[2025-November-NewBraindump2go MLA-C01 PDF Dumps Free Share[Q78-Q101

2025/November Latest Braindump2go MLA-C01 Exam Dumps with PDF and VCE Free Updated Today! Following are some new Braindump2go MLA-C01 Real Exam Questions!QUESTION 78A company is planning to use Amazon SageMaker to make classification ratings that are based on images. The company has 6 ? of training data that is stored on an Amazon FSx for NetApp ONTAP system virtual machine (SVM). The SVM is in the same VPC as SageMaker.An ML engineer must make the training data accessible for ML models that are in the SageMaker environment.Which solution will meet these requirements?A. Mount the FSx for ONTAP file system as a volume to the SageMaker Instance.B. Create an Amazon S3 bucket. Use Mountpoint for Amazon S3 to link the S3 bucket to the FSx for ONTAP file system.C. Create a catalog connection from SageMaker Data Wrangler to the FSx for ONTAP file system.D. Create a direct connection from SageMaker Data Wrangler to the FSx for ONTAP file system.Answer: AQUESTION 79A company regularly receives new training data from the vendor of an ML model. The vendor delivers cleaned and prepared data to the company's Amazon S3 bucket every 3-4 days.The company has an Amazon SageMaker pipeline to retrain the model. An ML engineer needs to implement a solution to run the pipeline when new data is uploaded to the S3 bucket.Which solution will meet these requirements with the LEAST operational effort?A. Create an S3 Lifecycle rule to transfer the data to the SageMaker training instance and to initiate training.B. Create an AWS Lambda function that scans the S3 bucket. Program the Lambda function to initiate the pipeline when new data is uploaded.C. Create an Amazon EventBridge rule that has an event pattern that matches the S3 upload. Configure the pipeline as the target of the rule.D. Use Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to orchestrate the pipeline when new data is uploaded.Answer: CQUESTION 80An ML engineer is developing a fraud detection model by using the Amazon SageMaker XGBoost algorithm. The model classifies transactions as either fraudulent or legitimate.During testing, the model excels at identifying fraud in the training dataset. However, the model is inefficient at identifying fraud in new and unseen transactions.What should the ML engineer do to improve the fraud detection for new transactions?A. Increase the learning rate.B. Remove some irrelevant features from the training dataset.C. Increase the value of the max_depth hyperparameter.D. Decrease the value of the max_depth hyperparameter.Answer: DQUESTION 81A company has a binary classification model in production. An ML engineer needs to develop a new version of the model.The new model version must maximize correct predictions of positive labels and negative labels. The ML engineer must use a metric to recalibrate the model to meet these requirements.Which metric should the ML engineer use for the model recalibration?A. AccuracyB. PrecisionC. RecallD. SpecificityAnswer: AQUESTION 82A company is using Amazon SageMaker to create ML models. The company's data scientists need fine-grained control of the ML workflows that they orchestrate. The data scientists also need the ability to visualize SageMaker jobs and workflows as a directed acyclic graph (DAG). The data scientists must keep a running history of model discovery experiments and must establish model governance for auditing and compliance verifications. Which solution will meet these requirements?A. Use AWS CodePipeline and its integration with SageMaker Studio to manage the entire ML workflows. Use SageMaker ML Lineage Tracking for the running history of experiments and for auditing and compliance verifications.B. Use AWS CodePipeline and its integration with SageMaker Experiments to manage the entire ML workflows. Use SageMaker Experiments for the running history of experiments and for auditing and compliance verifications.C. Use SageMaker Pipelines and its integration with SageMaker Studio to manage the entire ML workflows. Use SageMaker ML Lineage Tracking for the running history of experiments and for auditing and compliance verifications.D. Use SageMaker Pipelines and its integration with SageMaker Experiments to manage the entire ML workflows. Use SageMaker Experiments for the running history of experiments and for auditing and compliance verifications.Answer: DQUESTION 83A company wants to reduce the cost of its containerized ML applications. The applications use ML models that run on Amazon EC2 instances, AWS Lambda functions, and an Amazon Elastic Container Service (Amazon ECS) cluster. The EC2 workloads and ECS workloads use Amazon Elastic Block Store (Amazon EBS) volumes to save predictions and artifacts.An ML engineer must identify resources that are being used inefficiently. The ML engineer also must generate recommendations to reduce the cost of these resources.Which solution will meet these requirements with the LEAST development effort?A. Create code to evaluate each instance's memory and compute usage.B. Add cost allocation tags to the resources. Activate the tags in AWS Billing and Cost Management.C. Check AWS CloudTrail event history for the creation of the resources.D. Run AWS Compute Optimizer.Answer: DQUESTION 84A company needs to create a central catalog for all the company's ML models. The models are in AWS accounts where the company developed the models initially. The models are hosted in Amazon Elastic Container Registry (Amazon ECR) repositories.Which solution will meet these requirements?A. Configure ECR cross-account replication for each existing ECR repository. Ensure that each model is visible in each AWS account.B. Create a new AWS account with a new ECR repository as the central catalog. Configure ECR

