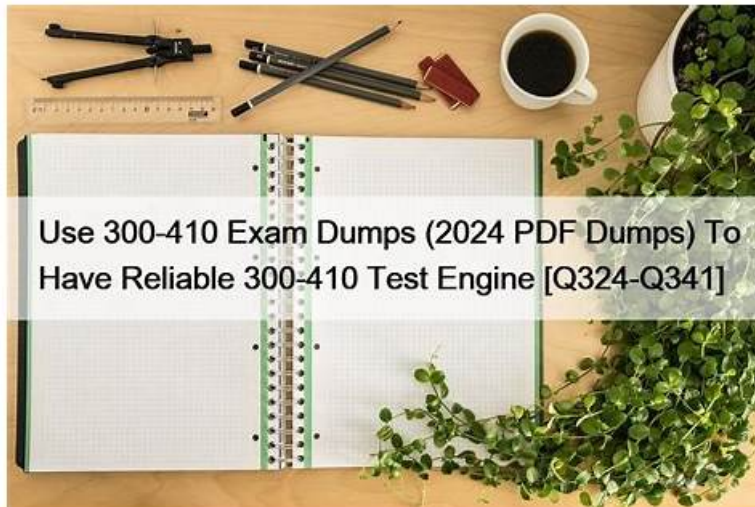


AI-300 Reliable Dumps Book, Reliable AI-300 Test Dumps



IT certification exam is very popular examination in the current society, especially in the IT industry. IT certification test qualification is widely recognized by the international community. Promotion, salary raise and improving your job skills, IT certification exam is your best choice. I believe that you must think so. Then, don't hesitate to take Microsoft AI-300 Exam which is the most popular test in the recent. If you have no idea how to prepare the certification materials for the exam, FreeCram serve you. FreeCram can provide you with everything you need.

There are different ways to achieve the same purpose, and it's determined by what way you choose. A lot of people want to pass Microsoft certification AI-300 exam to let their job and life improve, but people participated in the Microsoft Certification AI-300 Exam all knew that Microsoft certification AI-300 exam is not very simple. In order to pass Microsoft certification AI-300 exam some people spend a lot of valuable time and effort to prepare, but did not succeed.

>> AI-300 Reliable Dumps Book <<

Try a Free Demo and Then Buy Microsoft AI-300 Exam Dumps

The aim of FreeCram is help every candidates getting Microsoft certification easily and quickly. Comparing to attending expensive training institution, AI-300 dumps pdf is more suitable for people who are eager to passing actual test but no time and energy. If you decide to join us, you will receive valid AI-300 learning study materials with real questions and detailed explanations.

Microsoft Operationalizing Machine Learning and Generative AI Solutions Sample Questions (Q68-Q73):

NEW QUESTION # 68

A team deploys a model to a real-time endpoint in Azure Machine Learning. You deploy some updates to the endpoint. The endpoint returns errors after the new deployment is released. You need to restore the service as quickly as possible. What should you do first?

- A. Delete the endpoint and immediately redeploy it.
- B. Change the authentication type to Azure Machine Learning token-based authentication.
- C. Roll back traffic to the previous deployment.
- D. Increase the compute size.

Answer: C

Explanation:

To restore the service as quickly as possible, you can roll back traffic to the previous deployment by updating the traffic allocation

settings of your Azure Machine Learning online endpoint.

Azure Machine Learning managed online endpoints support multiple deployments under a single endpoint, allowing for blue-green deployment strategies where you can shift traffic between versions instantly.

Key Benefits of This Approach

Instant Recovery: Traffic shifting is a routing change and does not require redeploying the previous model's code or environment, making it the fastest recovery method.

No Downtime: Because the previous deployment remains "warm" (active but receiving no traffic), the switch happens without interrupting the service.

Isolation for Debugging: You can keep the failing deployment at 0% traffic to inspect its logs using `az ml online-deployment get-logs` without affecting end users. How to Roll Back Traffic If your new deployment (e.g., "green") is returning errors, you can reallocate 100% of the traffic back to the known stable deployment (e.g., "blue") using the following methods:

* Azure CLI: Use the `az ml online-endpoint update` command to set the traffic percentage:

```
az ml online-endpoint update --name <your-endpoint-name> --traffic "blue=100 green=0"
```

* Azure Machine Learning Studio:

Navigate to Endpoints in the left menu.

Select your specific real-time endpoint.

