

The heartbeat of modern marketing: Data activation & personalization

Julien Boudet, Brian Gregg, Jason Heller, and Caroline Tufft

Technology has finally advanced to the point where marketers can use real-time data in a way that is both meaningful to customers and profitable for companies.

We've come a long way from "People who bought this, also bought that."

Consider the experience of a representative customer we'll call Jane. An affluent, married mom and homeowner, Jane shops at a national clothing retailer online, in the store, and occasionally via the app. When visiting the retailer's website in search of yoga pants, she finds style choices based on previous purchases, the purchases of customers with profiles similar to hers, and the styles of yoga pants most frequently purchased on weekends. She adds one of the offered yoga pants to her shopping cart and checks out.

With the exception of a follow-up email, most interactions with the customer stop there. But here's what this example looks like when we activate Jane's data: Three days after her online purchase, the retailer sends Jane a health-themed email. Intrigued, she clicks the link and watches a video about raising healthy kids. One week later, she receives an iPhone message nudging her to use the store's mobile app to unlock a 15 percent one-day discount on workout equipment. Though she has never bought such items at this retailer, Jane takes advantage of the offer and purchases a new sports bag. What began as a simple task of buying yoga pants ended up being a much more engaged experience.

Such data-activated marketing based on a person's real-time needs, interests, and behaviors represents an important part of the new horizon of growth. It can boost total sales by 15 to 20 percent, and digital sales even more while significantly improving the ROI on marketing spend across marketing channels: from websites and mobile apps to—in the not-too-distant future—VR headsets and connected cars.

Customer data platform: Solving the ongoing challenge of true personalization

Companies regularly experiment with testing the impact of varied customer experiences, but they do it in isolation. When they do try to scale, they smack against the challenge of understanding what to prioritize. Going back to Jane, do marketers target her as a mom, a yoga enthusiast, or a homeowner? What happens when tests are running against all three segments? Is she part of a new microsegment that combines attributes and signals across all three segments?

This is a challenge that has continued to plague marketers, despite the promise of solutions such as customer-relationship management (CRM), master-data management (MDM), and marketing-resource management (MRM). These solutions can help companies consolidate and streamline data, manage segmentation, organize workflow, and improve customer relationships. But they don't take full advantage of digital signals customers provide. Instead, relying on antiquated "list pulls", basic segmentation, and campaigns, all lack the automated decision making, adaptive modeling, and nimble data utilization to scale personalized interactions.

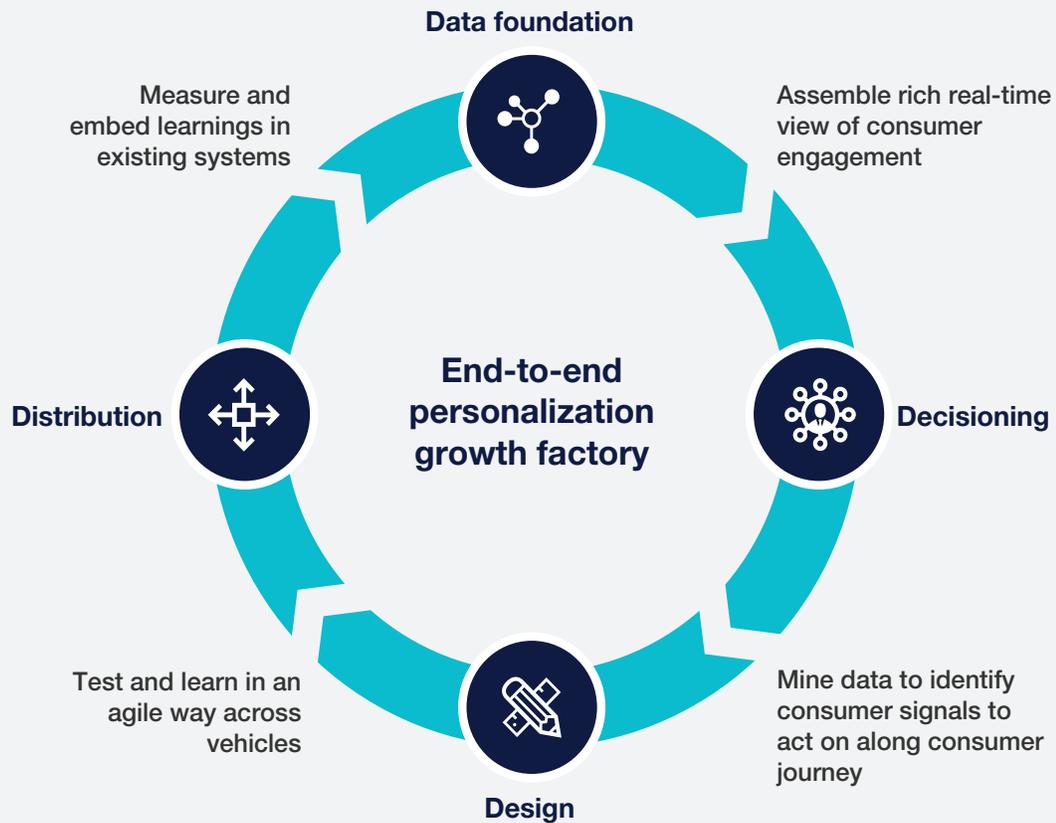
Enter the Customer Data Platform (CDP)—a data discovery and "decisioning" (i.e. automated decision making) platform. The CDP makes it possible for marketers to scale data-driven customer interactions in real time. And while CDP hasn't really broken into the Gartner Magic Quadrant or Forrester Wave, it is gradually becoming an industry-standard concept, with a small but growing cadre of third-party platforms emerging that will soon shape the category.

