


Free PDF Quiz 2026 DP-800: Developing AI-Enabled Database Solutions—Efficient Valid Test Papers



P.S. Free 2026 Microsoft DP-800 dumps are available on Google Drive shared by Exam4Docs: <https://drive.google.com/open?id=1CPhe33igiOHIzwYsyZiE6pOSaTOa4UWQ>

It is quite convenient to study with our DP-800 study materials. If you are used to study with paper-based materials you can choose the PDF version which is convenient for you to print. If you would like to get the mock test before the real DP-800 exam you can choose the software version, and if you want to study in anywhere at any time then our online APP version is your best choice since you can download it in any electronic devices. And the price of our DP-800 learning guide is favorable.

Microsoft DP-800 Exam Syllabus Topics:

Section	Objectives
Integrate AI capabilities with database systems	- Use Azure AI services with database workloads - Implement AI-assisted data processing
Develop and manage database solutions	- Optimize performance and scalability - Ensure security and compliance of data solutions
Design and implement data solutions	- Design database solutions using Azure data services - Implement data storage and data processing solutions
Monitor, troubleshoot, and maintain solutions	- Troubleshooting data pipeline issues - Monitoring database health and performance

DP-800 Valid Test Papers - 100% Marvelous Questions Pool

Exam4Docs is professional and is built for nearly all IT certification examinations. It not only ensures the quality, best service, but also the cheap price. Having Exam4Docs, you will not worry about DP-800 certification exams and answers. Moreover, Exam4Docs can provide DP-800 Latest Dumps demo and DP-800 study guide for you, which will help you pass DP-800 exam in a short time and let you be close to your dream to become an elite.

Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q81-Q86):

NEW QUESTION # 81

You have a SQL database in Microsoft Fabric that contains a table named WebSite.Logs. WebSite.Logs stores application telemetry data. Website.Logs contains a nvarchar(max) column named log that stores JSON documents. You have a daily report that filters by the \$.severity JSON property and returns LogId, LogDateTime, and log. The report frequently causes full table scans.

You need to modify Website.Logs to support efficient filtering by \$.severity and avoid key lookups for the columns returned by the report.

How should you complete the Transact-SQL code to avoid full table scans? 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:

Explanation:

The correct way to avoid full table scans here is to add a computed column that extracts the JSON scalar property with `JSON_VALUE`, and then create a nonclustered index on that computed column with the report's returned columns in the `INCLUDE` list. Microsoft's JSON indexing guidance specifically recommends creating a computed column that exposes the JSON property you filter on, using the same expression as in the query, and then indexing that computed column.

So the computed column must be:

`AS JSON_VALUE([log], '$.severity') PERSISTED`

This is correct because `$.severity` is a scalar JSON value, so `JSON_VALUE` is the proper function.

JSON_QUERY would be for extracting an object or array, not a scalar property. Microsoft also notes that persisted computed columns can improve access speed for JSON-derived values.

The index should then include:

```
INCLUDE (LogId, LogDateTime, [log])
```

That is the right covering strategy because the report filters by severity but returns LogId, LogDateTime, and log. Microsoft's guidance on included columns explains that nonkey included columns let a nonclustered index cover more queries and reduce extra lookups to the base table.

So the completed code is:

```
ALTER TABLE WebSite.Logs
ADD severity AS JSON_VALUE([log], '$.severity' ) PERSISTED;
GO
CREATE INDEX ix_severity
ON WebSite.Logs(severity)
INCLUDE (LogId, LogDateTime, [log]);
GO
```

NEW QUESTION # 82

You have a Microsoft SQL Server 2025 database that contains a table named products. products contains two columns named description and embedding. embedding is generated from description.

You have an application named App1. App1 has a feature that uses the following query.

```
VECTOR_DISTANCE('cosine', @query_vector, embedding)
```

Users report that the feature is slow during peak usage times.

You discover that during peak usage times, the semantic search latency increases and CPU utilization spikes.

You need to reduce latency and CPU utilization related to the App1 feature.

What should you do?

- A. Use CONTAINSTABLE against description only and remove the vector search.
- **B. Create a vector index on products.embedding and change the query to use VECTOR_SEARCH and METRIC = 'cosine'.**
- C. Add a nonclustered index on products.embedding.
- D. Change embedding to VARBINARY(8000).
- E. Create a vector index on products.embedding and change the query to use VECTOR_NORMALIZE and NORM_TYPE = 'norml'.

Answer: B

Explanation:

To optimize your Microsoft SQL Server 2025 environment, you should transition from a brute-force distance calculation to an Approximate Nearest Neighbor (ANN) search using a DiskANN-based vector index.

