Publication date:
22 Mar 2021
Author:
Adaora Okeleke
Principal Analyst, Telco IT & Operations

Enterprise Case Study: T-Mobile US Turns to Machine Learning to Enhance Customer Service

Brought to you by Informa Tech

Table of Contents :

Table of Figures :

Summary

Catalyst

The quality of interactions on the channels that customers choose for engaging with their communications service provider (CSP) to troubleshoot issues plays a critical role in determining the quality of their overall experience. Omdia's research shows that 42% of customers today prefer to use voice calls to engage with their service providers to resolve issues. Because of the large volume of requests that customer service representatives (CSRs) need to address, customers are deflected to self-service channels that don't fully resolve their issues. When in contact with a human agent, the average handling time (AHT) for calls is long as the CSRs are not fully equipped to resolve the issue right the first time.

Consequently, CSPs are deploying machine learning (ML) to enhance CSR operations. However, implementing ML-based solutions comes with several challenges, including the challenge of data labeling. Data labeling generates the training data for the model training phase of ML workflows. However, large training data sets are required to ensure the data models and results generated are accurate. Consequently, it is an important but costly and time-consuming process when executed manually. It is therefore important that CSPs identify efficient ways to perform data labeling to drive efficient ML workflows. This case study highlights how T-Mobile US implemented a ML-based data labeling solution from Amazon Web Service (AWS) to accelerate its CSR team's capabilities to address customer queries and get it right the first time.

Omdia view

Getting access to a customer agent and having them respond promptly to customer enquiries is critical to improving the customer journey. To achieve this, T-Mobile US transformed its customer service model, ensuring that, when needed, customers get immediate contact with an agent that understands their requirements and can serve them effectively.

T-Mobile's AI@T-Mobile team developed natural language understanding ML models that analyze large volumes of customer requests and data to predict why the customer has contacted T-Mobile and identify the information that would address the customer's needs. The data labeling process was inefficient, taking time and increasing costs. This led to the CSP using Amazon SageMaker Ground Truth to automate the data labeling process. Ground Truth is a fully managed data labeling service within Amazon SageMaker that uses ML to simplify the data labeling process.

The team saw time spent on labeling data decrease by a factor of seven, with the quality and volume of training data sets also increasing. Metrics such as customer calls per account reduced as a result of the improvements in the data labeling and model training processes. By deciding to run its customer service operations differently, maximizing the expertise of its data science teams to achieve this objective and moving non-differentiating tasks to the public cloud, T-Mobile US is able to achieve its vision of delivering enhanced customer experiences.

Key messages

- Providing customers with direct access to customer care agents is a critical element of the customer journey.

- T-Mobile US leverages ML to ensure customers connect with care agents that understand their needs and are equipped to address them right the first time.
- Automating the data labeling process is an important stage of ML workflows as it enhances the efficiency, quality, and accuracy of results obtained.

Recommendations for the telecoms industry

Recommendations for communications service providers

Leverage ML to facilitate customer care agent interactions

Use ML to provide self-service capabilities and to ensure customers can connect directly to humans if required. While tech-savvy customers are happy to address issues on their own, there are customers that require human assistance. Use ML to determine what the customers need, and with this insight connect them directly to CSRs that understand the context and have the right tools and information to address their needs quickly. This will save time and improve experiences delivered to customers.

Improving the care agent's experience is as important as improving the customer's

Reduce AHT and increase first-call resolution (FCR) by ensuring CSRs have immediate access to the right information to resolve customers' enquiries quickly. ML can enable this capability by analyzing customer requests, related customer information, and knowledge repositories to detect customer needs and to identify and provide care agents with relevant information to sort out the requests. Quality of experience and productivity levels of customer care agents will be enhanced, and agents can better serve customers. Attrition rates will also reduce, resulting in cost savings in recruiting and training new staff.

