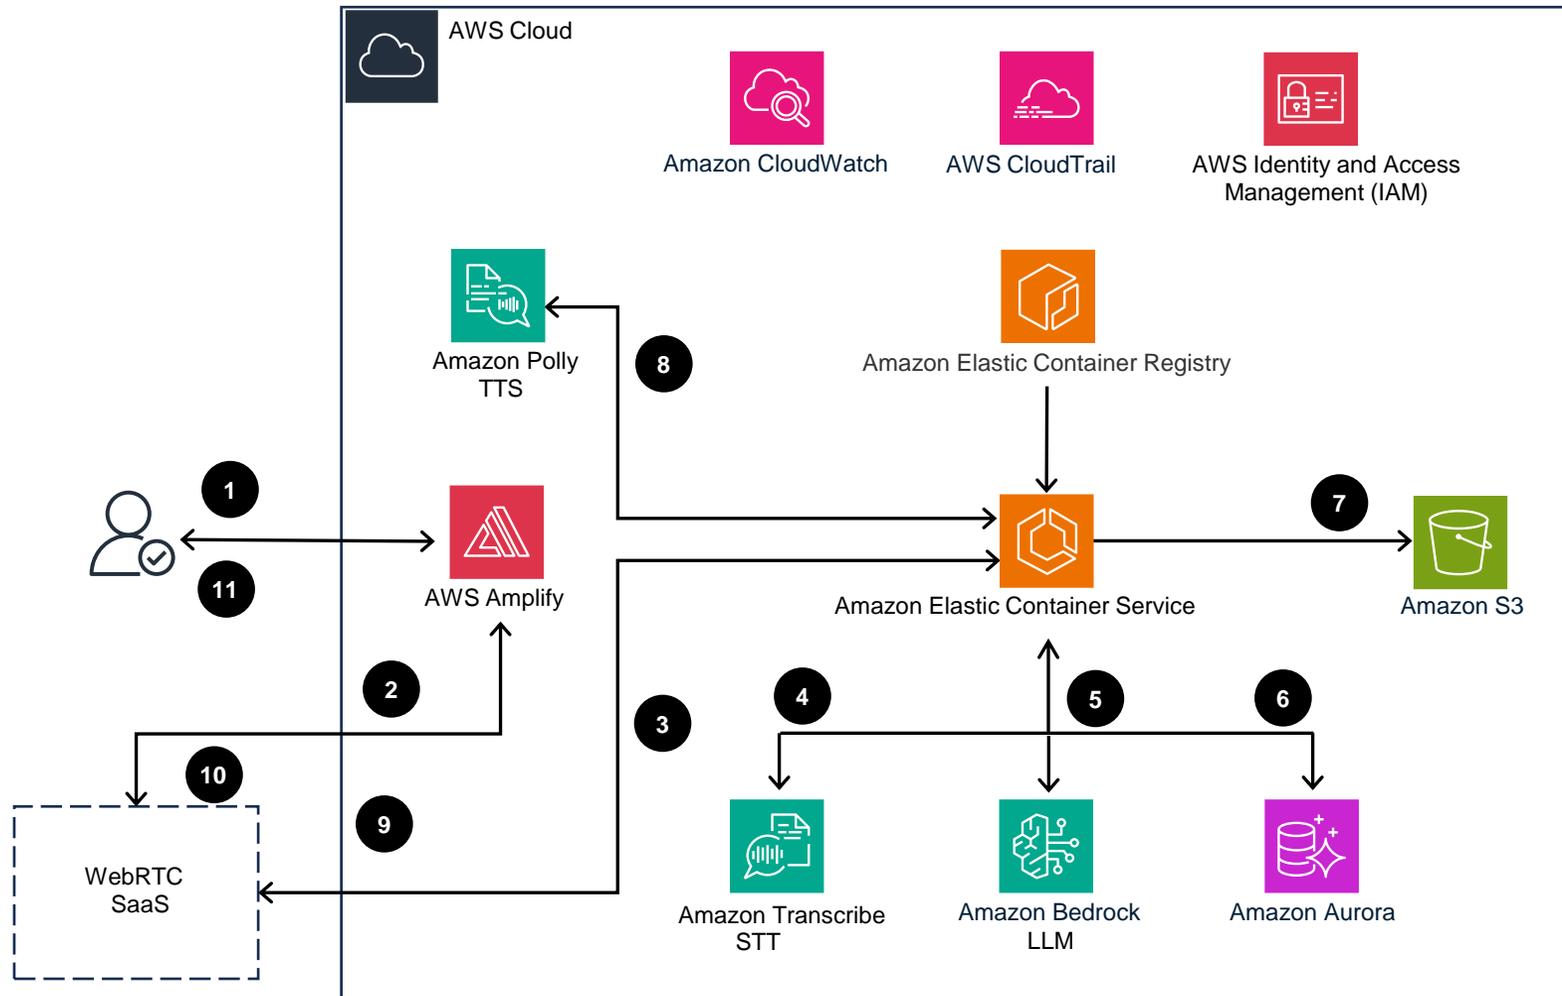


Guidance for Building a Virtual Car Showroom on AWS

This architecture diagram shows how to build a personalized and interactive Virtual Car Showroom using, Amazon Bedrock, Amazon Polly, and Amazon Transcribe, enabling customers to search car inventory, ask questions, and schedule test drives with natural language processing capabilities.



- 1 A car shopper starts by talking with their web or mobile app that is served via **AWS Amplify**.
- 2 The app connects to a WebRTC provider to start a user session.
- 3 The user session connects with an AI Voice Agent running as a Task in **Amazon Elastic Container Service (Amazon ECS)** based on a published container from **Amazon Elastic Container Registry (Amazon ECR)**.
- 4 The AI Voice Agent uses **Amazon Transcribe** to convert speech to text to be then processed, by a foundation model of your choice through Amazon Bedrock.
- 5 Your LLM from **Amazon Bedrock** responds to the prompt and utilises function calls to interact with external tools.
- 6 The function call initiates additional agents and LLMs to query **Amazon Aurora** and pull the car inventory via text-to-sql chaining.
- 7 The car images referenced in the data retrieved through RAG (Retrieval-Augmented Generation) in the previous steps are then pulled from **Amazon Simple Storage Service (Amazon S3)** and presented to the LLM to further enhance the context for describing the data set.
- 8 The text response from the LLM is then handed over to **Amazon Polly** to generate the response in natural language.
- 9 The generated voice and text are then transported back to the WebRTC layer.
- 10 The client app then relays the audio and visual content in the web or mobile app.
- 11 The user continues further interaction and Q&A via voice and visual displays as inputs.

