


Real Microsoft DP-800 Question | New DP-800 Exam Guide



P.S. Free & New DP-800 dumps are available on Google Drive shared by CramPDF: <https://drive.google.com/open?id=1yU4V31PoutiH2tKp-8AfYCBcOFBO5QYI>

As we all know, it is not easy to get promotion. For the first thing, you must be good at finishing your work excellently. At the same time, you must accumulate much experience and knowledge. If you urgently want to stand out in your company, our DP-800 exam guide can help you realize your aims in the shortest time. For not only that our DP-800 Study Materials can help you know more knowledge on the subject and our DP-800 practice engine can help you get your according certification.

Your final purpose is to get the DP-800 certificate. So it is important to choose good study materials. In fact, our aim is the same with you. Our DP-800 study materials have strong strengths to help you pass the exam. Maybe you still have doubts about our DP-800 exam materials. We have statistics to prove the truth. First of all, our sales volumes are the highest in the market. You can browse our official websites to check our sales volumes. At the same time, many people pass the exam for the first time under the guidance of our DP-800 Practice Exam.

>> Real Microsoft DP-800 Question <<

New DP-800 Exam Guide | Exam DP-800 Passing Score

To fit in this amazing and highly accepted exam, you must prepare for it with high-rank practice materials like our Developing AI-Enabled Database Solutions DP-800 study materials. Our DP-800 exam questions are the Best choice in terms of time and money. If you are a beginner, start with the learning guide of DP-800 Practice Engine and our products will correct your learning problems with the help of the Microsoft DP-800 training braindumps.

Microsoft DP-800 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">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 CICD pipelines using SQL Database Projects, and integrating SQL solutions with Azure services including Data API builder and monitoring tools.
Topic 2	<ul style="list-style-type: none">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 3

- 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.

Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q110-Q115):

NEW QUESTION # 110

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an SDK-style SQL database project stored in a Git repository. The project targets an Azure SQL database.

The CI build fails with unresolved reference errors when the project references system objects.

You need to update the SQL database project to ensure that dotnet build validates successfully by including the correct system objects in the database model for Azure SQL Database.

Solution: Build the project by running `dotnet build -bl -flp:v=diag`.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Correct:

* Add the `Microsoft.SqlServer.Dacpac.Azure.Master` NuGet package to the project.

To resolve system reference errors in an SDK-style SQL project targeting Azure SQL Database, you need to add a reference to the `Microsoft.SqlServer.Dacpac.Azure.Master` NuGet package.

In your `.sqlproj` file, include the following item group:

```
<ItemGroup>
<PackageReference Include="Microsoft.SqlServer.Dacpac.Azure.Master" Version="1.60.0" />
</ItemGroup>
```

Why this works:

System Objects: Standard SDK-style projects don't automatically include system views (like `sys.database_principals` or `sys.dm_db_resource_stats`). This package provides the necessary metadata for the compiler.

Azure Specifics: It includes Azure-only system objects that aren't present in the standard master database dacpac used for on-premises SQL Server.

CI/CD Friendly: Since it is a NuGet package, the `dotnet build` command will automatically restore it during the CI process without requiring manual file paths or local installations of Visual Studio.

Incorrect:

* Add an artifact reference to the Azure SQL Database master.dacpac file.

* Add the `Microsoft.SqlServer.Dacpac.Master` NuGet package to the project.

* Build the project by running `dotnet build -bl -flp:v=diag`.

Reference:

<https://learn.microsoft.com/en-us/sql/tools/sql-database-projects/concepts/system-objects>

NEW QUESTION # 111

Hotspot Question

You have a database named `db1`. The schema is stored in a Git repository as an SDK-style SQL database project. The repository contains the following GitHub Action workflow.

```
name: Database CI/CD
```

```
on:
```

```
  push:
```

```
    branches:
```

```
      - main
```

```
  pull_request:
```

```
    branches:
```

```
      - main
```

```
jobs:
```

```
  build-and-deploy:
```

```
    runs-on: ubuntu-latest
```

```
    steps:
```

```
      - name: Checkout code
```

```
        uses: actions/checkout@v3
```

```
      - name: Setup .NET
```

```
        uses: actions/setup-dotnet@v3
```

```
        with:
```

```
          dotnet-version: '8.x'
```

```
      - name: Build
```

```
        run: dotnet build db1.sqlproj --configuration Release
```

```
      - name: Deploy
```

```
        run: |
```

```
          sqlpackage /Action:Publish \
```

```
            /SourceFile:./bin/Release/db1.dacpac \
```

```
            /TargetConnectionString:"${{ secrets.Target_Connection_String }}"
```

```
  unit-tests:
```

```
    runs-on: ubuntu-latest
```

```
    needs: build-and-deploy
```

```
    if: github.ref == 'refs/heads/main'
```

```
    steps:
```

```
      - name: Checkout code
```

```
        uses: actions/checkout@v3
```

```
      - name: Setup .NET
```

```
        uses: actions/setup-dotnet@v3
```

```
        with:
```

```
          dotnet-version: '8.x'
```

```
      - name: Run unit tests
```

```
        run: dotnet test UnitTests.csproj --configuration Release
```



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
Unit tests run automatically whenever changes are pushed to main.	<input type="radio"/>	<input type="radio"/>
Schema validation occurs during the build step.	<input type="radio"/>	<input type="radio"/>
Schema validation occurs during the deploy step.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

Answer Area		
Statements	Yes	No
Unit tests run automatically whenever changes are pushed to main.	<input checked="" type="radio"/>	<input type="radio"/>
Schema validation occurs during the build step.	<input checked="" type="radio"/>	<input type="radio"/>
Schema validation occurs during the deploy step.	<input type="radio"/>	<input checked="" type="radio"/>

NEW QUESTION # 112

Hotspot Question

You have an Azure SQL database that contains a table named stores. stores contains a column named description and a vector column named embedding.

