


# **What Smart Companies Know About Integrating AI**

by Silvio Palumbo and David Edelman

From the Magazine (July–August 2023)



Edoardo Tresoldi

**Summary.** AI has the power to gather, analyze, and utilize enormous volumes of individual customer data to achieve precision and scale in personalization. The experiences of Mercury Financial, CVS Health, and Starbucks debunk the prevailing notion that extracting value... [more](#)

**Mercury Financial, founded in 2014,** is a small fintech with a big mission: helping customers restore their credit profiles. Jim Peterson, Mercury’s CEO and a finance veteran, knew from the outset that AI was crucial for creating the personalized customer

journeys that would be at the heart of Mercury's offering. So in 2021 the company began searching for an AI-driven engine that could give every customer the right nudge at the right time through the right channel and in the right sequence. The nudge might be a push to split up payments among multiple credit cards, or a gentle warning that the customer was nearing a credit limit. Some might respond to a text message, others to an email. Some might respond best two weeks before an action date, others two days before. Any one of those elements—or, more precisely, the particular combination of them—could spell the difference between a fully engaged customer and a deeply annoyed one.



**Listen to this article**

To hear more, download [Noa app](#) for iPhone or Android.

Mercury, a former BCG client, is not in the business of building technology, so its leaders decided to begin with open-source AI. Given that most such tools charge according to usage, the upfront fixed costs would be low. Mercury concentrated on how to integrate available AI solutions with its content-management, fraud, and eligibility systems, and many other front- and back-end systems. The company then automated its marketing processes, again drawing on available AI tools but using its own code for all the tests needed to learn what worked for whom and for tracking past results. The system it built focused on managing hundreds of variables for targeting purposes and creating content in a microgranular way. Within six months the pilot had generated a 10% improvement in actions taken as a result of the fintech's messages. Mercury knew it was on to something big.

AI is required to achieve precision and scale in personalization. It can gather, analyze, and use enormous volumes of individual customer data and tailor the customer journey at every touch point. Mercury's experience, and the experience of CVS and Starbucks (which we'll explore in detail), debunks the prevailing notion that extracting value from AI solutions is a complicated technology-building exercise. That thinking keeps companies from capturing the power of AI. They needn't build it; they just have to properly integrate it into a particular business context.


When you recognize the value of focusing your resources on integration and process change, it sharpens what you look for in an AI system. You begin to understand the importance of seeing your data and the design of your tech architecture as competitive assets. And you push the rest of your organization to drive more testing that can feed the intelligence of your AI system.

But AI is probably only about 10% of the secret sauce. The other 90% lies in the combination of data, experimentation, and talent that constantly activate and inform the intelligence behind the system. Personalization is the goal; it's what constitutes a company's strategic brawn. The technology is merely the tool for reaching it. In this article we'll present what it means to integrate AI tools and what it takes to continually experiment, constantly generate learning, and import fresh data to improve and refine customer journeys.

## **Rethink How You Acquire Technology**

In more conversations than we can count, we find ourselves disabusing executives of their notions about what creates an AI advantage. A company needn't aspire to be another AWS, Microsoft, Google, or Adobe—all builders of core AI tools that are, after all, in the business of selling them. Familiar masters of AI, such as Uber, Netflix, and Spotify, may research and design new solutions, but generally they do so to extend applications that accommodate their uniquely huge scale or to perform specific

functions not otherwise available (such as movie-frame analysis in the case of Netflix’s recommendation algorithm). But few companies outside the tech world are monetizing their own digital innovation. For them, innovation involves offering a fresh solution atop a base of digital capabilities. Competitors all have access to the same AI, yet business outcomes vary profoundly. One critical difference is the data a company feeds it. Competitive advantage hinges on unrelenting data collection, curated transformation or enrichment, and feeding the AI libraries that inform next-best-action capabilities. The marketer’s job is to creatively apply those AI-driven recommendations to marketing campaigns and iteratively learn from them.



Edoardo Tresoldi uses the form, framework, and transparency of wire mesh to illustrate a dialogue between architecture and the world.

