

Technology, Media & Telecommunications Practice

# Why AI-enabled customer service is key to scaling telco personalization

Telcos increasingly rely on customized commercial marketing. But integrating those efforts with service-related outreach into a comprehensive next-best-experience engine is critical for growth.

*This article is a collaborative effort by [Lars Engel Nielsen](#), [Lars Fiedler](#), Dominika Kampa, Stanislav Kucherenko, [Julian Raabe](#), Tim Schenk, and [Davide Schiavotto](#), representing views from McKinsey's Technology, Media & Telecommunications Practice.*



**A growing number of telcos** are harnessing commercial next-best-action engines that use real-time, data-based insights for objectives such as boosting cross-selling and upselling, reducing churn, improving retention, and acquiring new customers. But these marketing efforts are at risk of falling flat without [predictive, proactive, and customized outreach](#) about service-related issues and concerns. Put simply, customers frustrated with spotty connectivity or unexpectedly high bills are less likely to respond positively to a new offer or promotion.

To thrive in the [challenging market ahead](#), with disruptive new technologies, digital-native competition, and rising customer expectations, creating a next-best-action engine for service issues is critical. But service is still just one piece of what could be a more comprehensive approach. Rather than trying to use two separate next-best-action engines, telcos can work on integrating commercial and service-related capabilities into a single, comprehensive next-best-experience engine to communicate with customers in a highly personalized way—on the right issues, at the right time, and through the right channels.

In addition to the cascading benefits of superior customer experience, we estimate that within three years, this approach will allow operators to cut costs associated with certain use cases by up to 30 percent, boost overall B2C revenues by 2 to 4 percent, improve customer satisfaction by 10 to 20 percent, and reduce early-life churn by as much as 30 percent—on top of [gains from ongoing commercial customer value management activities](#).

There is a potentially significant upside for early movers, which will set new standards for what customers can expect from their telcos. At the same time, operators that are slow to venture into this new frontier risk lagging further behind over

the next three to five years as they struggle to satisfy customers' evolving demands.

## Service matters

Digital-native service and content providers are continually redefining customer expectations and, in the process, weakening telcos' relationships with their customers. This dynamic contributes to the commoditization of connectivity, which raises the stakes of getting telecom service right. Understandably, there is a strong correlation between customer experience and retention—with overlap as high as 80 to 90 percent in some markets (exhibit). Even if a telco's prices are the lowest available, customers reeling from poor customer experience, faulty connectivity, or surprise charges on their bills are unlikely to respond well to any marketing campaign.

Customers who spend long periods on hold, are transferred from agent to agent, or have to make multiple calls to resolve a single service issue are likely to remember the arduous process even if they ultimately get their desired outcome. One Western European telco discovered that customers who had to make two or more phone calls to resolve connectivity issues were almost twice as likely to walk away at the end of their contract.

