-37: A high-density report based on Google Analytics data.

Especially in a high-density report, you should focus on the overall quality and readability, reducing the number of distractions for the reader. There is already an overwhelming amount of information, so the focus of the reader should be entirely on data, not decorations or visualizations that are too complex and do not provide any added value.

Conclusions

In this chapter, you have seen a number of techniques to improve Power BI reports by choosing the best built-in visualizations and adding custom visualizations. Here are the main steps in this process:

- **Choose the right visualization type** You can choose between many visualization types, but usually you do not need too many of them in the same report. Do not be afraid of using the same visualization type many times if it is the one that best displays your data.
- **Customize visualization properties** You can customize every visualization by using format properties. Using a consistent color scheme is one of the most important aspects of a good report.
- **Consider custom visualizations when necessary** The Power BI custom visualizations gallery provides you with many visualizations that extend the set of standard ones available in Power BI. You should consider using them when there is a concrete advantage over the standard visualizations.

- **Use DAX to create measures and calculated columns** You also can use DAX expressions to achieve the desired visualization. Even a simple transformation, such as renaming a measure in a legend, is not always possible within the properties of a visualization, but you can overcome this by creating new measures and calculated columns.
- **Remove unnecessary elements in high density reports** In a high-density report, you need to remove any graphical element that is not necessary to communicate information to the users; you do not want to distract them by including details that do not provide any useful information.

These guidelines are just a starting point in your journey to create clear and useful reports. Your experience, the feedback from users of your reports, and the analysis of reports created by other people are the other important steps along this road.

Using Microsoft Power BI in your company

With Power BI, you can create a dashboard using some data in a Microsoft Excel workbook, or you can connect to existing structured databases in your company. The previous chapters in this book demonstrate how you can create data models, reports, and dashboards, gathering information from different data sources. They also show you how you can consume these results on different devices. This chapter shows you techniques to obtain a deeper integration with existing systems and applications.

You can get data from many existing data sources in your company. You can embed a dashboard visualization in Microsoft Office with minimal effort. If you have developer skills, you also can take advantage of the REST API to automate operations in Power BI by implementing real-time updates of dashboards, with no more latency than just a few seconds.

Power BI is a platform that exposes several services through a REST API that is easy to use by any application, including standard web-based applications and those running on mobile devices. In this chapter, we will look at some example results that will help you to understand why the presence of this extensibility option is so important. The goal is not to describe in detail how to use these APIs from a developer's point of view, but to make you aware of what you can obtain by using them. If you are a developer, you will find some links and references to additional material to understand how to use the existing API. If you are a business user or a BI architect, you will gain a fuller understanding and know what you can and cannot ask a developer. However, keep in mind that the API for Power BI

is undergoing constant evolution. If something is not possible as of this writing, it might become possible in a future release. When in doubt, check what new features are available in the API.

Getting data from existing systems

Any company has a number of existing data sources that you can use in Power BI. You have seen that you can create a data model in Power BI by copying the content of tables that exist in other databases or files. You also have the option of refreshing this content dynamically, or you can directly query the data source whenever you access a report. By querying directly, you avoid the need to create a copy of the data that you must then synchronize periodically. In this section, you will see the available options with which you can connect Power BI to either your on-premises database or a database in the cloud.

Before looking at the details, here are a few terms with which you should be familiar:

- **On-premises** If you get data from a database that is physically stored in a server managed by your company, we say that the database is on-premises (often shortened to on-prem).
- **Cloud** If you get data from a Microsoft Azure service, you are using data in the cloud. Cloud computing accesses and uses shared compute and storage resources on the Internet.
- **Relational database** This is a database that stores data using tables that have relationships with one another. Typically, you query this by using the SQL language. Examples of on-premises relational databases that Power BI supports are Microsoft SQL Server, Microsoft Access, Oracle, IBM DB2, MySQL, PostgreSQL, Sybase, and Teradata. Cloud-based relational databases that Power BI supports include Azure SQL Database and Azure SQL Data Warehouse.
- **Rich semantic model** This is a database that stores both data and metadata, simplifying navigation by using tools such as Excel PivotTables and Power BI reports. A typical example is Microsoft SQL Server Analysis Services. Other supported providers are SAP HANA and SAP Business Warehouse.
- **Power BI Personal Gateway** This is a component installed on the user's computer that makes it possible to perform data refreshes on models published using the Power BI service. (Chapter 3 explains how to install this.) A Personal Gateway serves only one user, and only when the user's computer is turned on.
- **Power BI Enterprise Gateway** This is a component similar to the Personal Gateway that a system administrator installs on a server in your company. A single Enterprise Gateway can serve all the users of a company, and it is also available as soon as the server is turned on (servers are usually active 24/7). You can find more technical details about how to install it at <https://powerbi.microsoft.com/documentation/powerbi-gateway-enterprise/>.

There is also another concept to clarify before moving forward, which is the difference between a model requiring data refresh and a live connection.

Understanding differences between data refresh and live connections

When you navigate in data via Power BI, you can read a copy of the data stored in Power BI (either the Power BI service or the Power BI Desktop application), or you can have a live connection that sends a real-time query to the data source without creating a copy of the data.

Chapter 4 shows that when you connect Power BI to a SQL Server database, you have two connection settings from which to choose: Import and DirectQuery. The DirectQuery option does not create a

copy of the data in Power BI, and it translates any user action made navigating on a report into one or more queries to SQL Server. In a more general classification, we can say that you can either import data in Power BI or connect to the data source via a “live connection.” With relational databases, DirectQuery is the tool used to obtain a live connection to the data source. As you will see later in this chapter, in Power BI you have similar options when you connect to a SQL Server Analysis Services data source: Connect Live and Import Data.

Regardless of the underlying database, when you create a model in Power BI by importing data, you have full access to the features of Power BI. However, you need to run or schedule a data refresh to keep data updated on Power BI.

On the other hand, when you use a live connection, your Power BI model can have only one data source, so a single Power BI report cannot mix visualizations connected to data coming from different data sources. To do that, you must import data into Power BI. In a dashboard, however, you can always include visualizations from different reports; thus, you can combine visualizations from different live connections in a dashboard only.

In the following sections, you will see in more detail how to use live and imported data sources using existing databases and models in your company.

Using relational databases on-premises

When you create a data model in Power BI Desktop, you often get data from an on-premises relational database. For example, Chapter 4 demonstrates how to create a new Power BI Desktop data model that gets data from a Microsoft SQL Server database. In this case, you use a database on-premises, and you can choose between the Import and DirectQuery connection types. The former creates a copy of the data in the Power BI Desktop model, which requires a refresh in order to synchronize the content of the Power BI data model with the source database. The latter does not create a copy of the data; instead, Power BI generates queries to SQL Server every time you navigate in a report.

In both cases, after you publish the Power BI Desktop file to the Power BI service, the refresh operation requires either a Personal Gateway or an Enterprise Gateway. If you use the Personal Gateway, you can only refresh datasets created by importing data, but you cannot use the DirectQuery option. To publish a Power BI Desktop data model created by using DirectQuery, you need the Power BI Enterprise Gateway.

