



Recipe for a Modern Marketing Dashboard

Table of Contents

Introduction/Background	02
Problem	03
Solution	03
Gather Requirements	04
Inventory of Systems and Data Sources	05
Backend and Frontend Technology Identification	06
Cube Development	07
Dashboard Building	08
Learnings	10

Introduction/Background

The Microsoft US Central Marketing Organization is responsible for planning, executing, and reporting on marketing efforts across the US market. We use a multitude of platforms and systems to run our business and support the many marketers across the sub. Each of these systems supports a marketing channel or a portion of a marketing channel. For example, our event registration and planning tool is used by marketers to collect registrations and track registrant status of in-person events. Most of these systems house their own reporting, and while that reporting is often robust, it only tells a single portion of the story. What happens once those registrants attend an in-person event and express their interest in learning more about a specific product?

This becomes especially complex in a matrixed organization like Microsoft, where every business group has its own internal systems that must be connected to chronicle the entire customer journey. In addition to working across business groups, our marketers also work across various customer segments, creating even more complexity with variations in marketing tactics.



Problem

Our marketers work with a wide range of complex systems and tools, like CRM and usage tracking. They also handle a myriad of post-acquisition, lower-funnel activities, like telemarketing and internal product purchase. Across these disparate systems and processes, it was impossible for them to get understand the entire customer journey—from anonymous visitor through the full selling lifecycle. When you can't see the full lifecycle of a lead, you don't know where to optimize, or what parts of the journey may need attention. It was almost like they were trying to follow a recipe with access to only some of the necessary ingredients.

To equip our marketers with the proper tools and ingredients, we needed a way to connect a convoluted ecosystem of internal and external tools and data sources, and create a comprehensive view of the entire marketing ecosystem. The ideal solution would give marketers a complete view of their programs, pinpoint where drop offs/issues were occurring, and provide insights into steps that should be taken to make improvements. Because the number of tools we used, and the ways in which we utilize existing tools were growing, we also needed something that would be extremely flexible.

Solution

Like perfecting a recipe, developing a robust standardized reporting solution that encompasses all marketing efforts and all stages of the customer lifecycle, and serves up the right data at the right level of granularity for the audience, is a process. It takes testing, obsessing over the results, adjusting accordingly, and, in some cases, thinking outside the box. **Below are the steps we took to perfect our recipe for the Modern Marketing Dashboard:**

1. **Gather Requirements**
2. **Take inventory of systems and data sources**
3. **Identify backend and frontend technology**
4. **Build backend infrastructure**
5. **Build frontend reports and dashboards**
6. **Release to users and provide support**



Step 1: Gather Requirements

Beef Wellington isn't the same without the puff pastry outer layer. Without cream, a White Russian is basically just alcohol. We knew our current reporting solutions weren't meeting most of our users' needs. Before knowing what kind of standardized solution to cook up, we needed to understand how our user base would leverage this reporting. In other words, we needed to know what ingredients our users needed to do their jobs the right way. We began this process by interviewing multiple marketers from various groups, and at various levels of the business. **Prior to conducting these interviews, we identified several key topics to touch on:**

- **Cadence:** How often will users be using the reporting? Will they be digging in daily or looking at high-level metrics during their monthly reviews? This would help us understand the refresh cadence we would need to support with the solution we were building.
- **Comfort level with data:** How sophisticated are users in working with this type of data (and data in general)? Reviewing any existing reporting was helpful to understand the level of data analysis marketers were already doing. This also helped us to understand which tools they were using, which would guide our platform decision.
- **Business metric familiarity:** Which performance metrics are they looking at, if any? How comfortable are they with these metrics (do they really understand how they are calculated) and do they feel that they are true indicators/predictors of the performance they purport the measure?

Asking questions about these items, and truly understanding the answers, informed our decisions about the architecture of both the backend and frontend reporting solutions. We also built personas to represent the key user groups we identified through the requirements gathering stage. This helped us identify two high level categories of users toward which we needed to fine tune our recipe:

1. **Executives**—These strategy-driven individuals typically look at reports on a weekly or monthly basis for key indicators that provide a high-level, broad view of marketing performance.
2. **Field Marketers**—This operational audience takes a more tactical approach to marketing performance. They want to understand which specific programs are working and which aren't, and where there might be issues that need attention. Field marketers look at data more often than executives, potentially even daily. As a result, this user needs near real-time data and the option to drill into the data for in-depth analysis.

After meeting with individuals across marketing organizations to understand what ingredients they needed, we built mockups and reviewed performance in another round of meetings. Here, we focused on which key metrics we would emphasize, the various views and cuts of data, and which visualizations were most interesting to the users. The result was a solid foundation for our modern marketing dashboard recipe.



