

Interactive AIF-C01 Questions & New AIF-C01 Exam Review

Introduction To AWS Certified AI Practitioner (AIF-C01) Practice Set

- Two sets
- Contains Practical set of questions
- Detailed answer with explanations
- 70 questions within 2 hours
- Created from my personal Experience
- Tried to check your knowledge with more Practical questions

DOWNLOAD the newest PassReview AIF-C01 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1qcpX4mVxaLgTE6chknBtr8JMbkkuZyTD>

There is no doubt that if you pass the AIF-C01 exam certification test, which means that your ability and professional knowledge are acknowledged by the authority field, we suggest that you can try our AIF-C01 reliable exam dumps. Although it is difficult to prepare the exam for most people, as long as you are attempting our AIF-C01 Exam Dumps, you will find that it is not as hard as you think. What you will never worry about is that the quality of AIF-C01 exam dumps, because once you haven't passed exam, we will have a 100% money back guarantee. You can easily pass the exam only if you spend some spare time studying our AIF-C01 materials.

Since the content of the examination is also updating daily, you will need real and latest Amazon AIF-C01 Exam Dumps to prepare successfully for the AIF-C01 certification exam in a short time. People who don't study from updated AWS Certified AI Practitioner (AIF-C01) questions fail the examination and loss time and money.

>> Interactive AIF-C01 Questions <<

New AIF-C01 Exam Review - Dumps AIF-C01 Collection

As the saying goes, an inch of gold is an inch of time. The more efficient the study guide is, the more our candidates will love and benefit from it. It is no exaggeration to say that you can successfully pass your AIF-C01 exams with the help our AIF-C01 learning torrent just for 20 to 30 hours even by your first attempt. And to cater to our customers' different study interests and hobbies, we have multiple choices on the AIF-C01 Exam Materials versions for you to choose: the PDF, the Software and the APP online.

Amazon AWS Certified AI Practitioner Sample Questions (Q64-Q69):

NEW QUESTION # 64

An online learning company with large volumes of educational materials wants to use enterprise search. Which AWS service meets these requirements?

- A. Amazon Personalize
- B. Amazon Textract
- C. Amazon Comprehend
- D. **Amazon Kendra**

Answer: D

Explanation:

* Amazon Kendra is an enterprise search service that uses machine learning to index and search unstructured data such as documents, manuals, and course materials.

* Amazon Comprehend performs NLP tasks like sentiment analysis and entity recognition, not enterprise search.

* Amazon Textract extracts structured text from scanned docs and PDFs.

* Amazon Personalize builds recommendation systems, not search engines.

Reference:

AWS Documentation - Amazon Kendra

NEW QUESTION # 65

A company wants to develop an educational game where users answer questions such as the following: "A jar contains six red, four green, and three yellow marbles. What is the probability of choosing a green marble from the jar?" Which solution meets these requirements with the LEAST operational overhead?

- A. Use reinforcement learning to train a model to return the probability.
- **B. Use code that will calculate probability by using simple rules and computations.**
- C. Use unsupervised learning to create a model that will estimate probability density.
- D. Use supervised learning to create a regression model that will predict probability.

Answer: B

Explanation:

The problem involves a simple probability calculation that can be handled efficiently by straightforward mathematical rules and computations. Using machine learning techniques would introduce unnecessary complexity and operational overhead.

* Option C (Correct): "Use code that will calculate probability by using simple rules and computations": This is the correct answer because it directly solves the problem with minimal overhead, using basic probability rules.

* Option A: "Use supervised learning to create a regression model" is incorrect as it overcomplicates the solution for a simple probability problem.

* Option B: "Use reinforcement learning to train a model" is incorrect because reinforcement learning is not needed for a simple probability calculation.

* Option D: "Use unsupervised learning to create a model" is incorrect as unsupervised learning is not applicable to this task.

AWS AI Practitioner References:

* Choosing the Right Solution for AI Tasks: AWS recommends using the simplest and most efficient approach to solve a given problem, avoiding unnecessary machine learning techniques for straightforward tasks.

