Pass Guaranteed Quiz Microsoft - DP-800 - High Hit-Rate Demo Developing AI-Enabled Database Solutions Test

Test4Cram provide you with the comprehensive Microsoft DP-800 Exam information to help you to succeed. Our training materials are the latest study materials which bring by experts. We help you achieve your success. You can get the most detailed and accurate exam questions and answers from us. Our Training Tools are updated in a timely manner in accordance with the changing of Exam Objectives. In fact, the success is not far away, go down along with Test4Cram, then you will come to the road to success.

Microsoft DP-800 Exam Syllabus Topics:

Topic	Details
Topic 1	• Implement AI capabilities in database solutions: This domain covers designing and managing external AI models and embeddings, implementing full-text, semantic vector, and hybrid search strategies, and building retrieval-augmented generation (RAG) solutions that connect database outputs with language models.
Topic 2	• Design and develop database solutions: This domain covers designing and building database objects such as tables, views, functions, stored procedures, and triggers, along with writing advanced T-SQL code and leveraging AI-assisted tools like GitHub Copilot and MCP for SQL development.
Topic 3	• Secure, optimize, and deploy database solutions: This domain focuses on implementing data security measures like encryption, masking, and row-level security, optimizing query performance, managing CI • CD pipelines using SQL Database Projects, and integrating SQL solutions with Azure services including Data API builder and monitoring tools.

>> Demo DP-800 Test <<

DP-800 Valid Exam Sample, DP-800 Accurate Answers

The Microsoft DP-800 certification exam is one of the hottest and career-oriented Developing AI-Enabled Database Solutions (DP-800) exams. With the Developing AI-Enabled Database Solutions (DP-800) exam you can validate your skills and upgrade your knowledge level. By doing this you can learn new in-demand skills and gain multiple career opportunities. To do this you just need to enroll in the Microsoft DP-800 Certification Exam and put all your efforts to pass this important Microsoft DP-800 Exam Questions. However, you should keep in mind that to get success in the Developing AI-Enabled Database Solutions (DP-800) exam is not an easy task.

Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q34-Q39):

NEW QUESTION # 34
You have an Azure SQL database that contains a table named dbo.Products. dbo.Products contains three columns named Embedding, Category, and Price. The Embedding column is defined as VECTOR(1536).
You use AI_GENERATE_EMBEDDINGS and VECTOR_SEARCH to support semantic search and apply additional filters on two columns named Category and Price.
You plan to change the embedding model from text-embedding-ada-002 to text-embedding-3- small. Existing rows already contain embeddings in the Embedding column.
You need to implement the model change. Applications must be able to use VECTOR_SEARCH without runtime errors.
What should you do first?

- A. Normalize the vector lengths before storing new embeddings.
- B. Create a vector index on dbo.Products.Embedding.
- C. Regenerate embeddings for the existing rows.
- D. Convert the Embedding column to nvarchar(max).

Answer: B

Explanation:
To ensure your applications can transition models without runtime errors while using VECTOR_SEARCH, you must first define a Vector Index that explicitly identifies the dimensions and distance metric.
Since you are moving from text-embedding-ada-002 to text-embedding-3-small, both models default to 1536 dimensions, which matches your existing column definition. To create the index as the first step, use the following SQL:
CREATE VECTOR INDEX idx_embedding ON YourTableName (Embedding)
WITH (DISTANCE_METRIC = 'COSINE');
Use code with caution.
Why this works:
Schema Consistency: Because both models use 1536 dimensions, you don't need to alter the VECTOR(1536) column type immediately.
Search Stability: Creating the index allows the engine to optimize the VECTOR_SEARCH function. As long as the incoming query vector (generated by the app) matches the dimensions of the stored vectors, the search will execute without a runtime dimension mismatch error.
Reference:
https://docs.couchbase.com/cloud/n1ql/n1ql-language-reference/vectorfun.html

NEW QUESTION # 35
Hotspot Question
You have an SDK-style SQL database project named MyDatabaseProject.sqlprojstored in a private GitHub repository. The repository contains the following GitHub Actions workflow.

```
n:
 push:
   branches:
     - main
 workflow_dispatch:
obs:
 build-and-deploy:
   runs-on: ubuntu-latest
   permissions:
     id-token: write
     contents: read
   steps:
     - uses: actions/checkout@v4
     - name: Build SQL project
       run: dotnet build MyDatabaseProject.sqlproj
     - name: Publish DACPAC
       uses: azure/sql-action@v2
       with:
         action: publish
         path: bin/Debug/MyDatabaseProject.dacpac
         connection-string: ${{ secrets.AZURE_SQL_CONNECTION_STRING }}
```

