

Chapter 1

Introducing Power BI

1.1 What is Microsoft Power BI? 1

1.2 Understanding the Power BI Capabilities 15

1.3 Understanding the Power BI Service Architecture 28

1.4 Power BI and You 31

1.5 Summary 38

Without supporting data, you are just another person with an opinion. But data is useless if you can't derive knowledge from it. And this is where Microsoft data analytics and Power BI can help! Power BI changes the way you gain insights from data; it brings you a cloud-hosted, business intelligence and analytics platform that democratizes and opens BI to everyone. Power BI makes data analytics pervasive and accessible to all users under a simple promise: "five seconds to sign up, five minutes to wow!"

This guide discusses the capabilities of Power BI, and this chapter introduces its innovative features. I'll start by explaining how Power BI fits into the Microsoft Data Platform and when to use it. You'll learn what Power BI can do for different types of users, including business users, data analysts, professionals, and developers. I'll also take you on a tour of the Power BI features and its toolset.

1.1 What is Microsoft Power BI?

Before I show you what Power BI is, I'll explain business intelligence (BI). You'll probably be surprised to learn that even BI professionals disagree about its definition. In fact, Forester Research offers two definitions (see https://en.wikipedia.org/wiki/Business_intelligence).



DEFINITION Broadly defined, BI is a set of methodologies, processes, architectures, and technologies that transform raw data into meaningful and useful information that's used to enable more effective strategic, tactical, and operational insights and decision-making. A narrower definition of BI might refer to just the top layers of the BI architectural stack, such as reporting, analytics, and dashboards.

Regardless of which definition you follow, Power BI can help you with your data analytics needs.

1.1.1 Understanding Business Intelligence

The definition above is a good starting point. but to understand BI better, you need to understand its flavors. First, I'll categorize who's producing the BI artifacts, and then I'll show you the different types of analytical tasks that these producers perform.

Self-service, team, and organizational BI

I'll classify BI by its main users and produced artifacts and divide it into self-service, team, and organizational BI.

- Self-service BI (or personal BI) – Self-service BI enables data analysts to offload effort from IT pros. For example, Maya is a business user and she wants to analyze CRM data from Salesforce. Maya can connect Power BI to Salesforce and get prepackaged dashboards and reports without

building a data model. In the more advanced scenario, Power BI empowers analysts to build data models for self-service data exploration and reporting. Suppose that Martin from the sales department wants to analyze some sales data that's stored in the corporate data warehouse and mash it up with some external data. With a few clicks, Martin can combine multiple tables from various data sources into a data model (like the one shown in **Figure 1.1**), build reports, and gain valuable insights. In other words, Power BI makes data analytics more pervasive because it enables more employees to perform BI tasks.



Figure 1.1 Power BI allows analysts to build data models whose features are on par with professional models.

- **Team BI** – Business users can share the reports and dashboards they've implemented with other team members without requiring them to install modeling or reporting tools. Suppose that Martin would like to share his sales model with his coworker, Maya. Once Martin has uploaded the model to Power BI, Maya can go online and view the reports and dashboards Martin has shared with her. She can even create her own reports and dashboards that connect to Martin's model.

- Organizational BI (or corporate BI) – BI professionals who implement organizational BI solutions, such as semantic models or real-time business intelligence, will find that they can use Power BI as a presentation layer. For example, as a BI pro, Elena has developed a Multidimensional or Tabular organizational semantic model layered on top of the company's data warehouse that is hosted on her company's network. Elena can install connectivity software called a data gateway on an on-premises computer so that Power BI can connect to her model. This allows business users to create instant reports and dashboards in Power BI by leveraging the existing infrastructure investment without moving data to the cloud! Or, if management asks for a real-time dashboard, Elena can write some code to push data to Power BI so that the dashboard updates itself as new data streams in.

Descriptive, predictive and prescriptive analytics

The main goal of BI is to get actionable insights that lead to smarter decisions and better business outcomes. Another way to classify BI is from a time perspective. Then we can identify three types of data analytics (descriptive, predictive, and prescriptive).

Descriptive analytics is retrospective. It focuses on what has happened in the past to understand the company's performance. This type of analytics is the most common and well understood. Coupled with a good data exploration tool, such as Power BI or Microsoft Excel, descriptive analytics helps you discover import trends and understand the factors that influenced these trends. You do descriptive analytics when you slice and dice data. For example, a business analyst can create a Power BI report to discover sale trends by year. Descriptive analytics can answer questions, such as "Who are my top 10 customers?", "What is the company's sales by year, quarter, month, and so on?", or "How does the company's profit compare against the predefined goal by business unit, product, time, and other subject areas?"

Predictive analytics is concerned with what will happen in the future. It uses machine learning algorithms to determine probable future outcomes and discover patterns that might not be easily discernible based on historical data. These hidden patterns can't be discovered with traditional data exploration since data relationships might be too complex, or because there's too much data for a human to analyze. Typical predictive tasks include forecasting, customer profiling, and basket analysis. Machine learning can answer questions, such as, "What are the forecasted sales numbers for the next few months?" and "What other products is a customer likely to buy along with the product he or she already chose?". Power BI includes many predictive features. Quick Insights applies machine learning algorithms to find hidden patterns, such as that the revenue for a product is steadily decreasing. You can use the Power BI clustering algorithms to quickly find groups of similar data points in a subset of data. Or, apply time-series forecasting to a line chart to predict sales for future periods or anomaly detection to discover outliers.

Addressing more involved requirements, a data analyst can build a self-service ML model in Power BI Service. Thanks to the huge investments that Microsoft has made in open-source software, a data analyst can also use R or Python scripts for data cleansing, statistical analysis, data mining, and visualizing data. Power BI can integrate with Azure Machine Learning experiments. For example, an analyst can build a predictive experiment with the Azure Machine Learning service and then visualize the results in Power BI. Or, if a BI pro has implemented a predictive model in R or Python and deployed to SQL Server, the analyst can simply query SQL Server to obtain the predictions.

Finally, *prescriptive analytics* goes beyond predictive analytics to not only attempt to predict the future but also recommend the best course of action and the implications of each decision option. Typical prescriptive tasks are optimization, simulation, and goal seek. While tools for descriptive and predictive needs have matured, prescriptive analytics is a newcomer and currently is in the realm of startup companies. Power BI includes certain features that might help, such as what-if analysis for simulation, Decomposition Tree visual that can automatically suggest the next dimension to drill down into, and key influencers that can help you find dimensions that correlate the most to a certain goal, such as increased revenue.


1.1.2 Introducing the Power BI Products

Now that you understand BI better, let's discuss what Power BI is. Power BI is a set of products and services that enable you to connect to your data, visualize it, and share insights with other users. Next, I'll introduce you to the Power BI product offerings.

What's behind the Power BI name?

At a high level, Power BI consists of several products (listed in the order they appear in the Products menu on the powerbi.com home page):

- **Power BI Desktop** – A freely available Windows desktop application that allows analysts to design self-service data models and for creating interactive reports connected to these models or to external data sources. For readers familiar with Power Pivot for Excel, Power BI Desktop offers similar self-service BI features in a standalone application (outside Excel) that updates every month.
- **Power BI Pro** – Power BI Pro is one of the licensing options of Power BI Service (the other two are Power BI Free and Power BI Premium). Power BI Service is a *cloud-based* business analytics service (powerbi.com) that allows you to host your data, reports, and dashboards online and share them with your coworkers. Because Power BI is hosted in the cloud and managed by Microsoft, your organization doesn't have to purchase, install, and maintain an on-premises infrastructure.
- **Power BI Premium** – Targeting large organizations, Power BI Premium offers a dedicated capacity environment, giving your organization more consistent performance without requiring you to purchase per-user licenses. Suppose you want to share reports with more than 500 users within your organizations and most of these users require read-only access. Instead of licensing each user, you could reduce cost by purchasing a Power BI Premium plan that doesn't require licenses for viewers and gives you predictable performance. Power BI Premium also adds features that are not available in Power BI Pro, such as larger dataset sizes and incremental data refresh.
- **Power BI Mobile** – A set of freely available mobile applications for iOS, Android, and Windows that allow users to use mobile devices, such as tablets and smartphones, to get data insights on the go. For example, a mobile user can view reports and dashboards deployed to Power BI.
- **Power BI Embedded** – Power BI Embedded is a collective name for a subset of the Power BI APIs for embedding content. Integrated with Power BI Service, Power BI Embedded lets developers embed interactive Power BI reports in custom apps for internal or external users. For example, Teo has developed a web application for external customers. Instead of redirecting to powerbi.com, Teo can use Power BI Embedded to let customers view interactive Power BI reports embedded in his app.
- **Power BI Report Server** – Evolving from Microsoft SQL Server Reporting Services (SSRS), Power BI Report Server allows you to deploy Power BI data models and reports to an on-premises server. This gives you a choice for deployment and sharing: cloud and/or on-premises. And the choice doesn't have to be exclusive. For example, you might decide to deploy some reports to Power BI to leverage all features it has to offer, such as natural queries, quick insights, and integration with Excel, while deploying the rest of the reports to a Power BI Report Server portal.

 **DEFINITION** Microsoft Power BI is a data analytics **platform** for self-service, team, and organizational BI that consists of several products. Although Power BI can access other Office 365 services, such as OneDrive and SharePoint, Power BI doesn't require an Office 365 subscription and it has no dependencies to Office 365. However, if your organization is on Office 365 E5 plan, you'll find that Power BI Pro is included in it.

Product usage scenarios

The Power BI product line has grown over time and a novice Power BI user might find it difficult to understand where each product mentioned above fits in. **Figure 1.2** should help you visualize the purpose of each product at a high level.

1. Power BI Desktop – The self-service BI journey typically starts with Power BI Desktop. As a data analyst, you can use Power BI Desktop to mash up data from various data sources and create a self-service data model. Or, you can use Power BI Desktop to connect directly to a data source, such as an organizational semantic model, and start analyzing data immediately without importing data and do any modeling.

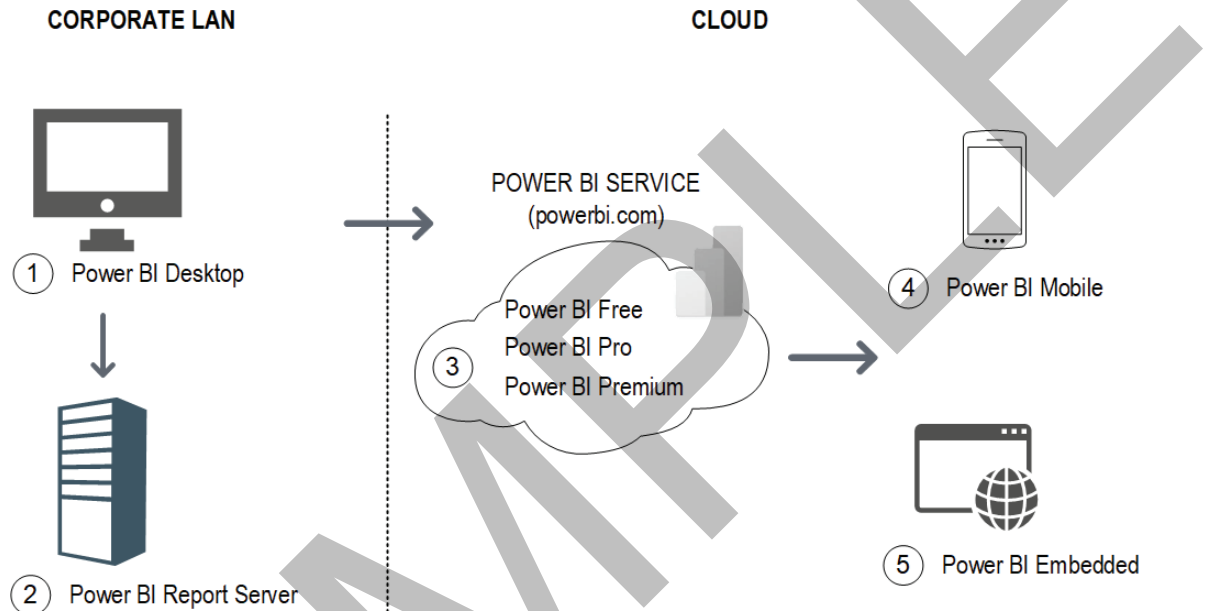


Figure 1.2 How Power BI products can be used for different tasks.

2. Power BI Report Server – One option to share your Power BI artifacts is to deploy them to on-premises Power BI Report Server. This is a good option if your organization needs an on-premises report portal that hosts not only Power BI reports but also operational SSRS reports and Excel reports, without requiring all Power BI features. However, Power BI Report Server lags in features compared to Power BI Service and is limited to only deploying and viewing Power BI reports online.
3. Power BI Service – Another sharing option is to deploy your data models and reports to the cloud Power BI Service (powerbi.com). Power BI Service has four licensing options:
 - Power BI Free – Any user can use Power BI service for personal data analytics for free! However, they can't share BI artifacts with other users.
 - Power BI Pro – Requiring per-user licensing, Power BI Pro gives you most of the Power BI Service features, including sharing BI content.
 - Power BI Premium – To avoid licensing per user for many users who will only view reports, a larger organization might decide to purchase a Power BI Premium plan. Besides cost savings, Power BI Premium is appealing from a performance standpoint as it offers a dedicated environment just for your organization and adds even more features.

- Premium per User – Targeting smaller organizations that can't afford a high monthly commitment, this option brings you premium features but retains licensing per user without having to provision and manage capacities.
4. Power BI Mobile – Although Power BI reports can render in any modern browser, your mobile workforce can install the Power BI Mobile apps on their mobile devices so that Power BI reports are optimized for the display capabilities of the device.
 5. Power BI Embedded – A developer can integrate a custom web app with Power BI Embedded to embed Power BI reports, so they render inside the app. Organizations typically use Power BI Embedded to provide reports for a third party, such as their external customers.

As you could imagine, Power BI is a versatile platform that enables different groups of users to implement a wide range of BI solutions depending on the task at hand.

1.1.3 How Did We Get Here?

Before I delve into the Power BI capabilities, let's step back for a moment and review what events led to its existence. **Figure 1.3** shows the major milestones in the Power BI journey.

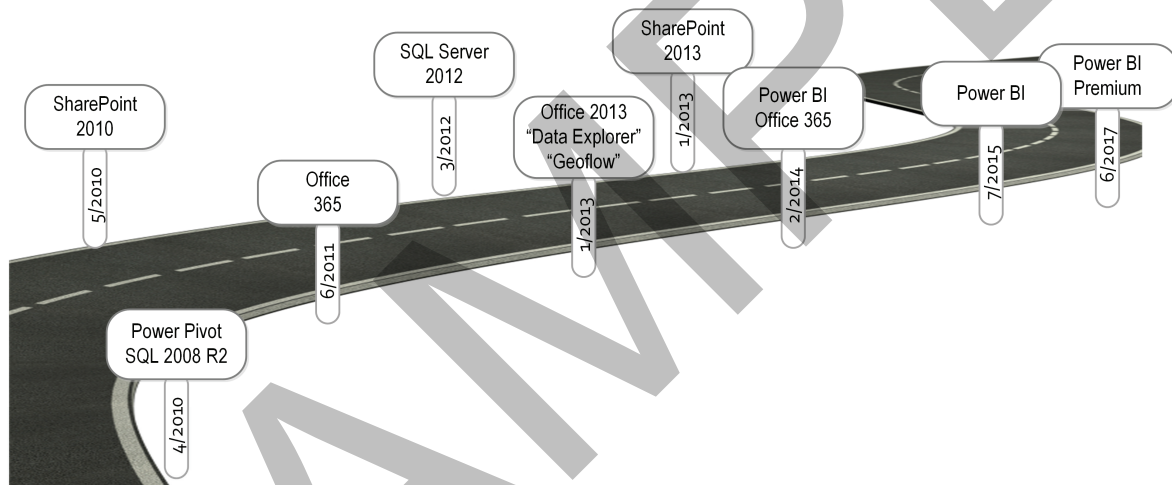


Figure 1.3 Important milestones related to Power BI.