NEW QUESTION # 66

Which AI technique combines large language models (LLMs) with external knowledge bases to improve response accuracy?

- **A. Retrieval Augmented Generation (RAG)**
- B. Natural language processing (NLP)
- C. Reinforcement learning (RL)
- D. Transfer learning

Answer: A

Explanation:

Comprehensive and Detailed Explanation From Exact AWS AI documents:

Retrieval Augmented Generation (RAG) enhances LLM responses by:

* Retrieving relevant information from external knowledge sources

* Injecting retrieved content into the prompt context

* Reducing hallucinations and improving factual accuracy

AWS generative AI guidance describes RAG as a best practice when models must use up-to-date or domain-specific knowledge that is not embedded in the model weights.

Why the other options are incorrect:

* RL (A) focuses on reward-based learning

* NLP (B) is a broad field, not a specific technique.

* Transfer learning (D) adapts model weights but does not retrieve external data at inference time.

AWS AI document references:

* Retrieval Augmented Generation on AWS

* Improving LLM Accuracy with External Knowledge

* Generative AI Architectures

NEW QUESTION # 67

A company is using Amazon SageMaker to develop AI models.

Select the correct SageMaker feature or resource from the following list for each step in the AI model lifecycle workflow. Each SageMaker feature or resource should be selected one time or not at all. (Select TWO.) SageMaker Clarify SageMaker Model Registry SageMaker Serverless Inference

Managing different versions of the model

Select...

Select...

SageMaker Clarify

SageMaker Model Registry

SageMaker Serverless Inference

Using the current model to make predictions

Select...

Select...

SageMaker Clarify

SageMaker Model Registry

SageMaker Serverless Inference

Answer:

Explanation:

Managing different versions of the model

Select...

Select...

SageMaker Clarify

SageMaker Model Registry

SageMaker Serverless Inference

Using the current model to make predictions

Select...

Select...

SageMaker Clarify

SageMaker Model Registry

SageMaker Serverless Inference

Explanation:

SageMaker Model Registry, SageMaker Serverless inference

This question requires selecting the appropriate Amazon SageMaker feature for two distinct steps in the AI model lifecycle. Let's break down each step and evaluate the options:

Step 1: Managing different versions of the model

The goal here is to identify a SageMaker feature that supports version control and management of machine learning models. Let's analyze the options:

SageMaker Clarify: This feature is used to detect bias in models and explain model predictions, helping with fairness and interpretability. It does not provide functionality for managing model versions.

SageMaker Model Registry: This is a centralized repository in Amazon SageMaker that allows users to catalog, manage, and track different versions of machine learning models. It supports model versioning, approval workflows, and deployment tracking, making it ideal for managing different versions of a model.

SageMaker Serverless Inference: This feature enables users to deploy models for inference without managing servers, automatically scaling based on demand. It is focused on inference (predictions), not on managing model versions.

Conclusion for Step 1: The SageMaker Model Registry is the correct choice for managing different versions of the model.

Exact Extract Reference: According to the AWS SageMaker documentation, "The SageMaker Model Registry allows you to catalog models for production, manage model versions, associate metadata, and manage approval status for deployment." (Source: AWS SageMaker Documentation - Model Registry, <https://docs.aws.amazon.com/sagemaker/latest/dg/model-registry.html>).

Step 2: Using the current model to make predictions

The goal here is to identify a SageMaker feature that facilitates making predictions (inference) with a deployed model. Let's evaluate the options:

SageMaker Clarify: As mentioned, this feature focuses on bias detection and explainability, not on performing inference or making predictions.

SageMaker Model Registry: While the Model Registry helps manage and catalog models, it is not used directly for making predictions. It can store models, but the actual inference process requires a deployment mechanism.

