

Reliable MLS-C01 Exam Review | Latest MLS-C01 Exam Book



What's more, part of that ITCertMagic MLS-C01 dumps now are free: <https://drive.google.com/open?id=1NADx-Aia0NwFK3jdxyoWb9cnxEgd6l2V>

We don't want you to prepare and practice the old questions and waste time. Therefore, our team of certified experts includes updated AWS Certified Machine Learning - Specialty MLS-C01 Exam Questions as soon as they are released. ITCertMagic provides up-to-date Amazon exam questions.

Amazon MLS-C01 Exam covers a wide range of topics related to machine learning, including data preparation, feature engineering, model selection, model training, and model deployment. It also assesses the candidate's understanding of AWS services such as Amazon SageMaker, Amazon Rekognition, Amazon Comprehend, and Amazon Translate, which are commonly used for machine learning applications.

>> Reliable MLS-C01 Exam Review <<

MLS-C01 testing engine training online | MLS-C01 test dumps

Now you do not need to worry about the relevancy and top standard of ITCertMagic AWS Certified Machine Learning - Specialty (MLS-C01) exam questions. These Amazon MLS-C01 dumps are designed and verified by qualified AWS Certified Machine Learning - Specialty (MLS-C01) exam trainers. Now you can trust ITCertMagic AWS Certified Machine Learning - Specialty (MLS-C01) practice questions and start preparation without wasting further time.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q150-Q155):

NEW QUESTION # 150

A company needs to deploy a chatbot to answer common questions from customers. The chatbot must base its answers on company documentation.

Which solution will meet these requirements with the LEAST development effort?

- A. Train an Amazon SageMaker BlazingText model based on past customer questions and company documents. Deploy the model as a real-time SageMaker endpoint. Integrate the model with the chatbot by using the SageMaker Runtime InvokeEndpoint API operation to answer customer questions.
- B. Index company documents by using Amazon OpenSearch Service. Integrate the chatbot with OpenSearch Service by using the OpenSearch Service k-nearest neighbors (k-NN) Query API operation to answer customer questions.
- C. Train a Bidirectional Attention Flow (BiDAF) network based on past customer questions and company documents. Deploy the model as a real-time Amazon SageMaker endpoint. Integrate the model with the chatbot by using the SageMaker Runtime InvokeEndpoint API operation to answer customer questions.
- D. Index company documents by using Amazon Kendra. Integrate the chatbot with Amazon Kendra by using the Amazon Kendra Query API operation to answer customer questions.

Answer: D

Explanation:

Explanation

The solution A will meet the requirements with the least development effort because it uses Amazon Kendra, which is a highly accurate and easy to use intelligent search service powered by machine learning. Amazon Kendra can index company documents from various sources and formats, such as PDF, HTML, Word, and more. Amazon Kendra can also integrate with chatbots by using the Amazon Kendra Query API operation, which can understand natural language questions and provide relevant answers from the indexed documents. Amazon Kendra can also provide additional information, such as document excerpts, links, and FAQs, to enhance the chatbot experience1.

The other options are not suitable because:

Option B: Training a Bidirectional Attention Flow (BiDAF) network based on past customer questions and company documents, deploying the model as a real-time Amazon SageMaker endpoint, and integrating the model with the chatbot by using the SageMaker Runtime InvokeEndpoint API operation will incur more development effort than using Amazon Kendra. The company will have to write the code for the BiDAF network, which is a complex deep learning model for question answering. The company will also have to manage the SageMaker endpoint, the model artifact, and the inference logic2.

Option C: Training an Amazon SageMaker BlazingText model based on past customer questions and company documents, deploying the model as a real-time SageMaker endpoint, and integrating the model with the chatbot by using the SageMaker Runtime InvokeEndpoint API operation will incur more development effort than using Amazon Kendra. The company will have to write the code for the BlazingText model, which is a fast and scalable text classification and word embedding algorithm. The company will also have to manage the SageMaker endpoint, the model artifact, and the inference logic3.