A wealth of open-source technology exists today, including most AI tools—broad ones, such as GPT-4 from OpenAI, and full-fledged libraries (applications written in open-source languages that are packaged for a specific use, such as XGBoost, for training a specific type of machine-learning model). Big tech makes many of its libraries or task-specific tools available—Meta (a BCG client), for instance, with its Prophet library for forecasting, and Airbnb with Airflow, a workflow management platform for data engineering pipelines. AI capabilities are embedded in many

common customer-experience tools, such as the “experience clouds” of Salesforce and Adobe (both companies are BCG partners). Those are constantly improving too: Thanks to application programming interfaces (APIs) and the architecture of modern tech systems, it has become easier to get systems to talk to one another, as we’ll explore later.

Another bit of debunking involves how to begin. Contrary to general thinking, a big-bang approach is not necessary. As with true agile approaches, it is smart to start out narrow, focusing on specific use cases—customer onboarding or early engagement, for example—rather than trying to redesign the entire customer experience from scratch. It is smart to choose an area in which you can get real traction with AI and then gradually expand its use.

None of this is meant to suggest that implementing an AI-based customer journey is easy. And three common shortcomings can make it especially challenging, even when you have the right AI solution.

**Limitations on data capture.** These can occur when a preponderance of the data is unstructured, as in health care, or unrecorded, as in the hospitality industry. No one inputs feedback from the cards left for guests in hotel rooms, which may contain immediately valuable and addressable information. If a hotel recorded your preferences regarding bed type, cleaning schedule, food allergies, facilities at each property that you’ve used, and so forth, your next stay could easily be tailored to your taste.

**Lack of connectivity between the AI solution and marketing.** A supermarket chain might assume which customers prefer nonmeat foods or which ones are bargain hunters, without explicit confirmation. AI provides access to such personalized predictions automatically, at scale. But to test the assumption, you need a martech system into which you can channel that data in the form of a dialogue. The insight itself can’t create value

unless it's activated.

**Inability of the system to automatically scale.** An abundance of customer insights is of little use if your systems can't convert every one into tailored actions for individual customers. And you can't rely on manual input or analysis to develop a perfect one-off personalized campaign. Without a robust, scalable system, insights will be suppressed or used too narrowly within just one function.

## **The Four Keys to Smart Integration**

In our experience guiding a variety of companies as they developed AI-based customer journeys, we've pinpointed four defining characteristics of a smart integrator: *clarity and alignment of goals, sound data instrumentation, a loosely connected tech architecture, and an experimental culture*. Notably, none of them includes having a better AI algorithm, though they do call for a better-trained one.

**1. Clarity and alignment of goals.** AI-based marketing requires clear optimization objectives for every use case, and those goals need to be reasonably narrow. Broad general objectives, such as "accelerate sales growth," make it impossible to know how to attribute results. A more appropriate objective for AI might be "minimize wait time," "lower the incentive cost per sale," or "make a suggestion the customer will accept." If there are multiple actions the customer might take, you'll need to decide in advance: Do we want the AI to optimize for the most valuable of the options or for the ones the customer is most likely to take?

AI applications make predictions primarily on the basis of historical data. If AI focuses too broadly, it loses the power to optimize; if it faces a situation for which there's no precedent, its predictions will be inaccurate. Rather than one AI engine to govern everything, you may need multiple AI engines to build a broader customer experience. One leading health care company, for example, has an engine to identify health plan members most

at risk for an acute incident; another to recommend how to approach a member; another to suggest a test design for experimentation; and still another to optimize the whole portfolio of outreach types within a specific budget.

**Personalization is the goal; it's what constitutes a company's strategic brawn. The technology is merely the tool for reaching it.**

One reason Starbucks has been successful in AI-based personalization is its goal clarity. In driving a sale from a marketing outreach, the company takes into account the fact that many of its product teams want their offerings promoted, so it has decided to optimize for the touch most likely to compel a customer to buy. (Along with actual response data, Starbucks captures implicit interest—for example, what the customer browses and whether she hovers over an image, clicks on a description, or returns to the same page three times in a week.) Beyond its constant experimentation, the company focuses on how marketing can boost total net incremental revenue rather than on optimizing a specific beverage's profit and loss. That allows for a holistic definition of success.