Step 2: Take Inventory of Systems and Data Sources

To deliver the reporting we had envisioned during the requirements gathering phase, we pulled data from a multitude different systems and integrated it into a single model. These sources included:

1. Data from internal tools and systems:
 - CRM systems.
 - Marketing Automation (multiple platforms used by different business groups).
 - Business group specific data sources.
2. Data directly from external 3rd party tools via API.
3. Marketing activity metadata from internal marketer Support Desk CRM system
4. Lookup tables and online forms where additional data points are tracked, including target metrics tracked by marketing teams.

In addition to integrating data sources, we also needed to perform augmentation on the data to improve the value of the information to users. This included conducting account identification and segmentation (matching of leads to existing accounts in Microsoft's database) and identification of individuals as either net new to the Microsoft ecosystem or already existing. These additional attributes required checking against additional internal databases.





Step 3: Identify Backend and Frontend Technology Identification

At this point, our recipe was taking shape and we were cooking with gas. After reviewing the data sources, we decided to build a cube to support our reporting. We knew we were dealing with a large amount of data since we were pulling transaction-level data from our marketing automation platform—for every lead, we had a row of data for each trackable action the lead takes.

Another option we explored was building a tabular model. Although tabular models are quicker to stand up, performance is typically slower because they are loaded in memory. Given the amount of data, and the high level of performance our users would require, we decided a cube was the best option. We also knew we would be making updates to the structure, and adding more data sources over time, and a cube would offer more flexibility and ease of maintenance.



A pinch of Power BI: Power BI is an excellent choice for visualizing data from a cube, but it also supports building data models within the Power BI Desktop application. This is a great option if you have a smaller data set and are not worried about performance. When conducting a POC for a new reporting offering, our team typically builds out an initial tabular model, before embarking on the more complex development of a cube. There are a multitude of options for connecting to existing data models or building out models within Power BI, which makes it a great tool that allows for lots of flexibility.

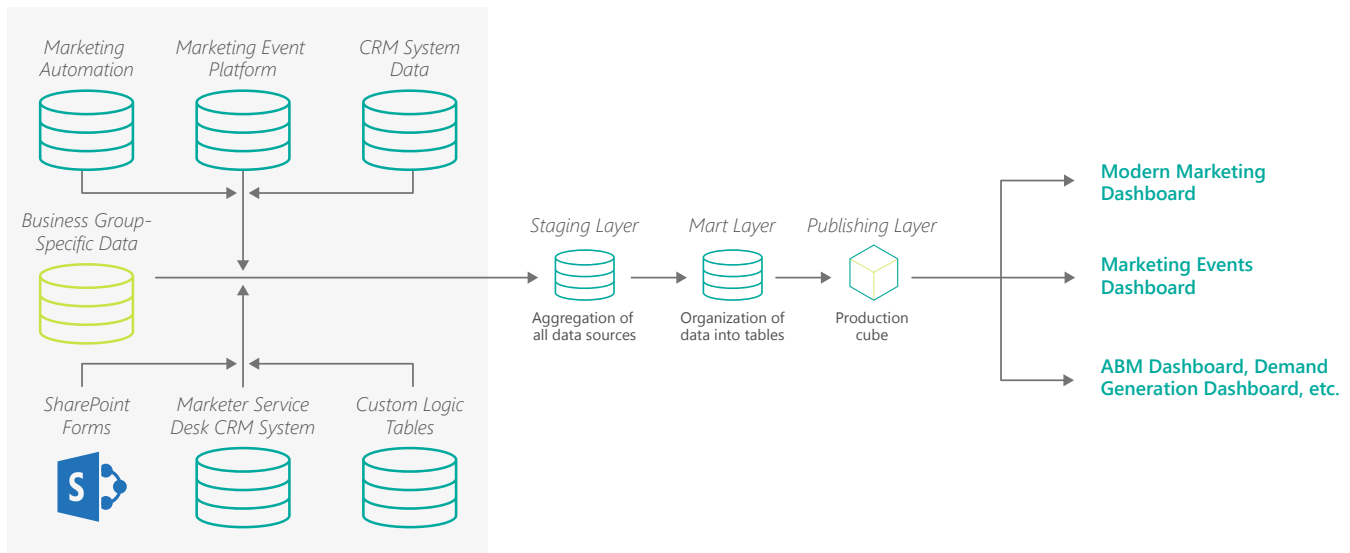
For these reasons, we chose to build our dashboard in Power BI. We knew that our cube would connect seamlessly to the Power BI desktop tool and allow for automated refresh. The existing functionality and visualizations in Power BI met our needs, but we also knew that new features and functionality were continuously being introduced to the platform.

We wanted to be able to take advantage of the new features that would roll out in the future, and knew we would be able to, given the incredible flexibility of the platform. This flexibility was also important as we were building a net new reporting infrastructure for a marketing organization that is constantly experimenting with new processes and tools. With Power BI, we can make revisions and updates as often as needed.



Step 4: Cube Development

System Architecture



We developed our system architecture in SQL servers hosted on Azure. We needed to set up APIs to call data from some sources, but most data were pulled from other databases owned by different teams throughout the company. Most of the databases we pulled from house data that support the entire company, so we applied filters to reduce the data set to what was relevant for our U.S. audience.