Leverage the public cloud to scale your ML workflows

Cloud service providers have invested heavily in providing enterprises with ML capabilities to embed intelligence within applications and processes. The cloud-native architecture of these solutions makes it relatively easy to add these capabilities to preexisting ML solutions. Where possible, leverage these tools to scale up ML processes.

Recommendations for vendors

Optimize the data labeling process

ML models are as good as the data used to train them. With more CSPs adopting ML to improve customer operations, vendors should simplify and optimize the data labeling process to ensure ML models are accurate and drive expected outcomes. Automating the data labeling process can improve the quality of training data sets and therefore enable ML models to drive accurate results.

Be flexible in accommodating the ML needs of customers

Take time to understand your CSP customers' competences and needs regarding ML, and craft solutions that align with them. CSPs have varying degrees of competences regarding ML, with some having in-house data scientists with the skills to develop proprietary solutions. However, to maximize these in-house skills, tasks such as data labeling are better outsourced. Solutions therefore will need to vary from fully packaged solutions providing all the required capabilities to support a fully developed ML-based solution to those that can address specific challenges within the ML workflow. Consequently, be flexible in accommodating the varying needs of CSP customers.

Using ML to improve customer experience

Setting the business context

Providing the human touch in customer service

Getting access to a customer agent and having them respond promptly to customer enquiries is a key issue plaguing today's CSP customer journey. A common solution is to deflect customers to chatbots, interactive voice response (IVR), and other self-service or automated tools to address customers' issues. While tech-savvy customers are happy to sort out problems on their own, most customers prefer to interact with humans to solve their problems. Omdia's Digital Consumer Insights 2020 survey shows that 42% of customers prefer to use voice call to address troubleshooting and technical support enquiries. With this proportion of customers using voice channels, CSPs need to have CSRs available to address customer needs.