Figure 8-1 illustrates what happens when you publish a Power BI data model connected to on-premises data sources via the Import connection type. A copy of the data and the description of the data model are stored in the Microsoft cloud. The data is always available to queries sent by any client, which sees data updated on the last data refresh. You need a Power BI gateway installed on-premises to complete the data refresh: either a Personal Gateway or an Enterprise Gateway, in this scenario.




Figure 8-1: A Power BI gateway is required to refresh a Power BI model, getting data from on-premises databases.

Figure 8-2 shows what happens when you publish a Power BI model created by using the DirectQuery option. The Power BI service does not store a copy of data in the Microsoft cloud; it has only a semantic description of the data model, with the information required to retrieve data from the original source database. Every time the Power BI service receives a query, it generates one or more queries in SQL language and sends these requests to the relational data source through the Power BI Enterprise Gateway.

More info The Microsoft cloud service does not preserve any data received from the relational databases on-premises; it might only keep a transient data cache on a volatile device in order to improve the performance of other queries sent by the same user. You can find more information in the Power BI Security whitepaper published by Microsoft at <http://download.microsoft.com/download/4/8/C/48CFCF8A-2025-4B97-B249-7B505E26E7ED/Power%20BI%20Security%20Whitepaper.docx>.




Figure 8-2: For on-premises data sources, you must have the Power BI Enterprise Gateway for a model using the DirectQuery connection setting.

As of April 2016, you can take advantage of DirectQuery on SQL Server, Oracle, or Teradata relational databases, which are all supported in the Power BI Enterprise Gateway.

Using relational databases in the cloud

If you create a Power BI model that uses a relational database stored in the cloud, you might not need the Power BI Gateway to refresh data. Power BI supports direct connection to Azure SQL Database and Azure SQL Data Warehouse data sources, so you can schedule a data refresh or you can use DirectQuery without the need to install and configure any gateway.

You will still have a different architecture, depending on which connection setting you use, Import or DirectQuery. Figure 8-3 illustrates that by using the Import setting you still have a copy of data owned by the Power BI service, but you can refresh that copy without any gateway if the data source is Azure SQL Database or Azure SQL Data Warehouse.




Figure 8-3: Power BI can connect directly to Azure SQL Database and Azure SQL Data Warehouse to refresh a Power BI model.

Figure 8-4 illustrates the behavior of Power BI using DirectQuery connected to Azure SQL Database or Azure SQL Data Warehouse. As with any DirectQuery connection, the Power BI service has only a semantic description of the data model, along with the information required to retrieve data from the original source database. It does not store a copy of data in the Microsoft cloud. Every time the Power BI service receives a query, Power BI generates one or more SQL queries and sends these requests to the relational data source, with no gateway required.

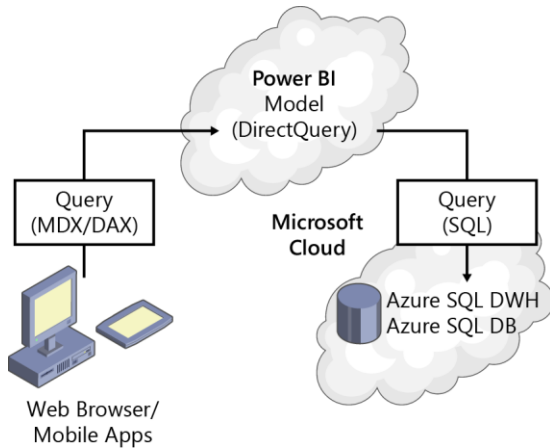


Figure 8-4: In DirectQuery mode, Power BI connects directly to Azure SQL Database and Azure SQL Data Warehouse.

If you have a cloud-based relational database, other than Azure SQL Database and Azure SQL Data Warehouse, you must use the architecture for on-premises data, and you need to install a gateway to complete the refresh operation or to implement DirectQuery.

Note Multiple requests on different servers in different locations might increase the latency of requests, so you might want to consider installing the Enterprise Gateway on a server hosted in Azure Virtual Machines to improve the performance.

Using live connections to Analysis Services

When you create a live connection to Analysis Services in Power BI Desktop, you do not create a data model, and you do not have a copy of the data in the PBIX file. Thus, when you publish the model on the Power BI service, the PBIX file contains only the definition of the reports, but the entities are defined in the Analysis Services file. When you edit a report in Power BI that is tied to a live connection to Analysis Services, all of the operations requested by the client are redirected to Analysis Services through the Power BI Enterprise Gateway, as shown in Figure 8-5.




Figure 8-5: When a model has a live connection, Power BI redirects a query to Analysis Services on-premises.

The Power BI service contains neither data nor metadata for the data model. Any change made to the data model in Analysis Services is automatically reflected in Power BI, without requiring a data refresh operation in Power BI. However, keep in mind that usually Analysis Services has a copy of the data read from one or more data sources (unless you create a data model in Analysis Services by using DirectQuery), so you need to refresh the model in Analysis Services to keep it up to date.

Integrating Power BI with Office

Showing dashboards and reports created in Power BI using one of the options available (web browser, dedicated apps on mobile devices, and Power BI Desktop) is very useful. However, perhaps you want to create a particular report in Excel or a presentation in Microsoft PowerPoint, which could benefit from a tighter integration with Power BI. As you will see in this section, there are several features in Power BI that can take advantage of such a service in certain applications of Office.

Publish Excel data models in Power BI

When you publish a Power BI Desktop file to Power BI, you are copying to the cloud a file containing a data model, a copy of the data, the query to import and refresh the data, and all the reports you created. If you have an Excel file with a data model, you have a similar file, and you can publish such a file on Power BI, as well. In fact, the following correspondence exists between features in Power BI and Excel:

- Power BI data model → Excel Data Model (also known as Power Pivot data model)
- Power BI Query Editor → Workbook Queries (formerly known as Power Query in Excel 2010 and Excel 2013)
- Power BI report → Power View

You can load into Power BI an XLSX file containing a data model instead of a .pbix file. In doing this, you keep all of the existing features and reports in your Excel file, but you also can then use the same data model in Power BI. All the existing reports in Power View are converted in equivalent reports in Power BI whenever possible (certain features of Power View might not have a corresponding

visualization or feature in Power BI). In this scenario, the PivotTables and PivotCharts you have in Excel continue to work with the Power Pivot data model. Figure 8-6 presents the Publish To Power BI feature that is available in Excel 2016, which guides you in publishing a Power Pivot data model to Power BI without even opening the Power BI website, similar to how you would publish within Power BI Desktop.





Figure 8-6: The Publish To Power BI feature in Excel 2016.

It is important to consider that, if you import the Power Pivot data model in Power BI Desktop, you will have a .pbix file instead of an .xlsx file. In this way, you can work locally with Power BI, but you will have two copies of the same data model and the same data. You will use Excel to navigate data with PivotTables and PivotCharts, and you will use Power BI to navigate using dashboards and reports. However, if you publish a .pbix file, you also can consume your data in Excel using the Analyze In Excel feature, which is described in the next section.

