Automated Safeguards for Sagemaker

Amazon Sagemaker is is a fully-managed service that covers the entire machine learning workflow to label and prepare your data, choose an algorithm, train the algorithm, tune and optimize it for deployment, make predictions, and take action. Your models get to production faster with much less effort and lower cost. Sagemaker is an easy way for customers to being to use machine learning models with their data, without a large or complicated investment. Please visit Amazon Sagemaker for details.

Amazon Sagemaker is priced based on usage across three categories: Machine Learning (ML) instance usage, Inference Acceleration with elastic GPUs, and storage. Please see the Amazon Sagemaker Pricing page for details.

Amazon SageMaker includes three modules: Build, Train, and Deploy. The Build module provides a hosted environment to work with your data, experiment with algorithms, and visualize your output. The Train module allows for one-click model training and tuning at high-scale and low cost. The Deploy module provides a managed environment for you to easily host and test models for inference securely and with low latency. Amazon SageMaker builds fully managed instances running Jupyter notebooks for training data exploration and preprocessing. Amazon SageMaker can automatically tune your model by adjusting thousands of different combinations of algorithm parameters, to arrive at the most accurate predictions the model is capable of producing. Amazon SageMaker manages your production compute infrastructure on your behalf to perform health checks, apply security patches, and conduct other routine maintenance, all with built-in Amazon CloudWatch monitoring and logging. Please see Amazon Sagemaker Features for details.

ClearDATA Automated Safeguards for Amazon Sagemaker help ensure data is encryption in transit when multiple instances are used for parallel model training jobs. ClearDATA Automated Safeguards for S3 are also used to ensure that the data used both inbound and outbound from the Sagemaker instance is encrypted at rest in the S3 bucket and in motion between the S3 and Sagemaker services.

Compliance Guidance

Encryption in Transit

Amazon Sagemaker ensures that any data that is processed on a compute instance is encrypted, while the ClearDATA Automated Safeguards ensure encryption in motion between the Sagemaker instance nodes and between Sagemaker and Amazon S3 buckets.

Remediation

If the encryption in transit settings are not properly configured when creating a Hyperparameter Tuning Job, the job will be terminated.

Shared Responsibility

There are no restrictions when using Amazon Sagemaker.

Please contact your ClearDATA team for a copy of the Responsibilities Matrix.

Exclusion

Exclusions are available on a per-instance basis. Please contact ClearDATA Support if you require a Safeguard exclusion.

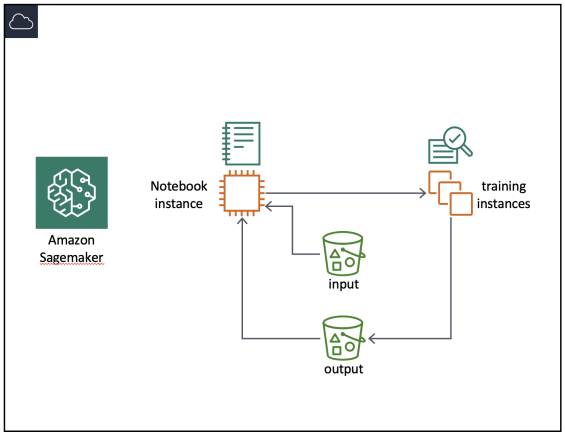

Users can be added to the Safeguard-Sagemaker IAM group in order to access the Amazon Sagemaker service.

Item	ClearDATA	Customer
Enforcement of Automated Safeguards	RA	IC
Creation, management, and configuration of all Sagemaker notebooks and features	C	RA
Ensure any service excluded from automated remediation does not contain any PHI/PII	IC	RA