1. Figure 1: Customers prefer using voice calls to engage with service providers

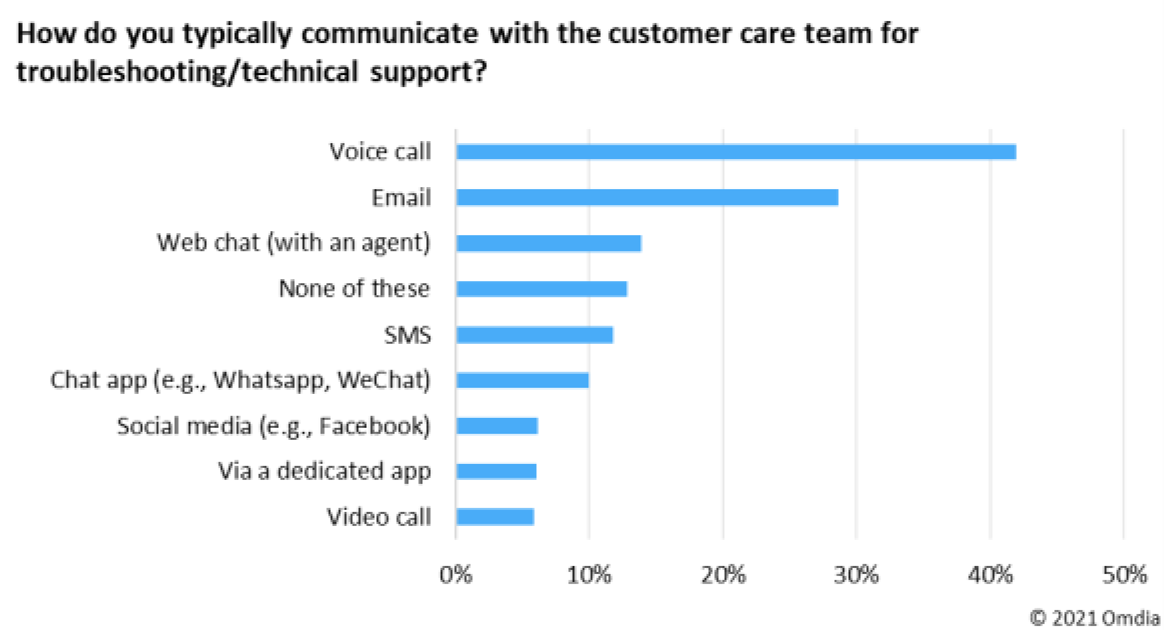

Source: Omdia's Digital Consumer Insights 2020 survey

For CSPs that aim to provide the human touch, getting CSRs to access the right information quickly is a daunting task. Why? CSRs must access customer data that is stored inconsistently across disparate systems. CSPs therefore need to optimize how CSRs access information to enhance customer interactions. Artificial intelligence (AI) presents an opportunity to address this challenge. By detecting what a customer needs and learning how to resolve it based on information in knowledge repositories, AI can help customer service agents be more proactive and efficient.

T-Mobile US improves customer service operations model

T-Mobile US is the third-largest US operator in terms of number of subscriptions, with over 140 million subscriptions at the end of 2020, up from 86 million subscriptions at end of 2019 (according to Omdia's World Cellular Information Service). The CSP recorded $68.4 billion in total revenue, growing by 52% year on year at the end of 2020. The large increase is due to T-Mobile's acquisition of Sprint in 2Q20. Prior to this acquisition, in 2013, T-Mobile US commenced its Un-carrier initiative with a view to prioritizing customers' needs to deliver better services and experiences.

In 2018, as part of its Un-carrier initiative, T-Mobile's then CEO, John Legere, announced a new approach to customer service, the Team of Experts (TEX). The TEX strategy was developed to provide customers with their own dedicated "team" of CSRs (with cross-functional roles) to offer personalized and efficient assistance. Each TEX team serves a named set of customer accounts and customers can reach an account rep in a team directly through a variety of channels: T-Mobile's messaging app, the customer portal on T-Mobile's website, and phone.

A key objective of the TEX teams is to allow agents to support multiple lines of business, reduce transfer rates, and improve metric, such as increasing the rate of first call resolution (FCR) and reducing average handling time (AHT). To achieve these objectives, agents needed to be trained and given permission to support multiple contact types (e.g., text and voice contact types) and services, addressing everything from billing, sales, and line activation to standard technical support enquiries. However, given the scale in terms of number of customers, it was difficult to achieve the set objectives because of the following:

- Large volumes of customer requests and data had to be processed daily.
- Customer messages were extremely varied and it was difficult to determine the customers' needs and match them to the right response.
- Customer information is held in several systems which agents need to query to address customer enquiries.

T-Mobile's TEX teams therefore required the right tools to automate CSRs' workflows and guarantee that customers' issues and needs were resolved quickly and accurately.

Providing customer care agents with contextual customer information in real time

T-Mobile US set up an intelligence-driven strategy focused on equipping customer care agents with contextual information in real time (target time within 250 msecs compared to the 45 secs it took an agent to respond to a message) to address customer enquiries right the first time. To achieve this, T-Mobile US had to find an efficient way to process data from large volumes of customer requests, related customer information, and knowledge repositories. This is essential for understanding why a customer is reaching out to T-Mobile and how to resolve their requests quickly. The TEX team worked with T-Mobile's AI@T-Mobile team to develop a solution that supports this strategy.

The role of ML in improving customer service operations

Existing landscape and solution selection

The AI@T-Mobile team developed natural language understanding ML models that analyze data from thousands of incoming customer requests daily and extract meaning from vast amounts of data coming from customers' text messages. It also analyzes knowledge repositories where potential answers to customer queries can be found and matched to address customer requests.

Based on this data, the ML models would then predict why the customer is contacting T-Mobile and identify the information that would address the customer's needs, such as helping them pay their bill or adding a new phone line. The relevant content will be presented to the CSR in the TEX team assigned to address that customer's needs. Figure 2 provides a summary of how the customer message and associated data is analyzed.