Consume Power BI content from Excel

In this book, you have seen how to import data from Excel to Power BI by using different techniques. However, you might want to move data in the opposite direction, consuming in Excel data that is published on Power BI. This is indeed possible, and you can do it by using the Analyze In Excel feature that is available in the Power BI service. Figure 8-7 depicts the Analyze In Excel action that is available for datasets and reports.



Click to let Excel consume data that's published on Power BI

Figure 8-7: The Analyze In Excel feature in Power BI.

When you click Analyze In Excel, you might be prompted to install an updated driver for Excel; if this happens, follow the instructions to install the suggested driver. In any case, this action downloads a small file on your computer (with an .odc extension) that Excel uses to open a connection to the corresponding Power BI model, creating a PivotTable on top of the model. Figure 8-8 shows the result if you request Analyze In Excel on the dataset Google Analytics.




Figure 8-8: A PivotTable in Excel that is connected to the Google Analytics model in Power BI by using Analyze In Excel.

Note By establishing a connection using Analyze In Excel, you are consuming your Power BI model as if it were an external analytical database. This is the same behavior you have in Excel when you connect to an Analysis Services database, or you connect Excel to a Power Pivot model published on SharePoint by using the Excel document URL as the server name. The Analyze In Excel feature uses Excel only as a client, without storing the data model within Excel, as you would do when you are using Power Pivot for Excel.

After you establish the connection with a PivotTable, you have all the tables, columns, and measures of the Power BI data model listed in the PivotTable fields. You can use only the measures that have been explicitly defined in the data model, so you cannot create measures during the navigation as you can do in Power BI. For this reason, it is important that you create all the measures that could be necessary to an Excel user, without assuming that any numeric column can be aggregated or that a measure can be created upon request. For example, the Sales 2015 – Analytics model that David created in Chapter 5 does not contain any explicit measures, and this makes it difficult to use in Excel. Figure 8-9 illustrates how the data of such a model is presented in the Excel PivotTable Fields pane.

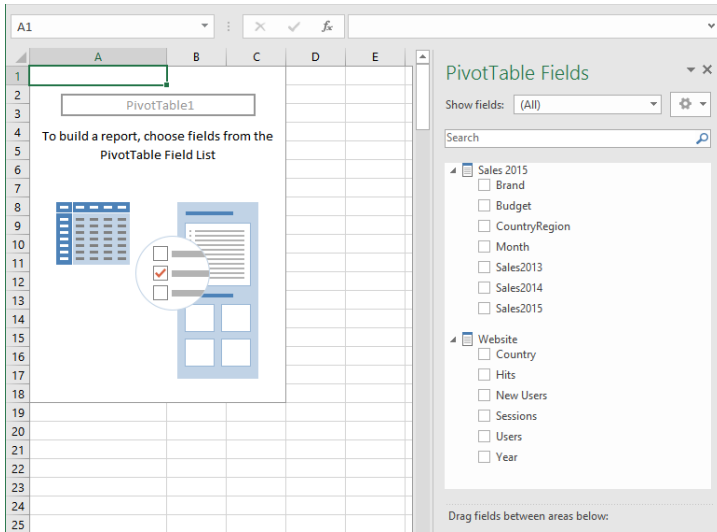


Figure 8-9: A PivotTable in Excel that is connected to the Sales 2015 - Analytics model in Power BI by using Analyze In Excel.

The model without explicit measures is not very useful in Excel, because you do not have any calculations to put in the Values area of the PivotTable. For example, Figure 8-10 shows the result of selecting the columns CountryRegion, Sales2013, and Sales2014 from the PivotTable Fields pane. These columns are placed in the Rows area of the PivotTable, and you cannot move them in the Values area. Thus, the result is a list of all the unique values of these columns, grouped by CountryRegion, Sales2013, and Sales2014.

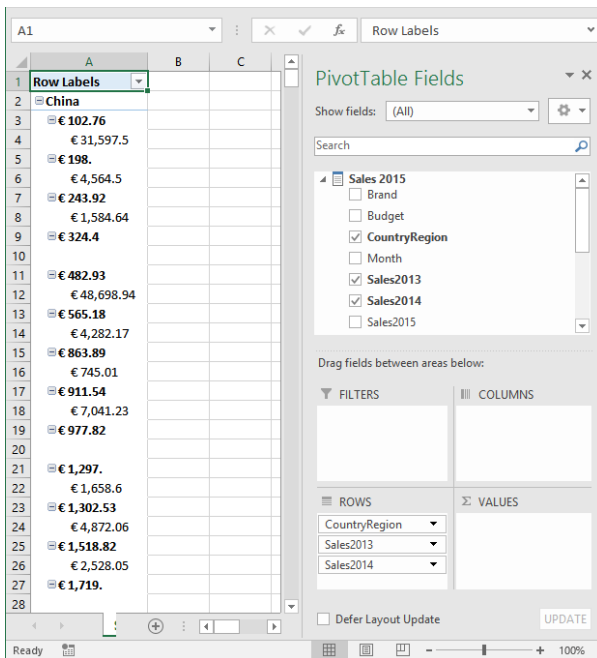


Figure 8-10: A PivotTable in Excel that is connected to the Sales 2015 - Analytics model in Power BI by using Analyze In Excel.

Excel ignores that the nature of the sales columns is that of numbers that can be aggregated. This information must be provided by using explicit measures, such as those demonstrated in Chapter 6. For example, Figure 8-11 shows a PivotTable obtained by using Analyze In Excel on the Budget data

model created in Chapter 6 and published on Power BI. In this case, the measure Sales Amount is aggregated by Brand and Year.