To achieve close to 100% reliability and a stable development environment, we built out several layers for processing the data. We started with a staging layer, where we pulled historical transaction level data from all our sources, then normalized the data in a data mart layer.

From here, we built a data mart and structured the data into dimension and fact tables. We built surrogate keys to join relevant tables together. Now, individual attendees of a marketing event from our event platform could be tied to individuals in our marketing automation system, then further tied to leads, and on to sales opportunities and wins in our CRM systems. We also built out our custom logic in the data mart layer to apply to the data. This allowed us to create categories and other descriptive metadata necessary for reporting.

For example, our account level data comes from a company-wide system, but we have field marketing teams that cover a specific group of accounts based on geography. These groupings are not relevant to the entire company, and therefore not integrated into the data source. Instead, we created mapping tables to categorize US accounts from the worldwide account data sources in the way we needed to surface them in reporting.

Finally, we created a publishing layer which houses our production cube.



Step 5: Dashboard Building

The requirements gathering informed how we structured our backend model, but we still needed to create the views to surface up this data. The actual building of the dashboards was the easiest part of the process. The challenge part was ensuring that all the user voices we captured during the requirements gathering guided the design of the dashboards.

Why Power BI was so easy to build on:

- Connects seamlessly to data sources and refreshes automatically.
- Can build and change visualizations by dragging and dropping onto the page.
- Many formatting options available for visualizations.
- Ability to pin key visualizations to a dashboard and create executive level views.
- Visualizations allow for export of data to excel. This is key for our field marketers who often need to do further analysis on the raw data.
- Many visualization options, including KPIs and custom visuals.

Even though this was the easiest step of the process, we still ran into several challenges that caused us to tweak our recipe. Ultimately, the following challenges helped us create a finished product that serves all our users:

- **Data granularity requirements differ amongst users.** Some users only need to see the top-level metrics, where others must be able to drill in and find a single marketing tactic. The reporting must allow for drilling up and down in the data. This was easily solved for with Power BI's drill functionality available on most visualizations.
- **Varying levels of familiarity with metric definitions across the user base.** The metric names and graph titles must to be user-friendly and appropriate for a wide-ranging audience. We created a data dictionary to help users. Power BI allows for embedding of hyperlinks so you can create a link to an external Data Dictionary, or add the Data Dictionary as a report directly in Power BI.
- **Maintaining focus on key metrics while allowing for deep data exploration.** For us, it was important to ensure the funnel view of leads from acquisition through the win persisted throughout the reports, so we added this to the top of most of the report pages. This helps keep the focus on the progression of the leads through the funnel, while the visualizations below these tiles provided detailed cuts of the same data.

- **Providing all necessary functionality without overloading the user visually.** There are so many options available in Power BI, we spent some time rationalizing how best to surface our data. One point of contention was the use of slicers vs. filters. Slicers are actually a type of visualization in Power BI that can be pinned directly to a page. They filter all the data on the page and are easy for the user to access. Filters show up on the right-hand side of the reports in Power BI and can be set to filter the entire report, a single page in the report, or a single visualization on a page. We made the key items, like customer segment and business priority available in both the slicers and as report-level filters so that a user could easily find them, but could also apply them to the entire report when necessary. Fields that we knew would be used less frequently, or by only a subset of users, were added to the filters only.



Learnings

Like any good chef, we learned a great deal about our business and our data in the process of building the recipe for this robust reporting platform. Most of the challenges we encountered were in aligning business definitions. Building out the data model and the Power BI dashboard was relatively easy.

Landing on agreed-upon definitions and taxonomy across business groups was more of a challenge. Bringing all the right people to the table as early as possible to land on metrics and business definitions are key ingredients to ensuring that you don't spend time building something that only works for a portion of your business.

We made tradeoffs between breadth and depth. To report across such a vast amount of systems and businesses, the data needed to be leveled-up at some points. We chose to cover all lead stages and business groups at a higher altitude. The decision of how granular reporting should be will depend on the audience.

For a marketing team that works across products and segments, the greatest value was in being able to see broadly versus in-depth. Depth reporting is supported in our underlying data model and we can conduct ad hoc analysis for users when necessary. We have also built subsequent dashboards that drill into specific businesses and account groups at a deeper level.

Having flexibility in the data visualization platform is key. We were able to update the dashboard and reports in real time as we reviewed with users. Having no development effort on the visualization side also allowed us to focus on creating a robust and reliable data cube. We can now easily stand up new reports and dashboards in Power BI of this same data cube.



After marinating with the results of the dashboard we've seen so far, and the feedback we've received from our internal teams, we're excited to see what the **modern marketing dashboard** can help your business achieve.





©2017 Microsoft Corporation. All rights reserved. This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY. This document is provided "as-is." Information and views expressed in this document, including URL and other Internet website references, may change without notice.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

microsoft.com