Power Pivot

Realizing the growing importance of self-service BI, in 2010 Microsoft introduced a new technology for personal and team BI called PowerPivot (renamed to Power Pivot in 2013). Power Pivot was initially implemented as a freely available add-in to Excel 2010 that had to be manually downloaded and installed. Office 2013 delivered deeper integration with Power Pivot, including distributing it with Excel 2013 and allowing users to import data directly into the Power Pivot data model.

NOTE I covered Excel and Power Pivot data modelling in my book "Applied Microsoft SQL Server 2012 Analysis Services: Tabular Modeling". If you prefer using Excel for self-service BI, the book should give you the necessary foundation to understand Power Pivot and learn how to use it to implement self-service data models. However, since the Microsoft premium tool for data analytics is now Power BI, I recommend you use Power BI Desktop instead. Just remember that many of the Power BI Desktop features are also available in Excel Power Pivot and Power Query. For example, instead of VLOOKUP, you can use Power Query to lookup values more efficiently from another spreadsheet or file.



The Power Pivot innovative engine, called xVelocity (initially named VertiPaq), transcended the limitations of the Excel native pivot reports. It allows users to load multiple datasets and import more than one million rows (the maximum number of rows that can fit in an Excel spreadsheet). xVelocity compresses the data efficiently and stores it in the computer's main memory.



DEFINITION xVelocity is a columnar data engine that compresses and stores data in memory. Originally introduced in Power Pivot, the xVelocity data engine has a very important role in Microsoft BI. xVelocity is now included in other Microsoft offerings, including SQL Server columnstore indexes, Tabular models in Analysis Services, Power BI Desktop, and Power BI.

For example, using Power Pivot, a business user can import data from a variety of data sources, relate the data, and create a data model. Then the user can create pivot reports or Power View reports to gain insights from the data model.

SQL Server

Originally developed as a relational database management system (RDBMS), Microsoft SQL Server is now a multi-product offering. In the context of organizational BI, SQL Server includes Analysis Services, which has traditionally allowed BI professionals to implement multidimensional cubes. SQL Server 2012 introduced another path for implementing organizational models called Tabular. Think of Analysis Services Tabular as Power Pivot on steroids. Just like Power Pivot, Tabular allows you to create in-memory data models but it also adds security and performance features to allow BI pros to scale these models and implement data security that is more granular.

SQL Server used to include Reporting Services (SSRS), which has been traditionally used to implement paper-oriented standard reports (also referred to as paginated reports) and it's now available as a separate download. SQL Server 2012 introduced a SharePoint 2010-integrated reporting tool, named Power View, for authoring ad hoc interactive reports. Power View targets business users without requiring query knowledge and report authoring experience. Suppose that Martin has uploaded his Power Pivot model to SharePoint Server. Now Maya (or anyone else who has access to the model) can quickly build a great-looking tabular or chart report in a few minutes to visualize the data from the Power Pivot model. Or, Maya can use Power View to explore data in a Multidimensional or Tabular organizational model. Microsoft used some of the Power View features in the Power BI reports.

In Office 2013, Microsoft integrated Power View with Excel 2013 to allow business users to create interactive reports from Power Pivot models and organizational Tabular models. And Excel 2016 extended Power View to connect to multidimensional cubes. Microsoft later deprecated Power View (it's disabled by default in Excel 2016) to encourage users to transition to Power BI Desktop, which is now the Microsoft premium data exploration tool.

SharePoint Server

Up to the release of Power BI, Microsoft BI has been intertwined with SharePoint. SharePoint Server is a Microsoft on-premises product for document storage and collaboration. In SharePoint Server 2010, Microsoft added new services, collectively referred to as Power Pivot for SharePoint, which allowed users to deploy Power Pivot data models to SharePoint and then share reports that connect to these data models. For example, a business user can upload the Excel file containing a data model and reports to SharePoint. Authorized users can view the embedded reports and create their own reports.

SharePoint Server 2013 brought better integration with Power Pivot and support for data models and reports created in Excel 2013. When integrated with SQL Server 2012, SharePoint Server 2013 offers other compelling BI features, including deploying and managing SQL Server Reporting Services (SSRS) reports, team BI powered by Power Pivot for SharePoint, and PerformancePoint Services dashboards.

Later, Microsoft realized that SharePoint presents adoption barriers for the fast-paced world of BI. Therefore, Microsoft deemphasized the role of SharePoint as a BI platform in SharePoint Server 2016 in favor of Power BI in the cloud and Power BI Report Server on premises. SharePoint Server can still be integrated with Power Pivot and Reporting Services but it's no longer a strategic on-premises BI platform.

Microsoft Excel

While prior to Power BI, SharePoint Server was the Microsoft premium server-based platform for BI, Microsoft Excel was their premium BI tool on the desktop. Besides Power Pivot and Power View, which I already introduced, Microsoft added other BI-related add-ins to extend the Excel data analytics features. To help end users perform predictive tasks in Excel, Microsoft released a Data Mining add-in for Microsoft Excel 2007, which is also available with newer Excel versions. For example, you could use this add-in to perform a market basket analysis and find which products customers tend to buy together.



NOTE In 2014, Microsoft introduced a cloud-based Azure Machine Learning Service (<https://ml.azure.com/>) to allow users to create predictive models in the cloud, such as a model that predicts the customer churn probability. SQL Server 2016 added integration with R and SQL Server 2017 added integration with Python. Azure ML in the cloud and R and Python on premise supersede the Data Mining add-in for self-service predictive analytics and Analysis Services data mining for organizational predictive analytics. It's unlikely that we'll see future Microsoft investments in these two technologies.

In January 2013, Microsoft introduced a freely available Data Explorer add-in for Excel, which was later renamed to Power Query. Power Query is now included in Excel and Power BI Desktop. Unique in the self-service BI tools market, Power Query allows business users to transform and cleanse data before it's imported. For example, Martin can use Power Query to replace wrong values in the source data or to unpivot a crosstab report. In Excel, Power Query is an optional path for importing data. If data doesn't require transformation, a business user can directly import the data using the Excel or Power Pivot data import capabilities. However, Power BI always uses Power Query when you import data so that its data transformation capabilities are there if you need them.

The screenshot shows the Microsoft Excel interface with the Power Query ribbon active. The ribbon includes tabs for File, Home, Transform, Add Column, View, Tools, and Help. The Transform tab is selected, showing options like Close & Apply, New Source, Recent Sources, Enter Data, Data source settings, Manage Parameters, Refresh Preview, Properties, Advanced Editor, Manage Columns, Reduce Rows, Sort, Split Column, Group By, Data Type: Text, Use First Row as Headers, Replace Values, and Combine. Below the ribbon, the Queries list on the left shows 'VendorParts' selected. The main area displays a table with columns 'Category', 'Manufacturer', and 'Month'. The table contains 15 rows of data, all with 'Wheels' in the Category column. The 'Manufacturer' column has 'Acme 1' for rows 1-12 and 'Acme 2' for rows 13-15. The 'Month' column lists months from Jan to Feb. The Query Settings pane on the right shows the 'APPLIED STEPS' section, which includes: Source, Navigation, Changed Type, Filtered Rows, Filtered Rows1, Promoted Headers, Filled Down, Filtered Rows2, Filtered Rows3, and Removed Columns. The status bar at the bottom indicates '7 COLUMNS, 872 ROWS' and 'Column profiling based on top 1000 rows'. A watermark 'PREVIEW DOWNLOADED AT 9:44 AM' is visible in the bottom right corner.

	Category	Manufacturer	Month
1	Wheels	Acme 1	Jan
2	Wheels	Acme 1	Feb
3	Wheels	Acme 1	Mar
4	Wheels	Acme 1	Apr
5	Wheels	Acme 1	May
6	Wheels	Acme 1	Jun
7	Wheels	Acme 1	Jul
8	Wheels	Acme 1	Aug
9	Wheels	Acme 1	Sep
10	Wheels	Acme 1	Oct
11	Wheels	Acme 1	Nov
12	Wheels	Acme 1	Dec
13	Wheels	Acme 2	Jan
14	Wheels	Acme 2	Feb
15

Figure 1.4 A data analyst can use Power Query to shape and transform data.

For example, **Figure 1.4** shows how I have applied several steps to cleanse and shape the data in Power Query. Power BI will sequentially apply these steps as the data is being imported from the data source. Power BI dataflows also leverage Power Query for self-service data staging to Azure data lake storage.

Another data analytics add-in that deserves attention is Power Map. Originally named Geoflow, Power Map is another freely available Excel add-in that's specifically designed for geospatial reporting. Power Map is included by default in Excel 2016. Using Power Map, a business user can create interactive 3D maps from Excel tables or Power Pivot data models. Power BI has several mapping visuals and Power Map is not included, but you can get a taste of it when you import the GlobeMap custom visual.

Power BI for Office 365

Unless you live under a rock, you know that one of the most prominent IT trends nowadays is cloud computing. Chances are that your organization is already using the Microsoft Azure Services Platform - a cloud platform for hosting and scaling applications and databases through Microsoft datacenters. Microsoft Azure gives you the ability to focus on your business and to outsource infrastructure maintenance to Microsoft.

In 2011, Microsoft unveiled its Office 365 cloud service to allow organizations to subscribe to and use a variety of Microsoft products online, including Microsoft Exchange and SharePoint. For example, at Prologika we use Office 365 for email, a subscription-based (click-to-run) version of Microsoft Office, OneDrive for Business, Microsoft Teams, Dynamics Online, and other products. From a BI standpoint, Office 365 allows business users to deploy Excel workbooks and Power Pivot data models to the cloud. Then they can view the embedded reports online, create new reports, and share BI artifacts.

In early 2014, Microsoft further extended SharePoint for Office 365 with additional BI features, including natural queries (Q&A), searching and discovering organizational datasets, and mobile support for Power View reports. Together with the "power" desktop add-ins (Power Pivot, Power View, Power Query, and Power Map), the service was marketed and sold under the name "Power BI for Office 365". While the desktop add-ins were freely available, Power BI for Office 365 required a subscription. Microsoft sold Power BI for Office 365 independently or as an add-on to Office 365 business plans.

Because of its dependency to SharePoint and Office, Power BI for Office 365 didn't gain wide adoption. One year after unveiling the new Power BI platform, Microsoft discontinued Power BI for Office 365. Power BI for Office 365 shouldn't be confused with the new Power BI platform, which was completely re-architected for agile and modern BI.

Power BI

Finally, the winding road brings us to Power BI, which is the subject of this book. In July 2015, after several months of public preview, Microsoft officially launched a standalone version of the cloud Power BI Service that had no dependencies on Office 365, SharePoint and Microsoft Office. What caused this change? The short answer is removing adoption barriers for both Microsoft and consumers. For Microsoft it became clear that to be competitive in today's fast-paced marketplace, its BI offerings couldn't depend on other product groups and release cycles. Waiting for new product releases on two and three-year cadences couldn't introduce the new features Microsoft needed to compete effectively with "pure" BI vendors (competitors who focus only on BI tools) who have entered the BI market in the past few years.

After more than a decade of working with different BI technologies and many customers, I do believe that Microsoft BI is the best and most comprehensive BI platform on the market! But it's not perfect. One ongoing challenge is coordinating BI features across product groups. Take for example SharePoint, which Microsoft promoted as a platform for sharing BI artifacts. Major effort underwent to extend SharePoint with SSRS in SharePoint integration mode, PerformancePoint, Power Pivot, and so on. But these products are owned by different product groups and apparently coordination has been problematic.

Seeking a stronger motivation for customers to upgrade, Excel added the "power" add-ins and was promoted as the Microsoft premium BI tool on the desktop. However, the Excel dependency turned out to be a double-edged sword. While there could be a billion Excel users worldwide, adding a new feature must be thoroughly tested to ensure that there are no backward compatibility issues or breaking changes, and that takes a lot of time. Case in point: we had to wait almost three years until Excel 2016 was able to

connect Power View reports to multidimensional cubes (only Tabular was supported before), although Analysis Services Multidimensional had a much broader adoption than Tabular.

For consumers, rolling out a Microsoft BI solution has been problematic. Microsoft BI has been traditionally criticized for its deployment complexity and steep price tag. Although SharePoint Server offers much more than just data analytics, having a SharePoint server integrated with SQL Server has been a cost-prohibitive proposition for smaller organizations. As many of you would probably agree, SharePoint Server adds complexity and troubleshooting it isn't for the faint of heart. Power BI for Office 365 alleviated some of these concerns by shifting maintenance to become Microsoft's responsibility, but many customers still find its "everything but the kitchen sink" approach too overwhelming and cost-prohibitive if all they want is the ability to deploy and share BI artifacts.


Going back to the desktop, Excel wasn't originally designed as a BI tool, leaving the end user with the impression that BI was something Microsoft bolted on top of Excel. For example, navigating add-ins and learning how to navigate the cornucopia of features has been too much to ask from novice business users.

How does the new Power BI address these challenges?

Power BI embraces the following design tenets to address the previous pain points:

- **Simplicity** – Power BI was designed for BI from the ground up. As you'll see, Microsoft streamlined and simplified the user interface to ensure that your experience is intuitive, and you aren't distracted by other non-BI features and menus.
- **No dependencies to SharePoint and Office** – Because it doesn't depend on SharePoint and Excel, Power BI can evolve independently. This doesn't mean that business users are now asked to forgo Excel. On the contrary, if you like Excel and prefer to create data models in Excel, you'll find that you can still deploy them to Power BI.
- **Frequent updates** – Microsoft delivers weekly updates for Power BI Service and monthly updates for Power BI Desktop. Hundreds of new features are added every year. This unprecedented speed of delivery allowed Microsoft to stay at the forefront of the BI market (Microsoft is a leader in the Gartner's Magic Quadrant for Analytics & BI Platforms).
- **Always up to date** – Because of its service-based nature, as a Power BI subscriber you're always on the latest and greatest version. In addition, because Power BI is a cloud service, you can get started with Power BI Pro or Premium in a minute as you don't have to provision servers and software.
- **Great value proposition** – As you'll see in "Power BI Editions and Pricing" (later in this chapter), Power BI has the best business model: most of it is free! Power BI Desktop and Power BI Mobile are free. Following a freemium model, Power BI is free for personal use and has subscription options that you could pay for if you need to share with other users. Cost was the biggest hindrance of Power BI, and it's now been turned around completely. You can't beat free!

No tool is a kingdom of its own and no tool should work in isolation. If you're tasked to evaluate BI tools, consider that one prominent strength of Power BI is that it's an integral part of a much broader Microsoft Data Platform that started in early 2004 with the powerful promise to bring "BI to the masses." Microsoft subsequently extended the message to "BI to the masses, by the masses" to emphasize its commitment to democratize. Indeed, a few years after Microsoft got into the BI space, the BI landscape changed dramatically. Once a domain of cost-prohibitive and highly specialized tools, BI is now within the reach of every user and organization!