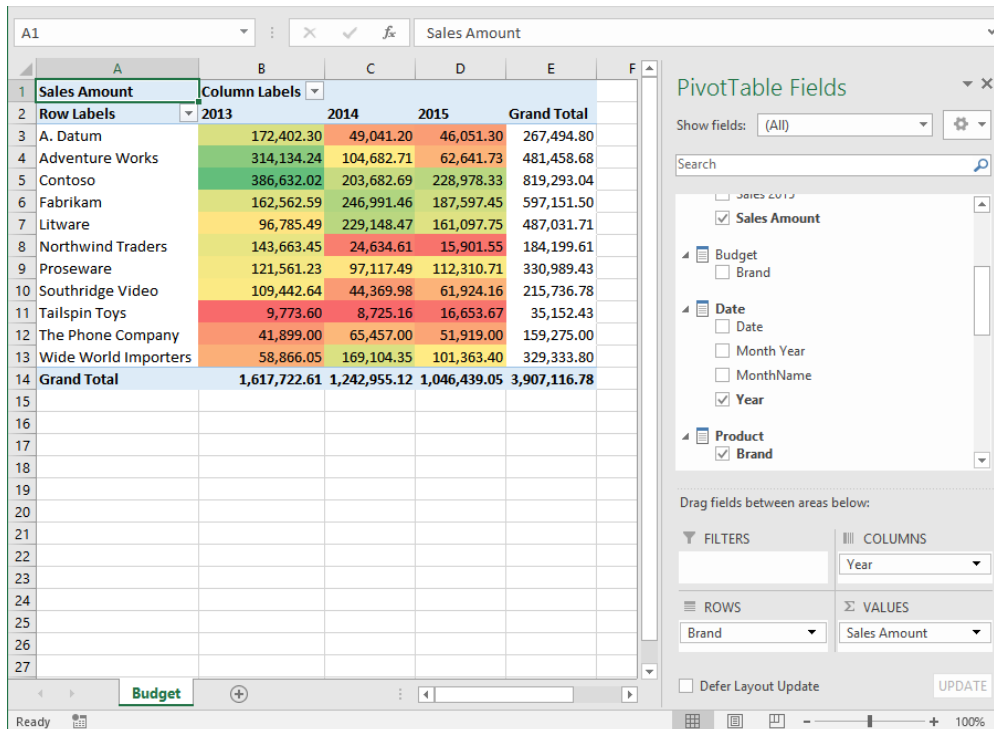


Figure 8-11: A PivotTable in Excel that is connected to the Budget model in Power BI by using Analyze In Excel.

One of the reasons to analyze data in Excel is to take advantage of specific Excel features. For example, in Figure 8-11 we applied a conditional formatting rule, so that higher values have a green background color, and smaller values have a red background color.

You might have many other reasons to use Excel to analyze a Power BI model. In general, Excel is a flexible application with which you can easily integrate data coming from different sources with data and/or calculations you have in the Excel file. Moreover, as of this writing, many of the features of a PivotTable in Excel are not yet available in the Power BI visualizations. Thanks to Analyze In Excel, you have the maximum flexibility to combine different clients (Power BI and Excel, for example) to analyze the data of the same data model.

Note The authentication used to connect Excel to the Power BI model requires a new version of the OLE DB driver that is used to establish a connection to an external Analysis Services database (OLE DB for OLAP). For this reason, you might be prompted to install such a driver the first time you use Analyze In Excel on a computer. The connection to Power BI uses claims-based authentication, which is a different technology than the Windows integrated security that you might be using to connect to Analysis Services. You might be prompted to provide a user name and password to connect to Power BI the first time you open such a connection. If you need to connect to different Power BI models using different users, you might need to modify the connections string manually, even if this behavior might change in the future, because the Analyze In Excel feature is in preview mode as of this writing. For updated documentation, go to <https://powerbi.microsoft.com/documentation/powerbi-service-analyze-in-excel/>.


Using Power BI Tiles from Office Store

You can create an Office document in Excel and PowerPoint in which you embed one or more Power BI visualizations. You can create a PowerPoint presentation in which you show live data from Power BI. In a similar way, you can create an Excel workbook in which you embed some visualizations from Power BI on the same page where you also have other data presented with standard Excel tools. This is possible thanks to a free third-party add-in called Power BI Tiles, which takes advantage of the Power BI APIs (these APIs will be explained in more detail later in this chapter). If you are not a developer, you still might be interested in the technical details; you just want to use the existing tool.

You can download Power BI Tiles from the Office Store. It is compatible with Microsoft Office 2013 Service Pack 1, or any following version, including Office 2016. If you have a subscription to Office 365 that includes the licensing of desktop applications, you should already have a compatible version of Office installed on your computer.

Note Power BI Tiles is a free add-in created by DevScope; it is not a Microsoft product, but as of this writing, there are no corresponding solutions produced by Microsoft. You do not need administrative rights to install Power BI Tiles.

Let's visit with David once again, and consider how he can create a PowerPoint presentation that embeds some of the visualizations he created while working on the budget. In PowerPoint, on the ribbon, on the Insert tab, David clicks the Store button, as shown in Figure 8-12.



Click to download Power BI Tiles from the Office Store

Figure 8-12: The Store button on the PowerPoint ribbon.

In the Office Add-Ins dialog box, in the search box, David types "power bi tiles," which presents him with the result depicted in Figure 8-13.




Figure 8-13: A list of available Office add-ins, filtered by the string “power bi tiles” in the search box.

Note If you do not have PowerPoint installed locally and you want to use Office online, you can still use the add-ins available in the Office Store by using the online version of the Office application.

After David installs the add-in, Power BI Tiles becomes available on the PowerPoint ribbon, on the Insert tab, in the My Add-Ins list, as demonstrated in Figure 8-14.

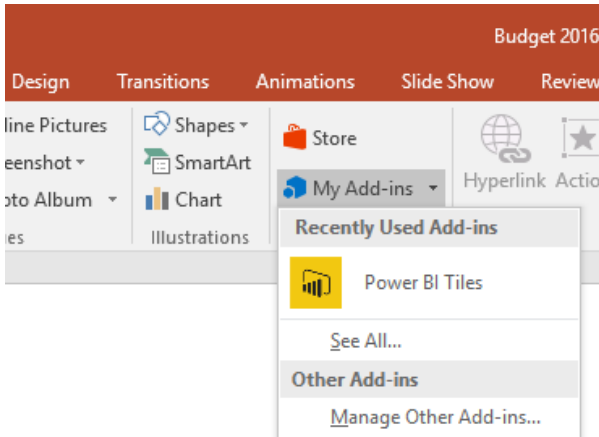


Figure 8-14: The Power BI Tiles add-in on the PowerPoint ribbon in the My Add-Ins list.

When David clicks Power BI Tiles in the list of My Add-Ins, PowerPoint inserts a new rectangular object in the current slide that will display the content of a report or a dashboard. Within this area, you are prompted to choose between connecting to your Power BI account or displaying a public report (Chapter 2 explains how to publish a report to a webpage for public access), as shown in Figure 8-15.

Add **Power BI** Tiles & Reports in your Office Documents



Figure 8-15: You can choose between connecting to a Power BI account or public reports for Power BI Tiles.

The first time David connects to Power BI, he is asked to grant authorization for Power BI Tiles to access certain Power BI features, as shown in Figure 8-16.

Power BI Tiles

App publisher website: localhost

Power BI Tiles needs permission to:

- View all Dashboards (preview) ?
- View all Groups ?
- View all Reports (preview) ?
- View all Datasets ?
- View content properties (preview) ?
- Sign you in and read your profile ?

You're signed in as: david.bradley@contoso-bi.com

[Show details](#)



Figure 8-16: Power BI authorization request for Power BI Tiles access.

David clicks Accept, which authorizes Power BI to accept requests coming from Power BI Tiles. He now can use dashboards and reports available in any Power BI workspaces to which he has access. David can select either a dashboard or a report, and he needs to consider that there are a few differences in the visualizations and user interactions between the two. For example, Figure 8-17 depicts the list of reports available to David in the Power BI Tiles visualization in a PowerPoint slide. In this example, he chose the Budget 2016 group workspace, using the middle button of the three in the upper-right corner of the area used by the Power BI add-in. The list of reports displays the only report published in the workspace: Budget Totals. If you want to see the list of available dashboards in the same workspace, click the Dashboards button, located directly to the left of Reports, above the report list.