Option D: Indexing company documents by using Amazon OpenSearch Service and integrating the chatbot with OpenSearch Service by using the OpenSearch Service k-nearest neighbors (k-NN) Query API operation will not meet the requirements effectively. Amazon OpenSearch Service is a fully managed service that provides fast and scalable search and analytics capabilities. However, it is not designed for natural language question answering, and it may not provide accurate or relevant answers for the chatbot. Moreover, the k-NN Query API operation is used to find the most similar documents or vectors based on a distance function, not to find the best answers based on a natural language query4.

References:

1: Amazon Kendra

2: Bidirectional Attention Flow for Machine Comprehension

3: Amazon SageMaker BlazingText

4: Amazon OpenSearch Service

NEW QUESTION # 151

An e-commerce company wants to launch a new cloud-based product recommendation feature for its web application. Due to data localization regulations, any sensitive data must not leave its on-premises data center, and the product recommendation model must be trained and tested using nonsensitive data only. Data transfer to the cloud must use IPsec. The web application is hosted on-premises with a PostgreSQL database that contains all the data. The company wants the data to be uploaded securely to Amazon S3 each day for model retraining.

How should a machine learning specialist meet these requirements?

- A. Create an AWS Glue job to connect to the PostgreSQL DB instance. Ingest all data through an AWS Site- to-Site VPN connection into Amazon S3 while removing sensitive data using a PySpark job.
- B. Use PostgreSQL logical replication to replicate all data to PostgreSQL in Amazon EC2 through AWS Direct Connect with a VPN connection. Use AWS Glue to move data from Amazon EC2 to Amazon S3.
- C. Create an AWS Glue job to connect to the PostgreSQL DB instance. Ingest tables without sensitive data through an AWS Site-to-Site VPN connection directly into Amazon S3.
- D. Use AWS Database Migration Service (AWS DMS) with table mapping to select PostgreSQL tables with no sensitive data through an SSL connection. Replicate data directly into Amazon S3.

Answer: D

Explanation:

Explanation

The best option is to use AWS Database Migration Service (AWS DMS) with table mapping to select PostgreSQL tables with no sensitive data through an SSL connection. Replicate data directly into Amazon S3.

This option meets the following requirements:

It ensures that only nonsensitive data is transferred to the cloud by using table mapping to filter out the tables that contain sensitive data1.

It uses IPsec to secure the data transfer by enabling SSL encryption for the AWS DMS endpoint2.

It uploads the data to Amazon S3 each day for model retraining by using the ongoing replication feature of AWS DMS3.

The other options are not as effective or feasible as the option above. Creating an AWS Glue job to connect to the PostgreSQL DB instance and ingest data through an AWS Site-to-Site VPN connection directly into Amazon S3 is possible, but it requires more steps and resources than using AWS DMS. Also, it does not specify how to filter out the sensitive data from the tables. Creating an AWS Glue job to connect to the PostgreSQL DB instance and ingest all data through an AWS Site-to-Site VPN connection into Amazon S3 while removing sensitive data using a PySpark job is also possible, but it is more complex and error-prone than using AWS DMS. Also, it does not use IPsec as required. Using PostgreSQL logical replication to replicate all data to PostgreSQL in Amazon EC2 through AWS Direct Connect with a VPN connection, and then using AWS Glue to move data from Amazon EC2 to Amazon S3 is not feasible, because PostgreSQL logical replication does not support replicating only a subset of data4. Also, it involves unnecessary data movement and additional costs.

References:

[Table mapping - AWS Database Migration Service](#)

[Using SSL to encrypt a connection to a DB instance - AWS Database Migration Service](#)

[Ongoing replication - AWS Database Migration Service](#)

[Logical replication - PostgreSQL](#)

NEW QUESTION # 152

A gaming company has launched an online game where people can start playing for free, but they need to pay if they choose to use certain features. The company needs to build an automated system to predict whether or not a new user will become a paid user within 1 year.