 **DEFINITION** The Microsoft Data Platform is a multi-service offering that addresses the data capturing, transformation, and analytics needs to create modern BI solutions. It's powered by Microsoft SQL Server on premises and Microsoft Azure in the cloud.

1.1.4 Power BI and the Microsoft Data Platform



Figure 1.5 The Microsoft Data Platform provides services and tools that address various data analytics and management needs on premises and in the cloud.

Understanding the Microsoft Data Platform

Figure 1.5 illustrates the most prominent services of the Microsoft Data Platform. Table 1.1 summarizes the various services of the Microsoft Data Platform and their purposes.

Table 1.1 The Microsoft Data Platform consists of many products and services, with the most prominent described below.


Category	Service	Audience	Purpose
Capture and manage	Relational	IT	Capture relational data in SQL Server, Analytics Platform System, Azure SQL Database, Azure Synapse Analytics, and others.
	Non-relational	IT	Capture Big Data in Azure HDInsight Service and Microsoft HDInsight Server.
	NoSQL	IT	Capture NoSQL data in cloud structures, such as Azure Table Storage, Cosmos DB, and others.
	Streaming	IT	Allow capturing of data streams from Internet of Things (IoT) with Azure StreamInsight.
Transform and analyze	Orchestration	IT/Business	Create data orchestration workflows with SQL Server Integration Services (SSIS), Azure Data Factory, Power Query, Power BI Desktop, and Data Quality Services (DQS).
	Information management	IT/Business	Allow IT to establish rules for information management and data governance using SharePoint, Azure Data Catalog, and Office 365, as well as manage master data using SQL Server Master Data Services.
	Complex event processing	IT	Process data streams with Azure Stream Analytics Service.

Category	Service	Audience	Purpose
Visualize and decide	Modelling	IT/Business	Transform data in semantic structures with Analysis Services Multidimensional, Tabular, Power Pivot, and Power BI.
	Machine learning	IT/Business	Create data mining models in SQL Server Analysis Services, Excel data mining add-in, and Azure Machine Learning Service.
	Cognitive services	IT/Business	Build intelligent algorithms into apps, websites, and bots so that they see, hear, speak, and learn.
	Applications	IT/Business	Analyze data with desktop applications, including Excel, Power BI Desktop, SSRS Designer, Report Builder, Power View, Power Map.
	Reports	IT/Business	Create operational and ad hoc reports with Power BI, SSRS, and Excel.
	Dashboards	IT/Business	Implement and share dashboards with Power BI and SSRS.
	Mobile	IT/Business	View reports and dashboards on mobile devices with Power BI Mobile.
	Power Apps	IT/Business	Implement low code/no code data-driven apps.
	Power Automate	IT/Business	Build data-driven workflows to automate processes.

About Microsoft Power Platform

Yet another way to appreciate the potential of Power BI for addressing your business needs is to consider it in the context of the Microsoft Power Platform (<https://powerplatform.microsoft.com>), which consists of four products:

- Power BI – The subject of this book.
- Power Apps – A tool that helps business users build no-code/low-code apps. Every organization has business automation needs which traditionally have been solved by developers writing custom code. However, just like Power BI democratizes BI, Power Apps empowers business users to create their own apps. Further, Power BI can integrate with Power Apps to redefine the meaning of reports. I demonstrate how you can leverage this integration to change the data behind a Power BI report (report writeback) in Chapter 10.
- Power Automate – Besides apps, automating business processes typically requires a workflow. Power Automate (formerly known as Microsoft Flow) lets you implement no-code/low-code workflows that react to conditions. For example, my company uses Power Automate to monitor leads posted to Dynamics Online and generate automatic replies.
- Power Virtual Agents – A tool for creating your own virtual agents (bots) that deliver conversational experiences with no coding required.

 **TIP** The Microsoft Business Applications Release Notes (<https://docs.microsoft.com/dynamics365/release-plans/>) provide the near-future roadmap of the Microsoft Power Platform and other applications, such as Microsoft Dynamics.

The role of Power BI in the Microsoft Data Platform

Microsoft has put a lot of effort into making Power BI a one-stop destination for your data analytics needs. Power BI plays an important role in the Microsoft Data Platform by providing services for getting, transforming, and visualizing your data. As far as data acquisition goes, it can connect to cloud and on-premises data sources so that you can import and relate data irrespective of its origin.

Capturing data is one thing but making dirty data suitable for analysis is quite another. However, you can use the data transformation capabilities of Power BI Desktop (or Power Query in Excel) to cleanse and enrich your data. For example, someone might give you an Excel crosstab report. If you import the

data as it is, you'll quickly find that you won't be able to relate it to the other tables in your data model. However, with a few clicks, you can unpivot your data and remove unwanted rows. Moreover, the transformation steps are recorded so that you can apply them later if you're given an updated file.

The main purpose and strength of Power BI is visualizing data in reports and dashboards without requiring any special skills. You can explore and understand your data by having fun with it. To summarize insights from these reports, you can then compile a dashboard. Or you can build the dashboard by asking natural questions. **Figure 1.6** shows a sample dashboard assembled from existing reports.

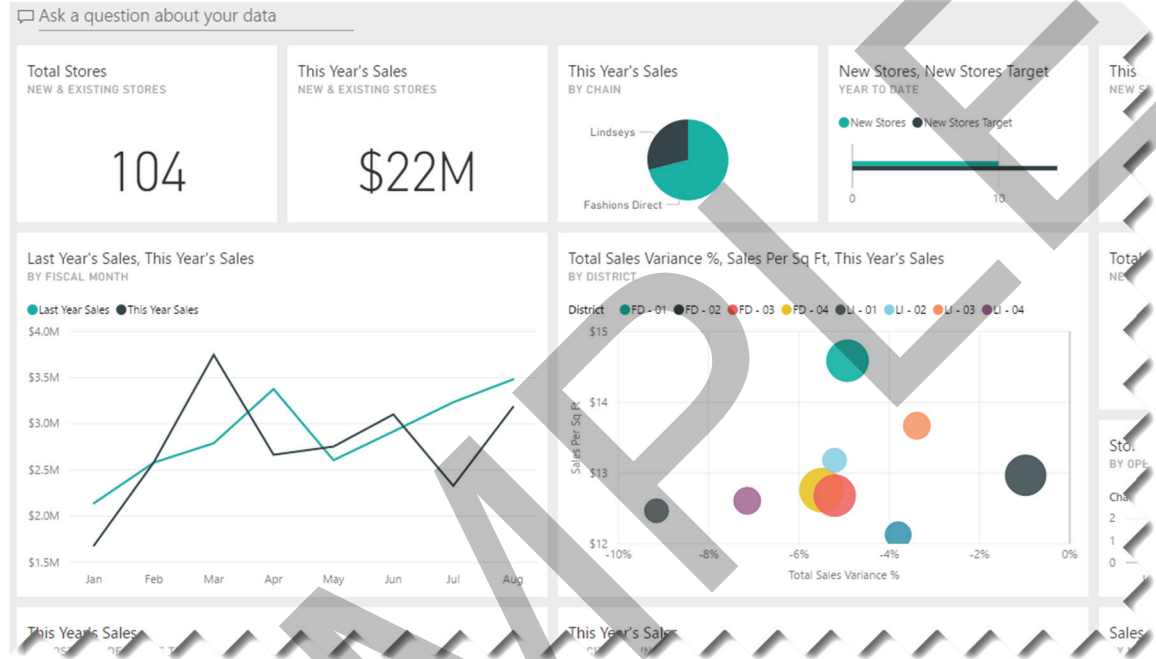


Figure 1.6 Power BI lets you assemble dashboards from existing reports or by asking natural questions.

1.1.5 Power BI Service Editions and Pricing

As it stands, Power BI Service has four licensing options (editions): Free, Power BI Pro, Power BI Premium, and Premium Per User (PPU).

NOTE These editions apply to the cloud Power BI Service (powerbi.com) only. Power BI Desktop and Power BI Mobile are freely available. Power BI Embedded (used by developers to embed reports in apps) has its own licensing and can be acquired by purchasing a Power BI Premium plan or Azure Power BI Embedded plan. Power BI Report Server can be licensed under Power BI Premium or with a SQL Server Enterprise Edition with a Software Assurance license.

Understanding the Free edition

The Power BI Free edition is a free offering and it includes most of the Power BI Service features, but it's licensed for personal use. "Personal" means that a Power BI Free user can't share BI artifacts deployed to the cloud with other users. Specifically, here are the most significant features that are not available in Power BI Free compared to Power BI Pro:

- Item sharing – A Power BI free user can't share reports and dashboards with other users.
- Workspaces – A Power BI Free user can't create workspaces or be a member of a workspace.

- Apps – A Power BI Free user can't create an app (Power BI apps are a mechanism to distribute prepackaged external or internal content).
- Subscriptions – Power BI supports report subscriptions so that reports are delivered via email to subscribed users when the data changes. Power BI Free users can't create subscriptions.
- Connect to published datasets – This feature allows users to connect Excel or Power BI Desktop to datasets published to Power BI and create pivot reports. This is conceptually like connecting directly to an Analysis Services semantic model.



NOTE Microsoft views Power BI Free as an experimental edition for testing Power BI features without requiring a formal approval or on-boarding process. Any user can sign up for Power BI Free using a work email and can keep on using it without time restrictions. Remember that any form of content sharing or collaboration requires a paid SKU.

Understanding the Power BI Pro edition

This paid edition of Power BI Service has a sticker price of \$9.99 per user per month but Microsoft offers discounts so check with your Microsoft reseller. Also, if your organization uses Office 365, you'll find that Power BI Pro is included in the E5 business plan. Power BI Pro offers all the features of Power BI Free, plus sharing and collaboration, and data integration with dataflows.



NOTE Not sure if the Power BI Pro edition is right for you? You can evaluate it for free for 60 days. To start the trial period, sign in to the Power BI portal, click the Settings menu in the top right corner, and then click "Manage Personal Storage". Then click the "Try Pro for free" link.

Understanding the Power BI Premium edition

Think of Power BI Premium as an add-on to Power BI Pro. It requires your organization to commit to a monthly plan. A Power BI Premium plan gives you preconfigured hardware (called a node) that is isolated from other organizations. A Power BI Premium plan has a fixed monthly cost irrespective of how many Power BI Free users you distribute content to. However, every user who will contribute content or change existing content requires a separate Power BI Pro license.

Microsoft provides a nice online calculator at <https://powerbi.microsoft.com/calculator/> to help you estimate your workload of contributors (Pro users), frequent users, and occasional users. Once you plug in the numbers, the online calculator recommends a Power BI Premium plan. For example, given 1,000 total users (200 Pro users, 350 frequent users, and 450 occasional users), the calculator recommends one node on P1 plan costing \$6,993 per month (\$1,998/month for the 200 Pro users and \$4,995 per month for the 1 P1 node). This is \$3,007 less compared to Power BI Pro per-user licensing!



NOTE To avoid overprovisioning, I suggest you start low, such as a P1 plan per environment (DEV, TEST, PRODUCTION), monitor utilization, and upgrade when needed. From a cost perspective alone, the break-even point between Power BI Pro and Power BI Premium is about 500 users. Above that number, Power BI Premium saves money, but cost is just one of the factors when deciding between Power BI Pro and Premium (the other two are features and performance).

From a feature standpoint and compared to Power BI Pro, Power BI Premium adds larger datasets (up to the maximum capacity memory), higher dataset refresh rates (Power BI Pro is limited to a maximum of 8 refreshes per day), incremental dataset refresh, dataset caching, more flexible dataflows, geo distribution, open connectivity, AutoML models, deployment pipelines, and deploying paginated (RDL) reports.

Understanding the Premium Per User edition

Currently in public preview, the Premium per User (PPU) edition targets smaller organizations that can't afford the Power BI Premium monthly commitment but need premium features. Think of it as a hybrid between Power BI Pro and Power BI Premium. Like Power BI Pro, it retains licensing per user without the overhead of managing a premium capacity (Microsoft manages the capacity for you) although at a higher

(currently undisclosed) price point. Like Power BI Premium, PPU provides access to most premium features, such as larger dataset sizes (up to 100GB), paginated reports, AutoML, and others.

Comparing editions and features

Table 1.2 summarizes how editions compare side by side.

Table 1.2 Comparing Power BI editions and features.

Feature	Power BI Free	Power BI Pro	Power BI Premium	Premium per User
Dashboard and report sharing	No	Yes (can't share to Power BI Free)	Yes (can share to Power BI Free)	Yes (all recipients require PPU license)
Workspaces	No	Yes	Yes	Yes (all members require PPU license)
Organizational apps	No	Yes (can't distribute to Power BI Free)	Yes (can distribute to Power BI Free)	Yes (all recipients require PPU license)
Subscriptions	No	Yes (can't distribute to Power BI Free)	Yes (can distribute to Power BI Free)	Yes (all recipients require PPU license)
Connect Excel and Power BI Desktop to published datasets	No	Yes	Yes	Yes
Maximum dataset size	1GB	1GB	Capacity maximum (up to 400GB)	100GB
Maximum workspace storage quota	1GB	10GB	100TB (across the entire capacity)	100TB
Incremental refresh	No	Yes	Yes	Yes
Dataset refresh frequency	8/day	8/day	48/day	48/day
Isolation with dedicated capacity	No	No	Yes	Yes
Data staging with dataflows	No	Serial ingestion, no incremental refresh, no linked entities	Parallel ingestion, incremental refresh, linked entities, calculation engine, Direct Query	Yes
SSRS paginated reports	No	No	Yes	Yes
Content geo distribution	No	No	Yes	No
XMLA endpoint connectivity	No	No	Yes	Yes
Power BI Report Server license	No	No	Yes	No
Deployment pipelines	No	No	Yes	Yes

1.2 Understanding the Power BI Capabilities

Now that I've introduced you to Power BI and the Microsoft Data Platform, let's take a closer look at the Power BI capabilities. I'll discuss them in the context of each of the Power BI products. As I mentioned in section 1.1, Power BI is an umbrella name that unifies several products: Power BI Desktop, Power BI Pro, Power BI Premium, Power BI Mobile, Power BI Report Server, and Power BI Embedded. Don't worry if you don't immediately understand some of these technologies or if you find this section too technical. I'll clarify them throughout the rest of this chapter and the book.

1.2.1 Understanding Power BI Desktop

Business analysts meet self-service BI needs by creating data models, such as to relate data from multiple data sources and then implement business calculations. With Power BI, the design tool for implementing such models is Power BI Desktop. Power BI Desktop is a freely available Windows app for implementing self-service data models and reports. You can download it for free from <https://powerbi.microsoft.com/desktop> or from the Downloads menu in the Power BI portal (powerbi.com) after you log in. Windows 10 users can also install it from the Microsoft Store to automatically keep it up to date.