Figure 8-17: A Power BI Tiles object inserted in a PowerPoint slide lists the available reports to display.

After David selects a report, it is rendered within the workspace of the PowerPoint slide. If the report size is larger than the available space, scrollbars will appear, as shown in Figure 8-18. The report embedded by Power BI Tiles is fully functional and interactive, so David can use filters and zoom single visualizations of the report, exactly as he can do on the Power BI website. He can even utilize this interaction in Slide Show mode.

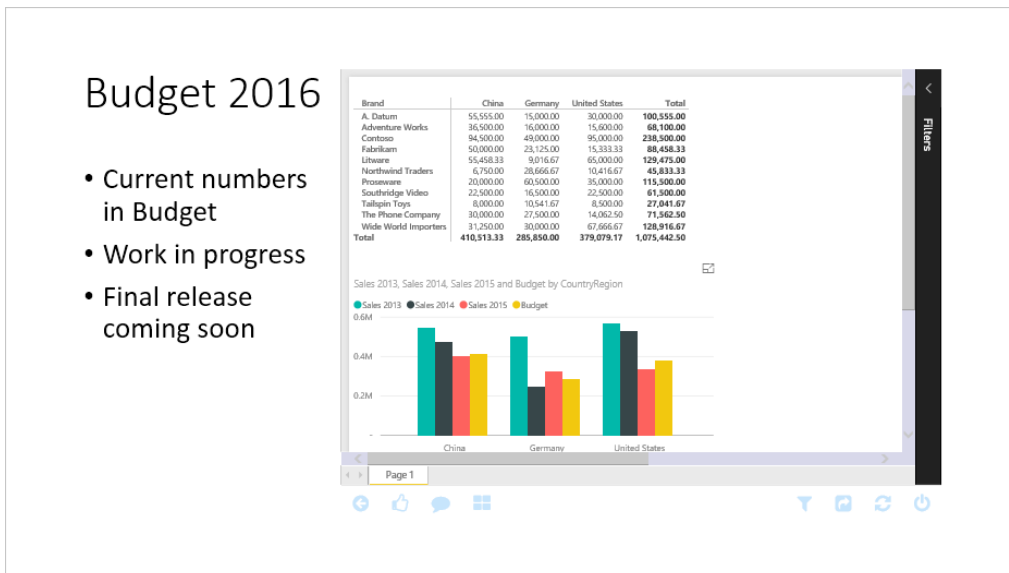


Figure 8-18: A slide in PowerPoint that embeds the content of the Budget Totals report.

You can change the objects displayed in the Power BI Tiles add-in by going back to the list of dashboards and reports (refer back to Figure 8-17). To do that, below the report, click the Back button (the left arrow), which is the one farthest left.

If you select a dashboard, you obtain a slightly different behavior: the Power BI Tiles add-in displays only one visualization from a dashboard at a time. If the dashboard contains two or more visualizations, arrows will appear on the left and right side of each one, which you can click to scroll through the visualizations, as illustrated in Figure 8-19. In this example, only one of the two available visualizations in the Budget Totals dashboard (which, by the way, are the same visualizations used in the report) is visible on the slide. Each visualization takes up the entire amount of space available, which makes them easier to view, but you do not have any interaction with the charts.

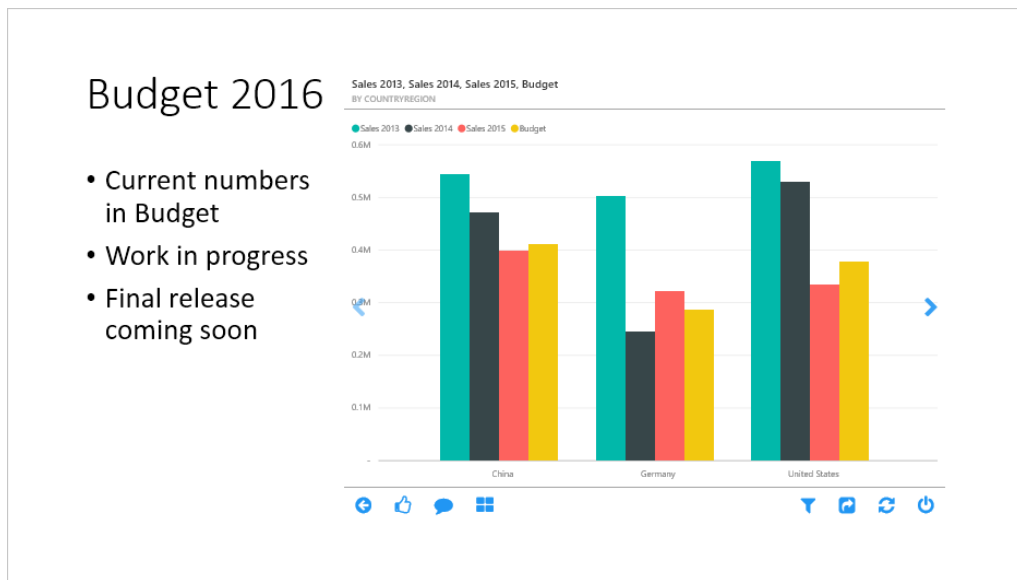


Figure 8-19: A slide in PowerPoint with embedded visualizations of the Budget Totals dashboard.

Whether you choose to embed a dashboard or a report depends on the type of presentation you are creating. If you want to interact with the data, you can modify either choice during the presentation, navigating in the list of dashboards and reports available. However, you should choose the visualization that is more readable and effective for your presentation.

If you want to refresh the data during the presentation, click the Refresh button, which is the second one from the right, below the Power BI Tiles add-in (see Figure 8-19).

You also can use the Power BI Tiles add-in in Excel, creating a worksheet that displays a visualization from Power BI next to data rendered in Excel; for example, using a PivotTable. The features of this add-in work well in Office online, too.

Managing security to access data

The previous chapters of this book demonstrated several ways to share data with other users, within and outside of your organization. In this section, you will review the features available in Power BI to share data with other users and to control access to data in a more granular way, up to the row level in each table.