Go to the Details or Live Traffic tab.

Adjust the traffic percentages so the previous deployment receives 100% and the failing deployment receives 0%.

Select Update or Save to apply the changes immediately

Reference:

<https://learn.microsoft.com/en-us/answers/questions/1275110/azure-ml-v2-yaml-code-for-live-traffic-allocation>

NEW QUESTION # 69

During training, pipelines occasionally fail due to schema mismatch caused by upstream data changes. You need a robust and automated solution that prevents invalid data from reaching training steps. What is the BEST approach?

- A. Ignore schema differences
- B. Add a data validation component in pipeline
- C. Retrain manually when failure occurs
- D. Use larger compute

Answer: B

Explanation:

A data validation component ensures that incoming data matches the expected schema before training begins. This prevents pipeline failures and avoids training on corrupted or incomplete data. Ignoring schema mismatches can introduce silent errors, making debugging difficult and compromising model quality.

NEW QUESTION # 70

You deploy a new model version to a managed online endpoint. You must test it with 10% traffic and automatically roll back if latency or error rate increases beyond threshold. What should you configure?

- A. Manual testing workflow
- B. Batch endpoint validation
- C. Separate endpoint for testing
- D. Traffic splitting with monitoring alerts

Answer: D

Explanation:

Traffic splitting enables controlled rollout of a new model version by directing a percentage of requests to it. Combined with monitoring alerts, it supports automated rollback when performance degrades. Separate endpoints lack built-in traffic management and do not provide seamless or automated rollback capabilities.

NEW QUESTION # 71

A team is experimenting with traditional models for a classification workflow in Azure Machine Learning. The team requires a consistent way to manage assets that are created during experimentation.

You need to ensure that artifacts can be reused and governed across projects.
Which asset should you register?

- A. Pipeline
- B. Environment
- C. Component
- **D. Model**

Answer: D

Explanation:

In an Azure Machine Learning classification workflow, you should register Models.

Registration creates a versioned asset in your workspace or a centralized registry, which is essential for ensuring that artifacts are reusable, governed, and trackable across different projects and environments.

Key Assets for Reuse and Governance

To maintain a consistent and governed workflow, you should focus on registering these specific assets:

Models: The primary artifact. Registering a model allows you to track its lineage (which experiment created it), version it, and deploy it consistently across environments.

Components: These are self-contained pieces of code that perform specific steps in a pipeline (e.g., data cleaning, training).

Registering them allows different teams to reuse the same "traditional" classification logic without rewriting code.

Environments: Encapsulates the software dependencies (Python packages, Docker images) required for your model to run.

Registering these ensures reproducibility across different compute targets.

Data Assets: Registering your training and testing datasets as versioned assets ensures that you can always audit exactly what data was used to train a specific model version.

Reference:

<https://learn.microsoft.com/en-us/azure/machine-learning/concept-azure-machine-learning-v2>

NEW QUESTION # 72

You manage an Azure Machine learning workspace. You develop a machine learning model.

You must deploy the model to use a low-priority VM with a pricing discount.

You need to deploy the model.

Which compute target should you use?

- **A. Azure Machine Learning compute clusters**
- B. Azure Container Instances (ACI)
- C. Local deployment
- D. Azure Kubernetes Service (AKS)

Answer: A

Explanation:

The best compute target for deploying a model using low-priority VMs (or their modern successor, Spot VMs) is an Azure Machine Learning compute cluster.

Best Compute Target: AML Compute Cluster

For low-priority/Spot pricing, you should use an Azure Machine Learning compute cluster configured with the LowPriority tier.

Primary Use Case: This target is specifically recommended for batch deployments. Batch inference is ideal for low-priority VMs because these jobs are asynchronous and can tolerate the interruptions (preemptions) inherent to discounted capacity.

Pricing Advantage: Low-priority VMs offer significant discounts-often up to 80% off standard rates-by utilizing unused Azure capacity.

Automatic Handling: When a node is preempted during a batch job, Azure Machine Learning automatically attempts to replace the lost capacity and re-queues failed tasks to the cluster.

Reference:

<https://learn.microsoft.com/en-us/azure/machine-learning/how-to-use-low-priority-batch>

NEW QUESTION # 73

.....

Are you planning to take the Operationalizing Machine Learning and Generative AI Solutions (AI-300) certification test and don't

www.stes.tyc.edu.tw, godata.co.in, Disposable vapes