SageMaker Serverless Inference: This feature allows users to deploy models for inference without managing infrastructure. It automatically scales based on traffic and is specifically designed for making predictions in a cost-efficient, serverless manner. Conclusion for Step 2: SageMaker Serverless Inference is the correct choice for using the current model to make predictions. Exact Extract Reference: The AWS documentation states, "SageMaker Serverless Inference is a deployment option that allows you to deploy machine learning models for inference without configuring or managing servers. It automatically scales to handle inference requests, making it ideal for workloads with intermittent or unpredictable traffic." (Source: AWS SageMaker Documentation - Serverless Inference, <https://docs.aws.amazon.com/sagemaker/latest/dg/serverless-inference.html>).

Why Not Use the Same Feature Twice?

The question specifies that each SageMaker feature or resource should be selected one time or not at all. Since SageMaker Model Registry is used for version management and SageMaker Serverless Inference is used for predictions, each feature is selected exactly once. SageMaker Clarify is not applicable to either step, so it is not selected at all, fulfilling the question's requirements.

References:

AWS SageMaker Documentation: Model Registry (<https://docs.aws.amazon.com/sagemaker/latest/dg/model-registry.html>) AWS SageMaker Documentation: Serverless Inference (<https://docs.aws.amazon.com/sagemaker/latest/dg/serverless-inference.html>)

AWS AI Practitioner Study Guide (conceptual alignment with SageMaker features for model lifecycle management and inference) Let's format this question according to the specified structure and provide a detailed, verified answer based on AWS AI Practitioner knowledge and official AWS documentation. The question focuses on selecting an AWS database service that supports storage and queries of embeddings as vectors, which is relevant to generative AI applications.

NEW QUESTION # 68

An AI practitioner is using an Amazon Bedrock base model to summarize session chats from the customer service department. The AI practitioner wants to store invocation logs to monitor model input and output data.

Which strategy should the AI practitioner use?

- A. Configure AWS Audit Manager as the logs destination for the model.
- B. Configure model invocation logging in Amazon EventBridge.
- **C. Enable invocation logging in Amazon Bedrock.**
- D. Configure AWS CloudTrail as the logs destination for the model.

Answer: C

NEW QUESTION # 69

.....

It is not hard to know that AWS Certified AI Practitioner torrent prep is compiled by hundreds of industry experts based on the syllabus and development trends of industries that contain all the key points that may be involved in the examination. Therefore, with AIF-C01 exam questions, you no longer need to purchase any other review materials, and you also don't need to spend a lot of money on tutoring classes. At the same time, AIF-C01 Test Guide will provide you with very flexible learning time in order to help you pass the exam.

New AIF-C01 Exam Review: https://www.passreview.com/AIF-C01_exam-braindumps.html

Amazon Interactive AIF-C01 Questions On the other hand we establish excellent relation with IT certification staff of international large companies so that we can always get the latest news about change or updates about real exam, Amazon Interactive AIF-C01 Questions Our PDF version is suitable for reading and printing requests, Before buying our AIF-C01 reliable study pdf, you can download AIF-C01 review free demo for try.

In Word, Excel, Outlook, and Publisher Publisher does not include SmartArt New AIF-C01 Exam Review graphics or the Screenshot command) the commands for inserting different graphic types are bundled into the Illustrations command group.

Interactive AIF-C01 Questions - Free PDF AIF-C01 - First-grade New AWS Certified AI Practitioner Exam Review

You become eligible for high-paying jobs and promotions in your current firm after earning the AIF-C01 AWS Certified AI Practitioner certification, On the other hand we establish excellent relation with IT certification staff of international AIF-C01 large companies so that we can always get the latest news about change or updates about real exam.

Our PDF version is suitable for reading and printing requests. Before buying our AIF-C01 reliable study pdf, you can download AIF-C01 review free demo for try.

But as long as you use the trial Latest AIF-C01 Exam Pdf version, you will believe what I say, You should run for it.

What's more, part of that PassReview AIF-C01 dumps now are free: <https://drive.google.com/open?id=1qcpX4mVxaLgTE6chknBtr8JMbkkUZyTD>