2. Figure 2: T-Mobile US using CNN to process customer text messages

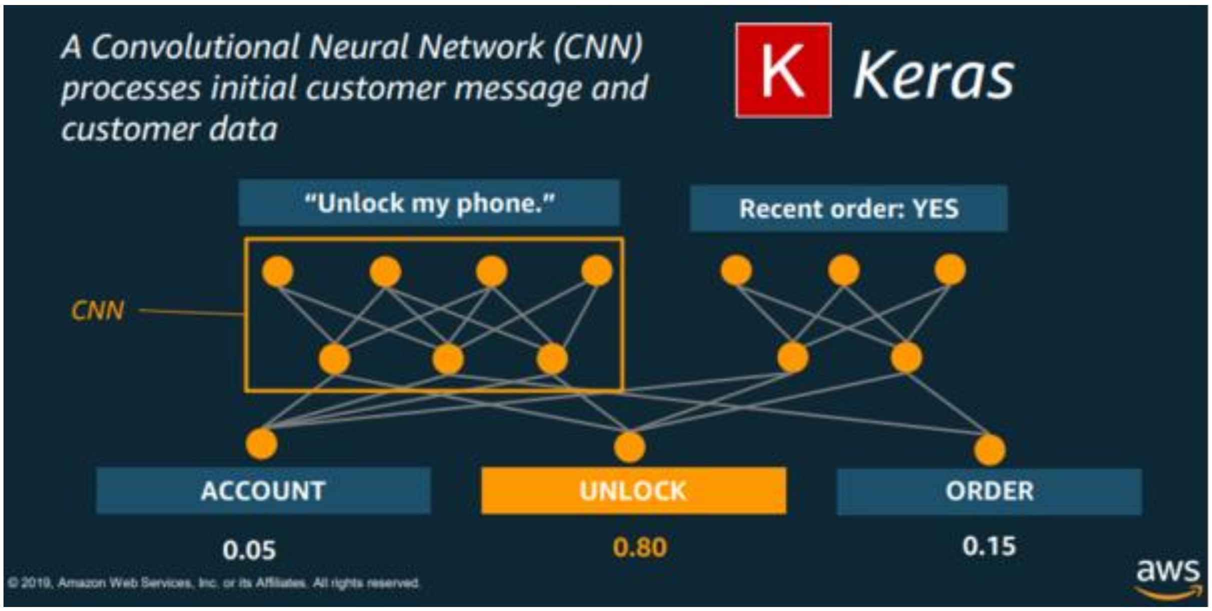

Source: AWS and T-Mobile US

A convolutional neural network (CNN) processes the initial customer message and related customer data to decide why the customer is calling (otherwise known as the topic). For example, if the customer says "unlock my phone," the CNN detects that the topic is "unlock." The algorithm also detects that the customer has just purchased a new phone (using purchase history) and so displays all related customer information to the agent to help unlock their phone.

The CNN required supervised learning and a large volume of labeled data to train it. Several learning techniques were explored, including unsupervised learning techniques such as clustering and latent Dirichlet allocation (LDA) analysis. However, the trained model and results generated were inaccurate. For example, 90% of text classifications were tagged T-Mobile. Rule-based supervised learning was also applied but did not result in a high-fidelity model as the rigid rules failed to capture the nuanced use of some keywords in customer messages.

After several trials, the team decided to use supervised learning which required between 500,000 and 1 million labeled conversations to build a robust neural network. The team commenced labeling the training data manually using excel spreadsheets. It required many hours for highly paid data scientists to label customer messages. Given how inefficient and expensive the data labeling process was, the AI@T-Mobile team turned to Amazon SageMaker Ground Truth (referred to as Ground Truth).

Ground Truth is a fully managed data labeling service within Amazon SageMaker that simplifies the data labeling process. It also leverages a ML model to automatically label data, thereby increasing the volume and quality of training data sets. The ML model learns from the labeling of simple messages and based on these learnings starts labeling similar messages. For example, the ML model will recognize that the conversation "I have a problem with my bill" was labeled "billing" and will start labeling similar conversations automatically, augmenting the functions of the human workforce.