Four steps to effectively activate your data

Incorporating a CDP into your organization—whether piggybacking on an existing master data-management or customer-relationship-management system or starting from scratch—requires mastery of four areas (see Exhibit):

Exhibit

Building deeper 1-to-1 relationships with consumers at scale



McKinsey&Company

1. Data foundation: Building a rich view of customer

Many companies have the elements of a relatively complete view of the customer already. But they reside in discrete pockets across the company. Just as a recipe does not come together until all the ingredients are combined, it is only when data is connected that it becomes ready to use. The CDP takes the data a company already has, combines it to create a meaningful customer profile, and makes it accessible across the organization.

“Feeding” the CDP starts by combining as much data as possible and building on it over time. Creating models that cluster customer profiles that behave and create value in similar ways requires advanced analytics to process the data and machine learning to refine it. Over time, as

the system “learns,” this approach generates ever-more-granular customer subsegments. Signals that the consumer leaves behind (e.g., a site visit, a purchase on an app, interest expressed on social media) can then expand the data set, enabling the company to respond in real time and think of new ways to engage yet again. Furthermore, the insights gleaned extend beyond a customer’s response to a specific campaign, for example by driving more targeted product development.

A number of companies we’re familiar with, struggling to truly understand their customers who make infrequent purchases, combine their own CRM data with Facebook consumer data to build look-alike models. This helps identify the highest-value prospects most likely to buy in their category. Increased targeting through display ads on and off Facebook can yield 50 to 100 percent higher returns than from the average Facebook audience. Mapping third-party data (when it exists) to customer segments via a data-management platform (DMP) can enhance the experience for both known and anonymous digital consumers, leading to improvements in engagement and conversion, measured in net promoter score, acquisition, and lifetime value.

2. Decisioning: Mine the data to act on the signals

The decisioning function enables marketers to decide what is the best content to send to a given customer for a given time and channel. Customers are scored based on their potential value. A set of business rules and regression models (increasingly done through machine learning) then matches specific messages, offers, and experiences to those customer scores, and prioritizes what gets delivered and when. This allows companies to make major improvements in how they engage with their customers by developing more relevant, personalized engagement, within a single channel or across channels, based on a customer’s behavioral cues. Those signals can be basic, such as “cart abandoned” or “browsed but didn’t buy,” or more nuanced, such as activity by segment and time of day, gleaned from mining customer data. In effect, these signals become triggers that invoke an action. A decisioning engine develops a set of triggers and outcomes based on signals and actions the company takes in response.

For example, one multichannel retailer discovered that many consumers made a purchase on the website just once per year. Further analysis revealed those same customers tended to return to browse the site a few days after purchase. The company now takes advantage of this window of opportunity to send tailored, targeted messaging, rather than risk losing the customer for another year. This approach doubled the open rate of its emails—from 10 to 15 percent for generic targeted communications to 25 to 35 percent for real-time, “trigger-based” communications acting on consumer signals.

More sophisticated companies build up a decisioning model that works across all distribution channels. That requires advanced modeling and analytics techniques to identify the impact of one channel on another as a customer proceeds along his/her decision journey. A travel company took this approach recently and saw coordinating messages across channels drive a 10 to 20 percent incremental boost in conversion rates and customer lifetime value.

Effective decisioning is based on repeated testing that validates and refines hypotheses and outcomes. Over time, these can become increasingly sophisticated as models and algorithms build on each other. One telecommunications company has been testing different offers to different groups: millennials, customers in specific cities, previous owners of a specific device, groups of relatives, and people who viewed a specific web page in the last three days. As complex as this may seem, a semi-automated decisioning engine prioritizes the offers and experiences proven to have the highest rate of return. This allows the telco to scale the results of dozens of tests without fear of inconsistent customer experiences or conflicting offers.