The following list explains the visibility options available to a Power BI user (other options through APIs are described later in this chapter):

- **My Workspace** All of the datasets, reports, and dashboards you save in My Workspace are visible only to you, unless you explicitly share them by using one of the sharing features that follow.
- **Group Workspace** All of the datasets, reports, and dashboards saved in a group workspace are visible to all the members of the group. The group in Power BI corresponds to a group in Office 365, so you can administer the group from both administrative user interfaces. A new member in the group automatically has access to the data available to the group.
- **Share dashboard** When you share a dashboard, you send a personal invitation to a single individual identified by an email address. This email must correspond to a Power BI account. If the invited user connects to Power BI for the first time, he can create a Power BI sign-in to access data. You can control who has access to a dashboard, adding and removing users from the allowed list at any time. Optionally, you can choose to delegate another user, who can then share the dashboard to other users. When you share a dashboard with a user, you also provide him read-only access to all the underlying reports (from pinned tiles) and datasets used in the dashboard. In fact, users can freely navigate in data sources by starting with a request in the Q&A question box and then customizing filters and visuals, selecting different slices of data and changing the measures displayed and attributes analyzed.
- **Create content pack** You can publish an organizational content pack, sharing datasets, reports, and dashboards, at the granularity that you prefer. When you include a report in a content pack, you also automatically share the underlying datasets. When you include a dashboard in a content pack, you also share the underlying reports and datasets. Those who use an organizational content pack have access to all of the data included in the content pack and can freely navigate to them, creating new reports and dashboards based on that data.
- **Publish to the web** You can publish a report to the web, embedding it within a custom webpage on a website you can edit, or by simply providing a URL containing just the published report. Any user can access that report through this URL, and she can interact with the report using only the visualizations displayed in the report itself. The user cannot gain access to the underlying data source, and cannot modify measures, slicers, filters, and visualizations used in the report. You cannot control who accesses the report, because the URL can be freely shared with anyone; in fact, Microsoft can publish the report in a public gallery. You should not use this feature to distribute reports that contain sensitive data.
- **Row-level security** You can restrict data access for specific users by defining filters at the row level in one or more tables of a dataset. This is an additional security level applied to users who already have access to the data because you shared a dashboard with them. For example, you might share the same dashboard with five different managers, one for each country/region, letting them see only the data of the country/region that they manage.

When you choose the method by which you want to share data with other users, you need to evaluate the visibility you want to provide to reports and datasets, and the type of restrictions you want to apply. In the previous chapters, we describe each of the aforementioned features except the last one, which is the topic of the next section.

Using row-level security

When you want to restrict the rows visible to a single user, you must apply a security rule to the dataset, so that regardless of the report displayed or edited, the user cannot access data that he is not

allowed to see. For example, the managers of China or Europe should see only data relevant to their area, even if all of them use the same report. This type of security is known as *row-level security*.

If you are using a live connection to Analysis Services or you created a model by using DirectQuery, you must implement row-level security on the source database, and you cannot modify its behavior in Power BI.

If you have a model that imported data in the Power BI service, you can apply row-level security to the dataset. You can manage row-level security by selecting the Security action available for datasets, as you can see in Figure 8-20.

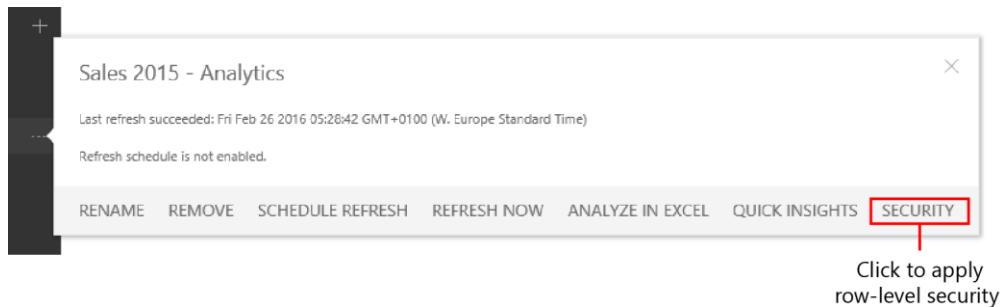


Figure 8-20: Selecting Security to activate row-level security on a dataset.

In the row-level security configuration, you create one or more security roles, which define the rows a user can see in each table. You can find a step-by-step guide to configuring row-level security at <https://powerbi.microsoft.com/documentation/powerbi-admin-rls/>.

Note As of this writing, this feature is in preview, and details might change very quickly. You might want to refer to the online documentation to see if there have been any updates to the user interface.

Figure 8-21 shows the final result of a security role (named China) providing access to only those rows corresponding to sales made in China.




Figure 8-21: The configuration for row-level security for limiting access to sales made only in China.

Each role contains one or more members; these are the users who can access the model through the role. The rules defined for the role are in the form of a DAX expression for each table. A row in a table is visible if the condition, for that row, is true. If a user belongs to more than one role, he will have access to all the rows that are visible in at least one of his roles. However, if a user does not belong to any security role and the dataset has row-level security active, he will not see any data for that

dataset, regardless of whether he can access the dashboard containing data from that dataset because it has been shared by another user.

When you restrict access to a table that has a one-to-many relationship with other tables, you restrict access to the related tables, too. Consider a model with two tables: Customers and Sales. Applying a security rule to Customers also restricts Sales, showing only the sales related to a visible customer.

Let's take a look at what happens when David creates the row-level security rule shown in Figure 8-21, assigning Wendy Kahn as a member of the role, and then he shares with her the dashboard depicted in Figure 8-22.




Figure 8-22: A dashboard shared by David, showing China, Germany, and the United States.

When Wendy opens the same dashboard shared by David, she sees only China; Germany and the United States are not visible, as illustrated in Figure 8-23.

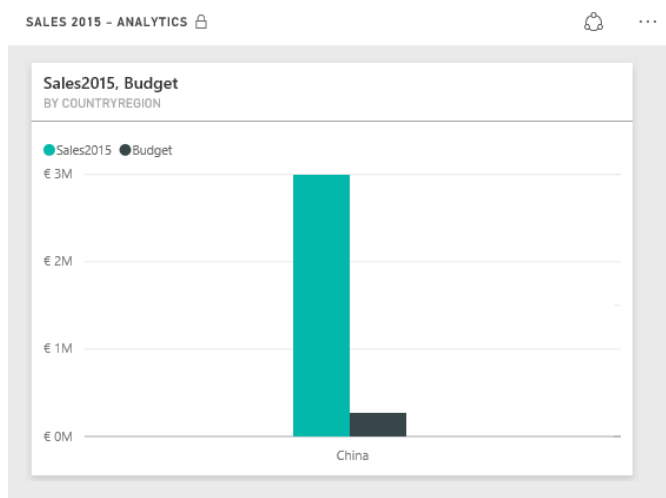


Figure 8-23: A dashboard displayed to Wendy, showing only China.

The security restrictions applied to the dashboard Wendy received from David are also applied to any other visualization shown to Wendy and are based on the same dataset. By applying security at the row level, you can easily customize the aggregations visible to each user. However, keep in mind that

you cannot prevent a user from viewing a particular table, column, or measure. The row-level security filters rows, not columns or other entities of a certain dataset. When you share a dashboard and its reports, in theory the user consuming the data cannot gain access to other entities (measures, columns) that were not published in the report, but this is not guaranteed by the row-level security filter and cannot be enforced at the dataset level. For this reason, if you need to ensure that certain measures are not visible to a group of users, you should consider creating a separate model for that, including only columns and measures that can be made visible to all of the users who can access a certain dataset.