3. Figure 3: How AWS Ground Truth operates

Source: AWS

T-Mobile US considered several criteria for its data labeling solution including being

- highly secure, limiting the need to retrain models within external cloud environment using customer data
- customizable, enabling the team to choose where they wanted to perform specific tasks and not be restricted to performing them within a specific platform
- able to deploy with a modern software stack in Kubernetes containers
- able to auto-scale as Kafka consumers and producers

T-Mobile US engaged with several vendors including Microsoft Azure (to leverage its Azure Databricks service), Google, and other smaller players.

T-Mobile US decided to go with AWS as they presented a solution that specifically addressed its data labeling challenge and could do so in an efficient way. AWS's solution was also flexible; for example, it allowed the team to export models and deploy them in their own systems. Furthermore, T-Mobile had a native Kubernetes environment that ran on AWS. T-Mobile's ML pipeline runs within the Kubernetes environment, executing data ingest and data engineering functions within AWS. Consequently, working with AWS resolved the integration challenge as the datasets to be labeled could be easily transferred to Amazon's cloud database, Amazon S3, and then to Ground Truth for labeling.

Bringing the strategy to life

Implementation of Ground Truth with T-Mobile's ML pipeline started with a proof of concept (POC) in August 2017, and by March 2018 the first solution was deployed to production. **Figure 4** provides a summary of the architecture driving the implementation of Ground Truth's data labeling service.