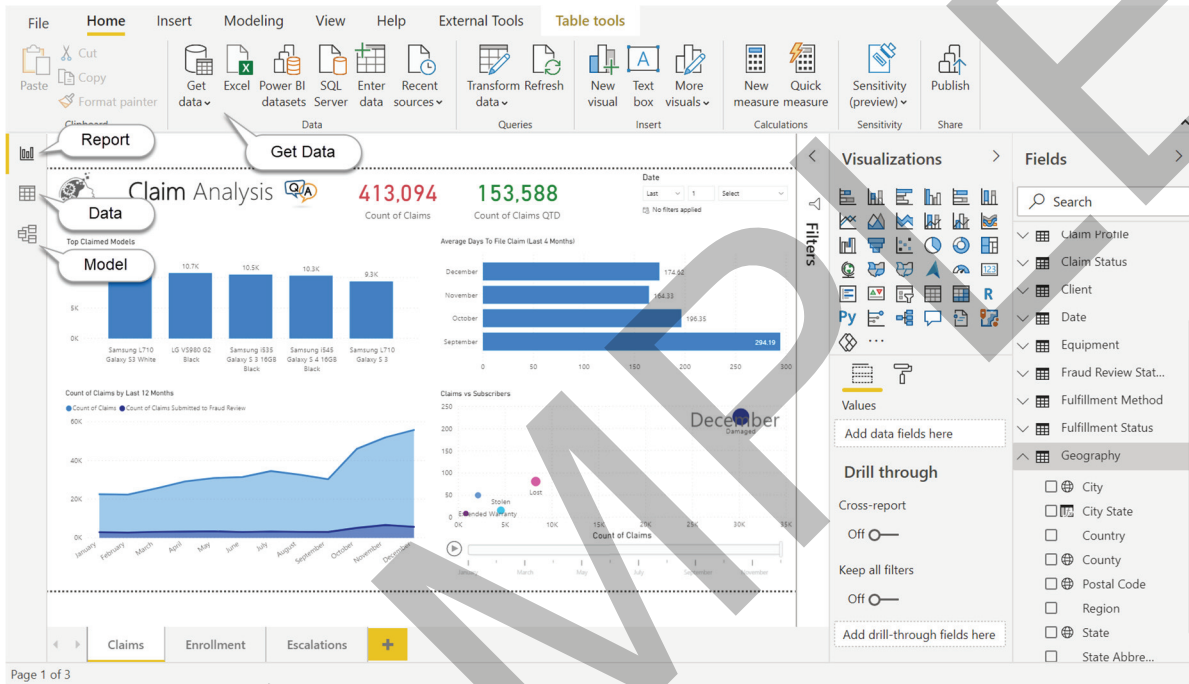


Figure 1.7 Power BI Desktop unifies the capabilities of Power Pivot, Power Query, and Power View.

Understanding Power BI Desktop features

Before Power BI, data analysts could implement data models in Excel. This option is still available, and you can upload your Excel data models to Power BI. However, to overcome the challenges associated with Excel data modeling (see section 1.1.3), Microsoft introduced Power BI Desktop.

If you are familiar with Excel self-service BI, think of Power BI Desktop as the unification of Power Pivot, Power Query, and Power View. Previously available as Excel add-ins, these tools now converge in a single tool. No more guessing which add-in to use and where to find it! At a high level, the data modeling experience in Power BI Desktop now encompasses the following steps (see **Figure 1.7**).

1. Former Power Query – Use the Get Data button in the ribbon to connect to and transform the data. This process is like using Excel Power Query. When you import a dataset, Power BI Desktop creates a table and loads the data. The data is stored in a highly compressed format and loaded in memory to allow you to slice and dice the data efficiently without querying the original data source. However, unlike Excel, Power BI Desktop allows you also to connect directly to a limited number of fast databases, such as Analysis Services and Azure Synapse Analytics (formerly SQL Data Warehouse), where it doesn't make sense to import the data.

2. Former Power Pivot – View and make changes to the data model using the Data and Model tabs in the left navigation bar. This is like Power Pivot in Excel.
3. Former Power View – Create interactive reports using the Report tab on the left, as you can do using Power View in Excel (version 2013 or higher).



NOTE Some data sources, such as Analysis Services, support live connectivity. Once you connect to a live data source, you can jump directly to the Report tab and start creating reports. There are no queries to edit and models to design. In this case, Power BI Desktop acts as a presentation layer that's directly connected to the data source.

Comparing design environments

Because there are many Power Pivot models out there, Power BI allows data analysts to deploy Excel files with embedded data models to Power BI Service and view the included pivot reports and Power View reports online. Power BI Desktop can also import a Power Pivot model if you prefer to migrate your model to Power BI Desktop. So, as a business analyst, you can choose which modeling tool to use:

- Microsoft Excel – Use this option if you prefer to work with Excel and you're familiar with the data modeling features delivered by Excel Power Pivot and Power Query.
- Power BI Desktop – Use this free option if you prefer a simplified tool that's specifically designed for data analytics and that's updated more frequently than Excel.

Table 1.3 compares these two tools side by side to help you choose a design environment. Let's go quickly through the list. While Excel supports at least three ways to import data, many users might struggle in understanding how they compare. By contrast, Power BI Desktop has only one data import option, which is the equivalent of Power Query in Excel. Similarly, Excel has various menus in different places that relate to data modelling. By contrast, if you use Power BI Desktop to import data, your data modeling experience is much more simplified.

Table 1.3 This table compares the data modeling capabilities of Microsoft Excel and Power BI Desktop.

Feature	Excel	Power BI Desktop
Data import	Excel native import, Power Pivot, Power Query	Power Query
Data transformation	Power Query	Power Query
Modeling	Power Pivot	Data and Models tabs
Reporting	Excel pivot reports, Power View, Power Map	Power BI reports (enhanced Power View reports)
Machine learning	Commercial and free add-ins, such as for integration with Azure Machine Learning	Built-in features, such as time series forecasting, clustering, Quick Insights, natural queries
Integration with R and Python	No	Yes
Update frequency	MS Office releases or more often with Office 365 click-to-run	Monthly
Server deployment	SharePoint, Power BI Service, and Power BI Report Server	Power BI Service and Power BI Report Server
Power BI deployment	Import data or connect to the Excel file	Deployed as Power BI Desktop (pbix) file
Convert models	Can't import Power BI Desktop models	Can import Excel Power Pivot models
Upgrade to Tabular	Yes	Not supported by Microsoft
Object model for automation	Yes	No
Cost	Excel license	Free

Excel allows you to create pivot, Power View (now deprecated), and Power Map reports from Power Pivot data models. At this point, Power BI Desktop supports interactive Power BI reports (think of Power View

reports on steroids) and some of the Power Map features (available as a GlobeMap custom visual), although it regularly adds more visualizations and features. Power BI Desktop includes features for machine learning and supports integration with the open-source R and Python languages for data preparation, statistical analysis, machine learning and data visualization.

The Excel update frequency depends on how it's installed. If you install it from a setup disk (MSI installation), you need to wait for the next version to get new features. Office 365 includes subscription-based Microsoft Office (click-to-run installation) which delivers new features as they become available. If you take the Power BI Desktop path, you'll need to download and install updates as they become available. Power BI Desktop is updated monthly so you're always on the latest!

As far as deployment goes, you can deploy Excel Power Pivot models to SharePoint, Power BI Report Server, or Power BI. Power BI Desktop models (files with extension *.pbix) can be deployed to Power BI and Power BI Report Server. Behind the scenes, both Excel and Power BI Desktop use the in-memory xVelocity engine to compress and store imported data.

Power BI Desktop supports importing Power Pivot models from Excel to allow you to migrate models from Excel to Power BI Desktop. Excel doesn't support importing Power BI Desktop models yet, so you can't convert your Power BI Desktop files to Excel data models. A BI pro can migrate Excel Power Pivot models to Analysis Services Tabular when professional features, such as scalability and source control, are desirable. Upgrading Power BI Desktop models to Analysis Services is not supported by Microsoft.



NOTE Power BI Desktop resonates well with business users and most data analysts prefer it over Excel Power Pivot. I recommend Power BI Desktop for self-service BI because it's designed from the ground up for business intelligence and has more data analytics features than Excel.

1.2.2 Understanding Power BI Service

At the heart of the Power BI cloud infrastructure is the *Power BI Service* (powerbi.com). Although not exactly technically accurate, this is what most people refer to when they say "Power BI". You use the service every time you utilize any of the powerbi.com features, such as connecting to online services, deploying and refreshing data models, viewing reports and dashboards, sharing content, or using Q&A (the natural language search feature). Recall that Power BI Service has four licensing options: Power BI Free, Power BI Pro, Power BI Premium, and Premium per User. Since Power BI Free doesn't let users share content, most organizations start with Power BI Pro. Next, I'll introduce you to some of Power BI Pro's most prominent features.

Connect to any data source

The BI journey starts with connecting to data that could be a single file or multiple data sources. Power BI allows you to connect to virtually any accessible data source, either hosted on the cloud or in your company's data center. Your self-service project can start small. If all you need is to analyze a single file, such as an Excel workbook, you might not need a data model. Instead, you can connect Power BI to your file, import its data, and start analyzing data immediately. However, if your data acquisition needs are more involved, such as when you relate data from multiple sources, you can use Power BI Desktop to build a data model whose capabilities can be on par with professional data models and cubes!

Some data sources, such as Analysis Services models, support direct connectivity. Because data isn't imported, direct connections allow reports and dashboards to always be up to date. In the case when you need to import data, you can specify how often the data will be refreshed to keep it synchronized with changes in the original data source. For example, Martin might have decided to import data from the corporate data warehouse and deploy the model to Power BI. To keep the published model up to date, Martin can schedule the data model to refresh daily.

Template apps

Continuing on data connectivity, chances are that your organization uses popular cloud services, such as Salesforce, Dynamics CRM, Google Analytics, Zendesk, and others. Power BI template apps are provided by Microsoft and partners to let you connect to such services and analyze their data without technical setup and data modeling. Apps include a curated collection of reports that continuously update with the latest data from these services. With a few clicks, you can connect to one of the supported online services and start analyzing data using prepackaged reports. **Figure 1.8** shows a prepackaged report for analyzing website traffic based on data imported from Google Analytics. This report is included in the Power BI Google Analytics template app by Heavens Consulting (a Microsoft partner).

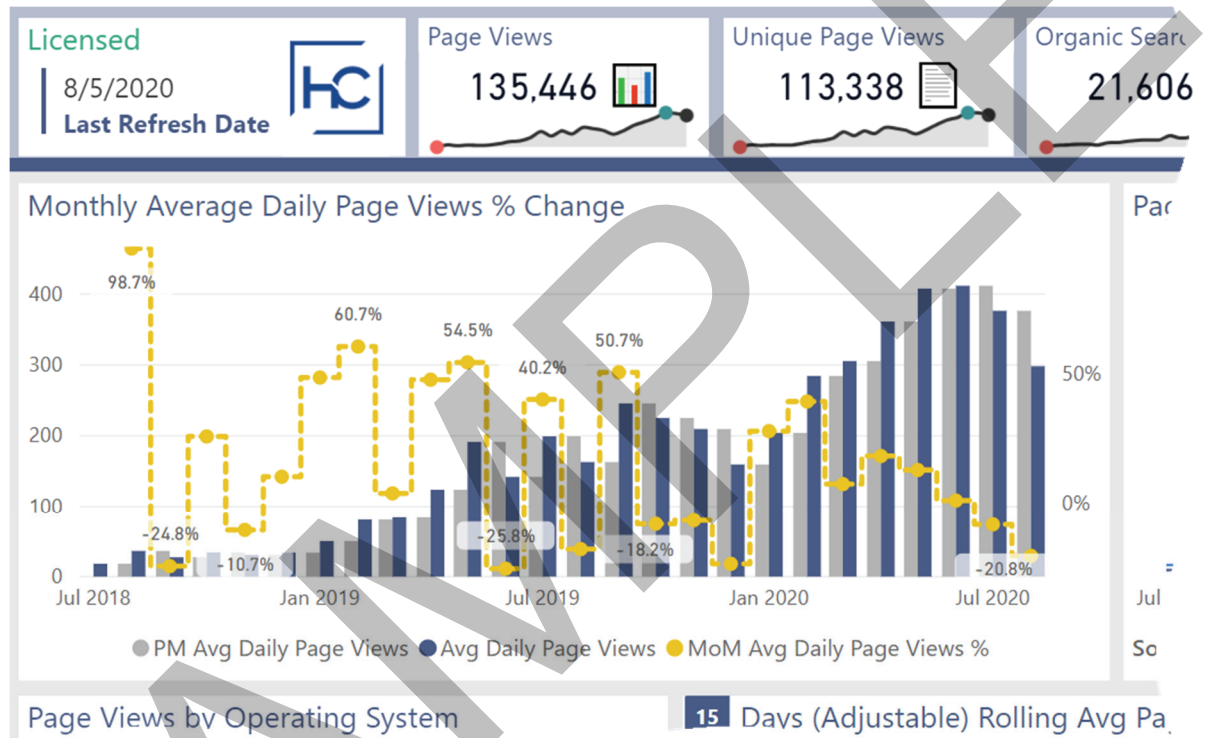


Figure 1.8 Template apps allow you to connect to online services and analyze data using prepackaged reports.

Dashboards and reports

Collected data is meaningless without useful reports. Insightful dashboards and reports are what Power BI Service is all about. To offer a more engaging experience and let users have fun with data while exploring it, Power BI reports are interactive. For example, the report in **Figure 1.9** demonstrates one of these interactive features. In this case, the user has selected Linda in the right bar chart. This action filtered the column chart on the left so that the user can see Linda's contribution to the overall sales. This feature is called cross highlighting.

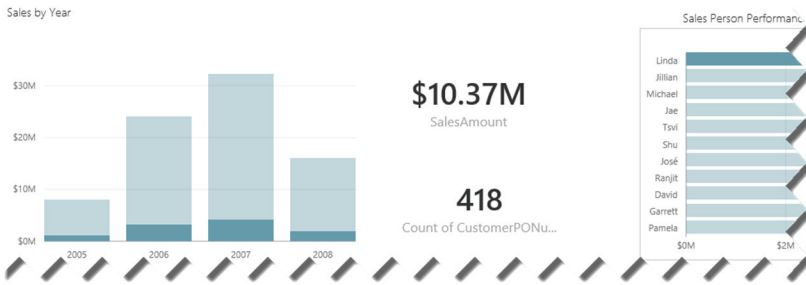


Figure 1.9 Interactive reports allow users to explore data in different ways.

Natural queries (Q&A)

One feature that might excite business users is Power BI natural queries or Q&A. End users are often overwhelmed when asked to create ad hoc reports from a data model. They don't know which fields to use and where to find them. The unfortunate "solution" by IT is to create new reports to answer new questions. This might result in a ton of reports that are replaced by new reports and are never used again. However, Power BI allows users to ask natural questions, such as "this year's sales by district in descending order by this year's sales" (see **Figure 1.10**).