**2. Sound data instrumentation.** The mechanisms that record, organize, and share data on customer interactions, the company's associated actions, and outcomes across touch points are the nuts and bolts of a company's AI personalization program. This data instrumentation includes everything from call center logs and data sourced from second- and third-party relationships (such as channel partners, media companies, and data brokers) to automation software that generates and tracks digital communications (such as that from Salesforce, HubSpot, and Illumin).

A company needs “receptors” to capture information about a customer’s every interaction across every channel, and that information must be remarkably granular. For example, in a call center interaction, the company would want to know more than the purpose of the call: Did the caller sound annoyed or confused? Was the caller’s computer open during the encounter? Was the caller following the rep’s instructions or being distracted by something else?

When a company sends a test email, it must be able to monitor the phrasing, the image that is embedded, the font size and colors used, and any other variable that could possibly affect the customer’s response. A tagging architecture—for capturing and classifying the metadata of these interactions and communications—allows for testing and learning.

When one of us (David) served as chief marketing officer at Aetna (now part of CVS Health), the primary goal was to get people to take health-promoting actions such as getting a flu shot and taking their medications regularly. The company knew who it wanted to take which actions, but it didn’t know how to get those individuals to comply. An enormous amount of experimentation was required: changing the message (both the offer and its creative content), testing an incentive, altering the time of day sent and the sequence of messages, and so forth—none of which would have been feasible without mechanisms in place to set up tests and track the microvariables that drove responses. Aetna had to make sure that every interaction was comprehensively tagged with descriptors of what was in it (color of the text, nature of the image, tone of the language, specific words) and that every point of interaction was captured in minute detail: when the customer interacted, how many clicks there were, whether the customer left and returned, where the journey went after the first click. This required embedding the code in all the places of interaction—web, app, email, SMS, and so forth—and integrating the data flows that came back from them. That was where the company focused its time—not on building new AI. As AI systems



do an increasingly better job of writing code to facilitate data integration and automatically tagging marketing assets, the work becomes easier.

**An abundance of customer insights is of little use if your systems can't convert every one into tailored actions for individual customers.**

Data from third-party sources, such as weather, power outages, demographic and psychographic data, and general health data about a zip code population, provides further context. The broader and more granular the information, the richer the models can be. That richness fuels your performance edge. Do customers respond right away or take days? Do they scroll down for further information and then click, or jump off immediately? Which incentives do they respond to most promptly? Data instrumentation enables the continual testing and experimentation that provide answers, help the system keep learning, and inform the company's understanding of what's needed to get customer A to respond versus what customers B, C, and D require.

**3. A loosely connected tech architecture.** The customer-experience technology stack consists of a prediction engine, a sequencing (or experience management) engine, a content engine, channel delivery engines, and an experimentation and analysis engine. In addition, the AI draws from five or more systems to stitch together a customer journey: marketing, customer service, product usage, billing, online channels, and sometimes a retail store. Given the likelihood that new capabilities will be added and that several AI engines may need to be plugged in, it is best to design the stack in a modular way.


AI provides the intelligence and the computational speed and scale to fuel operations that are increasingly driven by automation. Thus the technologies that carry out the automation must be able to accept the AI's signals and feed information back into it to help it improve. A loosely connected tech architecture, in which systems work with but are not dependent on one another, is ideal.

Publicly available application programming interfaces—which give developers access to proprietary software through a simple, versatile standard of communication—enable this modular architecture. APIs are the lingua franca of digital dialogue across disparate platforms. A simple example is an API that links a company's CRM system to a cloud-based phone system, allowing a call center agent to immediately phone a newly generated lead without having to exit the CRM software. With open APIs, information can be moved seamlessly, models updated easily, and new capabilities added in a modular way. A loosely connected architecture allows companies to mount faster competitive maneuvers, because they can easily swap out components the moment new capabilities become available, with minimal switching costs.

To see how a loosely connected architecture enables the integration of the various elements of the stack and supports personalization at scale—the whole point of smart integration—consider Comcast. The company realized it would be impossible to formally connect all its customer data systems. But the systems had to be linked somehow if AI decisions were to drive what to send to whom through which channel and feedback data were to be returned to the AI.