3. Design: Crafting the right offers, messages, and experiences at speed

Understanding your customers and how to engage them counts for little without the content to actually deliver to them. Designing great offers, however, is hampered by the fact that functions and departments within companies tend to operate as mini fiefdoms. The owners of each channel test and engage consumers exclusively within their own channel. Real benefits can only occur when companies shift to “war rooms” of people from relevant functions (marketing, digital, legal, merchandising, and IT/DevOps) who focus on specific consumer segments or journeys.

These teams have clear ownership of consumer priorities and responsibility for delivering on them. The cross-functional team continually develops new ideas, designs hypotheses for how to engage customers, devises experiments, and creates offers and assets. Analytics help size opportunities, test impact, and derive insights from tests. That content is then tagged, so that it can be associated with a trigger and be ready to go when needed. Just three months after launching its war room, one large multichannel retailer saw its testing speed go from 15 to 20 weeks to two to three weeks, and testing volume increase from four to six per month to 20 to 30 per month.

4. Distribution: Delivering experiences across platforms

Distribution systems are simply the “pipes” that deliver the ad or content to the end user (eg: ad server, DSP, or content management platform). Often they can be quite manual and just blast out communications to wide segments of people with little tailoring. But connect the CDP engine, with its predetermined triggers and tagged content, to these distribution systems, and a formerly blunt marketing instrument becomes a far more directed one sending specific messages to distinct customer subsegments across all addressable channels. Sophisticated businesses have developed a library of APIs to help tie the CDP into the “martech stack”—the marketing technologies that deliver and track experiences. Integrating the stack this way creates a feedback loop that sends customer response, engagement, and conversion data back into the CDP.

Implementing the data-activation framework

Not all data-activation efforts are created equal. We recommend using a case-driven approach, maintaining a backlog of tests ranked by opportunity, quantifying the impact of each potential use case, and balancing it with the level of effort required to implement it.

Data-activation self-assessment

This self-assessment can help company leaders develop a benchmark for measuring their progress on their data-activation journey.

Data foundation

How comprehensive is your view of the consumer across all your internal data sets, and how close to real time are those data feeds being updated?

Lagging: We do not use any data for personalization.

Basic: Data-driven personalization is mostly focused on transaction data, and/or anonymous third-party data. Data is manually updated daily or weekly.

Leading: Rich view of consumer across most touchpoints (e.g., transactions, media, clickstream, servicing/care). Data is actively used for personalization. Data is real time or refreshed multiple times per day.

Decisioning

What types of models are you activating across channels? Who manages your models?

Lagging: We are not using any propensity models to enhance targeting or to trigger personalized experiences.

Basic: We have basic propensity models that are used on a limited basis and not used widely in digital. We have limited or no dedicated data-science resources to manage models.

Leading: We have multiple propensity models to predict value creation or destruction for a given customer interaction, and most digital messaging is triggered by these propensity scores. Our models are managed by in-house data-science resources. We currently or soon will use machine learning to further fine-tune models.

Design

How often do you test offers and messages?

Lagging: We do limited tests and do not update our offers frequently.

Basic: Tests are set up and deployed manually. We analyze performance weekly or monthly and optimize periodically.

Leading: We run triggered A/B and/or multivariate tests daily.

Distribution

How are your marketing technology platforms integrated with your data systems?

Lagging: We have not optimized our martech stack and/or rely solely on the platforms our agency manages on our behalf.

Basic: We manually batch upload data to our martech systems, and we are able to deliver personalized experiences to broad customer segments in some channels. Response and transaction data is batch delivered back into the CDP.

Leading: We have API connections between our customer-data platform and our martech systems. All response and transaction data is fed back in a closed loop into our customer data platform.

Unlike a wholesale IT transformation, deploying a CDP isn't a replacement of current customer-data systems, but rather an operational solution that can piggyback on existing systems. In our experience, many marketers already have a large part of the marketing-technology equation in house; they're just not using it properly.



The promise of data-activated, one-to-one marketing is not only possible but is now increasingly expected by today's customers. It is now the key to transforming simple customer transactions into enduring relationships.

Julien Boudet is a partner in McKinsey's Seattle office, Brian Gregg is a partner in the San Francisco office, Jason Heller is a partner in the New York office, and Caroline Tufft is a partner in the London office.

Special thanks to Sean Flavin, McKinsey expert in personalization, for his valuable contributions to this article.