Note As of this writing, the row-level security feature is in preview. It has a number of restrictions that might be lifted or removed in subsequent releases. As of now, you can apply row-level security only in datasets included in My Workspace, but not in group workspaces and not in datasets published in content packs. You can add only single users as role members, not user groups or distribution lists. You can apply it only to datasets created by using Power BI Desktop, not to datasets created with Power Pivot for Excel (but you can import such a model in Power BI Desktop and then publish the Power BI Desktop file). However, when you publish a new version of the Power BI Desktop file, all the existing security roles are removed completely. DirectQuery is not supported for row-level security. Q&A and Cortana are not supported by row-level security, so Q&A input is not visible if role-level security is active for all the models related to a shared dashboard.

Extending and customizing Power BI

Power BI is not only a service that you can activate and use. Likewise, it is not only a product (Power BI Desktop) that you can download and install. You can extend and customize Power BI in many ways, because Power BI offers a number of extensibility points to developers who are interested in adding features, customizing the experience, or integrating existing applications with Power BI.

You can find a more detailed introduction oriented to developers at <https://powerbi.microsoft.com/documentation/powerbi-developer-overview-of-power-bi-rest-api/>. The goal of this section is to introduce you to what is possible and what the current limitations are. In this way, you will have a better understanding of the platform before looking for developers who might help you in this effort, if you do not have the required skills but are interested in achieving the results.

Creating custom visualizations for Power BI

Chapter 7 describes how to add custom visualizations in a Power BI report. As a user, you are likely interested in using existing custom visualizations, and you can find a public gallery of them at <https://visuals.powerbi.com>. If you need specific visualizations that are not available in the gallery, you can (or ask a programmer to) create a new visualization, following the guide available at <https://powerbi.microsoft.com/documentation/powerbi-custom-visuals/>. The main skills required to create a custom visualization are TypeScript (a typed version of JavaScript) and CSS. Thus, if you have programming skills in JavaScript and CSS, you will have a short learning curve to become proficient in writing code for a custom visualization. To become inspired to create new visualizations in Power BI, take a look at the examples of custom visualizations in Chapter 7.

Introducing the Power BI REST API

Previously in this chapter, we showed you how to use the Power BI Tiles add-in. Recall that this component interacts with the data and services provided by Power BI using a programming interface called REST API. In this book, we do not want to go into the details of a REST API or how to use it, but

you can find these details in the online documentation. Most of the information there is useful to developers who want to integrate Power BI services in their applications. The goal here is to explain the importance of this API and why it is the foundation of an ecosystem that makes the integration between Power BI and other applications and services possible, going beyond the features currently available to Power BI end users.

Using the REST API for Power BI, a developer can create a new application or extend an existing one so that it can publish or consume data, reports, and dashboards in Power BI. REST stands for *Representational State Transfer*, which is a protocol that allows any existing programming language to interact with the API, and it is widely adopted in modern programming platforms. REST facilitates interoperability across different languages and operating systems. There are particular specifications to handle authentication and authorization, and you can find all of the details and many examples at <https://msdn.microsoft.com/library/dn877544.aspx>.

Besides the REST protocol, the API exposes several features that can manipulate the following entities:

- **Dataset** You can create new datasets and read existing ones.
- **Table** You can create new tables and modify the schema of existing ones that you created before using the same API. You cannot modify tables that are a part of a data model that imports data from external data sources.
- **Row** You can add and delete rows in tables that you created in a dataset. The delete operation removes all the rows, and you cannot specify any filter, whereas the add operation works incrementally, adding new rows to the existing ones.
- **Group** You can access a particular group to create a dataset in a group instead of a personal workspace.
- **Import** You can import a Power BI Desktop model (.pbix file) or a Power Pivot for Excel data model (.xlsx file) in Power BI.
- **Dashboard** You can retrieve dashboards and tiles from dashboards from a particular workspace to which you have access. The Power BI Tiles add-in, for example, uses this API to retrieve the selected visualization from a dashboard.
- **Report** You can retrieve reports from a particular workspace to which you have access.

You can integrate the visualizations and report objects that you can access through the API in an existing application (this is called “embedding an IFrame”), which is actually what the Power BI add-in does. In the following sections, you will see two examples of applications that are possible thanks to the Power BI REST API.

Limitations

It is worth mentioning the current limitations, considering that the API will evolve and new features will be added, hopefully also to cover some of the scenarios that are not available today. In general, you cannot manipulate the content of a single object. For example, you cannot alter a published dataset, dashboard, or report. You cannot create a report or a dashboard programmatically. You can publish a Power BI Desktop file (.pbix), which can include particular reports, but you do not have an API to create a report or a .pbix file programmatically. This is a current limitation for the Power BI embedded scenario that is described in the next section; but keep in mind that it is still in preview as of this writing, and new API features certainly will be added in the future. Thus, you’ll probably want to

look at the updated documentation to check whether new features have been added to overcome the limitations described here.

By using the Power BI REST API, you can extend existing applications, integrating features available in the Power BI service. This API also opens myriad possibilities to third-party vendors to create components, applications, and services that extend the features available in Power BI.

Pushing real-time data to Power BI dashboards

One of the features available by using the Power BI API is the ability to “push” data in a dashboard in real time. By using this feature, the numbers and visualizations included in a dashboard are automatically updated almost every second, reflecting the changes received in the data. However, the datasets underlying these dashboards are designed in a particular way, and in this section, we want to give you an overall view of the features and limitations of this technique. If you’re a developer who is interested in creating these dashboards, a complete walkthrough is available at <https://powerbi.microsoft.com/documentation/powerbi-developer-walkthrough-push-data/>.

To obtain a real-time dashboard, you first need to programmatically create a dataset. Then, you need to build reports using this dataset, and finally you can pin report visualizations and Q&A visualizations in the dashboard. These tiles will be automatically updated as soon as the underlying dataset receives new data.

The first step also defines the biggest limitation that exists for real-time dashboards: You must first create a dataset programmatically, and this dataset can have tables and columns that will be filled with data sent by an application. The data refresh is not possible in these datasets. We call this a *push mode*, where an application sends data to Power BI, instead of having Power BI ask for data from the data source (which is the classic *pull mode* used by data refresh). You cannot create the data model by using Power BI Desktop, and you cannot add relationships and measures to the data model. You can obtain only standard aggregations for measures, such as sum, average, count distinct, and other predefined ones, but you cannot create either calculated columns or measures using custom DAX expressions. For this reason, it is difficult to display percentages and variations obtained by aggregating existing data. The dataset can be created in a personal workspace or in a group workspace.

The second step, which still requires that you or an application developer write custom code, is to insert rows in the tables of the dataset using the Power BI REST API. Every table has a limit of 5,000,000 rows, or 200,000 rows if you choose a storage model (also known as FIFO dataset) that automatically removes the older rows. The amount of rows written and the frequency of update depends entirely on the application that “pushes” data into the dataset in Power BI. A few limitations exist, based on the Power BI plan used (10,000 rows per hour for the free service plan, and 1,000,000 rows per hour for the paid service plan, Power BI Pro). You can write this code specifically for a single dataset, or you can take advantage of the Azure Stream Analytics service, which simplifies using the same stream of data in different datasets, as described at <https://azure.microsoft.com/documentation/articles/stream-analytics-power-bi-dashboard/>.