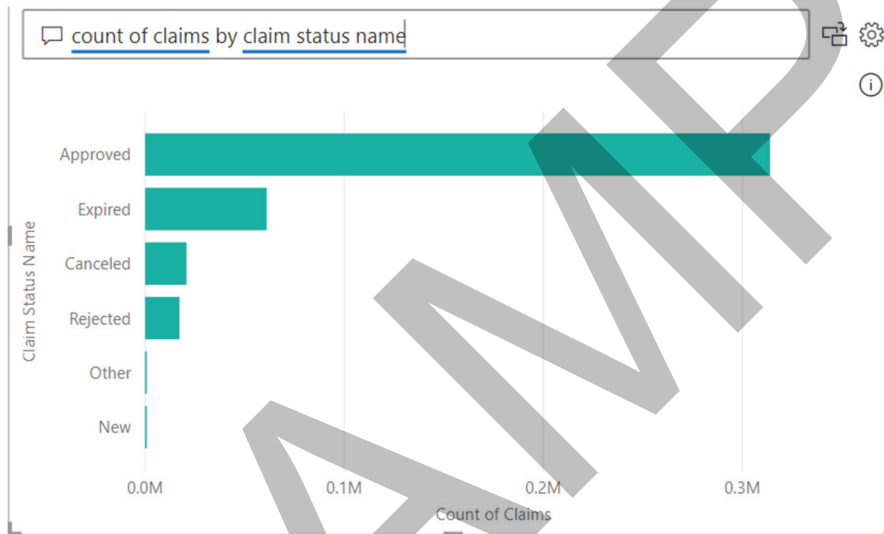


Figure 1.10 Q&A allows users to explore data by asking natural questions.

Not only can Power BI interpret natural questions, but it also chooses the best visualization! While in this case Q&A has decided to use a Bar Chart, it might have chosen a map if the question was phrased in a different way. Q&A is available in both Power BI Service and Power BI Desktop.

Sharing and collaboration

Once you've created informative reports and dashboards, you might want to share them with your coworkers. Power BI supports several sharing options but recall that all of them require Power BI Pro or Premium subscriptions. To start, you can share specific reports and dashboards as read-only with your coworkers. Or you can use Power BI Pro workspaces to allow groups of people to have access to the same workspace content and collaborate on it. For example, if Maya works in sales, she can create a Sales Department workspace and grant her coworkers access to the workspace. Then all content added to the Sales Department workspace will be shared among the group members.

Yet a third way to share content is to create an organizational app. Like a template app that you can use to analyze data from popular online services, you can create a Power BI organizational app to share

content from a workspace across teams or even with everyone from your organization. Users can discover and install template and organizational apps (see **Figure 1.11**). In this case, the user sees that someone has published a Sales Department app. The user can connect to the app and access its content as read-only.

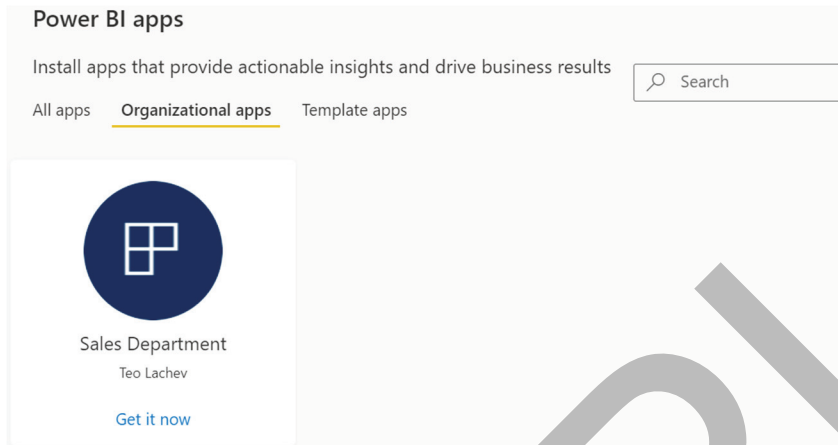


Figure 1.11 Users within your organization can use the Power BI AppSource to discover public online template apps or internal organizational apps.

Alerts and subscriptions

Do you want to be notified when your data changes beyond certain levels? Of course, you do! You can set up as many alerts as you want in both Power BI Service and Power BI Mobile. You can set up rules to be alerted when single number tiles in your dashboard exceed limits that you set. With data-driven alerts, you can gain insights and act wherever you're located.

Would you like Power BI to email you your favorite report when its data changes? Just view the report in Power BI Service and subscribe yourself and coworkers to a report page of interest. Power BI will regularly send a screenshot of that report page directly to your mail inbox and a link to the actual report.

Data staging and preparation

Data quality and integration is a big issue for many organizations. Although you can use Power Query that is included in Power BI Desktop to shape and transform data before it becomes available for reporting, some scenarios might require additional data staging and preparation. Suppose that your organization uses a cloud-based customer relationship management (CRM) system, such as Salesforce or Microsoft Dynamics Online. Since CRM data is so important to your company, many data analysts rely on this data. However, you might run into long data refresh times to synchronize your Power BI models with changes in the CRM system.

Instead of connecting directly to the CRM system, a better approach might be to stage the CRM data either by using the vendor-provided staging mechanism (Microsoft Dynamics CRM can stage its data to Azure Data Lake Store) or using Power BI dataflows. The former will require help from your IT department, while the latter opens the option for self-service data staging.

Think of a dataflow as "Power Query in the cloud". For example, **Figure 1.12** shows how I used a dataflow to connect to Microsoft Dynamics CRM and select which entity I'd like to stage. Once the dataflow is created, I can schedule it to extract and save the data periodically to a Microsoft-provided or your organization data lake. Then, data analysts can use Power BI Desktop to connect to the staged data and import it in their models.

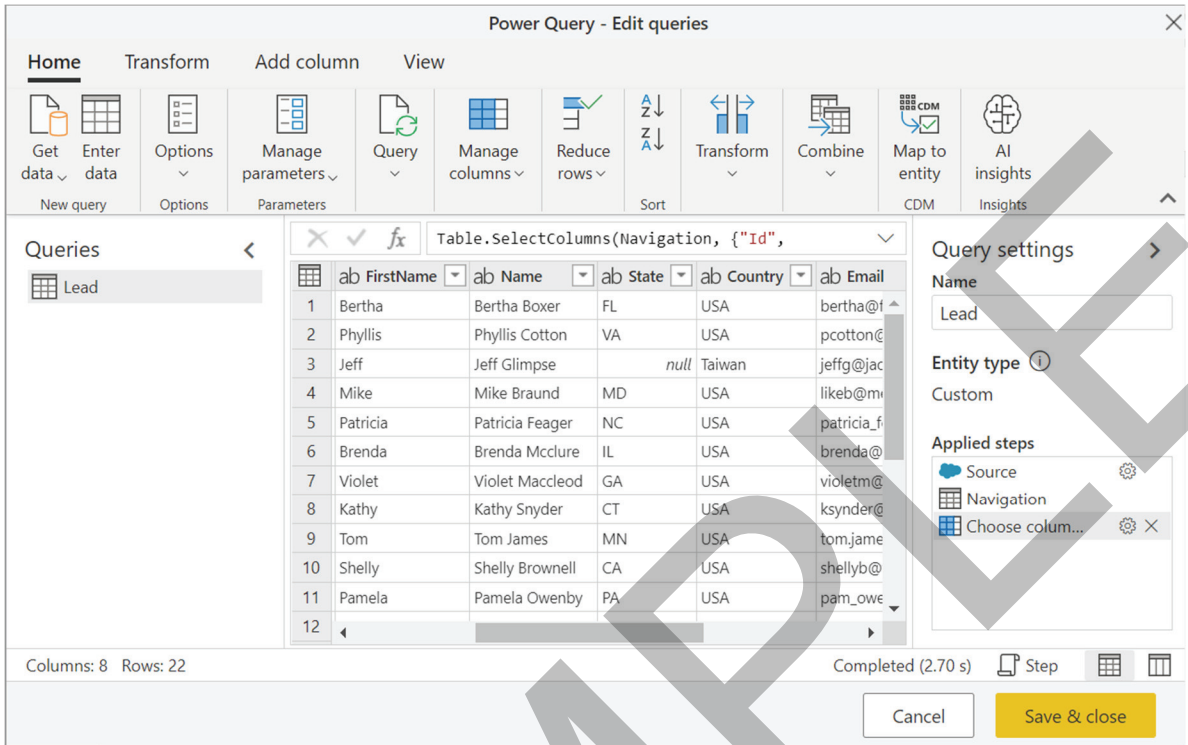


Figure 1.12 Using a dataflow to stage data from Salesforce.

1.2.3 Understanding Power BI Premium

I previously explained that Power BI Premium extends the Power BI Pro capabilities by providing a dedicated environment with more features and ability to reduce licensing cost by not requiring licenses for users who only require access to view reports. Let's take a quick look at some of the most prominent Power BI Premium features.

Understanding shared and dedicated capacities

Like how a Windows folder or network share is used to store related files, a Power BI workspace is a container of logically related Power BI artifacts. A workspace is in a shared capacity when its workloads run on computational resources shared by other customers. Power BI Free and Power BI Pro workspaces always run in a shared capacity. However, in Power BI Premium, a Power BI Pro user with special capacity admin permissions can move a workspace to a premium capacity. Premium capacity is a dedicated hardware provisioned just for your organization.

In Figure 1.13, the Sales workspace was initially created in a shared capacity. Its report performance could be affected by workloads from other Power BI customers. To avoid this, the admin might decide to move it to a premium capacity. Now the workspace is isolated, and its performance is not affected by other organizations that use Power BI. However, it's still dependent on the activity of other premium workspaces in *your* organization and the resourced constraints of the Power BI Premium plan that it's associated with.

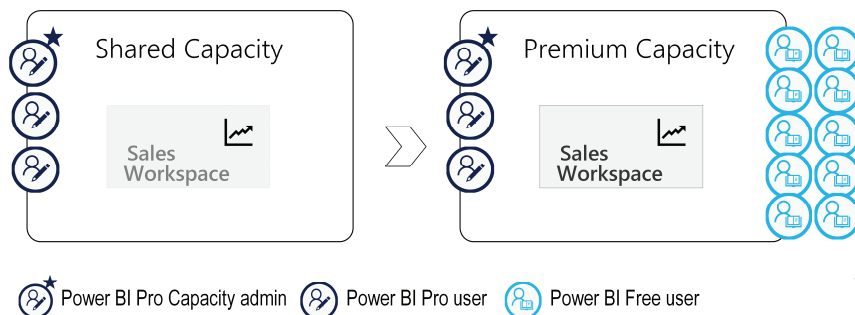


Figure 1.13 The Capacity Admin can move workspaces in and out of dedicated capacity.

The interesting detail is that the admin can move a workspace in and out of the premium capacity at any point in time. For example, increased seasonal workloads may prompt the admin to move some workspaces to a premium capacity for a certain duration and then move them back to shared capacity when the workloads are reduced. You control which workspaces are in what capacity.

Understanding content distribution

Glancing again at **Figure 1.13**, we can see that when the Sales workspace was in a shared capacity, only Power BI Pro members could access its content. Power BI Free users would need to upgrade to Power BI Pro to gain access as members. However, when the workspace is moved to a premium capacity, its content can be shared to Power BI Free users. This is how Power BI Premium helps large organizations reduce Power BI licensing cost and distribute content to many users when only read-only access is enough.

Understanding premium features

Power BI Premium includes all Power BI Pro capabilities and adds more features. At this point, here are the most prominent premium features:

- Large datasets – Power BI Premium increases the maximum dataset size up to the maximum capacity memory and the workspace storage quota across the entire capacity up to 100 TB.
- Dataset caching – A dataset with imported data can be configured for caching to speed up the first report page.
- More frequent refreshes – Datasets can be scheduled for refresh up to 48 times per day.
- More flexible dataflows – Dataflows can reference entities staged by other dataflows in the same or different workspaces. Entities within a dataflow are refreshed in parallel to speed up the overall refresh time. Moreover, Power BI Premium allows organizations to bring their own storage for storing dataflow entities. This enables interesting integration scenarios. For example, other applications can act upon the data, such as by applying machine learning algorithms, before the data is ingested in dataflows.
- XMLA endpoint – As discussed in more detail in section 1.3, the workhouse of Power BI Service is Analysis Services. Premium workspaces let you connect to the endpoint of the backend Analysis Services service. The main benefit is that you can connect third-party reporting tools to published datasets if you find the Power BI visualization capabilities lacking. The open XMLA endpoint also allows BI pros to deploy and monitor organizational semantic models using their tool of choice, such as SSDT, Tabular Editor, and SQL Profiler.
- Paginated (SSRS) reports – Although Power BI reports excel in interactivity, Reporting Services paginated reports excel in flexibility and customization. You can meet more advanced reporting requirements with paginated reports that can be deployed to a premium workspace.

- AI-powered self-service models – Data analysts can quickly put together Machine Learning models based on AutoML, Azure Cognitive Services, and Azure Machine Learning. For example, Martin can create a dataflow that integrates with Azure Cognitive Services for sentiment analysis.
- Multi-geo support – Larger organizations can distribute content to multiple data centers to meet regulatory and scalability requirements. For example, a US-based organization can have South Central US as its home region but configure a Power BI Premium capacity in a European region so that content deployed to that capacity stays in Europe.
- Deployment pipelines – Facilitate content deployment between environments, such as from Dev to Test to Production.

1.2.4 Understanding Power BI Mobile

Power BI Mobile is a set of native mobile applications for iOS, Windows and Android devices. You can access the download links from <https://powerbi.microsoft.com/mobile>. Why do you need these applications? After all, thanks to Power BI HTML5 rendering, you can view Power BI reports and dashboards in your favorite Internet browser. However, the native applications offer features that go beyond just rendering. Although there are some implementation differences, this section covers some of the most compelling features (Chapter 5 has more details).

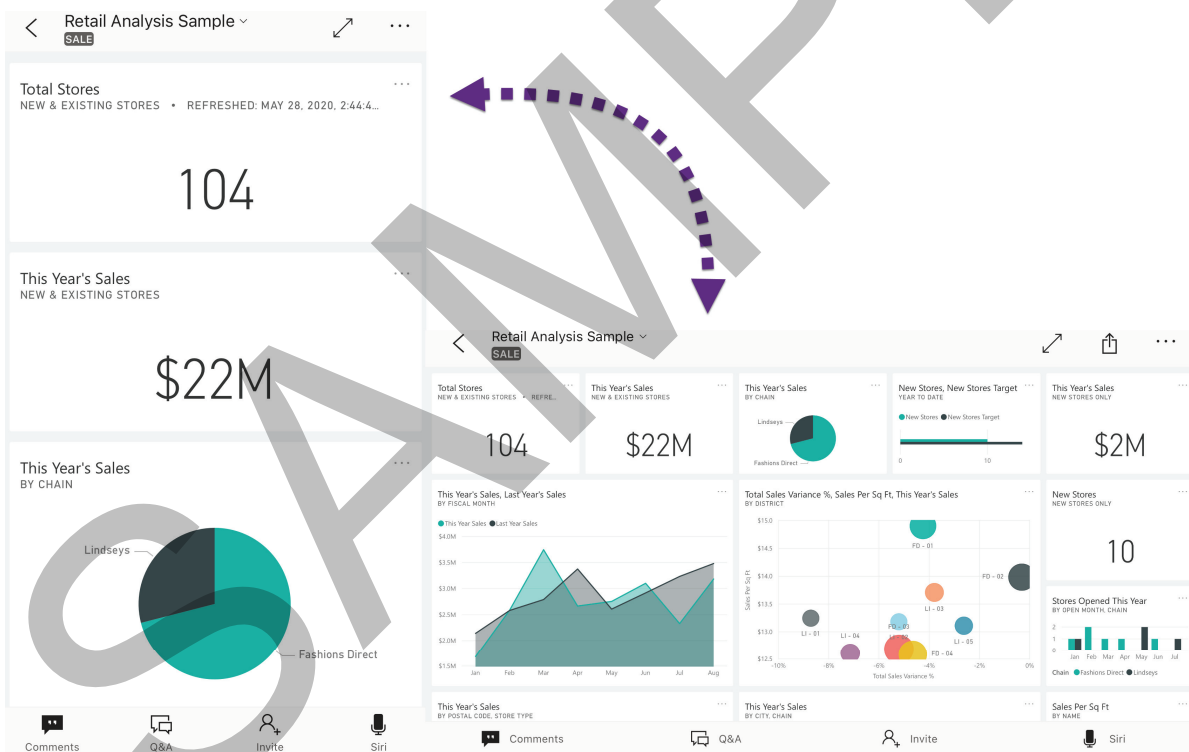


Figure 1.14 Power BI Mobile adjusts the dashboard layout when you rotate your phone from portrait to landscape.