The Solution: Vector Indexing

The VECTOR_DISTANCE function performs a full table scan, calculating the distance for every single row. During peak times, this consumes massive CPU cycles. A vector index allows SQL Server to navigate a pre-built graph to find matches, reducing complexity from O(n) to O(log n).

Implementation Steps

1. Create the Vector Index

You must specify the distance metric during index creation. For cosine similarity, use COSINE.

```
CREATE VECTOR INDEX idx_product_embedding
ON products (embedding)
WITH (METRIC = 'COSINE');
```

2. Update the Search Query

Replace the WHERE or ORDER BY clause that uses VECTOR_DISTANCE with the VECTOR_SEARCH function.

```
SELECT TOP (10) description, embedding
```

```
FROM products
```

```
WHERE VECTOR_SEARCH(
embedding,
```

```
@query_vector,
```

```
'COSINE',
```

```
10 -- Number of neighbors to return
```

```
);
```

Why This Solves Your Problem

CPU Relief: Avoids calculating distances for millions of rows.

Lower Latency: DiskANN is designed for sub-second retrieval even on large datasets.

Peak Performance: By offloading the math to a specialized index structure, the engine remains responsive for other transactional queries.

Reference:

<https://learn.microsoft.com/en-us/azure/postgresql/extensions/how-to-use-pgdiskann>

NEW QUESTION # 83

What is the primary purpose of Azure AI in SQL development?

- A. Enhance querying and analytics with AI capabilities
- B. Automate infrastructure deployment
- C. Eliminate the need for databases
- D. Replace SQL Server

Answer: A

Explanation:

Azure AI integrates with SQL to enable intelligent insights like natural language queries, anomaly detection, and predictive analytics.

NEW QUESTION # 84

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

Does this meet the goal?

- A. No
- B. Yes

Answer: A

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.

Reference:

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

NEW QUESTION # 85

You have an Azure SQL database that contains a table named `dbo.ManualChunks`. `dbo.ManualChunks` contains product manuals.

A retrieval query already returns the top five matching chunks as `nvarchar (max)text`.

You need to call an Azure OpenAI REST endpoint for chat completions. The request body must include both the user question and the retrieved chunks.

You write the following Transact-SQL code.

```
01 CREATE DATABASE SCOPED CREDENTIAL AzureOpenAIHeaders
02 WITH IDENTITY = 'HTTPEndpointHeaders',
03     SECRET = N'{"api-key":"<YOUR_AZURE_OPENAI_API_KEY>"}';
04 GO
05 CREATE OR ALTER PROCEDURE dbo.AskManuals
06     . . .
07     SELECT @chunks =
08     (
09         SELECT TOP (5)
10             mc.ChunkText AS [text]
11         FROM dbo.ManualChunks AS mc
12         ORDER BY mc.Score DESC
13         FOR JSON PATH
14     );
15     SET @payload =
16     (
17         SELECT
18             'system' AS [messages[0].role],
19             'Use only the provided manual chunks.' AS [messages[0].content],
20             'user' AS [messages[1].role],
21             CONCAT(@question, CHAR(10), JSON_QUERY(@chunks)) AS [messages[1].content]
22
23     );
24
25     EXEC @retval =
26     . . .
27 END;
28 GO
```

What should you insert at line 22?

- A. FOR XML AUTO, TYPE, XMLSCHEMA
- **B. FOR JSON PATH, WITHOUT_ARRAY_WRAPPER**
- C. FOR JSON AUTO, INCLUDE_NULL_VALUES
- D. FOR XML PATH, INCLUDE_NULL_VALUES

Answer: B

Explanation:

To construct a payload for the Azure OpenAI REST API within Azure SQL, you can use a Common Table Expression (CTE) to structure the messages and then collapse them into a single JSON object.

T-SQL Payload Generation

Key Implementation Details

Message Hierarchy: Azure OpenAI requires a messages array. By nesting a subquery with FOR JSON PATH inside the main SELECT, SQL Server creates the required JSON array of objects.

JSON Formatting: Using FOR JSON PATH, WITHOUT_ARRAY_WRAPPER on the outer selection ensures the root of your

@payload is a clean JSON object {...} rather than a single- element array [{...}].

Context Injection: It is standard practice to prepend the retrieved manual chunks to the system message or a dedicated user message to provide the model with the necessary grounding data.