cross-account replication between the initial ECR repositories and the central catalog.C. Use the Amazon SageMaker Model Registry to create a model group for models hosted in Amazon ECR. Create a new AWS account. In the new account, use the SageMaker Model Registry as the central catalog. Attach a cross-account resource policy to each model group in the initial AWS accounts.D. Use an AWS Glue Data Catalog to store the models. Run an AWS Glue crawler to migrate the models from the ECR repositories to the Data Catalog. Configure cross-account access to the Data Catalog.Answer: CQUESTION 85A company has developed a new ML model. The company requires online model validation on 10% of the traffic before the company fully releases the model in production. The company uses an Amazon SageMaker endpoint behind an Application Load Balancer (ALB) to serve the model.Which solution will set up the required online validation with the LEAST operational overhead?A. Use production variants to add the new model to the existing SageMaker endpoint. Set the variant weight to 0.1 for the new model. Monitor the number of invocations by using Amazon CloudWatch.B. Use production variants to add the new model to the existing SageMaker endpoint. Set the variant weight to 1 for the new model. Monitor the number of invocations by using Amazon CloudWatch.C. Create a new SageMaker endpoint. Use production variants to add the new model to the new endpoint. Monitor the number of invocations by using Amazon CloudWatch.D. Configure the ALB to route 10% of the traffic to the new model at the existing SageMaker endpoint. Monitor the number of invocations by using AWS CloudTrail.Answer: AQUESTION 86A company runs training jobs on Amazon SageMaker by using a compute optimized instance. Demand for training runs will remain constant for the next 55 weeks. The instance needs to run for 35 hours each week. The company needs to reduce its model training costs.Which solution will meet these requirements?A. Use a serverless endpoint with a provisioned concurrency of 35 hours for each week. Run the training on the endpoint.B. Use SageMaker Edge Manager for the training. Specify the instance requirement in the edge device configuration. Run the training.C. Use the heterogeneous cluster feature of SageMaker Training. Configure the instance_type, instance_count, and instance_groups arguments to run training jobs.D. Opt in to a SageMaker Savings Plan with a 1-year term and an All Upfront payment. Run a SageMaker Training job on the instance.Answer: DQUESTION 87A company deployed an ML model that uses the XGBoost algorithm to predict product failures. The model is hosted on an Amazon SageMaker endpoint and is trained on normal operating data. An AWS Lambda function provides the predictions to the company's application. An ML engineer must implement a solution that uses incoming live data to detect decreased model accuracy over time.Which solution will meet these requirements?A. Use Amazon CloudWatch to create a dashboard that monitors real-time inference data and model predictions. Use the dashboard to detect drift.B. Modify the Lambda function to calculate model drift by using real-time inference data and model predictions. Program the Lambda function to send alerts.C. Schedule a monitoring job in SageMaker Model Monitor. Use the job to detect drift by analyzing the live data against a baseline of the training data statistics and constraints. D. Schedule a monitoring job in SageMaker Debugger. Use the job to detect drift by analyzing the live data against a baseline of the training data statistics and constraints.Answer: CQUESTION 88A company has an ML model that uses historical transaction data to predict customer behavior. An ML engineer is optimizing the model in Amazon SageMaker to enhance the model's predictive accuracy. The ML engineer must examine the input data and the resulting predictions to identify trends that could skew the model's performance across different demographics.Which solution will provide this level of analysis?A. Use Amazon CloudWatch to monitor network metrics and CPU metrics for resource optimization during model training.B. Create AWS Glue DataBrew recipes to correct the data based on statistics from the model output.C. Use SageMaker Clarify to evaluate the model and training data for underlying patterns that might affect accuracy.D. Create AWS Lambda functions to automate data pre-processing and to ensure consistent quality of input data for the model.Answer: CQUESTION 89A company needs to use Retrieval Augmented Generation (RAG) to supplement an open source large language model (LLM) that runs on Amazon Bedrock. The company's data for RAG is a set of documents in an Amazon S3 bucket. The documents consist of .csv files and .docx files. Which solution will meet these requirements with the LEAST operational overhead?A. Create a pipeline in Amazon SageMaker Pipelines to generate a new model. Call the new model from Amazon Bedrock to perform RAG queries.B. Convert the data into vectors. Store the data in an Amazon Neptune database. Connect the database to Amazon Bedrock. Call the Amazon Bedrock API to perform RAG queries.C. Fine-tune an existing LLM by using an AutoML job in Amazon SageMaker. Configure the S3 bucket as a data source for the AutoML job. Deploy the LLM to a SageMaker endpoint. Use the endpoint to perform RAG queries.D. Create a knowledge base for Amazon Bedrock. Configure a data source that references the S3 bucket. Use the Amazon Bedrock API to perform RAG queries.Answer: DQUESTION 90A company plans to deploy an ML model for production inference on an Amazon SageMaker endpoint. The average inference payload size will vary from 100 MB to 300 MB. Inference requests must be processed in 60 minutes or less.Which SageMaker inference option will meet these requirements?A. Serverless inferenceB. Asynchronous inferenceC. Real-time inferenceD. Batch transformAnswer: BQUESTION 91An ML engineer notices class imbalance in an image classification training job.What should the ML engineer do to resolve this issue?A. Reduce the size of the dataset.B.

