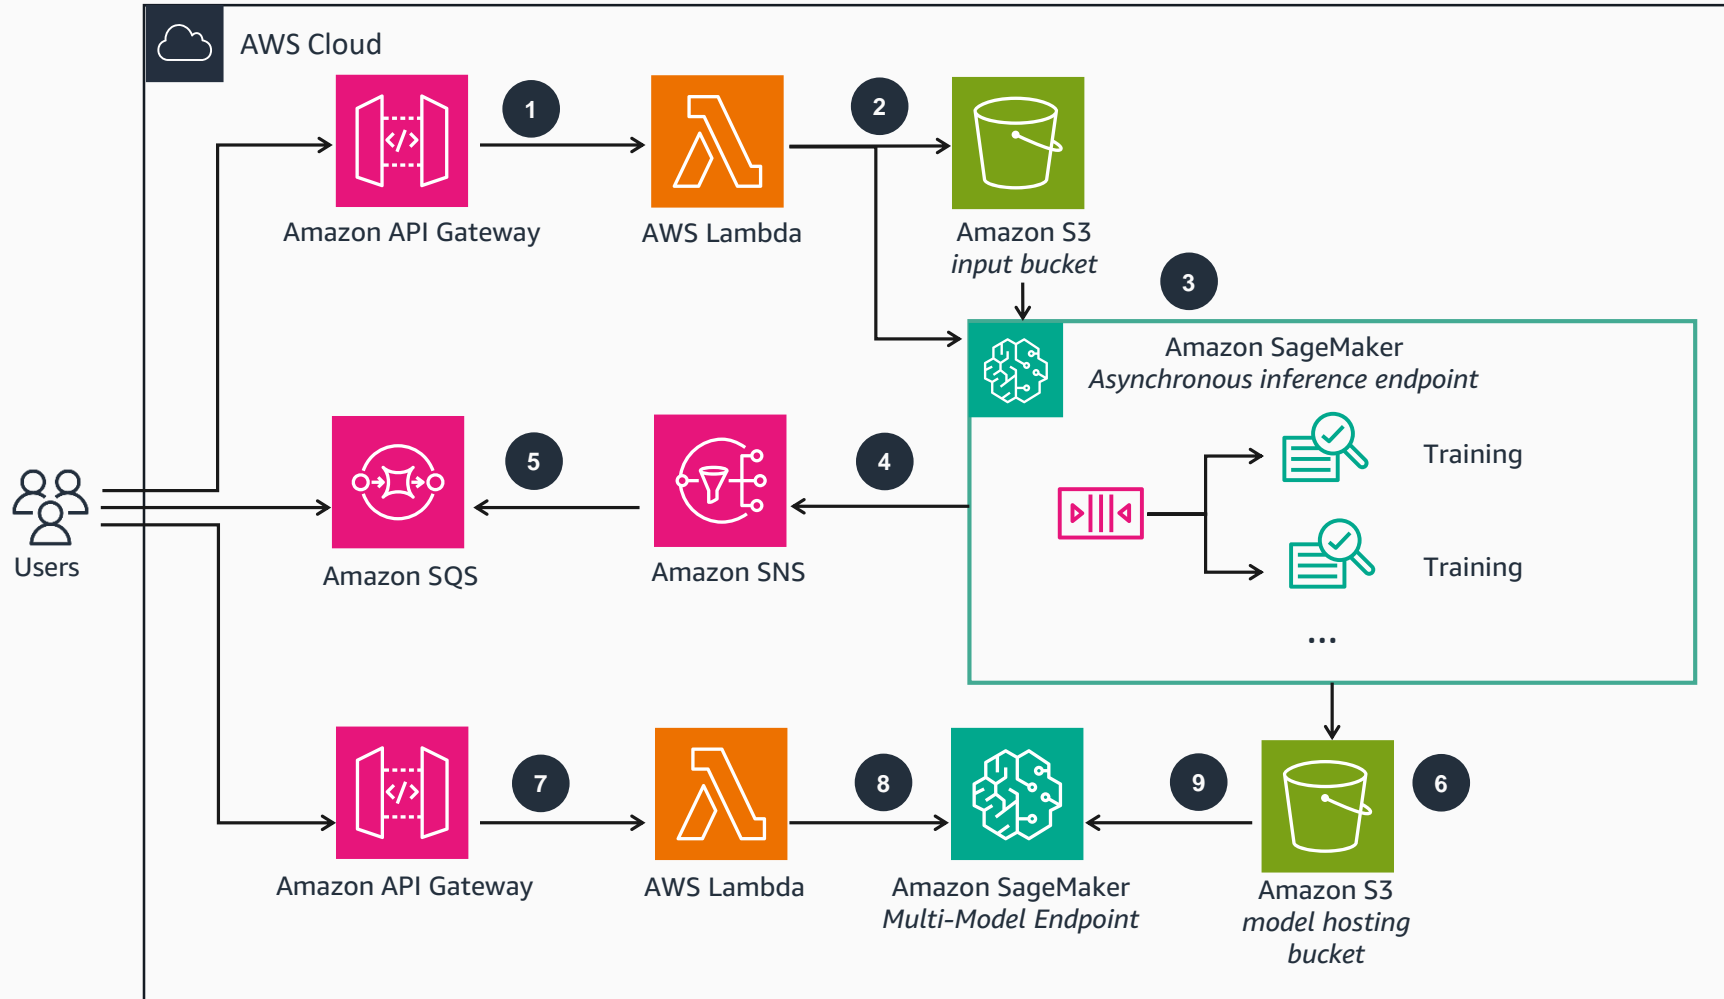


Guidance for Creating a Personalized Avatar with Amazon SageMaker

This architecture diagram shows how to calibrate and deploy a stable diffusion model at scale, and use the calibrated models to generate personalized avatars with a simple text prompt.



- 1 Initiate training through a call to an **Amazon API Gateway** RESTful API endpoint using **AWS Identity and Access Management (IAM)** authentication.
- 2 An **AWS Lambda** function packages user images and training configuration files, and uploads them to an **Amazon Simple Storage Service (Amazon S3)** bucket. It then invokes the training job.
- 3 An **Amazon SageMaker** asynchronous inference manages the training process. Training jobs are automatically queued before going through image preparation, calibration, and post-processing steps.
- 4 **SageMaker** publishes job status through **Amazon Simple Notification Service (Amazon SNS)** topics.
- 5 User application subscribes to **Amazon Simple Queue Service (Amazon SQS)** for update when a training job is completed.
- 6 Model artifacts are uploaded to **Amazon S3** model hosting bucket.
- 7 Initiate inference through a call to an **API Gateway** RESTful API endpoint using **IAM** authentication.
- 8 **Lambda** function invokes the model endpoint.
- 9 **SageMaker** Multi-Model Endpoints (MME) provide inference from dynamically loaded and cached personalized models from the **Amazon S3** model hosting bucket, based on the traffic pattern to each model.