Optimized viewing

Mobile devices have limited display capabilities. The Power BI mobile apps adjust the layout of dashboards and reports, so they display better on mobile devices. For example, by default, viewing a dashboard on a phone in portrait mode will position each dashboard tile after another. Rotating the phone to landscape will show the dashboard as it appears in Power BI Service (Figure 1.14). You can further tune the presentation by making changes to dashboards and reports in a special mobile portrait layout.

Alerts

Instead of going to powerbi.com to set up an alert on a dashboard tile, you can set up alerts directly in your mobile app. For example, Figure 1.15 shows that I've enabled an iPhone data alert to be notified when this year's sales exceed \$23 million. When the condition is met, I'll get a notification and email.

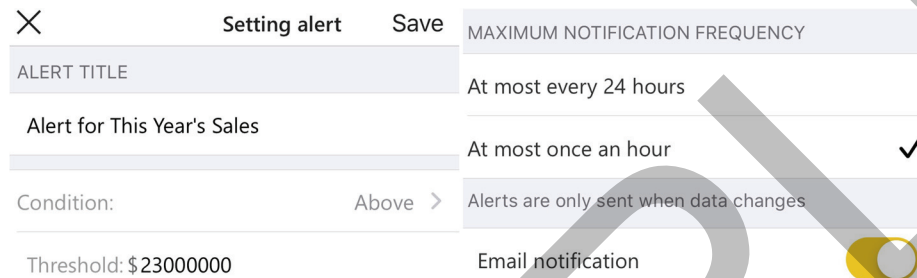


Figure 1.15 Alerts notify you about important data changes, such as when sales exceed a certain threshold.

Annotations and discussions

Annotations allow you to add lines, text, and stamps to dashboard tiles (see Figure 1.16). For example, you could use annotations to ask the person responsible to sign that the report is correct. Then you can mail a screen snapshot to recipients, such as to your manager.

Besides annotations, users can start a conversation at a dashboard, report, or even visual level. Think of a conversation as a discussion list. Users can type comments and reply to comments entered by other users.

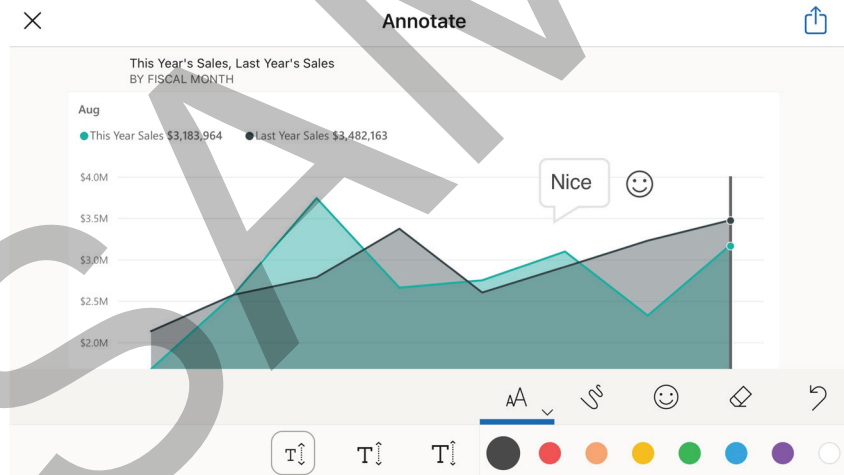


Figure 1.16 Annotations allow you to add comments to tiles and then send screenshots to your coworkers.

Sharing

Like Power BI simple sharing, you can use a mobile device to share a dashboard by inviting coworkers to access the dashboard. Dashboards shared by mail are read-only, meaning that the people you share with can only view the dashboard without making changes.

1.2.5 Understanding Power BI Embedded

Almost every app requires some reporting capabilities. Traditionally, developers would either use third-party widgets or embed Reporting Services reports using the Microsoft ReportViewer control. The first approach requires a lot of custom code. The latter limits your users to static (canned) reports. What if you want to deliver the Power BI interactive experience with your apps? Enter Power BI Embedded!

Introducing Power BI Embedded features

Power BI Embedded allows developers and Independent Software Vendors (ISV) to add interactive Power BI reports in their custom apps for internal or external users. Because Power BI Embedded uses the same APIs as Power BI Service, it has feature parity with Power BI Service. Suppose Teo has developed an ASP.NET MVP app for external customers. The app authenticates users any way it wants, such as by using Forms Authentication. Teo has created some nice reports in Power BI Desktop that connect directly to an Analysis Services semantic model or to data imported in Power BI Desktop.

With a few lines of code, Teo can embed these reports in his app (see **Figure 1.17**). If the app connects to a multi-tenant database (customers share the same database), the app can pass the user identity to Power BI Embedded, which in turn can pass it to the model. Then, row-level security (RLS) filters can limit access to data.

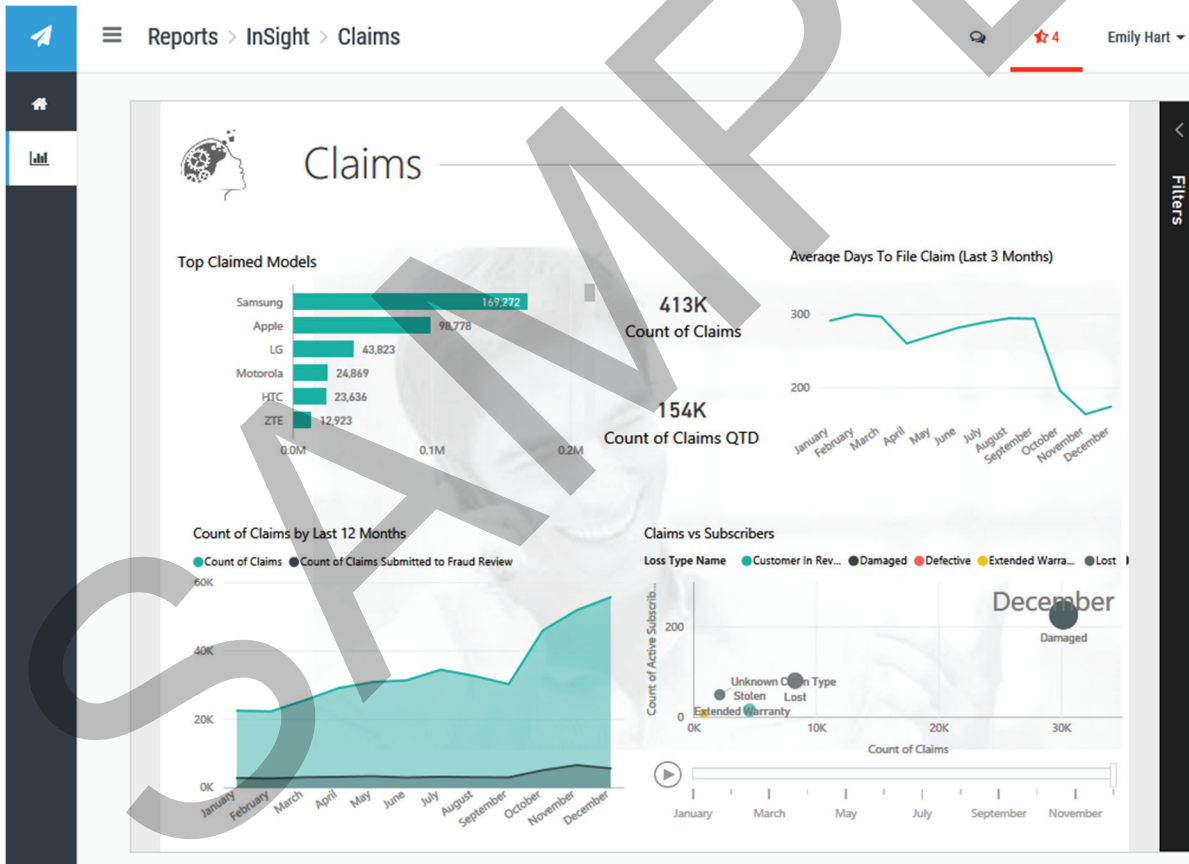


Figure 1.17 Power BI Embedded allows developers to embed Power BI reports in custom apps.

Power BI Embedded is extensible. Teo can use its JavaScript APIs to programmatically manipulate the client-side object model. For example, he can replace the Filters pane with a customized user interface to filter data. Or can navigate the user programmatically to a specific report page.

About Power BI Embedded licensing

Per-user, per-month licensing is not cost effective for delivering reports to many users. Like Power BI Premium, Power BI Embedded utilizes capacity-based pricing. Power BI Embedded can be acquired via Power BI Premium by purchasing a Power BI Premium P plan or EM plan. The Power BI Premium P plans give you access to both embedded and service deployments. The EM plans are mostly for embedded deployments.

Power BI Embedded can also be acquired outside of Power BI Premium by purchasing an Azure Power BI Embedded plan. This is the preferred and most cost-effective licensing option if you need only external reporting, such as if you work for an ISV that provides services for a third party. More information about these plans can be found at <https://azure.microsoft.com/pricing/details/power-bi-embedded/>. I'll also provide more details when I discuss Power BI Embedded in Chapters 13 and 16.

1.2.6 Understanding Power BI Report Server

Many organizations have investments in on-premises reporting with Microsoft SQL Server Reporting Services (SSRS). Starting with SQL Server 2017, SSRS doesn't ship with SQL Server anymore but can be downloaded separately from the Microsoft Download Center as two SKUs:

- Microsoft SQL Server Reporting Services – This is the SSRS SKU you are familiar with that continues to be licensed under SQL Server. It allows you to deploy operational (RDL) reports and SSRS mobile reports, but it doesn't support Power BI reports and Excel reports.
- Power BI Report Server – This SKU associates with the strong Power BI brand. It's still SSRS but in addition to operational and mobile reports, it also supports Power BI reports and Excel reports (the latter requires integration with Microsoft Office Online Server). With Power BI Report Server, you have full flexibility to decide what portions of the data and reports you want to keep on-premises and what portions should reside in the cloud.



NOTE Besides splitting SSRS into two products, decoupling SSRS from SQL Server allows Microsoft to deliver new features faster and be more competitive in the fast-changing BI world. Also, while there is nothing stopping you from deploying Power BI Desktop files and Excel files to SSRS, reports won't render online, and the user will be asked to download and open the file locally. So, when I said that Power BI Report Server supports Power BI and Excel reports, I meant that these reports render online, and that their management is integrated in the report portal.

Introducing Power BI Report Server features

As of this time, Power BI Report Server is limited to basic integration with Power BI reports, but Microsoft is actively working to deliver more features every two months or so. You can deploy Power BI Desktop files to the server and then view them online (see **Figure 1.18**). Report interactivity is supported. Power BI reports share the same security model as other items deployed to the report catalog. Power BI reports deployed to Power BI Report Server can also be viewed in Power BI mobile apps.

Understanding Power BI Report Server licensing

Power BI Report Server can be acquired in two ways:

- Power BI Premium P Plan – A P plan licenses the same number of on-premises cores as the number of v-cores licensed for cloud usage. Suppose your organization has purchased the Power BI Premium P1 plan. This plan licenses 8 v-cores of a premium capacity in Power BI Service. When

you install Power BI Report Server on premises, it will be licensed for 8 cores, giving you a total of 16 licensed cores!

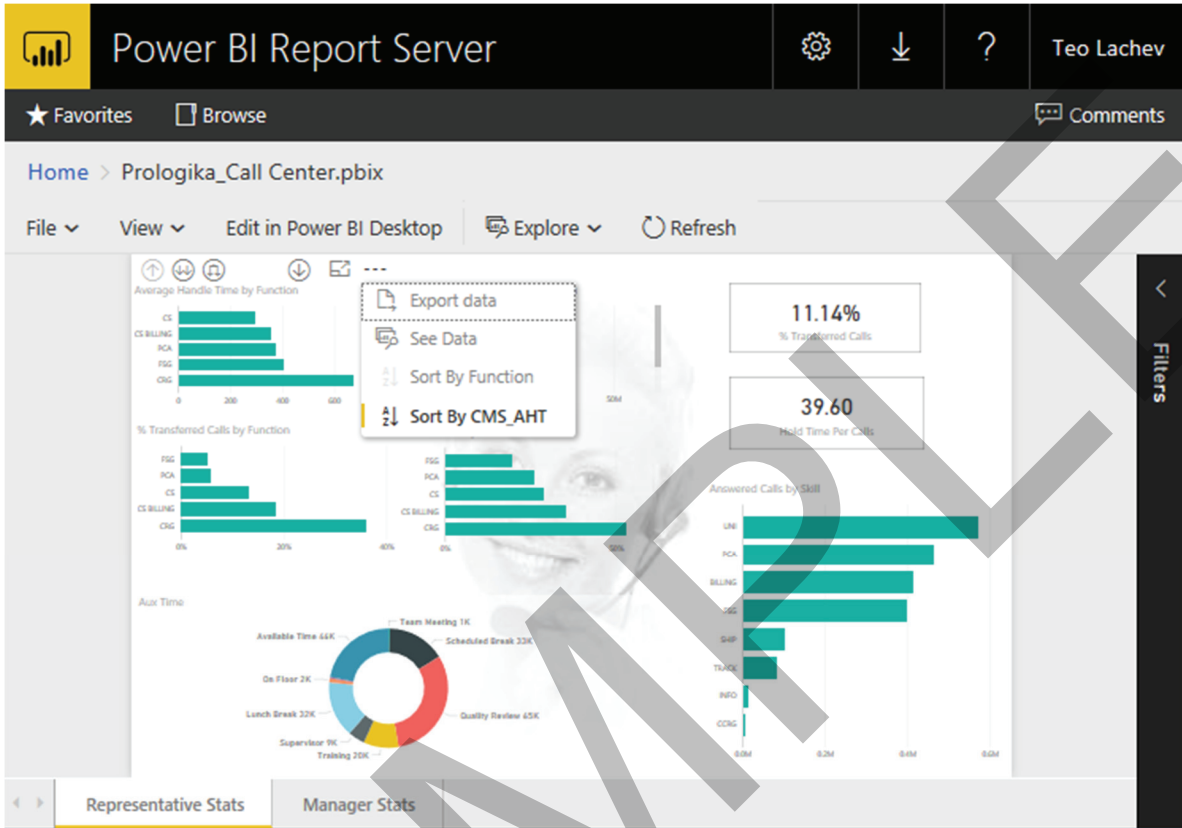


Figure 1.18 Power BI reports render online when deployed to the Power BI Report Server.

