PDF DP-800 VCE | DP-800 Certification Dump

You need to do something immediately to change the situation. For instance, the first step for you is to choose the most suitable DP-800 actual dumps for your coming exam. Just like the old saying goes, the little things will determine success or failure.so the study materials is very important for you exam, because the study materials will determine whether you can pass the DP-800 Exam successfully or not. However, how to choose the best and suitable study materials for yourself in a short time? It is very difficult for a lot of people to do a correct choice, especially these people who have no any experience about the DP-800 exam. We would like to tell you how to buy the most suitable and helpful study materials.

Microsoft DP-800 Exam Syllabus Topics:

Topic	Details
Topic 1	• 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.
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	• 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.

>> PDF DP-800 VCE <<

DP-800 Certification Dump, Test DP-800 Dumps

You will have the chance to renew your knowledge while getting trustworthy proof of your expertise with the Microsoft DP-800 exam. After passing the Microsoft DP-800 certification exam, you can take advantage of a number of extra benefits. The Microsoft DP-800 Certification test, however, is a valuable and difficult credential. But with the correct concentration, commitment, and DP-800 exam preparation, you could ace this test with ease.

Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q15-Q20):

NEW QUESTION # 15
You have an Azure SQL database that stores order data.
A reporting query aggregates monthly revenue per customer runs frequently.
You need to reduce how long it takes to retrieve the calculated values. The solution must NOT alter any underlying table structure.
What should you do?

- A. Create a view by using WITH SCHEMABINDING, include COUNT_BIG(*), and then create a unique clustered index on the view.
- B. Create a view without using WITH SCHEMABINDING, and then create a nonclustered index on the view.
- C. Create a view by using GROUP BY, and then create a unique clustered index on the view.
- D. Create a view by using ORDER BY without TOP, and then create a unique clustered index on the view.

Answer: A

Explanation:
Creating an indexed view using WITH SCHEMABINDING, including COUNT_BIG(*), and creating a unique clustered index is a valid and effective approach to significantly improve the performance of your frequent reporting query without altering the underlying table structure.
This process materializes the aggregated data and stores it physically in the database, so the query optimizer can read from the precomputed view rather than rescanning the base tables every time the query runs.
Reference:
https://learn.microsoft.com/en-us/sql/relational-databases/views/create-indexed-views

NEW QUESTION # 16
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 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: Identify the partition scheme for the oldest month, and then run the following Transact- SQL statement.
ALTER TABLE dbo.Orders
DROP PARTITION SCHEME (partition_scheme_name);
Does this meet the goal?

- A. No
- B. Yes

Answer: A

Explanation:

Correct:
* Identify the partition number for the oldest month, and then run the following Transact-SQL statement.
TRUNCATE TABLE dbo.Orders
WITH (PARTITIONS (partition number));
The best Transact-SQL statement to remove all rows for the oldest month while minimizing the impact on other queries is TRUNCATE TABLE with a WITH (PARTITIONS (...)) clause.
Why TRUNCATE TABLE ... WITH (PARTITIONS (...)) is Best
Efficiency: TRUNCATE TABLE is a Data Definition Language (DDL) operation that removes data by deallocating the data pages, which is a metadata operation and is very fast, regardless of the amount of data in the partition.
Minimal Logging: It uses less transaction log space compared to a DELETE statement, which logs each row deletion individually.
Low Impact on Concurrency: It performs a quick, partition-specific operation. A row-by-row DELETE would be a long-running transaction and could cause locking and blocking issues for other queries accessing the table.
Data Integrity: Because the table has a clustered index and is partitioned by the same column (aligned indexes), the TRUNCATE PARTITION operation is a fast, partition-level maintenance operation that targets only that specific data subset.
Incorrect:
* : Identify the partition scheme for the oldest month, and then run the following Transact-SQL statement.
ALTER TABLE dbo.Orders
DROP PARTITION SCHEME (partition_scheme_name);
The DROP PARTITION SCHEME statement removes the partition scheme object from the database but does not remove the data itself or free up the space, and it requires all tables to be moved off the scheme first, which is a complex operation. This does not meet the goal of removing the data efficiently.
* Run the following Transact-SQL statement.
DELETE FROM dbo.Orders
WHERE OrderDate < DATEADD(month, -36, SYSUTCDATETIME());
A standard DELETE statement, even with a WHERE clause that uses the partition column, can be a time-consuming, logged operation that causes locking and blocking on the main table, negatively impacting performance.
Reference:
https://stackoverflow.com/questions/63632963/truncate-partition-vs-drop-partition-performace- wise-which-one-is-efficient-an

NEW QUESTION # 17
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 the Microsoft.SqlServer.Dacpacs.Azure.Master NuGet package to the project.
Does this meet the goal?

- A. No
- B. Yes

Answer: B

Explanation:
Correct:
* Add the Microsoft.SqlServer.Dacpacs.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.Dacpacs.Azure.Master NuGet package.
In your .sqlproj file, include the following item group:
<ItemGroup>
<PackageReference Include="Microsoft.SqlServer.Dacpacs.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.Dacpacs.Master NuGet package to the project.
Reference:
https://learn.microsoft.com/en-us/sql/tools/sql-database-projects/concepts/system-objects

NEW QUESTION # 18
What is Retrieval-Augmented Generation (RAG)?

- A. Combining LLMs with external data sources
- B. A query optimization technique
- C. A SQL indexing method
- D. A backup strategy

Answer: A

NEW QUESTION # 19
You have a SQL database in Microsoft Fabric that contains a column named Payload. Payload stores customer data in JSON documents that have the following format.

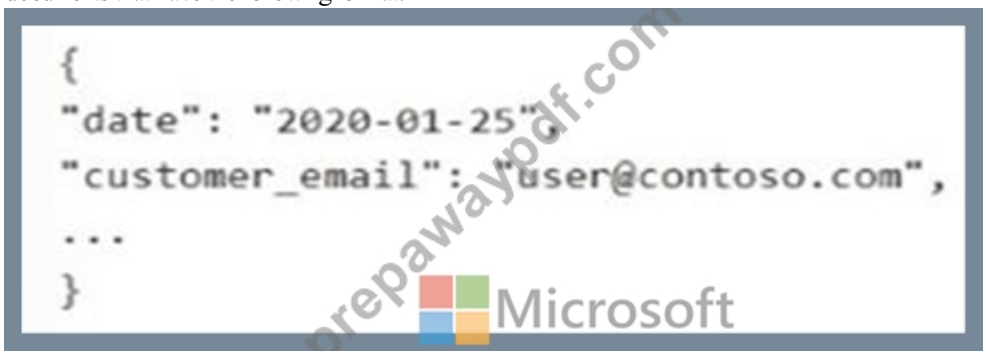

Data analysis shows that some customers have subaddressing in their email address, for example, user1+promo@contoso.com. You need to return a normalized email value that removes the subaddressing, for example, user1 +promo@contoso.com must be normalized to user1@contoso.com.
Which Transact-SQL expression should you use?

- A. REGEXP_REPLACE(JSON_VALUE(Payload, '$.customer_email'), '\+.*', '')
- B. REGEXP_REPLACE(JSON_VALUE(Payload, '$.customer_email'), '\+.*@', '@')
- C. REGEXP_REPLACE(JSON_VALUE(Payload, '$.customer_email'), '\+.*$', '')
- D. REGEXP_SUBSTR(JSON_VALUE(Payload, '$.customer_email'), '