The company has gathered a labeled dataset from 1 million users.

The training dataset consists of 1,000 positive samples (from users who ended up paying within 1 year) and 999,000 negative samples (from users who did not use any paid features). Each data sample consists of 200 features including user age, device, location, and play patterns.

Using this dataset for training, the Data Science team trained a random forest model that converged with over 99% accuracy on the training set. However, the prediction results on a test dataset were not satisfactory. Which of the following approaches should the Data Science team take to mitigate this issue?

(Choose two.)

- A. Change the cost function so that false positives have a higher impact on the cost value than false negatives.
- B. Add more deep trees to the random forest to enable the model to learn more features.
- C. Generate more positive samples by duplicating the positive samples and adding a small amount of noise to the duplicated data.
- D. **Include a copy of the samples in the test dataset in the training dataset.**
- E. **Change the cost function so that false negatives have a higher impact on the cost value than false positives.**

Answer: D,E**NEW QUESTION # 153**

A retail company is using Amazon Personalize to provide personalized product recommendations for its customers during a marketing campaign. The company sees a significant increase in sales of recommended items to existing customers immediately after deploying a new solution version, but these sales decrease a short time after deployment. Only historical data from before the marketing campaign is available for training.

How should a data scientist adjust the solution?

- A. Add event type and event value fields to the interactions dataset in Amazon Personalize.
- B. Implement a new solution using the built-in factorization machines (FM) algorithm in Amazon SageMaker.
- C. Add user metadata and use the HRNN-Metadata recipe in Amazon Personalize.
- D. **Use the event tracker in Amazon Personalize to include real-time user interactions.**

Answer: D

Explanation:

The best option is to use the event tracker in Amazon Personalize to include real-time user interactions. This will allow the model to learn from the feedback of the customers during the marketing campaign and adjust the recommendations accordingly. The event tracker can capture click-through, add-to-cart, purchase, and other types of events that indicate the user's preferences. By using the event tracker, the company can improve the relevance and freshness of the recommendations and avoid the decrease in sales.

The other options are not as effective as using the event tracker. Adding user metadata and using the HRNN-Metadata recipe in Amazon Personalize can help capture the user's attributes and preferences, but it will not reflect the changes in user behavior during the marketing campaign. Implementing a new solution using the built-in factorization machines (FM) algorithm in Amazon SageMaker can also provide personalized recommendations, but it will require more time and effort to train and deploy the model. Adding event type and event value fields to the interactions dataset in Amazon Personalize can help capture the importance and context of each interaction, but it will not update the model with the latest user feedback.

References:

Recording events - Amazon Personalize

Using real-time events - Amazon Personalize

NEW QUESTION # 154

A finance company needs to forecast the price of a commodity. The company has compiled a dataset of historical daily prices. A data scientist must train various forecasting models on 80% of the dataset and must validate the efficacy of those models on the remaining 20% of the dataset.

What should the data scientist split the dataset into a training dataset and a validation dataset to compare model performance?

- A. Pick a date so that 80% to the data points precede the date Assign that group of data points as the training dataset. Assign all the remaining data points to the validation dataset.
- B. Sample data points randomly without replacement so that 80% of the data points are in the training dataset. Assign all the remaining data points to the validation dataset.
- C. Starting from the earliest date in the dataset, pick eight data points for the training dataset and two data points for the validation dataset. Repeat this stratified sampling until no data points remain.
- D. Pick a date so that 80% of the data points occur after the date. Assign that group of data points as the training dataset. Assign all the remaining data points to the validation dataset.

Answer: A

Explanation:

A Comprehensive Explanation: The best way to split the dataset into a training dataset and a validation dataset is to pick a date so that 80% of the data points precede the date and assign that group of data points as the training dataset. This method preserves the temporal order of the data and ensures that the validation dataset reflects the most recent trends and patterns in the commodity price. This is important for forecasting models that rely on time series analysis and sequential data. The other methods would either introduce bias or lose information by ignoring the temporal structure of the data.