- SQL Server Enterprise with Software Assurance license – Not interested in the cloud yet? You can cover Power BI Report Server under the SQL Server Enterprise licensing model, just as you license SSRS Enterprise Edition.

NOTE Like Power BI, Power BI Report Server requires Power BI Pro licenses for content creators. For example, if you have 5 report developers that will deploy reports to Power BI Report Server, you will need 5 Power BI Pro licenses (recall that each Power BI license is \$9.99 per user, per month). Licensing content creators is honor-based as currently there is no mechanism to ensure that the user is licensed on deploying content to the Power BI Report Server.

1.3 Understanding the Power BI Service Architecture

Microsoft has put a significant amount of effort into building Power BI Service that consists of various Azure services that handle data storage, security, load balancing, disaster recovery, logging, tracing, and so on. Although it's all implemented and managed by Microsoft (that's why we like the cloud) and it's completely transparent for you, the following sections give you a high-level overview of these services to help you understand their value and Microsoft's decision to make Power BI a cloud service.

The Power BI Service is hosted on the Microsoft Azure cloud platform and it's deployed in various data centers around the world. **Figure 1.19** shows a summarized view of the overall technical architecture that consists of two clusters: a Web Front End (WFE) cluster and a Back End cluster.

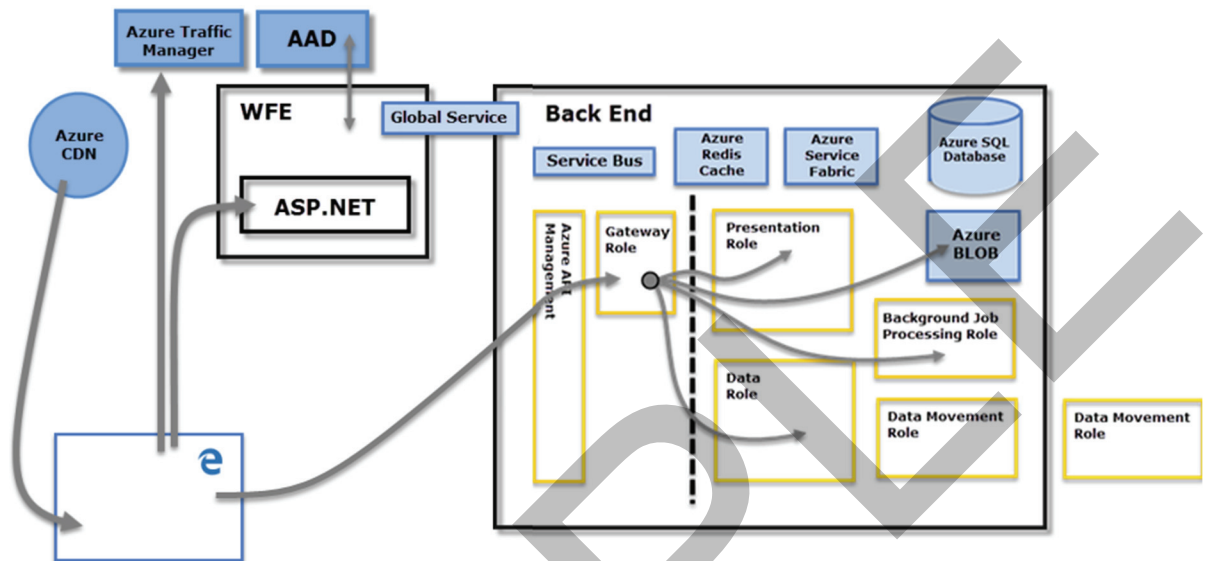


Figure 1.19 Power BI is powered by Microsoft Azure clusters.

1.3.1 The Web Front End (WFE) Cluster

The WFE cluster manages connectivity and authentication. Power BI relies on Azure Active Directory (AAD) to manage account authentication and management. Power BI uses the Azure Traffic Manager (ATM) to direct user traffic to the nearest datacenter. Which data center is used is determined by the DNS record of the client attempting to connect. The DNS Service can communicate with the Azure Traffic Manager to find the nearest datacenter with a Power BI deployment.

TIP To find where your data is stored, log in to Power BI and click the Help (?) menu in the top-right corner, and then click "About Power BI". Power BI shows a prompt that includes the Power BI version and the data center.

Power BI uses the Azure Content Delivery Network (CDN) to deliver the necessary static content and files to end users based on their geographical locale. The WFE cluster nearest to the user manages the user login and authentication and provides an access token to the user once authentication is successful. The ASP.NET component within the WFE cluster parses the request to determine which organization the user belongs to, and then consults the Power BI Global Service.

The Global Service is implemented as a single Azure Table that is shared among all worldwide WFE and Back End clusters. This service maps users and customer organizations to the datacenter that hosts their Power BI tenant. The WFE specifies to the browser which backend cluster houses the organization's tenant. Once a user is authenticated, subsequent client interactions occur with the backend cluster directly and the WFE cluster is not used.

1.3.2 The Backend Cluster

The backend cluster manages all actions the user does in Power BI Service, including visualizations, dashboards, datasets, reports, data storage, data connections, data refresh, and others. The Gateway Role acts as a gateway between user requests and the Power BI service. As you can see in the diagram, only the Gateway Role and Azure API Management (APIM) services are accessible from the public Internet.

When an authenticated user connects to the Power BI Service, the connection and any request by the client is accepted and managed by the Gateway Role, which then interacts on the user's behalf with the rest of the Power BI Service. For example, when a client attempts to view a dashboard, the Gateway Role accepts that request, and then sends a request to the Presentation Role to retrieve the data needed by the browser to render the dashboard.

Where is data stored?

As far as data storage in the cloud goes, Power BI uses two primary repositories for storing and managing data. Data that is uploaded from users or generated by dataflows is stored in Azure BLOB storage, but all the metadata definitions (dashboards, reports, recent data sources, workspaces, organizational information, tenant information) are stored in Azure SQL Database.

The working horse of the Power BI service is Microsoft Analysis Services in Tabular mode, which has been architected to fulfill the role of a highly scalable data engine where many servers (nodes) participate in a multi-tenant, load-balanced farm. For example, when you import some data into Power BI, the actual data is stored in Azure BLOB storage (or Azure Premium Files for large datasets deployed to a premium capacity), but an in-memory Tabular database is created to service queries.

Analysis Services Tabular enhancements

For BI pros who are familiar with Tabular, new components have been implemented so that Tabular is up to its new role. These components enable various cloud operations including tracing, logging, service-to-service operations, reporting loads and others. For example, Tabular has been enhanced to support the following features required by Power BI:

- Custom authentication – Because the traditional Windows NTLM authentication isn't appropriate in the cloud world, certificate-based authentication and custom security were added.
- Resource governance per database – Because databases from different customers (tenants) are hosted on the same server, Tabular ensures that any one database doesn't use all the resources.
- Diskless mode – For performance reasons, the data files aren't initially extracted to disk.
- Faster commit operations – This feature is used to isolate databases from each other. When committing data, the server-level lock is now only taken for a fraction of the time, although database-level commit locks are still taken, and queries can still block commits and vice versa.
- More Dynamic Management Views (DMVs) – For better status discovery and load balancing.
- Data refresh – From the on-premises data using a gateway.
- Additional features – Microsoft adds features first to Tabular in Power BI and later to Azure Analysis Services and SSAS. At this point, the following Analysis Services features are only available in Power BI: incremental refresh, composite models with hybrid storage, and aggregations.

1.3.3 Data on Your Terms

The increasing number of security exploits in the recent years has made many organizations cautious about protecting their data and skeptical about the cloud. You might be curious to know what is uploaded to the Power BI service and how you can reduce your risk for unauthorized access to your data. In

addition, you control where your data is stored. Although Power BI is a cloud service, this doesn't necessarily mean that your data must be uploaded to Power BI.

Live connections

In a nutshell, you have two options to access your data. If the data source supports direct connectivity, you can choose to leave the data where it is and only create reports and dashboards that connect live to your data. Currently, a subset of the supported data sources supports live connectivity, but that number is growing! Among them are Analysis Services, SQL Server (on premises and on Azure), Oracle, Azure Synapse Analytics (formerly Azure SQL Data Warehouse), Amazon Redshift, Snowflake, Google BigQuery, SAP Hana, and Spark/Databricks.

For example, if Elena has implemented an Analysis Services model and deployed to a server in her organization's data center, Maya can create reports and dashboards in Power BI Service by directly connecting to the model. In this case, the data remains on premises; only the report and dashboard definitions are hosted in Power BI. When Maya runs a report, the report generates a query and sends the query to the model. Then, the model returns the query results to Power BI. Finally, Power BI generates the report and sends the output to the user's web browser. Power BI always uses the Secure Sockets Layer (SSL) protocol to encrypt the traffic between the Internet browser and the Power BI Service so that sensitive data is protected.



NOTE Although in this case the data remains on premises, aggregated data displayed on reports and dashboards still travels from your data center to Power BI Service. This could be an issue for software vendors who have service level agreements prohibiting data movement. You can address such concerns by referring the customer to the Power BI Security document (<http://bit.ly/1SkEzTP>) and the accompanying Power BI Security whitepaper.

Importing data

The second option is to import and store the data in Power BI. For example, Martin might want to build a data model to analyze data from multiple data sources. Martin can use Power BI Desktop to import the data and analyze it locally. To share reports and allow other users to create reports, Martin decides to deploy the model to Power BI. In this case, the model and the imported data are uploaded to Power BI, where they're securely stored.

To synchronize data changes, Martin can schedule a data refresh. Martin doesn't need to worry about security because data transfer between Power BI and on-premises data sources is secured through Azure Service Bus. Azure Service Bus creates a secure channel between Power BI Service and your computer. Because the secure connection happens over HTTPS, there's no need to open a port in your company's firewall.



TIP If you want to avoid moving data to the cloud, one solution you can consider is implementing an Analysis Services model layered on top of your data source. Not only does this approach keep the data local, but it also offers other important benefits, such as the ability to handle larger datasets (millions of rows), a single version of the truth by centralizing business calculations, row-level security, and others. Finally, if you want to avoid the cloud completely, don't forget that you can deploy Power BI reports to an on-premises Power BI Report Server.

1.4 Power BI and You

Microsoft envisions that over time Power BI will become a one-stop destination for all BI needs. Now that I've introduced you to Power BI and its building blocks, let's see what Power BI means for you. As you'll see, Power BI has plenty to offer to anyone interested in data analytics, irrespective of whether you're a content producer or consumer, as shown in **Figure 1.20**.



Figure 1.20 Power BI supports the BI needs of business users, analysts, pros, and developers.

By the way, the book content follows the same organization so that you can quickly find the relevant information depending on what type of user you are. For example, if you're a business user, the first part of the book is for you and it has four chapters (chapters 2-5) for the first four features shown in the "For business users" section in the diagram.

1.4.1 Power BI for Business Users

To clarify the term, a business user is someone in your organization who is mostly interested in consuming BI artifacts, such as reports and dashboards. This group of users typically includes executives, managers, business strategists, and regular information workers. To get better and faster insights, some business users often become basic content producers, such as when they create reports to analyze simple datasets or data from online services.

For example, Maya is a manager in the Adventure Works Sales & Marketing department. She doesn't have skills to create sophisticated data models and business calculations. However, she's interested in monitoring the Adventure Works sales by using reports and dashboards produced by other users. She's also a BI content producer because she must create reports for analyzing data in Excel spreadsheets, website traffic, and customer relationship management (CRM) data.

Connect to your data without creating models

Thanks to the Power BI template apps, Maya can connect to popular cloud services, such as Google Analytics and Dynamics CRM, and get instant reports. She can also benefit from prepackaged content created jointly by Software as a Service (SaaS) partners and Microsoft. Power BI refers to these connectors with the prepackaged artifacts collectively as *template apps*.

For example, the Dynamics CRM template app provides an easy access to analyze data from the cloud-hosted version of Dynamics CRM. This app uses the Dynamics CRM OData feed to generate a model that contains the most important entities, such as Accounts, Activities, Opportunities, Products, Leads, and others. Similarly, if Maya uses Salesforce as a CRM platform, Power BI has a template app to allow Maya to connect to Salesforce in a minute. Power BI apps support data refresh, such as to allow Maya to refresh the CRM data daily.

Create reports

Power BI can also help Maya analyze simple datasets without data modeling. For example, if Maya receives an Excel file with some sales data, she can import the data into Power BI and create ad hoc reports with a few mouse clicks. The experience is not much different than creating Excel pivot reports.

Create and share content

Maya can easily assemble dashboards from her reports and from reports shared with her by her colleagues. She can also easily share her dashboards with coworkers. For example, Maya can select a dashboard and then click the Share button next to the dashboard name (see **Figure 1.21**).

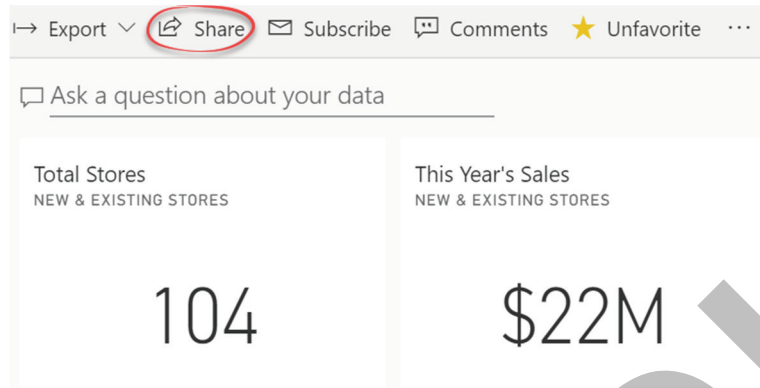


Figure 1.21 Business users can easily share dashboards and reports with coworkers using the Power BI portal or Power BI Mobile.

Go mobile

Some business users, especially managers, executives, and salespeople, would need access to BI reports on the go. These users would benefit from the Power BI Mobile native applications for iPad, iPhone, Android, and Windows. As I explained in section 1.2.4, Power BI Mobile allows users to not only view Power BI reports and dashboards, but to also receive alerts about important data changes, and to share and annotate dashboards.

For example, while Maya travels on business trips, she needs access to her reports and dashboards. Thanks to the cloud-based nature of Power BI, she can access them anywhere she has an Internet connection. Depending on what type of mobile device she uses, she can also install a Power BI app, so she can benefit from additional useful features, such as favorites, annotations, and content sharing.

1.4.2 Power BI for Data Analysts

A data analyst or BI analyst is a power user who has the skills and desire to create self-service data models. A data analyst typically prefers to work directly with the raw data, such as to relate corporate sales data coming from the corporate data warehouse with external data, like economic data, demographics data, weather data, or any other data purchased from a third-party provider.

For example, Martin is a BI analyst with Adventure Works. Martin has experience in analyzing data with Excel and Microsoft Access. To offload effort from IT, Martin wants to create his own data model by combining data from multiple data sources.

Acquire and mash up data from virtually everywhere

As I mentioned previously, to create data models, Martin can use Microsoft Excel and/or Power BI Desktop, which combines the best of Power Query, Power Pivot, and Power View in a single and simplified design environment. If he has prior Power Pivot experience, Martin will find Power BI Desktop easier to use and he might decide to switch to it to stay on top of the latest Power BI features. Irrespective of the design environment chosen, Martin can use either Excel or Power BI Desktop to connect to any accessible data source, such as a relational database, file, cloud-based services, SharePoint lists, Exchange servers, and many more.