Pointillist, Comcast's AI decision tool, does double duty. First it matches a customer's data from across all the company's systems—its app interactions, the call center, product use logs, and so forth—and creates an integrated view of the customer. Then, operating like middleware, it knits together multiple databases

into one integrated database, precluding the need to create yet another, formal database. Pointillist finds all relevant information about a customer and time stamps it to inform the company's view of the customer's journey. In real time the system finds the matches, builds journey maps, identifies anomalies that need attention, and drills down into root causes. Using the standards in Genesys, Comcast's central customer system (and a client of David's), Pointillist channels the information about the customer into all customer-facing interaction systems, which are compatible with Genesys's API.



Edoardo Tresoldi

Customer interactions are tracked in chronological, time-stamped order. The system will alert Comcast in real time that Jane Doe is on the mobile app and having difficulty accessing service plan information. Even before she calls Comcast, it can send her a text message suggesting a quick fix. If she doesn't act, or if that action doesn't resolve the problem, the software indicates that she will probably call—so if she does, the call center system already knows what her problem is. It can determine whether a particular automated response could help her, in which case it provides automated prompts, or whether she needs to

speak with a human—which, as the costlier solution, is always the last resort. The beauty of this system is its real-time responsiveness and the speed and frictionless experience it provides.

Over time, as Pointillist continues to gather information, it allows Comcast to constantly test new and more effective ways of handling any given problem: to determine the ideal message, to find the best way to intercept the customer, and even to track whether customers who experience problems actually contact the company or not. In its first 18 months this AI-driven system saved Comcast more than 10% in call center costs. The feedback the company gets and the improvements it has made to the online customer experience create fewer reasons to call. The system has helped Comcast's Xfinity Mobile achieve top ranking by J.D. Power for several consecutive years.

**4. An experimental culture.** AI stokes creativity by allowing a company to test ideas rapidly and to do more at scale. Furthermore, it learns from the past, across millions of data points, unlocking innovation quicker than a human could. But AI does not invent; it just predicts, on the basis of past patterns. Marketers invent, and the AI learns what works, for whom, when, and how. Invention requires a culture that values experimentation and risk-taking.

All else being equal, a company's experimental data is the source of its competitive advantage. Leaders must recognize that a test-and-learn mentality is essential for translating that data from raw material into currency. That means accepting that experimentation has opportunity costs and that by definition, some experiments won't work. But even failed experiments offer worthwhile lessons.

Agile teams are essential in this effort. They have the prowess and the creativity to pursue new ideas and keep improving on them over time. Using agile work methods means you plan for what you

can but leave room to pivot quickly according to what the experiments reveal. Leaders therefore need to examine their incentives and operating plans to ensure that they promote agile practices, allowing for tests that don't yield successful results and enabling the flexibility to change course. But leaders must also espouse an experimental mindset and have the financial license to take risks. Experimentation requires control groups to validate test results. You must be able to cut your losses and double down on winners. You need to finance operational changes and new incentives that have to be tested. The tighter the financial leash, the tougher it is to innovate. You cannot experiment if you can't take risks.

**When a company sends a test email, it must be able to monitor any variable that could possibly affect the customer's response.**

Starbucks's dominance in incentive-based marketing depends largely on its strategy of ongoing experimentation and its commitment to supporting it with the necessary resources. The company relies not only on digitally collected data and digital tech but also on its retail experience. Its tech is 100% composed of open-source tools and languages. Starbucks focuses on constantly enriching its data set and connecting it to its technology architecture, not on developing the algorithms.

Every week Starbucks runs randomized trials on a subset of its customers to test and learn and to reinforce or challenge marketing hypotheses. Customers who typically visit retail stores are incentivized to interact via the Starbucks app and leave a digital breadcrumb behind. The app gathers rich data: when, what, and where customers ordered; whether they placed their regular order or added an item; and details from the stimulus

message (wording, incentives, when it was sent, when it was opened). All those details are harnessed by the app to tailor marketing outreach at the individual level.

The company can experiment with existing products as well as with new ones. For instance, it might choose to improve the messaging or the formatting—changes that can be as granular as the language, color, or image embedded in the message. It might want to test pricing or product recommendations on customer segments or on the most-effective channels. The more parameters, the more test permutations. Testing can take time, which is yet another reason to make experimentation a continual effort.