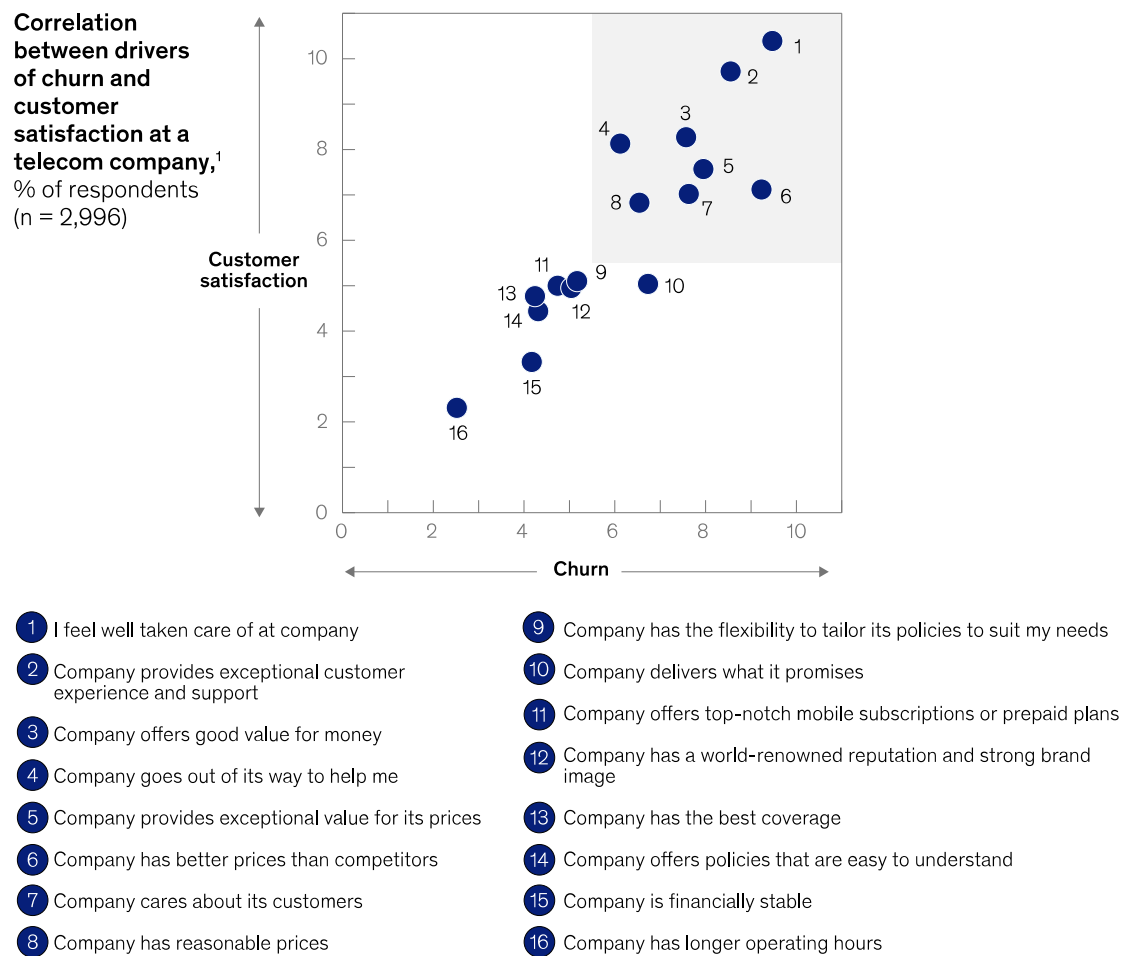
Moreover, when telcos' commercial and service capabilities operate in silos, as is generally the case, customers may be flooded with irrelevant, even irksome communications. As a result, some customers may go so far as to revoke their consent to receive marketing communications. One operator discovered that its various units were sending customers as many as five messages per week, with some customers receiving cross-selling offers and information about network outages within the same hour.

An integrated next-best-experience engine that takes a holistic approach to personalization can not only substantially improve the experience of

customers but also reduce the costs associated with serving them. This can be done by replacing reactive, inbound-focused customer interactions

Exhibit

**There is a strong correlation between telco customer satisfaction and retention—with overlap as high as 80 to 90 percent in some markets.**



<sup>1</sup>Derived importance methodology = Johnson's Relative Weights (JRW) using importance to churn (1–5) as the dependent variable and agreement with journey satisfaction (1–6) as the independent variable.  
Source: Japanese Telecom Journey Pulse, August 2022

McKinsey & Company

with highly automated outreach that is both proactive and preemptive, and by ensuring that a telco speaks to its customers with one voice.

## Driving the engine

The next-best-experience engine guides and powers all telco interactions with customers. It comprises four main layers: data, analytics, decisioning, and channel execution. Commercial and service-related factors are integrated into each layer, with every potential action a telco might take analyzed through the lens of customer lifetime value, a common metric for evaluating which action will generate the most value over the long term.

### The data layer

The data layer should be the single source of truth for the entire engine, encompassing all data points that may be relevant for addressing the full spectrum of customers' needs and desires. Commercial and service data, including data from external sources,<sup>1</sup> gets combined with data on demographics and location to create a consistent 360-degree view of each customer.

Commercial data might include tenure, historical receptivity to campaigns, customer satisfaction scores, and previous transactions. Service data might include network usage, network stability, device usage, and customers' propensity to call regarding faulty service.

While telcos understand that data is critical, very few have developed a sophisticated data layer that stretches across the entire organization and covers all use cases. At most telcos, service and commercial divisions create and collect data separately. Service data is generally far less robust, in part because it is often generated by sensors and therefore requires a great deal of effort to extract,

transform, and use effectively. Moreover, silos exist even within telcos' commercial divisions; churn-focused teams, for example, may compile and process data separately from teams focused on customer acquisition. This can duplicate work while limiting each team's insight into customers' attributes, demands, and behaviors.

Poor data governance also affects telcos' ability to make the most of the data they possess. Telcos tend to lack clarity on data quality and where different types of data reside, and they often resort to less productive, manual (or semimanual) means of preparing data for further analytical work.

