

Scaling Your Generative AI Solution with AWS

Accelerate and scale your Large Language Models on AWS

As you've advanced the development of your Large Language Models or other generative models, have you run into infrastructure limitations? Scaling generative AI can be resource-intensive, requiring significant compute capacity and advanced data architecture.

Mission Cloud has worked with several leading companies who are creating their own generative models using text, audio, and images. With our deep AWS expertise, we've helped them to design the proper AWS infrastructure for their models and to handle training datasets in excess of 50 terabytes.

Scale Your Model

We design your infrastructure based on AWS best practices and utilize common tooling like Amazon FSx for Lustre, AWS Sagemaker, AWS Sagemaker Pipelines, MLFlow, and others. We can also expedite the journey of your models from on premises or another cloud provider into the AWS ecosystem.

To do this, we'll work backward with you to understand all the requirements you need fulfilled by your Generative AI solution. We analyze things like:

Your current models

Business needs

Usage patterns for inference

UI for experimentation

Model Monitoring

Automation opportunities for QA, Training, and Parameter Validation

From these, we'll develop a best-fit architectural approach and design everything from model stores, experimentation, endpoint deployment, pipelines, and the data infrastructure. As a part of every engagement, we document this system from top to bottom and conduct extensive knowledge transfer so that your team can maintain it into the future.

Our Offer

Want to learn more about how to build your Generative AI with best practices and scale its development on AWS? Take a **FREE 60-minute consultation** with one of our AI specialists to discuss your ideas, concerns, and needs.



Mission Cloud is a next-generation cloud services provider, combining managed and professional services. We help businesses migrate, manage, modernize and optimize their AWS environments at every stage of their cloud journey.

1 (855) 647-7466 • www.missioncloud.com • sales@missioncloud.com