Assuming you have your retrieved chunks concatenated into a single @Context variable and the user's input in @UserQuestion, use the following logic:

```
DECLARE @UserQuestion NVARCHAR(MAX) = 'How do I reset the device?';
DECLARE @Context NVARCHAR(MAX) = 'Manual Chunk 1... Manual Chunk 2...'; -- Your top 5 chunks
DECLARE @payload NVARCHAR(MAX);
-- Define the message structure for Chat Completions
WITH ChatMessages AS (
SELECT 'system' AS [role],
'You are a helpful assistant. Use these manual excerpts: ' + @Context AS [content] UNION ALL
SELECT 'user' AS [role],
@UserQuestion AS [content]
)
-- Generate the final JSON body
SELECT @payload = (
SELECT
(SELECT [role], [content] FROM ChatMessages FOR JSON PATH) AS [messages],
0.7 AS [temperature],
800 AS [max_tokens]
FOR JSON PATH, WITHOUT_ARRAY_WRAPPER
);
-- Resulting @payload is now ready for sp_invoke_external_rest_endpoint PRINT @payload; Reference:
https://devblogs.microsoft.com/azure-sql/using-openai-rest-endpoints-with-azure-sql-database
```

NEW QUESTION # 86

.....

We know that the standard for most workers become higher and higher; so we also set higher goal on our DP-800 guide questions. Our training materials put customers' interests in front of other points, committing us to the advanced DP-800 learning materials all along. Until now, we have simplified the most complicated DP-800 Guide questions and designed a straightforward operation system, with the natural and seamless user interfaces of DP-800 exam question grown to be more fluent, we assure that our practice materials provide you a total ease of use.

DP-800 Customized Lab Simulation: <https://www.exam4docs.com/DP-800-study-questions.html>

- Quiz 2026 DP-800: Useful Developing AI-Enabled Database Solutions Valid Test Papers ☐ Open ➡ www.prep4away.com ☐ and search for > DP-800 < to download exam materials for free ☐ DP-800 Reliable Exam Tutorial
- Microsoft - High Pass-Rate DP-800 Valid Test Papers ☐ Search for ☐ DP-800 ☐ and obtain a free download on ➡ www.pdfvce.com ☐ ☐ DP-800 Latest Braindumps Ebook
- Latest Braindumps DP-800 Book ☐ DP-800 Labs ☐ DP-800 Labs ☐ Simply search for ☐ DP-800 ☐ for free download on ➡ www.troytecdumps.com ☐ ☐ DP-800 Related Certifications
- Latest Braindumps DP-800 Book ☐ DP-800 Reliable Exam Tutorial ☐ Reliable DP-800 Braindumps Ppt ☐ Search for ☐ DP-800 ☐ and download exam materials for free through (www.pdfvce.com) ☐ Reliable DP-800 Braindumps Ppt
- Perfect DP-800 Valid Test Papers - Leading Offer in Qualification Exams - Useful DP-800 Customized Lab Simulation ☐ Go to website ➡ www.torrentvce.com ☐ open and search for ➡ DP-800 ☐ to download for free ☐ Exam DP-800 Outline
- DP-800 Certification Materials ☐ DP-800 Reliable Study Notes ➡ DP-800 Certification Materials ☐ Search for ➡ DP-800 ☐ and download exam materials for free through ✓ www.pdfvce.com ☐ ✓ ☐ ☐ Reliable DP-800 Braindumps Ppt
- DP-800 Latest Demo ☐ Discount DP-800 Code ☐ DP-800 Reliable Test Topics ☐ Easily obtain free download of 「 DP-800 」 by searching on ☼ www.examcollectionpass.com ☐ ☼ ☐ ♥ Reliable DP-800 Test Pattern
- Actual DP-800 : Developing AI-Enabled Database Solutions ExamDumps Questions Is Easy to Understand - Pdfvce ☐ ➡ www.pdfvce.com ☐ is best website to obtain ✓ DP-800 ☐ ✓ ☐ for free download ☐ DP-800 Latest Braindumps Ebook
- DP-800 Labs ☂ DP-800 Certification Materials ☐ Valid DP-800 Test Objectives ☂ Search for ☐ DP-800 ☐ and obtain a free download on ☐ www.testkingpass.com ☐ ☐ Discount DP-800 Code
- Perfect DP-800 Valid Test Papers - Leading Offer in Qualification Exams - Useful DP-800 Customized Lab Simulation ☐ Go to website ➡ www.pdfvce.com ☐ ☐ open and search for ➡ DP-800 ☐ ☐ to download for free ☐ DP-800 Interactive Questions

- 2026 Latest Exam4Docs DP-800 PDF Dumps and DP-800 Exam Engine Free Share: <https://drive.google.com/open?id=1CPhe33igiOHIzwYsyZiE6pOSaTOa4UWQ>

2026 Latest Exam4Docs DP-800 PDF Dumps and DP-800 Exam Engine Free Share: <https://drive.google.com/open?id=1CPhe33igiOHIzwYsyZiE6pOSaTOa4UWQ>