

Guaranteed DP-800 Questions Answers, DP-800 Valid Exam Objectives



Our DP-800 quiz torrent boost 3 versions and they include PDF version, PC version, App online version. Different version boosts different functions and using method. For example, the PDF version is convenient for the download and printing our DP-800 exam torrent and is easy and suitable for browsing learning. It can be printed on the papers which are convenient for you to take notes and learn at any time and place. You can practice DP-800 Quiz prep repeatedly and there are no limits for the amount of the persons and times. And the PC version of DP-800 quiz torrent can stimulate the real exam's scenarios, is stalled on the Windows operating system and runs on the Java environment. You can use it any time to test your own Exam stimulation tests scores and whether you have mastered our DP-800 exam torrent.

Microsoft DP-800 Exam Syllabus Topics:

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

>> **Guaranteed DP-800 Questions Answers** <<

DP-800 Valid Exam Objectives - DP-800 New Study Questions

Having a good command of professional knowledge in this line, they devised our high quality and high effective DP-800 study materials by unremitting effort and studious research. They are meritorious and unsuspecting experts with professional background. By concluding quintessential points into DP-800 Preparation engine, you can pass the exam with the least time while huge progress. And our pass rate of the DP-800 exam questions is high as 98% to 100%.

Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q39-Q44):

NEW QUESTION # 39

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.

What should you insert at line 22?

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

Answer: B

Explanation:

The correct insertion at line 22 is `FOR JSON PATH, WITHOUT_ARRAY_WRAPPER`.

The request body for the Azure OpenAI chat completions call must be a single JSON object containing the messages array with both the system/user content and the retrieved chunks. Microsoft documents that `FOR JSON PATH` is the preferred way to shape JSON output, especially when you want precise control over nested property names like `messages[0].role` and `messages[1].content`.

The key detail is `WITHOUT_ARRAY_WRAPPER`. By default, `FOR JSON` returns results enclosed in square brackets as a JSON array. Microsoft documents that `WITHOUT_ARRAY_WRAPPER` removes those brackets so a single JSON object is produced instead. That is exactly what is needed here for `@payload`, because the stored procedure is building one request body, not an array of request bodies.

NEW QUESTION # 40

Your company has an ecommerce catalog in a Microsoft SQL Server 2022 database named `SalesDB`. `SalesDB` contains a table named `products`. `products` contains the following columns:

- * `product_id` (int)
- * `product_name` (nvarchar(200))
- * `description` (nvarchar(max))
- * `category` (nvarchar(50))
- * `brand` (nvarchar(W))
- * `price` (decimal)
- * `sku` (nvarchar(40))

The description fields are updated daily by a content pipeline, and price can change multiple times per day.

You want customers to be able to submit natural language queries and apply structured filters for brand and price. You plan to store embeddings in a new `VECTOR(1536)` column and use `VECTOR_SEARCH(...)`

METRIC=' cosine ' ...).

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:

Explanation:

Explanation:

The first statement is Yes . Embeddings are used to represent the semantic meaning of content, and vector search is for conceptually similar matches over that content. Here, the semantically meaningful fields are product_name, category, and description. Using those together supports natural-language search, while brand and price can be handled as structured filters outside the embedding itself. This is an inference from Microsoft's guidance that vector search works over embeddings representing content meaning, while filters remain part of the nonvector query pipeline.

The second statement is No . price changes multiple times per day and is a structured numeric attribute, not stable semantic content. Since the requirement already says customers can apply structured filters for brand and price , price does not need to be embedded into the text. Embedding volatile numeric values would also make embeddings stale faster without improving the semantic-search objective. This is again an inference grounded in Microsoft's distinction between vector similarity over content and filtering/sorting over nonvector fields.

The third statement is Yes . In SQL Server's vector type, the default underlying base type is float32 unless float16 is specified explicitly.

NEW QUESTION # 41

You are creating a table that will store customer profiles.

You have the following Transact-SQL code.

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:

Explanation:

Explanation:

* The schema meets the security requirements for PII data. # Yes

* Administrators of the Azure SQL server can see all the rows in dbo.CustomerProfiles when they use an application. # No

* The masking rules will apply even when row-level security (RLS) filters out rows. # No The first statement is Yes because the design combines two relevant SQL security controls for personally identifiable information: Dynamic Data Masking (DDM) on sensitive columns such as FullName, EmailAddress, and PhoneNumber, and Row-Level Security (RLS) to restrict which rows a user can access based on RegionCode. Microsoft documents that DDM limits sensitive data exposure for nonprivileged users , while RLS restricts row access according to the user executing the query. Together, these are valid and appropriate controls for protecting PII in Azure SQL Database.