Transform some of the images in the dataset.C. Apply random oversampling on the dataset.D. Apply random data splitting on the dataset.Answer: CQUESTION 92A company receives daily .csv files about customer interactions with its ML model. The company stores the files in Amazon S3 and uses the files to retrain the model. An ML engineer needs to implement a solution to mask credit card numbers in the files before the model is retrained.Which solution will meet this requirement with the LEAST development effort?A. Create a discovery job in Amazon Macie. Configure the job to find and mask sensitive data.B. Create Apache Spark code to run on an AWS Glue job. Use the Sensitive Data Detection functionality in AWS Glue to find and mask sensitive data.C. Create Apache Spark code to run on an AWS Glue job. Program the code to perform a regex operation to find and mask sensitive data.D. Create Apache Spark code to run on an Amazon EC2 instance. Program the code to perform an operation to find and mask sensitive data.Answer: BQUESTION 93A medical company is using AWS to build a tool to recommend treatments for patients. The company has obtained health records and self-reported textual information in English from patients. The company needs to use this information to gain insight about the patients.Which solution will meet this requirement with the LEAST development effort?A. Use Amazon SageMaker to build a recurrent neural network (RNN) to summarize the data.B. Use Amazon Comprehend Medical to summarize the data.C. Use Amazon Kendra to create a quick-search tool to query the data.D. Use the Amazon SageMaker Sequence-to-Sequence (seq2seq) algorithm to create a text summary from the data.Answer: BQUESTION 94A company needs to extract entities from a PDF document to build a classifier model.Which solution will extract and store the entities in the LEAST amount of time?A. Use Amazon Comprehend to extract the entities. Store the output in Amazon S3.B. Use an open source AI optical character recognition (OCR) tool on Amazon SageMaker to extract the entities. Store the output in Amazon S3.C. Use Amazon Textract to extract the entities. Use Amazon Comprehend to convert the entities to text. Store the output in Amazon S3.D. Use Amazon Textract integrated with Amazon Augmented AI (Amazon A2I) to extract the entities. Store the output in Amazon S3.Answer: CQUESTION 95A company shares Amazon SageMaker Studio notebooks that are accessible through a VPN. The company must enforce access controls to prevent malicious actors from exploiting presigned URLs to access the notebooks.Which solution will meet these requirements?A. Set up Studio client IP validation by using the aws:sourceIp IAM policy condition.B. Set up Studio client VPC validation by using the aws:sourceVpc IAM policy condition.C. Set up Studio client role endpoint validation by using the aws:PrimaryTag IAM policy condition.D. Set up Studio client user endpoint validation by using the aws:PrincipalTag IAM policy condition.Answer: AQUESTION 96An ML engineer needs to merge and transform data from two sources to retrain an existing ML model. One data source consists of .csv files that are stored in an Amazon S3 bucket. Each .csv file consists of millions of records. The other data source is an Amazon Aurora DB cluster.The result of the merge process must be written to a second S3 bucket. The ML engineer needs to perform this merge-and-transform task every week.Which solution will meet these requirements with the LEAST operational overhead?A. Create a transient Amazon EMR cluster every week. Use the cluster to run an Apache Spark job to merge and transform the data.B. Create a weekly AWS Glue job that uses the Apache Spark engine. Use DynamicFrame native operations to merge and transform the data.C. Create an AWS Lambda function that runs Apache Spark code every week to merge and transform the data. Configure the Lambda function to connect to the initial S3 bucket and the DB cluster.D. Create an AWS Batch job that runs Apache Spark code on Amazon EC2 instances every week. Configure the Spark code to save the data from the EC2 instances to the second S3 bucket.Answer: BQUESTION 97An ML engineer has deployed an Amazon SageMaker model to a serverless endpoint in production. The model is invoked by the InvokeEndpoint API operation.The model's latency in production is higher than the baseline latency in the test environment. The ML engineer thinks that the increase in latency is because of model startup time.What should the ML engineer do to confirm or deny this hypothesis?A. Schedule a SageMaker Model Monitor job. Observe metrics about model quality.B. Schedule a SageMaker Model Monitor job with Amazon CloudWatch metrics enabled.C. Enable Amazon CloudWatch metrics. Observe the ModelSetupTime metric in the SageMaker namespace.D. Enable Amazon CloudWatch metrics. Observe the ModelLoadingWaitTime metric in the