4. Figure 4: Architecture showing T-Mobile's implementation of Ground Truth

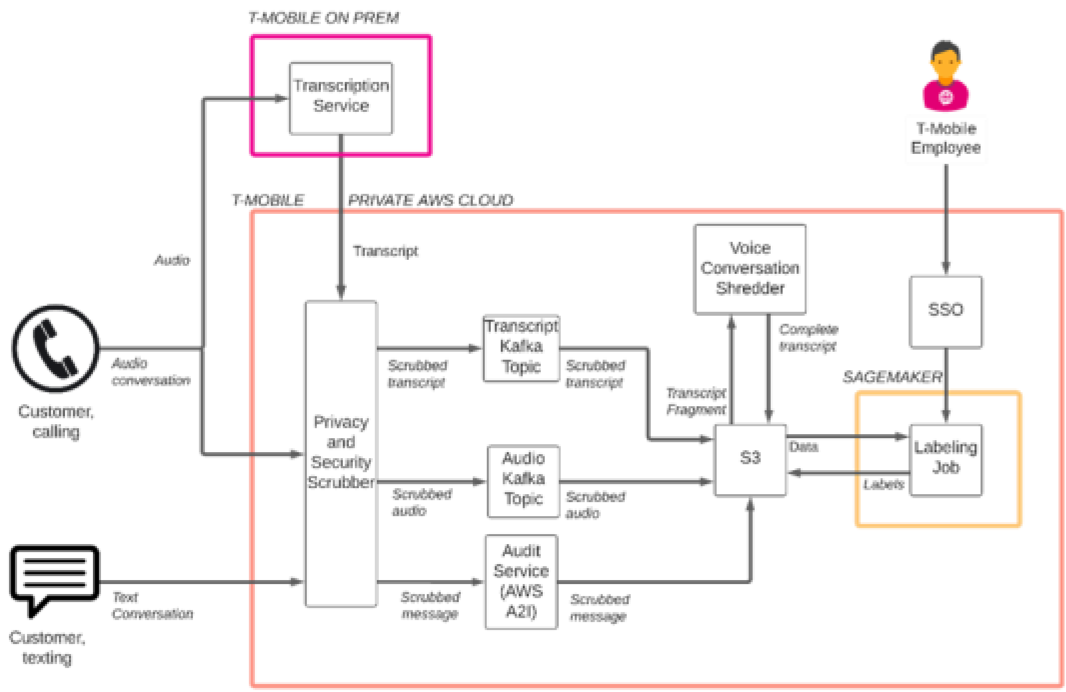

Source: T-Mobile US

When customers contact T-Mobile, their conversations are passed via Apache Kafka through services including scrubbing and audit (implemented using Amazon Augmented AI [A2I]) to be transformed to the correct format before they are stored in Amazon S3. This process differs based on the type of conversation. A Ground Truth labeling job for a voice conversation involves listening to the audio output from a call, correcting the transcription in one-minute intervals and then classifying the transcribed data. For a messaging conversation, it involves taking the text of a message and classifying it into intents or topics immediately.

With the accelerated pace in labeling data, T-Mobile was able to implement a multi-tiered labeling strategy. Once a conversation is labeled the first time, it would go through two other stages of labeling to ensure higher accuracy of the training data sets. Ground Truth helped to facilitate this strategy.

5. Figure 5: T-Mobile's multi-tiered data labeling strategy

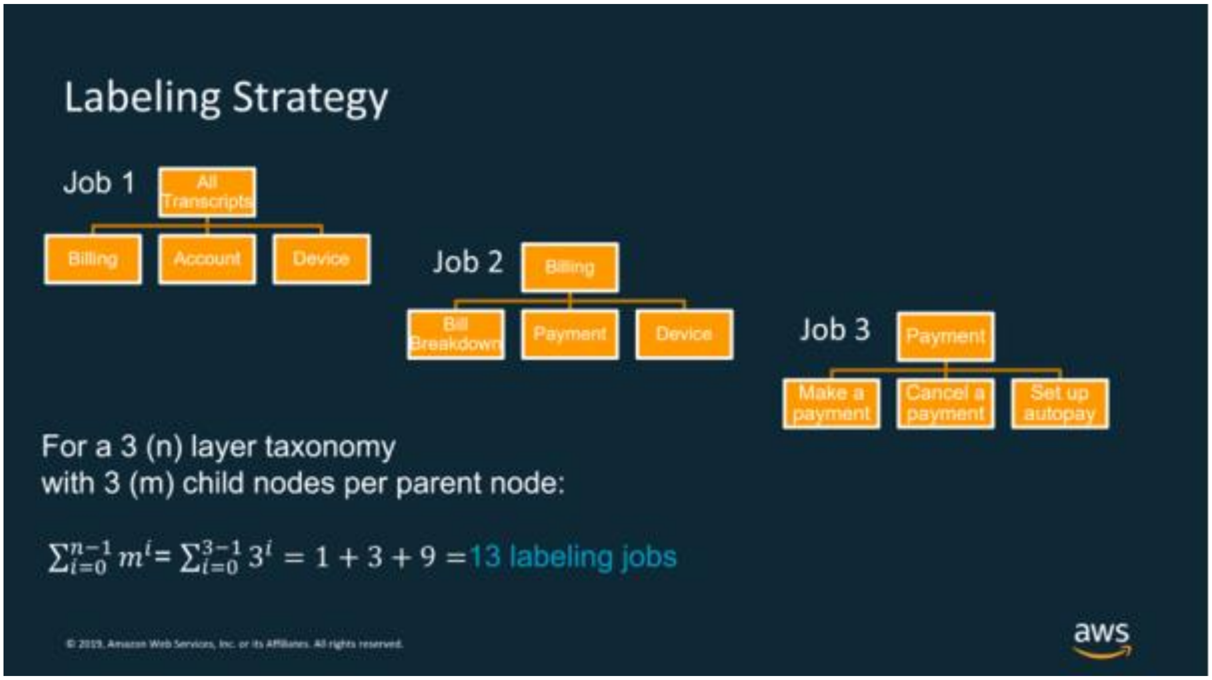

Source: AWS and T-Mobile US

Challenges encountered

T-Mobile encountered several challenges including the following:

- **Dealing with very large data sets in T-Mobile's ML pipeline**. T-Mobile's ML models required very large data sets, and using containers for data storage of these large data sets for ML isn't recommended. Usually, large datasets are stored in a storage system/service like Amazon S3 or RedShift or in a relational database like Oracle or MySQL prior to executing in a container. T-Mobile worked with AWS to design its AI/ML infrastructure with AWS for scale. The role of Docker containers in T-Mobile's DevOps strategy was to decouple the processing pipelines from the data processing layer. This allowed T-Mobile to scale both entities efficiently.
- **Addressing security concerns**. Based on T-Mobile's security policies, T-Mobile had to encrypt data before processing on Ground Truth. This challenge was resolved by creating an encrypted docker container to secure the labeling of the data sets.

Outcome assessment

Following the implementation of Amazon SageMaker Ground Truth, T-Mobile's AI@T-Mobile team saw the following benefits:

- **Time spent on labeling data decreased by nearly a factor of seven**. For example, where previously a human could complete 50 labels in one hour, they were able to complete 360 in the same period using Ground Truth.
- **Quality and volume of training data sets increased**. With Ground Truth supporting the multi-tiered labeling service, the team was able to improve the quality of training data sets as they could detect mislabeled data sets and label them again in future labeling jobs. T-Mobile US was able to generate more training data sets in less time with higher efficiency and accuracy.

In terms of operations of the TEX teams, the following improvement was noted:

- **Calls per account improved.** The number of calls per account fell by 35%. As a result of the improvements in the data labeling and model training processes, more messages from customer conversations are processed quickly, and customer care agents receive relevant information in time to help address customer requests, resulting in reduced calls per account.

T-Mobile US is currently conducting POCs with AWS to leverage its Transcribe service to offload current transcription processes so that the CSP can focus on other differentiated tasks that are centered around improving customer experience.

Lessons learned

Challenge the status quo

Using automated response systems such as IVR and chatbots has helped CSPs deflect thousands of calls from care agents to self-service tools. However, not accessing a care agent to address complex issues hampers the customer's experience. T-Mobile US decided to challenge the status quo and provide its customers with direct access to agents when they need it. Taking this approach has allowed the CSP to improve key metrics such as its net promoter score (NPS). It is important that CSPs weigh the impact of solutions that improve productivity against their impact on overall customer experience.

Utilize AI to augment the workforce, not to replace them

By leveraging ML, T-Mobile was able to offload some of the mundane and manual tasks involved in ML workloads to ML-based systems, releasing more time to its data scientists and engineers to focus on more creative tasks that improve the customer experience. This outcome contrasts with the view that AI would take away jobs. T-Mobile has proven instead that AI can drive increased productivity.

Offload non-differential functions to the public cloud

T-Mobile is a CSP with a strong in-house development team. However, the CSP understood the need to focus the efforts of its development resources on facilitating its customer experience Un-carrier initiatives. While labeling data is an important and necessary function, this function can be offloaded to third-party systems or services running in the public cloud (in this case AWS) that can perform them efficiently and at a reduced cost. By taking this approach, T-Mobile is better placed to maximize the skills of its data scientists and engineers.

Appendix

Methodology

This Omdia case study leveraged in-depth interviews with key stakeholders from T-Mobile US and Amazon Web Services as well as a review of available documentation in the public domain.

Further reading

T-Mobile US Update, November 2020 (January 2021)

Omdia Market Radar: Customer Engagement Solution for CSPs, 2020–21 (July 2020)

Enterprise Case Study: Enabling Intelligent Decisioning at Scale for Better Marketing Outcomes (April 2020)

Author

Adaora Okeleke, Principal Analyst, Service Provider Operations, and IT

adaora.okeleke@omdia.com

Citation policy

Request external citation and usage of Omdia research and data via citations@omdia.com.

Omdia consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help you. For more information about Omdia's consulting capabilities, please contact us directly at consulting@omdia.com.

Copyright notice and disclaimer

CONTACT US

omdia.com

askananalyst@omdia.com