The repository contains the AZURE_SQL_CONNECTION_STRING secrets.
The target is an Azure SQL database that allows access to Azure services and is configured to support mixed authentication.
The workflow runs successfully.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
The workflow relies on a Microsoft Entra workload identity to access the target Azure SQL database.	○	○
Changing the `Build SQL project` step to run: `dotnet build MyDatabaseProject.sqlproj -c Release` will result in a successful deployment.	○	○
The workflow can be triggered manually, without making changes to the repository.	○	○

Answer:

Explanation:

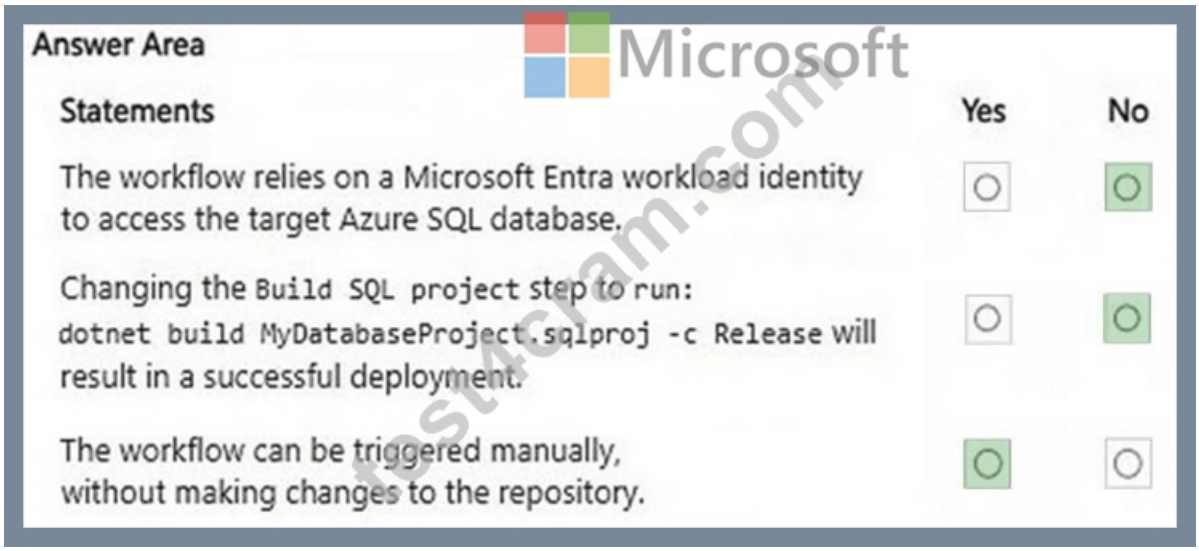

NEW QUESTION # 36
What challenge does AI introduce to SQL systems?

- A. Data privacy and hallucination risks
- B. Increased cost only
- C. Reduced performance always
- D. No challenges

Answer: A

Explanation:
AI may generate incorrect answers and requires strong governance and validation.

NEW QUESTION # 37
Drag and Drop Question
You have an Azure SQL database that contains a table named Sales.Orders. Sales.Orders contains the following columns.

Column	Data type
OrderId	int
CustomerId	int
OrderDate	datetime2
TotalAmount	decimal(18,2)

Reporting queries frequently repeat logic to calculate the number of days since an order was placed.
You need to create a scalar user-defined function (UDF) that returns the number of days between an input value of @OrderDate and the current date and time.
How should you complete the Transact-SQL code? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