Its AI-integration journey brought Starbucks a 45% increase in net incremental revenue (sales attributable to marketing, excluding discounts) within four months of running a simplified proof of concept. At the 12-month mark, after randomized trials, a steady stream of fresh data, further testing, and expanding the program to the full customer base, Starbucks saw a 150% increase. As more channels, more offer combinations, and more permutations were added, that number reached 300%.

### **Smart Integration in a Complex Customer Relationship**

Unlike Starbucks, CVS Health operates in a regulated industry in which offering incentives to customers is prohibited. Just as important (if not more so), the variations across its patient-customer base and their health situations, and in the behaviors and actions the company wants to elicit, are considerably more complicated than those of a retailer like Starbucks.

CVS's journey began within Aetna, which it acquired in 2018. Aetna wanted to motivate Medicare members to take more-healthy actions. Helping people get healthier would also reduce costs for all and enable the company to improve customer service.

## **AI does not invent; it just predicts, on the basis of past patterns. Marketers invent, and the AI learns what works, for whom, when, and how.**

Aetna thought it could encourage its customers to take dozens of health actions; however, the company understood that, for example, getting patients to refill their prescriptions requires a different path for each one, and even a different one over time for the same patient. The data that could predict the best way to influence behavior at a given moment might come from recent claims, responses to marketing campaigns, customers' use of digital tools, basic demographics and regional variations, and clinical protocols, among other sources. In order to tailor its communications, the company has invested in the ability to change every aspect of an outreach (channel, timing, frequency, message, language), thereby creating millions of potential permutations. Testing them means customizing common martech tools to be flexible enough to capture and use this expanding range of data. But to optimize across all the potential variants, commonly available AI tools needed more historical data than was available at the time. So Aetna developed an approach to test all those permutations on a gigantic scale, with controlled experiments.

What might a controlled experiment look like? One example would be testing the relative impact of a text message, an email, and an interaction at the pharmacy counter while also varying the message, the timing, and the incentive. Because the number of permutations adds up quickly, the company would use an AI tool to examine the range of test options and determine the minimum number of controlled tests sufficient to track the relative impact of each permutation.

Instead of creating a novel algorithm, Aetna chose to use open-source algorithms and assembled a team to conduct ongoing exploration. Business specialists sit alongside data scientists and engineers to feed the system novel ideas and variants—such as new ideas for copy or images, new incentives, or new ways to reach out to a customer—and then analyze the results and course-correct. Aetna also hired behavioral economists to help develop the contact strategies for various situational profiles: some emphasizing risk of loss, others emphasizing the importance of a more-healthy lifestyle; some reminding a customer of family, others reinforcing a health message using local community data.

---



Subscribe to our Monthly Newsletter

## **Strategy & Execution**

The tools you need to craft strategic plans — and how to make them happen.

[Sign Up](#)

---

Ultimately, this approach has created a combination of human beings and individual technical systems uniquely tailored to power CVS's business model. Downstream, this intelligence engine fuels an array of channel-specific systems where interactions are executed: email, CRM, call-center scripting, text messaging, app notifications, and even pop-ups on pharmacists' computer screens. Those execution systems funnel the results of the interactions back into the AI's algorithms to keep feeding their learning. Over time some of the data becomes more complex: For example, a call center discussion might be dissected using natural language processing to understand the caller's disposition, choice of words, comprehension, and so on. By focusing on data flows, the potential for innovation, experimentation, and architecture, CVS has turbocharged open AI



tools, generating substantial incremental annual profit margin from lower medical costs, better service ratings, improved health outcomes, and new cross-selling opportunities.

...

To drive competitive advantage with AI, you need to integrate your internal systems with external ones—first to collect accurate customer data and then to present the resulting insights as personalized offers. Both processes must be carried out with relentlessly expanding scale and scope, continually adding new variables and increasingly granular detail. For the executive who must begin thinking like a smart integrator, this approach calls for a new leadership model with new priorities to take advantage of the infinite possibilities. Getting integration right increasingly drives a superior customer experience, and it will be the decisive factor in how brand equity is built.

A version of this article appeared in the July–August 2023 issue of *Harvard Business Review*.

## SP

**Silvio Palumbo** is the managing director and a partner of BCG's AI and Advanced Analytics practice.

## DE

**David Edelman** is an executive adviser and a senior lecturer at Harvard Business School.