You need to implement a hybrid search query that meets the following requirements:

- Uses full-text search on description for the keyword portion
- Returns the top 20 results based on a combined score that uses a weighted formula of 60% vector distance and 40% full-text rank

How should you configure the query components? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Semantic query operator/function:

VECTOR_DISTANCE and order by distance ascending
VECTOR_SEARCH with METRIC cosine and TOP_N
VECTORPROPERTY to calculate the similarity between two vectors

Keyword retrieval operator/function:

CONTAINSTABLE on description and return ranked matches
FREETEXT TABLE on description for keyword scoring
JSON_VALUE extraction from description for keyword scoring

Final ranking expression:

order by (distance * 0.6) + ((1.0 - RANK/1000.0) * 0.4)
order by (distance * 0.6) + (RANK * 0.4)
order by (distance + RANK), and then apply TOP 20



Microsoft

Answer:

Explanation:

Semantic query operator/function:

VECTOR_DISTANCE and order by distance ascending
VECTOR_SEARCH with METRIC cosine and TOP_N
VECTORPROPERTY to calculate the similarity between two vectors



Microsoft

Keyword retrieval operator/function:

CONTAINSTABLE on description and return ranked matches
FREETEXT TABLE on description for keyword scoring
JSON_VALUE extraction from description for keyword scoring

Final ranking expression:

order by (distance * 0.6) + ((1.0 - RANK/1000.0) * 0.4)
order by (distance * 0.6) + (RANK * 0.4)
order by (distance + RANK), and then apply TOP 20

NEW QUESTION # 113

Which SQL feature is commonly used to integrate AI-generated insights?

- A. Stored Procedures
- B. Triggers
- C. Views
- D. External APIs

Answer: D

Explanation:

AI models (like OpenAI) are typically accessed via APIs, which SQL apps call externally.

NEW QUESTION # 114

You need to create a table in the database to store the telemetry data. You have the following Transact-SQL code.

```
TelemetryId BIGINT IDENTITY(1,1) NOT NULL,  
VehicleId NVARCHAR(50) NOT NULL,  
TelemetryTimeUtc DATETIME2(3) NOT NULL,  
BatteryPercent TINYINT NULL,  
SpeedKmh SMALLINT NULL,  
LocationJson JSON NULL,  
ErrorCodesJson JSON NULL,  
RawPayload NVARCHAR(MAX) NULL,  
SysStartTime DATETIME2(7) GENERATED ALWAYS AS ROW START NOT NULL,  
SysEndTime DATETIME2(7) GENERATED ALWAYS AS ROW END NOT NULL,  
PERIOD FOR SYSTEM_TIME (SysStartTime, SysEndTime),  
CONSTRAINT PK_VehicleTelemetry PRIMARY KEY CLUSTERED (TelemetryId)  
  
WITH  
  
SYSTEM_VERSIONING = ON  
(  
    HISTORY_TABLE = dbo.VehicleTelemetryHistory  
)  
;  
;  
;  
  
CREATE INDEX IX_VehicleTelemetry_Time ON dbo.VehicleTelemetry (TelemetryTimeUtc);  
CREATE JSON INDEX JI_VehicleTelemetry_Location ON dbo.VehicleTelemetry (LocationJson) ((  
    '$.location.heading',  
    '$.location.longitude',  
    '$.location.accuracy'
```

Answer Area

Statements	Yes	No
The code meets the database performance requirements for partitioning.	<input type="radio"/>	<input type="radio"/>
The code meets the database performance requirements for JSON property querying.	<input type="radio"/>	<input type="radio"/>
Queries that filter on \$.location.heading will use the JI_VehicleTelemetry_Location index.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

Answer Area

Statements	Yes	No
The code meets the database performance requirements for partitioning.	<input type="radio"/>	<input checked="" type="radio"/>
The code meets the database performance requirements for JSON property querying.	<input checked="" type="radio"/>	<input type="radio"/>
Queries that filter on \$.location.heading will use the JI_VehicleTelemetry_Location index.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

Answer Area

Statements	Yes	No
The code meets the database performance requirements for partitioning.	<input type="radio"/>	<input checked="" type="radio"/>
The code meets the database performance requirements for JSON property querying.	<input checked="" type="radio"/>	<input type="radio"/>
Queries that filter on \$.location.heading will use the JI_VehicleTelemetry_Location index.	<input type="radio"/>	<input checked="" type="radio"/>

The first statement is No . The requirement says telemetry data must be stored in a partitioned table to provide predictable performance for ingestion and retention operations. However, the shown CREATE TABLE statement does not define a partition function or partition scheme, and the table is created with a regular clustered primary key on TelemetryId. Microsoft's partitioning guidance states that creating a partitioned table requires a partition function , a partition scheme , and creating the table or index on that partition scheme using a partitioning column. None of that appears in the code, so the table is not partitioned. The second statement is Yes . The code creates a JSON index named JI_VehicleTelemetry_Location on LocationJson for these