To create actionable, data-driven insights, automation is key. This requires clear data accountability standards and consistent naming conventions to ensure that data from multiple sources is labeled and organized similarly. With these measures in place, it's possible to automate and streamline data-related work, accelerating the machine learning development process. Telcos can then create reliable data engineering pipelines that automate both underlying data manipulation routines and the process of creating features used in further modeling and decision making.

Operators have seen success in ensuring that their data lakes are rich reservoirs, rather than murky swamps, by appointing a chief data officer or head of data: a senior leader charged with overseeing all customer data, protecting it, and guaranteeing that its quality is as high as possible. This leader should be the go-to person for data-related questions from anywhere in the organization. They should have the authority to appoint a team to help them determine where data resides, evaluate its quality, understand how it's structured and consumed, and monitor its use.

---

<sup>1</sup> Subject to both internal data governance and architectural capabilities, as well as local regulations, legislation, and external data availability.

One European telecom group's experience illustrates the value that strategic data governance and data engineering can unlock. The group owns major telcos in several countries, with variable data and analytics maturity levels. To codify and share best practices, build economies of scale, and standardize and expedite service delivery, the group decided to create a single data and analytics center of excellence.

The group started by creating a common data model to serve as the foundation for analytical efforts across all countries. In nine weeks, it created a common abstraction layer linked to raw data in each country via standardized data manipulation and preparation routines. Analytical models were designed to work on top of this abstraction layer, making it easy to shift any model or use case from one country to another. The group expedited the timeline for deploying new “lighthouse” use cases from around five months to six weeks, with only two additional weeks required to deploy new use cases across multiple countries.

#### **The analytics layer**

The analytics layer can include a wide range of machine learning models for all potential use cases, including cross-selling, churn prevention, upselling, and [three categories of service-related issues](#): preemptive (in which telcos detect an impending problem before it occurs), proactive (in which telcos detect and correct faults before customers notice them), and reactive (in which telcos engage with customers who contact them about an issue).

The central aim of this layer is to intervene as early as possible, ideally by preventing errors altogether. This minimizes the cost of rectifying issues and boosts customers' perception of service quality.

Many telcos have machine learning models for commercial use cases such as churn and pricing. But very few have models for use cases such as

broadband connection faults, bill shocks, and repeated failures. As a result, these issues are generally handled reactively instead of proactively. Even fewer telcos have integrated machine learning models for both commercial and service use cases. This is a missed opportunity, as it robs telcos of the ability to factor service-related issues into their decision making.

A Western European telco developed a predictive analytics model to identify customers who were most likely to call regarding fault-related issues and found that about 10 percent of customers were almost 80 percent more likely than others to call about faults. The telco then conducted a second, overlapping analysis to pinpoint the most easily resolvable faults—those that could be fixed remotely or by giving customers simple instructions on how to restart their router or change their network settings. At that point, the telco began a series of pilots, focusing on the 20 percent of customers most likely to call about quickly resolvable issues. In early pilots, customers received phone calls offering information about service outages and help fixing problems. Drawing on lessons from these early efforts, subsequent pilots offered end-to-end automated solutions, such as overnight remote router resets.

The telco reduced the costs associated with serving this group of customers by 35 percent, significantly boosting customer experience in the process.

#### **The decisioning layer**

While machine learning models determine the likelihood that certain events will occur or that customers will react in a certain way, the decisioning layer determines which use case to activate, when to activate it, and how to do so for each customer, depending on the situation or issue that arises. At this stage, leading telcos gain their competitive edge by delivering true personalization at scale.

Each potential action (both service and commercial) should be assessed for how it will affect customer lifetime value. The decisioning layer factors in the expected long-term value from every possible action a telco might take, the likelihood of an individual customer taking certain actions (such as accepting a new offer or calling to complain about service), and the cost of serving the customer through each channel. Whenever a customer's status changes—for example, when they migrate to a new plan, activate a new device, or start using a different channel—the decisioning model incorporates this new information when predicting customers' responses and calculating the value and costs of each potential telco action.

Ultimately, the decisioning layer activates the use case that will generate the most value. If the model signals, for example, that a customer will likely leave the company, the churn-prevention use case may kick in. If a customer is deemed likely to place an inbound call, the proactive-service-support use case may be activated. The decisioning layer also covers telcos' contact policies, guarding against campaign oversaturation and factoring in the chances that customers who are included in a particular campaign at a particular time will opt out of future communications.