The second statement is No . Administrative users can view unmasked data because administrative roles effectively have CONTROL, which includes UNMASK. However, that does not mean they automatically see all rows through the application query path defined by the RLS policy. The security policy filters rows based on SUSER_SNAME() and matching RegionCode, so row visibility is governed by the predicate unless the policy is altered or bypassed administratively. DDM and RLS solve different problems: DDM affects how returned values are shown, while RLS affects which rows are returned at all.

The third statement is No because masking only applies to data that is actually returned in the query result set.

Microsoft describes DDM as hiding sensitive data in the result set of a query . If RLS filters a row out, that row is not returned, so there is nothing left for masking to act on. In other words, RLS eliminates inaccessible rows first from the user's perspective, and DDM masks sensitive column values only on rows the user is allowed to see.

NEW QUESTION # 42

You have an Azure SQL database that contains a table named dbo.orders, dbo.orders contains a column named createDate that stores order creation dates.

You need to create a stored procedure that filters Orders by CreateDate for a single calendar day. The solution must be SARGable. 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:

□ Explanation:

□ The correct SARGable pattern for filtering a single calendar day is to use a half-open date range :

o.CreateDate >= @StartDate

AND o.CreateDate < @EndDate

with:

SET @EndDate = DATEADD(day, 1, @StartDate)

This is the correct design because it keeps the function off the column and applies it only to the parameter.

That allows SQL Server and Azure SQL to use an index on CreateDate efficiently, which is the key requirement for a SARGable predicate. Microsoft documents DATEADD as the standard function for adding one day to a date value, which makes it the right way to derive the exclusive upper boundary for the next day.

The incorrect choices are the ones that wrap CreateDate in CONVERT(...), because expressions like:

CONVERT(char(10), CreateDate, 121) = ...

make the predicate non-SARGable and typically prevent efficient seeks on an index over CreateDate.

So the completed procedure is:

```
CREATE PROCEDURE dbo.usp_SearchOrders
```

```
@StartDate date
```

```
AS
```

```
BEGIN
```

```
SET NOCOUNT ON;
```

```
DECLARE @EndDate date;
```

```
SET @EndDate = DATEADD(day, 1, @StartDate);
```

```
SELECT o.CreateDate,
```

```
o.OrderId,
```

```
o.ShipDate
```

```
FROM dbo.Orders AS o
```

```
WHERE o.CreateDate >= @StartDate
```

```
AND o.CreateDate < @EndDate;
```

```
END;
```

NEW QUESTION # 43

You have a GitHub Codespaces environment that has GitHub Copilot Chat installed and is connected to a SQL database in Microsoft Fabric named DB1. DB1 contains tables named Sales.Orders and Sales.Customers.

You use GitHub Copilot Chat in the context of DB1 .

A company policy prohibits sharing customer Personally Identifiable Information (PII), secrets, and query result sets with any AI service.

You need to use GitHub Copilot Chat to write and review Transact-SQL code for a new stored procedure that will join Sales.Orders to sales .Customers and return customer names and email addresses. The solution must NOT share the actual data in the tables with GitHub Copilot Chat.

What should you do?

- A. Run a select statement that returns customer names and email addresses and provide the result set to GitHub Copilot Chat so that GitHub Copilot Chat can generate the stored procedure.
- B. From Sales.Customers, paste several rows that include email addresses into a chat, so that GitHub Copilot Chat can infer edge cases.
- C. Provide the database connection string to GitHub Copilot Chat so that GitHub Copilot Chat can validate the stored procedure.
- **D. Ask GitHub Copilot Chat to generate the stored procedure by using schema details only.**

Answer: D

Explanation:

The correct answer is D because the policy explicitly prohibits sharing customer PII, secrets, and query result sets with any AI service. The safe way to use GitHub Copilot Chat here is to provide only schema- level information such as table names, column names, relationships, and the required procedure behavior, without sharing actual table contents or result sets. That lets Copilot help generate and review the Transact- SQL while avoiding disclosure of customer data. This is consistent with Microsoft and GitHub guidance that content provided in prompts is what the AI can use, so avoiding real data in the prompt is the appropriate control.

The other options violate the requirement:

