Guidance for Subscriber Churn Prediction and Retention on AWS

Use machine learning to predict subscriber churn and initiate retention actions.

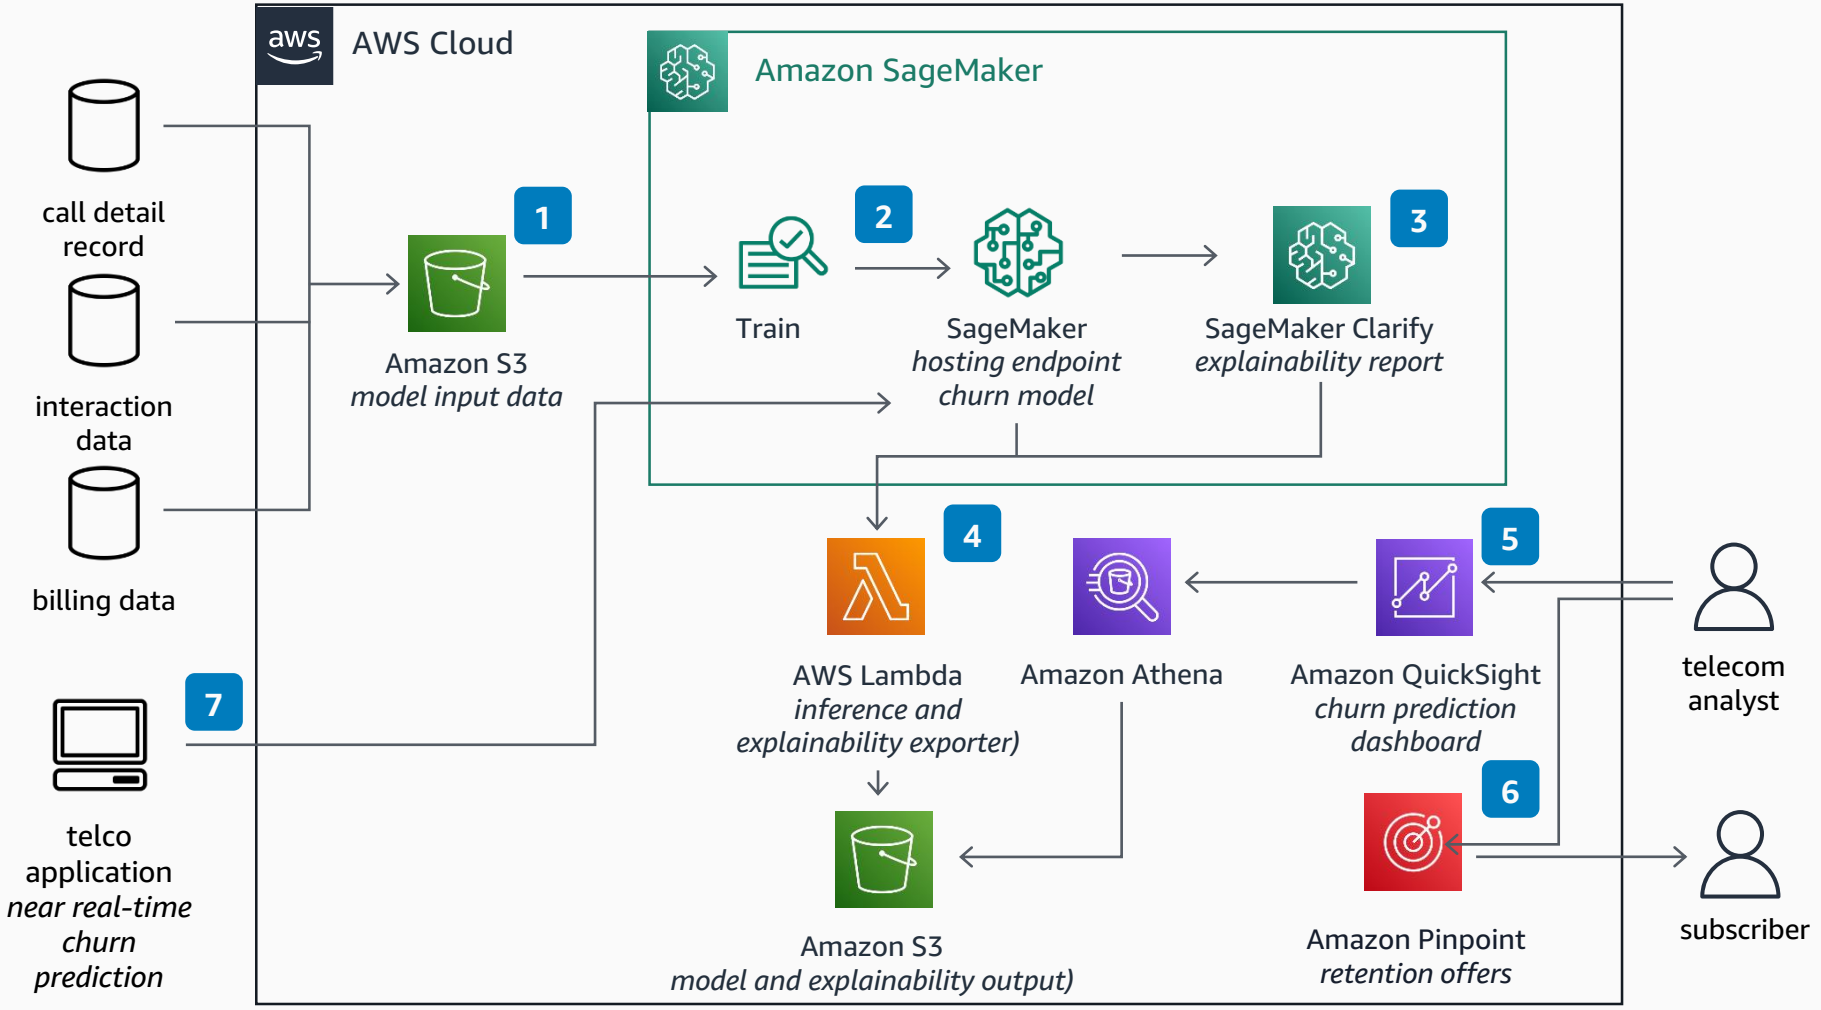

1. Telecom data is collected into an **Amazon Simple Storage Service** (Amazon S3) object storage bucket. Data includes call data records (CDRs), billing data, and data from customer care.
2. A churn model is trained on the labelled data set, tested and tuned, then deployed using **Amazon SageMaker.**
3. For every churn inference event, **Amazon SageMaker Clarify** identifies the important feature to the model to determine churn likelihood.
4. Churn model predictions and explainability reports are exported to an **Amazon S3** bucket by using an **AWS Lambda** function**.**
5. **Amazon Quicksight** visualizes the model and explainability data, allowing for interactive analysis and identification of trends and decision support of who to send a churn retention offer. **Amazon Athena** is used by **Amazon QuickSight** to access the **Amazon S3** data.
6. The telecom analyst can then decide how to act on the insight and can use **Amazon Pinpoint** to send out subscriber retention offers.
7. Telecom applications can also incorporate near real-time churn prediction by calling the **Amazon SageMaker** hosting endpoint.

Reviewed for technical accuracy October 19, 2022

AWS Reference Architecture