References:

Time Series Forecasting - Amazon SageMaker

Time Series Splitting - scikit-learn

Time Series Forecasting - Towards Data Science

NEW QUESTION # 155

.....

It never needs an internet connection. ITCertMagic's AWS Certified Machine Learning - Specialty practice exam software has several mock exams, designed just like the real exam. Amazon MLS-C01 practice exam software contains all the important questions which have a greater chance of appearing in the final exam. ITCertMagic always tries to ensure that you are provided with the most updated AWS Certified Machine Learning - Specialty (MLS-C01) Exam Questions to pass the exam on the first attempt.

Latest MLS-C01 Exam Book: <https://www.itcertmagic.com/Amazon/real-MLS-C01-exam-prep-dumps.html>

- MLS-C01 Online Test MLS-C01 Online Test MLS-C01 Latest Practice Questions Easily obtain  MLS-C01 for free download through www.prepawaypdf.com MLS-C01 Valid Guide Files
- Valid MLS-C01 Exam Dumps MLS-C01 Practice Test Online MLS-C01 Practice Test Online Search for  MLS-C01 and download exam materials for free through www.pdfvce.com Exam MLS-C01 Forum

- Free PDF Quiz Updated Amazon - Reliable MLS-C01 Exam Review □ Search for ⇒ MLS-C01 ⇄ and easily obtain a free download on ➡ www.prepawaypdf.com □□□ □New MLS-C01 Exam Testking
- MLS-C01 Latest Practice Questions □ MLS-C01 Question Explanations □ MLS-C01 Latest Practice Questions □ Search on ➡ www.pdfvce.com □ for ➡ MLS-C01 □ to obtain exam materials for free download □MLS-C01 Practice Test Online
- Free PDF Quiz Updated Amazon - Reliable MLS-C01 Exam Review □ Search for □ MLS-C01 □ and easily obtain a free download on 《 www.prepawayexam.com 》 □Valid MLS-C01 Exam Format
- MLS-C01 Online Test □ Actual MLS-C01 Test Answers □ MLS-C01 Question Explanations □ The page for free download of 《 MLS-C01 》 on ✓ www.pdfvce.com □✓ □ will open immediately □Valid MLS-C01 Exam Dumps
- MLS-C01 Online Test □ MLS-C01 Latest Test Sample □ MLS-C01 Trustworthy Exam Torrent □ Open ⇒ www.prep4sures.top ⇄ and search for ➡ MLS-C01 □ to download exam materials for free □MLS-C01 Exam Tests
- Pdfvce Real Amazon MLS-C01 Questions PDF □ Immediately open [www.pdfvce.com] and search for 《 MLS-C01 》 to obtain a free download □Test MLS-C01 Online
- www.vce4dumps.com Real Amazon MLS-C01 Questions PDF □ Search on ➡ www.vce4dumps.com □ for (MLS-C01) to obtain exam materials for free download □MLS-C01 Exam Certification
- Pdfvce Amazon MLS-C01 PDF □ Search for ➡ MLS-C01 □ and download it for free immediately on ▷ www.pdfvce.com ▲ □MLS-C01 Valid Guide Files
- Latest Updated Amazon Reliable MLS-C01 Exam Review - MLS-C01 Latest AWS Certified Machine Learning - Specialty Exam Book □ Copy URL 《 www.exam4labs.com 》 open and search for ⚡ MLS-C01 □⚡ □ to download for free □ □MLS-C01 Trustworthy Exam Torrent
- smartkidscampus.com, k12.instructure.com, bbs.t-firefly.com, bicyclebuysell.com, bbs.t-firefly.com, study.stcs.edu.np, class.most-d.com, bbs.t-firefly.com, www.stes.tyc.edu.tw, bbs.t-firefly.com, Disposable vapes

DOWNLOAD the newest ITCertMagic MLS-C01 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1NADx-Aia0NwFK3jdxyoWb9cnxEGd6l2V>