Currently, Power BI Desktop ships with more than 100 data connectors. Microsoft regularly adds new data sources and developers can create custom data sources using the Power BI Data Connector SDK.

Cleanse, transform, and shape data

Data is rarely cleaned. A unique feature of Power BI Desktop is cleansing and transforming data. Inheriting these features from Power Query, Power BI Desktop allows a data analyst to apply popular transformation tasks that save tremendous data cleansing effort, such as replacing values, un-pivoting data, combining datasets and columns, and many more tasks.

For example, Martin may need to import an Excel financial report that was given to him in a crosstab format where data is pivoted by months on columns. Martin realizes that if he imports the data as it is, he won't be able to relate it to a date table that he has in the model. However, with a couple of mouse clicks, Martin can use a Power BI Desktop query to un-pivot months from columns to rows. And once Martin gets a new file, the query will apply the same transformations so that Martin doesn't have to go through the steps again.

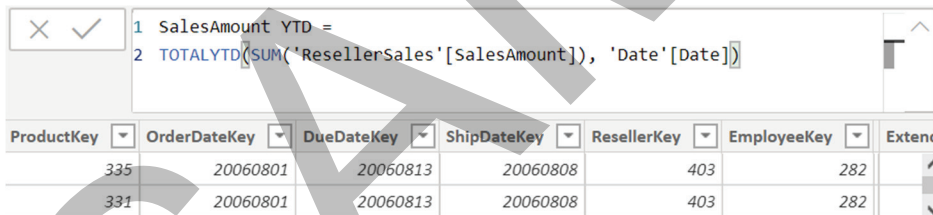
Implement self-service data models

Once the data is imported, Martin can relate the datasets to analyze the data from different angles by relating multiple datasets (see again **Figure 1.1**). No matter which source the data came from, Martin can use Power BI Desktop (or Excel) to relate tables and create data models whose features are on par with professional models. When doing so, Martin can also create a composite model spanning imported tables and tables with live connections.

For example, if some tables in an ERP system are frequently updated, Martin could decide to access the sales transactions via a live connection so that he always sees the latest data, while the rest of the data is imported. Further, Power BI supports flexible relationships with one-to-many and many-to-many cardinality, so Martin can model complex requirements, such as analyzing financial balances of joint bank accounts.

Create business calculations

Martin can also implement sophisticated business calculations, such as time calculations, weighted averages, variances, period growth, and so on. To do so, Martin will use the Data Analysis Expression (DAX) language and Excel-like formulas. To help you get started with common business calculations, Power BI includes *quick measures* (prepackaged DAX expressions). For example, the formula shown in **Figure 1.22** calculates the year-to-date (YTD) sales amount. As you can see, Power BI Desktop supports IntelliSense and color coding to help you with the formula syntax. IntelliSense offers suggestions as you type.



The screenshot shows a DAX formula bar with the following text:

```
1 SalesAmount YTD =  
2 TOTALYTD(SUM('ResellerSales'[SalesAmount]), 'Date'[Date])
```

Below the formula bar is a table with the following columns and data:

ProductKey	OrderDateKey	DueDateKey	ShipDateKey	ResellerKey	EmployeeKey	Extend
335	20060801	20060813	20060808	403	282	
331	20060801	20060813	20060808	403	282	

Figure 1.22 Business calculations are implemented in DAX.

Get insights

Once the model is created, the analyst can visualize and explore the data with interactive reports. If you come from using Excel Power Pivot and would like to give Power BI Desktop a try, you'll find that not only does it simplify the design experience, but it also supports many new visualizations, such as Funnel and Combo Charts, Treemap, Filled Map, and Gauge visualizations, as shown in **Figure 1.23**.

And when the Microsoft-provided visualizations aren't enough, Martin can use a custom visual contributed by Microsoft and the Power BI community. For example, Martin might need to present the most common words in surveys as a word cloud. Since Power BI doesn't include such a visual, Martin navigates to the Microsoft AppStore (<https://appssource.microsoft.com>) and picks the Word Cloud custom visual contributed by Microsoft to visualize data in awesome ways!

Once Martin is done with the report in Power BI Desktop, he can publish the model and reports to Power BI, so that he can share insights with other users. If they have permissions, his coworkers can view reports, gain more insights with natural query (Q&A) questions, and create dashboards. Martin can also schedule a data refresh to keep the imported data up to date.

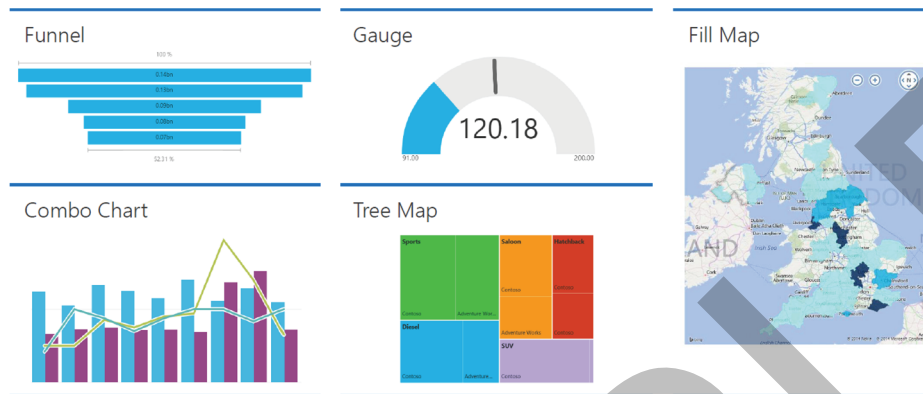


Figure 1.23 Power BI Desktop adds new visualizations.

1.4.3 Power BI for Pros

BI pros and IT pros have much to gain from Power BI. BI pros are typically tasked to create the backend infrastructure required to support organizational BI initiatives, such as data marts, data warehouses, cubes, ETL packages, operational reports, and dashboards. IT pros are also concerned with setting up and maintaining the necessary environment that facilitates self-service and organizational BI, such as providing access to data, managing security, data governance, and other services.

In a department or smaller organization, a single person typically fulfills both BI and IT pro tasks. For example, Elena has developed an Analysis Services model on top of the corporate data warehouse. She needs to ensure that business users can gain insights from the model without compromising security.

Enable team BI

Once she provides connectivity to the on-premises model, Elena must establish a trustworthy environment needed to facilitate content sharing and collaboration. To do so, she can use Power BI workspaces. As a first step, Elena would set up groups and add members to these groups. Then Elena can create workspaces for the organizational units interested in analyzing the SSAS model. For example, if the Sales department needs access to the organizational model, Elena can set up a Sales Department group. Next, she can create a Sales Department workspace and grant the group access to it. Finally, she can deploy to the workspace her sales-related dashboards and reports that connect to the model.

If Elena needs to distribute BI artifacts to a wider audience, such as the entire organization, she can create an app and publish it. Then her coworkers can search, discover, and use the app read-only.

Scale report workloads

No one likes to wait for a report to finish. If Elena works for a larger organization, she can scale report workloads by purchasing a Power BI Premium plan. She then decides which workspaces can benefit from a dedicated capacity and promotes them to premium workspaces. Not only does Power BI Premium deliver consistent performance but it also allows the organization to save on the Power BI licensing cost. Elena can now share out content in premium workspaces to "viewers" by sharing specific dashboards or distributing contents with apps.

Implementing BI solutions

Based on my experience, most organizations could benefit from what I refer to as a classic BI architecture that includes a data warehouse and semantic model (Analysis Services Multidimensional or Tabular mode) layered on top of the data warehouse. I'll discuss the benefits of this architecture in Part 3 of this book. If you already have or are planning such a solution, you can use Power BI as a presentation layer. This works because Power BI can connect to the on-premises Analysis Services, as shown in **Figure 1.24**.

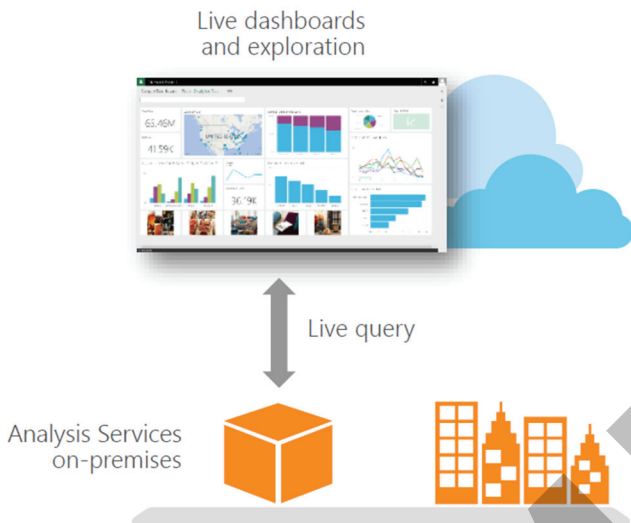


Figure 1.24 Power BI can directly connect to on-premises databases, such as Analysis Services semantic models.

So that Power BI can connect to on-premises SSAS models, Elena needs to download and install a component called a gateway to an on-premises computer that can connect to the semantic model. The gateway allows Elena to centralize management and access to on-premises data sources. Then Elena can implement reports and dashboards that connect live to Analysis Services and deploy them to Power BI. When users open a report, the report will generate a query and send it to the on-premises model via the gateway. Now you have a hybrid solution where data stays on premises but reports are hosted in Power BI.

If you're concerned about the performance of this architecture, you should know that Power BI only sends queries to the on-premises data source so there isn't much overhead on the trip from Power BI to the source. Typically, BI reports and dashboards summarize data. Therefore, the size of the datasets that travel back to Power BI probably won't be very large either. Of course, the speed of the connection between Power BI and the data center where the model resides will affect the duration of the round trip.

Another increasingly popular scenario that Power BI can help you implement is real-time BI. You've probably heard about Internet of Things (IoT) which refers to an environment of many connected devices, such as barcode readers, sensors, or cell phones, that transfer data over a network without requiring human-to-human or human-to-computer interaction. If your organization is looking for a real-time platform, you should seriously consider Power BI. Its streamed datasets allow an application to stream directly to Power BI with a few lines of code. If you need to implement Complex Event Processing (CEP) solutions, Microsoft Azure Stream Analytics lets you monitor event streams in real time and push results to a Power BI dashboard.

Finally, BI pros can implement predictive data analytics solutions that integrate with Power BI. For example, Elena can use the Azure Machine Learning Service to implement a data mining model that predicts the customer probability to purchase a product. Then she can easily set up a REST API web service, which Power BI can integrate with to display results. If all these BI pro features sound interesting, I'll walk you through these scenarios in detail in Part 3 of this book.

1.4.4 Power BI for Developers

Power BI has plenty to offer to developers as well because it's built on an open and extensible architecture. In the context of data analytics, developers are primarily interested in incorporating BI features in their applications or in providing access to data to support integration scenarios. For example, Teo is a developer with Adventure Works. Teo might be interested in embedding Power BI dashboards and reports in a web application that will be used by external customers. Power BI supports several extensibility options, including apps, real-time dashboards, custom visuals, and embedded reporting.

Automate management tasks

Power BI has a set of REST APIs to allow developers to programmatically manage certain Power BI resources, such as enumerating datasets, creating new datasets, and adding and removing rows to a dataset table. This allows developers to push data to Power BI, such as to create real-time dashboards. In fact, this is how Azure Stream Analytics integrates with Power BI. When new data is streamed, Azure Stream Analytics pushes the data to Power BI to update real-time dashboards.

The process for creating such applications is straightforward. First, you need to register your app. Then you write OAuth2 security code to authenticate your application with Power BI. Then you'd write code to manipulate the Power BI objects using REST APIs. For example, here's a sample method invocation for adding one row to a table:

```
POST https://api.powerbi.com/beta/myorg/datasets/2C0CCF12-A369-4985-A643-0995C249D5B9/Tables/Product/Rows HTTP/1.1
Authorization: Bearer {AAD Token}
Content-Type: application/json
{
  "rows": [
    {
      "ProductID":1,
      "Name":"Adjustable Race",
      "Category":"Components",
      "IsCompete":true,
      "ManufacturedOn":"07/30/2014"
    }
  ]
}
```

Microsoft supports a Power BI Developer Center website (<https://powerbi.microsoft.com/developers>) where you can read the REST API documentation and try the REST APIs.

Embed reports in custom apps

Many of you would like to embed beautiful Power BI dashboards and reports in custom applications. For example, your company might have a web portal to allow external customers to log in and access reports and dashboards that are included in the app.

For internal applications where users are already using Power BI, developers can call the Power BI REST APIs to embed dashboard tiles and reports. As I mentioned, external applications can benefit from Power BI Embedded. And, because embedded reports preserve interactive features, users can enjoy the same engaging experience, including report filtering, interactive sorting, and highlighting. I cover these integration scenarios in Chapter 16.

Implement custom visuals

Microsoft has published the required interfaces to allow developers to implement and publish custom visuals using any of the JavaScript-based visualization frameworks, such as D3.js, WebGL, Canvas, or SVG. Do you need visualizations that Power BI doesn't support to display data more effectively? With some coding wizardry, you can implement your own! You can use whatever tool you prefer to code the custom visual (visuals are coded in TypeScript), such as Microsoft Visual Code or Visual Studio. When the custom visual is ready, you can publish it to Microsoft AppSource at <https://appsource.microsoft.com> where Power BI users can search for it and download it.

Power BI is an extensible platform and there are other options for building Power BI solutions, including:

- Integrate Power BI with Microsoft Power Automate and Power Apps – For example, in Chapter 10, I'll show you how you can integrate Power BI with Power Apps to change the data behind a report.
- Implement custom data connectors – You can extend the Power BI data capabilities by implementing custom data connectors in M language (the programming language of Power Query). To learn more, see the M Extensions GitHub repo at <https://github.com/Microsoft/DataConnectors/blob/master/docs/m-extensions.md>.
- Implement template apps – I've already discussed how Power BI template apps can help you connect to popular online services, such as Dynamics CRM or Google Analytics. You can implement new apps to facilitate access to data and to provide prepackaged content. As a prerequisite, contact Microsoft and sign up for the Microsoft partner program which coordinates this initiative. Power BI partners and ISVs can also build Power BI template apps to provide out-of-the-box content for their customers and deploy them to any Power BI tenant.

1.5 Summary

This chapter has been a whirlwind tour of the innovative Power BI cloud data analytics service and its features. By now, you should view Power BI as a flexible platform that meets a variety of BI requirements. An important part of the Microsoft Data Platform, Power BI is a collective name of several products: Power BI, Power BI Desktop, Power BI Premium, Power BI Mobile, Power BI Embedded, and Power BI Report Server. You've learned about the major reasons that led to the release of Power BI. You've also taken a close look at the Power BI architecture and its components, as well as its editions and pricing model.

Next, this chapter discussed how Power BI can help different types of users with their data analytics needs. It allows business users to connect to their data and gain quick insights. It empowers data analysts to create sophisticated data models. It enables IT and BI pros to implement hybrid solutions that span on-premises data models and reports deployed to the cloud. Finally, its extensible and open architecture lets developers enhance the Power BI data capabilities and integrate Power BI with custom applications.

Having laid the foundation of Power BI, you're ready to continue the journey. Next, you'll witness the value that Power BI can deliver to business users.