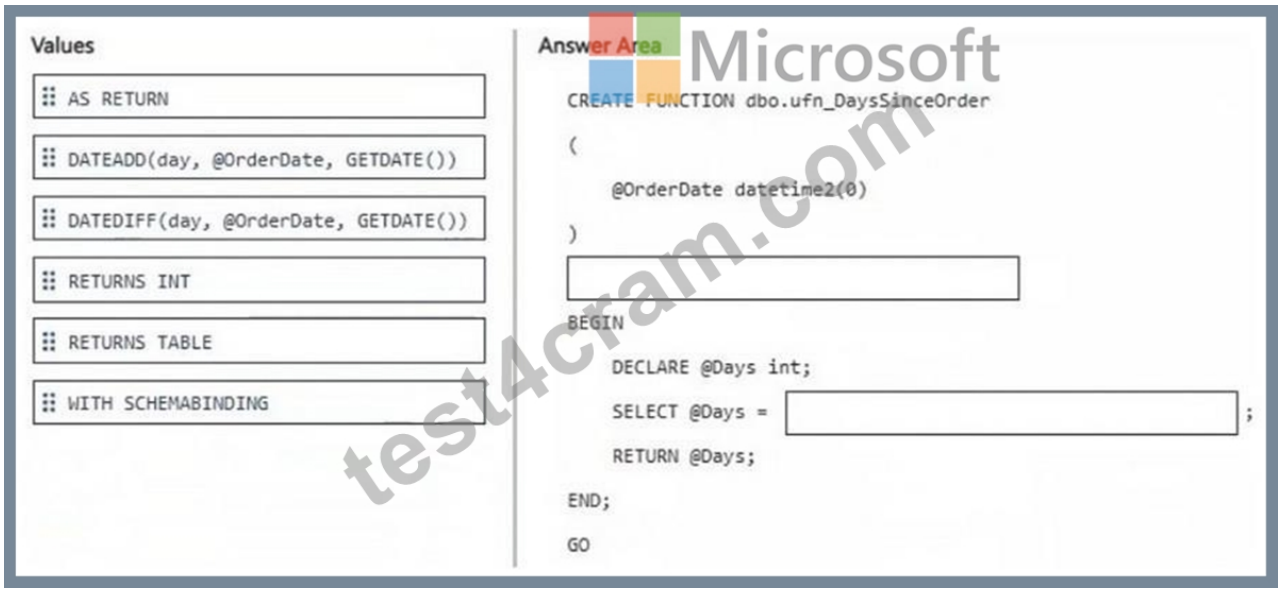

Answer:

Explanation:

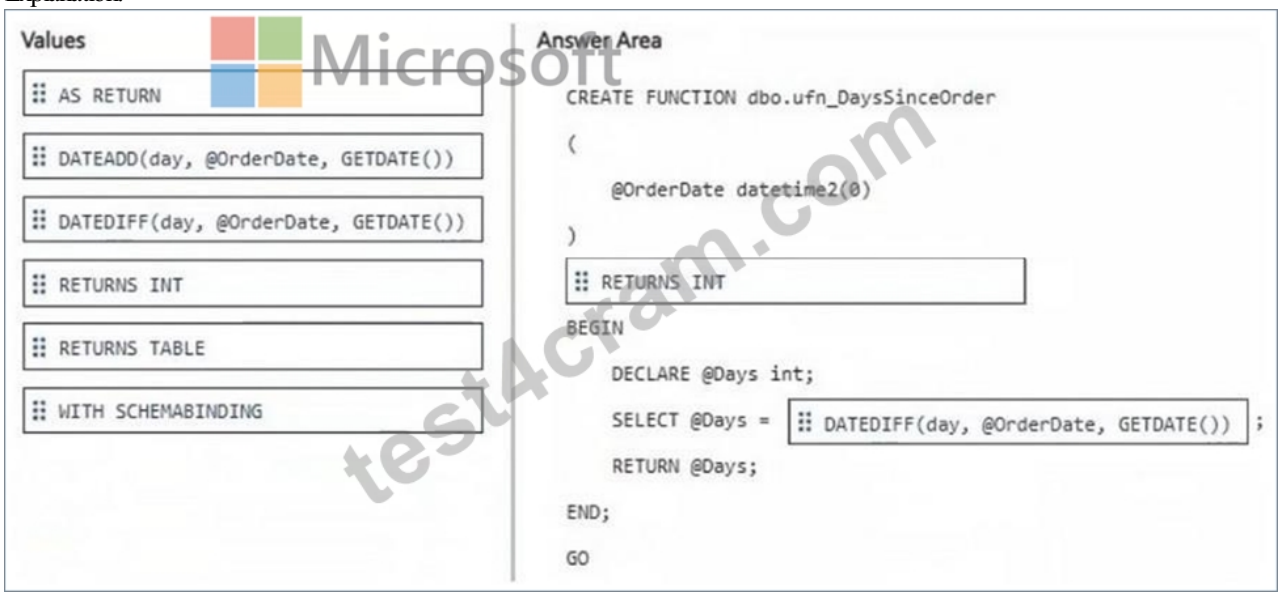

NEW QUESTION # 38

You have a SQL database in Microsoft Fabric that contains a table named dbo.Orders, dbo.Orders has a clustered index, contains three years of data, and is partitioned by a column named OrderDate by month.

You need to remove all the rows for the oldest month. The solution must minimize the impact on other queries that access the data in dbo.orders.

Solution: Run the following Transact-SQL statement.

DELETE FROM dbo.Orders

WHERE OrderDate < DATEADD(nonth, -36, SYSUTCDATETIME());

Does this meet the goal?

- A. No
- B. Yes

Answer: A

Explanation:

This does not meet the goal. A row-by-row DELETE against the oldest month is not the lowest-impact way to purge data from a monthly partitioned table. Microsoft's partitioning guidance specifically says partitioning lets you perform maintenance and retention operations more efficiently by targeting just the relevant partition, including the ability to truncate data in a single partition .

The proposed statement:
DELETE FROM dbo.Orders
WHERE OrderDate < DATEADD(month, -36, SYSUTCDATETIME());
would log row deletions and can hold locks longer, creating more overhead for other queries than a partition- level maintenance operation. Since the table is already partitioned by month , the expected low-impact approach is to operate on the oldest partition directly, not issue a broad delete predicate over rows. Microsoft explicitly highlights partition-targeted truncation as a faster, more efficient retention operation than working against the whole table or rowset.

NEW QUESTION # 39
.....