SageMaker namespace.Answer: DQUESTION 98An ML engineer needs to ensure that a dataset complies with regulations for personally identifiable information (PII). The ML engineer will use the data to train an ML model on Amazon SageMaker instances. SageMaker must not use any of the PII.Which solution will meet these requirements in the MOST operationally efficient way?A. Use the Amazon Comprehend DetectPiiEntities API call to redact the PII from the data. Store the data in an Amazon S3 bucket. Access the S3 bucket from the SageMaker instances for model training.B. Use the Amazon Comprehend DetectPiiEntities API call to redact the PII from the data. Store the data in an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system to the SageMaker instances for model training.C. Use AWS Glue DataBrew to cleanse the dataset of PII. Store the data in an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system to the SageMaker instances for model training.D. Use Amazon Macie for automatic discovery of PII in the data. Remove the PII. Store the data in an Amazon S3 bucket. Mount the S3 bucket to the SageMaker instances for model

training.Answer: AQUESTION 99A company must install a custom script on any newly created Amazon SageMaker notebook instances.Which solution will meet this requirement with the LEAST operational overhead?A. Create a lifecycle configuration script to install the custom script when a new SageMaker notebook is created. Attach the lifecycle configuration to every new SageMaker notebook as part of the creation steps.B. Create a custom Amazon Elastic Container Registry (Amazon ECR) image that contains the custom script. Push the ECR image to a Docker registry. Attach the Docker image to a SageMaker Studio domain. Select the kernel to run as part of the SageMaker notebook.C. Create a custom package index repository. Use AWS CodeArtifact to manage the installation of the custom script. Set up AWS PrivateLink endpoints to connect CodeArtifact to the SageMaker instance. Install the script.D. Store the custom script in Amazon S3. Create an AWS Lambda function to install the custom script on new SageMaker notebooks. Configure Amazon EventBridge to invoke the Lambda function when a new SageMaker notebook is initialized.Answer: AQUESTION 100A company is building a real-time data processing pipeline for an ecommerce application. The application generates a high volume of clickstream data that must be ingested, processed, and visualized in near real time. The company needs a solution that supports SQL for data processing and Jupyter notebooks for interactive analysis.Which solution will meet these requirements?A. Use Amazon Data Firehose to ingest the data. Create an AWS Lambda function to process the data. Store the processed data in Amazon S3. Use Amazon QuickSight to visualize the data.B. Use Amazon Kinesis Data Streams to ingest the data. Use Amazon Data Firehose to transform the data. Use Amazon Athena to process the data. Use Amazon QuickSight to visualize the data.C. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to ingest the data. Use AWS Glue with PySpark to process the data. Store the processed data in Amazon S3. Use Amazon QuickSight to visualize the data.D. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to ingest the data. Use Amazon Managed Service for Apache Flink to process the data. Use the built-in Flink dashboard to visualize the data.Answer: DQUESTION 101A medical company needs to store clinical data. The data includes personally identifiable information (PII) and protected health information (PHI).An ML engineer needs to implement a solution to ensure that the PII and PHI are not used to train ML models.Which solution will meet these requirements?A. Store the clinical data in Amazon S3 buckets. Use AWS Glue DataBrew to mask the PII and PHI before the data is used for model training.B. Upload the clinical data to an Amazon Redshift database. Use built-in SQL stored procedures to automatically classify and mask the PII and PHI before the data is used for model training.C. Use Amazon Comprehend to detect and mask the PII before the data is used for model training. Use Amazon Comprehend Medical to detect and mask the PHI before the data is used for model training.D. Create an AWS Lambda function to encrypt the PII and PHI. Program the Lambda function to save the encrypted data to an Amazon S3 bucket for model training.Answer: CResources From:1.2025 Latest Braindump2go MLA-C01 Exam Dumps (PDF & VCE) Free Share:**https://www.braindump2go.com/mla-c01.html**2.2025 Latest Braindump2go MLA-C01 PDF and MLA-C01 VCE Dumps Free Share:
https://drive.google.com/drive/folders/177LfePVjZ3YlBGmLxY4fy1Ft3ZTbtZm_?usp=sharing3.2025 Free Braindump2go MLA-C01 Exam Questions Download:https://www.braindump2go.com/free-online-pdf/MLA-C01-VCE-Dumps(78-101).pdf
Free Resources from Braindump2go,We Devoted to Helping You 100% Pass All Exams!