To enable this level of responsiveness and personalization, customers are grouped into hundreds of microsegments for each potential use case, both commercial and service related. Microsegments are dynamic, responsive, and automatically updated, constantly evolving based on new data, insights, and customer behaviors. Historically, telcos' commercial operations have grouped customers into ten to 15 relatively static segments, based on data related to such broad categories as demographics, usage, and tenure.

Service operations have taken an even more rudimentary approach to segmentation, typically classifying customers into a few large groups according to the severity and frequency of the problems they experience.

By combining these two operations, telcos position themselves to respond to customers in a way that aligns with their most urgent priorities and concerns. If a high-value customer experiencing service problems, for example, were to be contacted proactively with an offer to send an engineer to their home, all marketing campaigns would be paused until the customer's service issues were resolved.

The decisioning layer ideally includes three parallel flows: an always-on flow, which is the default; a regulatory flow, which is for communications that are legally required to be sent out regardless of commercial considerations; and a trigger-based flow, which is activated when circumstances change, necessitating a quick diversion from the machine learning model.

The regulatory flow is straightforward. It supersedes the other two flows, ensuring that telcos comply with all legal mandates (alerting customers two months before their contracts expire, for example). The trigger-based flow is activated when a customer's situation suddenly changes, calling for a different action than the always-on flow would have suggested. If a customer, for instance, explores a section of the telco's website with information about canceling contracts, the trigger-based flow might kick in to send offers enticing them to stay. If a customer explores information about iPhone upgrades, they may receive a text with upgrade offers. If they conduct a remote speed check on their broadband connection, they may receive an in-app message suggesting they reconnect their router.

The trigger-based flow should include at most a dozen triggers to monitor and act on. This helps to reduce any risk of compromising the accuracy or efficiency of the always-on flow—which is continually integrating data on customer behavior to determine the optimal cadence and content for marketing campaigns. While creating a single always-on flow that responds to triggers in real time is possible, it is a complex and expensive endeavor; establishing parallel decisioning systems is an efficient work-around.

Occasionally, business priorities or brand-new information that the model has not yet integrated may prompt leaders to manually override the always-on flow. If a competitor enters a new region, for example, the business may prioritize broadband repairs or offers in that specific region rather than the actions suggested by the always-on flow.

There may be a temptation to overuse the manual override feature, particularly during the transitional period when KPIs are still not fully aligned across departments or wholly focused on overall customer experience. A telco's broadband team could potentially advocate for sending out a campaign designed to boost broadband sales—regardless of its impact on long-term value. Telcos would be well advised to proceed with caution when exercising the manual override option, as its usage erodes the quality of the overall decisioning process.

#### **The channel execution layer**

The full breadth of telcos' automated and human-based channels drives the execution engine for personalization at scale. Inbound call centers, retail stores, websites, apps, SMS outreach, and email channels should ideally speak with one voice, so customers can switch channels seamlessly without being inundated with contradictory messaging. To ensure that telcos are sending the right message at the right time, on the most appropriate channel, all

channels should be able to seamlessly integrate recommendations emanating from the decisioning layer and send back new, outcome-related data.

Very few telcos have achieved true omnichannel capabilities, in which all inbound, outbound, automated, and human-assisted channels are completely aligned. And no more than a handful have integrated commercial and service-related interactions across channels.

When this integration takes place, customers can schedule engineering visits through their preferred channels and track their technicians' locations in real time, making visits much more convenient. Telcos can also gauge differences in how customers respond to interactions on various channels and at numerous times of day. When customers interact with the telco on any channel, the outcomes of these interactions are fed back into the decisioning layer to continually improve the model's decisioning logic.

A major integrated European telco implemented a comprehensive channel strategy to overhaul its service and repairs experience, which had long been a source of customer dissatisfaction and a driver of churn. The telco digitized the experience, creating an app and web-based platform that allow customers to schedule engineering visits easily, monitor how far a technician is from their home, see how many customers are ahead of them, and communicate with engineers before and after visits.

This may seem simple, and from the customer's perspective, it made the often frustrating process of fixing connectivity issues as easy as using a ridesharing or food-delivery app. Behind the scenes, though, the process involved a complex set of interactions among different systems. The effort was worth it, halving the telco's inbound calls related to service ticket management and reducing



the number of unsuccessful visits (those failing to yield a satisfactory resolution by 20 percent). As important, the telco converted what we call “boo moments” into “wow moments,” redefining its customers' expectations.

### **Enhancing the layers**

To unlock the most value from the next-best-experience engine, all four layers should be continually improved through a feedback loop that incorporates automated resolution tracking to ensure problems have been resolved. As new data becomes available, based on customers' responses to recent campaigns or the relative success of communications on different channels, it should be incorporated as quickly as possible. This will ensure, for example, that customers who reject a particular offer do not receive the same one a week later. As machine learning models are continually retrained using the best, most recent data available, accuracy and effectiveness should improve.