After a dataset is created and populated with data, you can create a report using Power BI online. You cannot use Power BI Desktop for such an operation. You can use all the visualizations and filters available in a report to create your visualization. A report is not updated in real time; you must always manually refresh it to display the updated data. You can pin every visualization used in the report to a dashboard visualization, and when you establish the connection between the dashboard and the

dataset (by pinning the first visualization), you can also begin using Q&A to navigate the data and to obtain other visuals that you can pin to the dashboard. All the tiles in the dashboard that are connected to a dataset that receive data in push mode, automatically refresh their content as soon as new data is received. You might observe a refresh almost every second, and the latency between updates of data and visualizations is typically only a few seconds.

Figure 8-24 presents an example of the real-time dashboard we created to monitor the usage of the DAX Formatter service, which is available at www.daxformatter.com. Every time a user formats a DAX query using this service, the application updates the dataset on Power BI, providing the date and time of the request and a flag that specifies whether the request was formatted correctly (increasing the Formatted counter) or contained a syntax error (increasing the Errors counter).

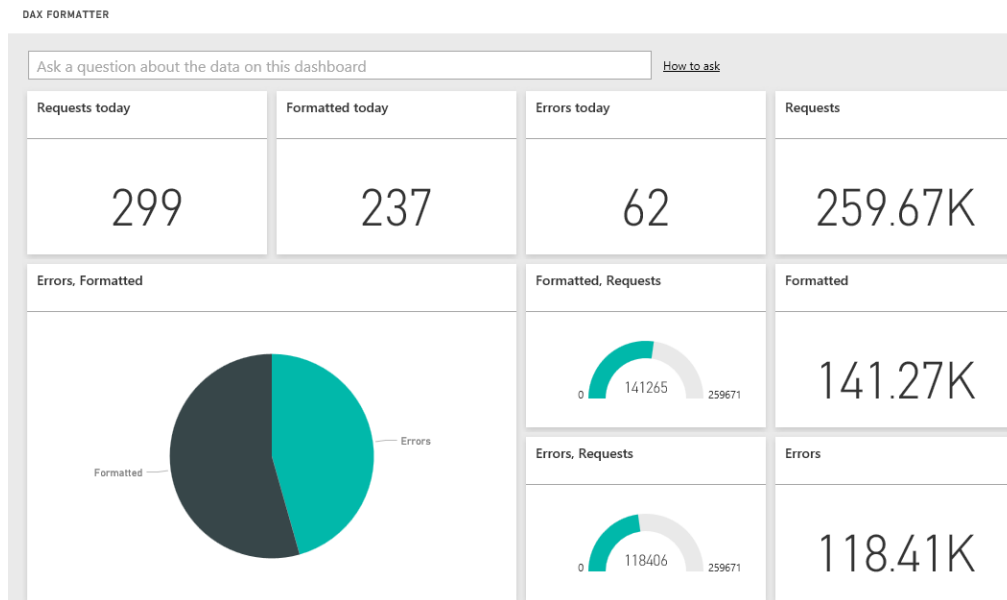



Figure 8-24: A dashboard updated in real time with data from www.daxformatter.com.

It is interesting to note that using Q&A is important to obtain certain dynamic filters, which are impossible to build in a report because you cannot create custom measures in DAX in this type of dataset. For example, the number of requests made today can be obtained by asking the question through Q&A, as illustrated in Figure 8-25. The word “today” is converted to the current date and is applied as a filter to the date column, returning the number of requests made in the current day. Note also that the number of requests in Figure 8-25 increased from Figure 8-24 because data is continuously updated in the dataset!



Source: DAX Formatter

Figure 8-25: Request made through Q&A for the number of requests made today.

Power BI embedded in applications

Another extension available through the Power BI REST API and specific libraries to manage authentication for custom application is Power BI Embedded, which is an Azure service with which you can set up integration between an application and Power BI services.

For example, consider a service that collects data about personal bicycle trips. This service has a web application to manage the configuration and manual upload of data, even if most of the information will be sent by specific devices and/or apps. When it comes time to analyze data, it would be nice for this application to use Power BI services. One way to obtain this integration is to export data to Power BI and create custom data models and reports. For the company that provides this service, it could be a good idea to create a service content pack to make it easy for existing Power BI users to import their data. However, this approach requires all users to create their own Power BI account (even if it is a free one), whereas the application instead should have some embedded solution to display standard reports containing personal data. By using Power BI Embedded, the developers who build the web application can design these reports and publish them within their application in a seamless way. In this way, the web application can show a report containing data of the user who signed in within the same webpage, and without requiring additional authentication to the user itself. The process is completely transparent to the user.

More info You can find documentation for Power BI Embedded and its pricing details at <https://azure.microsoft.com/services/power-bi-embedded/>.

Note It is worth to mention that we showed you another type of web publishing in Chapter 2, the Publish To Web feature. The main difference between Power BI Embedded and Publish To Web is that the former controls the authentication of the user and can display customized content within the same report, whereas the latter only shows the same content to anonymous users.

Conclusions

In this chapter, you learned how you can use Power BI in your company, what the architectural implications of the many options available for data refresh are, how to manage security, and what the options are to customize Power BI and/or integrate it with existing applications.

Here are the most important features you learned:

- Available options to update data with scheduled refresh or live connections
- Integration of Power BI with Microsoft Office
- Control data access for specific users with row-level security
- Possible extensibility options using the Power BI REST API

Power BI is an open ecosystem that is constantly growing, thanks to the features added by Microsoft and those additional options provided by third-party groups, which use the same API you can use to customize and extend Power BI according to your specific needs.

About the authors



Marco Russo and **Alberto Ferrari** are the founders of sqlbi.com, where they regularly publish articles about Microsoft Power BI, Power Pivot, DAX, and SQL Server Analysis Services. They have worked with DAX since the first beta version of Power Pivot in 2009, and, during these years, sqlbi.com became one of the major sources for DAX articles and tutorials.

They both provide consultancy and mentoring on business intelligence (BI), with a particular specialization in the Microsoft technologies related to BI. They have written several books and papers about Power Pivot, DAX, and Analysis Services. They also wrote popular white papers such as “The Many-to-Many Revolution” (about modeling patterns using many-to-many relationships) and “Using Tabular Models in a Large-Scale Commercial Solution” (a case study of Analysis Services adoption published by Microsoft). Marco and Alberto are also regular speakers at major international conferences, including Microsoft Ignite, PASS Summit, and SQLBits.

You can contact Marco at marco.russo@sqlbi.com, and Alberto at alberto.ferrari@sqlbi.com.



From technical overviews to drilldowns on special topics, get *free* ebooks from Microsoft Press at:

www.microsoftvirtualacademy.com/ebooks

Download your free ebooks in PDF, EPUB, and/or Mobi for Kindle formats.

Look for other great resources at Microsoft Virtual Academy, where you can learn new skills and help advance your career with free Microsoft training delivered by experts.

Microsoft Press