Once the four layers are in place, along with processes for continually refining them, organizations can establish an entirely new set of end-to-end processes that use the engine's insights to unlock value. This might involve creating hyperpersonalized marketing campaigns by combining inputs from the decisioning system with generative AI capabilities, or redesigning products and offerings based on analytical insights.

### **Starting the journey**

True personalization at scale requires a fundamental shift—away from disconnected decision making based on business units' individual KPIs and toward cohesive decision making based on what's best for the customer at a given point in time. Succeeding with that scale of transformation requires critical enablers to be in place, at both infrastructural and organizational levels.

At the infrastructural level, telcos should strongly consider having a modern tech stack in place, one fortified with best-in-class computational power, data engineering and software development tools, machine learning operations, DevOps (software development and IT operations), and real-time delivery capabilities. Also, as mentioned earlier, data governance and accountability, spanning all four layers, are essential; a single team or leader has to serve as the point of contact for all data-related questions and issues.

At an organizational level, people across dozens of commercial and service teams will need to work together on this effort, requiring complex workflows that stretch across the entire organization and new levels of collaboration among colleagues with diverse profiles, backgrounds, and experiences. Traditional “waterfall” ways of working won't be sufficient for the rapid pace of this kind of work; developing and managing models based on real-time data requires agile teams that can embrace a “fail fast” mentality and are empowered to make decentralized decisions. Finally, telcos will need to seek out and develop a new class of technical talent: workers skilled in scaling up machine learning—powered use cases, integrating technical and commercial considerations into decision making, and establishing effective pilots with solid feedback loops.

Organizations might start the journey with a diagnostic assessment to evaluate the maturity of their enablers and identify gaps that need to be filled immediately. The diagnostic can also help leaders align on the most appropriate deployment model. If certain enablers are particularly weak, starting with less complex models (using rule-based logic, for example) is preferable to waiting to develop the talent or tech stack required to build best-in-class models (which use neural networks).



Find more content like this on the  
**McKinsey Insights App**



Scan • Download • Personalize



To determine which use cases to focus on first, telcos might embark on a scoping exercise to evaluate each use case based on the difficulty of implementation and its potential to generate value for the business. Telcos can begin with a set of lighthouse use cases that are relatively easy to build and demonstrate the value of a robust, integrated next-best-experience engine. Lighthouse use cases are an effective way to persuade business leaders that creating new ways of working, measuring success, and allocating investment dollars will be worth the effort.

## Looking ahead

Over the next five years, customers will likely expect personalization at scale in all interactions with their telcos—including service, a critical aspect of any customer's experience. During that same period, the scope of those interactions could expand greatly.

As 5G (fifth generation) technology matures, telcos have an opportunity to maximize their share of a value pool that [could easily exceed \\$100 billion over the next five years](#) by venturing further beyond core connectivity. In the coming years, we are likely to see telcos create a range of new revenue

streams, from credit risk assessment products geared toward financial institutions to ecosystem offerings that bundle media and entertainment services. Personalization at scale, with a baked-in service component, has the potential to accelerate telcos' efforts to capture value in new ways.

Telcos can also derive new value from their new data and analytics capabilities to streamline the scale-up of use cases and move into adjacent businesses such as home security or financial services. They can even spin off separate companies or business units that sell these capabilities as services to telcos and other customer-facing businesses, tapping into new value pools.

The foundation of all this potential change is the next-best-experience engine that integrates commercial and service-related capabilities. Telcos that act early to create this powerful tool may have a distinct advantage, with the extended runway needed to experiment, iterate, refine, and monetize their efforts. In the process, they'll help establish the elevated customer expectations that all telcos will eventually have to satisfy—including those that may regret their decision not to act sooner.

[Lars Engel Nielsen](#) is a partner in McKinsey's Copenhagen office, [Lars Fiedler](#) is a partner in the Hamburg office, [Dominika Kampa](#) and [Stanislav Kucherenko](#) are associate partners in the London office, [Julian Raabe](#) is a partner in the Munich office, [Tim Schenk](#) is an associate partner in the Melbourne office, and [Davide Schiavotto](#) is a senior partner in the Milan office.

The authors wish to thank Anshul Goyal, Kalina Janevska, and Sokto Sultimov for their contributions to this article.

Copyright © 2023 McKinsey & Company. All rights reserved.